

**TRADEOFFS AND LIMITATIONS IN
STATISTICALLY BASED IMAGE
RECONSTRUCTION PROBLEMS**

by

Thomas J. Kragh

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in The University of Michigan
2002

Doctoral Committee:

Professor Alfred O. Hero, III, Chair
Emeritus Professor W. Leslie Rogers
Associate Professor Jeffery A. Fessler
Neal H. Clinthorne, Senior Associate Research Scientist

ABSTRACT

TRADEOFFS AND LIMITATIONS IN STATISTICALLY BASED IMAGE RECONSTRUCTION PROBLEMS

by
Thomas J. Kragh

Chair: Alfred O. Hero, III

Advanced nuclear medical imaging systems collect multiple attributes of a large number of photon events, resulting in extremely large datasets which present challenges to image reconstruction and assessment. This dissertation addresses several of these challenges. The image formation process in nuclear medical imaging can be posed as a parametric estimation problem where the image pixels are the parameters of interest. Since nuclear medical imaging applications are often ill-posed inverse problems, unbiased estimators result in very noisy, high-variance images. Typically, smoothness constraints and a priori information are used to reduce variance in medical imaging applications at the cost of biasing the estimator. For such problems, there exists an inherent tradeoff between the recovered spatial resolution of an estimator, overall bias, and its statistical variance; lower variance can only be bought at the price of decreased spatial resolution and/or increased overall bias. A goal of this dissertation is to relate these fundamental quantities in the analysis of imaging systems. A relationship between list-mode (single photon) measurements and binned measurements is shown. A model for the measurement statistics for a Compton Scatter Single Photon Emission Tomography (Compton SPECT) system is derived, and reconstructed images from both simulated and measured data are presented. In order

to reduce the computations involved in reconstruction, we explore lossy compression of the projection data using vector quantizers. Asymptotic expressions for the loss in the Kullback-Liebler divergence due to quantization for a low contrast lesion detection task are derived. A fast and efficient method of optimizing the measurement space partitioning using a lattice vector quantizer is presented, and results for a 2D Positron Emission Tomography (PET) imaging system show an optimal bit allocation. Finally, some fundamental limitations in image reconstruction are derived. In particular, the tradeoff between bias, resolution, and estimator variance is explored using a 2D image deconvolution problem as a motivating example.

© Thomas J. Kragh 2002
All Rights Reserved

To Patty.

ACKNOWLEDGEMENTS

The acknowledgments have to be the most difficult for me to write, mainly due to the uncountable number of people who have helped and guided me through these past five years.

I would like to start by extending my sincere gratitude to my advisor, Professor Al Hero, for his guidance, support and patience as I struggled with this material. Secondly, I would like to express my appreciation to my committee members, Professor Rogers, Professor Fessler, and Neal Clinthorne, whose theoretical and practical knowledge of medical imaging systems have been an invaluable resource. Furthermore, I would like to thank the EECS faculty, whose courses and lectures shaped my academic development. I would not be here today writing this dissertation if it was not for their patience and commitment to teaching, which gave me a solid foundation during my first few years of graduate school after a 5-year hiatus from academia.

I would also like to acknowledge the many close and dear friends I had made in Ann Arbor, who will always remain close to my heart. In particular, I would like to thank Kevin and Shauna, Joe, Dr. Bill and Jenny, Geof and Royann, Jason and Kari, Matt, Brad and Katharina, who have all helped me maintain balance by showing me that there is a side of life outside of school.

To my parents, Jim and Evalyn, for believing in me. Sometimes the surprise child can keep on surprising you. To my brothers, Jon and Rob, for pushing me to excel.

And finally, and most importantly, I would like to express my gratitude to my wife, Patty, whose love and support made this possible.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF TABLES	vii
LIST OF FIGURES	viii
LIST OF APPENDICES	xii
CHAPTERS	
1 Introduction	1
1.1 Dissertation Overview	1
1.2 Emission Tomography	2
1.3 Compton SPECT and Electronic Collimation	5
1.4 Image Reconstruction	7
1.5 Emission Tomography using Compressed Data	8
1.6 Resolution / Variance Tradeoffs	10
2 Statistical Model of Emission Tomography	13
2.1 Mathematical Source Model	13
2.2 Source Parameterization	15
2.3 Measurement Statistics	17
2.4 Continuous Measurement Likelihood Function	17
2.5 Continuous Measurement Fisher Information	18
2.6 Discrete Measurement Likelihood Function	19
2.7 Discrete Measurement Fisher Information	22
2.8 Source Intensity Estimator	22
3 Compton SPECT	25
3.1 Photon Transport and Measurement Statistics	27
3.2 Compton Scattering and Doppler Broadening	28
3.3 Energy Resolution	33
3.4 Position Resolution	35
3.5 Transmission and Attenuation	35

3.6	Detection	38
3.7	Simplifications	44
3.8	Sensitivity and Measurement Probability	48
3.9	Reconstructed Images from Simulated 3D Data	51
3.10	Reconstructed Images from Measured 2D Data.	52
3.11	Discussion	53
4	Simplifying Tomographic Reconstruction with Optimized Binning . .	64
4.1	Vector Quantization Background	70
4.2	MSE-Optimal Vector Quantizer	71
4.2.1	Normalized Moment of Inertia	72
4.2.2	Covariation Profile	73
4.2.3	Asymptotic MSE	75
4.3	A-optimal Vector Quantizer	76
4.3.1	Loss in the Fisher Information trace	77
4.3.2	Asymptotic loss in the Fisher Information trace . . .	78
4.3.3	Optimal point density	79
4.3.4	Loss in Fisher information trace for 1D example . . .	79
4.4	Minimum Discrimination Loss Vector Quantizer	81
4.4.1	Binary Hypothesis Testing	82
4.4.2	Neyman-Pearson Test and ROC Curve	83
4.4.3	Asymptotic Performance of Hypothesis Test	84
4.4.4	Small Signal Perturbation Analysis	85
4.4.5	Discrimination Loss	87
4.4.6	Asymptotic Discrimination Loss	89
4.4.7	Optimal point density	91
4.5	Quantizer Design	91
4.6	Lattice Quantizer	93
4.6.1	Discrimination Loss for Lattice Quantizer	94
4.6.2	Discrimination Loss Calculation for 1D example . . .	94
4.6.3	Discrimination Loss Calculation for 2D PET	96
4.6.4	Lattice Quantizer with Support Constraints	104
5	Uniform Cramèr-Rao Bound for Image Reconstruction	108
5.1	Cramèr-Rao Bound	110
5.2	Bias-gradient	112
5.3	Local Impulse Response and Mean-gradient	113
5.4	Constrained CR-Bound	115
5.4.1	Optimal bias-gradient	116
5.4.2	Critical 2nd-moment and minimum-variance surface .	118
5.5	UCRB for Image Deconvolution	119
5.6	Penalized weighted least-squares estimator	121
5.6.1	Matched PWLS estimator	123
5.6.2	Mis-matched PWLS estimator	124

6	Conclusions	128
	APPENDICES	132
	BIBLIOGRAPHY	135

LIST OF TABLES

Table

3.1	Sources of error present in each of the simulated data reconstruction cases, along with the resulting FWHM of the center point source. . .	52
4.1	Numerical values for the trace of the Fisher information and inverse Fisher information matrices, for three different quantizer point densities, for an $M = 150$ level quantizer.	80

LIST OF FIGURES

Figure	
1.1 Components of an Anger camera.	3
1.2 Diagram of an Anger camera with a converging fan-beam collimator.	4
1.3 Region of possible positron emission / annihilation location for a coincidence measurement in PET (left), TOF-PET (right).	6
2.1 Diagram of scanner field of view along with emission location.	13
3.1 Compton SPECT detection process.	26
3.2 Conical ambiguity of the γ -ray emission source location.	27
3.3 Geometry of a Compton scatter event.	29
3.4 Klein-Nishina Compton scattering angle distribution $p(\theta)$ (left) and associated energy distribution $p(e_1)$ (right) for an incident γ -ray energy of 140.4keV.	31
3.5 Compton scattering joint energy-angle distribution $p(e_1, \theta)$ of a 140.4keV γ -ray interacting with Silicon (top-left), and associated marginal distributions $p(e_1)$ (top-right), $p(\theta)$ (bottom-left), along with conditional scattering angle distributions for a nominal deposited energy of 17keV (bottom-right). The Klein-Nishina marginal distributions are overlaid in comparison.	33
3.6 Attenuation geometry through an arbitrary object.	37
3.7 Geometry of coincidence measurements in Compton SPECT.	40
3.8 Angle of incidence geometry relative to detector face.	47
3.9 Representative Compton backprojection onto a plane (gray) with the ideal conic section superimposed (white).	50
3.10 Reconstruction from simulated 3D data: No Doppler broadening, no collimator penetration.	55
3.11 Cross-section through reconstruction from simulated 3D data: No Doppler broadening, no collimator penetration.	55
3.12 Reconstruction from simulated 3D data: Doppler broadening, no collimator penetration.	56
3.13 Cross-section through reconstruction from simulated 3D data: Doppler broadening, no collimator penetration.	56

3.14	Reconstruction from simulated 3D data: Doppler broadening and collimator penetration.	57
3.15	Cross-section through reconstruction from simulated 3D data: Doppler broadening and collimator penetration.	57
3.16	Reconstruction from simulated 3D data: Doppler broadening, collimator penetration, and 2keV FWHM first-detector energy measurement error.	58
3.17	Cross-section through reconstruction from simulated 3D data: Doppler broadening, collimator penetration, and 2keV FWHM first-detector energy measurement error.	58
3.18	Reconstruction from simulated 3D data: Doppler broadening, collimator penetration, 2keV FWHM first-detector energy measurement error, and 1mm-offset in tomographic center of rotation.	59
3.19	Cross-section through reconstruction from simulated 3D data: Doppler broadening, collimator penetration, 2keV FWHM first-detector energy measurement error, and 1mm-offset in tomographic center of rotation.	59
3.20	Reconstruction from simulated 3D data: Doppler broadening, collimator penetration, 2keV FWHM first-detector energy measurement error, and 1mm vertical offset in collimator slot position.	60
3.21	Cross-section through reconstruction from simulated 3D data: Doppler broadening, collimator penetration, 2keV FWHM first-detector energy measurement error, and 1mm vertical offset in collimator slot position.	60
3.22	Reconstruction from measured 2D data of a point source.	61
3.23	Cross-section through reconstruction from measured 2D data of a point source.	61
3.24	Reconstruction from measured 2D data of a disk-shaped extended source.	62
3.25	Cross-section through reconstruction from measured 2D data of a disk-shaped extended source.	62
3.26	Reconstruction from combined measured 2D data of a point source and disk-shaped extended source.	63
3.27	Cross-section through reconstruction from combined measured 2D data of a point source and disk-shaped extended source.	63
4.1	Reconstructed image MSE vs. number of quantizer cells for simulated 2D PET data. The limiting case for 10^6 cells corresponds to list-mode reconstruction.	67
4.2	Block diagram of a vector quantizer.	70
4.3	Example quantizer in \mathbb{R}^1 , detailing the codevector locations and quantization regions.	71
4.4	Example quantizers in \mathbb{R}^2 . The quantizers on the left and middle are lattice quantizers and have a uniform point density, while the quantizer on the right has a non-uniform point density.	71
4.5	Example quantizer cells in \mathbb{R}^2 , along with their associated NMI.	72

4.6	Comparison of optimal point densities and bin location for a 32-bin, MSE-optimal and tr-FIM-optimal 1D quantizer.	80
4.7	Null hypothesis and alternative hypothesis for 1D example.	87
4.8	True and approximate discrimination functions for 1D example.	87
4.9	True and approximate ROC curve for 1D example.	88
4.10	Example 2D lattice with rectangular cells (left) and hexagonal cells (right), with codevector location marked with “x” at centroid of each cell.	94
4.11	Discrete measurement discrimination $\mathcal{D}(P_1 P_0)$ versus continuous measurement discrimination $\mathcal{D}(p_1 p_0)$ for a 1D uniform scalar quantizer.	95
4.12	Actual loss in discrimination $\Delta\mathcal{D}(p_1 p_0)$ due to quantized measurements versus asymptotic limit for uniform scalar quantizer.	96
4.13	Loss in ROC curve due to quantization for 1D example.	97
4.14	Simulated 2D PET emission source (left), emission source with 3cm circular lesion (right)	98
4.15	Line-of-Response for 2D PET with continuous (right) and Discrete (left) measurements.	99
4.16	PET Sinogram with $M = 5640$ samples	100
4.17	PET Sinogram with $M = 15700$ samples	100
4.18	Forward Projection of Lesion with $M = 15700$ samples	101
4.19	Discrimination versus number of sinogram cells	102
4.20	Discrimination Loss versus number of sinogram cells	103
4.21	Non-adaptive lattice quantization of 2D PET sinogram, $M = 1450$ bins.	104
4.22	Adaptive lattice quantization of 2D PET sinogram, $M = 984$ bins.	105
4.23	Adaptive lattice quantization of 2D PET sinogram, $M = 1450$ bins.	106
5.1	Tradeoff between reconstructed image noise and bias.	108
5.2	Example of three representative point response functions with identical bias-gradient norm and central lobe FWHM, but with varying sidelobe behavior.	115
5.3	Graphical interpretation of the constrained minimization problem involved in calculating the UCRB.	116
5.4	Original Image (left), Noisy/Blurred Image (right).	119
5.5	Bounds on single pixel estimation performance in presence of blur and additive gaussian noise.	121
5.6	Mean-gradients images for two different estimators with different bias-gradient norm. $\delta = 0.1$ (left), $\delta = 0.5$ (right).	122
5.7	Mean-gradient cross-sections for two different estimators with fixed bias-gradient norm $\delta = 0.5$	123
5.8	Minimum-variance surface, critical 2nd-moment, and active constraint regions.	125

5.9	PWLS estimator trajectories for both an identity and smoothing penalty. Both the measurements and estimator are matched with a 1.5pixel FWHM blur.	125
5.10	PWLS estimator trajectory for a smoothing penalty. The measurements had a 1.5pixel FWHM blur, while the estimator assumed a 1.75pixel FWHM blur.	126
5.11	PWLS estimator trajectory for an identity penalty. The measurements had a 1.5pixel FWHM blur, while the estimator assumed a 1.75pixel FWHM blur.	126
5.12	PWLS estimator trajectory for both an identity and smoothing penalty. The measurements had a 1.5pixel FWHM blur, while the estimator assumed a 1.75pixel FWHM blur.	127
5.13	PWLS estimator trajectory for both an identity and smoothing penalty. The measurements had a 1.75pixel FWHM blur, while the estimator assumed a 1.5pixel FWHM blur.	127

LIST OF APPENDICES

APPENDIX

A	List-mode ML-EM Estimator	133
---	-------------------------------------	-----

CHAPTER 1

Introduction

Advanced nuclear medical imaging systems, such as Compton SPECT, collect multiple attributes of a large number of photon events which result in extremely large high-dimensionality datasets. These in turn present challenges to image reconstruction and assessment. This dissertation introduces some novel approaches toward the goal of deriving tractable image reconstruction algorithms for advanced emission tomography systems with high-dimensional measurement spaces.

1.1 Dissertation Overview

We will start with a basic model of the measurement statistics for a generalized photon imaging system. The relationship between list-mode (i.e. single-photon) measurements and binned measurements will be shown, as well as their respective log-likelihood functions with respect to the observed measurements. A general method for parameterizing the source distribution will be given, along with a Maximum-Likelihood estimator of the source distribution. Next, we will focus on the specific problem of modeling the measurement statistics for a Compton SPECT camera and will show list-mode Maximum-Likelihood image reconstruction results for both measured 2D and simulated 3D Compton SPECT data. As a possible solution to the high-count list-mode reconstruction problem, we will investigate methods of adaptively binning the measurements in a manner driven by the measurements them-

selves. In particular, we will investigate the properties and applications of image reconstruction with vector quantized measurements, and develop asymptotic expressions for the performance loss due to quantization as compared to using unquantized measurements. Lastly, we will investigate fundamental limitations in image reconstruction. In particular, lower bounds of estimator variance will be derived for the general biased-estimator case. A fundamental tradeoff between bias, resolution, and estimator variance will be derived. An example 2D image deconvolution problem will be examined from the standpoint of estimating a single pixel value.

1.2 Emission Tomography

Nuclear medical imaging techniques such as Single Photon Emission Computed Tomography (SPECT) or Positron Emission Tomography (PET) are non-invasive methods of obtaining functional information of internal body processes. An example of functional imaging would be glucose uptake within the brain, which is a common method of determining relative brain activity. By labeling glucose with a radioisotope and injecting it into a patient, one can infer the internal metabolism by measuring the relative distribution of radiation emanating from the patient. However, one is not limited to simply measuring glucose uptake. There exists a wide-range of radioisotopes, as well as the compounds formed with them, that can be used to monitor internal body processes, detect tumors, etc. In comparison, imaging techniques such as Ultrasound or X-ray Computer Assisted Tomography (X-ray CT) [13] give structural and anatomical information, but do not give any information about the chemical processes occurring within the body.

Nuclear medical imaging relies upon detecting and measuring γ -ray intensity, where the γ -rays are emitted by the radioactive decay of a radioisotope injected within a patient's body. The typical SPECT measurement device is the Anger camera [9, 17, 67], which consists of a scintillating material such as sodium iodide (NaI)

coupled with an array of photomultiplier tubes. When an incoming γ -ray strikes the scintillating material, it creates visible light photons which are detected and amplified by an array of photomultiplier (PMT) tubes. By measuring the relative output intensity of the PMT tube array, one can estimate the position where the individual γ -ray struck the scintillation material. The components of a typical Anger camera are given in Figure 1.1.

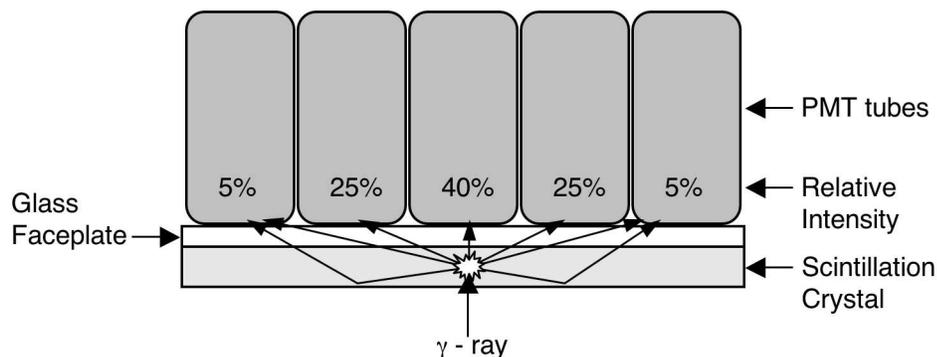


Figure 1.1: Components of an Anger camera.

In the above diagram, an incoming γ -ray can come from any direction and be detected. In order to obtain directional information as to where the γ -ray came from, a pin-hole or slit collimator is placed in front of the Anger camera. The collimator typically consists of an array of pin-holes or slits in a high atomic number material such as lead, tungsten, or gold. Only γ -rays traveling parallel to the pin-holes or slit array will pass through the collimator and be detected by the Anger camera.

In SPECT, a radioisotope is administered to the patient, which subsequently decays and results in the production of γ -rays that travel outside the patient's body and are detected externally by a collimated Anger camera. The Anger camera is rotated around the patient in some fashion to collect intensity measurements at various angles. By collecting the radiation intensity at various angles around the patient, the internal radiation distribution can then be estimated by solving an inverse problem that will be described in more detail in Chapter 2.

The main problem with SPECT imaging is its reliance upon mechanical collimation to determine the direction of the detected γ -ray. A mechanical collimator, a lead pinhole or slit, forms a hard physical constraint on where a detected photon could have originated. The incoming photon is constrained to have emanated from somewhere along a fan-beam passing through the patient, as illustrated in Figure 1.2.

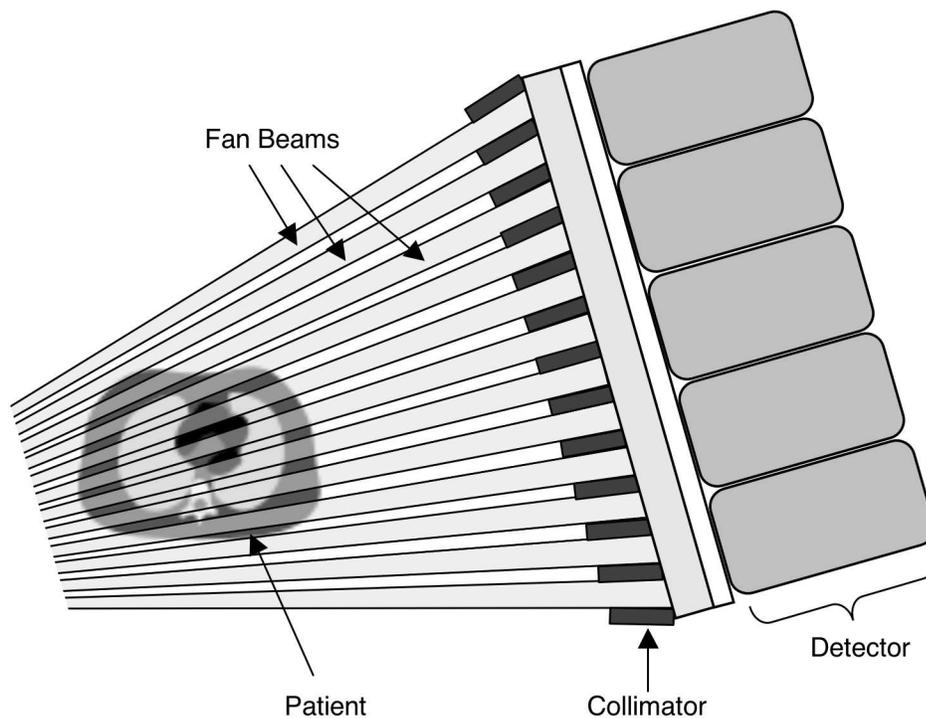


Figure 1.2: Diagram of an Anger camera with a converging fan-beam collimator.

Due to collimation, SPECT systems detect only a fraction of the emitted γ -rays, on order of 10^{-4} or less. Since the total radiation dose that can be administered is limited to avoid harming the patient, the only alternative for detecting more γ -rays is to increase the size of the collimator holes. However, there is an inverse relationship between collimator pinhole size and resolution. For a fixed geometry and fixed imaging time (equivalently, fixed radiation dose), one is forced to trade off resolution for signal intensity.

1.3 Compton SPECT and Electronic Collimation

In Compton SPECT imaging [71, 72, 95, 96, 101], the location uncertainty of the emitted photon is reduced by electronic collimation where additional measurements about the detected photon are used to infer where it came from rather than relying upon physical constraints such as pinholes or slits. These additional measurements about the detected photon event are obtained by *coincidence measurement*, which is defined as when multiple measurements of a γ -ray's interaction with a detection system occurs within some small time window so as to be considered occurring instantaneously. Since γ -ray photons travel at the speed of light, a set of measurements from two different detectors that occurs within some small time window (determined by the imaging volume viewed by the camera, and typically on order of nanoseconds) can be presumed to be different attribute measurements of the same γ -ray photon.

Compton SPECT is not the only imaging modality that relies on electronic collimation. In Positron Emission Tomography (PET) [13, 104], a radioisotope injected inside the patient emits a positron which in turn self-annihilates with a nearby electron. This produces a pair of 511keV γ -ray photons. Pairs of γ -ray photons then exit the body in opposite directions where they are subsequently detected by an external ring of detectors surrounding the patient. These pairs of detections form coincidence measurements; in that pairs of detections which occur within some small time-window are presumed to have come from the same positron annihilation. The positron emission / annihilation location is constrained to have occurred somewhere along a *line of response* between the two detections on the PET ring. In Time-of-Fight (TOF) PET [84, 98], the difference in the arrival times between the two detections of each coincidence measurement is also recorded. Since the speed of light is 30cm per nanosecond, a differential timing measurement accuracy on the order of 100ps (100×10^{-12} sec) can be used to reduce the uncertainty of the emission location to a region a few centimeters wide within the line-of-response, as shown in Figure 1.3.

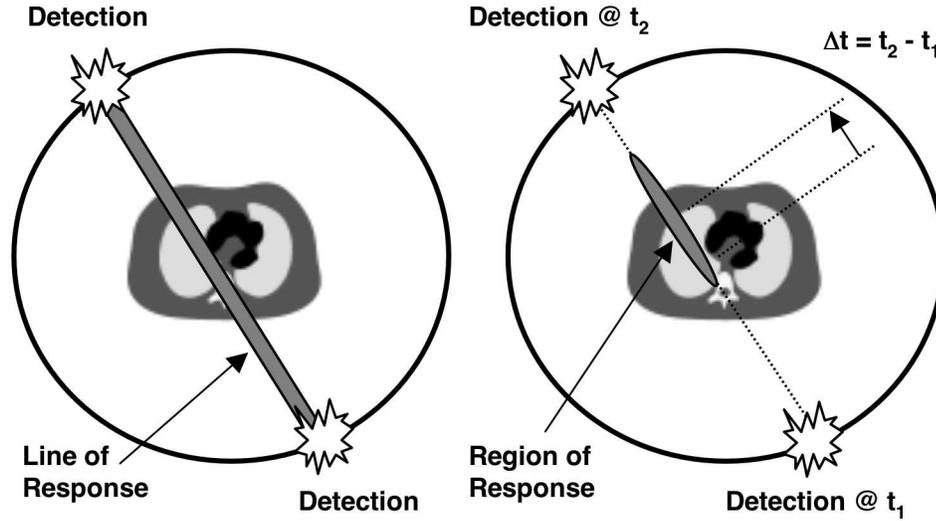


Figure 1.3: Region of possible positron emission / annihilation location for a coincidence measurement in PET (left), TOF-PET (right).

Currently, over 85% of all nuclear medicine procedures use some type of ^{99}Tc -labeled radio-pharmaceutical [90]. In order to be competitive as an imaging technique, a Compton SPECT camera must be able to exceed the performance of existing mechanically collimated SPECT systems. The advantage of applying electronic collimation to SPECT imaging is twofold: greatly increased photon collection efficiency and the decoupling of resolution vs. signal intensity. The disadvantage of Compton SPECT imaging is that the region of uncertainty of each detected photon is constrained to lie on the surface of a “fuzzy cone” rather than along a line or fan-beam. Intuitively, this conical region of uncertainty is larger than the region of uncertainty along a fan-beam. Thus, each measured γ -ray in a Compton SPECT camera gives less information as to the position of the emission source than traditional SPECT systems. Clinthorne et al [23, 61] have shown that on a per detected photon basis, a Compton SPECT camera has worse performance than standard SPECT imaging at ^{99}Tc energies. However, due to its vastly larger collection efficiency, a Compton SPECT camera is potentially superior to SPECT for a given imaging time because it

can collect far more of the emitted γ -rays due to its higher collection efficiency. Also, a Compton SPECT camera shows better performance than SPECT at higher γ -ray energies, since imaging at higher energies ($> 140\text{keV}$) requires thicker collimator septa which in turn results in decreased measurement resolution and detection sensitivity.

1.4 Image Reconstruction

Image reconstruction in nuclear medical imaging can be posed as a parametric estimation problem, where the internal radiation distribution inside the patient is discretized into N distinct pixels whose mean intensities are the parameters of interest. Traditional statistically-based image reconstruction methods rely on “binning” or quantizing the individual γ -ray detections into one of M discrete locations. Binning is a discretization method which has an associated discretization error due to quantizing the measurement space. Statistically-based estimation techniques are typically iterative, with a computational load of $O(MN)$ per iteration for binned measurements. For example, a PET detector ring with B detector elements will have at most $M = B(B-1)$ uniquely ordered detector element pairs. The corresponding angle and position measurements are quantized by the detection hardware, with the quantized values determined by the relative acceptance angles and locations of the finite-size detector element pairs in the PET ring.

In Compton SPECT, the image reconstruction problem is daunting due to the high dimensionality of each measurement. Binning of each measurement dimension individually results in many more possible bins than measurements. When fewer photons are collected than the possible number of measurements, an alternative is *list-mode* reconstruction techniques [10, 84]. In list-mode reconstruction, multiple attributes of each photon such as position, energy, scatter angle, etc., are recorded as a point in a multi-dimensional detector space. Each photon contributes individually to the likelihood function, and the underlying source distribution can be estimated using

maximum-likelihood based estimation techniques. A failing of list-mode reconstruction is that it scales directly with counts. Since each count independently contributes to the overall likelihood function, the computational complexity and storage requirements are proportional to the total number of counts. For PET or mechanically collimated SPECT systems, this is not a problem since the individual backprojections are relatively sparse. However, for imaging modalities which have non-sparse backprojections (such as Compton SPECT), this is critical since each count must be assigned a different likelihood value. For a typical Compton SPECT imaging problem involving $O(10^6)$ detections and $O(10^4)$ reconstruction pixels, the storage requirements for list-mode image reconstruction are in the gigabyte range, and it only gets worse as high efficiency fully 3D modalities which collect massive amounts of counts are developed. Thus, we are faced with the choice of either keeping all the measurements in their full resolution and accepting the storage requirements for the backprojections, or binning the data and accepting the loss in both measurement and reconstruction fidelity due to quantization error.

1.5 Emission Tomography using Compressed Data

The next part of this dissertation concerns acceleration of image reconstruction for emission tomography where equally spaced measurement bins are impractical due to the high dimensional detector space. Traditionally in imaging modalities such as 2D SPECT and PET, the detected photon measurements are quantized to one of a finite number of possible measurement bins. This quantization is either done mechanically by the geometry of the detectors (i.e., PET or SPECT detector elements) or electronically in software (i.e., Anger camera energy windows or position measurement binning). Typically, the measurements are binned uniformly with equal-sized bins along each measurement dimension. As pointed out in [10], when the dimensionality of the measurement space is greater than 4 or more, uniform quantization about each

measurement axis becomes impractical. One wants fine binning along each axis in order to minimize error due to quantization. However, the total number of quantizer bins grows exponentially with the number of dimensions.

A concept well known in the fields of information theory and communications is vector quantization [49, 52], which is a multivariate generalization of binning. For a k -dimensional measurement space, rather than distributing $O(M^{1/k})$ bins uniformly along each dimension, a vector quantizer places more bins in regions of the detector space that are more informative. A quantizer that has a non-uniform concentration of quantizer bins is referred to having a variable point density. Although vector quantizers with non-uniform point densities can achieve minimal distortion loss, they are difficult to design and implement in practice especially at the high dimension-rate product sizes required for image reconstruction.

For a fixed number of quantizer cells¹ about each axis, the total number of quantizer cells grows exponentially with the number of dimensions. Alternatively, for a fixed number M of quantizer cells, the average number of cells per measurement dimension is $O(M^{1/k})$ resulting in poorer measurement fidelity (or quantization error) with higher dimensional measurements when all other factors are kept fixed. However, there is merit to quantizing the components of a k -dimensional measurement independently along each of the k -axis due to the implementation simplicity of such so-called lattice-based vector quantizers [25, 92]. One possibility would be to optimize the number of quantizer levels along each axis with respect to a distortion measure under the constraint that the total number of quantizer levels (or rate of the quantizer) remains constant. One can constrain the support region of the lattice quantizer based on the observed measurements by only allocating quantizer cells to regions that have non-zero measurements. The effective volume of the support region can be estimated from the observed measurement entropy by the Asymptotic Equipartition

¹For the purposes of this dissertation we will use the terms “bins” and “cells” interchangeably.

Principle (AEP) [28]. From this one can estimate the savings in bit-rate or distortion by adapting the lattice to the observed measurement support region.

The most common distortion measure in the literature is the mean-squared error between the unquantized and quantized vectors. However, others exist which are more apropos to nuclear medical imaging. For example, in nuclear medical imaging the difference between a radiotracer image of a patient before and after the emergence of a tumor would manifest itself as a subtle variation of a few pixels against an similar background image. Thus, one is concerned not with dramatic changes between two images, but rather with sensitivity to detecting small changes in an initial image. In this dissertation, we investigate a strategy pioneered by Gupta [54] of designing optimal vector quantizers that minimize the loss in various metrics such as the Kullback-Liebler (KL) discrimination [28], along with the associated loss in performance as measured by the area under the Receiver Operating Characteristic (ROC) curve. The KL discrimination and ROC curve characterize the detection performance of the quantizer / detector combination, and are analogs to the mean-square quantization error metric commonly used to design vector quantizers. These KL-optimal quantizers depend on the source only through its projections, and can be quickly designed on-the-fly from the observed list-mode measurements.

1.6 Resolution / Variance Tradeoffs

The variance of an image, or alternatively the signal-to-noise ratio, is often used as a measure of image quality. Since nuclear medical imaging applications are often ill-posed inverse problems, unbiased estimators result in very noisy, high-variance images. Typically, smoothness constraints and a priori information are used to reduce variance in medical imaging applications at the cost of biasing the estimator. The next part of the dissertation deals with quantifying fundamental bias and variance tradeoffs for statistically based image reconstruction.

The classical Cramèr-Rao (CR) bound [28, 103] gives an estimator-independent lower bound on variance and can be used to find fundamental limitations of an imaging system. However, it only applies to unbiased estimators. The general form of the CR bound applies to biased estimators. However, it is not estimator-independent since it depends on the gradient of the estimator’s bias function. In order to measure the performance capability of an imaging system, a desired goal is to decouple the variance due to the inherent limitations of the imaging system from the choice of algorithm used to process the data.

The Uniform Cramèr-Rao bound (UCRB) [57] is a generalization of the CR bound that decouples the variance contributions of the biased estimator and the imaging system. The idea behind the UCRB is that the norm of the bias-gradient vector can be viewed as a measure of the total bias-error of an estimator. Among all possible estimators with a given bias-gradient norm, one can solve for the minimum possible variance of a hypothetical ideal estimator via a constrained optimization problem. The variance of this hypothetical ideal estimator is then a lower bound among all possible estimators with that choice of bias-gradient norm. Moreover, the variance of this hypothetical ideal estimator is only a function of the Fisher Information and the bias-gradient norm, and thus independent of any particular choice of estimation algorithm. For any given choice of estimator, the tradeoff between bias and variance can be expressed as a parametric curve indexed by the estimator’s regularization parameter. The variance of the hypothetical ideal estimator calculated using the UCRB forms a lower bound on the variance of this (or any other) choice of estimator. The margin between these two curves gives the performance margin between the particular regularized estimator in question and the fundamental performance limit of the imaging system². For an estimator whose variance lies on this curve, lower estimator variance can only be achieved at the price of increased estimator bias.

²In general, the Fisher Information and the resulting lower bound will be object dependent.

For single pixel intensity estimation, the estimator bias-gradient is equivalent to the error between the estimator's local impulse response and an ideal impulse response. Thus, the norm of the estimator bias-gradient would seem to be a natural measure of resolution in that local impulse response functions with small tails would have smaller bias-gradient norm. However, different impulse response functions can have identical bias-gradient norm but widely different resolution. This has led to some counter-intuitive results and interpretation difficulties when using the Uniform CR Bound in performance studies of imaging systems [81].

A major contribution made in this dissertation is the extension of this tradeoff concept to include resolution [68,69]. One can solve for an ideal hypothetical estimator that has minimum variance for a given amount of total bias and resolution, forming a parametric surface which is a function of the Fisher information, bias-gradient norm, and resolution parameter and thus independent of any particular choice of estimation algorithm. The variance of any given estimator traces out a trajectory, indexed by its regularization parameter, that is constrained to lie above this minimum-variance surface. The distance between this trajectory and surface gives the performance margin between the particular regularized estimator in question and the fundamental performance limit of the imaging system. For an estimator whose variance lies on this surface, lower estimator variance can only be achieved at the price of increased estimator bias and/or decreased resolution.

CHAPTER 2

Statistical Model of Emission Tomography

2.1 Mathematical Source Model

Let $\underline{x} \in \mathbb{R}^d$ be a random variable with distribution $p(\underline{x})$, and let \underline{X} be a realization of \underline{x} . It is common to model source emission locations as a time-ordered sequence $\{\underline{X}_1, \underline{X}_2, \dots\}$ of statistically independent random spatial locations (each in \mathbb{R}^d) drawn from a spatial Poisson point process with emission-rate function $\lambda(\underline{x})$ defined over some subset $\Omega \subset \mathbb{R}^d$, where $d = 2$ for planar imaging, $d = 3$ for volume imaging and $d = 4$ for dynamic volume imaging where the fourth dimension is time. The subset Ω is called the field-of-view of the scanner, the region of space over which emissions are recorded.

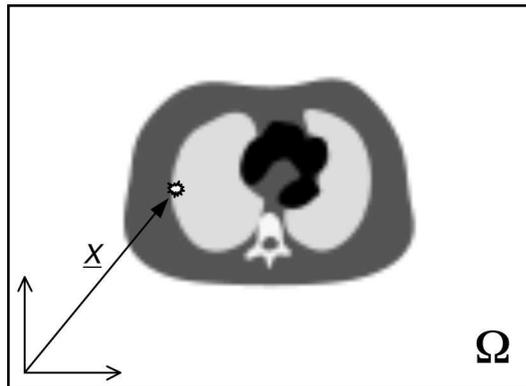


Figure 2.1: Diagram of scanner field of view along with emission location.

The emission-rate function $\lambda(\underline{x})$ has units of emissions per unit time per unit area

or volume, and the total emission-rate $\lambda_o = \int_{\Omega} \lambda(\underline{x}') d\underline{x}'$ has units of emissions per-unit-time. The emission rate function is proportional to the distribution $p(\underline{x})$, i.e. $\lambda(\underline{x}) = \lambda_o p(\underline{x})$.

Let the vector $\underline{y} \in \mathbb{R}^k$ be a measurement of k different attributes of a detected photon. Let $S \subset \mathbb{R}^k$ be the set of all possible detected measurements. Assuming a linear relationship between the measurements and emissions, the sequence $\{\underline{Y}_1, \underline{Y}_2, \dots\}$ of detected measurements is a set of random vectors (each in \mathbb{R}^k) drawn from a Poisson point process with detection-rate function $\mu(\underline{y})$ given by

$$\mu(\underline{y}) = \int_{\Omega} a(\underline{y}|\underline{x}') \lambda(\underline{x}') d\underline{x}'. \quad (2.1)$$

Equation (2.1) is referred to as the ideal measurement equation where $\mu(\underline{y})$ is the ideal (noiseless) detection-rate function and $\lambda(\underline{x})$ is the source intensity of interest.

The term $a(\underline{y}|\underline{x})$ is often referred to as the *transition probability* in that it is the probability density function of measurements \underline{y} given a detected emission from $\underline{x} \in \Omega$ combined with the probability of detecting an emission from $\underline{x} \in \Omega$ (depending on the detector efficiency),

$$a(\underline{y}|\underline{x}) = p(\underline{y}|D, \underline{x}) P(D|\underline{x}) \quad (2.2)$$

where the event D represents a detected photon emission. The term $P(D|\underline{x})$ is often referred to as the *survival probability* or *detection sensitivity* [10, 84] of a photon emission from the location \underline{x} , and will be denoted as $s(\underline{x})$.

From the definition of the transition probability function $a(\underline{y}|\underline{x})$, the integral over all measurements $\underline{y} \in S$ is the sensitivity function

$$s(\underline{x}) = \int_S a(\underline{y}|\underline{x}) d\underline{y}. \quad (2.3)$$

Similarly, the mean total detection-rate μ_D (in units of detections per unit time) is simply the integral over all detected measurements $\underline{y} \in S$ of the detection-rate

function $\mu(\underline{y})$,

$$\begin{aligned}
\mu_D &= \int_S \mu(\underline{y}) d\underline{y} \\
&= \int_S \int_{\Omega} a(\underline{y}|\underline{x}') \lambda(\underline{x}') d\underline{x}' d\underline{y} \\
&= \int_{\Omega} \lambda(\underline{x}') \left[\int_S a(\underline{y}|\underline{x}') d\underline{y} \right] d\underline{x}' \\
&= \int_{\Omega} \lambda(\underline{x}') s(\underline{x}') d\underline{x}'. \tag{2.4}
\end{aligned}$$

The ratio μ_D/λ_o is called the overall detection efficiency of the system.

The objective of image reconstruction is to estimate the source emission-rate intensity $\lambda = \lambda(\underline{x})$, $\underline{x} \in \mathbb{R}^d$ from a finite number p of observed measurements $\{\underline{Y}_1, \underline{Y}_2, \dots, \underline{Y}_p\}$. The most natural way to approach this problem is via the likelihood function $l(\lambda)$, which relates the measurements $\{\underline{Y}_1, \underline{Y}_2, \dots, \underline{Y}_p\}$ to λ through the joint density of $\{\underline{y}_1, \underline{y}_2, \dots, \underline{y}_p\}$ parameterized by λ . The function λ maximizing this likelihood function $l(\lambda)$ is called the *Maximum-Likelihood (ML)* estimator, and it enjoys many attractive properties [32, 103]. However, estimating the continuous function λ from a finite number of discrete measurements is an example of an *ill-posed* inverse problem [82]. Since λ is a continuous function on \mathbb{R}^d and the number of measurements is finite, there generally exists no unique solution even in the ideal situation where $\mu(\underline{y})$ is observed directly without additional noise due to Poisson counting statistics. To overcome this requires prior assumptions on the form of the emission-rate intensity function λ to constrain the estimation problem. Two common assumptions are *source parameterization* and *regularization*.

2.2 Source Parameterization

As pointed out in [41], it is natural to treat the image reconstruction problem as a statistical estimation problem due to the Poisson counting statistics of the measurements. Since we will want to estimate λ by maximizing a likelihood function, and estimating a continuous function from a finite number of observations is an ill-posed

problem, one solution is to parameterize λ by approximating it as a linear combination of N basis functions,

$$\lambda(\underline{x}) = \sum_{j=1}^N \lambda_j b_j(\underline{x}) \quad (2.5)$$

where $\boldsymbol{\lambda} = [\lambda_1, \dots, \lambda_N]^T$ is a coefficient vector of the basis functions $b_j(\underline{x}), j = 1, \dots, N$. In general there will be a truncation error between $\lambda(\underline{x})$ and its projection onto a finite set of basis functions parameterized by $\boldsymbol{\lambda}$. For this dissertation, we will restrict ourself to the pixel or voxel basis and assume that $\lambda(\underline{x})$ is smooth enough to be effectively constant over the support region of each basis function. The choice of best basis functions used to represent the image parameterization is open one. Some examples in the imaging literature include splines [4], gaussians [84], complex exponentials [102], and natural pixels [59]. However, the most common choice is a piece-wise constant function over the unit square or cube (i.e. the ‘‘pixel’’ or ‘‘voxel’’ basis).

The source-dependent terms in (2.1, 2.2, 2.3) can be made discrete in the source parameters by integrating against the basis functions. With $a_j(\underline{y}), \lambda_j$, and s_j defined as

$$a_j(\underline{y}) = \int_{\Omega} a(\underline{y}|\underline{x}') b_j(\underline{x}') d\underline{x}' \quad (2.6)$$

$$\lambda_j = \int_{\Omega} \lambda(\underline{x}') b_j(\underline{x}') d\underline{x}' \quad (2.7)$$

$$s_j = \int_{\Omega} s(\underline{x}') b_j(\underline{x}') d\underline{x}', \quad (2.8)$$

the finite truncated basis representation of $\mu(\underline{y})$ is

$$\mu(\underline{y}) = \sum_{j=1}^N a_j(\underline{y}) \lambda_j, \quad (2.9)$$

and the corresponding mean total detection-rate is

$$\mu_D = \sum_{j=1}^N \lambda_j s_j. \quad (2.10)$$

2.3 Measurement Statistics

In an emission tomography measurement, a random number p of time-ordered random measurement realizations $\{\underline{Y}_1, \underline{Y}_2, \dots, \underline{Y}_p\}$, $\underline{Y}_i \in \mathbb{R}^k$ are detected in some fixed, known time period $[0, T]$. This sequence of measurements is often referred to in the literature as *count-record* [97] or *list-mode* [10, 84]. This is contrasted to a *binned* measurement $Y_i = \sum_{l=1}^M I(\underline{Y}_i \in S_l)$, the (Poisson distributed) number of measurements collected in the fixed time interval $[0, T]$ within a bin or quantization cell $S_l \subset \mathbb{R}^k$, $l = 1, \dots, M$. Binned measurements will be discussed in more detail in Section 2.6.

One important consideration in emission tomography is the effect of detections from outside sources such as background radiation, cosmic rays, false coincidences, etc., otherwise known as *randoms* [70]. These are important effects in any γ -ray detection system and are typically modeled as an additive component to (2.1) independent of the emission source λ [40, 41, 45]. These effects are straightforward to include in the formulation of the likelihood function but will generally be neglected in this dissertation.

2.4 Continuous Measurement Likelihood Function

As shown in [10, 84], the joint density of the list-mode measurement vector \underline{y} and the number of measurements p (conditioned on being detected within some fixed time period $[0, T]$ from a source parameterized by λ) can be split into two components,

$$p(\underline{y}_1, \dots, \underline{y}_p, p | \lambda, D, T) = p(\underline{y}_1, \dots, \underline{y}_p | p, \lambda, D, T) P(p | \lambda, D, T)$$

where $p(\underline{y}_1, \dots, \underline{y}_p | p, \lambda, D, T)$ is the joint density of exactly p detected measurements, while $P(p | \lambda, T, D)$ is the probability of detecting p measurements in a fixed time T from a source parameterized by λ . Most authors drop this explicit conditioning on D , but we will keep it for notational clarity.

The density $p(\underline{y}_1, \dots, \underline{y}_p | p, \boldsymbol{\lambda}, D, T)$ can be further simplified by the fact that for a homogeneous¹ Poisson point process, the measurements are statistically independent of each other. As for the probability $P(p | \boldsymbol{\lambda}, T, D)$, the total number of measurements p detected in time T for a Poisson point process is itself Poisson distributed with parameter $(T\mu_D)$ where $\mu_D = \sum_{j=1}^N \lambda_j s_j$ is the total detected emission rate (in emissions per unit time) [97]. Using these results we obtain the following form for the joint density,

$$p(\underline{y}_1, \dots, \underline{y}_p, p | \boldsymbol{\lambda}, D, T) = \prod_{i=1}^p p(\underline{y}_i | \boldsymbol{\lambda}, D) \frac{e^{-(T\mu_D)} (T\mu_D)^p}{p!} \quad (2.11)$$

with the associated log-likelihood function of $\boldsymbol{\lambda}$ conditioned on a particular detected list-mode data sequence $\{\underline{Y}_1, \underline{Y}_2, \dots, \underline{Y}_p\}$ is given (after a bit of manipulation) by

$$l(\boldsymbol{\lambda}) = \sum_{i=1}^p \ln \left[T \sum_{j=1}^N a_j(\underline{Y}_i) \lambda_j \right] - T \sum_{j=1}^N \lambda_j s_j - \ln(p!). \quad (2.12)$$

2.5 Continuous Measurement Fisher Information

The Fisher Information Matrix (FIM) \mathbf{F} is defined as the expected value of the curvature of the log-likelihood function [103]

$$\mathbf{F}_{ij} = E_{\underline{y}} \left[-\frac{\partial^2}{\partial \lambda_i \partial \lambda_j} l(\boldsymbol{\lambda}) \right], \quad (2.13)$$

and quantifies the information efficiency of a measurement system [57]. The inverse of the FIM gives a lower bound on the variance of an unbiased estimate of the vector parameter $\boldsymbol{\lambda}$ from the measurements \underline{y} . The FIM will be used extensively in chapters 4 and 5.

For list-mode measurements, the expectation over all measurements \underline{y} is conditioned on those detected within fixed time T . Thus, the number of counts p is a r.v.

¹Strictly speaking this is not a homogeneous process due to radioactive decay. However for a single species we can rescale time to account for this.

which we take into account by averaging over all counts,

$$\begin{aligned}
\mathbf{F}_{ij} &= \sum_{p=0}^{\infty} P(p|\boldsymbol{\lambda}, D, T) E_{\underline{y}} \left[-\frac{\partial^2}{\partial \lambda_i \partial \lambda_j} l(\boldsymbol{\lambda}) | p, \boldsymbol{\lambda}, D, T \right] \\
&= E [p|\boldsymbol{\lambda}, D, T] \int_S p(\underline{y}|\boldsymbol{\lambda}, D, T) \frac{a_i(\underline{y})a_j(\underline{y})}{\mu(\underline{y})^2} d\underline{y} \\
&= T \int_S \frac{a_i(\underline{y})a_j(\underline{y})}{\mu(\underline{y})} d\underline{y}. \tag{2.14}
\end{aligned}$$

Here the last line comes from the fact that for a Poisson point process, the mean number of counts detected in a fixed time $E [p|\boldsymbol{\lambda}, D, T] = T\mu_D$, and the mean detected intensity rate $\mu(\underline{y}) = \mu_D p(\underline{y}|\boldsymbol{\lambda})$.

2.6 Discrete Measurement Likelihood Function

We have been assuming up to now that all measurements are continuous valued. Although it might be argued that physical processes are mostly continuous (to within quantum uncertainty), measurements will eventually be digitized by a finite precision analog-digital converter. Typically, the digitization is done further upstream in data acquisition by the hardware. For example, in a digital camera focal plane, the continuous photon position on the array is quantized to a discrete pixel row and column position while the total number of photons detected in a given pixel is measured indirectly by the total energy accumulated over a given time. Quantization can also be done in software by processing, as when the energy measurement of a single photon is acquired from an Anger camera photo-multiplier tube, resulting in an accumulation of photon counts in a *detection bin*.

We will generalize the detection bins as a set of cells $\{S_1, \dots, S_M\}$ that partition the set $S \subset \mathbb{R}^k$ of all possible detected measurements over the detector surfaces. This specifies a quantization of the measurements \underline{y} . The quantization operator can be written as

$$Q(\underline{y}) = \{ \underline{y}_i : \underline{y} \in S_i \}$$

corresponding to a quantization of the k -dimensional measurement \underline{y} into one of M disjoint cells. This will be covered in more detail in Chapter 4.

Let $\mathbf{y} = [y_1, \dots, y_M]^T$ be a $M \times 1$ random vector, where $y_i = \sum_{l=1}^p I(\underline{Y}_i \in S_l)$ is the integer number of list-mode measurement vectors detected in the fixed time interval $[0, T]$ in the quantization cell $S_l \subset \mathbb{R}^k$, $l = 1, \dots, M$. Since the continuous list-mode measurement vector \underline{Y} is a realization from a Poisson point process, the total detections in the fixed time interval $[0, T]$ in the l th quantizer cell is Poisson distributed with mean given by integral of the intensity rate function over the cell [97],

$$\begin{aligned} E[y_l] &= T \int_{S_l} \mu(\underline{y}) d\underline{y} \\ &= T\mu_l. \end{aligned} \quad (2.15)$$

The Poisson distribution gives the probability of collecting exactly m events in a fixed time T , assuming the events occur independently at a constant rate. The probability mass function of the integer number of measurements y_l detected in the cell S_l is

$$P(y_l = m) = \frac{(\mu_l T)^m e^{-\mu_l T}}{m!}. \quad (2.16)$$

The probability a_{ij} that an emission from the j th source voxel is detected and assigned to the i th quantization bin is called the transition probability. It is simply the integral of (2.6) over the i th cell S_i ,

$$a_{ij} = \int_{S_i} a_j(\underline{y}) d\underline{y}. \quad (2.17)$$

The sensitivity s_j of detecting a measurement given an emission from the j th source voxel is given by

$$\begin{aligned} s_j &= \int_S a_j(\underline{y}) d\underline{y} \\ &= \sum_{l=1}^M \int_{S_l} a_j(\underline{y}) d\underline{y} \\ &= \sum_{l=1}^M a_{lj}. \end{aligned} \quad (2.18)$$

Finally, the mean detection rate in the i th measurement bin is the integral of (2.9) over the cell S_i ,

$$\begin{aligned}
\mu_i &= \int_{S_i} \mu(\underline{y}) d\underline{y} \\
&= \int_{S_i} \sum_{j=1}^N a_j(\underline{y}) \lambda_j d\underline{y} \\
&= \sum_{j=1}^N \left[\int_{S_i} a_j(\underline{y}) d\underline{y} \right] \lambda_j \\
&= \sum_{j=1}^N a_{ij} \lambda_j,
\end{aligned} \tag{2.19}$$

with mean total detection-rate μ_D as per (2.10).

Let $\boldsymbol{\mu} = [\mu_1, \dots, \mu_M]^T$ be an $M \times 1$ vector. Let a_{ij} be the (i, j) th entry of the $M \times N$ matrix \mathbf{A} . Then in vector-matrix notation (ignoring randoms, which would introduce an additive $\boldsymbol{\lambda}$ -independent term to $\boldsymbol{\mu}$),

$$\boldsymbol{\mu} = \mathbf{A}\boldsymbol{\lambda}$$

which is a linear systems of equations for the parameter vector $\boldsymbol{\lambda}$ given the binned measurements mean intensity vector $\boldsymbol{\mu}$. The mean intensity vector $\boldsymbol{\mu}$ is the expected number of detections in each bin per unit time, i.e. $E[\mathbf{y}] = T\boldsymbol{\mu}$, which is simply the vector form of (2.15).

Let the vector $\mathbf{Y} = [Y_1, \dots, Y_M]^T$ be a realization of the integer number of quantized detected measurements in each of the cells $\{S_1, \dots, S_M\}$. By combining (2.16) and (2.19) along with property that the number of detections in each bin is independent, we arrive at the quantized-data likelihood function [94, 104],

$$\hat{l}(\boldsymbol{\lambda}) = \sum_{l=1}^M Y_l \ln \left(T \sum_{j=1}^N a_{lj} \lambda_j \right) - T \sum_{j=1}^N a_{lj} \lambda_j - \ln(Y_l!). \tag{2.20}$$

2.7 Discrete Measurement Fisher Information

By following the same procedure as in section (2.5), the (i, j) th-element of the quantized measurement Fisher Information Matrix $\hat{\mathbf{F}}$ is given by

$$\hat{\mathbf{F}}_{ij} = T \sum_{l=1}^M \frac{a_{li}a_{lj}}{\mu_l} \quad (2.21)$$

It can be shown that $\mathbf{F} - \hat{\mathbf{F}}$ is a non-negative definite matrix. This follows directly from the data processing theorem [28]. Intuitively, $\hat{\mathbf{F}}$ is less informative than \mathbf{F} since the measurements \underline{y} can always be partitioned in order to calculate the bin values \mathbf{y} , but the total number of measurements in each bin is not sufficient to recover \underline{y} . Note that the difference between the Fisher Information matrices \mathbf{F} and $\hat{\mathbf{F}}$ can be written as a summation over the M quantizer cells,

$$[\mathbf{F} - \hat{\mathbf{F}}]_{ij} = T \sum_{l=1}^M \left[\int_{S_l} \frac{a_i(\underline{y})a_j(\underline{y})}{\mu(\underline{y})} d\underline{y} - \frac{a_{li}a_{lj}}{\mu_l} \right] \quad (2.22)$$

where $\mu(\underline{y})$ and $\boldsymbol{\mu} = [\mu_1, \dots, \mu_M]^T$ are the continuous and discrete forward projections of the parameter vector $\boldsymbol{\lambda}$. Thus, for Poisson measurements, the difference between the Fisher Information matrices \mathbf{F} and $\hat{\mathbf{F}}$ will be a function of $\boldsymbol{\lambda}$ (i.e. image dependent) as well as the quantizer partitioning $\{S_1, \dots, S_M\}$. Analyzing as well as optimizing the quantizer partitions will be discussed at length in Chapter 4.

2.8 Source Intensity Estimator

From our earlier assumption that the source intensity function $\lambda(\underline{x})$ can be modeled as a linear combination of basis functions (2.5), estimating the source intensity $\lambda(\underline{x})$ can be posed as a parametric estimation problem where the source coefficients (equivalently, the image pixels) are the parameters of interest. This can be posed as an optimization problem. For a given list-mode measurement realization $\{\underline{Y}_1, \underline{Y}_2, \dots, \underline{Y}_p\}$, the maximum-likelihood estimate of the source distribution is the

value $\hat{\boldsymbol{\lambda}}_{ML}$ that maximizes the likelihood function (2.12),

$$\hat{\boldsymbol{\lambda}}_{ML} = \arg \max_{\boldsymbol{\lambda} \geq 0} l(\boldsymbol{\lambda}). \quad (2.23)$$

Because tomographic image reconstruction is an ill-posed problem, maximizing the likelihood function often leads to noisy images, even when the true source distribution is a smooth function. In order to reduce the noise and enforce smoothness in the estimated source distribution, one can include a *penalty* term $R(\boldsymbol{\lambda})$ to the likelihood function that biases the estimate towards smooth images [40,41,100]. The result is still an optimization problem, but for a modified *objective function* $\Phi(\boldsymbol{\lambda}) = l(\boldsymbol{\lambda}) - \beta R(\boldsymbol{\lambda})$, where the likelihood $l(\boldsymbol{\lambda})$ gives a measure of data fit, while the penalty $R(\boldsymbol{\lambda})$ enforces smooth estimates and $\beta \geq 0$ is a scalar parameter that allows for tradeoff between the two,

$$\begin{aligned} \hat{\boldsymbol{\lambda}}_{MAP} &= \arg \max_{\boldsymbol{\lambda} \geq 0} \Phi(\boldsymbol{\lambda}) \\ &= \arg \max_{\boldsymbol{\lambda} \geq 0} l(\boldsymbol{\lambda}) - \beta R(\boldsymbol{\lambda}). \end{aligned} \quad (2.24)$$

When the penalty $\beta = 0$, the penalized estimator $\hat{\boldsymbol{\lambda}}_{MAP}$ reduces to the maximum-likelihood estimator $\hat{\boldsymbol{\lambda}}_{ML}$. For the purposes of this dissertation we will restrict ourselves to estimators which maximize the likelihood function (2.12). There are multiple different numerical optimization algorithms available to maximize this objective function [16,29,31,41–44]. However, for this dissertation, we will restrict our attention to the Maximum-Likelihood Expectation-Maximization (ML-EM) algorithm [30,84].

The list-mode ML-EM estimator of $\boldsymbol{\lambda}$ for a finite number p of list-mode measurements $\{\underline{Y}_1, \underline{Y}_2, \dots, \underline{Y}_p\}$ is an iterative algorithm given by

$$\hat{\lambda}_j^{k+1} = \frac{\hat{\lambda}_j^k}{T s_j} \sum_{i=1}^p \left[\frac{a_j(\underline{Y}_i)}{\sum_{l=1}^N a_l(\underline{Y}_i) \hat{\lambda}_l^k} \right] \quad (2.25)$$

where $\hat{\boldsymbol{\lambda}}^k = [\hat{\lambda}_1^k, \dots, \hat{\lambda}_N^k]^T$ is the estimate of $\boldsymbol{\lambda}$ at the k iteration. A derivation of (2.25) is given in Appendix A. Derivations for the binned-data ML-EM estimator can be

found in multiple references [70, 94, 104]. The end result is

$$\hat{\lambda}_j^{k+1} = \frac{\hat{\lambda}_j^k}{T s_j} \sum_{i=1}^M \left[\frac{a_{ij} Y_i}{\sum_{l=1}^N a_{il} \hat{\lambda}_l^k} \right]. \quad (2.26)$$

Note the similarity with the list-mode ML-EM estimator (2.25), which is equivalent to (2.26) when the number of counts in each bin $Y_i = 1$.

CHAPTER 3

Compton SPECT

Compton SPECT is a form of emission tomography which uses *electronic collimation* by means of *coincidence measurements* to reduce the location uncertainty of detected γ -ray photons, as compared to traditional SPECT cameras which rely upon mechanical collimators. The advantage of Compton SPECT compared to traditional systems is extremely high detection sensitivity. For example, up to 4x increase in sensitivity for imaging technetium (^{99}Tc , 140.5keV) and up to 20x increase for iodine (^{131}I , 364.4keV) has been estimated for a Compton SPECT camera compared to a representative SPECT camera [23, 72].

A Compton SPECT camera consists of a pair of detectors, which measure where an emitted γ -ray scatters from one detector and where it is absorbed in a second detector. This forms a coincidence measurement, which is defined as when multiple measurements of a γ -ray's interaction with a detection system occur within some small time window so as to be considered occurring instantaneously. Since γ -ray photons travel at the speed of light, a set of measurements from two different detectors that occur within some small time window (determined by the imaging volume viewed by the camera, and typically on order of nanoseconds) can be presumed to be different attribute measurements of the same γ -ray photon. Figure 3.1 illustrates this detection process.

By the Compton scatter energy-angle relationship (3.3), measurement of the en-

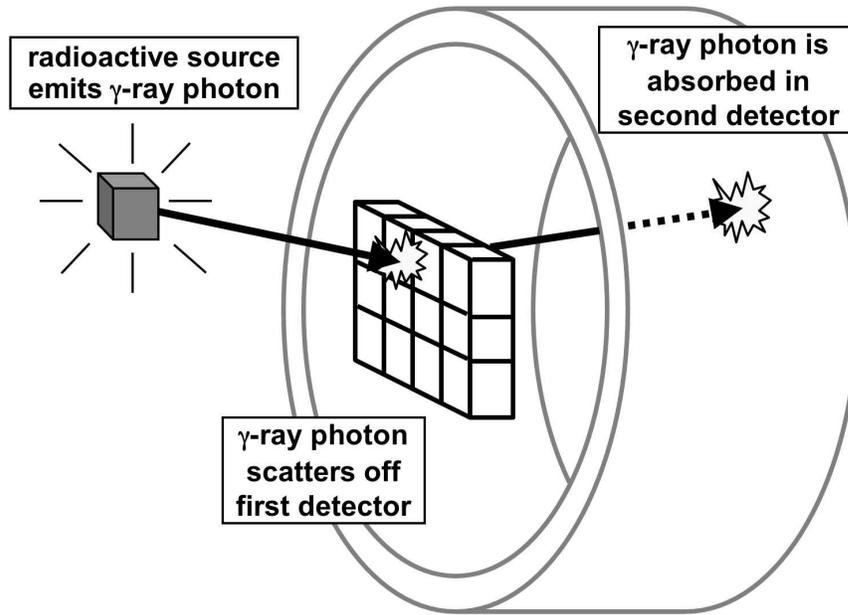


Figure 3.1: Compton SPECT detection process.

energy deposited in the first detector can be used to calculate the scattering angle relative to the initial γ -ray photon's direction. With the additional two position measurements, one can define the γ -ray's initial direction to within the surface of a cone defined by the scattering angle with apex located at the first detector position measurement, as shown in figure 3.2. If multiple γ -rays from the same source are measured, the “uncertainty cones” associated with each measurement would ideally intersect at the source location.

If noiseless detectors existed, the image reconstruction problem for a point source would conceptually reduce to determining the locus of intersections of every single cone from every single measurement. However, detectors are noisy and measurements are imperfect. Rather than having conic sections intersecting at points, there are intersecting “fuzzy” conical surfaces forming regions of where the incident photon most probably came from.

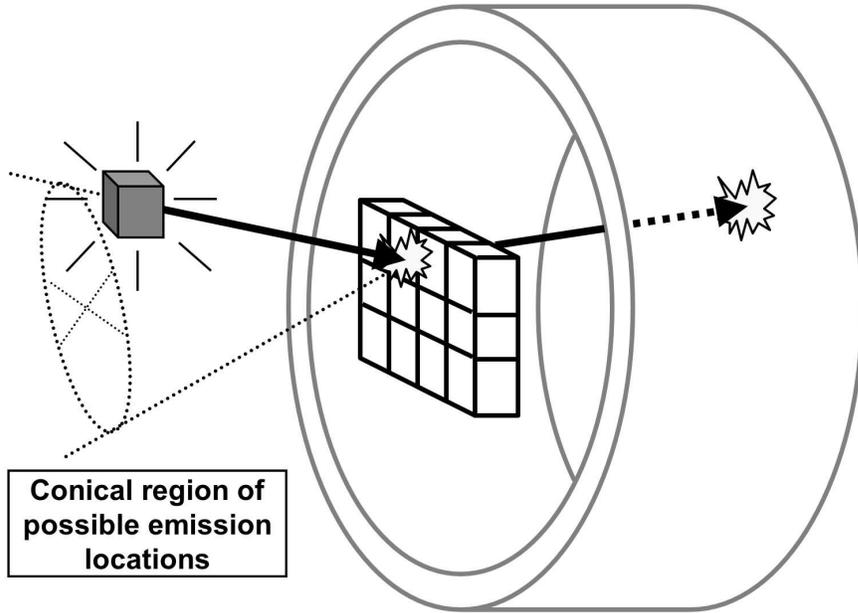


Figure 3.2: Conical ambiguity of the γ -ray emission source location.

3.1 Photon Transport and Measurement Statistics

A Compton SPECT camera coincidence measurement vector $\underline{y} \in \mathbb{R}^k$ consists of a collection of measurements $\{e_2, \underline{z}_2, e_1, \underline{z}_1\}$ of a scattered photon. The (scalar) energy measurements e_1, e_2 are of the energy deposited in the 1st and 2nd detector, and $\underline{z}_1, \underline{z}_2$ are the respective position measurement vectors resulting in a $k = 8$ dimensional measurement vector \underline{y} . Note that for a thin planar detector (such as a silicon chip) the measurement is often thought of as 2D since only the rows/columns of the chip are read out. In this case, the physical position of the chip can be used as an approximation of the missing 3rd dimensional component.

In Compton SPECT imaging, the dominant uncertainties are due to Compton scattering, Doppler broadening, detector energy efficiency and position uncertainty, and attenuation [71]. In this chapter we will introduce an overall detection probability and a measurement probability model that incorporates these dominant error sources. The technique involved is to model each step of the measurement process as a conditional probability, and then to combine them to determine the overall

probability. Note that no mention was made of scattering inside the source, false coincidences, multiple scatters, or random detections. These are important effects in any coincident-event detection system, but for now will be either assumed negligible or ignored outright.

Let $\underline{y}' \in \mathbb{R}^k$ be the noise-free value of the bias- and noise-corrupted measurement of \underline{y} . The statistics of \underline{y}' are determined by the underlying image or object of interest $\lambda(\underline{x})$ and the physics that describe the image formation process (i.e. collimators, photon transport, scatter, etc.). The statistics of \underline{y} , conditioned on \underline{y}' , depend on the noise characteristics of the sensors acquiring the measurements. Borrowing the notation of [10], the probability density function of an individual measurement $p(\underline{y}|\lambda, D)$ (2.11) can be expanded out as

$$p(\underline{y}|\lambda, D) = \int p_m(\underline{y}|\lambda, D, \underline{y}') p_i(\underline{y}'|\lambda, D) d\underline{y}', \quad (3.1)$$

where $p_m(\underline{y}|\lambda, D, \underline{y}')$ describes the statistics of the detected measurement \underline{y} of the noiseless photon attribute vector \underline{y}' and $p_i(\underline{y}'|\lambda, D)$ describes the image formation process. Models for these two terms are developed in the next four sections.

3.2 Compton Scattering and Doppler Broadening

Although position and energy measurement errors in the first and second detectors will add uncertainty to the position of the emitted γ -ray, by far the dominant error source for low energy Compton scatter imaging is Doppler broadening [71]. To understand the effect of Doppler broadening, we will first examine the classic Compton scatter angle-energy relationship. Compton scatter [13, 67] is when an incoming photon interacts with the electron shell of an atom, transfers energy e_1 to a recoil electron, and scatters off in some direction $\underline{\Omega} = \underline{\Omega}(\varphi, \theta)$ relative to its initial direction. The recoil electron is re-absorbed and releases energy e_1 , and by conservation of energy the scattered γ -ray has energy $e_\gamma = e_o - e_1$. The geometry of a Compton

scatter event is illustrated in figure 3.3.

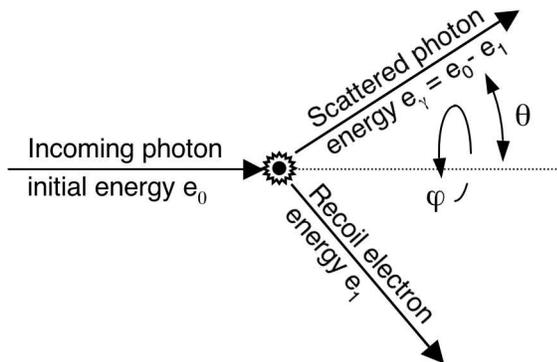


Figure 3.3: Geometry of a Compton scatter event.

Without any prior knowledge, the scattering angles and energy are random variables with joint probability density function (jpdf) of the form $p(\varphi, \theta, e_1)$. Under the assumptions of incoherent scattering of an unpolarized photon scattering off a free electron at rest, the out-of-plane scattering angle φ is independent of the in-plane scattering angle θ and uniformly distributed over the range $[-\pi, \pi)$ [67]. By conservation of energy and momentum, one can derive the relationship between the energy e_o of the incoming γ -ray photon, the energy deposited e_1 , the energy e_γ of the scattered γ -ray photon, the in-plane scattering angle θ , and the electron rest mass e_{e^-} (nominally 511keV).

$$\cos \theta = 1 + \frac{e_{e^-}}{e_0} - \frac{e_{e^-}}{e_\gamma} \quad (3.2)$$

$$= 1 + \frac{e_{e^-}}{e_0} - \frac{e_{e^-}}{e_o - e_1} \quad (3.3)$$

By measuring the energy e_1 deposited by the recoil electron into a detector, ideally one can calculate the scattering angle $\theta = \theta(e_1)$ via (3.3) to within a margin of error determined by the measurement accuracy of e_1 .

The marginal density $p(\theta)$ of the scattered γ -rays with respect to solid angle can be obtained from the Klein-Nishina (KN) Differential Cross Section (DCS) formula

$$f_{KN}(\theta) = Zr_e^2 \left(\frac{1}{1 + \kappa(1 - \cos \theta)} \right)^2 \left(\frac{1 + \cos^2 \theta}{2} \right)$$

$$\left(1 + \frac{\kappa^2(1 - \cos \theta)^2}{[1 + \cos^2 \theta](1 + \kappa(1 - \cos \theta))}\right), \quad (3.4)$$

where Z is the atomic number of the atom scattered off, r_e^2 is the classical electron radius, and $\kappa = e_o/e_{e^-}$ [9,67]. However, this formula gives the differential scattering cross-section with respect to solid angle relative to the *effective scattering area* of the electron shell. A simple normalization and a conversion from solid to planar angle gives the form of a density for θ ,

$$p(\theta) = K f_{KN}(\theta) \sin \theta \quad (3.5)$$

where K is a scale factor such that $\int_0^\pi p(\theta) d\theta = 1$. Note that there is an implicit conditioning on the event that a Compton scatter event occurred.

An interesting interpretation of (3.3) comes from the identity $p(\theta, e_1) = p(e_1|\theta)p(\theta)$. The conditional density $p(e_1|\theta)$ is the density of e_1 conditioned on a fixed value of θ . However, there is a deterministic relationship between e_1 and θ via equation (3.3), thus e_1 is not random given θ . With this in mind, the conditional density $p(e_1|\theta)$ can be modeled as a delta-function, $p(e_1|\theta) = \delta(e_1 - e_1(\theta))$, where $e_1 = e_1(\theta)$ is the unique solution to (3.3). The density $p(e_1)$ can be solved for by integrating $p(e_1, \theta)$ over θ ,

$$p(e_1) = \frac{e_{e^-}}{(e_o - e_1)^2} f_{KN}(\theta(e_1)). \quad (3.6)$$

Figure 3.4 shows the distribution of scattering angles and associated energy deposited for a Compton scattering event for an incident γ -ray energy of 140.4keV as predicted by the Klein-Nishina equations (3.5, 3.6). Note the maximum energy deposited of $e_1 = 49.8\text{keV}$ corresponding to a $\theta = 180$ deg scatter in (3.3). This feature is referred to as the *Compton edge* [67].

In the classical Klein-Nishina treatment of Compton scatter (3.2 - 3.6), the incident electron is assumed to be at rest. A more accurate description is given by the relativistic impulse response approximation (IA) [88]. The non-zero momentum distribution of the electron struck by the incoming γ -ray results in a random momen-

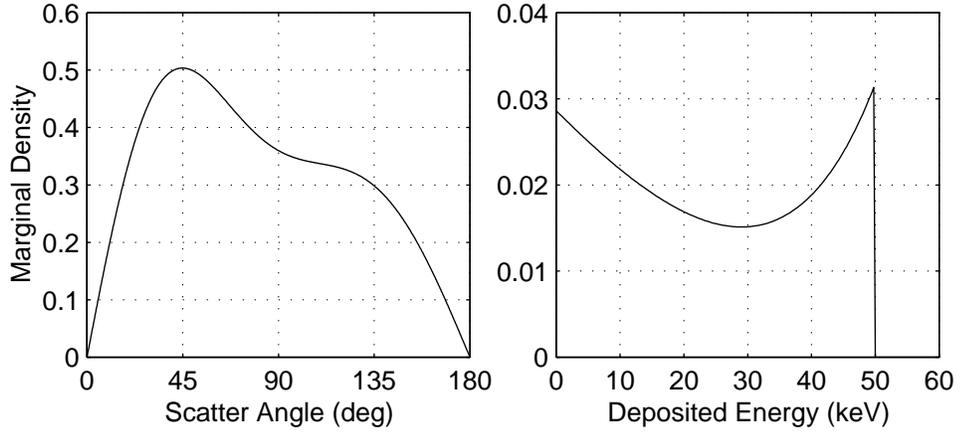


Figure 3.4: Klein-Nishina Compton scattering angle distribution $p(\theta)$ (left) and associated energy distribution $p(e_1)$ (right) for an incident γ -ray energy of 140.4keV.

tum transfer, and the Compton energy-angle relationships of (3.2, 3.3) are no longer deterministic.

The joint density $p(\theta, e_\gamma)$ of the scattered γ -ray energy and angle with respect to solid angle can be obtained from Compton Double Differential Cross Section (DDCS) model [19] based on the IA formula,

$$f_{IA}(\theta, e_\gamma) = \frac{1}{\alpha} \frac{r_e^2 e_\gamma}{2 e_o q c} \frac{X J'(p'_z)}{\sqrt{1 + (\alpha p'_z)^2}} \quad (3.7)$$

where

α *fine structure constant* $\approx 1/137$

r_e *classical electron radius*

e_o *initial photon energy*

e_γ *scattered photon energy*

θ *in-plane scattering angle*

J' *normalized Compton momentum profile*

p'_z *normalized projection of electron momentum*

$$= \frac{(e_o - e_\gamma) - \frac{(e_o e_\gamma)}{e_e} (1 - \cos \theta)}{\alpha q c}$$

$$\begin{aligned}
qc &= \sqrt{e_o^2 + e_\gamma^2 - 2e_o e_\gamma \cos \theta} \\
X &= (R/R') + (R'/R) + 2(1/R - 1/R') + (1/R - 1/R')^2 \\
R &= (e_o/e_{e^-})[\sqrt{1 + (\alpha p'_z)^2} + (1/qc)(e_o - e_\gamma \cos \theta)(\alpha p'_z)] \\
R' &= R - (e_o/e_{e^-})(e_\gamma/e_{e^-})(1 - \cos \theta)
\end{aligned}$$

Similar to (3.4), this formula gives the differential scattering cross-section with respect to energy and solid angle relative to the effective scattering area of the electron shell.

There are both analytical results and experimental measurements for the Compton momentum profile function J' [14,33,87]. However, one needs to use appropriate data because momentum distribution in molecular orbitals is important in Compton scatter imaging. Detailed models for the DDCS will take into account the binding energy of individual electron shells, where the overall profile is a weighted summation. However, they are difficult to calculate from first principles. In addition, the data available for typical detector materials such as Silicon or Germanium is from experimental measurements. They specify an overall profile function where the effect of binding energy is “lumped” into the overall profile. With these approximation in mind, one can model the joint density for θ and e_γ by appropriate normalization of (3.7)

$$p(\theta, e_\gamma) = K f_{IA}(\theta, e_\gamma) \sin \theta. \quad (3.8)$$

As in (3.5), K is a scale factor such that $\int_0^\infty \int_0^\pi p(\theta, e_\gamma) d\theta de_\gamma = 1$, along with an implicit conditioning on the event that a Compton scatter event occurred.

Figure 3.5 shows the joint density $p(e_1, \theta)$ for a 140.4 keV γ -ray scattering off Silicon, along with the associated marginal density functions and how they compare to the Klein-Nishina model. The lower-right panel of figure 3.5 shows the conditional scattering-angle distribution for a nominal energy deposited in a Compton interaction of 17keV. Even with perfect energy resolution in the first detector, there can still be significant uncertainty in the amount of energy deposited (and thus angle) of the scattered γ -ray.

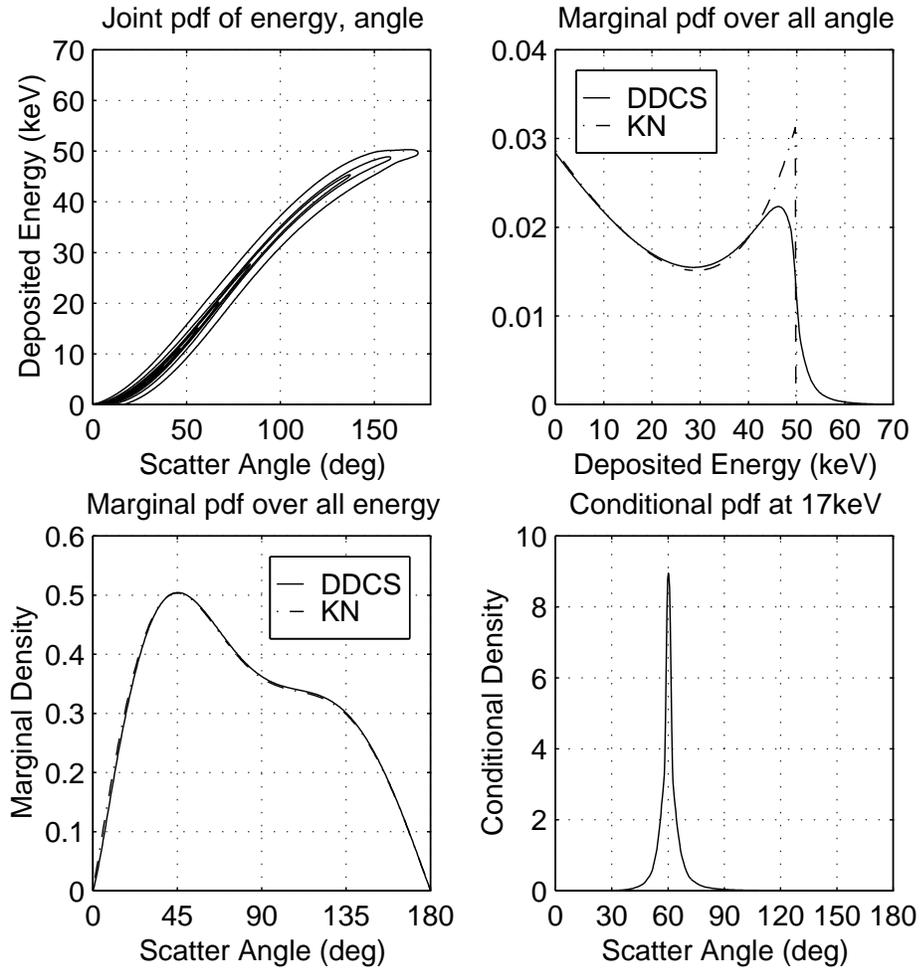


Figure 3.5: Compton scattering joint energy-angle distribution $p(e_1, \theta)$ of a 140.4keV γ -ray interacting with Silicon (top-left), and associated marginal distributions $p(e_1)$ (top-right), $p(\theta)$ (bottom-left), along with conditional scattering angle distributions for a nominal deposited energy of 17keV (bottom-right). The Klein-Nishina marginal distributions are overlaid in comparison.

3.3 Energy Resolution

For a solid state silicon detector such as that used in the experimental Compton camera of [71], a fairly simple model of the energy resolution takes into account the uncertainty in the number of electron hole-pairs generated by each γ -ray interacting with the detector [9]. The quantity of electrons N liberated created by each γ -ray photon interacting with the detector is proportional to the actual energy e' deposited by the recoil electron and the mean ionization energy w of creating an electron hole-

pair,

$$N = \frac{e'}{w}. \quad (3.9)$$

Ideally the quantity of electrons liberated would be exactly N . However, there is a statistical element to the number of electrons liberated with an observed variance proportional to N . The Fano factor f [9] is used to model the variance in the quantity of electrons liberated,

$$\sigma_N^2 = \frac{fe'}{w}. \quad (3.10)$$

The electronics of the silicon detector also have a read-out noise associated with the actual charge measurement, which is typically modeled as a zero-mean gaussian r.v. with variance σ_{ro}^2 . A reasonable assumption is that the read-out noise is independent of the ionization noise since they are due to completely separate physical processes. The total variance σ_T^2 is simply the sum of the component variances. Since σ_{ro}^2 and σ_T^2 are typically given in units of $(eV)^2$ whereas the ionization noise σ_N^2 given in (3.10) is in units of $(electrons)^2$, there is a conversion factor of w^2 ,

$$\begin{aligned} \sigma_T^2 &= w^2\sigma_N^2 + \sigma_{ro}^2 \\ &= fwe' + \sigma_{ro}^2. \end{aligned} \quad (3.11)$$

The measured energy e to a given noiseless energy e' can then be modeled as a gaussian r.v. with conditional density

$$p(e|e') = \frac{1}{\sqrt{2\pi\sigma_T^2}} e^{-\frac{1}{2} \frac{(e - e')^2}{\sigma_T^2}} \quad (3.12)$$

where σ_T^2 is given in (3.11). For silicon at room temperature, $w \approx 3.62 \left(\frac{eV}{e^-}\right)$ and $f \approx 0.14$.

The 2nd detector in the experimental Compton camera is an Anger camera, consisting of a grid of photo-multiplier tubes behind a slab of scintillating crystal. Since

the energy resolution is nearly an order of magnitude worse than a solid state detector, and there are large number of errors sources due to tube calibration, incomplete absorptions in the scintillator, etc., a simpler model is used. One is to calibrate the detector with a mono-energetic source and give the energy resolution as a percent-ratio of the photopeak width to the source energy, assuming an approximately Gaussian response of the detector energy measurement. The energy e measured for a noiseless energy e' deposited can also be modeled as a Gaussian r.v. with conditional density

$$p(e|e') = \frac{1}{\sqrt{2\pi}\sqrt{ke'}} e^{-\frac{1}{2} \frac{(e - e')^2}{ke'}} \quad (3.13)$$

where k is the energy resolution factor determined by calibration. For the NaI 2nd detector used in the experimental Compton camera, typical FWHM energy resolutions are $\approx 9\% - 15\%$ of the source energy [58, 71].

3.4 Position Resolution

For the experimental Compton camera [71], the 1st detector is a Silicon chip with pixel size $\approx 1.4mm \times 1.4mm \times 0.03mm$, while the 2nd detector is an Anger camera with spatial resolution $\approx 3mm \times 3mm$ with a scintillation crystal $1.27cm$ thick. We will assume that the spatial errors due to quantization are negligible, thus the measured position \underline{z} is approximately identical to the noiseless position \underline{z}' . A reasonable model for the position measurement conditional densities is

$$p(\underline{z}|\underline{z}') = \delta(\underline{z} - \underline{z}'). \quad (3.14)$$

3.5 Transmission and Attenuation

Given a photon pencil-beam with initial flux Φ_0 transmitting a distance r through a material or mixture of materials with path-dependent linear attenuation coefficient function $\mu(r)$, the resulting flux $\Phi(r)$ is given by Beer's law [78],

$$\Phi(r)/\Phi_0 = e^{-\int_0^r \mu(l) dl}$$

which models the *Probability of Transmission* of a photon through an attenuating medium. Note that the linear attenuation coefficient function is dependent on the photon energy e , i.e. $\mu = \mu(r, e)$. For notational simplicity we will implicitly assume this energy dependence in the following derivations.

Define T as the event that a photon transmits a distance r through a material with attenuation coefficient $\mu(r)$. The probability of transmission $P(T)$ as a function of r is given by

$$P_T(r) = e^{-\int_0^r \mu(l) dl} \quad (3.15)$$

for $r \in [0, \infty)$. Define A as the event that a photon is absorbed within (i.e. does not transmit through) the material. Since the event A is mutually exclusive of the event T , the probability of absorption $P(A)$ as a function of r is given by

$$\begin{aligned} P_A(r) &= 1 - P_T(r) \\ &= 1 - e^{-\int_0^r \mu(l) dl}. \end{aligned} \quad (3.16)$$

The absorption density $p_A(r) = \frac{d}{dr}P_A(r)$ is therefore

$$p_A(r) = e^{-\int_0^r \mu(l) dl} \mu(r), \quad (3.17)$$

and for any $r_o \in [0, r)$ the absorption density $p_A(r)$ can be rewritten as

$$\begin{aligned} p_A(r) &= e^{-\int_0^{r_o} \mu(l) dl} e^{-\int_{r_o}^r \mu(l) dl} \mu(r) \\ &= P_T(r_o) e^{-\int_{r_o}^r \mu(l) dl} \mu(r). \end{aligned}$$

For $r \geq r_o$ we can define the conditional density $p_A(r|r_o)$ as

$$p_A(r|r_o) = e^{-\int_{r_o}^r \mu(l) dl} \mu(r) \quad (3.18)$$

which gives

$$p_A(r) = P_T(r_o) p_A(r|r_o) \quad (3.19)$$

and

$$\begin{aligned}
\int_{r_o}^r p_A(r) dr &= P_T(r_o) \int_{r_o}^r p_A(r|r_o) dr \\
&= P_T(r_o) \left(1 - e^{-\int_{r_o}^r \mu(l) dl}\right) \\
&= P_T(r_o) P_A(r|r_o).
\end{aligned} \tag{3.20}$$

Thus the probability of absorption within the interval $[r_o, r)$ is simply the product of the probability of transmission $P_T(r_o)$ through the interval $[0, r_o)$ and the conditional probability of absorption $P_A(r|r_o)$ within the interval $[r_o, r)$.

In order to calculate the probability of absorption of a photon by an object, we need to integrate over all space subtended by the object. Consider a point-source photon emitter at spatial location \underline{x}_0 and an arbitrary object O as shown in Figure 3.6. Because of the form of the attenuation equations, spherical coordinates with origin at \underline{x}_0 are the most convenient way to solve this problem. For notational simplicity we will express the direction on the unit sphere as a unit vector $\underline{\Omega}$ rather than as a pair of angles.

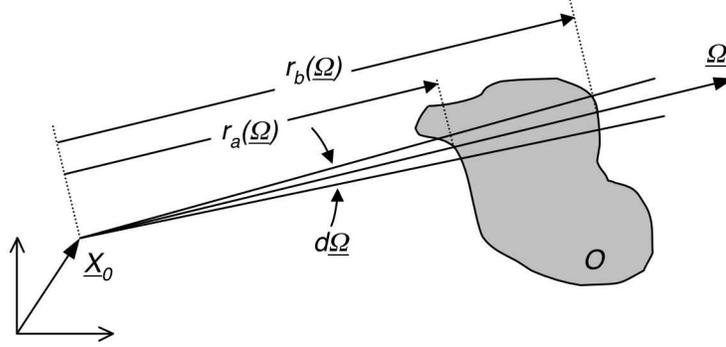


Figure 3.6: Attenuation geometry through an arbitrary object.

Let $S_{\underline{\Omega}}$ be the set of all unit direction vectors $\underline{\Omega}$ with origin at \underline{x}_0 subtended by the object O, let $d\underline{\Omega}$ be a differential solid angle about the unit vector $\underline{\Omega}$, and let $[r_a, r_b) = [r_a(\underline{\Omega}), r_b(\underline{\Omega}))$ be the interval intercepted by the object O in the direction $\underline{\Omega}$ relative to origin \underline{x}_0 . The dependence of the integration limits r_a, r_b on $\underline{\Omega}$ is due to

the different path lengths through the object as a function of path direction $\underline{\Omega}$. The total probability of absorption within the object O conditioned on an emission from \underline{x}_0 is thus given by integrating (3.19) over S and r ,

$$P(A|\underline{x}_0) = \int_{S_{\underline{\Omega}}} P_T(r_a(\underline{\Omega})) \int_{r_a(\underline{\Omega})}^{r_b(\underline{\Omega})} p_A(r|r_a(\underline{\Omega}), \underline{\Omega}, \underline{x}_0) p(\underline{\Omega}|\underline{x}_0) dr d\underline{\Omega} \quad (3.21)$$

where $p_A(r|r_a(\underline{\Omega}), \underline{\Omega}, \underline{x}_0)$ is simply (3.18), conditioned on a particular radial interval defined by $\underline{\Omega}$ and \underline{x}_0 . By substituting the inner integral over r in equation (3.21) with (3.20) and assuming a uniform distribution of emission directions over the unit sphere (i.e. $p(\underline{\Omega}|\underline{x}_0) = \frac{1}{4\pi}$) gives

$$P(A|\underline{x}_0) = \frac{1}{4\pi} \int_{S_{\underline{\Omega}}} P_T(r_a(\underline{\Omega})) P_A(r_b(\underline{\Omega})|r_a(\underline{\Omega}), \underline{\Omega}, \underline{x}_0) d\underline{\Omega}. \quad (3.22)$$

Solving for the outer integral over $\underline{\Omega}$ will (in general) not have a closed-form solution and must be solved using numerical techniques [1, 46]. More important, however, is the technique of modeling the overall probability of an event as a chain of conditional probabilities. This will be used extensively throughout this chapter.

3.6 Detection

Let A_1 be the event that a γ -ray is absorbed by the 1st detector, let C_1 be the event that the γ -ray Compton scatters off the 1st detector, and let E_1 be the event that an energy measurement occurs within the measurement threshold of the 1st detector. Let A_2 be the event that a γ -ray is absorbed by the 2nd detector and let E_2 be the event that an energy measurement occurs within the measurement threshold of the 2nd detector. Define the overall detection event D as the intersection of the events A_1, E_1, C_1, A_2, E_2 . The overall probability of detection, conditioned on an emission from \underline{x}_0 , is

$$P(D|\underline{x}_0) = P(E_2, A_2, C_1, E_1, A_1|\underline{x}_0). \quad (3.23)$$

Coincidence event measurements in a Compton SPECT camera are typically given as a vector \underline{y} with components $\{e_2, z_2, e_1, z_1\}$, where e_1, e_2 are the (scalar) energy

measurements from the 1st and 2nd detector and $\underline{z}_1, \underline{z}_2$ are the respective position measurement vectors. With this in mind, the previous expression for $P(D|\underline{x}_0)$ can now be expressed as an integral over all measurements \underline{y} within the *set of detectable measurements* S ,

$$P(D|\underline{x}_0) = \int_S P(C_1|\underline{y}, \underline{x}_0)p(\underline{y}|\underline{x}_0) d\underline{y}, \quad (3.24)$$

where $d\underline{y} = de_2 d\underline{z}_2 de_1 d\underline{z}_1$, $d\underline{x} = dx dy dz$, and $S = (S_{e_2}, S_{d_2}, S_{e_1}, S_{d_1})$ is the set of all possible measurements $(e_2, \underline{z}_2, e_1, \underline{z}_1)$ that are counted as valid measurements and/or within the detection threshold. Note that the event C_1 , the event that the γ -ray undergoes a Compton scatter, is not directly measurable.

Since the physics modeling photon scatter and attenuation are easier to describe in spherical coordinates, we will express the event measurement vector in terms of their spherical coordinate equivalents. Let $\underline{y}' = (e'_2, \underline{z}'_2, e'_1, \underline{z}'_1)$ be the noiseless event vector of the γ -ray's journey through the detectors. Given $\underline{z}'_1, \underline{z}'_2$, define two spherical coordinates with origins at \underline{x}_0 and \underline{z}'_1 . For now, we will express the direction on the unit sphere as a unit vector $\underline{\Omega}' = \underline{\Omega}'(\psi', \phi')$ rather than as a pair of angles (ψ', ϕ') where

$$r'_1 = \|\underline{z}'_1 - \underline{x}_0\|, \quad (3.25)$$

$$\underline{\Omega}'_1 = \frac{\underline{z}'_1 - \underline{x}_0}{r'_1}, \quad (3.26)$$

$$r'_2 = \|\underline{z}'_2 - \underline{z}'_1\|, \quad (3.27)$$

$$\underline{\Omega}'_2 = \frac{\underline{z}'_2 - \underline{z}'_1}{r'_2}, \quad (3.28)$$

as shown in Figure 3.7.

Let $\underline{b}' = (e'_2, r'_2, \underline{\Omega}'_2, e'_1, r'_1, \underline{\Omega}'_1)$ be the noiseless measurement vector \underline{y}' represented in spherical coordinates. Integrating over all possible noiseless measurement \underline{b}' gives

$$P(D|\underline{x}_0) = \int_S \int P(C_1|\underline{b}', \underline{x}_0)p(\underline{b}'|\underline{x}_0)p(\underline{y}|\underline{b}', \underline{x}_0) d\underline{b}' d\underline{y}, \quad (3.29)$$

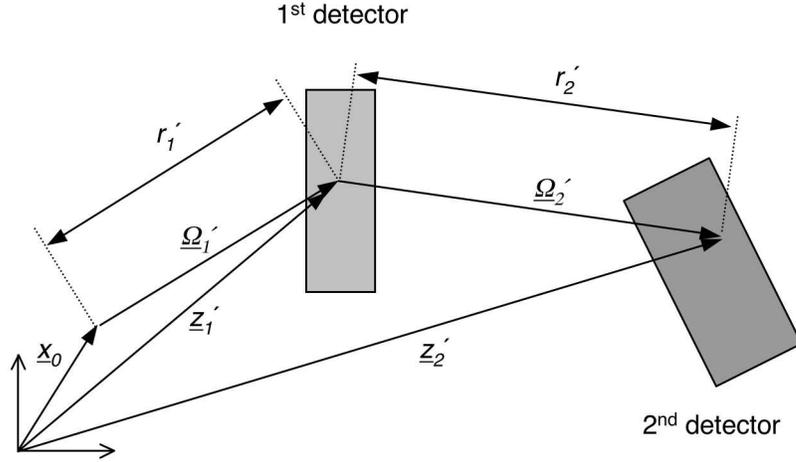


Figure 3.7: Geometry of coincidence measurements in Compton SPECT.

where $p(\underline{y}|\underline{b}', \underline{x}_0)$ is the response function of the measurement \underline{y} to a particular noiseless value \underline{b}' and initial emission point \underline{x}_0 (i.e. the measurement blur kernel). The expression $P(C_1|\underline{b}', \underline{x}_0)p(\underline{b}'|\underline{x}_0) d\underline{b}'$ inside (3.29) can be re-written by iterating conditional distributions, and noting that $d\underline{b}' = de_2' dr_2' d\underline{\Omega}_2' de_1' dr_1' d\underline{\Omega}_1'$.

$$\begin{aligned}
P(C_1|\underline{b}', \underline{x}_0)p(\underline{b}'|\underline{x}_0) d\underline{b}' = & \\
& p(e_2'|r_2', \underline{\Omega}_2', e_1', C_1, r_1', \underline{\Omega}_1', \underline{x}_0)p_A(r_2'|\underline{\Omega}_2', e_1', C_1, r_1', \underline{\Omega}_1', \underline{x}_0) \\
& p(\underline{\Omega}_2', e_1'|C, r_1', \underline{\Omega}_1', \underline{x}_0)P(C|r_1', \underline{\Omega}_1', \underline{x}_0)p_A(r_1'|\underline{\Omega}_1', \underline{x}_0) \\
& p(\underline{\Omega}_1'|\underline{x}_0)de_2' dr_2' d\underline{\Omega}_2' de_1' dr_1' d\underline{\Omega}_1'. \tag{3.30}
\end{aligned}$$

For a typical γ -ray detector of finite size, the integral over r_1' will be over some finite interval $[r_{1a}', r_{1b}')$, in which r_{1a}' is the radial distance where the pencil-beam in the direction of $\underline{\Omega}_1'$ first intercept the 1st detector, and likewise for the integral over r_2' . Substituting equation (3.19) into the absorption density terms in (3.30) gives

$$p_A(r_1'|\underline{\Omega}_1', \underline{x}_0) = p_A(r_1'|r_{1a}', \underline{\Omega}_1', \underline{x}_0)P_T(r_{1a}'|\underline{\Omega}_1', \underline{x}_0), \tag{3.31}$$

and

$$\begin{aligned}
p_A(r_2'|\underline{\Omega}_2', e_1', C_1, r_1', \underline{\Omega}_1', \underline{x}_0) = & \\
& p_A(r_2'|r_{2a}', \underline{\Omega}_2', e_1', C_1, r_1', \underline{\Omega}_1', \underline{x}_0)P_T(r_{2a}'|\underline{\Omega}_2', e_1', C_1, r_1', \underline{\Omega}_1', \underline{x}_0). \tag{3.32}
\end{aligned}$$

Each term can now be identified by inspection.

- $p(\underline{\Omega}'_1|\underline{x}_0) d\underline{\Omega}'_1$

Probability of a γ -ray emission from \underline{x}_0 emitted in direction $\underline{\Omega}'_1$. Nominally distributed uniformly over the unit sphere, i.e. $p(\underline{\Omega}'_1|\underline{x}_0) = \frac{1}{4\pi}$.

- $P_T(r'_{1a}|\underline{\Omega}'_1, \underline{x}_0)$

Probability of transmission from the point \underline{x}_0 in direction $\underline{\Omega}'_1$ for a distance r'_{1a} . This is also referred to as the *source attenuation factor*.

- $p_A(r'_1|r'_{1a}, \underline{\Omega}'_1, \underline{x}_0) dr'_1$

Conditional probability of absorption within interval $[r'_{1a}, r'_{1a} + dr'_1)$ inside the first detector, conditioned on a γ -ray emission from \underline{x}_0 emitted in direction $\underline{\Omega}'_1$.

- $P(C|r'_1, \underline{\Omega}'_1, \underline{x}_0)$

Probability of a Compton scatter, conditioned on a γ -ray emission from \underline{x}_0 emitted in direction $\underline{\Omega}'_1$ being absorbed by the 1st detector at position r'_1 . For the purposes of this paper, this probability is given by the ratio of the Compton scatter cross-section to the total material cross-section ($\frac{\mu_c}{\mu_t}$). Note that bulk properties are assumed, and any directionally-dependent crystalline effects are ignored.

- $p(\underline{\Omega}'_2, e'_1|C_1, r'_1, \underline{\Omega}'_1, \underline{x}_0) d\underline{\Omega}'_2 de'_1$

Probability of a γ -ray depositing energy e'_1 and scattering in direction $\underline{\Omega}'_2$. Nominally given by the Compton scattering jpdf (3.8).

- $P_T(r'_{2a}|\underline{\Omega}'_2, e'_1, C_1, r'_1, \underline{\Omega}'_1, \underline{x}_0)$

Probability of transmission from the point r'_1 in direction $\underline{\Omega}'_2$ for a distance r'_{2a} . When the path between the two detectors is through air, this term is negligible. However, for certain detector geometries where the path between the 1st and 2nd detector is through tissue [111], this term may also include attenuation.

Note the energy dependence on e'_1 , since the transmitted photon has energy $e_o - e'_1$, where e_o is the initial photon energy.

- $p_A(r'_2|r'_{2a}, \underline{\Omega}'_2, e'_1, C_1, r'_1, \underline{\Omega}'_1, \underline{x}_0) dr'_2$

Conditional probability of absorption within interval $[r'_{2a}, r'_{2a} + dr'_2)$ inside the second detector, conditioned on a γ -ray Compton scattering at location r'_1 and heading off in direction $\underline{\Omega}'_2$ with energy $e_o - e'_1$.

- $p(e'_2|r'_2, \underline{\Omega}'_2, e'_1, C_1, r'_1, \underline{\Omega}'_1, \underline{x}_0) de'_2$

Probability of depositing energy e'_2 in 2nd detector, conditioned on all the previous terms. Assuming complete absorption, it is given by $\delta(e'_2 - (e_o - e'_1))$.

The Compton scattering jpdf (3.8) is typically given as a function of the noiseless out-of- and in-plane scattering angles (φ', θ') . The noiseless in-plane scatter angle θ' can be calculated directly from the incoming and outgoing γ -ray direction vectors $\underline{\Omega}'_1$ and $\underline{\Omega}'_2$.

$$\cos \theta' = \underline{\Omega}'_1^T \underline{\Omega}'_2. \quad (3.33)$$

The expression for the out-of-plane scattering angle φ' is a bit involved. However, a reasonable assumption is that the incoming γ -ray is unpolarized¹ before scattering. Then the out-of-plane angle φ' is uniformly distributed over $[-\pi, \pi)$, and is independent of the in-plane scattering angle θ' and the energy deposited e'_1 . Under this condition the Compton scattering jpdf is equivalent to the following expression,

$$\begin{aligned} p(\underline{\Omega}'_2, e'_1|C_1, \dots) &= p(\varphi', \theta', e'_1|C_1, \dots) \\ &= p(\varphi'|C_1, \dots)p(\theta', e'_1|C_1, \dots) \\ &= \frac{1}{2\pi}p(\theta', e'_1|C_1, \dots), \end{aligned} \quad (3.34)$$

¹Compton scatter does polarize the γ -ray, which must be taken into account when modeling multiple scattering events. At the particle energies involved in gamma-ray astronomy [65] the probability of multiple scattering events is significant.

where $p(\theta', e'_1 | C_1, \dots)$ is the expression given earlier in (3.8) after substituting $e'_\gamma = e_o - e'_1$.

Spherical coordinates have been used to model the scattering physics up to now. Since most detectors can be modeled as blocks or thin plates, we will convert to cartesian coordinates as needed in order to simplify some expressions for the overall detection sensitivity and for the measurement probability expression to be derived later.

The unit vector $\underline{\Omega}'$ used throughout the derivation is composed of two spherical coordinate angles, the azimuthal angle $\psi' \in [-\pi, \pi)$ and elevation angle $\phi' \in [0, \pi)$ where

$$\underline{\Omega}' = [\sin \phi' \cos \psi', \sin \phi' \sin \psi', \cos \phi']^T, \quad (3.35)$$

and the differential solid angle $d\underline{\Omega}'$ is

$$d\underline{\Omega}' = \sin \phi' d\phi' d\psi'. \quad (3.36)$$

Since the Jacobian involved in converting between spherical and cartesian coordinates is $|J| = \sin \phi' r'^2$, the differential element $d\underline{b}'$ becomes

$$\begin{aligned} d\underline{b}' &= de'_2 dr'_2 d\underline{\Omega}'_2 de'_1 dr'_1 d\underline{\Omega}'_1 \\ &= \sin \phi'_1 \sin \phi'_2 de'_2 dr'_2 d\phi'_2 d\psi'_2 de'_1 dr'_1 d\phi'_1 d\psi'_1 \\ &= \sin \phi'_1 \sin \phi'_2 \frac{de'_2 dx'_2 dy'_2 dz'_2}{|J_2|} \frac{de'_1 dx'_1 dy'_1 dz'_1}{|J_1|} \\ &= \frac{de'_2 dz'_2}{r_2'^2} \frac{de'_1 dz'_1}{r_1'^2} \\ &= \frac{1}{r_2'^2} \frac{1}{r_1'^2} d\underline{y}', \end{aligned} \quad (3.37)$$

where $\underline{z}'_1 = (x'_1, y'_1, z'_1)^T$, $d\underline{z}'_1 = dx'_1 dy'_1 dz'_1$, and similarly for \underline{z}'_2 . Substituting these previous expressions into (3.30) results in

$$P(C_1 | \underline{b}', \underline{x}_0) p(\underline{b}' | \underline{x}_0) d\underline{b}' = P(C_1 | \underline{y}', \underline{x}_0) p(\underline{y}' | \underline{x}_0) \frac{1}{r_2'^2} \frac{1}{r_1'^2} d\underline{y}' \quad (3.38)$$

and

$$\begin{aligned}
P(C_1|\underline{y}', \underline{x}_0)p(\underline{y}'|\underline{x}_0) \frac{1}{r_2'^2} \frac{1}{r_1'^2} d\underline{y}' = \\
\frac{1}{8\pi^2} \left(\frac{\mu_c}{\mu_t} \right) p(e_2'|r_2', \underline{\Omega}'_2, \dots) p_A(r_2'|r_{2a}', \underline{\Omega}'_2, \dots) \\
P_T(r_{2a}'|\underline{\Omega}'_2, e_1', \dots) p(\theta', e_1'|C_1, r_1', \dots) p_A(r_1'|r_{1a}', \underline{\Omega}'_1, \dots) \\
P_T(r_{1a}'|\underline{\Omega}'_1, \underline{x}_0) \frac{1}{r_{2a}'^2} \frac{1}{r_{1a}'^2} d\underline{y}'. \tag{3.39}
\end{aligned}$$

The measurement conditional density $p(\underline{y}|\underline{y}', \underline{x}_0)$ is the measured response \underline{y} to a noiseless measurement \underline{y}' . A reasonable assumption is to assume the individual event measurements $(e_2, \underline{z}_2, e_1, \underline{z}_1)$ in \underline{y} are conditionally independent of each other and of the initial source location \underline{x}_0 ,

$$p(\underline{y}|\underline{y}', \underline{x}_0) = p(e_2|e_2') p(\underline{z}_2|\underline{z}_2') p(e_1|e_1') p(\underline{z}_1|\underline{z}_1'). \tag{3.40}$$

Combining (3.40), (3.39), and (3.29) gives

$$\begin{aligned}
P(D|\underline{x}_0) = \frac{1}{8\pi^2} \left(\frac{\mu_c}{\mu_t} \right) \int \int p(e_2'|r_2', \underline{\Omega}'_2, \dots) p_A(r_2'|r_{2a}', \underline{\Omega}'_2, \dots) \\
P_T(r_{2a}'|\underline{\Omega}'_2, e_1', \dots) p(\theta', e_1'|C_1, r_1', \dots) p_A(r_1'|r_{1a}', \underline{\Omega}'_1, \dots) \\
P_T(r_{1a}'|\underline{\Omega}'_1, \underline{x}_0) p(e_2|e_2') p(\underline{z}_2|\underline{z}_2') p(e_1|e_1') p(\underline{z}_1|\underline{z}_1') \frac{1}{r_{2a}'^2} \frac{1}{r_{1a}'^2} d\underline{y}' d\underline{y}. \tag{3.41}
\end{aligned}$$

After all is said and done, we are still left with the problem of integrating (3.41) over all spatial dimensions and two energy measurements along with their noiseless values, at minimum a $2 \times 8 = 16$ (!) dimensional numerical integral for a 2D source, evaluated for each source location point \underline{x}_0 of interest. The key is to make a few simplifying assumptions that allow a relatively straightforward solution to (3.41) but still capture the dominant measurement errors.

3.7 Simplifications

The following simplifications to (3.41) result in an expression for the detection probability that only involves a 4-dimensional numerical integral. These simplifying

assumptions will be violated in some form or another in Sections 3.9 and 3.10. However, the resulting expression retains the dominant error sources and can be calculated in a reasonable time using fast adaptive numerical quadrature techniques [46]:

- No attenuation or internal scattering within the source object.
- Attenuation in the open air between the detectors is negligible.

$$P_T(r'_{1a}|\underline{\Omega}'_1, \underline{x}_0) = 1 \quad (3.42)$$

$$P_T(r'_{2a}|\underline{\Omega}'_2, e'_1, \dots) = 1 \quad (3.43)$$

- Constant isotropic attenuation in the detectors, i.e. $\mu(l) = \mu$.
- Thin planar first detector.

Note that the term $P_T(r'_{2a}|\underline{\Omega}'_2, e'_1, \dots)$ (3.43) includes the probability of transmission of a photon exiting the first detector. We will assume this occurs with probability 1.

- Negligible edge effects when integrating over thin first detector.
- Thick planar second detector.
- All spatial measurements are uncorrupted by noise or quantization errors (3.14), i.e. no finite detector resolution or depth-of-interaction effects.
- Gaussian distributed energy measurement errors in the 1st detector (3.12).
- Gaussian distributed energy measurement errors in the 2nd detector (3.13).
- Complete energy deposition in the 2nd detector

$$p(e'_2|r'_2, \underline{\Omega}'_2, \dots) = \delta(e'_2 - (e_o - e'_1)). \quad (3.44)$$

- The set S_{e_2} of detectable 2nd detector energy measurements e_2 is sufficiently wide to capture all possible events,

$$\int_{S_{e_2}} p(e_2|\dots) = 1. \quad (3.45)$$

With these assumptions in mind, we start with the sifting property of the δ -function in which the integrals over all noiseless spatial locations z'_1, z'_2 in (3.41) simplify to

$$\begin{aligned} P(D|\underline{x}_0) &= \frac{1}{8\pi^2} \left(\frac{\mu_c}{\mu_t} \right) \int_S \int_{\forall e'_2} \int_{\forall e'_1} p(e'_2|r_2, \underline{\Omega}_2, \dots) p_A(r_2|r_{2a}, \underline{\Omega}_2, \dots) \\ &\quad P_T(r_{2a}|\underline{\Omega}_2, e'_1, \dots) p(\theta, e'_1|C_1, r_1, \dots) p_A(r_1|r_{1a}, \underline{\Omega}_1, \dots) \\ &\quad P_T(r_{1a}|\underline{\Omega}_1, \underline{x}_0) p(e_2|e'_2) p(e_1|e'_1) \frac{1}{r_{2a}^2} \frac{1}{r_{1a}^2} de'_1 de'_2 d\underline{y}. \end{aligned}$$

Since by assumption there is no attenuation outside the detectors, and thus no other dependencies on the noiseless energies e'_1, e'_2 besides the terms $p(e'_2|r_2, \underline{\Omega}_2, \dots), p(\theta, e'_1), p(e_2|e'_2)$ and $p(e_1|e'_1)$, we can integrate separately over e'_1 and e'_2 to account for the energy measurement error,

$$\begin{aligned} p(\theta, e_1|C_1, r_1, \dots) &= \int p(\theta, e'_1|C_1, r_1, \dots) p(e_1|e'_1) de'_1 \\ p(e_2|r_2, \underline{\Omega}_2, \dots) &= \int p(e'_2|r_2, \underline{\Omega}_2, \dots) p(e_2|e'_2) de'_2. \end{aligned}$$

Note that in actual implementation the most efficient method would be to precompute this integral numerically and evaluate by interpolation. Combining with the assumption of zero attenuation gives

$$\begin{aligned} P(D|\underline{x}_0) &= \frac{1}{8\pi^2} \left(\frac{\mu_c}{\mu_t} \right) \int_S p(e_2|r_2, \underline{\Omega}_2, \dots) p_A(r_2|r_{2a}, \underline{\Omega}_2, \dots) \\ &\quad p(\theta, e_1|C_1, r_1, \dots) p_A(r_1|r_{1a}, \underline{\Omega}_1, \dots) \frac{1}{r_{2a}^2} \frac{1}{r_{1a}^2} d\underline{y}. \end{aligned}$$

As for the attenuation term inside the first detectors, it can be simplified by the following: Let $l_1 = r_1 - r_{1a}$ be the path length through the first detector from the front face location r_{1a} to the detection location r_1 . Next, we will judiciously define the spherical coordinates (ψ_1, ϕ_1) relative to the cartesian coordinates (x_1, y_1, z_1) such

that ϕ_1 is the angle of incidence with the (x_1, y_1) plane of the detector face as in shown in Figure 3.8. With this definition, the path-length is related to the depth z_1 inside

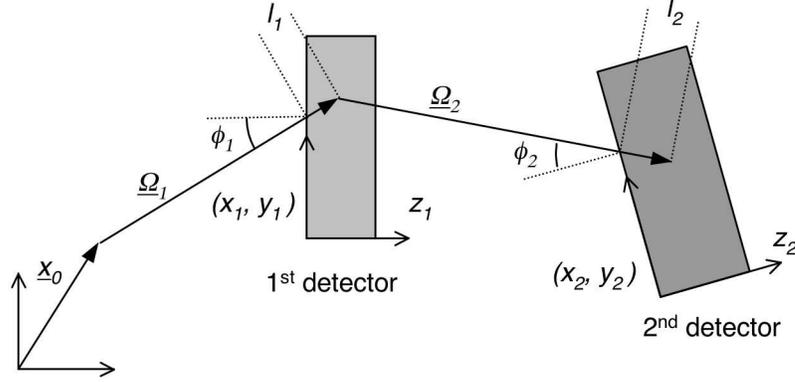


Figure 3.8: Angle of incidence geometry relative to detector face.

the detector by $z_1 = l_1 \cos \phi_1$,

$$\begin{aligned}
p_A(r_1 | r_{1a}, \underline{\Omega}_1, \dots) \frac{dx_1}{r_{1a}^2} &= \mu(r_1) e^{-\int_{r_{1a}}^{r_1} \mu(\zeta) d\zeta} dz_1 \frac{dx_1 dy_1}{r_{1a}^2} \\
&= \mu_1 e^{-\mu_1(r_1 - r_{1a})} dz_1 \frac{dx_1 dy_1}{r_{1a}^2} \\
&= \mu_1 e^{-\mu_1 l_1} dz_1 \frac{dx_1 dy_1}{r_{1a}^2} \\
&= \mu_1 e^{-\mu_1 l_1} \frac{dz_1}{\cos \phi_1} \frac{dx_1 dy_1 \cos \phi_1}{r_{1a}^2} \\
&= \mu_1 e^{-\mu_1 l_1} dl_1 \frac{dx_1 dy_1 \cos \phi_1}{r_{1a}^2}, \tag{3.46}
\end{aligned}$$

where μ_1 is the constant attenuation coefficient of the first detector at the initial particle energy e_0 . Noting that $d\underline{\Omega}_1 = \frac{dx_1 dy_1 \cos \phi_1}{r_{1a}^2}$, we can evaluate this integral in spherical coordinates along the radial l_1 direction with upper integration limit $l = \frac{t_z}{\cos \phi_1}$, where t_z is the thickness of the first detector, and then convert back to cartesian coordinates.

$$\begin{aligned}
\int_{S_{d1}} p_A(r_1 | r_{1a}, \dots) \frac{dx_1}{r_{1a}^2} &= \int_{S_{d1}} \int_{l_1=0}^{l_1=\frac{t_z}{\cos \phi_1}} \mu_1 e^{-\mu_1 l_1} dl_1 \frac{dx_1 dy_1 \cos \phi_1}{r_{1a}^2} \\
&= \int_{S_{d1}} (1 - e^{-\mu_1 \frac{t_z}{\cos \phi_1}}) \frac{dx_1 dy_1 \cos \phi_1}{r_{1a}^2}
\end{aligned}$$

where r_{1a} is the euclidean distance between the emission point \underline{x}_0 and the measured detection point \underline{x}_1 in the first detector. Assuming a thin first detector, $\frac{\mu_1 t_z}{\cos \phi_1} \ll 1$, we can linearize the exponential and use the approximation $1 - e^{-x} \approx x$,

$$\int_{S_{d1}} p_A(r_1|r_{1a}, \dots) \frac{dx_1}{r_{1a}^2} \approx \mu_1 t_z \int_{S_{d1}} \frac{dx_1 dy_1}{r_{1a}^2}, \quad (3.47)$$

where the $\cos \phi_1$ term cancels out.

A similar expression can be derived for the 2nd detector. By assuming a thick second detector that it is effectively an ideal attenuator at all scattered particle energies e_2 , the exponential term is approximately unity, $(1 - e^{-\mu_2 \frac{t_z}{\cos \phi_2}}) \approx 1$ and

$$\int_{S_{d2}} p_A(r_2|r_{2a}, \dots) \frac{dx_2}{r_{2a}^2} \approx \int_{S_{d2}} \frac{dx_2 dy_2 \cos \phi_2}{r_{2a}^2}, \quad (3.48)$$

which is simply an integral over all solid angle subtended by the 2nd detector surface (x_2, y_2) with respect to an initial scattering point in the first detector.

Note that there is an implicit assumption that the distance terms r_{1a} and r_{2a} do not vary over the (small) z-axis of the first detector and can be approximated by using the detector z-axis as an origin in order to avoid an extra integral in the sensitivity calculation. In addition, we are linearizing the depth dependency in the first detector and ignoring it in the second detector. This is justified by our assumption of higher sensitivity of the Compton scatter statistics to depth effects in the first detector than the absorption of photons in the secondary detector

3.8 Sensitivity and Measurement Probability

Combining these approximations together gives

$$P(D|\underline{x}_0) = \frac{1}{8\pi^2} \left(\frac{\mu_c}{\mu_t} \right) (\mu_t t_z) \int_{S_{d1}} \frac{1}{r_{1a}^2} \int_{S_{d2}} \frac{\cos \phi_2}{r_{2a}^2} \int_{S_{e1}} p(\theta, e_1|C_1, r_1, \dots) \left[\int_{S_{e2}} p(e_2|r_2, \underline{\Omega}_2, \dots) de_2 \right] de_1 dx_2 dy_2 dx_1 dy_2$$

By previous assumption, the energy window on the 2nd detector is sufficiently wide enough such that $\int_{S_{e2}} p(e_2|r_2, \underline{\Omega}_2, \dots) de_2 \approx 1$. This combined with the assumption of

negligible source attenuation when exiting the first detector and of ideal attenuation in the second detector results in no other dependency on e_1 besides $p(\theta, e_1|C_1, r_1, \dots)$. The integral over all $e_1 \in S_{e_1}$ can now be performed,

$$f(\theta) = \int_{S_{e_1}} p(\theta, e_1|C_1, r_1, \dots) de_1 \quad (3.49)$$

where $f(\theta)$ is the unnormalized marginal density of θ of all energy measurements e_1 within detection threshold S_{e_1} . Like before, $f(\theta)$ can be integrated ahead of time using numerical techniques and evaluated by interpolation.

The final result for the detection sensitivity is

$$P(D|\underline{x}_0) = \frac{1}{8\pi^2} \left(\frac{\mu_c}{\mu_t} \right) (\mu_t t_z) \int \frac{1}{r_{1a}^2} \int \frac{\cos \phi_2}{r_{2a}^2} f(\theta) dA_2 dA_1 \quad (3.50)$$

where the integral is over the surface areas A_1, A_2 of the first and second detectors. Note that the unnormalized marginal $f(\theta)$ can be equivalently expressed as $f(\cos \theta)$ since the scattering angle θ varies between $[0, \pi)$. This is a more efficient representation since $\cos \theta$ (along with $\cos \phi$) can be calculated directly from the incoming and outgoing γ -ray direction vectors $\underline{\Omega}'_1$ and $\underline{\Omega}'_2$, which avoids a computationally expensive calculation of a transcendental function.

The list-mode measurement probability (or backprojection) for some measurement vector $\underline{y} = (e_2, \underline{z}_2, e_1, \underline{z}_1)$ is solved for from the definition of the detection probability,

$$\int_S p(\underline{y}|D, \underline{x}_0) P(D|\underline{x}_0) d\underline{x}_0 \triangleq P(D|\underline{x}_0).$$

From the assumptions and simplifications outlined earlier, one can easily derive the following expression for the measurement probability by inspection:

$$p(\underline{y}|D, \underline{x}_0) P(D|\underline{x}_0) = \frac{1}{8\pi^2} \left(\frac{\mu_c}{\mu_t} \right) (\mu_t t_z) \left(\frac{\Delta z}{t_z} \right) \frac{1}{r_{1a}^2} \frac{\cos \phi_2}{r_{2a}^2} p(\theta, e_1) \quad (3.51)$$

where μ_t and (μ_c/μ_t) are the linear attenuation coefficient and Compton fraction for the first detector material at the particle initial energy. The Compton scattering jpdf $p(\theta, e_1)$ (3.8), the geometric terms $\cos \phi, r_{1a}, r_{2a}$, as well as the fractional depth of

interaction ($\Delta z/t_z$) in the first detector are calculated from the measured particle positions z_1 and z_2 and source pixel position x_0 . Note that the dependency on the second detector measured energy e_2 drops out since the event $E_2 = \{e_2 : e_2 \in S_{e_2}\}$ is assumed to occur with probability 1. Figure (3.9) shows a representative backprojection sliced by a 2D plane, with the ideal conic section superimposed.

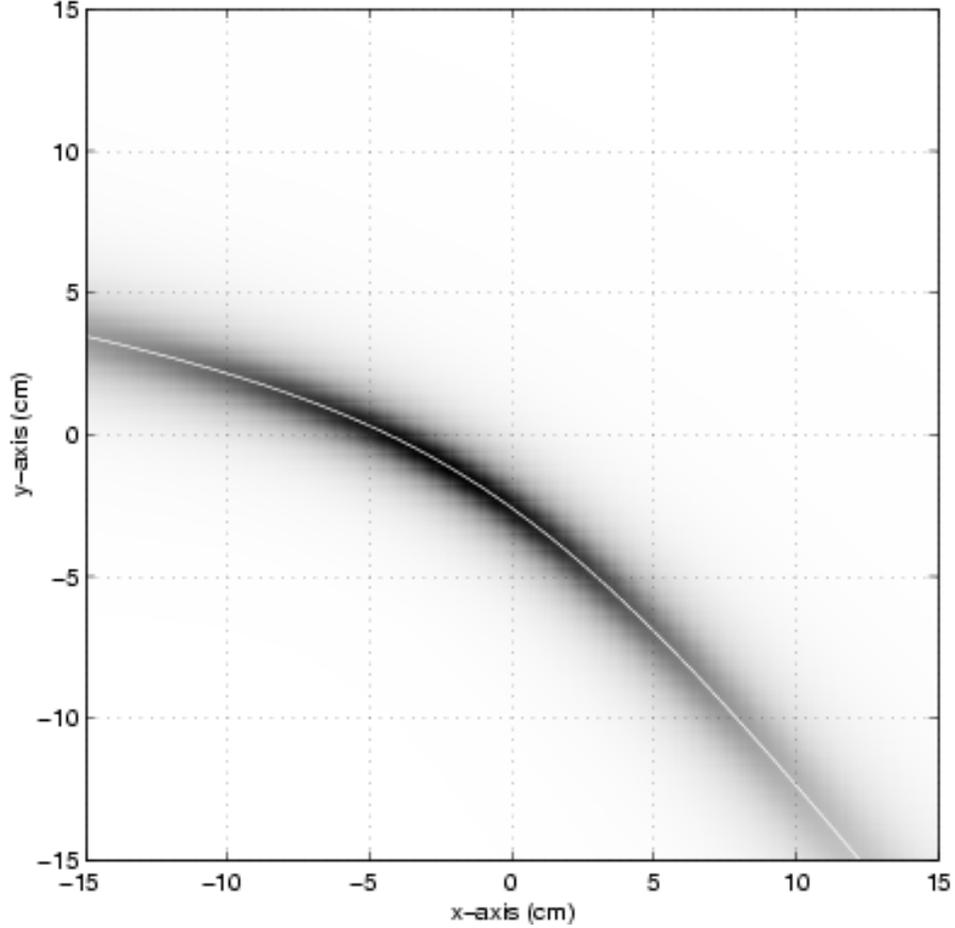


Figure 3.9: Representative Compton backprojection onto a plane (gray) with the ideal conic section superimposed (white).

In addition, the constant term $\frac{1}{8\pi^2} \left(\frac{\mu_c}{\mu_t} \right) (\mu_t t_z)$ can be left out of (3.50) and (3.51), since it will cancel out in the ML-EM update equations (2.25) or only contribute an additive constant to the likelihood function (2.12).

3.9 Reconstructed Images from Simulated 3D Data

In order to investigate the relative influence of various sources of error, a simulation study was performed where the various sources of error could be controlled. Compton SPECT list-mode coincidence data was generated by a Monte Carlo based simulation [107], which included the effects of Doppler broadening, collimator penetration, incomplete absorptions, detector measurement errors, geometric mis-alignments, and other phenomena. The simulated data was made up of six different cases. In the first case, there was no source of error or modeling uncertainty. For each subsequent case, additional sources of error were included in order to see how they would affect the reconstructed imagery. Table 3.1 gives an outline of the specific error sources as well as the reconstructed point-source full-width-half-maximum (FWHM). The simulated detector geometry was similar to that used in section 3.10. In addition, it included the effect of a lead collimator between the source and first detector in order to block any primary detections from occurring in the second detector.

Figures 3.10 - 3.21 show reconstructed γ -ray source distributions from a set of simulated emission sources. The simulated emission source distribution consisted of four identical point sources of ^{131}I (364keV). One point source was aligned 10.54cm in front of the first detector, with the remaining three point sources positioned 1.0cm away along different axis. The sources were rotated in 10 deg increments during data collection for a total of 360 deg, and the resulting reconstructed source intensity is fully tomographic. The reconstruction domain size was $4\text{cm} \times 2\text{cm} \times 4\text{cm}$ with $81 \times 41 \times 81$ pixels, for a sampling density of 0.05cm/sample. Over 50000 events were simulated, and all coincidence events are correctly matched (i.e. no incorrect coincidences). Only data with 1st-detector energy measurements between 20keV and 150keV were used, which corresponded to nominal scattering angles ranging from 23 deg to 90 deg.

Images were reconstructed using unregularized Maximum-Likelihood (ML) List-mode OSEM for 20 iterations, with 5-OSEM subsets per iteration, on 50000 counts.

Case	Error Sources	FWHM (cm)
1	None	0.062
2	Doppler broadening	0.091
3	Doppler broadening Collimator penetration	0.148
4	Doppler broadening Collimator penetration Error in 1st detector energy measurement (2keV)	0.225
5	Doppler broadening Collimator penetration Error in 1st detector energy measurement (2keV) Error in assumed tomographic rotation axis (1mm)	0.269
6	Doppler broadening Collimator penetration Error in 1st detector energy measurement (2keV) Error in assumed collimator position (1mm)	0.250

Table 3.1: Sources of error present in each of the simulated data reconstruction cases, along with the resulting FWHM of the center point source.

The backprojection values outside a cylindrical field-of-view centered about the tomographic (vertical) rotation axis as well as below a pre-determined threshold were set to 0 to enforce a 5%-sparsity in order to minimize disk storage requirements. The reconstructed images are horizontal planar slices at $y=0.0\text{cm}$, resulting in three point sources in the field-of-view.

3.10 Reconstructed Images from Measured 2D Data.

Figures 3.22 - 3.27 show reconstructed γ -ray source distributions for a ^{99}Tc (140keV) point-source and a disk extended-source. The sources were aligned vertically 11cm in front of the first detector, and data was collected separately for the two sources. For both sources, the rotation angle was kept fixed. Thus only 2D (i.e no tomographic) information was collected. The detection energy windows were 20keV–43keV on the first detector, and $\geq 114\text{keV}$ on the energy sum between the two detectors. Further details of the detector configuration can be found in [71].

A total of 10^5 detections were collected from the point-source, and 77590 detections were collected from the disk extended source. The two datasets were then combined and reconstructed as if both sources were present at the same time. List-mode ML-EM [10,84] was used for reconstruction with two additional approximations. First, to speed up the computation of the backprojections they were forced to a 10% sparsity by truncating small values to 0. Second, a 4-subset Ordered Subset EM (OS-EM) [62] routine was used to speed convergence. For the point-source alone, the observed FWHM was approximately 8.9mm, while for the combined point and disk extended source data the central peak FWHM was 11.2mm.

3.11 Discussion

From examination of figures 3.10 and 3.10, it is clear that when the model used for the backprojection calculation fits the data the resulting images are ideal. Indeed, the three point sources are almost perfect delta functions to within sampling error. Closer examination of figure 3.11 shows that the center point does not have the exact same amplitude as the point 1.0cm away, which could possible be due to statistical fluctuation in the estimate and/or not running the iterative algorithm to full convergence.

The addition of Doppler broadening (figures 3.12 - 3.13) appears to broaden the reconstructed point responses as well as introduce low-level noise. As more uncertainty and error source are introduced into the data, spike-like artifacts along the periphery of the reconstruction field-of-view begin to appear in the simulated 3D reconstructions (figures 3.14 - 3.21). A possible source of these artifacts could be *inconsistent data* in that the detected measurement was inconsistent with the valid field-of-view through the collimator. Thus the most “inconsistent” location for a possible source location would be as far away from the true source positions as allowed by reconstruction domain.

What is especially striking in the simulated data reconstructions is the sensitivity to model parameters such as tomographic center-of-rotation (figures 3.18, 3.19) and collimator position (figures 3.20, 3.21). A small shift on order of 1mm dramatically increases the artifacts, presumably due to inconsistent data, as compared to when the positions were modeled correctly (figures 3.10 - 3.15). The experimental Compton SPECT system used to collect the 2D data did not have any sort of precision alignment system, but instead relied upon the experimenter to align and measure by hand. Thus a tolerance on order of a few millimeters is within reason, which may go towards explaining the artifacts along the reconstruction domain periphery (figures 3.22 - 3.27).

There is one major discrepancy between the reconstructions of the measured and simulated data: over an order of magnitude difference in the point source resolution. Two possibilities come to mind to explain this. First, the “point” source used in the experiment was actually a finite-extent radioactive source placed in a small test tube. Second, the simulated data did not include the possibility of mis-matched coincidence detections.

A natural solution for dealing with inconsistent data would be regularization. Each spike-like artifacts only affect a pixel or two, a local smoothing penalty between neighboring pixels would help suppress these artifacts at the cost of smearing the true reconstructed point-source. However, exploring the benefits of different types of regularization was not pursued.

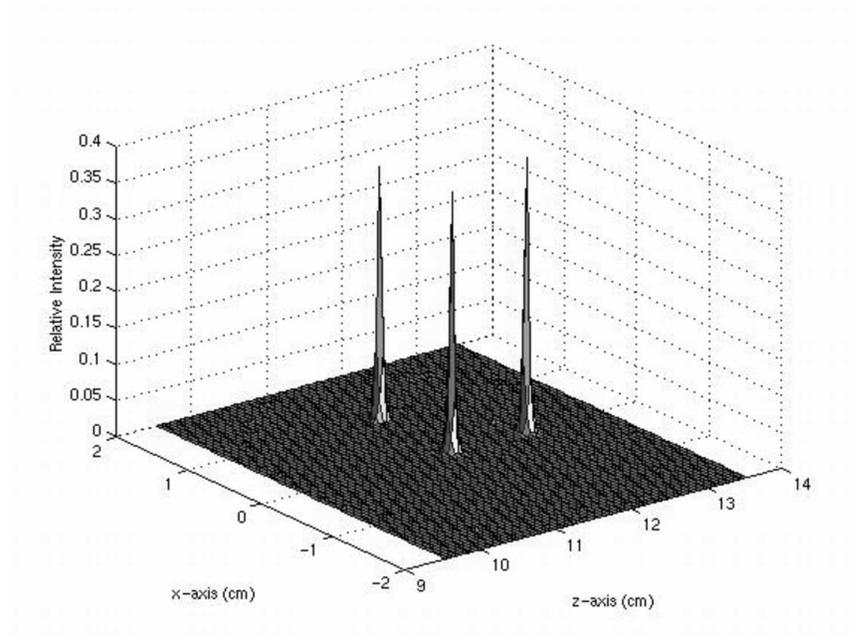


Figure 3.10: Reconstruction from simulated 3D data: No Doppler broadening, no collimator penetration.

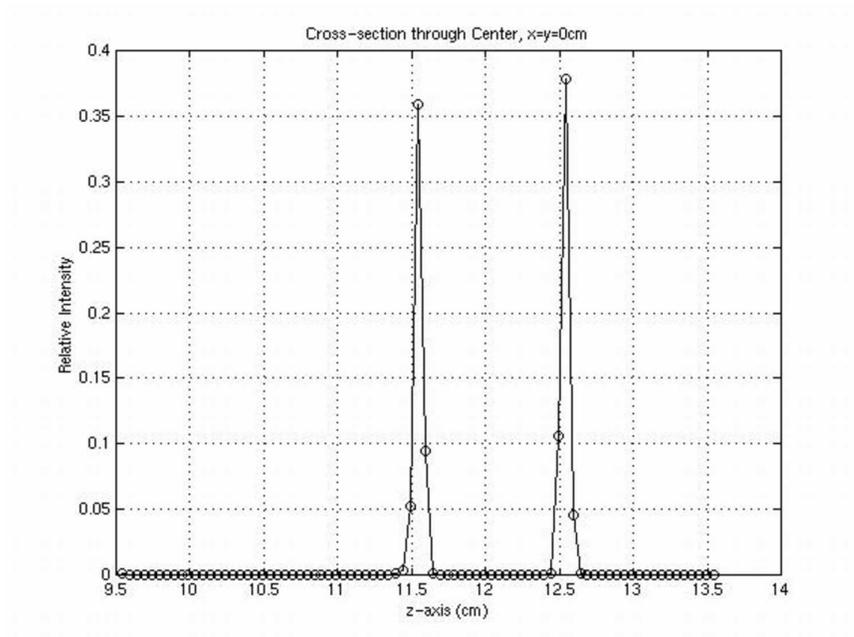


Figure 3.11: Cross-section through reconstruction from simulated 3D data: No Doppler broadening, no collimator penetration.

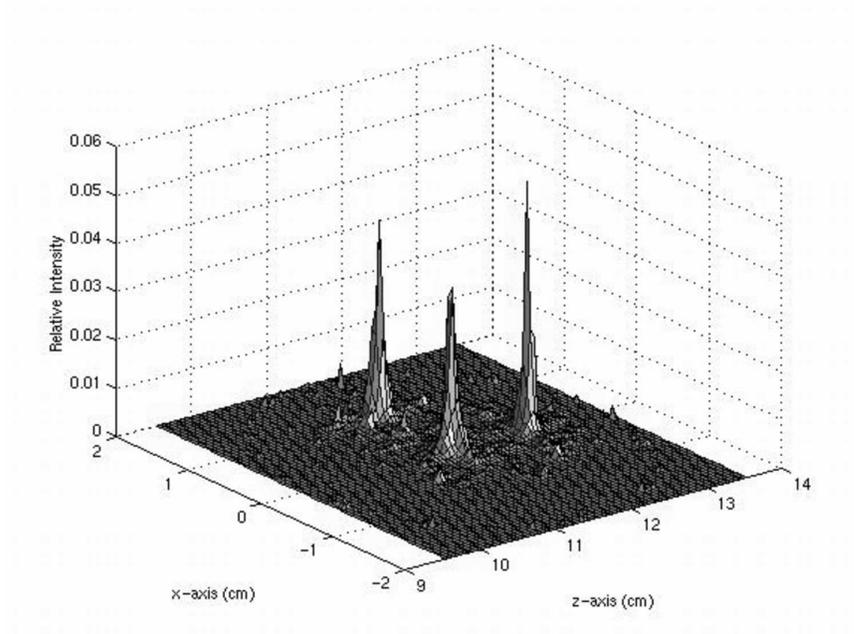


Figure 3.12: Reconstruction from simulated 3D data: Doppler broadening, no collimator penetration.

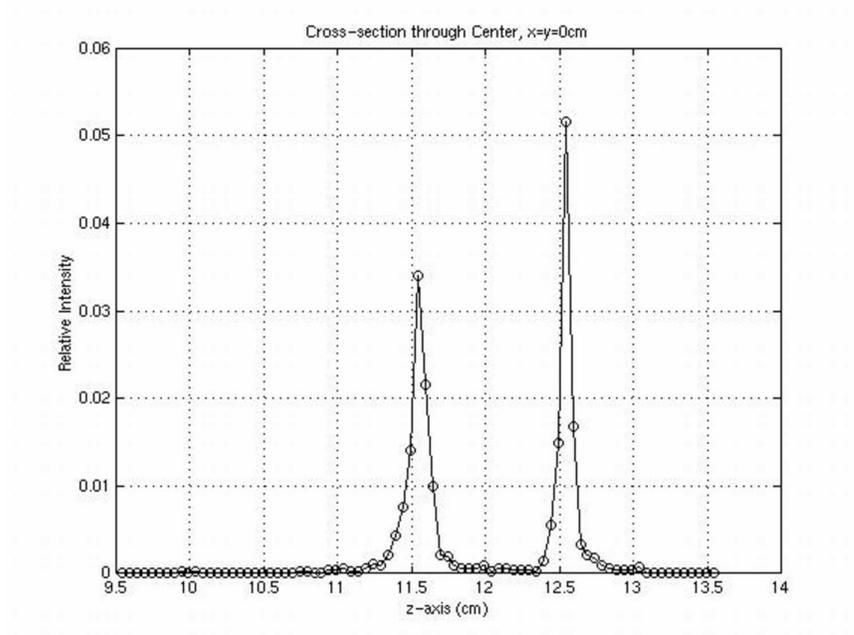


Figure 3.13: Cross-section through reconstruction from simulated 3D data: Doppler broadening, no collimator penetration.

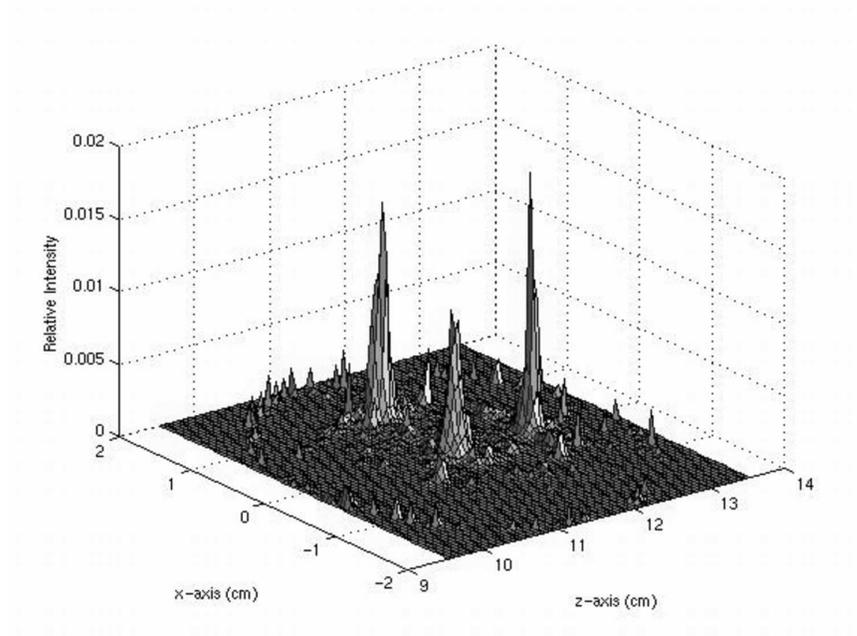


Figure 3.14: Reconstruction from simulated 3D data: Doppler broadening and collimator penetration.

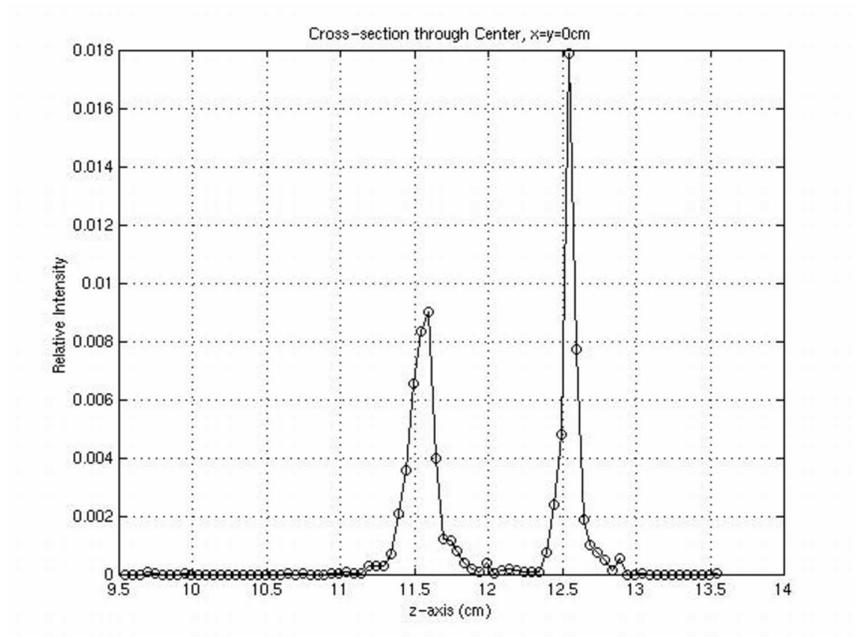


Figure 3.15: Cross-section through reconstruction from simulated 3D data: Doppler broadening and collimator penetration.

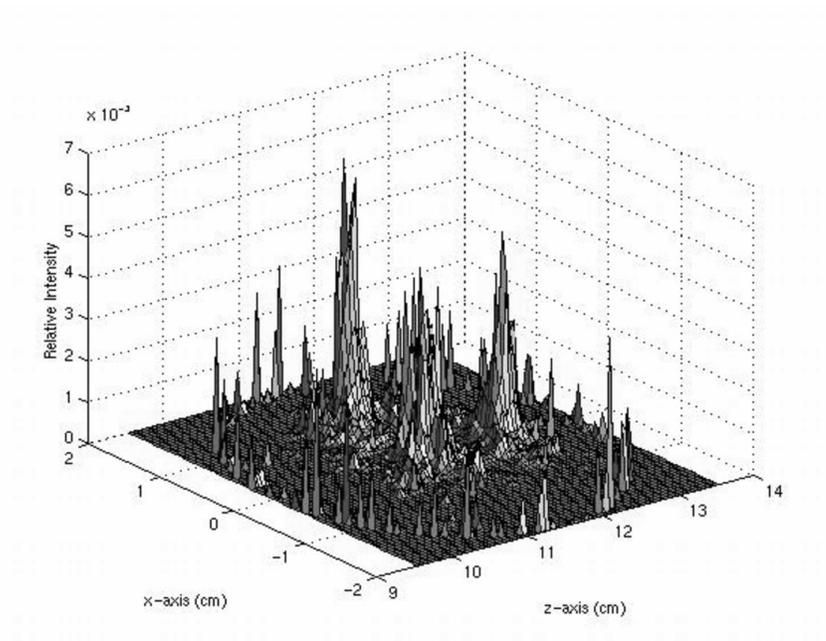


Figure 3.16: Reconstruction from simulated 3D data: Doppler broadening, collimator penetration, and 2keV FWHM first-detector energy measurement error.

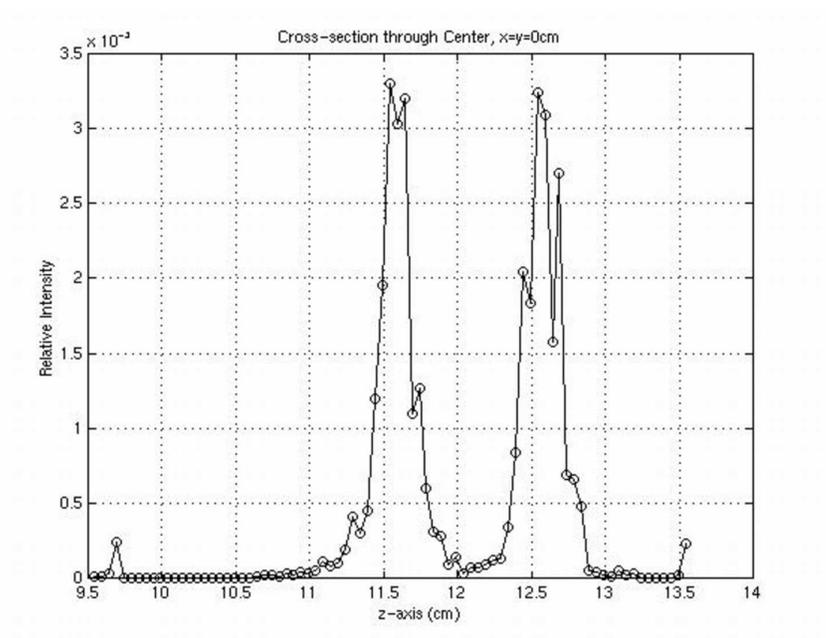


Figure 3.17: Cross-section through reconstruction from simulated 3D data: Doppler broadening, collimator penetration, and 2keV FWHM first-detector energy measurement error.

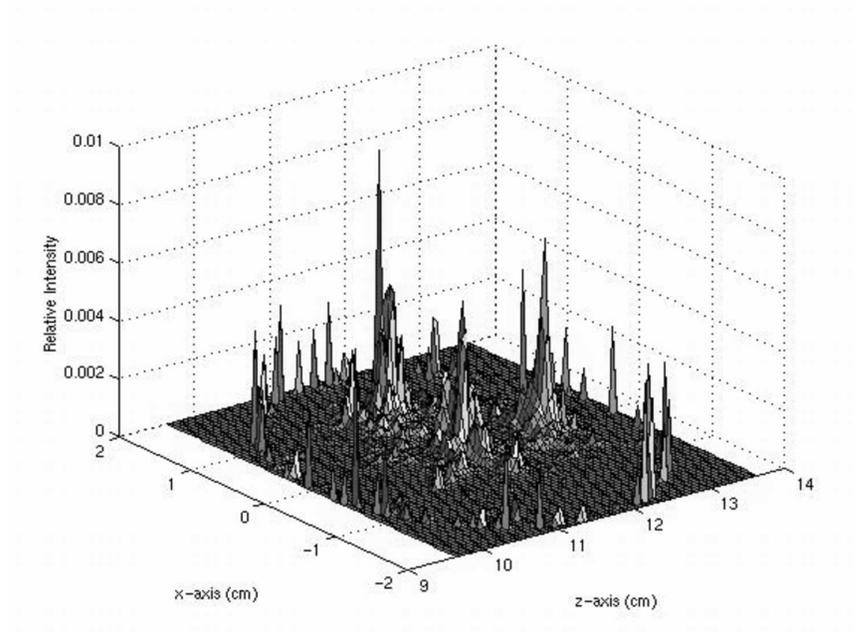


Figure 3.18: Reconstruction from simulated 3D data: Doppler broadening, collimator penetration, 2keV FWHM first-detector energy measurement error, and 1mm-offset in tomographic center of rotation.

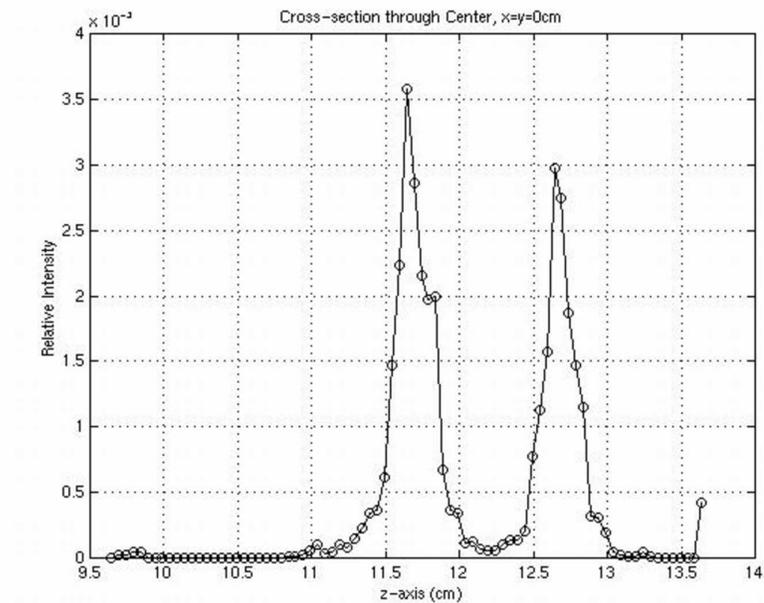


Figure 3.19: Cross-section through reconstruction from simulated 3D data: Doppler broadening, collimator penetration, 2keV FWHM first-detector energy measurement error, and 1mm-offset in tomographic center of rotation.

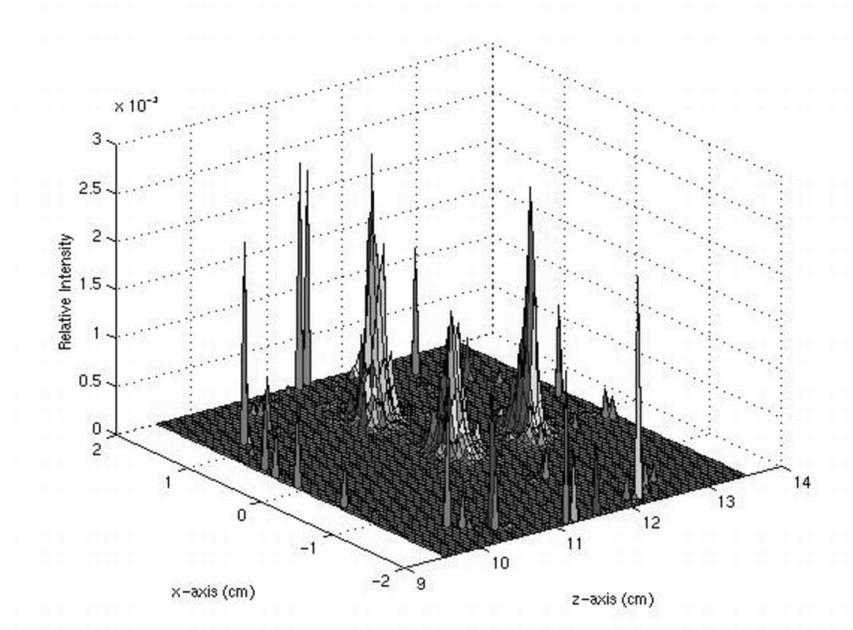


Figure 3.20: Reconstruction from simulated 3D data: Doppler broadening, collimator penetration, 2keV FWHM first-detector energy measurement error, and 1mm vertical offset in collimator slot position.

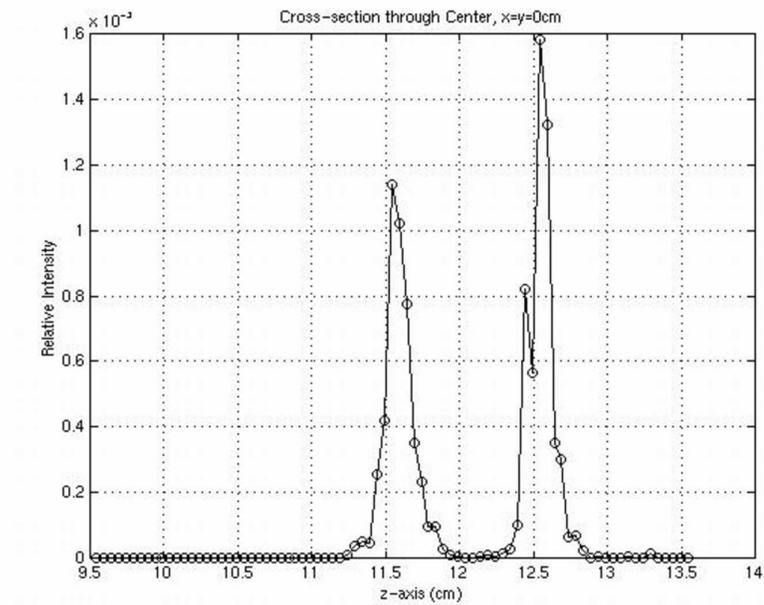


Figure 3.21: Cross-section through reconstruction from simulated 3D data: Doppler broadening, collimator penetration, 2keV FWHM first-detector energy measurement error, and 1mm vertical offset in collimator slot position.

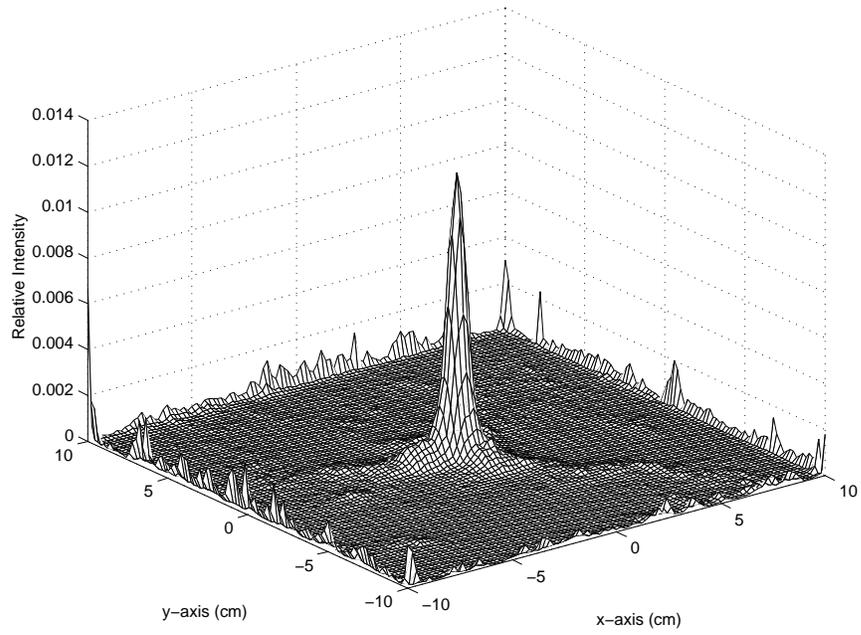


Figure 3.22: Reconstruction from measured 2D data of a point source.

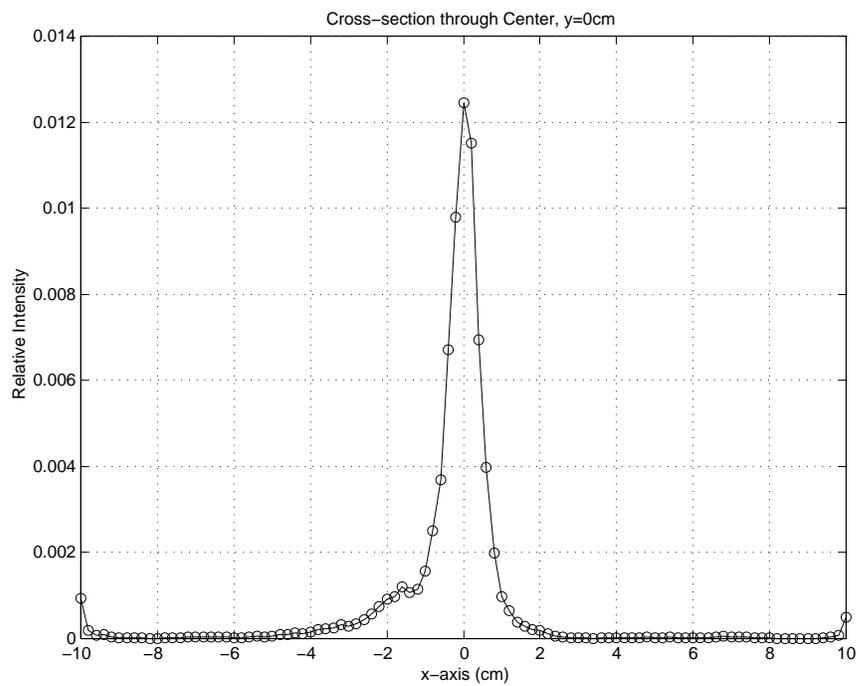


Figure 3.23: Cross-section through reconstruction from measured 2D data of a point source.

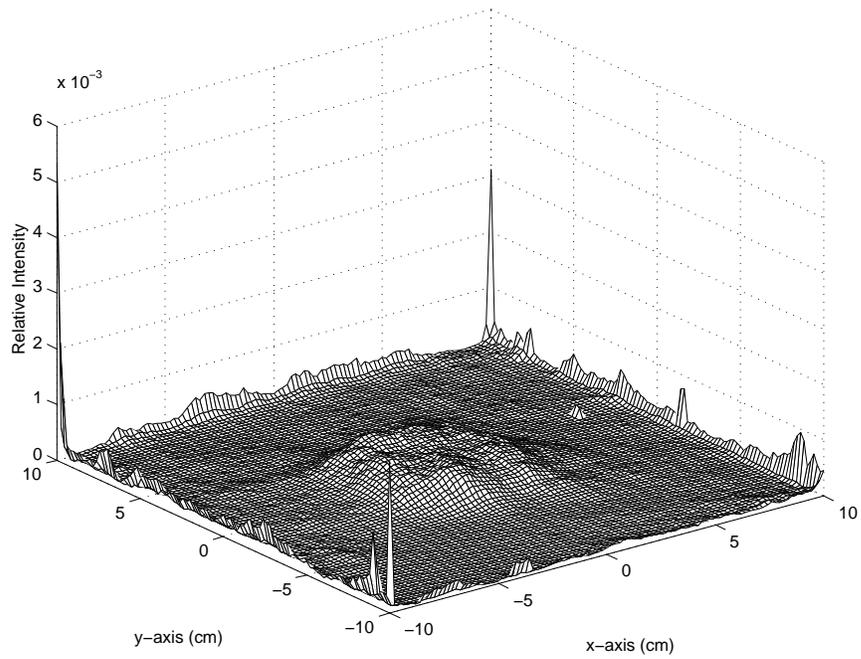


Figure 3.24: Reconstruction from measured 2D data of a disk-shaped extended source.

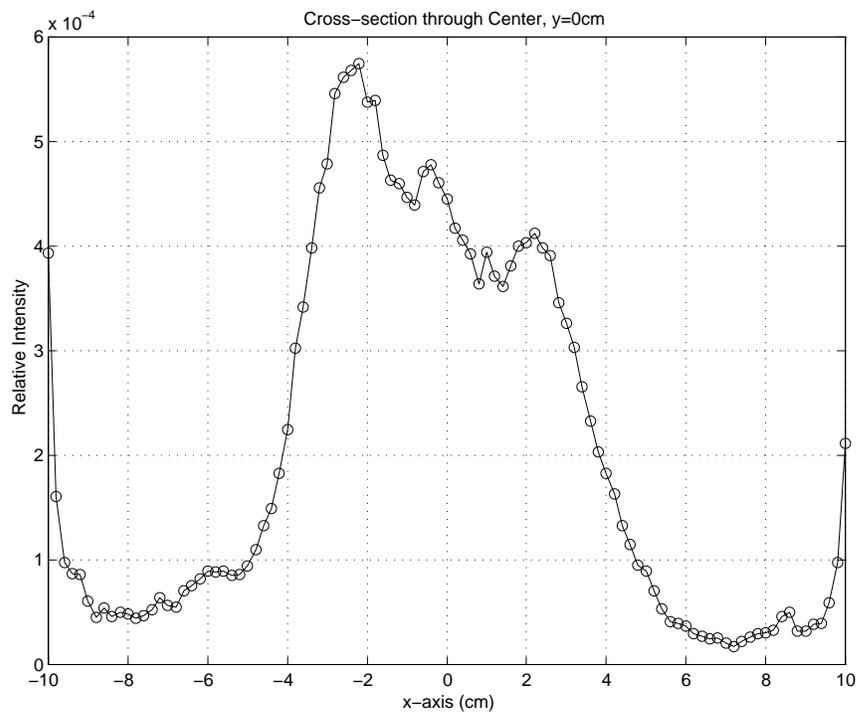


Figure 3.25: Cross-section through reconstruction from measured 2D data of a disk-shaped extended source.

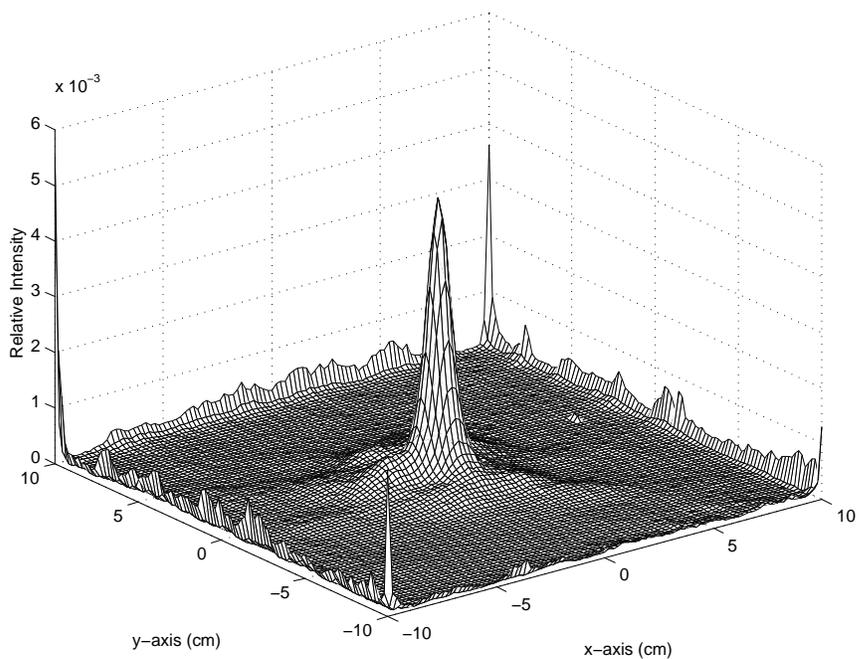


Figure 3.26: Reconstruction from combined measured 2D data of a point source and disk-shaped extended source.

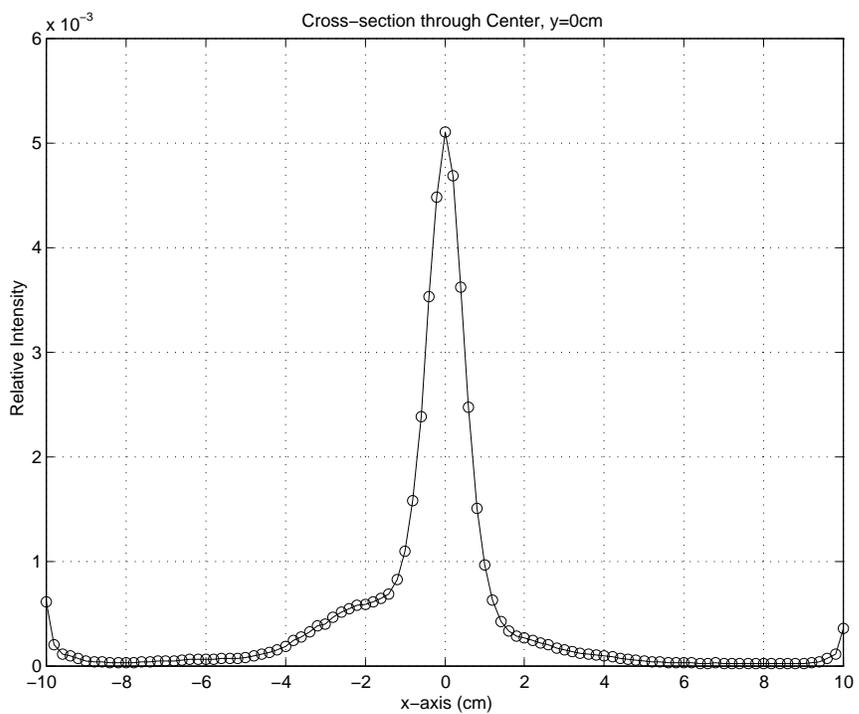


Figure 3.27: Cross-section through reconstruction from combined measured 2D data of a point source and disk-shaped extended source.

CHAPTER 4

Simplifying Tomographic Reconstruction with Optimized Binning

This chapter concerns acceleration of image reconstruction for emission tomographic modalities where equally-spaced measurement binning is impractical due to the very high dimensional detector measurement space. The techniques described herein are generally applicable to inverse problems with Poisson statistics.

Traditionally in emission tomographic modalities such as 2D SPECT and PET, the detected photon measurements are quantized to one of a finite number of possible measurement bins. This quantization is either done mechanically according to the geometry of the detectors (i.e. PET detector ring elements, SPECT detector arrays) or electronically in software (i.e. energy windows, binning of continuous position estimates). Typically, the measurements are binned uniformly with equal-sized bins along each measurement dimension. For example, a 2D PET sinogram will have bins spaced equally in angle and in radius, where the position of the bins are determined by the relative acceptance angles and locations of the finite-size detector element pairs in the PET ring. As pointed out in [10] and in previous chapters, when the dimensionality of the measurement space is greater than 4 or more, uniform quantization about each measurement axis becomes impractical. One wants fine binning along each axis in order to minimize measurement error due to quantization or to exploit the full resolution of the detectors. However, the total number of bins is geometrically

increasing with the number of dimensions. This makes storage impractical for large dimensional measurement spaces.

In list-mode reconstruction [10, 84], multiple attributes of each photon such as position, energy, scatter angle, etc., are recorded as a point in a multi-dimensional detector space. Each photon contributes individually to the likelihood function, and the underlying source distribution can be estimated using maximum-likelihood based estimation techniques. A failing of list-mode reconstruction is that it scales linearly with the number of counts. Since each individual detection contributes to the overall likelihood function, the computational complexity and storage requirements are proportional to the total number of counts. For example, given N source pixels to estimate and $M \gg N$ detections, list-mode ML-EM requires the calculation of an $M \times N$ system matrix requiring $B = o(\rho MN)$ bytes of random access memory and $o(MN^2)$ operations per iteration, assuming that the measurement back-projection have an average sparsity of $\rho < 1$. Alternatively, for a fixed amount B , ρ , and N , the maximum total number of measurement bins is $M = o(\frac{B}{\rho N})$.

For PET or mechanically collimated SPECT systems, the storage requirement is not so bad since the back-projection sparsity ρ is on the order of a few percent. However, for imaging modalities that have non-sparse back-projections (such as Compton SPECT), this is a major problem. For a “typical” 2-dimensional tomographic imaging problem involving $O(10^6)$ detections and $O(100^2)$ reconstruction pixels, the storage requirement for the system matrix can easily be in the gigabyte range and an order of magnitude larger for fully 3-dimensional tomographic imaging.

For Compton SPECT, several authors have proposed simplifying the system matrix by factoring [56], exploiting symmetries [91], or simplifying the forward- and back-projection operations [106]. Others have proposed variations on analytical cone-beam [39] reconstruction algorithms including inverse filtering in spherical coordinates [11, 83]. Although analytical reconstruction techniques promise fast reconstruc-

tion times, they suffer from the problems common to all analytical reconstruction techniques: they rely on specific geometric assumptions (such as hemispherical detectors), and they do not take into account the Poisson counting statistics and other noise sources. In other areas of emission tomography, authors have proposed different forward- and back-projection models in order to reduce the computational burden [64]. However, they suffer from noise artifacts precisely due to this mis-match [50,110].

As for statistically-based reconstruction, there has been much work in algorithms that quickly maximize the likelihood function in the first few iterations such as Ordered Subsets¹ EM (OSEM) based reconstruction techniques [34,62], block-iterative updates [20], and variations on algebraic reconstruction techniques [18,29] (of which OSEM is a special case). Although one still needs to pre-compute the system matrix or compute on-the-fly at each iteration, the reconstruction process can be sped up considerably by a block-by-block update of the estimated source distribution, resulting in an effective operation count of $O(\frac{M}{N_s}N^2)$ where N_s is the number of subsets or blocks used in the block-by-block update. However, convergence properties and image quality of the resulting estimates from some of these algorithms are still open to debate [99]. In either case, the scalability issue remains as high efficiency fully 3D modalities ($N \geq O(100^3)$) which collect massive amounts of counts ($M \geq O(10^6)$) are developed. Thus we are faced with the choice of either keeping all the measurements in their full resolution and accepting the storage requirements for the backprojections, or binning the data and accepting the loss in both measurement and reconstruction fidelity due to quantization.

As a study of how quantization affect reconstruction image quality, we investigated how various PET sinogram sampling rates would change reconstructed image quality. We simulated 10^6 coincidence detections for a 2D PET system, and reconstructed the emission source density using List-mode MLEM. As a comparison, we binned

¹Note that the idea of splitting the data into subsets appeared a few years earlier in [56].

the data at various levels and reconstructed using the binned data. The resulting emission density estimates were compared against the known true density and the mean-squared error was calculated. The result was that after a certain sampling density, there was no increase in image quality as measured by the reconstructed image MSE, as shown in figure 4.1. From this we see that there is a tradeoff between

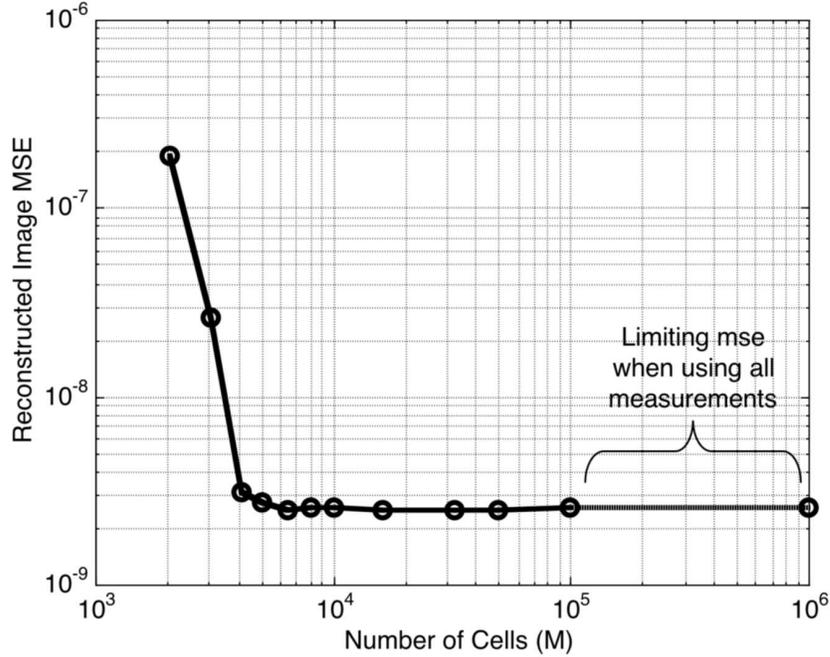


Figure 4.1: Reconstructed image MSE vs. number of quantizer cells for simulated 2D PET data. The limiting case for 10^6 cells corresponds to list-mode reconstruction.

the level of measurement quantization and the quality of the reconstructed image. If we can analytically predict this tradeoff, then we could choose the optimal binning level for reconstruction for a desired image quality. Since the reconstruction would be on binned data, there could be a dramatic decrease in the computation involved as compared to a list-mode reconstruction.

A concept well known in the fields of information theory and communications is Vector Quantization (VQ) [49,52]. Vector quantization is a multi-variable generalization of simple binning. When the measurement bins are equally spaced about each dimension, the intersection of the bins form equally-sized squares, cubical volumes,

or hyper-cubical quantization regions. This concept can be generalized by defining the quantization regions by the intersection of multiple hyperplanes at arbitrary (rather than orthogonal) orientations, resulting in non-uniform, variable size convex polytopes which partition each quantization region in the measurement space. For a k -dimensional measurement space with M -bins, rather than distributing $O(M^{1/k})$ bins uniformly along each dimension, one can conceivably place more bins in regions of the detector space that are “more important” in some sense using a vector quantizer. In the communications and source coding literature, this is referred to as a *bit allocation* problem [3, 21, 22, 51, 89], in that one has $\log_2(M)$ bits to allocate to various dimensions. Allocation has a measurable effect on the quality of digitally encoded measurements, and one would like to do so in an optimal manner. Vector quantization has appeared in the medical imaging literature in the context of image compression and classification [36, 37, 73] and the associated image quality loss [27] has been studied. However, to our knowledge, little has been published on the design of VQ’s specifically for image reconstruction tasks.

Just as in traditional uniform binning, the total number of detections in each partition under VQ binning is a Poisson random variable, and the underlying source distribution can be estimated using ML techniques. This concept applies to both electronic collimation, where the binning can be done in software, along with finite-size detectors that have already binned their measurements during data acquisition. Indeed, the detections from multiple detectors can be summed into a single “meta-bin”.

Since quantizing the continuous list-mode measurements into a finite number of discrete bins involves information loss and a resulting reduction in image quality, one obviously would want to find a quantizer that minimizes this loss. The mean squared-error (MSE) distortion $d(\underline{y}, Q(\underline{y})) = E[\|\underline{y} - Q(\underline{y})\|^2]$ is the most commonly used distortion measure for assessing the average effect of quantization. However, in

many applications minimizing the squared-error distortion is far from optimal. For example, in speech coding based on quantized linear-prediction coefficients, a popular distortion measure is the log-spectral distortion [47]. In imaging, there are various perceptual quality measures of interest besides squared-error [35].

In nuclear medical imaging, the difference between a radiotracer image of a patient before and after the emergence of a tumor would manifest itself as a subtle variation of a few pixels against an similar background image. Thus one is concerned not with dramatic changes between two images, but rather with sensitivity to detecting small changes in an initial image. Under this condition the fidelity of the image is not of direct interest, but rather how image fidelity affects the performance of the detection task [6–8, 79]. An appropriate distortion function in this case would be the Kullback-Liebler (KL) discrimination [15, 28] as well as the loss in performance as measured by the area under the Receiver Operating Characteristic (ROC) curve.

Previous authors have worked the problem of the loss in discrimination and generalized Ali-Silvey distances due to quantization [12, 85, 86], but only for the one-dimensional (scalar) quantizers. For the multi-dimensional case, when the (non-difference) distortion measure is a continuous non-negative locally-quadratic function, the distortion of a high-rate vector quantizer (equivalently, a large number of partitions) can be approximated by a weighted Euclidean norm of the true value \underline{y} and quantized value $Q(\underline{y})$ [74, 76, 77],

$$d(\underline{y}, Q(\underline{y})) \approx (\underline{y} - Q(\underline{y}))^T \mathbf{B}(\underline{y})(\underline{y} - Q(\underline{y}))$$

where $\mathbf{B}(\underline{y})$ is a $k \times k$ -dimensional matrix proportional to the Hessian of the distortion,

$$\mathbf{B}(\underline{y}) = \frac{1}{2} \nabla^2 d(\underline{y}, Q(\underline{y})) \Big|_{\underline{y}=\underline{y}_i}$$

which can be a function of the input vector \underline{y} .

Gupta [53, 54] solved for discrimination and ROC-curve optimal vector quantizers by introducing the concept of a *covariation profile* matrix which is related to the

normalized moment of inertia traditionally used in asymptotic quantization theory. In this chapter, we will apply and extend Gupta’s work. We will present an optimal vector quantizer that minimizes the reduction in the Kullback-Liebler discrimination of the source due to quantizing its projections. This quantizer is dependent on the source only through its projections, and can be quickly designed directly from the observed list-mode measurements. The derivation of a discrimination-optimal quantizer will follow the method outlined in [54]. In addition, we will derive optimal quantizers with a variable point-density along with a fast and more efficient but sub-optimal lattice-based quantizer.

4.1 Vector Quantization Background

An M -point, k -dimensional vector quantizer (VQ) [80] consists of a partition $S = \{S_1, \dots, S_M\}$ of \mathbb{R}^k into M unique quantization cells, and a codebook $C = \{\underline{y}_1, \dots, \underline{y}_M\}$ consisting of M quantizer values² in \mathbb{R}^k . A k -dimensional vector \underline{y} with probability density $p(\underline{y})$ is quantized to one of M codebook values by a quantization rule $Q(\underline{y})$,

$$Q(\underline{y}) = \left\{ \underline{y}_l : \underline{y} \in S_l, l = 1, \dots, M \right\} \quad (4.1)$$

A block diagram of the VQ operator is given in Figure 4.2.

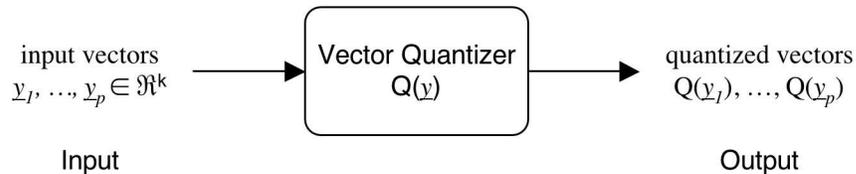


Figure 4.2: Block diagram of a vector quantizer.

The quantization rule is typically designed to minimize a function that measures the error or *distortion* between \underline{y} and $Q(\underline{y})$. An example of a one-dimensional quan-

²The quantized values are also referred to as codevectors or reproduction vectors, which we will use interchangeably.

tizer illustrating the codevector locations and quantization regions is shown in figure 4.3. Some example two-dimensional quantizers are shown in figure 4.4.

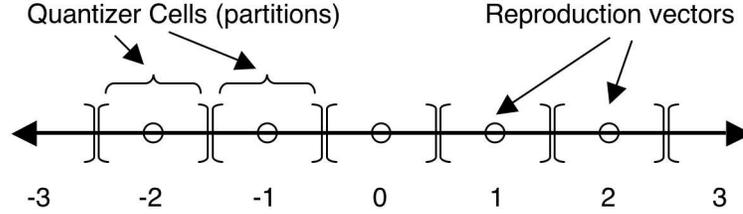


Figure 4.3: Example quantizer in \mathbb{R}^1 , detailing the codevector locations and quantization regions.

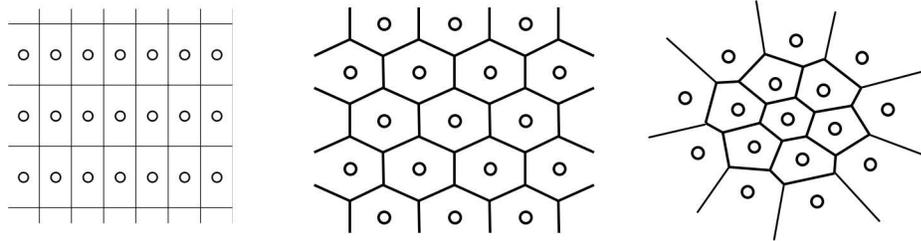


Figure 4.4: Example quantizers in \mathbb{R}^2 . The quantizers on the left and middle are lattice quantizers and have a uniform point density, while the quantizer on the right has a non-uniform point density.

4.2 MSE-Optimal Vector Quantizer

As mentioned above, the mean squared-error (MSE) $d(\underline{y}, Q(\underline{y})) = E[\|\underline{y} - Q(\underline{y})\|^2]$ is the most commonly used distortion measure for assessing the average effect of quantization. The mean-squared error normalized by the dimension k is given by

$$\begin{aligned}
 d(\underline{y}, Q(\underline{y})) &= \frac{1}{k} \int_S p(\underline{y}) \|\underline{y} - Q(\underline{y})\|^2 d\underline{y} \\
 &= \frac{1}{k} \sum_{l=1}^M \int_{S_l} p(\underline{y}) \|\underline{y} - \underline{y}_l\|^2 d\underline{y}
 \end{aligned} \tag{4.2}$$

When the number of cells M is large and the corresponding quantization cells are small, and assuming that the density $p(\underline{y})$ is smooth over the extent of each cell, then

the average distortion in (4.2) may be approximated by

$$\begin{aligned} d(\underline{y}, Q(\underline{y})) &\cong \frac{1}{k} \sum_{l=1}^M p(\underline{y}_l) \int_{S_l} \|\underline{y} - \underline{y}_l\|^2 d\underline{y} \\ &= \sum_{l=1}^M p(\underline{y}_l) m(S_l) V(S_l)^{1+2/k}, \end{aligned} \quad (4.3)$$

where “ \cong ” means to within an error of $o(\|\underline{y} - \underline{y}_l\|^2)$, $V(S_l) = \int_{S_l} d\underline{y}$ is the volume of the cell S_l , and $m(S_l)$ is the *normalized moment of inertia* of the cell.

4.2.1 Normalized Moment of Inertia

Following the example of [52], we define the normalized moment of inertia of the cell S_l about the point \underline{y}_l with respect to squared-error distortion as

$$m(S_l) = \frac{\int_{S_l} \|\underline{y} - \underline{y}_l\|^2 d\underline{y}}{kV(S_l)^{1+2/k}}. \quad (4.4)$$

The normalized moment of inertia (NMI) is proportional to the mean-squared quantization error in the cell S_l . The value of $m(S_l)$ depends on the shape of the cell and is invariant to its orientation or scale. For example, given a scaled version of a cell $\alpha S = \{\underline{u} : \exists \underline{y} \in S \text{ s.t. } \underline{u} = \alpha \underline{y}, \underline{u} \in \mathbb{R}^k\}$, it is trivial to show that $m(\alpha S) = m(S)$. With this definition the interval ($k = 1$), square ($k = 2$), and cube ($k = 3$) all have the same NMI of $m = 0.08\overline{33}$. In dimension $k = 2$, the cell shape with the smallest NMI that can partition \mathbb{R}^2 is the hexagon with $m_2^* \approx 0.080188$. In higher dimensions, the values of the smallest possible NMI is not known, but lower bounds exists [26]. Figure 4.5 gives some examples of quantizer cells in \mathbb{R}^2 along with their associated NMI. The key feature is that quantizer cells that are “rounder” have a smaller NMI.

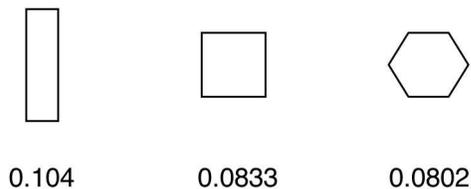


Figure 4.5: Example quantizer cells in \mathbb{R}^2 , along with their associated NMI.

When the components of \underline{y} have different dimensions or units, one may want to weight the MSE calculation to normalize or whiten \underline{y} . Thus, similar to the definitions given in [74], we will also define a weighted normalized moment of inertia of the k -dimensional quantization cell S_l with respect to a $k \times k$ -dimensional positive-definite symmetric weighting matrix \mathbf{B} as

$$m_{\mathbf{B}}(S_l) = \frac{\int_{S_l} (\underline{y} - \underline{y}_l)^T \mathbf{B} (\underline{y} - \underline{y}_l) d\underline{y}}{kV(S_l)^{1+2/k}}. \quad (4.5)$$

The relationship between (4.4) and (4.5) can be shown by making the following change of variables. Let $\underline{u} = \mathbf{B}^{-\frac{1}{2}}\underline{y}$. Define the transformation of the region S by the matrix \mathbf{B} as

$$\mathbf{B}[S] = \{\underline{u} : \exists \underline{y} \in S \text{ s.t. } \underline{u} = \mathbf{B}\underline{y}, \underline{u} \in \mathbb{R}^k\}.$$

Then

$$\begin{aligned} m_{\mathbf{B}}(S) &= \frac{\int_S (\underline{y} - Q(\underline{y}))^T \mathbf{B} (\underline{y} - Q(\underline{y})) d\underline{y}}{kV(S)^{1+2/k}} \\ &= |\mathbf{B}|^{\frac{1}{k}} \frac{\int_H (\underline{u} - Q(\underline{u}))^T (\underline{u} - Q(\underline{u})) d\underline{u}}{kV(H)^{1+2/k}} \\ &= |\mathbf{B}|^{\frac{1}{k}} m(H) \end{aligned} \quad (4.6)$$

where $H = \mathbf{B}^{\frac{1}{2}}[S]$.

Note that the matrix \mathbf{B} can be a function of the input vector \underline{y} . This results in quantizer cells whose orientation varies from location to location and aligned with the eigenvectors of $\mathbf{B}(\underline{y})$ (figure 3.6, pp. 74 of [54]). For the purposes of this dissertation we will restrict ourselves to the case where \mathbf{B} is a constant positive-definite matrix.

4.2.2 Covariation Profile

The normalized moment of inertia is invariant to cell orientation, and thus does not give any information about the alignment of the cell axis. Gupta [54] defined the normalized covariation profile of the cell to capture this orientation information.

Although the covariation profile does not appear in the asymptotic analysis for the quantizer MSE (section 4.2.3), it will be used extensively in the analysis of other distortion measures.

The normalized covariation profile $\mathbf{M}(S_l)$ of the cell S_l about the point \underline{y}_l is a $k \times k$ -dimensional matrix function,

$$\mathbf{M}(S_l) = \frac{\int_{S_l} (\underline{y} - \underline{y}_l)(\underline{y} - \underline{y}_l)^T d\underline{y}}{kV(S_l)^{1+2/k}}. \quad (4.7)$$

Whereas the normalized moment of inertia $m(S)$ is related to the mean-square error due to quantization, the normalized covariation profile $\mathbf{M}(S)$ is related to the quantization error autocorrelation matrix [109].

The normalized covariation profile with respect to a positive-definite symmetric weighting matrix \mathbf{B} is given by

$$\mathbf{M}_{\mathbf{B}}(S) = \frac{\mathbf{B}^{\frac{1}{2}} \int_S (\underline{y} - Q(\underline{y}))(\underline{y} - Q(\underline{y}))^T d\underline{y} \mathbf{B}^{\frac{1}{2}}}{kV(S)^{1+2/k}} \quad (4.8)$$

Similar to the normalized moment of inertia, the weighted and un-weighted normalized covariation profile can be related by the change of variables $\underline{u} = \mathbf{B}^{-\frac{1}{2}}\underline{y}$,

$$\begin{aligned} \mathbf{M}_{\mathbf{B}}(S) &= \frac{\mathbf{B}^{\frac{1}{2}} \int_S (\underline{y} - Q(\underline{y}))(\underline{y} - Q(\underline{y}))^T d\underline{y} \mathbf{B}^{\frac{1}{2}}}{kV(S)^{1+2/k}} \\ &= |\mathbf{B}|^{\frac{1}{k}} \frac{\int_H (\underline{u} - Q(\underline{u}))(\underline{u} - Q(\underline{u}))^T d\underline{u}}{kV(H)^{1+2/k}} \\ &= |\mathbf{B}|^{\frac{1}{k}} \mathbf{M}(H). \end{aligned} \quad (4.9)$$

As with the NMI, the normalized covariation profile $\mathbf{M}_{\mathbf{B}}(S)$ is invariant to scaling of the cell partition S and is related to the normalized moment of inertia via

$$\begin{aligned} m(S) &= \text{tr}(\mathbf{M}(S)) \\ &\geq k|\mathbf{M}(S)|^{1/k}. \end{aligned} \quad (4.10)$$

When the cell is symmetric the quantization error components are uncorrelated [109], resulting in covariation profile that is a scalar multiple of the identity matrix and equality in (4.10).

4.2.3 Asymptotic MSE

To explore the asymptotic distortion as the number of cells M grows large, we will define the *point density* $\zeta(\underline{y})$ of quantization points as

$$\zeta(\underline{y}) \cong \frac{1}{MV(S_l)} \text{ if } \underline{y} \in S_l, \quad (4.11)$$

where the fraction of quantization points in a small region A is approximately $\zeta(\underline{y})V(A)$. Lastly, define the continuous inertial profile function $m(\underline{y})$ as the normalized moment of inertia of cells in the neighborhood of \underline{y} . Combining this definition and (4.11) with (4.3) results in

$$d(\underline{y}, Q(\underline{y})) \cong \frac{1}{M^{2/k}} \sum_{l=1}^M p(\underline{y}_l) \frac{m(\underline{y}_l)}{\zeta(\underline{y}_l)^{2/k}} V(S_l). \quad (4.12)$$

Since (4.12) is an approximation of a Riemann integral, by using a few continuity and convergence assumptions [80] we can show that

$$\lim_{M \rightarrow \infty} M^{2/k} d(\underline{y}, Q(\underline{y})) = \int_S p(\underline{y}) \frac{m(\underline{y})}{\zeta(\underline{y})^{2/k}} d\underline{y}. \quad (4.13)$$

Equation (4.13) is known as *Bennett's Integral* and shows that asymptotically, distortion falls off as a rate of $M^{-2/k}$ with a constant term that depends on the distribution $p(\underline{y})$, the local quantizer cell NMI $m(\underline{y})$, and quantizer cell density $\zeta(\underline{y})$.

By Hölder's inequality, one can solve for the optimal point density function $\zeta^*(\underline{y})$ that minimizes the squared-error distortion in (4.13)

$$\zeta^*(\underline{y}) = \frac{[p(\underline{y})m(\underline{y})]^{k/2}}{\int_S [p(\underline{y})m(\underline{y})]^{k/2} d\underline{y}}, \quad (4.14)$$

and substituting (4.14) into (4.13) gives an expression for the distortion of an M -point quantizer in \mathbb{R}^k with the optimal cell density $\zeta^*(\underline{y})$,

$$d(\underline{y}, Q(\underline{y})) \cong \frac{1}{M^{2/k}} \left[\int_S [p(\underline{y})m(\underline{y})]^{k/2} d\underline{y} \right]^{k/2}. \quad (4.15)$$

As for the normalized moment of inertia $m(\underline{y})$, according to Gersho's conjecture [48] for large values of M the cells of an optimal k -dimensional vector quantizers are

congruent to the minimum NMI tessellating polytope in \mathbb{R}^k . The definition of NMI is scale invariant resulting in a constant $m(\underline{y}) = m_k^*$ for the optimal NMI in dimension k . Substituting m_k^* into (4.15) gives an expression for the least-possible mean-squared error distortion of any M -point quantizer in \mathbb{R}^k ,

$$d(\underline{y}, Q(\underline{y})) \cong \frac{m_k^*}{M^{2/k}} \left[\int_S p(\underline{y})^{\frac{k}{k+2}} d\underline{y} \right]^{\frac{k+2}{k}}. \quad (4.16)$$

4.3 A-optimal Vector Quantizer

For image reconstruction problems, the parameters of interest $\boldsymbol{\theta}$ are indirectly related to the measurements \underline{y} through a statistical measurement equation. MSE-optimal quantization of the measurements \underline{y} is not directly useful since it says nothing about the quality of the estimates of $\boldsymbol{\theta}$ based on quantized measurements $Q(\underline{y})$. However, the variance of the estimate of $\boldsymbol{\theta}$ is directly related to the quantized measurement Fisher information (FIM) (2.21) through the Cramèr-Rao Bound. As mentioned before in section 2.7, for Poisson measurements the difference between the continuous measurement FIM \mathbf{F} (2.14) and quantized measurement FIM $\hat{\mathbf{F}}$ (2.21) will be a function of the quantizer partitioning $\{S_1, \dots, S_M\}$.

In optimal experiment design [38,66], when the Fisher information is a function of some parameters, the optimal choice of those parameters would be the ones that maximize some function of the Fisher information. Examples of these include maximizing the determinate (D-optimal design), and maximizing the trace (A-optimal design). An example of optimal design for sensor placement in a manufacturing setting is given in [105].

For our purposes, the Fisher information parameter of interest is the quantizer partitioning $\{S_1, \dots, S_M\}$. Since the quantity we want to estimate $\boldsymbol{\theta} \in \mathbb{R}^N$, the Fisher information $\hat{\mathbf{F}} \in \mathbb{R}^{N \times N}$ where N is the number of pixel in the reconstructed image. Thus, evaluating the determinate $|\hat{\mathbf{F}}|$ is intractable. However, under the condition of a large number of partitions (equivalently, high-rate VQ), we can derive

a tractable expression for the asymptotic *difference* in the trace of the continuous and quantized measurement FIM. Optimizing the partitioning of S in order to minimize the difference in the trace results in an A-optimal quantizer, in that it minimizes the loss in the FIM due to quantization.

4.3.1 Loss in the Fisher Information trace

For any two non-negative functions f, g , Schwarz's inequality [1] states that

$$\int_{S_i} f(\underline{y})^2 d\underline{y} \int_{S_i} g(\underline{y})^2 d\underline{y} \geq \left[\int_{S_i} f(\underline{y})g(\underline{y}) d\underline{y} \right]^2.$$

If we define $f(\underline{y}) = \frac{a_j(\underline{y})}{\mu(\underline{y})^{\frac{1}{2}}}$ and $g(\underline{y}) = \mu(\underline{y})^{\frac{1}{2}}$, then Schwarz's inequality and a little algebra results in

$$\int_{S_i} \frac{a_j^2(\underline{y})}{\mu(\underline{y})} d\underline{y} \geq \frac{a_{lj}^2}{\mu_i}. \quad (4.17)$$

Evaluating the difference in the trace Δtr of Fisher information matrices \mathbf{F} and $\hat{\mathbf{F}}$ and incorporating (4.17) results in

$$\begin{aligned} \Delta \text{tr} &= \text{trace}(\mathbf{F}) - \text{trace}(\hat{\mathbf{F}}) \\ &= T \sum_{j=1}^N \sum_{l=1}^M \left[\int_{S_l} \frac{a_j^2(\underline{y})}{\mu(\underline{y})} d\underline{y} - \frac{a_{lj}^2}{\mu_l} \right] \\ &= \geq 0. \end{aligned} \quad (4.18)$$

Thus $\Delta \text{tr} \geq 0$ for any measurement partitioning $\{S_1, \dots, S_M\}$ of S . Therefore the difference in trace between the continuous measurement FIM (2.14) and quantized measurement FIM (2.21) corresponds to an overall loss in A-optimality due to measurement quantization.

When the number of cells M is large and the corresponding quantization cells are small, and assuming that the density $p(\underline{y})$ is smooth over the extent of each cell, the difference term in (4.18) can be expanded out in a Taylor series expansion about \underline{y}_l up to terms of order $o(\|\underline{y} - \underline{y}_l\|^2)$,

$$\begin{aligned}
& \int_{S_l} \frac{a_j^2(\underline{y})}{\mu(\underline{y})} d\underline{y} - \frac{a_{lj}^2}{\mu_l} \\
& \cong p(\underline{y}_l) \sum_{j=1}^N \left[\nabla_{\underline{y}}^T \left(\frac{a_j(\underline{y}_l)}{p(\underline{y}_l)} \right) \int_{S_l} (\underline{y} - \underline{y}_l)(\underline{y} - \underline{y}_l)^T d\underline{y} \nabla_{\underline{y}} \left(\frac{a_j(\underline{y}_l)}{p(\underline{y}_l)} \right) \right] \quad (4.19)
\end{aligned}$$

where the gradient operator $\nabla_{\underline{y}} = [\partial/\partial y_1, \dots, \partial/\partial y_k]^T$ is a k -dimensional column vector, and $\int_{S_l} (\underline{y} - \underline{y}_l)(\underline{y} - \underline{y}_l)^T d\underline{y}$ is the *covariation profile* of the cell S_l .

4.3.2 Asymptotic loss in the Fisher Information trace

In [54], the covariation profile was used to analyze the performance of a discrimination-optimal quantizer that had non-symmetric, elliptical cells that quantized the log-likelihood ratio of a binary hypothesis test. For this dissertation we will restrict our analysis to quantizer cells S that are stretched and transformed version of a symmetric cell H . This is similar to the assumption of ellipsoidal cells in [54]. Under these restrictions the covariation profile reduces to a multiple of the identity matrix, and the resulting asymptotic expression for the loss in the FIM-trace Δtr (4.18) has a form similar to Bennett's integral (4.13). Let $\mathbf{B}^{-\frac{1}{2}}$ be a positive definite symmetric matrix, and let $S = \mathbf{B}^{-\frac{1}{2}}[H]$. By combining (4.8 - 4.10) the weighted covariation profile of S reduces to the following,

$$\mathbf{M}_{\mathbf{B}}(S) = |\mathbf{B}|^{1/k} \frac{m(H)}{k} \mathbf{I} \quad (4.20)$$

which when combined with (4.7, 4.18, and 4.19) results in the following for the loss in the FIM trace,

$$\Delta \text{tr} \cong \frac{m(H)}{M^{2/k}} \sum_{l=1}^M p(\underline{y}_l) \frac{\sum_{j=1}^N \left\| \nabla_{\underline{y}} \left(\frac{a_j(\underline{y}_l)}{p(\underline{y}_l)} \right) \right\|_{\mathbf{U}}^2}{\zeta(\underline{y}_l)^{2/k}} V(S_l) \quad (4.21)$$

where the matrix $\mathbf{U} = \mathbf{B}^{-1}|\mathbf{B}|^{1/k}$, and $\| \cdot \|_{\mathbf{U}}^2$ is a weighted 2-norm with respect to \mathbf{U} . Since (4.21) is an approximation of a Riemann integral, taking the limit in the

number of quantizer cells M results in the following expression

$$\lim_{M \rightarrow \infty} M^{2/k} \Delta \text{tr} = m(H) \int_S p(\underline{y}) \frac{\sum_{j=1}^N \left\| \nabla_{\underline{y}} \left(\frac{a_j(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2}{\zeta(\underline{y})^{2/k}} d\underline{y}. \quad (4.22)$$

Similar to Bennett's Integral (4.13) for a MSE-optimal quantizer, equation (4.22) shows that asymptotically the loss Δtr falls off at a rate of $M^{-2/k}$ with a constant term that depends on the distribution $p(\underline{y})$, the quantizer cell shape $m(H)$, and quantizer cell density $\zeta(\underline{y})$.

4.3.3 Optimal point density

Using Hölder's inequality to solve for the optimal point density function $\zeta^*(\underline{y})$ of a quantizer that minimizes the distortion in (4.22) leads to

$$\zeta^*(\underline{y}) = \frac{\left\{ p(\underline{y}) \sum_{j=1}^N \left\| \nabla_{\underline{y}} \left(\frac{a_j(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2 \right\}^{\frac{k}{k+2}}}{\int_S \left\{ p(\underline{y}) \sum_{j=1}^N \left\| \nabla_{\underline{y}} \left(\frac{a_j(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2 \right\}^{\frac{k}{k+2}} d\underline{y}}. \quad (4.23)$$

Substituting (4.23) into (4.22) gives the following expression for least-possible loss in the FIM-trace of any M -point quantizer in \mathbb{R}^k whose cells are a scaled version of a symmetric cell H ,

$$\Delta \text{tr} \cong \frac{m(H)}{M^{2/k}} \left[\int_S \left\{ p(\underline{y}) \sum_{j=1}^N \left\| \nabla_{\underline{y}} \left(\frac{a_j(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2 \right\}^{\frac{k}{k+2}} d\underline{y} \right]^{\frac{k+2}{k}}. \quad (4.24)$$

As before, if the cell is the minimum NMI tessellating polytope in \mathbb{R}^k then substituting $m(H) = m_k^*$ into (4.24) gives an expression for the least-possible trace-FIM distortion of any M -point quantizer in \mathbb{R}^k .

4.3.4 Loss in Fisher information trace for 1D example

As a simple study, we investigated how different quantization strategies would affect the Fisher information for a simple 1D estimation problem. Given a continuous source corrupted by spatial blurring, the task was to estimate N discrete pixel values

of a discretized version of the source from M quantized measurements. A plot of the observed measurement density $p(\underline{y})$, the point-density function $\zeta(\underline{y})$, as well as the quantizer cell locations for a $M = 32$ level quantizer is shown in figure 4.6. Note that the trace-FIM quantizer places cells near regions of *change* in the density, rather than in proportion to the density as like the MSE-optimal quantizer. The quantized

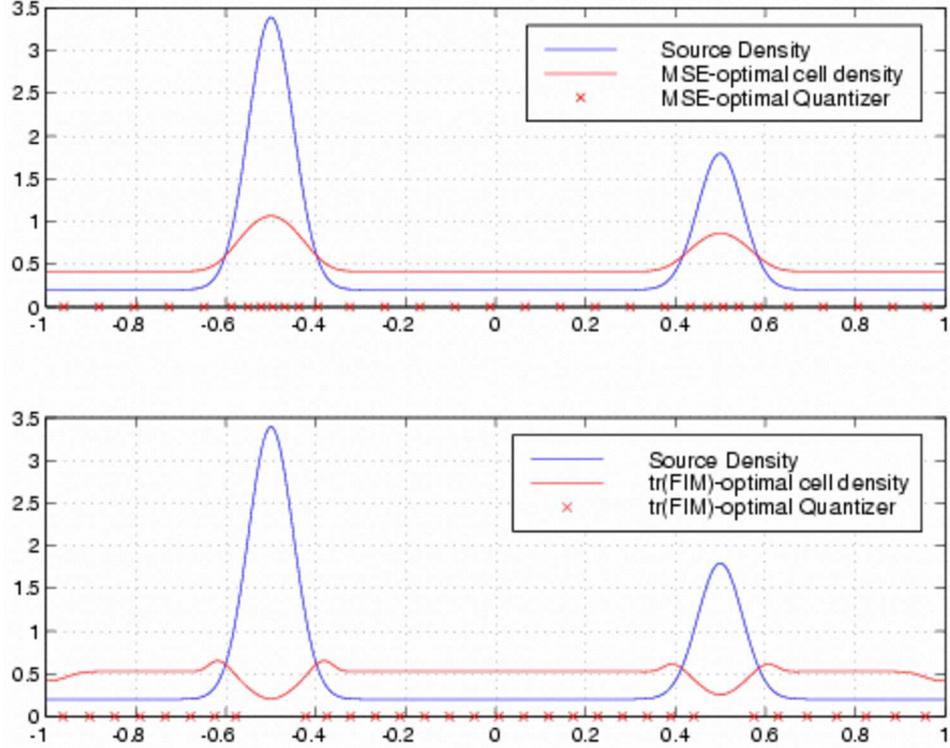


Figure 4.6: Comparison of optimal point densities and bin location for a 32-bin, MSE-optimal and tr-FIM-optimal 1D quantizer.

measurement Fisher information, its inverse, as well as the trace. The numerical values for a $M = 150$ level quantizer are given in table 4.1. The trace-FIM optimal

Quantizer	$\frac{1}{N}\text{tr}(\hat{\mathbf{F}})$	$\frac{1}{N}\text{tr}(\hat{\mathbf{F}}^{-1})$	
uniform	28.5283	3.350×10^6	3.00×10^{10}
mse	27.1602	∞	3.25×10^{17}
trace-FIM	28.6227	6.128×10^6	1.02×10^{12}

Table 4.1: Numerical values for the trace of the Fisher information and inverse Fisher information matrices, for three different quantizer point densities, for an $M = 150$ level quantizer.

quantizer does maximize the trace relative to the other quantizers, but at a cost. The condition number of the FIM is increased. All of the condition numbers are very large, corresponding to the ill-conditioned nature of the FIM for deconvolving significant blur. The worst quantizer, in terms of maximizing the trace, is the standard MSE-optimal quantizer with to numerical precision was singular.

4.4 Minimum Discrimination Loss Vector Quantizer

For image reconstruction problems, the parameters of interest θ are indirectly related to the measurements \underline{y} through a statistical measurement equation. MSE-optimal quantization of the measurements \underline{y} is not directly useful since it says nothing about the quality of the estimates of θ based on quantized measurements $Q(\underline{y})$. A more relevant metric for medical imaging would be the detection performance of a lesion detection task [7, 8, 60]. Gupta [54] developed a discrimination-optimal quantizer for testing a binary hypothesis that minimized the loss in discrimination due to quantization between the densities under the two hypothesis.

In medical imaging applications, one is certainly concerned with dramatic changes between two radically different measurement densities. However, such changes are trivial to detect. Here we focus on the more challenging problem of ensuring high sensitivity to small relevant changes in an initial density. For example, the difference between a radiotracer image of a patient before and after the sudden appearance of a tumor would manifest itself as a subtle variation of a few pixels against a similar background image. One could consider the background image as an image parameter vector θ and would want to be sensitive to a small change $\Delta\theta$ appearing somewhere in the background. The Kullback-Liebler discrimination function [15, 28] is related to the probability of error of deciding that the true parameter is θ when it is actually $\theta + \Delta\theta$ through the Chernoff bound. It is also related to the false alarm- and miss-

rate for binary hypothesis testing of presence or absence of the tumor. With this in mind, we will extend the work of Gupta by introducing a distortion measure based on the loss in discrimination due to quantization.

4.4.1 Binary Hypothesis Testing

In hypothesis testing, one must process a sequence of measurements in order to decide which out of a set of hypothesis is true given the measurements. In general, the observations are assumed to come from a family of probability distributions, where the particular form of the distribution is different depending on which of the hypotheses is true. For our purposes, we will consider the simpler problem of deciding between two hypothesis H_0 and H_1 (typically referred to as the *null* and *alternative* hypothesis, respectively), given a set of n i.i.d. k -dimensional observations $\mathbf{y} = [\underline{y}^{(1)}, \dots, \underline{y}^{(n)}]$ from one of two known probability distribution $p_0(\underline{y}), p_1(\underline{y})$. The probability densities associated with each hypothesis is given by

$$\begin{aligned} H_0 & : \mathbf{y} \sim p_0^{(n)}(\mathbf{y}) \\ H_1 & : \mathbf{y} \sim p_1^{(n)}(\mathbf{y}) \end{aligned}$$

where $p_0^{(n)}(\mathbf{y}) = \prod_{i=1}^n p_0(\underline{y}^{(i)})$ and $p_1^{(n)}(\mathbf{y}) = \prod_{i=1}^n p_1(\underline{y}^{(i)})$ are known functions of \underline{y}^i . We have assumed that the measurements \underline{y} are i.i.d.

Let the set $S \subset \mathbb{R}^k$ be the common support set of the two densities $p_0(\underline{y}), p_1(\underline{y})$, and let $S^{(n)}$ be the set of all n -tuples $[\underline{y}^{(1)}, \dots, \underline{y}^{(n)}]$ from either of the two densities. Our hypothesis test is a decision rule that partitions the set of all possible observations into two disjoint regions $S_0^{(n)}$ and $S_1^{(n)}$. If the observation $\mathbf{y} \in S_0^{(n)}$, then we say that the hypothesis H_0 is true. Similarly, if $\mathbf{y} \in S_1^{(n)}$ then H_1 is true.

There are two quantities that define the performance of any decision rule: the probability of false-alarm and the probability of detection (or equivalently, the probability of miss). Define the probability of false-alarm $\alpha = Pr(\text{decide } H_1 | H_0 \text{ true})$ as the probability of deciding that the observation vector \mathbf{y} comes from density under

hypothesis H_1 when it actually came from the density under H_0 . Similarly, define the probability of miss $\beta = Pr(\text{decide } H_0 | H_1 \text{ true})$ as the probability deciding that \mathbf{y} came from hypothesis H_0 when it actually came from the density under H_1 .

$$\alpha = \int_{S_1^{(n)}} p_0^{(n)}(\mathbf{y}) d\mathbf{y}$$

$$\beta = \int_{S_0^{(n)}} p_1^{(n)}(\mathbf{y}) d\mathbf{y}$$

4.4.2 Neyman-Pearson Test and ROC Curve

The Neyman-Pearson test [103] for deciding between two simple hypotheses is to compare the log of the ratios of the likelihood functions under the two hypothesis against a scalar threshold,

$$\Lambda(\underline{\mathbf{y}}) = \ln \frac{p_1(\underline{\mathbf{y}})}{p_0(\underline{\mathbf{y}})} \underset{H_0}{\overset{H_1}{\leq}} T$$

The Neyman-Pearson Theorem [15, 28, 103] establishes the optimality of the above test in that under the constraint of a fixed miss rate $\beta \leq \beta^*$, the best possible false alarm α is achieved by a likelihood-ratio test where the decision threshold T depends on the minimum allowable miss rate β^* . The likelihood-ratio also appears in the *ideal observer* [7, 79], but where the optimal threshold and associated false alarm and miss rate are determined by the prior probabilities of the two hypothesis.

By varying the value of the decision threshold T , one obtains a family of decision rules with specific values of (α, β) parameterized by the threshold T . The tradeoff between α and β can be shown graphically by plotting the *Probability of Detection* $(1 - \beta)$ and *Probability of False Alarm* (α) parametrically with the threshold T , and is referred to as the *receiver operating characteristic (ROC) curve*.

4.4.3 Asymptotic Performance of Hypothesis Test

The ROC curve can be approximated parametrically by large deviation representations of α and β using Sanov's Theorem [15, 28]. This parameterization depends on a quantity called the *Kullback-Liebler (KL) divergence*. The KL divergence is a measure of similarity between two densities. Given two continuous densities p and q defined over some common support set $S \subset \mathbb{R}^k$, the Kullback-Liebler discrimination or relative entropy between the two densities is given by

$$\mathcal{D}(q\|p) = \int_S q(\underline{y}) \ln \frac{q(\underline{y})}{p(\underline{y})} d\underline{y} \quad (4.25)$$

Similarly, the discrimination between two different discrete probability mass functions (PMF) $P = \{P_1, \dots, P_M\}$, $Q = \{Q_1, \dots, Q_M\}$ is given by

$$\mathcal{D}(Q\|P) = \sum_{i=1}^M Q_i \ln \frac{Q_i}{P_i} \quad (4.26)$$

The discrimination is always non-negative, and equals zero iff the two densities are the same. For the purposes of this paper we will be using the natural (base- e) logarithm in the definitions of equations (4.25, 4.26), resulting in a discrimination with units of *nats*.

Stein's lemma states that for a binary hypothesis test between two densities p_0 and p_1 with a fixed miss-rate β , the best possible false-alarm rate α^* among all possible decision rules of n -total i.i.d. observations is exponential in the discrimination between p_0 and p_1 ,

$$\lim_{n \rightarrow \infty} (\alpha_n^*)^{1/n} = e^{-\mathcal{D}(p_1\|p_0)}$$

Sanov's theorem gives a more general result for both the false alarm-rate α and miss-rate β as a function of the Neyman-Pearson test threshold T . Define the *tilted* density $p_\lambda(\underline{y})$ as the geometric mixture of $p_0(\underline{y})$ and $p_1(\underline{y})$,

$$p_\lambda(\underline{y}) = \frac{p_0^{1-\lambda}(\underline{y})p_1^\lambda(\underline{y})}{\int_S p_0^{1-\lambda}(\underline{y})p_1^\lambda(\underline{y})d\underline{y}}$$

where $\lambda \in [0, 1]$ is implicitly defined by the Neyman-Pearson threshold T ,

$$\begin{aligned} T &= \int_S p_\lambda(\underline{y}) \ln \frac{p_0(\underline{y})}{p_1(\underline{y})} d\underline{y} \\ &= \mathcal{D}(p_\lambda \| p_1) - \mathcal{D}(p_\lambda \| p_0). \end{aligned}$$

Then for large number of observations n the false alarm-rate α and miss-rate β are exponential in the discriminations $\mathcal{D}(p_\lambda \| p_0)$ and $\mathcal{D}(p_\lambda \| p_1)$ [15, 28],

$$\alpha \cong e^{-n\mathcal{D}(p_\lambda \| p_0)} \quad (4.27)$$

$$\beta \cong e^{-n\mathcal{D}(p_\lambda \| p_1)} \quad (4.28)$$

A similar result holds for the discrete case for the asymptotic false alarm- and miss-rates, where the integrals are replaced with summations.

4.4.4 Small Signal Perturbation Analysis

Let the density under the alternative hypothesis be a small additive perturbation $p_\Delta(\underline{y})$ of the density under the null hypothesis $p(\underline{y})$. Then the normalized densities under H_0, H_1 are given by

$$\begin{aligned} H_0 : p_0(\underline{y}) &= p(\underline{y}) \\ H_1 : p_1(\underline{y}) &= \frac{1}{1+\epsilon} p(\underline{y}) + \frac{\epsilon}{1+\epsilon} p_\Delta(\underline{y}) \end{aligned}$$

where $\epsilon \ll 1$. Under this small perturbation model, the divergence of the tilted density $p_\lambda(\underline{y})$ with respect to $p_0(\underline{y})$ and $p_1(\underline{y})$ simplify to

$$\mathcal{D}(p_\lambda \| p_0) = \frac{1}{2} \lambda^2 \epsilon^2 (2D - 1) + o(\epsilon^3) \quad (4.29)$$

$$\mathcal{D}(p_\lambda \| p_1) = \frac{1}{2} (\lambda - 1)^2 \epsilon^2 (2D - 1) + o(\epsilon^3) \quad (4.30)$$

where D is given by

$$D = \frac{1}{2} \int_S \frac{p_\Delta^2(\underline{y})}{p(\underline{y})} d\underline{y} \quad (4.31)$$

As for quantized measurements, given a Vector Quantizer $Q(\underline{y})$ defined over the partition $\{S_1, \dots, S_M\}$ of S , we can define the discrete probability mass function (PMF) of the fraction of quantized measurements in each cell $P = \{P_1, \dots, P_M\}$, where the probability of the continuous measurement \underline{y} being quantized to the l th value is given by the integral over the l th quantizer partition (i.e. $P_l = \int_{S_l} p(\underline{y}) d\underline{y}$). The PMFs under H_0, H_1 are

$$\begin{aligned} H_0 : P_0 &= P \\ H_1 : P_1 &= \frac{1}{1+\epsilon}P + \frac{\epsilon}{1+\epsilon}P_\Delta \end{aligned}$$

with the associated divergence of the tilted PMF P_λ with respect to P_0 and P_1 of

$$\mathcal{D}(P_\lambda \| P_0) = \frac{1}{2}\lambda^2\epsilon^2(2\bar{D} - 1) + o(\epsilon^3) \quad (4.32)$$

$$\mathcal{D}(P_\lambda \| P_1) = \frac{1}{2}(\lambda - 1)^2\epsilon^2(2\bar{D} - 1) + o(\epsilon^3) \quad (4.33)$$

where \bar{D} is the quantized version of D

$$\bar{D} = \frac{1}{2} \sum_{i=1}^M \frac{P_{\Delta i}^2}{P_i} \quad (4.34)$$

As an example calculation, Figure 4.7 shows the Null and Alternative hypothesis densities for the simple case of $p(y) \sim \text{uniform}[-4, 4]$ and the perturbation density $p_\Delta(y) \sim N(0, 1)$ on the interval $S = [-4, 4]$, with a perturbation weighting of $\epsilon = 10^{-2}$. The corresponding discriminations with respect to the tilted densities $\mathcal{D}(p_\lambda \| p_0)$, $\mathcal{D}(p_\lambda \| p_1)$ are shown in Figure 4.8, and the asymptotic ROC curve for $n = 10^5$ i.i.d. measurements is shown in Figure 4.9. Both Figures 4.8 and 4.9 show that the value of the discrimination functions $\mathcal{D}(p_\lambda \| p_0)$, $\mathcal{D}(p_\lambda \| p_1)$ as well as the resulting ROC curve calculated using (4.25) are well-approximated by (4.29) and (4.30).

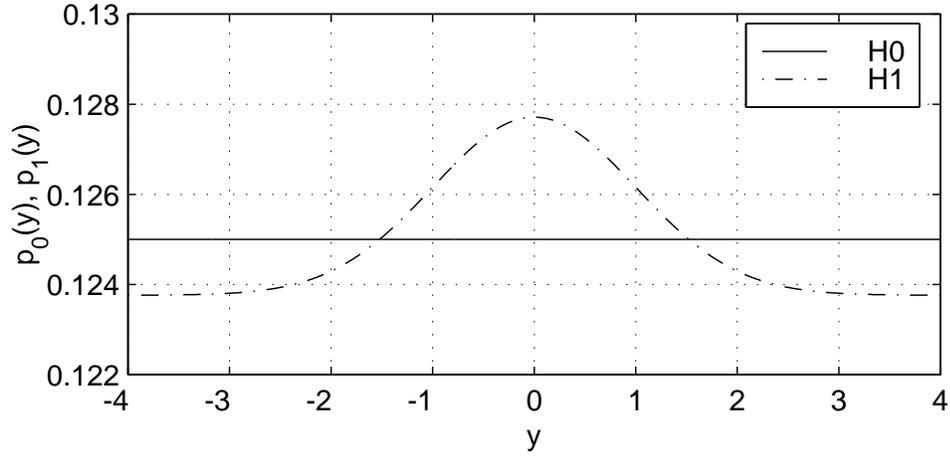


Figure 4.7: Null hypothesis and alternative hypothesis for 1D example.

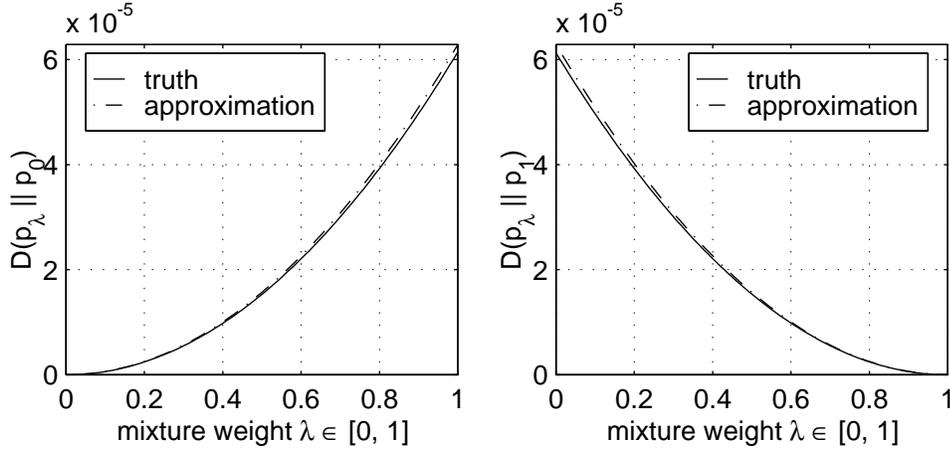


Figure 4.8: True and approximate discrimination functions for 1D example.

4.4.5 Discrimination Loss

Let $\Delta\mathcal{D}(p_\lambda\|p_0)$ and $\Delta\mathcal{D}(p_\lambda\|p_1)$ be the loss in discrimination of the tilted density due to quantization,

$$\begin{aligned}\Delta\mathcal{D}(p_\lambda\|p_0) &= \mathcal{D}(p_\lambda\|p_0) - \mathcal{D}(P_\lambda\|P_0) \\ &= \lambda^2\epsilon^2\Delta D + o(\epsilon^3)\end{aligned}\tag{4.35}$$

and

$$\begin{aligned}\Delta\mathcal{D}(p_\lambda\|p_1) &= \mathcal{D}(p_\lambda\|p_1) - \mathcal{D}(P_\lambda\|P_1) \\ &= (\lambda - 1)^2\epsilon^2\Delta D + o(\epsilon^3)\end{aligned}\tag{4.36}$$

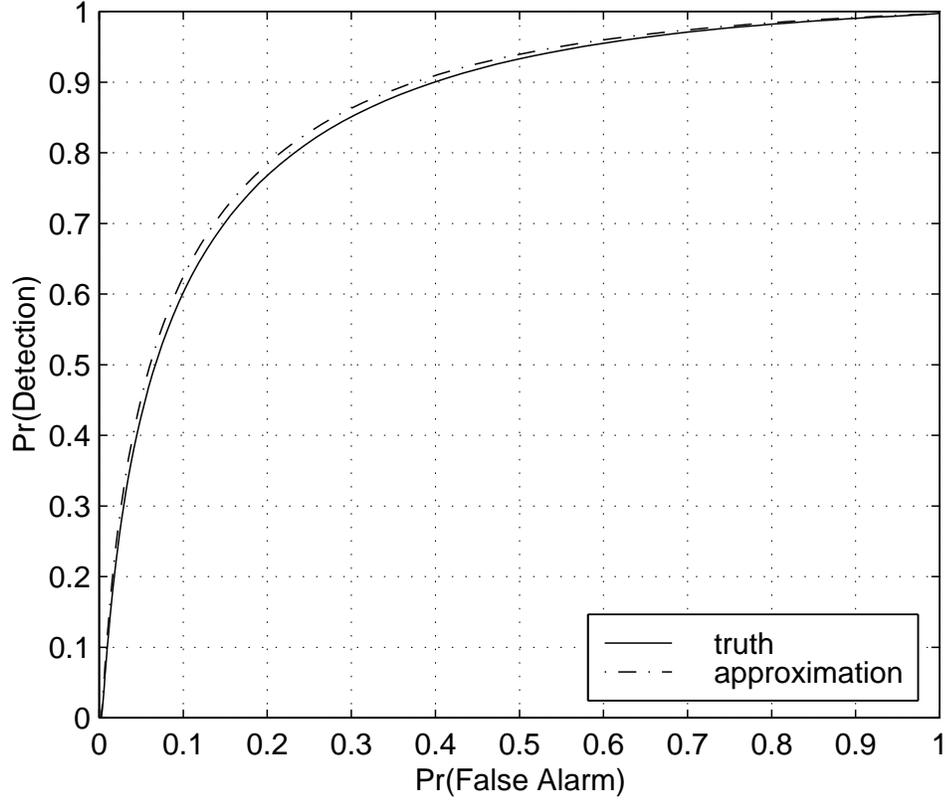


Figure 4.9: True and approximate ROC curve for 1D example.

where

$$\Delta D = \frac{1}{2} \left[\int_S \frac{p_{\Delta}^2(\underline{y})}{p(\underline{y})} d\underline{y} - \sum_{i=1}^M \frac{P_{\Delta i}^2}{P_i} \right]. \quad (4.37)$$

From the Cauchy-Schwartz inequality, it can be easily shown that $\Delta D \geq 0$. Thus there is a loss in discrimination due to quantization. Let $\hat{\alpha}$ and $\hat{\beta}$ be the false alarm- and miss-rate for a hypothesis test using quantized data. By substituting (4.35, 4.36) into (4.27, 4.28), we arrive at the following expressions relating the change in the ROC curve due to quantization,

$$\hat{\alpha} \cong \alpha e^{n\Delta\mathcal{D}(p_{\lambda}\|p_0)} \quad (4.38)$$

$$\hat{\beta} \cong \beta e^{n\Delta\mathcal{D}(p_{\lambda}\|p_1)}. \quad (4.39)$$

Since ΔD is non-negative, there is an increase due to quantization in the false alarm-rate α and miss-rate β , since they are exponential in $\Delta\mathcal{D}(p_{\lambda}\|p_0)$ and $\Delta\mathcal{D}(p_{\lambda}\|p_1)$,

which in turn are linear in ΔD . For a given quantizer partition, one can calculate the loss ΔD and the associated change in the ROC curve. Alternatively, by finding a quantizer partition that minimizes ΔD , one minimizes the loss in the ROC curve due to quantization. Since the quantizer cells $\{S_1, \dots, S_M\}$ are a partition of S , the integral in (4.37) can be written as a summation of integrals over each cell,

$$\int_S \frac{p_\Delta^2(\underline{y})}{p(\underline{y})} d\underline{y} = \sum_{l=1}^M \int_{S_l} \frac{p_\Delta^2(\underline{y})}{p(\underline{y})} d\underline{y}$$

When the number of cells M is large and the corresponding quantization cells are small, and assuming that the density $p(\underline{y})$ is smooth over the extent of each cell, the difference term in (4.37) can be expanded out in a Taylor series expansion about \underline{y}_l up to terms of order $o(\|\underline{y} - \underline{y}_l\|^2)$,

$$\begin{aligned} & \int_{S_l} \frac{p_\Delta^2(\underline{y})}{p(\underline{y})} d\underline{y} - \frac{P_{\Delta l}^2}{P_l} \\ & \cong p(\underline{y}_l) \nabla_{\underline{y}}^T \left(\frac{p_\Delta(\underline{y}_l)}{p(\underline{y}_l)} \right) \left[\int_{S_l} (\underline{y} - \underline{y}_l)(\underline{y} - \underline{y}_l)^T d\underline{y} \right] \nabla_{\underline{y}} \left(\frac{p_\Delta(\underline{y}_l)}{p(\underline{y}_l)} \right) \end{aligned} \quad (4.40)$$

where the gradient operator $\nabla_{\underline{y}} = [\partial/\partial y_1, \dots, \partial/\partial y_k]^T$ is a k -dimensional column vector, and $\int_{S_l} (\underline{y} - \underline{y}_l)(\underline{y} - \underline{y}_l)^T d\underline{y}$ is the *covariation profile* of the cell S_l .

4.4.6 Asymptotic Discrimination Loss

As before, for this dissertation we will restrict our analysis to quantizer cells S that are stretched and transformed version of a symmetric cell H . Under these restrictions the covariation profile reduces to a multiple of the identity matrix, and the resulting asymptotic expression for the discrimination loss ΔD (4.37) has a form similar to Bennett's integral (4.13). Let $\mathbf{B}^{-\frac{1}{2}}$ be a positive definite symmetric matrix, and let $S = \mathbf{B}^{-\frac{1}{2}}[H]$. By combining (4.8 - 4.10) the weighted covariation profile of S reduces to the following,

$$\mathbf{M}_{\mathbf{B}}(S) = |\mathbf{B}|^{1/k} \frac{m(H)}{k} \mathbf{I} \quad (4.41)$$

which when combined with (4.40, 4.7) results in the following for the discrimination loss,

$$\Delta D \cong \frac{1}{2} \frac{m(H)}{M^{2/k}} \sum_{l=1}^M p(\underline{y}_l) \frac{\left\| \nabla_{\underline{y}} \left(\frac{p_{\Delta}(\underline{y}_l)}{p(\underline{y}_l)} \right) \right\|_{\mathbf{U}}^2}{\zeta(\underline{y}_l)^{2/k}} V(S_l) \quad (4.42)$$

where the matrix $\mathbf{U} = \mathbf{B}^{-1} |\mathbf{B}|^{1/k}$ is the normalized weighting along each dimension and $\| \cdot \|_{\mathbf{U}}^2$ is a weighted 2-norm with respect to \mathbf{U} . Since (4.42) is an approximation of a Riemann integral, taking the limit in the number of quantizer cells M results in the following expression

$$\lim_{M \rightarrow \infty} M^{2/k} \Delta D = \frac{1}{2} m(H) \int_S p(\underline{y}) \frac{\left\| \nabla_{\underline{y}} \left(\frac{p_{\Delta}(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2}{\zeta(\underline{y})^{2/k}} d\underline{y}. \quad (4.43)$$

Similar to Bennett's Integral (4.13) for a MSE-optimal quantizer, equation (4.43) shows that asymptotically the discrimination loss ΔD falls off at a rate of $M^{-2/k}$ with a constant term that depends on the distribution $p(\underline{y})$, the quantizer cell shape $m(H)$, and quantizer cell density $\zeta(\underline{y})$. In addition, the discrimination optimal quantizer also depends on the magnitude-square of the gradient ratio between the background and perturbation density. For small signal perturbations, the magnitude of the density ratio gradient is equivalent to the magnitude of the log-likelihood ratio gradient,

$$\left\| \nabla_{\underline{y}} \ln \frac{p_1(\underline{y})}{p_0(\underline{y})} \right\|^2 = \left\| \nabla_{\underline{y}} \left(\frac{p_{\Delta}(\underline{y})}{p(\underline{y})} \right) \right\|^2 \epsilon^2 + o(\epsilon^3). \quad (4.44)$$

Gupta [54] referred to the squared magnitude of the log-likelihood ratio (4.44) as the *discriminability*, and noted that it is a measure of the usefulness of a measurement \underline{y} in deciding between the hypotheses H_0 and H_1 .

4.4.7 Optimal point density

Using Hölder's inequality to solve for the optimal point density function $\zeta^*(\underline{y})$ of a quantizer that minimizes the discrimination loss distortion in (4.43) leads to

$$\zeta^*(\underline{y}) = \frac{\left\{ p(\underline{y}) \left\| \nabla_{\underline{y}} \left(\frac{p_{\Delta}(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2 \right\}^{\frac{k}{k+2}}}{\int_S \left\{ p(\underline{y}) \left\| \nabla_{\underline{y}} \left(\frac{p_{\Delta}(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2 \right\}^{\frac{k}{k+2}} d\underline{y}}. \quad (4.45)$$

Substituting (4.45) into (4.43) gives the following expression for least-possible discrimination loss of any M -point quantizer in \mathbb{R}^k whose cells are a scaled version of a symmetric cell H ,

$$\Delta D \cong \frac{1}{2} \frac{m(H)}{M^{2/k}} \left[\int_S \left\{ p(\underline{y}) \left\| \nabla_{\underline{y}} \left(\frac{p_{\Delta}(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2 \right\}^{\frac{k}{k+2}} d\underline{y} \right]^{\frac{k+2}{k}}. \quad (4.46)$$

4.5 Quantizer Design

The fully optimal VQ has minimum discrimination loss ΔD (4.46) for a fixed large number of cells (equivalently, codebook size). However, without further constraints, it has a high implementation complexity in both storage and computational load. This high complexity is primarily due to the fact that each cell is uniquely defined as the voronoi partition about each codebook vector. Direct implementing of the quantization rule in (4.1) involves a *brute-force* search over all M cells for a given value of \underline{y} by calculating $\|\underline{y} - \underline{y}_l\|^2$, $l = 1, \dots, M$ and choosing the codevector with minimum distortion. A M -cell quantizer of a k -dimensional measurement \underline{y} with b -bits for each component of \underline{y} results in a storage requirement of $O(Mkb)$ bits to store the codebook, along with an computational complexity of $O(3Mk)$ per sample.

Optimal unstructured minimum-MSE quantizers can be generated via the well-known Linde-Buzo-Gray (LBG) [75] or generalized Lloyd algorithm. Given a *training set* of p realizations of random training vectors along with the desired number M of quantizer cells / codevectors, the LBG algorithm iteratively processes the training set

until it arrives at a solution where each training vector is associated with a codevector via a nearest-neighbor (voronoi) mapping, and that the codevectors are the centroids of each cell. The resulting quantizer partition will be MSE optimal, assuming it does not settle in a local minimum configuration.

Each iteration of the LBG algorithm involves quantizing p -total, k -dimensional training vectors to one of M -different codevectors, involving a brute-force search over all M codevectors for a resulting computational complexity of $o(3MNk)$ per iteration. The “rule of thumb” for designing large quantizers is for the number of training samples $p \approx 50N$ or more. Thus designing a quantizer becomes quite challenging for large data-sets (large N) with fine quantization (large M). The communications and source-coding literature refers to the *Rate* of a codebook $R = 1/k \log_2(M)$ as the average number of bits per component needed to enumerate M -total k -dimensional codevectors. The complexity of an unstructured quantizer is given by kR , the *dimension-Rate product*. Although with modern computers it is possible to store the codevectors associated with a large dimension-Rate product quantizer, a practical upper limit on designing an unstructured quantizer is for $kR \approx 14$ or equivalently a maximum codebook size $M \approx 16 \times 10^3$.

There are fast-search methods for searching over the M codevectors of an unstructured vector quantizer, of which [52] has an excellent survey. There are also sub-optimal algorithms such as *tree-structured* vector quantizers [5,22] that enforce a structure on the codevector locations, are much more efficient to implement, and can be searched in $O(\log(M))$ per sample. We will focus on the applications of *Lattice Quantizers* for fast quantization. As we will see, a Lattice Quantizer has a trivial computational complexity of $O(k^2)$ per sample with negligible design overhead and storage requirements, but comes at the expense of a sub-optimal uniform point density $\zeta(\underline{y})$.

4.6 Lattice Quantizer

Borrowing the notation of [2], a lattice $\mathcal{L} \subset \mathbb{R}^k$ can be defined as integer combinations of a set of basis vectors $\underline{b}_i \in \mathbb{R}^k$, $i = 1, \dots, k$,

$$\mathcal{L} = \{\underline{y} : \underline{y} = u_1 \underline{b}_1 + \dots + u_k \underline{b}_k\}$$

where $u_i \in \mathbb{Z}$ are integers. The matrix \mathbf{B} whose rows are the basis vectors is referred to as the *generator matrix* of the lattice,

$$\mathbf{B} = [\underline{b}_1, \underline{b}_2, \dots, \underline{b}_k]^T$$

Note that the basis vectors do not have to form a complete basis, resulting in what is known as a *degenerate* lattice. For this dissertation, we will only be concerned with non-degenerate lattices. The lattice \mathcal{L} can be expressed in terms of its $k \times k$ -dimensional generator matrix \mathbf{B} as

$$\mathcal{L} = \{\underline{y} \in \mathbb{R}^k : (\mathbf{B}^{-1})^T \underline{y} \in \mathbb{Z}^k\} \quad (4.47)$$

From the definition of a lattice (4.47), the quantization rule (4.1) for a lattice-based vector quantizer $Q_{\mathcal{L}}(\underline{y})$ that maps an input vector \underline{y} onto a point in \mathcal{L} can be expressed as a pair of matrix multiplications and a rounding operation,

$$\underline{y}_l = \mathbf{B}^T \text{round}\{(\mathbf{B}^{-1})^T \underline{y}\} \quad (4.48)$$

with a corresponding computational complexity of $o(k^2)$ per sample (and $o(k)$ for diagonal \mathbf{B}). Note that the only quantizer design issue involved with Lattice quantizers is choosing the generator matrix \mathbf{B} . The nearest-neighbor or voronoi region of the lattice defined by the generator matrix \mathbf{B} has volume $|\det \mathbf{B}|$, and the normalized moment of inertia and other properties for various lattices are well documented [24, 26].

The simplest possible lattice quantizer is the uniform scalar quantizer, with codevectors $\{\dots, -2\Delta, -\Delta, 0, \Delta, 2\Delta, \dots\}$ and associated quantization function

$$Q(y) = i\Delta, \text{ if } (i - 1/2)\Delta \leq y < (i + 1/2)\Delta$$

which corresponds to a 1-dimensional generator matrix $\mathbf{B} = \Delta$. Figure 4.10 shows a 2-dimensional lattice with rectangular cells ($\mathbf{B} = [2 \ 0; 0 \ 1]$) as well as hexagonal cells ($\mathbf{B} = [2 \ 0; 1 \ \sqrt{3}]$).

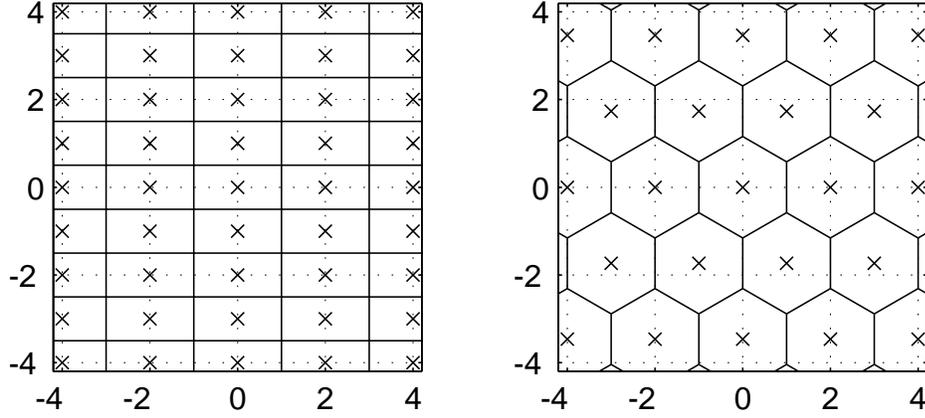


Figure 4.10: Example 2D lattice with rectangular cells (left) and hexagonal cells (right), with codevector location marked with “x” at centroid of each cell.

4.6.1 Discrimination Loss for Lattice Quantizer

Since the codevectors of a lattice quantizer are equally spaced, the associated point density is constant. Let S_L be the support set of the lattice quantizer. We will assume that the support set of the lattice contains the support set of the measurement density, i.e. $S \subseteq S_L$. Since the point density integrates to unity over the support set of the quantizer, the point density of a lattice quantizer is given by $\zeta(\underline{y}) = 1/V(S_L)$, with a resulting discrimination loss of

$$\Delta D = \frac{1}{2} m(H) \left(\frac{V(S_L)}{M} \right)^{2/k} \int_S p(\underline{y}) \left\| \nabla_{\underline{y}} \left(\frac{p_{\Delta}(\underline{y})}{p(\underline{y})} \right) \right\|_{\mathbf{U}}^2 d\underline{y}. \quad (4.49)$$

4.6.2 Discrimination Loss Calculation for 1D example

As in Section 4.4.4, let the null and alternative hypothesis densities be for the simple case of $p(y) \sim \text{uniform}[-4, 4]$ and the perturbation density $p_{\Delta}(y) \sim N(0, 1)$ on the interval $S = [-4, 4]$, with a perturbation weighting of $\epsilon = 10^{-2}$.

For a uniform quantizer in 1D, the cells consists of equal-size intervals over the real line. The $M = 2$ quantizer consists of the two quantizer cells $\{[-4, 0], [0, 4]\}$, the $M = 3$ quantizer consists of the cells $\{[-4, -2/3], [-2/3, 2/3], [2/3, 4]\}$, etc. For each quantizer (equivalently, each value of M), we calculate the discrimination term $\mathcal{D}(P_1||P_0)$ of the quantized measurements. Figure 4.11 shows the discrimination of each quantizer versus the number of cells, as well as the maximum possible discrimination $\mathcal{D}(p_1||p_0)$ of using continuous measurements.

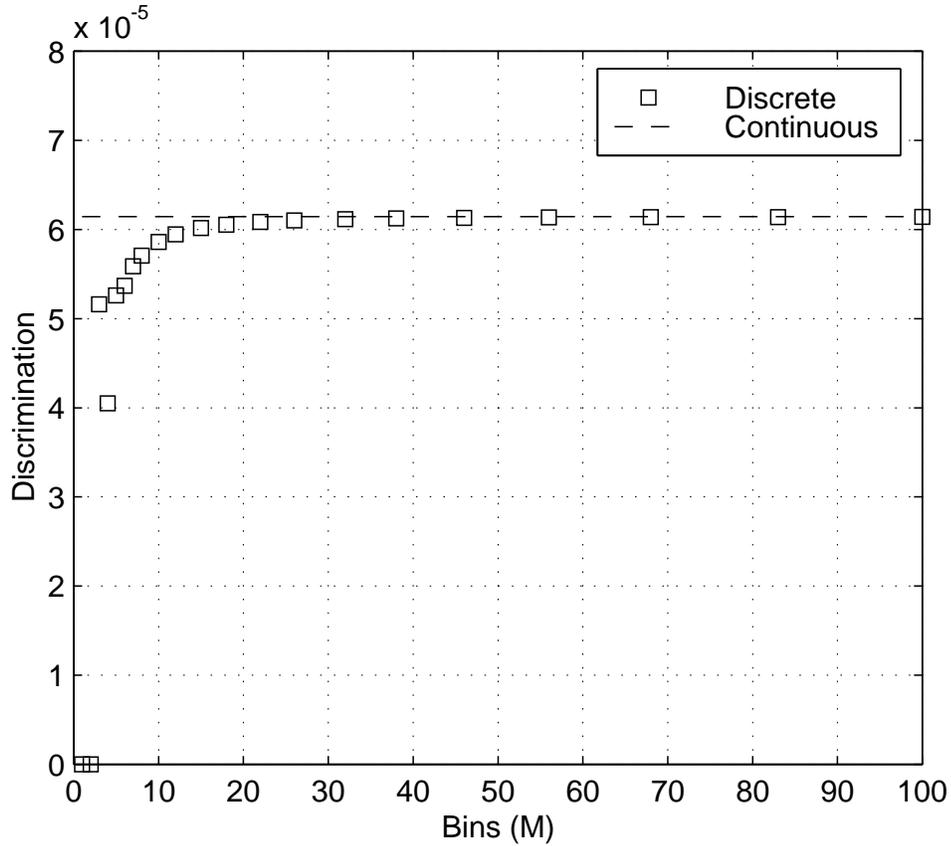


Figure 4.11: Discrete measurement discrimination $\mathcal{D}(P_1||P_0)$ versus continuous measurement discrimination $\mathcal{D}(p_1||p_0)$ for a 1D uniform scalar quantizer.

In Figure 4.12, the actual discrimination loss $\Delta\mathcal{D}(p_1||p_0) = \mathcal{D}(P_1||P_0) - \mathcal{D}(p_1||p_0)$ is plotted for each quantizer, as well as the asymptotic limit predicted by (4.49). Note that for $M \geq 6$ the discrimination loss $\Delta\mathcal{D}(p_1||p_0)$ is well predicted by the asymptotic expression.

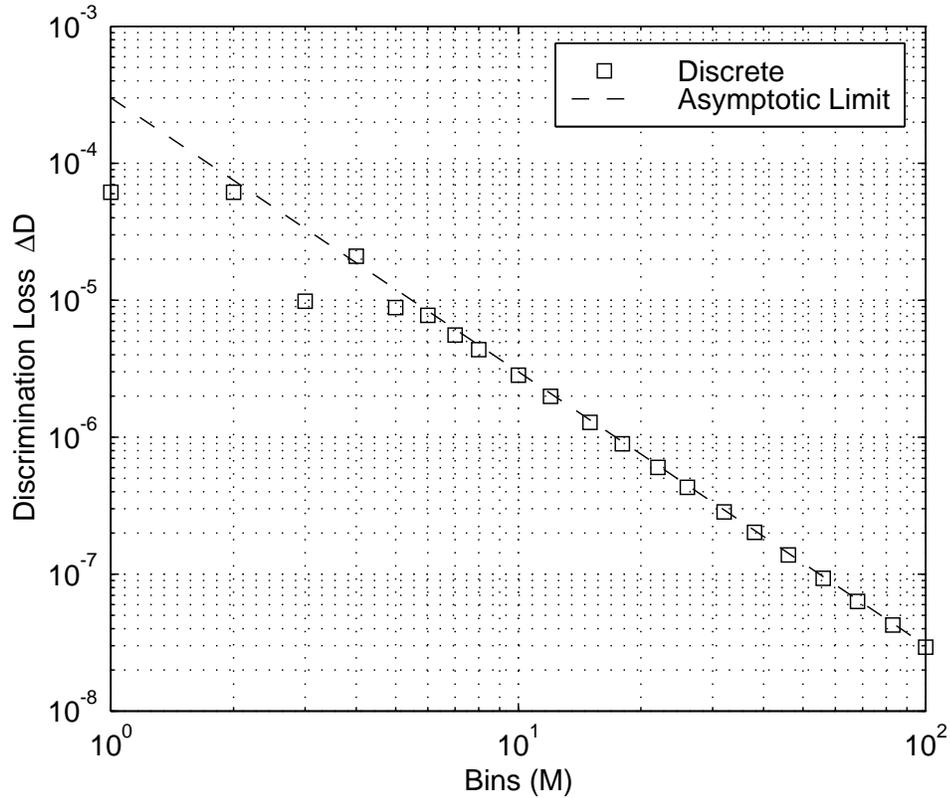


Figure 4.12: Actual loss in discrimination $\Delta\mathcal{D}(p_1||p_0)$ due to quantized measurements versus asymptotic limit for uniform scalar quantizer.

In Figure 4.13, the ROC curve for continuous measurements is compared against an $M = 6$, $M = 8$, and $M = 10$ level uniform scalar quantizer, with the resulting loss in false-alarm and detection-rate calculated from equations (4.38, 4.39)

4.6.3 Discrimination Loss Calculation for 2D PET

For illustration we will study the effect of measurement quantization in 2D PET for a lesion detection task. In PET, the source density $\lambda(\underline{x})$ is not directly observable, but rather indirectly by multiple measurements $\underline{y} \in \mathbb{R}^2$ of γ -ray photon pairs measured in coincidence by an external detector ring surrounding the patient. Each measurement $\underline{y} = \{\theta, r\}$ consists of the angular and radial offset of the line-of-response between each pair of detected γ -rays.

Let the unknown parameters under hypotheses H_0, H_1 be $\boldsymbol{\theta}$ and $\boldsymbol{\theta} + \Delta\boldsymbol{\theta}$, where

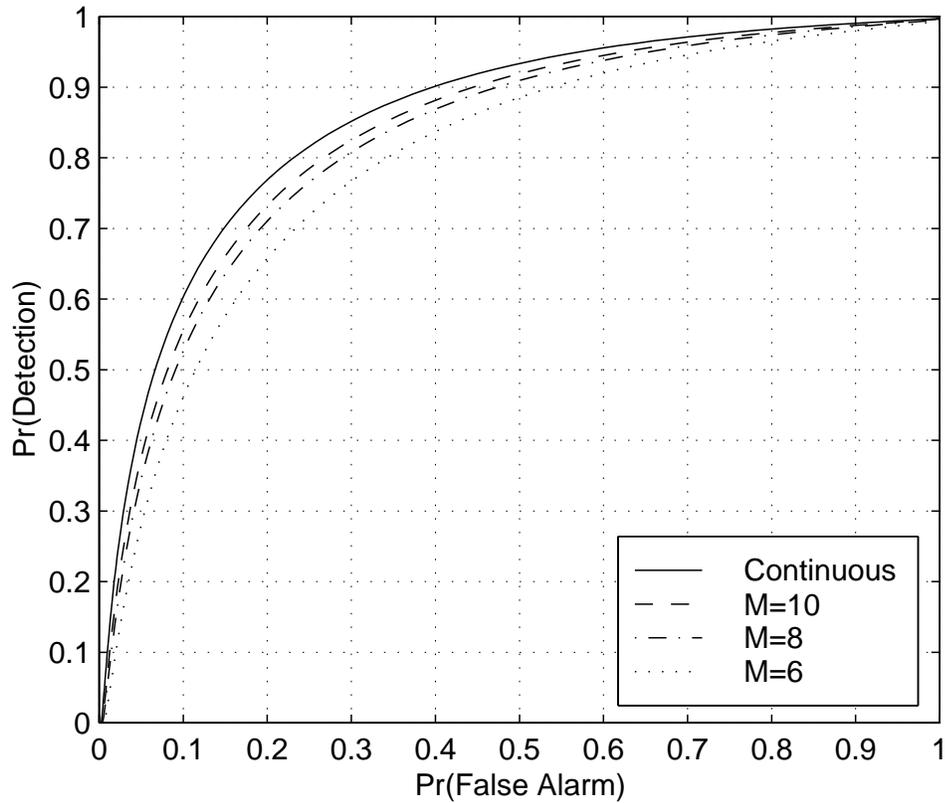


Figure 4.13: Loss in ROC curve due to quantization for 1D example.

$\boldsymbol{\theta} \in \mathbb{R}^N$ is some non-random parameter vector and $\Delta\boldsymbol{\theta}$ is a perturbation of $\boldsymbol{\theta}$. We will model the emission source density $\lambda(\underline{x})$ as an 100×100 -pixel, 2-dimensional (planar) image over a 50×50 cm FOV. The associated parameter vector $\boldsymbol{\theta}$ corresponds to a tomographic slice through a simulated anthropomorphic phantom [93] while the perturbation parameter $\Delta\boldsymbol{\theta}$ corresponds to a 3cm circular lesion in the bottom of the left lung, as shown in Figure 4.14.

The associated detection intensity rates under the two hypothesis are

$$\begin{aligned}
 H_0 : \mu_0(\underline{y}) &= \mu(\underline{y}|\boldsymbol{\theta}) \\
 H_1 : \mu_1(\underline{y}) &= \mu(\underline{y}|\boldsymbol{\theta}) + \mu(\underline{y}|\Delta\boldsymbol{\theta}),
 \end{aligned}$$

where the last line comes from the fact that in emission tomography the measurement intensity rate is linear in the parameter $\boldsymbol{\theta}$. Let $\mu_D = \int_S \mu_0(\underline{y}) d\underline{y}$ and $\mu_D + \mu_\Delta = \int_S \mu_1(\underline{y}) d\underline{y}$ be the total intensity rate under H_0 and H_1 . Let $\epsilon = \frac{\mu_\Delta}{\mu_D} \ll 1$ be the ratio

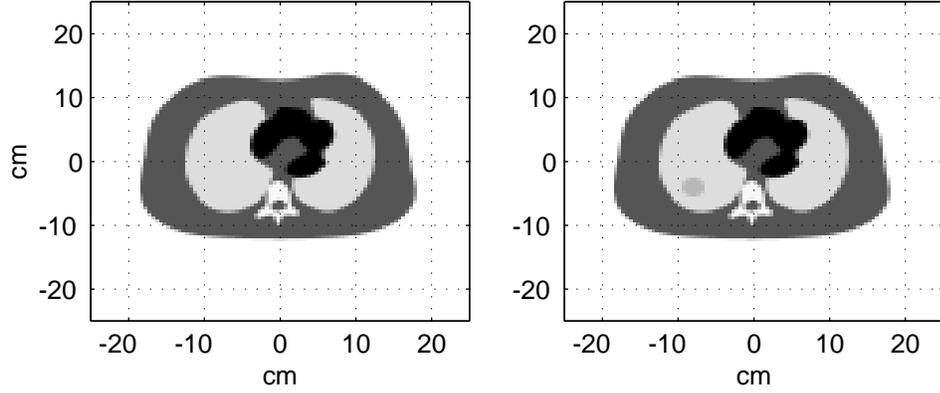


Figure 4.14: Simulated 2D PET emission source (left), emission source with 3cm circular lesion (right)

of detected measurements from the background parameter θ and the perturbation parameter $\Delta\theta$. Define the normalized background and perturbation densities $p(\underline{y})$, $p_{\Delta}(\underline{y})$ as

$$\begin{aligned}
 p(\underline{y}) &= \frac{1}{\mu_D} \mu(\underline{y}|\theta) \\
 p_{\Delta}(\underline{y}) &= \frac{1}{\mu_{\Delta}} \mu(\underline{y}|\Delta\theta)
 \end{aligned}$$

The PET detection ring is not a continuous measurement device, but rather consists of B discrete elements resulting in a measurement space S that is partitioned into (at most) $B(B - 1)$ uniquely ordered detector pair combinations. In a typical system, however, only some fraction of other blocks in an arc directly opposite will be used. The resulting line-of-response measurements $\underline{y} = \{\theta, r\}$ are quantized to one of M possible values with an inherent sampling in $\{\theta, r\}$ (Figure 4.15). Thus a PET detection ring can be modeled as a hardware-based vector quantizer where the inherent quantization as well as loss in discrimination will decrease as the number of detector blocks increase.

For this study we will assume that the total number of possible measurements $M = N_a * N_r$, where N_a is the number of parallel projection measurements over $[0, \pi]$ and N_r is the number of radial samples per projection angle. From the Fourier Slice

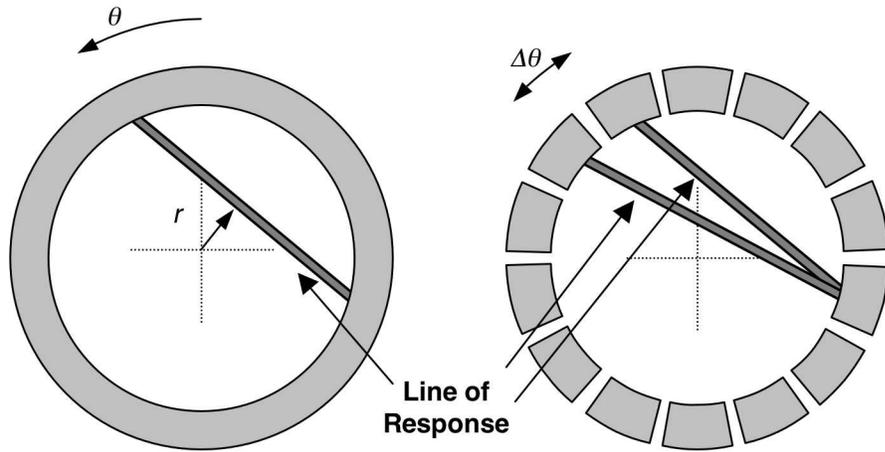


Figure 4.15: Line-of-Response for 2D PET with continuous (right) and Discrete (left) measurements.

Theorem [63] one can derive a Nyquist sampling condition of $N_a/N_r = \pi/2$ where for this study $N_r = 100$ (the number of pixels across the FOV) and $N_a = 157$, resulting in a measurement space with $M = 15700$ bins as the number required to estimate a 100×100 pixel parameter vector $\boldsymbol{\theta}$. For this study the total number of measurement bins M is varied by scaling the radial and angular sampling equally but keeping the ratio between them constant at $\pi/2$. The ratio was fixed so that the NMI was kept constant, and thus the asymptotic expression for discrimination loss (4.49) would apply.

Note that by decreasing the number of samples along each dimension we are not coarsely sub-sampling the measurement space, but rather each measurement bin is an integral over an increasingly larger partition area where the resulting sinogram bins are rectangular with constant normalized moment of inertia. Figure 4.16 is of a PET measurement space (or sinogram) sampled at $(N_a, N_r) = (94, 60)$ corresponding to $M = 5640$ total bins. Figure 4.17 is a sinogram sampled at $(N_a, N_r) = (157, 100)$ corresponding to $M = 15700$ bins. Whereas Figure 4.16 and 4.17 are of the measurement density $p(\underline{y})$, Figure 4.18 is the perturbation density $p_{\Delta}(\underline{y})$, otherwise known as the *forward projection* of $\Delta\boldsymbol{\theta}$.

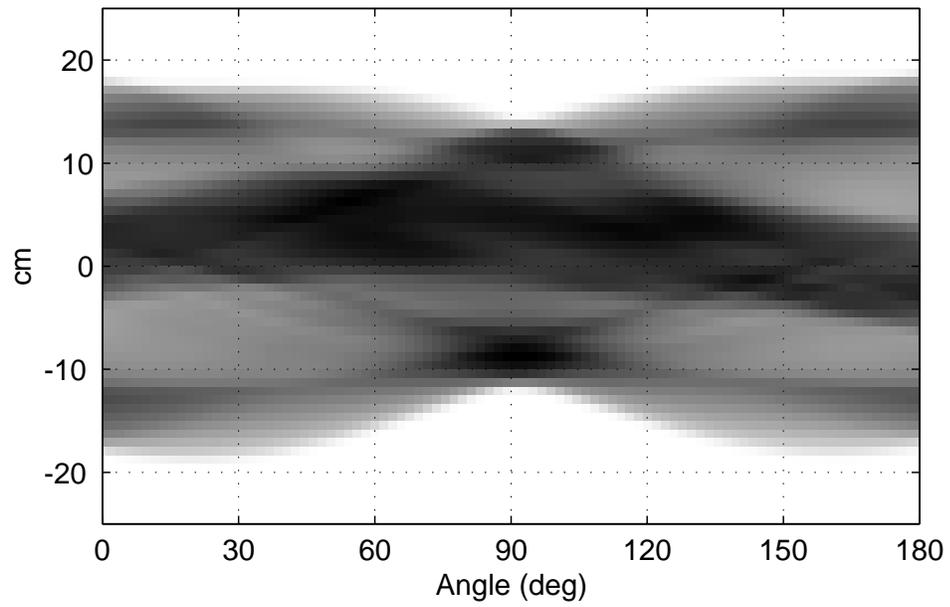


Figure 4.16: PET Sinogram with $M = 5640$ samples

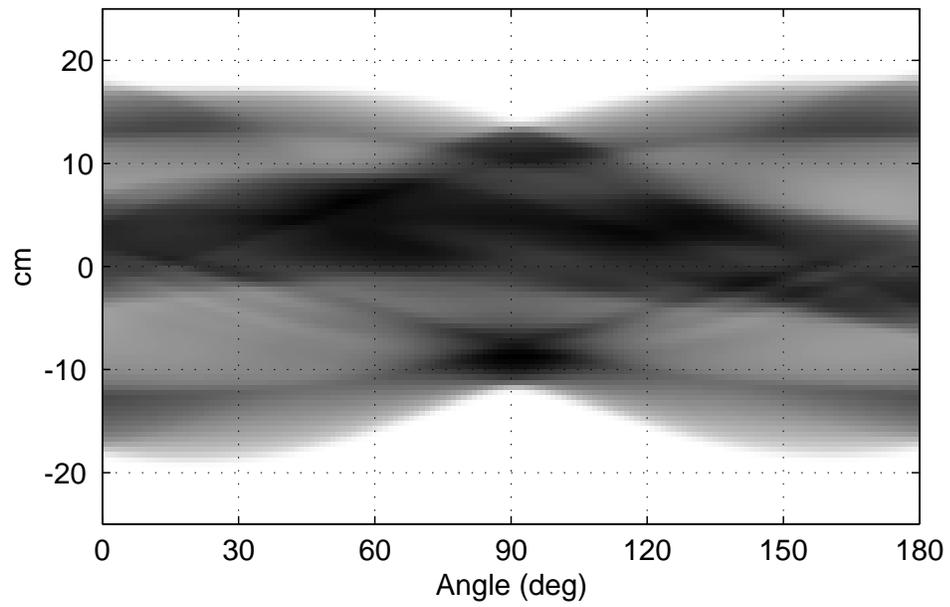


Figure 4.17: PET Sinogram with $M = 15700$ samples

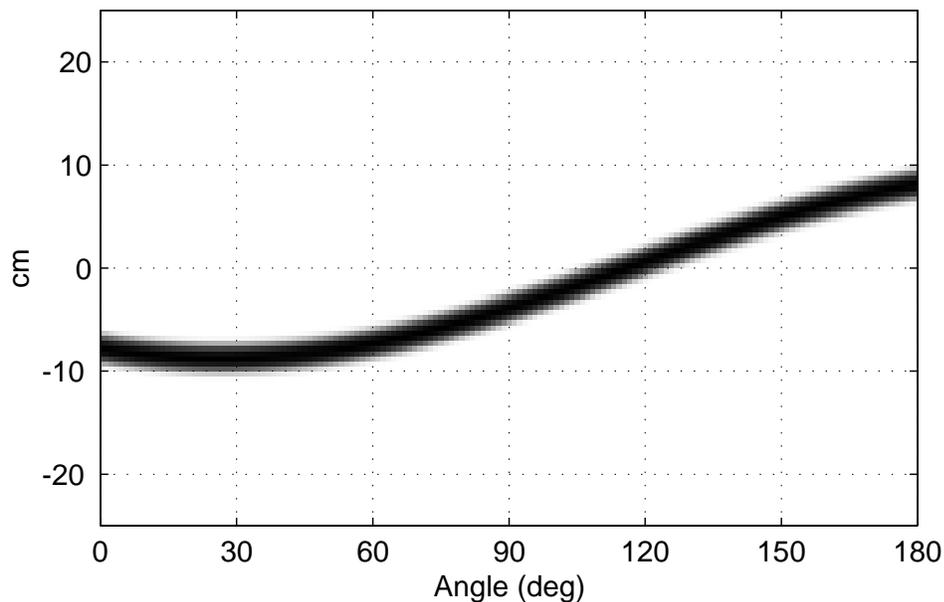


Figure 4.18: Forward Projection of Lesion with $M = 15700$ samples

The total detected counts is $N = 10^6$, with 10^3 counts from the lesion, for a perturbation parameter of $\epsilon = 10^{-3}$. Figure 4.19 shows the discrimination of quantized sinogram measurements with an extrapolated limit of the continuous measurement discrimination overlaid. Figure 4.20 shows the discrimination loss of quantized sinogram measurements with an extrapolated limit of the continuous measurement discrimination overlaid along with the discrimination loss predicted by equation (4.49).

Since these calculation depend on the ratio between the lesion density $p_{\Delta}(\underline{y})$ and the measurement density $p(\underline{y})$, it is dependent on both. By changing the measurement density or even shifting the lesion density, the resulting discrimination loss curve would be offset by a factor proportional to the discriminability 4.44 while the $M^{-2/k}$ trend would remain.

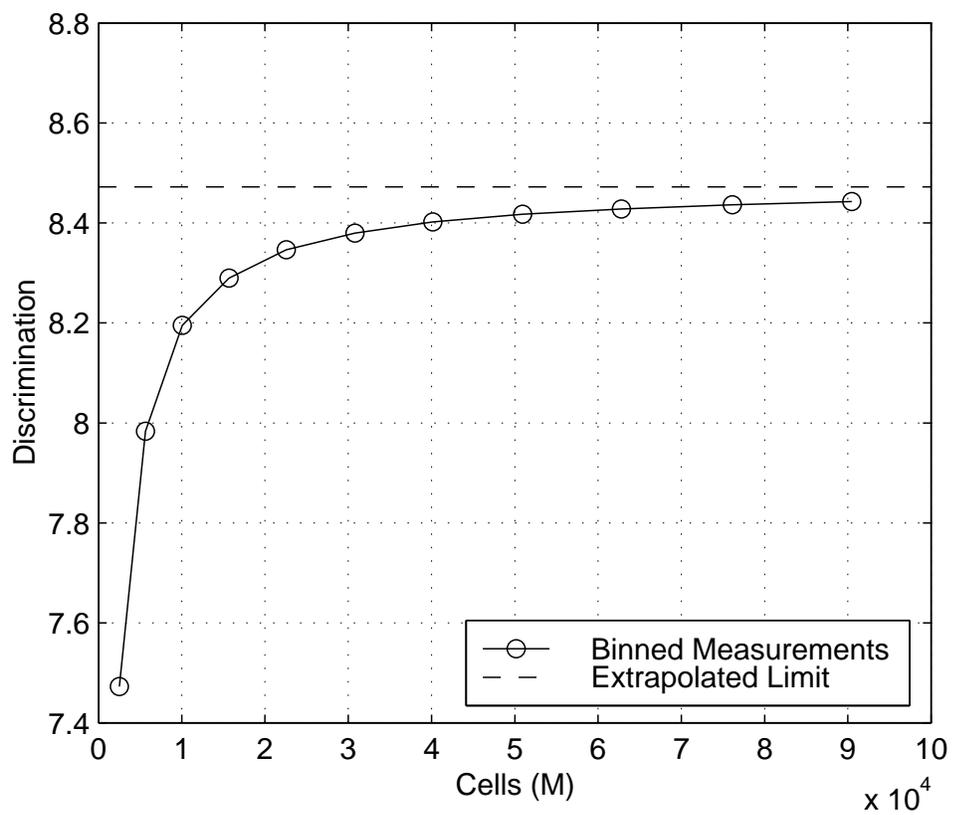


Figure 4.19: Discrimination versus number of sinogram cells

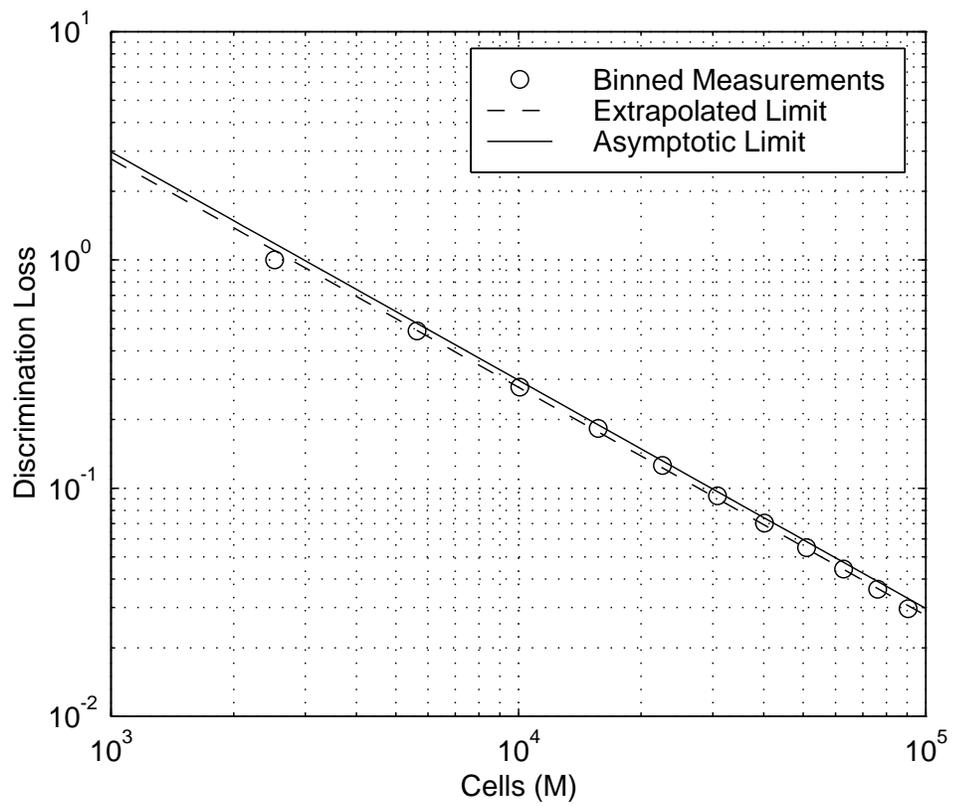


Figure 4.20: Discrimination Loss versus number of sinogram cells

4.6.4 Lattice Quantizer with Support Constraints

One argument against quantization is that at high dimensions, there will be more possible quantization bins than measurements [10]. As noted before though, designing a k -dimensional quantizer as a product code of k different scalar quantizers is extremely wasteful since most of the quantizer codepoints (and associated coding bits) will be wasted on regions of low or zero probability. Figure 4.21 shows a representative 2D PET sinogram sampled with $M = 1450$ bins along with a few thousand random measurement realizations. To a first order, a more efficient method would be

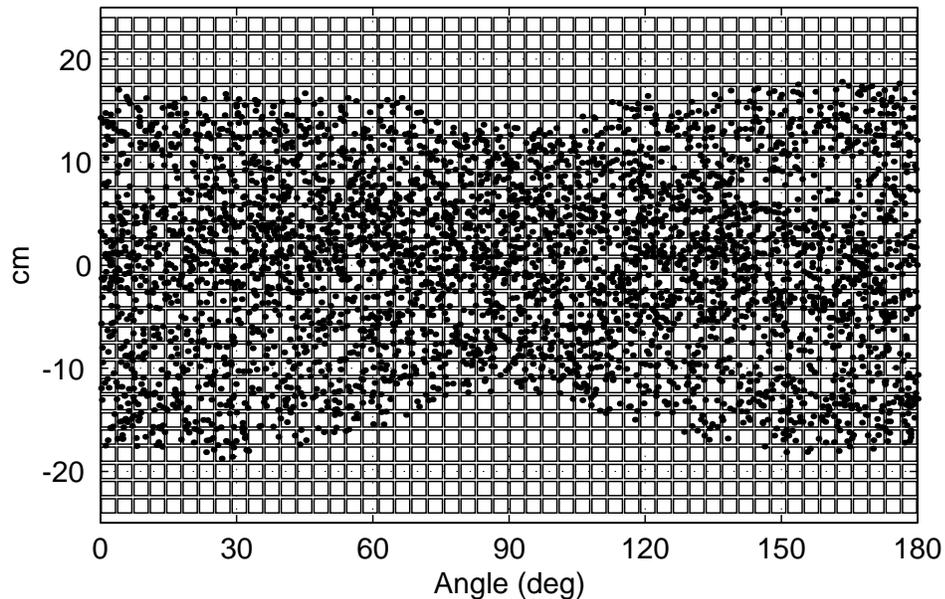


Figure 4.21: Non-adaptive lattice quantization of 2D PET sinogram, $M = 1450$ bins.

to simply quantize the measurements that actually occur and do not bother quantizing regions that do not have any measurements, as shown in Figure 4.22. For this particular case, the total number of non-zero measurement bins is $M = 984$ representing a storage savings of 68%. Alternatively, as shown in Figure 4.23, for a fixed number of bins $M = 1450$ one can use finer quantization³ which results in smaller quantization error and corresponding loss in discriminability. These quantizer support-constraints

³In PET the sampling is determined by the number of detector blocks and ring radius, however as a thought experiment for a system with a continuous measurement space the concept still applies.

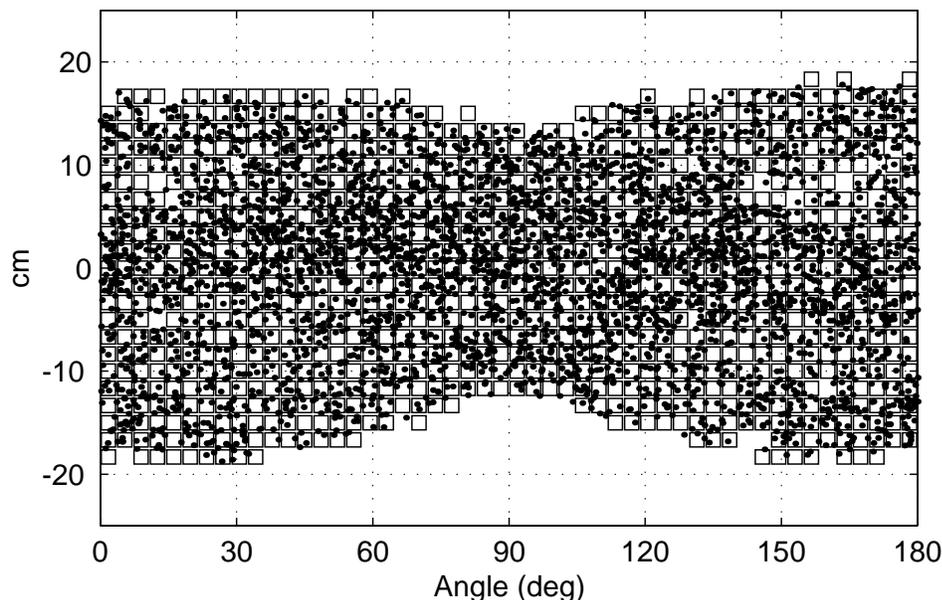


Figure 4.22: Adaptive lattice quantization of 2D PET sinogram, $M = 984$ bins.

are determined by the particular random realization of measurements $\{\underline{Y}_1, \underline{Y}_2, \dots\}$ and are thus *data dependent* and can not be pre-computed and stored ahead of time. However, the volume of the support set as well as the discrimination loss can be estimated directly from the measurements themselves without having to quantize.

A random variable $\underline{y} \sim p(\underline{y})$ can take all values within its support set $S \subset \mathbb{R}^k$, however in general some regions will be “more likely” than others. Realizations of a random variable with a highly localized density function will tend to cluster, while at the other extreme, realizations of a uniform distributed random variable will equally likely take all values in the support set S . More formally, the Asymptotic Equipartition Property [28] from information theory gives a bound on the probability that a random variable will occur within a volume.

The differential entropy $h(\underline{y})$ of the continuous random variable $\underline{y} \sim p(\underline{y})$ with support set $S \subset \mathbb{R}^k$ is defined as

$$h(\underline{y}) = - \int_S p(\underline{y}) \ln p(\underline{y}) d\underline{y} \quad (4.50)$$

Let $\underline{Y}_1, \underline{Y}_2, \dots, \underline{Y}_n$ be n iid realizations of the random variable \underline{y} . From the weak law

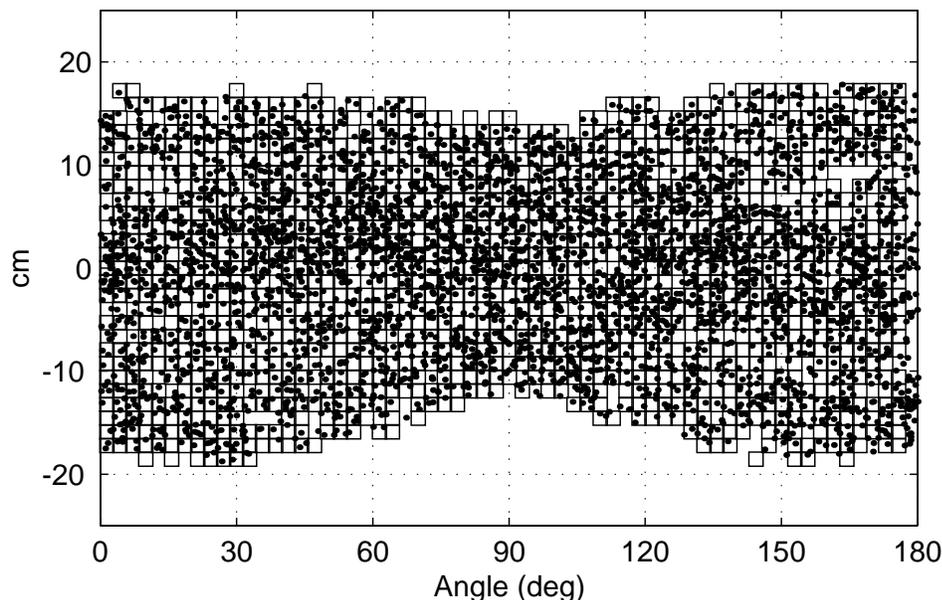


Figure 4.23: Adaptive lattice quantization of 2D PET sinogram, $M = 1450$ bins.

of large numbers

$$-\frac{1}{n} \ln p(\underline{Y}_1, \dots, \underline{Y}_n) \rightarrow E[-\ln p(\underline{y})] \text{ in probability}$$

Define the *typical set* $A_\epsilon^{(n)}$ of the above sequence as

$$A_\epsilon^{(n)} = \left\{ (\underline{y}_1, \dots, \underline{y}_n) \in S^n \text{ s.t. } \left| -\frac{1}{n} \ln p(\underline{y}_1, \dots, \underline{y}_n) - h(\underline{Y}) \right| \leq \epsilon \right\}$$

where $p(\underline{y}_1, \dots, \underline{y}_n) = \prod_{i=1}^n p(\underline{y}_i)$. The volume $V(A_\epsilon^{(n)})$ of the typical set can be interpreted as the “effective volume” of $\mathbb{R}^{k \otimes n}$ that will contain the sequence whose empirical entropy is within ϵ of the true entropy. Random variables with high clustered densities have low entropy, and small typical support volumes, while those with diffuse densities (such as uniform random variables) will have large entropy and correspondingly large typical support volumes. More formally, we can define bounds on the probability of the n realizations occurring within some volume by the following properties [28]

$$\begin{aligned} V(A_\epsilon^{(n)}) &\leq e^{n(h(\underline{y})+\epsilon)}, \\ V(A_\epsilon^{(n)}) &\geq (1-\epsilon)e^{n(h(\underline{y})-\epsilon)}, \end{aligned}$$

$$Pr(A_\epsilon^{(n)}) > 1 - \epsilon.$$

Thus we can bound the volume in $\mathbb{R}^{k \otimes n}$ of the typical set of n realizations to within a confidence of $1 - \epsilon$. By taking the n^{th} -root and manipulating the inequalities, we arrive at the following approximation for the effective volume of the support set S of the random variable \underline{y} ,

$$V(S) \approx e^{h(\underline{y})} \tag{4.51}$$

From this we can derive a measure of the performance gain of support-constrained lattice quantization versus a simple product scalar quantizer (i.e. binning independently along each measurement axis).

Let w_1, \dots, w_k be the range of the components of \underline{y} along each dimension. Then a product-code of uniform scalar quantizer will have a k -dimensional hyper-cube support set S_u with volume $V(S_u) = \prod_{i=1}^k w_i$. The ratio ρ of the effective support set volume $V(S)$ vs. that of a product-code uniform scalar quantizer is given by

$$\rho = \frac{e^{h(\underline{y})}}{V(S_u)}$$

which measures the packing efficiency of using a support constrained lattice vs. an uniform lattice over k -dimensions. Alternatively, for a fixed number of quantization bins M , the average quantization error is reduced by a factor of $\rho^{-2/k}$ by using an adaptive support-constrained lattice quantizer.

A quick numerical calculation of the measurement density $p(\underline{y})$ gives a differential entropy $h(\underline{y}) \sim 4.59$ and estimated support-set volume of $e^{4.59} \sim 98.7$. The unconstrained lattice quantizer (Figure 4.21) is over a support volume of $V(S_u) = \pi \times 50cm$, which results in a packing efficiency of $\rho \sim 63\%$ as compared to the value of $\sim 68\%$ found by taking the ratio of bins in Figures 4.21 and 4.22.

CHAPTER 5

Uniform Cramèr-Rao Bound for Image Reconstruction

The variance of an image, or alternatively the Signal-to-Noise Ratio (SNR), is often used as a measure of image quality. Since image reconstruction problems are ill-posed inverse problems, unbiased estimators result in very noisy, high-variance images. Typically, smoothness constraints and a priori information are used to reduce variance at the cost of biasing the estimator, as illustrated in Figure 5.1. For such

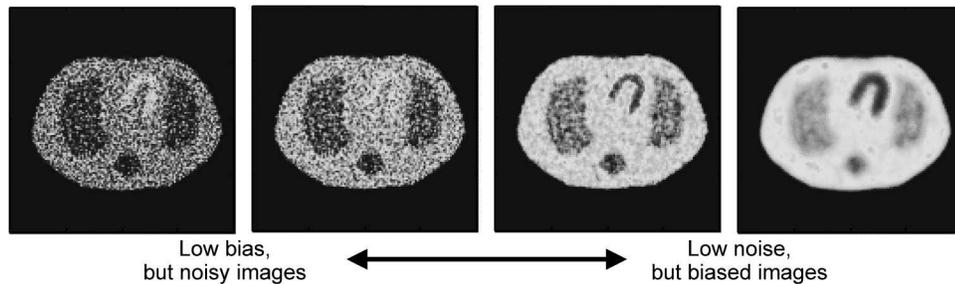


Figure 5.1: Tradeoff between reconstructed image noise and bias.

problems there exists an inherent tradeoff between the recovered spatial resolution of an estimator, overall bias, and its statistical variance; lower variance can only be bought at the price of decreased spatial resolution and/or increased overall bias. For a particular choice of regularized estimator, this tradeoff between bias and variance can be expressed as a parametric curve indexed by the estimator's regularization parameter.

However, total image variance depends not only on the imaging system but also

on the estimator used to create the image. In order to measure of the performance capability of an imaging system, a desired goal is to to decouple the variance due to the inherent limitations of the photon imaging system and the choice of algorithm used to process the data. The classical Cramèr-Rao (CR) bound [28, 103] gives an estimator-independent lower bound on variance and can be used find fundamental limitations of an imaging system. However, it only applies to unbiased estimators. An example bound calculation for an unbiased estimator in an imaging application is given in [108].

The general form of the CR bound for biased estimators (5.5) is not estimator independent, but is a function of the gradient of the estimator bias. Most estimators of interest in image reconstruction problems are biased, and thus have associated with them a bias-gradient vector. The bias-gradient vector is simply a measure of the sensitivity or coupling of a particular estimator’s bias function with respect to variations in the surrounding pixels of the image.

The Uniform Cramèr-Rao bound (UCRB) [57] is a generalization of the CR bound that decouples the variance contributions of the biased estimator and the imaging system. The idea behind the Uniform Cramèr-Rao Bound (UCRB) [57] is that the length or norm of the bias-gradient vector can be viewed as a measure of the total bias-error of an estimator. Among all possible estimators with a given bias-gradient length, there exists an ideal minimum variance estimator whose variance is a lower bound among all possible estimators with that choice of bias-gradient length.

In actual implementation, one does not actually solve for the ideal estimator function. Instead, one solves for the bias-gradient vector of an ideal estimator whose bias-gradient norm is fixed and has minimum variance via a constrained optimization problem. It turns out that this hypothesized ideal estimator’s bias-gradient vector and associated variance are a function of the Fisher Information of the measurements and thus independent of the estimator itself.

As mentioned before, for a particular regularized estimator, this tradeoff between bias and variance can be expressed as a parametric curve indexed by the estimator's regularization parameter. Underneath this parametric curve lies another curve given by the hypothesized ideal minimum variance estimator. The margin between these two curves give a measure of the performance margin between the particular regularized estimator in question, and the fundamental performance limit of the imaging system. For an estimator whose variance lies on this curve, lower estimator variance can only be achieved at the price of an increased estimator bias gradient norm and vice versa.

Hero et al [57] showed that a) the estimator bias gradient norm is an upper bound on the maximal squared variation of the estimator bias function over an ellipsoidal neighborhood, and b) for single pixel estimation, is equivalent to the error between the estimator's local impulse function to a point source and an ideal impulse response. Thus the norm of the estimator bias-gradient would seem to be a natural measure of resolution. However, for single pixel estimation, one can specify a variety of different estimator point response functions that have identical bias-gradient norm but with widely different resolution properties. This has lead to some counter-intuitive results and interpretation difficulties when using the Uniform CR Bound in performance studies of imaging systems. We now extend this tradeoff concept by introducing the 2nd-moment of the point response function as a measure of resolution for single-pixel estimation tasks.

5.1 Cramèr-Rao Bound

Let $\boldsymbol{\theta} = [\theta_1, \dots, \theta_N]^T \in \Theta$ be a column vector of unknown, nonrandom parameters that parameterize the conditional density $p_{\underline{Y}}(\underline{y}|\boldsymbol{\theta})$ of the observed random variable \underline{Y} . The parameter space Θ is assumed to be an open subset of the N -dimensional Euclidean space \mathbb{R}^N . For a fixed $\boldsymbol{\theta}$, let $\hat{\theta}_j = \hat{\theta}_j(\underline{Y})$ be a scalar estimator of the j th

component of $\boldsymbol{\theta}$. Let this estimator have mean value

$$m_{\boldsymbol{\theta}} = E_{\boldsymbol{\theta}}[\hat{\theta}_j], \quad (5.1)$$

bias

$$b_{\boldsymbol{\theta}} = m_{\boldsymbol{\theta}} - \theta_j, \quad (5.2)$$

and variance

$$\sigma_{\hat{\boldsymbol{\theta}}}^2 = E_{\boldsymbol{\theta}}[(\hat{\theta}_j - m_{\boldsymbol{\theta}})^2]. \quad (5.3)$$

In the context of image reconstruction and restoration, \underline{Y} corresponds to a noise and blur degraded measurement of the true image $\boldsymbol{\theta}$, and $\hat{\theta}_j$ is an estimate of the j th pixel of the true image $\underline{\theta}$. Bias $b_{\boldsymbol{\theta}}$ is due to mismatch between the estimation algorithm and truth. Variance $\sigma_{\hat{\boldsymbol{\theta}}}^2$ arises from statistical fluctuations due to statistical uncertainty in the observed data \underline{Y} . Resolution, in our context, is defined as the effective width of the estimation algorithm's local point response.

Note that the scalar estimator $\hat{\theta}_j$ can be expressed in terms of the vector estimator $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}(\underline{Y})$ by the inner-product $\hat{\theta}_j = \mathbf{e}_j^T \hat{\boldsymbol{\theta}}$, where $\mathbf{e}_j = [0, \dots, 1, \dots, 0]^T$ is the j th unit basis vector. Thus the gradient of the estimator bias function $\nabla b_{\boldsymbol{\theta}}$ and the gradient of the estimator mean function $\nabla m_{\boldsymbol{\theta}}$ are related by

$$\nabla m_{\boldsymbol{\theta}} = \nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j. \quad (5.4)$$

For the remainder of this chapter, the gradient $\nabla b_{\boldsymbol{\theta}}$ of the estimator bias-function $b_{\boldsymbol{\theta}}$ will be referred to as the “bias-gradient”, and the gradient $\nabla m_{\boldsymbol{\theta}}$ of the estimator mean-response function $m_{\boldsymbol{\theta}}$ as the “mean-gradient”.

For a biased estimator $\hat{\theta}_j$ of the j th pixel, the Cramèr-Rao Bound of estimator variance is given by

$$\begin{aligned} \sigma_{\hat{\boldsymbol{\theta}}}^2 &\geq \nabla m_{\boldsymbol{\theta}}^T \mathbf{F}_{\underline{Y}}^+ \nabla m_{\boldsymbol{\theta}} \\ &= (\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j)^T \mathbf{F}_{\underline{Y}}^+ (\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j) \end{aligned} \quad (5.5)$$

where the $n \times n$ Fisher Information matrix $\mathbf{F}_{\underline{Y}}$ is given by

$$\mathbf{F}_{\underline{Y}} = E_{\boldsymbol{\theta}} \left[[\nabla_{\boldsymbol{\theta}} \log p_{\underline{Y}}(\underline{Y}|\boldsymbol{\theta})] [\nabla_{\boldsymbol{\theta}} \log p_{\underline{Y}}(\underline{Y}|\boldsymbol{\theta})]^T \right] \quad (5.6)$$

and $\mathbf{F}_{\underline{Y}}^+$ is the Moore-Penrose pseudo-inverse of the (possibly singular) Fisher Information Matrix. Note that if the Fisher Information Matrix is full-rank and the estimator $\hat{\theta}_j$ is unbiased (equivalently, the bias-gradient vector $\nabla b_{\boldsymbol{\theta}} = \mathbf{0}$), then equation (5.5) reduces to the unbiased-estimator CR Bound $\sigma_{\theta}^2 \geq \mathbf{e}_j^T \mathbf{F}_{\underline{Y}}^{-1} \mathbf{e}_j$, i.e. the j th diagonal element of $\mathbf{F}_{\underline{Y}}^{-1}$.

5.2 Bias-gradient

An estimator $\hat{\theta}_j$ whose bias function is constant is as good as unbiased, in that the bias can be removed without knowledge of $\boldsymbol{\theta}$. Additionally, the integral of the bias-gradient specifies the overall bias to within an additive constant. Lastly, the variance lower-bound (5.5) is a function of the bias-gradient. Thus from the standpoint of fundamental tradeoffs, it is the bias-gradient $\nabla b_{\boldsymbol{\theta}}$ that is of interest.

Define the norm $\delta = \|\nabla b_{\boldsymbol{\theta}}\|_{\mathbf{C}}$ of the bias-gradient vector with respect to a positive-definite matrix \mathbf{C} as

$$\delta^2 = \nabla b_{\boldsymbol{\theta}}^T \mathbf{C} \nabla b_{\boldsymbol{\theta}}. \quad (5.7)$$

The norm of the bias-gradient at some point $\mathbf{u} = \boldsymbol{\theta}$ is a measure of the sensitivity of the estimator to changes in \mathbf{u} over some neighborhood about $\boldsymbol{\theta}$. More specifically, the maximum bias variation about some ellipsoidal neighborhood $\mathcal{C}(\boldsymbol{\theta}, \mathbf{C}) = \{\mathbf{u} : (\mathbf{u} - \boldsymbol{\theta})^T \mathbf{C}^{-1} (\mathbf{u} - \boldsymbol{\theta}) \leq 1\}$ centered about $\boldsymbol{\theta}$ is bounded above by the norm of the bias-gradient,

$$\max_{\mathbf{u} \in \mathcal{C}} |b_{\mathbf{u}} - b_{\boldsymbol{\theta}}| \leq \|\nabla b_{\boldsymbol{\theta}}\|_{\mathbf{C}}$$

and with equality when the estimator bias is linear in $\boldsymbol{\theta}$ [57].

5.3 Local Impulse Response and Mean-gradient

As before, let $\hat{\boldsymbol{\theta}}(\underline{Y})$ be a vector estimator of the parameter vector $\boldsymbol{\theta}$, given an observation vector \underline{Y} drawn from the conditional density $p_{\underline{Y}}(\underline{y}|\boldsymbol{\theta})$. Let $\boldsymbol{\mu}(\boldsymbol{\theta})$ be the expected value of the estimator $\hat{\boldsymbol{\theta}}(\underline{Y})$,

$$\begin{aligned}\boldsymbol{\mu}(\boldsymbol{\theta}) &= E_{\boldsymbol{\theta}}[\hat{\boldsymbol{\theta}}(\underline{Y})] \\ &= \int \hat{\boldsymbol{\theta}}(\underline{y})p_{\underline{Y}}(\underline{y}|\boldsymbol{\theta}) d\underline{y}.\end{aligned}\tag{5.8}$$

Let ϵ be a small perturbation in the p th pixel of the source $\boldsymbol{\theta}$. For an estimator with mean $\boldsymbol{\mu}(\boldsymbol{\theta})$, define the local impulse response \mathbf{h} of all reconstructed pixels due to a perturbation in the p th pixel of $\boldsymbol{\theta}$ as

$$\begin{aligned}\mathbf{h}^p(\boldsymbol{\theta}) &= \lim_{\epsilon \rightarrow 0} \frac{\boldsymbol{\mu}(\boldsymbol{\theta} + \epsilon \mathbf{e}_p) - \boldsymbol{\mu}(\boldsymbol{\theta})}{\epsilon} \\ &\triangleq \frac{\partial}{\partial \theta_p} \boldsymbol{\mu}(\boldsymbol{\theta}).\end{aligned}\tag{5.9}$$

As noted in [45], this definition of impulse response reflects the space-varying nature of nonlinear estimators. It is space-varying through its dependence on the perturbing pixel index p and object-dependent through $\boldsymbol{\theta}$. The mean-gradient and local impulse response are related by

$$\begin{aligned}\frac{\partial}{\partial \theta_p} m_{\boldsymbol{\theta}} &= \frac{\partial}{\partial \theta_p} E_{\boldsymbol{\theta}}[\hat{\theta}_j] \\ &= \mathbf{h}_j^p(\boldsymbol{\theta}),\end{aligned}\tag{5.10}$$

or more compactly as

$$\nabla m_{\boldsymbol{\theta}} = [\mathbf{h}_j^1(\boldsymbol{\theta}), \dots, \mathbf{h}_j^N(\boldsymbol{\theta})]^T.\tag{5.11}$$

Whereas the local impulse response describes the coupling to all reconstructed pixels due to a perturbation in a single source pixel, the mean-gradient describes the coupling into a single reconstructed pixel due to perturbations in all source pixels. In general they are not equivalent except under certain conditions [45,57]. Consider the

case of an estimator with mean linear in $\boldsymbol{\theta}$: $\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbf{L}\boldsymbol{\theta}$ for some square matrix \mathbf{L} . The local impulse response $\mathbf{h}^j(\boldsymbol{\theta})$ due to a perturbation in the j th source pixel is the j th column of \mathbf{L} , while the mean-gradient $\nabla m_{\boldsymbol{\theta}} = \nabla E_{\boldsymbol{\theta}}[\hat{\theta}_j]$ is the j th row of \mathbf{L} . Thus, when \mathbf{L} is symmetric, the mean-gradient is equivalent to the local impulse response.

A typical measure of resolution is the full-width half-maximum (FWHM) of the local impulse response. This resolution measure only takes into account the width of the center lobe while sidelobes or slowly decaying tails are ignored, resulting in different local impulse response functions possibly having the same FWHM but widely different resolution. A more discriminating measure of resolution is the 2nd-moment γ of the local impulse response at the j th pixel¹,

$$\gamma^2 = \frac{\sum_{k=1}^N (k-j)^2 (\nabla m_{\boldsymbol{\theta}})_k^2}{\sum_{k=1}^N (\nabla m_{\boldsymbol{\theta}})_k^2}$$

where we assume the mean-gradient and local impulse response are equivalent. Since the mean-gradient is the sum of the ideal response \mathbf{e}_j and the bias-gradient $\nabla b_{\boldsymbol{\theta}}$, the previous equation can be re-written as the ratio of two quadratic forms,

$$\gamma^2 = \frac{(\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j)^T \mathbf{M}_j (\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j)}{(\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j)^T (\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j)} \quad (5.12)$$

or equivalently,

$$(\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j)^T [\mathbf{M}_j - \gamma^2 \mathbf{I}] (\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j) = 0 \quad (5.13)$$

where \mathbf{M}_j is a positive semi-definite diagonal sparse matrix with elements proportional to the square of the distance of each pixel from the j th pixel. For the one-dimensional case, the k^{th} diagonal element of \mathbf{M}_j is given by

$$[\mathbf{M}_j]_{(k,k)} = (k-j)^2 \delta_{kj}$$

where δ_{kj} is (for this expression only) the discrete delta-function.

¹This definition is really the 2nd-moment of the *squared* local impulse response. However it is equivalent to within a constant factor of both the FWHM and 2nd-moment of gaussian- and laplacian-shaped impulse responses, and can be expressed as a quadratic form.

5.4 Constrained CR-Bound

Since the bias-gradient can be viewed as the residual error between the local impulse response and an ideal point response, $\nabla b_{\boldsymbol{\theta}} = \nabla m_{\boldsymbol{\theta}} - \mathbf{e}_j$, the norm of the bias gradient could be interpreted as a measure of resolution since it is the residual error due to side-lobes in the local impulse response. In general this does not hold. In fact, it is possible for local impulse response functions to have identical bias-gradient norm and FWHM yet have widely varying sidelobe behavior, as illustrated in figure 5.2. Thus, the 2nd-moment of the local impulse response has been proposed

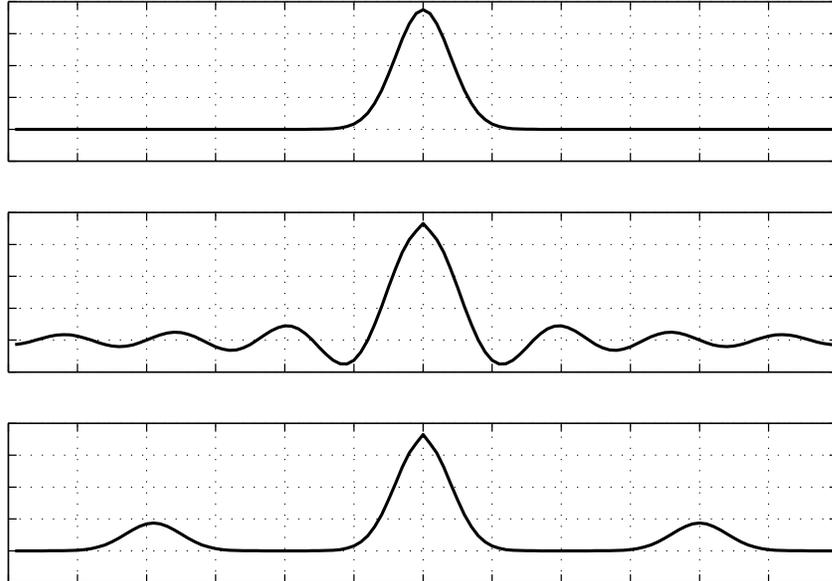


Figure 5.2: Example of three representative point response functions with identical bias-gradient norm and central lobe FWHM, but with varying sidelobe behavior.

as an additional constraint on the estimator bias-gradient [68].

For fixed constraints on allowable estimator bias-gradient norm δ and local impulse response 2nd-moment γ , the variance lower bound among all estimators with bias-gradient satisfying these constraints is found via a constrained optimization of (5.5)

$$\sigma_{\boldsymbol{\theta}}^2 \geq (\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j)^T \mathbf{F}_Y^+ (\nabla b_{\boldsymbol{\theta}} + \mathbf{e}_j) \quad (5.14)$$

$$\geq \min_{\mathbf{d} \in \mathcal{D}} Q(\mathbf{d}), \quad (5.15)$$

where

$$Q(\mathbf{d}) = (\mathbf{d} + \mathbf{e}_j)^T \mathbf{F}_Y^+ (\mathbf{d} + \mathbf{e}_j) \quad (5.16)$$

and the set \mathcal{D} of allowable bias-gradient vectors satisfy the constraints

$$\mathcal{D} = \left\{ \mathbf{d} : \begin{array}{l} \mathbf{d}^T \mathbf{C} \mathbf{d} \leq \delta^2 \\ (\mathbf{d} + \mathbf{e}_j)^T [\mathbf{M}_j - \gamma^2 \mathbf{I}] (\mathbf{d} + \mathbf{e}_j) \leq 0 \end{array} \right\}. \quad (5.17)$$

Figure 5.3 shows a graphical interpretation of this constrained minimization problem. Since the Fisher Information matrix \mathbf{F}_Y is at least positive semi-definite, the scalar function $Q(\mathbf{d})$ in (5.16) are level-sets of a hyper-ellipse centered at \mathbf{e}_j . As for the constraints in (5.17), the bias-gradient norm constraint has the form of an ellipse centered at the origin, while the resolution constraint on the mean-gradient is that of a hyper-cone centered at $-\mathbf{e}_j$.

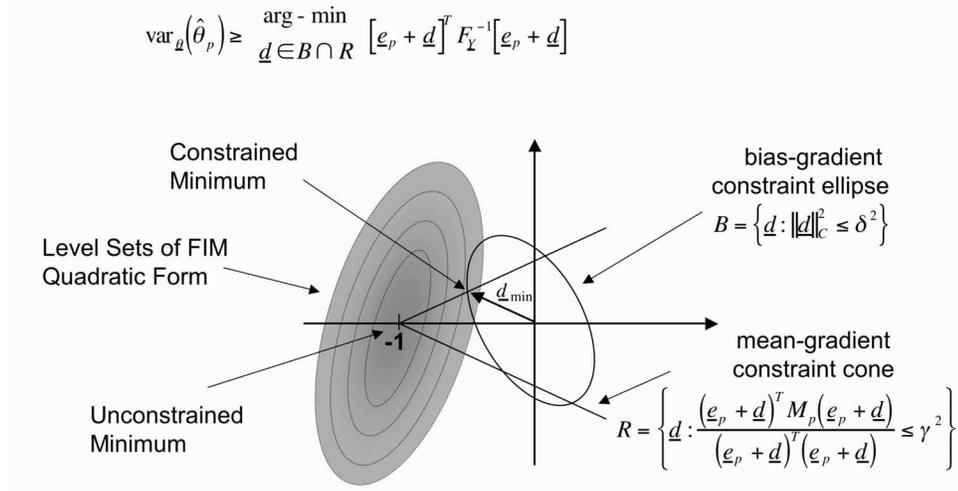


Figure 5.3: Graphical interpretation of the constrained minimization problem involved in calculating the UCRB.

5.4.1 Optimal bias-gradient

The constrained minimization problem

$$\mathbf{d}_{min} = \arg \min_{\mathbf{d} \in \mathcal{D}} (\mathbf{d} + \mathbf{e}_j)^T \mathbf{F}_Y^+ (\mathbf{d} + \mathbf{e}_j)$$

can be solved as an unconstrained minimization by the method of Lagrange multipliers,

$$\begin{aligned} \mathbf{d}_{min} &= \arg \min_{\mathbf{d}} (\mathbf{d} + \mathbf{e}_j)^T \mathbf{F}_{\underline{Y}}^+ (\mathbf{d} + \mathbf{e}_j) \\ &+ \lambda_1 (\mathbf{d}^T \mathbf{C} \mathbf{d} - \delta^2) \\ &+ \lambda_2 ((\mathbf{d} + \mathbf{e}_j)^T [\mathbf{M}_j - \gamma^2 \mathbf{I}] (\mathbf{d} + \mathbf{e}_j)). \end{aligned}$$

After some algebraic manipulations,

$$\begin{aligned} \mathbf{d}_{min} = \arg \min_{\mathbf{d}} & \quad \mathbf{d}^T \left[\mathbf{F}_{\underline{Y}}^+ + \lambda_1 \mathbf{C} + \lambda_2 [\mathbf{M}_j - \gamma^2 \mathbf{I}] \right] \mathbf{d} \\ &+ 2\mathbf{d}^T \left[\mathbf{F}_{\underline{Y}}^+ - \lambda_2 \gamma^2 \mathbf{I} \right] \mathbf{e}_j + \mathbf{e}_j^T \left[\mathbf{F}_{\underline{Y}}^+ - \lambda_2 \gamma^2 \mathbf{I} \right] \mathbf{e}_j - \lambda_1 \delta^2. \end{aligned}$$

Let

$$\begin{aligned} \mathbf{A} &= \left[\mathbf{F}_{\underline{Y}}^+ + \lambda_1 \mathbf{C} + \lambda_2 [\mathbf{M}_j - \gamma^2 \mathbf{I}] \right] \\ \mathbf{B} &= \left[\mathbf{F}_{\underline{Y}}^+ - \lambda_2 \gamma^2 \mathbf{I} \right]. \end{aligned}$$

Then

$$\mathbf{d}_{min} = \arg \min_{\mathbf{d}} \mathbf{d}^T \mathbf{A} \mathbf{d} + 2\mathbf{d}^T \mathbf{B} \mathbf{e}_j + \mathbf{e}_j^T \mathbf{B} \mathbf{e}_j - \lambda_1 \delta^2.$$

Assuming for now that the null-space of $\mathbf{F}_{\underline{Y}}^+$, \mathbf{C} , and $[\mathbf{M}_j - \gamma^2 \mathbf{I}]$ are disjoint, then \mathbf{A}^{-1} exists and we can re-arrange into terms either dependent or independent of \mathbf{d} ,

$$\mathbf{d}_{min} = \arg \min_{\mathbf{d}} \left[\mathbf{d} + \mathbf{A}^{-1} \mathbf{B} \mathbf{e}_j \right]^T \mathbf{A} \left[\mathbf{d} + \mathbf{A}^{-1} \mathbf{B} \mathbf{e}_j \right] + \mathbf{e}_j^T \left[\mathbf{B} - \mathbf{B} \mathbf{A}^{-1} \mathbf{B} \right] \mathbf{e}_j - \lambda_1 \delta^2.$$

Since the 2nd term is independent of \mathbf{d} , the optimal bias-gradient is

$$\begin{aligned} \mathbf{d}_{min} &= -\mathbf{A}^{-1} \mathbf{B} \mathbf{e}_j \\ &= - \left[\mathbf{F}_{\underline{Y}}^+ + \lambda_1 \mathbf{C} + \lambda_2 [\mathbf{M}_j - \gamma^2 \mathbf{I}] \right]^{-1} \left[\mathbf{F}_{\underline{Y}}^+ - \lambda_2 \gamma^2 \mathbf{I} \right] \mathbf{e}_j \quad (5.18) \\ &= \mathbf{d}(\lambda_1, \lambda_2). \end{aligned}$$

By adding a second constraint on the estimator resolution (5.12) to the UCRB, we can now define a minimum-variance *surface* above which that all estimators must

lie. This surface is composed of two regions corresponding to which constraints are active. By analyzing when the two constraints δ , γ are active, it can be shown that the parametric surface defined by (5.5) and (5.18),

$$\sigma^2(\delta, \gamma) = (\mathbf{d}_{min} + \mathbf{e}_j)^T \mathbf{F}_{\underline{Y}}^+ (\mathbf{d}_{min} + \mathbf{e}_j),$$

consists of two regions, where either a) only the bias-gradient constraint is active ($\lambda_1 > 0$, $\lambda_2 = 0$) or b) both constraints are active ($\lambda_1 > 0$, $\lambda_2 > 0$). These two regions are separated by a parametric curve $\sigma^2(\delta, \gamma) = \sigma^2(\delta, \gamma^*)$, where we will refer to γ^* as the *critical* 2nd-moment.

5.4.2 Critical 2nd-moment and minimum-variance surface

When the resolution constraint is not active, the Lagrange multiplier $\lambda_2 = 0$ and \mathbf{d}_{min} reduces to

$$\begin{aligned} \mathbf{d}_{min} &= - \left[\mathbf{F}_{\underline{Y}}^+ + \lambda_1 \mathbf{C} \right]^{-1} \mathbf{F}_{\underline{Y}}^+ \mathbf{e}_j \\ &= \mathbf{d}(\lambda_1), \end{aligned} \tag{5.19}$$

which is the minimum-variance bias-gradient with norm δ given in [57].

For a given bias-gradient norm δ , norm matrix \mathbf{C} and Fisher Information $\mathbf{F}_{\underline{Y}}$, there is a corresponding Lagrange multiplier $\lambda_1 \geq 0$ such that $\delta^2 = \|\mathbf{d}(\lambda_1)\|_{\mathbf{C}}^2$. For this value of λ_1 , define γ^* as the *critical* 2nd-moment constraint,

$$\gamma^{*2} = \frac{(\mathbf{d}(\lambda_1) + \mathbf{e}_j)^T \mathbf{M}_j (\mathbf{d}(\lambda_1) + \mathbf{e}_j)}{(\mathbf{d}(\lambda_1) + \mathbf{e}_j)^T (\mathbf{d}(\lambda_1) + \mathbf{e}_j)}. \tag{5.20}$$

For values of $\gamma \geq \gamma^*$, corresponding to a 2nd-moment greater than that of $\mathbf{d}(\lambda_1)$, only the bias-gradient constraint is active and $\mathbf{d}_{min} = \mathbf{d}(\lambda_1)$ (5.19). Otherwise, for $\gamma < \gamma^*$ the 2nd-moment is less than that of $\mathbf{d}(\lambda_1)$ and both constraints are active with the optimal bias-gradient given by $\mathbf{d}_{min} = \mathbf{d}(\lambda_1, \lambda_2)$ (5.18). The resulting parametric

surface $\sigma^2(\delta, \gamma)$ is given by the following expression:

$$\sigma^2(\delta, \gamma) = \begin{cases} 0 & \delta \geq \delta_{max} \\ (\mathbf{d}(\lambda_1, \lambda_2) + \mathbf{e}_j)^T \mathbf{F}_{\underline{Y}}^+ (\mathbf{d}(\lambda_1, \lambda_2) + \mathbf{e}_j) & 0 < \delta < \delta_{max}, 0 \leq \gamma < \gamma^* \\ (\mathbf{d}(\lambda_1) + \mathbf{e}_j)^T \mathbf{F}_{\underline{Y}}^+ (\mathbf{d}(\lambda_1) + \mathbf{e}_j) & 0 < \delta < \delta_{max}, \gamma \geq \gamma^* \\ \mathbf{e}_j^T \mathbf{F}_{\underline{Y}}^+ \mathbf{e}_j & \delta = 0 \end{cases} \quad (5.21)$$

where $\delta_{max} = [\mathbf{C}]_{(j,j)}$, the j^{th} -diagonal component of the norm matrix \mathbf{C} . Note that the third case is simply the UCRB of [57], and the last case is the unbiased Cramèr-Rao bound [28, 103] since $\delta = 0$ corresponds to an unbiased estimator.

5.5 UCRB for Image Deconvolution

Figure 5.4 shows a 64x64-pixel mandrill test image along with a noise- and blur-degraded simulated measurement. Image blur was simulated by convolving with a 5x5 pixel extent, shift-invariant, 1.5-pixel FWHM symmetric gaussian kernel, and measurement noise was simulated by adding i.i.d. gaussian noise (variance $\sigma^2 = 1$).



Figure 5.4: Original Image (left), Noisy/Blurred Image (right).

The measurement statistics of the observed blur- and noise-degraded image for this case can be expressed as

$$\mathbf{y} = \mathbf{A}\boldsymbol{\theta} + \mathbf{n}, \quad (5.22)$$

where $\boldsymbol{\theta} \in \mathbb{R}^N$ is the true (noiseless) values of the original pixel intensities, \mathbf{y} is the blur- and noise-degraded pixel intensities (both ordered lexicographically as a $N \times 1$ vector), \mathbf{A} is a block-Toeplitz matrix which models the 2D convolution operator, and $\mathbf{n} \sim N(0, \boldsymbol{\Sigma})$ is a vector of zero-mean i.i.d gaussian random noise with covariance $\boldsymbol{\Sigma} = \sigma^2 \mathbf{I}$. Therefore \mathbf{y} is also vector gaussian random variable with mean $\mathbf{A}\boldsymbol{\theta}$ and covariance $\boldsymbol{\Sigma}$, (i.e. $\mathbf{y} \sim N(\mathbf{A}\boldsymbol{\theta}, \boldsymbol{\Sigma})$), and the associated Fisher Information is

$$\begin{aligned} \mathbf{F}_Y &= \mathbf{A}^T \boldsymbol{\Sigma}^{-1} \mathbf{A} \\ &= \frac{1}{\sigma^2} \mathbf{A}^T \mathbf{A} \end{aligned} \quad (5.23)$$

Figure 5.5 shows the limiting square-root variance vs. bias-gradient norm of an estimate of pixel (32,32) in the presence of blur and additive gaussian noise calculated using the expressions in equations (5.6, 5.7, 5.19, and 5.23), otherwise known as the UCRB [57]. Two different cases are considered: a 1.5-pixel FWHM gaussian blur as in figure 5.4, along with a more challenging 1.75-pixel FWHM blur. The bias-gradient norm matrix \mathbf{C} used for the bias-gradient length δ (5.7) was the identity matrix \mathbf{I} . With this choice of norm matrix, the bias-gradient norm δ can be interpreted as the mean-squared estimator bias.

As can be seen in figure 5.5, the minimum-possible variance of pixel intensity when deconvolving a 1.75-pixel FWHM is larger than that achievable with a 1.5-pixel FWHM blur. This is to be expected in that estimating a pixel in the presence of larger blur is a more ill-conditioned problem, and would result in a noisier estimate for a given total bias.

Figure (5.6) shows two different estimator mean-gradient images for the 1.75-pixel FWHM blur case with bias-gradient norm $\delta = 0.1$ (left) and $\delta = 0.5$ (right). The mean-gradient is more spread out with increasing bias-gradient norm, which gives substance to the interpretation of bias-gradient norm as a measure of resolution.

One problem with the bias-gradient norm as a measure of estimator resolution is that it is possible for estimators with different mean-gradients to have the exact same

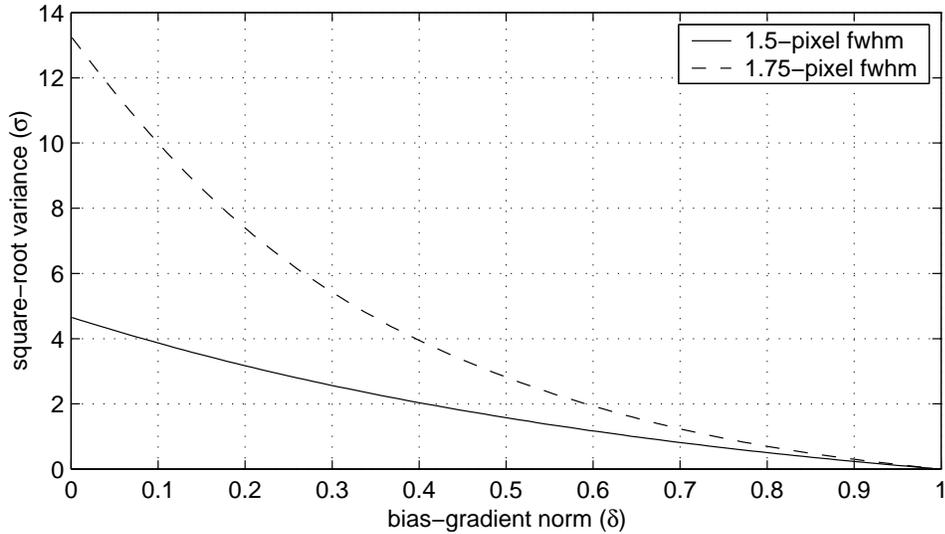


Figure 5.5: Bounds on single pixel estimation performance in presence of blur and additive gaussian noise.

bias-gradient norm, but with dramatically different resolution properties. Figure 5.7 shows cross-sectional slices through the mean-gradient of two representative estimators. Their associated bias-gradients both have the same norm $\delta = 0.5$, but their spread or FWHM are obviously different. By adding a second constraint on the estimator resolution (5.12) to the UCRB, we can now define a minimum-variance *surface* (5.21). Figure 5.8 illustrates this surface for the deconvolution problem presented earlier in figure 5.4.

The variance of a regularized estimator will follow a trajectory in (δ, γ) indexed by its regularization parameter. By analyzing the distance the particular estimator lies above the surface, one can determine how far from optimality the estimator is. For the purposes of this dissertation, we will focus on the penalized weighted least-squares (PWLS) estimator, mainly due to its analytical tractability.

5.6 Penalized weighted least-squares estimator

For the linear additive gaussian noise model in (5.22), the penalized weighted least-squares estimator (PWLS) for a measurement realization \mathbf{Y} is given by the

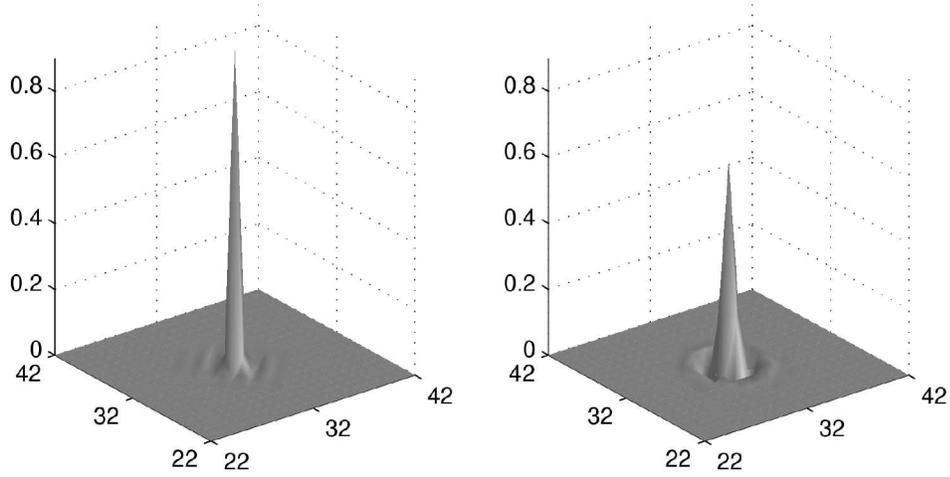


Figure 5.6: Mean-gradients images for two different estimators with different bias-gradient norm. $\delta = 0.1$ (left), $\delta = 0.5$ (right).

minimizer of the following objective function,

$$\Phi(\boldsymbol{\theta}) = [\mathbf{Y} - \mathbf{A}\boldsymbol{\theta}]^T \boldsymbol{\Sigma}^{-1} [\mathbf{Y} - \mathbf{A}\boldsymbol{\theta}] + \beta \boldsymbol{\theta}^T \mathbf{P} \boldsymbol{\theta} \quad (5.24)$$

where $\beta > 0$ is a regularization parameter, and \mathbf{P} is a positive definite penalty matrix.

The value of $\boldsymbol{\theta}$ that minimizes Φ (5.24) can be found explicitly,

$$\begin{aligned} \hat{\boldsymbol{\theta}}(\mathbf{Y}) &= [\mathbf{A}^T \boldsymbol{\Sigma}^{-1} \mathbf{A} + \beta \mathbf{P}]^{-1} \mathbf{A}^T \boldsymbol{\Sigma}^{-1} \mathbf{Y} \\ &= [\mathbf{F}_{\mathbf{Y}} + \beta \mathbf{P}]^{-1} \mathbf{A}^T \boldsymbol{\Sigma}^{-1} \mathbf{Y}. \end{aligned} \quad (5.25)$$

The penalty matrix \mathbf{P} serves to improve the numerical stability of the matrix inversion $[\mathbf{F}_{\mathbf{Y}} + \beta \mathbf{P}]^{-1}$ in (5.25) for systems with an ill-conditioned Fisher Information matrix. Common choices for \mathbf{P} include the identity matrix \mathbf{I} (often referred to as Tikonov regularization or energy penalty), as well as discrete approximations to the Laplacian operator to enforce smoothness constraints [40, 55].

In (5.25), it was assumed that that matrix \mathbf{A} and noise covariance $\boldsymbol{\Sigma}$ were known exactly. In general this will not be true, but rather only estimates $\tilde{\mathbf{A}}$ and $\tilde{\boldsymbol{\Sigma}}$. The PWLS estimator with these assumed values is then

$$\begin{aligned} \hat{\boldsymbol{\theta}}(\mathbf{Y}) &= [\tilde{\mathbf{A}}^T \tilde{\boldsymbol{\Sigma}}^{-1} \tilde{\mathbf{A}} + \beta \mathbf{P}]^{-1} \tilde{\mathbf{A}}^T \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{Y} \\ &= [\tilde{\mathbf{F}}_{\mathbf{Y}} + \beta \mathbf{P}]^{-1} \tilde{\mathbf{A}}^T \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{Y}. \end{aligned} \quad (5.26)$$

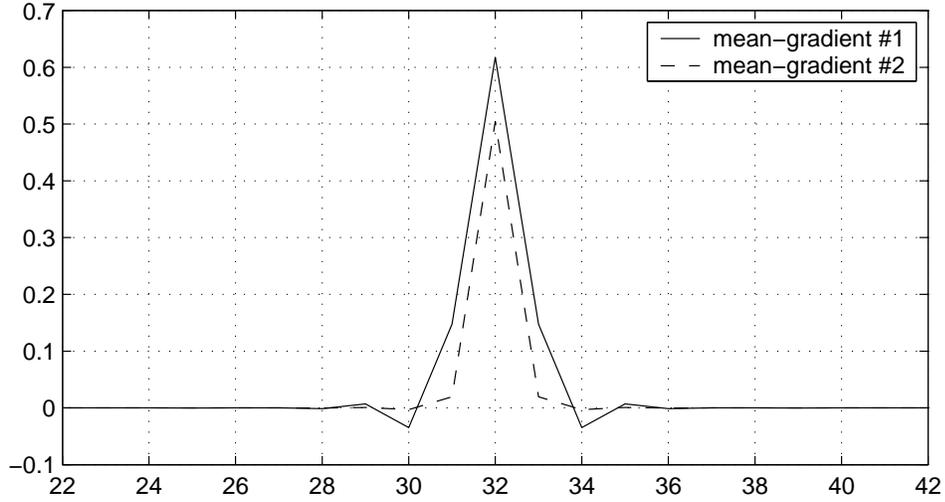


Figure 5.7: Mean-gradient cross-sections for two different estimators with fixed bias-gradient norm $\delta = 0.5$.

where $\tilde{\mathbf{F}}_{\underline{Y}}$ is the Fisher information matrix assumed by the estimator from its imperfect knowledge of the matrix \mathbf{A} and noise covariance Σ .

From the definitions given in (5.1, 5.2, 5.4, and 5.26), the PWLS estimator mean-gradient ∇m_{θ} and bias-gradient ∇b_{θ} function can be easily derived, resulting in a mean-gradient

$$\nabla m_{\theta} = \mathbf{A}^T \tilde{\Sigma}^{-1} \tilde{\mathbf{A}} [\tilde{\mathbf{F}}_{\underline{Y}} + \beta \mathbf{P}]^{-1} \mathbf{e}_j, \quad (5.27)$$

bias gradient

$$\begin{aligned} \nabla b_{\theta} &= \nabla m_{\theta} - \mathbf{e}_j \\ &= \left[\mathbf{A}^T \tilde{\Sigma}^{-1} \tilde{\mathbf{A}} [\tilde{\mathbf{F}}_{\underline{Y}} + \beta \mathbf{P}]^{-1} - \mathbf{I} \right] \mathbf{e}_j, \end{aligned} \quad (5.28)$$

and estimator variance found by the quadratic form in (5.5).

5.6.1 Matched PWLS estimator

In order to analyze the performance of a PWLS estimator whose parameters are matched to the data, i.e. $\tilde{\mathbf{A}} = \mathbf{A}$ and $\tilde{\Sigma} = \Sigma$, we calculated the estimator trajectory for both an identity penalty and a smoothing penalty. As can be seen in figure 5.9,

there is not a significant gap between the actual estimator variance and the bound surface for either choice of penalty. Upon closer inspection, the PWLS estimator with identity penalty ($\mathbf{P} = \mathbf{I}$) is found to *achieve* the bound. This is a property of PWLS estimators under gaussian statistics. When the choice of bias-gradient norm matrix is equal to the inverse of the penalty $\mathbf{C} = \mathbf{P}^{-1}$, the resulting trajectory achieves the bound [57].

5.6.2 Mis-matched PWLS estimator

To analyze the effect of when the estimator is mis-matched from the data, we calculated the estimator trajectory for an estimator with $\tilde{\mathbf{A}}$ corresponding to a 1.75pixel FWHM blur. Since the original image was convolved with a 1.5pixel FWHM blur, this corresponds to an estimator that is *over compensating*. As can be seen in figures 5.10 and 5.11, there is a significant gap between the actual estimator variance and the bound surface. Only after significant regularization do the estimator trajectories approach the bound surface, but at the cost of increased bias and resolution loss. In neither case do the estimator trajectories achieve the bound, which is due to the mis-match in the assumed value of \mathbf{A} . The difference between the estimator trajectories due to using different penalties can be seen in figure 5.12.

In figure 5.13, we calculated the estimator trajectory for both an identity- and smoothing-penalized estimator with $\tilde{\mathbf{A}}$ corresponding to a 1.5pixel FWHM blur, but for an original image that was convolved with a 1.75pixel FWHM blur. This corresponds to an estimator that is *under compensating*. Even with no regularization ($\beta = 0$) the estimator is biased, which seems logical because of the fundamental mis-match between the estimator's choice and actual value of \mathbf{A} .

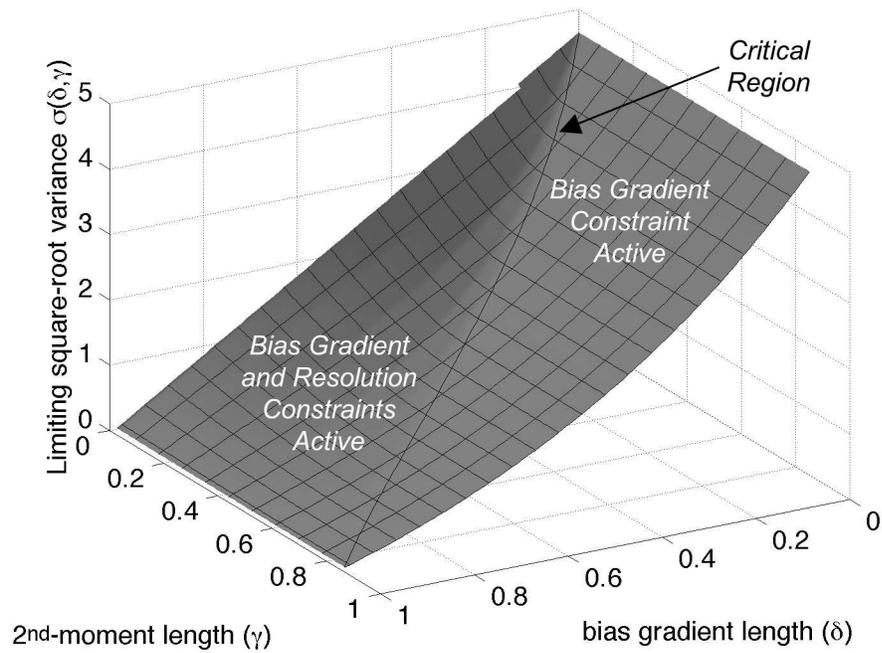


Figure 5.8: Minimum-variance surface, critical 2nd-moment, and active constraint regions.

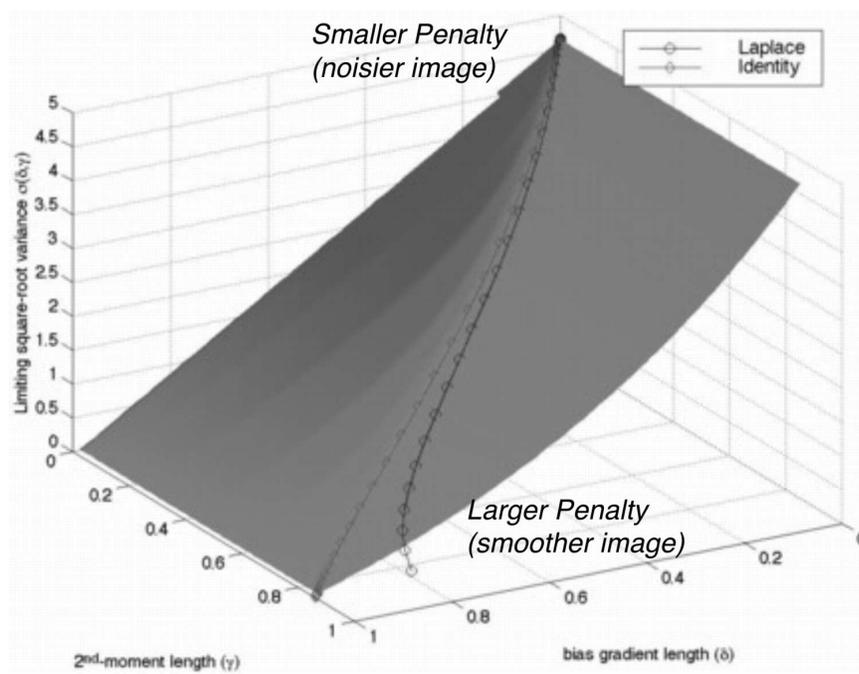


Figure 5.9: PWLS estimator trajectories for both an identity and smoothing penalty. Both the measurements and estimator are matched with a 1.5pixel FWHM blur.

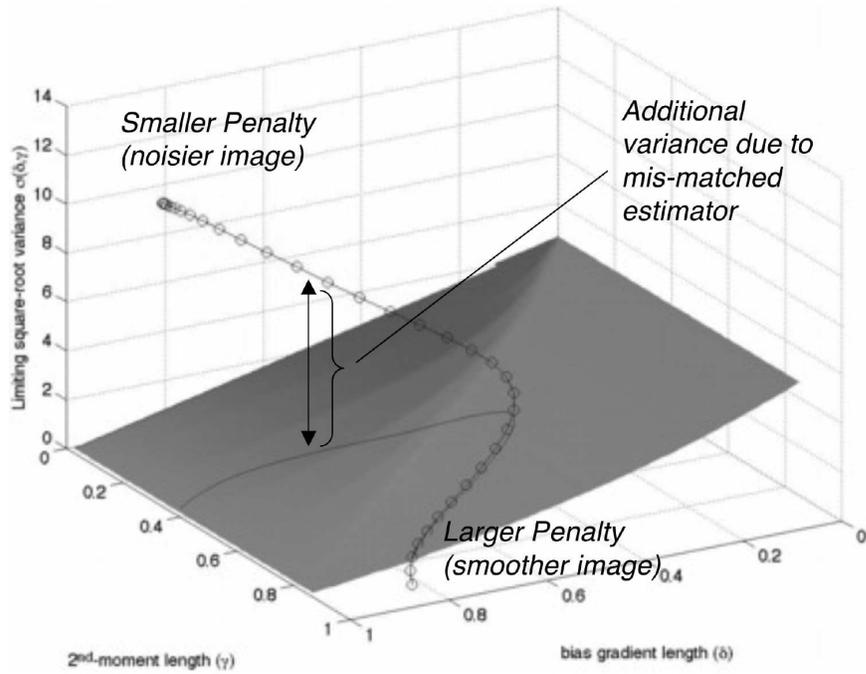


Figure 5.10: PWLS estimator trajectory for a smoothing penalty. The measurements had a 1.5pixel FWHM blur, while the estimator assumed a 1.75pixel FWHM blur.

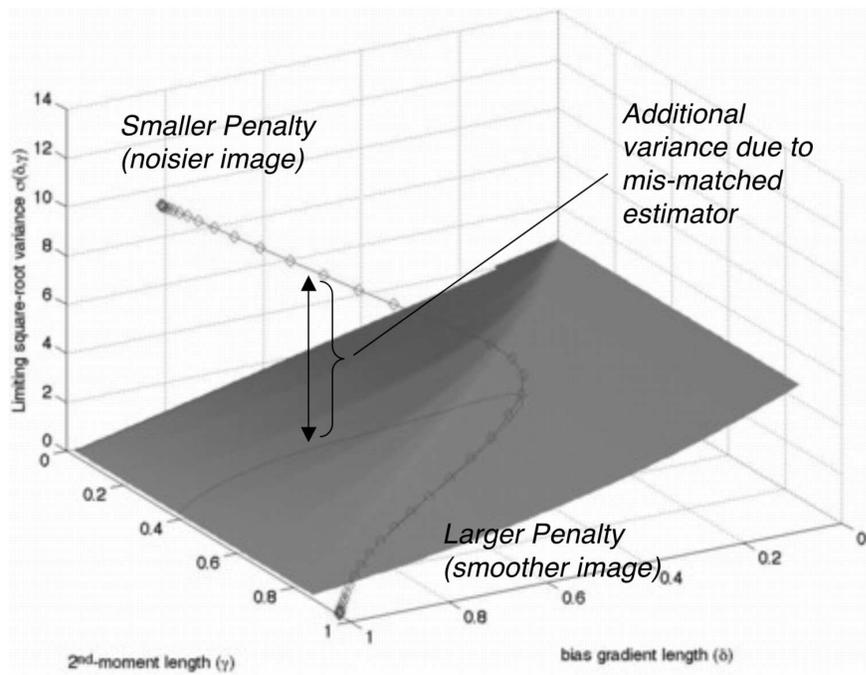


Figure 5.11: PWLS estimator trajectory for an identity penalty. The measurements had a 1.5pixel FWHM blur, while the estimator assumed a 1.75pixel FWHM blur.

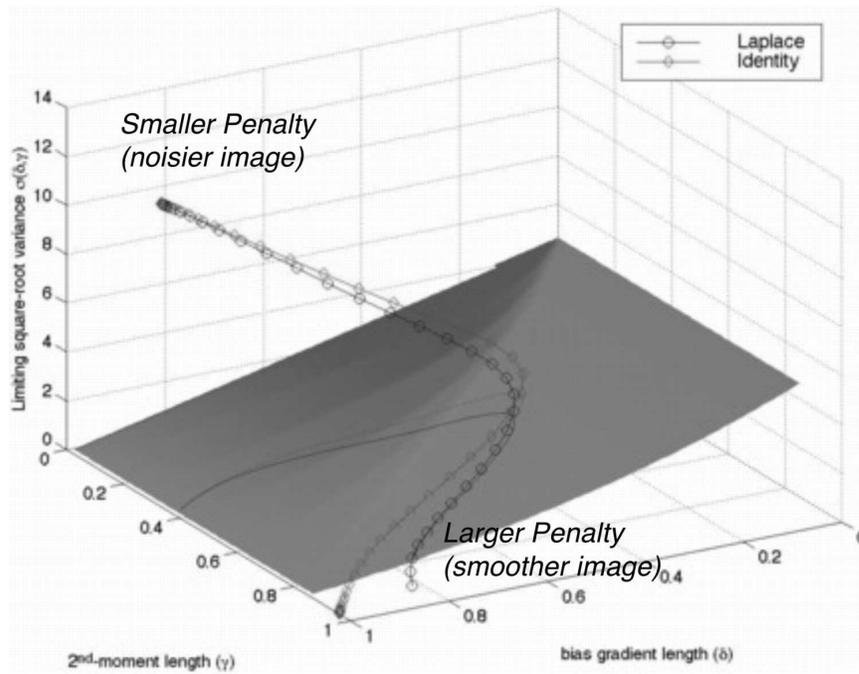


Figure 5.12: PWLS estimator trajectory for both an identity and smoothing penalty. The measurements had a 1.5pixel FWHM blur, while the estimator assumed a 1.75pixel FWHM blur.

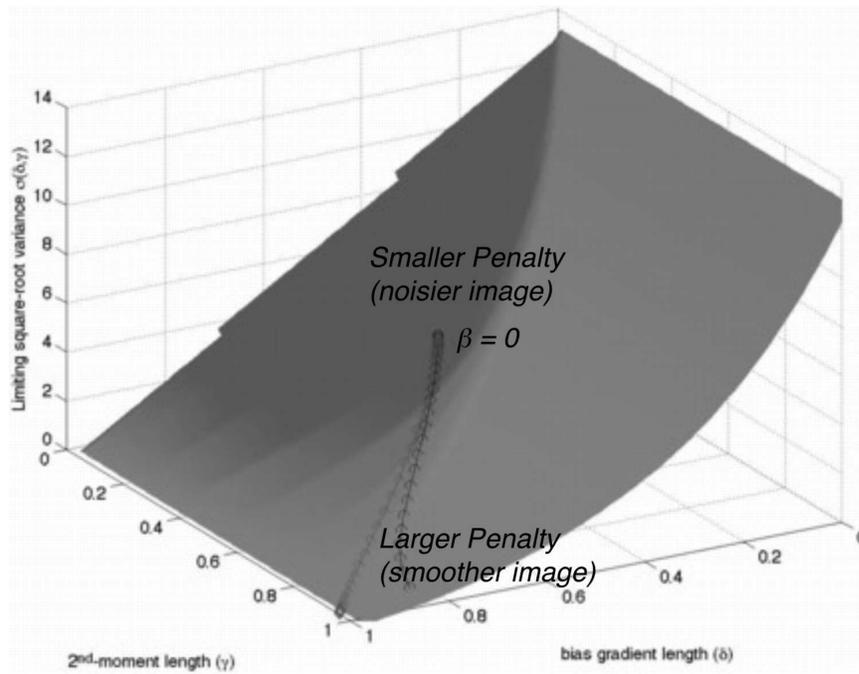


Figure 5.13: PWLS estimator trajectory for both an identity and smoothing penalty. The measurements had a 1.75pixel FWHM blur, while the estimator assumed a 1.5pixel FWHM blur.

CHAPTER 6

Conclusions

The theme of this dissertation has been the development of techniques for analyzing emission tomography systems with multi-dimensional measurements. In Chapter 2, the relationship between list-mode and quantized measurements was explicitly defined. Measurement binning, whether by hardware or software, was shown to be realistically modeled as a finite partitioning of the multi-dimensional continuous measurement space, and naturally falls into the framework of vector quantization.

In Chapter 3, a computationally tractable statistical model for Compton SPECT was developed that included the dominant error sources of Doppler broadening, energy resolution, and attenuation in a consistent manner. However, since with list-mode measurements each detection contributes independently to the likelihood function, the computational complexity of the maximum-likelihood estimator is linear in the number of measurements. Thus, in Chapter 4 vector quantizers were investigated as a possible method to minimize the reconstruction algorithm computation complexity. The vector quantizer was shown to be a valid model for discrete measurements in emission tomography systems. Traditionally, asymptotic expressions for distortion in a vector quantizer have been used to analyze the mean-squared error in the measurement domain due to quantization in the measurement domain. For this dissertation, we extended the analysis of loss in the Kullback-Leibler discrimination function due to quantization [54], as well as relating the loss to the asymptotic behavior of the

Receiver Operating Characteristic (ROC) curve, for a low-contrast lesion detection task.

Chapter 5 was concerned with the limiting resolution and variance of an emission tomography system. Once a partitioning of the list-mode measurements is defined, either by hardware configuration or software based binning, one can define the discrete likelihood function and associated Fisher Information Matrix. It is indirectly through the FIM that the quantizer partitioning limits the reconstructed resolution as well as limiting variance. One can solve for an ideal hypothetical estimator that has minimum variance for a given amount of total bias and resolution. This results in a parametric surface over bias, resolution and variance which specifies a lower bound on any estimator satisfying bias, resolution or variance constraints. This surface is a function of the Fisher information, bias-gradient norm, and resolution parameter and thus independent of any particular choice of estimation algorithm. The variance of any given estimator traces out a trajectory, indexed by its regularization parameter, that is constrained to lie above this minimum-variance surface. The margin between this trajectory and surface gives the performance margin between the particular regularized estimator in question and the fundamental performance limit of the imaging system. For an estimator whose variance lies on this surface, lower estimator variance can only be achieved at the price of increased estimator bias and/or decreased resolution.

Future Work

For Compton SPECT system modeling, future areas of work would include the following: In order to derive expressions for the survival probabilities and list-mode measurement back-projections given in Chapter 3, it was assumed that attenuation was negligible in order to arrive at tractable solutions for the survival probabilities and back-projections. However, under certain system geometries, such as those with

internal first detectors, this assumption is not valid and practical implementations to account for attenuation must be developed.

The binned measurement transition probabilities a_{ij} were shown to be the integral over the measurement bin of the list-mode transition probability. However, evaluating this integral is nontrivial and must be approximated by numerical methods. A possible fast method of integrating over each bin volume is by the following Monte Carlo technique: Given p list-mode measurement realizations, the p back-projections over each of the N pixels can be calculated once and selectively summed over each of the M bins. The rationale behind this is that for a fixed computational burden, the variance due to statistical fluctuations in a Monte Carlo based technique decays faster than the deterministic error due to multi-dimensional numerical integration, as the number of measurements p grows large. The resulting reconstruction algorithm would be $O(M)$ rather than $O(p)$ per iteration, with a one-time computational hit of initially calculating p back-projections. This assumes that the M bin locations have been determined ahead of time.

Another open question is the optimal bin distribution or measurement partitioning to minimize the Uniform CR Bound. As mentioned before in Chapter 4, the measurement partitioning that minimizes the MSE in the quantized measurements is known. However, it says nothing about the reconstructed image limiting variance, or of the distortion in the mean image due to under-sampling. One possible distortion metric motivated by optimal test design is to minimize the determinant or trace of the inverse Fisher Information \mathbf{F}^{-1} . However, this is an intractable expression when there are more than a handful of parameters.

As for lower bounds on estimator variance, a major difficulty in calculating bound curves for imaging applications is that they involve inverting a large system of equations involving the Fisher Information Matrix. For systems with localized impulse responses (and hence diagonally-dominant system matrices), the resulting FIM is

sparse and diagonally-dominant, with the system of equations easily solved for using matrix iteration techniques such as Gauss-Siedel. However, for tomography the impulse response is highly non-local, resulting in a non-sparse, non-diagonally dominate, poorly conditioned FIM that is difficult to invert. Thus, pre-conditioning the system of equations with a pre-conditioner matrix tuned explicitly for a tomographic system would aid in calculating bound curves.

APPENDICES

APPENDIX A

List-mode ML-EM Estimator

Maximizing (2.25) for the maximum-likelihood solution vector $\hat{\boldsymbol{\lambda}}$ directly is impractical due to the coupling inside the log function. As an alternative, we will fall back on the old trick of embedding the incomplete observed data into a larger, unobserved complete dataspace, derive the log-likelihood function as if we had access to the complete data, and then iteratively maximize the resulting surrogate function.

Our observed data $\{\underline{Y}_i\}_{i=1}^p$ is incomplete in that it is not known which voxel j emitted the photon that resulted in the measurement \underline{Y}_i . Define the random variable z_i as follows:

$$z_i = \begin{cases} j & \text{if photon emitted from voxel } j \text{ resulted in measurement } \underline{Y}_i \\ 0 & \text{otherwise} \end{cases}$$

The collection $\{\underline{Y}_i, z_i\}_{i=1}^p$ forms our unobserved complete dataspace. Assuming that we have access to the complete data and that each observation is iid, the joint density of the complete data is

$$\begin{aligned} & p(\underline{y}_1, z_1, \dots, \underline{y}_p, z_p, p | \boldsymbol{\lambda}, D, T) \\ &= \prod_{i=1}^p p(\underline{y}_i | z_i, \boldsymbol{\lambda}, D) p(z_i | \boldsymbol{\lambda}, D) P(p | \boldsymbol{\lambda}, D, T), \end{aligned}$$

with associated log-likelihood function of the complete data $l_c(\boldsymbol{\lambda})$

$$l_c(\boldsymbol{\lambda}) = \sum_{i=1}^p \log (p(\underline{Y}_i | z_i, \boldsymbol{\lambda}, D) p(z_i | \boldsymbol{\lambda}, D)) + \log P(p | T, \boldsymbol{\lambda}, D).$$

Taking the expected value of the log-likelihood with respect to z_i , conditioned on a fixed value $\boldsymbol{\lambda} = \hat{\boldsymbol{\lambda}}^k$ and on detecting the observed data sequence $\{\underline{Y}_i\}_{i=1}^p$,

$$E_{z_i} \left[l_c(\boldsymbol{\lambda}) | \hat{\boldsymbol{\lambda}}^k, \{\underline{Y}_i\} \right] = \sum_{j=1}^N P(z_i = j | \hat{\boldsymbol{\lambda}}^k, D, \{\underline{Y}_i\}) \sum_{i=1}^p \log(p(\underline{Y}_i | \boldsymbol{\lambda}, D, j) p(j | \boldsymbol{\lambda}, D)) - (T\lambda_D) + p \log(T\lambda_D) - \log(p!).$$

From the definition of z_i , the term $P(z_i = j | \hat{\boldsymbol{\lambda}}^k, D, \{\underline{Y}_i\}_{i=1}^p)$ can be brought into the inner summation since

$$P(z_i = j | \hat{\boldsymbol{\lambda}}^k, D, \{\underline{Y}_i\}_{i=1}^p) = \begin{cases} P(j | \hat{\boldsymbol{\lambda}}^k, D, \underline{Y}_i) & \text{if } z_i = j \\ 0 & \text{otherwise} \end{cases},$$

while Bayes' rule and a little algebra results in

$$P(j | \hat{\boldsymbol{\lambda}}^k, D, \underline{Y}_i) = \frac{a_j(\underline{Y}_i) \hat{\lambda}_j^k}{\sum_{l=1}^N a_l(\underline{Y}_i) \hat{\lambda}_l^k}$$

and

$$p(\underline{Y}_i | \boldsymbol{\lambda}, D, j) P(j | \boldsymbol{\lambda}, D) = \frac{a_j(\underline{Y}_i) \lambda_j}{\lambda_D}.$$

Plugging it all together gives the following expression ,

$$E_{z_i} \left[l_c(\boldsymbol{\lambda}) | \hat{\boldsymbol{\lambda}}^k, \{\underline{Y}_i\}_{i=1}^p \right] = \sum_{i=1}^M \sum_{j=1}^N \left[\frac{a_j(\underline{Y}_i) \hat{\lambda}_j^k}{\sum_{l=1}^N a_l(\underline{Y}_i) \hat{\lambda}_l^k} \log(T a_j(\underline{Y}_i) \lambda_j) \right] - (T\lambda_D) - \log(p!).$$

Maximizing with respect to the components of $\boldsymbol{\lambda}$ results in the set of list-mode ML-EM update equations,

$$\hat{\lambda}_j^{k+1} = \frac{\hat{\lambda}_j^k}{T s_j} \sum_{i=1}^p \left[\frac{a_j(\underline{Y}_i)}{\sum_{l=1}^N a_l(\underline{Y}_i) \hat{\lambda}_l^k} \right].$$

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*, Dover, New York, 1964.
- [2] E. Agrell and T. Eriksson, “Optimization of lattices for quantization,” *IEEE Tr. Info. Theory*, vol. 44, no. 5, pp. 1814–28, September 1998.
- [3] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, “Image coding using wavelet transform,” *IEEE Tr. Im. Proc.*, vol. 1, no. 2, pp. 205–20, April 1992.
- [4] E. Asma, T. E. Nichols, J. Qi, and R. M. Leahy, “4D pet image reconstruction from list mode data,” in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, 2000.
- [5] M. Balakrishnan, W. Pearlman, and L. Lu, “Variable-rate tree-structured vector quantizers,” *IEEE Tr. Info. Theory*, vol. 41, no. 4, pp. 917–930, July 1995.
- [6] H. H. Barrett, “Objective assessment of image quality: effects of quantum noise and object variability,” *J. Opt. Soc. Am. A*, vol. 7, no. 7, pp. 1266–1278, July 1990.
- [7] H. H. Barrett, C. K. Abbey, and E. Clarkson, “Objective assessment of image quality III: ROC metrics, ideal observers and likelihood-generating functions,” *J. Opt. Soc. Am. A*, vol. 15, no. 6, pp. 1520–35, June 1998.
- [8] H. H. Barrett, J. L. Denny, R. F. Wagner, and K. J. Myers, “Objective assessment of image quality. II. Fisher information, Fourier crosstalk, and figures of merit for task performance,” *J. Opt. Soc. Am. A*, vol. 12, no. 5, pp. 834–52, May 1995.
- [9] H. H. Barrett and W. Swindell, *Radiological Imaging: the Theory of Image Formation, Detection, and Processing*, Academic, New York, 1981.
- [10] H. H. Barrett, T. White, and L. C. Parra, “List-mode likelihood,” *J. Opt. Soc. Am. A*, vol. 14, no. 1, pp. 2914–23, November 1997.
- [11] R. Basko, G. L. Zeng, and G. T. Gullberg, “Application of spherical harmonics to image reconstruction for Compton camera,” in *Proc. of the 1997 Intl. Mtg. on Fully 3D Image Recon.*, 1997.

- [12] G. Benitz and J. Bucklew, "Asymptotically optimal quantizers for detection of I.I.D. data," *IEEE Tr. Info. Theory*, vol. 35, no. 2, pp. 316–25, March 1989.
- [13] J. Beutel, H. L. Kundel, and R. L. Van Metter, editors, *Handbook of Medical Imaging*, volume 1. Physics and Psychophysics, SPIE, Washington, 2000.
- [14] F. Biggs, L. B. Mendelsohn, and J. B. Mann, "Hartree-Fock Compton profiles for the elements," *Atomic Data and Nuclear Data Tables*, vol. 16, no. 3, pp. 201–309, September 1975.
- [15] R. Blahut, *Principles and Practise of Information Theory*, Addison-Wesley, 1987.
- [16] C. A. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *IEEE Tr. Im. Proc.*, vol. 5, no. 3, pp. 480–92, March 1996.
- [17] J. D. Bronzino, editor, *The Biomedical Engineering Handbook*, CRC, Florida, 1995.
- [18] J. A. Browne and A. R. D. Pierro, "A row-action alternative to the EM algorithm for maximizing likelihoods in emission tomography," *IEEE Tr. Med. Im.*, vol. 15, no. 5, pp. 687–99, October 1996.
- [19] D. Brusa, G. Stutz, J. A. Riveros, J. M. Fernández, and F. Salvat, "Fast sampling algorithms for the simulation of photon Compton scattering," *Nucl. Instr. Meth. Phys. Res. A.*, vol. 379, pp. 167–175, March 1996.
- [20] C. Byrne, E. Soares, T. S. Pan, S. Glick, M. King, and A. Del-Guerra, "Accelerating the EM algorithm using rescaled block-iterative methods," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pp. 1752–6, 1996.
- [21] C. Caini and A. Vanelli-Coralli, "Optimal bit allocation in subband coding with nonideal reconstruction filters," *IEEE Signal Proc. Letters*, vol. 8, no. 6, pp. 157–9, June 2001.
- [22] P. Chou, T. Lookabaugh, and R. Gray, "Optimal pruning with applications to tree-structured source coding and modeling," *IEEE Tr. Info. Theory*, vol. 35, no. 2, pp. 299–315, March 1989.
- [23] N. H. Clinthorne, C.-Y. Ng, C.-H. Hua, J. E. Gormley, J. W. Leblanc, S. J. Wilderman, and W. L. Rogers, "Theoretical performance comparison of a Compton-scatter aperture and parallel-hole collimator," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, 1996.
- [24] J. Conway and N. Sloane, "Voronoi regions of lattices, second moments of polytopes, and quantization," *IEEE Tr. Info. Theory*, vol. 28, no. 2, pp. 211–232, March 1982.

- [25] J. Conway and N. Sloane, "A fast encoding method for lattice codes and quantizers," *IEEE Tr. Info. Theory*, vol. 29, no. 6, pp. 820–824, November 1983.
- [26] J. H. Conway and N. J. A. Sloane, *Sphere Packing, Lattices, and Groups*, Springer-Verlag, New York, 2 edition, 1993.
- [27] P. C. Cosman, R. M. Gray, and R. A. Olshen, "Evaluating quality of compressed medical images: SNR, subjective rating, and diagnostic accuracy," *Proc. IEEE*, vol. 82, no. 6, pp. 919–32, June 1994.
- [28] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley, New York, 1991.
- [29] A. R. De Pierro and M. E. B. Yamagishi, "Fast EM-like methods for maximum 'a posteriori' estimates in emission tomography," *IEEE Tr. Med. Im.*, vol. 20, no. 4, pp. 280–8, April 2001.
- [30] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Royal Stat. Soc. Ser. B*, vol. 39, no. 1, pp. 1–38, 1977.
- [31] T. A. Driscoll, K.-C. Toh, and L. N. Trefethen, "From potential theory to matrix iterations in six steps," *SIAM Review*, vol. 40, no. 3, pp. 547–78, September 1998.
- [32] W. T. Eadie, D. Drijard, F. E. James, M. Roos, and B. Sadoulet, *Statistical Methods in Experimental Physics*, North Holland, Amsterdam, 1971.
- [33] P. Eisenberger and W. A. Reed, "Gamma-ray Compton scattering: Experimental Compton profiles for He, N₂, Ar, and Kr," *Phys. Rev. A.*, vol. 5, no. 5, pp. 2085–2093, May 1972.
- [34] H. Erdogan, G. Gualtieri, and J. A. Fessler, "An ordered subsets algorithm for transmission tomography," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, 1998. Inadvertently omitted from proceedings. Available from web page.
- [35] A. Eskicioglu and P. Fisher, "Image quality measures and their performance," *IEEE Trans. Commun.*, vol. 43, no. 12, pp. 2959–65, December 1995.
- [36] P. C. et al, "Tree-structured vector quantization of ct chest scans: image quality and diagnostic accuracy," *IEEE Tr. Med. Im.*, vol. 12, no. 4, pp. 727–39, December 1993.
- [37] S. P. et al, "Tree-structured vector quantization with region-based classification," in *Proceedings, IEEE Asilomar Conf on Signals, Systems, and Computers*, October 1992.
- [38] V. V. Federov, *Theory of Optimal Experiments*, Academic Press, New York, 1972.

- [39] L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone beam algorithm," *J. Opt. Soc. Am. A*, vol. 1, no. 6, pp. 612–9, 1984.
- [40] J. A. Fessler, "Penalized weighted least-squares image reconstruction for positron emission tomography," *IEEE Tr. Med. Im.*, vol. 13, no. 2, pp. 290–300, June 1994.
- [41] J. A. Fessler, "Statistical image reconstruction methods for transmission tomography," in *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis*, M. Sonka and J. M. Fitzpatrick, editors, pp. 1–70, SPIE, Bellingham, 2000.
- [42] J. A. Fessler and H. Erdogan, "A paraboloidal surrogates algorithm for convergent penalized-likelihood emission image reconstruction," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pp. 1132–5, 1998.
- [43] J. A. Fessler and A. O. Hero, "Space-alternating generalized expectation-maximization algorithm," *IEEE Tr. Sig. Proc.*, vol. 42, no. 10, pp. 2664–77, October 1994.
- [44] J. A. Fessler and A. O. Hero, "Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms," *IEEE Tr. Im. Proc.*, vol. 4, no. 10, pp. 1417–29, October 1995.
- [45] J. A. Fessler and W. L. Rogers, "Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs," *IEEE Tr. Im. Proc.*, vol. 5, no. 9, pp. 1346–58, September 1996.
- [46] W. Gander and W. Gautschi, "Adaptive quadrature - revisited," Technical Report 306, Institut für Wissenschaftliches Rechnen, Eidgenössische Technische Hochschule, CH-8092 Zürich, August 1998.
- [47] W. Gardner and B. D. Rao, "Theoretical analysis of the high-rate vector quantization of LPC parameters," *IEEE Tr. Speech and Audio Proc.*, vol. 3, no. 5, pp. 367–381, September 1995.
- [48] A. Gersho, "Asymptotically optimal block quantization," *IEEE Tr. Info. Theory*, vol. 25, no. 4, pp. 373–80, July 1979.
- [49] A. Gersho and R. Gray, *Vector Quantization and Signal Compression*, Kluwer, Norwell, MA, 1998.
- [50] S. J. Glick and E. J. Soares, "Noise characteristics of SPECT iterative reconstruction with a mis-matched projector-backprojector pair," *IEEE Tr. Nuc. Sci.*, vol. 45, no. 4, pp. 2183–8, August 1998.
- [51] J. R. Goldschneider and E. A. Riskin, "Optimal bit allocation and best-basis selection for wavelet packets and TSVQ," *IEEE Tr. Im. Proc.*, vol. 8, no. 9, pp. 1305–9, September 1999.

- [52] R. Gray and D. Neuhoff, "Quantization," *IEEE Trans. Inform. Theory*, vol. 44, no. 6, pp. 2325–2383, October 1998.
- [53] R. Gupta and A. O. Hero, "Performance limits of hypothesis testing from vector-quantized data," in *Proc. IEEE Int. Symp. Info. Theory*, June 2001.
- [54] R. Gupta, *Quantization Strategies for Low-Power Communications*, PhD thesis, University of Michigan, Ann Arbor, MI, 1999.
- [55] T. Hebert and R. Leahy, "A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors," *IEEE Tr. Med. Im.*, vol. 8, no. 2, pp. 194–202, June 1989.
- [56] T. Hebert, R. Leahy, and M. Singh, "Three-dimensional maximum-likelihood reconstruction for an electronically collimated single-photon-emission imaging system," *J. Opt. Soc. Am. A*, vol. 7, no. 7, pp. 1305–13, July 1990.
- [57] A. O. Hero, J. A. Fessler, and M. Usman, "Exploring estimator bias-variance tradeoffs using the uniform CR bound," *IEEE Tr. Sig. Proc.*, vol. 44, no. 8, pp. 2026–41, August 1996.
- [58] C. Ho Hua, *Compton Imaging System Development and Performance Assessment*, PhD thesis, University of Michigan, Ann Arbor, MI, August 2000.
- [59] Y. L. Hsieh, G. T. Gullberg, G. L. Zeng, and R. H. Huesman, "Image reconstruction using a generalized natural pixel basis," *IEEE Tr. Nuc. Sci.*, vol. 43, no. 4, pp. 2306–19, August 1996.
- [60] C. Hua, S. J. Wilderman, T. J. Kragh, N. H. Clinthorne, and W. L. Rogers, "The roles of observer models in evaluating Compton camera performance," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, October 1999.
- [61] C. H. Hua, N. H. Clinthorne, J. W. Leblanc, and W. L. Rogers, "Decoding penalty calculation for a ring Compton camera using uniform Cramèr-Rao bound," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, November 1998.
- [62] H. M. Hudson and R. S. Larkin, "Accelerated image reconstruction using ordered subsets of projection data," *IEEE Tr. Med. Im.*, vol. 13, no. 4, pp. 601–9, December 1994.
- [63] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*, IEEE Press, New York, 1988.
- [64] C. Kamphuis, F. J. Beekman, P. P. van Rijk, and M. A. Viergever, "Dual matrix ordered subsets reconstruction for accelerated 3D scatter compensation in single-photon emission tomography," *Eur. J. Nuc. Med.*, vol. 25, no. 1, pp. 8–18, January 1998.

- [65] G. Kanbach, R. Andritschke, P. F. Bloser, F. Schopper, V. Schönfelder, and A. Zoglauer, “Concept study for the next generation medium energy gamma-ray astronomy mission - MEGA,” in *Proc. SPIE, Astronomical Telescopes and Instrumentation*, August 2002.
- [66] J. Kiefer, “General equivalence theory for optimum designs (approximate theory),” *Annals of Statistics*, vol. 2, no. 5, pp. 849–879, September 1974.
- [67] G. F. Knoll, *Radiation Detection and Measurement*, Wiley, New York, 2 edition, 1989.
- [68] T. J. Kragh and A. O. Hero, “Bias-resolution-variance tradeoffs for single pixel estimation tasks using the uniform Cramer-Rao bound,” in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, 2000.
- [69] T. J. Kragh and A. O. Hero, “Image resolution-variance tradeoffs using the uniform Cramèr-Rao bound,” in *Proc. Euro. Sig. Proc. Conf. (EUSIPCO)*, 2002. submitted.
- [70] K. Lange and R. Carson, “EM reconstruction algorithms for emission and transmission tomography,” *J. Comp. Assisted Tomo.*, vol. 8, no. 2, pp. 306–316, April 1984.
- [71] J. W. LeBlanc, *A Compton Camera for Low Energy Gamma Ray Imaging in Nuclear Medicine Applications*, PhD thesis, University of Michigan, Ann Arbor, MI, 1999.
- [72] J. W. LeBlanc, N. H. Clinthorne, C.-H. Hua, E. Nygard, W. L. Rogers, D. K. Wehe, P. Weilhammer, and S. J. W. and, “C-SPRINT: a prototype Compton camera system for low energy gamma ray imaging,” *IEEE Tr. Nuc. Sci.*, vol. 46, no. 3, pp. 943–9, June 1998.
- [73] J. Li, N. Chaddha, and R. Gray, “Multiresolution tree structured vector quantization,” in *Proceedings, IEEE Asilomar Conf on Signals, Systems, and Computers*, November 1996.
- [74] J. Li, N. Chaddha, and R. Gray, “Asymptotic performance of vector quantizers with a perceptual distortion measure,” *IEEE Tr. Info. Theory*, vol. 45, no. 4, pp. 1082–91, May 1999.
- [75] Y. Linde, A. Buzo, and R. Gray, “An algorithm for vector quantizer design,” *IEEE Trans. Commun.*, vol. 26, pp. 840–847, June 1978.
- [76] T. Linder and R. Zamir, “High-resolution source coding for non-difference distortion measures: multidimensional companding,” *IEEE Tr. Info. Theory*, vol. 45, no. 2, pp. 548–61, March 1999.

- [77] T. Linder and R. Zamir, “High-resolution source coding for non-difference distortion measures: the rate-distortion function,” *IEEE Tr. Info. Theory*, vol. 45, no. 2, pp. 533–47, March 1999.
- [78] A. Macovski, *Medical Imaging Systems*, Prentice-Hall, New Jersey, 1983.
- [79] K. J. Myers, “Ideal observer models of visual signal detection,” in *Handbook of Medical Imaging, Volume 1. Physics and Pyschophysics*, J. Beutel, H. L. Kundel, and R. L. Van Metter, editors, pp. 558–92, SPIE, Bellingham, 2000.
- [80] S. Na and D. Neuhoff, “Bennett’s integral for vector quantizers,” *IEEE Trans. Inform. Theory*, vol. 41, no. 4, , July 1995.
- [81] C.-Y. Ng, N. H. Clinthorne, J. A. Fessler, A. O. Hero, Y. Zhang, and W. L. Rogers, “Preliminary studies on the feasibility of addition of vertex view to conventional brain SPECT,” in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pp. 1057–61, 1996.
- [82] J. A. O’Sullivan, R. E. Blahut, and D. L. Snyder, “Information-theoretic image formation,” *IEEE Tr. Info. Theory*, vol. 44, no. 6, pp. 2094–123, October 1998.
- [83] L. C. Parra, “Reconstruction of cone-beam projections from Compton scattered data,” *IEEE Tr. Nuc. Sci.*, vol. 47, no. 4, pp. 1543–50, August 2000.
- [84] L. C. Parra and H. H. Barrett, “List-mode likelihood - EM algorithm and noise estimation demonstrated on 2D-PET,” *IEEE Tr. Med. Im.*, vol. 17, no. 2, pp. 228–35, 1998.
- [85] B. Picinbono and P. Duvaut, “Optimum quantization for detection,” *IEEE Tr. Info. Theory*, vol. 36, no. 11, pp. 1254–58, November 1988.
- [86] H. V. Poor, “Fine quantization in signal detection and estimation,” *IEEE Tr. Info. Theory*, vol. 34, no. 5, pp. 960–972, September 1988.
- [87] W. A. Reed and P. Eisenberger, “Gamma-ray Compton profiles of diamond, silicon, and germanium,” *Phys. Rev. B.*, vol. 6, no. 12, pp. 4596–4604, December 1972.
- [88] R. Ribberfors, “Relationship of the relativistic compton cross section to the momentum distribution of bound electron states,” *Phys. Rev. B.*, vol. 12, no. 6, pp. 2067–2074, September 1975.
- [89] E. A. Riskin, “Optimal bit allocation via the generalized BFOS algorithm,” *IEEE Tr. Info. Theory*, vol. 37, no. 2, pp. 400–2, March 1991.
- [90] W. L. Rogers. *Personal communication*, July 2002.

- [91] A. Sauve, A. O. Hero, W. L. Rogers, and N. H. Clinthorne, "System modeling and spatial sampling techniques for simplification of transition matrix in 3D electronically collimated SPECT," in *Proc. of the 1997 Intl. Mtg. on Fully 3D Image Recon.*, 1997.
- [92] K. Sayood, J. Gibson, and M. Rost, "An algorithm for uniform vector quantizer design," *IEEE Tr. Info. Theory*, vol. 30, no. 6, pp. 805–814, November 1984.
- [93] W. P. Segars, *Development of a new dynamic NURBS-based cardiac-torso (NCAT) phantom*, PhD thesis, The University of North Carolina, Chapel Hill, NC, May 2001.
- [94] L. A. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Tr. Med. Im.*, vol. 1, no. 2, pp. 113–22, October 1982.
- [95] M. Singh, "An electronically collimated gamma camera for single photon emission computed tomography. Part I: Theoretical considerations and design criteria," *Med. Phys.*, vol. 10, no. 4, pp. 421–427, July 1983.
- [96] M. Singh and D. Doria, "An electronically collimated gamma camera for single photon emission computed tomography. Part II: Image reconstruction and preliminary experimental measurements," *Med. Phys.*, vol. 10, no. 4, pp. 428–435, July 1983.
- [97] D. L. Snyder and M. I. Miller, *Random Point Processes in Time and Space*, Springer-Verlag, New York, 1991.
- [98] D. L. Snyder, D. G. Politte, and M. I. Miller, "A case study in statistical image processing: Positron-emission tomography," in *Spatial Statistics and Imaging*, A. Possolo, editor, volume 20 of *Lecture Notes*, pp. 368–81, IMS, 1990.
- [99] E. J. Soares, C. L. Byrne, and S. J. G. and, "Noise characterization of block-iterative reconstruction algorithms. I. Theory," *IEEE Tr. Med. Im.*, vol. 19, no. 4, pp. 261–70, April 2000.
- [100] J. W. Stayman and J. A. Fessler, "Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction," *IEEE Tr. Med. Im.*, vol. 19, no. 6, pp. 601–15, June 2000.
- [101] R. W. Todd, J. M. Nightingale, and D. B. Everett, "A proposed gamma camera," *Nature*, vol. 251, no. 5471, pp. 132–4, September 1974.
- [102] R. Van de Walle, H. H. Barrett, K. J. Myers, M. I. A. I. Lemahieu, and A. F. Gmitro, "Reconstruction of MR images from data acquired on a general non-regular grid by pseudoinverse calculation," *IEEE Tr. Med. Im.*, vol. 19, no. 12, pp. 1160–7, December 2000.
- [103] H. L. Van Trees, *Detection, Estimation, and Modulation Theory*, Wiley, New York, 1968.

- [104] Y. Vardi, L. A. Shepp, and L. Kaufman, "A statistical model for positron emission tomography," *J. Am. Stat. Ass.*, vol. 80, no. 389, pp. 8–37, March 1985.
- [105] Y. M. Wang and S. Nagarkar, "Locators and sensors for automated coordinate checking fixtures," Technical Report T.R. 97-75, Dept. of Mech. Eng., Univ. of Maryland, College Park, MD, September 1997.
- [106] S. J. Wilderman, N. H. Clinthorne, J. A. Fessler, and W. L. Rogers, "Fast algorithm for list mode back-projection of Compton scatter camera data," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, November 1997.
- [107] S. J. Wilderman, *Vectorized algorithms for Monte Carlo simulation of kilovolt electron and photon transport (electron transport)*, PhD thesis, University of Michigan, Ann Arbor, MI, 1990.
- [108] K. A. Winick, "Cramer-Rao lower bounds on the performance of charge-coupled-device optical position estimators," *J. Opt. Soc. Am. A*, vol. 3, no. 11, pp. 1809–15, November 1986.
- [109] R. Zamir and M. Feder, "On lattice quantization noise," *IEEE Tr. Info. Theory*, vol. 42, no. 4, pp. 1152–1159, July 1996.
- [110] G. L. Zeng and G. T. Gullberg, "Unmatched projector/backprojector pairs in an iterative reconstruction algorithm," *IEEE Tr. Med. Im.*, vol. 19, no. 5, pp. 548–55, May 2000.
- [111] L. Zhang, S. J. Wilderman, N. H. Clinthorne, and W. L. Rogers, "An anthropomorphic phantom integrated EGS4 Monte Carlo code and its application in Compton probe," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, October 2000.