Published in final edited form as:

Phys Med Biol. 2010 September 21; 55(18): 5341–5361. doi:10.1088/0031-9155/55/18/006.

Evaluation of a 3D point cloud tetrahedral tomographic reconstruction method

N F Pereira and A Sitek

Department of Radiology, Brigham and Women's Hospital-Harvard Medical School Boston, MA, USA

N F Pereira: nfp4@bwh.harvard.edu; A Sitek: asitek@bwh.harvard.edu

Abstract

Tomographic reconstruction on an irregular grid may be superior to reconstruction on a regular grid. This is achieved through an appropriate choice of the image space model, the selection of an optimal set of points and the use of any available prior information during the reconstruction process. Accordingly, a number of reconstruction-related parameters must be optimized for best performance. In this work, a 3D point cloud tetrahedral mesh reconstruction method is evaluated for quantitative tasks. A linear image model is employed to obtain the reconstruction system matrix and five point generation strategies are studied. The evaluation is performed using the recovery coefficient, as well as voxel- and template-based estimates of bias and variance measures, computed over specific regions in the reconstructed image. A similar analysis is performed for regular grid reconstructions that use voxel basis functions. The maximum likelihood expectation maximization reconstruction algorithm is used. For the tetrahedral reconstructions, of the five point generation methods that are evaluated, three use image priors. For evaluation purposes, an object consisting of overlapping spheres with varying activity is simulated. The exact parallel projection data of this object are obtained analytically using a parallel projector, and multiple Poisson noise realizations of these exact data are generated and reconstructed using the different point generation strategies. The unconstrained nature of point placement in some of the irregular mesh-based reconstruction strategies has superior activity recovery for small, low-contrast image regions. The results show that, with an appropriately generated set of mesh points, the irregular grid reconstruction methods can out-perform reconstructions on a regular grid for mathematical phantoms, in terms of the performance measures evaluated.

1. Introduction

There has been interest in the development of content adaptive, irregular mesh methods for tomographic image reconstruction in nuclear medicine (Yang *et al* 2002, 2003, Brankov *et al* 2002, 2004, Sitek *et al* 2006, Yu *et al* 2009). This is partly because such methods result in a reduced size of the tomographic inverse problem which may result in improved computational performance, both in terms of execution times and memory requirements. In some cases, such benefits could be substantial, especially for fully 3D reconstruction of projection data. However, more importantly, due to having to estimate fewer number of parameters in the reconstruction process, these methods can lead to better conditioning of the inverse problem, better regularization properties and thus superior statistical performance of the final reconstructed data.

An irregular grid is obtained when based on local image content, the reconstruction domain is sampled at a specified number of arbitrarily located points. The tomographic reconstruction algorithm on such an irregular grid usually involves the (i) generation of an appropriate set of mesh points, (ii) generation of the mesh elements (area in 2D and volume

in 3D) from these points, (iii) generation of the reconstruction system matrix, and finally, (iv) the inversion of the projection data using a suitable inversion algorithm. Needless to say, a number of methods can be employed to accomplish each of these stages, each having its own set of tunable parameters. The final image quality, in whatever sense that may be defined, is likely to depend on these various factors. In order to characterize the overall method, it is therefore necessary to take some of these factors into consideration.

To date, a number of algorithms for the generation of mesh points from 2D (Yang *et al* 2003, Garcia 1997, Yu *et al* 2009) and 3D (Yang *et al* 2002) image priors have been developed. These methods use a fixed set of pre-determined mesh points obtained from an image prior, for the reconstruction. These points are also constrained to occupy locations defined by the spatial discretization of the prior. One of the 3D point generation algorithms evaluated in this work uses a modified version of the method proposed in Yang *et al* (2003), which uses the Floyd–Steinberg (Floyd and Steinberg 1976, Sandler *et al* 1997) error diffusion algorithm to generate points with minimum distances bounded by the pixel/voxel size. Our method, however, produces points at sub-voxel discretizations and at arbitrary locations in the image domain. Some evaluation studies for a 2D content adaptive mesh reconstruction methods were conducted in Brankov *et al* (2002, 2004) where mesh-based reconstructions were shown to have superior performance compared to the regular grid reconstructions.

In this work, we evaluate the quantitative performance of a 3D point cloud mesh reconstruction method that employs tetrahedral volume elements and analytically computed system matrix (Sitek *et al* 2006). Five point placement strategies are studied, three of which use image priors to generate an initial set of points for the point cloud. For evaluation purposes, we generate exact projection data of a simulated object and its multiple noise realizations. These are reconstructed using the maximum likelihood-expectation maximization (ML-EM) algorithm (Lange and Carson 1984, Shepp and Vardi 1982). To obtain the exact projection data, an analytical parallel projector is used. However, image degrading factors due to attenuation, scatter and finite system resolution are not modeled into the acquisition process.

Our hypothesis is that, due to the more involved image model employed in deriving the system matrix and reduction in the number of reconstruction unknowns, the tetrahedral method produces a better conditioned solution having superior noise properties compared to reconstruction on a regular grid. Also, due to content adaptive multi-resolution sampling, it does not introduce significant bias in the solution. Hence, although we employ an idealized data acquisition model, our approach is consistent with the hypothesis that we seek to establish. For the analysis of quantitative performance, we look at the bias and variance of the reconstructed data computed over a number of pre-specified image regions. To assess the effects of sub-voxel sampling on activity recovery, we also evaluate the recovery coefficient over an image region that represents a low-contrast lesion. For comparison purposes, a similar analysis is performed on reconstructions using the regular grid reconstruction method.

The rest of this paper is organized as follows. Section 2 introduces the regular and point cloud image representations and their reconstruction methods. The five point placement strategies are then discussed in section 3. Section 4 details the performance metrics based on which the reconstructions are evaluated. Section 5 discusses the data generation process and the study setup. Results and discussions are presented in section 6 and conclusions in section 7.

2. Image representation models

A standard method of reconstructing and representing 3D tomographic image data is to divide the image domain into a regular array of voxels within each of which the reconstructed image activity is assumed to be a constant. Such voxel-based representations can be described using constant and spatially limited basis functions on a regular grid. In order to accurately represent small lesions, this approach, however, must introduce a large number of small voxels, most of which also occupy regions having negligible activity or features of little diagnostic interest. For tomographic reconstruction, this leads to an increase in the number of unknowns that have to be solved for, thereby increasing the computational burden of the problem and degrading its conditioning. An alternative approach, as investigated here, is to proportionately place a greater number of points at unconstrained locations within image regions that have significant detail, or at pre-specified regions of interest. Given a collection of such point locations, the point cloud reconstruction method (Sitek *et al* 2006) assumes a linear model for image representation in between the points. In three-dimensional reconstruction space this assumption naturally leads to tetrahedral volume elements. The irregular nature of the point cloud results in irregular sizes and shapes for the tetrahedral volume elements. The point values are the unknowns that have to be solved for in the reconstruction process. The tetrahedral mesh is obtained from the initial set of points using a 3D Delauney algorithm (Delauney 1934, Si and Gartner 2004, Si 2005). This representation has the desirable feature that it partitions the reconstruction space into a set of non-overlapping volume elements and the union of all volume elements covers the entire reconstruction volume.

2.1. Point cloud reconstruction

On obtaining the tetrahedral mesh, the system matrix is generated. The process involves computing the impulse response of each point for each discrete angle of rotation. To this end, each tetrahedron containing a particular point is projected on to the detector plane and weights computed for detector bins that have non-zero intersection with the tetrahedron's projection. Details of the procedure can be found in Sitek *et al* (2006). The final reconstructed data consist of mesh point locations and the corresponding activity values, which are computed using the ML-EM algorithm.

2.2. Regular grid reconstruction

For the regular grid case, we consider a two-dimensional method which reconstructs the projection data slice by slice along the axial direction. The system matrix elements corresponding to a *pixel-bin* pair at a given detector angle are computed by evaluating the exact area of intersection between the beam path corresponding to that bin and the rectangular pixel area, at that angle. The ML-EM reconstruction algorithm is then used to reconstruct the projection data slice by slice.

3. Point cloud generation

We look at five methods for generating points for the tetrahedral mesh reconstructions. In all of these methods, the total number of points for a particular reconstruction is user specified. Four of the methods follow a *multi-stage* point addition strategy, whereby a predefined number of new points are adaptively introduced into the image domain based on local image content. The remaining strategy is a single-stage method in that all points determined from a suitable image prior are introduced into the reconstruction at the beginning and no points are added during the reconstruction. A *stage* here is defined as the sequential execution of the following three processes:

- i. point generation

- ii. evaluation of the system matrix
- iii. ML-EM iterations.

For a multi-stage process, an image estimate is obtained at the end of each stage using a predefined number of ML-EM iterations. This estimate is then used to generate points for the subsequent stage. The number of points added at the beginning of each stage is given by a fixed fraction, α , of the number of points already added to the reconstruction. If N_0 is the initial number of points, N_T the final number and D the number of stages, then α is determined by solving the polynomial equation given by

$$(1+\alpha)^D - \frac{N_T}{N_0} = 0. \quad (1)$$

The total number of ML-EM iterations for the entire reconstruction is also a user-defined parameter. This number is split across the total number of stages. For initial stages, the number of iterations per stage is kept low and is gradually increased in subsequent stages. Note that the system matrix must be updated after each point addition stage. Such a multi-stage strategy ensures that a greater number of points are introduced as the reconstructed image estimate gets better. We next discuss the different point addition strategies. Table 3 provides a summary of the values of the parameters related to these strategies that were used in this work.

3.1. Sorted variation method

This is a multi-stage point generation method that does not use an image prior for the mesh points. The tetrahedral reconstruction is initially started on a coarsely spaced set of mesh points on a regular grid of dimensions $10 \times 10 \times 10$. At each stage, a new point is introduced at the center of a given fraction, α , of tetrahedrons that have the largest variation. The variation of the j th tetrahedron, v_j , is determined from the intensity values at its points, I_i s, according to

$$v_j = \frac{1}{4} \sum_{i=1}^4 (I_i^j - \bar{I}^j)^2. \quad (2)$$

Here, $\bar{I}$ is the mean of the values of the j th tetrahedron points. The strategy is so named, since the variation of the tetrahedrons is sorted to determine the set tetrahedrons into which the additional points are introduced.

3.2. Histogram method

The sorted variation method described in section 3.1 can lead to point clustering as neighboring tetrahedrons can have similar high tetrahedral variation values. Since, a pre-specified number of points are used for a given reconstruction, such clustering can lead to other image regions being under-sampled causing artifacts in the reconstructed images. The histogram method partially alleviates this problem by determining a histogram of the variation coefficients, v_j s, into a predetermined number of bins. For a given bin k , a number of tetrahedrons, b_k , are selected. This number is proportional to the mean bin value and to the number of tetrahedrons in that bin. It may happen that the number of tetrahedrons B_k present in the bin k is less than b_k . In such cases, all of the B_k tetrahedrons are selected from the bin k and the difference $(b_k - B_k)$ is carried over to the bin $k - 1$. The notation here assumes that bins are numbered in descending order of mean bin values and that bins are traversed in the same order. The histogram method results in spreading points more

uniformly over the reconstruction region. Like the sorted variation method, this is a multi-stage method that does not use an image prior for the mesh-points.

3.3. Error diffusion (prior) method

This is a single-stage point placement strategy whereby mesh points generated from an image-based prior are used for the reconstruction. The prior itself consists of an edge-map derived from an image which may be obtained from a different imaging modality such as CT or MRI (Brankov *et al* 2002), or from a reconstruction of the projection data using an alternate but fast reconstruction method, such as FBP (Yang *et al* 2002). Given such prior data, mesh points are generated using the Floyd–Steinberg error diffusion algorithm as in (Yang *et al* 2003, Brankov *et al* 2004).

In our studies, we obtained the edge map from the true activity image using an implementation of the Canny edge detector (Monga *et al* 1991). Next, the mesh points were generated from this edge map using a modified 3D version of the Floyd–Steinberg error diffusion algorithm. For brevity, we henceforth term points generated in this manner as FS points. The main features in our implementation of the error diffusion process are as follows.

- i. The prior image is traversed row-wise in a serpentine order. This ensures left–right symmetry of the error diffusion process.
- ii. For left to right traversal, the error at a non-boundary voxel having coordinates (i, j, k) is diffused to the four neighboring points $(i + 1, j, k)$, $(i + 1, j + 1, k)$, $(i, j + 1, k)$ and $(i, j, k + 1)$ with weights $1/3.5$, $1/7.0$, $1/3.5$ and $1/3.5$, respectively. For right to left traversal, the four destination voxels are $(i - 1, j, k)$, $(i - 1, j + 1, k)$, $(i, j + 1, k)$ and $(i, j, k + 1)$ with the same weights as shown above. Note that the left–right and right–left diffusion directions are mirror images of each other.
- iii. For a voxel, (i, j, k) , located on the right boundary, the error is diffused to the three neighboring voxels $(i, j + 1, k)$, $(i - 1, j + 1, k)$ and $(i, j, k - 1)$ with weights $2/5$, $1/5$ and $2/5$, respectively. The left boundary case can be obtained by reflection.
- iv. For the lower boundary case, if traversal is from left to right, then the error at voxel (i, j, k) is diffused to the two voxels $(i + 1, j, k)$ and $(i, j, k + 1)$ with equal weights $1/2$ and $1/2$. If the lower boundary traversal is from right to left, the destination voxels, as before, are obtained using reflection.

We also investigated the effects of placing FS points at sub-voxel locations. To achieve this, the activity image is first interpolated to a larger grid having a higher spatial sampling frequency, f_{pr} . The error diffusion algorithm is then used to generate points on this image, and the integer voxel coordinates obtained for the FS points, multiplied by $1/f_{pr}$, to get the fractional coordinates on the original grid.

3.4. Histogram-prior method

This multi-stage point addition strategy starts the tetrahedral reconstruction with a set of FS points obtained from an image prior as described in section 3.3. Additional points are then introduced using the histogram method (section 3.2). The FS points ensure that information of important image detail from the prior is included in the tetrahedral reconstruction, while the histogram-based method allows point placement at arbitrary locations.

3.5. Special region of interest method

In this method, in addition to points generated by any one of the four previous point generation schemes, additional points are placed in a special region of interest (SROI). The

size and shape of this SROI are user specified. In our implementation, we allow either a (i) spherical, a (ii) cylindrical or a (iii) cubical shape region. Given the SROI, points within it are placed on a regular grid at a spatial sampling frequency f_{sr} , which also corresponds to the number of sample points per voxel. Note that this grid is limited to the space enclosed by the SROI. For this study, the *histogram-prior* method was then used to generate points for the entire reconstruction domain. We evaluate this strategy to determine if the activity of an otherwise *faint* lesion can be recovered by the tetrahedral reconstruction algorithm with increased sampling in that region.

4. Performance metrics

To access the quality of reconstructions for the various strategies, we evaluated voxel- and template-based measures of (i) *bias*, (ii) *variance*, (iii) *fractional error* and (iv) *recovery coefficient* for a low-contrast feature of interest. Note that all metrics were computed after converting the tetrahedral representation to a voxelized representation on a regular grid of size $64 \times 64 \times 64$ using trilinear interpolation. The interpolation was performed by sampling each voxel at $10 \times 10 \times 10$ points and the final voxel value computed as the discrete volume integral of these point values.

Let $\mathbf{I}_T$ and $\mathbf{I}^k$ respectively denote the discrete true activity data vector and the vector representing the reconstruction of the k th noisy realization in discrete representation. Let Ω_m be the m th region over which we wish to compute the performance metrics. Note that Ω_m can be represented as a set of voxel indices of voxels in this region. Further, let N_m be the number of voxel elements in Ω_m and R the number of reconstructed noise realizations.

4.1. Voxel metrics

The voxel-based performance metric for the region Ω_m is computed for each voxel in that region and the mean taken to obtain the metric estimate for that region. The mean voxel bias β_{vox} , voxel variance σ_{vox}^2 and voxel fractional error ε_{vox} are given by

$$\beta_{\text{vox}}(\Omega_m) = \frac{1}{N_m} \sum_{j \in \Omega_m} |\mathbf{I}_T[j] - \bar{\mathbf{I}}[j]|, \quad (3)$$

$$\sigma_{\text{vox}}^2(\Omega_m) = \frac{1}{N_m(R-1)} \sum_{i=1}^R \sum_{j \in \Omega_m} (\mathbf{I}^i[j] - \bar{\mathbf{I}}[j])^2, \quad (4)$$

$$\varepsilon_{\text{vox}}(\Omega_m) = \frac{1}{N_m R} \sum_{i=1}^R \sum_{j \in \Omega_m} \frac{|\mathbf{I}^i[j] - \mathbf{I}_T[j]|}{\mathbf{I}_T[j]}. \quad (5)$$

Here, $\mathbf{I}[j]$ is the j th voxel of the image vector $\mathbf{I}$, etc. Note that the voxel-based metrics are the measures of the fidelity with which voxel values are recovered in the region of interest.

4.2. Template metrics

The template-based performance metrics are motivated by the fact that, in quantitative tasks involving spatially or temporally distributed correlated random variables, a quantity of interest at a specific spatial/temporal location is usually estimated as a weighted integral

over coordinate space. The weight function, typically, assumes the form of a symmetric and transitionally invariant kernel/template, $k(\mathbf{x})$, over this space. If the correlation among the random variables can be expressed as a random function, $f(\mathbf{x})$, the template estimate is then given by

$$t(\mathbf{x}_0) = \int k(\mathbf{x} - \mathbf{x}_0) f(\mathbf{x}) d\mathbf{x}. \quad (6)$$

Here, $t(\mathbf{x}_0)$ is the template estimate of the random function at location $\mathbf{x}_0$. In the discrete case, the kernel is represented as a matrix and the template estimate at a discrete location index, i , is given by

$$t[i] = \sum_j k_{ij} f[j - i]. \quad (7)$$

For our analysis, we have chosen the template estimate to be the mean computed at a given image location over a cube-shaped region. The kernel then assumes the form of a cube-shaped template with uniform weights. Let Ω_m represents a cube having sides of length l_m , centered at a voxel with coordinates (x_m, y_m, z_m) . Note that Ω_m can be represented as a set of the voxel indices of the voxels in the cube. It follows that the values of the averaging kernel are all equal to $1/N_m$, where N_m is the number of voxels in the m th cube representing this kernel, or equivalently the size of the set of voxel indices, Ω_m . The value of the corresponding template for a given image vector, $\mathbf{I}$, is then given by

$$t^m = \frac{1}{N_m} \sum_{j \in \Omega_m} \mathbf{I}[j]. \quad (8)$$

Now let t_T^m and t_j^m , respectively, denote the value of the template in the m th cube, computed on the true image vector and on the j th noisy reconstruction. The mean, $\bar{t}^m$, of these noisy templates are given by

$$\bar{t}^m = \frac{1}{R} \sum_{j=1}^R t_j^m. \quad (9)$$

Assuming that the noise in the reconstructed dataset is conditionally independent given the mean of the reconstructions, the template bias β_{tmp} , template variance σ_{tmp}^2 and template error ϵ_{tmp} are given by the expressions

$$\beta_{\text{tmp}}(\Omega_m) = |t_T^m - \bar{t}^m| \quad (10)$$

$$\sigma_{\text{tmp}}^2(\Omega_m) = \sum_i k_{ii}^m \quad (11)$$

$$\varepsilon_{\text{tmp}}(\Omega_m) = \beta_{\text{tmp}}(\Omega_m) / t_T^m. \quad (12)$$

Here, K^m is the noise covariance matrix computed over the vector of mean deviates in region Ω_m (Gray and Davisson 2004).

4.3. Recovery coefficient

An often cited measure for quantitative tasks, the recovery coefficient ρ , is defined as the ratio of the observed to the true activity of the image (Saha 2005, Knoop *et al* 2002, Jentzen *et al* 2010, Srinivas *et al* 2009, Miller *et al* 1991). In practice, for a specified image region, Ω_r , it is computed as the ratio of the values of the reconstructed data, $\mathbf{I}$, to the value of the true data, $\mathbf{I}_T$, measured at a point, $\mathbf{x}_m$, that corresponds to the maximum of the true data value in that region:

$$\rho(\Omega_r) = \frac{\mathbf{I}[m]}{\mathbf{I}_T[m]}, \quad (13)$$

where $m = \arg \max_{j \in \Omega_r} \mathbf{I}_T[j]$.

5. Study setup

An analytical parallel projector was used to obtain noise-free, exact projection data of a simulated object. This object consists of a set of ten spheres of different radii and different activities. Within each sphere the activity distribution was constant. Due to the overlap of some of the spheres, the object has a relatively complex geometry and activity distribution. Figure 1 shows slices from the true activity object and table 1 shows the location of the center of the spheres, their radii and activities.

To obtain the true image in the voxel representation which was used for the calculation of fidelity measures, the analytically defined object was sampled to create an image of size $512 \times 512 \times 512$ and re-binned using averaging to an image having size $64 \times 64 \times 64$. Thus, the true image had the same size as the images reconstructed using the voxel-based approach. In order to create noiseless parallel projection data, line integrals were then calculated for each of the 64 projections over 360° creating a (512×512) high-resolution projection image. Each of the projection images was then re-binned using simple averaging to 64×64 resolution projection images. Twenty-five noisy realizations were then obtained by adding Poisson distributed noise to this projection data and scaled so that the resulting projection data had approximately 8.0 million counts. These were reconstructed using the different point placement strategies with point clouds and standard voxel-based image reconstruction approach.

Figure 2 shows five different regions in the image at which the voxel- and template-based performance metrics are evaluated. These regions are labeled as S1, S2, S3, S4 and S5. Four of the regions, S1–S4, are placed at the center of spheres that have a higher activity compared to the background. The fifth region, S5, is a location in the background. The voxel and template metrics were computed over cube-shaped regions with the center of the cubes being placed at the center of these spheres. Table 2 lists the location and side-lengths of these cubes. Note that the dimensions of the cubes for the voxel and template metrics are different. They were chosen such that, for the voxel metrics, they enclosed the spherical activity regions S1–S4 but were entirely within these regions for the template metrics. Figure 3(a) shows a set of prior points obtained using the Floyd–Steinberg 3D error

diffusion algorithm on the true activity image. Figure 3(b) shows a set of final reconstruction points obtained using the tetrahedral reconstruction method with the *histogram-prior* point placement strategy. The prior used for this reconstruction had 10 000 FS points and the *histogram-prior* strategy added 10 000 points over six stages.

6. Results and discussions

The voxel and template performance metrics for each point placement strategy were computed for tetrahedral reconstructions corresponding to different number of mesh points. Table 3 shows the number of points corresponding to each strategy at which the reconstructions were evaluated. However, in the plots shown in this section (figures 4– 6), we only present the best performance curves for each strategy. Accordingly, for a given strategy, the best bias–variance curve was determined by visual inspection to be the curve closest to the origin, while the best mean fractional error curve was selected to be the one that achieved the minimum. Note that, in the bias–variance plots, we actually show the plots of the squared bias versus the variance. This was done in consideration of that fact that the total error is composed of the squared bias and the variance term. Hence, with our choice the plot axis has the same units. Also note that the SROI-based point placement strategy was performed for region S4 only. This is because, as S4 represents a low-contrast region, recovering its activity with statistical fidelity is challenging for both tetrahedral and regular grid reconstruction methods. This region is therefore a good candidate to determine the effectiveness of the SROI point placement strategy. For each tetrahedral reconstruction strategy, the performance metrics were computed at 30, 90, 180, 300 and 600 iterations, while for the regular grid reconstructions, they were computed at 10, 30, 90, 180 and 300 iterations.

Figures 4(a)–(e) are the voxel-based bias–variance plots for the tetrahedral and regular grid reconstruction methods for regions S1–S5, respectively. Figures 5(a)–(e), similarly, are the plots of the template bias–variance curves for these regions. Note that, for the tetrahedral case, the number of points corresponding to the best performance curve is different in the different regions. For example, in region S1, the best bias–variance performance in the voxel metric is achieved by 20K points for the sorted variation strategy, by 100K points for the prior strategy and by 175K points for the histogram-prior strategy (figure 5(a)). Also note that for a given region, the reconstruction with the highest number of points does not necessarily achieve the least bias. The point placement strategy that achieves overall best bias–variance performance is also region dependent. From the voxel and template bias–variance curves, it is apparent that the tetrahedral reconstruction method out-performs the regular grid reconstruction in all regions. Also note that, for the regular grid reconstructions, the variance increases rapidly with increasing iteration numbers for both the voxel- and template-based metrics.

Figures 6(a), (c) and (e) are the plots of the voxel mean fractional error for regions S1, S2 and S4, respectively, while figures 6(b), (d) and (f) are the corresponding plots for the template mean fractional error. The fractional error plots corresponding to regions S3 and S5 have been omitted, as they are similar to the plots shown here. The error plots indicate that, though the minimum mean fractional are quite similar for the various reconstruction strategies, for the regular grid reconstruction method it increases rapidly with increasing iteration numbers. This is to be expected, since variance is the dominating source of error and increases rapidly for the regular grid reconstructions.

Tables 4 and 5 respectively list the voxel- and template-based estimates of the mean fractional bias and mean fractional error for the tetrahedral and regular grid reconstructions. These values were computed for regions S1–S5 (figure 2). For the tetrahedral case, we have

listed only the best performing strategies along with their corresponding reconstruction parameters. From table 4, note that the mean voxel error metrics for the tetrahedral and regular grid methods are similar in regions S1 and S3; however, the tetrahedral method is superior to the regular grid reconstruction in regions S2, S4 and S5. From table 5, note that the template fractional error metrics for the tetrahedral reconstructions are significantly smaller in magnitude compared to the regular grid reconstructions for all regions S1–S5. Also note that the template error metrics, overall, are significantly smaller than the voxel error metrics for both reconstruction methods.

Figure 7(a) shows slice images from the true activity image, while figures 7(b)–(f) show the same slices of the mean reconstructed data corresponding to the different reconstruction strategies. Note that for the tetrahedral case, the *histogram-prior* and the *prior* point placement strategy achieve the best boundary demarcation between sub-structures in the reconstructed image, while the *sorted-variance* (figure 7(c)) strategy is inferior to every other strategy in this regard. The bias–variance and error curves, however, indicate that the *sorted-variance* strategy can achieve comparable performance to other strategies in some regions. For example, in region S2, it out-performs the other strategies (figure 4(b)). This illustrates some of the shortcomings of bias–variance type of analysis to evaluate overall image quality. Note that the *prior* (figure 7(d)) method achieves comparable or even better performance in some regions, compared to other tetrahedral reconstruction methods. However, on closer inspection, we have observed that it has reconstruction artifacts similar to the images obtained from the regular grid reconstructions.

Figure 8 is a plot of the mean recovery coefficients for the different point placement strategies. These were computed over region S4 and correspond to reconstruction parameters that give the minimum mean voxel error for that region in the reconstructed images. To determine the significance of the differences in the mean recovery coefficient, we performed a one-way analysis of variance (ANOVA) (Faraway 2005) using the recovery coefficient values of each of the 25 noisy reconstructed datasets for each strategy. The analysis establishes an overall difference among the strategies at the ($p \ll 0.0001$) level of significance. To assess the pairwise difference in means between the strategies, we performed Tukey's HSD (Crawley 2007) test. Figure 9 is a plot of the 99% confidence interval obtained for the different pairs. According to this plot, if an interval does not contain zero, then the pairwise difference between the mean values of the recovery coefficients are significant at the $p = 0.001$ level of confidence. This plot shows that, while the regular grid reconstruction has better activity recovery compared to the tetrahedral method that uses the sorted variation, the histogram and the histogram-prior point placement strategies, it is significantly inferior to the prior, as well as the SROI point placement strategies. This illustrates the fact that the multi-resolution tetrahedral reconstruction method can achieve superior activity recovery in low-contrast regions.

Finally, figure 10(a) shows a slice from the true activity image that contains the maximum count in region S4, while figure 11 shows the activity profile on the line drawn through this image. Figures 10(b)–(d) show images of the same slice of the data reconstructed with the *prior* strategy, the *SROI* strategy and the *regular grid* strategy, respectively, while figures 11(b)–(d) show the corresponding activity profiles through the same line. From these images and plots we observe that the SROI-based point placement strategy achieves the better activity recovery of the low-contrast region S4. Note, however, that for the template metric case, the *histogram-prior* method achieves the best performance in region S4 (figure 5(d)).

7. Conclusions

We have evaluated a tomographic reconstruction method for a point cloud tetrahedral mesh reconstruction algorithm and compared it with the regular grid reconstruction method. The evaluation was performed using voxel- and template-based bias, variance and fractional error metrics and activity recovery, using the recovery coefficient. For the tetrahedral method, we evaluated several point placement strategies. Our analysis shows that these strategies have significant consequences on image quality as measured by the performance metrics.

We have seen that, for the tetrahedral reconstruction method, parameters corresponding to the best figure of merit for a given performance metric are region dependent and that the relation between model complexity, which in the tetrahedral reconstruction case is the number of points used, and the best figure of merit, is not a simple one. For regular grid reconstructions, an increase in the number of grid points results in a decrease in bias and an increase in variance. This behavior does not hold for the tetrahedral reconstructions. The bias–variance curves illustrate this fact. Our results strongly suggest that different diagnostic tasks require different point generation strategies for optimal performance.

For tetrahedral reconstructions, the use of a prior can enhance overall quality of reconstructions by improving boundary demarcation and ensuring that regions with significant details are appropriately sampled. However, points from a prior alone are not always sufficient and better reconstructions are obtained when points can be placed adaptively and at arbitrary locations in the image space. Specification of an SROI, in which points are placed more densely than in other regions, can result in superior activity recovery, especially when such regions have low-contrast features of interest. This might prove beneficial in a more realistic imaging situation when trying to recover a low-contrast lesion. A potential benefit of this could be in the development of an *interactive* image reconstruction application in which the user marks a suspicious region of an image for reconstruction with increased sampling.

For both the voxel- and template-based performance metrics we have noted that the tetrahedral reconstruction strategy is superior to reconstructions on a regular grid. And in both cases, the variance for the regular grid reconstructions increases rapidly with increasing iteration numbers. This observation is crucial, since in a realistic setting the optimal number of iterations for the ML-EM algorithm is not known *a priori*. As acquired data in nuclear medicine imaging are Poisson noise limited, this rapid increase in variance for the regular grid reconstruction method can result in significant uncertainty, especially in detection/estimation of low-contrast features. The variance of tetrahedral reconstructions, on the other hand, appears to be more stable in this regard. This suggests that the sparse grid reconstructions are better conditioned than the regular grid ones.

Due to the more elaborate image representation model employed, an appropriate tetrahedral reconstruction strategy can produce images with fewer reconstruction artifacts compared to reconstruction on a regular grid. However, for the tetrahedral method to achieve these features, it is important to have a reasonable prior image. This, however, is not a serious shortcoming as such priors can be obtained using multi-modality imaging, or from a fast reconstruction method that employs a less complex image model.

The multistage tetrahedral reconstruction strategies require that the system matrix be computed at each stage. This can significantly increase the total reconstruction time, especially when using single core processors or non-parallel methods on a multi-core processors. However, with the advent of GPU-enabled parallelization techniques in the field of medical imaging, this shortcoming can be overcome and we are working on

implementing a parallel version of this algorithm that utilizes such multiprocessor capabilities.

Finally, we note that an analysis involving bias–variance/error analysis is essentially limited, in the sense that it gives performance metrics that are pertinent to quantitative imaging and do not necessarily correlate with qualitative imaging tasks such as perceptual detection of lesions, defects, etc. We plan on pursuing such investigations in the future.

Acknowledgments

This work was supported in part by the National Cancer Institute, National Institutes of Health under grant R21 CA123057 and American Heart Association, Scientist Development Grant, grant 0735328N. The authors thank Jerry Xu and Dennis Gurgul of Partners High Performance Research Computing for providing support on the high performance cluster.

References

- Brankov JG, Yang Y, Wernick MN. Tomographic image reconstruction based on a content-adaptive mesh model. *IEEE Trans. Med. Imaging* 2004;23:202–212. [PubMed: 14964565]
- Brankov, JG.; Yang, Y.; Leahy, RM.; Wernick, MN. Multi-modality tomographic image reconstruction using mesh modeling. *Proc. of the 2002 IEEE Intl. Symp. on Biomedical Imaging*; 7–10 June 2002; Ritz-Carlton Hotel, Washington, DC, USA. 2002. p. 405-408.
- Crawley, MJ. *The R Book*. Chichester: Wiley; 2007.
- Delauney BN. Sur la sphère vide. *Bulletin of Academy of Sciences of the USSR* 1934;6:793–800.
- Floyd R, Steinberg L. An adaptive algorithm for spatial gray scale. *Proc. Soc. Image Display* 1976;17:75–77.
- Faraway, JJ. *Practical Regression and ANOVA using R*. London: Chapman and Hall; 2005. Available at <http://www.maths.bath.ac.uk/jjf23/book/>
- Garcia, MA. Efficient approximation of range images through data-dependent adaptive triangulations; 1997 IEEE Computer Society Conf. on Computer Vision and Pattern Recognition CVPR '97; 1997. p. 628-633.
- Gray, RM.; Davisson, LD. *An Introduction to Statistical Signal Processing*. Cambridge: Cambridge University Press; 2004.
- Jentzen W. Experimental investigation of factors affecting the absolute recovery coefficients in iodine-124 PET lesion imaging. *Phys. Med. Biol* 2010;55:2365–2398. [PubMed: 20360631]
- Knoop BO, Geworski L, Hofmann M, Munz DL, Knapp WH. Use of recovery coefficients as a test of system linearity of response in positron emission tomography (PET). *Phys. Med. Biol* 2002;47:1237–1254. [PubMed: 12030553]
- Lange K, Carson R. EM reconstruction algorithm for emission and transmission tomography. *J. Comput. Assist. Tomogr* 1984;8:306–316. [PubMed: 6608535]
- Monga O, Deriche R, Malandian G, Cocquerez JP. Recursive filtering and edge tracking: two primary tools for 3-d edge detection. *IEEE Image Vision Comput* 1991;4:203–214.
- Miller TR, Wallis JW, Grothe RA. Design and use of PET tomographs: the effect of slice spacing. *J. Nuclear Med* 1991;32:551–552.
- Sitek A, Huesman RH, Gulberg GT. Tomographic reconstruction using an adaptive tetrahedral mesh defined by a point cloud. *IEEE Trans. Med. Imaging* 2006;25:1172–1179. [PubMed: 16967802]
- Sandler EA, Gusev DA, Milman GY. Hybrid algorithms for digital halftoning and their applications to medical imaging. *Comput. Graph* 1997;21:69–78.
- Shepp LA, Vardi Y. Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imaging* 1982;1:113–122. [PubMed: 18238264]
- Si, H.; Gartner, K. An algorithm for three-dimensional constrained Delauney tetrahedralizations. *Proc. of 4th Intl. Conf. on Engineering Computational Technology*; Lisbon. 2004.
- Si H. TetGen version 1.3. 2005 Available at <http://tetgen.berlios.de>.

- Saha, GB. Basics of PET Imaging: Physics, Chemistry, and Regulations. 2nd edn. New York: Springer; 2005.
- Srinivas S, Dhurairaj T, Bural G, Surti S, Alavi A. A recovery coefficient method for partial volume correction of PET images. *Ann. Nucl. Med* 2009;23:341–348. [PubMed: 19367446]
- Yang Y, Wernick MN, Brankov JG. A fast approach for accurate content-adaptive mesh generation. *IEEE Trans. Image Process* 2003;8:866–881. [PubMed: 18237961]
- Yang Y, Brankov JG, Wernick MN. Content-adaptive mesh modeling for fully-3D tomographic image reconstruction. *ICIP* 2002;2:621–624.
- Yu B, Lu Y, Vogelsang L, Krol A, Xu Y, Feiglin D. Optimized mesh generation for OSEM SPECT reconstruction. *J. Nucl. Med* 2009;50 Suppl. 2:407.

Figure 1. Cuts through the true activity distribution of the simulated object consisting of ten spheres, each containing different activities. Some of the spheres overlap so the final object has a relatively complex activity distribution.

Figure 2.

Two slices of the true activity image voxelized to $64 \times 64 \times 64$ with marked regions of interest over which the performance metrics are computed. The exact locations of these regions are given in table 2.

Figure 3.

(a) Points obtained using a 3D Floyd–Steinberg error diffusion algorithm on the true activity image. Ten thousand points were generated for this illustration. (b) Final reconstruction points using the *histogram-prior* point placement strategy. The reconstruction was started using 10 000 FS points (a) and added 10 000 additional points during in six stages.

Figure 4.

Mean squared-bias versus variance curves for regions (a) S1, (b) S2, (c) S3, (d) S4 and (e) S5 in voxel measure. For each tetrahedral point placement strategy, the marker points correspond to 30, 90, 180, 300 and 600 iterations of the ML-EM reconstruction algorithm, while for the regular grid method, they correspond to 10, 30, 90, 180 and 300 iterations. The iteration numbers increase from left to right along the curves. For a given region, the bias and variance were computed for each voxel in the cube region. The mean of the voxel bias and the voxel variance was taken as an estimate for the bias and variance of the region.

Figure 5.

Mean squared-bias versus variance curves for regions (a) S1, (b) S2, (c) S3, (d) S4 and (e) S5 in template measure. For each tetrahedral point placement strategy, the marker points correspond to 30, 90, 180, 300 and 600 iterations of the ML-EM reconstruction algorithm, while for the regular grid method, they correspond to 10, 30, 90, 180 and 300 iterations. The iteration numbers increase from left to right along the curves. For a given region, the bias and variance were computed for each voxel in the cube region. The mean of the voxel bias and the voxel variance was taken as an estimate for the bias and variance of the region.

Figure 6. Voxel mean fractional error versus iteration number in regions (a) S1, (c) S2, (e) S4 and template mean fractional error in regions (b) S1, (d) S2 and (f) S4.

Figure 7.

(a) True activity image slices and corresponding slices of images reconstructed with (b) sorted variation strategy, (c) histogram strategy, (d) prior strategy, (e) histogram-prior strategy and (f) regular grid strategy. They correspond to the reconstructions, in the voxel-based bias–variance sense of the corresponding strategies.

Figure 8.
Mean values of the recovery coefficients computed over region S4.

Figure 9.

Ninety-nine percent confidence interval for the difference between mean values of the recovery coefficients computed over region S4 for the different reconstruction strategies. If this interval contains zero, then the difference in means between the corresponding pair of strategies are insignificant at the $p = 0.001$ level. For each tetrahedral point placement strategy, parameters corresponding to the minimum mean voxel error for region S4 were chosen for the reconstructions.

Figure 10. Slice containing region S4 from (a) true activity image and from data reconstructed using (b) prior strategy, (c) SROI strategy and (d) regular grid strategy. They correspond to the reconstructions, in the voxel-based bias–variance sense of the corresponding strategies.

Figure 11.

Activity profile from (a) true activity image and images reconstructed with (b) *prior*, (c) *SROI* and (d) *regular grid* reconstruction strategies. The profiles correspond to the line shown in figure 10(a).

Table 1

Location, size and activity of spheres in simulated object. Values are in units of projection pixel size.

Sphere ID	x_c	y_c	z_c	Radius	Activity vol
1	32	32	32	30	1
2	32	32	24	8	-1
3	32	24	40	2	1
4	32	40	40	2	5
5	24	32	40	1	1
6	40	32	40	1	5
7	32	24	24	8	1
8	32	40	24	8	5
9	24	32	24	4	1
10	40	32	24	4	10

Table 2

Center and side length of cubical regions of interest. Values are in units of projection pixel size. The size of the reconstruction volume is $64 \times 64 \times 64$ voxels.

Sphere ID	x_c	y_c	z_c	l_{vox}	l_{temp}
Ω_1	32	24	24	16	4
Ω_2	40	32	24	8	6
Ω_3	32	24	40	4	4
Ω_4	24	32	40	4	6
Ω_5	52	32	24	10	6

Table 3

Parameters used in this work for the point cloud generation strategies. The size of the reconstruction volume is $64 \times 64 \times 64$ voxels.

<i>D</i> = 6 for multi-stage strategies									
Hist-prior			Prior			SROI			
<i>N</i> _{hist}	<i>N</i> _{pr}	<i>f</i> _{pr}	Hist <i>N</i> _{hist} (<i>N</i> ₀ = 1000)	Sort <i>N</i> (<i>N</i> ₀ = 1000)	<i>N</i> _{pr}	<i>f</i> _{pr}	<i>N</i> _{hist}	<i>N</i> _{pr}	<i>f</i> _{sr}
30K	26K	1	10K	5K	10K	1	100K	26K	1
60K	26K	1	15K	7K	20K	1	100K	26K	2
100K	10K	1	30K	10K	30K	1			
100K	26K	1	40K	20K	30K	2			
150K	10K	1	100K	30K	60K	2			
150K	26K	1	150K	40K	60K	4			
				80K	80K	2			
					80K	4			
					100K	2			
					100K	4			

*N*_{hist} →number of histogram points, *N*_{pr} →number of prior points, *f*_{pr} →prior sampling frequency, *f*_{sr} →SROI sampling frequency.

Comparison of voxel mean-fractional bias and fractional error for irregular and regular grid reconstructions. Only the best strategies are shown for the irregular grid reconstructions.

Table 4

Tetrahedral grid						Regular grid (N = 262, 144)			
Region	Strategy	N _{tot}	Iter	% Bias	% Error	Iter	% Bias	% Error	
S1	Hst-Pr	176K	300	0.397	0.120	30	0.511	0.119	
S2	Hst-Pr	160K	600	0.533	0.083	90	0.450	0.116	
S3	Hst-Pr	126K	300	0.066	0.126	30	0.061	0.126	
S4	SROI	126K	300	0.026	0.088	30	0.029	0.117	
S5	Prior	100K	180	0.021	0.031	30	0.024	0.102	

Comparison of template mean-fractional bias and fractional error for irregular and regular grid reconstructions. Only the best strategies are shown for the irregular grid reconstructions.

Table 5

Tetrahedral grid						Regular grid (<i>N</i> = 262, 144)			
Region	Strategy	<i>N</i> _{tot}	Iter	% Bias	% Error	Iter	% Bias	% Error	
S1	Sort	20K	600	0.001	0.007	30	0.020	0.019	
S2	Prior	80K	600	0.007	0.008	30	0.006	0.016	
S3	Hst-Pr	126K	600	0.004	0.014	30	0.013	0.031	
S4	Hst-Pr	176K	600	0.009	0.014	10	0.020	0.021	
S5	Prior	20K	180	0.005	0.009	30	0.005	0.019	