Order Number 9513498

Biased and unbiased Cramer-Rao bounds: Computational issues and applications

Usman, Mohammad, Ph.D.

The University of Michigan, 1994

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

BIASED AND UNBIASED CRAMER-RAO BOUNDS: COMPUTATIONAL ISSUES AND APPLICATIONS

by

Mohammad Usman

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in The University of Michigan
1994

Doctoral Committee:

Associate Professor A.O. Hero III, Chair
Assistant Professor J.A. Fessler
Professor W.L. Rogers
Professor W. Williams
Associate Professor A.E. Yagle

Tomorrow is going to be wonderful, because tonight I do not understand anything.

Neils Bohr.

Dedicated to my parents and my sisters.

ACKNOWLEDGEMENTS

I am very grateful to Professor Alfred Hero for his guidance and encouragement. Our relationship took several turns in the last three and a half years. We started with some work on neural networks, then mutual information, and finally ended with the Cramer-Rao bound. Most of the ideas in this thesis originated during one of the several discussions that we had over the last few years. During this time Professor Jeffrey Fessler was also there to help. I am very thankful to him for letting me use his programs and in particular I appreciate his constructive comments. I would also like to acknowledge the physical insight in to SPECT imaging provided by Professor Leslie Rogers. I will always remember the manner in which he explained difficult concepts in easy and understandable way. Another source of encouragement was Professor Andrew Yagle, whose superb teaching skills helped to grow my interest in digital signal processing, and Professor William Williams who helped me to better understand the biological signals and their underlying processes. I would also like to thank Professor Noor Sheikh at the University of Engineering and Technology, Lahore, Pakistan, who has been a constant source of inspiration since my undergraduate years.

There is more to life than academia. Life would not be so great without Lily. She was always there when I needed her. I particularly enjoyed several trips that I took with Paul, Christine, Raja and Nasser, fixing cars with Gus, many interesting discussions with Vasu, the friday afternoon movies and the roller coaster rides with Robyn, and evening walks with Sudakshina. I would also like to thank my office mates for keeping up with me especially Jeff Frolik, Raghavan, Yi, Wonjin, Shue-Er, Bulent, Hamid, Nick Antoniadis, Nick Leung, Steve and Tal.

It would be very hard for me to imagine my life without cricket. I made several friends at the UM cricket team and I enjoyed every moment of it. Getting to the Midwest tournament final in 1993 was one of the most exciting moments of my stay here.

I am very thankful to the Ministry of Science and Technology, Government of Pakistan for awarding me the scholarship for postgraduate studies, and to the National Science Foundation under grant BCS-9024370.

Most of all, I would like to thank my parents and my sisters for their support and making it all worthwhile.

TABLE OF CONTENTS

LIST OF APPENDICES

Appendix

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

1.1 Problem Statement and Significance

Suppose that a signal θ is transmitted over a noisy channel, where the observed signal y is a noise corrupted version of θ, and the channel transfer characteristics are governed by the transition density $f(y;\theta)$. The *estimation problem* is to estimate θ given y. For example in emission tomography we are interested in estimating the image parameters or the pixel intensities. The channel model is specified by the system geometry. The set of observations are the number of detected γ-rays.

Associated with every estimation problem is a cost function which an estimator attempts to minimize. The cost is generally measured in terms of the error of estimation. The most common measure of error of an estimator is the mean square error (MSE). The MSE is simply expectation of the square of the difference between the actual parameter value θ and the estimated parameter value $\widehat{\theta}$. The MSE can be directly related to the estimator bias and variance. While the estimator bias measures a mismatch between θ and the mean of $\widehat{\theta}$, the variance determines the spread of the error around the mean of $\widehat{\theta}$. An estimator is called unbiased if θ and the mean of $\widehat{\theta}$ are identical. Clearly unbiasedness and low variance are desirable properties of an estimator. Unfortunately in most estimation problems, there is a trade-off between the estimator bias and variance, i.e. lower variance can only be bought at the price of increased bias. In this thesis we provide a framework to quantify the bias-variance trade-offs for an estimator, and to compare the performance of different estimators on the bias-variance trade-off plane. The principal tool for our methodology is a lower bound on the estimator error variance.

A bound on variance as a function of the estimator bias provides us with a

benchmark against which different estimators can be compared in terms of their closeness to the bound on variance, subject to the constraint that they do not exceed some preset threshold on allowable bias. If the bound is attained by a given estimator, then the designer can be confident that no other estimator exists with improved bias-variance performance. Furthermore, the variance bound can be used for channel/system optimization since it is expressed in terms of the channel/system parameters.

The conventional unbiased Cramer-Rao (CR) bound is our starting point. It determines a lower bound on the variance of any unbiased estimator. This CR bound has been widely applied in several different fields as a performance measure against which different estimation algorithms can be compared. The applications include the fields of array processing, parametric spectrum estimation, image processing, medical imaging, and numerous others [1, 2, 3, 4, 5, 6, 7, 8, 9].

Although the unbiased CR bound is known to be asymptotically achievable for large numbers of independent observations, many estimators are biased and their variance is not bounded by the unbiased CR bound. A performance comparison of estimation algorithms based on the unbiased CR bound is therefore meaningless unless the estimator bias is taken into consideration. A biased form of the CR bound is available, e.g. van Trees [33], but it only applies to a very restricted class of estimators with a fixed bias gradient function and it can not be used to simultaneously bound a pair of estimators which have different, but perhaps acceptable, bias gradients.

In [14] Hero presents a uniform CR bound that is applicable to all biased estimators whose bias gradient length is less than a small pre-specified threshold. The uniform CR bound specifies a curve over the bias-gradient-length/variance trade-off plane, denoted δ-σ, which separates the plane into achievable and unachievable regions of variance and bias gradient lengths. Different biased estimators can be placed on this trade-off plane and thus we can compare them with the same yardstick. However, the application of the uniform CR bound is limited due to the following factors:

- The uniform CR bound requires computing elements of the inverse of an $n \times n$ Fisher information matrix (FIM), where n is the length of the unknown parameter

vector. Since matrix inversion requires $\mathcal{O}(n^3)$ floating point operations (flops), the computation can become intractable if n is large. An example is image reconstruction where for a moderate size image of dimension 128×128, the FIM is of dimensions $128^2 \times 128^2$ requiring $\mathcal{O}(10^{12})$ flops. Even on a state of the art workstation, currently we require over 12 hours of computation time for calculating such a matrix inverse.

- The uniform CR bound requires computation of the bias gradient length of the estimator. For most cases, the computation of the bias gradient is analytically intractable and has to be empirically estimated. The standard method of moments could be used to estimate the bias gradient, but it is computationally inefficient for a large n.

- The uniform CR bound as presented in [14] is only applicable to well-posed estimation problems with well-conditioned FIM. Since many large dimensional estimation problems may be ill-posed, the uniform CR bound of [14] must be extended to allow singular, ill-conditioned, and rank deficient FIMs.

1.2 Overview of Thesis

Often only a few primary estimator components of the n-dimensional estimator are of direct interest. For example, in image reconstruction one is usually interested in a small image region-of-interest (ROI) corresponding to a tumor or a lesion. While, most of the results are easily extended to a larger region of interest, in this thesis we concentrate on the case when the region of interest consists only of one estimator component.

In [10] Hero and Fessler presents a recursive, monotone convergent algorithm, based on a geometric series argument, which computes a single column of the inverse of the FIM. The algorithm, called the geometric series (GS) algorithm in this thesis, is a special case of the more general inversion algorithms known as basic iterative methods in the literature [31, 22, 29]. The GS algorithm requires only $\mathcal{O}(n^2)$ flops per iteration and can result in considerable computational savings if it converges in significantly less than n iterations. An advantage of this algorithm is its monotone convergence, which guarantees a valid lower bound at the end of each iteration. However, the GS algorithm requires specification of a full rank sparse matrix F, called a splitting matrix, whose inverse is easy to compute and

is greater than the FIM in the sense of matrix inequality. In [10] a method of selecting an appropriate F was introduced based purely on statistical considerations. In this thesis we present methods for acceleration of the recursive GS algorithm based on purely algebraic considerations. This approach is motivated by optimizing the average and asymptotic rates of convergence for the GS algorithm over a class of structured splitting matrices. We also present a conjugate gradient (CG) type of algorithm tailored for the inversion of a positive definite and symmetric FIM, which has a significantly faster convergence rate, but has the disadvantage of non-monotone convergence.

For the case of singular FIM, the ROI CR bound requires a column of the pseudo-inverse of the FIM. Direct methods of pseudo-inversion are generally based on singular value decomposition of the FIM, which requires $\mathcal{O}(n^3)$ flops to compute [29, 22, 18, 19]. We present two recursive algorithms to approximate a column of the Moore-Penrose pseudo-inverse of the FIM with $\mathcal{O}(n^2)$ flops per iteration. These algorithms generalize our recursive algorithms to the case of singular FIMs. The first algorithm, a perturbation method, converges monotonically when implemented using the GS. The second algorithm, a differencing method, is non-monotonic for both the GS and the CG implementations, but is more accurate for most cases.

For an estimator to be placed on the bias gradient-variance trade-off plane, its variance and bias gradient length must be calculated. For most estimators, bias, bias gradient, and variance are analytically intractable and must be estimated. While bias and variance can be easily estimated using the standard method of moments approach, the standard approach to the estimation of the estimator bias gradient requires $n - 1$ times more computations as compared to the estimation of bias and variance. We present a more efficient and accurate method for experimentally determining the estimator bias gradient using a weighted sample average of the estimator bias. This method requires a single simulation of the same type as that commonly used to experimentally determine estimator bias and variance.

For ill-posed problems, such as super resolution in image restoration, where the number of parameters to be estimated are fewer than the number of independent observations, the FIM is singular. We extend the uniform CR bound of [14] to the case of singular

FIM via the pseudo-inverse.

The methods developed in this thesis have wide application. In this thesis we consider the following applications.

1. **Two-dimensional single photon emission computed tomography (SPECT) imaging.**

   The SPECT system used in most of the examples considered in this thesis is a ring tomograph called SPRINT [40]. The system was designed specifically for brain imaging and consists of a ring of detectors and a ring of collimators. During imaging time, the collimator ring is rotated through small steps about the source. A detailed description of the SPRINT II system is given in chapter 4.

   - **Comparison of the convergence rates for the GS and the CG recursive algorithms** (chapter 2).

     The rate of the convergence of the GS algorithm increases with the complexity of the splitting matrix. In particular, for a 32×32 image reconstruction, a splitting matrix with 50 off-diagonals achieves convergence in a similar number of iterations as the CG algorithm. However, the CG algorithm requires half as many flops.

   - **Comparison of penalized maximum-likelihood and weighted least-squares reconstruction algorithms** (chapter 4).

     The penalized maximum-likelihood (PML) and the weighted least-squares estimators (WLSE) were found to have similar bias gradient-variance trade-off curves in the δ–σ plane and followed the uniform CR bound closely. In particular the WLSE with optimal weights is shown to exactly achieve the uniform CR bound for all biases.

   - **Study of the effect of angular sampling using the uniform CR bound** (chapter 4).

     The uniform CR bound decreases as the number of collimator rotations of the ring increases. For a 32×32 image, the uniform CR bound tends to saturate

after 8 rotations, i.e. increasing the number of rotations does not decrease the uniform CR bound significantly. This indicates that the image is adequately sampled by 8 collimator rotations.

- **Design of optimal apertures using the uniform CR bound** (chapter 4). For a count normalized case, the uniform CR bound shows a sharp minimum for an aperture opening corresponding to a ray width of one pixel at all bias gradient lengths. Hence, this aperture opening is universally optimal for the pixel of interest for any biased or unbiased estimator.
- **Performance gain analysis of adding a vertex view anger camera to a standard SPRINT ring using the uniform CR bound** (chapter 4). We performed several experiments to analyze the gain in performance by adding an orthogonal vertex view to a SPRINT ring. The gain in performance was measured in terms of reduction in the uniform CR bound. The vertex view with the SPRINT ring shows a gain of at least 50% over the SPRINT ring alone in the absence of any background noise.

2. **One-dimensional discrete deconvolution.**

   - **Comparison of the error of approximation of the two recursive methods to estimate the pseudo-inverse of the FIM** (chapter 2) . The GS and the CG algorithms are compared in these analyses. For all the cases considered the differencing method gives a lower error of approximation when compared to the perturbation method. However, the GS implementation of the perturbation method converges monotonically, while the differencing method displays non-monotonic convergence both for the GS and the CG implementations.
   - **Performance evaluation of a WLSE for an ill-posed Gaussian kernel resulting in a singular FIM** (chapter 3). Even with the ideal weights, the WLSE fails to achieve the uniform CR bound. This is due to the fact that for singular problems uniform CR bound may not be achievable.

3. **One-dimensional discrete edge localization in the presence of Gaussian noise** (chapter 3).

   A Canny operator [39], which is similar to a Laplacian in one-dimension, was used for edge localization. The localization algorithm shows an optimum Canny filter bandwidth in terms of achieving minium MSE. The minimum MSE corresponds to the minimum bias and variance.

1.3 Contributions of This Thesis

This thesis makes the following specific contributions:

1. Specification of new banded diagonal splitting matrices to improve the average and the asymptotic rates of convergence of the geometric series algorithm (chapter 2).

2. Modification of the non-monotone convergent conjugate gradient algorithm, tailored for the inversion of positive-definite and symmetric FIM (chapter 2).

3. Specification of two new recursive algorithms to compute an approximation to the pseudo-inverse of the FIM requiring only $\mathcal{O}(n^2)$ flops per iteration. The recursive algorithms are based on the GS and the CG algorithms (chapter 2).

4. Extension of the uniform CR bound of [14] to singular FIM (chapter 3).

5. Specification of an efficient empirical approximation to the bias gradient, which has zero bias and low variance. The length of this bias gradient approximation can be used to compute the uniform CR bound (chapter 3).

6. A study of the SPECT imaging system performance evaluation and optimization via several examples. Comparison of image reconstruction algorithms, namely, the penalized maximum-likelihood and the weighted least squares, via the application of the uniform CR bound.

1.4 Outline of Chapters

This thesis is arranged in five chapters.

Chapter 1 is the introduction which you are currently reading.

Chapter 2 presents different methods to compute a column of the inverse and the pseudo-inverse of the FIM. First we specify a class of banded diagonal matrices to improve the rate of convergence of the GS algorithm (Section 2.3). We also modify CG algorithm to compute the columns of the inverse of the FIM (Section 2.5). Based on the GS and the CG algorithms we present two recursive algorithms to approximate the columns of the pseudo-inverse of the FIM (Section 2.6). The methods developed in this chapter are demonstrated through applications in two-dimensional image reconstruction and one-dimensional discrete deconvolution (Section 2.7).

Chapter 3 extends the uniform CR bound to allow for rank deficient FIM (Section 3.3). An efficient method to estimate the bias gradient is then given (Section 3.4). Using the methods developed in Chapter 2 we give a recipe to compute efficiently the uniform CR bound for non-singular FIM (Section 3.5). Applications to one-dimensional discrete deconvolution and one-dimensional edge localization are considered (Section 3.6).

Chapter 4 contains several application to SPECT for a single photon ring tomograph (SPRINT) system design, optimization, and performance analysis via the uniform CR bound.

Chapter 5 summarizes the thesis and indicates future directions of research.

CHAPTER 2

RECURSIVE ALGORITHMS FOR COMPUTING THE CRAMER-RAO BOUND

In this chapter we present recursive algorithms to compute columns of the inverse and the pseudo-inverse of a FIM required for the determination of the CR bound. The methods developed in this chapter will be used in Chapter 3 to efficiently compute the uniform CR bound.

2.1 Introduction

The Cramer-Rao (CR) bound determines a lower bound on the covariance of any estimator with a given bias gradient. For an n-dimensional parameter vector, and when the FIM F_Y is non-singular, calculation of the CR bound involves inversion of F_Y. Direct matrix inversion, requiring $\mathcal{O}(n^3)$ flops (floating point operations), is generally computationally intractable if the number of unknown parameters is large. For example, in image reconstruction where the pixel intensities are the parameters to be estimated, a moderate size image of 128×128 pixels has a FIM of dimension $(128)^2 \times (128)^2$ and will require $\mathcal{O}(10^{12})$ flops to compute its inverse. On a 100 MIPS workstation this would require over 12 hours of computation time for each object distribution to be considered.

Often only a few components of the n-dimensional estimator are of interest. For example in image reconstruction one may be primarily interested in a small image region of interest (ROI) corresponding to a tumor or a lesion. The method of sequential partitioning can be used in this case to calculate the bound on $\mathrm{cov}_{\underline{\theta}}(\underline{\hat{\theta}}^{ROI})$, where $\underline{\hat{\theta}}^{ROI}$ is an estimator of q parameters $\underline{\theta}^{ROI}$ of interest [16]. This still requires $\mathcal{O}(n^3)$ flops. Hero et al. [10, 13] present a recursive method, based on a linearly convergent geometric series decomposition of the

inverse, for calculating a small $q \times q$ sub-matrix of the matrix ${F_Y}^{-1}$ required for computing the CR bound on $\mathrm{cov}_{\underline{\theta}}(\underline{\hat{\theta}}^{ROI})$. The geometric series algorithm requires only $\mathcal{O}(n^2)$ flops per iteration so that if the convergence rate is fast, a significant computational saving is achieved. Furthermore, if the FIM is of the form $F_Y = A^T \Gamma A$, where Γ is an $m \times m$ diagonal, A is $m \times n$, and $n \ll m$ or when A is sparse, the geometric series algorithm has much reduced storage requirement as compared to the direct inversion methods. As presented in [10], the geometric series algorithm stipulates that we find a $n \times n$ matrix F, called the splitting matrix, such that F^{-1} is easy to compute ($\mathcal{O}(n^2)$ or less) and such that F dominates F_Y in the sense that $F - F_Y \geq 0$. In [10, 13] a method for selecting an appropriate matrix F was introduced based on purely statistical arguments; F is chosen as the FIM of any more informative data set such as a complete-data set used in the EM algorithm [12]. The main advantage of the algorithm of [10, 13] is its monotone convergence which generates a valid and improving lower bound on estimator covariance at each iteration of the algorithm. However, the price paid for monotone convergence is slow linear convergence rate, and its high sensitivity to the condition number of F_Y. For applications where a strict lower bound is not required at each iteration, a non-monotone algorithm may be acceptable. While the statistical approach of [10, 13] provides powerful guiding principle for finding splitting matrices which guarantee monotonic convergence, the approach requires finding a suitable reformulation of the estimation problem. Furthermore, unless the complete-incomplete data space is defined cleverly, the approach of [10, 13] does not guarantee rapid rate of convergence.

In this chapter we present methods for acceleration of the recursive algorithm of [10, 13] which are based on purely algebraic considerations. First an expression for an F_Y-dominating splitting matrix is given which optimizes the average convergence rates for regularly circulant FIM and gives a monotonically convergent series. When applied to a non-circulant FIM, our numerical results indicate that this splitting matrix provides very close to optimal average convergence rates. Second, we present an over-relaxed version of the geometric series which can greatly accelerate the convergence rate at the expense of non-monotonic convergence. Third, we present a conjugate gradient type of algorithm tailored for the inversion of positive-definite, symmetric FIM, which has a significantly faster

convergence rate and is less sensitive to the condition number of F_Y.

In many signal processing application, e.g. signal deconvolution, the parametric estimation problem is ill-posed which results in a singular FIM. The covariance of the estimates in this case are lower bounded by a function of the Moore-Penrose pseudo-inverse[1] of the FIM [11]. Direct calculation of the pseudo-inverse usually requires a singular value decomposition (SVD) of the FIM, which is an $\mathcal{O}(n^3)$ operation. We present two recursive algorithms to approximate columns of the Moore-Penrose pseudo-inverse of the FIM with $\mathcal{O}(n^2)$ flops per iteration. These algorithms generalize our recursive algorithms to the case of singular FIM. The first algorithm, a perturbation method, converges monotonically to the pseudo-inverse of F_Y when implemented using the geometric series approach. The other algorithm, a differencing method, is non-monotonic for both the geometric series and the conjugate gradient implementations.

We apply the methods developed in this chapter to: 1) Single Photon Emission Computed Tomography (SPECT) image reconstruction, and 2) one dimensional discrete deconvolution. In both these applications the FIM F_Y is of the factored form permitting a significant reduction in memory requirements.

2.2 Unbiased Cramer-Rao Bound

Given a measurement $\underline{Y}$ that is a vector of random variables and has probability distribution $f_{\underline{Y}}(\underline{y};\underline{\theta})$ dependent on an unknown parameter vector $\underline{\theta}$, it is desired to estimate a real parameter vector $\underline{\theta} = [\theta_1, ..., \theta_n]^T$. A vector function $\widehat{\underline{\theta}} = \widehat{\underline{\theta}}(\underline{y})$ is a parameter estimator based on the observation $\underline{Y} = \underline{y}$. The performance of $\widehat{\underline{\theta}}$ can be measured by its bias vector $\underline{b}(\underline{\theta}) = [\text{bias}_{\underline{\theta}}(\widehat{\theta}_1), ..., \text{bias}_{\underline{\theta}}(\widehat{\theta}_n)]^T$:

$$\underline{b}(\underline{\theta}) = \text{bias}_{\underline{\theta}}(\widehat{\underline{\theta}}) = E_{\underline{\theta}}(\widehat{\underline{\theta}}) - \underline{\theta},$$

and by its covariance matrix $S(\underline{\theta})$:

$$S(\underline{\theta}) = \text{cov}_{\underline{\theta}}(\widehat{\underline{\theta}}) = E_{\underline{\theta}}\left[(\widehat{\underline{\theta}} - E_{\underline{\theta}}(\widehat{\underline{\theta}}))(\widehat{\underline{\theta}} - E_{\underline{\theta}}(\widehat{\underline{\theta}}))^T\right],$$

[1] In this thesis we will use Moore-Penrose pseudo-inverse and pseudo-inverse interchangeably.

where $E_{\underline{\theta}}$ denotes statistical expectation with respect to the pdf $f_{\underline{Y}}(\underline{y};\underline{\theta})$. The CR lower bound on the covariance of any unbiased ($\text{bias}_{\underline{\theta}}(\underline{\hat{\theta}}) = 0$) parameter estimator, called the unbiased CR bound, is given by the Moore-Penrose pseudo-inverse $F_Y^+(\underline{\theta})$ of the estimator-independent Fisher information matrix (FIM) [11]:

$$S(\underline{\theta}) \geq F_Y^+(\underline{\theta}), \tag{2.1}$$

where $F_Y = F_Y(\underline{\theta})$ is the $n \times n$ FIM:

$$F_Y = -E_{\underline{\theta}}[\nabla_{\underline{\theta}}^T \ln f_{\underline{Y}}(\underline{Y};\underline{\theta})\, \nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y};\underline{\theta})],$$

and $\nabla_{\underline{\theta}}$ denotes the (row) gradient vector $\left[\frac{\partial}{\partial\theta_1}, \frac{\partial}{\partial\theta_2}, ..., \frac{\partial}{\partial\theta_n}\right]$. The matrix inequality (2.1) is simply an assertion that $S(\underline{\theta}) - F_Y^+(\underline{\theta})$ is positive semi-definite.

When F_Y is non-singular the unbiased CR lower bound (2.1) is given by the inverse of FIM:

$$S(\underline{\theta}) \geq F_Y^{-1}(\underline{\theta}). \tag{2.2}$$

The unbiased CR bound on the variance of a single component $\hat{\theta}_1$, corresponds to the top left entry of F_Y^+, denoted $[F_Y^+]_{11}$:

$$\text{var}_{\underline{\theta}}(\hat{\theta}_1) \geq \underline{e}_1^T F_Y^+(\underline{\theta})\underline{e}_1,$$

where $\underline{e}_1 = [1, 0, ..., 0]^T$ is the n-dimensional unit vector. Throughout this thesis a bound on $\text{var}_{\underline{\theta}}(\hat{\theta}_1)$ will be of interest. A bound on $\text{var}_{\underline{\theta}}(\hat{\theta}_k)$, $k \neq 1$ can be obtained by permuting the rows and columns of F_Y. The results to be presented are easily extended to q-elements: $q \leq n$.

2.3 Recursive CR Bound Algorithms

2.3.1 Monotone Convergent Algorithm: Geometric Series

For more details and proof of convergence of this method see [10]. Assume that F_Y is non-singular and that there exists an invertible $n \times n$ matrix F such that F dominates F_Y: $F \geq F_Y > 0$. It can be easily shown that all the eigenvalues of $I - F^{-1}F_Y$ lie in the interval $[0, 1)$. We can therefore apply the matrix form of the geometric series to F_Y^{-1}, also

known as the 'Neumann expansion' [32]:

$$\begin{aligned} F_Y^{-1} &= [F - (F - F_Y)]^{-1} \\ &= [I - (I - F^{-1}F_Y)]^{-1} F^{-1} \\ &= \left[\sum_{k=0}^{\infty}(I - F^{-1}F_Y)^k\right] F^{-1}. \end{aligned}$$

An approximation $B^{(m)}$ to F_Y^{-1} can be obtained by truncating the infinite series:

$$B^{(m)} = \left[\sum_{k=0}^{m}(I - F^{-1}F_Y)^k\right] F^{-1}.$$

Using induction on k it is straightforward to establish the following recursive relation:

$$B^{(k+1)} = [I - F^{-1}F_Y]\, B^{(k)} - F^{-1}. \tag{2.3}$$

The above expression converges monotonically to the CR bound in the sense that $B^{(k)} - B^{(k+1)} = [I - F^{-1}F_Y]^{k+1} F^{-1}$ is positive semi-definite for all k and $F_Y^{-1} - B^{(k)}$ converges to a matrix of zeros as $k \to \infty$. Note that the recursion (2.3) updates each column of $B^{(k)}$ independently.

Let $\beta^{(k)}$ be the first column of the matrix $B^{(k)}$. Then a recursive algorithm for computing the CR bound $[F_Y^{-1}(\underline{\theta})]_{11}$ on $\mathrm{cov}_{\underline{\theta}}(\hat{\theta}_1)$, which we call the geometric series (GS) algorithm, is given as follows:

GS Algorithm

INITIALIZE: $\underline{\beta}^{(0)} = \underline{0}$ an $n \times 1$ vector of zeroes.

RECURSION:

$$\underline{\beta}^{(k+1)} = \underline{\beta}^{(k)} - F^{-1}(\, F_Y\, \underline{\beta}^{(k)} + \underline{e}_1\,). \tag{2.4}$$

The recursion yields a sequence of vectors $\underline{\beta}^{(k)}$ which converges to the first column of F_Y^{-1}. The first element of $\beta^{(k)}$ is simply $\zeta^{(k)} = \underline{e}_1^T\, \underline{\beta}^{(k)} = \underline{e}_1^T B^{(k)} \underline{e}_1$. $\zeta^{(k)}$ monotonically converges to $[F_Y^{-1}(\underline{\theta})]_{11}$ as $k \to \infty$. This monotone convergence property guarantees that the scalar $\zeta^{(k)}$ will be a valid lower bound for all $k > 0$ and that this bound improves at each iteration [10].

The matrix recursion (2.3) is a special case of the 'matrix splitting method' [22, 28] for iteratively approximating a matrix inverse. When F is taken as the scaled identity γI, where γ is the Frobenius norm of F_Y, (2.3) reduces to the Frobenius normalization method for inversion of sparse matrices. When F is taken as the diagonal component of F_Y then (2.3) reduces to the Jacobi method, which however is generally non-monotone convergent as $F - F_Y$ may not be positive definite.

Since recursion (2.4) is a vector update, it requires only $\mathcal{O}(n^2)$ flops per iteration. Therefore if the algorithm is stopped for $k \ll n$, we obtain a bound on $\text{cov}_{\underline{\theta}}(\hat{\theta}_1)$ with significantly fewer flops than required for direct calculation of the (1,1) element of F_Y^{-1}. For this algorithm (2.4) to have computational advantage over direct inversion and also deliver a tight bound we require F^{-1} to be easy to compute and convergence to be fast.

In [10, 13] Hero et. al. exploited the properties of Fisher information matrices to specify an appropriate splitting matrix F. In this estimation theoretic approach one identifies a complete-incomplete data setting for $f_Y(y;\underline{\theta})$. Assume that we can find a hypothetical set of measurements $X \sim f_X(x;\underline{\theta})$ which is larger and more informative than Y in the sense that the conditional density $f_{Y/X}(y/x;\underline{\theta})$ is independent of $\underline{\theta}$. Let F_X and F_Y be the FIM for the 'complete data' X and 'incomplete data' Y, respectively, the data processing theorem asserts that $F_X - F_Y \geq 0$ [10]. This implies that F_X can be used as a splitting matrix. If X is chosen intelligently such that F_X is easily invertible, then GS can be implemented efficiently. In contrast to this estimation theoretic approach, in this chapter we will utilize purely algebraic approaches to specifying a splitting matrix F in order to improve the convergence rate of the GS. In section 2.7 we will compare the algebraic approach presented in this chapter to the estimation theoretic approach of [10, 13].

In the next section we show that the average rate of convergence of the GS algorithm increases as the distance between F and F_Y decreases. The objective will then be to select a simple sparse matrix F which satisfies $F - F_Y \geq 0$ and makes the distance small. It can be easily shown that the splitting matrix F can be viewed as a pre-conditioning matrix for the gradient search algorithm.

2.4 Convergence Rates for GS Algorithms

Two types of convergence criteria are commonly used in evaluating direct iterative methods: the average rate of convergence and the asymptotic rate of convergence. These are discussed below.

Defining $\Delta\beta^{(k)} = F_Y^{-1}\underline{e}_1 - \beta^{(k)}$, from (2.4) it is easily shown that

$$\Delta\beta^{(k+1)} = M^k \Delta\beta^{(0)},$$

where $M = I - F^{-1}F_Y$ and $\Delta\beta^{(0)} = F_Y^{-1}\underline{e}_1$. The Euclidean norm $\|\Delta\beta^{(k+1)}\|_2$ of the vector $\Delta\beta^{(k+1)}$ can be related to the matrix-2 and Frobenius norms $|||M|||_2$ and $|||M|||_{\mathcal{F}}$ of the matrix M and the Euclidean norm of $\Delta\underline{\beta}^{(0)}$ by [21]:

$$\begin{aligned} \|\Delta\beta^{(k+1)}\|_2 &\leq |||M^k|||_2 \|\Delta\beta^{(0)}\|_2 \\ &\leq |||M^k|||_{\mathcal{F}} \|\Delta\beta^{(0)}\|_2, \end{aligned} \tag{2.5}$$

where for any real vector $\underline{z}$, $\|\underline{z}\|_2$ is the Euclidean norm:

$$\|\underline{z}\|_2 = \sqrt{\sum_i |z_i|^2},$$

and for any real matrix $A = ((a_{ij}))$, $|||A|||_2$ is the matrix-2 norm and:

$$|||A|||_{\mathcal{F}} = \sqrt{\text{trace}(A^T A)} = \sqrt{\sum_i \sum_j |a_{ij}|^2}.$$

From (2.5), $|||M^k|||_{\mathcal{F}}$ bounds the relative improvement in the approximation error $\Delta\beta^{(k+1)}$ as compared to the initial approximation error $\Delta\beta^{(0)}$: $\frac{\|\Delta\beta^{(k+1)}\|_2}{\|\Delta\beta^{(0)}\|_2} \leq |||M^k|||_{\mathcal{F}}$. In general the Frobenius norm is preferable to the matrix-2 norm in applications since computation of $|||M|||_{\mathcal{F}}$ requires no expensive singular value decomposition of M.

We specify an absolute convergence criterion in terms of a threshold $\gamma \in [0, 1]$.

Definition *The GS algorithm is said to have converged after N iterations (to within a factor γ) if*

$$\frac{\|\Delta\beta^{(k+1)}\|_2}{\|\Delta\beta^{(0)}\|_2} < \gamma, \quad \forall k > N. \tag{2.6}$$

The minimum integer $N = N(\gamma)$ such that this inequality is satisfied is called 'the required number of iterations' for convergence.

Note that unless $\Delta\beta^{(0)} = 0$, $N(\gamma)$ increases as γ decreases to zero. Relation (2.5) and the convergence criterion (2.6) provide the following recipe for specifying an upper bound on the required number of iterations for convergence: *An upper bound on $N(\gamma)$ is any N such that*

$$|||M^k|||_{\mathcal{F}} < \gamma, \quad \forall k > N. \tag{2.7}$$

For specified γ the *average rate of convergence* is defined as [31]:

$$R_{ave}(\gamma) = \left(|||M^{N(\gamma)}|||_{\mathcal{F}}\right)^{\frac{1}{N(\gamma)}}. \tag{2.8}$$

In terms of $R_{ave}(\gamma)$ an upper bound N or $N(\gamma)$ is given by (2.7) :

$$N = \frac{\log\left(\frac{1}{\gamma}\right)}{\log\left(\frac{1}{R_{ave}(\gamma)}\right)}. \tag{2.9}$$

Unfortunately $R_{ave}(\gamma)$ is typically difficult to compute and (2.9) is not implementable. On the other hand, we can obtain a simpler upper bound on $N(\gamma)$ by using the matrix norm inequality $|||M^N|||_{\mathcal{F}} \leq |||M|||_{\mathcal{F}}^N$. Using this inequality in (2.8) and (2.9) we obtain the following upper bound N on $N(\gamma)$:

$$N(\gamma) \leq N = \frac{\log\left(\frac{1}{\gamma}\right)}{\log\left(\frac{1}{|||M|||_{\mathcal{F}}}\right)}. \tag{2.10}$$

Therefore, reducing the Frobenius norm $|||M|||_{\mathcal{F}}$ reduces the upper bound on the required number of iterations $N(\gamma)$. If the GS algorithm is stopped after N iterations, where N is given by (2.10), then convergence to within a factor γ is guaranteed.

Another measure of convergence rate is the asymptotic rate of convergence equal to the root convergence factor $\lim_{N\to\infty} \left(|||M^N|||_{\mathcal{F}}\right)^{\frac{1}{N}}$ [21]. The root convergence factor is associated with the 'worst case' rate of convergence of the matrix splitting methods [22, 31]. The root convergence factor is equal to the maximum singular value, i.e. the maximum magnitude eigenvalue $|\lambda^M|_{\max}$ of M, which is called the spectral radius $\rho(M)$.

Since $\rho(M) = \lim_{N\to\infty} R_{ave}(\gamma)$, in view of relation (2.9), the spectral radius is sometimes used to approximate the required number of iterations $N(\gamma)$:

$$N(\gamma) \approx \frac{\log\left(\frac{1}{\gamma}\right)}{\log\left(\frac{1}{\rho(M)}\right)}. \tag{2.11}$$

However, $\rho(M)$ can be significantly less than $R_{ave}(\gamma)$, and therefore, unless γ is very small, use of the expression (2.11) can result in severely underestimating the required number of iterations. Therefore, as seen from the bound (2.10), and as pointed out in [31], while the spectral radius is a good predictor of the asymptotic convergence rate of the GS algorithm, the Frobenius norm $|||M|||_{\mathcal{F}}$ is more reliable as a predictor of the actual number of iterations required for convergence. This will be quantitatively illustrated in Section 2.7.

2.4.1 Design of Splitting Matrix

To ensure monotone convergence and optimize the convergence rates, two worthwhile objectives would be to select the splitting matrix F such that F is simple to invert, $F - F_Y$ is non-negative definite and such that either: i) $\rho(M)$ is small; or ii) $|||M|||_{\mathcal{F}}$ is small, where as before $M = I - F^{-1}F_Y$. For large dimension n, computation of $\rho(M)$ is intractable since it requires performing an SVD or power iterations on M to find the maximum eigenvalue [21]. Optimization of the Frobenius norm $|||M|||_{\mathcal{F}}$ is easier, we consider an even simpler but related objective: find a symmetric p-diagonal $n \times n$ matrix F_p which is simply inverted, satisfies $F_p - F_Y \geq 0$, and makes $|||F - F_Y|||_{\mathcal{F}}$ small. A p-diagonal matrix F_p is defined as any banded matrix of the form

$$F_p = \begin{bmatrix} d_{11} & d_{12} & \cdots & d_{1p} & & \cdots & 0 \\ d_{21} & d_{22} & & & \ddots & & \vdots \\ \vdots & & \ddots & & & \ddots & \\ d_{p1} & & & \ddots & & & d_{n-p,n} \\ & \ddots & & & \ddots & & \vdots \\ \vdots & & \ddots & & & \ddots & d_{n-1,n} \\ 0 & \cdots & & d_{n,n-p} & \cdots & d_{n,n-1} & d_{nn} \end{bmatrix}. \tag{2.12}$$

Making $|||F - F_Y|||_{\mathcal{F}}$ small is a reasonable objective since the following inequalities relate the minimization of $|||F - F_Y|||_{\mathcal{F}}$ to minimization of $\rho(M) = \lim_{N\to\infty} |||(I - F^{-1}F_Y)^N|||^{\frac{1}{N}}$ and of $|||M|||_{\mathcal{F}}$, where $|||.|||$ is any matrix norm.

$$\begin{aligned} |||(I - F^{-1}F_Y)^N|||^{\frac{1}{N}} &\leq |||(I - F^{-1}F_Y)^N||| \\ &\leq |||F^{-1}||| \; |||F - F_Y||| \\ &\leq |||F_Y^{-1}||| \; |||F - F_Y||| \\ &\leq |||F_Y^{-1}|||_{\mathcal{F}} \; |||F - F_Y|||_{\mathcal{F}} \, C, \end{aligned}$$

for some constant C depending upon the $|||.|||$. For the spectral norm $|||.||| = |||.|||_2$ the constant C is equal to 1. The first inequality is well known [21], the second is the Cauchy-Schwartz inequality for matrix norms, and the third inequality follows from the restriction $F - F_Y \geq 0$.

Consider the following p-diagonal matrix, D_p, constructed from the elements a_{ij} of the matrix A:

$$D_p = \begin{bmatrix} \sum_{j\neq 2,\ldots,p} |a_{1j}| & a_{12} & \cdots & a_{1p} & & \cdots & 0 \\ a_{21} & \sum_{j\neq 1,3,\ldots,p+1} |a_{2j}| & & & \ddots & & \vdots \\ \vdots & & \ddots & & \ddots & & \\ & & & \sum_{\substack{j\neq k-p,\ldots,k-1 \\ j\neq k+1,\ldots,k+p-1}} |a_{kj}| & & \ddots & \\ a_{p1} & & & & \ddots & & a_{n-p,n} \\ & \ddots & & & & \ddots & \vdots \\ \vdots & & \ddots & & & \ddots & a_{n-1,n} \\ 0 & \cdots & & a_{n,n-p} & \cdots & a_{n,n-1} & \sum_{j\neq n-p,\ldots,n-1} |a_{nj}| \end{bmatrix}. \quad (2.13)$$

In particular D_1 and D_2 are diagonal and tridiagonal matrices:

$$D_1 = \begin{bmatrix} \sum_{j=1}^{n} |a_{1j}| & & \cdots & 0 \\ & \sum_{j=1}^{n} |a_{2j}| & & \vdots \\ \vdots & & \ddots & \\ 0 & \cdots & & \sum_{j=1}^{n} |a_{nj}| \end{bmatrix}. \quad (2.14)$$

$$D_2 = \begin{bmatrix} \sum_{j\neq 2} |a_{1j}| & a_{12} & & & \cdots & 0 \\ a_{21} & \sum_{j\neq 1,3} |a_{2j}| & \ddots & & & \vdots \\ & \ddots & \ddots & \ddots & & \\ & & \ddots & \sum_{j\neq k-1,k+1} |a_{kj}| & \ddots & \\ \vdots & & & \ddots & \ddots & a_{n-1,n} \\ 0 & \cdots & & & a_{n,n-1} & \sum_{j\neq n-1} |a_{nj}| \end{bmatrix}.$$

The following lemma is proven in Appendix A. The lemma could be alternatively proven by using Gershgorin circle theorem [22].

Lemma 1 *Assume that A is an $n\times n$ symmetric matrix. Then $D_p - A$ is PSD. Furthermore, if A has only non-negative entries then $D_p - A$ has rank at most $n-1$.*

One can show that a necessary condition for optimality of F_p is that $F_p - A$ is singular, i.e. the minimum of $F_p - A$ over $\{F_p : F_p - A \geq 0\}$ occurs at the boundary. Note

that while this condition is not sufficient for optimality, Lemma 1 asserts that D_p satisfies this necessary condition.

An $n \times n$ matrix A is regular circulant [21] if it has the form:

$$\begin{bmatrix} a_0 & a_1 & & \cdots & a_{n-1} \\ a_{n-1} & \ddots & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \ddots & a_1 \\ a_1 & \cdots & & a_{n-1} & a_0 \end{bmatrix},$$

and is symmetric regular circulant if it is regular circulant and it is symmetric about the diagonal, i.e. $a_{n-k} = a_k$, $k = 1, ..., n-1$. For a symmetric regular circulant matrix the vector $\frac{1}{\sqrt{n}}\underline{1}_+$ is an eigenvector, and $\sum_{i=1}^{n-1} a_i$ is the corresponding eigenvalue.

The following lemma is proven in Appendix B.

Lemma 2 *When $A = ((a_{ij}))$ is a regular symmetric circulant matrix and $a_{ij} \geq 0, \forall i, j$, the diagonal matrix D_1 minimizes the Frobenius norm $|||F - A|||_{\mathcal{F}}$ over all diagonal F which satisfy $F - A \geq 0$.*

While D_p may not always be optimal for non-circulant matrices A for $p \geq 1$, we can easily show that if $m > n$, then $|||D_m - A|||_{\mathcal{F}} = |||D_q - A|||_{\mathcal{F}} - 2\sum_{j=q+1}^{m}\sum_{k=1}^{n-q}(a_{k,k+q})^2$, so that $|||D_m - A|||_{\mathcal{F}} \leq |||D_q - A|||_{\mathcal{F}}$. Thus use of D_p for $p > 1$ should have as good or better average convergence rate than the diagonal matrix D_1.

Geometric Series with Over-Relaxation

The speed of convergence of the GS algorithm can be improved significantly by introducing an over-relaxation parameter ψ [31]. Assume that the splitting matrix F satisfies $F - A \geq 0$. The GS algorithm with over-relaxation (GSOR) is specified by:

$$\begin{aligned} \underline{\beta}_{GSOR}^{(k+1)} &= \underline{\beta}_{GSOR}^{(k)} - \psi F^{-1}(A\underline{\beta}_{GSOR}^{(k)} + \underline{e}_1) \\ &= (I - \psi F^{-1}A)\underline{\beta}_{GSOR}^{(k)} - \psi F^{-1}\underline{e}_1. \end{aligned} \tag{2.15}$$

The parameter $\psi \in (0, 2]$ can be chosen to optimize the asymptotic convergence rate by minimizing $\rho(I - \psi F^{-1}A)$. The optimal choice of ψ that results in the fastest asymptotic rate

of convergence depends on the maximum magnitude eigenvalue $|\lambda^M|_{\max}$ of $M = I - F^{-1}A$:

$$\psi_{asymptotic} = \frac{2}{2 - |\lambda^M|_{\max}}. \tag{2.16}$$

Note that $\psi_{asy} \in [1,2)$ so that the eigenvalues of $I - \psi F^{-1}A$ are spread over the interval (-1,1) resulting in a non-monotonic convergent GS algorithm. Since $|\lambda^M|_{\max}$ is expensive to compute, an alternative strategy for optimizing the GSOR (2.15) is to optimize the average convergence rate by minimizing $|||F - \psi A|||_{\mathcal{F}}$. This results in the easily computed value of ψ:

$$\psi_{average} = \frac{\sum_{i,j} f_{ij} a_{ij}}{\sum_{i,j} (a_{ij})^2}, \tag{2.17}$$

where f_{ij} and a_{ij} denote elements of the matrices F and A.

Computational Issues

The most expensive computations within the GS algorithm are the two matrix-vector multiplications in (2.4). We can rewrite (2.4) as $\underline{\beta}^{(k+1)} = \underline{\beta}^{(k)} + \underline{w}$, where $\underline{w} = F^{-1}(F_Y \underline{\beta}^{(k)} + \underline{e}_1)$. While computing $\underline{t} = F_Y \underline{\beta}^{(k)} + \underline{e}_1$ requires $2n^2 + n$ flops, computing $\underline{w} = F^{-1}\underline{t}$ requires only n flops if F is a diagonal splitting matrix. The use of banded diagonal splitting matrices $F = D_p$ does not result in a significant overhead in flops per iteration if we use the Cholesky LU decomposition [22, 29] for D_p:

$$D_p = L_p L_p^T,$$

where L_p denotes a lower triangular matrix. Since D_p is a banded diagonal, L_p has only p non-zero lower diagonals. In particular for $p \ll n$ the Cholesky decomposition requires only $\mathcal{O}(p^2 n)$ flops [29]. If we let $\underline{t} = F_Y \underline{\beta}^{(k)} + \underline{e}_1$ in (2.4), and $\underline{w} = D_p^{-1}\underline{t}$, then $\underline{w}$ can be computed in $\mathcal{O}(p^2 n)$ flops by solving

$$L_p L_p^T \underline{w} = \underline{t}$$

using back-substitution.

When F_Y is of the factored form $F_Y = A^T \Gamma A$ with Γ an $m \times m$ diagonal, and A is an $m \times n$ matrix, implementation of the recursion (2.4) requires $\mathcal{O}(n \times m)$ bytes of storage as compared to $\mathcal{O}(n^2)$ required for general F_Y. When $n \ll m$ or when A is sparse this corresponds to a significant memory savings.

2.5 Conjugate Gradient

The often slow convergence of the GS algorithm led us to consider using over-relaxation, which improves the convergence rate at the expense of monotonicity. Having sacrificed monotonicity, it is natural to ask whether there are other iterative methods with still faster convergence. The conjugate gradient (CG) method can be used to solve the linear system of equations:

$$A\,\underline{u} = \underline{b}, \tag{2.18}$$

for a non-singular matrix A. The CG algorithm is an iterative method, but for solving linear equations it converges to the exact solution in a finite number of iterations when run with infinite precision.

If we substitute $F_Y = A$, and $\underline{b} = \underline{e}_1 = [1, 0, 0, ..., 0]^T$ in (2.18), then the solution to the system of equations (2.18) will be the first column of the inverse of F_Y^{-1}. The following CG algorithm is only tailored to symmetric, positive definite matrices F_Y [29].

CG Algorithm

INITIALIZE: $\underline{u}^{(0)} = \underline{0}$; $\underline{r}^{(0)} = \underline{e}_1$

RECURSION:

$$\begin{aligned}
\alpha^{(i)} &= \begin{cases} 0 & i = 0 \\ \frac{<\underline{r}^{(i)}, \underline{r}^{(i)}>}{<\underline{r}^{(i-1)}, \underline{r}^{(i-1)}>} & i > 0 \end{cases} \\
\underline{p}^{(i)} &= \underline{r}^{(i)} + \alpha^{(i)}\, \underline{p}^{(i-1)} \\
\lambda^{(i)} &= \frac{<\underline{r}^{(i)}, \underline{r}^{(i)}>}{<\underline{p}^{(i)}, F_Y\, \underline{p}^{(i)}>} \\
\underline{u}^{(i+1)} &= \underline{u}^{(i)} + \lambda^{(i)}\, \underline{p}^{(i)} \\
\underline{r}^{(i+1)} &= \underline{r}^{(i)} - \lambda^{(i)}\, F_Y\, \underline{p}^{(i)}
\end{aligned}$$

BOUND APPROXIMATION:

$$\eta^{(i)} = \underline{e}_1^T\, \underline{u}^{(i)}. \tag{2.19}$$

The iterations are terminated when the residual error $\|r^{(i)}\|_2 = \|\underline{b} - F_Y\, \underline{u}^{(i)}\|_2$ is less than a user specified tolerance.

In the CG algorithm the direction vectors $\{\underline{p}^{(i)}\}_{i=1}^{n}$ are updated in such a way that $\underline{r}^i$ form an orthogonal basis for $\mathbb{R}^n$. With infinite precision arithmetic, this condition guarantees that $F_Y \underline{u}^{(n)} = \underline{b}$, where $\underline{u}^{(n)}$ is obtained at the n-th and final iteration in the CG algorithm. Each iteration of the algorithm involves a matrix-vector multiplication $F_Y \underline{p}^{(i)}$, requiring $\mathcal{O}(n^2)$ flops and therefore if the algorithm is stopped well before n iterations, an $\mathcal{O}(n^2)$ approximation to the CR bound is obtained. The convergence rate of the CG algorithm depends on the condition number of F_Y and its eigenvalue spread [22]. The speed of convergence of the CG algorithm can be further increased by pre-conditioning [29].

While the CG algorithm is guaranteed to give exact solution in a maximum of n iterations, the quantity $\eta^{(i)}$ (2.19) does not converge monotonically to the bound $[F_Y^{-1}]_{11}$. This means that if stopped at some iteration $k < n$ then the quantity $\eta^{(i)}$ calculated by the CG algorithm might not be a valid lower bound on the variance of an unbiased estimator $\widehat{\theta}_1$.

2.6 Recursive Methods to Approximate the Pseudo-Inverse of a Matrix

A square non-negative definite symmetric matrix F_Y of rank $r \leq n$ has an eigen-decomposition:

$$F_Y = \sum_{i=1}^{r} \sigma_i \, \underline{u}_i \underline{u}_i^T, \tag{2.20}$$

where $\sigma_1 \geq \sigma_2, ..., \geq \sigma_r > 0$ are the positive eigenvalues of F_Y arranged in decreasing order, and $\{\underline{u}_i\}_{i=1}^{r}$ is an orthonormal set of eigenvectors of F_Y. The Moore-Penrose pseudo-inverse, denoted by $F_Y^{\dagger}$, is given by [22, 29]:

$$F_Y^{\dagger} = \sum_{i=1}^{r} \frac{1}{\sigma_i} \underline{u}_i \underline{u}_i^T. \tag{2.21}$$

The Moore-Penrose pseudo-inverse provides an unbiased CR bound for the case of singular FIM F_Y [11]. Direct methods for calculation of the pseudo-inverse require performing an eigen-decomposition of F_Y. Recursive techniques for the calculation of the pseudo-inverse [17, 18, 19, 20] are computationally expensive: $\mathcal{O}(n^3)$ flops. In this section we give two iterative methods of approximation to the Moore-Penrose pseudo-inverse for use in approximating a 1×1 sub-block of $F_Y^{\dagger}$ requiring only $\mathcal{O}(n^2)$ flops per iteration.

2.6.1 Method 1: Perturbation

This method is based on adding a full rank matrix perturbation to F_Y to make it invertible. The objective is to keep the perturbation parameter ϵ as small as possible so that the approximation does not deviate too much from the actual solution. The approximation is based on the standard relation (2.22) below, given in [22]:

Theorem 1 *Let F_Y be an $n \times n$, non-negative definite symmetric matrix with rank $r \leq n$. Then the pseudo-inverse of F_Y given in (2.21), has the representation:*

$$F_Y^+ = \lim_{\epsilon \to 0^+} G_1(\epsilon),$$

where

$$G_1(\epsilon) \stackrel{def}{=} (F_Y^T F_Y + \epsilon I)^{-1} F_Y^T, \tag{2.22}$$

$\epsilon > 0$ is a scalar, and I is the $n \times n$ identity matrix [22]. Furthermore for all $\epsilon > 0$,

$$F_Y^+ \geq G_1(\epsilon). \tag{2.23}$$

Proof: We have,

$$\begin{aligned} F_Y^+ - G_1(\epsilon) &= \sum_{i=1}^{r} \frac{1}{\sigma_i} \underline{u}_i \underline{u}_i^T - \sum_{i=1}^{r} \frac{\sigma_i}{(\sigma_i^2 + \epsilon)} \underline{u}_i \underline{u}_i^T & (2.24) \\ &= \sum_{i=1}^{r} \frac{\epsilon}{\sigma_i(\sigma_i^2 + \epsilon)} \underline{u}_i \underline{u}_i^T & (2.25) \\ &\geq 0. \end{aligned}$$

Clearly, $\lim_{\epsilon \to 0} G_1(\epsilon) = F_Y^+$, and the theorem follows.

∎

From Theorem 1, $G_1(\epsilon)$ can be viewed as an approximation to the pseudo-inverse of F_Y based on Tikhonov regularization which has the property (2.23) that $G_1(\epsilon)$ is a valid lower bound on $\text{cov}_{\underline{\theta}}(\underline{\hat{\theta}})$ for all $\epsilon > 0$.

The choice of ϵ in Theorem 1 affects the error of approximation $F_Y^+ - G_1(\epsilon)$. For example if the trace norm $|||F_Y^+ - G_1(\epsilon)|||_T = \text{trace}(F_Y^+ - G_1(\epsilon))$ is small the approximation error will be small due to non-negative definiteness of $F_Y^+ - G_1(\epsilon)$. In the following discussion

we will study the choice of ϵ in terms of the trade-off between the error of approximation, and the condition number of $F_Y^T F_Y + \epsilon I$.

We want ϵ to be as small as possible for best approximation error $F_Y^+ - G_1(\epsilon)$. On the other hand for singular F_Y, we can not make ϵ too small, or else the matrix inversion $(F_Y^T F_Y + \epsilon I)^{-1}$ in (2.22) becomes ill-conditioned. Therefore, the choice of ϵ, and hence the error of approximation, must be driven by the maximum tolerable condition number c of $F_Y^T F_Y + \epsilon I$, where the condition number of a symmetric positive semi-definite matrix A is defined as the ratio of the maximum and minimum eigenvalues of A. The condition number of $F_Y^T F_Y + \epsilon I$ must satisfy:

$$\frac{\sigma_1^2 + \epsilon}{\epsilon} \approx \frac{\sigma_1^2}{\epsilon} \leq c. \tag{2.26}$$

One can find σ_1 using the power method [22, 29]. Thus choosing $\epsilon \geq \frac{\sigma_1^2}{c}$ guarantees that the condition number of $F_Y^T F_Y + \epsilon I$ lies below the maximum tolerable condition number c, and that the matrix can be inverted without running into precision problems.

Since in many cases the maximum eigenvalue σ_1 is not available, ϵ can be chosen more conservatively as:

$$\epsilon \geq \frac{\sum_{i=1}^{r} \sigma_i^2}{c} = \frac{\text{trace}(F_Y^T F_Y)}{c}. \tag{2.27}$$

The above choice of ϵ satisfies (2.26).

For $\epsilon = \frac{\sigma_1^2}{c}$, the trace error of approximation, denoted δ_p, can be easily computed.

$$\begin{aligned} \delta_p &\stackrel{def}{=} |||F_Y^+ - G_1(\epsilon)|||_T \\ &= \sum_{i=1}^{r} \frac{\epsilon}{\sigma_i(\sigma_i{}^2 + \epsilon)} \\ &= \sum_{i=1}^{r} \frac{\frac{\sigma_1^2}{c}}{\sigma_i(\sigma_i{}^2 + \frac{\sigma_1^2}{c})}. \end{aligned}$$

We can upper bound δ_p by $\delta_p^{\max}$:

$$\delta_p \leq \delta_p^{\max} = r \frac{\frac{\sigma_1^2}{c}}{\sigma_r(\sigma_r{}^2 + \frac{\sigma_1^2}{c})}.$$

2.6.2 Recursive Algorithm for Perturbation Method

Both the GS and the CG algorithms can be readily adapted to the perturbation method by the following steps:

1. Compute the column $\underline{\zeta}_1 = (F_Y^T F_Y + \epsilon I)^{-1}\underline{e}_1$ using either the GS or the CG method of Section 2.3 or Section 2.5 respectively.

2. Approximate the 1-1 element of the pseudo-inverse of F_Y: $[G_1(\epsilon)]_{11} = \underline{e}_1^T F_Y \underline{\zeta}_1$.

If GS is to be used in step 1, a splitting matrix F can be chosen such that $F \geq (F_Y^T F_Y + \epsilon\ I)$. The resultant un-relaxed GS algorithm has monotonic convergence since $\underline{e}_1^T(\underline{\zeta}_1^{(k)} - \underline{\zeta}_1^{(k+1)}) = \underline{e}_1^T \left([I - F^{-1}(F_Y^T F_Y + \epsilon I)]^{k+1} F^{-1}\right) \underline{e}_1 \geq 0$. Therefore, the algorithm generates a sequence of approximations $\underline{\zeta}_1^{(k)}$ which monotonically converge to $\underline{e}_1^T G_1(\epsilon)\underline{e}_1 \leq \underline{e}_1^T F_Y^+ \underline{e}_1$ as $k \to \infty$.

The GS recursions for the perturbation method are given by:

$$\underline{\beta}^{(k+1)} = \underline{\beta}^{(k)} - F^{-1}(\ F_Y^T F_Y\ \underline{\beta}^{(k)} + \epsilon \underline{\beta}^{(k)} + \underline{e}_1\ ).$$

If we denote $\underline{t} = F_Y\ \underline{\beta}^{(k)}$, then we have an additional matrix-vector multiplication ($2n^2$ flops), $F_Y^T \underline{t}$, when compared with (2.4) for the computation of F_Y^{-1}. Similarly, the CG algorithm has an additional $2n^2$ flops per iteration when applied to matrix perturbations via computation of the solutions $\underline{\zeta} = \underline{\zeta}_1$ to $(F_Y^T F_Y + \epsilon I)\underline{\zeta} = \underline{e}_1$.

The choice of ϵ directly affects the asymptotic rate of convergence of the GS and the CG algorithms. For the GS algorithm the rate of convergence depends on the matrix $M = M^* = [I - F^{-1}(F_Y^T F_Y + \epsilon I)]$. Since F_Y is rank deficient the spectral radius $\rho(M^*)$ is equal to $1 - O(\epsilon)$, which is close to unity for small ϵ resulting in slow asymptotic rate of convergence. Thus we see that there is a trade-off between the asymptotic error of approximation $|||F_Y^+ - G_1(\epsilon)|||_T$ and the asymptotic rate of convergence determined by $\rho(M^*)$. In particular we want ϵ to be small for good error of approximation, while ϵ should be large for fast asymptotic rate of convergence. Ideally, we would like the rate of convergence to be rapid for small ϵ. The following method reduces the effect of perturbation by differencing.

2.6.3 Method 2: Differencing Method

Theorem 2 *Let F_Y be an $n \times n$, non-negative definite symmetric matrix with rank $r \leq n$. Then the pseudo-inverse of F_Y given in (2.21), has the representation:*

$$F_Y^+ = \lim_{\epsilon \to 0+} G_2(\epsilon),$$

where,

$$G_2(\epsilon) \stackrel{def}{=} \frac{a}{a-b}(F_Y + a\,\epsilon I)^{-1} - \frac{b}{a-b}(F_Y + b\,\epsilon I)^{-1}, \tag{2.28}$$

$\epsilon > 0$ *is a scalar,* $a, b \in (0, \infty)$, $a \neq b$, *and* I *is the* $n \times n$ *identity matrix. Furthermore for all* $\epsilon > 0$,

$$F_Y^+ \geq G_2(\epsilon). \tag{2.29}$$

Note that due to (2.29) the CR bound approximated using Theorem 2 will be a valid lower bound.

Proof: Similar to the proof of Theorem 1, one can show:

$$\begin{aligned} F_Y^+ - G_2(\epsilon) &= \sum_{i=1}^{r} \frac{1}{\sigma_i}\underline{u}_i\underline{u}_i^T - \sum_{i=1}^{r} \frac{a}{(a-b)(\sigma_i + a\epsilon)}\underline{u}_i\underline{u}_i^T + \sum_{i=1}^{r} \frac{b}{(a-b)(\sigma_i + b\epsilon)}\underline{u}_i\underline{u}_i^T \\ &= \sum_{i=1}^{r} \frac{\epsilon\,(a\,\sigma_i + b\,\sigma_i + a\,b\,\epsilon)}{\sigma_i\,(\sigma_i + a\,\epsilon)\,(\sigma_i + b\,\epsilon)}\underline{u}_i\underline{u}_i^T \qquad (2.30) \\ &\geq 0. \qquad (2.31) \end{aligned}$$

Clearly $\lim_{\epsilon \to 0} G_2(\epsilon) = F_Y^+$, and the theorem follows.

■

Following the approach outlined for the perturbation method in section 2.6.1, we want to choose ϵ such that the condition numbers of $(F_Y + a\epsilon I)$ and $(F_Y + b\epsilon I)$ are less than the maximum tolerable condition number c. The condition numbers of $(F_Y + \epsilon I)$ and $(F_Y + b\epsilon I)$ must satisfy:

$$\frac{\sigma_1 + a\epsilon}{\epsilon}, \frac{\sigma_1 + b\epsilon}{\epsilon} \approx \frac{\sigma_1}{\epsilon} \leq c,$$

for $\frac{\sigma_1}{\epsilon} \gg a, b$, a situation satisfied in most practical applications. Hence by choosing $\epsilon \geq \frac{\sigma_1}{c}$, the matrix inversions $(F_Y + a\epsilon I)^{-1}$ and $(F_Y + b\epsilon I)^{-1}$ will be well conditioned. For $\epsilon = \frac{\sigma_1}{c}$, the asymptotic error of approximation δ_d associated with the differencing method is:

$$\begin{aligned} \delta_d &= |||F_Y^+ - G_2(\epsilon)|||_{\mathcal{T}} \\ &= \sum_{i=1}^{r} \frac{\epsilon\,((a+b)\,\sigma_i + a\,b\,\epsilon)}{\sigma_i\,(\sigma_i + a\,\epsilon)\,(\sigma_i + b\,\epsilon)}, \end{aligned}$$

and since δ_d is monotone in ϵ,

$$\delta_d \leq \delta_d^{\max} = r\,\frac{\frac{\sigma_1}{c}\,((a+b)\,\sigma_r + a\,b\,\frac{\sigma_1}{c})}{\sigma_r\,(\sigma_r + a\,\frac{\sigma_1}{c})\,(\sigma_r + b\,\frac{\sigma_1}{c})}. \tag{2.32}$$

Since in many cases σ_1 is not available ϵ can be chosen conservatively as:

$$\epsilon \geq \frac{\text{trace}(F_Y)}{c} = \frac{\sum_{i=1}^{r} \sigma_i}{c}. \tag{2.33}$$

Recursive Algorithm for Differencing Method

Using the GS or the CG algorithms, the approximation $[G_2(\epsilon)]_{11}$ to $[F_Y^+]_{11}$ can be computed by the following steps.

1. Approximate columns $\underline{\gamma}_1 = (F_Y + a\epsilon I)^{-1}\underline{e}_1$ and $\underline{\gamma}_2 = (F_Y + b\epsilon I)^{-1}\underline{e}_1$ using either the GS or the CG algorithms of Section 2.3 or Section 2.5 respectively.
2. Use approximation $[G_2(\epsilon)]_{11} \approx \underline{e}_1^T \left(\frac{a}{a-b}\underline{\gamma}_1 - \frac{b}{a-b}\underline{\gamma}_2\right)$ to approximate $[F_Y^+]_{11}$.

If a monotonic GS algorithm is to be used in step 1, two splitting matrices F_1 and F_2 must be chosen such that $F_1 > (F_Y + a\epsilon I)$, and $F_2 > (F_Y + b\epsilon I)$ in step 1. The asymptotic rate of convergence depends on the choice of ϵ in $M_1^* = I - F_1^{-1}(F_Y + \epsilon I)$, and $M_2^* = I - F_2^{-1}(F_Y + \epsilon I)$.

2.6.4 Comparison of Perturbation and Differencing Methods

If for a fixed condition number c we choose $\epsilon_d = \frac{\sigma_1}{c}$ for the differencing method, and $\epsilon_p = \frac{\sigma_1^2}{c}$ for the perturbation method, then the difference between the error bounds $\delta_p^{\max}$ and $\delta_d^{\max}$ is:

$$\begin{aligned}
\delta_p^{\max} - \delta_d^{\max} &= r\frac{\frac{\sigma_1^2}{c}}{\sigma_r({\sigma_r}^2 + \frac{\sigma_1^2}{c})} - r\frac{\frac{\sigma_1}{c}\left((a+b)\,\sigma_r + a\,b\,\frac{\sigma_1}{c}\right)}{\sigma_i\,(\sigma_r + a\,\frac{\sigma_1}{c})\,(\sigma_r + b\,\frac{\sigma_1}{c})} \\
&= \frac{r\sigma_1}{c}\left[\frac{\sigma_1}{{\sigma_r}^2 + \frac{\sigma_1^2}{c}} - \frac{\sigma_1}{\frac{{\sigma_r}^2}{(a+b)\frac{\sigma_r}{\sigma_1} + \frac{ab}{c}} + \frac{\sigma_1^2}{c}}\right].
\end{aligned}$$

From which it follows that if $(a+b)\frac{\sigma_r}{\sigma_1} + \frac{ab}{c} < 1$ then the asymptotic error of approximation of the differencing method is expected to be better than that of the perturbation method. In particular this will hold when the maximum tolerable condition number $c \gg 1$ and the ratio of the largest to smallest non-zero eigenvalues of F_Y satisfy $\frac{\sigma_1}{\sigma_r} \gg 1$.

While the GS implementation of the perturbation method gives a monotonically convergent algorithm, the GS implementation of the differencing method does not converge monotonically due to the differencing of the two monotone series. Thus unlike the other

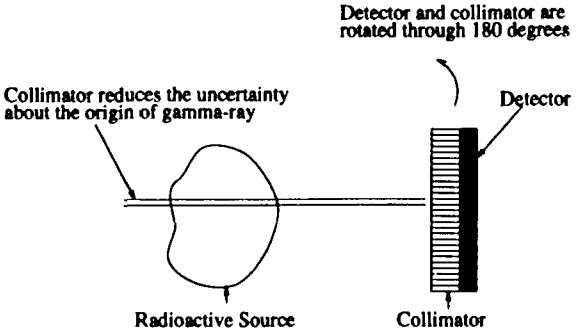

Figure 2.1: A SPECT system

recursive algorithms, the GS implementation of the perturbation method can be terminated at any iteration to obtain a valid lower bound on estimator variance.

2.7 Applications

In this section we apply the recursive algorithms described in the previous section to: 1) Single Photon Emission Computed Tomography (SPECT), and 2) one-dimensional discrete deconvolution. For more details on the SPECT system see chapter 4.

2.7.1 Single Photon Emission Computed Tomography

A SPECT system consists of three basic components: 1) a source of γ-ray photons, 2) a photon sensitive γ-ray detector and, 3) a γ-ray collimator. The function of the collimator is to reduce the uncertainty associated with the emission location of a detected γ-ray to a line or a strip in the field of view (Figure 2.1). During the image scan, the γ-ray detector is rotated through small steps 180^o around the source. A γ-ray photon passing through one of the collimator slits at one of the rotation angles is counted as an event acquired in one 'detector bin'. For reconstruction the source domain is divided into n small regions, called pixels. The detection process is governed by Poisson statistics:

$$f_{\underline{Y}}(\underline{y};\underline{\theta}) = \prod_{j=1}^{d} \frac{\mu_j^{y_j}}{y_j!} e^{\mu_j}, \tag{2.34}$$

where θ_i is the average γ-ray intensity at the i-th pixel; $i = 1, ...n$, and μ_j is the average γ-ray intensity at the j-th detector bin; $j = 1, ..., d$:

$$\underline{\mu} = A\,\underline{\theta}, \tag{2.35}$$

where A is the $d \times n$ matrix that depends on the tomographic system [35].

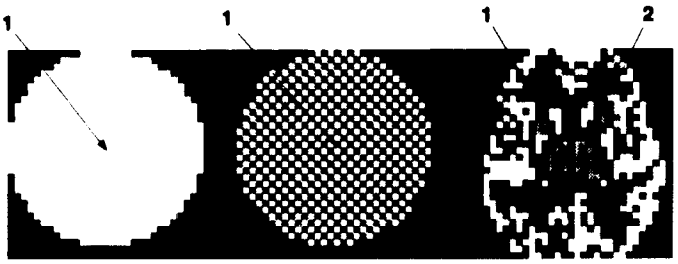

Figure 2.2: Objects used for numerical comparisons. The pixels of interest are the center pixel denoted by **1** and a neighboring pixel denoted by **2** in the brain image. The intensity ranges from 0 (black) to 2 (white). Objects are: Left: uniform, Middle: Checker board and, Right: Brain image.

The objective is to estimate the object intensity at each pixel $\underline{\theta} = [\theta_1, ..., \theta_n]^T$ given the set of observations Y. We make the standard assumption that $\theta_i \geq 0$ and $\mu_j \geq 0$, $i = 1, ..., n$, $j = 1, ..., d$. From (2.34) it can be easily shown that the FIM is of the factored form:

$$F_Y(\underline{\theta}) \;=\; A^T \left[\mathrm{diag}(\mu_i)\right]^{-1} A, \tag{2.36}$$

which ensures that the recursive CR bound algorithm requires only $\mathcal{O}(n \times d)$ units of storage. Furthermore, since A is generally sparse we obtain additional significant memory savings.

The FIM, denoted F_{EM}, for the complete-data set in emission tomography satisfies $F_{EM} - F_Y \geq 0$ and has the form [10]:

$$F_{EM} = \mathrm{diag}(\underline{1}^T A) \left[\mathrm{diag}(\underline{\theta})\right]^{-1}, \tag{2.37}$$

where $\mathrm{diag}(\underline{z})$ denotes a diagonal $n \times n$ matrix with the elements of the n-vector $\underline{z}$ along the diagonal. Note that since F_{EM} is a diagonal matrix, it is well suited for use as a splitting matrix in the GS algorithm.

2.7.2 Numerical Comparisons

All the simulations were done using MATLAB version 4.1 on an HP-755 computer running under the $HP - UX$ operating system.

GS Implemented with p-Diagonal Matrices Compared to CG

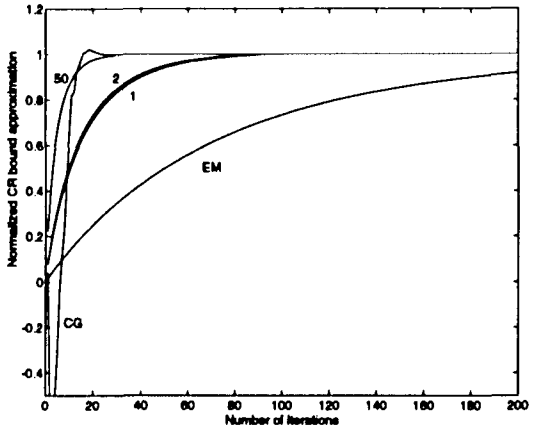

Figure 2.3: Effect of changing the number of diagonals for GS implemented with p-diagonal F for the low intensity pixel **1** for the brain phantom Figure 2.2. Numbers labeling each curve indicate the number of diagonals utilized in the splitting matrix D_p.

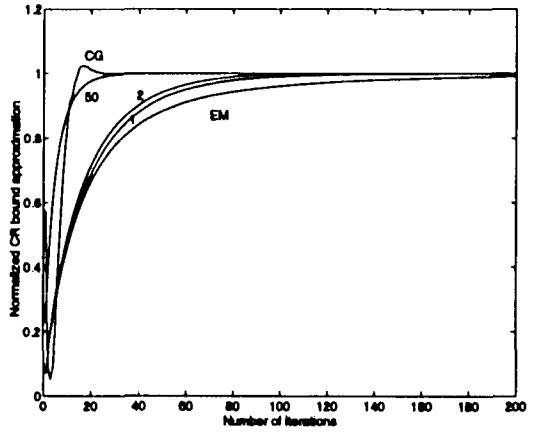

Figure 2.4: Same as Figure 2.3, except the pixel of interest was the high intensity pixel **2**.

Several two-dimensional objects were investigated (Figure 2.2). The high intensity pixels have a normalized intensity value of 2, the dark grey and the light grey pixels have normalized intensity 1 and 0.5 respectively. The black pixels are of zero intensity. The object lies in a field of view that is a disk of diameter 32 pixels. The total number of unknown parameters θ_i is $n = 716$ pixels.

Figure 2.3 and 2.4 display the effect of increasing the number of diagonals in the splitting matrix for the GS algorithm, and a comparison to CG. Figure 2.3 and 2.4 correspond to the brain image with two different pixels of interest, labeled 1 and 2 in Figure 2.2. In Table 2.1 the total number of iterations to converge, the Frobenius norm $|||\Delta|||_{\mathcal{F}} = |||F - F_Y|||_{\mathcal{F}}$ between the splitting matrix F and the FIM F_Y, the spectral norm $\rho(M) = \lambda^M_{\max}$ of $M = I - F^{-1}F_Y$, and the total number of Mflops to converge are given for the three image intensities and five different choices of the splitting matrices F for the GS algorithm. The splitting matrix D_λ is a diagonal matrix determined by numerically minimizing the spectral norm $\rho(M) = |\lambda^M|_{\max} = \lambda^M_{\max}$ for each F_Y corresponding to each image. Also shown is the convergence behavior of the CG algorithm. In all cases, the algorithms were initialized by the all zero vector corresponding to a zero image. For the results shown in Table 2.1 we used the convergence criterion (2.6) with a threshold $\gamma = 10^{-3}$, i.e. the algorithm was declared to have converged in N iterations when:

$$\frac{\|\underline{\beta}^{(k)} - F_Y^{-1}\underline{e}_1\|_2}{\|F_Y^{-1}\underline{e}_1\|_2} \leq \gamma = 10^{-3},$$

for more than 5 consecutive iterations $k = N, N+1, ..., N+4$, where $\underline{\beta}^{(k)}$ is the approximation to $F_Y^{-1}\underline{e}_1$ at the k-th iteration of either the GS or the CG algorithm. For the GS algorithm the total number of Mflops, shown in Table 2.1, was calculated taking account of a Cholesky LU decomposition of the corresponding splitting matrix F. The following comments on Table 2.1 are of interest:

- In each case the Frobenius norm $|||\Delta|||_{\mathcal{F}}$ is a much better predictor of the actual number of iterations required for GS convergence than is the root convergence factor $\lambda^M_{\max}$, which only predicts the asymptotic convergence rate. In particular, observe that while the use of $F = D_\lambda$ minimizes each $\lambda^M_{\max}$ for each of the three objects, our

Algorithm	Object											
	Uniform				Checker Board				Brain Image			
	$\|\|\|\Delta\|\|\|_{\mathcal{F}}$	$\lambda_{\max}^{M}$	No. of iter.	Mflops	$\|\|\|\Delta\|\|\|_{\mathcal{F}}$	$\lambda_{\max}^{M}$	No. of iter.	Mflops	$\|\|\|\Delta\|\|\|_{\mathcal{F}}$	$\lambda_{\max}^{M}$	No. of iter.	Mflops
$F = F_{EM}$	119.0	.9857	281	50.8	259.9	.9905	455	82.3	18634	.9999	583	105.4
$F = D_{\lambda}$	137.8	.9828	366	65.9	173.9	.9829	327	58.9	665.5	.9971	245	44.2
$F = D_1$	119.0	.9857	281	50.8	159.0	.9857	269	48.6	665.5	.9971	245	44.2
$F = D_2$	111.1	.9840	264	48.8	148.9	.9841	258	47.7	489.3	.9934	226	41.8
$F = D_{50}$	33.9	.9453	56	15.5	45.32	.9452	53	14.7	73.6	.9499	47	13.0
Conjugate Gradient			22	4.1			22	4.1			35	6.58

Table 2.1: Summary of convergence properties. $|||\Delta|||_{\mathcal{F}} = |||F - F_Y|||_{\mathcal{F}}$, $M = I - F^{-1}F_Y$.

choice of diagonal matrix $F = D_1$ introduced in (2.14) converges in fewer iterations and has smaller Frobenius error norm.

- For the complicated brain image the diagonal matrices D_1 and D_λ were found to be identical within the machine precision. On the other hand, for the simple uniform image, the diagonal matrices D_1 and F_{EM} are analytically identical, while D_1 and D_λ are quite different. For these cases F_{EM} and D_λ are not competitive with D_1 in terms of average convergence rate.

- As expected, since D_p gives smaller Frobenius norm $|||D_p - F_Y|||_{\mathcal{F}}$, for increasing p number of non-zero diagonals the number of iterations to converge decreases as p is increased. Unexpectedly, increasing p also decreases the root convergence factor $\lambda_{\max}^M$ resulting in more rapid asymptotic convergence rate. The use of 50 diagonals in the splitting matrix brings the required number of iterations of GS to within a factor of approximately 2 of the CG.

- The CG algorithm converges most rapidly, in under 35 iterations for all cases, and is approximately 2 to 3 times less computationally costly to implement as compared to the most rapid GS algorithm, indicated by $F = D_{50}$.

- The CG algorithm can have significant overshoots. In Figures 2.3 and 2.4 only 5% overshoot was observed. The magnitude of the overshoot was observed to depend on the initial condition. Thus if a monotone algorithm is desired, guaranteeing a valid bound at each iteration, the GS algorithm with $F = D_{50}$ is preferable to the CG algorithm.

2.7.3 Comparison of CG, GS and GS with Over-Relaxation

Figure 2.5 shows a comparison of the CG, the GS and the GSOR algorithms, the latter two implemented with the diagonal splitting matrix D_1 of (2.14), for pixel **1** in the brain phantom. To establish a benchmark for GSOR the over-relaxation parameter $\psi = 1.99$ was chosen using (2.16) for fastest asymptotic convergence. Since in general the use of (2.16) requires the impractical computation of the maximum eigenvalue of M, the asymptotic convergence of the implementable GSOR algorithm using (2.17) may be somewhat slower.

The GSOR exhibits a much faster initial rate of convergence when compared to the GS, however it oscillates around the point of convergence and fails to meet our convergence criterion of $\gamma = 10^{-3}$ even after the GS achieves convergence.

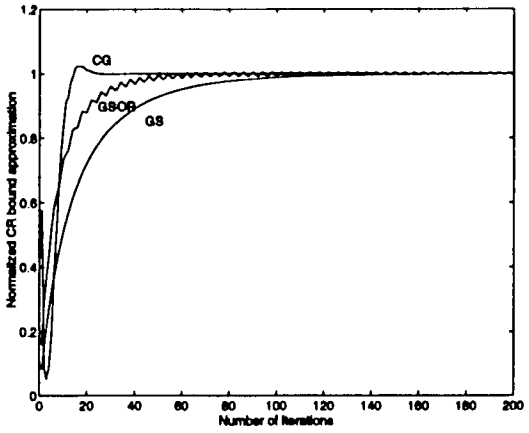

Figure 2.5: GS with over-relaxation, denoted GSOR, shows oscillatory effect.

2.7.4 One-Dimensional Discrete Deconvolution

We illustrate the ill-conditioned CR bound algorithms described in Section 2.6 for the problem of discrete deconvolution, specified by the model:

$$\underline{Y} = A\underline{\theta} + \underline{q}, \tag{2.38}$$

where $\underline{Y}$ is the blurred data, A corresponds to a convolution kernel, $\underline{\theta}$ is the parameter we want to deconvolve or estimate and $\underline{q}$ is additive white Gaussian noise with covariance matrix Σ. The blurring in the data can be due to several factors, e.g. imperfect measuring devices. We chose a discrete Gaussian kernel $a_{ij} = \frac{1}{\sqrt{2\pi}} e^{\frac{-(i-j)^2}{2w^2}}$ in these simulations. Philip [30] has remarked that for the discrete deconvolution problem the most ill-conditioned kernels have the form of discrete Gaussian kernels. The FIM F_Y is given by:

$$F_Y = A^T \mathrm{cov}(\underline{Y})^{-1} A = A^T \Sigma^{-1} A.$$

By choosing the width of the Gaussian kernel A we can vary the condition number of the FIM F_Y.

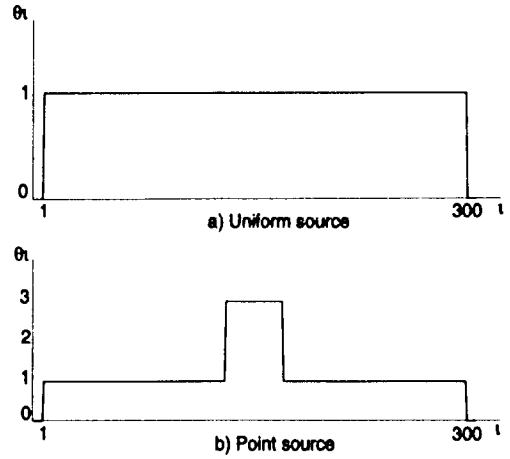

Figure 2.6: One dimensional sources used in discrete deconvolution problem.

2.7.5 Dependence of the Rate of Convergence on the Condition Number

First we consider a uniform source as shown in Figure 2.6 a). We used the same convergence criterion (2.6) as in the tomography example, i.e. $\gamma = 10^{-3}$, to terminate the iterations. The matrix A in (2.38) had dimension 300×300. Figure 2.7 indicates that

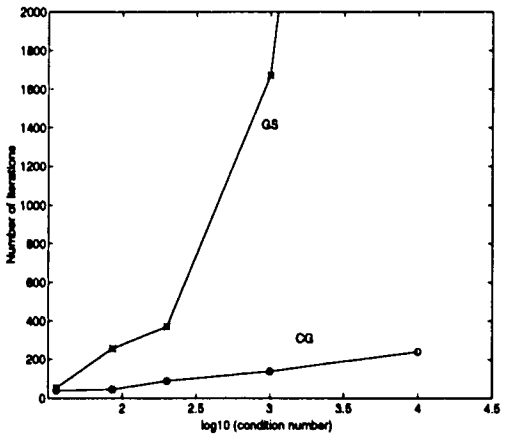

Figure 2.7: Effect of the condition number on the convergence rate for the GS and the CG algorithms for the uniform source. Matrices of dimension 300×300 were used.

the number of iterations required for convergence is more sensitive to the condition number for GS than for CG.

2.7.6 Performance of Pseudo-inverse Algorithm

In this section we consider the problem of super-resolution with an up-sampling factor of 2. This refers to a situation when the number of available independent observations are only a half the number of parameters we want to estimate. The A matrix has dimension 150×300. This problem is obviously under-determined. The width, w, of the Gaussian convolution kernel was varied from 5 to 32. The region of interest was the center pixel θ_{150} for the point source displayed in Figure 2.6. The value of ϵ was chosen using (2.27) and (2.33). The maximum allowable condition number used was $c = 10^{12}$. For the differencing method we chose a=1 and b=2 in (2.28). Figure 2.8 clearly indicates that differencing method

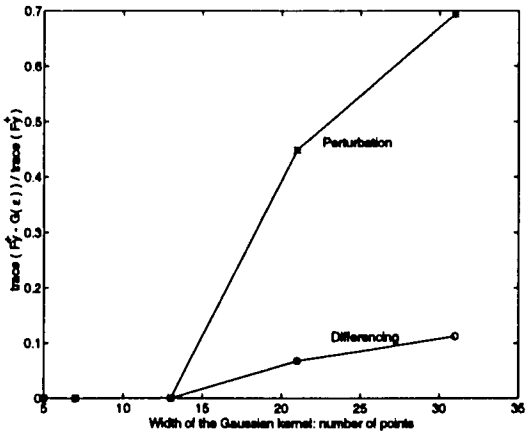

Figure 2.8: Accuracy of the perturbation method and differencing method for a singular FIM F_Y for the point source.

is more accurate for all resolutions as measured by the normalized error of approximation $\frac{\text{trace}(F_Y^+ - G(\epsilon))}{\text{trace}(F_Y^+)}$, where $G(\epsilon) = G_1(\epsilon)$ for the perturbation method and $G(\epsilon) = G_2(\epsilon)$ for the differencing method. The accuracy for both the methods degrades for wide convolution kernels, but the degradation is more pronounced in the perturbation method. Thus for ill-conditioned problems corresponding to wider convolution kernels and resulting in large but finite condition numbers, the differencing method is more accurate than the perturbation method.

To study the effect of ϵ on the normalized error of approximation $\frac{\text{trace}(F_Y^+ - G(\epsilon))}{\text{trace}(F_Y^+)}$ of the differencing and the perturbation method we fixed the width of the Gaussian kernel to 20 points. The accuracy of both the algorithms degrades gracefully over a wide range

of ϵ. In both cases the value of ϵ chosen using (2.27) and (2.33), indicated by vertical bars in Figure 2.9, is close to the knee of the curves. The differencing method is clearly more accurate than the perturbation method over the whole range of ϵ. The perturbation method fails to achieve the desired value of $\frac{\text{trace}(F_Y^+ - G(\epsilon))}{\text{trace}(F_Y^+)} \approx 0$ for all values of ϵ.

Non-monotone Convergence of Differencing Method

Figure 2.10 shows the non-monotone convergence curves of the differencing method. The width of the Gaussian kernel was 3 points.

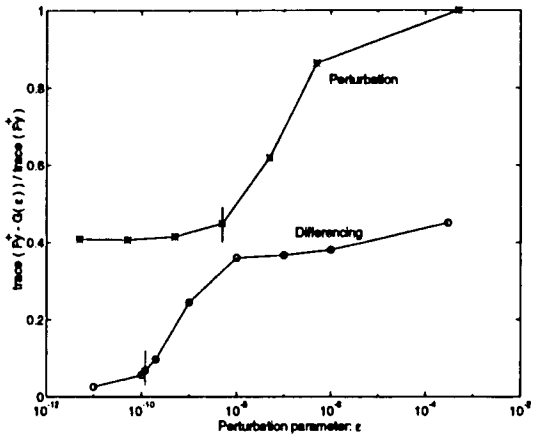

Figure 2.9: Performance comparison of the perturbation method and differencing method for varying ϵ for the point source. Vertical bar on both the plots corresponds to the choice of ϵ using (2.27) and (2.33).

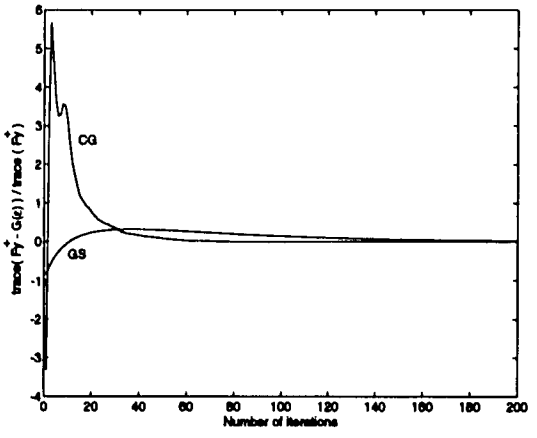

Figure 2.10: Differencing method has non-monotone convergence for both the GS and the CG.

CHAPTER 3

BIAS-VARIANCE TRADE-OFFS FOR PARAMETRIC ESTIMATION PROBLEMS USING THE UNIFORM CR BOUND

The computational algorithms for the unbiased CR bound presented in Chapter 2 are applicable to unbiased estimation studies. In this chapter we extend these algorithms to the uniform CR bound [14] which is applicable to biased estimators. It is then shown how to use the uniform CR bound to trace a curve separating the regions of achievable and unachievable performance in the bias-variance trade-off plane.

3.1 Introduction

As before we consider the problem of estimation of an n-dimensional parameter $\underline{\theta} = [\theta_1, ..., \theta_n]^T$ given an observation of a vector of random variables $\underline{Y}$ with probability density function (pdf) $f_{\underline{Y}}(\underline{y}; \underline{\theta})$. When an estimator $\underline{\widehat{\theta}}$ is biased, the mean-square error (MSE) is an important measure of precision of a scalar component $\widehat{\theta}_1$. It is well known that the MSE is a function of both the bias, denoted $\text{bias}_{\underline{\theta}}(\widehat{\theta}_1)$ and the variance, denoted $\text{var}_{\underline{\theta}}(\widehat{\theta}_1)$ of the scalar estimator:

$$\text{MSE}_{\underline{\theta}}(\widehat{\theta}_1) = \text{var}_{\underline{\theta}}(\widehat{\theta}_1) + \text{bias}^2_{\underline{\theta}}(\widehat{\theta}_1).$$

Obviously increases in MSE can be due to increases in either the bias or variance of $\widehat{\theta}_1$. Bias and variance are complementary in nature. While bias is due to 'mismatch' between the average value of the estimator and the true parameter, variance is due to statistical fluctuations in the estimator. There usually exists a tradeoff between bias and variance of the estimated parameter. For example in image reconstruction, implementation of the maximum likelihood algorithm with a smoothness penalty reduces the variance only at the

expense of introducing bias. Likewise for edge localization in image processing, spatial smoothing can reduce variance only at the price of increased localization bias. Different estimators can be effectively compared by plotting their performance on a bias-variance trade-off plane. This chapter provides a means for specifying achievable and unachievable regions in the bias-variance trade-off plane via an extension of the Cramer-Rao (CR) lower bound to biased estimators recently obtained in [14].

As explained in chapter 2, the conventional unbiased CR bound determines a lower bound on the variance of an unbiased estimator. Although this CR bound is known to be asymptotically achievable for a large number of independent observations, in practice, many parameter estimation algorithms are biased. In [14] a 'uniform' CR bound was obtained which is applicable to all biased estimators whose bias gradient length is less than a pre-specified threshold 'δ'. The uniform CR bound is more useful than the standard extension of the CR bound [33] in that the standard extension only applies to estimators with the same bias gradient function $\nabla_{\underline{\theta}}\text{bias}_{\underline{\theta}}(\underline{\hat{\theta}}) = \nabla_{\underline{\theta}} b$, where $\nabla_{\underline{\theta}} b$ is a fixed function. Since the uniform CR bound on the estimator variance σ^2 is indexed by the gradient length threshold δ, this bound specifies a curve over the δ–σ plane which separates the plane into achievable and unachievable regions of variance and bias gradient length. This plane is called the bias-gradient-length/variance trade-off plane. If a maximum tolerable bias gradient length is known or can be specified, e.g. based on estimation subsystem accuracy requirements, the uniform CR bound provides information about the minimum theoretically achievable variance of any estimator satisfying the constraint on the maximum tolerable bias gradient length. Alternatively a candidate estimator can be placed in the achievable region of the δ–σ plane where its closeness to the uniform CR bound curve provides an indication of the degree of optimality of the bias-variance trade-off characteristics of this estimator. Of course for an estimator to be placed on the δ–σ plane its variance and bias gradient length must be calculated or accurately estimated.

For most estimators computation of estimator variance, bias, and bias gradient are analytically intractable. For fixed $\underline{\theta}$, the standard method of moments technique for estimating bias and variance requires L repeated experiments, each generating a realization of $\underline{Y}$ from $f_{\underline{Y}}(\underline{y};\underline{\theta})$. However, the direct application of this method of bias gradient estimation

requires performing n sequences of L repeated experiments, each sequence being realizations of $\underline{Y}$ generated for one of n perturbations of different components of $\underline{\theta}$. This direct approach would be inefficient and impractical for a parameter space of large dimension n. We present a more efficient and accurate method for experimentally determining the bias gradient using a weighted sample average of the observed estimator bias, where the weights are equal to the score function $\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{y};\underline{\theta})$. The method only requires a single sequence of L experiments and can be the same one used to determine bias and variance of $\widehat{\theta}_1$.

We demonstrate the methods developed in this chapter for the following applications: 1) one-dimensional image restoration with a Gaussian convolution kernel, 2) one-dimensional edge localization in images.

3.2 Biased CR Bound

In Section 2.2 the unbiased CR bound on the variance $\widehat{\theta}_1$ was discussed. Under broad conditions the unbiased CR bound for non-singular FIM (2.2) is known to be asymptotically achievable for an increasing number of independent identically distributed observations $\underline{Y}_i$ [35]. Therefore the unbiased CR bound can provide useful information about the asymptotically achievable mean-square-error performance. However, the unbiased CR bound is not very useful for a small number of observations for which many estimation algorithms are biased and apparent bound violations may occur.

Motivating Example:

In Figure 3.1 the unbiased CR bound is plotted as a function of the collimator rotations for the single photon emission computed tomography (SPECT) tomograph described in Section 4.1.1. Also plotted is the variance of the weighted least-squares estimator (WLSE) [48, 43]. The variance of the WLSE is clearly less than the CR lower bound at all points. This is due to the fact that the weighted least-square estimator is biased. A biased version of the CR bound is needed to ensure that no apparent bound violation occurs in examples such as this.

For an estimator $\widehat{\theta}_1$ with mean function $m_1(\underline{\theta}) = E_{\underline{\theta}}(\widehat{\theta}_1)$ and bias function $b_1(\underline{\theta}) = m_1(\underline{\theta}) - \theta_1$, the following form of the so-called biased CR bound is available [11]:

$$\mathrm{var}_{\underline{\theta}}(\widehat{\theta}_1) \;\geq\; [\nabla_{\underline{\theta}} m_1(\underline{\theta})]\, F_Y^+(\underline{\theta})\, [\nabla_{\underline{\theta}} m_1(\underline{\theta})]^T$$

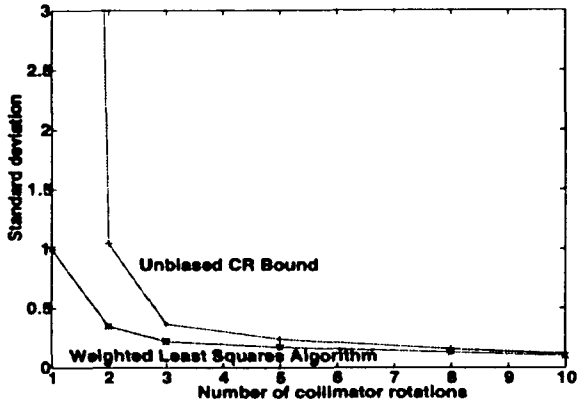

Figure 3.1: The regularized weighted least-squares estimator shows lower variance than the unbiased CR bound.

$$= \left[\underline{e}_1^T + \nabla_{\underline{\theta}} b_1(\underline{\theta})\right] F_Y^+ \left[\underline{e}_1^T + \nabla_{\underline{\theta}} b_1(\underline{\theta})\right]^T . \tag{3.1}$$

The biased CR bound (3.1) only applies to the class of estimators $\widehat{\theta_1}$ which have a particular bias gradient function $\nabla_{\underline{\theta}} b_1$. For example it can not be used to simultaneously bound the variance of a pair of estimators, which have different but perhaps acceptable biases. Thus the biased CR bound only applies to an unnatural class of estimators since seldom do two different estimators have the same bias gradient. A more natural class is the set of estimators whose bias gradient length is smaller than a specified threshold.

3.3 Uniform CR Bound

In [14] Hero presents a 'uniform' CR bound on the variance of $\widehat{\theta_1}$ for non-singular FIM F_Y. This bound is applicable to all biased estimators whose bias gradient length $\|\nabla_{\underline{\theta}} b_1\|$ is less than a pre-specified threshold,

$$\|\nabla_{\underline{\theta}} b_1\| \leq \delta < 1, \tag{3.2}$$

where $\|\underline{z}\| = \sum_{i=1}^n z_i^2$ for an n-element real vector $\underline{z}$. The following theorem extends the results of [14] by allowing both singular and non-singular FIM F_Y.

Theorem 3 *Let the pdf $f_{\underline{Y}}(\underline{y};\underline{\theta})$ have an associated FIM $F_Y = F_Y(\underline{\theta})$ and let $\widehat{\theta_1}$ be any estimator of θ_1 with bias $b_1(\underline{\theta})$ whose n-element bias gradient vector $\nabla_{\underline{\theta}} b_1 = \nabla_{\underline{\theta}} b_1(\underline{\theta})$ satisfies (3.2). Let the symmetric $n \times n$ projection $\mathcal{P}_{F_Y}$ be the matrix which orthogonally projects*

vectors in $\mathbb{R}^n$ onto the column space of the FIM F_Y, and define the n-element unit vector $\underline{e}_1 = [1, 0, ..., 0]^T$. Then the variance of $\hat{\theta}_1$ satisfies:

$$var_{\underline{\theta}}(\hat{\theta}_1) \geq B^o(\underline{\theta}, \delta), \tag{3.3}$$

where $B^o(\underline{\theta}, \delta)$ is equal to:

$$B^o(\underline{\theta}, \delta) = \begin{cases} 0, & \|\mathcal{P}_{F_Y}\underline{e}_1\| \leq \delta \\ [\underline{e}_1 + \underline{d}_{min}]^T F_Y^+ [\underline{e}_1 + \underline{d}_{min}], & \|\mathcal{P}_{F_Y}\underline{e}_1\| > \delta, \end{cases} \tag{3.4}$$

or equivalently,

$$B^o(\underline{\theta}, \delta) = \begin{cases} 0, & \|\mathcal{P}_{F_Y}\underline{e}_1\| \leq \delta \\ B^o(\underline{\theta}, 0) - \underline{e}_1^T [F_Y^+(\lambda I + F_Y^+)^{-1} F_Y^+] \underline{e}_1 - \lambda\delta^2, & \|\mathcal{P}_{F_Y}\underline{e}_1\| > \delta. \end{cases} \tag{3.5}$$

In (3.4)

$$\underline{d}_{min} = -[\lambda I + F_Y^+]^{-1} F_Y^+ \underline{e}_1, \tag{3.6}$$

where $\lambda > 0$ is given by non-negative solution of the following equation involving the monotone non-increasing, convex function $g(\lambda) \in [0, 1]$:

$$g(\lambda) = \underline{e}_1^T \left[F_Y^+ (\lambda I + F_Y)^{-2} F_Y^+\right] \underline{e}_1 = \underline{d}_{min}^T \underline{d}_{min} = \delta^2 \quad , \quad \lambda > 0. \tag{3.7}$$

If θ_1 is a parameter with finite non-zero Fisher information $[F_Y]_{11}$, then $g(\lambda)$ is a monotone decreasing and strictly convex function over $\lambda > 0$.

Proof: Our objective is to find a bound on the variance of $\hat{\theta}_1$ under the constraint $\|\nabla_{\underline{\theta}} b_1\|^2 \leq \delta^2 < 1$. We start with the bound (3.1) and minimize over the class of bias gradients (3.2).

$$\begin{aligned} \text{var}_{\underline{\theta}}(\hat{\underline{\theta}}_1) &\geq [\underline{e}_1 + \nabla_{\underline{\theta}}^T b_1]^T F_Y^+ [\underline{e}_1 + \nabla_{\underline{\theta}}^T b_1] \\ &\geq \min_{b_1 : \|\nabla_{\underline{\theta}} b_1\| \leq \delta} [\underline{e}_1 + \nabla_{\underline{\theta}}^T b_1]^T F_Y^+ [\underline{e}_1 + \nabla_{\underline{\theta}}^T b_1] \\ &= \min_{\underline{d} : \|\underline{d}\| \leq \delta} [\underline{e}_1 + \underline{d}]^T F_Y^+ [\underline{e}_1 + \underline{d}] \\ &= \min_{\underline{d} : \|\underline{d}\| \leq \delta} Q(\underline{d}), \end{aligned}$$

where $\underline{d} \in \mathbb{R}^n$ and,

$$Q(\underline{d}) = [\underline{e}_1 + \underline{d}]^T F_Y^+ [\underline{e}_1 + \underline{d}] \, .$$

An equivalent expression for the constrained minimization $\min_{\underline{d} : \|\underline{d}\| \leq \delta} Q(\underline{d})$ is obtained by performing an unconstrained minimization via introduction of an undetermined multiplier

$\lambda \geq 0$:

$$\min_{\underline{d}:\|\underline{d}\| \leq \delta} Q(\underline{d}) = \min_{\underline{d}} \left[Q(\underline{d}) + \lambda(\underline{d}^T \underline{d} - \delta^2) \right]. \tag{3.8}$$

First consider the class of solutions $\underline{d}$ to (3.8) for $\lambda = 0$, i.e. the case of inactive constraint for which $\min_{\underline{d}:\|\underline{d}\| \leq \delta} Q(\underline{d})$ is obviously zero. Since $\mathcal{P}_{F_Y}$ projects onto the column space of F_Y, $F_Y^+ = \mathcal{P}_{F_Y} F_Y^+ \mathcal{P}_{F_Y}$ and hence:

$$Q(\underline{d}) = \min_{\underline{d}} \left[\mathcal{P}_{F_Y} \underline{e}_1 + \mathcal{P}_{F_Y} \underline{d} \right]^T F_Y^+ \left[\mathcal{P}_{F_Y} \underline{e}_1 + \mathcal{P}_{F_Y} \underline{d} \right],$$

which is zero iff $\underline{d} = -\mathcal{P}_{F_Y} \underline{e}_1 + \underline{\phi}$, where $\underline{\phi}$ is an arbitrary vector such that $\mathcal{P}_{F_Y} \underline{\phi} = 0$. But for $\underline{d}$ to be the solution to the constrained minimization we must satisfy $\|\underline{d}\| \leq \delta$ so that, by orthogonality of $\mathcal{P}_{F_Y} \underline{e}_1$ and $\underline{\phi}$:

$$\delta \geq \|\underline{d}\| = \|\mathcal{P}_{F_Y} \underline{e}_1\| + \|\underline{\phi}\| \geq \|\mathcal{P}_{F_Y} \underline{e}_1\|.$$

Therefore $\min_{\underline{d}:\|\underline{d}\| \leq \delta} Q(\underline{d}) = 0$ iff $\|\mathcal{P}_{F_Y} \underline{e}_1\| \leq \delta$.

Next consider the class of solutions $\underline{d}$ to (3.8) for $\lambda > 0$, in which case the constraint $\|\nabla_{\underline{\theta}} b_1\| \leq \delta < 1$ is active, $\min_{\underline{d}:\|\underline{d}\| \leq \delta} Q(\underline{d}) > 0$, and $\|\mathcal{P}_{F_Y} \underline{e}_1\| > \delta$. Since $\lambda I + F_Y^+$ is positive definite, the completion of the square in (3.8) yields:

$$\begin{aligned} &\min_{\underline{d}} \left[Q(\underline{d}) + \lambda(\underline{d}^T \underline{d} - \delta^2) \right] \\ &\quad = \left[\underline{d} + (\lambda I + F_Y^+)^{-1} \underline{e}_1 \right]^T (\lambda I + F_Y^+) \left[\underline{d} + (\lambda I + F_Y^+)^{-1} \underline{e}_1 \right] \\ &\qquad + \underline{e}_1^T F_Y^+ \underline{e}_1 - \underline{e}_1^T \left[F_Y^+ (\lambda I + F_Y^+)^{-1} F_Y^+ \right] \underline{e}_1 - \lambda \delta^2. \end{aligned} \tag{3.9}$$

It is clear from (3.9) that

$$\underline{d} = \underline{d}_{min} = -(\lambda I + F_Y^+)^{-1} F_Y^+ \underline{e}_1$$

minimizes $Q(\underline{d})$, and the minimum is given by:

$$\begin{aligned} \min_{\underline{d}:\|\underline{d}\| \leq \delta} Q(\underline{d}) &= \left[\underline{e}_1 + \underline{d}_{min} \right]^T F_Y^+ \left[\underline{e}_1 + \underline{d}_{min} \right] \\ &= \underline{e}_1^T F_Y^+ \underline{e}_1 - \underline{e}_1^T \left[F_Y^+ (\lambda I + F_Y^+)^{-1} F_Y^+ \right] \underline{e}_1 - \lambda \delta^2 \\ &= B^o(\underline{\theta}; 0) - \underline{e}_1^T \left[F_Y^+ (\lambda I + F_Y^+)^{-1} F_Y^+ \right] \underline{e}_1 - \lambda \delta^2, \end{aligned}$$

as stated in (3.4) and (3.5). Now, since $\min_{\underline{d}:\|\underline{d}\| \leq \delta} Q(\underline{d})$ is monotone decreasing in δ, the minimum must occur at the boundary of the set $\{\underline{d} : \|\underline{d}\| \leq \delta\}$. Therefore for $\lambda > 0$, λ is

determined by the equality constraint:

$$\delta^2 = \underline{d}_{min}^T \underline{d}_{min} = \underline{e}_1^T \left[F_Y^\dagger (\lambda I + F_Y^\dagger)^{-2} F_Y^\dagger\right] \underline{e}_1 = g(\lambda). \tag{3.10}$$

Now the first and the second derivatives of $g(\lambda)$ are:

$$\begin{aligned} g'(\lambda) &= -2\underline{e}_1^T F_Y^\dagger (\lambda I + F_Y)^{-3} F_Y^\dagger \underline{e}_1 \\ g''(\lambda) &= 6\underline{e}_1^T F_Y^\dagger (\lambda I + F_Y)^{-4} F_Y^\dagger \underline{e}_1. \end{aligned}$$

Since $F_Y^\dagger (\lambda I + F_Y)^{-3} F_Y^\dagger$ is positive semi-definite for $\lambda > 0$, $g'(\lambda) \leq 0$ so that $g(\lambda)$ a is monotone non-increasing function over $\lambda > 0$. Also since $g''(\lambda) \geq 0$, $g(\lambda)$ is a convex cap function. Clearly, if $\underline{e}_1 \notin \mathcal{N}(F_Y)$, where $\mathcal{N}(A)$ denotes the null space of matrix A, then $g'(\lambda) < 0$, $g''(\lambda) > 0$ and $g(\lambda)$ is strictly monotone decreasing and strictly convex. Since $F_Y \underline{e}_1$ is the first column of F_Y this is satisfied whenever θ_1 has finite non-zero Fisher information.

∎

Comments:

- The orthogonal projection $\mathcal{P}_{F_Y}$ can be obtained either as $\mathcal{P}_{F_Y} = F_Y^\dagger F_Y$, or directly via the SVD of F_Y, as $\mathcal{P}_{F_Y} = \sum_{i=1}^r \underline{u}_i \underline{u}_i^T$, where r is the rank of F_Y and $\{\underline{u}_i\}_{i=1}^r$ are the orthonormal singular vectors associated with the non-zero singular values of F_Y.

- When F_Y is non-singular, $F_Y^\dagger = F_Y^{-1}$, $\|\mathcal{P}_{F_Y} \underline{e}_1\| = \|\underline{e}_1\| = 1 \geq \delta$ and Theorem 1 reduces to the result derived in [14]:

$$B(\underline{\theta}, \delta) = [\underline{e}_1 + \underline{d}_{min}]^T F_Y^{-1} [\underline{e}_1 + \underline{d}_{min}], \tag{3.11}$$

  where,

$$\underline{d}_{min} = -[I + \lambda F_Y]^{-1} \underline{e}_1,$$

  and $\lambda \geq 0$ is given by the unique non-negative solution of the following equation involving the monotone decreasing, strictly convex function $g(\lambda) \in [0, 1]$ over $\lambda \geq 0$:

$$g(\lambda) = \underline{d}_{min}^T \underline{d}_{min} = \underline{e}_1^T [I + \lambda F_Y]^{-2} \underline{e}_1 = \delta^2.$$

- A bias sensitivity index for the CR bound can be derived by studying the limit $\eta = \lim_{\delta \to 0} \frac{1}{\delta} \frac{B(\underline{\theta},0) - B(\underline{\theta},\delta)}{B(\underline{\theta};0)}$. For non-singular F_Y, it is shown in [14] that $\eta = 2\sqrt{1 + \underline{c}^T F_S^{-1} \underline{c}}$, where $\underline{c}$ is the first column of F_Y and F_S is the principal minor of $[F_Y]_{11}$.

- When F_Y is non-singular but has an ill-conditioned inverse, direct implementation of the bound $B(\underline{\theta}, \delta)$ in (3.11) is numerically unstable. A simple algebraic manipulation in (3.11) yields the equivalent form for $B(\underline{\theta}; \delta)$:

$$B(\underline{\theta}, \delta) = \lambda^2 \underline{e}_1^T \left[I + \lambda F_Y\right]^{-1} F_Y \left[I + \lambda F_Y\right]^{-1} \underline{e}_1. \tag{3.12}$$

  As contrasted with (3.11), whose computation requires inversion of the possibly ill-conditioned F_Y, the form (3.12) requires inversion of the matrix $I + \lambda F_Y$. Since $\lambda \geq 0$ and $F_Y \geq 0$, the inversion $[I + \lambda F_Y]^{-1}$ is well conditioned and (3.12) gives a numerically stable implementation of $B(\underline{\theta}; \delta)$.

- In Theorem 1, $\underline{d}_{min}$ defined in (3.6) is an optimal bias gradient in the sense that it minimizes the biased CR bound (3.1) over all vectors $\nabla_{\underline{\theta}} b_1$ under the constraint $\delta < 1$. Since the bound is specified in terms of the length of the bias gradient, it is independent of the particular estimator bias as long as the bias gradient length constraint is satisfied. From the proof of Theorem 1 we see that if $\|\mathcal{P}_{F_Y} \underline{e}_1\| \leq \delta$, then $\underline{d}_{min}$ can be taken as $\underline{d}_{min} = -\mathcal{P}_{F_Y} \underline{e}_1 + \underline{\phi}$, where $\underline{\phi}$ is any vector satisfying $F_Y \underline{\phi} = \underline{0}$, and $\|\underline{\phi}\| \leq \delta - \|\mathcal{P}_{F_Y} \underline{e}_1\|$. Thus for the singular case there exist a large number of optimal bias gradients for which the uniform CR bound is equal to the trivial bound $\mathrm{var}_{\underline{\theta}}(\widehat{\theta}_1) \geq 0$.

- By setting the bias gradient constraint $\|\nabla_{\underline{\theta}} b_1\| = \delta = 0$ we obtain the unbiased CR bound $B^o(\underline{\theta}, 0) = \underline{e}_1^T F_Y^+ \underline{e}_1$.

- Hero [14] has shown that if F_Y is non-singular, and the unbiased CR bound is achievable by an unbiased estimator $\widehat{\underline{\theta}}^*$, then one can construct an estimator that locally achieves the uniform bound by introducing a small amount of bias in $\widehat{\underline{\theta}}^*$. However, since unbiased estimators may not exist for singular F_Y, the uniform CR bound for singular F_Y may not be achievable.

- The range $\delta \in [0, 1)$ contains all the estimator bias gradient lengths of interest. In particular $\delta^2 = 1$ corresponds to $\nabla_{\underline{\theta}} b_1 = [-1, 0, ..., 0]$, identifiable as the bias gradient of the trivial estimator $\widehat{\theta_1}$ = constant, which has zero variance. On the other hand $\delta^2 = 0$ corresponds to an unbiased estimator.

- In most cases computation of the uniform CR bound requires numerical methods to find the root of the equation (3.7). Although the uniform bound is specified in terms of δ, it is usually simpler to sweep out the bound in terms of λ. Note that $g(\lambda) = \delta^2$ specifies a one-to-one map from $\lambda \in [0, \infty)$ to $\delta \in [0, 1]$ since $g(\lambda)$ is monotone over $\lambda \geq 0$. For each $\lambda \geq 0$ one can easily determine the corresponding δ by using the relation (3.7).

- For non-singular FIM, when an exact bound on the variance of $\widehat{\theta_1}$ is not necessary, a weaker bound can be derived on $\mathrm{var}_{\underline{\theta}}(\widehat{\theta_1})$, that requires only $\mathcal{O}(q^2 n)$ flops for each value of λ, where $1 \leq q \leq n$. Larger q gives a tighter bound in general. When the FIM is nearly block diagonal then, with little loss in tightness, q can be taken as the block dimension. In this case $q < n$ and this bound can result in significant computational savings. The $\mathcal{O}(q^2 n)$ uniform CR bound is derived in Appendix F.

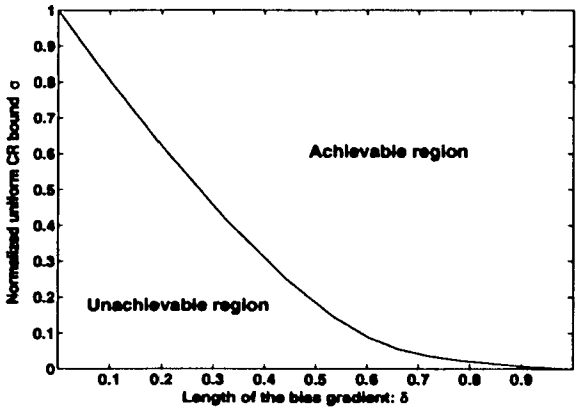

Figure 3.2: The Normalized Uniform CR bound on the δ–σ trade-off plane.

Figure 3.2 shows a typical bias-variance trade-off curve in terms of the bias gradient length and the normalized standard deviation $\sigma = \sqrt{\frac{\mathrm{var}_{\underline{\theta}}(\widehat{\theta_1})}{B(\underline{\theta};\delta)}}$. The region above and including the curve is the so-called 'achievable' region where all the realizable estimators must lie. Note

that if an estimator lies on the curve then lower variance can only be bought at the price of increased bias and vice versa.

As explained in Section 3.5, the curve in Figure 3.2 in the δ–σ trade-off plane can be used as a mapping to obtain a corresponding curve in the bias-variance trade-off plane. To accomplish this, we will require accurate estimation of the estimator bias gradient.

3.4 Estimation of the Bias Gradient

In order to be able to compare the performance of an estimator against the uniform CR bound of Theorem 3, we need to determine the estimator variance and the bias gradient length. In most cases the bias gradient can not be determined analytically and it is therefore important to have a computationally efficient method to estimate it either experimentally or via simulations.

The method of moment estimators for the variance, bias, and bias gradient are as follows. Assume L i.i.d. realizations of the estimator $\{\underline{\hat{\theta}}(\underline{Y}_i)\}_{i=1}^{L}$ are available from a sequence of i.i.d. realizations $\{\underline{Y}_i\}_{i=1}^{L}$ generated from the density $f_{\underline{Y}}(\underline{y};\underline{\theta})$. Consistent and unbiased estimates of bias $\underline{b}(\underline{\theta})$ and covariance $S(\underline{\theta})$ of the estimator $\underline{\hat{\theta}} = [\hat{\theta}_1, ..., \hat{\theta}_n]^T$ are given by:

$$\begin{aligned}
\widehat{\underline{b}(\underline{\theta})} &= \underline{\hat{m}} - \underline{\theta} = \frac{1}{L}\sum_{i=1}^{L} \underline{\hat{\theta}}(\underline{Y}_i) - \underline{\theta} \\
\widehat{S(\underline{\theta})} &= \frac{1}{L-1}\sum_{i=1}^{L} \left[\underline{\hat{\theta}}(\underline{Y}_i) - \underline{\hat{m}}(\underline{\theta})\right]\left[\underline{\hat{\theta}}(\underline{Y}_i) - \underline{\hat{m}}(\underline{\theta})\right]^T,
\end{aligned}$$

where $\underline{\hat{m}} = [\hat{m}_1, ..., \hat{m}_n]^T$ is the estimator of the n-element vector of means $\underline{m} = \underline{m}(\underline{\theta}) = [m_1(\underline{\theta}), ..., m_n(\underline{\theta})]^T$ of $\underline{\hat{\theta}}$.

The direct application of the method of moments to estimating the bias gradient would require performing n additional sequences of L repeated experiments, each for a different perturbation: $\underline{\theta} + \Delta\underline{e}_1, \ldots, \underline{\theta} + \Delta\underline{e}_n$, of $\underline{\theta}$; the n sequences are generated from densities $f_{\underline{Y}}(\underline{y};\underline{\theta} + \Delta\underline{e}_1), \ldots, f_{\underline{Y}}(\underline{y};\underline{\theta} + \Delta\underline{e}_n)$, respectively, where $\underline{e}_j$ is the j-th unit vector, and Δ is a small but non-zero scalar. An estimator of the bias gradient could then be obtained using the finite difference approximation:

$$\widehat{\nabla_{\underline{\theta}}\underline{b}} = \frac{1}{\Delta}\left[\underline{\hat{b}}(\underline{\theta} + \Delta\underline{e}_1) - \underline{\hat{b}}(\underline{\theta}), ..., \underline{\hat{b}}(\underline{\theta} + \Delta\underline{e}_n) - \underline{\hat{b}}(\underline{\theta})\right]$$

This provides a consistent and unbiased estimate of the finite difference approximation to $\nabla_{\underline{\theta}} \underline{b}$. However, the generation of nL realizations of $\underline{Y}$ is not practical when the dimension n of $\underline{\theta}$ is large. For example in image reconstruction a 32×32 image would require 1023 times additional simulations when compared to the computation of bias and variance.

Here we present a method for experimentally determining the bias gradient of an estimator $\widehat{\underline{\theta}}$ which does not require generating $n-1$ additional sequences of L repeated experiments nor does it require finite difference approximation. This method allows the variance, bias, and bias gradient to be determined from a set of simulations $\{\underline{Y}_i\}_{i=1}^{L}$ generated from the density $f_{\underline{Y}}(\underline{y}_i;\underline{\theta})$ for a single $\underline{\theta}$ value. The method is based on the fact that for any random variable Z:

$$\begin{aligned} E_{\underline{\theta}}\left[Z\frac{\partial \ln f_Z(Z;\underline{\theta})}{\partial \theta_j}\right] &= \int_Z z\frac{\partial \ln f_Z(z;\underline{\theta})}{\partial \theta_j} f_Z(z;\underline{\theta})\,dz \\ &= \int_Z z\frac{\partial f_Z(z;\underline{\theta})}{\partial \theta_j}\,dz \\ &= \frac{\partial}{\partial \theta_j}\int_Z z f_Z(z;\underline{\theta})\,dz \\ &= \frac{\partial}{\partial \theta_j} E_{\underline{\theta}}(Z). \end{aligned} \tag{3.13}$$

Thus in particular we have the following relation:

$$\nabla_{\underline{\theta}} \underline{m} = E_{\underline{\theta}}\left[\widehat{\underline{\theta}}(\underline{Y})\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y};\underline{\theta})\right],$$

or extracting the first row,

$$\nabla_{\underline{\theta}} m_1 = E_{\underline{\theta}}\left[\widehat{\theta}_1(\underline{Y})\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y};\underline{\theta})\right],$$

where $\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y};\underline{\theta})$ is called the 'score function' [23]. Since $E_{\underline{\theta}}\left[\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y};\underline{\theta})\right] = 0$, an equivalent relation is:

$$\nabla_{\underline{\theta}} m_1 = E_{\underline{\theta}}\left[(\widehat{\theta}_1(\underline{Y}) - \zeta_i)\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y};\underline{\theta})\right], \tag{3.14}$$

where ζ_i is an arbitrary random variable independent of $\underline{Y}$. As explained in the following discussion, the quantity ζ_i is introduced in (3.14) to control the covariance of the bias gradient estimate.

In view of (3.14) the matrix of gradients of the mean vector $\nabla_{\underline{\theta}}\underline{m}$ can be consistently estimated using the unbiased sample mean estimate:

$$\widehat{\nabla_{\underline{\theta}} m_1} = \frac{1}{L}\sum_{i=1}^{L}\left(\hat{\theta}_1(\underline{Y}_i) - \zeta_i\right)\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta}). \tag{3.15}$$

The estimate $\widehat{\nabla_{\underline{\theta}} b_1} = \widehat{\nabla_{\underline{\theta}} m_1} - \underline{e}_1^T$ of the bias gradient vector directly follows from (3.15), and since we are only interested in θ_1, we need only to generate the consistent and unbiased estimate $\hat{\theta}_1$ of the bias gradient vector:

$$\widehat{\nabla_{\underline{\theta}} b_1} = \widehat{\nabla_{\underline{\theta}} m_1} - \underline{e}_1^T = \frac{1}{L}\sum_{i=1}^{L}\left(\hat{\theta}_1(\underline{Y}_i) - \zeta_i\right)\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta}) - \underline{e}_1^T. \tag{3.16}$$

The estimate $\widehat{\nabla_{\underline{\theta}} b_1}$ converges to $\nabla_{\underline{\theta}} b_1$ with probability one as $L \to \infty$. One can show that the expression for the covariance matrix of $\widehat{\nabla_{\underline{\theta}} b_1}$ is:

$$S(\widehat{\nabla_{\underline{\theta}} b_1}) = \frac{1}{L}\mathrm{cov}_{\underline{\theta}}\left((\hat{\theta}_1(\underline{Y}_i) - m_1)\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta})\right) + \frac{1}{L^2}\sum_{j=1}^{L} E_{\underline{\theta}}\left[(m_1 - \zeta_j)^2\right] F_Y, \tag{3.17}$$

where $F_Y = E_{\underline{\theta}}\left([\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta})]^T[\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta})]\right)$ is the single trial FIM. Assuming ζ_i are i.i.d. for which $\frac{1}{L^2}\sum_{j=1}^{L} E_{\underline{\theta}}\left[(m_1 - \zeta_j)^2\right] F_Y = \frac{1}{L} E_{\underline{\theta}}\left[(m_1 - \zeta_j)^2\right] F_Y$ both terms on RHS of (3.17) decreases to zero at rate $\frac{1}{L}$. However, unless ζ_i is chosen appropriately, the second term on the RHS of (3.17) can be unbounded as a function of $\underline{\theta}$ in many cases since for a reasonably close to unbiased estimator m_1 should be approximately linear in $\underline{\theta}$. In particular we show unboundedness for the weighted least-squares estimator discussed in Appendix C. Thus it is critical to select ζ_i to control $S(\widehat{\nabla_{\underline{\theta}} b_1})$.

A significant reduction in $S(\widehat{\nabla_{\underline{\theta}} b_1})$ can be obtained by choosing $\zeta_i = m_1 =$ constant for all i. This choice of ζ_i minimizes the covariance of the bias gradient estimator by eliminating the possibly unbounded term in (3.17). However, since this optimal $\zeta_i = E_{\underline{\theta}}(\hat{\theta}_1)$ is unknown, we must approximate it with sample mean statistics even at the expense of introducing increased variance. The most obvious approximation of the optimal ζ_i is the sample mean estimate:

$$\zeta_i = \hat{\zeta} = \frac{1}{L}\sum_{i=1}^{L}\hat{\theta}_1(\underline{Y}_i). \tag{3.18}$$

However, since ζ_i is not statistically independent of $\underline{Y}_i$, $E_{\underline{\theta}}\left[\zeta_i \nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta})\right] \neq 0$ and this choice of ζ_i introduces bias in the bias gradient estimator, which is very undesirable. In

order not to introduce such bias, ζ_i should be chosen such that $E_{\underline{\theta}}\left[\zeta_i \nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta})\right] = 0$. Since by relation (3.13) this implies that $\nabla_{\underline{\theta}} E_{\underline{\theta}}(\zeta_i) = 0$ under the pdf $f_{\underline{Y}}(\underline{y}_i;\underline{\theta})$, ζ_i should be an unbiased estimate of a $\underline{\theta}$-independent constant e.g. an unbiased estimate of zero. Such a random variable ζ_i is called an ancillary statistic under $f_{\underline{Y}}(\underline{y}_i;\underline{\theta})$ [38]. The following statistic is easily shown to be ancillary under $f_{\underline{Y}}(\underline{y}_i;\underline{\theta})$ and also unbiased and a consistent estimator of m_1 under $\prod_{i=1}^{L} f_{\underline{Y}}(\underline{y}_i;\underline{\theta})$ for $L > 1$.

$$\zeta_i = \begin{cases} 0 & L = 1 \\ \frac{1}{L-1}\sum_{\substack{j=1 \\ j\neq i}}^{L} \underline{\hat{\theta}}(\underline{Y}_j) & L = 2,3,\ldots \end{cases} \tag{3.19}$$

Substitution of ζ_i defined in (3.16) into (3.15) gives the unbiased and consistent sample mean estimate:

$$\begin{aligned} \widehat{\nabla_{\underline{\theta}} b_1} &= \frac{1}{L}\sum_{i=1}^{L}\left(\hat{\theta}_1(\underline{Y}_i) - \sum_{\substack{j=1 \\ j\neq i}}^{L} \underline{\hat{\theta}}(\underline{Y}_j)\right) \nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta}) - \underline{e}_1^T \\ &= \frac{1}{L-1}\sum_{i=1}^{L}\left(\hat{\theta}_1(\underline{Y}_i) - \sum_{j=1}^{L} \underline{\hat{\theta}}(\underline{Y}_j)\right) \nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta}) - \underline{e}_1^T. \end{aligned} \tag{3.20}$$

Since ζ_i is statistically independent of $\underline{Y}_i$, formula (3.17) holds for the covariance of (3.20). A simple calculation shows that the covariance of the resultant bias gradient estimator is given by:

$$S(\widehat{\nabla_{\underline{\theta}} b_1}) = \frac{1}{L}\mathrm{cov}_{\underline{\theta}}\left((\hat{\theta}_1(\underline{Y}_i) - m_1)\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta})\right) + \frac{1}{L(L-1)}\mathrm{var}_{\underline{\theta}}(\hat{\theta}_1(\underline{Y}_i))F_Y. \tag{3.21}$$

Note that now the second term in (3.21) decreases to zero at a much faster rate of $\frac{1}{L^2}$ as compared to the rate $\frac{1}{L}$ in (3.17). Therefore, asymptotically the ancillary statistic ζ_i achieves nearly identical reduction in the bias gradient estimator covariance that is attained by the optimal ζ_i.

In Appendix C we specialize (3.17) and (3.21) to the case where the observations $\underline{Y}_i$ are assumed to be a set of i.i.d. Gaussian random variables with mean $A\underline{\theta}$ and covariance K, and $\underline{\hat{\theta}}$ is a weighted least-squares estimator of the form: $\underline{\hat{\theta}}(Y) = (A^T K^{-1} A + \alpha I)^{-1} A^T K^{-1} \underline{Y}$, where α is a regularization parameter. The covariance of $\widehat{\nabla_{\underline{\theta}} b_1}$ improves as the regularization parameter is increased. This is to be expected since more regularization implies less influence of the other parameters on the estimate $\hat{\theta}_1$.

3.5 Practical Implementation of the Uniform CR Bound

The estimate $\|\widehat{\nabla_{\underline{\theta}} b_1}\|$ of the bias gradient length $\|\nabla_{\underline{\theta}} b_1\|$, obtained using the expression (3.20), permits one to place an estimator $\widehat{\theta_1}$ of θ_1 on the δ–σ trade-off plane in Figure 3.2. Using the fact $\|\widehat{\nabla_{\underline{\theta}} b_1}\| \leq \delta$ we have an infinite number of choices for δ, namely using $B(\underline{\theta}, \delta)$ with any $\delta \geq \|\widehat{\nabla_{\underline{\theta}} b_1}\|$ provides a valid lower bound on the variance σ^2 of $\widehat{\theta_1}$. However, it is most natural to choose the smallest value $\delta = \|\widehat{\nabla_{\underline{\theta}} b_1}\|$, corresponding to the minimum acceptable bias gradient length for that particular estimator $\widehat{\theta_1}$. This choice of δ provides the tightest uniform CR bound for given $\|\widehat{\nabla_{\underline{\theta}} b_1}\|$ and enables us to effectively compare $\widehat{\theta_1}$ to all other estimators $\widehat{\theta_1}^*$ of θ_1 whose bias gradient length is less than or equal to that of $\widehat{\theta_1}$. At this point a few comments about the bias gradient are in order.

3.5.1 Bias Gradient

The bias gradient $\nabla_{\underline{\theta}} b_1$ is a measure of the influence of each component parameter $\theta_1, ..., \theta_n$ on the mean $m_1(\underline{\theta})$ of the estimator $\widehat{\theta_1}$. Ideally, to be close to unbiased, one would like $m_1(\underline{\theta}) \approx \theta_1$, i.e. $m_1(\underline{\theta})$ should be sensitive to the variations in θ_1 but should be insensitive to the variations in the other parameters $\theta_2, ..., \theta_n$. Alternatively, since $b_1(\underline{\theta}) = m_1(\underline{\theta}) - \theta_1$, it is desirable that the components $\frac{\partial}{\partial \theta_k} b_1(\underline{\theta})$ be of small magnitude, $k = 2, ..., n$. The bias gradient therefore provides important information about the parameter coupling to the estimator mean. However, the bias gradient is in general only indirectly related to the estimator bias, with the exception that $\nabla_{\underline{\theta}} b_1 = 0$ implies $b_1(\underline{\theta}) =$ constant. An estimator that has a constant bias independent of any estimator parameter is removable since it is independent of $\underline{\theta}$. On the other hand, $\nabla_{\underline{\theta}} b_1 \neq 0$ implies that there exists non-removable estimator bias that is dependent on the unknown parameters. A large bias gradient implies an abrupt change in the bias due a small perturbation about a nominal parameter $\underline{\theta}$, even though the actual bias at the nominal $\underline{\theta}$ could be very small. Therefore the bias gradient must be supplemented with information about the estimator bias in order to provide a more complete picture of the average estimator behavior.

3.5.2 Bias-Variance Trade-Off Plane

When accurate estimates $\widehat{b}_1$, $\widehat{\nabla_{\underline{\theta}} b_1}$ and $\widehat{\sigma}^2$ of the estimator bias, bias gradient, and variance are available for a given estimator $\widehat{\theta}_1$ of θ_1, the uniform CR bound lying in the δ–σ plane can be easily mapped into the b–σ plane of variance and biases. This is accomplished by using the ordered triplet $(\widehat{b}_1, \widehat{\nabla_{\underline{\theta}} b_1}, \widehat{\sigma}^2)$ as a mapping between the δ–σ and the b–σ planes. The CR bound on the variance as a function of bias is simply the ordered pair: $\left(\widehat{b}_1, \left[\underline{e}_1 + \widehat{\nabla^T_{\underline{\theta}} b_1}\right]^T F_Y^+ \left[\underline{e}_1 + \widehat{\nabla^T_{\underline{\theta}} b_1}\right]\right)$, a curve denoted $B(\underline{\theta}; b)$ in the sequel.

3.5.3 Recipe for Computing the Uniform Bound for Non-singular F_Y

Often we want to compute the bound at several different values of δ, equivalently λ, and therefore it is important to have a computationally efficient algorithm to compute (3.6). Furthermore, since we are mainly concerned with problems which have a large number of unknown parameters, it may not be practical to compute F_Y^{-1} or F_Y^+ explicitly both due to large computation time as well as large memory storage requirements. We use the method of conjugate gradient as presented in Section 2.5 for this purpose and give a recipe for efficiently calculating the uniform CR bound (3.11) with non-singular F_Y. Additional computational savings can be obtained by exploiting any sparseness of F_Y.

Recipe

Suppose we want to compute the uniform CR bound on the k-th estimator component $\widehat{\theta}_1$ of an $n \times 1$ estimator vector $\underline{\widehat{\theta}} = [\widehat{\theta}_1, \ldots, \widehat{\theta}_n]^T$. Then the numerically stable form of the uniform CR bound (3.12) for non-singular F_Y can be computed as follows.

1. Interchange 1-st and k-th row and column of F_Y.
2. Choose $\lambda \in [0, \infty)$.
3. Compute $\underline{t} = [I + \lambda F_Y]^{-1} \underline{e}_1$ by applying the CG algorithm to solve: $[I + \lambda F_Y]\, \underline{t} = \underline{e}_1$.
4. Compute the bound (3.12): $B(\underline{\theta}, \delta) = \lambda^2 \underline{t}^T F_Y \underline{t}$.

For each λ the above algorithm requires only one application of the conjugate gradient algorithm. Applying this algorithm for several values of λ allows one to trace the curve

$B(\underline{\theta}, \delta)$ in the δ–σ trade-off plane of Figure 3.2.

3.6 Applications

In this section we apply the uniform CR bound to: 1) one-dimensional discrete deconvolution, and 2) one-dimensional edge localization.

3.6.1 One-Dimensional Discrete Deconvolution

The problem of one-dimensional discrete deconvolution is described in Section 2.7.4. It is included here for convenience. We have the model:

$$\underline{Y} = A\underline{\theta} + \underline{q},$$

where $\underline{Y}$ is the blurred data, A corresponds to a convolution kernel, $\underline{\theta}$ is the parameter we want to deconvolve or estimate and $\underline{q}$ is additive white Gaussian noise with covariance matrix Σ. We chose a discrete Gaussian kernel $a_{ij} = \frac{1}{\sqrt{2\pi}} e^{\frac{-(i-j)^2}{2w^2}}$ in these simulations. Philip [30] has remarked that for the discrete deconvolution problem the most ill-conditioned kernels have the form of discrete Gaussian kernels. The FIM F_Y is given by:

$$F_Y = A^T \mathrm{cov}(\underline{Y})^{-1} A = A^T \Sigma^{-1} A.$$

By choosing the width of the Gaussian kernel A we can vary the condition number of the FIM F_Y.

The unconstrained weighted least squares estimator (WLSE) is:

$$\underline{\hat{\theta}}^{WLSE}(\underline{Y}) = (A^T \Sigma^{-1} A + \alpha I)^{-1} A^T \Sigma^{-1} \underline{Y}, \tag{3.22}$$

where α is a regularization parameter. The regularization parameter $\alpha > 0$ improves the numerical stability of the $(A^T \Sigma^{-1} A + \alpha I)^{-1}$ by lowering its condition number at the price of introducing bias, which may reduce the estimator variance. Note that when $\Sigma^{-1} = I$, $\underline{\hat{\theta}}^{WLSE}$ is an unweighted least squares estimator. Since the WLSE is linear it is easy to analytically derive the expressions for the bias of the estimator:

$$\underline{b}(\underline{\theta}) = ((A^T \Sigma^{-1} A + \alpha I)^{-1} A^T \Sigma^{-1} A - I)\, \underline{\theta},$$

the bias gradient:

$$\begin{aligned} \nabla_{\underline{\theta}} \underline{b} &= (A^T \Sigma^{-1} A + \alpha I)^{-1} A^T \Sigma^{-1} A - I \\ &= -(I + \tfrac{1}{\alpha} A^T K^{-1} A)^{-1}, \end{aligned} \tag{3.23}$$

and the covariance:

$$S(\hat{\underline{\theta}}^{WLSE}) = (A^T \Sigma^{-1} A + \alpha I)^{-1} A^T \Sigma^{-1} A (A^T \Sigma^{-1} A + \alpha I)^{-1}.$$

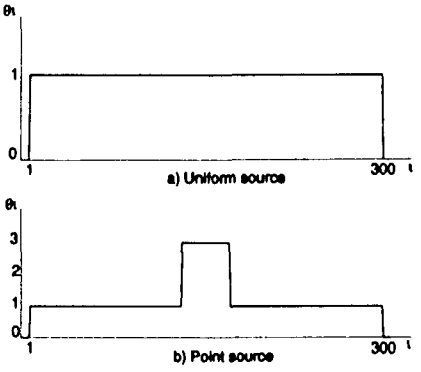

Figure 3.3: One dimensional sources used in discrete deconvolution problem.

Note that all these quantities are indexed by α. Using the above expressions we can analytically compute the uniform CR bound (3.11) for the case when A is 300×300 matrix of Gaussian kernels of width $w = 5$ so that the FIM F_Y is non-singular.

Both a uniform and a point source intensity were studied. As indicated in Figure 3.4, the WLSE with known ideal weights, $\Sigma = \text{cov}(\underline{Y})$, achieves the uniform CR bound for all biases, since $F_Y = A^T \Sigma^{-1} A$ and from expression (3.23) $\nabla_{\underline{\theta}}^T b_1 = \underline{d}_{min}$ for $\lambda = \frac{1}{\alpha}$. Regularization of the form (3.22) introduces bias into the estimator in an optimal manner, i.e. the bias gradient achieves the most variance reduction for its allowed amount of bias.

The uniform CR bound may not be achievable when the FIM is singular. We next consider the problem of super resolution with an up-sampling factor of 2. This refers to a situation when the number of available independent observations are only a half the number of parameters we want to estimate. The A matrix is now of dimension 150×300. This problem is under-determined due to up-sampling. The region of interest is the intensity θ_{150} for the uniform source (Figure 3.3a). The resulting FIM F_Y was singular. Since we do not have a recipe for computing the uniform CR bound $B^o(\underline{\theta}, \delta)$, based on the GS or the

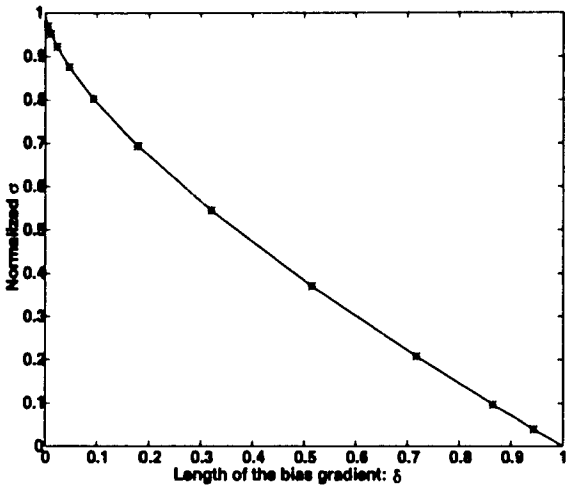

Figure 3.4: Unconstrained WLSE with ideal weights exactly achieves the uniform bound.

CG algorithms for singular F_Y, the uniform CR bound was computed using the singular value decomposition. The bound $B^o(\underline{\theta};\delta)$ is equal to zero for $\delta > 0.81$, for this particular A, since $\|\mathcal{P}_{F_Y}\underline{e}_1\| \leq \delta$ in that region. Note that the estimator (3.22) now fails to achieve the bound for any value of α.

One can show that when $\underline{e}_1$ lies in the range space of F_Y then the uniform CR bound can be achieved even for case of the singular F_Y. To see this consider the estimator:

$$\hat{\underline{\theta}}^{WLSE2} = (I + \alpha F_Y^+)^{-1} F_Y^+ A^T \Sigma^{-1} \underline{Y}, \tag{3.24}$$

where as before α is a non-negative constant. The bias gradient for $\hat{\theta}_1^{WLSE2}$ can easily be shown to be equal to the row vector:

$$\begin{aligned} \nabla_{\underline{\theta}} b_1 &= -\underline{e}_1^T \left(\frac{1}{\alpha} I + F_Y^+\right)^{-1} F_Y^+ - \underline{e}_1^T (I + \alpha F_Y^+)^{-1} (I - F_Y^+ F_Y) \\ &= -\underline{e}_1^T \left(\frac{1}{\alpha} I + F_Y^+\right)^{-1} F_Y^+ - \underline{e}_1^T (I + \alpha F_Y^+)^{-1} (I - \mathcal{P}_{F_Y}) \end{aligned} \tag{3.25}$$

We can recognize that the first term on the RHS of (3.25) is the optimal bias gradient $\underline{d}_{min}$, given by (3.6) for $\lambda = \frac{1}{\alpha}$, which by Theorem 3 achieves the uniform CR bound. The second term on the RHS of (3.25) is equal to zero under the condition that $\underline{e}_1$ is in the range space of F_Y. One case for which this condition holds is for F_Y a block diagonal system of the form:

$$F_Y = \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix},$$

where A ($p \times p$; $n \geq p \geq 1$) and B ($(n-p) \times (n-p)$) are sub-matrices of F_Y such that A is non-singular and B is any arbitrary singular matrix.

It has been verified that for deconvolution example of this subsection $\underline{e}_1$ does not lie in the range space of F_Y and therefore the bound is not achieved by $\widehat{\theta}_1^{WLSE2}$. We doubt that any estimator can achieve the lower bound on the variance $\text{var}_{\underline{\theta}}(\widehat{\theta}_1)$ for a single parameter $\widehat{\theta}_1$ for all biases. However, it should be possible to identify linear combinations $u = \underline{a}^T\underline{\theta}$ of $\theta_1, ..., \theta_n$ for which the bound on $\text{var}(\widehat{u})$ is achievable and the optimal estimator $\widehat{u}$ is identifiable.

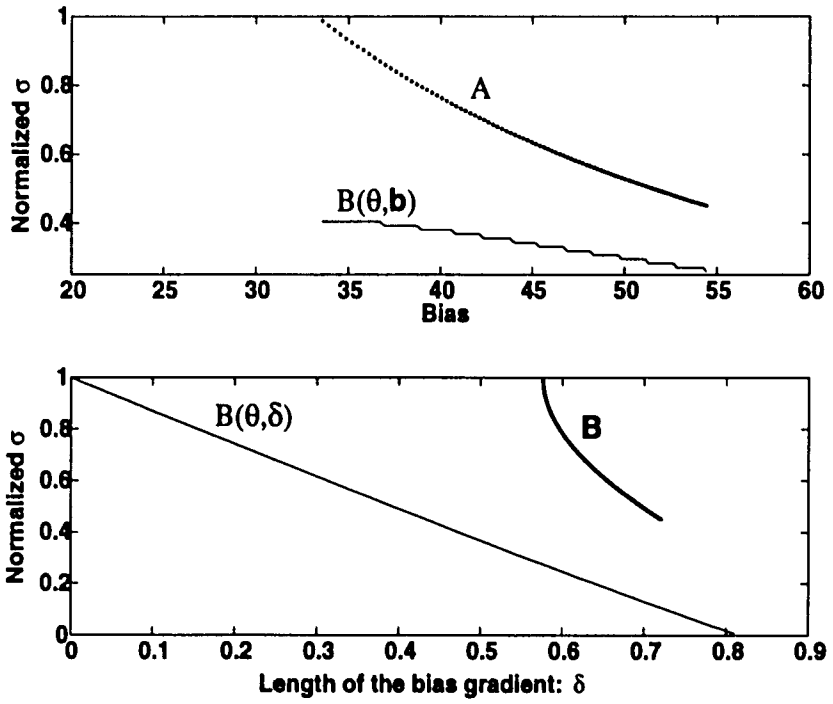

Figure 3.5: Super resolution problem: WLSE fails to achieve the uniform bound at all points. The unbiased CR bound $B(\theta, 0)$ is approximately 50% of the same in Figure 3.4

3.6.2 One-Dimensional Edge Localization

In many imaging applications it is important to determine the location l of an edge along an oriented line segment. In [39] the unbiased CR bound is derived on the localization accuracy of an edge estimator. As in [39] we define an edge by the following 3 parameters (Figure 3.6): 1) Intensity I, 2) location l, and 3) width σ_s. The edge is modelled

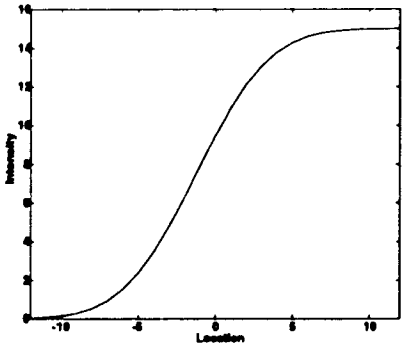

Figure 3.6: A typical edge profile: Intensity I=15, width σ_s = 6, and location l=0.

as the following function of position x along the oriented line segment:

$$R(x) = I\Phi\left(\frac{x-l}{\sigma_s}\right) + q(x),$$

where $q(x)$ is additive white Gaussian noise of variance n_o^2, and Φ is the cumulative distribution function of an $\mathcal{N}(0,1)$ Gaussian random variable. We assume that the width

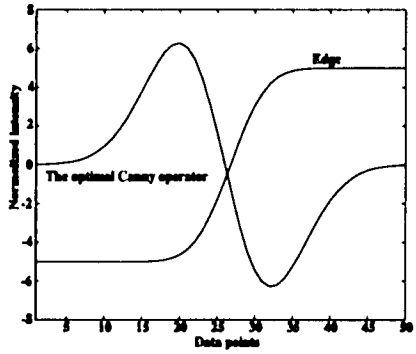

Figure 3.7: The optimal Canny operator in our example. The edge profile is shown here for clarity.

σ_s of the edge is known.

The FIM $F_R(\underline{\theta})$ for the parameter vector $\underline{\theta} = [l, I]^T$ based on the noisy edge observation R is given by [39]:

$$F_R(\underline{\theta}) = \begin{bmatrix} \frac{I^2}{n_o^2\sqrt{\pi}\sigma_s}\left[\Phi\left(\sqrt{2}\frac{T_x-l}{\sigma_s}\right) - \Phi\left(-\sqrt{2}\frac{T_x+l}{\sigma_s}\right)\right] & \frac{-I^2}{n_o^2}\left[\Phi^2\left(\frac{T_x-l}{\sigma_s}\right) - \Phi^2\left(-\frac{T_x+l}{\sigma_s}\right)\right] \\ \frac{-I^2}{n_o^2}\left[\Phi^2\left(\frac{T_x-l}{\sigma_s}\right) - \Phi^2\left(-\frac{T_x+l}{\sigma_s}\right)\right] & \frac{2}{n_o^2}\int_{-T_x}^{T_x}\Phi^2\left(\frac{x-l}{\sigma_s}\right)dx \end{bmatrix},$$

where T_x is the extent of the observation window.

An estimator of $\widehat{l}$ of the edge location was constructed using the Canny operator.

$$h_c(x) = \phi'_{\sigma_c}(x) = \left(\frac{-x}{\sqrt{2\pi}\sigma_c^3}e^{-\frac{x^2}{2\sigma_c^2}}\right) W(x),$$

where ϕ is a Gaussian function with scale σ_c, and $W(x)$ is a window function:

$$W(x) = \begin{cases} 1 & x \in [-\frac{T_x}{2}, \frac{T_x}{2}] \\ 0 & \text{otherwise.} \end{cases}$$

The precess $R(x)$ is filtered by the Canny operator to produce an output $f_c(x)$:

$$f_c(x) = R(x) * h_c(x),$$

where $*$ denotes discrete convolution. The minimum value of $f_c(x)$ determines the location of the edge.

It is shown in [39] that the optimal choice of the Canny width σ_c, determined by minimizing the unbiased CR bound, is $\sqrt{5}\sigma_s$ for an unbiased edge localization algorithm.

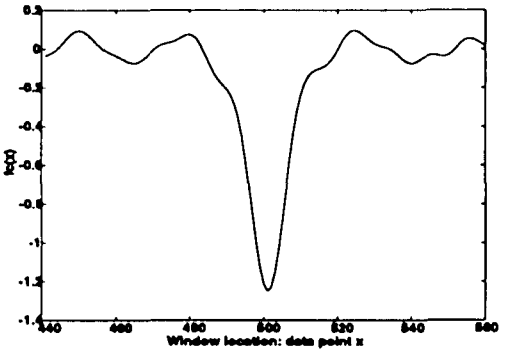

Figure 3.8: Observed output due to the application of the Canny operator. $\sigma_c = 4$. The minimum $f_c(x)$ determines the edge location.

The length of the data $R(x)$ containing the edge was 1000 points. The edge parameters used were: $I = 15$, $\sigma_s = 6$, and $l = 501$. We used a window T_x of 50 data points, $n_o^2 = 8$. We varied σ_c from 3 corresponding to a difference operator, to 31 corresponding to a ramp filter. For each value of σ_c investigated we generated 100 independent realizations of noisy edge profile $R(x)$. The bias, bias gradient and variance were estimated using the methods of Section 3.4 and 3.5.

The results are shown in Figure 3.9. The 95% confidence intervals are smaller than the size of the plotting symbol $*$.

The curve 'B' in Figure 3.9 shows a point of minimum variance at $\sigma_c = 16$, which also corresponds to minimum bias (curve 'A') on the b–σ plane, and hence a point of minimum MSE (Figure 3.10). Note that the minimum variance is achieved close to the optimal $\sigma_c = \sqrt{5}\sigma_s = 13.5$ determined by minimizing the unbiased CR bound.

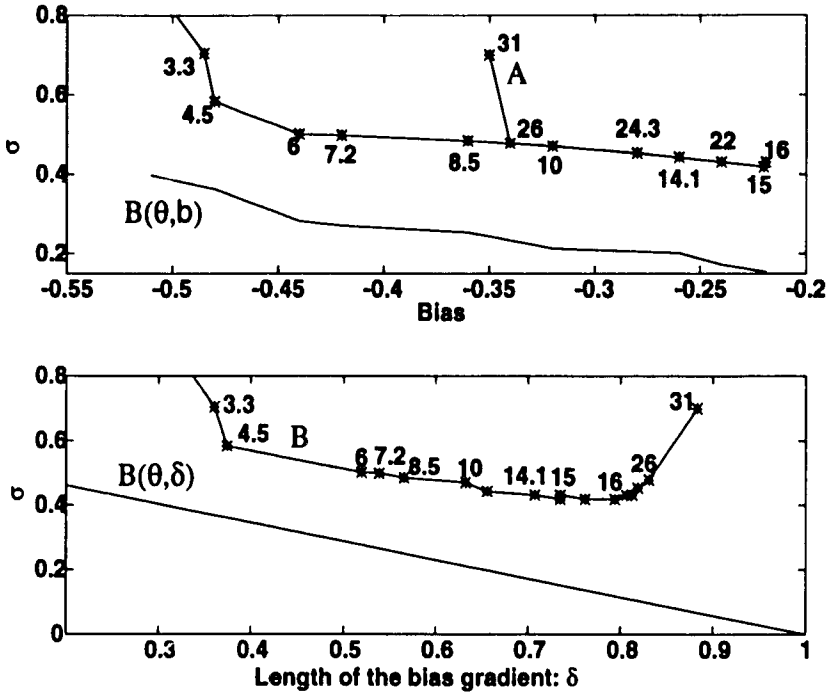

Figure 3.9: The uniform CR bound and the sample variance for varying σ_c. The numbers associated with the curves 'A' and 'B' indicate σ_c.

An interesting point to note is that although the bias and the variance vary non-monotonically with increasing σ_c, the bias gradient length increases monotonically. For σ_c between 4.5 and 16 the estimator standard deviation tracks the uniform CR bound $B(\underline{\theta}, \delta)$, however with an offset of approximately 0.2.

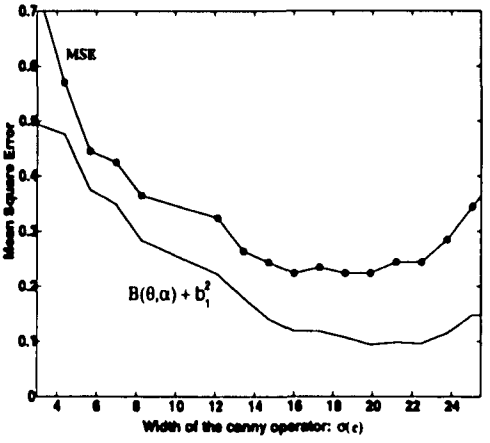

Figure 3.10: There is an optimal σ_c that gives the least MSE.

CHAPTER 4

APPLICATIONS OF UNIFORM CR BOUND FOR SPECT PERFORMANCE EVALUATION AND SYSTEM DESIGN

In this chapter we apply the methods developed in Chapters 2 and 3 to several different applications in SPECT including system design, system optimization, and performance analysis.

4.1 Introduction

Independent of the imaging modality the goodness of the image is eventually decided by human operators and is therefore subjective. Although some measures such as Hotlling trace norm have proven useful, to make a designer's job easier, different measures are defined to predict image quality. In single photon emission computed tomography (SPECT) image reconstruction, the following design measures are commonly used.

1. **Spatial or Geometric Resolution:** is a measure of the closest distance at which a projections geometry can distinguish between two point sources. It is usually measured at Full Width at Half Maximum (FWHM) of the geometric point spread function (psf) of the system [52].

2. **Sensitivity:** is the average number of γ-ray photons striking the detectors, per unit area, over a given period of time, for a given source.

3. **Mean Square Error (MSE):** is the square of the difference between the actual intensity $\underline{\theta}$ of the image, and the estimated intensity $\widehat{\underline{\theta}}$. The MSE can be directly related to the estimator bias $\text{bias}_{\underline{\theta}}(\widehat{\underline{\theta}})$:

$$\text{bias}_{\underline{\theta}}(\widehat{\underline{\theta}}) = E_{\underline{\theta}}(\widehat{\underline{\theta}}) - \underline{\theta},$$

and the estimator variance $\text{var}_{\underline{\theta}}(\hat{\underline{\theta}})$:

$$\text{var}_{\underline{\theta}}(\hat{\underline{\theta}}) = E_{\underline{\theta}}\left[\hat{\underline{\theta}} - E_{\underline{\theta}}(\hat{\underline{\theta}})\right]\left[\hat{\underline{\theta}} - E_{\underline{\theta}}(\hat{\underline{\theta}})\right]^T$$

by:

$$\text{MSE}_{\underline{\theta}}(\hat{\underline{\theta}}) = \text{var}_{\underline{\theta}}(\hat{\underline{\theta}}) + \text{bias}_{\underline{\theta}}(\hat{\underline{\theta}})\text{bias}_{\underline{\theta}}(\hat{\underline{\theta}})^T,$$

where $E_{\underline{\theta}}(\hat{\underline{\theta}})$ is the expected value of $\hat{\underline{\theta}}$.

In any imaging system there are inherent resolution-sensitivity and bias-variance trade-offs. For example in SPECT, a wider collimator opening results in increased sensitivity at the price of decreased spatial resolution. Similarly, in image reconstruction, lower variance can only be bought at the price of increased bias via image smoothing [48]. In this chapter we apply the uniform Cramer-Rao (CR) bound (3.12) to study the bias-variance and resolution-sensitivity trade-offs for several different examples.

The applications considered in this chapter are: i) effect of angular sampling on image reconstruction, ii) optimal aperture design for intensity estimation, iii) performance comparison of the biased weighted least-squares estimator (WLSE) and penalized maximum likelihood (PML) reconstruction algorithms, and iv) analysis of the image reconstruction performance gain by adding a vertex view Anger camera to a SPRINT (SPECT) system [40].

4.1.1 SPRINT System Description

The system investigated is shown in Figure 4.1 and is called an extended SPRINT system [40, 42]. It has two sets of detectors and collimators: ring and vertex. The system without the vertex detector is the standard SPRINT II system [40] which was designed specifically for brain imaging and consists of a ring of detectors and a rotating multi-slit collimator ring. The vertex view is obtained by placing an Anger camera perpendicular to the axis passing through the SPRINT ring.

The vertex detector has its own hexagonally shaped collimator. For increased resolution at a given collimator sensitivity, the vertex detector is placed as close to the ring opening as possible. The system parameters are given in Appendix E and unless otherwise specified are those used in the simulations.

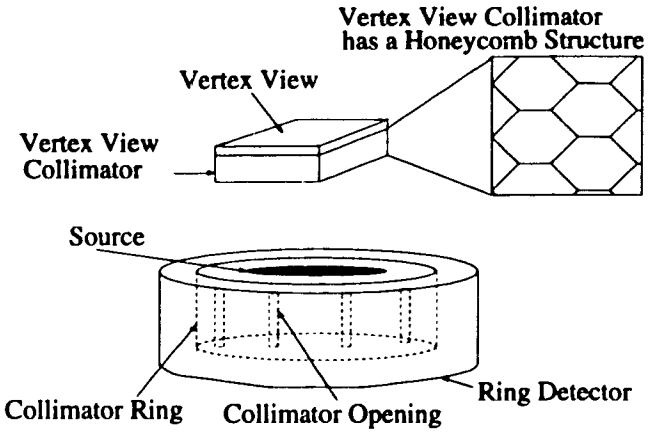

Figure 4.1: The Extended SPRINT system: Standard SPRINT with Vertex View. Not drawn to scale.

4.2 Image Reconstruction in SPECT

A radioactive source emits γ-rays in all directions with equal probability at spatial positions $\underline{V} = [V_1, ..., V_n]^T \in \mathcal{V}$. The source emissions are governed by a spatial intensity vector $\underline{\theta} = [\theta_1, ..., \theta_n]^T \in \Theta$, defined as the average number of γ-ray photons emitted by a pixel during the imaging time. The γ-rays photons are detected at positions $\underline{W} = [W_1, ..., W_d]^T \in \mathcal{W}$. The image reconstruction problem is stated as follows: given the projection data estimate the intensities $\underline{\theta} = [\theta_1, ..., \theta_n]^T \in \Theta$. The number of γ-ray emitted from each pixel, denoted by the vector $\underline{X} = [X_1, ..., X_n]^T$, and the number of detections in each detector bin, denoted by the vector $\underline{Y} = [Y_1, ..., Y_d]^T$ are Poisson distributed:

$$f_{\underline{X}}(\underline{x};\underline{\theta}) = \prod_{i=1}^{n} \frac{\theta_i^{X_i}}{X_i!} e^{\theta_i} \tag{4.1}$$

$$f_{\underline{Y}}(\underline{y};\underline{\theta}) = \prod_{j=1}^{d} \frac{\mu_j^{Y_j}}{Y_j!} e^{\mu_j}, \tag{4.2}$$

where μ_j is the γ-ray intensity at detector j,

$$\underline{\mu} = A\,\underline{\theta}, \tag{4.3}$$

and A is the so-called weights or system matrix consisting of the transition pdfs $f_{\underline{W}/\underline{V}}(\underline{w}/\underline{v})$ and depends upon the system geometry.

Shao [41] has derived the weights matrix, A_r, for the SPRINT ring. The weights matrix for the vertex view, A_v, without the collimator, is derived in the Appendix D.

The vertex view weights, used in the simulations in Section 2.7, incorporate the collimator response and attenuation (wherever specified).

4.3 Uniform Cramer-Rao Bound

In this chapter we will only consider non-singular FIM F_Y, however, the FIMs studied may be ill-conditioned. The following theorem, given here for convenience, gives the uniform CR bound for non-singular F_Y and is a special case of Theorem 3 of chapter 3 when (3.12) is substituted for (3.5).

Theorem 4 *Let $\widehat{\theta}_1$ be an estimator with bias $b_1(\underline{\theta})$ whose n-element bias gradient vector $\nabla_{\underline{\theta}} b_1$ satisfies (3.2). Assume that the FIM F_Y is non-singular. Then the variance of $\widehat{\theta}_1$ is given by:*

$$var_{\underline{\theta}}(\widehat{\theta}_1) \geq B(\underline{\theta}, \delta), \tag{4.4}$$

where $B(\underline{\theta}, \delta)$ is equal to:

$$B(\underline{\theta}, \delta) = [\underline{e}_1 + \underline{d}_{min}]^T F_Y^{-1} [\underline{e}_1 + \underline{d}_{min}], \tag{4.5}$$

$$= \lambda^2 \underline{e}_1^T [I + \lambda F_Y]^{-1} F_Y [I + \lambda F_Y]^{-1} \underline{e}_1 \tag{4.6}$$

where $\underline{e}_1 = [1, 0, ..., 0]^T$ is an n-element unit vector and:

$$\underline{d}_{min} = -[I + \lambda F_Y]^{-1} \underline{e}_1, \tag{4.7}$$

and λ is given by the unique non-negative solution of the following equation involving the monotone decreasing, strictly convex function $g(\lambda) \in [0, 1]$:

$$g(\lambda) = \underline{d}_{min}^T \underline{d}_{min} = \delta^2 \qquad \lambda \geq 0. \tag{4.8}$$

Note that the use of the expression (4.6) does not suffer from any ill-conditioning of the FIM F_Y.

4.4 Applications to Standard SPRINT

In all the following applications the uniform CR bound was efficiently computed using the recipe in Section 3.5.3.

The object used in the first two applications consists of two point sources within a disk of uniform intensity of radius 16 pixels (Figure 4.2). The high intensity white pixels have a normalized intensity value of 2, while the low intensity black pixels are set to 1. For the uniform CR bound computation the pixel of interest was the high intensity pixel at the top of the image marked '1'. Noise due to scatter was neglected in these studies.

It can be shown that for SPECT, the Fisher information matrix has the form [10]:

$$F_Y = A^T \, [\mathrm{diag}(\underline{\mu})]^{-1} \, A,$$

where A is a $d \times n$, $n \geq d$, weight matrix, $\underline{\mu} = A \, \underline{\theta}$, and F_Y is a $n \times n$ symmetric positive definite matrix. We will assume that A is full rank so that $F_Y > 0$. The system used

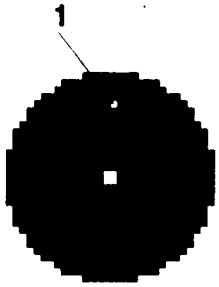

Figure 4.2: The object used in the simulations. The pixel of interest is the pixel at the top of the object. The image dimensions are 32x32 pixels.

in this study is the SPRINT II system, without the vertex view, as described in Section 4.1.1.

4.4.1 Spatial Sampling Study

The effect of attenuation was neglected in this experiment. In Figure 4.3 the bias-variance curves are displayed for varying amounts of spatial sampling. Spatial sampling was varied by rotating the SPRINT collimator ring through different numbers of equal angle increments over $[0, \pi]$ radians. The time for each step was kept constant. Note the monotonic nature of the curves as a function of the number of rotation increments. The curves do not intersect each other due to the fact that the projections from the lower sampled image are a subset of those of the higher sampled image. Also it can be seen that beyond 8 rotation increments increasing sampling rate does not reduce the bound significantly. At this stage

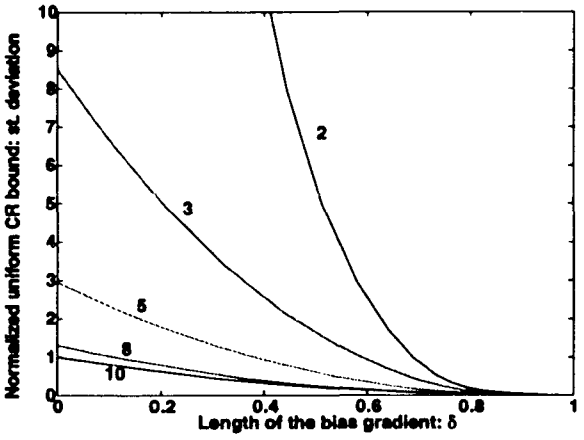

Figure 4.3: The effect of angular sampling on the unbiased CR bound. Curves denote lower bound for 2,3,5,8 and 10 rotation increments of the collimator.

it is not clear how this reduction in the uniform CR bound relates to the image quality.

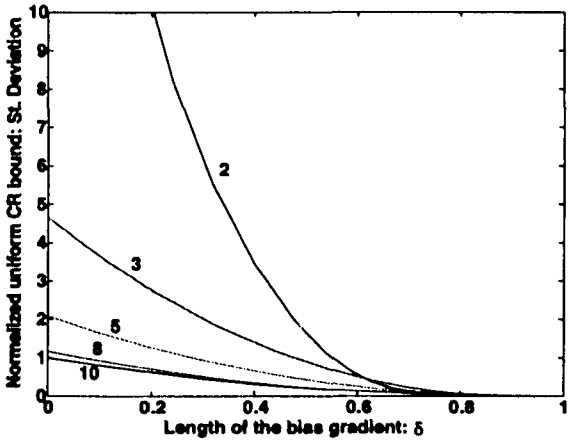

Figure 4.4: The effect of angular sampling on the unbiased CR bound. Time normalized case.

An interesting point to note is that the unbiased CR bound ($\nabla_{\underline{\theta}} b_1 = 0$) for 10 rotations is greater than the bound at a bias gradient length of 0.5 for only 5 rotations. This means that if the bias gradient length $\delta = 0.5$ is acceptable, then a biased algorithm with lower sampling can perform better than an unbiased algorithm with higher sampling.

For a time normalized case, when the total scan time for each curve is kept constant, the curves could intersect each other (Figure 4.4). They, however, display the

same trend as above. Since in this case the total sensitivity of the system is constant, any reduction in the bound in Figure 4.4 is due to a higher resolution.

Note that in experimental setting, for a 64×64 image, 54 uniform steps did not produce a similar saturation behavior. Further increase in resolution was observed by using non-uniform angular rotations for the collimator ring.

4.4.2 Optimal System Design

In this experiment we apply the uniform CR bound (4.5) to standard SPRINT in order to determine an aperture opening that optimizes the trade-off between resolution and sensitivity. The effect of attenuation was not included in the FIM F_Y. For these simulations the aperture opening was varied from narrow (ray width = 0.25 pixels) to wide (ray width = 30 pixels). The imaging time was adjusted so that the total number of detected counts are the same for all cases; hence smaller exposure time for wider openings. The

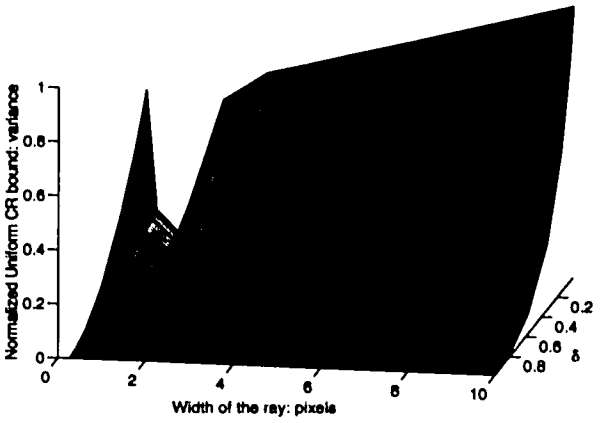

Figure 4.5: Optimal aperture design using the uniform CR bound. One pixel width shows minimum uniform bound.

bias-variance trade-offs are displayed in Figure 4.5. Remarkably the variance shows a sharp minimum over all δ when the width of the ray is approximately one pixel at the pixel of interest. Therefore for the object and the ROI studied, this one pixel aperture width is 'universally' optimal for estimation, irrespective of the bias gradient length of the estimator.

For a time normalized case, where the total imaging time is kept constant, we expect the minimum to be shallower as compared to the count normalized case. This is

due to normalization of the uniform CR bound with higher counts for wider collimator openings, which has an effect of reducing the uniform CR bound. Note that the variance is normalized by the reciprocal of the number of detected photons.

4.4.3 Performance Comparison Between Image Reconstruction Algorithms

We will apply the uniform CR bound to study the bias-variance trade-offs for a particular class of roughness penalized maximum-likelihood (PML) and penalized weighted least-squares (WLSE) image reconstruction algorithms. Since the penalty introduces bias in the estimator due to image smoothing, the standard unbiased CR bound is inapplicable in this study.

For the simulations in this section we used the standard SPRINT II ring geometry without the vertex view. The effect of attenuation was neglected. The total number of detected γ-ray counts were 10^9. Noise due to scatter were 5% of the total counts. In all the following cases, the algorithm was initialized by a uniform disk of intensity 1 and radius 16 pixels. Since both the algorithms considered in this section are non-linear, an analytic expression for the bias gradient is intractable, and therefore the bias gradient was estimated using (3.16). We used $L = 400$ realizations of the projection data $\underline{Y}$ to ensure the statistical accuracy of our estimator bias, bias gradient and variance.

The object is a disk of uniform intensity 1 with a high intensity region of 4 pixels in the center of uniform intensity 2, called the hot spot. The pixel of interest was the pixel at the upper edge of the hot spot, marked '1'. The diameter of the disk is 32 pixels.

Penalized Maximum Likelihood

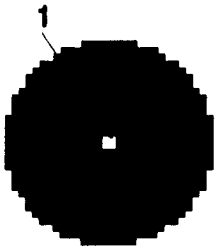

Figure 4.6: The object used in the simulations. The object dimensions are 32 × 32. The black pixels are of intensity 1 while the white pixels are of intensity 2.

The penalized maximum-likelihood (PML) is penalized for roughness and has the same functional form as a maximum-a posteriori (MAP) estimator of the image intensities $\underline{\theta}$. The general form of the PML is given by:

$$\hat{\underline{\theta}}(\underline{y}) = \underset{\underline{\theta} \in \Theta}{\text{argmax}} \left\{ \ln f_{\underline{Y}}(\underline{y}; \underline{\theta}) - \alpha P(\underline{\theta}) \right\},$$

where $P(\underline{\theta})$ is a roughness penalty and α is the smoothing parameter. We use a penalty function described in [36] which is imposed on the 8 neighboring pixels for each pixel of interest. Setting $\alpha = 0$ corresponds to no image smoothing while a large value of α corresponds to a significant amount of smoothing. We have implemented the recursive SAGE algorithm to maximize the PML objective function. SAGE, which stands for space alternating generalized EM, involves an intelligent choice of a 'complete data space' such that the E and M steps are analytically tractable. A detailed description of the PML-SAGE algorithm is given in [36].

It is easy to show that for the Poisson model

$$\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{y}; \underline{\theta}) = A^T \left[-\underline{1} + \underline{y} \oslash \underline{\mu} \right],$$

where $\oslash$ is a vector operation denoting element-by-element division, and $\underline{1} = [1, 1, ..., 1]^T$.

For the first set of simulations the smoothing parameter α was varied (Figure 4.7 (a)). Points on the curves in Figures 4.7 (a) and (b) are labeled by the exponent of α. The bias, bias gradient and variance were estimated and the uniform bound was plotted over the bias gradient length-variance trade-off plane, denoted δ–σ, and the bias-variance trade-off plane, denoted b–σ. The PML-SAGE algorithms were terminated after 100 iterations for each of the $L = 400$ trials. The ellipsoidal confidence regions are not shown in the figure since they are smaller than the size of the plotting symbol '$*$'. Note that the uniform CR bound, denoted by $B(\underline{\theta}; \delta)$ in Figure 4.7 (a), is achieved for large biases, i.e. large α. For α small, the curve 'B' tends to deviate more from the lower bound and saturate, i.e. lower α does not decrease the bias gradient. On the other hand the bias decreases to an asymptote near zero.

At points close to the unbiased point, i.e. the leftmost corner of the horizontal axis, in curve 'A', maximal reduction in bias is achieved at the price of significant increase

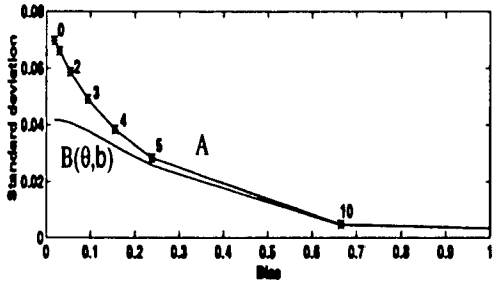

(a) PML-SAGE

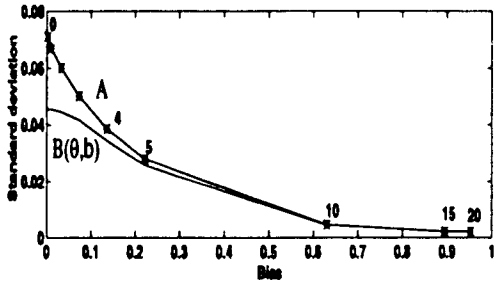

(b) WLSE

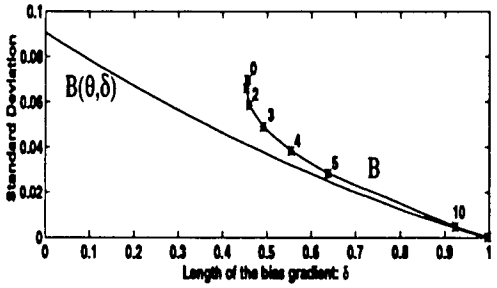

(a) PML-SAGE

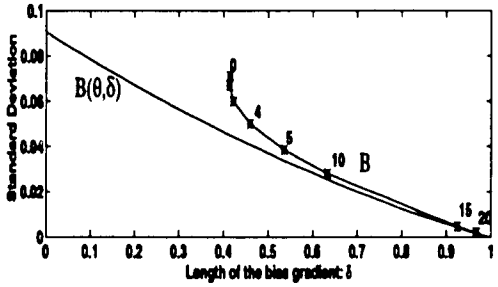

(b) WLSE

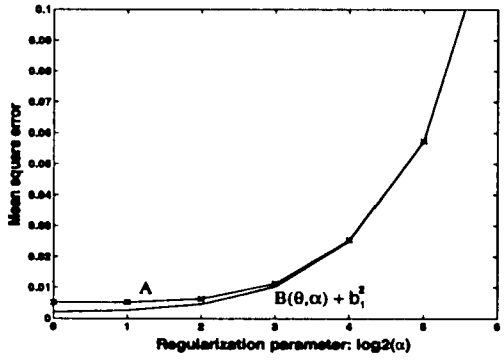

(c) MSE: PML-SAGE

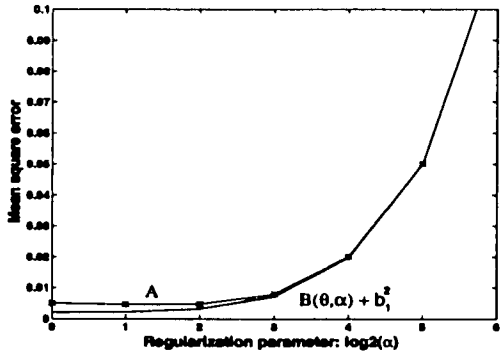

(d) MSE: WLSE

Figure 4.7: Comparison of the PML and WLSE

in the variance. An interesting question to ask is: what is the best choice of α in this case? Unfortunately there is not a simple answer to this question since it is application dependent. However, using the two plots in Figure 4.7 (a), the bound on the variance as a function of bias can be easily translated into a bound on the estimator MSE, denoted $B(\underline{\theta}, \alpha) + b_1^2$, as a function of the smoothing parameter α. The bound, along with the estimator MSE is shown in Figure 4.7 (c). The point of least MSE is the point of maximum variance but least smoothing. This supports the widespread opinion that when taken in isolation the pixel intensity MSE is not suitable as an image quality metric. However, a suitable trade-off between the smoothing, which directly corresponds to the image bias, and the variance is to choose a point close to the knee of the MSE curve 'A' in Figure 4.7 (c), corresponding to $\alpha = 2^3$, or $\alpha = 2^4$. This choice of α ensures that the reconstructed image is not overly smoothed, at the same time not having an unacceptably large variance.

Figures 4.8 and 4.9 show several graphs of reconstruction quantities for $\alpha = 2^4$, and $\alpha = 2^{10}$, respectively. For clarity in the figures, we down-sampled all the images by a factor of 2. For each image in Figures 4.8 and 4.9 the ordered pair at bottom indicates the minimum and maximum values for that image. In Figure 4.8, the mean reconstructed image is very close to the true image except around the edges. The correlation image, i.e. the column of F_Y^{-1} corresponding to the pixel of interest, θ_{ROI}, shows a strong correlation with the neighboring pixels. This implies that to estimate θ_{ROI} we must also estimate the strongly correlated neighboring pixels accurately, while the influence of the far pixels can be ignored. Ideally, one would like the correlation between the pixels to be zero so that the estimate of a certain pixel, θ_{ROI}, is independent of the estimates of all other pixels. The plot for the theoretically optimal bias gradient $\underline{d}_{min}$ shows a similar strong influence from the neighboring pixels.

The average bias gradient $\nabla_{\underline{\theta}} b_1$ for the reconstructed image is different from the theoretically optimal bias gradient $\underline{d}_{min}$. Thus the penalized SAGE image reconstruction algorithm does not take best advantage of its bias allocation since it is only by using the optimal bias gradient $\underline{d}_{min}$ given by (4.7) that the minimum bias gradient length is achieved.

Figure 4.9 shows the same set of images as in Figure 4.8 but for $\alpha = 2^{10}$. Due to very high regularization, the hot spot is almost entirely smoothed out. Also, neither $\underline{d}_{min}$

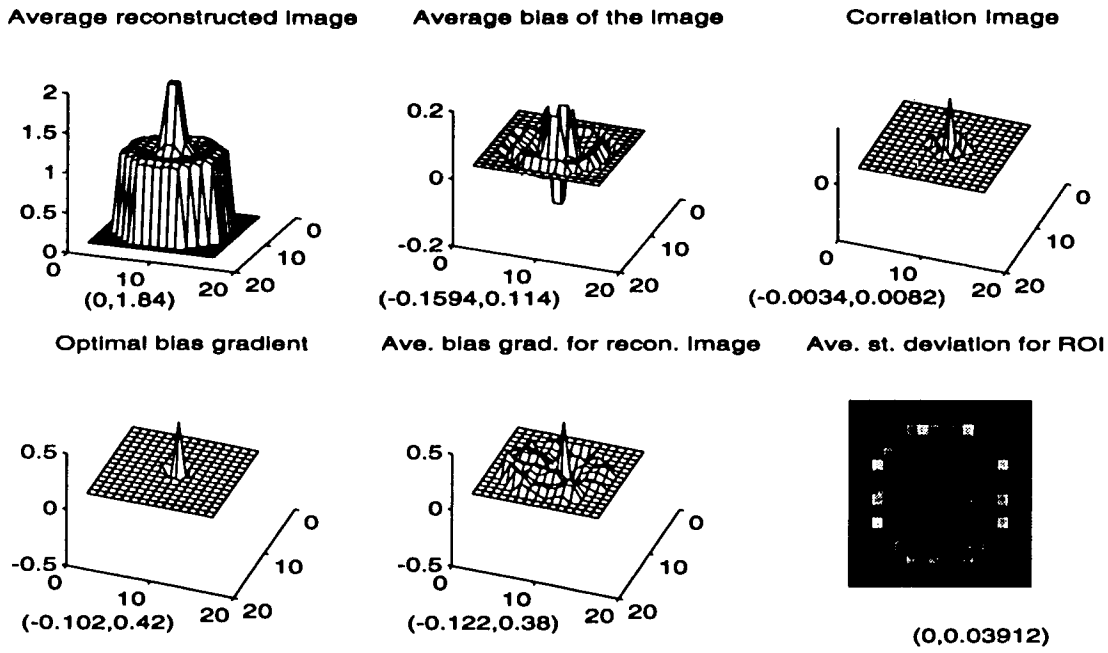

Figure 4.8: PML-SAGE: different graphs of reconstruction quantities for $\log_2(\alpha) = 4$. An ordered pair with each curve indicate the (minimum, maximum) value associated with that image.

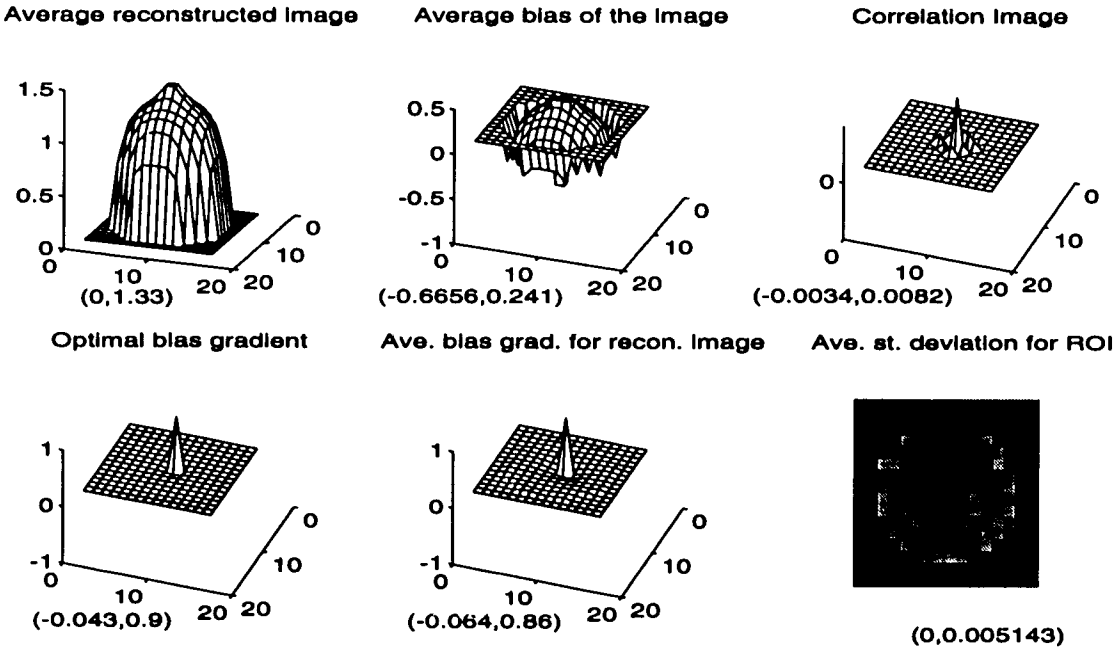

Figure 4.9: PML-SAGE: different graphs of reconstruction quantities for $\log_2(\alpha) = 10$.

nor the average bias gradient $\nabla_{\underline{\theta}} b_1$ for the reconstructed image show significant coupling between the pixel of interest and the neighboring pixels. This is to be expected since in the overly smoothed case the bias is principally determined by the smoothness penalty as opposed to the projection data.

Weighted Least-Squares Estimator

Similar to the PML, the WLSE is penalized for roughness, but minimized over a quadratic objective function. The WLSE is given by [35]:

$$\hat{\underline{\theta}}(\underline{y}) = \underset{\underline{\theta} \in \Theta,\ \underline{\theta} \geq 0}{\operatorname{argmin}} \left\{ \frac{1}{2} \left[(\hat{\underline{y}} - A\underline{\theta})^T \Sigma^{-1} (\hat{\underline{y}} - A\underline{\theta}) \right] + \alpha P(\underline{\theta}) \right\},$$

where Σ is a weight matrix, $P(\underline{\theta})$ is a regularization penalty, and A is the system matrix. We use a penalty function described in [35] which is imposed on the 8 neighboring pixels for each pixel of interest. The weight matrix Σ is diagonal, consisting of the covariance estimate of the observations. It is shown in [48] that a WLSE with an identity penalty function and ideal weight matrix $\Sigma = \text{diag}_i(\mu_i)$ exactly achieves the uniform CR bound for all biases.

Figure 4.7 (b) shows the δ–σ and b–σ plots for the WLSE. The WLSE estimator follows the uniform CR bound closely for high bias and low variance, but tends to deviate away from the bound for low biases. An interesting point to note is that the PML and the WLSE have similar bias-variance trade-off curves. However, the uniform bound on bias $B(\underline{\theta}, b)$ is different for PML than that for WLSE since the bound on bias is indexed by algorithm bias gradient which is obviously algorithm dependent.

Figures 4.10 and 4.11 show several graphs of reconstruction quantities using the WLSE for $\alpha = 2^4$ and $\alpha = 2^{10}$. The comments for Figures 4.8 and 4.9 are valid here. The only exception being that the WLSE fails to accurately estimate the edges for small α. This is due to the fact that the estimates of covariance involving the projections that graze the image edges

The PML and the WLSE have very similar bias gradient-variance and bias-variance trade-offs, and based on these simulations one can not prefer one algorithm over the other. However, the PML takes longer to converge when compares to the WLSE.

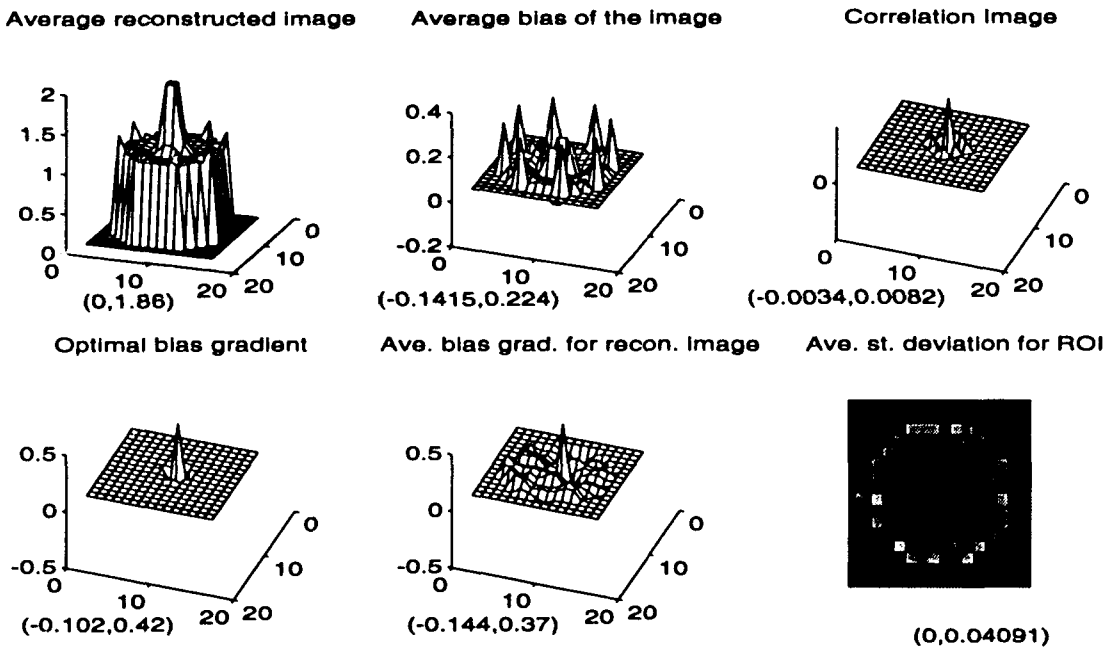

Figure 4.10: WLSE: different graphs of reconstruction quantities for $\log_2(\alpha) = 4$.

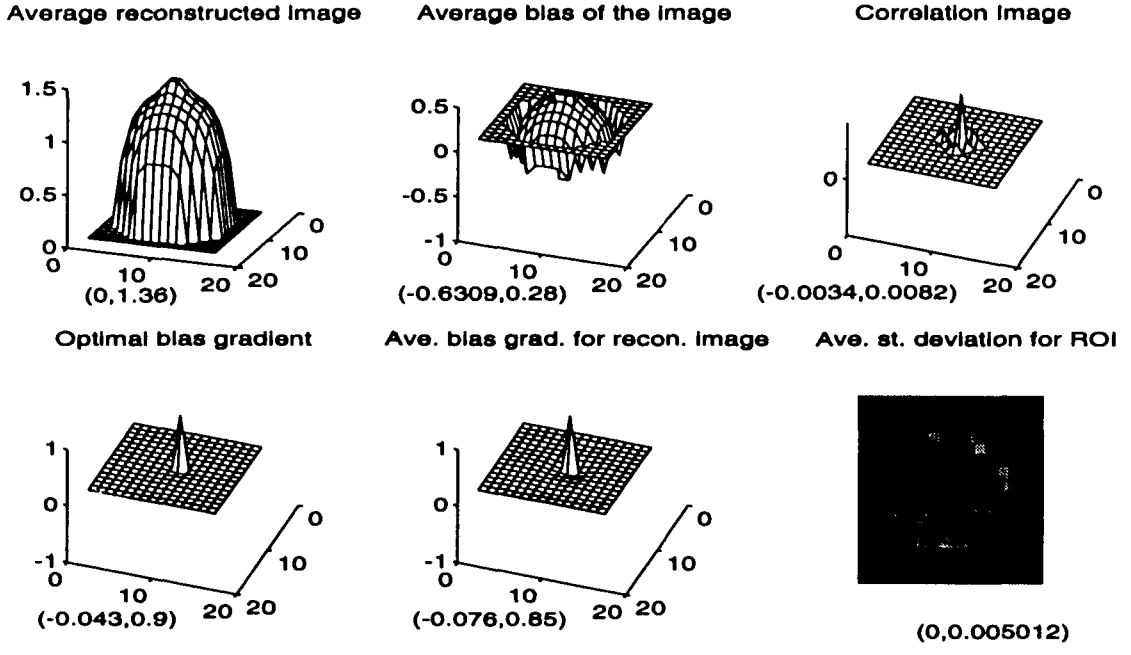

Figure 4.11: WLSE: different graphs of reconstruction quantities for $\log_2(\alpha) = 10$.

4.5 Vertex View

The aim of the following experiments is to determine a realistic estimate of the gain in performance, measured in terms of a reduction in the uniform CR bound, achievable by adding vertex view to a standard SPRINT ring. For these experiments we have used a cylindrical object 10 cms high with a radius of 11 cms. The object is divided into 6 slices, each 1.66 cms high. Each slice consists of 812 pixels (Figure 4.12). The slice diameter is 32 pixels and contains a pair of hot spots of normalized intensity 2, superimposed on a uniform disk of normalized intensity 1. The top slice is at a distance of 2.33 cms from the vertex view collimator.

All the following results are normalized for a total detected γ-ray count of 1 million (1M) for the SPRINT ring, unless otherwise stated. The vertex view spatially

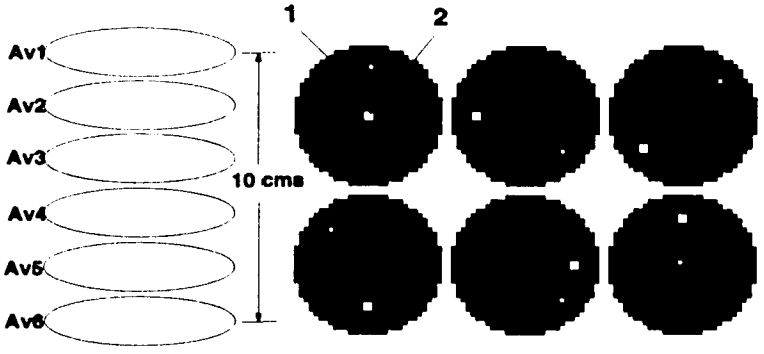

Figure 4.12: The object used in the experiments consists of six slices. Top left: slice 1; bottom right: slice 6.

varying solid angle resolution was taken into account in all of the following experiments.

4.5.1 Experiment 1: Effect of Increasing Number of Parameters

The resolution of the vertex view collimator decreases with distance, however, since the purpose is to study only the effect of increasing the number of parameters, we kept the vertex view collimator resolution constant. The collimator resolution is taken to be equal to the resolution associated with slice 3, which is close to the average resolution for our object. The number of unknown parameters were increased by successively adding identical slices, chosen to be slice 1, hence the object in this experiment is variable. The effect of attenuation was neglected. The total time of the experiment was kept constant for

Object	Experiment 1	Experiment 2	Experiment 3
Number of slices	1,2,...,6 (identical slices)	2 Top: fixed Bottom moved: 2,...,6	6 (non-identical slices)
Slices containing ROI	Top	Top and bottom	1,3, and 6
Attenuation (ring)	Off	Off	Off
Attenuation (vertex)	Off	Off	On
Vertex view resolution	Constant (slice 3)	Variable	Variable
Vertex collimator thickness	2 cms	2 and 4 cms	2 cms

Table 4.1: Objects used in the vertex view simulations.

each case, therefore, the most number of counts are detected when the object consists of all 6 slices including the vertex view. The pixels of interest θ_{ROI} were a low intensity pixel close to the top-left edge of slice 1, marked '1', and a high intensity pixel in slice 1 marked '2'.

The results are shown in Figure 4.13. The uniform bound for the ring alone remains the same for all objects since the ring projections for slices 2,...,6 are independent of the projections for slice 1 and do not effect the bound. When the object consists of only the top slice, addition of the vertex view shows an enormous gain. This is to be expected, since in the absence of any background noise, the vertex view directly estimate the object pixels with no need for image reconstruction. However, with the addition of another slice there is a sharp increase in the bound indicating poorer estimation performance. With all 6 slices present the gain due to adding the vertex view is significant with a variance reduction of approximately 80%. The edge pixel shows similar behavior, except that the bound for the edge pixel is lower than the center pixel since the SPRINT ring has higher resolution and increased sensitivity closer to the edge. Both of these factors, as we saw above, reduce the uniform CR bound.

The results from this experiment suggest that any performance gain offered by adding the vertex view is very sensitive to the increasing number of parameters to be estimated. Since the effect of distance of the slice from the vertex view is neglected in

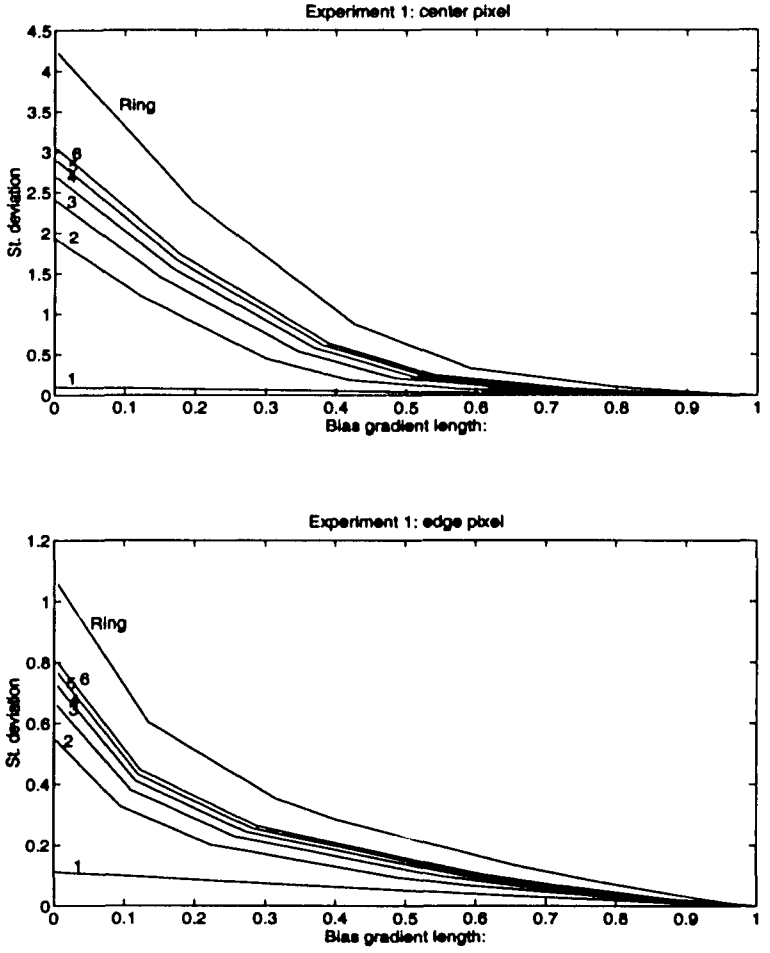

Figure 4.13: Experiment 1: the effect of increasing the number of parameters on the vertex view gain. The numbers on the curves indicate the number of slices used in the object. Ring denotes the SPRINT ring without the vertex view. Top: center pixel; bottom: edge pixel.

this experiment, due to a constant resolution vertex view collimator, we expect similar bias-variance trade-off curves for pixels chosen in slices other than slice 1.

4.5.2 Experiment 2: Resolution-Sensitivity Trade-offs

Slice number	High resolution (4 cms)	Low resolution (2 cms)
1	3.2 mm	4.4 mm
3	5.0 mm	7.8 mm
4	5.8 mm	9.6 mm
6	7.5 mm	13 mm

Table 4.2: Vertex collimator resolution.

This experiment studies the inherent resolution-sensitivity trade-offs due to the choice of different collimators. In this experiment we used two different vertex view collimators: 1) high resolution (4 cms thickness) and 2) moderate resolution (2 cms thickness). The resolution of the 4 cms collimator is roughly twice the resolution of the 2 cms collimator, at the expense of a factor 4 decrease in sensitivity. The effect of attenuation was turned off to isolate the effect of resolution. The scan time was kept constant. We used two slices, one slice was fixed at the top position, while the position of the other slice, chosen to be slice 2, was moved from position 2 to 6. The pixels of interest for slice 1 were the same as in experiment 1. We also chose 2 pixels in the moving slice corresponding to the positions of the pixels of interest in slice 1. Figures 4.14 and 4.15 shows the results. The numbers on the curves indicate the position of the second slice.

For slice 1 (Figure 4.14) the uniform bound for the 2 cms collimator is lower than that for the high resolution 4 cms collimator. This means that the gain in resolution by using the 4 cms collimator failed to compensate for the loss in sensitivity. This is to be expected since for slice 1 the low resolution (2 cms) collimator has better resolution than the pixel size (6.9 mm $\times$ 6.9 mm), as shown in Table 4.2. Therefore, there is no additional gain in resolution by using a high resolution collimator for slice 1. For the moving slice (Figure 4.15) the uniform CR bound for both the collimators are similar, which implies that at greater distance, the increase in the bound due to decreased resolution is compensated by increased sensitivity. Based on the results of this experiment alone, for a 32 $\times$ 32 image,

the 2 cms collimator is to be preferred since it gives a better performance for slice 1 when compared to the 4 cms collimator. Since for slice 1, the resolution of both the 2 cms and the 4 cms collimator is sub-pixel size, a lower uniform CR bound for the 23 cms collimator implies a better image quality. However, it is difficult to predict image quality for other cases.

In Figure 4.14 the uniform bound for slice 1 with the moving slice at position 2 is higher than that at position 6. This is due to the fact that a slice at position 2 offers a stronger background for the top slice, and therefore makes estimation more difficult. Similar reasoning explains the order of the curves in Figure 4.15.

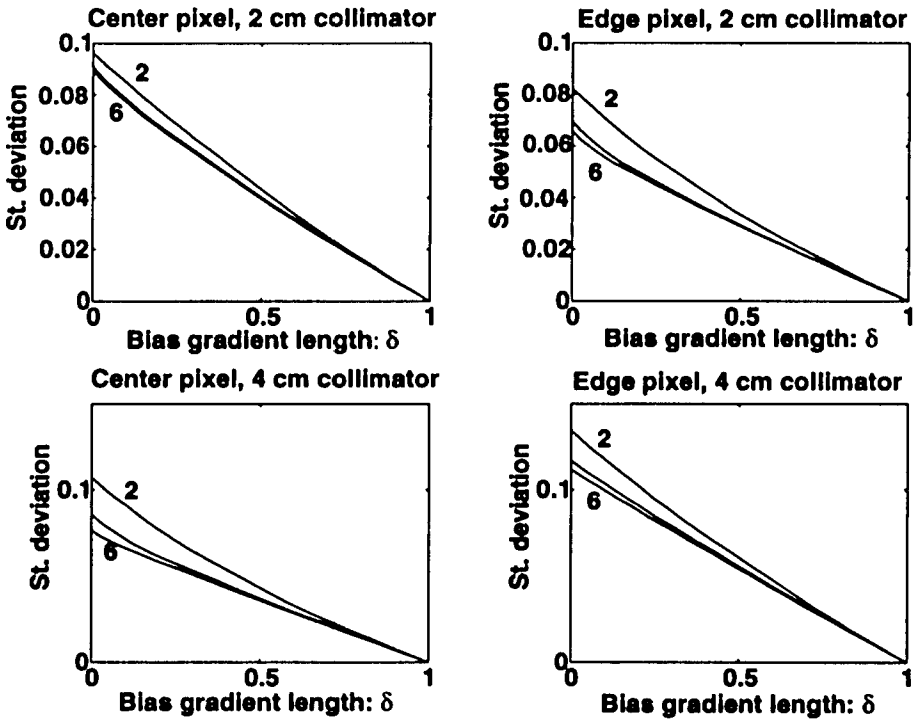

Figure 4.14: Experiment 2: Resolution-sensitivity trade-offs for the vertex view, slice 1. The numbers on the curves indicate the position of slice 2. Top left: 2 cms collimator thickness, center pixel; top right: 2 cms collimator thickness, edge pixel; bottom left: 4 cms collimator thickness, center pixel; bottom right: 4 cms collimator thickness, edge pixel.

4.5.3 Experiment 3: Source of the Vertex View Gain

Possible sources of gain are increased resolution, increased sensitivity, and orthogonality of the vertex view data set to the SPRINT data. The purpose of this experiment

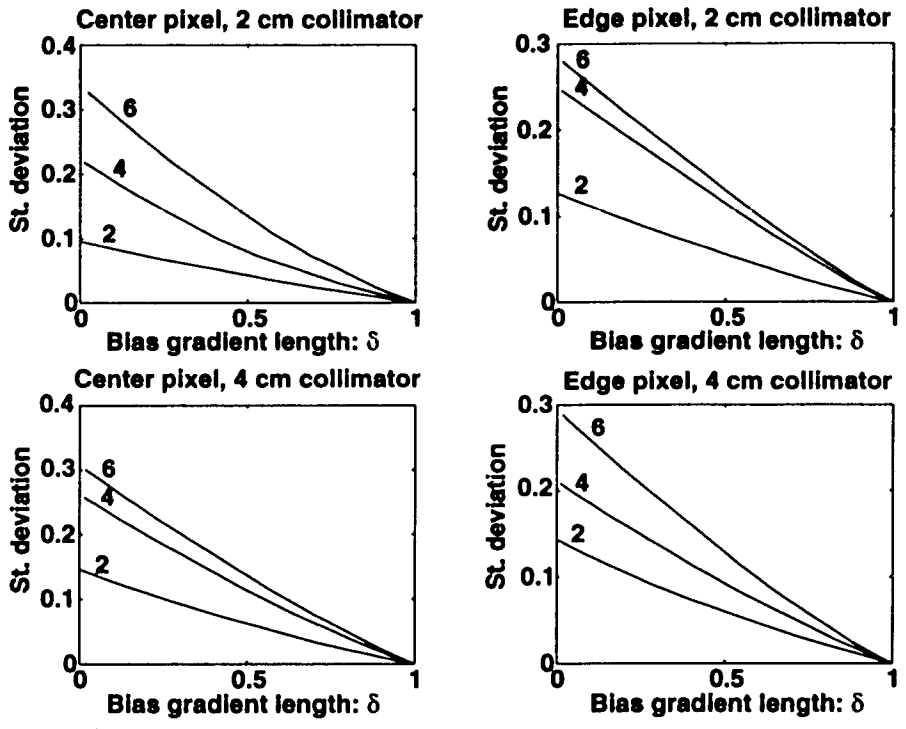

Figure 4.15: Experiment 2: Same as Figure 4.14 except the pixels of interest are in the moving slice 2.

is to determine the source of this gain. For this experiment the effect of attenuation in the vertex view was considered. All six slices were included, hence the object is fixed. Due to the results in the last experiment, we used 2 cm (moderate resolution) collimator. The object of interest was the center pixel in each of the slices 1, 3, and 6. The effect of increased sensitivity by adding the vertex view was neutralized by normalizing the total number of detected counts to 1M with and without the vertex view. Therefore, any gain observed in this case can only be attributed to the increased resolution due to the addition of the vertex view or the orthogonality of the vertex view data.

An interesting point to note is that due to the data-processing theorem, addition of vertex view necessarily means a better, or at the very least an equivalent performance when compared to the SPRINT ring. However , the count normalized case can result in a worse overall performance since it is possible that the increased resolution due to the vertex view or the orthogonal vertex view data is not able to compensate for the decreased ring sensitivity.

The results are shown in Figure 4.16. The curve marked 'Count normalized' refers to the total counts including the vertex view equal to 1M, while the curve marked 'Time normalized' refer to the system including the vertex view, and the total γ-ray counts for the whole system are $\approx$ 1.6M. The bound on the center pixel in the top plane is less than 10% near zero bias gradient length. This applies to both the count and time normalized cases. However, the gain decreases for the 3rd and the 6th slice. In particular, for the 6th slice, close to zero bias gradient length, the count normalized bound is $\approx$ 60% of the ring bound, which refers to the bound reduction due to orthogonality of the vertex view data only, since for slice 6 the vertex view resolution is approximately equal to the SPRINT ring resolution. Therefore the gain in performance can not be due to increased resolution. The time normalized bound, i.e. the combined reduction in variance due to increase in resolution, sensitivity as well as the orthogonal data set, is less than 50% of the SPRINT ring bound.

The very high gain for the top slice, even in the presence of all the other five slices, can be explained by the high detector count of the vertex view (73×73) compared to the image size (32×32). In the top slice the vertex view sees each pixel individually, even when the detector resolution is taken into account. Despite decreased resolution and sensitivity, the gain for the 6th slice is significant. This result implies that the 'information' contained in one vertex view detected photon is worth approximately 4 detected photons in the ring, even at a distance of 10.33 cms from the vertex view. As mentioned above, this gain is only due to the orthogonal vertex view data set. For slice 3, the vertex view photon contains close to 16 times more information than the SPRINT ring. However, this is due to a combined effect of increased resolution and orthogonal vertex view data set. In view of the fact that the SPRINT data does not include attenuation, and the results of experiment 1, this gain seems very optimistic and further experiments are needed to confirm it.

Based on the above 3 experiments we can conclude that the use of the vertex view will definitely improve the quality of the reconstructed image for the slices that are closest to the vertex view. Furthermore, the vertex view not only increases the overall sensitivity of the system, but it also adds to the resolution. In particular, the vertex view can improve the localization of the pixels that are farthest from the SPRINT ring detector, i.e. in the

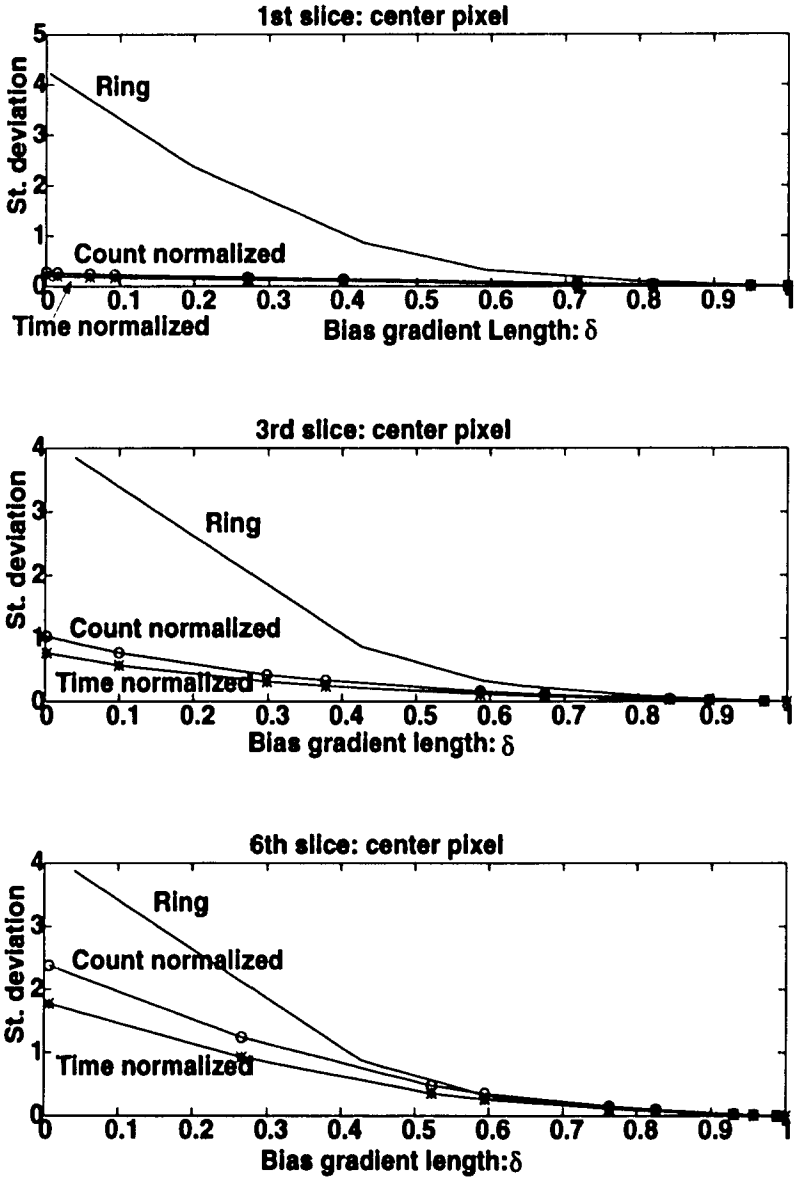

Figure 4.16: Experiment 3: The source of the vertex view gain. Top: top slice; middle: 3rd slice; bottom: 6th slice.

center of the object. Furthermore, even if we neglect the effect of increased resolution and sensitivity, the orthogonal data set appears to give significant gain when compared to the SPRINT ring only.

CHAPTER 5

SUMMARY AND FUTURE RESEARCH

A uniform CR bound applicable to biased estimators was presented in [14]. However the computation of the uniform CR bound required inversion of an FIM. Direct inversion algorithms require $\mathcal{O}(n^3)$ flops, where n is the length of the parameter vector. A geometric series based recursive and monotonically convergent algorithm, called the GS, was proposed in [10, 13] to compute the columns of the inverse of the FIM that requires only $\mathcal{O}(n^2)$ flops per iteration. The GS also requires specification of a splitting matrix which should be easy to invert. The choice of splitting matrix affects the rate of convergence and therefore a proper choice of these splitting matrices is crucial. Using purely algebraic techniques, we specified a class of banded diagonal matrices that result in much faster convergence rate of the GS when compared to a previous statistical approach (Section 2.3). We also presented a modified conjugate gradient (CG) algorithm which converges non-monotonically, but has a faster convergence rate than the GS (Section 2.5). The algorithms described here are only applicable to symmetric and positive definite matrices. Based on these algorithms we provided a 'recipe' to compute the uniform CR bound for full-rank FIMs (Section 3.5).

As presented in [14], the uniform CR bound is only applicable to full-rank FIM. For problems where the number of independent observations are less than the number of unknown, the FIM is rank-deficient. We extended the uniform CR bound to allow the rank-deficient FIM (Section 3.3). The uniform CR bound for the rank-deficient FIM requires pseudo-inverse of the FIM. We derived two recursive algorithms to approximate columns of the pseudo-inverse of the FIM (Section 2.6). These algorithms are based on the GS and the CG. The first method, called the perturbation method, is based on the well known

Tikhonov regularization [27], while the second method, called the differencing method, is new. In the examples considered in Chapter 2, the differencing method gives a significantly lower error of approximation when compared to the perturbation method.

The computation of the uniform CR bound requires the bias gradient length to be known at the parameter value of interest. In most cases the bias gradient of the estimator is analytically intractable and must be estimated. The method of moments approach to the estimation of the bias gradient, generally used to approximate the bias and the variance of an estimator [51], is computationally inefficient. We presented a computationally efficient method, based on sample averaging the score function, which has zero bias and low variance (Section 3.4).

The methods developed in this thesis were demonstrated for several different examples including two-dimensional SPECT image reconstruction, one-dimensional discrete deconvolution, and one-dimensional edge localization. A major contribution of this thesis is a better understanding of the SPECT system and image reconstruction algorithms through these examples. We have studied several SPECT system design and optimization examples, including the gain in performance in terms of the reduction in the uniform CR bound by adding a vertex view to a SPECT ring geometry. We also compared two popular image reconstruction algorithms, namely, penalized maximum likelihood and penalized weighted least-squares.

5.1 Future Research

5.1.1 Modified CG Algorithm

For our application, monotone convergence of the GS algorithm is an important feature since a valid approximation to the CR bound is obtained at the end of each iteration. However, the non-monotone convergent CG algorithm gives a faster rate of convergence when compared to the GS algorithm. One solution to this rate of convergence-monotonicity trade-off would be to design a hybrid algorithm which is controlled by a single parameter that controls this trade-off. We feel that the CG algorithm could be modified in this direction.

5.1.2 Recipe for Computing the Uniform CR Bound for Singular F_Y

In Section 3.3 we gave a recipe for computing the uniform CR bound for non-singular F_Y based on the GS and the CG algorithms. However, we do not have a similar recipe for the case of singular FIM F_Y. At present we compute the uniform CR bound for singular F_Y by performing SVD on F_Y. A recipe to compute the uniform CR bound for singular FIM F_Y based on the perturbation or the differencing method can help to significantly reduce the computation time.

5.1.3 Is There a Better Bound for Under-determined Problems

We saw that even with the ideal weights, the WLSE fails to achieve the uniform CR bound at all bias gradient lengths. The question is: Is the CR bound given by the pseudo-inverse the best possible lower bound for the under-determined problems, or can we do better?

5.1.4 Extension of the Uniform Bound to Multiple Estimator Components

The unbiased CR bound determines the lower bound on a single estimator component. A natural next step would be to extend the bound to multi-parameters.

5.1.5 Further Increase in the Rate of Convergence using D_p via Side Information

At present, if we want increased rate of convergence of the GS algorithm, we increase the number of diagonals of the D_p matrices. A more intelligent way of choosing the additional diagonals would be to exploit the structure of the FIM, i.e. introducing diagonals where the information contents are high. We expect that by using this information we can achieve faster convergence rates for less number of diagonals.

5.1.6 $\mathcal{O}(q^2n)$ Approximation of the Uniform Bound: How (What) Good is it?

We present a computationally efficient $\mathcal{O}(q^2n)$ approximation to the uniform CR bound in Appendix F. We expect the approximation to be 'good' if there is little

correlation among different estimator components. However, we need to derive a bound on the approximation error.

APPENDIX A

PROOF OF LEMMA 1

The matrix $D_p - A$ has the form:

$$\begin{bmatrix}
\sum_{j\neq 2,\ldots,p}|a_{1,j}| & 0 & \cdots & 0 & -a_{1,p+1} & \cdots & -a_{1n} \\
0 & \sum_{j\neq 1,3,\ldots,p+1}|a_{2,j}| & & & \ddots & & \vdots \\
\vdots & \ddots & & & & \ddots & -a_{n-p-1,n} \\
 & & \ddots & & & \ddots & \\
0 & & & \sum_{\substack{j\neq k-p,\ldots,k-1 \\ j\neq k+1,\ldots,k+p-1}}|a_{k,j}| & & & 0 \\
-a_{p+1,1} & \ddots & & & \ddots & & \\
 & & \ddots & & & \ddots & \vdots \\
\vdots & & & \ddots & & \ddots & 0 \\
-a_{n1} & \cdots & -a_{n,n-p-1} & 0 & \cdots & 0 & \sum_{j\neq n-p,\ldots,n-1}|a_{n,j}|
\end{bmatrix}.$$

Since $D_p - A$ is a special case of $D_1 - A$ for $a_{ij} = 0, |i-j| \geq p$, it suffices to prove that the following matrix is Positive Semi-Definite (PSD):

$$D_1 - A = \begin{bmatrix}
\sum_{j\neq 1} | a_{1j} | & -a_{12} & \cdots & -a_{1n} \\
-a_{21} & \sum_{j\neq 2} | -a_{2j} | & \cdots & -a_{2n} \\
\vdots & \vdots & \ddots & \vdots \\
-a_{n1} & -a_{n2} & \cdots & \sum_{j\neq n} | -a_{nj} |
\end{bmatrix} \tag{A.1}$$

Define:

$$M^{(q)} = \begin{bmatrix} N & \underline{b} \\ \underline{b}^T & a \end{bmatrix},$$

where,

$$\begin{aligned}
\underline{b} &= [-a_{q1}, -a_{q2}, \ldots, -a_{q,q-1}]^T, \\
a &= \sum_{i=1}^{q-1} |b_i|, \\
N &= M^{(q-1)} + \mathrm{diag}_i(|b_i|).
\end{aligned}$$

Since $D_1 - A$ is symmetric $M^{(q)}$ is the $q \times q$ upper left hand block of the $n \times n$ matrix $D_1 - A$. We proceed by induction. Assume that $M^{(q-1)}$ is positive semi-definite. Since $a = 0$ implies $\underline{b} = \underline{0}$, when $a = 0$, $M^{(q)}$ is obviously positive semi-definite. We therefore assume $a > 0$, in which case $\underline{b}$ is not identically zero. Consider the following factorization of $M^{(q)}$.

$$M^{(q)} = \begin{bmatrix} I & \frac{1}{a}\underline{b} \\ \underline{0}^T & 1 \end{bmatrix} \begin{bmatrix} N - \frac{1}{a}\underline{b}\underline{b}^T & \underline{0} \\ \underline{0}^T & a \end{bmatrix} \begin{bmatrix} I & \underline{0} \\ \frac{1}{a}\underline{b}^T & 1 \end{bmatrix}.$$

We only need to show that the $(q-1) \times (q-1)$ matrix,

$$N - \frac{1}{a}\underline{b}\underline{b}^T = M^{q-1} + \mathrm{diag}_i(|b_i|) - \frac{1}{a}\underline{b}\underline{b}^T$$

is PSD. Since the sum of PSD matrices is PSD, it is sufficient to show that

$$\mathrm{diag}_i(|b_i|) - \frac{1}{a}\underline{b}\underline{b}^T \geq 0.$$

Consider for any $\underline{x} = [x_1, x_2, ..., x_{q-1}]^T$:

$$\begin{aligned}
\underline{x}^T \left(\mathrm{diag}_i(|b_i|) - \frac{1}{a}\underline{b}\underline{b}^T\right)\underline{x} &= \sum_{i=1}^{q-1} x_i^2|b_i| - \frac{1}{a}\left(\sum_{i=1}^{q-1} x_i b_i\right)^2 \\
&= \sum_{j=1}^{q-1} |b_j| \left(\sum_{i=1}^{q-1} x_i^2 \frac{|b_i|}{\sum_{j=1}^{q-1}|b_j|} - \left(\sum_{i=1}^{q-1} x_i \frac{b_i}{\sum_{j=1}^{q-1}|b_j|}\right)^2\right) \\
&\geq a\left(\sum_{i=1}^{q-1} x_i^2 \frac{|b_i|}{\sum_{j=1}^{q-1}|b_j|} - \left(\sum_{i=1}^{q-1} |x_i| \frac{|b_i|}{\sum_{j=1}^{q-1}|b_j|}\right)^2\right) \\
&= a\left(\sum_{i=1}^{q-1} p_i|x_i|^2 - \left(\sum_{i=1}^{q-1} p_i|x_i|\right)^2\right) \qquad \text{(A.2)} \\
&= a\sum_{i=1}^{q-1} p_i \left(|x_i| - \sum_{j=1}^{q-1} p_j|x_j|\right)^2 \\
&\geq 0.
\end{aligned}$$

Where in (A.2) we have defined $p_i = \frac{|b_i|}{\sum_{j=1}^{q-1}|b_j|}$, $p_i \in [0,1]$, $\sum_{i=1}^{q-1} p_i = 1$. Furthermore, for $\underline{x} = \underline{1}$, $(D_p - A)\underline{1} = \underline{0}$, where $\underline{1} = [1, ..., 1]^T$. This implies that $D_p - A$ has at least one zero eigenvalue, therefore $D_p - A$ is rank deficient with rank at most $n - 1$.

∎

APPENDIX B

PROOF OF LEMMA 2

Define $S(A)$ to be a set of diagonal F such that $F - A \geq 0$. Let us denote $F_{\min} \in S(A)$ as the matrix F that minimizes $|||F - A|||_{\mathcal{F}}$, and $F_{\min}(\underline{x})$ as the matrix that minimizes $|||F - A|||_{\mathcal{F}}$ over the set of diagonal F such that $\underline{x}^T(F - A)\underline{x} \geq 0$. It can be easily shown using Lagrange multiplier theory that $F_{\min}(\underline{x})$ is given by:

$$F_{\min}(\underline{x}) = \text{diag}(A) - \frac{\underline{x}^T(\text{diag}(A) - A)\underline{x}}{\underline{x}^T \text{diag}(\underline{x}\underline{x}^T)\ \underline{x}} \text{diag}(\underline{x}\underline{x}^T).$$

Clearly $\mathcal{S}(A) \subset \mathcal{S}_{\underline{x}}(A)$, since the condition $F - A \geq 0$ implies $\underline{x}^T(F - A)\underline{x} \geq 0$ for all $\underline{x}$, so by the definition of $F_{\min}$:

$$|||F_{\min}(\underline{x}) - A|||_{\mathcal{F}} \leq |||F_{\min} - A|||_{\mathcal{F}}, \ \forall \underline{x}.$$

For $\underline{x} = \underline{1}$ and A symmetric regular circulant with non-negative entries one can verify that $F_{\min}(\underline{1}) = (a_{11} + ... + a_{1n})I = D_1$. By Lemma 1, $D_1 - A \geq 0$. Therefore, $F_{\min}(\underline{1}) \in S(A)$. Thus by the definition of $F_{\min}$ we must also have,

$$|||F_{\min} - A|||_{\mathcal{F}} \leq |||F_{\min}(\underline{1}) - A|||_{\mathcal{F}}$$

The above two inequalities together imply that

$$F_{\min}(\underline{1}) = F_{\min} = D_1.$$

∎

APPENDIX C

COVARIANCE OF THE BIAS-GRADIENT ESTIMATOR FOR A WEIGHTED LEAST-SQUARES ESTIMATOR

Theorem 5 *Assume: i)* $\underline{Y}_i$ *is a set of i.i.d. Gaussian random variables with mean* $A\underline{\theta}$ *and covariance* K, *and* $\widehat{\underline{\theta}}$ *is a weighted least-squares estimator of the form:* $\widehat{\underline{\theta}}(\underline{Y}) = (A^T K^{-1} A + \alpha I)^{-1} A^T K^{-1} \underline{Y}$, *where* α *is a smoothing parameter, then a bias gradient estimator* $\widehat{\nabla}_{\underline{\theta}} b_1 = \widehat{\nabla}_{\underline{\theta}} m_1 + \underline{e}_1^T$, *where* $\widehat{\nabla}_{\underline{\theta}} m_1$ *is given in (3.15), has the covariance:*

$$S(\widehat{\nabla}_{\underline{\theta}} b_1) = \frac{1}{L}\left[\nabla_{\underline{\theta}}^T m_1 \nabla_{\underline{\theta}} m_1 + var_{\underline{\theta}}(\widehat{\theta}_1) F_Y + \frac{1}{L}\sum_{i=1}^{L} E_{\underline{\theta}}(m_1 - \zeta_i)^2 F_Y\right],$$

where $F_Y = A^T K^{-1} A$ *is the Fisher information matrix.*

Proof: Since, $\underline{Y}_i \sim \mathcal{N}(A\underline{\theta}, K)$, $\nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{y};\underline{\theta}) = (\underline{Y} - A\underline{\theta})^T K^{-1} A = (\underline{Y} - \overline{\underline{Y}})^T K^{-1} A$, where $\overline{\underline{Y}} = E(\underline{Y}) = A\underline{\theta}$ is the mean of $\underline{Y}$. Since we are interested in the estimator component $\widehat{\theta}_1$:

$$\widehat{\theta}_1(\underline{Y}) = \underline{e}_1^T (A^T K^{-1} A + \alpha I)^{-1} A^T K^{-1} \underline{Y} = \underline{v}_1{}^T \underline{Y},$$

where $\underline{e}_1 = [1, 0, 0, ..., 0]^T$, and $\underline{v}_1 \stackrel{def}{=} K^{-1} A \, (A^T K^{-1} A + \alpha I)^{-1} \underline{e}_1$. Now,

$$\begin{aligned} S\left(\widehat{\nabla}_{\underline{\theta}} b_1\right) = \mathrm{cov}_{\underline{\theta}}\left(\widehat{\nabla}_{\underline{\theta}} b_1\right) &= \mathrm{cov}_{\underline{\theta}}\left(\frac{1}{L}\sum_{i=1}^{L}(\widehat{\theta}_1(\underline{Y}_i) - \zeta_i)\, \nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta}) - \underline{e}_1^T\right) \\ &= \mathrm{cov}_{\underline{\theta}}\left(\frac{1}{L}\sum_{i=1}^{L}(\widehat{\theta}_1(\underline{Y}_i) - \zeta_i)\, \nabla_{\underline{\theta}} \ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta})\right) \\ &= \frac{1}{L^2} A^T K^{-1} \underbrace{\mathrm{cov}_{\underline{\theta}}\left[\sum_{i=1}^{L}(\underline{Y}_i - \overline{\underline{Y}})(\widehat{\theta}_1(\underline{Y}_i) - \zeta_i)^T\right]}_{\mathrm{A}} K^{-1} A. \end{aligned}$$

Consider the underlined term A:

$$\mathrm{cov}_{\underline{\theta}}\left[\sum_{i=1}^{L}(\underline{Y}_i - \overline{\underline{Y}})(\widehat{\theta}_1(\underline{Y}_i) - \zeta_i)^T\right]$$

$$
\begin{aligned}
&= \operatorname{cov}_{\underline{\theta}}\left[\sum_{i=1}^{L}(\underline{Y}_i-\overline{\underline{Y}})(\underline{v}_1{}^T\underline{Y}_i-\zeta_i)^T\right]\\
&= \operatorname{cov}_{\underline{\theta}}\left[\sum_{i=1}^{L}(\underline{Y}_i-\overline{\underline{Y}})(\underline{v}_1{}^T\underline{Y}_i-\underline{v}_1{}^T\overline{\underline{Y}}+\underline{v}_1{}^T\overline{\underline{Y}}-\zeta_i)^T\right]\\
&= \underbrace{\operatorname{cov}_{\underline{\theta}}\left[\sum_{i=1}^{L}(\underline{Y}_i-\overline{\underline{Y}})\underline{v}_1{}^T(\underline{Y}_i-\overline{\underline{Y}})\right]}_{\mathrm{B}}+\underbrace{\operatorname{cov}_{\underline{\theta}}\left[\sum_{i=1}^{L}(\underline{Y}_i-\overline{\underline{Y}})(\underline{v}_1{}^T\overline{\underline{Y}}-\zeta_i)\right]}_{\mathrm{C}}. \qquad \text{(C.1)}
\end{aligned}
$$

Equality in (C.1) is due to the fact that $\underline{Y}-\overline{\underline{Y}}$ is zero mean and the cross term is zero since 3rd moment of a zero mean Gaussian is zero. We will solve Expression B and Expression C separately.

Expression B:

Note that, $\operatorname{cov}_{\underline{\theta}}\left[\sum_{i=1}^{L}(\underline{Y}_i-\overline{\underline{Y}})\underline{v}_1{}^T(\underline{Y}_i-\overline{\underline{Y}})\right]=L\operatorname{cov}_{\underline{\theta}}\left[(\underline{Y}-\overline{\underline{Y}})\underline{v}_1{}^T(\underline{Y}-\overline{\underline{Y}})\right]$. Let $\underline{Z}=\underline{Y}-\overline{\underline{Y}}, \underline{Z}\sim\mathcal{N}(\underline{0},K)$.

$$
\begin{aligned}
\operatorname{cov}_{\underline{\theta}}\left[(\underline{Y}-\overline{\underline{Y}})\underline{v}_1{}^T(\underline{Y}-\overline{\underline{Y}})\right] &= \operatorname{cov}_{\underline{\theta}}\left[\underline{Z}\underline{v}_1{}^T\underline{Z}\right]\\
&= \operatorname{cov}_{\underline{\theta}}\left[\underline{Z}\underline{Z}^T\underline{v}_1\right]\\
&= E_{\underline{\theta}}\left[(\underline{Z}\underline{Z}^T)\underline{v}_1\underline{v}_1{}^T(\underline{Z}\underline{Z}^T)^T\right]-E_{\underline{\theta}}\left[\underline{Z}\underline{Z}^T\underline{v}_1\right]E_{\underline{\theta}}\left[\underline{Z}\underline{Z}^T\underline{v}_1\right]^T\\
&= E_{\underline{\theta}}\left[(\underline{Z}\underline{Z}^T)\underline{v}_1\underline{v}_1{}^T(\underline{Z}\underline{Z}^T)^T\right]-K(\underline{v}_1\underline{v}_1{}^T)K\\
&= KWK+K\sum_{i=1}^{n}W_{ii}K_{ii}, \qquad \text{(C.2)}
\end{aligned}
$$

where $W=\underline{v}_1\underline{v}_1{}^T$. (C.2) is due to the following lemma.

Lemma 3 *Given an $n\times 1$ random vector $\underline{Z}=[Z_1,...,Z_n]^T$ of zero mean i.i.d Gaussian random variables, such that $\underline{Z}\sim\mathcal{N}(\underline{0},K)$, and an $n\times 1$ constant vector $\underline{v}_1$, then,*

$$
E_{\underline{\theta}}(\underline{Z}\underline{Z}^T\underline{v}_1\underline{v}_1{}^T\underline{Z}\underline{Z}^T) = 2E_{\underline{\theta}}(\underline{Z}\underline{Z}^T)\underline{v}_1\underline{v}_1{}^TE_{\underline{\theta}}(\underline{Z}\underline{Z}^T)+\underline{v}_1{}^TE_{\underline{\theta}}(\underline{Z}\underline{Z}^T)\underline{v}_1\,E_{\underline{\theta}}(\underline{Z}\underline{Z}^T).
$$

Proof: We have,

$$
\underline{Z}\underline{Z}^T=\begin{bmatrix} Z_1^2 & Z_1Z_2 & \cdots & Z_1Z_n\\ Z_1Z_2 & Z_2^2 & & Z_2Z_n\\ \vdots & & \ddots & \vdots\\ Z_1Z_n & Z_2Z_n & \cdots & Z_n^2\end{bmatrix},
$$

and $\underline{v}_1\underline{v}_1{}^T=W=((w))_{ij}:i=1,...,n,j=1,...,n$. From which,

$$
E_{\underline{\theta}}(\underline{Z}\underline{Z}^T\underline{v}_1\underline{v}_1{}^T\underline{Z}\underline{Z}^T)=\begin{bmatrix} E_{\underline{\theta}}(\sum_{i=1}^{n}w_{ii}Z_i^2Z_1^2) & 2E_{\underline{\theta}}(w_{12}Z_1^2Z_2^2) & \cdots & 2E_{\underline{\theta}}(w_{1n}Z_1^2Z_n^2)\\ 2E_{\underline{\theta}}(w_{12}Z_1^2Z_2^2) & E_{\underline{\theta}}(\sum_{i=1}^{n}w_{ii}Z_i^2Z_2^2) & & 2E_{\underline{\theta}}(w_{2n}Z_2^2Z_n^2)\\ \vdots & & \ddots & \vdots\\ 2E_{\underline{\theta}}(w_{1n}Z_1^2Z_n^2) & 2E_{\underline{\theta}}(w_{2n}Z_2^2Z_n^2) & \cdots & E_{\underline{\theta}}(\sum_{i=1}^{n}w_{ii}Z_i^2Z_n^2)\end{bmatrix}.
$$

Using the familiar expression for the 4-th moment of zero mean Gaussian random variables [22, 37]:

$$E_{\underline{\theta}}(Z_1Z_2Z_3Z_4) = E_{\underline{\theta}}(Z_1Z_2)E_{\underline{\theta}}(Z_3Z_4) + E_{\underline{\theta}}(Z_1Z_3)E_{\underline{\theta}}(Z_2Z_4) + E_{\underline{\theta}}(Z_1Z_4)E_{\underline{\theta}}(Z_2Z_3),$$

we immediately get the required result.

■

Expression C:

$$\begin{aligned}\mathrm{cov}_{\underline{\theta}}\left[\sum_{i=1}^{L}(\underline{Y}_i - \overline{\underline{Y}})(\underline{v}_1{}^T\overline{\underline{Y}} - \zeta_i)\right] &= \sum_{i=1}^{L} E_{\underline{\theta}}\left[(\underline{v}_1{}^T\overline{\underline{Y}} - \zeta_i)^2(\underline{Y}_i - \overline{\underline{Y}})(\underline{Y}_i - \overline{\underline{Y}})^T\right] \\ &= \sum_{i=1}^{L} E_{\underline{\theta}}\left[(\underline{v}_1{}^T\overline{\underline{Y}} - \zeta_i)^2\right] K. \qquad \text{(C.3)}\end{aligned}$$

Substituting (C.1) and (C.3) in (C.1) we get:

$$\begin{aligned}\mathrm{cov}_{\underline{\theta}}(\widehat{\nabla_{\underline{\theta}}\underline{b}_1}) &= \frac{1}{L}A^TK^{-1}\left[KWK + K\sum_{i=1}^{n}W_{ii}K_{ii} + \frac{1}{L}\sum_{i=1}^{L}E_{\underline{\theta}}(\underline{v}_1{}^T\overline{\underline{Y}} - \zeta_i)^2K\right]K^{-1}A \\ &= \frac{1}{L}\left[A^TWA + K\sum_i W_{ii}F_Y + \frac{1}{L}\sum_{i=1}^{L}E_{\underline{\theta}}(\underline{v}_1{}^T\overline{\underline{Y}} - \zeta_i)^2F_Y\right] \\ &= \frac{1}{L}\left[F_YF_{Y\alpha}^{-1}\underline{e}_1\underline{e}_1^TF_{Y\alpha}^{-1} + \underline{e}_1^TF_{Y\alpha}^{-1}F_YF_{Y\alpha}^{-1}\underline{e}_1\, I + \frac{1}{L}\sum_{i=1}^{L}E_{\underline{\theta}}\left(\underline{e}_1^TF_{Y\alpha}^{-1}F_Y\underline{\theta} - \zeta_i\right)^2 I\right]F_Y,\end{aligned}$$

where $F_Y = A^TK^{-1}A$ and $F_{Y\alpha}^{-1} = (F_Y + \alpha I)^{-1}$. It directly follows that:

$$S(\widehat{\nabla_{\underline{\theta}}\underline{b}_1}) = \mathrm{cov}_{\underline{\theta}}(\widehat{\nabla_{\underline{\theta}}\underline{b}_1}) = \frac{1}{L}\left[\nabla_{\underline{\theta}}^Tm_1\nabla_{\underline{\theta}}m_1 + \mathrm{var}_{\underline{\theta}}(\hat{\theta}_1)F_Y + \frac{1}{L}\sum_{i=1}^{L}E_{\underline{\theta}}(m_1 - \zeta_i)^2F_Y\right],$$

where $\nabla_{\underline{\theta}}m_1 = \underline{e}_1^TF_{Y\alpha}^{-1}F_Y$, $\mathrm{var}_{\underline{\theta}}(\hat{\theta}_1) = \underline{e}_1^TF_{Y\alpha}^{-1}F_YF_{Y\alpha}^{-1}\underline{e}_1$, and $m_1 = \underline{e}_1^TF_{Y\alpha}^{-1}F_Y\underline{\theta}$.

■

C.0.7 Bias of the Bias-Gradient Estimator with $\hat{\zeta}_i = \frac{1}{L}\sum_{j=1}^{L}\hat{\theta}_1(\underline{Y}_j)$

The bias of the bias gradient estimator is given by,

$$\underline{b}\left(\widehat{\nabla_{\underline{\theta}}\underline{b}}\right) = E_{\underline{\theta}}\left(\widehat{\nabla_{\underline{\theta}}\underline{b}}\right) + \nabla_{\underline{\theta}}\underline{b}. \qquad \text{(C.4)}$$

For WLSE it is easy to show that the bias gradient is;

$$\nabla_{\underline{\theta}}\underline{b} = -I + (F_Y + \alpha I)^{-1}F_Y. \qquad \text{(C.5)}$$

Now,

$$
\begin{aligned}
& E_{\underline{\theta}}\left(\widehat{\nabla_{\underline{\theta}}\underline{b}}\right) \\
= \; & E_{\underline{\theta}}\left\{\frac{1}{L}\sum_{i=1}^{L}\left[(F_Y+\alpha I)^{-1}A^TK^{-1}\underline{Y}_i - \frac{1}{L}\sum_{j=1}^{L}(F_Y+\alpha I)^{-1}A^TK^{-1}\underline{Y}_j\right](\underline{Y}_i-\overline{\underline{Y}})^TK^{-1}A\right\} - I \\
= \; & (F_Y+\alpha I)^{-1}A^TK^{-1}\left[E_{\underline{\theta}}(\underline{Y}_i\underline{Y}_i{}^T) - E_{\underline{\theta}}(\underline{Y}_i)\overline{\underline{Y}}^T\right]K^{-1}A \\
- \; & \frac{1}{L^2}\sum_{i=1}^{L}\sum_{j=1}^{L}(F_Y+\alpha I)^{-1}A^TK^{-1}\left[E_{\underline{\theta}}(\underline{Y}_j\underline{Y}_i{}^T) - E_{\underline{\theta}}(\underline{Y}_j)\overline{\underline{Y}}^T\right]K^{-1}A - I \\
= \; & (F_Y+\alpha I)^{-1}F_Y - \frac{1}{L}(F_Y+\alpha I)^{-1}F_Y - I \qquad \text{(C.6)}
\end{aligned}
$$

Substituting (C.6) and (C.5) in (C.4),

$$\underline{b}\left(\widehat{\nabla_{\underline{\theta}}\underline{b}}\right) = -\frac{1}{L}(F_Y+\alpha I)^{-1}F_Y.$$

For large L or for large α the bias goes to zero.

C.0.8 Covariance of the Bias-Gradient with $\hat{\zeta}_i = \frac{1}{L-1}\sum_{\substack{j=1\\ j\neq i}}^{L}\hat{\theta}_1(\underline{Y}_j)$

Theorem 6 *The bias gradient estimator given by:*

$$\widehat{\nabla_{\underline{\theta}}\underline{b}_1} = \frac{1}{L}\sum_{i=1}^{L}\left[\hat{\theta}_1(\underline{Y}_i) - \frac{1}{L-1}\sum_{\substack{j=1\\ j\neq i}}^{L}\hat{\theta}_1(\underline{Y}_j)\right]\nabla_{\underline{\theta}}\ln f_{\underline{Y}}(\underline{Y}_i;\underline{\theta}) - \underline{e}_1$$

achieves minimum covariance $S_{\min}(\widehat{\nabla_{\underline{\theta}}\underline{b}_1})$ *for large* L *and the assumptions given in Theorem 5.*

Proof:

Clearly for $L = 1, \zeta = 0$ and we get the same result as (3.17). In the sequel we will calculate the covariance for $L > 1$. Substituting $\nabla_{\underline{\theta}}\ln f_{\underline{Y}}(\underline{Y};\underline{\theta}) = (\underline{Y}-\overline{\underline{Y}})^T\ K^{-1}A$ we obtain:

$$\mathrm{cov}_{\underline{\theta}}\left(\widehat{\nabla_{\underline{\theta}}\underline{b}_1(\underline{\theta})}\right) = \frac{1}{L^2}A^TK^{-1}\underbrace{\mathrm{cov}_{\underline{\theta}}\left\{\sum_{i=1}^{L}\left[\hat{\theta}_1(\underline{Y}_i) - \frac{1}{L-1}\sum_{\substack{j=1\\ j\neq i}}^{L}\hat{\theta}_1(\underline{Y}_j)\right](\underline{Y}_i-\overline{\underline{Y}})^T\right\}}_{\mathrm{D}}K^{-1}A.$$

Consider the underlined term D:

$$\mathrm{cov}_{\underline{\theta}}\left\{\sum_{i=1}^{L}\left[\hat{\theta}_1(\underline{Y}_i) - \frac{1}{L-1}\sum_{\substack{j=1\\ j\neq i}}^{L}\hat{\theta}_1(\underline{Y}_j)\right](\underline{Y}_i-\overline{\underline{Y}})^T\right\}$$

$$= \text{cov}_{\underline{\theta}} \left\{ \underline{v}_1{}^T \sum_{i=1}^{L} \left[\underline{Y}_i - \frac{1}{L-1} \sum_{\substack{j=1 \\ j \neq i}}^{L} \underline{Y}_j \right] (\underline{Y}_i - \overline{\underline{Y}})^T \right\}$$

$$= \text{cov}_{\underline{\theta}} \left\{ \underline{v}_1{}^T \sum_{i=1}^{L} \left[\underline{Y}_i - \overline{\underline{Y}} + \overline{\underline{Y}} - \frac{1}{L-1} \sum_{\substack{j=1 \\ j \neq i}}^{L} \underline{Y}_j \right] (\underline{Y}_i - \overline{\underline{Y}})^T \right\}$$

$$= \text{cov}_{\underline{\theta}} \left\{ \left[\underline{v}_1{}^T \sum_{i=1}^{L} (\underline{Y}_i - \overline{\underline{Y}})(\underline{Y}_i - \overline{\underline{Y}})^T + \underline{v}_1{}^T \sum_{i=1}^{L} \left(\overline{\underline{Y}} - \frac{1}{L-1} \sum_{\substack{j=1 \\ j \neq i}}^{L} \underline{Y}_j \right) \right] (\underline{Y}_i - \overline{\underline{Y}})^T \right\}$$

$$= \text{cov}_{\underline{\theta}} \left\{ \underline{v}_1{}^T \sum_{i=1}^{L} (\underline{Y}_i - \overline{\underline{Y}})(\underline{Y}_i - \overline{\underline{Y}})^T - \frac{\underline{v}_1{}^T}{L-1} \sum_{i=1}^{L} \sum_{\substack{j=1 \\ j \neq i}}^{L} (\underline{Y}_j - \overline{\underline{Y}})(\underline{Y}_i - \overline{\underline{Y}})^T \right\}$$

$$= \text{cov}_{\underline{\theta}} \left\{ \underline{v}_1{}^T \sum_{i=1}^{L} (\underline{Y}_i - \overline{\underline{Y}})(\underline{Y}_i - \overline{\underline{Y}})^T \right\} + \underbrace{\text{cov}_{\underline{\theta}} \left\{ \frac{\underline{v}_1{}^T}{L-1} \sum_{i=1}^{L} \sum_{\substack{j=1 \\ j \neq i}}^{L} (\underline{Y}_i - \overline{\underline{Y}})(\underline{Y}_j - \overline{\underline{Y}})^T \right\}}_{\text{E}} \quad \text{(C.7)}$$

(C.7) is due to our assumption that different realizations of $\underline{Y}$ are i.i.d. and $\underline{Y}_i - \overline{\underline{Y}}$ is zero mean. The first term is already computed in the proof of Theorem 5. We will concentrate on the underlined term E.

$$\text{cov}_{\underline{\theta}} \left\{ \frac{\underline{v}_1{}^T}{L-1} \sum_{i=1}^{L} \sum_{\substack{j=1 \\ j \neq i}}^{L} (\underline{Y}_j - \overline{\underline{Y}})(\underline{Y}_i - \overline{\underline{Y}})^T \right\}$$

$$= \frac{1}{(L-1)^2} \text{cov}_{\underline{\theta}} \left\{ \underline{v}_1{}^T \sum_{i=1}^{L} \sum_{\substack{j=1 \\ j \neq i}}^{L} (\underline{Y}_j - \overline{\underline{Y}})(\underline{Y}_i - \overline{\underline{Y}})^T \right\}$$

$$= \frac{1}{(L-1)^2} \text{cov}_{\underline{\theta}} \left\{ \sum_{i=1}^{L} \sum_{\substack{j=1 \\ j \neq i}}^{L} (\underline{Y}_i - \overline{\underline{Y}})(\underline{Y}_j - \overline{\underline{Y}})^T \underline{v}_1 \right\}$$

$$= \frac{1}{(L-1)^2} E_{\underline{\theta}} \left\{ \sum_{i=1}^{L} \sum_{\substack{j=1 \\ j \neq i}}^{L} \sum_{k=1}^{L} \sum_{\substack{l=1 \\ l \neq k}}^{L} (\underline{Y}_i - \overline{\underline{Y}})(\underline{Y}_j - \overline{\underline{Y}})^T \underline{v}_1 \underline{v}_1{}^T (\underline{Y}_k - \overline{\underline{Y}})(\underline{Y}_l - \overline{\underline{Y}})^T \right\}$$

$$= \frac{2L(L-1)}{(L-1)^2} E_{\underline{\theta}} \left\{ (\underline{Y}_i - \overline{\underline{Y}})(\underline{Y}_j - \overline{\underline{Y}})^T \underline{v}_1 \underline{v}_1{}^T (\underline{Y}_j - \overline{\underline{Y}})(\underline{Y}_i - \overline{\underline{Y}})^T \right\} \quad \text{(C.8)}$$

$$= \frac{2L(L-1)}{(L-1)^2} \left[KWK + K \sum_{i=1}^{n} W_{ii} K_{ii} \, I - K \, \text{diag}(W) \, K \right] \quad \text{(C.9)}$$

(C.8) is due to the assumption that $\underline{Y}_i$ are i.i.d. and (C.9) is due to a similar derivation already given in the proof of Theorem 5. Substituting (C.9) in (C.7), we get:

$$\text{cov}_{\underline{\theta}} \left(\widehat{\nabla_{\underline{\theta}} \underline{b}_1} \right) = \frac{1}{L^2} A^T K^{-1} \left\{ L \left(KWK + K \sum_{i} W_{ii} K_{ii} \right) + \frac{2L}{L-1} \left[KWK + K \sum_{i} W_{ii} K_{ii} \, I - K \, \text{diag}(W) \, K \right] \right\} K^{-1} A$$

$$
\begin{aligned}
&= \frac{1}{L}\left[F_Y F_{Ya}^{-1}\underline{e}_1\underline{e}_1^T F_{Ya}^{-1}F_Y + \underline{e}_1^T F_{Ya}^{-1}F_Y F_{Ya}^{-1}\underline{e}_1\ F_Y\right] \\
&+ \frac{2}{L(L-1)}\left[F_Y F_{Ya}^{-1}\underline{e}_1\underline{e}_1^T F_{Ya}^{-1}F_Y + \left(\underline{e}_1^T F_{Ya}^{-1}F_Y F_{Ya}^{-1}\underline{e}_1\right) F_Y\right. \\
&- \left. A^T \mathrm{diag}\left(K^{-1}AF_{Ya}^{-1}\underline{e}_1\underline{e}_1^T F_{Ya}^{-1}A^T K^{-1}\right) A\right],
\end{aligned}
$$

where $\mathrm{diag}(W) = \mathrm{diag}(w_{11}, w_{22}, ..., w_{nn})$ is an $n \times n$ diagonal matrix. Therefore, we obtain:

$$
\begin{aligned}
&\mathrm{cov}_{\underline{\theta}}(\widehat{\nabla}_{\underline{\theta}} b_1) \\
&= \begin{cases}
\frac{1}{L}\left[F_Y F_{Ya}^{-1}\underline{e}_1\underline{e}_1^T F_{Ya}^{-1} + \underline{e}_1^T F_{Ya}^{-1}F_Y F_{Ya}^{-1}\underline{e}_1\ I + \left(\underline{e}_1^T F_{Ya}^{-1}F_Y\underline{\theta}\right)^2 I\right] F_Y & L = 1 \\
\frac{1}{L}\left[F_Y F_{Ya}^{-1}\underline{e}_1\underline{e}_1^T F_{Ya}^{-1}F_Y + \underline{e}_1^T F_{Ya}^{-1}F_Y F_{Ya}^{-1}\underline{e}_1\ F_Y\right] & \\
+\frac{2}{L(L-1)}\left[F_Y F_{Ya}^{-1}\underline{e}_1\underline{e}_1^T F_{Ya}^{-1}F_Y + \left(\underline{e}_1^T F_{Ya}^{-1}F_Y F_{Ya}^{-1}\underline{e}_1\right) F_Y\right. & \\
\left. - A^T \mathrm{diag}\left(K^{-1}AF_{Ya}^{-1}\underline{e}_1\underline{e}_1^T F_{Ya}^{-1}A^T K^{-1}\right) A\right]. & L > 1
\end{cases}
\end{aligned}
$$

Clearly for large L:

$$
\begin{aligned}
S(\widehat{\nabla}_{\underline{\theta}} b_1) = \mathrm{cov}_{\underline{\theta}}(\widehat{\nabla}_{\underline{\theta}} b_1) &= \frac{1}{L}\left[F_Y F_{Ya}^{-1}\underline{e}_1\underline{e}_1^T F_{Ya}^{-1} + \underline{e}_1^T F_{Ya}^{-1}F_Y F_{Ya}^{-1}\underline{e}_1\ I\right] F_Y \\
&= \frac{1}{L}\left(\nabla_{\underline{\theta}}^T m_1 \nabla_{\underline{\theta}} m_1 + \mathrm{var}_{\underline{\theta}}(\widehat{\theta}_1)\right) F_Y \\
&= S_{\min}(\widehat{\nabla}_{\underline{\theta}} b_1)
\end{aligned}
$$

■

APPENDIX D

DERIVATION OF TRANSITION DENSITY FOR THE VERTEX VIEW

In this section we derive the transition density $f_{\underline{W}/\underline{V}}(\underline{w}/\underline{v})$ for the vertex view. Consider the geometry in Figure D.1. R denotes the distance between the source plane and the vertex view. A γ-ray photon is emitted at A and detected at P. The path of the emitted γ-rays is completely determined by angles ζ and ψ. We assume that the γ-rays are emitted in all directions with equal probability: i.e.

$$f(\zeta/\underline{v}) = \begin{cases} \frac{1}{2\pi} & \zeta \in [0, 2\pi] \\ 0 & \text{Otherwise} \end{cases}$$

$$f(\psi/\underline{v}) = \begin{cases} \frac{1}{\pi} & \psi \in [0, \pi] \\ 0 & \text{Otherwise.} \end{cases}$$

We also assume independence of ζ and ψ. From the geometry in Figure D.1, we have,

$$AC = \frac{R}{\tan(\psi)}$$

Also,

$$\begin{aligned} w_1 - v_1 &= AC\cos(\zeta) \\ w_2 - v_2 &= AC\sin(\zeta) \end{aligned}$$

From which we get,

$$\begin{aligned} \psi &= \tan^{-1}\left(\frac{R}{\sqrt{(w_1 - v_1)^2 + (w_2 - v_2)^2}}\right) \\ \zeta &= \tan^{-1}\left(\frac{w_2 - v_2}{w_1 - v_1}\right) \end{aligned}$$

We use the Jacobians to make the transformation from ζ and ψ in the source plane $\mathcal{V}$, to the detector plane $\mathcal{W}$.

$$f_{\underline{W}/\underline{V}}(\underline{w}/\underline{v}) = f(\zeta/\underline{v})\, f(\psi/\underline{v}) \left|\frac{d\zeta}{dw_1}\frac{d\psi}{dw_2} - \frac{d\zeta}{dw_2}\frac{d\psi}{dw_1}\right| \qquad \text{(D.1)}$$

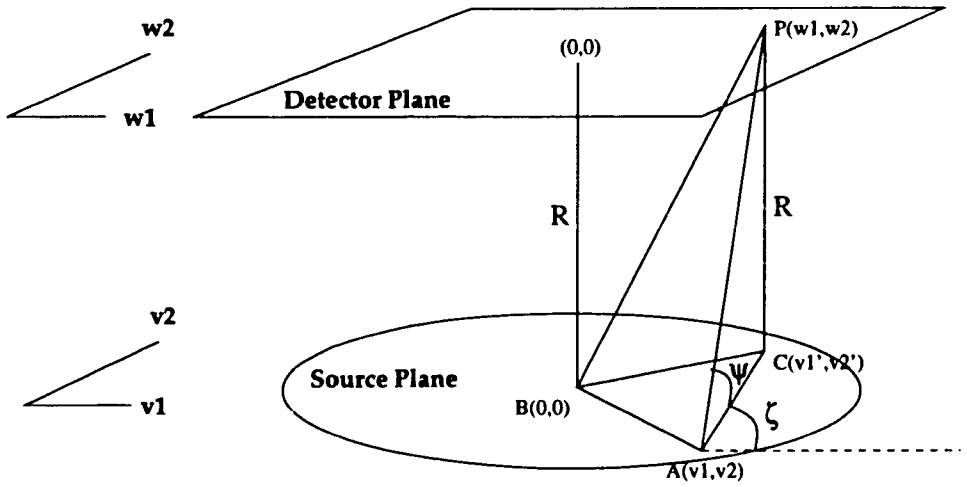

Figure D.1: System geometry

It is straightforward to show that:

$$\begin{aligned}
\frac{d\psi}{dw_1} &= \frac{-R[(w_1-v_1)^2+(w_2-v_2)^2]^{-1/2}\,(w_1-v_1)}{R^2+(w_1-v_1)^2+(w_2-v_2)^2} \\
\frac{d\psi}{dw_2} &= \frac{-R[(w_1-v_1)^2+(w_2-v_2)^2]^{-1/2}\,(w_2-v_2)}{R^2+(w_1-v_1)^2+(w_2-v_2)^2} \\
\frac{d\zeta}{dw_1} &= \frac{-(w_2-v_2)}{(w_1-v_1)^2+(w_2-v_2)^2} \\
\frac{d\zeta}{dw_2} &= \frac{-(w_1-v_1)}{(w_1-v_1)^2+(w_2-v_2)^2}.
\end{aligned}$$

Substituting the above differential in (D.1), after some algebraic manipulations, we obtain:

$$f_{\underline{W}/\underline{V}}(\underline{w}/\underline{v}) = \frac{1}{2\pi^2}\frac{R}{\sqrt{(w_1-v_1)^2+(w_2-v_2)^2}} \times \frac{1}{R^2+(w_1-v_1)^2+(w_2-v_2)^2}.$$

Which is the desired pdf. It can be easily verified that the integral of this density

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f_{\underline{W}/\underline{V}}(\underline{w}/\underline{v})\;\; dw_1\, dw_2 = \frac{1}{2}.$$

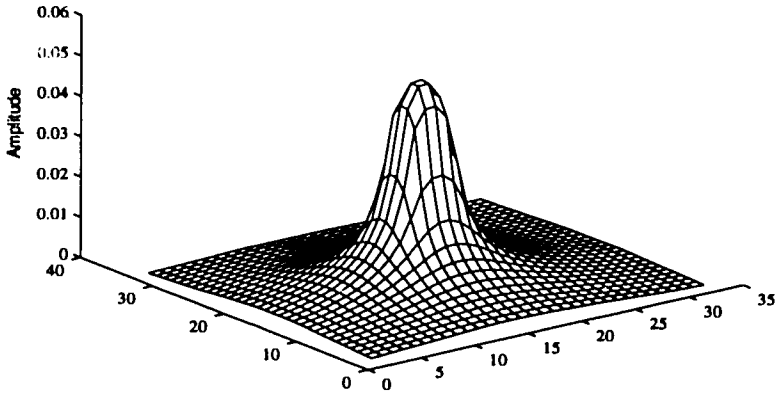

Figure D.2: Transition density f(w/v).

APPENDIX E

SYSTEM SPECIFICATIONS

Ring Detector

Radius of the detector ring	25 cms
Number of detectors	512
Radius of the collimator ring	17 cms
Number of collimator slits	10 (uniformly spaced)
Slit Width	2.4 mm

Table E.1: SPRINT ring parameters.

The SPRINT ring parameters are given in Table E.1. The SPRINT ring has a resolution of 13 mm at the center of the ring.

Vertex View Detector

Detector Dimensions	22 cms × 22 cms
Number of detectors	5476
Collimator thickness	2 cms and 4 cms
Number of collimator openings	31420
Distance between opposite faces	0.19 cms

Table E.2: Vertex view parameters.

The vertex view parameters are given in Table E.2. The 2 cms collimator has a resolution of 13 mm and detection probability of 5.25×10^{-4} at 10 cms from the vertex view collimator, while the same for the 4 cms collimator are 6.5 mm and 1.31×10^{-4}, respectively.

APPENDIX F

UNIFORM CR BOUND REQUIRING ONLY $\mathcal{O}(Q^2N)$ FLOPS

Our aim is to derive a bound on $\hat{\theta}_1$ by using only q columns of F_Y^{-1}, Rewriting (3.1):

$$\mathrm{cov}_{\underline{\theta}}(\underline{\hat{\theta}}) \geq [\nabla \underline{m}(\underline{\theta})]\, F_Y^{-1}(\underline{\theta})\, [\nabla \underline{m}(\underline{\theta})]^T$$

Define,

$$F_Y^{-1} = \begin{bmatrix} D & C^T \\ C & G \end{bmatrix},$$

$$\nabla \underline{m}(\underline{\theta}) = M = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix}$$

where D and M_{11} are $q \times q$ matrices, C and M_{21} are $(n-q) \times q$ matrices and G and M_{22} are $(n-q) \times (n-q)$ matrices. Then the biased bound using only q pixels of F_Y^{-1} is:

$$\begin{aligned} \mathrm{cov}_{\underline{\theta}}(\underline{\hat{\theta}}^{ROI}) &\geq [M_{11}, M_{12}] \begin{bmatrix} D & C^T \\ C & G \end{bmatrix} \begin{bmatrix} M_{11}^T \\ M_{12}^T \end{bmatrix} \\ &= (M_{11} + M_{12}\, C\, D^{-1})\, D\, (M_{11} + M_{12}\, C\, D^{-1})^T + M_{12}\, (G - C\, D^{-1}\, C^T)\, M_{12}^T. \end{aligned}$$

Since $G - C\, D^{-1}\, C^T$ is positive semi-definite, we have a weaker but valid bound if we neglect the second term.

$$\begin{aligned} \mathrm{cov}_{\underline{\theta}}(\underline{\hat{\theta}}^{ROI}) &\geq (M_{11} + M_{12}\, C\, D^{-1})\, D\, (M_{11} + M_{12}\, C\, D^{-1})^T \\ &= [M_{11}\ M_{12}]\, J^T\, D\, J\, [M_{11}\ M_{12}]^T, \end{aligned}$$

where $J \stackrel{def}{=} [I_q\,,\, D^{-1}\, C^T]$ is a constant $q \times n$ matrix, and I_q is a $q \times q$ identity matrix. Therefore a weaker biased CR bound on the variance of $\hat{\theta}_1$ that involves only the first q columns $[D, C^T]^T$ of F_Y^{-1} is given by:

$$\mathrm{var}_{\underline{\theta}}(\hat{\theta}_1) = \underline{e}_1^T\, \mathrm{cov}_{\underline{\theta}}(\underline{\hat{\theta}}^{ROI})\, \underline{e}_1^T$$

$$\geq \underline{e}_1^T [M_{11} M_{12}] J^T D J [M_{11} M_{12}]^T \underline{e}_1^T$$

$$= (\underline{e}_1 + \underline{d}^*)^T D (\underline{e}_1 + \underline{d}^*),$$

Where $\underline{d}^* = \underline{e}_1^T \nabla_{\underline{\theta}} \underline{b}(\underline{\theta})$. Following the method outlined in Section 3.3 we need to perform the following minimization.

$$\text{var}_{\underline{\theta}}(\hat{\theta}_1) \geq B^*(\underline{\theta}, \delta)$$

$$= \min_{\underline{d}^* : \|\underline{d}^*\|^2 \leq \delta^2} (\underline{e}_1 + \underline{d}^*)^T J^T D J (\underline{e}_1 + \underline{d}^*).$$

After a little algebra we get,

$$\underline{d}^*_{min} = -(J^T D J + \lambda I_n)^{-1} (J^T D J) \underline{e}_1, \tag{F.1}$$

where I_n is an $n \times n$ identity matrix. The form of (F.1) is not useful since it requires inversion of an $n \times n$ matrix. We use the following matrix inversion equality:

$$(J^T D J + \lambda I_n)^{-1} J^T D = J^T (J J^T + \lambda D^{-1})^{-1}. \tag{F.2}$$

The above equality can be easily proven by multiplying on the left by $(J^T D J + \lambda I_n)$ and on the right by $(J J^T + \lambda D^{-1})$. Using (F.2) we can write equation (F.1) as,

$$\underline{d}^*_{min} = -J^T (J J^T + \lambda D^{-1})^{-1} J \underline{e}_1.$$

We state the result in the form of the following theorem.

Theorem 7 *Let θ_1 be an estimator with bias $b_1(\underline{\theta})$ whose n-element bias gradient vector $\nabla_{\underline{\theta}} b_1 = \underline{g}^T$ satisfies $\underline{d}^{*T} \underline{d}^* \leq \delta^2$. Then the variance of θ_1 satisfies:*

$$var_{\underline{\theta}}(\hat{\theta}_1) \geq B^*(\underline{\theta}, \delta), \tag{F.3}$$

where $B^*(\underline{\theta}, \delta)$ is equal to:

$$B^*(\underline{\theta}, \delta) = [\underline{e}_1 + \underline{d}^*_{min}]^T J^T D J [\underline{e}_1 + \underline{d}^*_{min}]. \tag{F.4}$$

Here $\underline{e}_1 = [1, 0, ..., 0]^T$ is an n-element unit vector and:

$$\underline{d}^*_{min} = -J^T (J J^T + \lambda D^{-1})^{-1} J \underline{e}_1 \tag{F.5}$$

where λ is given by the unique non-negative solution of the following equation involving the monotone decreasing, convex function $g(\lambda) \in [0, 1]$:

$$g(\lambda) = \underline{d}^{*T}_{min} \underline{d}^*_{min} = \delta^2 \qquad \lambda \geq 0.$$

Obviously if the first column of F_Y has only diagonal entries then the above theorem yields exact bound for $q = 1$. The main load of computing $B^*(\underline{\theta}, \delta)$ is the matrix-matrix multiplication JJ^T, requiring $2q^2n$ flops.

BIBLIOGRAPHY

[1] S.V. Schell
The Cramer-Rao Lower Bound for Directions of Arrival of Gaussian Cyclo-stationary Signals, IEEE Transactions on Information Theory, 38, pp 1418-22, 1992.

[2] E. Caianiello
Systems Theory, Uncertainty and Quantum Mechanics., Selected Topics in Statistical Mechanics, pp 42-60, Dubna, 1988.

[3] D.J. Rossi, A.S. Willskey, and D.M. Spielman
Object Shape Estimation From Tomographic Measurements: A Performance Analysis, Signal Processing, 18, pp 63- 87, 1989.

[4] A.B. Baggeroer
Performance Bounds on Array Processing for Source Localization Using Full Wave Modeling of Signal and Noise Fields, Computational Acoustics, 2, pp 299-319, 1986.

[5] J.G.B. Beumee
Applications of Filtering Theory for the Inversion of Temporal Chemical Systems, The Journal of Chemical Physics 87:5, pp 2617-2629, 1987.

[6] B. Porat
On the Estimation of Variance for Autoregressive and Moving Average Processes, IEEE Transactions on Information Theory 32:1, pp 120-125, 1986.

[7] B. Friedlander
On the Cramer-Rao Bound for Time Delay and Doppler Estimation, IEEE Transactions on Information Theory 30, pp 575-580, 1984.

[8] B. Friedlander
On the Computation of the Cramer-Rao Bound for ARMA Parameter Estimation, IEEE Transactions on Acoustics, Speech, and Signal Processing 32:4, pp 721-727, 1984.

[9] P. Stoica and A. Nehorai
MUSIC, Maximum likelihood, and Cramer-Rao Bound, IEEE Transactions on Acoustics, Speech, and Signal Processing, 37:5, pp 720-741, May 1989.

[10] A.O. Hero and J.A. Fessler
A Recursive Algorithm for Computing CR-Type Bounds on Estimator Covariance, to appear in IEEE Transactions on Information Theory, 1994.

[11] J. D. Gorman and Alfred O. Hero
Lower bounds on parametric estimators with constraints. *IEEE transactions on Information Theory*, 36:6, pp. 1285-1301, 1991.

[12] A.P. Dempster, N.M. Laird and D.B. Rubin
Maximum Likelihood From Incomplete Data via the EM Algorithm; *Journal of the Royal Statistical Society*, Ser. B, Vol. 39, pp 1-37, 1977

[13] A.O. Hero, J.A. Fessler and W.L. Rogers
A Fast Recursive Algorithm for Computing CR-Type Bounds for Image Reconstruction Problems. *IEEE Nuclear Science Symposium and Medical Imaging Conference*, Orlando, 1993

[14] A. O. Hero
A Cramer-Rao Type Lower Bound for Essentially Unbiased parameter Estimation, Technical Report 890, Lincoln Laboratory, MIT 1992.

[15] M. Usman and A.O. Hero
The Analysis of a Single Photon Emission Computed Tomography (SPECT) System Using Mutual Information and Cramer-Rao Bound as Performance Criteria, CSPL Technical Report, The University of Michigan, 1993.

[16] A. Kuruc
Lower bounds on Multiple-source Direction Finding in the Presence of Direction-dependent Antenna-array-calibration errors, Technical Report 799, MIT Lincoln Laboratory, Lexington MA, 1989.

[17] T. N. E. Greville
Some Applications of the Pseudo-Inverse of a Matrix. *SIAM Review*, vol. 2, pp. 15-22, 1960

[18] P. A.Janakiraman and S. Renganathan
Recursive Computation of Pseudo-Inverse of Matrices. *Automatica*, vol. 18, pp. 631-633, 1982.

[19] W. Guo-rong and C. Yong-lin
A Recursive Algorithm for Computing the Weighted Moore-Penrose Inverse A^{+}_{MN}. *Journal of Computational Mathematics*, vol. 4:1, pp. 74-85, 1986

[20] R.C. K. Lee
Optimal Estimation, Identification, and Control, Research Monograph no. 28, The MIT Press, Cambridge Massachusetts 1964.

[21] R.A. Horn and C.R. Johnson
Matrix Analysis, Cambridge, 1985.

[22] G.H. Golub and C.F. van Loan
Matrix Computations (Second Edition), The Johns Hopkins University Press, 1989.

[23] L.L. Scharf
Statistical Signal Processing: Detection, Estimation, and Time Series Analysis, pp 211-222, Addison-Wesley Publishing Company, 1991.

[24] I.A. Ibragimov and R.Z. Has'minskii
Statistical Estimation: Asymptotic Theory, Springer Verlag, New York, 1981.

[25] D.G. Luenberger
Optimization by Vector Space Methods, pp 290-299, John Wiley and Sons, 1969.

[26] D. Rose and R.A. Willoughby
Sparse Matrices and their Applications, Plenum Press, 1972.

[27] A.N. Tikhonov and V.Y. Arsenin
Solutions of Ill-Posed Problems, Halsted Press, 1977.

[28] L.A. Hageman and D.M. Young
Applied Recursive Methods, pp 18-22, Academic Press, 1981.

[29] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B.P. Flannery
Numerical Recipes in FORTRAN, The Art of Scientific Computing, (Second Edition), Cambridge University Press, 1992.

[30] J. Philip
The Most Ill-posed Non-negative Kernels in Discrete Deconvolution, Inverse Problems, 3:2, pp 309-328, May 1987.

[31] D.M. Young
Iterative Solution of Large Linear Systems, Academic Press, New York, 1971.

[32] R. Barrett et. al.
Templates for the Solutions of Linear Systems Building Blocks for Iterative Methods, SIAM Press, Philadelphia, 1994.

[33] H.L. van Trees
Detection, Estimation and Modulation Theory, (Part I), John Wiley and Sons. 1968.

[34] J. A. Fessler and A. O. Hero
Cramer-Rao Bounds for Biased Estimators in Image Reconstruction, Midwest Symposium on Circuits and Systems, pp 253-256, August 1993, Detroit, MI.

[35] J. A. Fessler
Penalized Weighted Least-Squares Image Reconstruction for Positron Emission Tomography, IEEE Transactions on Medical Imaging, 13:2, June 1994, In Press.

[36] J.A. Fessler and A.O. Hero
Space-Alternating Generalized Expectation-Maximization Algorithm, To appear in IEEE Transactions on Signal Processing, Oct., 1994.

[37] S.M. Kay
Modern Spectral Estimation: Theory and Applications, Prentice Hall Signal Processing Series, 1988.

[38] E. L. Lehmann
Testing Statistical Hypothesis, 2nd Edition, pp 542-553, John Wiley and Sons, New York, 1986.

[39] R. Kakarala and A.O. Hero
On Achievable Accuracy in Edge Localization, IEEE Transactions on Pattern Analysis and Machine Intelligence, 14:7, pp. 777-781, July, 1992.

[40] W.L. Rogers, N.H. Clinthorne, L.Shao, P.Chiao, J.Stamos, and K.F.Koral
SPRINT II, A Second Generation Single Photon Ring Tomograph, IEEE Transactions on Medical Imaging, 7:4, pp. 291-297, 1988.

[41] L. Shao, A.O. Hero, W.L. Rogers and N.H. Clinthorne
The Mutual Information Criterion for SPECT Aperture Evaluation and Design, IEEE Transactions on Medical Imaging, 8:4, pp 322-336, Dec. 1989.

[42] M. Usman, A. O. Hero and W. L. Rogers
Performance Gain Analysis for Adding Vertex View to a Standard SPECT, Midwest Symposium on Circuits and Systems, pp 418-421, August 1993, Detroit, MI.

[43] M. Usman, A. O. Hero and J.A. Fessler
Bias-Variance Trade-offs Analysis Using Uniform CR Bound, IEEE Medical Imaging conference, San Francisco, Nov. 1993.

[44] M. Usman and A.O. Hero
Recursive CR bounds: Algebraic and statistical acceleration, ICASSP 94, Adelaide, Australia, Apr. 1994.

[45] M. Usman, A.O. Hero and J.A. Fessler
Bias-variance trade-offs analysis using the uniform CR bound for image reconstruction, Accepted for ICIP 94, Austin, Texas, Nov. 1994.

[46] M. Usman, A.O. Hero & J.A. Fessler
Uniform CR bound: Implementation issues and applications, Accepted for the IEEE Medical Imaging

[47] M. Usman, A.O. Hero and J.A. Fessler
Recursive Algorithms for Computing CR Bounds, To be submitted to IEEE Transactions on Signal Processing.

[48] M. Usman, A.O. Hero, and J.A. Fessler
Bias-Variance Trade-offs For Parametric Estimation Problems Using Uniform CR Bound, To be submitted to IEEE Transactions on Signal Processing.

[49] M. Usman, A.O. Hero, W.L. Rogers and J.A. Fessler
Applications of Uniform CR Bound for SPECT Performance Evaluation and System Design, To be submitted to IEEE Transactions on Medical Imaging.

[50] W.L. Rogers, N. Clinthorne, J.A. Fessler, Y. Zhang, L. Hua, C. Ng, M. Usman & A.O. Hero
Value of a vertex view for brain SPECT, Submitted to the World Federation of Nuclear Medicine and Biology Sixth World Congress, Sydney, Australia, Oct. 1994.

[51] Henry Stark and John W. Woods
Probability, Random Processes, and Estimation Theory for Engineers. Prentice-Hall, 1986.

[52] G.T. Herman
Image Reconstruction From Projections: Implementation and Applications Topics in Applied Physics, Vol. 32, pp 231-232, Springer-Verlag, 1979.