Graph-based Deformable Image Registration

Aristeidis Sotiras*, Yangming Ou*, Nikos Paragios and Christos Davatzikos

Abstract Deformable image registration is a field that has received considerable attention in the medical image analysis community. As a consequence, there is an important body of works that aims to tackle deformable registration. In this chapter we review one class of these techniques that use discrete optimization, and more specifically Markov Random Field models. We begin the chapter by explaining how one can formulate the deformable registration problem as a minimal cost graph problem where the nodes of the graph corresponds to the deformation grid, the graph connectivity encodes regularization constraints, and the labels correspond to 3D displacements. We then explain the use of discrete models in intensity-based volumetric registration. In the third section, we detail the use of Gabor-based attribute vectors in the context of discrete deformable registration, demonstrating the versatility of the graph-based models. In the last section of the chapter, the case of landmark-based registration is discussed. We first explain the discrete graphical model behind establishing landmark correspondences, and then continue to show how one can integrate it with the intensity-based model towards creating enhanced models that combine the best of both worlds.

———————————————

Aristeidis Sotiras

Section of Biomedical Image Analysis, Center for Biomedical Image Computing and Analytics, University of Pennsylvania, Philadelphia, USA, e-mail: aristeidis.sotiras@uphs.upenn.edu

Yangming Ou

Athinoula A. Martinos Center for Medical Imaging, Massachusetts General Hospital, Harvard Medical School, Boston, USA e-mail: yangming.ou@mgh.harvard.edu

Nikos Paragios

Center for Visual Computing, Department of Applied Mathematics, Ecole Centrale Paris, Paris, France e-mail: nikos.paragios@ecp.fr

Christos Davatzikos

Section of Biomedical Image Analysis, Center for Biomedical Image Computing and Analytics, University of Pennsylvania, Philadelphia, USA, e-mail: christos.davatzikos@uphs.upenn.edu

* The first two authors contributed equally to this work.

1 Introduction

Medical image analysis plays an increasingly important role in many clinical applications. The increased amount and complexity of medical image data, which often involve multiple 3D image modalities as well as multiple acquisitions in time, result in a challenging analysis setting. Image registration, as well as image segmentation, are the two principal tools that allow for automatic and timely data analysis.

Image registration consists of determining a spatial transformation that establishes meaningful anatomical, or functional, correspondences between different image acquisitions. The term deformable is used to specify that the transformation is allowed to spatially vary (in contrast to the case of linear or global registration). In general, registration can be performed between two or more images. Nonetheless, in this chapter, we will focus on registration methods that involve pairs of images. The pairs of images may consist of acquisitions that image either the same subject (intra-subject registration) or different subjects (inter-subject registration).

In intra-subject registration, the subject is typically imaged either under different protocols, or at different time points. In the first case, different imaging modalities are used to capture complementary anatomical or functional information, and image registration is used to fuse this information towards enhancing the analytical and diagnostic abilities of the clinicians. In the second case, one may study short- or long-term longitudinal processes that range from tumor perfusion properties to normal aging and development. Another application of image registration is surgical or treatment planning. The registration of pre-operative and interventional data allows the clinical experts to refine their planning and improve care-giving.

Inter-subject registration is the cornerstone of population studies. Mapping members of a population to a common domain allows the study of within-population variability and the quantitative analysis of the form of anatomical structures. On the other hand, when distinct populations are spatially aligned, it is possible to discover the focal differences that distinguish them by contrasting them in the common domain.

In general, an image registration algorithm involves three components (see Fig.1 [75]): i) a transformation model; ii) a similarity criterion; and iii) an optimization method. Image registration has been studied extensively during the past decades, leading to a rich body of works. These works differ mainly in their choices with respect to these three components. While an extensive overview of these components is beyond the scope of this chapter, let us briefly discuss some of the most common choices and models. For a more comprehensive review, we refer the interested reader to the books [26, 54], the surveys [53, 78, 93] and [75] that provide thorough overviews of the advances of the past decades in deformable registration.

The choice of the transformation model is usually dictated by the application at hand and is related to the nature of the deformation to be recovered. High-dimensional nonlinear models are necessary to cope with highly variable soft tissue, while low degrees of freedom models can represent the mapping between rigid bone structures. It is important to note that increasing the degrees of freedom of the
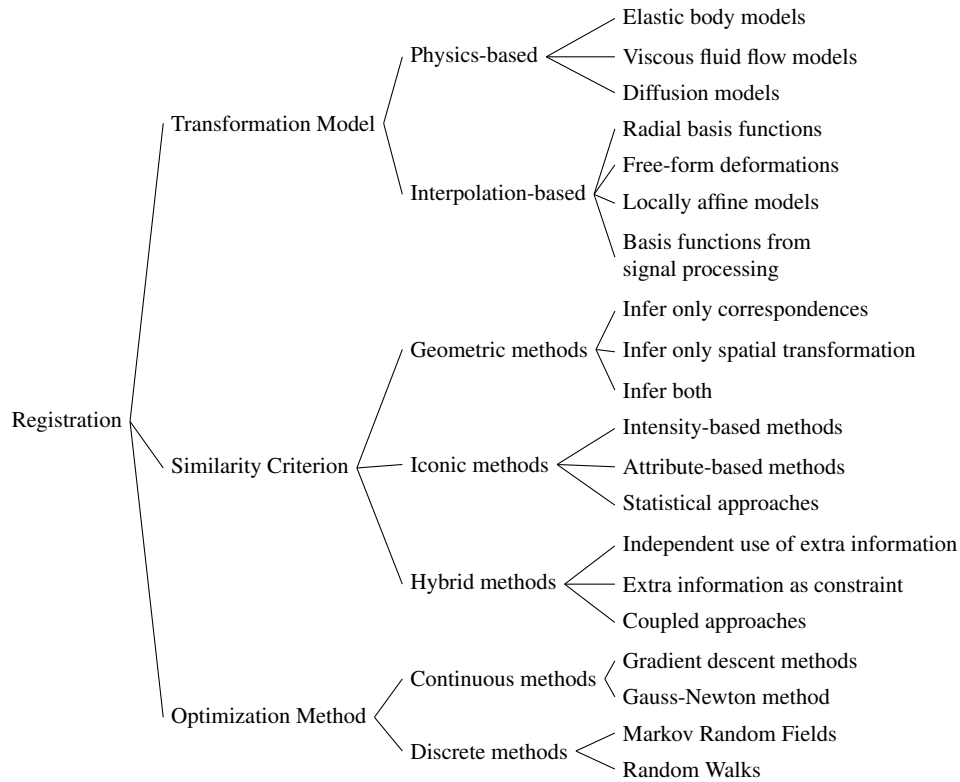
Fig. 1: Typical components of registration algorithms.

model, and thus enriching its descriptive power, often comes at the cost of increased computational efficiency.

Several transformation models have been introduced in medical imaging for non-rigid alignment. These models can be coarsely classified into two categories (see Fig.1 [30, 75]): i) models derived from physical models, and ii) models derived from the interpolation theory or geometric models. Among the most prominent choices of the first class, one may cite elastic [19, 20], fluid [14, 18] or diffusion models [22, 80, 85]. Whereas, the second class comprises radial basis functions [10, 67], free-form deformations [69, 68], locally affine [55] and poly-affine models [2], or models parametrized by Fourier [1, 4] or Wavelet basis functions [87].

The similarity criterion quantifies the degree of alignment between the images. Registration methods can be classified into three categories (see Fig.1) depending on the type of information that is utilized by the similarity criterion: i) geometric registration (*a.k.a.* landmark/feature-based registration); ii) iconic registration (*a.k.a.* voxel-wise registration); and iii) hybrid registration.

Geometric registration aims to align meaningful anatomical locations or salient landmarks, which are either automatically extracted from the images [51] or pro-

vided by an expert. Geometric information is typically represented as point-sets and registration is tackled by first estimating the point correspondences [43, 82] and then employing an interpolation strategy (*e.g.* thin-plate splines [10]) to determine a dense deformation field that will align the images. Alternatively, geometric methods may infer directly the transformation that aligns the images without explicitly estimating point correspondences. This is possible by representing geometric information either as probability distributions [23, 83] or through the use of signed distance transformations [32]. Last, there exist methods that opt to simultaneously solve for both the correspondences and the transformation [15].

Iconic methods employ a similarity criterion that takes into account the intensity information of all image elements. The difficulty of choosing an appropriate similarity criterion varies greatly depending on the problem. In the mono-modal case, where both images are acquired using the same device and one can assume that the intensity profiles for the two images differ only by Gaussian noise, the use of sum of squared differences can be sufficient. Nonetheless, in the multi-modal case, where images from different modalities are involved, the criterion should be able to account for the different principles behind the acquisition protocols and capture the relation between the distinct intensity profiles. Towards this end, criteria based on statistics and information theory have been proposed. Examples include correlation ratio [65], mutual-information [86, 52] and Kullback-Leibler divergence [16]. Last, attribute-based methods that construct rich descriptions by summarizing intensity information over local regions have been proposed for both mono-modal and multi-modal registration [74, 48, 62].

Hybrid methods opt to exploit both iconic and geometric information in an effort to leverage their complementary nature towards more robust and accurate registration. Depending on how one combines the two types of information, three subclasses can be distinguished. In the first case, geometric information is used to initialize the alignment, while intensity-based volumetric registration refines the results [35, 64]. In the second case, geometric information can be used to provide additional constraints that are taken into account during iconic registration [27, 29]. In the third case, iconic and geometric information are integrated in a single objective function that allows for the simultaneous solution of both problems [11, 76, 25].

Once the transformation model and a suitable similarity criterion have been defined, an optimization method is used in order to infer the optimal set of parameters by maximizing the alignment of the two images. Solving for the optimal parameters is particularly challenging in the case of image registration. The reason behind this lies in the fact that image registration is, in general, an ill-posed problem and the associated objective functions are typically non-linear and non-convex. The optimization methods that are typically used in image registration fall either under the umbrella of either continuous or discrete methods.

Typically, continuous optimization methods are constrained to problems where the variables take real values and the objective function is differentiable. This type of problems are common in image registration. As a consequence, these methods (typically gradient descent approaches) have been widely used in image registration [69, 8] because of the fact that they are rather intuitive and easy to implement.

Moreover, they can handle a wide class of objective functions allowing for complex modeling assumptions regarding the transformation model. Nonetheless, they are often sensitive to the initial conditions, while being non-modular with respect to the similarity criterion and the transformation model. What is more, they are often computationally inefficient [24].

On the other hand, discrete optimization methods tackle problems where the variables take discrete values. Discrete optimization methods based on the Markov Random Field theory have been recently investigated in the context of image registration [24, 25]. Discrete optimization methods are constrained by limited precision due to the necessary quantization of the solution space. Moreover, they can not efficiently model complex variable interactions due to increased difficulty in inference. However, recent advances in higher-order inference methods have allowed the modeling of more sophisticated regularization priors [42]. More importantly, discrete optimization methods are versatile and can handle a wide range of similarity metrics (including non-differentiable ones). What is more, they are more robust to the initial conditions due to the global search they perform, while often converging faster than continuous methods.

In this chapter, we review the application of Markov Random Fields (MRFs) in deformable image registration. We explain in detail how one can map image registration from the continuous domain to discrete graph structures. We first present graph-based deformable registration in the case of iconic registration and show how one can encode intensity-based and statistical approaches. We then present discrete attribute-based registration methods and complete the presentation by describing MRF models for geometric and hybrid registration. Throughout this chapter, we discuss the underlying assumptions as well as implementation details. Experimental results that demonstrate the value of graph-based registration are given at the end of every section.

2 Graph-based Iconic Deformable Registration

In this chapter, we focus on pairwise deformable registration. The two images are usually termed as source (or fixed) and target (or moving) images, respectively. The source image is denoted by $S : \Omega_S \subset \mathbb{R}^d \mapsto \mathbb{R}$, while the target image by $T : \Omega_T \subset \mathbb{R}^d \mapsto \mathbb{R}$, $d = \{2, 3\}$. Ω_S and Ω_T denote the image domain for the source and target image, respectively. The source image undergoes a transformation $\mathscr{T} : \Omega_S \mapsto \Omega_T$.

Image registration aims to estimate the transformation $\mathscr{T}$ such that the two images get aligned. This is typically achieved by means of an energy minimization problem:

$$\arg\min_{\theta} \mathscr{M}(T, S \circ \mathscr{T}(\theta)) + \mathscr{R}(\mathscr{T}(\theta)). \tag{1}$$

Thus, the objective function comprises two terms. The first term, $\mathscr{M}$, quantifies the level of alignment between a target image T and a source image S under the influence of the transformation $\mathscr{T}$ parametrized by θ. The second term, $\mathscr{R}$, regularizes

the transformation and accounts for the ill-posedness of the problem. In general, the transformation at every position $\mathbf{x} \in \Omega$ (Ω depicting the image domain) is given as $\mathscr{T}(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x})$ where $\mathbf{u}$ is the deformation field.

The previous minimization problem can be solved by adopting either continuous or discrete optimization methods. In this chapter, we focus on the application of discrete methods that exploit Markov Random Field theory.

2.1 Markov Random Fields

In discrete optimization settings, the variables take discrete values and the optimization is formulated as a discrete labeling problem where one searches to assign a label to each variable such that the objective function is minimized. Such problems can be elegantly expressed in the language of discrete Markov Random Field theory.

An MRF is a probabilistic model that can be represented by an undirected graph $\mathscr{G} = (\mathscr{V}, \mathscr{E})$. The set of vertices $\mathscr{V}$ encodes the random variables, which take values from a discrete set $\mathscr{L}$. The interactions between the variables are encoded by the set of edges $\mathscr{E}$. The goal is to estimate the optimal label assignment by minimizing an energy of the form:

$$E_{MRF} = \sum_{p \in \mathscr{V}} \mathscr{U}_p(l_p) + \sum_{pq \in \mathscr{E}} \mathscr{P}_{pq}(l_p, l_q). \tag{2}$$

The MRF energy also comprises two terms. The first term is the sum of all unary potentials $\mathscr{U}_p$ of the nodes $p \in \mathscr{V}$. This term typically corresponds to the data term since the unary terms are usually used to encode data likelihoods. The second term comprises the pairwise potentials $\mathscr{P}_{pq}$ modeled by the edges connecting nodes p and q. The pairwise potentials usually act as regularizers penalizing disagreements in the label assignment of tightly related variables.

Many algorithms have been proposed in order to perform inference in the case of discrete MRFs. In the works that are presented in this chapter, the fast-PD[2] algorithm [39, 40] has been used to estimate the optimal labeling. The main motivation behind this choice is its great computational efficiency. Moreover, the fast-PD algorithm is appropriate since it can handle a wide-class of MRF models allowing us to use different smoothness penalty functions and has good optimality guarantees.

In the continuation of this section, we detail how deformable registration is formulated in terms of Markov Random Fields. First, however, the discrete formulation requires a decomposition of the continuous problem into discrete entities. This is described below.

---

[2] Fast-PD is available at http://cvc-komodakis.centrale-ponts.fr/.

2.2 Decomposition into Discrete Deformation Elements

Without loss of generality, let us consider a grid-based deformation model that combines low degrees of freedom with smooth local deformations. Let us consider a set of k control points distributed along the image domain using a uniform grid pattern. Furthermore, let k be much smaller than the number of image points. One can then deform the embedded image by manipulating the grid of control points. The dense displacement field is defined as a linear combination of the control point displacements $\mathbf{D} = \{\mathbf{d}_1, ..., \mathbf{d}_k\}$, with $\mathbf{d}_i \in \mathbb{R}^d$, as:

$$\mathbf{u}(\mathbf{x}) = \sum_{i=1}^{k} \omega_i(\mathbf{x})\mathbf{d}_i, \tag{3}$$

and the transformation $\mathscr{T}$ becomes:

$$\mathscr{T}(\mathbf{x}) = \mathbf{x} + \sum_{i=1}^{k} \omega_i(\mathbf{x})\mathbf{d}_i. \tag{4}$$

ω_i corresponds to an interpolation or weighting function which determines the influence of a control point i to the image point x – the closer the image point the higher the influence of the control point. The actual displacement of an image point is then computed via a weighted sum of control point displacements. A dense deformation of the image can thus be achieved by manipulating these few control points.

The free-form deformation is a typical choice for such a representation [71]. This model employs a weighting scheme that is based on cubic B-splines and has found many applications in medical image registration [69] due to its efficiency and the local support of the control points. We also employ this model. Nonetheless, let us note that the discrete deformable registration framework is modular with respect to the interpolation scheme and one may use this preferred strategy.

The parametrization of the deformation field leads naturally to the definition of a set of discrete deformation elements. Instead of seeking a displacement vector for every single image point, now, only the displacement vectors for the control points need to be sought. If we take them into consideration, the matching term (see Eq. 1) can be rewritten as:

$$\mathscr{M}(S \circ \mathscr{T}, T) = \frac{1}{k} \sum_{i=1}^{k} \int_{\Omega_S} \hat{\omega}_i(\mathbf{x}) \rho(S \circ \mathscr{T}(\mathbf{x}), T(\mathbf{x})) d\mathbf{x}, \tag{5}$$

where $\hat{\omega}_i$ are weighting functions similar to the ones in Eq. 4 and ρ denotes a similarity criterion.

Here, the weightings determine the influence or contribution of an image point x onto the (local) matching term of individual control points. Only image points in the vicinity of a control point are considered for the evaluation of the intensity-based similarity measure with respect to the displacement of this particular control point. This is in line with the local support that a control point has on the deformation.

The previous is valid when point-wise similarity criteria are considered. When a criterion based on statistics or information theory is used, a different definition of $\hat{\omega}_i$ is adopted,

$$\hat{\omega}_i(\mathbf{x}) = \begin{cases} 1, & \text{if } \omega_i(\mathbf{x}) \geq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Thus, in both cases the criterion is evaluated on a patch. The only difference is that the patch is weighted in the first case. These local evaluations enhance the robustness of the algorithm to local intensity changes. Moreover, they allow for computationally efficient schemes.

The regularization term of the deformable registration energy (Eq. 1) can also be expressed on the basis of the set of control points as:

$$\mathscr{R} = \frac{1}{k} \sum_{i=1}^{k} \int_{\Omega_S} \hat{\omega}_i(\mathbf{x})\, \psi(\mathscr{T}(\mathbf{x})) d\mathbf{x}, \tag{7}$$

where ψ is a function that promotes desirable properties of the dense deformation field such as the smoothness and topology preservation.

2.3 Markov Random Field Registration Energy

Having identified the discrete deformation elements of our problem, we need to map them to MRF entities, *i.e.*, the graph vertices, the edges, the set of labels, and the potential functions.

Let $\mathscr{G}_{\text{ico}}$ denote the graph that represents our problem. In this case, the random variables of interest are the control point displacement updates. Thus, the set of vertices $\mathscr{V}_{\text{ico}}$ is used to encode them, *i.e.*, $|\mathscr{V}_{\text{ico}}| = |\Delta\mathbf{D}| = k$. Moreover, assigning a label $l_p \in \mathscr{L}_{\text{ico}}$ to a node $p \in \mathscr{V}_{\text{ico}}$ is equivalent to displacing the corresponding control point p by an update $\Delta\mathbf{d}_p$, or $l_p \equiv \Delta\mathbf{d}_p$. In other words, the label set for this set of variable is a quantized version of the displacement space ($L_{\text{ico}} \subset \mathbb{R}^d$). The edge system $\mathscr{E}_{\text{ico}}$ is constructed by following either a 6-connected neighborhood system in the 3*D* case, or a 4-connected system in the 2*D* case. The edge system follows the grid structure of the transformation model.

According to Eq. 5 we define the unary potentials as:

$$\mathscr{U}_{\text{ico},p}(l_p) = \int_{\Omega_S} \hat{\omega}_p(\mathbf{x})\, \rho(S \circ \mathscr{T}_{\text{ico},l_p}(\mathbf{x}), T(\mathbf{x}))\, d\mathbf{x}, \tag{8}$$

where $\mathscr{T}_{\text{ico},l_p}$ denotes the transformation where a control point p has been updated by l_p. Region-based and statistical measures are again encoded in a similar way based on a local evaluation of the similarity measure.

Conditional independence is assumed between the random variables. As a consequence, the unary potential that constitutes the matching term can only be an approximation to the real matching energy. That is because the image deformation,

and thus the local similarity measure, depends on more than one control point since their influence areas do overlap. Still, the above approximation yields very accurate registration as deomonstrated by the experimental validation results that are reported in latter sections (Sec. 2.4, Sec. 3.3 and Sec. 4.3). Furthermore, it allows an extremely efficient approximation scheme which can be easily adapted for parallel architectures yielding extremely fast cost evaluations.

Actually, the previous approximation results in a weighted block matching strategy encoded on the unary potentials. The smoothness of the transformation derives from the explicit regularization constraints encoded by the pairwise potentials and the implicit smoothness stemming from the interpolation strategy.

The evaluation of the unary potentials for a label $l \in \mathscr{L}_{\mathrm{ico}}$ corresponding to an update $\Delta \mathbf{d}$ can be efficiently performed as follows. First, a global translation according to the update $\Delta \mathbf{d}$ is applied to the whole image, and then the unary potentials for this label and for all control points are calculated simultaneously. This results in an one pass through the image to calculate the cost and distribute the local energies to the control points. The constrained transformation in the unary potentials is then simply defined as $\mathscr{T}_{\mathrm{ico},l_p}(\mathbf{x}) = \mathscr{T}_{\mathrm{ico}}(\mathbf{x}) + l_p$, where $\mathscr{T}_{\mathrm{ico}}(\mathbf{x})$ is the current or initial estimate of the transformation.

The regularization term defined in Eq. 7 could be defined as well in the above manner. However, this is not very efficient since the penalties need to be computed on the dense field for every variable and every label. If we consider an elastic-like regularization, we can employ a very efficient discrete approximation of this term based on pairwise potentials as:

$$\mathscr{P}_{\mathrm{ico,elastic},pq}(l_p,l_q) = \frac{\|(\mathbf{d}_p + \Delta\mathbf{d}_p) - (\mathbf{d}_q + \Delta\mathbf{d}_q)\|}{\|\mathbf{p} - \mathbf{q}\|}. \tag{9}$$

The pairwise potentials penalize deviations of displacements of neighboring control points $(p,q) \in \mathscr{E}_{\mathrm{ico}}$ which is an approximation to penalizing the first derivatives of the transformation. Recall that $l_p \equiv \Delta\mathbf{d}_p$. Note, we can also remove the current displacements $\mathbf{d}_p$ and $\mathbf{d}_q$ from the above definition yielding a term that only penalizes the updates on the deformation. This would change the behavior of the energy from an elastic-like to a fluid-like regularization.

Let us detail how the label set $\mathscr{L}_{ico}$ is constructed since that entails an important accuracy-efficiency trade-off. The smaller the set of labels, the more efficient is the inference. However, few labels result in a decrease of the accuracy of the registration. This is due to the fact that the registration accuracy is bounded by the range of deformations covered in the set of labels. As a consequence, it is reasonable to assume that the registration result is sub-optimal. In order to strike a satisfactory balance between accuracy and efficiency, we opt for an *iterative labeling* strategy combined with a *search space refinement* one. At each iteration, the optimal labeling is computed yielding an update on the transformation, *i.e.* $l_p \equiv \Delta\mathbf{d}_p$. This update is applied to the current estimate, and the subsequent iteration continues the registration based on the updated transformation and a refined label set. Thus, the error induced by the approximation stays small and incorrect matches can be corrected

in the next iteration. Furthermore, the overall domain of possible deformations is rather bounded by the number of iterations and not by the set of finite labels.

The iterative labeling allows us to keep the label set quite small. The refinement strategy on the search space is rather intuitive. In the beginning we aim to recover large deformations and as we iterate, finer deformations will be added refining the solution. In each iteration, a sparse sampling with a fixed number of samples s is employed. The total number of labels in each iteration is then $|\mathscr{L}_{\mathrm{ico}}| = g \cdot s + 1$ including the zero-displacement and g is the number of sampling directions. We uniformly sample displacements along certain directions up to a maximum displacement magnitude $\mathbf{d}_{\mathrm{max}}$. Initially, the maximum displacement corresponds to our estimation of the larger deformation to be recovered. In the subsequent iterations, it is decreased by a user-specified factor $0 < f < 1$ limiting and refining the search space.

The number and orientation of the sampling directions g depend on the dimensionality of the registration. One possibility is to sample just along the main coordinate axes, *i.e.* in positive and negative direction of the x-, y-, and z-axis (in case of 3*D*). Additionally, we can add samples for instance along diagonal axes. In 2*D* we commonly prefer a star-shape sampling, which turns out to be a good compromise between the number of samples and the sampling density. In our experiments we found that also very sparse samplings (*e.g.*, just along the main axes) gives very accurate registration results but might increase the total number of iterations that are needed until convergence. However, a single iteration is much faster to compute when the label set is small. In all our experiments we find that small label sets provide an excellent performance in terms of computational speed and registration accuracy.

The explicit control that one has over the creation of the label set $\mathscr{L}$ enables us to impose desirable properties on the obtained solution without further modifying the discrete registration model. Two interesting properties that can be easily enforced by adapting appropriately the discrete solution space are *diffeomorphisms* and *symmetry*. Both properties are of particular interest in medical imaging and have been the focus of the work of many researchers.

Diffeomorphic transformations preserve topology and both them and their inverse are differentiable. These transformations are of interest in the field of computational neuroanatomy. Moreover, the resulting deformation fields are, in general, more physically plausible since foldings, which would disrupt topology, are avoided. As a consequence, many diffeomorphic registration algorithms have been proposed [8, 68, 3, 5, 85].

In this discrete setting, it is straightforward to guarantee a diffeomorphic result through the creation of the label set. By bounding the maximum sampled displacement by 0.4 times the deformation grid spacing, the resulting deformation is guaranteed to be diffeomorphic [68].

The majority of image registration algorithms are asymmetric. As a consequence, when interchanging the order of input images, the registration algorithm does not estimate the inverse transformation. This asymmetry introduces undesirable bias upon any statistical analysis that follows registration because the registration result

depends on the choice of the target domain. Symmetric algorithms have been proposed in order to tackle this shortcoming [13, 5, 56, 84, 79].

Symmetry can also be introduced in graph-based deformable registration in a straightforward manner [77]. This is achieved by estimating two transformations, $\mathscr{T}^f$ and $\mathscr{T}^b$, that deform both the source and the target images towards a common domain that is constrained to be equidistant from the two image domains. In order for this to be true, the transformations, or equivalently the two update deformation fields, should sum up to zero. If one assumes a transformation model that consists of two isomoprhic deformation grids, this constraint translates to ensuring that the displacement updates of corresponding control points in the two grips sum to zero and can be simply mapped to discrete elements.

The satisfaction of the previous constraint can be easily guaranteed in a discrete setting by appropriately constructing the label set. More specifically, by letting the labels index pairs of displacement updates (one for each deformation field) that sum to zero, *i.e.* $l_p \equiv \{\Delta\mathbf{d}_p^f, -\Delta\mathbf{d}_p^b\}$. The extension of the unary terms is also straightforward, while the pairwise potentials and the graph construction are the same.

2.4 Experimental Validation

In this section, we present experimental results for the graph-based symmetric registration in 3D brain registration. The data set consists of 18 T1-weighted brain volumes that have been positionally normalized into the Talairach orientation (rotation only). The MR brain data set along with manual segmentations was provided by the Center for Morphometric Analysis at Massachusetts General Hospital and are available online[3]. The data set was rescaled and resampled so that all images have a size equal to $256 \times 256 \times 128$ and a physical resolution of approximately $0.9375 \times 0.9375 \times 1.5000$mm.

This set of experiments is based on intensity-based similarity metrics (for results using attribute-based similarity metrics, we refer the reader to the next section of this chapter). The results are compared with a symmetric registration method based on continuous optimization [5] that is considered to be the state of the art in continuous deformable registration [38]. Both methods use Normalized Cross Correlation as the similarity criterion.

A multiresolution scheme was used in order to harness the computational burden. A three-level image pyramid was considered while a deformation grid of four different resolutions was employed. The two finest grid resolutions operated on the finest image resolution. The two coarsest operated on the respective coarse image representations. The initial grid spacing was set to 40mm resulting in a deformation grid of size $7 \times 7 \times 6$. The size of the gird was doubled at each finer resolution. A number of 90 labels, 30 along each principal axis, were used. The maximum

---

[3] http://www.cma.mgh.harvard.edu/ibsr/data.html

displacement indexed by a label was bounded to 0.4 times the grid spacing. The pairwise potentials were weighted by a factor of 0.1.
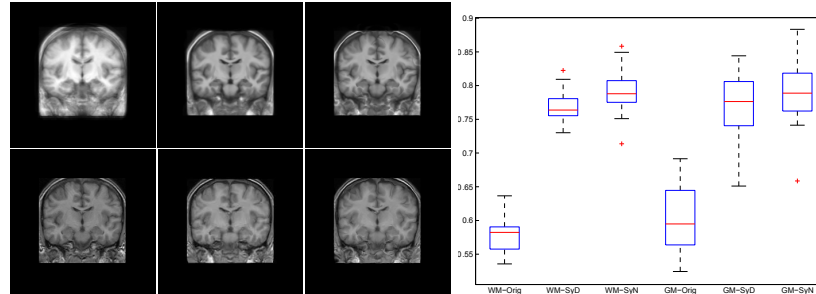
Fig. 2: a) In the first row, from left to right, the mean intensity image is depicted for the data set, after the graph-based symmetric registration method and after [5]. In the second row, from left to right, the target image is shown as well as a typical deformed image for the graph-based symmetric registration method and [5]. For all cases, the central slice is depicted. b) Boxplots for the DICE criterion initially, with the graph-based symmetric registration method and with [5]. On the left, the results for the WM. On the right, the results for the GM. The figure is reprinted from [77].

The qualitative results (sharp mean and deformed image) suggest that both methods successfully registered the images to the template domain. The results of [5] seem to have produced more aggressive deformation fields that have resulted to some unrealistic deformations in the top of the brain and can also be observed in the borders between white matter (WM) and gray matter (GM). This aggressive registration has also resulted in slightly more increased DICE coefficients for WM and GM. However, the results reported for the graph-based registration method were obtained in 10 min. On the contrary, 1 hour was necessary to register the images with [7] approximately. This important difference in the computational efficiency between the two methods can outweigh the slight difference in the quality of the solution in practice.

3 Graph-based Attribute-Based Deformable Registration

In the previous section, we studied the application of intensity-based deformable registration methods that involve voxel-wise and statistical similarity criteria. While these criteria are easy to compute and widely used, they suffer from certain shortcomings. First, they often have difficulties to reflect the underlying anatomy because pixels belonging to the same anatomical structure are often assigned different intensity values due to variabilities arising from scanners, imaging protocols, noise, partial volume effects, contrast differences, and image inhomogeneities. Moreover, single intensities are not informative enough to uniquely characterize image elements, and thus reliably guide image registration. For instance, hundreds of thousands of

gray matter voxels in a brain image share similar intensities; but they belong to different anatomical structures. As a consequence, matching ambiguities arise in the matching between two images.

In order to reduce matching, one needs to characterize each voxel more distinctively. This may be achieved by creating richer high-dimensional descriptors of image elements that capture texture or geometric regional attributes. Therefore, attribute-based similarity criteria have been increasingly used in image registration. Typical examples include the use of geometric-moment-invariant (GMI) attributes coupled with tissue membership attributes and boundary/edge attributes [74], neighborhood intensity profile attributes [28], local frequency attributes [49, 34], local intensity histogram attributes [73, 88], geodesic intensity histogram attributes [47, 44] and scale-invariant attributes [81].

3.1 Gabor Attributes

The versatility of graph-based deformable registration models allows the seamless integration of any of the previous attribute-based similarity criteria. Nonetheless, the previous approaches involve features that are either application-specific and fail to generalize to other applications, or require sophisticated pre-processing steps (*e.g.*segmentation). As a consequence, it is important to appropriately choose the attribute-based description so that, when coupled with the highly modular discrete approaches, a general-purpose registration method is possible.

Gabor-attributes, which involve image convolution with Gaussian filters at multiple scales and orientations, present an interesting choice for general-purpose deformable registration. The reason is threefold. First, all anatomical images have texture information, at some some scale and orientation, reflecting the underlying geometric and anatomical characteristics. As a results, Gabor features that are able to capture this information can be, and have been, applied in a variety of studies. Second, Gabor filters are able to capture edge information that is relatively encoded by various image modalities, thus making them suitable for both mono- and multi-modal registration tasks. Third, their multi-scale and multi-orientation nature render image elements more distinctive and better identifiable for establishing correspondences. For example, the scale information helps differentiate voxels that are the center of a small and a bigger plate, respectively. The orientation information can help distinguish, for example, a voxel on a left-facing edge from a voxel on a right-facing edge. Moreover, it is also possible to automatically select a subset of Gabor attributes such that the information redundancy is reduced and the distinctiveness of the descriptor is increased.

The effect of characterizing voxels using Gabor attributes (with and without optimal Gabor attribute subset selection) is presented in Fig. 3 (reprinted from [62]). These effects are contrasted to the effect of using only intensities and using Gray-Level-Cooccurance-Matrix (GLCM) texture attributes through the use of similarity maps between voxels from the source image (labeled under crosses) and all voxels
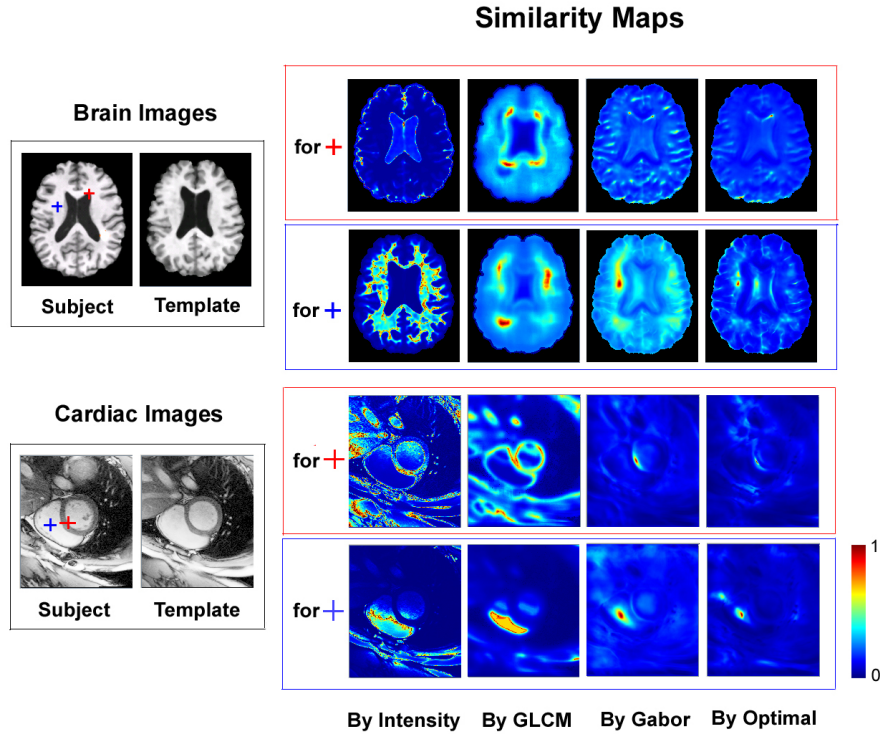
Similarity Maps

Fig. 3: The similarity maps between special/ordinary voxels (labeled by red/blue crosses) in the source (a.k.a, subject) images and all voxels in the target (a.k.a, template) images. As correspondences were sought based on voxel similarities (subject to spatial smoothness constraints), (optimal-)Gabor-attribute-based similarity maps returned a much smaller search range for correspondences. This figure is reprinted from [62].

.

in the target image. The similarity between two voxels, $\mathbf{x}$ in the source image and $\mathbf{y}$ in the target image, was defined as $\text{sim}(\mathbf{x}, \mathbf{y}) = \frac{1}{1+\|A(\mathbf{x})-A(\mathbf{y})\|^2}$, with $A(\cdot)$ being the attribute vector at each voxel. This similarity ranged from 0 (when the attributes between two voxels differed infinitely) and 1 (when the attributes between two voxels were identical). This figure shows that, as one replaced the intensity-based similarity to (optimal)-attribute-based similarities, even very ordinal voxels under the blue crosses were distinctively characterized or better localized in the space, therefore we only needed to search for their corresponding voxels within a much smaller range in the target image, largely removing matching ambiguities.

Let us detail in the next section how one can introduce Gabor-based attributes in the case of graph-based deformable registration [62]. More specifically, let us detail how the Markov Random Field energy (see Eq. 2) changed in this regard.

3.2 Markov Random Field Registration Energy

The ease with which one can adopt attribute-based similarity criteria in the case of graph-based formulations for deformable registration is evidence of their high versatility and modularity. The key elements of the graphical model (*i.e.*, graph construction, pairwise potentials, inference) need not change. One only needs to slightly change the definition of the unary potentials.

The unary potentials need only be modified in two regards: i) to evaluate the similarity criterion ρ over the attribute vectors $A(\cdot)$; and ii) to optionally, as suggested by [62], take into account a spatially-varying weighting parameter $ms(\mathbf{x})$, namely "mutual-saliency", which automatically quantified the confidence of each voxel $\mathbf{x}$ to establish reliable correspondences across images. Therefore, the modified unary potentials are defined as:

$$\mathscr{U}_{\text{ico},p}(l_p) = \int_{\Omega_S} ms(\mathbf{x}) \cdot \hat{\omega}_p(\mathbf{x}) \cdot \rho\left(A_S \circ \mathscr{T}_{\text{ico},l_p}(\mathbf{x}), A_T(\mathbf{x})\right) d\mathbf{x}. \tag{10}$$

3.3 Experimental Validation

In this section we present results obtained with an attributed-based discrete deformable registration termed DRAMMS (Deformable Registration via Attribute Matching and Mutual-Saliency) [62]. The presented results demonstrate the advantageous computational efficiency of graph-based registration method in comparison to the traditional gradient descent optimization strategy. Moreover, the results demonstrate the generality, accuracy and robustness of coupling attributed-based similarity criteria with graph-based formulations.

As far as the computational efficiency is concerned, Fig. 4 summarizes the computational time that is required to register brain, prostate, and cardiac images using a gradient descent optimization strategy and a discrete optimization strategy [39, 40], respectively. The discrete approach requires significantly reduced computational time.

In the second part of this section, we report results for DRAMMS in two different cases: i) skull-stripped brain MR images; and ii) brain MR images from the large-scale, multi-institutional, with-skull ADNI database.

In the first case, DRAMMS was compared to 11 other popular and publicly-available registration tools, all used with the optimized parameters as reported in [38] whenever applicable. In the public NIREP dataset containing T1-weighted MR images ($256 \times 300 \times 256$ voxels and $1.0 \times 1.0 \times 1.0 mm^3$/voxel) of 16 healthy subjects, each registration method was applied to all the possible 210 pair-wise registrations, leading to 2,520 registrations in total. DRAMMS had been shown to yield the highest average Jaccard overlap among 32 regions-of-interest (ROIs) annotated by human expert, indicating the high accuracy (Fig. 5). Such a trend had also been

	FFD + GradDes		FFD + DisOpt	
	AM w/ MS	AM w/o MS	AM w/ MS	AM w/o MS
Brain images (256 × 256 × 171)	534.67	268.82	115.78	36.49
Prostate images (256 × 256 × 34)	245.46	104.77	44.72	15.41
Cardiac images (256 × 256 × 20)	181.54	88.77	39.07	13.61

Fig. 4: The computational times (in minutes) when combing the MRF registration formulation with the discrete optimization strategy versus with the traditional gradient descent optimization strategy. The discrete optimization strategy on the MRF registration formulation helped significantly reduce the computational time. AM refers to attribute matching; MS refers to mutual-saliency weighting, which is a second component in DRAMMS but was not described in full detail in this section; basically it is a automatically computed weighting for adaptively utilizing voxels based on how much confidence we have for those voxels to find correspondences across images. FFD is the free form deformation transformation model as used in the MRF registration formulation. And DisOpt and GradDes are the discrete optimization and gradient descent optimization strategies. This figure is reprinted from [62].

observed in several other databases containing skull-stripped brain MR images from healthy subjects [57].
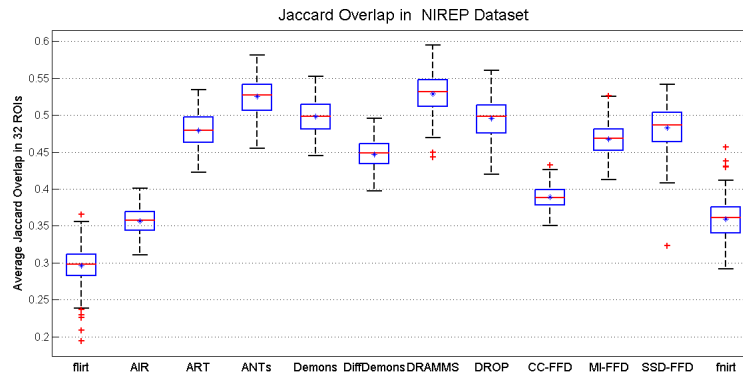

Fig. 5: The average Jaccard overlap among all ROIs in all possible pair-wise registrations within the NIREP database, for different registration tools. Reprinted from [57].
.

In the second case, DRAMMS was validated using brain MR images from the ADNI study. This study presents particular challenges because it contains data obtained at different sites and regions affected by pathologies. In Fig. 6 (re-printed from [57]) one can observe that DRAMMS can align largely variable ventricles,

whereas other registration tools encountered great challenges. This is characteristic of the accuracy and robustness of the attribute-based discrete deformable registration method.
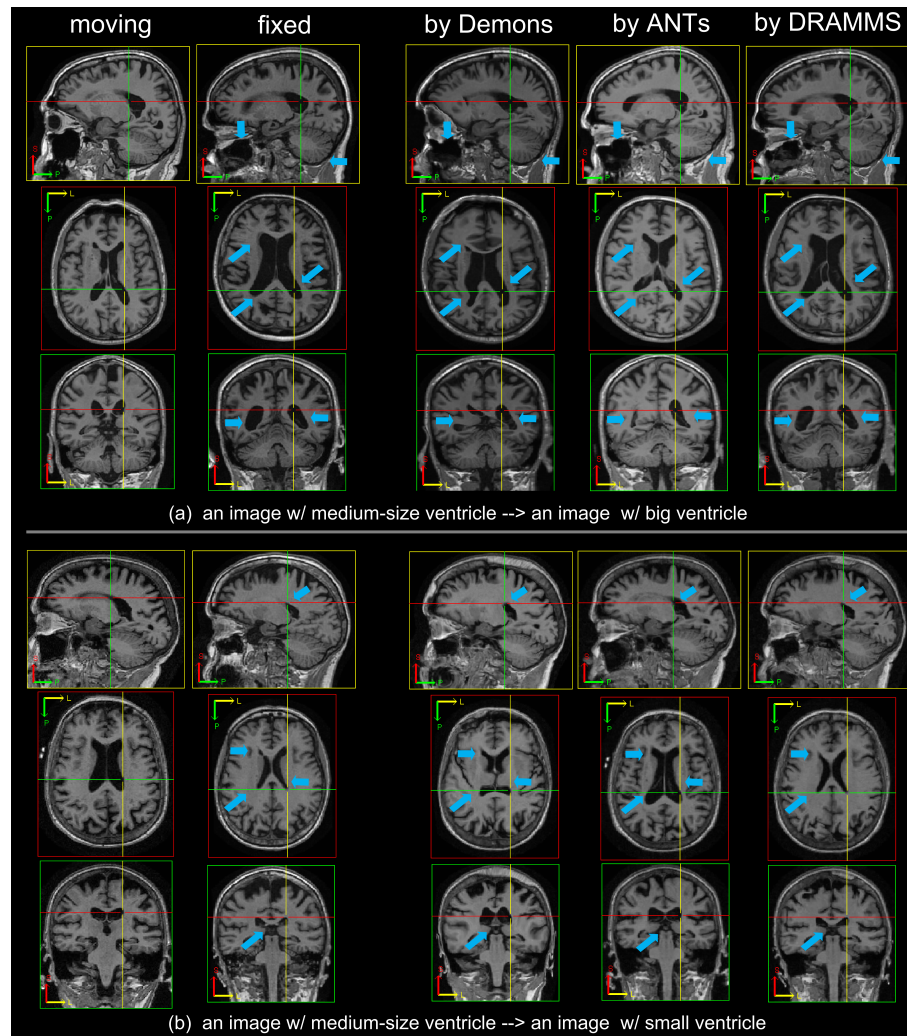
Fig. 6: Example registration results between subjects in the multi-site Alzheimer's Disease Neuroimaging Initiative (ADNI) database, by different registration methods. Blue arrows point out regions where the results from various registration methods differ. Reprinted from [57].

These results emphasize the generality, accuracy and robustness of the attribute-based discrete deformable registration. Because of these characteristics and its pub-

lic availability[4], DRAMMS has found application in numerous translational studies including neuro-degenerative studies [41, 90, 17, 72], neuro-developmental ones [70, 60, 21, 33] as well as oncology studies [6, 59]. These applications underline the versatility of combining attribute-based similarity criteria with graph-based formulations.

4 Graph-based Geometric and Hybrid Deformable Registration

The previous two sections presented MRF-based iconic (*a.k.a.* voxel-wise) registration using intensity- and attribute-based similarities. Typically, iconic approaches evaluate the similarity criteria over the whole image domain and have the potential to better quantify and represent the accuracy of the estimated dense deformation field, albeit at an important computational cost. Nonetheless, iconic approaches do not explicitly take into account salient image points, failing to fully exploit image information. Moreover, the performance of iconic methods, especially methods based on continuous optimization, is greatly influenced by the initial conditions.

On the other hand, geometric methods utilize only a sparse subset of image elements that correspond to salient geometry or anatomy. Exploiting relevant information results in increased robustness. Nonetheless, the quality of the estimated deformation field is high only on the vicinity of the landmarks.

Hybrid registration methods exploit both types of information towards bridging the gap between the two basic classes of registration and enjoying the advantages of both worlds. Iconic and geometric information are integrated in an unified objective function and the solutions of the two problems satisfy each other. In this setting, iconic methods may profit from geometric information in the cases they encounter difficulties arising, for example, from large deformations (*e.g.*, the largely different ventricle size in Alzheimer's Disease population), or from missing correspondences such as the existence of pathologies. At the same time, geometric correspondences can be refined based on the iconic information that is available throughout the image domain.

In this section, we consecutively study the graph-based formulation of geometric and hybrid deformable registration. Similarly to the previous sections, we first study the two problems in their continuous form and show how they can be decomposed in discrete entities. Then, we detail the graph-based formulation and present experimental results.

---

[4] DRAMMS is available at http://www.nitrc.org/projects/dramms/.

4.1 Decomposition into Discrete Deformation Elements

4.1.1 Geometric Registration

A prerequisite for geometric registration is the availability of landmarks that encode salient geometry or anatomy. Landmarks can be annotated by experts, or, to reduce intra-/inter-expert variability, by (semi-)automated methods. The latter is an open problem and an active topic of research.

Automatic approaches to detect landmarks include, but are not limited to, edge detection [66, 31], contour delineation [45], anatomical structure segmentation [9, 12], scale space analysis [63, 46, 36], and feature transformation (*e.g.*, SIFT [50, 89, 37], SURF [7, 92]). In [91, 61], for example, the authors used Laplacian operations to search for blob-like structures, and used the centers of the blobs as landmarks. In [58], the authors used regional centers or edges at various scales and orientations as landmarks, which were of strong response to Gabor filters. While a detailed survey of landmark detection is outside the scope of this section, we want to emphasize that the described graph-based formulation can seamlessly integrate landmark information coming from any algorithm or expert.

Given two sets of landmarks K ($\kappa \in K$) and Λ ($\lambda \in \Lambda$), one aims to estimate the transformation $\mathscr{T}_{\text{geo}}$ that will bring them into correspondence by minimizing an objective function of the form of Eq. 1. More specifically, the goal is to bring every landmark belonging the set K as close as possible to the landmark in the set Λ that is most similar to it. In other words, the matching term is expressed as:

$$\mathscr{M}_{\text{geo}}(K \circ \mathscr{T}_{\text{geo}}, \Lambda) = \frac{1}{n} \sum_{i=1}^{n} \delta(\mathscr{T}_{\text{geo}}(\kappa_i), \widetilde{\lambda}_i) \qquad (11)$$

where δ measures the Euclidean distance between two landmark positions, and

$$\widetilde{\lambda}_i = \arg\min_{\lambda_j} \rho(\mathscr{T}_{\text{geo}}(\kappa_i).\lambda_j). \qquad (12)$$

Note that the Euclidean position of the landmarks λ and κ is denoted in bold.

As far as the regularization term $\mathscr{R}_{\text{geo}}$ is concerned, it aims to preserve the smoothness of the transformation. More specifically, it aims to locally preserve the geometric distance between pairs of landmarks:

$$\mathscr{R}_{\text{geo}}(\mathscr{T}_{\text{geo}}) = \frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j=1, j \neq i}^{n} \|(\mathscr{T}_{\text{geo}}(\kappa_i) - \mathscr{T}_{\text{geo}}(\kappa_j)) - (\kappa_i - \kappa_j)\|. \qquad (13)$$

This implies the assumption that a linear registration step that has accounted for differences in scales has been applied prior to the deformable registration.

An equivalent way of formulating the geometric registration problem consists of first pairing landmarks $\kappa \in K$ with the most similar in appearance landmarks $\lambda \in \Lambda$ and then pruning the available pairs by keeping only those that are geometrically

consistent as quantified by the regularization term (Eq. 13). Let us note that, in both cases, the problem is inherently discrete.

4.1.2 Hybrid Registration

As discussed in the introduction, there are various ways of integrating geometric and iconic information. The most interesting, and potentially more accurate, is the one that allows both problems to be solved at the same time through the optimization of a universal energy that enforces the separate solutions to agree. This is possible by combining the previous energy terms for the iconic and geometric problem along with a hybrid term that acts upon the separate solutions:

$$\mathscr{H}(\mathscr{T}_{\text{ico}}, \mathscr{T}_{\text{geo}}) = \frac{1}{n} \sum_{i=1}^{n} \| \mathscr{T}_{\text{ico}}(\kappa_i) - \mathscr{T}_{\text{geo}}(\kappa_i) \|. \tag{14}$$

Note that we only need to enforce the agreement of the two solutions in the landmark positions. If we now also consider a connection between control point displacements $\mathbf{D}$ and landmark displacements, the previous relation can be rewritten as:

$$\mathscr{H}(\mathscr{T}_{\text{ico}}, \mathscr{T}_{\text{geo}}) = \frac{1}{n} \sum_{i=1}^{n} \| \kappa_i + \mathbf{u}_{\text{geo}}(\kappa_i) - \kappa_i - \sum_{j=1}^{k} \omega_j(\kappa_i) \mathbf{d}_j \|, \tag{15}$$

where $\mathbf{u}_{\text{geo}}(\kappa_i) = \widetilde{\lambda}_i - \kappa_i$, *i.e.* the displacement for the correspondence of the two landmarks κ_i and $\widetilde{\lambda}_i$. As a principle, we would like this displacement to be ideally equal to the one that is given as a linear combination of the displacements of the control points at the position of a landmark. However, we can relax the previous requirement in order to increase the computational efficiency of the method. If we apply the triangular inequality and exploit the fact that the coefficients ω_j are positive, the coupling constraint is redefined as:

$$\mathscr{H}(T_{\text{ico}}, T_{\text{geo}}) \leq \frac{1}{n} \sum_{i=1}^{n} \sum_{j=1}^{k} \omega_j(\kappa_i) \| \mathbf{u}_{\text{geo}}(\kappa_i) - \mathbf{d}_j \|. \tag{16}$$

The previous constraint comprises only pairwise interactions between discrete elements.

4.2 Markov Random Field Registration Energy

Having identified the discrete elements for both geometric and hybrid registration, let us map them to MRF entities.

4.2.1 Geometric Registration

Let us now introduce a second graph $\mathscr{G}_{\text{geo}} = \left(\mathscr{V}_{\text{geo}}, \mathscr{E}_{\text{geo}}\right)$ for the geometric entities K, Λ. We remind that they are two sets of landmarks having different cardinalities and we seek the transformation which will bring each landmark into correspondence with the best candidate. Equivalently, we may state that we are trying to solve for the correspondence of each landmark, which naturally results in a set of sparse displacements.

The second graph consists of a set of vertices $\mathscr{V}_{\text{geo}}$ corresponding to the set of landmarks extracted in the source image, *i.e.* $|\mathscr{V}_{\text{geo}}| = |K|$. A label assignment $l_p \in \mathscr{L}_{\text{geo}} := \Lambda$ (where $p \in \mathscr{V}_{\text{geo}}$) is equivalent to matching the landmark $\kappa_p \in K$ to a candidate point $l_p \equiv \lambda \in \Lambda$. Assigning a label l_p implicitly defines a displacement $\mathbf{u}_{\text{geo},l_p}(\kappa_p) = \lambda - \kappa_p$, since κ_p is mapped on the landmark l_p.

According to Eq. 12, the unary potentials are defined as:

$$\mathscr{U}_{\text{geo},p}(l_p) = \rho(\kappa_p, l_p). \tag{17}$$

The two different though equivalent ways to see the label assignment problem are depicted in the previous equation. Assigning a label l_p can interpreted as applying a transformation $\mathscr{T}_{\text{geo},l_p} = \kappa_p + \mathbf{u}_{\text{geo},l_p}(\kappa_p)$ or stating that the landmark κ_p corresponds to the l_p. Contrary to the iconic case, the set of transformations that can be applied is specified by the candidate landmarks and is sparse in its nature.

There is a number of ways to define the dissimilarity function ρ. One approach would be to consider neighborhood information. That can be easily done by evaluating the criterion over a patch centered around the landmarks,

$$\mathscr{U}_{\text{geo},p}(l_p) = \int_{\Omega_{S,p}} \rho(S \circ \mathscr{T}_{\text{geo},l_p}(\mathbf{x}), T(\mathbf{x})) d\mathbf{x}, \tag{18}$$

where $\Omega_{S,p}$ denotes a patch around the point κ_p. Another approach is to exploit attribute-based descriptors and mutual saliency [58] and define the potential as:

$$\mathscr{U}_{\text{geo},p}(l_p) = \exp\left(-\frac{ms(\kappa_p, l_p) \cdot sim(\kappa_p, l_p)}{2\sigma^2}\right). \tag{19}$$

where σ is a scaling factor, estimated as the standard deviation of the mutual saliency values of all the candidate pairs.

The regularization term defined in Eq. 13 can be encoded by the edge system $\mathscr{E}_{\text{geo}}$ of the graph. In this setting, the regularization term can be expressed as:

$$\mathscr{E}_{\text{geo},pq}(l_p, l_q) = \|(\mathscr{T}_{\text{geo},l_p}(\kappa_p) - \mathscr{T}_{\text{geo},l_q}(\kappa_q)) - (\kappa_p - \kappa_q)\|. \tag{20}$$

The pairwise potential will enforce an isometric constraint. Moreover, by considering the vector differences flipping of the point positions is penalized.

Last, it is interesting to note that the same graph $\mathscr{G}_{\text{geo}}$ is able to encode both ways of formulating the geometric registration problem that were presented in Sec. 4.1.1. This model was presented in [58] and [76].

4.2.2 Hybrid Registration

In this case, the graph-based formulation will consist of the discrete model for the iconic and geometric registration along with a coupling penalty (Eq. 16). Therefore, the graph that represents the problem comprises $\mathscr{G}_{\text{ico}}$ and $\mathscr{G}_{\text{geo}}$ along with a third set of edges $\mathscr{E}_{\text{hyb}}$ containing all possible connections between the iconic random variables and the geometric variables. The pairwise label assignment penalty on these coupling edges is then defined as:

$$\mathscr{P}_{\text{hyb},pq}(l_p, l_q) = \omega_q(\kappa_p) \left\| \mathbf{u}_{\text{geo},l_p}(\kappa_p) - (\mathbf{d}_q + l_q) \right\|, \tag{21}$$

where $p \in \mathscr{V}_{\text{geo}}$ and $q \in \mathscr{V}_{\text{ico}}$, $l_p \in \mathscr{L}_{\text{geo}}$ and $l_q \in \mathscr{L}_{\text{ico}}$, and $(p,q) \in \mathscr{E}_{\text{hyb}}$. Such a pairwise term couples the displacements given by the two registration processes and imposes consistency. To conclude, the coupled registration objective function is represented by an MRF graph $\mathscr{G}_{\text{hyb}} = (\mathscr{V}_{\text{geo}} \cup \mathscr{V}_{\text{ico}}, \mathscr{E}_{\text{geo}} \cup \mathscr{E}_{\text{ico}} \cup \mathscr{E}_{\text{hyb}})$ with its associated unary and pairwise potential functions. This model was presented in [76].

4.3 Experimental Validation

4.3.1 Geometric Registration

Fig. 7 shows a typical landmark pair detected by Gabor response and matched by the MRF formulation. Many such pairs found by the MRF formulation resulted in a deformation that was smoother with the MRF regularization rather than without, as can be seen in Fig. 8.

(a)   (b)
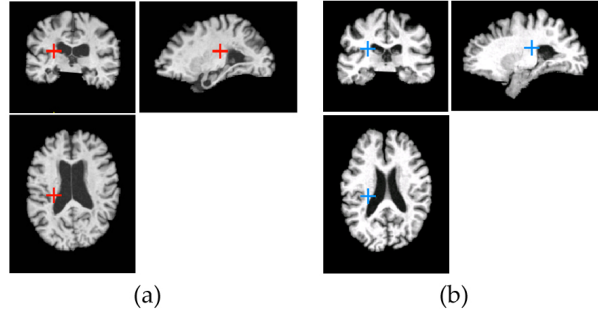
Fig. 7: An example landmark pair (denoted by red and blue crosses) detected based on the Gabor response-based similarity metric and the mutual-saliency measure. (a) Source and (b) target images. This figure is re-printed from [58].
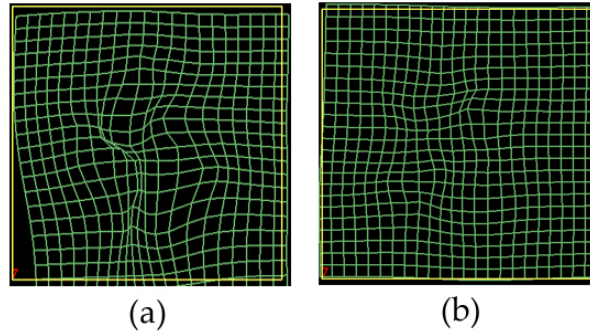
Fig. 8: The Dense deformation fields generated by (a) M1 – no MRF regularization and (b) M2 – with MRF regularization. This figure is re-printed from [58].

4.3.2 Hybrid Registration

In order to validate the coupled geometric registration method in a way that is invariant to landmark extraction, a multi-modal synthetic data set is used. In this setting, the ground truth deformation is known allowing for a quantitative analysis of the registration performance regarding both the dense deformation field accuracy and the quality of the established landmark correspondences.

The goal of this experiment is to demonstrate the added value from considering geometric information on top of standard iconic one. Thus, a comparison of the proposed framework with and without the geometric registration part takes place. Regarding the results, if we look at the registration accuracy in terms of end point error (Table 1), we see that the coupled iconic geometric registration method is able to further improve the results of the iconic one. This is evident, as the end point error has decreased by taking advantage of the geometric information.

As we expect the hybrid approach to be able to cope with large displacements better than the pure iconic one, we repeated the experiments by decreasing the initial control point spacing to 20mm and thus limiting the maximum amount of deformation that can be handled. The results are also reported in Table 1. In this case, we can observe a more significant difference between the performance of the two proposed approaches. Therefore, we should conclude that the additional computational cost demanded by the coupled approach can be compensated by the better quality of the results.

5 Conclusion

This chapter presents a comprehensive overview of graph-based deformable registration. Discrete models for the cases of deformable registration involving pointwise intensity-based similarity criteria, statistical intensity-based criteria, attribute-

#	Iconic (h=60mm)		Hybrid (h=60mm)		Iconic (h=20mm)		Hybrid (h=20mm)	
	mean	std	mean	std	mean	std	mean	std
1	1.33	0.69	**1.25**	0.59	1.38	1.21	**0.98**	0.61
2	1.32	0.75	**1.18**	0.53	2.46	3.21	**1.06**	0.68
3	1.44	0.97	**1.22**	0.56	2.05	2.40	**1.03**	0.67
4	1.40	0.74	**1.16**	0.50	1.40	1.02	**1.08**	0.69
5	1.23	0.60	**1.15**	0.56	1.38	1.01	**1.03**	0.67
6	1.35	0.74	**1.24**	0.62	1.58	1.39	**1.05**	0.71
7	1.16	0.56	**1.09**	0.50	1.45	1.18	**1.05**	0.67
8	1.29	0.68	**1.23**	0.58	1.93	2.61	**1.11**	0.79
9	1.23	0.62	**1.19**	0.53	1.72	1.89	**1.04**	0.71
10	1.54	1.08	**1.19**	0.58	2.60	3.43	**1.05**	0.73
all	1.33	0.11	**1.19**	0.05	1.79	0.45	**1.05**	0.03

Table 1: End point error (in millimeters) for the registration of the Synthetic MR Dataset. The grid spacing is denoted by h. This figure is reprinted from [25].

based ones as well as for geometric and hybrid registration were presented. The increased computational efficiency, accuracy and robustness of graph-based formulations was also demonstrated.

Acknowledgements

References

1. Amit, Y.: A nonlinear variational problem for image matching. SIAM Journal on Scientific Computing **15**(1), 207–224 (1994)
2. Arsigny, V., Pennec, X., Ayache, N.: Polyrigid and polyaffine transformations: A novel geometrical tool to deal with non-rigid deformations – Application to the registration of histological slices. Medical Image Analysis **9**(6), 507–523 (2005)
3. Ashburner, J.: A fast diffeomorphic image registration algorithm. NeuroImage **38**(1), 95–113 (2007)
4. Ashburner, J., Friston, K.J.: Nonlinear spatial normalization using basis functions. Human Brain Mapping **7**(4), 254–266 (1999)
5. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Medical image analysis **12**(1), 26–41 (2008)
6. Baumann, B.C., Teo, B.K., Pohl, K., Ou, Y., Doshi, J., Alonso-Basanta, M., Christodouleas, J., Davatzikos, C., Kao, G., Dorsey, J.: Multiparametric processing of serial mri during radiation therapy of brain tumors:finishing with flair?. International Journal of Radiation Oncology* Biology* Physics **81**(2), S794 (2011)
7. Bay, H., Ess, A., Tuytelaars, T., Van Gool, L.: Speeded-up robust features (surf). Computer vision and image understanding **110**(3), 346–359 (2008)
8. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision **61**(2), 139–157 (2005)

9. Betke, M., Hong, H., Thomas, D., Prince, C., Ko, J.P.: Landmark detection in the chest and registration of lung surfaces with an application to nodule registration. Medical Image Analysis **7**(3), 265–281 (2003)
10. Bookstein, F.L.: Principal warps: Thin-plate splines and the decomposition of deformations. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(6), 567–585 (1989)
11. Cachier, P., Mangin, J.F., Pennec, X., Rivière, D., Papadopoulos-Orfanos, D., Régis, J., Ayache, N.: Multisubject non-rigid registration of brain MRI using intensity and geometric features. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 734–742 (2001)
12. Can, A., Stewart, C.V., Roysam, B., Tanenbaum, H.L.: A feature-based, robust, hierarchical algorithm for registering pairs of images of the curved human retina. Pattern Analysis and Machine Intelligence, IEEE Transactions on **24**(3), 347–364 (2002)
13. Christensen, G.E., Johnson, H.J.: Consistent image registration. IEEE transactions on medical imaging **20**(7), 568–82 (2001)
14. Christensen, G.E., Rabbitt, R.D., Miller, M.I.: Deformable templates using large deformation kinematics. IEEE Transactions on Image Processing **5**(10), 1435–1447 (1996)
15. Chui, H., Rangarajan, A.: A new point matching algorithm for non-rigid registration. Computer Vision and Image Understanding **89**(2-3), 114–141 (2003)
16. Chung, A.C., Wells III, W.M., Norbash, A., Grimson, W.E.L.: Multi-modal image registration by minimizing Kullback-Leibler distance. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 525–532 (2002)
17. Da, X., Toledo, J.B., Zee, J., Wolk, D.A., Xie, S.X., Ou, Y., Shacklett, A., Parmpi, P., Shaw, L., Trojanowski, J.Q., et al.: Integration and relative value of biomarkers for prediction of mci to ad progression: Spatial patterns of brain atrophy, cognitive scores, apoe genotype and csf biomarkers. NeuroImage: Clinical **4**, 164–173 (2014)
18. D'Agostino, E., Maes, F., Vandermeulen, D., Suetens, P.: A viscous fluid model for multi-modal non-rigid image registration using mutual information. Medical Image Analysis **7**(4), 565–575 (2003)
19. Davatzikos, C.: Spatial transformation and registration of brain images using elastically deformable models. Computer Vision and Image Understanding **66**(2), 207–222 (1997)
20. Droske, M., Rumpf, M.: A variational approach to nonrigid morphological image registration. SIAM Journal on Applied Mathematics **64**(2), 668–687 (2004)
21. Erus, G., Battapady, H., Satterthwaite, T.D., Hakonarson, H., Gur, R.E., Davatzikos, C., Gur, R.C.: Imaging patterns of brain development and their relationship to cognition. Cerebral Cortex p. bht425 (2014)
22. Fischer, B., Modersitzki, J.: Fast diffusion registration. AMS Contemporary Mathematics, Inverse Problems, Image Analysis, and Medical Imaging **313**, 117–127 (2002)
23. Glaunès, J., Trouvé, A., Younes, L.: Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub-manifolds matching. In: International Conference on Computer Vision and Pattern Recognition, pp. 712–718 (2004)
24. 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)
25. Glocker, B., Sotiras, A., Komodakis, N., Paragios, N.: Deformable medical image registration: setting the state of the art with discrete methods. Annual Review of Biomedical Engineering **13**, 219–244 (2011)
26. Hajnal, J.V., Hill, D.L., Hawkes, D.J. (eds.): Medical image registration. CRC Press, Boca Raton, FL (2001)
27. Hartkens, T., Hill, D.L.G., Castellano-Smith, A., Hawkes, D.J., Maurer, C.R., Martin, A., Hall, W., Liu, H., Truwit, C.: Using points and surfaces to improve voxel-based non-rigid registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 565–572 (2002)
28. Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, S.M., Schnabel, J.A.: Mind: Modality independent neighbourhood descriptor for multi-modal deformable registration. Medical Image Analysis **16**(7), 1423–1435 (2012)

29. Hellier, P., Barillot, C.: Coupling dense and landmark-based approaches for nonrigid registration. IEEE Transactions on Medical Imaging **22**(2), 217–227 (2003)
30. Holden, M.: A review of geometric transformations for nonrigid body registration. IEEE Transactions on Medical Imaging **27**(1), 111–128 (2008)
31. Hsieh, J.W., Liao, H.Y.M., Fan, K.C., Ko, M.T., Hung, Y.P.: Image registration using a new edge-based approach. Computer Vision and Image Understanding **67**(2), 112–130 (1997)
32. Huang, X., Paragios, N., Metaxas, D.N.: Shape registration in implicit spaces using information theory and free form deformations. IEEE Transactions on Pattern Analysis and Machine Intelligence **28**(8), 1303–1318 (2006)
33. Ingalhalikar, M., Parker, D., Ghanbari, Y., Smith, A., Hua, K., Mori, S., Abel, T., Davatzikos, C., Verma, R.: Connectome and maturation profiles of the developing mouse brain using diffusion tensor imaging. Cerebral Cortex p. bhu068 (2014)
34. Jian, B., Vemuri, B., Marroquin, J.: Robust nonrigid multimodal image registration using local frequency maps. In: Information Processing in Medical Imaging (IPMI), pp. 504–515 (2005)
35. Johnson, H.J., Christensen, G.E.: Consistent landmark and intensity-based image registration. IEEE Transactions on Medical Imaging **21**(5), 450–461 (2002)
36. Kadir, T., Brady, M.: Saliency, scale and image description. International Journal of Computer Vision **45**(2), 83–105 (2001)
37. Ke, Y., Sukthankar, R.: Pca-sift: A more distinctive representation for local image descriptors. In: Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on, vol. 2, pp. II–506. IEEE (2004)
38. Klein, A., Andersson, J., Ardekani, B.A., Ashburner, J., Avants, B., Chiang, M.C., Christensen, G.E., Collins, D.L., Gee, J., Hellier, P., et al.: Evaluation of 14 nonlinear deformation algorithms applied to human brain mri registration. Neuroimage **46**(3), 786–802 (2009)
39. Komodakis, N., Tziritas, G.: Approximate labeling via graph cuts based on linear programming. IEEE transactions on pattern analysis and machine intelligence **29**(8), 1436–53 (2007)
40. Komodakis, N., Tziritas, G., Paragios, N.: Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal-dual strategies. Computer Vision and Image Understanding **112**(1), 14–29 (2008)
41. Koutsouleris, N., Davatzikos, C., Borgwardt, S., Gaser, C., Bottlender, R., Frodl, T., Falkai, P., Riecher-Rössler, A., Möller, H.J., Reiser, M., et al.: Accelerated brain aging in schizophrenia and beyond: a neuroanatomical marker of psychiatric disorders. Schizophrenia bulletin p. sbt142 (2013)
42. Kwon, D., Lee, K., Yun, I., Lee, S.: Nonrigid image registration using dynamic higher-order mrf model. In: European Conference on Computer Vision, pp. 373–386 (2008)
43. Leordeanu, M., Hebert, M.: A spectral technique for correspondence problems using pairwise constraints. In: International Conference on Computer Vision, pp. 1482–1489 (2005)
44. Li, G., Guo, L., Liu, T.: Deformation invariant attribute vector for deformable registration of longitudinal brain MR images. Computerized Medical Imaging and Graphics **33**(5), 273–297 (2009)
45. Li, H., Manjunath, B., Mitra, S.K.: A contour-based approach to multisensor image registration. Image Processing, IEEE Transactions on **4**(3), 320–334 (1995)
46. Lindeberg, T.: Detecting salient blob-like image structures and their scales with a scale-space primal sketch: a method for focus-of-attention. International Journal of Computer Vision **11**(3), 283–318 (1993)
47. Ling, H., Jacobs, D.: Deformation invariant image matching. In: The Tenth International Conference in Computer Vision (ICCV). Beijing, China. (2005)
48. Liu, J., Vemuri, B.C., Marroquin, J.L.: Local frequency representations for robust multimodal image registration. IEEE Transactions on Medical Imaging **21**(5), 462–469 (2002)
49. Liu, J., Vemuri, B.C., Marroquin, J.L.: Local frequency representations for robust multimodal image registration. IEEE Transactions on Medical Imaging **21**(5), 462–469 (2002)
50. Lowe, D.G.: Object recognition from local scale-invariant features. In: Computer vision, 1999. The proceedings of the seventh IEEE international conference on, vol. 2, pp. 1150–1157. Ieee (1999)

51. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60**(2), 91–110 (2004)
52. Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality image registration by maximization of mutual information. IEEE Transactions on Medical Imaging **16**(2), 187–198 (1997)
53. Maintz, J.A., Viergever, M.A.: A survey of medical image registration. Medical Image Analysis **2**(1), 1–36 (1998)
54. Modersitzki, J.: FAIR: Flexible algorithms for image registration. SIAM, Philadelphia (2009)
55. Narayanan, R., Fessler, J.A., Park, H., Meyer, C.R.: Diffeomorphic nonlinear transformations: a local parametric approach for image registration. In: International Conference on Information Processing in Medical Imaging, pp. 174–185 (2005)
56. Noblet, V., Heinrich, C., Heitz, F., Armspach, J.P.: Symmetric nonrigid image registration: application to average brain templates construction. In: Medical Image Computing and Computer-Assisted Intervention : MICCAI '08, no. Pt 2 in LNCS, pp. 897–904 (2008)
57. Ou, Y., Akbari, H., Bilello, M., Da, X., Davatzikos, C.: Comparative evaluation of dramms with 11 public tools in registering skull-stripped, raw, multi-site and tumor-recurrence brain mr images. IEEE Transactions on Medical Imaging p. Accept Pending Minor Revision (2014)
58. Ou, Y., Besbes, A., Bilello, M., Mansour, M., Davatzikos, C., Paragios, N.: Detecting mutually-salient landmark pairs with MRF regularization. In: International Symposium on Biomedical Imaging, pp. 400–403 (2010)
59. Ou, Y., Conants, E., Weinstein, S., Englander, S., Da, X., Gaonkar, B., Hsiao, M., Rosen, M., DeMichele, A., Davatzikos, C., Kontos, D.: Comparison of attribute- versus intensity-based methods for longitudinal breast mri registration: Application to quantification of tumor changes during neoadjuvant chemotherapy. Magnetic Resonance in Medicine p. In Press (2014)
60. Ou, Y., Reynolds, N., Gollub, R., Pienaar, R., Wang, Y., Wang, T., Sack, D., Andriole, K., Pieper, S., Herrick, C., Murphy, S., Grant, P., Zollei, L.: Developmental brain adc atlas creation from clinical images. In: Organization for Human Brain Mapping (OHBM) (2014)
61. Ou, Y., Shen, D., Feldman, M., Tomaszewski, J., Davatzikos, C.: Non-rigid registration between histological and MR images of the prostate: A joint segmentation and registration framework. In: Computer Vision and Pattern Recognition workshop, 2009. CVPR 2009. IEEE Conference on, pp. 125–132 (2009)
62. Ou, Y., Sotiras, A., Paragios, N., Davatzikos, C.: DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting. Medical Image Analysis **15**(4), 622–639 (2011)
63. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. Pattern Analysis and Machine Intelligence, IEEE Transactions on **12**(7), 629–639 (1990)
64. Postelnicu, G., Zollei, L., Fischl, B.: Combined volumetric and surface registration. IEEE Transactions on Medical Imaging **28**(4), 508–522 (2009)
65. Roche, A., Malandain, G., Pennec, X., Ayache, N.: The correlation ratio as a new similarity measure for multimodal image registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 1115–1124 (1998)
66. Rohr, K.: On 3d differential operators for detecting point landmarks. Image and Vision Computing **15**(3), 219–233 (1997)
67. Rohr, K., Stiehl, H.S., Sprengel, R., Buzug, T.M., Weese, J., Kuhn, M.: Landmark-based elastic registration using approximating thin-plate splines. IEEE Transactions on Medical Imaging **20**(6), 526–534 (2001)
68. Rueckert, D., Aljabar, P., Heckemann, R.A., Hajnal, J.V., Hammers, A.: Diffeomorphic registration using B-splines. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 702–709 (2006)
69. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: application to breast MR images. IEEE Transactions on Medical Imaging **18**(8), 712–721 (1999)

70. Satterthwaite, T.D., Elliott, M.A., Ruparel, K., Loughead, J., Prabhakaran, K., Calkins, M.E., Hopson, R., Jackson, C., Keefe, J., Riley, M., et al.: Neuroimaging of the philadelphia neurodevelopmental cohort. NeuroImage **86**, 544–553 (2014)

71. Sederberg, T.W., Parry, S.R.: Free-form deformation of solid geometric models. ACM Siggraph Computer Graphics **20**(4), 151–160 (1986)

72. Serpa, M.H., Ou, Y., Schaufelberger, M.S., Doshi, J., Ferreira, L.K., Machado-Vieira, R., Menezes, P.R., Scazufca, M., Davatzikos, C., Busatto, G.F., et al.: Neuroanatomical classification in a population-based sample of psychotic major depression and bipolar i disorder with 1 year of diagnostic stability. BioMed Research International **2014** (2014)

73. Shen, D.: Image registration by local histogram matching. Pattern Recognition **40**(4), 1166–1172 (1997)

74. Shen, D., Davatzikos, C.: HAMMER: hierarchical attribute matching mechanism for elastic registration. IEEE transactions on Medical Imaging **21**(11), 1421–39 (2002)

75. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: a survey. IEEE Transactions on Medical Imaging **32**(7), 1153–90 (2013)

76. Sotiras, A., Ou, Y., Glocker, B., Davatzikos, C., Paragios, N.: Simultaneous geometric–iconic registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 676–683 (2010)

77. Sotiras, A., Paragios, N.: Discrete symmetric image registration. In: IEEE International Symposium on Biomedical Imaging (ISBI), pp. 342–345 (2012)

78. Szeliski, R.: Image alignment and stitching: A tutorial. Foundations and Trends® in Computer Graphics and Vision **2**(1), 1–104 (2006)

79. Tagare, H., Groisser, D., Skrinjar, O.: Symmetric non-rigid registration: A geometric theory and some numerical techniques. Journal of Mathematical Imaging and Vision **34**(1), 61–88 (2009)

80. Thirion, J.P.: Image matching as a diffusion process: an analogy with Maxwell's demons. Medical Image Analysis **2**(3), 243–260 (1998)

81. Toews, M., Wells III, W.M.: Efficient and robust model-to-image alignment using 3d scale-invariant features. Medical image analysis **17**(3), 271–282 (2013)

82. Torresani, L., Kolmogorov, V., Rother, C.: Feature correspondence via graph matching: Models and global optimization. In: European Conference on Computer Vision, pp. 596–609 (2008)

83. Tsin, Y., Kanade, T.: A correlation-based approach to robust point set registration. In: European Conference on Computer Vision, pp. 558–569 (2004)

84. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric log-domain diffeomorphic Registration: a demons-based approach. In: Medical Image Computing and Computer-Assisted Intervention : MICCAI'08, no. Pt 1 in LNCS, pp. 754–61 (2008)

85. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic Demons: Efficient non-parametric image registration. NeuroImage **45**(1, Supplement 1), S61–S72 (2009)

86. Viola, P., Wells III, W.M.: Alignment by maximization of mutual information. International Journal of Computer Vision **24**(2), 137–154 (1997)

87. Wu, Y.T., Kanade, T., Li, C.C., Cohn, J.: Image registration using wavelet-based motion model. International Journal of Computer Vision **38**(2), 129–152 (2000)

88. Yang, J., Shen, D., Davatzikos, C.: Diffusion tensor image registration using tensor geometry and orientation features. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI), pp. 905–913 (2008)

89. Yi, Z., Zhiguo, C., Yang, X.: Multi-spectral remote image registration based on sift. Electronics Letters **44**(2), 107–108 (2008)

90. Zanetti, M.V., Schaufelberger, M.S., Doshi, J., Ou, Y., Ferreira, L.K., Menezes, P.R., Scazufca, M., Davatzikos, C., Busatto, G.F.: Neuroanatomical pattern classification in a population-based sample of first-episode schizophrenia. Progress in Neuro-Psychopharmacology and Biological Psychiatry **43**, 116–125 (2013)

91. Zhan, Y., Ou, Y., Feldman, M., Tomaszeweski, J., Davatzikos, C., Shen, D.: Registering histologic and mr images of prostate for image-based cancer detection. Academic Radiology **14**(11), 1367–1381 (2007)

92. ZHANG, R.j., Zhang, J.q., Yang, C.: Image registration approach based on surf [j]. Infrared and Laser Engineering **1**, 041 (2009)
93. Zitova, B., Flusser, J.: Image registration methods: a survey. Image and Vision Computing **21**(11), 977–1000 (2003)