Performance Bounds for Estimation of Structurally Constrained Signals

Zvika Ben-Haim

Performance Bounds for Estimation of Structurally Constrained Signals

Research Thesis

In Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy

Zvika Ben-Haim

Submitted to the Senate of the Technion—Israel Institute of Technology

Cheshvan 5771 Haifa November 2010

Acknowledgements

This Research Thesis was conducted under the supervision of Prof. Yonina C. Eldar in the Faculty of Electrical Engineering.

Chapter 5 represents joint work with A. Jung and Prof. F. Hlawatsch of the Institute of Communications and Radio-Frequency Engineering, Vienna University of Technology. Chapter 6 represents joint work with Prof. M. Elad of the Dept. of Computer Science at the Technion. Chapter 8 represents joint work with T. Michaeli of the Dept. of Electrical Engineering at the Technion. The collaboration with these colleagues is greatly appreciated.

The generous financial assistance of the Wolf Foundation, the Irwin and Joan Jacobs Fellowship, and the Viterbi Family Foundation is gratefully acknowledged.

Contents

Abstract **1**

Notation **3**

Abbreviations **5**

1 Introduction: Structured Estimation Problems **7**
1.1 Sparsity 9
1.2 Block Sparsity 13
1.3 Finite Rate of Innovation 14
1.4 Summary of Main Results 17

2 Background: Performance Bounds **21**
2.1 Frequentist Bounds 21
2.2 Bayesian Bounds 29

I Structural Constraints in the Frequentist Setting **37**

3 The CRB with Singular Fisher Information **39**
3.1 Introduction 39
3.2 Problem Statement 40
3.3 Cramér–Rao Bound 41
3.4 Example 46

4 The CRB for Sparse Estimation **49**
4.1 Introduction 49
4.2 Sparse Estimation Problems 52
4.3 The Constrained Cramér–Rao Bound 55
4.4 Bounds on Sparse Estimation 62
4.5 Numerical Results 70
4.6 Discussion 71

5 Bounds for Sparse Estimation with a Unitary Dictionary 73
5.1 Introduction . . . 73
5.2 The Sparse Signal in Noise Model . . . 76
5.3 Lower Bounds on the Minimum MSE . . . 80
5.4 Upper Bound on the Minimum MSE . . . 84
5.5 Numerical Results . . . 87
5.6 Conclusion . . . 93
5.A Proof of Theorem 5.1 . . . 94
5.B Proof of Theorem 5.3 . . . 98
5.C Proof of Proposition 5.4 . . . 100
5.D Obtaining the CRB from the HCRB . . . 102
5.E Proof of Theorem 5.5 . . . 102
5.F Proof of Lemma 5.6 . . . 108
5.G Proof of Lemma 5.7 . . . 110
5.H Proof of Equation (5.30) . . . 112
5.I MSE of the ML Estimator . . . 114

6 Performance Guarantees for Sparse Estimation 117
6.1 Introduction . . . 117
6.2 Preliminaries . . . 119
6.3 Performance Guarantees . . . 123
6.4 Comparison with Related Estimation Settings . . . 128
6.5 Numerical Results . . . 131
6.6 Conclusion . . . 136
6.A Proof of Theorem 6.3 . . . 137
6.B Proof of Theorem 6.4 . . . 138
6.C Proof of Theorem 6.6 . . . 141

7 Bounds and Guarantees for Block-Sparse Estimators 147
7.1 Introduction . . . 147
7.2 Problem Setting . . . 150
7.3 Techniques for Block-Sparse Estimation . . . 152
7.4 Guarantees for Adversarial Noise . . . 154
7.5 The Cramér–Rao Bound . . . 156
7.6 Guarantees for Gaussian Noise . . . 158
7.7 Numerical Experiments . . . 162
7.8 Conclusion . . . 167
7.A Proofs for Adversarial Noise . . . 168

7.B Proof of Theorem 7.4 . . . 173
7.C Proofs for Gaussian Noise . . . 175

8 Performance of Finite Rate of Innovation Signals 181
8.1 Introduction . . . 181
8.2 Definitions . . . 185
8.3 Types of FRI Signals . . . 187
8.4 Summary of Main Results . . . 190
8.5 Mathematical Prerequisites: CRB for General Parameter Spaces . . . 194
8.6 CRB for Continuous-Time Measurements . . . 197
8.7 CRB for Sampled Measurements . . . 199
8.8 Optimal Sampling Kernels: A Bayesian Viewpoint . . . 205
8.9 Application: Channel Estimation . . . 212
8.10 Conclusion . . . 224
8.A Proof of Theorem 8.1 . . . 225
8.B Maximization of (8.77) . . . 229

II Performance Analysis in the Bayesian Setting 233

9 The Optimal-Bias Bound 235
9.1 Introduction . . . 235
9.2 The Optimal-Bias Bound . . . 238
9.3 Mathematical Safeguards . . . 242
9.4 Calculating the Bound . . . 245
9.5 Properties . . . 247
9.6 Example: Uniform Prior . . . 250
9.7 Conclusion . . . 257
9.A Some Technical Lemmas . . . 257
9.B Proof of Proposition 9.1 . . . 261
9.C Proof of Theorem 9.3 . . . 262
9.D Proof of Theorem 9.5 . . . 263
9.E Proofs of Asymptotic Properties . . . 268

10 A Comment on the Weiss–Weinstein Bound 275
10.1 Introduction . . . 275
10.2 Background and Summary . . . 276
10.3 Scalar Case . . . 278
10.4 Vector Case . . . 278

10.5 Counterexample 280
10.6 Acknowledgement 281

11 Conclusion **283**

Bibliography **287**

List of Figures

1.1 Example of a wavelet transform. The existence of nonzero components in a given detail level are correlated with nonzero values in the corresponding entries of other detail levels. 13

4.1 In a locally balanced set such as a union of subspaces (a) and an open ball (b), each point is locally defined by a set of feasible directions along which an infinitesimal movement does not violate the constraints. The curve (c) is not characterized in this way and thus is not locally balanced. 56

4.2 MSE of various estimators compared with the unbiased CRB (4.44), for (a) varying SNR and (b) varying sparsity levels. 69

5.1 Lower bounds $\mathrm{HCRB}(\boldsymbol{x})$, $\mathrm{HCRB}_{\mathcal{V}}(\boldsymbol{x})$ and upper bounds $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})$, $\mathrm{BB}'_{\mathrm{c}}(\boldsymbol{x})$ on the MSE $\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}^{(x)})$ of the LMVU estimator at $\boldsymbol{x} = c\,(1\ \ 0\ \ 0\ \ 0\ \ 0)^T$, with c varied to obtain different values of $\mathrm{SNR}(\boldsymbol{x}) = \xi^2/\sigma^2$. The SSNM parameters are $N = 5$, $S = 1$, and $\sigma^2 = 1$. 88

5.2 MSE $\varepsilon(\boldsymbol{x}_r;\hat{\boldsymbol{x}}_{\mathrm{ML}})$ of the ML estimator for randomly generated parameter vectors $\boldsymbol{x}_r$ at four different SNRs ξ^2/σ^2, for SSNM parameters $N = 10$, $S = 4$, and $\sigma^2 = 1$. . 90

5.3 MSE of the ML and HT estimators compared with the performance bounds $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})$, $\mathrm{HCRB}(\boldsymbol{x})$, and CRB ($\equiv S\sigma^2$), as a function of the SNR ξ^2/σ^2, for SSNM parameters $N = 10$, $S = 4$, and $\sigma^2 = 1$. 92

5.4 Ratio $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})/\mathrm{HCRB}(\boldsymbol{x})$ versus the SNR ξ^2/σ^2 for different sets of parameter vectors $\boldsymbol{x}$. 93

6.1 Median estimation error for practical estimators (gray regions) compared with performance guarantees (solid lines) and the oracle estimator (dotted line). The gray regions report the range of performances observed for 8 different values of the unknown parameter vector $\boldsymbol{x}_0$. For the Dantzig selector, both the guarantee of Candes and Tao and that of Cai, Wang, and Xu are plotted. For OMP, performance is only guaranteed for $\sigma \leq 0.057$, while for thresholding, nothing can be guaranteed for the given problem dimensions. 132

6.2 MSE of various estimators as a function of the SNR. The sparsity level is $s = 5$ and the dictionary is a 256×512 two-ortho matrix. 134
6.3 MSE of various estimators as a function of the support size s. The noise standard deviation is $\sigma = 0.01$ and the dictionary is a 256×512 two-ortho matrix. 136

7.1 Median squared error as a function of the noise variance for block and scalar sparse estimation algorithms. The shaded region indicates the range of errors encountered for different parameter values. The dotted line plots the CRB. The thick solid line in Figs. 7.1(a) and 7.1(b) indicates the performance guarantees for the block sparse algorithms; no guarantee can be made for the scalar sparsity techniques in Figs. 7.1(c) and 7.1(d). 163
7.2 Median squared error as a function of the noise variance for block sparse estimation algorithms. The shaded region indicates the range of errors encountered for different parameter values. The dotted line plots the CRB. The thick solid line in Fig. 7.2(a) indicates the performance guarantee for BOMP; no guarantee can be made for BTH. The deteriorated performance of BTH is a result of the existence of low-magnitude blocks. 165

8.1 Comparison of the CRB and the performance of a practical estimator, as a function of the number of samples. 218
8.2 Comparison between the CRB and the performance of a practical estimator as a function of the pulse positions. The signal contains $L = 2$ pulses, the first of which is located at $t_1 = 0.5$. The MSE is plotted as a function of the position of the second pulse. 220
8.3 Demonstration of the different levels of overlap between pulses. 221
8.4 Comparison between the CRB for a periodic signal (8.16) and a semi-periodic signal (8.17). 222

9.1 A sequence of continuous functions for which both $|b(\theta)|^2$ and $|1 + b'(\theta)|^2$ tend to zero for almost every value of θ. 242
9.2 Comparison of the MSE bounds and the minimum achievable MSE in a one-dimensional setting for which $\theta \sim U[-r, r]$ and $x|\theta \sim N(\theta, \sigma^2)$. 253
9.3 Comparison of the MSE bounds and the minimum achievable MSE in a three-dimensional setting for which $\boldsymbol{\theta}$ is uniformly distributed over a ball of radius r and $\boldsymbol{x}|\boldsymbol{\theta} \sim N(\boldsymbol{\theta}, \sigma^2 \boldsymbol{I})$. 256

10.1 Plot of the corrected WWB (10.24), the incorrect version (10.6), and the actual MSE, for a scalar estimation setting with a disjoint set Θ. 279

Abstract

The field of signal processing has witnessed an ongoing quest for mathematical models describing natural phenomena, together with a search for techniques utilizing these models in applications such as signal quality improvement. Two common types of mathematical models for estimation problems are the Bayesian scenario, in which a prior probability distribution quantifies information about the unknown parameter, and structural constraints in a frequentist setting, where the parameter is deterministic, but belongs to a pre-specified set or family within a larger space of conceivable values. For example, the recently popular concept of sparsely representable signals can be modeled using either of these approaches: One can define a Bayesian model in which the prior probability of obtaining non-sparse signals is small or zero; alternatively, one can adopt a frequentist approach in which the unknown, deterministic parameter belongs to a set of sparse signals.

The goal of this work is to analyze the performance of various signal estimation techniques in both the Bayesian and structured frequentist settings. The theoretical foundations of our work are based on several extensions on the Cramér–Rao bound (CRB). Specifically, in the frequentist setting, our results rely on a novel interpretation of the effect of parametric constraints on the CRB, as well as a generalization of the CRB for estimating continuous-time functions. While the CRB is a bound for the frequentist (or deterministic) scenario, we show that with an appropriate adaptation, the bound can be applied in the Bayesian world as well, and derive a general lower bound on the minimum achievable MSE in Bayesian estimation settings.

The aforementioned results are complemented by upper bounds on the performance of specific estimators. Such upper bounds can be thought of as guarantees which ensure that a specific approach will never be worse than an analytically computable limit. The combination of upper and lower bounds identifies cases in which state-of-the-art techniques approach the theoretically optimal performance, while illuminating scenarios in which existing methods can be improved.

We demonstrate the applicability of our work in three practical estimation problems: sparsely representable signals, block-sparse parameters, and the finite rate of innovation (FRI) scenario. We demonstrate that under appropriate assumptions, sparse and block-sparse estimation algorithms come fairly close to the theoretical limit. The situation is less favorable in the case of FRI estimation, where there is often a large gap between existing techniques and the lower bound, implying that there is still room for improvement of FRI methods. This analysis allows us to distinguish fundamental limitations in the problem setting from flaws in existing algorithms, thus pinpointing directions for the development of improved estimation techniques.

Notation

$\boldsymbol{v}$	Vector
$\boldsymbol{M}$	Matrix
$\boldsymbol{M}^T$	Transpose
$\boldsymbol{M}^*$	Hermitian conjugate (for matrices)
$\boldsymbol{M}^{-1}$	Matrix inverse
$\boldsymbol{M}^\dagger$	Moore–Penrose pseudoinverse
$\mathrm{Tr}(\boldsymbol{M})$	Trace
$\boldsymbol{M} \succeq \boldsymbol{0}$	Positive semidefinite matrix
$\|\boldsymbol{v}\|_p$	ℓ_p norm ($p \geq 1$)
$\|\boldsymbol{v}\|_0$	Number of nonzero entries in $\boldsymbol{v}$
$\|\boldsymbol{M}\|$	Spectral norm (largest singular value of $\boldsymbol{M}$)
$\|\boldsymbol{M}\|_F$	Frobenius norm
$\mathcal{N}(\boldsymbol{M})$	Null space
$\mathcal{R}(\boldsymbol{M})$	Range space (column span)

Abbreviations

BB	Barankin bound
BOMP	block orthogonal matching pursuit
BPDN	basis pursuit denoising
BTH	block thresholding
CRB	Cramér–Rao bound
DS	Dantzig selector
ERC	exact recovery coefficient
FIM	Fisher information matrix
FRI	finite rate of innovation
GDS	Gauss–Dantzig selector
HCRB	Hammersley–Chapman–Robbins bound
IID	independent and identically distributed
KLT	Karhunen–Loève transform
LMVU	locally minimum variance unbiased
LS	least squares
MAP	maximum a posteriori
ML	maximum likelihood
MMSE	minimum mean squared error
MP	matching pursuit
MSE	mean squared error
OBB	optimal bias bound

OMP	orthogonal matching pursuit
pdf	probability density function
RIC	restricted isometry constant
RIP	restricted isometry property
ROP	restricted orthogonality property
SNR	signal-to-noise ratio
SSNM	sparse signal in noise model
UMVU	uniformly minimum variance unbiased
WWB	Weiss–Weinstein bound

Chapter 1

Introduction: Structured Estimation Problems

A central challenge of engineering is to formulate realistic models of physical phenomena which are amenable to mathematical analysis. Indeed, many of the great successes of signal processing can ultimately be attributed to the introduction of a novel mathematical technique which effectively describes particular types of signals. A classical example is the theory of wavelets, which has considerably improved the analysis and processing of signals containing transient effects [1]. A more recent example is the model of sparsely representable signals, which has attracted much research in the past decade [2]. This model is based on the observation that many signals can be approximated by a linear combination of a small number of atoms chosen from a large, suitably constructed dictionary.

An idea such as the sparse representability of signals can be modeled mathematically in various ways. In our work, we will be primarily interested in two types of model formalisms, namely, the Bayesian setting and the imposition of structural constraints in a frequentist scenario. In the Bayesian setting [3], the structure takes the form of a prior distribution on the signal to be estimated, as, for example, when sparse signals are assigned a higher prior probability than their non-sparse counterparts [4, 5]. By contrast, the frequentist (or deterministic) approach [6] models the unknown signal as a deterministic quantity having no prior distribution [7, 8]. The model then takes the form of a constraint set of possible signal values. For instance, one can adopt the assumption that the signal belongs to the set of all vectors having a sparse representation, while refraining from assigning a probability to each element within this set. The frequentist approach is thus more appropriate when one cannot make precise

statements about the *a priori* likelihood of different signals [6, § 1.1].

The goal of this dissertation is to analyze the performance achievable when various signal models are employed. Specifically, we focus on problems of signal estimation in the presence of noise. Our analysis is performed both in the Bayesian and constrained frequentist worlds. To this end, we develop general theoretical tools, which are suitable for a variety of models, and then apply our results to models of particular interest, such as sparsely representable signals.

To understand the performance of estimation algorithms, we develop two complementary types of results, namely, lower bounds and performance guarantees. Lower bounds are fundamental limits on the achievable performance in a given model. They provide a means for quantifying the difficulty of a model and, occasionally, for understanding how the model can be changed in order to facilitate estimation. In some cases, lower bounds can also be compared with the actual achievements of practical techniques. However, it is not always feasible to fully ascertain the performance of an estimator: The space of parameters or settings under which an estimator is required to function may be too large for an exhaustive analysis. In these cases it becomes necessary to derive upper bounds on the performance of the estimator under consideration, stating that under certain assumptions, a given technique performs no worse than a specified level. These bounds are specific to a particular technique and can be thought of as guarantees on its performance, in lieu of a complete description of the estimator's capabilities under all conceivable scenarios.

Ideally, one would like the lower bound to be close to the performance guarantee of a practical estimator: this indicates that the best achievable performance is known, and an estimator which approaches this optimal performance is available. Conversely, when a gap exists between the upper and lower bounds, this may indicate a flaw in the algorithm, or it may be the consequence of an overly weak bound. By exploring the situations in which such gaps occur, we will glean hints of the directions in which improvements can be made to both the estimator and the bound.

Several specific models will serve as practical examples of the theory developed in this dissertation. These include the aforementioned sparse representation setting; the extension of this scenario to block sparsity; and the related model of signals having a finite rate of innovation. Notably, many instances of these models are unions of subspaces: the set of feasible signals is a union of several subspaces within a larger space of conceivable values [9]. This structure will prove useful in developing a geometrical interpretation of the constraint sets. In the remainder of this chapter, we place our work in context by briefly summarizing prior work on these

estimation settings.

1.1 Sparsity

As is often the case with good ideas, precursors of the concept of sparse representation appeared some time before the explosion of research on sparsity in the last decade. Indeed, the idea of feature selection (finding a small set of features or variables which best explain observed phenomena), a fundamental concept in statistics and machine learning, can be framed in a way that is equivalent to the basic problem of sparsity pattern recovery [10]. However, as a model for signal representation, sparsity arose from the coming of age of the theory of wavelets in the 1990s [1, 11]. Practical work on wavelets led to the realization that real-world signals are best described as a linear combination of a small number of elements, or *atoms*, from a large and overcomplete *dictionary*. Thus, a signal $\boldsymbol{y} \in \mathbb{R}^m$ can be approximated as

$$\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x}_0 + \boldsymbol{w} \tag{1.1}$$

where $\boldsymbol{H} \in \mathbb{R}^{m \times p}$ is a dictionary, typically having far more columns than rows; $\boldsymbol{x}_0 \in \mathbb{R}^p$ is a sparse vector; and $\boldsymbol{w}$ represents noise or model mismatch. The sparsity requirement can be manifested, for example, as the structural constraint

$$\|\boldsymbol{x}_0\|_0 \leq s \tag{1.2}$$

where $\|\boldsymbol{x}_0\|_0$ is the number of nonzero entries in the vector $\boldsymbol{x}_0$, and s is a sparsity level which is usually assumed to be known. Thus, the support pattern of $\boldsymbol{x}_0$ effectively chooses columns of $\boldsymbol{H}$ to be included in the representation of $\boldsymbol{y}$, giving rise to the relation with the feature selection problem. It is not difficult to see that this structure forms a union of subspaces of $\mathbb{R}^p$: Each choice of an s-element support for $\boldsymbol{x}_0$ forms an s-dimensional subspace, and the set of allowed values of $\boldsymbol{x}_0$ is thus the union of all such subspaces.

In early sparsity applications, the dictionary $\boldsymbol{H}$ was chosen manually so as to combine different atoms which were thought to be likely candidates for describing the family of signals at hand. For example, $\boldsymbol{H}$ could be the concatenation of the Fourier basis, the trivial basis, and one or more wavelet bases. In time, the sparsity model (1.1) came to be used in a variety of settings unrelated to wavelet models. For example, much research has been devoted to the setting in which $\boldsymbol{H}$ is selected randomly, e.g., by choosing its entries independently from a Gaussian distribution [8, 12, 13]. As another example, $\boldsymbol{H}$ can be constructed empirically from a training

set of signals, resulting in a dictionary which has no pre-designed structure [14]. Regardless of the method by which $\boldsymbol{H}$ is chosen, we will assume throughout that the columns $\boldsymbol{h}_i$ of $\boldsymbol{H}$ are normalized, $\|\boldsymbol{h}_i\|_2 = 1$, which ensures that issues of SNR can be analyzed independently of the choice of the dictionary itself.

The fundamental estimation problem arising from the sparsity model (1.1) is that of recovering $\boldsymbol{x}_0$, either from $\boldsymbol{y}$ or from linearly transformed measurements thereof. Many important signal processing applications can be derived from this estimation problem. For example, $\boldsymbol{y}$ can be considered as a noisy version of the sparsely representable signal $\boldsymbol{H}\boldsymbol{x}_0$, and denoising can thus be performed by estimating $\boldsymbol{x}_0$ and then pre-multiplying it by $\boldsymbol{H}$. Interpolation, deblurring, inpainting, and super-resolution are further examples of applications in which sparsity has been successfully applied, often leading to state-of-the-art results [15].

How does one go about estimating $\boldsymbol{x}_0$ from the measurements $\boldsymbol{y}$? A standard approach is to assume a probabilistic model for the noise and then compute the maximum likelihood (ML) estimator [16]. In our case, assuming independent, unimodal, zero-mean noise model (such as white Gaussian noise) implies that the ML estimator is the solution of the optimization problem

$$\min_{\boldsymbol{x}} \|\boldsymbol{y} - \boldsymbol{H}\boldsymbol{x}\|_2^2 \quad \text{s.t. } \|\boldsymbol{x}\|_0 \leq s. \tag{1.3}$$

Unfortunately, solving (1.3) is NP-complete, i.e., to the best of our knowledge, finding the ML estimator is infeasible computationally, even for small dictionary sizes. In light of this, it is perhaps surprising that a number of practical algorithms can often accurately estimate $\boldsymbol{x}_0$ in many cases of interest. These practical approaches can be broadly divided into two classes: greedy approaches and ℓ_1 relaxation techniques.

Greedy approaches are iterative algorithms in which, at each iteration, one selects the "most likely" atom in $\boldsymbol{H}$ for inclusion in the estimated support of $\boldsymbol{x}_0$. For example, in matching pursuit (MP) [11], a residual $\boldsymbol{r}$ is maintained which represents the components of $\boldsymbol{y}$ not yet accounted for by the atoms already chosen from $\boldsymbol{H}$. In each iteration, the atom most highly correlated with $\boldsymbol{r}$ is added to the estimated support set. Other algorithms of the greedy variety include orthogonal matching pursuit (OMP) [17], thresholding, and CoSaMP [18].

As an alternative to the greedy approach, ℓ_1 relaxation methods attempt to change the NP-complete problem (1.3) into a convex (and thus efficiently solvable) optimization problem. This is done by replacing the ℓ_0 constraint $\|\boldsymbol{x}\|_0 \leq s$ with a constraint based on the ℓ_1 norm. The most straightforward approach is then to solve

$$\min_{\boldsymbol{x}} \|\boldsymbol{y} - \boldsymbol{H}\boldsymbol{x}\|_2^2 \quad \text{s.t. } \|\boldsymbol{x}\|_1 \leq \tau. \tag{1.4}$$

This technique is referred to as the lasso [19] or basis pursuit [20]. Note that the constraint $\|\boldsymbol{x}\|_1 \leq \tau$ no longer directly controls the sparsity level s of the outcome. However, tightening the constraint by decreasing the value of τ tends to promote sparsity in the solution of (1.4). A related technique, also based on replacing the ℓ_0 "norm" with an ℓ_1 norm, is known as the Dantzig selector [8].

With the availability of practical techniques for estimating $\boldsymbol{x}_0$ from the measurements (1.1), the question of performance analysis becomes pertinent. Consequently, various types of performance guarantees are available for the algorithms mentioned above. The strength of these results obviously hinges on the assumptions introduced in the model. For example, very powerful results can be obtained if one assumes that both $\boldsymbol{x}_0$ and $\boldsymbol{w}$ are chosen randomly from a known distribution [21], while much weaker guarantees are available if nothing is known about the distribution of $\boldsymbol{x}_0$ and $\boldsymbol{w}$ [7, 12, 22]. As an example of the latter type of guarantee, in the *adversarial* setting, nothing is assumed about the noise $\boldsymbol{w}$ except that it is bounded, $\|\boldsymbol{w}\|_2 \leq \varepsilon$. Thus, it is possible that the noise happens to align itself so as to maximally damage the measurements, an event which is highly improbable if the noise is random. Consequently, performance guarantees in the adversarial setting typically ensure that the estimation error is on the order of the noise power ε. As we will see shortly, when the noise is random, much better guarantees are possible.

In our work on sparse representations, we take the middle ground of structured frequentist estimation. Specifically, we assume that the signal $\boldsymbol{x}_0$ is deterministic (or that its distribution is unknown), but conforms to the sparsity requirement $\|\boldsymbol{x}_0\|_0 \leq s$. On the other hand, we assume that the noise $\boldsymbol{w}$ is random, and specifically follows a white Gaussian distribution with variance σ^2. This is a standard estimation model, which can be motivated by observing that the noise typically has multiple independent sources, to which the central limit theorem can be applied, whereas the signal itself results from a complex process for which any structural assumptions are tentative. Thus, one is often more comfortable assigning a probability distribution to the noise than to the signal.

Previous results on the performance of sparse estimation algorithms in the structured frequentist estimation setting can be summarized as follows. Under suitable conditions on the dictionary $\boldsymbol{H}$, and for sufficiently sparse vectors $\boldsymbol{x}_0$, various ℓ_1 relaxation techniques $\hat{\boldsymbol{x}}$ can be shown to satisfy [8, 23]

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}_0\|_2^2 \leq Cs\sigma^2 \log p \quad \text{with high probability} \tag{1.5}$$

where C is a constant depending on the properties of $\boldsymbol{H}$. Contrary to the case of adversarial noise, in this setting we observe a significant denoising capability: The total noise power $E\{\|\boldsymbol{w}\|_2^2\}$ is equal to $m\sigma^2$, which is typically much larger than the guarantee (1.5), since generally $s \ll m$. Notably, such a guarantee comes within a nearly constant factor of $s\sigma^2$. The expression $s\sigma^2$ is significant since, again up to a constant, it equals the error achievable by the oracle estimator, which knows the true support of $\boldsymbol{x}_0$. Results such as (1.5) are therefore referred to as "near-oracle" performance guarantees. As we will see in Chapter 4, the expression $s\sigma^2$ can also be given a more precise interpretation as the Cramér–Rao bound (CRB) for unbiased estimators in the sparse setting. Finally, we note in passing that the $\log p$ factor in (1.5) can be shown to be a necessary result of the lack of knowledge about the support of $\boldsymbol{x}_0$ [24, §7.4].

Performance guarantees of the form (1.5) must invariably make some assumptions on the matrix $\boldsymbol{H}$. For example, if $\boldsymbol{H}$ contains two identical atoms, then clearly no algorithm will be able to distinguish between the corresponding elements in $\boldsymbol{x}_0$. The assumptions on $\boldsymbol{H}$ typically take the form of requirements on the behavior of sets of atoms, in order to ensure that the subspaces spanned by sets of s atoms are distant from one another. The restricted isometry property (RIP) [25] is the most well-known of these measures. Such properties can be guaranteed (with high probability) for many types of randomly constructed dictionaries. However, if one must use a matrix $\boldsymbol{H}$ fixed by external constraints (for example, an empirically constructed dictionary), then it is NP-complete to compute the RIP (or related properties).

There are some properties of $\boldsymbol{H}$ which can be readily computed, and which are relevant for sparse estimation. The most useful of these is the mutual coherence μ, which is defined as

$$\mu \triangleq \max_{i \neq j} |\boldsymbol{h}_i^* \boldsymbol{h}_j| \tag{1.6}$$

where $\boldsymbol{h}_i$ is the ith column in $\boldsymbol{H}$. Thus, μ is the maximum correlation between dictionary atoms. The coherence can be calculated in $O(p^2)$ time by enumerating over all pairs of atoms. Furthermore, one can bound the value of various more complex properties, including the RIP, using the coherence; this is the standard technique for applying RIP-based performance guarantees when the dictionary is fixed *a priori*. However, as we show, the resulting guarantees lose much of their tightness. Moreover, we demonstrate that direct utilization of the coherence can lead to considerably improved performance guarantees, as well as novel guarantees for greedy algorithms, which were not possible using previous methods.

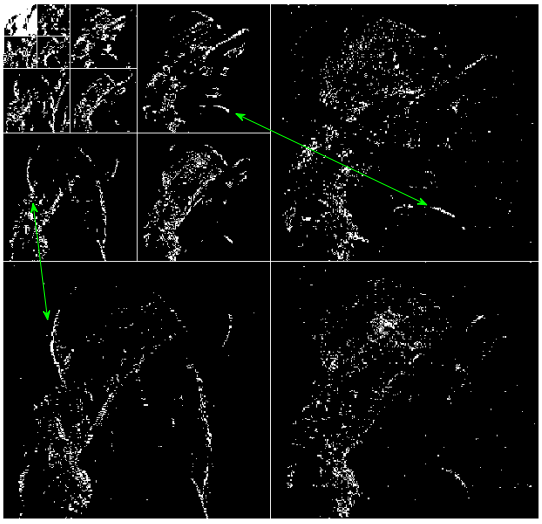

Figure 1.1: Example of a wavelet transform. The existence of nonzero components in a given detail level are correlated with nonzero values in the corresponding entries of other detail levels.

1.2 Block Sparsity

Incorporating knowledge about signal structure can justifiably be expected to improve estimation performance. Consequently, the more one assumes about a signal, the better one should be able to perform estimation tasks such as denoising and interpolation. In this respect, while sparsity is a powerful assumption which can be used to great advantage, it does not always capture all of the structure inherent in the signal. In particular, it has long been known that in many cases, the support pattern of signals tends to form clusters, with "similar" atoms appearing together in the support. For example, atoms representing nearby frequencies are often conjunctively used to represent signals whose precise frequency is not defined by a single atom. Another example is demonstrated in Fig. 1.1, where the wavelet transform of an image is shown. In this example, there is a clear correlation between corresponding entries in different detail levels.

A straightforward modification of the sparse representation model which can take such effects into account is the block sparsity approach [9, 26–28]. In this setting, the unknown parameter $\boldsymbol{x}_0 \in \mathbb{R}^p$ is assumed to be divided into blocks of size d, and the number of nonzero *blocks* is no greater than s. If a certain block is nonzero, however, then any or all of its constituent entries can have nonzero values. With an appropriate choice of blocks containing similar atoms,

this scenario can accurately model the clustering of nonzero entries described above. Note that this model is once again a union of subspaces: Each choice of s nonzero blocks defines a sd-dimensional subspace of $\mathbb{R}^p$, and the set of all block sparse signals is the union of all such subspaces.

Much of the development of the block sparse estimation problem can be performed in analogy to the ordinary (or "scalar") sparsity model of Section 1.1. As before, the problem of finding the maximum likelihood estimator is NP-complete. Practical techniques overcome this difficulty using either an ℓ_1 relaxation or a greedy approach. In particular, block sparse versions of the thresholding algorithm [29], MP [28], OMP [27,28], and the lasso [9,27] have been proposed in the literature.

Being a relatively new model, there have only been a handful of analytical guarantees on the performance of block sparse estimators. Successful recovery of $\boldsymbol{x}_0$ from noiseless measurements $\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x}_0$ has been demonstrated, under suitable conditions on $\boldsymbol{H}$, for both block-OMP and L-OPT (a block-sparse version of the lasso) [9, 27]. A performance guarantee for L-OPT has also been demonstrated in the presence of adversarial noise [9]. However, to the best of our knowledge, there have been no guarantees for any block sparse technique under random noise. Our work (see Chapter 7) begins to fill this gap by providing guarantees for greedy block sparse algorithms, which are generalizations of the corresponding results for non-block methods. These results not only provide numerical guarantees for specific estimation scenarios, but also illustrate the qualitative improvement over scalar sparsity which is achieved by the incorporation of the block sparse model.

1.3 Finite Rate of Innovation

The sparsity and block sparsity models of the previous sections defined a structure on signals in the finite-dimensional space $\mathbb{R}^p$. By contrast, in this section we consider structure in a space of continuous-time signals. Specifically, we investigate signals having a finite rate of innovation (FRI) [30, 31]. As we will see, the transition to infinite-dimensional spaces increases the complexity of the analysis, but it is still possible to obtain a performance analysis from which much insight is gained concerning the strengths and limitations of state-of-the-art techniques.

The power of the FRI model becomes evident when contrasted with the classical notion of Shannon–Nyquist sampling [32]. In classical sampling theory, a bandlimited signal whose maximum frequency is $f_{\max}$ is sampled at the Nyquist rate $2f_{\max}$ or greater. As is well known,

the signal can then be perfectly reconstructed from its samples. Unfortunately, real-world signals are rarely truly bandlimited, if only because most signals have finite duration in the time domain. Even signals which are approximately bandlimited often have a fairly high Nyquist rate, requiring expensive sampling hardware and high-throughput digital machinery.

Classical sampling theory necessitates a high sampling rate whenever a signal has a high bandwidth, even if the actual amount of information in the signal is fairly low. For instance, a piecewise linear signal is non-differentiable; it is therefore not bandlimited, and moreover, its Fourier transform decays at the fairly low rate $O(1/f^2)$. However, the signal is completely described by the positions of knots (transitions between linear segments) and the signal values at those positions. In this sense, as long as the knots are known to have a minimum separation, this signal has a finite information rate. The basic motivation of FRI is that it seems wasteful to sample such signals at the Nyquist rate.

To be specific, suppose that a function $x(t)$ has the following property: Any length-T_0 segment of $x(t)$ is completely determined by no more than K parameters. For example, the aforementioned piecewise linear signal has this property. In this case, the function $x(t)$ is said to have a T_0-local rate of innovation of K/T_0 [30]. In general, a signal is said to have a finite rate of innovation if its T_0-local rate of innovation is finite for any sufficiently large T_0.

There are many families of FRI signals with wide-ranging applications. For example, the piecewise linear model can be written as

$$x(t) = \sum_{n \in \mathbb{Z}} a_n (t - t_n) \mathbb{1}_{t \geq t_n} \tag{1.7}$$

where a_n and t_n are parameters, and $\mathbb{1}_E$ is an indicator function for the event E. To satisfy the FRI requirement, we must assume a minimum separation between the knots, such that $|t_n - t_m| > t_{\min}$ for some $t_{\min} > 0$. This is the simplest case of the class of signals known as splines, which are often used in image processing, and can also be described as FRI signals [33]. Another example of a family of FRI signals is the family of time-delayed pulses

$$x(t) = \sum_{n \in \mathbb{Z}} a_n h(t - t_n) \tag{1.8}$$

where $h(\cdot)$ is a known pulse shape, and again a minimum separation is assumed between the pulses. Such signals find uses in radar, ultrasound, and multipath channel estimation problems [34, 35]. Signals corresponding to the models (1.7) and (1.8) are examples of FRI families which are also unions of subspaces: Indeed, fixing t_n and varying a_n yields a subspace of the set of continuous-time functions, and the entire set of allowed signals is the union of all such

subspaces. However, the FRI model can also be used to describe non-union signals. For example, consider a generalization of (1.8) in which the pulses are controlled by a generic parameter θ, so that

$$x(t) = \sum_{n \in \mathbb{Z}} h(t, \theta_n) \tag{1.9}$$

where θ_n describes one or more parameters controlling the shape and position of the nth pulse. While the time-delay setting (1.8) is clearly a special case of (1.9), one can easily conceive of other families of FRI signals which cannot be described as unions of subspaces. Still, the union model is a useful scenario and will receive special attention in the sequel.

The power of the FRI model is summarized by the rule of thumb which states that FRI signals can typically be reconstructed without error from *samples taken at the rate of innovation* [30, 36]. The advantage of this result is self-evident: FRI signals need not be bandlimited, and even if they are, the Nyquist frequency can be much higher than the rate of innovation. There are, however, several caveats to this result. First, the sampling technique sometimes requires the use of analog filters designed for the specific family of signals, so some hardware flexibility is a prerequisite. Second, the reconstruction algorithm is again designed based on the family of FRI signals under consideration; there is no single algorithm guaranteeing reconstruction for any FRI model. Finally, and most importantly for our point of view, this rule of thumb only holds in the absence of noise. Indeed, empirical observations indicate that, for some noisy FRI signals, substantial performance improvements are achievable when the sampling rate is increased beyond the rate of innovation [34, 37].

The deterioration in observed performance in the presence of even minor levels of noise is clearly of practical relevance, and motivates the study of FRI techniques under noise. In particular, there has been some work in the past which derived the CRB on the ability to estimate the parameters of the FRI signal in the presence of noise [31, 38]. However, these results bounded the achievable performance from a given set of samples, and consequently were based on the assumption that the sampling rate and sampling technique are given. As we will see in Chapter 8, one can also derive a CRB which is independent of the sampling technique, and thus provides a fundamental performance limit which cannot be surpassed, regardless of the type of hardware employed.

1.4 Summary of Main Results

The remainder of this dissertation is organized as follows. We begin in Chapter 2 with background material on various types of performance bounds, in both deterministic and Bayesian settings. As we will see below, these bounds, and particularly the CRB, serve as the basis for many of the developments later in the thesis, and are thus discussed in detail.

The novel results in this dissertation are divided into two parts, corresponding to the frequentist and Bayesian points of view. Part I is the larger of the two and contains results which pertain to the frequentist scenario with structural constraints. Thus, in this part the assumption is that the parameter to be estimated is deterministic (has no prior distribution); prior knowledge about the signal is encapsulated in the assumption that the parameter belongs to a known constraint set. The arrangement within this part follows a deductive order, in which we first derive general estimation bounds and then apply them to specific scenarios. Specifically, our analysis begins in Chapter 3 with a derivation of the constrained CRB for estimation scenarios involving a singular Fisher information matrix (FIM). This case is of importance in the context of sparse estimation, where the matrix $\boldsymbol{H}$ of (1.1) is typically column rank deficient, leading to a singular FIM. For unconstrained problems, a singular FIM leads to an infinite CRB, meaning that no unbiased estimator exists in this case. However, as we show, the adoption of constraints implies that estimation may still be possible even when the FIM is singular, conforming to the fact that estimation is indeed possible in the sparse setting.

Following this general result, the CRB is derived for the specific case of sparsely representable signals in Chapter 4. As mentioned previously, the CRB in this case coincides, for almost all parameter values, with the error achieved by the oracle estimator. This result motivates the use of the oracle error as a standard against which the performance of practical estimators can be judged. In particular, the CRB is known to be achieved at high SNR by the ML estimator, and is therefore arguably more directly related to the performance of real-world techniques.

Despite being a tight bound at high SNR, the CRB is often unreliable when the noise level increases. To gain better understanding of the performance for a wide range of SNR values, in Chapter 5 we consider the Hammersley–Chapman–Robbins bound for the sparse setting. As this bound is analytically far more complicated, we restrict attention in this chapter to the case in which $\boldsymbol{H}$ is a unitary square matrix. This facilitates the derivation of an improved closed-form bound. We demonstrate the tightness of this bound by constructing a family of estimators

which jointly come very close to achieving it.

With the lower bounds for the sparsity setting developed in Chapters 4 and 5, the next challenge is to determine whether any practical technique comes close to achieving these bounds. This is accomplished in Chapter 6 through the use of estimator performance guarantees. As we have already mentioned in Section 1.1 above, some performance guarantees have already been developed for ℓ_1-relaxation techniques based on complex properties of the dictionary $\boldsymbol{H}$, such as the RIP. However, since the RIP is NP-hard to compute, we derive alternative bounds based directly on the mutual coherence. As we show, when the RIP is unknown, our coherence bounds provide tighter results than previously known bounds. Furthermore, our approach enables the derivation of performance guarantees for greedy techniques; such greedy guarantees have not previously been shown in the frequentist setting.

The combination of the performance guarantees of Chapter 6 and the lower bounds of Chapters 4 and 5 demonstrates that existing estimators come fairly close to ideal performance. As we show in Chapter 7, similar results can be obtained for the block sparse model as well: Both the CRB and greedy performance guarantees are supplied for this generalization of the sparse representation scenario. The comparison between the sparse and block sparse settings provides insight into the improvement obtainable by the addition of the block sparsity structure.

Next, in Chapter 8 we move on to the FRI setting. Here we derive two types of lower bounds, both based on the CRB. The first is a fundamental limit on the estimation quality achievable, regardless of the sampling technique, while the second limits the performance given a particular sampling configuration. Comparing these two bounds with the observed quality of actual FRI estimation techniques can answer two questions: First, does a given technique optimally utilize the information given to it by the available hardware? And, second, could an improvement in the hardware (e.g., an increased sampling rate) result in better performance? We demonstrate that the answers to both questions depend on the precise setup under consideration. When existing techniques fail to fully exploit the available information, we identify where these methods could be improved. Finally, we propose a technique for designing sampling schemes which optimally utilize a given sampling rate budget.

In Part II, we adopt the Bayesian point of view in which the unknown parameter is random and has a known prior distribution. In the Bayesian setting, it is well-known that the conditional expectation is the estimator minimizing the MSE. However, in many cases, the computational complexity of the minimum MSE (MMSE) estimator is prohibitive, requiring the use of

alternative techniques, such as the maximum a posteriori (MAP) approach. Furthermore, the performance of the MMSE method is itself often difficult to ascertain. It is therefore of interest to develop bounds on the MMSE performance. These bounds can then be compared with the achievements of practical techniques, so as to determine their proximity to the optimum. This goal has attracted a significant research effort over the years (see Section 2.2).

We provide two contributions to the field of Bayesian performance bounds. In Chapter 9, we consider the optimal bias bound (OBB) [39], an unfamiliar result dating from 1971 which uses the (frequentist) CRB to obtain a Bayesian performance bound. We rigorously demonstrate the soundness of the bound and extend it to the general case of a vector parameter. We also prove the asymptotic tightness of the bound both in the high SNR and in the low SNR regimes, a property which, to the best of our knowledge, has not been demonstrated for any other bound. Indeed, we demonstrate through examples that the OBB is often substantially more accurate than more standard approaches. Next, in Chapter 10, we consider the much more famous Weiss–Weinstein bound. We demonstrate that the standard statement of this bound is inaccurate in particular cases, specifically when the prior distribution is limited to a bounded support set. Indeed, we show examples in which this standard statement produces incorrect results: values which are not bounds on the MMSE performance, and may even be infinite. The source of this inaccuracy is identified and a corrected version is presented.

Chapter 2

Background: Performance Bounds

The main goal of this dissertation is to study the capabilities of estimation techniques in structured settings through the analysis of lower and upper performance bounds. As discussed in the Introduction, an upper bound is a performance guarantee for a given estimator. Consequently, proofs of upper bounds tend to be specialized to the peculiarities of the technique under consideration. By contrast, lower bounds are often based on more generic concepts. Indeed, many of the results in this dissertation are based on several extensions of the Cramér–Rao bound (CRB), be it an extension for the case of a singular Fisher information matrix (Chapter 3), various types of constraint sets (Chapters 4 and 7), infinite-dimensional spaces (Chapter 8), or applications to the Bayesian framework (Chapter 9). Therefore, to provide a basis for the upcoming chapters, we now review the CRB and other lower bounds on estimation performance. Specifically, Section 2.1 deals with bounds in the frequentist (or deterministic) estimation scenario, while Section 2.2 summarizes lower bounds in the Bayesian world.

2.1 Frequentist Bounds

The frequentist estimation setting is based on the assumption that an unknown deterministic parameter $\boldsymbol{x}$ is to be estimated from random measurements $\boldsymbol{y}$ [6, 16]. The probability density function (pdf) of $\boldsymbol{y}$ depends on $\boldsymbol{x}$ and is denoted $p(\boldsymbol{y};\boldsymbol{x})$, where the semicolon is intended to serve as a reminder that $\boldsymbol{x}$ is not random, merely a parameter which influences the distribution.

A standard objective in estimation is to achieve low mean-squared error (MSE), defined as $E\{\|\boldsymbol{x}-\hat{\boldsymbol{x}}\|^2\}$, where $\hat{\boldsymbol{x}}$ is the estimator under consideration. In the frequentist context, this quantity is a function of the parameter $\boldsymbol{x}$. Consequently, the MSE does not form a complete ordering of estimator performance, since one estimator may be better than another for some

values of $\boldsymbol{x}$ but worse for other values. In particular, the trivial estimator $\hat{\boldsymbol{x}} = \boldsymbol{x}_0$ (for some constant $\boldsymbol{x}_0$) achieves zero MSE when $\boldsymbol{x} = \boldsymbol{x}_0$, but is a very poor estimator for other parameter values.

Our goal in this section, and indeed in much of the dissertation, is to construct lower bounds on the MSE achievable in a given estimation problem. However, to construct these bounds, estimators such as $\hat{\boldsymbol{x}} = \boldsymbol{x}_0$ must be excluded from consideration. Otherwise, the only possible bound on the MSE is the zero bound, since there exists an estimator achieving zero MSE at any point $\boldsymbol{x}_0$. What is needed is therefore a mechanism for excluding estimators which are insensitive to changes in the parameter $\boldsymbol{x}$. There are several possible techniques for doing so, including various types of minimax approaches, in which one seeks the estimator whose worst-case error is minimal [40–42]. However, probably the most common technique for the expurgation of trivial approaches is the notion of unbiasedness, which will be adopted throughout the frequentist chapters in this dissertation. An estimator is said to be unbiased if it holds that

$$\boldsymbol{b}(\boldsymbol{x}) \triangleq E\{\hat{\boldsymbol{x}}\} - \boldsymbol{x} = \boldsymbol{0} \quad \text{for all } \boldsymbol{x}. \tag{2.1}$$

In this sense, the bias $\boldsymbol{b}(\boldsymbol{x})$ (or more accurately, the lack of bias) can be viewed as an assurance of the sensitivity of the estimator to changes in the parameter values. Conveniently, the MSE of an unbiased estimator is equal to its variance, so one generally searches for bounds on the variance of estimators $\hat{\boldsymbol{x}}$ satisfying (2.1).

Alternatively, frequentist bounds can be constructed for estimators which have any prespecified bias function $\boldsymbol{b}(\boldsymbol{x})$, not necessarily the zero function. It should be noted, though, that there apparently exists neither a rigorous technique or nor an intuitive approach for choosing desirable bias functions other than $\boldsymbol{b}(\boldsymbol{x}) = \boldsymbol{0}$. We will, however, find a use for biased bounds in Chapter 9, when exploiting the biased CRB in a Bayesian context. Suppose, then, that one desires a certain bias function $\boldsymbol{b}(\boldsymbol{x})$. This is equivalent to requiring an unbiased estimator of the quantity $\boldsymbol{x} + \boldsymbol{b}(\boldsymbol{x})$. Thus, instead of speaking of biased bounds, one can always seek bounds on the variance of estimators of functions of the parameter $\boldsymbol{x}$.

Beyond the desirable sensitivity of an estimator to changes in the parameter value, there are various other justifications for the popularity of unbiasedness as a design criterion. For example, the unbiased CRB is asymptotically achieved by the maximum likelihood (ML) estimator as the number of IID measurements increases. However, this statement, as well as all other formal justifications of unbiasedness, are based on asymptotic analyses. Indeed, in many cases biased estimators can perform better than any unbiased techniques, and indeed better

than bounds on unbiased techniques [42–44]. Consequently, bounds based on principles of unbiasedness should primarily be used to judge performance in the high SNR regime.

In the remainder of this section, we review the frequentist estimation bounds which will be exploited later in the dissertation. We begin with the well-known Cramér–Rao bound. We then discuss its non-local refinement, the Hammersley–Chapman–Robbins bound. Finally, we discuss prior work on applying the CRB to constrained estimation settings, in which the parameter is known to have a specified structure.

2.1.1 The Cramér–Rao Bound

The CRB was independently discovered by several mathematicians, including Cramér [45], Darmois [46], Fréchet [47], and Rao [48]. It is thus probably more historically correct to refer to this bound as the information inequality [6], although the term Cramér–Rao bound is more common in electrical engineering texts [16]. Since much of our work relies on various extensions of the CRB, we now take the time to review the bound in some detail.

Considering the elementary nature of the CRB and its widespread use, the bound is surprisingly difficult to formulate in a completely rigorous fashion. This is primarily due to the necessity of swapping integration over $\boldsymbol{y}$ and differentiation with respect to $\boldsymbol{x}$ in the proof of the CRB, a step which requires an application of the dominated convergence theorem, and thus necessitates domination conditions on the bias and pdf. Since these conditions are cumbersome, not very insightful, and hold for many of the most common distributions, in our discussion below we will simply assume that differentiation by $\boldsymbol{x}$ and integration over $\boldsymbol{y}$ can be swapped when necessary. The reader is referred to [6, Theorem 2.5.15] for a completely rigorous treatment in the case of a scalar parameter. Even with this simplification, readers not acquainted with the CRB will likely find our discussion somewhat challenging, and may find it easier to first follow the more friendly exposition in [16].

Theorem 2.1 (Cramér–Rao bound)**.** *Let $\mathcal{X} \subseteq \mathbb{R}^n$ be an open set and let $\{p(\boldsymbol{y};\boldsymbol{x})\}_{\boldsymbol{x}\in\mathcal{X}}$ be a family of pdfs having a common support, i.e., the set $\{\boldsymbol{y} : p(\boldsymbol{y};\boldsymbol{x}) > 0\}$ is independent of $\boldsymbol{x}$. Suppose that $p(\boldsymbol{y};\boldsymbol{x})$ and $\log p(\boldsymbol{y};\boldsymbol{x})$ are differentiable with respect to $\boldsymbol{x}$ and assume that integration over $\boldsymbol{y}$ and differentiation with respect to $\boldsymbol{x}$ can be interchanged whenever necessary. Define the Fisher information matrix (FIM)*

$$\boldsymbol{J}(\boldsymbol{x}) \triangleq E\left\{\left(\frac{\partial \log p}{\partial \boldsymbol{x}}\right)\left(\frac{\partial \log p}{\partial \boldsymbol{x}}\right)^T\right\} \tag{2.2}$$

and assume that $\boldsymbol{J}(\boldsymbol{x})$ exists and is finite and positive definite. Consider a finite-variance estimator $\hat{\boldsymbol{x}}$

having a differentiable bias function $\boldsymbol{b}(\boldsymbol{x})$*. Then,*

$$\mathrm{Cov}(\hat{\boldsymbol{x}}) \succeq \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{x}}\right) \boldsymbol{J}^{-1}(\boldsymbol{x}) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{x}}\right)^T \tag{2.3}$$

and thus

$$\mathrm{MSE}(\hat{\boldsymbol{x}}) \geq \mathrm{Tr}\left[\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{x}}\right) \boldsymbol{J}^{-1}(\boldsymbol{x}) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{x}}\right)^T\right]. \tag{2.4}$$

The bounds (2.3) *and* (2.4) *are achieved with equality if and only if there exists a matrix function* $\boldsymbol{A}(\boldsymbol{x})$ *such that*

$$\frac{\partial \log p(\boldsymbol{y};\boldsymbol{x})}{\partial \boldsymbol{x}} = \boldsymbol{A}(\boldsymbol{x}) \left(\hat{\boldsymbol{x}} - \boldsymbol{x} - \boldsymbol{b}(\boldsymbol{x})\right). \tag{2.5}$$

Specifically, for unbiased estimators ($\boldsymbol{b}(\boldsymbol{x}) = 0$*), the matrix function* $\boldsymbol{A}(\boldsymbol{x})$ *must equal the FIM* $\boldsymbol{J}(\boldsymbol{x})$*.*

Several remarks are in order concerning Theorem 2.1. First, we observe that the condition (2.5) for achieving the CRB with equality is rather strict. In particular, estimators achieving the unbiased CRB (referred to as efficient estimators) are only possible when the pdf satisfies a specific form. Rearranging terms in (2.5) and substituting the unbiasedness requirement, it can be shown that efficient estimators exist if and only if the class of pdfs forms an exponential family, and then only when estimating the natural parameter of that family [6]. Thus, achievability of the CRB is the exception rather than the rule. Various methods whereby the CRB can be tightened are discussed in Section 2.1.2, below.

It is interesting to note the requirement of a positive definite FIM in the conditions for Theorem 2.1. This condition is necessary: If the FIM is singular, then no finite-variance, unbiased estimator exists [49]. Intuitively, when $\boldsymbol{J}(\boldsymbol{x})$ is singular, its inverse is infinite, so that (2.4) explodes. However, the fact that the FIM is singular does not preclude the existence of unbiased estimators if the parameter is constrained to a sufficiently limited subset of $\mathbb{R}^n$ (rather than an open set, as in Theorem 2.1). We will return to this point in Section 2.1.3.

We next point out a common misconception concerning the CRB. The unbiased version of the bound is commonly quoted as follows. If an estimator $\hat{\boldsymbol{x}}$ is unbiased for all $\boldsymbol{x} \in \mathcal{X}$, then its MSE is bounded by

$$\mathrm{MSE} \geq \mathrm{Tr}(\boldsymbol{J}^{-1}(\boldsymbol{x})). \tag{2.6}$$

While this statement is correct, the assumption of unbiasedness *for all* $\boldsymbol{x}$ is stronger than required. Indeed, from the statement of Theorem 2.1, it is clear that one requires only that $\boldsymbol{b}(\boldsymbol{x}) = \boldsymbol{0}$ and $\partial \boldsymbol{b}/\partial \boldsymbol{x} = \boldsymbol{0}$ *at a particular value of* $\boldsymbol{x}$ in order for (2.6) to hold for that specific $\boldsymbol{x}$. The importance of this distinction is in the realization that the CRB is a local bound, in the

sense that it is based only on the statistics of the estimation problem in the neighborhood of the parameter $\boldsymbol{x}$ for which the bound is computed. This is manifested both in the requirements of the bound and in its guarantee. In terms of requirements, one must assume only unbiasedness at $\boldsymbol{x}$ and near it (the latter ensured by the specification of the bias gradient). Likewise, the guarantee provided by the CRB is based only on the local statistical properties of the setting (specifically, derivatives of the bias and the pdf).

2.1.2 The Hammersley–Chapman–Robbins Bound

The locality of the unbiasedness assumption will provide the theoretical basis for the derivation of the constrained CRB for sparse settings in Chapter 4. More generally, the locality of the CRB is the main reason for its simple closed form and, consequently, its widespread use. On the downside, though, this property is precisely the cause of the laxity of the bound, and many of the improvements of the CRB result from a deeper analysis based on non-local properties of the estimation setting.

One such improvement is the Hammersley–Chapman–Robbins bound (HCRB), which was derived independently by Hammersley [50] and Chapman and Robbins [51]. Whereas the CRB requires unbiasedness at the point $\boldsymbol{x}$ for which the bound is computed and in its local neighborhood, the HCRB is based on an assumption of unbiasedness at $\boldsymbol{x}$ and another arbitrary point $\boldsymbol{x}+\boldsymbol{\delta}$, which need not necessarily be close to $\boldsymbol{x}$. Ordinarily, the HCRB is applied to estimators which are assumed to be unbiased for *all* values of $\boldsymbol{x}$, and consequently the bound holds regardless of the choice of $\boldsymbol{\delta}$. One can therefore compute the bound for any $\boldsymbol{\delta}$ and choose the highest value thereof as the tightest possible HCRB.

In the following, we state the HCRB for a scalar parameter x [50,51]. A vector version of this bound is slightly more complicated, requiring the use of multiple test points. Such a derivation can be found in the work of Gorman and Hero [52].

Theorem 2.2 (Hammersley–Chapman–Robbins bound)**.** *Let $\boldsymbol{y}$ be a measurement vector whose pdf is given by $p(\boldsymbol{y};x)$ for $x \in \mathcal{X} \subseteq \mathbb{R}$. Suppose*[1] *that the set $\{\boldsymbol{y} : p(\boldsymbol{y};x) > 0\}$ is independent of x. Then, any unbiased estimator $\hat{x}$ of x satisfies*

$$\mathrm{MSE}(\hat{x}) \geq \sup_{\delta : x+\delta \in \mathcal{X}} \frac{\delta^2}{E\left\{\left(\frac{p(\boldsymbol{y};x+\delta)}{p(\boldsymbol{y};x)} - 1\right)^2\right\}}. \tag{2.7}$$

[1] This requirement can be replaced by a weaker, but more cumbersome, assumption on the support of $p(\boldsymbol{y};x)$ [6, p. 114].

The point $\boldsymbol{x} + \boldsymbol{\delta}$ is referred to as a *test point*, a concept which will reappear when discussing Bayesian bounds of the Weiss–Weinstein family in Section 2.2. As explained in the beginning of Section 2.1, the HCRB can easily be formulated for estimators having an arbitrary pre-specified bias function, but this will not be necessary for our purposes.

A striking feature of Theorem 2.2 when compared with the CRB is that the former requires far fewer regularity conditions. This, too, is a consequence of the non-locality of the HCRB. Indeed, the CRB exploits the local properties of the setting through the use of derivatives of the pdf and the bias. The many regularity conditions of the CRB result from the need to carefully handle the limits implied by the differentiation operation. By contrast, the HCRB in its basic form requires only two points of unbiasedness, $\boldsymbol{x}$ and $\boldsymbol{x} + \boldsymbol{\delta}$, and no limits are involved in the derivations. As an example of this advantage of the HCRB, note that settings involving the uniform distribution, such as $\boldsymbol{y} \sim U[x, x+1]$, cannot be analyzed with the CRB but are readily bounded by the HCRB.

It turns out that even more can be said: When a given estimation problem satisfies the regularity conditions of Theorem 2.1, the CRB is obtained as a special case of the HCRB. This can be seen in the one-dimensional case by taking the limit of (2.7) as $\delta \to 0$ and swapping the order of the expectation and limit, whereupon we obtain

$$\mathrm{MSE}(\hat{x}) \geq \frac{1}{E\left\{\left(\frac{1}{p(\boldsymbol{y};x)}\frac{\partial p(\boldsymbol{y};x)}{\partial x}\right)^2\right\}} = \frac{1}{J(x)} \tag{2.8}$$

where we used the definition (2.2) of the Fisher information. The resulting bound is precisely the CRB for unbiased estimators of a scalar parameter. Thus, Theorem 2.1 can be thought of as a HCRB with the test point chosen infinitesimally close to $\boldsymbol{x}$, again illustrating the local nature of the CRB.

We note that the HCRB is not the only possible extension of the CRB. Further improvements can be achieved by increasing the number of test points, at the cost of increased computational complexity. This technique is attributed to Barankin [53]. Moreover, the CRB is not even the tightest possible local bound: it can be improved locally by relying on higher-order derivatives of the pdf, as proposed by Bhattacharya [54]. However, for the purposes of this dissertation, familiarity with the CRB and HCRB will suffice.

2.1.3 The Constrained Cramér–Rao Bound

A central theme in our work involves the effect of structural constraints on estimation quality, as reflected in performance bounds. The constrained CRB is the name given to the performance bound when the parameter is known to belong to a structure set $\mathcal{X}$. Work on the constrained CRB was pioneered by Gorman and Hero [52], who considered primarily constraints of the form

$$\mathcal{X} = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{g}(\boldsymbol{x}) = \boldsymbol{0}, \boldsymbol{h}(\boldsymbol{x}) \leq \boldsymbol{0}\} \tag{2.9}$$

where $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^p$ and $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^q$ are continuously differentiable functions. It is assumed that the definition of the set using the functions $\boldsymbol{g}$ and $\boldsymbol{h}$ is not redundant, in the sense that it is impossible to define $\mathcal{X}$ using a smaller number of constraints. It is further assumed that the constraints are consistent, i.e., $\mathcal{X} \neq \varnothing$. We will refer to sets of the form (2.9) as continuously differentiable constraints.

The constrained CRB of Gorman and Hero [52] can now be stated as follows. Note that some of the regularity conditions stated below are omitted from [52], but are indeed necessary for the theorem to hold. These regularity conditions closely follow the requirements for the unconstrained CRB as discussed in Section 2.1.1 above.

Theorem 2.3 (Constrained CRB)**.** *Let $\mathcal{X}$ be a continuously differentiable set of the form* (2.9) *and let $\{p(\boldsymbol{y};\boldsymbol{x})\}_{\boldsymbol{x}\in\mathcal{X}}$ be a family of pdfs having a common support, i.e., the set $\{\boldsymbol{y} : p(\boldsymbol{y};\boldsymbol{x}) > 0\}$ is independent of $\boldsymbol{x}$. Suppose that $p(\boldsymbol{y};\boldsymbol{x})$ and $\log p(\boldsymbol{y};\boldsymbol{x})$ are differentiable with respect to $\boldsymbol{x}$, and assume that differentiation with respect to $\boldsymbol{x}$ and integration over $\boldsymbol{y}$ can be interchanged whenever necessary. Assume that the FIM* (2.2) *is well-defined, finite, and positive definite. Consider a finite-variance estimator $\hat{\boldsymbol{x}}$ having a differentiable bias function $\boldsymbol{b}(\boldsymbol{x})$. Then,*

$$\mathrm{Cov}(\hat{\boldsymbol{x}}) \succeq \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{x}}\right)\left(\boldsymbol{J}^{-1} - \boldsymbol{J}^{-1}\boldsymbol{G}^T(\boldsymbol{G}\boldsymbol{J}^{-1}\boldsymbol{G}^T)^{\dagger}\boldsymbol{G}\boldsymbol{J}^{-1}\right)\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{x}}\right)^T \tag{2.10}$$

where $\boldsymbol{G} \triangleq \partial \boldsymbol{g}/\partial \boldsymbol{x}$, and $\boldsymbol{G} = \boldsymbol{0}$ when no equality constraints are present.

Several key points arise from an examination of Theorem 2.3. Most strikingly, the bound (2.10) is completely indifferent to the existence of inequality constraints in the set $\mathcal{X}$. For example, if it is known that $\|\boldsymbol{x}\|_2 \leq \tau$, the resulting CRB is identical to the unconstrained bound, even if τ is very small. While a mathematical discussion of this result is given in [52], we believe that a more intuitive explanation arises from an analysis of the bias restrictions in the constrained CRB, which will be presented in Chapter 4. Whatever the interpretations, all versions of the

constrained CRB are insensitive to inequality constraints, even when the bound is computed at the boundary of the constraint set. Indeed, later studies of the constrained CRB often simply assume that no inequality constraints are present, as these have no effect on the bound [55–58].

Another observation follows from a comparison between the constrained bound (2.10) and the unconstrained CRB (2.3), whence it is evident that the constraints reduce the value of the bound. On a certain level, this seems unsurprising: More information is available concerning $\boldsymbol{x}$, so it should be easier to estimate it. This fact is stated as an intuitive argument for the power of the constrained CRB by Gorman and Hero [52, Remark 4]. On deeper inspection, however, the bound reduction is somewhat mysterious. Recall that the CRB is a bound on the performance of all unbiased estimators at the specific point $\boldsymbol{x}$. The bound thus applies to techniques designed specifically for $\boldsymbol{x}$ (as long as these remain unbiased). Such techniques should not benefit from the much less precise knowledge that $\boldsymbol{x} \in \mathcal{X}$. How, then, can knowledge of a constraint set $\mathcal{X}$ reduce the CRB? The answer lies in the class of estimators which are considered unbiased. As we will see in Chapters 3 and 4, unbiasedness in the constrained context has a wider meaning than the ordinary definition in the unconstrained case. Thus, the constrained CRB applies to a wider range of techniques, and this yields a reduction in the value of the bound.

Among the many papers following the work of [52], we mention two which are particularly relevant for our applications. The first is an analysis by Stoica and Ng [56] which examined the case of a singular FIM. Recall that both Theorem 2.1 and Theorem 2.3 assumed that $\boldsymbol{J}(\boldsymbol{x})$ is positive definite. As mentioned previously, in the unconstrained case, the invertibility of $\boldsymbol{J}(\boldsymbol{x})$ is necessary for the existence of finite-variance unbiased estimators [49]. However, in the constrained case, it is sometimes possible for unbiased estimators to exist even when the FIM is singular [56]. Intuitively, one can think of a singular FIM as indicative of a complete lack of measurement data concerning some parameters (or some linear combination of parameters). If the constraints provide the missing information, then the singularity of the FIM is no longer an obstacle. Chapter 3 provides a more detailed analysis of this situation, in which the interplay between bias, constraints, and singularity of the FIM is explored.

An alternative perspective on the constrained CRB is obtained by examining a reparametrization of the problem [52, 57]. Consider a continuously differentiable constraint set $\mathcal{X}$ containing p equality constraints (as we have seen, the inequality constraints can be ignored as they have no effect on the CRB). The constraint set can be locally parameterized using $n - p$ coordinates. The idea of reparametrization is to construct a CRB for the estimation of these $n - p$ coordinates, and transform the result into a bound on the original parameters. As it turns out,

the result is identical to the constrained CRB of Theorem 2.3. Thus, in theory, any result on the CRB for continuously differentiable constraints can equivalently be obtained using the classical CRB and a reparametrization. In practice, there are cases in which the reparametrization is quite cumbersome. Nonetheless, this technique is often useful, and becomes particularly powerful when the dimension of the parameter to be estimated is much higher than the dimension of the reparameterized coordinates. An extreme example of such a construction will occur in the case of bounds for FRI signals (Chapter 8), where the parameter is a continuous-time function (and thus belongs to an infinite-dimensional space), but the coordinates determining it belong to a finite-dimensional space. In this case we will utilize the reparametrization scheme in order to derive the constrained CRB.

As a final comment, we note that while continuously differentiable constraints of the form (2.9) describe many useful types of signal structures, the sparsity models of Sections 1.1 and 1.2 cannot be described in this manner. Thus, our derivation in Chapter 4 of the CRB for sparse estimation will require an extension of the CRB to non-differentiable constraints.

2.2 Bayesian Bounds

In this section, we adopt the Bayesian point of view, i.e., we assume that both $\boldsymbol{x}$ and $\boldsymbol{y}$ are random vectors whose joint pdf $p(\boldsymbol{y},\boldsymbol{x})$ is known. We denote by $p_{\boldsymbol{x}}(\boldsymbol{x})$ the prior pdf of $\boldsymbol{x}$, and by $p_{\boldsymbol{y}|\boldsymbol{x}}(\boldsymbol{y}|\boldsymbol{x})$ the conditional distribution of $\boldsymbol{y}$ given $\boldsymbol{x}$.

We are given a realization of $\boldsymbol{y}$ and wish to use it to estimate $\boldsymbol{x}$. In this case, the MSE $E\{\|\boldsymbol{x}-\hat{\boldsymbol{x}}\|^2\}$ is a scalar, rather than a function of $\boldsymbol{x}$, since the expectation is taken over both $\boldsymbol{x}$ and $\boldsymbol{y}$. Consequently, one can speak of the estimator minimizing the MSE. As is well known, this minimum MSE (MMSE) estimator is the conditional expectation, $\hat{\boldsymbol{x}}_{\text{MMSE}} = E\{\boldsymbol{x}|\boldsymbol{y}\}$.

While the conditional expectation provides the optimal performance in terms of MSE, in many practical cases it is difficult to calculate $\hat{\boldsymbol{x}}_{\text{MMSE}}$ in practice. There are various alternatives to the MMSE estimator which attempt to provide a suboptimal estimator whose computational cost is lower. Given such an alternative technique, one would like to know how much is lost with respect to the optimum performance. Unfortunately, determining the MSE of the conditional expectation is usually no easier than obtaining the MMSE estimate itself. Consequently, there is an interest in easily computable lower bounds on the MMSE performance. If a practical technique comes close to such a lower bound, then the proposed approach is known to also approach the performance of the MMSE method. We emphasize that contrary to the fre-

quentist case, Bayesian bounds are intended merely to provide a computational benefit, since the precise performance of the MMSE estimator can be accurately determined given sufficient computing power.

In this section, we review the most commonly used Bayesian performance bounds. The optimal bias bound (OBB) is described in detail in Chapter 9 and is thus omitted from the discussion herein. We focus in particular on the effectiveness of these bounds when the prior distribution $p_{\boldsymbol{x}}(\boldsymbol{x})$ is constrained, i.e., when there exists a bounded set S such that $p_{\boldsymbol{x}}(\boldsymbol{x}) = 0$ whenever $\boldsymbol{x} \notin S$. Such priors form a natural adaptation of the constraint sets of Section 2.1 to the Bayesian world, and will continue to play a central role in our Bayesian results in Part II of the dissertation.

2.2.1 Bayesian CRB

The Bayesian CRB, which is due to Van Trees [59, pp. 72–73], follows the derivation of the ordinary (frequentist) CRB. For simplicity of notation, in this subsection we assume that the parameter x to be estimated is a scalar; a vector extension of this bound is straightforward.

Theorem 2.4 (Bayesian CRB)**.** *Assume the following regularity conditions:*

1. *The first and second derivatives of $p(\boldsymbol{y}, x)$ with respect to x exist for all $x \in \mathbb{R}$, and the derivatives are absolutely integrable with respect to x and $\boldsymbol{y}$.*

2. *It holds that*

$$\lim_{x \to -\infty} E\{\hat{x} - x | x\}\, p_x(x) = \lim_{x \to +\infty} E\{\hat{x} - x | x\}\, p_x(x) = 0. \tag{2.11}$$

Any estimator $\hat{x}$ then satisfies

$$E\big\{(\hat{x} - x)^2\big\} \geq \left(E\left\{ \left(\frac{\partial \log p(\boldsymbol{y}, x)}{\partial x} \right)^2 \right\} \right)^{-1}. \tag{2.12}$$

Note the similarity of (2.12) to the ordinary unbiased CRB (2.6). The second regularity condition in Theorem 2.4 is the analog of the unbiasedness requirement for the deterministic CRB. However, while in the deterministic case unbiasedness excludes many interesting estimators, in the Bayesian case the condition (2.11) is usually not difficult to satisfy. This is because most probability distributions tend to zero very quickly as $x \to \pm\infty$, so that (2.11) typically holds for all estimators except those with exceptional deviations from the true parameter x.

That being said, the Bayesian CRB is typically inapplicable to situations where the prior $p_x(x)$ is constrained, in the sense that $p_x(x) = 0$ for values of x outside of a bounded set S. This

is because oftentimes the joint pdf will not be sufficiently smooth to satisfy the first regularity condition. For example, a uniform prior will result in a discontinuous pdf, which is clearly not differentiable.

As an alternative, one can derive the Bayesian CRB for a constrained set $x \in (a, b)$, where $a, b \in \mathbb{R}$. This yields a bound which is analogous to Theorem 2.4, with minor modifications to the regularity conditions. Specifically, the first condition now requires the derivatives to exist only in the range (a, b), which is satisfied by most constrained distributions. However, the second regularity condition becomes

$$\lim_{x \to a^+} E\{\hat{x} - x|x\}\, p_x(x) = \lim_{x \to b^-} E\{\hat{x} - x|x\}\, p_x(x) = 0. \tag{2.13}$$

Unfortunately, in many constrained cases (e.g., under a uniform prior), p_x will be nonzero on the boundaries a and b. The bound then applies only to estimators satisfying

$$E\{\hat{x} - x|x = a\} = E\{\hat{x} - x|x = b\} = 0. \tag{2.14}$$

However, it can be shown that any technique satisfying this requirement must produce estimates outside the range $[a, b]$ with nonzero probability. But the MMSE approach never produces estimates whose prior probability is zero. It follows that the bound is not applicable to the MMSE technique. In other words, we have obtained a lower bound on the MSE of *some* estimators, but not on the MSE of the optimal approach.

As mentioned above, Theorem 2.4 is derived in a manner analogous to the frequentist CRB. It was pointed out by Weinstein and Weiss [60] that it is similarly possible to obtain a Bayesian analog of the Bhattacharya bound, which is a tighter version of the CRB. However, this Bayesian Bhattacharya bound is also inapplicable to constrained estimation, as it too fails to meet the regularity conditions outlined above.

2.2.2 The Weinstein–Weiss Family of Bounds

As shown in Section 2.1.2, in the frequentist setting the CRB can be generalized by replacing the use of derivatives with the more general concept of test points, resulting in the HCRB. The elimination of derivatives also reduces the regularity conditions required by the HCRB. In an analogous manner, one can replace the derivatives in the Bayesian CRB with test points, resulting in the Weiss–Weinstein bound, below. This bound again essentially requires no regularity conditions, and is therefore appealing for situations in which the prior distribution is constrained to a bounded set. As before, we provide the bound for the case of a scalar parameter x.

Theorem 2.5 (Weiss–Weinstein bound [61]). *Let $s \in (0,1)$ and $h \in \mathbb{R}$ be two arbitrary constants. Then, any estimator $\hat{x}$ satisfies*

$$E\{(\hat{x}-x)^2\} \geq \frac{h^2 E^2\{L^{1-s}(\boldsymbol{y}; x-h, x)\}}{E\{(L^s(\boldsymbol{y}; x+h, x) - L^{1-s}(\boldsymbol{y}; x-h, x))^2\}} \tag{2.15}$$

where

$$L(\boldsymbol{y}; x_1, x_2) \triangleq \frac{p(\boldsymbol{y}, x_1)}{p(\boldsymbol{y}, x_2)}. \tag{2.16}$$

The points $x+h$ and $x-h$ are referred to as test points, and are analogous to the test points of the HCRB presented in Section 2.1.2. Since Theorem 2.5 holds for any s and h, ideally one would like to maximize over these two parameters so as to obtain the tightest possible bound; however, this can be computationally expensive. In many situations, the tightest bound is obtained for $s = 1/2$, and this value is sometimes used (even when it is not necessarily optimal) in order to reduce the computational load. In the limit as $h \to 0$, the bound tends to the Bayesian CRB, which again illustrates the locality of the CRB (see Section 2.1.1).

A useful extension of Theorem 2.5 is based on adding multiple test points. This generally tightens the bound, but increases the number of parameters which must be optimized. Eventually, adding too many test points will make the bound less practical than directly calculating the MMSE by numerical integration.

The extension to the vector case is straightforward. In the vector case, in order to obtain nontrivial results, one must use at least n test points, where n is the dimension of $\boldsymbol{x}$; furthermore, the test points must span the space $\mathbb{R}^n$. Thus, a common approach to reduce the complexity of the optimization is to choose test points $\{h\boldsymbol{e}_1, \ldots, h\boldsymbol{e}_n\}$, where the $\boldsymbol{e}_i$ are unit vectors. This leaves us with a single parameter h to optimize.

Using a similar technique, one can further generalize the Weiss–Weinstein bound, as shown below.

Theorem 2.6 (Weinstein–Weiss bound [60]). *Consider an integer m and assume we are given a set of functions $\psi_1(\boldsymbol{y}, \boldsymbol{x}), \ldots, \psi_m(\boldsymbol{y}, \boldsymbol{x})$ satisfying*

$$E\{\psi_i(\boldsymbol{y}, \boldsymbol{x}) | \boldsymbol{y}\} = 0 \quad \textit{for all } \boldsymbol{y}. \tag{2.17}$$

We then have the lower bound

$$E\left\{(\hat{\boldsymbol{x}} - \boldsymbol{x})(\hat{\boldsymbol{x}} - \boldsymbol{x})^T\right\} \geq \boldsymbol{V}\boldsymbol{P}^{-1}\boldsymbol{V}^T \tag{2.18}$$

where

$$V_{ij} = E\{x_i \psi_j(\boldsymbol{y}, \boldsymbol{x})\} \tag{2.19}$$

and

$$P_{ij} = E\{\psi_i(\boldsymbol{y}, \boldsymbol{x})\psi_j(\boldsymbol{y}, \boldsymbol{x})\}. \tag{2.20}$$

Intriguingly, many Bayesian bounds can be derived as special cases by an appropriate choice of functions $\psi_i(\boldsymbol{y}, \boldsymbol{x})$. These bounds include the aforementioned Bayesian CRB, Bayesian Bhattacharya, and Weiss–Weinstein bounds, which are therefore said to belong to the Weinstein–Weiss family of lower bounds [62].

The challenge in Theorem 2.6 is, of course, to find functions $\psi_i(\boldsymbol{y}, \boldsymbol{x})$ which satisfy (2.17) but yield new, hopefully tighter, bounds. This challenge can be addressed by exploiting a geometrical interpretation of Bayesian lower bounds [63, 64]. Specifically, consider the Hilbert space L_2 of finite-variance functions of both $\boldsymbol{x}$ and $\boldsymbol{y}$. The random variable representing the true parameter value $\boldsymbol{x}$ is an element of L_2, and the MMSE estimator $\hat{\boldsymbol{x}}_{\text{MMSE}}$ is the projection of $\boldsymbol{x}$ onto the subspace $\mathcal{Y}$ of L_2 containing only the functions of $\boldsymbol{y}$. Therefore, the estimation error is equal to the variance of the projection of $\boldsymbol{x}$ onto the orthogonal complement $\mathcal{Y}^\perp$ of $\mathcal{Y}$. Computing this projection is typically difficult, being in fact merely a geometrical statement of the requirement of computing $\hat{\boldsymbol{x}}_{\text{MMSE}}$. However, one can obtain lower bounds on the size of the MMSE error by projecting $\boldsymbol{x}$ onto any subspace $\mathcal{H}$ of $\mathcal{Y}^\perp$. Any such projection is a lower bound on the minimum MSE achievable in the given estimation setting. If the subspace $\mathcal{H}$ is low-dimensional, then computing the projection becomes simple, at least numerically.

It turns out that the Weinstein–Weiss family can be viewed precisely as projections onto specific subspaces of $\mathcal{Y}^\perp$ [63, 64]. Indeed, the requirement (2.17) is precisely a statement of the fact that $\psi_i \in \mathcal{Y}^\perp$. The lower bound of Theorem 2.6 is then the Bayesian bound corresponding to choosing the subspace

$$\mathcal{H} = \text{span}\{\psi_1, \ldots, \psi_m\} \subset \mathcal{Y}^\perp. \tag{2.21}$$

This observation can be used to great advantage in designing Hilbert subspaces $\mathcal{H}$ for specific types of signals. This is done by performing a transform on known functions $\{\psi_i(\boldsymbol{y}, \boldsymbol{x})\}$, such that the energy of the projection is concentrated in a small number of transformed coordinates. A lower bound can then be constructed by choosing those components having large magnitude. For example, measurements for which the energy is concentrated in specific frequency components can be exploited by performing a DFT on a standard set of test points, resulting in a bound which is substantially higher than previous bounds in the Weinstein–Weiss family [64].

2.2.3 The Ziv–Zakai Family of Bounds

The Ziv–Zakai bound was developed originally for the specific case of an unknown scalar parameter x which is distributed uniformly over $[0, T]$ [65]. The bound has since been substantially tightened, and also extended to vector parameters with arbitrary prior densities, albeit with some loss of the elegance of the original idea.

Given measurements $\boldsymbol{y}$, the bound is based on the following observation. Suppose it is known that x equals either x_1 or x_2, and that these two possibilities occur with equal probability. Then, an optimal detection approach would be to test the hypothesis $x = x_1$ vs. $x = x_2$. Let $P_{e,\min}(x_1, x_2)$ be the probability of error of this optimal technique. Alternatively, we can also use a suboptimal approach which is based on first estimating the value of x from the measurements, and then selecting x_1 or x_2 based on whichever is closer to $\hat{x}$. This will result in a particular error probability $P_e \geq P_{e,\min}$, which is given by

$$P_{e,\min} \leq P_e = \tfrac{1}{2}\Pr\Big\{\hat{x} - x_1 \geq \tfrac{\Delta}{2}\Big| x_1\Big\} + \tfrac{1}{2}\Pr\Big\{x_2 - \hat{x} \leq \tfrac{\Delta}{2}\Big| x_2\Big\} \tag{2.22}$$

where $\Delta = x_2 - x_1$.

The right-hand side of (2.22) can be bounded by using the Chebyshev inequality, which converts the probability terms into MSE terms. This results in a lower bound on the MSE which is based on the optimal probability of error. The bound holds for any two values of x_1 and x_2, so that obtaining the tightest bound requires maximizing over these two functions. In fact, it turns out that one need only optimize over their difference $\Delta = x_2 - x_1$. Working out the math, the Ziv–Zakai bound is given by

$$E\big\{(x - \hat{x})^2\big\} \geq \max_{\Delta \in [0,T]} \frac{\Delta^2}{4T} \int_0^{T-\Delta} P_{e,\min}(x, x + \Delta) dx. \tag{2.23}$$

The bound is loose in two places: First, in the transition from an optimal detection problem to a suboptimal approach based on estimation; and, second, in the application of the Chebyshev inequality. This latter point was improved by Chazan, Ziv, and Zakai [66] and by Bellini and Tartara [67]. Their approach avoids the use of Chebyshev's inequality, resulting in a bound which is higher by at least a factor of 2.

Bell et al. [68] extended the Bellini–Tartara bound to vector parameters and to arbitrary prior pdfs. This increases the complexity of the optimization problem, since in general one must now optimize over an n-dimensional vector $\boldsymbol{\Delta}$. They also demonstrated tightness at low SNR for a particular family of Gaussian-like location problems. However, it is unknown whether this

bound is tight for general distributions at low SNR, nor is tightness guaranteed for high SNR values.

Part I

Structural Constraints in the Frequentist Setting

Chapter 3

The CRB with Singular Fisher Information

This chapter is a reprint of the paper:

- Z. Ben-Haim and Y. C. Eldar, "On the constrained Cramér–Rao bound with a singular Fisher information matrix," *IEEE Signal Processing Letters*, vol. 16, no. 6, June 2009, pp. 453–456.

3.1 Introduction

A central goal in statistics and signal processing is to estimate unknown deterministic parameters from random measurements. The performance of estimators in such a setting is circumscribed by the well-known Cramér–Rao bound (CRB) [69]. Specifically, the CRB provides a lower limit on the variance obtainable by any technique as a function of the Fisher information matrix (FIM) and the estimator's bias gradient.

A variant of the CRB for constrained estimation problems was developed by Gorman and Hero [52]. They considered the setting in which the parameter vector belongs to a known set. When this information is incorporated into the estimator, performance can be improved. As a consequence, the constrained CRB can be lower than the unconstrained version.

The derivation of Gorman and Hero assumed that the FIM is positive definite. Stoica and Ng [56] later extended the constrained CRB to the case in which the FIM is positive *semi*-definite, and may thus be singular. In an unconstrained problem, a singular FIM implies that unbiased estimation of the entire parameter vector is impossible [49]. However, Stoica and Ng

demonstrated that, in some cases, one can obtain so-called constrained unbiased estimators, which are unbiased as long as the constraints hold.

The work of Stoica and Ng considers only unbiased estimation. Yet even when unbiased methods do not exist in a particular setting, biased techniques can still be found. As we will demonstrate, when the FIM is singular, estimators can be constructed by introducing a sufficient number of constraints, by specifying an appropriate bias function, or by a combination thereof.

More specifically, in this letter we generalize the above-mentioned bounds and obtain a biased CRB for constrained estimation with a positive semi-definite FIM. When an estimator achieving the CRB exists, we provide a closed form for it. We further derive a necessary and sufficient condition for the CRB to be infinite, indicating that no estimator exists in the given setting.

The following notation is used throughout the letter. Given a vector function $\boldsymbol{f} : \mathbb{R}^n \to \mathbb{R}^k$, we denote by $\partial \boldsymbol{f}/\partial \boldsymbol{\theta}$ the $k \times n$ matrix whose ijth element is $\partial f_i/\partial \theta_j$. Also, $\mathcal{R}(\boldsymbol{A})$, $\mathcal{N}(\boldsymbol{A})$, and $\boldsymbol{A}^\dagger$ are, respectively, the range space, null space, and Moore–Penrose pseudoinverse of a matrix $\boldsymbol{A}$, and $\mathcal{S}^\perp$ denotes the orthogonal complement of the subspace $\mathcal{S}$. Finally, $\boldsymbol{A} \succeq \mathcal{B}$ indicates that $\boldsymbol{A} - \mathcal{B}$ is positive semi-definite.

3.2 Problem Statement

Let $\boldsymbol{y}$ be a measurement vector with pdf $p(\boldsymbol{y}; \boldsymbol{\theta})$, for some deterministic unknown parameter vector $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^n$. Suppose that $p(\boldsymbol{y}; \boldsymbol{\theta})$ is differentiable with respect to $\boldsymbol{\theta}$. The FIM $\boldsymbol{J}(\boldsymbol{\theta})$ is then defined as

$$\boldsymbol{J}(\boldsymbol{\theta}) = E\Big\{\boldsymbol{\Delta}\boldsymbol{\Delta}^T\Big\} \tag{3.1}$$

where

$$\boldsymbol{\Delta} = \frac{\partial \log p(\boldsymbol{y}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}. \tag{3.2}$$

We assume throughout that $\boldsymbol{J}(\boldsymbol{\theta})$ is finite for all $\boldsymbol{\theta} \in \Theta$.

Suppose that $\boldsymbol{\theta}$ is known to belong to a constraint set

$$\Theta = \{\boldsymbol{\theta} \in \mathbb{R}^n : \boldsymbol{f}(\boldsymbol{\theta}) = \boldsymbol{0}\} \subseteq \mathbb{R}^n \tag{3.3}$$

where $\boldsymbol{f} : \mathbb{R}^n \to \mathbb{R}^k$ is a continuously differentiable function of $\boldsymbol{\theta}$ with $0 \leq k \leq n$. Note that we are assuming for simplicity that no inequality constraints are present, as it has been shown that such constraints have no effect on the CRB [52].

We further assume that the $k \times n$ matrix $\boldsymbol{F}(\boldsymbol{\theta}) = \partial \boldsymbol{f} / \partial \boldsymbol{\theta}$ has full row rank, which is equivalent to requiring that the constraints are not redundant. Thus, there exists an $n \times (n-k)$ matrix $\boldsymbol{U}(\boldsymbol{\theta})$ such that

$$\boldsymbol{F}(\boldsymbol{\theta})\boldsymbol{U}(\boldsymbol{\theta}) = \boldsymbol{0}, \qquad \boldsymbol{U}^T(\boldsymbol{\theta})\boldsymbol{U}(\boldsymbol{\theta}) = \boldsymbol{I}. \tag{3.4}$$

Intuitively, $\mathcal{R}(\boldsymbol{U}(\boldsymbol{\theta}))$ is the set of feasible directions at $\boldsymbol{\theta}$, i.e., the set of directions in which an infinitesimal change does not violate the constraints. For notational simplicity, in the sequel we will omit the dependence of $\boldsymbol{U}$ and $\boldsymbol{J}$ on $\boldsymbol{\theta}$.

Let $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}(\boldsymbol{y})$ be an estimator of $\boldsymbol{\theta}$. We are interested in the performance of $\hat{\boldsymbol{\theta}}$ under the assumption that $\boldsymbol{\theta} \in \Theta$. Specifically, we derive a lower bound on the covariance matrix

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) = E\Big\{ \left(\hat{\boldsymbol{\theta}} - E\{\hat{\boldsymbol{\theta}}\}\right) \left(\hat{\boldsymbol{\theta}} - E\{\hat{\boldsymbol{\theta}}\}\right)^T \Big\} \tag{3.5}$$

obtainable by any estimator $\hat{\boldsymbol{\theta}}$. The covariance matrix, as well as the CRB, are a function of $\boldsymbol{\theta}$; we are interested in bounding this matrix for all $\boldsymbol{\theta} \in \Theta$. To obtain a nontrivial bound, we assume that the desired bias $\boldsymbol{b}(\boldsymbol{\theta}) = E\{\hat{\boldsymbol{\theta}}\} - \boldsymbol{\theta}$ is specified for $\boldsymbol{\theta} \in \Theta$; the bias for $\boldsymbol{\theta} \notin \Theta$ is arbitrary.

Previous work on the constrained estimation setting [52, 56] assumed that the estimator $\hat{\boldsymbol{\theta}}$ satisfies the constraint $\hat{\boldsymbol{\theta}} \in \Theta$. However, it turns out that this requirement can be removed without altering the resulting bound. Furthermore, in some cases, the CRB can only be achieved by estimators violating the constraint. In this letter, the term "constrained estimator" refers to the situation in which the bias $\boldsymbol{b}(\boldsymbol{\theta})$ is specified only for $\boldsymbol{\theta} \in \Theta$, and the performance is evaluated when the true parameter value $\boldsymbol{\theta}$ belongs to the set Θ. The implications of this setting are discussed further in the next section.

3.3 Cramér–Rao Bound

3.3.1 Main Result

With the concepts developed in the previous section, our main result can be stated as follows.

Theorem 3.1. *Let Θ be a constraint set of the form* (3.3) *with a corresponding matrix $\boldsymbol{U}$ of* (3.4). *Let $\hat{\boldsymbol{\theta}}$ be an estimator of $\boldsymbol{\theta}$ whose bias is given by $\boldsymbol{b}(\boldsymbol{\theta})$ for all $\boldsymbol{\theta} \in \Theta$, and define*

$$\boldsymbol{A} = \boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}. \tag{3.6}$$

Assume that integration with respect to $\boldsymbol{y}$ *and differentiation with respect to* $\boldsymbol{\theta}$ *can be interchanged,*[1] *and suppose that*

$$\mathcal{R}\Big(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T\Big) \subseteq \mathcal{R}\Big(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\Big). \tag{3.7}$$

Then, the covariance of $\hat{\boldsymbol{\theta}}$ *satisfies*

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) \succeq \boldsymbol{A}\boldsymbol{U}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^{\dagger}\boldsymbol{U}^T\boldsymbol{A}^T \quad \textit{for all } \boldsymbol{\theta} \in \Theta. \tag{3.8}$$

Equality is achieved in (3.8) *if and only if*

$$\hat{\boldsymbol{\theta}} = \boldsymbol{b}(\boldsymbol{\theta}) + \boldsymbol{\theta} + \boldsymbol{A}\boldsymbol{U}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^{\dagger}\boldsymbol{U}^T\boldsymbol{\Delta} \tag{3.9}$$

in the mean square sense, for all $\boldsymbol{\theta} \in \Theta$*. Here,* $\boldsymbol{\Delta}$ *is given by* (3.2)*. Conversely, if* (3.7) *does not hold, then there exists no finite-variance estimator with the given bias function.*

It is illuminating to examine the influence of the constraints on the bound of Theorem 3.1. Recall that the CRB is a bound on the covariance of *all* estimators having a given bias function, at each specific point $\boldsymbol{\theta}$. The bound thus applies even to estimators which are designed for the specific point $\boldsymbol{\theta}$, a far more restrictive assumption than the knowledge that $\boldsymbol{\theta} \in \Theta$. How, then, can one expect to obtain a meaningful performance bound by imposing the constraint set Θ?

The answer stems from the fact that the bias is specified in Theorem 3.1 only for $\boldsymbol{\theta} \in \Theta$. For example, consider constrained unbiased estimators, for which $\boldsymbol{b}(\boldsymbol{\theta}) = \boldsymbol{0}$ for all $\boldsymbol{\theta} \in \Theta$; the bias when $\boldsymbol{\theta} \notin \Theta$ is irrelevant and unspecified. This is a far larger class of estimators than those which are unbiased for all $\boldsymbol{\theta} \in \mathbb{R}^n$. Consequently, the bound (3.8) is lower than the unconstrained CRB. The weakened bias specification is apparent in Theorem 3.1 from the fact that the matrix $\boldsymbol{A}$ only appears when multiplied by $\boldsymbol{U}$, which nullifies components in directions violating the constraints. Indeed, to calculate the bound, $\boldsymbol{A}$ only needs to be specified in directions consistent with Θ. This issue will be discussed further in a forthcoming paper [70].

Condition (3.7) succinctly describes the possibilities for estimation under various values of the FIM. If $\boldsymbol{J}$ is invertible, then (3.7) holds regardless of the constraint set and the bias gradient, implying that the CRB is always finite. The situation is more complicated when $\boldsymbol{J}$ is singular. In this case, one option is to choose a matrix $\boldsymbol{A}$ whose null space includes $\mathcal{N}(\boldsymbol{J})$; this implies that the estimator is insensitive to changes in elements of $\boldsymbol{\theta}$ for which there is no information. Another option is to provide external constraints for the unmeasurable elements of $\boldsymbol{\theta}$, thus

[1] This condition basically requires that the bounds of $p(\boldsymbol{y};\boldsymbol{\theta})$ do not depend on $\boldsymbol{\theta}$. Such regularity conditions are assumed in all forms of the CRB.

changing $\boldsymbol{U}$ in such a way as to ensure the validity of (3.7) for all $\boldsymbol{A}$. An example comparing these approaches will be presented in Section 3.4.

Theorem 3.1 encompasses several previous results as special cases. Most famously, when $\boldsymbol{J}$ is nonsingular and no constraints are imposed, we obtain the standard CRB

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) \succeq \boldsymbol{A}\boldsymbol{J}^{-1}\boldsymbol{A}^T. \tag{3.10}$$

Several prior extensions [49, 52, 56] of (3.10) are also special cases of Theorem 3.1.

3.3.2 Proof of Theorem 3.1

The proof of Theorem 3.1 is based on the following lemmas.

Lemma 3.2. *Assuming that integration with respect to $\boldsymbol{y}$ and differentiation with respect to $\boldsymbol{\theta}$ can be interchanged, we have*

$$E\Big\{ \left(\hat{\boldsymbol{\theta}} - E\{\hat{\boldsymbol{\theta}}\}\right) \boldsymbol{\Delta}^T \Big\} = \boldsymbol{A} \tag{3.11}$$

for any estimator $\hat{\boldsymbol{\theta}}$. Here, $\boldsymbol{\Delta}$ is defined by (3.2)*, and $\boldsymbol{A}$ is given by* (3.6)*.*

Proof. The proof is an extension of [55, Th. 1] to the case of a biased estimator. Using (3.2),

$$\begin{aligned} &E\Big\{ \left(\hat{\boldsymbol{\theta}} - E\{\hat{\boldsymbol{\theta}}\}\right) \boldsymbol{\Delta}^T \Big\} \\ &= \int \left(\hat{\boldsymbol{\theta}} - E\{\hat{\boldsymbol{\theta}}\}\right) \frac{1}{p(\boldsymbol{y};\boldsymbol{\theta})} \frac{\partial p(\boldsymbol{y};\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} p(\boldsymbol{y};\boldsymbol{\theta}) d\boldsymbol{y} \\ &= \frac{\partial}{\partial \boldsymbol{\theta}} \int \hat{\boldsymbol{\theta}} p(\boldsymbol{y};\boldsymbol{\theta}) d\boldsymbol{y} - E\{\hat{\boldsymbol{\theta}}\} \frac{\partial}{\partial \boldsymbol{\theta}} \int p(\boldsymbol{y};\boldsymbol{\theta}) d\boldsymbol{y} \end{aligned} \tag{3.12}$$

where we interchanged the order of differentiation and integration, and used the fact that $\hat{\boldsymbol{\theta}}$ is a function of $\boldsymbol{y}$ but not of $\boldsymbol{\theta}$. Noting that the second integral in (3.12) equals 1, we obtain

$$E\Big\{ \left(\hat{\boldsymbol{\theta}} - E\{\hat{\boldsymbol{\theta}}\}\right) \boldsymbol{\Delta}^T \Big\} = \frac{\partial E\{\hat{\boldsymbol{\theta}}\}}{\partial \boldsymbol{\theta}} = \boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \tag{3.13}$$

which completes the proof. □

The following lemma provides a family of bounds on $\mathrm{Cov}(\hat{\boldsymbol{\theta}})$ for any estimator $\hat{\boldsymbol{\theta}}$ having a specified bias function. Theorem 3.1 is obtained by choosing an optimal member from this class.

Lemma 3.3. *Let $\hat{\boldsymbol{\theta}}$ be an estimator of $\boldsymbol{\theta} \in \Theta$, and suppose its bias is $\boldsymbol{b}(\boldsymbol{\theta})$. Under the conditions of Lemma 3.2, for any $n \times n$ matrix $\boldsymbol{W}$, we have*

$$\begin{aligned} \mathrm{Cov}(\hat{\boldsymbol{\theta}}) \succeq\ & \boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T + \boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{W}^T \\ & - \boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{W}^T. \end{aligned} \tag{3.14}$$

Proof of Lemma 3.3. Let $\widetilde{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}} - E\{\hat{\boldsymbol{\theta}}\}$ and note that

$$\begin{aligned}\mathbf{0} &\preceq E\left\{\left(\widetilde{\boldsymbol{\theta}} - \boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{\Delta}\right)\left(\widetilde{\boldsymbol{\theta}} - \boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{\Delta}\right)^T\right\} \\ &= \mathrm{Cov}(\hat{\boldsymbol{\theta}}) - \boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T E\left\{\boldsymbol{\Delta}\widetilde{\boldsymbol{\theta}}^T\right\} - E\left\{\widetilde{\boldsymbol{\theta}}\boldsymbol{\Delta}^T\right\}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{W}^T \\ &\quad + \boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T E\left\{\boldsymbol{\Delta}\boldsymbol{\Delta}^T\right\}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{W}^T. \end{aligned} \tag{3.15}$$

Using (3.1) and Lemma 3.2, we obtain (3.14). □

We recall the following properties of the pseudoinverse, which will be required for some further developments.

Lemma 3.4. *Let $\boldsymbol{M}$ and $\boldsymbol{N}$ be arbitrary matrices and let $\boldsymbol{U}^T\boldsymbol{U} = \boldsymbol{I}$. Then*

$$(\boldsymbol{U}\boldsymbol{M}\boldsymbol{U}^T)^\dagger = \boldsymbol{U}\boldsymbol{M}^\dagger\boldsymbol{U}^T \tag{3.16}$$

$$(\boldsymbol{M}^T\boldsymbol{M})^\dagger = \boldsymbol{M}^\dagger\boldsymbol{M}^{T\dagger} \tag{3.17}$$

$$\boldsymbol{M}^\dagger = \boldsymbol{M}^\dagger\boldsymbol{M}^{T\dagger}\boldsymbol{M}^T \tag{3.18}$$

$$\boldsymbol{M} = \boldsymbol{M}^{T\dagger}\boldsymbol{M}^T\boldsymbol{M} \tag{3.19}$$

$$(\boldsymbol{M}\boldsymbol{N})^\dagger = (\boldsymbol{M}\boldsymbol{N})^\dagger\boldsymbol{M}^{T\dagger}\boldsymbol{M}^T. \tag{3.20}$$

Proof. Proofs for (3.16)–(3.19) can be found in [71, Theorem 1.2.1], while (3.20) can be demonstrated by showing that $(\boldsymbol{M}\boldsymbol{N})^\dagger\boldsymbol{M}^{T\dagger}\boldsymbol{M}^T$ satisfies the Moore–Penrose conditions for the pseudoinverse of $\boldsymbol{M}\boldsymbol{N}$. □

We are now ready to prove the main result.

Proof of Theorem 3.1. Our proof is based on that of Stoica and Ng [56]. Suppose first that (3.7) holds, and let

$$\boldsymbol{W} = \boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T)^\dagger. \tag{3.21}$$

Applying Lemma 3.3 and using the Moore–Penrose condition $\boldsymbol{M}^\dagger\boldsymbol{M}\boldsymbol{M}^\dagger = \boldsymbol{M}^\dagger$, we obtain

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) \succeq \boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T)^\dagger\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T. \tag{3.22}$$

It follows from (3.16) that

$$(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T)^\dagger = \boldsymbol{U}^{T\dagger}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^\dagger\boldsymbol{U}^\dagger = \boldsymbol{U}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^\dagger\boldsymbol{U}^T. \tag{3.23}$$

Substituting this into (3.22) yields (3.8), as required.

We now show that (3.9) holds if and only if

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) = \boldsymbol{A}\boldsymbol{U}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^\dagger\boldsymbol{U}^T\boldsymbol{A}^T \tag{3.24}$$

in the mean squared sense, for all $\boldsymbol{\theta} \in \Theta$. Note first that if $\hat{\boldsymbol{\theta}}$ satisfies (3.9), then the bias of $\hat{\boldsymbol{\theta}}$ is indeed $\boldsymbol{b}(\boldsymbol{\theta})$, since $E\{\boldsymbol{\Delta}\} = \boldsymbol{0}$. Furthermore

$$\mathrm{Cov}(\hat{\boldsymbol{\theta}}) = \boldsymbol{A}\boldsymbol{U}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^\dagger\boldsymbol{U}^T E\left\{\boldsymbol{\Delta}\boldsymbol{\Delta}^T\right\}\boldsymbol{U}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^\dagger\boldsymbol{U}^T\boldsymbol{A}^T \tag{3.25}$$

which yields (3.24). Conversely, suppose that (3.24) holds, and let $\widetilde{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}} - E\{\hat{\boldsymbol{\theta}}\}$. Using Lemma 3.2 and (3.24), it is straightforward to show that

$$\mathrm{Cov}\left(\widetilde{\boldsymbol{\theta}} - \boldsymbol{A}\boldsymbol{U}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^{-1}\boldsymbol{U}^T\boldsymbol{\Delta}\right) = \boldsymbol{0}. \tag{3.26}$$

Therefore, $\hat{\boldsymbol{\theta}} = E\{\hat{\boldsymbol{\theta}}\} + \boldsymbol{A}\boldsymbol{U}(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U})^{-1}\boldsymbol{U}^T\boldsymbol{\Delta}$ in the mean square sense, as required.

It remains to show that if

$$\mathcal{R}\left(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T\right) \nsubseteq \mathcal{R}\left(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\right) \tag{3.27}$$

then no finite variance estimator exists. Suppose that (3.27) holds. Since $\mathcal{R}(\boldsymbol{M}) = \mathcal{N}(\boldsymbol{M}^T)^\perp$ for any matrix $\boldsymbol{M}$, we have $\mathcal{N}(\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T)^\perp \nsubseteq \mathcal{N}(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T)^\perp$, or equivalently, $\mathcal{N}(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T) \nsubseteq \mathcal{N}(\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T)$. Thus, there exists a vector $\boldsymbol{v} \in \mathcal{N}(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T)$ for which $\boldsymbol{v} \notin \mathcal{N}(\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T)$. Now, let $\boldsymbol{W} = \alpha\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{v}\boldsymbol{v}^T$ for some scalar α to be defined below. From Lemma 3.3,

$$\begin{aligned}
&\mathrm{Tr}(\mathrm{Cov}(\hat{\boldsymbol{\theta}})) \\
&\geq 2\,\mathrm{Tr}(\boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T) - \mathrm{Tr}(\boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{W}^T) \\
&= 2\alpha\,\mathrm{Tr}(\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{v}\boldsymbol{v}^T\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T) \\
&\quad - \alpha^2\,\mathrm{Tr}(\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{v}\boldsymbol{v}^T\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{v}\boldsymbol{v}^T\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T).
\end{aligned} \tag{3.28}$$

The second term in (3.28) is zero since $\boldsymbol{v} \in \mathcal{N}(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T)$, whereas the first term equals $2\alpha\|\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{v}\|^2$, which is nonzero since $\boldsymbol{v} \notin \mathcal{N}(\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T)$. Thus, by choosing α appropriately, $\mathrm{Tr}(\mathrm{Cov}(\hat{\boldsymbol{\theta}}))$ can be shown to be larger than any finite number. Therefore, there does not exist a finite-variance estimator with the required bias. □

3.3.3 Choice of $\boldsymbol{W}$

The bound (3.8) of Theorem 3.1 is obtained from the more general Lemma 3.3 by choosing a specific value (3.21) for the matrix $\boldsymbol{W}$. We now show that this choice of $\boldsymbol{W}$ is optimal, in that

it results in the tightest bound obtainable from Lemma 3.3. Note that Lemma 3.3 provides a matrix inequality, so there does not necessarily exist a single maximum value of the bound (because the set of matrices is not totally ordered). However, in our case, such a maximum value does exist and results in the bound of Theorem 3.1.

The method of obtaining $\boldsymbol{W}$ used in [56] does not seem to generalize to the case of biased estimators. Instead, let $\boldsymbol{v}$ be an arbitrary vector in $\mathbb{R}^n$ and observe that

$$\boldsymbol{v}^T \left(\boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T + \boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{W}^T - \boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{W}^T \right) \boldsymbol{v} \tag{3.29}$$

is concave in $\boldsymbol{W}$. Therefore, to maximize (3.29), it suffices to find a point $\boldsymbol{W}$ at which the derivative is zero. Differentiating (3.29) with respect to $\boldsymbol{W}$, we obtain [72]

$$2\boldsymbol{v}\boldsymbol{v}^T\boldsymbol{A}\boldsymbol{U}\boldsymbol{U}^T - 2\boldsymbol{v}\boldsymbol{v}^T\boldsymbol{W}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T. \tag{3.30}$$

Thus, if there exists a matrix $\boldsymbol{W}$ such that

$$\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{W}^T = \boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T \tag{3.31}$$

then that value of $\boldsymbol{W}$ maximizes (3.29) simultaneously for any choice of $\boldsymbol{v}$. Note that (3.31) can be written as a set of n vector equations

$$\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{w}_i = \boldsymbol{u}_i, \quad i = 1, \ldots, n \tag{3.32}$$

where $\boldsymbol{w}_i$ is the ith row of $\boldsymbol{W}$ and $\boldsymbol{u}_i$ is the ith column of $\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{A}^T$. Clearly, (3.32) has a solution $\boldsymbol{w}_i$ if and only if $\boldsymbol{u}_i \in \mathcal{R}(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T)$. This does indeed occur under the condition (3.7) of Theorem 3.1, and one such solution is given by

$$\boldsymbol{w}_i = (\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T)^\dagger \boldsymbol{u}_i, \quad i = 1, \ldots, n. \tag{3.33}$$

Combining these n equations, we obtain that the matrix $\boldsymbol{W}$ chosen in (3.21) simultaneously maximizes (3.29) for all values of $\boldsymbol{v}$. Therefore, the bound of Theorem 3.1 is the tightest bound obtainable from Lemma 3.3.

3.4 Example

As an example of the applicability of Theorem 3.1, we consider an underdetermined linear regression setting. Let $\boldsymbol{\theta} \in \mathbb{R}^n$ be an unknown vector for which measurements $\boldsymbol{y} = \boldsymbol{H}\boldsymbol{\theta} + \boldsymbol{v}$ are available. Here, $\boldsymbol{v}$ is white Gaussian noise with variance σ^2 and $\boldsymbol{H}$ is a known $p \times n$ matrix

with $p < n$. Since there are fewer measurements than parameters, unbiased reconstruction of $\boldsymbol{\theta}$ is clearly impossible without additional assumptions. To see this formally, note that

$$\boldsymbol{J} = \sigma^{-2}\boldsymbol{H}^T\boldsymbol{H}. \tag{3.34}$$

Thus $\text{rank}(\boldsymbol{J}) \leq p < n$, so that the matrix $\boldsymbol{J}$ is singular, and by Theorem 3.1, unconstrained unbiased estimation is impossible. This also follows from earlier results [49, 56].

In order to enable reconstruction of $\boldsymbol{\theta}$, additional assumptions are required. One possibility is to restrict $\boldsymbol{\theta}$ to some subset $\Theta \subset \mathbb{R}^n$, and then seek an unbiased estimator over this set. An alternative is to choose a reasonable value for $E\{\hat{\boldsymbol{\theta}}\}$, taking into account the lack of information. As we will see, both approaches result in the same estimator, but the latter implies optimality under wider conditions.

Beginning with the first approach, let us assume that $\boldsymbol{\theta} = \boldsymbol{W}\boldsymbol{\alpha}$ for a given $n \times k$ matrix $\boldsymbol{W}$ and an unknown $\boldsymbol{\alpha} \in \mathbb{R}^k$. For example, $\boldsymbol{W}$ can define a smoothness requirement on $\boldsymbol{\theta}$. We seek an unbiased estimator for such $\boldsymbol{\theta}$.

Choosing $\Theta = \mathcal{R}(\boldsymbol{W})$ results in $\boldsymbol{U}(\boldsymbol{\theta}) = \boldsymbol{W}$. Thus, it follows from Theorem 3.1 that if there exists a constrained unbiased estimator $\hat{\boldsymbol{\theta}}_c$ which achieves the CRB, then $\hat{\boldsymbol{\theta}}_c$ must satisfy, for all $\boldsymbol{\theta} \in \mathcal{R}(\boldsymbol{W})$,

$$\begin{aligned}
\hat{\boldsymbol{\theta}}_c &= \boldsymbol{\theta} + \boldsymbol{W}(\boldsymbol{W}^T\boldsymbol{H}^T\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{W}^T\boldsymbol{H}^T(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{\theta}) \\
&= \boldsymbol{\theta} + \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger(\boldsymbol{H}\boldsymbol{W})^{T\dagger}(\boldsymbol{H}\boldsymbol{W})^T(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{\theta}) \\
&= \boldsymbol{\theta} + \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{\theta})
\end{aligned} \tag{3.35}$$

where we have used (3.17) in the first transition and (3.18) in the second. Since $\boldsymbol{\theta} \in \mathcal{R}(\boldsymbol{W})$, one may write $\boldsymbol{\theta} = \boldsymbol{W}\boldsymbol{d}$, for some vector $\boldsymbol{d}$. Thus

$$\hat{\boldsymbol{\theta}}_c = \boldsymbol{\theta} + \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{y} - \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}\boldsymbol{W}\boldsymbol{d}. \tag{3.36}$$

Suppose that

$$\mathcal{R}(\boldsymbol{W}) \cap \mathcal{N}(\boldsymbol{H}) = \{\boldsymbol{0}\}. \tag{3.37}$$

In this case, it is readily shown that $\mathcal{N}(\boldsymbol{H}\boldsymbol{W}) = \mathcal{N}(\boldsymbol{W})$, and consequently $\boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}\boldsymbol{W} = \boldsymbol{W}$. Thus

$$\hat{\boldsymbol{\theta}}_c = \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{y} \tag{3.38}$$

is the constrained unbiased estimator achieving the CRB. In other words, $\hat{\boldsymbol{\theta}}_c$ has minimum MSE among all estimators which are unbiased over Θ.

On the other hand, suppose that $\mathcal{R}(\boldsymbol{W}) \cap \mathcal{N}(\boldsymbol{H}) \neq \{\boldsymbol{0}\}$. This implies that the constraints on $\boldsymbol{\theta}$ do not sufficiently compensate for the lack of information in the measurements $\boldsymbol{y}$. Indeed, in this case we have $\mathcal{R}\big(\boldsymbol{W}\boldsymbol{W}^T\boldsymbol{H}^T\boldsymbol{H}\boldsymbol{W}\boldsymbol{W}^T\big) \subsetneq \mathcal{R}\big(\boldsymbol{W}\boldsymbol{W}^T\big)$, and it follows from Theorem 3.1 that no unbiased estimator exists. These conclusions can also be obtained from [56].

Observe that the expectation of $\hat{\boldsymbol{\theta}}_c$ is

$$E\big\{\hat{\boldsymbol{\theta}}_c\big\} = \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}\boldsymbol{\theta}. \tag{3.39}$$

If (3.37) holds, then (3.39) is the oblique projection of $\boldsymbol{\theta}$ along $\mathcal{N}(\boldsymbol{H})$ onto $\mathcal{R}(\boldsymbol{W})$ [71]. Thus, if $\boldsymbol{\theta} \in \mathcal{R}(\boldsymbol{W})$, then $\boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}\boldsymbol{\theta} = \boldsymbol{\theta}$, so that $\hat{\boldsymbol{\theta}}_c$ is indeed unbiased under this constraint. As a generalization, let us seek an estimator whose expectation is given by (3.39), while removing the constraint on $\boldsymbol{\theta}$ and the assumption (3.37). If such an estimator existed, then its bias would be given by

$$\boldsymbol{b}(\boldsymbol{\theta}) = \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}\boldsymbol{\theta} - \boldsymbol{\theta} \tag{3.40}$$

and therefore the matrix $\boldsymbol{A}$ of (3.6) would equal

$$\boldsymbol{A} = \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}. \tag{3.41}$$

Thus, we now seek an unconstrained but biased estimator. To find the minimum MSE estimator whose expectation is (3.39), we apply (3.9) of Theorem 3.1 with $\boldsymbol{U} = \boldsymbol{I}$ and $\boldsymbol{A}$ given by (3.41). This yields

$$\begin{aligned}
\hat{\boldsymbol{\theta}} &= \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}\boldsymbol{\theta} + \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}(\boldsymbol{H}^T\boldsymbol{H})^\dagger\boldsymbol{H}^T(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{\theta}) \\
&= \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}\boldsymbol{\theta} + \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}^{T\dagger}\boldsymbol{H}^T\boldsymbol{y} \\
&\qquad - \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}\boldsymbol{H}^\dagger\boldsymbol{H}\boldsymbol{\theta} \\
&= \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{H}^{T\dagger}\boldsymbol{H}^T\boldsymbol{y} \\
&= \boldsymbol{W}(\boldsymbol{H}\boldsymbol{W})^\dagger\boldsymbol{y} = \hat{\boldsymbol{\theta}}_c
\end{aligned} \tag{3.42}$$

where we used (3.17) and (3.18) in the second line, and (3.20) in the last line.

Thus, $\hat{\boldsymbol{\theta}}_c$ of (3.38) is the approach achieving minimum MSE among all estimators whose expectation is (3.39). This implies that $\hat{\boldsymbol{\theta}}_c$ is a useful estimator under a wider range of settings than suggested by the unbiased approach. Indeed, among estimators having the required expectation, $\hat{\boldsymbol{\theta}}_c$ is optimal even if $\boldsymbol{\theta}$ does not satisfy the constraint $\boldsymbol{\theta} \in \mathcal{R}(\boldsymbol{W})$, and, furthermore, its optimality is guaranteed even if the intersection between $\mathcal{R}(\boldsymbol{W})$ and $\mathcal{N}(\boldsymbol{H})$ is nontrivial.

Chapter 4

The CRB for Sparse Estimation

This chapter is an expanded version of the paper:

- Z. Ben-Haim and Y. C. Eldar, "The Cramér–Rao bound for estimating a sparse parameter vector," *IEEE Trans. Signal Processing*, vol. 58, no. 6, June 2010, pp. 3384–3389.

4.1 Introduction

The problem of estimating a sparse unknown parameter vector from noisy measurements has been analyzed intensively in the past few years [7, 8, 12, 22], and has already given rise to numerous successful signal processing algorithms [73–77]. In this paper, we consider the setting in which noisy measurements of a deterministic vector $\boldsymbol{x}_0$ are available. It is assumed that $\boldsymbol{x}_0$ has a sparse representation $\boldsymbol{x}_0 = \boldsymbol{D}\boldsymbol{\alpha}_0$, where $\boldsymbol{D}$ is a given dictionary and most of the entries of $\boldsymbol{\alpha}_0$ equal zero. Thus, only a small number of "atoms," or columns of $\boldsymbol{D}$, are required to represent $\boldsymbol{x}_0$. The challenges confronting an estimation technique are to recover either $\boldsymbol{x}_0$ itself or its sparse representation $\boldsymbol{\alpha}_0$. Several practical approaches turn out to be surprisingly successful in this task. Such approaches include the Dantzig selector (DS) [8] and basis pursuit denoising (BPDN), which is also referred to as the Lasso [7, 20, 22].

A standard measure of estimator performance is the mean-squared error (MSE). Several recent papers analyzed the MSE obtained by methods such as the DS and BPDN [8, 23, 78]. To determine the quality of estimation approaches, it is of interest to compare their achievements with theoretical performance limits: if existing methods approach the performance bound, then they are nearly optimal and further improvements in the current setting are impossible. This motivates the development of lower bounds on the MSE of estimators in the sparse setting.

Since the parameter to be estimated is deterministic, the MSE is in general a function of the parameter value. While there are lower bounds on the worst-case achievable MSE among all possible parameter values [24, §7.4], the actual performance for a specific value, or even for most values, might be substantially lower. Our goal is therefore to characterize the minimum MSE obtainable for each particular parameter vector. A standard method of achieving this objective is the Cramér–Rao bound (CRB) [16, 69].

The fact that $\boldsymbol{x}_0$ has a sparse representation is of central importance for estimator design. Indeed, many sparse estimation settings are underdetermined, meaning that without the assumption of sparsity, it is impossible to identify the correct parameter from its measurements, even without noise. In this paper, we treat the sparsity assumption as a deterministic prior constraint on the parameter. Specifically, we assume that $\boldsymbol{x}_0 \in \mathcal{S}$, where $\mathcal{S}$ is the set of all parameter vectors which can be represented by no more than s atoms, for a given integer s.

Our results are inspired by the well-studied theory of the constrained CRB [52, 55, 56, 58]. This theory is based on the assumption that the constraint set can be defined using the system of equations $\boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{0}$, $\boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}$, where $\boldsymbol{f}$ and $\boldsymbol{g}$ are continuously differentiable functions. The resulting bound depends on the derivatives of the function $\boldsymbol{f}$. However, sparsity constraints cannot be written in this form. This necessitates the development of a bound suitable for non-smooth constraint sets [79]. In obtaining this modified bound, we also provide new insight into the meaning of the general constrained CRB. In particular, we show that the fact that the constrained CRB is lower than the unconstrained bound results from an expansion of the class of estimators under consideration.

With the aforementioned theoretical tools at hand, we obtain lower bounds on the MSE in a variety of sparse estimation problems. Our bound limits the MSE achievable by any estimator having a pre-specified bias function, for each parameter value. Particular emphasis is given to the unbiased case; the reason for this preference is twofold: First, when the signal-to-noise ratio (SNR) is high, biased estimation is suboptimal. Second, for high SNR values, the unbiased CRB is achieved by the maximum likelihood (ML) estimator.

While the obtained bounds differ depending on the exact problem definition, in general terms and for unbiased estimation the bounds can be described as follows. For parameters having maximal support, i.e., parameters whose representation requires the maximum allowed number s of atoms, the lower bound equals the MSE of the "oracle estimator" which knows the locations (but not the values) of the nonzero representation elements. On the other hand, for parameters which do not have maximal support (a set which has Lebesgue measure zero in $\mathcal{S}$),

our lower bound is identical to the CRB for an unconstrained problem, which is substantially higher than the oracle MSE.

The correspondence between the CRB and the MSE of the oracle estimator (for all but a zero-measure subset of the feasible parameter set $\mathcal{S}$) is of practical interest since, unlike the oracle estimator, the CRB is achieved by the ML estimator at high SNR. Our bound can thus be viewed as an alternative justification for the common use of the oracle estimator as a baseline against which practical algorithms are compared. This gives further merit to recent results, which demonstrate that BPDN and the DS both achieve near-oracle performance [8, 23, 78]. However, the existence of parameters for which the bound is much higher indicates that oracular performance cannot be attained for *all* parameter values, at least using unbiased techniques. Indeed, as we will show, in many sparse estimation scenarios, one cannot construct *any* estimator which is unbiased for all sparsely representable parameters.

Our contribution is related to, but distinct from, the work of Babadi et al. [80], in which the CRB of the oracle estimator was derived (and shown to equal the aforementioned oracle MSE). Our goal in this work is to obtain a lower bound on the performance of estimators which are not endowed with oracular knowledge; consequently, as explained above, for some parameter values the obtained CRB will be higher than the oracle MSE. It was further shown in [80] that when the measurements consist of Gaussian random mixtures of the parameter vector, there exists an estimator which achieves the oracle CRB at high SNR; this is shown to hold on average over realizations of the measurement mixtures. The present contribution strengthens this result by showing that for any given (deterministic) well-behaved measurement setup, there exists a technique (namely, the ML estimator) achieving the CRB at high SNR. Thus, convergence to the CRB is guaranteed for all measurement settings, and not merely when averaging over an ensemble of such settings.

The rest of this paper is organized as follows. In Section 4.2, we review the sparse setting as a constrained estimation problem. Section 4.3 defines a generalization of sparsity constraints, which we refer to as locally balanced constraint sets; the CRB is then derived in this general setting. In Section 4.4, our general results are applied back to some specific sparse estimation problems. In Section 4.5, the CRB is compared to the empirical performance of estimators of sparse vectors. Our conclusions are summarized in Section 4.6.

Throughout the paper, boldface lowercase letters $\boldsymbol{v}$ denote vectors while boldface uppercase letters $\boldsymbol{M}$ denote matrices. Given a vector function $\boldsymbol{f} : \mathbb{R}^n \to \mathbb{R}^k$, we denote by $\partial \boldsymbol{f} / \partial \boldsymbol{x}$ the $k \times n$ matrix whose ijth element is $\partial f_i / \partial x_j$. The support of a vector, denoted $\mathrm{supp}(\boldsymbol{v})$, is the set of

indices of the nonzero entries in $\boldsymbol{v}$. The Euclidean norm of a vector $\boldsymbol{v}$ is denoted $\|\boldsymbol{v}\|_2$, and the number of nonzero entries in $\boldsymbol{v}$ is $\|\boldsymbol{v}\|_0$. Finally, the symbols $\mathcal{R}(\boldsymbol{M})$, $\mathcal{N}(\boldsymbol{M})$, and $\boldsymbol{M}^\dagger$ refer, respectively, to the column space, null space, and Moore–Penrose pseudoinverse of the matrix $\boldsymbol{M}$.

4.2 Sparse Estimation Problems

In this section, we describe several estimation problems whose common theme is that the unknown parameter has a sparse representation with respect to a known dictionary. We then review some standard techniques used to recover the unknown parameter in these problems. In Section 4.5 we will compare these methods with the performance bounds we develop.

4.2.1 The Sparse Setting

Suppose we observe a measurement vector $\boldsymbol{y} \in \mathbb{R}^m$, given by

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w} \tag{4.1}$$

where $\boldsymbol{x}_0 \in \mathbb{R}^n$ is an unknown deterministic signal, $\boldsymbol{w}$ is independent, identically distributed (IID) Gaussian noise with zero mean and variance σ^2, and $\boldsymbol{A}$ is a known $m \times n$ matrix. We assume the prior knowledge that there exists a sparse representation of $\boldsymbol{x}_0$, or, more precisely, that

$$\boldsymbol{x}_0 \in \mathcal{S} \triangleq \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{x} = \boldsymbol{D}\boldsymbol{\alpha}, \|\boldsymbol{\alpha}\|_0 \leq s\}\,. \tag{4.2}$$

In other words, the set $\mathcal{S}$ describes signals $\boldsymbol{x}$ which can be formed from a linear combination of no more than s columns, or atoms, from $\boldsymbol{D}$. The dictionary $\boldsymbol{D}$ is an $n \times p$ matrix with $n \leq p$, and we assume that $s < p$, so that only a subset of the atoms in $\boldsymbol{D}$ can be used to represent any signal in $\mathcal{S}$. We further assume that $\boldsymbol{D}$ and s are known.

Quite a few important signal recovery applications can be formulated using the setting described above. For example, if $\boldsymbol{A} = \boldsymbol{I}$, then $\boldsymbol{y}$ consists of noisy observations of $\boldsymbol{x}_0$, and recovering $\boldsymbol{x}_0$ is a denoising problem [73,74]. If $\boldsymbol{A}$ corresponds to a blurring kernel, we obtain a deblurring problem [75]. In both cases, the matrix $\boldsymbol{A}$ is square and invertible. Interpolation and inpainting can likewise be formulated as (4.1), but in those cases $\boldsymbol{A}$ is an underdetermined matrix, i.e., we have $m < n$ [77]. For all of these estimation scenarios, our goal is to obtain an estimate $\hat{\boldsymbol{x}}$ whose MSE is as low as possible, where the MSE is defined as

$$\text{MSE} \triangleq E\left\{\|\hat{\boldsymbol{x}} - \boldsymbol{x}_0\|_2^2\right\}. \tag{4.3}$$

Note that $\boldsymbol{x}_0$ is deterministic, so that the expectation in (4.3) (and throughout the paper) is taken over the noise $\boldsymbol{w}$ but not over $\boldsymbol{x}_0$. Thus, the MSE is in general a function of $\boldsymbol{x}_0$.

In the above settings, the goal is to estimate the unknown signal $\boldsymbol{x}_0$. However, it may also be of interest to recover the coefficient vector $\boldsymbol{\alpha}_0$ for which $\boldsymbol{x}_0 = \boldsymbol{D}\boldsymbol{\alpha}_0$, e.g., for the purpose of model selection [8, 22]. In this case, the goal is to construct an estimator $\widehat{\boldsymbol{\alpha}}$ whose MSE $E\{\|\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0\|_2^2\}$ is as low as possible. Unless $\boldsymbol{D}$ is unitary, estimating $\boldsymbol{\alpha}_0$ is not equivalent to estimating $\boldsymbol{x}_0$. Note, however, that when estimating $\boldsymbol{\alpha}_0$, the matrices $\boldsymbol{A}$ and $\boldsymbol{D}$ can be combined to obtain the equivalent problem

$$\boldsymbol{y} = \boldsymbol{H}\boldsymbol{\alpha}_0 + \boldsymbol{w} \tag{4.4}$$

where $\boldsymbol{H} \triangleq \boldsymbol{A}\boldsymbol{D}$ is an $m \times p$ matrix and

$$\boldsymbol{\alpha}_0 \in \mathcal{T} = \{\boldsymbol{\alpha} \in \mathbb{R}^p : \|\boldsymbol{\alpha}\|_0 \leq s\}. \tag{4.5}$$

Therefore, this problem can also be seen as a special case of (4.1) and (4.2). Nevertheless, it will occasionally be convenient to refer specifically to the problem of estimating $\boldsymbol{\alpha}_0$ from (4.4).

Signal estimation problems differ in the properties of the dictionary $\boldsymbol{D}$ and measurement matrix $\boldsymbol{A}$. In particular, problems of a very different nature arise depending on whether the dictionary is a basis or an overcomplete frame. For example, many approaches to denoising yield simple shrinkage techniques when $\boldsymbol{D}$ is a basis, but deteriorate to NP-hard optimization problems when $\boldsymbol{D}$ is overcomplete [81].

A final technical comment is in order. If the matrix $\boldsymbol{H}$ in (4.4) does not have full column rank, then there may exist different feasible parameters $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$ such that $\boldsymbol{H}\boldsymbol{\alpha}_1 = \boldsymbol{H}\boldsymbol{\alpha}_2$. In this case, the probability distribution of $\boldsymbol{y}$ will be identical for these two parameter vectors, and the estimation problem is said to be unidentifiable [6, §1.5.2]. A necessary and sufficient condition for identifiability is

$$\text{spark}(\boldsymbol{H}) > 2s \tag{4.6}$$

where $\text{spark}(\boldsymbol{H})$ is defined as the smallest integer k such that there exist k linearly dependent columns in $\boldsymbol{H}$ [82]. We will adopt the assumption (4.6) throughout the paper. Similarly, in the problem (4.1) we will assume that

$$\text{spark}(\boldsymbol{D}) > 2s. \tag{4.7}$$

4.2.2 Estimation Techniques

We now review some standard estimators for the sparse problems described above. These techniques are usually viewed as methods for obtaining an estimate $\widehat{\boldsymbol{\alpha}}$ of the vector $\boldsymbol{\alpha}_0$ in (4.4),

and we will adopt this perspective in the current section. One way to estimate $\boldsymbol{x}_0$ in the more general problem (4.1) is to first estimate $\boldsymbol{\alpha}_0$ with the methods described below and then use the formula $\hat{\boldsymbol{x}} = \boldsymbol{D}\widehat{\boldsymbol{\alpha}}$.

A widely-used estimation technique is the ML approach, which provides an estimate of $\boldsymbol{\alpha}_0$ by solving

$$\min_{\boldsymbol{\alpha}} \|\boldsymbol{y} - \boldsymbol{H}\boldsymbol{\alpha}\|_2^2 \quad \text{s.t. } \|\boldsymbol{\alpha}\|_0 \leq s. \tag{4.8}$$

Unfortunately, (4.8) is a nonconvex optimization problem and solving it is NP-hard [81], meaning that an efficient algorithm providing the ML estimator is unlikely to exist. In fact, to the best of our knowledge, the most efficient method for solving (4.8) for general $\boldsymbol{H}$ is to enumerate the $\binom{p}{s}$ possible s-element support sets of $\boldsymbol{\alpha}$ and choose the one for which $\|\boldsymbol{y} - \boldsymbol{H}\boldsymbol{\alpha}\|_2^2$ is minimal. This is clearly an impractical strategy for reasonable values of p and s. Consequently, several efficient alternatives have been proposed for estimating $\boldsymbol{\alpha}_0$. One of these is the ℓ_1-penalty version of BPDN [22], which is defined as a solution $\widehat{\boldsymbol{\alpha}}_{\mathrm{BP}}$ to the quadratic program

$$\min_{\boldsymbol{\alpha}} \tfrac{1}{2}\|\boldsymbol{y} - \boldsymbol{H}\boldsymbol{\alpha}\|_2^2 + \gamma\|\boldsymbol{\alpha}\|_1 \tag{4.9}$$

with some regularization parameter γ. More recently, the DS was proposed [8]; this approach estimates $\boldsymbol{\alpha}_0$ as a solution $\widehat{\boldsymbol{\alpha}}_{\mathrm{DS}}$ to

$$\min_{\boldsymbol{\alpha}} \|\boldsymbol{\alpha}\|_1 \quad \text{s.t. } \|\boldsymbol{H}^T(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{\alpha})\|_\infty \leq \tau \tag{4.10}$$

where τ is again a user-selected parameter. A modification of the DS, known as the Gauss–Dantzig selector (GDS) [8], is to use $\widehat{\boldsymbol{\alpha}}_{\mathrm{DS}}$ only to estimate the support of $\boldsymbol{\alpha}_0$. In this approach, one solves (4.10) and determines the support set of $\widehat{\boldsymbol{\alpha}}_{\mathrm{DS}}$. The GDS estimate is then obtained as

$$\widehat{\boldsymbol{\alpha}}_{\mathrm{GDS}} = \begin{cases} \boldsymbol{H}_{\widehat{\boldsymbol{\alpha}}_{\mathrm{DS}}}^{\dagger}\boldsymbol{y} & \text{on the support set of } \widehat{\boldsymbol{\alpha}}_{\mathrm{DS}} \\ \boldsymbol{0} & \text{elsewhere} \end{cases} \tag{4.11}$$

where $\boldsymbol{H}_{\widehat{\boldsymbol{\alpha}}_{\mathrm{DS}}}$ consists of the columns of $\boldsymbol{H}$ corresponding to the support of $\widehat{\boldsymbol{\alpha}}_{\mathrm{DS}}$.

Previous research on the performance of these estimators has primarily examined their worst-case MSE among all possible values of $\boldsymbol{\alpha}_0 \in \mathcal{T}$. Specifically, it has been shown [8] that, under suitable conditions on $\boldsymbol{H}$, s, and τ, the DS of (4.10) satisfies

$$\|\boldsymbol{\alpha}_0 - \widehat{\boldsymbol{\alpha}}_{\mathrm{DS}}\|_2^2 \leq Cs\sigma^2 \log p \quad \text{with high probability} \tag{4.12}$$

for some constant C. It follows that the MSE of the DS is also no greater than a constant times $s\sigma^2 \log p$ for all $\boldsymbol{\alpha}_0 \in \mathcal{T}$ [24]. An identical property was also demonstrated for BPDN (4.9) with

an appropriate choice of γ [78]. Conversely, it is known that the worst-case error of *any* estimator is at least a constant times $s\sigma^2 \log p$ [24, §7.4]. Thus, both BPDN and the DS are optimal, up to a constant, in terms of worst-case error. Nevertheless, the MSE of these approaches for specific values of $\boldsymbol{\alpha}_0$, even for a vast majority of such values, might be much lower. Our goal differs from this line of work in that we characterize the *pointwise* performance of an estimator, i.e., the MSE for specific values of $\boldsymbol{\alpha}_0$.

Another baseline with which practical techniques are often compared is the oracle estimator, given by

$$\widehat{\boldsymbol{\alpha}}_{\text{oracle}} = \begin{cases} \boldsymbol{H}_{\boldsymbol{\alpha}_0}^{\dagger} \boldsymbol{b} & \text{on the set } \operatorname{supp}(\boldsymbol{\alpha}_0) \\ \boldsymbol{0} & \text{elsewhere} \end{cases} \tag{4.13}$$

where $\boldsymbol{H}_{\boldsymbol{\alpha}_0}$ is the submatrix constructed from the columns of $\boldsymbol{H}$ corresponding to the nonzero entries of $\boldsymbol{\alpha}_0$. In other words, $\widehat{\boldsymbol{\alpha}}_{\text{oracle}}$ is the least-squares (LS) solution among vectors whose support coincides with $\operatorname{supp}(\boldsymbol{\alpha}_0)$, which is assumed to have been provided by an "oracle." Of course, in practice the support of $\boldsymbol{\alpha}_0$ is unknown, so that $\widehat{\boldsymbol{\alpha}}_{\text{oracle}}$ cannot actually be implemented. Nevertheless, one often compares the performance of true estimators with $\widehat{\boldsymbol{\alpha}}_{\text{oracle}}$, whose MSE is given by [8]

$$\sigma^2 \operatorname{Tr}((\boldsymbol{H}_{\boldsymbol{\alpha}_0}^T \boldsymbol{H}_{\boldsymbol{\alpha}_0})^{-1}). \tag{4.14}$$

Is (4.14) a bound on estimation MSE? While $\widehat{\boldsymbol{\alpha}}_{\text{oracle}}$ is a reasonable technique to adopt if $\operatorname{supp}(\boldsymbol{\alpha}_0)$ is known, this does not imply that (4.14) is a lower bound on the performance of practical estimators. Indeed, as will be demonstrated in Section 4.5, when the SNR is low, both BPDN and the DS outperform $\widehat{\boldsymbol{\alpha}}_{\text{oracle}}$, thanks to the use of shrinkage in these estimators. Furthermore, if $\operatorname{supp}(\boldsymbol{\alpha}_0)$ is known, then there exist biased techniques which are better than $\widehat{\boldsymbol{\alpha}}_{\text{oracle}}$ for *all* values of $\boldsymbol{\alpha}_0$ [42]. Thus, $\widehat{\boldsymbol{\alpha}}_{\text{oracle}}$ is neither achievable in practice, nor optimal in terms of MSE. As we will see, one can indeed interpret (4.14) as a lower bound on the achievable MSE, but such a result requires a certain restriction of the class of estimators under consideration.

4.3 The Constrained Cramér–Rao Bound

A common technique for determining the achievable performance in a given estimation problem is to calculate the CRB, which is a lower bound on the MSE of estimators having a given bias [16]. In this paper, we are interested in calculating the CRB when it is known that the parameter $\boldsymbol{x}$ satisfies sparsity constraints such as those of the sets $\mathcal{S}$ of (4.2) and $\mathcal{T}$ of (4.5).

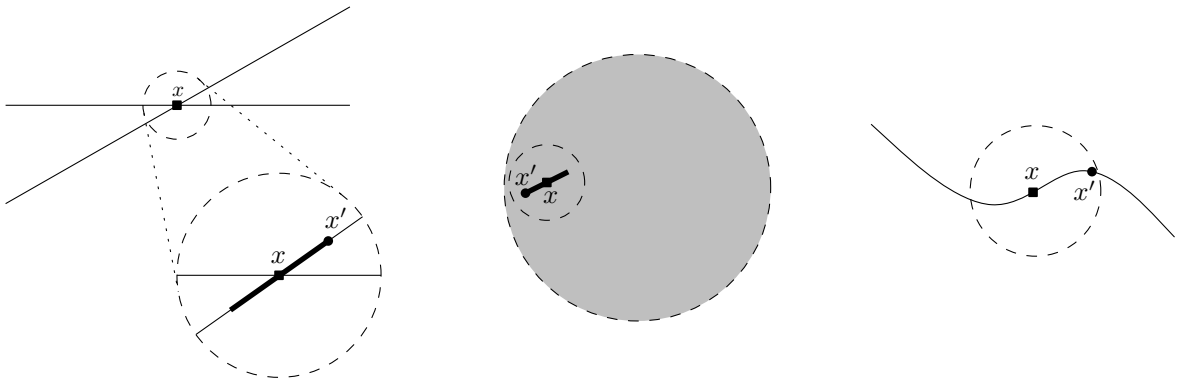

(a) (b) (c)

Figure 4.1: In a locally balanced set such as a union of subspaces (a) and an open ball (b), each point is locally defined by a set of feasible directions along which an infinitesimal movement does not violate the constraints. The curve (c) is not characterized in this way and thus is not locally balanced.

The CRB for constrained parameter sets has been studied extensively in the past [52, 55, 56, 58]. However, in prior work derivation of the CRB assumed that the constraint set is given by

$$\mathcal{X} = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{0},\ \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}\} \tag{4.15}$$

where $\boldsymbol{f}(\boldsymbol{x})$ and $\boldsymbol{g}(\boldsymbol{x})$ are continuously differentiable functions. We will refer to such $\mathcal{X}$ as continuously differentiable sets. As shown in prior work [52], the resulting bound depends on the derivatives of the function $\boldsymbol{f}$. Yet in some cases, including the sparse estimation scenarios discussed in Section 4.2, the constraint set cannot be written in the form (4.15), and the aforementioned results are therefore inapplicable. Our goal in the current section is to close this gap by extending the constrained CRB to constraint sets $\mathcal{X}$ encompassing the sparse estimation scenario.

We begin this section with a general discussion of the CRB and the class of estimators to which it applies. This will lead us to interpret the constrained CRB as a bound on estimators having an incompletely specified bias gradient. This interpretation will facilitate the application of the existing constrained CRB to the present context.

4.3.1 Bias Requirements in the Constrained CRB

In previous settings for which the constrained CRB was derived, it was noted that the resulting bound is typically lower than the unconstrained version [52, Remark 4]. At first glance, one

would attribute the reduction in the value of the CRB to the fact that the constraints add information about the unknown parameter, which can then improve estimation performance. On the other hand, the CRB separately characterizes the achievable performance for each value of the unknown parameter $\boldsymbol{x}_0$. Thus, the CRB at $\boldsymbol{x}_0$ applies even to estimators designed specifically to perform well at $\boldsymbol{x}_0$. Such estimators surely cannot achieve further gain in performance if it is known that $\boldsymbol{x}_0 \in \mathcal{X}$. Why, then, is the constrained CRB lower than the unconstrained bound? The answer to this apparent paradox involves a careful definition of the class of estimators to which the bound applies.

To obtain a meaningful bound, one must exclude some estimators from consideration. Unless this is done, the bound will be tarnished by estimators of the type $\hat{\boldsymbol{x}} = \boldsymbol{x}_{\mathrm{u}}$, for some constant $\boldsymbol{x}_{\mathrm{u}}$, which achieve an MSE of 0 at the specific point $\boldsymbol{x} = \boldsymbol{x}_{\mathrm{u}}$. It is standard practice to circumvent this difficulty by restricting attention to estimators having a particular bias $\boldsymbol{b}(\boldsymbol{x}) \triangleq E\{\hat{\boldsymbol{x}}\} - \boldsymbol{x}$. In particular, it is common to examine unbiased estimators, for which $\boldsymbol{b}(\boldsymbol{x}) = \boldsymbol{0}$.

However, in some settings, it is impossible to construct estimators which are unbiased for all $\boldsymbol{x} \in \mathbb{R}^n$. For example, suppose we are to estimate the coefficients $\boldsymbol{\alpha}_0$ of an overcomplete dictionary based on the measurements given by (4.4). Since the dictionary is overcomplete, its nullspace is nontrivial; furthermore, each coefficient vector in the nullspace yields an identical distribution of the measurements, so that an estimator can be unbiased for one of these vectors at most.

The question is whether it is possible to construct estimators which are unbiased for some, but not all, values of $\boldsymbol{x}$. One possible approach is to seek estimators which are unbiased for all $\boldsymbol{x} \in \mathcal{X}$. However, as we will see later in this section, even this requirement can be too strict: in some cases it is impossible to construct estimators which are unbiased for all $\boldsymbol{x} \in \mathcal{X}$. More generally, the CRB is a *local* bound, meaning that it determines the achievable performance at a particular value of $\boldsymbol{x}$ based on the statistics at $\boldsymbol{x}$ and at nearby values. Thus, it is irrelevant to introduce requirements on estimation performance for parameters which are distant from the value $\boldsymbol{x}$ of interest.

Since we seek a locally unbiased estimator, one possibility is to require unbiasedness at a single point, say $\boldsymbol{x}_{\mathrm{u}}$. As it turns out, it is always possible to construct such a technique: this is again $\hat{\boldsymbol{x}} = \boldsymbol{x}_{\mathrm{u}}$, which is unbiased at $\boldsymbol{x}_{\mathrm{u}}$ but nowhere else. To avoid this loophole, one can require an estimator to be unbiased in the neighborhood

$$\mathcal{B}_\varepsilon(\boldsymbol{x}_0) = \{\boldsymbol{x} \in \mathbb{R}^m : \|\boldsymbol{x} - \boldsymbol{x}_0\|_2 < \varepsilon\} \tag{4.16}$$

of $\boldsymbol{x}_0$, for some small ε. It follows that both the bias $\boldsymbol{b}(\boldsymbol{x})$ and the bias gradient

$$\mathcal{B}(\boldsymbol{x}) \triangleq \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{x}} \tag{4.17}$$

vanish at $\boldsymbol{x} = \boldsymbol{x}_0$. This formulation is the basis of the unconstrained unbiased CRB, a lower bound on the covariance at $\boldsymbol{x}_0$ which applies to all estimators whose bias gradient is zero at $\boldsymbol{x}_0$.

It turns out that even this requirement is too stringent in constrained settings. As we will see in Section 4.4.1, estimators of the coefficients of an overcomplete dictionary must have a nonzero bias gradient matrix. The reason is related to the fact that unbiasedness is required over the set $\mathcal{B}_\varepsilon(\boldsymbol{x}_0)$, which, in the overcomplete setting, has a higher dimension than the number of measurements.

However, it can be argued that one is not truly interested in the bias at all points in $\mathcal{B}_\varepsilon(\boldsymbol{x}_0)$, since many of these points violate the constraint set $\mathcal{X}$. A reasonable compromise is to require unbiasedness over $\mathcal{B}_\varepsilon(\boldsymbol{x}_0) \cap \mathcal{X}$, i.e., over the neighborhood of $\boldsymbol{x}_0$ restricted to the constraint set $\mathcal{X}$. This leads to a weaker requirement on the bias gradient $\mathcal{B}$ at $\boldsymbol{x}_0$. Specifically, the derivatives of the bias need only be specified in directions which do not violate the constraints. The exact formulation of this requirement depends on the nature of the set $\mathcal{X}$. In the following subsections, we will investigate various constraint sets and derive the corresponding requirements on the bias function.

It is worth emphasizing that the dependence of the CRB on the constraints is manifested through the class of estimators being considered, or more specifically, through the allowed estimators' bias gradient matrices. By contrast, the unconstrained CRB applies to estimators having a fully specified bias gradient matrix. Consequently, the constrained bound applies to a wider class of estimators, and is thus usually lower than the unconstrained version of the CRB. In other words, estimators which are unbiased in the constrained setting, and thus applicable to the unbiased constrained CRB, are likely to be biased in the unconstrained context. Since a wider class of estimators is considered by the constrained CRB, the resulting bound is lower, thus explaining the puzzling phenomenon described in the beginning of this subsection.

4.3.2 Locally Balanced Constraints

We now consider a class of constraint sets, called locally balanced sets, which encompass the sparsity constraints of Section 4.2. Roughly speaking, a locally balanced set is one which is locally defined at each point by the directions along which one can move without leaving the set. Formally, a metric space $\mathcal{X}$ is said to be locally balanced if, for all $\boldsymbol{x} \in \mathcal{X}$, there exists an

open set $\mathcal{C} \subset \mathcal{X}$ such that $\boldsymbol{x} \in \mathcal{C}$ and such that, for all $\boldsymbol{x}' \in \mathcal{C}$ and for all $|\lambda| \leq 1$, we have

$$\boldsymbol{x} + \lambda(\boldsymbol{x}' - \boldsymbol{x}) \in \mathcal{C}. \tag{4.18}$$

As we will see, locally balanced sets are useful in the context of the constrained CRB, as they allow us to identify the feasible directions along which the bias gradient must be specified.

An example of a locally balanced set is given in Fig. 4.1(a), which represents a union of two subspaces. In Fig. 4.1(a), for any point $\boldsymbol{x} \in \mathcal{X}$, and for any point $\boldsymbol{x}' \in \mathcal{X}$ sufficiently close to $\boldsymbol{x}$, the entire line segment between $\boldsymbol{x}$ and $\boldsymbol{x}'$, as well as the line segment in the opposite direction, are also in $\mathcal{X}$. This illustrates the fact that any union of subspaces is locally balanced, and, in particular, so are the sparse estimation settings of Section 4.2 [9, 36, 83]. As another example, consider any open set, such as the open ball in Fig. 4.1(b). For such a set, any point $\boldsymbol{x}$ has a sufficiently small neighborhood $\mathcal{C}$ such that, for any $\boldsymbol{x}' \in \mathcal{C}$, the line segment connecting $\boldsymbol{x}$ to $\boldsymbol{x}'$ is contained in $\mathcal{X}$. On the other hand, the curve in Fig. 4.1(c) is not locally balanced, since the line connecting $\boldsymbol{x}$ to any other point on the set does not lie within the set.[1]

Observe that the neighborhood of a point $\boldsymbol{x}$ in a locally balanced set $\mathcal{X}$ is entirely determined by the set of feasible directions $\boldsymbol{v}$ along which infinitesimal changes of $\boldsymbol{x}$ do not violate the constraints. These are the directions $\boldsymbol{v} = \boldsymbol{x}' - \boldsymbol{x}$ for all points $\boldsymbol{x}' \neq \boldsymbol{x}$ in the set $\mathcal{C}$ of (4.18). Recall that we seek a lower bound on the performance of estimators whose bias gradient is defined over the neighborhood of $\boldsymbol{x}_0$ restricted to the constraint set $\mathcal{X}$. Suppose for concreteness that we are interested in unbiased estimators. For a locally balanced constraint set $\mathcal{X}$, this implies that

$$\mathcal{B}\boldsymbol{v} = \boldsymbol{0} \tag{4.19}$$

for any feasible direction $\boldsymbol{v}$. In other words, all feasible directions must be in the nullspace of $\mathcal{B}$. This is a weaker condition than requiring the bias gradient to equal zero, and is thus more useful for constrained estimation problems. If an estimator $\hat{\boldsymbol{x}}$ satisfies (4.19) for all feasible directions $\boldsymbol{v}$ at a certain point $\boldsymbol{x}_0$, we say that $\hat{\boldsymbol{x}}$ is $\mathcal{X}$-unbiased at $\boldsymbol{x}_0$. This terminology emphasizes the fact that $\mathcal{X}$-unbiasedness depends both on the point $\boldsymbol{x}_0$ and on the constraint set $\mathcal{X}$.

Consider the subspace $\mathcal{F}$ spanned by the feasible directions at a certain point $\boldsymbol{x} \in \mathcal{X}$. We refer to $\mathcal{F}$ as the feasible subspace at $\boldsymbol{x}$. Note that $\mathcal{F}$ may include infeasible directions, if these are linear combinations of feasible directions. Nevertheless, because of the linearity of (4.19),

[1] We note in passing that since the curve in Fig. 4.1(c) is continuously differentiable, it can be locally approximated by a locally balanced set. Our derivation of the CRB can be extended to such approximately locally balanced sets in a manner similar to that of [52], but such an extension is not necessary for the purposes of this paper.

any vector $\boldsymbol{u} \in \mathcal{F}$ satisfies $\mathcal{B}\boldsymbol{u} = \mathbf{0}$, even if $\boldsymbol{u}$ is infeasible. Thus, $\mathcal{X}$-unbiasedness is actually a property of the feasible subspace $\mathcal{F}$, rather than the set of feasible directions.

Since $\mathcal{X}$ is a subset of a finite-dimensional Euclidean space, $\mathcal{F}$ is also finite-dimensional, although different points in $\mathcal{X}$ may yield subspaces having differing dimensions. Let $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_l$ denote an orthonormal basis for $\mathcal{F}$, and define the matrix

$$\boldsymbol{U} = [\boldsymbol{u}_1, \ldots, \boldsymbol{u}_l]. \tag{4.20}$$

Note that $\boldsymbol{u}_i$ and $\boldsymbol{U}$ are functions of $\boldsymbol{x}$. For a given function $\boldsymbol{x}$, different orthonormal bases can be chosen, but the choice of a basis is arbitrary and will not affect our results.

As we have seen, $\mathcal{X}$-unbiasedness at $\boldsymbol{x}_0$ can alternatively be written as $\mathcal{B}\boldsymbol{u} = \mathbf{0}$ for all $\boldsymbol{u} \in \mathcal{F}$, or, equivalently

$$\mathcal{B}\boldsymbol{U} = \mathbf{0}. \tag{4.21}$$

The constrained CRB can now be derived as a lower bound on all $\mathcal{X}$-unbiased estimators, which is a weaker requirement than "ordinary" unbiasedness.

Just as $\mathcal{X}$-unbiasedness was defined by requiring the bias gradient matrix to vanish when multiplied by any feasible direction vector, we can define $\mathcal{X}$-biased estimators by requiring a specific value (not necessarily zero) for the bias gradient matrix when multiplied by a feasible direction vector. In an analogy to (4.21), this implies that one must define a value for the matrix $\mathcal{B}\boldsymbol{U}$. Our goal is thus to construct a lower bound on the covariance at a given $\boldsymbol{x}$ achievable by any estimator whose bias gradient $\mathcal{B}$ at $\boldsymbol{x}$ satisfies $\mathcal{B}\boldsymbol{U} = \boldsymbol{P}$, for a given matrix $\boldsymbol{P}$. This is referred to as specifying the $\mathcal{X}$-bias of the estimator at $\boldsymbol{x}$.

4.3.3 The CRB for Locally Balanced Constraints

It is helpful at this point to compare our derivation with prior work on the constrained CRB, which considered continuously differentiable constraint sets of the form (4.15). It has been previously shown [52] that inequality constraints of the type $\boldsymbol{g}(\boldsymbol{x}) \leq \mathbf{0}$ have no effect on the CRB. Consequently, we will consider constraints of the form

$$\mathcal{X} = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{f}(\boldsymbol{x}) = \mathbf{0}\}. \tag{4.22}$$

Define the $k \times n$ matrix $\boldsymbol{F}(\boldsymbol{x}) = \partial \boldsymbol{f} / \partial \boldsymbol{x}$. For simplicity of notation, we will omit the dependence of $\boldsymbol{F}$ on $\boldsymbol{x}$. Assuming that the constraints are non-redundant, $\boldsymbol{F}$ is a full-rank matrix, and thus one can define an $n \times (n-k)$ matrix $\boldsymbol{W}$ (also dependent on $\boldsymbol{x}$) such that

$$\boldsymbol{F}\boldsymbol{W} = \mathbf{0}, \quad \boldsymbol{W}^T\boldsymbol{W} = \boldsymbol{I}. \tag{4.23}$$

The matrix $\boldsymbol{W}$ is closely related to the matrix $\boldsymbol{U}$ spanning the feasible direction subspace of locally balanced sets. Indeed, the column space $\mathcal{R}(\boldsymbol{W})$ of $\boldsymbol{W}$ is the tangent space of $\mathcal{X}$, i.e., the subspace of $\mathbb{R}^n$ containing all vectors which are tangent to $\mathcal{X}$ at the point $\boldsymbol{x}$. Thus, the vectors in $\mathcal{R}(\boldsymbol{W})$ are precisely those directions along which infinitesimal motion from $\boldsymbol{x}$ does not violate the constraints, up to a first-order approximation. It follows that if a particular set $\mathcal{X}$ is both locally balanced and continuously differentiable, its matrices $\boldsymbol{U}$ and $\boldsymbol{W}$ coincide. Note, however, that there exist sets which are locally balanced but not continuously differentiable (and vice versa).

With the above formulation, the CRB for continuously differentiable constraints can be stated as a function of the the matrix $\boldsymbol{W}$ and the bias gradient $\mathcal{B}$ [58]. In fact, the resulting bound depends on $\mathcal{B}$ only through $\mathcal{B}\boldsymbol{W}$. This is to be expected in light of the discussion of Section 4.3.1: The bias should be specified only for those directions which do not violate the constraint set. Furthermore, the proof of the CRB in [58, Theorem 1] depends not on the formulation (4.22) of the constraint set, but merely on the class of bias functions under consideration. Consequently, one can state the bound without any reference to the underlying constraint set. To do so, let $\boldsymbol{y}$ be a measurement vector with pdf $p(\boldsymbol{y};\boldsymbol{x})$, which is assumed to be differentiable with respect to $\boldsymbol{x}$. The Fisher information matrix (FIM) $\boldsymbol{J}(\boldsymbol{x})$ is defined as

$$\boldsymbol{J}(\boldsymbol{x}) = E\Big\{\boldsymbol{\Delta}\boldsymbol{\Delta}^T\Big\} \tag{4.24}$$

where

$$\boldsymbol{\Delta} = \frac{\partial \log p(\boldsymbol{y};\boldsymbol{x})}{\partial \boldsymbol{x}}. \tag{4.25}$$

We assume that the FIM is well-defined and finite. We further assume that integration with respect to $\boldsymbol{y}$ and differentiation with respect to $\boldsymbol{x}$ can be interchanged, a standard requirement for the CRB. We then have the following result.

Theorem 4.1. *Let $\hat{\boldsymbol{x}}$ be an estimator and let $\mathcal{B} = \partial \boldsymbol{b}/\partial \boldsymbol{x}$ denote the bias gradient matrix of $\hat{\boldsymbol{x}}$ at a given point $\boldsymbol{x}_0$. Let $\boldsymbol{U}$ be an orthonormal matrix, and suppose that $\mathcal{B}\boldsymbol{U}$ is known, but that $\mathcal{B}$ is otherwise arbitrary. If*

$$\mathcal{R}\Big(\boldsymbol{U}(\boldsymbol{U}+\mathcal{B}\boldsymbol{U})^T\Big) \subseteq \mathcal{R}\Big(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\Big) \tag{4.26}$$

then the covariance of $\hat{\boldsymbol{x}}$ at $\boldsymbol{x}_0$ satisfies

$$\mathrm{Cov}(\hat{\boldsymbol{x}}) \succeq (\boldsymbol{U}+\mathcal{B}\boldsymbol{U})\left(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\right)^{\dagger}(\boldsymbol{U}+\mathcal{B}\boldsymbol{U})^T. \tag{4.27}$$

Equality is achieved in (4.27) *if and only if*

$$\hat{\boldsymbol{x}} = \boldsymbol{x}_0 + \boldsymbol{b}(\boldsymbol{x}_0) + (\boldsymbol{U}+\mathcal{B}\boldsymbol{U})\left(\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\right)^{\dagger}\boldsymbol{U}^T\boldsymbol{\Delta} \tag{4.28}$$

in the mean square sense, where $\mathbf{\Delta}$ *is defined by* (4.25)*. Conversely, if* (4.26) *does not hold, then there exists no finite-variance estimator with the required bias gradient.*

As required, no mention of constrained estimation is made in Theorem 4.1; instead, partial information about the bias gradient is assumed. Apart from this restatement, the theorem is identical to [58, Theorem 1], and its proof is unchanged. However, the above formulation is more general in that it can be applied to any constrained setting, once the constraints have been translated to bias gradient requirements. In particular, Theorem 4.1 provides a CRB for locally balanced sets if the matrix $\boldsymbol{U}$ is chosen as a basis for the feasible direction subspace of Section 4.3.2.

4.4 Bounds on Sparse Estimation

In this section, we apply the CRB of Theorem 4.1 to several sparse estimation scenarios. We begin with an analysis of the problem of estimating a sparse parameter vector.

4.4.1 Estimating a Sparse Vector

Suppose we would like to estimate a parameter vector $\boldsymbol{\alpha}_0$, known to belong to the set $\mathcal{T}$ of (4.5), from measurements $\boldsymbol{y}$ given by (4.4). To determine the CRB in this setting, we begin by identifying the feasible subspaces $\mathcal{F}$ corresponding to each of the elements in $\mathcal{T}$. To this end, consider first vectors $\boldsymbol{\alpha} \in \mathcal{T}$ for which $\|\boldsymbol{\alpha}\|_0 = s$, i.e., vectors having maximal support. Denote by $\{i_1, \ldots, i_s\}$ the support set of $\boldsymbol{\alpha}$. Then, for all δ, we have

$$\|\boldsymbol{\alpha} + \delta \boldsymbol{e}_{i_k}\|_0 = \|\boldsymbol{\alpha}\|_0 = s, \quad k = 1, \ldots, s \tag{4.29}$$

where $\boldsymbol{e}_j$ is the jth column of the identity matrix. Thus $\boldsymbol{\alpha} + \delta \boldsymbol{e}_{i_k} \in \mathcal{T}$, and consequently, the vectors $\{\boldsymbol{e}_{i_1}, \ldots, \boldsymbol{e}_{i_s}\}$ are all feasible directions, as is any linear combination of these vectors. On the other hand, for any $j \notin \operatorname{supp}(\boldsymbol{\alpha})$ and for any nonzero δ, we have $\|\boldsymbol{\alpha} + \delta \boldsymbol{e}_j\|_0 = s + 1$, and thus $\boldsymbol{e}_j$ is not a feasible direction; neither is any other vector which is not in $\operatorname{span}\{\boldsymbol{e}_{i_1}, \ldots, \boldsymbol{e}_{i_s}\}$. It follows that the feasible subspace $\mathcal{F}$ for points having maximal support is given by $\operatorname{span}\{\boldsymbol{e}_{i_1}, \ldots, \boldsymbol{e}_{i_s}\}$, and a possible choice for the matrix $\boldsymbol{U}$ of (4.20) is

$$\boldsymbol{U} = [\boldsymbol{e}_{i_1}, \ldots, \boldsymbol{e}_{i_s}] \quad \text{for } \|\boldsymbol{\alpha}\|_0 = s. \tag{4.30}$$

The situation is different for points $\boldsymbol{\alpha}$ having $\|\boldsymbol{\alpha}\|_0 < s$. In this case, vectors $\boldsymbol{e}_i$ corresponding

to *any* direction i are feasible directions, since

$$\|\boldsymbol{\alpha} + \delta \boldsymbol{e}_i\|_0 \leq \|\boldsymbol{\alpha}\|_0 + 1 \leq s. \tag{4.31}$$

Because the feasible subspace is defined as the span of all feasible directions, we have

$$\mathcal{F} \supseteq \text{span}\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_p\} = \mathbb{R}^p. \tag{4.32}$$

It follows that $\mathcal{F} = \mathbb{R}^p$ and thus a convenient choice for the matrix $\boldsymbol{U}$ is

$$\boldsymbol{U} = \boldsymbol{I} \quad \text{for } \|\boldsymbol{\alpha}\|_0 < s. \tag{4.33}$$

Consequently, whenever $\|\boldsymbol{\alpha}\|_0 < s$, a specification of the $\mathcal{T}$-bias amounts to completely specifying the usual estimation bias $\boldsymbol{b}(\boldsymbol{x})$.

To invoke Theorem 4.1, we must also determine the FIM $\boldsymbol{J}(\boldsymbol{\alpha})$. Under our assumption of white Gaussian noise, $\boldsymbol{J}(\boldsymbol{\alpha})$ is given by [16, p. 85]

$$\boldsymbol{J}(\boldsymbol{\alpha}) = \frac{1}{\sigma^2} \boldsymbol{H}^T \boldsymbol{H}. \tag{4.34}$$

Using (4.30), (4.33), and (4.34), it is readily shown that

$$\boldsymbol{U}^T \boldsymbol{J} \boldsymbol{U} = \begin{cases} \frac{1}{\sigma^2} \boldsymbol{H}_{\boldsymbol{\alpha}}^T \boldsymbol{H}_{\boldsymbol{\alpha}} & \text{when } \|\boldsymbol{\alpha}\|_0 = s \\ \frac{1}{\sigma^2} \boldsymbol{H}^T \boldsymbol{H} & \text{when } \|\boldsymbol{\alpha}\|_0 < s \end{cases} \tag{4.35}$$

where $\boldsymbol{H}_{\boldsymbol{\alpha}}$ is the $p \times s$ matrix consisting of the columns of $\boldsymbol{H}$ indexed by $\text{supp}(\boldsymbol{\alpha})$.

We now wish to determine under what conditions (4.26) holds. Consider first points $\boldsymbol{\alpha}_0$ for which $\|\boldsymbol{\alpha}_0\|_0 = s$. Since, by (4.6), we have $\text{spark}(\boldsymbol{H}) > s$, it follows that in this case $\boldsymbol{U}^T \boldsymbol{J} \boldsymbol{U}$ is invertible. Therefore

$$\mathcal{R}\Big(\boldsymbol{U}\boldsymbol{U}^T \boldsymbol{J} \boldsymbol{U}\boldsymbol{U}^T\Big) = \mathcal{R}\Big(\boldsymbol{U}\boldsymbol{U}^T\Big). \tag{4.36}$$

Since

$$\mathcal{R}\Big(\boldsymbol{U}\boldsymbol{U}^T(\boldsymbol{I} + \mathcal{B}^T)\Big) \subseteq \mathcal{R}\Big(\boldsymbol{U}\boldsymbol{U}^T\Big) \tag{4.37}$$

we have that condition (4.26) holds when $\|\boldsymbol{\alpha}_0\|_0 = s$.

The condition (4.26) is no longer guaranteed when $\|\boldsymbol{\alpha}_0\|_0 < s$. In this case, $\boldsymbol{U} = \boldsymbol{I}$, so that (4.26) is equivalent to

$$\mathcal{R}\Big(\boldsymbol{I} + \mathcal{B}^T\Big) \subseteq \mathcal{R}\Big(\boldsymbol{H}^T \boldsymbol{H}\Big). \tag{4.38}$$

Using the fact that $\mathcal{R}(\boldsymbol{H}^T\boldsymbol{H}) = \mathcal{R}(\boldsymbol{H}^T)$ and that, for any matrix $\boldsymbol{Q}$, $\mathcal{R}(\boldsymbol{Q}^T) = \mathcal{N}(\boldsymbol{Q})^\perp$, we find that (4.38) is equivalent to

$$\mathcal{N}(\boldsymbol{H}) \subseteq \mathcal{N}(\boldsymbol{I} + \mathcal{B}). \tag{4.39}$$

Combining these conclusions with Theorem 4.1 yields the following CRB for the problem of estimating a sparse vector.

Theorem 4.2. *Consider the estimation problem* (4.4) *with* $\boldsymbol{\alpha}_0$ *given by* (4.5)*, and assume that* (4.6) *holds. For a finite-variance estimator* $\widehat{\boldsymbol{\alpha}}$ *of* $\boldsymbol{\alpha}_0$ *to exist, its bias gradient matrix* $\mathcal{B}$ *must satisfy* (4.39) *whenever* $\|\boldsymbol{\alpha}_0\|_0 < s$*. Furthermore, the covariance of any estimator whose* $\mathcal{T}$*-bias gradient matrix is* $\mathcal{B}\boldsymbol{U}$ *satisfies*

$$\begin{aligned} \mathrm{Cov}(\widehat{\boldsymbol{\alpha}}) &\succeq \sigma^2(\boldsymbol{I}+\mathcal{B})(\boldsymbol{H}^T\boldsymbol{H})^{\dagger}(\boldsymbol{I}+\mathcal{B}^T) \\ &\qquad\qquad \textit{when } \|\boldsymbol{\alpha}_0\|_0 < s, \\ \mathrm{Cov}(\widehat{\boldsymbol{\alpha}}) &\succeq \sigma^2(\boldsymbol{U}+\mathcal{B}\boldsymbol{U})(\boldsymbol{H}_{\boldsymbol{\alpha}_0}^T\boldsymbol{H}_{\boldsymbol{\alpha}_0})^{-1}(\boldsymbol{U}+\mathcal{B}\boldsymbol{U})^T \\ &\qquad\qquad \textit{when } \|\boldsymbol{\alpha}_0\|_0 = s. \end{aligned} \tag{4.40}$$

Here, $\boldsymbol{H}_{\boldsymbol{\alpha}_0}$ *is the matrix containing the columns of* $\boldsymbol{H}$ *corresponding to* $\mathrm{supp}(\boldsymbol{\alpha}_0)$*.*

Let us examine Theorem 4.2 separately in the underdetermined and well-determined cases. In the well-determined case, in which $\boldsymbol{H}$ has full row rank, the nullspace of $\boldsymbol{H}$ is trivial, so that (4.39) always holds. It follows that the CRB is always finite, in the sense that we cannot rule out the existence of an estimator having any given bias function. Some insight can be obtained in this case by examining the $\mathcal{T}$-unbiased case. Noting also that $\boldsymbol{H}^T\boldsymbol{H}$ is invertible in the well-determined case, the bound for $\mathcal{T}$-unbiased estimators is given by

$$\begin{aligned} \mathrm{Cov}(\widehat{\boldsymbol{\alpha}}) &\succeq \sigma^2(\boldsymbol{H}^T\boldsymbol{H})^{-1} && \text{when } \|\boldsymbol{\alpha}_0\|_0 < s, \\ \mathrm{Cov}(\widehat{\boldsymbol{\alpha}}) &\succeq \sigma^2\boldsymbol{U}(\boldsymbol{H}_{\boldsymbol{\alpha}_0}^T\boldsymbol{H}_{\boldsymbol{\alpha}_0})^{-1}\boldsymbol{U}^T && \text{when } \|\boldsymbol{\alpha}_0\|_0 = s. \end{aligned} \tag{4.41}$$

From this formulation, the behavior of the CRB can be described as follows. When $\boldsymbol{\alpha}_0$ has non-maximal support ($\|\boldsymbol{\alpha}_0\|_0 < s$), the CRB is identical to the bound which would have been obtained had there been no constraints in the problem. This is because $\boldsymbol{U} = \boldsymbol{I}$ in this case, so that $\mathcal{T}$-unbiasedness and ordinary unbiasedness are equivalent. As we have seen in Section 4.3.1, the CRB is a function of the class of estimators under consideration, so the unconstrained and constrained bounds are equivalent in this situation. The bound $\sigma^2(\boldsymbol{H}^T\boldsymbol{H})^{-1}$ is achieved by the unconstrained LS estimator

$$\widehat{\boldsymbol{\alpha}} = (\boldsymbol{H}^T\boldsymbol{H})^{-1}\boldsymbol{H}^T\boldsymbol{y} \tag{4.42}$$

which is the minimum variance unbiased estimator in the unconstrained case. Thus, we learn from Theorem 4.2 that for values of $\boldsymbol{\alpha}_0$ having non-maximal support, no $\mathcal{T}$-unbiased technique

can outperform the standard LS estimator, which does not assume any knowledge about the constraint set $\mathcal{T}$.

On the other hand, consider the case in which $\boldsymbol{\alpha}_0$ has maximal support, i.e., $\|\boldsymbol{\alpha}_0\|_0 = s$. Suppose first that $\mathrm{supp}(\boldsymbol{\alpha}_0)$ is known, so that one must estimate only the nonzero values of $\boldsymbol{\alpha}_0$. In this case, a reasonable approach is to use the oracle estimator (4.13), whose covariance matrix is given by $\sigma^2 \boldsymbol{U}(\boldsymbol{H}_{\boldsymbol{\alpha}_0}^T \boldsymbol{H}_{\boldsymbol{\alpha}_0})^{-1} \boldsymbol{U}^T$ [8]. Thus, when $\boldsymbol{\alpha}_0$ has maximal support, Theorem 4.2 states that $\mathcal{T}$-unbiased estimators can perform, at best, as well as the oracle estimator, which is equivalent to the LS approach when the support of $\boldsymbol{\alpha}_0$ is known.

The situation is similar, but somewhat more involved, in the underdetermined case. Here, the condition (4.39) for the existence of an estimator having a given bias gradient matrix no longer automatically holds. To interpret this condition, it is helpful to introduce the mean gradient matrix $\boldsymbol{M}(\boldsymbol{\alpha})$, defined as

$$\boldsymbol{M}(\boldsymbol{\alpha}) = \frac{\partial E\{\widehat{\boldsymbol{\alpha}}\}}{\partial \boldsymbol{\alpha}} = \boldsymbol{I} + \mathcal{B}. \tag{4.43}$$

The matrix $\boldsymbol{M}(\boldsymbol{\alpha})$ is a measure of the sensitivity of an estimator to changes in the parameter vector. For example, a $\mathcal{T}$-unbiased estimator is sensitive to any *feasible* change in $\boldsymbol{\alpha}$. Thus, $\mathcal{N}(\boldsymbol{M})$ denotes the subspace of directions to which $\widehat{\boldsymbol{\alpha}}$ is insensitive. Likewise, $\mathcal{N}(\boldsymbol{H})$ is the subspace of directions for which a change in $\boldsymbol{\alpha}$ does not modify $\boldsymbol{H}\boldsymbol{\alpha}$. The condition (4.39) therefore states that for an estimator to exist, it must be insensitive to changes in $\boldsymbol{\alpha}$ which are unobservable through $\boldsymbol{H}\boldsymbol{\alpha}$, at least when $\|\boldsymbol{\alpha}\|_0 < s$. No such requirement is imposed in the case $\|\boldsymbol{\alpha}\|_0 = s$, since in this case there are far fewer feasible directions.

The lower bound (4.40) is similarly a consequence of the wide range of feasible directions obtained when $\|\boldsymbol{\alpha}\|_0 < s$, as opposed to the tight constraints when $\|\boldsymbol{\alpha}\|_0 = s$. Specifically, when $\|\boldsymbol{\alpha}\|_0 < s$, a change to any component of $\boldsymbol{\alpha}$ is feasible and hence the lower bound equals that of an unconstrained estimation problem, with the FIM given by $\sigma^{-2}\boldsymbol{H}^T\boldsymbol{H}$. On the other hand, when $\|\boldsymbol{\alpha}\|_0 = s$, the bound is effectively that of an estimator with knowledge of the particular subspace to which $\boldsymbol{\alpha}$ belongs; for this subspace the FIM is the submatrix $\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}$ given in (4.35). This phenomenon is discussed further in Section 4.6.

Another difference between the well-determined and underdetermined cases is that when $\boldsymbol{H}$ is underdetermined, an estimator cannot be $\mathcal{T}$-unbiased for all $\boldsymbol{\alpha}$. To see this, recall from (4.21) that $\mathcal{T}$-unbiased estimators are defined by the fact that $\mathcal{B}\boldsymbol{U} = \boldsymbol{0}$. When $\|\boldsymbol{\alpha}\|_0 < s$, we have $\boldsymbol{U} = \boldsymbol{I}$ and thus $\mathcal{T}$-unbiasedness implies $\mathcal{B} = \boldsymbol{0}$, so that $\mathcal{N}(\boldsymbol{I} + \mathcal{B}) = \{\boldsymbol{0}\}$. But since $\boldsymbol{H}$ is underdetermined, $\mathcal{N}(\boldsymbol{H})$ is nontrivial. Consequently, (4.39) cannot hold for $\mathcal{T}$-unbiased

estimators when $\|\boldsymbol{\alpha}\|_0 < s$.

The lack of $\mathcal{T}$-unbiased estimators when $\|\boldsymbol{\alpha}_0\|_0 < s$ is a direct consequence of the fact that the feasible direction set at such $\boldsymbol{\alpha}_0$ contains all of the directions $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_p$. The conclusion from Theorem 4.2 is then that no estimator can be expected to be unbiased in such a high-dimensional neighborhood, just as unbiased estimation is impossible in the p-dimensional neighborhood $\mathcal{B}_\varepsilon(\boldsymbol{\alpha}_0)$, as explained in Section 4.3.1. However, it is still possible to obtain a finite CRB in this setting by further restricting the constraint set: if it is known that $\|\boldsymbol{\alpha}_0\|_0 = \tilde{s} < s$, then one can redefine $\mathcal{T}$ in (4.5) by replacing s with $\tilde{s}$. This will enlarge the class of estimators considered $\mathcal{T}$-unbiased, and Theorem 4.2 would then provide a finite lower bound on those estimators. Such estimators will not, however, be unbiased in the sense implied by the original constraint set.

While an estimator cannot be unbiased for *all* $\boldsymbol{\alpha} \in \mathcal{T}$, unbiasedness is possible at points $\boldsymbol{\alpha}$ for which $\|\boldsymbol{\alpha}\|_0 = s$. In this case, Theorem 4.2 produces a bound on the MSE of a $\mathcal{T}$-unbiased estimator, obtained by calculating the trace of (4.40) in the case $\mathcal{B}\boldsymbol{U} = \boldsymbol{0}$. This bound is given by

$$E\{\|\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0\|_2^2\} \geq \sigma^2 \operatorname{Tr}((\boldsymbol{H}_{\boldsymbol{\alpha}_0}^T \boldsymbol{H}_{\boldsymbol{\alpha}_0})^{-1}), \quad \|\boldsymbol{\alpha}_0\|_0 = s. \tag{4.44}$$

The most striking feature of (4.44) is that it is identical to the oracle MSE (4.14). However, the CRB is of additional importance because of the fact that the ML estimator achieves the CRB in the limit when a large number of independent measurements are available, a situation which is equivalent in our setting to the limit $\sigma \to 0$. In other words, an MSE of (4.44) is achieved at high SNR by the ML approach (4.8), as we will illustrate numerically in Section 4.5. While the ML approach is computationally intractable in the sparse estimation setting, it is still implementable in principle, as opposed to $\widehat{\boldsymbol{\alpha}}_{\text{oracle}}$, which relies on unavailable information (namely, the support set of $\boldsymbol{\alpha}_0$). Thus, Theorem 4.1 gives an alternative interpretation to comparisons of estimator performance with the oracle.

Observe that the bound (4.44) depends on the value of $\boldsymbol{\alpha}_0$ (through its support set, which defines $\boldsymbol{H}_{\boldsymbol{\alpha}_0}$). This implies that some values of $\boldsymbol{\alpha}_0$ are more difficult to estimate than others. For example, suppose the ℓ_2 norms of some of the columns of $\boldsymbol{H}$ are significantly larger than the remaining columns. Measurements of a parameter $\boldsymbol{\alpha}_0$ whose support corresponds to the large-norm columns of $\boldsymbol{H}$ will then have a much higher SNR than measurements of a parameter corresponding to small-norm columns, and this will clearly affect the accuracy with which $\boldsymbol{\alpha}_0$ can be estimated. To analyze the behavior beyond this effect, it is common to consider the situation in which the columns $\boldsymbol{h}_i$ of $\boldsymbol{H}$ are normalized so that $\|\boldsymbol{h}_i\|_2 = 1$. In this case, for

sufficiently incoherent dictionaries, $\mathrm{Tr}((\boldsymbol{H}_{\boldsymbol{\alpha}_0}^T \boldsymbol{H}_{\boldsymbol{\alpha}_0})^{-1})$ is bounded above and below by a small constant times s, so that the CRB is similar for all values of $\boldsymbol{\alpha}_0$. To see this, let μ be the coherence of $\boldsymbol{H}$ [22], defined (for $\boldsymbol{H}$ having normalized columns) as

$$\mu \triangleq \max_{i \neq j} \left| \boldsymbol{h}_i^T \boldsymbol{h}_j \right|. \tag{4.45}$$

By the Gershgorin disc theorem, the eigenvalues of $\boldsymbol{H}_{\boldsymbol{\alpha}_0}^T \boldsymbol{H}_{\boldsymbol{\alpha}_0}$ are in the range $[1 - s\mu, 1 + s\mu]$. It follows that the unbiased CRB (4.44) is bounded above and below by

$$\frac{s\sigma^2}{1 + s\mu} \leq \sigma^2 \,\mathrm{Tr}((\boldsymbol{H}_{\boldsymbol{\alpha}_0}^T \boldsymbol{H}_{\boldsymbol{\alpha}_0})^{-1}) \leq \frac{s\sigma^2}{1 - s\mu}. \tag{4.46}$$

Thus, when s is somewhat smaller than $1/\mu$, the CRB is roughly equal to $s\sigma^2$ for all values of $\boldsymbol{\alpha}_0$. As we have seen in Section 4.2.2, for sufficiently small s, the worst-case MSE of practical estimators, such as BPDN and the DS, is $O(s\sigma^2 \log p)$. Thus, practical estimators come almost within a constant of the unbiased CRB, implying that they are close to optimal for all values of $\boldsymbol{\alpha}_0$, at least when compared with unbiased techniques.

4.4.2 Denoising and Deblurring

We next consider the problem (4.1), in which it is required to estimate not the sparse vector $\boldsymbol{\alpha}_0$ itself, but rather the vector $\boldsymbol{x}_0 = \boldsymbol{D}\boldsymbol{\alpha}_0$, where $\boldsymbol{D}$ is a known dictionary matrix. Thus, $\boldsymbol{x}_0$ belongs to the set $\mathcal{S}$ of (4.2). We assume for concreteness that $\boldsymbol{D}$ has full row rank and that $\boldsymbol{A}$ has full column rank. This setting encompasses the denoising and deblurring problems described in Section 4.2.1, with the former arising when $\boldsymbol{A} = \boldsymbol{I}$ and the latter obtained when $\boldsymbol{A}$ represents a blurring kernel. Similar calculations can be carried out when $\boldsymbol{A}$ is rank-deficient, a situation which occurs, for example, in some interpolation problems.

Recall from Section 4.2.1 the assumption that every $\boldsymbol{x} \in \mathcal{S}$ has a *unique* representation $\boldsymbol{x} = \boldsymbol{D}\boldsymbol{\alpha}$ for which $\boldsymbol{\alpha}$ is in the set $\mathcal{T}$ of (4.5). We denote by $\boldsymbol{r}(\cdot)$ the mapping from $\mathcal{S}$ to $\mathcal{T}$ which returns this representation. In other words, $\boldsymbol{r}(\boldsymbol{x})$ is the unique vector in $\mathcal{T}$ for which

$$\boldsymbol{x} = \boldsymbol{D}\boldsymbol{r}(\boldsymbol{x}) \quad \text{and} \quad \|\boldsymbol{r}(\boldsymbol{x})\|_0 \leq s. \tag{4.47}$$

Note that while the mapping $\boldsymbol{r}$ is well-defined, actually calculating the value of $\boldsymbol{r}(\boldsymbol{x})$ for a given vector $\boldsymbol{x}$ is, in general, NP-hard.

In the current setting, unlike the scenario of Section 4.4.1, it is always possible to construct an unbiased estimator. Indeed, even without imposing the constraint (4.2), there exists an

unbiased estimator. This is the LS or maximum likelihood estimator, given by

$$\hat{\boldsymbol{x}} = (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{A}^T\boldsymbol{y}. \tag{4.48}$$

A standard calculation demonstrates that the covariance of $\hat{\boldsymbol{x}}$ is

$$\sigma^2(\boldsymbol{A}^T\boldsymbol{A})^{-1}. \tag{4.49}$$

On the other hand, the FIM for the setting (4.1) is given by

$$\boldsymbol{J} = \frac{1}{\sigma^2}\boldsymbol{A}^T\boldsymbol{A}. \tag{4.50}$$

Since $\boldsymbol{A}$ has full row rank, the FIM is invertible. Consequently, it is seen from (4.49) and (4.50) that the LS approach achieves the CRB $\boldsymbol{J}^{-1}$ for unbiased estimators. This well-known property demonstrates that in the unconstrained setting, the LS technique is optimal among all unbiased estimators.

The LS estimator, like any unbiased approach, is also $\mathcal{S}$-unbiased. However, with the addition of the constraint $\boldsymbol{x}_0 \in \mathcal{S}$, one would expect to obtain improved performance. It is therefore of interest to obtain the CRB for the constrained setting. To this end, we first note that since $\boldsymbol{J}$ is invertible, we have $\mathcal{R}\big(\boldsymbol{U}\boldsymbol{U}^T\boldsymbol{J}\boldsymbol{U}\boldsymbol{U}^T\big) = \mathcal{R}\big(\boldsymbol{U}\boldsymbol{U}^T\big)$ for any $\boldsymbol{U}$, and consequently (4.26) holds for any matrix $\mathcal{B}$. The bound (4.27) of Theorem 4.1 thus applies regardless of the bias gradient matrix.

For simplicity, in the following we derive the CRB for $\mathcal{S}$-unbiased estimators. A calculation for arbitrary $\mathcal{S}$-bias functions can be performed along similar lines. Consider first values $\boldsymbol{x} \in \mathcal{S}$ such that $\|\boldsymbol{r}(\boldsymbol{x})\|_0 < s$. Then, $\|\boldsymbol{r}(\boldsymbol{x}) + \delta\boldsymbol{e}_i\|_0 \leq s$ for any δ and for any $\boldsymbol{e}_i$. Therefore,

$$\boldsymbol{x} + \delta\boldsymbol{D}\boldsymbol{e}_i \in \mathcal{S} \tag{4.51}$$

for any δ and $\boldsymbol{e}_i$. In other words, the feasible directions include all columns of $\boldsymbol{D}$. Since it is assumed that $\boldsymbol{D}$ has full row rank, this implies that the feasible subspace $\mathcal{F}$ equals $\mathbb{R}^n$, and the matrix $\boldsymbol{U}$ of (4.20) can be chosen as $\boldsymbol{U} = \boldsymbol{I}$.

Next, consider values $\boldsymbol{x} \in \mathcal{S}$ for which $\|\boldsymbol{r}(\boldsymbol{x})\|_0 = s$. Then, for sufficiently small $\delta > 0$, we have $\|\boldsymbol{r}(\boldsymbol{x}) + \delta\boldsymbol{v}\|_0 \leq s$ if and only if $\boldsymbol{v} = \boldsymbol{e}_i$ for some $i \in \mathrm{supp}(\boldsymbol{r}(\boldsymbol{x}))$. Equivalently,

$$\boldsymbol{x} + \delta\boldsymbol{v} \in \mathcal{S} \text{ if and only if } \boldsymbol{v} = \boldsymbol{D}\boldsymbol{e}_i \text{ and } i \in \mathrm{supp}(\boldsymbol{r}(\boldsymbol{x})). \tag{4.52}$$

Consequently, the feasible direction subspace in this case corresponds to the column space of the matrix $\boldsymbol{D}_{\boldsymbol{x}}$ containing the s columns of $\boldsymbol{D}$ indexed by $\mathrm{supp}(\boldsymbol{r}(\boldsymbol{x}))$. From (4.7) we have

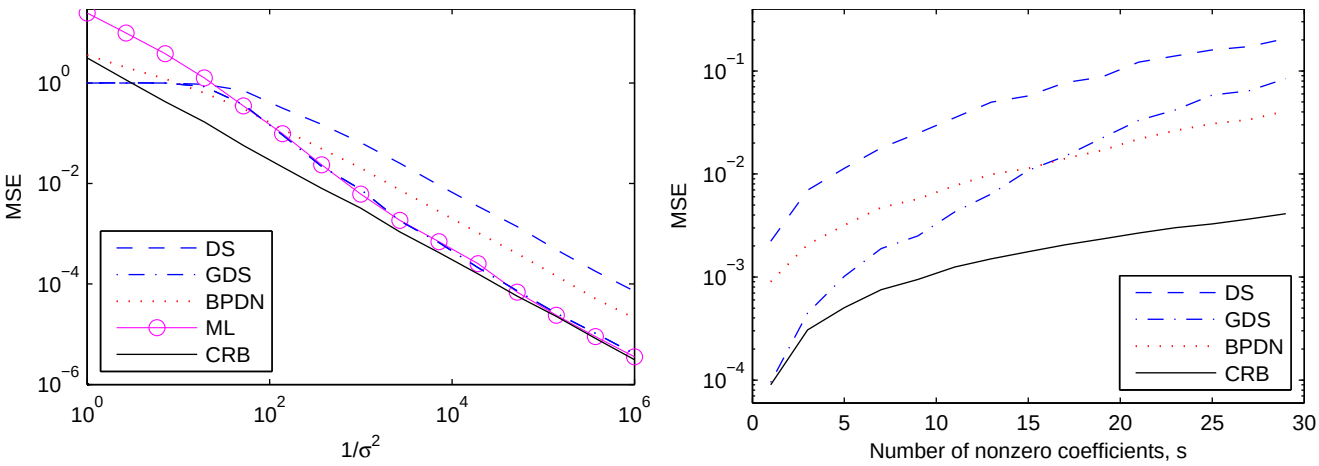

Figure 4.2: MSE of various estimators compared with the unbiased CRB (4.44), for (a) varying SNR and (b) varying sparsity levels.

$\text{spark}(\boldsymbol{D}) > s$, and therefore the columns of $\boldsymbol{D}_x$ are linearly independent. Thus the orthogonal projector onto $\mathcal{F}$ is given by

$$\boldsymbol{P} \triangleq \boldsymbol{U}\boldsymbol{U}^T = \boldsymbol{D}_x(\boldsymbol{D}_x^T\boldsymbol{D}_x)^{-1}\boldsymbol{D}_x^T. \tag{4.53}$$

Combining these calculations with Theorem 4.1 yields the following result.

Theorem 4.3. *Consider the estimation setting* (4.1) *with the constraint* (4.2)*, and suppose* $\text{spark}(\boldsymbol{D}) > 2s$*. Let* $\hat{\boldsymbol{x}}$ *be a finite-variance* $\mathcal{S}$*-unbiased estimator. Then,*

$$\begin{aligned} &\text{Cov}(\hat{\boldsymbol{x}}) \succeq \sigma^2(\boldsymbol{A}^T\boldsymbol{A})^{-1} && \text{when } \|\boldsymbol{r}(\boldsymbol{x})\|_0 < s, \\ &\text{Cov}(\hat{\boldsymbol{x}}) \succeq \sigma^2\left(\boldsymbol{P}\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{P}\right)^\dagger && \text{when } \|\boldsymbol{r}(\boldsymbol{x})\|_0 = s. \end{aligned} \tag{4.54}$$

Here, $\boldsymbol{P}$ *is given by* (4.53)*, in which* $\boldsymbol{D}_x$ *is the* $n \times s$ *matrix consisting of the columns of* $\boldsymbol{D}$ *participating in the (unique)* s*-element representation* $\boldsymbol{D}\boldsymbol{\alpha}$ *of* $\boldsymbol{x}$*.*

As in Theorem 4.2, the bound exhibits a dichotomy between points having maximal and non-maximal support. In the former case, the CRB is equivalent to the bound obtained when the support set is known, whereas in the latter the bound is equivalent to an unconstrained CRB. This point is discussed further in Section 4.6.

4.5 Numerical Results

In this section, we demonstrate the use of the CRB for measuring the achievable MSE in the sparse estimation problem (4.4). To this end, a series of simulations was performed. In each simulation, a random 100×200 dictionary $\boldsymbol{H}$ was constructed from a zero-mean Gaussian IID distribution, whose columns $\boldsymbol{h}_i$ were normalized so that $\|\boldsymbol{h}_i\|_2 = 1$. A parameter $\boldsymbol{\alpha}_0$ was then selected by choosing a support uniformly at random and selecting the nonzero elements as Gaussian IID variables with mean 0 and variance 1. Noisy measurements $\boldsymbol{y}$ were obtained from (4.4), and $\boldsymbol{\alpha}_0$ was then estimated using BPDN (4.9), the DS (4.10), and the GDS (4.11). The regularization parameters were chosen as $\tau = 2\sigma\sqrt{\log p}$ and $\gamma = 4\sigma\sqrt{\log(p-s)}$, rules of thumb which are motivated by a theoretical analysis [78]. The MSE of each estimate was then calculated by repeating this process with different realizations of the random variables. The unbiased CRB was calculated using (4.44). In this case, the unbiased CRB equals the MSE of the oracle estimator (4.13), but as we will see below, interpreting (4.44) as a bound on unbiased estimators provides further insight into the estimation problem.

A first set of experiments was conducted to examine the CRB at various SNR levels. In this simulation, the ML estimator (4.8) was also computed, in order to verify its convergence to the CRB at high SNR. Since the ML approach is computationally prohibitive when p and s are large, this necessitated the selection of the rather low support size $s = 3$. The MSE and CRB were calculated for 15 SNR values by changing the noise standard deviation σ between 1 and 10^{-3}. The MSE of the ML approach, as well as the other estimators of Section 4.2.2, is compared with the CRB in Fig. 4.2(a). The convergence of the ML estimator to the CRB is clearly visible in this figure. The performance of the GDS is also impressive, being as good or better than the ML approach. Apparently, at high SNR, the DS tends to correctly recover the true support set, in which case GDS (4.11) equals the oracle (4.13). Perhaps surprisingly, applying a LS estimate on the support set obtained by BPDN (which could be called a "Gauss–BPDN" strategy) does not work well at all, and in fact results in higher MSE than a direct application of BPDN. (The results for the Gauss–BPDN method are not plotted in Fig. 4.2.)

Note that some estimation techniques outperform the oracle MSE (or CRB) at low SNR. It may appear surprising that a practical technique such as the DS outperforms the oracle. The explanation for this stems from the fact that the CRB (4.44) is a lower bound on the MSE of *unbiased* estimators. The bias of most estimators tends to be negligible in low-noise settings, but often increases with the noise variance σ^2. Indeed, when σ^2 is as large as $\|\boldsymbol{\alpha}_0\|_2^2$, the measure-

ments carry very little useful information about $\boldsymbol{\alpha}_0$, and an estimator can improve performance by shrinkage. Such a strategy, while clearly biased, yields lower MSE than a naive reliance on the noisy measurements. This is indeed the behavior of the DS and BPDN, since for large σ^2, the ℓ_1 regularization becomes the dominant term, resulting in heavy shrinkage. Consequently, it is to be expected that such techniques will outperform even the best unbiased estimator at low SNR, as indeed occurs in Fig. 4.2(a).

The performance of the estimators of Section 4.2.2, excluding the ML method, was also compared for varying sparsity levels. To this end, the simulation was repeated for 15 support sizes in the range $1 \leq s \leq 30$, with a constant noise standard deviation of $\sigma = 0.01$. The results are plotted in Fig. 4.2(b). While a substantial gap exists between the CRB and the MSE of the practical estimators in this case, the general trend in both cases describes a similar rate of increase as s grows. Interestingly, a drawback of the GDS approach is visible in this setting: as s increases, correct support recovery becomes more difficult, and shrinkage becomes a valuable asset for reducing the sensitivity of the estimate to random measurement fluctuations. The LS approach practiced by the GDS, which does not perform shrinkage, leads to gradual performance deterioration.

Results similar to Fig. 4.2 were obtained for a variety of related estimation scenarios, including several deterministic, rather than random, dictionaries $\boldsymbol{H}$.

4.6 Discussion

In this paper, we extended the CRB to constraint sets satisfying the local balance condition (Theorem 4.1). This enabled us to derive lower bounds on the achievable performance in various estimation problems (Theorems 4.2 and 4.3). In simple terms, Theorems 4.2 and 4.3 can be summarized as follows. The behavior of the CRB differs depending on whether or not the parameter has maximal support (i.e., $\|\boldsymbol{\alpha}\|_0 = s$). In the case of maximal support, the bound equals that which would be obtained if the sparsity pattern were known; this can be considered an "oracle bound". On the other hand, when $\|\boldsymbol{\alpha}\|_0 < s$, performance is identical to the unconstrained case, and the bound is substantially higher. We now discuss some practical implications of these conclusions. To simplify the discussion, we consider the case of unbiased estimators, though analogous conclusions can be drawn for any bias function.

When $\|\boldsymbol{\alpha}\|_0 = s$ and all nonzero elements of $\boldsymbol{\alpha}$ are considerably larger than the standard deviation of the noise, the support set can be recovered correctly with high probability (at least

if computational considerations are ignored). Thus, in this case an estimator can mimic the behavior of the oracle, and the CRB is expected to be tight. Indeed, in the high SNR limit, the ML estimator achieves the unbiased CRB. On the other hand, when the support of $\boldsymbol{\alpha}$ is not maximal, the unbiasedness requirement demands sensitivity to changes in all components of $\boldsymbol{\alpha}$, and consequently the bound coincides with the unconstrained CRB. Thus, as claimed in Section 4.3, in underdetermined cases no estimator is unbiased for all $\boldsymbol{\alpha} \in \mathcal{S}$.

An interesting observation can also be made concerning maximal-support points $\boldsymbol{\alpha}$ for which some of the nonzero elements are close to zero. The CRB in this "low-SNR" case corresponds to the oracle MSE, but as we will see, the bound is loose for such values of $\boldsymbol{\alpha}$. Intuitively, at low-SNR points, any attempt to recover the sparsity pattern will occasionally fail. Consequently, despite the optimistic CRB, it is unlikely that the oracle MSE can be achieved. Indeed, the covariance matrix of any finite-variance estimator is a continuous function of $\boldsymbol{\alpha}$ [6], and the fact that performance is bounded by the (much higher) unconstrained bound when $\|\boldsymbol{\alpha}\|_0 < s$ implies that performance must be similarly poor for low SNR.

This excessive optimism is a result of the local nature of the CRB: The bound is a function of the estimation setting only in an ε-neighborhood of the parameter itself. Indeed, the CRB depends on the constraint set only through the feasible directions, which were defined in Section 4.3.2 as those directions which do not violate the constraints for *sufficiently small* deviations. Thus, for the CRB, it is entirely irrelevant if some of the components of $\boldsymbol{\alpha}$ are close to zero, as long as $\operatorname{supp}(\boldsymbol{\alpha})$ is held constant.

A tighter bound for sparse estimation problems may be obtained using the Hammersley–Chapman–Robbins (HCR) approach [50–52], which depends on the constraints at points beyond the local neighborhood of $\boldsymbol{x}$. Such a bound is likely to yield tighter results for low SNR values, and will create a smooth transition between the regions of maximal and non-maximal support. However, the bound will depend on more complex properties of the estimation setting, such as the distance between $\boldsymbol{D\alpha}$ and feasible points with differing supports. The derivation of such a bound is a subject for further research.

Acknowledgement

The authors would like to thank Yaniv Plan for helpful discussions.

Chapter 5

Bounds for Sparse Estimation with a Unitary Dictionary

This chapter has been submitted for publication as:

- A. Jung, Z. Ben-Haim, F. Hlawatsch, and Y. C. Eldar, "Unbiased estimation of a sparse vector in white Gaussian noise," submitted to *IEEE Trans. Information Theory*, May 2010.

5.1 Introduction

Research in the past few years has led to a recognition that the performance of signal processing algorithms can be boosted by exploiting the tendency of many signals to have sparse representations. Applications of this principle include signal reconstruction (e.g. in the context of compressed sensing [84,85]) and signal enhancement (e.g. in the context of image denoising and deblurring [75,76,86]).

In this work, we consider the estimation of an S-sparse, finite-dimensional vector $\boldsymbol{x} \in \mathbb{R}^N$. By "$S$-sparse" we mean that the vector $\boldsymbol{x}$ has at most S nonzero entries, which is denoted by $\|\boldsymbol{x}\|_0 \triangleq |\mathrm{supp}(\boldsymbol{x})| \leq S$, where $\mathrm{supp}(\boldsymbol{x})$ denotes the set of indices of the nonzero entries of $\boldsymbol{x}$. The "sparsity" S is assumed to be known, and typically $S \ll N$. However, the positions of the nonzero entries (i.e., $\mathrm{supp}(\boldsymbol{x})$) as well as the values of the nonzero entries are unknown. We investigate how much we can gain in estimation accuracy by knowing *a priori* that the vector $\boldsymbol{x}$ is S-sparse. We will use the frequentist setting [6] of estimation theory, i.e., we will model $\boldsymbol{x}$ as unknown but deterministic. This is in contrast to Bayesian estimation theory, where one treats $\boldsymbol{x}$ as a random vector whose probability density function (pdf) or certain moments thereof are

assumed to be known. In the Bayesian setting, the sparsity can be modeled by using a pdf that favors sparse vectors, see e.g. [87–89].

A fundamental concept in the frequentist setting is that of unbiasedness [6, 16, 90]. An unbiased estimator is one whose expectation always equals the true underlying vector $\boldsymbol{x}$. The restriction to unbiased estimators is important as it excludes trivial and practically useless estimators, and it allows us to study the difficulty of the estimation problem using established techniques such as the Cramér–Rao bound (CRB) [16, 45, 90]. Another justification of unbiasedness is that for typical estimation problems, when the variance of the noise is low, it is necessary for an estimator to be unbiased in order to achieve a small mean-squared estimation error (MSE) [6].

These reasons notwithstanding, there is no guarantee that unbiased estimators are necessarily optimal. In fact, in many settings, including the scenario described in this paper, there exist biased estimators which are strictly better than any unbiased technique in terms of MSE [42, 44, 91]. Nevertheless, for simplicity and because of the reasons stated above, we focus on bounds for unbiased estimation in this work. As we will see, bounds on unbiased techniques give some indication of the general difficulty of the setting, and as such some of our conclusions will be shown empirically to characterize biased techniques as well.

Our main contribution is a characterization of the optimal performance of unbiased estimators $\hat{\boldsymbol{x}}(\boldsymbol{y})$ that are based on observing

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{n} \tag{5.1}$$

where $\boldsymbol{A} \in \mathbb{R}^{M\times N}$ ($M \geq N$) is a known matrix with orthonormal columns, i.e., $\boldsymbol{A}^T\boldsymbol{A} = \boldsymbol{I}_N$, and $\boldsymbol{n} \sim \mathcal{N}(\boldsymbol{0}, \sigma^2\boldsymbol{I}_M)$ denotes zero-mean white Gaussian noise with known variance σ^2 (here, $\boldsymbol{I}_N$ denotes the identity matrix of size $N \times N$). Note that without loss of generality we can then assume that $\boldsymbol{A} = \boldsymbol{I}_N$ and $M = N$, i.e., $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{n}$, since premultiplication of the model (5.1) by $\boldsymbol{A}^T$ will reduce the estimation problem to an equivalent problem $\boldsymbol{y}' = \boldsymbol{A}'\boldsymbol{x} + \boldsymbol{n}'$ in which $\boldsymbol{A}' = \boldsymbol{A}^T\boldsymbol{A} = \boldsymbol{I}_N$ and the noise $\boldsymbol{n}' = \boldsymbol{A}^T\boldsymbol{n}$ is again zero-mean white Gaussian with variance σ^2. Such a sparse signal model can be used, e.g., for channel estimation [92] when the channel consists only of few significant taps and an orthogonal training signal is used [93]. Another application that fits our scope is image denoising using an orthonormal wavelet basis [86]. We note that parts of this work were previously presented in [94].

The estimation problem (5.1) with $\boldsymbol{A} = \boldsymbol{I}_N$ was studied by Donoho and Johnstone [95, 96]. Their work was aimed at demonstrating asymptotic minimax optimality, i.e., they considered

estimators having optimal worst-case behavior when the problem dimensions N, S tend to infinity. By contrast, we consider the finite-dimensional setting, and attempt to characterize the performance at each value of $\boldsymbol{x}$, rather than analyzing worst-case behavior. Such a "pointwise" approach was also advocated by the authors of [70, 79], who studied the CRB for the sparse linear model (5.1) with arbitrary $\boldsymbol{A}$. However, the CRB is a local bound, in the sense that the performance characterization it provides is only based on the statistical properties in the neighborhood of the specific value of $\boldsymbol{x}$ being examined. In particular, the CRB for a given $\boldsymbol{x}$ is only based on a *local* unbiasedness assumption, meaning that the estimator is only required to be unbiased at $\boldsymbol{x}$ and in its infinitesimal neighborhood. Our goal in this paper is to obtain performance bounds for the more restrictive case of globally unbiased estimators, i.e., estimators whose expectation equals the true $\boldsymbol{x}$ for each S-sparse vector $\boldsymbol{x}$. Since any globally unbiased estimator is also locally unbiased, our lower bounds will be tighter than those of [70, 79].

Our contributions and the organization of this paper can be summarized as follows. In Section 5.2, we show that whereas only one unbiased estimator exists for the ordinary (nonsparse) signal in noise model, there are infinitely many unbiased estimators for the sparse signal in noise model; on the other hand, none of them has uniformly minimum variance. In Sections 5.3 and 5.4, we characterize the performance of *locally* minimum variance unbiased estimators by providing, respectively, lower and upper bounds on their mean-squared error (MSE). These bounds can equivalently be viewed as lower and upper bounds on the *Barankin bound* [52, 53]. Finally, numerical studies exploring and extending our performance bounds and comparing them with established estimator designs are presented in Section 5.5.

Notation: Throughout the paper, boldface lowercase letters (e.g., $\boldsymbol{x}$) denote column vectors while boldface uppercase letters (e.g., $\boldsymbol{M}$) denote matrices. We denote by $\mathrm{tr}(\boldsymbol{M})$, $\boldsymbol{M}^T$, and $\boldsymbol{M}^\dagger$ the trace, transpose, and Moore-Penrose pseudoinverse of $\boldsymbol{M}$, respectively. The identity matrix of size $N \times N$ is denoted by $\boldsymbol{I}_N$. The notation $\boldsymbol{M} \succeq \boldsymbol{N}$ indicates that $\boldsymbol{M} - \boldsymbol{N}$ is a positive semidefinite matrix. The set of indices of the nonzero entries of a vector $\boldsymbol{x}$ is denoted by $\mathrm{supp}(\boldsymbol{x})$, and $\|\boldsymbol{x}\|_0$ is defined as the size of this set. The kth entry of $\boldsymbol{x}$ is written x_k. We also use the signum function of a real number y, $\mathrm{sgn}(y) \triangleq y/|y|$. The sets of nonnegative, nonpositive, and positive real numbers will be denoted by $\mathbb{R}_+$, $\mathbb{R}_-$, and $\mathbb{R}_{++}$, respectively.

5.2 The Sparse Signal in Noise Model

5.2.1 Problem Setting

Let $\boldsymbol{x} \in \mathbb{R}^N$ be an unknown deterministic vector which is known to be S-sparse, i.e.,

$$\boldsymbol{x} \in \mathcal{X}_S\,, \qquad \text{with } \mathcal{X}_S \triangleq \{\boldsymbol{x} \in \mathbb{R}^N : \|\boldsymbol{x}\|_0 \leq S\}\,.$$

The vector $\boldsymbol{x}$ is to be estimated based on the observation of a vector $\boldsymbol{y}$ which is the sum of $\boldsymbol{x}$ and zero-mean white Gaussian noise. Thus

$$\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{n}\,, \qquad \text{with } \boldsymbol{x} \in \mathcal{X}_S\,,\ \boldsymbol{n} \sim \mathcal{N}(\boldsymbol{0}, \sigma^2 \boldsymbol{I}_N) \tag{5.2}$$

where the noise variance σ^2 is assumed to be nonzero and known. It follows that the pdf of $\boldsymbol{y}$, parameterized by the deterministic but unknown parameter $\boldsymbol{x} \in \mathcal{X}_S$, is

$$f(\boldsymbol{y}; \boldsymbol{x}) \;=\; \frac{1}{(2\pi\sigma^2)^{N/2}} \exp\left(-\frac{1}{2\sigma^2}\|\boldsymbol{y}-\boldsymbol{x}\|_2^2\right). \tag{5.3}$$

We refer to (5.2) as the *sparse signal in noise model* (SSNM). As explained previously, settings of the form (5.1) with an orthonormal matrix $\boldsymbol{A}$ can be converted to the SSNM (5.2). The case $S = N$ corresponds to the situation in which no sparsity assumption is made. As we will see, this case is fundamentally different from the sparse setting $S < N$, which is our focus in this paper.

An *estimator* $\hat{\boldsymbol{x}}(\boldsymbol{y})$ of the parameter $\boldsymbol{x}$ is a function that maps (a realization of) the observation $\boldsymbol{y}$ to (a realization of) the estimated vector $\hat{\boldsymbol{x}}$, i.e.,

$$\hat{\boldsymbol{x}}(\cdot) : \mathbb{R}^N \to \mathbb{R}^N : \boldsymbol{y} \mapsto \hat{\boldsymbol{x}}.$$

With an abuse of notation, we will use the symbol $\hat{\boldsymbol{x}}$ for both the estimator (which is a function) and the estimate (a specific function value). The meaning should be clear from the context. The question now is how we can exploit the information that $\boldsymbol{x}$ is S-sparse in order to construct "good" estimators. Our measure of the quality of an estimator $\hat{\boldsymbol{x}}(\cdot)$ for a given parameter value $\boldsymbol{x} \in \mathcal{X}_S$ will be the estimator's MSE, which is defined as

$$\varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}}) \;\triangleq\; \mathsf{E}_{\boldsymbol{x}}\big\{\|\hat{\boldsymbol{x}}(\boldsymbol{y}) - \boldsymbol{x}\|_2^2\big\}\,.$$

Here, the notation $\mathsf{E}_{\boldsymbol{x}}\{\cdot\}$ means that the expectation is taken with respect to the pdf $f(\boldsymbol{y}; \boldsymbol{x})$ of the observation $\boldsymbol{y}$ parameterized by $\boldsymbol{x}$. Note that even though $\boldsymbol{x}$ is known to be S-sparse, the estimates $\hat{\boldsymbol{x}}$ are not constrained to be S-sparse.

The MSE can be written as the sum of a bias term and a variance term, i.e.,

$$\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \;=\; \|\boldsymbol{b}(\boldsymbol{x};\hat{\boldsymbol{x}})\|_2^2 \;+ V(\boldsymbol{x};\hat{\boldsymbol{x}})$$

where the bias $\boldsymbol{b}(\boldsymbol{x};\hat{\boldsymbol{x}}) \triangleq \mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}(\boldsymbol{y})\} - \boldsymbol{x}$ accounts for systematic estimation errors and the variance $V(\boldsymbol{x};\hat{\boldsymbol{x}}) \triangleq \mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y}) - \mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}(\boldsymbol{y})\}\|_2^2\}$ accounts for errors due to random fluctuations of the estimate. Thus, for unbiased estimators ($\boldsymbol{b}(\boldsymbol{x};\hat{\boldsymbol{x}}) = \boldsymbol{0}$ for all $\boldsymbol{x} \in \mathcal{X}_S$), the MSE is equal to the variance, i.e., $\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) = V(\boldsymbol{x};\hat{\boldsymbol{x}})$.

We will also consider the mean power (second moment) of an estimator,

$$P(\boldsymbol{x};\hat{\boldsymbol{x}}) \;\triangleq\; \mathsf{E}_{\boldsymbol{x}}\big\{\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\big\} \;=\; V(\boldsymbol{x};\hat{\boldsymbol{x}}) + \|\mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}(\boldsymbol{y})\}\|_2^2\,. \tag{5.4}$$

For unbiased estimators, $\|\mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}(\boldsymbol{y})\}\|_2^2 = \|\boldsymbol{x}\|_2^2$; thus, minimizing the variance $V(\boldsymbol{x};\hat{\boldsymbol{x}})$ at a fixed $\boldsymbol{x} \in \mathcal{X}_S$ among all unbiased estimators is equivalent to minimizing $P(\boldsymbol{x};\hat{\boldsymbol{x}})$.

5.2.2 Estimator Design

Two well-established estimator designs are the least squares (LS) estimator defined by

$$\hat{\boldsymbol{x}}_{\mathrm{LS}}(\boldsymbol{y}) \;\triangleq\; \operatorname*{arg\,min}_{\boldsymbol{x}' \in \mathcal{X}_S} \|\boldsymbol{y} - \boldsymbol{x}'\|_2^2 \tag{5.5}$$

and the maximum likelihood (ML) estimator defined by

$$\hat{\boldsymbol{x}}_{\mathrm{ML}}(\boldsymbol{y}) \;\triangleq\; \operatorname*{arg\,max}_{\boldsymbol{x}' \in \mathcal{X}_S} f(\boldsymbol{y};\boldsymbol{x}'). \tag{5.6}$$

For the SSNM, due to (5.3), the LS and ML estimators coincide; they are easily seen to be given by

$$\hat{\boldsymbol{x}}_{\mathrm{LS}}(\boldsymbol{y}) = \hat{\boldsymbol{x}}_{\mathrm{ML}}(\boldsymbol{y}) = \mathsf{P}_S(\boldsymbol{y}) \tag{5.7}$$

where P_S is an operator that retains the S largest (in magnitude) components and zeros out all others. The LS/ML estimator is biased unless $S = N$. Note that this estimator is not based on a direct minimization of the MSE. Indeed, if the sparsity constraint is removed ($S = N$) and $N \geq 3$, it has been shown [42, 44, 91] that there exist estimators which yield a better MSE performance than that of the LS/ML estimator.

The MSE $\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}})$ of a specific estimator $\hat{\boldsymbol{x}}(\cdot)$ depends on the value of the parameter $\boldsymbol{x}$. This makes it difficult to define optimality in terms of minimum MSE. For example, if an estimator $\hat{\boldsymbol{x}}(\cdot)$ performs well (i.e., has a small MSE) for a specific parameter value $\boldsymbol{x}_1$, it may still exhibit poor performance (i.e., a large MSE) for a different parameter value $\boldsymbol{x}_2$. Ideally, an optimal estimator should have minimum MSE for all parameter values *simultaneously*. However, such an

optimality criterion is unobtainable since the minimum MSE achievable at a specific parameter value $\boldsymbol{x}_1$ is zero; it is achieved by the trivial estimator $\hat{\boldsymbol{x}}(\boldsymbol{y}) \equiv \boldsymbol{x}_1$ which is constant and completely ignores the observation $\boldsymbol{y}$. Therefore, if there were a *uniformly minimum MSE* estimator, it would have to achieve zero MSE for all parameter values, which is obviously impossible. Thus, requiring the estimator to minimize the MSE at all parameter values simultaneously makes no sense.

One useful optimality criterion is the minimax approach, which considers the worst-case MSE

$$\sup_{\boldsymbol{x} \in \mathcal{X}_S} \varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}})$$

of an estimator $\hat{\boldsymbol{x}}(\cdot)$. An optimal estimator in the minimax sense minimizes the worst-case MSE, i.e., is a solution of the optimization problem

$$\inf_{\hat{\boldsymbol{x}}(\cdot)} \sup_{\boldsymbol{x} \in \mathcal{X}_S} \varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}}) \,.$$

Considerable effort has been spent to identify minimax estimators for sparse models such as the SSNM in (5.2); see, e.g., [95–97]. However, these results only apply in the asymptotic regime, i.e., when $N, S \to \infty$. By contrast, our goal is to analyze estimator performance for finite problem dimensions. There are no known closed-form expressions of the minimax risk or of minimax-optimal estimators for the SSNM in this case.

In this work, rather than pursuing the minimax criterion, we consider *unbiased* estimators $\hat{\boldsymbol{x}}(\cdot)$ for the SSNM. An unbiased estimator is one for which the bias $\boldsymbol{b}(\boldsymbol{x}; \hat{\boldsymbol{x}})$ is zero for all S-sparse parameter vectors i.e.,

$$\mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}(\boldsymbol{y})\} = \boldsymbol{x} \qquad \text{for all } \boldsymbol{x} \in \mathcal{X}_S \,. \tag{5.8}$$

Let $\mathcal{U}$ denote the set of all unbiased estimators $\hat{\boldsymbol{x}}(\cdot)$ for the SSNM. Constraining an estimator to be unbiased excludes such trivial estimators as $\hat{\boldsymbol{x}}(\boldsymbol{y}) \equiv \boldsymbol{x}_1$ where $\boldsymbol{x}_1 \in \mathcal{X}_S$ is some fixed S-sparse parameter vector.

5.2.3 Unbiased Estimation for the SSNM

We now study the set $\mathcal{U}$ of unbiased estimators for the SSNM in more detail. In particular, we will show that with the exception of the case $S = N$, this set is uncountably large, i.e., there are infinitely many unbiased estimators. We will also show that there exists no uniformly minimum variance unbiased estimator unless $S = N$. In what follows, we will say that an

estimator $\hat{\boldsymbol{x}}$ has a bounded MSE if $\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \leq C$ for all $\boldsymbol{x} \in \mathbb{R}^N$, where C is a constant which may depend on N, S, and σ^2.

Theorem 5.1. *Consider the SSNM* (5.2) *with $S = N$, i.e., without a sparsity constraint, in which case $\mathcal{X}_S = \mathbb{R}^N$. Then, there exists exactly one unbiased estimator having bounded MSE (up to deviations having zero measure). This estimator is given by $\hat{\boldsymbol{x}}(\boldsymbol{y}) = \boldsymbol{y}$, which equals the LS/ML estimator in* (5.5)–(5.7).

The proof of this result can be found in Appendix 5.A. By contrast with Theorem 5.1, when sparsity constraints are imposed there exists a large family of unbiased estimators, as we now show.

Theorem 5.2. *For $1 \leq S < N$, there are uncountably infinitely many unbiased estimators for the SSNM.*

Proof. Consider the class of estimators defined by

$$\hat{x}(\boldsymbol{y}) = \boldsymbol{y} + a y_1 \left[\prod_{k=2}^{S+1} h^{(c,d)}(y_k) \right] \begin{pmatrix} 1 & 0 & \cdots & 0 \end{pmatrix}^T, \quad a \in \mathbb{R},\ c, d \in \mathbb{R}_{++} \tag{5.9}$$

where

$$h^{(c,d)}(y) \triangleq \begin{cases} \operatorname{sgn}(y), & |y| \in [c, c+d] \\ 0, & \text{else.} \end{cases} \tag{5.10}$$

A straightforward calculation shows that each estimator of this uncountably infinite class is an unbiased estimator for the SSNM. □ □

This (constructive) proof points at a noteworthy fact. Consider a particular parameter value $\boldsymbol{x}$. By an appropriate choice of the parameters a, c, d in (5.9), one can reduce the magnitude of the estimate $\hat{\boldsymbol{x}}(\boldsymbol{y})$ for sets of realizations $\boldsymbol{y}$ with high probability, i.e., for which $f(\boldsymbol{y};\boldsymbol{x})$ is large. This results in a reduced mean power and (since the estimator is unbiased) in a reduced variance and MSE at the specific parameter value $\boldsymbol{x}$. One can thus construct an unbiased estimator that performs better than the (biased) LS/ML estimator at the given $\boldsymbol{x}$.

In view of Theorems 5.1 and 5.2, we will only consider the case $S < N$ in the following. Since in this case there are infinitely many unbiased estimators, we would like to find an unbiased estimator having minimum variance (and, thus, minimum MSE) among all unbiased estimators. If there exists an unbiased estimator $\hat{\boldsymbol{x}}(\cdot) \in \mathcal{U}$ which minimizes the variance *simultaneously* for all S-sparse parameter vectors $\boldsymbol{x} \in \mathcal{X}_S$, then this estimator is called a *uniformly minimum variance unbiased* (UMVU) estimator [6]. In other words, a UMVU estimator for the SSNM solves the

optimization problem

$$\arg\min_{\hat{\boldsymbol{x}}(\cdot)\in\mathcal{U}} V(\boldsymbol{x};\hat{\boldsymbol{x}}) \tag{5.11}$$

simultaneously for all $\boldsymbol{x}\in\mathcal{X}_S$. In the nonsparse case $S=N$, it is well known that the LS estimator is the UMVU estimator [90]; however, in light of Theorem 5.1, this is not a very strong result, since $\hat{\boldsymbol{x}}_{\mathrm{LS}}$ is the *only* unbiased estimator in that case. On the other hand, for the sparse case $S<N$, the following negative result is shown in Appendix 5.B.

Theorem 5.3. *For the SSNM with $S<N$, there exists no UMVU estimator, i.e., there is no unbiased estimator $\hat{\boldsymbol{x}}\in\mathcal{U}$ that minimizes $V(\boldsymbol{x};\hat{\boldsymbol{x}})$ simultaneously for all parameter vectors $\boldsymbol{x}\in\mathcal{X}_S$.*

Despite the fact that a UMVU estimator does not exist for the SSNM, one can still attempt to solve the optimization problem (5.11) separately for each value of $\boldsymbol{x}\in\mathcal{X}_S$. An unbiased estimator which solves (5.11) for a specific value of $\boldsymbol{x}$ is said to be *locally minimum variance unbiased* (LMVU) [6]. The MSE of this estimator at $\boldsymbol{x}$ cannot be improved upon by any unbiased estimator. When viewed as a function of $\boldsymbol{x}$, this minimum MSE is known as the *Barankin bound* (BB) [52,53]. Thus, the BB characterizes the minimum MSE achievable by any unbiased estimator for each value of $\boldsymbol{x}\in\mathcal{X}_S$; it is the highest and tightest lower bound on the MSE of unbiased estimators. As such, the BB serves as a measure of the difficulty of estimating $\boldsymbol{x}$.

Computing the BB is equivalent to calculating $\min_{\hat{\boldsymbol{x}}(\cdot)\in\mathcal{U}} V(\boldsymbol{x};\hat{\boldsymbol{x}})$ for each parameter vector $\boldsymbol{x}\in\mathcal{X}_S$ separately. Unfortunately, there does not appear to be a simple closed-form expression of the BB, and the numerical computation of the BB seems to be difficult as well. Therefore, in the remainder of this paper, we will provide lower and upper bounds on the BB. When these bounds are close to one another, they provide an accurate characterization of the BB.

5.3 Lower Bounds on the Minimum MSE

In this section, we will develop a lower bound on the BB (which is thus a lower bound on the MSE of any unbiased estimator) by calculating a limiting case of the Hammersley–Chapman–Robbins bound [52] for the SSNM.

5.3.1 Review of the CRB

A variety of techniques exist for developing lower bounds on the MSE of unbiased estimators. The simplest is the CRB [16, 45, 48], which was previously derived for a more general sparse

estimation setting in [70,79]. In the current setting, i.e., for the SSNM (5.1), the CRB is given by

$$\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \geq \begin{cases} S\sigma^2, & \|\boldsymbol{x}\|_0 = S \\ N\sigma^2, & \|\boldsymbol{x}\|_0 < S \end{cases} \tag{5.12}$$

where $\hat{\boldsymbol{x}} \in \mathcal{U}$, i.e., $\hat{\boldsymbol{x}}(\cdot)$ is any unbiased estimator for the SSNM.

In the case of parameter values $\boldsymbol{x} \in \mathcal{X}_S$ with non-maximal support, i.e., $\|\boldsymbol{x}\|_0 < S$, the CRB is $N\sigma^2$. This is the MSE of the trivial unbiased estimator $\hat{\boldsymbol{x}}(\boldsymbol{y}) = \boldsymbol{y}$. Since the CRB is thus achieved by an unbiased estimator, we conclude that the CRB is a *maximally tight* lower bound for $\|\boldsymbol{x}\|_0 < S$; no other lower bound can be tighter (higher). We also conclude that for $\|\boldsymbol{x}\|_0 < S$, the trivial estimator $\hat{\boldsymbol{x}}(\boldsymbol{y}) = \boldsymbol{y}$ is the LMVU estimator; no other unbiased estimator can have a smaller MSE.

For parameter values $\boldsymbol{x} \in \mathcal{X}_S$ with maximal support, i.e., $\|\boldsymbol{x}\|_0 = S$, we will see that the CRB is not maximally tight, and the trivial estimator $\hat{\boldsymbol{x}}(\boldsymbol{y}) = \boldsymbol{y}$ is not the LMVU estimator. Indeed, one problem with the CRB in (5.12) is that it is discontinuous in the transition between $\|\boldsymbol{x}\|_0 = S$ and $\|\boldsymbol{x}\|_0 < S$. Since the MSE of any estimator is continuous [6], this discontinuity implies that the CRB is not the tightest lower bound obtainable for unbiased estimators. In order to obtain tighter bounds for $\|\boldsymbol{x}\|_0 = S$, it is important to realize that the CRB is a local bound, which assumes unbiasedness only in a neighborhood of $\boldsymbol{x}$. Since we are interested in estimators that are unbiased for all $\boldsymbol{x} \in \mathcal{X}_S$, which is a more restrictive constraint than local unbiasedness, tighter (i.e., higher) lower bounds can be expected for unbiased estimators in the case $\|\boldsymbol{x}\|_0 = S$.

5.3.2 Hammersley–Chapman–Robbins Bound

An alternative lower bound for unbiased estimators is the Hammersley–Chapman–Robbins bound (HCRB) [50–52], which can be stated, in our context, as follows.

Proposition 5.4. *Given a parameter value $\boldsymbol{x} \in \mathcal{X}_S$, consider a set of p "test points" $\{\boldsymbol{v}_i\}_{i=1}^p$ such that $\boldsymbol{x} + \boldsymbol{v}_i \in \mathcal{X}_S$ for all $i = 1, \ldots, p$. Then, the covariance of any unbiased estimator $\hat{\boldsymbol{x}}(\cdot)$, $C(\boldsymbol{x};\hat{\boldsymbol{x}}) \triangleq \mathsf{E}_{\boldsymbol{x}}\{[\hat{\boldsymbol{x}}(\boldsymbol{y}) - \mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}(\boldsymbol{y})\}][\hat{\boldsymbol{x}}(\boldsymbol{y}) - \mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}(\boldsymbol{y})\}]^T\}$, satisfies*

$$C(\boldsymbol{x};\hat{\boldsymbol{x}}) \succeq \boldsymbol{V}\boldsymbol{J}^{\dagger}\boldsymbol{V}^T \tag{5.13}$$

where

$$\boldsymbol{V} \triangleq (\boldsymbol{v}_1 \cdots \boldsymbol{v}_p) \in \mathbb{R}^{N \times p} \tag{5.14}$$

and the (i,j)th entry of the matrix $\boldsymbol{J} \in \mathbb{R}^{p\times p}$ is given by

$$(\boldsymbol{J})_{i,j} \triangleq \exp\left(\frac{\boldsymbol{v}_i^T \boldsymbol{v}_j}{\sigma^2}\right) - 1\,. \tag{5.15}$$

In particular, the MSE of $\hat{\boldsymbol{x}}(\cdot)$ satisfies

$$\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \geq \operatorname{tr}\left(\boldsymbol{V}\boldsymbol{J}^{\dagger}\boldsymbol{V}^T\right). \tag{5.16}$$

The proof of Proposition 5.4, which can be found in Appendix 5.C, involves an application of the multivariate HCRB of Gorman and Hero [52] to the SSNM setting. Note that both the number of test points p and their values $\boldsymbol{v}_i$ are arbitrary and can depend on $\boldsymbol{x}$. In general, including additional test points $\boldsymbol{v}_i$ will result in a tighter HCRB [52]. Our goal in this section is to choose test points $\boldsymbol{v}_i$ which result in a tight but analytically tractable bound.

Before attempting to derive a bound which is tighter than the CRB, we first observe that the CRB itself can be obtained as the limit of a sequence of HCRBs with appropriately chosen test points. Indeed, consider the specific test points given by[1]

$$\{t\boldsymbol{e}_i\}_{i\in\operatorname{supp}(\boldsymbol{x})}\,, \quad \|\boldsymbol{x}\|_0 = S \tag{5.17a}$$

$$\{t\boldsymbol{e}_i\}_{i\in\{1,\ldots,N\}}\,, \quad \|\boldsymbol{x}\|_0 < S \tag{5.17b}$$

where $t > 0$ is a constant and $\boldsymbol{e}_i$ represents the ith column of the $N \times N$ identity matrix. Note that $p = S$ in (5.17a) and $p = N$ in (5.17b). Each value of t yields a different set of test points and, via Proposition 5.4, a different lower bound on the MSE of unbiased estimators. We show in Appendix 5.D that the CRB in (5.12) is the limit of a sequence of such bounds as $t \to 0$, and that it is tighter than any bound that can be obtained via Proposition 5.4 using the test points (5.17) for a fixed $t > 0$.

Can a set of test points different from (5.17) yield a lower bound that is tighter than the CRB? As discussed above, this is only possible for parameter values $\boldsymbol{x}$ having maximal support, i.e., $\|\boldsymbol{x}\|_0 = S$, because for $\|\boldsymbol{x}\|_0 < S$ the CRB is already maximally tight. Therefore, let us consider a parameter $\boldsymbol{x}$ with $\|\boldsymbol{x}\|_0 = S$. Suppose one of the entries within the support, x_j for some $j \in \operatorname{supp}(\boldsymbol{x})$, has a small magnitude. Such a parameter $\boldsymbol{x}$ just barely qualifies as having maximal support, so it makes sense to adapt the optimal test points (5.17b) from the non-maximal support case. However, including a test point $t\boldsymbol{e}_i$ with $i \notin \operatorname{supp}(\boldsymbol{x})$ is not allowed, since in this case $\boldsymbol{x} + t\boldsymbol{e}_i$ is not in $\mathcal{X}_S$. Instead, one could include the test point $\boldsymbol{v}_i = t\boldsymbol{e}_i - x_j\boldsymbol{e}_j$,

[1] Note that, with a slight abuse of notation, the index i of the test points is now allowed to take on non-sequential values from the set $\{1,\ldots,N\}$.

which satisfies the requirement $\boldsymbol{x}+\boldsymbol{v}_i \in \mathcal{X}_S$ and is still close to $t\boldsymbol{e}_i$ if x_j is small. More generally, for any maximal-support parameter $\boldsymbol{x}$, we propose the set of N test points given by

$$\boldsymbol{v}_i = \begin{cases} t\boldsymbol{e}_i\,, & i \in \operatorname{supp}(\boldsymbol{x}) \\ t\boldsymbol{e}_i - \xi \boldsymbol{e}_{(S)}\,, & i \notin \operatorname{supp}(\boldsymbol{x}) \end{cases} \tag{5.18}$$

for $i = 1, \ldots, N$. Here, ξ denotes the smallest (in magnitude) of the S nonzero components of $\boldsymbol{x}$ and $\boldsymbol{e}_{(S)}$ denotes the corresponding unit vector. These test points $\boldsymbol{v}_i$ satisfy the condition $\boldsymbol{x}+\boldsymbol{v}_i \in \mathcal{X}_S$. Note that the test points in (5.17a), which yield the CRB, are a subset of the test points in (5.18). It can be shown [52] that this implies that the bound induced by (5.18) will always be at least as tight as that obtained from (5.17a). It is important to note that (5.18) uses N test points for parameter values with maximal support, just as (5.17b) does for parameter values with non-maximal support. In fact, there is a smooth transition between the optimal test points (5.17b) for non-maximal support and the proposed test points (5.18) for maximal support.

While an expression of the HCRB can be obtained by simply plugging (5.18) into (5.16), the resulting bound is extremely cumbersome and not very insightful. Instead, in analogy to the derivation of the CRB above, one can obtain a simple result by taking the limit for $t \to 0$. This leads to the following theorem, which combines the cases of maximal support ((5.16) using (5.18) for $t \to 0$) and non-maximal support ((5.16) using (5.17b) for $t \to 0$), and whose proof can be found in Appendix 5.E.

Theorem 5.5. *The MSE of any unbiased estimator $\hat{\boldsymbol{x}} \in \mathcal{U}$ for the SSNM satisfies*

$$\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \;\geq\; \mathrm{HCRB}(\boldsymbol{x}) \triangleq \begin{cases} S\sigma^2 + (N-S-1)e^{-\xi^2/\sigma^2}\sigma^2, & \|\boldsymbol{x}\|_0 = S \\ N\sigma^2, & \|\boldsymbol{x}\|_0 < S\,, \end{cases} \tag{5.19}$$

where, in the case $\|\boldsymbol{x}\|_0 = S$, ξ is the smallest (in magnitude) of the S nonzero entries of $\boldsymbol{x}$.

For simplicity, we will continue to refer to (5.19) as an HCRB, even though it was obtained as a limit of HCRBs. Note that when $\|\boldsymbol{x}\|_0 < S$, the HCRB in (5.19) is identical to the CRB in (5.12), since in that case the CRB is maximally tight and cannot be improved. The HCRB also approaches the CRB when $\|\boldsymbol{x}\|_0 = S$ and all components of $\boldsymbol{x}$ are much larger than σ: here $e^{-\xi^2/\sigma^2}$ is negligible and the respective bound in (5.19) converges to $S\sigma^2$, which is equal to the CRB in (5.12). This is due to the fact that the CRB is achieved by the ML estimator asymptotically[2] as $\xi^2/\sigma^2 \to \infty$, and is therefore also maximally tight when $\xi \gg \sigma$. Furthermore, if we

[2]This can be explained by the fact that according to (5.7), the ML estimator for the SSNM retains the S largest

define the "worst-case component SNR" (briefly denoted as SNR) as ξ^2/σ^2, then Theorem 5.5 hints that the convergence to the high-SNR limit is exponential in the SNR.

One of the motivations for improving the CRB (5.12) was that (5.12) is discontinuous in the transition between $\|\boldsymbol{x}\|_0 = S$ and $\|\boldsymbol{x}\|_0 < S$. While the HCRB (5.19) is still discontinuous in this transition, the discontinuity is much smaller than that of the CRB. Indeed, the transition from $\|\boldsymbol{x}\|_0 = S$ to $\|\boldsymbol{x}\|_0 < S$ corresponds to $\xi \to 0$, in which case the first bound in (5.19) tends to $(N-1)\sigma^2$, whereas the second bound, valid for $\|\boldsymbol{x}\|_0 < S$, is $N\sigma^2$; thus, the difference between the two bounds in (5.19) is σ^2. By contrast, the difference between the two bounds in (5.12) is $(N-S)\sigma^2$, which is typically much larger. Again, the discontinuity of (5.19) implies that (5.19) is not the tightest lower bound obtainable for unbiased estimators. In Section 5.5, we will demonstrate experimentally that this discontinuity can be eliminated altogether by using a much larger number of test points. However, in that case the resulting bound no longer has a simple closed-form expression and can only be evaluated numerically.

5.4 Upper Bound on the Minimum MSE

As pointed out in the previous section, the lower bound HCRB($\boldsymbol{x}$) on the BB is not maximally tight since it is discontinuous in the transition between parameter vectors with $\|\boldsymbol{x}\|_0 = S$ and those with $\|\boldsymbol{x}\|_0 < S$. In other words, there is a gap between the HCRB and the BB. How large is this gap? We will address this issue by deriving an *upper* bound on the BB. This will be done by finding a constrained solution of (5.11). If this upper bound is close to the lower bound HCRB($\boldsymbol{x}$), we can conclude that both bounds are fairly tight and thus provide a fairly accurate characterization of the BB. As before, we consider the nontrivial case $\|\boldsymbol{x}\|_0 = S$.

We first note (cf. (5.4)) that (5.11) is equivalent to the optimization problem

$$\operatorname*{arg\,min}_{\hat{\boldsymbol{x}}(\cdot)\in\mathcal{U}} \mathsf{E}_{\boldsymbol{x}}\big\{\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\big\} = \operatorname*{arg\,min}_{\hat{\boldsymbol{x}}(\cdot)\in\mathcal{U}} \sum_{k=1}^{N} \mathsf{E}_{\boldsymbol{x}}\big\{(\hat{x}_k(\boldsymbol{y}))^2\big\}, \tag{5.20}$$

where $\hat{x}_k$ denotes the kth entry of $\hat{\boldsymbol{x}}$. This, in turn, is equivalent to the N individual scalar optimization problems

$$\operatorname*{arg\,min}_{\hat{x}_k(\cdot)\in\mathcal{U}^k} \mathsf{E}_{\boldsymbol{x}}\big\{(\hat{x}_k(\boldsymbol{y}))^2\big\}, \qquad k = 1,\ldots,N \tag{5.21}$$

components in $\boldsymbol{y}$ and zeros out all other components. For noise variances σ^2 that are extremely small compared to the nonzero entries, i.e., for $\xi^2/\sigma^2 \to \infty$, the probability that the ML estimator selects the true components becomes very close to one. Therefore, for high ξ^2/σ^2, the ML estimator behaves like an oracle estimator which knows the support of $\boldsymbol{x}$ and whose MSE is equal to $S\sigma^2$.

where $\mathcal{U}^k$ denotes the set of unbiased estimators of the kth entry of $\boldsymbol{x}$, i.e.,

$$\mathcal{U}^k \triangleq \left\{ \hat{x}_k(\cdot) \,\middle|\, \mathsf{E}_{\boldsymbol{x}}\{\hat{x}_k(\boldsymbol{y})\} = x_k \text{ for all } \boldsymbol{x} \in \mathcal{X}_S \right\}.$$

By combining the unbiased estimators $\hat{x}_k(\cdot)$ for $k = 1, \ldots, N$ into a vector, we obtain an unbiased estimator of the parameter $\boldsymbol{x}$.

It will be convenient to write the kth scalar estimator as

$$\hat{x}_k(\boldsymbol{y}) = y_k + \hat{x}'_k(\boldsymbol{y}) \tag{5.22}$$

with $\hat{x}'_k(\boldsymbol{y}) \triangleq \hat{x}_k(\boldsymbol{y}) - y_k$. Since for any $\hat{x}_k(\cdot) \in \mathcal{U}^k$ we have $\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_k(\boldsymbol{y})\} = \mathsf{E}_{\boldsymbol{x}}\{y_k\} + \mathsf{E}_{\boldsymbol{x}}\{\hat{x}'_k(\boldsymbol{y})\} = x_k + \mathsf{E}_{\boldsymbol{x}}\{\hat{x}'_k(\boldsymbol{y})\}$, the unbiasedness condition $\hat{x}_k(\cdot) \in \mathcal{U}^k$ is equivalent to

$$\mathsf{E}_{\boldsymbol{x}}\{\hat{x}'_k(\boldsymbol{y})\} = 0 \qquad \text{for all } \boldsymbol{x} \in \mathcal{X}_S .$$

For $k \in \operatorname{supp}(\boldsymbol{x})$, the solution of the optimization problem (5.21) is stated in the following lemma, which is proved in Appendix 5.F. In what follows, it will be convenient to denote by $\hat{\boldsymbol{x}}^{(\boldsymbol{x})}(\boldsymbol{y})$ a solution of the optimization problem (5.11) for a given parameter vector $\boldsymbol{x} \in \mathcal{X}_S$. We recall that the estimator $\hat{\boldsymbol{x}}^{(\boldsymbol{x})}(\boldsymbol{y})$ is an LMVU at the parameter value $\boldsymbol{x}$, and its MSE, $\varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}}^{(\boldsymbol{x})}) = \min_{\hat{\boldsymbol{x}}(\cdot) \in \mathcal{U}} V(\boldsymbol{x}; \hat{\boldsymbol{x}})$, equals the BB at $\boldsymbol{x}$.

Lemma 5.6. *Consider a parameter vector $\boldsymbol{x} \in \mathcal{X}_S$ with maximal support, i.e., $\|\boldsymbol{x}\|_0 = S$. Then, for any $k \in \operatorname{supp}(\boldsymbol{x})$, the solution of the optimization problem* (5.21) *is given by*

$$\hat{x}_k^{(\boldsymbol{x})}(\boldsymbol{y}) = y_k , \qquad k \in \operatorname{supp}(\boldsymbol{x}) .$$

Moreover, this is the LMVU for $k \in \operatorname{supp}(\boldsymbol{x})$. The MSE of this estimator is σ^2.

Because Lemma 5.6 describes the scalar LMVU estimators for all indices $k \in \operatorname{supp}(\boldsymbol{x})$, it remains to consider the scalar problem (5.21) for $k \notin \operatorname{supp}(\boldsymbol{x})$. Since $\varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}}^{(\boldsymbol{x})})$ is the minimum of $\varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}})$ as defined by the optimization problem (5.11), we can obtain an upper bound on $\varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}}^{(\boldsymbol{x})})$ by placing further constraints on the estimator $\hat{\boldsymbol{x}}(\cdot)$ to be optimized. We will thus consider the modified optimization problem

$$\operatorname*{arg\,min}_{\hat{\boldsymbol{x}}(\cdot) \in \mathcal{U} \cap \mathcal{A}_{\boldsymbol{x}}} V(\boldsymbol{x}; \hat{\boldsymbol{x}}) \tag{5.23}$$

where the set $\mathcal{A}_{\boldsymbol{x}}$ is chosen such that a simpler problem is obtained. We will define $\mathcal{A}_{\boldsymbol{x}}$ in a componentwise fashion. More specifically, the kth component $\hat{x}_k(\boldsymbol{y})$ of $\hat{\boldsymbol{x}}(\boldsymbol{y})$, where $k \notin \operatorname{supp}(\boldsymbol{x})$, is said to belong to the set $\mathcal{A}_{\boldsymbol{x}}^k$ if the correction term $\hat{x}'_k(\boldsymbol{y}) = \hat{x}_k(\boldsymbol{y}) - y_k$ (see (5.22)) satisfies the following two properties.

- *Odd symmetry* with respect to k and all indices in $\mathrm{supp}(\boldsymbol{x})$:

$$\hat{x}'_k(\ldots,-y_l,\ldots) \;=\; -\,\hat{x}'_k(\ldots,y_l,\ldots)\,, \qquad \text{for all } \; l \in \{k\} \cup \mathrm{supp}(\boldsymbol{x})\,. \tag{5.24}$$

- *Independence* with respect to all other indices:

$$\hat{x}'_k(\ldots,y_l,\ldots) \;=\; \hat{x}'_k(\ldots,0,\ldots)\,, \qquad \text{for all } \; l \notin \{k\} \cup \mathrm{supp}(\boldsymbol{x})\,. \tag{5.25}$$

We then define $\mathcal{A}_{\boldsymbol{x}}$ as the set of estimators $\hat{\boldsymbol{x}}(\boldsymbol{y})$ such that $\hat{x}_k(\boldsymbol{y}) \in \mathcal{A}_{\boldsymbol{x}}^k$ for all $k \notin \mathrm{supp}(\boldsymbol{x})$. Note that any function $\hat{x}(\boldsymbol{y}) \in \mathcal{A}_{\boldsymbol{x}}^k$ is fully specified by its values for all arguments $\boldsymbol{y}$ such that $\mathrm{supp}(\boldsymbol{y}) = \{k\} \cup \mathrm{supp}(\boldsymbol{x})$ and all entries of $\boldsymbol{y}$ are nonnegative. The values of $\hat{x}(\boldsymbol{y})$ for all other $\boldsymbol{y}$ follow by the decomposition (5.22) and the properties (5.24) and (5.25).

To solve the modified optimization problem (5.23), we consider the equivalent scalar form

$$\underset{\hat{x}_k(\cdot)\in\mathcal{U}^k\cap\mathcal{A}_{\boldsymbol{x}}^k}{\arg\min}\; \mathsf{E}_{\boldsymbol{x}}\big\{(\hat{x}_k(\boldsymbol{y}))^2\big\}\,, \qquad k \notin \mathrm{supp}(\boldsymbol{x})\,. \tag{5.26}$$

The resulting minimum MSE is stated by the following lemma, whose proof can be found in Appendix 5.G.

Lemma 5.7. *Consider a parameter vector $\boldsymbol{x} \in \mathcal{X}_S$ with maximal support, i.e., $\|\boldsymbol{x}\|_0 = S$. Then, for any $k \notin \mathrm{supp}(\boldsymbol{x})$, the minimum MSE of any estimator $\hat{x}_k(\cdot) \in \mathcal{U}^k \cap \mathcal{A}_{\boldsymbol{x}}^k$, denoted by $\mathrm{BB}_{\mathrm{c}}^k(\boldsymbol{x})$, is given by*

$$\mathrm{BB}_{\mathrm{c}}^k(\boldsymbol{x}) = \left[1 - \prod_{l\in\mathrm{supp}(\boldsymbol{x})} g(x_l;\sigma^2)\right]\sigma^2 \tag{5.27}$$

with

$$g(x;\sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty e^{-(x^2+y^2)/(2\sigma^2)} \sinh\left(\frac{xy}{\sigma^2}\right) \tanh\left(\frac{xy}{\sigma^2}\right) dy\,. \tag{5.28}$$

Lemma 5.7 identifies the minimum MSE of any unbiased estimator of the kth component of $\boldsymbol{x}$ (where $k \notin \mathrm{supp}(\boldsymbol{x})$) that is also constrained to be an element of $\mathcal{A}_{\boldsymbol{x}}^k$. Note that $\mathrm{BB}_{\mathrm{c}}^k(\boldsymbol{x})$ does not depend on k. It provides an upper bound on the minimum MSE of any unbiased estimator of the kth component of $\boldsymbol{x}$, for any $k \notin \mathrm{supp}(\boldsymbol{x})$.

The total MSE of a vector estimator $\hat{\boldsymbol{x}}(\cdot)$ can be decomposed as $\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) = \sum_{k\in\mathrm{supp}(\boldsymbol{x})} \varepsilon(\boldsymbol{x};\hat{x}_k) + \sum_{k\notin\mathrm{supp}(\boldsymbol{x})} \varepsilon(\boldsymbol{x};\hat{x}_k)$ with the component MSE $\varepsilon(\boldsymbol{x};\hat{x}_k) \triangleq \mathsf{E}_{\boldsymbol{x}}\big\{(\hat{x}_k(\boldsymbol{y}) - x_k)^2\big\}$. Inserting the minimum component MSE for $k \in \mathrm{supp}(\boldsymbol{x})$ (which is σ^2 according to Lemma 5.6) in the first sum and the upper bound $\mathrm{BB}_{\mathrm{c}}^k(\boldsymbol{x})$ on the minimum component MSE for $k \notin \mathrm{supp}(\boldsymbol{x})$ in the second sum, we obtain the following upper bound on the minimum total MSE of any unbiased vector estimator.

Theorem 5.8. *The minimum MSE achievable by any unbiased estimator for the SSNM at a parameter vector $\boldsymbol{x} \in \mathcal{X}_S$ with $\|\boldsymbol{x}\|_0 = S$ satisfies*

$$\varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}}^{(x)}) \;\leq\; \mathrm{BB_c}(\boldsymbol{x}) \;\triangleq\; S\sigma^2 + (N-S)\,\mathrm{BB_c^k}(\boldsymbol{x}) \tag{5.29}$$

with $\mathrm{BB_c^k}(\boldsymbol{x})$ *given by* (5.27).

Depending on the parameter vector $\boldsymbol{x}$, the upper bound $\mathrm{BB_c}(\boldsymbol{x})$ varies between two extreme values. For decreasing SNR ξ^2/σ^2, it converges to the low-SNR value $N\sigma^2$ (because the factor $g(\xi, \sigma^2)$ in (5.27) vanishes for $\xi^2/\sigma^2 \to 0$). On the other hand, we will show below that for increasing SNR, $\mathrm{BB_c}(\boldsymbol{x})$ converges to its high-SNR value, which is given by $S\sigma^2$.

The lower bound $\mathrm{HCRB}(\boldsymbol{x})$ in (5.19) for the case $\|\boldsymbol{x}\|_0 = S$, i.e., $S\sigma^2 + (N - S - 1)e^{-\xi^2/\sigma^2}\sigma^2$, exhibits an exponential transition between the low-SNR and high-SNR regimes. More specifically, when considering a sequence of parameter vectors $\boldsymbol{x} \in \mathcal{X}_S$ with increasing SNR ξ^2/σ^2, the bound transitions from the low-SNR value $(N-1)\sigma^2$ (obtained for $\xi^2/\sigma^2 = 0$) to the high-SNR value $S\sigma^2$ (obtained for $\xi^2/\sigma^2 \to \infty$); this transition is exponential in the SNR. The upper bound $\mathrm{BB_c}(\boldsymbol{x})$ in (5.29) also exhibits a transition that is exponential in ξ^2/σ^2. In fact, it is shown in Appendix 5.H that

$$\mathrm{BB_c}(\boldsymbol{x}) \;\leq\; S\sigma^2 + (N-S)\,3^S e^{-\xi^2/(2\sigma^2)}\sigma^2. \tag{5.30}$$

This shows that for increasing ξ^2/σ^2, the upper bound $\mathrm{BB_c}(\boldsymbol{x})$—just like the lower bound $\mathrm{HCRB}(\boldsymbol{x})$—decays exponentially to its asymptotic value $S\sigma^2$, which is also the asymptotic value of $\mathrm{HCRB}(\boldsymbol{x})$. It follows that the BB itself also converges exponentially to $S\sigma^2$ as ξ^2/σ^2 increases. This result will be further explored in Section 5.5.3.

5.5 Numerical Results

In this section, we describe several numerical studies which explore and extend the theoretical bounds developed above. These include a numerical improvement of the bounds, a comparison with practical (biased) estimation techniques, an analysis of the performance at high SNR, and an examination of the ability to estimate the threshold region in which the transition from low to high SNR occurs.

We will first show that it is possible to obtain significantly tighter versions of the lower and upper bounds developed in Sections 5.3 and 5.4. These tightened versions can only be computed numerically and no longer have a simple form; consequently, they are less convenient

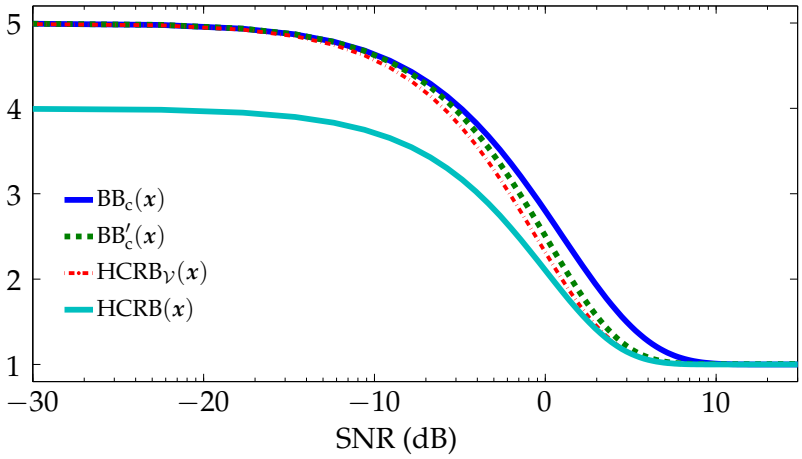

Figure 5.1: Lower bounds HCRB$(\boldsymbol{x})$, HCRB$_{\mathcal{V}}(\boldsymbol{x})$ and upper bounds BB$_c(\boldsymbol{x})$, BB$'_c(\boldsymbol{x})$ on the MSE $\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}^{(x)})$ of the LMVU estimator at $\boldsymbol{x} = c\,(1\ \ 0\ \ 0\ \ 0\ \ 0)^T$, with c varied to obtain different values of SNR$(\boldsymbol{x}) = \xi^2/\sigma^2$. The SSNM parameters are $N=5$, $S=1$, and $\sigma^2=1$.

for theoretical analyses. Nevertheless, they characterize the BB very accurately and therefore also provide an indication of the accuracy of the simpler, closed-form bounds.

5.5.1 Numerical Lower Bound

For a parameter vector $\boldsymbol{x}$ with $\|\boldsymbol{x}\|_0 = S$, let us reconsider the HCRB in (5.16). We will show that by using an increased number of appropriately chosen test points, we can obtain a lower bound that is higher (thus, tighter) than (5.19). Specifically, assume without loss of generality that $\mathrm{supp}(\boldsymbol{x}) = \{1,\dots,S\}$, and consider the set of test points

$$\mathcal{V} \triangleq \mathcal{V}_0 \cup \bigcup_{k=1}^{S} (\mathcal{V}_k \cup \mathcal{W}_k)$$

with the component sets

$$\begin{aligned}
\mathcal{V}_0 &\triangleq \bigcup_{l\in\,\mathrm{supp}(\boldsymbol{x})} \{\alpha \boldsymbol{e}_l\} \\
\mathcal{V}_k &\triangleq \bigcup_{l\in\{S+1,\dots,N\}} \{\alpha \boldsymbol{e}_l - x_k \boldsymbol{e}_k\}\,, \qquad k = 1,\dots,S \\
\mathcal{W}_k &\triangleq \bigcup_{l\in\{S+1,\dots,N\}} \{x_k \boldsymbol{e}_l - x_k \boldsymbol{e}_k\}\,, \qquad k = 1,\dots,S
\end{aligned}$$

where $\alpha = 0.02\sigma$. In Fig. 5.1, the HCRB (5.16) for the new set $\mathcal{V}$ of test points—denoted HCRB$_{\mathcal{V}}(\boldsymbol{x})$—is displayed versus the SNR and compared with HCRB$(\boldsymbol{x})$. For this figure, we chose $N = 5$, $S = 1$, $\sigma^2 = 1$, and $\boldsymbol{x} = c\,(1\ \ 0\ \ 0\ \ 0\ \ 0)^T$, where the parameter $c \in \mathbb{R}$ is varied

to obtain different SNR values.[3] As before, the SNR is defined as $\mathrm{SNR}(\boldsymbol{x}) = \xi^2/\sigma^2$, where ξ is the S-largest (in magnitude) component of $\boldsymbol{x}$ (in our example with $S = 1$, ξ is simply the single nonzero component). It can be seen from Fig. 5.1 that the numerical lower bound $\mathrm{HCRB}_{\mathcal{V}}(\boldsymbol{x})$ computed from the above test points is indeed tighter than the closed-form lower bound $\mathrm{HCRB}(\boldsymbol{x})$ in (5.19).

5.5.2 Numerical Upper Bound

It is also possible to find upper bounds on the BB that are tighter (lower) than the upper bound $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})$ in (5.29). Consider a parameter vector $\boldsymbol{x}$ with $\|\boldsymbol{x}\|_0 = S$. We recall that $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})$ was derived by constructing, for all $k \notin \mathrm{supp}(\boldsymbol{x})$, unbiased estimators $\hat{x}_k(\boldsymbol{y}) = y_k + \hat{x}'_k(\boldsymbol{y})$ with $\hat{x}'_k(\boldsymbol{y})$ constrained by (5.24) and (5.25). We will now investigate how much we can improve on $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})$ if we remove the constraint (5.24). Thus, in the optimization problem (5.23), the constraint set $\mathcal{A}_{\boldsymbol{x}}$ is hereafter considered to correspond only to the constraint (5.25).

In order to numerically solve this modified optimization problem (5.23), a discrete approximation for $\hat{x}'_k(\boldsymbol{y})$ was used. More specifically, we defined $\hat{x}'_k(\boldsymbol{y})$ to be piecewise constant in each of the components y_l with $l \in \{k\} \cup \mathrm{supp}(\boldsymbol{x})$, and constant in the remaining components y_l (the latter being required by (5.25)). We used Q piecewise constant segments for each $l \in \{k\} \cup \mathrm{supp}(\boldsymbol{x})$, with each segment of length $\Delta = 10\,\sigma/Q$. These arrays of constant segments were centered about $\boldsymbol{y} = \boldsymbol{x}$. The remaining values of $\hat{x}'_k(\boldsymbol{y})$ were set to 0. Thus, we obtained a function $\hat{x}'_k(\boldsymbol{y})$ with linear dependence on a finite number Q^{S+1} of parameters. For functions of this form, the optimization problem (5.23) becomes a finite-dimensional quadratic program with linear constraints, which can be solved efficiently [98]. The MSE of the resulting estimator, denoted by $\mathrm{BB}'_{\mathrm{c}}(\boldsymbol{x})$, is an upper bound on the BB. This bound is tighter than the closed-form upper bound $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})$ in (5.29) if Q is large enough. In Fig. 5.1, we compare $\mathrm{BB}'_{\mathrm{c}}(\boldsymbol{x})$ for $Q = 20$ with $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})$ as a function of the SNR. The improved accuracy of $\mathrm{BB}'_{\mathrm{c}}(\boldsymbol{x})$ relative to $\mathrm{BB}_{\mathrm{c}}(\boldsymbol{x})$ is evident, particularly at high SNR values. Moreover, the proximity of the numerical upper bound $\mathrm{BB}'_{\mathrm{c}}(\boldsymbol{x})$ to the numerical lower bound $\mathrm{HCRB}_{\mathcal{V}}(\boldsymbol{x})$ indicates that these two bounds achieve an accurate characterization of the BB, since the BB lies between them.

[3] The use of a low-dimensional model is mandated by the complexity of the numerical approximation to the upper bound on the BB which will be described in Section 5.5.2.

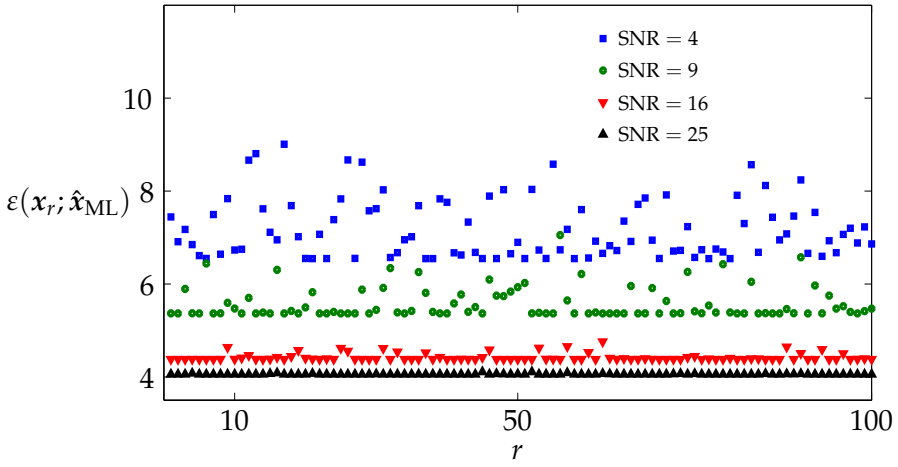

Figure 5.2: MSE $\varepsilon(\boldsymbol{x}_r;\hat{\boldsymbol{x}}_{\mathrm{ML}})$ of the ML estimator for randomly generated parameter vectors $\boldsymbol{x}_r$ at four different SNRs ξ^2/σ^2, for SSNM parameters $N=10$, $S=4$, and $\sigma^2=1$.

5.5.3 The Role of ξ

We have seen in Section 5.4 that for $\|\boldsymbol{x}\|_0 = S$, the MSE of the LMVU estimator at high SNR is given by $S\sigma^2$, and furthermore, convergence to this value is exponential in the quantity ξ^2/σ^2. A remarkable aspect of this conclusion is the fact that convergence to the high-SNR regime depends solely on ξ, the smallest nonzero component of $\boldsymbol{x}$, rather than having a more complex dependency on all the S nonzero components of $\boldsymbol{x}$. For example, one might imagine the behavior of an estimator to be rather different when all nonzero components have the same value ξ, as opposed to the situation in which one component equals ξ and the others are much larger. However, our analysis shows that when $\xi \gg \sigma$, the remaining components of $\boldsymbol{x}$ have no effect on the performance of the LMVU estimator. We will next investigate whether practical estimators also exhibit such an effect.

To answer this question, we examined the MSE of the ML estimator (5.7) for a wide range of parameter vectors $\boldsymbol{x}$ having a predetermined smallest component ξ. More specifically, for a given value of ξ, we randomly generated 100 parameter vectors $\boldsymbol{x}_r$, $r = 1,\ldots,100$, with $\boldsymbol{x}_r \in \mathcal{X}_S$ and $\|\boldsymbol{x}_r\|_0 = S$, whose minimum nonzero component was equal to ξ. The other nonzero components were generated as independent, identically distributed realizations of the random variable $x = \xi(1+3\sigma|q|)$, where $q \sim \mathcal{N}(0,1)$ is a standard Gaussian random variable and σ is the standard deviation of the noise. The MSE $\varepsilon(\boldsymbol{x}_r;\hat{\boldsymbol{x}}_{\mathrm{ML}})$ of the ML estimator is shown in Fig. 5.2 for $N=10$, $S=4$, and four different SNRs ξ^2/σ^2, with the horizontal axis representing the different choices of $\boldsymbol{x}_r$ in arbitrary order. It is seen that for large ξ, the performance of

the ML estimator, like that of the LMVU, depends almost exclusively on ξ. This suggests that the performance guarantees of Sections 5.3 and 5.4, while formally valid only for unbiased estimators, can still provide general conclusions which are relevant to biased techniques such as the ML estimator. Moreover, this result also justifies our definition of the SNR as the ratio ξ^2/σ^2, since this is the most significant factor determining estimation performance for the SSNM.

5.5.4 Threshold Region Identification

In Sections 5.3 and 5.4, we characterized the performance of unbiased estimators as a means of quantifying the difficulty of estimation for the SSNM. A common use of this analysis is in the identification of the threshold region, a range of SNR values which constitutes a transition between low-SNR and high-SNR behavior [99–101]. Specifically, in many cases the performance of estimators can be calculated analytically when the SNR is either very low or very high. It is then important to identify the threshold region which separates these two regimes. Although the analysis is based on bounds for unbiased estimators, the result is often heuristically assumed to approximate the threshold region for biased techniques as well [99, 101].

For $\|\boldsymbol{x}\|_0 = S$, the lower and upper bounds on the BB (HCRB($\boldsymbol{x}$) in (5.19), $\mathrm{BB_c}(\boldsymbol{x})$ in (5.29)) exhibit a transition between a low-SNR region, where both bounds are on the order of $N\sigma^2$, and a high-SNR region, for which both bounds converge to $S\sigma^2$. The BB therefore also displays such a transition. One can define the threshold region of the SSNM (for unbiased estimation) as the range of values of ξ^2/σ^2 in which this transition takes place. Since the BB is itself a lower bound on the performance of unbiased estimators, one would expect the transition region of actual estimators to occur at slightly higher SNR values than that of the BB.

To test this hypothesis, we compared the bounds of Sections 5.3 and 5.4 with the MSE of two well-known estimation schemes, namely, the ML estimator in (5.7) and the hard-thresholding (HT) estimator $\hat{\boldsymbol{x}}_{\mathrm{HT}}(\boldsymbol{y})$, which is given componentwise as

$$\hat{x}_{\mathrm{HT},k}(\boldsymbol{y}) = \begin{cases} y_k, & |y_k| \geq T \\ 0, & \text{else} \end{cases}$$

for a given threshold $T > 0$. In our simulations, we chose the commonly used value $T = \sigma\sqrt{2\log N}$ [2]. Note that since the ML and HT estimators are biased, their MSE is not bounded by $\mathrm{BB_c}(\boldsymbol{x})$, HCRB($\boldsymbol{x}$), and the CRB. Assuming SSNM parameters $N = 10$ and $S = 4$, we generated a number of parameter vectors $\boldsymbol{x}$ from the set $\mathcal{R} \triangleq \left\{c\,(1\ 1\ 1\ 1\ 0\ 0\ 0\ 0\ 0\ 0)^T\right\}_{c\in\mathbb{R}}$, where c was varied to obtain a range of SNR values. For these $\boldsymbol{x}$, we calculated the MSE of the two

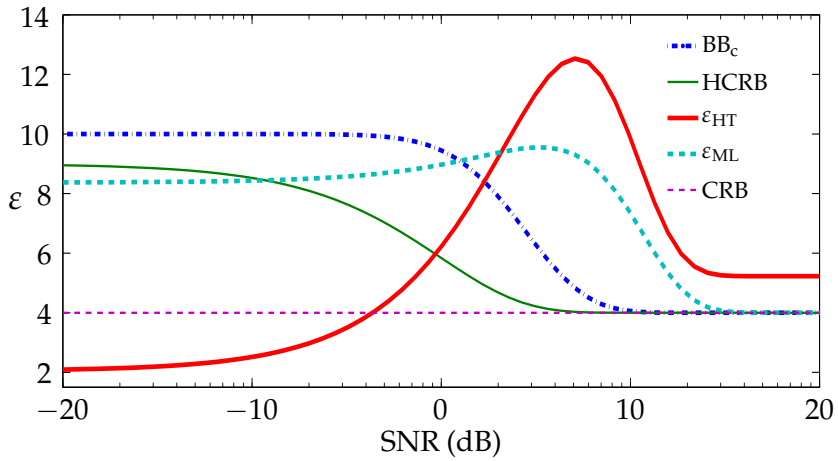

Figure 5.3: MSE of the ML and HT estimators compared with the performance bounds $\mathrm{BB_c}(\boldsymbol{x})$, $\mathrm{HCRB}(\boldsymbol{x})$, and CRB ($\equiv S\sigma^2$), as a function of the SNR ξ^2/σ^2, for SSNM parameters $N = 10$, $S = 4$, and $\sigma^2 = 1$.

estimators $\hat{\boldsymbol{x}}_{\mathrm{ML}}$ and $\hat{\boldsymbol{x}}_{\mathrm{HT}}$ by means of numerical integration (see Appendix 5.I for a discussion of the computation of $\varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}}_{\mathrm{ML}})$).

The results are displayed in Fig. 5.3 as a function of the SNR ξ^2/σ^2. Although there is some gap between the lower bound (HCRB) and the upper bound ($\mathrm{BB_c}$), a rough indication of the behavior of the BB is conveyed. As expected, the threshold region exhibited by the ML and HT estimators is somewhat higher than that predicted by the bounds. Specifically, the threshold region of the BB (as indicated by the bounds) can be seen to occur at SNR values between -5 and 5 dB, while the threshold region of the ML and HT estimators is at SNR values between 5 and 12 dB. Another effect which is visible in Fig. 5.3 is the convergence of the ML estimator to the BB at high SNR; this is a manifestation of the well-known fact that the ML estimator is asymptotically unbiased and asymptotically optimal. Finally, at low SNR, both the ML and HT estimators are better than the best unbiased approach. This is because unbiased methods generally perform poorly at low SNR, so that even the best unbiased technique is outperformed by the biased ML and HT estimators. On the other hand, for medium SNR, the MSE of the ML and HT estimators is significantly higher than the BB. Thus, there is a potential for unbiased estimators to perform better than biased estimators in the medium-SNR regime.

One may argue that considering only parameter vectors $\boldsymbol{x}$ in the set $\mathcal{R}$ is not representative, since $\mathcal{R}$ covers only a small part of the parameter space $\mathcal{X}_S$. However, the choice of $\mathcal{R}$ is conservative in that the maximum deviation between $\mathrm{HCRB}(\boldsymbol{x})$ and $\mathrm{BB_c}(\boldsymbol{x})$ is largest when the nonzero entries of $\boldsymbol{x}$ have approximately the same magnitude, which is the case for each ele-

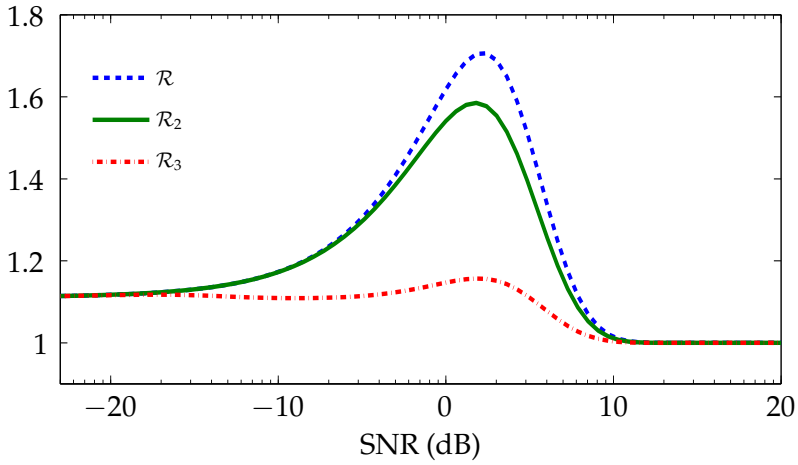

Figure 5.4: Ratio $\mathrm{BB_c}(\boldsymbol{x})/\mathrm{HCRB}(\boldsymbol{x})$ versus the SNR ξ^2/σ^2 for different sets of parameter vectors $\boldsymbol{x}$.

ment of $\mathcal{R}$. This is illustrated in Fig. 5.4, which shows the ratio between the two bounds versus the SNR ξ^2/σ^2 for three different configurations of the nonzero entries in the parameter vector. Specifically, we considered the two additional sets $\mathcal{R}_2 \triangleq \left\{c\,(10\ 1\ 1\ 1\ 0\ 0\ 0\ 0\ 0\ 0)^T\right\}_{c\in\mathbb{R}}$ and $\mathcal{R}_3 \triangleq \left\{c\,(0.1\ 1\ 1\ 1\ 0\ 0\ 0\ 0\ 0\ 0)^T\right\}_{c\in\mathbb{R}}$, in which the nonzero entries have different magnitudes. It can be seen from Fig. 5.4 that the ratio $\mathrm{BB_c}(\boldsymbol{x})/\mathrm{HCRB}(\boldsymbol{x})$ is indeed highest when $\boldsymbol{x}$ is in $\mathcal{R}$.

5.6 Conclusion

In this paper, we have studied unbiased estimation of a sparse vector in white Gaussian noise within a frequentist setting. As we have seen, without the assumption of sparsity, there exists only a single unbiased estimator. However, the addition of a sparsity assumption yields a rich family of unbiased estimators. The analysis of the performance of these estimators has been the primary goal of this paper. We first demonstrated that there exists no uniformly minimum variance unbiased estimator, i.e., no single unbiased estimator is optimal for all parameter values. Consequently, we focused on analyzing the Barankin bound (BB), i.e., the MSE of the locally minimum variance unbiased estimator, or equivalently, the smallest MSE achievable by an unbiased estimator for each value of the sparse vector.

For the sparse estimation problem considered, as for most estimation problems, the BB cannot be computed precisely. However, we demonstrated that it can be characterized quite accurately using numerical lower and upper bounds. Furthermore, we derived simple closed-

form lower and upper bounds which are somewhat looser than the numerical bounds. These closed-form bounds allow an estimation of the threshold region separating the low-SNR and high-SNR regimes, and they indicate the asymptotic behavior of the BB at high SNR. In particular, a notable conclusion is that the high-SNR behavior of the BB depends solely on the value of the smallest nonzero component of the sparse vector.

While the unbiasedness property is intuitively appealing and related to several desirable asymptotic features of an estimator [6], one can often obtain biased estimators which outperform any unbiased estimator [42, 44, 91]. Thus, it is interesting to note that some of the conclusions obtained from our analysis of unbiased estimators appear to provide insight into the behavior of standard biased estimators. In particular, we saw that the behavior of two commonly used biased estimators at high SNR corresponds to the predictions of our unbiased bounds, not only in terms of the asymptotically achievable MSE but also in certain finer details, such as the SNR range of the threshold region and the fact that the convergence to the high-SNR regime depends primarily on the value of the smallest nonzero component of the sparse vector, rather than on the entire vector. This gives additional merit to the analysis of achievable estimation performance within the unbiased setting.

5.A Proof of Theorem 5.1

We wish to show that for $S = N$, the only unbiased estimator with bounded MSE is the trivial estimator $\hat{\boldsymbol{x}}(\boldsymbol{y}) = \boldsymbol{y}$. We will first show that a bounded MSE implies that $\hat{\boldsymbol{x}}(\boldsymbol{y})$ is equivalent to a tempered distribution. This will allow us to reformulate the unbiasedness condition in the Fourier transform domain.

Using (5.3), the unbiasedness condition in (5.8) for $S = N$ reads

$$\frac{1}{(2\pi\sigma^2)^{N/2}} \int_{\mathbb{R}^N} \hat{\boldsymbol{x}}(\boldsymbol{y}) \, \exp\left(-\frac{1}{2\sigma^2}\|\boldsymbol{y}-\boldsymbol{x}\|_2^2\right) d\boldsymbol{y} \;=\; \boldsymbol{x} \qquad \text{for all } \boldsymbol{x} \in \mathbb{R}^N . \tag{5.31}$$

The integral in (5.31) is the convolution of $\hat{\boldsymbol{x}}(\boldsymbol{y})$ with $\exp\left(-\frac{1}{2\sigma^2}\|\boldsymbol{y}\|_2^2\right)$. The result of this convolution, viewed as a function of $\boldsymbol{x}$, must equal $(2\pi\sigma^2)^{N/2}\,\boldsymbol{x}$ for all parameter vectors $\boldsymbol{x}$. For absolutely integrable functions, the Fourier transform maps a convolution onto a pointwise product, and consequently it seems natural to consider the Fourier transform of condition (5.31) in order to simplify the analysis. However, typically, the estimator function $\hat{\boldsymbol{x}}(\boldsymbol{y})$ will be neither absolutely integrable nor square integrable, and thus its Fourier transform can only exist in the sense of a tempered distribution [102]. From a practical point of view, the class of tempered

distributions is large enough so that it does not exclude reasonable estimators such as the LS estimator (5.7). The following lemma states that $\hat{\boldsymbol{x}}(\boldsymbol{y})$ can be viewed as a tempered distribution if it has a bounded MSE.

Lemma 5.9. *Consider an estimator $\hat{\boldsymbol{x}}$ for the SSNM (5.2) with $S = N$. If $\hat{\boldsymbol{x}}$ has a bounded MSE, i.e., $\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \leq C$ for all $\boldsymbol{x} \in \mathbb{R}^N$ (where C is a constant which may depend on N, S, and σ^2), then $\hat{\boldsymbol{x}}$ is equivalent to a tempered distribution.*

Proof. The proof of Lemma 5.9 is based on the following result which gives a sufficient condition for a function $\hat{\boldsymbol{x}}(\boldsymbol{y})$ to be (equivalent to) a tempered distribution.

Proposition 5.10 ([102]). *If there exist constants $B, n, R_0 \in \mathbb{R}_+$ such that*

$$\int_{\|\boldsymbol{y}\|_2 \leq R} \|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2 \, d\boldsymbol{y} \;\leq\; BR^n \qquad \text{for all } \; R \geq R_0 \tag{5.32}$$

then $\hat{\boldsymbol{x}}(\boldsymbol{y})$ is equivalent to a tempered distribution.

Let $\hat{\boldsymbol{x}}(\boldsymbol{y})$ be an estimator function with bounded MSE, i.e., there exists a constant C such that

$$\mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y}) - \boldsymbol{x}\|_2^2\} \;\leq\; C \qquad \text{for all } \; \boldsymbol{x} \in \mathcal{X}_S \, . \tag{5.33}$$

Defining the usual norm $\|\cdot\|_{\mathrm{RV}}$ on the space of of random vectors by $\|\boldsymbol{y}\|_{\mathrm{RV}} \triangleq \sqrt{\mathsf{E}_{\boldsymbol{x}}\{\|\boldsymbol{y}\|_2^2\}}$, we can use the (reverse) triangle inequality $\|\hat{\boldsymbol{x}}(\boldsymbol{y}) - \boldsymbol{x}\|_{\mathrm{RV}} \geq \|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_{\mathrm{RV}} - \|\boldsymbol{x}\|_{\mathrm{RV}}$ to obtain

$$\sqrt{\mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y}) - \boldsymbol{x}\|_2^2\}} \;\geq\; \sqrt{\mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\}} - \sqrt{\mathsf{E}_{\boldsymbol{x}}\{\|\boldsymbol{x}\|_2^2\}} = \sqrt{\mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\}} - \|\boldsymbol{x}\|_2 \, .$$

From this, it follows that

$$\sqrt{\mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\}} \;\leq\; \sqrt{\mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y}) - \boldsymbol{x}\|_2^2\}} + \|\boldsymbol{x}\|_2 \;\leq\; \sqrt{C} + \|\boldsymbol{x}\|_2 \qquad \text{for all } \; \boldsymbol{x} \in \mathcal{X}_S \, ,$$

where (5.33) has been used. Squaring both sides and using the inequality $(x+y)^2 \leq 2(x^2+y^2)$, we obtain

$$\mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\} \;\leq\; (\sqrt{C} + \|\boldsymbol{x}\|_2)^2 \;\leq\; 2\,(C + \|\boldsymbol{x}\|_2^2) \qquad \text{for all } \; \boldsymbol{x} \in \mathcal{X}_S$$

or equivalently

$$\frac{1}{(2\pi\sigma^2)^{N/2}} \int_{\mathbb{R}^N} \|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2 \, e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)} d\boldsymbol{y} \;\leq\; 2(C + \|\boldsymbol{x}\|_2^2) \qquad \text{for all } \; \boldsymbol{x} \in \mathcal{X}_S. \tag{5.34}$$

We will now show that (5.32) holds for $R_0 = 1$, i.e., $R \geq 1$. We define the N-dimensional grid

$$\mathcal{G} \;\triangleq\; \{-m\Delta, -(m-1)\Delta, \ldots, -\Delta, 0, \Delta, \ldots, m\Delta\}^N$$

where $0 < \Delta \leq R$ (hence, $R/\Delta \geq 1$) and $m = \lfloor R/\Delta \rfloor \leq R/\Delta$. The number of grid points in any single dimension satisfies

$$2m+1 \;\leq\; \frac{2R}{\Delta} + 1 \tag{5.35}$$

so that

$$|\mathcal{G}| = (2m+1)^N \leq \left(\frac{2R}{\Delta}+1\right)^N. \tag{5.36}$$

We thus have

$$\begin{aligned}\sum_{\boldsymbol{x}\in\mathcal{G}} \|\boldsymbol{x}\|_2^2 = \sum_{\boldsymbol{x}\in\mathcal{G}}\sum_{k=1}^{N} x_k^2 = \sum_{k=1}^{N}\sum_{\boldsymbol{x}\in\mathcal{G}} x_k^2 = \sum_{k=1}^{N}\left[(2m+1)^{N-1}\sum_{l=-m}^{m}(l\Delta)^2\right] = N(2m+1)^{N-1}\sum_{l=-m}^{m}(l\Delta)^2 \\ \leq\; N(2m+1)^{N-1}\Delta^2\int_{x=-R/\Delta}^{R/\Delta} x^2 dx \;\leq\; N\left(\frac{2R}{\Delta}+1\right)^{N-1}\frac{2}{3}\frac{R^3}{\Delta}\end{aligned} \tag{5.37}$$

where (5.35) was used in the last step. Furthermore, for $c \triangleq \frac{1}{(2\pi\sigma^2)^{N/2}} e^{-N\Delta^2/(2\sigma^2)}$, we have

$$\frac{1}{c}\,\frac{1}{(2\pi\sigma^2)^{N/2}}\sum_{\boldsymbol{x}\in\mathcal{G}} e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)} \geq 1\,, \qquad \text{for all } \boldsymbol{y} \text{ with } \|\boldsymbol{y}\|_2 \leq R \tag{5.38}$$

In order to verify this inequality, consider an arbitrary $\boldsymbol{y}\in\mathbb{R}^N$ with $\|\boldsymbol{y}\|_2 \leq R$. Since $0 < \Delta \leq R$, and since $\|\boldsymbol{y}\|_2 \leq R$ implies that no component y_k of $\boldsymbol{y}$ can be larger than R, there always exists a grid point $\tilde{\boldsymbol{x}} \in \mathcal{G}$ (dependent on $\boldsymbol{y}$) such that $|y_k - \tilde{x}_k| \leq \Delta$ for all $k \in \{1,\ldots,N\}$. It follows that $\|\boldsymbol{y}-\tilde{\boldsymbol{x}}\|_2^2 \leq N\Delta^2$ and, in turn,

$$e^{-N\Delta^2/(2\sigma^2)} \;\leq\; e^{-\|\boldsymbol{y}-\tilde{\boldsymbol{x}}\|_2^2/(2\sigma^2)} \leq \sum_{\boldsymbol{x}\in\mathcal{G}} e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)}\,, \qquad \|\boldsymbol{y}\|_2 \leq R$$

which is equivalent to (5.38).

Successively using (5.38), (5.34), (5.36), (5.37), and $1 \leq 2R/\Delta$, we obtain the following sequence of inequalities:

$$\begin{aligned}\int_{\|\boldsymbol{y}\|_2\leq R} \|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\, d\boldsymbol{y} \;&\leq\; \int_{\|\boldsymbol{y}\|_2\leq R} \|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2 \left[\frac{1}{c}\,\frac{1}{(2\pi\sigma^2)^{N/2}}\sum_{\boldsymbol{x}\in\mathcal{G}} e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)}\right] d\boldsymbol{y} \\ &\leq\; \frac{1}{c}\sum_{\boldsymbol{x}\in\mathcal{G}}\frac{1}{(2\pi\sigma^2)^{N/2}}\int_{\mathbb{R}^N}\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\, e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)} d\boldsymbol{y} \\ &\leq\; \frac{1}{c}\sum_{\boldsymbol{x}\in\mathcal{G}} 2(C+\|\boldsymbol{x}\|_2^2) \\ &\leq\; \frac{2}{c}\left[\left(\frac{2R}{\Delta}+1\right)^N C + N\left(\frac{2R}{\Delta}+1\right)^{N-1}\frac{2}{3}\frac{R^3}{\Delta}\right] \\ &\leq\; \frac{2}{c}\left[\left(\frac{4R}{\Delta}\right)^N C + N\left(\frac{4R}{\Delta}\right)^{N-1}\frac{2}{3}\frac{R^3}{\Delta}\right].\end{aligned} \tag{5.39}$$

It then follows from (5.39) that for $R \geq 1$

$$\begin{aligned}\int_{\|\boldsymbol{y}\|_2 \leq R} \|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2 \, d\boldsymbol{y} \;&\leq\; \frac{2}{c}\left[\left(\frac{4}{\Delta}\right)^N R^{N+2} C + N\left(\frac{4}{\Delta}\right)^{N-1} \frac{2}{3}\frac{R^{N+2}}{\Delta}\right] \\ &\leq\; \frac{2}{c}\frac{R^{N+2}}{\Delta^N}\left(4^N C + N 4^N \frac{2}{3}\right) \\ &=\; \frac{2^{2N+1}}{c\,\Delta^N}\left(C + \frac{2N}{3}\right) R^{N+2}.\end{aligned}$$

Thus, we have established that under the conditions of Lemma 5.9 (bounded MSE), the bound (5.32) holds with $R_0 = 1$, $B = \frac{2^{2N+1}}{c\,\Delta^N}(C + 2N/3)$, and $n = N + 2$. Therefore, it follows from Proposition 5.10 that an estimator with bounded MSE is equivalent to a tempered distribution. This concludes the proof of Lemma 5.9. □ □

We now continue our proof of Theorem 5.1. Any estimator $\hat{\boldsymbol{x}}(\boldsymbol{y})$ for the SSNM (5.2) can be written as

$$\hat{\boldsymbol{x}}(\boldsymbol{y}) = \boldsymbol{y} + \hat{\boldsymbol{x}}'(\boldsymbol{y}) \tag{5.40}$$

with the correction term $\hat{\boldsymbol{x}}'(\boldsymbol{y}) \triangleq \hat{\boldsymbol{x}}(\boldsymbol{y}) - \boldsymbol{y}$. Because $\mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}(\boldsymbol{y})\} = \mathsf{E}_{\boldsymbol{x}}\{\boldsymbol{y}\} + \mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}'(\boldsymbol{y})\} = \boldsymbol{x} + \mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}'(\boldsymbol{y})\}$, $\hat{\boldsymbol{x}}(\boldsymbol{y})$ is unbiased if and only if

$$\boldsymbol{b}(\boldsymbol{x};\hat{\boldsymbol{x}}) \;=\; \mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}'(\boldsymbol{y})\} \;\equiv\; \frac{1}{(2\pi\sigma^2)^{N/2}} \int_{\mathbb{R}^N} \hat{\boldsymbol{x}}'(\boldsymbol{y})\, e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)} d\boldsymbol{y} \;=\; \boldsymbol{0} \qquad \text{for all } \boldsymbol{x} \in \mathcal{X}_S\,. \tag{5.41}$$

Remember that we assume that $\hat{\boldsymbol{x}}$ has a bounded MSE, so that according to our above proof of Lemma 5.9, the estimator function $\hat{\boldsymbol{x}}(\boldsymbol{y})$ satisfies condition (5.32) with $n = N + 2$, i.e.,

$$\int_{\|\boldsymbol{y}\|_2 \leq R} \|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2 \, d\boldsymbol{y} \;\leq\; B R^{N+2} \qquad \text{for all } \; R \geq 1 \tag{5.42}$$

with B as given at the end of the proof of Lemma 5.9. We will also need the following bound, in which $\mathcal{R} \triangleq [-R, R]^N$:

$$\int_{\|\boldsymbol{y}\|_2 \leq R} \|\boldsymbol{y}\|_2^2 \, d\boldsymbol{y} \;\leq \int_{\mathcal{R}} \|\boldsymbol{y}\|_2^2 \, d\boldsymbol{y} \;= \sum_{k=1}^{N} \int_{\mathcal{R}} y_k^2 \, d\boldsymbol{y} \;= \sum_{k=1}^{N} (2R)^{N-1} \frac{2}{3} R^3 = \frac{N}{3} 2^N R^{N+2}. \tag{5.43}$$

We then have for the correction term $\hat{\boldsymbol{x}}'(\boldsymbol{y})$, for all $R \geq 1$,

$$\begin{aligned}
\int_{\|\boldsymbol{y}\|_2 \leq R} \|\hat{\boldsymbol{x}}'(\boldsymbol{y})\|_2^2 \, d\boldsymbol{y} &= \int_{\|\boldsymbol{y}\|_2 \leq R} \|\hat{\boldsymbol{x}}(\boldsymbol{y}) - \boldsymbol{y}\|_2^2 \, d\boldsymbol{y} \\
&\leq \int_{\|\boldsymbol{y}\|_2 \leq R} 2\big(\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2 + \|\boldsymbol{y}\|_2^2\big) d\boldsymbol{y} \\
&= 2\left(\int_{\|\boldsymbol{y}\|_2 \leq R} \|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2 \, d\boldsymbol{y} + \int_{\|\boldsymbol{y}\|_2 \leq R} \|\boldsymbol{y}\|_2^2 \, d\boldsymbol{y} \right) \\
&\leq 2\left(BR^{N+2} + \frac{N}{3} 2^N R^{N+2} \right) \\
&= \left(2B + \frac{N}{3} 2^{N+1} \right) R^{N+2}
\end{aligned}$$

where (5.42) and (5.43) have been used. Therefore, the correction term $\hat{\boldsymbol{x}}'(\boldsymbol{y})$ also satisfies (5.32) and thus, according to Proposition 5.10, it is equivalent to a tempered distribution.

The bias function $\boldsymbol{b}(\boldsymbol{x}, \hat{\boldsymbol{x}})$ in (5.41) is the convolution of $\hat{\boldsymbol{x}}'(\boldsymbol{y})$ with the Gaussian function $(2\pi\sigma^2)^{-N/2} e^{-\|\boldsymbol{y}\|_2^2/(2\sigma^2)}$. Because $S = N$, we have $\mathcal{X}_S = \mathbb{R}^N$, and thus (5.41) holds for all $\boldsymbol{x} \in \mathbb{R}^N$. Since $\hat{\boldsymbol{x}}'(\boldsymbol{y})$ is a tempered distribution and the Gaussian function is in the Schwartz class, it follows that the Fourier transform of the convolution product (5.41) is a smooth function which can be calculated as the pointwise product $\bar{\boldsymbol{x}}'(\bar{\boldsymbol{y}}) \, e^{-\|\bar{\boldsymbol{y}}\|_2^2/(2\sigma^2)}$, where $\bar{\boldsymbol{x}}'(\bar{\boldsymbol{y}})$ denotes the Fourier transform of $\hat{\boldsymbol{x}}'(\boldsymbol{y})$ [102]. Therefore, (5.41) is equivalent to $\bar{\boldsymbol{x}}'(\bar{\boldsymbol{y}}) \, e^{-\|\bar{\boldsymbol{y}}\|_2^2/(2\sigma^2)} = \boldsymbol{0}$ for all $\bar{\boldsymbol{y}} \in \mathbb{R}^N$. This can only be satisfied if $\bar{\boldsymbol{x}}'(\bar{\boldsymbol{y}}) \equiv \boldsymbol{0}$, which in turn implies that $\hat{\boldsymbol{x}}'(\boldsymbol{y}) \equiv \boldsymbol{0}$ (up to deviations of zero measure) and further, by (5.40), that $\hat{\boldsymbol{x}}(\boldsymbol{y}) = \boldsymbol{y}$. Recalling that $\mathcal{X}_S = \mathbb{R}^N$, it is clear from (5.5) that $\hat{\boldsymbol{x}}(\boldsymbol{y}) = \boldsymbol{y}$ is the LS estimator. Thus, we have shown that $\hat{\boldsymbol{x}}_{\mathrm{LS}}(\boldsymbol{y}) = \boldsymbol{y}$ is the unique unbiased estimator for the SSNM with $S = N$.

5.B Proof of Theorem 5.3

We must show that there exists no UMVU estimator for the SSNM with $S < N$. The outline of our proof is as follows. We first demonstrate that the unique solution of the optimization problem (5.11) at the parameter value $\boldsymbol{x} = \boldsymbol{0}$, i.e., $\arg\min_{\hat{\boldsymbol{x}}(\cdot) \in \mathcal{U}} V(\boldsymbol{0}; \hat{\boldsymbol{x}})$, is the estimator $\hat{\boldsymbol{x}}^{(\boldsymbol{0})}(\boldsymbol{y}) = \boldsymbol{y}$. We then show that there exist unbiased estimators which have lower variance than $\hat{\boldsymbol{x}}^{(\boldsymbol{0})}$ at other points $\boldsymbol{x}$. This implies that neither $\hat{\boldsymbol{x}}^{(\boldsymbol{0})}$ nor any other estimator uniformly minimizes the variance for all $\boldsymbol{x}$ among all unbiased estimators.

The estimator $\hat{\boldsymbol{x}}^{(\boldsymbol{0})}(\boldsymbol{y}) = \boldsymbol{y}$ is a solution of (5.11) when $\boldsymbol{x} = \boldsymbol{0}$ because the minimum variance at $\boldsymbol{x} = \boldsymbol{0}$ of any unbiased estimator is bounded below by $N\sigma^2$ and $\hat{\boldsymbol{x}}^{(\boldsymbol{0})}(\boldsymbol{y}) = \boldsymbol{y}$ achieves this

lower bound [70]. To show that $\hat{\boldsymbol{x}}^{(\mathbf{0})}$ is the unique solution of (5.11) for $\boldsymbol{x} = \mathbf{0}$, suppose by contradiction that there exists a second unbiased estimator $\hat{\boldsymbol{x}}_a$ different from $\hat{\boldsymbol{x}}^{(\mathbf{0})}$, also having variance $N\sigma^2$ at $\boldsymbol{x} = \mathbf{0}$. Consider the estimator $\hat{\boldsymbol{x}}_{\text{new}} \triangleq (\hat{\boldsymbol{x}}^{(\mathbf{0})} + \hat{\boldsymbol{x}}_a)/2$. Since $\hat{\boldsymbol{x}}^{(\mathbf{0})}$ and $\hat{\boldsymbol{x}}_a$ are unbiased, $\hat{\boldsymbol{x}}_{\text{new}}$ is unbiased as well. Thus, its variance is (see (5.4)) $V(\boldsymbol{x}; \hat{\boldsymbol{x}}_{\text{new}}) = P(\boldsymbol{x}; \hat{\boldsymbol{x}}_{\text{new}}) - \|\boldsymbol{x}\|_2^2$. In particular, we obtain for $\boldsymbol{x} = \mathbf{0}$

$$\begin{aligned} V(\mathbf{0}; \hat{\boldsymbol{x}}_{\text{new}}) = P(\mathbf{0}; \hat{\boldsymbol{x}}_{\text{new}}) &= \mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\left\{ \left\| \frac{1}{2}(\hat{\boldsymbol{x}}^{(\mathbf{0})} + \hat{\boldsymbol{x}}_a) \right\|_2^2 \right\} \\ &= \frac{1}{4}\left[\mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\{\|\hat{\boldsymbol{x}}^{(\mathbf{0})}\|_2^2\} + \mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\{\|\hat{\boldsymbol{x}}_a\|_2^2\} + 2\,\mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\{(\hat{\boldsymbol{x}}^{(\mathbf{0})})^T \hat{\boldsymbol{x}}_a\} \right] \\ &\overset{(*)}{<} \frac{1}{4}\left[\mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\{\|\hat{\boldsymbol{x}}^{(\mathbf{0})}\|_2^2\} + \mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\{\|\hat{\boldsymbol{x}}_a\|_2^2\} + 2\sqrt{\mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\{\|\hat{\boldsymbol{x}}^{(\mathbf{0})}\|_2^2\}\mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\{\|\hat{\boldsymbol{x}}_a\|_2^2\}} \right] \\ &= \frac{1}{4} \cdot 4N\sigma^2 = N\sigma^2 \end{aligned}$$

where the strict inequality $(*)$ follows from the Cauchy-Schwarz inequality applied to the inner product $\mathsf{E}_{\boldsymbol{x}=\mathbf{0}}\{(\hat{\boldsymbol{x}}^{(\mathbf{0})})^T \hat{\boldsymbol{x}}_a\}$, combined with the fact that $\hat{\boldsymbol{x}}^{(\mathbf{0})}$ and $\hat{\boldsymbol{x}}_a$ are not linearly dependent (indeed, $\hat{\boldsymbol{x}}_a \neq c\hat{\boldsymbol{x}}^{(\mathbf{0})}$ since $\hat{\boldsymbol{x}}^{(\mathbf{0})}$ and $\hat{\boldsymbol{x}}_a$ were assumed to be different unbiased estimators). This inequality means that the variance of $\hat{\boldsymbol{x}}_{\text{new}}$ at $\boldsymbol{x} = \mathbf{0}$ is lower than $N\sigma^2$. But this is impossible, as $N\sigma^2$ is the minimum variance at $\boldsymbol{x} = \mathbf{0}$ achieved by any unbiased estimator. Thus, we have shown that $\hat{\boldsymbol{x}}^{(\mathbf{0})}$ is the unique solution of (5.11) for $\boldsymbol{x} = \mathbf{0}$.

Next, still for $S < N$, we consider the specific parameter value $\boldsymbol{x}' \in \mathcal{X}_S$ whose components are given by

$$x'_k = \begin{cases} 1, & k = 2, \ldots, S+1, \\ 0, & \text{else.} \end{cases}$$

The estimator $\hat{\boldsymbol{x}}^{(\mathbf{0})}$ has variance $V(\boldsymbol{x}'; \hat{\boldsymbol{x}}^{(\mathbf{0})}) = N\sigma^2$ at $\boldsymbol{x}'$ (and at all other $\boldsymbol{x} \in \mathcal{X}_S$). We will now construct an unbiased estimator $\hat{\boldsymbol{x}}_b(\boldsymbol{y})$ whose variance at $\boldsymbol{x}'$ is smaller than $N\sigma^2$. The components of this estimator are defined as

$$\hat{x}_{b,k}(\boldsymbol{y}) \triangleq \begin{cases} y_1 + A y_1 \prod_{l=2}^{S+1} h(y_l), & k = 1 \\ y_k\,, & k = 2, \ldots, N \end{cases} \tag{5.44}$$

where

$$h(y) \triangleq \begin{cases} \operatorname{sgn}(y), & |y| \in [0.4, 0.6] \\ 0, & \text{else} \end{cases}$$

and $A \in \mathbb{R}$ is a parameter to be determined shortly.[4] A direct calculation shows that $\hat{\boldsymbol{x}}_b(\boldsymbol{y})$ is

[4]The interval $[0.4, 0.6]$ in the definition of $h(y)$ is chosen rather arbitrarily. Any interval which ensures that β in (5.45) is nonzero can be used.

unbiased for all $\boldsymbol{x} \in \mathcal{X}_S$. Note that $\hat{\boldsymbol{x}}_b(\boldsymbol{y})$ is identical to $\hat{\boldsymbol{x}}^{(\mathbf{0})}(\boldsymbol{y}) = \boldsymbol{y}$ except for the first component, $\hat{x}_{b,1}(\boldsymbol{y})$.

We recall that for unbiased estimators, minimizing the variance $V(\boldsymbol{x}; \hat{\boldsymbol{x}})$ is equivalent to minimizing the mean power $P(\boldsymbol{x}; \hat{\boldsymbol{x}}) = \mathsf{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}}(\boldsymbol{y})\|_2^2\}$ (see (5.4)); furthermore, $P(\boldsymbol{x}; \hat{\boldsymbol{x}}) = \sum_{k=1}^{N} P(\boldsymbol{x}; \hat{x}_k)$ with $P(\boldsymbol{x}; \hat{x}_k) \triangleq \mathsf{E}_{\boldsymbol{x}}\{(\hat{x}_k(\boldsymbol{y}))^2\}$. For the proposed estimator $\hat{\boldsymbol{x}}_b$, $P(\boldsymbol{x}'; \hat{x}_{b,k}) = P\big(\boldsymbol{x}'; \hat{x}_k^{(\mathbf{0})}\big)$ except for $k = 1$. Therefore, our goal is to choose A such that $P(\boldsymbol{x}'; \hat{x}_{b,1})$ is smaller than $P\big(\boldsymbol{x}'; \hat{x}_1^{(\mathbf{0})}\big) = \sigma^2 + (x_1')^2 = \sigma^2$. We have

$$P(\boldsymbol{x}'; \hat{x}_{b,1}) = \mathsf{E}_{\boldsymbol{x}'}\left\{\left(y_1 + A y_1 \prod_{l=2}^{S+1} h(y_l)\right)^2\right\} = \alpha A^2 + \beta A + \gamma \tag{5.45}$$

with

$$\alpha = \mathsf{E}_{\boldsymbol{x}'}\left\{y_1^2 \prod_{l=2}^{S+1} h^2(y_l)\right\}, \qquad \beta = \mathsf{E}_{\boldsymbol{x}'}\left\{2y_1^2 \prod_{l=2}^{S+1} h(y_l)\right\}, \qquad \gamma = \mathsf{E}_{\boldsymbol{x}'}\{y_1^2\} = \sigma^2.$$

Note that $\gamma = P\big(\boldsymbol{x}'; \hat{x}_1^{(\mathbf{0})}\big)$. From (5.45), the A minimizing $P(\boldsymbol{x}'; \hat{x}_{b,1})$ is obtained as $-\beta/(2\alpha)$; the associated minimum $P(\boldsymbol{x}'; \hat{x}_{b,1})$ is given by $\gamma - \beta^2/(4\alpha^2)$. It can be shown that β is nonzero due to the construction of $h(y)$. It follows that β is positive, and therefore $P(\boldsymbol{x}'; \hat{x}_{b,1})$ is smaller than $\gamma = P\big(\boldsymbol{x}'; \hat{x}_1^{(\mathbf{0})}\big)$. Thus, using $A = -\beta/(2\alpha)$ in (5.44), we obtain an estimator $\hat{\boldsymbol{x}}_b$ which has a smaller component power $P(\boldsymbol{x}'; \hat{x}_{b,1})$ than $\hat{\boldsymbol{x}}^{(\mathbf{0})}$. Since $P(\boldsymbol{x}'; \hat{x}_{b,k}) = P\big(\boldsymbol{x}'; \hat{x}_k^{(\mathbf{0})}\big)$ for $k = 2, \ldots, N$, it follows that the overall mean power of $\hat{\boldsymbol{x}}_b$ at $\boldsymbol{x}'$ is smaller than that of $\hat{\boldsymbol{x}}^{(\mathbf{0})}$, i.e., $P(\boldsymbol{x}'; \hat{\boldsymbol{x}}_b) < P(\boldsymbol{x}'; \hat{\boldsymbol{x}}^{(\mathbf{0})})$. Since both estimators are unbiased, this moreover implies that at $\boldsymbol{x}'$, the variance of $\hat{\boldsymbol{x}}_b$ is smaller than that of $\hat{\boldsymbol{x}}^{(\mathbf{0})}$. Thus, $\hat{\boldsymbol{x}}^{(\mathbf{0})}$ cannot be the LMVU estimator at $\boldsymbol{x} = \boldsymbol{x}'$. On the other hand, as we have seen, $\hat{\boldsymbol{x}}^{(\mathbf{0})}$ is the unique LMVU estimator at $\boldsymbol{x} = \mathbf{0}$. We conclude that there does not exist a single unbiased estimator which simultaneously minimizes the variance for all parameters $\boldsymbol{x} \in \mathcal{X}_S$.

5.C Proof of Proposition 5.4

We begin by stating the multivariate HCRB.

Proposition 5.11 (Gorman and Hero [52])**.** *Let $f(\boldsymbol{y}; \boldsymbol{x})$ be a family of pdf's of $\boldsymbol{y}$ indexed by $\boldsymbol{x} \in \mathcal{X}_S$, and let $\boldsymbol{x} + \boldsymbol{v}_1, \ldots, \boldsymbol{x} + \boldsymbol{v}_p$ be a set of points in $\mathcal{X}_S$. Given an estimator $\hat{\boldsymbol{x}}$, define*

$$\begin{aligned} \boldsymbol{m}_{\boldsymbol{x}} &\triangleq \mathsf{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}\} \\ \delta_i \boldsymbol{m}_{\boldsymbol{x}} &\triangleq \boldsymbol{m}_{\boldsymbol{x}+\boldsymbol{v}_i} - \boldsymbol{m}_{\boldsymbol{x}} \\ \boldsymbol{\delta}\boldsymbol{m}_{\boldsymbol{x}} &\triangleq (\delta_1 \boldsymbol{m}_{\boldsymbol{x}} \cdots \delta_p \boldsymbol{m}_{\boldsymbol{x}})^T \end{aligned}$$

and

$$\begin{aligned}\delta_i f &\triangleq f(\boldsymbol{y};\boldsymbol{x}+\boldsymbol{v}_i)-f(\boldsymbol{y};\boldsymbol{x})\\ \delta f &\triangleq (\delta_1 f \cdots \delta_p f)^T\\ \boldsymbol{Q} &\triangleq \mathsf{E}_{\boldsymbol{x}}\left\{\frac{\delta f}{f}\frac{\delta f^T}{f}\right\}.\end{aligned} \tag{5.46}$$

Then, the covariance matrix of $\hat{\boldsymbol{x}}$ *satisfies*

$$C(\boldsymbol{x};\hat{\boldsymbol{x}}) \succeq \delta\boldsymbol{m}_{\boldsymbol{x}}^T \boldsymbol{Q}^{\dagger}\delta\boldsymbol{m}_{\boldsymbol{x}}. \tag{5.47}$$

We will now prove Proposition 5.4 by applying the multivariate HCRB (5.47) to the case of unbiased estimation under Gaussian noise. For an unbiased estimator $\hat{\boldsymbol{x}}$, we have $\boldsymbol{m}_{\boldsymbol{x}} = \boldsymbol{x}$, so $\delta_i \boldsymbol{m}_{\boldsymbol{x}} = \boldsymbol{v}_i$ and further

$$\delta\boldsymbol{m}_{\boldsymbol{x}} = \boldsymbol{V} \triangleq (\boldsymbol{v}_1 \cdots \boldsymbol{v}_p) \tag{5.48}$$

(see (5.14)). We next show that the matrix $\boldsymbol{Q}$ in (5.46) coincides with $\boldsymbol{J}$ in (5.15). Because of the Gaussian noise, $f(\boldsymbol{y};\boldsymbol{x}) = (2\pi\sigma^2)^{-N/2}\exp\big(-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)\big)$, and thus we obtain by direct calculation

$$\frac{\delta_i f}{f} = \exp\left(\frac{2\boldsymbol{v}_i^T(\boldsymbol{y}-\boldsymbol{x}) - \|\boldsymbol{v}_i\|_2^2}{2\sigma^2}\right) - 1$$

and consequently

$$\begin{aligned}(\boldsymbol{Q})_{i,j} &= \mathsf{E}_{\boldsymbol{x}}\left\{\frac{\delta_i f}{f}\frac{\delta_j f}{f}\right\}\\ &= 1 - \exp\left(-\frac{\|\boldsymbol{v}_i\|_2^2}{2\sigma^2}\right)\mathsf{E}_{\boldsymbol{x}}\left\{\exp\left(\frac{\boldsymbol{v}_i^T(\boldsymbol{y}-\boldsymbol{x})}{\sigma^2}\right)\right\} - \exp\left(-\frac{\|\boldsymbol{v}_j\|_2^2}{2\sigma^2}\right)\mathsf{E}_{\boldsymbol{x}}\left\{\exp\left(\frac{\boldsymbol{v}_j^T(\boldsymbol{y}-\boldsymbol{x})}{\sigma^2}\right)\right\}\\ &\qquad + \exp\left(-\frac{\|\boldsymbol{v}_i\|_2^2+\|\boldsymbol{v}_j\|_2^2}{2\sigma^2}\right)\mathsf{E}_{\boldsymbol{x}}\left\{\exp\left(\frac{(\boldsymbol{v}_i+\boldsymbol{v}_j)^T(\boldsymbol{y}-\boldsymbol{x})}{\sigma^2}\right)\right\}.\end{aligned}$$

Now $\mathsf{E}_{\boldsymbol{x}}\{\exp(\boldsymbol{a}^T(\boldsymbol{y}-\boldsymbol{x}))\}$ is the moment-generating function of the zero-mean Gaussian random vector $\boldsymbol{y}-\boldsymbol{x}$, which equals $\exp\big(\|\boldsymbol{a}\|_2^2\sigma^2/2\big)$. We thus have

$$\begin{aligned}(\boldsymbol{Q})_{i,j} &= 1 - \exp\left(-\frac{\|\boldsymbol{v}_i\|_2^2}{2\sigma^2}\right)\exp\left(\frac{\|\boldsymbol{v}_i\|_2^2}{2\sigma^2}\right) - \exp\left(-\frac{\|\boldsymbol{v}_j\|_2^2}{2\sigma^2}\right)\exp\left(\frac{\|\boldsymbol{v}_j\|_2^2}{2\sigma^2}\right)\\ &\qquad + \exp\left(-\frac{\|\boldsymbol{v}_i\|_2^2+\|\boldsymbol{v}_j\|_2^2}{2\sigma^2}\right)\exp\left(\frac{\|\boldsymbol{v}_i+\boldsymbol{v}_j\|_2^2}{2\sigma^2}\right)\\ &= -1 + \exp\left(\frac{\boldsymbol{v}_i^T\boldsymbol{v}_j}{\sigma^2}\right)\end{aligned} \tag{5.49}$$

which equals $(\boldsymbol{J})_{i,j}$ in (5.15). Inserting (5.48) and (5.49) into (5.47), we obtain (5.13). Finally, taking the trace of both sides of (5.13) yields (5.16).

5.D Obtaining the CRB from the HCRB

We will demonstrate that the CRB (5.12) can be obtained as a limit of HCRBs (5.16) by choosing the test points $\boldsymbol{v}_i$ according to (5.17) and letting $t \to 0$. Since the test points (5.17) are orthogonal vectors, it follows from (5.15) that the matrix $\boldsymbol{J}$ is diagonal. More specifically, we have

$$\boldsymbol{J} = \begin{cases} \left[\exp(t^2/\sigma^2) - 1\right] \boldsymbol{I}_S\,, & \|\boldsymbol{x}\|_0 = S \\ \left[\exp(t^2/\sigma^2) - 1\right] \boldsymbol{I}_N\,, & \|\boldsymbol{x}\|_0 < S\,. \end{cases}$$

Thus, both for $\|\boldsymbol{x}\|_0 = S$ and for $\|\boldsymbol{x}\|_0 < S$, the pseudoinverse of $\boldsymbol{J}$ is obtained simply by inverting the diagonal entries of $\boldsymbol{J}$. From (5.16), we then obtain

$$\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \;\geq\; \begin{cases} \dfrac{St^2}{\exp(t^2/\sigma^2) - 1}\,, & \|\boldsymbol{x}\|_0 = S \\ \dfrac{Nt^2}{\exp(t^2/\sigma^2) - 1}\,, & \|\boldsymbol{x}\|_0 < S\,. \end{cases} \tag{5.50}$$

We now use the third-order Taylor series expansion

$$\exp\left(\frac{t^2}{\sigma^2}\right) \;=\; 1 + \frac{t^2}{\sigma^2} + \frac{\tau^4}{2\sigma^4}\,, \qquad \text{where } \tau \in [0, t]\,. \tag{5.51}$$

Substituting (5.51) into (5.50) yields

$$\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \;\geq\; \begin{cases} \dfrac{St^2}{t^2/\sigma^2 + \tau^4/(2\sigma^4)}\,, & \|\boldsymbol{x}\|_0 = S \\ \dfrac{Nt^2}{t^2/\sigma^2 + \tau^4/(2\sigma^4)}\,, & \|\boldsymbol{x}\|_0 < S\,. \end{cases} \tag{5.52}$$

In the limit as $t \to 0$, $\tau^4 \in [0, t^4]$ decays faster than t^2, and thus the bound (5.52) converges to the CRB (5.12).

The CRB can also be obtained by formally replacing $\exp\left(t^2/\sigma^2\right)$ with $1 + t^2/\sigma^2$ in (5.50). From (5.51), we have $\exp\left(t^2/\sigma^2\right) \geq 1 + t^2/\sigma^2$ for all $t > 0$. This shows that for any $t > 0$, the bound (5.50) is lower than the CRB (5.12). Thus, the CRB (which, as shown above, is obtained using the test points (5.17) in the limit $t \to 0$) is tighter than any bound that is obtained using the test points (5.17) for any fixed $t > 0$.

5.E Proof of Theorem 5.5

We will prove the HCRB-type bound in (5.19). For $\|\boldsymbol{x}\|_0 < S$, (5.19) was already demonstrated by the CRB (5.12), and thus it remains to show (5.19) for $\|\boldsymbol{x}\|_0 = S$. This will be done by plugging

the test points (5.18) into the HCRB (5.16), calculating the resulting bound for an arbitrary constant $t > 0$, and then taking the limit as $t \to 0$. We will use the following lemma, whose proof is provided at the end of this appendix.

Lemma 5.12. *Let $\boldsymbol{P}$ be an $(r+1) \times (r+1)$ matrix with the following structure:*

$$\boldsymbol{P} = \begin{pmatrix} a & b\mathbf{1}^T \\ b\mathbf{1} & \boldsymbol{M} \end{pmatrix} = \begin{pmatrix} a & b & b & b & \cdots & b \\ b & d & c & c & \cdots & c \\ b & c & d & c & \cdots & c \\ b & c & c & \ddots & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \ddots & c \\ b & c & c & \cdots & c & d \end{pmatrix} \tag{5.53}$$

where $\mathbf{1}$ is the column vector of dimension r whose entries all equal 1, and

$$\boldsymbol{M} = (d-c)\boldsymbol{I}_r + c\mathbf{1}\mathbf{1}^T. \tag{5.54}$$

Let

$$q \triangleq rb^2 - ad - (r-1)ac \tag{5.55}$$

and assume that

$$d-c \neq 0, \qquad d+(r-1)c \neq 0, \qquad q \neq 0. \tag{5.56}$$

Then, $\boldsymbol{P}$ is nonsingular and its inverse is given by

$$\boldsymbol{P}^{-1} = \begin{pmatrix} a' & b'\mathbf{1}^T \\ b'\mathbf{1} & \boldsymbol{M}' \end{pmatrix} = \begin{pmatrix} a' & b' & b' & b' & \cdots & b' \\ b' & d' & c' & c' & \cdots & c' \\ b' & c' & d' & c' & \cdots & c' \\ b' & c' & c' & \ddots & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \ddots & c' \\ b' & c' & c' & \cdots & c' & d' \end{pmatrix} \tag{5.57}$$

where $\boldsymbol{M}' = (d'-c')\boldsymbol{I}_r + c'\mathbf{1}\mathbf{1}^T$ and

$$a' = -\frac{d+(r-1)c}{q}, \qquad b' = \frac{b}{q}, \qquad c' = \frac{ac-b^2}{(d-c)q}, \qquad d' = \frac{(r-1)b^2 - (r-2)ac - ad}{(d-c)q}. \tag{5.58}$$

Let $\|\boldsymbol{x}\|_0 = S$, and assume for concreteness and without loss of generality that $\mathrm{supp}(\boldsymbol{x}) = \{1, \ldots, S\}$ and that ξ, the smallest (in magnitude) nonzero component of $\boldsymbol{x}$, is the Sth entry. A direct calculation of the matrix $\boldsymbol{J}$ in (5.15) based on the test points (5.18) then yields

$$\boldsymbol{J} = \begin{pmatrix} a\boldsymbol{I}_{S-1} & \mathbf{0}_{(S-1)\times(r+1)} \\ \mathbf{0}_{(r+1)\times(S-1)} & \boldsymbol{P} \end{pmatrix}.$$

Here, $\boldsymbol{P}$ is an $(r+1)\times(r+1)$ matrix, where $r = N-S$, having the structure (5.53) with entries

$$a = e^{t^2/\sigma^2}-1\,, \qquad b = e^{-t\xi/\sigma^2}-1\,, \qquad c = e^{\xi^2/\sigma^2}-1\,, \qquad d = e^{(t^2+\xi^2)/\sigma^2}-1\,. \tag{5.59}$$

We now apply Lemma 5.12 in order to show that $\boldsymbol{J}$ is nonsingular and to calculate its inverse. More precisely, it suffices to calculate the inverse for all but a finite number of values of t, since any finite set of values can simply be excluded from consideration when t tends to 0. When applying Lemma 5.12, we first have to verify that the conditions (5.56) hold for all but a finite number of values of t. By substituting (5.59), it is seen that the left-hand sides of (5.56) are nonconstant entire functions of t, and thus have a finite number of roots on any compact set of values of t. By Lemma 5.12, this implies that $\boldsymbol{J}$ is nonsingular for all but a finite number of values of values of t, and that the inverse (if it exists) is given by

$$\boldsymbol{J}^{-1} = \begin{pmatrix} \frac{1}{a}\boldsymbol{I}_{S-1} & \boldsymbol{0}_{(S-1)\times(r+1)} \\ \boldsymbol{0}_{(r+1)\times(S-1)} & \boldsymbol{P}^{-1} \end{pmatrix} \tag{5.60}$$

where $\boldsymbol{P}^{-1}$ is given by (5.57) and (5.58), again with $r = N-S$. Next, we observe that for our choice of test points (5.18),

$$\boldsymbol{V}^T\boldsymbol{V} = \begin{pmatrix} t^2\boldsymbol{I}_{S-1} & \boldsymbol{0}_{(S-1)\times(r+1)} \\ \boldsymbol{0}_{(r+1)\times(S-1)} & \widetilde{\boldsymbol{P}} \end{pmatrix} \tag{5.61}$$

where $\widetilde{\boldsymbol{P}}$ is an $(r+1)\times(r+1)$ matrix having the structure (5.53) with entries

$$\tilde{a} = t^2, \qquad \tilde{b} = -t\xi\,, \qquad \tilde{c} = \xi^2, \qquad \tilde{d} = t^2+\xi^2.$$

Using (5.16) together with (5.60) and (5.61), a direct calculation yields

$$\begin{aligned} \varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \;\geq\; \mathrm{tr}\big(\boldsymbol{V}\boldsymbol{J}^{\dagger}\boldsymbol{V}^T\big) = \mathrm{Tr}\Big(\boldsymbol{V}^T\boldsymbol{V}\boldsymbol{J}^{-1}\Big) &= \sum_{i=1}^{N}\sum_{j=1}^{N}(\boldsymbol{V}^T\boldsymbol{V})_{i,j}(\boldsymbol{J}^{-1})_{i,j} \\ &= (S-1)\frac{t^2}{a} + t^2a' - 2rt\xi b' + r(r-1)\xi^2c' + r(t^2+\xi^2)d'. \end{aligned} \tag{5.62}$$

We now take the limit $t\to 0$ in (5.62). For the first term, we obtain

$$(S-1)\frac{t^2}{a} = (S-1)\frac{t^2}{e^{t^2/\sigma^2}-1} = (S-1)\frac{t^2}{t^2/\sigma^2+o(t^2)} \;\longrightarrow\; (S-1)\sigma^2 \tag{5.63}$$

where we have expanded e^{t^2/σ^2} into a second-oder Taylor series. Here, $o(f(t))$ indicates terms which are negligible compared with $f(t)$ when $t\to 0$, i.e., $\lim_{t\to 0} o(f(t))/f(t) = 0$. To find the limit of the second term in (5.62), $t^2a' = -(t^2/q)[d+(r-1)c]$, we first consider the reciprocal of the first factor, t^2/q. We have

$$\frac{q}{t^2} = \frac{1}{t^2}\big[r\big(e^{-t\xi/\sigma^2}-1\big)^2 - \big(e^{t^2/\sigma^2}-1\big)\big(e^{(t^2+\xi^2)/\sigma^2}-1\big) - (r-1)\big(e^{t^2/\sigma^2}-1\big)\big(e^{\xi^2/\sigma^2}-1\big)\big]\,.$$

Expanding some of the t-dependent exponentials into Taylor series, dropping higher-order terms, and simplifying, we obtain

$$\frac{q}{t^2} = \frac{1}{t^2}\left[r\left(\frac{-t\xi}{\sigma^2}+o(t)\right)^2 - \left(\frac{t^2}{\sigma^2}+o(t^2)\right)\left(e^{(t^2+\xi^2)/\sigma^2}-1\right) - (r-1)\left(\frac{t^2}{\sigma^2}+o(t^2)\right)\left(e^{\xi^2/\sigma^2}-1\right)\right]$$
$$\longrightarrow\ r\,\frac{\xi^2}{\sigma^4} - \frac{1}{\sigma^2}\left(e^{\xi^2/\sigma^2}-1\right) - (r-1)\,\frac{1}{\sigma^2}\left(e^{\xi^2/\sigma^2}-1\right) = \frac{r}{\sigma^4}\left[\xi^2 - \sigma^2\left(e^{\xi^2/\sigma^2}-1\right)\right]. \tag{5.64}$$

For the second factor, we obtain

$$d+(r-1)c = e^{(t^2+\xi^2)/\sigma^2} - 1 + (r-1)\left(e^{\xi^2/\sigma^2}-1\right) \longrightarrow\ r\left(e^{\xi^2/\sigma^2}-1\right). \tag{5.65}$$

Then, using (5.64) and (5.65), it is seen that the second term in (5.62) converges to

$$t^2a' = -\frac{t^2}{q}[d+(r-1)c] \longrightarrow\ -\frac{r\left(e^{\xi^2/\sigma^2}-1\right)}{\frac{r}{\sigma^4}\left[\xi^2-\sigma^2\left(e^{\xi^2/\sigma^2}-1\right)\right]} = \sigma^2\left[1+\frac{\xi^2}{\sigma^2\left(e^{\xi^2/\sigma^2}-1\right)-\xi^2}\right]. \tag{5.66}$$

Next, we consider the third term in (5.62), $-2rt\xi b'$, which can be written as $-2r\xi\frac{b/t}{q/t^2}$. We have

$$\frac{b}{t} = \frac{1}{t}\left(e^{-t\xi/\sigma^2}-1\right) = \frac{1}{t}\left(\frac{-t\xi}{\sigma^2}+o(t)\right) \longrightarrow\ -\frac{\xi}{\sigma^2}.$$

Combining with (5.64), we obtain

$$-2rt\xi b' \longrightarrow\ 2r\xi\,\frac{\xi/\sigma^2}{\frac{r}{\sigma^4}\left[\xi^2-\sigma^2\left(e^{\xi^2/\sigma^2}-1\right)\right]} = \frac{2\sigma^2\xi^2}{\xi^2-\sigma^2\left(e^{\xi^2/\sigma^2}-1\right)}. \tag{5.67}$$

The fourth and fifth terms in (5.62) have to be calculated together because each of them by itself diverges. The sum of these terms is

$$\begin{aligned} r(r-1)\xi^2c' + r(t^2+\xi^2)d' &= \frac{r}{(d-c)q}\left[(r-1)\xi^2(ac-b^2) + (t^2+\xi^2)[(r-1)b^2-(r-2)ac-ad]\right] \\ &= \frac{r}{(d-c)q}\left[-\xi^2a(d-c) + t^2[(r-1)b^2-(r-2)ac-ad]\right] \\ &= -\frac{r\xi^2a}{q} + \frac{rt^2}{(d-c)q}\,(q+ac-b^2) \\ &= \underbrace{-\frac{r\xi^2a}{q}}_{z_1} + \underbrace{\frac{rt^2}{d-c}}_{z_2} + \underbrace{\frac{rt^2}{(d-c)q}(ac-b^2)}_{z_3}. \end{aligned} \tag{5.68}$$

Using (5.64), z_1 in (5.68) becomes

$$\begin{aligned} z_1 &= -\frac{r\xi^2a/t^2}{q/t^2} = -r\xi^2\frac{\left(e^{t^2/\sigma^2}-1\right)/t^2}{q/t^2} \\ &\longrightarrow\ -r\xi^2\frac{1/\sigma^2}{\frac{r}{\sigma^4}[\xi^2-\sigma^2\left(e^{\xi^2/\sigma^2}-1\right)]} = -\frac{\sigma^2\xi^2}{\xi^2-\sigma^2\left(e^{\xi^2/\sigma^2}-1\right)}. \end{aligned} \tag{5.69}$$

Furthermore, a direct calculation yields

$$z_2 = \frac{rt^2}{e^{(t^2+\xi^2)/\sigma^2} - e^{\xi^2/\sigma^2}} = re^{-\xi^2/\sigma^2}\frac{t^2}{e^{t^2/\sigma^2}-1} \;\longrightarrow\; r\sigma^2 e^{-\xi^2/\sigma^2}. \tag{5.70}$$

To take the limit of z_3, first note that

$$\frac{ac-b^2}{d-c} = \frac{(e^{t^2/\sigma^2}-1)(e^{\xi^2/\sigma^2}-1) - (e^{-t\xi/\sigma^2}-1)^2}{e^{(t^2+\xi^2)/\sigma^2} - e^{\xi^2/\sigma^2}}$$
$$\longrightarrow\; \frac{(t^2/\sigma^2)(e^{\xi^2/\sigma^2}-1) - (-t\xi/\sigma^2)^2}{e^{\xi^2/\sigma^2}\, t^2/\sigma^2} = \frac{\sigma^2(e^{\xi^2/\sigma^2}-1) - \xi^2}{\sigma^2 e^{\xi^2/\sigma^2}}\,.$$

Together with (5.64), we thus have

$$z_3 = r\frac{t^2}{q}\,\frac{ac-b^2}{d-c} \;\longrightarrow\; r\frac{1}{\frac{r}{\sigma^4}[\xi^2 - \sigma^2(e^{\xi^2/\sigma^2}-1)]}\,\frac{\sigma^2(e^{\xi^2/\sigma^2}-1)-\xi^2}{\sigma^2 e^{\xi^2/\sigma^2}} = -\sigma^2 e^{-\xi^2/\sigma^2}. \tag{5.71}$$

Adding the limits of z_1, z_2, and z_3 in (5.69)–(5.71), we find that the sum of the fourth and fifth terms in (5.62) converges to

$$z_1 + z_2 + z_3 \;\longrightarrow\; \frac{-\sigma^2\xi^2}{\xi^2 - \sigma^2(e^{\xi^2/\sigma^2}-1)} + (r-1)\sigma^2 e^{-\xi^2/\sigma^2}. \tag{5.72}$$

Finally, adding the limits of all terms in (5.62) as given by (5.63), (5.66), (5.67), and (5.72) and simplifying, we obtain the following result for the limit of the bound (5.62) for $t \to 0$:

$$\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}) \;\geq\; S\sigma^2 + (r-1)\sigma^2 e^{-\xi^2/\sigma^2}.$$

This equals (5.19), as claimed.

Proof of Lemma 5.12: We first calculate the inverse of $\boldsymbol{M}$ in (5.54). Applying the Sherman–Morrison–Woodbury formula [103, §2.8]

$$\left(\boldsymbol{A} + c\boldsymbol{u}\boldsymbol{v}^T\right)^{-1} = \boldsymbol{A}^{-1} - \frac{c}{1+c\boldsymbol{v}^T\boldsymbol{A}^{-1}\boldsymbol{u}}\,\boldsymbol{A}^{-1}\boldsymbol{u}\boldsymbol{v}^T\boldsymbol{A}^{-1}$$

to (5.54) and simplifying yields

$$\boldsymbol{M}^{-1} = \frac{1}{d-c}\boldsymbol{I}_r - \frac{c}{(d-c)[d+(r-1)c]}\,\mathbf{1}\mathbf{1}^T. \tag{5.73}$$

Next, we invoke the block inversion lemma [103, §2.8]

$$\begin{pmatrix} \boldsymbol{A} & \mathcal{B}^T \\ \mathcal{B} & \boldsymbol{M} \end{pmatrix}^{-1} = \begin{pmatrix} \boldsymbol{E}^{-1} & -\boldsymbol{E}^{-1}\mathcal{B}^T\boldsymbol{M}^{-1} \\ -\boldsymbol{M}^{-1}\mathcal{B}\boldsymbol{E}^{-1} & \boldsymbol{M}^{-1} + \boldsymbol{M}^{-1}\mathcal{B}\boldsymbol{E}^{-1}\mathcal{B}^T\boldsymbol{M}^{-1} \end{pmatrix}, \quad \text{with } \boldsymbol{E} \triangleq \boldsymbol{A} - \mathcal{B}^T\boldsymbol{M}^{-1}\mathcal{B}\,.$$

Specializing to $\boldsymbol{A} = a$ and $\mathcal{B} = b\mathbf{1}$ as is appropriate for $\boldsymbol{P}$ in (5.53), we obtain for the inverse of $\boldsymbol{P}$

$$\boldsymbol{P}^{-1} = \begin{pmatrix} 1/e & -(b/e)\mathbf{1}^T\boldsymbol{M}^{-1} \\ -(b/e)\boldsymbol{M}^{-1}\mathbf{1} & \boldsymbol{M}^{-1} + (b^2/e)\boldsymbol{M}^{-1}\mathbf{1}\mathbf{1}^T\boldsymbol{M}^{-1} \end{pmatrix}, \quad \text{with } e \triangleq a - b^2\mathbf{1}^T\boldsymbol{M}^{-1}\mathbf{1}\,. \tag{5.74}$$

We now develop the various blocks of $\boldsymbol{P}^{-1}$ by using the expression of $\boldsymbol{M}^{-1}$ in (5.73). We first consider the upper-left block, $1/e$. We have

$$e = a - \frac{b^2}{d-c}\,\mathbf{1}^T\left[\boldsymbol{I}_r - \frac{c}{d+(r-1)c}\mathbf{1}\mathbf{1}^T\right]\mathbf{1} = a - \frac{b^2}{d-c}\left[r - \frac{cr^2}{d+(r-1)c}\right] = \frac{ad+(r-1)ac-rb^2}{d+(r-1)c}\,.$$

Thus, using the definitions in (5.55) and (5.58) yields

$$\frac{1}{e} = -\frac{d+(r-1)c}{q} = a' \tag{5.75}$$

which proves the validity of the upper-left entry of $\boldsymbol{P}^{-1}$ in (5.57). Next, using (5.73) and (5.75) and simplifying, the upper-right block in (5.74) becomes

$$-\frac{b}{e}\,\mathbf{1}^T\boldsymbol{M}^{-1} = -ba'\left[\frac{1}{d-c} - \frac{rc}{(d-c)[d+(r-1)c]}\right]\mathbf{1}^T = -\frac{ba'}{d+(r-1)c}\,\mathbf{1}^T = \frac{b}{q}\mathbf{1}^T = b'\mathbf{1}^T.$$

Thus, we have shown the validity of the first row and first column of $\boldsymbol{P}^{-1}$ in (5.57). Finally, to develop the remaining block $\boldsymbol{M}^{-1} + (b^2/e)\boldsymbol{M}^{-1}\mathbf{1}\mathbf{1}^T\boldsymbol{M}^{-1}$ in (5.74), we first calculate

$$\boldsymbol{u} \;\triangleq\; \boldsymbol{M}^{-1}\mathbf{1} = \frac{1}{d-c}\left[1 - \frac{rc}{d+(r-1)c}\right]\mathbf{1} = \frac{1}{d+(r-1)c}\,\mathbf{1}\,. \tag{5.76}$$

We then have

$$\boldsymbol{M}^{-1} + \frac{b^2}{e}\boldsymbol{M}^{-1}\mathbf{1}\mathbf{1}^T\boldsymbol{M}^{-1} = \boldsymbol{M}^{-1} + b^2a'\boldsymbol{u}\boldsymbol{u}^T = \frac{1}{d-c}\,\boldsymbol{I}_r - \frac{1}{d+(r-1)c}\left[\frac{c}{d-c} + \frac{b^2}{q}\right]\mathbf{1}\mathbf{1}^T \tag{5.77}$$

where (5.73), (5.76), and the definition of a' in (5.58) were used. Using the definition of q in (5.55) and simplifying, the factor in brackets can be written as

$$\frac{c}{d-c} + \frac{b^2}{q} = \frac{cq+(d-c)b^2}{(d-c)q} = \frac{[d+(r-1)c](b^2-ac)}{(d-c)q}\,.$$

Substituting back into (5.77), we obtain

$$\boldsymbol{M}^{-1} + \frac{b^2}{e}\boldsymbol{M}^{-1}\mathbf{1}\mathbf{1}^T\boldsymbol{M}^{-1} = \frac{1}{d-c}\boldsymbol{I}_r - \frac{b^2-ac}{(d-c)q}\mathbf{1}\mathbf{1}^T = \frac{1}{d-c}\boldsymbol{I}_r + c'\mathbf{1}\mathbf{1}^T.$$

Thus, within the $r\times r$ lower-right block of $\boldsymbol{P}^{-1}$, the off-diagonal entries all equal c', as required. Furthermore, the diagonal entries in this block are given by

$$\frac{1}{d-c} - \frac{b^2-ac}{(d-c)q} = \frac{(r-1)b^2 - ad - (r-2)ac}{(d-c)q} = d'$$

which completes the proof of the lemma. □

5.F Proof of Lemma 5.6

Let $\boldsymbol{x} \in \mathcal{X}_S$ with $\|\boldsymbol{x}\|_0 = S$ and consider a fixed $k \in \operatorname{supp}(\boldsymbol{x})$. We have to show that a solution of (5.21), i.e.,

$$\arg\min_{\hat{x}(\cdot)\in\mathcal{U}^k} \mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\}\,, \qquad \text{with } \mathcal{U}^k = \{\hat{x}(\cdot) \,|\, \mathsf{E}_{\tilde{\boldsymbol{x}}}\{\hat{x}(\boldsymbol{y})\} = \tilde{x}_k \text{ for all } \tilde{\boldsymbol{x}} \in \mathcal{X}_S\} \tag{5.78}$$

is given by $\hat{x}_k^{(\boldsymbol{x})}(\boldsymbol{y}) = y_k$. Let $\varepsilon_0 \triangleq \min_{\hat{x}(\cdot)\in\mathcal{U}^k} \mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\}$ denote the mean power of the LMVU estimator defined by (5.78). We will show that $\varepsilon_0 \geq \sigma^2 + x_k^2$ and, furthermore, that $\sigma^2 + x_k^2$ is achieved by the estimator $\hat{x}_k^{(\boldsymbol{x})}(\boldsymbol{y}) = y_k$.

Let $\mathcal{C}_{\boldsymbol{x}}^k$ denote the set of all S-sparse vectors $\tilde{\boldsymbol{x}}$ which equal $\boldsymbol{x}$ except possibly for the kth component, i.e., $\mathcal{C}_{\boldsymbol{x}}^k \triangleq \{\tilde{\boldsymbol{x}} \in \mathcal{X}_S \,|\, \tilde{x}_l = x_l \text{ for all } l \neq k\}$. Consider the modified optimization problem

$$\arg\min_{\hat{x}(\cdot)\in\mathcal{U}_{\boldsymbol{x}}^k} \mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\}\,, \qquad \text{with } \mathcal{U}_{\boldsymbol{x}}^k \triangleq \{\hat{x}(\cdot) \,|\, \mathsf{E}_{\tilde{\boldsymbol{x}}}\{\hat{x}(\boldsymbol{y})\} = \tilde{x}_k \text{ for all } \tilde{\boldsymbol{x}} \in \mathcal{C}_{\boldsymbol{x}}^k\} \tag{5.79}$$

and let $\varepsilon_0' \triangleq \min_{\hat{x}(\cdot)\in\mathcal{U}_{\boldsymbol{x}}^k} \mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\}$ denote the mean power of the estimator defined by (5.79). Note the distinction between $\mathcal{U}^k$ and $\mathcal{U}_{\boldsymbol{x}}^k$: $\mathcal{U}^k$ is the set of estimators of x_k which are unbiased for all $\tilde{\boldsymbol{x}} \in \mathcal{X}_S$ whereas $\mathcal{U}_{\boldsymbol{x}}^k$ is the set of estimators of x_k which are unbiased for all $\tilde{\boldsymbol{x}} \in \mathcal{X}_S$ which equal a given, fixed $\boldsymbol{x}$ except possibly for the kth component. Therefore, the unbiasedness requirement expressed by $\mathcal{U}^k$ is more restrictive than that expressed by $\mathcal{U}_{\boldsymbol{x}}^k$, i.e., $\mathcal{U}^k \subseteq \mathcal{U}_{\boldsymbol{x}}^k$, which implies that

$$\varepsilon_0' \leq \varepsilon_0\,. \tag{5.80}$$

We will use the following result, which is proved at the end of this appendix.

Lemma 5.13. *Given an arbitrary estimator $\hat{x}(\boldsymbol{y}) \in \mathcal{U}_{\boldsymbol{x}}^k$, the estimator*

$$\hat{x}_c(y_k) \triangleq \mathsf{E}_{\boldsymbol{x}}\{\hat{x}(\boldsymbol{y})|y_k\} \tag{5.81}$$

also satisfies the constraint $\hat{x}_c(y_k) \in \mathcal{U}_{\boldsymbol{x}}^k$, and its mean power does not exceed that obtained by $\hat{x}$, i.e., $\mathsf{E}_{\boldsymbol{x}}\{(\hat{x}_c(y_k))^2\} \leq \mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\}$.

Thus, to each estimator $\hat{x}(\boldsymbol{y}) \in \mathcal{U}_{\boldsymbol{x}}^k$ which depends on the entire observation $\boldsymbol{y}$, we can always find at least one estimator $\hat{x}_c(y_k) \in \mathcal{U}_{\boldsymbol{x}}^k$ which depends only on the observation component y_k and is at least as good. Therefore, with no loss in optimality, we can restrict the optimization problem (5.79) to estimators $\hat{x}(y_k) \in \mathcal{U}_{\boldsymbol{x}}^k$ which depend on $\boldsymbol{y}$ only via its kth component y_k. This

means that (5.79) can be replaced by

$$\arg\min_{\hat{x}(\cdot)\in\widetilde{\mathcal{U}}^k} \mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(y_k))^2\}\,, \qquad \text{with } \widetilde{\mathcal{U}}^k \triangleq \{\hat{x}(\cdot)\,|\,\mathsf{E}_{\tilde{\boldsymbol{x}}}\{\hat{x}(y_k)\}=\tilde{x}_k \text{ for all } \tilde{\boldsymbol{x}}\in\mathbb{R}^N\}\,. \tag{5.82}$$

Note that in the definition of $\widetilde{\mathcal{U}}^k$, we can use the requirement $\tilde{\boldsymbol{x}}\in\mathbb{R}^N$ instead of $\tilde{\boldsymbol{x}}\in\mathcal{C}_{\boldsymbol{x}}^k$ since the expectation $\mathsf{E}_{\tilde{\boldsymbol{x}}}\{\hat{x}(y_k)\}$ does not depend on the components $\tilde{x}_l$ with $l\neq k$. The corresponding minimum mean power $\min_{\hat{x}(\cdot)\in\widetilde{\mathcal{U}}^k}\mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(y_k))^2\}$ is still equal to ε_0'. However, the new problem (5.82) is equivalent to the classical problem of finding the LMVU estimator of a scalar x_k based on the observation $y_k = x_k + n_k$, with $n_k\sim\mathcal{N}(0,\sigma^2)$. A solution of this latter problem is the estimator $\hat{x}(y_k)=y_k$, whose variance and mean power are σ^2 and $\sigma^2+x_k^2$, respectively [90]. Thus, a solution of (5.82) or, equivalently, of (5.79) is the trivial estimator $\hat{x}(y_k)=y_k$, and

$$\varepsilon_0' = \sigma^2 + x_k^2\,. \tag{5.83}$$

Combining (5.80) and (5.83), we see that the minimum mean power for our original optimization problem (5.78) satisfies

$$\varepsilon_0 \;\geq\; \sigma^2 + x_k^2\,.$$

As we have shown, this lower bound is achieved by the estimator $\hat{x}(y_k)=y_k$. In addition, $\hat{x}(y_k)=y_k$ is an element of $\mathcal{U}^k$, the constraint set of (5.78). Therefore, it is a solution of (5.78).

Proof of Lemma 5.13: Consider a fixed $\boldsymbol{x}\in\mathcal{X}_S$ and an estimator $\hat{x}(\boldsymbol{y})\in\mathcal{U}_{\boldsymbol{x}}^k$. In order to show the first statement of the lemma, $\hat{x}_c(y_k)\in\mathcal{U}_{\boldsymbol{x}}^k$, we first note that

$$\mathsf{E}_{\boldsymbol{x}}\{\hat{x}(\boldsymbol{y})|y_k\} = \mathsf{E}_{\tilde{\boldsymbol{x}}}\{\hat{x}(\boldsymbol{y})|y_k\}\,, \qquad \text{for any } \tilde{\boldsymbol{x}}\in\mathcal{C}_{\boldsymbol{x}}^k\,. \tag{5.84}$$

We now have for $\tilde{\boldsymbol{x}}\in\mathcal{C}_{\boldsymbol{x}}^k$

$$\mathsf{E}_{\tilde{\boldsymbol{x}}}\{\hat{x}_c(y_k)\} \overset{(a)}{=} \mathsf{E}_{\tilde{\boldsymbol{x}}}\{\mathsf{E}_{\boldsymbol{x}}\{\hat{x}(\boldsymbol{y})|y_k\}\} \overset{(b)}{=} \mathsf{E}_{\tilde{\boldsymbol{x}}}\{\mathsf{E}_{\tilde{\boldsymbol{x}}}\{\hat{x}(\boldsymbol{y})|y_k\}\} \overset{(c)}{=} \mathsf{E}_{\tilde{\boldsymbol{x}}}\{\hat{x}(\boldsymbol{y})\} \overset{(d)}{=} \tilde{x}_k$$

where we used the definition (5.81) in (a), the identity (5.84) in (b), the law of total probability [104] in (c), and our assumption $\hat{x}(\boldsymbol{y})\in\mathcal{U}_{\boldsymbol{x}}^k$ in (d). Thus, $\hat{x}_c(y_k)\in\mathcal{U}_{\boldsymbol{x}}^k$.

Next, the inequality $\mathsf{E}_{\boldsymbol{x}}\{(\hat{x}_c(y_k))^2\}\leq\mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\}$ is proved as follows:

$$\mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\} \overset{(a)}{=} \mathsf{E}_{\boldsymbol{x}}\{\mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2|y_k\}\} \overset{(b)}{\geq} \mathsf{E}_{\boldsymbol{x}}\{(\mathsf{E}_{\boldsymbol{x}}\{\hat{x}(\boldsymbol{y})|y_k\})^2\} \overset{(c)}{=} \mathsf{E}_{\boldsymbol{x}}\{(\hat{x}_c(y_k))^2\}$$

where we used the law of total probability in (a), Jensen's inequality for convex functions [98] in (b), and the definition (5.81) in (c). □

5.G Proof of Lemma 5.7

We wish to solve the componentwise optimization problem (5.26), i.e.,

$$\underset{\hat{x}(\cdot)\in\mathcal{U}^k\cap\mathcal{A}_{\boldsymbol{x}}^k}{\arg\min}\ \mathsf{E}_{\boldsymbol{x}}\big\{(\hat{x}(\boldsymbol{y}))^2\big\}, \tag{5.85}$$

for $k\notin\operatorname{supp}(\boldsymbol{x})$. Note that $x_k=0$ and, thus, the variance equals the mean power $\mathsf{E}_{\boldsymbol{x}}\big\{(\hat{x}(\boldsymbol{y}))^2\big\}$.

We first observe that the constraint $\hat{x}\in\mathcal{A}_{\boldsymbol{x}}^k$ implies that the estimator $\hat{x}$ is unbiased, and thus $\mathcal{U}^k\cap\mathcal{A}_{\boldsymbol{x}}^k=\mathcal{A}_{\boldsymbol{x}}^k$. Indeed, using (5.22) and $x_k=0$, we have

$$\begin{aligned}
\mathsf{E}_{\boldsymbol{x}}\{\hat{x}(\boldsymbol{y})\} &= \underbrace{\mathsf{E}_{\boldsymbol{x}}\{y_k\}}_{x_k\,(=0)}+\mathsf{E}_{\boldsymbol{x}}\{\hat{x}'(\boldsymbol{y})\}\\
&= x_k+\frac{1}{(2\pi\sigma^2)^{N/2}}\int_{\mathbb{R}^N}\hat{x}'(\boldsymbol{y})\,e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)}d\boldsymbol{y}\\
&= x_k+\frac{1}{(2\pi\sigma^2)^{N/2}}\int_{\mathbb{R}^{N-1}}e^{-\|\boldsymbol{y}_{\sim k}-\boldsymbol{x}_{\sim k}\|_2^2/(2\sigma^2)}\Bigg[\underbrace{\int_{-\infty}^{\infty}\hat{x}'(\boldsymbol{y})\,e^{-(y_k-0)^2/(2\sigma^2)}dy_k}_{0}\Bigg]d\boldsymbol{y}_{\sim k}\\
&= x_k
\end{aligned}\tag{5.86}$$

where $\boldsymbol{x}_{\sim k}$ and $\boldsymbol{y}_{\sim k}$ denote the $(N-1)$-dimensional vectors obtained from $\boldsymbol{x}$ and $\boldsymbol{y}$ by removing the kth component x_k and y_k, respectively, and the result in (5.86) follows because $\int_{-\infty}^{\infty}\hat{x}'(\boldsymbol{y})\,e^{-y_k^2/(2\sigma^2)}dy_k=0$ due to the odd symmetry assumption (5.24). Thus, we can replace the constraint $\hat{x}(\cdot)\in\mathcal{U}^k\cap\mathcal{A}_{\boldsymbol{x}}^k$ in (5.26) by $\hat{x}(\cdot)\in\mathcal{A}_{\boldsymbol{x}}^k$.

A solution of (5.26) can now be found by noting that for any $\hat{x}(\cdot)\in\mathcal{A}_{\boldsymbol{x}}^k$, we have

$$\begin{aligned}
\mathsf{E}_{\boldsymbol{x}}\big\{(\hat{x}(\boldsymbol{y}))^2\big\} &= \frac{1}{(2\pi\sigma^2)^{N/2}}\int_{\mathbb{R}^N}\big(y_k+\hat{x}'(\boldsymbol{y})\big)^2\,e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)}d\boldsymbol{y}\\
&= \frac{1}{(2\pi\sigma^2)^{N/2}}\int_{\mathbb{R}^N}y_k^2\,e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)}d\boldsymbol{y}\\
&\qquad+\frac{1}{(2\pi\sigma^2)^{N/2}}\int_{\mathbb{R}^N}\big[2y_k\hat{x}'(\boldsymbol{y})+(\hat{x}'(\boldsymbol{y}))^2\big]\,e^{-\|\boldsymbol{y}-\boldsymbol{x}\|_2^2/(2\sigma^2)}d\boldsymbol{y}.
\end{aligned}$$

The first term is equal to $\sigma^2+x_k^2=\sigma^2$. Regarding the second term, let $\boldsymbol{y}_k$ be the length-$(S+1)$ subvector of $\boldsymbol{y}$ that comprises all y_l with $l\in\{k\}\cup\operatorname{supp}(\boldsymbol{x})$. Due to (5.25), $\hat{x}'(\boldsymbol{y})$ depends only on $\boldsymbol{y}_k$ and can thus be written (with some abuse of notation) as $\hat{x}'(\boldsymbol{y}_k)$. Let $\bar{\boldsymbol{y}}_k$ denote the complementary subvector of $\boldsymbol{y}$, i.e., the length-$(N-S-1)$ subvector comprising all y_l with $l\notin\{k\}\cup\operatorname{supp}(\boldsymbol{x})$. Furthermore, let $\boldsymbol{x}_k$ and $\bar{\boldsymbol{x}}_k$ denote the analogous subvectors of $\boldsymbol{x}$. The second integral can then be written as the product

$$\begin{aligned}
&\frac{1}{(2\pi\sigma^2)^{(S+1)/2}}\int_{\mathbb{R}^{S+1}}\big[2y_k\hat{x}'(\boldsymbol{y}_k)+(\hat{x}'(\boldsymbol{y}_k))^2\big]\,e^{-\|\boldsymbol{y}_k-\boldsymbol{x}_k\|_2^2/(2\sigma^2)}d\boldsymbol{y}_k\\
&\qquad\qquad\times\frac{1}{(2\pi\sigma^2)^{(N-S-1)/2}}\int_{\mathbb{R}^{N-S-1}}e^{-\|\bar{\boldsymbol{y}}_k-\bar{\boldsymbol{x}}_k\|_2^2/(2\sigma^2)}d\bar{\boldsymbol{y}}_k\,.
\end{aligned}$$

The second factor is 1, and thus we have

$$\mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\} = \sigma^2 + \frac{1}{(2\pi\sigma^2)^{(S+1)/2}} \int_{\mathbb{R}^{S+1}} \left[2y_k\hat{x}'(\boldsymbol{y}_k) + (\hat{x}'(\boldsymbol{y}_k))^2\right] e^{-\|\boldsymbol{y}_k-\boldsymbol{x}_k\|_2^2/(2\sigma^2)} d\boldsymbol{y}_k\,. \quad (5.87)$$

Using the symmetry property (5.24), this can be written as

$$\mathsf{E}_{\boldsymbol{x}}\{(\hat{x}(\boldsymbol{y}))^2\} = \sigma^2 + \frac{2}{(2\pi\sigma^2)^{(S+1)/2}} \int_{\mathbb{R}_+^{S+1}} \left[2\hat{x}'(\boldsymbol{y}_k)b(\boldsymbol{y}_k) + (\hat{x}'(\boldsymbol{y}_k))^2 c(\boldsymbol{y}_k)\right] d\boldsymbol{y}_k\,, \quad (5.88)$$

with

$$b(\boldsymbol{y}_k) \triangleq y_k\, e^{-y_k^2/(2\sigma^2)} \prod_{l\in\mathrm{supp}(\boldsymbol{x})} \left[e^{-(y_l-x_l)^2/(2\sigma^2)} - e^{-(y_l+x_l)^2/(2\sigma^2)}\right] \quad (5.89)$$

$$c(\boldsymbol{y}_k) \triangleq e^{-y_k^2/(2\sigma^2)} \prod_{l\in\mathrm{supp}(\boldsymbol{x})} \left[e^{-(y_l-x_l)^2/(2\sigma^2)} + e^{-(y_l+x_l)^2/(2\sigma^2)}\right]. \quad (5.90)$$

We sketch the derivation of expressions (5.89) and (5.90) by showing the first of $S+1$ similar sequential calculations. For simplicity of notation and without loss of generality, we assume for this derivation that $k=1$ and $\mathrm{supp}(\boldsymbol{x}) = \{2,\ldots,S+1\}$. The integral in (5.87) then becomes

$$\begin{aligned}\int_{\mathbb{R}^{S+1}} &\left[2y_k\hat{x}'(\boldsymbol{y}_k) + (\hat{x}'(\boldsymbol{y}_k))^2\right] e^{-\|\boldsymbol{y}_k-\boldsymbol{x}_k\|_2^2/(2\sigma^2)} d\boldsymbol{y}_k \\ &= \int_{\mathbb{R}^{S+1}} \left[2y_1\hat{x}'(\boldsymbol{y}_1) + (\hat{x}'(\boldsymbol{y}_1))^2\right] \left[\prod_{l=1}^{S+1} e^{-(y_l-x_l)^2/(2\sigma^2)}\right] d\boldsymbol{y}_1\,. \end{aligned} \quad (5.91)$$

The $\int_{\mathbb{R}^{S+1}}$ integration can now be represented as $\int_{\mathbb{R}^S\times(\mathbb{R}_+\cup\mathbb{R}_-)}$, where the component $\int_{\mathbb{R}^S}$ refers to $y_1,\ldots,y_S$ and the component $\int_{\mathbb{R}_+\cup\mathbb{R}_-}$ refers to y_{S+1}. Then (5.91) can be further processed as

$$\begin{aligned}&\int_{\mathbb{R}^S\times\mathbb{R}_+} \left[2y_1\hat{x}'(\boldsymbol{y}_1) + (\hat{x}'(\boldsymbol{y}_1))^2\right] \left[\prod_{l=1}^{S+1} e^{-(y_l-x_l)^2/(2\sigma^2)}\right] d\boldsymbol{y}_1 \\ &\qquad + \int_{\mathbb{R}^S\times\mathbb{R}_-} \left[2y_1\hat{x}'(\boldsymbol{y}_1) + (\hat{x}'(\boldsymbol{y}_1))^2\right] \left[\prod_{l=1}^{S+1} e^{-(y_l-x_l)^2/(2\sigma^2)}\right] d\boldsymbol{y}_1 \\ &\overset{(*)}{=} \int_{\mathbb{R}^S\times\mathbb{R}_+} \Big[2y_1\hat{x}'(\boldsymbol{y}_1)\left(e^{-(y_{S+1}-x_{S+1})^2/(2\sigma^2)} - e^{-(y_{S+1}+x_{S+1})^2/(2\sigma^2)}\right) \\ &\qquad + (\hat{x}'(\boldsymbol{y}_1))^2\left(e^{-(y_{S+1}-x_{S+1})^2/(2\sigma^2)} + e^{-(y_{S+1}+x_{S+1})^2/(2\sigma^2)}\right)\Big] \left[\prod_{l=1}^{S} e^{-(y_l-x_l)^2/(2\sigma^2)}\right] d\boldsymbol{y}_1 \end{aligned}$$

where the odd symmetry property (5.24) was used in $(*)$. After performing this type of manipulation S times, the integral is obtained in the form

$$\begin{aligned}\int_{\mathbb{R}\times\mathbb{R}_+^S} \Bigg[&2y_1\hat{x}'(\boldsymbol{y}_1) \prod_{l=2}^{S+1} \left(e^{-(y_l-x_l)^2/(2\sigma^2)} - e^{-(y_l+x_l)^2/(2\sigma^2)}\right) \\ &+ (\hat{x}'(\boldsymbol{y}_1))^2 \prod_{l=2}^{S+1} \left(e^{-(y_l-x_l)^2/(2\sigma^2)} + e^{-(y_l+x_l)^2/(2\sigma^2)}\right)\Bigg] e^{-y_1^2/(2\sigma^2)}\, d\boldsymbol{y}_1 \end{aligned}$$

where $x_1 = 0$ was used. With $y_1\hat{x}'(y_1,\ldots) = (-y_1)\hat{x}'(-y_1,\ldots)$, this becomes further

$$\int_{\mathbb{R}_+\times\mathbb{R}_+^S}\Bigg[2y_1\hat{x}'(\boldsymbol{y}_1)2e^{-y_1^2/(2\sigma^2)}\prod_{l=2}^{S+1}\left(e^{-(y_l-x_l)^2/(2\sigma^2)}-e^{-(y_l+x_l)^2/(2\sigma^2)}\right) \\ +(\hat{x}'(\boldsymbol{y}_1))^2 2e^{-y_1^2/(2\sigma^2)}\prod_{l=2}^{S+1}\left(e^{-(y_l-x_l)^2/(2\sigma^2)}+e^{-(y_l+x_l)^2/(2\sigma^2)}\right)\Bigg]d\boldsymbol{y}_1\,.$$

Finally, removing our "notational simplicity" assumptions $k=1$ and $\mathrm{supp}(\boldsymbol{x}) = \{2,\ldots,S+1\}$, this can be written for general k and $\mathrm{supp}(\boldsymbol{x})$ as

$$2e^{-y_k^2/(2\sigma^2)}\int_{\mathbb{R}_+^{S+1}}\Bigg[2y_k\hat{x}'(\boldsymbol{y}_k)\prod_{l\in\mathrm{supp}(\boldsymbol{x})}\left(e^{-(y_l-x_l)^2/(2\sigma^2)}-e^{-(y_l+x_l)^2/(2\sigma^2)}\right) \\ +(\hat{x}'(\boldsymbol{y}_k))^2\prod_{l\in\mathrm{supp}(\boldsymbol{x})}\left(e^{-(y_l-x_l)^2/(2\sigma^2)}+e^{-(y_l+x_l)^2/(2\sigma^2)}\right)\Bigg]d\boldsymbol{y}_k\,. \tag{5.92}$$

Inserting (5.92) into (5.87) yields (5.88).

The integral $\int_{\mathbb{R}_+^{S+1}}\left[2\hat{x}'(\boldsymbol{y}_k)b(\boldsymbol{y}_k)+(\hat{x}'(\boldsymbol{y}_k))^2c(\boldsymbol{y}_k)\right]d\boldsymbol{y}_k$ is minimized with respect to $\hat{x}'(\boldsymbol{y}_k)$ by minimizing the integrand $2\hat{x}'(\boldsymbol{y}_k)b(\boldsymbol{y}_k)+(\hat{x}'(\boldsymbol{y}_k))^2c(\boldsymbol{y}_k)$ pointwise for each value of $\boldsymbol{y}_k\in\mathbb{R}_+^{S+1}$. This is easily done by completing the square in $\hat{x}'(\boldsymbol{y}_k)$, yielding the optimization problem $\min_{\hat{x}'(\boldsymbol{y}_k)}\left[\hat{x}'(\boldsymbol{y}_k)+b(\boldsymbol{y}_k)/c(\boldsymbol{y}_k)\right]^2$. Thus, the optimal $\hat{x}'(\boldsymbol{y}_k)$ is obtained as

$$\hat{x}'_{k,\boldsymbol{x}}(\boldsymbol{y}_k) \triangleq -\frac{b(\boldsymbol{y}_k)}{c(\boldsymbol{y}_k)} = -y_k\prod_{l\in\mathrm{supp}(\boldsymbol{x})}\tanh\left(\frac{x_l y_l}{\sigma^2}\right) \qquad \text{for all } \boldsymbol{y}_k\in\mathbb{R}_+^{S+1}$$

and the corresponding pointwise minimum of the integrand is given by $-(b(\boldsymbol{y}_k))^2/c(\boldsymbol{y}_k)$. The extension $\hat{x}'_{k,\boldsymbol{x}}(\boldsymbol{y})$ to all $\boldsymbol{y}\in\mathbb{R}^N$ is then obtained using the properties (5.24) and (5.25), and the optimal component estimator solving (5.26) follows as $\hat{x}_{k,\boldsymbol{x}}(\boldsymbol{y}) = y_k + \hat{x}'_{k,\boldsymbol{x}}(\boldsymbol{y})$. The corresponding minimum variance, denoted by $\mathrm{BB}_{\mathrm{c}}^k(\boldsymbol{x})$, is obtained by substituting the minimum value of the integrand, $-(b(\boldsymbol{y}_k))^2/c(\boldsymbol{y}_k)$, in (5.88). This yields

$$\mathrm{BB}_{\mathrm{c}}^k(\boldsymbol{x}) \triangleq \mathsf{E}_{\boldsymbol{x}}\left\{(\hat{x}_{k,\boldsymbol{x}}(\boldsymbol{y}))^2\right\} = \sigma^2 - \frac{2}{(2\pi\sigma^2)^{(S+1)/2}}\int_{\mathbb{R}_+^{S+1}}\frac{(b(\boldsymbol{y}_k))^2}{c(\boldsymbol{y}_k)}d\boldsymbol{y}_k\,. \tag{5.93}$$

Inserting (5.89) and (5.90) into (5.93) and simplifying gives (5.27).

5.H Proof of Equation (5.30)

To show (5.30), we consider $g(x;\sigma^2)$ for $x\geq 0$ (this is sufficient since $g(-x;\sigma^2) = g(x;\sigma^2)$), and we use the simple bound $\tanh(x)\geq 1-e^{-x}$, which can be verified using elementary calculus.

We then obtain from (5.28), for $x \geq 0$,

$$
\begin{aligned}
g(x;\sigma^2) \;&\geq\; \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty e^{-(x^2+y^2)/(2\sigma^2)} \sinh\!\left(\frac{xy}{\sigma^2}\right) \big(1 - e^{-xy/\sigma^2}\big) dy \\
&= \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty \big[e^{-(x-y)^2/(2\sigma^2)} - e^{-(x+y)^2/(2\sigma^2)}\big] \big(1 - e^{-xy/\sigma^2}\big) dy \\
&= \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty \big[e^{-(x-y)^2/(2\sigma^2)} - e^{-(x^2+y^2)/(2\sigma^2)} - e^{-(x+y)^2/(2\sigma^2)} + e^{-(x+y)^2/(2\sigma^2)} e^{-xy/\sigma^2}\big] dy \\
&\geq \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty \big[e^{-(x-y)^2/(2\sigma^2)} - e^{-(x^2+y^2)/(2\sigma^2)} - e^{-(x+y)^2/(2\sigma^2)}\big] dy \\
&= \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty e^{-(x-y)^2/(2\sigma^2)} dy \;-\; \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty \big[e^{-(x^2+y^2)/(2\sigma^2)} + e^{-(x+y)^2/(2\sigma^2)}\big] dy \,.
\end{aligned}
$$

The first integral can be written as $\frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty e^{-(x-y)^2/(2\sigma^2)} dy = 1 - \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^0 e^{-(x-y)^2/(2\sigma^2)} dy = 1 - \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty e^{-(x+y)^2/(2\sigma^2)} dy$. The bound thus becomes

$$
\begin{aligned}
g(x;\sigma^2) \;&\geq\; 1 - \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty \big[2e^{-(x+y)^2/(2\sigma^2)} + e^{-(x^2+y^2)/(2\sigma^2)}\big] dy \\
&= 1 - \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty \big[2e^{-2xy/(2\sigma^2)} \,+\, 1\big]\, e^{-(x^2+y^2)/(2\sigma^2)} dy \\
&\overset{(*)}{\geq}\; 1 - \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^\infty 3e^{-(x^2+y^2)/(2\sigma^2)} dy \\
&= 1 - \frac{3}{\sqrt{2\pi\sigma^2}}\, e^{-x^2/(2\sigma^2)} \int_0^\infty e^{-y^2/(2\sigma^2)}\, dy \\
&= 1 - \frac{3}{2}\, e^{-x^2/(2\sigma^2)}
\end{aligned}
$$

where $e^{-2xy/(2\sigma^2)} \leq 1$ was used in $(*)$. This bound on $g(x;\sigma^2)$ is actually valid for all $x \in \mathbb{R}$ because $g(-x;\sigma^2) = g(x;\sigma^2)$. Inserting it in (5.27), we obtain

$$
\mathrm{BB}_{\mathrm{c}}^k(\boldsymbol{x}) \;\leq\; \left[1 - \prod_{l \in \mathrm{supp}(\boldsymbol{x})} \left(1 - \frac{3}{2}\, e^{-x_l^2/(2\sigma^2)}\right)\right] \sigma^2 \,. \tag{5.94}
$$

The statement in (5.30) follows since we have (note that $\sum_{\mathcal{I}\subseteq\mathrm{supp}(\boldsymbol{x})}$ denotes the sum over all possible subsets $\mathcal{I}$ of $\mathrm{supp}(\boldsymbol{x})$, including $\mathrm{supp}(\boldsymbol{x})$ and the empty set $\emptyset$)

$$\begin{aligned}
1-\prod_{l\in\mathrm{supp}(\boldsymbol{x})}\left(1-\frac{3}{2}\,e^{-x_l^2/(2\sigma^2)}\right) &= 1-\sum_{\mathcal{I}\subseteq\mathrm{supp}(\boldsymbol{x})}\prod_{l\in\mathcal{I}}\left(-\frac{3}{2}\,e^{-x_l^2/(2\sigma^2)}\right)\\
&= -\sum_{\mathcal{I}\subseteq\mathrm{supp}(\boldsymbol{x}),\,\mathcal{I}\neq\emptyset}\prod_{l\in\mathcal{I}}\left(-\frac{3}{2}\,e^{-x_l^2/(2\sigma^2)}\right)\\
&\leq \sum_{\mathcal{I}\subseteq\mathrm{supp}(\boldsymbol{x}),\,\mathcal{I}\neq\emptyset}\prod_{l\in\mathcal{I}}\left(\frac{3}{2}\,e^{-x_l^2/(2\sigma^2)}\right)\\
&\leq \sum_{\mathcal{I}\subseteq\mathrm{supp}(\boldsymbol{x}),\,\mathcal{I}\neq\emptyset}\prod_{l\in\mathcal{I}}\left(\frac{3}{2}\,e^{-\xi^2/(2\sigma^2)}\right)\\
&= \sum_{\mathcal{I}\subseteq\mathrm{supp}(\boldsymbol{x}),\,\mathcal{I}\neq\emptyset}\left(\frac{3}{2}\,e^{-\xi^2/(2\sigma^2)}\right)^{|\mathcal{I}|}\\
&\leq \sum_{\mathcal{I}\subseteq\mathrm{supp}(\boldsymbol{x}),\,\mathcal{I}\neq\emptyset}\left(\frac{3}{2}\right)^{S}e^{-\xi^2/(2\sigma^2)}\\
&\leq 2^S\left(\frac{3}{2}\right)^{S}e^{-\xi^2/(2\sigma^2)}\\
&= 3^S e^{-\xi^2/(2\sigma^2)}
\end{aligned}$$

where we have used the fact that the number of different subsets $\mathcal{I}\subseteq\mathrm{supp}(\boldsymbol{x})$ is $2^{|\,\mathrm{supp}(\boldsymbol{x})|}=2^S$. Inserting the last bound in (5.94) and, in turn, the resulting bound on $\mathrm{BB}_{\mathrm{c}}^k(\boldsymbol{x})$ in (5.29) yields (5.30).

5.I MSE of the ML Estimator

We calculate the MSE $\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}_{\mathrm{ML}})$ of the ML estimator $\hat{\boldsymbol{x}}_{\mathrm{ML}}$ in (5.7). Let $\hat{x}_{\mathrm{ML},k}$ denote the kth component of $\hat{\boldsymbol{x}}_{\mathrm{ML}}$. We have

$$\begin{aligned}
\varepsilon(\boldsymbol{x};\hat{\boldsymbol{x}}_{\mathrm{ML}}) &= \sum_{k=1}^{N}\mathsf{E}_{\boldsymbol{x}}\big\{(\hat{x}_{\mathrm{ML},k}-x_k)^2\big\}\\
&= \sum_{k=1}^{N}\big[\mathsf{E}_{\boldsymbol{x}}\big\{\hat{x}_{\mathrm{ML},k}^2\big\}-2\,\mathsf{E}_{\boldsymbol{x}}\big\{\hat{x}_{\mathrm{ML},k}\big\}x_k+x_k^2\big]\\
&= \sum_{k=1}^{N}\big[\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}^2\}+\big(\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}\}-x_k\big)^2-\big(\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}\}\big)^2\big]\,.
\end{aligned}\tag{5.95}$$

Thus, we have to calculate the quantities $\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}\}$ and $\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}^2\}$.

We recall that $\hat{x}_{\mathrm{ML},k}(\boldsymbol{y}) = \big(\mathsf{P}_S(\boldsymbol{y})\big)_k$, where P_S is an operator that retains the S largest (in magnitude) components and zeros out all others. Let $\mathcal{L}_k$ denotes the set of vectors $\boldsymbol{y}$ for which

y_k is not among the S largest (in magnitude) components. We then have

$$\hat{x}_{\mathrm{ML},k}(\boldsymbol{y}) = \begin{cases} y_k\,, & \boldsymbol{y} \notin \mathcal{L}_k \\ 0\,, & \boldsymbol{y} \in \mathcal{L}_k\,. \end{cases}$$

Equivalently, $\hat{x}_{\mathrm{ML},k}(\boldsymbol{y}) = y_k[1 - \mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)]$, where $\mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)$ is the indicator function of the event $\{\boldsymbol{y} \in \mathcal{L}_k\}$ (i.e., $\mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)$ is 1 if $\boldsymbol{y} \in \mathcal{L}_k$ and 0 else). Thus, we obtain $\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}\}$ as

$$\begin{aligned} \mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}\} &= \mathsf{E}_{\boldsymbol{x}}\big\{y_k[1 - \mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)]\big\} \\ &= x_k - \mathsf{E}_{\boldsymbol{x}}\big\{y_k \mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)\big\} \\ &\overset{(a)}{=} x_k - \mathsf{E}_{\boldsymbol{x}}^{(y_k)}\big\{\mathsf{E}_{\boldsymbol{x}}^{(\boldsymbol{y}_{\sim k})}\big\{y_k \mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)\big|y_k\big\}\big\} \\ &\overset{(b)}{=} x_k - \mathsf{E}_{\boldsymbol{x}}^{(y_k)}\big\{y_k\, \mathsf{E}_{\boldsymbol{x}}^{(\boldsymbol{y}_{\sim k})}\big\{\mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)\big|y_k\big\}\big\} \\ &= x_k - \mathsf{E}_{\boldsymbol{x}}^{(y_k)}\big\{y_k\, \mathrm{P}_{\boldsymbol{x}}(\boldsymbol{y} \in \mathcal{L}_k|y_k)\big\} \end{aligned} \tag{5.96}$$

where the notations $\mathsf{E}_{\boldsymbol{x}}^{(y_k)}$ and $\mathsf{E}_{\boldsymbol{x}}^{(\boldsymbol{y}_{\sim k})}$ indicate that the expectation is taken with respect to the random quantities y_k and $\boldsymbol{y}_{\sim k}$, respectively (here, $\boldsymbol{y}_{\sim k}$ denotes $\boldsymbol{y}$ without the component y_k) and $\mathrm{P}_{\boldsymbol{x}}(\boldsymbol{y} \in \mathcal{L}_k|y_k)$ is the conditional probability that $\boldsymbol{y} \in \mathcal{L}_k$, given y_k. Furthermore, we used the law of total probability in (a) and the fact that y_k is held constant in the conditional expectation $\mathsf{E}_{\boldsymbol{x}}\big\{y_k \mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)\big|y_k\big\}$ in (b). Similarily,

$$\begin{aligned} \mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}^2\} &= \mathsf{E}_{\boldsymbol{x}}\big\{y_k^2\,[1 - \mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)]^2\big\} \\ &= \mathsf{E}_{\boldsymbol{x}}\big\{y_k^2\,[1 - \mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)]\big\} \\ &= \sigma^2 + x_k^2 - \mathsf{E}_{\boldsymbol{x}}\big\{y_k^2\, \mathrm{I}(\boldsymbol{y} \in \mathcal{L}_k)\big\} \\ &= \sigma^2 + x_k^2 - \mathsf{E}_{\boldsymbol{x}}^{(y_k)}\big\{y_k^2\, \mathrm{P}_{\boldsymbol{x}}(\boldsymbol{y} \in \mathcal{L}_k|y_k)\big\}\,. \end{aligned} \tag{5.97}$$

Calculating $\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}\}$ and $\mathsf{E}_{\boldsymbol{x}}\{\hat{x}_{\mathrm{ML},k}^2\}$ is thus reduced to calculating the conditional probability $\mathrm{P}_{\boldsymbol{x}}(\boldsymbol{y} \in \mathcal{L}_k|y_k)$.

Let $\mathcal{M}_k \triangleq \{1, \ldots, N\} \setminus \{k\}$, and let $\mathcal{P}$ denote the set of all binary partitions $(\mathcal{A}, \mathcal{B})$ of the set $\mathcal{M}_k$, where $\mathcal{A}$ is at least of cardinality S:

$$\mathcal{P} \triangleq \big\{(\mathcal{A}, \mathcal{B})\big|\, \mathcal{A} \subseteq \mathcal{M}_k, \mathcal{B} \subseteq \mathcal{M}_k, \mathcal{A} \cap \mathcal{B} = \emptyset, \mathcal{A} \cup \mathcal{B} = \mathcal{M}_k, |\mathcal{A}| \geq S\big\}\,.$$

In order to evaluate the conditional probability $\mathrm{P}_{\boldsymbol{x}}(\boldsymbol{y} \in \mathcal{L}_k|y_k)$ of the event $\{\boldsymbol{y} \in \mathcal{L}_k\}$, i.e., of the event that a given y_k is not among the S largest (in magnitude) components of $\boldsymbol{y}$, we split the event $\{\boldsymbol{y} \in \mathcal{L}_k\}$ into several elementary events. More specifically, let $\mathcal{E}_{\mathcal{A},\mathcal{B}}$ denote the event that every component y_l with $l \in \mathcal{A}$ satisfies $|y_l| > |y_k|$ and every component y_l with $l \in \mathcal{B}$ satisfies

$|y_l| \leq |y_k|$. The events $\mathcal{E}_{\mathcal{A},\mathcal{B}}$ for all $(\mathcal{A},\mathcal{B}) \in \mathcal{P}$ are mutually exclusive, i.e., $(\mathcal{A},\mathcal{B}) \neq (\mathcal{A}',\mathcal{B}') \Rightarrow \mathcal{E}_{\mathcal{A},\mathcal{B}} \cap \mathcal{E}_{\mathcal{A}',\mathcal{B}'} = \emptyset$, and their union corresponds to the event $\{\boldsymbol{y} \in \mathcal{L}_k\}$, i.e., $\bigcup_{(\mathcal{A},\mathcal{B})\in\mathcal{P}} \mathcal{E}_{\mathcal{A},\mathcal{B}} = \{\boldsymbol{y} \in \mathcal{L}_k\}$. Consequently,

$$
\begin{aligned}
\mathrm{P}_{\boldsymbol{x}}(\boldsymbol{y} \in \mathcal{L}_k | y_k = y) &= \sum_{(\mathcal{A},\mathcal{B})\in\mathcal{P}} \mathrm{P}_{\boldsymbol{x}}(\mathcal{E}_{\mathcal{A},\mathcal{B}} | y_k = y) \\
&= \sum_{(\mathcal{A},\mathcal{B})\in\mathcal{P}} \prod_{l\in\mathcal{A}} \mathrm{P}_{\boldsymbol{x}}(|y_l| > |y_k| | y_k = y) \prod_{m\in\mathcal{B}} \mathrm{P}_{\boldsymbol{x}}(|y_m| \leq |y_k| | y_k = y) \\
&= \sum_{(\mathcal{A},\mathcal{B})\in\mathcal{P}} \prod_{l\in\mathcal{A}} \mathrm{P}_{\boldsymbol{x}}(|y_l| > |y|) \prod_{m\in\mathcal{B}} \mathrm{P}_{\boldsymbol{x}}(|y_m| \leq |y|) \\
&= \sum_{(\mathcal{A},\mathcal{B})\in\mathcal{P}} \prod_{l\in\mathcal{A}\cap\mathrm{supp}(\boldsymbol{x})} \mathrm{P}_{\boldsymbol{x}}(|y_l| > |y|) \prod_{m\in\mathcal{B}\cap\mathrm{supp}(\boldsymbol{x})} \mathrm{P}_{\boldsymbol{x}}(|y_m| \leq |y|) \\
&\qquad \times \prod_{n\in\mathcal{A}\setminus\mathrm{supp}(\boldsymbol{x})} \mathrm{P}_{\boldsymbol{x}}(|y_n| > |y|) \prod_{p\in\mathcal{B}\setminus\mathrm{supp}(\boldsymbol{x})} \mathrm{P}_{\boldsymbol{x}}(|y_p| \leq |y|) \\
&= \sum_{(\mathcal{A},\mathcal{B})\in\mathcal{P}} \prod_{l\in\mathcal{A}\cap\mathrm{supp}(\boldsymbol{x})} \left[Q\left(\frac{|y| - x_l}{\sigma}\right) + 1 - Q\left(\frac{-|y| - x_l}{\sigma}\right) \right] \\
&\qquad \times \prod_{m\in\mathcal{B}\cap\mathrm{supp}(\boldsymbol{x})} \left[-Q\left(\frac{|y| - x_m}{\sigma}\right) + Q\left(\frac{-|y| - x_m}{\sigma}\right) \right] \\
&\qquad \times \prod_{n\in\mathcal{A}\setminus\mathrm{supp}(\boldsymbol{x})} 2Q\left(\frac{|y|}{\sigma}\right) \prod_{p\in\mathcal{B}\setminus\mathrm{supp}(\boldsymbol{x})} \left[1 - 2Q\left(\frac{|y|}{\sigma}\right) \right] \qquad (5.98)
\end{aligned}
$$

where we have used the fact that the y_l are independent and $k \notin \mathcal{M}_k$; furthermore, $Q(y) \triangleq \frac{1}{\sqrt{2\pi}} \int_y^\infty e^{-x^2/2} dx$ is the right tail probability of a standard Gaussian random variable. Plugging (5.98) into (5.96) and (5.97) and, in turn, the resulting expressions into (5.95) yields a (very complicated) expression of $\varepsilon(\boldsymbol{x}; \hat{\boldsymbol{x}}_{\mathrm{ML}})$. This expression is evaluated numerically in Section 5.5.

Chapter 6

Performance Guarantees for Sparse Estimation

This chapter is a reprint of the paper:

- Z. Ben-Haim, Y. C. Eldar, and M. Elad, "Coherence-based performance guarantees for estimating a sparse vector under random noise," *IEEE Trans. Signal Proc.*, vol. 58, no. 10, pp. 5030-5043, Oct. 2010.

6.1 Introduction

Estimation problems with sparsity constraints have attracted considerable attention in recent years because of their potential use in numerous signal processing applications, such as denoising, compression, and sampling [15]. In a typical setup, an unknown deterministic parameter $\boldsymbol{x}_0 \in \mathbb{R}^m$ is to be estimated from measurements $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w}$, where $\boldsymbol{A} \in \mathbb{R}^{n\times m}$ is a deterministic matrix and $\boldsymbol{w}$ is a noise vector. Typically, the dictionary $\boldsymbol{A}$ consists of more columns than rows (i.e., $m > n$), so that without further assumptions, $\boldsymbol{x}_0$ is unidentifiable from $\boldsymbol{b}$. The impassé is resolved by assuming that the parameter vector is sparse, i.e., that most elements of $\boldsymbol{x}_0$ are zero. Under the assumption of sparsity, several estimation approaches can be used. These include greedy algorithms, such as thresholding and orthogonal matching pursuit (OMP) [17], and ℓ_1 relaxation methods, such as the Dantzig selector [8] and basis pursuit denoising (BPDN) [19,20] (also known as the Lasso). A comparative analysis of these techniques is crucial for determining the appropriate strategy in a given situation.

There are two standard approaches to modeling the noise $\boldsymbol{w}$ in the sparse estimation prob-

lem. The first is to assume that $\boldsymbol{w}$ is deterministic and bounded [7, 12, 105]. This leads to a worst-case analysis in which an estimator must perform adequately even when the noise maximally damages the measurements. The noise in this case is thus called adversarial. By contrast, if one assumes that the noise is random, then the analysis aims to describe estimator behavior for typical noise values [8, 22, 23]. In this paper, we focus on the random noise scenario. As one might expect, stronger performance guarantees can be obtained in this setting.

It is common to judge the quality of an estimator by comparing its mean-squared error (MSE) with the Cramér–Rao bound (CRB) [16]. In the case of sparse estimation under Gaussian noise, it has recently been shown that the unbiased CRB is identical (for almost all values of $\boldsymbol{x}_0$) to the MSE of the "oracle" estimator, which knows the locations of the nonzero elements of $\boldsymbol{x}_0$ [70]. Thus, a gold standard for estimator performance is the MSE of the oracle. Indeed, it can be shown that ℓ_1 relaxation algorithms come close to the oracle when the noise is Gaussian. Results of this type are sometimes referred to as "oracle inequalities." Specifically, Candès and Tao [8] have shown that, with high probability, the ℓ_2 distance between $\boldsymbol{x}_0$ and the Dantzig estimate is within a constant times $\log m$ of the performance of the oracle. Recently, Bickel et al. [23] have demonstrated that the performance of BPDN is similarly bounded, with high probability, by $C \log m$ times the oracle performance, for a constant C. However, the constant involved in this analysis is considerably larger than that of the Dantzig selector. Interestingly, it turns out that the $\log m$ gap between the oracle and practical estimators is an unavoidable consequence of the fact that the nonzero locations in $\boldsymbol{x}_0$ are unknown [24].

The contributions [8, 23] state their results using the restricted isometry constants (RICs). These measures of the dictionary quality can be efficiently approximated in specific cases, e.g., when the dictionary is selected randomly from an appropriate ensemble. However, in general it is NP-hard to evaluate the RICs for a given matrix $\boldsymbol{A}$, and they must then be bounded by efficiently computable properties of $\boldsymbol{A}$, such as the mutual coherence [106]. In this respect, coherence-based results are appealing since they can be used with arbitrary dictionaries [28, 107, 108].

In this paper, we seek performance guarantees for sparse estimators based directly on the mutual coherence of the matrix $\boldsymbol{A}$ [109]. While such results are suboptimal when the RICs of $\boldsymbol{A}$ are known, the proposed approach yields tighter bounds than those obtained by applying coherence bounds to RIC-based results. Specifically, we demonstrate that BPDN, OMP and thresholding all achieve performance within a constant times $\log m$ of the oracle estimator, under suitable conditions. In the case of BPDN, our result provides a tighter guarantee than the

coherence-based implications of the work of Bickel et al. [23]. To the best of our knowledge, there are no prior performance guarantees for greedy approaches such as OMP and thresholding when the noise is random.

It is important to distinguish the present work from Bayesian performance analysis, as practiced in [4, 5, 21, 110, 111], where on top of the assumption of stochastic noise, a probabilistic model for $\boldsymbol{x}_0$ is also used. Our results hold for any specific value of $\boldsymbol{x}_0$ (satisfying appropriate conditions), rather than providing results on average over realizations of $\boldsymbol{x}_0$; this necessarily leads to weaker guarantees. It also bears repeating that our results apply to a fixed, finite-sized matrix $\boldsymbol{A}$; this distinguishes our work from asymptotic performance guarantees for large m and n, such as [112, 113].

The rest of this paper is organized as follows. We begin in Section 6.2 by comparing dictionary quality measures and reviewing standard estimation techniques. In Section 6.3, our main results are presented in the form of performance guarantees for various sparse estimation algorithms. Subsequently, in Section 6.4, these results are compared with previous findings for different sparse models, including the case of adversarial noise and the random design matrix setting. Finally, in Section 6.5, the validity of our results is examined by simulation in practical estimation scenarios.

The following notation is used throughout the paper. Vectors and matrices are denoted, respectively, by boldface lowercase and boldface uppercase letters. The set of indices of the nonzero entries of a vector $\boldsymbol{x}$ is called the support of $\boldsymbol{x}$ and denoted $\operatorname{supp}(\boldsymbol{x})$. Given an index set Λ and a matrix $\boldsymbol{A}$, the notation $\boldsymbol{A}_\Lambda$ refers to the submatrix formed from the columns of $\boldsymbol{A}$ indexed by Λ. The identity matrix is denoted by $\boldsymbol{I}$. Furthermore, for any matrix $\boldsymbol{A}$, $\boldsymbol{A}^T$ refers to the matrix transpose, $\boldsymbol{A}^\dagger$ is the Moore–Penrose pseudoinverse, $\operatorname{Tr}(\boldsymbol{A})$ is the trace, and $\operatorname{span}(\boldsymbol{A})$ is the column span, while $\lambda_{\min}(\boldsymbol{A})$ and $\lambda_{\max}(\boldsymbol{A})$ denote the smallest and largest eigenvalues of $\boldsymbol{A}$. Finally, the ℓ_p norm of a vector $\boldsymbol{x}$, for $1 \leq p \leq \infty$, is denoted $\|\boldsymbol{x}\|_p$, while $\|\boldsymbol{x}\|_0$ denotes the number of nonzero elements in $\boldsymbol{x}$.

6.2 Preliminaries

6.2.1 Characterizing the Dictionary

Let $\boldsymbol{x}_0 \in \mathbb{R}^m$ be an unknown deterministic vector, and denote its support set by $\Lambda_0 = \operatorname{supp}(\boldsymbol{x}_0)$. Let $s = \|\boldsymbol{x}_0\|_0$ be the number of nonzero entries in $\boldsymbol{x}_0$. In our setting, it is typically assumed that s is much smaller than m, i.e., that most elements in $\boldsymbol{x}_0$ are zero. Suppose we obtain noisy

measurements

$$\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w} \tag{6.1}$$

where $\boldsymbol{A} \in \mathbb{R}^{n\times m}$ is a known deterministic overcomplete dictionary ($m > n$). We refer to the columns $\boldsymbol{a}_i$ of $\boldsymbol{A}$ as the *atoms* of the dictionary, and assume throughout our work that the atoms are normalized such that

$$\|\boldsymbol{a}_i\|_2 = 1. \tag{6.2}$$

We further assume that $\boldsymbol{w}$ is zero-mean white Gaussian noise with covariance $E\{\boldsymbol{w}\boldsymbol{w}^T\} = \sigma^2\boldsymbol{I}_{m\times m}$.

For $\boldsymbol{x}_0$ to be identifiable, one must guarantee that different values of $\boldsymbol{x}_0$ produce significantly different values of $\boldsymbol{b}$. One way to ensure this is to examine all possible *subdictionaries*, or s-element sets of atoms, and verify that the subspaces spanned by these subdictionaries differ substantially from one another.

More specifically, several methods have been proposed to formalize the notion of the suitability of a dictionary for sparse estimation. These include the mutual coherence [106], the cumulative coherence [22], the exact recovery coefficient (ERC) [22], the spark [7], and the RICs [8, 12]. Except for the mutual coherence and cumulative coherence, none of these measures can be efficiently calculated for an arbitrary given dictionary $\boldsymbol{A}$. Since the values of the cumulative and mutual coherence are quite close, our focus in this paper will be on the mutual coherence $\mu = \mu(\boldsymbol{A})$, which is defined as

$$\mu \triangleq \max_{i\neq j} \left|\boldsymbol{a}_i^T\boldsymbol{a}_j\right|. \tag{6.3}$$

While the mutual coherence can be efficiently calculated directly from (6.3), it is not immediately clear in what way μ is related to the requirement that subdictionaries must span different subspaces. Indeed, μ ensures a lack of correlation between single atoms, while we require a distinction between s-element subdictionaries. To explore this relation, let us recall the definitions of the RICs, which are more directly related to the subdictionaries of $\boldsymbol{A}$. We will then show that the mutual coherence can be used to bound the constants involved in the RICs, a fact which will also prove useful in our subsequent analysis. This strategy is inspired by earlier works, which have used the mutual coherence to bound the ERC [22] and the spark [5, 7, 28]. Thus, the coherence can be viewed as a tractable proxy for more accurate measures of the quality of a dictionary, which cannot themselves be calculated efficiently.

By the RICs we refer to two properties describing "good" dictionaries, namely, the restricted isometry property (RIP) and the restricted orthogonality property (ROP), which we now define.

A dictionary $\boldsymbol{A}$ is said to satisfy the RIP [12] of order s with parameter δ_s if, for every index set Λ of size s, we have

$$(1-\delta_s)\|\boldsymbol{y}\|_2^2 \leq \|\boldsymbol{A}_\Lambda \boldsymbol{y}\|_2^2 \leq (1+\delta_s)\|\boldsymbol{y}\|_2^2 \tag{6.4}$$

for all $\boldsymbol{y} \in \mathbb{R}^s$. Thus, when δ_s is small, the RIP ensures that any s-atom subdictionary is nearly orthogonal, which in turn implies that any two disjoint $(s/2)$-atom subdictionaries are well-separated.

Similarly, $\boldsymbol{A}$ is said to satisfy the ROP [8] of order (s_1, s_2) with parameter θ_{s_1,s_2} if, for every pair of disjoint index sets Λ_1 and Λ_2 having cardinalities s_1 and s_2, respectively, we have

$$\left|\boldsymbol{y}_1^T \boldsymbol{A}_{\Lambda_1}^T \boldsymbol{A}_{\Lambda_2} \boldsymbol{y}_2\right| \leq \theta_{s_1,s_2}\|\boldsymbol{y}_1\|_2\|\boldsymbol{y}_2\|_2 \tag{6.5}$$

for all $\boldsymbol{y}_1 \in \mathbb{R}^{s_1}$ and for all $\boldsymbol{y}_2 \in \mathbb{R}^{s_2}$. In words, the ROP requires any two disjoint subdictionaries containing s_1 and s_2 elements, respectively, to be nearly orthogonal to each other. These two properties are therefore closely related to the requirement that distinct subdictionaries of $\boldsymbol{A}$ behave dissimilarly.

In recent years, it has been demonstrated that various practical estimation techniques successfully approximate $\boldsymbol{x}_0$ from $\boldsymbol{b}$, if the constants δ_s and θ_{s_1,s_2} are sufficiently small [8, 9, 12, 13]. This occurs, for example, when the entries in $\boldsymbol{A}$ are chosen randomly according to an independent, identically distributed Gaussian law, as well as in some specific deterministic dictionary constructions.

Unfortunately, in the standard estimation setting, one cannot design the system matrix $\boldsymbol{A}$ according to these specific rules. In general, if one is given a particular dictionary $\boldsymbol{A}$, then there is no known algorithm for efficiently determining its RICs. Indeed, the very nature of the RICs seems to require enumerating over an exponential number of index sets in order to find the "worst" subdictionary. While the mutual coherence μ of (6.3) tends to be far less accurate in capturing the accuracy of a dictionary, it is still useful to be able to say something about the RICs based only on μ. Such a result is given in the following lemma, whose proof can be found in [114, 115].

Lemma 6.1 (Cai, Xu, and Zhang)**.** *For any matrix $\boldsymbol{A}$, the RIP constant δ_s of* (6.4) *and the ROP constant θ_{s_1,s_2} of* (6.5) *satisfy the bounds*

$$\delta_s \leq (s-1)\mu, \tag{6.6}$$

$$\theta_{s_1,s_2} \leq \mu\sqrt{s_1 s_2} \tag{6.7}$$

where μ is the mutual coherence (6.3)*.*

We will apply this lemma in Section 6.3 in order to obtain coherence-based settings from results which use the RICs.

6.2.2 Estimation Techniques

To fix notation, we now briefly review several approaches for estimating $\boldsymbol{x}_0$ from noisy measurements $\boldsymbol{b}$ given by (6.1). The two main strategies for efficiently estimating a sparse vector are ℓ_1 relaxation and greedy methods. The first of these involves solving an optimization problem wherein the nonconvex constraint $\|\boldsymbol{x}_0\|_0 = s$ is relaxed to a constraint on the ℓ_1 norm of the estimated vector $\boldsymbol{x}_0$. Specifically, we consider the ℓ_1-penalty version of BPDN, which estimates $\boldsymbol{x}_0$ as a solution $\hat{\boldsymbol{x}}_{\text{BP}}$ to the quadratic program

$$\min_{\boldsymbol{x}} \tfrac{1}{2}\|\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}\|_2^2 + \gamma\|\boldsymbol{x}\|_1 \tag{6.8}$$

for some regularization parameter γ. We refer to the optimization problem (6.8) as BPDN, although it should be noted that some authors reserve this term for the related optimization problem

$$\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_1 \quad \text{s.t. } \|\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}\|_2^2 \leq \nu \tag{6.9}$$

where ν is a given constant.

Another estimator based on the idea of ℓ_1 relaxation is the Dantzig selector [8], defined as a solution $\hat{\boldsymbol{x}}_{\text{DS}}$ to the optimization problem

$$\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_1 \quad \text{s.t. } \|\boldsymbol{A}^T(\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x})\|_\infty \leq \tau \tag{6.10}$$

where τ is again a user-selected parameter. The Dantzig selector, like BPDN, is a convex relaxation method, but rather than penalizing the ℓ_2 norm of the residual $\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}$, the Dantzig selector ensures that the residual is weakly correlated with all dictionary atoms.

Instead of solving an optimization problem, greedy approaches estimate the support set Λ_0 from the measurements $\boldsymbol{b}$. Once a support set Λ is chosen, the parameter vector $\boldsymbol{x}_0$ can be estimated using least-squares (LS) to obtain

$$\hat{\boldsymbol{x}} = \begin{cases} \boldsymbol{A}_\Lambda^\dagger \boldsymbol{b} & \text{on the support set } \Lambda, \\ \boldsymbol{0} & \text{elsewhere.} \end{cases} \tag{6.11}$$

Greedy techniques differ in the method by which the support set is selected. The simplest method is known as the thresholding algorithm. This technique computes the correlation of

the measured signal $\boldsymbol{b}$ with each of the atoms $\boldsymbol{a}_i$ and defines Λ as the set of indices of the s atoms having the highest correlation. Subsequently, the LS technique (6.11) is applied to obtain the thresholding estimate $\hat{\boldsymbol{x}}_{\text{th}}$.

A somewhat more sophisticated greedy algorithm is OMP [17]. This iterative approach begins by initializing the estimated support set Λ^0 to the empty set and setting a residual vector $\boldsymbol{r}^0$ to $\boldsymbol{b}$. Subsequently, at each iteration $i = 1, \ldots, s$, the algorithm finds the single atom which is most highly correlated with $\boldsymbol{r}^{i-1}$. The index of this atom, say k_i, is added to the support set, so that $\Lambda^i = \Lambda^{i-1} \cup \{k_i\}$. The estimate $\hat{\boldsymbol{x}}^i_{\text{OMP}}$ at the ith iteration is then defined by the LS solution (6.11) using the support set Λ^i. Next, the residual is updated using the formula

$$\boldsymbol{r}^i = \boldsymbol{b} - \boldsymbol{A}\hat{\boldsymbol{x}}^i_{\text{OMP}}. \tag{6.12}$$

The residual thus describes the part of $\boldsymbol{b}$ which has yet to be accounted for by the estimate. The counter i is now incremented, and s iterations are performed, after which the OMP estimate $\hat{\boldsymbol{x}}_{\text{OMP}}$ is defined as the estimate at the final iteration, $\hat{\boldsymbol{x}}^s_{\text{OMP}}$. A well-known property of OMP is that the algorithm never chooses the same atom twice [7]. Consequently, stopping after s iterations guarantees that $\|\hat{\boldsymbol{x}}_{\text{OMP}}\|_0 = s$.

Finally, we also mention the so-called oracle estimator, which is based both on $\boldsymbol{b}$ and on the true support set Λ_0 of $\boldsymbol{x}_0$; the support set is assumed to have been provided by an "oracle". The oracle estimator $\hat{\boldsymbol{x}}_{\text{or}}$ calculates the LS solution (6.11) for Λ_0, and is often used as a gold standard against which the performance of practical algorithms can be compared.

6.3 Performance Guarantees

Under the setting (6.1), it cat be shown [70] that the MSE of any unbiased estimator $\hat{\boldsymbol{x}}$ of $\boldsymbol{x}_0$ satisfies the Cramér–Rao bound [44]

$$\text{MSE}(\hat{\boldsymbol{x}}) \triangleq E\{\|\hat{\boldsymbol{x}} - \boldsymbol{x}_0\|_2^2\} \geq \sigma^2 \,\text{Tr}((\boldsymbol{A}_{\Lambda_0}^T \boldsymbol{A}_{\Lambda_0})^{-1}) \triangleq \text{CRB}. \tag{6.13}$$

Interestingly, CRB is also the MSE of the oracle estimator mentioned above [8]. It follows from the Gershgorin disc theorem [116] that all eigenvalues of $\boldsymbol{A}_{\Lambda_0}^T \boldsymbol{A}_{\Lambda_0}$ are between $1 - (s-1)\mu$ and $1 + (s+1)\mu$. Therefore, for reasonable sparsity levels, $\text{Tr}((\boldsymbol{A}_{\Lambda_0}^T \boldsymbol{A}_{\Lambda_0})^{-1})$ is not much larger than s, and consequently (6.13) is on the order of $s\sigma^2$. Considering that the mean power of $\boldsymbol{w}$ is $n\sigma^2$, it is evident that the oracle estimator has substantially reduced the noise level. In this section, we will demonstrate that comparable performance gains are achievable using practical methods, which do not have access to the oracle.

6.3.1 ℓ_1-Relaxation Approaches

Review of previous results Historically, performance guarantees under random noise were first obtained for the Dantzig selector (6.10) by Candès and Tao [8]. Their result is derived using the RICs (6.4)–(6.5); by applying the bounds of Lemma 6.1, it is possible to obtain from their work a coherence-based guarantee. More recently, Cai, Wang, and Xu [107] have shown that one can obtain a tighter performance guarantee for the Dantzig selector directly from the mutual coherence, as shown in the following theorem.

Theorem 6.2 (Cai, Wang, and Xu)**.** *Under the setting* (6.1)*, assume that*

$$s < \frac{1}{2}\left(1 + \frac{1}{\mu}\right) \tag{6.14}$$

and consider the Dantzig selector (6.10) *with parameter*

$$\tau = \sigma\sqrt{2\log m} + \tfrac{3}{2}\sigma. \tag{6.15}$$

Then, with probability exceeding

$$1 - \frac{1}{\sqrt{\pi \log m}}, \tag{6.16}$$

the Dantzig selector $\hat{\boldsymbol{x}}_{\mathrm{DS}}$ satisfies

$$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{DS}}\|_2^2 \leq \frac{8\left(\frac{3}{2} + \sqrt{2\log m}\right)^2}{\left(1 - (2s-1)\mu\right)^2}(s+1)\sigma^2. \tag{6.17}$$

Thus, while $\hat{\boldsymbol{x}}_{\mathrm{DS}}$ does not quite reach the performance of the oracle estimator, it does come within a constant factor multiplied by $\log m$, with high probability. Interestingly, the $\log m$ factor is an unavoidable result of the fact that the locations of the nonzero elements in $\boldsymbol{x}_0$ are unknown (see [24, §7.4] and the references therein).

Recently, a performance guarantee has also been demonstrated for BPDN [23]. Once again, this result is based on RIC-like properties. Its translation to a coherence-based guarantee is given in the following theorem, whose derivation from [23] is described in Appendix 6.A.

Theorem 6.3 (Bickel, Ritov, and Tsybakov)**.** *Under the setting* (6.1)*, assume that*

$$s < \frac{1+\mu}{(2+3\sqrt{2})\mu} \tag{6.18}$$

and consider the BPDN estimator (6.8) *with parameter*

$$\gamma = \sqrt{8\sigma^2(1+\alpha)\log m} \tag{6.19}$$

for some $\alpha > 0$. Then, with probability exceeding

$$1 - \frac{1}{m^\alpha}, \tag{6.20}$$

BPDN satisfies

$$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{BP}}\|_2^2 \leq 2048(1+\alpha)\frac{1}{\kappa^4} s\sigma^2 \log m \tag{6.21}$$

where

$$\kappa = \sqrt{1-(2s-1)\mu} - \frac{s\mu 3\sqrt{2}}{\sqrt{1-(2s-1)\mu}} < 1. \tag{6.22}$$

Coherence-based guarantee for BPDN The constant in the BPDN performance guarantee (6.21) is typically much larger than that given in (6.17) for the Dantzig selector. The necessary condition (6.18) is also more stringent than the requirements for the Dantzig selector. However, following the experience in the case of the Dantzig selector, one may hope for substantially better guarantees to be obtained by directly relying on the mutual coherence. This approach indeed bears fruit, as we now show. We begin by stating the following somewhat more general result, whose proof is found in Appendix 6.B.

Theorem 6.4. *Let $\boldsymbol{x}_0$ be an unknown deterministic vector with known sparsity $\|\boldsymbol{x}_0\|_0 = s$, and let $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w}$, where $\boldsymbol{w} \sim N(\boldsymbol{0}, \sigma^2 \boldsymbol{I})$ is a random noise vector. Suppose that*[1]

$$s < \frac{1}{3\mu}. \tag{6.23}$$

Then, with probability exceeding

$$\left(1 - (m-s)\exp\left(-\frac{\gamma^2}{8\sigma^2}\right)\right)\left(1 - e^{-s/7}\right), \tag{6.24}$$

the solution $\hat{\boldsymbol{x}}_{\mathrm{BP}}$ of BPDN (6.8) *is unique, its support is contained in the true support Λ_0, and*

$$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{BP}}\|_2^2 \leq \left(\sigma\sqrt{3} + \tfrac{3}{2}\gamma\right)^2 s. \tag{6.25}$$

To compare this result with the previous theorems, we now derive from Theorem 6.4 a result which holds with a probability on the order of (6.20). Observe that for (6.24) to be a high probability, we require $\exp(-\gamma^2/(8\sigma^2))$ to be substantially smaller than $1/(m-s)$. This requirement can be used to select a value for the regularization parameter γ. In particular, one

[1] As in [22], analogous findings can also be obtained under the weaker requirement $s < 1/(2\mu)$, but the resulting expressions are somewhat more involved.

requires γ to be at least on the order of $\sqrt{8\sigma^2 \log(m-s)}$. However, γ should not be much larger than this value, as this will increase the error bound (6.25). Therefore, it is natural to use

$$\gamma = \sqrt{8\sigma^2(1+\alpha)\log(m-s)} \tag{6.26}$$

for some fairly small $\alpha > 0$; note the encouraging similarity of this value of γ to (6.19). Substituting γ of (6.26) into Theorem 6.4 yields the following result.

Corollary 6.5. *Under the conditions of Theorem 6.4, let $\hat{\boldsymbol{x}}_{\mathrm{BP}}$ be a solution of BPDN* (6.8) *with γ given by* (6.26)*. Then, with probability exceeding*

$$\left(1 - \frac{1}{(m-s)^{\alpha}}\right)\left(1 - e^{-s/7}\right), \tag{6.27}$$

the solution $\hat{\boldsymbol{x}}_{\mathrm{BP}}$ is unique, its support is a subset of Λ_0, and

$$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{BP}}\|_2^2 \leq \left(\sqrt{3} + 3\sqrt{2(1+\alpha)\log(m-s)}\right)^2 s\sigma^2. \tag{6.28}$$

To compare the BPDN guarantees of Corollary 6.5 and Theorem 6.3, we first examine the probability (6.27). This expression consists of a product of two terms, both of which converge to 1 as the problem dimensions increase. The right-hand term may seem odd because it appears to favor non-sparse signals; however, this is an artifact of the method of proof, which requires a sufficient number of nonzero coefficients for large number approximations to hold. This right-hand term converges to 1 exponentially and therefore typically has a negligible effect on the overall probability of success; for example, for $s \geq 50$ this term is larger than 0.999.

The left-hand term in (6.27) tends to 1 polynomially as $m - s$ increases. This is a slightly lower rate than that of Theorem 6.3; however, this difference is compensated for by a correspondingly lower multiplicative factor of $\log(m-s)$ in the error bound (6.28), as opposed to the $\log m$ factor in previous results. In any case, for any of the theorems to hold, m must increase much more quickly than s, so that these differences are negligible.

Whereas the probability of success of Corollary 6.5 is comparable to that of Theorem 6.3, the required sparsity level (6.23) in the proposed guarantee is substantially better than that of previous results. Furthermore, the constant (6.25) in the new result is also much smaller than that of the previous BPDN guarantee (6.21). Thus, it appears that the direct application of the mutual coherence is successful in obtaining tighter performance guarantees.

What can be learned by comparing the guarantees for BPDN and the Dantzig selector? In some respects, the BPDN result appears stronger; in particular, the probability of success in Corollary 6.5 is better than that of Theorem 6.2, and the resulting guarantee has somewhat

smaller constants. On the other hand, the sparsity requirements of Theorem 6.2 are somewhat less stringent. Choosing the most accurate guarantee (and, consequently, the preferred estimator) will thus depend on the specifics of the setting under consideration. One such specific example will be presented in Section 6.5.

6.3.2 Greedy Approaches

The performance guarantees obtained for the ℓ_1-relaxation techniques required only the assumption that $\boldsymbol{x}_0$ is sufficiently sparse. By contrast, for greedy algorithms, successful estimation can only be guaranteed if one further assumes that all nonzero components of $\boldsymbol{x}_0$ are somewhat larger than the noise level. The reason is that greedy techniques are based on a LS solution for an estimated support, an approach whose efficacy is poor unless the support is correctly identified. Indeed, when using the LS technique (6.11), even a single incorrectly identified support element may cause the entire estimate to be severely incorrect. To ensure support recovery, all nonzero elements must be large enough to overcome the noise.

To formalize this notion, denote $\boldsymbol{x}_0 = (x_{0,1}, \ldots, x_{0,m})^T$ and define

$$\begin{aligned} |x_{\min}| &= \min_{i \in \Lambda_0} |x_{0,i}|, \\ |x_{\max}| &= \max_{i \in \Lambda_0} |x_{0,i}|. \end{aligned} \tag{6.29}$$

A performance guarantee for both OMP and the thresholding algorithm is then given by the following theorem.

Theorem 6.6. *Let $\boldsymbol{x}_0$ be an unknown deterministic vector with known sparsity $\|\boldsymbol{x}_0\|_0 = s$, and let $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w}$, where $\boldsymbol{w} \sim N(\boldsymbol{0}, \sigma^2 \boldsymbol{I})$ is a random noise vector. Suppose that*

$$|x_{\min}| - (2s-1)\mu |x_{\min}| \geq 2\sigma \sqrt{2(1+\alpha)\log m} \tag{6.30}$$

for some constant $\alpha > 0$. Then, with probability at least

$$1 - \frac{1}{m^{\alpha} \sqrt{\pi(1+\alpha)\log m}}, \tag{6.31}$$

the OMP estimate $\hat{\boldsymbol{x}}_{\mathrm{OMP}}$ identifies the correct support Λ_0 of $\boldsymbol{x}_0$ and, furthermore, satisfies

$$\|\hat{\boldsymbol{x}}_{\mathrm{OMP}} - \boldsymbol{x}_0\|_2^2 \leq \frac{2(1+\alpha)}{(1-(s-1)\mu)^2} s\sigma^2 \log m \tag{6.32a}$$

$$\leq 8(1+\alpha) s\sigma^2 \log m. \tag{6.32b}$$

If the stronger condition

$$|x_{\min}| - (2s-1)\mu|x_{\max}| \geq 2\sigma\sqrt{2(1+\alpha)\log m} \tag{6.33}$$

holds, then with probability exceeding (6.31)*, the thresholding algorithm also correctly identifies* Λ_0 *and satisfies* (6.32)*.*

The performance guarantee (6.32) is better than that provided by Theorem 6.2 and Corollary 6.5. However, this result comes at the expense of requirements on the magnitude of the entries of $\boldsymbol{x}_0$. Our analysis thus suggests that greedy approaches may outperform ℓ_1-based methods when the entries of $\boldsymbol{x}_0$ are large compared with the noise, but that the greedy approaches will deteriorate when the noise level increases. As we will see in Section 6.5, simulations also appear to support this conclusion.

It is interesting to compare the success conditions (6.30) and (6.33) of the OMP and thresholding algorithms. For given problem dimensions, the OMP algorithm requires $|x_{\min}|$, the smallest nonzero element of $\boldsymbol{x}_0$, to be larger than a constant multiple of the noise standard deviation σ. This is required in order to ensure that all elements of the support of $\boldsymbol{x}_0$ will be identified with high probability. The requirement of the thresholding algorithm is stronger, as befits a simpler approach: In this case $|x_{\min}|$ must be larger than the noise standard deviation plus a constant times $|x_{\max}|$. In other words, one must be able to separate $|x_{\min}|$ from the combined effect of noise and interference caused by the other nonzero components of $\boldsymbol{x}_0$. This results from the thresholding technique, in which the entire support is identified simultaneously from the measurements. By comparison, the iterative approach used by OMP identifies and removes the large elements in $\boldsymbol{x}_0$ first, thus facilitating the identification of the smaller elements in later iterations.

6.4 Comparison with Related Estimation Settings

The difficulty of an estimation problem naturally depends on the strength of the assumptions in the underlying model. We now compare the results of the previous section with prior performance guarantees for sparse estimation algorithms in different settings.

The equation $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w}$ describes a variety of situations, depending on what is assumed to be known about $\boldsymbol{A}$, $\boldsymbol{x}_0$, and $\boldsymbol{w}$. The focus in this paper is on the frequentist estimation setting, wherein the known matrix $\boldsymbol{A}$ and the unknown parameter vector $\boldsymbol{x}_0$ are both deterministic, and the noise $\boldsymbol{w}$ is Gaussian and white. In the following, we compare the performance guarantees

obtained for this setting with known results for other scenarios. This comparison emphasizes that better performance can be guaranteed if a more detailed model is assumed. We note that many of the results cited in this section are obtained up to an unknown numerical constant, and consequently the discussion herein likewise performs qualitative comparisons based on asymptotic rates.

Adversarial noise In the adversarial or bounded noise scenario, the noise $\boldsymbol{w}$ is assumed to be deterministic and unknown, and to have a bounded ℓ_2 norm, $\|\boldsymbol{w}\|_2 \leq \varepsilon$. The matrix $\boldsymbol{A}$ and the unknown vector $\boldsymbol{x}_0$ are also deterministic. Since no probabilistic information is assumed about $\boldsymbol{w}$, one must perform a worst-case analysis, namely, in this case guarantees for an estimator $\hat{\boldsymbol{x}}$ ensure that $\|\hat{\boldsymbol{x}} - \boldsymbol{x}_0\|$ is small for all noise realizations and all feasible values of $\boldsymbol{x}_0$. Consequently, adversarial performance guarantees are relatively weak, and indeed no denoising capability can be ensured for any known algorithm.

Typical "stability" results under adversarial noise guarantee that if the mutual coherence μ of $\boldsymbol{A}$ is sufficiently small, and if $\boldsymbol{x}_0$ is sufficiently sparse, then the distance between $\boldsymbol{x}_0$ and its estimate is on the order of the noise magnitude. Such results can be derived for algorithms including BPDN, OMP, and thresholding. Consider, for example, the following theorem, which is based on the work of Tropp [22, §IV-C].[2]

Theorem 6.7 (Tropp)**.** *Let $\boldsymbol{x}_0$ be an unknown deterministic vector with known sparsity $\|\boldsymbol{x}_0\|_0 = s$, and let $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w}$, where $\|\boldsymbol{w}\|_2 \leq \varepsilon$. Suppose the mutual coherence μ of the dictionary $\boldsymbol{A}$ satisfies $s < 1/(3\mu)$. Let $\hat{\boldsymbol{x}}_{\mathrm{BP}}$ denote a solution of BPDN* (6.8) *with regularization parameter $\gamma = 2\varepsilon$. Then, $\hat{\boldsymbol{x}}_{\mathrm{BP}}$ is unique, the support of $\hat{\boldsymbol{x}}_{\mathrm{BP}}$ is a subset of the support of $\boldsymbol{x}_0$, and*

$$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{BP}}\|_\infty < \left(3 + \sqrt{\tfrac{3}{2}}\right)\varepsilon \approx 4.22\varepsilon. \tag{6.34}$$

Results similar to Theorem 6.7 have also been obtained [7,12,13,107] for the related ℓ_1-error estimation approach (6.9), as well as for the OMP algorithm [7]. Furthermore, the technique used in the proof for the OMP [7] can also be applied to demonstrate a (slightly weaker) performance guarantee for the thresholding algorithm.

In all of the aforementioned results, the only guarantee is that the distance between $\hat{\boldsymbol{x}}_{\mathrm{BP}}$ and $\boldsymbol{x}_0$ is on the order of the noise power ε. As expected, these results are much weaker than those

[2]Tropp considers only the case in which the entries of $\boldsymbol{x}_0$ belong to the set $\{0, \pm 1\}$. However, since the analysis performed in [22, §IV-C] can readily be applied to the general setting considered here, we omit the proof of Theorem 6.7.

obtained in Section 6.3 for the frequentist setting. Indeed, in the frequentist setting, which differs only in the adoption of the Gaussian distribution for the noise, typical performance guarantees ensure that $\|\hat{\boldsymbol{x}} - \boldsymbol{x}_0\|_2$ is on the order of $s\sigma^2 \log m$, while the average noise power is $n\sigma^2$, which is much larger. This difference is due to the fact that results in the adversarial context must take into account values of $\boldsymbol{w}$ which are chosen so as to cause maximal damage to the estimation algorithm.

Random design matrix Whereas the adversarial setting assumes less about the signal and consequently provides weak guarantees, incorporating further probabilistic assumptions can provide much more optimistic assurances of the achievable performance than those obtained in our setting. For example, a common setting in the compressed sensing literature assumes that the matrix $\boldsymbol{A}$ is known, but chosen randomly from an appropriate ensemble. The primary advantage of such an assumption is that with high probability, the RICs of the resulting matrix will be low, and performance guarantees obtained based on the use of the RICs are typically sharper than those based on the mutual coherence.

Specifically, suppose that the entries of $\boldsymbol{A}$ are chosen from a white Gaussian distribution and are then normalized to satisfy (6.2). Also suppose that $\boldsymbol{x}_0$ is deterministic and that $\boldsymbol{w}$ is white Gaussian noise, as in our setting. Then it can be shown that, with high probability, the Dantzig selector achieves an error on the order of $s\sigma^2 \log m$ as long as the sparsity level s is on the order of $n/\log(m/n)$ [8]. Comparable results can also be demonstrated for OMP,[3] and in this case near-oracle performance is asymptotically guaranteed with high probability when s is no larger than about $n/\log n$ [113, 117].

The guarantees of Section 6.3 are much weaker than these random-matrix results. Indeed, to obtain near-oracle performance guarantees, the results of Section 6.3 all require the sparsity level s to be on the order of $1/\mu$. However, for any matrix $\boldsymbol{A}$ we have [118, Thm. 2.3]

$$\mu \geq \sqrt{\frac{m-n}{n(m-1)}}. \tag{6.35}$$

Making the reasonable assumption that $m \geq 2n$, it follows that, at best, our frequentist results hold with sparsity levels s on the order of $\sqrt{n}$. This is not nearly as strong as the random matrix results, for which the number of nonzero entries in $\boldsymbol{x}_0$ is allowed to come within a log factor of the number of measurements n.

[3] As in any analysis of a greedy algorithm, the OMP result requires some further assumptions on the SNR and, in particular, on the value of $|x_{\min}|$ relative to the noise power.

The difference between the random and deterministic matrix guarantees should not, however, be interpreted to mean that the proposed performance guarantees can be substantially improved. Rather, it indicates that truly better performance should be expected in the random matrix setting. Indeed, there exist suitably chosen deterministic matrices $\boldsymbol{A}$ and vectors $\boldsymbol{x}_0$ for which s is on the order of $1/\mu$, wherein reasonable algorithms completely fail to recover $\boldsymbol{x}_0$ [21]. In these constructions, the value of s fails to meet the requirements for the theorems of Section 6.3 by no more than a small multiplicative constant. It follows that any performance guarantee based solely on the mutual coherence of $\boldsymbol{A}$ cannot guarantee performance for sparsity levels higher than $O(1/\mu)$.

6.5 Numerical Results

In this section, we describe a number of numerical experiments comparing the performance of various estimators to the guarantees of Section 6.3. Our first experiment measured the median estimation error, i.e., the median of the squared ℓ_2 distance between $\boldsymbol{x}_0$ and its estimate. The median error is intuitively appealing as it characterizes the "typical" estimation error, and it can be readily bounded by the performance guarantees of Section 6.3.

Specifically, we chose the two-ortho dictionary $\boldsymbol{A} = [\boldsymbol{I} \;\; \boldsymbol{H}]$, where $\boldsymbol{I}$ is the 512×512 identity matrix and $\boldsymbol{H}$ is the 512×512 Hadamard matrix with normalized columns. The RICs of this dictionary are unknown, but the coherence can be readily calculated and is given by $\mu = 1/\sqrt{512}$. Consequently, the theorems of Section 6.3 can be used to obtain performance guarantees for sufficiently sparse vectors. In particular, in our simulations we chose parameters $\boldsymbol{x}_0$ having a support of size $s = 7$. The smallest nonzero entry in $\boldsymbol{x}_0$ was $|x_{\min}| = 0.1$ and the largest entry was $|x_{\max}| = 1$. To obtain guarantees on the median error, for each of the theorems of Section 6.3 a value of α was chosen such that the resulting error bound holds with probability $1/2$ or greater.[4] Under these conditions, applying the theorems of Section 6.3 yields the bounds

$$\begin{aligned} \|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{OMP}}\|_2^2 &\leq 3.7 s\sigma^2 \log m && \text{w.p. } \tfrac{3}{4}, \text{ if } \sigma \leq 0.057; \\ \|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{BP}}\|_2^2 &\leq 22.1 s\sigma^2 \log m && \text{w.p. } \tfrac{1}{2}; \\ \|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{DS}}\|_2^2 &\leq 198.4 s\sigma^2 \log m && \text{w.p. } \tfrac{3}{4}. \end{aligned} \tag{6.36}$$

[4] In particular, the results for the Dantzig selector (Theorem 6.2) and OMP (Theorem 6.6) can only be used to yield guarantees holding with probabilities of approximately $3/4$ and higher. These are, of course, also bounds on the median error.

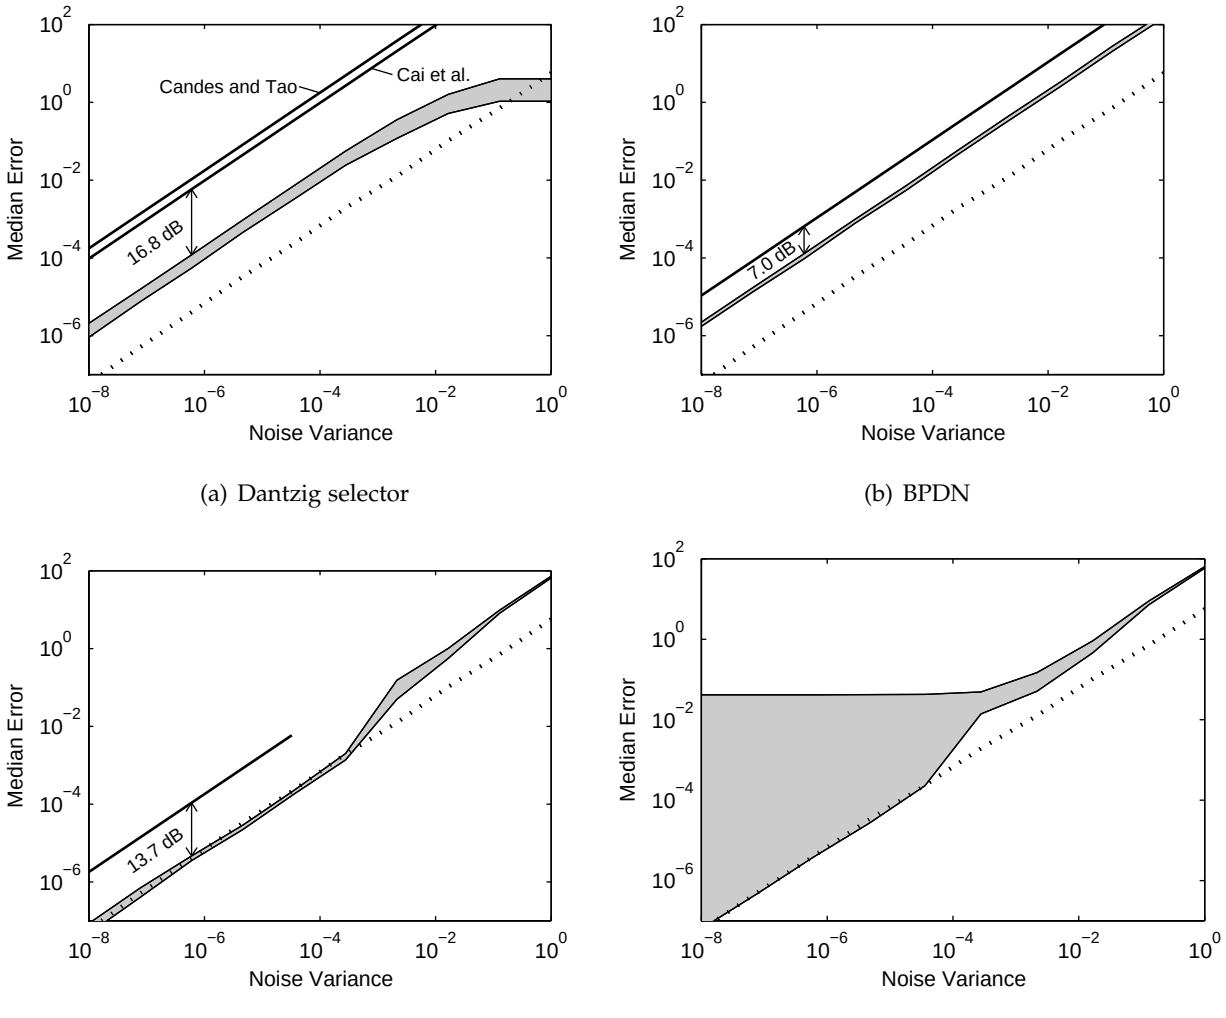

(a) Dantzig selector

(b) BPDN

(c) OMP

(d) Thresholding

Figure 6.1: Median estimation error for practical estimators (gray regions) compared with performance guarantees (solid lines) and the oracle estimator (dotted line). The gray regions report the range of performances observed for 8 different values of the unknown parameter vector $\boldsymbol{x}_0$. For the Dantzig selector, both the guarantee of Candes and Tao and that of Cai, Wang, and Xu are plotted. For OMP, performance is only guaranteed for $\sigma \leq 0.057$, while for thresholding, nothing can be guaranteed for the given problem dimensions.

We have thus obtained guarantees for the median estimation error of the Dantzig selector, BPDN, and OMP. Under these settings, no guarantee can be made for the performance of the thresholding algorithm. Indeed, as we will see, for some choices of $\boldsymbol{x}_0$ satisfying the above requirements, the performance of the thresholding algorithm is not proportional to $s\sigma^2 \log m$. To obtain thresholding guarantees, one requires a narrower range between $|x_{\min}|$ and $|x_{\max}|$.

The RIC-based BPDN guarantee of Theorem 6.3 is not applicable in the present setting; however, by comparing the BPDN guarantees of Theorems 6.3 and 6.4, it is evident that the latter is substantially tighter, even when the conditions for both theorems hold. It is also interesting to note that by applying Lemma 6.1 to the (RIC-based) Dantzig selector guarantee of Candès and Tao [8], one can obtain for the current setting

$$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{DS}}\|_2^2 \leq 361.8 s\sigma^2 \log m \quad \text{w.p. } \tfrac{3}{4}. \tag{6.37}$$

This result is again looser than (6.36), once again demonstrating that considerable improvement in performance guarantees is possible by direct use of the mutual coherence.

To measure the actual median error obtained by various estimators, 8 different parameter vectors $\boldsymbol{x}_0$ were selected. These differed in the distribution of the magnitudes of the nonzero components within the range $[|x_{\min}|, |x_{\max}|]$ and in the locations of the nonzero elements. For each parameter $\boldsymbol{x}_0$, a set of measurement vectors $\boldsymbol{b}$ were obtained from (6.1). The estimation algorithms of Section 6.2.2 were then applied to each measurement realization. For the Dantzig selector, τ was selected using (6.15), and for BPDN, γ was chosen as the smallest value such that the probability of success (6.27) would exceed $1/2$. The median over noise realizations of the distance $\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}\|_2^2$ was then computed for each estimator. This process was repeated for 10 values of the noise variance σ^2 in the range $10^{-8} \leq \sigma^2 \leq 1$. The results are plotted in Fig. 6.1 as a function of σ^2. The performance guarantees (6.36)–(6.37) are also plotted.[5]

It is evident from Fig. 6.1 that some parameter vectors are more difficult to estimate than others. Indeed, there is a large variety of parameters $\boldsymbol{x}_0$ satisfying the problem requirements, and it is likely that some of them come closer to the theoretical limits than the parameters chosen in our experiment. This highlights the importance of performance guarantees in ensuring adequate performance for *all* parameter values. On the other hand, it is quite possible that further improvements of the constants in the performance bounds are possible. For example, the Dantzig selector guarantee is almost 50 times higher than the worst of the examined parameter

[5]The guarantee (6.37) actually requires a slightly different value of τ, but this difference has a negligible effect on performance and is ignored.

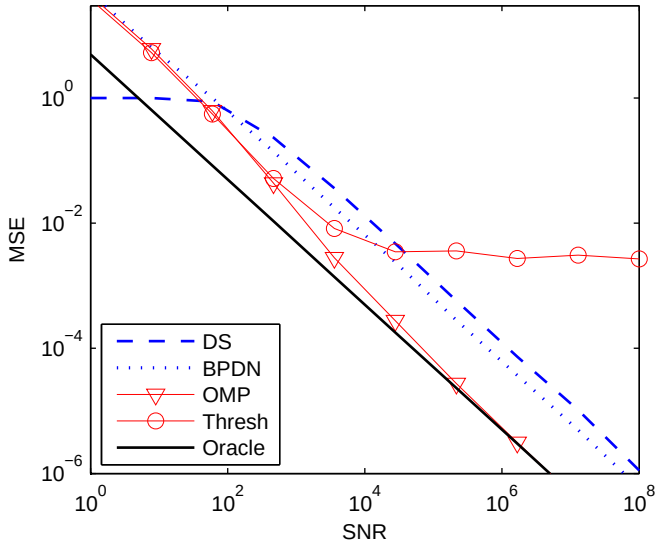

Figure 6.2: MSE of various estimators as a function of the SNR. The sparsity level is $s = 5$ and the dictionary is a 256×512 two-ortho matrix.

values.

In practice, it is more common to measure the MSE of an estimator than its median error. Our next goal is to determine whether the behavior predicted by our theoretical analysis is also manifested in the MSE of the various estimators. To this end, we conducted an experiment in which the MSEs of the estimators of Section 6.2.2 were compared. In this simulation, we chose the two-ortho dictionary $\boldsymbol{A} = [\boldsymbol{I} \;\; \boldsymbol{H}]$, where $\boldsymbol{I}$ is the 256×256 identity matrix and $\boldsymbol{H}$ is the 256×256 Hadamard matrix with normalized columns.[6] Once again, the RICs of this dictionary are unknown. However, the coherence in this case is given by $\mu = 1/16$, and consequently, the ℓ_1 relaxation guarantees of Section 6.3.1 hold for $s \leq 5$.

We obtained the parameter vector $\boldsymbol{x}_0$ for this experiment by selecting a 5-element support at random, choosing the nonzero entries from a white Gaussian distribution, and then normalizing the resulting vector so that $\|\boldsymbol{x}_0\|_2 = 1$. The regularization parameters τ and γ of the Dantzig selector and BPDN were chosen as recommended by Theorem 6.2 and Corollary 6.5, respectively; for the latter, a value of $\alpha = 1$ was chosen. The MSE of each estimate was then calculated by averaging over repeated realizations of $\boldsymbol{x}_0$ and the noise. The experiment was conducted for 10 values of the noise variance σ^2 and the results are plotted in Fig. 6.2 as a

[6]Similar experiments were performed on a variety of other dictionaries, including an overcomplete DCT and a matrix containing Gaussian random entries. The different dictionaries yielded comparable results, which are not reported here.

function of the signal-to-noise ratio (SNR), which is defined by

$$\text{SNR} = \frac{\|\boldsymbol{x}_0\|_2^2}{n\sigma^2} = \frac{1}{n\sigma^2}. \tag{6.38}$$

To compare this plot with the theoretical results of Section 6.3, observe first the situation at high SNR. In this case, OMP, BPDN, and the Dantzig selector all achieve performance which is proportional to the oracle MSE (or CRB) given by (6.13). Among these, OMP is closest to the CRB, followed by BPDN and, finally, the Dantzig selector. This behavior matches the proportionality constants given in the theorems of Section 6.3. Indeed, for small σ, the condition (6.30) holds even for large α, and thus Theorem 6.6 guarantees that OMP will recover the correct support of $\boldsymbol{x}_0$ with high probability, explaining the convergence of this estimator to the oracle. By contrast, the performance of the thresholding algorithm levels off at high SNR; this is again predicted by Theorem 6.6, since, even when $\sigma = 0$, the condition (6.33) does not always hold, unless $|x_{\min}|$ is not much smaller than $|x_{\max}|$. Thus, for our choice of $\boldsymbol{x}_0$, Theorem 6.6 does not guarantee near-oracle performance for the thresholding algorithm, even at high SNR.

With increasing noise, Theorem 6.6 requires a corresponding increase in $|x_{\min}|$ to guarantee the success of the greedy algorithms. Consequently, Fig. 6.2 demonstrates a deterioration of these algorithms when the SNR is low. On the other hand, the theorems for the relaxation algorithms make no such assumptions, and indeed these approaches continue to perform well, compared with the oracle estimator, even when the noise level is high. In particular, the Dantzig selector outperforms the CRB at low SNR; this is because the CRB is a bound on unbiased techniques, whereas when the noise is large, biased techniques such as an ℓ_1 penalty become very effective. Robustness to noise is thus an important advantage of ℓ_1-relaxation techniques.

It is also interesting to examine the effect of the support size s on the performance of the various estimators. To this end, 15 support sizes in the range $2 \leq s \leq 30$ were tested. For each value of s, random vectors $\boldsymbol{x}_0$ having s nonzero entries were selected as in the previous simulation. The dictionary $\boldsymbol{A}$ was the 256×512 two-ortho matrix defined above; as in the previous experiment, other matrices were also tested and provided similar results. The standard deviation of the noise for this experiment was $\sigma = 0.01$. The results are plotted in Fig. 6.3.

As mentioned above, the mutual coherence of the dictionary $\boldsymbol{A}$ is $1/16$, so that the proposed performance guarantees apply only when $\boldsymbol{x}_0$ is quite sparse ($s \leq 5$). Nevertheless, Fig. 6.3 demonstrates that the estimation algorithms (with the exception of the thresholding approach) exhibit a graceful degradation as the support of $\boldsymbol{x}_0$ increases. At first sight this would appear to mean that the performance guarantees provided are overly pessimistic. For example, it is pos-

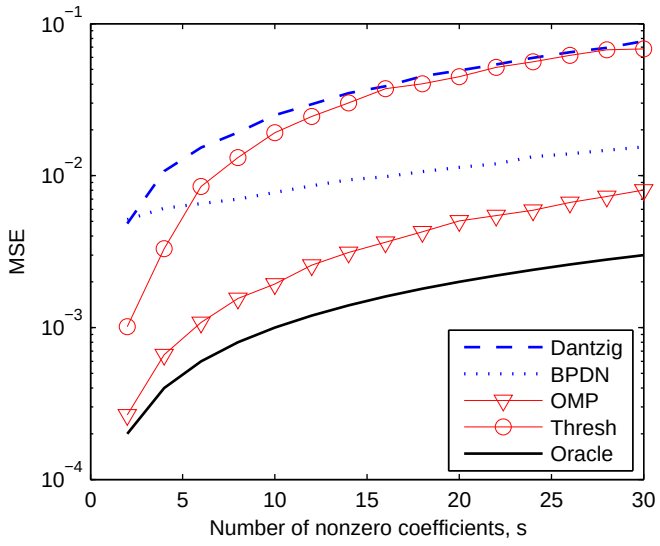

Figure 6.3: MSE of various estimators as a function of the support size s. The noise standard deviation is $\sigma = 0.01$ and the dictionary is a 256×512 two-ortho matrix.

sible that the RICs in the present setting, while unknown, are fairly low and permit a stronger analysis than that of Section 6.3. It is also quite reasonable to expect, as mentioned above, that some improvement in the theoretical guarantees is possible. However, it is worth recalling that the performance guarantees proposed in this paper apply to all sparse vectors, while the numerical results describe the performance averaged over different values of $\boldsymbol{x}_0$. Thus it is possible that there exist particular parameter values for which the performance is considerably poorer than that reported in Fig. 6.3. Indeed, there exist values of $\boldsymbol{A}$ and $\boldsymbol{x}_0$ for which BPDN yields grossly incorrect results even when $\|\boldsymbol{x}_0\|_0$ is on the order of $1/\mu$ [21]. However, identifying such worst-case parameters numerically is quite difficult; this is doubtlessly at least part of the reason for the apparent pessimism of the performance guarantees.

6.6 Conclusion

The performance of an estimator depends on the problem setting under consideration. As we have seen in Section 6.4, under the adversarial noise scenario, the estimation error of any algorithm can be as high as the noise power; in other words, the assumption of sparsity has not yielded any denoising effect. On the other hand, when both the noise and the design matrix $\boldsymbol{A}$ are random, practical estimators come close to the performance of the oracle estimator. In this paper, we examined a middle ground between these two cases, namely the setting in which $\boldsymbol{x}_0$

and $\boldsymbol{A}$ are deterministic but the noise is random. As we have shown, despite the fact that less information is available in this case, a variety of estimation techniques are still guaranteed to achieve performance close to that of the oracle estimator.

Our theoretical and numerical results suggest some conclusions concerning the choice of an estimator. In particular, at high SNR values, it appears that the greedy OMP algorithm has an advantage over the other algorithms considered herein. In this case the support set of $\boldsymbol{x}_0$ can be recovered accurately and OMP thus converges to the oracle estimator; by contrast, ℓ_1 relaxations have a shrinkage effect which causes a loss of accuracy at high SNR. This is of particular interest since greedy algorithms are also computationally more efficient than relaxation methods. On the other hand, the ℓ_1 relaxation techniques, and particularly the Dantzig selector, appear to be more effective than the greedy algorithms when the noise level is significant: in this case, shrinkage is a highly effective denoising technique. Indeed, as a result of the bias introduced by the shrinkage, ℓ_1-based approaches can even perform better than the oracle estimator and the Cramér–Rao bound.

6.A Proof of Theorem 6.3

To prove the theorem, we will adapt Theorem 7.2 and Lemma 4.1 of [23] to the present setting, and then apply the coherence bounds of Lemma 6.1. Note that the normalization of the matrix $\boldsymbol{A}$ differs between the present paper and that of [23], and consequently the results stated herein differs by a factor of n from the original.

We begin by recalling the definition of the "restricted eigenvalue condition" $RE(s, q, 3)$ of [23]. Let

$$\kappa = \min_{J_0 \subseteq \{1,\ldots,m\} : |J_0| \leq s} \quad \min_{\boldsymbol{x} \neq 0 : \|\boldsymbol{x}_{J_0^c}\|_1 \leq 3\|\boldsymbol{x}_{J_0}\|_1} \frac{\|\boldsymbol{A}\boldsymbol{x}\|_2}{\left\|\boldsymbol{x}_{J_0 \cup J_1}\right\|_2} \tag{6.39}$$

where $\boldsymbol{x}_J$ denotes the subvector of $\boldsymbol{x}$ indexed by the elements of the set J, J_0^c is the complement of the set J_0, and J_1 contains the indices of the q largest elements in $\boldsymbol{x}$ which are not contained in J_0. If $\kappa > 0$, then the matrix $\boldsymbol{A}$ is said to satisfy the restricted eigenvalue condition $RE(s, q, 3)$ with parameter κ. Note that $\kappa \leq 1$, since one could, for example, choose $J_0 = \{1\}$ and $\boldsymbol{x} = [1, 0, \ldots, 0]^T$, whereupon it follows from (6.2) that the objective function in (6.39) equals 1.

Using this definition, the following result follows directly from [23, Theorem 7.2].

Theorem 6.8 (Bickel, Ritov and Tsybakov)**.** *Under the setting* (6.1)*, assume that the restricted eigenvalue condition $RE(s, q, 3)$ holds for some q. Consider the BPDN estimator* (6.8) *with γ given by* (6.19)

and choose $a > 2\sqrt{2}$. Then, with probability exceeding $1 - m^{1-a^2/8}$, we have

$$\|\hat{\boldsymbol{x}}_{\mathrm{BP}} - \boldsymbol{x}_0\|_2^2 \leq 16 \left(1 + 3\sqrt{\frac{s}{q}}\right)^2 \frac{a^2}{\kappa^4} s\sigma^2 \log m. \tag{6.40}$$

From [23, Lemma 4.1], it follows that a sufficient condition for $RE(s, s, 3)$ to hold is $\delta_{2s} + 3\theta_{s,2s} < 1$, where δ_{2s} and $\theta_{s,2s}$ are the RICs defined in Section 6.2.1. Applying the bounds of Lemma 6.1, it follows that $RE(s, s, 3)$ holds if (6.18) holds. Furthermore, from [23, Lemma 4.1] we have that in this case,

$$\kappa \geq \sqrt{1 - \delta_{2s}} - \frac{3\theta_{s,2s}}{\sqrt{1 - \delta_{2s}}} \tag{6.41}$$

from which, again by applying Lemma 6.1, we obtain

$$\kappa \geq \sqrt{1 - (2s-1)\mu} - \frac{s\mu 3\sqrt{2}}{\sqrt{1 - (2s-1)\mu}}. \tag{6.42}$$

Since decreasing κ does not violate the bound (6.40), our coherence-based bound must choose κ given by (6.22). Defining $1 + \alpha = a^2/8$ and choosing $q = s$, we obtain the probability of success (6.20) and the bound (6.21). Thus we have translated the results of Bickel et al. to the coherence-based Theorem 6.3.

6.B Proof of Theorem 6.4

The proof is based closely on the work of Tropp [22]. From the triangle inequality,

$$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{BP}}\|_2 \leq \|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{or}}\|_2 + \|\hat{\boldsymbol{x}}_{\mathrm{or}} - \hat{\boldsymbol{x}}_{\mathrm{BP}}\|_2 \tag{6.43}$$

where $\hat{\boldsymbol{x}}_{\mathrm{or}}$ is the oracle estimator. Our goal is to separately bound the two terms on the right-hand side of (6.43). Indeed, as we will see, the two constants $\sigma\sqrt{3}$ and $\frac{3}{2}\gamma$ in (6.25) arise, respectively, from the two terms in (6.43).

Beginning with the term $\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\mathrm{or}}\|_2$, let $\boldsymbol{x}_{0,\Lambda}$ denote the s-vector containing the elements of $\boldsymbol{x}_0$ indexed by Λ_0, and similarly, let $\hat{\boldsymbol{x}}_{\mathrm{or},\Lambda}$ denote the corresponding subvector of $\hat{\boldsymbol{x}}_{\mathrm{or}}$. We then have

$$\begin{aligned} \boldsymbol{x}_{0,\Lambda} - \hat{\boldsymbol{x}}_{\mathrm{or},\Lambda} &= \boldsymbol{x}_{0,\Lambda} - \boldsymbol{A}_{\Lambda_0}^{\dagger}(\boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w}) \\ &= \boldsymbol{x}_{0,\Lambda} - \boldsymbol{A}_{\Lambda_0}^{\dagger}(\boldsymbol{A}_{\Lambda_0}\boldsymbol{x}_{0,\Lambda} + \boldsymbol{w}) \\ &= -\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{w}, \end{aligned} \tag{6.44}$$

where we have used the fact that $\boldsymbol{A}_{\Lambda_0}$ has full column rank, which is a consequence [82] of the condition (6.23). Thus, $\boldsymbol{x}_{0,\Lambda} - \hat{\boldsymbol{x}}_{\mathrm{or},\Lambda}$ is a Gaussian random vector with mean $\boldsymbol{0}$ and covariance $\sigma^2 \boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{A}_{\Lambda_0}^{\dagger T} = \sigma^2(\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{A}_{\Lambda_0})^{-1}$.

For future use, we note that the cross-correlation between $\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{w}$ and $(\boldsymbol{I}-\boldsymbol{A}_{\Lambda_0}\boldsymbol{A}_{\Lambda_0}^{\dagger})\boldsymbol{w}$ is

$$\begin{aligned} E\Big\{\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{w}\boldsymbol{w}^T(\boldsymbol{I}-\boldsymbol{A}_{\Lambda_0}\boldsymbol{A}_{\Lambda_0}^{\dagger})^T\Big\} &= \sigma^2\boldsymbol{A}_{\Lambda_0}^{\dagger}(\boldsymbol{I}-\boldsymbol{A}_{\Lambda_0}\boldsymbol{A}_{\Lambda_0}^{\dagger})^T \\ &= \boldsymbol{0}, \end{aligned} \tag{6.45}$$

where we have used the fact [71, Th. 1.2.1] that for any matrix $\boldsymbol{M}$

$$\boldsymbol{M}^{\dagger}\boldsymbol{M}^{\dagger T}\boldsymbol{M}^T = (\boldsymbol{M}^T\boldsymbol{M})^{\dagger}\boldsymbol{M}^T = \boldsymbol{M}^{\dagger}. \tag{6.46}$$

Since $\boldsymbol{w}$ is Gaussian, it follows that $\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{w}$ and $(\boldsymbol{I}-\boldsymbol{A}_{\Lambda_0}\boldsymbol{A}_{\Lambda_0}^{\dagger})\boldsymbol{w}$ are statistically independent. Furthermore, because $\boldsymbol{x}_{0,\Lambda}-\hat{\boldsymbol{x}}_{\mathrm{or},\Lambda}$ depends on $\boldsymbol{w}$ only through $\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{w}$, we conclude that

$$\boldsymbol{x}_0-\hat{\boldsymbol{x}}_{\mathrm{or}} \text{ is statistically independent of } (\boldsymbol{I}-\boldsymbol{A}_{\Lambda_0}\boldsymbol{A}_{\Lambda_0}^{\dagger})\boldsymbol{w}. \tag{6.47}$$

We now wish to bound the probability that $\|\boldsymbol{x}_0-\hat{\boldsymbol{x}}_{\mathrm{or}}\|_2^2 > 3s\sigma^2$. Let $\boldsymbol{z}$ be a normalized Gaussian random variable, $\boldsymbol{z}\sim N(\boldsymbol{0},\boldsymbol{I}_s)$. Then

$$\begin{aligned} &\Pr\big\{\|\boldsymbol{x}_0-\hat{\boldsymbol{x}}_{\mathrm{or}}\|_2^2 > 3s\sigma^2\big\} \\ &= \Pr\left\{\left\|\sigma(\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{A}_{\Lambda_0})^{-1/2}\boldsymbol{z}\right\|_2^2 \geq 3s\sigma^2\right\} \\ &\leq \Pr\left\{\left\|(\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{A}_{\Lambda_0})^{-1/2}\right\|^2\|\boldsymbol{z}\|_2^2 \geq 3s\right\} \end{aligned} \tag{6.48}$$

where $\|\boldsymbol{M}\|$ denotes the maximum singular value of the matrix $\boldsymbol{M}$. Thus, $\|(\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{A}_{\Lambda_0})^{-1/2}\| = 1/s_{\min}$, where $s_{\min}$ is the minimum singular value of $\boldsymbol{A}_{\Lambda_0}$. From the Gershgorin disc theorem [116, p. 320], it follows that $s_{\min}\geq\sqrt{1-(s-1)\mu}$. Using (6.23), this can be simplified to $s_{\min}\geq\sqrt{2/3}$, and therefore

$$\left\|(\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{A}_{\Lambda_0})^{-1/2}\right\| \leq \sqrt{\frac{3}{2}}. \tag{6.49}$$

Combining with (6.48) yields

$$\Pr\big\{\|\boldsymbol{x}_0-\hat{\boldsymbol{x}}_{\mathrm{or}}\|_2^2 > 3s\sigma^2\big\} \leq \Pr\big\{\|\boldsymbol{z}\|_2^2\geq 2s\big\}. \tag{6.50}$$

Observe that $\|\boldsymbol{z}\|_2^2$ is the sum of s independent normalized Gaussian random variables. The right-hand side of (6.50) is therefore $1-F_{\chi_s^2}(2s)$, where $F_{\chi_s^2}(\cdot)$ is the cumulative distribution function of the χ^2 distribution with s degrees of freedom. Using the formula [119, §16.3] for $F_{\chi_s^2}(\cdot)$, we have

$$\Pr\big\{\|\boldsymbol{x}_0-\hat{\boldsymbol{x}}_{\mathrm{or}}\|_2^2 > 3s\sigma^2\big\} \leq Q\big(\tfrac{1}{2}s, s\big) \tag{6.51}$$

where $Q(a,z)$ is the regularized Gamma function

$$Q(a,z) \triangleq \frac{\int_z^{\infty}t^{a-1}e^{-t}dt}{\int_0^{\infty}t^{a-1}e^{-t}dt}. \tag{6.52}$$

$Q\left(\frac{1}{2}s, s\right)$ decays exponentially as $s \to \infty$, and it can be seen that

$$Q\left(\tfrac{1}{2}s, s\right) < e^{-s/7} \quad \text{for all } s. \tag{6.53}$$

We thus conclude that the event

$$\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\text{or}}\|_2^2 \leq 3s\sigma^2 \tag{6.54}$$

occurs with probability no smaller than $1 - e^{-s/7}$. Note that the same technique can be applied to obtain bounds on the probability that $\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_{\text{or}}\|_2^2 > \alpha s\sigma^2$, for any $\alpha > \frac{2}{3}$. The only difference will be the rate of exponential decay in (6.53). However, the distance between $\boldsymbol{x}_0$ and $\hat{\boldsymbol{x}}_{\text{or}}$ is usually small compared with the distance between $\hat{\boldsymbol{x}}_{\text{or}}$ and $\hat{\boldsymbol{x}}_{\text{BP}}$, so that such an approach does not significantly affect the overall result.

The above calculations provided a bound on the first term in (6.43). To address the second term $\|\hat{\boldsymbol{x}}_{\text{or}} - \hat{\boldsymbol{x}}_{\text{BP}}\|_2$, define the random event

$$G: \ \max_i \left| \boldsymbol{a}_i^T (\boldsymbol{I} - \boldsymbol{A}_{\Lambda_0} \boldsymbol{A}_{\Lambda_0}^{\dagger}) \boldsymbol{b} \right| \leq \tfrac{1}{2}\gamma \tag{6.55}$$

where $\boldsymbol{a}_i$ is the ith column of $\boldsymbol{A}$. It is shown in [22, App. IV-A] that

$$\Pr\{G\} \geq 1 - (m - s) \exp\left(-\frac{\gamma^2}{8\sigma^2} \right). \tag{6.56}$$

If G indeed occurs, then the portion of the measurements $\boldsymbol{b}$ which do not belong to the range space of $\boldsymbol{A}_{\Lambda_0}$ are small, and consequently it has been shown [22, Cor. 9] that, in this case, the solution $\hat{\boldsymbol{x}}_{\text{BP}}$ to (6.8) is unique, the support of $\hat{\boldsymbol{x}}_{\text{BP}}$ is a subset of Λ_0, and

$$\|\hat{\boldsymbol{x}}_{\text{BP}} - \hat{\boldsymbol{x}}_{\text{or}}\|_\infty \leq \tfrac{3}{2}\gamma. \tag{6.57}$$

Since both $\hat{\boldsymbol{x}}_{\text{BP}}$ and $\hat{\boldsymbol{x}}_{\text{or}}$ are nonzero only in Λ_0, this implies that

$$\|\hat{\boldsymbol{x}}_{\text{BP}} - \hat{\boldsymbol{x}}_{\text{or}}\|_2 \leq \tfrac{3}{2}\gamma\sqrt{s}. \tag{6.58}$$

The event G depends on the random variable $\boldsymbol{w}$ only through $(\boldsymbol{I} - \boldsymbol{A}_{\Lambda_0}\boldsymbol{A}_{\Lambda_0}^{\dagger})\boldsymbol{w}$. Thus, it follows from (6.47) that G is statistically independent of the event (6.48). The probability that both events occur simultaneously is therefore given by the product of their respective probabilities. In other words, with probability exceeding (6.24), both (6.58) and (6.54) hold. Using (6.43) completes the proof of the theorem.

6.C Proof of Theorem 6.6

The claims concerning both algorithms are closely related. To emphasize this similarity, we first provide several lemmas which will be used to prove both results. These lemmas are all based on an analysis of the random event

$$B = \left\{ \max_{1 \le i \le m} |\boldsymbol{a}_i^T \boldsymbol{w}| < \tau \right\} \tag{6.59}$$

where

$$\tau \triangleq \sigma \sqrt{2(1+\alpha) \log m} \tag{6.60}$$

and $\alpha > 0$. Our proof will be based on demonstrating that B occurs with high probability, and that when B does occur, both thresholding and OMP achieve near-oracle performance.

Lemma 6.9. *Suppose that $\boldsymbol{w} \sim N(\boldsymbol{0}, \sigma^2 \boldsymbol{I})$. Then, the event B of* (6.59) *occurs with a probability of at least* (6.31).

Proof. The random variables $\{\boldsymbol{a}_i^T \boldsymbol{w}\}_{i=1}^m$ are jointly Gaussian. Therefore, by Šidák's lemma [120, Th. 1]

$$\Pr\{B\} = \Pr\left\{ \max_{1 \le i \le m} |\boldsymbol{a}_i^T \boldsymbol{w}| < \tau \right\} \ge \prod_{i=1}^{m} \Pr\left\{ |\boldsymbol{a}_i^T \boldsymbol{w}| \le \tau \right\}. \tag{6.61}$$

Since $\|\boldsymbol{a}_i\|_2 = 1$, each random variable $\boldsymbol{a}_i^T \boldsymbol{w}$ has mean zero and variance σ^2. Consequently,

$$\Pr\left\{ |\boldsymbol{a}_i^T \boldsymbol{w}| < \tau \right\} = 1 - 2Q\left(\frac{\tau}{\sigma}\right) \tag{6.62}$$

where $Q(x) = (1/\sqrt{2\pi}) \int_x^\infty e^{-z^2/2} dz$ is the Gaussian tail probability. Using the bound

$$Q(x) \le \frac{1}{x\sqrt{2\pi}} e^{-x^2/2} \tag{6.63}$$

we obtain from (6.62)

$$\Pr\left\{ |\boldsymbol{a}_i^T \boldsymbol{w}| < \tau \right\} \ge 1 - \eta \tag{6.64}$$

where

$$\eta \triangleq \sqrt{\frac{2}{\pi}} \cdot \frac{\sigma}{\tau} e^{-\tau^2/2\sigma^2}. \tag{6.65}$$

When $\eta > 1$, the bound (6.31) is meaningless and the theorem holds vacuously. Otherwise, when $\eta \le 1$, we have from (6.61) and (6.64)

$$\Pr\{B\} \ge (1-\eta)^m \ge 1 - m\eta \tag{6.66}$$

where the final inequality holds for any $\eta \le 1$ and $m \ge 1$. Substituting the values of η and τ and simplifying, we obtain that B holds with a probability no lower than (6.31), as required. □

The next lemma demonstrates that, under suitable conditions, correlating $\boldsymbol{b}$ with the dictionary atoms $\boldsymbol{a}_i$ is an effective method of identifying the atoms participating in the support of $\boldsymbol{x}_0$.

Lemma 6.10. *Let $\boldsymbol{x}_0$ be a vector with support $\Lambda_0 = \operatorname{supp}(\boldsymbol{x}_0)$ of size $s = |\Lambda_0|$, and let $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w}$ for some noise vector $\boldsymbol{w}$. Define $|x_{\min}|$ and $|x_{\max}|$ as in* (6.29)*, and suppose that*

$$|x_{\max}| - (2s-1)\mu|x_{\max}| \geq 2\tau. \tag{6.67}$$

Then, if the event B of (6.59) *holds, we have*

$$\max_{j\in\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}| > \max_{j\notin\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}|. \tag{6.68}$$

If, rather than (6.67)*, the stronger condition*

$$|x_{\min}| - (2s-1)\mu|x_{\max}| \geq 2\tau \tag{6.69}$$

is given, then, under the event B, we have

$$\min_{j\in\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}| > \max_{j\notin\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}|. \tag{6.70}$$

Proof. The proof is an adaptation of [7, Lemma 5.2]. Beginning with the term $\max_{j\notin\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}|$, we have, under the event B,

$$\begin{aligned}
\max_{j\notin\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}| &= \max_{j\notin\Lambda_0} \left| \boldsymbol{a}_j^T\boldsymbol{w} + \sum_{i\in\Lambda_0} x_i \boldsymbol{a}_j^T\boldsymbol{a}_i \right| \\
&\leq \max_{j\notin\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{w}| + \max_{j\notin\Lambda_0} \sum_{i\in\Lambda_0} \left| x_i \boldsymbol{a}_j^T\boldsymbol{a}_i \right| \\
&< \tau + s\mu|x_{\max}|.
\end{aligned} \tag{6.71}$$

On the other hand, when B holds,

$$\begin{aligned}
\max_{j\in\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}| &= \max_{j\in\Lambda_0} \left| x_j + \boldsymbol{a}_j^T\boldsymbol{w} + \sum_{i\in\Lambda_0\setminus\{j\}} x_i \boldsymbol{a}_j^T\boldsymbol{a}_i \right| \\
&\geq |x_{\max}| - \max_{j\in\Lambda_0} \left| \boldsymbol{a}_j^T\boldsymbol{w} + \sum_{i\in\Lambda_0\setminus\{j\}} x_i \boldsymbol{a}_j^T\boldsymbol{a}_i \right| \\
&> |x_{\max}| - \tau - (s-1)\mu|x_{\max}| \\
&= |x_{\max}| - (2s-1)\mu|x_{\max}| - \tau + s\mu|x_{\max}|.
\end{aligned} \tag{6.72}$$

Together with (6.71), this yields

$$\max_{j\in\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}| > |x_{\max}| - (2s-1)\mu|x_{\max}| - 2\tau + \max_{j\notin\Lambda_0} |\boldsymbol{a}_j^T\boldsymbol{b}|. \tag{6.73}$$

Thus, under the condition (6.67), we obtain (6.68). Similarly, when B holds, we have

$$\begin{aligned}\min_{j\in\Lambda_0}\left|\boldsymbol{a}_j^T\boldsymbol{b}\right| &= \min_{j\in\Lambda_0}\left|x_j+\boldsymbol{a}_j^T\boldsymbol{w}+\sum_{i\in\Lambda_0\setminus\{j\}}x_i\boldsymbol{a}_j^T\boldsymbol{a}_i\right| \\ &> |x_{\min}|-\tau-(s-1)\mu|x_{\max}| \\ &= |x_{\min}|-(2s-1)\mu|x_{\max}|-\tau+s\mu|x_{\max}|. \end{aligned} \tag{6.74}$$

Again using (6.71), we obtain

$$\min_{j\in\Lambda_0}\left|\boldsymbol{a}_j^T\boldsymbol{b}\right| > |x_{\min}|-(2s-1)\mu|x_{\max}|-2\tau+\max_{j\notin\Lambda_0}|\boldsymbol{a}_j^T\boldsymbol{b}|. \tag{6.75}$$

Consequently, under the assumption (6.69), we conclude that (6.70) holds, as required. □

The following lemma bounds the performance of the oracle estimator under the event B. The usefulness of this lemma stems from the fact that, if either OMP or the thresholding algorithm correctly identify the support of $\boldsymbol{x}_0$, then their estimate is identical to that of the oracle.

Lemma 6.11. *Let $\boldsymbol{x}_0$ be a vector with support $\Lambda_0 = \operatorname{supp}(\boldsymbol{x}_0)$, and let $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{w}$ for some noise vector $\boldsymbol{w}$. If the event B of* (6.59) *occurs, then*

$$\|\hat{\boldsymbol{x}}_{\text{or}}-\boldsymbol{x}_0\|_2^2 \le 2s\sigma^2(1+\alpha)\log m\frac{1}{(1-(s-1)\mu)^2}. \tag{6.76}$$

Proof. Note that both $\hat{\boldsymbol{x}}_{\text{or}}$ and $\boldsymbol{x}_0$ are supported on Λ_0, and therefore

$$\|\hat{\boldsymbol{x}}_{\text{or}}-\boldsymbol{x}_0\|_2^2 = \|\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{b}-\boldsymbol{x}_{0,\Lambda_0}\|_2^2 \tag{6.77}$$

where $\boldsymbol{x}_{0,\Lambda_0}$ is the subvector of nonzero entries of $\boldsymbol{x}_0$. We thus have, under the event B,

$$\begin{aligned}\|\hat{\boldsymbol{x}}_{\text{or}}-\boldsymbol{x}_0\|_2^2 &= \|\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{A}_{\Lambda_0}\boldsymbol{x}_{0,\Lambda_0}+\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{w}-\boldsymbol{x}_{0,\Lambda_0}\|_2^2 \\ &= \|\boldsymbol{A}_{\Lambda_0}^{\dagger}\boldsymbol{w}\|_2^2 \\ &= \left\|(\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{A}_{\Lambda_0})^{-1}\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{w}\right\|_2^2 \\ &\le \left\|(\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{A}_{\Lambda_0})^{-1}\right\|^2\sum_{i\in\Lambda_0}(\boldsymbol{a}_i^T\boldsymbol{w})^2 \\ &\le \frac{1}{(1-(s-1)\mu)^2}s\sigma^2 2(1+\alpha)\log m \end{aligned} \tag{6.78}$$

where, in the last step, we used the definition (6.59) of B and the fact that $\|\boldsymbol{A}_{\Lambda_0}^T\boldsymbol{A}_{\Lambda_0}\| \ge 1-(s-1)\mu$, which was demonstrated in Appendix 6.B. This completes the proof the lemma. □

We are now ready to prove Theorem 6.6. The proof for the thresholding algorithm is obtained by combining the three lemmas presented above. Indeed, Lemma 6.9 ensures that the event B occurs with probability at least as high as the required probability of success (6.31). Whenever B occurs, we have by Lemma 6.10 that the atoms corresponding to Λ_0 all have strictly higher correlation with $\boldsymbol{b}$ than the off-support atoms, so that the thresholding algorithm identifies the correct support Λ_0, and is thus equivalent to the oracle estimator $\hat{\boldsymbol{x}}_{\text{or}}$ as long as B holds. Finally, by Lemma 6.11, identification of the true support Λ_0 guarantees the required error (6.32).

We now prove the OMP performance guarantee. Our aim is to show that when B occurs, OMP correctly identifies the support of $\boldsymbol{x}_0$; the result then follows by Lemmas 6.9 and 6.11. To this end we employ the technique used in the proof of [7, Th. 5.1]. We begin by examining the first iteration of the OMP algorithm, in which one identifies the atom $\boldsymbol{a}_i$ whose correlation with $\boldsymbol{b}$ is maximal. Note that (6.30) implies (6.67), and therefore, by Lemma 6.10, the atom having the highest correlation with $\boldsymbol{b}$ corresponds to an element in the support Λ_0 of $\boldsymbol{x}_0$. Consequently, the first step of the OMP algorithm correctly identifies an element in Λ_0.

The proof now continues by induction. Suppose we are currently in the ith iteration of OMP, with $1 < i \leq s$, and assume that atoms from the correct support were identified in all $i-1$ previous steps. Referring to the notation used in the definition of OMP in Section 6.2.2, this implies that $\text{supp}(\hat{\boldsymbol{x}}^{i-1}_{\text{OMP}}) = \Lambda^{i-1} \subset \Lambda_0$. The ith step consists of identifying the atom $\boldsymbol{a}_i$ which is maximally correlated with the residual $\boldsymbol{r}^i$. By the definition of $\boldsymbol{r}^i$, we have

$$\boldsymbol{r}^i = \boldsymbol{A}\tilde{\boldsymbol{x}}^{i-1} + \boldsymbol{w} \tag{6.79}$$

where $\tilde{\boldsymbol{x}}^{i-1} = \boldsymbol{x}_0 - \hat{\boldsymbol{x}}^{i-1}_{\text{OMP}}$. Thus $\text{supp}(\tilde{\boldsymbol{x}}^{i-1}) \subseteq \Lambda_0$, so that $\boldsymbol{r}^i$ is a noisy measurement of the vector $\boldsymbol{A}\tilde{\boldsymbol{x}}^{i-1}$, which has a sparse representation consisting of no more than s atoms. Now, since

$$\|\hat{\boldsymbol{x}}^{i-1}_{\text{OMP}}\|_0 = i - 1 < s = \|x_0\|_0, \tag{6.80}$$

it follows that at least one nonzero entry in $\tilde{\boldsymbol{x}}^{i-1}$ is equal to the corresponding entry in $\boldsymbol{x}_0$. Consequently

$$\max_i |\tilde{x}^{i-1}_i| \geq |x_{\min}|. \tag{6.81}$$

Note that the model (6.79) is precisely of the form (6.1), with $\boldsymbol{r}^i$ taking the place of the measurements $\boldsymbol{b}$ and $\tilde{\boldsymbol{x}}^{i-1}$ taking the place of the sparse vector $\boldsymbol{x}_0$. It follows from (6.81) and (6.30) that this model satisfies the requirement (6.67). Consequently, by Lemma 6.10, we have that under

the event B,

$$\max_{i \in \Lambda_0} |\boldsymbol{a}_i^T \boldsymbol{r}^i| > \max_{i \notin \Lambda_0} |\boldsymbol{a}_i^T \boldsymbol{r}^i|. \tag{6.82}$$

Therefore, the ith iteration of OMP will choose an element within Λ_0 to add to the support. By induction it follows that the first s steps of OMP all identify elements in Λ_0, and since OMP never chooses the same element twice, the entire support Λ_0 will be identified after s iterations. This completes the proof of Theorem 6.6.

Chapter 7

Bounds and Guarantees for Block-Sparse Estimators

This chapter has been submitted for publication as:

- Z. Ben-Haim and Y. C. Eldar, "Near-oracle performance of greedy block-sparse estimation techniques from noisy measurements," submitted to *IEEE J. of Selected Topics in Signal Processing*, Sep. 2010.

7.1 Introduction

The success of signal processing techniques depends to a large extent on the availability of an appropriate model which captures our knowledge of the system under consideration and translates it to a productive mathematical framework. There is consequently an ongoing search for mathematical models which can accurately describe real-world signals. In recent years, much research has been devoted to the sparse representation model, which stems from the observation that many signals can be approximated using a small number of elements, or "atoms," chosen from a large dictionary [12, 15, 84]. Thus, we may write $\boldsymbol{y} = \boldsymbol{D}\boldsymbol{x} + \boldsymbol{w}$, where the signal $\boldsymbol{y}$ is a linear combination of a small number of columns of the dictionary matrix $\boldsymbol{D}$, corrupted by noise $\boldsymbol{w}$. Since only a small number of elements of $\boldsymbol{D}$ are required for this representation, the vector $\boldsymbol{x}$ is sparse, i.e., most of its entries equal 0. It turns out that the sparsity assumption can be used to accurately estimate $\boldsymbol{x}$ from $\boldsymbol{y}$, even when the number of possible atoms (and thus, the length of $\boldsymbol{x}$) is greater than the number of measurements in $\boldsymbol{y}$ [7, 12, 22]. This model has been used to great advantage in many fundamental fields of signal processing, including

compressed sensing [12, 84], denoising [73], deblurring [121], and interpolation [122].

The assumption of sparsity is an example of a much more general class of signal models which can be described as a union of subspaces [9, 123, 124]. Indeed, each support pattern defines a subspace of the space of possible parameter vectors. Saying that the parameter contains no more than k nonzero entries is equivalent to stating that $\boldsymbol{x}$ belongs to the union of all such subspaces. Unions of subspaces are proving to be a powerful generalization of the sparsity model. Apart from ordinary sparsity, unions of subspaces have been applied to estimate signals as diverse as pulse streams [36, 125], multi-band communications [126–128], and block sparse vectors [9, 26–28], the latter being the focus of this paper. The common thread running through these applications is the ability to exploit the union of subspaces structure in order to achieve accurate reconstruction of signals from a very low number of measurements.

The block sparsity model is based on the realization that in many practical sparse representation settings, not all support patterns are equally likely. Specifically, if a particular element of $\boldsymbol{x}$ is nonzero, then in many cases "similar" elements in $\boldsymbol{x}$ are also nonzero. The precise definition of similarity is context-dependent. For example, in Fourier-based dictionaries, neighboring frequency bins are often jointly nonzero, while in wavelet-based dictionaries, nonzero entries in a certain detail level are likely to be correlated with nonzeros in higher detail levels. Consequently, the sparsity model does not incorporate all of the structure present in the signal. The block sparsity approach aims to partially overcome this drawback by partitioning the vector $\boldsymbol{x}$ into blocks, each of which contains a small number of elements. The structure imposed by the block sparsity model is that no more than a small number k of blocks are nonzero. The model thus favors the use of related atoms, rather than sporadic dictionary columns. Consequently, block sparsity is well-suited for those situations described above, in which specific atoms tend to be used together.

The usefulness of a model depends on the existence of efficient and effective methods for estimating a signal $\boldsymbol{x}$ from its measurements. Fortunately, estimators designed for the ordinary sparsity model can be readily adapted to the block sparse setting. Thus, previous work has described techniques such as block orthogonal matching pursuit (BOMP) [28] and the mixed ℓ_2/ℓ_1-optimization (L-OPT) [9, 27], the latter being a block version of the Lasso. In this paper, we also describe a block-sparse version of the thresholding algorithm, which we refer to as block-thresholding (BTH). The BOMP and BTH approaches are representatives of a class of so-called greedy algorithms, which attempt to identify the support of $\boldsymbol{x}$ by choosing at each step the most likely candidate. In this paper we restrict attention to these greedy techniques,

which are simpler (and more naive) than convex relaxation techniques such as L-OPT, and are therefore more suitable for implementation in large-scale or computationally parsimonious settings.

Having described various estimation algorithms, it is natural to ask what can be guaranteed analytically about the performance of these methods in practice. For example, in the ordinary (non-block) sparsity setting, a rich collection of performance guarantees exists for various algorithms under different noise models. In particular, a distinction is made between adversarial and random noise models. In the former case, nothing is known about $\boldsymbol{w}$ except that it is bounded, $\|\boldsymbol{w}\|_2 \leq \varepsilon$; in particular, $\boldsymbol{w}$ might be chosen so as to maximally harm a given estimation algorithm. Consequently, guarantees in this case are relatively weak, ensuring only that the error in $\boldsymbol{x}$ is on the order of ε [7, 12, 22]. By contrast, when the noise is random, estimation performance is considerably improved for most noise realizations [8, 22, 78].

It is natural to seek an extension of these results to the block sparsity model. In the absence of noise, successful recovery of a block sparse parameter $\boldsymbol{x}$ from measurements $\boldsymbol{y} = \boldsymbol{D}\boldsymbol{x}$ has been demonstrated in the past for both BOMP and L-OPT [9, 28]. However, to the best of our knowledge, the only result providing analytical guarantees for a block sparse estimator under noise was given in [9], where the performance of L-OPT was analyzed under adversarial noise. The goal of this paper is to analyze the performance of the greedy algorithms BOMP and BTH under both adversarial and random noise models. As we will see, despite the fact that these greedy algorithms are simpler and more efficient to implement, their performance is close to the optimal achievable results.

Specifically, we first analyze the adversarial noise model, and show that both BOMP and BTH achieve an error on the order of ε when the noise is bounded by $\|\boldsymbol{w}\|_2 \leq \varepsilon$. These results generalize previous guarantees in several ways: First, when each block contains one element, we recover the non-block sparsity guarantee of Donoho et al. [7]. Second, when the noise bound ε equals 0, we obtain the noise-free guarantees of Eldar et al. [28].

We next turn to the random noise model, and examine in particular the case in which $\boldsymbol{w}$ is white Gaussian noise. We derive the Cramér–Rao bound (CRB) for estimating $\boldsymbol{x}$ from its measurements, and show that this bound equals the error of the "oracle estimator" which knows the locations of the nonzero blocks of $\boldsymbol{x}$. However, while the oracle estimator relies on information which is unavailable in practice, the CRB is known to be achievable by the maximum likelihood (ML) technique at high SNR. Unfortunately, the ML approach is NP-complete, and thus can probably not be implemented efficiently. Nevertheless, we proceed

to show that both BOMP and BTH come within a nearly constant factor of the CRB at high SNR, for dictionaries satisfying suitable requirements. Once again, when each block contains one element, we can recover previously known guarantees for non-block sparsity [78] from our results. Furthermore, we show that in typical block sparse situations, the performance guarantees of block algorithms is substantially better than that of non-block techniques.

The rest of this paper is organized as follows. The block sparse setting is defined in Section 7.2, and the BOMP and BTH techniques are described in Section 7.3. The adversarial noise model is then analyzed in Section 7.4. The treatment of random noise begins with the derivation of the CRB in Section 7.5, while performance guarantees for this case appear in Section 7.6. Finally, the guarantees and the CRB are compared with the actual performance of BOMP and BTH in a numerical study in Section 7.7.

7.2 Problem Setting

7.2.1 Notation

The following notation is used throughout the paper. Matrices and vectors are denoted by boldface uppercase letters $\boldsymbol{M}$ and boldface lowercase letters $\boldsymbol{v}$, respectively. The ℓ_2 norm of a vector $\boldsymbol{v}$ is $\|\boldsymbol{v}\|_2$ and the spectral norm of a matrix $\boldsymbol{M}$ is $\|\boldsymbol{M}\|$. The expectation of a random vector $\boldsymbol{v}$ will be denoted $E\{\boldsymbol{v}\}$ or, occasionally, $\mathbb{E}_{\boldsymbol{x}}\{\boldsymbol{v}\}$, where the subscript is intended to emphasize the fact that the expectation is a function of the deterministic quantity $\boldsymbol{x}$. The adjoint and the Moore–Penrose pseudoinverse of a matrix $\boldsymbol{M}$ are denoted, respectively, by $\boldsymbol{M}^*$ and $\boldsymbol{M}^\dagger$, while the column space of $\boldsymbol{M}$ is $\mathcal{R}(\boldsymbol{M})$. We denote by $\boldsymbol{v}[i]$ the ith d-element block of a vector $\boldsymbol{v}$ of length $N = Md$. Thus

$$\boldsymbol{v}[i] \triangleq [v_{(i-1)d+1}, v_{(i-1)d+2}, \ldots, v_{id}]^T, \quad 1 \leq i \leq M. \tag{7.1}$$

Consequently, we may write

$$\boldsymbol{v} = \left[\boldsymbol{v}^T[1], \ldots, \boldsymbol{v}^T[M]\right]^T. \tag{7.2}$$

Similarly, given a matrix $\boldsymbol{M}$ having N columns, the submatrix $\boldsymbol{M}[i]$ contains the columns $(i-1)d+1, (i-1)d+2, \ldots, id$ of $\boldsymbol{M}$, i.e., those columns of $\boldsymbol{M}$ which correspond to the ith block. The support $\mathrm{supp}(\boldsymbol{v})$ of $\boldsymbol{v}$ is defined as the set of indices of nonzero blocks of $\boldsymbol{v}$; formally

$$\mathrm{supp}(\boldsymbol{v}) \triangleq \{i : \boldsymbol{v}[i] \neq \boldsymbol{0}\}. \tag{7.3}$$

Given an index set I, the vector $\boldsymbol{v}_I$ is constructed as the subvector of $\boldsymbol{v}$ containing the blocks indexed by I; in other words, if $I = \{i_1, \ldots, i_p\}$, then

$$\boldsymbol{v}_I = \left[\boldsymbol{v}^T[i_1], \ldots, \boldsymbol{v}^T[i_p]\right]^T. \tag{7.4}$$

Likewise, the submatrix $\boldsymbol{M}_I$ contains the column blocks indexed by I, so that

$$\boldsymbol{M}_I = \left[\boldsymbol{M}[i_1], \ldots, \boldsymbol{M}[i_p]\right]. \tag{7.5}$$

To uniquely define $\boldsymbol{v}_I$ and $\boldsymbol{M}_I$, we will assume as a convention that the elements of I are sorted, i.e., $i_1 < i_2 < \cdots < i_p$.

7.2.2 Problem Definition

Let $\boldsymbol{x} \in \mathbb{C}^N$ be a deterministic block-sparse vector, i.e., $\boldsymbol{x}$ consists of M blocks $\boldsymbol{x}[1], \ldots, \boldsymbol{x}[M]$ of size d, of which at most k are nonzero [28]. The maximum support size k is assumed to be known. The block sparsity restriction can then be written as

$$\boldsymbol{x} \in \mathcal{X} \triangleq \{\boldsymbol{v} \in \mathbb{R}^N : |\operatorname{supp}(\boldsymbol{v})| \leq k\}. \tag{7.6}$$

For convenience, let $S \triangleq \operatorname{supp}(\boldsymbol{x})$ be the support of the parameter $\boldsymbol{x}$, and let $s = |S|$. Note the distinction between k and s: It is known that at most k blocks are nonzero, but the actual number of nonzero blocks s is unknown and may be smaller than k. In the sequel, it will be useful to define

$$\begin{aligned} |x_{\max}| &\triangleq \max_{i \in S} \|\boldsymbol{x}[i]\|_2, \\ |x_{\min}| &\triangleq \min_{i \in S} \|\boldsymbol{x}[i]\|_2. \end{aligned} \tag{7.7}$$

The block sparse model differs from the more common non-block sparsity setting: in the latter, it is assumed that a small number of entries (rather than blocks) in the vector $\boldsymbol{x}$ are nonzero. To emphasize this difference, we will occasionally refer to the non-block sparsity model as "ordinary" or "scalar" sparsity.

We are given noisy observations

$$\boldsymbol{y} = \boldsymbol{D}\boldsymbol{x} + \boldsymbol{w} \tag{7.8}$$

where $\boldsymbol{D} \in \mathbb{C}^{L \times N}$ is a known, deterministic dictionary, and $\boldsymbol{w}$ is a noise vector. Our goal is to estimate $\boldsymbol{x}$ from the measurements $\boldsymbol{y}$. It will be convenient to denote the ith column (or "atom") of $\boldsymbol{D}$ as $\boldsymbol{d}_i$. Thus we have

$$\boldsymbol{D} = [\underbrace{\boldsymbol{d}_1, \ldots, \boldsymbol{d}_d}_{\boldsymbol{D}[1]}, \underbrace{\boldsymbol{d}_{d+1}, \ldots, \boldsymbol{d}_{2d}}_{\boldsymbol{D}[2]}, \ldots, \underbrace{\boldsymbol{d}_{N-d+1}, \ldots, \boldsymbol{d}_N}_{\boldsymbol{D}[M]}]. \tag{7.9}$$

We assume for simplicity that the dictionary atoms are normalized, $\|\boldsymbol{d}_i\|_2 = 1$. We also assume that the measurement system is underdetermined, i.e., the number of measurements L is less than the number of parameters N; thus, we must utilize the structure $\mathcal{X}$, for otherwise we have no hope of recovering $\boldsymbol{x}$ from its measurements. Finally, we require that for any index set I of size $|I| \leq k$, the subdictionary $\boldsymbol{D}_I$ has full column rank. This latter assumption is needed to ensure that after a support set is chosen, one may estimate $\boldsymbol{x}$ using standard techniques for inverting an overcomplete set of linear equations, e.g., the least-squares approach.

We will provide performance guarantees for two separate noise models. First, we consider the adversarial setting, in which the noise is unknown but bounded,

$$\|\boldsymbol{w}\|_2 \leq \varepsilon \tag{7.10}$$

for a known constant $\varepsilon > 0$. In this case the goal is to provide performance guarantees which hold for all values of $\boldsymbol{w}$ satisfying (7.10). Second, we treat additive white Gaussian noise, in which

$$\boldsymbol{w} \sim N(\boldsymbol{0}, \sigma^2 \boldsymbol{I}). \tag{7.11}$$

In this case $\boldsymbol{w}$ is unbounded, and the goal will be to provide guarantees which hold with high probability.

Following [28], we define the block coherence of $\boldsymbol{D}$ as

$$\mu_B \triangleq \max_{i \neq j} \frac{1}{d} \|\boldsymbol{D}^*[i]\boldsymbol{D}[j]\|. \tag{7.12}$$

We also define the sub-coherence

$$\nu = \max_{1 \leq \ell \leq M} \max_{(\ell-1)d+1 \leq i \neq j \leq \ell d} |\boldsymbol{d}_i^* \boldsymbol{d}_j|. \tag{7.13}$$

The block coherence and sub-coherence are generalizations of the concept of the coherence, which is defined as

$$\mu = \max_{1 \leq i \neq j \leq N} |\boldsymbol{d}_i^* \boldsymbol{d}_j| \tag{7.14}$$

and applies to dictionaries regardless of whether they have a block structure.

7.3 Techniques for Block-Sparse Estimation

For reference and in order to fix notation, we now describe the two greedy algorithms for which we provide performance guarantees.

Block-Thresholding (BTH) We propose the following straightforward extension of the well-known thresholding algorithm. Given a measurement vector $\boldsymbol{y} \in \mathbb{C}^L$, perform the following steps:

1. Compute the correlations

$$\rho_i = \|\boldsymbol{D}^*[i]\boldsymbol{y}\|_2, \quad i = 1, \ldots, M. \tag{7.15}$$

2. Find the k largest correlations and denote their indices by $i_1, \ldots, i_k$. In other words, find a set of indices $\widehat{S} = \{i_1, \ldots, i_k\}$ such that $\rho_i \geq \rho_j$ for all $i \in \widehat{S}$ and $j \notin \widehat{S}$.

3. The reconstructed signal is given by

$$\hat{\boldsymbol{x}}_{\text{BTH}} = \underset{\widetilde{\boldsymbol{x}}:\text{supp}(\widetilde{\boldsymbol{x}})=\widehat{S}}{\arg\min} \|\boldsymbol{y} - \boldsymbol{D}\widetilde{\boldsymbol{x}}\|_2. \tag{7.16}$$

Block Orthogonal Matching Pursuit (BOMP) The BOMP algorithm, based on the OMP algorithm [17], was first proposed in [28].

Given a measurement vector $\boldsymbol{y} \in \mathbb{C}^L$, perform the following steps:

1. Define $\boldsymbol{r}^0 = \boldsymbol{y}$.

2. For each $\ell = 1, \ldots, k$, do the following:

   (a) Set

$$i_\ell = \arg\max_i \|\boldsymbol{D}^*[i]\boldsymbol{r}^{\ell-1}\|_2. \tag{7.17}$$

   (b) Set

$$\boldsymbol{x}^\ell = \underset{\widetilde{\boldsymbol{x}}:\text{supp}(\widetilde{\boldsymbol{x}})\subseteq\{i_1,\ldots,i_\ell\}}{\arg\min} \|\boldsymbol{y} - \boldsymbol{D}\widetilde{\boldsymbol{x}}\|_2. \tag{7.18}$$

   (c) Set $\boldsymbol{r}^\ell = \boldsymbol{y} - \boldsymbol{D}\boldsymbol{x}^\ell$.

3. The estimate is given by $\hat{\boldsymbol{x}}_{\text{BOMP}} = \boldsymbol{x}^k$.

Oracle Estimator We will find it useful to analyze the oracle estimator, which is defined as the least-squares solution within the true support set, i.e.,

$$\hat{\boldsymbol{x}}_{\text{or}} = \underset{\widetilde{\boldsymbol{x}}:\text{supp}(\widetilde{\boldsymbol{x}})\subseteq S}{\arg\min} \|\boldsymbol{x} - \widetilde{\boldsymbol{x}}\|_2^2. \tag{7.19}$$

Using the notation introduced above, we have

$$(\hat{\boldsymbol{x}}_{\text{or}})_S = (\boldsymbol{D}_S^* \boldsymbol{D}_S)^{-1} \boldsymbol{D}_S^* \boldsymbol{y},$$
$$(\hat{\boldsymbol{x}}_{\text{or}})_{S^C} = \boldsymbol{0} \tag{7.20}$$

where $S^C = \{1, \ldots, M\} \backslash S$ is the complement of the support set S. Note that the term "oracle estimator" is somewhat misleading, since $\hat{\boldsymbol{x}}_{\text{or}}$ relies on knowledge of the true support set S, and is therefore not a true estimator.

7.4 Guarantees for Adversarial Noise

We begin by stating our performance guarantees in the case of adversarial noise. The proofs of these results are quite technical and can be found in Appendix 7.A.

Theorem 7.1. *Consider the setting of Section 7.2 with adversarial noise* (7.10)*. Suppose that*

$$(1 - (d-1)\nu)|x_{\min}| > 2\varepsilon\sqrt{1 + (d-1)\nu} + (2k-1)d\mu_B|x_{\max}|. \tag{7.21}$$

Then, the BTH algorithm correctly identifies all elements of the support of $\boldsymbol{x}$*, and its error is bounded by*

$$\|\hat{\boldsymbol{x}}_{\text{BTH}} - \boldsymbol{x}\|_2^2 \leq \frac{\varepsilon^2}{1 - (d-1)\nu - (k-1)d\mu_B}. \tag{7.22}$$

Theorem 7.2. *Consider the setting of Section 7.2 with adversarial noise* (7.10)*. Suppose that*

$$(1 - (d-1)\nu)|x_{\min}| > 2\varepsilon\sqrt{1 + (d-1)\nu} + (2k-1)d\mu_B|x_{\min}|. \tag{7.23}$$

Then, the BOMP algorithm identifies all elements of $\text{supp}(\boldsymbol{x})$*, and its error is bounded by*

$$\|\hat{\boldsymbol{x}}_{\text{BOMP}} - \boldsymbol{x}\|_2^2 \leq \frac{\varepsilon^2}{1 - (d-1)\nu - (k-1)d\mu_B}. \tag{7.24}$$

The following remarks should be made concerning Theorems 7.1 and 7.2.

• *Scalar sparsity:* The scalar sparsity setting, in which $\boldsymbol{x}$ has no more than k nonzero elements, can be recovered by choosing $d = 1$. In this case, BOMP and BTH reduce to their scalar versions, which are called OMP and thresholding, respectively, and the block-coherence μ_B equals the coherence μ of (7.14). Theorems 7.1 and 7.2 then coincide with the well-known results of Donoho et al. [7] for performance of scalar sparse signals under adversarial noise. As an example (and for future reference), the OMP performance guarantee is given below.

Corollary 7.3 (Donoho et al. [7]). *Let $\boldsymbol{y} = \boldsymbol{D}\boldsymbol{x} + \boldsymbol{w}$ be a measurement vector of a signal $\boldsymbol{x}$ having sparsity $\|\boldsymbol{x}\|_0 \leq k$. Suppose that the coherence μ of the dictionary $\boldsymbol{D}$ satisfies*

$$|x_{\min}|(1 - (2k-1)\mu) > 2\varepsilon. \tag{7.25}$$

Then, OMP recovers the correct support pattern of $\boldsymbol{x}$ and achieves an error bounded by

$$\|\hat{\boldsymbol{x}}_{\text{OMP}} - \boldsymbol{x}\|_2^2 \leq \frac{\varepsilon^2}{1-(k-1)\mu}. \tag{7.26}$$

Note that in the case of ordinary sparsity, $d = 1$, and therefore $|x_{\min}|$ can be defined simply as the magnitude of the smallest nonzero element in $\boldsymbol{x}$.

• *Benefits and limitations of block sparsity:* It is interesting to compare the achievable performance guarantees when one utilizes the block-sparse structure, as opposed to merely using ordinary (scalar) sparsity information. For concreteness, we focus in this discussion on a comparison between OMP and BOMP, but identical conclusions can be drawn by comparing the thresholding algorithm with its block-sparse version BTH.

Consider a block sparse signal $\boldsymbol{x}$ as defined in Section 7.2. Such a signal can also be viewed as a scalar sparse signal of length $N = Md$, having no more than sd nonzero elements. It is readily shown that the coherence μ satisfies $\nu \leq \mu$ and $\mu_B \leq \mu$ [28]. Consequently,

$$\frac{\varepsilon^2}{1-(d-1)\nu-(k-1)d\mu_B} \leq \frac{\varepsilon^2}{1-(sd-1)\mu} \tag{7.27}$$

which implies that if the conditions for the performance guarantees of both BOMP and OMP hold, then the performance guarantee (7.24) for BOMP will be at least as good as that of OMP (7.26). Moreover, in typical block-sparse settings, both ν and μ_B will be substantially smaller than μ [28], and the guarantees for BOMP will then be considerably better.

These results notwithstanding, it should be noted that BOMP should not automatically be preferred over OMP in any setting. This is because the condition (7.23) of Theorem 7.2 can sometimes be weaker than that of OMP. Specifically, the factor $2\varepsilon\sqrt{1+(d-1)\nu}$ in (7.23) is larger than the analogous term 2ε in (7.25).[1] This implies that if the sub-coherence ν is large, block sparse algorithms will not perform as well as their scalar counterparts. Such a result is to be expected: Highly correlated dictionary blocks may cause noise amplification, and in such cases, it may be preferable to separately correlate each atom with the measurements, rather than relying on the combined correlation of the entire block. Indeed, it would be quite surprising if a partition of *any* dictionary $\boldsymbol{D}$ into arbitrary blocks could be shown to perform as well as

[1] The remaining terms in (7.23) are always no worse than the corresponding terms in (7.25).

a scalar sparsity algorithm, since the former adds a restriction on the possible support patterns of the vector $\boldsymbol{x}$. The lesson to be learned from this analysis is that block sparsity techniques are effective when the dictionary can be separated into blocks whose elements are orthogonal or nearly orthogonal.

• *Noiseless case:* The situation in which $\boldsymbol{y} = \boldsymbol{D}\boldsymbol{x}$, i.e., no noise is present in the system, has been previously analyzed in the context of block sparsity in [28]. This setting can be recovered by choosing the noise bound $\varepsilon = 0$. In this case, the condition (7.24) simplifies to

$$(d-1)\nu + (2k-1)d\mu_B < 1 \tag{7.28}$$

and Theorem 7.2 then amounts to a guarantee for perfect recovery of $\boldsymbol{x}$ if (7.28) holds. This result for the noise-free setting has been previously demonstrated in [28, Thm. 3].

Similarly, by substituting $\varepsilon = 0$ into Theorem 7.1, one obtains a perfect recovery condition for BTH in the noiseless setting. Specifically, if the condition

$$(d-1)\nu\frac{|x_{\max}|}{|x_{\min}|} + (2k-1)d\mu_B < 1 \tag{7.29}$$

is satisfied, then BTH correctly recovers $\boldsymbol{x}$ from its noiseless measurements $\boldsymbol{y} = \boldsymbol{D}\boldsymbol{x}$.

Since BTH is a much simpler algorithm than BOMP, it is not surprising that the necessary condition (7.29) for BTH is somewhat stronger than the corresponding condition (7.28) for BOMP. This difference between the conditions is indicative of the different strategies employed by the two techniques, and will be further discussed in Section 7.6.

• *Severity of the error:* As in the scalar sparsity scenario, the presence of adversarial noise severely limits the ability of any algorithm to perform denoising. This is evident from Theorems 7.1 and 7.2, which guarantee only that the distance between the estimates and the true value of $\boldsymbol{x}$ is on the order of the noise magnitude ε. Given our detailed knowledge of the structure of the signal $\boldsymbol{x}$, one would expect more powerful denoising capabilities for typical noise realizations. Consequently, in the remainder of this paper, we adopt the assumption of random noise, which cannot align itself so as to maximally interfere with the recovery algorithms.

7.5 The Cramér–Rao Bound

A central goal in assessing the quality of an estimator is to check its proximity to the best possible performance in the given setting. To this end, it is common practice to compute the CRB for unbiased estimators [16], i.e., those techniques $\hat{\boldsymbol{x}}$ for which the bias $\boldsymbol{b}(\boldsymbol{x}) \triangleq \mathbb{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}\} - \boldsymbol{x}$

equals zero. The CRB is a lower bound on the mean-squared error $\mathrm{MSE}(\hat{\boldsymbol{x}}, \boldsymbol{x}) = \mathbb{E}_{\boldsymbol{x}}\{\|\hat{\boldsymbol{x}} - \boldsymbol{x}\|_2^2\}$ for any unbiased estimator $\hat{\boldsymbol{x}}$.

To utilize the information inherent in the block sparsity structure, we apply the constrained CRB [52,56,58,70] to the present setting. In the constrained estimation scenario, one often seeks estimators which are unbiased for all parameter values in the constraint set [52,56]. However, as we will see below, this requirement is too strict in the block sparse setting. Indeed, in Theorem 7.4 we show that it is not possible to construct *any* method which is unbiased for all feasible parameter values. Consequently, a weaker, local definition of unbiasedness is called for, which we refer to as $\mathcal{X}$-unbiasedness [70].

Intuitively, an estimator $\hat{\boldsymbol{x}}$ is said to be $\mathcal{X}$-unbiased at a point $\boldsymbol{x} \in \mathcal{X}$ if $\mathbb{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}\} = \boldsymbol{x}$ holds at the point $\boldsymbol{x}$ and at all points $\widetilde{\boldsymbol{x}}$ in $\mathcal{X}$ which are sufficiently close to $\boldsymbol{x}$. To formally define $\mathcal{X}$-unbiasedness, we first recall the concept of a feasible direction. A vector $\boldsymbol{v} \in \mathbb{C}^N$ is said to be a feasible direction at $\boldsymbol{x}$ if, for any sufficiently small α, we have $\boldsymbol{x} + \alpha\boldsymbol{v} \in \mathcal{X}$. We then say that $\hat{\boldsymbol{x}}$ is $\mathcal{X}$-unbiased at $\boldsymbol{x}$ if $\mathbb{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}\} = \boldsymbol{x}$ and if

$$\left.\frac{\partial \boldsymbol{b}(\boldsymbol{x} + \alpha\boldsymbol{v})}{\partial \alpha}\right|_{\alpha=0} = 0 \tag{7.30}$$

for any feasible direction $\boldsymbol{v}$. In other words, the bias is zero at $\boldsymbol{x}$ and remains unchanged, up to a first-order approximation, when moving away from $\boldsymbol{x}$ along feasible directions. This definition yields the following result, whose proof can be found in Appendix 7.B.

Theorem 7.4 (Cramér–Rao bound for block-sparse signals)**.** *Consider the setting of Section 7.2 in which the block sparse parameter vector $\boldsymbol{x}$ is to be estimated from measurements corrupted by Gaussian noise* (7.11)*.*

(a) Suppose $\boldsymbol{x}$ contains fewer than k nonzero blocks, i.e., $s < k$. Then, no finite-variance estimator is $\mathcal{X}$-unbiased at $\boldsymbol{x}$.

(b) Suppose $\boldsymbol{x}$ contains precisely k nonzero blocks, i.e., $s = k$. Then, any estimator which is $\mathcal{X}$-unbiased at $\boldsymbol{x}$ satisfies

$$\mathrm{MSE}(\hat{\boldsymbol{x}}, \boldsymbol{x}) \geq \sigma^2 \,\mathrm{Tr}\left((\boldsymbol{D}_S^* \boldsymbol{D}_S)^{-1}\right). \tag{7.31}$$

We recall that both the MSE and the CRB are functions of the unknown vector $\boldsymbol{x}$, as is generally the case when estimating a deterministic parameter. It follows immediately from Theorem 7.4 that no finite-variance estimator can satisfy $\mathbb{E}_{\boldsymbol{x}}\{\hat{\boldsymbol{x}}\} = \boldsymbol{x}$ for all $\boldsymbol{x} \in \mathcal{X}$, which explains why we previously avoided this simpler definition of unbiasedness in the constrained

setting. Instead, restricting attention to a local unbiasedness requirement led to a finite CRB for almost all parameter values in $\boldsymbol{x}$: specifically, those parameters whose support is maximal, $|\operatorname{supp}(\boldsymbol{x})| \triangleq s = k$.

For maximal-support values of $\boldsymbol{x}$, it is not difficult to show that the CRB (7.31) coincides with the MSE of the oracle estimator (7.20). In this case it is possible to get a sense for the value of the bound, as follows. From (7.44) of Lemma 7.8 (see Appendix 7.A), we have that none of the eigenvalues of $(\boldsymbol{D}_S^*\boldsymbol{D}_S)^{-1}$ are larger than $1/(1-(d-1)\nu-(k-1)d\mu_B)$. Thus

$$\sigma^2 \operatorname{Tr}\left((\boldsymbol{D}_S^*\boldsymbol{D}_S)^{-1}\right) \leq \frac{1}{1-(d-1)\nu-(k-1)d\mu_B} kd\sigma^2. \tag{7.32}$$

In other words, when the block coherence and sub-coherence of $\boldsymbol{D}$ are low, the bound of Theorem 7.4 will be close to $kd\sigma^2$. This value is typically much lower than the total noise variance $E\{\|\boldsymbol{w}\|_2^2\} = L\sigma^2$. Thus, at least according to the CRB, it is possible to achieve substantial denoising in the presence of random noise. This stands in contrast to the rather disappointing guarantees presented for adversarial noise in the previous section. We may thus hope that the performance will be improved when considering random noise.

As opposed to the oracle estimator, which cannot be implemented in practice, it is well-known that the CRB can be asymptotically achieved at high SNR by the maximum likelihood (ML) estimator [16]. However, in the present setting, computing the ML estimator is NP-hard, and thus impractical. Consequently, it is of interest to determine whether there exist *efficient* techniques which come close to the performance bound (7.31), at least for high SNR values. As we will show in the next section, this question is answered in the affirmative: greedy block sparsity techniques do indeed approach the CRB for sufficiently high SNR.

7.6 Guarantees for Gaussian Noise

In this section, we analyze the performance of block sparse algorithms when the noise $\boldsymbol{w}$ is a Gaussian random variable having mean zero and covariance $\sigma^2\boldsymbol{I}$. Our main performance guarantees are summarized in Theorems 7.5 and 7.6. The proofs of these theorems are found in Appendix 7.C.

Theorem 7.5. *Consider the setting of Section 7.2 with additive white Gaussian noise $\boldsymbol{w} \sim N(\boldsymbol{0}, \sigma^2\boldsymbol{I})$. Suppose it is known that*

$$\begin{aligned}(1-(d-1)\nu)|x_{\min}| - (2k-1)d\mu_B|x_{\max}| \\ \geq 2\sigma\sqrt{2\alpha d(1+(d-1)\nu)\log N}\end{aligned} \tag{7.33}$$

for some constant $\alpha \geq 1/(2d \log N)$. *Then, with probability exceeding*

$$1 - \frac{0.8d(2\alpha d \log N)^{d/2-1}}{N^{\alpha d-1}} \tag{7.34}$$

the BTH algorithm identifies the correct support of $\boldsymbol{x}$ *and achieves an error bounded by*

$$\|\hat{\boldsymbol{x}}_{\text{BTH}} - \boldsymbol{x}\|_2^2 \leq \frac{2\alpha(1+(d-1)\nu)}{(1-(d-1)\nu-(k-1)d\mu_B)^2} dk\sigma^2 \log N. \tag{7.35}$$

Theorem 7.6. *Consider the setting of Section 7.2 with additive white Gaussian noise* $\boldsymbol{w} \sim N(\boldsymbol{0}, \sigma^2 \boldsymbol{I})$. *Suppose it is known that*

$$\begin{aligned}(1-(d-1)\nu)|x_{\min}| - (2k-1)d\mu_B|x_{\min}| \\ \geq 2\sigma\sqrt{2\alpha d(1+(d-1)\nu)\log N}\end{aligned} \tag{7.36}$$

for some constant $\alpha \geq 1/(2d \log N)$. *Then, with probability exceeding* (7.34)*, the BOMP algorithm identifies the correct support of* $\boldsymbol{x}$ *and achieves an error bounded by*

$$\|\hat{\boldsymbol{x}}_{\text{BOMP}} - \boldsymbol{x}\|_2^2 \leq \frac{2\alpha(1+(d-1)\nu)}{(1-(d-1)\nu-(k-1)d\mu_B)^2} dk\sigma^2 \log N. \tag{7.37}$$

We now provide some insights into the performance of block-sparse algorithms under random noise.

• *Random noise vs. adversarial noise:* As noted in Section 7.4, performance guarantees in the case of adversarial noise can ensure a recovery error on the order of the total noise magnitude. This is a result of the fact that the noise could, in principle, be concentrated in a single nonzero component of $\boldsymbol{x}$, whereupon it would be indistinguishable from the signal. However, for random noise, such an event is highly unlikely. Consequently, Theorems 7.5 and 7.6 provide much tighter performance guarantees: both theorems demonstrate that, with high probability, the estimation error is on the order of $dk\sigma^2 \log N$, i.e., within a constant times $\log N$ of the CRB presented in Section 7.5. Since the noise variance $E\{\|\boldsymbol{w}\|^2\}$ is given by $N\sigma^2$, and since typically $dk \log N \ll N$, we conclude that the block sparse algorithms have successfully removed a large portion of the noise, owing to the utilization of the union-of-subspaces structure.

• *BOMP vs. BTH:* Comparing Theorems 7.5 and 7.6 leads to an important insight concerning the advantage of the more sophisticated BOMP algorithm over its simpler counterpart. Indeed, the guarantee for BOMP requires condition (7.36), which basically states that $|x_{\min}|$ must be larger than a constant multiplied by the standard deviation of the noise. By contrast, for the BTH guarantee one requires the stronger condition (7.33), which can be interpreted as requiring $|x_{\min}|$ to be larger than a small constant times $|x_{\max}|$, plus another constant times the noise standard deviation.

To explain this difference, recall from Section 7.3 that the BTH approach relies on a single support-identification stage in which the blocks most highly correlated with the measurements are chosen as the estimated support set $\widehat{S}$. Thus, for BTH to correctly identify the support, each block in S must be sufficiently large in magnitude to overcome interference from the noise and from the remaining blocks. Condition (7.33) can therefore be interpreted as a requirement that the magnitude $|x_{\min}|$ of the smallest nonzero block must be larger than the sum of the interference from the large nonzero blocks (the $|x_{\max}|$ term) and the noise. By contrast, the BOMP algorithm iteratively identifies support elements, maintaining a residual vector $\boldsymbol{r}^\ell$ containing the components of the measurement vector which have yet to be identified. Thus, BOMP requires only the ability to separately isolate each nonzero block, and hence its weaker condition (7.36), which necessitates only that $|x_{\min}|$ be larger than the noise.

Finally, it should be noted that when BTH and BOMP both identify the correct support set, the estimates of the two algorithms coincide, explaining the identical bounds on their performance. The conclusion from this analysis is that BOMP should be preferred if a wide dynamic range of block magnitudes is possible, but that when all blocks have roughly the same size, the simpler and more efficient BTH technique can be used.

• *Scalar sparsity:* It is interesting to note that known results for scalar sparsity algorithms can be recovered from our block sparsity guarantees, by substituting $d = 1$ into Theorems 7.5 and 7.6. For example, consider the BOMP guarantee (Theorem 7.6). In the scalar case, this algorithm is known as OMP, and its performance guarantee can be written as follows.

Corollary 7.7. *Let $\boldsymbol{y} = \boldsymbol{D}\boldsymbol{x} + \boldsymbol{w}$ be a measurement vector of a signal $\boldsymbol{x}$ having sparsity $\|\boldsymbol{x}\|_0 \leq k$. Suppose the coherence μ of $\boldsymbol{D}$ satisfies*

$$|x_{\min}|(1 - (2k-1)\mu) \geq 2\sigma\sqrt{2\alpha \log N} \tag{7.38}$$

for some $\alpha > 1$. Then, with probability exceeding

$$1 - \frac{0.8/\sqrt{2}}{N^{\alpha-1}\sqrt{\alpha \log N}} \tag{7.39}$$

the OMP algorithm recovers the correct support of $\boldsymbol{x}$, and achieves an error bounded by

$$\|\hat{\boldsymbol{x}}_{\text{OMP}} - \boldsymbol{x}\|_2^2 \leq \frac{2\alpha}{(1-(k-1)\mu)^2} k\sigma^2 \log N. \tag{7.40}$$

Corollary 7.7 is nearly identical to [78, Thm. 4], with the only difference being that the constant $0.8/\sqrt{2} \approx 0.566$ in (7.39) is replaced in [78] with the slightly better constant $1/\sqrt{\pi} \approx 0.564$. This slight discrepancy can be resolved if the more accurate version (7.89a) of Lemma 7.11 is

Problem Dimensions				Coherence		OMP		Block-OMP		Cramér–Rao
M	d	L	k	μ	μ_B	Guar./σ^2	$\sigma_{\max}$	Guar./σ^2	$\sigma_{\max}$	CRB/σ^2
1200	5	3000	1	0.10	0.026	301.0	0.033	37.0	0.160	5.0
1200	5	3000	2	0.10	0.026	—	—	98.8	0.110	10.0
1200	5	3000	3	0.10	0.026	—	—	204.4	0.063	15.1
1200	5	3000	4	0.10	0.026	—	—	417.0	0.010	20.1
1200	5	3000	5	0.10	0.026	—	—	—	—	25.2
1200	5	3000	3	0.10	0.026	—	—	204.4	0.063	15.1
600	10	3000	3	0.10	0.015	—	—	364.3	0.049	30.2
300	20	3000	3	0.10	0.010	—	—	879.1	0.008	60.8
200	30	3000	3	0.10	0.007	—	—	—	—	91.8
1200	5	3000	1	0.10	0.026	301.0	0.033	37.0	0.160	5.0
1200	5	1000	1	0.17	0.043	—	—	37.0	0.144	5.0
1200	5	500	1	0.25	0.060	—	—	37.0	0.128	5.0
1200	5	100	1	0.51	0.133	—	—	37.0	0.062	5.0
1200	5	50	1	0.71	0.165	—	—	37.0	0.032	5.0
1200	5	20	1	0.90	0.197	—	—	37.0	0.003	5.0
1200	5	10	1	0.98	0.200	—	—	—	—	5.0

Table 7.1: Performance Guarantees for OMP and Block-OMP

used in the proof of Theorem 7.6, but the resulting expression becomes much more cumbersome in the block sparse case.

• *Block sparsity vs. scalar sparsity:* A legitimate question is whether the incorporation of the block sparsity structure substantially assists estimation algorithms. In other words, do the performance guarantees of the block algorithms BOMP and BTH compare favorably with the results achievable on identical signals using scalar sparsity algorithms, such as OMP and thresholding? This question is examined numerically in the next section.

7.7 Numerical Experiments

From a practical point of view, it is important to determine whether the use of block sparse algorithms contributes significantly to the performance of estimation algorithms. After all, any block sparse signal containing k nonzero blocks of size d can also be viewed as a sparse signal containing kd nonzero elements. Is there a significant benefit in using the block algorithms rather than the ordinary scalar versions?

There are two possible approaches to answering this question. First, one may compare the performance achieved in practice by block sparse and scalar sparse algorithms. This requires a complete specification of the problem setting, including a choice of the parameter value $\boldsymbol{x}$, which is unknown in practice. Alternatively, one can compare the performance guarantees for block sparse techniques, which were derived in Section 7.6, to the previously known guarantees for scalar approaches [109]. The performance guarantees apply to all parameter values having a specified sparsity level, and are therefore more general. However, there may be a gap between the guarantee and the performance observed in practice. In order to take advantage of both approaches, in the following we compare both the actual performance and the guarantees of the various algorithms discussed in this paper.

In our experiments, we used dictionaries containing orthonormal blocks. Such dictionaries were constructed by first generating a random $L \times N$ matrix containing IID, zero-mean Gaussian random variables, and then performing a Gram–Schmidt procedure separately on the columns of each block. As a first experiment, we generated a variety of such dictionaries, and computed their coherence μ and block coherence μ_B. (The sub-coherence of dictionaries generated in this manner is necessarily $\nu = 0$.) These values were used to compute performance guarantees for BOMP (using Theorem 7.6) and for OMP (using Corollary 7.7). We assumed throughout that the minimum norm $|x_{\min}|$ among nonzero blocks equals 1 and that

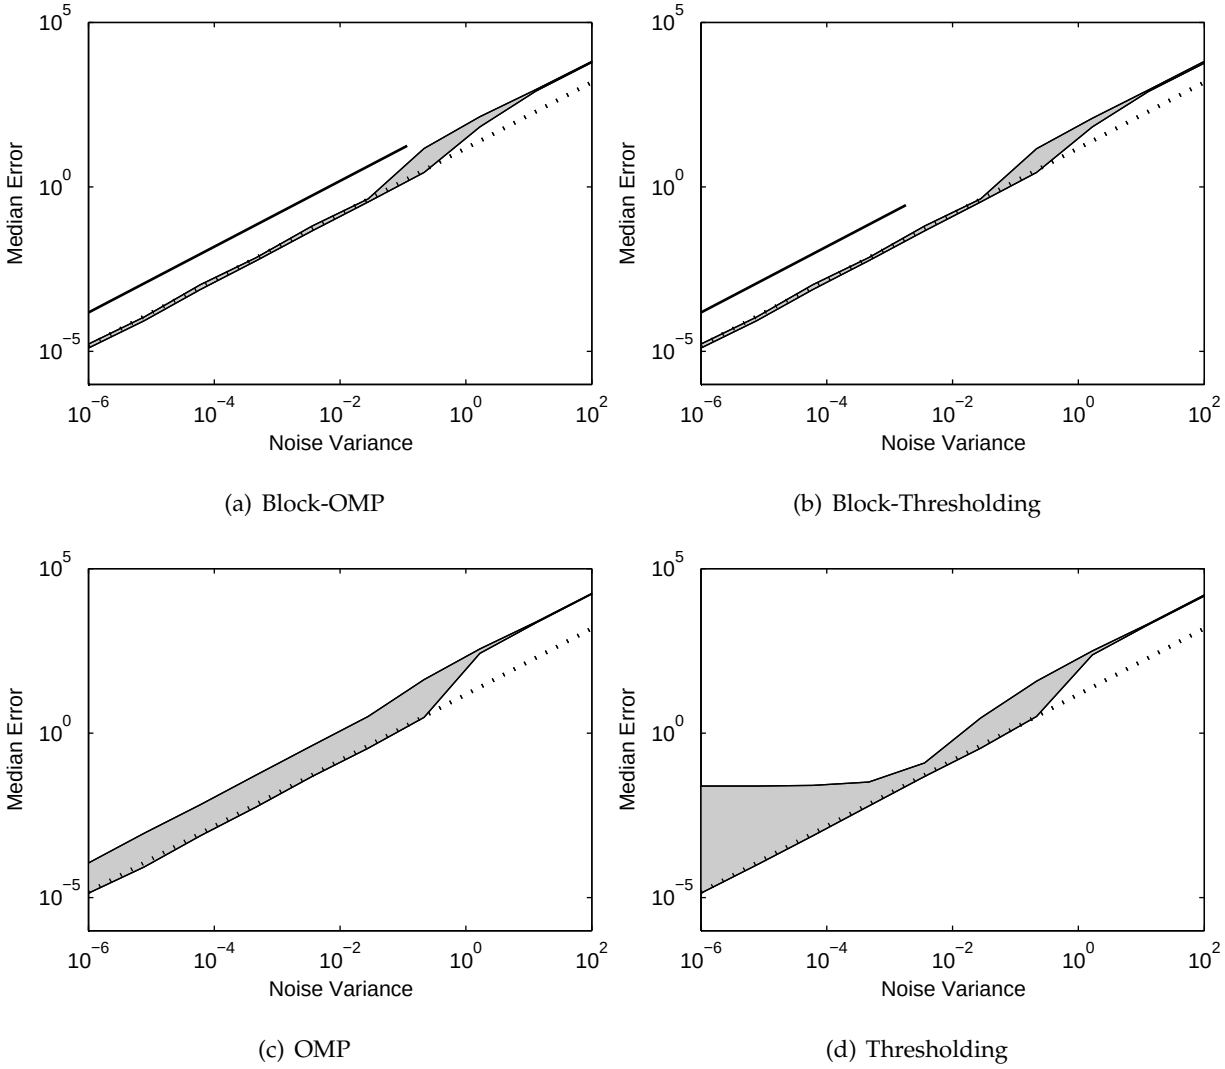

(a) Block-OMP

(b) Block-Thresholding

(c) OMP

(d) Thresholding

Figure 7.1: Median squared error as a function of the noise variance for block and scalar sparse estimation algorithms. The shaded region indicates the range of errors encountered for different parameter values. The dotted line plots the CRB. The thick solid line in Figs. 7.1(a) and 7.1(b) indicates the performance guarantees for the block sparse algorithms; no guarantee can be made for the scalar sparsity techniques in Figs. 7.1(c) and 7.1(d).

the minimum nonzero element equals $1/\sqrt{d}$. Some typical results are listed in Table 7.1. To compute the guarantees in this table, the smallest value of α yielding a 99% probability of success was chosen. The resulting guarantee is listed in multiples of σ^2. For example, a value of Guarantee$/\sigma^2 = 100$ means that $\|\hat{\boldsymbol{x}} - \boldsymbol{x}\|_2^2 \leq 100\sigma^2$ for 99% of the noise realizations. Also listed in Table 7.1 are the maximum noise standard deviations $\sigma_{\max}$ for which the performance guarantees still hold. A dash (—) indicates that no guarantee can be made for the given setting even in the noise-free case.

It is evident from Table 7.1 that the block sparse algorithm BOMP is guaranteed to perform over a much wider range of problem settings than the scalar OMP approach. Furthermore, even when performance guarantees are provided for both techniques, those for BOMP are substantially stronger. To provide merely one striking example from Table 7.1, note that 50 measurements suffice for BOMP to identify a signal composed of a single 5-element block among a set of 1200 possible blocks, whereas for OMP to identify such a signal at the same noise level, as many as 3000 measurements are required. The reason for this advantage is clear: the OMP algorithm must separately identify each nonzero component of the signal, and must therefore choose among a total of $\binom{1200}{5} \approx 2.1 \cdot 10^{13}$ possible support sets. This is obviously more challenging than identifying one nonzero block among a set of 1200 possibilities. Clearly, then, knowledge of a block-sparse structure can substantially improve performance if it is correctly utilized.

Table 7.1 also compares the performance guarantees with the CRB of Theorem 7.4. The CRB is listed for a random choice of support set S containing precisely k nonzero blocks; however, choosing different sets S only has a small effect on the value of the bound. The gap between these lower and upper bounds is not inconsiderable, and is typically on the order of a factor of 10. There are several reasons for this gap. First, the performance guarantees plotted above indicate an error which is obtained with 99% confidence, whereas the CRB is a bound on the MSE. By its very nature, the MSE averages out unusually disruptive noise realizations, and thus tends to be more optimistic. Second, different values of $\boldsymbol{x}$ may yield significantly different performance; the performance guarantees apply to *all* values of $\boldsymbol{x}$, whereas the CRB is plotted for a single, typical parameter value. Third, some loss of tightness undoubtedly results from the derivations of the theorems, i.e., there may still be room for improved bounds.

To measure the relative influence of these factors, we performed another experiment, in which the guarantees were compared with the actual performance of the various algorithms. To overcome the aforementioned pessimistic effect of a guarantee which holds with overwhelming

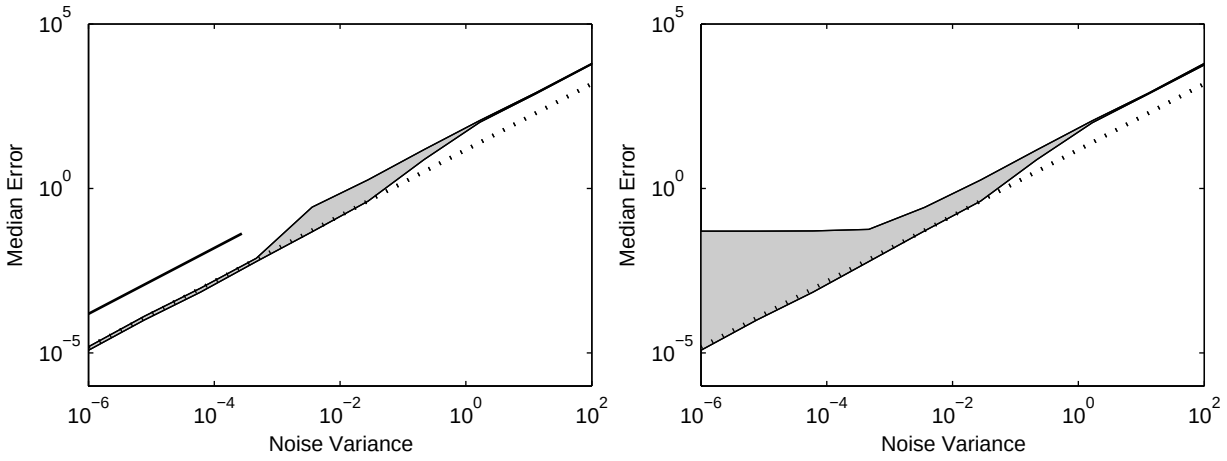

(a) Block-OMP (b) Block-Thresholding

Figure 7.2: Median squared error as a function of the noise variance for block sparse estimation algorithms. The shaded region indicates the range of errors encountered for different parameter values. The dotted line plots the CRB. The thick solid line in Fig. 7.2(a) indicates the performance guarantee for BOMP; no guarantee can be made for BTH. The deteriorated performance of BTH is a result of the existence of low-magnitude blocks.

probability, in this second experiment we computed guarantees with a 50% confidence level. In other words, these are assurances on the median of the distance between $\boldsymbol{x}$ and its estimate, which captures the typical estimation error. We also computed the actual median error of the various algorithms for a variety of parameter values.

The details of this experiment are as follows. We constructed a 3000×6000 dictionary $\boldsymbol{D}$ containing $M = 1200$ blocks of $d = 5$ atoms each, using the orthogonalization algorithm described above. The resulting coherence of $\boldsymbol{D}$ was $\mu = 0.094$, the block coherence was $\mu_B = 0.026$, and since each block was orthonormal, the sub-coherence was $\nu = 0$. We then constructed a variety of block sparse vectors $\boldsymbol{x}$, each having $s = 3$ nonzero blocks, with $|x_{\min}| = 2\sqrt{d}$ and $|x_{\max}| = 3\sqrt{d}$. We chose the parameter vectors so as to cover as wide a range of scenarios as possible, within the aforementioned requirements. For example, some parameter vectors contained a block with a single nonzero component whose value was $|x_{\max}|$, while other vectors contained a block with each of the d elements receiving a value of $|x_{\max}|/\sqrt{d}$. Although it is clearly not feasible to cover the full range of possible parameter vectors, it is hoped that in this way some sense is given of the variability in performance for different parameter values. Indeed, as shown below, different parameters often yield widely differing estimation errors.

For each choice of a parameter vector, 20 noise realizations were generated and the resulting measurement vector $\boldsymbol{y}$ was computed using (7.8). The BOMP, BTH, OMP, and thresholding algorithms were then applied to each of the measurement vectors. For every technique and each parameter vector, the median estimation error (among the noise realizations) was computed. The range of median estimation errors obtained for different choices of $\boldsymbol{x}$ is plotted as a shaded area in Fig. 7.1.

In the present setting, neither of the scalar sparsity algorithms was capable of providing a performance guarantee. For BOMP and BTH, performance guarantees were available, and these are plotted as a solid line in Fig. 7.1. These guarantees are valid only up to a certain maximal noise variance, at which point the solid line in Fig. 7.1 stops. The results are also compared with the CRB of Theorem 7.4. It should be emphasized that the CRB is a bound on the MSE, rather than the median error, although in practice the differences between these two quantities appear to be quite small. It is also worth recalling that the CRB is a bound on unbiased estimators, while all of the techniques discussed herein are biased; nevertheless, it is evident that the CRB still provides a rough measure of the optimal performance of the proposed algorithms.

Several comments are in order concerning Fig. 7.1. First, the performance of both block sparse algorithms exhibits a transition: near-CRB performance for low noise levels deteriorates substantially when the noise level crosses a certain threshold. This behavior qualitatively matches the predictions of the performance guarantees, which ensure support recovery and near-CRB performance for sufficiently low noise levels. The threshold at which this transition occurs is identified fairly accurately for BOMP, and less so for BTH, although it is possible that there exist some (untested) parameter values for which the BTH transition occurs at lower noise levels. However, the numeric value of the performance guarantee is somewhat pessimistic: while the observed performance is close to the CRB for all parameter values, analytically one can guarantee only that the median error will not be larger than approximately 10 times the CRB. This result is most likely due to the various inequalities employed in the proofs of Theorems 7.5 and 7.6. Indeed, since the correct support is identified with high probability for most noise realizations, the BTH and BOMP algorithms will likely tend to coincide with the oracle estimator, whose error equals that of the CRB. The question of formally proving such a claim remains a topic for further research.

The advantages of the block sparse approach become evident when compared with scalar sparsity algorithms (Figs. 7.1(c) and 7.1(d)). For the scalar techniques, no performance guar-

antees can be made in the present setting. Unlike the block sparsity algorithms, the scalar approaches fail to recover the correct parameter vector even when the noise is negligible, and for some parameter values, their error does not converge to the CRB. The thresholding algorithm, in particular, ceases to improve (for some parameter values) as the noise is reduced, while the OMP approach, although significantly better than thresholding, does not converge to the CRB as do the block sparse techniques. This demonstrates the advantages of utilizing the fact that the signal is known to have a block-sparse structure.

The performance of BOMP (Fig. 7.1(a)) is quite similar to that of BTH (Fig. 7.1(b)) in the experiment above. This is not surprising when one compares our problem setting with the guarantees of Section 7.6. Indeed, as we have seen, the primary difference between the BOMP and BTH algorithms is that the one-shot support estimation employed by BTH causes large-magnitude blocks to overshadow small-magnitude nonzero blocks. In the setting of Fig. 7.1, the range of magnitudes between $|x_{\max}| = 3\sqrt{d}$ and $|x_{\min}| = 2\sqrt{d}$ is not very large, and therefore BTH performs nearly as well as BOMP. The advantages of BOMP become readily apparent if one considers a wider dynamic range. This is illustrated in Fig. 7.2, in which the setup is identical to that of the previous experiment, except that parameter vectors having $|x_{\min}| = 0.1\sqrt{d}$ and $|x_{\max}| = \sqrt{d}$ were chosen, yielding a 10-fold dynamic range in the block magnitudes. In this case, while the guarantee for BOMP is hardly changed, the conditions for Theorem 7.5 no longer hold, so that nothing can be ensured concerning the BTH technique. Indeed, in Fig. 7.2 we see that BTH performs poorly for some parameter values even when the noise level is low, and its performance is no longer proportional to the CRB.

7.8 Conclusion

In this paper, we analyzed the performance of the greedy block algorithms BOMP and BTH under the adversarial and Gaussian noise models. In the adversarial setting $\|\boldsymbol{w}\|_2 \leq \varepsilon$, we showed that the estimation error equals a constant times the noise bound ε, which shows that performance in this case will not necessarily reduce the noise power. The situation is much better in the presence of random noise, where we saw that, under suitable conditions, greedy techniques obtain an error on the order of $dk\sigma^2 \log N$ with high probability; this is substantially lower than the input noise power $N\sigma^2$. Indeed, the BTH and BOMP algorithms come close to the CRB and the error of the oracle estimator.

There remain many open questions concerning the performance of block sparse techniques

under random noise. For example, for scalar sparsity, performance guarantees for convex relaxation techniques do not require assumptions on the SNR. An important challenge is to determine whether similar SNR-independent results can be demonstrated for block convex relaxation techniques such as L-OPT. Furthermore, it is well-known that scalar sparsity guarantees can be strengthened if the restricted isometry constants of the dictionary $\boldsymbol{D}$ are known, as is the case, for example, when $\boldsymbol{D}$ is chosen from an appropriate random ensemble. Thus, it is also of interest to provide guarantees for block techniques under random noise based on an extension of the RIP to the block sparse setting. One such extension has already been proposed in [9], and its application to the Gaussian noise model may provide tighter bounds for some performance algorithms.

7.A Proofs for Adversarial Noise

We begin by providing several lemmas which will prove useful for the analysis under both the adversarial and the Gaussian noise models.

Lemma 7.8. *Given a dictionary $\boldsymbol{D}$ having block coherence μ_B and sub-coherence ν, we have*

$$\|\boldsymbol{D}^*[i]\boldsymbol{D}[j]\| \leq d\mu_B \quad \text{for all } i \neq j \tag{7.41}$$

and

$$\|\boldsymbol{D}[i]\|^2 = \|\boldsymbol{D}^*[i]\boldsymbol{D}[i]\| \leq 1 + (d-1)\nu. \tag{7.42}$$

If $1-(d-1)\nu > 0$, then

$$\|(\boldsymbol{D}^*[i]\boldsymbol{D}[i])^{-1}\| \leq \frac{1}{1-(d-1)\nu}. \tag{7.43}$$

Suppose $1-(d-1)\nu-(k-1)d\mu_B > 0$ and let I be an index set with $|I| \leq k$. Then

$$\|(\boldsymbol{D}_I^*\boldsymbol{D}_I)^{-1}\| \leq \frac{1}{1-(d-1)\nu-(k-1)d\mu_B}. \tag{7.44}$$

Proof. The bound (7.41) follows directly from the definition (7.12) of block coherence. To prove (7.42)–(7.43), observe that the diagonal elements of the matrix $\boldsymbol{D}^*[i]\boldsymbol{D}[i]$ equal 1, while the off-diagonal elements are bounded in magnitude by ν. Therefore, by the Gershgorin circle theorem [116], all eigenvalues of $\boldsymbol{D}^*[i]\boldsymbol{D}[i]$ are in the range $[1-(d-1)\nu, 1+(d-1)\nu]$, demonstrating (7.42). Furthermore, it follows that the eigenvalues of $(\boldsymbol{D}^*[i]\boldsymbol{D}[i])^{-1}$ are in the range $[(1+(d-1)\nu)^{-1}, (1-(d-1)\nu)^{-1}]$, leading to (7.43).

It remains to prove (7.44). To this end, let $|I| = \ell \leq k$ and write $\boldsymbol{D}_I^* \boldsymbol{D}_I$ as

$$\boldsymbol{D}_I^* \boldsymbol{D}_I = \begin{pmatrix} \boldsymbol{M}[1,1] & \boldsymbol{M}[1,2] & \cdots & \boldsymbol{M}[1,\ell] \\ \boldsymbol{M}[2,1] & \boldsymbol{M}[2,2] & \cdots & \boldsymbol{M}[2,\ell] \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{M}[\ell,1] & \boldsymbol{M}[\ell,2] & \cdots & \boldsymbol{M}[\ell,\ell] \end{pmatrix} \tag{7.45}$$

where each $\boldsymbol{M}[i,j]$ is a $d \times d$ matrix containing the correlations between two blocks of dictionary atoms. From the definition of block coherence, we have

$$\|\boldsymbol{M}[i,j]\| \leq d\mu_B, \quad \text{for all } i \neq j. \tag{7.46}$$

By a generalization of the Gershgorin circle theorem [129, Thm. 2], it follows that all eigenvalues λ of $\boldsymbol{D}_I^* \boldsymbol{D}_I$ satisfy

$$\begin{aligned} \|\boldsymbol{M}[i,i] - \lambda \boldsymbol{I}\| \leq \sum_{j \neq i} \|\boldsymbol{M}[i,j]\| &\leq (\ell - 1) d\mu_B \\ &\leq (k-1) d\mu_B. \end{aligned} \tag{7.47}$$

Now, from the definition of sub-coherence, the off-diagonal elements of $\boldsymbol{M}[i,i]$ are no larger in magnitude than ν, while the diagonal elements of $\boldsymbol{M}[i,i]$ all equal 1. Therefore, by the Gershgorin circle theorem, given an arbitrary constant λ, all eigenvalues of the $d \times d$ matrix $\boldsymbol{M}[i,i] - \lambda \boldsymbol{I}$ are in the range $[1 - \lambda - (d-1)\nu, 1 - \lambda + (d-1)\nu]$. Consequently

$$\|\boldsymbol{M}[i,i] - \lambda \boldsymbol{I}\| \geq 1 - \lambda - (d-1)\nu. \tag{7.48}$$

Combining with (7.47) and rearranging, we conclude that all eigenvalues of $\boldsymbol{D}_I^* \boldsymbol{D}_I$ satisfy

$$\lambda \geq 1 - (d-1)\nu - (k-1) d\mu_B. \tag{7.49}$$

Consequently, the eigenvalues of $(\boldsymbol{D}_I^* \boldsymbol{D}_I)^{-1}$ are no larger than $(1 - (d-1)\nu - (k-1)d\mu_B)^{-1}$, establishing (7.44). □

Lemma 7.9. *Consider the setting of Section 7.2, and suppose it is known that*

$$\max_{1 \leq j \leq M} \|\boldsymbol{D}^*[j] \boldsymbol{w}\|_2 < \tau \tag{7.50}$$

for a given value $\tau > 0$. If the dictionary $\boldsymbol{D}$ satisfies

$$(1 - (d-1)\nu)\, |x_{\max}| > 2\tau + (2s-1) d\mu_B |x_{\max}| \tag{7.51}$$

then

$$\max_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 > \max_{j\notin S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 \tag{7.52}$$

where $S = \operatorname{supp}(\boldsymbol{x})$.

If (7.51) *is replaced by the stronger condition*

$$(1-(d-1)\nu)\,|x_{\min}| > 2\tau + (2s-1)d\mu_B|x_{\max}| \tag{7.53}$$

then

$$\min_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 > \max_{j\notin S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2. \tag{7.54}$$

Proof. The proof is an extension of [78, Lemma 3] to the block-sparse case, and is ultimately inspired by [7]. We first note that

$$\begin{aligned}\max_{j\notin S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 &= \max_{j\notin S} \left\| \boldsymbol{D}^*[j]\boldsymbol{w} + \sum_{i\in S} \boldsymbol{D}^*[j]\boldsymbol{D}[i]\boldsymbol{x}[i] \right\|_2 \\ &\leq \max_{j\notin S} \|\boldsymbol{D}^*[j]\boldsymbol{w}\|_2 + \max_{j\notin S} \sum_{i\in S} \|\boldsymbol{D}^*[j]\boldsymbol{D}[i]\|\,|x_{\max}|. \end{aligned} \tag{7.55}$$

By (7.50), the first term in (7.55) is smaller than τ. Together with (7.41), we obtain

$$\max_{j\notin S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 < \tau + sd\mu_B|x_{\max}| \leq \tau + kd\mu_B|x_{\max}|. \tag{7.56}$$

On the other hand,

$$\begin{aligned}\max_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 &= \max_{j\in S} \left\| \boldsymbol{D}^*[j]\boldsymbol{w} + \sum_{i\in S} \boldsymbol{D}^*[j]\boldsymbol{D}[i]\boldsymbol{x}[i] \right\|_2 \\ &\geq \max_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{D}[j]\boldsymbol{x}[j]\|_2 \\ &\quad - \max_{j\in S} \left\| \boldsymbol{D}^*[j]\boldsymbol{w} + \sum_{i\in S\setminus\{j\}} \boldsymbol{D}^*[j]\boldsymbol{D}[i]\boldsymbol{x}[i] \right\|_2 . \end{aligned} \tag{7.57}$$

As we have seen in the proof of Lemma 7.8, the eigenvalues of $\boldsymbol{D}^*[j]\boldsymbol{D}[j]$ are bounded in the range $[1-(d-1)\nu, 1+(d-1)\nu]$. Consequently

$$\begin{aligned}\max_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{D}[j]\boldsymbol{x}[j]\|_2 &\geq \max_{j\in S} (1-(d-1)\nu)\|\boldsymbol{x}[j]\|_2 \\ &= (1-(d-1)\nu)|x_{\max}|. \end{aligned} \tag{7.58}$$

Combining this result with (7.57), we have

$$\begin{aligned}\max_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 &\geq (1-(d-1)\nu)|x_{\max}| \\ &\quad - \max_{j\in S} \sum_{i\in S\setminus\{j\}} \|\boldsymbol{D}^*[j]\boldsymbol{D}[i]\boldsymbol{x}[i]\|_2 - \max_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{w}\|_2. \end{aligned} \tag{7.59}$$

Together with (7.50) and (7.41), this implies that

$$\begin{aligned}
&\max_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 \\
&\quad > (1-(d-1)\nu)|x_{\max}| - (k-1)|x_{\max}|d\mu_B - \tau \\
&\quad = (1-(d-1)\nu)|x_{\max}| - (2k-1)|x_{\max}|d\mu_B - 2\tau \\
&\qquad + k|x_{\max}|d\mu_B + \tau. \qquad (7.60)
\end{aligned}$$

Merging the results (7.56) and (7.60) yields

$$\begin{aligned}
\max_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 &> \max_{j\notin S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 \\
&+ (1-(d-1)\nu)|x_{\max}| - (2k-1)|x_{\max}|d\mu_B - 2\tau. \qquad (7.61)
\end{aligned}$$

Consequently, if (7.51) holds, then (7.52) follows, as required.

In a similar fashion, observe that

$$\begin{aligned}
\min_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 &= \min_{j\in S} \left\| \sum_{i\in S} \boldsymbol{D}^*[j]\boldsymbol{D}[i]\boldsymbol{x}[i] + \boldsymbol{D}^*[j]\boldsymbol{w} \right\|_2 \\
&\geq \min_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{D}[j]\boldsymbol{x}[j]\|_2 \\
&\quad - \max_{j\in S} \sum_{i\in S\setminus\{j\}} \|\boldsymbol{D}^*[j]\boldsymbol{D}[i]\boldsymbol{x}[i]\|_2 - \|\boldsymbol{D}^*[j]\boldsymbol{w}\|_2. \qquad (7.62)
\end{aligned}$$

As noted previously, all eigenvalues of $\boldsymbol{D}^*[j]\boldsymbol{D}[j]$ are larger than or equal to $1-(d-1)\nu$, and therefore

$$\min_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{D}[j]\boldsymbol{x}[j]\|_2 \geq (1-(d-1)\nu)|x_{\min}|. \qquad (7.63)$$

Furthermore, using (7.41) we have, for $i \neq j$,

$$\|\boldsymbol{D}^*[j]\boldsymbol{D}[i]\boldsymbol{x}[i]\|_2 \leq \|\boldsymbol{D}^*[j]\boldsymbol{D}[i]\|\,|x_{\max}| \leq d\mu_B|x_{\max}|. \qquad (7.64)$$

Substituting (7.50), (7.63), and (7.64) into (7.62) provides us with

$$\begin{aligned}
&\min_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 \\
&\quad > (1-(d-1)\nu)|x_{\min}| - (k-1)d\mu_B|x_{\max}| - \tau \\
&\quad = (1-(d-1)\nu)|x_{\min}| - (2k-1)d\mu_B|x_{\max}| - 2\tau \\
&\qquad + kd\mu_B|x_{\max}| + \tau. \qquad (7.65)
\end{aligned}$$

Finally, using (7.56) we obtain

$$\begin{aligned}
\min_{j\in S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 &> \max_{j\notin S} \|\boldsymbol{D}^*[j]\boldsymbol{y}\|_2 \\
&+ (1-(d-1)\nu)|x_{\min}| - (2k-1)d\mu_B|x_{\max}| - 2\tau. \qquad (7.66)
\end{aligned}$$

Therefore, if the condition (7.53) is satisfied, then (7.54) holds, completing the proof. □

We are now ready to prove Theorems 7.1 and 7.2.

Proof of Theorem 7.1. Using (7.10) and (7.42), we have for all j

$$\|\boldsymbol{D}^*[j]\boldsymbol{w}\|_2 \leq \|\boldsymbol{D}[j]\| \cdot \|\boldsymbol{w}\|_2 \leq \varepsilon\sqrt{1+(d-1)\nu}. \tag{7.67}$$

Thus, (7.50) holds with $\tau = \varepsilon\sqrt{1+(d-1)\nu}$.

In light of (7.21), the condition (7.53) for the second part of Lemma 7.9 holds, and therefore, by Lemma 7.9, we conclude that (7.54) holds. It follows that all blocks $\boldsymbol{D}[i]$ with $i \in S$ are more highly correlated than the off-support blocks $\boldsymbol{D}[i], i \notin S$. Thus, the estimated support $\widehat{S}$ contains the true support set S (with the possible addition of superfluous indices if $s < k$). It follows from the definition (7.16) of $\hat{\boldsymbol{x}}_{\text{BTH}}$ that $(\hat{\boldsymbol{x}}_{\text{BTH}})_{\widehat{S}} = D_{\widehat{S}}^{\dagger}\boldsymbol{y}$, and thus

$$\begin{aligned}
\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\text{BTH}}\|_2^2 &= \|\boldsymbol{x}_{\widehat{S}} - (\hat{\boldsymbol{x}}_{\text{BTH}})_{\widehat{S}}\|_2^2 \\
&= \|\boldsymbol{D}_{\widehat{S}}^{\dagger}\boldsymbol{D}_{\widehat{S}}\boldsymbol{x}_{\widehat{S}} - \boldsymbol{D}_{\widehat{S}}^{\dagger}\boldsymbol{y}\|_2^2 \\
&\leq \|\boldsymbol{D}_{\widehat{S}}^{\dagger}\|^2 \cdot \|\boldsymbol{y} - \boldsymbol{D}_{\widehat{S}}\boldsymbol{x}\|_2^2 \\
&= \|\boldsymbol{D}_{\widehat{S}}^{\dagger}\|^2 \cdot \|\boldsymbol{w}\|_2^2
\end{aligned} \tag{7.68}$$

where we have used the fact that $\boldsymbol{D}_{\widehat{S}}^{\dagger}\boldsymbol{D}_{\widehat{S}} = \boldsymbol{I}$, which follows from our assumption that $\boldsymbol{D}_I$ has full row rank for any set I of size s (see Section 7.2).

Since $|x_{\min}| \leq |x_{\max}|$, it follows from (7.21) that

$$1 - (d-1)\nu > (2k-1)d\mu_B. \tag{7.69}$$

Therefore, we may apply (7.44), yielding

$$\begin{aligned}
\|\boldsymbol{D}_{\widehat{S}}^{\dagger}\|^2 &= \|(\boldsymbol{D}_S^*\boldsymbol{D}_S)^{-1}\| \\
&\leq \frac{1}{1-(d-1)\nu-(k-1)d\mu_B}.
\end{aligned} \tag{7.70}$$

Combining this result with (7.68) and using (7.10), we obtain (7.22), as required. □

Proof of Theorem 7.2. As shown in the proof of Theorem 7.1, it follows from (7.10) that (7.50) holds with $\tau = \varepsilon\sqrt{1+(d-1)\nu}$. From (7.23) we then have

$$(1-(d-1)\nu)|x_{\min}| > 2\tau + (2k-1)d\mu_B|x_{\min}|. \tag{7.71}$$

Since $|x_{\max}| \geq |x_{\min}|$, this implies the condition (7.51) for the first part of Lemma 7.9. Thus, by Lemma 7.9, the dictionary block most highly correlated with $\boldsymbol{y}$ is a block within the support S

of $\boldsymbol{x}$. In other words, the first iteration in the BOMP algorithm correctly identifies an element within the support S.

The proof continues by induction. Assume we have reached the ℓth iteration with $2 \leq \ell \leq s$ and that all previous iterations have correctly identified elements of S. In other words, using the notation of Section 7.3, we have $i_1, \ldots, i_{\ell-1} \in S$.

By definition, we now have

$$\boldsymbol{r}^\ell = \boldsymbol{y} - \boldsymbol{D}\boldsymbol{x}^{\ell-1} = \boldsymbol{D}\widetilde{\boldsymbol{x}}^{\ell-1} + \boldsymbol{w} \tag{7.72}$$

where $\widetilde{\boldsymbol{x}}^{\ell-1} \triangleq \boldsymbol{x} - \boldsymbol{x}^{\ell-1}$ is the estimation error after $\ell - 1$ iterations. Since $\operatorname{supp}(\boldsymbol{x}) = S$ and, by induction, $\operatorname{supp}(\boldsymbol{x}^{\ell-1}) \subset S$, we have $\operatorname{supp}(\widetilde{\boldsymbol{x}}^{\ell-1}) \subset S$. Furthermore, $\ell - 1 < s$, so that $\operatorname{supp}(\boldsymbol{x}^{\ell-1})$ contains less than s elements, and is thus a strict subset of S. It follows that at least one nonzero block in $\widetilde{\boldsymbol{x}}^{\ell-1}$ is equal to the corresponding block in $\boldsymbol{x}$. Therefore

$$\max_j \|\widetilde{\boldsymbol{x}}^{\ell-1}[j]\|_2 \geq |x_{\min}|. \tag{7.73}$$

To summarize, by (7.72), $\boldsymbol{r}^\ell$ can be thought of as a noisy measurement of the block sparse vector $\widetilde{\boldsymbol{x}}^{\ell-1}$, which contains a block whose norm is at least $|x_{\min}|$. Using (7.73) and (7.23), we find that the condition (7.51) holds for this modified estimation problem. Consequently, by Lemma 7.9, we have

$$\max_{j \in S} \|\boldsymbol{D}^*[j]\boldsymbol{r}^{\ell-1}\|_2 > \max_{j \notin S} \|\boldsymbol{D}^*[j]\boldsymbol{r}^{\ell-1}\|_2. \tag{7.74}$$

Therefore, by (7.17), the ℓth iteration of the BOMP algorithm will choose an index i_ℓ belonging to the correct support set S, as long as $\ell \leq s$.

Since the BOMP algorithm never chooses the same support element twice, we conclude that precisely the s elements of S will be identified in the first s iterations. If $s < k$, then the remaining iterations will identify some additional elements not in S, so that ultimately the estimated support set $\widehat{S} = \{i_1, \ldots, i_k\}$ will satisfy $\widehat{S} \supseteq S$. The estimate $\hat{\boldsymbol{x}}_{\text{BOMP}}$ therefore satisfies $(\hat{\boldsymbol{x}}_{\text{BOMP}})_{\widehat{S}} = \boldsymbol{D}_{\widehat{S}}^\dagger \boldsymbol{y}$. Following the procedure (7.68)–(7.70) in the proof of Theorem 7.1, we obtain in an identical manner the required result (7.24). □

7.B Proof of Theorem 7.4

To compute the CRB, we must first determine the Fisher information matrix $\boldsymbol{J}(\boldsymbol{x})$ for estimating $\boldsymbol{x}$ from $\boldsymbol{y}$ of (7.8). This can be done using a standard formula [16, p. 85] and yields

$$\boldsymbol{J}(\boldsymbol{x}) = \frac{1}{\sigma^2}\boldsymbol{D}^*\boldsymbol{D}. \tag{7.75}$$

We now identify, for each $\boldsymbol{x} \in \mathcal{X}$, an orthonormal basis for the feasible direction subspace, which is defined as the smallest subspace of $\mathbb{C}^N$ containing all feasible directions at $\boldsymbol{x}$. To this end, denote by $\boldsymbol{e}_i$ the ith column of the $N \times N$ identity matrix. Consider first points $\boldsymbol{x} \in \mathcal{X}$ for which $s < k$. In other words, these are parameter values whose support S contains fewer than k elements. For such values of $\boldsymbol{x}$, we have, for any ε and any $1 \leq i \leq N$,

$$|\operatorname{supp}(\boldsymbol{x} + \varepsilon \boldsymbol{e}_i)| \leq |S| + 1 < k + 1 \leq k \tag{7.76}$$

and therefore $\boldsymbol{x} + \varepsilon \boldsymbol{e}_i \in \mathcal{X}$ for any ε and for any i. Consequently, the set of feasible directions at $\boldsymbol{x}$ includes $\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_N\}$, and the feasible direction subspace is therefore $\mathbb{C}^N$ itself. Thus, for values $\boldsymbol{x}$ containing fewer than k nonzero blocks, a convenient choice of a basis for the feasible direction subspace consists of the columns of the identity matrix.

Next, consider maximal-support parameter values, i.e., vectors $\boldsymbol{x}$ for which $s = k$. It is now no longer possible to add any vector $\boldsymbol{e}_i$ to $\boldsymbol{x}$ without violating the constraints. Indeed, it is not difficult to see that the only feasible directions are linear combinations of the unit vectors $\boldsymbol{e}_i$ for which i belongs to one of the blocks in S. These unit vectors can thus be chosen as a basis for the feasible direction subspace.

Let $\boldsymbol{U}(\boldsymbol{x})$ be a matrix whose columns comprise the chosen orthonormal basis for the feasible direction subspace at $\boldsymbol{x}$. Note that the dimensions of $\boldsymbol{U}(\boldsymbol{x})$ change with $\boldsymbol{x}$; specifically, $\boldsymbol{U}(\boldsymbol{x}) = \boldsymbol{I}_{N \times N}$ when $|S| < k$, and $\boldsymbol{U}(\boldsymbol{x})$ is an $N \times sd$ matrix otherwise. A necessary condition for a finite-variance $\mathcal{X}$-unbiased estimator to exist at a point $\boldsymbol{x}$ is [70, Thm. 1]

$$\mathcal{R}(\boldsymbol{U}(x)\boldsymbol{U}^*(x)) \subseteq \mathcal{R}(\boldsymbol{U}(\boldsymbol{x})\boldsymbol{U}^*(\boldsymbol{x})\boldsymbol{J}(\boldsymbol{x})\boldsymbol{U}(\boldsymbol{x})\boldsymbol{U}^*(\boldsymbol{x})). \tag{7.77}$$

When $s < k$, we have $\boldsymbol{U}(\boldsymbol{x}) = \boldsymbol{I}$. In this case, using (7.75), the condition (7.77) becomes

$$\mathbb{C}^N \subseteq \mathcal{R}(\boldsymbol{J}(\boldsymbol{x})) = \mathcal{R}(\boldsymbol{D}^*\boldsymbol{D}). \tag{7.78}$$

Since the dimensions of $\boldsymbol{D}$ are $L \times N$ with $L < N$, the rank of $\boldsymbol{D}^*\boldsymbol{D}$ is at most L, and thus $\mathcal{R}(\boldsymbol{D}^*\boldsymbol{D})$ cannot include the entire space $\mathbb{C}^N$. We conclude that in this case, (7.77) does not hold, and therefore no $\mathcal{X}$-unbiased estimator exists at points $\boldsymbol{x}$ for which $|S| < s$, proving part (a) of the theorem.

Let us now turn to maximal-support parameter values $\boldsymbol{x}$. As we have seen above, in this case the matrix $\boldsymbol{U}(\boldsymbol{x})$ consists of the columns $\boldsymbol{e}_i$ for which i is an element of a block within the support of $\boldsymbol{x}$. Therefore, the product $\boldsymbol{D}\boldsymbol{U}(\boldsymbol{x})$ selects those atoms of $\boldsymbol{D}$ belonging to blocks within S, i.e., $\boldsymbol{D}\boldsymbol{U}(\boldsymbol{x}) = \boldsymbol{D}_S$. Using (7.75), this leads to

$$\boldsymbol{U}^*(\boldsymbol{x})\boldsymbol{J}(\boldsymbol{x})\boldsymbol{U}(\boldsymbol{x}) = \frac{1}{\sigma^2}\boldsymbol{D}_S^*\boldsymbol{D}_S \tag{7.79}$$

which is invertible by assumption (see Section 7.2). It follows that the condition (7.77) holds for maximal-support parameters $\boldsymbol{x}$. One can therefore apply [70, Thm. 1], which states that for such values of $\boldsymbol{x}$,

$$\mathrm{MSE}(\hat{\boldsymbol{x}}, \boldsymbol{x}) \geq \mathrm{Tr}\Big(\boldsymbol{U}(\boldsymbol{x}) \left(\boldsymbol{U}^*(\boldsymbol{x}) \boldsymbol{J}(\boldsymbol{x}) \boldsymbol{U}(\boldsymbol{x})\right)^\dagger \boldsymbol{U}^*(\boldsymbol{x})\Big) . \tag{7.80}$$

Combining with (7.79) and using the fact that $\boldsymbol{U}^*(\boldsymbol{x})\boldsymbol{U}(\boldsymbol{x}) = \boldsymbol{I}$, we obtain (7.31), proving part (b) of the theorem.

7.C Proofs for Gaussian Noise

We begin with two lemmas which prove some useful properties of the Gaussian distribution. The first of these is a generalization of a result due to Šidák [120].

Lemma 7.10. *Let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_M$ be a set of M jointly Gaussian random vectors. Suppose that $E\{\boldsymbol{v}_i\} = \boldsymbol{0}$ for all i, but that the covariances of the vectors are unspecified and that the vectors are not necessarily independent. We then have*

$$\begin{aligned} &\Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1, \|\boldsymbol{v}_2\|_2 \leq c_2, \ldots, \|\boldsymbol{v}_M\|_2 \leq c_M\} \\ &\quad \geq \Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1\} \cdot \Pr\{\|\boldsymbol{v}_2\|_2 \leq c_2\} \cdots \\ &\qquad\qquad \cdots \Pr\{\|\boldsymbol{v}_M\|_2 \leq c_M\} . \end{aligned} \tag{7.81}$$

Proof. We will demonstrate that

$$\begin{aligned} &\Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1, \|\boldsymbol{v}_2\|_2 \leq c_2, \ldots, \|\boldsymbol{v}_M\|_2 \leq c_M\} \\ &\quad \geq \Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1\} \Pr\{\|\boldsymbol{v}_2\|_2 \leq c_2, \ldots, \|\boldsymbol{v}_M\|_2 \leq c_M\} . \end{aligned} \tag{7.82}$$

The result then follows by induction. For simplicity of notation, we will prove that (7.82) holds for the case $M = 2$; the general result can be shown in the same manner.

Denote by $f(\boldsymbol{v}_1|\boldsymbol{v}_2)$ the pdf of $\boldsymbol{v}_1$ conditioned on $\boldsymbol{v}_2$. Observe that, for a deterministic value $\boldsymbol{w}$, the pdf $f(\boldsymbol{v}_1|\boldsymbol{w})$ defines a Gaussian random vector whose mean depends linearly on $\boldsymbol{w}$, but whose covariance is constant in $\boldsymbol{w}$. Therefore, using a result due to Anderson [130], it follows that

$$\Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1 | \boldsymbol{v}_2 = \alpha \boldsymbol{w}\} = \int_{\|\boldsymbol{u}_1\|_2 \leq c_1} f(\boldsymbol{u}_1 | \alpha \boldsymbol{w}) d\boldsymbol{u} \tag{7.83}$$

is a non-increasing function of α.

Next, denoting by $f(\boldsymbol{v}_2)$ the marginal pdf of $\boldsymbol{v}_2$, we have

$$\begin{aligned} a(c_1, c_2) &\triangleq \Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1 \big| \, \|\boldsymbol{v}_2\|_2 \leq c_2\} \\ &= \frac{\int_{\|\boldsymbol{u}\|_2 \leq c_1} \int_{\|\boldsymbol{w}\|_2 \leq c_2} f(\boldsymbol{u}|\boldsymbol{w}) f(\boldsymbol{w}) \, d\boldsymbol{w} \, d\boldsymbol{u}}{\Pr\{\|\boldsymbol{v}_2\|_2 \leq c_2\}} \\ &= \frac{\int_{\|\boldsymbol{w}\|_2 \leq c_2} \Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1 | \boldsymbol{v}_2 = \boldsymbol{w}\} \, f(\boldsymbol{w}) \, d\boldsymbol{w}}{\int_{\|\boldsymbol{w}\|_2 \leq c_2} f(\boldsymbol{w}) \, d\boldsymbol{w}}. \end{aligned} \tag{7.84}$$

Thus, the function $a(c_1, c_2)$ is a weighted average of expressions of the form

$$\Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1 | \boldsymbol{v}_2 = \boldsymbol{w}\} \tag{7.85}$$

for values of $\boldsymbol{w}$ satisfying $\|\boldsymbol{w}\|_2 \leq c_2$. However, as we have shown, (7.85) is non-increasing in $\|\boldsymbol{w}\|_2$. Consequently, $a(c_1, c_2)$ is non-increasing in c_2.

On the other hand, observe that as $c_2 \to \infty$, the probability of the event $\|\boldsymbol{v}_2\|_2 \leq c_2$ converges 1. Thus we have

$$\lim_{c_2 \to \infty} a(c_1, c_2) = \Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1\}. \tag{7.86}$$

Combined with the fact that $a(c_1, c_2)$ is non-increasing in c_2, we find that

$$a(c_1, c_2) \geq \Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1\} \quad \text{for all } c_1, c_2. \tag{7.87}$$

Using the definition of $a(c_1, c_2)$ and applying Bayes's rule, we obtain

$$\begin{aligned} &\Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1, \|\boldsymbol{v}_2\|_2 \leq c_2\} \\ &\quad \geq \Pr\{\|\boldsymbol{v}_1\|_2 \leq c_1\} \Pr\{\|\boldsymbol{v}_2\|_2 \leq c_2\} \end{aligned} \tag{7.88}$$

and thus complete the proof. □

Our next lemma bounds the tail probability of the chi-squared distribution.

Lemma 7.11. *Let $\boldsymbol{u}$ be a d-dimensional Gaussian random vector having mean zero and covariance $\boldsymbol{I}$. Then, for any $t \geq 1$, we have*

$$\Pr\{\|\boldsymbol{u}\|_2^2 \geq t^2\} \leq \frac{(d-2)!!\lceil d/2 \rceil}{2^{d/2-1}\Gamma(d/2)} t^{d-2} e^{-t^2/2} \tag{7.89a}$$

$$\leq 0.8 d t^{d-2} e^{-t^2/2} \tag{7.89b}$$

where $\Gamma(z) \triangleq \int_0^\infty t^{z-1} e^{-t} dt$ is the Gamma function and

$$n!! \triangleq \prod_{0 \leq i < n/2} (n - 2i) \tag{7.90}$$

is the double factorial operator.

Of the two bounds provided in (7.89), the first is somewhat tighter, but obviously more cumbersome. For analytical tractability, we will use the latter bound in the sequel.

Proof of Lemma 7.11. The expression $\|\boldsymbol{u}\|_2^2$ is distributed as a chi-squared random variable with d degrees of freedom. Therefore, its tail probability is given by [119, §16.3]

$$\Pr\{\|\boldsymbol{u}\|_2^2 \geq t^2\} = \frac{\Gamma(d/2, t^2/2)}{\Gamma(d/2)} \tag{7.91}$$

where $\Gamma(a,z)$ is the incomplete Gamma function $\Gamma(a,z) \triangleq \int_z^\infty t^{a-1}\, e^{-t}\, dt$. It follows from the series expansion of $\Gamma(a,z)$ that [131, §6.5.32]

$$\begin{aligned}\Gamma\left(\frac{d}{2}, \frac{t^2}{2}\right) \leq \frac{e^{-t^2/2}}{2^{d/2-1}t^2}\big[&t^d + (d-2)t^{d-2} \\ &+ (d-2)(d-4)t^{d-4} + \cdots + (d-2)!!t^m\big]\end{aligned} \tag{7.92}$$

where $m = 1$ when d is odd and $m = 2$ when d is even. Note that (7.92) holds with equality for even d, but the inequality is strict for odd d. Since $t \geq 1$, we can enlarge each of the terms in the square brackets in (7.92) by replacing it with $(d-2)!!t^d$. The total number of terms in brackets is $\lceil d/2 \rceil$, yielding

$$\Gamma\left(\frac{d}{2}, \frac{t^2}{2}\right) \leq \frac{e^{-t^2/2}}{2^{d/2-1}} t^{d-2}(d-2)!! \left\lceil \frac{d}{2} \right\rceil. \tag{7.93}$$

Substituting into (7.91) demonstrates (7.89a).

To prove (7.89b), we distinguish between even and odd values of d. Assume first that d is even and denote $d = 2p$. We then have

$$\Gamma(d/2) = \Gamma(p) = (p-1)! \tag{7.94}$$

and

$$(d-2)!! = (2p-2)!! = 2^{p-1}(p-1)!. \tag{7.95}$$

Substituting these values into (7.89a) and simplifying yields

$$\Pr\{\|\boldsymbol{u}\|_2^2 \geq t^2\} \leq \frac{d}{2} t^{d-2} e^{-t^2/2} \tag{7.96}$$

which clearly satisfies (7.89b).

Similarly, assume that d is odd and write $d = 2p+1$. Substituting the formula

$$\Gamma(d/2) = \Gamma(p+1/2) = \frac{(2p-1)!!\sqrt{\pi}}{2^p} \tag{7.97}$$

into (7.89a), we obtain

$$\Pr\{\|\boldsymbol{u}\|_2^2 \geq t^2\} \leq \sqrt{\frac{2}{\pi}}\, \frac{d+1}{2} t^{d-2} e^{-t^2/2}. \tag{7.98}$$

It is easily verified that

$$\sqrt{\frac{2}{\pi}}\,\frac{d+1}{2} \leq 0.8d \quad \text{for all } d \geq 1. \tag{7.99}$$

Substituting back into (7.98) yields the required result. □

Our next result applies more specifically to the block sparse estimation setting. Following [22,78], we consider the event

$$B = \left\{ \max_{1 \leq i \leq M} \|\boldsymbol{D}^*[i]\boldsymbol{w}\|_2^2 \leq \tau^2 \right\} \tag{7.100}$$

where

$$\tau^2 = 2d\sigma\alpha(1 + (d-1)\nu)\log N \tag{7.101}$$

for a given $\alpha > 1/(2d\log N)$. We then have the following lemma.

Lemma 7.12. *Under the setting of Section 7.2, assume that $\boldsymbol{w}$ is a Gaussian random vector with mean zero and covariance $\sigma^2\boldsymbol{I}$. Then, the probability of the event B of* (7.100) *is bounded by*

$$\Pr\{B\} \geq 1 - \frac{0.8(2\alpha d\log N)^{d/2-1}}{N^{\alpha d-1}}. \tag{7.102}$$

Proof. Observe that $\boldsymbol{D}^*[i]\boldsymbol{w}$ is a d-dimensional Gaussian random vector with mean zero and covariance $\sigma^2\boldsymbol{D}^*[i]\boldsymbol{D}[i]$. Therefore, the random vector

$$\boldsymbol{u} = \frac{1}{\sigma}(\boldsymbol{D}^*[i]\boldsymbol{D}[i])^{-1/2}\boldsymbol{D}^*[i]\boldsymbol{w} \tag{7.103}$$

is a d-dimensional Gaussian random vector with mean zero and covariance $\boldsymbol{I}$. We thus have

$$\begin{aligned}
\Pr\{\|\boldsymbol{D}^*[i]\boldsymbol{w}\|_2^2 \leq \tau^2\} &= \Pr\Big\{\sigma^2\|(\boldsymbol{D}^*[i]\boldsymbol{D}[i])^{1/2}\boldsymbol{u}\|_2^2 \leq \tau^2\Big\} \\
&\geq \Pr\{\sigma^2\|\boldsymbol{D}^*[i]\boldsymbol{D}[i]\| \cdot \|\boldsymbol{u}\|_2^2 \leq \tau^2\} \\
&\geq \Pr\left\{\|\boldsymbol{u}\|_2^2 \leq \frac{\tau^2}{\sigma^2(1+(d-1)\nu)}\right\}
\end{aligned} \tag{7.104}$$

where, in the last step, we used (7.42). Using Lemma 7.11 and substituting the value (7.101) of τ^2, we obtain

$$\Pr\{\|\boldsymbol{D}^*[i]\boldsymbol{w}\|_2^2 \leq \tau^2\} \geq 1 - \eta \tag{7.105}$$

where

$$\begin{aligned}
\eta &\triangleq 1 - 0.8d(2\alpha d\log N)^{d/2-1}\exp(-d\alpha\log N) \\
&= 1 - \frac{0.8d(2\alpha d\log N)^{d/2-1}}{N^{\alpha d}}.
\end{aligned} \tag{7.106}$$

Using Lemma 7.10, we have

$$\begin{aligned}\Pr\{B\} &\geq \prod_{i=1}^{M} \Pr\big\{ \|\boldsymbol{D}^*[i]\boldsymbol{w}\|_2^2 \leq \tau^2 \big\} \\ &= (1-\eta)^M. \end{aligned} \tag{7.107}$$

When $\eta > 1$, the bound (7.102) is meaningless and the theorem holds vacuously. Otherwise, when $\eta \leq 1$, we have

$$\Pr\{B\} \geq 1 - M\eta \tag{7.108}$$

where we used the fact that $(1-\eta)^M \geq 1 - M\eta$ whenever $\eta \leq 1$ and $M \geq 1$. Substituting the value of η from (7.106) and recalling that $N = Md$ yields the required result. □

We are now ready to prove Theorems 7.5 and 7.6.

Proof of Theorem 7.5. By Lemma 7.12, the event B of (7.100) occurs with probability exceeding (7.34). Furthermore, using (7.33), it follows from Lemma 7.9 that under the event B, all blocks in the correct support set S are more highly correlated with $\boldsymbol{y}$ than the off-support blocks. Consequently, when B occurs, we have $S \subseteq \widehat{S}$, where $\widehat{S}$ is the support estimated by the BTH algorithm. Note, however, that the estimated set $\widehat{S}$ will contain additional blocks not in S if $s < k$. It follows that

$$\begin{aligned}\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\text{BTH}}\|_2^2 &= \|\boldsymbol{x}_{\widehat{S}} - (\hat{\boldsymbol{x}}_{\text{BTH}})_{\widehat{S}}\|_2^2 \\ &= \|\boldsymbol{D}_{\widehat{S}}^{\dagger}\boldsymbol{D}_{\widehat{S}}\boldsymbol{x}_{\widehat{S}} - \boldsymbol{D}_{\widehat{S}}^{\dagger}\boldsymbol{y}\|_2^2 \\ &\leq \|(\boldsymbol{D}_{\widehat{S}}^*\boldsymbol{D}_{\widehat{S}})^{-1}\|^2 \cdot \|\boldsymbol{D}_{\widehat{S}}^*\boldsymbol{w}\|_2^2 \\ &\leq \|(\boldsymbol{D}_{\widehat{S}}^*\boldsymbol{D}_{\widehat{S}})^{-1}\|^2 \cdot \sum_{i\in\widehat{S}} \|\boldsymbol{D}^*[i]\boldsymbol{w}\|_2^2 \end{aligned} \tag{7.109}$$

where we have used the fact that $\boldsymbol{D}_{\widehat{S}}^{\dagger}\boldsymbol{D}_{\widehat{S}} = \boldsymbol{I}$, which is a consequence of the assumption that $\boldsymbol{D}_{\widehat{S}}$ has full row rank (see Section 7.2). Using (7.44) and (7.100), we have that when B occurs

$$\|\boldsymbol{x} - \hat{\boldsymbol{x}}_{\text{BTH}}\|_2^2 \leq \frac{k\tau^2}{(1-(d-1)\nu-(k-1)d\mu_B)^2}. \tag{7.110}$$

Substituting the value (7.101) of τ yields the required result (7.35). □

Proof of Theorem 7.6. It follows from Lemma 7.12 that the event B occurs with probability exceeding (7.34). Our goal in this proof will thus be to show that, if B does occur, then the BOMP algorithm correctly identifies all elements of the support S of $\boldsymbol{x}$ (although some off-support elements may be identified as well if $s < k$). The remainder of the proof will then follow the steps of the proof of Theorem 7.5.

To demonstrate that the correct support is recovered, we begin by analyzing the first iteration of the BOMP algorithm. This iteration chooses a block i_1 having maximal correlation $\|\boldsymbol{D}^*[i_1]\boldsymbol{y}\|_2$ with the measurements $\boldsymbol{y}$. Now, since $|x_{\max}| \geq |x_{\min}|$, the condition (7.36) implies (7.51), with τ given by (7.101). Consequently, by Lemma 7.9, under the event B we find that the first iteration of BOMP identifies an element i_1 in the correct support set S.

To show that the next $s-1$ iterations of the BOMP algorithm also identify support elements, we proceed by induction. Specifically, assume that $\ell - 1 < s$ iterations have correctly identified elements $i_1, \ldots, i_{\ell-1}$, all of which are in the support set S. As in the proof of Theorem 7.2, define the estimation error after $\ell - 1$ iterations as $\widetilde{\boldsymbol{x}}^{\ell-1} \triangleq x - x^{\ell-1}$. By the induction hypothesis, $\operatorname{supp}(\widetilde{\boldsymbol{x}}) \subset S$, and clearly $\operatorname{supp}(\boldsymbol{x}) = S$. Thus $\operatorname{supp}(\widetilde{\boldsymbol{x}}) \subset S$, i.e., the support of $\widetilde{\boldsymbol{x}}$ is a strict subset of S. Using the same arguments as in the proof of Theorem 7.2, we find that $\widetilde{\boldsymbol{x}}^{\ell-1}$ contains a block whose norm is at least $|x_{\min}|$. Therefore, we can consider a modified estimation problem, in which $\boldsymbol{r}^\ell$ is a noisy measurement vector of the block sparse signal $\widetilde{\boldsymbol{x}}^{\ell-1}$. Together with (7.36), this implies that (7.51) holds for the modified setting. Therefore, by (7.52), the block in $\boldsymbol{r}^\ell$ having maximal correlation with the measurements is an element of S. Consequently, BOMP will correctly identify a support element in the ℓth iteration. Since the BOMP algorithm never selects a previously chosen support element, we find by induction that the support set S will be identified in full after s iterations. If $s < k$, then the remaining $k - s$ iterations will identify arbitrary off-support elements.

Denoting by $\widehat{S}$ the complete k-element support set identified by the BOMP approach, we thus have $S \subseteq \widehat{S}$. Following the technique (7.109)–(7.110) used in the proof of Theorem 7.5 thus yields the required result (7.37). □

Chapter 8

Performance of Finite Rate of Innovation Signals

This chapter has been submitted for publication as:

- Z. Ben-Haim, T. Michaeli, and Y. C. Eldar, "Performance bounds and design criteria for estimating finite rate of innovation signals," submitted to *IEEE Trans. Information Theory*, Sep. 2010.

8.1 Introduction

The field of digital signal processing hinges on the availability of techniques for sampling analog signals, thus converting them to discrete measurements. The sampling mechanism aims to preserve the information present in the analog domain, ideally permitting flawless recovery of the original signal. For example, one may wish to recover a continuous-time signal $x(t)$ from a discrete set of samples. The archetypical manifestation of this concept is the Shannon sampling theorem, which states that a B-bandlimited function can be reconstructed from samples taken at the Nyquist rate $2B$ [32].

Recently, considerable attention has been devoted to the extension of sampling theory to functions having a finite rate of innovation (FRI). These are signals determined by a finite number ρ of parameters per time unit [30]. Such a definition encompasses a rich variety of signals, including splines, shift-invariant signals, multiband signals, and pulse streams. In many FRI settings, several existing algorithms are guaranteed to recover the signal $x(t)$ from samples taken at rate ρ [30,31,34,36,37,125]. In other words, signals which correspond to the FRI model

can be reconstructed from samples taken at the rate of innovation, which is potentially much lower than their Nyquist rate.

The set of signals described by an FRI model can often be viewed as a union of subspaces [9, 34, 36, 125]. For example, consider a stream of pulses parameterized by pulse locations and amplitudes. The set of all pulses having a given location is a subspace of the space of continuous-time functions. Thus, the set of all pulses having arbitrary locations is a union of such subspaces. As we will see, this point of view yields a flexible and productive framework for understanding the types of constraints implied by the model.

Real-world signals are often contaminated by continuous-time noise and thus do not conform precisely to the FRI model. Furthermore, like any mathematical model, the FRI framework is an approximation which does not precisely hold in practical scenarios, an effect known as mismodeling error [9]. It is therefore of interest to quantify the effect of noise and mismodeling errors on FRI techniques [31, 34, 38, 125]. In the noisy case, it is no longer possible to perfectly recover the original signal from its samples. Nevertheless, one might hope for an appropriate finite-rate technique which achieves the best possible estimation accuracy, in the sense that increasing the sampling rate confers no further performance benefits. For example, to recover a B-bandlimited signal contaminated by continuous-time white noise, one can use an ideal low-pass filter with cutoff B prior to sampling at a rate of $2B$. This strategy removes all noise components with frequencies larger than B, while leaving all signal components intact. Consequently, any alternative method which does not zero out frequencies above B can be improved upon, whereas methods which zero out some of the signal frequencies can suffer from an arbitrarily large reconstruction error. Thus, sampling at a rate of $2B$ is indeed optimal in the case of a B-bandlimited signal, if the signal is corrupted by continuous-time noise prior to sampling. Sampling at a rate higher than $2B$ can be beneficial only when the sampling process itself introduces additional noise into the system, e.g., as a result of quantization.

By contrast, empirical observations indicate that, for some noisy FRI signals, substantial performance improvements are achievable when the sampling rate is increased beyond the rate of innovation [34, 37]. Thus, in some cases, there appears to be a fundamental difference between the noiseless and noise-corrupted settings, in terms of the required sampling rate. Our first goal in this paper will be to provide an analytical justification and quantification of these empirical findings. As we will see, the fact that oversampling improves performance is not merely indicative of flaws in existing algorithms; rather, it is a consequence of the inherent difficulty of reconstructing FRI signals under noise. Indeed, we will demonstrate that for

some FRI signals, unless considerable oversampling is employed, performance will necessarily deteriorate by several orders of magnitude relative to the optimal achievable reconstruction capability. Such effects occur even when the noise level is exceedingly low. Our analysis will also enable us to identify and characterize the types of signals for which oversampling is necessary.

To demonstrate these results, we first derive the Cramér–Rao bound (CRB) for estimating a finite-duration segment of an FRI signal $x(t)$ directly from continuous-time measurements $y(t) = x(t) + w(t)$, where $w(t)$ is a Gaussian white noise process. This yields a lower bound on the accuracy whereby $x(t)$ can be recovered by any technique, regardless of its sampling rate. This setting is to be distinguished from previous bounds in the FRI literature [38, 132] in three respects. First and most importantly, the measurements are a continuous-time process $y(t)$ and the bound therefore applies regardless of the sampling method. Second, in our model, the noise is added prior to sampling. Thus, as will be shown below, even sampling at an arbitrarily high rate will not completely compensate for the noise. Third, we bound the mean-squared error (MSE) in estimating $x(t)$ and not the parameters defining it, since we seek to determine the accuracy with which $x(t)$ itself can be recovered. Such a bound does not depend on the specific parametrization of the signal, and consequently, possesses a simpler analytical expression.

In practice, rather than processing the continuous-time signal $y(t)$, it is typically desired to estimate $x(t)$ from a discrete set of samples $\{c_n\}$ of $y(t)$. In this scenario, in addition to the continuous-time noise $w(t)$, digital noise may arise from the sampling process itself, for example due to quantization. To quantify the extent to which sampling degrades the ability to recover the signal, we next derive the CRB for estimating $x(t)$ from the measurements $\{c_n\}$. This analysis depends on the relative power of the two noise factors. When only digital noise is present, oversampling can be used to completely overcome its effect. On the other hand, when there exists only continuous-time noise, the bound converges to the continuous-time CRB as the sampling rate increases. In some cases, these bounds coincide at a finite sampling rate, which implies that the sampling scheme has captured all of the information present in the continuous-time signal, and any further increase in the sampling rate is useless. Conversely, when the continuous-time and sampled CRBs differ, the gap between these bounds is indicative of the degree to which information is lost in the sampling process. Our technique can then be used to plot the best possible performance as a function of the sampling rate, and thus provide the practitioner with a tool for evaluating the benefits of oversampling.

When a certain sampling technique achieves the performance of continuous-time measurements, it can be identified using the method described above. However, in some cases no such

technique exists, or the sampling rate it requires may be prohibitive. In these cases, it is desirable to determine the optimal sampling scheme having an allowed rate. Since different signals are likely to perform successfully with different sampling kernels, a Bayesian or average-case analysis is well-suited for this problem. Specifically, we assume that the signal $x(t)$ has a known prior distribution over the class of signals, and determine the linear sampling and reconstruction technique which minimizes the MSE for recovering $x(t)$ from its measurements. While nonlinear reconstruction techniques are commonly used and typically outperform the best linear estimator, this approach provides a simple means for identifying an appropriate sampling method. The resulting method can then be used in conjunction with standard nonlinear FRI algorithms.

We demonstrate our results via the problem of estimating a finite-duration sequence of pulses having unknown positions and amplitudes [30, 34, 36, 37]. In this case, a simple sufficient condition is obtained for the existence of a sampling scheme whose performance bound coincides with the continuous-time CRB. This scheme is based on sampling the Fourier coefficients of the pulse shape, and is reminiscent of recent time-delay estimation algorithms [34]. However, while the sampling scheme is theoretically sufficient for optimal recovery of $x(t)$, we show that in some cases there is room for substantial improvement in the reconstruction stage of these algorithms. Finally, we demonstrate that the Fourier domain is also optimal (in the sense of minimizing the reconstruction MSE) when the sampling budget is limited. Specifically, given an allowed number of samples N, the reconstruction MSE is minimized by sampling the N highest-variance Fourier coefficients of the signal $x(t)$.

The rest of this paper is organized as follows. The problem setting is defined in Section 8.2, and some examples of signals conforming to this model are presented in Section 8.3. We then briefly summarize our main results in Section 8.4. In Section 8.5, we provide a technical generalization of the CRB to general spaces. This result is used to obtain bounds on the achievable reconstruction error from continuous-time measurements (Section 8.6) and using a sampling mechanism (Section 8.7). Next, in Section 8.8 a Bayesian viewpoint is introduced and utilized to determine the optimal sampling kernels having a given rate budget. The results are demonstrated for the specific signal model of time-delay estimation in Section 8.9.

8.2 Definitions

8.2.1 Notation

The following notation is used throughout the paper. A boldface lowercase letter $\boldsymbol{v}$ denotes a vector, while a boldface uppercase letter $\boldsymbol{M}$ denotes a matrix. $\boldsymbol{I}_N$ is the $N \times N$ identity matrix. For a vector $\boldsymbol{v}$, the notation $\|\boldsymbol{v}\|$ indicates the Euclidean norm. Given a complex number $z \in \mathbb{C}$, the symbols z^* and $\Re\{z\}$ denote the complex conjugate and the real part of z, respectively. For an operator P, the range space and null space are $\mathcal{R}(P)$ and $\mathcal{N}(P)$, respectively, while the trace and adjoint are denoted, respectively, by $\mathrm{Tr}(P)$ and P^*. The Kronecker delta, denoted by $\delta_{m,n}$, equals 1 when $m = n$ and 0 otherwise. The expectation of a random variable v is written as $E\{v\}$.

The Hilbert space of square-integrable complex-valued functions over $[0, T_0]$ is denoted $L_2[0, T_0]$ or simply L_2. The corresponding inner product is

$$\langle f, g \rangle \triangleq \int_0^{T_0} f(t) g^*(t) dt \tag{8.1}$$

and the induced norm is $\|f\|_{L_2}^2 \triangleq \langle f, f \rangle$. For an ordered set of K functions $g_1, \ldots, g_K$ in L_2, we define the associated *set transformation* $G : \mathbb{C}^K \to L_2$ as

$$(G\boldsymbol{v})(t) = \sum_{k=1}^{K} \boldsymbol{v}_k g_k(t). \tag{8.2}$$

By the definition of the adjoint, it follows that

$$G^* f = \left(\langle f, g_1 \rangle, \ldots, \langle f, g_K \rangle \right)^T. \tag{8.3}$$

8.2.2 Setting

In this work, we are interested in the problem of estimating FRI signals from noisy measurements. To define FRI signals formally, let the T_0-local number of degrees of freedom $N_{T_0}(t)$ of a signal $x(t)$ at time t be the number of parameters defining the segment $\{x(t) : t \in [t - T_0/2, t + T_0/2]\}$. The T_0-local rate of innovation of $x(t)$ is then defined as [30]

$$\rho_{T_0} = \max_{t \in \mathbb{R}} \frac{N_{T_0}(t)}{T_0}. \tag{8.4}$$

We then say that $x(t)$ is an FRI signal if ρ_{T_0} is finite for all sufficiently large values of T_0. In Section 8.3, we will give several examples of FRI signals and compute their rates of innovation.

For concreteness, let us focus on the problem of estimating the finite-duration segment $\{x(t) : t \in [0, T_0]\}$, for some constant T_0, and let $K \triangleq N_{T_0}(T_0/2)$ denote the number of parameters defining this segment. We then have

$$x \in \mathcal{X} \triangleq \{h_{\boldsymbol{\theta}} \in L_2[0, T_0] : \boldsymbol{\theta} \in \Theta\} \tag{8.5}$$

where $h_{\boldsymbol{\theta}}$ is a set of functions parameterized by the vector $\boldsymbol{\theta}$, and Θ is an open subset of $\mathbb{R}^K$.

We wish to examine the random process

$$y(t) = x(t) + w(t), \quad t \in [0, T_0] \tag{8.6}$$

where $w(t)$ is continuous-time white Gaussian noise. Recall that formally, it is not possible to define Gaussian white noise over a continuous-time probability space [133]. Instead, we interpret (8.6) as a simplified notation for the equivalent set of measurements

$$z(t) = \int_0^t x(\tau)d\tau + \sigma_c b(t), \quad t \in [0, T_0] \tag{8.7}$$

where $b(t)$ is a standard Wiener process (also called Brownian motion) [134]. It follows that $w(t)$ can be considered as a random process such that, for any $f, g \in L_2$, the inner products $a = \langle f, w\rangle$ and $b = \langle g, w\rangle$ are zero-mean jointly Gaussian random variables satisfying $E\{ab^*\} = \sigma_c^2 \langle f, g\rangle$ [133]. The subscript c in σ_c is meant as a reminder of the fact that $w(t)$ is continuous-time noise. By contrast, when examining samples of the random process $y(t)$, we will also consider digital noise which is added during the sampling process.

In this paper, we consider estimators which are functions either of the entire continuous-time process (8.6) or of some subset of the information present in (8.6), such as a discrete set of samples of $y(t)$. To treat these two cases in a unified way, let $(\Omega, \mathscr{F})$ be a measurable space and let $\{P_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \Theta\}$ be a family of probability measures over $(\Omega, \mathscr{F})$. Let $(\mathcal{Y}, \mathscr{U})$ be a measurable space, and let the random variable $y : \Omega \to \mathcal{Y}$ denote the measurements. This random variable can represent either $y(t)$ itself or samples of this quantity.

An estimator can be defined in this general setting as a measurable function $\hat{x} : \mathcal{Y} \to L_2$. The MSE of an estimator $\hat{x}$ at x is defined as

$$\mathrm{MSE}(\hat{x}, x) \triangleq E\{\|\hat{x} - x\|_{L_2}^2\} = E\left\{\int_0^{T_0} |\hat{x}(t) - x(t)|^2 dt\right\}. \tag{8.8}$$

An estimator $\hat{x}$ is said to be unbiased if

$$E\{\hat{x}(t)\} = x(t) \text{ for all } x \in \mathcal{X} \text{ and almost all } t \in [0, T_0]. \tag{8.9}$$

In the next section, we demonstrate the applicability of our model by reviewing several scenarios which can be formulated using the FRI framework. Some of these settings will also be used in the sequel to exemplify our theoretical results.

8.3 Types of FRI Signals

Numerous FRI signal structures have been proposed and analyzed in the sampling literature. Whereas most of these can be treated within our framework, some FRI structures do not conform exactly to our problem setting. Thus, before delving into the derivation of the CRB, we first provide examples for scenarios that can be analyzed via our model and discuss some of its limitations.

8.3.1 Shift-Invariant Spaces

Consider the class of signals that can be expressed as

$$x(t) = \sum_{m\in\mathbb{Z}} a[m]g(t - mT) \tag{8.10}$$

with some arbitrary square-integrable sequence $\{a[m]\}_{m\in\mathbb{Z}}$, where $g(t)$ is a given pulse in $L_2(\mathbb{R})$ and $T > 0$ is a given scalar. This set of signals is a linear subspace of $L_2(\mathbb{R})$, which is often termed a *shift-invariant* (SI) space [135, 136]. The class of functions that can be represented in the form (8.10) is quite large. For example, choosing $g(t) = \mathrm{sinc}(t/T)$ leads to the subspace of π/T-bandlimited signals. Other important examples include the space of spline functions (obtained by letting $g(t)$ be a B-spline function) and communication signals such as pulse-amplitude modulation (PAM) and quadrature amplitude modulation (QAM). Reconstruction in SI spaces from noiseless samples has been addressed in [137, 138] and extended to the noisy setting in [139–141].

Intuitively, every signal lying in a SI space with spacing T has one degree of freedom per T seconds (corresponding to one coefficient from the sequence $\{a[m]\}$). It is thus tempting to regard the rate of innovation of such signals as $1/T$. However, this is only true in an asymptotic sense and for compactly supported pulses $g(t)$. For any finite window size T_0, the T_0-local rate of innovation ρ_{T_0} is generally larger. Specifically, suppose that the support of $g(t)$ is contained in $[t_a, t_b]$ and consider intervals of the form $[t, t + MT]$, where M is an integer. Then, due to the overlaps of the pulses, for any such interval we can only assure that there are no more than $M + \lceil (t_b - t_a)/T \rceil$ coefficients affecting the values of $x(t)$. Thus, the MT-local rate of

innovation of signals of the form (8.10) is given by

$$\rho_{MT} = \frac{1}{T}\left(1 + \frac{\lceil \frac{t_b - t_a}{T} \rceil}{M}\right). \tag{8.11}$$

In particular, signals of the form (8.10) having a generator $g(t)$ which is not compactly supported have an infinite T_0-local rate of innovation, for any finite T_0. This is the case, for example, with bandlimited signals, which are therefore not FRI functions under our definition. As will be discussed in the sequel, this is not a flaw of the definition we use for the rate of innovation. Rather, it reflects the fact that it is impossible to recover any finite-duration segment $[T_1, T_2]$ of such signals from a finite number of measurements.

8.3.2 Nonlinearly-Distorted Shift-Invariant Spaces

In certain communication scenarios, nonlinearities are introduced in order to avoid amplitude clipping, an operation known as companding [142]. When the original signal lies in a SI space, the resulting transmission takes the form

$$x(t) = r\left(\sum_{m\in\mathbb{Z}} a[m]g(t - mT)\right), \tag{8.12}$$

where $r(\cdot)$ is a nonlinear, invertible function. Clearly, the MT-local rate of innovation ρ_{MT} of this type of signals is the same as that of the underlying SI function, and is thus given by (8.11). The recovery of nonlinearly distorted SI signals from noiseless samples was treated in [142–145]. We are not aware of research works treating the noisy case.

8.3.3 Union of Subspaces

Much of the FRI literature treats signal classes which are unions of subspaces [9, 36, 125, 126]. We now give examples of a few of these models.

Finite Union of Subspaces

There are various situations in which a continuous-time signal is known to belong to one of a finite set of spaces. One such signal model is described by

$$x(t) = \sum_{m\in\mathbb{Z}}\sum_{k=1}^{K} a_k[m]g_k(t - mT), \tag{8.13}$$

where $\{g_k(t)\}_{k=1}^{K}$ are a set of generators. In this model it is assumed that only $L < K$ out of the K sequences $\{a_1[m]\}_{m\in\mathbb{Z}}, \ldots, \{a_K[m]\}_{m\in\mathbb{Z}}$ are not identically zero [83]. Therefore, the signal

$x(t)$ is known to reside in one of $\binom{K}{L}$ spaces, each of which is spanned by an L-element subset of the set of generators $\{g_k(t)\}_{k=1}^K$. This class of functions can be used to describe multiband signals [126, 127]. However, the discrete nature of these models precludes analysis using the differential tools employed in the remainder of this paper. Therefore, in this work we will focus on infinite unions of subspaces.

Single-Burst Channel Sounding

In certain medium identification and channel sounding scenarios, the echoes of a transmitted pulse $g(t)$ are analyzed to identify the positions and reflectance coefficients of scatterers in the medium [34, 146]. In these cases, the received signal has the form

$$x(t) = \sum_{\ell=1}^{L} a_\ell g(t - t_\ell), \tag{8.14}$$

where L is the number of scatterers and the amplitudes $\{a_\ell\}_{\ell=1}^L$ and time-delays $\{t_\ell\}_{\ell=1}^L$ correspond to the reflectance and location of the scatterers. Such signals can be thought of as belonging to a union of subspaces, where the parameters $\{t_\ell\}_{\ell=1}^L$ determine an L-dimensional subspace, and the coefficients $\{a_\ell\}_{\ell=1}^L$ describe the position within the subspace. In contrast with the previous example, however, in this setting we have a union of an infinite number of subspaces, since there are infinitely many possible values for the parameters $t_1, \ldots, t_L$.

In this case, for any window of size $T_0 > \max_\ell\{t_\ell\} - \min_\ell\{t_\ell\}$, the T_0-local rate of innovation is given by

$$\rho_{T_0} = \frac{2L}{T_0}. \tag{8.15}$$

Periodic Channel Sounding

Occasionally, channel sounding techniques consist of repeatedly probing the medium [147]. Assuming the medium does not change throughout the experiment, the result is a periodic signal

$$x(t) = \sum_{m\in\mathbb{Z}} \sum_{\ell=1}^{L} a_\ell g(t - t_\ell - mT). \tag{8.16}$$

As before, the set $\mathcal{X}$ of feasible signals is an infinite union of finite-dimensional subspaces in which $\{t_\ell\}_{\ell=1}^L$ determine the subspace and $\{a_\ell\}_{\ell=1}^L$ define the position within the subspace. The T_0-local rate of innovation in this case coincides with (8.15).

Semi-Periodic Channel Sounding

There are situations in which a channel consists of L paths whose amplitudes change rapidly, but the time delays can be assumed constant throughout the duration of the experiment [35,36, 147]. In these cases, the output of a channel sounding experiment will have the form

$$x(t) = \sum_{m\in\mathbb{Z}} \sum_{\ell=1}^{L} a_\ell[m] g(t - t_\ell - mT), \tag{8.17}$$

where $a_\ell[m]$ is the amplitude of the ℓth path at the mth probing experiment. This is, once again, a union of subspaces, but here each subspace is infinite-dimensional, as it is determined by the infinite set of parameters $\{a_\ell[m]\}$. In this case, the MT-local rate of innovation can be shown to be

$$\rho_{MT} = \frac{L}{T}\left(1 + \frac{1 + \lceil \frac{t_b - t_a}{T} \rceil}{M}\right). \tag{8.18}$$

Multiband Signals

Multiuser communication channels are often characterized by a small number of utilized subbands interspersed by large unused frequency bands [127]. The resulting signal can be described as

$$x(t) = \sum_{n\in\mathbb{Z}} \sum_{\ell=1}^{L} a_\ell[n] g(t - nT) e^{j\omega_\ell t}, \tag{8.19}$$

where $\{a_\ell[n]\}_{n\in\mathbb{Z}}$ is the data transmitted by the ℓth user, and ω_ℓ is the corresponding carrier frequency. In some cases the transmission frequencies are unknown [126, 127], resulting in an infinite union of infinite-dimensional subspaces. This setting is analogous in many respects to the semi-periodic channel sounding case; in particular, the MT-local rate of innovation can be shown to be the same as that given by (8.18).

8.4 Summary of Main Results

Before delving into the mathematical details, we provide in this section a high-level overview of our main contributions and summarize the resulting conclusions.

The overarching objective of this paper is to design and analyze sampling schemes for reconstructing FRI signals from noisy measurements. This goal is accomplished in three stages. First, we identify the best achievable MSE for estimating an FRI signal $x(t)$ from its continuous-time measurements $y(t) = x(t) + w(t)$, providing a fundamental lower bound which is independent of the sampling method. We then compare this continuous-time bound with the

lowest possible MSE for a given sampling scheme, thus measuring the loss entailed in any particular technique. Finally, we provide a mechanism for choosing the optimal sampling kernels (in a specific Bayesian sense) utilizing a pre-specified sampling rate budget. Our results can be applied to specific families of FRI signals, but they also yield some general conclusions as to the relative difficulty of various classes of estimation problems. These general observations are summarized below.

8.4.1 Continuous-Time Bound

Our first goal in this paper is to derive the continuous-time CRB, which defines a fundamental limit on the accuracy with which an FRI signal can be estimated, regardless of the sampling technique. This bound turns out to have a particularly simple closed form expression which depends on the rate of innovation, but not on the class $\mathcal{X}$ of FRI signals being estimated. Specifically, under suitable regularity conditions, the MSE of any unbiased estimator $\hat{x}$ satisfies

$$\frac{1}{T_0}\text{MSE}(\hat{x}, x) \geq \rho_{T_0}\sigma_c^2. \tag{8.20}$$

Thus, the rate of innovation can be given a new interpretation as the ratio between the best achievable MSE and the noise variance σ_c^2. This is to be contrasted with the characterization of the rate of innovation in the noise-free case as the lowest sampling rate allowing for perfect recovery of the signal; indeed, when noise is present, perfect recovery is no longer possible.

8.4.2 Bound for Sampled Measurements

We next consider lower bounds for estimating $x(t)$ from *samples* of the signal $y(t)$. In this setting, the samples inherit the noise $w(t)$ embedded in the signal $y(t)$, and may suffer from additional discrete-time noise, for example, due to quantization. We derive the CRB for estimating $x(t)$ from sampled measurements in the presence of both types of noise. However, the combination of the two noise models complicates the mathematical analysis. Consequently, since the sampling noise model has been previously analyzed [38, 132], we focus in this paper on the assumption that the discrete-time noise is negligible.

In this setting, the sampled CRB can be designed so as to converge to the continuous-time bound as the sampling rate increases. Moreover, if the family $\mathcal{X}$ of FRI signals is contained in a finite-dimensional subspace $\mathcal{M}$ of L_2, then a sampling scheme achieving the continuous-time CRB can be constructed. Such a sampling scheme is obtained by choosing kernels which span the subspace $\mathcal{M}$, and yields samples which fully capture the information present in the signal

$y(t)$. Contrariwise, if $\mathcal{X}$ is not contained in a finite-dimensional subspace, then no finite-rate sampling method achieves the continuous-time CRB. In this case, any increase in the sampling rate can improve performance, and the continuous-time bound is obtained only asymptotically.

It is interesting to examine this distinction from a union of subspaces viewpoint. Suppose that, as in the examples of Section 8.3.3, the family $\mathcal{X}$ can be described as a union of an infinite number of subspaces $\{\mathcal{U}_\alpha\}$ indexed by the continuous parameter α, so that

$$\mathcal{X} = \bigcup_\alpha \mathcal{U}_\alpha. \tag{8.21}$$

In this case, a finite sampling rate captures all of the information present in the signal if and only if

$$\dim\left(\sum_\alpha \mathcal{U}_\alpha\right) < \infty \tag{8.22}$$

where $\dim(\mathcal{M})$ is the dimension of the subspace $\mathcal{M}$. By contrast, in the noise-free case, it has been previously shown [123] that the number of samples required to recover $x(t)$ is given by

$$\max_{\alpha_1,\alpha_2} \dim(\mathcal{U}_{\alpha_1} + \mathcal{U}_{\alpha_2}), \tag{8.23}$$

i.e., the largest dimension among sums of *two* subspaces belonging to the union. In general, the dimension of (8.22) will be much higher than (8.23), illustrating the qualitative difference between the noisy and noise-free settings. For example, if the subspaces $\mathcal{U}_\alpha$ are finite-dimensional, then (8.23) is also necessarily finite, whereas (8.22) need not be.

Nevertheless, one may hope that the structure embodied in $\mathcal{X}$ will allow *nearly* optimal recovery using a sampling rate close to the rate of innovation. This is certainly the case in many noise-free FRI settings. For example, there exist techniques which recover the pulse stream (8.14) from samples taken at the rate of innovation, despite the fact that in this case $\mathcal{X}$ is typically not contained in a finite-dimensional subspace. However, this situation often changes when noise is added, in which case standard techniques improve considerably under oversampling. This empirical observation can be quantified using the CRB: as we show, the CRB for samples taken at the rate of innovation is substantially higher in this case than the optimal, continuous-time bound. This demonstrates that the sensitivity to noise is a fundamental aspect of estimating signals of the form (8.14), rather than a limitation of existing algorithms. On the other hand, other FRI models, such as the semi-periodic pulse stream (8.17), exhibit considerable noise resilience, and indeed in these cases the CRB converges to the continuous-time value much more quickly.

As we discuss in Section 8.9.5, the different levels of robustness to noise can be explained when the signal models are examined in a union of subspaces context. In this case, the parameters $\boldsymbol{\theta}$ defining $x(t)$ can be partitioned into parameters defining the subspace $\mathcal{U}_\alpha$ and parameters pinpointing the position within the subspace. Our analysis hints that estimation of the position within a subspace is often easier than estimation of the subspace itself. Thus, when most parameters are used to select an intra-subspace position, estimation at the rate of innovation is successful, as occurs in the semi-periodic case (8.17). By contrast, when a large portion of the parameters define the subspace in use, a sampling rate higher than the rate of innovation is necessary; this is the case in the non-periodic pulse stream (8.14), wherein $\boldsymbol{\theta}$ is evenly divided among subspace-selecting parameters $\{t_\ell\}$ and intra-subspace parameters $\{a_\ell\}$. Thus we see that the CRB, together with the union of subspaces viewpoint, provide valuable insights into the relative degrees of success of various FRI estimation techniques.

8.4.3 Choosing the Sampling Kernels

In some cases, one may choose the sampling system according to design specifications, leading to the question: What sampling kernels should be chosen given an allotted number of samples? We tackle this problem by adopting a Bayesian framework, wherein the signal $x(t)$ is a random process distributed according to a known prior distribution. We further assume that both the sampling and reconstruction techniques are linear. While nonlinear reconstruction methods are often used for estimating FRI signals, this assumption is required for analytical tractability, and is used only for the purpose of identifying sampling kernels. Once these kernels are chosen, they can be used in conjunction with nonlinear reconstruction algorithms.

Under these assumptions, we identify the sampling kernels yielding the minimal MSE. An additional advantage of our assumption of linearity is that in this case, the optimal kernels depend only on the autocorrelation

$$R_X(t,\tau) = E\{x(t)x^*(\tau)\} \tag{8.24}$$

of the signal $x(t)$, rather than on higher-order statistics. Indeed, given a budget of N samples, the optimal sampling kernels are given by the N eigenfunctions of R_X corresponding to the N largest eigenvalues. This is reminiscent of the Karhunen–Loève transform (KLT), which can be used to identify the optimal sampling kernels in the noiseless setting. However, in our case, shrinkage is applied to the measurements prior to reconstruction, as is typically the case with Bayesian estimation of signals in additive noise.

A setting of particular interest occurs when the autocorrelation R_X is cyclic, in the sense that

$$R_X(t,\tau) = R_X((t-\tau) \bmod T) \tag{8.25}$$

for some T. This scenario occurs, for example, in the periodic pulse stream (8.16) and the semi-periodic pulse stream (8.17), assuming a reasonable prior distribution on the parameters $\boldsymbol{\theta}$. It is not difficult to show that the eigenfunctions of R_X are given, in this case, by the complex exponentials

$$\psi_n(t) = \frac{1}{\sqrt{T}} e^{j\frac{2\pi}{T}nt}, \quad n \in \mathbb{Z}. \tag{8.26}$$

Furthermore, in the case of the periodic and semi-periodic pulse streams, the magnitudes of the eigenvalues of R_X are directly proportional to the magnitudes of the respective Fourier coefficients of the pulse shape $g(t)$. It follows that the optimal sampling kernels are the exponentials (8.26) corresponding to the largest Fourier coefficients of $g(t)$. This result is encouraging in light of recently proposed FRI reconstruction techniques which utilize exponential sampling kernels [125], and demonstrates the suitability of the Bayesian approach for designing practical estimation kernels.

8.5 Mathematical Prerequisites: CRB for General Parameter Spaces

In statistics and signal processing textbooks, the CRB is typically derived for parameters belonging to a finite-dimensional Euclidean space [6, 16, 69]. However, this result is insufficient when it is required to estimate a parameter x belonging to other Hilbert spaces, such as the L_2 space defined above. When no knowledge about the structure of the parameter x is available, a bound for estimating $x(t)$ from measurements contaminated by *colored* noise was derived in [59]. However, this bound does not hold when the noise $w(t)$ is white. Indeed, in the white noise case, it can be shown that no finite-MSE unbiased estimators exist, unless further information about $x(t)$ is available. For example, the naive estimator $\hat{x}(t) = y(t)$ has an error $\hat{x}(t) - x(t)$ equal to $w(t)$, whose variance is infinite.

In our setting, we are given the additional information that x belongs to the constraint set $\mathcal{X}$ of (8.5). To the best of our knowledge, the CRB has not been previously defined for any type of constraint set $\mathcal{X} \in L_2$, a task which will be accomplished in the present section. As we show below, a finite-valued CRB can be constructed by requiring unbiasedness only within the constraint set $\mathcal{X}$, as per (8.9). As we will see, the CRB increases linearly with the dimension of the manifold $\mathcal{X}$. Thus, in particular, the CRB is infinite when $\mathcal{X} = L_2$. However, for FRI signals,

the dimension of $\mathcal{X}$ is finite by definition, implying that a finite-valued CRB can be constructed. Although the development of this bound invokes some deep concepts from measure theory, it is a direct analog of the CRB for finite-dimensional parameters [6, Th. 2.5.15].

To derive the bound in the broadest setting possible, in this section we temporarily generalize the scenario of Section 8.2, and consider estimation of a parameter x belonging to an arbitrary measurable and separable Hilbert space $\mathcal{H}$. The MSE of an estimator $\hat{x}$ in this setting is defined as $\mathrm{MSE}(\hat{x}, x) = E\{\|\hat{x} - x\|_{\mathcal{H}}^2\}$. The concept of bias can similarly be extended if one defines the expectation $E\{v\}$ of a random variable $v : \Omega \to \mathcal{H}$ as an element $k \in \mathcal{H}$ such that $\langle k, \varphi \rangle = E\{\langle v, \varphi \rangle\}$ for any $\varphi \in \mathcal{H}$. If no such element exist, the expectation is said to be undefined.

The derivation of the CRB requires the existence of a "probability density" $p_{\boldsymbol{\theta}}(y)$ (more precisely, a Radon–Nikodym derivative) which is differentiable with respect to $\boldsymbol{\theta}$, and such that its differentiation with respect to $\boldsymbol{\theta}$ can be interchanged with integration with respect to y. The CRB also requires the mapping $h_{\boldsymbol{\theta}}$ between $\boldsymbol{\theta}$ and x to be non-redundant and differentiable. The formal statement of these regularity conditions is specified below. For the measurement setting (8.6), with reasonable mappings $h_{\boldsymbol{\theta}}$, these conditions are guaranteed to hold, as we will demonstrate in the sequel; in this section, however, we list these conditions in full so that a more general statement of the CRB will be possible.

P1) There exists a value $\boldsymbol{\theta}_0 \in \Theta$ such that the measure $P_{\boldsymbol{\theta}_0}$ dominates $\{P_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \Theta\}$. In other words, there exists a Radon–Nikodym derivative $p_{\boldsymbol{\theta}}(y) \triangleq dP_{\boldsymbol{\theta}}/dP_{\boldsymbol{\theta}_0}$ such that, for any event $A \in \mathscr{Y}$,

$$P_{\boldsymbol{\theta}}(A) = \int_A p_{\boldsymbol{\theta}}(y) P_{\boldsymbol{\theta}_0}(dy). \tag{8.27}$$

P2) For all y such that $p_{\boldsymbol{\theta}}(y) > 0$, the functions $p_{\boldsymbol{\theta}}(y)$ and $\log p_{\boldsymbol{\theta}}(y)$ are continuously differentiable with respect to $\boldsymbol{\theta}$. We denote by $\partial p_{\boldsymbol{\theta}}(y)/\partial \boldsymbol{\theta}$ and $\partial \log p_{\boldsymbol{\theta}}(y)/\partial \boldsymbol{\theta}$ the column vectors of the gradients of these two functions.

P3) The support $\{y \in \mathcal{Y} : p_{\boldsymbol{\theta}}(y) > 0\}$ of $p_{\boldsymbol{\theta}}(y)$ is independent of $\boldsymbol{\theta}$.

P4) There exists a measurable function $q : \mathcal{Y} \times \Theta \to \mathbb{R}$ such that for all sufficiently small $\Delta > 0$, for all $i = 1, \ldots, K$, for all y, and for all $\boldsymbol{\theta}$,

$$\frac{1}{\Delta} \left| p_{\boldsymbol{\theta} + \Delta \boldsymbol{e}_i}(y) - p_{\boldsymbol{\theta}}(y) \right| \leq q(y, \boldsymbol{\theta}) \tag{8.28}$$

and such that for all $\boldsymbol{\theta}$,

$$\int q^2(y, \boldsymbol{\theta}) P_{\boldsymbol{\theta}_0}(dy) < \infty. \tag{8.29}$$

In (8.28), $\boldsymbol{e}_i$ represents the ith column of the $K \times K$ identity matrix.

P5) For each $\boldsymbol{\theta}$, the $K \times K$ Fisher information matrix (FIM)

$$\boldsymbol{J}_{\boldsymbol{\theta}} \triangleq E\left\{ \left(\frac{\partial \log p_{\boldsymbol{\theta}}(y)}{\partial \boldsymbol{\theta}} \right) \left(\frac{\partial \log p_{\boldsymbol{\theta}}(y)}{\partial \boldsymbol{\theta}} \right)^* \right\} \tag{8.30}$$

is finite and invertible.

P6) $h_{\boldsymbol{\theta}}$ is Fréchet differentiable with respect to $\boldsymbol{\theta}$, in the sense that for each $\boldsymbol{\theta}$, there exists a continuous linear operator $\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta} : \mathbb{R}^K \to L_2$ such that, for any sufficiently small $\boldsymbol{\delta} \in \mathbb{R}^K$,

$$\frac{h_{\boldsymbol{\theta}+\boldsymbol{\delta}} - h_{\boldsymbol{\theta}}}{\|\boldsymbol{\delta}\|} = \frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \boldsymbol{\delta} + o(\|\boldsymbol{\delta}\|) \quad \text{as } \|\boldsymbol{\delta}\| \to 0. \tag{8.31}$$

P7) The null space of the mapping $\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta}$ contains only the zero vector. This assumption is required to ensure that the mapping from $\boldsymbol{\theta}$ to x is non-redundant, in the sense that there does not exist a parametrization of $\mathcal{X}$ in which the number of degrees of freedom is smaller than K.

We are now ready to state the CRB for the estimation of a parameter $x \in L_2[0, T]$ parameterized by a finite-dimensional vector $\boldsymbol{\theta}$. The proof of this theorem is given in Appendix 8.A.

Theorem 8.1. *Let $\boldsymbol{\theta} \in \Theta$ be a deterministic parameter, where Θ is an open set in $\mathbb{R}^K$. Let $\mathcal{H}$ be a measurable, separable Hilbert space and let $h_{\boldsymbol{\theta}}$ be a mapping from Θ to $\mathcal{H}$. Let $\{P_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \Theta\}$ be a family of probability measures over a measurable space $(\Omega, \mathscr{F})$, and let $y : \Omega \to \mathcal{Y}$ be a random variable, where $\mathcal{Y}$ is a measurable Hilbert space. Assume regularity conditions P1–P6. Let $\hat{x} : \mathcal{Y} \to \mathcal{H}$ be an unbiased estimator of x from the measurements y such that*

$$E\{ \|\hat{x}(y)\|_{\mathcal{H}}^2 \} < \infty. \tag{8.32}$$

Then, the MSE of $\hat{x}$ satisfies

$$\mathrm{MSE}(\hat{x}, x) \geq \mathrm{Tr}\left[\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}_{\boldsymbol{\theta}}^{-1} \right] \tag{8.33}$$

where $\boldsymbol{J}_{\boldsymbol{\theta}}$ is the FIM (8.30).

Theorem 8.1 enables us to obtain a lower bound on the estimation error of x based on the FIM for estimating $\boldsymbol{\theta}$. The latter can often be computed relatively easily since $\boldsymbol{\theta}$ is a finite-dimensional vector. Even more conveniently, the trace on the right-hand side of (8.33) is taken over a $K \times K$ matrix, despite the involvement of continuous-time operators. Thus, the computation of (8.33) is often possible either analytically or numerically, a fact which will be used extensively in the sequel.

8.6 CRB for Continuous-Time Measurements

We now apply Theorem 8.1 to the problem of estimating a deterministic signal x from continuous-time measurements y given by (8.6).

Theorem 8.2. *Let x be a deterministic function defined by* (8.5)*, where $\boldsymbol{\theta} \in \Theta$ is an unknown deterministic parameter and Θ is an open subset of $\mathbb{R}^K$. Suppose that Assumptions P6–P7 are satisfied. Then, the MSE of any unbiased, finite-variance estimator $\hat{x}$ is bounded by*

$$\mathrm{MSE}(\hat{x}, x) \geq K\sigma_c^2. \tag{8.34}$$

The bound of Theorem 8.2 can be translated to units of the rate of innovation ρ_{T_0} if we assume that the segment $[0, T_0]$ under analysis achieves the maximum (8.4), i.e., this is a segment containing the maximum allowed number of degrees of freedom. In this case $\rho_{T_0} = K/T_0$, and any unbiased estimator $\hat{x}(t)$ satisfies

$$\frac{\frac{1}{T_0} E\left\{ \int_0^{T_0} |x(t) - \hat{x}(t)|^2 dt \right\}}{\sigma_c^2} \geq \rho_{T_0}. \tag{8.35}$$

In the noisy setting, ρ_{T_0} loses its meaning as a lower bound on the sampling rate required for perfect recovery, since the latter is no longer possible at any sampling rate. On the other hand, it follows from (8.35) that the rate of innovation gains an alternative meaning; namely, ρ_{T_0} is a lower bound on the ratio between the average MSE achievable by any unbiased estimator and the noise variance σ_c^2, *regardless of the sampling method*.

Before formally proving Theorem 8.2, note that (8.34) has an intuitive geometric interpretation. Specifically, the constraint set $\mathcal{X}$ is a K-dimensional differential manifold in $L_2[0, T]$. In other words, for any point $x \in \mathcal{X}$, there exists a K-dimensional subspace $\mathcal{U}$ tangent to $\mathcal{X}$ at x. We refer to $\mathcal{U}$ as the feasible direction subspace [70]: any perturbation of x which remains within the constraint set $\mathcal{X}$ must be in one of the directions in $\mathcal{U}$. Formally, $\mathcal{U}$ can be defined as the range space of $\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta}$.

If one wishes to use the measurements y to distinguish between x and its local neighborhood, then it suffices to observe the projection of y onto $\mathcal{U}$. Projecting the measurements onto $\mathcal{U}$ removes most of the noise, retaining only K independent Gaussian components, each having a variance of σ_c^2. Thus we have obtained an intuitive explanation for the bound of $K\sigma_c^2$ in Theorem 8.2. To formally prove this result, we apply Theorem 8.1 to the present setting, as follows.

Proof of Theorem 8.2. The problem of estimating the parameters $\boldsymbol{\theta}$ from a continuous-time signal $y(t)$ of the form (8.6) was examined in [134, Example I.7.3], where the validity of Assumptions P1–P4 was demonstrated. It was further shown that the FIM $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}}$ for estimating $\boldsymbol{\theta}$ from $y(t)$ is given by [134, *ibid.*]

$$\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}} = \frac{1}{\sigma_c^2} \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right). \tag{8.36}$$

Our goal will be to use (8.36) and Theorem 8.1 to obtain a bound on estimators of the continuous-time function $x(t)$. To this end, observe that the FIM $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}}$ is finite since, by Assumption P6, the operator $\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta}$ is a bounded operator into L_2. Furthermore, by Assumption P7, $\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta}$ has a trivial null space, and thus $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}}$ is invertible. Therefore, Assumption P5 has been demonstrated. We may consequently apply Theorem 8.1, which yields

$$\begin{aligned} &\text{MSE}(\hat{x}, x) \\ &\quad \geq \sigma_c^2 \,\text{Tr} \left[\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right) \left(\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right) \right)^{-1} \right] \\ &\quad = \sigma_c^2 \,\text{Tr} \left(\boldsymbol{I}_K \right) \\ &\quad = K \sigma_c^2 \end{aligned} \tag{8.37}$$

thus completing the proof. □

To illustrate the use of Theorem 8.2 in practice, let us consider as a simple example a signal $x(t)$ belonging to a finite-dimensional subspace $\mathcal{G}$. Specifically, assume that

$$x(t) = \sum_{k=1}^{K} a_k g_k(t) \tag{8.38}$$

for some coefficient vector $\boldsymbol{\theta} = (a_1, \ldots, a_K)^T$ and a given set of linearly independent functions $\{g_k\}$ spanning $\mathcal{G}$. This includes, for example, families of shift-invariant subspaces with a compactly supported generator (see Section 8.3.1). From Theorem 8.2, the MSE of any unbiased estimator of x is bounded by $K\sigma_c^2$, where K is the dimension of the subspace $\mathcal{G}$. We now demonstrate that this bound is achieved by the unbiased estimator

$$\hat{x} = \boldsymbol{P}_{\mathcal{G}} \, y \tag{8.39}$$

where $\boldsymbol{P}_{\mathcal{G}}$ is the orthogonal projector onto the subspace $\mathcal{G}$.

To verify that (8.39) achieves the CRB, let G denote the set transformation (8.2) associated with the functions $\{g_k\}_{k=1}^K$. One may then write $x = G\boldsymbol{\theta}$ and $\boldsymbol{P}_{\mathcal{G}} = G(G^*G)^{-1}G^*$. Thus (8.39)

becomes

$$\begin{aligned}\hat{x} &= G(G^*G)^{-1}G^*G\boldsymbol{\theta} + G(G^*G)^{-1}G^*w \\ &= G\boldsymbol{\theta} + \boldsymbol{P}_{\mathcal{G}}\, w \end{aligned} \tag{8.40}$$

and therefore

$$E\big\{\|\hat{x} - x\|_{L_2}^2\big\} = E\big\{\|\boldsymbol{P}_{\mathcal{G}}\, w\|_{L_2}^2\big\}\,. \tag{8.41}$$

Since $\mathcal{G}$ is a K-dimensional subspace, it is spanned by a set of K orthonormal[1] functions $u_1, \ldots, u_K \in L_2$. Thus

$$E\big\{\|\boldsymbol{P}_{\mathcal{G}}\, w\|_{L_2}^2\big\} = \sum_{k=1}^{K} E\Big\{|\langle w, u_k\rangle|^2\Big\} = K\sigma_c^2 \tag{8.42}$$

which demonstrates that $\hat{x}$ indeed achieves the CRB in this case.

In practice, a signal is not usually estimated directly from its continuous-time measurements. Rather, the signal $y(t)$ is typically sampled and digitally manipulated. In the next section, we will compare the results of Theorem 8.2 with the performance achievable from sampled measurements, and demonstrate that in some cases, a finite-rate sampling scheme is sufficient to achieve the continuous-time bound of Theorem 8.2.

8.7 CRB for Sampled Measurements

In this section, we consider the problem of estimating $x(t)$ of (8.5) from a finite number of samples of the process $y(t)$ given by (8.6). Specifically, suppose our measurements are given by

$$c_n = \langle y, s_n\rangle + v_n = \int_0^{T_0} y(t)s_n^*(t)dt + v_n, \quad n = 1, \ldots, N \tag{8.43}$$

where $\{s_n\}_{n=1}^N \subset L_2[0, T_0]$ are sampling kernels, and v_n is a discrete white Gaussian noise process, independent of $w(t)$, having mean zero and variance σ_d^2. Note that the model (8.43) includes both continuous-time noise, which is present in the signal $y(t) = x(t) + w(t)$ prior to sampling, and digital noise v_n, which arises from the sampling process, e.g., as a result of quantization. In this section, we will separately examine the effect of each of these noise components.

From (8.6) and (8.43), it can be seen that the measurements $c_1, \ldots, c_N$ are jointly Gaussian with mean

$$\mu_n \triangleq E\{c_n\} = \langle x, s_n\rangle \tag{8.44}$$

[1]We require the new functions $u_1, \ldots, u_K$ since the functions $g_1, \ldots, g_K$ are not necessarily orthonormal. The choice of non-orthonormal functions $g_1, \ldots, g_K$ will prove useful in the sequel.

and covariance

$$\Gamma_{ij} \triangleq \mathrm{Cov}(c_i, c_j) = \sigma_c^2 \langle s_i, s_j \rangle + \sigma_d^2 \delta_{ij}. \tag{8.45}$$

A somewhat unusual aspect of this estimation setting is that the choice of the sampling kernels $s_n(t)$ affects not only the measurements obtained, but also the statistics of the noise. One example of the impact of this fact is the following. Suppose first that no digital noise is present, i.e., $\sigma_d = 0$, and consider a modified set of sampling kernels $\{\tilde{s}_n(t)\}_{n=1}^N$ which are an invertible linear transformation of $\{s_n(t)\}_{n=1}^N$, so that

$$\tilde{s}_n(t) = \sum_{i=1}^N \mathcal{B}_{ni} s_i(t) \tag{8.46}$$

where $\mathcal{B} \in \mathbb{R}^{N \times N}$ is an invertible matrix. Then, the resulting measurements $\tilde{\boldsymbol{c}}$ are given by $\tilde{\boldsymbol{c}} = \mathcal{B}\boldsymbol{c}$, and similarly the original measurements $\boldsymbol{c}$ can be recovered from $\tilde{\boldsymbol{c}}$. It follows that these settings are equivalent in terms of the accuracy with which x can be estimated. In particular, the FIM for estimating x in the two settings is identical [134, Th. I.7.2].

When digital noise is present in addition to continuous-time noise, the sampling schemes $\{s_n(t)\}$ and $\{\tilde{s}_n(t)\}$ are no longer necessarily equivalent, since the gain introduced by the transformation $\mathcal{B}$ will alter the ratio between the energy of the signal and the digital noise. The two estimation problems are then equivalent if and only if $\mathcal{B}$ is a unitary transformation.

How should one choose the space $\mathcal{S} = \mathrm{span}\{s_1, \ldots, s_N\}$ spanned by the sampling kernels? Suppose for a moment that there exist elements in the range space of $\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta}$ which are orthogonal to $\mathcal{S}$. This implies that one can perturb x in such a way that the constraint set $\mathcal{X}$ is not violated, without changing the distribution of the measurements $\boldsymbol{c}$. This situation occurs, for example, when the number of measurements N is smaller than the dimension K of the parametrization of $\mathcal{X}$. While it may still be possible to reconstruct some of the information concerning x from these measurements, this is an undesirable situation from an estimation point of view. Thus we will assume henceforth that

$$\mathcal{R}\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right) \cap \mathcal{S}^\perp = \{\mathbf{0}\}. \tag{8.47}$$

As an example of the necessity of the condition (8.47), consider again the signal (8.38), which belongs to a K-dimensional subspace $\mathcal{G} \subset L_2$ spanned by the functions $g_1, \ldots, g_K$. In this case it is readily seen that for any vector $\boldsymbol{v}$

$$\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \boldsymbol{v} = \sum_{k=1}^K \boldsymbol{v}_k g_k(t). \tag{8.48}$$

Since the functions $\{g_k\}$ span the space $\mathcal{G}$, this implies that $\mathcal{R}(\partial h_{\boldsymbol{\theta}}/\partial \boldsymbol{\theta}) = \mathcal{G}$, and therefore the condition (8.47) can be written as

$$\mathcal{G} \cap \mathcal{S}^{\perp} = \{\mathbf{0}\} \tag{8.49}$$

which is a standard requirement in the design of a sampling system for signals belonging to a subspace $\mathcal{G}$ [148].

By virtue of Theorem 8.1, a lower bound on unbiased estimation of x can be obtained by first computing the FIM $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ for estimating $\boldsymbol{\theta}$ from $\boldsymbol{c}$. This yields the following result. For simplicity of notation, in this theorem we assume that the function $h_{\boldsymbol{\theta}}$ and the sampling kernels s_n are real. If complex sampling kernels are desired (as will be required in the sequel), the result below can still be used by translating each measurement to an equivalent pair of real-valued samples.

Theorem 8.3. *Let x be a deterministic real function defined by* (8.5)*, where $\boldsymbol{\theta} \in \Theta$ is an unknown deterministic parameter and Θ is an open subset of $\mathbb{R}^K$. Assume regularity conditions P6–P7, and let $\hat{x}$ be an unbiased estimator of x from the real measurements $\boldsymbol{c} = (c_1, \ldots, c_N)^T$ of* (8.43)*. Then, the FIM $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ for estimating $\boldsymbol{\theta}$ from $\boldsymbol{c}$ is given by*

$$\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}} = \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right)^* S \left(\sigma_c^2 S^* S + \sigma_d^2 \boldsymbol{I}_N\right)^{-1} S^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right) \tag{8.50}$$

where S is the set transformation corresponding to the functions $\{s_n\}_{n=1}^N$. If (8.47) *holds, then $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ is invertible. In this case, any finite-variance, unbiased estimator $\hat{x}$ for estimating x from $\boldsymbol{c}$ satisfies*

$$\text{MSE}(\hat{x}, x) \geq \text{Tr}\left[\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right)^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right) (\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}})^{-1}\right]. \tag{8.51}$$

Proof. In the present setting, the FIM $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ is given by [16]

$$\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}} = \left(\frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\theta}}\right)^* \boldsymbol{\Gamma}^{-1} \left(\frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\theta}}\right) \tag{8.52}$$

where the matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{N \times N}$ is defined by (8.45) and the matrix $\partial \boldsymbol{\mu}/\partial \boldsymbol{\theta} \in \mathbb{R}^{N \times K}$ is given by

$$\left(\frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\theta}}\right)_{nk} = \frac{\partial \mu_n}{\partial \theta_k} \tag{8.53}$$

with μ_n defined in (8.44).

By the definition of the set transformation, the ijth element of the $N \times N$ matrix S^*S is given by

$$(S^*S)_{ij} = \langle S\boldsymbol{e}_j, S\boldsymbol{e}_i \rangle = \langle s_j, s_i \rangle \tag{8.54}$$

where $\boldsymbol{e}_i$ is the ith column of the $N \times N$ identity matrix. Therefore, we have

$$\boldsymbol{\Gamma} = \sigma_c^2 S^* S + \sigma_d^2 \boldsymbol{I}_N. \tag{8.55}$$

Similarly, observe that

$$\left(S^*\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right)_{nk} = \left\langle \frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\tilde{\boldsymbol{e}}_k, S\boldsymbol{e}_n \right\rangle = \left\langle \frac{\partial h_{\boldsymbol{\theta}}}{\partial \theta_k}, s_n \right\rangle = \frac{\partial \mu_n}{\partial \theta_k} \tag{8.56}$$

where $\tilde{\boldsymbol{e}}_k$ is the kth column of the $K \times K$ identity matrix. Thus

$$\frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\theta}} = S^*\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}. \tag{8.57}$$

Substituting (8.55) and (8.57) into (8.52) yields the required expression (8.50).

We next demonstrate that if (8.47) holds, then $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ is invertible. To see this, note that from (8.50) we have

$$\mathcal{N}(\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}) = \mathcal{N}\left(S^*\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right)\right). \tag{8.58}$$

Now, consider an arbitrary function $f \in \mathcal{R}(\partial h_{\boldsymbol{\theta}}/\partial \boldsymbol{\theta})$. If (8.47) holds, then f is not orthogonal to the subspace $\mathcal{S}$. Therefore, $\langle f, s_n \rangle \neq 0$ for at least one value of n, and thus by (8.3), $S^*f \neq 0$. This implies that

$$\mathcal{N}\left(S^*\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right)\right) = \mathcal{N}\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right) = \{\boldsymbol{0}\}. \tag{8.59}$$

Combined with (8.58), we conclude that $\mathcal{N}(\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}) = \{\boldsymbol{0}\}$. This demonstrates that $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ is invertible, proving Assumption P5. Moreover, in the present setting, Assumptions P1–P4 are fulfilled for any value of $\boldsymbol{\theta}_0$ [6]. Applying Theorem 8.1 yields (8.51) and completes the proof. □

In the following subsections we draw several conclusions from Theorem 8.3.

8.7.1 Discrete-Time Noise

Suppose first that $\sigma_c^2 = 0$, so that only digital noise is present. This setting has been analyzed previously [38,132], and therefore only briefly examine the contrast with continuous-time noise. When only digital noise is present, its effects can be surmounted either by increasing the gain of the sampling kernels, or by increasing the number of measurements. These intuitive conclusions can be verified from Theorem 8.3 as follows. Assume that condition (8.47) holds, and consider the modified kernels $\tilde{s}_n(t) = 2s_n(t)$. The set transformation $\tilde{S}$ corresponding to the modified kernels is $\tilde{S} = 2S$, and since $\sigma_c^2 = 0$, this implies that the FIM obtained from the modified kernels is given by $\tilde{\boldsymbol{J}}_{\boldsymbol{\theta}}^{\text{samp}} = 4\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$. Thus, a sufficient increase in the sampling gain can arbitrarily increase $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ and consequently reduce the bound (8.51) arbitrarily close to zero.

Of course, from a practical point of view, increasing the gain also increases the likelihood that the sampled signal will exceed the dynamic range of the quantizer. It is therefore not feasible to arbitrarily increase the sampling gain. As an alternative, it is possible to increase

the number of measurements. For example, suppose one simply repeats each measurement twice. Let S and $\tilde{S}$ denote the transformations corresponding to the original and doubled sets of measurements. It can then readily be seen from the definition of the set transformation (8.2) and its adjoint (8.3) that $\tilde{S}\tilde{S}^* = 4SS^*$. Consequently, by the same argument, in the absence of continuous-time noise one can achieve arbitrarily low error by repeated measurements, or more generally, by increasing the sampling rate.

8.7.2 Continuous-Time Noise

As we have seen, sampling noise can be mitigated by increasing the sampling rate. Furthermore, digital noise is inherently dependent on the sampling scheme being used. Since our goal is to determine the fundamental performance limits regardless of the sampling technique, we will focus here and in subsequent sections on continuous-time noise. Thus, suppose that $\sigma_d^2 = 0$, so that only continuous-time noise is present. In this case, as we now show, it is generally impossible to achieve arbitrarily low reconstruction error, regardless of the sampling kernels used; indeed, it is never possible to outperform the continuous-time CRB of Section 8.6, which is typically nonzero. To see this formally, observe first that in the absence of digital noise, the FIM for estimating $\boldsymbol{\theta}$ can be written as

$$\begin{aligned} \boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}} &= \frac{1}{\sigma_c^2} \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^* S \left(S^* S \right)^{-1} S^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right) \\ &= \frac{1}{\sigma_c^2} \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^* \boldsymbol{P}_{\mathcal{S}} \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right) \end{aligned} \tag{8.60}$$

where $\boldsymbol{P}_{\mathcal{S}}$ is the orthogonal projection onto the subspace $\mathcal{S}$. It is insightful to compare this expression with the FIM $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}}$ obtained from continuous-time measurements in (8.36). In both cases, a lower bound on the MSE for unbiased estimation of x was obtained from $\boldsymbol{J}_{\boldsymbol{\theta}}$ by applying Theorem 8.1. Consequently, if it happens that $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}} = \boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$, then the continuous-time bound of Theorem 8.2 and the sampled bound of Theorem 8.3 coincide. Thus, if no digital noise is added, then it is possible (at least in terms of the performance bounds) that estimators based on the samples $\boldsymbol{c}$ will suffer no degradation compared with the "ideal" estimator based on the entire set of continuous-time measurements. This occurs if and only if $\mathcal{R}(\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta}) \subseteq \mathcal{S}$; in this case, the projection $\boldsymbol{P}_{\mathcal{S}}$ will have no effect on the FIM $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$, which will then coincide with $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}}$ of (8.36). In the remainder of this section, we will discuss several cases in which this fortunate circumstance arises.

8.7.3 Example: Sampling in a Subspace

One situation in which samples provide all of the information present in the continuous-time signal is the case in which $x(t)$ belongs to a K-dimensional subspace $\mathcal{G}$ of L_2. This is the case, for example, when the signal lies in a shift-invariant subspace having a compactly supported generator (see Section 8.3.1). As we have seen above (cf. (8.48)), in this scenario $\partial h_{\boldsymbol{\theta}}/\partial \boldsymbol{\theta}$ is a mapping onto the subspace $\mathcal{G}$. Assuming that there is no discrete-time noise, it follows from (8.60) that the optimal choice of a sampling space $\mathcal{S}$ is $\mathcal{G}$ itself. Such a choice requires $N = K$ samples and yields $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}} = \boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$. Of course, such an occurrence is not possible if the sampling process contributes additional noise to the measurements.

In some cases, it may be difficult to implement a set of sampling kernels spanning the subspace $\mathcal{G}$. It may then be desirable to choose a K-dimensional subspace $\mathcal{S}$ which is close to $\mathcal{G}$ but does not equal it. We will now compute the CRB for this setting and demonstrate that it can be achieved by a practical estimation technique. This will also demonstrate achievability of the CRB in the special case $\mathcal{S} = \mathcal{G}$. We first note from (8.2) and (8.48) that $\partial h_{\boldsymbol{\theta}}/\partial \boldsymbol{\theta} = G$, where G is the set transformation corresponding to the generators $\{g_k\}_{k=1}^K$. Furthermore, it follows from (8.49) that S^*G and G^*S are invertible $K \times K$ matrices [148]. Using Theorem 8.3, we thus find that the CRB is given by

$$\begin{aligned}\text{MSE}(\hat{x}, x) &\geq \sigma_c^2 \,\text{Tr}\Big(G\left(G^*S(S^*S)^{-1}S^*G\right)^{-1}G^*\Big) \\ &= \sigma_c^2 \,\text{Tr}\Big(G(S^*G)^{-1}S^*S(G^*S)^{-1}G^*\Big). \end{aligned} \tag{8.61}$$

It is readily seen that when $\mathcal{S} = \mathcal{G}$, the bound (8.61) reduces to $K\sigma_c^2$, which is (as expected) the continuous-time bound of Theorem 8.2. When $\mathcal{S} \neq \mathcal{G}$, the bound (8.61) will generally be higher than $K\sigma_c^2$, since $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ of (8.60) will exceed $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{cont}}$ of (8.36). In this case, it is common to use the consistent, unbiased estimator [137, 148]

$$\hat{x} = G(S^*G)^{-1}\boldsymbol{c}. \tag{8.62}$$

As we now show, the bound (8.61) is achieved by this estimator. Indeed, observe that $\boldsymbol{c} = S^*y = S^*G\boldsymbol{\theta} + S^*w$, and thus

$$\begin{aligned}E\big\{\|\hat{x} - x\|_{L_2}^2\big\} &= E\Big\{\|G(S^*G)^{-1}S^*w\|_{L_2}^2\Big\} \\ &= E\Big\{\text{Tr}\Big(G(S^*G)^{-1}S^*ww^*S(G^*S)^{-1}G^*\Big)\Big\} \\ &= \text{Tr}\Big(G(S^*G)^{-1}\,\text{Cov}(S^*w)(G^*S)^{-1}G^*\Big). \end{aligned} \tag{8.63}$$

Note that $\mathrm{Cov}(S^*w) = \mathrm{Cov}(\boldsymbol{c})$, which by (8.55) is equal to $\sigma_c^2 S^*S$. Substituting this result into (8.63) and comparing with (8.61) verifies that $\hat{x}$ achieves the CRB.

8.7.4 Nyquist-Equivalent Sampling

We refer to situations in which the dimension of the sampling space equals the dimension of the signal space as "Nyquist-equivalent" sampling schemes. In the previous section, we saw that Nyquist-equivalent sampling is possible using K samples when the signal lies in a K-dimensional subspace $\mathcal{X}$, and that the resulting system achieves the continuous-time CRB. A similar situation occurs when the set of possible signals $\mathcal{X}$ is a subset of an M-dimensional subspace $\mathcal{M}$ of L_2 with $M > K$. In this case, it can be readily shown that $\mathcal{R}(\partial h_{\boldsymbol{\theta}}/\partial \boldsymbol{\theta}) \subseteq \mathcal{M}$. Thus, by choosing $N = M$ sampling kernels such that $\mathcal{S} = \mathcal{M}$, we again achieve $\boldsymbol{J}_{\boldsymbol{\theta}}^{\mathrm{cont}} = \boldsymbol{J}_{\boldsymbol{\theta}}^{\mathrm{samp}}$, demonstrating that all of the information content in x has been captured by the samples. This is again a Nyquist-equivalent scheme, but the number of samples it requires is higher than the number of parameters K defining the signals. Therefore, in this case it is not possible to sample at the rate of innovation without losing some of the information content of the signal.

In general, the constraint set $\mathcal{X}$ will not be contained in any finite-dimensional subspace of L_2. In such cases, it will generally not be possible to achieve the performance of the continuous-time bound using *any* finite number of samples, even in the absence of digital noise. This implies that in the most general setting, sampling above the rate of innovation can often improve the performance of estimation schemes. This conclusion will be verified by simulation in Section 8.9.

8.8 Optimal Sampling Kernels: A Bayesian Viewpoint

In this section, we address the problem of designing a sampling method which minimizes the MSE. One route towards this goal could be to minimize the sampled CRB of Theorem 8.3 with respect to the sampling space $\mathcal{S}$. However, the CRB is a function of the unknown parameter vector $\boldsymbol{\theta}$. Consequently, for each value of $\boldsymbol{\theta}$, there may be a different sampling space $\mathcal{S}$ which minimizes the bound. To obtain a sampling method which is optimal on average over all possible choices of $\boldsymbol{\theta}$, we now make the additional assumption that the parameter vector $\boldsymbol{\theta}$ is random and has a known distribution. Our goal, then, is to determine the sampling space $\mathcal{S}$ that minimizes the MSE $E\left\{\|\hat{x} - x\|_{L_2}^2\right\}$ within a class of allowed estimators. Note that the mean is now taken over realizations of both the noise $w(t)$ and the parameter $\boldsymbol{\theta}$.

Since $\boldsymbol{\theta}$ is random, the signal $x(t)$ is random as well. To make our discussion general, we will derive the optimal sampling functions for estimating a general random process $x(t)$ (not necessarily having realizations in $\mathcal{X}$ of (8.5)) from samples of the noisy process $y(t) = x(t) + w(t)$. We will then specialize the results to several specific types of FRI signals and obtain explicit expressions for the optimal sampling kernels in these scenarios.

Let $x(t)$ denote a zero-mean random process defined over $t \in [0, T_0]$, and suppose that its autocorrelation function

$$R_X(t, \eta) \triangleq E\{x(t)x^*(\eta)\} \tag{8.64}$$

is continuous in t and η. Our goal is to estimate $x(t)$ based on a finite number N of samples of the signal $y(t) = x(t) + w(t)$, $t \in [0, T_0]$, where $w(t)$ is a white noise process (not necessarily Gaussian) with variance σ_c^2 which is uncorrelated with $x(t)$. We focus our attention on *linear sampling* schemes, i.e., we assume the samples are given by

$$c_n = \langle y, s_n \rangle . \tag{8.65}$$

Finally, we restrict the discussion to *linear estimation* methods, namely those techniques in which the estimate $\hat{x}(t)$ is constructed as

$$\hat{x}(t) = \sum_{n=1}^{N} c_n v_n(t), \tag{8.66}$$

for some set of reconstruction functions $\{v_n(t)\}_{n=1}^N$. It is important to note that for any given set of sampling functions $\{s_n(t)\}_{n=1}^N$, the minimum MSE (MMSE) estimator of $x(t)$ is often a nonlinear function of the measurements $\{c_n\}_{n=1}^N$. Indeed, typical FRI reconstruction techniques involve a nonlinear stage. Consequently, restricting the discussion to linear recovery schemes may seem inadequate. However, this choice has two advantages. First, as we will see, the optimal linear scheme is determined only by the second-order statistics of $x(t)$ and $w(t)$, whereas the analysis of nonlinear methods necessitates exact knowledge of their entire distribution functions. Second, it is not the final estimate $\hat{x}(t)$ that interests us in this discussion, but merely the set of optimal sampling functions. Once such a set is determined, albeit from a linear recovery perspective, it can be used in conjunction with existing nonlinear FRI techniques. As we will see in Section 8.9, the conclusions obtained through our analysis appear to apply to FRI techniques in general. Under the above assumptions, our goal is to design the sampling kernels $\{s_n(t)\}_{n=1}^N$ and reconstruction functions $\{v_n(t)\}_{n=1}^N$ such that the MSE (8.8) is minimized.

As can be seen from (8.65), we assume henceforth that only continuous-time noise is present in the sampling system. The situation is considerably more complicated in the presence of digital noise. First, without digital noise, one must choose only the subspace spanned by the sampling kernels, as the kernels themselves do not affect the performance; this is no longer the case when digital noise is added. Second, digital noise may give rise to a requirement that a particular measurement be repeated in order to average out the noise. This is undesirable in the continuous noise regime, since the repeated measurement will contain the exact same noise realization.

8.8.1 Relation to the Karhunen–Loève Expansion and Finite-Dimensional Generalizations

The problem posed above is closely related to the Karhunen–Loève transform (KLT) [149, 150], which is concerned with the reconstruction of a random signal $x(t)$ from its noiseless samples. Specifically, one may express $x(t)$ in terms of a complete orthonormal basis $\{\psi_k(t)\}_{k=1}^{\infty}$ for L_2 as

$$x(t) = \sum_{k=1}^{\infty} \langle x, \psi_k \rangle \, \psi_k(t). \tag{8.67}$$

The goal of the KLT is to choose the functions $\{\psi_k(t)\}_{k=1}^{\infty}$ such that the MSE resulting from the truncation of this series after N terms is minimal. It is well known that the solution to this problem is given by the N-term truncation of the Karhunen–Loève expansion [149, 151].

Since $R_X(t, \eta)$ is assumed to be continuous in our setting, by Mercer's theorem [151] it possesses a discrete set of eigenfunctions $\{\psi_k(t)\}_{k=1}^{\infty}$, which constitute an orthonormal basis for L_2. These functions satisfy the equations

$$\lambda_k \psi_k(t) = \int_0^{T_0} R_X(t, \eta) \psi_k(\eta) d\eta, \tag{8.68}$$

in which the corresponding eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$ are nonnegative and are assumed to be arranged in descending order. With these functions, (8.67) is known as the Karhunen–Loève expansion. It can be easily shown that the first N terms in this series constitute the best N-term approximation of $x(t)$ in an MSE sense [151]. In other words, in the noiseless case, the optimal sampling and reconstruction functions are $s_n(t) = v_n(t) = \psi_n(t)$.

In our setting, we do not have access to samples of $x(t)$ but rather only to samples of the noisy process $y(t)$. In this case, it is not clear *a priori* whether the optimal sampling and reconstruction filters coincide or whether they match the Karhunen–Loève expansion of $x(t)$.

The finite-dimensional analogue of our problem, in which x, y, and w are random vectors taking values in $\mathbb{C}^M$, was treated in [152,153]. The derivation in these works, however, relied on the low-rank approximation property of the singular-value decomposition (SVD) of a matrix. The generalization of this concept to infinite-dimensional operators is subtle and will thus be avoided here. Instead, we provide a conceptually simple (if slightly cumbersome) derivation of the optimal linear sampling and reconstruction method for noisy signals. As we will see, it still holds that $s_n(t) = v_n(t) = \psi_n(t)$.

8.8.2 Optimal Sampling in Noisy Settings

As explained in Section 8.7, in the absence of discrete-time noise, the MSE is not affected by modifications of the sampling kernels which leave the set $\mathcal{S} = \text{span}\{s_1(t), \ldots, s_N(t)\}$ unchanged. Thus, without loss of generality, we constrain $\{s_n(t)\}_{n=1}^N$ to satisfy

$$\left\langle s_n \, , \, \sigma_c^2 s_m + \int_0^T R_X(\cdot, \tau) s_m(\tau) d\tau \right\rangle = \delta_{m,n} \tag{8.69}$$

for every $m, n = 1, \ldots, N$. This can always be done since the operator $R_Y : L_2 \to L_2$ defined by $(R_Y f)(t) = \int_0^T R_X(t, \tau) f(\tau) d\tau + \sigma_c^2 f(t)$ is positive definite. This choice is particularly convenient as it results in a set of uncorrelated samples $\{c_n\}$. Indeed

$$\begin{aligned} E\{c_m c_n^*\} &= E\left\{ \left(\int_0^T s_m^*(\tau) y(\tau) d\tau \right) \left(\int_0^T s_n^*(\eta) y(\eta) d\eta \right)^* \right\} \\ &= \iint_0^T s_m^*(\tau) E\{y(\tau) y^*(\eta)\} \, s_n(\eta) d\tau d\eta \\ &= \iint_0^T s_m^*(\tau) R_X(\tau, \eta) s_n(\eta) d\tau d\eta \\ &\qquad + E\left\{ \langle s_n, w \rangle \langle s_m, w \rangle^* \right\} \\ &= \iint_0^T s_m^*(\tau) R_X^*(\eta, \tau) s_n(\eta) d\tau d\eta + \sigma_c^2 \langle s_n, s_m \rangle \\ &= \delta_{m,n}. \end{aligned} \tag{8.70}$$

We are now ready to determine the optimal sampling method. We begin by expressing the MSE (8.8) as

$$\begin{aligned} \int_0^T E\left\{ |x(t) - \hat{x}(t)|^2 \right\} dt &= \int_0^T E\left\{ |x(t)|^2 \right\} dt \\ &\quad - 2 \int_0^T \Re \left\{ E\{x^*(t) \hat{x}(t)\} \right\} dt + \int_0^T E\left\{ |\hat{x}(t)|^2 \right\} dt. \end{aligned} \tag{8.71}$$

The first term in this expression does not depend on the choice of $\{s_n(t)\}_{n=1}^N$ and $\{v_n(t)\}_{n=1}^N$, and is therefore irrelevant for our purpose. Substituting (8.66) and (8.65), and using the fact

that $w(t)$ is uncorrelated with $x(t)$, the second term can be written as

$$\begin{aligned}
&\int_0^T 2\Re\left\{E\left\{x^*(t)\sum_{n=1}^N c_n v_n(t)\right\}\right\}dt \\
&= \sum_{n=1}^N 2\int_0^T \Re\left\{E\left\{x^*(t)\int_0^T y(\tau)s_n^*(\tau)d\tau\right\}v_n(t)\right\}dt \\
&= \sum_{n=1}^N 2\iint_0^T \Re\left\{E\{x^*(t)(x(\tau)+w(\tau))\}\, s_n^*(\tau)v_n(t)\right\}d\tau dt \\
&= \sum_{n=1}^N 2\iint_0^T \Re\left\{s_n^*(\tau)R_X(\tau,t)v_n(t)\right\}d\tau dt \\
&= \sum_{n=1}^N 2\Re\left\{\left\langle v_n\,,\int_0^T R_X(\cdot,\tau)s_n(\tau)d\tau\right\rangle\right\}. \qquad (8.72)
\end{aligned}$$

Similarly, using the fact that $\{c_n\}_{n=1}^N$ are uncorrelated and have unit variance (see (8.70)), the last term in (8.71) becomes

$$\begin{aligned}
\int_0^T E\left\{\left|\sum_{n=1}^N c_n v_n(t)\right|^2\right\}dt &= \sum_{m=1}^N\sum_{n=1}^N E\{c_m^* c_n\}\,\langle v_m, v_n\rangle \\
&= \sum_{n=1}^N \|v_n\|^2. \qquad (8.73)
\end{aligned}$$

Substituting (8.72) and (8.73) back into (8.71), we conclude that minimization of the MSE is equivalent to minimization of

$$\sum_{n=1}^N\left(\|v_n\|^2 - 2\Re\left\{\left\langle v_n, \int_0^T R_X(\cdot,\tau)s_n(\tau)d\tau\right\rangle\right\}\right) \qquad (8.74)$$

with respect to $\{s_n(t)\}_{n=1}^N$ and $\{v_n(t)\}_{n=1}^N$, subject to the set of constraints (8.69).

As a first stage, we minimize (8.74) with respect to the reconstruction functions $\{v_n(t)\}_{n=1}^N$. To this end, we note that the nth summand in (8.74) is lower bounded by

$$\begin{aligned}
&\|v_n\|^2 - 2\|v_n\|\left\|\int_0^T R_X(\cdot,\tau)s_n(\tau)d\tau\right\| \\
&\qquad\geq -\left\|\int_0^T R_X(\cdot,\tau)s_n(\tau)d\tau\right\|^2, \qquad (8.75)
\end{aligned}$$

where we used the Cauchy–Schwarz inequality and the fact that $\min_z\{z^2 - 2bz\} = -b^2$. This bound is achieved by choosing

$$v_n(t) = \int_0^T R_X(t,\tau)s_n(\tau)d\tau, \qquad (8.76)$$

thus identifying the optimal reconstruction functions.

Substituting (8.76) into (8.74), our goal becomes to maximize

$$\sum_{n=1}^{N} \left\| \int_0^T R_X(\cdot,\tau) s_n(\tau) d\tau \right\|^2 \tag{8.77}$$

with respect to the sampling functions $\{s_n(t)\}_{n=1}^N$. As we show in Appendix 8.B, the maximum of this expression is achieved by any set of kernels of the form

$$s_n(t) = \sum_{k=1}^{N} \boldsymbol{A}_{k,n} \left(\lambda_k + \sigma_c^2\right)^{-\frac{1}{2}} \psi_k(t), \tag{8.78}$$

where $\boldsymbol{A}$ is a unitary $N \times N$ matrix and λ_k and $\psi_k(t)$ are the eigenvalues and eigenfunctions of $R_x(t,\eta)$ respectively (see (8.68)). In particular, we can choose $\boldsymbol{A} = \boldsymbol{I}_N$, leading to

$$s_n(t) = \frac{1}{\sqrt{\lambda_n + \sigma_c^2}} \psi_n(t), \quad n = 1, \ldots, N. \tag{8.79}$$

From (8.76), the optimal reconstruction kernels are given by

$$v_n(t) = \frac{\lambda_n}{\sqrt{\lambda_n + \sigma_c^2}} \psi_n(t), \quad n = 1, \ldots, N. \tag{8.80}$$

The following theorem summarizes the result.

Theorem 8.4. *Let $x(t)$, $t \in [0,T]$ be a random process whose autocorrelation function $R_X(t,\eta)$ is jointly continuous in t and η. Assume that $y(t) = x(t) + w(t)$, where $w(t)$ is a white noise process uncorrelated with $x(t)$. Then, among all estimates $\hat{x}(t)$ of $x(t)$ having the form*

$$\hat{x}(t) = \sum_{n=1}^{N} v_n(t) \int_0^T s_n^*(\tau) y(\tau) dt \tag{8.81}$$

the MSE (8.8) *is minimized with $\{s_n(t)\}_{n=1}^N$ and $\{v_n(t)\}_{n=1}^N$ of* (8.79) *and* (8.80) *respectively. In these expressions, λ_n and $\psi_n(t)$ are the eigenvalues and eigenfunctions of $R_x(t,\eta)$ respectively (see* (8.68)*).*

Interestingly, the optimal sampling and reconstruction functions in our noisy setting are similar to those dictated by the KLT. The only difference is that in the present scenario, the nth sample is shrunk by a factor of $\lambda_n/(\lambda_n + \sigma_c^2)$ prior to reconstruction. This ensures that the low-SNR measurements do not contribute to the recovery as much as their high-SNR counterparts. From the viewpoint of designing the sampling mechanism, however, this difference is of no importance.

As stated above, in practice one would generally favor nonlinear processing of the samples (namely, applying standard nonlinear FRI techniques) rather than a simple element-wise shrinkage. Thus, the importance of Theorem 8.4 for our purposes is in identifying that the eigenfunctions of $R_X(t,\tau)$ remain the optimal sampling kernels even in the noisy setting.

8.8.3 Example: Sampling in a Subspace

To demonstrate the utility of Theorem 8.4, we now revisit the situation in which $x(t)$ is given by (8.38) for some set of linearly independent functions $\{g_k(t)\}_{k=1}^K$ spanning a subspace $\mathcal{G} \in L_2$. We assume that the coefficients $\boldsymbol{\theta} = \{a_1, \ldots, a_K\}^T$ form a zero-mean random vector and denote its autocorrelation matrix by $\boldsymbol{R}_{\boldsymbol{\theta}}$. In this case, the signal's autocorrelation function is given by

$$\begin{aligned} R_X(t,\eta) &= E\{x(t)x^*(\eta)\} \\ &= E\left\{ \sum_{k=1}^K a_k g_k(t) \sum_{\ell=1}^K a_\ell^* g_k^*(\eta) \right\} \\ &= \sum_{k=1}^K \sum_{\ell=1}^K g_k(t) g_\ell^*(\eta) (\boldsymbol{R}_{\boldsymbol{\theta}})_{k,\ell}. \end{aligned} \tag{8.82}$$

Consequently, the operator $R_X : L_2 \to L_2$ defined by $(R_x h)(t) = \int_0^T R_x(t,\eta) h(\eta) d\eta$ can be expressed as

$$R_x = G \boldsymbol{R}_{\boldsymbol{\theta}} G^*, \tag{8.83}$$

where G is the set transformation (8.2) associated with $\{g_k\}_{k=1}^K$.

Now, let $\boldsymbol{U}$ be a unitary matrix and let $\boldsymbol{D}$ be a diagonal matrix, such that

$$\boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^* = (G^*G)^{1/2} \boldsymbol{R}_{\boldsymbol{\theta}} (G^*G)^{1/2}. \tag{8.84}$$

Since the dimension of $\mathcal{R}(G)$ is K, the operator R_X has at most K nonzero eigenvalues $\{\lambda_k\}_{k=1}^K$. Let Ψ denote the set transformation associated with the N eigenfunctions $\{\psi_n\}_{n=1}^N$ corresponding to the N largest eigenvalues, for some $N \leq K$. Then, it can be shown that

$$\Psi = G(G^*G)^{-1/2} \boldsymbol{U} \tag{8.85}$$

and the corresponding eigenvalues are

$$\lambda_n = \boldsymbol{D}_{n,n}. \tag{8.86}$$

To see this, note that according to (8.85), Ψ is an isometry, since

$$\Psi^*\Psi = \boldsymbol{U}^*(G^*G)^{-1/2} G^*G (G^*G)^{-1/2} \boldsymbol{U} = \boldsymbol{U}^*\boldsymbol{U} = \boldsymbol{I}_K. \tag{8.87}$$

Furthermore, (8.83) and (8.84) imply that $R_X = \Psi \boldsymbol{D} \Psi^*$. Consequently

$$R_X \Psi = \Psi \boldsymbol{D} \Psi^* \Psi = \Psi \boldsymbol{D}, \tag{8.88}$$

which proves the claim.

It is important to emphasize that the K functions $\{\psi_n(t)\}_{n=1}^K$ span $\mathcal{G}$. Therefore, if one is allowed to take $N = K$ samples, then the optimal choice is a set of kernels that span $\mathcal{G}$. This conclusion is compatible with the CRB analysis of the previous sections. However, the advantage of the Bayesian viewpoint is that it allows us to identify the optimal sampling space when less than K samples are allowed. For example, suppose that $\{g_n(t)\}$ are orthonormal, and the coefficients $\{a_n\}$ are uncorrelated. Then the optimal sampling space is the one spanned by the N functions $\{g_n(t)\}$ corresponding to the N largest-variance coefficients $\{a_n\}$.

A second example demonstrating the derivation of the optimal sampling kernels will be given in the next section.

8.9 Application: Channel Estimation

In this section, we focus on a specific application of FRI signals, namely, that of estimating a signal consisting of a number of pulses having unknown positions and amplitudes [34,36,125]. More precisely, we consider periodic signals $x(t)$ of the form (8.16), which were discussed in Section 8.3.3. These are T-periodic pulse sequences, in which the pulse shape $g(t)$ is known, but the amplitudes $\{a_\ell\}$ and delays $\{t_\ell\}$ are unknown. After analyzing periodic signals of this type, we will also compare estimation performance in this case with the semi-periodic family (8.17), and attempt to explain the empirically observed differences in stability between these two cases.

By defining the T-periodic function $h(t) = \sum_{n\in\mathbb{Z}} g(t - nT)$, we can write $x(t)$ of (8.16) as

$$x(t) = \sum_{\ell=1}^{L} a_\ell h(t - t_\ell). \tag{8.89}$$

Our goal is now to estimate $x(t)$ from samples of the noisy process $y(t)$ of (8.6). As before, we will assume that only continuous-time noise is present in the system. Since $x(t)$ is T-periodic, it suffices to recover the signal in the region $[0, T]$. In particular, we would like to identify the optimal sampling kernels for this setting, and to compare existing algorithms with the resulting CRB in order to determine when the optimal estimation performance is achieved.

Let

$$\tilde{h}_k = \frac{1}{T}\langle h, \varphi_k\rangle, \quad k \in \mathbb{Z} \tag{8.90}$$

be the Fourier series of $h(t)$, where $\varphi_k(t) = e^{j2\pi kt/T}$. The Fourier series of $x(t)$ is then given by

$$\tilde{x}_k \triangleq \frac{1}{T}\langle x, \varphi_k\rangle = \tilde{h}_k \sum_{\ell=1}^{L} a_\ell e^{-j\frac{2\pi}{T}kt_\ell}, \quad k \in \mathbb{Z}. \tag{8.91}$$

Let $\mathcal{K} = \{k \in \mathbb{Z} : \tilde{h}_k \neq 0\}$ denote the indices of the nonzero Fourier coefficients of $h(t)$. Suppose for a moment that $\mathcal{K}$ is finite. It then follows from (8.91) that $x(t)$ also has a finite number of nonzero Fourier coefficients. Consequently, the set $\mathcal{X}$ of possible signals $x(t)$ is contained in the finite-dimensional subspace $\mathcal{M} = \text{span}\{\varphi_k\}_{k\in\mathcal{K}}$. Therefore, as explained in Section 8.7.4, choosing the $N = |\mathcal{K}|$ sampling kernels $\{s_n(t) = e^{-j2\pi nt/T}\}_{n\in\mathcal{K}}$ results in a sampled CRB which is equivalent to the continuous-time bound. This result is compatible with recent work demonstrating successful performance of FRI recovery algorithms using exponentials as sampling kernels [125].

Note, however, that this is a "Nyquist-equivalent" sampling scheme, i.e., the number of samples required $N = |\mathcal{K}|$ is potentially much higher than the number of degrees of freedom $2L$ in the signal $x(t)$ (see Section 8.7.4). This provides a theoretical explanation of the empirically recognized fact that sampling above the rate of innovation improves the performance of FRI techniques in the presence of noise [34], a fact which stands in contrast to the noise-free performance guarantees of many FRI algorithms.

Moreover, if there exists an infinite number of nonzero coefficients $\tilde{h}_k$, then in general the set $\mathcal{X}$ will not belong to any finite-dimensional subspace. Consequently, it will not be possible in this case for an algorithm based on a finite number of samples to achieve the performance obtainable from the complete signal $y(t)$. This occurs, for example, whenever the pulse $g(t)$ of (8.16) is time-limited. In such cases, any increase in the sampling rate will potentially continue to reduce the CRB, although the sampled CRB will converge to the asymptotic value of $\rho_{T_0}\sigma_c^2$ in the limit as the sampling rate increases.

8.9.1 Choosing the Sampling Kernels

An important question in the current setting is how to choose the sampling kernels so as to achieve the best possible performance under a limited budget of samples. This can be done via the Bayesian analysis provided in Section 8.8. Assume, for example, that the time delays $\{t_\ell\}_{\ell=1}^L$ are independently drawn from a uniform distribution over the interval $[0, T]$. Furthermore, suppose that the amplitudes $\{a_\ell\}_{\ell=1}^L$ are mutually uncorrelated zero-mean random

variables which are independent of the time delays and have variance σ_a^2. Then,

$$\begin{aligned} R_X(t,\tau) &= E\{x(t)x^*(\tau)\} \\ &= \sum_{k=1}^{L}\sum_{\ell=1}^{L} E\{a_k a_\ell^*\}\, E\{h(t-t_k)h^*(\tau-t_\ell)\} \\ &= \sigma_a^2 \sum_{\ell=1}^{L} E\{h(t-t_\ell)h^*(\tau-t_\ell)\} \\ &= \sigma_a^2 L \frac{1}{T}\int_0^T h(t-t_\ell)h^*(\tau-t_\ell)dt_\ell \\ &= \sigma_a^2 L \sum_{k\in\mathbb{Z}} |\tilde{h}_k|^2 e^{j\frac{2\pi}{T}k(t-\tau)}, \end{aligned} \tag{8.92}$$

where we used Parseval's theorem. It is easily verified that the eigenfunctions of $R_X(t,\tau)$ are given by

$$\psi_n(t) = \frac{1}{\sqrt{T}}\, e^{j\frac{2\pi}{T}nt}, \quad n \in \mathbb{Z} \tag{8.93}$$

and the corresponding eigenvalues are

$$\lambda_n = L\sigma_a^2 T|\tilde{h}_n|^2, \quad n \in \mathbb{Z}. \tag{8.94}$$

Therefore, the optimal set of N sampling functions is

$$s_n(t) = e^{j\frac{2\pi}{T}p_n t}, \quad n = 1,\ldots,N \tag{8.95}$$

where p_n is the index of the nth largest Fourier coefficient $|\tilde{h}_{p_n}|$. The optimal linear recovery of $x(t)$ from the resulting samples is given by

$$\hat{x}(t) = \sum_{n=1}^{N} c_n \frac{L\sigma_a^2|\tilde{h}_{p_n}|^2}{L\sigma_a^2 T|\tilde{h}_{p_n}|^2 + \sigma_c^2}\, e^{j\frac{2\pi}{T}p_n t}. \tag{8.96}$$

The performance of this estimator is poorer than state-of-the-art techniques, due to the restriction to linear reconstruction schemes. We recall that this technique is intended only for selecting the sampling kernels.

The above analysis again lends credence to the recently proposed time-delay estimation technique of Gedalyahu et al. [125], which makes use of complex exponentials as sampling functions. A disadvantage of this algorithm is that it can only handle a set of exponents with successive frequencies, while for general pulses, the indices of the N largest Fourier coefficients may be sporadic. As we will see in Section 8.9.3, this limitation may result in deteriorated performance of the algorithm in some cases.

8.9.2 Computing the CRB

Having identified the optimal sampling kernels (8.95), we would now like to compute the CRB for estimating $x(t)$ from the resulting samples. In order to compare these results with the continuous-time CRB, we assume that no digital noise is added in the sampling process. However, the calculations described below can be adapted without difficulty to situations containing both continuous-time and digital noise.

We assume for simplicity that $h(t)$ and $\{a_\ell\}$ are real-valued. Nonetheless, the sampling kernels chosen above are complex-valued, implying that Theorem 8.3 cannot be directly applied. Yet since $h(t)$ is real-valued, we have $|\tilde{h}_k| = |\tilde{h}_{-k}|$, and consequently the optimal sampling kernels consist of complex conjugate pairs $e^{\pm j2\pi nt/T}$. Recall that the sampling kernels can be changed without affecting the CRB, as long as the subspace they span remains constant. Consequently, the CRB can be computed for the equivalent sampling kernels $\sin(2\pi nt/T)$ and $\cos(2\pi nt/T)$, which are real and can therefore be used in conjunction with the results of Section 8.7. We note that since the transition to these real-valued kernels is unitary, the CRB will not change even if digital noise is added. To be specific without complicating the notation, we assume that N is odd and that the DC component is included among the sampling kernels chosen in (8.95). We can then define the equivalent set of kernels

$$\begin{aligned} \tilde{s}_0(t) &= 1, \\ \tilde{s}_n(t) &= \cos(2\pi p_n t/T), && n = 1, \ldots, \frac{N-1}{2}, \\ \tilde{s}_{n+\frac{N+1}{2}}(t) &= \sin(2\pi p_n t/T), && n = 1, \ldots, \frac{N-1}{2}. \end{aligned} \tag{8.97}$$

We further define the parameter vector

$$\boldsymbol{\theta} = (a_1, \ldots, a_L, t_1, \ldots, t_L)^T \tag{8.98}$$

whose length is $K = 2L$.

Theorem 8.3 provides a two-step process for computing the CRB of the signal $x(t)$ from its samples. First, the FIM $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ for estimating $\boldsymbol{\theta}$ is determined. Second, the formula (8.51) is applied to compute the CRB. While Theorem 8.3 also provides a means for calculating $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$, it is more convenient in the present setting to derive the FIM directly. This can be done by calculating the expectations μ_n of (8.44) and applying (8.52). In our setting, $\mu_n = \langle x, \tilde{s}_n \rangle$ are

given by

$$\mu_n = \frac{\tilde{x}_{p_n} + \tilde{x}_{-p_n}}{2}, \qquad n = 0, \ldots, \frac{N-1}{2}$$
$$\mu_{n+\frac{N+1}{2}} = \frac{\tilde{x}_{p_n} - \tilde{x}_{-p_n}}{2j}, \qquad n = 1, \ldots, \frac{N-1}{2} \tag{8.99}$$

where $\{\tilde{x}_n\}$ are the Fourier coefficients of $x(t)$. These coefficients depend in turn on the parameter vector $\boldsymbol{\theta}$, as shown in (8.91). Substituting μ_n into (8.52) yields a closed-form expression for $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$. Since the resulting formula is cumbersome and not very insightful, it is not explicitly written herein.

To obtain the sampled CRB, our next step is to compute the $2L \times 2L$ matrix

$$\boldsymbol{M} \triangleq \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right)^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right). \tag{8.100}$$

The function $h_{\boldsymbol{\theta}} : \mathbb{R}^{2L} \to L_2$ maps a given parameter vector $\boldsymbol{\theta}$ to the resulting signal $x(t)$ as defined by (8.89). Differentiating this function with respect to $\boldsymbol{\theta}$, we find that the operator $\partial h_{\boldsymbol{\theta}}/\partial \boldsymbol{\theta} : \mathbb{R}^{2L} \to L_2$ is defined by

$$\begin{aligned}\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right) \boldsymbol{v} = {} & \boldsymbol{v}_1 h(t - t_1) + \cdots + \boldsymbol{v}_L h(t - t_L) \\ & - \boldsymbol{v}_{L+1} a_1 h'(t - t_1) - \cdots - \boldsymbol{v}_{2L} a_L h'(t - t_L)\end{aligned} \tag{8.101}$$

for any vector $\boldsymbol{v} \in \mathbb{R}^{2L}$.

One may now compute the ikth element of $\boldsymbol{M}$ as

$$M_{ik} = \boldsymbol{e}_i^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right)^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}}\right) \boldsymbol{e}_k = \left\langle \frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \boldsymbol{e}_i, \frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \boldsymbol{e}_k \right\rangle. \tag{8.102}$$

Thus, each element of $\boldsymbol{M}$ is an inner product between two of the terms in (8.101). To calculate this inner product numerically for a given function $h(t)$, it is more convenient to use Parseval's theorem in order to convert the (continuous-time) inner product to a sum over Fourier coefficients. For example, in the case $1 \leq i, k \leq L$, we have

$$\begin{aligned} M_{ik} &= \int_0^T h(t - t_i) h^*(t - t_k) dt \\ &= T \sum_{n \in \mathbb{Z}} \tilde{h}_n e^{-j2\pi t_i n/T} \tilde{h}_n^* e^{j2\pi t_k n/T} \\ &= T \sum_{n \in \mathbb{Z}} |\tilde{h}_n|^2 e^{-j2\pi (t_k - t_i) n/T}, \quad 1 \leq i, k \leq L. \end{aligned} \tag{8.103}$$

An analogous derivation can be carried out when i or k are in the complementary range $L + 1, \ldots, 2L$.

Finally, having calculated the matrices $\boldsymbol{J}_{\boldsymbol{\theta}}^{\text{samp}}$ and $\boldsymbol{M}$, the CRB for sampled measurements is obtained using (8.51). We are now ready to compare this bound to the performance of practical estimators in some specific scenarios.

8.9.3 Effect of the Pulse Shape

In Fig. 8.1, we document several experiments comparing the CRB with the time-delay estimation technique of Gedalyahu et al. [125]. Specifically, we sampled the signal $x(t)$ of (8.16) using a set of exponential kernels, and used the matrix pencil method [154] to estimate $x(t)$ from the resulting measurements. Since we are considering only continuous-time noise, applying an invertible linear transformation to the sampling kernels has no effect on our performance bounds (see Section 8.7). The various kernels suggested in [125] amount to precisely such an invertible linear transformation, and the same performance bound applies to all of these approaches. Moreover, under the continuous-time noise model, it can be shown that these techniques also exhibit the same performance. For the same reason, the performance reported here is also identical to the method of Vetterli et al. [30].

In our experiments, a signal containing $L = 2$ pulses was constructed. The delays and amplitudes of the pulses were chosen randomly and are given by

$$\begin{aligned} a_1 &= 0.3204, & t_1 &= 0.6678, \\ a_2 &= 0.6063, & t_2 &= 0.9863. \end{aligned} \tag{8.104}$$

Modifications of these parameters does not appear to significantly affect the reported results, except when the time delays are close to one another, a situation which will be discussed in depth in Section 8.9.4. The pulse $h(t)$ consisted of $|\mathcal{K}| = 401$ nonzero Fourier coefficients at positions $\mathcal{K} = \{-200, \ldots, 200\}$. The CRB is plotted as a function of the number of samples N, where the sampling kernels are given by $s_n(t) = e^{j2\pi nt/T}$ with $n \in \{-\lfloor N/2 \rfloor, \ldots, \lfloor N/2 \rfloor\}$. This is done because the matrix pencil method requires the sampling kernels to have contiguous frequencies.

In Fig. 8.1(a), we chose $\tilde{h}_k = 1$ for $-200 \leq k \leq 200$ and $\tilde{h}_k = 0$ elsewhere; these are the low-frequency components of a Dirac delta function. The noise standard deviation was $\sigma_c = 10^{-5}$. In this case, for a fixed budget of N samples, any choice of N exponentials having frequencies in the range $-200 \leq k \leq 200$ is optimal according to the criterion of Section 8.8. As expected, the sampled CRB achieves the continuous-time bound $K\sigma_c^2$ when $N \geq |\mathcal{K}|$. However, the CRB obtained at low sampling rates is higher by several orders of magnitude than the continuous-time limit. This indicates that the maxim of FRI theory, whereby sampling at the rate of innovation suffices for reconstruction, may not always hold in the presence of mild levels of noise. Indeed, if no noise is added in the present setting, then perfect recovery can be guaranteed using as few as $N = 4$ samples; yet even in the presence of mild noise, our

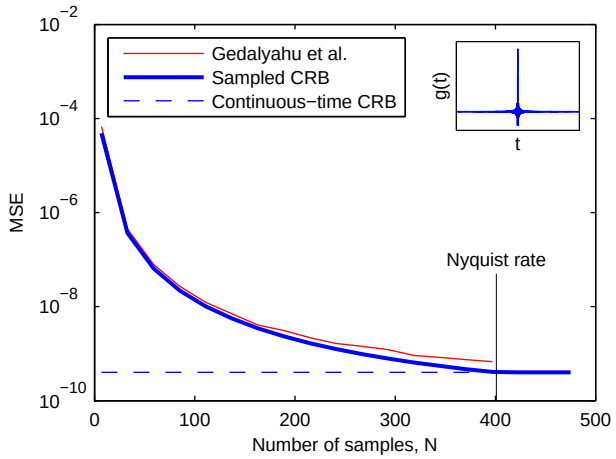

(a) The pulse $g(t)$ is a filtered Dirac with 401 Fourier coefficients.

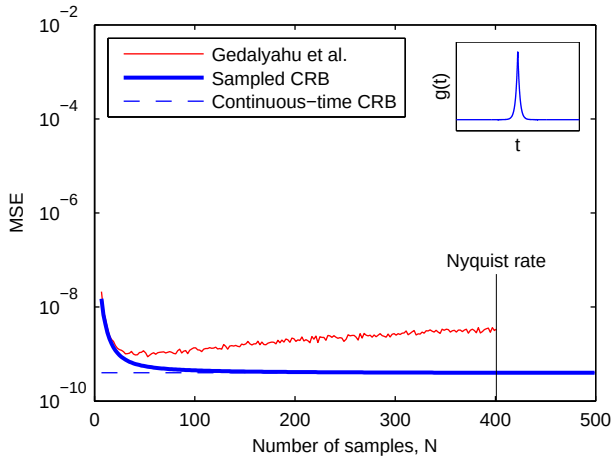

(b) The pulse $g(t)$ contains 401 nonzero Fourier coefficients which decrease monotonically with frequency.

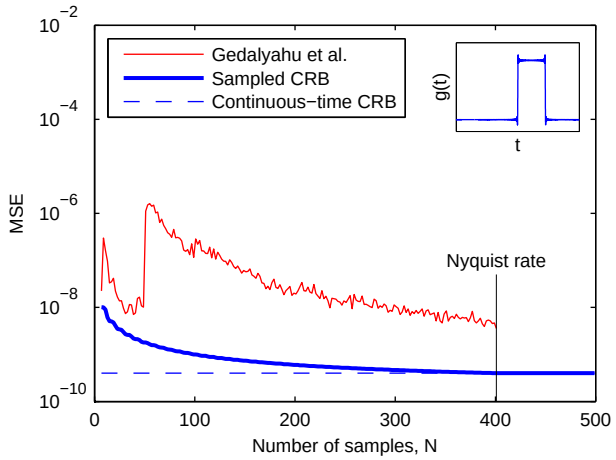

(c) The pulse $g(t)$ is a filtered rect$(\cdot)$ with 401 Fourier coefficients.

Figure 8.1: Comparison of the CRB and the performance of a practical estimator, as a function of the number of samples.

bounds demonstrate that performance is quite poor unless the number of samples is increased substantially. This result may provide an explanation for the previously observed numerical instability of FRI techniques [30, 34].

As a further observation, we note that in this scenario, existing algorithms come very close to the CRB. Thus, the previously observed improvements achieved by oversampling are a result of fundamental limitations of low-rate sampling, rather than drawbacks of the specific technique used.

The same experiment is repeated in Fig. 8.1(b) with a pulse having Fourier coefficients $\tilde{h}_k = 1/(1 + 0.01k^2)$. Since the Fourier coefficients decrease with $|k|$, in this case our choice of low-frequency sampling kernels is optimal. However, the SNR of the measurements c_n decreases with n. As can be seen, this has a negative effect on the performance of the algorithm, which is not designed for high noise levels. Indeed, including low-SNR measurements causes the MSE not only to depart from the CRB, but eventually even to increase as more noisy samples are provided. In other words, one would do better to ignore the high-frequency measurements than to feed them to the recovery algorithm. Yet information is clearly present in these high-frequency samples, as indicated by the continual decrease of the CRB with increasing N. Thus, our analysis indicates that improved estimation techniques should be achievable in this case, in particular by careful utilization of low-SNR measurements.

The adverse effect of low-SNR measurements is exacerbated if, for a given N, one does not choose the N largest Fourier coefficients. This is demonstrated in Fig. 8.1(c). Here, the results of a similar experiment are plotted, in which $\tilde{h}_k = P\,\mathrm{sinc}(nP/T)$, $-200 \leq k \leq 200$. These are the 401 lowest-frequency Fourier coefficients of a rectangular pulse having width P. In this case, the Fourier coefficients are no longer monotonically decreasing with $|k|$. Consequently, the sampling kernels $s_n(t) = e^{j2\pi nt/T}$ with $n \in \{-\lfloor N/2 \rfloor, \ldots, \lfloor N/2 \rfloor\}$ do not correspond to the N largest Fourier coefficients, and thus are not optimal. In particular, for the chosen parameters, $|\tilde{h}_{25}| = |\tilde{h}_{-25}|$ are considerably smaller than the rest of the coefficients. When $N \geq 50$, the corresponding measurements are included, causing the MSE to deteriorate significantly.

8.9.4 Closely-Spaces Pulses

It is well-known that the estimation of pulse positions becomes ill-conditioned when several of the pulses are located close to one another. Intuitively, this is a consequence of the overlap between the pulses, which makes it more difficult to identify the precise location of each pulse. However, our goal is to estimate the signal $x(t)$ itself, rather than the positions of its constituent

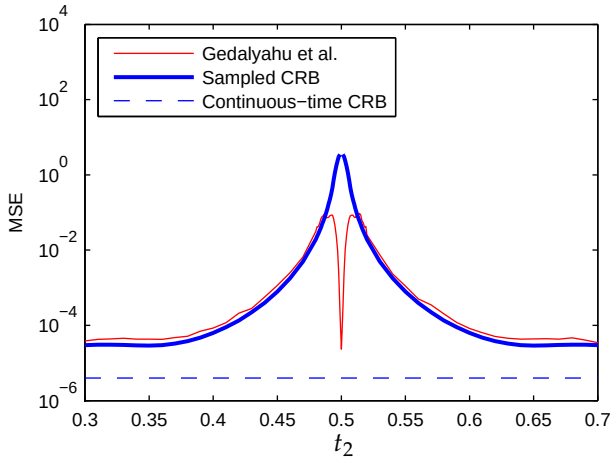

Figure 8.2: Comparison between the CRB and the performance of a practical estimator as a function of the pulse positions. The signal contains $L = 2$ pulses, the first of which is located at $t_1 = 0.5$. The MSE is plotted as a function of the position of the second pulse.

pulses. As we will see, for this purpose the effect of closely-spaced pulses is less clear-cut.

To study the effect of pulse position on the estimation error, we used a setup similar to the one of Fig. 8.1(b), with the following differences. First, a higher noise level of $\sigma_c = 10^{-3}$ was chosen. Second, the signal consisted of $L = 2$ pulses, with the first pulse at position $t_1 = 0.5$. The position of the second pulse was varied in the range $[0.3, 0.7]$ to demonstrate the effect of pulse proximity on the performance. The setting was otherwise identical to that of Section 8.9.3. In particular, recall that $T = 1$.

The results of this experiment are plotted in Fig. 8.2, which documents both the values of the sampled CRB and the actual MSE obtained by the estimator of Gedalyahu et al. [125]. The continuous-time CRB is also plotted, although, as is evident from Theorem 8.2, this bound is a function only of the number of parameters determining the signal, and is therefore unaffected by the proximity of the pulses.

Several different effects are visible in Fig. 8.2. First, as the two pulses begin to come closer, both the CRB and the observed MSE increase by several orders of magnitude; this occurs when $|t_1 - t_2|$ is between about 0.15 and 0.03. (Of course, the precise distances at which these effects occur depend on the pulse width and other parameters of the experiment.) This level of proximity is demonstrated in Fig. 8.3(a). At this stage, the overlap between the pulses is sufficient to make it more difficult to estimate their positions accurately, but the separation between the pulses is still large, so that they are not mistaken for a single pulse.

As the pulses draw nearer each other, they begin to resemble a single pulse located at

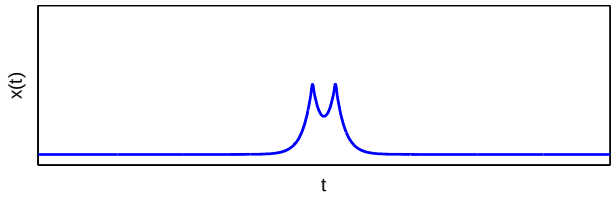

(a) The spacing between the pulses is 0.04.

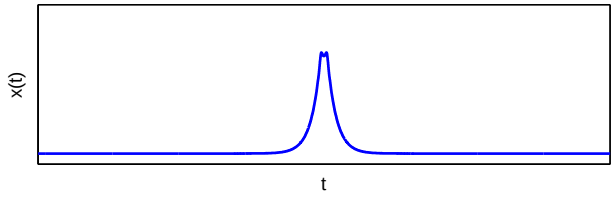

(b) The spacing between the pulses is 0.01.

Figure 8.3: Demonstration of the different levels of overlap between pulses.

$(t_1 + t_2)/2$ (see Fig. 8.3(b)). Depending on the noise level, at some point the estimation algorithm will indeed identify the two pulses as one. Since our goal is to estimate $x(t)$ and not the pulse positions, such an "error" causes little deterioration in MSE. This is visible in Fig. 8.2 as the region in which the MSE of the practical algorithm ceases to deteriorate and ultimately decreases.

Interestingly, the CRB does not capture this improvement in performance. This failure is due to the fact that the CRB applies only to unbiased estimators, while the strategy utilized in [125] becomes biased for closely-spaced pulses. For an estimator to be unbiased, it is required that the mean estimate, averaged over noise realizations, will converge to the true value of $x(t)$, which has a form similar to that of Fig. 8.3(b). The expectation of an estimator reconstructing a single pulse will not have the form of two closely-spaced pulses; such an estimator is thus necessarily biased. In other words, the discrepancy observed here results from the fact that in this case, biased techniques outperform the best unbiased approach.

8.9.5 Non-Periodic and Semi-Periodic Signal Models

As we have seen above, the reconstruction of signals of the form (8.16) in the presence of noise is often severely hampered when sampled at or slightly above the rate of innovation. Rather than indicating a lack of appropriate algorithms, in many cases this phenomenon results from fundamental limits on the ability to recover such signals from noisy measurements. A similar effect was demonstrated [34] in the non-periodic (or finite) pulse stream model (8.14). In fact, if one is allowed to sample a non-periodic pulse stream with arbitrary sampling kernels, then by

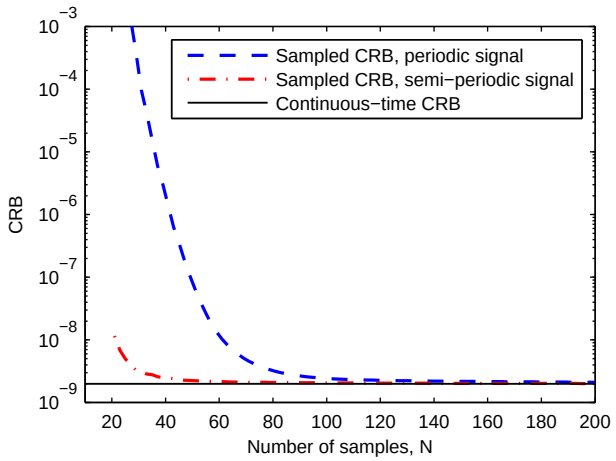

Figure 8.4: Comparison between the CRB for a periodic signal (8.16) and a semi-periodic signal (8.17).

designing kernels having sufficiently large time-domain support, one can capture all or most of the energy in the signal. This setting then essentially becomes equivalent to a periodic signal model (8.16) in which the period is larger than the effective support of the pulse stream: One can imagine that the signal repeats itself beyond the sampled region, as this would not affect the measurements. Consequently, it is not surprising that the non-periodic model demonstrates substantial improvement in the presence of oversampling [34].

On the other hand, some types of FRI and union of subspace signals exhibit remarkable noise resilience, and do not appear to require substantial oversampling in the presence of noise [36, 126]. As we now show, the CRB can be used to verify that such phenomena arise from a fundamental difference between families of FRI signals.

As an example, we compare the CRB for reconstructing the periodic signal (8.16) with the semi-periodic signal (8.17). Recall that in the former case, each period consists of pulses having unknown amplitudes and time shifts. By contrast, in the latter signal, the time delays are identical throughout all periods, but the amplitudes can change from one period to the next.

While these are clearly different types of signals, an effort was made to form a fair comparison between the reconstruction capabilities in the two cases. To this end, we chose an identical pulse $g(t)$ in both cases. We selected the signal segment $[0, T_0]$, where $T_0 = 1$, and chose the signal parameters so as to guarantee an identical T_0-local rate of innovation. We also used identical sampling kernels in both settings: specifically, we chose the kernels (8.97) which measure the N lowest frequency components of the signal.

To simplify the analysis and focus on the fundamental differences between these settings,

we will assume in this section that the pulses $g(t)$ are compactly supported, and that the time delays are chosen such that pulses from one period do not overlap with other periods. In other words, if the support of $g(t)$ is given by $[t_a, t_b]$, then we require

$$\min_{\ell}\{t_\ell\} > t_a \quad \text{and} \quad \max_{\ell}\{t_\ell\} < T - t_b. \tag{8.105}$$

Specifically, we chose the pulse $g(t)$ used in Fig. 8.1(b), which is compactly supported to a high approximation.

For the periodic signal, we chose $L = 10$ pulses with random delays and amplitudes, picked so as to satisfy the condition (8.105). A period of $T = 1$ was selected. This implies that the signal of interest is determined by $2L = 20$ parameters (L amplitudes and L time delays).

To construct a semi-periodic signal with the same number of parameters, we chose a period of $T = 1/9$ containing $L = 2$ pulses. The segment $[0, T_0]$ then contains precisely $M = 9$ periods, for a total of 20 parameters. While it may seem plausible to require the same number of periods for both signals, this would actually disadvantage the periodic approach, as it would require the estimation of much more closely-spaced pulses.

The CRB for the periodic signal was computed as explained in Section 8.9.2, and the CRB for the semi-periodic signal can be calculated in a similar fashion. The results are compared with the continuous-time CRB in Fig. 8.4. Note that since the number of parameters to be estimated is identical in both signal models, the continuous-time CRB for the two settings coincides. Consequently, for a large number of measurements, the sampled bounds also converge to the same values. However, when the number of samples is closer to the rate of innovation, the bound on the reconstruction error for the semi-periodic signal is much lower than that of the periodic signal. As mentioned above, this is in agreement with previously reported findings for the two types of signals [30, 36, 125].

To find an explanation for this difference, it is helpful to recall that both signals can be described using the union of subspaces viewpoint (see Section 8.3.3). Each of the signals in this experiment is defined by precisely 20 parameters, which determine the subspace to which the signal belongs and the position within this subspace. Specifically, the values of the time delays select the subspace, and the pulse amplitudes define a point within this subspace. Thus, in the above setting, the periodic signal contains 10 parameters for selecting the subspace and 10 additional parameters determining the position within it; whereas for the semi-periodic signal, only 2 parameters determine the subspace while the remaining 18 parameters set the location in the subspace. Evidently, identification of the subspace is challenging, especially in the presence

of noise, but once the subspace is determined, the remaining parameters can be estimated using a simple linear operation (a projection onto the chosen subspace). Consequently, if many of the unknown parameters identify the position within a subspace, estimation can be performed more accurately. This may provide an explanation for the difference between the two examined signal models.

As further evidence in support of this explanation, we recall from Section 8.3.3 that the multiband signal model (8.19) can also be viewed as a union of subspaces. Here, again, the parameters $\{\omega_\ell\}_{\ell=1}^{L}$ determining the subspace (i.e., the utilized frequency bands) are far fewer than the parameters $\{a_\ell[n]\}$ selecting the point within the subspace (i.e., the content of each frequency band). In support of the proposed explanation, highly noise resistant algorithms can be constructed for the recovery of multiband signals [126, 127]. An even more extreme case is the single subspace setting, exemplified by shift-invariant signals (Section 8.3.1). In this case, all of the signal parameters are used to determine the position within the one possible subspace. As we have seen in Section 8.7.3, in this case Nyquist-equivalent sampling at the rate of innovation achieves the continuous-time CRB.

8.10 Conclusion

In this paper, we studied the inherent limitations in recovering FRI signals from noisy measurements. We derived a continuous-time CRB which provides a lower bound on the achievable MSE of any unbiased estimation method, regardless of the sampling mechanism. We showed that the rate of innovation ρ_{T_0} is a lower bound on the ratio between the average MSE achievable by any unbiased estimator and the noise variance σ_c^2, *regardless of the sampling method*. This stands in contrast to the noise-free interpretation of ρ_{T_0} as the minimum sampling rate required for perfect recovery.

We next examined the CRB for estimating an FRI signal from a discrete set of noisy samples. We showed that the sampled bound is in general higher than the continuous-time CRB, and approaches it as the sampling rate increases. In general, the rate which is needed in order to achieve the continuous-time CRB is equal to the rate associated with the smallest subspace that encompasses all possible signal realizations. In particular, if a signal belongs to a union of subspaces, then the rate required to achieve the continuous-time bound is that associated with the sum of the subspaces. In some cases, this rate is finite, but in other cases the sum covers the entire space L_2 and no finite-rate technique achieves the CRB.

A consequence of these results is that oversampling can generally improve estimation performance. Indeed, our experiments demonstrate that sampling rates much higher than ρ_{T_0} are required in certain settings in order to approach the optimal performance. Furthermore, these gains can be substantial: In some cases, oversampling can improve the MSE by several orders of magnitude. We showed that the CRB can be used to determine which estimation problems require substantial oversampling to achieve stable performance. As a rule of thumb, it appears that for union of subspace signals, performance is improved at low rates if most of the parameters identify the position within the subspace, rather than the subspace itself. Our analysis can also be used to identify cases in which no existing algorithm comes close to the CRB, implying that better approaches can be constructed. In particular, it seems that existing algorithms do not deal well with measurement sets having a wide dynamic range.

Lastly, we addressed the problem of choosing the sampling kernels. This was done by adopting a Bayesian framework, so that an optimality criterion can be rigorously defined. Using a generalization of the KLT, we showed that the optimal kernels are the eigenfunctions of the autocorrelation function of the signal. In the context of time-delay estimation, these kernels are exponentials with appropriately chosen frequencies. This choice coincides with recent FRI techniques [36].

8.A Proof of Theorem 8.1

The following notation will be used within this appendix. Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be two measurable Hilbert spaces, and let $(\Omega, \mathscr{F}, P)$ be a probability space. Consider two random variables $u : \Omega \to \mathcal{H}_1$ and $v : \Omega \to \mathcal{H}_2$. Then, the notation $E\{uv^*\}$ will be used to denote the linear operator $\mathcal{H}_2 \to \mathcal{H}_1$ such that, for any $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$,

$$\langle h_1, E\{uv^*\}\, h_2\rangle_{\mathcal{H}_1} = E\Big\{ \langle h_1, u\rangle_{\mathcal{H}_1} \langle v, h_2\rangle_{\mathcal{H}_2} \Big\} \tag{8.106}$$

if the expectation exists for all h_1 and h_2.

We begin by stating two general lemmas which will be of use in the proof of Theorem 8.1.

Lemma 8.5. *Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be two Hilbert spaces, and consider the operators*

$$\begin{aligned} A &: \mathcal{H}_1 \to \mathcal{H}_1, \\ B &: \mathcal{H}_2 \to \mathcal{H}_1, \\ C &: \mathcal{H}_2 \to \mathcal{H}_2. \end{aligned} \tag{8.107}$$

Suppose C is self-adjoint and invertible. Define the product Hilbert space $\mathcal{H}_1 \times \mathcal{H}_2$ in the usual manner, and suppose the operator $M : \mathcal{H}_1 \times \mathcal{H}_2 \to \mathcal{H}_1 \times \mathcal{H}_2$ defined by

$$M \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} = \begin{pmatrix} Ah_1 + Bh_2 \\ B^* h_1 + Ch_2 \end{pmatrix} \tag{8.108}$$

is positive semidefinite (psd). Then,

$$A \succeq BC^{-1}B^* \tag{8.109}$$

in the sense that the $\mathcal{H}_1 \to \mathcal{H}_1$ operator $A - BC^{-1}B^*$ is psd.

Proof. Since M is psd, we have for any $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$

$$\left\langle \begin{pmatrix} h_1 \\ h_2 \end{pmatrix}, M \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} \right\rangle_{\mathcal{H}_1 \times \mathcal{H}_2} \geq 0 \tag{8.110}$$

which implies

$$\langle h_1, Ah_1 \rangle_{\mathcal{H}_1} + 2\Re \left[\langle h_1, Bh_2 \rangle_{\mathcal{H}_2} \right] + \langle b, Cb \rangle_{\mathcal{H}_2} \geq 0. \tag{8.111}$$

Choosing $h_2 = -C^{-1}B^*h_1$, we have that $\langle h_1, Bh_2 \rangle_{\mathcal{H}_2} = -\left\langle B^*h_1, C^{-1}B^*h_1 \right\rangle_{\mathcal{H}_1}$, which is real since C^{-1} is self-adjoint. It follows from (8.111) that

$$\langle h_1, Ah_1 \rangle_{\mathcal{H}_1} - \left\langle h_1, BC^{-1}B^*h_1 \right\rangle_{\mathcal{H}_1} \geq 0 \tag{8.112}$$

which leads to (8.109), as required. □

Lemma 8.6. *Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be two Hilbert spaces and let $(\Omega, \mathcal{F}, P)$ be a probability space. Let $u : \Omega \to \mathcal{H}_1$ and $v : \Omega \to \mathcal{H}_2$ be random variables, and suppose the expectations $E\{uu^*\}$, $E\{uv^*\}$, and $E\{vv^*\}$ exist as linear operators as defined in* (8.106)*. If $E\{vv^*\}$ is invertible, then*

$$E\{uu^*\} \succeq E\{uv^*\} \left(E\{vv^*\}\right)^{-1} E\{vu^*\}. \tag{8.113}$$

Proof. Let us denote $A = E\{uu^*\}$, $B = E\{uv^*\}$, and $C = E\{vv^*\}$ and define the linear operator

$M : \mathcal{H}_1 \times \mathcal{H}_2 \to \mathcal{H}_1 \times \mathcal{H}_2$ as in (8.108). From (8.106), for any $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$ we have

$$\begin{aligned}
&\left\langle \begin{pmatrix} h_1 \\ h_2 \end{pmatrix}, M \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} \right\rangle_{\mathcal{H}_1 \times \mathcal{H}_2} \\
&= \langle h_1, Ah_1 \rangle_{\mathcal{H}_1} + 2\Re\left[\langle h_1, Bh_2 \rangle_{\mathcal{H}_1} \right] + \langle b, Cb \rangle_{\mathcal{H}_2} \\
&= \mathbb{E}\Big\{ \left| \langle h_1, u \rangle_{\mathcal{H}_1} \right|^2 + 2\Re\left[\langle h_1, u \rangle_{\mathcal{H}_1} \langle v, h_2 \rangle_{\mathcal{H}_2} \right] \\
&\qquad + \left| \langle h_2, v \rangle_{\mathcal{H}_2} \right|^2 \Big\} \\
&= E\left\{ \left| \langle h_1, u \rangle_{\mathcal{H}_1} + \langle v, h_2 \rangle_{\mathcal{H}_2} \right|^2 \right\} \\
&\geq 0. \qquad (8.114)
\end{aligned}$$

Thus M is a psd operator. Invoking Lemma 8.5 yields (8.113), as required. □

We are now ready to prove Theorem 8.1.

Proof of Theorem 8.1. Throughout the proof, let $\boldsymbol{\theta}$ be a fixed parameter and consider all functions as implicitly dependent on $\boldsymbol{\theta}$. Define the random variables

$$u : \Omega \to \mathcal{H} : \qquad u(\omega) = \hat{x}(y(\omega)) - h_{\boldsymbol{\theta}}, \qquad (8.115)$$

$$v : \Omega \to \mathbb{R}^K : \qquad v(\omega) = \frac{\partial \log p_{\boldsymbol{\theta}}(y(\omega))}{\partial \boldsymbol{\theta}}. \qquad (8.116)$$

We then have the linear operators $E\{vv^*\} : \mathbb{R}^K \to \mathbb{R}^K$, $E\{uu^*\} : \mathcal{H} \to \mathcal{H}$, and $E\{uv^*\} : \mathbb{R}^K \to \mathcal{H}$, which satisfy

$$E\{vv^*\} = \boldsymbol{J}_{\boldsymbol{\theta}}, \qquad (8.117)$$

$$\left\langle \varphi_i, E\{uu^*\}\, \varphi_j \right\rangle = E\left\{ \langle \varphi_i, u \rangle \langle u, \varphi_j \rangle \right\}, \qquad (8.118)$$

$$\left\langle \varphi_i, E\{uv^*\}\, \boldsymbol{e}_j \right\rangle = E\left\{ \langle \varphi_i, u \rangle \frac{\partial \log p_{\boldsymbol{\theta}}(y)}{\partial \theta_j} \right\}, \qquad (8.119)$$

where $\{\varphi_n\}_{n \in \mathbb{Z}}$ denotes a complete orthonormal basis for $\mathcal{H}$. The operator $E\{uu^*\}$ can be thought of as the covariance of $\hat{x}$, and is well-defined since, by (8.32), $\hat{x}$ has finite variance. Indeed, we have

$$\sum_{i \in \mathbb{Z}} \langle \varphi_i, E\{uu^*\}\, \varphi_i \rangle = E\left\{ \|\hat{x} - h_{\boldsymbol{\theta}}\|_{L_2}^2 \right\} < \infty \qquad (8.120)$$

so that $E\{uu^*\}$ is not only well-defined, but a trace class operator. Furthermore, $E\{vv^*\} = \boldsymbol{J}_{\boldsymbol{\theta}}$ is well-defined and invertible by Assumption P5. The operator $E\{uv^*\}$ is thus also well-defined by virtue of the Cauchy–Schwarz inequality.

To prove the theorem, we will show that

$$E\{uv^*\} = \frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \tag{8.121}$$

and then obtain the required result by applying Lemma 8.6. To demonstrate (8.121), observe that

$$\begin{aligned}
\langle \varphi_i, E\{uv^*\}\, \boldsymbol{e}_j \rangle &= E\left\{ \langle \varphi_i, u \rangle \, \frac{\partial \log p_{\boldsymbol{\theta}}(y)}{\partial \theta_j} \right\} \\
&= \int \langle \varphi_i, \hat{x}(y) - h_{\boldsymbol{\theta}} \rangle \, \frac{1}{p(y;\boldsymbol{\theta})} \frac{\partial p(y;\boldsymbol{\theta})}{\partial \theta_j} p(y;\boldsymbol{\theta}) P_{\boldsymbol{\theta}_0}(dy) \\
&= \int \langle \varphi_i, \hat{x}(y) - h_{\boldsymbol{\theta}} \rangle \lim_{\Delta \to 0} \frac{p(y;\boldsymbol{\theta} + \Delta \boldsymbol{e}_j) - p(y;\boldsymbol{\theta})}{\Delta} P_{\boldsymbol{\theta}_0}(dy).
\end{aligned} \tag{8.122}$$

By Assumption P4, for any sufficiently small $\Delta > 0$ we have

$$\begin{aligned}
&\left| \langle \varphi_i, \hat{x}(y) - h_{\boldsymbol{\theta}} \rangle \, \frac{p(y;\boldsymbol{\theta} + \Delta \boldsymbol{e}_j) - p(y;\boldsymbol{\theta})}{\Delta} \right| \\
&\quad \le |\langle \varphi_i, \hat{x}(y) - h_{\boldsymbol{\theta}} \rangle| \, q(y, \boldsymbol{\theta}).
\end{aligned} \tag{8.123}$$

Let us demonstrate that the right-hand side of (8.123) is absolutely integrable. By the Cauchy–Schwarz inequality,

$$\begin{aligned}
&\left(\int |\langle \varphi_i, \hat{x}(y) - h_{\boldsymbol{\theta}} \rangle| \, q(y, \boldsymbol{\theta}) P_{\boldsymbol{\theta}_0}(dy) \right)^2 \\
&\quad \le \int |\langle \varphi_i, \hat{x}(y) - h_{\boldsymbol{\theta}} \rangle|^2 \, P_{\boldsymbol{\theta}_0}(dy) \cdot \int q^2(y, \boldsymbol{\theta}) P_{\boldsymbol{\theta}_0}(dy).
\end{aligned} \tag{8.124}$$

The rightmost integral in (8.124) is finite by virtue of (8.29). As for the remaining integral, we have

$$\begin{aligned}
&\int |\langle \varphi_i, \hat{x}(y) - h_{\boldsymbol{\theta}} \rangle|^2 \, P_{\boldsymbol{\theta}_0}(dy) \\
&\quad \le \int \|\hat{x}(y) - h_{\boldsymbol{\theta}}\|^2 P_{\boldsymbol{\theta}_0}(dy) \\
&\quad \overset{\text{(a)}}{\le} \int \left(\|\hat{x}(y)\| + \|h_{\boldsymbol{\theta}}\| \right)^2 P_{\boldsymbol{\theta}_0}(dy) \\
&\quad \overset{\text{(b)}}{\le} \int \|\hat{x}(y)\|^2 P_{\boldsymbol{\theta}_0}(dy) + \|h_{\boldsymbol{\theta}}\|^2 \int P_{\boldsymbol{\theta}_0}(dy) \\
&\qquad\quad + 2\|h_{\boldsymbol{\theta}}\| \left(\int \|\hat{x}(y)\|^2 P_{\boldsymbol{\theta}_0}(dy) \right)^{1/2} \\
&\quad \overset{\text{(c)}}{<} \infty
\end{aligned} \tag{8.125}$$

where we have used the triangle inequality in (a), the Cauchy–Schwarz inequality in (b), and the assumption (8.32) that $\hat{x}$ has finite energy in (c). We conclude that (8.123) is bounded by an

absolutely integrable function, and we can thus apply the dominated convergence theorem to (8.122), obtaining

$$\begin{aligned}\left\langle \varphi_i, E\{uv^*\}\, \boldsymbol{e}_j \right\rangle = \frac{\partial}{\partial \theta_j} \int \left\langle \varphi_i, \hat{x}(y) \right\rangle p(y;\boldsymbol{\theta}) P_{\theta_0}(dy) \\ - \left\langle \varphi_i, h_{\boldsymbol{\theta}} \right\rangle \frac{\partial}{\partial \theta_j} \int p(y;\boldsymbol{\theta}) P_{\theta_0}(dy).\end{aligned} \tag{8.126}$$

The second integral in (8.126) equals 1 and its derivative is therefore 0. Thus we have

$$\left\langle \varphi_i, E\{uv^*\}\, \boldsymbol{e}_j \right\rangle = \frac{\partial E\{\langle \varphi_i, \hat{x}(y)\rangle\}}{\partial \theta_j} = \frac{\partial \left\langle \varphi_i, h_{\boldsymbol{\theta}} \right\rangle}{\partial \theta_j}. \tag{8.127}$$

On the other hand, note that the Fréchet derivative $\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta}$ of (8.31) coincides with the Gâteaux derivative of $h_{\boldsymbol{\theta}}$. In other words, for any vector $\boldsymbol{v} \in \mathbb{R}^K$, we have

$$\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \boldsymbol{v} = \lim_{\varepsilon \to 0} \frac{h_{\boldsymbol{\theta}+\varepsilon \boldsymbol{v}} - h_{\boldsymbol{\theta}}}{\varepsilon}. \tag{8.128}$$

It follows that

$$\left\langle \varphi_i, \frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \boldsymbol{e}_j \right\rangle = \frac{\partial \left\langle \varphi_i, h_{\boldsymbol{\theta}} \right\rangle}{\partial \theta_j}. \tag{8.129}$$

Since $E\{uv^*\}$ and $(\partial h_{\boldsymbol{\theta}} / \partial \boldsymbol{\theta})^*$ are both linear operators, (8.127) and (8.129) imply that the two operators are equal, demonstrating (8.121). Applying Lemma 8.6 and using the results (8.117) and (8.121), we have

$$E\{(\hat{x} - h_{\boldsymbol{\theta}})(\hat{x} - h_{\boldsymbol{\theta}})^*\} \succeq \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}_{\boldsymbol{\theta}}^{-1} \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^*. \tag{8.130}$$

As we have seen, the left-hand side of (8.130) is trace class, and thus so is the right-hand side. Taking the trace of both sides of the equation, we obtain

$$E\left\{ \|\hat{x} - h_{\boldsymbol{\theta}}\|^2 \right\} \geq \operatorname{Tr}\left(\left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}_{\boldsymbol{\theta}}^{-1} \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^* \right) \tag{8.131}$$

which is equivalent to (8.33), as required. □

8.B Maximization of (8.77)

The task of maximizing (8.77) is most easily accomplished by optimizing the coordinates of $s_n(t)$ in the orthonormal basis for $L_2[0,T]$ generated by the eigenfunctions of $R_X(t,\eta)$. Specifically, the function $s_n(t)$ can be written as

$$s_n(t) = \sum_{k=1}^{\infty} \alpha_k^n \left(\lambda_k + \sigma_c^2 \right)^{-\frac{1}{2}} \psi_k(t), \tag{8.132}$$

with $\{\psi_k(t)\}_{k=1}^{\infty}$ and $\{\lambda_k\}_{k=1}^{\infty}$ of (8.68). (The coefficients $(\lambda_k+\sigma_c^2)^{-1/2}$ are inserted since they simplify the subsequent analysis.) Now, by Mercer's theorem, $R_X(t,\eta)$ can be expressed as

$$R_X(t,\eta)=\sum_{\ell=1}^{\infty}\lambda_\ell\psi_\ell(t)\psi_\ell^*(\eta), \tag{8.133}$$

where the convergence is absolute and uniform. Therefore

$$\begin{aligned}&\int_0^T R_X(t,\tau)s_n(\tau)d\tau\\&\quad=\int_0^T\sum_{k=1}^{\infty}\alpha_k^n\left(\lambda_k+\sigma_c^2\right)^{-\frac{1}{2}}\psi_k(\tau)\sum_{\ell=1}^{\infty}\lambda_\ell\psi_\ell(t)\psi_\ell^*(\tau)d\tau\\&\quad=\sum_{k=1}^{\infty}\alpha_k^n\frac{\lambda_k}{(\lambda_k+\sigma_c^2)^{\frac{1}{2}}}\psi_k(t),\end{aligned} \tag{8.134}$$

and consequently, by Parseval's theorem, (8.77) is given by

$$\begin{aligned}&\sum_{n=1}^{N}\int_0^T\left|\int_0^T R_X(t,\tau)s_n(\tau)d\tau\right|^2dt\\&\qquad=\sum_{n=1}^{N}\int_0^T\left|\sum_{k=1}^{\infty}\alpha_k^n\frac{\lambda_k}{(\lambda_k+\sigma_c^2)^{\frac{1}{2}}}\psi_k(t)\right|^2dt\\&\qquad=\sum_{n=1}^{N}\sum_{k=1}^{\infty}|\alpha_k^n|^2\frac{\lambda_k^2}{\lambda_k+\sigma_c^2}.\end{aligned} \tag{8.135}$$

Similarly, using (8.132) and (8.134), we have

$$\iint_0^T s_m^*(t)R_X(t,\tau)s_n(\tau)d\tau dt=\sum_{k=1}^{\infty}\alpha_k^n(\alpha_k^m)^*\frac{\lambda_k}{\lambda_k+\sigma_c^2} \tag{8.136}$$

and by (8.132)

$$\sigma_c^2\int_0^T s_m^*(t)s_n(\tau)dt=\sum_{k=1}^{\infty}\alpha_k^n(\alpha_k^m)^*\frac{\sigma_c^2}{\lambda_k+\sigma_c^2}. \tag{8.137}$$

Combining (8.136) and (8.137), the set of constraints (8.69) is translated to

$$\sum_{k=1}^{\infty}\alpha_k^m(\alpha_k^n)^*=\delta_{m,n} \tag{8.138}$$

for every $m,n=1,\ldots,N$. Consequently, our problem has now been reduced to

$$\max_{\{\alpha_k^n\}}\sum_{n=1}^{N}\sum_{k=1}^{\infty}|\alpha_k^n|^2\frac{\lambda_k^2}{\lambda_k+\sigma_c^2}\quad\text{s.t.}\ \sum_{k=1}^{\infty}\alpha_k^m(\alpha_k^n)^*=\delta_{m,n}. \tag{8.139}$$

We now show that the sequences $\{\alpha_k^n\}$ which solve (8.139) must satisfy $\alpha_k^n=0$ for every $k>N$ and $n=1,\ldots,N$. To see this, assume to the contrary that the nth sequence satisfies

$\alpha_\ell^n \neq 0$ for some $\ell > N$. We can then replace this sequence by a sequence $\{\tilde{\alpha}_k^n\}_{k\in\mathbb{Z}}$ satisfying

$$|\tilde{\alpha}_k^n|^2 = \begin{cases} |\alpha_k^n|^2 + a_k^2 & 1 \leq k \leq N \\ 0 & k = \ell \\ |\alpha_k^n|^2 & N < k \text{ and } k \neq \ell \end{cases} \tag{8.140}$$

where $\sum_{k=1}^N a_k^2 = |\alpha_\ell^n|^2$ (to ensure that $\sum_{k\in\mathbb{Z}} |\tilde{\alpha}_k^n|^2 = 1$). Such a set of coefficients $\{a_k\}_{k=1}^N$ can always be found since the N-term truncation of the remaining $N-1$ sequences cannot span $\mathbb{C}^N$. With this sequence, the nth summand in the objective of (8.139) becomes

$$\begin{aligned} \sum_{k=1}^\infty |\tilde{\alpha}_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} &= \sum_{k=1}^\infty |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} \\ &\quad + \left(\sum_{k=1}^N a_k^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} - |\alpha_\ell^n|^2 \frac{\lambda_\ell^2}{\lambda_\ell + \sigma_c^2} \right) \\ &\geq \sum_{k=1}^\infty |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} + \frac{\lambda_\ell^2}{\lambda_\ell + \sigma_c^2} \left(\sum_{k=1}^N a_k^2 - |\alpha_\ell^n|^2 \right) \\ &= \sum_{k=1}^\infty |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2}, \end{aligned} \tag{8.141}$$

where we used the fact that $\lambda_k \geq \lambda_\ell$ for every $k < \ell$ and that $z^2/(a+z)$ is a monotone increasing function of z for all $z > 0$. This contradicts the optimality of $\{\alpha_k^n\}_{k\in\mathbb{Z}}$. Therefore, the set of sequences maximizing (8.139) satisfy $\alpha_k^n = 0$ for every $k > N$ and $n = 1, \ldots, N$.

It remains to determine the optimal values of the first N elements of each of the N sequences $\{\alpha_k^n\}_{k\in\mathbb{Z}}$, $n = 1, \ldots, N$. For this purpose, let $\boldsymbol{A}$ denote the $N \times N$ matrix whose entries are $\boldsymbol{A}_{k,n} = \alpha_k^n$ and let $\boldsymbol{\Lambda}$ be a diagonal matrix with $\boldsymbol{\Lambda}_{k,k} = \lambda_k^2/(\lambda_k + \sigma_c^2)$. Then, the constraint (8.138) can be written as $\boldsymbol{A}^*\boldsymbol{A} = \boldsymbol{I}_N$, which is equivalent to $\boldsymbol{A}\boldsymbol{A}^* = \boldsymbol{I}_N$. Now, the objective in (8.139) can be expressed as

$$\begin{aligned} \sum_{n=1}^N \sum_{k=1}^\infty |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} &= \mathrm{Tr}\{\boldsymbol{A}^*\boldsymbol{\Lambda}\boldsymbol{A}\} \\ &= \mathrm{Tr}\{\boldsymbol{A}\boldsymbol{A}^*\boldsymbol{\Lambda}\} \\ &= \mathrm{Tr}\{\boldsymbol{\Lambda}\}, \end{aligned} \tag{8.142}$$

which is independent of $\boldsymbol{A}$. Therefore, we conclude that any set of orthonormal sequences $\{\alpha_k^n\}_{k\in\mathbb{Z}}$, $n = 1, \ldots, N$, whose elements vanish for every $k > N$ is optimal.

Part II

Performance Analysis in the Bayesian Setting

Chapter 9

The Optimal-Bias Bound

This chapter is a reprint of the paper:

- Z. Ben-Haim and Y. C. Eldar, "A lower bound on the Bayesian MSE based on the optimal bias function," *IEEE Trans. Inform. Theory*, vol. 55, no. 11, Nov. 2009, pp. 5179–5196.

9.1 Introduction

The goal of estimation theory is to infer the value of an unknown parameter based on observations. A common approach to this problem is the Bayesian framework, in which the estimate is constructed by combining the measurements with prior information about the parameter [3]. In this setting, the parameter $\boldsymbol{\theta}$ is random, and its distribution describes the *a priori* knowledge of the unknown value. In addition, measurements $\boldsymbol{x}$ are obtained, whose conditional distribution, given $\boldsymbol{\theta}$, provides further information about the parameter. The objective is to construct an estimator $\hat{\boldsymbol{\theta}}$, which is a function of the measurements, so that $\hat{\boldsymbol{\theta}}$ is close to $\boldsymbol{\theta}$ in some sense. A common measure of the quality of an estimator is its mean-squared error (MSE), given by $E\{\|\boldsymbol{\theta}-\hat{\boldsymbol{\theta}}\|^2\}$.

It is well-known that the posterior mean $E\{\boldsymbol{\theta}|\boldsymbol{x}\}$ is the technique minimizing the MSE. Thus, from a theoretical perspective, there is no difficulty in finding the minimum MSE (MMSE) estimator in any given problem. In practice, however, the complexity of computing the posterior mean is often prohibitive. As a result, various alternatives, such as the maximum *a posteriori* (MAP) technique, have been developed [16]. The purpose of such methods is to approach the performance of the MMSE estimator with a computationally efficient algorithm.

An important goal is to quantify the performance degradation resulting from the use of

these suboptimal techniques. One way to do this is to compare the MSE of the method used in practice with the MMSE. Unfortunately, computation of the MMSE is itself infeasible in many cases. This has led to a large body of work seeking to find simple lower bounds on the MMSE in various estimation problems [39, 59–61, 65–68, 155, 156].

Generally speaking, previous bounds can be divided into two categories. The Weiss–Weinstein family is based on a covariance inequality and includes the Bayesian Cramér–Rao bound [59], the Bobrovski–Zakai bound [155], and the Weiss–Weinstein bound [60, 61]. The Ziv–Zakai family of bounds is based on comparing the estimation problem to a related detection scenario. This family includes the Ziv–Zakai bound [65] and its improvements, notably the Bellini–Tartara bound [67], the Chazan–Zakai–Ziv bound [66], and the generalization of Bell *et al.* [68]. Recently, Renaux *et al.* have combined both approaches [156].

The accuracy of the bounds described above is usually tested numerically in particular estimation settings. Few of the previous results provide any sort of analytical proof of accuracy, even under asymptotic conditions. Bellini and Tartara [67] briefly discuss performance of their bound at high signal-to-noise ratio (SNR), and Bell *et al.* [68] prove that their bound converges to the true value at low SNR for a particular family of Gaussian-like probability distributions. To the best of our knowledge, there are no other results concerning the asymptotic performance of Bayesian bounds.

A different estimation setting arises when one considers $\boldsymbol{\theta}$ as a *deterministic* unknown parameter. In this case, too, a common goal is to construct an estimator having low MSE. However, the term MSE has a very different meaning in the deterministic setting, since in this case, the expectation is taken only over the random variable $\boldsymbol{x}$. One elementary difference with far-reaching implications is that in the Bayesian case, the MSE is a single real number, whereas the deterministic MSE is a function of the unknown parameter $\boldsymbol{\theta}$ [6, 44, 157].

Many lower bounds have been developed for the deterministic setting, as well. These include classical results such as the Cramér–Rao [45, 48], Hammersley–Chapman–Robbins [50, 51], Bhattacharya [54], and Barankin [53] bounds, as well as more recent results [99, 100, 158–161]. By far the simplest and most commonly used of these approaches is the Cramér–Rao bound (CRB). Like most other deterministic bounds, the CRB deals explicitly with unbiased estimators, or, equivalently, with estimators having a specific, pre-specified bias function. Two exceptions are the uniform CRB [158, 159] and the minimax linear-bias bound [160, 161]. The CRB is known to be *asymptotically* tight in many cases, even though many later bounds are sharper than it [44, 134, 159].

Although the deterministic and Bayesian settings stem from different points of view, there exist insightful relations between the two approaches. The basis for this connection is the fact that by adding a prior distribution for $\boldsymbol{\theta}$, any deterministic problem can be transformed to a corresponding Bayesian setting. Several theorems relate the performance of corresponding Bayesian and deterministic scenarios [6]. As a consequence, numerous bounds have both a deterministic and a Bayesian version [59, 60, 156, 162].

The simplicity and asymptotic tightness of the deterministic CRB motivate its use in problems in which $\boldsymbol{\theta}$ is random. Such an application was described by Young and Westerberg [39], who considered the case of a scalar θ constrained to the interval $[\theta_0, \theta_1]$. They used the prior distribution of θ to determine the optimal bias function for use in the biased CRB, and thus obtained a Bayesian bound. It should be noted that this result differs from the Bayesian CRB of Van Trees [59]; the two bounds are compared in Section 9.2.3. We refer to the result of Young and Westerberg as the optimal-bias bound (OBB), since it is based on choosing the bias function which optimizes the CRB using the given prior distribution.

This paper provides an extension and a deeper analysis of the OBB. Specifically, we generalize the bound to an arbitrary n-dimensional estimation setting [163]. The bound is determined by finding the solution to a certain partial differential equation. Using tools from functional analysis, we demonstrate that a unique solution exists for this differential equation. Under suitable symmetry conditions, it is shown that the method can be reduced to the solution of an ordinary differential equation and, in some cases, presented in closed form.

The mathematical tools employed in this paper are also used for characterizing the performance of the OBB. Specifically, it is demonstrated analytically that the proposed bound is asymptotically tight for both high and low SNR values. Furthermore, the OBB is compared with several other bounds; in the examples considered, the OBB is both simpler computationally and more accurate than all relevant alternatives.

The remainder of this paper is organized as follows. In Section 9.2, we derive the OBB for a vector parameter. Section 9.3 discusses some mathematical concepts required to ensure the existence of the OBB. In Section 9.4, a practical technique for calculating the bound is developed using variational calculus. In Section 9.5, we demonstrate some properties of the OBB, including its asymptotic tightness. Finally, in Section 9.6, we compare the performance of the bound with that of other relevant techniques.

9.2 The Optimal-Bias Bound

In this section, we derive the OBB for the general vector case. To this end, we first examine the relation between the Bayesian and deterministic estimation settings (Section 9.2.1). Next, we focus on the deterministic case and review the basic properties of the CRB (Section 9.2.2). Finally, the OBB is derived from the CRB (Section 9.2.3).

The focus of this paper is the Bayesian estimation problem, but the bound we propose stems from the theory of deterministic estimation. To avoid confusion, we will indicate that a particular quantity refers to the deterministic setting by appending the symbol $;\boldsymbol{\theta}$ to it. For example, the notation $E\{\cdot\}$ denotes expectation over *both* $\boldsymbol{\theta}$ and $\boldsymbol{x}$, i.e., expectation in the Bayesian sense, while expectation solely over $\boldsymbol{x}$ (in the deterministic setting) is denoted by $E\{\cdot;\boldsymbol{\theta}\}$. The notation $E\{\cdot \mid \boldsymbol{\theta}\}$ indicates Bayesian expectation conditioned on $\boldsymbol{\theta}$.

Some further notation used throughout the paper is as follows. Lowercase boldface letters signify vectors and uppercase boldface letters indicate matrices. The ith component of a vector $\boldsymbol{v}$ is denoted v_i, while $\boldsymbol{v}^{(1)}, \boldsymbol{v}^{(2)}, \ldots$ signifies a sequence of vectors. The derivative $\partial f/\partial \boldsymbol{v}$ of a function $f(\boldsymbol{v})$ is a vector function whose ith element is $\partial f/\partial v_i$. Similarly, given a vector function $\boldsymbol{b}(\boldsymbol{\theta})$, the derivative $\partial \boldsymbol{b}/\partial \boldsymbol{\theta}$ is defined as the matrix function whose (i,j)th entry is $\partial b_i/\partial \theta_j$. The squared Euclidean norm $\boldsymbol{v}^T\boldsymbol{v}$ of a vector $\boldsymbol{v}$ is denoted $\|v\|^2$, while the squared Frobenius norm $\mathrm{Tr}(\boldsymbol{M}\boldsymbol{M}^T)$ of a matrix $\boldsymbol{M}$ is denoted $\|\boldsymbol{M}\|_F^2$. In Section 9.3, we will also define some functional norms, which will be of use later in the paper.

9.2.1 The Bayesian–Deterministic Connection

We now review a fundamental relation between the Bayesian and deterministic estimation settings. Let $\boldsymbol{\theta}$ be an unknown random vector in $\mathbb{R}^n$ and let $\boldsymbol{x}$ be a measurement vector. The joint probability density function (pdf) of $\boldsymbol{\theta}$ and $\boldsymbol{x}$ is $p_{\boldsymbol{x},\boldsymbol{\theta}}(\boldsymbol{x},\boldsymbol{\theta}) = p_{\boldsymbol{x}|\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{\theta})p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$, where $p_{\boldsymbol{\theta}}$ is the prior distribution of $\boldsymbol{\theta}$ and $p_{\boldsymbol{x}|\boldsymbol{\theta}}$ is the conditional distribution of $\boldsymbol{x}$ given $\boldsymbol{\theta}$. For later use, define the set Θ of feasible parameter values by

$$\Theta = \{\boldsymbol{\theta} \in \mathbb{R}^n : p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) > 0\}. \tag{9.1}$$

Suppose $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}(\boldsymbol{x})$ is an estimator of $\boldsymbol{\theta}$. Its (Bayesian) MSE is given by

$$\mathrm{MSE} = E\{\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2\} = \int \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2 p_{\boldsymbol{x},\boldsymbol{\theta}}(\boldsymbol{x},\boldsymbol{\theta}) d\boldsymbol{x} d\boldsymbol{\theta}. \tag{9.2}$$

By the law of total expectation, we have

$$\begin{aligned}\mathrm{MSE} &= \int \left(\int \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2 p_{\boldsymbol{x}|\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{\theta}) d\boldsymbol{x} \right) p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= E\Big\{ E\Big\{ \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2 \Big| \boldsymbol{\theta} \Big\} \Big\}. \end{aligned} \tag{9.3}$$

Now consider a deterministic estimation setting, i.e., suppose $\boldsymbol{\theta}$ is a deterministic unknown which is to be estimated from random measurements $\boldsymbol{x}$. Let the distribution $p_{\boldsymbol{x};\boldsymbol{\theta}}$ of $\boldsymbol{x}$ (as a function of $\boldsymbol{\theta}$) be given by $p_{\boldsymbol{x};\boldsymbol{\theta}}(\boldsymbol{x};\boldsymbol{\theta}) = p_{\boldsymbol{x}|\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{\theta})$, i.e., the distribution of $\boldsymbol{x}$ in the deterministic case equals the conditional distribution in the corresponding Bayesian problem.

The estimator $\hat{\boldsymbol{\theta}}$ defined above is simply a function of the measurements, and can therefore be applied in the deterministic case as well. Its deterministic MSE is given by

$$E\big\{ \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2; \boldsymbol{\theta} \big\} = \int \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2 p_{\boldsymbol{x};\boldsymbol{\theta}}(\boldsymbol{x};\boldsymbol{\theta}) d\boldsymbol{x} \tag{9.4}$$

Since $p_{\boldsymbol{x};\boldsymbol{\theta}}(\boldsymbol{x};\boldsymbol{\theta}) = p_{\boldsymbol{x}|\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{\theta})$, we have

$$E\big\{ \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2; \boldsymbol{\theta} \big\} = E\Big\{ \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2 \Big| \boldsymbol{\theta} \Big\}. \tag{9.5}$$

Combining this fact with (9.3), we find that the Bayesian MSE equals the expectation of the MSE of the corresponding deterministic problem, i.e.

$$E\big\{ \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2 \big\} = E\big\{ E\big\{ \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2; \boldsymbol{\theta} \big\} \big\}. \tag{9.6}$$

This relation will be used to construct the OBB in Section 9.2.3.

9.2.2 The Deterministic Cramér–Rao Bound

Before developing the OBB, we review some basic results in the deterministic estimation setting. Suppose $\boldsymbol{\theta}$ is a deterministic parameter vector and let $\boldsymbol{x}$ be a measurement vector having pdf $p_{\boldsymbol{x};\boldsymbol{\theta}}(\boldsymbol{x};\boldsymbol{\theta})$. Denote by $\Theta \subseteq \mathbb{R}^n$ the set of all possible values of $\boldsymbol{\theta}$. We assume for technical reasons that Θ is an open set.[1]

Let $\hat{\boldsymbol{\theta}}$ be an estimator of $\boldsymbol{\theta}$ from the measurements $\boldsymbol{x}$. We require the following regularity conditions to ensure that the CRB holds [69, §3.1.3].

[1]This is required in order to ensure that one can discuss differentiability of $p_{\boldsymbol{x};\boldsymbol{\theta}}$ with respect to $\boldsymbol{\theta}$ at any point $\boldsymbol{\theta} \in \Theta$. In the Bayesian setting to which we will return in Section 9.2.3, Θ is defined by (9.1); in this case, adding a boundary to Θ essentially leaves the setting unchanged, as long as the prior probability for $\boldsymbol{\theta}$ to be on the boundary of Θ is zero. Therefore, this requirement is of little practical relevance.

1. $p_{\boldsymbol{x};\boldsymbol{\theta}}(\boldsymbol{x};\boldsymbol{\theta})$ is continuously differentiable with respect to $\boldsymbol{\theta}$. This condition is required to ensure the existence of the Fisher information.

2. The Fisher information matrix $\boldsymbol{J}(\boldsymbol{\theta})$, defined by

$$[\boldsymbol{J}(\boldsymbol{\theta})]_{ij} = E\left\{ \frac{\partial \log p_{\boldsymbol{x};\boldsymbol{\theta}}}{\partial \theta_i} \frac{\partial \log p_{\boldsymbol{x};\boldsymbol{\theta}}}{\partial \theta_j} ; \boldsymbol{\theta} \right\} \tag{9.7}$$

   is bounded and positive definite for all $\boldsymbol{\theta} \in \Theta$. This ensures that the measurements contain data about the unknown parameter.

3. Exchanging the integral and derivative in the equation

$$\int t(\boldsymbol{x}) \frac{\partial}{\partial \theta_i} p_{\boldsymbol{x};\theta}(\boldsymbol{x};\theta) d\boldsymbol{x} = \frac{\partial}{\partial \theta_i} \int t(\boldsymbol{x}) p_{\boldsymbol{x};\theta}(\boldsymbol{x};\theta) d\boldsymbol{x} \tag{9.8}$$

   is justified for any measurable function $t(\boldsymbol{x})$, in the sense that, if one side exists, then the other exists and the two sides are equal. A sufficient condition for this to hold is that the support of $p_{\boldsymbol{x};\boldsymbol{\theta}}$ does not depend on $\boldsymbol{\theta}$.

4. All estimators $\hat{\boldsymbol{\theta}}$ are Borel measurable functions which satisfy

$$\left\| \frac{\partial p_{\boldsymbol{x};\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \hat{\boldsymbol{\theta}}^T \right\|_F \leq g(\boldsymbol{x}) \text{ for all } \boldsymbol{\theta} \tag{9.9}$$

   for some integrable function $g(\boldsymbol{x})$. This technical requirement is needed in order to exclude certain pathological estimators whose statistical behavior is insufficiently smooth to allow the application of the CRB.

The bias of an estimator $\hat{\boldsymbol{\theta}}$ is defined as

$$\boldsymbol{b}(\boldsymbol{\theta}) = E\{\hat{\boldsymbol{\theta}};\boldsymbol{\theta}\} - \boldsymbol{\theta}. \tag{9.10}$$

Under the above assumptions, it can be shown that the bias of any estimator is continuously differentiable [39, Lemma 2]. Furthermore, under these assumptions, the CRB holds, and thus, for any estimator having bias $\boldsymbol{b}(\boldsymbol{\theta})$, we have

$$\begin{aligned} E\{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2;\boldsymbol{\theta}\} &\geq \text{CRB}[\boldsymbol{b},\boldsymbol{\theta}] \\ &\triangleq \text{Tr}\left[\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}^{-1}(\boldsymbol{\theta}) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right)^T \right] + \|\boldsymbol{b}(\boldsymbol{\theta})\|^2. \end{aligned} \tag{9.11}$$

A more common form of the CRB is obtained by restricting attention to unbiased estimators (i.e., techniques for which $\boldsymbol{b}(\boldsymbol{\theta}) = \boldsymbol{0}$). Under the unbiasedness assumption, the bound simplifies to $\text{MSE} \geq \text{Tr}(\boldsymbol{J}^{-1}(\boldsymbol{\theta}))$. However, in the sequel we will make use of the general form (9.11).

9.2.3 A Bayesian Bound from the CRB

The OBB of Young and Westerberg [39] is based on applying the Bayesian–deterministic connection described in Section 9.2.1 to the deterministic CRB (9.11). Specifically, returning now to the Bayesian setting, one can combine (9.6) and (9.11) to obtain that, for any estimator $\hat{\boldsymbol{\theta}}$ with bias function $\boldsymbol{b}(\boldsymbol{\theta})$,

$$E\{\|\boldsymbol{\theta}-\hat{\boldsymbol{\theta}}\|^2\} \geq Z[\boldsymbol{b}] \triangleq \int_\Theta \mathrm{CRB}[\boldsymbol{b},\boldsymbol{\theta}]\, p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \tag{9.12}$$

where the expectation is now performed over both $\boldsymbol{\theta}$ and $\boldsymbol{x}$. Note that (9.12) describes the Bayesian MSE as a function of a deterministic property (the bias) of $\hat{\boldsymbol{\theta}}$. Since any estimator has *some* bias function, and since all bias functions are continuously differentiable in our setting, minimizing $Z[\boldsymbol{b}]$ over all continuously differentiable functions $\boldsymbol{b}$ yields a lower bound on the MSE of any Bayesian estimator. Thus, under the regularity conditions of Section 9.2.2, a lower bound on the Bayesian MSE is given by

$$s = \inf_{\boldsymbol{b}\in C^1} \int_\Theta \left[\|\boldsymbol{b}(\boldsymbol{\theta})\|^2 + \mathrm{Tr}\left(\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}^{-1}(\boldsymbol{\theta}) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right)^T \right) \right] p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \tag{9.13}$$

where C^1 is the space of continuously differentiable functions $\boldsymbol{f}:\Theta\to\mathbb{R}^n$.

Note that the OBB differs from the Bayesian CRB of Van Trees [59]. Van Trees' result is based on applying the Cauchy–Schwarz inequality to the joint pdf $p_{\boldsymbol{x},\boldsymbol{\theta}}$, whereas the deterministic CRB is based on applying a similar procedure to $p_{\boldsymbol{x};\boldsymbol{\theta}}$. As a consequence, the regularity conditions required for the Bayesian CRB are stricter, requiring that $p_{\boldsymbol{x},\boldsymbol{\theta}}$ be twice differentiable with respect to $\boldsymbol{\theta}$. By contrast, the OBB requires differentiability only of the conditional pdf $p_{\boldsymbol{x}|\boldsymbol{\theta}}$. An example in which this difference is important is the case in which the prior distribution $p_{\boldsymbol{\theta}}$ is discontinuous, e.g., when $p_{\boldsymbol{\theta}}$ is uniform. The performance of the OBB in this setting will be examined in Section 9.6.

In the next section, we will see that it is advantageous to perform the minimization (9.13) over a somewhat modified class of functions. This will allow us to prove the unique existence of a solution to the optimization problem, a result which will be of use when examining the properties of the bound later in the paper.

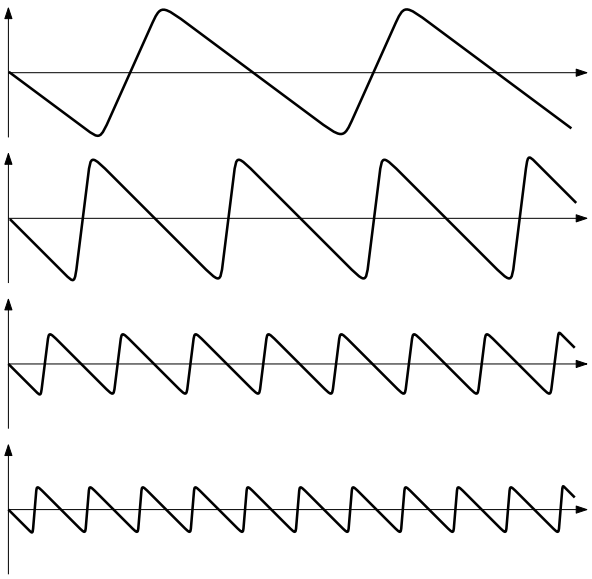

Figure 9.1: A sequence of continuous functions for which both $|b(\theta)|^2$ and $|1 + b'(\theta)|^2$ tend to zero for almost every value of θ.

9.3 Mathematical Safeguards

In the previous section, we saw that a lower bound on the MMSE can be obtained by solving the minimization problem (9.13). However, at this point, we have no guarantee that the solution s of (9.13) is anywhere near the true value of the MMSE. Indeed, at first sight, it may appear that $s = 0$ for any estimation setting. To see this, note that $Z[\boldsymbol{b}]$ is a sum of two components, a bias gradient part and a squared bias part. Both parts are nonnegative, but the former is zero when the bias gradient is $-\boldsymbol{I}$, while the latter is zero when the bias is zero. No differentiable function $\boldsymbol{b}$ satisfies these two constraints simultaneously for all $\boldsymbol{\theta}$, since if the squared bias is everywhere zero, then the bias gradient is also zero. However, it is possible to construct a sequence of functions $\boldsymbol{b}^{(i)}$ for which both the bias gradient and the squared bias norm tend to zero for *almost* every value of $\boldsymbol{\theta}$. An example of such a sequence in a one-dimensional setting is plotted in Fig. 9.1. Here, a sequence $\boldsymbol{b}^{(i)}$ of smooth, periodic functions is presented. The function period tends to zero, and the percentage of the cycle in which the derivative equals -1 increases as i increases. Thus, the pointwise limit of the function sequence is zero almost everywhere, and the pointwise limit of the derivative is -1 almost everywhere.

In the specific case shown in Fig. 9.1, it can be shown that the value of $Z[\boldsymbol{b}^{(i)}]$ does not tend to zero; in fact, $Z[\boldsymbol{b}^{(i)}]$ tends to infinity in this situation. However, our example illustrates that care must be taken when applying concepts from finite-dimensional optimization problems to

variational calculus.

The purpose of this section is to show that $s > 0$, so that the bound is meaningful, for any problem setting satisfying the regularity conditions of Section 9.2.2. (This question was not addressed by Young and Westerberg [39].) While doing so, we develop some abstract concepts which will also be used when analyzing the asymptotic properties of the OBB in Section 9.5.

As often happens with variational problems, it turns out that the minimum of (9.13) is not necessarily achieved by any continuously differentiable function. In order to guarantee an achievable minimum, one must instead minimize (9.13) over a slightly modified space, which is defined below. As explained in Section 9.2.2, all bias functions are continuously differentiable, so that the minimizing function ultimately obtained, if it is not differentiable, will not be the bias of any estimator. However, as we will see, the minimum value of our new optimization problem is identical to the infimum of (9.13). Furthermore, this approach allows us to demonstrate several important theoretical properties of the OBB.

Let L^2 be the space of $p_{\boldsymbol{\theta}}$-measurable functions $\boldsymbol{b} : \Theta \to \mathbb{R}^n$ such that

$$\int_{\Theta} \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) < \infty. \tag{9.14}$$

Define the associated inner product

$$\left\langle \boldsymbol{b}^{(1)}, \boldsymbol{b}^{(2)} \right\rangle_{L^2} \triangleq \sum_{i=1}^{n} \int_{\Theta} b_i^{(1)}(\boldsymbol{\theta}) b_i^{(2)}(\boldsymbol{\theta}) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \tag{9.15}$$

and the corresponding norm $\|\boldsymbol{b}\|_{L^2}^2 \triangleq \langle \boldsymbol{b}, \boldsymbol{b} \rangle_{L^2}$. Any function $\boldsymbol{b} \in L^2$ has a derivative in the distributional sense, but this derivative might not be a function. For example, discontinuous functions have distributional derivatives which contain a Dirac delta. If, for every i, the distributional derivative $\partial b_i / \partial \boldsymbol{\theta}$ of $\boldsymbol{b}$ is a function in L^2, then $\boldsymbol{b}$ is said to be weakly differentiable [164], and its weak derivative is the matrix function $\partial \boldsymbol{b} / \partial \boldsymbol{\theta}$. Roughly speaking, a function is weakly differentiable if it is continuous and its derivative exists almost everywhere.

The space of all weakly differentiable functions in L^2 is called the first-order Sobolev space [164], and is denoted H^1. Define an inner product on H^1 as

$$\left\langle \boldsymbol{b}^{(1)}, \boldsymbol{b}^{(2)} \right\rangle_{H^1} \triangleq \left\langle \boldsymbol{b}^{(1)}, \boldsymbol{b}^{(2)} \right\rangle_{L^2} + \sum_{j=1}^{n} \left\langle \frac{\partial b_j^{(1)}}{\partial \boldsymbol{\theta}}, \frac{\partial b_j^{(2)}}{\partial \boldsymbol{\theta}} \right\rangle_{L^2}. \tag{9.16}$$

The associated norm is $\|\boldsymbol{b}\|_{H^1}^2 \triangleq \langle \boldsymbol{b}, \boldsymbol{b} \rangle_{H^1}$. An important property which will be used extensively in our analysis is that H^1 is a Hilbert space.

Note that since Θ is an open set, not all functions in C^1 are in H^1. For example, in the case $\Theta = \mathbb{R}^n$, the function $\boldsymbol{b}(\boldsymbol{\theta}) = \boldsymbol{k}$, for some nonzero constant $\boldsymbol{k}$, is continuously differentiable

but not integrable. Thus $\boldsymbol{b}$ is in C^1 but not in H^1, nor even in L^2. However, any measurable function which is not in H^1 has $\|\boldsymbol{b}\|_{H^1} = \infty$, meaning that either $\boldsymbol{b}$ or $\partial\boldsymbol{b}/\partial\boldsymbol{\theta}$ has infinite L^2 norm. Consequently, either the bias norm part or the bias gradient part of $Z[\boldsymbol{b}]$ is infinite. It follows that performing the minimization (9.13) over $C^1 \cap H^1$, rather than over C^1, does not change the minimum value. On the other hand, $C^1 \cap H^1$ is dense in H^1, and $Z[\boldsymbol{b}]$ is continuous, so that minimizing (9.13) over H^1 rather than $C^1 \cap H^1$ also does not alter the minimum. Consequently, we will henceforth consider the problem

$$s = \inf_{\boldsymbol{b} \in H^1} Z[\boldsymbol{b}]. \tag{9.17}$$

The advantage of including weakly differentiable functions in the minimization is that a unique minimizer can now be guaranteed, as demonstrated by the following result.

Proposition 9.1. *Consider the problem*

$$\bar{\boldsymbol{b}} = \arg\min_{\boldsymbol{b} \in H^1} Z[\boldsymbol{b}] \tag{9.18}$$

where $Z[\boldsymbol{b}]$ is given by (9.12) *and $\boldsymbol{J}(\boldsymbol{\theta})$ is positive definite and bounded with probability 1. This problem is well-defined, i.e., there exists a unique $\bar{\boldsymbol{b}} \in H^1$ which minimizes $Z[\boldsymbol{b}]$. Furthermore, the minimum value $s = Z[\bar{\boldsymbol{b}}]$ is finite and nonzero.*

Proving the unique existence of a minimizer for (9.17) is a technical exercise in functional analysis which can be found in Appendix 9.B. However, once the existence of such a minimizer is demonstrated, it is not difficult to see that $0 < s < \infty$. To see that $s < \infty$, we must find a function $\boldsymbol{b}$ for which $Z[\boldsymbol{b}] < \infty$. One such function is $\boldsymbol{b} = \boldsymbol{0}$, for which $Z[\boldsymbol{b}]$ is finite since $\boldsymbol{J}(\boldsymbol{\theta})$ is bounded. Now suppose by contradiction that $s = 0$, which implies that there exists a function $\bar{\boldsymbol{b}} \in H^1$ such that $Z[\bar{\boldsymbol{b}}] = 0$. Therefore, both the bias gradient and the squared bias parts of $Z[\bar{\boldsymbol{b}}]$ are zero. In particular, since the squared bias part equals zero, we have $\|\bar{\boldsymbol{b}}\|_{L^2} = 0$. Hence, $\bar{\boldsymbol{b}} = \boldsymbol{0}$, because L^2 is a normed space. But then, by the definition (9.12) of $Z[\cdot]$,

$$Z[\bar{\boldsymbol{b}}] = \int_\Theta \mathrm{Tr}(\boldsymbol{J}^{-1}(\boldsymbol{\theta})) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \tag{9.19}$$

which is positive; this is a contradiction.

Note that functions in H^1 are defined up to changes on a set having zero measure. In particular, the fact that $\boldsymbol{b}^{(0)}$ is unique does not preclude functions which are identical to $\boldsymbol{b}^{(0)}$ almost everywhere (which obviously have the same value $Z[\boldsymbol{b}]$).

Summarizing the discussion of the last two sections, we have the following theorem.

Theorem 9.2. *Let $\boldsymbol{\theta}$ be an unknown random vector with pdf $p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) > 0$ over the open set $\Theta \subseteq \mathbb{R}^n$, and let $\boldsymbol{x}$ be a measurement vector whose pdf, conditioned on $\boldsymbol{\theta}$, is given by $p_{\boldsymbol{x}|\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{\theta})$. Assume the regularity conditions of Section 9.2.2 hold. Then, for any estimator $\hat{\boldsymbol{\theta}}$,*

$$E\{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2\} \geq \min_{\boldsymbol{b} \in H^1} \int_{\Theta} \mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}] p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) d\boldsymbol{\theta}. \tag{9.20}$$

The minimum in (9.20) *is nonzero and finite. Furthermore, this minimum is achieved by a function $\bar{\boldsymbol{b}} \in H^1$, which is unique up to changes having zero probability.*

Two remarks are in order concerning Theorem 9.2. First, the function $\boldsymbol{b}$ solving (9.20) might not be the bias of any estimator; indeed, under our assumptions, all bias functions are continuously differentiable, whereas $\boldsymbol{b}$ need only be weakly differentiable. Nevertheless, (9.20) is still a lower bound on the MMSE. Another important observation is that Theorem 9.2 arises from the deterministic CRB; hence, there are no requirements on the prior distribution $p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$. In particular, $p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ can be discontinuous or have bounded support. By contrast, many previous Bayesian bounds do not apply in such circumstances.

9.4 Calculating the Bound

In finite-dimensional convex optimization problems, the requirement of a vanishing first derivative results in a set of equations, whose solution is the global minimum. Analogously, in the case of convex functional optimization problems such as (9.20), the optimum is given by the solution of a set of differential equations. The following theorem, whose proof can be found in Appendix 9.C, specifies the differential equation relevant to our optimization problem.

In this section and in the remainder of the paper, we will consider the case in which the set $\Theta = \{\boldsymbol{\theta} : p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) > 0\}$ is bounded. From a practical point of view, even when Θ consists of the entire set $\mathbb{R}^n$, it can be approximated by a bounded set containing only those values of $\boldsymbol{\theta}$ for which $p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) > \epsilon$.

Theorem 9.3. *Under the conditions of Theorem 9.2, suppose Θ is a bounded subset of $\mathbb{R}^n$ with a smooth boundary Λ. Then, the optimal $\boldsymbol{b}(\boldsymbol{\theta})$ of* (9.20) *is given by the solution to the system of partial differential equations*

$$\begin{aligned} p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) b_i(\boldsymbol{\theta}) &= p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \sum_{j,k} \frac{\partial^2 b_i}{\partial \theta_j \partial \theta_k} (\boldsymbol{J}^{-1})_{jk} \\ &+ \sum_{j,k} \left(\delta_{ik} + \frac{\partial b_i}{\partial \theta_k} \right) \left((\boldsymbol{J}^{-1})_{jk} \frac{\partial p_{\boldsymbol{\theta}}}{\partial \theta_j} + p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \frac{\partial (\boldsymbol{J}^{-1})_{jk}}{\partial \theta_j} \right) \end{aligned} \tag{9.21}$$

for $i = 1, \dots n$, within the range $\boldsymbol{\theta} \in \Theta$, which satisfies the Neumann boundary condition

$$\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}^{-1} \boldsymbol{\nu}(\boldsymbol{\theta}) = \boldsymbol{0} \tag{9.22}$$

for all points $\boldsymbol{\theta} \in \Lambda$. Here, $\boldsymbol{\nu}(\boldsymbol{\theta})$ is a normal to the boundary at $\boldsymbol{\theta}$. All derivatives in this system of equations are to be interpreted in the weak sense.

Note that Theorem 9.2 guarantees the existence of a unique solution in H^1 to the differential equation (9.21) with the boundary conditions (9.22).

The bound of Young and Westerberg [39] is a special case of Theorem 9.3, and is given here for completeness.

Corollary 9.4. *Under the settings of Theorem 9.2, suppose $\Theta = (\theta_0, \theta_1)$ is a bounded interval in $\mathbb{R}$. Then, the bias function $b(\theta)$ minimizing* (9.20) *is a solution to the second-order ordinary differential equation*

$$J(\theta) b(\theta) = b''(\theta) + (1 + b'(\theta)) \left(\frac{d \log p_{\boldsymbol{\theta}}}{d\theta} - \frac{d \log J}{d\theta} \right) \tag{9.23}$$

within the range $\theta \in \Theta$, subject to the boundary conditions $b'(\theta_0) = b'(\theta_1) = -1$.

Theorem 9.3 can be solved numerically, thus obtaining a bound for any problem satisfying the regularity conditions. However, directly solving (9.21) becomes increasingly complex as the dimension of the problem increases. Instead, in many cases, symmetry relations in the problem can be used to simplify the solution. As an example, the following spherically symmetric case can be reduced to a one-dimensional setting equivalent to that of Corollary 9.4. The proof of this theorem can be found in Appendix 9.D.

Theorem 9.5. *Under the setting of Theorem 9.2, suppose that $\Theta = \{\boldsymbol{\theta} : \|\boldsymbol{\theta}\| < r\}$ is a sphere centered on the origin, $p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) = q(\|\boldsymbol{\theta}\|)$ is spherically symmetric, and $\boldsymbol{J}(\boldsymbol{\theta}) = J(\|\boldsymbol{\theta}\|)\boldsymbol{I}$, where $J : \mathbb{R} \to \mathbb{R}$ is a scalar function. Then, the optimal-bias bound* (9.20) *is given by*

$$\begin{aligned} E\{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2\} \geq \frac{2\pi^{n/2}}{\Gamma(n/2)} \int_0^r \Bigg[& b^2(\rho) + \frac{(1 + b'(\rho))^2}{J(\rho)} \\ & + \frac{n-1}{J(\rho)} \left(1 + \frac{b(\rho)}{\rho} \right)^2 \Bigg] q(\rho) \rho^{n-1} d\rho. \end{aligned} \tag{9.24}$$

Here, $\Gamma(\cdot)$ is the Gamma function, and $b(\rho)$ is a solution to the ODE

$$\begin{aligned} J(\theta) b(\theta) = b''(\theta) & + (n-1) \left(\frac{b'(\theta)}{\theta} - \frac{b(\theta)}{\theta^2} \right) \\ & + (1 + b'(\theta)) \left(\frac{d \log q}{d\theta} - \frac{d \log J}{d\theta} \right) \end{aligned} \tag{9.25}$$

subject to the boundary conditions $b(0) = 0$, $b'(r) = -1$. *The bias function for which the bound is achieved is given by*

$$\boldsymbol{b}(\boldsymbol{\theta}) = b(\|\boldsymbol{\theta}\|)\frac{\boldsymbol{\theta}}{\|\boldsymbol{\theta}\|}. \tag{9.26}$$

In this theorem, the requirement $\boldsymbol{J}(\boldsymbol{\theta}) = J(\|\boldsymbol{\theta}\|)\boldsymbol{I}$ indicates that the Fisher information matrix is diagonal and that its components are spherically symmetric. Parameters having a diagonal matrix $\boldsymbol{J}$ are sometimes referred to as *orthogonal*. The simplest case of orthogonality occurs when, to each parameter θ_i, there corresponds a measurement x_i, in such a way that the random variables $x_i|\boldsymbol{\theta}$ are independent. Other orthogonal scenarios can often be constructed by an appropriate parametrization [165].

The requirement that $\boldsymbol{J}$ have spherically symmetric components occurs, for example, in location problems, i.e., situations in which the measurements have the form $\boldsymbol{x} = \boldsymbol{\theta} + \boldsymbol{w}$, where $\boldsymbol{w}$ is additive noise which is independent of $\boldsymbol{\theta}$. Indeed, under such conditions, $\boldsymbol{J}$ is constant in $\boldsymbol{\theta}$ [69, §3.1.3]. If, in addition, the noise components are independent, then this setting also satisfies the orthogonality requirement, and thus application of Theorem 9.5 is appropriate. Note that this estimation problem is not separable, since the components of $\boldsymbol{\theta}$ are correlated; thus, the MMSE in this situation is lower than the sum of the components' MMSE. An example of such a setting is presented in Section 9.6.

9.5 Properties

In this section, we examine several properties of the OBB. We first demonstrate that the optimal bias function has zero mean, a property which also characterizes the bias function of the MMSE estimator. Next, we prove that, under very general conditions, the resulting bound is tight at both low and high SNR values. This is an important result, since a desirable property of a Bayesian bound is that it provides an accurate estimate of the ambiguity region between high and low SNR [68]. Reliable estimation at the two extremes increases the likelihood that the transition between these two regimes will be correctly identified.

9.5.1 Optimal Bias Has Zero Mean

In any Bayesian estimation problem, the bias of the MMSE estimator $\hat{\boldsymbol{\theta}}_{\text{opt}} = E\{\boldsymbol{\theta}|\boldsymbol{x}\}$ has zero mean:

$$E\{\hat{\boldsymbol{\theta}}_{\text{opt}}\} = E\{E\{\boldsymbol{\theta}|\boldsymbol{x}\}\} = E\{\boldsymbol{\theta}\} \tag{9.27}$$

so that

$$E\{\boldsymbol{b}(\hat{\boldsymbol{\theta}}_{\text{opt}})\} = E\{E\{\boldsymbol{\theta}|\boldsymbol{x}\} - \boldsymbol{\theta}\} = \mathbf{0}. \tag{9.28}$$

Thus, it is interesting to ask whether the optimal bias which minimizes (9.20) also has zero mean. This is indeed the case, as shown by the following theorem.

Theorem 9.6. *Let* $\boldsymbol{b}(\boldsymbol{\theta})$ *be the solution to* (9.20)*. Then,*

$$E\{\boldsymbol{b}(\boldsymbol{\theta})\} = \mathbf{0}. \tag{9.29}$$

Proof. Assume by contradiction that $\boldsymbol{b}(\boldsymbol{\theta})$ has nonzero mean $E\{\boldsymbol{b}(\boldsymbol{\theta})\} = \boldsymbol{\mu} \neq \mathbf{0}$. Define $\boldsymbol{b}_0(\boldsymbol{\theta}) \triangleq \boldsymbol{b}(\boldsymbol{\theta}) - \boldsymbol{\mu}$. From (9.11), we then have

$$\begin{aligned} \text{CRB}[\boldsymbol{b}_0, \boldsymbol{\theta}] - \text{CRB}[\boldsymbol{b}, \boldsymbol{\theta}] &= \|\boldsymbol{b}_0(\boldsymbol{\theta})\|^2 - \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 \\ &= \|\boldsymbol{\mu}\|^2 - 2\boldsymbol{\mu}^T\boldsymbol{b}(\boldsymbol{\theta}). \end{aligned} \tag{9.30}$$

Using the functional $Z[\cdot]$ defined in (9.12), we obtain

$$\begin{aligned} Z[\boldsymbol{b}_0] - Z[\boldsymbol{b}] &= E\Big\{\|\boldsymbol{\mu}\|^2 - 2\boldsymbol{\mu}^T\boldsymbol{b}(\boldsymbol{\theta})\Big\} \\ &= \|\boldsymbol{\mu}\|^2 - 2\boldsymbol{\mu}^T E\{\boldsymbol{b}(\boldsymbol{\theta})\} \\ &= -\|\boldsymbol{\mu}\|^2 < 0. \end{aligned} \tag{9.31}$$

Thus $Z[\boldsymbol{b}_0] < Z[\boldsymbol{b}]$, contradicting the fact that $\boldsymbol{b}(\boldsymbol{\theta})$ minimizes (9.20). □

9.5.2 Tightness at Low SNR

Bell *et al.* [68] examined the performance of the extended Ziv–Zakai bound at low SNR and demonstrated that, for a particular family of distributions, the extended Ziv–Zakai bound achieves the MSE of the optimal estimator as the SNR tends to 0. We now examine the low-SNR performance of the OBB, and demonstrate tightness for a much wider range of problem settings.

Bell *et al.* did not define the general meaning of a low SNR value, and only stated that "[a]s observation time and/or SNR become very small, the observations become useless ... [and] the minimum MSE estimator converges to the *a priori* mean." This statement clearly does not apply to all estimation problems, since it is not always clear what parameter corresponds to the observation time or the SNR. We propose to define the zero SNR case more generally as any situation in which $\boldsymbol{J}(\boldsymbol{\theta}) = \mathbf{0}$ with probability 1. This definition implies that the measurements do not contain information about the unknown parameter, which is the usual informal meaning

of zero SNR. In the case $\boldsymbol{J}(\boldsymbol{\theta}) = \mathbf{0}$, it can be shown that the MMSE estimator is the prior mean, so that our definition implies the statement of Bell *et al.*

The OBB is inapplicable when $\boldsymbol{J}(\boldsymbol{\theta}) = \mathbf{0}$, since the CRB is based on the assumption that $\boldsymbol{J}(\boldsymbol{\theta})$ is positive definite. To avoid this singularity, we consider a sequence of estimation settings which converge to zero SNR. More specifically, we require all eigenvalues of $\boldsymbol{J}(\boldsymbol{\theta})$ to decrease monotonically to zero for $p_{\boldsymbol{\theta}}$-almost all $\boldsymbol{\theta}$. The following theorem, the proof of which can be found in Appendix 9.E, demonstrates the tightness of the OBB in this low-SNR setting.

Theorem 9.7. *Let $\boldsymbol{\theta}$ be a random vector whose pdf $p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ is nonzero over an open set $\Theta \subseteq \mathbb{R}^n$. Let $\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}, \ldots$ be a sequence of observation vectors having finite Fisher information matrices $\boldsymbol{J}^{(1)}(\boldsymbol{\theta}), \boldsymbol{J}^{(2)}(\boldsymbol{\theta}), \ldots$, respectively. Suppose that, for all N, the matrix $\boldsymbol{J}^{(N)}(\boldsymbol{\theta})$ is positive definite for $p_{\boldsymbol{\theta}}$-almost all $\boldsymbol{\theta}$, and that all eigenvalues of $\boldsymbol{J}^{(N)}(\boldsymbol{\theta})$ decrease monotonically to zero as $N \to \infty$ for $p_{\boldsymbol{\theta}}$-almost all $\boldsymbol{\theta}$. Let β_N denote the optimal-bias bound for estimating $\boldsymbol{\theta}$ from $\boldsymbol{x}^{(N)}$. Then,*

$$\lim_{N\to\infty} \beta_N = E\Big\{ \|\boldsymbol{\theta} - E\{\boldsymbol{\theta}\}\|^2 \Big\}. \tag{9.32}$$

9.5.3 Tightness at High SNR

We now examine the performance of the OBB for high SNR values. To formally define the high SNR regime, we consider a sequence of measurements $\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}, \ldots$ of a single parameter vector $\boldsymbol{\theta}$. It is assumed that, when conditioned on $\boldsymbol{\theta}$, all measurements $\boldsymbol{x}^{(i)}$ are identically and independently distributed (IID). Furthermore, we assume that the Fisher information matrix of a single observation $\boldsymbol{J}(\boldsymbol{\theta})$ is well-defined, positive definite and finite for $p_{\boldsymbol{\theta}}$-almost all $\boldsymbol{\theta}$. We consider the problem of estimating $\boldsymbol{\theta}$ from the set of measurements $\{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(N)}\}$, for a given value of N. The high SNR regime is obtained when N is large.

When N tends to infinity, the MSE of the optimal estimator tends to zero. An important question, however, concerns the rate of convergence of the minimum MSE. More precisely, given the optimal estimator $\hat{\boldsymbol{\theta}}^{(N)}$ of $\boldsymbol{\theta}$ from $\{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(N)}\}$, one would like to determine the asymptotic distribution of $\sqrt{N}(\hat{\boldsymbol{\theta}}^{(N)} - \boldsymbol{\theta})$, conditioned on $\boldsymbol{\theta}$. A fundamental result of asymptotic estimation theory can be loosely stated as follows [134, §III.3], [6, §6.8]. Under some fairly mild regularity conditions, the asymptotic distribution of $\sqrt{N}(\hat{\boldsymbol{\theta}}^{(N)} - \boldsymbol{\theta})$, conditioned on $\boldsymbol{\theta}$, does not depend on the prior distribution $p_{\boldsymbol{\theta}}$; rather, $\sqrt{N}(\hat{\boldsymbol{\theta}}^{(N)} - \boldsymbol{\theta}) \,|\, \boldsymbol{\theta}$ converges in distribution to a Gaussian random vector with mean zero and covariance $\boldsymbol{J}^{-1}(\boldsymbol{\theta})$. It follows that

$$\lim_{N\to\infty} NE\Big\{ \|\hat{\boldsymbol{\theta}}^{(N)} - \boldsymbol{\theta}\|^2 \Big\} = E\Big\{ \mathrm{Tr}[\boldsymbol{J}^{-1}(\boldsymbol{\theta})] \Big\}. \tag{9.33}$$

Since the minimum MSE tends to zero at high SNR, any lower bound on the minimum MSE must also tend to zero as $N \to \infty$. However, one would further expect a good lower bound to follow the behavior of (9.33). In other words, if β_N represents the lower bound for estimating $\boldsymbol{\theta}$ from $\{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(N)}\}$, a desirable property is $N\beta_N \to E\{\mathrm{Tr}[\boldsymbol{J}^{-1}(\boldsymbol{\theta})]\}$. The following theorem, whose proof is found in Appendix 9.E, demonstrates that this is indeed the case for the OBB.

Except for a very brief treatment by Bellini and Tartara [67], no previous Bayesian bound has shown such a result. Although it appears that the Ziv–Zakai and Weiss–Weinstein bounds may also satisfy this property, this has not been proven formally. It is also known that the Bayesian CRB is *not* asymptotically tight in this sense [166, Eqs. (37)–(39)].

Theorem 9.8. *Let $\boldsymbol{\theta}$ be a random vector whose pdf $p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ is nonzero over an open set $\Theta \subseteq \mathbb{R}^n$. Let $\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}, \ldots$ be a sequence of measurement vectors, such that $\boldsymbol{x}^{(1)}|\boldsymbol{\theta}, \boldsymbol{x}^{(2)}|\boldsymbol{\theta}, \ldots$ are IID. Let $\boldsymbol{J}(\boldsymbol{\theta})$ be the Fisher information matrix for estimating $\boldsymbol{\theta}$ from $\boldsymbol{x}^{(1)}$, and suppose $\boldsymbol{J}(\boldsymbol{\theta})$ is finite and positive definite for $p_{\boldsymbol{\theta}}$-almost all $\boldsymbol{\theta}$. Let β_N be the optimal-bias bound* (9.20) *for estimating $\boldsymbol{\theta}$ from the observation sequence $\{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(N)}\}$. Then,*

$$\lim_{N\to\infty} N\beta_N = E\Big\{\mathrm{Tr}\big(\boldsymbol{J}^{-1}(\boldsymbol{\theta})\big)\Big\}. \tag{9.34}$$

Note that for Theorem 9.8 to hold, we require only that $\boldsymbol{J}(\boldsymbol{\theta})$ be finite and positive definite. By contrast, the various theorems guaranteeing asymptotic efficiency of Bayesian estimators all require substantially stronger regularity conditions [134, §III.3], [6, §6.8]. One reason for this is that asymptotic efficiency describes the behavior of $\hat{\boldsymbol{\theta}}$ conditioned on each possible value of $\boldsymbol{\theta}$, and is thus a stronger result than the asymptotic Bayesian MSE of (9.33).

9.6 Example: Uniform Prior

The original bound of Young and Westerberg [39] predates most Bayesian bounds, and, surprisingly, it has never been cited by or compared with later results. In this section, we measure the performance of the original bound and of its extension to the vector case against that of various other techniques. We consider the case in which $\boldsymbol{\theta}$ is uniformly distributed over an n-dimensional open ball $\Theta = \{\boldsymbol{\theta} : \|\boldsymbol{\theta}\| < r\} \subseteq \mathbb{R}^n$, so that

$$p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) = \frac{1}{V_n(r)} \mathbb{1}_{\Theta} \tag{9.35}$$

where $\mathbb{1}_S$ equals 1 when $\boldsymbol{\theta} \in S$ and 0 otherwise, and

$$V_n(r) = \frac{\pi^{n/2} r^{n-1}}{\Gamma(1 + n/2)} \tag{9.36}$$

is the volume of an n-ball of radius r [167]. We further assume that

$$\boldsymbol{x} = \boldsymbol{\theta} + \boldsymbol{w} \tag{9.37}$$

where $\boldsymbol{w}$ is zero-mean Gaussian noise, independent of $\boldsymbol{\theta}$, having covariance $\sigma^2\boldsymbol{I}$. We are interested in lower bounds on the MSE achievable by an estimator of $\boldsymbol{\theta}$ from $\boldsymbol{x}$.

We begin by developing the OBB for this setting, as well as some alternative bounds. We then compare the different approaches in a one-dimensional and a three-dimensional setting.

The Fisher information matrix for the given estimation problem is given by $\boldsymbol{J}(\boldsymbol{\theta}) = \sigma^{-2}\boldsymbol{I}$, so that the conditions of Theorem 9.5 hold. It follows that the optimal bias function is given by $\boldsymbol{b}(\boldsymbol{\theta}) = b(\|\boldsymbol{\theta}\|)\boldsymbol{\theta}/\|\boldsymbol{\theta}\|$, where $b(\cdot)$ is a solution to the differential equation

$$\frac{b}{\sigma^2} = b'' + (n-1)\left(\frac{b'}{\theta} - \frac{b}{\theta^2}\right) \tag{9.38}$$

with boundary conditions $b(0) = 0$, $b'(r) = -1$. The general solution to this differential equation is given by

$$b(\theta) = C_1\theta^{1-n/2}I_{n/2}\left(\frac{\theta}{\sigma}\right) + C_2\theta^{1-n/2}K_{n/2}\left(\frac{\theta}{\sigma}\right) \tag{9.39}$$

where $I_\alpha(z)$ and $K_\alpha(z)$ are the modified Bessel functions of the first and second types, respectively [131]. Since $K_\alpha(z)$ is singular at the origin, the requirement $b(0) = 0$ leads to $C_2 = 0$. Differentiating (9.39) with respect to θ, we obtain

$$b'(\theta) = C_1\theta^{-n/2}\left(I_{n/2}\left(\frac{\theta}{\sigma}\right) + \frac{\theta}{\sigma}I_{1+n/2}\left(\frac{\theta}{\sigma}\right)\right) \tag{9.40}$$

so that the requirement $b'(r) = -1$ leads to

$$C_1 = -\frac{r^{n/2}}{I_{n/2}(r/\sigma) + r/\sigma I_{1+n/2}(r/\sigma)}. \tag{9.41}$$

Substituting this value of $b(\cdot)$ into (9.24) yields the OBB, which can be computed by evaluating a single one-dimensional integral. Alternatively, in the one-dimensional case, the integral can be computed analytically, as will be shown below.

Despite the widespread use of finite-support prior distributions [60, 65], the regularity conditions of many bounds are violated by such prior pdf functions. Indeed, the Bayesian CRB of Van Trees [59], the Bobrovski–Zakai bound [155], and the Bayesian Abel bound [156] all assume that $p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ has infinite support, and thus cannot be applied in this scenario.

Techniques from the Ziv–Zakai family are applicable to constrained problems. An extension of the Ziv–Zakai bound for vector parameter estimation was developed by Bell *et al.* [68]. From

[68, Property 4], the MSE of the ith component of $\boldsymbol{\theta}$ is bounded by

$$E\{(\theta_i - \hat{\theta}_i)^2\} \geq \int_0^\infty V\left\{\max_{\boldsymbol{\delta}:\boldsymbol{e}_i^T\boldsymbol{\delta}=h} A(\boldsymbol{\delta})P_{\min}(\boldsymbol{\delta})\right\} h\,dh \tag{9.42}$$

where $\boldsymbol{e}_i$ is a unit vector in the direction of the ith component, $V\{\cdot\}$ is the valley-filling function defined by

$$V\{f(h)\} = \max_{\eta\geq 0} f(h+\eta), \tag{9.43}$$

$$A(\boldsymbol{\delta}) \triangleq \int_{\mathbb{R}^n} \min\left(p_{\boldsymbol{\theta}}(\boldsymbol{\theta}), p_{\boldsymbol{\theta}}(\boldsymbol{\theta}+\boldsymbol{\delta})\right) d\boldsymbol{\theta}, \tag{9.44}$$

and $P_{\min}(\boldsymbol{\delta})$ is the minimum probability of error for the problem of testing hypothesis $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ vs. $H_1 : \boldsymbol{\theta} = \boldsymbol{\theta}_0 + \boldsymbol{\delta}$. In the current setting, $P_{\min}(\boldsymbol{\delta})$ is given by $P_{\min}(\boldsymbol{\delta}) = Q(\|\boldsymbol{\delta}\|/2\sigma)$, where $Q(z) = (2\pi)^{-1/2}\int_z^\infty e^{-t^2/2}dt$ is the tail function of the normal distribution. Also, we have

$$A(\boldsymbol{\delta}) = \frac{V_n^C(r, \|\boldsymbol{\delta}\|)}{V_n(r)} \tag{9.45}$$

where

$$V_n^C(r,h) = \int_{\mathbb{R}^n} \mathbb{1}_\Theta \mathbb{1}_{\Theta+h\boldsymbol{e}_1} d\boldsymbol{\theta} \tag{9.46}$$

and $\Theta + h\boldsymbol{e}_1 = \{\boldsymbol{\theta} + h\boldsymbol{e}_1 : \boldsymbol{\theta} \in \Theta\}$. Thus, $V_n^C(r,h)$ is the volume of the intersection of two n-balls whose centers are at a distance of h units from one another. Substituting these results into (9.42), we have

$$\begin{aligned} &E\{(\theta_i - \hat{\theta}_i)^2\} \\ &\quad\geq \int_0^\infty V\left\{\max_{\boldsymbol{\delta}:\boldsymbol{e}_i^T\boldsymbol{\delta}=h} \frac{V_n^C(r,\|\boldsymbol{\delta}\|)}{V_n(r)} Q\left(\frac{\|\boldsymbol{\delta}\|}{2\sigma}\right)\right\} h\,dh. \end{aligned} \tag{9.47}$$

Note that both $V_n^C(r,\|\boldsymbol{\delta}\|)$ and $Q(\|\boldsymbol{\delta}\|/2\sigma)$ decrease with $\|\boldsymbol{\delta}\|$. Therefore, the maximum in (9.47) is obtained for $\boldsymbol{\delta} = h\boldsymbol{e}_i$. Also, since the argument of $V\{\cdot\}$ is monotonically decreasing, the valley-filling function has no effect and can be removed. Finally, since $V_n^C(r,h) = 0$ for $h > 2r$, the integration can be limited to the range $[0, 2r]$. Thus, the extended Ziv–Zakai bound is given by

$$E\{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2\} \geq \int_0^{2r} n\frac{V_n^C(r,h)}{V_n(r)} Q\left(\frac{h}{2\sigma}\right) h\,dh. \tag{9.48}$$

We now compute the Weiss–Weinstein bound for the setting at hand. This bound is given by

$$E\{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2\} \geq \mathrm{Tr}(\boldsymbol{H}\boldsymbol{G}^{-1}\boldsymbol{H}^T) \tag{9.49}$$

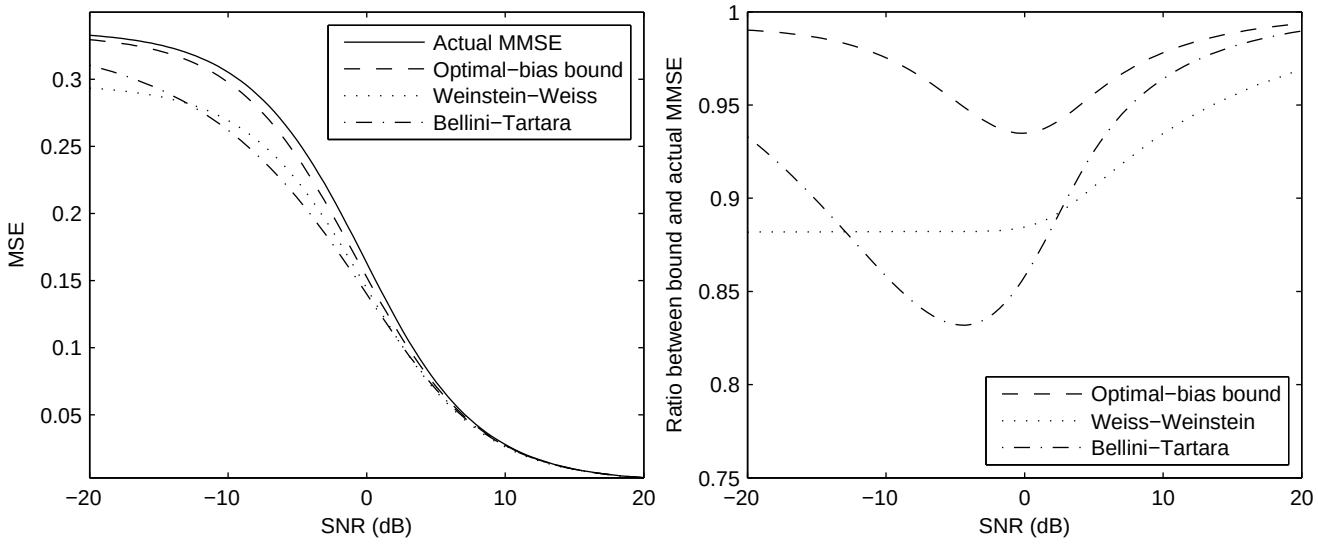

(a) (b)

Figure 9.2: Comparison of the MSE bounds and the minimum achievable MSE in a one-dimensional setting for which $\theta \sim U[-r, r]$ and $x|\theta \sim N(\theta, \sigma^2)$.

where $\boldsymbol{H} = [\boldsymbol{h}_1, \ldots, \boldsymbol{h}_m]$ is a matrix containing an arbitrary number m of test vectors and $\boldsymbol{G}$ is a matrix whose elements are given by

$$G_{ij} = \frac{E\{r(\boldsymbol{x}, \boldsymbol{\theta}; \boldsymbol{h}_i, s_i) r(\boldsymbol{x}, \boldsymbol{\theta}; \boldsymbol{h}_j, s_j)\}}{E\{L^{s_i}(\boldsymbol{x}; \boldsymbol{\theta} + \boldsymbol{h}_i, \boldsymbol{\theta})\} E\{L^{s_j}(\boldsymbol{x}; \boldsymbol{\theta} + \boldsymbol{h}_j, \boldsymbol{\theta})\}} \tag{9.50}$$

in which

$$r(\boldsymbol{x}, \boldsymbol{\theta}; \boldsymbol{h}_i, s_i) \triangleq L^{s_i}(\boldsymbol{x}; \boldsymbol{\theta} + \boldsymbol{h}_i, \boldsymbol{\theta}) - L^{1-s_i}(\boldsymbol{x}; \boldsymbol{\theta} - \boldsymbol{h}_i, \boldsymbol{\theta}) \tag{9.51}$$

and

$$L(\boldsymbol{x}; \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \triangleq \frac{p_{\boldsymbol{\theta}}(\boldsymbol{\theta}_1) p_{\boldsymbol{x}|\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{\theta}_1)}{p_{\boldsymbol{\theta}}(\boldsymbol{\theta}_2) p_{\boldsymbol{x}|\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{\theta}_2)}. \tag{9.52}$$

The vectors $\boldsymbol{h}_1, \ldots, \boldsymbol{h}_m$ and the scalars $s_1, \ldots, s_m$ are arbitrary, and can be optimized to maximize the bound (9.49). To avoid a multidimensional nonconvex optimization problem, we restrict attention to $m = n$, $\boldsymbol{h}_i = h\boldsymbol{e}_i$, and $s_i = 1/2$, as suggested by [60]. This results in a dependency on a single scalar parameter h.

Under these conditions, G_{ij} can be written as

$$G_{ij} = \frac{1}{M(\boldsymbol{h}_i) M(\boldsymbol{h}_j)} \big[\tilde{M}(\boldsymbol{h}_i - \boldsymbol{h}_j, -\boldsymbol{h}_j) + \tilde{M}(\boldsymbol{h}_i - \boldsymbol{h}_j, \boldsymbol{h}_i) - \tilde{M}(\boldsymbol{h}_i + \boldsymbol{h}_j, \boldsymbol{h}_j) - \tilde{M}(\boldsymbol{h}_i + \boldsymbol{h}_j, \boldsymbol{h}_i)\big] \tag{9.53}$$

where

$$M(\boldsymbol{h}) \triangleq E\Big\{L^{1/2}(\boldsymbol{x}; \boldsymbol{\theta} + \boldsymbol{h}, \boldsymbol{\theta})\Big\} \tag{9.54}$$

and

$$\tilde{M}(\boldsymbol{h}_1, \boldsymbol{h}_2) \triangleq E\Big\{ L^{1/2}(\boldsymbol{x}; \boldsymbol{\theta} + \boldsymbol{h}_1, \boldsymbol{\theta}) \mathbb{1}_{\Theta + \boldsymbol{h}_2} \Big\}. \tag{9.55}$$

Note that we have used the corrected version of the Weiss–Weinstein bound [168]. Substituting the probability distribution of $\boldsymbol{x}$ and $\boldsymbol{\theta}$ into the definitions of $M(\boldsymbol{h})$ and $\tilde{M}(\boldsymbol{h}_1, \boldsymbol{h}_2)$, we have

$$\begin{aligned} M(\boldsymbol{h}) &= E\Big\{ e^{-\|\boldsymbol{\theta}+\boldsymbol{h}-\boldsymbol{x}\|^2/4\sigma^2} e^{\|\boldsymbol{\theta}-\boldsymbol{x}\|^2/4\sigma^2} \mathbb{1}_{\Theta+\boldsymbol{h}} \Big\} \\ &= \frac{V_n^C(r, \|\boldsymbol{h}\|)}{V_n(r)} e^{-\|\boldsymbol{h}\|^2/8\sigma^2} \end{aligned} \tag{9.56}$$

and, similarly,

$$\tilde{M}(\boldsymbol{h}_1, \boldsymbol{h}_2) = \frac{e^{-\|\boldsymbol{h}_1\|^2/8\sigma^2}}{V_n(r)} \int \mathbb{1}_{\Theta} \mathbb{1}_{\Theta+\boldsymbol{h}_1} \mathbb{1}_{\Theta+\boldsymbol{h}_2} d\boldsymbol{\theta}. \tag{9.57}$$

Thus, $M(\boldsymbol{h})$ is a function only of $\|\boldsymbol{h}\|$, and $\tilde{M}(\boldsymbol{h}_1, \boldsymbol{h}_2)$ is a function only of $\|\boldsymbol{h}_1\|$, $\|\boldsymbol{h}_2\|$, and $\|\boldsymbol{h}_1 - \boldsymbol{h}_2\|$. Since $\boldsymbol{h}_i = h\boldsymbol{e}_i$, it follows that, for $i \neq j$, the numerator of (9.53) vanishes. Thus, $\boldsymbol{G}$ is a diagonal matrix, whose diagonal elements equal

$$G_{ii} = 2\frac{\tilde{M}(0, h\boldsymbol{e}_1) - \tilde{M}(2h\boldsymbol{e}_1, h\boldsymbol{e}_1)}{M^2(h\boldsymbol{e}_1)}. \tag{9.58}$$

The Weiss–Weinstein bound is given by substituting this result into (9.49) and maximizing over h, i.e.,

$$E\big\{ \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2 \big\} \geq \max_{h \in [0, 2r]} \frac{nh^2 M^2(h\boldsymbol{e}_1)}{2[\tilde{M}(0, h\boldsymbol{e}_1) - \tilde{M}(2h\boldsymbol{e}_1, h\boldsymbol{e}_1)]}. \tag{9.59}$$

The value of h yielding the tightest bound can be determined by performing a grid search.

To compare the OBB with the alternative approaches developed above, we first consider the one-dimensional case in which θ is uniformly distributed in the range $\Theta = (-r, r)$. Let $x = \theta + w$ be a single noisy observation, where w is zero-mean Gaussian noise, independent of θ, with variance σ^2. We wish to bound the MSE of an estimator of θ from x.

The optimal bias function is given by (9.39). Using the fact that $I_{1/2}(t) = \sqrt{2/\pi} \sinh(t)/\sqrt{t}$, we obtain

$$b(\theta) = -\sigma \frac{\sinh(\theta/\sigma)}{\cosh(r/\sigma)} \tag{9.60}$$

which also follows [39] from Corollary 9.4. Substituting this expression into (9.20), we have that, for any estimator $\hat{\theta}$,

$$E\big\{ (\theta - \hat{\theta})^2 \big\} \geq \sigma^2 \left(1 - \frac{\tanh(r/\sigma)}{r/\sigma} \right). \tag{9.61}$$

Apart from the reduction in computational complexity, the simplicity of (9.61) also emphasizes several features of the estimation problem. First, the dependence of the problem on the

dimensionless quantity r/σ, rather than on r and σ separately, is clear. This is to be expected, as a change in units of measurement would multiply both r and σ by a constant. Second, the asymptotic properties demonstrated in Theorems 9.7 and 9.8 can be easily verified. For $r \gg \sigma$, the bound converges to the noise variance σ^2, corresponding to an uninformative prior whose optimal estimator is $\hat{\theta} = x$; whereas, for $\sigma \gg r$, a Taylor expansion of $\tanh(z)/z$ immediately shows that the bound converges to $r^2/3$, corresponding to the case of uninformative measurements, where the optimal estimator is $\hat{\theta} = 0$. Thus, the bound (9.61) is tight both for very low and for very high SNR, as expected.

In the one-dimensional case, we have $V_1(r) = 2r$ and $V_1^C(r,h) = \max(2r-h, 0)$, so that the extended Ziv–Zakai bound (9.48) and the Weiss–Weinstein bound (9.59) can also be simplified somewhat. In particular, the extended Ziv–Zakai bound (9.48) can be written as

$$E\{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2\} \geq \int_0^{2r} \left(1 - \frac{h}{2r}\right) hQ\left(\frac{h}{2\sigma}\right) dh. \tag{9.62}$$

Using integration by parts, (9.62) becomes

$$E\{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2\} \geq \frac{2r^2}{3} Q\left(\frac{r}{\sigma}\right) + \sigma^2 \left[\Gamma_{3/2}\left(\frac{r^2}{2\sigma^2}\right) - \frac{8}{3\sqrt{2\pi}} \frac{\sigma}{r} \Gamma_2\left(\frac{r^2}{2\sigma^2}\right)\right] \tag{9.63}$$

where $\Gamma_a(z) = (1/\Gamma(a)) \int_0^z e^{-t} t^{a-1} dt$ is the incomplete Gamma function. Like the expression (9.61) for the OBB, this bound can be shown to converge to the noise variance σ^2 when $r \gg \sigma$ and to the prior variance $r^2/3$ when $\sigma \gg r$. However, while the convergence of the OBB to these asymptotic values has been demonstrated in general in Theorems 9.7 and 9.8, the asymptotic tightness of the Ziv–Zakai bound in the general case remains an open question.

The Weiss–Weinstein bound (9.59) can likewise be simplified further in the one-dimensional case, yielding

$$E\{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2\} \geq \max_{h \in [0,2r]} \frac{h^2 e^{-h^2/4\sigma^2} \left(1 - \frac{h}{2r}\right)^2}{2\left(1 - \frac{h}{2r} - \max\left(0, 1 - \frac{h}{r}\right) e^{-h^2/2\sigma^2}\right)}. \tag{9.64}$$

However, calculating this bound still requires a numerical search for the optimal value of h.

These bounds are compared with the exact value of the MMSE in Fig. 9.2. In this figure, the SNR is defined as

$$\mathrm{SNR(dB)} = 10 \log_{10}\left(\frac{\mathrm{Var}(\theta)}{\mathrm{Var}(w)}\right) = 10 \log_{10}\left(\frac{r^2}{3\sigma^2}\right). \tag{9.65}$$

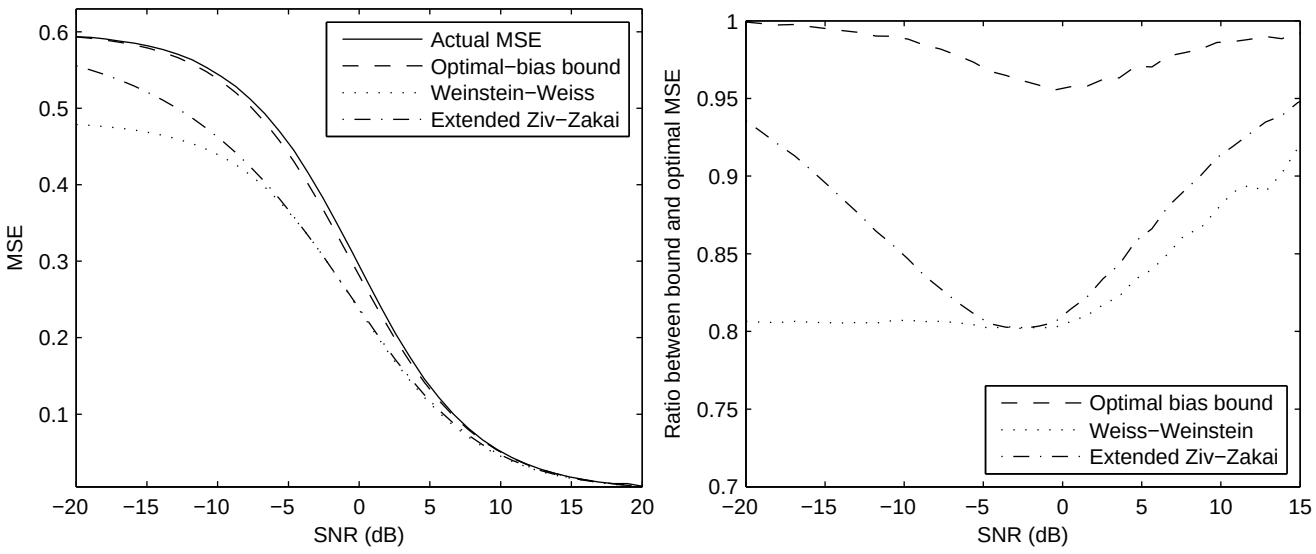

(a) (b)

Figure 9.3: Comparison of the MSE bounds and the minimum achievable MSE in a three-dimensional setting for which $\boldsymbol{\theta}$ is uniformly distributed over a ball of radius r and $\boldsymbol{x}|\boldsymbol{\theta} \sim N(\boldsymbol{\theta}, \sigma^2 \boldsymbol{I})$.

The MMSE was computed by Monte Carlo approximation of the error of the optimal estimator $E\{\theta|x\}$, which was itself computed by numerical integration. Fig. 9.2(a) plots the MMSE and the values obtained by the aforementioned bounds, while Fig. 9.2(b) plots the ratio between each of the bounds and the actual MMSE in order to emphasize the difference in accuracy between the various bounds. As can be seen from this figure, the OBB is closer to the true MSE than all other bounds, for all tested SNR values.

The improvements provided by the OBB continue to hold in higher dimensions as well, although in this case it is not possible to provide a closed form for any of the bounds. For example, Fig. 9.3 compares the aforementioned bounds with the true MMSE in the three-dimensional case. In this case the SNR is given by

$$\mathrm{SNR}(\mathrm{dB}) = 10\log_{10}\left(\frac{\mathrm{Var}(\boldsymbol{\theta})}{\mathrm{Var}(\boldsymbol{w})}\right) = 10\log_{10}\left(\frac{r^2}{5\sigma^2}\right). \tag{9.66}$$

Here, computation of the minimum MSE requires multi-dimensional numerical integration, and is by far more computationally complex than the calculation of the bounds. Again, it is evident from this figure that the OBB is a very tight bound in all ranges of operation, and is considerably closer to the true value than either of the alternative approaches.

9.7 Conclusion

Although often considered distinct settings, there are insightful connections between the Bayesian and deterministic estimation problems. One such relation is the use of the deterministic CRB in a Bayesian problem. The application of this deterministic bound to the problem of estimating the minimum Bayesian MSE results in a Bayesian bound which is provably tight at both high and low SNR values. Numerical simulation of the location estimation problem demonstrates that the technique is both simpler and tighter than alternative approaches.

Acknowledgement

The authors are grateful to Dr. Volker Pohl for fruitful discussions concerning many of the mathematical aspects of the paper. The authors would also like to thank the anonymous reviewers for their many constructive comments.

9.A Some Technical Lemmas

The proof of several theorems in the paper relies on the following technical results.

Lemma 9.9. *Consider the minimization problems*

$$M_\ell = \inf_{\boldsymbol{b}\in S} Z_\ell[\boldsymbol{b}], \quad \ell = 1,2,3 \tag{9.67}$$

where $\boldsymbol{J}(\boldsymbol{\theta})$ *is positive definite and bounded a.e.* $(p_{\boldsymbol{\theta}})$*,*

$$\begin{aligned} Z_1[\boldsymbol{b}] &\triangleq \int_\Theta \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ Z_2[\boldsymbol{b}] &\triangleq \int_\Theta \mathrm{Tr}\left(\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right)\boldsymbol{J}^{-1}(\boldsymbol{\theta})\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right)^T\right) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ Z_3[\boldsymbol{b}] &\triangleq Z_1[\boldsymbol{b}] + Z_2[\boldsymbol{b}] \end{aligned} \tag{9.68}$$

and $S \subset H^1$ *is convex, closed, and bounded under the* H^1 *norm* (9.16)*. Then, for each* ℓ*, there exists a function* $\boldsymbol{b}^{(0)} \in S$ *such that* $Z[\boldsymbol{b}^{(0)}] = M_\ell$*. If* $\ell = 1$ *or* $\ell = 3$*, then the minimizer of* (9.67) *is unique.*

Note that $Z_3[\boldsymbol{b}]$ equals $Z[\boldsymbol{b}]$ of (9.12); the notation $Z_3[\boldsymbol{b}]$ is introduced for simplicity. Also note that under mild regularity assumptions on $\boldsymbol{J}(\boldsymbol{\theta})$, uniqueness can be demonstrated for $\ell = 2$ as well, but this is not necessary for our purposes.

Proof. The space H^1 is a Cartesian product of n Sobolev spaces $H^1(\Theta)$, each of which is a separable Hilbert space [169, §3.7.1]. Therefore, H^1 is also a separable Hilbert space. It follows from the Banach–Alaoglu theorem [170, §3.17] that all bounded sequences in H^1 have weakly convergent subsequences [164, §2.18]. Recall that a sequence $\boldsymbol{f}^{(1)}, \boldsymbol{f}^{(2)}, \ldots \in H^1$ is said to converge weakly to $\boldsymbol{f}^{(0)} \in H^1$ (denoted $\boldsymbol{f}^{(i)} \rightharpoonup \boldsymbol{f}^{(0)}$) if

$$L[\boldsymbol{f}^{(j)}] \to L[\boldsymbol{f}^{(0)}] \tag{9.69}$$

for all continuous linear functionals $L[\cdot]$ [164, §2.9].

Given a particular value $\ell \in \{1,2,3\}$, let $\boldsymbol{b}^{(i)}$ be a sequence of functions in S such that $Z_\ell[\boldsymbol{b}^{(i)}] \to M_\ell$. This is a bounded sequence since S is bounded, and therefore there exists a subsequence $\boldsymbol{b}^{(i_k)}$ which converges weakly to some $\boldsymbol{b}^{(\ell)}_{\text{opt}} \in H^1$. Furthermore, since S is closed,[2] we have $\boldsymbol{b}^{(\ell)}_{\text{opt}} \in S$. We will now show that $Z_\ell[\boldsymbol{b}^{(\ell)}_{\text{opt}}] = M_\ell$.

To this end, it suffices to show that $Z_\ell[\cdot]$ is weakly lower semicontinuous, i.e., for any sequence $\boldsymbol{f}^{(i)} \in H^1$ which converges weakly to $\boldsymbol{f}^{(0)} \in H^1$, we must show that

$$Z_\ell[\boldsymbol{f}^{(0)}] \leq \liminf_{i\to\infty} Z_\ell[\boldsymbol{f}^{(i)}]. \tag{9.70}$$

Consider a weakly convergent sequence $\boldsymbol{f}^{(j)} \rightharpoonup \boldsymbol{f}^{(0)}$. Then, (9.69) holds for any continuous linear functional $L[\cdot]$. Specifically, choose the continuous linear functional

$$L_1[\boldsymbol{f}] = \int_\Theta \boldsymbol{f}^{(0)}(\boldsymbol{\theta}) \boldsymbol{f}(\boldsymbol{\theta}) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}). \tag{9.71}$$

We then have

$$\begin{aligned}
Z_1[\boldsymbol{f}^{(0)}] &= L_1[\boldsymbol{f}^{(0)}] \\
&= \lim_{j\to\infty} L_1[\boldsymbol{f}^{(j)}] \\
&= \lim_{j\to\infty} \int_\Theta \sum_{i=1}^n f_i^{(0)}(\boldsymbol{\theta}) f_i^{(j)}(\boldsymbol{\theta}) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\
&\leq \liminf_{j\to\infty} \sqrt{\int_\Theta \|\boldsymbol{f}^{(0)}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \cdot \int_\Theta \|\boldsymbol{f}^{(j)}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta})} \\
&= \sqrt{Z_1[\boldsymbol{f}^{(0)}]} \liminf_{j\to\infty} \sqrt{Z_1[\boldsymbol{f}^{(j)}]}
\end{aligned} \tag{9.72}$$

where we have used the Cauchy–Schwarz inequality. It follows that

$$\sqrt{Z_1[\boldsymbol{f}^{(0)}]} \leq \liminf_{j\to\infty} \sqrt{Z_1[\boldsymbol{f}^{(j)}]} \tag{9.73}$$

[2] In fact, we require that S be "weakly closed" in the sense that weakly convergent sequences in S converge to an element in S. However, since S is convex, this notion is equivalent to the ordinary definition of closure [170, §3.13].

and therefore $Z_1[\boldsymbol{f}^{(0)}] \le \liminf_{j\to\infty} Z_1[\boldsymbol{f}^{(j)}]$, so that $Z_1[\cdot]$ is weakly lower semicontinuous.

Similarly, consider the continuous linear functional

$$L_2[\boldsymbol{f}] = \int_\Theta \operatorname{Tr}\left(\left(\boldsymbol{I} + \frac{\partial \boldsymbol{f}^{(0)}}{\partial \boldsymbol{\theta}}\right)\boldsymbol{J}^{-1}(\boldsymbol{\theta})\left(\boldsymbol{I} + \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{\theta}}\right)^T\right) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \tag{9.74}$$

for which we have

$$\begin{aligned} Z_2[\boldsymbol{f}^{(0)}] &= L_2[\boldsymbol{f}^{(0)}] \\ &= \lim_{j\to\infty} L_2[\boldsymbol{f}^{(j)}] \\ &= \lim_{j\to\infty} \int_\Theta \operatorname{Tr}\left[\left(\boldsymbol{I} + \frac{\partial \boldsymbol{f}^{(0)}}{\partial \boldsymbol{\theta}}\right)\boldsymbol{J}^{-1}(\boldsymbol{\theta}) \right. \\ &\qquad \left. \cdot \left(\boldsymbol{I} + \frac{\partial \boldsymbol{f}^{(j)}}{\partial \boldsymbol{\theta}}\right)^T\right] p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}). \end{aligned} \tag{9.75}$$

Note that, for any positive definite matrix $\boldsymbol{W}$, $\operatorname{Tr}(\boldsymbol{AWB}^T)$ is an inner product of the two matrices $\boldsymbol{A}$ and $\boldsymbol{B}$. Therefore, by the Cauchy–Schwarz inequality,

$$\operatorname{Tr}(\boldsymbol{AWB}^T) \le \sqrt{\operatorname{Tr}(\boldsymbol{AWA}^T)\operatorname{Tr}(\boldsymbol{BWB}^T)}. \tag{9.76}$$

Applying this to (9.75), we have

$$\begin{aligned} Z_2[\boldsymbol{f}^{(0)}] \le \liminf_{j\to\infty} & \\ \int_\Theta & \sqrt{\operatorname{Tr}\left(\left(\boldsymbol{I} + \frac{\partial \boldsymbol{f}^{(0)}}{\partial \boldsymbol{\theta}}\right)\boldsymbol{J}^{-1}(\boldsymbol{\theta})\left(\boldsymbol{I} + \frac{\partial \boldsymbol{f}^{(0)}}{\partial \boldsymbol{\theta}}\right)^T\right)} \\ & \cdot \sqrt{\operatorname{Tr}\left(\left(\boldsymbol{I} + \frac{\partial \boldsymbol{f}^{(j)}}{\partial \boldsymbol{\theta}}\right)\boldsymbol{J}^{-1}(\boldsymbol{\theta})\left(\boldsymbol{I} + \frac{\partial \boldsymbol{f}^{(j)}}{\partial \boldsymbol{\theta}}\right)^T\right)} p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}). \end{aligned} \tag{9.77}$$

Once again using the Cauchy–Schwarz inequality results in

$$Z_2[\boldsymbol{f}^{(0)}] \le \liminf_{j\to\infty} \sqrt{Z_2[\boldsymbol{f}^{(0)}] Z_2[\boldsymbol{f}^{(j)}]} \tag{9.78}$$

and therefore $Z_2[\boldsymbol{f}^{(0)}] \le \liminf_{j\to\infty} Z_2[\boldsymbol{f}^{(j)}]$, so that $Z_2[\cdot]$ is weakly lower semicontinuous. Since $Z_3[\boldsymbol{f}] = Z_1[\boldsymbol{f}] + Z_2[\boldsymbol{f}]$, it follows that $Z_3[\cdot]$ is also weakly lower semicontinuous.

Now recall that $\boldsymbol{b}^{(i_k)} \rightharpoonup \boldsymbol{b}_{\text{opt}}^{(\ell)}$ and $Z_\ell[\boldsymbol{b}^{(i_k)}] \to M_\ell$. By the definition (9.70) of lower semicontinuity, it follows that

$$Z_\ell[\boldsymbol{b}_{\text{opt}}^{(\ell)}] \le \liminf_{k\to\infty} Z_\ell[\boldsymbol{b}^{(i_k)}] = M_\ell \tag{9.79}$$

and since M_ℓ is the infimum of $Z_\ell[\boldsymbol{b}]$, we obtain $Z[\boldsymbol{b}_{\text{opt}}^{(\ell)}] = M$. Thus $\boldsymbol{b}_{\text{opt}}^{(\ell)}$ is a minimizer of (9.67).

It remains to show that for $\ell \in \{1, 3\}$, the minimizer of (9.67) is unique. To this end, we first show that $Z_1[\cdot]$ is strictly convex. Let $b^{(0)}, b^{(1)} \in S$ be two essentially different functions, i.e.,

$$p_{\boldsymbol{\theta}}\Big(\Big\{\boldsymbol{\theta} \in \Theta : \boldsymbol{b}^{(0)}(\boldsymbol{\theta}) \neq \boldsymbol{b}^{(1)}(\boldsymbol{\theta})\Big\}\Big) > 0. \tag{9.80}$$

Let $\boldsymbol{b}^{(2)}(\boldsymbol{\theta}) = \lambda \boldsymbol{b}^{(0)}(\boldsymbol{\theta}) + (1-\lambda)\boldsymbol{b}^{(1)}(\boldsymbol{\theta})$ for some $0 < \lambda < 1$, so that $\boldsymbol{b}^{(2)} \in S$ by convexity. We then have

$$\begin{aligned}
Z_1[\boldsymbol{b}^{(2)}] &= \int_Q \left\|\lambda \boldsymbol{b}^{(0)}(\boldsymbol{\theta}) + (1-\lambda)\boldsymbol{b}^{(1)}(\boldsymbol{\theta})\right\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\
&\quad + \int_{\Theta\setminus Q} \left\|\lambda \boldsymbol{b}^{(0)}(\boldsymbol{\theta}) + (1-\lambda)\boldsymbol{b}^{(1)}(\boldsymbol{\theta})\right\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\
&< \int_Q \left[\lambda \|\boldsymbol{b}^{(0)}(\boldsymbol{\theta})\|^2 + (1-\lambda)\|\boldsymbol{b}^{(1)}(\boldsymbol{\theta})\|^2\right] p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \\
&\quad + \int_{\Theta\setminus Q} \left[\lambda \|\boldsymbol{b}^{(0)}(\boldsymbol{\theta})\|^2 + (1-\lambda)\|\boldsymbol{b}^{(1)}(\boldsymbol{\theta})\|^2\right] p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \\
&= \lambda Z_1[\boldsymbol{b}^{(0)}] + (1-\lambda) Z_2[\boldsymbol{b}^{(1)}]
\end{aligned} \tag{9.81}$$

where the inequality follows from strict convexity of the squared Euclidean norm $\|\boldsymbol{x}\|^2$. Thus $Z_1[\cdot]$ is strictly convex, and hence has a unique minimum.

Note that $Z_3[\boldsymbol{b}] = Z_1[\boldsymbol{b}] + Z_2[\boldsymbol{b}]$. Since $Z_1[\cdot]$ is strictly convex and $Z_2[\cdot]$ is convex, it follows that $Z_3[\cdot]$ is strictly convex, and thus also has a unique minimum. This completes the proof. □

The following lemma can be thought of as a triangle inequality for a normed space of matrix functions over Θ.

Lemma 9.10. *Let $p_{\boldsymbol{\theta}}$ be a probability measure over Θ, and let $\boldsymbol{M} : \Theta \to \mathbb{R}^{n\times n}$ be a matrix function. Suppose*

$$\int_\Theta \|\boldsymbol{I} + \boldsymbol{M}(\boldsymbol{\theta})\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq \alpha \tag{9.82}$$

for some constant α. It follows that

$$\int_\Theta \|\boldsymbol{M}(\boldsymbol{\theta})\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq (\sqrt{\alpha} + \sqrt{n})^2. \tag{9.83}$$

Proof. By the triangle inequality,

$$\|\boldsymbol{M}(\boldsymbol{\theta})\|_F = \|\boldsymbol{M}(\boldsymbol{\theta}) + \boldsymbol{I} - \boldsymbol{I}\|_F \leq \|\boldsymbol{M}(\boldsymbol{\theta}) + \boldsymbol{I}\|_F + \|\boldsymbol{I}\|_F. \tag{9.84}$$

Since $\|\boldsymbol{I}\|_F^2 = n$, we have

$$\begin{aligned}
&\int_\Theta \|\boldsymbol{M}(\boldsymbol{\theta})\|_F^2 \, p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\
&\leq \int_\Theta \left[\|\boldsymbol{I} + \boldsymbol{M}(\boldsymbol{\theta})\|_F^2 + n + 2\sqrt{n}\,\|\boldsymbol{I} + \boldsymbol{M}(\boldsymbol{\theta})\|_F\right] p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}).
\end{aligned} \tag{9.85}$$

Using the fact that

$$\int_\Theta \|\boldsymbol{I} + \boldsymbol{M}(\boldsymbol{\theta})\|_F \, p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq \sqrt{\int_\Theta \|\boldsymbol{I} + \boldsymbol{M}(\boldsymbol{\theta})\|_F^2 \, p_{\boldsymbol{\theta}}(d\boldsymbol{\theta})} \tag{9.86}$$

and combining with (9.82), it follows that

$$\int_\Theta \|\boldsymbol{M}(\boldsymbol{\theta})\|_F^2 \, p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq \alpha + n + 2\sqrt{n\alpha} \tag{9.87}$$

which completes the proof. □

9.B Proof of Proposition 9.1

The following proof of Proposition 9.1 makes use of the results developed in Appendix 9.A.

Proof of Proposition 9.1. Recall that $Z_3[\boldsymbol{b}]$ of (9.68) equals $Z[\boldsymbol{b}]$. Thus, we would like to apply Lemma 9.9 (with $\ell = 3$) to prove the unique existence of a minimizer of (9.17). However, Lemma 9.9 requires that the minimization be performed over a closed, bounded, and convex set S, whereas (9.17) is performed over the unbounded set H^1. To resolve this issue, we must show that the minimization (9.17) can be reformulated as a minimization over a closed, bounded, and convex set S.

To this end, note that

$$Z[\boldsymbol{0}] = \int_\Theta \mathrm{Tr}(\boldsymbol{J}^{-1}(\boldsymbol{\theta})) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \triangleq U \tag{9.88}$$

and therefore $M \leq U < \infty$. Thus, it suffices to perform the minimization (9.17) over those functions for which $Z[\boldsymbol{b}] \leq U$. We now show that this can be achieved by minimizing over a closed, bounded, and convex set S. First, note that $Z[\boldsymbol{b}] \geq \|\boldsymbol{b}\|_{L^2}^2$, so that one may choose to minimize (9.17) only over functions $\boldsymbol{b}$ for which

$$\|\boldsymbol{b}\|_{L^2}^2 \leq U. \tag{9.89}$$

Similarly, we have

$$Z[\boldsymbol{b}] \geq \int_\Theta \mathrm{Tr}\left(\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}^{-1}(\boldsymbol{\theta}) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right)^T \right) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \tag{9.90}$$

so that it suffices to minimize (9.17) over functions $\boldsymbol{b}$ for which

$$\int_\Theta \mathrm{Tr}\left(\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}^{-1}(\boldsymbol{\theta}) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right)^T \right) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq U. \tag{9.91}$$

Note that $\boldsymbol{J}(\boldsymbol{\theta})$ is bounded a.e., and therefore $\lambda_{\min}(\boldsymbol{J}^{-1}) \geq 1/K$ a.e., for some constant K. It follows that

$$\operatorname{Tr}\left(\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right)\boldsymbol{J}^{-1}(\boldsymbol{\theta})\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right)^{T}\right) \geq \frac{1}{K}\left\|\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right\|_{F}^{2} \text{ a.e.}(p_{\boldsymbol{\theta}}). \quad (9.92)$$

Combining with (9.91) yields

$$\int_{\Theta}\left\|\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right\|_{F}^{2} p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq KU. \quad (9.93)$$

From Lemma 9.10, we then have

$$\int_{\Theta}\left\|\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right\|_{F}^{2} p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq \left(\sqrt{n}+\sqrt{KU}\right)^{2}. \quad (9.94)$$

From (9.89) and (9.94) it follows that the minimization (9.17) can be limited to the closed, bounded, convex set

$$S=\left\{\boldsymbol{b}\in H^{1} : \|\boldsymbol{b}\|_{H^{1}}^{2} \leq U+\left(\sqrt{KU}+\sqrt{n}\right)^{2}\right\}. \quad (9.95)$$

Applying Lemma 9.9 proves the unique existence of a minimizer of (9.17). The proof that $0<s<\infty$ appears immediately after the statement of Proposition 9.1. □

9.C Proof of Theorem 9.3

The following is the proof of Theorem 9.3 concerning the calculation of the OBB.

Proof of Theorem 9.3. Consider the more general problem of minimizing the functional

$$Z[\boldsymbol{b}]=\int_{\Theta} F[\boldsymbol{b},\boldsymbol{\theta}]d\boldsymbol{\theta} \quad (9.96)$$

where $F[\boldsymbol{b},\boldsymbol{\theta}]$ is smooth and convex in $\boldsymbol{b}:\Theta\to\mathbb{R}^{n}$, and $\Theta\subset\mathbb{R}^{n}$ is a bounded set with a smooth boundary Λ. Then, $Z[\boldsymbol{b}]$ is also smooth and convex in $\boldsymbol{b}$, so that $\boldsymbol{b}$ is a global minimum of $Z[\boldsymbol{b}]$ if and only if the differential $\delta Z[\boldsymbol{h}]$ equals zero at $\boldsymbol{b}$ for all admissible functions $\boldsymbol{h}:\Theta\to\mathbb{R}^{n}$ [171].

By a standard technique [171, §35], it can be shown that

$$\delta Z[\boldsymbol{h}]=\epsilon\sum_{i}\int_{\Theta}\left(\frac{\partial F}{\partial b_{i}}-\sum_{j}\frac{\partial}{\partial\theta_{j}}\frac{\partial F}{\partial b_{i}^{(j)}}\right)h_{i}(\boldsymbol{\theta})d\boldsymbol{\theta} + \epsilon\sum_{i}\int_{\Lambda}\left(\frac{\partial F}{\partial b_{i}^{(1)}},\ldots,\frac{\partial F}{\partial b_{i}^{(n)}}\right)^{T}\boldsymbol{\nu}(\boldsymbol{\theta})\,h_{i}(\boldsymbol{\theta})\,d\sigma \quad (9.97)$$

where ϵ is an infinitesimal quantity, $b_i^{(j)} = \partial b_i / \partial \theta_j$, and $\boldsymbol{\nu}(\boldsymbol{\theta})$ is an outward-pointing normal at the boundary point $\boldsymbol{\theta} \in \Lambda$. We now seek conditions for which $\delta Z[\boldsymbol{h}] = 0$ for all $\boldsymbol{h}(\boldsymbol{\theta})$. Consider first functions $\boldsymbol{h}(\boldsymbol{\theta})$ which equal zero on the boundary Λ. In this case, the second integral vanishes, and we obtain the Euler–Lagrange equations

$$\forall i, \ \frac{\partial F}{\partial b_i} - \sum_j \frac{\partial}{\partial \theta_j} \frac{\partial F}{\partial b_i^{(j)}} = 0. \tag{9.98}$$

Substituting this result back into (9.97), and again using the fact that $\delta Z[\boldsymbol{h}] = 0$ for all $\boldsymbol{h}$, we obtain the boundary condition

$$\forall i, \ \forall \boldsymbol{\theta} \in \Lambda, \ \left(\frac{\partial F}{\partial b_i^{(1)}}, \ldots, \frac{\partial F}{\partial b_i^{(n)}} \right)^T \boldsymbol{\nu}(\boldsymbol{\theta}) = 0. \tag{9.99}$$

Plugging $F[\boldsymbol{b}, \boldsymbol{\theta}] = \mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}] p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ into (9.98) and (9.99) provides the required result. □

9.D Proof of Theorem 9.5

Before proving Theorem 9.5, we provide the following two lemmas, which demonstrate some symmetry properties of the CRB.

Lemma 9.11. *Under the conditions of Theorem 9.5, the functional $Z[\boldsymbol{b}]$ of* (9.12) *is rotation and reflection invariant, i.e., $Z[\boldsymbol{b}] = Z[\boldsymbol{U}\boldsymbol{b}]$ for any unitary matrix $\boldsymbol{U}$.*

Proof. We first demonstrate that $Z[\boldsymbol{b}]$ is rotation invariant. From the definitions of $Z[\boldsymbol{b}]$ and $\mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]$, we have

$$\begin{aligned} Z[\boldsymbol{b}] = &\int_\Theta \mathrm{Tr}\left[\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right)^T \right] \frac{q(\|\boldsymbol{\theta}\|)}{J(\|\boldsymbol{\theta}\|)} d\boldsymbol{\theta} \\ &+ \int_\Theta \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 q(\|\boldsymbol{\theta}\|) d\boldsymbol{\theta}. \end{aligned} \tag{9.100}$$

The second integral is clearly rotation invariant, since a rotation of $\boldsymbol{b}$ does not alter its norm. It remains to show that the first integral, which we denote by $I_1[\boldsymbol{b}]$, does not change when $\boldsymbol{b}$ is rotated. To this end, we begin by considering a rotation about the first two coordinates, such that $\boldsymbol{b}$ is transformed to $\tilde{\boldsymbol{b}} \triangleq \boldsymbol{R}_\phi \boldsymbol{b}$, where the rotation matrix $\boldsymbol{R}_\phi$ is defined such that

$$\begin{aligned} \boldsymbol{R}_\phi \boldsymbol{b} = (&b_1 \cos\phi + b_2 \sin\phi, \\ &- b_1 \sin\phi + b_2 \cos\phi, b_3, \ldots, b_n)^T. \end{aligned} \tag{9.101}$$

We must thus show that $I_1[\boldsymbol{b}] = I_1[\tilde{\boldsymbol{b}}]$. Let us perform the change of variables $\boldsymbol{\theta} \mapsto \tilde{\boldsymbol{\theta}}$, where $\tilde{\boldsymbol{\theta}} = \boldsymbol{R}_{(-\phi)}\boldsymbol{\theta}$. Rewriting the trace in (9.100) as a sum, we have

$$I_1[\tilde{\boldsymbol{b}}] = \int_\Theta \sum_{i,j} \left(\delta_{ij} + \frac{\partial \tilde{b}_i}{\partial \theta_j}\right)^2 \frac{q(\|\tilde{\boldsymbol{\theta}}\|)}{J(\|\tilde{\boldsymbol{\theta}}\|)} d\tilde{\boldsymbol{\theta}} \tag{9.102}$$

where we have used the facts that $\|\boldsymbol{\theta}\| = \|\tilde{\boldsymbol{\theta}}\|$ and that Θ does not change under the change of variables.

We now demonstrate some properties of the transformation of $\boldsymbol{b}$ and $\boldsymbol{\theta}$. First, we have, for any j,

$$\begin{aligned}\left(\frac{\partial \tilde{b}_1}{\partial \theta_j}\right)^2 + \left(\frac{\partial \tilde{b}_2}{\partial \theta_j}\right)^2 &= \left(\frac{\partial b_1}{\partial \theta_j}\cos\phi + \frac{\partial b_2}{\partial \theta_j}\sin\phi\right)^2 \\ &\quad + \left(-\frac{\partial b_1}{\partial \theta_j}\sin\phi + \frac{\partial b_2}{\partial \theta_j}\cos\phi\right)^2 \\ &= \left(\frac{\partial b_1}{\partial \theta_j}\right)^2 + \left(\frac{\partial b_2}{\partial \theta_j}\right)^2. \end{aligned} \tag{9.103}$$

Also, for any i,

$$\begin{aligned}\left(\frac{\partial b_i}{\partial \tilde{\theta}_1}\right)^2 + \left(\frac{\partial b_i}{\partial \tilde{\theta}_2}\right)^2 &= \left(\frac{\partial b_i}{\partial \theta_1}\frac{\partial \theta_1}{\partial \tilde{\theta}_1} + \frac{\partial b_i}{\partial \theta_2}\frac{\partial \theta_2}{\partial \tilde{\theta}_1}\right)^2 \\ &\quad + \left(\frac{\partial b_i}{\partial \theta_1}\frac{\partial \theta_1}{\partial \tilde{\theta}_2} + \frac{\partial b_i}{\partial \theta_2}\frac{\partial \theta_2}{\partial \tilde{\theta}_2}\right)^2 \\ &= \left(\frac{\partial b_i}{\partial \theta_1}\right)^2 + \left(\frac{\partial b_i}{\partial \theta_2}\right)^2 \end{aligned} \tag{9.104}$$

where we used the fact that $\boldsymbol{\theta} = \boldsymbol{R}_\phi \tilde{\boldsymbol{\theta}}$. Third, we have

$$\begin{aligned}\frac{\partial \tilde{b}_1}{\partial \theta_1} &= \frac{\partial b_1}{\partial \tilde{\theta}_1}\cos^2\phi + \frac{\partial b_1}{\partial \tilde{\theta}_2}\sin\phi\cos\phi \\ &\quad + \frac{\partial b_2}{\partial \tilde{\theta}_1}\sin\phi\cos\phi + \frac{\partial b_2}{\partial \tilde{\theta}_2}\sin^2\phi, \\ \frac{\partial \tilde{b}_2}{\partial \theta_2} &= \frac{\partial b_1}{\partial \tilde{\theta}_1}\sin^2\phi - \frac{\partial b_1}{\partial \tilde{\theta}_2}\sin\phi\cos\phi \\ &\quad - \frac{\partial b_2}{\partial \tilde{\theta}_1}\sin\phi\cos\phi + \frac{\partial b_2}{\partial \tilde{\theta}_2}\cos^2\phi, \end{aligned} \tag{9.105}$$

so that

$$\frac{\partial \tilde{b}_1}{\partial \theta_1} + \frac{\partial \tilde{b}_2}{\partial \theta_2} = \frac{\partial b_1}{\partial \tilde{\theta}_1} + \frac{\partial b_2}{\partial \tilde{\theta}_2}. \tag{9.106}$$

We now show that

$$\sum_{i,j}\left(\delta_{ij} + \frac{\partial \tilde{b}_i}{\partial \theta_j}\right)^2 = \sum_{i,j}\left(\delta_{ij} + \frac{\partial b_i}{\partial \tilde{\theta}_j}\right)^2. \tag{9.107}$$

For terms with $i, j \geq 3$, we have $b_i = \tilde{b}_i$ and $\theta_j = \tilde{\theta}_j$, so that replacing $\tilde{\boldsymbol{b}}$ with $\boldsymbol{b}$ and $\boldsymbol{\theta}$ with $\tilde{\boldsymbol{\theta}}$ does not change the result. The terms with $i = 1, 2$ and $j \geq 3$ do not change because of (9.103), while the terms with $i \geq 3$ and $j = 1, 2$ do not change because of (9.104). It remains to show that the terms $i, j = 1, 2$ do not modify the sum. To this end, we write out these four terms as

$$\begin{aligned}
&\left(1 + \frac{\partial \tilde{b}_1}{\partial \theta_1}\right)^2 + \left(1 + \frac{\partial \tilde{b}_2}{\partial \theta_2}\right)^2 + \left(\frac{\partial \tilde{b}_1}{\partial \theta_2}\right)^2 + \left(\frac{\partial \tilde{b}_2}{\partial \theta_1}\right)^2 \\
&\quad = 2 + 2\frac{\partial \tilde{b}_1}{\partial \theta_1} + 2\frac{\partial \tilde{b}_2}{\partial \theta_2} \\
&\qquad + \left(\frac{\partial \tilde{b}_1}{\partial \theta_1}\right)^2 + \left(\frac{\partial \tilde{b}_1}{\partial \theta_2}\right)^2 + \left(\frac{\partial \tilde{b}_2}{\partial \theta_1}\right)^2 + \left(\frac{\partial \tilde{b}_2}{\partial \theta_2}\right)^2 \\
&\quad = 2 + 2\frac{\partial b_1}{\partial \tilde{\theta}_1} + 2\frac{\partial b_2}{\partial \tilde{\theta}_2} \\
&\qquad + \left(\frac{\partial b_1}{\partial \tilde{\theta}_1}\right)^2 + \left(\frac{\partial b_1}{\partial \tilde{\theta}_2}\right)^2 + \left(\frac{\partial b_2}{\partial \tilde{\theta}_1}\right)^2 + \left(\frac{\partial b_2}{\partial \tilde{\theta}_2}\right)^2 \\
&\quad = \left(1 + \frac{\partial b_1}{\partial \tilde{\theta}_1}\right)^2 + \left(1 + \frac{\partial b_2}{\partial \tilde{\theta}_2}\right)^2 + \left(\frac{\partial b_1}{\partial \tilde{\theta}_2}\right)^2 + \left(\frac{\partial b_2}{\partial \tilde{\theta}_1}\right)^2
\end{aligned} \tag{9.108}$$

where, in the second transition, we have used (9.103), (9.104), and (9.106). It follows that $I_1[\tilde{\boldsymbol{b}}]$ of (9.102) is equal to $I_1[\boldsymbol{b}]$, and hence $Z[\boldsymbol{b}] = Z[\tilde{\boldsymbol{b}}]$. The result similarly holds for rotations about any other two coordinates. Since any rotation can be decomposed into a sequence of two-coordinate rotations, we conclude that $Z[\boldsymbol{b}]$ is rotation invariant.

Next, we prove that $Z[\boldsymbol{b}]$ is invariant to reflections through hyperplanes containing the origin. Since $Z[\boldsymbol{b}]$ is invariant to rotations, it suffices to choose a single hyperplane, say $\{\boldsymbol{\theta} : \theta_1 = 0\}$. Let

$$\tilde{\boldsymbol{b}} \triangleq (-b_1(\boldsymbol{\theta}), b_2(\boldsymbol{\theta}), \ldots, b_n(\boldsymbol{\theta}))^T \tag{9.109}$$

be the reflection of $\boldsymbol{b}$, and consider the corresponding change of variables

$$\tilde{\boldsymbol{\theta}} \triangleq (-\theta_1, \theta_2, \ldots, \theta_n)^T. \tag{9.110}$$

By the symmetry assumptions, $p_{\boldsymbol{\theta}}$ and $\boldsymbol{J}$ are unaffected by the change of variables; furthermore, $\partial \tilde{\boldsymbol{b}} / \partial \tilde{\boldsymbol{\theta}} = \partial \boldsymbol{b} / \partial \boldsymbol{\theta}$. It follows that $\mathrm{CRB}[\tilde{\boldsymbol{b}}, \tilde{\boldsymbol{\theta}}] = \mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]$, and therefore $Z[\boldsymbol{b}] = Z[\tilde{\boldsymbol{b}}]$. □

Lemma 9.12. *Suppose $\boldsymbol{b}(\boldsymbol{\theta})$ is radial and rotation invariant, i.e., $\boldsymbol{b}(\boldsymbol{\theta}) = t(\|\boldsymbol{\theta}\|^2)\boldsymbol{\theta}$ for some function $t \in H^1$. Also suppose that $\boldsymbol{J}(\boldsymbol{\theta}) = J(\|\boldsymbol{\theta}\|)\boldsymbol{I}$, where $J(\cdot)$ is a scalar function. Then,* $\mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]$ *of* (9.11) *is rotation invariant in $\boldsymbol{\theta}$, i.e.,* $\mathrm{CRB}[\boldsymbol{b}, \boldsymbol{R}\boldsymbol{\theta}] = \mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]$ *for any rotation matrix $\boldsymbol{R}$.*

Proof. We will show that $\mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]$ depends on $\boldsymbol{\theta}$ only through $\|\boldsymbol{\theta}\|^2$, and is therefore rotation

invariant. For the given value of $\boldsymbol{b}(\boldsymbol{\theta})$ and $\boldsymbol{J}(\boldsymbol{\theta})$, we have

$$\begin{aligned}
&\mathrm{CRB}[\boldsymbol{b},\boldsymbol{\theta}] \\
&= \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 + \mathrm{Tr}\left[\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right)\boldsymbol{J}^{-1}(\boldsymbol{\theta})\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right)^T\right] \\
&= t^2\|\boldsymbol{\theta}\|^2 + \frac{1}{J(\|\boldsymbol{\theta}\|)}\,\mathrm{Tr}\left[\left(\boldsymbol{I}+\frac{\partial t\boldsymbol{\theta}}{\partial \boldsymbol{\theta}}\right)\left(\boldsymbol{I}+\frac{\partial t\boldsymbol{\theta}}{\partial \boldsymbol{\theta}}\right)^T\right] \qquad (9.111)
\end{aligned}$$

where, for notational convenience, we have omitted the dependence of t on $\|\boldsymbol{\theta}\|^2$. It remains to show that the trace in the above expression is a function of $\boldsymbol{\theta}$ only through $\|\boldsymbol{\theta}\|^2$. To this end, we note that

$$\frac{\partial b_i}{\partial \theta_j} = t\delta_{ij} + t'\theta_i\frac{\partial\|\boldsymbol{\theta}\|^2}{\partial\theta_j} = t\delta_{ij} + 2t'\theta_i\theta_j \qquad (9.112)$$

where δ_{ij} is the Kronecker delta. It follows that

$$\left(\delta_{ij}+\frac{\partial b_i}{\partial\theta_j}\right)^2 = (1+t)^2\delta_{ij} + 4(1+t)t'\theta_i\theta_j\delta_{ij} + 4t'^2\theta_i^2\theta_j^2. \qquad (9.113)$$

Therefore

$$\begin{aligned}
\mathrm{Tr}\left[\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right)\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right)^T\right] &= \sum_{i,j}\left(\delta_{ij}+\frac{\partial b_i}{\partial\theta_j}\right)^2 \\
&= n(1+t)^2 + 4t'^2\sum_{i,j}\theta_i^2\theta_j^2 + 4(1+t)t'\sum_i\theta_i^2 \\
&= n(1+t)^2 + 4t'^2\|\boldsymbol{\theta}\|^4 + 4(1+t)t'\|\boldsymbol{\theta}\|^2. \qquad (9.114)
\end{aligned}$$

Thus, $\mathrm{CRB}[\boldsymbol{b},\boldsymbol{\theta}]$ depends on $\boldsymbol{\theta}$ only through $\|\boldsymbol{\theta}\|^2$, completing the proof. □

Proof of Theorem 9.5. We have seen in Theorem 9.3 that the solution of (9.20) is unique. Now suppose that the optimum $\boldsymbol{b}$ is not rotation invariant, i.e., there exists a rotation matrix $\boldsymbol{R}$ such that $\boldsymbol{R}\boldsymbol{b}(\boldsymbol{\theta})$ is not identical to $\boldsymbol{b}(\boldsymbol{\theta})$. By Lemma 9.11, $\boldsymbol{R}\boldsymbol{b}(\boldsymbol{\theta})$ is also optimal, which is a contradiction.

Furthermore, suppose that $\boldsymbol{b}$ is not radial, i.e., for some value of $\boldsymbol{\theta}$, $\boldsymbol{b}(\boldsymbol{\theta})$ contains a component perpendicular to the vector $\boldsymbol{\theta}$. Consider a hyperplane passing through the origin, whose normal is the aforementioned perpendicular component. By Lemma 9.11, The reflection of $\boldsymbol{b}$ through this hyperplane is also an optimal solution of (9.20), which is again a contradiction. Therefore, the optimum $\boldsymbol{b}$ is spherically symmetric and radial, so that it can be written as

$$\boldsymbol{b}(\boldsymbol{\theta}) = b(\|\boldsymbol{\theta}\|)\frac{\boldsymbol{\theta}}{\|\boldsymbol{\theta}\|} \qquad (9.115)$$

where $b(\cdot)$ is a scalar function.

To determine the value of $b(\cdot)$, it suffices to analyze the differential equation (9.21) along a straight line from the origin to the boundary. We choose a line along the θ_1 axis, and begin by calculating the derivatives of $b_1(\boldsymbol{\theta})$, $q(\|\boldsymbol{\theta}\|)$, and $\boldsymbol{J}(\|\boldsymbol{\theta}\|)$ along this axis. The derivative of $q(\|\boldsymbol{\theta}\|)$ is given by

$$\frac{\partial q}{\partial \theta_j} = q'(\rho)\frac{\theta_j}{\rho} \tag{9.116}$$

where we have denoted $\rho = \|\boldsymbol{\theta}\|$, so that ρ is weakly differentiable and

$$\frac{\partial \rho}{\partial \theta_j} = \frac{\theta_j}{\rho}. \tag{9.117}$$

Along the θ_1 axis, we have $\theta_1 = \rho$ while $\theta_2 = \cdots = \theta_n = 0$, so that

$$\left.\frac{\partial q}{\partial \theta_j}\right|_{\boldsymbol{\theta}=\rho\boldsymbol{e}_1} = q'(\rho)\delta_{j1}. \tag{9.118}$$

Similarly, since $\boldsymbol{J}(\boldsymbol{\theta}) = J(\rho)\boldsymbol{I}$,

$$\frac{\partial (\boldsymbol{J}^{-1})_{jk}}{\partial \theta_j} = -\frac{J'(\rho)}{J^2(\rho)}\frac{\theta_j}{\rho}\delta_{jk} \tag{9.119}$$

so that along the θ_1 axis

$$\left.\frac{\partial (\boldsymbol{J}^{-1})_{jk}}{\partial \theta_j}\right|_{\boldsymbol{\theta}=\rho\boldsymbol{e}_1} = -\frac{J'(\rho)}{J^2(\rho)}\delta_{jk}\delta_{j1}. \tag{9.120}$$

From (9.115), we have

$$\frac{\partial b_i}{\partial \theta_j} = b'(\rho)\frac{\theta_i\theta_j}{\rho^2} + \frac{b(\rho)}{\rho}\left(\delta_{ij} - \frac{\theta_i\theta_j}{\rho^2}\right). \tag{9.121}$$

Thus, on the θ_1 axis, we have

$$\left.\frac{\partial b_1}{\partial \theta_j}\right|_{\boldsymbol{\theta}=\rho\boldsymbol{e}_1} = b'(\rho)\delta_{j1}. \tag{9.122}$$

The second derivative of $b_i(\boldsymbol{\theta})$ can be shown to equal

$$\begin{aligned}\frac{\partial^2 b_i}{\partial \theta_j \partial \theta_k} &= b''(\rho)\frac{\theta_i\theta_j\theta_k}{\rho^3} \\ &+ \left(\frac{b'(\rho)}{\rho} - \frac{b(\rho)}{\rho^2}\right)\left(\frac{\theta_i}{\rho}\delta_{jk} + \frac{\theta_j}{\rho}\delta_{ik} + \frac{\theta_k}{\rho}\delta_{ij} - 3\frac{\theta_i\theta_j\theta_k}{\rho^3}\right).\end{aligned} \tag{9.123}$$

Therefore, on the θ_1 axis

$$\begin{aligned}\left.\frac{\partial^2 b_1}{\partial \theta_1^2}\right|_{\boldsymbol{\theta}=\rho\boldsymbol{e}_1} &= b''(\rho) \\ \left.\frac{\partial^2 b_1}{\partial \theta_j^2}\right|_{\boldsymbol{\theta}=\rho\boldsymbol{e}_1} &= \frac{b'(\rho)}{\rho} - \frac{b(\rho)}{\rho^2} && (j \neq 1) \\ \left.\frac{\partial^2 b_1}{\partial \theta_j \partial \theta_k}\right|_{\boldsymbol{\theta}=\rho\boldsymbol{e}_1} &= 0 && (j,k \neq 1).\end{aligned} \tag{9.124}$$

Substituting these derivatives into (9.21), we obtain

$$q(\rho)b(\rho) = \frac{q(\rho)}{J(\rho)}\left(b''(\rho) + (n-1)\frac{b'(\rho)}{\rho} - (n-1)\frac{b(\rho)}{\rho^2}\right) + (1+b'(\rho))\left(\frac{q'(\rho)}{J(\rho)} - q(\rho)\frac{J'(\rho)}{J^2(\rho)}\right) \tag{9.125}$$

which is equivalent to (9.25).

To obtain the boundary conditions, observe that Lemma 9.11 implies $\boldsymbol{b}(\mathbf{0}) = \mathbf{0}$, whence we conclude that $b(0) = 0$. Next, evaluate the boundary condition (9.22) at boundary point $\boldsymbol{\theta} = r\boldsymbol{e}_1$, where the surface normal $\boldsymbol{\nu}(\boldsymbol{\theta})$ equals $\boldsymbol{e}_1$, so that

$$1 + b'(\rho) = 1 + \frac{\partial b_1}{\partial \theta_1} = 0, \quad \boldsymbol{\theta} = r\boldsymbol{e}_1 \tag{9.126}$$

which is equivalent to the boundary condition $b'(r) = -1$.

To find the OBB (9.24), we must now calculate $Z[\boldsymbol{b}]$ for the obtained bias function (9.115). To this end, note that, by Lemma 9.12, $\mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]$ is rotation invariant in $\boldsymbol{\theta}$ for the required $\boldsymbol{b}(\boldsymbol{\theta})$. Thus, the integrand $\mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]q(\|\boldsymbol{\theta}\|)$ is constant on any $(n-1)$-sphere centered on the origin, so that

$$Z[\boldsymbol{b}] = \int_0^r \mathrm{CRB}[\boldsymbol{b}, \rho\boldsymbol{e}_1]q(\rho)S_n(\rho)d\rho \tag{9.127}$$

where

$$S_n(\rho) = \frac{2\pi^{n/2}}{\Gamma(n/2)}\rho^{n-1} \tag{9.128}$$

is the hypersurface area of an $(n-1)$-sphere of radius ρ [167]. It thus suffices to calculate the value of $\mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]$ at points along the θ_1 axis. From (9.121), it follows that

$$\left.\frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}}\right|_{\boldsymbol{\theta}=\rho\boldsymbol{e}_1} = \mathrm{diag}\left(b'(\rho), \frac{b(\rho)}{\rho}, \ldots, \frac{b(\rho)}{\rho}\right). \tag{9.129}$$

Substituting this into the definition of $\mathrm{CRB}[\boldsymbol{b}, \boldsymbol{\theta}]$, we obtain

$$\mathrm{CRB}[\boldsymbol{b}, \rho\boldsymbol{e}_1] = b^2(\rho) + \frac{1}{J(\rho)}(1+b'(\rho))^2 + \frac{n-1}{J(\rho)}\left(1 + \frac{b(\rho)}{\rho}\right)^2. \tag{9.130}$$

Combining (9.130) with (9.127) yields (9.24), as required. □

9.E Proofs of Asymptotic Properties

Theorems 9.7 and 9.8 demonstrate asymptotic tightness of the OBB. The proofs of these two theorems follow.

Proof of Theorem 9.7. We begin the proof by studying a certain optimization problem, whose relevance will be demonstrated shortly. Let $t \geq 0$ be a constant and consider the problem

$$\begin{aligned} u(t) = \inf_{\boldsymbol{b} \in H^1} & \int_\Theta \left\| \boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ & \text{s.t.} \int_\Theta \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq t. \end{aligned} \tag{9.131}$$

Notice that $u(t) \leq n$ for all t, since an objective having a value of n is achieved by the function $\boldsymbol{b}(\boldsymbol{\theta}) = \boldsymbol{0}$. Thus, it suffices to perform the minimization (9.131) over functions $\boldsymbol{b} \in H^1$ satisfying

$$\int_\Theta \left\| \boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq n. \tag{9.132}$$

It follows from Lemma 9.10 that such functions also satisfy

$$\int_\Theta \left\| \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq (2\sqrt{n})^2 = 4n. \tag{9.133}$$

Therefore, (9.131) is equivalent to the minimization

$$u(t) = \inf_{\boldsymbol{b} \in S_t} \int_\Theta \left\| \boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \tag{9.134}$$

where

$$\begin{aligned} S_t = \Bigg\{ \boldsymbol{b} \in H^1 : & \int_\Theta \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq t, \\ & \int_\Theta \left\| \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq 4n \Bigg\}. \end{aligned} \tag{9.135}$$

The set S_t is convex, closed, and bounded in H^1. Applying Lemma 9.9 (with $\ell = 2$) implies that there exists a function $\boldsymbol{b}_{\text{opt}} \in S_t$ which minimizes (9.134), and hence also minimizes (9.131).

Note that the objective in (9.131) is zero if and only if

$$\frac{\partial \boldsymbol{b}_{\text{opt}}}{\partial \boldsymbol{\theta}} = -\boldsymbol{I} \quad \text{a.e. } (p_{\boldsymbol{\theta}}). \tag{9.136}$$

The only functions in H^1 satisfying this requirement are the functions

$$\boldsymbol{b}(\boldsymbol{\theta}) = \boldsymbol{k} - \boldsymbol{\theta} \quad \text{a.e. } (p_{\boldsymbol{\theta}}) \tag{9.137}$$

for some constant $\boldsymbol{k} \in \mathbb{R}^n$. Let $\boldsymbol{\mu} \triangleq E\{\boldsymbol{\theta}\}$ and define

$$v \triangleq E\left\{ \|\boldsymbol{\theta} - E\{\boldsymbol{\theta}\}\|^2 \right\}. \tag{9.138}$$

For functions of the form (9.137), the constraint of (9.131) is given by

$$\begin{aligned}\int_\Theta \|\boldsymbol{k}-\boldsymbol{\theta}\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) &= \int_\Theta \|\boldsymbol{k}-\boldsymbol{\mu}+\boldsymbol{\mu}-\boldsymbol{\theta}\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ &= \|\boldsymbol{k}-\boldsymbol{\mu}\|^2 + v \\ &\geq v. \end{aligned} \tag{9.139}$$

In (9.139), equality is obtained if and only if $\boldsymbol{k}=\boldsymbol{\mu}$. Therefore, if $t < v$, no functions satisfying (9.136) are feasible, and thus

$$\begin{aligned} u(t) &= 0 \quad \text{if } t \geq v, \\ u(t) &> 0 \quad \text{if } t < v. \end{aligned} \tag{9.140}$$

We now return to the setting of Theorem 9.7. We must show that $\beta_N \to v$ as $N \to \infty$. We denote functions corresponding to the problem of estimating $\boldsymbol{\theta}$ from $\boldsymbol{x}^{(N)}$ with a superscript (N). Thus, for example, $Z^{(N)}[\boldsymbol{b}]$ denotes the functional $Z[\boldsymbol{b}]$ of (9.12) for the problem corresponding to the measurement vector $\boldsymbol{x}^{(N)}$.

Since all eigenvalues of $\boldsymbol{J}^{(N)}(\boldsymbol{\theta})$ decrease monotonically with N for $p_{\boldsymbol{\theta}}$-almost all $\boldsymbol{\theta}$, we have

$$\mathrm{CRB}^{(N)}[\boldsymbol{b},\boldsymbol{\theta}] \leq \mathrm{CRB}^{(N+1)}[\boldsymbol{b},\boldsymbol{\theta}] \tag{9.141}$$

for any $\boldsymbol{b} \in H^1$, for $p_{\boldsymbol{\theta}}$-almost all $\boldsymbol{\theta}$, and for all N. Therefore

$$Z^{(N)}[\boldsymbol{b}] \leq Z^{(N+1)}[\boldsymbol{b}]. \tag{9.142}$$

for any $\boldsymbol{b} \in H^1$ and for all N. It follows that for all N

$$\beta_N = \min_{\boldsymbol{b}\in H^1} Z^{(N)}[\boldsymbol{b}] \leq \min_{\boldsymbol{b}\in H^1} Z^{(N+1)}[\boldsymbol{b}] = \beta_{N+1} \tag{9.143}$$

so that β_N is a non-decreasing sequence. Furthermore, note that

$$Z^{(N)}[\boldsymbol{\mu}-\boldsymbol{\theta}] = v \quad \text{for all } N \tag{9.144}$$

where v is given by (9.138). Therefore, $\beta_N \leq v$ for all N. Thus β_N converges to some value q, and we have

$$\beta_N \leq q \leq v \quad \text{for all } N. \tag{9.145}$$

To prove the theorem, it remains to show that $q = v$.

Let $\boldsymbol{b}^{(N)}$ be the minimizer of (9.17) when $\boldsymbol{\theta}$ is estimated from $\boldsymbol{x}^{(N)}$; this minimizer exists by virtue of Proposition 9.1. We then have

$$\beta_N = Z^{(N)}[\boldsymbol{b}^{(N)}] \leq q \tag{9.146}$$

and therefore

$$\int_\Theta \|\boldsymbol{b}^{(N)}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq q. \tag{9.147}$$

It follows that $\boldsymbol{b}^{(N)}$ satisfies the constraint of the optimization problem (9.131) with $t = q$. As a consequence, we have

$$\int_\Theta \left\| \boldsymbol{I} + \frac{\partial \boldsymbol{b}^{(N)}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \geq u(q). \tag{9.148}$$

Define

$$\lambda_N \triangleq \operatorname*{ess\,sup}_{\boldsymbol{\theta}\in\Theta} \lambda_{\max}(\boldsymbol{J}^{(N)}(\boldsymbol{\theta})) \tag{9.149}$$

and note that $\lambda_N > 0$ for all N, since $\boldsymbol{J}^{(N)}(\boldsymbol{\theta})$ is positive definite. Thus

$$\begin{aligned} Z^{(N)}[\boldsymbol{b}^{(N)}] &\geq \int_\Theta \operatorname{Tr}\left[\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}^{(N)}}{\partial \boldsymbol{\theta}} \right) \left(\boldsymbol{J}^{(N)}(\boldsymbol{\theta}) \right)^{-1} \right. \\ &\qquad \left. \cdot \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}^{(N)}}{\partial \boldsymbol{\theta}} \right)^T \right] p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ &\geq \frac{1}{\lambda_N} \int_\Theta \left\| \boldsymbol{I} + \frac{\partial \boldsymbol{b}^{(N)}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ &\geq \frac{u(q)}{\lambda_N}. \end{aligned} \tag{9.150}$$

Assume by contradiction that $q < v$. From (9.140), it then follows that $u(q) > 0$. Since all eigenvalues of $\boldsymbol{J}^{(N)}(\boldsymbol{\theta})$ decrease to zero, we have $\lambda_N \to 0$, and thus

$$\beta_N \geq \frac{u(q)}{\lambda_N} \to \infty. \tag{9.151}$$

This contradicts the fact (9.145) that $\beta_N \leq v$. We conclude that $q = v$, as required. □

Proof of Theorem 9.8. The proof is analogous to that of Theorem 9.7. We begin by considering the optimization problem

$$\begin{aligned} &\inf_{\boldsymbol{b}\in H^1} \int_\Theta \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ &\text{s.t.} \int_\Theta \operatorname{Tr}\left(\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}^{-1}(\boldsymbol{\theta}) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right)^T \right) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq t \end{aligned} \tag{9.152}$$

for some constant $t \geq 0$. Denote the minimum value of (9.152) by $w(t)$. Let $\boldsymbol{\mu} = E\{\boldsymbol{\theta}\}$ and note that $\boldsymbol{b}(\boldsymbol{\theta}) = \boldsymbol{\mu} - \boldsymbol{\theta}$ satisfies the constraint in (9.152) for any $t \geq 0$, and has an objective equal to

v of (9.138). Thus, to determine $w(t)$, it suffices to minimize (9.152) over the set

$$S_t = \Bigg\{ \boldsymbol{b} \in H^1 : \int_\Theta \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \le v, \\ \int_\Theta \mathrm{Tr}\left(\left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right) \boldsymbol{J}^{-1}(\boldsymbol{\theta}) \left(\boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right)^T \right) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \le t \Bigg\}.$$

Define

$$\lambda \triangleq \operatorname*{ess\,sup}_{\boldsymbol{\theta} \in \Theta} \lambda_{\max}(\boldsymbol{J}(\boldsymbol{\theta})). \tag{9.153}$$

Since $\boldsymbol{J}(\boldsymbol{\theta})$ is positive definite almost everywhere, we have $\lambda > 0$. For any $\boldsymbol{b} \in S_t$, we have

$$\frac{1}{\lambda} \int_\Theta \left\| \boldsymbol{I} + \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \le t \tag{9.154}$$

and therefore, by Lemma 9.10,

$$\int_\Theta \left\| \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \le \left(\sqrt{t\lambda} + \sqrt{n} \right)^2. \tag{9.155}$$

Hence, for any $\boldsymbol{b} \in S_t$,

$$\begin{aligned} \|\boldsymbol{b}\|_{H^1}^2 &= \int_\Theta \|\boldsymbol{b}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) + \int_\Theta \left\| \frac{\partial \boldsymbol{b}}{\partial \boldsymbol{\theta}} \right\|_F^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ &\le v + \left(\sqrt{t\lambda} + \sqrt{n} \right)^2. \end{aligned} \tag{9.156}$$

Thus S_t is bounded for all t. It is straightforward to show that S_t is also closed and convex. Therefore, employing Lemma 9.9 (with $\ell = 1$) ensures that there exists a (unique) $\boldsymbol{b}_{\mathrm{opt}} \in S_t$ minimizing (9.152).

Note that the objective in (9.152) is 0 if and only if $\boldsymbol{b}_{\mathrm{opt}}(\boldsymbol{\theta}) = \boldsymbol{0}$ almost everywhere. So, if $\boldsymbol{0} \in S_t$, we have $w(t) = 0$, and otherwise $w(t) > 0$. Let us define

$$s \triangleq E\Big\{ \mathrm{Tr}(\boldsymbol{J}^{-1}(\boldsymbol{\theta})) \Big\} \tag{9.157}$$

and note that $\boldsymbol{0} \in S_t$ if and only if $t \ge s$. Thus

$$\begin{aligned} w(t) &= 0 \quad \text{for } t \ge s \\ w(t) &> 0 \quad \text{otherwise.} \end{aligned} \tag{9.158}$$

Let us now return to the setting of Theorem 9.8. For simplicity, we denote functions corresponding to the problem of estimating $\boldsymbol{\theta}$ from $\{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(N)}\}$ with a superscript (N). For example, from the additive property of the Fisher information [16, §3.4], we have

$$\boldsymbol{J}^{(N)}(\boldsymbol{\theta}) = N\boldsymbol{J}(\boldsymbol{\theta}). \tag{9.159}$$

It follows that

$$(N+1)\mathrm{CRB}^{(N+1)}[\boldsymbol{b},\boldsymbol{\theta}] \geq N\mathrm{CRB}^{(N)}[\boldsymbol{b},\boldsymbol{\theta}] \tag{9.160}$$

for all $\boldsymbol{b} \in H^1$, all $\boldsymbol{\theta} \in \Theta$, and all N. Therefore

$$(N+1)Z^{(N+1)}[\boldsymbol{b}] \geq NZ^{(N)}[\boldsymbol{b}] \tag{9.161}$$

for all $\boldsymbol{b} \in H^1$, and hence

$$\begin{aligned}(N+1)\beta_{N+1} &= \min_{\boldsymbol{b}\in H^1}\left((N+1)Z^{(N+1)}[\boldsymbol{b}]\right) \\ &\geq \min_{\boldsymbol{b}\in H^1}\left(NZ^{(N)}[\boldsymbol{b}]\right) \\ &= N\beta_N. \end{aligned} \tag{9.162}$$

Thus $\{N\beta_N\}$ is a non-decreasing sequence. Furthermore, we have

$$NZ^{(N)}[\mathbf{0}] = s \tag{9.163}$$

so that $N\beta_N \leq s$ for all N. It follows that $\{N\beta_N\}$ is non-decreasing and bounded, and therefore converges to some value r such that

$$N\beta_N \leq r \leq s \quad \text{for all } N. \tag{9.164}$$

To prove the theorem, we must show that $r = s$.

Let $\boldsymbol{b}^{(N)} \in H^1$ denote the minimizer of (9.17) when $\boldsymbol{\theta}$ is estimated from $\{\boldsymbol{x}^{(1)},\ldots,\boldsymbol{x}^{(N)}\}$ (the existence of $\boldsymbol{b}^{(N)}$ is guaranteed by Proposition 9.1). We then have $N\beta_N = NZ^{(N)}[\boldsymbol{b}^{(N)}] \leq r$, so that

$$\int_\Theta \mathrm{Tr}\left(\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}^{(N)}}{\partial\boldsymbol{\theta}}\right)\boldsymbol{J}^{-1}(\boldsymbol{\theta})\left(\boldsymbol{I}+\frac{\partial \boldsymbol{b}^{(N)}}{\partial\boldsymbol{\theta}}\right)^T\right) p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \leq r. \tag{9.165}$$

Thus, $\boldsymbol{b}^{(N)}$ satisfies the constraint of (9.152) with $t = r$. As a consequence, we have

$$\int_\Theta \|\boldsymbol{b}^{(N)}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \geq w(r) \tag{9.166}$$

and therefore

$$\begin{aligned} N\beta_N &= NZ^{(N)}[\boldsymbol{b}^{(N)}] \\ &\geq N\int_\Theta \|\boldsymbol{b}^{(N)}(\boldsymbol{\theta})\|^2 p_{\boldsymbol{\theta}}(d\boldsymbol{\theta}) \\ &\geq Nw(r). \end{aligned} \tag{9.167}$$

Now suppose by contradiction that $r < s$. It follows from (9.158) that $w(r) > 0$. Hence, by (9.167), $N\beta_N \to \infty$, which contradicts the fact that $N\beta_N$ is bounded. We conclude that $r = s$, as required. □

Chapter 10

A Comment on the Weiss–Weinstein Bound

This chapter is a reprint of the paper:

- Z. Ben-Haim and Y. C. Eldar, "A comment on the use of the Weiss-Weinstein bound for constrained parameter sets," *IEEE Trans. Inform. Theory*, vol. 54, no. 10, October 2008, pp. 4682-4684.

10.1 Introduction

We consider the problem of estimating a random vector $\boldsymbol{\theta}$ from observations $\boldsymbol{x}$, where the quality of an estimator $\boldsymbol{g}(\boldsymbol{x})$ is measured by its mean-squared error (MSE) $E\{\|\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{\theta}\|^2\}$. It is well-known that the posterior mean $E\{\boldsymbol{\theta}|\boldsymbol{x}\}$ is the technique minimizing the MSE. However, in many cases, determining the posterior mean is computationally prohibitive, and various approaches have been developed as alternatives. It is therefore of interest to determine the degradation in accuracy resulting from the use of suboptimal methods. Unfortunately, computation of the optimal MSE is itself infeasible in many cases. This has led to a large body of work seeking to find simple lower bounds for the minimum MSE in a given estimation problem [59–61, 65].

In a landmark paper, Weiss and Weinstein developed a general technique for deriving lower bounds on the minimum achievable MSE [61]. A noteworthy feature of their method is that it requires almost no regularity assumptions on the problem setting. The bound was further developed in [60], and has been used in a variety of practical applications [172–174].

A common setting in which the Weiss–Weinstein bound (WWB) is often applied concerns constrained parameter sets, i.e., situations in which the parameter $\boldsymbol{\theta}$ occurs with probability 1 in a subset Θ of $\mathbb{R}^m$. For example, in time-delay estimation, the delay is sometimes assumed to be uniformly distributed in a given interval [61]. While the WWB continues to hold in the constrained setting, some of the simplifications presented in [61] and [60] are not valid in this case.

In this correspondence, we point out the versions of the WWB which do not necessarily hold for constrained parameter sets. We discuss the regularity conditions under which these simplified versions are valid, and conclude with an example in which the simplified version of the bound yields incorrect (and unreasonable) results.

10.2 Background and Summary

We begin by recalling some definitions and results relating to the WWB [60, 61]. Let $\boldsymbol{\theta}$ and $\boldsymbol{x}$ be finite-variance random vectors whose joint probability density function (pdf) is $f(\boldsymbol{x}, \boldsymbol{\theta})$. Suppose that $f(\boldsymbol{x}, \boldsymbol{\theta})$ is nonzero only for values of $\boldsymbol{\theta}$ in a subset Θ of $\mathbb{R}^m$, and let Θ' be the set of values of $\boldsymbol{\theta}$ for which $f(\boldsymbol{x}, \boldsymbol{\theta})$ is positive a.e. in $\boldsymbol{x}$. We are interested in estimating $\boldsymbol{\theta}$ using a function $\boldsymbol{g}(\boldsymbol{x})$ of the measurements. The error covariance matrix is defined as

$$\boldsymbol{R} \triangleq E\Big\{(\boldsymbol{\theta} - \boldsymbol{g}(\boldsymbol{x}))(\boldsymbol{\theta} - \boldsymbol{g}(\boldsymbol{x}))^T\Big\}. \tag{10.1}$$

The goal is to find a lower bound on $\boldsymbol{R}$, i.e., a matrix $\mathcal{B}$ such that $\boldsymbol{R} \geq \mathcal{B}$, where the matrix inequality means that $\boldsymbol{R} - \mathcal{B}$ is positive semidefinite.

Denote the likelihood ratio by

$$L(\boldsymbol{x}; \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \triangleq \frac{f(\boldsymbol{x}, \boldsymbol{\theta}_1)}{f(\boldsymbol{x}, \boldsymbol{\theta}_2)}, \tag{10.2}$$

where the function is defined only for values of $\boldsymbol{x}$ and $\boldsymbol{\theta}_2$ such that $f(\boldsymbol{x}, \boldsymbol{\theta}_2) \neq 0$. Let

$$\mu(s, \boldsymbol{h}) \triangleq \ln E\{L^s(\boldsymbol{x}; \boldsymbol{\theta} + \boldsymbol{h}, \boldsymbol{\theta})\} \tag{10.3}$$

and note that the expectation is calculated only over points $(\boldsymbol{x}, \boldsymbol{\theta})$ such that $f(\boldsymbol{x}, \boldsymbol{\theta}) > 0$, i.e., those points for which $L(\boldsymbol{x}; \boldsymbol{\theta} + \boldsymbol{h}, \boldsymbol{\theta})$ is defined.

In the case of a scalar parameter $\theta \in \Theta \subseteq \mathbb{R}$, the error covariance (10.1) equals the MSE. The WWB in this setting is given by

$$\text{MSE} \geq \frac{h^2 E^2\{L^s(\boldsymbol{x}; \theta + h, \theta)\}}{E\Big\{\big[L^s(\boldsymbol{x}; \theta + h, \theta) - L^{1-s}(\boldsymbol{x}; \theta - h, \theta)\big]^2\Big\}} \tag{10.4}$$

for any h and s such that

$$0 < E\left\{ \left[L^s(\boldsymbol{x};\theta+h,\theta) - L^{1-s}(\boldsymbol{x};\theta-h,\theta) \right]^2 \right\} < \infty. \tag{10.5}$$

If Θ' is a connected subset of $\mathbb{R}$, then (10.4) can also be written as [60, 61, 174]

$$\text{MSE} \geq \frac{h^2 e^{2\mu(s,h)}}{e^{\mu(2s,h)} + e^{\mu(2-2s,-h)} - 2e^{\mu(s,2h)}}. \tag{10.6}$$

In these cases, calculation of $\mu(s,h)$ is sufficient for evaluation of the bound. However, the equivalence between the original bound (10.4) and the simplified version (10.6) does not necessarily hold if Θ' is a disjoint subset of $\mathbb{R}$; specifically, the cross-term $2e^{\mu(s,2h)}$ does not always equal $2E\{L^s(\boldsymbol{x};\theta+h,\theta)L^{1-s}(\boldsymbol{x};\theta-h,\theta)\}$. As we will see, when Θ' is disjoint, (10.6) can be larger than the minimum estimation MSE, and may even be infinite. Thus, care must be used when applying the bound to disjoint parameter sets Θ'.

In the case of a vector parameter $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^m$, the WWB on the error covariance is given by

$$\boldsymbol{R} \geq \boldsymbol{H}\boldsymbol{G}^{-1}\boldsymbol{H}^T. \tag{10.7}$$

Here, $\boldsymbol{H} = [\boldsymbol{h}_1, \boldsymbol{h}_2, \ldots, \boldsymbol{h}_m]$ is a matrix consisting of "test vectors" $\boldsymbol{h}_i \in \mathbb{R}^m$, and $\boldsymbol{G}$ is the $m \times m$ matrix whose elements are given by

$$G_{ij} = \frac{E\{r(\boldsymbol{x},\boldsymbol{\theta};\boldsymbol{h}_i,s_i)r(\boldsymbol{x},\boldsymbol{\theta};\boldsymbol{h}_j,s_j)\}}{E\{L^{s_i}(\boldsymbol{x};\boldsymbol{\theta}+\boldsymbol{h}_i)\}\,E\{L^{s_j}(\boldsymbol{x};\boldsymbol{\theta}+\boldsymbol{h}_j)\}} \tag{10.8}$$

where

$$r(\boldsymbol{x},\boldsymbol{\theta};\boldsymbol{h}_i,s_i) \triangleq L^{s_i}(\boldsymbol{x};\boldsymbol{\theta}+\boldsymbol{h}_i,\boldsymbol{\theta}) - L^{1-s_i}(\boldsymbol{x};\boldsymbol{\theta}-\boldsymbol{h}_i,\boldsymbol{\theta}). \tag{10.9}$$

The bound holds for any $\{\boldsymbol{h}_i, s_i\}_{i=1}^m$ such that $\boldsymbol{G}$ is well-defined and invertible.

Weinstein and Weiss [60] suggest that attention be restricted to the case $s_1 = \cdots = s_m = 1/2$. In this case, and under the additional assumption that $\Theta' = \mathbb{R}^m$, the bound (10.8) simplifies to[1]

$$G_{ij} = 2\frac{e^{\mu(1/2,\boldsymbol{h}_i-\boldsymbol{h}_j)} - e^{\mu(1/2,\boldsymbol{h}_i+\boldsymbol{h}_j)}}{e^{\mu(1/2,\boldsymbol{h}_i)}e^{\mu(1/2,\boldsymbol{h}_j)}}. \tag{10.10}$$

While this simplification is valid if $f(\boldsymbol{x},\boldsymbol{\theta})$ is positive almost everywhere (i.e., if $\Theta' = \mathbb{R}^m$), it does not necessarily hold in other cases. Furthermore, when $\boldsymbol{\theta}$ is a scalar, (10.10) does not necessarily reduce to (10.6).

[1] A slightly different version of this equation appears in [60, eq. (42)], the result of an obvious typographical error.

10.3 Scalar Case

To demonstrate equivalence between (10.4) and (10.6), it is required to show that

$$E\Big\{L^s(\boldsymbol{x};\theta+h,\theta)L^{1-s}(\boldsymbol{x};\theta-h,\theta)\Big\} \stackrel{?}{=} e^{\mu(s,2h)}. \tag{10.11}$$

When Θ' is a connected set, this statement can be verified as follows.

$$\begin{aligned}
&E\Big\{L^s(\boldsymbol{x};\theta+h,\theta)L^{1-s}(\boldsymbol{x};\theta-h,\theta)\Big\} \\
&= \int \frac{f^s(\boldsymbol{x},\theta+h)}{f^s(\boldsymbol{x},\theta)}\frac{f^{1-s}(\boldsymbol{x},\theta-h)}{f^{1-s}(\boldsymbol{x},\theta)}f(\boldsymbol{x},\theta)d\boldsymbol{x}\,d\theta && (10.12)\\
&= \int f^s(\boldsymbol{x},\theta+h)f^{1-s}(\boldsymbol{x},\theta-h)\,d\boldsymbol{x}\,d\theta && (10.13)\\
&= \int \frac{f^s(\boldsymbol{x},\theta'+2h)}{f^s(\boldsymbol{x},\theta')}f(\boldsymbol{x},\theta')\,d\boldsymbol{x}\,d\theta' && (10.14)\\
&= e^{\mu(s,2h)} && (10.15)
\end{aligned}$$

where we used the change of variables $\theta' = \theta - h$.

As noted previously, it is implicitly assumed in (10.11) that the expectation is calculated only over those values of $\boldsymbol{x}$ and θ for which $f(\boldsymbol{x},\theta) > 0$, otherwise $L(\boldsymbol{x};\theta+h,\theta)$ is undefined. Thus, the integral (10.12) is taken only over those values of $\boldsymbol{x}$ and θ for which $f(\boldsymbol{x},\theta) > 0$. However, after $f(\boldsymbol{x},\theta)$ is canceled out in (10.13), this fact is ignored. The integral (10.14) is taken over the range $f(\boldsymbol{x},\theta') > 0$, which corresponds to $f(\boldsymbol{x},\theta+h) > 0$. If there exist points $(\boldsymbol{x},\theta)$ for which $f(\boldsymbol{x},\theta) = 0$, $f(\boldsymbol{x},\theta+h) > 0$, and $f(\boldsymbol{x},\theta-h) > 0$, then (10.12) is not taken over those points, whereas (10.13) is positive at those points.

If Θ' is a connected set (i.e., a finite or infinite interval), then, for any $\boldsymbol{x}$ and θ such that $f(\boldsymbol{x},\theta-h) > 0$ and $f(\boldsymbol{x},\theta+h) > 0$, we also have $f(\boldsymbol{x},\theta) > 0$. In this case, the range of integration in (10.12) equals that of (10.13), so that the simplified version (10.6) is correct. This occurs, for instance, in the example given in [61], where Θ' is a closed interval.

However, if Θ' is disjoint, then (10.13) can be greater than (10.12), so that (10.6) can be larger than (10.4), and is not necessarily a lower bound on the MSE. Indeed, as we will see, in some cases (10.6) is higher than the minimum MSE; in other cases, (10.6) is infinite, as a result of a division by zero.

10.4 Vector Case

When $\Theta' = \mathbb{R}^m$, the simplified equation (10.10) can be derived from the WWB (10.8) as follows. Substituting $s_1 = \cdots = s_m = 1/2$ in (10.8), the denominator equals that of (10.10). The

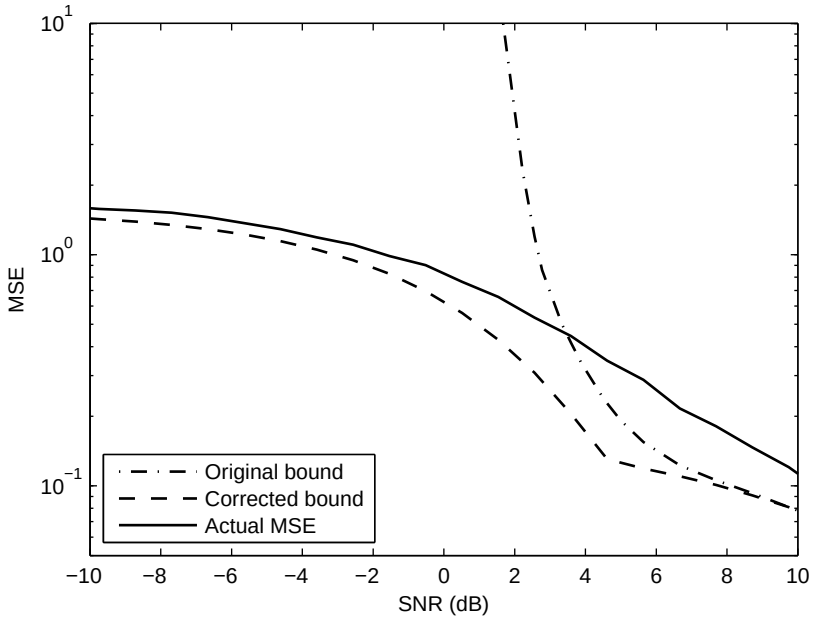

Figure 10.1: Plot of the corrected WWB (10.24), the incorrect version (10.6), and the actual MSE, for a scalar estimation setting with a disjoint set Θ.

numerator consists of a sum of four expressions of the type

$$E\Big\{ L^{1/2}(\boldsymbol{x};\boldsymbol{\theta}\pm\boldsymbol{h}_i,\boldsymbol{\theta})L^{1/2}(\boldsymbol{x};\boldsymbol{\theta}\pm\boldsymbol{h}_j,\boldsymbol{\theta})\Big\}. \tag{10.16}$$

These can be simplified by writing

$$\begin{aligned}
&E\Big\{ L^{1/2}(\boldsymbol{x};\boldsymbol{\theta}+\boldsymbol{h}_i,\boldsymbol{\theta})L^{1/2}(\boldsymbol{x};\boldsymbol{\theta}+\boldsymbol{h}_j,\boldsymbol{\theta})\Big\} \\
&\quad= \int \left(\frac{f(\boldsymbol{x},\boldsymbol{\theta}+\boldsymbol{h}_i)}{f(\boldsymbol{x},\boldsymbol{\theta})} \frac{f(\boldsymbol{x},\boldsymbol{\theta}+\boldsymbol{h}_j)}{f(\boldsymbol{x},\boldsymbol{\theta})} \right)^{1/2} f(\boldsymbol{x},\boldsymbol{\theta})\,d\boldsymbol{x}\,d\boldsymbol{\theta} && (10.17)\\
&\quad= \int \left(f(\boldsymbol{x},\boldsymbol{\theta}+\boldsymbol{h}_i) f(\boldsymbol{x},\boldsymbol{\theta}+\boldsymbol{h}_j) \right)^{1/2} d\boldsymbol{x}\,d\boldsymbol{\theta} && (10.18)\\
&\quad= \int \left(\frac{f(\boldsymbol{x},\boldsymbol{\theta}'+\boldsymbol{h}_i-\boldsymbol{h}_j)}{f(\boldsymbol{x},\boldsymbol{\theta}')} \right)^{1/2} f(\boldsymbol{x},\boldsymbol{\theta}')\,d\boldsymbol{x}\,d\boldsymbol{\theta}' && (10.19)\\
&\quad= e^{\mu(1/2,\boldsymbol{h}_i-\boldsymbol{h}_j)} && (10.20)
\end{aligned}$$

where a change of coordinates $\boldsymbol{\theta}' = \boldsymbol{\theta}+\boldsymbol{h}_j$ was performed. Analogous results can be obtained for the remaining expressions of the type (10.16). Substituting these into (10.8) yields (10.10).

When Θ' does not consist of the entire space $\mathbb{R}^m$, the reasoning above is not valid. The integration in (10.17) must be carried out only over those values of $\boldsymbol{x}$ and $\boldsymbol{\theta}$ for which $f(\boldsymbol{x},\boldsymbol{\theta}) > 0$, but this restriction is dropped in the transition to (10.18). If there exist values $(\boldsymbol{x},\boldsymbol{\theta})$ such that $f(\boldsymbol{x},\boldsymbol{\theta}) = 0$, $f(\boldsymbol{x},\boldsymbol{\theta}+\boldsymbol{h}_i) > 0$, and $f(\boldsymbol{x},\boldsymbol{\theta}+\boldsymbol{h}_j) > 0$, then those values are included in (10.18),

but not in (10.17). This will always occur for some values of $\boldsymbol{h}_i$ and $\boldsymbol{h}_j$, unless $f(\boldsymbol{x},\boldsymbol{\theta}) > 0$ for all $\boldsymbol{\theta}$. Therefore, the value $E\{L^{1/2}(\boldsymbol{x};\boldsymbol{\theta}'+\boldsymbol{h}_i-\boldsymbol{h}_j,\boldsymbol{\theta}')\}$ can, in fact, be larger than $E\{L^{1/2}(\boldsymbol{x};\boldsymbol{\theta}+\boldsymbol{h}_i,\boldsymbol{\theta})L^{1/2}(\boldsymbol{x};\boldsymbol{\theta}+\boldsymbol{h}_j,\boldsymbol{\theta})\}$. We conclude that (10.10) does not necessarily hold unless Θ' includes the entire parameter space $\mathbb{R}^m$.

10.5 Counterexample

We now present an example which illustrates that (10.6) provides incorrect and even impossible results in some cases.

Consider the problem of estimating a scalar θ based on a single measurement x. Suppose that θ is uniformly distributed over the set $\Theta = [a,b] \cup [-b,-a]$, where $b > a > 0$ are given constants. Also suppose that the distribution of x conditioned on θ is Gaussian with mean θ and variance σ^2. The joint pdf of x and θ is thus given by

$$f(x,\theta) = \frac{e^{-(x-\theta)^2/2\sigma^2}}{2(b-a)\sqrt{2\pi\sigma^2}}\mathbb{1}_\Theta \tag{10.21}$$

where $\mathbb{1}_\Theta$ is an indicator function, which equals 1 when $\theta \in \Theta$ and 0 otherwise.

From (10.3),

$$e^{\mu(s,h)} = \frac{1}{2(b-a)}e^{-h^2 s(1-s)/2\sigma^2}\int \mathbb{1}_\Theta \mathbb{1}_{\Theta+h}\,d\theta \tag{10.22}$$

where $\Theta + h = \{\theta + h : \theta \in \Theta\}$, and the integral equals the length of the intersection of the sets Θ and $\Theta + h$. On the other hand,

$$\begin{aligned}\tilde{M}(s,h) &\triangleq E\Big\{L^s(x;\theta+h,\theta)L^{1-s}(x;\theta-h,\theta)\Big\}\\ &= \frac{1}{2(b-a)}e^{-2s(1-s)h^2/\sigma^2}\int \mathbb{1}_\Theta \mathbb{1}_{\Theta+h}\mathbb{1}_{\Theta-h}\,d\theta.\end{aligned} \tag{10.23}$$

Thus, contrary to (10.11), $e^{\mu(s,2h)}$ does not equal $\tilde{M}(s,h)$, since the latter depends on the length of the intersection of the three sets Θ, $\Theta + h$, and $\Theta - h$. Indeed, $\tilde{M}(s,h)$ is often substantially smaller than $e^{\mu(s,2h)}$, and as a result, use of (10.6) results in a "lower bound" which may exceed the true MSE.

This problem is illustrated in Fig. 10.1, where the incorrect bound (10.6) is compared with the original Weiss–Weinstein bound (10.4), which can be written as

$$\frac{h^2 e^{2\mu(s,h)}}{e^{\mu(2s,h)} + e^{\mu(2-2s,-h)} - 2\tilde{M}(s,h)}. \tag{10.24}$$

The actual MSE obtained by the optimal estimator can be calculated using Monte Carlo simulations, and is also plotted. In the figure, values of $a = 1/2$ and $b = 2$ were used. The variance σ^2 was modified to obtain various signal-to-noise ratios (SNRs), where $\mathrm{SNR} = \mathrm{Var}(\theta)/\sigma^2$.

It is evident from Fig. 10.1 that the value (10.6) becomes exceedingly high at low SNR. Indeed, for SNR values below approximately 0 dB, there always exist values of s and h such that the denominator of (10.6) is arbitrarily small, and thus the bound tends to infinity. For SNR values around 2–4 dB, (10.6) yields finite values which are larger than the actual MSE obtained by the optimal estimator. The original version (10.24), by contrast, closely follows the true MSE value.

10.6 Acknowledgement

The authors are grateful to Prof. A. J. Weiss and Prof. E. Weinstein for carefully reviewing an early version of this correspondence.

Chapter 11

Conclusion

The central goal of this thesis is to quantify the performance achieved by practical techniques in structured estimation problems, and to determine whether and to what extent current estimation capabilities are close to the theoretical optimum. This goal is attained by the combined use of lower bounds and performance guarantees, with the former providing a limit on the feasible achievements in a given scenario, and the latter determining whether actual techniques come close to this performance.

Several structured estimation settings were considered in the course of our analysis. First, we examined the setting of sparsely representable signals, wherein the parameter vector is a linear combination of a small number of atoms from a large dictionary. A lower bound for this setting was derived in Chapter 4 by developing a version of the constrained CRB for non-differentiable constraints. This bound was shown to equal the oracle error for almost all parameter values, a fact which gives further credibility to the oracle as a benchmark against which practical estimators can be compared. Such a comparison was performed in Chapter 6, where we demonstrated that several common methods do indeed come within a nearly constant factor of the oracle, thus providing performance guarantees for sparse estimators. Our approach provided tighter bounds than those previously available for ℓ_1 relaxation techniques, as well as the first frequentist performance guarantees for greedy sparse estimation methods. Nevertheless, there is still some gap between these upper and lower bounds, and empirical observations suggest that both can still be improved somewhat.

Bounds and guarantees were also developed for the related model of block sparsity (Chapter 7). In this setting, it was again shown that the CRB coincides with the oracle error for nearly all parameter values. Moreover, performance guarantees for greedy block sparse estimators

were derived and shown to be close to the CRB or oracle bound. Performance guarantees for ℓ_1 relaxation techniques in the frequentist block sparse setting have not been demonstrated, and their existence remains an open question. Comparing the guarantees for scalar and block sparsity, we demonstrated that the block model is particularly advantageous when the coherence between atoms within a single block is low; indeed, highly correlated blocks may cause the performance of block sparse algorithms to fall below standard sparsity approaches. This example illustrates the power of performance guarantees in pinpointing the pros and cons of estimation models.

We next investigated the finite rate of innovation (FRI) model (see Chapter 8). In this setting, one seeks to estimate a continuous-time signal from low-rate samples. This analysis necessitated an extension of the CRB for the estimation of parameters in arbitrary Hilbert spaces, such as the space of finite-energy continuous-time functions. We applied this extended CRB to the FRI setting in two ways. First, a bound was obtained for the ultimate performance achievable in the absence of any restrictions on the allowed sampling technique or rate. Next, a lower bound was derived for a given sampling scheme. Comparing these two bounds with the actual performance achieved by practical estimators can identify both situations in which the estimator fails to utilize all of the information present in the given samples, and cases in which the available information is fully exploited, but better performance could have been achieved by sampling at a higher rate. Examples of state-of-the-art algorithms with both types of failures were identified. Thus, unlike the sparse setting, in the FRI case there are clearly situations in which existing algorithms can be improved considerably.

Interestingly, our analysis helps explain the apparent discrepancy between the successful performance of low-rate FRI techniques in the absence of noise and the empirically observed difficulty in achieving similar results when even low levels of noise are introduced. Indeed, we showed that while some models are amenable to accurate estimation from low-rate noisy samples, other models are necessitate sampling at the Nyquist rate for optimal recovery. Specifically, in the context of the common union of subspaces structure, high sampling rates tend to be necessary, unless the number of parameters determining the subspace under consideration is small relative to the dimension of the subspaces involved.

Finally, we examined performance bounds for Bayesian models in Chapters 9 and 10. In this context, we analyzed the optimal bias bound, a technique utilizing the biased CRB for obtaining performance bounds in the Bayesian setting. We extended this bound to general vector parameters and demonstrated that the bound is tight at both low and high SNR values,

a result which has not been demonstrated, to the best of our knowledge, for any other bound. We also numerically compared this result with other Bayesian bounds, and showed that it tends to be tighter than both the Ziv–Zakai and Weiss–Weinstein techniques.

Bibliography

[1] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA: SIAM, 1992.

[2] S. Mallat, *A Wavelet Tour of Signal Processing – The Sparse Way*, 3rd ed. San Diego, CA: Academic Press, 2009.

[3] J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. New York, NY: Springer-Verlag, 1985.

[4] K. Schnass and P. Vandergheynst, "Average performance analysis for thresholding," *IEEE Signal Process. Lett.*, vol. 14, no. 11, pp. 828–831, Nov. 2007.

[5] Y. C. Eldar and H. Rauhut, "Average case analysis of multichannel sparse recovery using convex relaxation," *IEEE Trans. Inf. Theory*, to appear. [Online]. Available: http://arxiv.org/abs/0904.0494

[6] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, 2nd ed. New York: Springer, 1998.

[7] D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 52, no. 1, pp. 6–18, 2006.

[8] E. Candes and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n," *Ann. Statist.*, vol. 35, no. 6, pp. 2313–2351, 2007, with discussion.

[9] Y. C. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Inf. Theory*, vol. 55, no. 11, pp. 5302–5316, Nov. 2009.

[10] H. Liu and H. Motoda, *Feature selection for knowledge discovery and data mining*. Dordrecht, The Netherlands: Kluwer, 1998.

[11] S. Mallat and Z. Zhang, "Matching pursuit in a time-frequency dictionary," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.

[12] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math.*, vol. LIX, pp. 1207–1223, 2006.

[13] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Compte Rendus de l'Academie des Sciences, Paris, Serie I*, vol. 346, pp. 589–592, 2008. [Online]. Available: http://www.acm.caltech.edu/~emmanuel/papers/RIP.pdf

[14] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.

[15] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Review*, vol. 51, no. 1, pp. 34–81, Feb. 2009.

[16] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ: Prentice Hall, 1993.

[17] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Asilomar Conf. Signals, Systems, and Computers*, Nov. 1993, pp. 40–44.

[18] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Computat. Harmonic Anal.*, vol. 26, no. 3, pp. 301–321, 2009.

[19] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Statist. Soc. B*, vol. 58, no. 1, pp. 267–288, 1996.

[20] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, pp. 33–61, 1998.

[21] E. J. Candès and Y. Plan, "Near-ideal model selection by ℓ_1 minimization," *Ann. Statist.*, vol. 37, no. 5A, pp. 2145–2177, Oct. 2009.

[22] J. A. Tropp, "Just relax: Convex programming methods for identifying sparse signals in noise," *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 1030–1051, 2006.

[23] P. J. Bickel, Y. Ritov, and A. Tsybakov, "Simultaneous analysis of Lasso and Dantzig selector," *Ann. Statist.*, vol. 37, no. 4, pp. 1705–1732, 2009.

[24] E. J. Candès, "Modern statistical estimation via oracle inequalities," *Acta Numerica*, pp. 1–69, 2006.

[25] E. J. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

[26] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. Royal Statist. Soc. B*, vol. 68, pp. 49–67, 2006.

[27] M. Stojnic, F. Parvaresh, and B. Hassibi, "On the reconstruction of block-sparse signals with an optimal number of measurements," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 3075–3085, Aug. 2009.

[28] Y. C. Eldar, P. Kuppinger, and H. Bölcskei, "Block-sparse signals: Uncertainty relations and efficient recovery," *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 3042–3054, Jun. 2010.

[29] Z. Ben-Haim and Y. C. Eldar, "Near-oracle performance of greedy block-sparse estimation techniques from noisy measurements," *IEEE J. Sel. Topics Signal Process.*, 2010, submitted. [Online]. Available: http://arxiv.org/abs/1009.0906

[30] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1417–1428, 2002.

[31] I. Maravic and M. Vetterli, "Sampling and reconstruction of signals with finite rate of innovation in the presence of noise," *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 2788–2805, Aug. 2005.

[32] C. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379423, 623656, 1948.

[33] M. Unser, "Splines: A perfect fit for signal/image processing," *IEEE Signal Process. Mag.*, vol. 16, no. 6, pp. 22–38, Nov. 1999.

[34] R. Tur, Y. C. Eldar, and Z. Friedman, "Low rate sampling of pulse streams with application to ultrasound imaging," *IEEE Trans. Signal Process.*, to appear. [Online]. Available: http://arxiv.org/abs/1003.2822

[35] W. U. Bajwa, K. Gedalyahu, and Y. C. Eldar, "Identification of underspread linear systems with application to super-resolution radar," *IEEE Trans. Signal Process.*, to appear. [Online]. Available: http://arxiv.org/abs/1008.0851

[36] K. Gedalyahu and Y. C. Eldar, "Time delay estimation from low rate samples: A union of subspaces approach," *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 3017–3031, Jun. 2010.

[37] P. L. Dragotti, M. Vetterli, and T. Blu, "Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets Strang–Fix," *IEEE Trans. Signal Process.*, vol. 55, no. 7, pp. 1741–1757, 2007.

[38] J. Berent, P. L. Dragotti, and T. Blu, "Sampling piecewise sinusoidal signals with finite rate of innovation methods," *IEEE Trans. Signal Process.*, vol. 58, no. 2, pp. 613–625, Feb. 2010.

[39] T. Y. Young and R. A. Westerberg, "Error bounds for stochastic estimation of signal parameters," *IEEE Trans. Inf. Theory*, vol. 17, no. 5, pp. 549–557, Sep. 1971.

[40] Y. C. Eldar, A. Ben-Tal, and A. Nemirovski, "Robust mean-squared error estimation in the presence of model uncertainties," *IEEE Trans. Signal Process.*, vol. 53, no. 1, pp. 168–181, Jan. 2005.

[41] ——, "Linear minimax regret estimation of deterministic parameters with bounded data uncertainties," *IEEE Trans. Signal Process.*, vol. 52, no. 8, pp. 2177–2188, Aug. 2004.

[42] Z. Ben-Haim and Y. C. Eldar, "Blind minimax estimation," *IEEE Trans. Inf. Theory*, vol. 53, no. 9, pp. 3145–3157, Sep. 2007.

[43] W. James and C. Stein, "Estimation with quadratic loss," in *Proc. Fourth Berkeley Symp. Math. Statist. Prob.*, vol. 1, 1961, pp. 361–379. [Online]. Available: http://projecteuclid.org/euclid.bsmsp/1200512173

[44] Y. C. Eldar, "Rethinking biased estimation: Improving maximum likelihood and the Cramér–Rao bound," *Foundations and Trends in Signal Processing*, vol. 1, no. 4, pp. 305–449, 2008.

[45] H. Cramér, "A contribution to the theory of statistical estimation," *Skand. Akt. Tidskr.*, vol. 29, pp. 85–94, 1945.

[46] G. Darmois, "Sur les limites de la dispersion de certains estimations," *Rev. Inst. Int. Statist.*, vol. 13, pp. 9–15, 1945.

[47] M. Fréchet, "Sur l'extension de certaines evaluations statistiques au cas de petits echantillons," *Rev. Inst. Int. Statist.*, vol. 11, pp. 182–205, 1943.

[48] C. R. Rao, "Information and accuracy attainable in the estimation of statistical parameters," *Bull. Calcutta Math. Soc.*, vol. 37, pp. 81–91, 1945.

[49] P. Stoica and T. L. Marzetta, "Parameter estimation problems with singular information matrices," *IEEE Trans. Signal Process.*, vol. 49, no. 1, pp. 87–90, Jan. 2001.

[50] J. M. Hammersley, "On estimating restricted parameters," *J. Roy. Statist. Soc. B*, vol. 12, no. 2, pp. 192–240, 1950.

[51] D. G. Chapman and H. Robbins, "Minimum variance estimation without regularity assumptions," *Ann. Math. Statist.*, vol. 22, no. 4, pp. 581–586, Dec. 1951.

[52] J. D. Gorman and A. O. Hero, "Lower bounds for parametric estimation with constraints," *IEEE Trans. Inf. Theory*, vol. 26, no. 6, pp. 1285–1301, Nov. 1990.

[53] E. W. Barankin, "Locally best unbiased estimates," *Ann. Math. Statist.*, vol. 20, no. 4, pp. 477–501, Dec. 1949.

[54] P. K. Bhattacharya, "Estimating the mean of a multivariate normal population with general quadratic loss function," *Ann. Math. Statist.*, vol. 37, no. 6, pp. 1819–1824, Dec. 1966.

[55] T. L. Marzetta, "A simple derivation of the constrained multiple parameter Cramér–Rao bound," *IEEE Trans. Signal Process.*, vol. 41, no. 6, pp. 2247–2249, Jun. 1993.

[56] P. Stoica and B. C. Ng, "On the Cramér–Rao bound under parametric constraints," *IEEE Signal Process. Lett.*, vol. 5, no. 7, pp. 177–179, 1998.

[57] T. J. Moore, R. J. Kozick, and B. M. Sadler, "The constrained Cramér–Rao bound from the perspective of fitting a model," *IEEE Signal Process. Lett.*, vol. 14, no. 8, pp. 564–567, Aug. 2007.

[58] Z. Ben-Haim and Y. C. Eldar, "On the constrained Cramér-Rao bound with a singular Fisher information matrix," *IEEE Signal Process. Lett.*, vol. 16, no. 6, pp. 453–456, Jun. 2009.

[59] H. L. Van Trees, *Detection, Estimation, and Modulation Theory*. New York: Wiley, 1968, vol. 1.

[60] E. Weinstein and A. J. Weiss, "A general class of lower bounds in parameter estimation," *IEEE Trans. Inf. Theory*, vol. 34, no. 2, pp. 338–342, Mar. 1988.

[61] A. J. Weiss and E. Weinstein, "A lower bound on the mean-square error in random parameter estimation," *IEEE Trans. Inf. Theory*, vol. 31, no. 5, pp. 680–682, Sep. 1985.

[62] K. L. Bell and H. L. Van Trees, *Bayesian bounds for parameter estimation and nonlinear filtering / tracking*. IEEE, 2007.

[63] K. Todros and J. Tabrikian, "A new Bayesian lower bound on the mean square error of estimators," in *Proc. 16th European Signal Processing Conference (EUSIPCO 2008)*, Lausanne, Switzerland, Aug. 2008.

[64] K. Todros and J. Tabirikian, "General classes of performance lower bounds for parameter estimation—part ii: Bayesian bounds," *IEEE Trans. Inf. Theory*, vol. 56, no. 10, pp. 5064–5082, Oct. 2010.

[65] J. Ziv and M. Zakai, "Some lower bounds on signal parameter estimation," *IEEE Trans. Inf. Theory*, vol. 15, no. 3, pp. 386–391, May 1969.

[66] D. Chazan, M. Zakai, and J. Ziv, "Improved lower bounds on signal parameter estimation," *IEEE Trans. Inf. Theory*, vol. 21, no. 1, pp. 90–93, 1975.

[67] S. Bellini and G. Tartara, "Bounds on error in signal parameter estimation," *IEEE Trans. Commun.*, vol. 22, no. 3, pp. 340–342, 1974.

[68] K. L. Bell, Y. Steinberg, Y. Ephraim, and H. L. Van Trees, "Extended Ziv–Zakai lower bound for vector parameter estimation," *IEEE Trans. Inf. Theory*, vol. 43, no. 2, pp. 624–637, 1997.

[69] J. Shao, *Mathematical Statistics*, 2nd ed. New York: Springer, 2003.

[70] Z. Ben-Haim and Y. C. Eldar, "The Cramér–Rao bound for estimating a sparse parameter vector," *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 3384–3389, Jun. 2010.

[71] S. L. Campbell and C. D. Meyer, Jr., *Generalized Inverses of Linear Transformations*. London: Pitman, 1979.

[72] K. B. Petersen and M. S. Pedersen. (2008, Feb.) The matrix cookbook. [Online]. Available: http://matrixcookbook.com/

[73] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, Dec. 2006.

[74] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.

[75] ——, "Image restoration by sparse 3D transform-domain collaborative filtering," in *Proc. SPIE Electronic Imaging '08*, no. 6812-07, San Jose, CA, Jan. 2008.

[76] M. Protter and M. Elad, "Image sequence denoising via sparse and redundant representations," *IEEE Trans. Image Process.*, vol. 18, no. 1, pp. 27–36, Jan. 2009.

[77] M. Elad, J.-L. Starck, P. Querre, and D. Donoho, "Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA)," *J. Applied and Computational Harmonic Analysis*, vol. 19, pp. 340–358, Nov. 2005.

[78] Z. Ben-Haim, Y. C. Eldar, and M. Elad, "Coherence-based performance guarantees for estimating a sparse vector under random noise," *IEEE Trans. Signal Process.*, vol. 58, no. 10, pp. 5030–5043, Oct. 2010.

[79] Z. Ben-Haim and Y. C. Eldar, "Performance bounds for sparse estimation with random noise," in *Proc. IEEE Workshop on Statistical Signal Processing*, Cardiff, Wales, UK, Sep. 2009.

[80] B. Babadi, N. Kalouptsidis, and V. Tarokh, "Asymptotic achievability of the Cramér–Rao bound for noisy compressive sampling," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 1233–1236, 2009.

[81] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Computing*, vol. 24, no. 2, pp. 227–234, 1995.

[82] D. L. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization," *Proc. Nat. Acad. Sci. USA*, vol. 100, no. 5, pp. 2197–2202, March 4, 2003.

[83] Y. C. Eldar, "Compressed sensing of analog signals in shift-invariant spaces," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 2986–2997, Aug. 2009.

[84] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[85] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.

[86] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, pp. 425–455, 1994.

[87] E. G. Larsson and Y. Selen, "Linear regression with a sparse parameter vector," *IEEE Trans. Signal Process.*, vol. 55, no. 2, pp. 451–460, Feb. 2007.

[88] M. E. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. Machine Learn. Res.*, vol. 1, pp. 211–244, 2001.

[89] P. Schniter, L. C. Potter, and J. Ziniel, "Fast Bayesian matching pursuit," in *Proc. Inf. Theory Appl. Workshop*, San Diego, CA, Jan./Feb. 2008, pp. 326–333.

[90] L. L. Scharf, *Statistical Signal Processing*. Reading, MA: Addison-Wesley, 1991.

[91] C. Stein, "Inadmissibility of the usual estimator for the mean of a multivariate distribution," in *Proc. Third Berkeley Symp. Math. Statist. Prob.*, vol. 1, 1956, pp. 197–206. [Online]. Available: http://projecteuclid.org/euclid.bsmsp/1200501656

[92] C. Carbonelli, S. Vedantam, and U. Mitra, "Sparse channel estimation with zero tap detection," vol. 6, no. 5, pp. 1743–1763, May 2007.

[93] M. Dong and L. Tong, "Optimal design and placement of pilot symbols for channel estimation," *IEEE Trans. Signal Process.*, vol. 50, no. 12, pp. 3055–3069, Dec. 2002.

[94] A. Jung, Z. Ben-Haim, F. Hlawatsch, and Y. C. Eldar, "On unbiased estimation of sparse vectors corrupted by Gaussian noise," in *Proc. Int. Conf. Acoustics, Speech, and Signal Processing (ICASSP 2010)*, Dallas, TX, Mar. 2010, pp. 3990–3993.

[95] D. L. Donoho and I. M. Johnstone, "Minimax risk over ℓ_p balls for ℓ_q-error," *Probab. Theory Relat. Fields*, vol. 99, pp. 277–303, 1994.

[96] D. Donoho and I. M. Johnstone, "Minimax estimation via wavelet shrinkage," *Ann. Statist.*, vol. 26, pp. 879–921, 1992.

[97] I. M. Johnstone and B. W. Silverman, "Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences," *Ann. Statist.*, vol. 32, pp. 1594–1649, 2004.

[98] S. Boyd and L. Vanderberghe, *Convex Optimization*. Cambridge, UK: Cambridge Univ. Press, 2004. [Online]. Available: http://www.stanford.edu/~boyd/cvxbook.html

[99] J. S. Abel, "A bound on mean-square-estimate error," *IEEE Trans. Inf. Theory*, vol. 39, no. 5, pp. 1675–1680, 1993.

[100] P. Forster and P. Larzabal, "On lower bounds for deterministic parameter estimation," in *Proc. Int. Conf. Acoust., Speech and Signal Processing (ICASSP 2002)*, vol. 2, Orlando, FL, May 2002, pp. 1137–1140.

[101] K. Todros and J. Tabrikian, "A new lower bound on the mean-square error of unbiased estimators," in *Proc. 2008 IEEE Int. Conf. on Acoust., Speech, and Signal Processing (ICASSP 2008)*, Las Vegas, NV, 2008, pp. 3913–3916.

[102] R. Strichartz, *A Guide to Distribution Theory and Fourier Transforms*. Boca Raton, FL: CRC Press, 1993.

[103] D. S. Bernstein, *Matrix Mathematics*. Princeton, NJ: Princeton Univ. Press, 2005.

[104] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd ed. New York: McGraw-Hill, 1991.

[105] J. J. Fuchs, "Recovery of exact sparse representations in the presence of bounded noise," *IEEE Trans. Inf. Theory*, vol. 51, no. 10, pp. 3601–3608, 2005.

[106] D. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.

[107] T. Cai, L. Wang, and G. Xu, "Stable recovery of sparse signals and an oracle inequality," U. Penn., Tech. Rep., 2009. [Online]. Available: http://www-stat.wharton.upenn.edu/~tcai/paper/Stable-Recovery-MIP.pdf

[108] Y. C. Eldar, "Uncertainty relations for shift-invariant analog signals," *IEEE Trans. Inf. Theory*, vol. 55, no. 12, pp. 5742–5757, Dec. 2009.

[109] Z. Ben-Haim, Y. C. Eldar, and M. Elad, "Coherence-based near-oracle performance guarantees for sparse estimation under Gaussian noise," in *Proc. Int. Conf. Acoustics, Speech, and Signal Processing (ICASSP 2010)*, Dallas, TX, Mar. 2010, pp. 3590–3593.

[110] J. A. Tropp, "On the conditioning of random subdictionaries," *Applied and Computational Harmonic Analysis*, vol. 25, no. 1, pp. 1–24, 2008.

[111] J. S. Turek, I. Yavneh, M. Protter, and M. Elad, "On MMSE and MAP denoising under sparse representation modeling over a unitary dictionary," *Applied and Computational Harmonic Analysis*, 2010, submitted.

[112] M. J. Wainwright, "Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ_1-constrained quadratic programming (Lasso)," *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2183–2202, May 2009.

[113] A. K. Fletcher and S. Rangan, "Orthogonal matching pursuit from noisy measurements: A new analysis," in *Proc. Neural Information Processing Systems Conf. (NIPS)*, Vancouver, Canada, Dec. 2009.

[114] T. T. Cai, G. Xu, and J. Zhang, "On recovery of sparse signals via ℓ_1 minimization," *IEEE Trans. Inf. Theory*, vol. 55, no. 7, pp. 3388–3397, Jul. 2009.

[115] M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. New York: Springer, 2010, in print.

[116] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Baltimore, MD: Johns Hopkins University Press, 1996.

[117] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.

[118] T. Strohmer and R. W. Heath, Jr., "Grassmannian frames with applications to coding and communication," *Applied and Computational Harmonic Analysis*, vol. 14, no. 3, pp. 257–275, May 2003.

[119] A. Stuart and J. K. Ord, *Kendall's Advanced Theory of Statistics*, 6th ed. London: Edward Arnold, 1994, vol. 1.

[120] Z. Šidák, "Rectangular confidence regions for the means of multivariate normal distributions," *J. Amer. Statist. Assoc.*, vol. 62, no. 318, pp. 626–633, Jun. 1967.

[121] M. M. Bronstein, A. M. Bronstein, M. Zibulevsky, and Y. Y. Zeevi, "Blind deconvolution of images using optimal sparse representations," *IEEE Trans. Image Process.*, vol. 14, no. 6, pp. 726–736, Jun. 2005.

[122] M. J. Fadili, J.-L. Starck, and F. Murtagh, "Inpainting and zooming using sparse representations," *The Computer Journal*, vol. 52, no. 1, pp. 64–79, 2009.

[123] Y. M. Lu and M. N. Do, "A theory for sampling signals from a union of subspaces," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2334–2345, Jun. 2008.

[124] T. Blumensath and M. E. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *IEEE Trans. Inf. Theory*, vol. 55, no. 4, pp. 1872–1882, Apr. 2009.

[125] K. Gedalyahu, R. Tur, and Y. C. Eldar, "Multichannel sampling of pulse streams at the rate of innovation," *IEEE Trans. Signal Process.*, submitted. [Online]. Available: http://arxiv.org/abs/1004.5070

[126] M. Mishali and Y. C. Eldar, "From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 375–391, Apr. 2010.

[127] ——, "Blind multi-band signal reconstruction: Compressed sensing for analog signals," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 993–1009, Mar. 2009.

[128] M. Mishali, Y. C. Eldar, O. Dounaevsky, and E. Shoshan, "Xampling: Analog to digital at sub-Nyquist rates," *IET J. Circuits, Devices and Systems*, 2010, to appear. [Online]. Available: http://arxiv.org/abs/0912.2495

[129] D. G. Feingold and R. S. Varga, "Block diagonally dominant matrices and generalizations of the Gerschgorin circle theorem," *Pacific J. Math.*, vol. 12, no. 4, pp. 1241–1250, 1962.

[130] T. W. Anderson, "The integral of a symmetric unimodal function over a symmetric convex set and some probability inequalities," *Proc. Am. Math. Soc.*, vol. 6, pp. 170–176, 1955.

[131] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover, 1964.

[132] P. Stoica and A. Nehorai, "MUSIC, maximum likelihood, and Cramer-Rao bound," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 5, pp. 720–741, 1989.

[133] H.-H. Kuo, "White noise theory," in *Handbook of Stochastic Analysis and Applications*, D. Kannan and V. Lakshmikantham, Eds. New York: Marcel Dekker, 2002, ch. 3.

[134] I. A. Ibragimov and R. Z. Has'minskii, *Statistical Estimation: Asymptotic Theory*. New York: Springer, 1981.

[135] M. Unser, "Sampling—50 years after Shannon," *Proc. IEEE*, vol. 88, pp. 569–587, Apr. 2000.

[136] Y. C. Eldar and T. Michaeli, "Beyond bandlimited sampling," *IEEE Signal Process. Mag.*, vol. 26, no. 3, pp. 48–68, May 2009.

[137] M. Unser and A. Aldroubi, "A general sampling theory for nonideal acquisition devices," *IEEE Trans. Signal Process.*, vol. 42, no. 11, pp. 2915–2925, Nov. 1994.

[138] Y. C. Eldar and T. G. Dvorkind, "A minimum squared-error framework for generalized sampling," *IEEE Trans. Signal Process.*, vol. 54, no. 6, pp. 2155–2167, Jun. 2006.

[139] M. Unser and T. Blu, "Generalized smoothing splines and the optimal discretization of the Wiener filter," *IEEE Trans. Signal Process.*, vol. 53, no. 6, pp. 2146–2159, 2005.

[140] Y. C. Eldar and M. Unser, "Nonideal sampling and interpolation from noisy observations in shift-invariant spaces," *IEEE Trans. Signal Process.*, vol. 54, no. 7, pp. 2636–2651, Jul. 2006.

[141] S. Ramani, D. Van De Ville, T. Blu, and M. Unser, "Nonideal Sampling and Regularization Theory," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 1055–1070, 2008.

[142] H. J. Landau and W. L. Miranker, "The recovery of distorted band-limited signals," *Journal of Math. Anal. and Appl.*, vol. 2, pp. 97–104, Feb. 1961.

[143] I. W. Sandberg, "On the properties of some systems that distort signals I," *Bell Syst. Tech. J.*, vol. 42, pp. 2033–2047, September 1963.

[144] T. G. Dvorkind, Y. C. Eldar, and E. Matusiak, "Nonlinear and nonideal sampling: theory and methods," *IEEE Trans. Signal Process.*, vol. 56, no. 12, pp. 5874–5890, Dec. 2008.

[145] T. Faktor, T. Michaeli, and Y. C. Eldar, "Nonlinear and nonideal sampling revisited," *IEEE Signal Process. Lett.*, vol. 17, no. 2, pp. 205–208, February 2010.

[146] J. G. Proakis, *Digital Communications*, 3rd ed. New York: McGraw-Hill, 1995.

[147] A. Bruckstein, T. J. Shan, and T. Kailath, "The resolution of overlapping echos," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 6, pp. 1357–1367, 1985.

[148] Y. C. Eldar, "Sampling and reconstruction in arbitrary spaces and oblique dual frame vectors," *J. Fourier Analys.*, vol. 1, no. 9, pp. 77–96, Jan. 2003.

[149] M. Loève, "Fonctions aléatoires du second ordre," in *Processes Stochastiques et Mouvement Brownien*, P. Lévy, Ed. Paris: Gauthier-Villars, 1948.

[150] ——, *Probability Theory*. Springer, 1978, vol. II.

[151] O. P. Le Maître and O. M. Knio, *Spectral Methods for Uncertainty Quantification*. Springer, 2010.

[152] Y. Yamashita and H. Ogawa, "Relative Karhunen–Loève transform," *IEEE Trans. Signal Process.*, vol. 44, no. 2, pp. 371–378, 1996.

[153] Y. Hua and W. Liu, "Generalized Karhunen–Loève transform," *IEEE Signal Process. Lett.*, vol. 5, no. 6, pp. 141–142, 1998.

[154] Y. Hua and T. K. Sarpar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 5, pp. 814–824, May 1990.

[155] B. Z. Bobrovski and M. Zakai, "A lower bound on the estimation error for certain diffusion problems," *IEEE Trans. Inf. Theory*, vol. 22, no. 1, pp. 45–52, Jan. 1976.

[156] A. Renaux, P. Forster, P. Larzabal, and C. Richmond, "The Bayesian Abel bound on the mean square error," in *Proc. Int. Conf. Acoust., Speech and Signal Processing (ICASSP 2006)*, vol. III, Toulouse, France, May 2006, pp. 9–12.

[157] S. M. Kay and Y. C. Eldar, "Rethinking biased estimation," *IEEE Signal Process. Mag.*, vol. 25, no. 3, pp. 133–136, May 2008.

[158] A. O. Hero, J. A. Fessler, and M. Usman, "Exploring estimator bias-variance tradeoffs using the uniform CR bound," *IEEE Trans. Signal Process.*, vol. 44, no. 8, pp. 2026–2041, 1996.

[159] Y. C. Eldar, "Minimum variance in biased estimation: Bounds and asymptotically optimal estimators," *IEEE Trans. Signal Process.*, vol. 52, no. 7, pp. 1915–1930, 2004.

[160] ——, "Uniformly improving the Cramér-Rao bound and maximum-likelihood estimation," *IEEE Trans. Signal Process.*, vol. 54, no. 8, pp. 2943–2956, 2006.

[161] ——, "MSE bounds with affine bias dominating the Cramér–Rao bound," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3824–3836, Aug. 2008.

[162] A. Renaux, "Contribution à l'analyse des performances d'estimation en traitement statistique du signal," Ph.D. dissertation, Ècole Normale Superieure de Cachan, 2006. [Online]. Available: http://tel.archives-ouvertes.fr/tel-00129527/

[163] Z. Ben-Haim and Y. C. Eldar, "A Bayesian estimation bound based on the optimal bias function," in *Proc. 2nd Int. Workshop on Computational Adv. in Multi-Sensor Adapt. Process. (CAMSAP 2007)*, St. Thomas, U.S. Virgin Islands, Dec. 2007.

[164] E. H. Lieb and M. Loss, *Analysis*, 2nd ed. American Mathematical Society, 2001.

[165] D. R. Cox and N. Reid, "Parameter orthogonality and approximate conditional inference," *J. Roy. Statist. Soc. B*, vol. 49, no. 1, pp. 1–39, 1987.

[166] H. L. Van Trees and K. L. Bell, *Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking*. New York: Wiley, 2007.

[167] I. M. Vinogradov, Ed., *Encyclopaedia of Mathematics*. Dordrecht, The Netherlands: Kluwer, 1995.

[168] Z. Ben-Haim and Y. C. Eldar, "A comment on the use of the Weiss–Weinstein bound with constrained parameter sets," *IEEE Trans. Inf. Theory*, vol. 54, no. 10, pp. 4682–4684, Oct. 2008.

[169] L. P. Lebedev and M. J. Cloud, *The Calculus of Variations and Functional Analysis*. New Jersey: World Scientific, 2003.

[170] W. Rudin, *Functional Analysis*. New York: McGraw-Hill, 1973.

[171] I. M. Gelfand and S. V. Fomin, *Calculus of Variations*. Mineola, NY: Dover, 2000.

[172] T. J. Nohara and S. Haykin, "Application of the Weiss–Weinstein bound to a two-dimensional antenna array," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 9, pp. 1533–1534, Sep. 1988.

[173] D. F. DeLong, "Use of the Weiss–Weinstein bound to compare the direction-finding performance of sparse arrays," Lincoln Lab, Mass. Inst. Tech., Cambridge, MA, Tech. Rep. 982, 1993.

[174] K. L. Bell, "Performance bounds in parameter estimation with application to bearing estimation," Ph.D. dissertation, George Mason University, 1995.

אלגוריתמים פשוטים למדי. מסתבר שמצב דומה מתרחש גם כאשר אנו עוברים למודל מעט יותר מורכב של דלילות בבלוקים. במודל זה ההנחה היא כי מספר מועט של אטומים מייצגים את הסיגנל, ויתר על כן, אטומים אלה מופיעים בקבוצות בעלות מאפיינים דומים. חסמים תחתונים ועליונים תחת מודל זה מוצגים בפרק 7, וגם כאן אנו מראים כי החסמים קרובים זה לזה, משמע, אלגוריתמים פרקטיים יכולים להגיע לביצועים המתקרבים לאופטימום.

בפרק 8 אנו חוקרים את הביצועים של אותות FRI תחת רעש. לשם כך אנו מפתחים שני סוגי חסמים תחתונים: חסם יסודי, שאינו רגיש למערכת הדגימה של האות, וחסם שמתאר את הביצועים שניתן להשיג באמצעות מערכת דגימה נתונה. השוואת החסמים האלה עם הביצועים של אלגוריתמים פרקטיים מאפשרת לענות על שתי שאלות: ראשית, האם משערך נתון מנצל באופן מלא את המידע הניתן לו על ידי מערכת הדגימה? שנית, האם ניתן להשיג ביצועים טובים עוד יותר על ידי שדרוג החומרה? אנו מראים כי התשובות לשתי השאלות משתנות בהתאם לאופי הבעיה, ומציגים כללי אצבע לזיהוי מצבים שבהם אלגוריתמים קיימים אינם מנצלים באופן מיטבי את המידע הנתון להם, וניתנים לפיכך לשיפור.

החלק השני של העבודה עוסק בחסמי ביצועים בעולם הבייסיאני, שבהם האות המשוערך הוא בעל התפלגות א-פריורי ידועה. כידוע, במצב הבייסיאני, משערך התוחלת המותנית משיג את השגיאה הריבועית הנמוכה ביותר האפשרית. אך במקרים רבים, סיבוכיות החישוב של התוחלת המותנית גבוהה, ולכן מקובל להשתמש במשערכים תת-אופטימליים, כגון משערך ההסתברות הפוסטריורית המירבית (MAP). במצב זה יש עניין לבדוק מהם הביצועים האופטימליים האפשריים, כדי שיהיה ניתן לדעת האם משערך ה-MAP מתקרב אליהם. ואולם, גם חישוב השגיאה של משערך התוחלת המותנית עלול להיות קשה חישובית. כתוצאה מכך התפתח ענף מחקר המנסה לקרב את השגיאה האופטימלית באמצעות חסמים קלים לחישוב.

במסגרת תחום מחקר זה אנו נתאר שתי תוצאות. ראשית, בפרק 9, אנו מציגים את חסם ההטיה האופטימלית. זהו חסם שפותח לראשונה בשנת 1971, אך כמעט ולא זכה להתייחסות. אנו מרחיבים את החסם למקרים כלליים יותר, וכן מוכיחים כי הוא הדוק הן ב-SNR נמוך והן ב-SNR גבוה, תוצאה שלא הודגמה עבור חסמים אחרים. בנוסף, בפרק 10 אנו סוקרים את חסם וייס-ויינשטיין (Weiss-Weinstein) המוכר, ומראים כי צורת ההצגה הנפוצה שלו אינה מדויקת במצבים מסוימים, בפרט כאשר תמך ההסתברות א-פריורי של האות מוגבלת לקבוצה חסומה בתוך מרחב האותות האפשריים. אנו מציעים תיקון לחסם ומדגימים מצבים שבהם הגרסה המקורית מספקת תוצאות שגויות.

פרקטי. כאשר ישנו פער בין החסמים, הדבר יכול להצביע על חוסר דיוק באחד החסמים, או שהוא עשוי להיות תוצאה של פגם באלגוריתם. מתוך ניתוח של המצבים בהם פערים כאלה נוצרים, נוכל לדלות רמזים לגבי הדרכים בהן ניתן לשפר הן את המשערך והן את החסם.

במהלך העבודה, נחקור מספר מודלים פרקטיים וניישם את החסמים התיאורטיים בהן. המודלים בהם נתרכז יכללו את הייצוגים הדלילים שהוזכרו לעיל, וכן הרחבה של מודל הדלילות למצב של דלילות בבלוקים (block sparsity). כמו כן נתעניין במודל של אותות בעלי קצב חידוש סופי (finite rate of innovation או FRI). מעניין לציין כי רבים מהמודלים האלה מגדירים מבנה המהווה איחוד של תתי מרחבים בתוך מרחב גדול יותר של אותות אפשריים. מבנה זה יספק אינטואיציה גיאומטרית ויסייע, במקרים רבים, בבניית החסמים.

המבנה של עבודה זו הוא כדלקמן. פרק 1 מציג את המודלים הפרקטיים שהוזכרו זה עתה, וסוקר בקצרה את הספרות הענפה העוסקת בהם. אחריו, פרק 2 סוקר את יסודות התחום התיאורטי של חסמי ביצועים, הן בשערוך דטרמיניסטי והן בשערוך בייסיאני. חסמים אלה, ובמיוחד הגרסאות השונות של חסם קרמר-ראו (Cramér-Rao), יהוו את הבסיס התיאורטי לפיתוחים שבכל העבודה.

יתרת העבודה מחולקת לשני חלקים, המתייחסים בהתאמה לשערוך דטרמיניסטי עם אילוצים מבניים, ולשערוך ביסייאני. כך, בחלק הראשון והארוך יותר, ההנחה היא כי לא ידועה הסתברות א-פריורי של הסיגנל המשוערך, אך ידוע כי הוא משתייך לקבוצה נתונה המגדירה את מבנה הסיגנל הצפוי. החלק מתחיל בפרק 3, אשר מתאר הרחבה של חסם קרמר-ראו למצב שבו מטריצת האינפורמציה של פישר היא סינגולרית. מצב זה מתעורר, למשל, בבעיות שערוך יתירות.

בהמשך, פרק 4 מציג פיתוח של חסם קרמר-ראו עבור בעיות שערוך דלילות ויתירות. חסם קרמר-ראו המתקבל כאן מפתיע במקצת, מכיוון שהוא מתלכד, כמעט בכל מקום, עם שגיאת האורקל. תוצאה זו נותנת משנה תוקף לנוהג הקיים של השוואת ביצועים של אלגוריתמי שערוך עם האורקל. זאת, למשל, מכיוון שידוע שחסם קרמר-ראו מושג על ידי משערך הנראות המרבית כאשר יחס האות לרעש (SNR) הוא גבוה, ולכן לפחות במצב זה ניתן להגיע לביצועים קרובים לאלה של האורקל.

על אף היותו של חסם קרמר-ראו הדוק כאשר ה-SNR גבוה, במקרים רבים ניתן לשפר את החסם ב-SNR נמוך. זו מטרתו של פרק 5, בו מוצג חסם Hammersley-Chapman-Robbins לבעיות שערוך דלילות. חסם זה הוא מורכב יותר אנליטית, אך במצבים מסוימים ניתן עדיין להשיג ביטוי סגור עבורו, וללמוד ממנו על הביצועים של משערכים בתחומי SNR נמוכים יותר.

לאחר פיתוח החסמים התחתונים, בפרק 6 אנו עוברים לחקר הבטחות ביצועים בבעיית הייצוגים הדלילים. נושא זה נחקר רבות בשנים האחרונות, וקיימים מספר אלגוריתמים עבורם ניתן לספק חסמי ביצועים עליונים תחת תנאים סטטיסטיים מתאימים. תרומתנו בפרק זה היא במציאת הבטחות ביצועים המבוססות על מדדים פשוטים כגון הקוהרנטיות, אשר ניתנים לחישוב באופן יעיל. במקרים רבים הבטחות הביצועים המתקבלות הדוקות יותר מאלה של חסמים קודמים.

השילוב של הבטחות הביצועים של פרק 6 עם החסמים התחתונים של פרקים 4 ו-5 מראה כי ניתן להגיע קרוב למדי לביצועים אופטימליים, בעזרת משערכים קיימים, ולעיתים אף בעזרת

תקציר

אתגר מרכזי בכל תחומי ההנדסה הוא מציאת מודלים נוחים מתמטית המתארים במדויק תופעות פיסיקליות. רבות מבין ההצלחות החשובות בתחום עיבוד האותות בעשורים האחרונים נזקפות לזכות פיתוחי טכניקות מתמטיות חדשות שהיטיבו לתאר סוג מסוים של סיגנלים. דוגמה קלאסית לכך הינה תורת הגלונים (wavelets), ששיפרה באופן משמעותי את היכולת לנתח ולעבד אותות בעלי תופעות טרנזיינטיות. דוגמה נוספת היא המודל של ייצוגים דלילים (sparse) ויתירים (overcomplete), שזכתה להתעניינות מחקרית רבה בעשור האחרון. מודל זה מבוסס על האבחנה לפיה אותות רבים ניתנים לייצוג מדויק כצירוף ליניארי של מספר קטן של "אטומים" מתוך מילון מתאים.

רעיון כמו הייצוג של אותות באופן דליל ניתן למידול מתמטי בדרכים שונות. בעבודה זו, נתרכז בשתי שיטות מידול: השיטה הבייסיאנית (Bayesian) והגישה של הוספת אילוצים מבניים בשערוך דטרמיניסטי. בעולם הבייסיאני, המודל מקבל ייצוג מתמטי כהתפלגות א-פריורי על הסיגנל המשוערך; כך, לדוגמה, ניתן לייחס הסתברות א-פריורי גבוהה יותר לאותות דלילים מאשר לאותות שאינם ניתנים לייצוג דליל. לעומת זאת, בבעיות שערוך דטרמיניסטיות, ההנחה היא כי הסיגנל המשוערך אינו כולל הסתברות א-פריורי. בשערוך מסוג זה ניתן להכניס מודל על ידי הוספת אילוץ, לפיו הסיגנל הלא-ידוע שייך לקבוצה מוגדרת מראש של ערכים אפשריים. כך נוכל, למשל, להניח כי הסיגנל שייך לקבוצת האותות בעלי ייצוגים דלילים, וזאת מבלי להקצות הסתברות לאיברי הקבוצה. לפיכך, הגישה הדטרמיניסטית מתאימה יותר כאשר אין אנו מעוניינים לספק אמירות מדויקות לגבי הסבירות א-פריורי של הופעת סיגנלים שונים.

מטרת עבודה זו היא לנתח את הביצועים שניתן להשיג תחת מודלים מבניים מסוגים שונים. בפרט, אנו נתרכז בבעיות שערוך אותות מתוך מדידות רועשות, וזאת הן בעולם הבייסיאני והן בעולם השערוך הדטרמיניסטי. לשם כך יהיה עלינו לפתח כלים תיאורטיים-כלליים, שיתאימו למגוון מודלים. לאחר מכן, ניישם את הכלים האלה במספר מודלים ספציפיים, כגון אותות דלילים, שהם בעלי עניין עכשווי בקרב קהילת עיבוד האותות.

כדי לאפיין את הביצועים האפשריים של אלגוריתמי שערוך, נפתח שני סוגים משלימים של חסמי ביצועים: חסמים תחתונים, המגדירים מהם הביצועים הטובים ביותר אותם ניתן להשיג, וחסמים עליונים (או הבטחות ביצועים), המספקים רמת ביצועים נתונה ומבטיחים שאלגוריתם מסוים לעולם לא יהיה גרוע ממנה. החסמים התחתונים מַכְמֶתים את קושי הבעיה, ולעיתים מאפשרים לזהות כיצד ניתן לשנות את המודל על מנת להקל על השערוך. לעיתים ניתן להשוות בצורה ישירה את החסם התחתון לביצועים בפועל של אלגוריתם ספציפי. אולם, במקרים רבים מרחב הסיגנלים האפשריים רחב מכדי שניתן יהיה למדוד אמפירית את הביצועים של אלגוריתם בכל מצב. כאן נכנס תפקידו של החסם העליון, אשר מבטיח כי ביצועי המשערך לעולם לא יהיו גרועים יותר מגודל נתון וקל לחישוב.

אידיאלית, היינו רוצים לגלות כי החסם התחתון והחסם העליון קרובים זה לזה: מצב זה מעיד על כך שהביצועים האופטימליים ידועים וכי ניתן (כמעט) להשיג אותם בעזרת אלגוריתם

תודות

המחקר נעשה בהנחיית פרופ׳ יונינה אלדר בפקולטה להנדסת חשמל.
פרק 5 בוצע במשותף עם אלכסנדר יונג ופרופ׳ פרנץ לוואטש מהמכון להנדסת תקשורת ותדרי רדיו באוניברסיטה הטכנית של וינה. פרק 6 בוצע במשותף עם פרופ׳ מיכאל אלעד מהפקולטה למדעי המחשב בטכניון. פרק 8 בוצע במשותף עם תומר מיכאלי מהפקולטה להנדסת חשמל בטכניון. אני מודה לכל השותפים למחקר על תרומתם.
אני מודה לקרן וולף, לקרן אירווין וג׳ואן ג׳ייקובס, ולקרן משפחת ויטרבי על התמיכה הכספית הנדיבה בהשתלמותי.

חסמי ביצועים לשערוך אותות בעלי אילוצים מבניים

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת התואר
דוקטור לפילוסופיה

צביקה בן-חיים

הוגש לסנאט הטכניון – מכון טכנולוגי לישראל
חשוון תשע״א חיפה נובמבר 2010

חסמי ביצועים לשערוך אותות בעלי אילוצים מבניים

צביקה בן-חיים