

Tomographic Reconstruction using Information-Weighted Spline Smoothing

Jeffrey A. Fessler

Division of Nuclear Medicine, University of Michigan
Ann Arbor, MI 48109-0552

January 22, 2002

Abstract

The conventional method for tomographic image reconstruction, convolution backprojection (CBP), attempts to reduce the effects of measurement noise by radial smoothing with a spatially-invariant filter. Spatially-invariant smoothing is suboptimal when the measurement statistics are nonstationary, and often leads to a choice between oversmoothing or streak artifacts. In this paper, we describe a nonstationary sinogram smoothing method that accounts for the relative variances between different detector measurements and for the finite width of tomographic detectors. The method is based on an information-weighted smoothing spline, where the weights are determined from the calibration factors and from the measurements themselves. This weighting diminishes the influence of high variance measurements, such as detectors with relatively poor efficiency, which is shown to reduce streak artifacts. Simulations of emission and transmission tomography applications demonstrate qualitatively improved image noise structure and quantitative improvements in the tradeoffs between bias and variance.

I. INTRODUCTION

The CBP method for tomographic reconstruction is derived from a mathematical idealization of tomographic imaging without consideration of statistical measurement errors. This idealization leads to the well-known ramp filter [1], whose frequency response has the unfortunate effect of amplifying high-frequency measurement noise. The conventional approach to reducing the effects of measurement noise is to apodize the ramp filter with a window function that attenuates the high frequencies [1]. This windowing is equivalent to radially smoothing the projection measurements with a *spatially-invariant* low-pass filter. In high resolution tomographs, such as positron emission tomography (PET) systems based on block detectors, large variations in detector efficiency lead to nonstationary variances. The equality of Poisson mean and variance also leads to an object-dependent variance nonuniformity. Spatially-invariant smoothing is suboptimal when the measurement statistics are nonstationary, since it treats all measurements equally. In this paper, we describe an information-weighted spline smoothing approach that accounts for the relative variances of different detectors.

When applying spatially-invariant filters to projection data with nonuniform variance, one faces two unattractive alternatives: oversmooth most of the measurements so as to minimize the noise due to the more variable data, or undersmooth the more variable data. The latter option often leads to streak artifacts in the reconstructed images, due to the backprojection of noise that has been amplified by the ramp filter. (See Fig. 8 for an example.) To remove these streak artifacts after reconstruction, Abidi and Davis [2] proposed a filtering method based on transforming the reconstructed images from rectangular to polar coordinates. Sauer and Liu [3] describe a related method for images composed of objects with known location, extent, and density. A fundamental difference between these approaches and the spline method described below is that we perform nonstationary smoothing of the projection data *before* ramp-filtering and backprojecting. This approach has the advantage of “nipping in the bud” potential streaks before they are created. It is also computationally less demanding since only one-dimensional operations are required.

The method we describe is not optimal; a more principled approach to remedying the statistical shortcomings of CBP is iterative image reconstruction based on statistical models. In fact, our motivation for this study developed from experiences with penalized weighted least-squares reconstruction [4]. In that work we have compared regularized least-squares reconstruction with and without information weighting. The noise properties of information-weighted

reconstructed images appeared superior to those of CBP, but the noise properties of unweighted reconstructions were no better than those of CBP [4]. This suggests that accounting for the measurement statistics, or in particular the second moments of those statistics, is central to the improvements in iterative methods. The spline method described below is an attempt to capture part of the benefits of statistical modeling in a non-iterative algorithm.

Since computational considerations seem to still inhibit routine use of fully iterative methods, several investigators have turned to partially iterative methods in which the projection measurements are first iteratively “restored” in sinogram space, and then CBP is applied to form the image [5–7]. These methods are particularly useful for systems with spatially-variant detector response. Such restoration will require at least a partial deconvolution of the system response, which leads to noise amplification. In this paper we focus on applications with poor measurement statistics, for which noise amplification will generally be unacceptable. Therefore, our goal is somewhat more modest than attempting to restore lost frequency components. We will accept reconstructed images whose resolution is somewhat below the intrinsic system resolution, but as the image resolution decreases, the aim is to achieve greater variance reductions than are yielded by ordinary CBP.

Section 2 describes the theory of the method, and the nonstationary characteristics of the method are illustrated in Section 3. Section 4 summarizes the simulation results. Section 5 outlines some future directions for this study. The details of the computation of the spline coefficients are described in the Appendix.

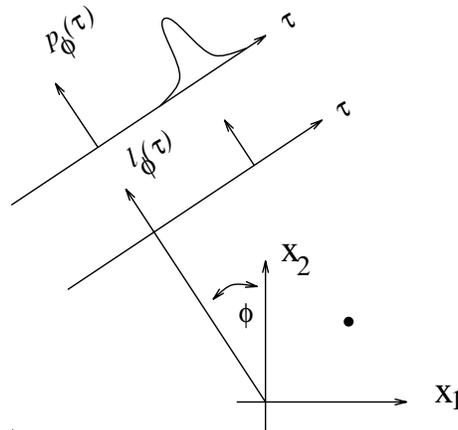


Figure 1: Idealized 2D projection geometry showing the ideal response $l_\phi(\tau)$ and blurred response $p_\phi(\tau)$ to a point source as a function of angle ϕ and radial position τ .

II. THEORY

A. Projection Model

Let $g(x_1, x_2)$ denote the object distribution being imaged, restricted to two dimensions for simplicity. The ideal line-integral projection of this object at an angle ϕ and radial offset τ is given by

$$l_\phi(\tau) = \int g(\tau \cos \phi - t \sin \phi, \tau \sin \phi + t \cos \phi) dt,$$

as shown in Fig. 1. Assume that the tomographic system has a detector response that is approximately depth independent, and for the remainder drop the dependence on ϕ . The mean response of the i th detector pair is approximately:

$$p_i = L_i l$$

where

$$L_i l = \int h_i(\tau) l(\tau) d\tau \quad (1)$$

and $h_i(\tau)$ is the line response of the detector pair.

Actual detector measurements will fluctuate around the ideal value p_i according to a statistical model that depends on the imaging modality. For PET, an approximate model for the number of coincidences measured by the i th detector pair is:

$$y_i \sim \text{Poisson} \{c_i p_i\},$$

where c_i is product of the effects of attenuation, detector efficiency, scan time, and deadtime losses. We refer to $\{c_i\}$ as “calibration factors” since they are determined through normalization scans, blank scans, transmission scans, etc. For simplicity we assume the c_i ’s are known¹. For transmission tomography, an approximate model is:

$$y_i \sim \text{Poisson} \{c_i \exp(-p_i)\},$$

where c_i is calibrated by a blank scan and is also assumed known. The next section describes a method for estimating $l(\tau)$ from the measurements $\{y_i\}$. After estimating $l - \phi$ for each projection angle ϕ , we then apply ramp-filtered CBP to reconstruct an estimate of the image g .

B. Objective Function

In the interest of simplicity, we seek a quadratic objective function, so the first step is to transform and precorrect the measurements. For emission tomography, let $z_i = y_i/c_i$ be the corrected sinogram measurement. Then z_i is an unbiased estimate of p_i , and the variance of z_i is $\sigma_i^2 = p_i/c_i$. Note that if c_i is small, then the variance of z_i can be very large. For transmission scans, let $z_i = \log(c_i) - \log(y_i + 1/4)$, then z_i is a nearly unbiased estimate of p_i , and the variance of z_i is approximately $\sigma_i^2 = e^{p_i}/c_i$. Again, if c_i is small (or if p_i is large), then the variance of z_i can be large. The statistical dependence of the measurement variance on both the calibration factors $\{c_i\}$ and the projections $\{p_i\}$ is disregarded by conventional linear smoothing methods.

We would like to recover $l(\tau)$ from $\{z_i\}$. Since direct deconvolution would be amplify the noise, we include a smoothness constraints. Such smoothness constraints were also used by Hutchins *et al.* [8,9], although without variance weighting.

Smoothing splines are naturally suited to problems with second-order statistics and smoothness constraints. Therefore, we propose estimating $l(\tau)$ using the the following penalized least-squares objective function:

$$\hat{l} = \arg \min_{l \in C_m} \sum_{i=1}^n \left(\frac{1}{w_i} \right)^2 (z_i - L_i l)^2 + \beta \int |\partial^m l(\tau) / \partial \tau^m|^2 d\tau, \quad (2)$$

where C_m is the class of functions on $[\tau_1^l, \tau_n^r]$ whose m th derivatives are square integrable. This objective trades off smoothness (measured by a squared derivative) and a weighted agreement with the measurements. The smoothing parameter β controls that tradeoff, and is analogous to the cutoff frequency of a CBP filter. In fact, one can show that if the weights $\{w_i\}$ are identical, and the system is radially-invariant, then this generalize spline smoothing corresponds to a Wiener filter with Butterworth regularization of order $2m$ and half-power frequency that is a function of β [10]. In the tomographic applications of interest, the weights $\{w_i\}$ will be nonuniform, so the “impulse response” of the smoothing spline will be spatially variant, as illustrated in the next section.

Ideally, we would like to have $w_i^2 = \sigma_i^2$, in which case w_i^{-2} corresponds to the Fisher information for z_i . In practice, the variances $\{\sigma_i^2\}$ of $\{z_i\}$ are not known exactly, so we base the weights $\{w_i^2\}$ on estimates of the variance computed from the measured data and the calibration factors. For emission tomography, an unbiased variance estimate is given by $\hat{\sigma}_i^2 = y_i/c_i^2$. Since this estimate may be unreliable for small y_i , a sensible weighting is $w_i^{-2} = c_i^2 / \max\{y_i, K\}$, for some $K > 0$. For transmission tomography, a Taylor’s expansion [11] gives the following variance estimate: $1/\hat{\sigma}_i^2 = y_i$. For robustness to small projection measurements, we use the following weights: $w_i^{-2} = \max\{y_i, K\}$. These data-based weightings introduce a slight nonlinearity into the estimation. If one desired linearity, for example in order to apply a generalization of Huesman’s ROI formulae [12], then one could use weights determined only from the relative calibration factors $\{c_i\}$, and from the projection of a “prototypical object” such as a uniform disk. For very noisy data, one could conceivably improve on this weighting scheme by applying iteratively reweighted least-squares [13].

¹Since we are interested primarily in the second-order statistics, the effects of accidental coincidences, scatter subtraction, noisy blank scans, etc. could all easily be incorporated into this framework as well.

As described in the Appendix, it follows from the Euler-Lagrange formulae for the variational problem (2) that its minimizer is a spline of order $2m$. The spline has an even order because of our use of the integral model $p_i = \int_{\tau_i^-}^{\tau_i^+} p(\tau) d\tau$ in (1). In the nonparametric regression literature, a point-sample model $p_i = p(\tau_i)$ is more conventional (although see [14] for an exception). The point-sample model results in odd order splines, such as the popular cubic spline for $m = 2$. The integral approach more naturally captures the finite width of tomographic detectors.

As outlined in the Appendix, the coefficients of the smoothing spline are easily computed noniteratively using fast banded-matrix operations [15]. The computation generally increases with m , so we focus our attention on the case $m = 1$. Despite this restriction, the simulations described in Section 4 demonstrate that the smoothing spline outperforms Butterworth filters over a range of cutoff frequencies and orders.

The spline approach has two side benefits. In principle, one could replace the conventional linear interpolation associated with pixel-driven backprojection with exact interpolations of $p(\tau)$ computed from the spline coefficients. However, in high noise applications in which significant smoothing will be required, a linear interpolator will often suffice. More importantly, the spline approach directly accommodates unevenly spaced detector samples. This is desirable for circular PET systems and many nonstandard tomographic systems [16]. It is perhaps worth emphasizing in this context that we are *not* simply performing spline interpolation. The function that minimizes (2) will in general not pass through any of the measurements, except in the limit as $\beta \rightarrow 0$.

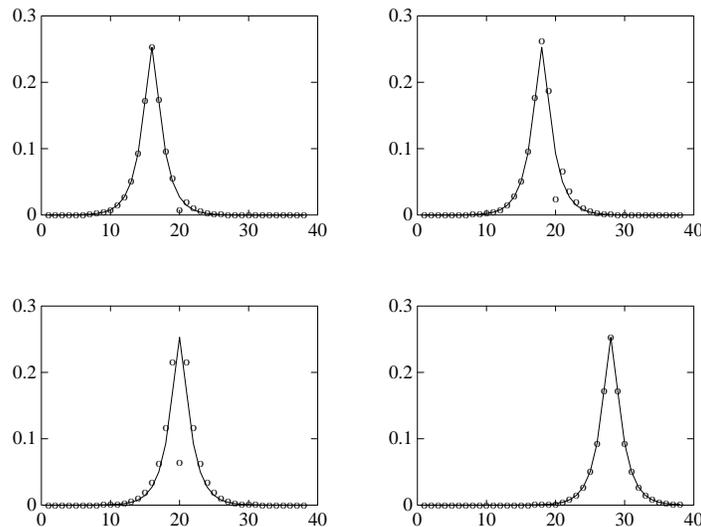


Figure 2: Kernel response functions for two sets of weights. The solid lines represent the case when all weights are unity. The open circles represent the case when all weights are unity except $w_{20}^{-2} = 0.2$. Note diminished response to the more variable detector.

III. SPATIALLY VARIANT RESPONSE KERNEL

Although information-weighted spline smoothing is nonlinear and nonstationary, the nonlinearity enters only through the dependence of the weights on the measurements. It follows from the derivation in the appendix that given β and the diagonal weight matrix \mathbf{W} , the estimate of $p(\tau)$ given the measurement vector $\mathbf{y} = [y_1, \dots, y_n]'$ is a linear function:

$$\hat{p}(\tau) = \langle \mathbf{k}(\tau; \beta, \mathbf{W}), \mathbf{y} \rangle,$$

where $\langle \cdot, \cdot \rangle$ denotes the standard inner product, and the kernel function $\mathbf{k}(\tau)$ is a vector in \mathfrak{R}^n for each τ . Therefore, by fixing the weights and examining $\mathbf{k}(\tau; \beta, \mathbf{W})$ for various τ , one can obtain some insight into the nonstationarity behavior.

Figure 2 illustrates two cases: 1) where the weights are all unity, and 2) where one weight is 20% of the others. Except near the edges, the response for uniform weights is nearly spatially invariant. In the second case the response is

nonstationary, since the contribution of the relatively more variable measurement is consistently diminished. In actual practice, when each measurement has its own weight, the kernel response will be even more variable than those in Fig. 2, with each kernel tailored to the variances of the neighboring detectors. This tailored response is the essence of our approach.

IV. SIMULATION

Both ordinary CBP and information-weighted spline smoothing include parameters that control the tradeoff between contrast and noise, or equivalently bias and variance. These parameters also influence the noise correlation structure. Therefore, a meaningful comparison requires examining both methods over a range of the parameters. In this section, we present simulation results that make qualitative and quantitative comparisons of the methods.

A. Emission reconstruction - qualitative

To illustrate the drawbacks of conventional CBP and the potential advantages of using nonstationary smoothing, we performed the preliminary comparison summarized in Figures 3 through 8. The brain-sized software phantom shown in Fig. 3 was forward projected including the effects of attenuation and nonuniform detector efficiency, and pseudo-random Poisson noise (600,000 counts) was added. The sinogram dimensions were 90 bins by 100 angles, with a 3mm bin spacing. Figure 3 also shows an unwindowed ramp-filtered CBP reconstructed image, which displays severe streak artifacts due to variations in detector efficiency. The reconstructed images shown in Figs. 4 and 5 were computed using CBP with a 3rd order Butterworth apodizing window for various cutoff frequencies. The streak artifacts persist even when one uses a low cutoff frequency.

Figures 6 and 7 show images that were reconstructed using ramp-filtered CBP of sinograms that were smoothed using the proposed information-weighted spline smoothing method. The streak artifacts appear to be virtually eliminated without loss of resolution. The noise structure in these images is markedly different than the Butterworth-apodized reconstructions, but whether this noise structure actually improves any visual clinical task is an open question.

The origin of the streak artifacts is apparent from examining Fig. 8, which displays the ramp-filtered ideal sinogram, a Butterworth-apodized ramp-filtered sinogram, and a spline-smoothed ramp-filtered sinogram. The bright noise spikes in the Butterworth-apodized ramp-filtered sinogram are due to a few detectors with relatively poor efficiency (small c_i), and consequently poor precision. These spikes produce streaks when back-projected. The spline approach uses the information in the normalization scan to diminish the effects of poor detectors by weighting each measurement appropriately. Since the influence of relatively poor detectors is reduced, the streaks are diminished.

B. Transmission reconstruction - quantitative

To corroborate the above qualitative comparison, we performed the following comparison of bias and variance in a PET transmission simulation based on the attenuation distribution shown in Fig. 9. The attenuation coefficient of the two squares is twice that of the elliptical background. Although this distribution is non-biological, the total attenuation is approximately that of a human thorax. It was chosen so that the small square could be used to examine contrast and noise, and the large square could be used to assess the spatial distribution of variance.

The calibration factors c_i were drawn from a log-normal distribution with variance 0.3, which was based on an empirical fit to the distribution of detector efficiencies of an ECAT 931. Using those calibration factors, 50 Poisson realizations of transmission scans were generated, each having a mean total of 1M events. The sinogram size was 128 bins by 128 angles, and the image size was 128 by 64 pixels. Each of the 50 realizations was reconstructed by Butterworth-apodized ramp-filtered CBP for several orders and cutoff frequencies, and by spline-smoothed ramp-filtered CBP for several values of β . Typical reconstructions are shown in Fig. 9.

Figure 10 shows standard-deviation images for typical Butterworth and spline based reconstructions. Interestingly, the variance of the spline-based estimate is fairly uniform, while the variance of the Butterworth-based estimate is strongly correlated with the object density. We hypothesize that information-weighting acts as a “pre-whitening” transformation, thereby making the the variance more uniform.

We have examined the bias-variance tradeoff as a function of smoothing parameter β and Butterworth order and cutoff frequency for several regions of interest (ROI). Since the bias of a point source is directly related to resolution, we focus on the small 3x3 square on the left of Fig. 9. For each reconstructed image, we applied the following mask

centered on the small square:

$$\frac{1}{\mu_{\text{background}}} \begin{bmatrix} \frac{-1}{16} & \frac{-1}{16} & \frac{-1}{16} & \frac{-1}{16} & \frac{-1}{16} \\ \frac{-1}{16} & \frac{1}{9} & \frac{1}{9} & \frac{1}{9} & \frac{-1}{16} \\ \frac{-1}{16} & \frac{1}{9} & \frac{1}{9} & \frac{1}{9} & \frac{-1}{16} \\ \frac{-1}{16} & \frac{1}{9} & \frac{1}{9} & \frac{1}{9} & \frac{-1}{16} \\ \frac{-1}{16} & \frac{-1}{16} & \frac{-1}{16} & \frac{-1}{16} & \frac{-1}{16} \end{bmatrix},$$

where $\mu_{\text{background}}$ is the attenuation coefficient in the background. Using this mask, the computed value will decrease from 1 to 0 as the resolution decreases. The variance of this value is directly related to the image variance as well. Figure 11 shows a plot of the bias versus variance for this feature of interest for both spline smoothing and 1st and 9th order Butterworth filters. (The curves for Butterworth filters with orders higher than 2 were all indistinguishable.) Not surprisingly, at the extremes of very high variance or very large bias, there is less difference between the methods. In the typical operating range of moderate biases, the variance of the spline smoothing approach is significantly less than that of conventional Butterworth filtering. Other features such as the contrast at the edges of the larger square showed similar trends.

For this study, the ramp-filtered or Butterworth-apodized CBP required approximately 3 seconds per slice on a Kubota Pacific Titan 3000 computer. The spline smoothing computation increased the computing time per slice by less than 1 second.

V. DISCUSSION

We have described an information-weighting spline smoothing method that accounts for the relative variance of each detector measurement, in contrast to conventional CBP. The weighting diminishes the influence of relatively poor detectors, resulting in reduced streak artifacts and an improved bias-variance tradeoff in the examples studied thus far.

This preliminary study leaves open several questions that could be investigated further. A comparison to iterative methods is in order, since most iterative methods can incorporate nonnegativity constraints and spatially-variant detector response. Such a comparison would be particularly meaningful if both methods were also compared to lower bounds. We are currently investigating an approach for computing bounds that is practical for tomography [17]. Furthermore, comparisons of bias and variance only explore one dimension of an estimator's statistics. Since our images suggest that the noise structure of ordinary CBP and spline-smoothing CBP are different, the noise autocorrelations should be examined and compared.

There is also a message in this work for investigators studying iterative algorithms: *realistic computer simulations should include nonuniform detector efficiencies*. A few investigators have reported only small improvements using iterative reconstruction when compared to CBP in spatially invariant systems. Some of those conclusions might have been different if nonuniform detector efficiencies had been included.

There are a plethora of reconstruction algorithms available; where does this method fit in? For applications where the system has stationary noise, such as in "projection-reconstruction" magnetic resonance imaging, information-weighted spline smoothing will probably not be beneficial. Since no attempt is made to deconvolve a spatially-variant system response, we would expect that regularized iterative algorithms should outperform the spline method for systems with spatially-variant detector response. However, for spatially-invariant systems, it should be possible to narrow the performance difference by following the spline smoothing with a Metz-type filter that boosts some of the attenuated midrange spatial frequencies.

In conclusion, we expect that this method will be most appropriate for systems that have nearly spatially-invariant point responses but have nonstationary variances. Single photon emission tomography (SPECT) and X-ray transmission tomography exhibit nonstationary object-dependent measurement variances; in PET, as the detector crystal sizes continue to shrink in the quest for higher resolution, the detector efficiencies will become increasingly nonuniform. In such cases, information-weighted spline smoothing is a noniterative method that accounts for that nonuniformity.

VI. APPENDIX: SPLINE CALCULATIONS

In this appendix, we show that the solution to the minimization problem posed in (2) is a spline composed of polynomials of order at most $2m$, and then describe the banded matrices that are used to compute the polynomial coefficients.

We first describe the objective function slightly more precisely. In low-count transmission applications it is possible to have detector measurements with zero information, in which case the corresponding weight w_i would need to be infinite. Using a heuristically determined “large” w_i could lead to numerical instabilities, so a more appropriate approach is simply to discard the measurements with zero information. Let n denote the number of measurements with nonzero information, and permute the measurements so that the first n are those measurements. Then the objective function of interest is:

$$\Phi(f) = \sum_{i=1}^n \left(\frac{1}{w_i} \right)^2 \left(z_i - \int_{\tau_i^l}^{\tau_i^r} f(\tau) d\tau \right)^2 + \beta \int |\partial^m f(\tau)/\partial\tau^m|^2 d\tau. \quad (3)$$

In the following, we assume the measurement intervals $[\tau_i^l, \tau_i^r]$ do not overlap ($\tau_i^r \leq \tau_{i+1}^l$), but they are allowed to be non-adjacent, i.e. when ($\tau_i^r < \tau_{i+1}^l$), which adjacency accommodates discarded measurements having zero information.

By setting the Euler-Lagrange formula for the variational problem (3) to zero, one can show that:

$$\hat{f}^{(2m)}(\tau) = (-1)^m \beta^{-1} \sum_{i=1}^n \left(\frac{1}{w_i} \right)^2 \left(z_i - \int_{\tau_i^l}^{\tau_i^r} \hat{f}(t) dt \right) I_i(\tau), \quad (4)$$

where $I_i(\tau)$ is one for $\tau \in [\tau_i^l, \tau_i^r]$ and zero elsewhere. Therefore, \hat{f} is a polynomial of order $2m$ on the measurement intervals $[\tau_i^l, \tau_i^r]$, and is a polynomial of order $2m - 1$ in the gaps $[\tau_i^r, \tau_{i+1}^l]$ between intervals (if any).

In the remainder, we focus our attention on the case $m = 1$, in which case the spline is quadratic on the measurement intervals and linear in the gaps. There are a multitude of potential parameterizations for this spline; for simplicity, we adopt the following piecewise-polynomials:

$$\hat{f}(t) = \begin{cases} a_i/h_i + b_i(\tau - m_i) + c_i((\tau - \tau_i^l)^2/2 - h_i^2/6), & \tau \in [\tau_i^l, \tau_i^r] \\ A_i + B_i(\tau - \tau_i^l) & \tau \in [\tau_i^r, \tau_{i+1}^l] \end{cases}, \quad (5)$$

where $h_i = \tau_i^r - \tau_i^l$ is the interval width, and $m_i = (\tau_i^l + \tau_i^r)/2$ is the interval midpoint. The parameterization (5) is convenient because

$$\int_{\tau_i^l}^{\tau_i^r} \hat{f}(\tau) d\tau = a_i.$$

This parameterization has $3n + 2(n - 1)$ unknown coefficients. The Euler equation (4) provides $4(n - 1)$ continuity conditions, and end conditions give 2 more constraints. This leaves n degrees of freedom that are determined from the n measurements $\{z_i\}$, as described below. For tomography, a natural end condition is that \hat{f} is a constant outside of the measurement interval, which implies that $b_1 = 0$ and $b_n + c_n h_n = 0$. Continuity of \hat{f} at the right and left sides of each interval implies respectively that:

$$a_i/h_i + b_i h_i/2 + c_i h_i^2/3 = A_i$$

and

$$A_i + B_i h_i = a_{i+1}/h_{i+1} - b_{i+1} h_{i+1}/2 - c_{i+1} h_{i+1}^2/6.$$

Continuity of the derivative of \hat{f} at the right and left sides of each interval implies respectively that:

$$b_i + c_i h_i = B_i \text{ and } B_i = b_{i+1}.$$

Each of these equalities holds for $i = 1, \dots, n$. Combining these equations to eliminate A_i and B_i yields

$$c_i = (b_{i+1} - b_i)/h_i, \quad i = 1, \dots, n - 1,$$

and

$$\mathbf{Qa} = \mathbf{Tb},$$

where $\mathbf{a} = [a_1, \dots, a_n]'$, $\mathbf{b} = [b_2, \dots, b_n]'$,

$$\mathbf{Q} = \begin{bmatrix} -1 & 1 & & & \\ & \ddots & \ddots & & \\ & & -1 & 1 & \\ & & & & \end{bmatrix} \begin{bmatrix} h_1^{-1} & & & \\ & \ddots & & \\ & & & h_n^{-1} \end{bmatrix}$$

and \mathbf{T} is an $n - 1 \times n - 1$ symmetric banded matrix whose i th diagonal element is $(h_i + h_{i+1})/3 + \tau_{i+1}^l - \tau_i^r$, and whose first upper diagonal is given by $h_2/6, \dots, h_{n-1}/6$. One can also show that

$$\int |\partial \hat{f}(\tau) / \partial \tau|^2 d\tau = \mathbf{b}' \mathbf{T} \mathbf{b}.$$

Substituting these relationships into (3), our goal is to minimize

$$(\mathbf{z} - \mathbf{a})' \mathbf{W}^{-1} (\mathbf{z} - \mathbf{a}) + \beta \mathbf{b}' \mathbf{T} \mathbf{b}$$

where \mathbf{W} is diagonal with elements $\{w_i^2\}$, subject to the constraint $\mathbf{Q}\mathbf{a} = \mathbf{T}\mathbf{b}$. This constrained minimization is easily performed using a trick due to Reinsch [18], giving the following expression for the spline coefficients:

$$\hat{\mathbf{a}} = \mathbf{z} - \beta \mathbf{W} \mathbf{Q}' \hat{\mathbf{b}},$$

where

$$(\mathbf{T} + \beta \mathbf{Q} \mathbf{W} \mathbf{Q}') \hat{\mathbf{b}} = \mathbf{Q} \mathbf{y}.$$

This latter system of equations is solved using a Cholesky decomposition for banded matrices [15], which requires only $O(n)$ operations. Having solved for the coefficients of the piecewise-polynomial representation of \hat{f} , one can then evaluate \hat{f} at any point τ using (5).

VII. ACKNOWLEDGEMENT

The author gratefully acknowledges helpful discussions with Neal Clinthorne, Les Rogers, and Al Hero, and the support of a DOE Alexander Hollaender Postdoctoral Fellowship and DOE Grant DE-FG02-87ER65061.



Figure 3: Left: ideal emission distribution. Right: ramp-filtered CBP reconstruction from 600K event scan.



Figure 4: 3rd order Butterworth apodization with cutoff frequency 100% of Nyquist (left) and 70% of Nyquist (right).



Figure 5: As above but with cutoff frequency 50% of Nyquist (left) and 30% of Nyquist (right). Note that even with a cutoff frequency of .15 cycles per detector, there are still prominent streak artifacts.

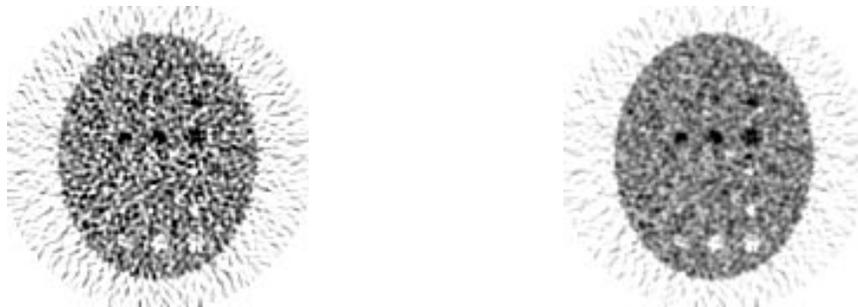


Figure 6: Ramp-filtered CBP reconstruction after information-weighted spline smoothing of the sinogram. Left: $\beta = e^{-10}$, right: $\beta = e^{-8}$.



Figure 7: As above but with larger smoothing parameters: left: $\beta = e^{-6}$, right: $\beta = e^{-5}$. Note the absence of streak artifacts and the different noise structure.

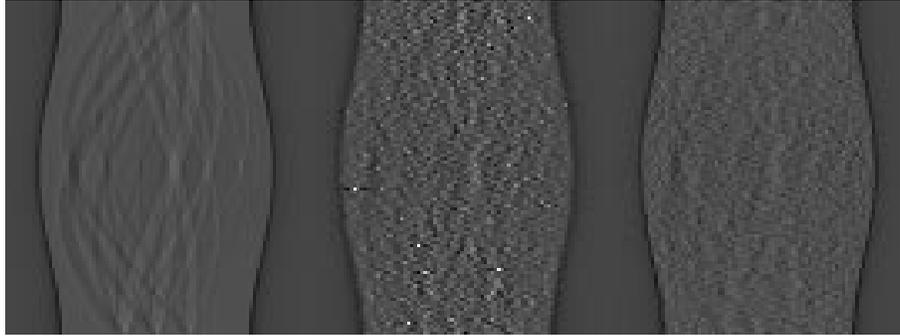


Figure 8: Ramp-filtered sinograms immediately prior to backprojection. Left: ramp-filtered ideal noiseless sinogram, middle: Butterworth-apodized ramp-filtered noisy sinogram, right: spline-smoothed ramp-filtered noisy sinogram. The bright spikes in the middle sinogram are created when the ramp filter amplifies high-variance noise in poor efficiency detectors.



Figure 9: Typical attenuation image reconstructions, left: 3rd order Butterworth with cutoff frequency 40% of Nyquist, right: spline-smoothed with $\beta = e^2$.

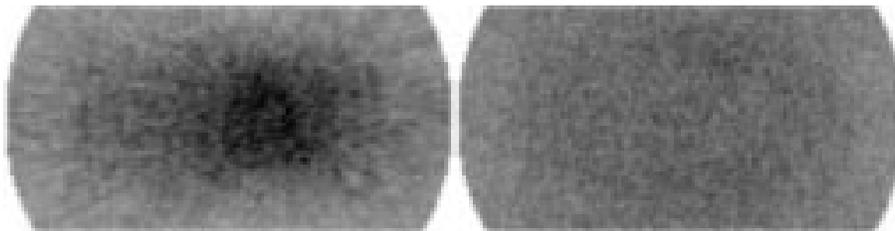


Figure 10: Pixel-by-pixel standard deviation images, left: 3rd order Butterworth with cutoff frequency 60% of Nyquist, right: spline-smoothed with $\beta = e^1$. The variance of the spline-smoothed estimate is more uniform than the Butterworth estimates.

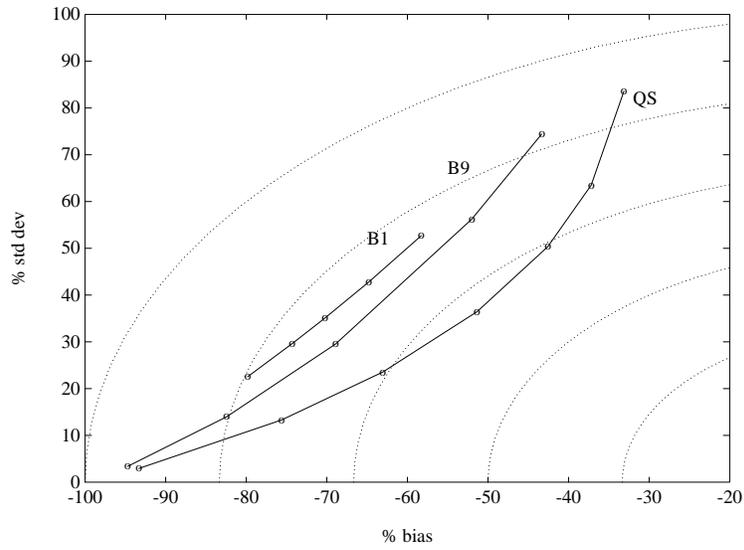


Figure 11: Plot of bias versus standard deviation in estimating the contrast of the small ROI (see text). B1 and B9 are respectively with 1st and 9th order Butterworth filters of various cutoff frequencies. QS is with spline-smoothing with $m = 1$ and various β . The dotted lines represent contours of constant RMS error; closer to the lower right hand corner means lower RMS error.

REFERENCES

- [1] A C Kak and M Slaney. *Principles of computerized tomographic imaging*. IEEE Press, New York, 1988.
- [2] M A Abidi and P B Davis. Radial noise filtering in positron emission tomography. *Optical Engineering*, 29(5):567–74, May 1990.
- [3] K Sauer and B Liu. Nonstationary filtering of transmission tomograms in high photon counting noise. *IEEE Tr. Med. Im.*, 10(3):445–52, September 1991.
- [4] J A Fessler. Penalized weighted least-squares image reconstruction for positron emission tomography. *IEEE Tr. Med. Im.*, 13(2):290–300, June 1994.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] G D Hutchins, W L Rogers, N H Clinthorne, R A Koeppe, and R D Hichwa. Constrained least-squares projection filtering: a new method for the reconstruction of emission computed tomographic images. *IEEE Tr. Nuc. Sci.*, 34(1):379–83, February 1987.
- [9] G D Hutchins, W L Rogers, P Chiao, R Raylman, and B W Murphy. Constrained least-squares projection filtering in high resolution PET and SPECT imaging. *IEEE Tr. Nuc. Sci.*, 37(2):647–51, April 1990.
- [10] G Wahba. *Spline models for observational data*. SIAM CBMS-NSF, Philadelphia, 1990.
- [11] K Sauer and C Bouman. A local update strategy for iterative reconstruction from projections. *IEEE Tr. Sig. Proc.*, 41(2):534–48, February 1993.
- [12] R H Huesman. A new fast algorithm for the evaluation of regions of interest and statistical uncertainty in computed tomography. *Phys. Med. Biol.*, 29(5):543–552, 1984.
- [13] P J Green. Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives. *J. Royal Stat. Soc. Ser. B*, 46(2):149–192, 1984.
- [14] P M Anselone and P J Laurent. A general method for the construction of interpolating or smoothing spline-functions. *Numerische Mathematik*, 12:66–82, 1968.
- [15] J A Fessler. Nonparametric fixed-interval smoothing with vector splines. *IEEE Tr. Sig. Proc.*, 39(24):852–9, April 1991.
- [16] W L Rogers, N H Clinthorne, L Shao, P Chiao, Y Ding, J A Stamos, and K F Koral. SPRINT II: A second generation single photon ring tomograph. *IEEE Tr. Med. Im.*, 7(4):291–7, December 1988.
- [17] A O Hero, J A Fessler, and W L Rogers. A fast recursive algorithm for computing CR-type bounds for image reconstruction problems. In *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, volume 2, pages 1188–1190, 1992.
- [18] C H Reinsch. Smoothing by spline functions. *Numerische Mathematik*, 10:177–183, 1967.