

Continuous Time Dynamic PET Imaging Using List Mode Data

Thomas E. Nichols¹, Jinyi Qi², and Richard M. Leahy²

¹ Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213-3890, USA

nicholst@stat.cmu.edu

² Signal & Image Processing Institute, University of Southern California, Los Angeles, CA 90089-2564, USA

{jqi, leahy}@sipi.usc.edu

Abstract. We describe a method for computing a continuous time estimate of dynamic changes in tracer density using list mode PET data. The tracer density in each voxel is modeled as an inhomogeneous Poisson process whose rate function can be represented using a cubic B-spline basis. An estimate of these rate functions is obtained by maximizing the likelihood of the arrival times of each detected photon pair over the control vertices of the spline. By resorting the list mode data into a standard sinogram plus a “timogram” that retains the arrival times of each of the events, we are able to perform efficient computation that exploits the symmetry inherent in the ordered sinogram. The maximum likelihood estimator uses quadratic temporal and spatial smoothness penalties and an additional penalty term to enforce non-negativity. Corrections for scatter and randoms are described and the results of studies using simulated and human data are included.

1 Introduction

Dynamic PET imaging usually involves the collection of a series of frames of sinogram data over contiguous time intervals that can range in duration from 10 seconds to over 20 minutes. Data from each of the frames is independently reconstructed to form a set of images. These images can then be used to estimate physiological parameters [8]. This approach involves selection of the set of acquisition times, where one must choose between collecting longer scans with good counting statistics but poor temporal resolution, or shorter scans that are noisy but preserve temporal resolution. List mode data acquisition avoids this problem by allowing frame durations to be determined after acquisition. Alternatively, the problem of temporal binning can be avoided entirely by directly using the arrival times in the list mode data to estimate a dynamic image.

Snyder [23] developed a list mode maximum likelihood (ML) method for estimation of dynamic PET images using inhomogeneous Poisson processes. Each voxel has an associated time-varying tracer density that is modeled using basis functions that are based on assumptions about the physiological processes generating the data, e.g. blood activity curves convolved with a basis of exponentials.

The observed list mode PET data are then inhomogeneous Poisson processes whose rate functions are linear combinations of the dynamic voxel tracer densities. Here we follow a similar approach but instead work with rate functions formed as a linear combination of known basis functions. Not only does the linearity of the model lend itself to efficient computation of the estimates, but also we can better represent the dynamic activity seen in experimental data that is not well modeled by the more restrictive physiological models.

A second advantage of using list mode data arises in cases where the number of detected photon pairs in a particular study is far less than the total number of detector pairs. This is often the case in modern 3D PET systems which can have in excess of 10^8 sinogram elements in a single frame. To reduce this number to manageable proportions, the data are often rebinned by adding nearby elements together. Alternatively, the raw list mode data can be stored and the need for rebinning is avoided. Barrett *et al.* [1,18] describe a list mode maximum likelihood method for estimation of a temporally stationary image. While this method will often reduce storage costs and avoid the need for rebinning, the random spatial ordering of the detected events in the list mode data does not lend itself to fast forward and backprojection and exploitation of the many symmetries in 3D projection matrices [10,19]. To avoid this problem we use a hybrid combination of the standard sinogram and list mode formats that allows the reconstruction algorithm to exploit the same matrix symmetries used in our static imaging work [19]. All events in a dynamic study are collected into a single standard sinogram; this is then augmented by a “timogram” which contains the arrival times of each event stored so that they are indexed using the values in the associated sinogram.

In this paper we present a method for reconstructing a continuous time estimate of a dynamic PET image using list mode data and the theory of inhomogeneous Poisson processes. A general B-spline model is used to represent the dynamic activity in each voxel so that the dynamic image is parameterized by a sequence of control vertex “images” where the control vertices are the coefficients for the spline basis. Tomographic projections of these control vertices produce the control vertices for the rate functions of the inhomogeneous Poisson processes representing coincidence detections between each detector pair. A maximum likelihood estimate of the control vertices for each voxel can then be computed using the standard likelihood function for inhomogeneous Poisson processes [21,22]. The final result is a temporally continuous representation of the PET image that utilizes the temporal resolution of list mode data.

Our parameterization of the inhomogeneous Poisson rate function is applicable to any linear combination of basis functions. This form encompasses the parametric imaging work of Matthews [14], Snyder [21] and mixture models of O’Sullivan [17]. We also note that Ollinger [16] used list mode data to reconstruct rate functions as histograms with adaptive bin-widths; our work could be viewed as a continuous-time extension of this. For this paper we consider only cubic B-splines. The key advantage of B-splines are that they have systematic compact support. In particular, for any point on a cubic spline only 4 basis func-

tions are nonzero. Also, simple closed forms exist for all derivatives and integrals of a polynomial spline.

Since inhomogeneous Poisson rate functions are unnormalized densities, we note that the density estimation literature using splines is closely related to our work (e.g. [24,6]). The standard methods involve exponentiated splines or squared splines. While these implicitly constrain the rate function to be positive, they cannot be represented with a linear basis. As there are substantial computational savings to having a common basis for all voxels and projections, we did not pursue these approaches.

The paper is organized as follows. We describe the model and maximum likelihood method in Sections 2 and 3, respectively. Methods for selecting the spline knot points and methods for randoms and scatter correction are included in Section 4. Computational considerations including resorting data into a timogram format and the details of the algorithm used for computing the ML estimate are given in Section 5. In Section 6 we demonstrate the performance of the method with some preliminary simulation and experimental results.

2 Dynamic Modeling Using Inhomogeneous Poisson Processes

We model the positron emissions from each voxel in the volume as an inhomogeneous Poisson process. The rate function for the voxel represents, to within a scalar calibration factor, the time varying PET tracer density. We parameterize the rate functions using a cubic B-spline basis:

$$\eta_j(t) = \sum_{\ell} w_{j\ell} B_{\ell}(t), \quad \eta_j(t) \geq 0 \forall t,$$

where $\eta_j(\cdot)$ is the rate function for voxel j , $w_{j\ell}$ is the ℓ th basis weight (control vertex) for voxel j , and $B_{\ell}(t)$ is the ℓ th spline basis function. The problem of reconstructing the dynamic PET image is then reduced to estimating the control vertices for each voxel.

We denote by p_{ij} the probability of detecting at detector pair i a photon pair produced by emission of a positron from voxel j . The probabilities p_{ij} are identical to those used in static PET imaging. Here we use the factored matrix forms developed in [19]. Assuming that the detection probabilities are independent and time invariant, it follows that coincidence detection at detector pair i is also an inhomogeneous Poisson process with rate function

$$\lambda_i(t) = \sum_j p_{ij} \sum_{\ell} w_{j\ell} B_{\ell}(t) = \sum_{\ell} \left(\sum_j p_{ij} w_{j\ell} \right) B_{\ell}(t) \quad (1)$$

where the right-most term demonstrates that the rate functions for the data are also B-splines.

The Poisson process observed at the detectors is corrupted by random and scatter components that can also be modeled as inhomogeneous Poisson processes. Combining the three components, we have the model:

$$\lambda_i^*(t) = \lambda_i(t) + r_i(t) + s_i(t)$$

where $r_i(\cdot)$ and $s_i(\cdot)$ are the randoms and scatter rate functions for detector pair i and $\lambda_i^*(t)$ is the rate function for the process actually observed at detector pair i . In estimating the rate function parameters $w_{j\ell}$ we will assume that the rate functions for the random and scatter components have been determined through a calibration procedure and can be treated as known processes.

For a Poisson process with rate function $\lambda(t)$, with N events observed from time T_0 to T_1 and event arrival times $a_1, \dots, a_k, \dots, a_N$, the likelihood function is [22]

$$P(a_1, \dots, a_k, \dots, a_N | \lambda(t)) = \left(\prod_{k=1}^N \lambda(a_k) \right) \exp \left\{ - \int_{T_0}^{T_1} \lambda(u) du \right\}. \quad (2)$$

For $N = 0$, the product is defined as unity.

For the set of independent events recorded in the list mode data the log likelihood is therefore given by

$$L(\mathcal{D} | \mathcal{W}) = \sum_i \sum_k \log \lambda_i^*(a_{ik}) - \sum_i \int \lambda_i^*(u) du, \text{ s.t. } \lambda_i^*(t) \geq 0 \forall t. \quad (3)$$

where \mathcal{D} denotes the list mode data and \mathcal{W} the set of parameters for the rate functions. We represent the data as $\mathcal{D} = (\mathbf{x}, \mathbf{a}_1, \dots, \mathbf{a}_i, \dots, \mathbf{a}_I)$, where $\mathbf{x} = (x_1, \dots, x_i, \dots, x_I)$ are the sinogram count data, and $\mathbf{a}_i = (a_{i1}, \dots, a_{ik}, \dots, a_{ix_i})$, the set of x_i event arrival times at detector pair i . For the B-spline basis, $\mathcal{W} = (w_{j\ell} | \ell = 1, \dots, L, j = 1, \dots, J)$ are the set of basis coefficients. While \mathbf{x} is a function of \mathbf{a} and hence redundant, we use the sinogram counts to index the arrival times, as described in section 5.1.

3 Penalized Maximum Likelihood Estimation

We estimate the image control vertex values that define our dynamic image using penalized maximum likelihood. The objective function of the statistical model is modified with three regularizing terms

$$L^*(\mathcal{D} | \mathcal{W}) = L(\mathcal{W} | \mathcal{D}) - \alpha \rho(\mathcal{W}) - \beta \phi(\mathcal{W}) - \gamma \nu(\mathcal{W}). \quad (4)$$

The terms $\rho(\mathcal{W})$ and $\phi(\mathcal{W})$ regularize temporal and spatial roughness, respectively; $\nu(\mathcal{W})$ penalizes negativity of the image rate functions; α , β and γ are the tuning parameters. We now describe each of these terms.

We employ a temporal smoothing term to control the roughness of the spline rate functions [3]. The form of the roughness penalty is the integrated squared

curvature. For voxel j this is

$$\int \left\{ \frac{\partial^2}{\partial u^2} \eta_j(u) \right\}^2 du.$$

Fortunately, for cubic splines this quantity has a simple expression, a quadratic form of the control vertices ([3], pg. 238). We denote the symmetric, banded matrix of this quadratic form Q . Thus the temporal roughness penalty is given by

$$\rho(\mathcal{W}) = \sum_j \sum_{\ell_1} \sum_{\ell_2} w_{j\ell_1} Q_{\ell_1\ell_2} w_{j\ell_2}.$$

We regularize the estimates of the control vertices using a spatial smoothing function equivalent to the pair-wise quadratic penalty used previously in penalized ML [4] and Bayesian estimation [19] of static PET images:

$$\phi(\mathcal{W}) = \sum_{\ell} \sum_j \sum_{j' \in \mathcal{N}_j, j' > j} \kappa_{jj'} (w_{j\ell} - w_{j'\ell})^2.$$

where \mathcal{N}_j denotes a set of neighbors of voxel j and $\kappa_{jj'}$ is the reciprocal of the Euclidean distance between voxel j and j' . Other possible choices of the penalty function include the discrete approximation of the thin plate spline bending energy [12] or a non-quadratic edge preserving function such as that described in [5].

We now justify applying the same regularization to the control vertices as has previously been applied to images. First, the spline basis is the same for all voxels, so the control vertices have the same meaning for all voxels. Second, each member of the spline basis has limited support so that the effect of spatial smoothing is localized in time. Lastly, the B-spline basis we use is well conditioned [3], meaning that small changes in the control vertices produce small changes in the spline function. Hence if we expect two rate functions to be similar, then it is sufficient to constrain their control vertices to be similar.

The optimization method must account for the non-negativity of the image rate functions $\eta_j(t)$. We use unconstrained optimization with a penalty function [13]. The problem is complicated somewhat in that the control vertices themselves are not necessarily non-negative; instead we need to ensure that the corresponding spline does not become negative. The local extrema of a cubic spline have a closed form, so we initially tried to penalizing negative local minima. This approach complicated the gradient and Hessian and made their evaluation prohibitively slow.

Instead we simply penalize negative values computed at a finite number of time points. The vector \mathbf{z} contains the locations at which we enforce positivity. It is constructed by uniformly spacing d_z points in each inter-knot interval. Any elements of \mathbf{z} for which the spline is negative are penalized with the square of the spline value, resulting in the penalty:

$$\nu(\mathcal{W}) = \sum_j \sum_m \min \left(0, \sum_{\ell} w_{j\ell} B_{\ell}(z_m) \right)^2.$$

This approach does not necessarily ensure that the spline is non-negative everywhere. However, we have found that when used in combination with the temporal roughness penalty, the resulting estimates do not become negative, except possibly in the intervals just preceding a large increase in activity.

It is straightforward to show that each of the four terms in the penalized likelihood (4) have negative semi-definite Hessian matrices. Their null spaces only intersect at the zero vector. Therefore, the objective function is strictly concave and has a unique global maximum which can be found by a gradient-based search algorithm.

4 Calibration Procedures

4.1 Selection of Knot Spacing

Before proceeding to the estimation we must decide on the spacing between knots in the B-spline basis. A cubic B-spline basis is defined by knot locations, $\mathbf{u} = (u_1, \dots, u_{L+4})$, where $L \geq 4$ is the number of basis elements and the first and last 4 knots are identical, to allow discontinuity at the end points. Uniformly spaced knots will not be efficient for most tracer studies since early changes in concentration have much greater magnitude than those later in the study. While we do not attempt to adaptively place the knots, in a modest attempt to optimize knot placement, we use the head curve to define knots that produce approximately equal arc lengths, as suggested in [3]. The head curve is a temporal histogram using all of the list mode data and it serves as an estimate of the average rate function. Once the knot locations are determined, the actual basis functions are computed using the recurrence relations as described in [2,3].

4.2 Randoms and Scatter Rate Functions

To apply the penalized likelihood estimation procedure described above, we should first apply calibration procedures to account for the presence of scattered and random events in the list mode data. We note that the simple randoms subtraction method that is used in static imaging is not applicable to list mode data. While neither randoms or scatter are included in the preliminary results presented here, they are described for completeness and will be essential in extracting accurate quantitative dynamic information from our results.

The randoms rate varies approximately as the square of the true coincidence rate. We can model the randoms rate for each detector pair using an inhomogeneous Poisson process:

$$r_i(t) = \sum_{\ell} \gamma_{i\ell} B_{\ell}(t),$$

where $\gamma_{i\ell}$, $\ell = 1, \dots, L$ are the control vertices for the randoms component in the i th line of response (LOR). The list mode data produced on the ECAT HR+ (CTI Systems, Knoxville Tennessee) contains both prompt (on-time) and delayed

events. We can use the delayed events to compute an ML estimate of these control vertices. The number of counts per LOR in the delayed events is typically quite small so that these estimates would probably exhibit high variance. However, after scaling for variations in individual detector sensitivities, there is a high degree of spatial smoothness in the mean randoms sinogram [15]. Consequently we can use a penalized ML estimate in which substantial spatial smoothing is used to regularize the estimator. By choosing the knot spacing for the randoms rate functions to be at the same locations as for the image rate functions, the separate treatment of randoms in the estimation algorithm below produces little increase in the computational cost.

The spatio-temporal scatter distribution is a function of both the dynamic tracer distribution and the object. We assume no interaction between the temporal and spatial distribution and scale a fixed spatial scatter estimate over time. While this is a rather crude approximation, we anticipate that it will be reasonably accurate due to the very smooth nature of the scatter contribution to the sinogram. However, for certain ligand studies of the brain, where the tracer eventually binds solely to subcortical structures, this approximation may perform poorly.

Integrating the coincidence detections over time yields a sinogram from which we estimate the spatial scatter distribution using the simulation method in [25]. Let S_i denote the estimated scatter contribution at the i th LOR. Next we calculate a least-squares spline estimate of the head curve using the same B-spline basis of the dynamic study; we normalize this spline to integrate to unity. Denote this estimate as $h(t) = \sum_{\ell} h_{\ell} B_{\ell}(t)$ where h_{ℓ} are the control vertices of the head curve spline fit. The estimated scatter rate function is then

$$s_i(t) = S_i h(t) = S_i \sum_{\ell} h_{\ell} B_{\ell}(t).$$

Note that when computing S_i and $h(t)$ we subtract the delayed events from the prompts to correct for randoms.

5 Computational Considerations and Image Estimation

5.1 The “Timogram”

The raw list mode data is in a form that is inconvenient for computing the gradient of the penalized likelihood function. The list mode events arrive in random spatial order and hence require random rather than sequential access to the control vertices that define the rate functions in the sinogram domain. We have therefore developed a means to store list mode data in sinogram form while preserving the temporal information. This is achieved using a single standard sinogram that contains all detected events augmented by a second file listing the arrival times of all events sorted in projection order. We call this second file the “timogram”. The timogram simply consists of the arrival times of each event. The sinogram is required to indicate how many arrival times to read for

each bin. The resulting pair of files can be substantially smaller than either the original list mode data file or the set of sinograms that would be stored in a conventional dynamic study. We note that Ollinger [16] also resorted list mode data prior to reconstruction, though his format did not completely eliminate the random spatial order.

ECAT HR+ list mode data consists of a sequence of 4-byte event words, each either a coincident event or a timing event. The coincident events record the sinogram bin, optional gating information, and are identified as “prompt” or “delay”. The timing events are inserted in the list mode stream every 1 millisecond, and they also record time with a 27 bit integer. By re-encoding the arrival time of each coincidence event using 16 bits, we can retain a temporal resolution of 256ms and a maximum acquisition time of 4.6 hours. Using this format we need only 2 bytes per event in the timogram. Thus we can discard all of the timing events in the list mode data and save a factor of two in the space required to store the remaining coincidence arrival times. The space savings from discarding the timing events are significant. For example, in a 90 minute scan, the timing events take more space than a 3D sinogram set and hence the raw list mode data will always take more space than the sinogram-timogram, even if no coincidences are detected!

The sinogram-timogram format will also be more space efficient than a multi-frame sinogram when the space required to store the event arrival times in the timogram is less than the 2nd through n th sinograms. For example, an 11 frame acquisition is 10 frames larger ($\sim 200\text{MB}$ larger) than a sinogram-timogram with no events; only after 200MB-worth of events, or 100 million counts are stored will the sinogram-timogram be less space efficient.

The sinogram-timogram format could be made even more compact by storing inter-arrival times and then performing entropy-based compression [9]. The motivation for this is that LOR’s with high activity will tend to have short inter-arrival times, hence will have many high bits consistently zero, a property that compression can exploit.

5.2 Preconditioned Conjugate Gradient Based Reconstruction

A preconditioned conjugate-gradient method was used to maximize the objective function. The particular method closely follows our previous work on static reconstructions [15,19], so we only describe the method briefly here. We use the following preconditioned Polak-Ribiere form of the conjugate gradient method.

$$\begin{aligned}\mathcal{W}^{(n+1)} &= \mathcal{W}^{(n)} + \alpha^{(n)} \mathbf{s}^{(n)} \\ \mathbf{s}^{(n)} &= \mathbf{d}^{(n)} + \beta^{(n-1)} \mathbf{s}^{(n-1)} \\ \mathbf{d}^{(n)} &= \mathbf{C}^{(n)} \mathbf{g}^{(n)} \\ \beta^{(n-1)} &= \frac{(\mathbf{g}^{(n)} - \mathbf{g}^{(n-1)})' \mathbf{d}^{(n)}}{\mathbf{g}^{(n-1)}' \mathbf{d}^{(n-1)}}\end{aligned}$$

where $\mathbf{g}^{(n)}$ is the gradient vector of the penalized likelihood (4) at $\mathcal{W} = \mathcal{W}^{(n)}$, $\mathbf{C}^{(n)}$ is a preconditioner, and the step size $\alpha^{(n)}$ is found using a Newton-Raphson line search.

In this study $\mathbf{C}^{(n)}$ was chosen analogously to the static PET reconstruction [20] as

$$\mathbf{C}^{(n)} = \text{diag} \left\{ \frac{|w_{j\ell}^{(n)}| + \delta}{\sum_i p_{ij}} \right\},$$

where δ is a small positive number to ensure that $\mathbf{C}^{(n)}$ is positive definite. Here we set δ equal to $0.01 \max_{j\ell} \{w_{j\ell}^{(n)}\}$.

The algorithm was initialized with a constant image for which the forward projected rate function matches the average rate of the data after subtracting scatters and randoms. The search vector is initialized by setting $\mathbf{s}^{(0)} = \mathbf{d}^{(0)}$. At each iteration we test whether the search vector is an ascent direction, i.e. $\mathbf{g}^{(n)'} \mathbf{s}^{(n)} > 0$. If not, then we reinitialize the PCG algorithm with $\mathbf{s}^{(n)} = \mathbf{d}^{(n)}$.

The logarithm in the likelihood function requires that the line search in (5) is performed with the hard constraint that the forward projected rate function at any arrival time is non-negative, i.e.

$$\lambda_i(a_{ik}) \geq 0, \quad \forall i, k.$$

The negativity penalty in (4) is soft allowing small negative values. The hard constraint can be satisfied by limiting the step size in the update step of the conjugate gradient algorithm. To minimize the effect of this constraint on the convergence rate, we use a bent, rather than truncated, line search [11].

6 Simulation Studies and Performance Evaluation

6.1 Simulation Study

We evaluated our method with simulated and real data. We simulated a blood flow data set using a single slice of the Hoffman brain phantom. We evaluated the simulated data on the basis of instantaneous rate accuracy as described below. The real data consisted of one 2D subset of a 10 minute 3D ^{15}O -water list mode brain study. Our subjective evaluations focused on tissues that are known to have distinctly different dynamics with this tracer.

The simulated data was a simplified model of the dynamics of a bolus injection of ^{15}O -water using tissue time activity curves generated by the Kety autoradiographic model ([7], Figure 3B). We chose two extreme curves, one corresponding to very high blood flow, one to very low blood flow. White matter voxels were assigned to have low blood flow, gray matter voxels to have high blood flow. We used an 11 element B-spline basis with support from 0 to 140 seconds; the spacing of the knot locations were determined by equally spacing 7 points along a medium blood flow curve. We used 7 negativity penalty points (d_z) in each knot interval. Approximately 5 million counts were generated for this data set.

As a preliminary evaluation we computed the mean squared error (MSE) between the true source and the instantaneous rate estimate at three times, $t =$

10, 23 and 60 seconds. We compared this MSE to that obtained by estimating the instantaneous rate with a static sinogram based on events arriving in the interval $[t - d/2, t + d/2]$, for $d = 1, 2, 4, 10,$ and 20 seconds. In both cases the MSE's are based on one realization, the mean taken over voxels. While this comparison of instantaneous rate accuracy could be regarded as an unfair since the static data has no information on the nonstationarity of the tracer distribution, it is comparable to existing methods. We did not attempt to match the spatial smoothness (bias) of the two methods; for each d , the static data sets were reconstructed with an ML estimate ($\beta = 0$) with 25 iterations.

Figure 1 shows the results of the blood flow simulation for 120 iterations. The rate functions for six voxels are shown in top left; there is generally good agreement between true and estimated functions. The plot of instantaneous mean squared error is shown top right. The spline estimates (horizontal lines) have appreciably lower MSE than all static estimates with frame durations less than 10 seconds. Note that of the 20 second static estimates, the one with largest MSE occurs at $t = 23$ seconds, corresponding to the mode of the high-flow curve. This is expected since averaging across greater durations from the mode will bias the static estimate downward; at the other time points the rate function is approximately linear and there will be less bias in the static estimate.

6.2 Human Study

For the real data we used a 15 element B-spline basis with support over the whole acquisition duration, 0 to 600 seconds; knot spacing was determined by approximate equal spacing of 11 points along the head curve; again $d_z = 7$. The subject was injected with a 5 mCi (~ 200 MBq) bolus of ^{15}O -water approximately 30 seconds after the start of 3D data acquisition. To create the 2D data set we rebinned data from eight ring pairs into a single dataset with about 400,000 counts, using only the prompt events.

Figure 2 shows the results of the human study after 40 iterations. A three panel image shows the tracer distribution at 20, 60 and 120 seconds post injection. At 20 seconds the carotid arteries are visible, especially the right one; at 60 seconds the water has perfused the brain and surrounding tissue and the carotids are still visible; at 120 seconds the carotids are indistinguishable from background tissue though the brain still has increased activity. This differing temporal character is clear from the plot of selected voxels. The carotid artery shows a sharply peaked distribution, while brain tissue rises later and more smoothly; the sinus region has much lower flow though it's rate function shows a similar character to that of the brain tissue.

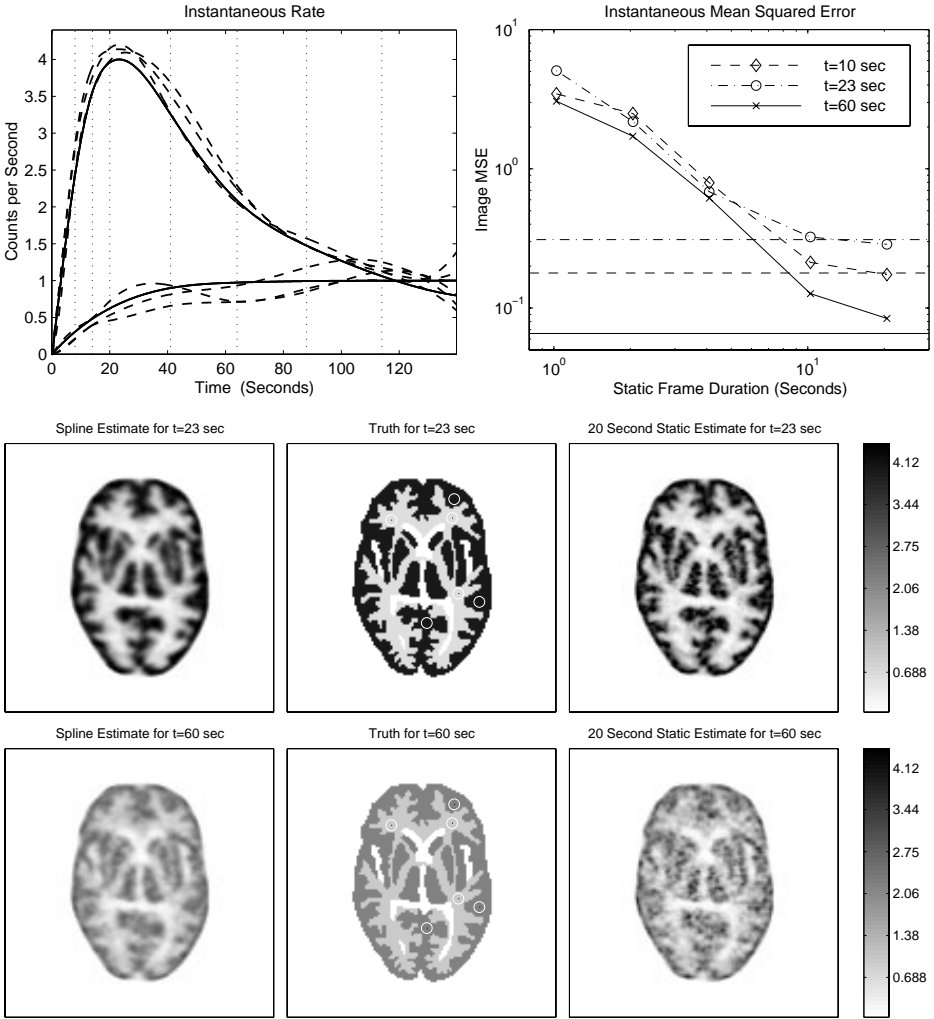


Fig. 1. This figure shows the results of the blood flow simulation. The top left plot shows the estimated (*dashed*) and true (*solid*) rate functions; the vertical lines (*dotted*) indicate knot locations. The top right plot shows the mean squared error over the image for estimating the instantaneous rate at 3 time points; the horizontal lines are for the spline estimate, the decreasing curves show mean squared error for the static estimates of different frame lengths. The bottom two rows show instantaneous rate images; the top row is for $t=23$ seconds, the bottom row for $t=60$ seconds. The left column is the spline estimate, the center column is the truth and the right column is the estimate from the longest static acquisition; the truth images have circles noting the location of the voxels plotted top left

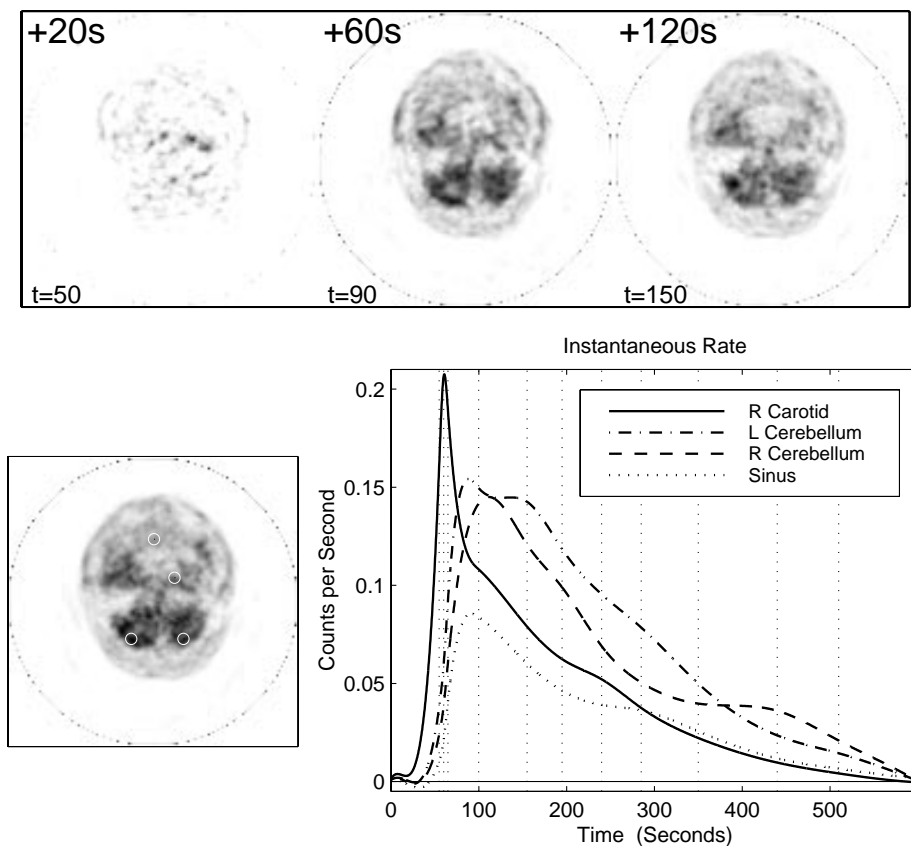


Fig. 2. This figure shows the result of our method using real data. The top row shows a sequence of instantaneous rate images for 20, 60 and 120 seconds post-injection. (Injection occurred at $t \approx 30$ seconds). The early arrival and fast clearance of the tracer in the carotid artery is apparent, as the carotid is visible in the left and center images, but not in the right. The bottom right shows the estimated rate functions for 4 individual voxels; the vertical dotted lines are the knot-locations. The image on left is the total counts image; circles indicate the location of the 4 voxels plotted; the bilateral circles mark the right and left cerebellum, the circle outside of brain tissue is the sinus and the other circle is the right carotid artery

7 Discussion and Conclusions

We have presented preliminary results on estimating continuous time dynamic PET images from list mode PET data. We modeled the dynamic tracer density as an inhomogeneous Poisson process and parameterized the rate functions with a B-spline basis. We introduced the timogram as a means to compactly represent the temporal information of list mode data. The B-spline basis and the timogram's spatial ordering both contribute to an efficient implementation that makes the creation of continuous time reconstructions feasible.

We have presented basic performance analysis with arbitrarily chosen tuning parameters for spatial and temporal regularization. While these results are encouraging in general, Monte Carlo simulations are needed to assess bias and variance in ROI's and in the image at different parameter values.

Estimating images of physiological parameters is a possible extension of this work. This could be accomplished either through embedding the physiological model in the rate function (as in [16]) or estimating parameters with the spline functions. The standard to compare these results to would be estimates from the temporally binned data. In fact, temporally binned data could also be applied in this inhomogeneous Poisson framework, as there is a counterpart to equation (2) for binned data [22].

Acknowledgments

This work was conducted while T.E.N. was a visiting scholar at the University of Southern California in the summer of 1998. The authors would like to thank David Townsend for providing the list mode data and for information on the HR+ list mode format. We also thank Peter Bloomfield for generously sharing his list mode decoder source code. This work was supported by the National Cancer Institute under grant R01 CA59794.

References

1. Barrett, H.H., White, T., Parra, L.C.: List-mode likelihood. *Journal of the Optical Society of America*, **14** (1997) 2914–2923
2. Bartels, R.H., Beatty, J.C., Barsky, B.A.: An introduction to splines for use in computer graphics and geometric modeling. M. Kaufmann Publishers, Los Altos, CA (1986)
3. de Boor, C.: A Practical Guide to Splines. Vol. 27 of Applied Mathematical Sciences. Springer-Verlag, New York (1978)
4. Fessler, J.A.: Penalized weighted least-squares image reconstruction for PET. *IEEE Transactions on Medical Imaging* **13** (1994) 290–300
5. Geman, S., McClure, D.E.: Statistical methods for tomographic image reconstruction. In *Proceedings of The 46th Session of The ISI, Bulletin of The ISI* **52** (1987)
6. Gu, C., Qiu, C.: Smoothing spline density estimation: Theory. *The Annals of Statistics* **21** (1993) 217–234

7. Herscovitch, P., Markham, J., Raichle, M.E.: Brain blood flow measured with intravenous $H_2^{15}O$. I. Theory and error analysis. *Journal of Nuclear Medicine* **24** (1983) 782–789
8. Huang, S.C., Phelps, M.E.: Principles of Tracer Kinetic Modeling in Positron Emission Tomography and Autoradiography. In: *Positron Emission Tomography and Autoradiography. Principles and Applications for the Brain and Heart*. Raven Press, New York (1986)
9. Huffman, D.A.: A method for the construction of minimum-redundancy codes. *Proceedings of the Institute of Radio Engineers* **40** (1952) 1098–1101
10. Johnson, C., Yan, Y., Carson, R., Martino, R., Daube-Witherspoon, M.: A system for the 3D reconstruction of retracted-septa PET data using the EM algorithm. *IEEE Transactions on Nuclear Science* **42** (1995) 1223–1227
11. Kaufman, L.: Maximum likelihood, least squares, and penalized least squares for PET. *IEEE Transactions on Medical Imaging* **12** (1993) 200–214
12. Lee, S.-J., Rangarajan, A., Gindi, G.: Bayesian image reconstruction in SPECT using higher order mechanical models as priors. *IEEE Transactions on Medical Imaging* **14** (1995) 669–680
13. Luenberger, D.: *Linear and nonlinear programming*. Addison-Wesley, Reading, Mass (1989)
14. Matthews, J., Bailey, D., Price, P., Cunningham, V.: The direct calculation of parametric images from dynamic PET data using maximum-likelihood iterative reconstruction. *Physics in Medicine and Biology* **42** (1997) 1155–1173
15. Mumcuoglu, E.U., Leahy, R., Cherry, S.R., Zhou, Z.: Fast gradient-based methods for Bayesian reconstruction of transmission and emission PET images. *IEEE Transactions on Medical Imaging* **13** (1994) 687–701
16. Ollinger, J.M.: Algorithms for parameter estimation in dynamic tracer studies using positron emission tomography. PhD thesis, Washington University School of Medicine, St. Louis, MO (1986)
17. O’Sullivan, F. Image radiotracer model parameters in PET: A mixture analysis approach. *IEEE Transactions on Medical Imaging* **12** (1993) 399–412
18. Parra, L., Barrett, H.H.: List-mode likelihood: EM algorithm and image quality estimation demonstrated on 2D PET. *IEEE Transactions on Medical Imaging* **17** (1998) 228–235
19. Qi, J., Leahy, R.M., Cherry, S.R., Chatzioannou, A., Farquhar, T.H.: High resolution 3D bayesian image reconstruction using the microPET small animal scanner. *Physics in Medicine and Biology* **43** (1998) 1001–1013
20. Qi, J., Leahy, R.M., Hsu, C., Farquhar, T.H., Cherry, S.R.: Fully 3D Bayesian image reconstruction for ECAT EXACT HR+. *IEEE Transactions on Nuclear Science* **45** (1998) 1096–1103
21. Snyder, D.: Parameter estimation for dynamic studies in emission-tomography systems having list-mode data. *IEEE Transactions on Nuclear Science* **31** (1984) 925–931
22. Snyder, D., Miller, M.: *Random Point processes in time and space*, 2nd edition. Springer-Verlag, New York (1991)
23. Snyder, D.L.: Utilizing side information in emission tomography. *IEEE Transactions on Nuclear Science* **31** (1984) 533–537
24. Wahba, G.: Interpolating spline methods for density estimation. I: Equi-spaced knots. *The Annals of Statistics* **3** (1975) 30–48
25. Watson, C.C., Newport, D., Casey, M.E., deKemp, R.A., Beanlands, R.S., Schmand, M.: Evaluation of simulation based scatter correction for 3D PET cardiac imaging. *IEEE Transactions on Nuclear Science* **44** (1997) 90–97