Published in final edited form as:

Phys Med Biol. 2008 June 21; 53(12): 3201–3216. doi:10.1088/0031-9155/53/12/009.

Representation of photon limited data in emission tomography using origin ensembles

A Sitek

Radiology Department, Brigham and Women's Hospital and Harvard Medical School, 75 Francis Street, Boston, MA 02115, USA, E-mail: asitek@bwh.harvard.edu

Abstract

Representation and reconstruction of data obtained by emission tomography scanners are challenging due to high noise levels in the data. Typically, images obtained using tomographic measurements are represented using grids. In this work, we define images as sets of origins of events detected during tomographic measurements; we call these origin ensembles (OEs). A state in the ensemble is characterized by a vector of $3N$ parameters $\mathbf{Y}$, where the parameters are the coordinates of origins of detected events in a three-dimensional space and N is the number of detected events. The $3N$ -dimensional probability density function (PDF) for that ensemble is derived, and we present an algorithm for OE image estimation from tomographic measurements. A displayable image (e.g. grid based image) is derived from the OE formulation by calculating ensemble expectations based on the PDF using the Markov chain Monte Carlo method. The approach was applied to computer-simulated 3D list-mode positron emission tomography data. The reconstruction errors for a 10 000 000 event acquisition for simulated ranged from 0.1 to 34.8%, depending on object size and sampling density. The method was also applied to experimental data and the results of the OE method were consistent with those obtained by a standard maximum-likelihood approach. The method is a new approach to representation and reconstruction of data obtained by photon-limited emission tomography measurements.

1. Introduction

Emission tomography is an important diagnostic imaging modality, as it allows assessment of physiological function. Tomographic reconstruction of emission data is challenging because of high levels of noise in the acquired data, due to limits on radiation dose and the geometrical complexity of the scanners. In this paper, we propose a new representation of the data, and present a method, based on statistical ensembles (Nummelin 1984, Meyn and Tweedie 1993), to accomplish estimation applicable to the new data representation. In this work, we introduce several new concepts to the field of emission tomography.

With the latest hardware developments, the reconstruction of images from emission tomographic data acquired by new scanners is becoming more complex. We will use the term, 'region of response (ROR)', which we define as the region in the imaged volume from which a detected event may have originated. For example, for single-photon-computed emission tomography (SPECT) a simple ROR would be a line determined by collimator geometry intersecting detection location, for positron emission tomography (PET), a simple ROR would be a line joining two detectors that have recorded photon interactions in coincidence; in this case, the ROR is equivalent to the 'line of response (LOR).' For time of flight (TOF)-PET it would be a segment of that line, weighted by the probabilities implied by the TOF information; for a camera based on generation and detection of Compton scattered photons (Singh 1983, Singh and Doria 1983) (Compton cameras) it would be a surface of a cone. When the scanner detects an event by recording photons interacting with detector elements, a corresponding ROR

is assigned to this event. Modern scanners have large numbers of detector elements, which provide many ways the event can be detected, and as a result, there are large numbers of possible RORs. For example, for PET scanners with large numbers of detector elements, such as the ECAT system (high-resolution research tomography, Siemens, Knoxville, TN), the number of possible ROR is on the order of 10^9 , while for Compton cameras (Singh 1983, Singh and Doria 1983) the number of the ROR could reach 10^{18} . For systems with high numbers of ROR, the number of acquired events is considerably less than the number of ROR. The number of acquired events for PET examination is on the order of 10^7 . Therefore, no binning is performed and the acquired data are recorded in a list-mode format. In list-mode data acquisition, the detection parameters (generally, position, energy and time) of every detected event are stored for further processing. Without going into detail, it is sufficient to say that by using the detection parameter saved in list-mode format the ROR can be constructed from saved data. In this work, we focus on emission data saved in list-mode format. This is not, however, a limitation of our approach, because all types of emission tomographic data can be considered as special cases of a general list-mode approach; thus, the methods presented here apply to all variants of emission imaging.

Approaches based on likelihood functions are typically used for the statistical tomographic reconstruction of list-mode data (Snyder 1984, Barrett *et al* 1997, Nichols *et al* 2002) and binned data (Carson 1986, Shepp and Vardi 1982, Qi and Leahy 2006, Defrise and Gullberg 2006). Monte Carlo methods can also be employed for optimization (Webb 1989) or analysis of posterior likelihood (Higdon *et al* 2002). All of the above approaches require projection and backprojection operations. In the projection operation, estimates of acquired data are obtained based on the current image estimate. These estimates are compared to the actual measurements (usually by evaluating likelihood function), and the ratio is recorded. In ML reconstruction the ratio of the forward projected data and the recorded projections is backprojected and that result is multiplied with the current estimate of the reconstructed image. Repeating these steps some number of times yields the estimate.

In standard approaches, the image is represented as a set of voxels. A voxel is a volume element, representing a value on a regular grid in a three-dimensional space. Other 3D image representations have also been proposed, e.g. blobs (Matej *et al* 1994) or point clouds (Sitek *et al* 2006), that use different bases for representing volume data.

In this work, we present a new approach to representation and reconstruction of data acquired by an emission tomography scanner. Although our reconstruction method is stochastic in nature, we do not use any of the features of standard iterative methods. In particular, in order to reconstruct the image, we do not use the likelihood function defined in projection space. It follows that the projection/backprojection operations are also not used since the likelihood of data in projection space is never evaluated. We consider geometrical origins of detected events in continuous three-dimensional (3D) space and define the image as a set of these origins. It follows that the standard image representations (voxels, blobs, point clouds) are not used. Possible configurations of these origins are considered and expectations of some quantities of interests are determined. For example, one such quantity of interest may be the spatial density of origins which, after correcting for scanner sensitivity (Qi and Leahy 2006) can be considered an estimate of emission density, a typical end point of tomographic reconstruction.

In this work, the origin ensemble concept is introduced and the algorithm for tomographic reconstruction based on the origin ensemble is presented. The results of a computer simulation study for 3D list-mode PET and an experimental study results for SPECT are presented.

2. Methods

In the first two sections, the derivation of the theory of origin ensembles and the Markov chain-based approach used to analyze these ensembles are presented. The derivation is based on the assumption that the underlying event emission density function is known. For tomographic reconstruction, this density is unknown, and the goal of the reconstruction is to estimate it. We address this difficulty by using an approximation to find an estimate of event emission density based on the geometrical origins of the events. This is described in section 3, which also includes details of the computer implementation of the methods developed in section 1 and section 2, as well as an outline of an algorithm for tomographic reconstruction of these data. In the final section, the simulation studies and experiments validating the derived theory and algorithm are described.

2.1. Origin ensembles in emission tomography

We assume that for a radiographic tomographic scan, the data are acquired in a list-mode format. During the examination, N events are detected using an emission tomography imaging camera. In general, we define an event as a radioactive disintegration that results in direct or indirect (e.g. a positron that annihilates, creating two photons) emission of photons. For every detected event i , based on the geometry of the tomographic apparatus and physics of the detection, the probabilities $\alpha_i(y)$ are defined. These values describe a probability density that the event was detected in projection bin i assuming it was emitted at a location y , where y is a three-dimensional vector describing spatial coordinates. This corresponds to the definition of the system matrix used in standard approaches with exception that functions $\alpha_i(\cdot)$ are defined over a continuous 3D space. Note that the event index i is a pointer to all parameters stored in the list-mode file for this event. The above description of $\alpha_i(y)$ is the equivalent to the ROR concept discussed in the introduction.

Next, we define the event emission rate as $f(y)$. This function describes the number of events emitted from spatial location y per second per unit volume, where y is a three-dimensional vector. Note that the function $f(y)$ is unknown and the determination of this function is the ultimate goal of the tomographic reconstruction. The probability density function that a given detected event i was emitted at y for a given $f(y)$ can be expressed as

$$\pi_i(y) \sim \alpha_i(y)f(y). \quad (1)$$

It follows that the probability density for a system of N detected events presumably emitted at $y_1, y_2, \dots, y_N$ is

$$\Pi(\mathbf{Y}) \sim \prod_{i=1}^N \alpha_i(y_i)f(y_i), \quad (2)$$

where by the capital bold $\mathbf{Y}$ we denote a $3N$ -dimensional vector built from vector components of y_1 through y_N . For a given $\mathbf{Y}$, we define a probability density $\Pi(\mathbf{Y})$ such that the value $\Pi(\mathbf{Y}) d\mathbf{Y}$ is proportional to a probability that the observed (detected) events have originated from locations described by the vector denoted by $\mathbf{Y}$ with $d\mathbf{Y}$ being an infinitesimal $3N$ -dimensional volume element. We define the origin ensemble (OE) Ω as a set of all possible $\mathbf{Y}$ s. For example, Ω can be equal to the entire $3N$ -dimensional Cartesian space, but some other choices for Ω can be used (for example see the end of section 2.3). Note that equation (2) shows a proportionality

relation. The normalization constant $Z = \int_{\Omega} \prod_{i=1}^N \alpha_i(y_i)f(y_i) d\mathbf{Y}$ is omitted for clarity of notation. A normalization constant (partition function) is a key concept in thermodynamic ensembles (Huang 1967) but here it is not used (it cancels in final equations) and it will not be discussed further. This normalization constant ensures that $\Pi(\mathbf{Y})$ integrated over Ω is equal to unity and that the units of $\Pi(\mathbf{Y})$ are correct.

In order to proceed, it is very important to understand that each state in the ensemble could be interpreted as describing the locations of all detected events with the probability described by equation (2). Unlike maximum-likelihood approaches, where a single maximum-likelihood solution is sought (point estimation) in the OE approach an ensemble average of some quantity that we would like to measure is determined. For example, in order to measure the emission density, D , in some region of the reconstructed image, an ensemble average $D^{\wedge}$ is calculated using the following (Newman and Barkema 1999):

$$\widehat{D} = \int_{\Omega} D(\mathbf{Y}) \Pi(\mathbf{Y}) d\mathbf{Y}, \quad (3)$$

where the integration is performed over the entire Ω . The value of $D(\mathbf{Y})$ corresponds to the value of D for a given $\mathbf{Y}$. Estimation of any quantity can be performed by averaging this quantity over the ensemble as in equation (3). Although equation (3) gives a straightforward method of calculating the ensemble averages, it is impractical due to the dimensionality of the integration space that can reach billions of dimensions, which makes the integral described by equation (3) impossible to evaluate.

2.2. Markov chains Monte Carlo (MCMC)

One would want to reduce computational burden in equation (3) by integrating only where the function $\Pi(\mathbf{Y})$ is large and skipping integrations over the integration space for small values of $\Pi(\mathbf{Y})$. This can be achieved by use of MCMC methods. In this approach a starting configuration $\mathbf{Y}_0$ with $\Pi(\mathbf{Y}_0) > 0$ is first determined, followed by the construction of a Markov Chain of states. The Markovian process is constructed as a sequence of states, in which the next state in a sequence $\mathbf{Y}_{s+1}$ depends only on the current state $\mathbf{Y}_s$ and is independent of the history of how the current state was reached. In theory, if sufficiency conditions are satisfied, the chain of states will reach an equilibrium. When in equilibrium, states in the Markovian sequence are sampled from the ensemble with a chance proportional to $\Pi(\mathbf{Y})$. Thus, the estimation of the ensemble average state can be determined by averaging the Markov states while in equilibrium (Howard 1971):

$$\widehat{D} \approx \frac{1}{C} \sum_{c=1}^C D(\mathbf{Y}_c), \quad (4)$$

where C is the number of states to average.

Two conditions are sufficient for ensuring that the Markov process reaches the equilibrium state. One is the condition of detailed balance:

$$\Pi(\mathbf{Y}_s) P(\mathbf{Y}_s \rightarrow \mathbf{Y}_{s+1}) = \Pi(\mathbf{Y}_{s+1}) P(\mathbf{Y}_{s+1} \rightarrow \mathbf{Y}_s), \quad (5)$$

where $P(\mathbf{Y}_s \rightarrow \mathbf{Y}_{s+1})$ is the probability of transition from state $\mathbf{Y}_s$ to $\mathbf{Y}_{s+1}$ in a Markov chain of states. We will consider $\mathbf{Y}_s$ and $\mathbf{Y}_{s+1}$ as two subsequent states (points in the phase space) in the Markov chain sequence that differ only by the location of a single event k . In particular, without losing generality, at state $\mathbf{Y}_s$, an event k is located at the spatial position described by vector $y_{k,s}$ and for $\mathbf{Y}_{s+1}$ event k is at the spatial location described by vector $y_{k,s+1}$. Thus, using equation (2) we see that

$$\frac{\Pi(\mathbf{Y}_{s+1})}{\Pi(\mathbf{Y}_s)} = \frac{\alpha_k(y_{k,s+1}) f(y_{k,s+1})}{\alpha_k(y_{k,s}) f(y_{k,s})}. \quad (6)$$

Comparing the above with the detailed balance equation (equation (5)) we arrive at

$$\frac{P(\mathbf{Y}_s \rightarrow \mathbf{Y}_{s+1})}{P(\mathbf{Y}_{s+1} \rightarrow \mathbf{Y}_s)} = \frac{\alpha_k(y_{k,s+1}) f(y_{k,s+1})}{\alpha_k(y_{k,s}) f(y_{k,s})}. \quad (7)$$

This equation is very important. It states that if we start at any state in the ensemble with a non-zero probability and change that point stochastically by modifying locations of single events in accordance with the transition probabilities described by equation (7), the detailed balance condition will be satisfied. The other sufficiency condition required for the Markov chain to reach equilibrium is the condition of ergodicity that requires that any state of the system must be reachable from any other state using Markov chain. For the Markov chain defined above, the condition of ergodicity is satisfied because each new state in the Markov chain is constructed by changing the location of a single event. We can easily imagine reaching any configuration by constructing a chain of N new states in each, changing the location of every event to a new location equal to the location of that event in the target configuration. The conditions of detailed balance and ergodicity ensure that Markov chain will reach an equilibrium (steady state).

2.3. Tomographic reconstruction for origin ensembles

The considerations described above were done based on the assumption that $f(y)$ is known. In the case of tomographic reconstruction, this function is unknown and the goal of the reconstruction is to determine this function. We start by noting that the spatial density of emitted events $p(y)$, divided by the acquisition time T , is by definition equal to $f(y)$. Assuming that for a given $\mathbf{Y}$, the emission density $p(y)$ can be estimated $p(y) \approx \hat{p}(y; \mathbf{Y}) + e(y; \mathbf{Y})$ with $e(y; \mathbf{Y})$ describing the error of the estimation we have that

$$\hat{f}(y; \mathbf{Y}) = \begin{cases} \hat{p}(y; \mathbf{Y})/T \in (y) & \text{for } \epsilon > 0 \\ 0 & \text{for } \epsilon \leq 0, \end{cases} \quad (8)$$

where $\epsilon(y)$ is the sensitivity of the imaging device that describes the probability that the event emitted at y is detected by the scanner. Notation $\hat{f}(y; \mathbf{Y})$ indicates the estimate of f evaluated at y which is a function of the vector $\mathbf{Y}$. The same notation applies to $\hat{p}(y; \mathbf{Y})$. Note also that although y is a vector in three dimensions the bold font is not used throughout the paper to clearly distinguish y from $3N$ -dimensional vector $\mathbf{Y}$. In our approach, the deterministic estimation of $\hat{p}(y; \mathbf{Y})$ was used based on the spatial distribution of events described by $\mathbf{Y}$. It is important to mention here that ultra-fast computer implementation of the estimation of $\hat{p}(y; \mathbf{Y})$ (density of events at y for a given $\mathbf{Y}$) needs to be considered since we expect that the final reconstruction process will require long Markov chain and will be computationally demanding. In this work we used the most straightforward and fastest way to estimate the $\hat{p}(y; \mathbf{Y})$. We voxelized the reconstruction volume and counted the events within each voxel. This number of events divided by the voxel volume and acquisition time was used as an estimation of $\hat{p}(y; \mathbf{Y})$ inside any given voxel. It follows that $\hat{p}(y; \mathbf{Y})$ within each voxel was constant. Note that this approximation is very inaccurate for configurations $\mathbf{Y}$ far from the truth. However, once the algorithm reaches equilibrium the estimates $\hat{p}(y; \mathbf{Y})$ are more accurate. The procedure described above describes the biased estimator.

Constructing the Markov chain that will be used for the reconstruction we first define that two subsequent states in the Markov chain in OE reconstruction (s and $s + 1$) differ by a location of an event k . This event is explicitly at locations $y_{k,s}$ and $y_{k,s+1}$, respectively, for states s and $s + 1$. These locations are contained within voxels that are unambiguously described by pairs of indices k, s and $k, s + 1$. The voxel index can be implicitly determined from s and k , since for a given state s locations of all events are known, and in particular the location of the event k . We chose this formalism in order to simplify the notation. Let $n_{k,s}$ and $n_{k,s+1}$ denote the total number of events in these voxels. Thus, by using the above and equation (7) and equation (8) we arrive at

$$\frac{P(\mathbf{Y}_s \rightarrow \mathbf{Y}_{s+1})}{P(\mathbf{Y}_{s+1} \rightarrow \mathbf{Y}_s)} = \frac{\alpha_k(y_{k,s+1})(n_{k,s} - 1)^{n_{k,s}-1}(n_{k,s+1} + 1)^{n_{k,s+1}+1} \in (y_{k,s})}{\alpha_k(y_{k,s})(n_{k,s})^{n_{k,s}}(n_{k,s+1})^{n_{k,s+1}} \in (y_{k,s+1})}. \quad (9)$$

Equation (9) can be derived from equation (7) and equation (8), but to intuitively understand equation (9) one needs to consider that moving the event k from out of the voxel containing $n_{k,s}$ events lowers the count density in this voxel, thus the estimate of f for this voxel will be lowered by $(n_{k,s} - 1)/n_{k,s}$ for all other events in this voxel. Similarly moving an event to voxel containing $n_{k,s} + 1$ events would increase the estimate of f for these events by $(n_{k,s+1} + 1)/n_{k,s+1}$. Note that the above considerations do not destroy the detailed balance or the ergodicity conditions, and thus the Markov chain governed by the transition probabilities described above will certainly reach equilibrium. When in equilibrium, the estimation of the emission density $\hat{f}(y)$ can be done by calculating the ensemble average (equation (4)) by

$$\hat{f}(y) = \frac{1}{CT \in (y)} \sum_{c=1}^C \hat{\rho}_c(y), \quad (10)$$

where $\rho_c(y)$ is the estimated event spatial density at y for the state c sampled from a Markov chain in equilibrium and C is the number of samples.

We will use standard Monte Carlo Markov chain procedure to implement the algorithm. Given the state s , we randomly select new state $s + 1$ from all states for which $\Pi(\mathbf{Y}_{s+1}) > 0$, and accept this state with acceptance probability $A(\mathbf{Y}_s \rightarrow \mathbf{Y}_{s+1})$. In order for detailed balance to hold

$$\frac{A(\mathbf{Y}_s \rightarrow \mathbf{Y}_{s+1})}{A(\mathbf{Y}_{s+1} \rightarrow \mathbf{Y}_s)} = \frac{P(\mathbf{Y}_s \rightarrow \mathbf{Y}_{s+1})}{P(\mathbf{Y}_{s+1} \rightarrow \mathbf{Y}_s)}. \quad (11)$$

The above equation does not put any constraints on the values of the acceptance probability, and only constrains the ratio of the acceptance probabilities. We will exploit that fact to maximize the efficiency of the algorithm. This follows the work of Metropolis *et al* (1953), Hastings (1970), and has been widely used in statistical physics for several decades. Ideally, we would like the acceptance probability to be as high as possible, so that there are no wasted moves; yet at the same time the acceptance probability must be smaller than 1. With that in mind, using equation (9) and equation (11) we define

$$A(\mathbf{Y}_s \rightarrow \mathbf{Y}_{s+1}) = \min \left(1, \frac{\alpha_k(y_{k,s+1})(n_{k,s} - 1)^{n_{k,s}-1} (n_{k,s+1} + 1)^{n_{k,s+1}+1} \in (y_{k,s})}{\alpha_k(y_{k,s})(n_{k,s})^{n_{k,s}} (n_{k,s+1})^{n_{k,s+1}} \in (y_{k,s+1})} \right). \quad (12)$$

Using all of the above, the entire reconstruction algorithm is defined as follows:

- i. Create the starting state $\mathbf{Y}_0$ by randomly selecting the origins of all events in the reconstruction area. The final results will not depend on the starting point so the algorithm used for the creation of the starting point does not affect the final result. However, the time needed for the algorithm to reach equilibrium may depend on the starting point. The probability of the initial state $\Pi(\mathbf{Y}_0)$ has to be greater than zero.
- ii. Randomly select an event k .
- iii. Randomly select the new origin for the event $y_{k,s+1}$ with $\alpha_k(y_{k,s+1}) > 0$.
- iv. Determine the acceptance ratio using equation (12).
- v. Move the event k to the new origin with a probability equal to the acceptance ratio determined in step (iv).
- vi. Go to (ii).

We define one sweep of the algorithm if steps (ii) through (vi) are repeated N times. In order to increase efficiency of the algorithm and to reduce the bias due to the approximation, we randomly selected new origins only inside the object (step iii). In other words smaller than $\mathbb{R}^{3N}$ ensemble was used. The outline of the object is typically available in real studies because

the large majority of modern scanners acquire images of attenuating medium needed for attenuation correction.

2.4. Computer simulations and experimental studies

We simulated a list-mode 3D PET acquisition. A scanner with perfect detection efficiency was simulated as a cylinder with 44.61 cm radius and 16 cm axial span. Objects emitting positrons that annihilate—thereby creating two 511 keV photons—were simulated as 3D ellipsoids. It was assumed that annihilation photons travel in exactly opposite directions, and the positron range was ignored. The ellipsoids representing objects were superimposed to create a more complex geometrical object (figure 1). No photon attenuation was simulated. To achieve a statistically realistic simulation of the event emission, for each object the decays were simulated with exponentially distributed time between the decays with the following exponential distribution $D(t) = \lambda \exp - \lambda t$. The constant λ was specified separately for every ellipsoid. The spatial density of the event emissions within each ellipsoid was assumed to be constant. Thus, the location of each event was randomly generated within the ellipsoid volume. For every event, a random direction over a solid 4π angle was generated. Next, it was determined whether this line intersects the detector surface. If the direction of the given photon pair intersected the detector for both photons the detection was recorded in a list-mode file. The record of each detection consisted of the coordinates of the two detection points (one for each of the two photons) and the time of the detection. Simulation was performed until the number of prescribed photons N to be detected was reached. In this work we used $N = 10\,000\,000$, which is a reasonable number of events that would be expected in an actual acquisition. For the numerical experiments, we used the phantom presented in figure 1, and used λ 's for the small spheres inside the large ellipsoid equal to 2, 0.5, 5, 10, 20 times the λ corresponding to the largest ellipsoid. In other words the activity concentration in the smaller spheres was assumed to have the same ratios of that of the large ellipsoid. Reconstruction was performed using three setups. In the first setup, the known emission distribution f was assumed instead of voxel information. This setup was used to investigate convergence properties of the algorithm. This numerical experiment was possible in the 3D PET computer simulation in which the distributions f were known. In the other two numerical experiment matrices of sizes $128 \times 128 \times 128$ and $384 \times 384 \times 384$ describing the $\hat{p}$ distributions were implemented and used as described in section 2.3. Cubic voxels of sizes 0.55 cm and 0.18 cm, respectively, were used to represent the matrices. The method was also used for reconstruction of torso-phantom data experimentally acquired with a cardiac insert (Data Spectrum, Hillsborough, North Carolina) using SPECT. Projections of the phantom were generated using methodology developed by (Sitek *et al* 2006). A set of 60 2D 128 bin projections acquired over 180° was created. The sinogram consisting of 111 000 counts was created including the cardiac region and reconstructed without attenuation or detector response corrections. Images were reconstructed using both the OE and ML-EM approaches. A 128×128 reconstruction matrix was used for the ML-EM and a 128×128 matrix was used for p approximation for the OE approach. The elements of the system matrix (ROR definitions) were calculated as intersections of parallel strips with square pixels, and the same system matrix was used for both approaches. Reconstructions were carried out by 1500 iterations of ML-EM and 50 000 sweeps of OE approach. The last 10 000 sweeps were used to calculate the ensemble average image of p . The final results obtained by both methods were filtered by the Gaussian filter with an FWHM of 2 pixels.

3. Results

We found that the algorithm applied to the reconstruction with known f rapidly reaches equilibrium for simulated 3D list-mode PET. Based on the information in figure 3 we estimated that about 50 sweeps were sufficient to reach equilibrium. The results of the simulations

performed with known values of f are summarized in table 1. Not surprisingly, we found excellent agreement between the true numbers of emitted events in all objects with the ensemble average numbers estimated by OE reconstruction since no approximations were used. We found errors in the order of 1%. For reconstructions with the density matrices the accuracy was in the range of 2–35% and 4–10% for matrices of $128 \times 128 \times 128$ and $384 \times 384 \times 384$ size, respectively. The ensemble average and standard deviation of numbers of events for each object were estimated by sampling the numbers of events in each object during the MCMC while in equilibrium. We varied the sampling frequency and found no significant differences in values of standard deviation for sampling frequencies larger than 50 sweeps. The reconstruction speed for not optimized code was 350 min per 1000 sweeps for an Intel Xeon 2 GHz processor.

Figure 2 presents visualizations of event locations for the central section of the simulated 3D objects. This section was a 1 cm thick axial slab (figure 1). Events are displayed as semitransparent (50% of opacity) points. For clarity of the images, 50 000 randomly selected points within the slab are displayed. Image 2(B) presents the starting point for the reconstruction. The events are randomly distributed on their LOR as previously described, which results in slightly increased event densities in the center of the reconstruction area seen in figure 2(B). Note that not all of the 50 000 events are displayed in figure 2(B) because the images were cropped to the region containing the object. Figure 2(C) presents visualization of the slab after 6000 sweeps of OE reconstruction with a $128 \times 128 \times 128$ density matrix. The effect of voxelization can clearly be seen. The events tend to group in certain voxels leaving some other voxels empty. This is a similar effect to noise increase in ML solutions. Similarly, the image presented in figure 2(D) for a $384 \times 384 \times 384$ shows the tendency for event grouping in some voxels.

Figure 3 presents the number of events in each object versus the number of sweeps. Reaching equilibrium is much slower for both the $128 \times 128 \times 128$ and $384 \times 384 \times 384$ approaches, in comparison with the approaches with known f due to the enormous size of the phase space (3×10^7 dimensions for the example used in this paper). We empirically found the equilibrium by selecting a point after which the number of events in regions was not changing considerably. We assumed that after 5000 sweeps the system is in equilibrium, although that is not obvious in figure 3. We also found that increasing the number of sweeps does not significantly affect the ensemble average image once the algorithm reached the equilibrium or proximity of the equilibrium, so although we are not certain when looking at the figure 3 about the point at which the algorithm reaches equilibrium, the effect in the resulting reconstructed image due to this uncertainty is minor.

Figure 4 presents the ensemble average images of the emission densities ($\hat{\rho}$) that obtained from the reconstruction. Although these images were used to implement the OE algorithm they correspond to the voxel-based images used in the standard approaches. We provide these images due to lack of other means of displaying images or millions of points. Figure 4(A) is a central axial slice from the $128 \times 128 \times 128$ matrix. The gray scale of these images corresponds to the number of events per voxel, so it corresponds to estimates of emission density. The images on the right are smoothed by a Gaussian filter with a 2 pixel FWHM. The image presented in figure 4(C) is a sum of the three slices at axial locations corresponding to a slice presented in figure 4(A). Summing over the three slices makes the axial thickness of the image presented in figure 4(C) the same as the one in figure 4(A). Note that both ensemble average images (figures 4(A) and (C)) contain voxels that are empty, making the reconstructed image very noisy. Keeping in mind that the presented images are averages over 4000 sweeps, it can be concluded that once in equilibrium the pixels that have a high number of events which vary only slightly, and empty pixels remain empty during the running of the OE in equilibrium.

The quantitative analysis of the algorithm is presented in table 1. We obtained an agreement between the number of events emitted and detected when estimating the number of events using OE with the known f approach. Again, this is because no approximations were used. When density matrices were used to approximate the function f , the quantitative accuracy dropped significantly, especially for small objects comparable in size with the size of voxels. The accuracy was worse for larger voxels. In contrast, for a larger object (object 1 figure 1(A)) the quantitative accuracy lessens with the increased matrix size.

Figure 5 presents reconstructions of experimentally acquired torso phantom data and provides a comparison of images reconstructed by the ML-EM approach with those reconstructed by the method described in this work. We found remarkable similarity between the images. Not only are the major structures similar but also the noise ‘blobs’ are the same. The reconstruction time for 1500 iterations of not optimized code of ML-EM was 63 s and for 50 000 sweeps of OE algorithm 2300 s on Intel Core 2 CPU with 3.46 GHz clock. The random access memory (RAM) requirements for OE are identical to those of the ML-EM.

4. Discussion and summary

We have presented a novel approach to the estimation of emission density from emission tomography data based on origin ensembles. In statistical approaches to reconstruction, the detected data are modeled as random variables and the reconstructed image is defined as a set of parameters (usually voxel values representing emission rates) that are optimized so that the likelihood function is maximized. A variation of this approach is to optimize the likelihood augmented by some priors using Bayesian inference. All of these approaches are point estimations that yield a single estimate that is interpreted as a result of the estimation. In this paper we describe an approach where the system is described by an ensemble. Rather than postulating a single estimate we use an infinite set of possible solutions (states in ensemble) considered all at once. This was investigated for standard image representations in (Higdon *et al* 1997). Probability of each state is defined and in order to estimate some property of the system this property is integrated over the entire ensemble with weighting provided by the probability. Such property could be for example an emission intensity in some region of the imaged space or for that matter emission density in the entire 3D volume of interest. Since we sample from the PDF the uncertainty of the average can also be determined. The PDF for the ensemble can be defined very accurately if f is known and can be approximated for the purpose of tomographic reconstruction when f is unknown.

Our method converges to a stationary state. A point at which the system reaches equilibrium needs to be determined, which sometimes may not be obvious. The reconstruction is considered complete once equilibrium is reached, and enough samples acquired at equilibrium have been generated. The uncertainties of estimated quantities are computed during the estimation with almost no additional computational cost.

One interesting question pertains to how the method would perform under conditions of insufficient angular sampling. During the course of the reconstruction, projection operations are not used; thus, explicit agreement with the projection data is not guaranteed. As it turns out, our approach faces the same challenges and creates artifacts very similar to those that occur with standard reconstruction when insufficient data are used. Under conditions of insufficient tomographic sampling, the probabilities of some states in the phase space are artificially increased or decreased (because of the decreased angular sampling), resulting in incorrect sampling frequency of states by the Markov chain. Thus, in the calculation of the ensemble average, the contribution of these states is weighted by the wrong likelihood, resulting in artifacts in the image.

With known value of f , our method gives unbiased estimates. However, for real experiment with unknown f our estimator is biased because for a given state from the ensemble the $\hat{\rho}(y; \mathbf{Y})$ may be quite different from the true value $\rho(y)$. It follows that the assumed state probabilities are biased as well and so is the estimator. However, as shown in this paper integration over the ensemble gives still remarkably accurate (considering ill conditioning and noise) estimate of the underlying unknown emission density function, which is comparable to maximum likelihood estimate.

One of the very interesting features of our approach is the lack of voxels or any other binned image representation. This represents a major advantage over the standard methods because the geometrical intersections of voxels with ROR are not calculated. This advantage may not be an important feature for standard acquisition geometries such as PET or SPECT, but may play a crucial role in the implementation of tomographic reconstructions for complex acquisition geometries such as those of TOF-PET or Compton camera. In our approach, where the image is defined by x , y and z coordinates of event origins $\mathbf{Y}$, there are no partial volume effects caused by voxelization. However, although the image representation is free of these artifacts, the results of our approach are still biased, because of the voxelization used for $\hat{\rho}(y)$ estimation in the current implementation. In the example of tomographic reconstruction used in this paper, we utilized voxels to estimate the emission density at a given 3D coordinate. We used this simple, deterministic and computationally fast approach but we speculate that the algorithm used in this paper may not be the optimal. Various reconstruction algorithms can be studied by investigating different algorithms for estimation of emission density from spatial origins of the events. For issues related to accuracy and precision of the approach, tomographic regularization of the solution, and partial volume effects, resolution noise trade-offs can be regulated just by changing the matrix size or by devising new algorithms for estimation of $\hat{\rho}(y)$ from x , y and z positions of origins of the events. Obviously, this opens almost endless possibilities for modifications of our algorithm and at this stage of research it is difficult to foresee all possible variations of implementation and advantages that may be discovered.

Our approach allows easy implementation of complex acquisition geometries since we use parameterized ROR. We used lines as ROR for the 3D PET and strips for the SPECT examples presented in this paper. Lines were parameterized by eight independent parameters. Six parameters were used for coordinates of photon interaction locations and the additional two parameters were used to describe the object boundary. This ROR parameterization will be even more important for more complex geometries such as TOF 2D/3D PET or Compton camera. For example, in TOF-PET the time difference between the detection of each pair of photons in coincidence is estimated, and the location of the event on the line of response is modeled by a Gaussian function. For the 3D case, the ROR can be described by only nine parameters. For in which where the exact detector geometric response for PET or SPECT or other complex modalities is used, ROR can also be easily parameterized and used in the approach. The feature of ROR parameterization constitutes a major advantage in implementation for very complex geometries compared to standard iterative methods and is the main reason for this development.

The computational cost in this method is high and billions of Markovian steps will likely be needed to solve any actual 3D problem. Implementation of a more complex algorithm for event-density estimation would require greater computational power. At this stage of research, the reconstruction time is a secondary concern; however, we point out that the algorithm is easily parallelizable and we will parallelize it in the future. Implementation of attenuation correction is straightforward. For example, for PET, the entire correction can be factored to the sensitivity matrix based on known attenuation maps with no other modifications required. For a SPECT implementation, this needs to be incorporated within the sensitivity matrix and the ROR definition, which is also straightforward. The computational difficulty of implementing model-

based corrections for Compton scattering and randoms (for PET) is a drawback to the approach presented here.

The visualization of images reconstructed by our method is an important aspect of our work. For the purpose of visualization of the results presented in this paper we used the simplest approach where we counted numbers of events in each voxel, creating a standard voxelized image with the number of events in each voxel as the image value. This approach though may not be optimal and we plan to investigate other direct methods of visualization, without the need for voxelization, using stereoscopy. We consider this a very important extension of this research. Using developed visualization approaches a comparative analysis of the method presented in this paper with standard approaches to reconstruction, namely the ML-EM and OS-EM algorithms. This comparison will be done using objective task-based image assessment (Barrett and Myers 2004).

In summary, we have presented an approach for image representation in emission tomography. We provided the theory of the origin ensembles and derived an estimator applicable to tomographic reconstruction using OE. Using computer simulations, we have shown that the method performs well for synthetic data simulating 3D PET. Reconstruction of experimentally acquired SPECT data by ML-EM and the approach presented here show remarkable similarities.

Acknowledgments

The author would like to acknowledge Dr Stephen C Moore for help in clarifying the ideas presented here and helpful discussions. Also, he would like to thank Drs Anna Celler, Michel Defrise, Grant T Gullberg and Marie Foley Kijewski for assistance in various aspects of this work. Lastly, he would like to thank Sean Webb for editing this manuscript. This work was supported by the National Cancer Institute, National Institutes of Health under Grant R21 CA123057 and Grant R01 EB00121 and American Heart Association, Scientist Development Grant, Grant 0735328N.

References

- Barrett, HH.; Myers, KJ. Foundations of Image Science (Wiley Series in Pure and Applied Optics). Hoboken: Wiley; 2004.
- Barrett HH, White T, Parra LC. List-mode likelihood. J. Opt. Soc. Am. A Opt. Image Sci. Vis 1997;14:2914–2923. [PubMed: 9379247]
- Carson RE. A maximum likelihood method for region-of-interest evaluation in emission tomography. J. Comput. Assist. Tomogr 1986;10:654–663. [PubMed: 3488338]
- Defrise M, Gullberg GT. [Image reconstruction. Phys. Med. Biol 2006;51:R139–R154. [PubMed: 16790900]
- Hastings W. Monte Carlo sampling methods using Markov chains and their applications. Biometrika 1970;97–109.
- Higdon DM, Bowsher JE, Johnson VE, Turkington TG, Gilland DR, Jaszczak RJ. Fully Bayesian estimation of Gibbs hyperparameters for emission computed tomography data. IEEE Trans. Med. Imaging 1997;16:516. [PubMed: 9368107]
- Higdon D, Lee H, Zhuoxin B. A Bayesian approach to characterizing uncertainty in inverse problems using coarse and fine-scale information. IEEE Trans. Signal Process 2002;50:389.
- Howard, R. Dynamic Probabilistic Systems. New York: Wiley; 1971.
- Huang, K. Statistical Mechanics. New York: Wiley; 1967.
- Matej S, Herman GT, Narayan TK, Furuie SS, Lewitt RM, Kinahan PE. Evaluation of task-oriented performance of several fully 3D PET reconstruction algorithms. Phys. Med. Biol 1994;39:355–367. [PubMed: 15551585]
- Metropolis N, Rosenbluth A, Rosenbluth M, Teller A, Teller E. Equation of state calculations by fast computing machines. J. Chem. Phys 1953;21:1087–1092.
- Meyn, S.; Tweedie, R. Markov chains and Stochastic Stability. London: Springer; 1993.

- Newman, M.; Barkema, G. Monte Carlo Methods in Statistical Physics. New York: Oxford University Press; 1999.
- Nichols TE, Qi J, Asma E, Leahy RM. Spatiotemporal reconstruction of list-mode PET data. *IEEE Trans. Med. Imaging* 2002;21:396–404. [PubMed: 12022627]
- Nummelin, E. General Irreducible Markov Chains and Non-negative Operators. Cambridge: Cambridge University Press; 1984.
- Qi J, Leahy RM. Iterative reconstruction techniques in emission computed tomography. *Phys. Med. Biol* 2006;51:R541–R578. [PubMed: 16861768]
- Shepp L, Vardi Y. Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imaging* 1982;1:113–122. [PubMed: 18238264]
- Singh M. An electronically collimated gamma camera for single photon emission computed tomography: part I. Theoretical considerations and design criteria. *Med. Phys* 1983;10:421–427. [PubMed: 6604216]
- Singh M, Doria D. An electronically collimated gamma camera for single photon emission computed tomography: part II. Image reconstruction and preliminary experimental measurements. *Med. Phys* 1983;10:428–435. [PubMed: 6604217]
- Sitek A, Huesman RH, Gullberg GT. Tomographic reconstruction using an adaptive tetrahedral mesh defined by a point cloud. *IEEE Trans. Med. Imaging* 2006;25:1172. [PubMed: 16967802]
- Sitek A, Reutter BW, Huesman RH, Gullberg GT. Method of generating multiple sets of experimental phantom data. *J. Nucl. Med* 2006;47:1187–1192. [PubMed: 16818954]
- Snyder D. Parameter estimation for dynamic studies in emission-tomography systems having list-mode data. *IEEE Trans. Nucl. Sci* 1984;31:925–931.
- Webb S. SPECT reconstruction by simulated annealing. *Phys. Med. Biol* 1989;34:259. [PubMed: 2784573]

Figure 1.

A three-dimensional numerical phantom used for computer simulations. (A) a view along axial (z) dimension and (B) along the x -direction. The phantom consists of six ellipsoids. The largest ellipsoid (number 6) constitutes a body outline with diameters: 30, 20 and 40 cm in x , y , and z directions, respectively. All other objects (numbers 1 through 5) were simulated as spheres of 3, 1, 1, 1, 1, cm in diameter. The relative emission intensity of the spheres was simulated as follows: 2, 0.5, 5, 10, 20, 1, respectively for objects 1 through 6. The lines in (B) define borders of the 1 cm slab used to create figure 2.

Figure 2.

A section in the 3D image corresponding to a 1 cm thickness slab. For clarity of the image, only 50 000 randomly selected events located within this slab are depicted. (A) corresponds to the true origins of the events. (B) represents the starting point for the reconstruction. (C) and (D) are reconstructions by 5000 sweeps with $128 \times 128 \times 128$ and $384 \times 384 \times 384$ matrices used for the approximation of $\hat{\rho}(y; \mathbf{Y})$, respectively.

Figure 3. The number of events in each object versus the number of sweeps of the MCMC. Data for the first 6000 sweeps are presented. Each graph corresponds to objects 1 through 6, respectively.

Figure 4.

Ensemble average images of the event density. A central slice for the $128 \times 128 \times 128$ reconstruction is depicted in (A). Image (B) corresponds to the Gaussian filter smoothing (FWHM = 2 voxels) of image (A). The gray scale corresponds to the number of events per pixel. (C) represents the sum of three axial slices corresponding to the same axial location as images (A) and (B), and (D) is the image depicted in (C) smoothed by the Gaussian filter (FWHM = 6 voxles). The images were constructed by averaging between sweep 2000 and sweep 6000 every 50 sweeps.

Figure 5.

Reconstruction of SPECT data acquired experimentally for the cardiac torso phantom. Image (A) shows results of ML-EM reconstruction done with 1500 ML-EM iterations, and (B) shows reconstruction of the same data done with 50 000 sweeps of the OE algorithm. The presented image corresponds to the ensemble average image calculated using the last 10 000 sweeps sampled every 50 sweeps. Images (C) and (D) show the same reconstructed images as (A) and (B) but after filtering with a 2-pixel FWHM Gaussian filter. (E) shows the profiles through the images corresponding to lines on images (C) and (D). Apparent increased activity is due to 180° orbit and lack attenuation correction.

(This figure is in colour only in the electronic version)

Table 1
Quantitative analysis of the OE algorithm. Numbers of events corresponding to OE reconstruction are ensemble averages calculated from Markov states while in equilibrium.

	Object 1	Object 2	Object 3	Object 4	Object 5	Background
True emitted	423 617	3909	39 473	78 080	157 599	187 822 935
True emitted & detected	68 765	636	6530	12 885	26 101	9 885 083
OE reconstruction	(-0.4%)	(-1.6%)	(-0.4%)	(-0.2%)	(-0.5%)	(0.0%)
known distribution	69 055	646	6504	12 859	25 984	9 884 976
OE std. dev	(0.4%)	(3.7%)	(1.0%)	(0.8%)	(0.5%)	(0.0%)
known distribution ^a	247	24	66	105	141	314
OE reconstruction	(-2.4%)	(11.5%)	(-25.6%)	(33.4%)	(34.8%)	(0.2%)
128 × 128 × 128	67 108	709	4862	8585	17 010	99 017 47
OE std. dev.	(0.8%)	(13.4%)	(3.3%)	(1.8%)	(1.1%)	(0.0%)
128 × 128 × 128 ^a	536	95	161	155	189	684
OE reconstruction	(-4.4%)	(-4.1%)	(-6.4%)	(-10.1%)	(-6.9%)	(-0.1%)
384 × 384 × 384	65 737	610	6113	11 583	24 297	9 891 659
OE std. dev	(0.8%)	(11.0%)	(2.5%)	(2.0%)	(0.9%)	(0.0%)
384 × 384 × 384 ^a	558	67	153	234	221	584

^a Calculated from data in equilibrium sampled every 50 sweeps. Values in brackets indicate percentage of the measured events.