Deformable Registration with Discontinuity Preservation using Multi-Scale MRF

Dohyung Seo and Jeroen van Baar

MERL - Mitsubishi Electric Research Labs, Cambridge MA, USA.
jeroen@merl.com,

Abstract. Deformable (2D or 3D) medical image registration is a challenging problem. Existing approaches assume that the underlying deformation is smooth. This smoothness assumption allows for solving the deformable registration at a coarse resolution and interpolate for finer resolutions. However, sliding of organs and breathing motion, exhibit discontinuities. We propose a discrete optimization approach to preserve these discontinuities. Solving continuous deformations using discrete optimization requires a fine distribution of the discrete labels. Coupled with the typical size of medical image datasets, this poses challenges to compute solutions efficiently. In this paper we present a practical, multi-scale formulation. We describe how discontinuities can be preserved, and how the optimization problem is solved. Results on synthetic 2D, and real 3D data show that we can well approximate the smoothness of continuous optimization, while accurately maintaining discontinuities.

Keywords: MRF, discrete optimization, non-rigid registration

1 Introduction

Deformable image registration is an important component in medical applications. Examples include alignment of a patient's MRI to a reference MRI for diagnosis, image-guided therapy and adaptive radiation therapy [15].

Although individual organs typically exhibit smooth deformations, sliding motions along organ and tissue boundaries present discontinuities in the deformation fields. For example, in 4D CT breathing motion data, there is a deformation with discontinuity between the lung and surrounding tissue. However, the emphasis of previous approaches has been on underlying smooth deformations, e.g. [6, 16, 17]. An advantage of these methods is that the deformation can be solved on a coarse control points grid, and then interpolated to a more fine resolution. Additional constraints can be incorporated into the optimization, e.g., no negative values in the Jacobian map of the deformation to prevent folding. On the other hand, these methods cannot accurately compute solutions in the presence of discontinuities in the deformation field.

One approach to circumvent this discontinuity issue could be to segment each organ and perform registration on each organ separately. However, this requires accurate segmentation as a preprocessing step, which is not an easy task. In

addition, different organs may require different segmentation strategies. Finally, this approach would not take any deformations due to organ interactions into account.

In this paper, we propose a deformable image registration framework based on a discrete approach with appropriate regularization which can handle local discontinuities in deformations efficiently, and accurately preserve smoothness globally. We consider the registration problem as a labeling problem, where each label represents a deformation vector. A challenge with discrete approaches is the large number of labels that has to be considered in order to obtain sufficiently accurate deformations. We present a multi-scale Graph Cuts (GC) approach which drastically reduces the number of labels to consider at each scale. We discuss the properties of such a multi-scale formulation compared to *regular* Graph Cuts, and how we solve the corresponding optimization problem. We will show that we can obtain solutions which are comparable to continuous approaches, but outperform them in the case of discontinuities. We show results for both 2D and 3D images. In the remainder of this paper we simply refer to pixels, which could mean either 2D pixels or 3D voxels.

2 Related Work

Non-rigid registration is a well-studied field. A complete list of previous work is beyond the scope of this paper. Broadly, non-rigid registration can be divided into continuous and discrete approaches. Examples of continuous approaches are described in [6, 17, 16]. These methods assume a smooth underlying deformation field. Discontinuity preserving approaches on the other hand explicitly aim to handle discontinuities in the deformation field. Several variational approaches have been proposed [3, 2, 11]. We propose a discrete optimization approach which preserves discontinuities, but is much simpler to formulate and achieves better results.

Non-rigid registration using discrete optimization has been studied in [14, 5, 10, 4, 12]. In most cases, computing accurate deformation vectors requires large label sets, which translates into high computation time and memory consumption.

Tang and Hamarneh [13] aim to reduce the set of labels by proposing to start with a coarse label set. By examining the corresponding random walk solution, they determine where labels need to be refined, recompute a solution and iterate. If, however, the data contains some mixture of deformations, this approach will also result in the need to solve a problem with a large set of labels.

We propose the following multi-scale approach: At each scale we consider a fixed set of labels. For each subsequent scale the discrete step is refined until desired accuracy is reached. Glocker et al. [5] discuss a refinement scheme for registration using grid of control points. We instead define our multi-scale approach directly on the pixels in order to preserve discontinuities. In the next section we discuss our approach in more detail, and how to solve the non-submodular problem that results from the multi-scale approach.

3 Multi-Scale Graph-Cuts

The goal of non-rigid or deformable registration is to register a moving N–dimensional image M to a fixed (reference) N–dimensional image F. The result of the registration procedure yields a deformation vector field $\mathcal{T} : \mathrm{u} \in \mathbb{R}^N$. Rather than solving the *forward* registration problem: $M(\mathcal{T}(\mathrm{x})) = \tilde{F}$, we solve for the *backward* registration, i.e., $\forall \mathrm{x} \in \Omega_F : \mathcal{T}(\mathrm{x}) = \mathrm{u}_{F \to M}$. Here, Ω_F is the domain of the fixed image.

For the ease of formulation, the ability to preserve discontinuities with appropriate regularization, and obtaining strong local minima, we have chosen to use discrete optimization using Graph Cuts. The solution space is discretized into a set of possible deformations $\mathcal{T} = \{\mathrm{t}_1, \ldots, \mathrm{t}_n\}$, and each deformation is associated with a discrete label: $L = \{\forall i \in \mathcal{T} : \mathrm{l}_i \leftrightarrow \mathrm{t}_i\}$. Optimization is thus formulated as a labeling (f) problem and written as:

$$E(f) = \sum_{\mathrm{x} \in \Omega} E_d(\mathrm{f_x}) + \sum_{\mathrm{x} \in \Omega, \mathrm{y} \in \mathcal{N}(\mathrm{x})} E_{sm}(\mathrm{f_x}, \mathrm{f_y}). \tag{1}$$

A strong local minimum to Equation 1 can be found using α-expansion [12]. To obtain accurate results, sub-pixel deformation values should be considered. A fine sampling of the solution space leads to a large number of labels. As complexity grows linearly with the number of pixels and the number of labels, we propose a multi-scale approach. We employ multiple scales in both image space, as well as label space. We next describe the steps involved in our approach.

We first build an image pyramid and start at the coarsest *image* level I_i. For image level I_i we start with a coarse *label* level $I_{i,l}$, in which the deformation range for $I_{i,l}$ is defined by $\left[r_{1,d}^{I_{i,l}} : \Delta^{I_{i,l}} : r_{2,d}^{I_{i,l}}\right]$ pixels, where d is the dimension index, e.g., $d = 1, 2$ for 2D, and $\Delta^{I_{i,l}}$ is the step size at level $I_{i,l}$. For discontinuity preservation, we choose our pairwise potential to be a log student-T distribution: $\rho_V(|u_p - u_q|) = log(1 + \frac{1}{2v^2}(|u_p - u_q|)^2)$ [8]. To help avoid bad local minima, for the current set of labels at level $I_{i,l}$, we solve the optimization problem n times, and perform fusion on the resulting solutions [9]. We then proceed to the next level $I_{i,l+1}$, with deformation range $\left[r_{1,d}^{I_{i,l+1}} : \Delta^{I_{i,l+1}} : r_{2,d}^{I_{i,l+1}}\right] = \left[r_{1,d}^{I_{i,l}}/\delta : \Delta^{I_{i,l}}/\delta : r_{2,d}^{I_{i,l}}/\delta\right]$, where $\delta > 1$. Scaling by δ redefines the search space in image matching: it simultaneously reduces the search space while obtaining higher resolution w.r.t. labeled deformations. At each level $I_{i,l+1}$, the deformation vectors are aggregated so that the data and smoothness term of Equation 1 are redefined on the aggregated deformation vectors as follows:

$$E = \sum_{\mathrm{p} \in \Omega} E_d(\bar{\mathrm{u}}_\mathrm{p} + \mathrm{u}_\mathrm{p}^{I_{\mathrm{i,l}}}) + \sum_{\mathrm{p} \in \Omega, \mathrm{q} \in \mathcal{N}(\mathrm{p})} E_{sm}(\bar{\mathrm{u}}_\mathrm{p} + \mathrm{u}_\mathrm{p}^{I_{\mathrm{i,l}}}, \bar{\mathrm{u}}_\mathrm{q} + \mathrm{u}_\mathrm{q}^{I_{\mathrm{i,l}}}). \tag{2}$$

Here, $\bar{u}$ and $u^{I_{i,l}}$ represent the aggregated deformation vector up until level $I_{i,l-1}$, and the current set of labeled deformation vectors, respectively. We repeat this

re-scaling process until a certain amount of smoothness and accuracy is achieved, and we set the final deformation vector as $\bar{u}_{final} = \bar{u} + u^{I_{i,l}}$ at the current image pyramid level I_i. The resulting deformation vector fields are then up-scaled for the next level of the image pyramid.

In Equation 2, we should remark that due to the aggregated deformation vector, the smoothness terms are now non-submodular. Quadratic pseudo-boolean optimzation (QPBO) [7] can compute solutions to problems with non-submodular terms. However, QPBO may return some pixels with *unknown* label. We employ α-expansion using QPBO and observe that pixels with *unknown* label occur only for a few iterations and just a small (typically < 30) number of pixels. For each pixel with *unknown* label, we determine which label gives lower energy: α or current label. The label that yields lowest energy is then assigned to that pixel. We also employ QPBO for the fusion of the n solutions computed at each level.

4 Results

We can choose from different dissimilarity metrics for the data term in Equation 1, e.g., Normalized Cross Correlation or Mutual Information. For the results in this paper we found that the Sum of Squared Differences (SSD) dissimilarity metric between moving and fixed images works well. Figure 1 shows a comparison for two synthetic examples: one with smooth deformation, and one with discontinuities in the deformation. The top row are source, smooth target, ground truth smooth deformation field colormap, discontinuous target, ground truth discontinuous deformation field colormap respectively. The second through fourth rows show the results for our multi-scale GC (column 1), FastPD [5] (column2), TV-L1 [2] (column 3), Diffeomorphic Demons [16] (column 4), and FFD using B-splines [17] (column 5). Our method performs well for both smooth and discontinuous deformations. See also Table 1 for a comparison of angular errors (AE) and relative errors (RE).

Table 1. Error metrics for synthetic smooth and discontinuous deformations. For angular and relative errors, we report: mean (standard deviation). For the discontinuous case we list error values computed for all pixels, and for only the region near the discontinuities. Our method performs well for both cases, unlike the other methods.

	Multi-scale GC	**FastPD**	**TV-L1**	**B-splines**	**Diff. demons**
AE	0.17 (0.25)	0.12 (0.28)	0.32 (0.514)	0.09 (0.20)	0.18 (0.30)
RE	0.58 (1.02)	0.39 (0.94)	0.93 (1.80)	0.29 (0.64)	0.63 (1.11)
AE (all)	0.17 (0.25)	0.22 (0.33)	0.22 (0.28)	0.18 (0.29)	0.18 (0.29)
RE (all)	0.60 (1.00)	0.76 (1.39)	0.79 (1.25)	0.61 (0.85)	0.63 (1.11)
AE (region)	0.33 (0.36)	0.64 (0.52)	0.29 (0.30)	0.51 (0.46)	0.46 (0.43)
RE (region)	0.98 (3.922)	1.96 (7.73)	0.94 (3.75)	1.51 (1.18)	1.38 (1.19)

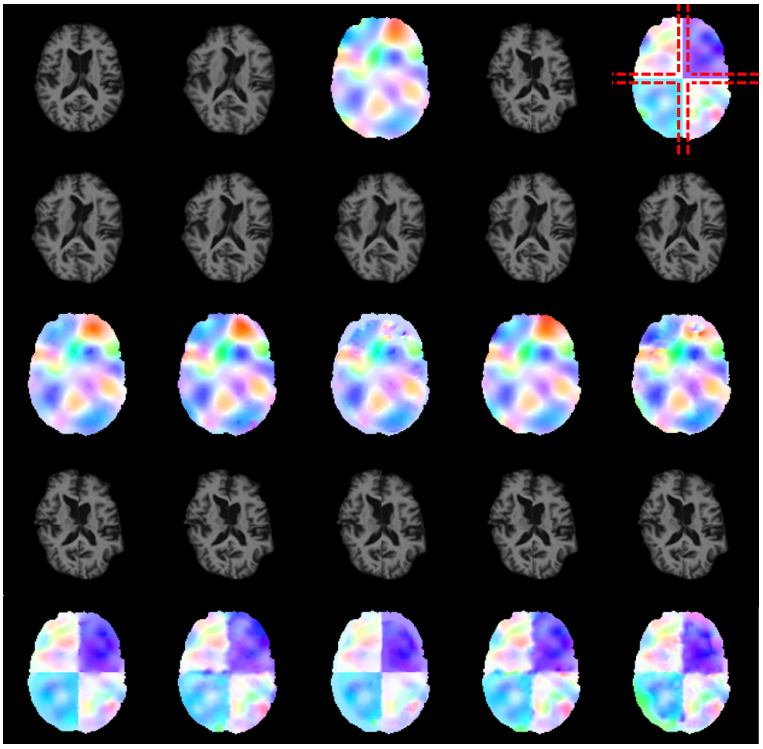

Fig. 1. Top Row source, smooth target, ground truth smooth deformation field colormap, discontinuous target, ground truth discontinuous deformation field colormap. **Rows 2–5** Results for our multi-scale GC (column 1), FastPD [5] (column2), TV-L1 [2] (column 3), Diffeomorphic Demons [16] (column 4), and FFD using B-splines [17] (column 5). Our method performs well for both smooth and discontinuous deformations.

In our next 2D example we register a coronal and sagittal slice between full inhale (T00) and full exhale (T50) from case 1 of the DIR-lab dataset [1]. We compare TV-L1 with our method. Although both methods are able to preserve discontinuities, our method also recovers smooth deformation within the lungs.

Figure 3 shows a 3D example using case 5 of the DIR-lab dataset. We now register the 3D volumes of full inhale (T00) and full exhale (T50). In this figure we compare our method to FastPD [5]. The right image in the top row is the FastPD result. The highlighted area shows that discontinuities were not well recovered. In rows two, three and four we show the target image slice, multi-scale GC, and FastPD registration respectively. The highlighted areas in the images in row four show areas where FastPD failed to register the images correctly. Finally, the bottom row shows the deformation field color maps for multi-scale GC (left) and FastPD (right). Highlighted areas again show where discontinuities were incorrectly recovered. We also computed registration errors for the first five cases

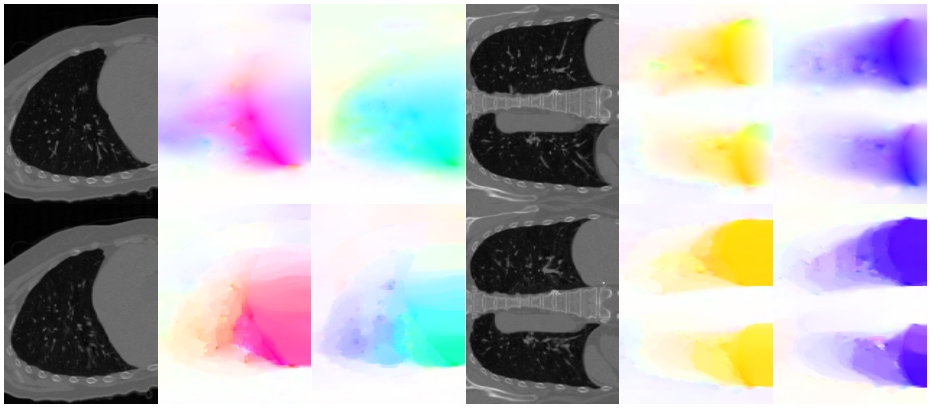

Fig. 2. Column 1, 4 Coronal and sagittal source (top row) and target (bottom row) slices with breathing motion. The registration for our multi-scale GC (top row) source to target (column 2, 5) and target to source (column 3, 6) compared to TV-L1 (bottom row). Although TV-L1 can preserve discontinuity along the lung-chest boundary, deformations within lungs are not smooth.

of the DIR-lab dataset, see Table 2. Our registration accuracy for the landmarks within the lungs is similar to that of FastPD, and somewhere in the middle compared to the best and worst reported results[1]. However, the top performing methods require segmentation prior to registration, whereas our method does not.

Table 2. Comparison of registration errors (in mm), mean (standard deviation), for ground-truth landmark points.

	Multi-scale GC	**FastPD**	**Best**	**Worst**
Case 1	1.03 (1.09)	1.07 (0.62)	0.76 (0.89)	1.58 (1.30)
Case 2	0.99 (1.00)	1.07 (0.74)	0.72 (0.87)	1.49 (1.35)
Case 3	1.33 (1.23)	1.59 (1.34)	0.91 (1.05)	2.27 (1.40)
Case 4	1.85 (1.83)	2.25 (2.36)	1.22 (1.24)	3.27 (4.09)
Case 5	1.82 (1.91)	1.91 (2.05)	1.07 (1.46)	3.59 (2.83)

5 Conclusion

We have presented a multi-scale GC approach which can achieve both smooth results as well as preserve discontinuities in the deformations. Although other methods rely on the assumption that the underlying deformation is always smooth, practical cases show that this is not necessarily the case. Our method achieves accurate results with well behaved deformations. Our multi-scale approach significantly reduces the size of the problem, making it more practical.

[1] http://www.dir-lab.com/Results.html

With our current (naive) C++ implementation, on a single-core processor, registration of a $256 \times 256 \times 100$ volume takes about 120 minutes. Future work includes a GPU implementation and further evaluation, including clinical cases.

References

1. Castillo, R., Castillo, E., Guerra, R., Johnson, V.E., McPhail, T., Garg, A.K., Guerrero, T.: A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. Physics in Medicine and Biology 54(7), 1849–1870 (2009)
2. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. Journal of Mathematical Imaging and Vision 40(1), 120–145 (2011)
3. Chumchob, N., Chen, K.: A variational approach for discontinuity-preserving image registration. In: Proceedings of ICMA-CU. pp. 266–282 (2010)
4. Cobzas, D., Sen, A.: Random walks for deformable image registration. In: Proceedings of MICCAI. vol. 2, pp. 557–565 (2011)
5. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through {MRFs} and efficient linear programming. Medical Image Analysis 12(6), 731–741 (2008)
6. Haker, S., Tannenbaum, A., Kikinis, R.: Mass preserving mappings and image registration. In: Proceedings of MICCAI, vol. 2208, pp. 120–127 (2001)
7. Kolmogorov, V., Rother, C.: Minimizing nonsubmodular functions with graph cuts-a review. IEEE Pattern Analysis and Machine Intelligence 29(7), 1274–1279 (July 2007)
8. Lempitsky, V., Roth, S., Rother, C.: Fusionflow: Discrete-continuous optimization for optical flow estimation. In: IEEE Computer Vision and Pattern Recognition, 2008. pp. 1–8 (June 2008)
9. Lempitsky, V., Rother, C., Roth, S., Blake, A.: Fusion moves for markov random field optimization. Tech. Rep. MSR-TR-2009-60 (May 2009), http://research.microsoft.com/apps/pubs/default.aspx?id=80790
10. Mahapatra, D., Sun, Y.: Joint registration and segmentation of dynamic cardiac perfusion images using mrfs. In: Proceedings of MICCAI, vol. 6361, pp. 493–501 (2010)
11. Panin, G.: Mutual information for multi-modal, discontinuity-preserving image registration. In: Advances in Visual Computing, Lecture Notes in Computer Science, vol. 7432, pp. 70–81. Springer Berlin Heidelberg (2012)
12. So, R.W., Tang, T.W., C.S.Chung, A.: Non-rigid image registration of brain magnetic resonance images using graph-cuts. Pattern Recognition 44, 2450–2467 (2011)
13. Tang, L., Hamarneh, G.: Random walks with efficient search and contextually adapted image similarity for deformable registration. In: Proceedings of MICCAI 2013, vol. 8150, pp. 43–50 (2013)
14. Tang, T.W.H., Chung, A.C.S.: Non-rigid image registration using graph-cuts. In: Proceedings of MICCAI. pp. 916–924 (2007)
15. Timmerman, R., Xing, L.: Image-Guided and Adaptive Radiation Therapy. Wolters Kluwers Health (2010)
16. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. In: Neuroimage. vol. 45 (2009)
17. Xie, Z., Farin, G.: Image registration using hierarchical b-splines. In: Visualization and Computer Graphics, IEEE Transactions on. vol. 10, pp. 85–94 (2004)

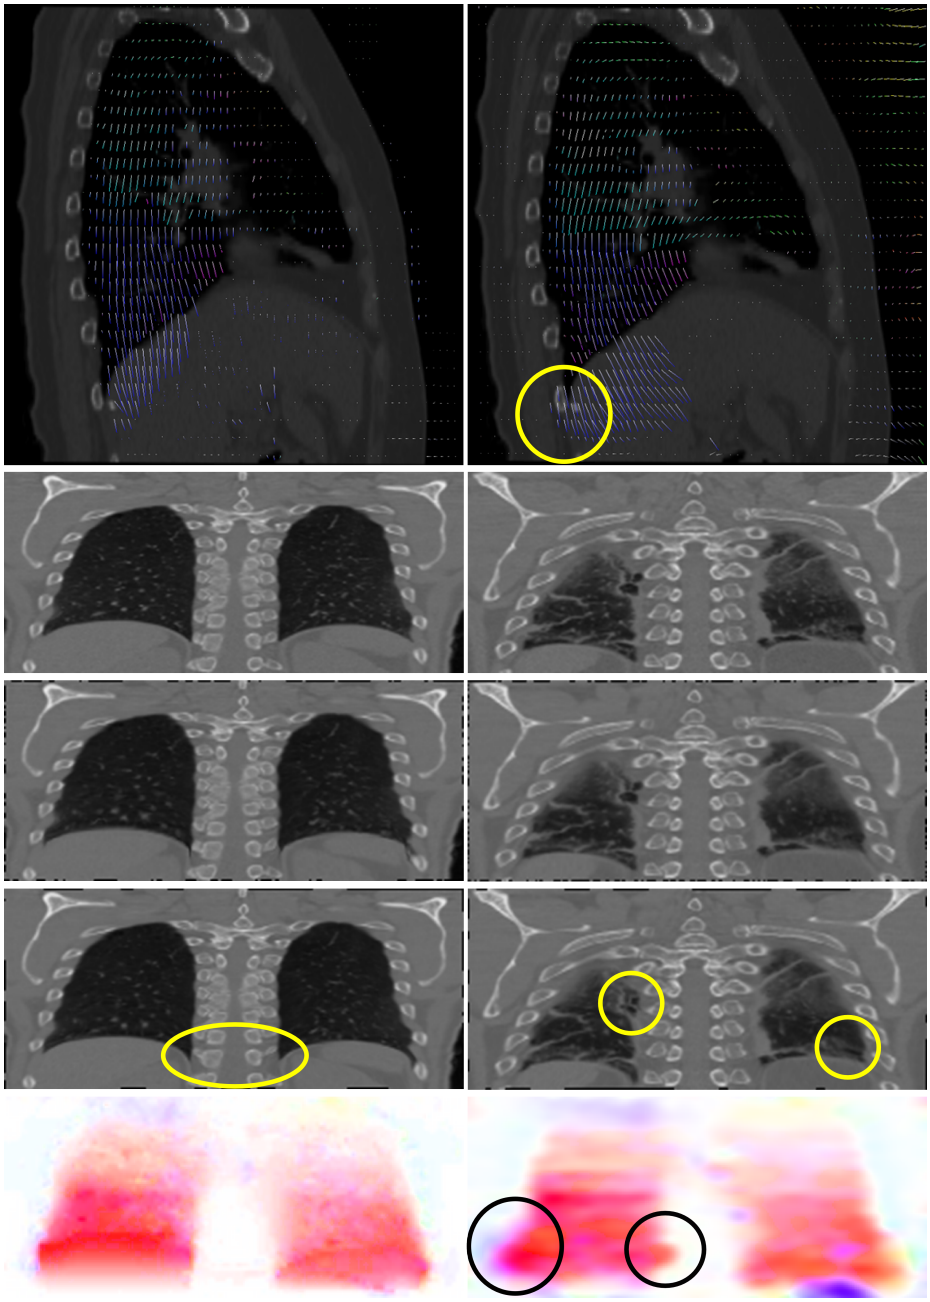

Fig. 3. (Best viewed electronically) **Top Row** Multi-scale GC vs FastPD. Highlighted area shows that FastPD did not preserve discontinuity. **Rows 2,3,4** Shown are target slice, multi-scale GC, and FastPD registration respectively. Highlighted areas in row 4 show the failure to preserve discontinuities by FastPD. **Bottom Row** The deformation field color maps for multi-scale GC (left) and FastPD (right). The highlighted areas for the FastPD result show where discontinuities were properly preserved.