

Three-dimensional reconstructions from low-count SPECT data using deformable models

G.S. Cunningham and K.M. Hanson

MS P940, Los Alamos National Laboratory, Los Alamos, NM 87545 USA

cunning@lanl.gov and kmh@m4next.lanl.gov

X.L. Battle

Faculte de Medecine, Universite de Bretagne Occidentale, 29200 Brest, FR

Xavier.Battle@univ-brest.fr

Abstract: We demonstrate the reconstruction of a 3D, time-varying bolus of radiotracer from first-pass data obtained by the dynamic SPECT imager, FASTSPECT, built by the University of Arizona. The object imaged is a CardioWest Total Artificial Heart. The bolus is entirely contained in one ventricle and its associated inlet and outlet tubes. The model for the radiotracer distribution is a time-varying closed surface parameterized by 162 vertices that are connected to make 960 triangles, with uniform intensity of radiotracer inside. The total curvature of the surface is minimized through the use of a weighted prior in the Bayesian framework. MAP estimates for the vertices, interior intensity and background count level are produced for diastolic and systolic frames, the only two frames analyzed. The strength of the prior is determined by finding the corner of the L-curve. The results indicate that qualitatively pleasing results are possible even with as few as 1780 counts per time frame (total after summing over all 24 detectors). Quantitative estimates of ejection fraction and wall motion should be possible if certain restrictions in the model are removed, e.g., the spatial homogeneity of the radiotracer intensity within the volume defined by the triangulated surface, and smoothness of the surface at the tube/ventricle join.

©1998 Optical Society of America

OCIS codes: (100.3010) Image reconstruction techniques; (100.6960) Tomography

References and links

1. W.P. Klein, H.H. Barrett, I. W. Pang, D.D. Patton, M.M. Rogulski, J.J. Sain, and W. Smith "FASTSPECT: Electrical and mechanical design of a high resolution dynamic SPECT imager," *Conference Record of the 1995 IEEE Nucl. Sci. Symp. & Med. Imaging Conf.* (IEEE, Los Alamitos, 1996), Vol. 2, pp. 931-933.
2. K.M. Hanson, "Bayesian reconstruction based on flexible prior models," *J. Opt Soc. Amer. A* **10**, pp. 997-1004 (1993).
3. Y. Bresler, J.A. Fessler, and A. Macovski, "A Bayesian approach to reconstruction from incomplete projections of a multiple object 3D domain," *IEEE Trans. Pattern Anal. Mach. Intell.* **11**, pp. 840-858 (1989).
4. P.C. Chiao, W.L. Rogers, N.H. Clinthorne, J.A. Fessler, and A.O. Hero, "Model-based estimation for dynamic cardiac studies using ECT," *IEEE Trans. Med. Imaging* **13**, pp. 217-226 (1994).
5. X.L. Battle, G.S. Cunningham, and K.M. Hanson, "3D tomographic reconstruction using geometrical models," *Proc. SPIE* (SPIE, Bellingham, 1996), Vol. 3034, pp. 346-357.
6. X.L. Battle, G.S. Cunningham, and K.M. Hanson, "Tomographic reconstruction using 3D deformable models," to appear in *Phys. Med. Biol.*, 1998.
7. <http://www.radiology.arizona.edu/~fastspec/detectors.html>
8. M. Kass, A. Witkin, and D. Terzopolous, "Snakes: active contour models," *Int. J. Comput. Vis.*, pp. 321-331 (1988).

9. K.M. Hanson and G.S. Cunningham, "A computational approach to Bayesian inference," M.M. Meyer and J.L. Rosenberger, eds., *Computing Science and Statistics* (Interface Foundation, Fairfax Station, VA 22039-7460, 1996) Vol. 27, pp. 202-211.
 10. K.M. Hanson, R.L. Bilisoly, and G.S. Cunningham, "Kinky tomographic reconstruction," *Proc. SPIE* (SPIE, Bellingham, 1996), Vol. 2710, pp. 156-166.
 11. D.J.C. MacKay, "Bayesian interpolation," *Neural Comput.* **4**, pp. 415-447 (1992).
 12. P.C. Hansen and D.P. O'Leary, "The use of the L-curve in the regularization of discrete ill-posed problems," *SIAM J. Sci. Comput.* **14**, pp. 1487-1503 (1993).
-

1. Introduction

The FASTSPECT imaging system [1], developed at the University of Arizona, has been used for first-pass tomographic imaging of the time-varying distribution of a bolus of Tc-99m pertechnetate radiotracer infused into a CardiacWest Total Artificial Heart. The FASTSPECT machine simultaneously provides 24 pinhole views of the bolus distribution evolving in time, and is unique in its ability to perform this type of dynamic imaging. The goal in obtaining first-pass tomographic data is to demonstrate that clinically important measures of heart function, such as ejection fraction and wall motion, can be quantitatively estimated without having to gate and average over many cardiac cycles, an approach necessarily utilized by single- or dual-head cardiac SPECT systems. If ejection fraction and wall motion can be estimated from first-pass data during the first few cardiac cycles, then later cycles can be used to estimate myocardial perfusion, another important indicator of heart function. If successful, the FASTSPECT approach would mean that a single, relatively cheap instrument could perform multiple diagnostic tests of cardiac function with a single bolus of radiotracer. This type of capability would be clinically valuable and affordable for use in emergency rooms across the country to do initial assessment of cardiac patients.

A traditional approach to reconstruction of the 24-view tomographic data might employ the EM method to produce the maximum likelihood estimate of the activity in each voxel of the volume being imaged. A voxel-based reconstruction can be further processed by manual or automated segmentation to yield an estimate of an isosurface of the radiotracer distribution. If the radiotracer distribution is homogeneously mixed throughout a ventricular chamber of interest, then the time-sequenced estimated isosurface yields an estimate of the ventricular volume as a function of time, and hence ejection fraction, but it also provides a great deal more information of potential clinical value, since the entire interior surface of the ventricle is revealed. It is the dynamic information about the interior surface of the ventricle that distinguishes FASTSPECT and its potential follow-up (a clinical device) from other cardiac SPECT implementations, wherein time-averaging over different cardiac cycles is needed to provide similar information.

The principal attraction of the traditional approach to tomographic reconstruction and segmentation is that it is fast, and provides nearly optimal solutions when used in medical imaging modalities, e.g., PET and MRI, where the number of views and SNR are quite high. However, for data with very low SNR and including only a limited number of views, the traditional approach will perform very poorly at estimating the surface that defines the spatial extent of the radiotracer distribution. This is due to the fact that the non-parametric, voxel-based model used in the tomographic reconstruction is underconstrained. Surface estimates performed in this way will be unacceptably noisy for the type of data analyzed in this article. For such sparse and noisy data, an approach that directly estimates the shape parameters of a time-varying surface from the raw projection data is advantageous, although potentially very time-consuming.

We discuss the direct estimation approach in this article and apply it for the first time to real data obtained by FASTSPECT. We formulate a Bayesian estimation problem for first-pass tomographic imaging using FASTSPECT that directly estimates the time-varying (x,y,z) components of the vertices of a triangulated surface within which it is assumed the bolus is uniformly distributed [2-6]. The estimation is performed by doing a quasi-Newton minimization of the minus-log posterior over a 486-dimensional parameter space. The

computation time associated with our direct estimation approach is comparable to the time required by the traditional approach, due to the fact that we utilize the adjoint differentiation technique to compute search directions needed by the optimization.

First, we describe the FASTSPECT imaging system and the real data that were analyzed. Next we formulate the Bayesian estimation problem and, finally, we present some results and conclusions.

2. The data

2.1 FASTSPECT

FASTSPECT is a dynamic SPECT imager that has been used for brain, heart and bone imaging [1]. Two circular arrays with a total of 24 pinhole apertures are arranged on a hemispherical dome that is roughly 35 cm in diameter. The hemisphere surrounds the volume of interest. Each pinhole is mapped to an Anger detector, and the voltages from the 4 photomultipliers that see a monolithic scintillating crystal for each detector are converted to an estimate of the position of each detected photon that must lie on a 64x64 uniformly-binned image grid. If a reliable estimate of the position of a photon cannot be determined, then the photon hit is discarded. There is no explicit windowing in energy to discriminate against scattered photons. This detection system may eventually be replaced by a semiconductor-based system [7].

Pinholes of various diameter can be inserted into the dome surrounding the object volume; 2.5 mm diameter pinholes were used to generate the data analyzed in this article. The system is characterized by a matrix, \mathbf{H} , that is measured by passing a small volume element of radiotracer throughout the volume being imaged, and measuring the response of every detector pixel to that source, producing an enormous amount of information, even when compressed to take advantage of the sparsity of the matrix (150 MB of disk space after compression). The system matrix used in this article was obtained by passing a $[5\text{mm}]^3$ volume element through a 43x57x39 grid. The system matrix is noisy since only a finite number of counts are obtained for each location of the source. Given enough patience and time, though, this noise could presumably be made as low as is needed.

Note that attenuation through the dome is included in the measurement of \mathbf{H} . If information is available concerning attenuating material between the radiotracer distribution and the pinholes, it can also be incorporated into \mathbf{H} , and this was done for the \mathbf{H} used to analyze the data discussed in this article. A phantom was defined, using simple geometric shapes, that approximately characterizes the torso surrounding the artificial heart system described in the next section. The geometric shapes used to simulate the lungs were assigned an attenuation length of 0.07 cm^{-1} and the geometric shapes used to simulate the rest of the system were assigned an attenuation length of 0.15 cm^{-1} (the attenuation length of water at 140 keV). For each (object voxel, pinhole center) pair, the exponential attenuation $e^{-(0.15*L_1+0.07*L_2)}$ was computed, where L_1 =distance in cm along the line connecting the object voxel to the pinhole center which passes through the water, and L_2 =distance in cm along the same line which passes through the simulated lung material. The attenuation factor was multiplied by the \mathbf{H} entries for that voxel and all of the detector pixels associated with that pinhole to produce new \mathbf{H} entries.

2.2 The imaged object and raw data

The object that was imaged is a CardioWest Total Artificial Heart. Only the left and right ventricles, each about 120 ml in volume, were used. The lungs were simulated using bottles filled with water and styrofoam beads. A 20 mCi bolus of Tc-99m pertechnetate was injected into the input tube of the right ventricle at a site outside the field of view about 17 cm from the input valve of the right ventricle. The 20 mCi is first placed into 0.5 mL of a syringe and injected into a very small tube that leads to the input tube of the right ventricle. The bolus is eventually "flushed" into the input tube using a volume of water that is greater than the volume enclosed by the small tube. This process is thought to parallel the method of

injection for humans, in which the injection is followed by a saline flush. Ultimately, the fluid flows into and out of a Donovan mock circulatory system, which is out of the field of view.

First-pass diastolic and systolic frames were analyzed. The diastolic frame (Fig. 1a) contained a total of 2133 counts, of which about 40 appear to be inconsistent with the assumption that the radiotracer distribution is contained within the right ventricle and its associated input and output tubes. We assume that these counts are from photons that were scattered but still accepted, although other explanations may be possible. We refer to these counts as “background” counts, and will later model and estimate this background as a single constant, different in value for each time frame, but the same for all detector pixels at a given time. The systolic frame (Fig. 1c) contained a total of 1780 counts, with a similar number of background counts.




Fig. 1. Raw data for a) diastolic frame and c) systolic frame. Predicted detector Poisson rates for b) diastolic frame and d) systolic frame using MAP solution with $\alpha=0.2$.

3. The Bayesian estimation problem

We have implemented a general tool for Bayesian estimation in the context of image analysis using geometric models that we call the Bayes Inference Engine (BIE). We conceive of the Bayesian estimation problem as consisting of three parts: the object model, the measurement model, and the probability model. In the BIE, the user constructs a graphical program that transforms object and measurement system parameters into predicted data. The predicted data are compared with real data to produce a minus-log-probabilistic goal function, and an

optimizer is connected to the goal function and any parameters that are to be estimated by minimizing the goal function. See Fig. 2 for the graphical program that was used to analyze the data discussed in this article.

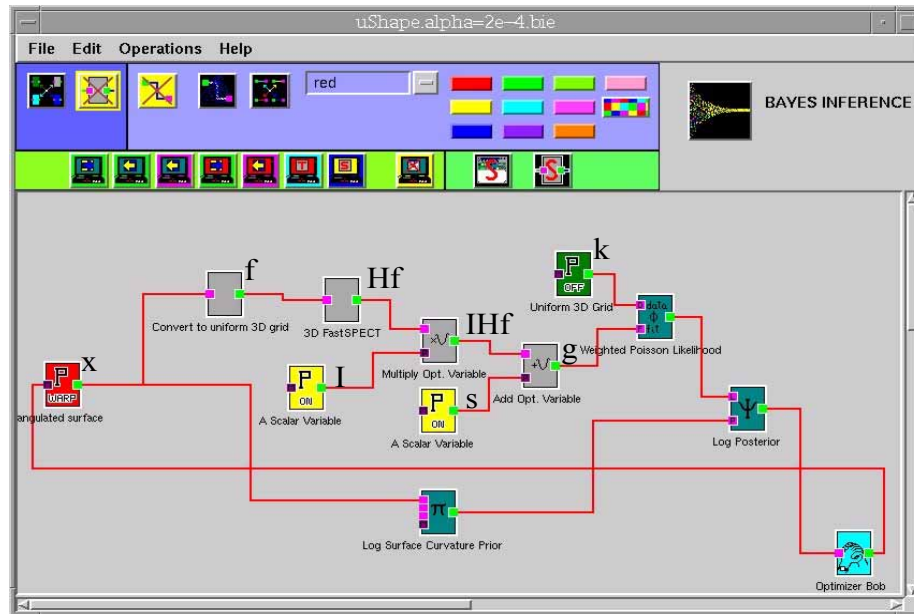


Fig. 2. BIE canvas used to analyze SPECT data. The triangulated surface box at the far left contains \mathbf{x} . It is transformed into a voxelated grid, \mathbf{f} , and then by the matrix, \mathbf{H} , to produce $\mathbf{Hf}(\mathbf{x})$. Finally it is multiplied by the intensity, I , and the additive background constant, s , is added, to produce the predicted detector pixel rates $\mathbf{g}=\mathbf{IHf}(\mathbf{x})+s$.

3.1 The object model

The object model is the parametric model of whatever spatio-temporal physical quantity we are interested in; in this case, it is the parametric model for the 3D, time-varying, radiotracer intensity distribution. In the BIE, we always convert parametric models to non-parametric ones (uniformly sampled grids) so that complex models can easily be built through combination of parametric models after conversion to a non-parametric form.

The parametric model we use here is a triangulated surface that evolves in time, defined by a set of 162 vertices with (x,y,z) components, and a connectivity network that creates 960 triangles by connecting vertices together. The 486-point vector that lists the (x,y,z) components for each of the 162 vertices is denoted \mathbf{x} . All 486 values in \mathbf{x} are estimated for each time. This parametric model is converted to a non-parametric uniformly-voxellated grid, \mathbf{f} , by setting the value of each voxel in \mathbf{f} to the fraction of that voxel that is contained within the volume described by the triangulated surface. We assume that the radiotracer is homogeneously distributed throughout the volume enclosed by the triangulated surface, so that only a single parameter, I , is needed for the activity level. Thus, \mathbf{If} is a uniformly-voxellated grid with intensity I interior to the triangulated surface defined by the vertices \mathbf{x} . Obviously, \mathbf{f} is functionally dependent on the parameters, $\mathbf{f}=\mathbf{f}(\mathbf{x})$.

More complicated models for the 3D distribution within the volume defined by the surface can be accommodated in the BIE, e.g., a voxelated grid of values with lower and upper bounds or an unconstrained voxelated grid with an associated prior penalizing high-frequency variations could be used.

Many alternative geometric models have been used in medical imaging, including implicitly-defined surfaces, snakes, and line processes. Implicitly-defined surfaces are ideal in situations where the topology of the surface is unknown, which is not the case here since we

expect the volume containing radiotracer to be path-connected (except perhaps at the valves). Similarly, reconstruction using line processes can produce disconnected lines rather than polygons, and there are issues regarding convergence to the global minimum. "Snakes" and surface snakes have been used in medical imaging, and in other situations, to segment image or volume data. The segmentation step is often preceded by a non-parametric tomographic reconstruction technique, e.g., EM or filtered backprojection, to transform the data from projection space to object space. The spatial derivative of the reconstructed image or volume is interpreted as a force that acts on the geometrical model, driving it toward spatial locations of maximum image gradient (edges). The approach adopted in this article can be formulated as a surface snake problem if the forces on the surface snake are defined to be equal to the gradient of the minus-log posterior (defined below) w.r.t. the surface parameters [8]. However, even if this is done, the update equation for snakes is to simply step in the direction of the "force" (gradient) at each iteration. This yields an optimization algorithm that is essentially gradient descent, which is much slower than the quasi-Newton approach we use.

3.2 The measurement model

The measurement model uses as input the nonparametric version of the object model, $\mathbf{I}\mathbf{f}$, and produces a set of predicted data elements, \mathbf{g} , in this case a Poisson rate for each detector pixel. The measurement model might, in general, contain many components. For example, in an x-ray radiographic system, one would expect to have line integral transformations (parallel- or divergent-beam), convolutions, exponential point transforms, etc.

For the FASTSPECT machine, though, the measurement model is merely the very large matrix, \mathbf{H} , along with a single additive constant that models the background (the same background constant is used for all 24 detectors), so that $\mathbf{g}=\mathbf{H}\mathbf{f}+s$. The background, s , must be jointly estimated from the data along with the object model parameters. Much more complicated 2D spatial field models exist within the BIE, but the very low number of background counts probably make more complex models impossible to estimate well. One extension that is worth investigating is a different background constant for each detector that varies in time in a plausible way.

The nature of our object model allows us to speed up the calculation $\mathbf{H}\mathbf{f}$ dramatically since only a few percent of the voxels in the object model are nonzero. Simply skipping over $\mathbf{H}\mathbf{f}$ for values of \mathbf{f} that are zero allows us to calculate $\mathbf{H}\mathbf{f}$ in about 300 msec on a DEC Alpha 500/500. The same speedup applies in the adjoint direction, wherein derivatives are propagated according to the chain rule in the direction opposite to the path that transforms object parameters into predicted data [9].

3.3 The probability model

The probability model for the object penalizes a discrete approximation to the local curvature at every edge shared by two triangles on the surface in order to enforce smoothness of the estimated surface. Let \mathbf{n}_i be the normal to the i^{th} triangle. We define θ_{ij} to be the angle between \mathbf{n}_i and \mathbf{n}_j . Then, if A_i is the area of the i^{th} triangle, and l_{ij} is 1/3 the height of the i^{th} triangle relative to the edge shared by triangles i and j , the curvature prior is defined as

$$\pi(\mathbf{x}) = \sum_i A_i (\sum_j [\tan(\theta_{ij}/2) / l_{ij}]^2), \quad (1)$$

where \mathbf{x} is the list of (x,y,z) components for each of the 162 vertices, i indexes over all triangles on the surface and j indexes over all triangles that share an edge with the i^{th} triangle.

The form in (1) is slightly different from the form used during our first attempts to analyze simulated FASTSPECT data [5,6]. First, the tangent-squared term in (1) is divided by a length over which it is assumed that half the angular difference occurred [10]. As the angle gets small, the total term looks like an angular velocity w.r.t. arclength, which is the definition of curvature for a curve. Note that this term is not symmetric in the vertices of triangles i and j , and that we did not use a simpler definition for l_{ij} that calculates the distance

from the midpoint of one triangle to the midpoint of the shared edge. These subtleties are intentional, and may be important in keeping the triangulation evenly distributed on the surface during the course of the gradient-based optimization so that no remeshing of the surface is needed. Second, the total surface area no longer appears in the denominator in (1). This is also intentional since such a term would make the prior favor surfaces with larger total area (with the new numerator definition).

The definition in (1) is invariant to isotropic scale changes in the object and also to the number of triangles used in the discrete representation of the surface, as long as the angles are small. This could be an important feature if re-meshing is needed, since one can easily re-mesh to a new set of triangles that does not affect the value of the prior, if all of the angles are small.

The probability model for the likelihood is the Poisson distribution with mean value equal to the predicted data (predicted detector pixel Poisson rates, $\mathbf{g}=\mathbf{H}\mathbf{f}(\mathbf{x})+s$) and count values equal to the raw data:

$$\begin{aligned}\phi(\mathbf{x},\mathbf{I},s) &= -\ln \text{Prob}[\text{data} \mid \text{predicted data}] \\ &= \sum_i [-k_i \ln g_i + g_i],\end{aligned}\tag{2}$$

where we have ignored terms in (2) that depend only on the data \mathbf{k} . The dependence of the predicted data \mathbf{g} on the underlying parameters \mathbf{x} , \mathbf{I} , and s , is understood. Note that because the background constant in the measurement model is additive, the predicted detector pixel rates \mathbf{g} can never be equal to or less than zero as long as the activity level and background level are greater than or equal to zero, which makes the form in (2) well-defined, and makes the derivative of $\phi(\mathbf{x},\mathbf{I},s)$ w.r.t. \mathbf{g} (and ultimately \mathbf{x}) well-behaved.

3.4 The estimation problem

The Bayesian estimation problem is to find the values for the object model parameters, \mathbf{x} , that produce the maximum *a posteriori* (MAP) probability, or the minimum minus-log posterior:

$$\mathbf{x}^{\text{MAP}}(\alpha) = \arg \min_{\mathbf{x}} [\phi(\mathbf{x},\mathbf{I},s) + \alpha \pi(\mathbf{x})],\tag{3}$$

for some fixed value of the hyperparameter, α . The higher-order problem is to determine the value of α from the data. The Bayesian solution to the higher-order problem is to determine the α that yields the greatest evidence for the data, where the evidence is the integral w.r.t. parameters over the joint posterior distribution of parameters and data (leaving just the probability of the data, called the evidence) [10]. However, evaluation of the evidence is computationally nontrivial, so for now we use an alternative, heuristic approach.

We determine α using the L-curve [11], the continuum of 2D points,

$$(\phi(\mathbf{x}^{\text{MAP}}(\alpha)), \pi(\mathbf{x}^{\text{MAP}}(\alpha))),\tag{4}$$

parameterized by α . The L-curve approach is traditionally used for linear least-squares with quadratic regularization, but, to the degree that the minus-log posterior is quadratic, an L-curve approach should be reasonable.

The value of α chosen for the final estimate is the one that yields the point on the L-curve that is closest to the ‘‘corner’’. For very large values of α , the MAP solution is over-regularized. As α decreases, much better fits to the data are allowed (decreasing minus-log likelihood) at very little cost to the prior. This is the vertical line part of the L-curve (Fig. 3). For small values of α , the MAP solution is under-regularized, and the data is well-fit (small value of minus-log likelihood). Now, very large increases in the minus-log prior are needed to allow sufficient freedom in the model so that any decrease in the minus-log likelihood can result. This is the horizontal line part of the L-curve (Fig. 3). Obviously, we’d like to choose an α that is somewhere between these two extremes, thus the ‘‘corner’’ criterion. A

value of $\alpha=0.2$ produces a point on the L-curve that is approximately halfway between the endpoints of the two extreme regions just described and yields a qualitatively pleasing result.

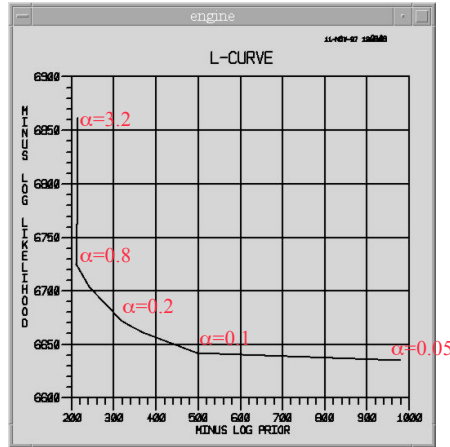


Fig. 3. The L-curve for the diastolic frame.

Since we have no prior information regarding the activity level, I , within the heart, or the background level, s , on the detectors, we simply find the maximum likelihood estimates of those parameters, $I^{ML}(\alpha) = \arg \min_I [\phi(\mathbf{x}, I, s)]$, and $s^{ML}(\alpha) = \arg \min_s [\phi(\mathbf{x}, I, s)]$.

4. Results

The MAP solutions for a range of α using the data in Fig. 1a (diastolic frame) are shown in Fig. 4. As discussed above, large values of α produce over-regularized MAP solutions, while small values of α produce under-regularized, noisy solutions. The MAP solution using $\alpha = 0.2$ is chosen as the “best” MAP estimate. This value of α was also used to find the MAP solution for the data in Fig. 1c (systolic frame).




Fig. 4. MAP reconstructions of the bolus boundary surface using a) $\alpha=3.2$, b) $\alpha=0.2$, and c) $\alpha=0.1$.

Comparison between the MAP reconstructions of the bolus boundary for the diastolic and systolic frames (see Figs. 5 and 6) shows the type of behavior that we expect in most regions. The inlet tube shows very little change in distribution between diastolic and systolic frames presumably because the valve that regulates flow between the tube and the ventricle is closed during that time (Figs. 5a-b and 6a-d, top tube). The diaphragm at the bottom of the ventricle moves toward the inlet and outlet tubes, as expected, causing the ventricular volume to decrease in that region (Figs. 5a and 6d-e, right side). The bolus boundary expands along the direction of blood flow through the outlet tube (Figs. 5b and 6b-f, bottom tube).

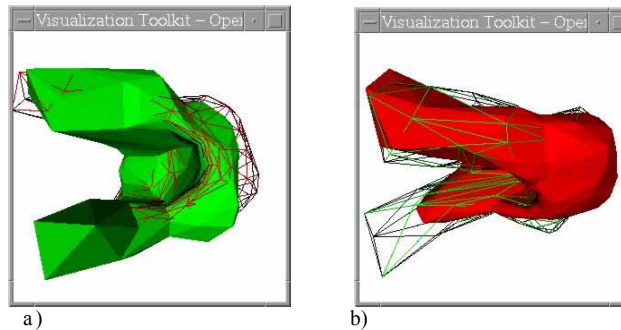


Fig. 5. Comparison between diastolic and systolic frame reconstructions: a) diastolic frame is wireframe and systolic frame is solid surface, b) diastolic frame is solid surface and systolic frame is wireframe. See <http://planck.lanl.gov/~cunning/3D> for an interactive Java display of the reconstructions.




Fig. 6. Cut planes through the reconstructions in Fig. 5. Red lines are for diastolic frame and green lines are for systolic frame. Z-slices are a) -8 mm, b) -4 mm, c) 0 mm, d) 4 mm, e) 8 mm, and f) 12 mm.

5. Discussion

There are several features of the reconstructions that require further investigation. The bolus boundary appears to expand rather than contract within the ventricular volume on the side closest to the outlet tube during the transition from diastoli to systoli (Figs. 5b and 6f, bottom right). This feature tends to make the apparent ventricular volume increase in that region from diastoli to systoli, rather than decrease. This problem could be due to inhomogeneity in the mixing of the radiotracer. If the inhomogeneity is not severe, then simple low-frequency models for the spatial distribution of the radiotracer within the volume defined by the closed surface might be able to eliminate this feature and result in a surface estimate at systoli that is nearer to the ventricular wall.

Another puzzling feature of the reconstructions is the ring connecting the two tubes that lies on top of one side of the ventricle (Fig. 4b, back surface of ventricle between the two tubes). This feature may be an artifact of the curvature prior since there is a high-curvature region that connects the tubes to the ventricles whose true configuration may be disallowed by the prior. This feature might be eliminated if the curvature prior is de-weighted at the join between tube and ventricle, allowing a “kink” to develop at the join [10]. Until these features are eliminated, an estimate of the time-varying volume and ejection fraction using the model proposed in this article would underestimate the quantitative accuracy achievable by a more capable model applied to the same data, and so it is not attempted.

6. Conclusion

We have formulated the analysis of very low-count, first-pass cardiac SPECT data in a Bayesian framework using deformable geometric models. In particular, the model assumes that the radiotracer distribution within the tubes and ventricle is uniformly distributed inside a volume defined by a closed, triangulated surface with 162 vertices and 960 triangles. We jointly estimate the intensity of the distribution as well as the positions of the vertices of the surface from the raw data. The optimization process takes, on a DEC Alpha 500/500 (500 MB memory, 500 MHz processor), 10-15 minutes per frame after the first frame has been

analyzed, approximately the time required to complete 20-30 iterations of ML-EM if the full \mathbf{H} matrix is applied at each iteration.

We use the system matrix for FASTSPECT and an unknown constant additive background to model the predicted rates at the detector as a function of volumetric distributions of radiotracer parameterized by the surface. The system matrix and raw data were provided to us by the University of Arizona. The raw data consist of 24 pinhole views of the radiotracer intensity distribution at diastoli and systoli. The total number of counts, integrated over all 24 detectors, is 2133 at diastoli and 1780 at systoli.

The results are qualitatively quite pleasing, and one can conceive of a number of extensions to the model that will only improve the performance of the approach, e.g., surface "kinks", 3D low-frequency spatial variation of the radiotracer distribution within the volume defined by the closed surface, and truly time-evolving models for the radiotracer distribution which utilize priors on the spatio-temporal nature of the surface velocity field. These extensions should help eliminate the few undesirable features of the current reconstructions that prevent us from producing highly quantitative estimates of ventricular volume and wall motion at this time.

Acknowledgments

This work was supported by the U.S. Department of Energy under contract 7405-ENG-36. We thank Irene Pang and Harry Barrett of the University of Arizona for supplying us with the FASTSPECT data and associated system matrix, complete with attenuation correction.