

SYMMETRIC NON-RIGID IMAGE REGISTRATION VIA AN ADAPTIVE QUASI-VOLUME-PRESERVING CONSTRAINT

Iman Aganj, Martin Reuter, Mert R. Sabuncu, and Bruce Fischl

Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School
Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology

ABSTRACT

The standard implementation of non-rigid image registration is asymmetric, even though symmetry might be an intrinsic attribute of the particular application, e.g., pairwise image alignment. Current approaches to restore symmetry to non-rigid registration, although successful in achieving inverse-consistency, generally alter the objective function through implicit inclusion of a non-uniform weight in the integral that is computed on the native space of an input image. This inhomogeneous integral measure, which varies through the course of the registration, results in regional biases by allowing image regions to contribute differently to the objective function. In this work, instead of symmetrizing the objective function, we address the root of the problem: the non-uniformity of the integral in both the asymmetric and the symmetrized implementations. We introduce a new quasi-volume-preserving constraint that keeps the forward and backward objective functions arbitrarily close to each other – hence the registration symmetry – without compromising the uniformity of the integrals. We show the advantages of our method through experiments on synthetic images and real X-ray and MRI data.

Index Terms— Non-rigid image registration, symmetry, inverse-consistency, volume-preserving constraints.

1. INTRODUCTION

In many image registration applications, such as the alignment of two serial scans, the ordering of the images should not influence the registration results. Pairwise image registration is an inherently symmetric problem, in the sense that the correspondence between the regions in two aligned images is naturally expected to be independent of the order that the images are given to the registration algorithm. Consequently, such “inverse-consistency” is necessary for a registration method to be considered reliable. Most standard non-rigid image registration approaches [1] optimize a cost function (CF) defined as the integral of a distance measure taken uniformly on the space of one of the input images. Due to deformations, however, the equivalent integral on the space of the other image is not uniform anymore. The bias resulting from the arbitrary choice of the reference image, on whose *native* space (the space in which it is not distorted) the integral is defined uniformly, breaks the symmetry of registration. In longitudinal studies in particular, favoring one time point over another may result in errors dominating the subtle changes that we seek to measure [2].

Symmetrization of the CF with respect to the two input images, i.e. adding the CF to itself with the order of the images reversed, has so far been the leading approach to inverse-consistent non-rigid image registration [3-7]. In a second class of symmetric

registration in the literature, the CF integral is taken, not in the native space of the images, but in an abstract “mid-space” chosen to be “in between” them [8-11]. These symmetrization approaches are quite effective in achieving inverse-consistency. Nonetheless, they come at the price of altering the original CF, causing the aforementioned non-uniform integral measure in the asymmetric registration (appearing in only one of the image spaces, thereby being the reason for the asymmetry in the first place) to continue to exist in both image spaces (see Sec. 2 for details). Non-uniform integration of images – intrinsic to asymmetric and the symmetrization methods – is undesirable, as it causes different image regions to contribute with arbitrarily different weights to the CF, an inconsistent regional bias that is not alleviated by eliminating the *symptom* of asymmetry through symmetrization.

In this work, instead of symmetrizing the CF, we address the underlying cause of the asymmetry: non-uniform integration of images in their native spaces. We propose a new non-rigid registration algorithm that constrains the deformation so that the integrals in both forward and backward (native-space) CFs remain uniform except for the regions where non-uniformity does not contribute to asymmetry error, which, as will be clear in Sec. 2, are areas where the images match. An immediate result of the new quasi-volume-preserving (QVP) constraint is the registration symmetry, as expected. Furthermore, the forward and backward native CFs, i.e. those with a uniform integral on the native space of the image, will be close to each other throughout the registration process, a property that we call *native symmetry*, which is a characteristic of our method. Native symmetry ensures that *both* native CFs agree on the solution of the registration, as opposed to only one of them (asymmetric registration) or only the sum of them (symmetrization).

An additional advantage of restricting the deformation in dissimilar regions is helping to avoid entrapment of the iterative algorithm in local minima as a result of too much flexibility, thereby guiding it towards a good overall QVP fit before relaxing the constraints and achieving an optimum warp. This is particularly important in registration across time in medical imaging, where large changes in anatomy may be present.

In the remainder of this paper, we first describe the proposed method in detail (Sec. 2), and then present and discuss experimental results (Sec. 3), along with some concluding remarks.

2. METHODS

The standard asymmetric forward and backward CFs, to be minimized in non-rigid registration, are defined as:

$$I_{fw}(f, g, T) := \int_{\Omega} D(f(x), g \circ T(x)) dx, \quad (1)$$

$$\begin{aligned}
I_{bw}(f, g, T) &:= I_{fw}(g, f, T^{-1}) \\
&= \int_{\Omega} D(f \circ T^{-1}(y), g(y)) dy \\
&= \int_{\Omega} D(f(x), g \circ T(x)) J(x) dx,
\end{aligned} \tag{2}$$

where $f, g: \Omega \rightarrow \mathbb{R}$ are two d -dimensional images ($\Omega \subseteq \mathbb{R}^d$), $T: \Omega \rightarrow \Omega$ is the invertible transformation to be determined, with $J(x) := \det(\partial T(x))$ its Jacobian determinant resulting from the change of variable $y = T(x)$, and $D: \mathbb{R}^2 \rightarrow \mathbb{R}$ is the positive and symmetric image distance measure; e.g., $D(s, t) = (s - t)^2$ for the common sum-of-squared-difference (SSD) CF, also used in our experiments. We call I_{fw} and I_{bw} the *native* CFs, since their integrals have uniform representations in the *native* spaces of f and g respectively. By native space of an image, we mean the space where the image is not distorted; for instance, when g appears as $g(y)$ in an integral, it is undistorted and the integral is taken in the native space of g . However, if it appears as $g \circ T(x)$, then it is generally distorted and the integral is not taken in its native space.

One can see that the factor $J(x)$ in I_{bw} makes the two forward and backward CFs in general dissimilar, and the registration inverse-inconsistent. In other words, the uniform integral in the native space of one of the two images is generally non-uniform in the native space of the other image. CF symmetrization [3-7] gives:

$$\begin{aligned}
I_{ave}(f, g, T) &:= (I_{fw}(f, g, T) + I_{bw}(f, g, T))/2 \\
&= \int_{\Omega} D(f(x), g \circ T(x)) \frac{1 + J(x)}{2} dx.
\end{aligned} \tag{3}$$

This average CF satisfies the inverse-consistency property, $I_{ave}(g, f, T^{-1}) = I_{ave}(f, g, T)$. Nevertheless, the images are integrated in their native spaces with the non-uniform measure $(1 + J(x))/2$, which varies for each voxel and at each iteration, resulting in inconsistently different contributions of different image regions toward the CF. In the alternative methods that integrate the CF uniformly on a ‘‘mid-space’’ [8-11], the equivalent integrals on the native spaces of the images are, once again, non-uniform.

It should be noted that minimizing I_{ave} results in transformations that are not necessarily minimizers of the native CFs, and for which, I_{fw} and I_{bw} are not necessarily equal. Here, we define the *native symmetry* property of an image registration algorithm, as the ability to produce transformations minimizing the native CFs, while keeping them equal ($I_{fw} = I_{bw}$). Rigid, and more generally volume-preserving [12-14] registration is an example of natively symmetric registration (disregarding the asymmetry due to resampling artifacts [2, 15, 16]), since with $J(x) = 1$ everywhere, the two native CFs are always equal. Yet, it can be seen that symmetrized non-rigid registration (minimizing I_{ave}) is symmetric, but not guaranteed to be natively symmetric.

In this work, we seek to achieve native symmetry in non-rigid image registration, which will offer symmetry without compromising the uniformity of the integrals in the native CFs. The idea is based on the fact that the difference between I_{fw} and I_{bw} emerges from regions for which both $J(x) \neq 1$ and $D(f(x), g \circ T(x)) > 0$. We propose to keep the deformation quasi-volume-preserving (QVP) by restricting $J(x)$ to remain close to 1 in regions where the difference between the two images is large, but relax this restriction where the images are similar. Mathematically speaking, we make sure that,

$$D(f(x), g \circ T(x)) |J(x) - 1| < \epsilon, \quad \forall x \in \Omega \tag{4}$$

with ϵ being a user-defined positive threshold. This way, the integrals in I_{fw} , I_{bw} , and I_{ave} – when all represented in the native space of one image, as in Eqs. (1,2,3) – remain (almost) uniform, except for areas where non-uniformity barely perturbs the CF. In fact, one can verify that choosing a small ϵ makes the global native asymmetry error arbitrarily small, since it will be bounded as,

$$|I_{fw}(f, g, T) - I_{bw}(f, g, T)| < \epsilon |\Omega|, \tag{5}$$

$$|I_{ave}(f, g, T) - I_{fw/bw}(f, g, T)| < \epsilon |\Omega|/2.$$

Moreover, the local nature of the QVP constraint also ensures regional native symmetry; i.e., satisfying Inequality (4) keeps native CFs defined in any local $\Omega' \subset \Omega$ within $\epsilon |\Omega'|$ of each other.

In our Demons-based [17] implementation, we minimize I_{ave} via gradient descent in a similar manner as in [6]. We have two types of regularization. We perform a general Tikhonov regularization by including the standard term $\lambda \|\partial T(x)\|_f^2$ inside the integral of I_{fw} , and adjusting I_{bw} and I_{ave} accordingly. (The constant λ determines the amount of this general regularization.) Furthermore, we account for the QVP constraint by smoothing the transformation field T inhomogeneously until Inequality (4) is satisfied. At each iteration, we solve the following diffusion equation on the transformation with an inhomogeneous diffusion coefficient that depends on the local asymmetry error:

$$T_{t+1}(x) = T_t(x) + \gamma \nabla \cdot [K_t(x) \nabla T_t(x)], \tag{6}$$

where t and γ are the diffusion iteration number and step size, respectively, and the diffusion coefficient $K_t(x)$ is a function of the (low-pass filtered) local asymmetry error $e_t(x) := D(f(x), g \circ T_t(x)) |J_t(x) - 1|$ (the left-hand side of Inequality (4)). In order to only penalize regions where the error is above the threshold ϵ , we hard-threshold K by making it grow linearly with the error when it is larger than ϵ , and stay small and convex otherwise, as follows:

$$K(x) := \begin{cases} e(x) & , \quad e(x) > \epsilon \\ \epsilon \cdot (e(x)/\epsilon)^4 & , \quad 0 \leq e(x) \leq \epsilon \end{cases} \tag{7}$$

The diffusion iterations are repeated until Inequality (4) is satisfied. This inequality is the major contrast between our proposed approach and the volume-preserving methods in the literature [12-14]; our algorithm preserves volume *only* in regions where it finds no good match between the images. This means that the transformation may still contain compression and expansion in areas where the two images are locally similar, allowing good matching in regions where it is feasible.

3. RESULTS AND DISCUSSION

We validate our method by comparing the performance of the following three approaches (without diffeomorphism constraints): minimization of I_{fw} , minimization of I_{ave} , and minimization of I_{ave} with the proposed QVP constraint (with the heuristically-determined optimum ϵ). For each experiment, we plot I_{fw} and I_{bw} with respect to the iteration number to assess the registration native symmetry by the extent to which these two native CFs agree with each other. We also evaluate the overall results by the number of iterations that each algorithm takes to converge, and the ultimate value of the CFs. For each method and experiment, we choose the optimum general regularization parameter λ heuristically to achieve the best convergence with no observable topology break.

We first compare the three methods on synthetic data, by registering two letters ‘B’ with noticeably different shapes. As the

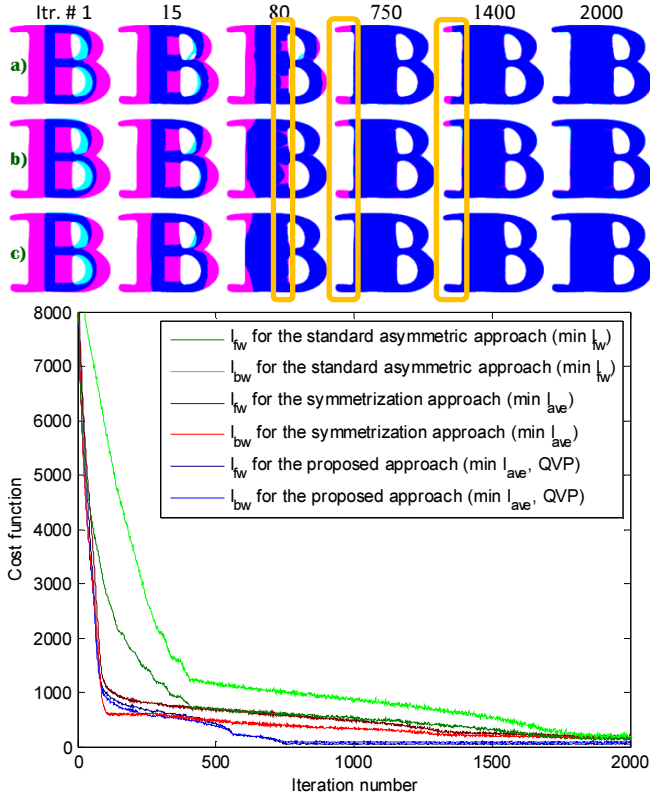


Fig. 1. Registration of synthetic data by (a) the standard asymmetric, (b) symmetrization, and (c) the proposed QVP approaches. Top: Cyan and magenta indicate respectively the reference and the interpolated images, with blue being their intersection. Bottom: The native CFs for each method.

plot in Fig. 1 demonstrates, minimizing I_{fw} (dark and light green) results in a relatively large difference $I_{bw} - I_{fw} > 0$, which is expected, since I_{bw} is not in any form accounted for in the optimization. Minimizing I_{ave} (dark and light red) reduces the gap between I_{fw} and I_{bw} , as equal weight is given to both of them in the optimization. Nevertheless, this approach – not being natively symmetric – does not guarantee the equality of the two native CFs. The proposed QVP approach (dark and light blue), however, enforces a bound on $|I_{fw} - I_{bw}|$ and therefore, as illustrated in the plot, produces transformations on the cost of which both native CFs agree (within the error margin ϵ). In addition, this method converges at an earlier gradient descent iteration (\sim Itr. 750, see the orange rectangles in Fig. 1) compared to the two others (\sim Itr. 1800),¹ while resulting in lower final native CF values ($I_{fw} = 56, I_{bw} = 87$) than the standard asymmetric ($I_{fw} = 152, I_{bw} = 202$) and the symmetrized ($I_{fw} = 188, I_{bw} = 162$) approaches do.

The better convergence of the QVP method may be related to the fact that its optimum general regularization parameter λ is naturally lower compared to the other two techniques, given that it also performs a separate inhomogeneous regularization (Eq. (6)). This method allows the transformation to freely absorb one ‘B’ into the other, since the corresponding volume change happens in regions where the intensities match, for which the inhomogeneous regularization, and consequently the total regularization is weak.

¹ Each iteration is, however, on average more expensive in the QVP method than in the other two methods, given the extra step of diffusion.

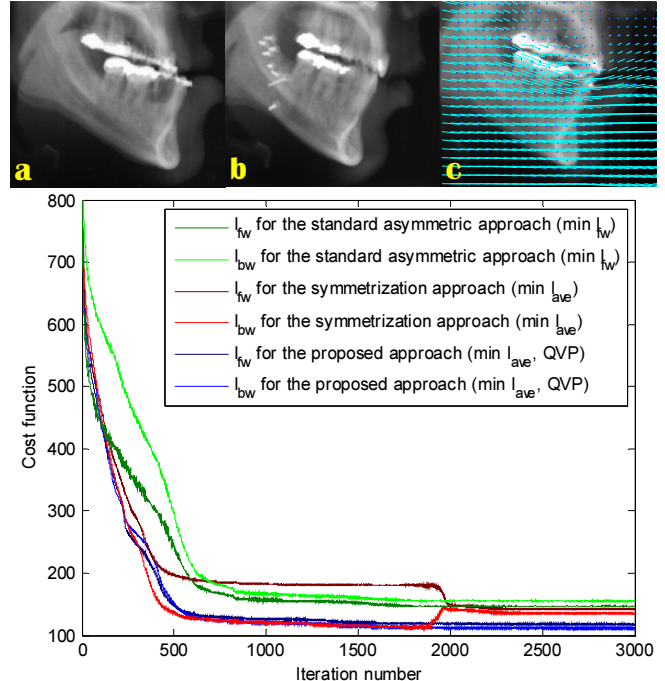


Fig. 2. Top: X-ray images taken (a) before and (b) after surgery, with (c) registration results using the proposed QVP approach. Bottom: Native CFs for the three methods.

The same inhomogeneous regularization, in contrast, prevents a topology break by regularizing more aggressively in areas where the difference between the two images (and therefore the SSD force) is large. On the contrary, the other two methods need a higher λ to avoid topology break, and yet apply this strong general regularization to the entire image (including regions with matching intensities and no gradient force), hence a slower convergence.

Next, we compared the three methods on a pair of public jaw X-ray images acquired before (Fig. 2a) and after (Fig. 2b) an orthognathic surgery. The images were initially aligned rigidly so their upper jaws overlap completely. As in the previous case, minimization of the standard asymmetric forward CF, I_{fw} (Fig. 2, the two green curves), soon creates a considerable gap between the native CFs. Minimizing I_{ave} (Fig. 2, the two red curves) results in a steady optimization of the average of the native CFs. However, since the native CFs are not individually considered, not only does the gap between them still mostly exist, but a jump that abruptly changes them – yet not their average – occurs around iteration 2000, which is a possible sign of the potential instability of the algorithm. In fact, an optimization algorithm might decide that it has converged and stop the procedure way before iteration 2000, thus leaving the teeth region misaligned (as was actually the case here for iterations before 2000). Conversely, the proposed QVP method (Fig. 2, the two blue curves), produces the smallest native CF gap, while obtaining the lowest final CF values ($I_{fw} = 118, I_{bw} = 112$) compared to the standard asymmetric ($I_{fw} = 147, I_{bw} = 155$) and the symmetrized ($I_{fw} = 142, I_{bw} = 135$) approaches. The warped image obtained by the QVP approach is depicted in Fig. 2c, along with the computed deformation field. Results by the two other methods were visually similar, except for the fact that the displacement in the upper jaw, which is supposed to be small because of the initial rigid alignment of the upper jaws, was lowest in the deformation obtained by the QVP approach (2.5 pixels, averaged in an upper jaw mask), and therefore least affected

by the general regularization, compared to the standard (3.9 pixels) and the symmetrization (3.2 pixels) approaches. This is likely to be due to the same reason as explained in the previous paragraph.

Lastly, we tested the three algorithms on the mid-sagittal planes of 20 brain images taken from the publicly available OASIS database [18], which we pre-processed in FreeSurfer [19]. The intensity-normalized and resampled volumes (1-mm³ isotropic voxel size) were made upright by robust rigid registration [2] of each volume to its left-right mirrored version. The sagittal slice located four voxels right to the mid-sagittal plane was extracted from each volume, and to adjust for any nodding rotation, was rigidly registered to that slice of the first volume, and resampled to the size 128×128. Out of the 20 sagittal slices corresponding to the 20 subjects, the one closest (in L₂ norm) to the rest was chosen as the reference, and non-rigidly registered to the rest of them individually. For each subject and method, we ran the registration with 21 different values for general regularization parameter λ , and then inspected the results and chose the one with lowest λ containing no visible topology break. The optimally chosen λ was the same for the symmetrization and the QVP methods in 14 subjects; however, for the other 6 subjects, QVP passed the visual inspection test at a lower λ (see Fig. 3 for examples). We computed the native asymmetry error as the mean absolute value of the difference of the two native CFs through all iterations. When averaged across subjects, this error was 161% and 16% higher for the asymmetric and the symmetrization techniques, respectively, compared to the proposed QVP approach. We also compared the final values of the native CFs (I_{fw}, I_{bw}) among the methods, which were (30%, 56%) and (5%, 4%) higher for the asymmetric and the symmetrization approaches, respectively, compared to the QVP method. We hypothesized that the QVP approach results in ultimate native CF values that are lower than those of the two other techniques. A left-tailed Student's *t*-test with a 0.05 significance level rejected the null hypothesis when comparing the QVP with both the asymmetric ($p=10^{-6}$) and the symmetrization ($p=0.02$) approaches. It should be noted, however, that our implementation did not include any explicit diffeomorphism constraints. Validation of the proposed methodology using state-of-the-art 3D diffeomorphic implementations, and exploring quasi-rigidity (as opposed to QVP) constraints, are subjects of ongoing research.

Acknowledgments: Support for this research was provided in part by the National Center for Research Resources (P41-RR14075, and the NCRR BIRN Morphometric Project BIRN002, U24 RR021382), the National Institute for Biomedical Imaging and Bioengineering (R01EB006758), the National Institute on Aging (AG022381), the National Center for Alternative Medicine (RC1 AT005728-01), the National Institute for Neurological Disorders and Stroke (R01 NS052585-01, 1R21NS072652-01, 1R01NS070963), and was made possible by the resources provided by Shared Instrumentation Grants 1S10RR023401, 1S10RR019307, and 1S10RR023043. Additional support was provided by The Autism & Dyslexia Project funded by the Ellison Medical Foundation, and by the NIH Blueprint for Neuroscience Research (5U01-MH093765), part of the multi-institutional Human Connectome Project. Dr. Sabuncu received support from an NIH K25 grant (NIBIB 1K25EB013649-01) and an AHAF AD pilot grant A2012333.

REFERENCES

[1] A. Sotiras, C. Davatzikos, and N. Paragios, *Deformable medical image registration: A survey*, RR-7919, INRIA, 2012.
 [2] M. Reuter, H. D. Rosas, and B. Fischl, "Highly accurate inverse consistent registration: A robust approach," *NeuroImage*, vol. 53, 2010.
 [3] G. E. Christensen, and H. J. Johnson, "Consistent image registration," *Medical Imaging, IEEE Trans.*, vol. 20, no. 7, pp. 568–582, 2001.




Fig. 3. (a) QVP registration results on four subjects with optimal values of λ . (b) Results of the symmetrization method using the same λ as in (a). Note the distortion due to under-regularization in the encircled areas. (c) Results of the symmetrization method using optimal (higher) values of λ , thus with higher CF values.

[4] P. Cachier, and D. Rey, "Symmetrization of the non-rigid registration problem using inversion-invariant energies: Application to multiple sclerosis," *Proc. MICCAI*, pp. 697–708, 2000.
 [5] A. Trounev, and L. Younes, "Diffeomorphic matching problems in one dimension: Designing and minimizing matching functionals," *ECCV*, 2000.
 [6] A. Leow, S.-C. Huang, A. Geng, J. Becker, S. Davis, A. Toga, and P. Thompson, "Inverse consistent mapping in 3D deformable image registration: Its construction and statistical properties," *IPMI*, 2005.
 [7] H. Tagare, D. Groisser, and O. Skrinjar, "Symmetric non-rigid registration: A geometric theory and some numerical techniques," *J. Mathematical Imaging and Vision*, vol. 34, no. 1, pp. 61–88, 2009.
 [8] P. Lorenzen, B. Davis, and S. Joshi, "Model based symmetric information theoretic large deformation multi-modal image registration," *Proc. ISBI*, pp. 720–723, 2004.
 [9] M. F. Beg, and A. Khan, "Symmetric data attachment terms for large deformation image registration," *Medical Imaging, IEEE Transactions on*, vol. 26, no. 9, pp. 1179–1189, 2007.
 [10] B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee, "Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain," *Medical Image Analysis*, vol. 12, no. 1, pp. 26–41, 2008.
 [11] V. Noblet, C. Heinrich, F. Heitz, and J.-P. Armpach, "Symmetric nonrigid image registration: Application to average brain templates construction," *Proc. MICCAI*, pp. 897–904, 2008.
 [12] G. Hirota, R. Maheshwari, and M. C. Lin, "Fast volume-preserving free-form deformation using multi-level optimization," *Computer-Aided Design*, vol. 32, no. 8–9, pp. 499–512, 2000.
 [13] T. Rohlfing, C. R. Maurer, Jr., D. A. Bluemke, and M. A. Jacobs, "Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint," *Medical Imaging, IEEE Transactions on*, vol. 22, no. 6, pp. 730–741, 2003.
 [14] E. Haber, and J. Modersitzki, "Image registration with guaranteed displacement regularity," *International Journal of Computer Vision*, vol. 71, no. 3, pp. 361–372, 2007.
 [15] I. Aganj, B. Yeo, M. Sabuncu, and B. Fischl, "On removing interpolation and resampling artifacts in rigid image registration," *Image Processing, IEEE Transactions on*, in press, 2012.
 [16] M. Sabuncu, B. Yeo, K. Van Leemput, T. Vercauteren, and P. Golland, "Asymmetric image-template registration," *MICCAI*, 2009.
 [17] J. P. Thirion, "Image matching as a diffusion process: An analogy with Maxwell's demons," *Medical Image Analysis*, vol. 2, 1998.
 [18] D. S. Marcus, T. H. Wang, J. Parker, J. G. Csernansky, J. Morris, and R. L. Buckner, "Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults," *J. Cog. Neurosci.*, vol. 19, no. 9, 2007.
 [19] B. Fischl, "FreeSurfer," *NeuroImage*, vol. 62, pp. 774–781, 2012.