

Topology Preservation and Anatomical Feasibility in Random Walker Image Registration

Shawn Andrews, Lisa Tang, and Ghassan Hamarneh

Medical Image Analysis Lab, Simon Fraser University, Canada

Abstract. The random walker image registration (RWIR) method is a powerful tool for aligning medical images that also provides useful uncertainty information. However, it is difficult to ensure topology preservation in RWIR, which is an important property in medical image registration as it is often necessary for the anatomical feasibility of an alignment. In this paper, we introduce a technique for determining spatially adaptive regularization weights for RWIR that ensure an anatomically feasible transformation. This technique only increases the run time of the RWIR algorithm by about 10%, and avoids over-smoothing by only increasing regularization in specific image regions. Our results show that our technique ensures topology preservation and improves registration accuracy.

1 Introduction

Medical image registration (MIR), which is finding a spatial transformation that maps anatomical objects in one image to corresponding objects in another image, is a key step in many medical analysis tasks, including tracking disease progression, multi-modal fusion, shape analysis, and atlas construction. To encourage anatomical feasibility, regularization is imposed on the transformation to encourage smoothness and topology preservation (TP). However, regularizing sufficiently to ensure anatomical feasibility without over-smoothing and losing accuracy can be difficult or computationally expensive. A variety of techniques have been proposed to ensure TP, including implicitly, e.g. by expressing transformations as integrals over vector fields [1], or explicitly, e.g. by penalizing the lack of TP in the regularization objective term [9].

Recently, registration approaches that represent transformations in the discrete domain have arisen [4,3,11,12]. In discrete approaches every pixel (or voxel) is assigned a displacement vector from a predefined set, referred to as a *discrete transformation*, allowing the image registration energy to be formulated as a Markov random field (MRF) and well established optimization techniques such as graph cuts [2] to be utilized.

The discrete random walker image registration (RWIR) algorithm [3] uses a globally optimizable Gaussian MRF energy for regularization and has been shown to achieve results comparable to other state-of-the-art registration schemes [3,10,7]. RWIR provides a probabilistic registration that can be leveraged to calculate uncertainty. Registration uncertainty can be utilized to direct clinicians to possible registration errors or image abnormalities, and can influence diagnostic and therapeutic decisions. [8,6].

A useful feature of RWIR is that it seamlessly allows for spatially adaptive regularization weights, which have been shown to improve registration accuracy compared to

a constant regularization weight [11]. However, it is not clear a priori how much regularization is required in RWIR to ensure topology preservation. One option to deal with this is time consuming trial and error. Alternatively, post-processing techniques such as scaling and squaring [1] can convert a transformation to one that preserves topology, but such methods discard the uncertainty information provided by RWIR.

In this paper, we introduce a technique that determines if a set of regularization weights for RWIR will result in a topology preserving transformation without having to run the RWIR algorithm. We use our technique to iteratively increase regularization in regions where topology is not preserved, avoiding over-smoothing in other regions.

2 Anatomical Feasibility in RWIR

2.1 RWIR Review

In RWIR, we wish to align two images, I_1 and I_2 , defined over Ω , a set of n pixels (or voxels). We define $\hat{\Omega}$ as the n by d matrix of pixel locations, where d is the dimensionality of the images. Transformations are represented using a discrete set of K displacements, $V = [\mathbf{v}_1, \dots, \mathbf{v}_K]$, where $\mathbf{v}_k \in \mathbb{R}^d$. A *probabilistic transformation* assigns, to each pixel a , probabilities for the displacement vectors, $\mathbf{u}^a = [u_1^a, \dots, u_K^a] \in \mathbb{P}_K$. $\mathbb{P}_K \subset \mathbb{R}^K$ is the unit simplex of positive vectors that are normalized to sum to one. \mathbf{u}_k denotes a vector of length n of the probabilities of \mathbf{v}_k at each pixel.

The first step in RWIR is to calculate *prior probabilities* for each pixel a and displacement \mathbf{v}_k from a data similarity term $f_{sim}(I_1, I_2, a, \mathbf{v}_k)$. The prior probabilities at a are given by $\mathbf{p}^a = [p_1^a, \dots, p_K^a] \in \mathbb{P}_K$, $p_k^a = f_{sim}(I_1, I_2, a, \mathbf{v}_k)/Z^a$, where Z^a normalizes.

The second step is to construct a weighted image graph, with a node for each pixel and a weighted edge between neighboring pixels, and use it to regularize the displacement labels. The edge weights are stronger for pixels of similar intensities; for pixels a and b , and $\beta = 50$, we use edge weight $w_{a,b} = \exp(-\beta|I(a) - I(b)|)$. Defining W as the n by n matrix of edge weights and D as the n by n diagonal matrix of the row sums of W , then $L = D - W$ is known as the combinatorial Laplacian of the image graph. The probabilities \mathbf{u}_k are then calculated by minimizing the energy:

$$E_{RW}(\mathbf{u}_k) = \mathbf{u}_k^\top L \mathbf{u}_k + (\mathbf{u}_k - \mathbf{p}_k)^\top \Gamma^{-1} (\mathbf{u}_k - \mathbf{p}_k), \quad k \in \{1, \dots, K\}, \quad (1)$$

where Γ is an n by n diagonal matrix of spatially varying weights $\gamma_1, \dots, \gamma_n$, controlling the trade-off between regularization and image similarity. We note the larger the value for γ_a , the more regularization is applied at pixel a . We must solve (1) for every $k \in \{1, \dots, K\}$, but denoting $U = [\mathbf{u}_1, \dots, \mathbf{u}_K]$ and $P = [\mathbf{p}_1, \dots, \mathbf{p}_K]$, n by K matrices, we can combine the K problems into one linear system of equations, which we solve for U to get the RWIR probabilistic transformation:

$$(L + \Gamma^{-1})U = \Gamma^{-1}P. \quad (2)$$

2.2 Anatomical Feasibility for Probabilistic Transformations

TP primarily requires that the Jacobian¹ of T , $J(T)$, is positive everywhere. We adopt the convention that a probabilistic transformation is TP if its *expected transformation*,

¹ For brevity, we refer to the Jacobian determinant as simply the Jacobian.

with the expected displacement taken at each pixel, has a positive Jacobian everywhere. Formally, the expected transformation corresponding to a probabilistic transformation U is given by:

$$T^* = \hat{\Omega} + UV^\top. \quad (3)$$

T^* is a discrete transformation, whose a^{th} row corresponds to the location pixel a is mapped to. $J(T)$ is defined in the continuous case as a function of the derivatives of the components of T . So, to discretize the Jacobian, we must choose a discrete approximation of the derivative operators. It has been shown that if all combinations of forward and backward difference operators along the d axis directions, Δ_i^F and Δ_i^B , $i \in \{1, \dots, d\}$, are used to construct 2^d approximate Jacobians, and if they are all positive everywhere, then the continuous bi- or tri-linear interpolated version of T^* will also have positive Jacobian everywhere [5]. We thus define the discrete Jacobian at pixel a , $J_a(T)$, to be the minimum of the 2^d approximate Jacobians.

We note that while TP is often a necessary condition for anatomical feasibility, it is not always sufficient: transformations exhibiting excessive stretching may be anatomically infeasible. These conditions can be identified by pixels with very large or small (but still positive) Jacobians.

Our goal in the following sections is to choose spatially adaptive weights $\gamma_1, \dots, \gamma_n$ in such a way that U , calculated from (2), corresponds to an anatomically feasible expected transformation T^* with respect to its Jacobian values.

2.3 Efficient Jacobian Calculation

In this section, we introduce a technique for calculating the Jacobian of the expected transformation $J(T^*)$ *without* needing to calculate the probabilistic transformation U . This technique is orders of magnitude faster than the calculation of U via (2), allowing regularization to be increased where needed, *prior* to running the RWIR algorithm.

In order to calculate $J(T^*)$, where $T^* = [\mathbf{t}_1^*, \dots, \mathbf{t}_d^*]$, we must calculate $\Delta_i \mathbf{t}_j^*$ for $i, j \in \{1, \dots, d\}$, where $\Delta_i \in \{\Delta_i^F, \Delta_i^B\}$. Combining (2) and (3) gives

$$\Delta_i \mathbf{t}_j^* = \Delta_i \left(\hat{\Omega} \mathbf{e}_j + (L + \Gamma^{-1})^{-1} \Gamma^{-1} P V^\top \mathbf{e}_j \right), \quad (4)$$

where \mathbf{e}_j is the j^{th} standard basis vector for \mathbb{R}^d . We define $\mathbf{b}_j = \Gamma^{-1} P V^\top \mathbf{e}_j$ and $\mathbf{x}_j = (L + \Gamma^{-1})^{-1} \mathbf{b}_j$, and note that \mathbf{x}_j can be calculated without performing an expensive matrix inversion by solving the system of equations

$$(L + \Gamma^{-1}) \mathbf{x}_j = \mathbf{b}_j. \quad (5)$$

We note the similarity between (2) (used to solve for U) and (5). The difference is that in (5), the right hand side has 1 column, whereas in (2) the right hand side has K columns, which is on the order of hundreds or even thousands in 3D registration.

Solving for \mathbf{x}_j allows us to rewrite (4) as:

$$\Delta_i \mathbf{t}_j^* = \Delta_i \hat{\Omega} \mathbf{e}_j + \Delta_i \mathbf{x}_j. \quad (6)$$




Fig. 1. A 1D example with circles representing pixels, unit distance apart, and arrows representing their displacement vectors. (i) The red pixels have crossing displacements and thus a negative Jacobian. (ii) After increasing regularization at the red pixels, the orange and purple pairs of pixels now have crossing displacements instead. (iii) Increasing the regularization of pixels based on their proximity to the location of the negative Jacobian in (i) results in a TP transformation.

Once (6) is solved for each (i, j) pair, we can calculate $J(T^*)$. By repeating this process for each combination of forward and backward operators, we will determine if and where the discrete Jacobian of the expected transformation is negative.

The calculations in this section consist mainly of efficient multiplications between sparse matrices, with the exception being the need to solve (5) for each $j \in \{1, \dots, d\}$; this is the computational bottleneck. Since d is usually only 2 or 3, this bottleneck is still orders of magnitude faster than the full RWIR algorithm. This leads us to the strategy of repeatedly calculating $J(T^*)$ with different settings for Γ until we determine how much regularization is needed at each pixel in order to ensure an anatomically feasible Jacobian, as we detail in the next section.

2.4 Determining Regularization Weights

Using the technique from the previous section, one could increase regularization slightly wherever the Jacobian is negative, recalculate the Jacobians, and iterate. This would ensure TP and avoid over-regularization, but may require many iterations. We reduce the number of iterations by using larger increases in regularization each iteration and preemptively increasing regularization in a neighborhood around pixels with negative Jacobian (as the non-invertibility may be “pushed out” to nearby pixels, see Fig. 1).

To make this concept concrete, we define a function $\phi_1 : \Omega \rightarrow \mathbb{R}^+$ evaluating how much additional regularization each pixel needs:

$$\phi_1(a) = \lambda_1 \max_{b \in \Omega} \left(\underbrace{\max(0, -J_b(T^*))}_{\text{Amount of Violation}} - \lambda_2 \underbrace{\|a - b\|_2}_{\text{Proximity}} \right), \quad (7)$$

where λ_1 and λ_2 are positive scalar parameters. Larger values for λ_1 and λ_2 speed convergence, but may result in over-regularization if too large. We found $\lambda_1 = 0.3$ and $\lambda_2 = 5$ provide a good trade-off. We increase the regularization at pixel a by updating the trade-off weight γ_a :

$$\gamma_a \leftarrow \gamma_a \cdot (1 + \phi_1(a)). \quad (8)$$

The update values for each pixel can be efficiently computed using dynamic programming. With the updated Γ , $J(T^*)$ can be recalculated, and the update (8) performed

again until all Jacobian values are positive. A TP probabilistic transformation U can then be calculated using the updated matrix Γ .

While this method will ensure TP, it can easily be extended to enforce other restrictions on the Jacobian. For example, one may want a volume preserving transformation that has Jacobian values close to 1. We can use $J(T^*)$ to identify pixels with Jacobians that are deemed too small or too large, e.g. outside of the range $[\lambda_3, \lambda_4]$, $\lambda_3 < \lambda_4$. We define a second function to capture this idea:

$$\phi_2(a) = \lambda_1 \max_{b \in \Omega} \left(\max \left(0, \underbrace{\lambda_3 - J_b(T^*)}_{J \text{ Too Small}}, \underbrace{J_b(T^*) - \lambda_4}_{J \text{ Too Large}} \right) - \lambda_2 \underbrace{\|a - b\|_2}_{\text{Proximity}} \right). \quad (9)$$

By replacing ϕ_1 with ϕ_2 in (8) we not only regularize away negative Jacobian values, but also Jacobian values that we deem anatomically infeasible. This targeted approach avoids over-regularizing regions that already have feasible transformations, allowing those regions to stay loyal to their prior probabilities.

3 Results

In this section, we perform experiments to demonstrate the benefits of our spatially adaptive regularization scheme. We compare registration results from RWIR with and without our technique for choosing Γ . We also establish the correlation between uncertainty and error in probabilistic transformations, indicating that uncertainty information should not be discarded by post-processing techniques designed for non-probabilistic transformations. We focus on comparisons to RWIR instead of other techniques for ensuring TP since previous MICCAI papers [3,10,7] have shown RWIR to be comparable to other state-of-the-art techniques, and this comparison removes confounding factors.

Our experiments are performed on 40 T1-MR volumetric thigh images of dimension $250 \times 250 \times 40$, each segmented into 16 regions, including 11 different muscles (Fig. 3). This data set was used because it contains regions with rich details leading to accurate priors (e.g. around the bone) and other regions that are largely devoid of detail (e.g. the homogeneous thigh muscles), so we expect well chosen spatially adaptive regularization weights to be important for accurate registration. f_{sim} is defined as the negative exponential of the sum of squared intensity differences in a patch of size $5 \times 5 \times 5$. We used unoptimized MATLAB code run on a machine with 2 Quad Core Intel Xeon 2.33 GHz CPUs.

3.1 Synthetic Warpings

We applied known warps to each image and attempted to recover these warps using RWIR. We generated 5 warps for each image by randomly displacing B-spline control points, spaced 30 mm apart, where the displacements were sampled uniformly from vectors up to 8 mm in magnitude.

For each of the 200 image/warp pairs, we compare the results of RWIR run using 3 different matrices Γ . First, we use a spatially constant $\Gamma = cI$, where I is the identity

Table 1. The results of running RWIR with different regularization matrices. RWIR-TP achieves less error than the other registrations while also ensuring topology preservation, while only requiring marginally longer computation time.

	TRE (mm)	% of Pixels with $J \leq 0$	Run Time (sec)
Priors P	2.83 ± 0.18	16.6 ± 1.0	—
RWIR-C	2.56 ± 0.24	1.67 ± 0.54	4957 ± 533
RWIR-U	1.70 ± 0.19	0.36 ± 0.11	5056 ± 561
RWIR-TP	1.64 ± 0.19	0.0 ± 0.0	5657 ± 577




Fig. 2. Mean TRE for different uncertainty percentiles. The mean TRE of the top $k\%$ most uncertain pixels was calculated, for $k = \{1, \dots, 99\}$. The mean TRE consistently increases across percentiles, indicating uncertainty could be used to identify errors in a registration. Particularly, we see a sharp increase in TRE for the top 10% most uncertain pixels.

matrix and c is a scalar. We refer to the results of this registration as RWIR-C. Second, we use uncertainty information from the prior probabilities P to construct Γ (based on an idea presented in previous works [10]):

$$\gamma_a = c \cdot \exp\left(\frac{-H(\mathbf{p}^a)}{H_{max}}\right), \quad H(\mathbf{p}^a) = \sum_{k=1}^K \sum_{\ell=1}^K p_k^a p_\ell^a \|\mathbf{v}_k - \mathbf{v}_\ell\|^2, \quad (10)$$

where H is a measure of uncertainty found to correlate well with registration error [6], and H_{max} is the maximum possible value for H . Using (10) results in stronger regularization for pixels with uncertain prior probabilities. We refer to the results of this registration as RWIR-U. Third, we use (8) with ϕ_2 to iteratively update Γ , initialized using (10) and iterating 10 times, which was found to be sufficient to ensure TP. We refer to the results of this registration as RWIR-TP. In both RWIR-C and RWIR-U, c was empirically set to minimize target registration error (TRE). The Jacobian is restricted to the range $[0.5, 1.5]$, chosen based on the minimum and maximum Jacobian values of a random synthetic warping, constructed as stated above.

A comparison of the results achieved by the 3 registrations is shown in Table 1. The regularization used in RWIR-TP ensures a positive Jacobian, while only requiring about 10% longer to run. Further, RWIR-TP achieves less error than RWIR techniques.

To demonstrate the benefits of a probabilistic registration, for each test we calculated uncertainty values for each pixel by applying H from (10) to the displacement probabilities generated by RWIR-U and RWIR-TP. We found a Pearson correlation coefficient between the uncertainty and the TRE of 0.45 for RWIR-U and 0.50. Fig. 2 demonstrates the relationship between uncertainty and TRE by taking the top $k\%$ most uncertain pixels, for $k = \{1, \dots, 99\}$, and calculating their mean error. We see that error

Table 2. TOs from registering pairs of images, averaged across all regions and tests. RWIR-TP and RWIR-U achieve similar TOs, but RWIR-TP provides TP transformations. Note the target overlap of RWIR-U is slightly higher than RWIR-TP because TO is only an approximate measure of accuracy - an erroneous displacement (e.g. one causing negative Jacobian) may still map a voxel to the correct label.

	Original	Priors P	RWIR-U	RWIR-TP
Target Overlap (TO)	0.632 ± 0.108	0.762 ± 0.134	0.841 ± 0.071	0.833 ± 0.062
% of Pixels with $J \leq 0$	–	38.2 ± 8.10	3.37 ± 0.88	0.0 ± 0.0




Fig. 3. An example on real data demonstrating how our technique increases regularization locally where necessary without over-regularizing other areas. While the registration results are unchanged in most areas, RWIR-TP smooths out certain problem areas in the RWIR-U registration.

steadily increases for more uncertain pixels, and that the top 10% most uncertain pixels have significantly more error than an average pixel. Using post-processing techniques to ensure topological invertibility would discard these uncertainty values.

3.2 Real Data

In this section, we demonstrate the usefulness of our method in a real medical imaging application by taking each pair of thigh images and registering them to each other, using both RWIR-U and RWIR-TP with the same bounds on the Jacobian as in Sec. 3.1. We evaluate registration results using the target overlap, the number of correctly mapped foreground pixels divided by the total number of foreground pixels. We evaluate anatomical feasibility by looking at the number of pixels with negative Jacobian.

The results of this experiment are summarized in Table 2, and an example slice of the registered 3D images is seen in Fig. 3. We see that while the transformations resulting from RWIR-U and RWIR-TP have very similar target overlaps, the RWIR-U transformations have $3.37 \pm 0.88\%$ of pixels with negative Jacobian on average, and thus are not anatomically feasible transformations, whereas none of the RWIR-TP transformations have negative Jacobians, and thus represent anatomically feasible alignments (again, with only about 10% increased run time).

4 Conclusion

Properly selecting regularization parameters is often key to accurate registration, yet in general can be very difficult without expensive trial and error approaches, involving a registration being performed multiple times and the results examined. Our method efficiently adjusts spatially adaptive regularization in order to provide explicit guarantees regarding topology preservation without over-regularizing regions with accurate prior probabilities. For future work, we will more rigorously examine how the values of λ_1 and λ_2 affect the speed and accuracy of the registration for different applications.

References

1. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-Euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006, Part I. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006)
2. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. *IEEE TPAMI* 23(11), 1222–1239 (2001)
3. Cobzas, D., Sen, A.: Random walks for deformable image registration. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011, Part II. LNCS, vol. 6892, pp. 557–565. Springer, Heidelberg (2011)
4. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. *MedIA* 12(6), 731–741 (2008)
5. Karaçalı, B., Davatzikos, C.: Topology preservation and regularity in estimated deformation fields. In: Taylor, C.J., Noble, J.A. (eds.) IPMI 2003. LNCS, vol. 2732, pp. 426–437. Springer, Heidelberg (2003)
6. Lotfi, T., Tang, L., Andrews, S., Hamarneh, G.: Improving probabilistic image registration via reinforcement learning and uncertainty evaluation. In: Wu, G., Zhang, D., Shen, D., Yan, P., Suzuki, K., Wang, F. (eds.) MLMI 2013. LNCS, vol. 8184, pp. 187–194. Springer, Heidelberg (2013)
7. Popuri, K., Cobzas, D., Jägersand, M.: A variational formulation for discrete registration. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013, Part III. LNCS, vol. 8151, pp. 187–194. Springer, Heidelberg (2013)
8. Risholm, P., Pieper, S., Samset, E., Wells III, W.M.: Summarizing and visualizing uncertainty in non-rigid registration. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010, Part II. LNCS, vol. 6362, pp. 554–561. Springer, Heidelberg (2010)
9. Rohlfing, T., Maurer Jr., C.R., Bluemke, D.A., Jacobs, M.A.: Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE TMI* 22(6), 730–741 (2003)
10. Tang, L., Hamarneh, G.: Random walks with efficient search and contextually adapted image similarity for deformable registration. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013, Part II. LNCS, vol. 8150, pp. 43–50. Springer, Heidelberg (2013)
11. Tang, L., Hamarneh, G., Abugharbieh, R.: Reliability-driven, spatially-adaptive regularization for deformable registration. In: Fischer, B., Dawant, B.M., Lorenz, C. (eds.) WBIR 2010. LNCS, vol. 6204, pp. 173–185. Springer, Heidelberg (2010)
12. Tang, T., Chung, A.: Non-rigid image registration using graph-cuts. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 916–924. Springer, Heidelberg (2007)