

SIGNAL PROCESSING PITFALLS IN POSITRON EMISSION TOMOGRAPHY

Jeffrey A. Fessler and John M. Ollinger

COMMUNICATIONS & SIGNAL PROCESSING LABORATORY
Department of Electrical Engineering and Computer Science
The University of Michigan
Ann Arbor, Michigan 48109-2122

Sep. 1996

Technical Report No. 302
Approved for public release; distribution unlimited.

Signal Processing Pitfalls in Positron Emission Tomography

Jeffrey A. Fessler and John M. Ollinger
4240 EECS Bldg., University of Michigan, Ann Arbor, MI 48109-2122
Email: fessler@umich.edu, Voice: 734-763-1434,

Published as part of “Positron emission tomography,”
IEEE Signal Processing Magazine, 14(1):43-55, Jan. 1997.

I. PET MEASUREMENT STATISTICS

The previous sections have outlined the physics of PET. Now we describe the statistics. This summary is condensed from [1]. The measurement statistics are actually quite complex, so any treatment (including ours) must make simplifying assumptions. However, many papers in the signal processing and statistics literature considerably over-simplify the problem, e.g. [2], so we attempt to be somewhat more complete here.

Since PET measurements are based on a counting process, a reasonable statistical model is that the measurements have independent Poisson distributions:

$$Y_i \sim \text{Poisson}\{\bar{Y}_i(\lambda)\}, \quad i = 1, \dots, n,$$

where n is the number of coincident detector pairs, $\lambda(\underline{x})$ is the spatial distribution of radiotracer (typical units are counts/s/cm³), and \bar{Y}_i is the mean of the i th measurement. (Note that each i corresponds to a unique $d\theta$ pair in the notation used above.) The measurement means depend on the radiotracer distribution $\lambda(\underline{x})$ through the physical model described above; for low to moderate counting rates, the dependence is predominantly linear in λ :

$$\bar{Y}_i(\lambda) = T \left(\int p_i(\underline{x}) \lambda(\underline{x}) d\underline{x} + s_i(\lambda) + r_i(\lambda) \right), \quad (17)$$

where T is the scan time, $p_i(\underline{x})$ is the (scatter-free) point-response function (unitless) of the i th detector pair $s_i(\lambda)$ is mean rate of scattered events for the i th detector pair, $r_i(\lambda)$ is the mean rate of random coincidences for the i th detector pair, and the integral is over the scanner field of view. Although the scatter contribution $s_i(\lambda)$ is linear in λ , the random coincidences $r_i(\lambda)$ depend nonlinearly on λ (roughly related to the square of λ , as described above. [4]). This nonlinear dependence is not easily modeled directly. Thus, rather using modeling, most PET centers use the separate “delayed coincidence window” measurements [4] to obtain information about r_i as described above.

For moderate counting rates, the linearity in λ implied by the first term in (17) is reasonable. However, for high count rates, the measurement means are highly nonlinear functions (in fact nonmonotonic functions) of the activity in the patient due to scanner deadtime [5]. Specifically $p_i(\underline{x})$ should be replaced by something like $p_i(\underline{x}; \lambda)$, since as λ increases, the detection probability will decrease since there will be increasing probability of multiple events within the coincidence timing window (as well as other factors). In routine practice, the effect of this nonlinearity is reduced to a single global “deadtime correction factor”, based on the sum of the Y_i 's, or, more accurately, by different correction

This work was supported in part by NIH grants CA-60711 and CA-54362.

If a deterministic finite number of nuclei are injected into the patient, then strictly speaking a multinomial distribution would be more precise than the Poisson assumption. However, in practice the exact number of nuclei is unknown and may well be considered a random variable with a Poisson distribution. In this case the radioactive decay will be a Poisson process; furthermore, a Poisson process “thinned” by Bernoulli trials remains Poisson [3], all of which leads to the Poisson model.

$p_i(\underline{x})$ is probability that a positron emitted from a nuclei at position \underline{x} will produce a pair of annihilation photons that are detected by the i th detector pair without scattering (including geometric effects, attenuation, and detector efficiencies).

factors for different detector pairs or detector blocks. This type of correction implicitly separates the nonlinear deadtime loss from the ideal linear relationship between λ and $\{\bar{Y}_i\}$. Thus a representation of the spirit of this post-correction method in model form would be as follows:

$$\bar{Y}_i(\lambda) = d_i(\lambda) T \left(\int p_i(\underline{x}) \lambda(\underline{x}) d\underline{x} + s_i(\lambda) + r_i(\lambda) + b_i \right),$$

where $d_i(\lambda)$ is a positive, unitless deadtime loss function that monotonically decreases away from unity as its argument increases. We are unaware of any attempts to estimate λ directly from such a model; most model-based methods have ignored this nonlinearity term completely without comment, or have included a single data-based correction term ($d_i = \hat{d}_i(\lambda)$) in the p_i 's. We also take the latter approach here.

II. PET RECONSTRUCTION PITFALLS

The goal of image reconstruction is to recover the radiotracer concentration λ from the measurements $\{Y_i\}$. This inverse problem is not unlike the classical signal processing problem of *deconvolution* [6]. However, we believe that straightforward application of “off-the-shelf” signal processing and image restoration methods yield suboptimal results for PET image reconstruction.

In this section we summarize some of the methods that have been proposed for PET image reconstruction, with a particular emphasis on those with origins in signal and image processing. This review is by no means complete, and is primarily intended to describe the potential “pitfalls” of each approach, in the hope of steering newcomers to PET away from over-simplified methods that are unlikely to lead to significant improvements. Most of the discussion also applies to SPECT image reconstruction.

A. Deterministic Measurement Models

One way to greatly (over)simplify the problem is to ignore the measurement noise altogether, and to approximate the detector point-response functions $p_i(\underline{x})$ by equally-spaced parallel lines. This leads to the classical filtered-backprojection (FBP) method for tomographic image reconstruction [7]. The FBP method is used routinely for X-ray CT, as well as for PET and SPECT, for historical reasons of computational simplicity. Since the FBP method is derived without any statistical information, it is unsurprising that a “pure” use of the FBP method leads to unacceptable noise amplification (due to the ramp filter) in PET.

B. Sinogram Preprocessing

To remedy this problem, one usually apodizes the ramp filter by a classical window such as the Hanning or Parzen window. Such apodization is equivalent to space-invariant smoothing of the sinogram data. Although this smoothing does reduce the noise variance, it is suboptimal since PET measurement statistics are very *nonstationary* due the Poisson distribution. There have been a few attempts to improve the sinogram smoothing using both iterative [8–12] and noniterative [13] nonstationary methods. While requiring less computation than the iterative methods described below, these preprocessing methods are still suboptimal since object constraints such as nonnegativity and piecewise smoothness are not naturally expressed in the sinogram domain.

C. Tomogram Post-processing

The radiotracer distribution estimate computed by any reconstruction method is typically represented by a discrete image. This certainly invites the application of many an image processing method, both those classical (such as Wiener filtering) as well as those trendy (such as wavelets, neural nets, etc.). Unfortunately, most image processing methods are based on the (often implicit) assumption that the noise is Gaussian, or at least independent from pixel to pixel. The noise in tomographic images is generally highly correlated between neighboring pixels (since each measurement “ray” transects many pixels). For the (linear) FBP method the correlation function can be calculated [14–16]. In principle the noise correlation function can also be determined for some statistical image reconstruction methods [17–19], although the correlation functions may be expensive to compute. In our experience, classical image processing methods perform

The reason for the prevalence of use of FBP is certainly not due to any widely accepted advantage in image quality.

poorly for images with such correlated noise. Furthermore, the correlation structure is often nonstationary, so noise prewhitening is usually impractical. On the other hand, post-processing methods that specifically account for the correlation structure have shown some promise, e.g. [20].

D. Classical Estimation Methods

Since a PET scanner collects only a finite number of measurements, it is essentially inevitable that one must also represent the radiotracer distribution $\lambda(\underline{x})$ by a finite parameterization, e.g. in terms of a set of basis functions:

$$\lambda(\underline{x}) \approx \sum_{j=1}^p \theta_j b_j(\underline{x}),$$

where $\theta = [\theta_1 \dots \theta_p]'$ is the vector of unknown coefficients that must be computed from the Y_i 's. (Typically $b_j(\underline{x})$ is just the indicator function for the j th voxel, so we will refer to θ_j as the j th pixel value hereafter.) With such a discretization, the reconstruction problem is equivalent to a parameter estimation problem. If one assumes the scatter and random contributions are assumed to be “predetermined” values s_i and r_i respectively (i.e. if they are determined separately), and if the deadtime nonlinearity is approximated by a single “known” loss factor d_i , then the measurement mean is linear in θ :

$$\bar{Y}_i(\theta) = \sum_{j=1}^p a_{ij} \theta_j + s_i + r_i$$

where

$$a_{ij} = d_i T \int p_i(\underline{x}) b_j(\underline{x}) d\underline{x}.$$

Dozens of papers have been published based on this model, most of which have ignored the d_i , r_i , and s_i terms, and have used very simple approximations for $p_i(\underline{x})$. The linear form above invites application of the two most common tools from statistical signal processing: maximum likelihood estimation and linear least-squares estimation. The linear-least squares estimate is easily written on paper:

$$\hat{\theta}_{\text{LS}} = (\mathbf{A}'\mathbf{A})^{-1} \mathbf{A}'(\underline{Y} - \underline{s} - \underline{r}),$$

but this expression is impractical for computation due to the large size of the matrix $\mathbf{A} = \{a_{ij}\}$. Furthermore, the conventional linear least-squares estimate produces negative pixel values, which are physically impossible. Both the size of \mathbf{A} and incorporation of the nonnegativity constraint necessitate iterative algorithms.

Many PET centers use the delayed-window method for correcting for random coincidences, which destroys the Poisson statistics. For such measurements, estimates based on (weighted) least-squares may be suitable [21]. (Also see [22] for more accurate approaches.) For scans that are not precorrected for randoms, the least-squares methods are suboptimal since they do not fully accommodate the Poisson distribution. (Often the number of counts per ray is sufficiently low that the central-limit theorem based Gaussian approximation to the Poisson distribution is inapplicable.) Furthermore, data-based weighted least-squares methods lead to systematic biases for low-count Poisson measurements [19, 23]. The use of the measurement log-likelihood $L(\theta)$ rather than the weighted least-squares criterion avoids these problems, where

$$L(\theta) = \sum_{i=1}^n (Y_i \log \bar{Y}_i(\theta) - \bar{Y}_i(\theta) - \log Y_i!).$$

Unfortunately, there is no closed-form expression for the estimate $\hat{\theta}_{\text{ML}}$ that maximizes the likelihood, which again necessitates iterative algorithms. (Unfortunately, such algorithms require computation roughly comparable to the FBP method for *each iteration*, which has hampered their clinical acceptance.) The oldest of these algorithms (for PET) is an expectation-maximization (EM) algorithm [24], which converges very slowly to $\hat{\theta}_{\text{ML}}$. This slow convergence has not greatly diminished the popularity of the EM algorithm, however, because usually a few of the intermediate images generated during the iterations toward $\hat{\theta}_{\text{ML}}$ are more appealing than $\hat{\theta}_{\text{ML}}$ itself. (Determining which of the many iterates is the best one is nontrivial however.) The problem of determining $\lambda(\underline{x})$ from $\{Y_i\}$ is inherently ill-posed,

so, after parameterization, the problem of estimating θ from $\{Y_i\}$ is generally very ill-conditioned. Thus $\hat{\theta}_{\text{ML}}$ is usually extremely noisy [25].

Naturally, one simple way to reduce this noise is to post-smooth $\hat{\theta}_{\text{ML}}$. Such post-smoothing is a special case of the more general *method of sieves* [25]. Nevertheless, simple post-smoothing is by far the most popular version of sieves. Although post-smoothing reduces noise, the problem of slow convergence of the EM algorithm remains, and hundreds to thousands of EM iterations may be required for the post-smoothed images to converge [26]. (This problem has spawned a variety of methods for accelerating the EM algorithm, which vary in the extent to which convergence is guaranteed, see [27, 28].) Another disadvantage of the usual form of space-invariant post-smoothing is that the nonstationary measurement statistics are not incorporated.

E. Classical Regularization Methods

Another way to overcome the problems of slow convergence and to reduce the image noise is to replace the log-likelihood criterion by a penalized-likelihood objective function:

$$\hat{\theta}_{\text{PL}} = \arg \max_{\theta} \Phi(\theta) \quad \text{where} \quad \Phi(\theta) = L(\theta) - \beta R(\theta),$$

where $R(\theta)$ is a measure of image roughness. Larger values of β encourage smoother images with less noise. When first investigated for PET, the penalty function posed a computational challenge since the M-step of the EM algorithm has no closed form [29–31]. However, now there are a variety of fast algorithms (compared to EM) available for maximizing such objective functions, e.g. [21, 27, 28, 32, 33]. These algorithms converge rapidly in part because the penalty function greatly improves the conditioning of the reconstruction problem.

In the context of least-squares problems, such regularization methods date at least to the early 70's [34], so now may well be considered “classical.” The most classical penalty function simply measures the norm of the image:

$$R(\theta) = \|\theta\|^2 = \sum_{j=1}^p \theta_j^2,$$

which has its origins in ridge-regression. This simple penalty leads to images that are “squashed down” since even the DC component is penalized. For reducing noise, a more suitable penalty is to discourage neighboring pixels from having disparate values:

$$R(\theta) = \sum_{j=1}^p \sum_{k \in \mathcal{N}_j} \psi(\theta_j - \theta_k),$$

where \mathcal{N}_j is the set of pixel indices in the neighborhood of pixel j , and $\psi(t)$ is a symmetric function typically chosen to be nondecreasing for $t \geq 0$. Such penalty functions (or “priors” in the Bayesian terminology) have yielded good results in image restoration and image segmentation problems. However, in PET the nonstationary noise statistics again complicate the problem. Although $R(\theta)$ above is a shift-invariant function, recent analysis shows that images reconstructed by maximizing $\Phi(\theta)$ have nonuniform spatial resolution, due to interactions between the log-likelihood and penalty terms [35,36]. (Such effects are absent in image restoration problems with white Gaussian noise.) Although modified penalty functions have been proposed that reduce the resolution nonuniformity, these modifications cause more nonuniform noise variance [36, 37].

Another challenge in penalized-likelihood methods is choosing β . This problem is comparable to that of choosing the width of the apodizing window in FBP or the resolution of the filter used when post-smoothing ML images. However, in the latter two problems the parameter that one varies to tradeoff resolution and noise is one that is naturally related to spatial resolution, whereas β has essentially arbitrary units. Automatic or data-based methods for choosing β , e.g. [38, 39] have shown some potential, but may also be unstable in imaging problems [40].

There is also no consensus on the best choice for ψ . Quadratic penalties lead to oversmoothing, and nonquadratic penalties require additional parameter(s) that must be chosen. Nonconvex penalties cause additional problems with algorithm convergence, but have led to impressive results in image restoration problems in images with sharply defined regions [41]. However, in medical images one must take care to avoid turning smooth transitions into stair steps [42].

F. Model Errors

Nearly all papers on model-based methods for PET image reconstruction assume that the measurement model is known, particularly the “system matrix” \mathbf{A} . In practice this matrix occasionally measured, or more commonly simply computed based on an approximate geometry. In either case \mathbf{A} contains errors, and the effect of this model mismatch on $\hat{\theta}$ is poorly understood. The errors in \mathbf{A} might invite the application of the *total least-squares* (TLS) estimation method, e.g. [43]. However, TLS essentially assumes that the errors in \mathbf{A} are normally distributed, which is very questionable in PET. Furthermore, \mathbf{A} usually includes attenuation factors that are determined from separate noisy transmission scans. Understanding the effects of both deterministic and random errors in the model remains an important problem.

G. Attenuation correction

As described above, the conventional attenuation correction method in PET uses the ratio of the measurements in the blank and transmission scans. The transmission measurements can be very noisy, and with random subtraction can even take negative or zero values. This noise is usually reduced by smoothing with a space-invariant filter. However, such smoothing introduces bias and again ignores the nonstationary statistics. More accurate attenuation correction factors can be computed by first using statistical methods to reconstruct an attenuation image, while incorporating nonlinear constraints such as nonnegativity and piecewise smoothness, and then reprojecting this image along all of the lines-of-response [23, 44–46].

H. SPECT

Most of the above discussion also applies to SPECT imaging. Statistical methods are perhaps even more useful in SPECT, because the SPECT physical model is considerably further away from the Radon idealization than in PET, due to single-photon attenuation and the space-variant resolution of collimators, e.g. [47]. In fact, for SPECT cardiac studies, statistical methods are now in routine use at some centers, e.g. [48], and the EM algorithm is available commercially.

I. Computing Speed and the Future

Since computers are continually increasing in speed and memory, it might seem at first that it is only a matter of time before iterative reconstruction methods become used routinely. However, the same advances in technology that lead to faster computers also lead to bigger and harder problems! For example, although computing speed certainly has reached the point where iterative methods are clinically feasible for 2D problems, the focus is now on 3D PET where the size of \mathbf{A} is 11-15 times larger than in 2D (after exploiting symmetries). (Similar considerations apply to cone-beam SPECT, or even to parallel collimator SPECT with 3D compensation for detector response.) Thus there is continuing need for new ideas in image reconstruction algorithm development. Although some of those ideas will undoubtedly be borrowed from signal and image processing work, the physics and statistics of PET will need to be incorporated for the methods to be fully effective. Convincingly demonstrating that new methods are truly more effective than previous methods requires careful matching of the resolution or noise properties of the methods compared. The medical imaging community is generally unconvinced by the type of anecdotal single-image comparisons often found in image processing papers. There is increasing emphasis on formal statistical evaluations of different image reconstruction methods [49–51], which are also being applied to image processing [52].

J. To do (for tech report)

bias-variance for ML-SAGE post-smoothed (ala sieves) vs PML-SAGE.

REFERENCES

- [1] J. A. Fessler and J. M. Ollinger. Signal processing pitfalls in positron emission tomography. Technical Report 302, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, September 1996. Available from <http://www.eecs.umich.edu/~fessler>.
- [2] Y. Vardi, L. A. Shepp, and L. Kaufman. A statistical model for positron emission tomography. *J. Am. Stat. Ass.*, 80(389):8–37, March 1985.
- [3] A. Macovski. *Medical imaging systems*. Prentice-Hall, New Jersey, 1983.
- [4] E. J. Hoffman, S. C. Huang, M. E. Phelps, and D. E. Kuhl. Quantitation in positron emission computed tomography: 4 Effect of accidental coincidences. *J. Comp. Assisted Tomo.*, 5(3):391–400, 1981.
- [5] M. E. Daube-Witherspoon and R. E. Carson. Unified deadtime correction model for PET. *IEEE Tr. Med. Im.*, 10(2):267–275, September 1991.
- [6] G. Demoment. Image reconstruction and restoration: overview of common estimation structures and problems. *IEEE Tr. Acoust. Sp. Sig. Proc.*, 37(12):2024–36, December 1989.
- [7] A. C. Kak and M. Slaney. *Principles of computerized tomographic imaging*. IEEE Press, New York, 1988.
- [8] T. J. Hebert and S. S. Gopal. An improved filtered back-projection algorithm using pre-processing. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 2068–72, 1991.
- [9] T. J. Hebert. A union of deterministic and stochastic methods for image reconstruction. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1117–9, 1992.
- [10] M. N. Wernick and C. T. Chen. Superresolved tomography by convex projections and detector motion. *J. Opt. Soc. Am. A*, 9(9):1547–53, September 1992.
- [11] M. N. Wernick, C. M. Kao, and C. T. Chen. Evaluation of PET reconstruction with sinogram restoration and detector motion. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 3, pages 1454–7, 1993.
- [12] Z. Liang and R. E. Coleman. Restoration for detector response in high resolution PET image reconstruction. *J. Nuc. Med. (Abs. Book)*, 33(5):872, May 1992.
- [13] J. A. Fessler. Tomographic reconstruction using information weighted smoothing splines. In H H Barrett and A F Gmitro, editors, *Information Processing in Medical Im.*, volume 687 of *Lecture Notes in Computer Science*, pages 372–86. Springer Verlag, Berlin, 1993.
- [14] S. J. Riederer, N. J. Pelc, and D. A. Chelser. The noise power spectrum in computed X-ray tomography. *Phys. Med. Biol.*, 23(3):446–454, May 1978.
- [15] N. M. Alpert, D. A. Chesler, J. A. Correia, R. H. Ackerman, J. Y. Chang, S. Finklestein, S. M. Davis, G. L. Brownell, and J. M. Taveras. Estimation of the local statistical noise in emission computed tomography. *IEEE Tr. Med. Im.*, 1(2):142–146, October 1982.
- [16] M. F. Kijewski et al. The noise spectrum of CT images. *Phys. Med. Biol.*, 32(5):565–575, 1987. See letters p. 1047.
- [17] H. H. Barrett, D. W. Wilson, and B. M. W. Tsui. Noise properties of the EM algorithm: I. Theory. *Phys. Med. Biol.*, 39(5):833–46, May 1994.
- [18] D. W. Wilson, B. M. W. Tsui, and H. H. Barrett. Noise properties of the EM algorithm: II. Monte Carlo simulations. *Phys. Med. Biol.*, 39(5):847–72, May 1994.

- [19] J. A. Fessler. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Tr. Im. Proc.*, 5(3):493–506, March 1996.
- [20] F. O’Sullivan, Y. Pawitan, and D. Haynor. Reducing negativity artifacts in emission tomography: post-processing filtered backprojection solutions. *IEEE Tr. Med. Im.*, 12(4):653–63, December 1993.
- [21] J. A. Fessler. Penalized weighted least-squares image reconstruction for positron emission tomography. *IEEE Tr. Med. Im.*, 13(2):290–300, June 1994.
- [22] M. Yavuz and J. A. Fessler. Objective functions for tomographic reconstruction from randoms-precorrected PET scans. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1067–71, 1996.
- [23] J. A. Fessler. Hybrid Poisson/polynomial objective functions for tomographic image reconstruction from transmission scans. *IEEE Tr. Im. Proc.*, 4(10):1439–50, October 1995.
- [24] K. Lange and R. Carson. EM reconstruction algorithms for emission and transmission tomography. *J. Comp. Assisted Tomo.*, 8(2):306–16, April 1984.
- [25] D. L. Snyder, M. I. Miller, L. J. Thomas, and D. G. Polite. Noise and edge artifacts in maximum-likelihood reconstructions for emission tomography. *IEEE Tr. Med. Im.*, 6(3):228–38, September 1987.
- [26] T. R. Miller and J. W. Wallis. Clinically important characteristics of maximum-likelihood reconstruction. *J. Nuc. Med.*, 33(9):1678–84, September 1992.
- [27] J. A. Fessler and A. O. Hero. Space-alternating generalized expectation-maximization algorithm. *IEEE Tr. Sig. Proc.*, 42(10):2664–77, October 1994.
- [28] J. A. Fessler and A. O. Hero. Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms. *IEEE Tr. Im. Proc.*, 4(10):1417–29, October 1995.
- [29] P. J. Green. Bayesian reconstructions from emission tomography data using a modified EM algorithm. *IEEE Tr. Med. Im.*, 9(1):84–93, March 1990.
- [30] T. Hebert and R. Leahy. A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors. *IEEE Tr. Med. Im.*, 8(2):194–202, June 1989.
- [31] K. Lange. Convergence of EM image reconstruction algorithms with Gibbs smoothing. *IEEE Tr. Med. Im.*, 9(4):439–46, December 1990. Corrections, T-MI, 10:2(288), June 1991.
- [32] K. Sauer and C. Bouman. A local update strategy for iterative reconstruction from projections. *IEEE Tr. Sig. Proc.*, 41(2):534–48, February 1993.
- [33] E. U. Mumcuoglu, R. Leahy, S. R. Cherry, and Z. Zhou. Fast gradient-based methods for Bayesian reconstruction of transmission and emission PET images. *IEEE Tr. Med. Im.*, 13(3):687–701, December 1994.
- [34] A. Tikhonov and V. Arsenin. *Solution of ill-posed problems*. Wiley, New York, 1977.
- [35] J. A. Fessler and W. L. Rogers. Uniform quadratic penalties cause nonuniform image resolution (and sometimes vice versa). In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 4, pages 1915–1919, 1994.
- [36] J. A. Fessler and W. L. Rogers. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Tr. Im. Proc.*, 5(9):1346–58, September 1996.
- [37] J. A. Fessler. Resolution properties of regularized image reconstruction methods. Technical Report 297, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, August 1995.

- [38] Z. Zhou, R. M. Leahy, and E. U. Mumcuoglu. Maximum likelihood hyperparameter estimation for Gibbs priors with applications to PET. In Y. Bizais, C. Barillot, and R. D. Paola, editors, *Information Processing in Medical Im.*, pages 39–51. Kluwer, Netherlands, 1995.
- [39] Y. Pawitan and F. O’Sullivan. Data-dependent bandwidth selection for emission computed tomography reconstruction. *IEEE Tr. Med. Im.*, 12(2):167–172, June 1993.
- [40] J. W. Hilgers and W. R. Reynolds. Instabilities in the optimal regularization parameter relating to image recovery problems. *J. Opt. Soc. Am. A*, 9(8):1273–1279, August 1992.
- [41] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and Bayesian restoration of images. *IEEE Tr. Patt. Anal. Mach. Int.*, 6(6):721–41, November 1984.
- [42] S-J. Lee, A. Rangarajan, and G. Gindi. Bayesian image reconstruction in SPECT using higher order mechanical models as priors. *IEEE Tr. Med. Im.*, 14(4):669–80, December 1995.
- [43] V Z Mesarović, N. P. Galatsanos, and A. K. Katsaggelos. Regularized constrained total least squares image restoration. *IEEE Tr. Im. Proc.*, 4(8):1096–109, August 1995.
- [44] J. M. Ollinger. Maximum likelihood reconstruction of transmission images in emission computed tomography via the EM algorithm. *IEEE Tr. Med. Im.*, 13(1):89–101, March 1994.
- [45] C. A. Bouman and K. Sauer. A unified approach to statistical tomography using coordinate descent optimization. *IEEE Tr. Im. Proc.*, 5(3):480–92, March 1996.
- [46] J. A. Fessler, E. P. Ficaro, N. H. Clinthorne, and K. Lange. Grouped-coordinate ascent algorithms for penalized-likelihood transmission image reconstruction. *IEEE Tr. Med. Im.*, 16(2):166–75, April 1997.
- [47] S. H. Manglos, F. D. Thomas, and R. B. Capone. Attenuation compensation of cone beam SPECT images using maximum likelihood reconstruction. *IEEE Tr. Med. Im.*, 10(1):66–73, March 1991.
- [48] E. P. Ficaro, J. A. Fessler, P. D. Shreve, J. N. Kritzman, P. A. Rose, and J. R. Corbett. Simultaneous transmission/emission myocardial perfusion tomography: Diagnostic accuracy of attenuation-corrected 99m-Tc-Sestamibi SPECT. *Circulation*, 93(3):463–73, February 1996.
- [49] C. E. Metz. ROC methodology in radiological imaging. *Invest. Radiol.*, 21(9):720–733, September 1986.
- [50] H. H. Barrett, J. L. Denny, R. F. Wagner, and K. J. Myers. Objective assessment of image quality. II. Fisher information, Fourier crosstalk, and figures of merit for task performance. *J. Opt. Soc. Am. A*, 12(5):834–52, May 1995.
- [51] A. O. Hero, J. A. Fessler, and M. Usman. Exploring estimator bias-variance tradeoffs using the uniform CR bound. *IEEE Tr. Sig. Proc.*, 44(8):2026–41, August 1996.
- [52] T. Kanungo, M. Y. Jaisimha, J. Palmer, and R. M. Haralick. A methodology for quantitative performance evaluation of detection algorithms. *IEEE Tr. Im. Proc.*, 4(12):1667–1674, December 1995.