Forward Deformation of PET Volumes Using Material Constraints

Gregory J Klein (GJKlein@lbl.gov)
Center For Functional Imaging
Lawrence Berkeley National Laboratory,
Berkeley, CA 94720

Abstract

A method for non-rigidly deforming 3D PET datasets is described. The method uses a Lagrangian motion field description and a forward deformation mapping which conserves total voxel intensities. To regularize the deformation, a large-deformation isotropic strain energy function is used that models the material properties of cardiac tissue. The method is applied to motion compensation in PET to combine different time frames, or gates, of a cardiac sequence.

1. Introduction

In gated acquisitions of cardiac Positron Emission Tomography (PET), the motion of the heart is stopped in the images by dividing the data obtained during the cardiac cycle into a number of different time frames, or gates. Typical PET cardiac acquisitions require 20-30 minutes to obtain adequate statistics for acceptable image reconstructions. To direct the tomograph events obtained during this time to the appropriate time frame, the electrocardiograph (ECG) is monitored and used to obtain a trigger on the R-wave. Tomograph events are distributed to the appropriate gate based upon the time since the last R-wave occurrence. An unfortunate effect of distributing the data to many different time frames is that the statistical quality of the reconstructed volume suffers. Ideally, one would like to correct each image for motion, then add them back together to obtain a composite image with less motion blur and better signal to noise properties.

We describe here a deformable motion technique that allows motion compensation for subsequent combination of PET datasets. Unique in the approach are two aspects. First a forward deformation mapping is used which conserves total PET activity between a target and a deformed volume. Second, a smoothing constraint incorporating large deformation strain energy is used to model the underlying cardiac tissue being deformed.

2. Background

Three dimensional deformable motion models find their roots in the two dimensional optical flow literature. Horn [1] was one of the first to develop an optical flow algorithm which calculated a velocity vector for each pixel in an image from a time sequence of 2D images. Integrated over short time periods, this 2D motion field, or flow field could roughly be thought of as the vector field mapping corresponding pixels in two temporally separated images. The technique relied on two constraints: first a constraint assuming that pixel intensity representing a specific image feature would be approximately conserved and second a motion field smoothness constraint based on the assumption that the objects in the images had continuous surfaces which would induce a smoothly varying motion field.

Similar works in the 3D deformation world parallel the constraints of Horn. Song [2],[3] used a direct 3D extension of the Horn algorithm to calculate the motion field in ultrafast computed tomography (CT) images of the heart. Zhou [4] used a similar formulation to calculate larger deformations in CT images of deformed asphalt test structures. Both these techniques relied upon a voxel matching constraint as the driving force to the deformation. A simple smoothness and an incompressibility constraint was imposed to restrict the set of possible particle deformations.

Bajcsy and Kovacic [5] were one of the first to incorporate a more realistic elastic material model as a regularization constraint. Though they used their deformation technique to match 3D CT volumes of brains from different people, the brains appeared similar enough so that a material model of continuous media gave acceptable results. Unfortunately, unpredictable results can occur using the linear elastic model of Bajcsy for large deformations because the elastic models assumes infinitesimal displacements. The model will still provide a smoothing constraint, and hence it has been used by numerous other authors [6, 7, 8, 9] but it's validity as a material model is questionable. Christensen [10],[11] tried to overcome this problem by introducing a viscous fluid model capable of tracking large deformations. He used this technique to match largely differing brain datasets from different patients.

Material model based algorithms appear as elegant solutions to deformable motion problems, however, they are computationally intensive. Essentially, they make use of properties from real world elastic materials that enforce a smooth deformation of some 3D dataset. This is necessary because without such constraints, nearly any arbitrary volume can be warped to match voxels in another volume in many ways. For warping brain datasets from different patients, it is known that the two brains are simply not the same piece of matter that has been warped. Rather, they are similar enough that a linear elastic or viscous fluid model can be used to approximate the required smooth deformation. However, there are a number of other, perhaps simpler, smoothness constraints that could similarly constrain the solution. For these reasons, a number of researchers have found success in the brain warping community using simpler models, such as linear or affine [12], spline [13, 14, 15, 16] or other non-linear [17, 18] deformations. Alternatively, the flow field may be smoothed using a low pass filtering during the evolution of the solution, and this too will impose on overall smoothness constraint on the final solution [19].

Modeling deformations of the heart is distinguished from the brain matching application in two ways. First, the image volume of the heart at two different times actually represents a deforming piece of physical elastic media. Therefore, in this case, as opposed to the brain matching literature, it may make more sense to use an accurate physical model to impose smoothness constraints on the motion field. Second, assuming the isotope has cleared the blood pool in PET images, it is a safe assumption that total radioactivity in the cardiac walls is conserved between time frames. For this reason, a resampling scheme that conserves total voxel intensities is warranted.

For modeling cardiac deformations, it is known that large displacements take place during the cardiac cycle, yet we know that the heart tissue is not a viscous fluid. Our formulation tries to more accurately describe the strain energy seen by the cardiac tissue based on finite deformation mechanics [20],[21],[22].

3. Motion Estimation

As is the case with most 3D deformable algorithms, the constraint driving the deformation is an image matching constraint. A motion field is sought which will deform one volume to best match a reference volume. Because numerous matching transformations exists using only this constraint, the solution is often regularized by imposing an additional smoothness constraint. This latter requirement makes the assumption that the material within the volume is a continuously deforming elastic medium, which can only deform as is consistent with elastic material models.

The motion estimation framework is described as follows. Define two 3D density fields, $f_1(\boldsymbol{r})$, and $f_2(\boldsymbol{r})$, $\boldsymbol{r} = (x, y, z)$, in a discrete domain,

$$\boldsymbol{r} \in \Omega = \{[1, N_x], [1, N_y], [1, N_z]\},$$

where N_x, N_y, N_z are the dimensions of the image volume. We will call $f_1(\boldsymbol{r})$ the target volume, and $f_2(\boldsymbol{r})$ the reference volume. A Lagrangian motion field is defined as,

$$\boldsymbol{m}(x, y, z) = (u(x, y, z), v(x, y, z), w(x, y, z))$$

and the deformed volume of f_1 is defined as,

$$\hat{f}_1(\boldsymbol{r} + \boldsymbol{m}) = f_1(\boldsymbol{r}).$$

Assuming f_1 and f_2 are perfect measurements of a conserved medium, the goal of the motion estimation technique is to find the motion field such that

$$\hat{f}_1(\boldsymbol{r}) = f_2(\boldsymbol{r}).$$

With these definitions, we can define an error term at each voxel location r, as follows:

image matching:

$$e_I(\boldsymbol{r}) = \lambda_I(f_2(\boldsymbol{r}) - \hat{f}_1(\boldsymbol{r}))^2$$

and strain energy

$$e_S(\boldsymbol{r}) = \lambda_S A(\boldsymbol{r})$$

where $A(\boldsymbol{r})$ is the material strain energy density at the voxel location defined by $\boldsymbol{r}$ and where λ_I, λ_S are global scalars used to weight the different error terms. The overall minimization problem is to find a motion field consistent with elastic material properties that best match the deformed and reference via a minimization of:

$$E_{tot} = \sum_{\boldsymbol{r}} [e_I(\boldsymbol{r}) + e_S(\boldsymbol{r})]$$

3. 1. Smoothness Constraint

An isotropic large deformation strain energy function is used to model the smoothness properties of cardiac tissue deformation. Given a displacement field, $\boldsymbol{m}$, the large, or finite displacement strain tensor at each voxel can be calculated as

$$\varepsilon = \begin{bmatrix} e_{11} & e_{21} & e_{31} \\ e_{12} & e_{22} & e_{23} \\ e_{13} & e_{23} & e_{33} \end{bmatrix}$$

where

$$e_{11} = u_x + \frac{1}{2}(u_x^2 + v_x^2 + w_x^2)$$

$$e_{22} = v_y + \frac{1}{2}(u_y^2 + v_y^2 + w_y^2)$$

$$e_{33} = w_z + \frac{1}{2}(u_z^2 + v_z^2 + w_z^2)$$

$$2e_{12} = 2e_{21} = v_x + u_y + (u_x u_y + v_x v_y + w_x w_y)$$

$$2e_{13} = 2e_{31} = u_z + w_x + (u_z u_x + v_z v_x + w_z w_x)$$
$$2e_{23} = 2e_{32} = v_z + w_y + (u_z u_y + v_z v_y + w_z w_y)$$

and where derivatives of the motion field are denoted as $u_x = \frac{du}{dx}$. We distinguish the above formulation from an infinitesimal description of the strain components as described in [5],[23]. The infinitesimal description assumes that the displacement components are small and thus neglect the second order terms. For example, the first term of the tensor would equal $e_{11} = u_x$ in the infinitesimal formulation. Waldeman [24] has shown that because of the large-distance deformations of the heart during the cardiac cycle, the infinitesimal approximation can induce errors of at least 16% in the strain components. Therefore, our analysis uses the finite distance components.

Numerous strain energy functions have been proposed for cardiac tissue. The most realistic models to date incorporate cardiac fiber direction and non-isotropic material properties[22],[25]. However, such models are computationally intensive and for this application would require segmentation of the cardiac tissue in a 3D density volume. To implement a proper model for cardiac fiber direction, recognition of landmarks to establish proper cardiac orientation would be required as well. We do not wish to require such preprocessing at this point, so we make the simplifying assumption that the cardiac tissue material properties are isotropic. Even with this assumption, there are a number of different strain energy functions that can be used. We choose to define the strain energy at a given voxel location, A, as follows [20]:

$$A = \frac{\lambda}{2}(e_{11} + e_{22} + e_{33})^2 + \mu(e_{11}^2 + e_{22}^2 + e_{22}^2)$$
$$+ \frac{\mu}{2}(e_{23}^2 + e_{12}^2 + e_{13}^2)$$

where λ, μ are scalars called the Lame constants. These can be written in terms of the more intuitive constants, E, called the Young's elasticity modulus and ν, called the Poisson ratio:

$$E = \mu(3\lambda + 2\mu)/(\lambda + \mu)\,,\ \nu = \lambda/(2(\lambda + \mu))$$

E relates the tension of the object and it's stretch in the same direction, and ν is the ratio between lateral contraction and axial extension [26].

As a comparison, past efforts [2],[4],[27] have used simple smoothness terms such as

$$\tilde{e}_S = \lambda_S(u_x^2 + u_y^2 + u_z^2 + v_x^2 + v_y^2 + v_z^2 + w_x^2 + w_y^2 + w_z^2)\,,$$

which globally penalize discontinuities in the motion field. The stress model provides regularization constraints more characteristic of true material deformations.

3. 2. Forward Deformation Sampling

Though the motion field describing the volume deformation is a one-to-one mapping in a continuous domain, implementation on a discrete domain involves some subtleties that are important to recognize in the deformation of PET datasets. Past efforts implementing 3D deformations such as the work of Christensen, Bajseky, and Zhou [10],[5],[4] have used a reverse transformation to morph voxels from the deformed volume. In this Eulerian formulation, the motion vectors describe a particle's motion with respect to its final position. That is, the motion vectors are defined in the space of the reference volume, f_2 and the "heads" of all the motion vectors terminate at a discrete voxel locations in the reference volume. The motion vector "tails" effectively sample from a continuously defined location in the target volume, f_1. Thus to obtain the value of each voxel in the deformed volume, $\hat{f}_1(\boldsymbol{r}) = f_1(\boldsymbol{r} - \boldsymbol{m})$, eight voxels from the deformation volume are sampled at the location, ***r-m***, and weighted according to trilinear interpolation. Such *backward* sampling does not guarantee that each voxel in the deformation volume will contribute to the deformed volume. Therefore the total activity in a deformed target PET volume would not generally be conserved using this sampling technique.

We propose a forward sampling technique which defines the location of the motion vector at its starting position in the space of the target volume. In this Lagrangian formulation, the motion of each discrete voxel in the target volume is well-defined, though the motion vector generally terminates on non-integer spatial locations within the reference volume.

To calculate the deformed volume, the value of each voxel in the target volume is simply distributed into the eight voxels at the location, ***r+m***, via trilinear interpolation. More exactly, the value of the deformed volume at some discrete location on the voxel grid, $\tilde{\boldsymbol{r}}$ may be expressed as

$$\hat{f}(\tilde{\boldsymbol{r}}) = \sum_{\boldsymbol{r} + \boldsymbol{m} \in R(\tilde{\boldsymbol{r}})} \gamma_r f_1(\boldsymbol{r})$$

where $R(\tilde{\boldsymbol{r}})$ represents the region where a morphed voxel would contribute via trilinear interpolation to a voxel at location, $\tilde{\boldsymbol{r}}$, and γ_r represents the interpolation weighting factor. In general the deformed volume, $\hat{f}_1(\boldsymbol{r})$, is calculated in a single pass by first initializing the volume to zero, then forward projecting each voxel in the target volume, adding to the appropriate eight voxels in $\hat{f}_1(\boldsymbol{r})$ for each projection.

The forward deformation technique guarantees that each voxel in the target volume is represented in the deformed volume, and that the total PET density is conserved. The technique also allows a voxel intensity increase at regions of true material compression in the deformed volume. In contrast, a backward sampling scheme with a converging motion field could not increase the voxel intensity greater than values in the target volume, since the value of a voxel in the deformation volume is only a single weighted sample in the target volume.

3. 3. Solution Formulation

The energy criterion to be minimized over all voxels in the PET volume is as follows:

$$E_{tot} = \sum_{\boldsymbol{r}} [e_I(\boldsymbol{r}) + e_S(\boldsymbol{r})]$$

For each of the $M = 3N_x, N_y, N_z$ parameters in the motion vector field, we can define the derivative of the energy function with respect to that parameter as:

$$g_k(u_1, v_1, w_1, \ldots, w_{N_xN_yN_z}) = g_k(x_1, x_2, \ldots, x_M) = \frac{\partial E_{tot}}{\partial x_k}$$

We use a nonlinear successive overrelaxation (NLOR) method to obtain the minimization [28]. In this minimization technique, for the solution to a function of the form,

$$g_k(x_1, x_2, \ldots, x_M) = 0, k = 1, 2, \ldots, M$$

a NLOR iteration is defined as

$$x_k^{n+1} = x_k^n - \omega \frac{g_k(x_1^{n+1}, \ldots, x_{k-1}^{n+1}, x_k^n \ldots, x_M)}{\frac{\partial g_k}{\partial x_k}(x_1^{n+1}, \ldots, x_{k-1}^{n+1}, x_k^n \ldots, x_M)}$$

The term, ω, is a scalar typically set between $0 < \omega < 2$. Normally, each iteration of the NLOR technique requires evaluation of the entire objective function and its derivative, however for the PET deformation problem, a small change in a single motion vector component affects only a very small fraction of all the terms in the function. Therefore the NLOR technique can be carried out relatively efficiently. In our implementation, we initialize the motion field to zero, then use a checkerboard update to proceed with the iteration. Additionally, we found that convergence was improved if a limit equal to the voxel size was imposed on the maximum step size.

3. 4. Implementation Details

Because of fairly large displacements with respect to the voxel size, a multiscale approach was useful to obtain suitable convergence. For example, a 128 × 128 × 47 volume is subsampled into 64 × 64 × 23, 32 × 32 × 23 and 16 × 16 × 23 datasets using a uniform cubic B-spline approximation to a Gaussian pyramid [29]. A motion flow field solution was found at the lowest resolution, then is propagated at the next level as the initial condition of the flow field. This technique speeds the overall convergence, and in many cases was found necessary to avoid solutions at incorrect local minima.

Once the deformed volume, $\hat{f}_1(\boldsymbol{r})$ matching $f_2(\boldsymbol{r})$ is obtained, subsequent processing to obtain a composite PET dataset is straightforward. The composite sum is computed as

$$f_{sum}(\boldsymbol{r}) = \hat{f}_1(\boldsymbol{r}) + f_2(\boldsymbol{r})$$

Because the deformed volume conserves the total counts present in the original volume (except at the volume borders), the composite volume represents the total PET counts acquired in the two gates. In general, data from all cardiac gates could be combined to form a single composite image. This paper will only consider the summing of two frames.

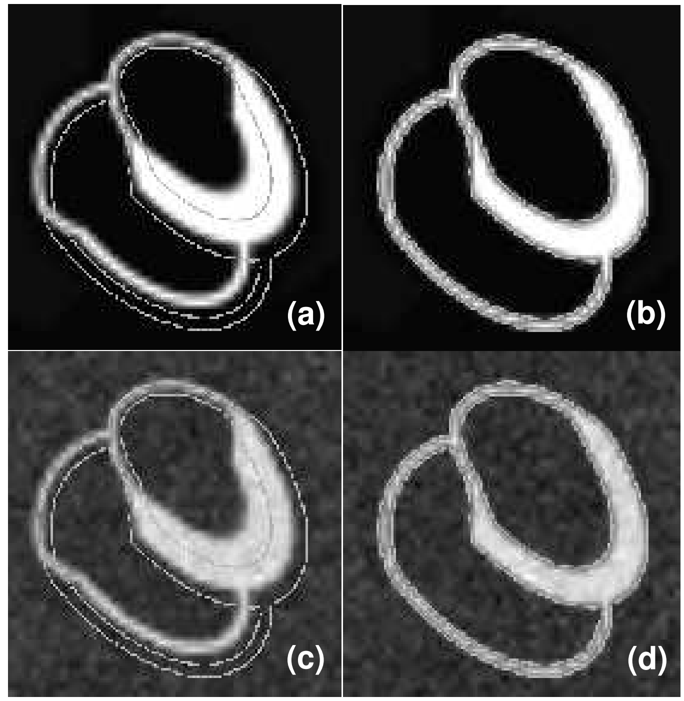

Figure 1. MCAT Heart Phantom. Transverse slice of the target image (a) and reference image (b). Noisy versions of the phantom target and reference images obtained by adding Gaussian noise (c,d) Outline of the reference image is displayed on all images to show the motion between frames.

4. RESULTS

4. 1. Simulated Cardiac Phantom

A realistic cardiac PET phantom was used to test the algorithm. The phantom was obtained using the Mathematical CArdiac Torso (MCAT) software implemented at the University of North Carolina [30], [31], and provides a realistic dataset of the emission PET image of a beating heart against a zero background. Figure 1 shows a single transverse slice through the dataset at the reference and target time frames for both a noise-free case, and for a case where Gaussian noise is added. Though the phantom heart changes between the two time frames, material and intensity is conserved. The results of the deformation algorithm for the noiseless case are shown in Figure 2. Results obtained using backward sampling and a smoothness constraint penalizing flow field discontinuities and divergence are also shown. Both cases show that the target heart at systole can be warped to result in a image matching the reference image at diastole. The important distinction here is that the forward mapping technique conserves voxel intensities

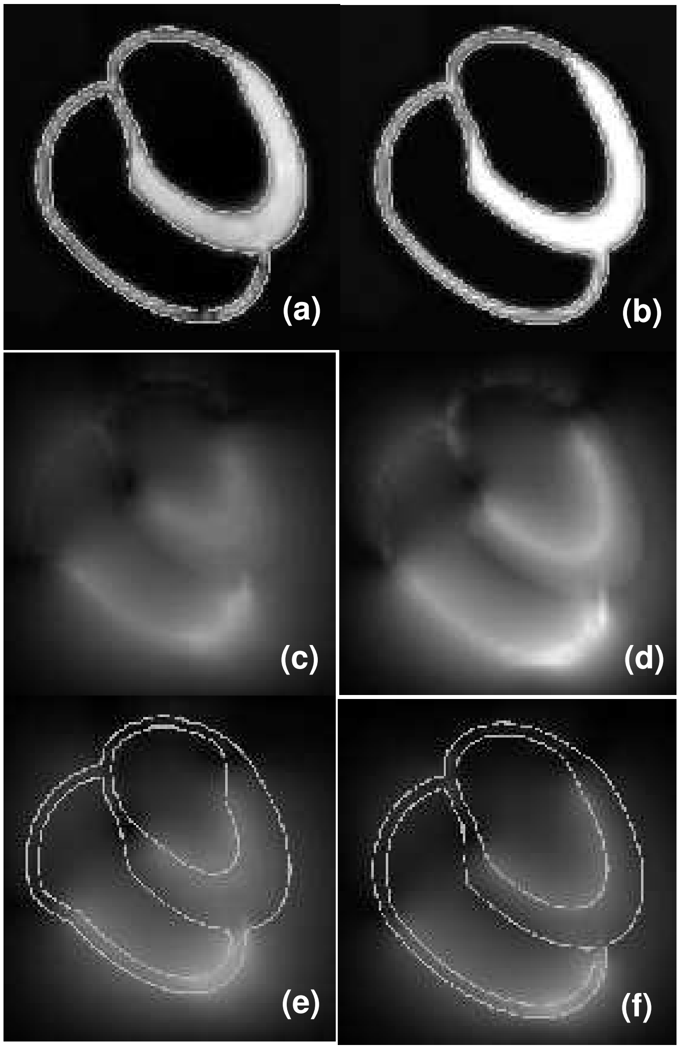

Figure 2. Deformed MCAT Images (noiseless case). Deformed target using forward sampling and strain energy (a) verses backward sampling (b) both are warped to match the reference quite well, but voxels intensities are conserved only in the forward sampling case. Comparison of the flow magnitude for each (c) and (d) shows that the backward sampling technique permits concentration of motion along the edges despite an incompressibility constraint in the formulation, and therefore may not uniformly sample the target image. Note that the overlay of the target edge map on the forward magnitude image(e) and the edge map of the reference image on the backward magnitude (d) show non-zero motion outside the cardiac boundaries, but since the outside voxels are zero-valued, they do not adversely affect the warp.

of the target volume. If the grey level of all voxels in the volume are summed, it is seen that the volumes displayed in Figure 2a conserve the sum (target grey sum = 9091, deformed volumes a sum = 9091), whereas the backward sampled volume Fig 2b does not conserve the sum (voxel grey sum = 8718). Additionally, a close examination of the corresponding motion magnitude image, (Fig. 2d, 2f), shows that despite an incompressibility constraint in the matching criteria, the backward sampling technique allowed a greater concentration to be sampled from the edges of the ventricle walls, and allowed the central portion of the walls to be undersampled. The figure also points out in that the highest flow magnitude for the forward sampling occurs in the portion of the volume corresponding to the shape of the heart in the target (e), whereas for the backward sampling, the motion magnitude is effectively showing the motion at the vector “heads” and are thus concentrated at the voxels corresponding to the heart shape in the reference volume. Note that for both the forward and backward sampling cases, regions outside the boundaries of the heart have non-zero motion magnitudes. Though these vectors may not represent true motion of material adjacent to the cardiac walls, the voxel values of this material in PET imagery are close to zero, so they do not adversely affect the deformed image. It is anticipated that images with a significant non-zero background may require a segmented material model to be effective.

Results for the MCAT phantom in a noisy case are shown in Figure 3. The same forward deformation techniques and material model were again applied in Fig. 3a, 3c. For comparison, results of the algorithm using the simple smoothing criterion, $\tilde{e}_S$, are shown in Fig. 3b, 3d. Close examination shows that the technique using the strain material model achieved slightly superior results. Comparison of the flow magnitude maps indicate that the strain model may more accurately represent motion along object boundaries.

4. 2. Data From Human PET Study

Transverse slices from two successive gates acquired from a human cardiac PET study are shown in Figure 4. The images were obtained using the CTI/Siemens ECAT EXACT HR scanner using the radiotracer, ^{18}F-fluorodeoxyglucose, and 100 msec cardiac gates [27]. The resulting image seen in Fig. 4c show that though there is considerable noise and background, the deformed target image appears to match the reference well.

Image (4e) shows the result of a simple sum of the two gates without motion compensation; the adjacent image (4d) shows the sum, $f_{sum}(\boldsymbol{r}) = \hat{f}_1(\boldsymbol{r}) + f_2(\boldsymbol{r})$. Comparison of the edge map of $f_2(\boldsymbol{r})$ overlaid on the uncompensated sum shows how the motion compensation reduces blur.

5. DISCUSSION

The need for motion compensation in cardiac PET is becoming more apparent as the resolution of conventional PET scanners improves. Blur due to heart motion is now probably one of the limiting factors for resolution of cardiac features. Gating the PET acquisition into different reconstructed time

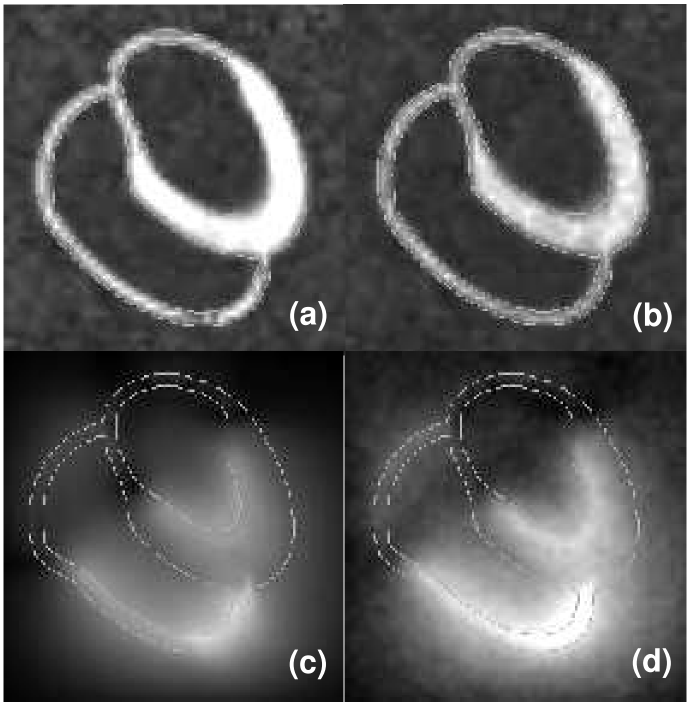

Figure 3. Deformed MCAT Images (noisy case). Deformed target using strain energy model (a) verses simple smoothness (b) shows a slight improvement when compared to the edge map of the reference volume overlaid on each image. Comparison of the motion field magnitude map for the strain model (c) and the simple smoothness model (d) indicates that the strain model may more accurately represent motion along object boundaries.

frames will continue to be a useful technique for stopping the motion. However, because patient time in the scanner and isotope dose will always be at a premium, it will always be desired to combine all gates for optimal signal to noise characteristics.

The deformation model described here appears to be a promising technique for providing motion compensation before recombining PET data. Considerable work needs to be done before this can be done in practice. Validation will be a main focus of this work. To date, the extent of our validation is the appearance of the match between a reference and a deformed dataset. As is always the problem with processing biomedical imaging, there is a lack of a gold standard that expresses a "true" motion field that we could use to evaluate the accuracy of our algorithm. The MCAT phantom is a step in the right direction, but given this model is described by a set of ellipses, true point correspondences are still difficult to establish.

The appropriateness of an isotropic material model is another question that needs further research. We recognize that the heart tissue does not display isotropic material properties.

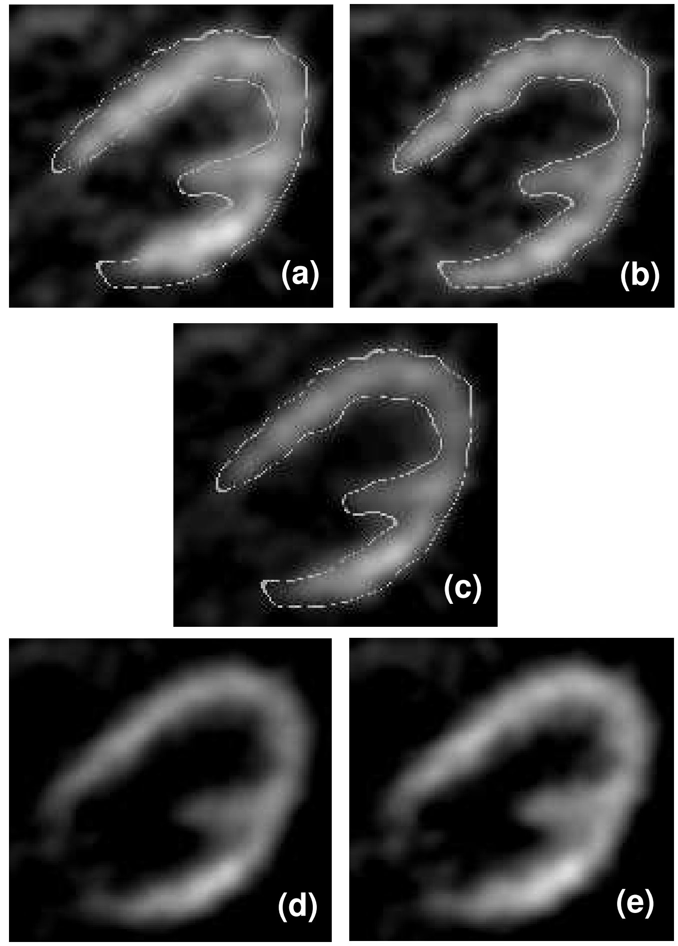

Figure 4. Human Cardiac PET. Target (a) and reference (b) image obtained from a single transverse slice through the data volume. Deformed volume show in (c) closely matches the shape of the reference, and conserves the voxel intensities of the target. An edge map of the reference is displayed on (a,b,c). Result of summing the deformed and reference volume is seen in (d). For comparison, a summed image of the target and reference without motion compensation is seen in (e). Note the decrease in motion blur in the compensated sum.

Likewise, the blood pool within the chambers of the heart and the surrounding image of lung tissue are not continuously joined to the heart tissue, and they obviously have different material properties that the myocardium. It is our hope that the image matching constraint provides suitable information to overcome these material model simplifications. As mentioned earlier, detailed material models of the heart using finite element techniques are available at the expense of considerable computation and pattern recognition requirements. It is a research question as to whether these more detailed models would produce "better" motion fields.

Another topic imploring further research is the convergence properties of the overall energy function using a forward deformation mapping. Because an increment in a motion vector not only moves the deformed location that voxel, but also affects the summed voxel value at that location, we have found that the energy function is quite complicated, and subject to numerous local minima. A technique different from nonlinear successive overrelaxation may be warranted.

6. Acknowledgments

The author would like to thank Ronald Huesman for his valuable comments.

This work was supported in part by the National Heart, Lung and Blood Institute of the U.S. Department of Health and Human Services under grant P01-HL25840; and in part by the Director, Office of Energy Research, Office of Biological and Environmental Research, Medical Applications and Biophysical Research Division of the U.S. Department of Energy under Contract no. DE-AC03-76SF00098.

7. References

[1] B. Horn and B. G. Schunck. Determining optical flow. *Artificial Intell.*, 17:185–203, 1981.

[2] S. M. Song and R. M. Leahy. Computation of 3-D velocity fields from 3-D cine CT images of a human heart. *IEEE Trans Med Imag*, 10(1):295–306, 1991.

[3] S. M. Song, R. M. Leahy, D. P. Boyd, B. H. Brundage, and S. Napel. Determining cardiac velocity fields and intraventricular pressure distribution from a sequence of Ultrafast CT cardiac images. *IEEE Trans Med Imag*, 13(2):386–397, 1994.

[4] Z. Zhou, C. E. Synolakis, R. M. Leahy, and S. M. Song. Calculation of 3D internal displacement fields from 3D X-ray computer tomographic images. *Proc R Soc Lond A*, 449(1937):537–554, 1995.

[5] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Comput. Vision, Graph., Image Proc.*, 46:1–21, 1989.

[6] M. Moshfeghi, S. Ranganath, and K. Nawyn. Three-dimensional elastic matching of volumes. *IEEE Trans Med Imag*, 3(2):128–138, 1994.

[7] C. Davatzikos. Spatial normalization of 3d brain images using deformable models. *J Comput Assist Tomogr*, 20(4):656–665, 1996.

[8] P. Thompson and A. W. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans Med Imag*, 15(4):402–417, 1996.

[9] M. H. Davis, A. Khotanzad, D. P. Flamig, and S. E. Harms. A physics-based coordinate transformation for 3-D image matching. *IEEE Trans Med Imag*, 16(3):317–328, 1997.

[10] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. 3d brain mapping using a deformable neuroanatomy. *Phys Med Biol*, 39:609–618, 1994.

[11] G. E. Christensen, M. I. Miller, and M. W. Vannier. Individualizing neuroanatomical atlases using a massively parallel computer. *Computer*, 29(1):32–38, 1996.

[12] C. A. Pelizzari, G. T. Chen, D. R. Spelbring, R. R. Weichselbaum, and C. T. Chen. Accurate three-dimensional registration of CT, PET and/or MR images of the brain. *J Comput Assist Tomogr*, 13(1):20–26, 1989.

[13] J. Declerck, G. Subsol, J. Thirion, and N. Ayache. Automatic retrieval of anatomical structures in 3D medical images. In *Computer Vision, Virtual Reality and Robotics in Medicine, Lecture Notes in Computer Science*, volume 905, pages 153–162. Springer-Verlag, INRIA, Sophia-Antipolis, France, 1995.

[14] J. Declerck, J. Feldmar, M. L. Goris, and F. Betting. Automatic registration and alignment on a template of cardiac stress and rest reoriented SPECT images. *IEEE Trans Med Imag*, 16(6):727–737, 1997.

[15] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans Patt Anal Machine Intell*, 11(6):567–585, 1989.

[16] S. Minoshima, R. A. Koeppe, K. A. Frey, M. Ishihara, and D. E. Kuhl. Stereotactic PET atlas of the human brain: Aid for visual interpretation of functional brain images. *J Nucl Med*, 35(6):949–954, 1994.

[17] K. J. Friston, C. D. Frith, P. F. Liddle, and R. Fackowiak. Plastic transformation of PET images. *J Comput Assist Tomogr*, 15(4):634–639, 1991.

[18] T. Greitz, C. Bohm, S. Holte, and L. Eriksson. A computerized brain atlas: Construction, anatomical content, and some applications. *J Comput Assist Tomogr*, 15(1):26–38, 1991.

[19] P. Thirion. Fast non-rigid matching of 3d medical images. In *Medical Robotics and Computer Aided Sutgery*, pages 47–54, Baltimore, 1995.

[20] K. Washizu. *Variational Methods in Elasticity and Plasticity*. Pergamon Press, Oxford, 1982.

[21] A. E. Green and J. E. Atkins. *Large Elastic Deformations*. Clarendon Press, Oxford, 1970.

[22] P. Nielsen, I. Le Grice, B. H. Smaill, and P. J. Hunter. Mathematical model of geometry and fibrous structure of the heart. *Am. J. Physiol.*, 29:H1365–H1378, 1991.

[23] G. E. Christensen. *Deformable Shape Models for Anatomy*. PhD thesis, Washington University, 1994.

[24] L. K. Waldeman, Y. C. Fung, and J. W. Covell. Transmural myocardial deformation in the canine left ventricle. *Circulation Research*, 57(1):152–163, 1985.

[25] P. J. Hunter, M. P. Nash, and G. B. Sands. Computational electromechanics of the heart. In A. V. Panfilov and A. V. Holden, editors, *Computational Biology of the Heart*, chapter 12, pages 345–404. John Wiley and Sons, Chichester, 1997.

[26] R. J. Atkin and N. Fox. *An Introduction to the Theory of Elas-*

ticity. Longman, London, 1980.

[27] G. J. Klein and R. H. Huesman. A 3D optical flow approach to addition of deformable PET volumes. In *Proceedings of the IEEE Nonrigid and Articulated Motion Workshop, June 15-16, Puerto Rico*, pages 136–143. IEEE Computer Society, 1997.

[28] G. Sewell. *The Numerical Solution of Ordinary and Partial Differential Equations*. Academic Press, Inc., San Diego, 1988.

[29] B. W. Reutter, G. J. Klein, K. M. Brennan, and R. H. Huesman. Acquisition and automated 3-D segmentation of respiratory/cardiac-gated PET transmission images. In A. Del Guerra, editor, *1996 IEEE Nuclear Science Symposium Conference Record, Anaheim, CA November 2-9*, pages 1357–1361, 1996.

[30] J. A. Terry, B. M. W. Tsui, J. R. Perry, J. L. Hendericks, and G. T. Gullberg. The design of a mathematical phantom of the upper human torso for use in 3D SPECT imaging research. In D. C. Mikulecky and A. M. Clarke, editors, *Biomedical Engineering: Opening New Doors (Proceedings of the 1990 Fall Meeting of the Biomedical Engineering Society, Blacksburg, VA)*, pages 185–190, New York, 1990. New York University Press.

[31] J. A. Terry, B. M. W. Tsui, J. R. Perry, and G. T. Gullberg. A three-dimensional mathematical phantom of the human torso for use in SPECT imaging research studies. *J Nucl Med*, 31(5):868, 1990.