

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

U·M·I

University Microfilms International
A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
313/761-4700 800/521-0600

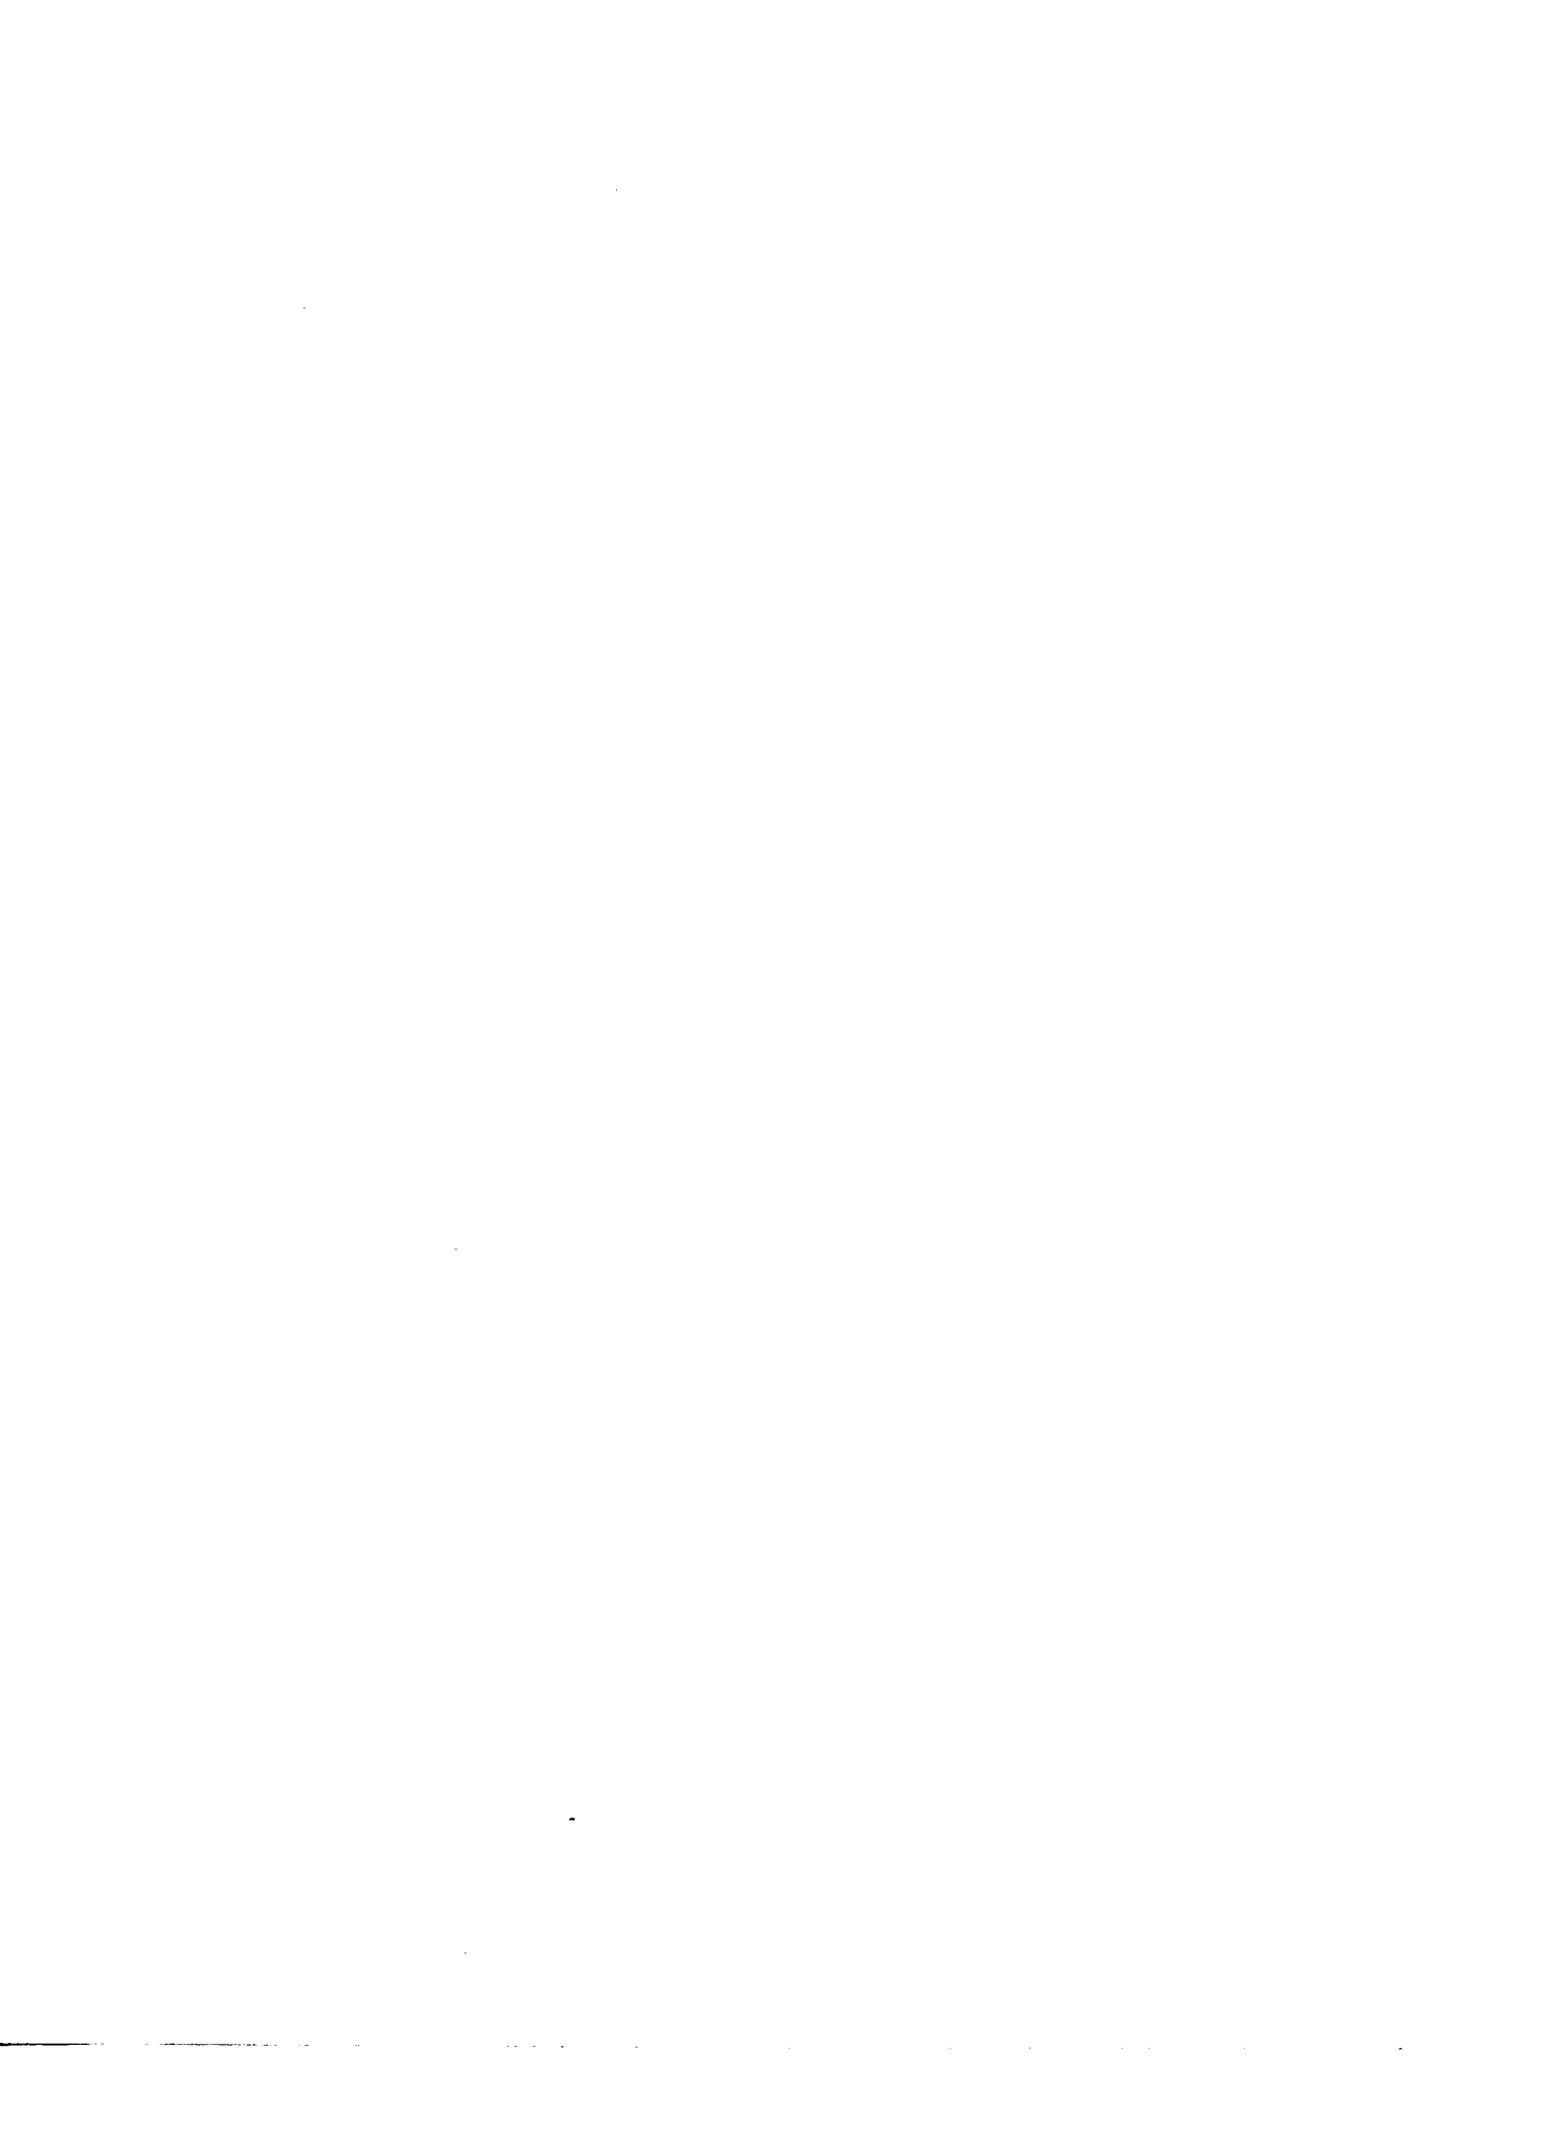
Order Number 1353137

**Expectation maximization methods for processing SPECT
images**

Marcotte, Hope Ann, M.S.

The University of Arizona, 1993

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



**EXPECTATION MAXIMIZATION METHODS
FOR PROCESSING SPECT IMAGES**

by

Hope Ann Marcotte

A Thesis Submitted to the Faculty of the
COMMITTEE ON OPTICAL SCIENCES
In Partial Fulfillment of the Requirements
For the Degree of
MASTER OF SCIENCE
In the Graduate College
THE UNIVERSITY OF ARIZONA

1993

ACKNOWLEDGEMENTS

I wish to thank the many people whose support, inspiration, and humor helped me to complete this work. Dr. Harry Barrett continuously guided my work while giving me the freedom to discover what interested me. I deeply appreciate his enthusiastic and conscientious help in all matters. I am also grateful to the talented individuals among the staff, students and faculty, many of whom helped me in specific ways. Don Richards wrote the software to collect, compress and process the brain imager data, and offered great advice on how to improve my code. John Sain, Bill Klein, Irene Pang and Jyh Chen were essential for their help calibrating the cameras and collecting data. Bob Rowe aided me with his enthusiasm and knowledge and saved me the trauma of learning PDOS by writing the code to collect the super-MDRF. Tim White answered countless questions and asked others that help direct my research. Michel Rogulski's simulations, suggestions, and general disposition were tremendously valuable. Thanks to all whose involvement in the development of the brain imager made this work possible.

I'd also like to thank my parents, Anita and Mike Marcotte, for their encouragement in reaching all my goals. Thanks to Monty Queener, whose love and trust reached across the many miles between us to sustain me. I especially thank God, who gives life to us all, for the glimpse of divinity brought to me through all these people.

This work was supported by NCI grant PO1-CA23417.

TABLE OF CONTENTS

1. INTRODUCTION	12
Single-Photon Emission Computed Tomography	12
Organization of this Thesis	16
2. THE UNIVERSITY OF ARIZONA BRAIN IMAGER	18
Overview of the Brain Imager	18
Projection Data from the Modular Gamma Cameras	20
Estimating Interaction Coordinates with the Poisson Model	25
Single Modular Camera Response	30
The Brain Imager System Matrix	33
3. THE EXPECTATION MAXIMIZATION ALGORITHM	37
Motivation	37
Mathematical Basis	37
Properties of the EM Algorithm	44
Methods for Suppressing Noise and Edge Artifacts	47
Comparisons between EM Reconstruction and Convolution	
Backprojection	49
4. THE POISSON MODEL FOR MODULAR CAMERA RESPONSE	52
Modelling the Cameras for EM Processing	52
Estimated Camera Coordinate as the Measurement Index	54
Direct PMT Signals as the Measurement Index	57
Simulation Studies	60
Flood Image Results	75
Blur Image Results	78
Experiments with Low Counts	86
Results of EM Processing the Measured Camera Response	90
Calibration Image Results	92
5. THE MULTIVARIATE NORMAL MODEL FOR MODULAR CAMERA	
RESPONSE	103
Modelling the Cameras for EM Processing	103
Direct PMT Signals as the Measurement Index	108
Estimated Camera Coordinate as the Measurement Index	110
Simulation Studies	113
Point Array Results	114
Flood Image Results	117
Blur Image Results	122
Summary	129
Calibration Image Results	129

TABLE OF CONTENTS - *Continued*

6. SUMMARY AND RECOMMENDATIONS FOR FUTURE WORK	137
Improvements to the Camera Response Model	138
Reconstruction Approaches	140
Reconstructions with Measured Centroids and Averaged Blurs ...	141
Reconstructions with Measured Centroids and Theoretical Blurs ..	141
REFERENCES	143

LIST OF ILLUSTRATIONS

Figure 2.1	The University of Arizona Brain Imager. (Illustration by Wally Varner.)	19
Figure 2.2	Arrangement of the modular gamma cameras in the University of Arizona brain imager.	21
Figure 2.3	Coordinate system of the modular gamma camera. A, B, C and D identify the four PMT positions.	24
Figure 2.4	Interpolated mean detector response at $q=14$ for a typical modular gamma camera. Note that $q=14$ crosses over the center of PMTs C and D.	29
Figure 2.5	Example of a flood image detected by the modular gamma camera.	32
Figure 4.1	Measured flood image mapped by the LUT that included uphill (\hat{p}, \hat{q}) estimates.	58
Figure 4.2	Measured flood image mapped by the LUT that excluded uphill (\hat{p}, \hat{q}) estimates.	59
Figure 4.3	Point array (top) and blur(bottom) images that were used in simulations as the ideal flux vector f	62
Figure 4.4	Top: Simulated response to an ideal point array. Bottom: Measured response to a point array calibration image.	65
Figure 4.5	Profile of simulated and measured response to point array flux at $\hat{q}=34$	67
Figure 4.6	Profile of simulated and measured response to point array flux at $\hat{q}=2$	68
Figure 4.7	Measurement energy $E(g)$ of unregularized and smoothed EM estimate of simulated point array image.	69
Figure 4.8	Estimate energy $E(f)$ of unregularized and smoothed EM estimate of simulated point array image.	71

LIST OF ILLUSTRATIONS - *Continued*

Figure 4.9	Image of EM estimate of the simulated point array after 500 iterations.	72
Figure 4.10	Profile of the ideal flux, simulated response and EM estimate of simulated point array at $\hat{q}=34$. The EM result shown is after 500 iterations without smoothing.	73
Figure 4.11	Profile of the ideal flux, simulated response, and EM estimate of point array at $\hat{q}=34$. The EM result shown is after 500 iterations with smoothing.	74
Figure 4.12	Top: Simulated response to the ideal uniform flux. Bottom: Measured response to the flood calibration.	76
Figure 4.13	Profile of the simulated and measured response to uniform flux at $\hat{q}=1$	77
Figure 4.14	Measurement energy $E(g)$ and estimate energy $E(f)$ of EM estimate of simulated flood image. Results from unregularized and smoothed EM processing are shown.	79
Figure 4.15	Profile of the ideal flux, simulated response, and EM estimate of a flood image at $\hat{q}=14$. The EM result shown is after 10 iterations without smoothing.	80
Figure 4.16	Profile of the ideal flux, simulated response and EM estimate of flood image at $\hat{q}=1$. The EM result shown is after 10 iterations with smoothing enabled.	81
Figure 4.17	Image of the simulated response to the ideal blur flux vector. ...	83
Figure 4.18	Measurement energy $E(g)$ and estimate energy $E(f)$ of the EM estimate of the simulated blur image. Results from unregularized and smoothed EM processing are shown.	84
Figure 4.19	Image of the EM estimate of simulated blurs after 10 iterations. ...	85
Figure 4.20	Profile of the ideal flux, simulated response and EM estimate of the simulated blur image at $\hat{q}=1$. The EM result shown is after 10 iterations without smoothing.	87

LIST OF ILLUSTRATIONS - *Continued*

Figure 4.21	Profile of the ideal flux, simulated response and EM estimate of the simulated blur image at $\hat{q}=1$. The EM result shown is after 10 iterations with smoothing.	88
Figure 4.22	Profile of the ideal flux and EM estimate of blur images with decreased total counts. EM result shown is after 10 iterations with smoothing.	89
Figure 4.23	EM estimate of flood image using direct PMT signals.	91
Figure 4.24	Profile of the measured response and EM estimate of the point array calibration image at $\hat{q}=34$. The EM result shown is after 10 iterations with smoothing.	93
Figure 4.25	Profile of the measured response and EM estimate of the point array calibration image at $\hat{q}=2$. EM result shown is after 10 iterations with smoothing.	94
Figure 4.26	Profile of the measured response and EM estimate of the flood image at $\hat{q}=1$. EM result shown is after 1 and 10 iterations without smoothing.	95
Figure 4.27	Profile through EM estimate of flood calibration image at $\hat{q}=1$, with and without corner triangles set to zero. EM results shown are after 10 iterations with smoothing enabled.	97
Figure 4.28	Projection images of the bottle phantom recorded by the brain imager for cameras 0-11.	99
Figure 4.29	Projection images of the bottle phantom recorded by the brain imager for cameras 12-23.	100
Figure 4.30	Projection images of the bottle phantom for cameras 0-11 after EM processing for 10 iterations with smoothing enabled.	101
Figure 4.31	Projection images of the bottle phantom for cameras 12-23 after EM processing for 10 iterations with smoothing enabled.	102
Figure 5.1	Mean 5-bit response computed from the "super-MDRF" measurement. The four large blocks represent, from top left to lower right, PMTs C, D, A and B. Each smaller block represents a position in the 16 x 16 array of measured points.	105

LIST OF ILLUSTRATIONS - *Continued*

Figure 5.2	PMT covariance matrix computed from the 16 x 16 point "super-MDRF". Diagonal blocks are variances as a function of camera position for PMTs A, B, C and D. Sub-diagonal blocks are the covariance values between PMTs.	106
Figure 5.3	Top: Simulated response to an ideal point array using the MV model. Bottom: Measured response to a point array calibration image using the MV model.	115
Figure 5.4	Measurement energy $E(g)$ of the EM estimate of the simulated point array using the multivariate (MV) and Poisson (P) models. Results from unregularized and smoothed EM processing are shown.	116
Figure 5.5	Estimate energy $E(f)$ for EM estimates of simulated point array images using the multivariate (MV) and Poisson (P) models. Results from unregularized and smoothed EM processing are shown.	118
Figure 5.6	Profile of the ideal flux, simulated response and EM estimate of a point array through $\hat{q}=2$ using the MV model. The EM result shown is after 10 passes without smoothing.	119
Figure 5.7	Profile of the ideal flux, simulated response and EM estimate of a point array through $\hat{q}=2$ using the MV model. The EM result shown is after 10 passes with smoothing.	120
Figure 5.8	Top: Simulated response to an ideal flood image using the MV model. Bottom: Measured response to the flood calibration using the MV model LUT.	121
Figure 5.9	Measurement energy $E(g)$ of the EM estimate of the simulated flood image using the multivariate (MV) and Poisson (P) models. Results from unregularized and smoothed EM processing are shown.	123
Figure 5.10	Estimate energy $E(f)$ of the EM estimate of the flood image using the multivariate (MV) and Poisson (P) models. Results from unregularized and smoothed EM processing are shown.	124
Figure 5.11	Profile of the ideal flux, simulated response and EM estimate of flood image at $\hat{q}=1$ using the MV model. EM result shown is after 10 iterations without smoothing.	125

LIST OF ILLUSTRATIONS

Figure 5.12	Profile of the ideal flux, simulated response and EM estimate of flood image at $\hat{q}=1$ using the MV model. EM result shown is after 10 iterations with smoothing.	126
Figure 5.13	Measurement energy $E(g)$ for the EM estimate of the simulated blur image using the multivariate (MV) and Poisson (P) models. Results from unregularized and smoothed EM processing are shown.	127
Figure 5.14	Estimate energy $E(f)$ for the EM estimate of the blur image using multivariate (MV) and Poisson (P) models. Results from unregularized and smoothed EM processing are shown.	129
Figure 5.15	Profile of the ideal flux, simulated response and EM estimate of blur image at $\hat{q}=1$ using the MV model. EM result shown is after 10 iterations without smoothing.	130
Figure 5.16	Profile of the ideal flux, simulated response and EM estimate of blur image at $\hat{q}=1$ using the MV model. EM result shown is after 10 iterations with smoothing.	131
Figure 5.17	Profile of the measured response and EM estimate of the point array calibration at $\hat{q}=2$ using the MV model. EM result shown is after 10 iterations without smoothing.	132
Figure 5.18	Profile of measured response and EM estimate of point array calibration at $\hat{q}=2$ using the MV model. The EM result shown is after 10 iterations with smoothing.	133
Figure 5.19	Profile of measured response and EM estimate of flood calibration at $\hat{q}=1$ using the MV model. EM result shown is after 10 iterations with smoothing enabled.	135

ABSTRACT

A method is developed for pre-processing projection images for a SPECT brain imaging system. The projection images are recorded by modular gamma cameras that exhibit noisy response before processing. The image acquisition process is modeled so that the mean of the detected gamma-ray emissions is a linear transformation of the actual flux. Two models for detection are examined, one based on independent Poisson distributions and the other based on a multivariate distribution. The Expectation Maximization (EM) algorithm is used to invert the forward model to obtain a Maximum Likelihood estimate of the flux. Simulations using uniform, Gaussian and point flux patterns demonstrated that EM processing recovered improved estimates of these patterns. Processing measured images yielded improved estimates, but also revealed that both forward models are incomplete.

1. INTRODUCTION

Medical imaging systems assist physicians in making diagnoses by showing them the patient's anatomical structure or physiological function. Every imaging system includes a source, a detector, and some way of displaying the detected quantity in a way that is useful to the physician. The source used to produce nuclear medicine and radiology images is high-energy radiation to which the patient is exposed. The detector records the radiation after it passes through or emerges from the patient's body. The display may be a two-dimensional representation such as an X-ray film, or a slice through a three-dimensional space produced after computer processing (tomography). All such imaging systems have some tradeoff in the amount of radiation they can detect without subjecting the patient to undue exposure. A primary objective in developing better systems is to extract more information with less patient exposure and inconvenience. One approach to this objective is to increase the efficiency and signal-to-noise ratio of the detectors in the imaging system. The research described in this thesis targets the recovery of previously unused information available at the detector of a single-photon emission computed tomography system.

Single-Photon Emission Computed Tomography

Single-photon emission computed tomography (SPECT) is a medical imaging modality used in nuclear medicine to provide physicians with information about how

organs are functioning. Other modalities such as magnetic resonance imaging (MRI) and transmission computed tomography (CT) provide anatomical maps of the shape and locations of different materials within the patient's body. Nuclear medicine images instead show where biochemical processes are taking place. Anatomical structures that are not associated with the selected process are not imaged, and therefore, the anatomical features are unclear. SPECT images also suffer from poor resolution and low signal-to-noise ratio.

Another functional imaging modality used in nuclear medicine is positron emission tomography (PET). PET is appealing because it can use isotopes of carbon, nitrogen, oxygen and fluorine to produce images. These isotopes allow virtually any biologically significant molecule to be made radioactive and studied (Barrett 1981). The main limitation of PET is its financial cost, which arises because the desired isotopes must be produced by a cyclotron facility. The isotopes have short half-lives, so that they must be delivered from the cyclotron frequently. Transporting the isotopes in a timely manner can be expensive if the nearest cyclotron is in another city. Obtaining isotopes may be prohibitively expensive for small communities with a limited need for PET studies. PET is mentioned here for completeness, so that the reader will be aware that SPECT is not the only method for performing functional studies.

In SPECT imaging the patient is administered a radiopharmaceutical (a chemical compound tagged with a radionuclide and prepared in a form suitable for

human use, Chandra, 1987). The radiopharmaceutical is more likely to accumulate in a particular type of tissue than in the surrounding tissues. Alternatively, the radionuclide may be attached to particles of a size that traps them within an organ, like the lungs or liver. The radionuclide emits gamma rays in all directions so that the patient becomes the radiating source in the imaging system. The resulting image shows the distribution of the radionuclide in organs of the body. For example, the distribution of a particular tracer can be used to study perfusion and ventilation of the lungs, blood flow, cardiac and renal function. SPECT is of particular interest in identifying and locating cancerous tissues. Tumors in the liver, spleen, brain and bones will show characteristically higher or lower activity in a SPECT image (Barrett 1981).

To detect gamma-ray emissions, one or more gamma-ray cameras are positioned around the patient. The cameras collect data for a period of time and then they are rotated about the patient to gather projections from many angles. The mechanical motion of the rotating camera can produce uncertainties in the measurements. A typical SPECT study takes 30 minutes, which is tiring for the patient. Also, the patient cannot stay perfectly still for such a period. The degradation of the measurements due to patient and mechanical motion can cause artifacts in the reconstructed image.

The gamma camera used in most commercial systems is the Anger camera, which was developed in 1958 by H. Anger. The detector face of the Anger camera is

a scintillation crystal, typically sodium iodide, which emits light photons when a gamma ray gives up its energy in the crystal. An array of photomultiplier tubes (PMTs) collects the light photons and produces a set of electrical signals. The position of the gamma ray interaction is estimated by how the event contributes to the entire array of PMT signals.

Typically a collimator is placed between the source and the gamma camera to define a field of view for each position on the camera face. The collimator is an array of holes bored into a lead plate. Gamma rays arriving in a line parallel to the bore pass through the bore and are detected. Other gamma rays are absorbed by the lead.

The long imaging time required to perform a SPECT study results from several factors that limit the number of gamma-rays that can arrive at the scintillation crystal. These factors include the size of the detector, the absorption of photons by the lead between the bores of the collimator, and the limited number of photons that can be emitted by the patient while keeping the dose of radionuclide as low as possible. Temporal resolution better than 10 minutes is very unlikely in conventional single-camera SPECT studies. Studies of how the chemical and mechanical distribution of the radiopharmaceutical changes over shorter periods of time are not possible. (Singh 1989). Such dynamic studies could provide currently unavailable information and improve diagnosis and treatment. For example, imaging the regional kinetics of cisplatin, a cancer chemotherapeutic drug, would improve understanding of the patient's response to chemotherapy (Singh 1989).

To address the need for shorter imaging times and for dynamic functional studies, researchers have used a number of approaches to collect more projection data simultaneously. The most prominent approach is to increase detector area. Singh *et al.* proposed a system of 19 scintillation cameras based on individual position-sensitive photomultiplier tubes (PSPMTs). The PSPMTs of the proposed system exhibited good linearity within the 10 cm diameter inner region of the 13 cm diameter photocathode. The design attaches a 10 cm diameter cone-beam collimator to a scintillation crystal over this inner 10 cm region of the photocathode. Computer simulation studies showed that the proposed system could provide a factor of 4 improvement in sensitivity with 19% worse spatial resolution than a single conventional scintillation camera SPECT instrument (Singh 1989). Researchers at Osaka University Medical School and Hitachi Medial Corporation have developed a four-head rotating gamma camera for SPECT (Kimura 1990). The four detectors each have an area of 26 cm x 20.8 cm. The detectors have high sensitivity and are rotated at high speeds to allow dynamic studies. The ASPECT camera (Digital Scintigraphics, Inc.) includes an annular crystal which is 31 cm in diameter, 13 cm in length. The PRISM 3000 (Picker International, Ohio Imaging) has three detectors each with crystal area 28 x 44 cm. The Strichman 810 scanner (Strichman Medical Equipment, Inc.) utilizes twelve detectors, each with a 20 cm x 12.5 cm crystal. The Strichman 810 scanner has sufficiently high sensitivity that it can acquire and reconstruct tomographic data as fast as one slice every 30 seconds (Zubal 1993).

Organization of this Thesis

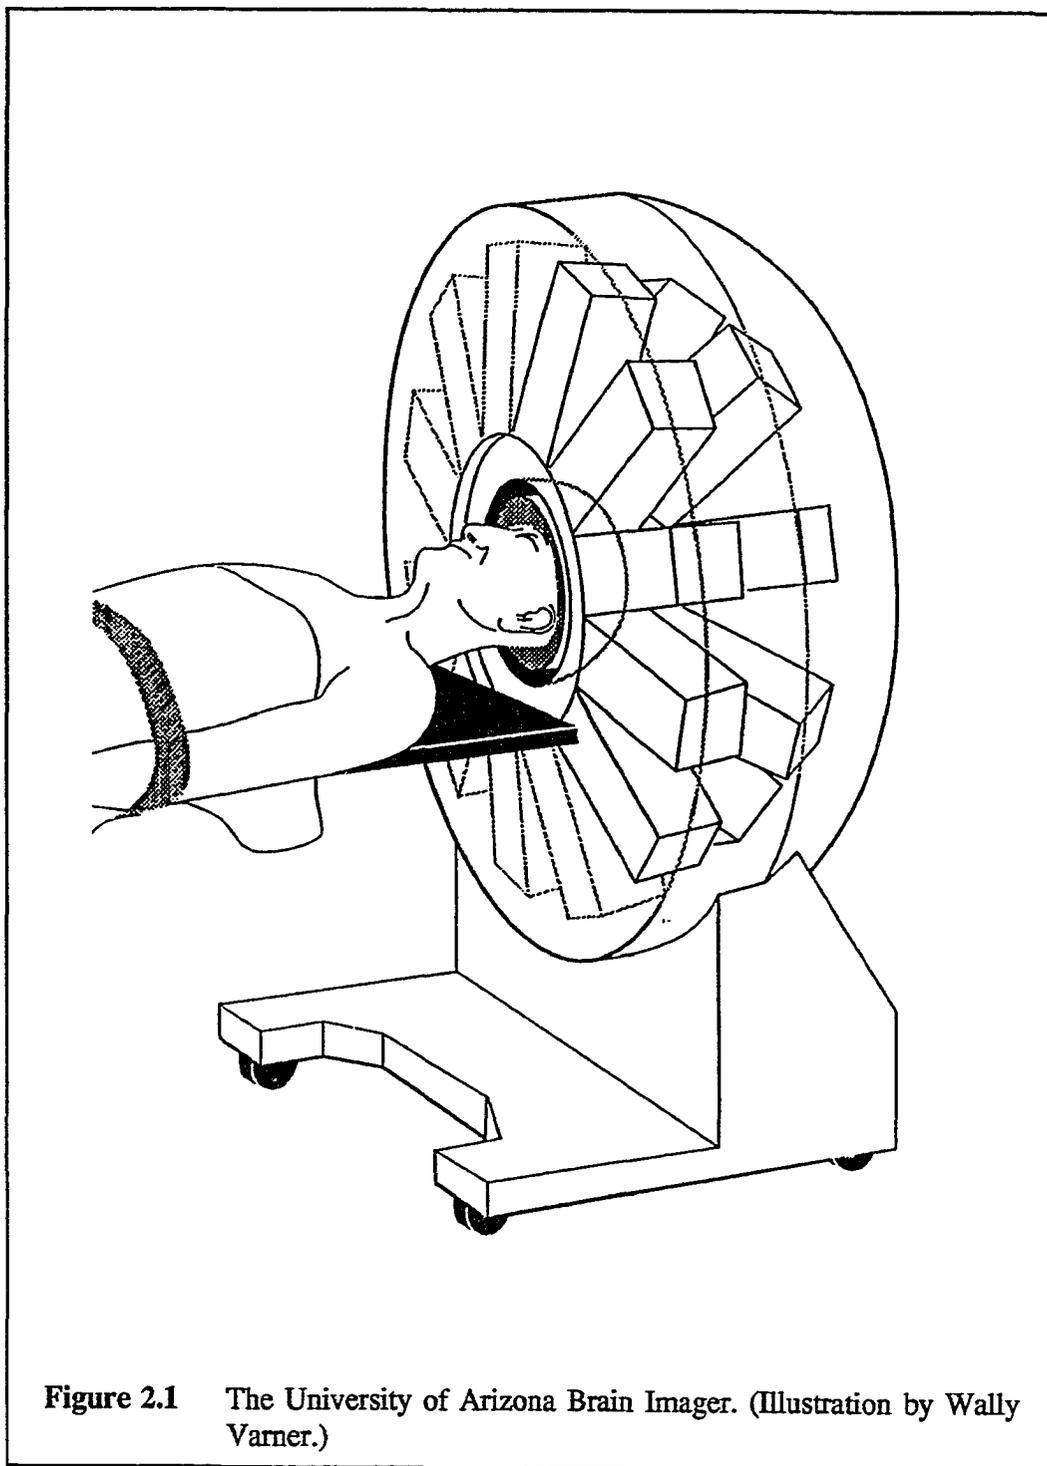
The imaging systems discussed above begin to address the need to reduce imaging time by increasing detector area around the patient. Chapter 2 of this thesis describes the University of Arizona brain imager, which also surrounds the patient with detectors. The modular gamma cameras that comprise the detectors of this system are the focus of this thesis. The work described here proposes an additional approach to minimizing imaging times. We know that detection of scintillation events is imperfect and that not all events are recorded. Errors also occur in estimating the position of the events. We model these properties statistically and attempt to recover an improved estimate of the number and position of scintillation events. Rather than collecting more photons, we seek to extract more information out of the photons we collect already. The algorithm used to recover the improved estimate is the Expectation Maximization (EM) algorithm. Chapter 3 describes the mathematical basis and properties of the EM algorithm. The model used to describe the statistics of the modular gamma camera has historically been Poisson. Chapter 4 describes this model and the results of using the Poisson model with the EM algorithm to improve the projection data. An alternative model based on a multivariate probability distribution is proposed and evaluated in Chapter 5. Finally, a summary of this work and suggestions for extending this research are described in Chapter 6.

2. THE UNIVERSITY OF ARIZONA BRAIN IMAGER

The radiology research group at the University of Arizona has developed a SPECT imaging system to perform dynamic studies of the brain. This imaging system addresses the problems of long collection times, camera rotation artifacts, and loss of data due to the collimator. A full discussion of system can be found in the references (Rowe 1991). This chapter describes the design and performance of the brain imager at the onset of the research, with emphasis on the role of the modular gamma cameras.

Overview of the Brain Imager

The brain imager, depicted in Figure 2.1, uses 24 modular gamma cameras arranged in two rings as the detector array. These modular cameras were designed so that they could be arranged in complex geometries around the patient. Such arrangements are not possible with the larger Anger cameras. The scintillation crystal of the modular camera is 10 cm x 10 cm. (Typical Anger cameras have circular crystals, 30 to 51 cm in diameter.) The two-ring geometry approximates a hemispherical detector surface, so that the imager collects all projections simultaneously, eliminating camera motion and reducing collection time. The first ring of cameras forms a short cylinder around the sides of the head, and the second ring forms a truncated cone.



The University of Arizona brain imager does not use a lead collimator positioned next to each camera crystal to limit the field of view of the accepted gamma radiation for each detector element. Instead, the imager employs a lead multiple-pinhole coded aperture positioned between the object and detectors. The aperture shape is a hemisphere extended to a cylinder to accommodate the patient's head. Each pinhole in the lead aperture allows a cone of gamma radiation to fall on the modular camera face, which is divided into small area elements. The cone of radiation projects onto many detector elements. If a detector element field of view includes lines that project through more than one pinhole, the projection data are coded or multiplexed. The apertures used with the brain imager include multiplexed (146 total pinholes) and unmultiplexed (1 pinhole per camera) pinhole patterns.

Each gamma ray interaction in the camera crystal produces electrical signals from four PMTs. These signals are sampled, digitized, compressed and stored for further processing. The PMT signals are mapped to an image of the locations of the interactions on the crystal face. The images from the 24 cameras are used to reconstruct an estimate of the emitting object. Next we discuss specifically how the PMT signals were used to produce projection images at the onset of this research.

Projection Data from the Modular Gamma Cameras

The detectors of the brain imaging system consist of 24 modular gamma cameras. The cameras are arranged, as shown in Figure 2.2, in two rings around the

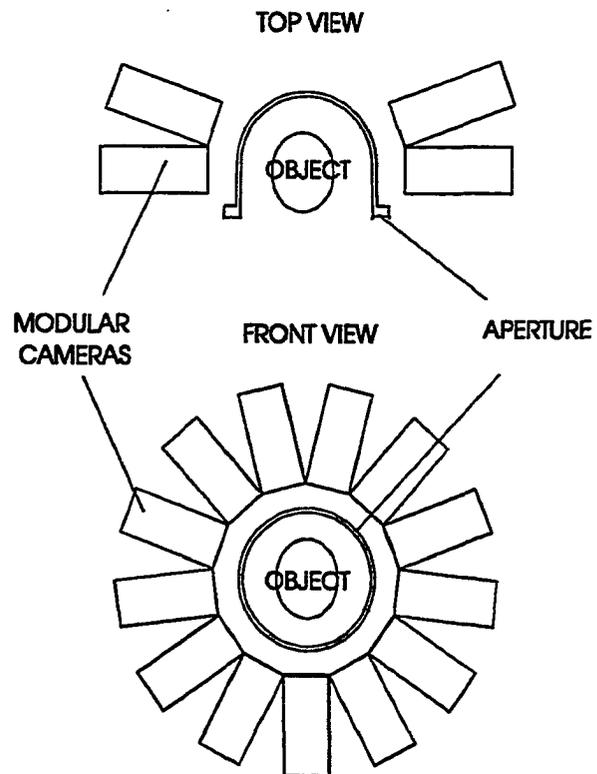


Figure 2.2 Arrangement of the modular gamma cameras in the University of Arizona brain imager.

aperture. Thirteen cameras form the first ring, with the camera faces parallel to the axis of rotation of the hemisphere. Eleven cameras form the second ring, with the crystal faces 20 degrees from parallel, towards the aperture. The modular gamma camera, introduced by Milster *et. al.* (Milster 1983), is a small, independently operating detector that allows a geometry such as this one to be implemented easily. Details regarding its design, construction and operation have been presented (Arendt 1980, Milster 1984, Selberg 1984, Milster 1985, Milster 1987, Aarsvold 1988, Milster 1990). We present here a summary of how the camera detects gamma ray emissions to produce a projection data set.

Gamma rays passing through the aperture interact in the crystal and generate light photons. The light photons pass through a light guide towards four PMTs, which generate electrical signals that depend upon the number of detected photons. The PMT signals are then acquired with a custom analog card. The analog card samples and digitizes each PMT signal to 8 bits. To minimize storage requirements, the 8-bit signal is then compressed to 5 bits. The four signals are concatenated to form a 20-bit value, referred to as an ABCD value, because the PMTs are labelled A, B, C and D. With 20 bits, over one million different ABCD values can be represented.

We wish to associate each possible ABCD value with a gamma ray interacting within the crystal. A convenient way to categorize the events is with a coordinate location where the interaction took place. The approach taken here is to estimate the coordinate which has the maximum likelihood for producing the ABCD value. To

compute a maximum likelihood (ML) estimate, we must have a suitable probability model for the PMT signals as a function of coordinate position.

The arrangement of the PMTs of the modular camera is shown in Figure 2.3. The coordinate system for the camera face uses the variables p and q to identify the axes. Variable p has values from 0 to 63, from left to right across the crystal face. Variable q has values from 0 to 63 from top to bottom. Coordinate (0,0) is located at the top, left of tube C when facing the crystal side of the modular camera. The notation for an estimate of the coordinate is (\hat{p}, \hat{q}) .

Two probability models for the modular camera response are presented in this research. The Poisson model includes the assumption that the PMTs are statistically independent, and has been used in all studies prior to this report. The second model is based on a multivariate distribution. The multivariate model is employed for the first time in this research and is presented in Chapter 5.

A brief note regarding notation will clarify the following discussions. The probability for a random variable x conditional on some quantity y is expressed $\Pr(x|y)$. The mean of a random variable is indicated by an overscore ($\bar{}$). Vectors are expressed in lower-case boldface letters. Matrices are expressed as upper-case boldface letters. Elements of vectors or matrices are indicated by a lower-case subscript. Estimated quantities are indicated by a hat ($\hat{}$).

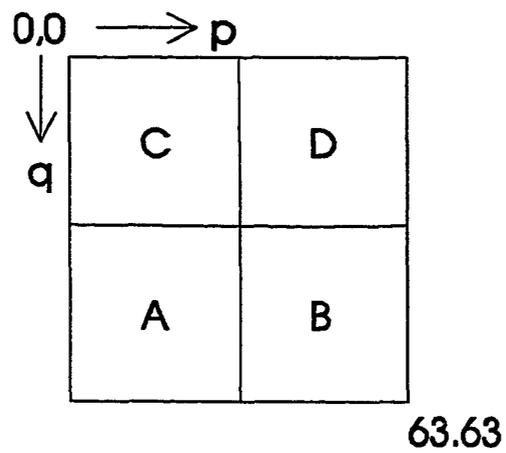


Figure 2.3 Coordinate system of the modular gamma camera. A, B, C and D identify the four PMT positions.

Estimating Interaction Coordinates with the Poisson Model

Given a PMT signal value ABCD, we seek the coordinate estimate (\hat{p}, \hat{q}) that maximizes the probability $\Pr(ABCD|\hat{p}, \hat{q})$. To describe this probability, first we must describe the random processes that contribute to the ABCD signal. We consider the emission of gamma rays over a interval of time to be a Poisson process. The detection process can be modelled as having two outcomes, detected or not detected, so it is considered to be binomial. The binomial probability $\Pr(k|n)$ of k detected events is conditional on n , the number of events which occur. A binomial process that depends upon a Poisson random variable results in another Poisson distribution (Barrett 1981). We express this Poisson probability as follows.

$$\Pr(x|\bar{x}) = \frac{\bar{x}^x e^{-\bar{x}}}{x!} \quad (2.1)$$

The quantity x represents the random number of recorded photoelectrons, and \bar{x} represents the mean of x . Each PMT records some value x for each gamma photon interaction. The number of photoelectrons is recorded as an 8-bit signal, a . This signal includes the gain of the PMT expressed as G , such that $a = Gx$. If we assume that a unique mean signal from each PMT (\bar{a} , \bar{b} , \bar{c} or \bar{d}) exists for each coordinate of the detector, we can express the conditional probability as

$$\Pr(abcd|p, q) = \Pr(a|\bar{a}(p, q)) \Pr(b|\bar{b}(p, q)) \Pr(c|\bar{c}(p, q)) \Pr(d|\bar{d}(p, q)) \quad (2.2)$$

The four PMTs are modelled as statistically independent, because it can be shown that Poisson sources cannot produce correlated counts in two different non-overlapping detectors (Barrett 1981). The joint probability function is therefore the product of the individual probabilities.

The 8-bit signal, a , (ranging from 0 to 255) is compressed to a 5-bit value, A , during acquisition. We refer to this 5-bit value (ranging from 0 to 31) as a channel. The compression follows a square-root mapping so that if a is small, few values of a will be assigned to a single channel. When a is large, a larger range of a maps to one channel. The mapping follows the formula

$$A = \begin{cases} 0 & \text{if } a=0 \\ \frac{31}{\sqrt{255}} \sqrt{a-0.5} & \text{otherwise} \end{cases} \quad (2.3)$$

This many-to-one mapping is incorporated into the model by taking the sum of a values that map to each channel. The total probability of recording a particular 5-bit value is then

$$\Pr(A|\bar{a}(p,q)) = \sum_{a \in A} \Pr(a|\bar{a}(p,q)) \quad (2.4)$$

To implement the search for the ML coordinate estimate, we require knowledge of the mean number of events, \bar{x} , in response to a source at each camera coordinate (p,q). This response could be measured directly if we could acquire the full 8-bit signal a by simply dividing \bar{a} by the gain. With 5-bit channels, the measurement is collected at the lower precision and used to estimate what the mean response would be at higher precision. The measurement is performed by placing a collimated source at a known crystal coordinate and recording in a histogram the number of times each 20-bit signal value occurred until the total number of events is 25000. For a typical PMT, 7 to 9 channels will record most counts and the other channels will record nearly zero. The measurement is repeated at 256 different locations in a 16 x 16 pixel grid and is referred to as the mean-detector-response-function (MDRF).

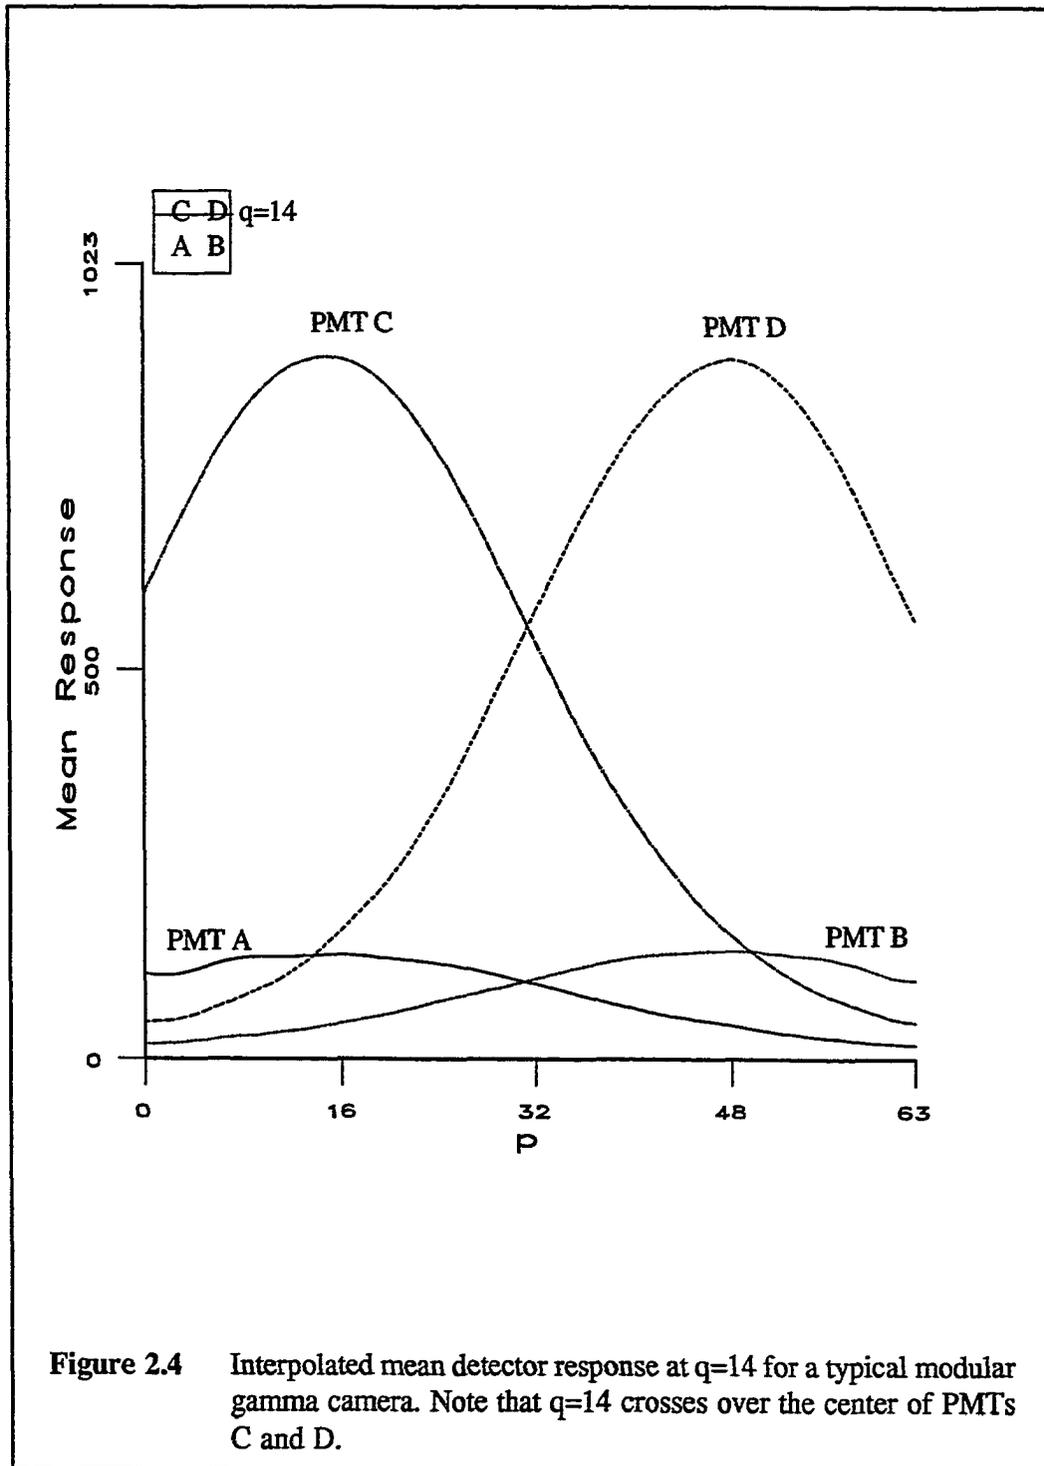
Estimating the mean of the distribution at higher precision is performed by a method called logarithmic matched filtering (Aarsvold 1988). To determine the mean of the Poisson distribution recorded, a finer-resolution Poisson distribution curve is fitted to the histogram, taking into account the mapping from 8-bits to 5-bits. The mean of the fitted curve is stored as the mean detector response at that measured source location. The fine-resolution mean for each tube is stored as a 10-bit value. The MDRF calibration includes both the above measurement and the estimation of the mean.

We treat the crystal face as a 64 x 64 pixel detector. To obtain mean values for all pixels, the 16 x 16 MDRF is interpolated to obtain mean response values for

the remaining points. The interpolation results in a file that associates every (p,q) coordinate on the face of the camera with a 10-bit mean of a Poisson distribution for each PMT. In notation to follow, the 10-bit value from each PMT is referred to in overscored lowercase letters. The 10-bit value corresponds to the 8-bit signal described above, multiplied by 4. The 64×64 file of \bar{a} , \bar{b} , \bar{c} and \bar{d} values is called the "abar" file. The abar file, addressed by coordinate, stores the 10-bit mean value associated with that coordinate for the four PMTs. The abar file supplies the maximum likelihood search program with the mean values needed to describe the Poisson distributions. An example of the mean values for the four PMTs is shown for constant q in Figure 2.4.

Given the mean of the Poisson distribution, we can determine the maximum likelihood estimate for the coordinate for each ABCD value. During data collection, a fast way to map the ABCD value to a crystal coordinate is to use a look-up table (LUT). The ABCD value determines the LUT address from which to retrieve the coordinate estimate for the event. During data acquisition, an image is formed by retrieving the coordinate for each ABCD event, then incrementing the event count at that coordinate in an image.

The LUT is generated by finding the (p,q) pair that maximizes equation (2.2). The LUT based on maximum probabilities can then be adjusted to promote uniformity in the response to a flood source. This adjustment is referred to as "flood correction". The LUT is adjusted by associating some ABCD values with a (\hat{p},\hat{q}) pair that has



lower (but non-zero) probability than the original (\hat{p}, \hat{q}) estimate. The ABCD value is reassigned if the reassignment removes counts from an above-average pixel, adds counts to a below-average pixel, and moves the counts no more than one pixel.

Between 2% and 3% of the possible ABCD values map to a coordinate. The remaining values are unlikely to occur at all, or are associated with low-energy scattered events or background radiation. These events do not map into the image. If the data were acquired as a histogram for each 20-bit value, one million locations would be required. Mapping to (\hat{p}, \hat{q}) reduces the size of the data file to an image with 4096 locations. The image represents the estimated number of times an interaction occurred at each crystal coordinate. The mapping to camera coordinates provides data compression as well as a familiar coordinate system for dealing with projection data.

Once each camera has been characterized by the abar and LUT files, the data collected from the brain imager can be stored as an image of the 24 camera faces. The intensity of each pixel in this image represents the estimated number of gamma interactions recorded at that pixel.

Single Modular Camera Response

To evaluate the performance of each modular camera, two data sets are collected soon after the MDRF is collected. The first data set is for a uniform or flood flux resulting from an uncollimated point source placed 1 meter from the camera face. The ABCD histogram is collected until the total number of events is 4 million.

The second is an 8 x 8 point array collected by moving a collimated point source in a grid pattern and collecting at each position for 25,000 events. The histogram file contains the data for all points in the array, so that the resulting image is a sum over time of the complete array.

If the LUT mapped PMT signals to true flux across the camera face, the image for the flood data set would be uniform. The image in Figure 2.5 shows a typical flood image. Only the central 62 x 62 pixel area is shown here because the edges and corners accumulate much higher counts than do the central region of the camera. The four corner points (not shown) are about 22 times greater than the flood average. (The reasons for this build up and ways to reduce it are addressed in Chapter 4.) The central flood image shows great variation, especially towards the corners of the crystal. A typical flood image has an average of 580 counts per pixel in this 62 x 62 region. The standard deviation among 3844 pixels and 24 cameras is 194 counts. The point array image is also degraded in the corners of the crystal.

At the onset of this research, we asked how well the modular gamma camera response represented the actual flux of gamma radiation across the crystal face. The purpose of this research is to improve how the modular camera PMT signals are used to produce projection data. With an accurate model of how the modular camera response alters the incoming signal, it is possible recover a more complete representation of the projection data. An improved estimate of the projections at the data acquisition stage would allow reconstructions to be more accurate.

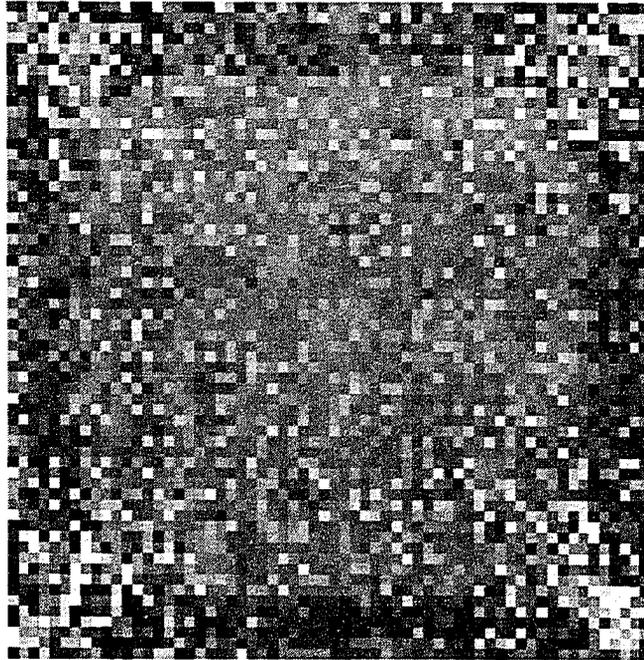


Figure 2.5 Example of a flood image detected by the modular gamma camera.

The Brain Imager System Matrix

Another motivation for pre-processing the projection images is related to the reconstruction method used in the brain imager. The algorithm used to reconstruct objects from the projection images is a weighted least-squares iterative algorithm. A full description of the algorithm can be found in Rowe, 1991. Briefly, the algorithm begins with an initial estimate of the object. With each iteration, it perturbs each element in the estimate, then computes an energy function that is based on the estimate. The energy function is chosen so that it decreases when agreement increases between the estimate and measured data. The equation for the forward imaging situation can be represented as $\mathbf{g} = \mathbf{H}\mathbf{f} + \mathbf{n}$. Here, \mathbf{g} is a vector of the measured camera response at each pixel when the three-dimensional object vector \mathbf{f} is imaged by the brain imager. The matrix \mathbf{H} maps discrete points in the object space to a response vector in detector space. The system response to a vector in object space is treated as the sum of responses from the discrete volume elements (voxels) that make up the object. The measurement vector includes this sum, plus additive noise \mathbf{n} due to photon noise and estimation errors.

Finding the estimate of the three-dimensional object is an inverse problem that requires an accurate description of the system matrix \mathbf{H} . Historically, the system matrix for the brain imager has been measured by stepping a 5 mm diameter source through a three-dimensional grid in the volume where the object would be placed. For each voxel in object space, the response of the 24 cameras is recorded. The gamma

radiation from the point source projects through the pinholes in the aperture, generating a set of blurs in the plane of the detectors. Some projections miss the cameras, partially miss the cameras, or fall in a region of the camera where the response is noisy. For example, in the corner regions of the crystal, adjacent pixels will have very different values, giving the appearance of a checkerboard. When the blur falls on these regions, the boundaries of the blur are difficult to distinguish. The size of the blur is typically less than 9×9 pixels. This measurement is also called the point-spread function (PSF) of the system, because a point in object space is not recorded as a collection of points in detector space, but as blurred projections. The magnification of the blur depends on the distance between the voxel and pinhole, and between the pinhole and the detector. The width and shape of the blur depend on the position on the camera where the blur is recorded.

The PSF measurement requires a large amount of storage. Using the complete PSF data as H in the reconstruction process requires long computation times. For a 20 cm diameter, 20 cm tall cylindrical volume measured at 5 mm intervals, for example, the PSF requires more than 250 Mbytes for storage, about 24 hours for measurement, and more than 24 hours for reconstruction. One way to decrease the storage and computation requirements is to represent each blur in the measured PSF in a more compact way. The method employed historically was to represent each blur with a centroid and an associated blur function. The centroid represents the magnitude and location of each pinhole projection. The blur function represents the shape of the blur,

and the weights of the pixels surrounding the centroid, whenever a centroid falls on that detector. All blurs in the PSF file that were associated with a given centroid pixel were averaged together to represent the response of the camera at that pixel.

This format reduces the storage requirements greatly. The average blur is non-zero in a 9×9 pixel region. The blur function file is always the same length: 81 pixels per detector, for 98304 detectors in the system. Using an aperture with 146 pinholes, typically four pinholes project onto each camera per voxel position. The PSF file is reduced from about 324 pixels per camera per voxel, to 4 pixels per camera per voxel.

Another advantage of the centroid-and-blur storage format is that the H matrix can be measured on a coarse grid and interpolated to a fine grid. Once centroids are extracted, a trilinear interpolation can be used to determine where intermediate voxels would project onto the camera face. Measuring the 20 cm diameter, 20 cm tall cylindrical volume on a 1 cm grid takes less than four hours. Measuring the same volume on a 5 mm grid takes more than 24 hours. During reconstruction of the phantom, the PSF for each voxel is generated from the centroid and average blur file, as an approximation to the original image. The time required for regenerating the blurs is comparable to the time needed to read in the raw data for each voxel.

Another approach to reconstruction employs only the centroids, with the blurring at each centroid ignored. This reconstruction approach is appealing because it greatly reduces computation time. Since only non-zero values in H are used, the

product \mathbf{Hf} can be computed much more quickly with 4 pixels per voxel, rather than the complete blurs. This approach is not satisfying in practice, because we know that the point in object space projects onto more than one pixel.

If we could pre-process the projection images and remove the blurring from the modular camera response, reconstruction from centroids only would be more reasonable. If blur functions could be eliminated from \mathbf{H} , storage requirements and reconstruction times could be reduced. Performing reconstructions from centroids only, or from a modified version of the centroid and blur PSF, was beyond the scope of this thesis. Anticipation of these reconstruction methods did provide motivation for pre-processing the modular camera images.

3. THE EXPECTATION MAXIMIZATION ALGORITHM

Motivation

The motivation for using the expectation-maximization algorithm in this research emerges from the success of its use in reconstructions from SPECT and PET projection data. The algorithm incorporates the statistical relationship between the source distribution and the measured data. Since the emission of gamma photons is a random process, methods based on a deterministic geometric model fail to incorporate this knowledge about the data.

Mathematical Basis

The mathematical basis for the EM algorithm was described by Dempster, Laird and Rubin in 1977 (Dempster 1977). They proposed it as a general approach to iterative computation of maximum likelihood estimates, given a set of incomplete observations. Shepp and Vardi (Shepp 1982) and Lange and Carson (Lange 1984) independently introduced the EM algorithm for the specific case of Poisson distributions in emission tomography. In this situation, the observed incomplete data are the projection data measured with the gamma cameras. The unknown data are the distribution of radionuclide within the object imaged by the system. The derivation below is based on the Shepp and Vardi approach, with notation from Barrett (Barrett 1993).

The EM algorithm is used to compute an estimate of an unknown vector, \mathbf{f} , from an observed data set, \mathbf{g} . It assumes that a known statistical relationship maps the mean of the true vector to the observable \mathbf{g} . This relationship is given by the matrix \mathbf{H} . We consider \mathbf{g} to be the incomplete data. The complete data, \mathbf{c} , contains information that clearly identifies how each element of \mathbf{f} contributed to the element of \mathbf{g} .

The observed data are digitized as photon counts in each detector element. The true object is continuous, but the estimate of the object must have a finite number of elements so that it can be computed. The object is therefore divided into discrete elements of regular size and spacing. The reconstruction is an estimate of the activity within each object element, with the activity assumed uniform over the extent of the element. The components of the derivation are as follows.

g_m is the measurement at the element m in detector space. The measurement vector has elements $\{0, \dots, M-1\}$ and is denoted \mathbf{g} .

H_{mn} is the probability that an element n in object space will be detected by element m in image space. The transformation matrix has dimension $M \times N$ and is denoted \mathbf{H} .

f_n is the unknown vector we wish to estimate. It is the mean of the distribution of activity at the element n in object space. The object vector has elements $\{0, \dots, N-1\}$ and is denoted \mathbf{f} .

c_{mn} is the complete data that cannot be observed. Each element represents the number of photons recorded in element m that originated at object element n . The complete matrix is denoted c and has dimension $M \times N$.

y_n is the true distribution of activity at the element n in object space. The true vector has elements $\{0, \dots, N-1\}$ and is denoted y . The elements of the vector f are the expected values of the elements of y .

The vectors y , f and g consist of sums of the complete data. Given an object vector f , the mean measurement at detector m can be represented as elements of the mean of the complete data as follows.

$$\bar{c}_{mn} = H_{mn} f_n \quad (3.1)$$

The mean of the measurement vector can be computed by summing over object space the product of H_{mn} with f_n .

$$\bar{g}_m = \sum_{n=0}^{N-1} \bar{c}_{mn} = \sum_{n=0}^{N-1} H_{mn} f_n \quad (3.2)$$

$$g_m = \sum_{n=0}^{N-1} c_{mn} \quad (3.3)$$

Note that the measurement vector is the sum of the complete data over all object indices. The true distribution y_n , is the sum of the complete data over the measurement space.

$$y_n = \sum_{m=0}^{M-1} c_{mn} \quad (3.4)$$

The derivation ignores the emissions that are not detected, so that both y and f represent detected emissions only. We estimate the mean of the object distribution as \hat{f}^k , where k is the number of iterations executed. The EM algorithm computes the estimate for each iteration in two steps. The first step computes the expectation, $\bar{y}_n^k = E[y_n | \hat{f}^k, g]$, of the (random) number of emissions at element n (y_n), based on the mean of the previous object estimate, f^k and the measurement vector, g . The second step computes a new estimate of the mean number of emissions, f^{k+1} , that maximizes the probability for the observed data, given the old estimate y^k .

Consider the probability of an emission being detected by any element m in detector space as

$$\Pr(m) = \sum_{n=0}^{N-1} \Pr(n, m) = \sum_{n=0}^{N-1} H_{mn} = 1 \quad (3.5)$$

We wish to maximize the $\Pr(g | \hat{f})$ or the probability of the measured g , given an object estimate \hat{f} . The number of emissions from the object in an interval of time occurs as a Poisson random variable. The emissions are also uncorrelated (Barrett

1981). Since the number of detected emissions is the sum of independent Poisson random variables, each g is also Poisson, and depends on f . The likelihood function is defined by the product of the independent Poisson variables.

$$L(f) = \Pr(g | f) = \sum_A \prod_{m=0}^{M-1} \prod_{n=0}^{N-1} \frac{(\bar{c}_{mn})^{c_{mn}} e^{-\bar{c}_{mn}}}{c_{mn}!} \quad (3.6)$$

The measurement vector is expressed in terms of the complete data. The sum is over all arrays A of c_{mn} such that equation (3.3) holds true. By direct differentiation with respect to a particular element of the object vector, n_o , we obtain

$$\frac{\partial L(f)}{\partial f_{n_o}} = \sum_A \left\{ \prod_{m=0}^{M-1} \prod_{n=0}^{N-1} \frac{(\bar{c}_{mn})^{c_{mn}} e^{-\bar{c}_{mn}}}{c_{mn}!} \right\} \left\{ -1 + \frac{1}{f_{n_o}} \sum_{n=0}^{N-1} c_{mn_o} \right\} \quad (3.7)$$

The conditional expectation for the emissions at object element n_o can be expressed as

$$E[y_{n_o} | g, f] = \frac{1}{\Pr(g | f)} \sum_A \prod_{m=0}^{M-1} \prod_{n=0}^{N-1} \frac{(\bar{c}_{mn})^{c_{mn}} e^{-\bar{c}_{mn}}}{c_{mn}!} y_{n_o} \quad (3.8)$$

Define the log of the likelihood $l(f) = \log(L(f))$. Differentiation of the log of the likelihood with respect to an element f_n produces the following relation as a function of the complete data. The equation uses the expressions in equation (3.7) and (3.8).

$$\begin{aligned} \frac{\partial l(\mathbf{f})}{\partial f_n} &= \frac{1}{L(\mathbf{f})} \frac{\partial L(\mathbf{f})}{\partial f_n} = \frac{1}{\Pr(\mathbf{g}|\mathbf{f})} \left\{ -\Pr(\mathbf{g}|\mathbf{f}) + \frac{\Pr(\mathbf{g}|\mathbf{f})}{f_n} \mathbb{E}[y_n|\mathbf{g},\mathbf{f}] \right\} \quad (3.9) \\ &= -1 + \frac{1}{f_n} \sum_{m=1}^{M-1} \mathbb{E}[c_{mn}|\mathbf{g},\mathbf{f}] \end{aligned}$$

Since the c_{mn} are all mutually independent, we can express the expectation in (3.9) as

$$\begin{aligned} \mathbb{E}[c_{mn}|\mathbf{g},\mathbf{f}] &= \mathbb{E}[c_{mn}|\mathbf{g}_m,\mathbf{f}] \\ &= \frac{g_m \bar{c}_{mn}}{\sum_{n'=0}^{N-1} \bar{c}_{mn'}} \quad (3.10) \end{aligned}$$

Substituting equation (3.1) for \bar{c}_{mn} and (3.10) for the expectation term in (3.9)

produces the relation that forms the basis of the EM algorithm.

$$\frac{\partial l(\mathbf{f})}{\partial f_n} = -1 + \sum_{m=0}^{M-1} \frac{g_m H_{mn}}{\sum_{n'=0}^{N-1} f_{n'} H_{mn'}} \quad (3.11)$$

We set the derivative equal to zero, add one to both sides and multiply by f_n to

determine an expression for f_n in terms of \mathbf{g} and \mathbf{H} .

$$f_n = f_n \sum_{m=0}^{M-1} \frac{g_m H_{mn}}{\sum_{n'=0}^{N-1} f_{n'} H_{mn'}} \quad (3.12)$$

To obtain an iterative relation between the old estimate (k), and the new estimate (k+1), let the left side of the equation be estimate (k+1), and the estimates in the right side of the equation be from iteration (k).

$$\hat{f}_n^{(k+1)} = \hat{f}_n^{(k)} \frac{\sum_{m=0}^{M-1} g_m H_{mn}}{\sum_{n'=0}^{N-1} \hat{f}_{n'}^{(k)} H_{mn'}} \quad (3.13)$$

After each iteration, the likelihood function $L(f^{(k+1)})$ is greater than the likelihood function of the previous estimate $L(f^{(k)})$. (The interested reader is referred to Shepp 1984 for the proof). Note in (3.13) that if $\hat{g} = g$, the new estimate equals the old estimate because of the normalization of equation (3.5). If the normalization of (3.5) is not possible, the algorithm requires division by the sum of H_{mn} over detector space. The EM algorithm therefore takes the following form:

$$\begin{aligned} \hat{g}_m^{(k)} &= \sum_{n'=0}^{N-1} H_{mn'} \hat{f}_{n'}^{(k)} \\ \hat{f}_n^{(k+1)} &= \frac{\hat{f}_n^{(k)} \sum_{m=0}^{M-1} g_m H_{mn}}{\sum_{n=0}^{N-1} H_{mn} \hat{g}_m^{(k)}} \end{aligned} \quad (3.14)$$

Properties of the EM Algorithm

The EM algorithm has a number of properties that are well suited to emission tomography. It also suffers from limitations that have prevented it from being used in commercial systems. These properties are as follows.

Positivity

One appealing property is positivity. It is known that the object source distribution cannot take on negative values. The EM algorithm starts with some positive, non-zero estimate of f and remains positive or zero for all subsequent iterations. This is easily seen since g is positive and H (a probability distribution) is positive.

Self-Normalization

The total counts in the object estimate remain constant with each iteration, once the initial pass is complete. If H is normalized to 1, the counts in g equal the counts in f after the first pass. This property prevents the object estimate from increasing beyond the known maximum radioactivity administered to the patient. When the total counts in g equals the counts in \hat{f} , the object can be thought of as the distribution of detected emissions, rather than the distribution of all (detected or missed) emissions.

Convergence

The EM algorithm converges monotonically to a global maximum of the likelihood function. While the object boundaries and low-frequency features are defined early in the algorithm, convergence towards the maximum likelihood (ML) solution progresses very slowly (Kaufman 1987, Lange 1984). Important practical issues related to convergence are speed per iteration and stopping criteria. EM reconstructions must produce an image quickly if they are to function in a clinical setting. A number of approaches to accelerating the algorithm include parallelization (Magee 1990, Miller 1991), beginning the reconstruction with coarse grid estimation (Ranganath 1988), and orthogonalization (Li 1993). Also, since convergence to the ML solution is only asymptotic, some measure is required to determine the best iteration at which to stop the EM algorithm. Approaches to stopping the algorithm include the use of statistical hypothesis testing (Veklerov 1987, Veklerov 1988) and Pearson's chi-squared statistic (Hebert 1988). Alternatively, cross-validation (Coakley 1991) can be applied by splitting the measurement data into two parts.

Noise and Edge Artifacts

Each iteration of the EM algorithm increases the log likelihood function. It has been observed (Snyder 1987, Chornoboy 1990) that as iteration number increases and the solution approaches the maximum likelihood solution, uniform regions of the estimate develop irregular peaks and valleys. These regions look noisy or

"checkerboarded" in the reconstructed image. These effects are characteristic of ML solutions and are not unique to the EM ML algorithm.

If we view the situation in the frequency domain, the unconstrained EM algorithm performs no band-limiting. With each iteration, the solution is converging towards an inverse filter solution, with amplification of high frequencies. Since the measured data include Poisson noise due to low counts and other noise due to imperfect detectors, the EM algorithm is trying to recover an estimate \hat{f} from a corrupted g . If the errors in g are high frequency, they will be amplified as the iteration tends toward the ML solution.

Snyder *et. al.* also describe a condition of overshoot at edges. At an edge, the estimated counts will first overshoot the higher value, then oscillate around the true value in the uniform regions bounded by the edge (Snyder 1987).

Barrett (1993) and Wilson (1993) have investigated the noise properties of the EM algorithm. They developed expressions for the statistical properties of EM estimates as a function of iteration number. Their expressions showed that the variance in the estimate is a function of the square of the estimate's mean and of the covariance matrix of the estimate. Monte Carlo simulations demonstrated agreement with theoretical predictions within 10% for a large number of images with different noise realizations.

Initial Estimate Can Influence the Solution

The choice of the initial estimate $\hat{\mathbf{f}}$ determines the path which the algorithm takes towards the ML solution. Kaufman observed that when the initial estimate $\hat{\mathbf{f}}^0$ was chosen to be one value for half of the estimate space, and another value for the remaining estimate space, the edge defined in $\hat{\mathbf{f}}^0$ remained in all subsequent iterations. This edge was accentuated by the overshoot described above. Similarly, if the initial estimate is constant, the early estimates will be smooth, and variations will increase with iteration number (Kaufman 1987). If the mapping function \mathbf{H} maps part of the object to the null space, many realizations of \mathbf{f} will result in the same measurement vector \mathbf{g} (Barrett 1981). In this case, the choice of the initial estimate can determine which realization the ML solution will favor.

Methods for Suppressing Noise and Edge Artifacts

A number of approaches to regularizing or suppressing noise in the EM estimate have been studied. Snyder *et. al.* proposed the use of Grenander's method of sieves to suppress the noise artifact (Snyder 1987). This method constrains the estimate to a subset, called a sieve, of the space of non-negative functions. To suppress the edge effect, they proposed blurring the estimate by constraining to a sieve of blurred functions. This was unsuccessful at suppressing the edge artifact. A second approach blurred the quantity to be estimated and succeeded in suppressing the edge effect.

Another approach to regularizing the EM algorithm is to use knowledge about the quantity to be estimated to limit the space of possible estimates. In this approach, the EM algorithm is derived to include knowledge about the object statistics. If the algorithm assumes nothing about the object, so that the prior probability is uniform, the EM ML algorithm results. If some other distribution can be assumed, the solution becomes a maximum *a posteriori* (MAP) estimation. Levitan and Herman considered a Gaussian (multivariate) *a priori* distribution for the object, using a diagonal covariance matrix with all variances equal to 1, and a scale factor to effect how heavily the prior influenced the reconstructions. To determine the *a priori* mean value of the object, they used a smoothed, filtered backprojection of the measurement data. Their simulations showed the MAP EM algorithm to produce images that exhibited less checkerboarding in uniform regions and converged to a stable estimate (Levitan 1987).

Hebert and Leahy (1989) and Hebert and Gopal (1992) developed a generalized EM (GEM) MAP algorithm based on Gibbs priors. The generalized EM seeks to increase the likelihood function, without requiring that it be maximized. Their purpose was to include the local correlations in the object neglected by the diagonal covariance assumed by Levitan and Herman. Gibbs distributions also do not require specification of the object mean, removing the need to estimate it by backprojection. Use of the Gibbs prior penalized separation of neighboring pixels up to a threshold (to provide smoothness) and allowed separations beyond the threshold (to allow well-defined boundaries). Their simulations showed that the GEM MAP reconstruction converged

at the same rate as the EM algorithm, but as the EM estimate began to diverge from agreement with the original object (increased noisiness), the GEM MAP solution continued to converge. They also showed that the GEM MAP algorithm was more robust with respect to camera modelling errors. The images exhibited correspondingly less variation in uniform regions of the object.

Green used a Bayesian approach to produce a modified EM algorithm aimed at maximum posterior probability estimation (Green 1990). The prior depended on a weighted function of the neighboring pixel differences. Empirical results of the algorithm, called one-step-late (OSL) showed it converged faster and could produce smoother reconstructions for certain parameters, chosen by trial and error. Hebert and Gopal note, however, that the OSL algorithm has not been shown to converge in theory, and may oscillate or diverge for some parameters (Hebert 1992). Miller and Roysam incorporated Good's roughness prior to address noise suppression in their parallel EM implementation. In this approach, additions of Good's prior corresponds to smoothing the unconstrained EM algorithm solution at each iteration (Miller 1991).

Comparisons between EM Reconstruction and Convolution Backprojection

The EM algorithm has been examined in recent years as an alternative to convolution backprojection (CBP) or filtered backprojection (FBP) methods for reconstruction of the source distribution in emission tomography. If a single element of the measurement vector can be treated as the line integral of all emissions

originating along a ray through the object, perpendicular to the detector plane, then the projection can be treated as the Radon transform of the object into detector space.

The Radon transform of the object is recorded as a function of projection angle, and distance from an origin in object space. Inversion of the Radon transform, to obtain an estimate of the source distribution, is achieved using the convolution-backprojection or filtered-backprojection method (Macovski 1983, Barrett 1981).

Physically, radioactive emissions are photon limited. Therefore, the measured projection is actually a poor estimate of the emissions truly present in the object space. The projections are also degraded by photon attenuation through the object, photon scattering, and the effect of the solid angle from the collimator. Compensation for these effects are not easily incorporated into FBP reconstructions (Chornoboy 1990). The FBP approach is used in clinical systems because of its speed, but FBP reconstructions exhibit ring-like and radial artifacts, and may underestimate the isotope concentration in the middle of the body (Green 1990).

Shepp and Vardi emphasized that the EM algorithm is superior to the traditional reconstruction approach of CBP. CBP is successful in transmission tomography, but does not account for the physics of emission tomography. CBP treats the variations due to low count rates as noise and amplifies the noise in the reconstructions. Maximum likelihood methods may produce less amplification of this statistical noise. The EM model can include known counting time per projection, known radioactive decay, variation in spatial resolution, attenuation and scatter in the

system matrix. With CBP, additional processing is required to correct for scatter and attenuation. Also, CBP treats three-dimensional volumes as slices, whereas the physical model of the EM algorithm generalizes directly to three dimensions. However, the increased computation time required for the iterative EM method has limited its applicability in clinical settings (Lange 1984).

Benefits of EM reconstructions over CBP and FBP methods have been described in many studies. Chornoboy *et. al.* compared FBP to EM with simulations and acrylic phantoms. Their EM reconstructions assumed that the point source response function was known. They found that improved signal-to-noise ratios and image contrast with their ML reconstructions resulted in better preservation of object boundaries. However, with increased iteration, their unconstrained EM algorithm developed and amplified peaks in uniform areas of the object estimate. With known activities in phantom objects their CBP images had increased counts because negative values in the CBP estimate were set to zero, to enforce positivity. The EM estimates were slightly lower (0.992) for the entire image. Additionally, they found using a simulated object with non-uniform attenuation that the EM algorithm recovered a cold region where the FBP reconstruction placed activity in the region. Important is their conclusion that EM maintains the above advantages, even with decreased photon counts. With decreased count requirements, patient dosage and collection times can be minimized (Chornoboy 1990).

4. THE POISSON MODEL FOR MODULAR CAMERA RESPONSE

The objective of this research was to improve the estimate of the gamma-ray flux incident on the modular camera face, given the modular camera response. The unknown data vector is represented as a two-dimensional array of planar elements of incident radiation. The geometry of the imaging problem is different from the previously described emission tomography studies that reconstructed from projections. In this research, EM is used as a pre-processing step. Nevertheless, the estimation of the incident flux follows the same approach as other EM solutions. The unknown object distribution is the plane of incident flux. The measured data are the response of the camera to the incident flux.

Modelling the Cameras for EM Processing

The EM algorithm requires an accurate forward model of the system that generates the observed data. Using a linear model, the observed data vector can be thought of as the result of multiplying the system matrix by the input flux vector. The notation we will use is as follows:

f_n is the flux of the gamma emissions incident on the modular camera crystal, at area element n . This index runs from 0 to $N-1$.

g_m is the measured modular camera response to flux across the camera face at detector m . This index in measurement space can be the (\hat{p}, \hat{q}) detector element estimated from the look-up table, or the binned PMT signal. The index m runs from 0 to $M-1$.

H_{mn} is the probability that a source at coordinate n will be detected at measurement element m . H has dimension $N \times M$.

ϵ_m is the noise added to the measurement component at m , independent of the system matrix.

Using this model, we can think of the elements of the measurement vector using the equation

$$g_m = \sum_{n=0}^{N-1} H_{mn} f_n + \epsilon_m \quad (4.1)$$

Since we know nothing about the additive noise, we approximate the solution by finding f from only g and H . The presence of noise is not ignored, however. Since noise is amplified at higher iterations of the EM algorithm, the presence of noise influences how we implement stopping criteria and regularization.

Estimated Camera Coordinate as the Measurement Index

First we describe the forward model when the measurement data are recorded as counts at each (\hat{p}, \hat{q}) coordinate. Each event recorded in the ABCD histogram is either mapped by the LUT to a coordinate in the camera response image or discarded. The camera response, $H(\hat{p}, \hat{q}|p, q)$, to a gamma photon interaction can be described as the probability $\Pr(\hat{p}, \hat{q}|p, q)$ that a gamma ray actually incident at coordinate (p, q) will result in a PMT signal response that maps to the estimated detector element at (\hat{p}, \hat{q}) .

A source located at (p, q) will produce a distribution of PMT responses. These responses may map to the same coordinate in detector space, or to some nearby element. To compute the conditional probability $\Pr(\hat{p}, \hat{q}|p, q)$, we consider the conditional probability that a detector estimate will result from a given PMT signal, and the probability that the same PMT signal will result from a source located at coordinate (p, q) . For simplicity of notation, (\hat{p}, \hat{q}) will be designated as m , and (p, q) will be designated as n as we continue.

Recall that the PMT signals consist of four 5-bit signals, one from each PMT. The composite 20-bit signal is described as an ABCD value. We compute $\Pr(\hat{p}, \hat{q}|p, q)$ by the summing the conditional probability over all possible ABCD values.

$$\begin{aligned} H_{mn} = \Pr(m | n) &= \sum_{ABCD} \Pr(m | ABCD) \Pr(ABCD | n) \\ &= \sum_{ABCD} \Pr(m | ABCD) \Pr(A|n) \Pr(B|n) \Pr(C|n) \Pr(D|n) \end{aligned} \quad (4.2)$$

Recall from the development of the LUT in Chapter 2 that we compute the joint probability $p(\text{ABCD}|p,q)$ from the product of the independent probabilities for each PMT. The probability $\Pr(m|\text{ABCD})$ results from the direct mapping of the look-up table. If the given ABCD value is recorded by the acquisition electronics, it will always be mapped to the same detector, or to no detector at all. If the ABCD value maps to element m in the look-up table, the probability equals 1 for element m , and 0 for all other elements.

The individual PMT probabilities, $\Pr(A|n)$, etc., are computed from theory and measurements. Recall from the description of the modular cameras that the mean response of the PMTs is available from the MDRF measurements and the interpolated abar file. For each coordinate in the flux plane, the abar file yields $\bar{a}(n)$, the mean 10-bit value for the response of PMT A. The means for all four PMTs are retrieved in this way. Each 10-bit unit can be mapped to an equivalent 5-bit channel number by dividing the 10-bit unit by 4 to obtain an 8-bit quantity equivalent to the PMT response signal. We obtain the 5-bit channel by mapping the 8-bit signal through the square root compression algorithm used by the acquisition hardware. This channel value for each PMT is used as the center of the window of ABCD values to consider when computing $\Pr(\hat{p},\hat{q}|p,q)$. We determine the size of the window in 5-bit signal space to limit the size of the columns of \mathbf{H} , and to eliminate extremely small values in \mathbf{H} .

We previously computed the probability $\Pr(A|\bar{a}(n))$ that a given 10-bit abar mean value will result in a certain 5-bit PMT signal, for the purpose of finding the maximum likelihood coordinate for each ABCD value. This probability function is stored in a table to allow fast access while generating the LUT. From the values of $\Pr(A|\bar{a}(n))$ in the table, we compute the joint probability $\Pr(ABCD|n)$ by taking the product of the individual probabilities for each PMT.

Use of the $\Pr(\hat{p},\hat{q}|p,q)$ relationship for the system matrix has some advantages over the raw PMT signal. The primary advantage is that the measured image resembles the recovered flux image and can be qualitatively assessed and understood. The PMT signal is a four-dimensional vector that cannot be easily visualized in relation to the flux pattern. Also, the dimension of the PMT signal space is one million (2^{20}) while the dimension of position estimates is 4096. Computing \mathbf{H} and processing with \mathbf{H} are faster with the smaller measurement vector.

Additionally, the use of the look-up table provides a way to eliminate the build-up of counts that occurs on the edge of the crystal. The look-up table determines the maximum probability coordinate location, (\hat{p},\hat{q}) that would produce an ABCD value. If an ABCD value maps to an edge or corner (\hat{p},\hat{q}) , it is possible that a higher probability exists for the PMT signal to result if the incident coordinate was outside of the camera space. While the expected PMT response to such an event would be no response at all, the likelihood can be computed by extrapolating the abar curve to one coordinate past the camera edge, (\hat{p}',\hat{q}') . From the extrapolated abar

value, $\Pr(A|\bar{a}(\hat{p}',\hat{q}'))$ can be compared with $\Pr(A|\bar{a}(\hat{p},\hat{q}))$. If the probability is higher for the extrapolated coordinate, the mapping to that edge coordinate is removed from the LUT. This process is called "discarding uphill (\hat{p},\hat{q}) estimates."

Without this rejection of uphill estimates, a typical histogram of PMT signals resulting from a flood source will produce an image in (\hat{p},\hat{q}) coordinates that has very high intensity along the edges, and even higher intensity in the corners. The corners are higher because they can include uphill estimates from both sides and the diagonal. An example of an image that includes uphill estimates is shown in Figure 4.1. The camera response is more uniform when the uphill estimates are rejected, as shown in Figure 4.2.

When the look-up table excludes uphill estimates, \mathbf{H} shows a falloff in sensitivity along the edge of the camera. This is exhibited when each column n of \mathbf{H} is summed over all detectors m . The sum for each flux position is close to 1, except along the edges, where it drops off rapidly.

Direct PMT Signals as the Measurement Index

The directly recorded PMT signals can be used as the measurement vector as an alternative to estimating where each gamma-ray interaction took place. In this case, the camera system matrix is the conditional probability $\Pr(ABCD|n)$, computed in the same manner as described above. The same ABCD window is considered, but instead of summing over the ABCDs that contribute to each (\hat{p},\hat{q}) , the individual

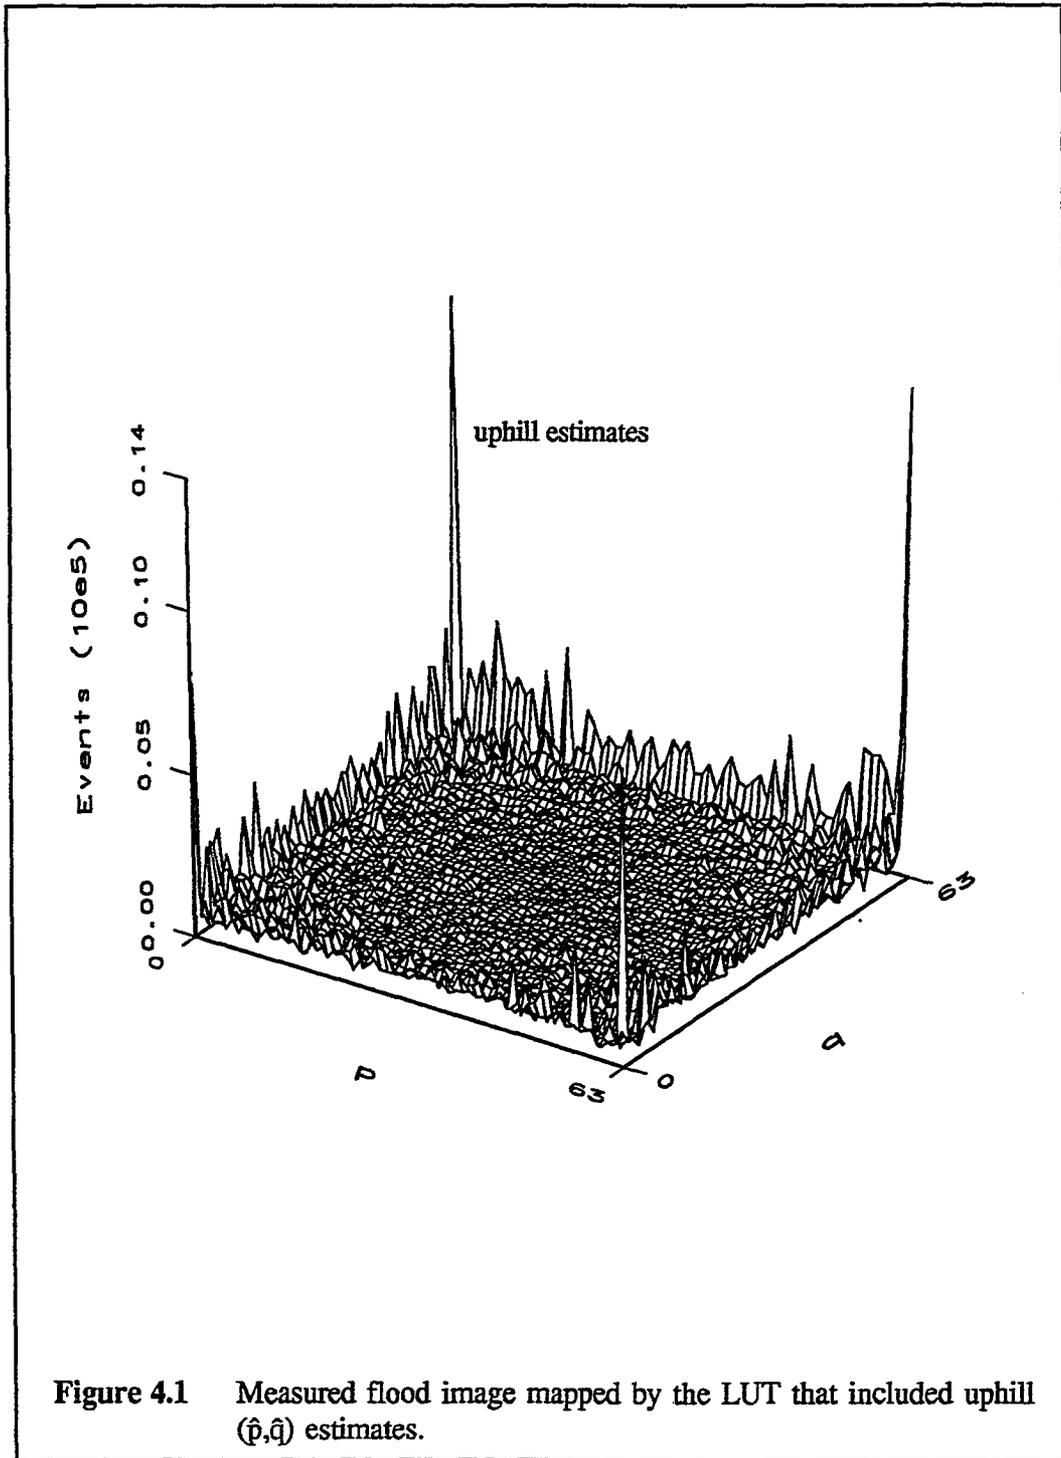
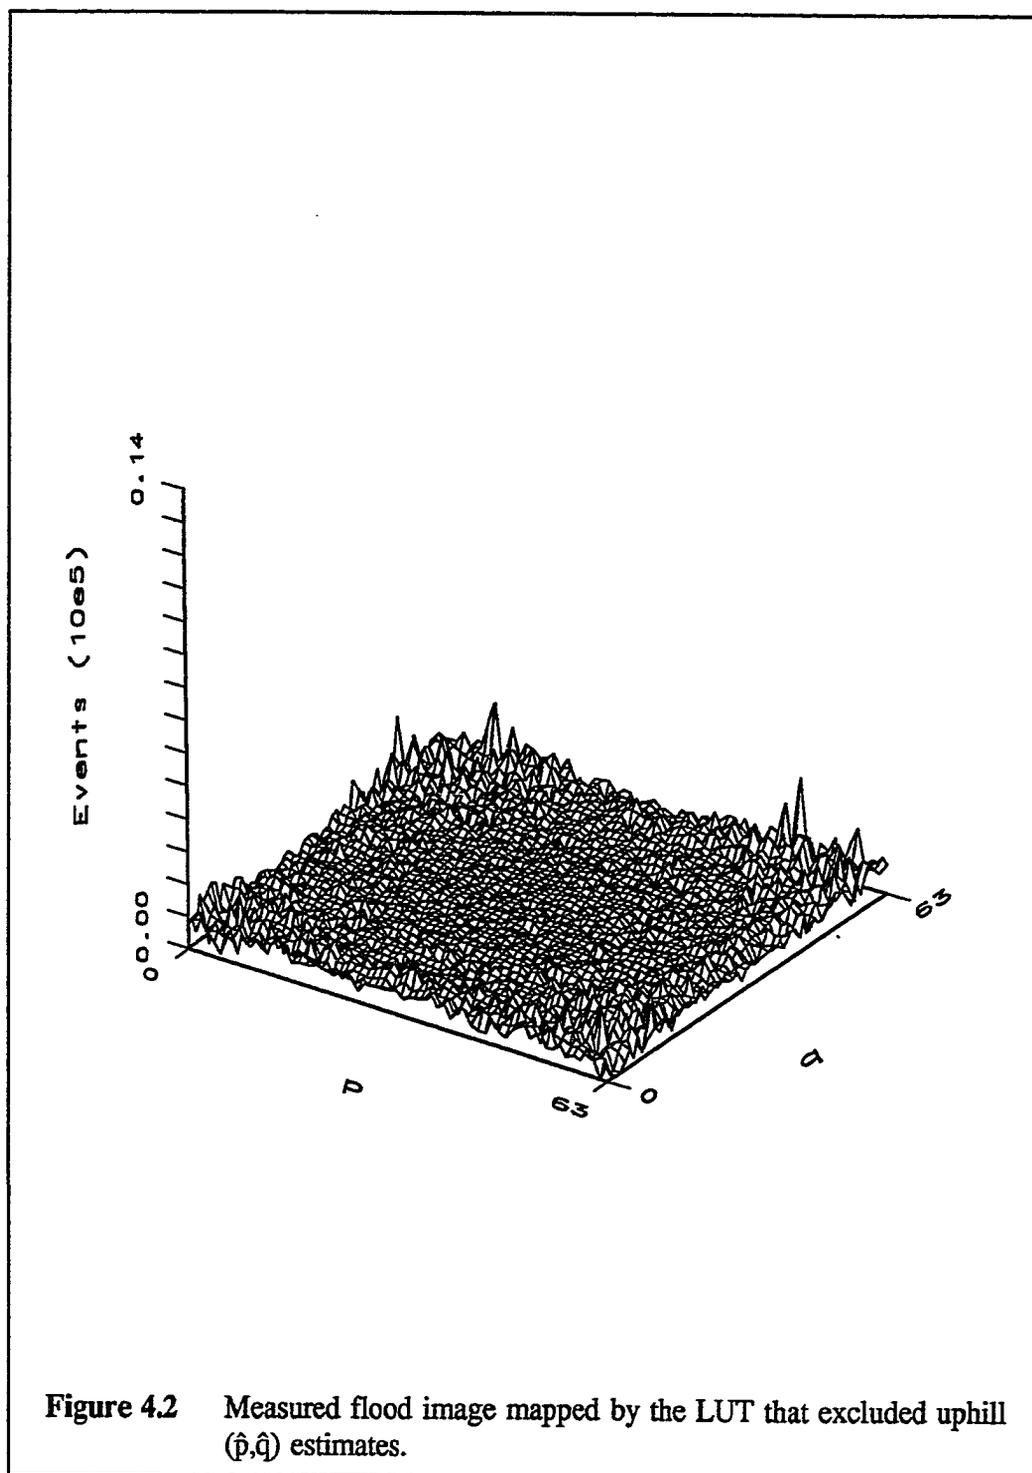


Figure 4.1 Measured flood image mapped by the LUT that included uphill (\hat{p}, \hat{q}) estimates.



$\Pr(\text{ABCD}|\mathbf{n})$ elements are stored separately within each column of \mathbf{H} . Without the look-up table mapping, \mathbf{H} includes the ABCD values that map to uphill estimates because it does not identify the ABCD values with regions of the camera. Hence, the sensitivity map columns all sum to nearly 1. A threshold can be employed to limit the size of the \mathbf{H} matrix, while maintaining the sum of each column close to 1.

Using direct PMT signals as the measurement index has some disadvantages, including the inclusion of uphill estimates. More elements per flux position are required in \mathbf{g} because many ABCD values map to each (\hat{p}, \hat{q}) . Also, the LUT acts to thin the number of elements added to \mathbf{H} when the product

$$\Pr(\hat{p}, \hat{q} | \text{ABCD}) \Pr(\text{ABCD} | p, q) \quad (4.3)$$

is computed. Larger files require longer processing times and larger storage capacity. If direct PMT signals could be shown to significantly improve the EM estimate, however, we could justify using them instead of the (\hat{p}, \hat{q}) image.

Simulation Studies

As the first step in developing the EM processing algorithm, we studied the performance of the algorithm on recovering an ideal known flux vector \mathbf{f} . The matrix \mathbf{H} was computed for a single modular camera using the Poisson model described above. The vector \mathbf{g} was indexed by camera coordinates, so that the elements of \mathbf{H} were $\Pr(\hat{p}, \hat{q} | p, q)$. Using the forward model $\mathbf{g} = \mathbf{H}\mathbf{f} + \mathbf{n}$, we simulated the camera

response to the known flux vector. The simulated response vectors were then compared with the measured vectors. To obtain the measured vectors, we used the LUT based on the Poisson model (with uphill estimates discarded) to map the histogram of the flood and point array measurements to images in (\hat{p}, \hat{q}) coordinates. The EM algorithm was used to recover an estimate of the flux vector, $\hat{\mathbf{f}}$.

Three ideal flux vectors were examined: a flood image, a point array and a "blur" image consisting of four two-dimensional Gaussian functions. The ideal flux vectors were chosen to emulate the images that are collected during calibration of the cameras. The point array was an 8×8 array of identical point sources spaces at regular intervals in (p, q) coordinates. A measured point array typically includes 25000 events for each position in the 8×8 grid (1.6 million total). After mapping by the LUT, a typical number of the total events is about 1 million. The ideal point array was therefore chosen so that the simulated response would be close to one million. The flood image had constant flux at all positions. The measured flood image typically includes a total of 2.5 million mapped events. The ideal flood was chosen so that the simulated response was close to 2.5 million counts. The blur flux vector consisted of four Gaussian blurs. It was selected to simulate the type of image collected when a radiating object projects through the pinhole aperture of the brain imaging system onto the camera face. The total number of counts in the ideal blur flux vector was one million. The two-dimensional images of the point array and blur flux vectors are shown in Figure 4.3.

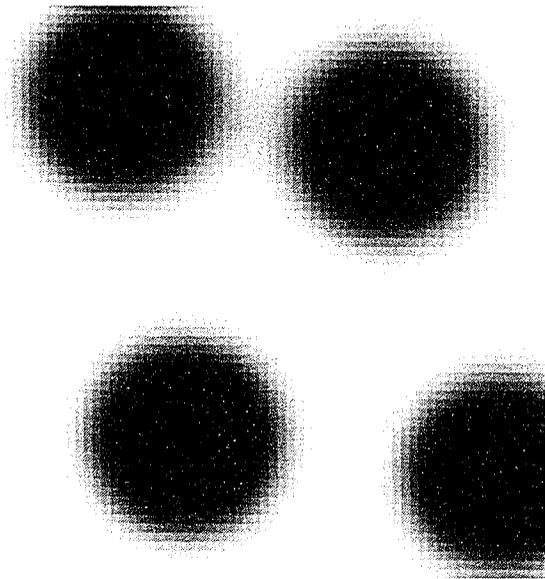
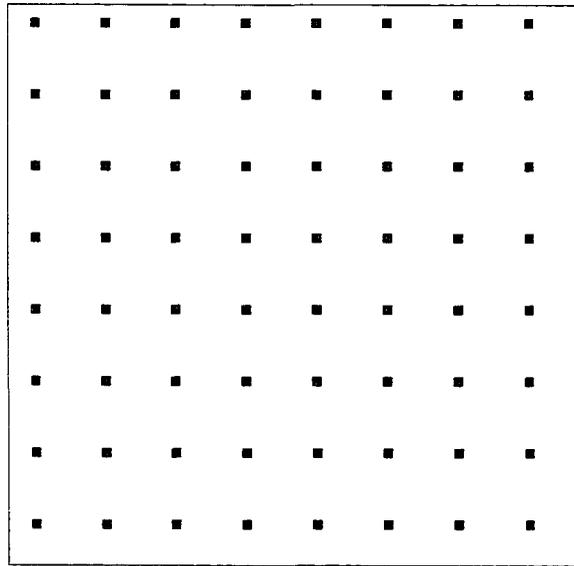


Figure 4.3 Point array (top) and blur(bottom) images that were used in simulations as the ideal flux vector f .

The measurement vector was computed by multiplying the \mathbf{H} matrix by one of the above flux vectors \mathbf{f} . Noise was added to each element of the resulting vector to produce \mathbf{g} . The noise was computed as a normal approximation to the Poisson distribution. This approximation is appropriate provided the mean number of events is greater than about 10 (Barrett 1981).

The simulated response images were processed using the EM algorithm. For all cases, the initial estimate was a constant value at all positions equal to the average value in the image. A constant initial estimate was selected to promote smoothness in the processed image. In addition, an option was included to smooth the estimate after every iteration. With smoothing enabled, each pixel was smoothed by taking the weighted average of its surrounding neighbors. The weighting function was Gaussian with standard deviation of 1 pixel.

The EM processing was repeated for a specified number of iterations. With each iteration, the program computed two energy functions. These functions were used to measure the difference between the estimate and ideal flux. The first energy function, termed the measurement energy $E(\mathbf{g})$, was computed using the equation

$$E(\mathbf{g}) = \sum_{m=0}^{M-1} \frac{(g_m - \hat{g}_m)^2}{(g_m + 1)} \quad (4.4)$$

$$\hat{g}_m = \sum_{n=0}^{N-1} H_{mn} \hat{f}_n$$

where \hat{g} is the estimate of the measured data that would result by multiplying the estimated flux by the system matrix. The denominator includes the constant 1 to prevent division by zero. The second energy function, the estimate energy $E(\mathbf{f})$, was computed from the known flux and the estimated flux as

$$E(\mathbf{f}) = \sum_{n=0}^{N-1} \frac{(f_n - \hat{f}_n)^2}{(f_n + 1)} \quad (4.5)$$

The energy functions were used as a guide for determining how many iterations to compute. During the first few iterations, $E(\mathbf{g})$ decreases rapidly. As processing continues, the change in energy between iterations decreases. Since we wish to process for a minimal amount of time, we stopped iterating when the rate of change in $E(\mathbf{g})$ was small compared to the previous rates. The estimate energy depended on the ideal flux vector and the regularization. From these simulations, we sought to determine a stopping rule for measured blur data.

Point Array Results

An example of measured and simulated point array images for one modular

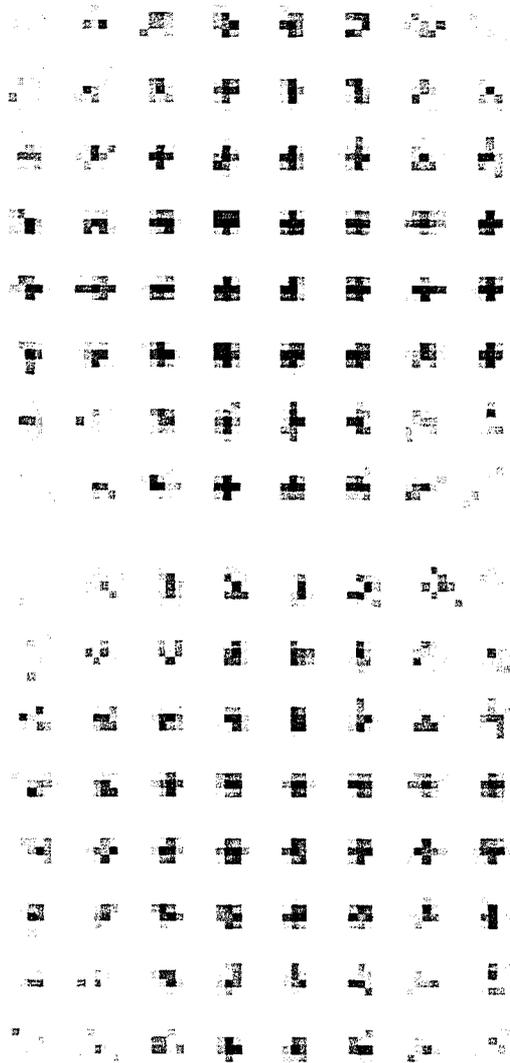
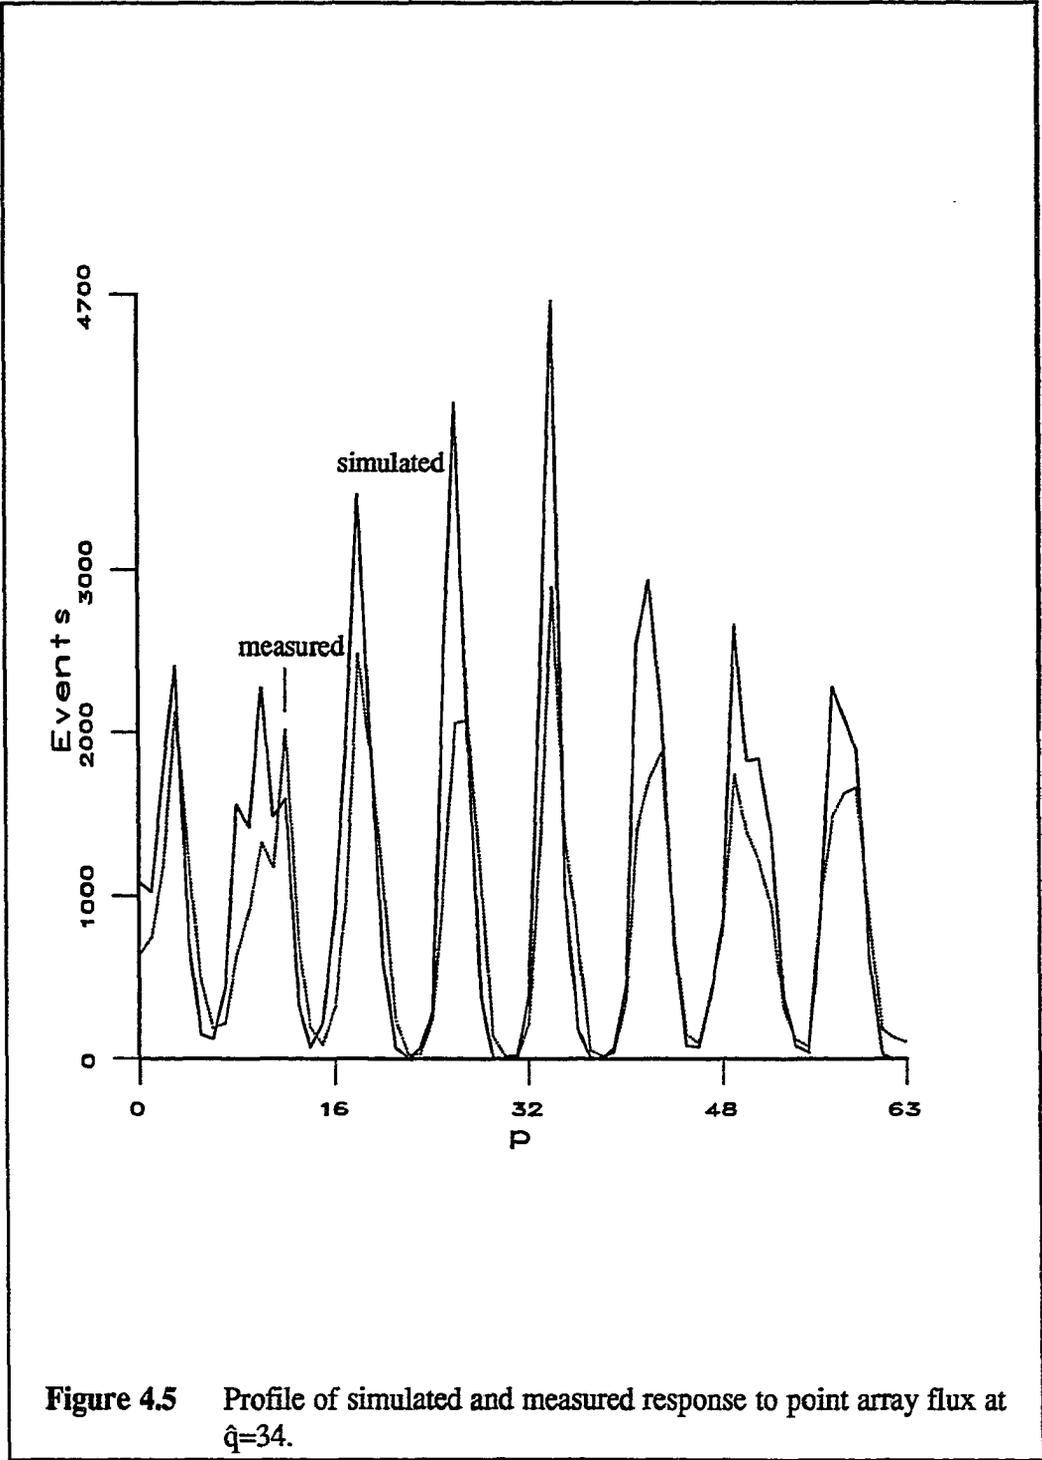
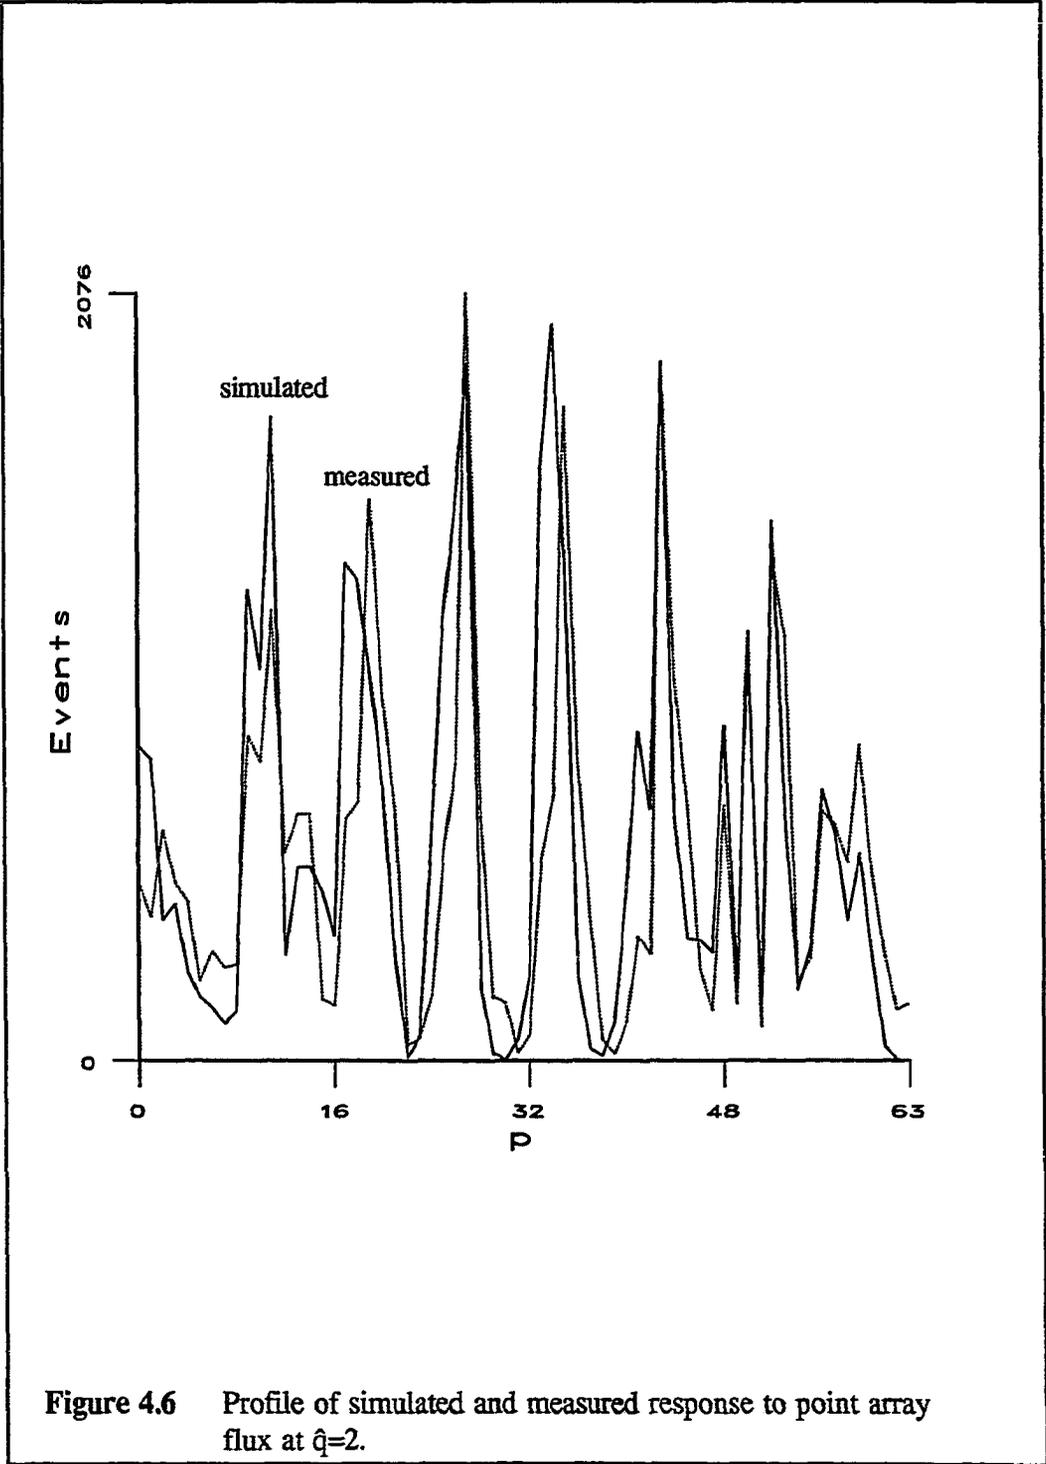


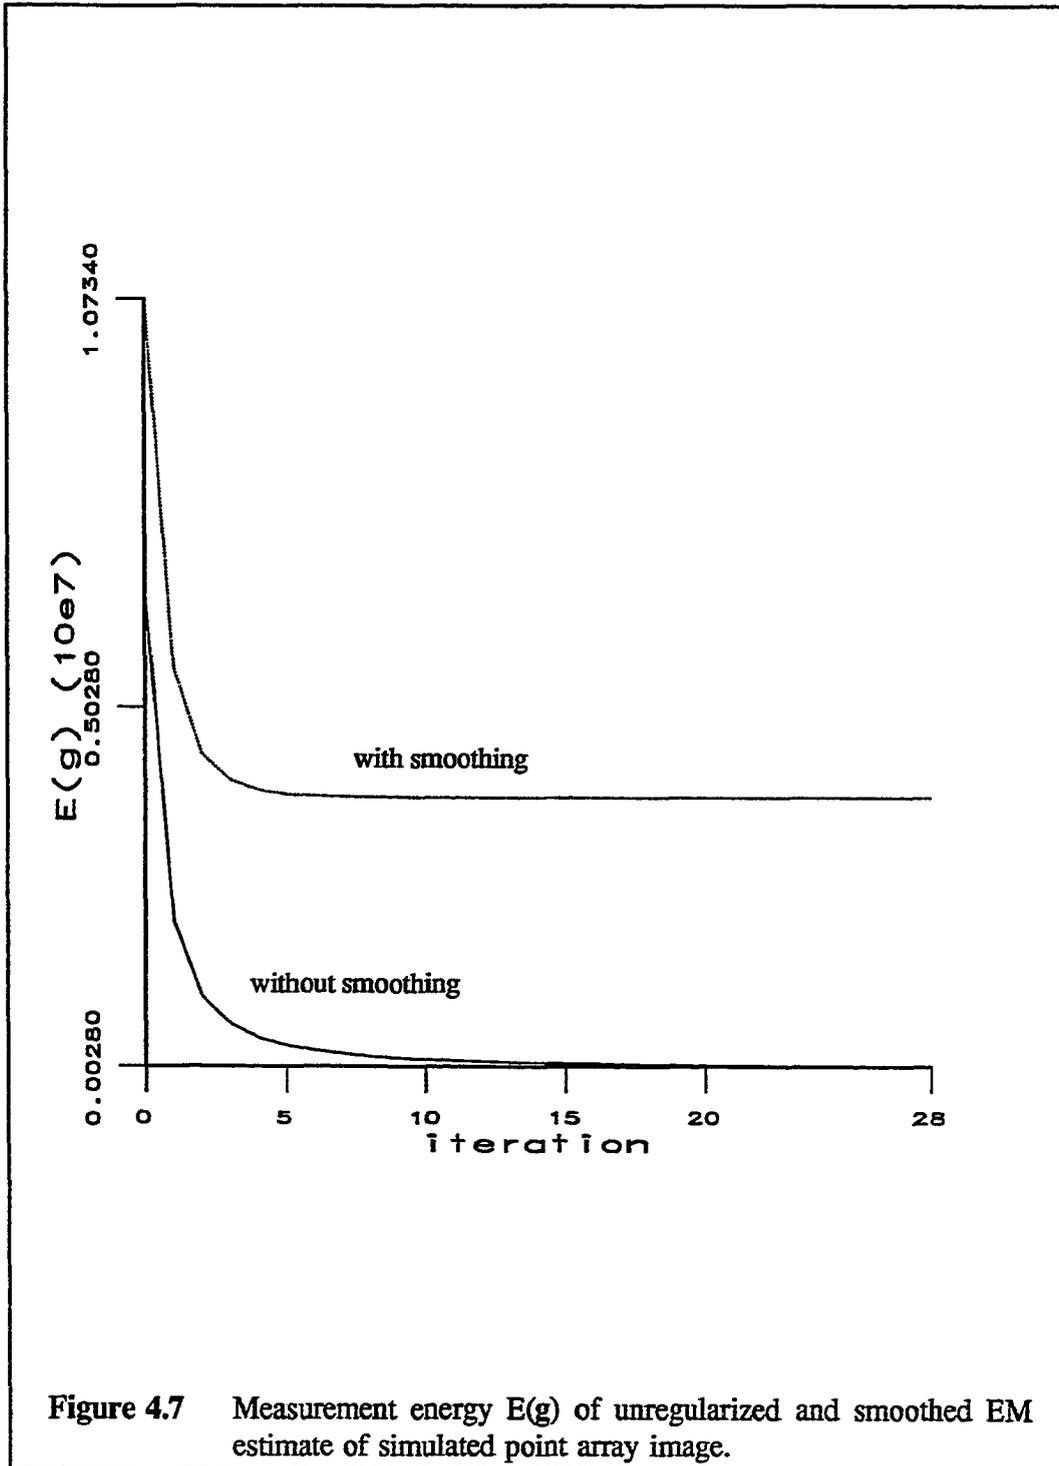
Figure 4.4 Top: Simulated response to an ideal point array.
Bottom: Measured response to a point array calibration image.

camera is shown in Figure 4.4. The simulation program generated noisy point array images that resemble measured images, especially in the center of the camera crystal. Note that both are well resolved in the center of the camera, but towards the corners and edges, the amplitude decreases and the counts are distributed over more detectors. The corners also show large variations between adjacent pixels. Figure 4.5 shows the profile of the point array along of one row of detector elements at $\hat{q}=34$. The solid line is the simulated vector and the dashed line is the measured vector. The amplitudes of the simulated points are higher than the measured points, but the widths are similar. The resemblance of the two images is encouraging. If the two images differed greatly it would suggest that the forward model was inadequate. The model performs best in the central region of the crystal. Note the difference between the measured and simulated point array at coordinate $\hat{q}=2$ in Figure 4.6. Here we see that the model does not perform as well. The widths remain similar, but the location of each peak is not matched as well as for $\hat{q}=34$. The differences in amplitude at $\hat{q}=2$ are proportionally smaller than the differences at $\hat{q}=34$.

Next we show the results after processing the simulated point array image with the EM algorithm. Here we compare the unregularized EM algorithm with one regularized by smoothing the estimate after each iteration. Figure 4.7 is a plot of $E(\mathbf{g})$ as a function of iteration number. The agreement between the projected estimate and the measurement is best for the unregularized EM algorithm. The energy is very high in either case. Both converge quickly, but smoothing causes the energy to converge



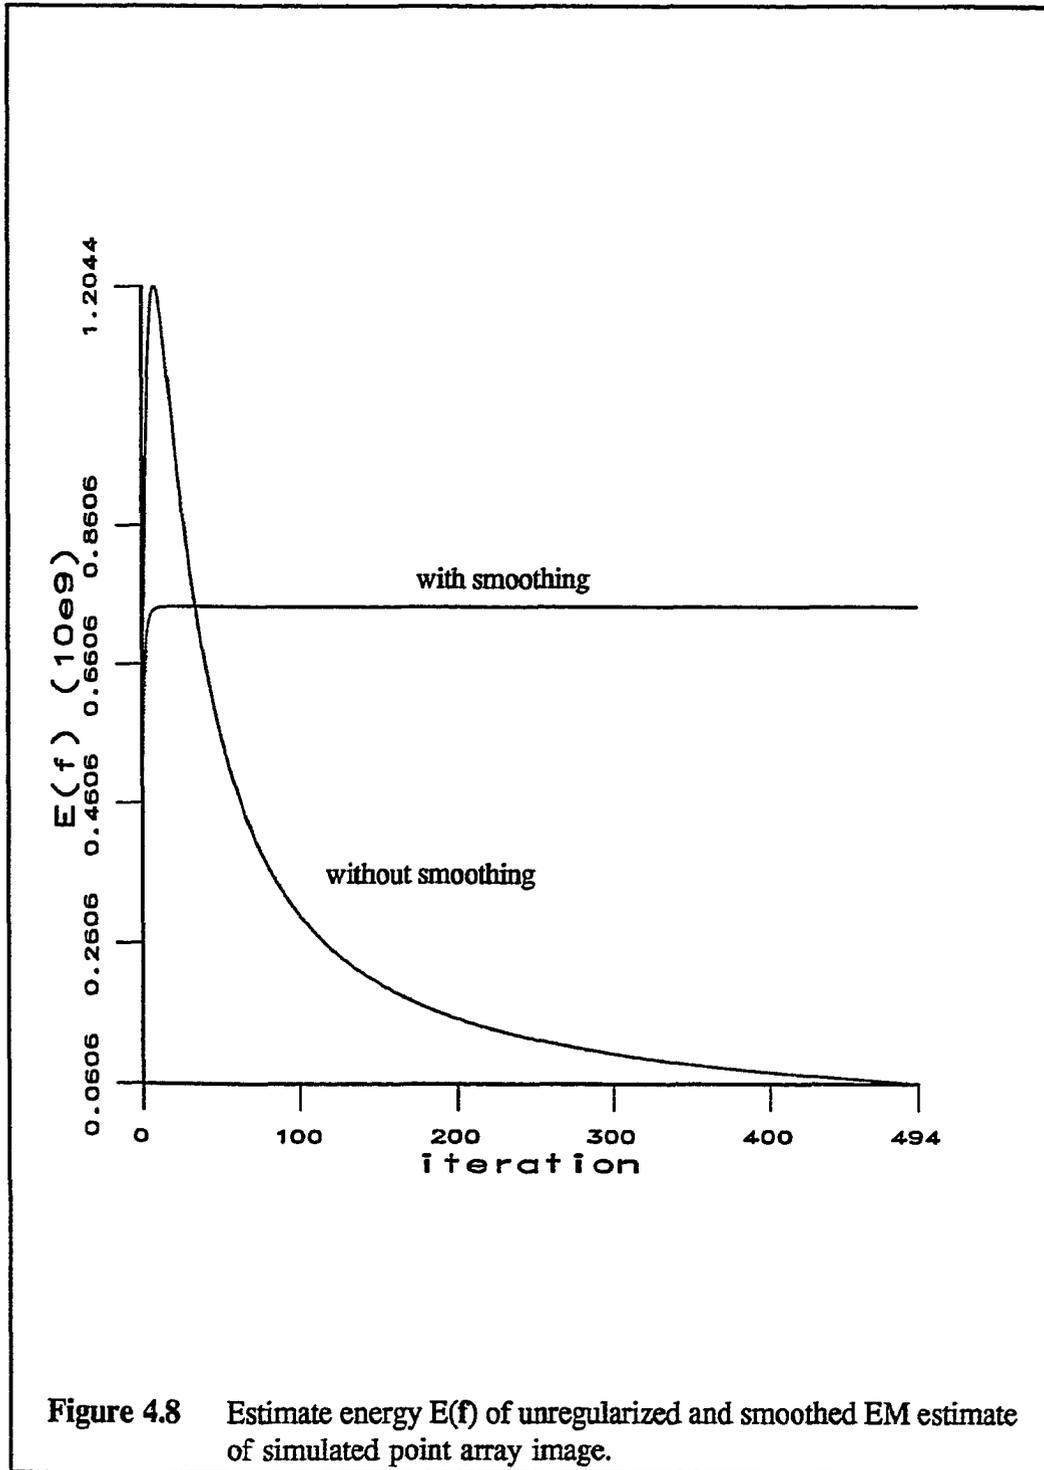




even more rapidly. The next plot in Figure 4.8 shows the great difference between $E(\mathbf{f})$ for the two EM methods. Using unregularized EM, $E(\mathbf{f})$ peaks by the tenth iteration, then decreases through iteration 494. The smoothed EM image $E(\mathbf{f})$ converges to a level lower than the unregularized peak $E(\mathbf{f})$, but higher than the unregularized $E(\mathbf{f})$ at higher iterations. From these results we see that even though $E(\mathbf{g})$ is changing slowly by the 30th iteration, large changes in $E(\mathbf{f})$ continue through many more iterations. The EM algorithm yields lower values for both energy terms when smoothing is disabled for images with high spatial frequency.

The processed image after 500 iterations without smoothing is shown in Figure 4.9. Note that the points in the corners are spread along a diagonal line, but the other points are well resolved. A profile of the processed image at $\hat{q}=34$ after iteration 500 is shown in Figure 4.10. The ideal flux vector and the simulated \mathbf{g} vector are also shown. The differences between the ideal flux and the EM result are best seen at the base of each peak. Note how the EM algorithm recovers the amplitude and narrow width of the points. The simulated response vector \mathbf{g} demonstrates the spread of the system response. When the point array is processed with smoothing enabled, the image maintains the broader image of each point, as shown in the profile through $\hat{q}=34$ in Figure 4.11. A smaller standard deviation for the smoothing function would alleviate this deficiency.

The point array study showed that the best estimate of the true flux vector is obtained from the central region of the camera. For a high spatial frequency flux



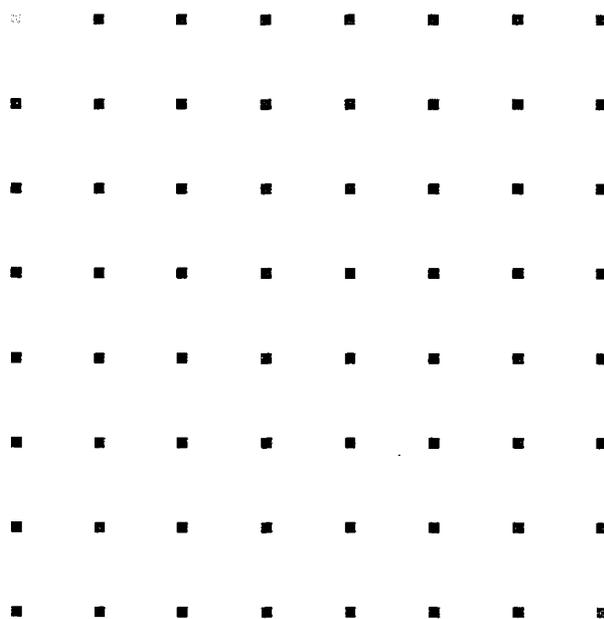
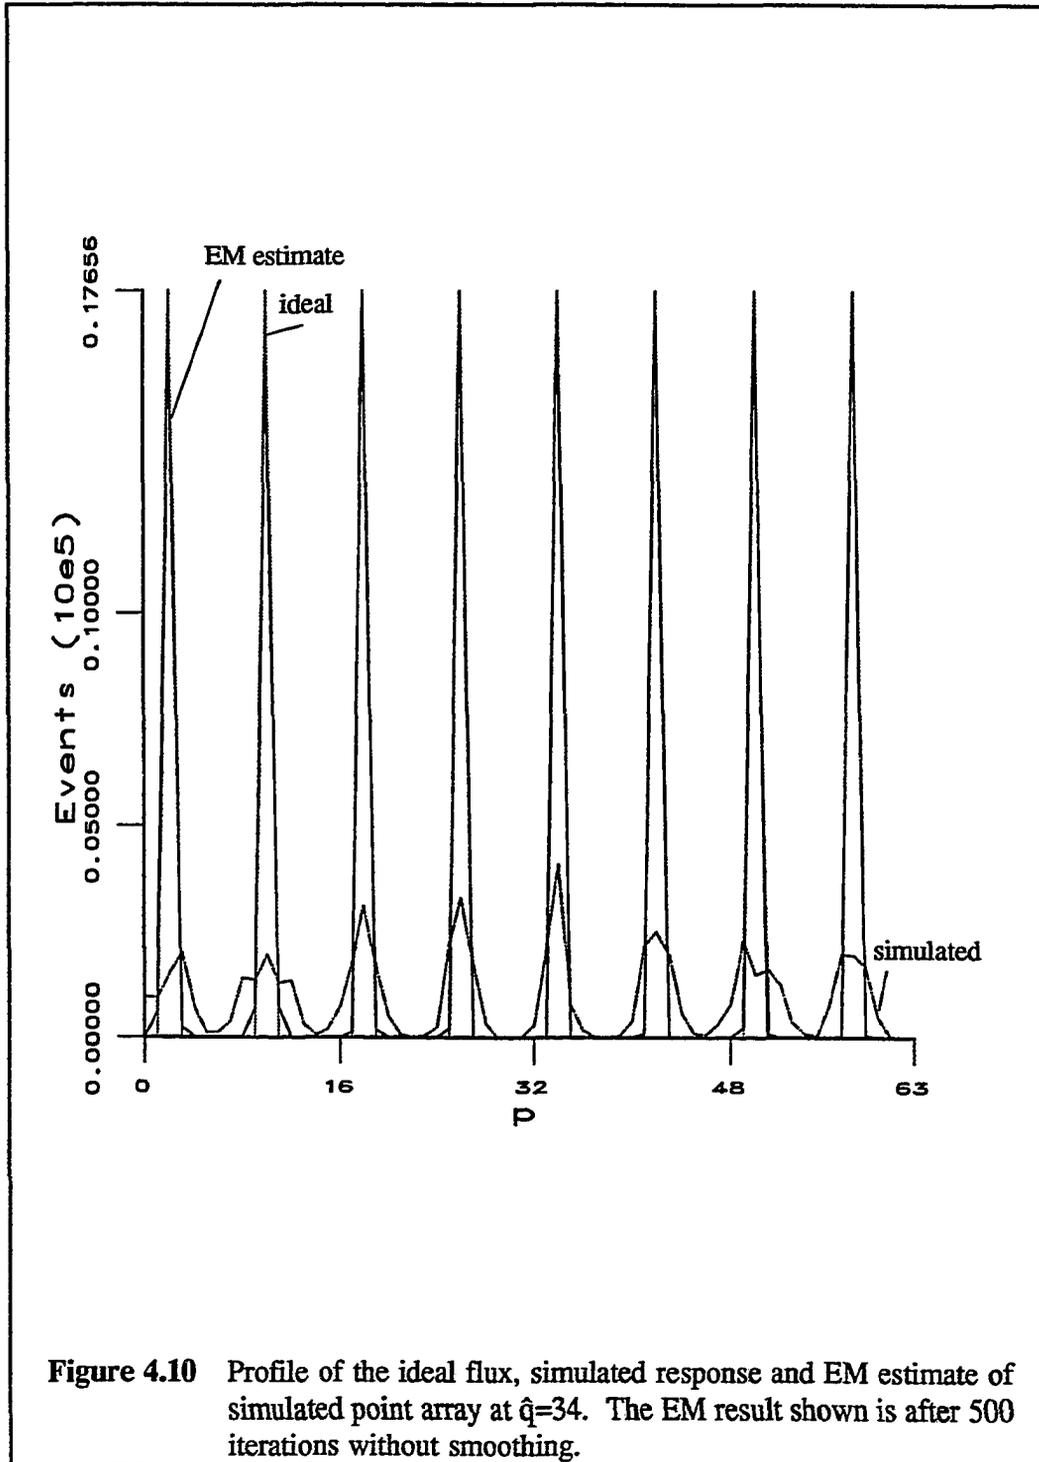


Figure 4.9 Image of the EM estimate of the simulated point array after 500 iterations.



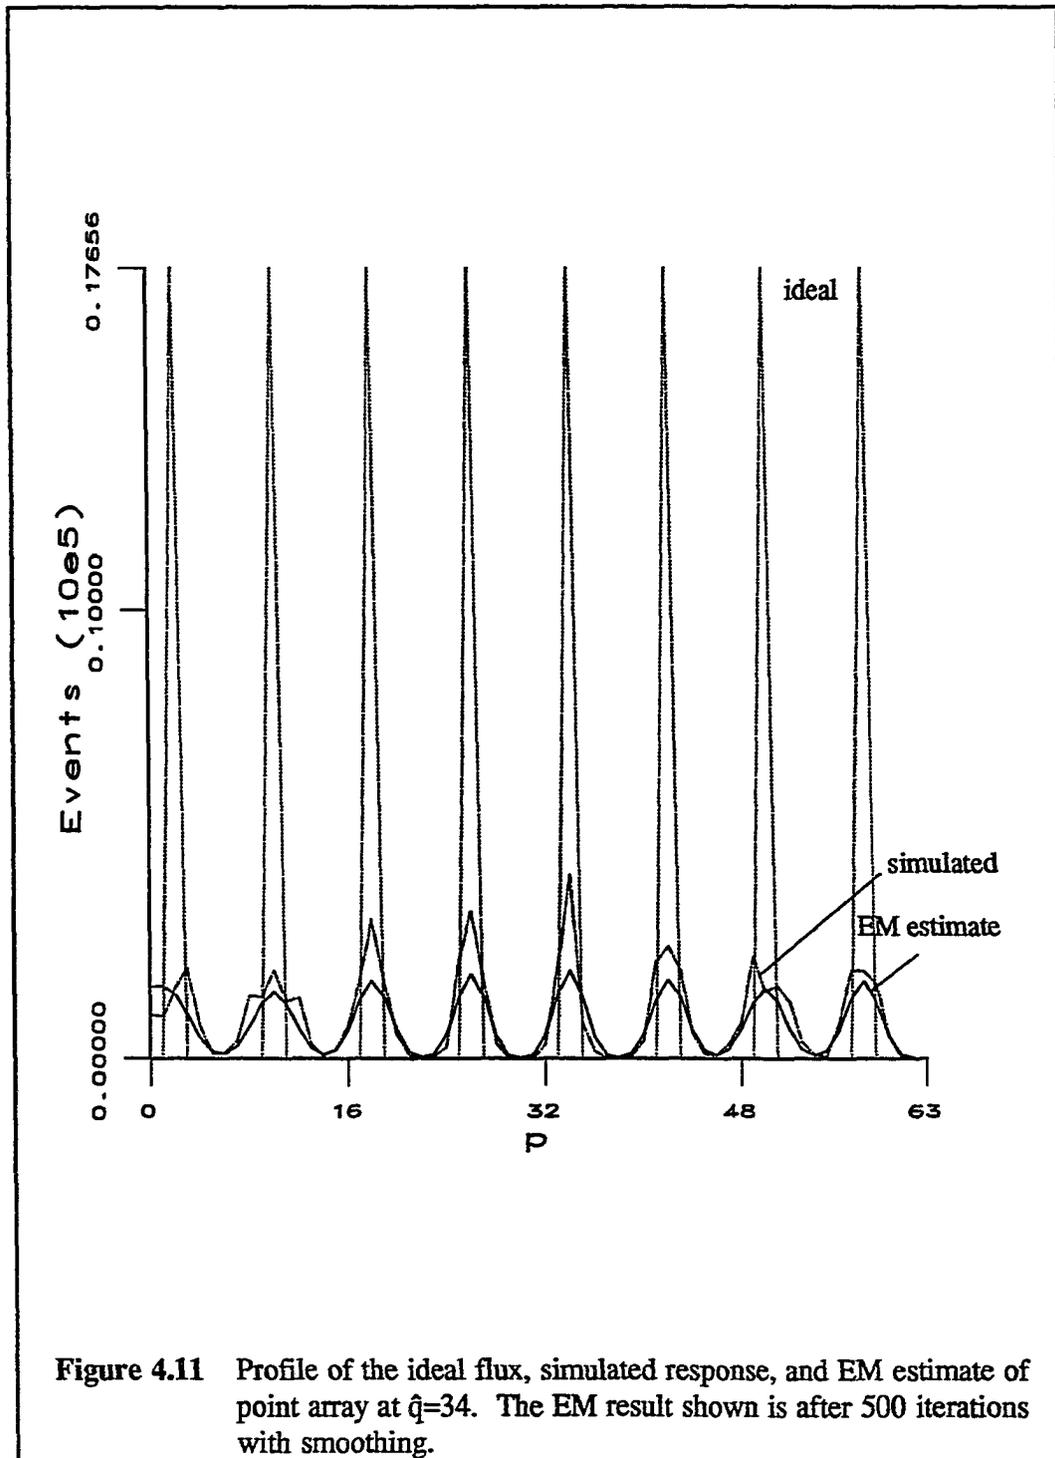


Figure 4.11 Profile of the ideal flux, simulated response, and EM estimate of point array at $\hat{q}=34$. The EM result shown is after 500 iterations with smoothing.

pattern, unregularized EM performed better than the smoothed EM. The measurement energy converges quickly, but the estimate energy continues to decrease after the initial peak.

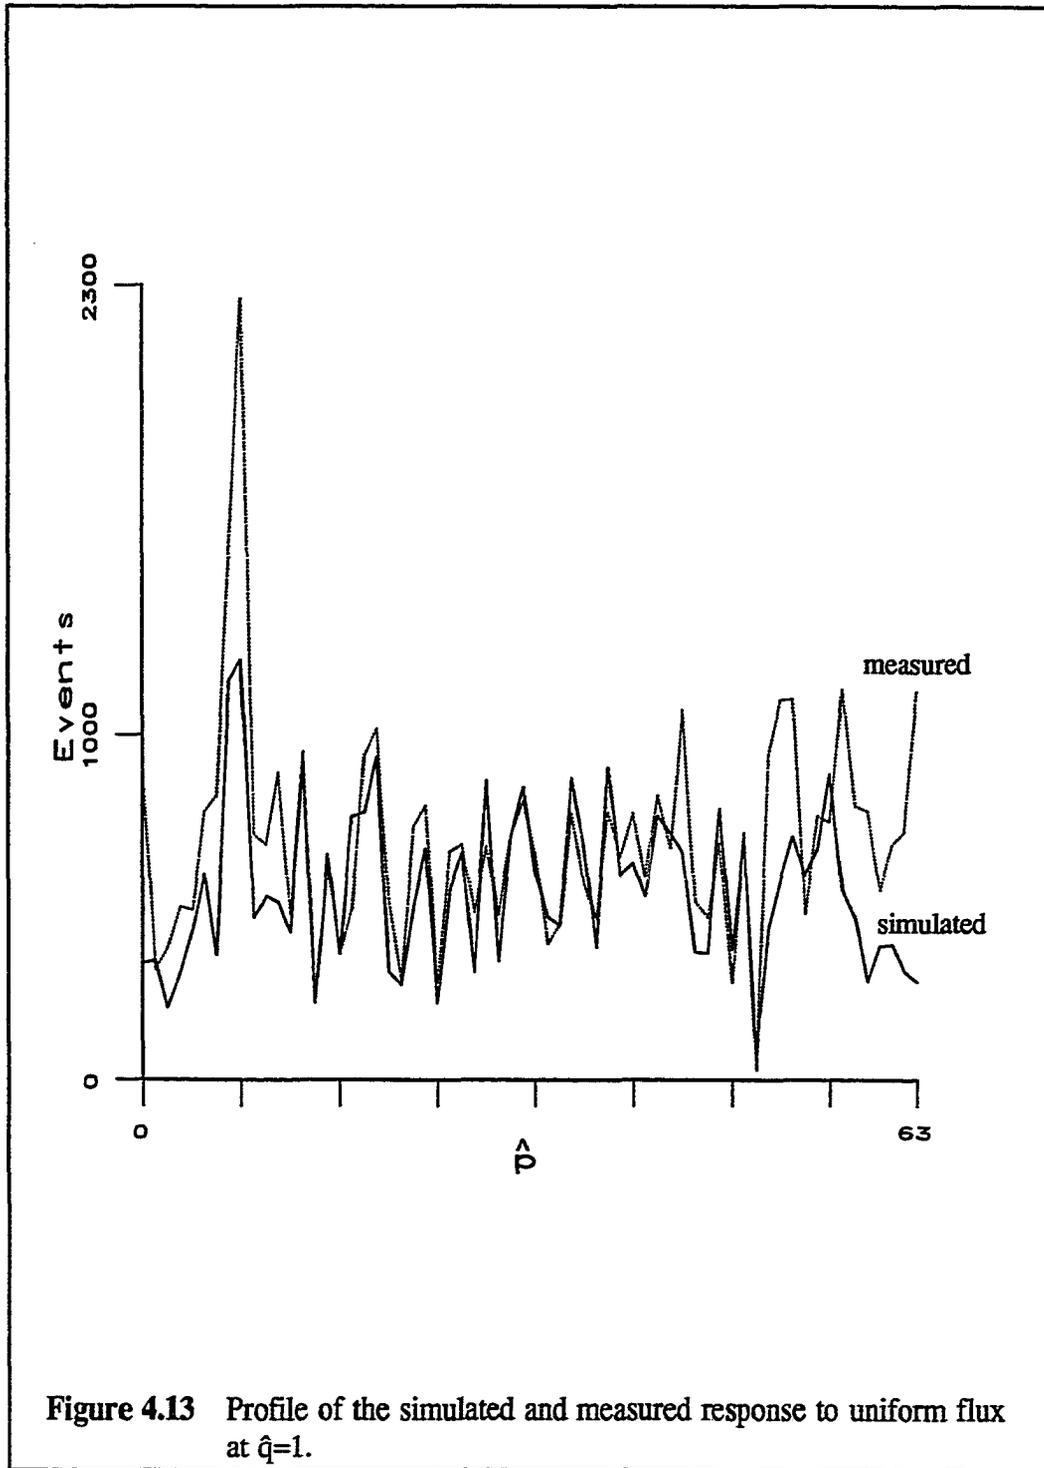
Flood Image Results

The flood image study examined the EM algorithm for processing images with low spatial frequency. Flood images demonstrate well the deficiencies in the modular camera response, as shown in Figure 4.12. Note that the variations in the image increase in amplitude and frequency in the corners and along the edges of the crystal. A camera response that was a true representation of the flux would not demonstrate these position-dependent variations.

We computed the simulated measurement vector in Figure 4.12 by multiplying the matrix \mathbf{H} by a constant flux vector, then adding noise as with the point array. The simulated image variance increases toward the corners in a pattern that is similar to that of the measured flood image. The amplitude of the variations, however, is lower in the simulation. This deficiency of the model can be observed by examining a profile through both images at coordinate $\hat{q}=1$, which is near the top edge, as shown in Figure 4.13. The measured flood image has greater amplitude at $\hat{p}=0$ than at $\hat{p}=32$. Note the similarities between the measured and simulated response when \hat{p} is close to 32. These similarities encouraged us to proceed with simulations, to determine how well the EM algorithm would recover the ideal flux.



Figure 4.12 Top: Simulated response to the ideal uniform flux.
Bottom: Measured response to the flood calibration.

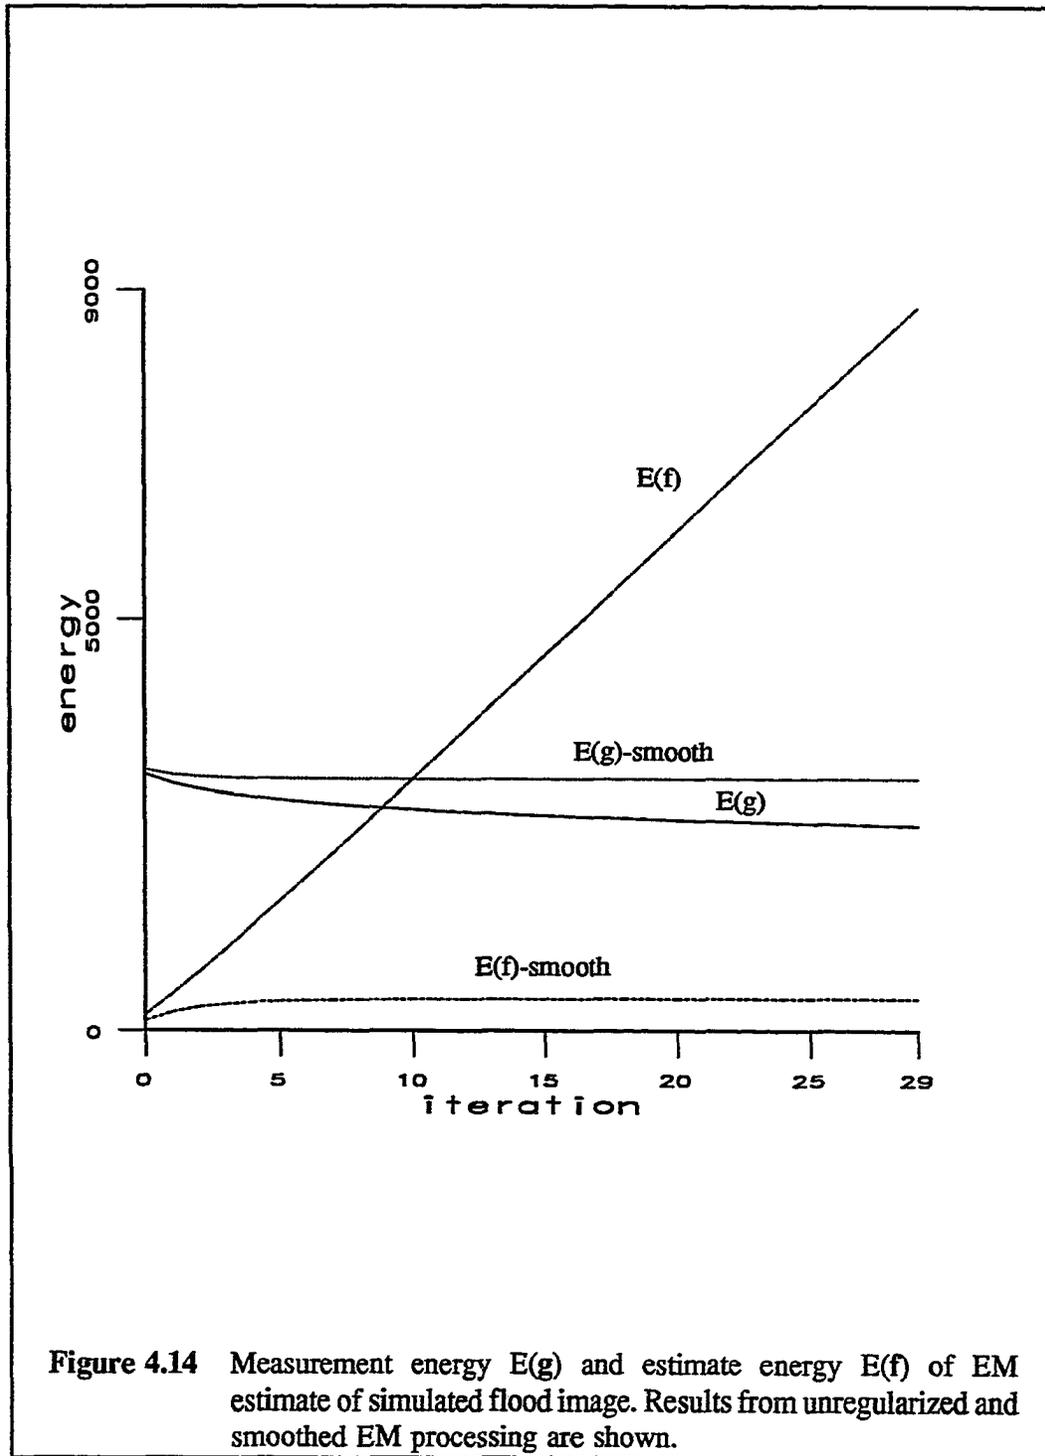


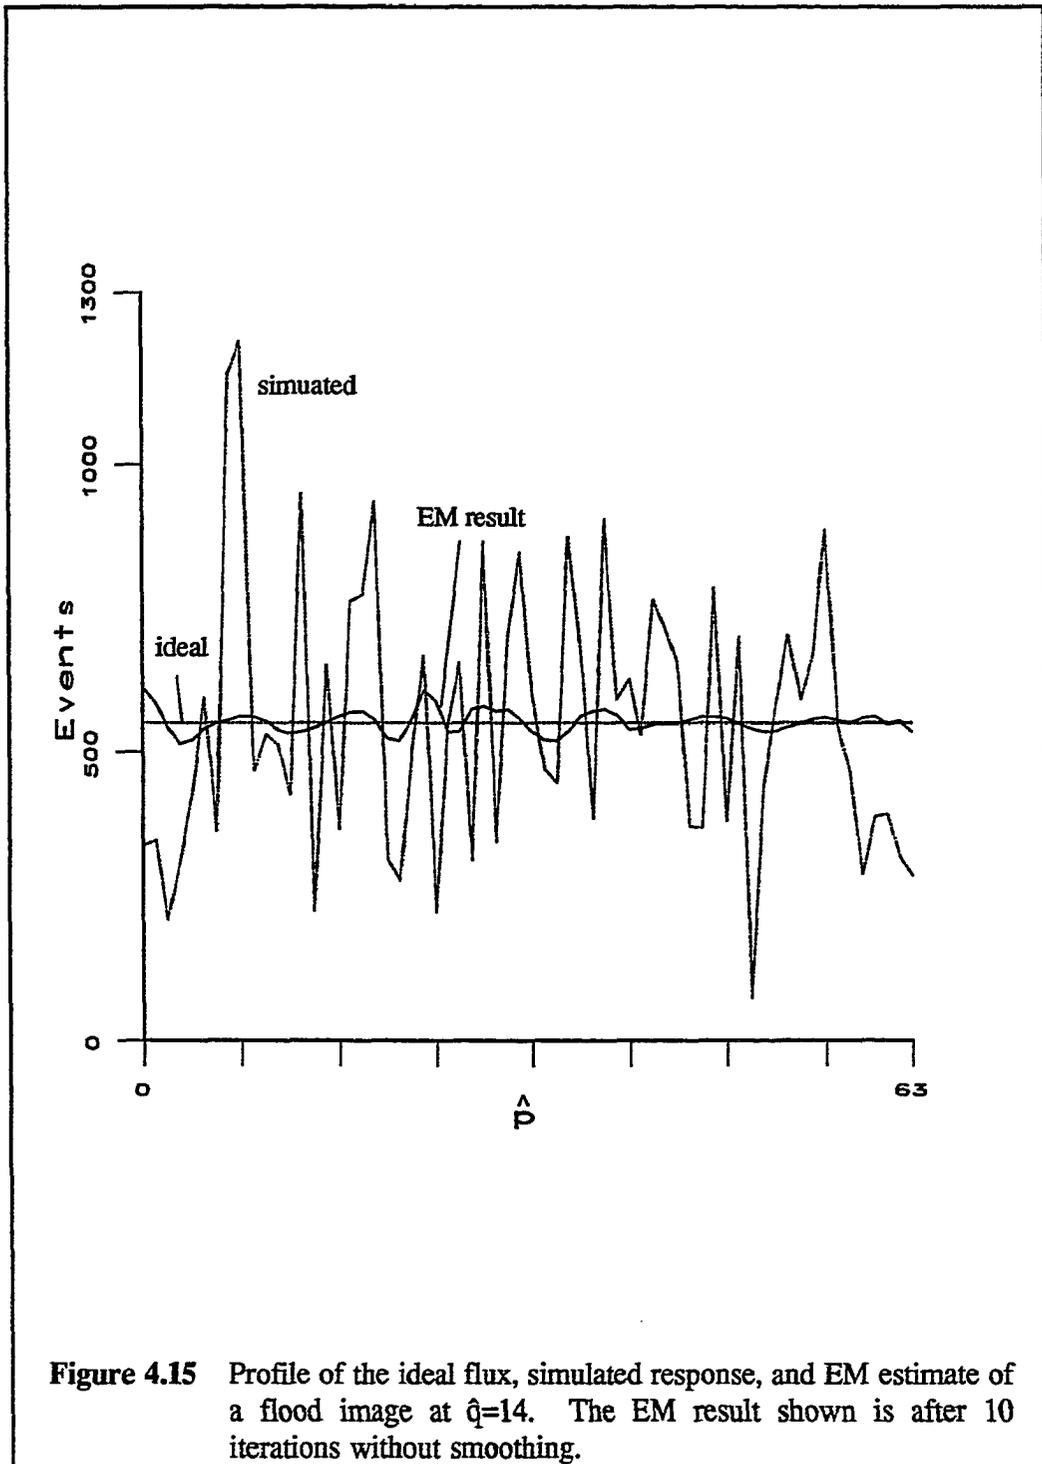
The energy functions for the flood images are shown in Figure 4.14. The measurement energy is higher and converges more quickly when smoothing is enabled. The estimate energy demonstrates the advantage of smoothing the flood image. For this object the unregularized EM estimate diverges from uniform linearly with iteration number.

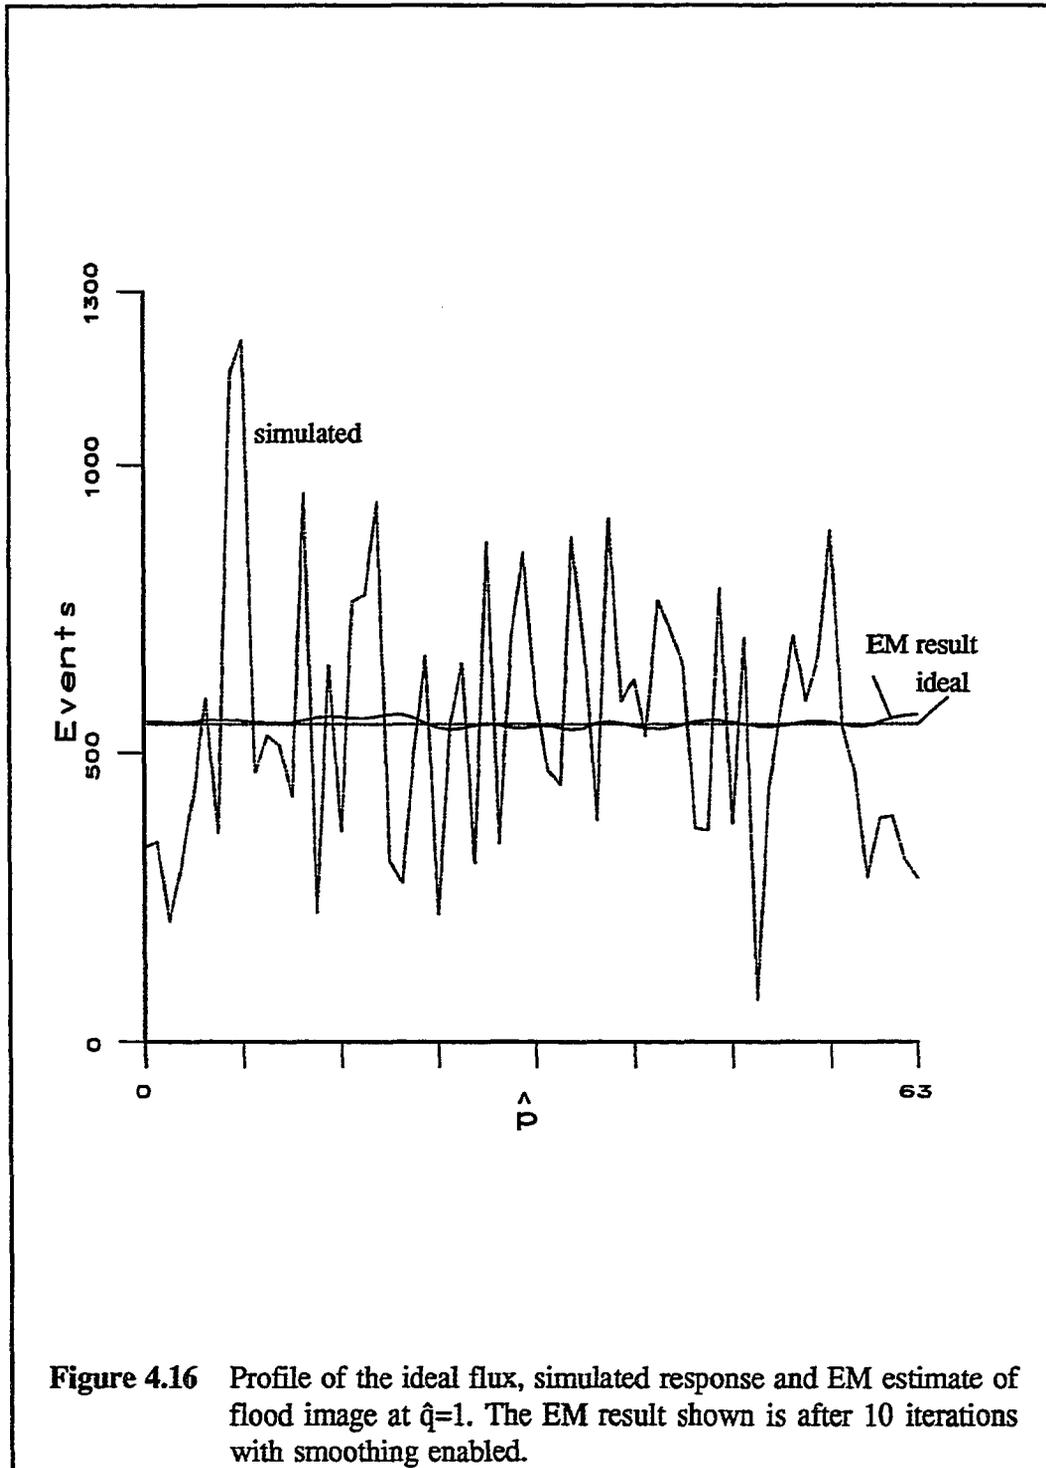
The smoothed EM algorithm recovered the uniform flux from the simulated image very well. The ideal flux was a 549 at all pixels. The mean of the simulated response was 516, with standard deviation of 116.6. Note that this standard deviation is significantly greater than the value expected from Poisson statistics (23.4). After EM processing 10 iterations, the mean in the estimate was 549 for both unregularized and smooth algorithms. For the unregularized algorithm, the standard deviation increased with each iteration. For example, the standard deviation is 19.2 at iteration 10 and 27.8 at iteration 20. The standard deviation between iteration 10 and 20 is close to the predicted value. A profile through the flood image at $\hat{q}=14$ after 10 iterations is plotted in Figure 4.15. Here we show the ideal constant flux, the simulated camera response and the EM result. When smoothing is enabled, the standard deviation is 7.1 after 10 iterations. The smoothed EM result after 10 iterations is shown in Figure 4.16.

Blur Image Results

The blur case was simulated so that an image with more typical spatial







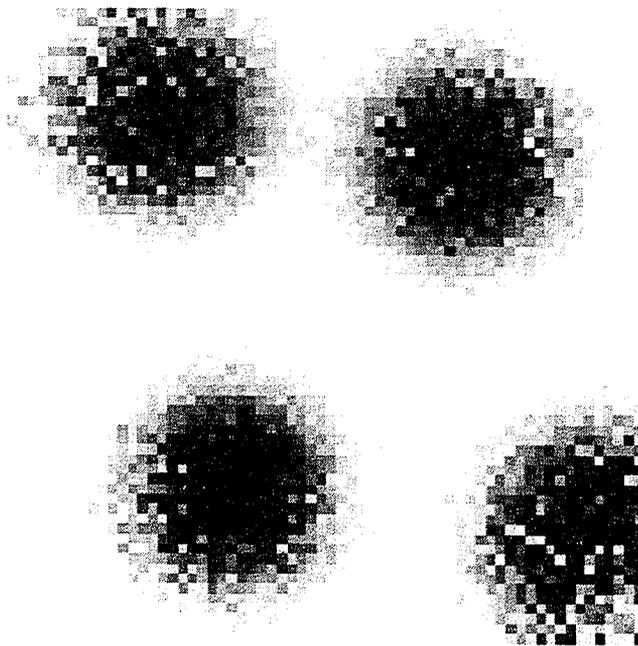
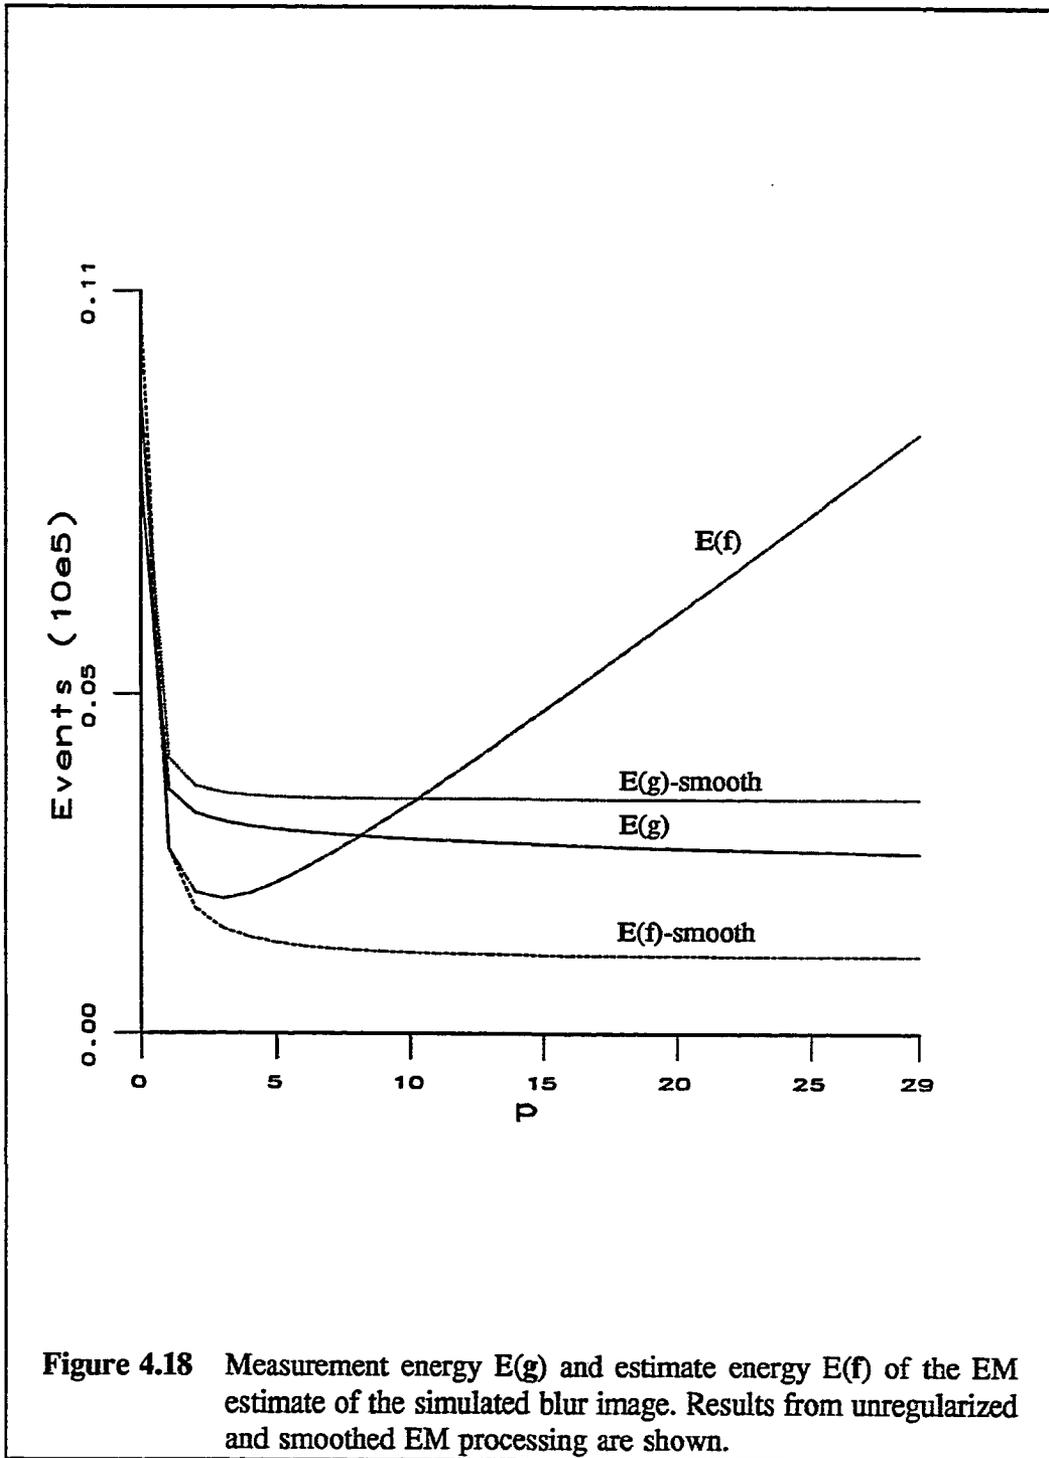


Figure 4.17 Image of the simulated response to the ideal blur flux vector.

frequency could be examined. The point array and flood images represent images with spatial frequencies that are very different from the projection data collected with the brain imaging system. We have demonstrated the dependence of estimate energy on smoothing and the spatial frequency of the object. With a point array, smoothing prevented the EM algorithm from recovering the true high-frequency, high-amplitude point images. Best results for the point array were found by disabling smoothing. With the flood image, smoothing prevented unwanted oscillations in the estimate from being amplified. In both cases smoothing decreased the number of iterations needed before the energy terms became nearly constant. Because smoothing sacrifices resolution, however, we wished to examine its effect on an image that was more typical of projections in the system. The simulated response to the ideal blur vector is shown in Figure 4.17. The checkerboarding pattern in the corners is typical of such an image. The energy functions are plotted as a function of iteration number in Figure 4.18. The measurement energy quickly reaches a point where it changes slowly, with or without smoothing. The estimate energy is higher for the unregularized EM, and diverges after the fourth iteration. With smoothing enabled, the estimate energy decreases monotonically. These results suggest that for projection data, we should enable smoothing, and should be able to stop the EM algorithm by the 10th pass.

The recovered estimate after 10 iterations is shown in Figure 4.19. The checkerboarding is removed and the blurs are well defined. The profile of the EM



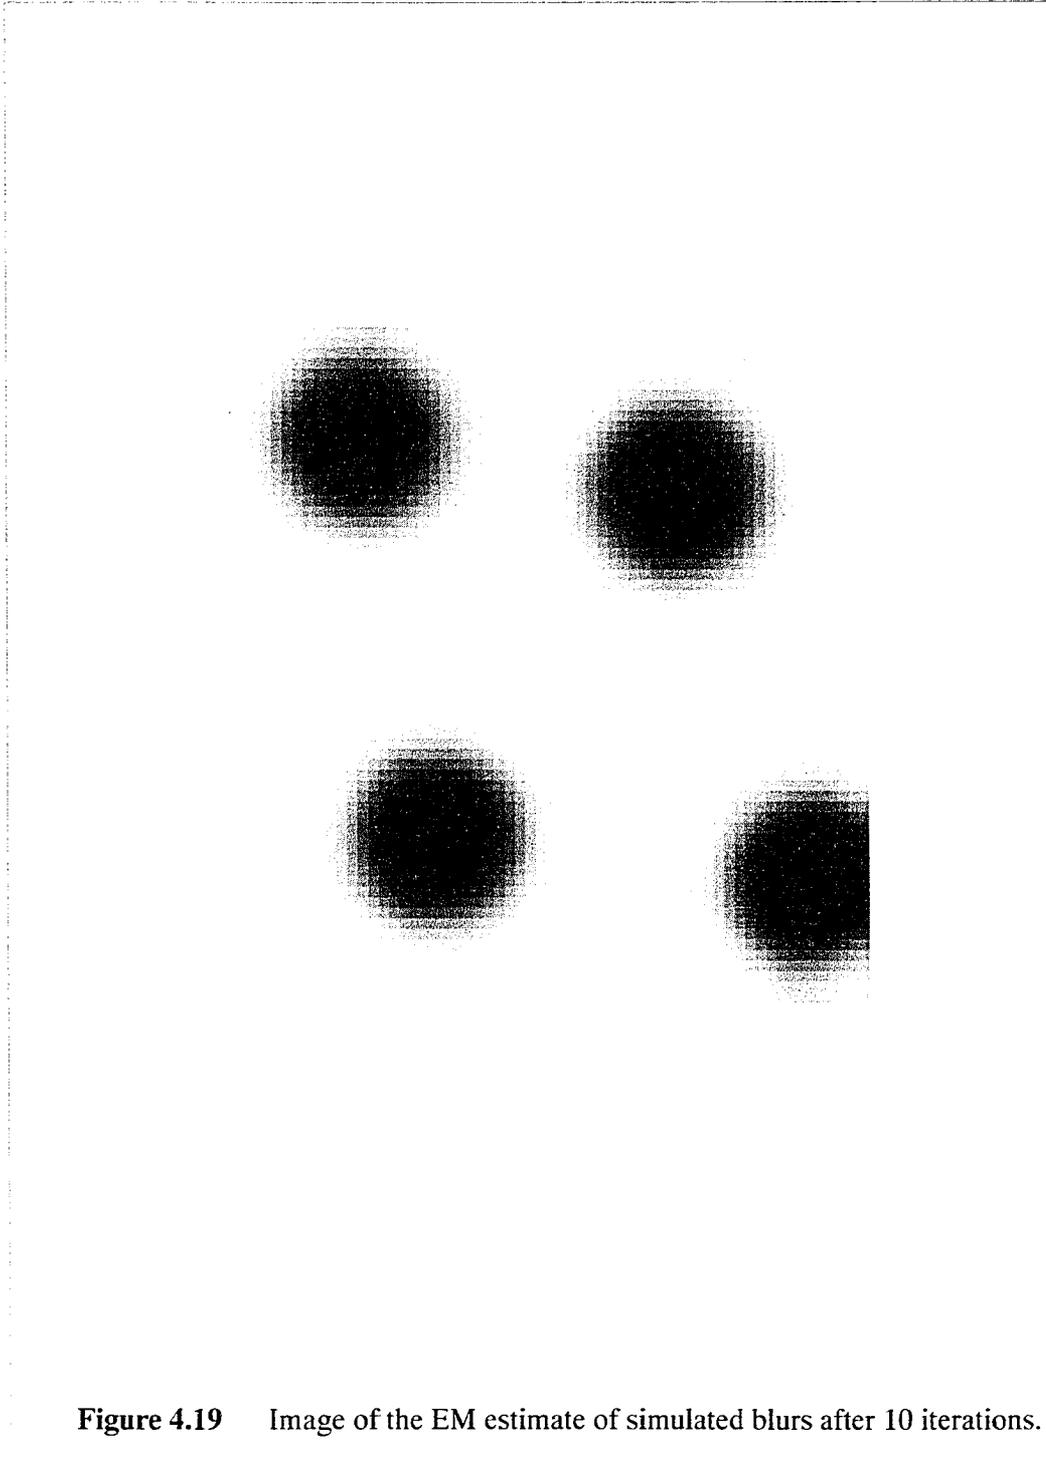


Figure 4.19 Image of the EM estimate of simulated blurs after 10 iterations.

estimate computed without smoothing at $\hat{q}=2$ is shown in Figure 4.20. The ideal flux and simulated response are also shown for comparison. The profile at $\hat{q}=2$ of the EM estimate computed with smoothing enabled is shown in Figure 4.21.

Experiments with Low Counts

We are interested in using the EM algorithm to recover a better estimate of the flux incident on the crystal. One of the degradations in the measured data is photon noise. While measurements with one million or more counts per camera yield good results, we would like to reduce the number of counts needed to reconstruct objects with the brain imager. Reduced counting requirements would allow us to have shorter imaging times for dynamic studies, and allow us to minimize the dose of radiopharmaceutical to the patient.

We examined the performance of the EM algorithm by repeating the simulations with the blur flux vectors with fewer counts in the ideal image. For each example, we enabled smoothing and iterated 15 times. Profiles of the processed images are shown for two cases in Figure 4.22. Decreasing total counts from 1,000,000 to 50,000 did not prevent the algorithm from recovering the Gaussian flux estimate.

EM Processing Direct PMT Signal Histograms

To evaluate the potential for using direct PMT signals rather, than the (\hat{p}, \hat{q}) estimate, as the index for g , we simulated the flood response using the forward model

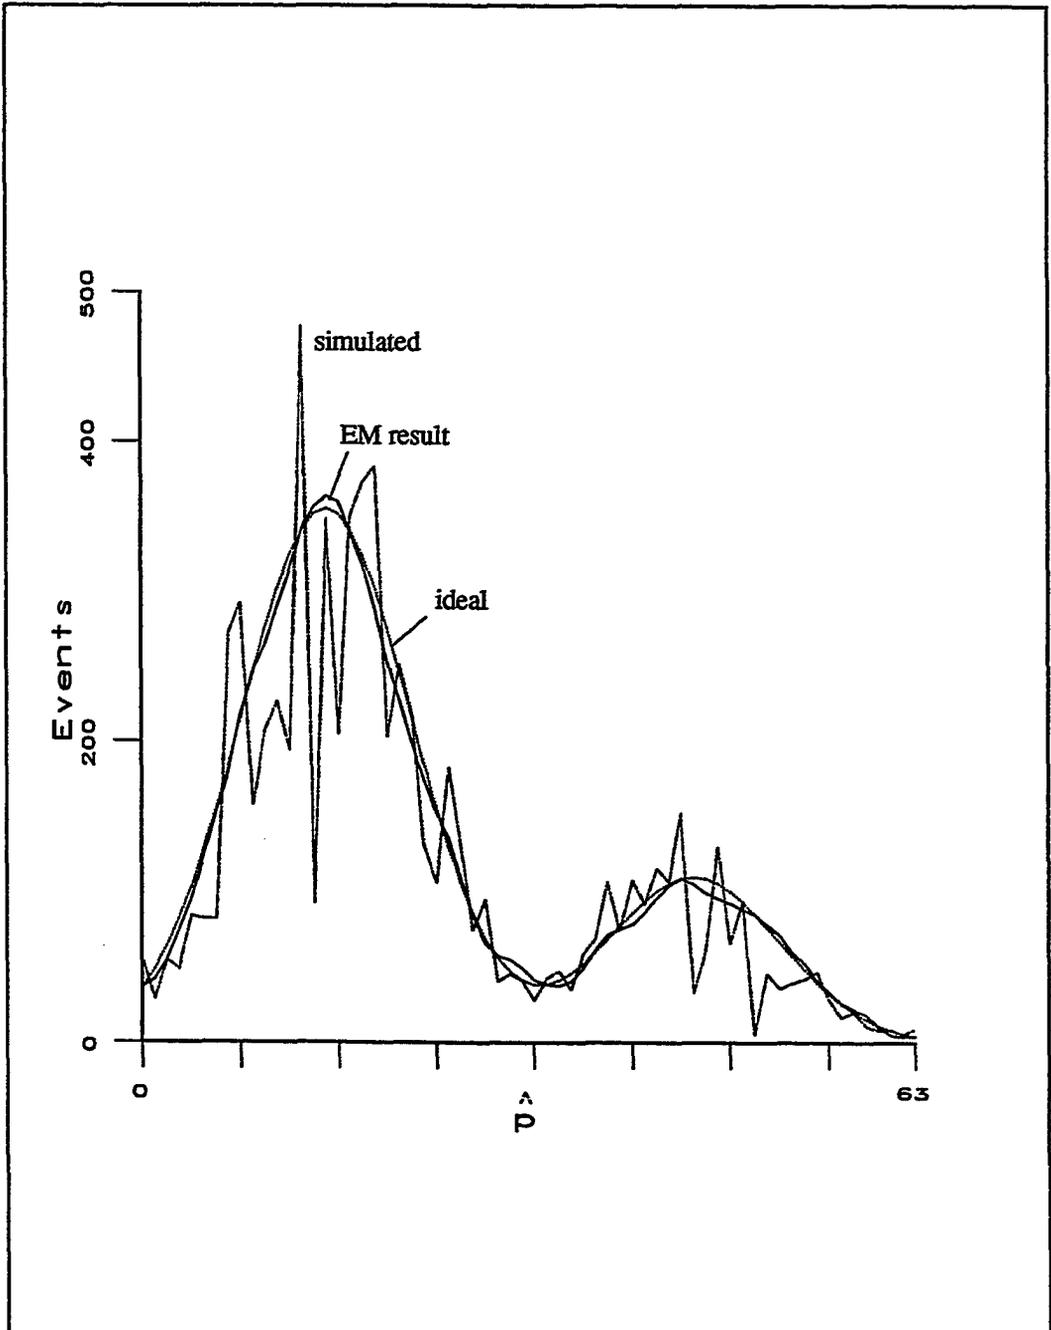
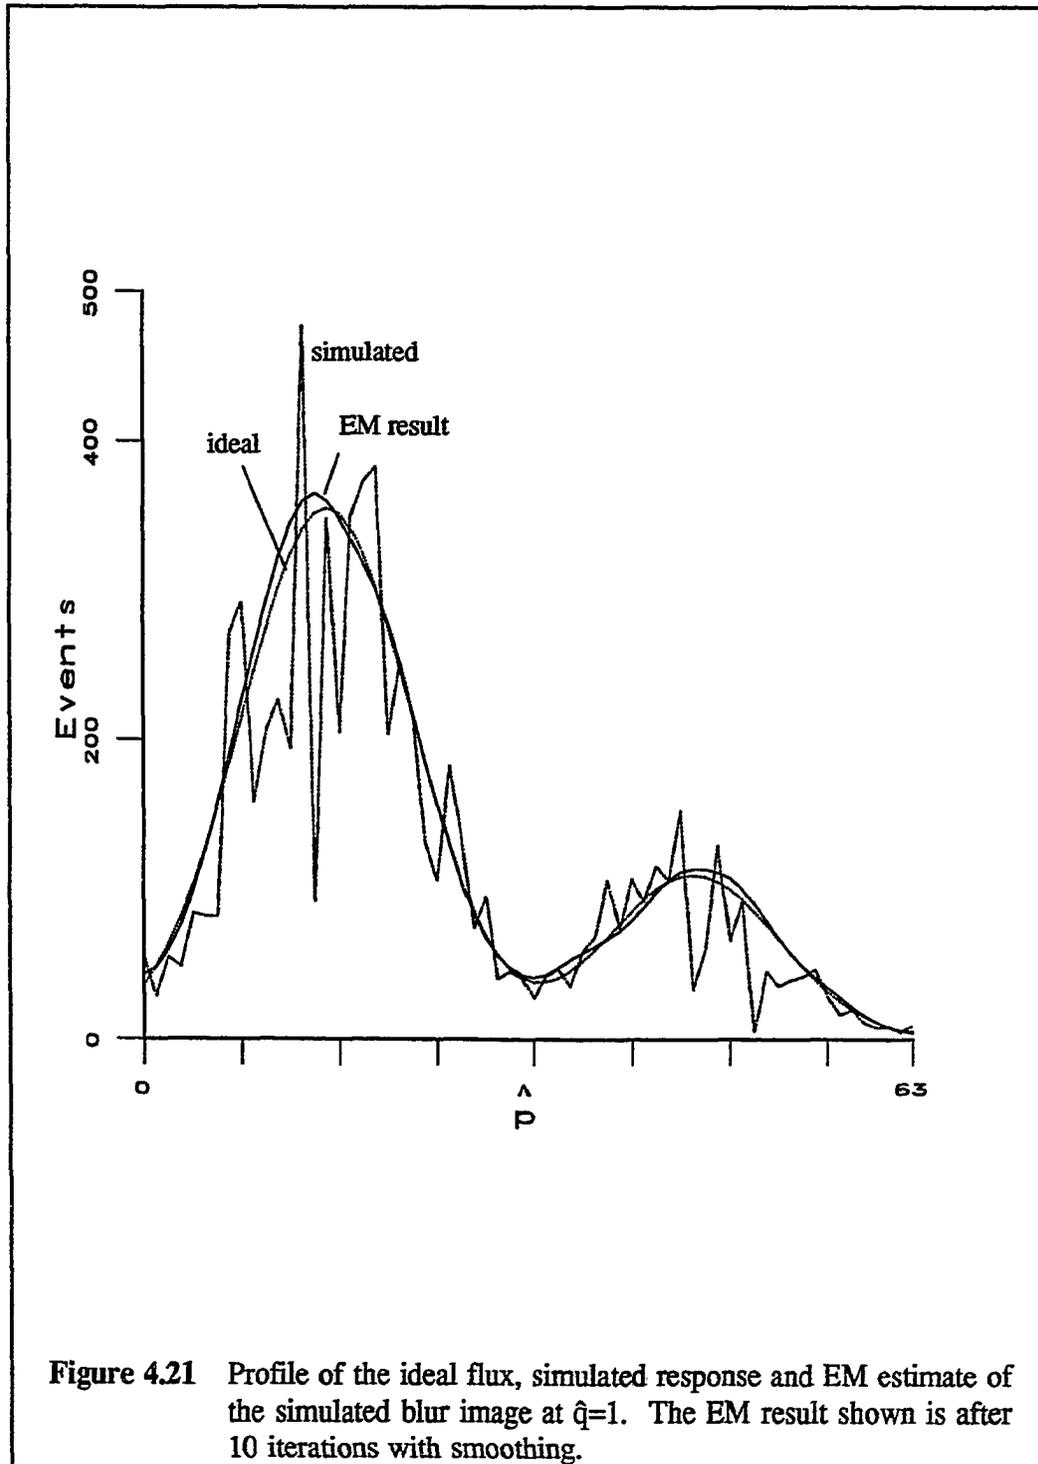


Figure 4.20 Profile of the ideal flux, simulated response and EM estimate of the simulated blur image at $\hat{q}=1$. The EM result shown is after 10 iterations without smoothing.



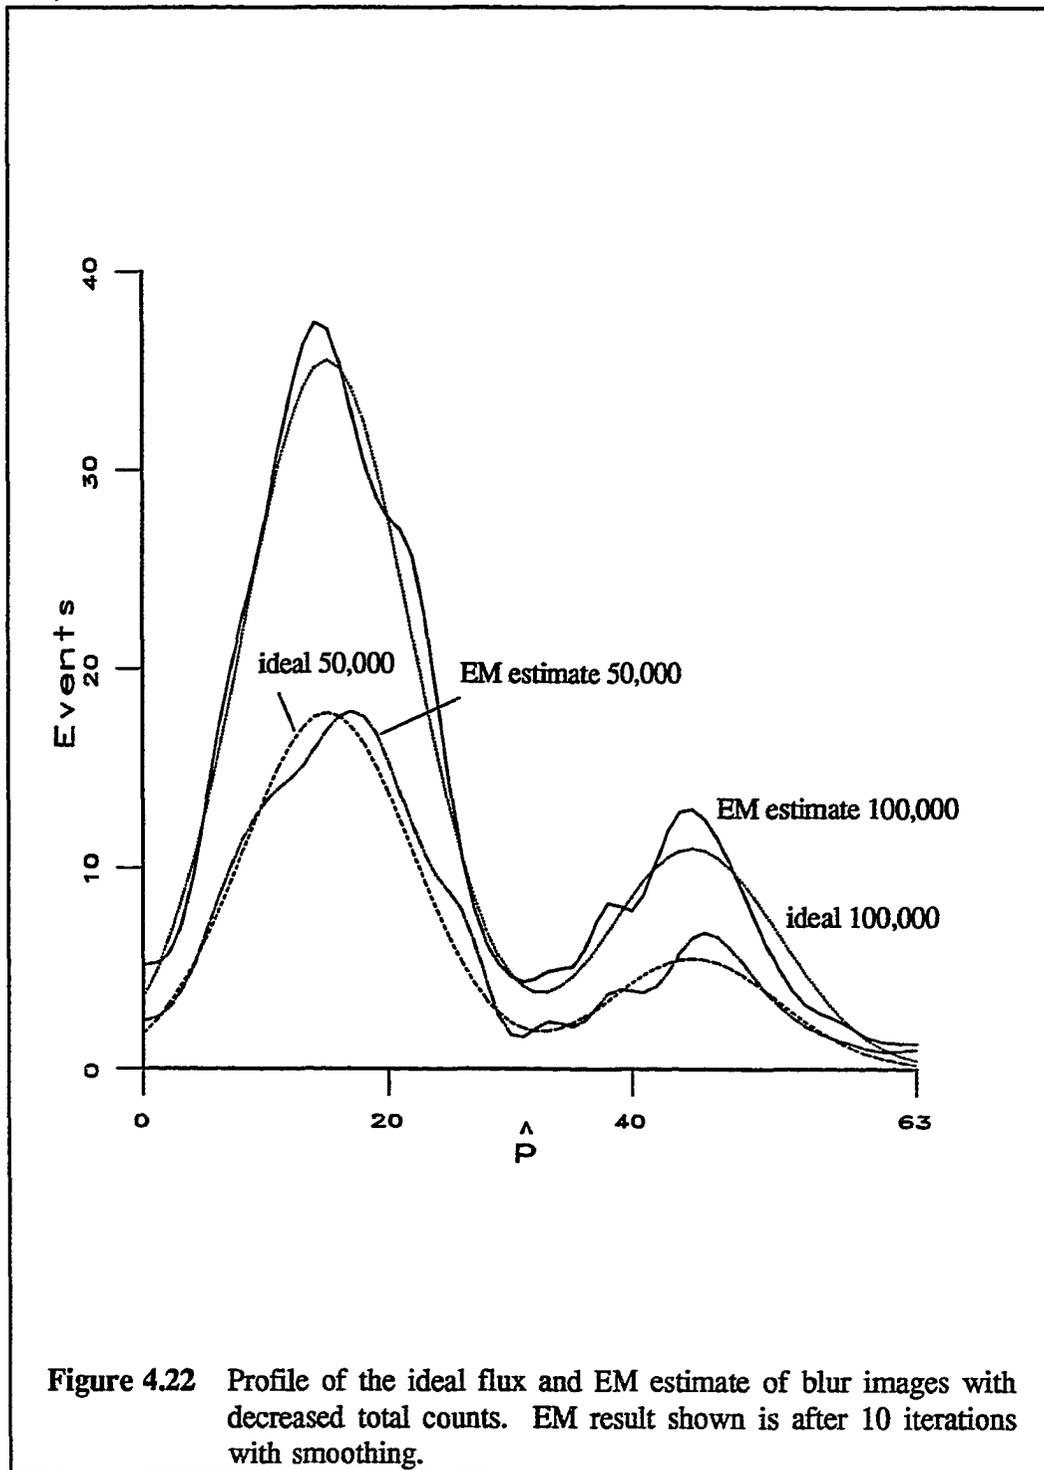


Figure 4.22 Profile of the ideal flux and EM estimate of blur images with decreased total counts. EM result shown is after 10 iterations with smoothing.

$$g(ABCD) = \sum_{p=0}^{63} \sum_{q=0}^{63} H(ABCD,(p,q)) f(p,q) + n(ABCD)$$

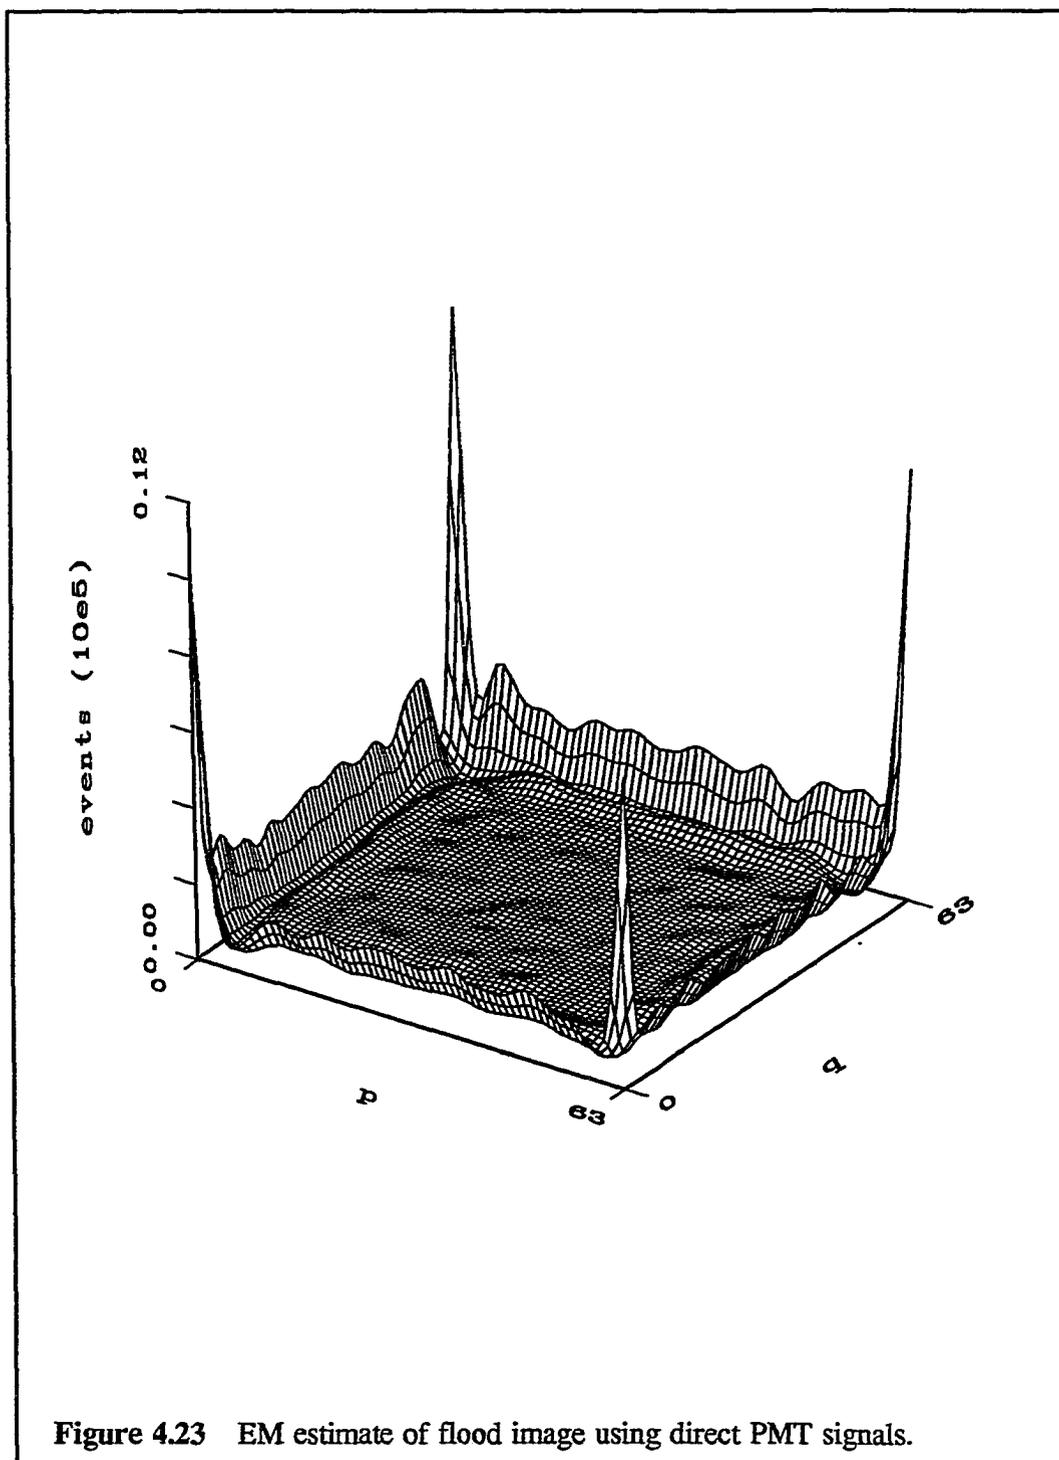
where each element of $H(ABCD,(p,q))$ is computed as $\Pr(ABCD|p,q)$ using the method shown in equation (4.2).

We processed the simulated flood histogram to determine how well we could estimate the uniform flux without the intermediate step of LUT mapping. The image estimated by EM processing after 10 iterations is shown in Figure 4.23. Note that the uphill estimates accumulate along the edges of the crystal. If we compare this estimate to the one found using the crystal coordinates as the measurement index, no advantage is apparent. We opted to concentrate on the forward model which uses crystal coordinates for further studies, since the direct PMT signal model required longer computation time and retained the unwanted uphill estimates in the processed image.

Results of EM Processing the Measured Camera Response

Encouraged by the simulation studies, we examined how EM processing would perform on images measured by the modular cameras during calibration. The calibration images included the flood and point array images described above. We also used the algorithm to process projection images from all 24 modular cameras.

The EM estimates described below were computed with smoothing enabled. While we know that the high spatial frequency in the point array is best recovered without smoothing, we chose to use one algorithm for all types of measured data.

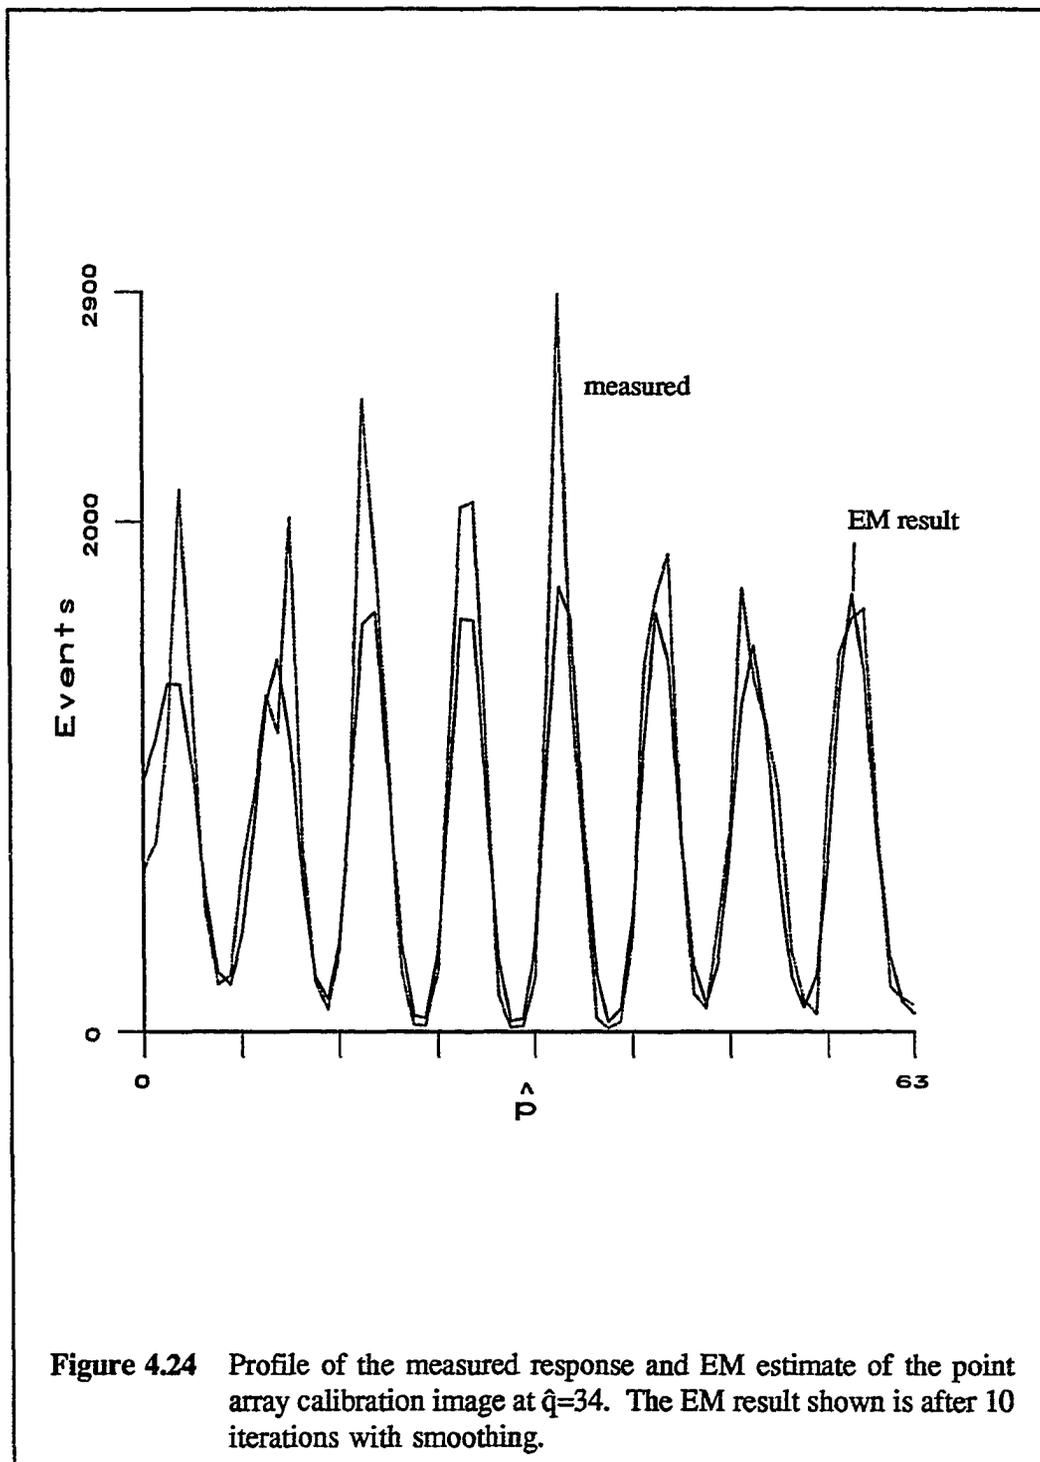


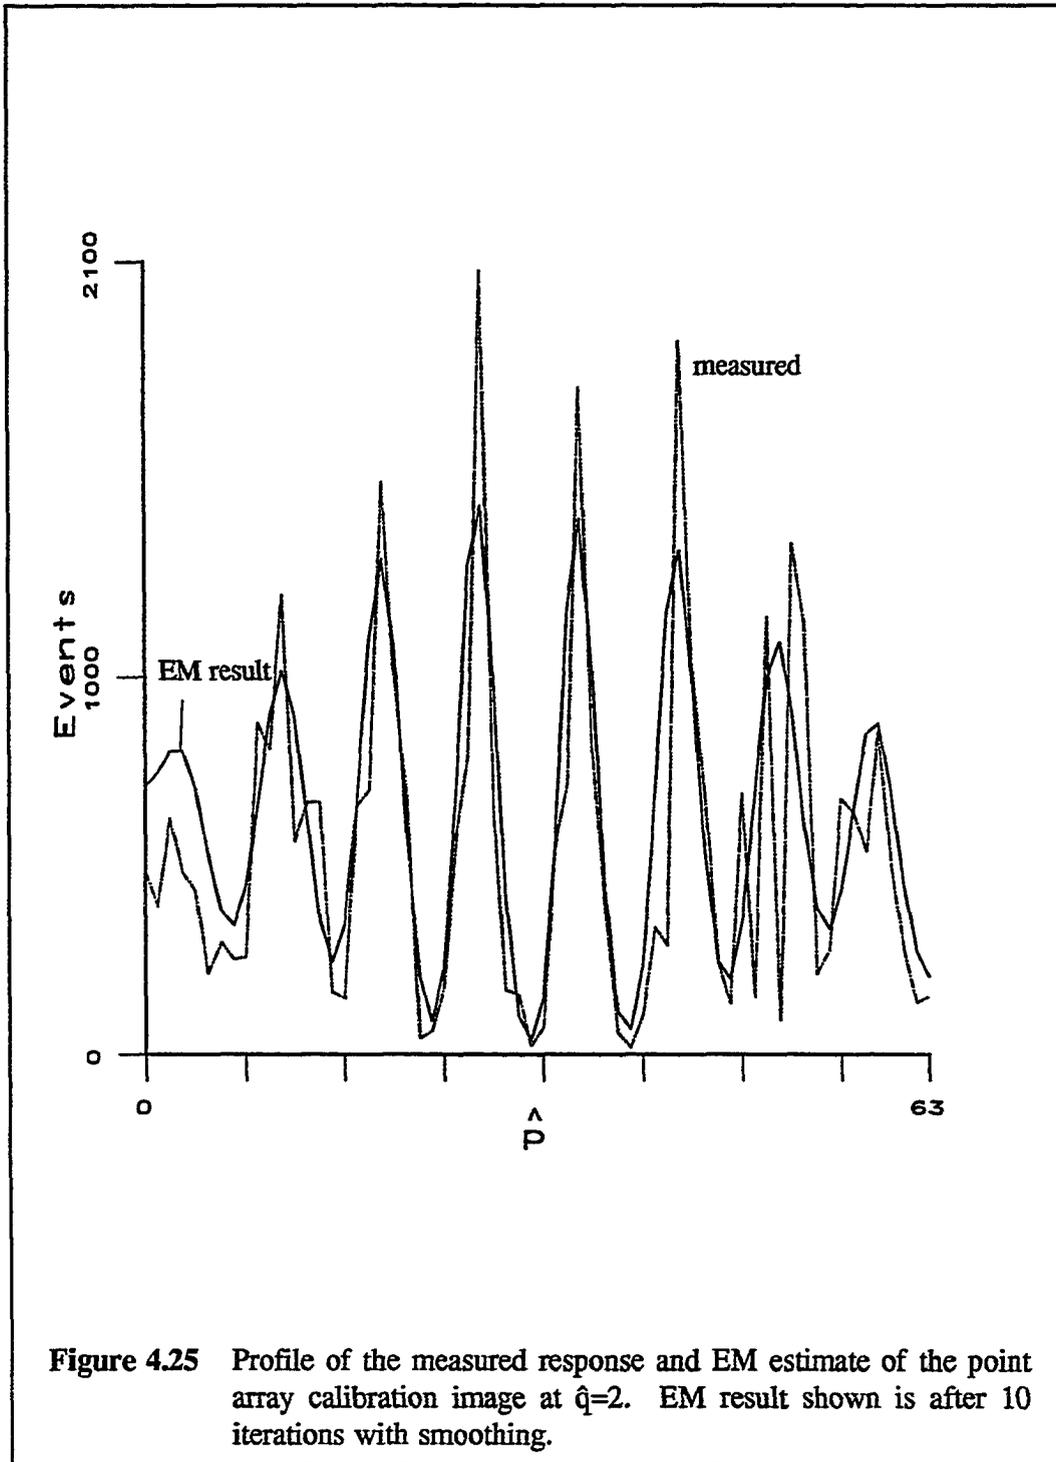
Since the smoothing option decreased the estimate energy for both the flood and blur simulations, and since the spatial frequency of the projection data is similar to that of the blur simulation, we enabled smoothing for all the measured images.

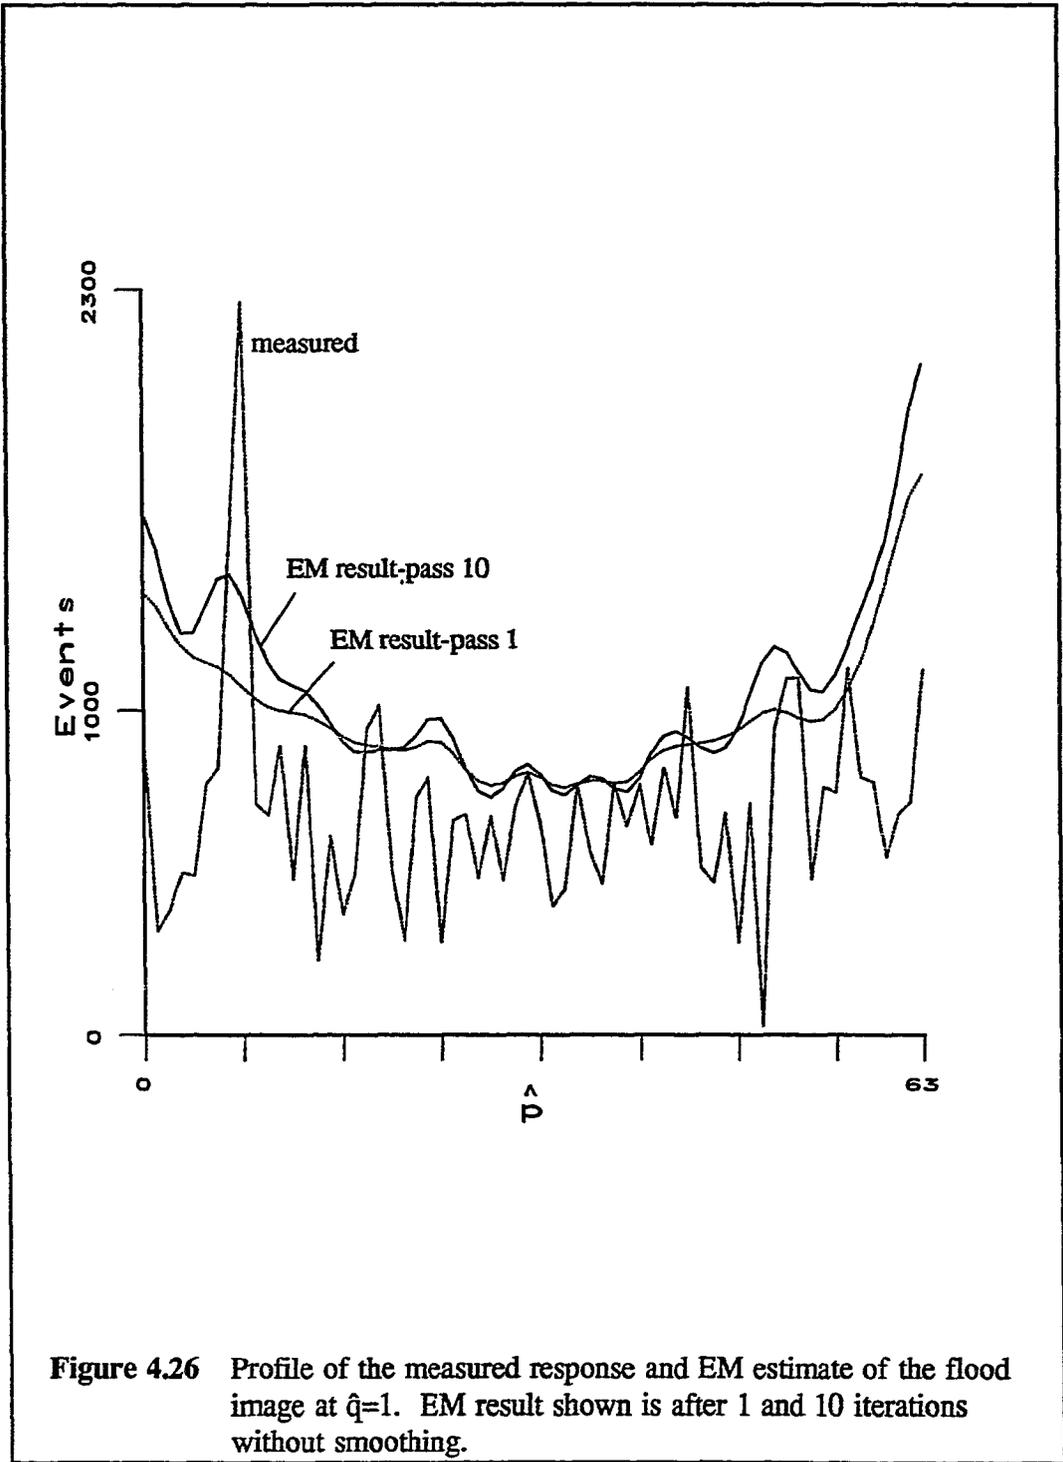
Calibration Image Results

The point array images were processed for 30 iterations. The profile through $\hat{q}=34$ is shown for the measured and recovered point array in Figure 4.24. The profile through the points at $\hat{q}=34$ shows that the amplitude of the peaks are more uniform after processing. The width of the peaks near the edge of the camera changes more than the width of the central peaks. The profile through $\hat{q}=2$ in Figure 4.25 demonstrates that the algorithm locates the regularly spaced peaks from the noisy corners and edges. The amplitudes of the peaks decrease toward the edge of the crystal.

We applied the EM algorithm to the measured flood image for 30 iterations to recover the estimate of the flood flux. A profile through $\hat{q}=1$ is shown in Figure 4.26 for the measured response image, the estimate after the first pass and the estimate after pass 10. (The estimate after pass 10 does not change enough to show in the profile.) Note that the amplitude of the flux estimate increases at the edges of the crystal. Since we know that the flux incident on the crystal is uniform, the error near the edges indicates that the forward model for camera response is inadequate in those regions. The discrepancy between the modelled and measured response was also apparent from the comparison with the simulated flood in Figure 4.13.





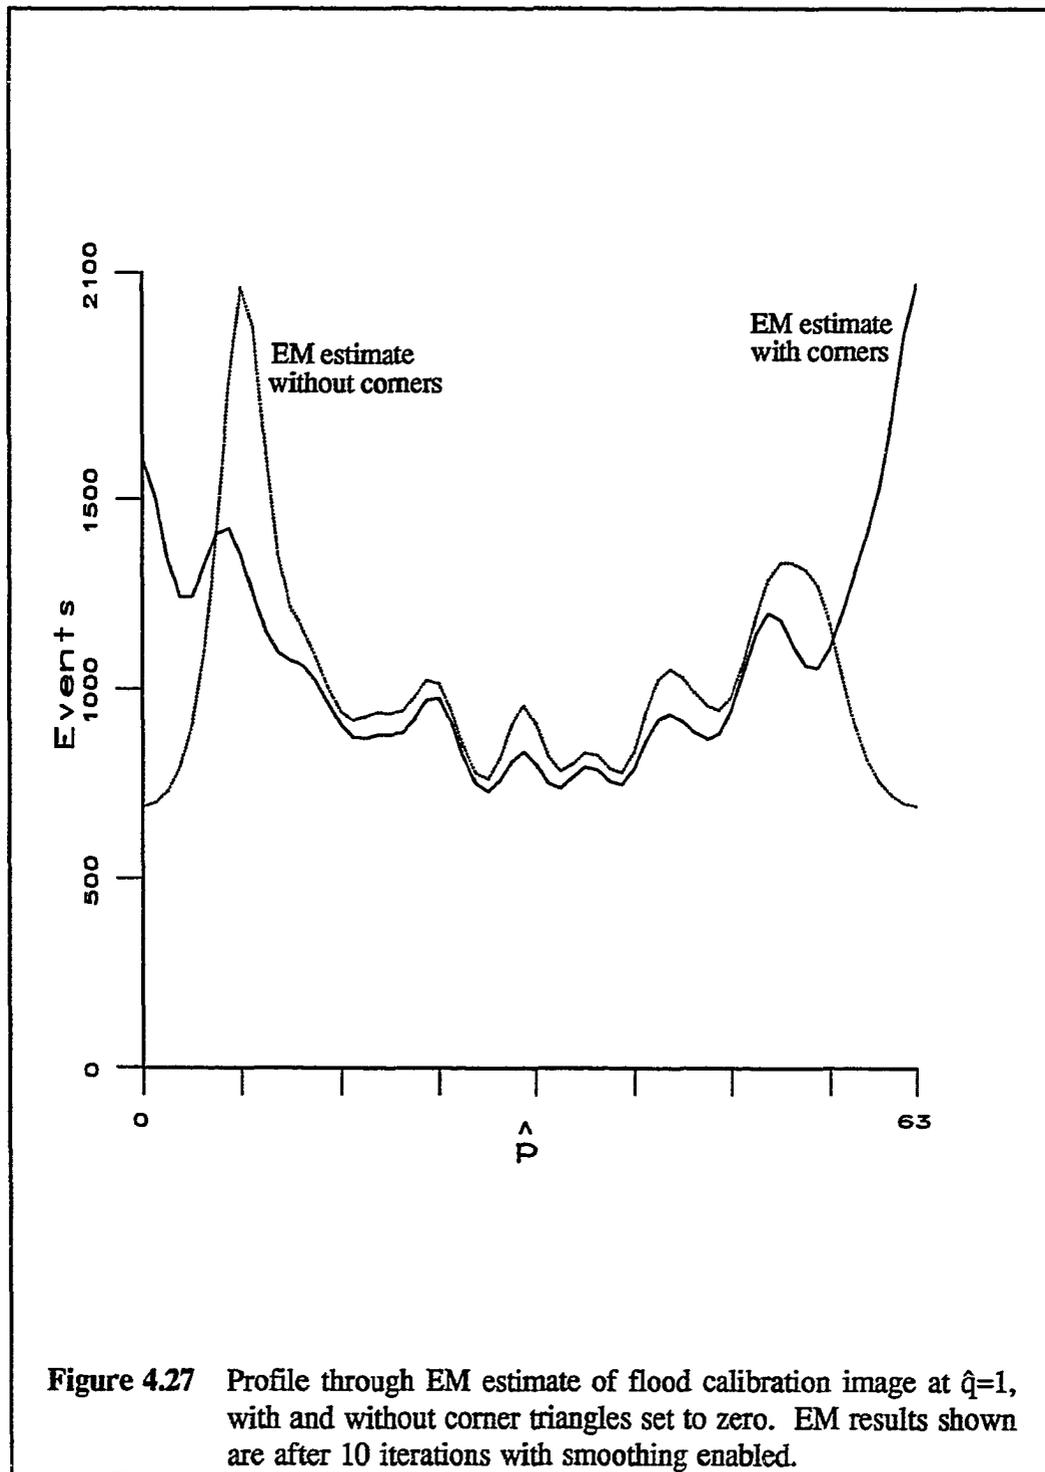


The standard deviation in the central 62 x 62 pixel region of the flood image for this camera was 175 counts, for a mean value of 580 events per pixel. The standard deviation predicted by the Poisson statistics is 24.1. The processed image standard deviation is 112 counts in the same 62 x 62 pixel region. If we consider only the central 56 x 56 pixel region, the standard deviation is reduced from 153 to 67 counts by EM processing.

The high intensity in the corner of the flood estimate is known to be in error, but the source of the error is at this time unclear. One way to deal with the error could be to redefine the projection data to exclude the high amplitude data from the corners of the measured image prior to processing. We excluded a triangular region, 8 pixels on each side, from the four corners of the crystal and applied the EM algorithm with smoothing enabled. The excluded area was 3.1% of the crystal area. A profile through $\hat{q}=1$ is shown in Figure 4.27 for the EM result with and without the corner pixels. Note how the rise in amplitude begins closer to the center when the corner pixels are excluded. From these results we decided to process the camera images with the corner regions included.

Projection Image Results

The goal of this work was to remove errors from the projection images that are used to reconstruct three-dimensional objects with the brain imager. The studies of the calibration images on individual modular cameras showed that EM processing does improve the estimate of the response in the central region of the crystal. Simulations with Gaussian blurs also showed that the EM algorithm could recover the original



flux patterns consisting of intermediate spatial frequencies if an accurate model of the camera response was used. Encouraged by these results, we next sought to evaluate how the EM algorithm would perform on projection images.

The projection images were recorded by the 24 cameras of the University of Arizona brain imager. The phantom imaged was a plastic bottle, 10 cm in diameter filled with water and 5 millicuries of ^{99m}Tc . Three plastic spheres extended from the top of the bottle, displacing the radioactive solution. Two of the spheres were 26 mm in diameter and the other was 33 mm in diameter. The measured projection images are shown in Figure 4.28 and Figure 4.29. The size and activity of the bottle are typical of objects imaged by the system.

We collected the image using the original Poisson LUTs that included uphill estimates. With the uphill estimates included, the four corner pixels can contain 22 times the number of events in the central pixels. These four pixels were set equal to zero prior to EM processing to reduce errors caused by the uphill estimates. We computed the EM estimate for 10 iterations. The processed images are shown in Figure 4.30 and Figure 4.31. The outermost pixels have been removed to show the detail in the lower amplitude regions of the camera. Note how the variations in the corners of the cameras are reduced. One deficiency that is apparent is the band of decreased counts near the edges of some cameras. The estimate is also higher immediately on the edge. The results demonstrate the need for a model which is a better representation of the camera response near the edges of the crystal.

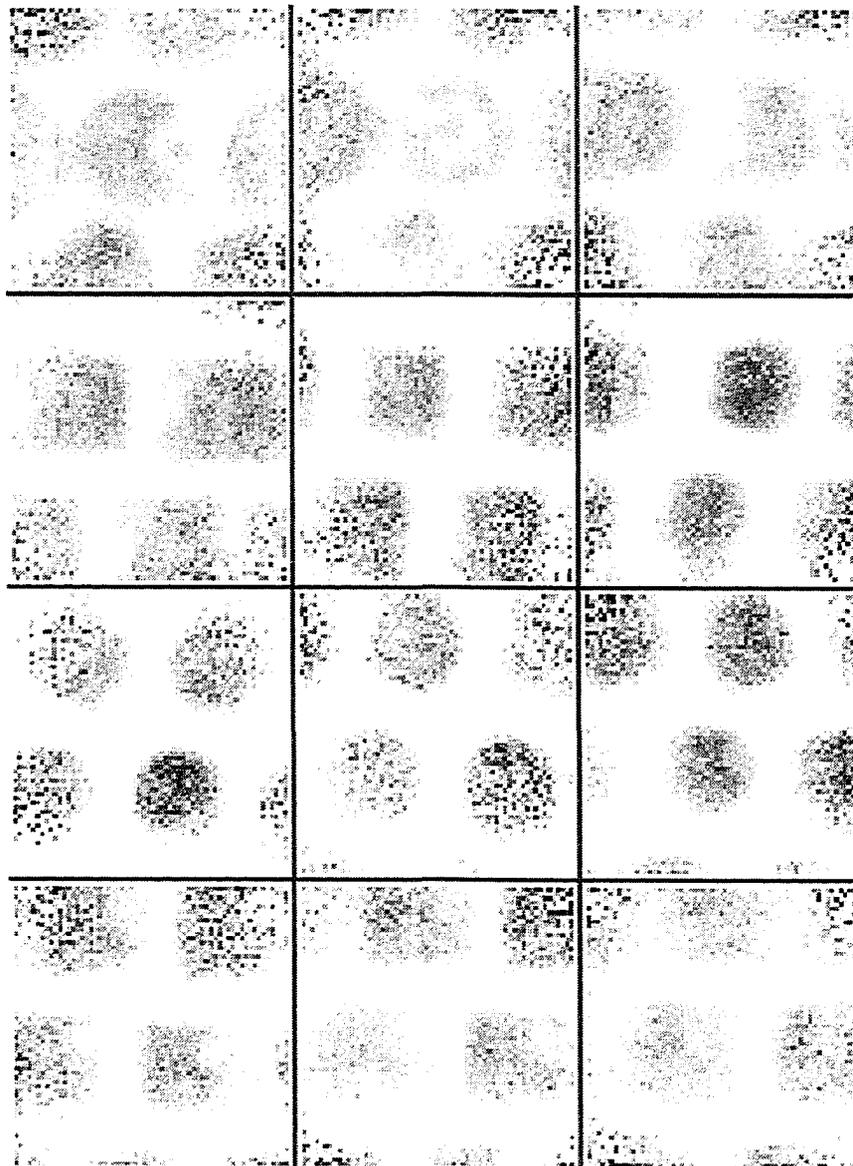


Figure 4.28 Projection images of the bottle phantom recorded by the brain imager for cameras 0-11.

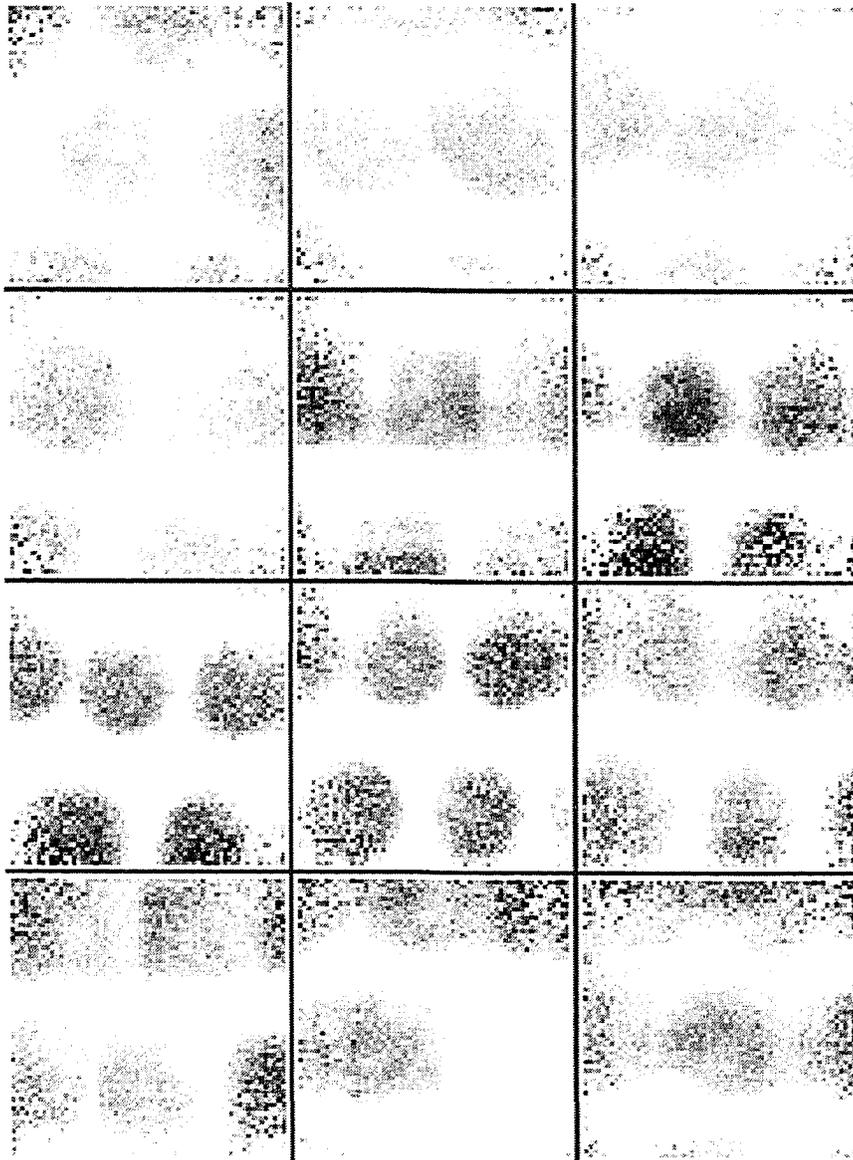


Figure 4.29 Projection images of bottle phantom recorded by the brain imager for cameras 12-23.

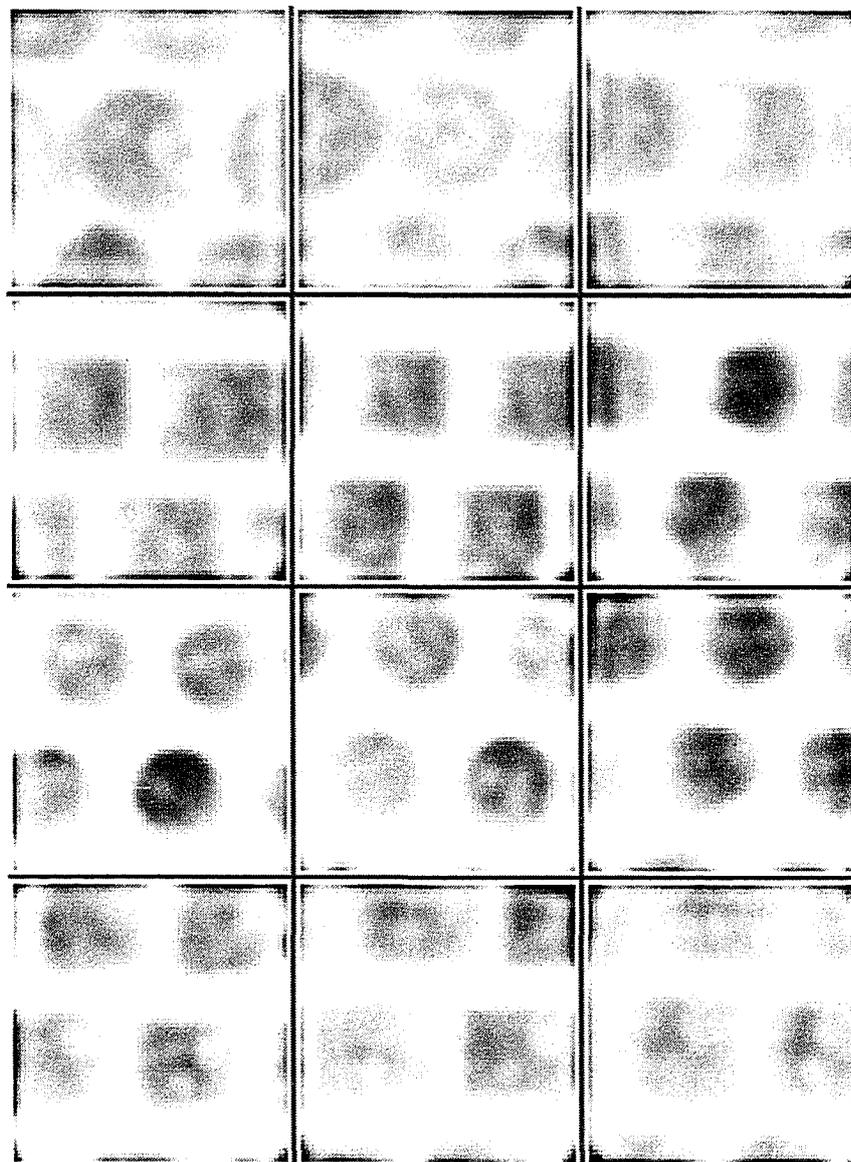


Figure 4.30 Projection images of the bottle phantom for cameras 0-11 after EM processing for 10 iterations with smoothing enabled.

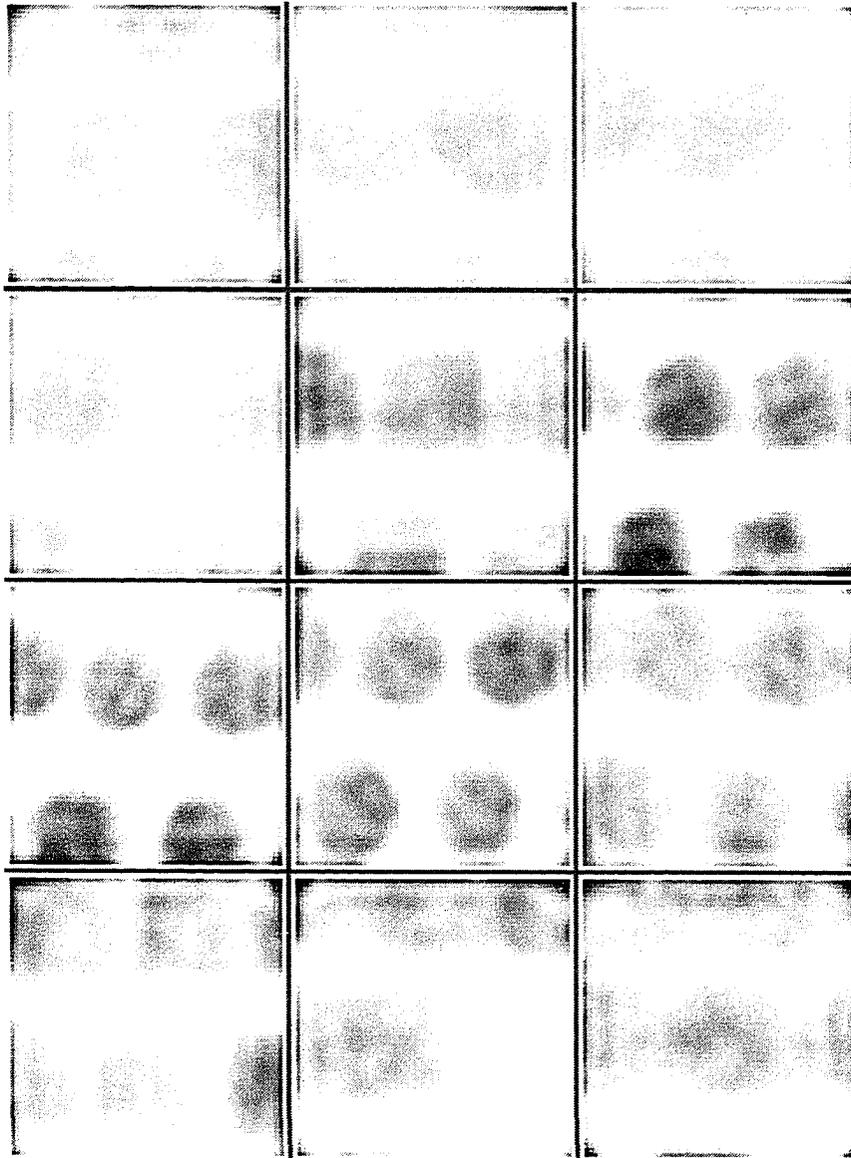


Figure 4.31 Projection images of the bottle phantom for cameras 12-23 after EM processing for 10 iterations with smoothing enabled.

5. THE MULTIVARIATE NORMAL MODEL FOR MODULAR CAMERA RESPONSE

EM processing modular camera images using the Poisson model recovered uniform flux estimates well in the central region of the crystal. However, the estimate of the response was amplified in the corners of the crystal, indicating that the model did not accurately represent the true camera response. Results from a previous study of the modular cameras (Hansen 1988) suggested that the probability distribution that characterizes the modular camera response does not have equal mean and variance. The distribution could therefore not be truly Poisson. The Fano factor (variance/mean) determined in these measurements was greater than 1 and varied with source position. Therefore, we sought a model that would more accurately represent the camera response.

Modelling the Cameras for EM Processing

To investigate the statistical performance of the modular cameras, we measured the response of the cameras to a collimated source as a four-dimensional vector. Recall that the Poisson model treats the PMTs as statistically independent and uses logarithmic matched filtering to estimate the mean from the individual PMT histograms. We sought to determine if the assumption that the PMT signals are statistically independent is valid by measuring the four-dimensional response directly.

In the four-dimensional measurement, the response of the four PMTs was recorded for a window of 20-bit ABCD values. The window included seven 5-bit channels for each PMT, or 7^4 ABCD values. The seven channels included the peak (the channel recording the most counts), and the 3 channels above and below the peak. The sample mean for each PMT and the 4×4 sample covariance matrix were computed from this "super-MDRF" measurement. The measurements were recorded for a 16×16 array of points on the same grid used for MDRF calibration.

The mean 5-bit value for the 16×16 array for each PMT is shown in Figure 5.1. The four blocks are arranged according to tube position, with tube C in the top left block. Within each block is the response for that tube as a function of crystal coordinate. Note that the response is highest (darkest) when the point source is placed directly over the PMT. The range of values plotted is from 2 to 29.

The covariances are plotted for each PMT pairing in Figure 5.2. The figure is arranged so that each smaller block is the covariance k_{ij} for the 16×16 array. The matrix is symmetric ($k_{ij} = k_{ji}$) so only the lower triangular elements are shown. In the upper left, k_{AA} is the variance for tube A. Note that the variance is higher in the corner nearest tube A, decreasing radially, and increasing again at the diagonal associated with tube D. The range of the variance values is from 0.45 to 1.68.

The variance (σ^2_{PMT}) of the measured data agreed with the variances reported by Hansen over most of the crystal area. Both experiments showed high variance in the corner closest to the PMT. The variance changed slowly when the source was

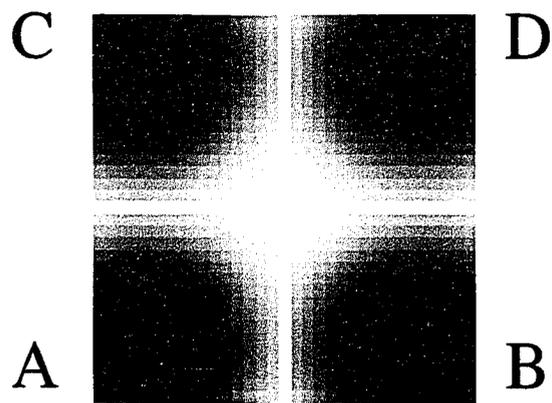


Figure 5.1 Mean 5-bit response computed from the "super-MDRF" measurement. The four large blocks represent, from top left to lower right, PMTs C, D, A and B.

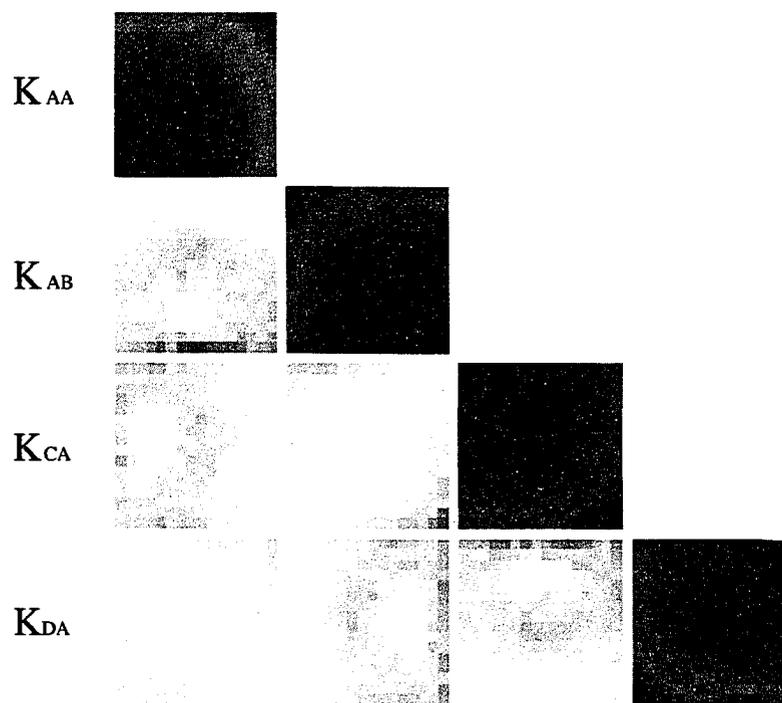


Figure 5.2. PMT covariance matrix computed from the 16 x 16 point "super-MDRF". Diagonal blocks are variances as a function of camera coordinate for PMTs A, B, C and D. Sub-diagonal blocks are the covariance values between PMTs.

directly over the PMT and decreased more rapidly as distance from the PMT increased. However, the rise in variance at the corner diagonal to the PMT was not exhibited in Hansen's experiments.

The covariances of adjacent PMTs show an elliptical pattern centered on the line between the two PMTs. The covariances of the diagonally oriented tubes (k_{AD}, k_{BC}) are highest along the edges and minimum at the center. The range of the covariance values is from -0.094 to 0.528. Since these values are not zero, and can be about 30% of the variance, the PMTs cannot be considered statistically independent.

The Poisson model cannot accommodate the statistical dependence of the PMTs, so we chose a multivariate normal distribution to model the camera response. The multivariate normal (MVN) distribution can be written

$$\Pr(\mathbf{x}|\bar{\mathbf{x}}) = \frac{1}{(2\pi)^{M/2} \sqrt{\det(\mathbf{K})}} \exp\left[-\frac{1}{2}(\mathbf{x} - \bar{\mathbf{x}})' \mathbf{K}^{-1} (\mathbf{x} - \bar{\mathbf{x}})\right] \quad (5.1)$$

where M is the dimension of the MVN random vector, in this case 4. The elements of the vector $(\mathbf{x} - \bar{\mathbf{x}})$ are the differences between the mean PMT signal level (\bar{x}_j) at the given coordinate, and the signal level being considered (x_j) for each PMT ($j=A, B, C$ or D). The matrix \mathbf{K} is the covariance matrix, determined experimentally from the super-MDRF measurement.

Measurement of the super MDRF for all 64 x 64 positions was not possible due to time, software, and memory limitations. The 16 x 16 point data set was

therefore interpolated to 64×64 points. As shown in Figure 5.2, the covariances exhibited a clear pattern, but the original covariance data were noisy. To produce a smooth data set suitable for interpolation, the four adjacent covariances were rotated to produce the same elliptical orientation. The four covariances were averaged and smoothed for each measured location. The two diagonal covariances were likewise averaged and smoothed. The variance of the four PMTs were also reoriented so they could be averaged. The averaged variance values were interpolated without prior smoothing. These three 16×16 images were used to generate values for the variance, adjacent covariance and diagonal covariance for each point in a 64×64 array by bilinear interpolation. The interpolated values were then reoriented to produce the 4×4 covariance matrix for each of the 4096 points.

Direct PMT Signals as the Measurement Index

As with the Poisson model, we can describe the measurement vector as events per ABCD value, or as events per estimated coordinate. Mapping from ABCD space to coordinate space through the LUT introduces an estimation step in the measurement process. Mapping to coordinate space has the advantage of allowing us to use our knowledge of the geometry of the camera to adjust the estimate. ABCD values with higher probability of mapping outside the crystal can be eliminated as described in Chapter 4. The distribution of ABCD values can be adjusted to improve uniformity using the flood correction technique described in Chapter 2. Considering these

advantages, we decided to focus the following studies on measurements in coordinate space. Nevertheless, computing the MVN camera response matrix $H(\hat{p}, \hat{q} | p, q)$ requires that we compute the probability $\Pr(ABCD | \bar{A}\bar{B}\bar{C}\bar{D}(p, q))$ in PMT space, so we describe the PMT space model first.

To compute the probabilities for the camera system matrix, we consider a window of ABCD values around the mean value determined by the super-MDRF for that source location. For the MVN model, we treat the ABCD value as a four-dimensional vector \mathbf{x} . The mean of the PMT signals is denoted with an overscore.

$$\begin{aligned} \Pr(ABCD | p, q) &= \sum_{ABCD} \Pr(ABCD | \bar{A}\bar{B}\bar{C}\bar{D}(p, q)) \\ &= \sum_{\mathbf{x}} \frac{1}{(2\pi)^2 \det(\mathbf{K}(p, q))} \cdot \\ &\quad \exp\left[-\frac{1}{2}(\mathbf{x}(p, q) - \bar{\mathbf{x}}(p, q))^t \mathbf{K}^{-1}(p, q) (\mathbf{x}(p, q) - \bar{\mathbf{x}}(p, q))\right] \end{aligned} \quad (5.2)$$

The joint probability $\Pr(ABCD | \bar{A}\bar{B}\bar{C}\bar{D}(p, q))$ is computed using the MVN probability density function from the mean values and covariance matrix \mathbf{K} resulting from the super-MDRF. The computed probability is stored as an element of \mathbf{H} if it exceeds a threshold. The threshold is used to eliminate very small probabilities from each column of \mathbf{H} . The sum of any column of \mathbf{H} will equal 1 if all probabilities are included. The threshold is chosen so that the sum of each column is close to 1 and the number of elements per column is less than 1000. We used a threshold of 0.0001

so that each column of \mathbf{H} contained a maximum of 513 elements and summed to a value between 0.978 and 0.990.

A limited range of ABCD values was considered when computing the $\Pr(\hat{p}, \hat{q} | p, q)$ matrix to reduce computation time. The window of ABCD values for each position was plus and minus 4, in four dimensions, from the mean ABCD value. Extending this range did not generate probabilities higher than the threshold.

Estimated Camera Coordinate as the Measurement Index

When the measurement vector is indexed as camera coordinates, each element is recorded as the number of events per area element on the camera face. The histogram of events per PMT signal is mapped to camera coordinates using an LUT that is based on the MVN model. The LUT contains the coordinate (\hat{p}, \hat{q}) which maximizes the probability function $\Pr(\text{ABCD} | \hat{p}, \hat{q})$.

The camera response matrix H_{mn} or $H(\hat{p}, \hat{q} | p, q)$, is the probability $\Pr(\hat{p}, \hat{q} | p, q)$, where (\hat{p}, \hat{q}) is the detector element and (p, q) is the coordinate where the gamma interaction took place. H_{mn} is computed from the LUT and the probability $\Pr(\text{ABCD} | \bar{A}\bar{B}\bar{C}\bar{D}(p, q))$ as

$$H_{mn} = \Pr(\hat{p}, \hat{q} | p, q) = \sum_{\text{ABCD}} \Pr(\hat{p}, \hat{q} | \text{ABCD}) \Pr(\text{ABCD} | \bar{A}\bar{B}\bar{C}\bar{D}(p, q))$$

To compute $\Pr(\hat{p}, \hat{q}|p, q)$, we consider the conditional probability that a detector estimate m at (\hat{p}, \hat{q}) will result from a given PMT signal, ABCD, and the probability that the same PMT signal will result from a source element n located at (p, q) .

The probability $\Pr(\hat{p}, \hat{q}|ABCD)$ is determined by a MVN model LUT. This LUT is generated by finding the flux element (p, q) that maximizes the probability $\Pr(ABCD|p, q)$ for each ABCD. To prevent low-probability events (corresponding to scattered photons of low energy) from contributing to the image, the maximum probability $\Pr(ABCD|p, q)$ must exceed a threshold value. The threshold is set to be some percentage of the maximum (when the exponential term equals 1). For the sum of the elements in a column of H to be above 0.90, this threshold must be 0.05 or lower. This threshold differs from the threshold used in the Poisson model, which is one value for all coordinates. We chose to use the variable threshold with the MVN model because the width and maximum value of the distribution varies as a function of camera coordinate. A fixed threshold would exclude a greater percentage of ABCD values when the maximum probability $p(ABCD|p, q)$ for a particular element was lower than average.

When the mean PMT signal values were calculated from the super-MDRF and mapped to the nearest integer value, more than one detector was associated with each mean value. The detector which maximized the probability at a given mean ABCD value would be determined only by the determinant of $K(p, q)$. This lack of unique mean values resulted in a LUT which favored one pixel in the group of pixels

associated with that mean ABCD. To alleviate this problem, we allowed the mean values to be floating point so that they would be smoothly varying.

We intend to use the MVN model on all cameras, but we measured the super-MDRF for one camera (number 21) only. Since the covariance matrix is symmetrical and slowly varying, we assume that the measured \mathbf{K} closely approximates \mathbf{K} for any camera. We also know from past calibration of the cameras that the mean detector response function is unique for each camera. Since we did not have super-MDRF measurements for all cameras, we used the abar file for each camera to generate the mean PMT values as a function of (p,q) for each camera. We computed $\bar{x}_j(p,q)$ by dividing the 10-bit abar value by 4, mapping the 8-bit result through the square-root compression algorithm, and storing the floating-point result for each PMT.

Comparison between this operation on the abar data and super-MDRF calculation for camera 21 indicated that the root-mean-square error between the resulting values for $\bar{x}_j(p,q)$ was 0.39. Since the the abar data and super-MDRF mean values showed good agreement, the probabilities for \mathbf{H} were computed using the abar data, to facilitate use of the same software with all cameras in the future.

After the maximum probability LUT is computed, the program that generates the LUT checks the edges to find ABCD values that have higher probability of being mapped outside of the camera space. These "uphill" estimates are discarded from the LUT. Then we apply the flood correction algorithm described in Chapter 2 to the LUT to promote uniformity in the image.

Each ABCD value mapped by the LUT is associated with only one detector or to no detector. The probability $\Pr(\hat{p}, \hat{q} | ABCD) = 1$ for the mapped detector coordinate and zero for all other coordinates. Similarly, if ABCD maps to no coordinate, $\Pr(\hat{p}, \hat{q} | ABCD) = 0$ for all coordinates. Consider a column of $H(ABCD | p, q)$. Each ABCD in the column maps to some (\hat{p}, \hat{q}) , and the probability $\Pr(ABCD | p, q)$ is added to the corresponding element of $H(\hat{p}, \hat{q} | p, q)$. The elements of $H(ABCD | p, q)$ that map to no (\hat{p}, \hat{q}) are discarded. Since many of the elements of the PMT-based matrix are discarded or below the LUT threshold, the sum of each column is lower for $H(\hat{p}, \hat{q} | p, q)$ than for $H(ABCD | p, q)$. Due to removal of uphill estimates, the sum of a column associated with an edge coordinate is lower than the sum of a column for a central coordinate.

Simulation Studies

We simulated the camera response with the multivariate normal model to initially evaluate if the model produced plausible results. We used the procedure described in Chapter 4 to generate simulated response images. The ideal flux vectors included the flood, point array and blur images, as before. \mathbf{H} was computed using the MVN model from above. Poisson noise was added to the product $\mathbf{H}\mathbf{f}$ to represent the fluctuations in the number of detected photons. The energy functions $E(\mathbf{f})$ and $E(\mathbf{g})$ were computed for each case and compared with the energy functions from the Poisson model simulations. The initial energy, $E(\mathbf{f}^0)$, when the estimate is a constant,

is much larger than the subsequent energy values. The energy plots therefore begin with the energy values after the first iteration, $E(f^1)$ and $E(g^1)$ so that the shape of the energy function curve is apparent as the energy becomes asymptotic. The simulated response images are compared with the measured response that results from mapping the calibration histograms by the MVN model LUT.

Point Array Results

The MVN model produced point array images that are similar in appearance to the images produced by the Poisson model. Figure 5.3 shows the simulated and measured point array images when the MVN model LUT is used. The response in the corners is spread out among many detectors and shows high-frequency fluctuations. The response near the edges is elliptical. The best response is near the center of the crystal.

EM processing the point array demonstrated some differences between the Poisson and MV models. The measurement energy, $E(g)$, is shown in Figure 5.4 for the multivariate normal (MV) and Poisson (P) model simulations of the point array. Both unregularized and smooth results are shown. The energy values are very high for the point array case because the difference between the ideal and the measured images is large everywhere, especially at the peaks in the ideal image. The events for each point in the array are distributed among the surrounding detectors. The unregularized $E(g)$ values drop from initial energies of order 10^7 to about 30000 by the

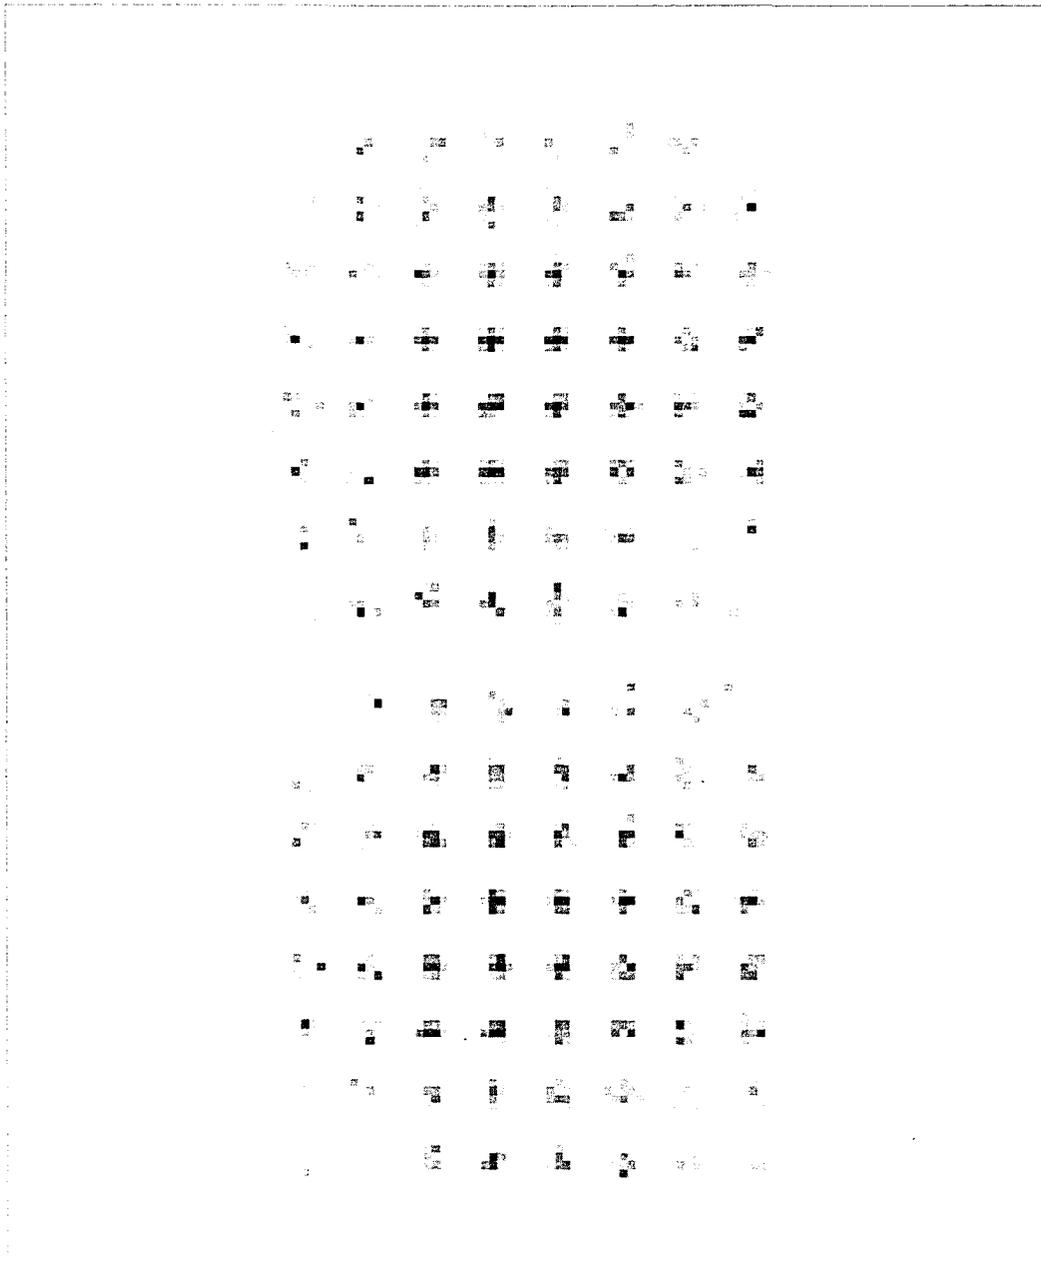
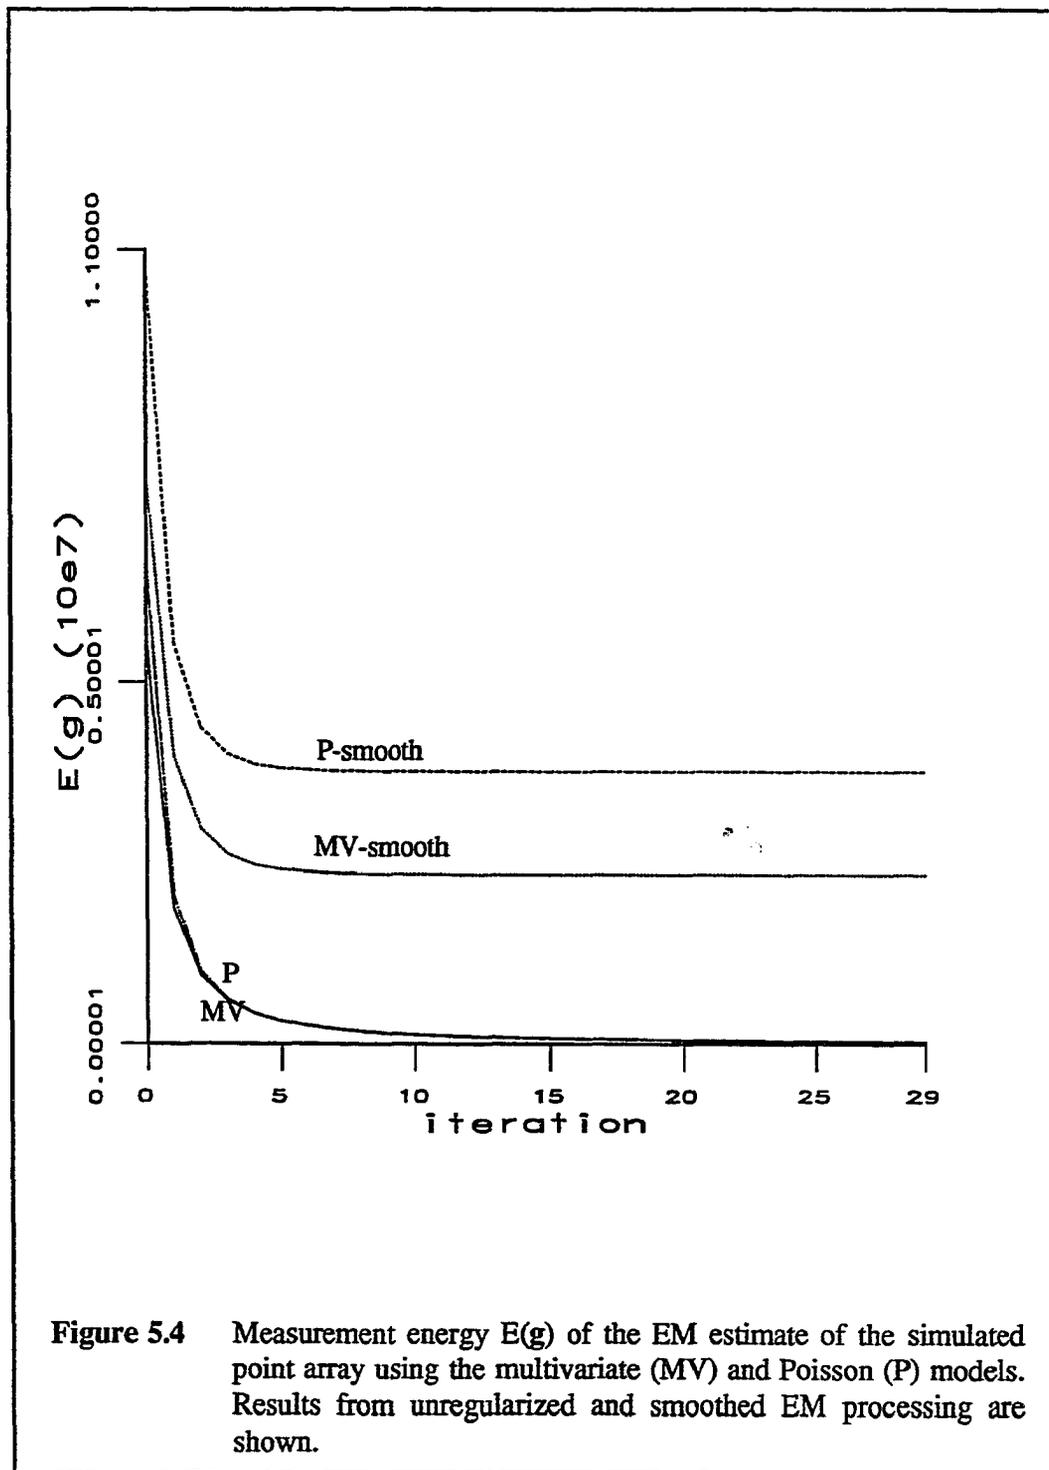


Figure 5.3 Top: Simulated response to an ideal point array using the MV model.
Bottom: MEasured response to a point array calibration image using the MV model.

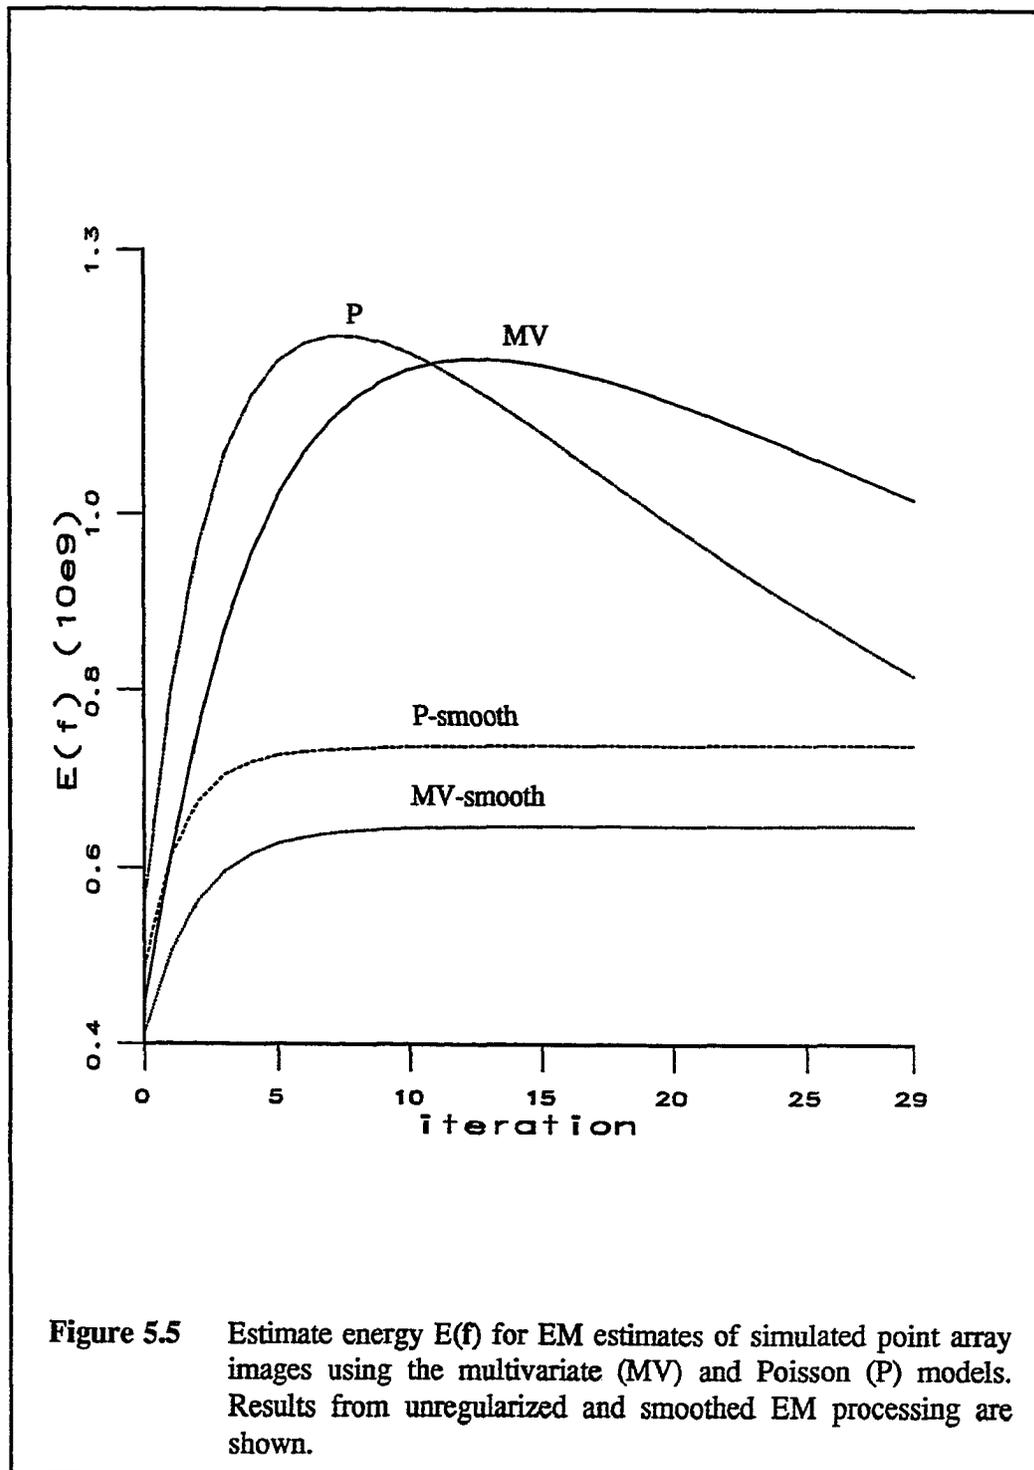


thirtieth iteration. In the plot, both models appear to have the same $E(\mathbf{g})$. Actually the Poisson model $E(\mathbf{g})$ value is 80.8% of the MV $E(\mathbf{g})$. With smoothing enabled the energy for both models changes little after the 16th iteration. The smoothed EM algorithm produced values of $E(\mathbf{g})$ for the MV model that are 60% of $E(\mathbf{g})$ for the Poisson model. The estimate energy, $E(\mathbf{f})$, is shown in Figure 5.5. $E(\mathbf{f})$ is minimum (between 4.1×10^8 and 5.6×10^8) after the first iteration for all cases. The unregularized Poisson model $E(\mathbf{f})$ has the most negative slope and is lower than $E(\mathbf{f})$ for either MV case at higher iterations. Of the smoothed results, the MV model produces the lower energy.

A profile at $\hat{q}=2$ of the simulated point array processed with the unregularized EM algorithm after only 10 passes is shown in Figure 5.7. The ideal and simulated response profiles are also shown for comparison. The EM result is a better estimate than the simulated response even at this early iteration. The point array processed with the smoothed EM algorithm is shown in Figure 5.6 for pass 10 at $\hat{q}=2$ for comparison. As with the smoothed EM result using the Poisson model, the widths of the peaks have changed little, but the peaks are more symmetrical and uniform in amplitude.

Flood Image Results

Figure 5.8 shows a profile of the simulated and measured flood images when the MV model H and LUT are used. The most notable difference between these



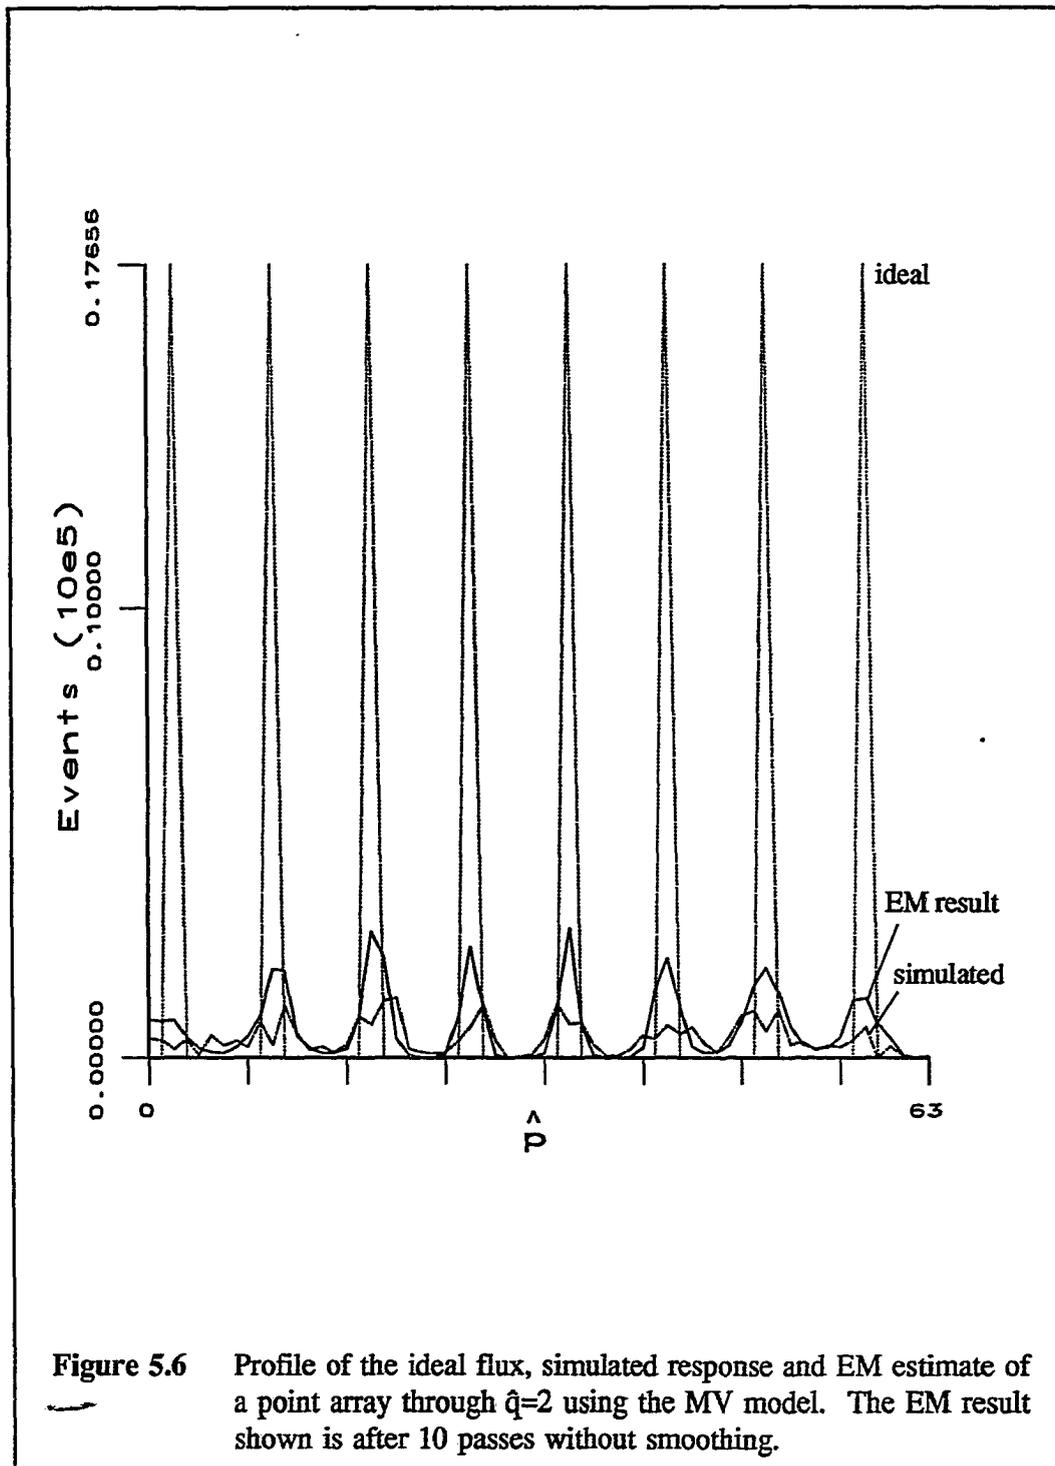
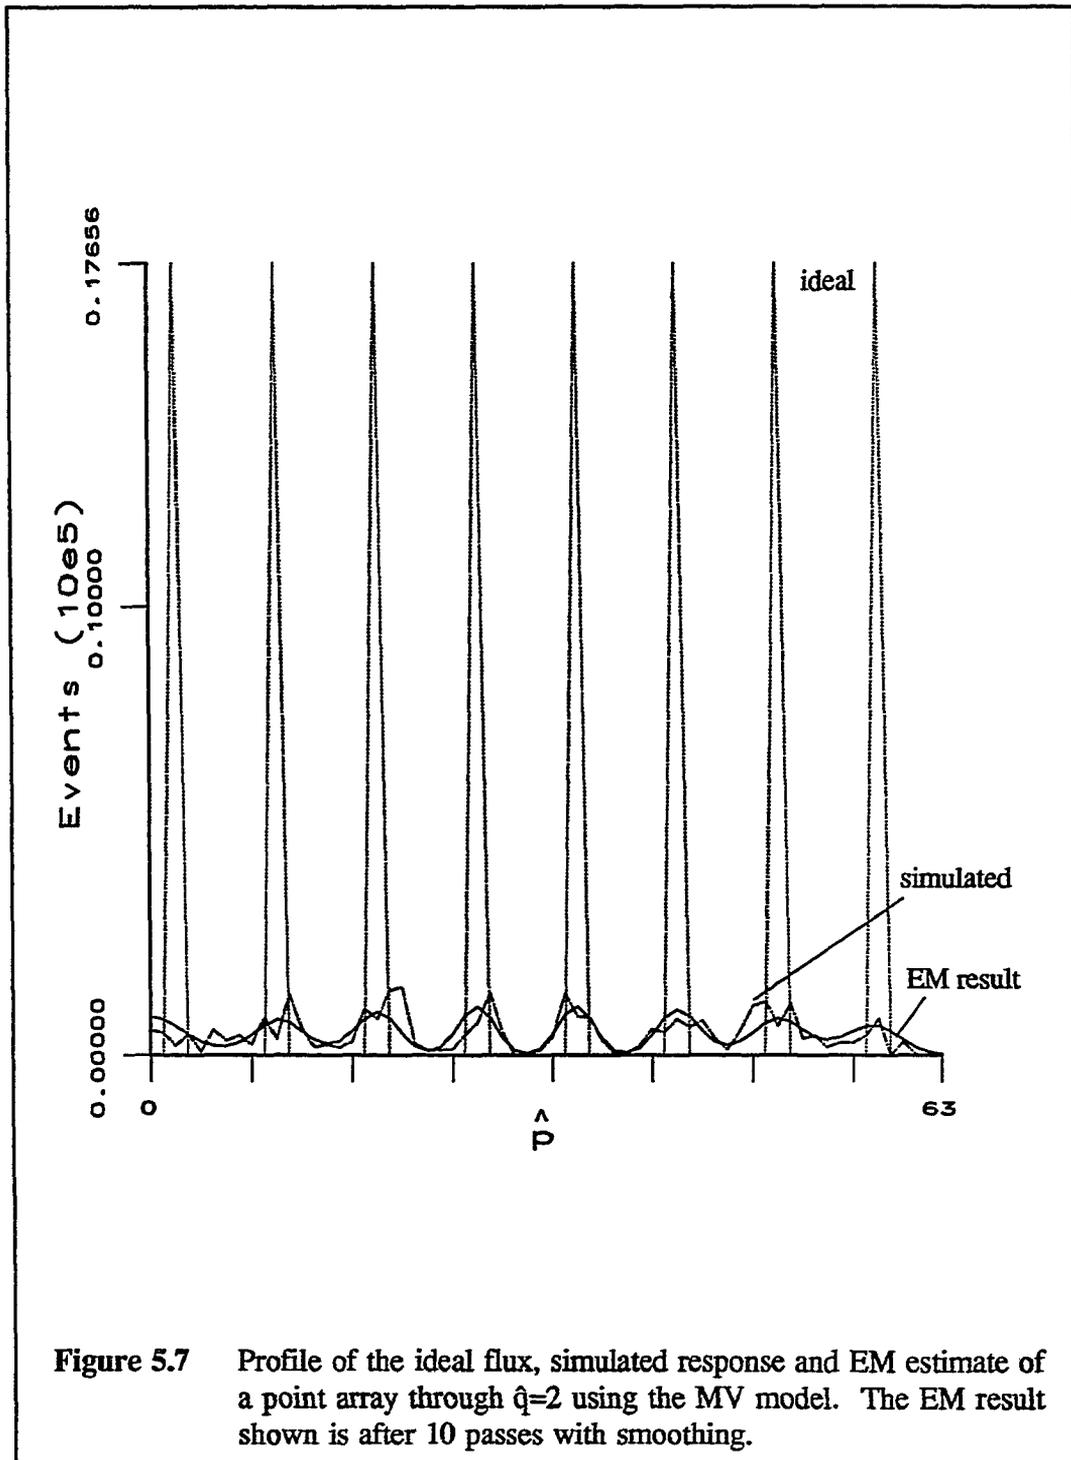


Figure 5.6 Profile of the ideal flux, simulated response and EM estimate of a point array through $\hat{q}=2$ using the MV model. The EM result shown is after 10 passes without smoothing.



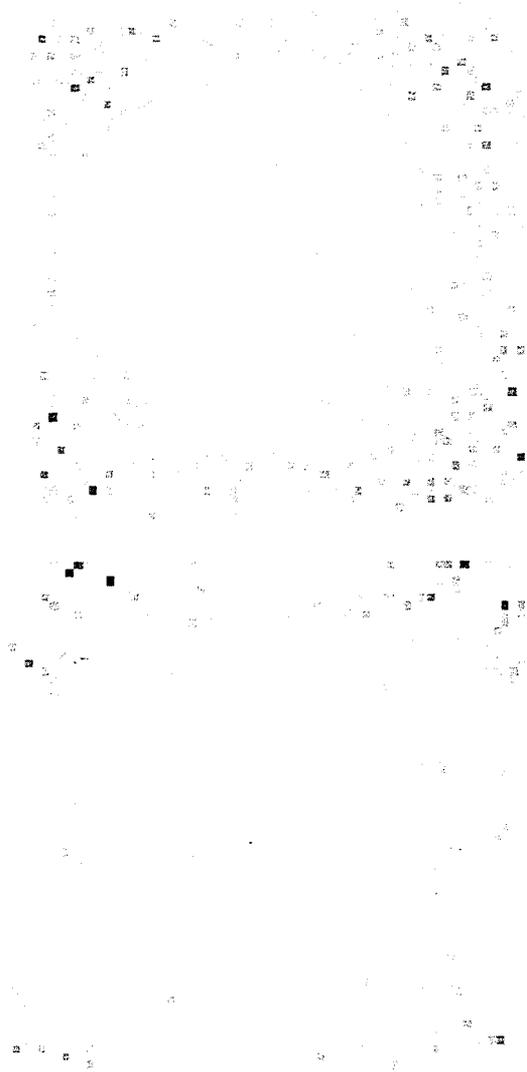


Figure 5.8 Top: Simulated response to an ideal flood image using the MV model.
Bottom: Measured response to the flood calibration using the the MV model LUT.

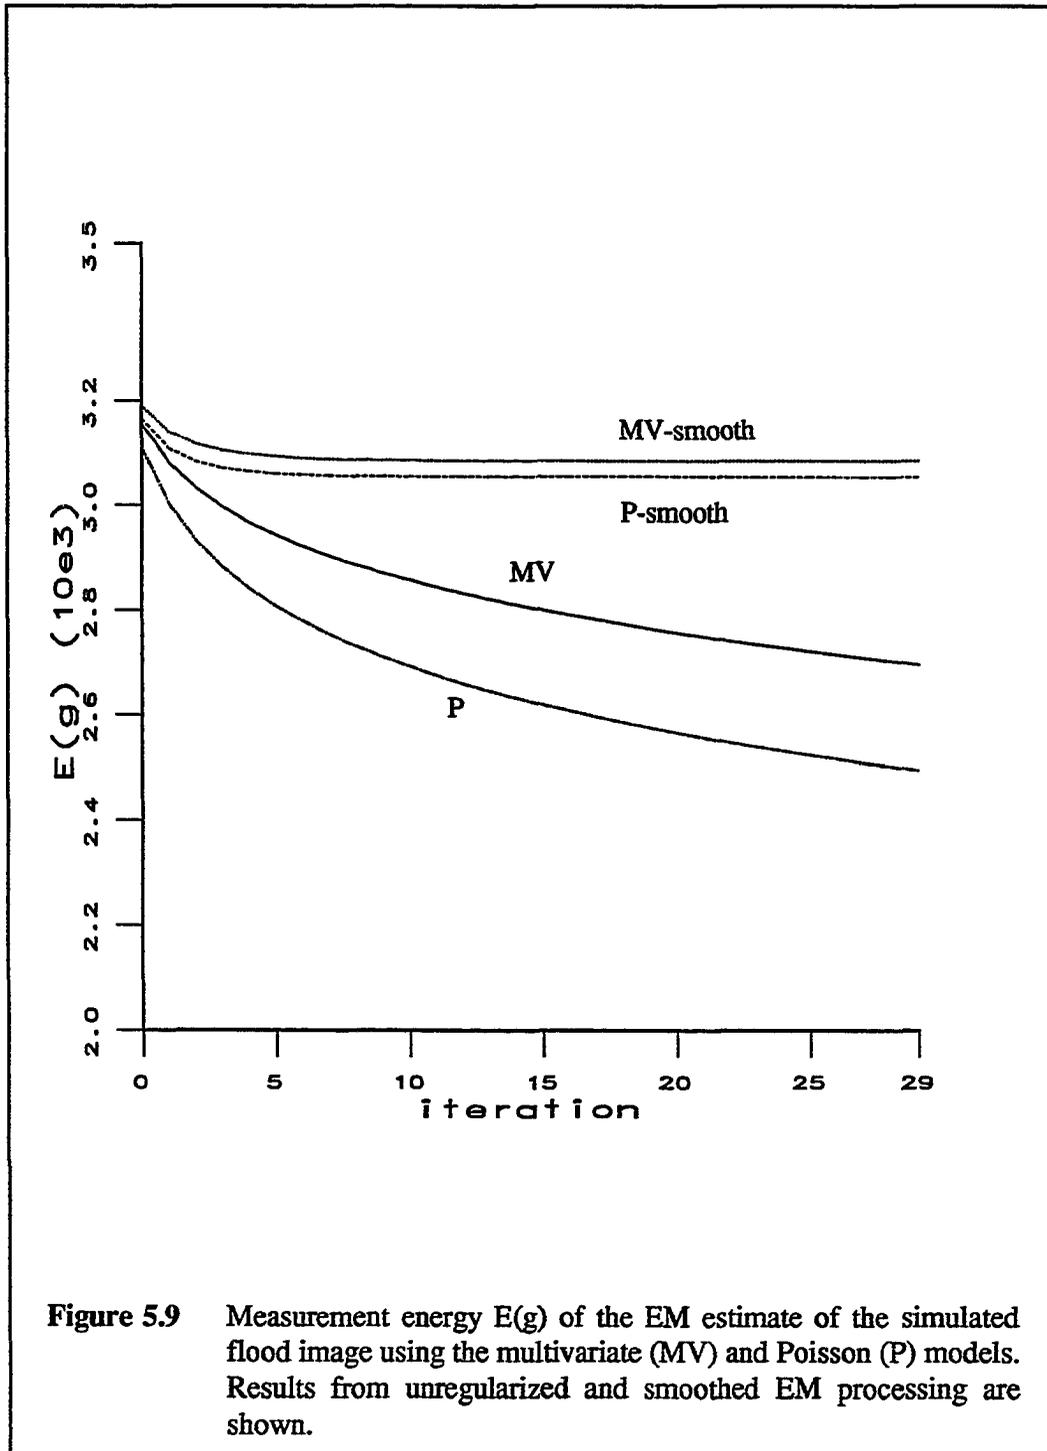
images is that the response along the edges of the crystal in the simulated image is typically lower than the measured response mapped by the LUT.

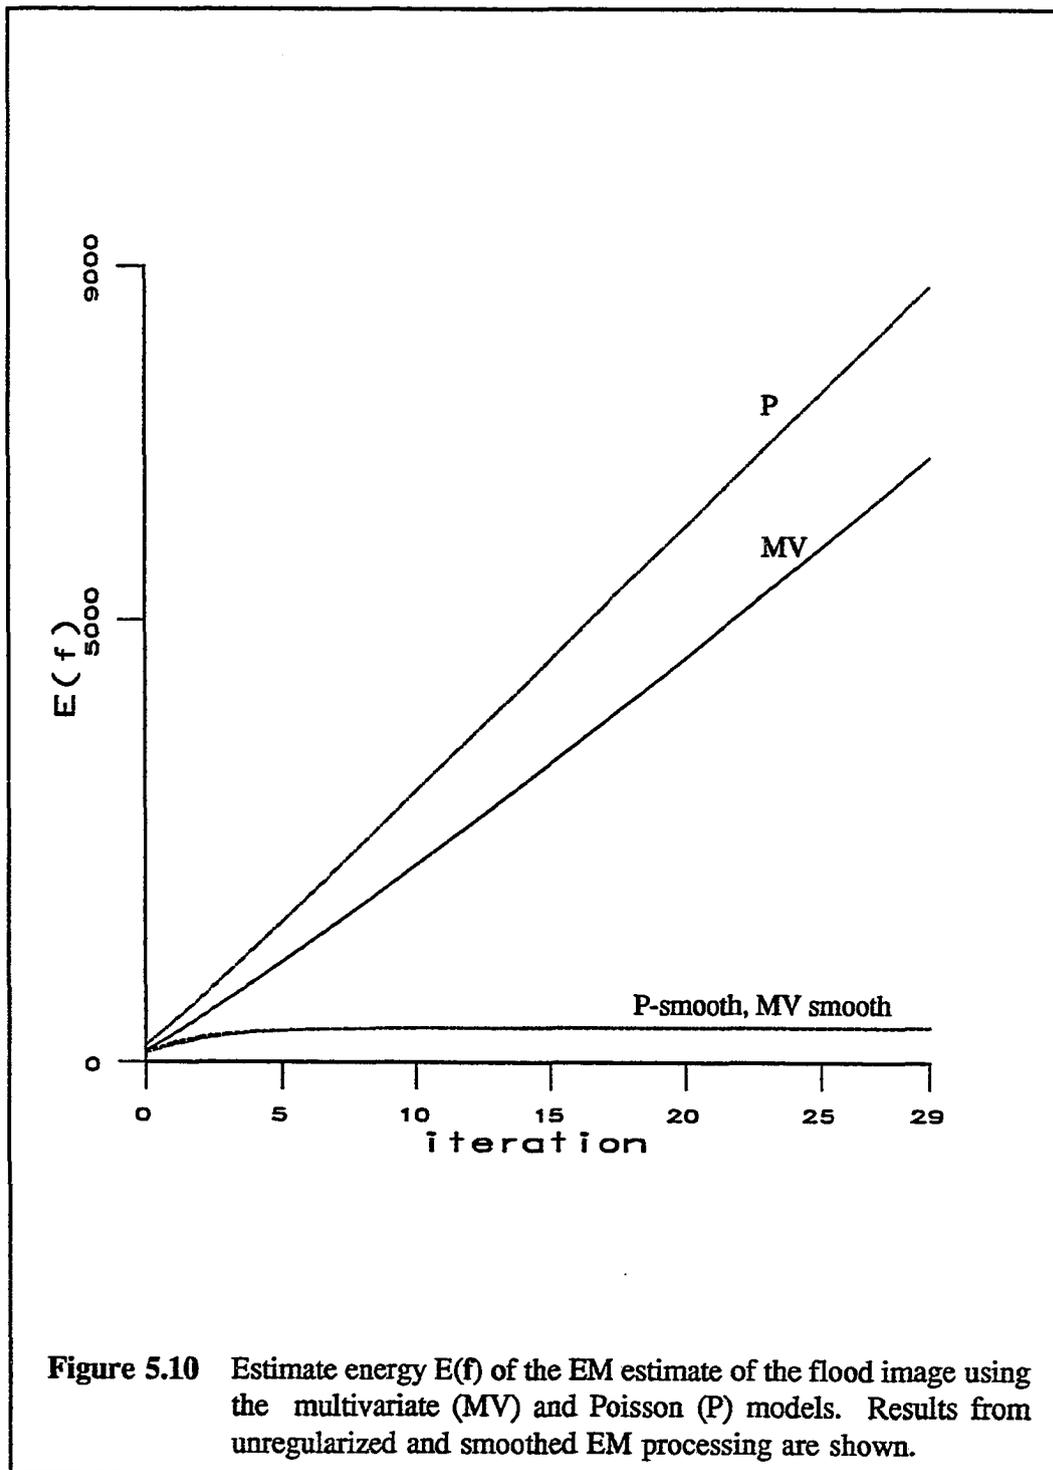
The measurement energy is shown in Figure 5.9 for the multivariate normal (MV) and Poisson (P) model simulations of the flood response. Both unregularized and smooth results are shown. The energy values are lower than for the point array because the initial estimate is very close to the ideal image. The unregularized $E(g)$ for the MV model is higher than $E(g)$ for the Poisson model. With smoothing, the difference in $E(g)$ between the two models is 0.98%. The estimate energy, shown in Figure 5.8, is similar for the two models. The unregularized energy is minimum after the first iteration and increases with iteration number for both models. The MVN model energy increases less rapidly. The smoothed $E(f)$ values are also very similar, differing by 2%.

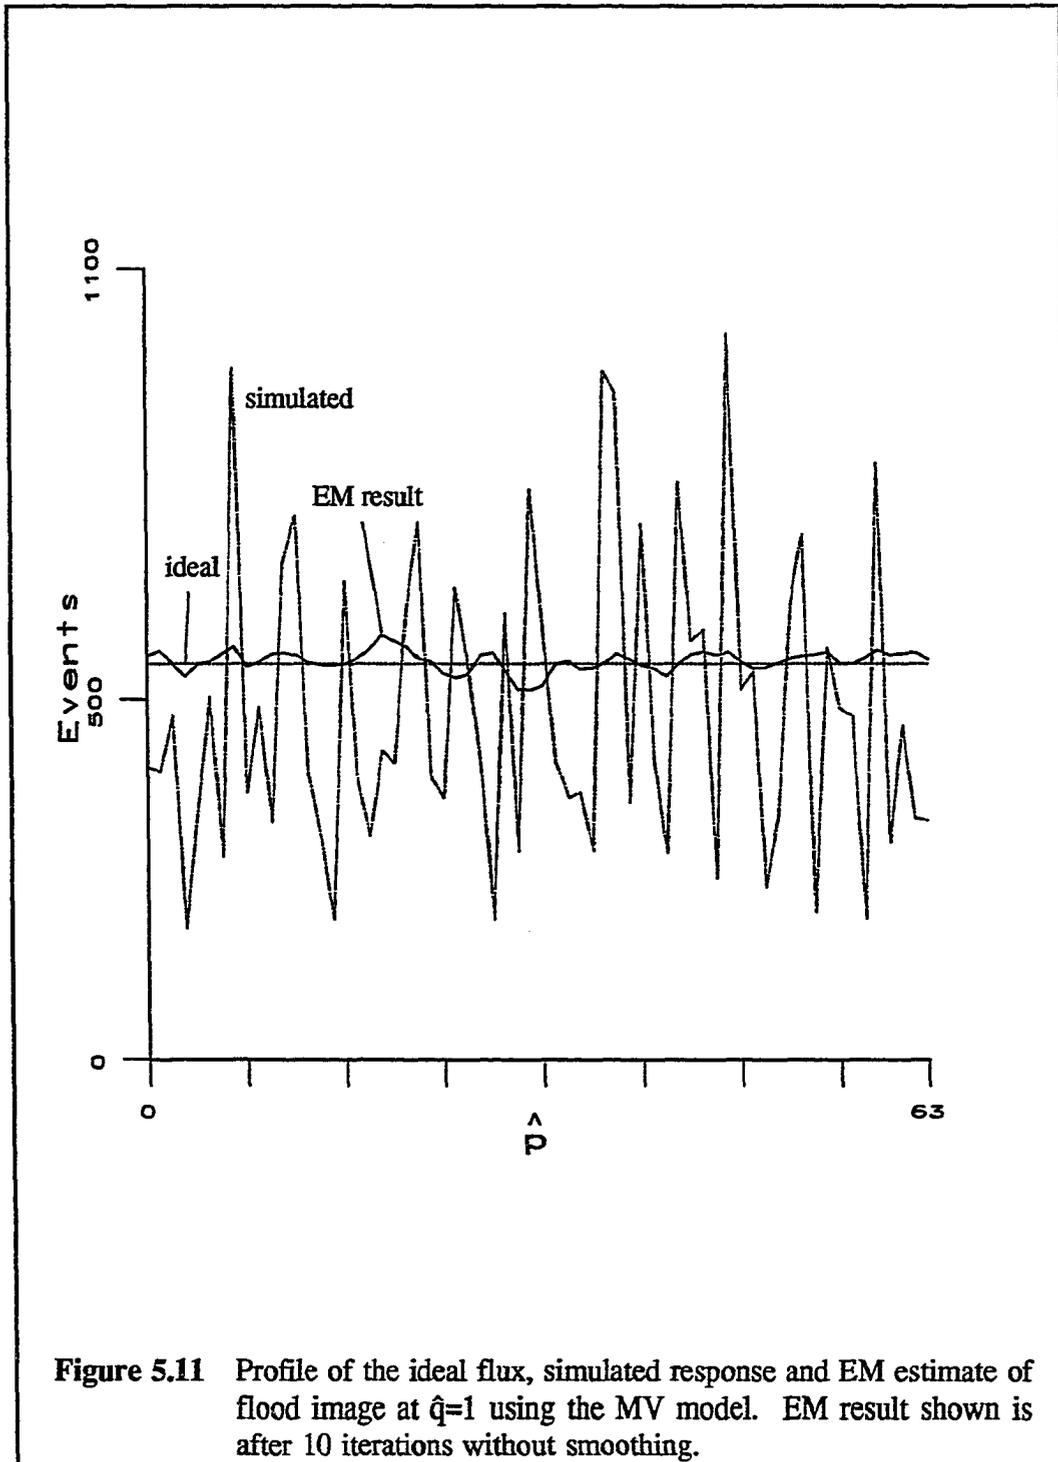
A profile through $\hat{q}=2$ of the simulated flood image, processed with the unregularized EM algorithm after 10 passes, is shown in Figure 5.11. The ideal flux and simulated response profiles are also shown for comparison. The standard deviation in the simulated response was 152.5, and decreased to 7.3 in the processed image. The smoothed EM result for the same profile is shown in Figure 5.12.

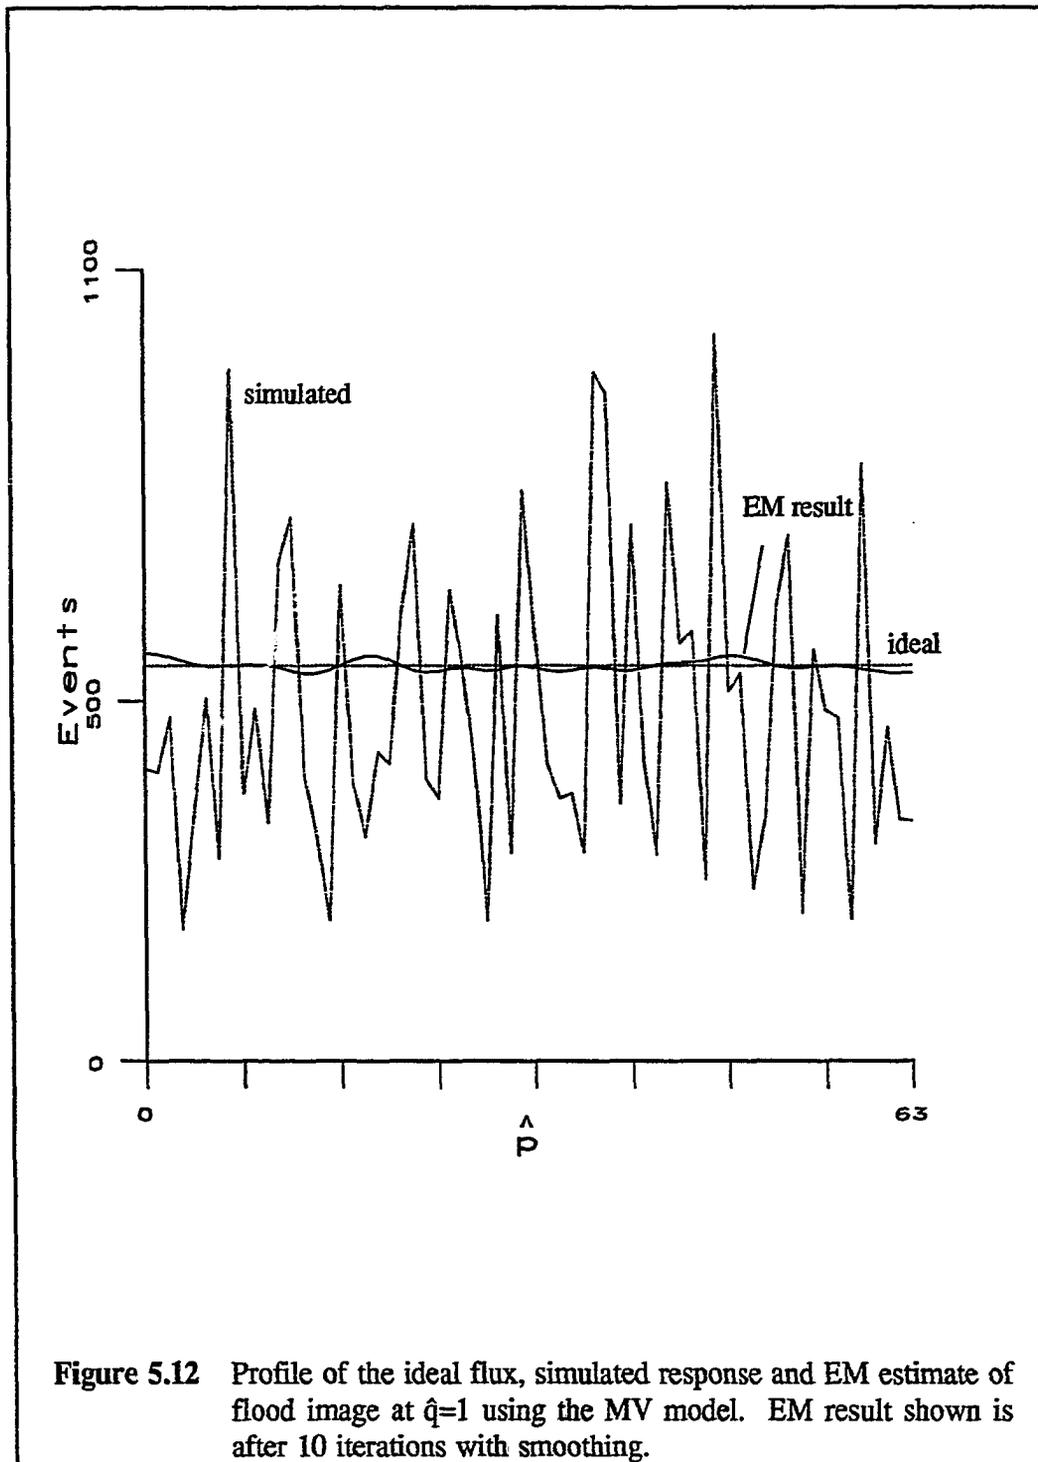
Blur Image Results

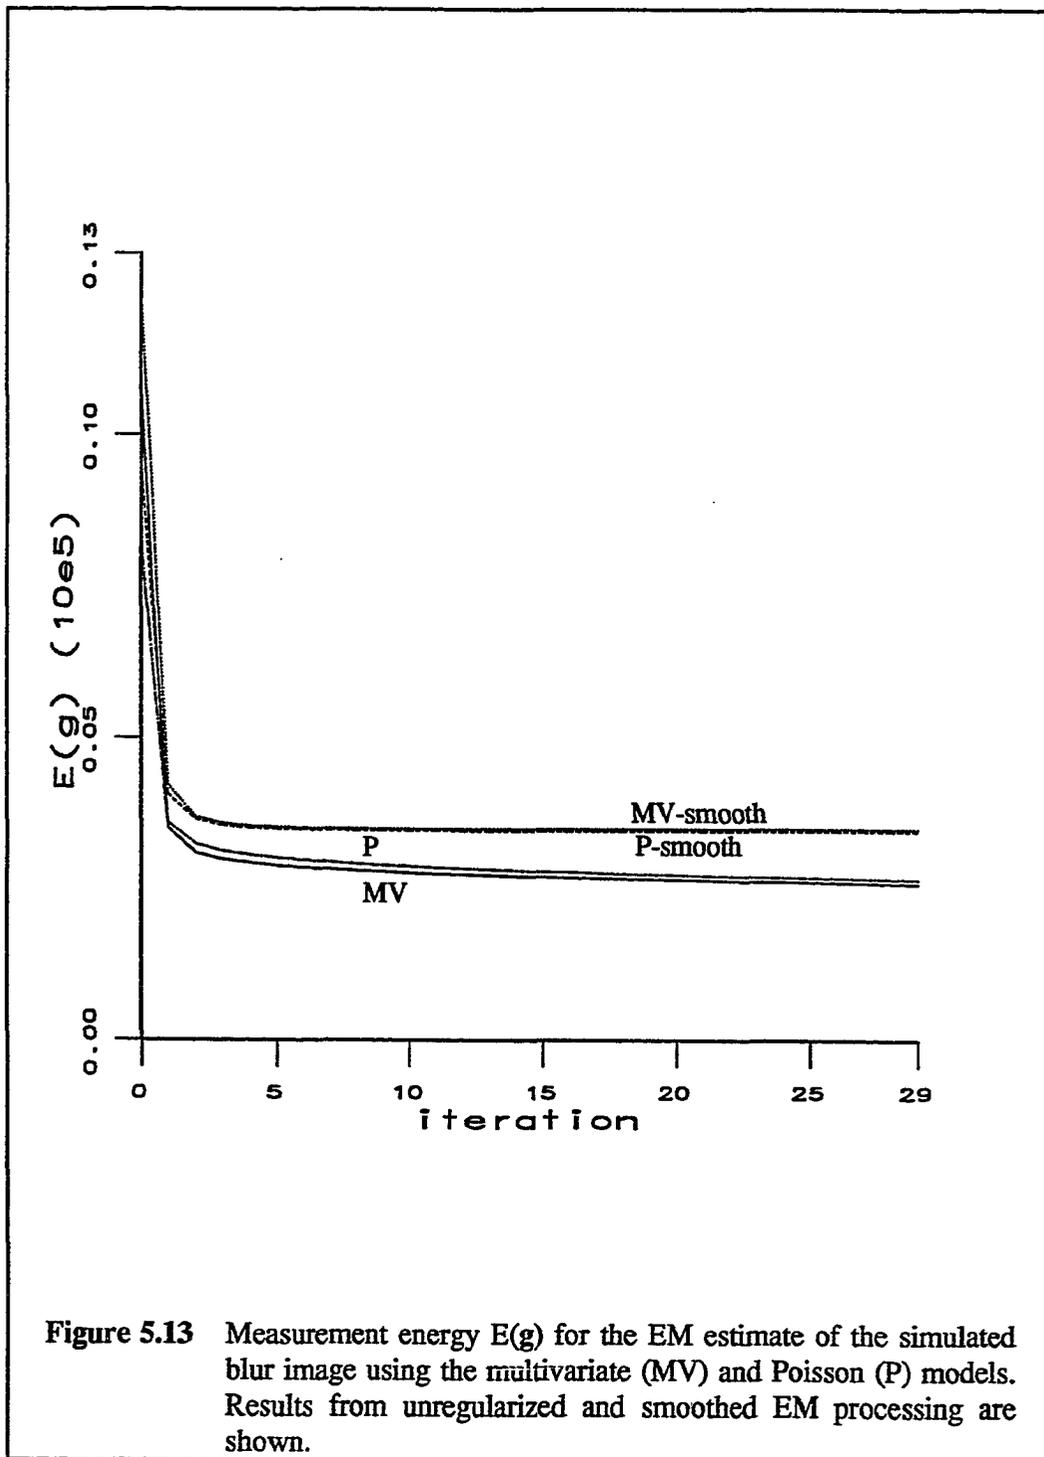
The measurement energy functions for the blur simulations are shown in Figure 5.13. The asymptotic value for $E(g)$ differs by 0.8% between the models in the











smoothed EM results, and by 15% between the models in the unregularized results. The estimate energy, $E(\mathbf{f})$, is plotted for both models in Figure 5.14. Again, the results for the two models are similar. The value of $E(\mathbf{f})$ for the smoothed result is lower for the Poisson model than for the MV model by 21.2%. Using both models, the unregularized $E(\mathbf{f})$ increases with iteration number after the minimum at iteration 4. The profile through the EM processed simulated blur at $\hat{q}=1$ is shown for the unregularized and smoothed algorithm in Figure 5.15 and Figure 5.16 respectively.

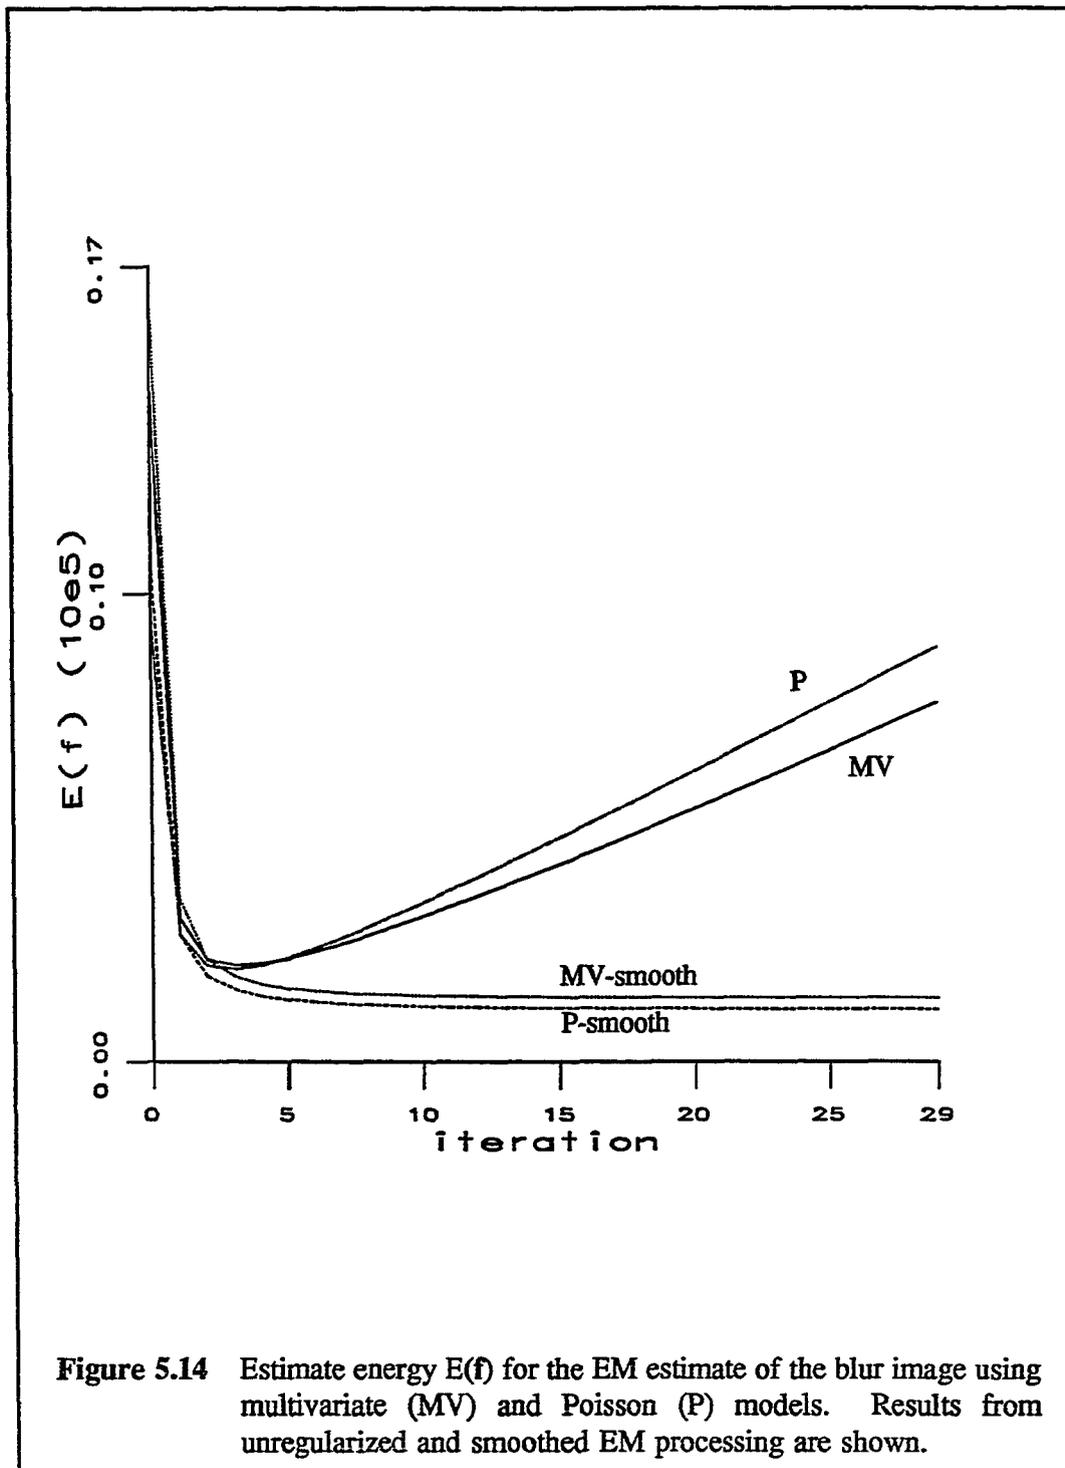
Summary

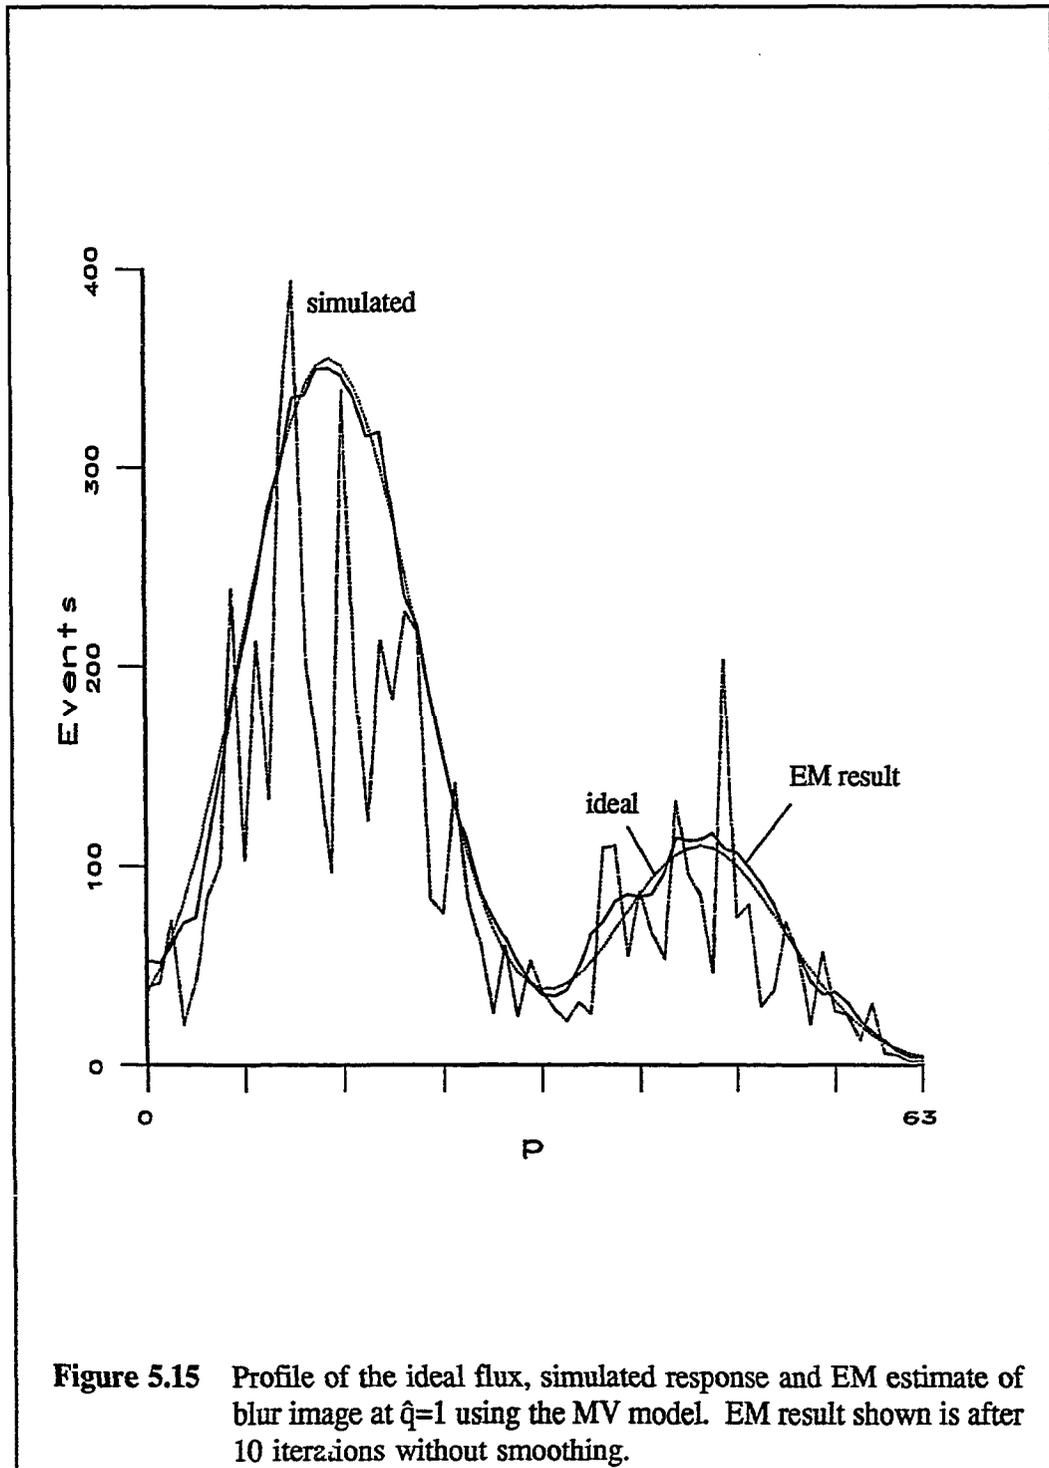
The EM results with above simulated images show that the MVN model for the modular camera response produces results that are similar to the results from the Poisson model. Close approximations to the ideal flux vectors can be obtained from simulated response images using the EM algorithm and the MVN model.

Calibration Image Results

The measured point array and flood histograms were mapped to (p,q) coordinates by the MVN model LUT to produce the measured images previously shown in Figure 5.3 and Figure 5.8. These measured images were processed using the EM algorithm to determine if we could produce a better estimate of these vectors.

The results from processing the MVN mapped point array are shown in Figure 5.17 (unregularized) and Figure 5.18 (smoothed). Both profiles have lower





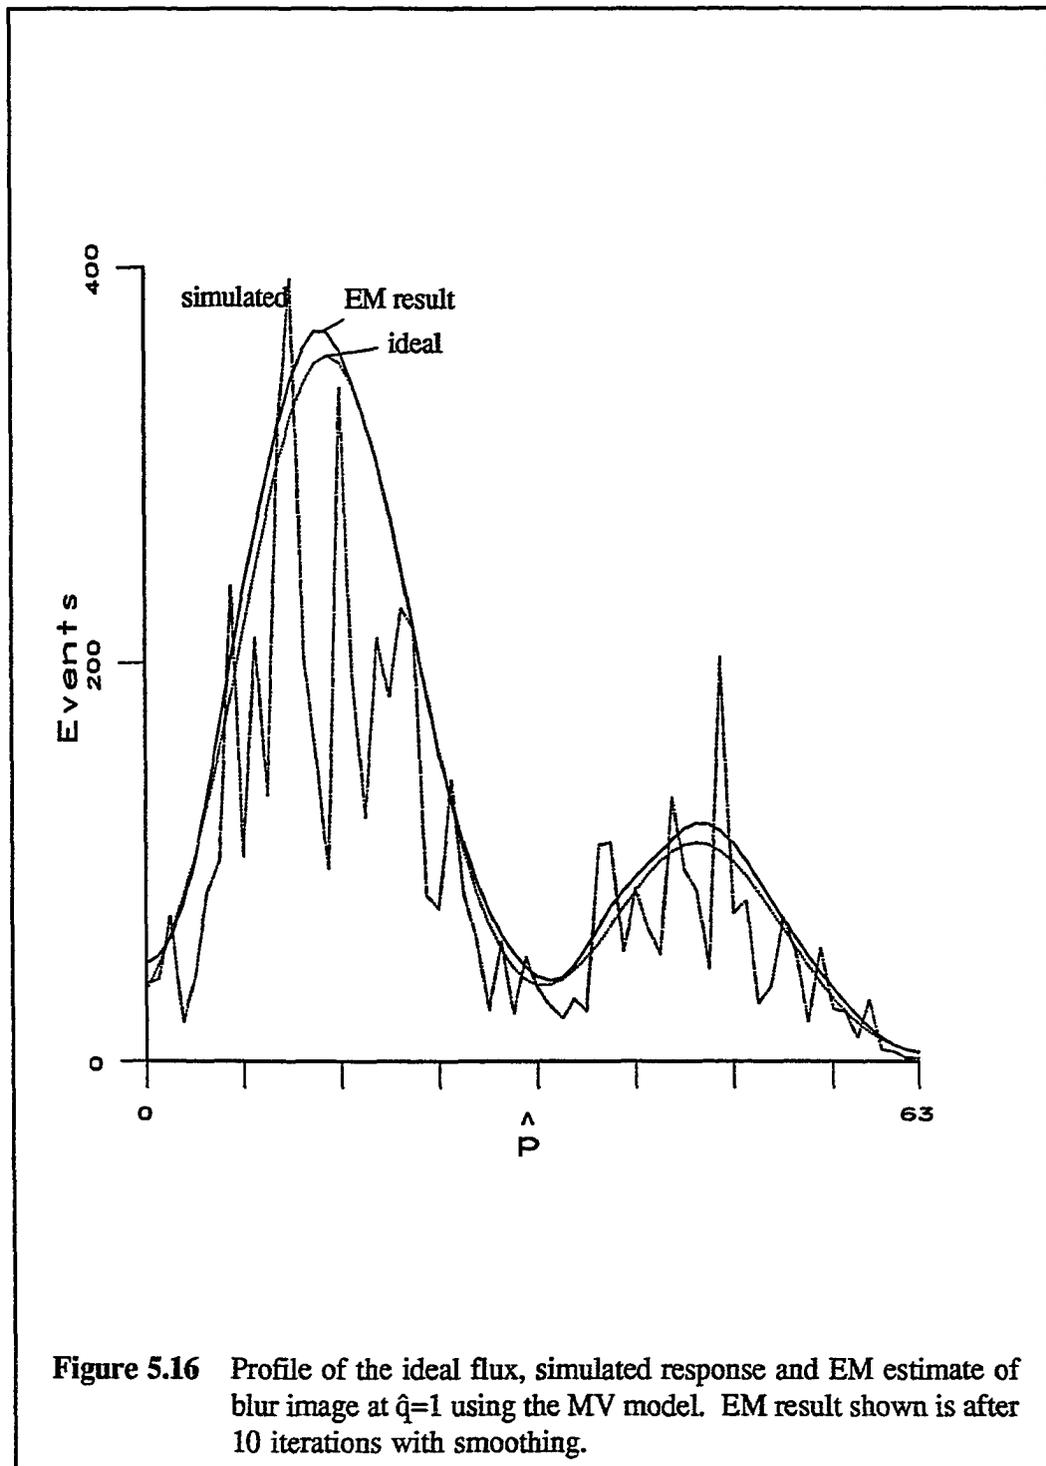
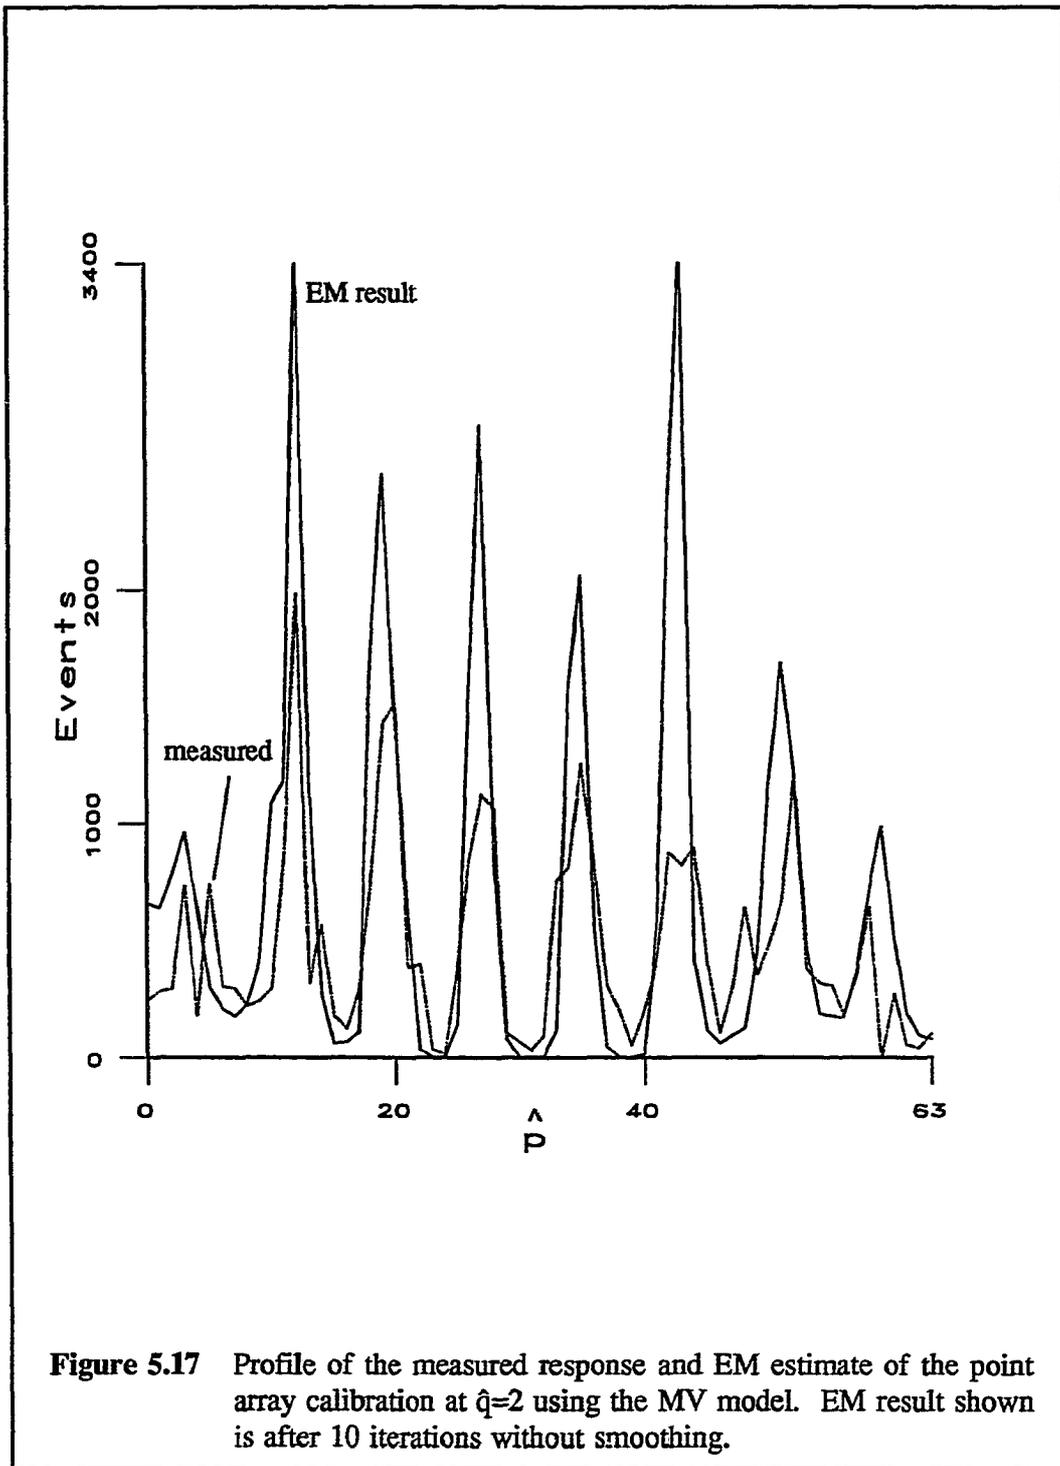
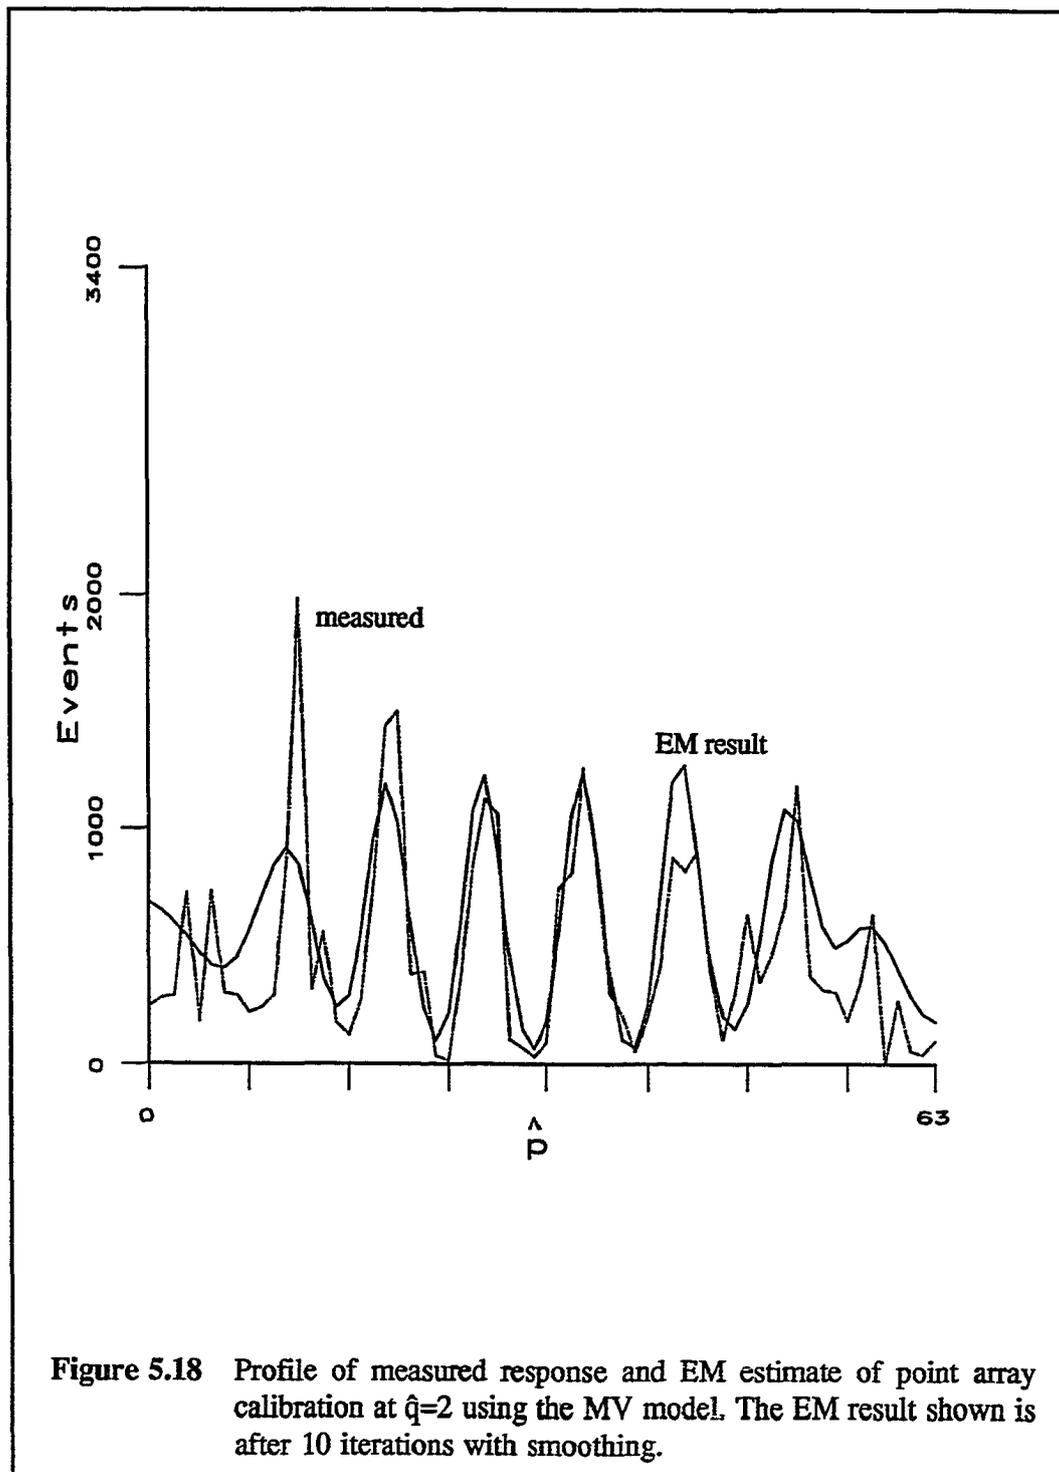


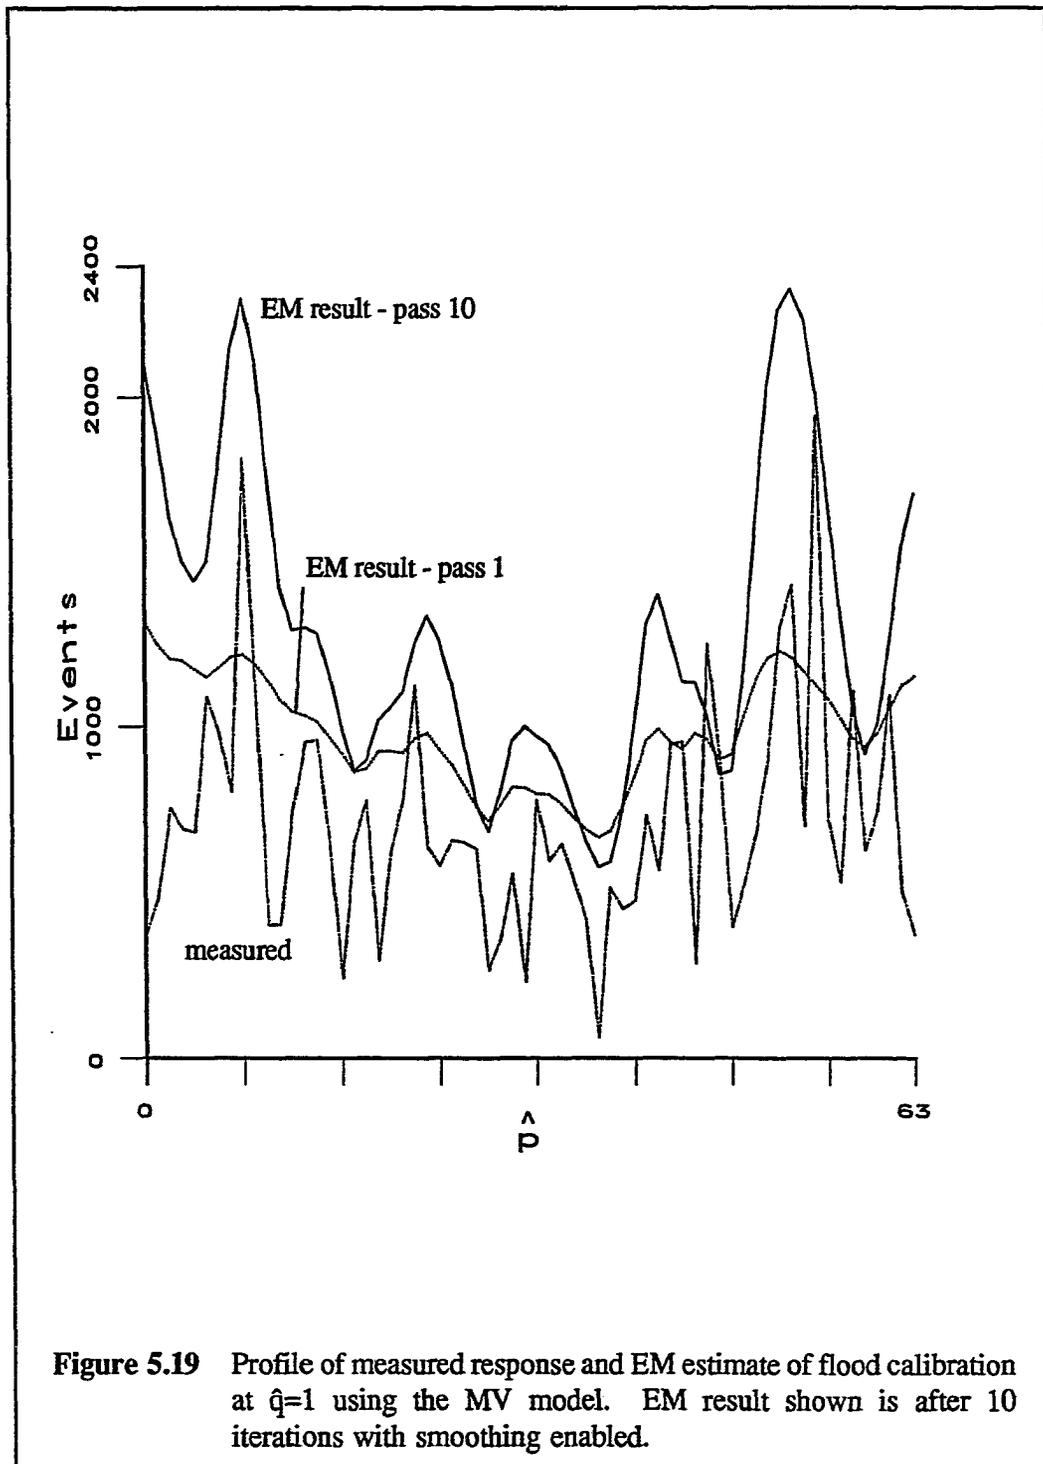
Figure 5.16 Profile of the ideal flux, simulated response and EM estimate of blur image at $\hat{q}=1$ using the MV model. EM result shown is after 10 iterations with smoothing.





contrast than the measured point array image. The valleys between points do not go to zero and the peaks are much lower than the measured peaks. The smoothing operation after each iteration prevents the estimate from resolving into higher spatial frequencies, and shifts events from the peaks into the surrounding pixels. The MVN EM estimate is blurred more than the Poisson EM estimate after the same number of iterations. If the weighting function for the smoothing operation was a narrower Gaussian, the EM algorithm would produce a better estimate. Without smoothing, the EM estimate is closer to a point array.

The measured MVN flood, previously shown in Figure 5.8, was processed using the EM algorithm with smoothing enabled. The standard deviation in the 56 x 56 central region is 157 events, which is close to the standard deviation of the Poisson flood for the same region (153 events). After EM processing, the MVN estimate has standard deviation of 90.8 in this same area, which is higher than the standard deviation of the EM processed Poisson model flood image (67). The mean value of the unprocessed image was 654 for the Poisson flood and 668 for the MVN model. The higher mean number of events in the MVN flood could be caused by a greater number of accepted ABCD values in the MVN model LUT. The increased standard deviation in the processed MVN model is not accounted for entirely by the increased mean, however, if we consider the Fano factor (variance/mean). The Fano factor for the processed Poisson image is 6.95 and is lower than the Fano factor for the processed MVN image of 11.79.



The calibration image results show that images generated and EM processed with the multivariate model are not superior to the images produced using the Poisson model. The flood images are less uniform and the point array images have less contrast. Nevertheless, the super-MDRF measurement indicates that the Poisson model assumption that the PMT signals are independent is in error. Another approach to the multivariate model may be required to produce results that are superior to results from the Poisson model. Suggestions for improving the multivariate model are presented in Chapter 6.

6. SUMMARY AND RECOMMENDATIONS FOR FUTURE WORK

We have shown in this work that the EM algorithm can be used to improve estimates of projection data collected by a SPECT imaging system. Two models for the camera response were examined. The Poisson model followed from the model used historically for the modular cameras and included the assumption that the PMTs were statistically independent. The EM algorithm was applied to simulated images, measured calibration images and projection images, employing the Poisson camera response matrix. EM processing decreased variance in flood images and resolved points that were distorted by noise in point array images. The most notable difference in all cases was the removal of high-frequency checkerboard patterns in the camera response. Simulations showed that the EM algorithm performs well if the camera response vector \mathbf{g} is consistent with the camera response matrix \mathbf{H} . Differences between simulated and measured camera response indicated that the model is in error primarily along the pixels close to the edge of the crystal.

The camera response was measured directly as a function of the four-dimensional, compressed (5-bit) PMT signal response to determine if the assumption of PMT independence was valid. The measurement indicated that the covariance between PMTs is up to 30% of the variance and cannot be considered negligible. The multivariate camera response model was applied for the first time in this research. Processing the simulated and calibrated images using the multivariate model produced

estimates that suggest that the model has potential for future use. The model does require revision in order to obtain estimates that are more accurate than those of the EM processed Poisson model.

A number of experiments would increase our understanding of how the EM algorithm can be used to pre-process projection data. Three objectives are especially of interest to the author: improving the camera response model, implementing the EM algorithm, and reconstructing object distributions from the EM processed projections.

Improvements to the Camera Response Model

Initial experiments with the multivariate model show that the model produces reasonable estimates, but it does not appear to be superior to the Poisson model. Both models have weaknesses. The Poisson model fails to incorporate correlation of the PMT signals. The current multivariate model is based on one measurement from a single modular camera. It is possible that the multivariate model could be improved by incorporating more empirical knowledge.

One weakness in the multivariate model is that the super-MDRF measurements could only be made using 5-bit signals per PMT. With 5-bit signals, each pixel is not associated with a unique 20-bit peak, therefore, the mean of the distribution was assumed to be floating point and to vary smoothly as a function of (p,q) to improve the LUT histogram. With 8-bit PMT signals, the distribution function could be measured more accurately. At completion of this work, the data acquisition hardware

was limited to 5-bit PMT signals, but use of 8-bit signals is planned for the near future. The increased precision of the measurements will clarify the relationship between the assumed statistical model and the true performance.

The current multivariate model represents the covariance matrix for all cameras by the covariance matrix of camera number 21. Measuring the super-MDRF from several cameras is necessary to determine if this assumption is reasonable. If all the cameras have unique covariance matrices, a new calibration process will be required. The covariance matrix used for the multivariate model was approximated by averaging adjacent PMT covariances and interpolating the results from a 16 x 16 pixel array to the 64 x 64 pixel array. Measurement of the super-MDRF for all pixels in a 16 x 16 square at the corner of the crystal would allow us to determine how the variance and covariance changes at the edges of the crystal. The bilinear interpolation of the covariance matrix near the edges of the crystal may result in poor estimation of (\hat{p}, \hat{q}) if the covariance values change nonlinearly at the edge.

Implementing the EM Algorithm

In this research, the EM algorithm was implemented with the intent of showing that improved estimates could be obtained, without planning how to implement the algorithm as part of the data acquisition process for the brain imager. One aspect of implementation is the method used to regularize the algorithm. The EM algorithm used in these studies was either unregularized, or regularized simply by smoothing the

estimate. In point array images, smoothing degrades the estimate. More sophisticated regularization methods such as those referenced in Chapter 3 could be examined. Alternatively, the unregularized EM algorithm could be applied and stopped before variance in low-frequency images becomes unacceptable.

At this writing, the data acquisition hardware is being redesigned to include a processor for each of the modular cameras. The 8-bit signals from the four PMTs could be combined to form a 32-bit integer that is supplied to the processor. The processor could perform various functions such as LUT mapping and flood correction. Future work will involve implementing the EM algorithm with the dedicated modular camera processor to compute an estimate of the projection data .

Reconstruction Approaches

The simulation studies showed that EM processing can improve the estimate of the camera response when the response vector is consistent with the system matrix. It is not clear, however, that the improved estimate yields better results in the object reconstruction. The current reconstruction method employs a system response matrix for the brain imager that includes the noisiness of each point projection. When the projection data are processed to remove this noisiness, the system matrix and projection data are no longer consistent. The need exists for a new representation for the brain imager system matrix that is consistent with the processed projection data.

Reconstructions with Measured Centroids and Averaged Blurs

The problem of checkerboarding in the blur function used in the reconstruction algorithm could be addressed by using an averaged, position-dependent blur for the blur function. Rather than averaging the blurs for a single centroid pixel in all of detector space, we could average a group of pixels in subregions of the camera, over all cameras, to produce a position dependent blur. The averages would not contain the checkerboarding pattern, because they would include the patterns from all cameras, and all centroids in the subregion. The blur function would retain the position dependent shape of the blur, which is elliptical near the edges, broader in the corners, and narrow in the center. This approach alone would not account for magnification of the blur, which depends on the distances between the voxel, pinhole and detector. To implement this approach would require only that the blur function file be rewritten. Existing reconstruction methods could be used with little modification.

Reconstructions with Measured Centroids and Theoretical Blurs

If a theoretical function could be developed to represent the blur associated with each centroid, the magnification of the blur could be incorporated as a function of voxel, pinhole and centroid positions. One approach could be to use a two dimensional Gaussian function, with variable widths in both directions. The width ratio would depend on camera position to represent the elliptical or symmetrical shape seen empirically. The theoretical function could be independent of the PSF

measurement and require minimal storage. This approach has been suggested for use with the unprocessed projections to account for magnification, but the difficulty in representing the checkerboarding in the PSF was previously a disadvantage to this approach. With the checkerboarding in the projection image removed, the smooth theoretical blur model holds more promise.

REFERENCES

Aarsvold J N, Barrett H H, Chen J, Landesman A L, Milster T D, Patton D D, Roney T J, Rowe R K, Seacat R H, III, Strimbu L M, 1988, "Modular Scintillation Cameras: A Progress Report", *Medical Imaging II: Image Formation Detection, Processing, and Interpretation*, SPIE, 914: 319-325.

Arendt J, 1980, "The Design and Analysis of a 1D Scintillation Detector," M.S. Thesis, University of Arizona.

Barrett H H and Swindell W, 1981, Radiological Imaging: Theory of Image Formation, Detection and Processing, Academic Press, New York.

Barrett H H, Wilson D W and Tsui B M W, 1993, "Noise Properties of the EM Algorithm: I. Theory", *Phys. Med. Biol.*, submitted for publication.

Coakley K J, 1991, "A Cross-Validation Procedure for Stopping the EM Algorithm and Deconvolution of Neutron Depth Profiling Spectra", *IEEE Trans. Nuc. Sci.* 38(1): 9-15.

Chandra R., 1987, Introductory Physics of Nuclear Medicine, Lea & Febiger, Philadelphia.

Chornoboy E S, Chen C J, Miller M I, Miller T R, Snyder D L, 1990, "An Evaluation of Maximum Likelihood Reconstruction for SPECT", *IEEE Trans. Med. Img.*, 9(1): 99-110.

Dempster A P, Laird N M, and Rubin D B, 1977, "Maximum Likelihood from Incomplete Data via the EM Algorithm". *J. Roy. Statistical Soc., Series B*, 39: 1-38.

Green P J, 1990, "Bayesian Reconstructions form Emission Tomography Data Using a Modified EM Algorithm", *IEEE Trans. Med. Img.*, 9(1): 84-93.

Hansen E W, 1988, "Statistics of the Modular Scintillation Camera", unpublished report submitted to the Department of Radiology, Arizona Health Science Center.

Hebert T J, Leahy R and Singh M, 1988, "Fast MLE for SPECT using and Intermediate Polar Representation and a Stopping Criterion", *IEEE Trans. Nuc. Sci.*, 35(1): 615-619.

- Hebert T J and Leahy R, 1989, "A Generalized EM Algorithm for 3-D Bayesian Reconstruction from Poisson Data Using Gibbs Priors", *IEEE Trans. Med. Img.*, 8(2): 194-202.
- Hebert T J, 1990, "Statistical Stopping Criteria for Iterative Maximum Likelihood Reconstruction of Emission Images", *Phys. Med. Biol.*, 35(9): 1221-1232.
- Hebert T J and Gopal S S, 1992, "The GEM MAP Algorithm with 3-D SPECT System Response", *IEEE Trans. Med. Img.*, 11(1): 81-90.
- Kaufman L, 1987, "Implementing and Accelerating the EM Algorithm for Positron Emission Tomography", *IEEE Trans. Med. Img.* MI-6(1): 37-51.
- Kimura K, Hashikawa K, Etani H, Uehara A, Kozuka T, Moriwaki H, Isaka Y, Matsumoto M, Kamada T, Moriyama H, and Tabuchi H, 1990, "A New Apparatus for Brain Imaging: Four-Head Rotating Gamma Camera Single-Photon Emission Computed Tomograph", *J Nuc Med.*, 31: 603-609.
- Lange K and Carson R, 1984, "EM Reconstruction Algorithms for Emission and Transmission Tomography", *J. Comp. Assist. Tomo.*, 8(2): 306-316.
- Levitan E and Herman G T, 1987, "A Maximum *A Posteriori* Probability Expectation Maximization Algorithm for Image Reconstruction in Emission Tomography", *IEEE Trans. Med. Img.*, MI-6(3): 185-192.
- Li J, Jaszczak R., Greer K, and Coleman E, 1993, "Implementation of an Accelerated EM Algorithm for Cone Beam SPECT", to be published.
- Macovski A, 1983, Medical Imaging Systems, Prentice Hall, Englewood Cliffs, New Jersey.
- Magee K A, 1990, "Parallel Implementation of Image Reconstruction Algorithms for Emission Tomography", M.S. Thesis, University of Arizona.
- Mar L, 1989, "Estimation and Control of the Mean Response of Photomultiplier Tubes", M.S. Thesis, University of Arizona.
- Miller M I, and Roysam B, 1991, "Bayesian Image Reconstruction for Emission Tomography Incorporating Good's Roughness Prior on Massively Parallel Processors", *Proc. Nat. Acad. Sci.*, 88: 3223-3227.

Milster T D, Selberg, L A, Barrett H H, Easton, R L, Rossi G R, Arendt, J, 1983, "A Modular Imaging System for Use in Nuclear Medicine," Scientific Exhibit, Annual Meeting, Soc. of Nucl. Med.

Milster T D, Selberg L A, Barrett H H, Easton R L, Rossi G R, Arendt J, Simpson R G, 1984, "A Modular Scintillation Camera for Use in Nuclear Medicine," IEEE Trans. Nuc. Sci., NS-31: 578-580.

Milster T D, Selberg L A, Barrett H H, Landesman A L, Seacat R H III, 1985, "Digital Position Estimation for the Modular Scintillation Camera," IEEE Trans. Nuc. Sci., NS-32: 748-752.

Milster T D, 1987, "Design and Construction of a Modular Gamma Camera," Ph.D. Dissertation, University of Arizona.

Milster T D, Aarsvold J N, Barrett H H, Landesman A L, Mar L S, Patton D D, Roney T J, Rowe R K, Seacat R H III, "A Full -ield Modular Gamma Camera", 1990, J. Nuc. Med., 31(5): 632-639.

Ranganath M V, Dhawan A P, and Mullani N, 1988, "A Multigrid Expectation Maximization Reconstruction Algorithm for Positron Emission Tomography", IEEE Trans. Med. Img., 7(4): 273-278.

Rowe R K, 1991, "A System for Three Dimensional SPECT without Motion", Ph.D. Dissertation, University of Arizona, Tucson.

Selberg L A, 1984, "Design Studies for a Modular Scintillation Camera," M.S. Thesis, University of Arizona.

Shepp L A and Vardi Y, 1982, "Maximum Likelihood Reconstruction for Emission Tomography", IEEE Trans. Med. Img., MI-1(2): 112-122.

Singh M, Leahy R, Brechner R and Yan X, 1989, "Design and Imaging Studies of a Position Sensitive Photomultiplier Based Dynamic SPECT System", IEEE Trans. Nuc. Sci., NS-36: 1132-1137.

Snyder D L, Miller M I, Thomas L J, Polite D G, 1987, "Noise and Edge Artifacts in Maximum-Likelihood Reconstructions for Emission Tomography", IEEE Trans. Med. Img., MI-6(3): 228-238.

Wilson D W, Tsui B M W, and Barrett H H, 1993, "Noise Properties of the EM algorithm: II. Monte Carlo Simulations", *Phys. Med. Biol.*, submitted for publication.

Veklerov E and Llacer J, 1987, "Stopping Rule for the MLE Algorithm Based on Statistical Hypothesis Testing", *IEEE Trans. Med. Img.*, MI-6(4): 313-319.

Veklerov E, Llacer J and Hoffman E J, 1988, "MLE Reconstruction of a Brain Phantom Using a Monte Carlo Transition Matrix and a Statistical Stopping Rule", *IEEE Trans. Nuc. Sci.*, 35(1): 603-607.

Zubal G, 1993, "Imaging and Quantitation Characteristics for Three Commercial SPECT Instruments", *Computer and Instrumentation Council Newsletter, Society of Nuclear Medicine*, 3-5.