Discrete Image Registration : a Hybrid Paradigm

Aristeidis Sotiras

ECOLE CENTRALE PARIS

P H D T H E S I S

to obtain the title of

Doctor of Ecole Centrale Paris
Specialty : APPLIED MATHEMATICS

Defended by

Aristeidis SOTIRAS

Discrete Image Registration:
a *Hybrid* Paradigm

prepared at Ecole Centrale de Paris

defended on November 04, 2011

Committee:

Reviewers :	Daniel RUECKERT	-	Imperial College London
	William WELLS III	-	Harvard Medical School
Advisor :	Nikos PARAGIOS	-	Ecole Centrale Paris
Examiners :	Charles-Andre CUENOD	-	Hôpital Européen Georges Pompidou
	Christos DAVATZIKOS	-	University of Pennsylvania
	Xavier PENNEC	-	INRIA - Sophia-Antipolis
	Bertrand THIRION	-	INRIA - Saclay

Acknowledgments

I'll start by reciting the beginning of the once official in-office favorite poem:

> When you set sail for Ithaca,
> wish for the road to be long,
> full of adventures, full of knowledge.

Well, the road was pretty long. Four years of PhD studies and in total five years at Ecole Centrale Paris. There were the Lestrygonians and the Cyplopes as well as the Phoenician markets and the Egyptian cities. All in all, it has been a nice journey. A journey that would not be possible to come to an end, were it not the support of a great number of people, whom I would like to sincerely thank.

First, I would like to express my gratitude to my advisor Prof. Nikos Paragios for his trust, support and guidance during the whole course of this doctoral work. His investment was great and his door was always open to discuss any matter that preoccupied me, be it personal, scientific or just the amazing performance of Panathinaikos. Without his expertise and continuous encouragement this work would not be possible.

It was a great honor for me to have Prof. Charles-Andre Cuenod, Prof. Christos Davatzikos, Dr. Xavier Pennec, Prof. Daniel Rueckert, Dr. Bertrand Thirion and Prof. William Wells III as members of my thesis committee. I greatly appreciate their finding time to travel to Paris to attend my defense despite their great work load and providing me with insightful feedback. I am grateful to Prof. Daniel Rueckert and William Wells III for their fruitful reviews on my thesis manuscript, their fairness,scrutiny and generous enthusiastic remarks. It was a privilege to have such well-known researchers, whose work has been an important source of inspiration, consider my work.

I would also like to thank Prof. Cristos Davatzikos for hosting me at the SBIA lab in the University of Pennsylvania. My stay there was very inspiring, fruitful and fun. I still remember the warmth with which people welcomed me that made my stay far more easy. Many thanks to all SBIA members and especially to Evi, Ragini, Ying, Luke, Bilwaj, Kayhan, Stathis K., Yangming, Guray, Steffen and Ali for the friendly atmosphere in the

lab and the time spent outside it. I should also thank Katerina F., AK, Scot, Rebekka, Stathis A., Thanasis, Sevda and Selma for having shared lots of funny moments with me.

During this doctoral work, I had the honor and privilege to collaborate with a number of people whose input enriched both the content of the work and me as person. I particularly acknowledge Dr. Nikos Komodakis and Ben Glocker for discussions and for kindly sharing their code. My thanks also go to Yangming Ou for the fruitful collaboration that started at SBIA and all the energy that he put towards making our projects advance.

During the first steps of my thesis, I took great profit from the discussions with Georg. I greatly enjoyed working with Radhouene who taught me a lot about research and literature. Our collaboration was a great source for inside jokes. Working with Mickael was also fun. I learned a lot about optical imaging and mice. I greatly appreciate his commitment and his tireless efforts to understand the code I wrote and correct my precious bugs. My thanks also go to Stavros A. for letting me play the boss though I was filled with disappointment when the squirtle duo did not manage to eat all the food!

I am also grateful to the unofficial reviewers of my work. Ahmed, Loic, Stavros A. and AK made sure to spot all typos, inconsistencies in the formulation while at the same time made a number of recommendations that improved greatly the presentation of the content of the thesis.

I would like to acknowledge the members of the Applied Mathematics Laboratory of Ecole Centrale Paris for providing a nice environment to work. Many thanks to the greatest Snoopy fan Sylvie and Annie for helping out with all administrative issues and for their understanding and patience. Life in lab could be far more difficult if not for them. Geraldine, Catherine and Christine were also of great help with numerous bureaucratic issues related to the PhD thesis and of course the Greek army.

I am particular indebted to my colleagues that bore me in the same room during these years: Panos (my precious debugger, debugging teddy bear and geek advisor), Loic (double diplome, ecole normale, aggregation and above all papa of beautiful Lilah), Olivier (Olifou that eats iron like koutalianos), Chaohui (our Chinese spy 007, Jiji - Xiao Jiji), Salma (our office queen - Lela), Radhouene (Radadan show me the target), Ahmed (the great pants terrorist), for supporting and enhancing my extreme sense of humor (Je chauffe bien after all). They were my supporting vectors during the difficult moments and great companions during the good ones. I feel lucky that I can consider them as true friends.

On a larger scale, I am thankful to all members of the computer vision and medical imaging team for injecting so much fun in a working environment. Basket, football, blooby volley, armagetron, pranks, running sessions to the closest RER B station, long discussions about Tibet and Taiwan, manga reading, eating exotic delicacies, were few of the many extracurricular activities that we did together. All lab mates, Maxime, Noura, Regis, Olivier J., Charlotte, Martin, Daniel, Sarah, Pierre-Yves, BoBo, Katerina, Nicolas,

Stavros T., Haithem, Fabrice, Fatima were a pleasure to be around on a daily basis. Finally, I do not forget people who passed a shorter period of time at the lab: Kostas who always finds some time to treat me when I am in Athens, Mihai with whom we passed a great time exploring Romania, Dimitris who always showed interest on what I was doing, Rola, Jose, Thanos, Fahim, Wanping, Sylvain, Pascale, Samuel, Alexandros, Jean, Chris, Olfa, Michalis and I hope I didn't forget anyone.

Outside the scope of work, I enjoyed the support of many people that either brightened my stay at Paris or made me feel still at home whenever back to Greece or sometimes both. I would like to thank Maria for somehow managed to survive an overdose of some of my most "interesting" work focused-on moments combined with a certain degradation of my omnipresent jokes. One crevette would like to thank Francois, Komporis family and company. Life at Paris was fun due to the presence of the Italian mafia leaded by Mara, Valeria, Alessandro, Paola and Laura that were kind to host me and show me some of Italy, Sylvia, crazy Maria P., Nathalie, Nina, Patrick and Angeliki. All the parisaki friends, Manos, Anthi and their young devils, Didi, the Unbeatable one, Alex, Foteini, Evi, Dimitris Geitonas, Undo, Eleni, Varkouda & bro, Orestis, Pantelis M. & Eri for having explored the great variety of french wines with me. Arctouros would like to thank Anastase, Vasw, Nikolas and Dialecti. Christos, Marina, Michaela for the poem and chocolate sessions. George and Georgia for many hours of many things that were sympa, the best duo of blog-writers, poets and philosophers, sun and stars and the list goes on, Tolis, Kostas, Vasw, Panos, Vagelis, Pantelis T., my distinguished Tripoli company and more and more that either have heard me complain without stop or making jokes (frankly I am not sure which one is worse). And as this thanks giving journey seems to take so long before passing to the next stop, thank you my friends and sorry if I failed to write down everything here.

Last but not least, I would like to thank all my Greek family. And Greek in the sense that family is considered really in its most extended sense. My cousins, aunts and uncles that not only expressed their support, interest and love whenever I am back at Greece but at many occasions traveled here to either help or share my joy. I really appreciate and acknowledge their help throughout many years far exceeding the ones of the PhD. It goes without saying that a tribute should be paid to my brother for his love and support and my parents for all their sacrifices, unconditional love and trust and for so many things that duly acknowledging goes beyond the scope of this section. To them goes my everlasting gratitude.

This thesis is dedicated to my parents and brother.

Abstract

This thesis is devoted to dense deformable image registration/fusion using discrete methods. The main contribution of the thesis is a principled registration framework coupling iconic/geometric information through graph-based techniques. Such a formulation is derived from a pair-wise MRF view-point and solves both problems simultaneously while imposing consistency on their respective solutions. The proposed framework was used to cope with pair-wise image fusion (symmetric and asymmetric variants are proposed) as well as group-wise registration for population modeling. The main qualities of our framework lie in its computational efficiency and versatility. The discrete nature of the formulation renders the framework modular in terms of iconic similarity measures as well as landmark extraction and association techniques. Promising results using a standard benchmark database in optical flow estimation and 3D medical data demonstrate the potentials of our methods.

keywords: Markov Random Fields, registration, group-wise, iconic, geometric

Résumé

La présente thèse est consacrée au recalage et à la fusion d'images de façon dense et déformable via des méthodes d'optimisation discrète. La contribution majeure consiste en un principe de couplage entre recalage géométrique et iconique via l'utilisation de méthodes dites graphiques. Une telle formulation peut être obtenue à partir d'un Champ de Markov Aléatoire binaire et permet de résoudre les deux problèmes simultanément tout en imposant une cohérence à leurs solutions respectives. La méthodologie s'applique à la fusion de paires d'images (dans ses versions symétrique et asymétrique), ainsi qu'au recalage simultané de groupes d'images nécessaire à l'étude de populations. Les qualités principales de notre approche résident dans sa faible complexité algorithmique et sa versatilité. L'utilisation d'une formulation discrète assure une grande modularité concernant tant la mesure de similarité iconique que l'extraction et l'association de points d'intérêt. Les résultats prometteurs obtenus sur les bases de données de référence en flot optique et sur des données médicales tridimensionnelles démontrent tout le potentiel de notre méthodologie.

mots clés: Champs de Markov Aléatoire, recalage, analyse simultanée, iconique, géométrique

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The present work lies at intersection of Computer Vision and Medical Image Analysis. It aims to address the problem of establishing accurate dense correspondences between images. When correspondences in space are considered, the problem is usually referred to as *image registration*. This is a problem that arises predominantly in medical image processing under three scenarios. Under the first scenario, we aim to establish correspondences between successive exams of the same subject (longitudinal modeling). Under the second scenario, the goal is to fuse different image modalities to perform better diagnosis. Finally, under the third scenario, correspondences between different subjects are sought towards population modeling. When correspondences in time are considered, the problem is referred to either as *optical flow* or *motion estimation*. In this case, a sequence of images is given as input and the goal is to estimate the motion from one time instance to the next under the assumption of preservation of intensity values over time. Correspondence estimation, along with image segmentation, is one of the most important problems in the field of image processing. Despite important advances, registration is still considered a challenging problem.

In this introductory chapter, in section 1.1 we give a brief description of the context of this work. Next, in section 1.2, after having introduced the main challenges posed by the current technological and methodological advances, we present the main ideas of this work. Section 1.3 concludes the chapter giving an overview of the work done in the context of this thesis.

1.1 Thesis Context

There are many computer vision and medical imaging tasks that call for the estimation of correspondences. The span of applications stretches from numerous clinical to industrial

ones. It is by considering this extreme variety of applications that one can understand the importance of the problem.

From the clinical point of view, one can imagine that the goal of automating medical diagnosis with the aid of computers would be out of consideration unless reliable registration methods exist. Diagnosis may be assisted by permitting the doctors to quantify the changes appearing to the patient's anatomical structures on the course of time or in a differentiation setting to compare a patient to a normal subject and localize the differences of pathological origin. In the context of diagnosis, registration delivered quantified information like myocardial strain estimation may be considered as an important biomarker for the monitoring of cardiac diseases and the estimation of cardiac injuries.

Computer assisted diagnosis may be the most challenging objective sought in medical imaging but still there exist numerous clinical applications upon which registration has a great impact. Spatial normalization allows for longitudinal or population modeling of anatomical structures extending our knowledge with respect the anatomical variability. Knowing the normal anatomical variability allows for group studies that reveal statistical significant regions influenced by pathogenies.

Image registration plays a vital role in treatment planning. One can imagine registration between healthy subjects and ones with brain tumors in order to localize important brain structures to be taken into consideration for surgical planning. Registration is important even during the surgery itself as it allows for accurate localization of anatomical structures accounting for position shifts induced by surgical operations. In terms of treatment planning, we should also note the impact that registration has in radiotherapy by localizing tumorous cells and thus limiting the destruction of healthy ones.

From the industrial point of view, motion estimation has been an important component of many systems. Video surveillance systems constitute one of the most prominent directions to boost security. Cameras capturing video sequences are installed in sensitive public areas towards automating the detection and tracking of activity. Both tasks rely heavily on motion estimation.

Motion estimation is also important for more low-level video processing tasks. Vintage film restoration can profit from motion estimation to restore a frame based not only on spatial but also temporal information. That way artifacts or even subtitles can be removed accurately. Two problems that are closely related to the hardware and can benefit from motion estimation algorithms are the de-interlacing and frame interpolation ones. In the first case, the recording device acquires the video with alternating fields of even and odd lines and special care is taken to convert it to a signal that contains both. In the second case, new frames are created by interpolating between two successive ones i.e. to match an incoming video to the display's refresh rate. Last but not least, in order to compress video information to either store or transmit it efficiently, motion estimation is indispensable and

currently an important component of MPEG encoding.

1.2 Thesis Statement

The vast range of the potential applications as well as the technological advances create a dynamically changing environment that poses important challenges. Being mostly in the front end of an application, correspondence estimation algorithms must be computational efficient, operate even on real-time constraints, and highly versatile. Moreover, especially in clinical settings the solution they provide should exhibit certain desired properties to facilitate or even allow for reliable subsequent processing.

Two main reasons fuel the need for efficiency. On one hand, the volume of acquired data grows with a fast rate. Video monitoring devices are installed constantly. Filming videos is an important industry but also an everyday activity. Most importantly, medical imaging has become standard clinical practice. On the other hand, the time constraints imposed either in an industrial or a clinical setting become more and more important. One can only imagine the importance of time during a surgical operation where the surgeon needs spatially normalized data to perform the surgery.

A second challenge is revealed if we consider either the vast range of problems to be tackled or the new imaging devices introduced in the clinical practice resulting in images with different properties. In its essence, the problem is always the same, establishing correspondences between two or more images. Thus, one should expect that an appropriate solution to it should be versatile enough to be used for any of the previous applications or image modalities. This underlines the need for modularity with respect to the main components of the algorithms.

The broad purpose of this thesis is to develop novel methods to tackle the problem of recovering correspondences between images, especially ones of medical interest, in an efficient and modular way. To respond to the needs advocated earlier, efficient approximate schemes have been devised to achieve a good balance for the important accuracy versus efficiency trade off. In an effort to render the proposed algorithms both more efficient and modular, discrete optimization techniques based on a Markov Random Field formulation of the problem have been preferred. Finally, the particularities of the application of such methods in statistical studies have been accounted for by thoughtful design of the proposed algorithms.

1.3 Thesis Roadmap

The remainder of this thesis is organized into four chapters. In Chapter 2 an extensive state-of-the-art review of deformable registration is provided. Additional emphasis has be given to techniques applied to medical images towards positioning this thesis while demonstrating the important challenges of our work.

Chapter 3 covers the introduction of our novel hybrid registration framework. The proposed method is able to provide diffeomorphic deformation fields by considering in a coupled way both iconic and geometric information. The problem is solved in an efficient one-shoot optimization.

Following, in Chapter 4 the previous method is extended and endowed with the symmetric property. In this setting, larger deformations can be recovered by allowing both images to deform towards a common domain. Moreover, the result of the algorithm is the same regardless the order the input images are given.

The previous symmetric algorithm is then extended in Chapter 5 to the group-wise case, where a set of images are considered simultaneously and registered towards a common domain that is not explicitly chosen. That way, statistical studies may be performed without caring for the bias that is introduced by the choice of a template domain.

The thesis is concluded in Chapter 6 with a discussion over the work presented in this manuscript, while suggesting different research directions for the future.

Chapter 2

Deformable Registration: A Survey

2.1 Introduction

In general, in image registration two images are considered. One is usually referred to as source or moving image, while the other is referred to as target or fixed image. In this doctoral study, the source image is denoted by $S : \Omega_S \subset \mathbb{R}^d \mapsto \mathbb{R}$, while the target image by $T : \Omega_T \subset \mathbb{R}^d \mapsto \mathbb{R}$, $d = \{2, 3\}$. The source image undergoes a transformation $\mathcal{T} : \Omega_S \mapsto \mathbb{R}^d$.

The goal of registration is to estimate the optimal transformation. This is often achieved by means of an energy minimization problem:

$$\arg\min_{\boldsymbol{\theta}} \mathcal{M}(T, S \circ \mathcal{T}(\boldsymbol{\theta})) + \mathcal{R}(\mathcal{T}(\boldsymbol{\theta})). \tag{2.1}$$

The previous energy (Eq. 2.1) comprises two terms. The first term, $\mathcal{M}$, quantifies the level of alignment between a target image T and a source image S under the influence of transformation $\mathcal{T}$ parametrized by $\boldsymbol{\theta}$. The notation $S \circ \mathcal{T}$ will be used interchangeably with $S \circ \mathbf{u}$ to denote that a moving image is deformed. The second term, $\mathcal{R}$, 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 $\mathcal{T}(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x})$ where $\mathbf{u}$ is the deformation field. The velocity field, denoted as $\mathbf{v}$, is defined as $\mathbf{v} = \partial_t \mathbf{u} + v^T \nabla \mathbf{u}$. An effort to keep the notation consistent throughout the document has been made. Hereafter, vectors are denoted in bold, while the upper-script is used to index their components. Due to the diversity of the models described, the meaning of some symbols may change. Whenever a change occurs, the signification of the symbol will be stated explicitly.

Image registration involves three main components: (i) a deformation model; (ii) an objective function; and (iii) an optimization strategy. Image registration is a problem that

has been studied in great detail during the past few decades and many innovative ideas regarding these three main aspects have been proposed. General reviews of the field may be found in [2, 3, 4, 5, 6, 7, 8]. However due to the rapid progress of the field such reviews are to a certain extent outdated.

The aim of this chapter is to provide a thorough overview of advances of the past decade in deformable registration. Nevertheless, some *classic* papers that have greatly influenced the advance of the ideas in the field are mentioned. Even though our primary interest is deformable registration, for the completeness of the presentation, references to linear methods are also included as many problems have been treated in this low degrees of freedom setting before being extended to the deformable case.

The chapter is divided following the previous structural separation of registration algorithms in three components: (i) a deformation model; (ii) an objective function; and (iii) an optimization strategy.

2.2 Deformation Models

The choice of deformation model is of great importance for the registration process as it entails an important compromise between computational efficiency and richness of description. The parameters that registration estimates through the optimization strategy correspond to the degrees of freedom of the deformation model. Their number varies greatly, from around 10, in the case of global linear transformations, to millions, when non-parametric dense transformations are considered. The greater the dimension of the state space, the richer the descriptive power of the model though at the cost of more challenging and computational demanding inference. Furthermore, the choice of the deformation field often implies an assumption regarding the nature of the deformation to be recovered.

The rest of the section is organized following the classification of deformation models given by Holden [9]. More emphasis is put on aspects that were not covered by that review. According to [9], geometric transformations can be classified into two main categories:(i) those that are inspired by physical models; and (ii) those inspired by interpolation and approximation theory.

Of great importance for biomedical applications are the constraints that may be applied to the transformation such that it exhibits special properties. Such properties include, but are not limited to, inverse consistency, symmetry, topology preservation and diffeomorphicity. The value of these properties was made apparent to the research community and were gradually introduced as extra constraints.

- **Inverse consistency:** Despite common intuition, interchanging the order of input

images could produce transformation changes, making the choice of template an important task and biasing the subsequent statistical analysis. Moreover, the composition of the forward and backward transforms (or the transforms that we obtain by interchanging the role of the images) is not the identity one. In some cases, the transformations may not even be invertible. To account for that counter intuitive result, inverse consistent algorithms have been developed.

- **Symmetry:** Taking one step further the previous idea, symmetric algorithms have been conceived that guarantee that the result does not change by swapping the inputs. The distinction between symmetric and inverse consistent algorithms lies in their primary aim. The first guarantee that the result is invariant to the order of the images, while the second approximates symmetry by imposing the inverse consistency constraint.

- **Topology preservation:** The previous properties are desired and are usually imposed as soft constraints. Thus, the resulting transformation is not always one-to-one and crossings may appear in the deformation field. Such an effect in undesirable because it is quite often the case (especially when dealing with normal subjects) that we know that true correspondences can be obtained for every point of the image domain. What is more, the mapping from one anatomical structure to another should conserve the relative positions of points. Such a demand has motivated topology preserving algorithms whose deformation field is one-to-one or has a Jacobian determinant that is not zero.

- **Diffeomorphicity:** On top of the topology preservation constraint, ensures that the transformation field exhibits a certain level of smoothness.

In the following sections, the most important methods of the two classes are presented with emphasis on the approaches that endow the model under consideration with the previous desirable properties. One subsection is devoted to methods that employ task-specific constraints.

2.2.1 Geometric Transformations Inspired by Physical Models

Following [3], physical models can be further separated in five categories:

- Elastic model.

- Viscous fluid flow model.

- Diffusion model.

- Curvature registration.
- Flows of diffeomorphisms.

Elastic Model

In this case, the image under deformation is modeled as an elastic body. The Navier-Cauchy equation describes the deformation, or

$$\mu\nabla^2\mathbf{u} + (\mu + \lambda)\nabla(\nabla \cdot \mathbf{u}) + \mathbf{F} = 0, \tag{2.2}$$

where $\mathbf{F}(\mathbf{x})$ is the force that drives the registration based on an image matching criterion, μ refers to the rigidity that quantifies the stiffness of the material and λ is Lamé's first coefficient.

Image grid was modeled initially as an elastic membrane in [10]. Two different types of forces compete to reach an equilibrium. The external force tries to deform the image such that matching is achieved while the internal enforces the elastic properties of the material. The resulting equation is solved by deploying a finite difference scheme in an iterative manner.

Bajcsy and Kovacic [11] extended this approach in a hierarchical fashion where the solution of the coarsest scale is up-sampled and used as initialization for the finer one. Linear registration is used at the lowest resolution.

The linear elastostatic problem can be formulated in a variational setting where the displacement is given by minimizing:

$$E = \frac{1}{2}\int_V \boldsymbol{\sigma} : \boldsymbol{\epsilon}\, dV - \int_V \mathbf{F} \cdot \mathbf{u}\, dV - \int_S \mathbf{t} \cdot \mathbf{u}\, dS, \tag{2.3}$$

where $\boldsymbol{\epsilon}$ denotes the strain tensor, $\boldsymbol{\sigma}$ the stress tensor and $\mathbf{t}$ the tractions that are distributed over a surface S of a volume V. The inner product between two tensors is denoted as $\cdot : \cdot$. The first term in the equation entails the internal forces while the other two, the external body and surface ones.

In [12] the previous equation 2.3 was considered in a Bayesian setting allowing for the computation of the uncertainty of the solution as well as for confidence intervals. The Finite Element Method (FEM) is used where the displacements are inferred for the element nodes and interpolated elsewhere.

Linear elastic models have been also used when registering brain images based on sparse correspondences. In [13], a mapping between the cortical surface is established, two images are brought into correspondence by estimating a global transformation modeling

images as inhomogeneous elastic objects. The elasticity parameters vary as a function of location in the brain image to compensate for the fact that certain structures tend to deform more than others. In addition, a non-zero initial strain was considered so that some structures expand or contract naturally.

In general, an important drawback of registration process is that when source and target volumes are interchanged, the obtained transformation is not the inverse of the previous solution. Christensen and Johnson [14] tried to overcome this shortcoming by estimating simultaneously both transformations while penalizing inconsistent transformations by adding a constraint to the objective function. The regularization constraint is that of linear elasticity though the transformation is parametrized by a $3D$ Fourier series.

To tackle the inconsistency problem Leow *et al.* [15] took a different approach. Instead of adding a constraint that penalizes the inconsistency error, they proposed a unidirectional approach that couples the forward and backward transformation and provides inverse consistent transformations by construction. This method considers only the forward mapping to minimize the symmetric energy:

$$E(\mathcal{T}) = \underbrace{\int_{\Omega_T} \|S(\mathcal{T}(\mathbf{x})) - T(\mathbf{x})\|^2 d\mathbf{x} + \omega\mathcal{R}(\mathcal{T})}_{E_1} + \underbrace{\int_{\Omega_S} \|T(\mathcal{T}^{-1}(\mathbf{x})) - S(\mathbf{x})\|^2 d\mathbf{x} + \omega\mathcal{R}(\mathcal{T}^{-1})}_{E_2}. \tag{2.4}$$

Minimizing the previous energy $E(\mathcal{T})$ in a gradient descent fashion would lead to the following update rule:

$$\mathcal{T} \rightarrow \mathcal{T} + \delta\eta_1 + \delta\eta_2,\ \mathcal{T}^{-1} \rightarrow \mathcal{T}^{-1} + \delta\xi_1 + \delta\xi_2 \tag{2.5}$$

where δ is the step, η_1 and ξ_1 are the gradient descent directions of E_1 when considering the forward and backward warping respectively. Similarly, η_2 and ξ_2 are defined for E_2. By taking into account that the composition of the two maps should give the identity, η_2 can be expressed as:

$$\eta_2 = -\nabla \cdot (\mathcal{T}(\mathbf{x}))\xi_2((\mathcal{T}(\mathbf{x})). \tag{2.6}$$

The previous formulation permits to minimize the symmetric energy by considering only the forward mapping but at the same time imposing that the transformation is inverse consistent.

However, linear elastic models is are limited and can efficiently handle the case of small deformations. In [16], a way to recover large deformations, that are inverse consistent is given by considering a sequence of small deformation transformations. Each

incremental transformation is modeled by a linear elastic model. A periodic sequence of images is considered where the first (or last) and middle image are the source and target respectively. Each incremental transformation maps an image of the sequence to the one that proceeds it. The concatenation of all transformations maps one input image to another. The objective function to be minimized is the summation of the intensity similarity cost evaluated over two successive pairs of images.

A different way to produce inverse consistent transformations that can account for large deformations is proposed in [17]. They modeled the deformation process through the St Venant-Kirchoff elasticity energy:

$$\mathcal{R} = \int \frac{\mu}{4} Tr((\boldsymbol{\sigma} - \mathbf{I})^2) + \frac{\lambda}{8} Tr(\boldsymbol{\sigma} - \mathbf{I})^2. \tag{2.7}$$

Deformations that result in strain tensors that differ from the identity (or rigid transformation) are penalized by measuring the Euclidean matrix distance. Such an approach considers the trace of the squared difference of the strain tensor with the identity $Tr((\boldsymbol{\sigma} - \mathbf{I})^2)$. However, the Euclidean metric has been proven not good for the tensor space while its asymmetry results in the inverse-inconsistency problem. Such a limitation was dealt with the use of log-Euclidean metrics in the place of the Euclidean ones, resulting in the Riemannian elasticity energy which is inverse consistent:

$$\mathcal{R} = \frac{1}{4} \int Tr(\log(\boldsymbol{\sigma})^2). \tag{2.8}$$

In [18], fast diffeomorphic image registration was proposed based on either membrane, bending or linear elastic energy. The velocity field is assumed to be constant over time. The solution is thus given through integration over time by composing successive solutions. Given a pair number of steps, this can be performed efficiently by a scaling and squaring approach. Furthermore, the exponential of the flow field is considered to guarantee that then the mapping is diffeomorphic. The energy is optimized by using the Levenberg-Marquardt algorithm and a full multi-grid approach in order to compute efficiently its update step.

Viscous Flow Model

In this case, the image under deformation is modeled as a viscous fluid. The transformation is governed by the Navier-Stokes equation:

$$\mu_f \nabla^2 \mathbf{v} + (\mu_f + \lambda_f) \nabla(\nabla \cdot \mathbf{v}) + \mathbf{F} = 0. \tag{2.9}$$

Operating on velocities, such a model permits to recover large deformations. The first term of the Navier-Stokes equation (Eq. 2.9), constrains neighboring points to deform similarly

by spatially smoothing the velocity field. The second term allows structures to change in mass while μ_f and λ_f are the viscosity coefficients.

Viscous fluid flow transformations were introduced in medical image registration in [19, 20]. In [19], a template is modeled as a viscous fluid allowing for large magnitude non-linear deformations. The Partial Differential Equation (PDE) is solved for small time intervals and the complete solution is given by an integration over time. For each time interval a successive over-relaxation scheme is used. The Jacobian is monitored and each time its value falls under 0.5, the template is re-gridded and a new one is generated that is subsequently used to estimate a transformation. The final solution is the concatenation of all successive transformations occurring for each re-gridding step. That way, the topology is guaranteed to be preserved. In a subsequent work [20], a hierarchical way to recover the transformations for brain anatomy is presented. Initially, global affine transformation is performed followed by a landmark transformation model. The result is refined by fluid transformation preceded by an elastic registration step.

Computational inefficiency is the main drawback of viscous fluid models and the reason they have not gained the popularity of the rest of transformation models. To circumvent this shortcoming, a massive parallel computer implementation has been proposed in [19] while Bro *et al.* [21] proposed a technique based on a convolution filter in scale-space. The filter is designed as the impulse response of the linear operator $L = \mu_f \Delta \mathbf{u} + (\mu_f + \lambda_f)\nabla(\nabla \cdot \mathbf{v})$ defined in its eigen-function basis. A multi-grid approach was proposed in [22] towards handling anisotropic data along with a multi-resolution scheme opting for recovering first coarse velocity estimations and refining them in a subsequent step.

An inverse consistent variant of fluid registration was proposed in [23] to register Diffusion Tensor images. Symmetrized Kullback-Leibler (KL) divergence was used as matching criterion. Inverse consistency was achieved by evaluating the matching and regularization criterion towards both directions. Fluid deformation models have also been used to tackle multi-modal registration [24] or in an atlas-enhanced registration setting [25].

Diffusion Model

In this case, the deformation is modeled by the diffusion equation:

$$\Delta \mathbf{u} + \mathbf{F} = 0, \tag{2.10}$$

that corresponds to the following regularization energy:

$$\mathcal{R} = \int \sum_{j=1}^{d} \|\nabla u^j(\mathbf{x})\|^2 d\mathbf{x}. \tag{2.11}$$

In the previous equation, the upper-script denotes order of the spatial dimension. Algorithms based on this transformation model exploit the fact that the Gaussian kernel is the Green's function of the diffusion equation Eq. 2.10 to provide for an efficient regularization step.

Thirion [26], inspired by Maxwell's Demons, was the first to propose diffusion models in image registration. In this work, object boundaries are modeled as membranes and the image is diffused through it under the influence of the Demons placed inside the membranes. The algorithm operates in two steps: (i) estimation of the displacements for every point, and (ii) a regularization step. The Demon force is calculated by considering the optical flow constraint that is valid for small displacements. Regularization is achieved through Gaussian smoothing.

Demons, as initially introduced, is an efficient algorithm able to provide dense correspondences but lacks a sound theoretical justification. Due to the success of the algorithm, a number of papers tried to give theoretical insight into its workings. Fischer *et al.* [27] gave a fast algorithm for image registration. The result was given as the solution of a diffusion PDE. An efficient scheme for its solution was proposed while a connection to Thirion's Demons algorithm was drawn.

The most successful attempt to shed light on Demons is described in *et al.* [28]. Image registration is formulated as an energy minimization problem and the connection of the Demons algorithm with gradient descent schemes is shown.

Vercauteren *et al.* [29] adopted the alternate optimization framework introduced by Cachier [30] to relate symmetric Demons forces with the Efficient Second-order Minimization (ESM)[31]. The transformation was computed by optimizing the global energy

$$E = \underbrace{\underbrace{\mathcal{M}(T, S \circ \mathbf{u}) + \mathcal{D}(\mathbf{u}, \mathbf{w})^2}_{Matching} + \mathcal{R}(\mathbf{w})}_{Regularization} . \tag{2.12}$$

Optimization can be achieved efficiently as the introduction of the auxiliary variable v decouples the problem into two parts that can be solved separately very fast. Matching was performed by minimizing the first term through ESM optimization while regularization was achieved by Gaussian smoothing.

In [32] Thirion's algorithm was endowed with the diffeomorphic property. Diffeomorphic transformations were parametrized by using stationary speed vector fields allowing for a fast computation of exponentials based on the scaling and squaring method.

To further facilitate the use of the Demons algorithm in anatomical computational studies, Demons were provided with invertibility and symmetry properties in [33]. This was achieved by representing the complete spatial transformation in the log-domain. A sym-

metric extension was also proposed by considering in the energy the difference between the two images under the influence of the forward transform and its inverse.

Recently, efforts opting for a mathematical justification of the smoothing step in order to enable for deformations bearing different physical properties [34, 35, 36] have been considered.

Cahil *et al.* [34] showed that curvature and fluid registration can be formulated as two coupled diffusion equations. Their stationary solution may be approached via successive Gaussian convolutions, thus yielding a Demons algorithm for these cases. In a subsequent work [35], they showed how to extend the curvature regularization to consider local image gradient content. Again, a coupled PDE system was proposed whose stationary solution can be attained by consecutive convolutions with the Green's function of the diffusion equation. Another way to perform adaptive smoothing was presented in [37] where a non-stationary diffusion filter was used exploiting available knowledge regarding the deformability of tissues. In areas where greater deformations were expected, diffusion was scaled down. On the contrary, inside objects where coherence should be preserved, diffusion was scaled up.

Mansi et al. [36] introduced a physical constraint in the registration process to estimate the myocardium strain from Cine-MRI. The logDemons algorithm [33] was endowed with the incompressibility constraint by making the velocity field divergence-free. This was achieved by solving the Poisson equation under 0-Dirichlet boundary conditions within a subdomain of the image showing the myocardium.

A main drawback of this family of methods is that the Demons forces are usually calculated based on a SSD criterion and thus are appropriate for mono-modal image registration. Numerous efforts have been made towards extending Demons for multi-modal registration. Guimond *et al.* [38] proposed a method that alternates between Demons based registration and intensity correction. Other efforts include the encoding of similarity metrics such as Normalized Mutual Information [39, 40].

The application of the Demons algorithm is not limited to scalar images but has been extended to diffusion tensor images [41], multi-channel ones [42] as well as different geometries [43]. In [41], Demons forces were derived by considering the squared difference between each element of the Log-Euclidean transformed tensors while taking into account the reorientation introduced by the transformation. Multi-channel Demons were considered to register $4D$ Time-Series of Cardiac Images by enforcing trajectory constraints in [42]. Each time instance was considered as a different channel. In addition, the transformation between successive channels was considered as given. Finally, the Demons framework was employed to register cortical surfaces parametrized as spheres in [41]. To generalize Demons on the sphere, a way to measure the distance between two transformations and to regularize the transformation was introduced.

Curvature Registration

In this case, the deformation is modeled by the following equilibrium equation:

$$\Delta^2 \mathbf{u} + \mathbf{F} = 0. \tag{2.13}$$

The regularization part of the energy to be minimized in this specific case, is given by

$$\mathcal{R} = \int \sum_{j=1}^{d} (\Delta u^j(\mathbf{x}))^2 d\mathbf{x}. \tag{2.14}$$

The advantage of such a regularization scheme is that it does not penalizes affine linear transformations thus removing the need for a an additional affine linear pre-registration step.

Fischer and Modersitzki were the first to introduce this constraint in [44, 45]. To solve equation Eq. 2.13, the Gâteaux derivatives with respect to the data and regularization term were calculated:

$$\int_{\Omega} \left\langle \mathbf{F}(\mathbf{x}, \mathbf{u}(\mathbf{x})) + \Delta^2 \mathbf{u}, \mathbf{v}(\mathbf{x}) \right\rangle_{\mathbb{R}^d} d\mathbf{x}. \tag{2.15}$$

To solve the resulting PDE a finite difference scheme was used by introducing an artificial time parameter while the following boundary conditions were considered:

$$\nabla u^j(\mathbf{x}) = \nabla \Delta u^j(\mathbf{x}) = 0, \ \mathbf{x} \in \partial\Omega, \ j = 1, \cdots d. \tag{2.16}$$

The use of the previous Neumann boundary conditions was motivated by the fact that the resulting matrix problem is highly structured and thus can be solved fast. While this is true, the resulting underlying function space penalizes the affine linear displacements as pointed out by Henn in [46]. Thus, Henn proposed to include second-order terms as boundary conditions in the energy and applied a semi-implicit time discretization scheme to solve the full curvature registration problem.

In [47], another way to solve the curvature based registration problem was proposed. Instead of devising a numerical scheme to solve the PDE that results from the equilibrium equation (Eq. 2.13), recursive convolutions with an appropriate Green's function were used following [21].

Flows of Diffeomorphisms

Last but not least, flows of Diffeomorphisms have been proposed to model the deformation. In this case, the deformation is modeled by considering its velocity over time

according to the Lagrange transport equation [19, 48, 49]. The regularization constraint constrains the velocity field to be smooth:

$$\mathcal{R} = \int_0^1 \|\mathbf{v}_t\|_V^2 dt. \tag{2.17}$$

$\|\cdot\|_V$ is a norm on the space V of smooth velocity vector fields defined as $\|f\|_V = \|Lf\|_{L_2}$, where L is a differential operator and $\|\cdot\|_{L_2}$ is the standard L_2 norm of square integrable functions. By letting the velocity field vary over time, it is feasible to estimate large deformations [50].

The advantage of this framework, known as Large Deformation Diffeomorphic Metric Mapping (LDDMM), is that it allows for the definition of a distance between images or sets of points [51, 52]. The distance is defined as a geodesic according to a metric and can be used for studies of anatomical variability [53]. A number of theoretical aspects of this framework and especially the ones related with computational analysis were further developed in [54, 55, 56, 57, 58]. The interested reader is referred to [59] for an overview of the evolutions and the corresponding equations.

The LDDMM framework has been proven extremely versatile and has been extended to treat a number of problems, *e.g.* volume registration for scalar, vector- and tensor-valued data [50, 60, 61, 62, 63, 64], point-matching [51], point-matching on spheres [65], matching sets of unlabeled points [66, 67, 68], shape-matching [69, 70], curve-mapping [71, 72, 73, 74] and hybrid registration [75, 76].

Even though the LDDMM framework provides diffeomorphic transformations, it is not symmetric. To encode the symmetric property a number of approaches have been proposed [77, 60, 61]. Beg *et al.* [60] focused on providing symmetric data terms. Younes also discussed ways to render the alignment process symmetric in [77] while a symmetric LDDMM registration process driven by cross-correlation was presented in [61].

The mathematical rigorousness of the LDDMM framework comes at an important cost. The fact that the velocity field has to be integrated over time results in high computational and memory demands. Moreover, the gradient descent scheme that is usually employed to solve the optimization problem of the geodesic path estimation converges slowly. More efficient optimization techniques for the LDDMM have been investigated in [78, 79, 62].

Marsland *et al.* [78] formulated the problem in a PDE framework and a particle method was used to solve for the diffeomorphism. A similar approach that involves a particle mesh method was presented in [79] by Cotter *et al.* More recently, a Gauss-Newton implementation of the previous algorithm was given in [62]. These approaches are based upon the fact that knowing the initial velocity field suffices to calculate the intermediate and final deformation.

Instead of devising more sophisticated optimization schemes to calculate efficiently

diffeomorphisms, one could simplify the problem by decreasing its degrees of freedom. Stationary velocity fields [80] have been used towards this direction. Despite being limited with respect to the diffeomorphisms that they can capture, stationary velocity fields are a common choice among many researchers [18, 81, 82, 83]. Hernandez *et al.* followed this approach and considered stationary Ordinary Differential Equations (ODEs) in the LDDMM framework [84].

2.2.2 Geometric Transformations Inspired by Interpolation Theory

Rather than being motivated by a physical model, models of this class are derived either by interpolation theory or approximation theory. In interpolation theory, displacements, considered to be known for a restricted set of locations in the image, are interpolated for the rest of the image domain. In approximation theory, we assume that there is an error in the estimation of displacements. Thus, the transformation approximates smoothly the known displacements rather than taking the exact same values. The success of these models lies in the fact that they are rich enough to describe the transformations present in image registration problems while having low degrees of freedom facilitating the inference of the parameters.

Radial Basis Functions

One of the most important families of interpolation strategies are the Radial Basis Functions (RBFs), where the value at an interpolation point $\mathbf{x}$ is given as function of its distance r from the known sample $\mathbf{p}$, or

$$\mathbf{u}(\mathbf{x}) = \sum_{i=1}^{N} \omega_i \phi(\|\mathbf{x} - \mathbf{p}_i\|). \tag{2.18}$$

An evaluations study comparing RBFs used as transformation functions in non-rigid image registration was first given in [85]. A more thorough presentation and an analysis with respect to their property to preserve topology was given more recently in [86]. The advantage of the RBFs is that they are able to interpolate a deformation field from irregularly placed known values. Their main disadvantage lies in the fact that they often have a global support. As a result, knowing the displacement at one point influences the values that points in the whole image domain will take. A behavior that is undesirable, especially when local transformations are to be recovered.

One of the first RBFs that has been used in the field, and still of great importance, is Thin-Plate Splines (TPS) [87, 88]. Given N control points p whose displacement is

known, the displacement field is given by:

$$\mathbf{u}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b} + \sum_{i=1}^{N} \xi_i \phi(\|\mathbf{x} - \mathbf{p}_i\|), \tag{2.19}$$

where $\phi(r) = r^2 \log r$ in the $2D$ case and $\phi(r) = r$ in the $3D$ case. r is the distance $r = \|\mathbf{x} - \mathbf{p}_i\|$. The weighting factors ξ_i as well as the matrices $\mathbf{A}$ and $\mathbf{b}$ are calculated by solving a linear equations system. In the $2D$ case, TPS minimize the bending energy:

$$\mathcal{R} = \int \frac{\partial^2 \mathbf{u}}{\partial x^2} + 2\frac{\partial^2 \mathbf{u}}{\partial x \partial y} + \frac{\partial^2 \mathbf{u}}{\partial y^2} dx dy, \tag{2.20}$$

assuming infinite boundary conditions.

TPS are known to exhibit certain shortcomings. The transformation from one image domain to another is not inverse consistent. Moreover, their support is global not allowing to recover local image warping. Furthermore, TPS do not take into consideration possible errors in the estimation of the displacements in the landmark positions. Lastly, as the number of points increase, the interpolation becomes computationally demanding. A number of researchers have worked towards lessening the importance of these drawbacks [89, 90, 91, 92, 93].

In [89], Johnson and Christensen tackled the inverse inconsistency problem. They considered the minimization of the energy Eq. 2.20 under cyclic boundary conditions in an effort to account for the great consistency error that they observed in the boundary of the images. Additionally, a term that penalizes the consistency error was introduced in the objective function to render the registration inverse consistent.

The problem with the global nature of TPS was treated in [91]. TPS were constructed in such a way that their support is restricted locally. In a subsequent work, Yang *et al.* [92] defined the support of each point in an adaptive way by taking into consideration the distribution of the points in the image domain.

Rohr *et al.* [90] proposed to take into consideration the landmark localization error when estimating the dense deformation field through the use of approximating Thin-Plate Splines. The proposed method is able to handle both isotropic and anisotropic errors in the estimation of the landmark position.

Three ways to address the computational problems related with the presence of a great number of landmarks were studied by Donato *et al.* [93]. The straightforward approach of sub-sampling the points was compared to more elaborated ones that use either a subset of the basis functions or matrix approximation techniques.

Clamped-Plate Splines [52, 94] minimize the same energy as TPS though under specific boundary conditions. Following the LDDMM framework, the energy is solved so

that the resulting dense deformation field is diffeomorphic and thus this type of splines is referred to as Geodesic Interpolating Splines [95]. An extension to combine them with affine transformations was given in [96] while two ways to calculate them were presented by Mills *et al.* [97].

Another family of RBFs that has global support is the multi-quadratics [98]

$$\phi(r) = \left(r^2 + d^2\right)^{\frac{\mu}{2}}. \tag{2.21}$$

d is a parameter that controls the smoothness of the deformation and may vary for different points allowing for an adaptive smoothness based on their spatial distribution. μ a non-negative scalar. The previous approach has been extended in the case that a rigid object is present by Little *et al.* [99].

Gaussian functions is another class of RBFs that can be used to parametrize the deformation [100]. They are defined as:

$$\phi(r) = e^{-\frac{r^2}{\sigma^2}}, \tag{2.22}$$

Their advantage is that, despite having a global support, their spatial influence may be controlled by choosing appropriately the Gaussian kernel σ. By choosing a small size for the Gaussian kernel, their influence can be restricted greatly and thus local displacements may be recovered. A recent example of the use of this deformation model can be found in [101], where it has been used in the brain registration problem.

An approximative method to create a dense deformation field from a set of sparse displacements is by calculating their weighted average [85]. The weights at each interpolation point for all control points sum to one. Their support is global, though by choosing appropriately the function that determines the weights, the locality of their support can be adapted.

In medical imaging it is pretty often the case that we search to recover local deformations. In such cases, the previous functions that have global support, and thus all landmarks influence the estimation of the transformation, are not well suited. To be able to treat successfully such cases, interpolation methods where control points have spatially limited influence in the image domain are appropriate.

Fornefett *et al.* [102] investigated the use of Wendland functions [103, 104], that exhibit the desired locality property, for elastic registration.

Other local support radial basis functions are the C^2 smooth Wu functions [105] or the functions proposed by Buhmann [106]. Rohde *et al.* [107] applied the Wu functions in image registration and derived bounds for the basis function's coefficients so that the Jacobian of the computed transformation remains positive.

More recently, a new radial basis function was proposed in [108]. It is defined by using the cosine function and contrary to what claimed in the paper, it is not positive

definite [109]. Lowitzsch [110] introduced a special class of matrix valued RBFs that are divergence free.

In [86], the previous locally constrained radial basis functions were compared using transformations on random point sets, artificial images, and medical images. In most cases, in the presence of both local and global deformation, the locally constrained TPS and Buhmann functions were found to perform better than the rest.

Elastic Body Splines

Splines, though mainly inspired by interpolation and approximation theory, may also be inspired by physical models. Such is the case of Elastic Body Splines (EBS) that were introduced by Davis *et al.* [111]. As the name implies, they are solutions of the Navier-Cauchy equilibrium equation for an homogeneous isotropic elastic body subjected to forces. In the case that the applied forces are either polynomial or rational one can solve analytically the equation.

The previous work was extended in [112] by considering a different type of forces acting upon the elastic body. Instead of assuming forces given as a polynomial or rational function, a Gaussian function was considered. As a result, the transformation model can cope better with local deformations. Moreover, the size of the kernel of the Gaussian can be used to parametrize the compactness of the model's support. Again, the authors were able to obtain an analytic solution of the equilibrium equation.

Wörz and Rohr further extended Gaussian Elastic Body Splines in [113]. Instead of opting for an exact interpolation, an approximation strategy was employed to account for errors in the landmark displacements. The PDE was extended to incorporate Gaussian forces that were weighted by the localization uncertainty. The uncertainties, depending on their isotropic or anisotropic nature, were represented as either scalar weights or matrices. An analytic solution was obtained for the extended equation. The error-aware method performed better than the two previous ones.

Free Form Deformations

Free-form Deformations (FFDs) is one of the most common transformation models in medical imaging. A rectangular grid $G = K \times L$ is superimposed upon the image (size $M \times N$, $K \ll M,\ L \ll N$) that gets deformed under the influence of the control points. The dense deformation is given as a summation of tensor products of univariate splines. FFDs were first popularized in the computer graphics community [114, 115] but gained vast acceptance in the medical imaging community when coupled with cubic-B splines [116, 117, 118, 119].

The deformation in $2D$ is given as:

$$\mathbf{u}(\mathbf{x}) = \sum_{k=1}^{K} \sum_{l=1}^{L} B_k(\mu) B_l(\nu) \mathbf{d}_{kl}, \tag{2.23}$$

where $\mathbf{d}$ denotes displacement, $k = \lfloor x/\delta_x \rfloor$, $l = \lfloor y/\delta_y \rfloor$, $\mu = x/\delta_x - \lfloor x/\delta_x \rfloor$, and $\nu = y/\delta_y - \lfloor y/\delta_y \rfloor$. B_l represents the lth basis function of the B-spline and $\delta_x = \frac{M}{K-1}, \delta_y = \frac{N}{L-1}$ denote the control point spacing. The advantage of the transformation model lies in its simplicity, smoothness, efficiency and ability to describe with few degrees of freedom local deformations.

While in general the transformations that result from cubic B-spline FFDs are smooth, the preservation of topology is not guaranteed. Rueckert *et al.* [120] imposed the hard constrains proven in [121] to produce diffeomorphic deformation fields. The required condition is that the maximum displacement should not be greater than 0.4 the grid spacing.

Many extensions of FFDs have been proposed in the literature. For example in [122], an inverse consistent method based on FFDs is proposed that does not need the inversion of the deformation field. Moreover, the fact that the gradient and Hessian can be computed analytically in this framework can be exploited by the optimizer to achieve faster and more accurate convergence. While, FFDs are usually uniform, non-uniform rational B-splines have been used in medical image registration in an adaptive focus approach [123].

FFDs have been extended to treat multiple image registration where hard constraints are employed to define a reference domain [124, 1, 125]. Last but not least, the transformation model has been extended to the spatio-temporal domain [126, 127, 128] where B-splines are also used for the temporal axis.

Basis functions from signal representation

Inspired by the mathematical tools that are available to represent and analyze signals, many researchers have used Fourier and Wavelet analysis to model the transformation. Maybe the most important reason to use them is the fact that they can provide naturally a multi-resolution decomposition of the displacement field. A necessary property for the coarse-to-fine schemes that are commonly applied in medical image registration to ease the computations and handle large deformations.

Maybe one of the most well known registration algorithms that employs a Fourier based transformation scheme is the consistent registration framework introduced by Christensen [14]. The advantage of the Fourier series representation of the transformation is that it results in a simplification of the used linear elasticity constraint. Moreover, it allows for an efficient numerical implementation.

Fourier basis functions are well localized in the frequency domain. On the contrary, they are not localized at all in the spatial domain. Wavelet basis functions, being localized in both domains, bear the important advantage that they can model local deformations often present in deformable image registration.

A wavelet is a square integrable function $\psi \in \mathbf{L^2}(\mathbb{R})$ with zero mean $\int_{-\inf}^{+\inf} \psi(t)dt = 0$, normalized $\|\psi\| = 1$ and centered at zero. A family of wavelet functions that generate an orthonormal base in $\mathbf{L^2}(\mathbb{R})$ can be obtained by scaling and translation. The construction of this basis is related to the multi-resolution signal approximation that can obtained by an orthogonal projection to a family of nested subspaces $V_{j+1} \subset V_j$, such that $V_{j+1} \oplus W_{j+1} = V_j$, where W_{j+1} is the necessary additional detail. V_j and W_j are created by projecting the function to orthogonal scaling and wavelet functions. For higher dimensions, tensor products of $1D$ wavelet functions are used.

Wu *et al.* [129] used a wavelet-based deformation model. The Cai-Wang wavelet was employed to generate a multi-resolution description in Sobolev space yielding intrinsically smooth deformations. Based on this model, the authors were able to treat global and local information simultaneously in a coarse-to-fine approach. Gefen *et al.* [130] modeled the deformation field by a finite-supported, semi-orthogonal wavelet. The Marquardt-Levenberg optimization scheme was used to minimize a functional that comprises of two terms. The first term penalizes surface distances while the second term regularizes the deformation according to a linear elasticity model.

Musse *et al.* [131] presented a topology preserving multi-resolution approach. Non-orthogonal Riesz basis of polynomial splines were used due to their compactness. The topology was preserved by controlling the Jacobian through hard linear constraints. The approach was extended to the $3D$ domain in [132] and was further validated in [133]. In the $3D$ case, the same multi-resolution framework was used, though no linear constraints can be derived to preserve the topology. To guarantee the desired property, a constrained optimization problem, where the Jacobian is enclosed between two user-defined bounds, was solved. Optimization was performed with a block-wise descent scheme. Cathier [134] used the same wavelet basis as in [129] to decompose the transformation in a multi-resolution fashion. An L_1 penalty on the wavelet coefficients was used to regularize the registration problem. This regularization leads to sparse transformations with respect to the wavelet basis and thus facilitates their storage in the disk.

Piecewise Affine

One of the simplest ways to deform an image is to consider a piecewise linear deformation model. The image is mosaicked by a set of triangles or tetrahedra depending on its dimension. The deformation is parametrized by the nodes of the mesh and an affine interpolation

takes place inside each region. Efficiency and invertibility are the main strengths of this method, while lack of smoothness is its main limitation.

Some of the most recent approaches where a piecewise affine model was used include, but are not limited to, the following. Hellier *et al.* [135] proposed a multi-resolution and multi-grid approach. The image was partitioned adaptively to cubes and an affine transformation was inferred for each one. A regularization energy term encouraged neighboring pairs to deform similarly. Zhang *et al.* A similar approach was employed in [136] to tackle diffusion tensor registration by taking into consideration tensor reorientation. The images were separated into contiguous blocks and an affine transformation was recovered for each one of them. Regularization on the interface of regions ensured the global smoothness of the transformation.

Pitiot *et al.* [137] reconstructed $3D$ volumes of biological images by employing a piecewise affine transformation model. The images were separated into independent components through hierarchical clustering. In a subsequent step, affine registration was performed for each pair of regions. The final transformation was estimated by considering the affine transformation for each region and applying a non-linear interpolation in-between the regions. A similar approach was presented in [138] with the difference that a regularization step followed to improve the smoothness in the interpolated areas. The final transformation was composed in such a way that its invertibility was ensured.

Two more recent application of piecewise affine model were presented in [139, 140]. Cootes [139] favored the use of piecewise affine transformations as they can be easily inverted. Buerger *et al.* proposed a hierarchical framework to separate adaptively the images in regions. Splitting was formulated as an energy minimization problem.

Most of the approaches that employ piecewise linear strategies, consider the affine transformations independently. As a result, singularities may occur and the transformation is not globally invertible. To account for the previous drawback, sophisticated methods have been introduced. A transformation model that is affine at the center of a region and reduces to identity as the distance from the center increases was introduced in [141]. The novel transformation model has a closed form and can be computed efficiently. Moreover, constraints were given in the form of bounds on the translation so that invertibility is ensured.

Arsigny *et al.* [142] introduced a poly-rigid/affine transformation model. Based on a finite set of points whose transformation is known, the global transformation at each point is given by integrating over time a distance-weighted sum of infinitesimal velocities at the known points. No closed form exists but a computationally more expensive integration of ODEs is necessary. [143] extended the poly-affine transformation so that its inverse is also poly-affine. Moreover, the fusion of affine transformations was rendered invariant to affine changes of the coordinate system.

2.2.3 Constraints

Image registration is principally an ill-posed problem as the number of unknowns is greater than the number of available constraints. In order to account for that, regularization is necessary. Moreover, regularization allows us to introduce any prior-knowledge we may have for the physical properties of the underlying anatomical structure and helps optimization to avoid local minima.

There are two possible ways to regularize the problem, either explicitly or implicitly. Explicit regularization may be achieved through the addition of a smoothing term in the energy that penalizes non-regular configurations. This is mostly the case when physically motivated models are used. Implicit regularization may be achieved by parameterizing the deformation field with smooth functions. That does not exclude the use of explicit regularization but that may be used now to achieve complementary goals, usually tailored for the specific problem at hand.

One of the most important properties that a registration algorithm should exhibit is the preservation of topology. Apart from the framework of flows of diffeomorphisms that delivers naturally such results, the rest of deformation models don not exhibit these properties. In that case, smoothness and invertibility of the resulting deformation field may be guaranteed through the use of constraints. The Jacobian of the deformation field is very informative regarding the local properties of the deformation field. Thus, by tracking its values, singularities may be avoided by creating intermediate templates and reinitializing the process [19].

This technique is efficient but the way the preservation of topology is guaranteed seems rather awkward. A more elegant way would be to incorporate an appropriate term acting upon the Jacobian in the objective function. In [14], a term that penalizes small and large Jacobian values for both the forward and backward transformation was added to the objective function. Similarly, Droske and Rumpf [144] used a regularization term that considers the length, area and volume deformation. This approach has the disadvantage that it depends greatly upon the way the different energy terms are weighted. If the weight for the regularization term is not important enough, singularities may appear. Whereas, if the weight for the regularization is too great, the optimization may be hindered.

A different and probably more appropriate strategy is to cast the problem as a constrained optimization problem. In [131], linear inequality constraints were derived so that the topology is preserved. The optimization was solved by employing a fast method that bears a resemblance to sequential linear programming. In its $3D$ variant [132], the energy was optimized under the constraint that the Jacobian will stay in-between user specified bounds. Interval analysis techniques were used in order to solve the optimization problem. Inequality constraints were also used by Haber and Modersitzki [145]. A variant of a log-

barrier method was used to solve the optimization problem. Instead of solving the initial constrained problem, a sequence of unconstrained ones was considered. The weight for the barrier terms increased gradually for each unconstrained problem that was optimized by applying a variant of the Gauss-Newton's method.

Sdika [119] also proposed a constrained optimization framework to ensure that the transformation, parametrized by cubic B-splines, is invertible. Two constraints were proposed. The first constrains the Jacobian of every pixel to be greater than a threshold. As this constraint does not control the value of the Jacobian between the voxels, a second constraint was proposed that relates the Jacobian with its derivative. In that way, the Jacobian is restricted to be within a range of values. Moreover, when approaching values close to the bounds, its derivatives are constrained to be close to zero. In [146], a simpler penalty was devised for the case of B-splines. The penalty takes into consideration the difference between two adjacent nodes and is memory efficient.

In many applications, on top of ensuring that the topology will be preserved, we are also interested in volume preservation. Such a constraint is of particular interest in cases that we know that the imaged anatomical structure is not compressible and all changes are due to either motion or intensity changes provoked by the action of a contrast agent. The simplest case would be to consider a rigid part of the body such as bone structures. More complicated cases would include deformable structures that preserve their volume such as breast, myocardium and liver.

In [147], such a strategy was employed to register contrast-enhanced MR breast images. Along with the image matching term a second term that penalizes volume changes was considered. The penalty integrates the absolute logarithm of the Jacobian determinant and is zero only when local volume is preserved. A sequential approach was proposed in [148]. First, a standard registration was performed. Based on its result, areas whose volume should be preserved were identified. Once found, the displacement of the control points of the FFD model that influence these areas were fixed to the mean of their previous ones and the registration was solved again for the rest of the variables. A constrained optimization approach to preserve volume was presented by Haber *et al.* [149]. An energy comprising a matching and regularization term was minimized under the constraint that the determinant of the transformation is equal to one ($det(\mathbf{I} + \nabla\mathbf{u}) - 1 = 0$). Staggered grids were used to discretize the problem and Sequential Quadratic Programming to solve it.

Another type of problems that call for an incompressibility constraint are those where the deformation of the myocardium is studied as it is known to be a nearly incompressible material. Bistoquet *et al.* [150] approximated the previous constraint by $\nabla \cdot \mathbf{u} = 0$. This constraint was enforced by the use of divergence-free radial basis functions as deformation model [110]. In addition, a hard constraint was introduced in the objective function to pe-

nalize deviations from incompressibility. The determinant of the Jacobian was constrained to be close to one in a predefined region by using Lagrange multipliers in [151]. A different approach was taken by Mansi *et al.* [36] where the velocity field $\mathbf{v}$ was constrained to be divergence-free. The method was based upon the fact that the integration over time of divergence-free velocities results in incompressible deformations.

Last but not least, as in medical images rigid structures are present, it would be beneficial for the quality of the registration result to incorporate rigidity constraints in order to treat them according to its physical properties. Loeckx *et al.* [152] constrained locally a non-rigid FFD registration method by penalizing deviations of the Jacobian from orthogonality. In [153] rigidity was imposed by introducing three conditions. The first condition enforces the second derivatives of the transformation to be zero. The second condition enforces the orthonormality of the rotation matrix while the third condition imposes the determinant of the Jacobian to be equal to one. Local rigidity has also been considered in a variational setting by Modersitzki [154]. Along with the matching and the regularization term, a third one was introduced in the objective function. The third term controls the rigidity of the transformation by enforcing its Jacobian to be linear, orthogonal and orientation preserving. A Gauss-Newton optimizer was used.

2.3 Objective Functions

Images can be aligned either by evaluating a criterion based on intensities over the whole image domain (iconic methods) or by establishing point correspondences (geometric methods).

2.3.1 Iconic Methods

Iconic registration methods offer accuracy by providing dense correspondences between the considered image domains at the cost of considerable computational expense. Due to the fact that all image points are considered equally, salient points might fail to get the importance they deserve. In addition, initial conditions influence greatly the quality of the obtained result. In the presence of large deformations, typical in longitudinal and population studies, the quality of the solution is often degraded.

The matching term integrates the evaluation of a dissimilarity criterion over all image elements:

$$\mathcal{M} = \int_{\Omega_T} \rho(S \circ \mathcal{T}(\mathbf{x}), T(\mathbf{x}))d\mathbf{x}. \tag{2.24}$$

Devising an appropriate criterion ρ is an important and difficult task as the criterion should be able to account for different physical principles behind the acquisition of the two images

and thus for the intensity relation between them. Moreover, the properties of the objective function (*e.g.* its convexity) may influence the difficulty of the inference and thus the quality of the obtained result. An ideal criterion would take low values, when points belonging to the same tissue class are considered and great values, when points from different tissue classes are compared. Moreover, it should be convex allowing for accurate inference.

At this point, two cases should be discerned: i) the mono-modal case, considering images from one modality, and ii) the multi-modal case, considering images from multiple modalities.

Mono-modal registration

The mono-modal case is the easiest one and historically the first to be studied. Since the same imaging device is used to image both volumes it is often assumed that same anatomical structures correspond to similar intensity values. This assumption leads naturally to the use of Sum of Squared or Absolute Differences (SSD and SAD respectively) as a matching criterion. The choice between the two depends on the assumption on the noise that corrupts the image intensities. In a more sophisticated setting when a linear relation is assumed between the signal intensities, the optimal criterion is Cross Correlation (CCor) and Correlation Coefficient (CCoef) [5, 155, 61].

Intensity information by itself is considered a poor feature that often leads to ambiguous matching and is one of the main reasons for the presence of local minima in the objective function. Trying to minimize the effect of these local minima and establish more accurate correspondences, a number of researchers have proposed to increase the dimensionality of the feature space by introducing local information through the use of attributes that represent the geometric structure of the underlying anatomy. These approaches are often referred to as feature- or attribute-based approaches.

In their seminal work, Shen and Davatzikos [101] proposed the use of an attribute vector including Geometric Moment Invariants in an attempt to capture local anatomical information at different spatial scales. The motivation is that a rich enough attribute vector will be able to differentiate between voxels that would be considered the same based only on their intensity information. Thus, fewer local minima will be present and better accuracy may be achieved. To further reduce the effect of the local minima, they proposed a hierarchical scheme that approximates successively the objective function through the use of an increasing number of voxels where the matching is evaluated.

The success and importance of the previous method becomes evident in view of the number of approaches that have been proposed to improve its performance. One drawback of the previous method is that it requires a pre-segmentation step in order to intro-

duce local spatial information. To remove the previous requirement, the use of Daubechies wavelets to populate the attribute vector was proposed in [156]. The attribute vector was constructed in a multiscale fashion and is translation and rotation invariant. It was shown that the wavelet-based attribute vector is more discriminative when compared to the geometric moments. Another approach to tackle the previous shortcoming is by using local histograms and boundary information as attributes [157]. In [158], a learning approach was proposed to improve the result in two ways. First, the optimal scale for the geometric features for each voxel was determined leading to increased discriminative power. The hierarchical scheme was improved by deciding upon which voxels drive the registration process based on their saliency and the consistency of their description across the training data.

Another way to incorporate local information using attribute vectors is through the use of local frequency representations being the responses of Gabor filters [159, 160]. Gabor features have proven successful for both mono-modal and multi-modal image registration as they are able to capture information across different scales and orientations. Ou *et al.* [159] optimized the Gabor features to be more distinctive and introduced the notion of mutual saliency to let the most reliable points drive the registration process. Liao and Chung [161] however argued that frequency spectrums of MRI brain images often exhibit non-Gaussian behavior and thus the choice of Gabor filters is not optimal. They proposed the use of symmetric alpha stable filters and showed experimentally that they outperform Gabor features in non-rigid MRI brain registration. A new feature for non-rigid registration, named uniform spherical region descriptor, was proposed in [162]. It is invariant with respect to rotation as well as monotonic gray-level transformation and thus is able to account for the presence of bias field.

Multi-modal registration

Multi-modal registration is more challenging as the choice of an appropriate objective function is harder a task. Two main approaches have been proposed to solve the problem:

1. Use information theoretic measures.
2. Reduction of multi-modal problem to a mono-modal problem by:
   (a) simulating one modality from another, or
   (b) mapping both modalities to a common domain.

Here, we are going to focus primarily on information theoretic approaches as they constitute the most frequently used way to tackle the challenges posed by multi-modal registration. Reduction techniques will also be briefly discussed.

Information theoretic approaches: Information theoretic approaches were popularized by Viola, Wells, Colignon and Maes [163, 164, 165, 166]. They proposed to use as objective criterion, the mutual information (MI):

$$I(XY) = H(X) + H(Y) - H(X,Y). \tag{2.25}$$

That is, the MI between two random variables X and Y is defined as the difference of their marginal entropies $H(X)$ and $H(Y)$ with their joint entropy $H(X,Y)$. The difference between the two approaches is the way entropy is estimated. In [163, 164] a non-parametric estimator was used while in [165, 166] histograms were applied. The advantage of MI is its generality since it does not assume any relationship between the image intensities. For a survey on MI-based registration methods, the interested reader is referred to [167].

The widespread use and study of MI has revealed some of its shortcomings. One of the main shortcomings is that it is not overlap invariant. Thus, in certain cases it may be possible for mutual information to be maximized when the images get misaligned. To remedy that, a Normalized version of Mutual Information (NMI) was proposed in [168]. Recently, Cahill and co-workers elaborated upon the idea of overlap invariance and showed that neither NMI, MI, CR, CCor nor CCoef are invariant to changes of overlap and introduced appropriate invariant versions of the previous similarity measures in [169].

The success of MI paved the way for the introduction of an important number of statistical criteria in image registration. Roche *et al.*. [170] argued that the generality of mutual information can be a drawback when a reasonable hypothesis can be made regarding the relationship between the intensities. They proposed to use the Correlation Ratio (CR) as appropriate similarity measure when the assumption of functional dependence between the image intensities is valid.

Pluim *et al.*. [171] compared the performance of a number of f-information measures in medical image registration. Given two probability distributions P and Q, the quantification of their difference can be given by a measure of divergence. An f-divergence measure is defined as:

$$f(P\|Q) = \sum_i q_i f(\frac{p_i}{q_i}). \tag{2.26}$$

For different functions f, different divergence measures may be defined. In the context of registration, what is measured is the difference of the joint distribution of the intensities with respect to the joint distribution that would arise when images are independent.

The idea to use divergence measures to compare joint intensity distributions has attracted significant attention and numerous divergence measures have been proposed for multi-modal image registration. Kullback-Leibler divergence (KLD) was used in [172, 173] to register multi-modal images. The joint intensity distribution is either learned from aligned pairs of images or by segmenting corresponding anatomical structures. Images get

aligned by minimizing the divergence between the observed and estimated distributions. In a similar setting [174], in the sense that learned distributions are used, Jensen-Shannon Divergence (JSD) was used and shown to be better that KLD being symmetric, bounded and true metric.

Another family of information theoretic approaches is built upon Renyi entropy (RE),

$$RE_\alpha(P) = \frac{1}{1-\alpha} \log \sum_i p_i^\alpha,\ \alpha > 0\ and\ \alpha \neq 1. \tag{2.27}$$

Based on this entropy, the Jensen-Renyi divergence can be defined. Its advantage lies in the fact that it is symmetric and generalizable to any finite number of probability distributions. The Jensen-Renyi divergence is convex for $\alpha \in (0, 1)$ and is maximum when the distributions are degenerate. Its use for image registration was proposed in [175, 176]. In [177, 178], a Minimum Spanning Tree (MST) was used to estimate the RE and the optimal transformation was estimated by minimizing it. Spanning graphs were also used by Sabuncu and Ramadge [179]. The latter demonstrated the superiority of the entropic graph estimator with respect to standard estimators. A generalization of KLD was introduced in [180]. The new divergence measure is based on modified Bessel functions of the second kind and allows for an efficient recursive computation. The generalization of KLD was shown to perform better than the standard measures of divergence.

Most of the previous approaches share a common drawback; they are based upon a single pixel joint probability model. As a consequence, by changing the positions of the pixels in a random way and evaluating the statistical criterion, the same similarity is obtained [181]. To rectify the previous shortcoming, local context should be introduced in the used criterion.

One way to relax the global way the statistical criteria are considered is by means of computing them locally and thus cope with the fact that the relation between the intensities of the two images is non stationary. This approach was investigated by Hermosillo *et al.* and Karaçali. Hermossilo *et al.* [182] derived the Euler-Lagrange equations for MI, CR and CCoef based on locally estimated probability distribution functions. Karaçali [183] followed a deterministic rationale to express in closed form mutual information, joint entropy and the sum of marginal entropies over small spherical regions.

Local evaluation of mutual information has also been proposed by other researchers. For instance, [184] introduced Regional Mutual Information (RMI). The proposed objective function is a linear weighted sum of local evaluations of MI and aims to reduce the error caused by local intensity changes. Loeckx [185] *et al.* proposed to condition the evaluation of MI upon the position. More recently, locally evaluated MI in combination with standard global MI was used in [186]. Moreover, the local evaluation of the probability distribution function considers pixels relatively to their distance with respect to the

FFD control points.

An alternative way to introduce local context is by inserting spatial information. This has been mainly achieved by considering additional features that capture local geometric information resulting in higher order entropic measures.

In one of the first attempts to exploit spatial information Pluim *et al.* [187] used intensity image gradient as an additional cue. The proposed algorithm sought not only to maximize NMI but also intensity gradient information. This was simply achieved by multiplying NMI with a measure that takes into consideration both the intensity gradient magnitude and its orientation in an effort to encourage the alignment of strong intensity gradients. Intensity gradient information drove the registration to more accurate results but most importantly rendered it more robust.

Approximately at the same time, Rueckert *et al.* [181] proposed to use second-order MI encoding local information by considering co-occurrences of intensities between neighboring voxels. That approach requires a $4D$-histogram to estimate the information measures. To account for the high dimension of the histogram and the curse of dimensionality, the number of bins was reasonably small. More robust and accurate registration with respect to standard MI was obtained.

Russakoff *et al.* [188] introduced the Regional Mutual Information that pushed forward the previous idea by considering co-occurrences between regions. Moreover, an efficient way to deal with the curse of dimensionality was presented. Assuming a high-dimensional distribution, the data points were transformed so that they were independent in each dimension. Then, the entropy was estimated by summing the distributed $1D$ entropies. In [189], NMI between blocks of image elements was studied. The high-dimensional NMI was estimated by using random lines and reducing the number of bins.

Instead of taking into account explicitly neighboring voxels, a more compact way to consider local information would be by extracting features that describe concisely regional characteristics. In [190], Gaussian scale space derivatives were employed and considered as an additional information channel in a higher dimensional MI criterion. On top of the intensity, one more channel of information was also used in [191] resulting in a multi-dimensional NMI. A novel spatial field, named Maximum Distance-Gradient (MDG), was introduced. First, a set of special points located in important gradient areas was created. Based on that, a vector field that contained both local and global information was created. For every voxel, in addition to the local gradient, the distance, the direction and the intensity difference with respect to its source was retained. The magnitude of the MDG vector field formed the supplementary channel while its orientation was used as a second element in the objective function.

The curse of dimensionality is an important limitation of the previous approaches as it hinders the evaluation of higher dimensional statistical criteria. To be able to handle

such calculations, most researchers resort to crude implementation approximations such as limiting the number of histogram bins. Nevertheless, ways to estimate high dimensional entropies have been proposed and used to perform image registration.

In [192], spatial information through the construction of feature vectors was adopted. The resulting high dimensional entropy was estimated with the use of the MST estimator. Entropic graphs were also used in [193] to tackle high dimensional α-MI registration of ultrasound images. Both previous approached coped with global linear registration. Most recently, deformable registration of Cervical MRI using high-dimensional MI was presented in [194]. Features were used to describe local geometric information and a k-nearest neighbor graph to estimate the multi-dimensional MI.

Spatial information is not the only type of information that can be used to endow registration with increased robustness and accuracy. Assuming that a prior step of segmentation has been performed, tissue classification information may also help disambiguate between voxels that belong to different tissues but share common appearance properties.

In [195], regions were segmented by thresholding and connected component labeling. The labels were used as an additional channel and the objective function considered the difference between the three entropies and the joint entropy.

Knops *et al.* [196] performed a k-means clustering before registration. Based on this clustering, voxels that share similar intensity profiles but belong to different anatomical structures were mapped to different intensity bins during the construction of the histogram. The new remapped intensities along with the initial one contributed to an NMI based similarity criterion. The new objective function performed better than the standard MI approach.

Voxel class probabilities were considered in the objective criterion by D'Agostino *et al.* [197] for the labeled-to-labeled and intensity-to-labeled image registration. For the labeled-to-labeled case, KLD was used to compare the distribution of the joint classes whereas for the intensity-to-label registration, a version of MI was used with the difference that one of the features is a class probability and not intensity.

Reduction to mono-modal registration: An alternative way to proceed with multi-modal registration is by reducing the problem to a mono-modal one where the solution can be obtained in a simpler and more accurate way. There are two possible ways to perform such a task, either to simulate one modality from another so that at the end both images come from the same modality, or to map them both to a third domain where the registration will take place.

Simulating one modality from another can be achieved by taking advantage of the knowledge that is available with respect to the physical properties of the imaging device and trying to model the imaging process. An alternative way is to exploit available co-

registered pairs of images and use machine learning techniques to capture the relation between the intensities.

The first approach, that is task specific, was first proposed by Roche *et al.* [198] towards registering rigidly US images to MR ones. In [198], US intensities are predicted by exploiting MR intensities and MR gradient magnitude information. Complex phenomena such as US signal attenuation and speckle are neglected, leading to images that only roughly resemble to actual US images. A more sophisticated model was proposed in [199] to tackle the problem of CT-to-US registration problem. Based on the physical principles of ultrasound, the authors were able to simulate an US image that was then used along with a locally evaluated statistical criterion to drive the registration.

In [200], mixture of experts were used to learn the conditional probability of the target intensity given a source patch. The conditional probability was then used to drive a Markov Random Field to regularize the simulated image. An SSD criterion, evaluated by considering the simulated and target image, was shown to outperform MI.

The most common approach out of the two is to map both modalities to a common space. As both modalities image the same anatomical structure, it is logical to assume that the local geometry would be helpful to establish meaningful correspondences. Thus, in principle, most methods apply filters that extract geometrical information and then use it in a mono-modal registration setting.

For instance, Maintz *et al.* [201] used morphological tools to create new gray-value intensity images. Their method basically uses morphological opening and closing to extract edge information and then cross-correlation to align the images. It resembles to a surface registration with the difference that instead of having binary values, real ones are used. Haber *et al.* [202] assumed that borders of anatomical structures correspond to intensity changes and thus opted to exploit intensity gradient information. An intermediate image domain was created by considering the normalized intensity gradient field that conveys purely geometric information and accounts for the fact that the gradient magnitude may vary among different modalities. As objective criterion, the difference in angles between the normalized gradient vectors was considered. In [144], following the mathematical morphology theory that states that an image can be characterized uniquely by the entity of its level sets, the Gauss map of the images was considered as common space. The registration was formulated in a variational framework where a morphological, contrast invariant, matching criterion was minimized under the influence of an appropriate regularization term.

Butz *et al.* [203] also experimented with edge related information. They used an edgeness operator that considers the local edge variance to map both images to a common space. Mutual information driven registration was then performed coupled with a multi-scale genetic optimization. Depending on the nature of the images, other operators may

be applied. In [204], the probability of vessel presence was used along with normalized cross-correlation to register MRI with ultrasound images.

Richer descriptions of local structure based on Gabor filtering can also be used to perform registration [160, 159]. In [160], local frequency that is robust to edge strength and contrast differences was used. It was estimated by calculating the local phase gradient of the most significant Gabor filter response. Then, the integral squared error was chosen as matching criterion. The responses of Gabor filters were used in [159] to construct a rich vector descriptor. The images were aligned by minimizing a weighted sum of the vector differences. In [205], the authors took a different direction to estimate the local frequency information. Instead of using Gabor filters, the Riesz transform was used.

In [206], a pseudo-modality was created by remapping intensity values based on the conditional distributions of the intensities. This step was performed in order to apply simpler criteria than MI when the size of patches is small and its estimation compromised. Recently, Wachinger *et al.* [207] proposed two techniques that derive from information theory and manifold learning to create the intermediate structural representation. The first one considered the entropy of a patch centered around the voxel to assign a new intensity value. The second method used Laplacian Eigenmaps to embed the patches in a lower-dimensional manifold that preserves local distances.

A supervised technique was presented in [208] to learn the similarity measure for multi-modal image registration. The approach was formulated in a discriminative setting where the goal is to optimize a similarity function so that correct correspondences are assigned high values and erroneous ones low. Support vector machine regression was employed to learn the metric. The optimal metric performed better than the standard NMI.

Another supervised technique was presented in [209, 210] to learn a similarity metric that discerns between corresponding and non-corresponding points. This technique maps both modalities to a Hamming metric space where true correspondences are more probable to have the same code while wrong ones are not likely to. The embedding was constructed by using AdaBoost. The proposed method was proven experimentally to outperform CR, MI and NMI.

It should also be noted that some of the techniques that were previously presented under the *information theoretic* class of methods learn a similarity measure. The difference is that a generative framework is employed. Given co-registered data, the joint distribution of the intensities is learned. Then, either a maximum likelihood approach [211] or a divergence criterion [172, 173, 174, 179] is used to compare the estimated and learned distribution.

2.3.2 Geometric Methods

Geometric registration establishes sparse correspondences between a subset of the image voxels. The voxels are placed in salient image locations considered to correspond to meaningful anatomical locations. The underlying assumption is that saliency in the image level is equivalent to anatomical regions of interest. Geometric registration is able to overcome some of the limitations of the iconic registration being robust with respect to the initial conditions and the existence of large deformations. The solution of the registration problem is obtained in a relative straightforward way once landmarks have been extracted. However, locating reliable landmarks is an open problem and active field of research. Most importantly, the sparse set of directly obtained correspondences gives rise to the need for interpolation which results in decrease of the accuracy as the distance from the landmarks increases. This, as well as the advance of technology permitting to meet both the memory and computational demands of iconic registration have boosted research towards that direction stripping off the geometric methods from the interest they enjoyed during previous decades. Nevertheless, for some task specific applications geometric registration is the most reliable choice.

Detecting points of interest

The first step in geometric registration is to detect points of interest. Images that contain sufficient details facilitate point detection. Medical images are not as rich in details as natural images. That is why, point detection has mainly drawn the interest of the computer vision community and thus is well-studied in the case of $2D$ images and not in the case of $3D$ images mostly found in medical imaging.

A full overview of the point detectors that have been proposed for computer vision related problems is out of the scope of this review. Nonetheless, let us give a brief description of the most important ones. Maybe the most well-known, point of interest detector is the Harris corner detector proposed in [212]. Harris *et al.* proposed to identify corners by exploiting the information conveyed by the structure tensor A. Specifically, points of interest are determined by considering the following quantity: $det(A) - \alpha\, Tr(A)^2$. In similar lines, Shi and Tomasi [213] proposed the use of the minimal eigenvalue of the structure tensor in order to tract points of interest.

Many extensions to the Harris detector have been proposed in the literature. Their main aim is to impose a certain invariance. One may cite the approach proposed in [214] and affine-invariant Harris and Hessian [215]. Affine invariance is important as it enables the detection of points under affine transformations and a lot of efforts have been concentrated in defining such detectors. An evaluation study comparing the most important methods was presented by Mikolajczyk *et al.* [216]. A historical review of point detection methods

can be found in [217, 214], while evaluation studies of point and corner detectors can be found in [217] and [218], respectively.

An alternative way to determine point of interest is by performing scale-space analysis and detecting blob-like regions. One of the first methods to perform such a task is the Laplacian of Gaussian. The image is convolved with different scales of a Gaussian kernel and at each level the Laplacian operator is applied. By tracking across scales the local maxima/minima of its response, key-points are detected [219]. Lowe [220] proposed to use the Difference of Gaussians that is an approximation of the Laplacian. From this scale-space representation, local minima/maxima are extracted in order to detect feature points and the local Hessian information is used to reject spurious ones. Lowe's Scale Invariant Feature transform has been proven extremely successful and has been extended in the $3D$ domain so as to be applied in medical imaging in [221]. A more recent technique for blob detection was proposed by Matas *et al.* [222]. Image regions, created by thresholding in the intensity domain, are tracked and selected based on their area's stability as the threshold value varies.

In medical imaging, such generic approaches have not been explored. On the contrary, feature detection is performed in a task specific manner, usually as product of a segmentation preprocessing step. In brain image registration, sulci information has been used in [223, 73, 224, 225]. The cortical surface information has also served as feature in [225, 226, 227].

Another case where geometric registration has been proven successful is the retina image registration. In this case, intensities are homogeneous while important information is conveyed by the vasculature. In [228] branching and crossover points of the blood vessel structure were used as feature points. While Stewart *et al.* [229] used in addition the centerlines of the segmented vasculature. For each centerline point, its location, tangent direction and width were retained. Vascular structures are also important in brain sift correction [230], pulmonary CT images [231] and liver registration [232]. That is why a number of task-tailored detectors have been devised [233, 234, 235, 236].

In a more general setting, points of interest can be found based on the maximal response of Gabor features weighted by a mutual saliency criterion in [237]. Lastly, markers that are often used in medical imaging such as fiducial ones can be used to guide image registration. Some resent studies regarding the errors in the process are given in [238, 239, 240].

Following the previous, two sets of points ($K = \{\kappa_1, \cdots, \kappa_n\}$ and $\Lambda = \{\lambda_1, \cdots, \lambda_m\}$, $m \geq n$) are created. The first contains points belonging to the source domain Ω_S, while the second points that belong to the target one Ω_T. These constitute the set of known variables. The set of unknown variables usually comprises: i) the correspondence, and ii) the transformation. Based on how these variables are treated, three different classes of

methods can be discerned:

- Methods that infer only the spatial transformation.
- Methods that infer only the correspondence.
- Methods that infer both variables.

Methods that infer only the spatial transformation

Two categories of methods should be considered. The first one assumes that the correspondences are known in an exact or inexact way. This problem is known as *exact or inexact landmark matching*. In the exact case, a smooth transformation is sought so that the correspondence are respected exactly or a regularization energy is optimized under correspondence constraints. In the inexact case, a compromise between matching and smoothing the deformation is preferred.

Procrustes analysis is a popular method for shape analysis and is useful when homologies between point-sets are given [241, 242, 243, 139]. In Procrustes analysis, a least-squares distance is minimized. Given the correspondences, an analytical solution, that consists of translating, rotating and scaling [242], exists. In the affine case, the solution is usually given by numerical optimization [244]. Recently, an algebraic way to solve for the affine registration of planar sets based on complex numbers was presented in [245].

During the last decade, it became possible to estimate non-rigid transformation based on point correspondences. As we saw in the previous section (Sec. 2.2.2), radial basis functions are able to produce dense deformation fields for any spatial distribution of points. Moreover, inspired by approximation theory, non-interpolating splines are able to account for the uncertainty in the estimated correspondences [90, 113]. What is more, both the exact and inexact landmark matching have been solved for the case of diffeomorphic deformations [246]. This is achieved by introducing a time parameter t and considering the transport equation:

$$\frac{\partial \mathcal{T}(\mathbf{x}, t)}{\partial t} = \mathbf{v}(\mathcal{T}(\mathbf{x}, t), t), \tag{2.28}$$

under the constraint that $\mathcal{T}(\mathbf{p}_i, 1) = \mathbf{q}_i$ or that the given correspondence is respected. The previous method has been extended to the case where the domain is a sphere [65]. This is of interest as surfaces are often mapped to spheres in order to facilitate their study.

The second subclass opts for estimating the transformation without concerning itself with the establishment of correspondences. These methods are more robust to missing correspondences and outliers. One of the most recent methods belonging to this category is the one proposed in [247]. The point sets are transformed into their canonical forms by

considering second order statistics. This results in reducing the estimation of the affine matrix to determining an orthonormal one which is subsequently estimated by using third order moments. Affine transformation of point sets without correspondences can also be achieved by considering the convex hull and Hausdorff distance. For instance, Gope *et al.* [248] proposed to obtain an affine invariant representation of the point sets based on the convex hull. The transformation was optimized based on a variant of the Hausdorff distance that takes into consideration some pairing distance. Another way to perform such a task is to consider robust ways to calculate the distance between the points [249].

Initially, such methods were able to estimate only linear transformations. Recently, methodological advances have permitted the handling of non-rigid transformations. The breakthrough was achieved by representing the point sets as probability distributions and minimizing a distance measure between the two distributions. One of the first attempts in that direction was proposed by Glaunes *et al.* [67]. The authors extended the large diffeomorphic deformation framework in the case of distributions and unlabeled point sets. Point sets were modeled as a weighted sum of Dirac measures and a kernel-based error measure was used. Tsin and Kanade [250] proposed to register point sets based on a measure called *kernel correlation*. The proposed measure is proportional to the correlation of two kernel density estimates. A similar approach based on kernel density correlation was presented in [251].

Gaussian Mixture Models (GMMs) are a common way to model distributions. In [252], each point set was modeled using GMMs and a L_2 distance was used to compare them. In [253], registration was recast as a probability density estimation problem. The points of the first set were considered as the centroids of the GMMs which were fitted to the data (or points of the second set) by likelihood maximization. Special care was taken so that the centroids move in a coherent way. Roy *et al.* [254] modeled each feature of each shape as GMM. A mixture model was used to represent the shape by assuming that features are independent and identically distributed. A closed-form distance between the two distributions was used along with a TPS parametrization of the transformation. A similar model was used in [255] where the problem of the simultaneous registration of multiple point sets is tackled. Jensen-Shannon divergence was used as similarity metric. The drawback of this approach is that no closed-form solution exists. Thus, a computationally and memory demanding estimation based on the law of large numbers is required. In a subsequent work, Wang *et al.* [256] alleviated the previous shortcoming by using the generalized L_2-divergence that allows for a closed-form solution. Lastly, a GMM was also used in [257] with the difference that the Gaussians are not isotropic. The Havrda-Charvat-Tsallis (HCT) divergence was used along with directly manipulated free form deformation.

It should be noted that another way to perform non-rigid registration of shapes and

points without caring to establish correspondences, is to embed them in a higher dimensional space and perform classic image-based registration there [66, 68, 258, 259].

Methods that infer only the correspondences

Inferring the correspondences or matching the features is a task that is inherently coupled with the way the features are described. To better disambiguate between close potential candidates rich descriptors should be used. Moreover, as the image undergoes some sort of deformation, it is to be expected that the appearance of the features will vary between images. To account for this fact, descriptors should be invariant to such changes.

Constructing rich invariant descriptors is an active field of research though mainly oriented towards computer vision applications. We are going to present briefly some of the most important feature descriptors that have been proposed and discuss briefly some of the ones that have been applied in medical imaging. For a comparison of the performance of different feature descriptors, the interested reader is referred to [260, 261].

Lowe in his seminar article [220] proposed a feature descriptor (SIFT) based on the gradient information at the scale a point of interest was detected. For every pixel in a neighborhood of the key-point the gradient magnitude is computed. Its value is weighted depending on its distance from the key-point. From these values, gradient orientation histograms are computed and normalized to account for photometric variations. Many variants of SIFT have been proposed. Among them PCA-SIFT [262], Gradient Location and Orientation Histogram (GLOH) [260] and Speeded-Up Robust Features (SURF) [263]. For a comparison between the original SIFT and its variants see [264]. An affine invariant version of SIFT was proposed by Morel and Guoshen [265]. Invariance is introduced by simulating latitude and longitude angles, the two parameters of the camera for which the original SIFT is not invariant.

A number of descriptors have been applied in medical imaging. The SIFT descriptor for example has been used to match points in a hybrid registration framework in [266]. Gabor features have proven their value in a number of applications [159, 237]. Wavelet filters and Steerable filters [267, 268] have been used in [156] and [269] respectively. The list is not complete but indicative of the importance of the feature descriptors in medical imaging.

Having established a discriminative and ideally deformation invariant description of the key-points, correspondences may be established either by i) relying solely on the closeness of the descriptions; or ii) by incorporating structural constraints.

In the first case, the information contained by the descriptor is used to determine the correspondences. There is an implicit assumption that the descriptors are constructed so that the use of the Euclidean distance is sufficient to rank potential matches. This con-

struction can be achieved by appropriate rescaling of the feature vector values. Based on an established ranking, different matching strategies may be considered. The simplest one is by thresholding. The definition of the threshold can be achieved through ROC analysis. A better strategy would be to assign each point to its closest candidate. As the probability of detecting a false positive is significant, a threshold is still needed to control it. The third strategy is to consider the ratio between the distance with the nearest and the second nearest neighbor in the feature space. For an evaluation of these strategies, the interested reader is referred to [260].

While being intuitive and efficient, the previous approaches discard any information regarding the spatial location of the key-points in the image. The incorporation of such a knowledge aims to constrain better the matching problem and reduce further the number of erroneous correspondences.

A popular way to introduce structural constraints is by formulating the problem as graph matching. Leordeanu *et al.* [270] proposed a spectral technique to solve the matching problem. Pairwise constraints were used to preserve pairwise geometry. Berg *et al.* [271] formulated the problem of recovering feature correspondences as an integer quadratic programming problem. Changes in the length and the direction of vectors defined by pairs of features were penalized. Pairwise constraints were also employed in [272] to model local spatial coherence. Moreover, the authors showed that is possible to handle outliers during the optimization. In medical imaging, [237] presented recently a method to detect mutually-salient pairs in brain images.

Despite the success pairwise constraints have enjoyed in many applications, they are limited with respect to the relations they can model. Recently, a number of researchers have tried to tackle the graph matching problem with higher order constraints. The spectral matching method of [270] was generalized to higher order constraints by Duchenne *et al* [273]. A tensor power iteration method was employed to solve the matching problem. A similar formulation was also proposed by Zass and Shashua in [274] favoring the use of a different optimization method. Wang *et al.* [275] proposed a higher-order graph matching formulation incorporating learned structural constraints in a segmentation framework. The inference was performed by a dual decomposition based method [276].

Methods that infer both the correspondences and the transformation

The last class of methods aims to estimate the correspondences and the transformation at the same time. This is usually performed in an iterative way, where one component is estimated first and then, based on this estimation, the second is refined.

On of the most well known approaches is the Iterative Closest Point (ICP) method [277]. The advantage of this algorithm lies is its simplicity and speed. Correspondences

are defined based on a closest (in a geometric sense) neighbor principle. Based on this estimation, the transformation is calculated. Then, a new closest neighbor is assigned to each key-point and the process continues till convergence.

ICP has drawn a lot of attention and a number of researchers have tried to ameliorate the method over the years ([278] presents an overview of the improvements over ICP). The performance of the ICP depends greatly on the initial conditions. In general, a good overlap of the two point-sets is necessary for the method to converge to a good minimum. This stems from the assumption that the closest point is a good approximation of the true correspondence. Thus, Sharp *et al.* [279] investigated the use of invariant features in addition to the positional information in order to make the method more robust.

Penney *et al.* [280] also tried to improve on the precision and robustness of the algorithm. To achieve that, Gaussian noise was added to the positions of the points in one set before each iteration of the original ICP. The magnitude of the noise was decreased as the process advanced. [281] tried to render the algorithm more robust by allowing for anisotropic noise in both target and source point sets. The problem was cast in the form of a Generalized Total Least Square problem. [278] improved on the robustness and accuracy of the algorithm for free form deformation shapes by considering collinearity and closeness constraints.

One of the most important works in the domain was introduced by Chui *et al.* [282]. The proposed TPS-RPM algorithm iterates between estimating the correspondence with the softassign method and computing the transformation field with a TPS model. [225] further refined the latter approach by iteratively solving a clustering and matching problem. Stewart *et al.* [229] proposed a dual-bootstrap ICP method to register retinal images. The method operates initially on small regions where accurate correspondences can be obtained. Based on these correspondences low order transformations are estimated. In the subsequent steps, the size of the regions as well as the order of the transformation model are refined. The region refinement is based on the uncertainty of the transformation.

The iterative refinement between the estimation of the correspondences and transformation can be naturally formulated in an Expectation-Maximization fashion. In [283], an approach named multi-scale EM-ICP was proposed. The method is similar to standard ICP with a Mahalanobis distance. The principal difference lies in the estimation of the transformation step where multiple matches weighted by Gaussian weights are considered. [284] aimed to render spectral methods for matching more robust to noise and outliers. Towards this end, the use of an extension of EM [285] along with a spectral method to compute the correspondence probabilities was investigated. Finally, an EM approach was presented in [286] where correspondences are found between scale-invariant salient region features. Based on the set of all pairs of correspondences, a global transformation is estimated. Then, the set of correspondences is refined by considering the increase of

global aligned-ness.

As previously seen, Procrustes analysis is a useful tool when the correspondences are known [241]. When that is not the case, iterative methods can be used to enhance the performance of the Procrustes method. In [287], the use of the Dual-Step EM algorithm is studied. Rangarajan *et al.* [288] proposed the softassign Procrustes method treating the problem from an optimization point of view.

2.4 Optimization Strategies

While most research efforts have been concentrated on devising better similarity metrics as well as transformation models and regularization terms, little attention has been paid to the last component of registration, optimization. It is usually treated as a black-box and techniques are used in a plug and play fashion. Even though few methods have been designed specifically for the registration problem, it would not be an exaggeration to say that almost all optimization methods have been tested to tackle it.

Optimization methods may be discerned based on the nature of the variables that try to infer in two categories: i) *continuous*, and ii) *discrete*. The first class of methods treats real valued variables while in the second case, the variables take values from a discrete set. The two previous classes of methods are constrained with respect to the nature of the objective function as well as the structure to be optimized. *Heuristic* and *metaheuristic* methods bear not the previous limitations. Following, some typical exemplars of each class of methods that have found applications in image registration are going to be presented. Particular emphasis will be put on Markov Random Field (MRF) formulations for image registration. An approach that has gained significant attention during the past few years.

2.4.1 Continuous Optimization

Typically, continuous optimization methods are constrained to problems where the variable take real values and the objective function is differentiable. In medical image registration, it is often the case that both the previous constraints are satisfied. Moreover, being rather intuitive and easy to implement, continuous optimization methods have been applied to numerous registration problems. Let $\boldsymbol{\theta}$ denote the vector of parameters of the transformation and t index the number of iteration. Continuous optimization methods estimate the optimal parameters following an update rule of the following form:

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha_t \mathbf{g}_t(\boldsymbol{\theta}_t), \tag{2.29}$$

where α_t is generally referred to as step size or gain factor, while g defines the search direction.

There are various ways to define the previous parameters. For example, the step size may be constant, decrease with the iterations or so that it minimizes the objective function along the search direction (exact or inexact line search). The search direction can be specified by considering only first-order information or, for example, by also taking into consideration second-order one. It is the choice with respect to the previous parameters that discriminates different methods between them. The approaches the most frequently used in medical image registration are:

- Gradient Descent (GD).
- Conjugate Gradient (CG).
- Powell's conjugate directions.
- Newton-type methods.
- Levenberg-Marquardt (LM).
- Stochastic gradient descent.

A comparison study between some algorithms of the previous categories along with their description was presented in [289].

The simplest approach is to optimize the objective function by following the direction that decreases the energy or its negative gradient. In other words, the direction is given as $\mathrm{g} = -\nabla_\theta(\boldsymbol{\theta})$. In [289], two variants of gradient descent are tested. The first employs a decaying with the iterations function of the step size while the second is based upon the inexact line search algorithm of Moré and Thuente [290]. Other line strategies include keeping the step size fixed, monotone line search [291], line search and golden section search [292]. Gradient descent has been used by numerous researchers to solve a registration problem. In the LDDMM framework, usually posed in a variational setting, gradient descent is often used to solve the problem [50, 58, 63]. Johnson's consistent registration approach [293] as well as Rueckert's FFD registration algorithm [117] were also based upon a gradient descent optimization scheme. Without trying to give a full account for all registration methods that employ gradient descent, a task that can be deemed impossible, let us also cite two more variational approaches [144, 258] .

While gradient descent is intuitive and easy to compute it is known to suffer from slow convergence. Therefore, techniques that have better convergence rates have been tested. Conjugate gradient descent methods try to exploit the knowledge conveyed by the previous gradients so that to proceed not down to the new gradient but instead, towards a direction that is conjugate to the previous. Thus, the direction now is given as $\mathrm{g}_t = f(\nabla_\theta(\boldsymbol{\theta}_t), \mathrm{g}_{t-1})$, where f usually depicts a linear combination $\mathrm{g}_t = \nabla_\theta(\boldsymbol{\theta}_t) + \beta_t \mathrm{g}_{t-1}$. Different ways to

define the weighting factor β_t have been proposed. Some examples of registration methods that use conjugate gradient descent as an optimizer are [54, 294, 65, 295]. An interesting tailored approach for FFD image registration using a preconditioned gradient scheme was presented in [296].

A similar in spirit optimization approach that has been used in an important number of registration problems is Powell's or Direction Set method [292]. Powell's method aims to minimize the objective function by following conjugate directions but contrary to the previous method, no gradient information is used to produce them. The basic procedure that Powell proposed sets the initial direction to the basis vectors $\mathbf{g}^i = \mathbf{e}^i,\ i = 1, \cdots, N$, optimizes along each parameter axis independently from the rest, performs the replacement $\mathbf{g}_t^i = \mathbf{g}_{t-1}^{i+1}$ while adding $\mathbf{g}_t^N = \boldsymbol{\theta}_{t-1} - \boldsymbol{\theta}_0$ and iterates till convergence. Despite being less efficient than the previous method, being gradient free, Powell's method has been applied in many registration tasks [166, 171, 172, 187, 191].

Another class of optimization methods that has been tested in an important number of registration applications is the Quasi-Newton (QN) methods [292]. This class of methods, like the two previous ones, aims to accumulate information from the previous iterations and take advantage of it in order to achieve better convergence. More specifically, their goal is to estimate the inverse Hessian matrix $H^{-1}(\boldsymbol{\theta})$ and use it to define the search direction. Thus, the search direction is defined as $\mathbf{g} = -\hat{H}^{-1}(\boldsymbol{\theta})\nabla_\theta(\boldsymbol{\theta})$, where the $\hat{}$ denotes that an approximation is used (the true Hessian is used in the case of Newton's or Newton-Raphson method). Two main algorithms exist in this category, the *Davidon-Fletcher-Powell (DFP)* and the *Broyden-Fletcher-Goldfarb-Shanno (BFGS)*. BFGS is considered to be slightly better than DFP. A version of BFGS that uses less memory (L-BFGS) was tested in [289]. Other efforts where researchers have experimented with Quasi-Newton methods can be found in [255, 82, 185, 128]

An optimization method of the same family is the Gauss-Newton (GN) algorithm. It is devised to solve optimization problems where sum of squared function values are considered. This is of particular interest for image registration as such objective functions are common when aligning images from the same modality. The advantage of this algorithm is that it does not require the computation of second derivatives. Instead, the Hessian is approximated by ignoring derivatives higher than first order as $\hat{H} = 2J^TJ$ where J denotes the Jacobian. The search direction is now given as $\mathbf{g} = -(J^T(\boldsymbol{\theta})J(\boldsymbol{\theta}))^{-1}\nabla_\theta(\boldsymbol{\theta})$. The Gauss-Newton optimizer has been used in [62, 145, 154]. The Gauss-Newton algorithm is the optimizer that is most frequently used in the demons registration framework to tackle mono-modal registration [32, 29, 41, 81, 43]. In the demons registration setting, an extension of Gauss-Newton [31] was employed to derive the symmetric demons forces [33, 36]. This algorithm exploits more knowledge with respect to the problem at hand. More specifically, it takes advantage of the fact that when the images are aligned,

the gradient of the source can be approximated by the gradient of the target.

A related to the previous method that has been successfully applied to the problem of image registration is the Levenberg-Marquardt algorithm. The search direction in this case is given by: $\mathbf{g} = -\left(\hat{H}^{-1}(\boldsymbol{\theta}) + \zeta\mathbf{I}\right)\nabla_\theta(\boldsymbol{\theta})$. $\mathbf{I}$ is the identity matrix and ζ is a weighting factor that regulates the performance of the optimizer with respect to its speed and stability. By decreasing its value, greater speed may be achieved. At the limit, when ζ equals to zero, we fall to the previous algorithm. On the contrary, when its value increases, the stability increases as well. For some applications of the LM approach the interested reader is referred to [18, 118, 129, 130]. Based on the LM algorithm, Thevanez *et al.* [297] proposed an efficient optimizer for mutual information driven registration. [118] compared the LM algorithm with GD, GD with a quadratic step size estimation and CG to find that it performs the best for a FFD registration task.

The previous techniques cover the *deterministic* gradient methods that are used most often to solve the optimization problems that arise when tackling image registration. In medical image registration, the computation of the derivative information can be computationally demanding. Thus, in order to alleviate the computational burden, many researchers have experimented with *stochastic* gradient approaches. The update rule for the stochastic gradient approaches is based upon an approximation of the gradient,

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha_t\hat{\mathbf{g}}_t(\boldsymbol{\theta}_t). \tag{2.30}$$

The only difference with respect to the deterministic approaches is that an approximation $\hat{\mathbf{g}}_t(\boldsymbol{\theta})$ of the gradient is used.

The variants of the stochastic gradient approach differ with respect to the way the gradient is approximated. In [289], three approaches were considered. The first one, referred to as Kiefer-Wolfowitz (KW) [298], approximates the gradient by a finite difference scheme. The second one, known as Simultaneous Perturbation (SP) [299], estimates the gradient perturbing it not along the basis axis but instead along a random perturbation vector $\boldsymbol{\Delta}$ whose elements are independent and symmetrically Bernoulli distributed. The last method that was studied is the one proposed by Robbins and Monro (RM) [300]. It is more general, in the sense that it only assumes that an approximation of the gradient exists. This method uses a step-size that decreases with time in order to decrease the inaccuracy. [289] estimated the gradient by using a subset of the image voxels sampled uniformly. Their conclusion is that the RM method performs better than the rest. This method was extended in two subsequent works [301, 302] by considering adaptive image-driven strategies. In [301], an adaptive step mechanism was presented while in [302] an edge-driven importance sampling was proposed to improve the gradient approximation. For some applications of stochastic gradient see [1, 164, 163, 194].

All the previous approaches aim to solve an *unconstrained* optimization problem. As we have seen in Sec. 2.2.3, *constrained* optimization problems often arise when trying to impose task-specific conditions on the deformation field. The solution of such optimization problems is more challenging. The optimization strategies that are usually employed transform the constrained to an unconstrained one that can be solved efficiently. For example, in [145], a log-barrier method was used. Most often, the transformation is achieved by augmenting the dimensionality of the problem using the method of Lagrange multipliers [149, 151].

2.4.2 Discrete Optimization

The main drawback of the previous methods is that they proceed by performing local search in the parameter spacing. As a consequence, they are sensitive to the initial conditions and tend to get trapped in local minima. Moreover, the fact that they rely on the computation of the gradient of the objective function limits their use in two ways. On the one hand, not all objective functions are differentiable. On the other hand, their practical use is hindered by the fact that they are not modular with respect to the objective criterion and the transformation model. Last but not least, they are often computationally inefficient.

On the contrary, discrete methods, performing a global search, are robust to initial conditions while exhibiting better convergence rates in comparison to the continuous methods. Furthermore, their modularity permits their use in various settings. Their limitation stems from the quantization of the search space and the lack of precision that is thus introduced. This implies an important precision versus computational efficiency trade-off. The more densely the solution space is sampled, the better accuracy we may achieve, though at a higher computational cost. We should also note that the quantization of the search space allows us to introduce our knowledge with respect to where the solution should lie by sampling appropriately. Such a control mechanism is not available in the continuous optimization methods.

Let us attempt a brief historical overview of the methods used to solve the discrete optimization problems arising from Markov Random Field (MRF) formulations that are typically employed in the case of image registration. Even though a thorough presentation of all methods is neither possible nor of interest here, the principal classes of methods along with their basic ideas are going to be traced. More emphasis will be put on the methods that were used in a registration framework and the registration methods themselves.

Before describing any methods, let us introduce some notation. In general, a graph $\mathcal{G}$, consisting of set of vertices $\mathcal{V}$ and a set of edges $\mathcal{E}$ ($\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$), is going to be considered. Depending on the case, the graph may be directed or undirected. The corresponding energy

is the sum of all unary potentials $\mathcal{U}_p$ of the nodes p (*i.e.* data cost) along with the pairwise potentials $\mathcal{P}_{pq}$ (*i.e.* regularization cost) modeled by the edges connecting nodes p and q.

$$E_{MRF} = \sum_{p \in \mathcal{V}} \mathcal{U}_p(l_p) + \sum_{pq \in \mathcal{E}} \mathcal{P}_{pq}(l_p, l_q) \tag{2.31}$$

The random variables take values in a discrete label set $\mathcal{L}$. Minimizing the previous energy result in an optimal labeling $\mathrm{l}^\star$ or in assigning to each random variable p a label l_p.

The first attempts to solve the difficult combinatorial problem were based on *heuristic* and *metaheuristic* strategies. In the first case, one may cite Iterated Conditional Modes (ICM) [303] and Highest Confidence First (HCF) [304, 305]. In the second case, a method inspired by metallurgy, Simulated Annealing [306] should be referred. These methods suffer from lack of optimality guarantees and slow convergence.

The interest on the field was revitalized during the nineties with the introduction of the graph-cut [307] and belief propagation methods [308]. If we would like to class the methods according to the techniques they employ, three different categories should be discerned:

- Graph-based methods.
- Message passing methods.
- Linear-programming (LP) approaches.

The first class of methods is based on the max-flow min-cut principle [309] that states that the maximum amount of flow that can pass from the source to the sink is equal to the minimum cut that separates the two terminal nodes. The two terminal nodes are defined as source and sink depending on the direction of their edges. The cost of a cut is given by the sum of the weights of the edges that were removed.

Among the most efficient implementations one can distinct the push-relabel [310] and the graph-cut method [307]. Of particular interest due to its application in medical image registration is the α-expansion [311]. α-expansion is the extension to the multi-label case of the method presented in [307]. In [307], it is shown how to calculate the exact maximum a posteriori estimation for the case of the Ising model through a single graph cut computation. Boykov et al. [311] extend the previous technique in the multi-dimensional case by iteratively applying binary graph cuts. The algorithm starts from an initial labeling and then checks every label to see if the energy may be decreased by allowing any set of nodes to change their label to the one under study. The optimal labeling is estimated each time by performing a single graph cut. The method is guaranteed to finish in a finite number of iterations though, most often, it terminates earlier as the energy stops decreasing.

The advantage of these methods is their computational efficiency, especially when regular grids are concerned. Unfortunately, they are limited with respect to the type of energies that they can solve [312]. Several variants of these methods, like fusion moves [313] or dynamic graph-cuts [314] were introduced to further improve convergence at the expense of becoming problem specific.

In medical image registration, α-expansion is the optimizer used in [315, 316, 317, 318, 162]. A graph the size of the image is constructed assuming a 6-connectivity scheme and the solution space is sampled densely resulting in a large set of candidate solutions. The size of the graph as well as the large label set result in computational times that almost prohibit the use of this method in any reasonable clinical setting. Moreover, the registration setting that is used is not able to guarantee the diffeomorphicity of the transformation as the regularization only penalizes the first derivatives of the transformation.

The second class of methods, the so-called belief propagation methods [308], is based on local message exchange between the nodes of the graph and then backtracking towards recovering the best solution to the problem. The message, a vector of size equal to the cardinality of the set of solutions or labels, conveys the belief of a node regarding each solution to its neighboring one. For each label, the message transmitted from a node p to a node q is equal to the minimum of an energy given the label. Belief propagation methods can provide exact inference for chain and tree-structured graphs. Though such a property is not true in the case of graphs that contain loops, they have been shown to converge to satisfactory solutions [319, 320] (in this case, normally they are referred to as Loopy Belief Propagation). Their main strength is their ability to copy with non-submodular pairwise interactions.

Maybe, the most notable drawback of these methods is the computational burden of the message calculation. To alleviate the computational burden, efficient techniques to perform belief propagation were proposed in [321]. The most interesting of the proposed techniques shows how to decrease the complexity of the message computation by reinterpreting the process as a distance transform calculation. The methods proposed in [321] apply to the linear and quadratic case of pairwise interactions. An extension to arbitrary pairwise functions appropriate for parallel architectures was most recently introduced in [322].

The previous methods are rather generic as they are appropriate for any MRF problem. On the contrary, Shekhovtsov et al. [323] proposed an efficient MRF deformation model for non-rigid $2D$ image matching by decomposing the original graph into two layers. The latent variables of each layer model the displacement along each axis. Nodes placed at corresponding positions at each layer are connected with an edge that models the data matching term. In each layer, nodes are connected with each-other following a 4-connectivity scheme. These edges encode the regularization term. For this model, the

operations needed to update the messages are greatly reduced.

In [324], the previous decomposed model was used along with loopy belief propagation to calculate the SIFT-flow. SIFT-flow, similarly to optical flow, aims to match SIFT-descriptors along the flow vectors. Moreover, a combination of the techniques presented in [321, 322] have been tested in graphical processing units to perform fast non-rigid image registration.

More recently, techniques based on Linear Programming, endowed with better theoretical properties, have been proposed to solve the optimization problem. These techniques instead of trying to solve the original problem, that is in general NP-hard, opt for a solution of its LP relaxation. We are going to present briefly two methods that have been used in registration tasks, Fast-PD [325, 326] and TRW-S [327].

Fast-PD casts the original problem as a linear integer program. Then, solutions are derived based on the primal and dual LP relaxations so that the primal-dual gap decreases. The advantages of this optimization technique lie in its generality (it only requires that the pairwise potentials are non-negative), its optimality guarantees, the fact that it provides per-instance approximation factors and, last but not least, its extreme speed.

TRW-S or sequential tree-reweighted message passing is also based on an LP relaxation. The algorithm aims to solve the dual of the relaxation that provides a lower bound of the optimal MRF energy. Hence, the goal is to maximize this lower bound. The lower bound in this case is given by a convex combination of trees. Practically, the algorithm first decomposes the graph into a set of trees. Then, in a sequential order for every node, performs belief propagation in each tree that contains it, followed by an averaging operation. That way, the lower bound is guaranteed to increase. The algorithm iterates till the lower bound ceases to increase (within some precision) or a user-defined maximum number of iterations is reached.

FastPD was first used for image registration in [328]. There, it was used to infer the displacements of a grid-based deformation model. Hard constraints on the set of solutions imposed the diffeomorphic property on the deformation field despite the use of a simple first-order regularization term. On top of the efficient optimizer, a computationally efficient, though approximative scheme, was used for the calculation of the data cost. As a consequence, it was made possible to obtain results for both intra- and inter-modal registration tasks in low computational times. Thereafter, the optimizer have been applied to deformable registration with deformation priors [329], feature-based deformable registration [159] and linear registration [330].

TRW-S has also been used in a number of registration tasks. For example, it was the optimizer that was preferred to optimize the efficient decomposed MRF deformation model in [323]. Discrete graphical models and the algorithms that are used to perform inference on them are efficient when first-order relations between the variables are consid-

ered. On the contrary, their efficiency is more limited when more complex interactions are concerned as that entails inference in higher-order graphs. Nevertheless, research efforts have begun to bear fruit and registration algorithms with a higher-order spatial prior have presented in [331, 332].

2.4.3 Miscellaneous

The previous methods are limited regarding to what objective function and structure they can optimize. *Heuristic* and *metaheuristic* methods, on the contrary, are able to handle a wide range of problems and explore large solution spaces. Nevertheless, they are not able to provide any guarantee with respect to the optimality of the solution.

A heuristic often used in image registration is to make at each step the locally optimal choice. This greedy strategy needs at each step the definition of a set of plausible solutions and a score function. Being gradient free and intuitive, it has been applied to tackle the problem of feature-driven image registration where the cost function is based on the comparison of feature vectors and thus is not differentiable. The candidate sets are constructed in a multi-resolution fashion. Moreover, the features as well as the score function can also be designed to reflect information at different resolutions. More information about the practical implementation of this strategy can be found in [101, 156, 158, 157, 226].

Evolutionary algorithms is a strategy that has been used quite often in medical image registration [333], though most of the times limited to linear registration. These algorithms derive from the theory of evolution and natural selection. These algorithms start from an initial set of solutions. The individual solutions are ranked according to a fitness measure and a subset of them is chosen in a stochastic fashion to generate a new set of solutions. The new set of solutions is generated by adapting the current set following a nature-motivated strategy such as mutation. In [289], the covariance matrix adaptation method was considered [334] and found to converge slowly. For a more elaborated presentation and comparison of state-of-the-art evolutionary methods for image registration the interested reader is referred to [333].

2.5 Discussion

In this chapter, a detailed overview of the current image registration methods was given. The previous analysis can act as a road map to guide our efforts towards proposing novel registration algorithms with emphasis on computational efficiency and versatility.

Let us first discuss how our goals influence our choices regarding the deformation model, the objective function and the optimization strategy. All three choices have an impact upon the qualities of the resulting registration algorithm.

Efficiency depends upon the number of deformation parameters, the computational burden of the evaluation of the objective function and the convergence properties of the optimization strategy. One could imagine that the fewer considered parameters, the lesser the computational load to infer them would be. The complexity of the objective function relates directly to the computational performance of the algorithm. Complex functions requiring computationally expensive calculations will slow down the response. A fast converging optimization strategy is required to restrict the number of cost evaluations.

Versatility is also conditioned upon the number of deformation parameters, the nature of the objective function and the variety of energy forms the optimizer can handle. The number of degrees of freedom bounds the expressive power of the deformation model having a direct impact upon the deformations that may be recovered and thus restricts the problems the registration algorithm can treat successfully. Different problems call for different objective functions to account for the relation between different intensity distributions. A universally applicable objective function does not exist. The optimization method is constrained by the nature of the objective function it can handle. The ideal optimizer should be able to deal with a variety of objective functions.

Let us postpone for now the choice of objective function and focus on the justification of our choices regarding the deformation model and the optimization strategy. As far as the deformation model is concerned, we favor interpolation based models. By controlling directly the parameterization of the deformation model, we can choose the minimum number of degrees of freedom that is needed to recover the deformations of the problem at hand. Thus, we can achieve the best compromise between computational efficiency and versatility. In addition, the local spatial influence of the deformation parameters allows for the design of efficient approximative schemes for the cost calculations. Finally, hard constraints in grid-based models ensure the diffeomorphicity of the deformation, a key property that we desire our algorithm to exhibit.

As far as the optimization strategy is concerned, we opt for discrete optimization methods. On the one hand, they exhibit good convergence rates. On the other hand, they are modular with respect to the objective criterion enlarging the range of applications, the registration algorithm can cope with. Discrete optimization methods are limited with respect to the structures they can optimize. Being able to cope only with discrete structures, they may lack precision. Nonetheless, smart sampling of the solution space may alleviate this precision issue while and in addition boost their efficiency.

In the following chapters, based on these two choices, we will derive novel registration algorithms to tackle the pairwise and group-wise registration problems.

Chapter 3

Hybrid Registration

3.1 Introduction

While the generic lines regarding our choices with respect to the deformation model and the optimization scheme have been drawn, the question of an appropriate objective function is still open.

The previous study enables us to draw two important conclusions. First of all, while in the mono-modal case objective functions that model well the intensity relation between pixels have been devised, that is not the case for multi-modal tasks. There, despite recent advances based on machine learning approaches that come at the cost of being task-specific, information theoretic metrics seem limited as they lack spatial context and their estimation is problematic. Second, local evaluations of similarity measures are advantageous as they render the optimization more robust to local intensity changes provoked by imaging artifacts or the presence of a contrast agent. Our first remark urges us to believe that the matching criterion cannot be based only on intensity information. We firmly believe that geometric information should also be taken into consideration along with spatially localized intensity information. Thus, we opt for a matching strategy that considers in a coupled way a weighted block-matching criterion and a geometric matching one. Such a strategy allows us to deliver efficiently solutions under different settings.

In the remainder of this chapter, a review of the existing methods is presented in section 3.2. Then, in section 3.3, we are going to discuss how to formulate the problem as energy minimization by considering discrete deformation elements, while in section 3.4 we show how to optimize the previous energy in a discrete optimization setting. Then, the experimental validation is presented in section 3.5. In section 3.6, a discussion concludes the chapter.

3.2 Prior Work

As stated in the previous chapter, both iconic and geometric registration methods bear certain advantages while suffering from some shortcomings. Hybrid methods try to capitalize on their advantages by using both complementary information in an effort to get the best of both worlds. Among the methods that belong to this class, the following distinct subclasses may be distinguished based on the way the geometric information is exploited, that is:

- As initialization.
- As constraint.
- Coupled.

Geometric information as initialization

The first subclass treats each type of information separately and sequentially. Registration is decomposed into two independent steps, each one acting principally on a different type of information. Typically, geometric registration precedes providing a rough alignment of the two images. Subsequently, iconic registration is performed to refine the result. Using the same notation as in the previous chapter, the energy comprises two terms, a matching term, $\mathcal{M}$, and a regularization term, R, or:

$$E_{\text{hyb}}(\mathcal{T}) = \mathcal{M}_{\text{ico}}(S \circ \mathcal{T} \circ \mathcal{T}_{\text{geo}}, T) + \mathcal{R}_{\text{ico}}(\mathcal{T} \circ \mathcal{T}_{\text{geo}}), \tag{3.1}$$

Notice that the estimate transformation $\mathcal{T}$ acts as an update upon the transformation $\mathcal{T}_{\text{geo}}$ that has been estimated in a preceding step.

Johnson and Christensen initialized their consistent intensity algorithm with the result of a consistent landmark approach in [293]. The landmark and intensity registration were solved independently in an iterative way till a criterion was met. Landmark information was also used in [335] to provide a coarse registration that was used as initialization for a multi-scale deformable image registration. In [226], a hybrid algorithm that combines surface and volume information to register cortical structures was proposed. The algorithm was initialized with the result of a volumetric approach [101] and was subsequently refined using a surface warping method. Postelnicu *et al.* [295] on the contrary, started from the geometric registration, propagated the result to the whole volume using a biophysical model of the brain and refined with a non-linear optical flow registration algorithm. Recently, a similar approach was presented in [336]. Moreover, Auzias *et al.* in [76], tested their diffeomorphic sulcal-based cortical registration (DISCO) in collaboration with an intensity method (DARTEL [18]) in a sequential manner to further improve their results. In

[337], geometric registration of anatomical structures preceded the iconic registration in order to make CT-PET registration more robust to initialization and local minima. Similar methods have also been proposed in the case of CT abdominal [338] and retinal [339] images.

Geometric information as constraint

Using one type of information independently of the other to initialize the following step usually results in an increase of the robustness of the registration procedure. However, there is no guarantee that the correspondences that were established during the previous step will be preserved. To overcome this limitation, a number of researchers have proposed to use the correspondences that are recovered during a first step of geometric registration as a constraint in the objective function that the iconic registration seeks to optimize. In general, the objective function takes the following form:

$$E_{\text{hyb}}(\mathcal{T}) = \mathcal{M}_{\text{ico}}(S \circ \mathcal{T}, T) + \mathcal{R}_{\text{ico}}(\mathcal{T}) + \mathcal{D}(\mathcal{T}, \mathcal{T}_{\text{geo}}), \tag{3.2}$$

where $\mathcal{T}_{\text{geo}}$ denotes the geometric transformation that can be either sparse or dense and $\mathcal{D}$ measures a distance between the two. The influence of the constraint varies from point-wise to global.

Hellier and Barillot proposed to couple dense and landmark-based approaches for non-rigid brain registration in [224]. In a first step, sulci were extracted and modeled as active ribbons. Then, a matching point algorithm was used to establish geometric correspondences that were subsequently used in a robust function as constraints with local spatial support. In [340], sulcal constraints were also used. A robust point matching method was used to account for outliers. The objective function was enhanced by an additional term. The latter ensured that the estimated deformation field adhered to the point correspondences as well as minimized the normalized mutual information. Normalized mutual information combined with geometric cues was used to tackle brain registration in [341]. Two kinds of geometric cues were employed, landmarks and surfaces. The correspondences for the landmarks were fixed while the surface correspondences were estimated in an ICP fashion. The ratio between the iconic and geometric terms was calculated automatically based on their derivatives.

Adherence to point correspondences was also considered by Rohr *et al.*[342]. The local correlation coefficient was combined with landmarks, detected using a model fitting approach, to register pairs of electrophoresis images. The mean landmark distance was also incorporated in a multi-modal diffeomorphic demons approach to tackle the problem of diffusion weighted imaging distortion in [343]. Registration between preoperative and intra-operative images has also been attempted with the use of hybrid approaches [344].

Avants *et al.* [75] added a landmark inexact matching term in the LDDMM framework in order to compare human and chimpanzee cortices. Landmarks were provided manually to establish either anatomical or functional correspondences between the two species.

A dense deformation field was created in [345] by considering both landmark correspondences and their localization uncertainties. The solution of the registration problem was a compromise between matching the image data, being regular and close to the landmark-based deformation field. A similar approach was presented in [346]. The difference was that a local measure of mutual information was used as an intensity criterion. A dense deformation field was also created from user provided landmark correspondences and TPS interpolation in [347]. The transformation was given as an adaptive combination of intensity- and landmark-fitting. Point information was weighted more in landmarks's vicinity. The advantage of this method is its ability to incorporate any intensity-based algorithm though it cannot guarantee convergence.

While most methods establish geometric correspondences and then encourage the intensity driven deformation field to comply with them without guaranteeing their preservation, Joshi *et al.* [294] imposed geometric correspondences as hard constraints. First correspondences were established between the cortical gray/white matter and gray/CSF surfaces using sulcal constraints. The correspondences were then propagated to the whole cortical volume with the use of an harmonic map. Following, the dense deformation field was refined by considering image intensity information under the hard constraint that the deformation is zero for the previously registered surfaces.

Coupled approach

The disadvantage of the previous approaches is that the flow of information is towards one direction. By treating the problems in a decoupled way, iconic registration may profit from geometric either by being initialized closer to the solution or by being driven by an extra force of adherence to correspondences. However, the geometric registration, being treated independently, does not benefit from the iconic one. In an ideal case, the two problems should be considered in a unified way and solved by minimizing a single objective function simultaneously. The advantage of such a setting would be a more consistent and accurate registration relieved from the previous limitations and endowed with the advantages of both classes of methods. A coupled registration objective function may be defined as

follows:

$$E_{\text{uni}}(\mathcal{T}_{\text{ico}}, \mathcal{T}_{\text{geo}}) = \underbrace{\mathcal{M}_{\text{ico}}(S \circ \mathcal{T}_{\text{ico}}, T) + \mathcal{R}_{\text{ico}}(\mathcal{T}_{\text{ico}}) + \mathcal{H}(\mathcal{T}_{\text{ico}}, \mathcal{T}_{\text{geo}})}_{iconic} + \underbrace{\mathcal{M}_{\text{geo}}(K \circ \mathcal{T}_{\text{geo}}, \Lambda) + \mathcal{R}_{\text{geo}}(\mathcal{T}_{\text{geo}})}_{geometric}, \quad (3.3)$$

where K (respectively Λ) is the set of landmarks in the source S (respectively target T) image domain. $\mathcal{H}$ is a consistency term defined on $\mathcal{T}_{\text{ico}}$ and $\mathcal{T}_{\text{geo}}$. It has the following form:

$$\mathcal{H}(\mathcal{T}_{\text{ico}}, \mathcal{T}_{\text{geo}}) = \int_{\Omega} \delta(\mathcal{T}_{\text{ico}}(\mathbf{x}), \mathcal{T}_{\text{geo}}(\mathbf{x}))d\mathbf{x}. \quad (3.4)$$

The minimization of the coupled objective function results in the estimation of both the iconic and geometric registration, which should ideally be identical upon convergence.

Such a universal energy for the problem of deformable registration was first proposed by Cachier *et al.* [223]. However, they considered a slightly different energy: the coupling was performed through the introduction of an auxiliary smooth deformation field. The authors proposed to extract sulci modeled as point distributions and use them in the coupled formulation to solve brain registration. The previous universal formulation (Eq. 3.3) was solved by iterating between three steps: i) solve for the deformation that minimizes the iconic criterion; ii) solve the geometric one by establishing correspondences between the closest points of the geometrical structures; iii) and finally opt for a smooth deformation that respects both iconic and geometric constraints.

In the same context, the method proposed by Joshi *et al.* [348] should also be referred to. In surface registration, it is often the case that both surfaces are mapped to a common domain, usually a sphere, where registration is performed. Joshi *et al.* proposed to also map the interior brain volumes to the interior of the spheres through harmonic maps. Then, correspondences may be established by considering the complete sphere domain or both the surface and iconic information at the same time.

3.3 Continuous Domain

In this chapter, we opt for the universal approach as defined in Eq. 3.3 combined with our previous choices, that is an interpolation-based transformation model and a discrete optimization strategy. In this section, we introduce in the energy formulation the deformation model and decompose the continuous problem in discrete entities. These entities are modeled subsequently by an MRF and solved by an efficient discrete optimization technique.

3.3.1 Iconic Registration

Let us first detail the iconic part of the registration algorithm which is based upon Glocker *et al.*'s work [328]. As argued earlier, the use of an interpolation-based transformation model is preferred. More specifically, a grid-based one is chosen. Let us consider a set of k control points distributed along the image domain using a grid pattern. Furthermore, let k be much smaller than the number of image points. The dense displacement field can be 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

$$\mathbf{u}(\mathbf{x}) = \sum_{i=1}^{k} \omega_i(\mathbf{x})\mathbf{d}_i, \tag{3.5}$$

and the iconic transformation $\mathcal{T}_{\text{ico}}$ becomes:

$$\mathcal{T}_{\text{ico}}(\mathbf{x}) = \mathbf{x} + \sum_{i=1}^{k} \omega_i(\mathbf{x})\mathbf{d}_i. \tag{3.6}$$

As we have seen in section 2.2.2, 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 a few control points.

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 iconic objective matching term (Eq. 2.24) can be rewritten as:

$$\mathcal{M}_{\text{ico}}(S \circ \mathcal{T}_{\text{ico}}, T) = \frac{1}{k} \sum_{i=1}^{k} \int_{\Omega_S} \hat{\omega}_i(\mathbf{x})\, \rho(S \circ \mathcal{T}_{\text{ico}}(\mathbf{x}), T(\mathbf{x}))d\mathbf{x}, \tag{3.7}$$

where $\hat{\omega}_i$ are weighting functions similar to the ones in Eq. 3.6.

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 dissimilarity measure with respect to the displacement of this particular control point. This is in line with the local support which a control point has on the deformation. The previous is

valid when point-wise similarity criteria are considered. When an information theoretic criterion is to be used, a different definition of $\hat{\omega}_i$ is adopted,

$$\hat{\omega}_i(\mathbf{x}) = \begin{cases} 1, & \text{if } \omega_i(\mathbf{x}) \geq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{3.8}$$

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 energy can also be expressed on the basis of the set of control points as:

$$\mathcal{R}_{\text{ico}} = \frac{1}{k} \sum_{i=1}^{k} \int_{\Omega_S} \hat{\omega}_i(\mathbf{x})\, \psi(\mathcal{T}_{\text{ico}}(\mathbf{x}))d\mathbf{x}, \tag{3.9}$$

where ψ is a function that penalizes certain properties of the dense deformation field and in practice, can be any of the functions that were studied in section 2.2.

3.3.2 Geometric Registration

As far as the geometric registration is concerned, the case where both the transformation and the correspondences are sought is considered. On the one hand, the aim is to infer the transformation that will establish correspondences between the two sets of landmarks K and Λ such that a dissimilarity criterion based upon their appearances ϱ is minimized, or:

$$\mathcal{M}_{\text{geo}}(K \circ \mathcal{T}_{\text{geo}}, \Lambda) = \frac{1}{n} \sum_{i=1}^{n} \delta(\mathcal{T}_{\text{geo}}(\boldsymbol{\kappa}_i), \widetilde{\boldsymbol{\lambda}}_i) \tag{3.10}$$

where

$$\widetilde{\lambda}_i = \arg\min_{\lambda_j} \varrho(\mathcal{T}_{\text{geo}}(\boldsymbol{\kappa}_i), \boldsymbol{\lambda}_j). \tag{3.11}$$

In other words, the transformation will try to minimize the distance between a landmark and the candidate that it is most similar in appearance to it. We remind that vectors (here the Euclidean position of the landmarks λ and κ) are denoted in bold.

On the other hand, the transformation is required to be smooth or to preserve locally the geometric distance between pairs of landmarks,

$$\mathcal{R}_{\text{geo}}(\mathcal{T}_{\text{geo}}) = \frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j=1,\, j\neq i}^{n} \|(\mathcal{T}_{\text{geo}}(\boldsymbol{\kappa}_i) - \mathcal{T}_{\text{geo}}(\boldsymbol{\kappa}_j)) - (\boldsymbol{\kappa}_i - \boldsymbol{\kappa}_j)\|. \tag{3.12}$$

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

It should be noted that the geometric matching and regularization terms are already in discrete form as they are based on the finite set of landmarks.

3.3.3 Coupling Term

The transformation consistency term (Eq. 3.4) is of great importance as it will couple the solutions of the previous two problems. It can be rewritten in a more efficient way solely based on the landmarks as:

$$\mathcal{H}(\mathcal{T}_{\text{ico}}, \mathcal{T}_{\text{geo}}) = \frac{1}{n} \sum_{i=1}^{n} \|\mathcal{T}_{\text{ico}}(\boldsymbol{\kappa}_i) - \mathcal{T}_{\text{geo}}(\boldsymbol{\kappa}_i)\|. \tag{3.13}$$

If we now also consider a connection between control point displacements $\mathbf{D}$ and landmark displacements, the previous relation can be rewritten as:

$$\mathcal{H}(\mathcal{T}_{\text{ico}}, \mathcal{T}_{\text{geo}}) = \frac{1}{n} \sum_{i=1}^{n} \|\boldsymbol{\kappa}_i + \mathbf{u}_{\text{geo}}(\boldsymbol{\kappa}_i) - \boldsymbol{\kappa}_i - \sum_{j=1}^{k} \omega_j(\boldsymbol{\kappa}_i)\mathbf{d}_j\|, \tag{3.14}$$

where $\mathbf{u}_{\text{geo}}(\boldsymbol{\kappa}_i) = \widetilde{\boldsymbol{\lambda}}_i - \boldsymbol{\kappa}_i$, *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.

At this point, we should redefine with greater precision the goal of this section. The goal is to decompose the energy terms into discrete deformation entities under a certain constraint. Only first-order relations between the entities should arise. The reason behind this constraint is that such a type of relations can be modeled with the use of first-order MRFs and thus solved by standard discrete optimization algorithms. The previous term (Eq. 3.14) includes more than two discrete entities and thus does not comply with the previous constraint.

To account for that, we are going to approximate Eq. 3.14 by its upper bound. Let us detail the calculation. By taking into account that the sum of the coefficients ω_j is equal to one, Eq. 3.14 can be rewritten as:

$$\begin{aligned}\mathcal{H}(\mathcal{T}_{\text{ico}}, \mathcal{T}_{\text{geo}}) &= \frac{1}{n} \sum_{i=1}^{n} \| \sum_{j=1}^{k} \omega_j(\boldsymbol{\kappa}_i)\mathbf{u}_{\text{geo}}(\boldsymbol{\kappa}_i) - \sum_{j=1}^{k} \omega_j(\boldsymbol{\kappa}_i)\mathbf{d}_j\| \\ &= \frac{1}{n} \sum_{i=1}^{n} \| \sum_{j=1}^{k} \omega_j(\boldsymbol{\kappa}_i)(\mathbf{u}_{\text{geo}}(\boldsymbol{\kappa}_i) - \mathbf{d}_j)\|. \end{aligned} \tag{3.15}$$

If we now apply the triangular inequality and exploit the fact that the coefficients ω_j are positive, the coupling constraint is redefined as:

$$\begin{aligned}\mathcal{H}(T_{\text{ico}}, T_{\text{geo}}) &\leq \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{k}\|\omega_j(\boldsymbol{\kappa}_i)(\mathbf{u}_{\text{geo}}(\boldsymbol{\kappa}_i) - \mathbf{d}_j)\| \\ &\leq \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{k}\omega_j(\boldsymbol{\kappa}_i)\|\mathbf{u}_{\text{geo}}(\boldsymbol{\kappa}_i) - \mathbf{d}_j\|. \end{aligned} \tag{3.16}$$

The above redefinitions allow us to introduce our main contribution which is the mapping of the deformable registration problem into a discrete labeling one. The discrete entities of this problem are the set of control point displacements $\mathbf{D}$, and the two sets of landmarks K and Λ.

3.4 Discrete Domain

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

3.4.1 Iconic Registration

Let us first introduce a graph $\mathcal{G}_{\text{ico}} = (\mathcal{V}_{\text{ico}}, \mathcal{E}_{\text{ico}})$ consisting of a set of vertices $\mathcal{V}_{\text{ico}}$ corresponding to the set of control point displacement updates V, *i.e.* $|\mathcal{V}_{\text{ico}}| = |\Delta\mathbf{D}| = k$. Further, we say that a label assignment $l_p \in \mathcal{L}_{\text{ico}}$ (with $p \in \mathcal{V}_{\text{ico}}$) is equivalent to an update $\Delta\mathbf{d}_p$ of the control point p, *i.e.* $l_p \equiv \Delta\mathbf{d}_p$. So, the label set for this set of variables is $L_{\text{ico}} \subset \mathbb{R}^d$, *i.e.* a discrete version of the displacement space. We also introduce a neighborhood system $\mathcal{E}_{\text{ico}}$ on the vertices, which follows a common 6-connected neighborhood in case of $3D$ and a 4-connected one in $2D$ corresponding to the uniform grid like structure of the transformation model.

According to Eq. 3.7 we define the unary potentials as:

$$\mathcal{U}_{\text{ico},p}(l_p) = \int_{\Omega_S} \hat{\omega}_p(\mathbf{x})\,\rho(S \circ \mathcal{T}_{\text{ico},l_p}(\mathbf{x}), T(\mathbf{x}))\,d\mathbf{x}, \tag{3.17}$$

where $\mathcal{T}_{\text{ico},l_p}$ 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 as 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. 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 $l \in \mathcal{L}_{\text{ico}}$ 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 $\mathcal{T}_{\text{ico},l_p}(\mathbf{x}) = \mathcal{T}_{\text{ico}}(\mathbf{x}) + l_p$, where $\mathcal{T}_{\text{ico}}(\mathbf{x})$ is the current or initial estimate of the transformation.

The compact regularization term defined in Eq. 3.9 could be defined as well in the above manner. However, this 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:

$$\mathcal{P}_{\text{elastic},pq}(l_p, l_q) = \frac{\|(\mathbf{d}_p + \Delta\mathbf{d}_p) - (\mathbf{d}_q + \Delta\mathbf{d}_q)\|}{\|\mathbf{p} - \mathbf{q}\|}. \tag{3.18}$$

The pairwise potentials penalize deviations of displacements of neighboring control points $(p, q) \in \mathcal{E}_{\text{ico}}$ 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.

Before moving to the geometric registration case, let us detail how the label set $\mathcal{L}_{ico}$ is constructed so that the accuracy-efficiency trade-off mentioned in the previous chapter is faced with success. We remind that a small set of labels is desirable so that inference can be performed efficiently. The cost to pay is the decrease of the accuracy of the registration as it is bounded by the range of deformation covered in the set of labels. As a consequence, it is reasonable to assume that the registration result is sub-optimal. In order to tackle

the previous shortcomings in an efficient way, 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. 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 $|\mathcal{L}_{\text{ico}}| = g \cdot s + 1$ 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 $\mathbf{d}_{\text{max}}$. 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 depends 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 $3D$). Additionally, we can add samples for instance along diagonal axes. In $2D$ we commonly prefer a star-shape sampling, which turns out to be a good compromise between the number of samples and 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 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 labels sets provide an excellent performance in terms of computational speed and registration accuracy.

3.4.2 Geometric Registration

Let us now introduce a second graph $\mathcal{G}_{\text{geo}} = (\mathcal{V}_{\text{geo}}, \mathcal{E}_{\text{geo}})$ 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 $\mathcal{V}_{\text{geo}}$ corresponding to the set of land-

marks extracted in the source image, *i.e.* $|\mathcal{V}_{\text{geo}}| = |K|$. A label assignment $l_p \in \mathcal{L}_{\text{geo}} := \Lambda$ (where $p \in \mathcal{V}_{\text{geo}}$) 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 $\mathbf{u}_{\text{geo},l_p}(\boldsymbol{\kappa}_p) = \boldsymbol{\lambda} - \boldsymbol{\kappa}_p$, since κ_p is mapped on the landmark l_p.

According to Eq. 3.11, the unary potentials are defined as:

$$\mathcal{U}_{\text{geo},p}(l_p) = \varrho(\kappa_p, l_p). \tag{3.19}$$

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 $\mathcal{T}_{\text{geo},l_p} = \boldsymbol{\kappa}_p + \mathbf{u}_{\text{geo},l_p}(\boldsymbol{\kappa}_p)$ 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. As we have seen in section 2.3.2, a common way is to construct descriptors that are able to capture local geometric information and compare them in a point-wise manner. That would be the case expressed by the previous equation (Eq. 3.19). Let us rewrite the equation taking into account the descriptor information Desc:

$$\mathcal{U}_{\text{geo},p}(l_p) = \varrho(\text{Desc}_S(\kappa_p), \text{Desc}_T(l_p)). \tag{3.20}$$

A more robust way would be to consider neighborhood information. That can be easily done by evaluating the criterion over a patch centered around the landmarks,

$$\mathcal{U}_{\text{geo},p}(l_p) = \int_{\Omega_{S,p}} \varrho(S \circ \mathcal{T}_{\text{geo},l_p}(\mathbf{x}), T(\mathbf{x}))d\mathbf{x}, \tag{3.21}$$

where $\Omega_{S,p}$ denotes a patch around the point κ_p.

The regularization term defined in Eq. 3.12 can be encoded by the edge system $\mathcal{E}_{\text{geo}}$ of the graph. In this setting, the regularization term can be expressed as:

$$\mathcal{E}_{\text{geo},pq}(l_p, l_q) = \|(\mathcal{T}_{\text{geo},l_p}(\boldsymbol{\kappa}_p) - \mathcal{T}_{\text{geo},l_q}(\boldsymbol{\kappa}_q)) - (\boldsymbol{\kappa}_p - \boldsymbol{\kappa}_q)\|. \tag{3.22}$$

The pairwise potential will enforce an isometric constraint. Having performed an affine registration step to account for scale differences as preprocessing, it is natural to suppose that the geometric distance between homologous points in the two images does not vary too much. Moreover, by considering the vector differences flipping of the point positions is penalized. A similar graph has been proposed in [237] to tackle a different problem. There, the aim is to decide which correspondences to keep in different parts of the image so that the resulting transformation will be smooth.

The previous edge system consists of edges between all nodes (complete graph). As shown in [349], there is a certain redundancy in the information modeled by such graphs. On top of that, the great number of edges adds computational burden. In [349], the authors propose to obtain a more compact model by performing a clustering step [350] based on shape map distances [351]. Here, we take a simpler approach and consider edges only between a node and a user-specified number of its neighbors. The underlying assumption is that close points (in terms of Euclidean distance belong most probably to the same anatomical structure and thus are likely to undergo similar deformations.

3.4.3 Coupling Term

The last missing term of the universal registration energy is the coupling penalty (Eq. 3.16). We introduce a third set of edges $\mathcal{E}_{\text{hyb}}$ 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:

$$\mathcal{P}_{\text{hyb},pq}(l_p, l_q) = \omega_q(\boldsymbol{\kappa}_p) \left\| \mathbf{u}_{\text{geo},l_p}(\boldsymbol{\kappa}_p) - (\mathbf{d}_q + l_q) \right\|, \tag{3.23}$$

where $p \in \mathcal{V}_{\text{geo}}$ and $q \in \mathcal{V}_{\text{ico}}$, $l_p \in \mathcal{L}_{\text{geo}}$ and $l_q \in \mathcal{L}_{\text{ico}}$, and $(p, q) \in \mathcal{E}_{\text{hyb}}$. Such a pairwise term couples the displacements given by the two registration processes and imposes consistency. The coupled registration objective function (Eq. 3.3) is finally represented by an MRF graph $\mathcal{G}_{\text{hyb}} = (\mathcal{V}_{\text{geo}} \cup \mathcal{V}_{\text{ico}}, \mathcal{E}_{\text{geo}} \cup \mathcal{E}_{\text{ico}} \cup \mathcal{E}_{\text{hyb}})$ with its associated unary and pairwise potential functions.

3.5 Experimental Validation

In order to validate the coupled geometric registration method and due to the dependence of the performance of the proposed algorithm on a landmark extraction preprocessing step, 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.

Starting from a co-registered pair of T1-T2 brain MRI images of size $256 \times 256 \times 48$ and of physical resolution equal to $0.9375 \times 0.9375 \times 3$mm, 10 target images were generated from the 3D T2 MRI image. The deformation fields are created by composing two different configurations of FFD grids. The first one has a control spacing equal to 48mm resulting in coarse deformations, while the second one, with a control spacing of 18mm, is responsible for more local deformations. The different configurations of the

FFD grids are produced by moving randomly the grid nodes in all three dimensions. The maximum displacement that the grid nodes can perform is bounded such that the resulting deformation is diffeomorphic. That is, for the first grid it is equal to 19,2mm and for the second one is 7,2mm. In each of the 10 generated images, 300 to 400 landmarks positioned in salient image regions were manually selected. The landmark set K consists of these points. Given the deformation fields, for each point of interest in the target image domain its homologous in the source image domain can be determined. Around each homologous point, 9 points are selected randomly. These points along with the true correspondence form the set of the candidate landmarks Λ.

In this experimental setting, the initial T1 image is used as source and is registered to the generated T2 ones. Once the registration task is performed, the estimated deformation field is compared to the one that was used to create the images. The difference between the two deformation fields is quantified by measuring the end point and the angular error. Voxels that belong to the background are ignored during these calculations. The end point error (EP) is measured in millimeters and is simply the Euclidean distance between the two $3D$ motion vectors. The angular (AE) between the two motion vectors is the angle between them. It can be computed by calculating the inner product between the normalized vectors and then calculate the inverse cosine. The use of these errors measurements is common in the literature of optical flow [352].

Given the multi-modal nature of the data, the mutual information is used as the iconic criterion. Regarding the regularization term of the iconic graph, the diffusion regularization with fluid-like behavior is preferred as it leads to faster computations. The pairwise terms modeling the diffusion regularization are weighted by a factor of 0.01. As our main interest is to outline the potential of the proposed coupled approach, the matching criterion for the case of the geometric registration is based on the Euclidean distance between the true correspondence and the point indexed by the label at hand. The unary potentials of the geometric graph are multiplied by 10. Regarding the pairwise potential for the geometric case, the one presented in (Eq. 3.22) is used as is. Last but not least, the weight for the coupling term (Eq. 3.23) is set to 0.01.

Our implementation considers FFDs with cubic B-spline basis functions and control points uniformly distributed over the image domain as the compact transformation model. We employ a common hierarchical coarse-to-fine strategy where the resolution of both the images and the FFD control grid is subsequently refined during the registration process. For the images, we use Gaussian image pyramids with a scale factor of 2 and standard deviation of 1. The initial control point spacing is set in millimeters according to the expected maximum amount of deformation (here it was set to 60mm). With increasing image resolution, we halve the FFD spacing by inserting new control points, which allows us to recover larger deformations in the beginning and smaller, more subtle ones in later

iterations. Each registration takes around 15 minutes on an Intel Xeon X5680 3.33 GHz CPU.

Regarding the discrete displacement label space $\mathcal{L}_{\text{ico}}$ for the iconic part we used a sparse sampling of the deformation space where samples are taken only along the positive and negative main axes of the 3D coordinate system. We set the number of samples equal to 5, yielding 31 labels in total (including the zero displacement). The initial maximum displacement magnitude on each grid resolution level is coupled to the control point spacing by a factor of 0.4. This is motivated by the fact that composing the single updates on the transformation bounded to this maximum displacement yields diffeomorphic transformations. After each iteration, which corresponds to one MRF labeling computation, we rescale the set of labels with a fixed factor of 0.66. On each level, we commonly perform five iterations before increasing the image and grid resolution. This strategy allows for small label sets that can be optimized efficiently, whereas the successive refinement allows sub-millimeter registration accuracy as demonstrated in the following experiments.

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 and angular error (Table 3.1, 3.2), we see that our coupled iconic geometric registration method is able to further improve the results of the iconic one. This is evident, as both the end point and angular error have 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 Tables 3.1, 3.2. 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. Visual results for one specific case are shown in Fig. 3.1. A checkerboard representation is created by interchanging parts of the two images. It is evident from the difference between the neighboring tiles that the images differ greatly before the registration. The continuity of the patterns across neighboring tiles suggests that the images have been aligned. No significant visual difference can be observed between the result obtained with the pure iconic method and the one that is obtained by using the coupled approach.

#	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 3.1: End point error (in millimeters) for the registration of the Synthetic MR Dataset. The grid spacing is denoted by h.

#	Iconic (h=60mm)		Hybrid (h=60mm)		Iconic (h=20mm)		Hybrid (h=20mm)	
	mean	std	mean	std	mean	std	mean	std
1	8.70	6.66	**8.18**	6.16	8.25	8.92	**6.08**	5.12
2	6.85	4.37	**6.43**	4.32	13.89	23.73	**5.13**	3.93
3	10.40	10.52	**8.50**	7.76	14.02	17.12	**7.34**	7.29
4	14.46	11.50	**11.69**	10.24	14.61	14.34	**11.49**	11.72
5	7.65	7.79	**7.23**	7.93	8.31	9.32	**6.17**	6.97
6	8.20	7.08	**7.21**	6.16	9.85	13.64	**6.13**	6.13
7	10.19	11.20	**9.50**	10.19	12.01	15.06	**8.67**	9.69
8	8.67	9.75	**8.01**	8.49	10.87	15.14	**6.69**	7.58
9	9.60	9.19	**8.84**	8.014	13.62	19.06	**8.03**	8.95
10	13.12	14.43	**11.03**	12.64	17.42	20.26	**9.24**	11.15
all	9.78	2.39	**8.66**	1.68	12.29	2.97	**7.50**	1.90

Table 3.2: Angular error (in degrees) for the registration of the Synthetic MR Dataset. The grid spacing is denoted by h.

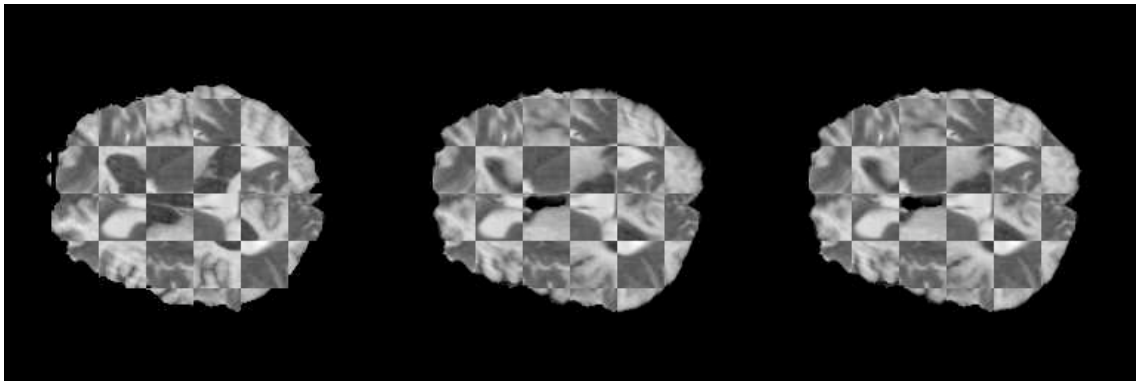

Figure 3.1: Visual results before (left column) and after registration of image 2 using the proposed iconic (middle column) and hybrid approach (right column). The results are given in the form of a checkerboard where neighboring tiles come from different images.

3.6 Discussion

In this chapter, we proposed a framework to couple the point correspondence and the dense registration problem into a unified objective function where the two problems are solved simultaneously in a one-step optimization. A Markov Random Field formulation was employed towards introducing individual costs for the family of parameters and their interactions. The advantages of the proposed framework can be summed up to:

- Its ability to use any intensity-based criterion due to the discrete nature of the formulation. To the best of our knowledge, only one other hybrid method can claim that [347], but it cannot guarantee the convergence as a ping-pong effect is possible between iterations.
- The diffeomorphicity of the obtained deformation field by controlling appropriately the set of solutions.
- The influence of the landmarks is defined in a local way and in respect to the deformation model without introducing heuristics or any assumption on the number or the nature of the landmarks.

The proposed framework exhibit some shortcomings. One of the limitations lies in the use of an upper bound of the energy for the coupling term. This shortcoming was tackled in [353] where the exact L_2 distance was used to couple the geometric and iconic information, modeled by pairwise relations. The need for a global linear registration was alleviated in [354] by using learned higher-order relations for the graph that models the geometric

problem. Moreover, the landmarks are matched by using an optimized similarity metric. Last but not least, the proposed registration algorithm is asymmetric. In the next chapter, a *symmetric* extension of the previous algorithm is detailed.

Chapter 4

Symmetric Hybrid Registration

4.1 Introduction

The result of the previous algorithm depends upon the selection of the template. Swapping the order of the images does not result in an inversion of the estimated deformation field. This is not only counter-intuitive but also undesirable. To illustrate better the significance of this effect, let us consider the case of a group study where voxel-wise comparisons take place. By changing the template image, the failure to produce consistent deformation fields would result in comparisons that are biased upon the choice of the template. Lastly, in various cases symmetric algorithms have been shown to perform better than asymmetric ones, even in the case where the asymmetric one delivers diffeomorphic results [33]. The above shortcomings lead us naturally to consider endowing our hybrid registration algorithm with the symmetric property.

In the remainder of this chapter, a review of the existing methods is presented in section 4.2. Following, in section 4.3, we discuss how to formulate the problem as energy minimization by considering discrete deformation elements, while in section 4.4 we show how to optimize the previous energy in a discrete optimization setting. Then, the experimental validation is presented in section 4.5. In section 4.6, a discussion concludes the chapter.

4.2 Prior Art

In the first chapter, we had the opportunity to come across some symmetric algorithms in the course of the bibliographical study. In this section, we are going to attempt a more detailed account of the symmetric methods available in the literature. On a side note, let us

state that these algorithms are usually referred to as either *inverse consistent* or *symmetric*. There is an important difference between the two. The first try to approximate symmetry by incorporating penalty terms in the objective function, while the second aim to guarantee the exact symmetry of the results. Here, we favor the use of the term *symmetric* in order to put more emphasis on exact symmetry being the goal of this work. Moreover, this term seems more appropriate as it is commonly applied in mathematics to describe similar situations.

The interest towards symmetric registration algorithms was popularized by Christensen and Johnson's seminal work [14]. There, both the forward and backward transformation were estimated by minimizing an objective function comprising of terms evaluated on both source and target image domains. The two transformations were deduced separately, while the coupling was achieved by imposing a constraint that penalized the difference of one from the inverse of the other.

This framework has been proven of great importance and has been extended in various registration settings. In [293], it was extended to the case where both landmark and intensity information are considered. In [16], an extension of the algorithm that can account for large deformations by considering a concatenation of transformations was provided. Landmark-initialized, segmentation and intensity-based inverse-consistent linear elastic image registration was investigated in [355]. In [356, 357], object matching by considering object boundary information and the enclosed intensity patterns has been investigated. In a similar context, manifold registration that exhibits the inverse consistency and transitivity properties was presented in [358]. In the case of multi-modal registration, an inverse consistent algorithm driven by mutual information was presented in [359, 360] while an FFD inverse consistent method was presented recently in [122].

Despite their success, this family of algorithms has a number of disadvantages. First and foremost, the algorithms are only asymptotically symmetric when the weight given to the inverse consistency penalty tends to infinity. Moreover, they require the computation of an inverse transformation despite the fact that its existence is not ensured. A number of methods were devised in order to ameliorate upon the previous algorithms.

In [101], in the hierarchical attribute matching algorithm forces are estimated from the template to the subject and vice versa. Consistency is enforced only on the driving voxels and not the whole volume by letting the inverse transformation act as an additional constraint on the forward one. An ad hoc symmetrization of the image forces inspired by the Newton's action-reaction law was presented in [361].

In the symmetric log-domain diffeomorphic Demons framework [33], the force is rendered symmetric by a simple projection of the forward and backward forces separately calculated by exploiting a log-domain representation of the deformation fields. In the diffeomorphic framework proposed by Hernandez *et al.* [84], the authors symmetrize the

cost by evaluating it over both image domains by considering the transformation and its inverse, that now is guaranteed to exist. In a similar approach, in the LDDMM framework, the cost was evaluated in both image domains by considering the transformation at time 1 and its inverse [68].

Leow *et al.* [15] consider only the forward and the backward transformation. No inverse transformations are calculated but instead the backward transformation is constrained to be the inverse of the forward during the gradient descent optimization. The previous algorithm and its embedded consistency constraint was extended in the fluid registration framework to tackle the registration of diffusion tensor images driven by an information theoretic criterion [23]. Another algorithm that imposes the inverse consistency property during the optimization, claiming great efficiency, was presented most recently in [362]. There, the consistency property is imposed explicitly. Each step of a gradient descent based interleaved optimization that solves in turns for the forward and the backward transformation is initialized by the inverse of the transformation that was estimated during the previous step.

While adding an extra term evaluated in the second image domain is sufficient to symmetrize the energy, it is more computational demanding and more importantly does not shed any light on the causes of the asymmetry. The need for a better insight into the reasons of the inherent asymmetry of registration algorithms was fulfilled in [363, 364]. According to these studies, there are four main sources of asymmetry: i) Order non-preservation of the energies due to the use of standard differential volume forms on the domain of the images. ii) Non-stable deformation space. iii) Local minima. iv) Numerical implementation reasons.

Cachier and Rey [363], probably in the first work to study symmetry, restrict themselves in tackling the first reason. They propose the symmetrization of the energy with the use of the Jacobian of the transformation. Regarding the regularization energy, they propose to multiply the standard term with $(1+\frac{1}{det(J)})$. As for the matching term, the multiplication with the quantity $(1+det(J))$ is recommended. Tagare *et al.* [364] based on a thorough mathematical analysis conclude that the term that Cachier and Ray proposed for the matching term is the optimal one. Moreover, they show how to guarantee numerically the symmetry.

As for the regularization term, they show that the one proposed in [365, 366] and defined as $\int(det(J)-1)\log(det(J))dx$ satisfies the required mathematical properties. Pennec *et al.* also presented a symmetric version of the elastic regularized in [17]. In [367], a symmetric prior was devised for the triangles that govern the deformation of a $2D$ image domain. The prior is based on the Jacobian of the determinant and the eigenvalues of the deformation tensor.

While the previous showed how to construct order preserving energies defined in one

image domain assuming that only one of the images is deformed, there is a great body of work in the diffeomorphic framework that has constructed such energies by letting both images deform towards a common domain. Such algorithms require an inversion and a composition step in order to provide the mapping from one image domain to another.

Joshi *et al.* proposed a method for unbiased atlas construction in the LDDMM framework in [368]. When only two images are considered, they are deformed towards a common domain where the matching term is evaluated and each transformation is regularized by a separate term. Having estimated two forward mappings towards the common domain, one of them has to be inverted and composed to the other in order to estimate the full flow. To get the transformation from one image domain to another, an inversion and a composition step are needed. An extension of the previous framework to the model based multi-modal case is given in [369] while Avants *et al.* devise a similar algorithm using Cross Correlation as a similarity metric [61]. The same framework has also been used recently to register fibre orientation distributions in [370]. Tagare et al [364] argue that the methods that are based upon a composition of different transformations lack a natural geometric interpretation.

The previous methods consider two different flows, each from one image domain to the common one. It is possible though to consider only one. In [60], two such symmetric data terms were reported for the LDDMM framework. For the first data term, called Consistent-Integral-Cost, the matching of the template is not only enforced to the target at the end of the flow, but also at each point along the flow. The second data term is called Consistent-Midpoint-Cost (CMC) (initially suggested and discussed in [77]) and suggests that the matching takes place at the midpoint of the flow.

Studholme *et al.* [184] used a similar scheme to tackle serial brain MRI registration in the presence of local tissue contrast change. The proposed regional Mutual Information that is able to account for such local intensities changes was used to drive a viscous fluid deformation model. The evaluation of the cost takes place at the common domain that is placed halfway between the images. As a result, the number of parameters that have to be estimated is equal to a standard asymmetric pairwise registration.

Inspired from the previous approach, Noblet *et al.* [371] proposed a similar strategy for mono-modal registration. The difference is that an FFD deformation model is used while the flows from the common domain towards the image ones are constructed so that the topology is preserved. This is an important property that ensures that the whole flow from one image to the other can be estimated.

In this chapter, we investigate the symmetric property in a hybrid registration framework. To the best of our knowledge, little efforts have been exerted towards this direction. Apart from Christiensen and his coworkers that have exploited both types of information in the inverse consistency framework, we have not come across any other method

performing a similar task. We have placed great emphasis on establishing accurate correspondences without undermining the symmetric property of the overall approach. The proposed method is formulated in a discrete setting and our effort was concentrated on delivering an efficient and versatile algorithm. We have to remark that the iconic part of our method bears a strong resemblance to the one proposed by Noblet [371]. The main difference is the discrete nature of the proposed method that allows us to tackle both monomodal and multi-modal registration.

4.3 Continuous Domain

Similarly to the previous chapter, we are going to begin with the description of the objective function in the continuous domain. Subsequently, the symmetric energy for the iconic part, the geometric and the coupling term are going to be given. The previous energy will be mapped to discrete deformation elements to facilitate the description of the MRF energy that will take place in the next subsection.

4.3.1 Iconic Registration

As we stated in the previous section, there is a number of ways to construct a symmetric energy. The simplest is by considering both forward and backward transformation and evaluating the cost function over both image domains:

$$E_{\text{ico}} = \mathcal{M}_{\text{ico}}(S \circ \mathcal{T}_{\text{ico},ST}, T) + \mathcal{R}_{\text{ico}}(\mathcal{T}_{\text{ico},ST}) + \mathcal{M}_{\text{ico}}(T \circ \mathcal{T}_{\text{ico},TS}, S) + \mathcal{R}_{\text{ico}}(\mathcal{T}_{\text{ico},TS}). \quad (4.1)$$

The subscripts denote the direction of the transformation. In [14], the energy was minimized separately for $\mathcal{T}_{\text{ico},ST}$ and $\mathcal{T}_{\text{ico},TS}$ and two penalties were added to impose the inverse consistency constraint (ICC):

$$\text{ICC} = \mathcal{D}(T_{\text{ico},ST}, T^{-1}_{\text{ico},TS}) + \mathcal{D}(T_{\text{ico},TS}, T^{-1}_{\text{ico},ST}). \quad (4.2)$$

$\mathcal{D}$ denotes a function that measures the difference between the two transformations. As stated before, this framework is only asymptotically symmetric. A slight different strategy would be to use as backward transformation directly the inverse of the forward, or

$$E_{\text{ico}} = \mathcal{M}_{\text{ico}}(S \circ \mathcal{T}_{\text{ico},ST}, T) + \mathcal{R}_{\text{ico}}(\mathcal{T}_{\text{ico},ST}) + \mathcal{M}_{\text{ico}}(T \circ \mathcal{T}^{-1}_{\text{ico},ST}, S) + \mathcal{R}_{\text{ico}}(\mathcal{T}^{-1}_{\text{ico},ST}). \quad (4.3)$$

In that way, the extra penalizing term is no longer needed.

The previous strategies are not able to serve our goal of efficiency. The major drawback with respect to that end is that the computational burden has doubled due to the evaluation of the criteria in both image domains. An ideal strategy would be one that would

need a minimal amount of additional operations, when compared to the standard pairwise registration, to calculate the cost function.

There are two possible schemes that lead to cost evaluations over one domain. The first would be to eliminate the asymmetry by considering forms defined on the graph of the warping function [363, 364], i.e.

$$\mathcal{M}_{\text{ico}} = \int_{\Omega_T} (S \circ \mathcal{T}_{\text{ico},ST}(\mathbf{x}) - T(\mathbf{x}))^2 (1 + det(J_{\mathcal{T}_{\text{ico},ST}}))d\mathbf{x}. \tag{4.4}$$

The disadvantage of this transformation is that we still need to calculate the Jacobian of the transformation. The other plausible way is to let both images deform and evaluate the matching term over the common domain Ω_C as done in [368, 61],

$$\mathcal{M}_{\text{ico}} = \int_{\Omega_C} \rho(S \circ \mathcal{T}_{\text{ico},SC}(\mathbf{x}), T \circ \mathcal{T}_{\text{ico},TC}(\mathbf{x}))d\mathbf{x}. \tag{4.5}$$

On the one hand, this strategy restricts the evaluation of the matching criterion over only the common domain. On the other hand, the number of parameters has doubled since two transformations have to be estimated.

There are two possible ways to keep the number of parameters stable. The first one is to consider one flow as proposed in [60, 77]. The second one is to employ two flows that obey such a relation that permits us to consider them as one [371]. Here, we take an approach similar to the one proposed in [371].

The symmetric energy is defined as follows:

$$E_{\text{ico}} = \mathcal{M}_{\text{ico}}(S \circ \mathcal{T}_{\text{ico},CS}, T \circ \mathcal{T}_{\text{ico},CT}) + \mathcal{R}_{\text{ico}}(\mathcal{T}_{\text{ico},CS}, \mathcal{T}_{\text{ico},CT}). \tag{4.6}$$

There are two important aspects of this scheme that have to be underlined. First of all, the transformations from the common domain and back to the respective image domains are considered. That is important as it will allow us under conditions to express the two flow problem in terms of only one flow. The second regards a requirement upon the properties of the mappings. They should be both invertible so that we can estimate at the end the total mapping from one image to the other. In this setting, the matching component of the energy can be expressed as

$$\mathcal{M}_{\text{ico}} = \int_{\Omega_C} \rho(S \circ \mathcal{T}_{\text{ico},CS}(\mathbf{x}), T \circ \mathcal{T}_{\text{ico},CT}(\mathbf{x}))d\mathbf{x}, \tag{4.7}$$

while the regularization one as

$$\mathcal{R}_{\text{ico}} = \frac{1}{2} \int_{\Omega_C} \psi(\mathcal{T}_{\text{ico},CS}(\mathbf{x})) + \psi(\mathcal{T}_{\text{ico},CT}(\mathbf{x}))d\mathbf{x}. \tag{4.8}$$

Let us now introduce the deformation model and show how the previous remarks can be taken into account before mapping the previous element to discrete deformation elements.

Similarly to the previous chapter and the registration algorithm proposed there, an interpolation based transformation model is employed. The difference is that now two transformations need to be estimated. As a consequence, two grids (G_{CS} and G_{CT}), one isomorphic to the other (copies of each other), are going to be used to model the transformation. Again, each grid consists of k control points distributed uniformly along the image domain. The dense displacement field is given by Eq. 3.5. Thus, the transformation fields are given in the following form:

$$\mathcal{T}_{\text{ico},CS}(\mathbf{x}) = \mathbf{x} + \sum_{i=1}^{k} \omega_i(\mathbf{x})\mathbf{d}_{i,CS} \text{ and } \mathcal{T}_{\text{ico},CT}(\mathbf{x}) = \mathbf{x} + \sum_{i=1}^{k} \omega_i(\mathbf{x})\mathbf{d}_{i,CT}. \tag{4.9}$$

As the grids are isomorphic, the respective nodes of each grid have the same spatial support, hence the spatial weight ω_i is the same. Their displacements however are different. The subscripts CT and CS are used to emphasize this difference.

As it has been discussed earlier, the requirement for invertibility can be easily satisfied by imposing hard constraints on the maximum displacement grid nodes are allowed to take. It is also with ease that we can halve the number of parameters to infer. The idea is to impose that each point in the common domain is halfway from the corresponding points in the image domains, or

$$\mathcal{T}_{\text{ico},CS}(\mathbf{x}) + \mathcal{T}_{\text{ico},CT}(\mathbf{x}) = 0 \Rightarrow \mathbf{u}_{\text{ico},CS}(\mathbf{x}) + \mathbf{u}_{\text{ico},CT}(\mathbf{x}) = 0 \Rightarrow \mathbf{d}_{i,CS} = -\mathbf{d}_{i,CT}. \tag{4.10}$$

As the previous equation suggests, one parameter has to be estimated for each node and depending on which deformation field is considered its positive or negative value will be used.

Assuming the previous deformation model, the discrete entities of the problem are the k pairs $(\mathbf{d}_{i,CS}, -\mathbf{d}_{i,CT})$. Expressing the symmetric matching term (Eq. 4.7) on the basis of these discrete entities will result in the following equation

$$\mathcal{M}_{\text{ico}} = \frac{1}{k} \sum_{i=1}^{k} \int_{\Omega_C} \hat{\omega}_i(\mathbf{x})\, \rho(S \circ \mathcal{T}_{\text{ico},CS}(\mathbf{x}), S \circ \mathcal{T}_{\text{ico},CT}(\mathbf{x}))d\mathbf{x}. \tag{4.11}$$

As for the regularization term (Eq. 4.8), it is expressed as

$$\mathcal{R}_{\text{ico}} = \frac{1}{2k} \sum_{i=1}^{k} \int_{\Omega_C} \hat{\omega}_i(\mathbf{x})\, \psi(\mathcal{T}_{\text{ico},CS}(\mathbf{x})) + \psi(\mathcal{T}_{\text{ico},CT}(\mathbf{x}))d\mathbf{x}, \tag{4.12}$$

or,

$$\mathcal{R}_{\text{ico}} = \frac{1}{k} \sum_{i=1}^{k} \int_{\Omega_C} \hat{\omega}_i(\mathbf{x})\, \psi(\mathcal{T}_{\text{ico},CS}(\mathbf{x}))d\mathbf{x}. \tag{4.13}$$

4.3.2 Geometric Registration

Opting for a symmetric registration framework, the geometric registration algorithm should also be endowed with such a property. Unfortunately, the scheme used in the previous chapter is not symmetric and cannot easily be adapted to meet our ends. Thus, we have to devise a new way to establish correspondences that will be independent of the order of the input images without compromising the overall accuracy. Moreover, given that the ultimate goal is to use the correspondences as an extra driving force for the hybrid registration, we request that they exhibit certain geometric constraints (smoothness of the implied dense deformation and no flipping of positions) and that they are dispersed in image space. Last but not least, only the most important correspondences should be kept, alleviating the computational burden and boosting the efficiency.

Here, we detail a three-step algorithm that is able to deliver results that conform with previous principles. The three steps that compose the proposed algorithm are i) landmark extraction; ii) matching by uniqueness constraints; and iii) refinement by considering geometric constraints.

Regarding the first step, for now, we are not going to concern ourselves with the detection step. We are going to assume, as before, that two set of landmarks (K and Λ) have been created either by employing a general-purpose algorithm such as SIFT [220] or a task-specific segmentation algorithm. We assume again that these landmark are placed in prominent locations that correspond to anatomical structures of interest. The detection step is by construction symmetric as the same algorithm treats both images equally.

On the contrary, matching is asymmetric. To account for this fact, we consider the following objective function

$$E = \mathcal{M}_{\text{geo}}(K \circ \mathcal{T}_{\text{geo}}, \Lambda) + \mathcal{M}_{\text{geo}}(\Lambda \circ \mathcal{T}_{\text{geo}}^{-1}, K). \tag{4.14}$$

Note that the transformation $\mathcal{T}_{\text{geo}}$ maps points from the source image domain to the target one. To compute it, a greedy search is employed. We first search exhaustively among the candidates belonging to the set Λ to find the best match $\lambda^\star$ for each $\kappa \in K$. Then, in a second pass, we search for the optimal matching $\kappa^\star$ for each $\lambda \in \Lambda$. At the end, we keep only these pairs that constitute a unique matching in both forward and backward fashion (uniqueness constraint). This strategy, apart from being symmetric, is robust to outliers.

The most interesting and novel part of this three-step process is the last refining one. So far, we have been able to establish correspondences by considering only a matching criterion. Considering only appearance information may not be sufficient to tackle ambiguous cases. To render the matching more robust, geometric information should be exploited. The third step of the proposed aims to introduce such information in order to refine the matching and keep the most reliable of the pairs.

The objective is threefold. First and foremost, we are interested in keeping only the most reliable of the previous correspondences. Fewer pairs of points lessen the computational load. Nonetheless, we have to ensure that meaningful correspondences are distributed along the whole image domain to provide the necessary constraints that will drive the registration. Last but not least, a geometric constraint to disambiguate erroneous matches is considered. Assuming that no scale differences exist between the two images, a distance preservation constraint is imposed.

The problem is cast as a clustering one. Let us assume that during the previous steps ζ pairs have been retained composing the set $Z = \{z_i = (\kappa_i, \lambda_i), i = 1, \cdots, \zeta\}$. The goal is to choose a number of Ξ cluster exemplars $\{c_\xi\}_{\xi=1}^{\Xi}$. Ideally, the number of clusters should be determined automatically by the clustering algorithm itself. That can be provided by state-of-the-art clustering algorithm such as Affinity-Propagation [372] and Linear-Programming based stabilities algorithm [350]. Here, we are going to use the second as its convergence is always guaranteed.

The problem is formulated as follows:

$$\min_{\Xi, \{c_\xi\}_{\xi=1}^{\Xi}} \left(\sum_{z \in Z} \min_{\xi} d^A(z, c_\xi) + \sum_{r=1}^{\Xi} \left(d^B(c_r, c_r) + \sum_{s=1}^{\Xi} d^\Gamma(c_r, c_s) \right) \right) \tag{4.15}$$

Each one of the three terms of the previous equation manifest a geometric constraint.

The first term, $d^A(p, c_\xi)$, requires all pairs belonging to the same cluster to be geometrically close to the exemplar of this cluster, encouraging nearby landmark pairs to have smooth displacements. It is defined as

$$d^A(z, c_\xi) = \frac{1}{2}\|\boldsymbol{\kappa}_z - \boldsymbol{\kappa}_{c_\xi}\| + \frac{1}{2}\|\boldsymbol{\lambda}_z - \boldsymbol{\lambda}_{c_\xi}\| + \|(\boldsymbol{\lambda}_z - \boldsymbol{\kappa}_z) - (\boldsymbol{\lambda}_{c_\xi} - \boldsymbol{\kappa}_{c_\xi})\|. \tag{4.16}$$

The subscript indexes pairs of correspondence. The previous energy can be interpreted as follows. Source points $\boldsymbol{\kappa}_z$ and $\boldsymbol{\kappa}_{c_\xi}$ between a pair z and its exemplar c_ξ should be spatially close. The same goes for the points belonging to the target domain, that is points $\boldsymbol{\lambda}_z$ and $\boldsymbol{\lambda}_{c_\xi}$ for the previous case. Moreover, the displacements $(\boldsymbol{\lambda}_z - \boldsymbol{\kappa}_z)$ and $(\boldsymbol{\lambda}_{c_\xi} - \boldsymbol{\kappa}_{c_\xi})$ that are implied by a pair z and its exemplar c_ξ should be similar.

The second term, $d^B(c_r, c_r)$, encourages cluster exemplars to be chosen as those pairs having high similarities. The assumption is that higher similarity implies higher reliability

for the matching and at the end is the most reliable correspondences that we would like to keep. Therefore it is defined as

$$d^B(c_r, c_r) = \varrho(\boldsymbol{\kappa}_{c_r}, \boldsymbol{\lambda}_{c_r}). \tag{4.17}$$

The third term, $d^\Gamma(c_r, c_s)$, encourages any two different exemplars c_r and c_s to be spatially distributed. To achieve such an effect, we opt to maximize their distance

$$d^\Gamma(c_r, c_s) = \frac{1}{\|(\boldsymbol{\kappa}_{c_r} - \boldsymbol{\kappa}_{c_s})\| + \|(\boldsymbol{\lambda}_{c_r} - \boldsymbol{\lambda}_{c_s})\|}. \tag{4.18}$$

Note that the number of clusters (hence the number of finally selected landmark pairs) Ξ is also optimized rather than being fixed. Contrary to the previous chapter that the iconic and the geometric registration problem were treated simultaneously, here geometric registration precedes the iconic one providing a set of correspondences to act as constraints. This set of correspondences, or displacements $\{\mathcal{T}_{\text{geo},\xi}, \xi = 1, \cdots, \Xi\}$, constitutes the set of discrete elements. Next, we will see how such an information can be taken advantage of during registration.

4.3.3 Coupling Term

Again, we would like to introduce the geometric correspondences as constraints to urge the dense deformation field to be locally consistent with the sparse displacements. The fact that now the correspondences are pre-estimated is a limiting factor. We argued that it is important to let the solution of both problems interact and profit from each other. Thus, it is important to add more degrees of freedom to our formulation to render it more robust to erroneous correspondences. This can be achieved by letting the algorithm decide upon whether to try to be consistent with the sparse correspondences or not.

Let us suppose Ξ pairs of landmark correspondences have been established composing a set $Z^\star = \{(\boldsymbol{\kappa}_i^\star, \boldsymbol{\lambda}_i^\star), i = 1, \cdots, \Xi\}$. For ease of notation, $\star$ symbol will be dropped hereafter. The displacements that occurred during the last step map a point in one image to a point in the other one. In this framework though, we aim to estimate displacements from the common domain to the image ones. Knowing the complete mapping and keeping in mind that the common domain is halfway from the corresponding points in the two images according to the Euclidean distance, the corresponding point $\mu \in \Omega_C$ as well as the displacements $\mathbf{d}_{\mu,\kappa}$ and $\mathbf{d}_{\mu,\kappa}$ can be calculated.

$$\mathbf{d}_{\mu\kappa} = \frac{\boldsymbol{\kappa} - \boldsymbol{\mu}}{2} \text{ and } \mathbf{d}_{\mu\lambda} = \frac{\boldsymbol{\lambda} - \boldsymbol{\mu}}{2}. \tag{4.19}$$

Keeping the previous in mind, the following holds

$$\begin{aligned}\mathcal{H}(\mathcal{T}_{\text{ico}}, \mathcal{T}_{\text{geo}}) &= \frac{1}{2}\mathcal{H}(\mathcal{T}_{\text{ico},CS}, \mathcal{T}_{\text{geo},CS}) + \frac{1}{2}\mathcal{H}(\mathcal{T}_{\text{ico},CT}, \mathcal{T}_{\text{geo},CT}) \\ &= \mathcal{H}(\mathcal{T}_{\text{ico},CS}, \mathcal{T}_{\text{geo},CS}) = \mathcal{H}(\mathcal{T}_{\text{ico},CT}, \mathcal{T}_{\text{geo},CT}).\end{aligned} \tag{4.20}$$

Thus, the coupling term can be defined as

$$\mathcal{H}(\mathcal{T}_{\text{ico},CS}, \mathcal{T}_{\text{geo},CS}) = \frac{1}{\Xi}\sum_{i=1}^{\Xi} \min(\|\mathcal{T}_{\text{ico},CS}(\boldsymbol{\mu}_i) - \mathcal{T}_{\text{geo},CS}(\boldsymbol{\mu}_i)\|, \Upsilon). \tag{4.21}$$

Equivalently, we can express the previous in terms of the transformation that maps from the common to the target domain. In the previous equation, Υ stands for a threshold whose goal is to discard those correspondences that are not at all consistent with the estimated dense deformation field. Its value varies depending on our certainty upon the solution of the previous geometric registration step.

Taking into account the grid-based transformation model, we can derive as previously a relation that consists of only two discrete entities

$$\mathcal{H}(T_{\text{ico},CS}, T_{\text{geo},CS}) = \frac{1}{\Xi}\sum_{i=1}^{\Xi} \min\left(\omega_j(\boldsymbol{\mu}_i)\|\mathbf{u}_{\text{geo},CS}(\boldsymbol{\mu}_i) - \mathbf{d}_{j,CS}\|, \Upsilon\right). \tag{4.22}$$

The above allows us to map the symmetric hybrid deformable registration problem into a discrete labeling one. The discrete entities of this problem are the set of control point displacements $\bar{\mathbf{D}} = (\mathbf{D}_{CS}, \mathbf{D}_{CT})$, the set of correspondences $Z^\star$ as well as the option to consider a correspondence or not.

4.4 Discrete Domain

To sum up the previous section, the energy that is optimized is the following:

$$E = \mathcal{M}_{\text{ico}}(S \circ \mathcal{T}_{\text{ico},CS}, T \circ \mathcal{T}_{\text{ico},CT}) + \mathcal{R}_{\text{ico},CS}(\mathcal{T}_{\text{ico},CS}) + \mathcal{H}(\mathcal{T}_{\text{ico},CS}, \mathcal{T}_{\text{geo},CS}). \tag{4.23}$$

The geometric registration is solved separately. In this section, the previously identified discrete entities are going to be mapped to MRF entities.

4.4.1 Iconic Registration

Let us denote again as $\mathcal{G}_{\text{ico}} = (\mathcal{V}_{\text{ico}}, \mathcal{E}_{\text{ico}})$ the graph modeling the symmetric iconic registration problem. The set of vertices $\mathcal{V}_{\text{ico}}$ corresponds again to the set of control point displacement updates, that is $|\mathcal{V}_{\text{ico}}| = |\Delta\bar{\mathbf{D}}| = |(\Delta\mathbf{D}_{CS}, \Delta\mathbf{D}_{CT})|$. The assignment of a label

$l_p \in \mathcal{L}_{\text{ico}}$ (with $p \in \mathcal{V}_{\text{ico}}$) is equivalent to an update $\Delta\mathbf{d}_{p,CS}$ on the control point $p \in G_{CS}$ as well as an update $\Delta\mathbf{d}_{p,CT}$ on the control point $p \in G_{CT}$, i.e. $l_p \equiv (\Delta\mathbf{d}_{p,CS}, \Delta\mathbf{d}_{p,CT})$ ($l_{p,CS} \equiv \Delta\mathbf{d}_{p,CS}$, $l_{p,CT} \equiv \Delta\mathbf{d}_{p,CT}$). So, the label set for this set of variables is $L_{\text{ico}} \subset \mathbb{R}^d \times \mathbb{R}^d$, i.e. a discrete version of the displacement space for both transformations so that the common space lies in equal distance from both image domains. The neighborhood system $\mathcal{E}_{\text{ico}}$ on the vertices is the same as the one used in the previous chapter. According to Eq. 4.7 we define the unary potentials as:

$$\mathcal{U}_{\text{ico},p}(l_p) = \int_{\Omega_C} \hat{\omega}_p(\mathbf{x})\, \rho(S \circ \mathcal{T}_{\text{ico},CS,l_{p,CS}}(\mathbf{x}), T \circ \mathcal{T}_{\text{ico},CT,l_{p,CT}}(\mathbf{x}))\, d\mathbf{x}, \tag{4.24}$$

where $\mathcal{T}_{\text{ico},CS,l_{p,CS}}$ denotes the transformation where a control point $p \in G_{CS}$ has been updated by $l_{p,CS}$. $\mathcal{T}_{\text{ico},CT,l_{p,CT}}$ is defined similarly. In the same way as in the previous chapter, region-based and statistical measures can be encoded based on local evaluation of the similarity measure in the common domain.

It goes without saying that the random variables are assumed to be conditionally independent resulting in the same efficient approximative scheme for the unary potential cost calculation. The unaries can be seen as encoding an overlapping blocks matching strategy. There is little difference with respect to the evaluation of the unary potential described in the previous chapter. The evaluation of the unary potentials for a label $l \in \mathcal{L}_{\text{ico}}$ corresponding to the updates $\Delta\mathbf{d}_{CS}$ and $\Delta\mathbf{d}_{CT}$ can be efficiently performed as follows. First, each image is globally translated by applying the respective update. In the overlapping domain, the unary potentials for this label and for all control points are calculated simultaneously. This result in a one pass through the common domain to calculate the cost and distribute the local energies to the graph nodes. The constrained transformation in the unary potentials is then simply defined as $\mathcal{T}_{\text{ico},CS,l_{p,CS}}(\mathbf{x}) = \mathcal{T}_{\text{ico},CS}(\mathbf{x}) + l_{p,CS}$, where $\mathcal{T}_{\text{ico},CS}(\mathbf{x})$ is the current or initial estimate of the transformation. From a computational point of view, the extra load is minimal and results from the fact that both images are transformed simultaneously. As a result, the number of interpolation operations that are performed is doubled. An insignificant cost to pay to symmetrize the algorithm.

If we compare Eq. 4.13 with Eq. 3.9, we will note that they are identical. That is why the exactly same pairwise potentials that were used in the previous discrete model can be used here too.

$$\mathcal{P}_{\text{elastic},CS,pq}(l_p, l_q) = \frac{\|(\mathbf{d}_{p,CS} + \Delta\mathbf{d}_{p,CS}) - (\mathbf{d}_{q,CS} + \Delta\mathbf{d}_{q,CS})\|}{\|\mathbf{p} - \mathbf{q}\|}. \tag{4.25}$$

The fluid-like behavior can be obtained by adopting the following pairwise potential that penalizes only the difference between the updates of the deformation field:

$$\mathcal{P}_{\text{fluid},CS,pq}(l_p, l_q) = \frac{\|l_{p,CS} - l_{q,CS}\|}{\|\mathbf{p} - \mathbf{q}\|}. \tag{4.26}$$

Recall that $l_{p,CS} \equiv \Delta \mathbf{d}_{p,CS}$.

Before moving to the coupling constraint, let us underline one more time the advantage of controlling the solution space that this framework provides. In order to construct the label set, the same iterative labeling scheme and search space refinement strategies as the ones in the previous chapter were used. The difference is that now the estimated transformations are required to be diffeomorphic. The control that we have over the solution space allows us to comply with such a constraint in a straight-forward manner by sampling displacements that don't exceed 0.4 of the grid spacing. As it has been discussed in various points in this manuscript, this is a sufficient hard constraint that guarantees that the deformation is both smooth and invertible.

4.4.2 Coupling Term

In order to model the coupling constraint a second graph $\mathcal{G}_{\text{cp}} = (\mathcal{V}_{\text{cp}}, \mathcal{E}_{\text{cp}})$ is considered. The set of the vertices $\mathcal{V}_{\text{cp}}$ corresponds to the decision whether to keep the correspondence or not (i.e. $|\mathcal{V}_{\text{cp}}| = \Xi$). Thus, the label set here is binary $\mathcal{L}_{\text{cp}} = \{0, 1\}$. Zero for no correspondence and one for the contrary. If we were to design the graph, each vertex would be placed in the location of its respective point μ in the common domain.

Such a decision can be based solely on the difference between the geometric and iconic transformation. Nonetheless, our certainty regarding the correspondence, as it is quantified by the matching criterion, can be taken into account. This leads to the following unary potential:

$$\mathcal{U}_{\text{cp},p}(l_p) = \begin{cases} \Upsilon_{\mathcal{U}} & \text{if } l_p = 0, \\ \varrho(\kappa_p, \lambda_p) & \text{if } l_p = 1. \end{cases} \tag{4.27}$$

The same edge system $\mathcal{E}_{\text{cp}}$, as in the asymmetric case, is employed to connect the iconic random variables to the geometric ones. The coupling penalty (Eq. 4.22) can be modeled by the following pairwise label assignment or:

$$\mathcal{P}_{\text{cp},pq}(l_p, l_q) = \begin{cases} \Upsilon_{\mathcal{P}} & \text{if } l_p = 0, \\ \omega_q(\boldsymbol{\mu}_p) \left\| \mathbf{u}_{\text{geo},CS}(\boldsymbol{\mu}_p) - (\mathbf{d}_{q,CS} + l_{q,CS}) \right\| & \text{if } l_p = 1. \end{cases} \tag{4.28}$$

$p \in \mathcal{V}_{\text{cp}}$ and $q \in \mathcal{V}_{\text{ico}}$, $l_p \in \mathcal{L}_{\text{cp}}$ and $l_q \in \mathcal{L}_{\text{ico}}$, and $(p, q) \in \mathcal{E}_{\text{cp}}$. Such a pairwise term couples the displacements given by the two registration processes and imposes consistency. However, if the difference is greater than a threshold, it would prefer not to impose any consistency and discard this correspondence. The coupled registration objective function (Eq. 4.23) is finally represented by an MRF graph $\mathcal{G}_{\text{hyb}} = (\mathcal{V}_{\text{cp}} \cup \mathcal{V}_{\text{ico}}, \mathcal{E}_{\text{cp}} \cup \mathcal{E}_{\text{ico}})$ with its associated unary and pairwise potential functions.

4.5 Experimental Validation

Validating the performance of the proposed framework is accomplished through the use of a commonly available database for optical flow[1]. This dataset contains image sequences for which the flow between successive frames is known. There are three different categories of image sequences. The fist category comprises challenging real data. The estimation of the ground truth flow is achieved through the use of hidden fluorescent texture. The second category consists of data that were synthesized through computer graphic techniques. The third category consists of modified stereo sequences. The interested reader may find more information concerning the data set in [352]. Four typical image sequences belonging to the data set along with the ground truth flow are depicted in Fig. 4.1.

4.5.1 Geometric Registration

As the geometric correspondence establishment is treated in an initial step and in a sense separately from the subsequent dense correspondence establishment step, we present in this section typical results obtained with our proposed framework.

Landmark detection is performed by using the SIFT detector [220]. Our choice was motivated by the widespread acknowledgement that the specific algorithm has gained in the Computer Vision community. Moreover, there is a number of efficient implementations of the algorithm available online[2] facilitating experimentation. Regarding the clustering step, a modified version of the Linear Programming clustering method presented in [350] was used.

In the first step of our algorithm, key-points are detected in both image domains. Subsequently, they are matched to their best candidate in the other image domain constructing symmetrically the initial set of correspondences. In the first row of Fig. 4.2 and Fig. 4.3 the initial set of correspondences for 4 different images of the testing data set is shown. The initial sets of correspondences comprise a great number of key-point pairs located in the proximity of each other. Moreover, there is an important number of false correspondences.

During the second step of the algorithm, we try to rid the set of correspondences from the false ones. This is achieved by imposing the unique matching constraint. In the second row of Fig. 4.2 and Fig. 4.3, the correspondences that have been kept after enforcing the uniqueness constraint are shown. While an important number of erroneous correspondences have been discarded, there are some that still remain. In the first row of the Fig. 4.2

[1] http://vision.middlebury.edu/flow/
[2] http://www.vlfeat.org/~vedaldi/code/sift.html

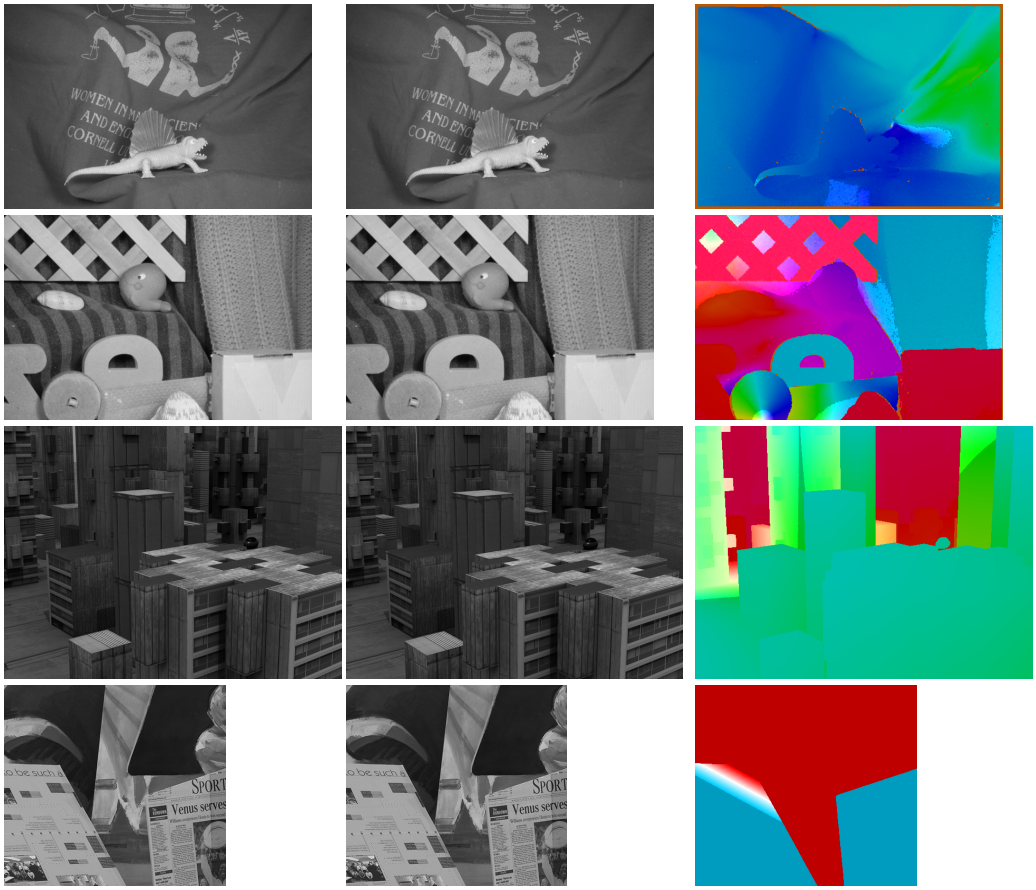

Figure 4.1: Four representative images of the data set. From left to right: the first frame of the sequence, the second frame of the sequence and the ground truth flow.

and Fig. 4.3 we have marked with a red square some examples of false correspondences that have been rejected during the second step of the process.

We remark that the correspondences that have remained after the previous step tend to be localized in space. Therefore, there is a certain redundancy in the information that is conveyed by the geometric correspondences. The third step of the proposed process aims to discard this redundancy and keep the best (according to a matching criterion) correspondences. The result of the clustering step is shown in the third line of Fig. 4.2 and Fig. 4.3. Let us point-out that the clustering step assigns the outliers that were not rejected after the second step to respective clusters. As the outliers are significantly different from the rest of key-point pairs, there is no redundancy in the information they convey and thus are treated as individual clusters. This fact motivates the use of a robust coupling term

during the registration. In the third row Fig. 4.2 and Fig. 4.3 a red square has been used to enclose some of the remaining erroneous correspondences.

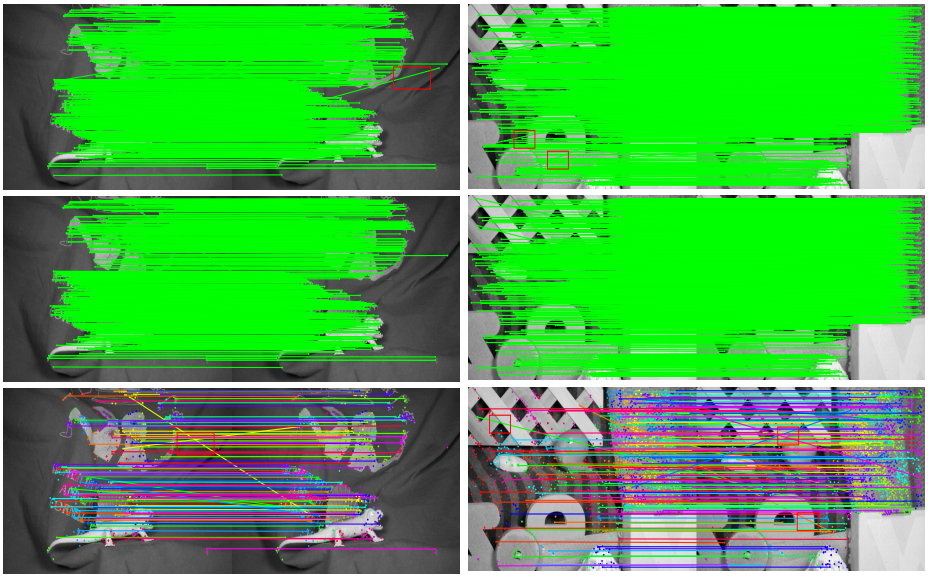

(a) Dimetrodon (b) RubberWhale

Figure 4.2: The result of the proposed three-step process to establish geometric correspondences. Top row: initial correspondences. Middle row: correspondences after enforcing the uniqueness constraints. Bottom row: correspondences after the clustering step.

4.5.2 Coupled Registration

The goal of this experiment is twofold. First and foremost, we would like to demonstrate the potential of including geometric information during the registration process. In other words, we would like to show that a hybrid approach may achieve better results than a purely iconic method even in this specific setting where iconic methods are known to perform with extreme accuracy. Moreover, we would like to show that we are able to discard erroneous correspondences through the proposed outlier handling strategy.

As the assumption that intensities are preserved in time is true in this specific case, we chose the Sum of Absolute Differences as similarity criterion. The weight for the pairwise term modeling the smoothness constraint was set equal to 30. The values the smoothing pairwise potential can take were truncated to 0.4 in order to allow abrupt changes in the labeling. This is essential due to the presence of different objects in the images that move

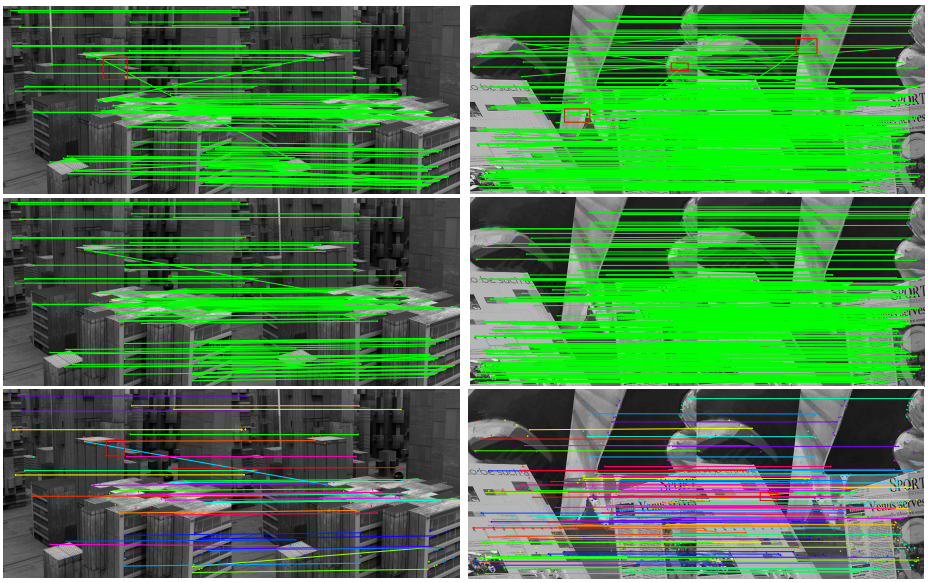

(a) Urban2 (b) Venus

Figure 4.3: The result of the proposed three-step process to establish geometric correspondences. Top row: initial correspondences. Middle row: correspondences after enforcing the uniqueness constraints. Bottom row: correspondences after the clustering step.

independently from one another. The parameters were kept the same for both methods. The weight for the pairwise potential encoding the coupling term was set to 17. The threshold Υ_p was defined in an adaptive way and was to 5 $\times$ the sampling step. Taking into account that the SIFT pairs that were considered during the registration exhibited high similarity, the term described in Eq. 4.27 was not used.

Cubic B-spline basis functions were used as interpolation functions in our grid-based deformation model. A coarse-to-fine strategy was considered only for the deformation grid. The initial spacing of the uniform grid was set equal to 16 and it was reduced to 8 and eventually 4 pixels as three levels were considered. The image resolution was kept to its finest level so that the cost calculation could be carried out with the highest possible precision.

As far as the label set is concerned, a dense sampling strategy was preferred and 10 samples were taken from zero to the maximum allowed displacement, resulting in a total number of 441 labels. In the $2D$ optical flow case, we can easily use a dense sampling strategy without being restricted by memory demands as opposed to the $3D$ case. Therefore, it is logical to opt for such a strategy that will enable us to approach the true solution

as much as possible. On each grid level, 5 iterations were performed considering different label sets. Each label set was a rescaled version of the previous one. The scaling factor that was used was set to 0.66.

As far as the results are concerned, the performance of the iconic method is in line with the results presented in [373] and in [374]. [373] presents the asymmetric version of the algorithm proposed here while [374] proposes an MRF formulation for the optical flow problem. Given the results presented here and in [373], we have reason to believe that they are close to the limit of accuracy one may expect for a parametric deformation model. The errors are within sub-pixel accuracy and mainly in the border of areas that move in a different sense. In these areas, the result of the algorithm is limited by the interpolation strategy we have chosen.

As far as the performance of the coupled method is concerned, it performed slightly better than the iconic method. In the cases where the intensity information was able to guide the registration with accuracy, its performance was either slightly better or equivalent to the performance of the iconic method. When that was not the case, the coupled method was able to achieve a significant improvement over both angular and end point error measures.

The difference between the quality of the solution of the two methods is evident in Fig. 4.6 where the areas of erroneous flow in the lower part of the image has greatly decreased. In Fig. 4.4, Fig. 4.5 and Fig. 4.7, we are not able to observe a notable difference between the flow fields of the iconic and coupled method as their performance is practically equivalent. The quantitative results for all image sequences are summarized in Table 4.1. We should also remark that in all cases, the proposed coupled method was able to discard the erroneous correspondences.

To conclude, geometric information may help improve the performance of a registration algorithm. Especially, in the cases where the iconic information fail partially to drive the registration process towards the desired result. The optical flow benchmark that we consider here may not be the best setting to evaluate the performance of the proposed approach as the errors are within a sub-pixel accuracy. We strongly believe that the impact of the geometric information will be more significant in the case of $3D$ medical registration where it could help remove the $3D$ ambiguity and recover large deformations.

4.6 Discussion

In this chapter, we proposed a framework that takes advantage of both geometric and iconic information to tackle the registration problem in a symmetric way. The problem was formulated as an energy minimization one and was modeled by a Markov Random

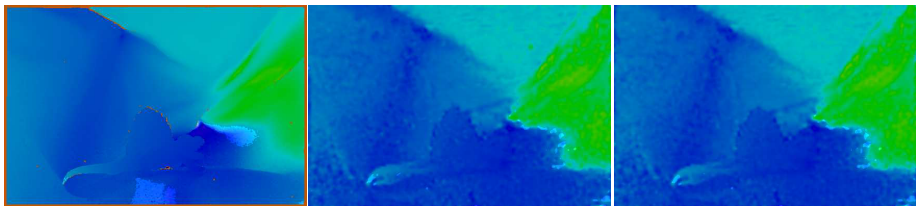

Figure 4.4: Optical flow visualization (Dimetrodon). Left column: ground truth flow. Middle column: flow obtained with the coupled method. Right column: flow obtained with the iconic method.

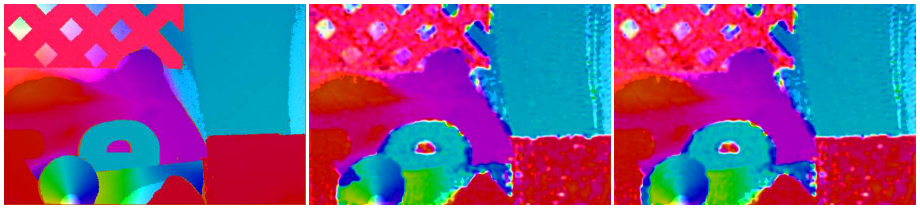

Figure 4.5: Optical flow visualization (RubberWhale). Left column: ground truth flow. Middle column: flow obtained with the coupled method. Right column: flow obtained with the iconic method.

Field. The interest of the proposed algorithm lies in the following:

- Its extreme efficiency. The problem is formulated so that as little as possible computational burden is added with respect to the asymmetric case. Moreover, an efficient approximation scheme and optimization strategy are used.
- Its versatility. Being based on a discrete framework the algorithm is modular with respect to the matching criterion and transformation model.
- To the best of our knowledge, no other hybrid registration method is endowed with

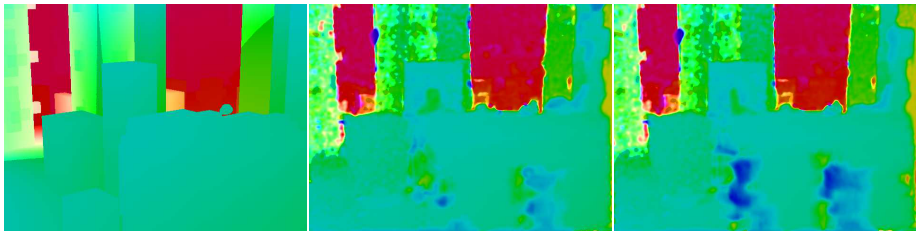

Figure 4.6: Optical flow visualization (Urban2). Left column: ground truth flow. Middle column: flow obtained with the coupled method. Right column: flow obtained with the iconic method.

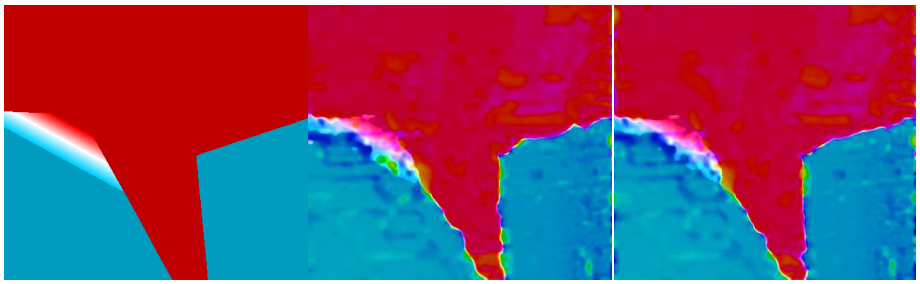

Figure 4.7: Optical flow visualization (Venus). Left column: ground truth flow. Middle column: flow obtained with the coupled method. Right column: flow obtained with the iconic method.

	Angular error (in degrees)				End point error (in mm)			
Image sequence	Iconic		Coupled		Iconic		Coupled	
	mean	std	mean	std	mean	std	mean	std
Dimetrodon	5.71	4.70	**5.68**	4.71	**0.28**	0.23	**0.28**	0.24
Grove2	3.92	6.84	**3.90**	6.92	**0.28**	0.44	**0.28**	0.44
Grove3	**7.88**	15.88	7.97	16.01	**0.82**	1.52	0.83	1.54
Hydrangea	3.73	6.55	**3.63**	6.45	**0.33**	0.49	0.33	0.51
RubberWhale	**6.65**	12.70	7.05	13.91	**0.20**	0.36	0.22	0.45
Urban2	7.95	12.60	**7.46**	12.50	1.51	3.01	**1.27**	2.50
Urban3	9.82	25.89	**8.16**	22.38	1.43	3.11	**1.21**	2.54
Venus	9.00	16.80	**8.97**	9.99	0.58	0.76	**0.56**	0.73

Table 4.1: Angular and end point error for the optical flow testing data set.

the symmetric property.

- The discrete formulation of the symmetric iconic registration is novel.
- The proposed strategy to establish geometric correspondences is novel itself and its demonstrated qualities are of great interest. Of particular interest is the third step as:
  - Decreases the number of correspondences keeping only the most important ones.
  - The number of the correspondences that are kept is not pre-defined but subject to optimization.
  - The chosen correspondences satisfy important geometric constraints (i.e. isometry).

  - Last but not least, the correspondences are chosen so that the whole image domain is covered in order to provide reliable guidance to the registration.

- The algorithm is robust to erroneous correspondences allowing for a decision upon whether to consider a given correspondence or not.

Despite all the interesting qualities of this algorithm, it is not easily extended to the case where a group of images is considered in order to perform a statistical study. In the next chapter, such a group-wise registration algorithm is going to be detailed.

Chapter 5

Group-wise Registration

5.1 Introduction

The previous algorithm is able to register pairs of images in an unbiased way. However, quite often, groups of images should be considered either to perform a group study or to learn a statistical model of anatomical variability. In such cases, a registration pre-processing step is essential to establish correspondences between all images. This problem is defined as *population* or *group-wise* registration.

Often, to solve the aforementioned problem, a reference frame is chosen and all population members are mapped to this pose using pair-wise registration algorithms. As it has been discussed already discussed in Chapter 4, the explicit selection of the reference image biases inherently the registration towards the chosen reference frame and influences its performance. Such a behavior goes in the opposite direction to the one expected towards appropriate representation of the population. Therefore, we would also like to perform such a task in an unbiased way. This chapter investigates the extension of the previous algorithm in the case of group-wise registration.

In the remainder of this chapter, we present first a review of the existing methods in section 5.2. Then, the proposed algorithm is detailed in two parts. In section 5.3, the first part details the formulation of the problem as an energy minimization one along with its decomposition to discrete deformation elements. Then, in the second part (section 5.4) the previous energy is modeled by a Markov Random Field. The experimental validation is presented in section 5.5, while section 5.6 concludes the chapter with a discussion.

5.2 Prior Art

Methods that try to tackle the group-wise registration problem in an unbiased way can be divided in two classes, *template-driven* and *template-free* ones.

The first class of methods can be further subdivided in two subclasses. In the first case, one image is selected as a template, either randomly or based on a criterion, and the rest of the population is registered to it. The template is initialized by the population but is not optimized afterwards. The bias is reduced either by the appropriate selection of the template or by post-processing the estimated deformation fields.

For example, Guimond *et al.* [375] register all images to a template belonging to the group. Once all images have been registered, an averaging is performed to yield an average deformation and an average intensity image. A model is constructed by applying the average deformation to the mean intensity image. Christensen *et al.* [376] proposed a variant of the previous method where inverse consistent pairwise registration is used and no averaging of the intensities is performed. The approach proposed by Seghers *et al.* [377] is closer in spirit to [375]. Instead of using only one image as a template, all images are used as templates. For each template, the rest of the images are registered to it resulting in a mean shape image after an averaging process. Then, all mean shape images are averaged in order to generate an atlas with mean intensity and mean morphology.

Bhatia and coworkers [124] may not choose explicitly a spatial reference domain but a member of the population is used as an intensity reference for the evaluation of an information theoretic similarity measure. The reference domain is defined by constraining the average deformation to be the identity transform.

The selection of the template can be performed in a more rigorous way than the one used in the previous methods. Park *et al.* [378] opt to decrease the bias by choosing the least biased target. In order to achieve that, they perform all possible pairwise registrations between the members of the population. Based on a distance defined as the bending energy of the estimated transformation, multidimensional scaling is performed to find an embedding of the data such that their distances are preserved. Then, the member of the population that is closest to the mean geometry is chosen as the optimal template.

In a similar setting, Hamm *et al.* [379] learn the empirical manifold that represents anatomical variation by performing a registration between all pairs of images. Having learned the manifold, the optimal template for group-wise registration can be defined as the median of the geodesic distances. Then, registration may be performed by following the geodesic path along the empirical manifold that connects the image with the template.

The choice regarding the template needs not be stable during the registration. In [94] an iterative group-wise algorithm is presented. All images are registered to a template with the use of polyharmonic clamped-plate splines transformation model. Then, based

on the warping distances, a new template is selected and the process continues till convergence. Clamped-plate splines were also used in [380] along with an Minimum Description Length criterion to perform group-wise registration. There, the reference image selection is regarded as a model choice.

In the second case, instead of choosing as template a member of the population, a template is constructed during the process. In the simplest case, the template is a mean model, however it can be more complex than that. These methods are based upon an iterative scheme where model estimation and pairwise registrations are interleaved.

A typical method of this category is the one proposed by Joshi *et al.* [368]. The algorithm iterates between the Fréchet mean estimation and deforming each image towards the template in the LDDMM framework. One drawback of this approach is that it is based on the mean estimation that is not robust to outliers. To tackle this shortcoming, the use of geometric median was proposed in [381]. A similar approach proposed by Avants and coworkers [382] estimates the template geodesic averaging in the manifold of diffeomorphisms. Instead of using a voxel-wise criterion to compare each image with the template, an information theoretic one is proposed in [383]. There, the matching is driven by the normalized mutual information between each image and the average intensity image of the population.

A similar approach where the average template is constructed during the registration process was proposed in [371]. Similar to [124], the reference domain is defined by enforcing the sum of all deformation fields to vanish as a hard constraint. The images are considered as an ordered set and each one is aligned with the use of a symmetric algorithm to the average intensity image of all the images that precede it. The previous scheme and the fact that the pairwise symmetric algorithm ensures that the common domain is halfway from the image domains guarantee that the constraint is satisfied.

One drawback of the previous approaches is that the constructed template is fuzzy. This fact that may hinder the performance of the algorithm. In [384], a group-wise registration method guided by a sharp mean image is presented. The group mean is created by adopting an adaptive weighting strategy where weights vary not only across subjects but also across spatial locations. As a consequence, the constructed template is sharp. Once it has been estimated, a minimum spanning tree is created with it as root and registration is performed in a pairwise fashion following the path from the leafs to the root. This step was inspired from [385] where shape instances were organized into a minimum spanning tree. The tree is constructed so that similar shape instances are connected by edges. Group correspondences can be estimated by performing fast and accurate pairwise registrations between neighboring instances towards the root. A similar algorithm was presented in [386] to tackle the diffusion tensor registration problem. The volume that is closest to the population center is chosen as root.

Apart from an intensity image, more complicated models may be used. In [387], an algorithm for automatic construction of active appearance models is proposed. The aim is to construct a model which can encode the set of images as efficiently as possible. The algorithm now iterates between estimating the model parameters and fitting it to the data. The model fitting takes account of the statistics of intensities and shapes across the group. At the end, implicit correspondences have been defined across the data set. In [139], similar to the previous, a model comprising appearance and shape information is used. The model is updated across the iterations and each image is deformed to match the model. In an earlier work, Cootes and coworkers [388] used a minimum message length criterion. An image is chosen as reference. Each of the rest of the images is left out in turn and a model is built by considering $N-1$ images. Then the transformation that minimizes the criterion is estimated and applied before the model is re-estimated.

Studholme *et al.* [389] inspired from multi-modal registration proposes to estimate the joint intensity distribution of all images and an entropy-based matching criterion to drive the registration. Moreover, geometric constraints penalizing displacement from the average shape along with a regularization constraint are considered as additional terms in the objective function.

One drawback of the previous methods lies in the selection of the template or model. Methods that don't require the explicit estimation of a model have been proposed in the literature. Maybe the most well-known representative of this class of methods is the congealing framework introduced in [390, 391]. There, the population of images is considered as a stack and a pixel-stack entropy measure is minimized to bring into correspondences the members of the population. In [392] the previous framework was applied to the group-wise affine registration of $3D$ volumes. Balci *et al.* [1] extended the congealing framework to the deformable case assuming a free-form deformation model. An efficient stochastic gradient descent optimization scheme along with a multi-resolution technique to diminish the computational burden were used. The univariate entropy was estimated with the use of a Parzen estimator. In [393], attribute vectors were used to guide the registration process. Huang *et al.* [394] also used richer descriptors (SIFT) to align complex images. Soft clustering was used to decrease the dimensionality of the problem.

The congealing framework is not the only option to tackle the problem in a template free fashion. For the problem of $3D$ ultrasound mosaicking, Wachinger *et al.* [395] proposed a method for the simultaneous rigid registration of all images using multivariate similarity measures defined by summing up bivariate ones. Later on, in [396] the definition of multivariate similarity measures by considering all possible combination of bivariate ones was studied further and connections to the congealing framework were drawn. The Efficient Second Order [31] was used to optimize the proposed energy. A similar implicit reference-based group-wise image registration algorithm was proposed by Geng *et al.* in

[397]. In their approach, the summation of the intensity difference between each pair of deformed images is minimized.

In [398], group-wise registration is solved in a hierarchical iterative way. Images are grouped and pair-wise registrations are performed between local neighbors. Once the registrations are performed, each subject is warped along an average direction defined by considering its neighbors. Then, clustering is performed, followed by pair-wise iterations. At each step, subjects move towards a common domain. The algorithm iterates till only one cluster is recovered or a maximum number of levels has been reached.

Baloch and Davatzikos in [399] proposed an approach to perform morphometric analysis based on a lossless descriptor comprising residual and transformation information. Each image is registered first to a random set of intermediate templates which are then mapped to a final one. For each registration, different parameters are used. This way, anatomical equivalence classes are constructed for each member of the population by considering the morphological appearance manifold. The images are registered in a group-wise fashion by moving along the manifolds while minimizing the group variance of the morphological descriptor. This approach has been extended in [400] by using non-linear approximations of the manifolds and allowing regionally optimal solutions. Another extension was presented in [401]. There, an affinity propagation clustering method was performed to identify close members of the population. The difference with respect to the previous methods lies in the suppression of random templates. Instead, representative intermediate ones are used, while the final template is chosen by learning the empirical manifold. This approach relies on a k-NN graph where any edge holds the geodesic distance between two images. Given this data structure, the template can be estimated as the geodesic median.

While the previous methods investigate the establishment of dense correspondences, there is also a growing literature on estimating sparse correspondences across a group of point sets or shapes. Davies *et al.* [402] perform group-wise registration across a set of shapes. The objective function is based upon the minimum description length principle. The idea is that the quality of the correspondences can be judged by considering the quality of the model built by these correspondences. A similar approach, for point-correspondences was proposed by Langs *et al.* [403].

Regarding the simultaneous registration of multiple point-sets, there is a number of methods that represent the point sets as finite mixtures of densities [255, 256, 404]. These methods aim to register the point sets without establishing explicitly correspondences through the minimization of a divergence measure upon the estimated densities. In [255] the Jensen-Shannon divergence is used, in [256] the generalized L_2-divergence is used while in [404] the Jensen-Renyi one. The last two bear the advantage to yield closed-form solutions.

As we have seen, most research efforts have been concentrated in the dense correspondence group-wise problem while fewer, though in increasing number, concern the sparse correspondence one. To the best of our knowledge, little attention have been paid in combining intensity and geometric information in a group-wise setting. Moreover, some remarks should be made regarding the methods dealing with group-wise dense registration. On the one hand, template-driven methods introduce bias to the process through the selection of the reference, and treat the examples of the population on an individual basis. On the other hand, template-free population-registration methods suffer from the lack of modularity with respect to the registration criterion and the deformation model. In addition, they are sensitive to initial conditions while being computationally inefficient.

In this chapter, we propose a graphical model approach to population registration. The discrete nature of the problem formulation enables different deformation models and matching criteria to be used in a straightforward manner. This fact allows us to introduce a principled objective function that considers both geometric and iconic information. The template-free method that is presented here bears resemblance with the methods proposed in [397, 396]. The transformation model is grid-based and the geometric constraint that is used to remove the bias is similar to the one used in [124, 1, 371].

5.3 Continuous Domain

As in the previous chapters, first we are going to describe the objective function in a continuous setting. The group-wise objective function is composed of two terms. The first term takes iconic information into account, while geometry is handled by the second. The coupling of the two terms is carried out through an adaptive weighting of the energy components. Once the energy construction is detailed, it will be mapped to discrete deformation elements

5.3.1 Iconic Registration

In order to introduce the concept of our approach, it is necessary to introduce some notation given that the nature of the problem is significantly different from the pairwise ones that we have studied up to this point. Let us consider n images $\{I_1, ..., I_n\}$. We can assume that each image is described by intensity values $I_i(\mathbf{x}_i)$ for different image domains $\Omega_i, \mathbf{x}_i \in \Omega_i$. Without loss of generality, a common reference frame can be defined (one should note that this assumption refers to the transformation domain definition and not to an image template). The aim of the mutual population deformation is to determine a transformation

$\bar{\mathcal{T}}$ which maps corresponding points from the common domain Ω_C to all n image domains

$$\bar{\mathcal{T}} = \{\mathcal{T}_i : \mathbf{x}_i = \mathcal{T}_i(\mathbf{x}_C), i = 1, ..., n\}. \tag{5.1}$$

such as an optimality criterion is satisfied. We are going to consider two optimality criteria that can be used either simultaneously or separately.

Notation-wise, we should remark that in this chapter, we are not going to use different subscripts to index the transformation for the iconic and the geometric case. The reason is that we take here an approach that is substantially different from the ones in the previous chapters. Instead of coupling the iconic and the geometric problem through a regularization term, the linking is accomplished by considering a hybrid matching term. Therefore, subscripts will be used only for the matching terms to differentiate them based on the type of information they exploit.

Population-wise Global Comparisons: Inspired from the congealing framework [390, 391], the first criterion that we consider is a global statistical one. Let $\pi(\mathbf{i}(\mathbf{x}))$ be the distribution of the corresponding sample intensities of the deformed images at the reference domain, or $\mathbf{i}(\mathbf{x}) = \{I_1 \circ \mathcal{T}_1(\mathbf{x}), \cdots, I_n \circ \mathcal{T}_n(\mathbf{x})\}$. In statistics, one can associate various measures of compactness with this density. We introduce the following global measurement towards population registration

$$\mathcal{M}_{\text{global,ico}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \int_{\Omega_C} \gamma(\pi(\mathbf{i}(\mathbf{x})))d\mathbf{x}, \tag{5.2}$$

with γ being a monotonic function inversely proportional to the compactness of the distribution at $\mathbf{x}$. For example, we can consider the entropy being estimated by means of histograms of intensity values.

The justification behind this criterion lies in the fact that as the images are mutually deformed towards alignment, the compactness of the probability distribution shall increase and the intensity values at corresponding coordinate locations from all the images will form a low entropy distribution. Examples of such global measurements can be the standard deviation, the skewness, the kurtosis, the Shannon entropy as considered in the case of congealing, or more advanced statistical compactness measurements. We have to note that employing this criterion requires a significant number of samples (images).

Pairwise Local Comparisons: The second criterion is based upon the comparison of pairs of images. Let $\rho_{i,j}(\cdot)$ be a similarity measurement between images I_i and I_j based on their intensities. Then, if we consider for example pixel-wise measurements, the pair

of deformations $\mathcal{T}_i$, $\mathcal{T}_j$, should minimize the distance between the two individual images in the reference domain:

$$\mathcal{M}_{\text{local.ico}}(\mathcal{T}_i, \mathcal{T}_j) = \int_{\Omega_C} \rho_{i,j}(I_i \circ \mathcal{T}_i(\mathbf{x}), I_i \circ \mathcal{T}_j(\mathbf{x}))d\mathbf{x}. \tag{5.3}$$

This criterion simply measures the similarity for every point of the reference pose between the two deformed images. It is actually the same matching criterion as the one that was used in the previous chapter for the symmetric iconic registration (Eq. 4.7).

This criterion can be extended to deal with the case of a group of n images simply by considering all possible pairs of images, or :

$$\mathcal{M}_{\text{local,ico}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \sum_{i=1}^{n} \sum_{j>i}^{n} \mathcal{M}_{\text{local,ico}}(\mathcal{T}_i, \mathcal{T}_j). \tag{5.4}$$

The intuition behind this term is that as the images are mutually aligned, the difference between any pair of population members will decrease. This term will introduce consistent pairwise deformations but will not introduce compactness in terms of statistical behavior for the observations once brought to the reference pose. Similar approaches have been investigated in [397, 396].

Likewise, the regularization term can be extended in the group-wise case in a straightforward manner:

$$\mathcal{R}(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \sum_{i=1}^{n} \psi(\mathcal{T}_i(\mathbf{x})) \tag{5.5}$$

where ψ, as previously, is a convex function imposing smoothness.

In order to map the previous energies to discrete deformation entities, a deformation model needs to be introduced. It is our firm choice to consider interpolation based transformation models. As a consequence, we opt for an extension of the grid-based model, that was used for the pairwise problem, to the group-wise case. A grid G_i superimposed over the image I_i will govern its deformation by manipulating the underlying mesh of k control points (k is assumed to be much smaller than the number of image points). Thus, the set of all grids $\bar{G} = \{G_i, \cdots, G_n\}$ is defined as the deformation model in this case. The central idea of our approach is to deform the grids simultaneously by considering different combinations of each grid's control point displacements $\mathbf{D}_i = \{\mathbf{d}_{i,1}, \ldots, \mathbf{d}_{i,k}\}$ so that the optimality criterion is minimized. In this case, the transformation of an image pixel can be written as

$$\mathcal{T}_i(\mathbf{x}) = \mathbf{x} + \mathbf{u}_i(\mathbf{x}) = \mathbf{x} + \sum_{j=1}^{k} \omega_j(\mathbf{x})\mathbf{d}_{i,j}. \tag{5.6}$$

We assume that the transformation is one-to-one and invertible. Moreover, to remove the bias, the following constraint should hold [124, 1]

$$\sum_{i=1}^{n} \mathcal{T}_i = 0. \tag{5.7}$$

The deformation grids are by construction isomorphic. Thus, the previous equation is equivalent to

$$\forall i : \sum_{j=1}^{k} \mathbf{d}_{i,j} = 0. \tag{5.8}$$

The control point displacements form the set of discrete deformation elements. The energies can now be expressed by taking them into account.

Population-wise Global Comparisons: The criterion described in Eq. 5.2 depends on the simultaneous deformation of all images resulting in higher-order relations between the discrete entities.

$$\mathcal{M}_{\text{global,ico}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \frac{1}{k} \sum_{j=1}^{k} \int_{\Omega_C} \hat{\omega}_j(\mathbf{x}) \gamma(\pi(I_1 \circ \mathcal{T}_1(\mathbf{x}), \cdots, I_n \circ \mathcal{T}_n(\mathbf{x})) d\mathbf{x}. \tag{5.9}$$

The discrete entity here is the set of the n simultaneous displacements of the grid nodes indexed by j in all deformation grids. Each node governs the displacement of the same point $\mathbf{x} \in \Omega_C$ but for a different image.

Pairwise Local Comparisons: Mapping this criterion to discrete entities is straightforward. In the same way as in the previous chapter, the energy can be expressed as

$$\mathcal{M}_{\text{local,ico}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \frac{1}{n(n-1)k} \sum_{i=1}^{n} \sum_{j>i}^{n} \sum_{r=1}^{k} \int_{\Omega_C} \hat{\omega}_r(\mathbf{x})\, \rho_{i,j}(I_i \circ \mathcal{T}_i(\mathbf{x}), I_j \circ \mathcal{T}_j(\mathbf{x})) d\mathbf{x}. \tag{5.10}$$

As for the regularization term, it is re-expressed as

$$\mathcal{R}(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \frac{1}{nk} \sum_{i=1}^{n} \sum_{j=1}^{k} \int_{\Omega_C} \hat{\omega}_j(\mathbf{x})\, \psi(\mathcal{T}_i(\mathbf{x})) d\mathbf{x}. \tag{5.11}$$

5.3.2 Geometric Registration

Establishing geometric correspondences between a group of images in an explicit way is not a trivial task. That is why we adopt an approach where sets of points are aligned without defining explicit correspondences between them. Such approaches have emerged either by representing point sets as probability distributions [255, 256, 404] or by using implicit functions [258, 259, 66, 68, 405].

There are numerous ways to represent geometric information. A popular one is to embed it in in higher dimensional space using a distance transform. Without loss of generality, let us consider that the available geometric information is given at the form of a shape. Then, the shape can be represented as

$$\Phi_g(\mathbf{x}) = \begin{cases} 0, & \mathbf{x} \in g, \\ +\mathcal{D}(\mathbf{x}, g), & \mathbf{x} \in \Omega_g, \\ -\mathcal{D}(\mathbf{x}, g), & \mathbf{x} \in \Omega - \Omega_g. \end{cases} \tag{5.12}$$

$\mathcal{D}$ refers to the minimum distance between the point $\mathbf{x}$ and the shape g, while Ω_g denotes the region enclosed by the shape.

The advantage of such an embedding is that it makes standard iconic approaches well-adapted to the geometric registration problem. Moreover, such an approach is robust to outliers and missing parts. Last but not least, it can be easily adapted to the case of landmarks or even more complex geometric information.

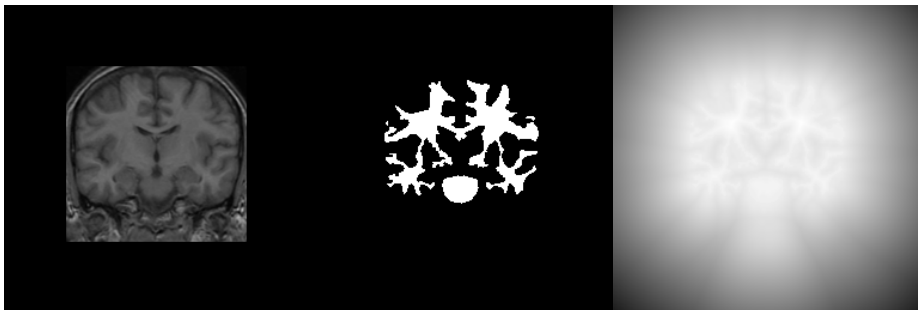

Figure 5.1: On the left, a slice of an MRI brain volume is depicted. In the middle, the corresponding white matter mask is shown. On the right, one can see its distance transform representation.

Applying the previous transformation to the geometric component of every image of the group will result in a second set of images $\bar{\Phi} = \{\Phi_1, \cdots, \Phi_n\}$ conveying all the geometric information of the images. Then, similar to the iconic case, the group of the geometric representation may be aligned simultaneously if we consider the pairwise local

comparisons

$$\mathcal{M}_{\text{local,geo}}(\mathcal{T}_i, \mathcal{T}_j) = \int_{\Omega_C} \varrho_{i,j}(\Phi_i \circ \mathcal{T}_i(\mathbf{x}), \Phi_j \circ \mathcal{T}_j(\mathbf{x}))d\mathbf{x}. \tag{5.13}$$

When considering all possible pairs of images and the discrete deformation elements (the same deformation model as in the iconic case is employed), the previous equation is rewritten as

$$\mathcal{M}_{\text{local,geo}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \frac{1}{n(n-1)k} \sum_{i=1}^{n} \sum_{j>i}^{n} \sum_{r=1}^{k} \int_{\Omega_C} \hat{\omega}_r(\mathbf{x})\, \varrho_{i,j}(\Phi_i \circ \mathcal{T}_i(\mathbf{x}), \Phi_j \circ \mathcal{T}_j(\mathbf{x}))d\mathbf{x}. \tag{5.14}$$

Essentially, the previous equation projects the similarity cost to pairs of simultaneous node displacements.

The global criterion could also be used here. Nonetheless, the nature of the cues and especially the fact that information is degraded away from the geometrical features, make its use inappropriate in this context.

5.3.3 Coupling Term

The two problems can be solved simultaneously, if the iconic and geometric terms are considered simultaneously.

$$\mathcal{M} = \mathcal{M}_{\text{ico}} + \mathcal{M}_{\text{geo}}, \tag{5.15}$$

or

$$\begin{aligned}\mathcal{M}_{\text{hyb}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) =& \frac{1}{n(n-1)k} \sum_{i=1}^{n} \sum_{j>i}^{n} \sum_{r=1}^{k} \int_{\Omega_C} \hat{\omega}_r(\mathbf{x})(\rho_{i,j}(I_i \circ \mathcal{T}_i(\mathbf{x}), I_j \circ \mathcal{T}_j(\mathbf{x})) + \\ & \varrho_{i,j}(\Phi_i \circ \mathcal{T}_i(\mathbf{x}), \Phi_j \circ \mathcal{T}_j(\mathbf{x})))d\mathbf{x}.\end{aligned} \tag{5.16}$$

There is nothing particularly wrong with the previous coupled energy. However, the distance transformation offers more information that can be exploited.

As we commented earlier, as the distance from the geometric feature increases, our confidence regarding the distance information decreases. Thus, it is only natural to trust more the geometric information in the vicinity of the geometric features or when the distance is close to zero. When the distance increases, more emphasis should be put on the iconic information.

We can account for this remark by deriving a weighted version of the previous equation as follows

$$\mathcal{M}_{\text{hyb}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \frac{1}{n(n-1)k} \sum_{i=1}^{n} \sum_{j>i}^{n} \sum_{r=1}^{k} \int_{\Omega_C} \hat{\omega}_r(\mathbf{x}) (\delta_\alpha(\mathbf{x}) \varrho_{i,j}(\Phi_i \circ \mathcal{T}_i(\mathbf{x}), \Phi_j \circ \mathcal{T}_j(\mathbf{x})) + (1 - \delta_\alpha(\mathbf{x})) \rho_{i,j}(I_i \circ \mathcal{T}_i(\mathbf{x}), I_j \circ \mathcal{T}_j(\mathbf{x}))) d\mathbf{x} \quad (5.17)$$

where $\delta_\alpha(\mathbf{x}) = \delta_\alpha(\min(\Phi_i \circ \mathcal{T}_i(\mathbf{x}), \Phi_j \circ \mathcal{T}_j(\mathbf{x})))$ is a smooth indicator function. The intuition behind the weighting scheme is that when at least one of the two pixels is close to the geometric feature we can really count heavily on the geometric term of the energy to drive the registration. In the rest of the cases, we assume that iconic information is richer and more reliable so we let it drive the process. The proposed framework can be thought of as a multi-channel registration one, where one channel contains intensity information and the other geometric one. The formulation presented here bears a strong resemblance to the approach we presented in [406] to fuse optical imaging and cinematic video.

Before proceeding with the discrete model, let us summarize this section. The group-wise registration problem was formulated both for the iconic and the geometric case. As far as the iconic case is concerned, two optimality criteria were discussed. The geometric case was treated similarly to the iconic after having embedded the geometric information in a higher dimensional space through a distance transformation. The coupling was achieved through an adaptive weighting of the two terms. Last but not least, the complete energy

$$E(\mathcal{T}_1, \cdots, \mathcal{T}_n) = \mathcal{M}_{\text{hyb}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) + \mathcal{M}_{\text{global,ico}}(\mathcal{T}_1, \cdots, \mathcal{T}_n) + \mathcal{R}(\mathcal{T}_1, \cdots, \mathcal{T}_n) \quad (5.18)$$

was expressed in terms of discrete deformation entities, that is the control point displacements of all the grids $\bar{\mathbf{D}} = \{\mathbf{D}_1, \cdots, \mathbf{D}_n\}$.

5.4 Discrete Domain

In this section, we detail the MRF model for the group-wise case. Based on the previous analysis, we are going to explain how to construct a graph to model the group-wise registration problem, what is the label set and how to the potential functions are defined.

5.4.1 Graph Construction

Our goal is to find which are the optimal displacements that should be applied to every node of every grid so that the optimality criterion is satisfied. As the discrete entities are

the same for both the iconic and the geometric case, a common graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ shall be used. The question is how create this graph?

The most evident way would be to apply the approach taken in the previous chapter to the group-wise case. In other words, create a graph whose vertices $\mathcal{V}$ correspond to the set of control point displacements of all n deformation grids, or $|\mathcal{V}| = |\bar{\Delta\mathbf{D}}|$. The assignment of a label $l_p \in \mathcal{L}$ would correspond to an update $\Delta\mathbf{d}_{i,p},\ p \in G_i,\ i = 1, \cdots, n$. The label set in this case is $\mathcal{L} \subset \underbrace{\mathbb{R}^d \times \cdots \times \mathbb{R}^d}_{n \text{ times}}$, or a discrete version of the displacement space of the transformations of all images. Moreover, it would even be possible to sample the space so that every label satisfies the constraint Eq. 5.8, i.e. $\sum_{i=1}^{n} \mathbf{d}_{i,p} = 0$.

There is one important shortcoming regarding this product model. It is inefficient. In order to sample the high dimensional space densely enough, a very large dataset has to be created. This has a direct impact on the computational load of the pairwise potentials calculations. Without loss of generality, let us consider the $2D$ case. Let us also assume that the same number of displacements L is sampled for every image. As no member of the population should be favored, this is a reasonable assumption. In this case the label set will be defined as $\mathcal{L} = \underbrace{L \times \cdots \times L}_{n \text{ times}}$. Given this label set, the complexity of the calculation of a pairwise potential is $O(|L|^{2n})$. It is possible to reduce this complexity by considering a decomposed model [323]. Let us note, however, that this advantage comes at the cost of a weaker relaxation.

The decomposed model in [323] is introduced for the $2D$ pairwise registration problem. Here, we extended it to deal with the group-wise case. Let us now detail the graph construction (see Fig. 5.2). The graph comprises n layers (as many as there are members in the population). Every layer i is composed of a set of nodes $\mathcal{V}_i$, with $\mathcal{V}_i$ isomorphic to G_i. Thus, the set of vertices is defined as $\mathcal{V} = \mathcal{V}_1 \cup \cdots \mathcal{V}_n$. As far as the edge system is concerned, there are two different types of edges that should be discerned, *intra*-layer edges and *inter*-layer ones. The intra-layer edge system consists of edges between nodes that belong to the same layer i or set of vertices $\mathcal{V}_i$. It includes multiple copies of the same neighborhood system. Each duplicate $\mathcal{E}_i$ is the same as the edge systems used in the previous chapters. The second type is composed of edges connecting nodes from different set of vertices, $\mathcal{E}_{ij} = \{(p, q) | p \in \mathcal{V}_i,\ q \in \mathcal{V}_j,\ p \sim q\}$ ($p \sim q$ denotes that these elements correspond to one another in the isomorphism $\sim$). The complete edge system is given as the composition of all the previous ones, or $\mathcal{E} = (\cup_{i=1}^{n} \mathcal{E}_i) \cup (\cup_{i=1}^{n} \cup_{j>i} \mathcal{E}_{ij})$. Last but not least, the label set is defined as $\mathcal{L} = L \subset \mathbb{R}^d$. As a result, the constraint Eq. 5.8 cannot be imposed through the construction of the solution set. To impose it, once the optimal labeling is estimated, it will be projected to the constraint surface.

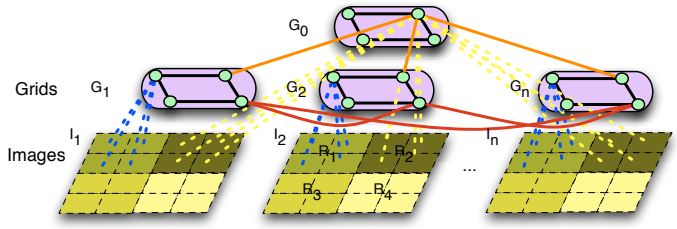

Figure 5.2: The node and the edge system of the constructed graph. With blue color the relationship between the grid nodes and the images is depicted (deformation model). The black edges represent the smoothness terms while the red ones encode the local dissimilarity measure. The global relationship between all the nodes at respective places in the grids is shown by the orange edges. Solid lines represent relations between the latent variables while the dashed lines depict the observed variables. (For clarity a fraction of the edges is shown.)

5.4.2 Iconic Criterion

Let us now define the unary and pairwise potentials that approximate the group-wise energy Eq. 5.18. First, we will explain how the global statistical compactness optimality criterion can be handled. The criterion described in Eq. 5.9 depends on the simultaneous deformation of all images resulting in higher-order relation between the discrete entities.

Inference in Markov Random Fields (MRFs) with higher order cliques is possible and a lot of research efforts have been devoted recently to investigate it (see for example [407, 408, 409]). However, the gain of the faithful modeling comes at the expense of significantly increased computational burden.

In order to spare ourselves the inefficiency of higher-order MRFs we will limit our approach to an approximation of the energy (Eq. 5.9). This approximation consists of assuming that when one image is deformed the rest stay still or in other words, come from a previous iteration. (Eq. 5.9). In practice, a *semi*-simultaneous scheme is employed.

Let us introduce a time variable t in order to express the energy in terms of the discrete entities while taking into account the previous assumption

$$\mathcal{M}_{\text{global,ico}}(\mathcal{T}_1^{t-1}, \cdots, \mathcal{T}_i^{t}, \cdots, \mathcal{T}_n^{t-1}) = \frac{1}{nk} \sum_{i=1}^{n} \sum_{j=1}^{k} \int_{\Omega_C} \hat{\omega}_j(\mathbf{x}) \gamma(\pi(\mathbf{i}^i(\mathbf{x}))) d\mathbf{x}, \quad (5.19)$$

where $\mathbf{i}^i(\mathbf{x}) = \{I_1 \circ \mathcal{T}_1^{t-1}(\mathbf{x}), \cdots, I_i \circ \mathcal{T}_i^{t}(\mathbf{x}), \cdots, I_n \circ \mathcal{T}_n^{t-1}(\mathbf{x})\}$. To ease the notation, the time variable t will be used only regarding this term. For the rest, we are able to derive simultaneous schemes and thus we are going to assume that all the considered quantities belong to the same iteration. The previous energy term can now be modeled with the use

of unary potentials defined as

$$\mathcal{U}_{\text{global,ico},i,p}(l_p) = \int_{\Omega_C} \hat{\omega}_p(\mathbf{x})\gamma(\pi(I_1 \circ \mathcal{T}_1^{t-1}(\mathbf{x}), \cdots, I_i \circ \mathcal{T}_{i,l_p}^{t}(\mathbf{x}), \cdots, I_n \circ \mathcal{T}_n^{t-1}(\mathbf{x}))d\mathbf{x}. \tag{5.20}$$

We should remark that in order to derive the previous unary potential, the conditional independence assumption has been extended to cover not only random variables within the same layer but also variables from different layers. This results in a cruder, with respect to the previous chapters, approximation. If the conditional independence was referring only to random variables within the same layer that would have as a consequence a higher-order relation between the random variables that belong to different layers. Whereas now, the cost of a deformation depends only on the label of the node p.

It should be noted that this assumption is not restrictive and is made quite often in the context of both group-wise registration and discrete optimization. On the one hand, in iterative registration schemes, models derived from the previous iteration are used to drive the (usually pairwise) deformation at the current iteration. In our case, the iterative scheme is devised in order to lessen the computational burden and does not occur as a part of an EM approach. On the other hand, in discrete optimization, it resembles the expansion moves used to minimize graphical models.

The most important advantage of this assumption is that it allows for an efficient approximative scheme for the unary cost calculation. The computations are performed following a strategy similar to the one presented in Chapter 3. The computation of the unary potentials of nodes belonging to the same layer is performed simultaneously. In order to evaluate the cost for a label $l \in \mathcal{L}$, the image that corresponds to the layer under consideration is translated by the displacement $\Delta\mathbf{d} \equiv l$. For every image element, the stack-wise statistical compactness measure is calculated. Then, the values, weighted by the $\hat{\omega}_p(\mathbf{x})$ are distributed to each control point p where they are summed to form the local energy. The previous process is iterated for all layers.

According to the previous assumption, the constrained transformation in the unary potentials is then simply defined as $\mathcal{T}_{i,l_p} = \mathcal{T}_i + l_p$, where $\mathcal{T}_i$ is the current or initial estimate of the transformation. Once more, the previous energy term can be interpreted as a weighted block matching one. The blocks being matched to the population of the images through the congealing framework. This term corresponds to the G_0 graphical model variables (see Fig. 5.2).

The case of the local matching criterion is easier to treat. Based on the Eq. 5.10, the pairwise potentials modeling the local iconic matching can be defined as

$$\mathcal{P}_{\text{ico},pq}(l_p, l_q) = \int_{\Omega_C} \hat{\omega}_p(\mathbf{x})\rho_{i,j}(I_i \circ \mathcal{T}_{i,l_p}(\mathbf{x}), I_j \circ \mathcal{T}_{j,l_q}(\mathbf{x})))d\mathbf{x}. \tag{5.21}$$

Thus, the matching cost is modeled by a pairwise relation, more specifically, by the inter-layer edges. Pairwise potentials are usually employed in order to model regularization terms. Nonetheless, there is no reason why they could not be used for a different purpose. This term is identical to the unary term Eq. 4.24. In order to calculate them, we consider the transformation $\mathcal{T}_{i,l_p} = \mathcal{T}_i + l_p$ and $\mathcal{T}_{j,l_q} = \mathcal{T}_j + l_q$. Images I_i and I_j are translated globally by the displacements $\Delta\mathbf{d}_i \equiv l_p$ and $\Delta\mathbf{d}_j \equiv l_q$, respectively. The overlapping domain is scanned calculating the cost that is then weighted and attributed to the respective edges.

5.4.3 Unified Criterion

The geometric term, also based on pairwise local comparisons (Eq. 5.14), is given by the following MRF pairwise potentials

$$\mathcal{P}_{\text{geo},pq}(l_p, l_q) = \int_{\Omega_C} \hat{\omega}_p(\mathbf{x}) \varrho_{i,j}(\Phi_i \circ \mathcal{T}_{i,l_p}(\mathbf{x}), \Phi_j \circ \mathcal{T}_{j,l_q}(\mathbf{x}))) d\mathbf{x}. \tag{5.22}$$

The pairwise potentials for the iconic and geometric case, Eq. 5.21 and Eq. 5.22, are modeled by the same set of edges. If we also consider the adaptive weight, a single pairwise potential can be used to model both

$$\begin{aligned}\mathcal{P}_{\text{hyb},pq}(l_p, l_q) = (1 - \delta_\alpha(\min(\Phi_i \circ \mathcal{T}_{i,l_p}(\mathbf{x}), \Phi_j \circ \mathcal{T}_{j,l_q})))\mathcal{P}_{\text{ico},pq}(l_p, l_q) \\ + \delta_\alpha(\min(\Phi_i \circ \mathcal{T}_{i,l_p}(\mathbf{x}), \Phi_j \circ \mathcal{T}_{j,l_q})\mathcal{P}_{\text{geo},pq}(l_p, l_q)\end{aligned} \tag{5.23}$$

The intra-layer edges model the regularization term Eq. 5.11. These pairwise potentials are defined as

$$\mathcal{P}_{\text{fluid},pq}(l_p, l_q) = \frac{\|l_p - l_q\|}{\|\mathbf{p} - \mathbf{q}\|}. \tag{5.24}$$

In the group-wise case we strongly encourage the use of this type of pairwise relation instead of the ones that have an elastic-like performance,

$$\mathcal{P}_{\text{elastic},pq}(l_p, l_q) = \frac{\|(\mathbf{d}_{i,p} + l_p) - (\mathbf{d}_{j,q} + l_q)\|}{\|\mathbf{p} - \mathbf{q}\|}. \tag{5.25}$$

This choice is dictated by memory and performance constraints. The former pairwise function depends only on the label set. Thus, it is the same for all edges belonging to the intra-layers. It can be computed only once and stored to be used for all such edges. On the contrary, since the latter depends on the current displacements $\mathbf{d}_{i,p}$ and $\mathbf{d}_{j,q}$, it is specific to each edge. In practice, we would face an important computational challenge if we were

to compute these terms on the fly. And alternatively, storing pre-computed values would pose serious memory issues.

To conclude, the total MRF energy that approximates the group-wise energy introduced in the previous setting (Eq. 5.18) has the following form

$$E_{MRF}(\mathbf{l}) = \sum_{i=1}^{n} \sum_{p \in \mathcal{V}_i} \mathcal{U}_{\text{global,ico},i,p}(l_p) + \sum_{i=1}^{n} \sum_{j>i} \sum_{pq \in \mathcal{E}_{i,j}} \mathcal{P}_{\text{hyb},pq}(l_p, l_q) + \sum_{i=1}^{n} \sum_{pq \in \mathcal{E}_i} \mathcal{P}_{\text{fluid},pq}(l_p, l_q). \tag{5.26}$$

5.5 Experimental Validation

Our validation setting is constructed upon the comparison of the proposed framework with the one proposed in [1]. The reason behind our choice is threefold. First, the entropy based deformable group-wise framework is considered state-of-the-art and has demonstrated repeatedly its effectiveness in population analysis [410]. Second, the source code is publicly available[1] allowing for a straightforward comparison. Third, it is based on a B-spline deformation model as our own framework. This is of utmost importance as it enables us to draw fair conclusions.

Two different settings were considered to compare the performance of the two algorithms. In the first place, a simulated data set was considered. The goal of the first set of experiments was principally to familiarize ourselves with the congealing group-wise framework in order to guarantee the fairness of the comparison. Thus, our discrete framework utilized only the available iconic information for this experiment. The simulated data set was constructed by applying random transformations to a T1-weighted brain MRI image of size $256 \times 256 \times 128$ and of physical resolution $0.9375 \times 0.9375 \times 1.5000$mm. The simulated data set is depicted in Fig. 5.3.

In the second set of experiments, a 3D brain MRI data set is used to validate the performance of the proposed framework. 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[2].

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. Images

[1] http://www.insight-journal.org/browse/publication/173

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

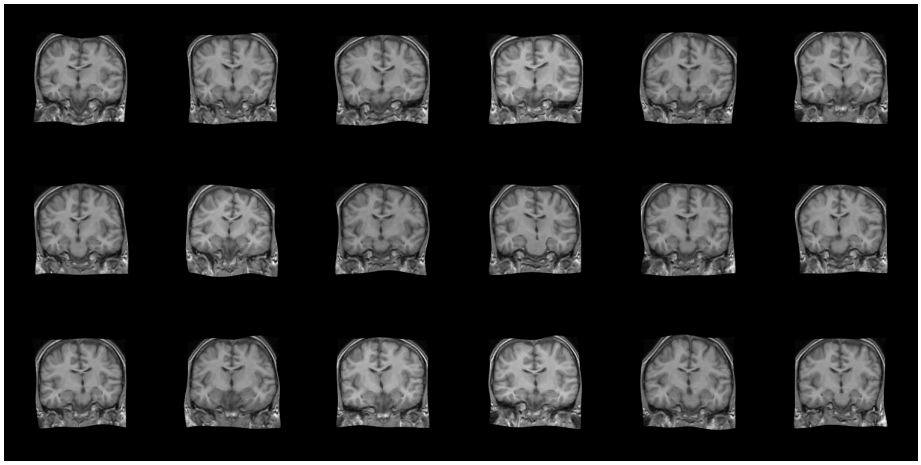

Figure 5.3: Simulated data set. The central slice from each member of the population is depicted.

intensities were normalized by applying the midway histogram equalization method [411]. In Fig. 5.4 and Fig. 5.5 the data set before and after the preprocessing step is depicted. The intensity normalization step is of great importance as it accounts for the great intensity differences that are present in the data set. For the hybrid framework, we used the available segmentation masks to derive a representation of the surface of white matter through the use of implicit functions. The derived distance representation of the white matter surfaces is shown in Fig. 5.6.

For all the experiments presented here, the Sum of Absolute Differences was used as similarity criterion be it intensities or distance values that were considered. As far as the regularization term is considered, as detailed in Sec. 5.4, the pairwise term with a fluid-like behavior was used. The weighting between the data term and the regularization one was set equal to 1.

As far as congealing group-wise method is concerned, stack-wise entropy was used as similarity criterion, estimated with the use of a Parzen window of width equal to 10. The default parameters that the authors of [1] propose were used. We experimented mainly with the number of samples used to evaluate the objective criterion. The affine registration step that precedes the deformable one was not used.

Cubic B-spline basis functions were used as interpolation strategy for our grid-based deformation model. While other interpolation strategies have been implemented and tested, the need to let the two competing methods share as many features as possible dictated this choice. For both methods, an hierarchical coarse-to-fine strategy was employed. Despite slight implementation differences, both methods start from coarse FFD control grids that are refined during subsequent iterations. For the proposed framework, the initial grid was

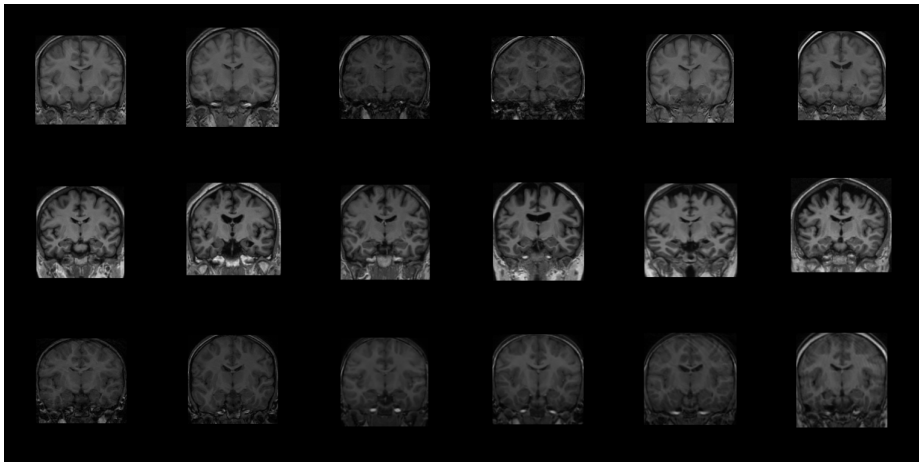

Figure 5.4: Data set before intensity normalization. The central slice from each member of the population is depicted.

of size $6 \times 6 \times 5$. Its size was approximately doubled at each finer level. For [1], a uniform deformation grid of size $6 \times 6 \times 6$ was employed. Its size was doubled at each finer level. For both methods, three levels were considered in total.

A multi-resolution method based on image pyramids was employed by both methods. The size of the images was scaled by a factor of 2 when passing from a coarse level to finer one. In our case, a standard Gaussian image pyramid was created by convolving the images with a Gaussian kernel with a standard deviation of 1. Both methods considered three pyramid levels.

The same strategy to define the discrete label set as the one presented in Sec. 3.5 was used. The deformation space was sampled sparsely along the positive and negative main axes of the 3D coordinate system. 5 samples along each direction were taken resulting in a total number of 31 labels. The maximum allowed displacement is sampled so that the resulting transformation is diffeomorphic. At each level, we iterated 5 times by considering the previous solution and refined label sets before continuing with the next, finer level. Each refined label set is created by rescaling the set of labels by a fixed factor of 0.66.

Let us first start with the results obtained for the simulated data set. In Fig. 5.7 the result for [1] and the proposed framework are shown for the central slice of 6 different volumes that were chosen randomly. For each volume, the respective central slice before registration is depicted to facilitate the comparison. The visual assessment of the result allows to conclude that both methods have achieved to register successfully the group of images. To further comprehend the quality of the result, we depict the central slice for both the mean intensity image and the standard deviation one for each case in Fig.

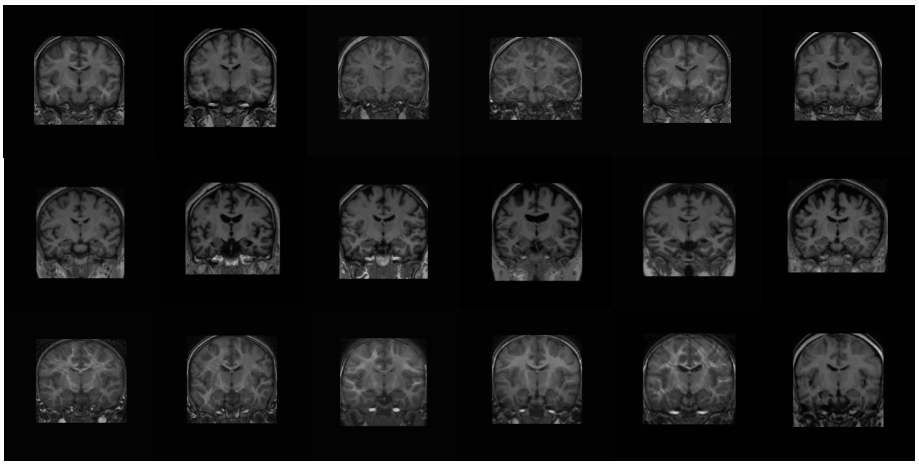

Figure 5.5: Data set after intensity normalization. The central slice from each member of the population is depicted.

5.8. The sharpness of the mean images suggests that both methods achieved a satisfactory result. When comparing the standard deviation images, we observe that both methods have improved with respect to the baseline. In general, [1] has achieved overall lower standard deviation values. In our case, comparatively greater errors can be observed in the borders of the object while lower ones in its interior. The differences are small and not significant. We have to remark that [1] is able to achieve such results in lower computational time using a low sampling rate for its stochastic gradient optimization scheme. Nonetheless, we have to underline that there is a certain bias as the data set has been created assuming exactly the same deformation model as in [1] rendering this experiment the best case scenario one may hope for [1].

In our second experiment that is a realistic scenario, we have compared the performance of the proposed framework (both iconic and unified variants) with [1] using two different sampling rates. One is a low sampling rate similar to the one used for the previous experiment while the second considers all image domain (sampling rate equal to 1). This case was considered towards rendering the comparison of the two methods unbiased with respect to the used image information.

The results of the registration for 6 different volumes are shown in Fig. 5.9. The first thing one can remark is that [1] with a low sampling rate fails to provide satisfactory results. In particular, the results for the forth and sixth volume prove the validity of the previous comment. Both volumes have been deformed in an unrealistic manner. In the forth volume, important anatomical structures have disappeared while in the sixth volume we observe a violent deformation of the skull. Increasing the sampling rate to consider

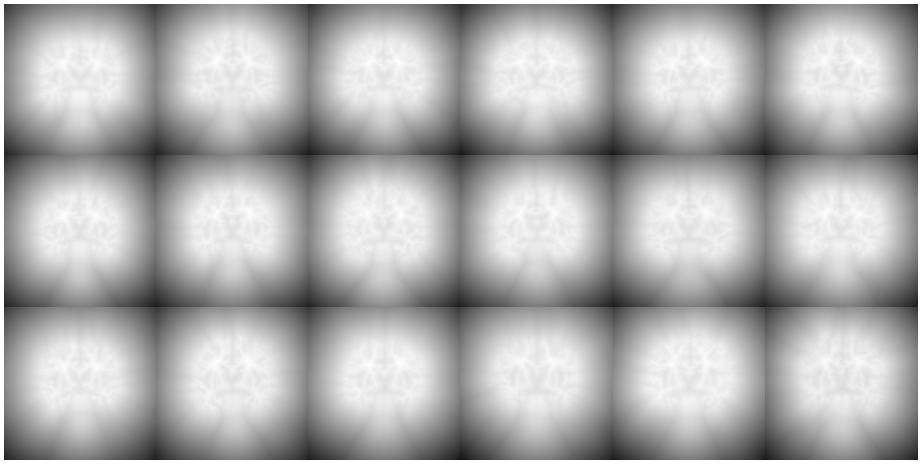

Figure 5.6: Distance maps created from the available white matter segmentation mask. The central slice from each member of the population is depicted.

all image domain did not produce unrealistic deformation fields but does not seem to register the group of images towards a common domain. One can argue that the parameters used for [1] were not the optimal ones which is probably true. Nonetheless, we kept the parameters of the proposed framework and [1] during both experiments.

Regarding the proposed framework, we observe an improvement of the results when taking into consideration the geometric information conveyed by the distance representation of the white matter surface. This improvement is more obvious when considering the mean and standard deviation intensity images depicted in Fig. 5.10. The mean intensity image of the result of our framework with the iconic criterion is sharper than the mean intensity image produced by [1]. We should also underline the increase of the sharpness of the mean intensity image when the unified criterion is used.

5.6 Discussion

In this chapter, we proposed an algorithm that exploits both iconic and geometric information to towards guiding group-wise registration. The novel framework to unite a population to an optimal (unknown) pose through their mutual deformation was cast as an energy minimization problem.

The registration criterion comprises three terms. The first imposes compactness on the appearance of the registered population at the pixel level. The second aspires to minimize the individual distances between all possible pairs of images based on an adaptively weighted hybrid criterion, while the last is a regularization one imposing smoothness on

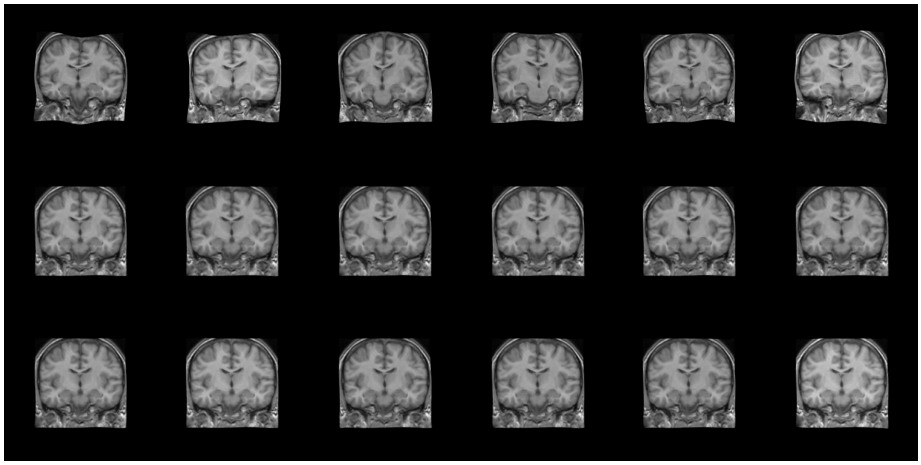

Figure 5.7: Central slice for 6 simulated volumes. Top row: images before registration. Middle row: images after registration with [1]. Bottom row: images after registration with the proposed framework.

the deformation fields. The problem is reformulated as a graphical model that consists of hidden (deformation fields) and observed variables (intensities and distances). This graphical model is expressed in the form of a first-order MRF.

The novelty of the approach lies principally in the formulation of the problem in a discrete setting. Therefore, our approach is gradient free, modular in terms of the image and smoothness components and can encode global population criteria and pair-wise comparisons. This modularity allows to investigate with ease a hybrid formulation for the group-wise registration problem.

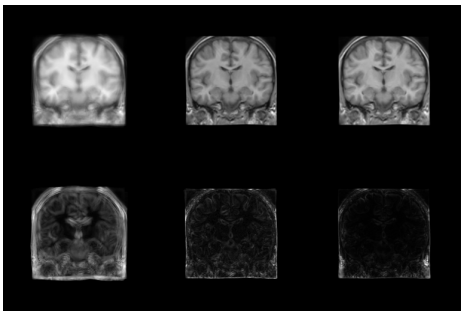

Figure 5.8: Central slice for the mean (top row) and standard deviation intensity image (bottom row). Left column: images before registration. Middle column: images after registration with [1]. Right column: images after registration with the proposed framework.

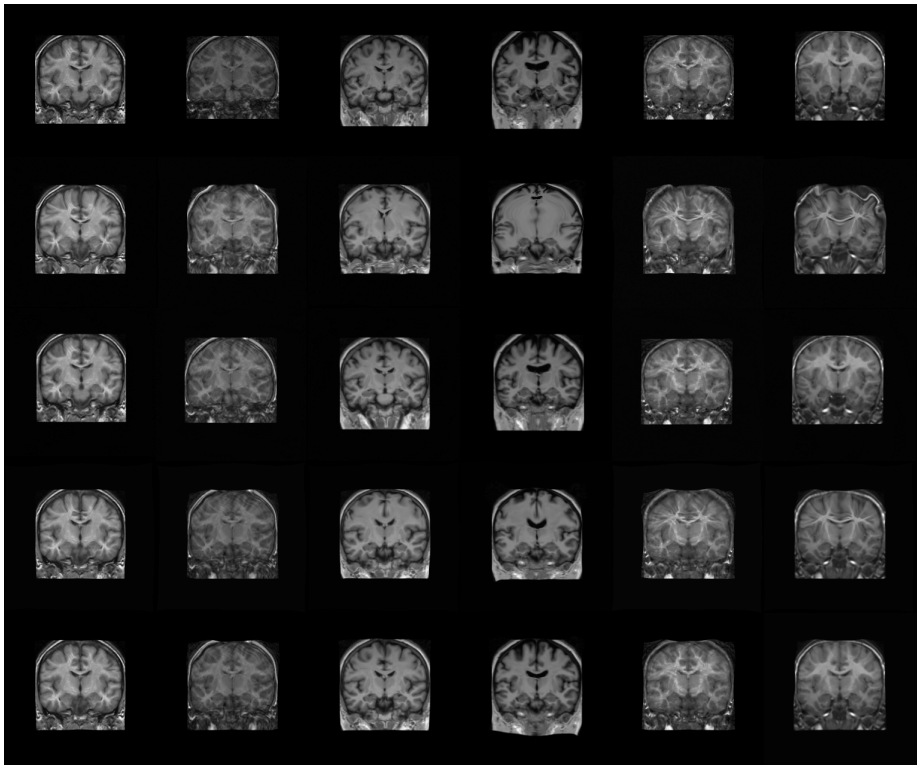

Figure 5.9: Central slice for the mean (top row) and standard deviation intensity image (bottom row). First row: images before registration. Second row: images after registration with [1] (low sampling rate). Third row: images after registration with [1] (sampling rate equal to 1). Forth row: images after registration with the proposed framework (only iconic criterion). Fifth row: images after registration with the proposed framework (unified criterion).

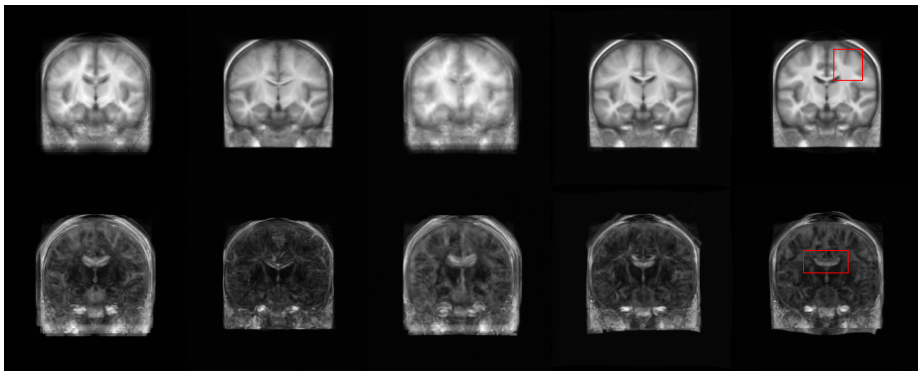

Figure 5.10: Central slice for the mean (top row) and standard deviation intensity image (bottom row). First column: images before registration. Second column: images after registration with [1] (low sampling rate). Third column: images after registration with [1] (sampling rate equal to 1). Forth column: images after registration with the proposed framework (only iconic criterion). Fifth column: images after registration with the proposed framework (unified criterion).

Chapter 6

Conclusion

We conclude this thesis by summarizing the majors contributions and drawing some potential directions for future work.

6.1 Contributions

The principal contribution of this thesis is the mapping of the iconic/geometric deformable registration problem into a unified discrete labeling one with interconnected variables. In this doctoral work, we provided discrete formulations for the problem of symmetric and group-wise registration. We detailed the construction of the graphical model for each case and presented efficient ways to evaluate the cost functions.

The hybrid registration framework was formulated as a two-layer graphical model. Intra-layer edges encoded smoothness constraints for each problem while inter-layer ones imposed a consistency constraint between the two, resulting in a solution that is a compromise of both worlds. The resulting formulation inherits modularity and scalability. Both the geometric and the iconic terms are modular in the sense that various detectors/similarity metrics can be considered and addressed under a common optimization framework.

The main message of this dissertation is twofold. On the one hand, the introduction of geometric information in an iconic dense deformable registration framework renders it more robust with respect to the initial conditions, large deformations and choice of matching criterion. On the other hand, the discrete formulations coupled with efficient but not restrictive approximate schemes could lead to powerful and versatile algorithms that can tackle with accuracy the registration problem.

Apart from the previous major contribution, the careful inspection of the proposed algorithms would reveal their advantages with respect to the state-of-the art. Such a dis-

cussion has taken place at each chapter. Thus, reciting them here would be redundant. Nonetheless, we would like to underline the interest of the geometric framework presented in Chapter 4 for it states and solves a problem little studied.

6.2 Future Work

During this thesis we have provided promising responses, yet various extensions are under consideration to push forward the scope of our work. Let us discuss briefly, three of them.

I) The success of our hybrid framework depends heavily on the quality of the geometric information. Refining further the geometric part of the framework seems a promising way to improve the obtained results. There are two possible extensions that could be attempted. Machine learning techniques can be applied to improve the localization of the points of interest. Better correspondences could be established by either devising discriminative task-specific matching criteria and richer descriptors, or by introducing more knowledge regarding the shape or the region of interest by adopting higher-order interactions.

II) Combined segmentation and registration is another promising direction. Atlas-based registration methods exhibit iconic similarity but fail to account for the expected geometric properties of the structure of interest. Such a constraint could be easily introduced within our formulation. Shapes/Organs of interest could be represented in the atlas space as point distribution models with pair-wise [349] or higher-order interactions [275]. Such a models can be endowed with statistical behavior through a learning step applied on their relative distances withing the training set. Segmentation and tracking can be coupled then through an introduction of an additional constraints that aims at maximizing the posterior of the point distribution model according to the pair-wise/higher order distributions. Such an idea can be further extended into a multi-atlas setting

III) Statistical interpretation of the results is another promising direction. Graph-based optimization could produce measures of uncertainties [314]. Their statistical interpretation could lead to more efficient ways of sampling the search space as suggested in [329]. However, we can imagine even going further towards dynamic construction of the uncertainties map and its direct use to define the transformation itself. In such a context, control-based deformations will refer to multi-dimensional Gaussian functions and the optimal estimation of the transformation will correspond to the mean and covariance matrix of this Gaussian. The use of kernel-density estima-

tion methods with variable bandwidth can then be used as basis for the deformation function.

IV) Coupling linear registration with deformable fusion is another promising direction of our work. Such a coupling requires the adoption of higher order interactions that can be easily encoded at the level of geometric registration. One can imagine a limited number of such interactions and a master-slave formulation for the resulting MRF, where one of the sub-problems will correspond to the pair-wise registration while the second to the higher order one. The aim of such a perspective will be to further eliminate the bias of the initial alignment and improve the robustness of the method. Such an action could also produce a natural separation between global pose differences and local deformations, a necessary step when aiming statistical modeling of variations within populations.

V) Finally, having conceived and realized unbiased registration algorithms, it is only natural to conduct group studies where their properties would be exploited to the maximum.

Publications of the Author

Journal Articles

- Yangming Ou, **Aristeidis Sotiras**, Nikos Paragios and Christos Davatzikos. *DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting.* Medical Image Analysis (MedIA), 2010.

- Ben Glocker[*], **Aristeidis Sotiras**[*], Nikos Komodakis and Nikos Paragios. *Deformable Medical Image Registration: Setting The State Of The Art With Discrete Methods.* Annual Reviews in Biomedical Engineering (AR), 2011.

Conference Articles

- **Aristeidis Sotiras**, Nikos Komodakis, Georg Langs and Nikos Paragios. *Atlas-based deformable mutual population segmentation.* IEEE International Symposium on Biomedical Imaging (ISBI), 2009.

- **Aristeidis Sotiras**, Nikos Komodakis, Ben Glocker, Jean-François Deux and Nikos Paragios. *Graphical Models and Deformable Diffeomorphic Population Registration Using Global and Local Metrics.* International Conference, Medical Image Computing and Computer Assisted Intervention (MICCAI), 2009.

- Mickaël Savinaud, **Aristeidis Sotiras**, Serge Maitrejean and Nikos Paragios. *Bioluminescence Enhancement Through Fusion of Optical Imaging and Cinematic Video Flow.* IEEE International Symposium on Biomedical Imaging (ISBI), 2010. **Oral presentation.**

- **Aristeidis Sotiras**[*], Radhouène Neji[*], Jean-François Deux, Nikos Komodakis, Gilles Fleury and Nikos Paragios. *A Kernel-based Graphical Model for Diffusion Tensor*

[*]The first two authors contributed equally to this article.

Registration. IEEE International Symposium on Biomedical Imaging (ISBI), 2010. **Best student paper.**

- **Aristeidis Sotiras**, Yangming Ou, Ben Glocker, Christos Davatzikos and Nikos Paragios. *Simultaneous Geometric - Iconic Registration.* International Conference, Medical Image Computing and Computer Assisted Intervention (MICCAI), 2010.

- Stavros Alchatzidis, **Aristeidis Sotiras** and Nikos Paragios. *Efficient Parallel Message Computation for MAP Inference.* IEEE International Conference on Computer Vision (ICCV), 2011.

Bibliography

[1] Serdar K Balci, Polina Golland, Martha Shenton, and William M. Wells III, "Free-Form B-spline Deformation Model for Groupwise Registration.", in *Statistical Registration Workshop: Medical Image Computing and Computer-Assisted Intervention : MICCAI'07*, Jan. 2007, number WS, pp. 23–30. 14, 15, 38, 62, 112, 114, 117, 125, 126, 127, 128, 129, 130, 131, 132

[2] Hajnal J, Hill DLG, and Hawkes DJ, Eds., *Medical Image Registration*, CRC Press, 2001. 24

[3] Jan Modersitzki, *Numerical Methods for Image Registration*, Oxford University Press Series: Numerical Mathematics and Scientific Computation, 2004. 24, 25

[4] Ardeshir Goshtasby, *2-D and 3-D Image Registration for Medical, Remote Sensing, and Industrial Applications*, Wiley, 2005. 24

[5] Lisa Gottesfeld Brown, "A Survey of Image Registration Techniques", *ACM Computing Surveys*, vol. 24, no. 4, pp. 1–60, 1992. 24, 44

[6] J.B. Antoine Maintz and Max A. Viergever, "A survey of medical image registration", *Medical Image Analysis*, vol. 2, no. 1, pp. 1 – 36, 1998. 24

[7] Barbara Zitova and Jan Flusser, "Image registration methods: a survey", *Image and Vision Computing*, vol. 21, no. 11, pp. 977 – 1000, 2003. 24

[8] Richard Szeliski, "Image alignment and stitching: A tutorial", Tech. Rep., Microsoft Research, 2006. 24

[9] M. Holden, "A review of geometric transformations for nonrigid body registration", *IEEE Transactions on Medical Imaging : TMI*, vol. 27, no. 1, pp. 111 –128, Jan. 2008. 24

[10] Chaim Broit, *Optimal Registration of Deformed Images*, PhD thesis, University of Pennsylvania, 1981. 26

[11] R Bajscy and S Kovacic, "Multiresolution Elastic Matching", *Computer Vision, Graphics and Image Processing*, vol. 46, pp. 1–21, 1989. 26

[12] James C Gee and Ruzena Bajcsy, "Elastic Matching : Continuum Mechanical and Probabilistic Analysis", *Brain Warping*, pp. 183–197, 1999. 26

[13] C Davatzikos, "Spatial transformation and registration of brain images using elastically deformable models.", *Computer Vision and Image Understanding : CVIU*, vol. 66, no. 2, pp. 207–22, May 1997. 26

[14] G E Christensen and Hans J Johnson, "Consistent image registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 20, no. 7, pp. 568–82, Jul. 2001. 27, 38, 41, 88, 91

[15] Alex Leow, Sung-Cheng Huang, Alex Geng, James Becker, Simon Davis, Arthur Toga, and Paul Thompson, "Inverse consistent mapping in 3D deformable image registration: its construction and statistical properties.", in *Information Processing in Medical Imaging : IPMI*, Jan. 2005, vol. 19, pp. 493–503. 27, 89

[16] Jianchun He and Gary E Christensen, "Large deformation inverse consistent elastic image registration.", in *Information Processing in Medical Imaging : IPMI*, C.J. Taylor and J.A. Noble, Eds. Jul. 2003, vol. 18, pp. 438–49, Springer-Verlag. 27, 88

[17] Xavier Pennec, Radu Stefanescu, Vincent Arsigny, Pierre Fillard, and Nicholas Ayache, "Riemannian elasticity: a statistical regularization framework for non-linear registration.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'05*, Jan. 2005, number 2 in LNCS, pp. 943–50. 28, 89

[18] John Ashburner, "A fast diffeomorphic image registration algorithm.", *NeuroImage*, vol. 38, no. 1, pp. 95–113, Oct. 2007. 28, 34, 62, 70

[19] Gary E. Christensen, Richard D. Rabbitt, and Michael I. Miller, "Deformable templates using large deformation kinematics.", *IEEE Transactions on Image Processing : TIP*, vol. 5, no. 10, pp. 1435–47, Jan. 1996. 29, 33, 41

[20] Gary E. Christensen, Sarang C. Joshi, and Michael I. Miller, "Volumetric transformation of brain anatomy.", *IEEE Transactions on Medical Imaging : TMI*, vol. 16, no. 6, pp. 864–77, Dec. 1997. 29

[21] M. Bro-Nielsen and Claus Gramkow, "Fast fluid registration of medical images", in *Visualization in Biomedical Computing*. 1996, pp. 265–276, Springer. 29, 32

[22] W R Crum, C Tanner, and D J Hawkes, "Anisotropic multi-scale fluid registration: evaluation in magnetic resonance breast imaging.", *Physics in medicine and biology*, vol. 50, no. 21, pp. 5153–74, Nov. 2005. 29

[23] Ming-Chang Chiang, Alex D. Leow, Andrea D. Klunder, Rebecca A. Dutton, Marina Barysheva, Stephen E. Rose, Katie L. McMahon, Greig I. De Zubicaray, Arthur W. Toga, and Paul M. Thompson, "Fluid registration of diffusion tensor images using information theory", *IEEE Transactions on Medical Imaging : TMI*, vol. 27, no. 4, pp. 442–456, 2008. 29, 89

[24] Emiliano D'Agostino, Frederik Maes, Dirk Vandermeulen, and Paul Suetens, "A viscous fluid model for multimodal non-rigid image registration using mutual information.", *Medical Image Analysis*, vol. 7, no. 4, pp. 565–75, Dec. 2003. 29

[25] Yongmei Wang and Lawrence H. Staib, "Physical model-based non-rigid registration incorporating statistical shape information.", *Medical Image Analysis*, vol. 4, no. 1, pp. 7–20, Mar. 2000. 29

[26] Jean-Philippe Thirion, "Image matching as a diffusion process: an analogy with Maxwell's demons", *Medical Image Analysis*, vol. 2, no. 3, pp. 243–260, 1998. 30

[27] Bernd Fischer and Jan Modersitzki, "Fast Diffusion Registration", in *Inverse problems, image analysis and medical imaging*, 2001, pp. 117–127. 30

[28] Xavier Pennec, Pascal Cachier, and Nicholas Ayache, "Understanding the "Demon's Algorithm": 3D Non-Rigid registration by Gradient Descent", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'99*, 1999, pp. 597–606. 30

[29] Tom Vercauteren, Xavier Pennec, Ezio Malis, Aymeric Perchant, and Nicholas Ayache, "Insight into efficient image registration techniques and the demons algorithm.", in *Information Processing in Medical Imaging : IPMI*, Jan. 2007, vol. 20, pp. 495–506. 30, 61

[30] Pascal Cachier, Eric Bardinet, Didier Dormont, Xavier Pennec, and Nicholas Ayache, "Iconic feature based nonrigid registration: the PASHA algorithm", *Computer Vision and Image Understanding*, vol. 89, no. 2-3, pp. 272–298, Mar. 2003. 30

[31] Ezio Malis, "Improving vision-based control using second-order minimization techniques", in *Internation Conference on Robotics & Automation*, Apr. 2004, pp. 1843–1848. 30, 61, 112

[32] Tom Vercauteren, Xavier Pennec, Aymeric Perchant, and Nicholas Ayache, "Non-parametric diffeomorphic image registration with the demons algorithm.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'07*, Jan. 2007, number Pt 2 in LNCS, pp. 319–26. 30, 61

[33] Tom Vercauteren, Xavier Pennec, Aymeric Perchant, and Nicholas Ayache, "Symmetric log-domain diffeomorphic Registration: a demons-based approach.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'08*, Jan. 2008, number Pt 1 in LNCS, pp. 754–61. 30, 31, 61, 87, 88

[34] Nathan D Cahill, J Alison Noble, and David J Hawkes, "Demons Algorithms for Fluid and Curvature Registration", in *2008 5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2009, pp. 730–733. 31

[35] Nathan D Cahill, J Alison Noble, and David J Hawkes, "A Demons algorithm for image registration with locally adaptive regularization.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'09*, Jan. 2009, number Pt 1 in LNCS, pp. 574–81. 31

[36] Tommaso Mansi, Xavier Pennec, Maxime Sermesant, Hervé Delingette, and Nicholas Ayache, "iLogDemons: A Demons-Based Registration Algorithm for Tracking Incompressible Elastic Biological Tissues", *International Journal of Computer Vision*, vol. 92, no. 1, pp. 92–111, Nov. 2011. 31, 43, 61

[37] Radu Stefanescu, Xavier Pennec, and Nicholas Ayache, "Grid powered nonlinear image registration with locally adaptive regularization.", *Medical Image Analysis*, vol. 8, no. 3, pp. 325–42, Sep. 2004. 31

[38] Alexandre Guimond, Alexis Roche, Nicholas Ayache, and Jean Meunier, "Three-Dimensional Multimodal Brain Warping Corrections", *IEEE Transactions on Medical Imaging : TMI*, vol. 20, no. 1, pp. 58–69, 2001. 31

[39] Marc Modat, Gerard R. Ridgway, Zeike A. Taylor, Manja Lehmann, Josephine Barnes, David J. Hawkes, Nick C. Fox, and Sébastien Ourselin, "Fast free-form deformation using graphics processing units", *Computer Methods and Programs in Biomedicine*, vol. 98, no. 3, pp. 278 – 284, 2010, HP-MICCAI 2008. 31

[40] Antonio Tristán-Vega, Gonzalo Vegas-Sánchez-Ferrero, and Aja-Fernández Santiago, "Local Similarity Measures for Demons-like Registration Algorithms", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2008, pp. 1087–1090. 31

[41] B T Thomas Yeo, Tom Vercauteren, Pierre Fillard, Jean-Marc Peyrat, Xavier Pennec, Polina Golland, Nicholas Ayache, and Olivier Clatz, "DT-REFinD: diffusion tensor registration with exact finite-strain differential.", *IEEE Transactions on Medical Imaging : TMI*, vol. 28, no. 12, pp. 1914–28, Dec. 2009. 31, 61

[42] Jean-Marc Peyrat, Hervé Delingette, Maxime Sermesant, Xavier Pennec, Chenyang Xu, and Nicholas Ayache, "Registration of 4D time-series of cardiac images with multichannel Diffeomorphic Demons.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'08*, Jan. 2008, number Pt 2 in LNCS, pp. 972–9. 31

[43] B T Thomas Yeo, Mert R Sabuncu, Tom Vercauteren, Nicholas Ayache, Bruce Fischl, and Polina Golland, "Spherical demons: fast diffeomorphic landmark-free surface registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 29, no. 3, pp. 650–68, Mar. 2010. 31, 61

[44] Bernd Fischer and Jan. Modersitzki, "Curvature based image registration", *Journal of Mathematical Imaging and Vision*, vol. 18, no. 1, pp. 81–85, 2003. 32

[45] Bernd Fischer and Jan. Modersitzki, "A unified approach to fast image registration and a new curvature based registration technique", *Linear Algebra and its Applications*, vol. 380, pp. 107–124, Mar. 2004. 32

[46] Stefan Henn, "A Full Curvature Based Algorithm for Image Registration", *Journal of Mathematical Imaging and Vision*, vol. 24, no. 2, pp. 195–208, Jan. 2006. 32

[47] Björn Beuthien, Ali Kamen, and Bernd Fischer, "Recursive Green's Function Registration", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'10*, 2010, pp. 546–553. 32

[48] Paul Dupuis and Ulf Grenander, "Variational problems on flows of diffeomorphisms for image matching", *Quarterly of applied mathematics*, vol. LVI, no. 3, pp. 1–20, 1998. 33

[49] A Trouvé, "Diffeomorphisms groups and pattern matching in image analysis", *International Journal of Computer Vision*, vol. 28, no. 3, pp. 213–221, 1998. 33

[50] M. Faisal Beg, Michael I. Miller, Alain Trouvé, and Laurent Younes, "Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms", *International Journal of Computer Vision*, vol. 61, no. 2, pp. 139–157, Feb. 2005. 33, 60

[51] S C Joshi and M I Miller, "Landmark matching via large deformation diffeomorphisms.", *IEEE Transactions on Image Processing : TIP*, vol. 9, no. 8, pp. 1357–70, Jan. 2000. 33

[52] Stephen Marsland and Carole J. Twining, "Constructing diffeomorphic representations for the groupwise analysis of nonrigid registrations of medical images.", *IEEE Transactions on Medical Imaging : TMI*, vol. 23, no. 8, pp. 1006–20, Aug. 2004. 33, 35

[53] Michael I Miller, Alain Trouve, and Laurent Younes, "On the metrics and Euler-Lagrange equations of computational anatomy.", *Annual review of biomedical engineering*, vol. 4, pp. 375–405, Jan. 2002. 33

[54] Michael I. Miller and Laurent Younes, "Group Actions , Homeomorphisms , and Matching : A General Framework", *International Journal of Computer Vision*, vol. 41, no. 1, pp. 61–84, 2001. 33, 61

[55] Alain Trouvé and Laurent Younes, "Metamorphoses Through Lie Group Action", *Foundations of Computational Mathematics*, vol. 5, no. 2, pp. 173–198, Feb. 2005. 33

[56] Michael I Miller, Alain Trouvé, and Laurent Younes, "Geodesic Shooting for Computational Anatomy.", *Journal of mathematical imaging and vision*, vol. 24, no. 2, pp. 209–228, Jan. 2006. 33

[57] Laurent Younes, Anqi Qiu, Raimond L Winslow, and Michael I Miller, "Transport of Relational Structures in Groups of Diffeomorphisms.", *Journal of mathematical imaging and vision*, vol. 32, no. 1, pp. 41–56, Sep. 2008. 33

[58] M Vaillant, M I Miller, L Younes, and A Trouvé, "Statistics on diffeomorphisms via tangent space representations.", *NeuroImage*, vol. 23 Suppl 1, pp. S161–9, Jan. 2004. 33, 60

[59] Laurent Younes, Felipe Arrate, and Michael I Miller, "Evolutions equations in computational anatomy.", *NeuroImage*, vol. 45, no. 1 Suppl, pp. S40–50, Mar. 2009. 33

[60] Mirza Faisal Beg and Ali Khan, "Symmetric data attachment terms for large deformation image registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 26, no. 9, pp. 1179–89, Sep. 2007. 33, 90, 92

[61] Brian B. Avants, C L Epstein, M Grossman, and James C. Gee, "Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain.", *Medical Image Analysis*, vol. 12, no. 1, pp. 26–41, Feb. 2008. 33, 44, 90, 92

[62] John Ashburner and Karl J Friston, "Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation.", *NeuroImage*, vol. 55, no. 3, pp. 954–67, Apr. 2011. 33, 61

[63] Yan Cao, Michael I Miller, Raimond L Winslow, and Laurent Younes, "Large deformation diffeomorphic metric mapping of vector fields.", *IEEE Transactions on Medical Imaging : TMI*, vol. 24, no. 9, pp. 1216–30, Sep. 2005. 33, 60

[64] Yan Cao, Michael I Miller, Susumu Mori, Raimond L Winslow, and Laurent Younes, "Diffeomorphic Matching of Diffusion Tensor Images.", in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, Jul. 2006, p. 67. 33

[65] Joan Glaunès, Marc Vaillant, and Michael I. Miller, "Landmark Matching via Large Deformation Diffeomorphisms on the Sphere", *Journal of Mathematical Imaging and Vision*, vol. 20, no. 1/2, pp. 179–200, Jan. 2004. 33, 54, 61

[66] A Leow, R. Thompson, H. Protas, and S.C. Huang, "Brain warping with implicit representations", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2004, pp. 603–606, IEEE. 33, 56, 118

[67] J. Glaunes, A. Trouve, and L. Younes, "Diffeomorphic matching of distributions: a new approach for unlabelled point-sets and sub-manifolds matching", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. 2004, pp. 712–718, Ieee. 33, 55

[68] Alex Leow, C L Yu, S J Lee, S C Huang, H Protas, R Nicolson, K M Hayashi, Arthur W. Toga, and Paul M. Thompson, "Brain structural mapping using a Novel hybrid implicit/explicit framework based on the level-set method.", *NeuroImage*, vol. 24, no. 3, pp. 910–27, Feb. 2005. 33, 56, 89, 118

[69] Marc Vaillant and Joan Glaunès, "Surface matching via currents.", in *Information Processing in Medical Imaging : IPMI*, Jan. 2005, vol. 19, pp. 381–92. 33

[70] L Risser, F Vialard, R Wolz, M Murgasova, D Holm, and D Rueckert, "Simultaneous Multiscale Registration using Large Deformation Diffeomorphic Metric Mapping.", *IEEE Transactions on Medical Imaging : TMI*, , no. c, pp. 1–14, Apr. 2011. 33

[71] Anqi Qiu, Laurent Younes, Lei Wang, J Tilak Ratnanather, Sarah K Gillepsie, Gillian Kaplan, John Csernansky, and Michael I Miller, "Combining anatomical manifold information via diffeomorphic metric mappings for studying cortical thinning of the cingulate gyrus in schizophrenia.", *NeuroImage*, vol. 37, no. 3, pp. 821–33, Sep. 2007. 33

[72] Joan Glaunès, Anqi Qiu, Michael I Miller, and Laurent Younes, "Large Deformation Diffeomorphic Metric Curve Mapping.", *International journal of computer vision*, vol. 80, no. 3, pp. 317–336, Dec. 2008. 33

[73] Stanley Durrleman, Xavier Pennec, Alain Trouvé, Paul Thompson, and Nicholas Ayache, "Inferring brain variability from diffeomorphic deformations of currents: an integrative approach.", *Medical Image Analysis*, vol. 12, no. 5, pp. 626–37, Oct. 2008. 33, 53

[74] Stanley Durrleman, Xavier Pennec, Alain Trouvé, and Nicholas Ayache, "Statistical models of sets of curves and surfaces based on currents.", *Medical Image Analysis*, vol. 13, no. 5, pp. 793–808, Oct. 2009. 33

[75] Brian B. Avants, P Thomas Schoenemann, and James C. Gee, "Lagrangian frame diffeomorphic image registration: Morphometric comparison of human and chimpanzee cortex.", *Medical Image Analysis*, vol. 10, no. 3, pp. 397–412, Jun. 2006. 33, 72

[76] Guillaume Auzias, Olivier Colliot, Joan Alexis Glaunes, Matthieu Perrot, Jean-François Mangin, Alain Trouve, and Sylvain Baillet, "Diffeomorphic brain registration under exhaustive sulcal constraints.", *IEEE Transactions on Medical Imaging : TMI*, vol. 30, no. 6, pp. 1214–27, Jun. 2011. 33, 70

[77] Laurent Younes, "Jacobi fields in groups of diffeomorphisms and applications", *Quarterly of applied mathematics*, vol. 65, no. 1, pp. 113–134, 2007. 33, 90, 92

[78] Stephen Marsland and Robert McLachlan, "A Hamiltonian particle method for diffeomorphic image registration.", in *Information Processing in Medical Imaging : IPMI*, Jan. 2007, vol. 20, pp. 396–407. 33

[79] Colin J. Cotter and Darryl D. Holm, "Singular solutions, momentum maps and computational anatomy", *Arxiv preprint nlin/0605020*, 2006. 33

[80] Vincent Arsigny, Olivier Commowick, Xavier Pennec, and Nicholas Ayache, "A log-Euclidean framework for statistics on diffeomorphisms.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'06*, Jan. 2006, number Pt 1 in LNCS, pp. 924–31. 34

[81] Tom Vercauteren, Xavier Pennec, Aymeric Perchant, and Nicholas Ayache, "Diffeomorphic demons: efficient non-parametric image registration.", *NeuroImage*, vol. 45, no. 1 Suppl, pp. S61–S72, Mar. 2009. 34, 61

[82] Mathieu De Craene, Oscar Camara, Bart H. Bijnens, and Alejandro F. Frangi, "Large diffeomorphic FFD registration for motion and strain quantification from 3D-US sequences", in *Functional Imaging and Modeling of the Heart*. 2009, pp. 437–446, Springer. 34, 61

[83] Matias Bossa, Monica Hernandez, and Salvador Olmos, "Contributions to 3D diffeomorphic atlas estimation: application to brain images.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'07*, Jan. 2007, number Pt 1 in LNCS, pp. 667–74. 34

[84] Monica Hernandez, Matias N. Bossa, and Salvador Olmos, "Registration of Anatomical Images Using Paths of Diffeomorphisms Parameterized with Stationary Vector Field Flows", *International Journal of Computer Vision*, vol. 85, no. 3, pp. 291–306, Feb. 2009. 34, 88

[85] Lyubomir Zagorchev and Ardeshir Goshtasby, "A comparative study of transformation functions for nonrigid image registration.", *IEEE Transactions on Image Processing : TIP*, vol. 15, no. 3, pp. 529–38, Mar. 2006. 34, 36

[86] Xuan Yang, Zhong Xue, Xia Liu, and Darong Xiong, "Topology preservation evaluation of compact-support radial basis functions for image registration", *Pattern Recognition Letters*, vol. 32, no. 8, pp. 1162–1177, Jun. 2011. 34, 37

[87] F.L. Bookstein, "Principal warps: thin-plate splines and the decomposition of deformations", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 11, no. 6, pp. 567–585, Jun. 1989. 34

[88] F.L. Bookstein, "Thin-plate splines and the atlas problem for biomedical images", in *Information Processing in Medical Imaging : IPMI*, 1991. 34

[89] Hans J. Johnson and Gary E. Christensen, "Landmark and Intensity-Based, Consistent Thin-Plate Spline Image Registration", in *Information Processing in Medical Imaging : IPMI*, 2001, pp. 329–343. 35

[90] Karl Rohr, H S Stiehl, R Sprengel, T M Buzug, J Weese, and M H Kuhn, "Landmark-based elastic registration using approximating thin-plate splines.", *IEEE Transactions on Medical Imaging : TMI*, vol. 20, no. 6, pp. 526–34, Jun. 2001. 35, 54

[91] Jing Li, Xuan Yang, and Jianping Yu, "Compact support Thin Plate Spline algorithm", *Journal of Electronics (China)*, vol. 24, no. 4, pp. 515–522, Jul. 2007. 35

[92] X Yang and Z Zhang, "Elastic image deformation using adaptive support radial basic function", *International Conference on Wavelet Analysis and Pattern Recognition, 2008.*, pp. 158–162, 2008. 35

[93] Gianluca Donato and Serge Belongie, "Approximate thin plate spline mappings", in *European Conference on Computer Vision : ECCV'02*. 2002, pp. 13–31, Springer. 35

[94] Stephen Marsland, C.J. Twining, and C.J. Taylor, "Groupwise non-rigid registration using polyharmonic clamped-plate splines", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'03*. 2003, LNCS, pp. 771–779, Springer. 35, 110

[95] Vincent Camion and Laurent Younes, "Geodesic interpolating splines", in *Energy Minimization Methods in Computer Vision and Pattern Recognition*. 2001, pp. 513–527, Springer. 36

[96] Laurent Younes, "Combining Geodesic Interpolating Splines and Affine Transformations", *IEEE Transactions on Image Processing : TIP*, vol. 15, no. 5, pp. 1111–1119, 2006. 36

[97] Anna Mills, Tony Shardlow, and Stephen Marsland, "Computing the geodesic interpolating spline", *Biomedical Image Registration*, vol. 4057, no. 2, pp. 169–177, 2006. 36

[98] Detlef Ruprecht, Ralf Nagel, and Heinrigh Müller, "Free form deformation with scattered data interpolation methods", *Computer & Graphics*, vol. 19, no. 1, pp. 63–71, 1995. 36

[99] JA Little, DLG Hill, and DJ Hawkes, "Deformations Incorporating Rigid Structures", *Computer Vision and Image Understanding*, vol. 66, no. 2, pp. 223–232, May 1996. 36

[100] Y. Arad, N., Dyn, N., Reisfeld, D., Yeshurun, "Warping by radial basis functions â?? application to facial expressions.", *CVGIP Graph. Models Image Process.*, vol. 56, no. 2, pp. 161–172, 1994. 36

[101] Dinggang Shen and Christos Davatzikos, "HAMMER: hierarchical attribute matching mechanism for elastic registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 21, no. 11, pp. 1421–39, Nov. 2002. 36, 44, 67, 70, 88

[102] M. Fornefett, K. Rohr, and H.S. Stiehl, "Radial basis functions with compact support for elastic registration of medical images", *Image and Vision Computing*, vol. 19, no. 1-2, pp. 87–96, Jan. 2001. 36

[103] Holger Wendland, "Piecewise polynomial, positive definite, and compactly supported radial basis functions of minimal degree.", *Advances in Computational Mathematics*, vol. 4, pp. 389–396, 1995. 36

[104] Holger Wendland, "Error estimates for interpolation by compactly supported radial functions of minimal degree", 1997. 36

[105] Z. M. Wu, "Compactly supported positive definite radial functions.", *Advances in Computational Mathematics*, vol. 4, no. 1, pp. 283–292, 1995. 36

[106] M D Buhmann, "A New Class of Radial Basis Functions with Compact Support", *Mathematics of Computation*, vol. 70, no. 233, pp. 307–318, 2000. 36

[107] Gustavo K Rohde, Akram Aldroubi, and Benoit M Dawant, "The adaptive bases algorithm for intensity-based nonrigid image registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 22, no. 11, pp. 1470–9, Nov. 2003. 36

[108] Adil Masood Siddiqui, Asif Masood, and Muhammad Saleem, "A locally constrained radial basis function for registration and warping of images", *Pattern Recognition Letters*, vol. 30, no. 4, pp. 377–390, Mar. 2009. 36

[109] Antonio Tristán-Vega and Verónica García-Pérez, "Comments on: A locally constrained radial basis function for registration and warping of images", Mar. 2011. 37

[110] Svenja Lowitzsch, *Approximation and Interpolation Employing Divergence-free Radial Basis Functions with Applications*, PhD thesis, Texas A&M University, 2002. 37, 42

[111] M H Davis, a Khotanzad, D P Flamig, and S E Harms, "A physics-based coordinate transformation for 3-D image matching.", *IEEE Transactions on Medical Imaging : TMI*, vol. 16, no. 3, pp. 317–28, Jun. 1997. 37

[112] Jan. Kohlrausch, Karl Rohr, and H. Siegfried Stiehl, "A New Class of Elastic Body Splines for Nonrigid Registration of Medical Images", *Journal of Mathematical Imaging and Vision*, vol. 23, no. 3, pp. 253–280, Nov. 2005. 37

[113] Stefan Wörz and Karl Rohr, "Physics-based elastic registration using non-radial basis functions and including landmark localization uncertainties", *Computer Vision and Image Understanding*, vol. 111, no. 3, pp. 263–274, 2008. 37, 54

[114] Thomas W. Sederberg and Scott R. Parry, "Free-form deformation of solid geometric models", *ACM Siggraph Computer Graphics*, vol. 20, no. 4, pp. 151–160, 1986. 37

[115] W.M. Hsu, J.F. Hughes, and Henry Kaufman, "Direct manipulation of free-form deformations", in *ACM Siggraph Computer Graphics*. Jul. 1992, number 2, pp. 177–184, ACM. 37

[116] Jérôme Declerck, Jacques Feldmar, Michael L. Goris, and Fabienne Betting, "Automatic registration and alignment on a template of cardiac stress and rest reoriented SPECT images.", *IEEE Transactions on Medical Imaging : TMI*, vol. 16, no. 6, pp. 727–37, Dec. 1997. 37

[117] Daniel Rueckert, L .I. Sonoda, C. Hayes, Derek L. G. Hill, M. O. Leach, and David J. Hawkes, "Nonrigid registration using free-form deformations: application to breast MR images.", *IEEE Transactions on Medical Imaging : TMI*, vol. 18, no. 8, pp. 712–721, Aug. 1999. 37, 60

[118] Jan. Kybic and Michael Unser, "Fast parametric elastic image registration.", *IEEE Transactions on Image Processing : TIP*, vol. 12, no. 11, pp. 1427–42, Jan. 2003. 37, 62

[119] Michaël Sdika, "A Fast Nonrigid Image Registration With Constraints on the Jacobian Using Large Scale Constrained Optimization", *IEEE Transactions on Medical Imaging : TMI*, vol. 27, no. 2, pp. 271–281, 2008. 37, 42

[120] Daniel Rueckert, Paul Aljabar, Rolf A Heckemann, Joseph V. Hajnal, and Alexander Hammers, "Diffeomorphic registration using B-splines.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'06*, Jan. 2006, number Pt 2 in LNCS, pp. 702–9. 38

[121] Yongchoel Choi and Seungyong Lee, "Injectivity Conditions of 2D and 3D Uniform Cubic B-Spline Functions", *Graphical Models*, vol. 62, no. 6, pp. 411–427, Nov. 2000. 38

[122] Wei Feng, Stanley J. Reeves, Thomas S. Denney Jr, Steven Lloyd, Louis Dell'Italia, and Himanshu Gupta, "A NEW CONSISTENT IMAGE REGISTRATION FORMULATION WITH A B-SPLINE DEFORMATION MODEL", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2009, pp. 2–5. 38, 88

[123] J Wang and T Jiang, "Nonrigid registration of brain MRI using NURBS", *Pattern Recognition Letters*, vol. 28, no. 2, pp. 214–223, Jan. 2007. 38

[124] K.K. Bhatia, Joseph V. Hajnal, B.K. Puri, A.D. Edwards, and Daniel Rueckert, "Consistent groupwise non-rigid registration for atlas construction", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2004, pp. 908–911, IEEE. 38, 110, 111, 114, 117

[125] C T Metz, Stefan Klein, M. Schaap, T van Walsum, and Wiro J. Niessen, "Nonrigid registration of dynamic medical imaging data using nD + t B-splines and a groupwise optimization approach.", *Medical Image Analysis*, vol. 15, no. 2, pp. 238–49, Apr. 2011. 38

[126] Dimitrios Perperidis, Raad H Mohiaddin, and Daniel Rueckert, "Spatio-temporal free-form registration of cardiac MR image sequences.", *Medical Image Analysis*, vol. 9, no. 5, pp. 441–56, Oct. 2005. 38

[127] María J Ledesma-Carbayo, Jan. Kybic, Manuel Desco, Andrés Santos, Michael Sühling, Patrick Hunziker, and Michael Unser, "Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation.", *IEEE Transactions on Medical Imaging : TMI*, vol. 24, no. 9, pp. 1113–26, Sep. 2005. 38

[128] Jef Vandemeulebroucke, Simon Rit, Jan. Kybic, Patrick Clarysse, and David Sarrut, "Spatiotemporal motion estimation for respiratory-correlated imaging of the lungs", *Medical Physics*, vol. 38, no. 1, pp. 166, 2011. 38, 61

[129] Y.T. Wu, Takeo Kanade, C.C. Li, and Jeffrey Cohn, "Image registration using wavelet-based motion model", *International Journal of Computer Vision*, vol. 38, no. 2, pp. 129–152, 2000. 39, 62

[130] Smadar Gefen, Oleh Tretiak, and Jonathan Nissanov, "Elastic 3-D alignment of rat brain histological images.", *IEEE Transactions on Medical Imaging : TMI*, vol. 22, no. 11, pp. 1480–9, Nov. 2003. 39, 62

[131] O Musse, F Heitz, and J P Armspach, "Topology preserving deformable image matching using constrained hierarchical parametric models.", *IEEE Transactions on Image Processing : TIP*, vol. 10, no. 7, pp. 1081–93, Jan. 2001. 39, 41

[132] Vincent Noblet, Christian Heinrich, Fabrice Heitz, and Jean-Paul Armspach, "3-D deformable image registration: a topology preservation scheme based on hierarchical deformation models and interval analysis optimization.", *IEEE Transactions on Image Processing : TIP*, vol. 14, no. 5, pp. 553–66, May 2005. 39, 41

[133] V Noblet, C Heinrich, F Heitz, and J-P Armspach, "Retrospective evaluation of a topology preserving non-rigid registration method.", *Medical Image Analysis*, vol. 10, no. 3, pp. 366–84, Jun. 2006. 39

[134] Pascal Cathier, "Iconic feature registration with sparse wavelet coefficients.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'06*, Jan. 2006, number Pt 2 in LNCS, pp. 694–701. 39

[135] Pierre Hellier, Christian Barillot, E. Mémin, and P. Pérez, "Hierarchical estimation of a dense deformation field for 3-D robust registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 20, no. 5, pp. 388–402, May 2001. 40

[136] Hui Zhang, Paul a Yushkevich, Daniel C Alexander, and James C Gee, "Deformable registration of diffusion tensor MR images with explicit orientation optimization.", *Medical Image Analysis*, vol. 10, no. 5, pp. 764–85, Oct. 2006. 40

[137] Alain Pitiot, Eric Bardinet, Paul M Thompson, and Grégoire Malandain, "Piecewise affine registration of biological images for volume reconstruction.", *Medical Image Analysis*, vol. 10, no. 3, pp. 465–83, Jun. 2006. 40

[138] O Commowick, V Arsigny, a Isambert, J Costa, F Dhermain, F Bidault, P-Y Bondiau, N Ayache, and G Malandain, "An efficient locally affine framework for the smooth registration of anatomical structures.", *Medical Image Analysis*, vol. 12, no. 4, pp. 427–41, Aug. 2008. 40

[139] Timothy F Cootes, Carole J Twining, Vladimir S Petrovi, Kolawole O Babalola, and Christopher J Taylor, "Compution Accurate Correspondences across Groups of Images", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 32, no. 11, pp. 1994–2005, 2010. 40, 54, 112

[140] Christian Buerger, Tobias Schaeffter, and Andrew P King, "Hierarchical adaptive local affine registration for fast and robust respiratory motion estimation.", *Medical Image Analysis*, vol. 15, no. 4, pp. 551–564, Mar. 2011. 40

[141] Ramkrishnan Narayanan, Jeffrey A. Fessler, H. Park, and C. R. Meyerl, "Diffeomorphic nonlinear transformations: a local parametric approach for image registration.", in *Information processing in medical imaging : IPMI*, Jan. 2005, vol. 19, pp. 174–85. 40

[142] Vincent Arsigny, Xavier Pennec, and Nicholas Ayache, "Polyrigid and polyaffine transformations: a Novel geometrical tool to deal with non-rigid deformations - application to the registration of histological slices.", *Medical Image Analysis*, vol. 9, no. 6, pp. 507–23, Dec. 2005. 40

[143] Vincent Arsigny, Olivier Commowick, Nicholas Ayache, and Xavier Pennec, "A Fast and Log-Euclidean Polyaffine Framework for Locally Linear Registration", *Journal of Mathematical Imaging and Vision*, vol. 33, no. 2, pp. 222–238, Jan. 2009. 40

[144] M. Droske and M. Rumpf, "A variational approach to nonrigid morphological image registration", *SIAM Journal on Applied Mathematics*, vol. 64, no. 2, pp. 668–687, 2003. 41, 50, 60

[145] Eldad Haber and Jan. Modersitzki, "Image Registration with Guaranteed Displacement Regularity", *International Journal of Computer Vision*, vol. 71, no. 3, pp. 361–372, Jul. 2006. 41, 61, 63

[146] Se Young Chun and Jeffrey A. Fessler, "Regularized Methods for Topology-Preserving Smooth Nonrigid Image Registration Using B-Spline Basis", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2008, pp. 1099–1102, IEEE. 42

[147] Torsten Rohlfing, Calvin R Maurer, David A. Bluemke, and Michael A. Jacobs, "Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint.", *IEEE Transactions on Medical Imaging : TMI*, vol. 22, no. 6, pp. 730–41, Jun. 2003. 42

[148] C. Tanner, J. A. Schnabel, D. Chung, M. J. Clarkson, D. Rueckert, and D. L. G. Hill, "Volume and Shape Preservation of Enhancing Lesions when Applying Non-rigid Registration to a Time Series of Contrast Enhancing MR Breast Images", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'00*, 2000, pp. 327–337. 42

[149] Eldad Haber and Jan. Modersitzki, "Numerical methods for volume preserving image registration", *Inverse Problems*, vol. 20, no. 5, pp. 1621–1638, Oct. 2004. 42, 63

[150] Arnaud Bistoquet, John Oshinski, and Oskar Skrinjar, "Myocardial deformation recovery from cine MRI using a nearly incompressible biventricular model.", *Medical Image Analysis*, vol. 12, no. 1, pp. 69–85, Feb. 2008. 42

[151] Julien Dauguet, A.S. Herard, J. Declerck, and Thierry Delzescaux, "Locally Constrained Cubic B-Spline Deformations to Control Volume Variations", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2009, pp. 983–986, IEEE. 43, 63

[152] Dirk Loeckx, Frederik Maes, Dirk Vandermeulen, and Paul Suetens, "Nonrigid image registration using free-form deformations with a local rigidity constraint", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'04*, 2004, pp. 639–646. 43

[153] Marius Staring, Stefan Klein, and Josien P. W. Pluim, "A rigidity penalty term for nonrigid registration", *Medical Physics*, vol. 34, no. 11, pp. 4098–4108, 2007. 43

[154] Jan. Modersitzki, "FLIRT with Rigidity–Image Registration with a Local Non-rigidity Penalty", *International Journal of Computer Vision*, vol. 76, no. 2, pp. 153–163, Sep. 2007. 43, 61

[155] Jeongtae Kim and Jeffrey A. Fessler, "Intensity-based image registration using robust correlation coefficients.", *IEEE Transactions on Medical Imaging : TMI*, vol. 23, no. 11, pp. 1430–44, Nov. 2004. 44

[156] Zhong Xue, Dinggang Shen, and Christos Davatzikos, "Determining correspondence in 3-D MR brain images using attribute vectors as morphological signatures of voxels.", *IEEE Transactions on Medical Imaging : TMI*, vol. 23, no. 10, pp. 1276–91, Oct. 2004. 45, 56, 67

[157] Dinggang Shen, "Image registration by local histogram matching", *Pattern Recognition*, vol. 40, no. 4, pp. 1161–1172, Apr. 2007. 45, 67

[158] Guorong Wu, Feihu Qi, and Dinggang Shen, "Learning-based deformable registration of MR brain images.", *IEEE Transactions on Medical Imaging : TMI*, vol. 25, no. 9, pp. 1145–57, Sep. 2006. 45, 67

[159] Yangming Ou, Aristeidis Sotiras, Nikos Paragios, and Christos Davatzikos, "DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting.", *Medical Image Analysis*, vol. 15, pp. 622–639, Jul. 2010. 45, 51, 56, 66

[160] J Liu, B C Vemuri, and J L Marroquin, "Local frequency representations for robust multimodal image registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 21, no. 5, pp. 462–469, May 2002. 45, 51

[161] Shu Liao and Albert C. S. Chung, "Feature based nonrigid brain MR image registration with symmetric alpha stable filters.", *IEEE Transactions on Medical Imaging : TMI*, vol. 29, no. 1, pp. 106–19, Jan. 2010. 45

[162] S Liao and Albert C. S. Chung, "Non-rigid Brain MR Image Registration using Uniform Spherical Region Descriptor.", *IEEE Transactions on Image Processing : TIP*, vol. XX, no. XX, pp. 1–14, Jun. 2011. 45, 65

[163] William M. Wells III, Paul Viola, Hideki Atsumi, Shin Nakajima, and Ron Kikinis, "Multi-modal volume registration by maximization of mutual information", *Medical Image Analysis*, vol. 1, no. 1, pp. 35–51, 1996. 46, 62

[164] Paul Viola and William M. Wells III, "Alignment by maximization of mutual information", *International Journal of Computer Vision*, vol. 24, no. 2, pp. 137–54, 1997. 46, 62

[165] André Collignon, Frederik Maes, D. Delaere, D. Vandermeulen, Paul Suetens, and Guy Marchal, "Automated Multi-Modality Image Registration Based on Information Theory", in *Information Processing in Medical Imaging : IPMI*. 1995, number 6, p. 263, Dordrecht, The Netherlands: Kluwer. 46

[166] F Maes, a Collignon, D Vandermeulen, G Marchal, and P Suetens, "Multimodality image registration by maximization of mutual information.", *IEEE Transactions on Medical Imaging : TMI*, vol. 16, no. 2, pp. 187–98, Apr. 1997. 46, 61

[167] Josien P. W. Pluim, J.B. Antoine Maintz, and Max A. Viergever, "Mutual-information-based registration of medical images: a survey.", *IEEE Transactions on Medical Imaging : TMI*, vol. 22, no. 8, pp. 986–1004, Aug. 2003. 46

[168] C Studholme, D. L. G. Hill, and David J. Hawkes, "An overlap invariant entropy measure of 3D medical image alignment", *Pattern Recognition*, vol. 32, no. 1, pp. 71–86, Jan. 1999. 46

[169] Nathan D. Cahill, Julia a. Schnabel, J. Alison Noble, and David J. Hawkes, "Revisiting overlap invariance in medical image alignment", in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. Jun. 2008, pp. 1–8, Ieee. 46

[170] Alexis Roche, G. Malandain, Xavier Pennec, and Nicholas Ayache, "The correlation ratio as a new similarity measure for multimodal image registration", in *Medical Image Computing and Computer-Assisted Intervention: MICCAI'98*. 1998, vol. 1496, pp. 1115 –1124, Springer. 46

[171] Josien P W Pluim, J B Antoine Maintz, and Max A. Viergever, "F-information measures in medical image registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 23, no. 12, pp. 1508–16, Dec. 2004. 46, 61

[172] Albert C.S. Chung, William M. Wells III, Alexander Norbash, and W. Eric L. Grimson, "Multi-modal image registration by minimizing Kullback-Leibler distance", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'02*. 2002, pp. 525–32, Springer-Verlag Berlin Heidelberg. 46, 51, 61

[173] Christoph Guetter, Chenyang Xu, Frank Sauer, and Joachim Hornegger, "Learning based non-rigid multi-modal image registration using Kullback-Leibler divergence.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'05*, Jan. 2005, number Pt 2 in LNCS, pp. 255–62. 46, 51

[174] Rui Liao, Christoph Guetter, Chenyang Xu, Yiyong Sun, Ali Khamene, and Frank Sauer, "Learning-based 2D/3D rigid registration using Jensen-Shannon divergence for image-guided surgery", in *Medical Imaging and Augmented Reality*. 2006, pp. 228–235, Springer. 47, 51

[175] A. Ben Hamza and Hamid Krim, "Image Registration and Segmentation by Maximizing the Jensen-Rényi Divergence", in *Energy Minimization Methods in Computer Vision and Pattern Recognition*. 2003, pp. 147–163, Springer. 47

[176] A. Ben Hamza and H. Krim, "A generalized divergence measure for robust image registration", *IEEE Transactions on Signal Processing*, vol. 51, no. 5, pp. 1211–1220, May 2003. 47

[177] Huzefa Neemuchwala, Alfred O. Hero III, and Paul Carson, "Image registration using entropic graph-matching criteria", in *Conference Record of the Thirty-Sixth Asilomar Conference on Signals, Systems and Computers, 2002.* 2002, pp. 134–138, IEEE. 47

[178] Alfred O. Hero III, Bing Ma, Olivier J. J. Michel, and John Gorman, "Applications of entropic spanning graphs", *IEEE Signal Processing Magazine*, vol. 19, no. 5, pp. 85–95, Sep. 2002. 47

[179] Mert R. Sabuncu and Peter Ramadge, "Using spanning graphs for efficient image registration.", *IEEE Transactions on Image Processing : TIP*, vol. 17, no. 5, pp. 788–97, May 2008. 47, 51

[180] Stefan Martin and Tariq S. Durrani, "A New Divergence Measure for Medical Image Registration", *IEEE Transactions on Image Processing : TIP*, vol. 16, no. 4, pp. 957–966, Apr. 2007. 47

[181] Daniel Rueckert, M.J. Clarkson, D. L. G. Hill, and David J. Hawkes, "Non-rigid registration using higher-order mutual information", *Proceedings of SPIE*, vol. 3979, pp. 438–447, 2000. 47, 48

[182] Gerardo Hermosillo, Christophe Chefd'Hotel, and Olivier Faugeras, "Variational methods for multimodal image matching", *International Journal of Computer Vision*, vol. 50, no. 3, pp. 329–343, 2002. 47

[183] Bilge Karaçali, "Information Theoretic Deformable Registration Using Local Image Information", *International Journal of Computer Vision*, vol. 72, no. 3, pp. 219–237, Jul. 2007. 47

[184] Colin Studholme, Corina Drapaca, Bistra Iordanova, and Valerie Cardenas, "Deformation-based mapping of volume change from serial brain MRI in the presence of local tissue contrast change.", *IEEE Transactions on Medical Imaging : TMI*, vol. 25, no. 5, pp. 626–39, May 2006. 47, 90

[185] D. Loeckx, P. Slagmolen, F. Maes, D. Vandermeulen, and P. Suetens, "Nonrigid Image Registration Using Conditional Mutual Information", *IEEE Transactions on Medical Imaging : TMI*, vol. 29, no. 1, pp. 19–29, 2010. 47, 61

[186] Xiahai Zhuang, Simon Arridge, D. Hawkes, and Sebastien Ourselin, "A Nonrigid Registration Framework Using Spatially Encoded Mutual Information and Free-Form Deformations.", *IEEE Transactions on Medical Imaging : TMI*, , no. c, pp. 1–10, May 2011. 47

[187] J P Pluim, J B Maintz, and M A Viergever, "Image registration by maximization of combined mutual information and gradient information.", *IEEE Transactions on Medical Imaging : TMI*, vol. 19, no. 8, pp. 809–14, Aug. 2000. 48, 61

[188] Daniel Russakoff, Carlo Tomasi, Torsten Rohlfing, and Calvin R. Maurer Jr, "Image similarity using mutual information of regions", in *European Conference on Computer Vision : ECCV'04*. 2004, pp. 596–607, Springer. 48

[189] A Bardera, M Feixas, I Boada, and M Sbert, "High-dimensional normalized mutual information for image registration using random lines", in *Biomedical Image Registration*. 2006, pp. 264–271, Springer. 48

[190] Mark Holden, L.D. Griffin, Nadeem Saeed, and D.L.G. Hill, "Multi-channel mutual information using scale space", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'04*. 2004, pp. 797–804, Springer. 48

[191] Rui Gan, Albert C. S. Chung, and Shu Liao, "Maximum distance-gradient for robust image registration.", *Medical Image Analysis*, vol. 12, no. 4, pp. 452–68, Aug. 2008. 48, 61

[192] M.R. Sabuncu and P.J. Ramadge, "Spatial information in entropy-based image registration", in *Biomedical Image Registration*. 2003, pp. 132–141, Springer. 49

[193] Huzefa Neemuchwala, Alfred O. Hero III, and Paul Carson, "Image matching using alpha-entropy measures and entropic graphs", *Signal Processing*, vol. 85, no. 2, pp. 277–296, Feb. 2005. 49

[194] Marius Staring, Uulke a van der Heide, Stefan Klein, Max a Viergever, and Josien P W Pluim, "Registration of cervical MRI using multifeature mutual information.", *IEEE Transactions on Medical Imaging : TMI*, vol. 28, no. 9, pp. 1412–21, Sep. 2009. 49, 62

[195] C. Studholme, D.L.G. Hill, and D.J. Hawkes, "Incorporating connected region labelling into automated image registration using mutual information", in *Proceedings of the Workshop on Mathematical Methods in Biomedical Image Analysis*. 1996, pp. 23–31, Ieee. 49

[196] Z.F. Knops, JBA Maintz, MA Viergever, and JPW Pluim, "Registration using segment intensity remapping and mutual information", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'04*. 2004, pp. 805–812, Springer. 49

[197] Emiliano D'Agostino, Frederik Maes, Dirk Vandermeulen, and Paul Suetens, "An information theoretic approach for non-rigid image registration using voxel class probabilities.", *Medical Image Analysis*, vol. 10, no. 3, pp. 413–31, Jun. 2006. 49

[198] a Roche, X Pennec, G Malandain, and N Ayache, "Rigid registration of 3-D ultrasound with MR images: a new approach combining intensity and gradient information.", *IEEE Transactions on Medical Imaging : TMI*, vol. 20, no. 10, pp. 1038–49, Oct. 2001. 50

[199] Wolfgang Wein, Shelby Brunke, Ali Khamene, Matthew R Callstrom, and Nassir Navab, "Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention.", *Medical Image Analysis*, vol. 12, no. 5, pp. 577–85, Oct. 2008. 50

[200] Fabrice Michel and Nikos Paragios, "Image Transport Regression Using Mixture of Experts and Discrete Markov Random Fields", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2010, pp. 1229–1232, IEEE. 50

[201] J.B. Antoine Maintz, P.A. van den Elsen, and Max A. Viergever, "3D multimodality medical image registration using morphological tools", *Image and Vision Computing*, vol. 19, no. 1-2, pp. 53–62, Jan. 2001. 50

[202] Eldad Haber and Jan. Modersitzki, "Intensity gradient based registration and fusion of multi-modal images.", *Methods of information in medicine*, vol. 46, no. 3, pp. 292–9, Jan. 2007. 50

[203] Torsten Butz and Jean-Philippe Thiran, "Affine registration with feature space mutual information", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'01*, Wiro J. Niessen and Max A. Viergever, Eds. 2001, pp. 549–556, Springer-Verlag Berlin Heidelberg. 50

[204] G.P. Penney, J.M. Blackall, M.S. Hamady, T. Sabharwal, A. Adam, and D.J. Hawkes, "Registration of freehand 3D ultrasound and magnetic resonance liver images", *Medical Image Analysis*, vol. 8, no. 1, pp. 81–91, Mar. 2004. 51

[205] Bing Jian, Baba C. Vemuri, and José L. Marroquin, "Robust nonrigid multimodal image registration using local frequency maps.", in *Information Processing in Medical Imaging : IPMI*, Jan. 2005, vol. 19, pp. 504–15. 51

[206] a Andronache, M von Siebenthal, G Székely, and Ph Cattin, "Non-rigid registration of multi-modal images using both mutual information and cross-correlation.", *Medical Image Analysis*, vol. 12, no. 1, pp. 3–15, Feb. 2008. 51

[207] Christian Wachinger and Nassir Navab, "Entropy and Laplacian images : Structural representations for multi-modal registration", *Medical Image Analysis*, 2011. 51

[208] Daewon Lee, Matthias Hofmann, Florian Steinke, Yasemin Altun, Nathan D. Cahill, and Bernhard Scholkopf, "Learning similarity measure for multi-modal 3D image registration", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2009, pp. 186–193, Ieee. 51

[209] Michael M. Bronstein, Alexander M. Bronstein, Fabrice Michel, and Nikos Paragios, "Data Fusion through Cross-modality Metric Learning using Similarity-Sensitive Hashing", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*, 2010, pp. 3594–3601. 51

[210] Fabrice Michel, Michael M. Bronstein, Alexander M. Bronstein, and Nikos Paragios, "Boosted Metric Learning for 3D Multi-modal Deformable Registration", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2011, pp. 1209–1214, IEEE. 51

[211] M. Leventon and W. Grimson, "Multi-modal volume registration using joint intensity distributions", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'98*. 1998, pp. 1057–1066, Springer. 51

[212] Chris Harris and Mike Stephens, "A Combined Corner and Edge Detector", in *Alvey vision conference*. 1988, vol. 15, p. 50, Manchester, UK. 52

[213] Carlo Tomasi, "Good features to track", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. 1994, pp. 593–600, IEEE Comput. Soc. Press. 52

[214] Bill Triggs, "Detecting keypoints with stable position, orientation, and scale under illumination changes", in *European Conference on Computer Vision : ECCV'04*. 2004, pp. 100–113, Springer. 52, 53

[215] Krystian Mikolajczyk and C. Schmid, "Scale & affine invariant interest point detectors", *International journal of computer vision*, vol. 60, no. 1, pp. 63–86, Oct. 2004. 52

[216] K. Mikolajczyk, T. Tuytelaars, C. Schmid, a. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool, "A Comparison of Affine Region Detectors", *International Journal of Computer Vision*, vol. 65, no. 1-2, pp. 43–72, Oct. 2005. 52

[217] Cordelia Schmid, Roger Mohr, and Christian Bauckhage, "Evaluation of Interest Point Detectors", *International Journal of Computer Vision*, vol. 37, no. 2, pp. 151–172, 2000. 53

[218] F Mokhtarian and F Mohanna, "Performance evaluation of corner detectors using consistency and accuracy measures", *Computer Vision and Image Understanding*, vol. 102, no. 1, pp. 81–94, Apr. 2006. 53

[219] Tony Lindeberg, "Feature detection with automatic scale selection", *International Journal of Computer Vision*, vol. 30, no. 2, pp. 79–116, 1998. 53

[220] D.G. Lowe, "Distinctive image features from scale-invariant keypoints", *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, Nov. 2004. 53, 56, 94, 100

[221] Warren Cheung and Ghassan Hamarneh, "n-SIFT: n-dimensional scale invariant feature transform.", *IEEE Transactions on Image Processing : TIP*, vol. 18, no. 9, pp. 2012–21, Sep. 2009. 53

[222] Jiri Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide-baseline stereo from maximally stable extremal regions", *Image and Vision Computing*, vol. 22, no. 10, pp. 761–767, Sep. 2004. 53

[223] Pascal Cachier, J.F. Mangin, Xavier Pennec, D. Rivière, D. Papadopoulos-Orfanos, J. Régis, and Nicholas Ayache, "Multisubject non-rigid registration of brain MRI using intensity and geometric features", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'01*. 2001, vol. 16, pp. 734–742, Springer-Verlag. 53, 73

[224] Pierre Hellier and Christian Barillot, "Coupling dense and landmark-based approaches for nonrigid registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 22, no. 2, pp. 217–227, Feb. 2003. 53, 71

[225] Haili Chui, Lawrence Win, Robert Schultz, James S. Duncan, and Anand Rangarajan, "A unified non-rigid feature registration method for brain mapping", *Medical Image Analysis*, vol. 7, no. 2, pp. 113–130, Jun. 2003. 53, 58

[226] Tianming Liu, Dinggang Shen, and Christos Davatzikos, "Deformable registration of cortical structures via hybrid volumetric and surface warping.", *NeuroImage*, vol. 22, no. 4, pp. 1790–801, Aug. 2004. 53, 67, 70

[227] Anqi Qiu and Michael I Miller, "Cortical hemisphere registration via large deformation diffeomorphic metric curve mapping.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'07*, Jan. 2007, number Pt 1 in LNCS, pp. 186–93. 53

[228] Ali Can, Charles V. Stewart, Badrinath Roysam, and Howard L. Tanenbaum, "A feature-based, robust, hierarchical algorithm for registering pairs of images of the curved human retina", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 24, no. 3, pp. 347–364, Mar. 2002. 53

[229] Charles V. Stewart, Chia-Ling Tsai, and Badrinath Roysam, "The dual-bootstrap iterative closest point algorithm with application to retinal image registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 22, no. 11, pp. 1379–1394, Nov. 2003. 53, 58

[230] I Reinertsen, M Descoteaux, K Siddiqi, and D L Collins, "Validation of vessel-based registration for correction of brain shift.", *Medical Image Analysis*, vol. 11, no. 4, pp. 374–88, Aug. 2007. 53

[231] Changyan Xiao, Marius Staring, Denis Shamonin, Johan H C Reiber, Jan. Stolk, and Berend C Stoel, "A strain energy filter for 3D vessel enhancement with application to pulmonary CT images.", *Medical Image Analysis*, vol. 15, no. 1, pp. 112–24, Feb. 2011. 53

[232] Klaus Drechsler and Cristina Oyarzun Laura, "Comparison of vesselness functions for multiscale analysis of the liver vasculature", in *Proceedings of the 10th IEEE International Conference on Information Technology and Applications in Biomedicine*. Nov. 2010, pp. 1–5, IEEE. 53

[233] Y Sato, S Nakajima, N Shiraga, H Atsumi, S Yoshida, T Koller, G Gerig, and R Kikinis, "Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images.", *Medical Image Analysis*, vol. 2, no. 2, pp. 143–68, Jun. 1998. 53

[234] Alejandro F. Frangi, Wiro J. Niessen, Koen L. Vincken, and Max A. Viergever, "Multiscale Vessel Enhancement Filtering", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'98*, William M. Wells III, Alan Colchester, and Scott Delp, Eds. Jun. 1998, number 3 in LNCS, pp. 130–137, Springer Berlin / Heidelberg. 53

[235] C. Cañero and P. Radeva, "Vesselness enhancement diffusion", *Pattern Recognition Letters*, vol. 24, no. 16, pp. 3141–3151, Dec. 2003. 53

[236] Rashindra Manniesing, Max a Viergever, and Wiro J Niessen, "Vessel enhancing diffusion: a scale space representation of vessel structures.", *Medical Image Analysis*, vol. 10, no. 6, pp. 815–25, Dec. 2006. 53

[237] Yangming Ou, A Besbes, M Bilello, M Mansour, Christos Davatzikos, and Nikos Paragios, "Detecting mutually-salient landmark pairs with MRF regularization", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2010, pp. 400–403. 53, 56, 57, 80

[238] J M Fitzpatrick and J B West, "The distribution of target registration error in rigid-body point-based registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 20, no. 9, pp. 917–27, Sep. 2001. 53

[239] Andrew D Wiles, Alexander Likholyot, Donald D Frantz, and Terry M Peters, "A statistical model for point-based target registration error with anisotropic fiducial localizer error.", *IEEE Transactions on Medical Imaging : TMI*, vol. 27, no. 3, pp. 378–90, Mar. 2008. 53

[240] Mehdi Hedjazi Moghari and Purang Abolmaesumi, "Distribution of fiducial registration error in rigid-body point-based registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 28, no. 11, pp. 1791–801, Nov. 2009. 53

[241] Colin Goodall, "Procrustes methods in the statistical analysis of shape", *Journal of the Royal Statistical Society B*, vol. 53, no. 2, pp. 285–339, 1991. 54, 59

[242] T.F. Cootes and C.J. Taylor, "A mixture model for representing shape variation", *Image and Vision Computing*, vol. 17, no. 8, pp. 567–573, Jun. 1999. 54

[243] N. Duta, A.K. Jain, and M.-P. Dubuisson-Jolly, "Automatic construction of 2D shape models", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 23, no. 5, pp. 433–446, May 2001. 54

[244] K. Voss and H. Suesse, "Affine point pattern matching", in *Pattern Recognition*, Bernd Radig and Stefan Florczyk, Eds., vol. 2191 of *Lecture Notes in Computer Science*, pp. 155–162. Springer Berlin / Heidelberg, 2001. 54

[245] Jeffrey Ho and Ming-Hsuan Yang, "On affine registration of planar point sets using complex numbers", *Computer Vision and Image Understanding*, vol. 115, no. 1, pp. 50–58, Jan. 2011. 54

[246] Hongyu Guo, Anand Rangarajan, and Sarang C. Joshi, "Diffeomorphic point matching", in *Handbook of Mathematical Models in Computer Vision*, Nikos Paragios, Yunmei Chen, and Olivier Faugeras, Eds., chapter 13, pp. 205–219. Springer, 2006. 54

[247] J Heikkila, "Pattern matching with affine moment descriptors", *Pattern Recognition*, vol. 37, no. 9, pp. 1825–1834, Sep. 2004. 54

[248] C Gope and N Kehtarnavaz, "Affine invariant comparison of point-sets using convex hulls and Hausdorff distances", *Pattern Recognition*, vol. 40, no. 1, pp. 309–320, Jan. 2007. 55

[249] A. Fitzgibbon, "Robust registration of 2D and 3D point sets", *Image and Vision Computing*, vol. 21, no. 13-14, pp. 1145–1153, Dec. 2003. 55

[250] Yanghai Tsin and Takeo Kanade, "A correlation-based approach to robust point set registration", in *European Conference on Computer Vision : ECCV'04*. 2004, pp. 558–569, Springer. 55

[251] Maneesh Singh, Arora Himanshu, and Narendra Ahuja, "Robust registration and tracking using kernel density correlation", in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. 2004, Published by the IEEE Computer Society. 55

[252] Bing Jian and Baba Vemuri, "Robust Point Set Registration Using Gaussian Mixture Models", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 33, no. 8, pp. 1633–45, 2011. 55

[253] Andriy Myronenko and Xubo Song, "Point set registration: coherent point drift.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 32, no. 12, pp. 2262–75, Dec. 2010. 55

[254] Arunabha S Roy, Ajay Gopinath, and Anand Rangarajan, "Deformable density matching for 3D non-rigid registration of shapes.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'07*, Jan. 2007, number Pt 1 in LNCS, pp. 942–9. 55

[255] Fei Wang, Baba C Vemuri, Anand Rangarajan, and Stephan J Eisenschenk, "Simultaneous nonrigid registration of multiple point sets and atlas construction.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 30, no. 11, pp. 2011–22, Nov. 2008. 55, 61, 113, 118

[256] Fei Wang, Baba Vemuri, and Tanveer Syeda-Mahmood, "Generalized L2-divergence and its application to shape alignment.", in *Information processing in medical imaging : IPMI*, Jan. 2009, vol. 21 of *LNCS*, pp. 227–38. 55, 113, 118

[257] Nicholas J Tustison, Suyash P Awate, Gang Song, Tessa S Cook, and James C Gee, "Point set registration using Havrda-Charvat-Tsallis entropy measures.", *IEEE Transactions on Medical Imaging : TMI*, vol. 30, no. 2, pp. 451–60, Feb. 2011. 55

[258] Xiaolei Huang, Nikos Paragios, and Dimitris N Metaxas, "Shape registration in implicit spaces using information theory and free form deformations.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 28, no. 8, pp. 1303–18, Aug. 2006. 56, 60, 118

[259] Mickael Savinaud, Nikos Paragios, and Serge Maitrejean, "Motion-based enhancement of optical imaging", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. Jun. 2009, pp. 738–741, Ieee. 56, 118

[260] Krystian Mikolajczyk and Cordelia Schmid, "Performance evaluation of local descriptors.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 27, no. 10, pp. 1615–30, Oct. 2005. 56, 57

[261] Arturo Gil, Oscar Martinez Mozos, Monica Ballesta, and Oscar Reinoso, "A comparative evaluation of interest point detectors and local descriptors for visual SLAM", *Machine Vision and Applications*, vol. 21, no. 6, pp. 905–920, Apr. 2009. 56

[262] Yan Ke and Rahul Sukthankar, "PCA-SIFT: a more distinctive representation for local image descriptors", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. 2004, pp. 506–513, IEEE. 56

[263] H Bay, a Ess, T Tuytelaars, and L Vangool, "Speeded-Up Robust Features (SURF)", *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346–359, Jun. 2008. 56

[264] Luo Juan and O. Gwun, "A Comparison of SIFT , PCA-SIFT and SURF", *International Journal of Image Processing (IJIP)*, vol. 3, no. 5, pp. 143–152, 2010. 56

[265] Jean-Michel Morel and Guoshen Yu, "ASIFT: A New Framework for Fully Affine Invariant Image Comparison", *SIAM Journal on Imaging Sciences*, vol. 2, no. 2, pp. 438, 2009. 56

[266] Xiao Han, "Feature-constrained Nonlinear Registration of Lung CT Images", in *MICCAI EMPIRE10 Challenge*, 2010. 56

[267] W.T. Freeman, E.H. Adelson, Massachusetts Institute Technology. Media Laboratory. Vision, and Modeling Group, "The design and use of steerable filters", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 13, no. 9, pp. 891–906, 1991. 56

[268] Mathews Jacob and Michael Unser, "Design of steerable filters for feature detection using canny-like criteria.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 26, no. 8, pp. 1007–19, Aug. 2004. 56

[269] Nicolas Honnorat, Régis Vaillant, and Nikos Paragios, "Guide-wire extraction through perceptual organization of local segments in fluoroscopic images.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'10*, Jan. 2010, number Pt 3 in LNCS, pp. 440–8. 56

[270] M. Leordeanu and M. Hebert, "A spectral technique for correspondence problems using pairwise constraints", in *Tenth IEEE International Conference on Computer Vision (ICCV'05)*. 2005, pp. 1482–1489 Vol. 2, Ieee. 57

[271] Alexander C. Berg, Tamara L. Berg, and Jitendra Malik, "Shape Matching and Object Recognition Using Low Distortion Correspondences", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. 2005, pp. 26–33, Ieee. 57

[272] Lorenzo Torresani, Vladimir Kolmogorov, and Carsten Rother, "Feature correspondence via graph matching: Models and global optimization", in *European Conference on Computer Vision : ECCV'08*. 2008, pp. 596–609, Springer. 57

[273] Olivier Duchenne, Francis Bach, In-So Kweon, and Jean Ponce, "A Tensor-Based Algorithm for High-Order Graph Matching.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. V, pp. 1–13, May 2010. 57

[274] Ron Zass and Amnon Shashua, "Probabilistic graph and hypergraph matching", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2008, pp. 1–8, Ieee. 57

[275] Chaohui Wang, Olivier Teboul, Fabrice Michel, Salma Essafi, and Nikos Paragios, "3D knowledge-based segmentation using pose-invariant higher-order graphs.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'10*, Jan. 2010, number Pt 3 in LNCS, pp. 189–96. 57, 134

[276] Nikos Komodakis, Nikos Paragios, and Georgios Tziritas, "MRF energy minimization and beyond via dual decomposition.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 33, no. 3, pp. 531–52, Mar. 2011. 57

[277] Paul J. Besl and Neil D. McKay, "A method for registration of 3-D shapes", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 14, no. 2, pp. 239–256, 1992. 57

[278] Yonghuai Liu, "Improving ICP with easy implementation for free-form surface matching", *Pattern Recognition*, vol. 37, no. 2, pp. 211–226, Feb. 2004. 58

[279] G.C. Sharp, S.W. Lee, and D.K. Wehe, "ICP registration using invariant features", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 24, no. 1, pp. 90–102, 2002. 58

[280] G. Penney, P. Edwards, A. King, J. Blackall, P. Batchelor, and D. Hawkes, "A stochastic iterative closest point algorithm (stochastICP)", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'01*. 2001, pp. 762–769, Springer. 58

[281] R.S.J. Estépar, Anders Brun, and C.F. Westin, "Robust generalized total least squares iterative closest point registration", pp. 234–241, 2004. 58

[282] Haili Chui and Anand Rangarajan, "A new point matching algorithm for non-rigid registration", *Computer Vision and Image Understanding*, vol. 89, no. 2-3, pp. 114–141, Mar. 2003. 58

[283] S. Granger and X. Pennec, "Multi-scale EM-ICP: A fast and robust approach for surface registration", in *European Conference on Computer Vision : ECCV'02*. 2002, pp. 69–73, Springer. 58

[284] M Carcassoni and E Hancock, "Spectral correspondence for point pattern matching", *Pattern Recognition*, vol. 36, no. 1, pp. 193–204, Jan. 2003. 58

[285] Andrew D.J. Cross and Edwin R. Hancock, "Graph matching with a dual-step EM algorithm", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 20, no. 11, pp. 1236–1253, 1998. 58

[286] Xiaolei Huang, Yiyong Sun, Dimitris Metaxas, Frank Sauer, and Chenyang Xu, "Hybrid image registration based on configural matching of scale-invariant salient region features", in *IEEE Computer Society Conference on Computer Vision and*

Pattern Recognition Workshops. 2004, Published by the IEEE Computer Society. 58

[287] B Luo and E. R. Hancock, "Iterative Procrustes alignment with the EM algorithm", *Image and Vision Computing*, vol. 20, no. 5-6, pp. 377–396, Apr. 2002. 59

[288] Anand Rangarajan, Haili Chui, and F.L. Bookstein, "The softassign procrustes matching algorithm", in *Information Processing in Medical Imaging : IPMI*. 1997, pp. 29–42, Springer. 59

[289] Stefan Klein, Marius Staring, and Josien P. W. Pluim, "Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines.", *IEEE Transactions on Image Processing : TIP*, vol. 16, no. 12, pp. 2879–90, Dec. 2007. 60, 61, 62, 67

[290] J.J. Moré and D.J. Thuente, "Line search algorithms with guaranteed sufficient decrease", *ACM Transactions on Mathematical Software (TOMS)*, vol. 20, no. 3, pp. 286–307, 1994. 60

[291] L. Grippo, F. Lampariello, and S. Lucidi, "A nonmonotone line search technique for Newton's method", *SIAM Journal on Numerical Analysis*, vol. 23, no. 4, pp. 707–716, 1986. 60

[292] William H Press, Saul A Teukolsky, William T Vetterling, and Brian P Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 1992. 60, 61

[293] Hans J Johnson and Gary E Christensen, "Consistent landmark and intensity-based image registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 21, no. 5, pp. 450–461, May 2002. 60, 70, 88

[294] Anand a Joshi, David W Shattuck, Paul M Thompson, and Richard M Leahy, "Surface-constrained volumetric brain registration using harmonic mappings.", *IEEE Transactions on Medical Imaging : TMI*, vol. 26, no. 12, pp. 1657–69, Dec. 2007. 61, 72

[295] Gheorghe Postelnicu, Lilla Zollei, and Bruce Fischl, "Combined volumetric and surface registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 28, no. 4, pp. 508–22, Apr. 2009. 61, 70

[296] Nicholas J. Tustison, Brian B. Avants, and James C. Gee, "Directly manipulated free-form deformation image registration.", *IEEE Transactions on Image Processing : TIP*, vol. 18, no. 3, pp. 624–35, Mar. 2009. 61

[297] P Thévenaz and M Unser, "Optimization of mutual information for multiresolution image registration.", *IEEE Transactions on Image Processing : TIP*, vol. 9, no. 12, pp. 2083–99, Jan. 2000. 62

[298] J. Kiefer and J. Wolfowitz, "Stochastic estimation of the maximum of a regression function", *The Annals of Mathematical Statistics*, vol. 23, no. 3, pp. 462–466, 1952. 62

[299] James C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation", *IEEE Transactions on Automatic Control*, vol. 37, no. 3, pp. 332–341, Mar. 1992. 62

[300] H. Robbins and S. Monro, "A stochastic approximation method", *The Annals of Mathematical Statistics*, vol. 22, no. 3, pp. 400–407, 1951. 62

[301] Stefan Klein, Josien P. W. Pluim, Marius Staring, and Max A. Viergever, "Adaptive Stochastic Gradient Descent Optimisation for Image Registration", *International Journal of Computer Vision*, vol. 81, no. 3, pp. 227–239, Aug. 2008. 62

[302] Roshni Bhagalia, Jeffrey a Fessler, and Boklye Kim, "Accelerated nonrigid intensity-based image registration using importance sampling.", *IEEE Transactions on Medical Imaging : TMI*, vol. 28, no. 8, pp. 1208–16, Aug. 2009. 62

[303] Julian Besag, "On the statistical analysis of dirty pictures", *Journal of the Royal Statistical Society, Series B*, vol. 48, no. 3, pp. 259–302, 1986. 64

[304] Paul Bao-Luo Chou and C. M. Brown, "The Theory and Practice of Bayesian Image Labeling", *International Journal of Computer Vision*, vol. 4, no. 3, pp. 185–210, 1990. 64

[305] P. B. Chou, P. R. Cooper, M. J. Swain, C. M. Brown, and L. E. Wixson, "Probabilistic network inference for cooperative high and low level vision", in *Markov Random Fields: Theory and Applications*, R. Chellappa and A. Jain, Eds., pp. 211–43. Academic Press, Boston, MA, 1993. 64

[306] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by Simulated Annealing", *Science*, vol. 220, no. 4598, pp. 671–680, 1983. 64

[307] D. M. Greig, B. T. Porteous, and A. H. Seheult, "Exact maximum a posteriori estimation for binary images", *Journal of the Royal Statistical Society Series B (Methodological)*, vol. 51, no. 2, pp. 271–279, 1989. 64

[308] Judea Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, The Morgan Kaufmann Series in Representation and Reasoning. Morgan Kaufmann, San Francsico, 1988. 64, 65

[309] L.R. Ford and D.R. Fulkerson, *Flows in Networks*, Princeston University Press, New Jersey, 1962. 64

[310] Andrew V Goldberg and Robert E Tarjan, "A new approach to the maximum-flow problem", *Journal of the ACM*, vol. 35, no. 4, pp. 921–940, 1988. 64

[311] Yuri Boykov, Olga Veksler, and Ramin Zabih, "Fast approximate energy minimization via graph cuts", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 23, no. 11, pp. 1222–1239, 2001. 64

[312] Vladimir Kolmogorov and Ramin Zabih, "What energy functions can be minimized via graph cuts?", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 26, no. 2, pp. 147–59, Mar. 2004. 65

[313] Oliver Woodford, Philip Torr, Ian Reid, and Andrew Fitzgibbon, "Global stereo reconstruction under second-order smoothness priors.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 31, no. 12, pp. 2115–28, Dec. 2009. 65

[314] Pushmeet Kohli and Philip H S Torr, "Dynamic graph cuts for efficient inference in Markov Random Fields.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 29, no. 12, pp. 2079–88, Dec. 2007. 65, 134

[315] Tommy W. H. Tang and Albert C. S. Chung, "Non-rigid image registration using graph-cuts.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'07*, Jan. 2007, number Pt 1 in LNCS, pp. 916–24. 65

[316] Ronald W. K. So and Albert C. S. Chung, "Multi-level non-rigid image registration using graph-cuts", in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*. Apr. 2009, pp. 397–400, IEEE. 65

[317] Ronald W. K. So and Albert C. S. Chung, "Non-rigid image registration by using graph-cuts with mutual information", in *2010 IEEE International Conference on Image Processing*. Sep. 2010, pp. 4429–4432, IEEE. 65

[318] Ronald W. K. So, Tommy W. H. Tang, and Albert C. S. Chung, "Non-rigid image registration of brain magnetic resonance images using graph-cuts", *Pattern Recognition*, vol. 44, pp. 2450–2467, Apr. 2011. 65

[319] Brendan J. Frey and David J. C. MacKay, "A revolution: Belief propagation in graphs with cycles", in *In Neural Information Processing Systems*, Michael I Jordan, M J Kearns, and S A Solla, Eds. 1998, vol. 10, pp. 479–485, MIT Press. 65

[320] Kevin Murphy, Yair Weiss, and Michael I Jordan, "Loopy belief propagation for approximate inference: An empirical study", in *Proceedings of Uncertainty in AI*, Kathryn B Laskey and Henri Prade, Eds., 1999, vol. 9, pp. 467–475. 65

[321] Pedro F. Felzenszwalb and Daniel P. Huttenlocher, "Efficient Belief Propagation for Early Vision", *International Journal of Computer Vision*, vol. 70, no. 1, pp. 41–54, May 2006. 65, 66

[322] Stavros Alchatzidis, Aristeidis Sotiras, and Nikos Paragios, "Efficient parallel message computation for MAP inference", in *IEEE International Conference on Computer Vision*. Nov. 2011, Ieee. 65, 66

[323] Alexander Shekhovtsov, Ivan Kovtun, and Václav Hlaváč, "Efficient MRF deformation model for non-rigid image matching", *Computer Vision and Image Understanding*, vol. 112, no. 1, pp. 91–99, Oct. 2008. 65, 66, 121

[324] Ce Liu, Jenny Yuen, and Antonio Torralba, "SIFT flow: dense correspondence across scenes and its applications.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 33, no. 5, pp. 978–94, May 2011. 66

[325] Nikos Komodakis and Georgios Tziritas, "Approximate labeling via graph cuts based on linear programming.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 29, no. 8, pp. 1436–53, Aug. 2007. 66

[326] Nikos Komodakis, Georgios Tziritas, and Nikos Paragios, "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*, vol. 112, no. 1, pp. 14–29, Oct. 2008. 66

[327] Vladimir Kolmogorov, "Convergent tree-reweighted message passing for energy minimization.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 28, no. 10, pp. 1568–83, Oct. 2006. 66

[328] Ben Glocker, Nikos Komodakis, Georgios Tziritas, Nassir Navab, and Nikos Paragios, "Dense image registration through MRFs and efficient linear programming.", *Medical Image Analysis*, vol. 12, no. 6, pp. 731–741, Dec. 2008. 66, 74

[329] Ben Glocker, Nikos Komodakis, Nassir Navab, Georgios Tziritas, and Nikos Paragios, "Dense registration with deformation priors.", in *Information Processing in Medical Imaging : IPMI*, Jan. 2009, vol. 21, pp. 540–551. 66, 134

[330] Darko Zikic, Ben Glocker, Oliver Kutter, Martin Groher, Nikos Komodakis, Ali Kamen, Nikos Paragios, and Nassir Navab, "Linear intensity-based image registration by Markov random fields and discrete optimization.", *Medical Image Analysis*, vol. 14, no. 4, pp. 550–562, Aug. 2010. 66

[331] Dongjin Kwon, Kyong Lee, Il Yun, and Sang Lee, "Nonrigid Image Registration Using Dynamic Higher-Order MRF Model", in *European Conference on Computer Vision : ECCV'08*, 2008. 67

[332] Dongjin Kwon, Kyong Joon Lee, Il Dong Yun, and Sang Uk Lee, "Nonrigid Image Registration Using Higher-Order MRF Model with Dense Local Descriptor", in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 2011. 67

[333] J. Santamaría, O. Cordón, and S. Damas, "A comparative study of state-of-the-art evolutionary image registration methods for 3D modeling", *Computer Vision and Image Understanding*, vol. 115, no. 9, pp. 1340–1354, Sep. 2011. 67

[334] N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies.", *Evolutionary computation*, vol. 9, no. 2, pp. 159–95, Jan. 2001. 67

[335] Dana Paquin, Doron Levy, and Lei Xing, "Hybrid Multiscale Landmark and Deformable Image Registration.", *Mathematical Biosciences and Engineering*, vol. 4, no. 4, pp. 711–737, 2007. 70

[336] Eli Gibson, Ali R Khan, and Mirza Faisal Beg, "A combined surface and volumetric registration (SAVOR) framework to study cortical biomarkers and volumetric imaging data.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'09*, Jan. 2009, number Pt 1 in LNCS, pp. 713–20. 70

[337] Oscar Camara, Gaspar Delso, Olivier Colliot, Antonio Moreno-Ingelmo, and Isabelle Bloch, "Explicit incorporation of prior anatomical information into a nonrigid registration of thoracic and abdominal CT and 18-FDG whole-body emission

PET images.", *IEEE Transactions on Medical Imaging : TMI*, vol. 26, no. 2, pp. 164–78, Feb. 2007. 71

[338] Xiuying Wang and D D Feng, "Automatic hybrid registration for 2-dimensional CT abdominal images", in *Image and Graphics, 2004. Proceedings. Third International Conference on*, 2004, pp. 208–211. 71

[339] Thitiporn Chanwimaluang, Guoliang Fan, and Stephen R. Fransen, "Hybrid retinal image registration.", *IEEE Transactions on Information Technology in Biomedicine : TITB*, vol. 10, no. 1, pp. 129–42, Jan. 2006. 71

[340] Xenophon Papademetris, Andrea P Jackowski, Robert T Schultz, Lawrence H Staib, and James S Duncan, "Integrated Intensity and Point-Feature Nonrigid Registration", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'04*, 2004, LNCS, pp. 763–770. 71

[341] Thomas Hartkens, Derek L. G. Hill, A. Castellano-Smith, David J. Hawkes, Calvin R. Maurer, A. Mar.tin, W. Hall, H. Liu, and C. Truwit, "Using points and surfaces to improve voxel-based non-rigid registration", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'02*. 2002, pp. 565–572, Springer. 71

[342] Karl Rohr, Pascal Cathier, and Stefan Wörz, "Elastic registration of electrophoresis images using intensity information and point landmarks", *Pattern Recognition*, vol. 37, no. 5, pp. 1035–1048, May 2004. 71

[343] Huanxiang Lu, Philippe C. Cattin, Lutz-Peter Nolte, and Mauricio Reyes, "Diffusion weighted imaging distortion correction using hybrid multimodal image registration", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. Mar. 2011, pp. 594–597, IEEE. 71

[344] Christine DeLorenzo, Xenophon Papademetris, Kun Wu, Kenneth P Vives, Dennis Spencer, and James S Duncan, "Nonrigid 3D brain registration using intensity/feature information.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'06*, Jan. 2006, number Pt 1 in LNCS, pp. 932–939. 71

[345] Stefan Wörz and Karl Rohr, "Hybrid Spline-Based Elastic Image Registration Using Analytic Solutions of the Navier Equation", in *Bildverarbeitung für die Medizin 2007*, 2007, pp. 151–155. 72

[346] Andreas Biesdorf, Stefan Wörz, Hans-Jürgen Kaiser, Christoph Stippich, and Karl Rohr, "Hybrid spline-based multimodal registration using local measures for joint entropy and mutual information.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'09*, Jan. 2009, number Pt 1 in LNCS, pp. 607–15. 72

[347] A. Azar, C. Xu, Xavier Pennec, and Nicholas Ayache, "An Interactive Hybrid Non-Rigid Registration Framework for 3D Medical Images", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2006, pp. 824–827, Ieee. 72, 85

[348] Anand Joshi, Richard Leahy, Arthur W Toga, and David Shattuck, "A framework for brain registration via simultaneous surface and volume flow.", in *Information Processing in Medical Imaging : IPMI*, Jan. 2009, vol. 21, pp. 576–88. 73

[349] Ahmed Besbes, Nikos Komodakis, Georg Langs, and Nikos Paragios, "Shape priors and discrete MRFs for knowledge-based segmentation", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2009, pp. 1295–1302, Ieee. 81, 134

[350] Nikos Komodakis, Nikos Paragios, and Georgios Tziritas, "Clustering via lp-based stabilities", in *Neural Information Processing Systems (NIPS)*, 2008. 81, 95, 100

[351] Georg Langs and Nikos Paragios, "Modeling the structure of multivariate manifolds: Shape maps", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2008, pp. 1–8, Ieee. 81

[352] Simon Baker, Daniel Scharstein, J. P. Lewis, Stefan Roth, Michael J. Black, and Richard Szeliski, "A Database and Evaluation Methodology for Optical Flow", *International Journal of Computer Vision*, vol. 92, no. 1, pp. 1–31, Nov. 2010. 82, 100

[353] Nicolas Honnorat, Régis Vaillant, and Nikos Paragios, "Graph-based Geometric-Iconic Guide-wire Tracking", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'11*, 2011. 85

[354] Uday Kurkure, Yen H. Le, Nikos Paragios, James P. Carson, Tao Ju, and Ioannis A. Kakadiaris, "Landmark/Image-based Deformable Registration of Gene Expression Data", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*, 2011, pp. 1089–96. 85

[355] Vincent A. Magnotta, H.Jeremy Bockholt, Hans J. Johnson, Gary E. Christensen, and Nancy C. Andreasen, "Subcortical, cerebellar, and magnetic resonance based consistent brain image registration", *NeuroImage*, vol. 19, no. 2, pp. 233–245, Jun. 2003. 88

[356] Gary E. Christensen, "Inverse consistent registration with object boundary constraints", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2004, pp. 591–594, IEEE. 88

[357] Dinesh Kumar, Xiujuan Geng, Eric A. Hoffman, and Gary E. Christensen, "BICIR:Boundary-Constrained Inverse Consistent Image Registration Using WEB-Splines", in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. 2006, pp. 68–68, Ieee. 88

[358] Xiujuan Geng, Dinesh Kumar, and Gary E. Christensen, "Transitive inverse-consistent manifold registration.", in *Information Processing in Medical Imaging : IPMI*, Jan. 2005, vol. 19, pp. 468–79. 88

[359] Zhijun Zhang, Yifeng Jiang, and Hungtat Tsui, "Consistent multi-modal non-rigid registration based on a variational approach", *Pattern Recognition Letters*, vol. 27, no. 7, pp. 715–725, May 2006. 88

[360] Guozhi Tao, Renjie He, Sushmita Datta, and Ponnada A. Narayana, "Symmetric inverse consistent nonlinear registration driven by mutual information.", *Computer methods and programs in biomedicine*, vol. 95, no. 2, pp. 105–15, Aug. 2009. 88

[361] Peter Rogelj and Stanislav Kovacic, "Symmetric image registration.", *Medical Image Analysis*, vol. 10, no. 3, pp. 484–93, Jul. 2006. 88

[362] Christoph Guetter, H. Xue, Christophe Chefd'Hotel, and J. Guehring, "Efficient symmetric and inverse-consistent deformable registration through interleaved optimization", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2011, pp. 590–593, IEEE. 89

[363] Pascal Cachier and David Rey, "Symmetrization of the non-rigid registration problem using inversion-invariant energies: application to multiple sclerosis", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'00*. 2000, vol. 1935 of *LNCS*, pp. 697–708, Springer Berlin / Heidelberg. 89, 92

[364] Hemant Tagare, David Groisser, and Oskar Skrinjar, "Symmetric non-rigid registration: A geometric theory and some numerical techniques", *Journal of Mathematical Imaging and Vision*, vol. 34, no. 1, pp. 61–88, 2009. 89, 90, 92

[365] Alex D. Leow, Igor Yanovsky, Ming-Chang Chiang, Agatha D. Lee, Andrea D. Klunder, Allen Lu, James T Becker, Simon W. Davis, Arthur W. Toga, and Paul M. Thompson, "Statistical properties of Jacobian maps and the realization of unbiased large-deformation nonlinear image registration.", *IEEE Transactions on Medical Imaging : TMI*, vol. 26, no. 6, pp. 822–32, Jun. 2007. 89

[366] Igor Yanovsky, Paul M. Thompson, Stanley Osher, and Alex D. Leow, "Topology Preserving Log-Unbiased Nonlinear Image Registration: Theory and Implementation", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2007, pp. 1–8, Ieee. 89

[367] John Ashburner, Jesper L. R. Andersson, and Karl J. Friston, "High-dimensional image registration using symmetric priors.", *NeuroImage*, vol. 9, no. 6 Pt 1, pp. 619–28, Jun. 1999. 89

[368] Sarang C. Joshi, Brad Davis, Matthieu Jomier, and Guido Gerig, "Unbiased diffeomorphic atlas construction for computational anatomy.", *NeuroImage*, vol. 23 Suppl 1, pp. S151–60, Jan. 2004. 90, 92, 111

[369] Peter Lorenzen, Marcel Prastawa, Brad Davis, Guido Gerig, Elizabeth Bullitt, and Sarang C. Joshi, "Multi-modal image set registration and atlas formation.", *Medical Image Analysis*, vol. 10, no. 3, pp. 440–51, Jun. 2006. 90

[370] David Raffelt, J-Donald Tournier, Jurgen Fripp, Stuart Crozier, Alan Connelly, and Olivier Salvado, "Symmetric diffeomorphic registration of fibre orientation distributions.", *NeuroImage*, vol. 56, no. 3, pp. 1171–80, Jun. 2011. 90

[371] Vincent Noblet, Christian Heinrich, Fabrice Heitz, and Jean-Paul Armspach, "Symmetric nonrigid image registration: application to average brain templates construction.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI '08*, Jan. 2008, number Pt 2 in LNCS, pp. 897–904. 90, 91, 92, 111, 114

[372] Brendan J. Frey and Delbert Dueck, "Clustering by passing messages between data points.", *Science*, vol. 315, no. 5814, pp. 972–976, 2007. 95

[373] Ben Glocker, Nikos Paragios, Nikos Komodakis, Georgios Tziritas, and Nassir Navab, "Optical flow estimation with uncertainties through dynamic MRFs", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2008, pp. 1–8, Ieee. 104

[374] Tristrom Cooke, "Two Applications of Graph-Cuts to Image Processing", in *2008 Digital Image Computing: Techniques and Applications*. 2008, pp. 498–504, Ieee. 104

[375] Alexandre Guimond, Jean Meunier, and Jean-Philippe Thirion, "Average Brain Models: A Convergence Study", *Computer Vision and Image Understanding*, vol. 77, no. 2, pp. 192–210, Feb. 2000. 110

[376] Gary E Christensen, Hans J Johnson, and Michael W Vannier, "Synthesizing average 3D anatomical shapes.", *NeuroImage*, vol. 32, no. 1, pp. 146–58, Aug. 2006. 110

[377] Dieter Seghers, Emiliano D'Agostino, Frederik Maes, Dirk Vandermeulen, and Paul Suetens, "Construction of a brain template from MR images using state-of-the-art registration and segmentation techniques", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI '04*. 2004, vol. 3216 of *LNCS*, pp. 696–703, Springer. 110

[378] Hyunjin Park, Peyton H. Bland, Alfred O. Hero III, and Charles R. Meyer, "Least biased target selection in probabilistic atlas construction.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'05*, Jan. 2005, number Pt 2 in LNCS, pp. 419–26. 110

[379] Jihun Hamm, Dong Hye Ye, Ragini Verma, and Christos Davatzikos, "GRAM: A framework for geodesic registration on anatomical manifolds.", *Medical Image Analysis*, vol. 14, no. 5, pp. 633–42, Oct. 2010. 110

[380] Stephen Marsland, Carole J. Twining, and Christopher J. Taylor, "A minimum description length objective function for groupwise non-rigid image registration", *Image and Vision Computing*, vol. 26, no. 3, pp. 333–346, Mar. 2008. 111

[381] P. Thomas Fletcher, Suresh Venkatasubramanian, and Sarang C. Joshi, "The geometric median on Riemannian manifolds with application to robust atlas estimation.", *NeuroImage*, vol. 45, no. 1 Suppl, pp. S143–52, Mar. 2009. 111

[382] Brian Avants and James C. Gee, "Geodesic estimation for large deformation anatomical shape averaging and interpolation", *NeuroImage*, vol. 23, pp. S139–S150, 2004. 111

[383] Kanwal K Bhatia, Jo Hajnal, Alexander Hammers, and Daniel Rueckert, "Similarity metrics for groupwise non-rigid registration.", in *Medical Image Computing and*

Computer-Assisted Intervention : MICCAI'07, Jan. 2007, number Pt 2 in LNCS, pp. 544–52. 111

[384] Guorong Wu, Hongjun Jia, Qian Wang, and Dinggang Shen, "SharpMean: groupwise registration guided by sharp mean image and tree-based registration.", *NeuroImage*, vol. 56, no. 4, pp. 1968–81, Jun. 2011. 111

[385] B.C. Munsell and A. Temlyakov, "Fast multiple shape correspondence by preorganizing shape instances", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2009, pp. 840–847, Ieee. 111

[386] Hongjun Jia, Pew-Thian Yap, Guorong Wu, Qian Wang, and Dinggang Shen, "Intermediate templates guided groupwise registration of diffusion tensor images.", *NeuroImage*, vol. 54, no. 2, pp. 928–939, Sep. 2010. 111

[387] Simon Baker, Iain Matthews, and Jeff Schneider, "Automatic construction of active appearance models as an image coding problem.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 26, no. 10, pp. 1380–4, Oct. 2004. 112

[388] Timothy F Cootes, Stephen Marsland, Carole J. Twining, K. Smith, and Christopher J. Taylor, "Groupwise diffeomorphic non-rigid registration for automatic model building", in *European Conference on Computer Vision : ECCV'04*, T. Pajdla and Jiri Matas, Eds., vol. 3024, pp. 316–327. Springer Berlin / Heidelberg, 2004. 112

[389] Colin Studholme and Valerie Cardenas, "A template free approach to volumetric spatial normalization of brain anatomy", *Pattern Recognition Letters*, vol. 25, no. 10, pp. 1191–1202, Jul. 2004. 112

[390] Erik G Miller, N.E Matsakis, and Paul A. Viola, "Learning from one example through shared densities on transforms", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. 2000, number 1, pp. 464–471, IEEE Comput. Soc. 112, 115

[391] Erik G Learned-Miller, "Data driven image models through continuous joint alignment.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 28, no. 2, pp. 236–50, Feb. 2006. 112, 115

[392] Lilla Zöllei, E Learned-Miller, Eric Grimson, and W. Wells, "Efficient population registration of 3D data", in *Computer Vision for Biomedical Image Applications*,

Changshui Liu, Yanxi and Jiang, Tianzi and Zhang, Ed., vol. 3765, pp. 291–301. Springer Berlin / Heidelberg, 2005. 112

[393] Qian Wang, Guorong Wu, Pew-Thian Yap, and Dinggang Shen, "Attribute vector guided groupwise registration.", *NeuroImage*, vol. 50, no. 4, pp. 1485–96, May 2010. 112

[394] Gary B. Huang, Vidit Jain, and Erik Learned-Miller, "Unsupervised Joint Alignment of Complex Images", in *International Conference on Computer Vision*. 2007, pp. 1–8, Ieee. 112

[395] Christian Wachinger, Wolfgang Wein, and Nassir Navab, "Three-dimensional ultrasound mosaicing.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI '07*, Jan. 2007, number Pt 2 in LNCS, pp. 327–35. 112

[396] Christian Wachinger and Nassir Navab, "Similarity metrics and efficient optimization for simultaneous registration", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2009, pp. 667–674, IEEE. 112, 114, 116

[397] Xiujuan Geng, Gary E Christensen, Hong Gu, Thomas J Ross, and Yihong Yang, "Implicit reference-based group-wise image registration and its application to structural and functional MRI.", *NeuroImage*, vol. 47, no. 4, pp. 1341–51, Oct. 2009. 113, 114, 116

[398] Hongjun Jia, Guorong Wu, Qian Wang, and Dinggang Shen, "ABSORB: Atlas Building by Self-organized Registration and Bundling.", *NeuroImage*, vol. 51, no. 3, pp. 1057–70, Jul. 2010. 113

[399] Sajjad Baloch and Christos Davatzikos, "Morphological appearance manifolds in computational anatomy: groupwise registration and morphological analysis.", *NeuroImage*, vol. 45, no. 1 Suppl, pp. S73–85, Mar. 2009. 113

[400] Nai-Xiang Lian and Christos Davatzikos, "Morphological Appearance Manifolds for Group-wise Morphometric Analysis", *Medical Image Analysis*, Jul. 2011. 113

[401] Dong H. Ye, Killian M. Pohl, Harold Litt, and Christos Davatzikos, "Groupwise morphometric analysis based on high dimensional clustering", in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. 2010, pp. 47–54, IEEE. 113

[402] Rhodri H. Davies, Carole J. Twining, and Christopher J. Taylor, "Groupwise surface correspondence by optimization: representation and regularization.", *Medical Image Analysis*, vol. 12, no. 6, pp. 787–96, Dec. 2008. 113

[403] Georg Langs, René Donner, Philipp Peloschek, and Horst Bischof, "Robust autonomous model learning from 2D and 3D data sets.", in *Medical Image Computing and Computer-Assisted Intervention : MICCAI'07*, Jan. 2007, number Pt 1 in LNCS, pp. 968–76. 113

[404] Fei Wang, Tanveer Syeda-Mahmood, Baba C Vemuri, David Beymer, and Anand Rangarajan, "Closed-form Jensen-Renyi divergence for mixture of Gaussians and applications to group-wise shape registration.", , no. Pt 1, pp. 648–55, Jan. 2009. 113, 118

[405] Maxime Taron, Nikos Paragios, and Marie-Pierre Jolly, "Registration with uncertainties and statistical modeling of shapes with variable metric kernels.", *IEEE Transactions on Pattern Analysis and Machine Intelligence : TPAMI*, vol. 31, no. 1, pp. 99–113, Jan. 2009. 118

[406] Mickael Savinaud, Aristeidis Sotiras, Serge Maitrejean, and Nikos Paragios, "Bioluminescence enhancement through fusion of optical imaging and cinematic video flow", in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2010, pp. 688–691, Ieee. 120

[407] Hiroshi Ishikawa, "Higher-order clique reduction in binary graph cut", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2009, pp. 2993–3000, IEEE. 122

[408] Carsten Rother, Pushmeet Kohli, Wei Feng, and Jiaya Jia, "Minimizing sparse higher order energy functions of discrete variables", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2009, pp. 1382–1389, IEEE. 122

[409] Nikos Komodakis and Nikos Paragios, "Beyond pairwise energies: Efficient optimization for higher-order MRFs", in *IEEE Conference on Computer Vision and Pattern Recognition : CVPR*. Jun. 2009, pp. 2985–2992, IEEE. 122

[410] Mert R Sabuncu, Serdar K Balci, Martha E Shenton, and Polina Golland, "Image-driven population analysis through mixture modeling.", *IEEE Transactions on Medical Imaging : TMI*, vol. 28, no. 9, pp. 1473–87, Sep. 2009. 125

[411] Julie Delon, "Midway Image Equalization", *Journal of Mathematical Imaging and Vision*, vol. 21, no. 2, pp. 119–134, Sep. 2004. 126