

Unified Reconstruction-Segmentation for PET Attenuation Correction

Jeffrey A. Fessler*

Division of Nuclear Medicine, University of Michigan 3480 Kresge III, Box 0552

University of Michigan Medical Center

Ann Arbor, MI 48109-0552

email: fessler@umich.edu

(734) 763-1434

(Originally prepared circa 1992)

July 2, 2003

ABSTRACT

We describe and compare two hybrid measured/calculated methods for attenuation correction in positron-emission tomography (PET). Both are iterative estimation methods based on penalized least-squares objective functions. One method is sequential: first reconstruct the attenuation map, then segment it. The other is a unified reconstruction/segmentation method. Simulations demonstrate that both methods can significantly reduce PET transmission scan time, as well as nearly eliminate the additional emission image variance typically introduced by noisy attenuation correction factors (ACF). The unified method is shown to be superior to the sequential method for very low count transmission scans.

When applied to a 1M event simulated transmission scan, the unified method reduced the ACF variance contribution in a 1M event simulated FDG thorax emission scan from 90% to 1%.

The sequential method and conventional linear smoothing of the transmission scan only reduced the contribution to 33% and 45% respectively. In other words, the unified method nearly eliminates the noise contribution due to attenuation correction, whereas the sequential and linear-smooth methods leave a substantial ACF component. The algorithms typically converge in about 10 iterations, making the method practical.

I. INTRODUCTION

Correction for attenuation in PET is essential for both quantitative and visual tasks. In principle, attenuation correction factors (ACFs) computed from the ratio of a blank scan to a patient transmission scan should accurately com-

pensate for attenuation. In current practice, however, measurement noise in both scans produce noisy ACFs. This noise propagates additional variance into the reconstructed emission images. The ACF contribution to emission image noise is often very substantial.

The conventional solution to this noise problem is to acquire disproportionately lengthy transmission scans and to smooth the blank and transmission scans [1]. Indiscriminant smoothing can produce systematic errors in the emission image [2].

statistical and structural image segmentation methods

almost exclusively applied in image domain

statistical methods naturally generalize to projection domain

Nevertheless, some sort of smoothing is both necessary and desirable. To illustrate, a high-resolution system such as a CTI 931-08/12 PET system has about 50,000 detectors. The transmission coincidence events are divided among this large number of detectors, so the relative accuracy of each measurement is low. Since it is unlikely that the attenuation map requires 50,000 degrees of freedom to describe adequately, there must be redundancy in the measurements. Linear smoothing is one simple but suboptimal attempt to exploit this redundancy. The reconstruction-reprojection method [3] is a somewhat more sophisticated approach that also reduces the degrees of freedom. (If the attenuation map is reconstructed within a circle in a 128^2 image, then there can be at most about 13,000 degrees of freedom in the reprojected ACFs.) One hopes to reduce the noise by eliminating degrees of freedom without inducing the systematic biases inherent to linear smoothing.

The two methods proposed in this paper reduce the degrees of freedom by exploiting two properties of attenuation maps: (i) they are composed of a relatively small number of fairly homogeneous tissue classes, such as air,

*This work was supported in part by a DOE Alexander Hollaender Postdoctoral Fellowship.

lung, soft tissue, and bone, and (ii) except at boundaries between tissues, neighboring voxels tend to be composed of the same tissue class.

These properties have been used implicitly and explicitly by several investigators, almost exclusively through a sequential three-step process: (1) reconstruct using FBP an attenuation map from the logarithm of the ACFs, (2) process that attenuation map, and (3) reproject the attenuation map to form new ACFs. Huang *et al.* [4] laid the groundwork for these methods by demonstrating the efficacy of segmenting an attenuation map into discrete classes. However, the manual adjustments of that segmentation method are impractical for routine use.

More recent attempts to automate the segmentation [5] still suffer from the fundamental limitation of such sequential approaches: FBP produces streak artifacts when applied to low-count transmission data. Therefore, we propose a unified reconstruction/segmentation method that iteratively estimates a segmented attenuation map directly from the transmission data. This method makes better use of the statistical information in the transmission measurements, and therefore can significantly reduce PET transmission scan time, as well as nearly eliminate the additional emission image variance typically introduced by noisy attenuation correction factors (ACF).

II. THEORY

A statistical approach to image reconstruction requires five components: (i) a finite parameterization of the object (attenuation map), (ii) a system model that relates the attenuation map to ideal measurement values, (iii) a statistical model that describes how the actual measurements vary about their ideal values, (iv) an objective function that is to be maximized to estimate the attenuation map, and (v) an algorithm, typically iterative, for maximizing the objective function, including specification of the initial estimate and stopping criterion. This section describes the unified reconstruction/segmentation method in terms of these five components.

A. Object Parameterization

Let $\mu(\underline{x})$ denote the spatial distribution of attenuation coefficients within the patient. We assume this distribution can be approximately decomposed into rectangular voxels:

$$\mu(\underline{x}) \approx \sum_j \mu_j I_j(\underline{x}), \quad (1)$$

where μ_j denotes the mean attenuation coefficient in the j th voxel, and $I_j(\underline{x})$ is the indicator function with the j th voxel as its support [6]. For this approximation to be

valid, the voxels must be sufficiently small. A natural generalization of this parameterization to account for inhomogeneity within voxels would be to adopt a ‘‘mixel’’ model [7, 8].

B. System Model

For a transmission scan, the patient is surrounded by a ring containing a positron-emitting radioisotope. If the ring emits a pair of annihilation photons along a chord that intersects the i th detector pair, then a coincidence event occurs provided that both photons escape unabsorbed. If the detectors are small relative to attenuation map inhomogeneities, then this survival probability is approximately

$$\alpha_i \approx \exp(-l_i), \quad (2)$$

where l_i is the effective attenuation path length along the chord between the i th pair of detectors. We assume this length can be expressed:

$$l_i = \int_{S_i} \mu(\underline{x}) d\underline{x} \quad (3)$$

where S_i denotes the strip integral over the i th detector chord. Other authors have used line integrals rather than strip integrals, but the finite width of PET detectors makes strip integrals more plausible. An even more accurate system model would account for noncollinearity of the annihilation photon pairs and for inter-crystal mispositioning errors in block detectors. Note that for a circular PET geometry, the strip widths vary across the image.

Under the discretization (1), we can rewrite (3) as

$$l_i = \sum_j a_{ij} \mu_j, \quad (4)$$

where

$$a_{ij} = \int_{S_i} I_j(\underline{x}) d\underline{x}. \quad (5)$$

e precompute and store the nonzero elements of the system matrix $\mathbf{A} = \{a_{ij}\}$ using the ASPIRE software library [9], which significantly reduces the reconstruction time.

C. Statistical Model

Ideally, the statistical model describes the distribution of each measurement about its mean, and consequently determines a measure of similarity between the actual measurements and the calculated projections of image estimates from (4). Under a simplified statistical model, Poisson likelihood has been proposed as a similarity measure [10], and was maximized using an expectation-maximization (EM) [11] algorithm for PET

[12–14]. These papers ignored the effects of accidental coincidence (AC) events [15, 16]. The system¹ at our institution uses real-time subtraction of delayed-window coincidences [15, 17] to correct for AC events. This subtraction process produces negative sinogram bins that clearly invalidate the usual Poisson model. We present a more realistic Poisson-difference model below. Because the likelihood for this Poisson-difference model is intractable, we propose a simple weighted, transformed, least-squares similarity measure.

Let B_i denote the i th detector measurement for the blank scan. For a system employing real-time delayed-window AC event correction, the following statistical model is reasonable:

$$B_i \sim \text{Poisson}\{\tau^b(b_i + r_i^b)\} - \text{Poisson}\{\tau^b r_i^b\}, \quad (6)$$

where τ^b denotes the time duration of the blank scan, b_i denotes the product of detector efficiency and photon flux on the i th detector pair, and r_i^b denotes the AC event rate during the blank scan (typically very small). Similarly, after the patient is put into the PET scanner, a reasonable model for the transmission scan measurement is:

$$T_i \sim \text{Poisson}\{\tau^t(b_i \alpha_i (1 + s_i) + r_i^t)\} - \text{Poisson}\{\tau^t r_i^t\}, \quad (7)$$

where τ^t denotes the time duration of the transmission scan, α_i was defined by (2), and s_i denotes the fractional increase in direct coincidence events due to photon scatter within the patient.

Unfortunately the individual AC event rates r_i^b and r_i^t are not available, which appears to preclude using an exact likelihood as a similarity measure. For this reason, and for computational simplicity, we propose to forgo a likelihood approach in favor of a quadratic approach based on the first two moments. The result is an approximate similarity measure that nevertheless performs remarkably well.

The first step is to apply a logarithmic transformation [18]. Let $\{b_i\}$ and $\{t_i\}$ denote the measured realizations of the random variables $\{B_i\}$ and $\{T_i\}$. Then a (noisy) estimate of l_i , the strip integral of attenuation between the i th detector pair, is given by

$$y_i = \begin{cases} \log\left(\frac{b_i}{\tau^b}\right) - \log\left(\frac{t_i}{\tau^t}\right) + \log(1 + \hat{s}_i), & t_i > 0, b_i > 0 \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

where \hat{s}_i is an estimate of the scatter fraction [19, 20]. For simplicity, in the remainder of this paper we assume $s_i = 0$. This assumption is reasonable for PET systems employing orbiting rod sources with sinogram masking

[21–24]. Using a Taylor’s expansion for the logarithmic transformation of a Poisson variate, one can show that

$$\sigma_i^{-2} = \begin{cases} (b_i t_i)/(b_i + t_i), & t_i > 0, b_i > 0 \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

is an estimate of the Fisher information (1/variance) of y_i . Let $\boldsymbol{\mu} = [\mu_1, \mu_2, \dots]'$ denote the vector of unknown attenuation coefficients. We propose to use the following weighted least-squares (WLS) similarity measure:

$$\frac{1}{2}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu}) \quad (10)$$

where $\mathbf{y} = [y_1, y_2, \dots]'$, and $\boldsymbol{\Sigma}^{-1}$ is a diagonal matrix with elements σ_i^{-2} . The similarity measure proposed by Sauer [18] is similar, except that we have included the blank scan variance b_i . That variance can be non-negligible for poor efficiency detectors.

D. Objective Function

Although one could minimize (10) to estimate $\boldsymbol{\mu}$ from the transformed data \mathbf{y} , it is well known that objective functions based solely on similarity measures such as (10) perform poorly due to the ill-conditioned nature of tomographic reconstruction. To remedy this problem, we propose a regularization method that exploits two properties of attenuation maps: (i) the attenuation coefficients can be grouped into discrete classes, e.g. air, lung, water, bone, and (ii) neighboring voxels tend to be of the same class.

To exploit the first property, we restrict the domain of $\boldsymbol{\mu}$ by reparameterizing the attenuation map. Assume that there are K classes of attenuation coefficients with nominal values $\boldsymbol{\theta} = [\theta_1, \dots, \theta_K]'$. We assume that μ_j , the attenuation coefficient in the j th voxel, is one of the K values $\{\theta_1, \dots, \theta_K\}$. Let x_j indicate the class of the j th voxel, i.e., $\mu_j = \theta_{x_j}$, where x_j takes values $1, 2, \dots, K$. An explicit notation for this parameterization is $\boldsymbol{\mu}(\mathbf{x}, \boldsymbol{\theta})$ where

$$\mu_j(\mathbf{x}, \boldsymbol{\theta}) = \theta_{x_j}. \quad (11)$$

We assume K is known, although information theoretic approaches can in principle be used to determine K [25, 26].

To exploit the second property, we use a penalty function that encourages neighboring pixels to be of the same class. Specifically,

$$R_1(\mathbf{x}) = \frac{1}{2} \sum_j \sum_{k \in N_j} w_{jk} 1_{\{x_j \neq x_k\}}, \quad (12)$$

where N_j is the set of eight neighbors of the j pixel. The weights w_{jk} equal 1 for horizontal and vertical neighbors,

¹CTI ECAT 931-08, see [17]

and $1/\sqrt{2}$ for diagonal neighbors. This type of penalty function is often presented as a Bayesian “prior” for the ensemble of voxel classes [27]. We do not adopt that philosophy here since we have no evidence that the Gibb’s distribution corresponding to (12) resembles the ensemble statistics of attenuation maps.

Depending on the scatter fraction and the accuracy of the scatter correction method, the attenuation coefficients of the K classes may be known exactly, or their effective values may vary somewhat from the “narrow beam” case. In the latter case, it is desirable to also estimate θ from \mathbf{y} . A penalty function for this parameter estimation is also useful:

$$R_2(\theta) = \frac{1}{2} \sum_{k=1}^K w_k^2 (\theta_k - \bar{\theta}_k)^2,$$

where $\bar{\theta}_k$ denotes the ideal attenuation coefficients and w_k denotes weights that reflect their uncertainties.

Combining the WLS similarity measure with the above penalty functions yields the following *penalized, weighted least-squares* (PWLS) objective function for the unified method:

$$\begin{aligned} \Phi_u(\mathbf{x}, \theta) = & \frac{1}{2} (\mathbf{y} - \mathbf{A}\boldsymbol{\mu}(\mathbf{x}; \theta))' \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \mathbf{A}\boldsymbol{\mu}(\mathbf{x}; \theta)) \\ & + \beta R_1(\mathbf{x}) + R_2(\theta), \end{aligned} \quad (13)$$

where β controls the influence of the smoothness penalty, in analogy with the filter window that must be chosen for FBP reconstruction.

Having defined this objective, our goal is to estimate \mathbf{x} and θ from \mathbf{y} :

$$(\hat{\mathbf{x}}, \hat{\theta}) = \arg \min_{\mathbf{x}, \theta} \Phi_u(\mathbf{x}, \theta).$$

E. Iterative Algorithm

The objective function Φ_u is nonconvex and nondifferentiable because of the discrete parameterization (11). Therefore, conventional gradient maximization methods are inapplicable. A natural algorithm is iterative coordinate-descent (ICD) [18]. The ICD algorithm updates each image parameter individually by minimizing the objective function (13) over that parameter while holding the other parameters fixed. The method is closely related to the Gauss-Siedel algorithm for differential equations [28, 29]. (If we considered (12) to be a prior, then ICD would be equivalent to ICM [27]). Since the number of classes is small, the objective is evaluated for each possible class, and the minimizing class chosen. In the event of a tie, the most physically abundant class is chosen. One iteration consists of updating every voxel value in some

sequence. If θ is unknown, we alternate between updating \mathbf{x} and θ .

Although a derivation for the case where θ is known is given in [18], we summarize the algorithm here for completeness. Let $\hat{\mathbf{x}}$ denote the current estimate of \mathbf{x} , and let \mathbf{a}_j denote the j th column of \mathbf{A} . The PWLS-ICD algorithm for known θ is as follows.

Initialization:

$$\begin{aligned} \hat{\mathbf{f}} & := \text{FBP}\{\mathbf{y}\} \\ \hat{x}_j & := \arg \min_k \|f_j - \theta_k\| \end{aligned} \quad (14)$$

$$\begin{aligned} \hat{\mathbf{r}} & := \mathbf{y} - \mathbf{A}\boldsymbol{\mu}(\hat{\mathbf{x}}, \theta) \\ s_j & := \mathbf{a}_j' \boldsymbol{\Sigma}^{-1} \mathbf{a}_j, \forall j. \end{aligned} \quad (15)$$

For each j :

begin

$$k^{\text{old}} := \hat{x}_j$$

$$\delta_k = \theta_k - \theta_{k^{\text{old}}}, k = 1, \dots, K$$

$$\hat{x}_j := \arg \min_k \frac{1}{2} s_j \delta_k^2 - \mathbf{a}_j' \boldsymbol{\Sigma}^{-1} \hat{\mathbf{r}} \delta_k + \beta \sum_{i \in N_j} w_{ji} 1_{\{\hat{x}_i \neq k\}}$$

$$\hat{\mathbf{r}} := \hat{\mathbf{r}} - \mathbf{a}_j \delta_{\hat{x}_j}$$

end.

We have found that the convergence rate is improved by updating the image voxels in four different raster scan orderings. The sparse inner-product operations are computed using ASPIRE [9].

Since Φ_u is quadratic in the attenuation values θ , updating θ in the case where θ is estimated is performed using standard least-squares methods. Given the previous estimate for $\hat{\mathbf{x}}$, the estimate for θ is given by

$$\hat{\theta} = (\mathbf{B}' \mathbf{A}' \boldsymbol{\Sigma}^{-1} \mathbf{A} \mathbf{B} + \mathbf{W}^{-1})^{-1} (\mathbf{B}' \mathbf{A}' \boldsymbol{\Sigma}^{-1} \mathbf{y} + \mathbf{W}^{-1} \bar{\theta}), \quad (16)$$

where the indicator matrix \mathbf{B} is given by $b_{jk} = 1_{\{\hat{x}_j = k\}}$, and \mathbf{W} is diagonal with elements w_k^2 . There is a small but nonzero probability that this update may produce negative estimate for some component of $\hat{\theta}$. Such an event would be a strong indication of model mismatch in \mathbf{A} , β , $\boldsymbol{\Sigma}$, or \mathbf{W} , and is therefore reported as an “error condition” in our current implementation. For concreteness, if a negative value occurs, that component of $\hat{\theta}$ is left at its previous value, and the other components are recomputed using (16). This approach preserves the monotonicity of the algorithm.

The ICD algorithm monotonically decreases Φ_u , and since the parameter space for \mathbf{x} is discrete it will converge

in a finite number of iterations. Typically it will converge to a local minima of Φ_u , so the initialization step is important. Following Sauer’s intuitive suggestion [18], we have initialized the iteration by performing a FBP reconstruction of the attenuation map, and then classifying each voxel using the nearest-neighbor rule. If we desired complete optimality and invariance to the initial condition, an annealing approach would be required [30]. The ICD approach is considerably more practical, and appears to be adequate for our purposes here.

After the final iteration of the above algorithm, we have computed the voxel classes estimate $\hat{\mathbf{x}}$, as well as the residual $\hat{\mathbf{r}} = \mathbf{y} - \mathbf{A}\boldsymbol{\mu}(\hat{\mathbf{x}}, \boldsymbol{\theta})$. Subtracting this residual from \mathbf{y} gives the final estimate of the path lengths l_i , which after exponentiation (2) gives the desired ACFs. In other words, the “reprojection” step is intrinsic to the algorithm.

Other investigators have reported that it is desirable to smooth the reprojected ACFs to match the system response. Since our system model uses strip integrals rather than line integrals, some smoothing is inherent in the method, further smoothing seemed unnecessary.

III. SEQUENTIAL METHOD

In the sequential reconstruct-then-segment method, we first perform a FBP reconstruction of the attenuation map from \mathbf{y} . Let \mathbf{m} denote this (noisy) map. We then segment this map in image space using the following objective function:

$$\Phi_s(\mathbf{x}, \boldsymbol{\theta}) = \frac{1}{2} \|\mathbf{m} - \boldsymbol{\mu}(\mathbf{x}; \boldsymbol{\theta})\|^2 + \beta R_1(\mathbf{x}) + R_2(\boldsymbol{\theta}), \quad (17)$$

where R_1 and R_2 were defined above. This objective function is of the same basic form as (13), so the same ICD procedure applies, with \mathbf{A} and $\boldsymbol{\Sigma}$ replaced by an identity matrix. After the iterates have converged, the estimates must be reprojected ($\mathbf{A}\boldsymbol{\mu}(\hat{\mathbf{x}}, \hat{\boldsymbol{\theta}})$) and exponentiated to form the ACFs.

The essential difference (and limitation) of this approach is that it uses processed data (the reconstructed map \mathbf{m}) rather than the original measurements. The limitations of the FBP method, such as streak artifacts, reduce the accuracy of the segmentation for low count transmission scans. However, because the system matrix \mathbf{A} is not used iteratively, it uses less computation.

IV. SIMULATION

This section describes the simulation methods used to compare three methods for determining the ACFs: 1) simple linear smoothing of the transmission scan and blank

scan measurements *prior* to taking their ratio (this is the conventional approach), 2) the sequential reconstruct-then-segment approach described in Section IV, 3) the unified reconstruction/segmentation method described in Section II.

Figure 1 displays the simulated thorax attenuation map. The assigned attenuation coefficients were 0, 0.025/cm, 0.096/cm, and 0.165/cm for air, lung, soft tissue, and bone respectively. This high resolution image is 384×192 pixels, with a 1.5 mm pixel dimension. Figure 2 displays the simulated cardiac emission distribution with relative uptakes 0, 1, 2, and 4, for bone, lung, soft tissue, and myocardium respectively.

The high-resolution attenuation map was forward projected (4) to compute the effective path lengths using strip integrals appropriately spaced for the CTI 931-08. The system geometry is a 1020 mm diameter circle of 512 detectors with width approximately 6.25 mm, for which the sinogram dimension was 96×512 . We simulated an un-wobbled acquisition.

The high-resolution emission image was forward projected using the 2D joint angle between opposing pairs of detectors [31]. These projections were scaled by the nonuniform attenuation factors and by a global scale factor so that the resulting sinogram summed to 10^6 events. An emission sinogram was generated using pseudo-random Poisson-difference variates with 1.1% precorrected AC events [31].

A 32M event blank scan with about 1% AC events, was simulated using (6). The b_i ’s were generated by drawing from a uniform distribution on $[1, 10]$, which represents a typical range of detector efficiencies. Both a 3M event and 1M event transmission scan were simulated from the effective path lengths using (7), also with about 1% AC events.

A. Linear Smoothing

The blank and transmission scans were linearly smoothed using 2D Gaussian kernels with 1, 2, 3, 4, and 5 pixel FWHM. ACFs were calculated as the ratio of the smoothed scans.

B. Sequential Method

The logarithm \mathbf{y} was computed using (8) from both the 3M and 1M event transmission scans. Noisy attenuation maps were then reconstructed from these logarithms using FBP with a ramp filter onto 128×64 matrices. The ICD algorithm was applied in image space to segment these attenuation maps. The smoothing parameter β was chosen to be e^{-12} and e^{-11} for the 3M and 1M event cases, re-

spectively, by the visual appearance of the segmentation results. (Therefore the values of β chosen may not be optimal in terms of the measures evaluated below.) The final segmented attenuation maps were reprojected to compute ACFs.

C. Unified method

The logarithm \mathbf{y} and the covariance Σ were computed using (8) and (9) from both the 3M and 1M event transmission scans. The unified reconstruction/segmentation algorithm described in Section II was applied using \mathbf{y} and Σ in order to iteratively reconstruct a segmented attenuation map within a 128×64 matrix. Note that because this grid size is coarser than the simulated attenuation map, there will be partial-volume model mismatch, increasing the realism of the simulation. To save computations, only voxels within a support ellipse having short axes 64 and 32 pixels were estimated. The values for β were again visually chosen to be e^2 and e^0 for the 3M and 1M event cases respectively. The ACFs were computed from the final residual as described in Section II-E.

D. Emission reconstruction

Each of the above three methods produces its own set of ACFs. Each set of ACFs was multiplied by the noisy emission projections, and an emission image was reconstructed using FBP with a ramp filter onto a 128×64 matrix with 4.5 mm pixels. An emission image was also reconstructed using ideal ACFs 4, so that the emission component of the noise could be evaluated. Finally, an emission image was reconstructed using noise-free emission data and ideal ACFs 3; this image served as the reference image for computing errors, as described below.

E. Evaluation

In virtually all other papers on image segmentation, the primary interest is in the accuracy of the segmentation algorithm in image space. In this paper, our primary interest is the emission image, which is computed using ACFs derived from the segmented attenuation map.

Let $\hat{\lambda}_j^n$

Let λ denote the emission distribution, and $\hat{\lambda}$ denote an estimate of λ obtained by reconstructing the product of a noisy emission sinogram with an estimate \widehat{ACF} . Define the percent ACF (PACF) contribution to the emission error by

$$PAC \triangleq \frac{\|\hat{\lambda}(\widehat{ACF}) - \lambda\|^2 - \|\hat{\lambda}(\text{true ACF}) - \lambda\|^2}{\|\hat{\lambda}(\widehat{ACF}) - \lambda\|^2} \cdot 100\%.$$

This measure clearly reflects what fraction of the emission noise is due to ACF errors. Note that we are using error norm, *not variance* as other authors have. Using the variance measure neglects the systematic effects of over-smoothing.

V. RESULTS

Tables 1 and 2 report the PACF contribution for the simulated transmission scans.

When applied to a 3M event (600 seconds on CTI 931 with ring source) simulated transmission scan, the new method reduced the PAC in a 1M event simulated FDG thorax scan reconstructed by filtered back-projection from 90% to 3%. For comparison, image-domain segmentation and conventional linear smoothing of the transmission scan only reduced the contribution to 12% and 32% respectively. Thus, for 3M events, image-domain and projection-based segmentation both work adequately, and significantly outperform simple smoothing. When applied to a 1M event (200 seconds) simulated transmission scan, the new method reduced the PAC from 90% to 1%, while image-domain segmentation and transmission scan smoothing only reduced the PAC to 33% and 45% respectively. Here the image domain segmentation method has deteriorated due to the large noise in the FBP attenuation image.

VI. DISCUSSION

We have reported our results in terms of number of transmission events, but it may be more useful to translate this figure into a time duration.

In summary, even for short transmission scan times, the new unified reconstruction/segmentation method summarized by (13) can provide almost noise free ACF's from short transmission scans. This is possible due to the power of the homogeneity constraint. The ICM iterations typically converge in about 10 iterations, making the method practical for routine use.

VII. ACKNOWLEDGEMENT

The author gratefully acknowledges helpful discussions with Neal Clinthorne, Ken Koral, and Les Rogers, and thanks Gary Hutchins for providing the software phantom.

REFERENCES

- [1] M. R. Palmer, J. G. Rogers, M. Bergstrom, M. P. Beddoes, and B. D. Pate. Transmission profile filtering for positron emission tomography. *IEEE Tr. Nuc. Sci.*, 33(1):478–481, February 1986.

| Method | $\ \hat{\lambda} - \lambda\ $ | %ACF |
|-------------------|-------------------------------|------|
| Ideal ACF | 19.3 | |
| Measured ACF | 192.3 | 90% |
| Smooth FWHM=1 | 10111.8 | 99% |
| Smooth FWHM=2 | 28.8 | 33% |
| Smooth FWHM=3 | 28.9 | 33% |
| Smooth FWHM=4 | 32.4 | 40% |
| Smooth FWHM=5 | 36.5 | 47% |
| Sequential Seg. | 21.9 | 12% |
| Unified Rec./Seg. | 19.9 | 3% |

Table 1: Percent ACF contribution to emission image noise for the 3M event transmission scan.

| Method | $\ \hat{\lambda} - \lambda\ $ | %ACF |
|-------------------|-------------------------------|------|
| Ideal ACF | 19.3 | |
| Measured ACF | 186.0 | 90% |
| Smooth FWHM=2 | 72.5 | 73% |
| Smooth FWHM=3 | 35.1 | 45% |
| Smooth FWHM=4 | 35.5 | 46% |
| Smooth FWHM=5 | 38.7 | 50% |
| Sequential Seg. | 28.8 | 33% |
| Unified Rec./Seg. | 19.4 | 1% |

Table 2: Percent ACF contribution to emission image noise for the 1M event transmission scan.

- [2] S. C. Huang, E. J. Hoffman, M. E. Phelps, and D. E. Kuhl. Quantitation in positron emission computed tomography: 2 Effects of inaccurate attenuation correction. *J. Comp. Assisted Tomo.*, 3(6):804–814, December 1979.
- [3] K. Kearfott, L. Carroll, and P. Kretz. PET attenuation correction using reprojection of measured transmission data. *J. Nuc. Med. (Abs. Book)*, 25:47, May 1984.
- [4] S. C. Huang, R. E. Carson, M. E. Phelps, E. J. Hoffman, H. R. Schelbert, and D. E. Kuhl. A boundary method for attenuation correction in positron computed tomography. *J. Nuc. Med.*, 22(1):627–637, 1981.
- [5] E. Z. Xu, N. A. Mullani, K. L. Gould, and W. L. Anderson. A segmented attenuation correction for PET. *J. Nuc. Med.*, 32(1):161–, January 1991.
- [6] G. T. Herman. *Image reconstruction from projections: The fundamentals of computerized tomography*. Academic, New York, 1980.
- [7] J. T. Kent and K. V. Mardia. Spatial classification using fuzzy membership models. *IEEE Tr. Patt. Anal. Mach. Int.*, 10(5):659–71, September 1988.
- [8] H. S. Choi, D. R. Haynor, and Y. Kim. Partial volume tissue classification of multichannel magnetic resonance images—a mixel model. *IEEE Tr. Med. Im.*, 10(3):395–407, September 1991.
- [9] J. A. Fessler. ASPIRE 3.0 user’s guide: A sparse iterative reconstruction library. Technical Report 293, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, July 1995. Available from <http://www.eecs.umich.edu/~fessler>.
- [10] A. J. Rockmore and A. Macovski. A maximum likelihood approach to transmission image reconstruction from projections. *IEEE Tr. Nuc. Sci.*, 24(3):1929–35, June 1977.
- [11] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Stat. Soc. Ser. B*, 39(1):1–38, 1977.
- [12] K. Lange and R. Carson. EM reconstruction algorithms for emission and transmission tomography. *J. Comp. Assisted Tomo.*, 8(2):306–16, April 1984.

- [13] J. A. Browne and T. J. Holmes. Developments with maximum likelihood X-ray computed tomography. *IEEE Tr. Med. Im.*, 12(2):40–52, March 1992.
- [14] J. T. Kent and C. Wright. Some suggestions for transmission tomography based on the EM algorithm. In M Piccioni P Barone, A Frigessi, editor, *Stochastic Models, Statistical Methods, and Algorithms in Im. Analysis*, volume 74 of *Lecture Notes in Statistics*, pages 219–32. Springer, New York, 1992.
- [15] 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.
- [16] M. E. Casey and E. J. Hoffman. Quantitation in positron emission computed tomography: 7 a technique to reduce noise in accidental coincidence measurements and coincidence efficiency calibration. *J. Comp. Assisted Tomo.*, 10(5):845–850, 1986.
- [17] T. J. Spinks, T. Jones, M. C. Gilardi, and J. D. Heather. Physical performance of the latest generation of commercial positron scanner. *IEEE Tr. Nuc. Sci.*, 35(1):721–725, February 1988.
- [18] K. Sauer and C. Bouman. A local update strategy for iterative reconstruction from projections. *IEEE Tr. Sig. Proc.*, 41(2):534–48, February 1993.
- [19] M Bergström, L. Eriksson, C. Bohm, and J. Litton. Correction for scattered radiation in a ring detector positron camera by integral transformation of the projections. *J. Comp. Assisted Tomo.*, 7(1):42–50, February 1983.
- [20] B. Chan, M Bergström, M. R. Palmer, C. Sayre, and B. D. Pate. Scatter distribution in transmission measurements with positron emission tomography. *J. Comp. Assisted Tomo.*, 10(2):296–301, March 1986.
- [21] L. R. Carroll, P. Kretz, and G. Orcutt. The orbiting rod source: Improving performance in PET transmission correction scans. In Esser, editor, *Emission Computed Tomography: Current Trends*, pages 235–47. Society of Nuclear Medicine, New York, 1983.
- [22] C. J. Thompson et al. A technique to reject scatter radiation in PET transmission scans. In *Proc. SPIE 671, Physics and Engineering of Computerized Multidimensional Im. and Processing*, pages 244–253, 1986.
- [23] 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 Tr. Nuc. Sci.*, 35(1):735–739, February 1988.
- [24] W K Kübler, H. Ostertag, H. Hoverath, J. Doll, S. Ziegler, and W. J. Lorenz. Scatter suppression by using a rotating pin source in PET transmission measurements. *IEEE Tr. Nuc. Sci.*, 35(1):749–752, February 1988.
- [25] J. Rissanen. Modeling by shortest data description. *Automatica*, 14:465–471, 1978.
- [26] H. Akaike. A new look at the statistical model identification. *IEEE Tr. Auto. Control*, 19(6):716–723, December 1974.
- [27] J. Besag. On the statistical analysis of dirty pictures. *J. Royal Stat. Soc. Ser. B*, 48(3):259–302, 1986.
- [28] D. M. Young. *Iterative solution of large linear systems*. Academic Press, New York, 1971.
- [29] G. Gullberg and B. M. W. Tsui. Maximum entropy reconstruction with constraints: iterative algorithms for solving the primal and dual programs. In C N de Graaf and M A Viergever, editors, *Proc. Tenth Intl. Conf. on Information Processing in Medical Im.*, pages 181–200. Plenum Press, New York, 1987.
- [30] 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.
- [31] J. A. Fessler. Penalized weighted least-squares image reconstruction for positron emission tomography. *IEEE Tr. Med. Im.*, 13(2):290–300, June 1994.



Figure 1: Simulated attenuation map representing a thorax cross-section including arms, spine, lungs, and soft tissue.

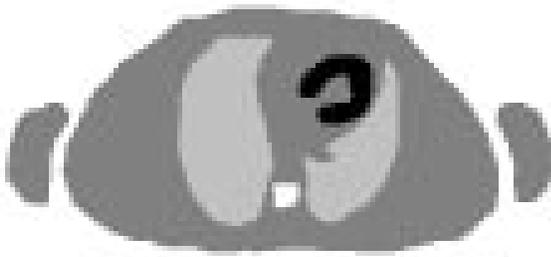


Figure 2: Simulated emission distribution.

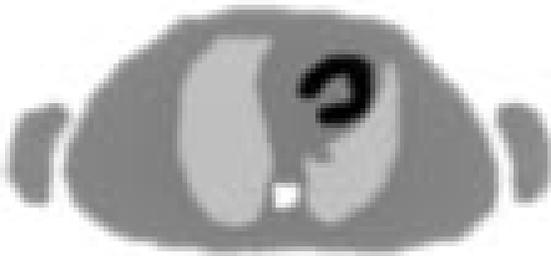


Figure 3: Emission distribution reconstructed from noise-free measurements. This image is the standard to which the reconstructions below should be compared.

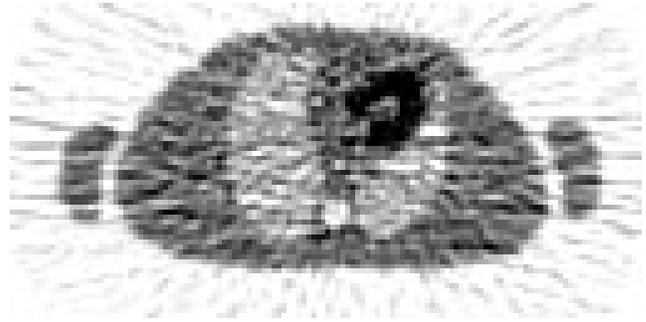


Figure 4: Reconstructed emission image (1M event) using ideal ACFs. The noise in this image is due solely to the emission measurements.

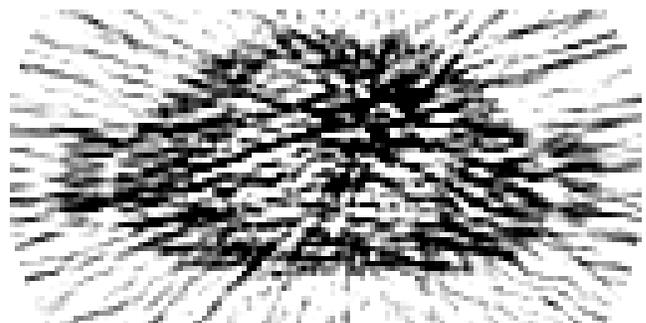


Figure 5: Reconstructed emission image (1M event) using ACFs from unprocessed measured transmission data (3M event). With no smoothing, the transmission noise dominates the image noise.

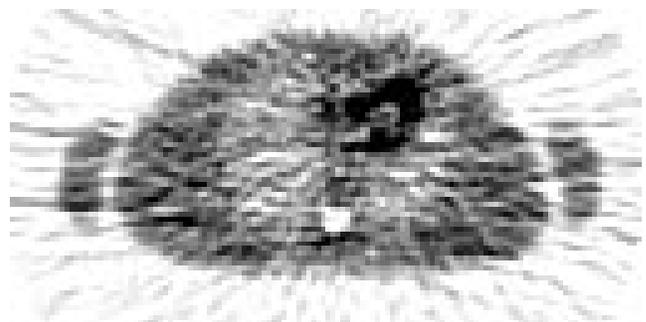


Figure 6: Reconstructed emission image (1M event) using ACFs from linearly smoothed transmission data. This smoothing introduces artifacts at the boundaries between tissue types.

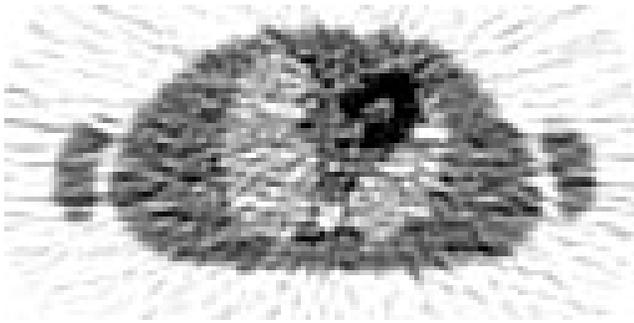


Figure 7: Reconstructed emission image (1M event) using ACFs computed by the “sequential” segmentation method.

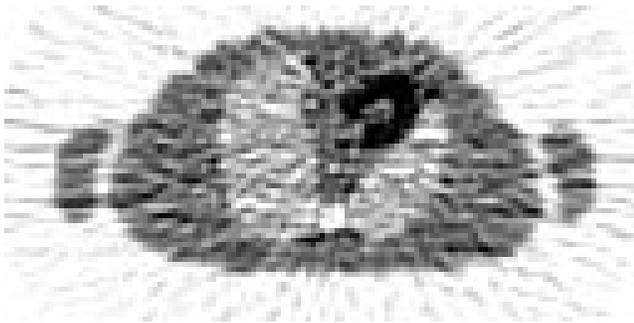


Figure 8: Reconstructed emission image (1M event) using ACFs computed by the unified reconstruction/segmentation method. The transmission component of the image noise has been virtually eliminated.



Figure 10: Segmented attenuation map computed by the unified reconstruction/segmentation method.



Figure 9: Segmented attenuation map computed by the sequential reconstruct-then-segment method.