A Unified Approach to Statistical Tomography Using Coordinate Descent Optimization [†‡]

Charles A. Bouman
School of Electrical Engineering
Purdue University
West Lafayette, IN 47907-0501
(317) 494-0340

Ken Sauer
Laboratory for Image and Signal Analysis
Department of Electrical Engineering
University of Notre Dame
Notre Dame, IN 46556
(219) 631-6999

December 4, 1996

Abstract [1]

Over the past ten years there has been considerable interest in statistically optimal reconstruction of image cross-sections from tomographic data. In particular, a variety of such algorithms have been proposed for maximum *a posteriori* (MAP) reconstruction from emission tomographic data. While MAP estimation requires the solution of an optimization problem, most existing reconstruction algorithms take an indirect approach based on the expectation maximization (EM) algorithm.

In this paper we propose a new approach to statistically optimal image reconstruction based on direct optimization of the MAP criterion. The key to this direct optimization approach is greedy pixel-wise computations known as iterative coordinate decent (ICD). We show that the ICD iterations require approximately the same amount of computation per iteration as EM based approaches, but the new method converges much more rapidly (in our experiments typically 5 iterations). Other advantages of the ICD method are that it is easily applied to MAP estimation of transmission tomograms, and typical convex constraints, such as positivity, are simply incorporated.

IP EDICS #2.3 Computed Imaging: Tomography

[†]This work was supported by the National Science Foundation under Grant no. MIP93-00560, Electricité de France under Grant no. P21L03/2K3208/EP542, and an NEC Faculty Fellowship

[‡]*IEEE Trans. on Image Processing,* vol. 5, no. 3, pp. 480-492, March 1996.

[1]Permission to publish this abstract separately is granted.

1 Introduction

In the past decade, emission tomography and other photon-limited imaging problems have benefited greatly from the introduction of statistical methods of reconstruction. Unlike the relatively rigid deterministically-based methods such as filtered backprojection, statistical methods can be applied without modification to data with missing projections or low signal-to-noise ratios. This makes statistical reconstruction methods well suited to emission problems, or transmission tomograms of dense materials. The Poisson processes in emission and transmission tomography invite the application of maximum-likelihood (ML) estimation, simply choosing the parameters in the discretized reconstruction which best match the data. Due to the typical limits in fidelity of data, however, ML estimates are unstable, and have been improved upon by regularized estimation, such as maximum *a posteriori* probability (MAP) estimation [1], or the method of sieves [2].

Both the ML and MAP reconstructions may be formulated as the solution to an optimization problem. However, this optimization problem is a formidable numerical task due to both the number of parameters in the estimate (pixels or voxels) and the number of observations (photon counting measurements). Since the work of Shepp and Vardi[3], the method of choice for finding ML estimates in emission tomography has been the expectation-maximization (EM) algorithm. The EM algorithm is based on the notion of a set of "complete" data which, if available, would make the problem easier.

For the emission tomography problem, the complete data are usually the number of photons emitted from each cell in the discretized reconstruction. While the use of such a large complete data set appears to simplify the computation of the ML estimate, Fessler and Hero[4] have shown that a large or "informative" complete data space also slows down convergence. This suggests that the fastest convergence should be achieved by using the actual observations as the complete data set, which is equivalent to direct optimization of the ML or MAP functionals.

In the Bayesian problem, the EM approach is less simple to apply to emission tomographic reconstruction. This is because the maximization step has no closed form and itself requires the use of an iterative optimization technique. To address this problem, many approaches have been proposed, but all of them approximate the otherwise intractable maximization step. For the case of Gaussian prior densities, Liang and Hart[5, 6], Herman and Odhner[7] and Herman, De Pierro, and Gai [8] have modified the EM approach to include Bayesian estimation. A variety of methods have also been proposed for adapting the EM algorithm to more general Markov random field priors. These methods include the Generalized EM (GEM) algorithm proposed by Hebert and Leahy[9, 10], the one-step-late (OSL) method proposed by Green[11], and a more general form of De Pierro's method[12]. Most recently, Fessler and Hero have proposed SAGE [13], a collection of methods designed to minimize the complete data with each pixel update in EM reconstruction.

For the transmission problem, the EM methods are more difficult to apply because there is not such an obvious choice for the complete data space[14]. Ollinger used the EM approach to solve the transmission problem and found that convergence required from 200 to 2000 iterations[15].

In this paper, we expand on the work first presented in [16] and take a direct optimization approach to the problem of MAP image reconstruction of emission and transmission tomograms. It is interesting to note that in spite of their relative simplicity, such direct optimization methods do not seem to have been investigated for exact statistical image reconstruction with Poisson measurement noise. However, we show that direct optimization is quite tractable for this problem. Moreover, a direct approach to optimization without the use of EM allows one to use a wide range of efficient optimization algorithms to achieve fast convergence. It may also be viewed as maximizing convergence speed by employing a minimally informative complete data set.

Our approach to optimization is based on the sequential greedy optimization of pixel (or

voxel) values in the reconstruction. This method, which we refer to as iterative coordinate descent (ICD) [2] goes by a variety of other names including iterated conditional modes (ICM) for MAP segmentation [17], and Gauss-Seidel (GS) for partial differential equations [18]. All these algorithms work by iteratively updating individual pixels or coordinates to minimize a cost functional. The ICD method was first applied to least squares Bayesian tomographic reconstruction by Sauer and Bouman [19]. More recently, Fessler has applied ICD to the least squares emission problem and shown that underrelaxation methods can speed ICD convergence[20].

The ICD method has a number of important advantages which make it a good choice for direct optimization. First, ICD can be efficiently applied to the log likelihood expressions resulting from photon limited imaging systems. In fact, each ICD iteration is similar in nature and computational complexity to an iteration of the EM algorithm; however, the function we attack is the posterior probability function rather than the Q function of EM. Second, the ICD algorithm is demonstrated to converge very rapidly (in our experiments typically 5-10 iterations) when initialized with the filtered back projection (FBP) reconstruction. This is not surprising since the FBP reconstruction is accurate at low spatial frequencies, and in [19] we showed that the ICD method (also known as Gauss-Seidel) has rapid convergence at high spatial frequencies. The third important advantage of the ICD algorithm is that it easily incorporates convex constraints, and non-Gaussian prior distributions [3]. In particular, positivity is an important convex constraint which can both improve the quality of reconstructions and significantly speed numerical convergence[21]. Non-Gaussian prior distributions are also important since they can substantially reduce noise while preserving edge detail[11, 22, 23, 24].

Our analysis starts with an approximate analysis of the optimization problem based on a Taylor expansion of the log likelihood function. This analysis is important for two

[2]We choose to use the name ICD since it is most descriptive of the algorithmic approach.

[3]We note that convex constraints, such as positivity, may be thought of as a special type of non-Gaussian prior information, and is therefore consistent with the Bayesian problem formulation.

reasons. It provides a common conceptual framework for both the emission and transmission problems, and it gives insight into the best choice of numerical techniques for the exact optimization problem. We also show that this Taylor series approximation leads to an expression similar to one proposed by Fessler[20] for the modeling of accidental coincidences in emission reconstructions.

We next develop the framework for the optimization of the exact posterior distribution. Based on the approximate analysis, we implement the ICD pixel updates using a variant on the classical Newton-Raphson root finding method. The per-iteration computational cost of ICD/Newton-Raphson is found to be similar to that of the EM algorithm, but unlike EM, the ICD/Newton-Raphson algorithm is easily adapted to the Bayesian problem.

2 Formulation of Statistical Problem

In this section, we will develop the statistical framework for the MAP reconstruction problem for both the emission and transmission case. We will also review the conventional EM approach for later comparison.

For the emission problem, λ is the N dimensional vector of emission rates, Y is the M dimensional vector of Poisson-distributed photon counts, λ_j represents the emission rate from pixel (voxel) j, and P_{ij} is the probability that an emission from pixel j is registered by the i^{th} detector. Thus, according to the standard emission tomographic model, the random vector Y has the distribution

$$\mathcal{P}(Y = y|\lambda) = \prod_{i=1}^{M} \frac{\exp\{-P_{i*}\lambda\}\{P_{i*}\lambda\}^{y_i}}{y_i!} \tag{1}$$

where the matrix P contains the probabilities $\{P_{ij}\}$, and P_{i*} denotes the vector formed by its i-th row. This formulation is general enough to include a wide variety of photon-limited imaging problems, and the entries of P may also incorporate the effects of detector response

and attenuation. Using (1), the log likelihood may be computed.

$$\text{(emission)} \qquad \log \mathcal{P}(Y = y|\lambda) = \sum_{i=1}^{M} (-P_{i*}\lambda \ + y_i \log\{P_{i*}\lambda\} - \log(y_i!)) \tag{2}$$

The corresponding result for the transmission case is discussed in [19]. In order to emphasize the similarity of the transmission problem, we use the same notation but interpret λ as the attenuation density of a pixel, and y as the observed photon counts. The log likelihood is then given by

$$\text{(transmission)} \qquad \log \mathcal{P}(Y = y|\lambda) = \sum_{i=1}^{M} \left(-y_T e^{-P_{i*}\lambda} + y_i(\log y_T - P_{i*}\lambda) - \log(y_i!)\right), \tag{3}$$

where y_T is the photon dosage per ray. Both log likelihood functions have the form

$$\log \mathcal{P}(Y = y|\lambda) = -\sum_{i=1}^{M} f_i(P_{i*}\lambda) \tag{4}$$

where $f_i(\cdot)$ are convex and differentiable functions. This common form will lead to similar methods of solving these two problems.

For the emission problem, maximum-likelihood (ML) estimation of λ from y yields the optimization problem

$$\hat{\lambda}_{ML} = \arg\min_{\lambda} \sum_{i=1}^{M} (P_{i*}\lambda \ - y_i \log\{P_{i*}\lambda\}) \ .$$

Probably the most widely applied algorithm for finding $\hat{\lambda}_{ML}$ is expectation-maximization (EM)[25], which was first applied by Shepp and Vardi[3] to the emission tomographic problem. EM solves the ML estimation by hypothesising the existence of *complete data*, which would allow very simple estimation of λ if available. For the emission tomographic problem, these observations are the number of photons actually emitted from each discretized cell of the reconstruction region. The iteration resulting from the EM formulation is

$$\lambda_j^{n+1} = \frac{\lambda_j^n}{\sum_{m=1}^{M} P_{mj}} \sum_{i=1}^{M} \frac{y_i P_{ij}}{\sum_{l=1}^{N} P_{il}\lambda_l^n} \tag{5}$$

where n indicates the number of the iteration, updating the entire reconstruction. Because the log likelihood is concave, this approach can be shown to converge to the ML estimate[3].

In Sections 3 and 4, we will show that the exact ML or MAP reconstruction may also be computed through direct optimization using the ICD algorithm. The ICD algorithm works by sequentially optimizing the log likelihood with respect to each pixel (or voxel) value, λ_j. In [19], we introduced a fast algorithm for implementing ICD in the transmission problem when the log likelihood is approximated by a single quadratic function. This basic ICD algorithm exploits the sparse nature of the projection matrix by maintaining a state vector $\tilde{p}^n = P\lambda^n$ of the projected values.

In this paper, we will investigate three new techniques for applying ICD to emission and transmission MAP reconstruction. The first technique is as in [19], but relies on a new quadratic approximation to the log likelihood derived for the emission problem. The second technique, which we will refer to as ICD/Half Interval, finds the solution through greedy sequential pixel updates, using a half-interval search to solve the problem exactly at each step. Finally, the method we call ICD/Newton Raphson (ICD/NR) forms a revised quadratic approximation to the log likelihood at each new pixel update. The ICD/Newton Raphson method solves the MAP reconstruction problem exactly, with the optimum estimate being its only fixed point. The details of computation will follow in Sections 3 and 4.

In order to compare these various algorithms, we will need an objective measure of computational complexity which is independent of the specific hardware or software implementation. For this purpose, we use two figures of merit: the approximate number of multiplies plus divides per full iteration, and the number of complete reads of the P matrix. In practice, we have found the number of matrix reads to be a good predictor of algorithm speed since computation is often dominated by memory access time and indexing overhead. Table 1 lists these two performance measures for the EM algorithm in terms of M_0 the average number of nonzero projection values associated with each pixel. [4] NM_0 is then the number of nonzero entries in the sparse projection matrix P. Notice that one iteration of the EM algorithm

[4]The entries in Table 1 assumes that $M_0 >> 1$, the sparse matrix P is precomputed and stored, and sums independent of the data are precomputed.

Algorithm	# of multi./div.	# of matrix reads
Filtered Backprojection	M_0N	1
EM (maximum likelihood)	$2M_0N$	2
ICD (Gauss-Seidel) with quadratic approximation	$3M_0N$	2
ICD/Half Interval exact	$3KM_0N$	2
ICD/Newton-Raphson exact	$4M_0N$	2

Table 1: Two measures of computational complexity which are independent of hardware/software implementation. The first measure is the number of multiplies plus divides per full iteration of each method. The second measure is the number times the P matrix is read from memory. N is the number of points in the image, M_0 is the average number of nonzero projections associated with each image pixel, and K is the average number of iterations required for the half interval search.

requires the computation of two iterations of filtered back projection.

For low signal-to-noise ratio medical imaging problems, the shortcomings of ML estimation are well documented[2, 26, 27]. Therefore, many researchers have resorted to some form of regularized estimation for tomographic inversion. Maximum *a posteriori* probability (MAP) estimation addresses this problem by adding regularization in the form an *a priori* density for λ. The MAP estimate has been shown to substantially improve performance in many image reconstruction and estimation problems. In addition, the computation of the MAP estimate is not prohibitively difficult provided that the log of the prior density is a concave function of λ.

A frequent choice for a prior model is the Gaussian MRF, but the quadratic penalty extracted for the Gaussian often causes excessive smoothing of edges. Several prior models have been developed which include desirable edge-preserving properties, and which maintain convexity in their log prior densities[11, 28, 23, 24]. Provided we choose one of these models, the MAP problem is also convex, and iterations converge to the unique global minimum solution. We use the Generalized Gaussian MRF (GGMRF) model proposed in [23] with the density function

$$p_\lambda(\lambda) = \frac{1}{z} \exp\left\{ -\gamma^q \sum_{\{j,k\} \in C} b_{j-k} |\lambda_j - \lambda_k|^q \right\}$$

where C is the set of all neighboring pixel pairs, b_{j-k} is the coefficient linking pixels j and

k, γ is a scale parameter, and $1 \leq q \leq 2$ is a parameter which controls the smoothness of the reconstruction. This model includes a Gaussian MRF for $q = 2$, and an absolute-value potential function with $q = 1$. In general, smaller values of q allow sharper edges to form in reconstructed images.

Prior information may also be available in the form of constraints on the reconstructed solution. We will assume that the set of feasible reconstructions Ω is convex, and in all experiments we will choose Ω to be the set of positive reconstructions[5]. Combining this prior model with the log likelihood expression of (4) yields the expression for the MAP reconstruction.

$$\hat{\lambda}_{MAP} = \arg\min_{\lambda \in \Omega} \left[\sum_{i=1}^{M} f_i(P_{i*}\lambda) \; + \; \gamma^q \sum_{\{j,k\} \in C} b_{j-k} |\lambda_j - \lambda_k|^q \right] \tag{6}$$

While the EM algorithm is not difficult to implement or understand in the ML case, it is not directly and simply applicable to MAP estimation when the complete data is taken to be the number of photons emitted from each pixel. This is because there is no closed form solution for the maximization step of the iteration. Hebert and Leahy[9, 10] have developed the GEM algorithm to cope with these effects. The GEM algorithm takes the form of coordinate gradient ascent of the MAP EM cost functional with a heuristic step size which can be adjusted to guarantee convergence. De Pierro's majorization method for MAP reconstruction is also guaranteed to converge[12]. The OSL method proposed by Green[11] uses an approximate maximization step based on the previous values for neighboring pixels. For all three of these algorithms, computation per iteration is approximately the same as listed for EM in Table 1.

[5]This is consistent with the Bayesian formulation since we may condition on the information that $\lambda \in \Omega$.

3 Computation of Approximate MAP estimate

The first step toward efficient direct optimization of (6) will be to develop a quadratic approximation to the log likelihood functions for the emission and transmission problems. This approximation is useful because it will guide the design of efficient optimization techniques for the exact problem, and because it gives important intuition into the method and its relationship to existing reconstruction algorithms.

In the Appendix, we compute the first two terms in a Taylor expansion to find a quadratic approximation to the emission log likelihood of (2). In [19], a similar quadratic approximation was derived for the transmission problem. Both approximations have the form

$$\log \mathcal{P}(Y = y|\lambda) \;\approx\; -\frac{1}{2}(\hat{p} - P\lambda)^T D(\hat{p} - P\lambda) + c(y) \tag{7}$$

where $\hat{p}$ is a vector of projection measurements, D is a diagonal matrix, and $c(y)$ is some function of the data. For the purposes of MAP estimation $c(y)$ may be ignored since it does not depend on λ. For the emission case, $\hat{p}$, and D are given by [6]

$$\begin{aligned} (emission) \qquad \hat{p} &= y \\ D &= \mathrm{diag}\{y_i^{-1}\} \end{aligned}$$

while for the transmission case they are given by

$$\begin{aligned} (transmission) \qquad \hat{p}_i &= \ln(y_T/y_i) & (8) \\ D &= \mathrm{diag}\{y_i\}\ . & (9) \end{aligned}$$

The placement and character of the diagonal matrix D gives immediate insight into both problems. For the transmission problem, D more heavily weights those projections which correspond to large photon counts, and therefore have high signal-to-noise ratios. The emission problem is the reverse, with larger photon counts resulting in greater measurement variance.

[6] For notational simplicity, we assume that all $y_i > 0$. The Appendix gives a more general quadratic approximation obtained by treating the terms corresponding to $y_i = 0$ separately.

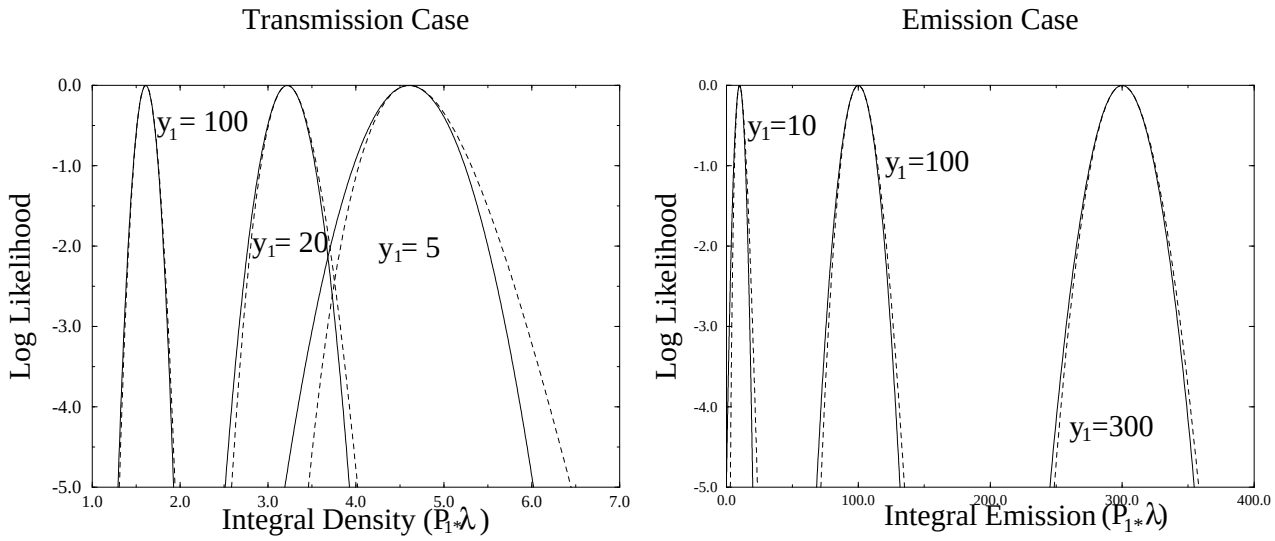

Figure 1: Plots of the relative log-likelihoods for (a) transmission and (b) emission tomography as a function of a single projection across the reconstructed image. The exact Poisson model (dashed lines) and the quadratic approximation (solid lines) are each plotted for three values of y_1, the observed photon count. In the transmission case, $y_T = 500$. Each of these plots covers a confidence interval of 0.99 for a maximum-likelihood estimate of the projection value.

Fig. 1 compares the exact log likelihood functions and their quadratic approximations over a range of photon counts. Both plots assume a single projection y_1 and use a 99% confidence interval. In the Appendix, we show that for both transmission and emission cases the approximation error bound is proportional to $\sum_{i=1}^{M} 1/\sqrt{y_i}$ for large y_i. This is consistent with the plots which show better approximations for larger y_i. In practice, the accuracy of this least squares approximation to the log likelihood function will depend on the dosage (transmission) or emission rates (emission), and the particular application.

Fessler[20] has independently developed a closely related approximation to the log likelihood for reconstruction of PET imagery from data precompensated for accidental coincidences. In this case, the Poisson model is violated since the observed counts of the compensated data can be negative, Fessler employed a Gaussian approximation for the density of corrected counts $\hat{y}$, yielding the objective function

$$\Phi(\lambda) = (\hat{y} - P\lambda)^T \mathrm{diag}\{\sigma_i^{-2}\}(\hat{y} - P\lambda).$$

Here the variance estimates are given by

$$\sigma_i^2 = n_i a_i^2 (a_i^{-1} \tilde{y}_i + 2r_i).$$

where n_i represents detector efficiency and a_i represent attenuation correction factors, $\tilde{y}$ is a smoothed version of $\hat{y}$, constrained to be positive, and r_i are accidental coincidence counts made in an independent measurement. Under this least squares formulation, Fessler found ICD to converge relatively quickly, improving on the speed of EM iterations.

Our approximate formulation of the log likelihood of (9) is similar to Fessler's approximation without accidental coincidences, except for the smoothing of the entries in $\hat{y}$. In section 5, we will show how the basic emission model of (2) can also be used to optimally account for accidental coincidences when calibration measurements are available.

4 Computation of Exact MAP estimate

To compute the exact MAP reconstruction, we must perform the optimization of (6). Of course, there is a wide variety of possible optimization techniques from which to choose, but we will use the approximate quadratic structure derived in section 3 to guide our selection.

Optimization techniques such as gradient ascent are undesirable because of their slow convergence [19, 29]. Alternatively, conjugate gradient[30] or preconditioned gradient descent[31, 32] techniques may be used since they are known to have rapid convergence for quadratic optimization problems. However, the performance of these techniques is less predictable for the nonquadratic problems resulting from non-Gaussian prior distributions. In fact, as q approaches 1, the log prior distribution is not differentiable, and gradient based optimization methods can become unreliable[33]. Possibly the most significant drawback of conjugate gradient and preconditioning methods is the difficulty of incorporating positivity constraints. The incorporation of positivity constraints for these methods requires the use of "bending" techniques which are potentially very computationally intensive[34].

We choose to use the ICD algorithm for a number of reasons. First, it was shown in [19] that the greedy pixel-wise updates of the ICD algorithm produce rapid convergence of the high spatial frequency components. Since the FBP can be used as a starting point of the algorithm, the convergence of the low spatial frequencies is less important. Second, the ICD updates work well with non-Gaussian prior models. In fact, the ICM algorithm, which is functionally equivalent to ICD, was developed for the MAP segmentation problem with a discrete prior model. Finally, the ICD algorithm is easily implemented along with convex constraints such as positivity. For this example, each pixel update is simply constrained to be positive.

The ICD algorithm is implemented by sequentially updating each pixel of the image. With each update the current pixel is chosen to minimize the MAP cost function. For emission tomography, the ICD update of the j^{th} pixel is given by

$$\lambda_j^{n+1} = \arg\min_{x \geq 0} \left[\sum_{i=0}^{M} \left[P_{ij} x - y_i \log \left(P_{ij}(x - \lambda_j^n) + P_{i*} \lambda^n \right) \right] + \gamma^q \sum_{k \in \mathcal{N}_j} b_{j-k} |x - \lambda_k^n|^q \right], \quad (10)$$

where $\mathcal{N}_j$ is the set of pixels neighboring j. Notice that in this case λ^n and λ^{n+1} differ at a single pixel, so a full update of the image requires that (10) be applied sequentially at each pixel.

No simple closed-form expression for λ_j^{n+1} results from (10), but there are many optimization techniques which can be employed to find its minimum. We will describe two strategies to the solution of (10). The first strategy is a direct application of half interval search. However, each iteration of this direct approach is significantly more computationally expensive than an iteration of EM based methods. The second strategy uses a technique similar to Newton-Raphson search to substantially reduce computation by exploiting the approximately quadratic nature of the log likelihood function.

Half interval search may be directly applied to solve (10) by searching for a zero in the derivative of the cost function. The cost function may be analytically differentiated to yield

the expression

$$\sum_i P_{ij}\left(1-\frac{y_i}{P_{ij}(x-\lambda_j^n)+P_{i*}\lambda^n}\right)+q\gamma^q\sum_{k\in\mathcal{N}_j}b_{j-k}|x-\lambda_k^n|^{q-1}\text{sign}[x-\lambda_k^n]\ .$$

The disadvantage of this direct approach is that it requires the repeated computation of the derivative. By maintaining a state vector for $P_{i*}\lambda^n$, the derivative may be evaluated with approximately $3M_0$ multiplies and divides. Therefore, the computation required for a complete update is given in Table 1 as $3KM_0N$ where K is the average number of half interval iterations required for convergence to the desired precision.

We may reduce computation of ICD updates by exploiting the fact that the log likelihood function is approximately quadratic. Newton-Raphson optimization works by applying a second order Taylor series approximation to the function being maximized. For example, if $g(x)$ is the function being maximized and x^n is the current value, then the new value x^{n+1} is given by

$$\begin{aligned}x^{n+1} &= \arg\min_x\left\{g(x^n)+(x-x^n)g'(x^n)+\frac{1}{2}(x-x^n)^2g''(x^n)\right\}\\ &= x^n-g'(x^n)/g''(x^n)\end{aligned}$$

where $g'(x)$ and $g''(x)$ are the first and second derivatives of $g(\cdot)$ respectively. Should $g(\cdot)$ be quadratic, we find the exact solution in a single step.

We apply the Newton-Raphson approach to the ICD updates by locally approximating the log likelihood function as quadratic. However, our method will deviate from conventional Newton-Raphson because we retain the exact expression for the log likelihood of the prior distribution. This is because the prior term is generally not well approximated by a quadratic function. Let θ_1 and θ_2 be the first and second derivatives of the log likelihood function evaluated for the current pixel value λ_j^n. Using our Newton-Raphson update, the new pixel value is given by

$$\lambda_j^{n+1}=\arg\min_{x\geq 0}\left[\theta_1(x-\lambda_j^n)+\frac{\theta_2}{2}(x-\lambda_j^n)^2+\gamma^q\sum_{k\in\mathcal{N}_j}b_{j-k}|x-\lambda_k^n|^q\right]\ . \tag{11}$$

This equation may be solved by analytically calculating the derivative and then numerically computing the derivative's root using the half interval method. This root finding computation is relatively inexpensive since θ_1 and θ_2 are precomputed and the number of neighboring pixels is generally small.

The complete set of ICD/Newton-Raphson update equations for emission data is then given by

$$\theta_1 = \sum_{i=1}^{M} P_{ij} \left(1 - \frac{y_i}{\tilde{p}_i^n}\right) \tag{12}$$

$$\theta_2 = \sum_{i=1}^{M} y_i \left(\frac{P_{ij}}{\tilde{p}_i^n}\right)^2 \tag{13}$$

$$0 = \theta_1 + \theta_2(x - \lambda_j^n) + q\gamma^q \sum_{k \in \mathcal{N}_j} b_{j-k} |x - \lambda_k^n|^{q-1} \text{sign}(x - \lambda_k^n) \Big|_{x=\lambda_j^{n+1}} \tag{14}$$

$$\tilde{p}^{n+1} = P_{*j}(\lambda_j^{n+1} - \lambda_j^n) + \tilde{p}^n \tag{15}$$

The root finding operation of (14) is usually computationally inexpensive, since the neighborhood $\mathcal{N}_j$ typically contains only a few pixels. [7] Therefore, the computation is dominated by the $4M_0$ total multiplies and divides required to compute θ_1 and θ_2. A full iteration consists of applying a single Newton-Raphson update to each pixel in λ. This results in $4M_0N$ operations per full image update. In practice, the computation is often dominated by the time required to index through the data. By this measure, the ICD/Newton-Raphson and EM algorithms are computationally equivalent, both requiring two indexings through the projection matrix P.

It should be noted that the distinction between the ICD/Newton-Raphson method and the approximate method of section 3, is that for ICD/Newton-Raphson the parameters of the quadratic approximation are recomputed for each new update. This guarantees that the exact MAP reconstruction is the only fixed point of the algorithm. While we have not yet shown that the ICD/Newton-Raphson method is guaranteed to converge to its fixed point,

[7] In some cases, computation time can be reduced by replacing the power function x^{q-1} by a linearly interpolated lookup table.

there are a number of reasons to believe that its convergence should be very stable. Since the true log likelihood is very close to quadratic the approximation of the Newton Raphson updates is quite accurate. Of course in the quadratic case, the ICD/Newton-Raphson method is very stable. In fact, ICD/Newton-Raphson method has an intrinsic safety factor since it remains stable in the quadratic case even with over relaxation by a factor of two [18]. In practice, we have never observed ICD/Newton-Raphson to have unstable or unreliable convergence.

The ICD/Newton-Raphson method may also be applied to the transmission tomography problem. In this case, the parameters θ_1 and θ_2 of (12) and (13) are then given by the following equations.

$$\text{(transmission)} \qquad \begin{aligned} \theta_1 &= \sum_{i=1}^{M} P_{ij} \left(y_i - y_T e^{-\tilde{p}_i^n} \right) \\ \theta_2 &= \sum_{i=1}^{M} P_{ij}^2 y_T e^{-\tilde{p}_i^n} \ . \end{aligned}$$

We note that these updates require the evaluation of exponential functions.

5 Attenuation and Accidental Coincidence Effects

Typical emission tomographic systems include imperfections which have been included into EM-based estimation approaches. For example, emitted photons have some non-zero probability of attenuation before they are registered. This probability is related to the transmission characteristics of the object, and can typically be measured by a preliminary transmission scan. In this case, the attenuation probability can be included directly in the P_{ij} values. Another possibility is the joint estimation of emission and attenuation properties[35]. For this work, we have assumed attenuation probabilities are included into P.

In positron emission tomography (PET) imaging, a significant fraction of registered photons are caused by "accidental coincidences," i.e. readings from distinct positron/electron annihilations which are counted as having arisen from a single event. This increases the

expected value of the count y_i by some number $\bar{r}_i$. If this value is assumed known, as in [36], then the log likelihood values may simply be chosen to reflect a Poisson distribution with the adjusted mean.

Another possibility is that independent calibration measurements are made which result in Poisson random variables R_i with mean $\bar{r}_i$. In principle this measurement can be made for a PET imaging system by counting the number of coincidences that occur between delayed time intervals[20, 37]. In this case, the optimal reconstruction can be performed by simply augmenting the vector of projection measurements, y, to include the independent Poisson calibration measurements $r = [r_1, \cdots, r_M]^t$. We define the quantities

$$\begin{aligned} \tilde{\lambda} &= \begin{bmatrix} \lambda \\ \bar{r} \end{bmatrix} \\ \tilde{y} &= \begin{bmatrix} y \\ r \end{bmatrix} \\ \tilde{P} &= \begin{bmatrix} P & I \\ 0 & I \end{bmatrix} \end{aligned}$$

where I is the identity matrix. Since $\tilde{y}$ is a vector of independent Poisson measurements with mean $\tilde{P}\tilde{\lambda}$, all the techniques of the previous sections may be used directly. We note that updates to the components of $\bar{r}$ are simple to compute since only two "projection" measurements are associated with each value $\bar{r}_i$.

6 Experimental Results

Fig. 2 shows the phantoms we used for our emission and transmission tomography experiments together with the FBP reconstructions. The emission phantom in Fig. 2(a) represents higher emission rates with higher image intensity, having zero emission from the background. Rates are scaled to yield a total count of approximately 5×10^4 for the section in Fig. 2(a), with readings taken at 64 equally spaced angles, and 64 perfectly collimated detectors at each angle. The transmission phantom represents an object of diameter 20 cm and density 0.2 cm^{-1}, with higher density regions of up to 0.48 cm^{-1} added. The dosage per ray (y_T)

of only 500 in this experiment results in zero counts at many detectors. Though this is well below typical medical transmission CT imaging rates, low dosage reconstructions are useful in reconstructing an approximate attenuation map for emission imaging[38]. The FBP reconstructions of Fig. 2 show significant noise and streaking artifacts. This is typical of FBP reconstructions since they do not account for the relative accuracy of projection measurements.

We choose an 8-point neighborhood system for the GGMRF, with normalization of weights $\{b_{j-k}\}$ to a total of 1.0 for each j, and $b_{j-k} = (2\sqrt{2}+4)^{-1}$ for nearest neighbors and $b_{j-k} = (4+4\sqrt{2})^{-1}$ for diagonal neighbors. In order to illustrate the effect of the Bayesian prior, we will compute reconstructions for both $q = 2$ and $q = 1.1$. The first case is equivalent to the common Gaussian prior, and the second does a better job of preserving edges. Since the log prior is strictly concave and differentiable in both cases, convergence of numerical algorithms can be guaranteed. We choose scale parameters yielding the qualitatively best results for comparison of reconstructions under the exact and approximate likelihoods. The parameters of the prior model were $(q = 2.0, \gamma = 1.0)$, $(q = 1.1, \gamma = 3.0)$ for the emission problem, and $(q = 2.0, \gamma = 15.0)$, $(q = 1.1, \gamma = 40.0)$ for the transmission problem.

Fig. 3 and Fig. 4 show the results of MAP reconstruction for the emission and transmission problems. In each figure, the exact and approximate MAP reconstructions are shown. The exact MAP reconstructions were computed using a large number of iterations of the ICD/Newton-Raphson algorithm of section 4. In each case the cost function being minimized is convex and the algorithm converges to the global minimum. Therefore, the MAP reconstruction computed using the EM algorithm will be identical. Since the form of the approximate log likelihood (7) is the same for the emission and transmission problems, the ICD algorithm (called Gauss-Seidel) of [19] was used to calculate both approximate MAP estimates. This algorithm has been shown to converge rapidly and requires approximately $3M_0N$ operations per iterations as listed in Table 1.

In both examples of Fig. 3, there are small but perceptible differences between the exact and approximate reconstructions. Fessler has found that the quadratic approximation can introduce significant bias into transmission reconstructions under low dosages[39]. So exact MAP transmission reconstruction may be of value if precise and absolute density measurements are required.

We concentrate on the exact emission reconstruction problem for comparison of convergence rates to previously proposed algorithms, since it is in this arena that the majority of recent research activity has taken place. (Convergence of ICD in transmission tomography was treated in [19].) Three alternatives to ICD/NR appear in the plots. Green's one-step-late (OSL) algorithm allows EM to be applied to MAP problems by adding a regularizing term to each EM maximization step which is based on the previous iteration's pixel values[11]. The update is made by setting to zero the sum of the gradient of the EM functional, and the gradient of the log of the prior density, evaluated at pixel values from the previous iteration. The OSL is very simple to compute, but may fail to converge. The generalized expectation-maximization (GEM) of Hebert and Leahy [9] substitutes an increase in the MAP/EM objective function for the more difficult maximization. GEM features an adjustable step size for the update accompanied by evaluation of the cost functional to guarantee increase in *a posteriori* likelihood. While the vector $\tilde{p}$ is updated after all pixels have been visited, the pixel values used in evaluating the log of the prior density are updated sequentially. Finally, we include DePierro's method[8], which guarantees convergence through a MAP/EM approach which decouples the computation of pixel updates in the maximization step. This technique may therefore be applied to complete parallel updates. While DePierro's method was originally designed for the case of a Gaussian prior density, it applies for other convex penalties as well[12], such as the GGMRF with $q = 1.1$.

The three alternative algorithms were implemented without modification to their originally proposed forms. All methods could include a one-dimensional search for each pixel's

update, which is assumed available in both De Pierro's method and ICD/NR. Both of these techniques require the minimization of a non-quadratic function at each pixel in the case of a non-Gaussian prior. The adjustable step size of GEM may be replaced by a minimization, and OSL may be augmented to vary the influence of the derivative of the log of the prior at each step to guarantee convergence as well[11]. While computational costs are often dominated by accessing the data and projection matrix, the expense of the 1-D minimizations may be appreciable for certain priors.

We will plot convergence performance in terms of complete updates of the image, since like the computational cost measures of Table 1, it is independent of implementation. Each of the four methods was initialized in all cases with an FBP reconstruction, which is of negligible cost relative to the ensuing computation. Since the *a posteriori* log likelihood is strictly convex, the solution will not be influenced by this choice of initial condition. Because low frequency components in the error between the FBP image and the MAP reconstruction will converge most slowly[19], we correct the zero-frequency component of the initial condition with a least-squares estimate directly from the data.

Figure 5 shows the convergence rates for the maximum-likelihood problem ($\gamma = 0$). In this case, all three alternative methods reduce to the EM algorithm. The ICD/NR estimate has, for practical purposes, converged after 5 or 6 iterations, while EM appears to require over an order of magnitude more. This behavior is at least partly explained by the similarity of EM to gradient ascent, which is particularly slow for this type of problem[19].

Figures 6 and 7 illustrate similar results for the MAP problem. With the Gaussian prior model, the three EM-based algorithms perform similarly in early iterations, but OSL fails to converge for this case, settling into an oscillation significantly below the maximum *a posteriori* likelihood. For a scale of $\gamma = 2.0$ with the same data, OSL diverged badly from the solution. Both GEM and DePierro's method approach the optimum, but as in Fig. 5, the convergence is much slower than ICD/NR.

Non-Gaussian models allow better preservation of abrupt transitions in MAP estimates. One example which possesses this advantage along with convexity of the potential function is the GGMRF with values of q near 1.0. We use $q = 1.1$, which affords good edge preservation with tractable optimization. While this likelihood has first derivatives which are well behaved, the second derivative of the function is unbounded, which may have varying effects on the optimization approaches. Although convergence is similar in this case also, interesting differences exhibit themselves in Fig. 7. OSL again fails to converge, which is not surprising given the character of the derivative of $|x|^{1.1}$ near the origin. GEM is the fastest of the three EM-type methods in this problem, reaching parity with the ICD/NR solution at about 50 iterations. There is also some potentially interesting asymptotic behavior. After about 60 iterations, when the estimate is undergoing very minor changes, the log likelihood of the GEM estimate slightly exceeds that of ICD/NR. Asymptotic characteristics of ICD/NR in nonlinear problems may require further study and improvement.

7 Conclusion

We have presented a new method of computing MAP reconstructions using direct optimization of the log likelihood function. Each iteration of our proposed ICM/Newton-Raphson algorithm has computation comparable to an iteration of the EM algorithm. However, the new method works well with Bayesian prior distributions and converges rapidly. The direct optimization approach also gives a common framework for solving both the emission and transmission tomography problems. Experiments indicated that while a fixed quadratic approximation is adequate for some transmission problems, optimization of the exact likelihood appeared to yield improved results in the experiments presented here.

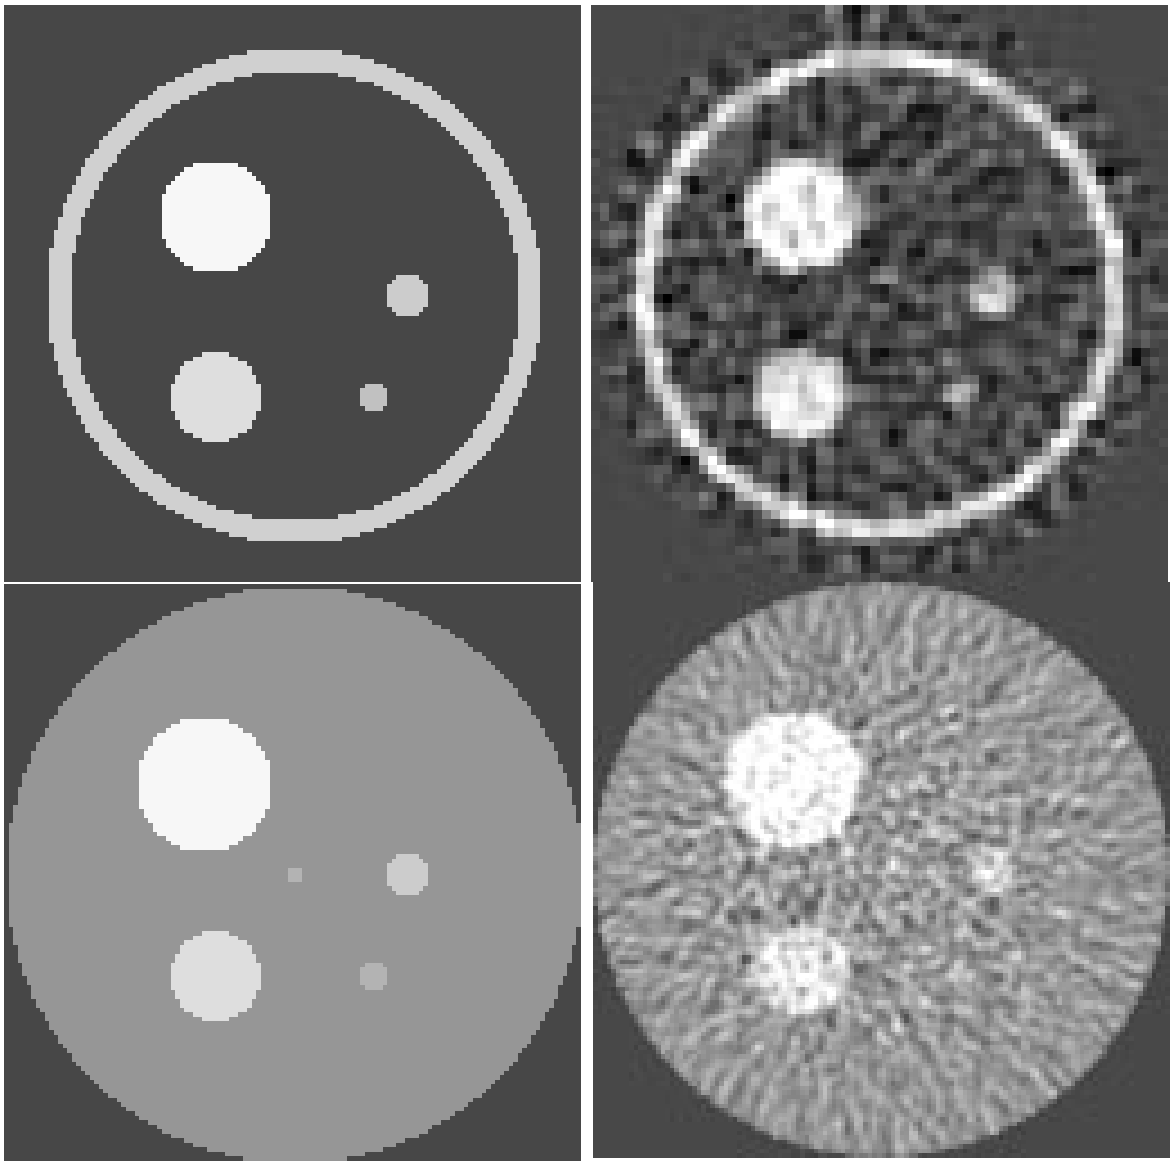

a	b
c	d

Figure 2: Original synthetic phantoms and their FBP reconstructions for emission and transmission examples. (a) Emission phantom with higher emission intensities in lighter areas, and (b) FBP emission reconstruction. (c) Transmission phantom with higher density in lighter areas, and (d) FBP transmission reconstruction. FBP reconstructions were computed using a raised cosine rolloff filter, and served as the initial estimate for the iterative statistical methods.

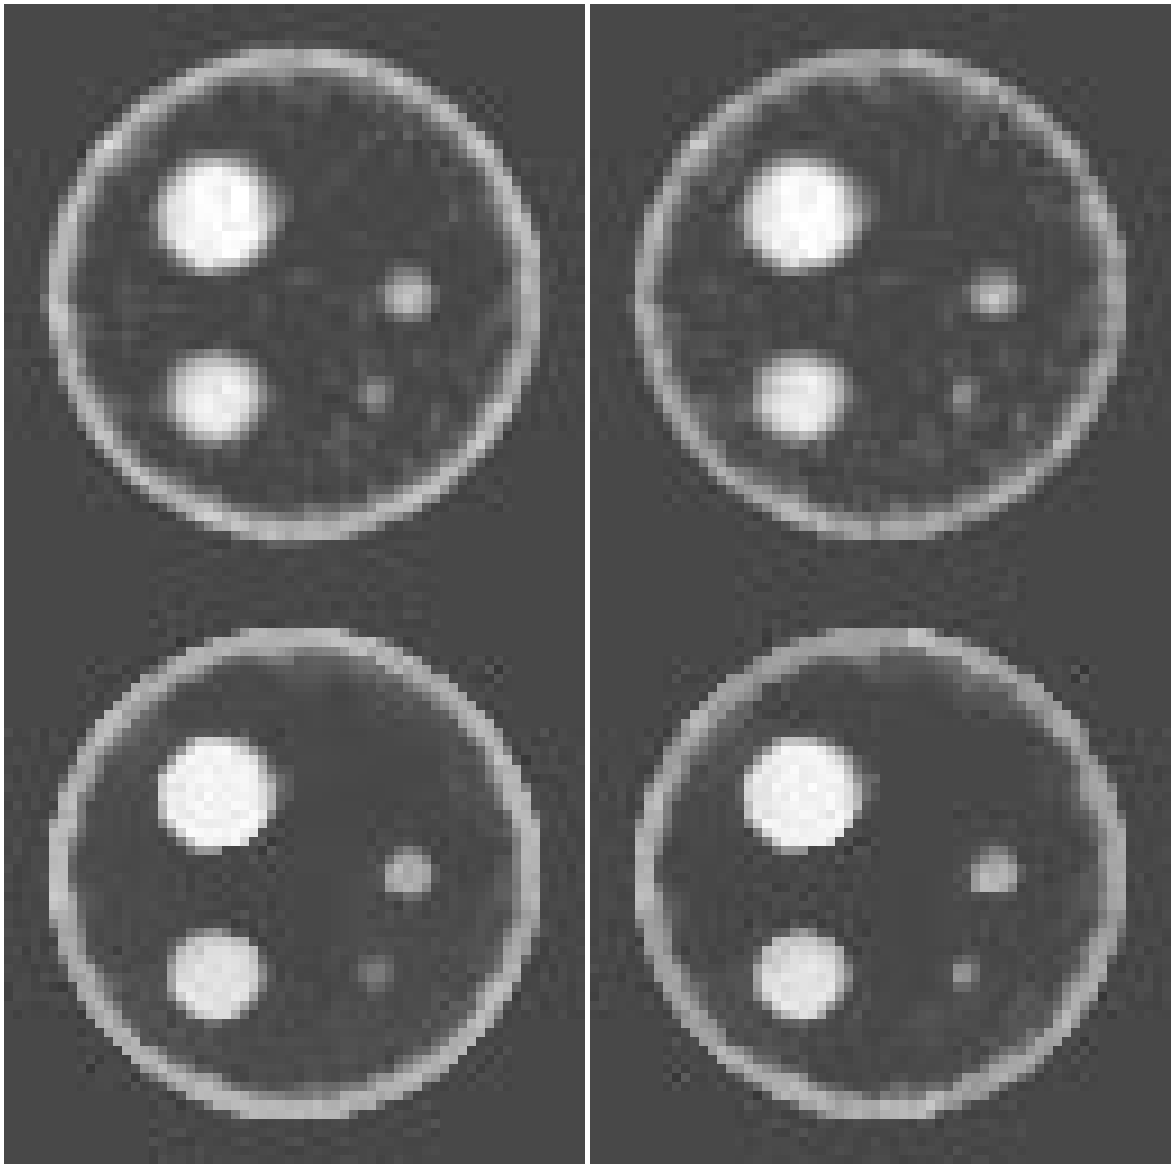

a	b
c	d

Figure 3: Emission MAP reconstructions with a Gaussian MRF prior, and (a) exact reconstruction using ICD/Newton-Raphson; (b) quadratic approximation. MAP estimates resulting from GGMRF model with $q = 1.1$, and (c) exact reconstruction using ICD/Newton-Raphson; (d) quadratic approximation.

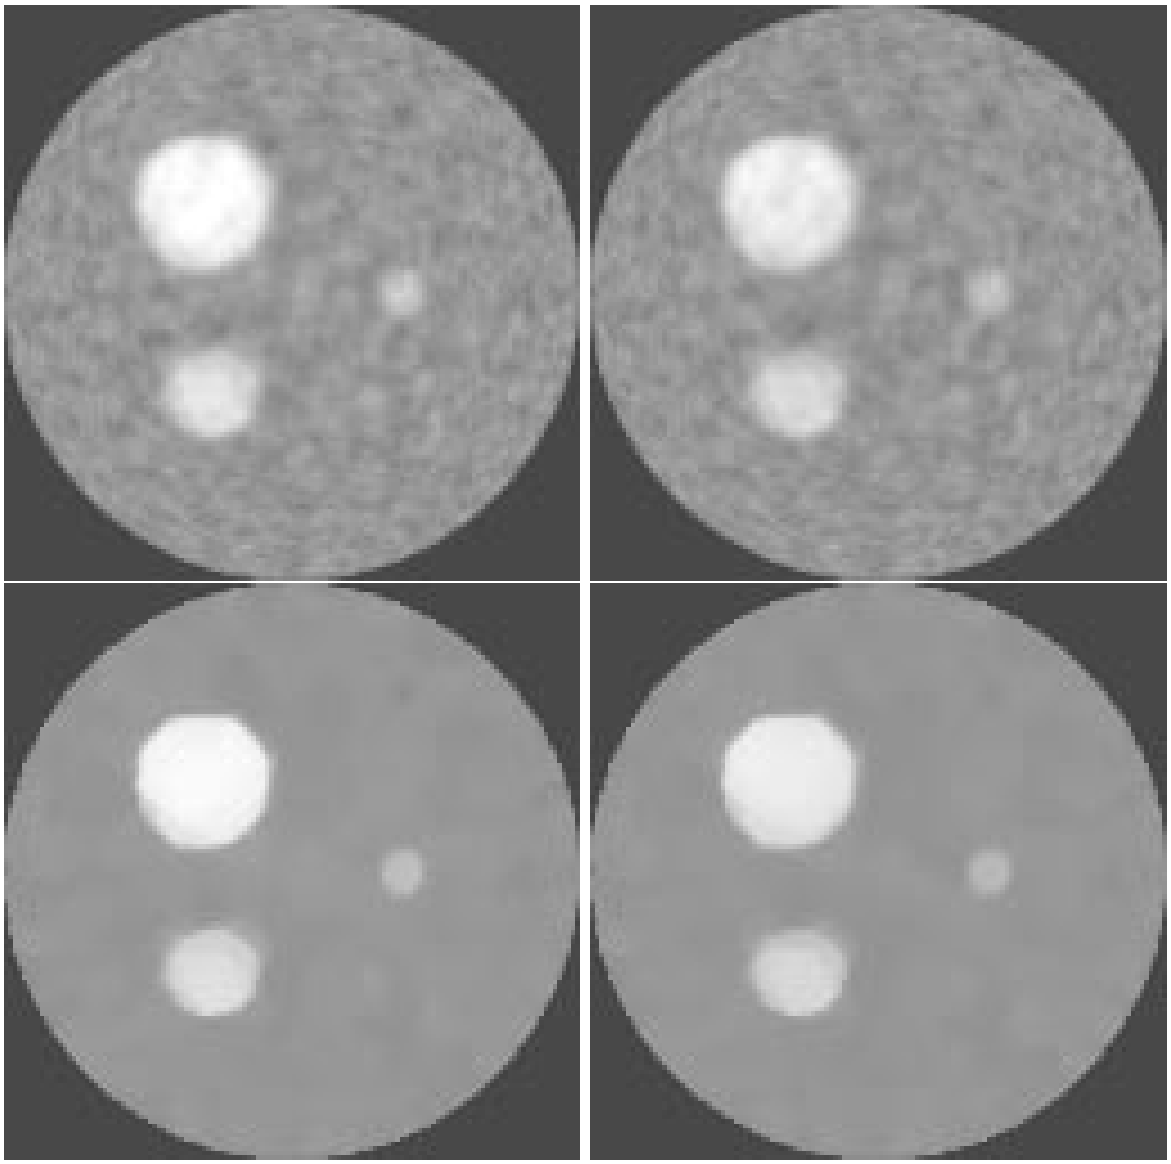

a	b
c	d

Figure 4: Transmission MAP reconstructions with a Gaussian MRF prior, and (a) exact reconstruction using ICD/Newton-Raphson; (b) quadratic approximation. MAP estimates resulting from GGMRF model with $q = 1.1$, and (c) exact reconstruction using ICD/Newton-Raphson; (d) quadratic approximation.

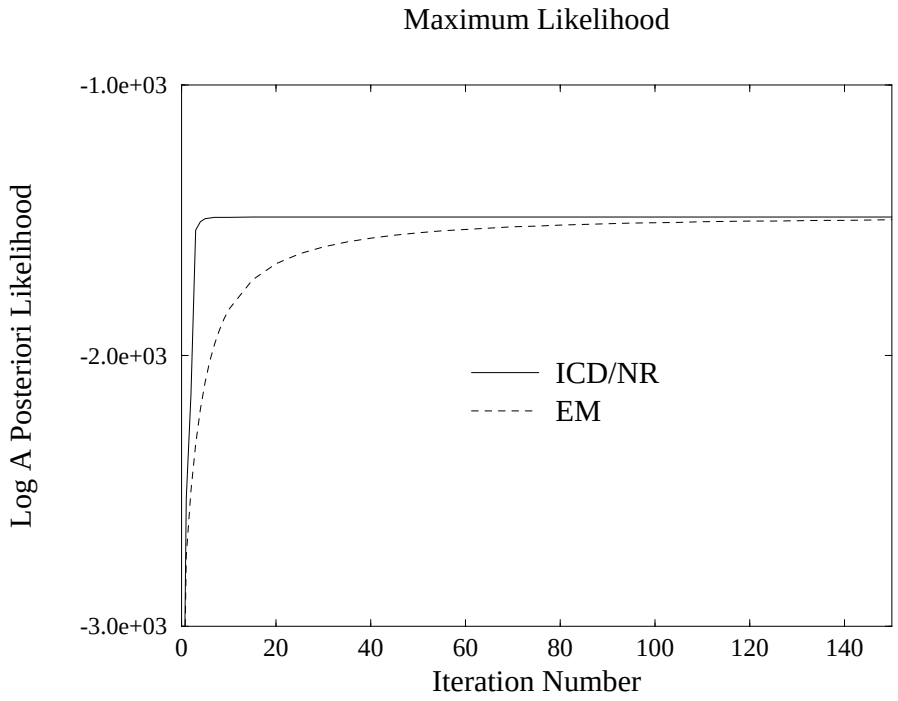

Figure 5: Convergence of ML estimates using ICD/Newton-Raphson updates and EM. The *a posteriori* likelihood function values are plotted as a function of full iterations.

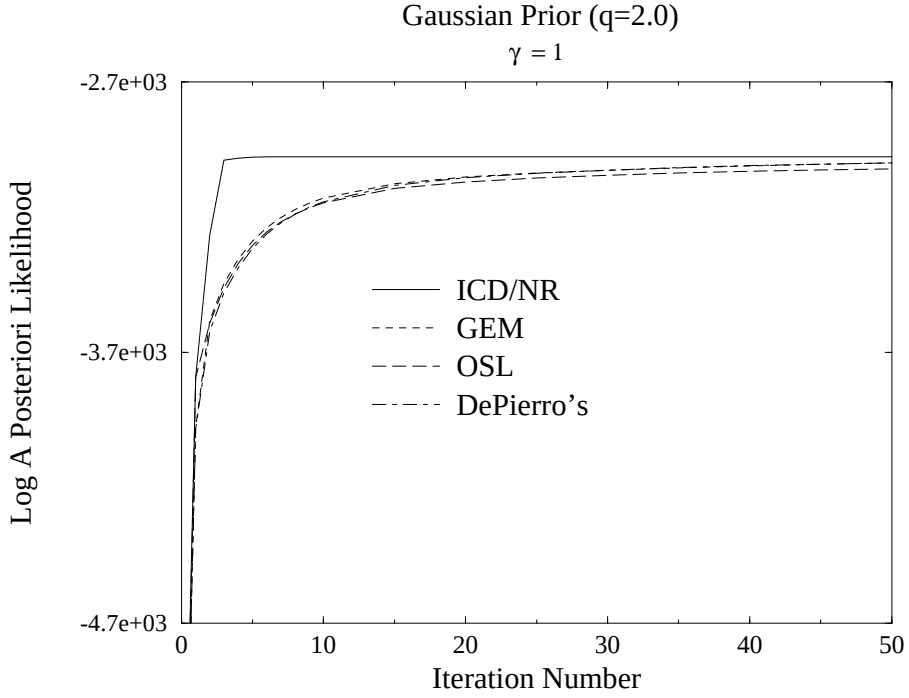

Figure 6: Convergence of MAP estimates using ICD/Newton-Raphson updates, Green's (OSL), and Hebert/Leahy's GEM, and De Pierro's method, and a Gaussian prior model with $\gamma = 1.0$.

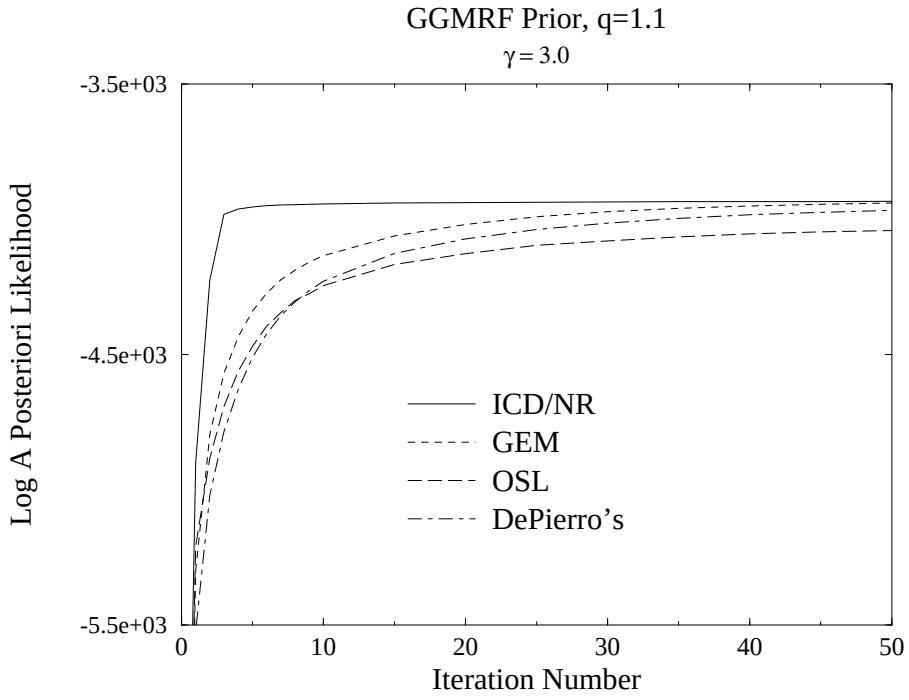

Figure 7: Convergence of MAP estimates with a generalized Gaussian prior model with $q = 1.1$ and $\gamma = 3.0$.

8 Appendix: Taylor Series Approximation of Emission Log Likelihood

In this Appendix, we derive the Taylor series approximation for the emission log likelihood of (2). We will also determine the convergence behavior of the quadratic approximation for both the transmission and emission case.

Using the convention that $\tilde{p}_i = P_{i*}\lambda$, the log likelihood may be expressed as

$$\log \mathcal{P}(Y = y|\lambda) = -\sum_i f_i(\tilde{p}_i)$$

where for the emission case

$$f_i(\tilde{p}_i) = \tilde{p}_i - y_i \log\{\tilde{p}_i\} + \log(y_i!) \ .$$

If we assume that $y_i > 0$, then the gradient of the likelihood function evaluated at $\tilde{p} = y$ has entries

$$\left.\frac{\partial \log \mathcal{P}(Y = y|\lambda)}{\partial \tilde{p}_i}\right|_{\tilde{p}=y} = -1 + \left.\frac{y_i}{\tilde{p}_i}\right|_{\tilde{p}=y}$$

$$= \quad 0$$

The Hessian is diagonal, with

$$\begin{aligned} \left.\frac{\partial^2 \log \mathcal{P}(Y=y|\lambda)}{\partial \tilde{p}_i^2}\right|_{\tilde{p}=y} &= \left.\frac{-y_i}{\tilde{p}_i^2}\right|_{\tilde{p}=y} \\ &= -\frac{1}{y_i} \end{aligned}$$

We may account for the terms with $y_i = 0$ by including them separately. Let $S_0 = \{i : y_i = 0\}$ and let $S_1 = S - S_0$. Then the approximation for the log likelihood is given by

$$\log \mathcal{P}(Y=y|\lambda) \approx \sum_{i \in S_1} -\frac{1}{2y_i}(y_i - P_{i*}\lambda)^2 + \sum_{i \in S_0} -(y_i - P_{i*}\lambda) + c(y) \ .$$

If we ignore the terms in S_0, then the log likelihood may be simply written as

$$\begin{aligned} \log \mathcal{P}(Y=y|\lambda) &\approx -\frac{1}{2}(y - P\lambda)^T D(y - P\lambda) + c(y), \\ \text{with} \quad D &= \text{diag}\{y_i^{-1}\}. \end{aligned}$$

We next show that the quadratic approximation for the log likelihood function converges as $\sum_i 1/\sqrt{y_i} \to \infty$ for both the transmission and emission cases. For the Taylor series approximation of a function $g(x)$,

$$g(x) = g(a) + g'(a)(x-a) + \frac{g''(a)(x-a)^2}{2!} + R_3,$$

we have the Lagrange form of the remainder

$$R_3 = \frac{g^{(3)}(\epsilon)(x-a)^3}{3!}, \tag{16}$$

where ϵ is between a and x. We will show in both the transmission and emission tomographic cases, that (16) goes to zero on arbitrarily large confidence intervals as the photon counts, y_i, become large.

In order to cover an arbitrarily large confidence interval for the parameter $\tilde{p}_i$, we will assume that $\tilde{p}_i \in [y_i - k\sqrt{y_i}, y_i + k\sqrt{y_i}]$ where k is any fixed positive integer. For large y_i,

the total error is then bounded by

$$\begin{aligned} E &= \sum_i \frac{f_i^{(3)}(\epsilon_i)(\tilde{p}_i - y_i)^3}{3!} \\ &= \sum_i \frac{2y_i}{\epsilon_i^3} \frac{(\tilde{p}_i - y_i)^3}{6} \\ &\leq \sum_i \frac{y_i (k\sqrt{y_i})^3}{3(y_i - k\sqrt{y_i})^3} \\ &\approx \frac{k^3}{3} \sum_i \frac{1}{\sqrt{y_i}} \ . \end{aligned}$$

Thus the error goes to zero as $\sum_i 1/\sqrt{y_i}$.

For the transmission problem

$$f_i(\tilde{p}_i) = y_T e^{-\tilde{p}_i} - y_i(\ln y_T - \tilde{p}_i) + \log y_i!$$

To cover an arbitrary confidence interval, we assume $\tilde{p}_i \in [\hat{p} - k/\sqrt{y_i}, \hat{p} + k/\sqrt{y_i}]$ where $\hat{p}_i = \log(y_T/y_i)$. This results in the following error bound for large y_i.

$$\begin{aligned} E &= \frac{f_i^{(3)}(\epsilon_i)(\tilde{p}_i - \hat{p}_i)^3}{3!} \\ &= \sum_i y_T e^{-\epsilon_i} \frac{(\tilde{p}_i - \hat{p}_i)^3}{6} \\ &\leq \sum_i \frac{k^3 e^{\frac{k}{\sqrt{y_i}}}}{6\sqrt{y_i}} \\ &\approx \frac{k^3}{6} \sum_i \frac{1}{\sqrt{y_i}} \end{aligned}$$

Thus in both cases, the bound on the error magnitude $\to 0$ as $\sum_i 1/\sqrt{y_i}$.

References

[1] S. Geman and D. McClure, "Bayesian Image Analysis: An Application to Single Photon Emission Tomography," in *Proc. Statist. Comput. Sect. Amer. Stat. Assoc.*, Washington, DC, pp. 12-18, 1985.

[2] D. Snyder and M. Miller, "The use of Sieves to Stabilize Images Produced with the EM Algorithm for Emission Tomography," *IEEE Trans. on Nuclear Science*, vol. NS-32, no. 5, pp. 3864-3871, Oct. 1985.

[3] L. Shepp and Y. Vardi, "Maximum Likelihood Reconstruction for Emission Tomography," *IEEE Trans. Med. Im.,* vol. MI-1, no. 2, Oct. 1982.

[4] J.A. Fessler and A.O. Hero, "Complete-Data Spaces and Generalized EM Algorithms," *Proc. ICASSP 1993*, vol. 5, pp. 1-4, April 27-30, 1993, Minneapolis MN.

[5] H. Hart and Z. Liang, "Bayesian Image Processing in Two Dimensions," *IEEE Trans. Med. Im.,* vol. MI-6, no. 3, pp. 201-208, Sept. 1987.

[6] Z. Liang and H. Hart, "Bayesian Image Processing of Data from Constrained Source Distributions; I:Non-valued, uncorrelated and correlated constraints," *Bull. Math. Biol.,* vol. 49, no. 1, pp. 51-74, 1987.

[7] G.T. Herman and D. Odhner, "Performance Evaluation of an Iterative Image Reconstruction Algorithm for Positron Emission Tomography," *IEEE Trans. Med. Im.,* vol. 10, no. 3, Sept. 1991.

[8] G. T. Herman, A. R. De Pierro, and N. Gai, "On Methods for Maximum A Posteriori Image Reconstruction with a Normal Prior," *J. Visual Comm. Image Rep.,* vol. 3, no. 4, pp. 316-324, Dec. 1992.

[9] T. Hebert and R. Leahy, "A Generalized EM Algorithm for 3-D Bayesian Reconstruction from Poisson data Using Gibbs Priors," *IEEE Trans. Med. Im.*, vol. 8, no. 2, pp. 194-202, June 1989.

[10] T. Hebert and S. Gopal, "The GEM MAP Algorithm with 3-D SPECT System Response," *IEEE Trans. Med. Im.*, vol. 11, no. 1, pp. 81-90, March 1992.

[11] P. J. Green "Bayesian Reconstructions from Emission Tomography Data Using a Modified EM Algorithm," *IEEE Trans. Med. Im.*, vol. 9, no. 1, pp. 84-93, March 1990.

[12] A.R. De Pierro, "A Modified Expectation Maximization Algorithm for Penalized Likelihood Estimation in Emission Tomography," to appear in *IEEE Trans. on Med. Im.*, Dec. 1994.

[13] J. A. Fessler and A. O. Hero, "Space-Alternating Generalized EM Algorithm for Penalized Maximum-Likelihood Reconstruction," University of Michigan, Technical Report No. 286, Feb. 1994.

[14] K. Lange and R. Carson,, "EM Reconstruction Algorithms for Emission and Transmission Tomography," *Journal of Computer Assisted Tomography*, vol. 8, no. 2, pp. 306-316, April 1994.

[15] J. M. Ollinger, "Maximum-Likelihood Reconstruction of Transmission Images in Emission Computed Tomography via the EM Algorithm," *IEEE Trans. Med. Im.*, vol. 13, no. 1, pp. 89-101, March 1994.

[16] C. Bouman and K. Sauer, "Fast Numerical Methods for Emission and Transmission Tomographic Reconstruction," *Proceedings of the Twenty-seventh Annual Conference on Information Sciences and Systems*, pp. 611-616, The Johns Hopkins University, Baltimore, MD, March, 1993.

[17] J. Besag, "On the Statistical Analysis of Dirty Pictures," *J. Roy. Statist. Soc. B,* vol. 48, no. 3, pp. 259-302, 1986.

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

[19] K. Sauer and C. Bouman, "A Local Update Strategy for Iterative Reconstruction from Projections," *IEEE Trans. on Sig. Proc.*, vol. 41, no. 2, pp. 533-548, Feb. 1993.

[20] J.A. Fessler, "Penalized Weighted Least-Squares Image Reconstruction for Positron Emission Tomography," *IEEE Trans. Med. Im.*, vol. 13, no. 2, pp. 290-300, June 1994.

[21] P. Oskoui-Fard and H. Stark, "Tomographic Image Reconstruction Using the Theory of Convex Projections," *IEEE Trans. Med. Im.*, vol. 7, no. 1, pp. 45-58, March 1988.

[22] T. A. Gooley and H. H. Barrett, "Evaluation of Statistical Methods of Image Reconstruction through ROC Analysis," *IEEE Transactions on Medical Imaging,* vol. 11, no. 2, June 1992.

[23] C. Bouman and K. Sauer, "A Generalized Gaussian Image Model for Edge-Preserving MAP Estimation," *IEEE Trans. Image Proc.,* vol. 2, no. 3, pp. 296-310, July 1993.

[24] R. Stevenson, B. Schmitz, and E. Delp, "Discontinuity Preserving Regularization of Inverse Visual Problems," *IEEE Transactions on Systems, Man, and Cybernetics,* vol. 24, no. 3, March, 1994.

[25] A.P. Dempster, N.M. Laird, and D.B. Rubin, "Maximum Likelihood from Incomplete Data via the EM Algorithm," *J. Royal Stat. Soc. B,* vol. 39, pp. 1-38, 1977.

[26] E. Veklerov and J. Llacer, "Stopping Rule for the MLE Algorithm Based on Statistical Hypothesis Testing," *IEEE Trans. Med. Im.,* vol. MI-6, pp. 313-319, 1987.

[27] T. Hebert, R. Leahy, and M. Singh, "Fast MLE for SPECT Using an Intermediate Polar Representation and a Stopping Criterion," *IEEE Trans. Nucl. Sci.,* vol. NS-35, pp. 615-619, Feb. 1988.

[28] K. Lange, "Convergence of EM Image Reconstruction Algorithms with Gibbs Priors," *IEEE Trans. Med. Im.*, vol. 9, no. 4, pp. 439-446, Dec. 1990.

[29] W. Press, B. Flannery, S. Teukolsky and Vetterling, *Numerical Recipes in C: The Art of Scientific Computing,* Cambridge University Press, Cambridge, 1988.

[30] F. Beckman, "The Solution of Linear Equations by the Conjugate Gradient Method," in eds. A. Ralston, H. Wilf and K. Enslein, *Mathematical Methods for Digital Computers,* Wiley, 1960.

[31] T.-S. Pan, A. E. Yagle, N. H. Clinthorne, and W. L. Rogers, "Acceleration and Filtering in the Generalized Landweber Iteration Using a Variable Shaping Matrix," *IEEE Transactions on Medical Imaging,* vol. 12, no. 2, pp. 278-286, June 1993.

[32] N. H. Clinthorne, T.-S. Pan, P.-C. Chiao, W. L. Rogers, and J. A. Stamos, "Preconditioning Methods for Improved Covergence Rates in Iterative Reconstruction," *IEEE Trans. on Medical Imaging,* vol. 12, no. 1, March 1993.

[33] K. Sauer and C. Bouman, "Bayesian Estimation of Transmission Tomograms Using Segmentation Based Optimization," *IEEE Trans. on Nuclear Science*, vol. 39, no. 4, pp. 1144-1152, Aug. 1992.

[34] D. M. Goodman, "Appling Conjugate Gradient to Image Processing Problems" *Proceedings of the Seventh Workshop on Multidimensional Signal Processing,* p. 1.1, September 23-25, 1991, Lake Placid, New York.

[35] N.H. Clinthorne, J.A. Fessler, G.D. Hutchins, and W.L. Rogers, "Joint Maximum Likelihood Estimation of Emission and Attenuation Densities in PET," *Proc. 1991 IEEE Nucl. Sci. Symp. & Med. Im. Conf.,* Nov. 2-9, 1991, Sante Fe, NM, pp. 1927-1932.

[36] D.G. Politte and D.L. Snyder, "Corrections for Accidental Coincidences and Attenuation in Maximum-Likelihood Image Reconstruction for Positron-Emission Tomography," *IEEE Trans. Med. Im.,* vol. 10, no. 1, pp. 82-89, March 1991

[37] T. Spinks, T. Jones, M. Gilardi and J. Heather, "Physical Performance of the Latest Generation of Commercial Positron Scanner," *IEEE Trans. Nucl. Sci.*, vol. 35, no. 1, pp. 721-725, Feb. 1988.

[38] R.H. Huesman, S.E. Derenzo, J.L. Cahoon, A.B. Geyer, W.W. Moses, D.C. Uber, T. Vuletich, and T.F. Budinger, "Orbiting Transmission Source for Positron Tomography," *IEEE Trans. Nucl. Sci.,* vol. NS-35, vol. 1, pp. 735-739, Feb. 1988.

[39] J. A. Fessler, "Sequential Iterative Algorithms for Image Reconstruction," presented at the *Midwest Workshop on Iterative Image Reconstruction,* June 17-18, 1994, Washington University St. Louis, Missouri.