Reconstruction of multiple-pinhole micro-SPECT data using origin ensembles

Morgan C. Lyon^{a)}

Division of Nuclear Medicine and Molecular Imaging, Department of Radiology, Brigham and Women's Hospital, Boston, Massachusetts 02115 and Department of Radiology, Harvard Medical School, Boston, Massachusetts 02115

Arkadiusz Sitek

Philips Research North America, Cambridge, Massachusetts 02141

Scott D. Metzler

Department of Radiology, University of Pennsylvania, Philadelphia, Pennsylvania 19104

Stephen C. Moore

Division of Nuclear Medicine and Molecular Imaging, Department of Radiology, Brigham and Women's Hospital, Boston, Massachusetts 02115 and Department of Radiology, Harvard Medical School, Boston, Massachusetts 02115

(Received 17 March 2016; revised 24 August 2016; accepted for publication 28 August 2016; published 14 September 2016)

Purpose: The authors are currently developing a dual-resolution multiple-pinhole microSPECT imaging system based on three large NaI(Tl) gamma cameras. Two multiple-pinhole tungsten collimator tubes will be used sequentially for whole-body “scout” imaging of a mouse, followed by high-resolution (hi-res) imaging of an organ of interest, such as the heart or brain. Ideally, the whole-body image will be reconstructed in real time such that data need only be acquired until the area of interest can be visualized well-enough to determine positioning for the hi-res scan. The authors investigated the utility of the origin ensemble (OE) algorithm for online and offline reconstructions of the scout data. This algorithm operates directly in image space, and can provide estimates of image uncertainty, along with reconstructed images. Techniques for accelerating the OE reconstruction were also introduced and evaluated.

Methods: System matrices were calculated for our 39-pinhole scout collimator design. SPECT projections were simulated for a range of count levels using the MOBY digital mouse phantom. Simulated data were used for a comparison of OE and maximum-likelihood expectation maximization (MLEM) reconstructions. The OE algorithm convergence was evaluated by calculating the total-image entropy and by measuring the counts in a volume-of-interest (VOI) containing the heart. Total-image entropy was also calculated for simulated MOBY data reconstructed using OE with various levels of parallelization.

Results: For VOI measurements in the heart, liver, bladder, and soft-tissue, MLEM and OE reconstructed images agreed within 6%. Image entropy converged after ~2000 iterations of OE, while the counts in the heart converged earlier at ~200 iterations of OE. An accelerated version of OE completed 1000 iterations in <9 min for a 6.8M count data set, with some loss of image entropy performance, whereas the same dataset required ~79 min to complete 1000 iterations of conventional OE. A combination of the two methods showed decreased reconstruction time and no loss of performance when compared to conventional OE alone.

Conclusions: OE-reconstructed images were found to be quantitatively and qualitatively similar to MLEM, yet OE also provided estimates of image uncertainty. Some acceleration of the reconstruction can be gained through the use of parallel computing. The OE algorithm is useful for reconstructing multiple-pinhole SPECT data and can be easily modified for real-time reconstruction. © 2016 American Association of Physicists in Medicine. [<http://dx.doi.org/10.1118/1.4962480>]

Key words: SPECT, multi-pinhole, origin ensemble, iterative reconstruction

1. INTRODUCTION

High-resolution (hi-res) SPECT images of a mouse may be obtained by imaging the mouse in a collimator tube that provides multiple pinhole projections.^{1,2} Previously, we described a design for a microSPECT imaging system based on a Trionix Triad triple-camera SPECT system with

two multiple-pinhole tungsten collimator tubes placed end-to-end and centered within the three heads.^{3,4} Briefly, a “scout tube” containing 39 rectangular-aperture pinholes will first be used for high-sensitivity, low-resolution (~2.4 mm system resolution), whole-body mouse imaging in order to locate an organ of interest, such as the heart or the brain. Additionally, the scout tube will acquire sufficient data from

outside the organ of interest to permit accurate tomographic reconstruction in a short period of time. After determining the location of the organ of interest from the scout image, a hi-res pinhole tube will be moved over the mouse, which will also be moved in order to center the organ of interest within the “sweet spot” of the hi-res pin hole tube. The hi-res imaging collimator tube will also be made of tungsten, but it will contain 87 platinum pinhole inserts all focused on the volume-of-interest (VOI). Alternatively, the collimator tubes can be thought of as a single tube with two sections.

For the scout acquisition to be most effective, it is desirable to obtain whole-body images that can be used to precisely locate the target organ in a short period of time and accurately depict some features from outside the organ of interest. Ideally, we would like to be able to rapidly visualize reconstructed images of the whole body of the mouse while acquiring list-mode scout data, and continuously update those images until the organ of interest can be well localized. We have implemented Sitek's origin ensemble (OE) reconstruction algorithm^{3,4} with simulated scout list-mode data in order to evaluate its utility for our application.

The OE reconstruction algorithm lends itself to this application because all events can be processed in image space, thereby avoiding computationally expensive repeated reprojection and back projection. The reconstruction algorithm is also straightforward to implement with list-mode data as new events can easily be incorporated as they are acquired, for online reconstruction. Finally, the OE reconstruction algorithm automatically provides a standard-deviation image that represents the uncertainty of all reconstructed voxel values, in addition to the mean voxel values.

2. METHODS

2.A. System description and modeling

To provide accurate reconstructions, system matrices modeling the geometric response were calculated. Figure 1 illustrates the detector and collimator configurations. The

scout-tube and the high-resolution tube [Fig. 1(b)] were supported by motor-driven rails and centered between the detectors. Three NaI(Tl) detectors [Fig. 1(a)], with active detector areas of 170×224 pixels (2.24 mm/pixel), were positioned 17.2 cm from the center of the scout tube. Lead shielding, in the form of truncated pyramids, was mounted on each camera head to shield the detector from background radiation. Thirty-nine pinholes were modeled for the scout collimator tube with inner diameter 8.0 cm and outer diameter 11.0 cm; 13 of these pinholes projected images onto each camera. Ray tracing was used to create the geometric system model. Therefore, the effects of scatter, attenuation, and pinhole penetration were not incorporated into the system matrices. The depth-of-interaction distribution in the NaI crystals was modeled and included in the system matrices.

For each detector pixel on all three camera heads, the location of possible emission origin voxels within the field of view of the scout tube was first determined. The locations of all such origin voxels were first saved in a precomputed table along with the probability of emission from each voxel and detection in the given detector pixel (i.e., system-matrix elements). The list of image voxels that could contribute to each detector pixel was then sorted from the most probable to least probable origin voxel for each detector pixel. Finally, this list was converted to a list of cumulative probabilities by computing a running sum over the sorted list, and then renormalizing the cumulative probability list to obtain a cumulative density function (cdf) with values ranging from 0 to 1, which was later directly used by the OE reconstruction code. The system sensitivity map, representing the probability that a photon emitted from a given source voxel could be detected by any of the detectors, was also computed from the system matrix and stored in a different table for use by the OE reconstruction program.

2.B. Origin ensemble reconstructions

The OE reconstruction algorithm^{5,6} determines the expected number of photons emitted and later detected from

FIG. 1. Illustration of the dual-resolution imaging system. (a) The transaxial view shows the lead environmental shielding covering the three detector crystals, the aluminum rails used to support the collimator assembly, and the collimator tubes. (b) The axial view shows the tungsten-epoxy collimator with whole-body “scout” imaging section containing 39 rectangular pinholes and a high-resolution section containing 87 platinum-iridium square pinhole inserts. (In this paper, we are only considering OE reconstruction of scout pinhole data.)

FIG. 2. Simplified 2D schematic showing one portion of the collimator and one of the three detectors. In this example, a single event is detected. From the system matrix, it is known that the event could have originated from any of the voxels within the light shaded area of the emission array. The “current origin voxel” represents the emission location currently assigned to the list-mode event. A new “candidate voxel,” representing another possible origin voxel for the list-mode event, is then selected randomly, as described in Sec. 2.B. Finally, to see if the event origin will be moved, a new random number is drawn; if P [by Eq. (1)] is greater than this random number, the candidate voxel then becomes the new origin voxel. Otherwise, the event stays in its current origin voxel. This process is repeated once for every list-mode event each iteration.

each voxel by moving detected photons among allowed voxels from which the emission could have originated (origin voxels). After a statistical equilibrium is reached, the algorithm continues to move photons stochastically from voxel to voxel within the entire region of response. Each combination of origin voxels for all detected events represents one of many possible states of the image; images are typically averaged over many such states to obtain estimates of the ensemble mean and standard deviation (uncertainty) images.

Using the precalculated list of voxels contributing to each detector pixel, a starting image was first created from a dataset by choosing a random possible emission origin voxel for each detected event in the list-mode data with uniform probability. Thereafter, with every iteration (loop over all list-mode events), it was decided whether the origin of every detected emission event would remain in its current voxel or be moved to a new candidate voxel (Fig. 2).

Candidate origin voxels were selected from the precalculated table by drawing a random number on the interval $[0,1)$ which corresponded to the cumulative probability of emission from an origin voxel, given the location of the detected event. In this way, more probable emission locations were chosen with higher frequency. An interpolation search algorithm was then used to efficiently find the position in the table that contained a cumulative probability value just greater than or equal to the chosen random number. To illustrate this concept, consider a simple case in which just five image voxels could be “seen” by a given detector pixel. Assuming lexicographic ordering of all voxels throughout the image volume, these might be, for example, voxel numbers 370, 165, 47, 813, and 610. Table I shows the precomputed cumulative probability of emission from each of the voxels and detection in this event’s detected pixel. If we happened to choose a random number

of, for example, 0.81, the table search algorithm would then rapidly choose voxel number 813 as the next candidate origin voxel for the current event.

The statistical decision to accept or reject a candidate voxel was made in accordance with the Metropolis–Hastings acceptance rule,^{8,9} with probability,

$$P = \min\left(\frac{\varepsilon_i(c_i + 1)}{\varepsilon_{i'}c_i}, 1\right), \quad (1)$$

where c_i is the total number of events in the current origin voxel, i , ε_i is the sensitivity of the origin voxel, and primed symbols refer to the corresponding new candidate voxel properties. For every event, a new random number was drawn and compared with the calculated probability P , which determined whether the new origin voxel would be accepted or rejected. When P was greater than the random number, the event was moved to the candidate voxel. When P was less than the random number, the candidate voxel was rejected and the event remained at its current position.

One iteration of the OE algorithm therefore consisted of one attempt, for every detected photon, to move its origin to a new candidate voxel using the rule above. In our implementation,

TABLE I. Example of a simple cumulative probability density table for a given detector pixel.

Cumulative probability	Voxel number
0.3	370
0.6	165
0.8	47
0.9	813
1.0	610

FIG. 3. The percentage of detected list-mode events that were moved to candidate origin voxels within each iteration of OE.

we elected to process the events in the order of list-mode events; however, the events can be processed in any random order. Figure 3 shows the percentage of events moved as a function of the number of iterations for a simulated scout imaging application using the cardiac tracer, ^{99m}Tc -Sestamibi. Finally, after completing a desired number of iterations, the entire reconstructed image (mean number of events per voxel) was divided by the system-sensitivity map, which compensates for variations in the overall system detection efficiency and converts the values of image voxels from photon-origin counts to the number of the total emissions per voxel. This quantity is directly proportional to the tracer concentration.

2.C. Simulated data and reconstructions

All scout data were reconstructed into a $42 \times 42 \times 88$ voxel (0.81 mm/voxel) image array on an Apple MacPro system with a 12-core Intel Xeon E5 2.7 GHz processor. SPECT list-mode data were simulated for our prototype scout pinhole collimator using the MOBY phantom. First, a list-mode dataset containing approximately 9×10^6 simulated events (including a small number of multiplexed events) was reconstructed using MLEM (Ref. 10) and OE, separately. This number of list-mode events could be acquired in less than 2.5 min using the scout collimator, assuming a ~ 1 mCi injection of a Tc^{99m} tracer (e.g., Sestamibi) into a mouse. VOIs were drawn in the heart, liver, bladder, and soft-tissue of these reconstructed images. The mean counts per voxel were calculated for each VOI along with the average standard deviation from the estimated noise image.

Next, to evaluate the convergence of OE for whole-body reconstructions, we simulated 6.8×10^6 events using the MOBY phantom. Data were reconstructed for up to 6000 iterations of OE. The whole-body image entropy,

$$H = \sum_i \frac{n_i}{N} \log\left(\frac{n_i}{N}\right), \quad (2)$$

was calculated for N total detected events every five iterations, where n_i is the current number of event origins in voxel i .

Because we wanted to be able to localize the organ of interest, we also determined the number of iterations beyond which the counts in a region of interest approached a stationary state. Therefore, we also measured the average counts per voxel in a region encompassing the heart.

2.D. Accelerated and real-time reconstructions

For the application of scout imaging, we would like to be able to obtain images of an organ of interest as quickly as possible for localization. We developed an accelerated multi-threaded (parallelized) version of the OE reconstruction algorithm to improve the reconstruction time for real-time imaging. In order to avoid race conditions in our parallel code, some modifications had to be made to the OE reconstruction routine.

The probability with which moves are accepted [Eq. (1)] is, in part, based on how many photon origins are in the current voxel (c_i), as well as how many are in the candidate voxel ($c_{i'}$). It follows that only one origin of an emission photon can be moved at a time. Therefore, in our accelerated version of OE, the candidate origin voxel was selected and the decision to move or not move an event was made for all detected events within the parallel section of code (Fig. 4). After decisions were made on some number of events, n , the accepted moves from all threads were then processed sequentially within a single thread to update the image array before proceeding to the next group of n events. This “synchronization” procedure (periodic updating of the image array) was performed j times per iteration, where $j = N_{\text{tot}}/n$, N_{tot} is the total number of list-mode events, and an iteration is defined as accepting or rejecting a move for every list-mode event. Each list-mode event was processed once per iteration, i.e., the N_{tot} list-mode events were partitioned across the threads, such that each thread processed its own unique set of events. This differed from the serial version in which events were moved sequentially. In this work, we briefly investigated how the number of image updates per iteration affects the convergence behavior of parallel OE.

Finally, we anticipated that the convergence of the OE reconstruction would behave differently for real-time image reconstruction, where new list-mode events would be incorporated almost continuously. To study this real-time convergence behavior, we calculated the entropy from reconstructions of simulated MOBY phantom data for which additional new events were incorporated during the reconstruction. Beginning with a 1×10^6 event simulated dataset, a million new list-mode events were incorporated every 100 iterations of parallel OE until all $\sim 9 \times 10^6$ events were acquired/reconstructed.

3. RESULTS

3.A. Comparison with MLEM

The images resulting from MLEM and OE reconstructions of identical datasets were visually very similar (Fig. 5). The MLEM image was reconstructed for 300 iterations.

FIG. 4. Flowchart illustrating parallel implementation of the OE reconstruction algorithm for the case when all list-mode events are being processed with one image update. N_{tot} is the total number of list-mode events, while N_{thread} is the number of processors or threads. As described in the text, it is also possible to process a smaller number of events, n , over all threads, using j image updates per iteration, where $j \sim N_{\text{tot}}/n$.

An ensemble-mean image was calculated during the OE reconstruction, using 500 states (sampled every 10 iterations) between 2000 and 7000 iterations. We also calculated the standard deviation over the ensemble to obtain an estimate of the noise image. Similar mean VOI values were obtained

for both MLEM and OE reconstructions (Table II). The total counts in the OE image measurements differed from the MLEM measurements by 5.8%, 0.7%, 5.4%, and -1.8% , respectively, for VOIs in the heart, liver, bladder, and soft-tissue.

FIG. 5. Coronal and axial slices (0.81 mm thick) of a MOBY phantom simulation with 9M events reconstructed by MLEM [(a) and (b)] and OE [(c) and (d)], including the OE uncertainty images [(e) and (f)]. All images were smoothed with a 0.6 mm 3D-Gaussian filter. The MLEM reconstruction time was 75 min. The OE reconstruction [(c) and (d)] time was 201 min using the serial method. Using parallel OE with 32 updates/iteration, the same reconstruction could be obtained in 27 min.

3.B. Steady state convergence

For OE reconstructions of 6.8×10^6 simulated events, the entropy, calculated over the entire image, asymptotically approached a steady-state convergence after approximately 2000 iterations (Fig. 6). Additionally, we found that while the parallel version, with only one image update per iteration, completed ~ 9 times more iterations than the serial version for equivalent time, the two versions of OE did not converge to the same image entropy. In the same reconstructions, the mean counts per voxel in a VOI within the heart stabilized after ~ 200 iterations of OE (Fig. 7) for both versions. The left and right ventricles could be visualized reasonably well in transaxial views of the heart between 100 and 400 iterations with even better delineation at 2000 iterations. More frequent updates (Fig. 8) improved the reconstruction performance at the cost of longer reconstruction times.

TABLE II. MOBY phantom simulation VOI measurements.

Region	MLEM mean (Cts/voxel)	OE mean (Cts/voxel)	Mean of uncertainty image (Cts/voxel)
Heart	330	349	92.4
Liver	293	295	112
Bladder	462	487	111
Soft tissue	27.2	26.7	20.6

FIG. 6. Total image entropy vs iteration number for serial version (dotted-line), parallel version (dashed-line), and combined version (solid-line) of OE, for a 6.8M simulated event.

We also considered the use of a combination of the two methods to optimize speed and image quality in whole-body imaging (Fig. 6). In the combined version, the data were reconstructed using the parallel version of OE (with $j = 1$) for the first 1000 iterations after which we continued to iterate using the serial version of OE until the minimum image entropy was reached. A total-image entropy equivalent to that of an image reconstructed by conventional OE for 2000 iterations (155.3 min), could be obtained in 88.9 min of reconstruction time using the combined version (1000 iterations parallel, followed by an additional 1015 iterations serial).

Another way in which the performance of parallel OE could be improved was by adjusting the number of synchronizations, j , per iteration. Identical 6.8×10^6 events simulated MOBY datasets were reconstructed using parallel OE with 1, 2, 4, and 32 updates/iteration. We observed that while more frequent updates improved the algorithm's performance in terms of lower total-image entropy (Fig. 8), it also decreased the

FIG. 8. The number of image array updates (thread synchronizations) per iteration was evaluated for the conditions of 1 ("parallel"), 2, 4, and 32 updates/iteration.

reconstruction speed. The system overhead time associated with the multi-threaded code slowed the reconstruction such that parallel OE with more than ~ 2000 updates/iteration was slower than the serial OE (Table III).

3.C. Real-time reconstruction

For simulated "real-time" reconstructions using parallel OE, the image entropy initially increased each time new events were incorporated, which caused sharp increases in the entropy curve (Fig. 9). However, the entropy recovered within a few iterations. As the total number of acquired events increased, the increase in image entropy was less. The first image, 100 iterations of 1×10^6 counts, was reconstructed in 8.7 s and would be more than adequate for scout localization of the heart in preparation for high-resolution imaging.

FIG. 7. Mean voxel count values within a VOI in the heart for the serial and parallel versions of OE. A transaxial slice through the heart is also shown at 50, 100, 500, and 2000 iterations. 3D Gaussian smoothing (0.6-voxel FWHM) was applied to all reconstructions.

TABLE III. OE reconstruction times for 1000 iterations simulated MOBY phantom data with 6.8×10^6 events.

Method	OE parallel	OE parallel	OE parallel	OE parallel	OE serial
No. of updates/iteration	1	32	400	2000	Continuous
Time (min)	8.9	10.1	23.0	79.7	79.2

4. DISCUSSION

We have demonstrated that the OE algorithm, originally developed for PET reconstruction, can be used to reconstruct multiple-pinhole list-mode SPECT data to obtain images quantitatively similar to those reconstructed by MLEM. In this work, we compared the mean counts per voxel for four VOIs within simulated MOBY data reconstructed with OE and MLEM, demonstrating quantitative agreement within $\sim 6\%$. The results of the comparison between MLEM and OE are also consistent with those of Andreyev *et al.* for a Compton-camera reconstruction,⁷ as well as those of Sitek for PET reconstruction.¹¹ In addition to achieving quantitatively similar data to MLEM, we have shown that OE can be used to reconstruct data while it is being acquired. Conventional MLEM cannot be used for real-time imaging applications.

Excellent OE-reconstructed images can be obtained when the algorithm is allowed to run to convergence for the whole body and for a region encompassing the heart. Image

quality can be improved further by averaging images over an ensemble of many allowed states of the system after statistical equilibrium has been reached. Since the decision to move events is made probabilistically, there will always be some number of events moved each iteration, even after the equilibrium state has been reached (Fig. 3). For the conditions presented in this paper, the probability of moving an event in the equilibrium state was $\sim 20\%$. However, this value will be dependent on the object and the activity distribution within the object.

In addition, image uncertainty can be estimated in a straightforward manner by computing the standard deviation over the same ensemble of images. Although not shown here, it is also possible to calculate the covariance between neighboring voxels or between any ROIs evaluated within the image dataset. While we have not yet explored the use of the uncertainty image for a particular scout imaging application, we expect that in the vicinity of the targeted organ of interest, a “signal-to-noise” image in which each voxel’s mean value is

FIG. 9. The image entropy vs iteration number, where a million new events were added every 100 iterations, beginning with 1×10^6 events at the first iteration (dashed-line). The same $\sim 9 \times 10^6$ count dataset was also reconstructed using all the events for all iterations (solid-line). Below: Single-slice views of reconstructed images with 0.5 mm Gaussian smoothing.

divided by its standard deviation may eventually prove useful for estimating the reliability of organ localization and for helping to determine whether an adequate number of scout events have been acquired so that the imaging study could then proceed to the high-resolution phase of data acquisition. It should be noted that the list-mode likelihood algorithm of Parra and Barrett¹² was also developed to permit direct image reconstruction from list-mode data; however, this algorithm does not automatically provide an estimate of the image uncertainty.

To accelerate reconstructions for scout imaging, a modified parallel version of OE can be implemented. However, since the parallel and serial versions of OE do not converge to the same state (due to an approximation made in the parallel version), a combined hybrid version of the two was introduced, which provides some acceleration without loss of accuracy. We propose that the combined version could be useful for reconstruction of scout image data in real time because the initial parallel-OE reconstruction would allow the microSPECT user to quickly localize the organ of interest. Afterward, the program could continue to reconstruct the scout data using serial OE while the hi-res scan is being acquired.

An alternative approach would be to increase the frequency with which the image array is updated and synchronized across the multiple threads. The fully parallelized version of OE was synchronized only once per iteration (at the end); however, as the number of synchronizations approached the number of detected events, the algorithm's convergence behavior approached that of the serial version (Fig. 6). Correspondingly, the degree of acceleration diminished with increasing numbers of synchronization events (Table III). Reconstruction times did not scale linearly with the number of threads due to system overhead time associated with creating the threads, as well as the time cost in updating the image array after each move. Likewise, while increasing the number of updates per iteration improved the algorithm's performance as measured by entropy, there was also an associated time cost in waiting for all threads to complete before applying the update to the image array. Therefore, the use of updates with the parallel version of OE became increasingly prohibitive up to around 2000 updates/iteration. Thereafter, more updates produced a version of OE that was slower than the serial version.

In the current implementation of OE, a flat prior is used. Other priors can also be implemented; however, not all priors can be efficiently implemented with OE and more research is needed to find new algorithms that can be used to estimate posteriors with general priors. For a list of priors that can be implemented efficiently, as well as a discussion of computational aspects, readers should refer to Sitek.⁵

A limitation of this work was that we did not correct for attenuation in either the MLEM or OE reconstruction. It should be noted, however, that an attenuation map, by use of either a CT attenuation map or a map estimated by other means (e.g., atlas matching), could easily be incorporated into the OE system matrix. The probability of the photon not being attenuated or scattered within the mouse could be computed along each possible straight-line trajectory from a given source voxel to the detector pixel. The attenuation

information could then be included by multiplying the system matrix values by the probability of transmission through the attenuation map from the appropriate voxel to the detector pixel of interest. While this approach is straightforward in principle, it may of course be nontrivial in practice. For example, if the mouse must be moved to a separate CT scanner to obtain the attenuation map, the emission and attenuation images must later be spatially registered in order to perform accurate attenuation correction. For our application of small animal imaging, the effects of attenuation are small. Therefore, microSPECT researchers commonly ignore attenuation.

5. CONCLUSIONS

We have shown, for the first time, that the origin ensemble reconstruction can be used for reconstructing multiple-pinhole SPECT data. The same dataset led to quantitatively and qualitatively similar results when reconstructed by MLEM and by OE. Finally, we have also shown that the OE algorithm can easily handle real-time reconstructions of list-mode SPECT data by incorporating new events with each iteration every few iterations which is of particular interest for certain applications, e.g., those for which data need only be acquired until an area of interest can be visualized well-enough to position the mouse for the high-resolution scan.

ACKNOWLEDGMENTS

This research was supported by the National Heart, Lung, and Blood Institute of the National Institutes of Health (NIH) under Award No. R01HL111883 (PIs: Metzler and Moore). The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. The authors would like to thank Dr. Dan Xia and Dr. Mi-Ae Park for their contributions to the design of the microSPECT system on which our simulations were based.

CONFLICT OF INTEREST DISCLOSURE

The authors have no COI to report.

^{a)} Author to whom correspondence should be addressed. Electronic mail: mlyon@invicro.com

¹F. P. DiFilippo, M. J. Riffe, K. M. Harsch, N. P. McCabe, and W. D. Heston, "Detached multi-pinhole small animal SPECT device with real-time calibration," *IEEE Trans. Nucl. Sci.* **53**, 2605–2612 (2006).

²F. P. DiFilippo, "Design and performance of a multi-pinhole collimation device for small animal imaging with clinical SPECT and SPECT-CT scanners," *Phys. Med. Biol.* **53**, 4185–4201 (2008).

³S. D. Metzler, S. C. Moore, and M.-A. Park, "Design of a new small animal imaging system based on rectangular pinhole aperture," *Nuclear Science Symposium and Medical Imaging Conference, M21-43, Anaheim, CA, November 2012*.

⁴S. C. Moore, M.-A. Park, D. Xia, and S. D. Metzler, "Dual-resolution MicroSPECT mouse imaging using a triple-head SPECT system," *Nuclear Science Symposium and Medical Imaging Conference, M25-1, Seoul, South Korea, November 2013*.

⁵A. Sitek, "Data analysis in emission tomography using emission-count posteriors," *Phys. Med. Biol.* **57**, 6779–6795 (2012).

- ⁶A. Sitek, "Representation of photon limited data in emission tomography using origin ensembles," *Phys. Med. Biol.* **53**, 3201–3216 (2008).
- ⁷A. Andreyev, A. Sitek, and A. Celler, "Fast image reconstruction of Compton camera using stochastic origin ensemble approach," *Med. Phys.* **38**, 429–438 (2011).
- ⁸N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equations of state calculations by fast computing machines," *J. Chem. Phys.* **21**, 1087–1092 (1953).
- ⁹W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika* **57**, 97–109 (1970).
- ¹⁰L. A. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Trans. Med. Imaging* **1**, 113–122 (1982).
- ¹¹A. Sitek, "Reconstruction of emission tomography data using origin ensembles," *IEEE Trans. Med. Imaging* **30**, 946–956 (2011).
- ¹²L. Parra and H. H. Barrett, "List-mode likelihood: EM algorithm and image quality estimation demonstrated on 2-D PET," *IEEE Trans. Med. Imaging* **17**, 228–235 (1998).