

Geodesic Active Fields: A Geometric Framework for Image Registration

THÈSE N° 5175 (2011)

PRÉSENTÉE LE 11 NOVEMBRE 2011

À LA FACULTÉ SCIENCES ET TECHNIQUES DE L'INGÉNIEUR
LABORATOIRE DE TRAITEMENT DES SIGNAUX 5
PROGRAMME DOCTORAL EN GÉNIE ÉLECTRIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Dominique ZOSSO

acceptée sur proposition du jury:

Dr J.-M. Vesin, président du jury
Prof. J.-Ph. Thiran, directeur de thèse
Prof. X. Bresson, rapporteur
Prof. N. Sochen, rapporteur
Prof. M. Unser, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2011

Abstract

Image registration is the concept of mapping homologous points in a pair of images. In other words, one is looking for an underlying deformation field that matches one image to a target image. The spectrum of applications of image registration is extremely large: It ranges from bio-medical imaging and computer vision, to remote sensing or geographic information systems, and even involves consumer electronics.

Mathematically, image registration is an inverse problem that is ill-posed, which means that the exact solution might not exist or not be unique. In order to render the problem tractable, it is usual to write the problem as an energy minimization, and to introduce additional regularity constraints on the unknown data. In the case of image registration, one often minimizes an image mismatch energy, and adds an additive penalty on the deformation field regularity as smoothness prior.

Here, we focus on the registration of the human cerebral cortex. Precise cortical registration is required, for example, in statistical group studies in functional MR imaging, or in the analysis of brain connectivity. In particular, we work with spherical inflations of the extracted hemispherical surface and associated features, such as cortical mean curvature. Spatial mapping between cortical surfaces can then be achieved by registering the respective spherical feature maps. Despite the simplified spherical geometry, inter-subject registration remains a challenging task, mainly due to the complexity and inter-subject variability of the involved brain structures.

In this thesis, we therefore present a registration scheme, which takes the peculiarities of the spherical feature maps into particular consideration. First, we realize that we need an appropriate hierarchical representation, so as to coarsely align based on the important structures with greater inter-subject stability, before taking smaller and more variable details into account. Based on arguments from brain morphogenesis, we propose an anisotropic scale-space of mean-curvature maps, built around the Beltrami framework.

Second, inspired by concepts from vision-related elements of psycho-physical Gestalt theory, we hypothesize that anisotropic Beltrami regularization better suits the requirements of image registration regularization, compared to traditional Gaussian filtering. Different objects in an image should be allowed

to move separately, and regularization should be limited to within the individual Gestalts. We render the regularization feature-preserving by limiting diffusion across edges in the deformation field, which is in clear contrast to the indifferent linear smoothing. We do so by embedding the deformation field as a manifold in higher-dimensional space, and minimize the associated Beltrami energy which represents the hyperarea of this embedded manifold as measure of deformation field regularity. Further, instead of simply adding this regularity penalty to the image mismatch in lieu of the standard penalty, we propose to incorporate the local image mismatch as weighting function into the Beltrami energy. The image registration problem is thus reformulated as a weighted minimal surface problem. This approach has several appealing aspects, including (1) invariance to re-parametrization and ability to work with images defined on non-flat, Riemannian domains (e.g., curved surfaces, scale-spaces), and (2) intrinsic modulation of the local regularization strength as a function of the local image mismatch and/or noise level. On a side note, we show that the proposed scheme can easily keep up with recent trends in image registration towards using diffeomorphic and inverse consistent deformation models.

The proposed registration scheme, called Geodesic Active Fields (GAF), is non-linear and non-convex. Therefore we propose an efficient optimization scheme, based on splitting. Data-mismatch and deformation field regularity are optimized over two different deformation fields, which are constrained to be equal. The constraint is addressed using an augmented Lagrangian scheme, and the resulting optimization problem is solved efficiently using alternate minimization of simpler sub-problems. In particular, we show that the proposed method can easily compete with state-of-the-art registration methods, such as Demons.

Finally, we provide an implementation of the fast GAF method on the sphere, so as to register the triangulated cortical feature maps. We build an automatic parcellation algorithm for the human cerebral cortex, which combines the delineations available on a set of atlas brains in a Bayesian approach, so as to automatically delineate the corresponding regions on a subject brain given its feature map. In a leave-one-out cross-validation study on 39 brain surfaces with 35 manually delineated gyral regions, we show that the pairwise subject-atlas registration with the proposed spherical registration scheme significantly improves the individual alignment of cortical labels between subject and atlas brains, and, consequently, that the estimated automatic parcellations after label fusion are of better quality.

Keywords: Image Registration, Computer Vision, Inverse Problem, Regularization, Human Cortex, Gestalt Theory, Scale-Space, Beltrami Framework, Polyakov Energy, Minimal Surface, Differential Geometry, Computational Geometry, Sphere, Optimization, Splitting, Augmented Lagrangian, Diffeomorphism, Classification.

Résumé

Le recalage d'images est le concept de mettre en relation les points homologues dans une paire d'images. En d'autres termes, on est à la recherche d'un champ de déformation sous-jacent qui fait correspondre une image à une autre image cible. Le spectre d'applications du recalage d'images est extrêmement important : il s'étend de l'imagerie bio-médicale et vision par ordinateur jusqu'à la télédétection ou les systèmes d'information géographique, et implique même l'électronique grand public.

Mathématiquement, le recalage d'images est un problème inverse qui est mal posé, ce qui signifie que la solution exacte peut ne pas exister ou ne pas être unique. Afin de rendre le problème docile, il est habituel de le récrire comme une minimisation d'énergie, et d'introduire des contraintes de régularité supplémentaire sur les données inconnues. Ainsi on minimise souvent la différence entre les images après déformation, et ajoute une pénalité additive sur la régularité du champ de déformation.

Ici, nous nous intéressons au recalage du cortex cérébral humain. Un recalage cortical précis est nécessaire, par exemple, dans les études statistiques de groupes en IRM fonctionnelle, ou dans l'analyse de la connectivité cérébrale. En particulier, nous travaillons avec des gonflements sphériques de la surface hémisphérique extraite, et de caractéristiques associées, telles que la courbure moyenne corticale. La relation spatiale entre les surfaces corticales peut alors être établie en recalant les cartes sphériques respectives. Malgré la géométrie sphérique simplifiée, le recalage reste une tâche difficile, principalement en raison de la complexité des structures cérébrales impliquées et de leur variabilité d'une personne à l'autre.

Dans cette thèse, nous présentons donc un système de recalage, qui prend les particularités des cartes sphériques en considération. Tout d'abord, nous avons besoin d'une représentation hiérarchique appropriée, afin de pouvoir aligner grossièrement sur la base des structures importantes, ayant une plus grande stabilité entre sujets, avant de prendre en compte des plus petits détails. A partir d'arguments basés sur la morphogenèse du cerveau, nous proposons un espace-échelle anisotrope de la carte de courbure moyenne.

Deuxièmement, inspiré par les concepts de vision liés aux éléments psychophysique de la théorie de la Gestalt, nous faisons l'hypothèse que le filtrage

anisotrope Beltrami répond mieux aux exigences de régularisation du recalage d'images, que le traditionnel filtrage gaussien. Différents objets dans une image sont autorisés à se déplacer séparément, et la régularisation est confinée dans l'intérieur de la Gestalt individuelle. Nous rendons la régularisation préservatrice aux caractéristiques en limitant la diffusion à travers des bords du champ de déformation. Nous obtenons ceci en interprétant le champ de déformation comme une variété dans un espace de dimension supérieure, et en minimisant l'énergie de Beltrami associée. Cette dernière représente l'hyper-surface de cette variété en tant que mesure de la régularité du champ de déformation. De plus, au lieu de simplement ajouter cette pénalité de régularité à la place du terme Gaussien standard, nous proposons d'intégrer le décalage local des images comme fonction de pondération dans l'énergie Beltrami. Le recalage est ainsi reformulé comme un problème de surface pondérée minimale. Cette approche a plusieurs aspects attrayants, y compris (1) l'invariance à la re-paramétrisation et la capacité de travailler avec des images définies sur de variétés riemanniennes non-planes (par exemple, des surfaces courbées, des espaces-échelle), et (2) la modulation intrinsèque et locale de la force de régularisation en fonction du décalage image local et / ou du niveau de bruit local.

Le système de recalage proposé, appelé champs actifs géodésiques, est non linéaire et non-convexe. C'est pourquoi nous proposons un schéma d'optimisation efficace, basé sur le fractionnement. Le terme de décalage et la régularité du champ de déformation sont optimisés sur deux champs de déformation différents, qui sont contraints à être égaux. La contrainte est adressée en utilisant un schéma Lagrangien augmenté, et le problème d'optimisation résultant est résolu efficacement en utilisant la minimisation alternante de sous-problèmes simples. En particulier, nous montrons que la méthode proposée peut facilement rivaliser avec les méthodes de recalage de l'état de l'art, comme les démons.

Enfin, nous fournissons une implémentation de la méthode rapide sur la sphère, afin de recalculer les cartes corticales triangulées. En outre, nous construisons un algorithme pour la parcellisation automatique du cortex cérébral humain, qui combine les étiquettes disponibles sur un ensemble de cerveaux atlas dans une approche bayésienne, afin de délimiter automatiquement les régions correspondantes sur un sujet, dont seule la carte sphérique est donnée. Dans une étude de validation croisée sur 39 hémisphères, chacune avec 35 régions gyrales délimitées manuellement, nous montrons que le recalage sphérique améliore significativement l'alignement individuel des étiquettes corticales entre le sujet et les cerveaux de référence. Par conséquent, les parcellations automatiques après la fusion d'étiquettes sont de meilleure qualité, elles aussi.

Mots-clés: Recalage d'images, vision par ordinateur, problème inverse, régularisation, cortex humain, théorie de la Gestalt, espace-échelle, méthode Beltrami, énergie Polyakov, surface minimale, géométrie différentielle, géométrie par ordinateur, sphère, optimisation, fractionnement, Lagrangien augmenté, difféomorphisme, classification.

Zusammenfassung

Bildregistrierung ist das Konzept der gegenseitigen Zuordnung homologer Punkte in einem Paar von Bildern. In anderen Worten, es wird ein zugrunde liegendes Deformationsfeld gesucht, welches ein Bild bestmöglich auf ein Zielbild abbildet. Das Anwendungsspektrum von Bildregistrierung ist extrem gross: Es reicht von der biomedizinischen Bildverarbeitung und Computervision bis zu Fernerkundung und geografischen Informationssystemen.

Mathematisch ist Bildregistrierung ein schlecht gestelltes inverses Problem, d.h. die exakte Lösung ist möglicherweise nicht vorhanden oder nicht eindeutig. Um ein solches Problem lösbar zu machen, formuliert man es üblicherweise als eine Energieminimierung und führt weitere Einschränkungen bezüglich Regularität der unbekannten Daten ein. Im Fall von Bildregistrierung, minimiert man oft die nach Abbildung verbleibende Bilddifferenz und wählt eine additive Strafe für Unregelmässigkeit der Verformung.

Hier konzentrieren wir uns auf die Registrierung der menschlichen Grosshirnrinde. Präzise kortikale Registrierung ist z.B. in statistischen Gruppen-Studien der funktionellen MR-Bildgebung erforderlich, oder ebenso in der Analyse der Gehirn-Konnektivität. Insbesondere arbeiten wir mit sphärischen Aufblähungen der extrahierten hemisphärischen Kortex-Oberfläche und damit verbundenen Features. Räumliche Zuordnung zwischen kortikalen Flächen verschiedener Gehirne kann demnach mittels Registrierung der jeweiligen sphärischen Karten erreicht werden. Trotz der vereinfachten sphärischen Geometrie bleibt dies eine anspruchsvolle Aufgabe, vor allem aufgrund der Komplexität und interindividuellen Variabilität der beteiligten Hirnstrukturen.

In dieser Arbeit haben wir ein Registrierungssystem entwickelt, das die Besonderheiten der sphärischen Karten speziell berücksichtigt. Erstens erkennen wir, dass wir eine entsprechende hierarchische Darstellung benötigen, um die Karten zuerst grob anhand der wichtigen Strukturen mit grösserer interindividueller Stabilität ausrichten zu können, bevor kleinere und variabelere Details miteinbezogen werden. Basierend auf Argumenten der Gehirn-Morphogenese schlagen wir einen anisotropen Skalen-Raum der mittleren Krümmungen vor.

Zweitens vermuten wir, inspiriert durch Konzepte aus Elementen der Gestalt-Theorie, dass anisotrope Beltrami-Regularisierung verglichen mit der herkömmlichen Gauss-Glättung besser den Anforderungen der Bildregistrierung

entspricht. Verschiedene Objekte in einem Bild sollten sich unabhängig voneinander bewegen können, und Regularisierung auf das Innere der jeweiligen Gestalten beschränkt sein. Wir machen die Regularisierung kantenerhaltend, indem die Diffusion über Kanten des Deformationsfeldes eingeschränkt wird, was in deutlichem Kontrast zur gleichgültigen linearen Glättung steht. Wir tun dies mittels Einbettung des Deformationsfeldes als eine Mannigfaltigkeit in einen höherdimensionalen Raum, und minimieren die damit verbundene Beltrami-Energie, welche die Oberfläche dieser eingebetteten Mannigfaltigkeit als Regularitätsmass des Deformationsfeldes darstellt. Weiter, anstatt einfach nur den standard-mässigen Gauss-Term durch diese Beltrami-Energie zu ersetzen, schlagen wir vor, den lokalen Abbildungsfehler als Gewichtungsfunktion in die Beltrami-Energie zu integrieren. Das Bildregistrierungs-Problem wird somit zu einem Problem minimaler gewichteter Oberfläche umformuliert. Dieser Ansatz hat mehrere attraktive Eigenschaften, einschliesslich (1.) der Invarianz bezüglich Parametrisierung und der Fähigkeit, mit Bildern in nicht-ebenen, Riemannschen Räumen arbeiten zu können (z.B. gekrümmte Oberflächen und Skalen-Räume), und (2.) der intrinsischen Modulation der lokalen Regularisierungsstärke als Funktion des lokalen Ausrichtungsfehlers und Rausch-Pegels. Wir zeigen, dass die vorgeschlagene Methode mit jüngsten Trends in Richtung diffeomorpher und invers-konsistenter Deformationsmodelle mithalten kann.

Das vorgeschlagene Registrierungssystem, genannt Aktive Geodätische Felder, ist nicht-linear und nicht-konvex. Daher schlagen wir eine effiziente Optimierung vor, die auf Variablen-Duplikation basiert. Ausrichtungsfehler und Deformationsfeld-Regularität werden über zwei verschiedene Deformationsfelder optimiert, die aber gleich sein müssen. Diese Einschränkung wird unter Verwendung eines erweiterten Lagrange-Schemas angepackt und das sich daraus ergebende Optimierungs-Problem kann effizient mit alternierender Minimierung der einfacheren Teilprobleme gelöst werden. Insbesondere erreichen wir so, dass das vorgeschlagene Verfahren den Stand der Technik übertrifft.

Schliesslich erarbeiten wir eine Implementierung auf der Kugel um so die triangulierten kortikalen Karten zu registrieren. Wir entwickeln einen automatischen Parzellierungs-Algorithmus, der die manuellen Kartierungen auf einer Reihe von Atlas-Gehirnen in einem bayesianischen Ansatz kombiniert, um so aufgrund der sphärischen Karte eines spezifischen Gehirnes seine automatische Beschriftung der entsprechenden Regionen zu erzielen. In einer Kreuzvalidierungs-Studie mit 39 Gehirn-Oberflächen zeigen wir, dass die paarweise Registrierung die individuelle Ausrichtung der kortikalen Regionen deutlich verbessert, und folglich, dass die darauf beruhende automatische Parzellierung an Qualität gewinnt.

Stichwörter: Bildregistrierung, Computervision, Inverses Problem, Regularisierung, Menschliche Grosshirnrinde, Gestalt Theorie, Skalen-Raum, Beltrami Ansatz, Polyakov Energie, Minimale Oberfläche, Differentialgeometrie, Computergeometrie, Kugeloberfläche, Optimierung, Variablen-Duplikation, erweitertes Lagrange Schema, Diffeomorphismus, Klassifizierung.

Acknowledgments

Research is—to a greater extent than commonly expected—the fruit of serendipity. Despite all plans and intents, targeted studying and focused work, the most interesting findings seem to simply pop up incidentally; suddenly and out of nowhere. In reality, however, they require the accumulation of knowledge, experience and skills on the one hand, and creativity and inspiration on the other hand, most often in a series of apparently unrelated events. Eventually, after some period of “silent fermentation”, a small spark is enough to trigger the emergence of what might turn out to be an interesting idea. In this sense, although I have crafted the work presented in this thesis on my own, it was not without the valuable contribution of many people who provided teaching, time, tools and inspiration. Therefore, a few acknowledgments are more than just in order.

First, I like to thank my supervisor, Prof. Jean-Philippe Thiran, for letting me “ferment” in his lab: I appreciated very much your gentle guidance and the large freedom you granted me. I believe, that your confidence has not been disappointed.

In addition, I am grateful to Prof. Xavier Bresson (City University, Hong Kong), for participating in much of the research reported in this thesis. In all the discussions we had, you have always been a great “catalyzer”: You quickly triage ideas and you are always able to bring another valuable perspective into the game.

Then, I wish to thank the other members of my thesis Jury—Prof. Michael Unser (EPFL), Prof. Nir Sochen (University of Tel Aviv), MER Dr. Jean-Marc Vesin (EPFL)—for the careful reading and evaluation of the present work. Your challenging questions and the discussions were an interesting feedback and your encouraging comments are an honor to me.

I also want to thank the members of the “variational group,” especially Virginia Estellers and Subrahmanyam Gorthi, for the inspiring discussions on the white board and around coffee/tea/hot water.

More generally, I wish to express my gratitude to all the ones, who shared time and space with me during the last few years. This obviously includes all the members of the various Signal Processing Labs in ELD/ELE—of which I will not list all names in order to avoid any omissions. You have been a great

crowd! Especially, many thanks to all contributors and customers of “what is not a cafeteria.” In particular, I would like to mention Dr. Alessandro Daducci, who had the arguable honor to share the office with me during what is traditionally the worst phase of the thesis: its writing. Thank you for enduring me, thank you for your encouraging comments nonetheless.

I further want to thank Dr. Samuel Stutz, who regularly spent lunch-time with us and whom we always had odd and interesting discussions with. I’ll get to the other part in “we”, in a minute.

A very warm thank you to all other friends out there. At this point, I apologize to all the people I wrongfully did not explicitly mention so far.

During the four years spent in this lab, the present thesis was not my only “baby”. In fact, our boy twins were born at the very beginning of my dissertation enterprise. Obviously, it was only possible to dedicate so much time to my thesis thanks to the Nursery and Garderie Polychinelle at EPFL. I therefore wish to warmly thank all the personnel of Polychinelle for the care you gave to our twins.

Moreover, I am grateful to my parents. Obviously, they contributed essentially to my education by making it possible to study at EPFL. Beyond, they also had their share of “twin-care,” both on a regular and an “emergency” basis. Without their help, we could not have realized this important career step.

Most importantly, I wish to thank my significant other, the other part in “we,” soon-to-be-Dr. Anja Kunze, for all the time we have together. Research in general—and thesis writing in particular—is not always easy, and yet you supported me throughout the whole journey. So will I. We have come this far—why would things not work out for the rest of our plans?

Finally, I thank our two little boys, Valentin and Florentin, for all what they are to us. I wish to be a good parent for them and hope that, some day, they will be as proud of their Papa as Papa is proud of them.

Contents

Abstract	i
Résumé	iii
Zusammenfassung	v
Acknowledgments	vii
Contents	ix
List of Figures	xv
List of Tables	xix
List of Symbols and Notations	xxi
1 Introduction	1
1.1 Biomedical Context and Motivation	2
1.1.1 Functional mapping	4
1.1.2 Brain connectivity	5
1.2 A Quick Review on Image Registration	6
1.2.1 Image dissimilarity metric	6
1.2.2 Deformation model and regularization	6
1.3 Registration of the Cerebral Cortex	7
1.3.1 Voxel-based versus landmark-based	7
1.3.2 Surface-based	8
1.3.3 What next?	8
1.4 Inverse Problems in Computer Vision	9
1.4.1 Ill-posedness	9
1.4.2 Wiener Filtering	10
1.4.3 Non-linear Regularization	11
1.5 Regularization in Motion Estimation	11
1.5.1 Aperture Problem	12
1.5.2 Variational Model	13

1.5.3	Anisotropic Regularization	13
1.6	Elements of Gestalt Theory and Motion Perception	14
1.6.1	Gestalt Theory	14
1.6.2	Gestalt principles	15
1.6.3	Prägnanz and the Gestalt Laws	15
1.6.4	Motion coherence	17
1.7	A New Registration Framework	18
1.7.1	Gaussian Regularization Shortcomings	18
1.7.2	Drawbacks of Additive Regularization	19
1.7.3	Gestalt and the Beltrami Framework	20
1.7.4	Weighted Beltrami	20
1.8	Organization of this thesis	21
1.8.1	Contributions	21
1.8.2	Structure	22
2	The Beltrami Framework	25
2.1	Preliminaries (1) – Definitions from Differential Geometry	26
2.1.1	First fundamental form I	26
2.1.2	Second fundamental form II	27
2.1.3	On curvature	28
2.1.4	Higher dimensions and co-dimensions	29
2.1.5	First variation of area	29
2.2	Preliminaries (2) – On Scale and Space	29
2.2.1	Linear Scale Space	30
2.2.2	Relation to regularization	30
2.2.3	Anisotropy	31
2.3	Beltrami Embedding, Geodesic Distance and Edge Detection	33
2.3.1	Embedding	33
2.3.2	Induced Metric and Edge Detection	34
2.3.3	Polyakov Energy	34
2.3.4	Area Minimization, Anisotropic Diffusion and Feature Preservation	35
2.3.5	Examples	36
2.3.6	Weighted Beltrami Energy	37
2.3.7	Parametrization invariance	38
2.4	Some Applications	39
2.4.1	Image Denoising, Wiener Filtering, and Blind Deconvolution	40
2.4.2	Image Segmentation	40
2.4.3	Motion Estimation Regularization	42
2.4.4	Registration with Topological Defects	43
2.4.5	Non-standard Image Geometries	44
2.5	Numerical Schemes	44
2.5.1	Explicit, semi-implicit and implicit	44
2.5.2	Operator splitting techniques	45

2.5.3	Vector extrapolation	46
2.5.4	Decoupling schemes	47
3	Cortical Scale space	49
3.1	Introduction	50
3.1.1	From gyrification to scale-space	50
3.1.2	Beltrami flow on triangulated manifolds	51
3.1.3	Discretizations on 2D triangulations in 3D	52
3.1.4	Generalization to arbitrary co-dimension	53
3.2	Proposed Beltrami scale space	54
3.2.1	Comparison with Lopez' discretization	54
3.2.2	Numerical PDE solution	55
3.2.3	Histogram equalization	55
3.2.4	Scale space	56
3.3	Results	57
3.3.1	Equalization	57
3.3.2	Scale-space	57
3.3.3	Choice of β	58
3.3.4	Morphogenic signification	58
3.4	Conclusions	58
4	Geodesic Active Fields	63
4.1	Motivation	64
4.2	Geodesic Active Fields	65
4.3	Stereo Vision and Image Registration	68
4.3.1	The general Euclidean case	68
4.3.2	Stereo vision	68
4.3.3	2-D image registration	69
4.3.4	Registration on non-flat manifolds	70
4.4	Multiscale Image Registration	71
4.4.1	Motivation	71
4.4.2	Multiscale active deformation fields	72
4.4.3	Multiscale 2-D image registration	73
4.5	Weighting Function for the Registration Problem	73
4.5.1	Deformation model	73
4.5.2	Squared error	74
4.5.3	Local joint entropy	75
4.5.4	Absolute error	76
4.5.5	Data term and regularization balancing	76
4.6	Results	77
4.6.1	Mean curvature estimation	77
4.6.2	A few words on β and regularization	77
4.6.3	Application to stereo vision	78
4.6.4	Application to medical imaging	79
4.6.5	Registration on non-flat manifolds	81

4.6.6	Multiscale image registration	81
4.7	Discussion and Conclusion	81
4.7.1	Contributions	84
4.7.2	Conclusions	86
5	A Fast Scheme	87
5.1	Geodesic Active Fields – Direct Implementation	88
5.1.1	Discretization	88
5.1.2	Direct implementation for Cartesian images	89
5.1.3	Discretized Laplace-Beltrami operator	89
5.2	Lagrangian Multiplier	90
5.2.1	Classical Lagrange multiplier	90
5.2.2	Augmented Lagrangian	90
5.2.3	Dealing with multiple constraints	91
5.3	FastGAF Energy Minimization	91
5.3.1	Splitting	92
5.3.2	Augmented Lagrangian	92
5.3.3	Subminimization w.r.t. \mathbf{u}	93
5.3.4	Balancing the computational complexity	93
5.3.5	Subminimization w.r.t. \mathbf{v}	94
5.3.6	A Jacobi scheme for approximate inversion	94
5.3.7	Complete approximate-inverse FastGAF algorithm	96
5.4	Experiments	96
5.4.1	Multiresolution optimization and performance metrics	96
5.4.2	Stereo vision	96
5.4.3	2D registration	97
5.4.4	Hemi-spherical registration	97
5.5	Discussion	99
5.5.1	Algorithms and performance comparison	100
5.5.2	The roles of the parameters r , L and J	102
5.6	Conclusions	103
6	Fast Scheme on the Sphere	107
6.1	GAF on the sphere	107
6.1.1	Triangulated spherical images	108
6.1.2	Local Coordinates and Parallel Transport	108
6.1.3	Deformation Field	110
6.1.4	Data term	110
6.1.5	Deformation Field Embedding and Energy	111
6.2	Minimizing Scheme	113
6.2.1	Discretization and Splitting	113
6.2.2	Data-term optimization	114
6.2.3	Regularity-term optimization	115
6.2.4	Spherical tools	116
6.3	Results	116

6.3.1	Impulse response and denoising	116
6.3.2	Synthetic image pair	117
6.4	Discussion and Conclusions	118
6.4.1	Trivial solutions and rigid registration	118
6.4.2	Matrix inversion	118
6.4.3	Data interpolation	119
7	Mapping the Human Brain	121
7.1	Background and Concept	122
7.1.1	Proposed parcellation scheme	122
7.1.2	Leave-one-out cross-validation	122
7.2	Materials and Methods	123
7.2.1	Brain Data	123
7.2.2	Registration scheme	125
7.2.3	Classification Scheme	126
7.2.4	Scripting and Process Monitoring	129
7.3	Results and Discussion	130
7.3.1	Measuring label consensus: Dice's similarity coefficient	130
7.3.2	Registration	131
7.3.3	Label fusion scheme	133
7.3.4	Parcellation Cross-validation	133
7.4	Conclusions and Outlook	140
7.4.1	Comparison to the state-of-the-art	140
7.4.2	Potential improvements of the registration scheme . . .	140
7.4.3	Potential improvements of the label-fusion	141
8	Diffeomorphic Deformations	143
8.1	Background	144
8.1.1	Exponential maps	144
8.1.2	Diffeomorphic Demons	144
8.1.3	Deformation model in GAF	145
8.2	Diffeomorphic Deformation Model	146
8.2.1	Exponential map	146
8.2.2	Diffeomorphic weighting functions	147
8.2.3	Inversibility	148
8.3	Inverse Consistent Registration Model	149
8.4	Results	149
8.5	Discussion and Conclusion	150
9	Conclusions	153
9.1	Summary and Main Contributions	153
9.1.1	Scale-space of cortical feature maps	154
9.1.2	Geodesic Active Fields	154
9.1.3	Fast Scheme	155
9.1.4	Registration on the Sphere	155

9.1.5	Automatic cerebro-cortical parcellation	156
9.1.6	Diffeomorphic Deformations	156
9.2	Future Work	157
9.2.1	Short term	157
9.2.2	And beyond	158
A	Some Spherical Tools	161
A.1	Spherical Mesh Generation	161
A.1.1	Regular polyhedra	162
A.1.2	Subdivision	162
A.2	Mesh smoothing	163
A.3	Gradient estimation	165
A.3.1	Definitions	165
A.3.2	Gradient estimation on triangulated maps	166
A.4	Spherical point location and data interpolation	167
A.4.1	Triangle containment	167
A.4.2	Classical point location schemes	168
A.4.3	Spherical triangle walk	169
A.4.4	Vertex binning	170
A.4.5	Octree binning	172
A.4.6	Interpolation	172
	Bibliography	175
	Curriculum Vitae	195
	List of Publications	197

List of Figures

1 Introduction

1.1	Image registration is the concept of mapping homologous points in different images	2
1.2	The human cerebral cortex: location and structure	3
1.3	Human cerebral cortex: nomenclature and areas of localization . .	4
1.4	The aperture problem	12
1.5	Gestalt principles	16
1.6	Gestalt laws	17
1.7	Motion perception experiment according to Newsome	18
1.8	Motion perception self-experiment	19

2 The Beltrami Framework

2.1	The Swiss view of the Beltrami framework	26
2.2	Beltrami framework in a nutshell	27

3 Cortical Scale Space

3.1	Creation of spherical attribute image	50
3.2	Histogram equalization of cortical mean curvature maps	59
3.3	Scale space of cortical feature maps	60
3.4	Anisotropy parameter β^2 and label coherence	60

4 Geodesic Active Fields

4.1	Omnidirectional bunny image obtained from a catadioptric system	67
4.2	Linear and Beltrami scale space	72
4.3	Regularization for different embedding scales β^2	78
4.4	Stereo vision depth recovery	79
4.5	2D registration of brain monomodal MRI slices	80
4.6	Multimodal 2D registration of brain MRI slices	82
4.7	Diffusion on non-flat image domain	83
4.8	Registration of a spherical patch of the earth's topographical map	83

4.9	Multiscale registration of a brain MRI slice	84
5 A Fast Scheme		
5.1	Depth recovery from stereo: FastGAF versus Demons	98
5.2	2D registration of brain MRI slice: FastGAF versus Demons . . .	99
5.3	Spherical registration through stereographic projections	100
5.4	Performance comparison of registration schemes for stereo depth recovery	101
5.5	Performance comparison of registration schemes for 2D image reg- istration	101
5.6	Influence of penalty term and Lagrangian multiplier on registration performance	104
5.7	Computational load balancing L	104
5.8	Number of iterations J of the Jacobi approximate inversion scheme	105
6 Fast Scheme on the Sphere		
6.1	Nomenclature and deformation on the sphere	108
6.2	Neighborhood and Voronoi region on triangular meshes	112
6.3	Deformation field regularization on the sphere	117
6.4	Registration of a synthetic image pair on the sphere	118
7 Mapping the Human Brain		
7.1	Manual parcellation scheme in the Buckner cerebral cortex data set	124
7.2	Spherical multiresolution scheme for faster registration	127
7.3	Multiple scales of cortical feature used as geometry descriptor . .	129
7.4	DSC of pairwise labelmap agreement, with and without registration	132
7.5	Visual evaluation of different label fusion schemes	135
7.6	Effect of spatial smoothing for varying λ	136
7.7	Manual and automatic parcellation for an average subject	137
7.8	Manual vs. automatic parcellation for best and worst subject . .	138
7.9	Dice similarity coefficients for individual structures, with and with- out registration	139
8 Diffeomorphic Deformations		
8.1	Convergence of exponential maps with increasing scaling and squar- ing depths N	147
8.2	Warping a mesh with a vector field and its exponential	148
8.3	Diffeomorphic registration of a star to a disc	151
8.4	Additive and diffeomorphic deformation on the sphere	152
A Some Spherical Tools		
A.1	Polyhedra subdivision for regular point distribution on the sphere	164

A.2	Detail of the original and smoothed mesh, using implicit Laplacian smoothing	165
A.3	Gradient estimation on the cortical map	167
A.4	Triangle walk for spherical point location	171
A.5	Different aspects of octree vertex binning	173

List of Tables

1 Introduction

1.1	Structures of the central nervous system	3
1.2	Inverse problems	10

7 Mapping the Human Brain

7.1	List of parcellated structures	125
7.2	Pairwise labelmap agreement prior to and after registration	131
7.3	Parcellation accuracy for different atlas fusion schemes	137

List of Symbols and Notations

Basic algebra and analysis

$A, [a_{i,j}]$	$\mathbb{R}^{n \times m}$	matrix with entries $a_{i,j}$
A^T	$\mathbb{R}^{m \times n}$	transpose of A
A^{-1}	$\mathbb{R}^{m \times m}$	inverse of A
$h(x)$	$\mathbb{R} \mapsto \mathbb{R}$	scalar 1D function
$\mathbf{h}(\mathbf{x})$	$\mathbb{R}^n \mapsto \mathbb{R}^m$	vector-valued, n -D function
$\mathbf{h}^S(\mathbf{x})$	$\mathbb{R}^n \mapsto \mathbb{R}^m$	reflection: $\mathbf{h}^S(\mathbf{x}) = \mathbf{h}(-\mathbf{x})$
$\hat{h}(\omega)$	$\mathbb{R}^n \mapsto \mathbb{C}$	Fourier transform of $h(\mathbf{x})$
a^*	\mathbb{C}	Complex conjugate of a : $a^* = \Re(a) - i\Im(a)$
$\Delta h(\mathbf{x})$	\mathbb{R}	Laplacian of h : $\Delta h(\mathbf{x}) = \sum_{i=1}^n \frac{\partial^2 h}{\partial x_i^2}(x)$
h_{xx}	$\mathbb{R}^2 \mapsto \mathbb{R}$	second order partial derivative of h : $h_{xx} = \frac{\partial^2 h}{\partial x^2}(x)$
h_{xy}	$\mathbb{R}^2 \mapsto \mathbb{R}$	mixed second order partial derivative of h : $h_{xy} = \frac{\partial^2 h}{\partial x \partial y}(x)$
h_{yy}	$\mathbb{R}^2 \mapsto \mathbb{R}$	second order partial derivative of h : $h_{yy} = \frac{\partial^2 h}{\partial y^2}(x)$
$(f * g)(x)$	$\mathbb{R} \mapsto \mathbb{R}$	convolution: $(f * g)(x) = \int f(t)g(x-t)dt$
$\ f\ _p$	\mathbb{R}_+	L^p -norm: $\ f\ _p = \left(\int f(x) ^p dx\right)^{1/p}$
$\ f\ _1$	\mathbb{R}_+	L^1 -norm: $\ f\ _1 = \int f(x) dx$
$\ f\ _2^2$	\mathbb{R}_+	squared L^2 -norm: $\ f\ _2^2 = \int f(x) ^2 dx$

(Differential) geometry

\vec{b}	\mathbb{R}^n	(column) vector in n -D, mostly 2D or 3D
$\vec{e}_{1 2 3}$	\mathbb{R}^3	three basis vectors in 3D
E	$SO(3)$	local basis in 3D, $E = [\vec{e}_1, \vec{e}_2, \vec{e}_3]$
$\vec{r}(u^1, u^2)$	\mathbb{R}^n	2D manifold in n -D
$\vec{r}_i(u^1, u^2)$	\mathbb{R}^n	i -th component of 2D manifold in n -D
$\vec{r}_{,\alpha}$	\mathbb{R}^n	local tangent vector, $\vec{r}_{,\alpha} = \frac{\partial \vec{r}}{\partial u^\alpha}$ for $\alpha = 1, 2$
\vec{n}	\mathbb{R}^n	local unit normal to the manifold
κ_1, κ_2	\mathbb{R}	local principal curvature of manifold
\mathcal{S}		shape operator (second order tensor)
H	\mathbb{R}	local mean curvature
\vec{H}	\mathbb{R}^n	mean curvature vector: $\vec{H} = H\vec{n}$
I, II	\mathbb{R}	first and second fundamental form

Images		
$\mathcal{I}(\vec{x})$	\mathbb{R}	image at \vec{x}
$\mathcal{F}(\vec{x})$	\mathbb{R}	fixed image at \vec{x}
$\mathcal{M}(\vec{x})$	\mathbb{R}	moving image at \vec{x}
$\vec{\nabla}\mathcal{M}(\vec{x})$	\mathbb{R}^n	moving image gradient
$f(\vec{x}, \vec{x}')$	\mathbb{R}_+	local image mismatch between $\mathcal{F}(\vec{x})$ and $\mathcal{M}(\vec{x}')$
Beltrami embedding		
X	\mathbb{R}^m	e.g. $(x^1, \dots, x^2) \rightarrow (x^1, \dots, x^2, u^1, \dots, u^p)$
X_i	\mathbb{R}^m	local embedding around \vec{x}_i
$[h_{ij}]$	$\mathbb{R}^{m \times m}$	metric tensor, embedding space \mathbb{R}^m
$[g_{\mu\nu}]$	$\mathbb{R}^{n \times n}$	metric tensor, image coordinates \mathbb{R}^n
g	\mathbb{R}_+	determinant of metric tensor $[g_{\mu\nu}]$
β	\mathbb{R}_+	aspect ratio
S	\mathbb{R}_+	Polyakov energy
S_f	\mathbb{R}_+	weighted Polyakov energy
Numerical scheme		
$F(\mathbf{u})$	\mathbb{R}^N	disparity term: $F_i = f(\vec{x}_i, \vec{x}'_i(\mathbf{u}))$
$G(\mathbf{u})$	\mathbb{R}^N	regularity term: $G_i = \sqrt{g(\vec{x}_i, \vec{u}_i)}$
$\mathbb{F}(\mathbf{u})$	\mathbb{R}^{2N}	forked disparity term: $\mathbb{F}_{2i 2i+1} = F_i$
$\mathbb{F}'(\mathbf{u})$	\mathbb{R}^{2N}	disparity flow: $\mathbb{F}'_{2i 2i+1} = \frac{\partial f(\vec{x}_i, \vec{x}'_i(\mathbf{u}))}{\partial u_{2i 2i+1}}$
$\mathbb{G}(\mathbf{u})$	\mathbb{R}^{2N}	forked regularity term: $\mathbb{G}_{2i 2i+1} = G_i$
λ	\mathbb{R}^{2N}	Lagrangian multiplier
r	\mathbb{R}_+	penalty weight
W	$\mathbb{R}^{2N \times 2N}$	quasi-Laplacian matrix
Spherical mesh specific		
\vec{x}_i	S^2	coordinates of node i , $i = [0, N-1]$
\vec{x}_j	S^2	one-ring neighbor of \vec{x}_i : $\vec{x}_j \in N_1(\vec{x}_i)$
α_{ij}, β_{ij}	\mathbb{R}	patch angles, cf. figure
\vec{n}_i	S^2	unit-norm, outward normal at \vec{x}_i
${}_i\vec{e}_{1 2}$	S^2	local tangent basis vectors attached to i
E_i	$SO(3)$	local basis at i , $E_i = [{}_i\vec{e}_1, {}_i\vec{e}_2, \vec{n}_i]$
${}_i\vec{\theta}_j$	\mathbb{R}^3	relative position of neighbors in E_i : ${}_i\vec{\theta}_j = E_i^T(\vec{x}_j - \vec{x}_i)$
$\vec{\omega}_{ij}$	\mathbb{R}^3	axis of great circle rotation: $\vec{\omega}_{ij} = \frac{\vec{x}_j \wedge \vec{x}_i}{\ \vec{x}_j \wedge \vec{x}_i\ }$
Ω_{ij}	$SO(3)$	parallel transport around $\vec{\omega}_{ij}$
\vec{t}_j	\mathbb{R}^3	tangent vector at \vec{x}_j
${}_i\vec{t}_j$	\mathbb{R}^3	\vec{t}_j parallelly transported to \vec{x}_i : ${}_i\vec{t}_j = \Omega_{ij}\vec{t}_j$
t_j^1, t_j^2	\mathbb{R}	non-zero components of a tangent vector at \vec{x}_j in E_j
${}_i t_j^1, {}_i t_j^2$	\mathbb{R}	non-zero components of a tangent vector at \vec{x}_j transported to \vec{x}_i , in E_i
Λ_{ij}	$SO(2)$	mapping $({}_i t_j^1, {}_i t_j^2)^T = \Lambda_{ij} (t_j^1, t_j^2)^T$

∂V_{ij}	\mathbb{R}_+	Voronoi edge between \vec{x}_i and \vec{x}_j , and its length
V_{ij}	\mathbb{R}_+	Voronoi sub-patch spanned by \vec{x}_i and ∂V_{ij} , and its area
V_i	\mathbb{R}_+	Voronoi patch around \vec{x}_i : $V_i = \sum_{j \in N_1(i)} V_{ij}$
∂V_i	\mathbb{R}_+	boundary of V_i : $\partial V_i = \sum_{j \in N_1(i)} \partial V_{ij}$
\vec{u}_i	\mathbb{R}^3	tangent vector describing local displacement
\vec{x}'_i	\mathbb{R}^3	displaced node: $\vec{x}'_i = \vec{x}'_i(\vec{x}_i, \vec{u}_i)$
\mathbf{u}	\mathbb{R}^{2N}	non-zero components of N local displacement vectors: $\vec{u}_i = E_i(u_{2i}, u_{2i+1}, 0)$
S_i	\mathbb{R}_+	Polyakov energy approximation for Voronoi region V_i
g_{ij}	\mathbb{R}_+	metric tensor determinant representative for V_{ij}

Abbreviations

AL	Augmented Lagrangian
AOS	Additive Operator Splitting
CDF	Cumulative Density Function
CT	Computed Tomography
DTI	Diffusion Tensor Imaging
DSC	Dice Similarity Coefficient
DSI	Diffusion Spectrum Imaging
FastGAF	Fast scheme for GAF
GAC	Geodesic Active Contours
GAF	Geodesic Active Fields
HAC	Harmonic Active Contours
ICM	Iterated Conditional Modes
LOD	Locally One-Dimensional
LSF	Level Set Function
MAC	Multiscale Active Contours
MAP	Maximum <i>A Posteriori</i>
MPE	Minimal Polynomial Extrapolation
MRF	Markov Random Field
MRI	Magnetic Resonance Imaging
PDE	Partial Differential Equation
PDF	Probability Density Function
PSF	Point Spread Function
RRE	Reduced Rank Extrapolation
SDF	Signed Distance Function
TLA	Three Letter Abbreviation
TV	Total Variation

Introduction

1

“The whole is different from the sum of its parts.”

Aristotle, *Metaphysics*.

IMAGES AND IMAGING are ubiquitous and many related research questions have not been answered to full satisfaction, yet. In contrast to visual arts that may be the result of pure imagination, interpretation and perceptive distortion, medical images always are a depiction of a precise physical reality. Even though represented in most cases as a discrete array of quantified digital values, they reflect a real body part of given dimensions and in a continuous space.

As such, an image is just a *projection* of the infinitely dimensional reality (in terms of both space and features), and it is often useful to jointly work on different images at the same time to get a maximum of information. As an example, one can easily imagine to work on images of a same body part but coming from different subjects. It is obvious, that images will most probably not be identical and that information contained at given image or physical coordinates will not correspond. In such a case, it is intuitive to previously match homologous points of the two images. Even if images are issue of an identical physical reality they might still be different. They can have a different spatial resolution, another field of view, modified orientations or been acquired using different imaging techniques. The direct link between corresponding points in both images can be derived when the exact relation between image and physical space is known. As most of the time the latter information is not available, the spatial relation has to be established *a posteriori*, and by other




Figure 1.1: Image registration is the concept of mapping homologous points in different images. **(a)** The skull of a human is registered to chimpanzee and baboon by finding the deformation fields $\vec{u}_1(\vec{x})$ and $\vec{u}_2(\vec{x})$, such that human features, e.g., the mandible (red), at \vec{x} match those of **(b)** chimpanzee at $\vec{x} + \vec{u}_1(\vec{x})$ and **(c)** baboon at $\vec{x} + \vec{u}_2(\vec{x})$. Skull sketches reproduced from (Thompson, 1917).

means.

In all those cases, one is interested in the spatial transform that links the two images, as illustrated in figure 1.1. The process that establishes this spatial transform is called image registration.

The goal of this thesis is to develop a new solution for high-precision registration of spherical representations of the human cerebral cortex. In the process, we propose a novel geometric framework for image registration, based on the very versatile Beltrami framework, and called Geodesic Active Fields (GAF). In the following paragraphs we give a short introduction on the **biomedical context and motivation**, and provide a brief **review of image registration** in general and **brain registration** in particular. We then contextualize the image registration problem in the more general scope of **inverse problems encountered in computer vision**, and briefly discuss the related standard solution schemes, in particular with respect to regularization. We pay more detailed attention to the role of **regularization in the context of motion estimation**, and provide a small excursus on **psycho-physical Gestalt theory** in visual perception. Having these ideas in mind, we will then formulate the central idea behind the research presented in this thesis, i.e., the development of a parametrization-invariant **image registration framework with anisotropic Beltrami-regularization** of the estimated deformation field. Finally, the **contributions** of this thesis will be highlighted, and the **structure** of the following chapters will be outlined.

1.1 Biomedical Context and Motivation

The brain and the spinal cord form the central nervous system. The latter consists of several structures, as detailed in table 1.1 and illustrated in figure 1.2. The *telencephalon* (*cerebrum*, endbrain) is the topmost and biggest

Table 1.1: *Structures of the central nervous system, adapted from (Nieuwenhuys et al., 2008).* The endbrain and the optic nuclei form the forebrain. The diencephalon (interbrain) includes the thalamic structures, followed by the brainstem, consisting of the mid- and hindbrain. The latter includes the pons, the cerebellum and the medulla oblongata, which finally leads into the spinal cord.

Central nervous system

Encephalon (brain)

Prosencephalon (forebrain)

Truncus cerebri (brainstem)

Rhombenc. (hindbrain)

Metencephalon

<i>Telencephalon</i> (endbrain)	<i>Telencephalon</i> <i>impar</i>	<i>Diencephalon</i> (interbrain)	<i>Mesencephalon</i> (midbrain)	<i>Pons</i>	<i>Cerebellum</i> (little brain)	<i>Medulla oblongata</i>	<i>Medulla spinalis</i> (spinal cord)
------------------------------------	--------------------------------------	-------------------------------------	------------------------------------	-------------	-------------------------------------	--------------------------	--




Figure 1.2: *The human cerebral cortex: location and structure.* (a) Medial view and section of the human encephalon. The cortex is the topmost convoluted structure. (b) Coronal frontal section through the human brain. The outer surface of the cortical gray-matter shell is the pial surface (since it is wrapped in the *pia mater*), whereas its inner surface is referred to as white-matter-gray-matter interface. (Figures reproduced from (Gray, 1918).)




Figure 1.3: *Human cerebral cortex: nomenclature and areas of localization.* (a) “Areas of localization on lateral surface of hemisphere. Motor area in red. Area of general sensations in blue. Auditory area in green. Visual area in yellow. The psychic portions are in lighter tints.” (b) “Areas of localization on medial surface of hemisphere. Motor area in red. Area of general sensations in blue. Visual area in yellow. Olfactory area in purple. The psychic portions are in lighter tints.” (Figures and caption reproduced from (Gray, 1918).)

of these structures. The endbrain contains the cerebral cortex and subcortical structures such as the basal ganglia and the limbic system. Together with the hypothalamus, the endbrain is responsible for all voluntary control of the human body. The cortex is split into two hemispheres, linked by the corpus callosum, as well as the anterior and posterior/fornix commissure. It consists of white and gray matter. While white matter essentially contains the connecting fibers, the gray matter builds the outer shell of the cortex, contains most of the “computing neurons”, and is itself organized in different layers. The cortical surface, i.e., the gray matter sheet, is highly convoluted. The folding pattern consists of gyri (singular: gyrus, the visible ridges) and sulci (singular: sulcus, the buried valleys).

1.1.1 Functional mapping

Ever since the phrenologists’ era, elucidating the actual organization of the cortex, and in particular the functional regions of the gray matter, are an active area of research. Coarse sketches of functional localization are shown in figure 1.3. For example, the mapping of the sensori-motor regions on the pre- and post-central gyrus are today very well understood. Other regions are still more uncertain and current research relentlessly strives for even more precise localization of specific functional areas.

To this end, today, statistical group studies are a recognized methodology in brain imaging, in particular in functional imaging. Functional Magnetic Resonance (fMRI) is one of the most important tools in this domain, generating an impressive amount of research as attested by the hundreds of papers published every year. One of the key steps in the statistical analysis of fMRI

data is the registration of the images of the different subjects into a common reference, aiming at compensating the inter-individual anatomical variability. Standard methods in this domain largely involve global voxel-based registration, i.e. registration methods aiming at deforming the whole image volume to make it similar to a reference brain, e.g., (Lancaster et al., 1995; Ashburner and Friston, 1999; Zosso et al., 2006). Unfortunately, those volume registration techniques are not able to cope with the large complexity of the cortical surface, so that there is no guarantee that specific gyri or sulci will be properly registered by these methods. Though, such a very precise cortical surface registration is needed for precise brain functional mapping by fMRI group studies.

As an example, in a recent work, it has been shown that precise cortical registration can reveal very important activation zones that were never described by fMRI studies using global state-of-the-art registration methods. Indeed the presence of several non-primary auditory areas specialized for sound recognition and sound localization around the primary auditory cortex has been demonstrated in humans (Viceic, 2007; Viceic et al., 2006). The registration algorithm used in this work involved the manual placing of landmarks to guide the registration. Achieving the same level of registration precision but **fully automatically** will obviously open new perspectives in brain functional mapping. This is one of the motivations of this thesis.

1.1.2 Brain connectivity

Moreover, recently there has been an increasing interest for brain connectivity analysis thanks to the development of effective diffusion MRI sequences, such as Diffusion Tensor MRI (DTI) or Diffusion Spectrum MRI (DSI), and associated image analysis algorithms, namely tractography. These new tools allow inferring information about the brain architecture (at least in the white matter) (Hagmann et al., 2003, 2006b; Honey et al., 2009). The research in this domain is now moving from individual subject analysis to group studies (Hagmann et al., 2006a, 2008). In this context, a methodology was proposed to study the global brain circuitry by a labeled connectivity matrix (Hagmann et al., 2010). The entries of this matrix correspond to small regions of interest covering the whole cortical surface and its cells contain connectivity measures between the corresponding regions of interest, obtained with diffusion MR tractography. By statistical analysis of the connectivity matrix of patients and control subjects, this approach has an important potential in revealing for instance connectivity abnormalities in the early stages of some important pathologies such as schizophrenia (Cammoun et al., 2009; Lynall et al., 2010), multiple sclerosis, epilepsy, etc. But to be effective, this approach needs a precise matching between the regions of interest defined on the cortical surface of each subject in the study. Here again, a precise cortical surface registration is thus the key challenge, that we will try to tackle in this thesis.

To cut a long story short: “Registration algorithms also enable the pooling

and comparison of experimental findings across laboratories, the construction of population-based brain atlases, and the creation of systems to detect group patterns in structural and functional imaging data” (Toga and Thompson, 2001).

1.2 A Quick Review on Image Registration

Image registration is the concept of mapping homologous points of different images, representing a same object, see figure 1.1. Homology, in turn, is defined as the relation between “organs deriving from the same embryonic blanks”¹. In practice, however, it is highly difficult to establish homology in images strictly based on this definition, and a broad variety of registration methods exist, see (Brown, 1992; Maintz and Viergever, 1998; Audette et al., 2000; Zitová et al., 2003) and references therein.

1.2.1 Image dissimilarity metric

In expert-based registration, the necessary homology cues can be provided by the expert, but this approach becomes unfeasible if a higher number of images is to be registered. On the other hand, for automatic image registration, it is commonplace to substitute homology by a measurable criterion of image dissimilarity, which is to be minimized. Depending on the nature of the images to be registered, different metrics are used to assess image distances. If the images have been acquired using similar sensors, one can generally assume that the same entities are pictured at the same feature intensity in both images, subject to noise. An intuitive and simple choice for monomodal image registration subject to additive Gaussian noise is the squared error metric (sum of squared intensity differences) (Toga, 1999). In different settings, in particular in presence of impulse noise, the non-differentiable absolute error metric may be more suitable (Lucas and Kanade, 1981; Pock et al., 2007; Feigin and Sochen, 2009). In the multimodal case, other metrics, such as the Kullback-Leibler distance or Mutual Information are more appropriate (Wells III et al., 1996; Viola and Wells III, 1997; Maes et al., 1997; Rueckert et al., 1999; Thévenaz and Unser, 2000; Mattes, 2001; Pluim et al., 2003).

1.2.2 Deformation model and regularization

Registration methods also differ in the deformation model and constraints that are applied on the resulting deformation field. Simple parametric deformation models, including rigid and affine transformations, which are defined globally for the whole image space, restrict the degrees of freedom to a small number of

¹*Homologue: “sont homologues les organes dérivant des mêmes ébauches embryonnaires”.*
Le trésor de la langue française informatisé.

parameters. Semi-local models are parametrized by a coarse grid of deformation vectors, which are then interpolated using appropriate basis functions to obtain the dense deformation field, e.g. (Ashburner and Friston, 1999; Rueckert et al., 1999; Kybic and Unser, 2003; Rueckert et al., 2006). Registration becomes more flexible the finer those grid-points are seeded. Ultimately, each pixel/voxel is allowed to move individually, which is then called model-free or free-form deformation, as e.g., with the Demons framework (Thirion, 1998).

1.3 Registration of the Cerebral Cortex

Inter-subject anatomical registration is a difficult task due to the complexity and variability of brain structures, such as the brain cortex: sulci and gyri, vary a lot between subjects. There are, among all the non-rigid registration techniques, several matching criteria and several types of transformation, that have been applied in the context of brain registration (Lester and Arridge, 1999; Toga, 1999; Hellier et al., 2003; Ardekani et al., 2005; Gholipour et al., 2007).

1.3.1 Voxel-based versus landmark-based

Two main approaches can be distinguished for non-rigid registration: voxel-based and feature-based approaches. In voxel-based approaches the objective of the transformation is to optimize some global intensity-based correspondence measure like gray-level correlation or mutual information, see (Lancaster et al., 1995; Ashburner and Friston, 1999; Zosso et al., 2006) to name just a few. Several such registration methods have been compared in (West et al., 1997). The main limitation of these methods is usually that they lead to a compromise between the accuracy of the registration and the smoothness of the deformation. Such approaches work well for sub-cortical structure registration, such as (Christensen et al., 1997), but they are usually not appropriate for aligning the highly convoluted pattern of sulci and gyri (Collins et al., 1998; Hellier et al., 2003).

On the other hand, feature-based approaches rely on landmarks or segmentation: identifiable anatomical elements (point landmarks, lines or surfaces) have to be extracted in both reference and moving image and their correspondence determines the registration transformation. They use high-level anatomical information (sulcal lines, functional surfaces, important landmarks), which explicitly guarantees the physical validity required in the brain registration process. As an example of model based registration, point-based sulcus extraction allows refining cortical alignments with respect to volumetric methods (Chui et al., 1999). Vaillant and Davatzikos extracted representations of cortical sulci with active contours as landmarks for non-rigid brain image registration (Vaillant and Davatzikos, 1997). Thompson and Toga base their deformation field on embryologist-motivated brain structure surfaces (Thomp-

son and Toga, 1996). Miga et al. used iterated closest point registration (ICP) in conjunction with mutual information of colored surface point clouds to align MRI images to intra-operative laser range-scan images (Miga et al., 2003).

1.3.2 Surface-based

Several approaches have been developed that focus on the reconstructed cortical surface, as postulated by Van Essen and colleagues (Van Essen et al., 2000). It is, however, extremely challenging to directly align cortical surfaces due to the highly convoluted and irregular structure of reconstructed cortical sheet meshes. Projection of the cortical surface on a spherical manifold is therefore commonplace, preserving both topology and connectivity (Hebert et al., 1995). Mesh inflation allows coarse registration after partial cortical flattening (PCF) (Tosun et al., 2004), and gives rise to spherical feature maps, either after further inflation (Fischl et al., 1999a,b) or conformal mapping (Angenent et al., 1999). Landmark-based large deformation diffeomorphic registration on the sphere was then proposed in (Bakircioglu, 1999). Further, these spherical maps contain features of either the complete cortical surface or only one hemisphere, and lend themselves to registration by powerful intensity-based tools on a simplified shape. The features used include information about mesh geometry such as surface normals (Gu et al., 2004) or curvature (MacDonald et al., 2000), node convexity (Fischl et al., 1999b) or clamp histograms (Liu et al., 2004), geodesic distances to gyral crown vertices (Robbins et al., 2004) or stereotaxic coordinates (Toga, 1999), as well as anatomical features such as cortical thickness or sulcal depth. Flattening of local sections of spherical parametrizations into a 2D representation allows using classical 2D image registration methods in a first approximation, as done by several authors. However, this introduces major irregular dilations and contractions, for which partial solutions have been proposed (Thompson and Toga, 1999). The most recent achievement, based on the spherical feature maps provided by (Fischl et al., 1999a,b), is the *Spherical Demons* framework (Yeo et al., 2010), which extends the *Diffeomorphic Demons* (Thirion, 1998; Vercauteren et al., 2009) to the sphere.

1.3.3 What next?

In summary, it is today commonly accepted that purely volumetric voxel-based approaches are not suitable to address the problem of high-precision cortical surface registration. Whereas, in the present context, point-based registration methods are also unlikely to yield the necessary precision, surface-based techniques are today most promising (Van Essen, 1998). In this thesis, we will therefore focus our efforts on the registration of spherical feature maps, much along the lines of (Fischl et al., 1999a,b; Yeo et al., 2010).

1.4 Inverse Problems in Computer Vision

The discovery of X-rays, commonly attributed to Wilhelm Conrad Röntgen in 1895, marks the beginning of the extension of imaging beyond the boundaries of normal human visual perception. From then on, it became possible to “see” information from objects that was otherwise hidden and inaccessible. The spectrum of imaging modalities has not stopped broadening ever since, each modality picturing a different physical property of the underlying object.

Modern day imaging is mainly challenged by the *inverse* problem that consists of gaining insight on the physical reality of an object, given one or several derived images. Think of clinical healthcare without the assistance of non-invasive imaging techniques s.a. computed tomography or magnetic resonance imaging, or modern live science research without appropriate imaging instrumentation. Thanks to the exploding number of image acquisition devices and the availability of cheap (electronic) computing power, not only the process of imaging itself, but increasingly also the interpretation of acquired images is to a large extent left to non-human automata. The computational problems related to this task are generally referred to as image processing and computer vision. Typical tasks include, but are not limited to: Image restoration, segmentation, registration and classification, stereo and multiview scene reconstruction, and many others. The range of applications is enormous: medical imaging, biological imaging, non-destructive testing, remote sensing, surveillance and monitoring, robotics, consumer electronics, and many others.

1.4.1 Ill-posedness

Most of these problems, in particular the image registration problem, are **ill-posed** (Bertero et al., 1988), which means that not all of the following criteria apply (Hadamard, 1902; Morozov, 1975):

1. A solution exists
2. The solution is unique
3. The solution depends smoothly on the data

The classical resolution scheme is the following (Poggio et al., 1985):

- **Formulate the inverse problem in terms of the corresponding forward problem:** we look for an unknown signal u , that under the action of an operator Φ and possibly affected by some noise n yields the observed signal f .

$$\text{find } u \text{ s.t. } f = \Phi u + n \quad (1.1)$$

- **Rewrite as an energy minimization problem:** Typically choose $p = 2$ (least squares error), corresponding to a Gaussian additive noise

Table 1.2: Inverse problems. Some examples of inverse problems in image processing and computer vision. Note that in this list, the registration problem is the most complicated one, since its operator Φ (warping the moving image) is non-linear.

Task	u	Φ	f
Scale-space Produce coarse-scale image given a fine-grained version.	Coarse-scale image	0	n/a
Denoising Retrieve original image given a noisy acquisition.	Noise-free original	\mathcal{I}	Degraded image
Super-resolution Retrieve high-resolution image given low-resolution samples.	High-res image	Subsampling $\downarrow 2$	Low-res image
Inpainting Retrieve full image given a version with missing patches.	Full image	Sampling/masking	Holey image
Deconvolution Reconstruct original (sharp) image given a degraded version affected by spatial blur h and noise.	Sharp image	Blur H	Blurred image
Registration Find the deformation field matching homologous points in a pair of images	Deformation field	Warp image \mathcal{M}	Fixed image \mathcal{F}

model.

$$\min_u \{E = \|\Phi u - f\|_p^p\} \quad (1.2)$$

- **Add regularity penalty to render problem well-posed** (Morozov, 1975; Tichonov, 1963):. Here, $q = 2$ is generally preferred because of its convexity and differentiability.

$$\min_u \{E = \|\Phi u - f\|_p^p + \alpha \|\Gamma u\|_q^q\} \quad (1.3)$$

where typically in image processing, one chooses $\Gamma := D$, the first order differential operator, in order to impose smoothness of the solution.

An incomprehensive list of some inverse problems that fit into this scheme is given in table 1.2.

1.4.2 Wiener Filtering

The classical Wiener filtering result can be derived from (1.3) when $p = q = 2$. Take for example an image deconvolution problem: $\Phi = H$ represents some

spatial blur corresponding to the convolution with a kernel h . The minimization problem writes (Feigin and Sochen, 2009):

$$\min_u \|u * h - f\|_2^2 + \alpha \|\nabla u\|_2^2, \quad (1.4)$$

of which the Euler-Lagrange equations are (Feigin and Sochen, 2009; Bar et al., 2006)

$$(u * h - f) * h^S + \alpha \Delta u = 0, \quad (1.5)$$

where $h^S(\mathbf{x}) = h(-\mathbf{x})$ and Δ denotes the Laplacian. In Fourier domain, and after reordering, this yields the classical Wiener filter result (Gonzalez and Woods, 1992; Pratt, 2001):

$$\hat{u} = \frac{\hat{h}^*}{|\hat{h}|^2 + \alpha|\omega|^2} \hat{f}, \quad (1.6)$$

where \hat{u} is the Fourier transform of u , \hat{h}^* is the complex conjugate of \hat{h} , $|\hat{h}|^2 = \hat{h}^* \hat{h}$, and ω is the frequency vector.

1.4.3 Non-linear Regularization

Although widely used today thanks to its simplicity and optimality, this resolution scheme has some important shortcomings. Despite their popularity due to computational tractability, it is well-known, that these Gaussian regularizers are not feature-preserving and that the obtained solutions are typically overly smooth. Different anisotropic, feature-preserving regularization schemes have therefore been researched, e.g., (Nordström, 1990; Alvarez et al., 1992; You and Kaveh, 1999; Forsberg et al., 2010; Hong and Park, 2010). Beyond, the Rudin-Osher-Fatemi (ROF) or Total Variation (TV) model (Rudin et al., 1992), corresponding to $q = 1$ has particularly gained in importance. TV penalizes the norm of the gradient directly, instead of its square. This model addresses the fact that natural images are typically composed of different, homogeneous regions, separated by more or less steep transitions, i.e., the image gradient is sparse in the L^1 rather than the L^2 -sense. Most recently, in the context of compressed sensing, this model has been employed to recover signals sampled well beyond the Nyquist criterion, thanks to strong sparsity assumptions. The TV norm is not differentiable, however, and sophisticated methods need to be employed in order to optimize it efficiently.

1.5 Regularization in Motion Estimation

In image registration and motion estimation, the determination of the underlying deformation field between two images is an instance of an ill-posed inverse problem (Bertero et al., 1988), requiring additional prior knowledge to make it well-posed. Ergo, to restrict such deformation fields to what is believed to




Figure 1.4: *The aperture problem.* **(a)** The motion of a single contour is not sufficient to determine the global motion of the entire object. At three different positions, only portions of the motion can be detected: only components normal to the contour are perceivable. **(b)** In this example, three different motions produce the same physical stimulus (Movshon et al., 1985). Conversely, each stimulus spans a subspace of admissible underlying motion, and true motion is at the intersection of the subspaces of all edges of an object. This intersection is to be found by regularization.

be “physically meaningful” deformations, constraints on the field regularity are introduced. Typical regularization constraints reduce the variations of the deformation field by defining an additional penalty, based on some norm on the gradients of the deformation field.

1.5.1 Aperture Problem

In addition to the mathematical aspects of ill-posedness, motion regularization also addresses a more intrinsic shortcoming of the perception of motion, namely the well-known aperture problem (Stumpf, 1911; Marr, 1982; Adelson and Movshon, 1982; Hildreth, 1984; Hildreth and Koch, 1987), illustrated in figure 1.4. Indeed, due to the localization of the individual motion detector, motion can only be perceived at image contours, i.e., object boundaries. Even worse, only its components normal to that contour can be directly estimated. In contrast, in homogeneous regions or tangential to contours, motion cannot be perceived. Indeed, every local motion detector spans an entire subspace of underlying motions that are compatible with the perceived motion cue. This not only affects (first-order) detectors in computer vision, but has equally been shown experimentally for visual and tactile perception in living beings, e.g. (Gizzi et al., 1990; Pei et al., 2008). While in the living case, the individual motion cues are supposedly combined in subsequent neuron layers, in computer vision it is the role of regularization to intersect the different admissible motion subspaces, and thus to propagate motion information from object boundaries into their homogeneous inner regions, and to complete tangential components at boundaries from nearby data.

1.5.2 Variational Model

Typically, image distance metric and regularization penalty are incorporated into a single energy minimization model, a.k.a. variational model, e.g., (Hermosillo et al., 2002). The energy functionals are commonly of the form

$$E = E_{\text{data}} + r \cdot E_{\text{regularization}}, \quad r > 0, \quad (1.7)$$

which generalizes the Tikhonov-regularization approach (1.3) to arbitrary data and regularity measures.

In 1981, Horn and Schunck defined the term *optical flow* as “the distribution of apparent velocities of movement of brightness patterns in an image,” which “can arise from relative motion of objects and the viewer” (Horn and Schunck, 1981). In particular, one is interested in retrieving the underlying motion, given two subsequent image frames, \mathcal{I}_1 and \mathcal{I}_2 . This inverse problem is a special case of the concept of image registration, where we require the two images to be subsequent (monomodal) frames of the same scene.

Horn and Schunck introduced a fundamental method to determine optical flow, based on a first order brightness constancy assumption and a Gaussian smoothness constraint:

$$\min_{\vec{u}} \left\{ \int \left\langle \vec{\nabla} I_1, \vec{u} \right\rangle + (\mathcal{I}_2(\vec{x}) - \mathcal{I}_1(\vec{x} - \vec{u})) d\vec{x} + \alpha \int \left(\|\vec{\nabla} u_1\|^2 + \|\vec{\nabla} u_2\|^2 \right) d\vec{x} \right\}, \quad (1.8)$$

where $\vec{u}(\vec{x}) = (u_1(\vec{x}), u_2(\vec{x}))$ denote the two components of the flow-field. Because of its simple quadratic structure, this regularization penalty is differentiable and easy to optimize, which made it so popular.

1.5.3 Anisotropic Regularization

It was quickly realized, again, that this Gaussian regularization overly smoothens the flow field across object boundaries, such as occluding edges. Nagel and Enkelmann replaced the data-term by the squared intensity difference, thus accounting for higher-order terms, and simply wrote the same regularization term differently (Nagel and Enkelmann, 1986):

$$\min_{\vec{u}} \left\{ \int (\mathcal{I}_2(\vec{x}) - \mathcal{I}_1(\vec{x} - \vec{u}))^2 d\vec{x} + \alpha \int \text{tr}\{(\nabla \vec{u}^T)^T (\nabla \vec{u}^T)\} d\vec{x} \right\}. \quad (1.9)$$

Note, that this model now corresponds to the standard resolution scheme (1.3), and can also be found in the Demons framework (Thirion, 1998).

Driven by the realization that individual objects move independently and that regularization should be limited at their respective boundaries, Nagel and Enkelmann now tuned the regularization term by intercalating an edge-detecting matrix C :

$$\min_{\vec{u}} \left\{ \int (\mathcal{I}_2(\vec{x}) - \mathcal{I}_1(\vec{x} - \vec{u}))^2 d\vec{x} + \alpha \int \text{tr}\{(\nabla \vec{u}^T)^T C^{-1} (\nabla \vec{u}^T)\} d\vec{x} \right\}, \quad (1.10)$$

where the eigenvectors of C are oriented along and normal to the local image gradient, respectively, and the eigenvalues indicate gradient strength in that direction. Several similar results have been developed since, a summary and generalization of which can be found in (Weickert and Schnörr, 2001).

In a completely different, less fundamental and more ad-hoc approach, regularization was delimited by explicitly modeled boundaries (Pitiot and Guimond, 2008). Clearly, this second approach is unsuitable for a general formulation of an automatic registration framework.

1.6 Elements of Gestalt Theory and Motion Perception

Nature has a long-standing experience in vision and developed very sophisticated mechanisms for the efficient processing of visual and motion information. It might be helpful in the course of improving our current computational models for vision, to revisit this archetype and consider it as some sort of “natural state-of-the-art”. In this section, therefore, we want to shed some light on a few fundamental concepts derived from studies of psycho-physical vision and motion perception, without considering anatomical and implementational details, however. In particular, we briefly present selected elements of *Gestalt-Theorie* that deal with vision. For a more complete theory of vision, see for example (Metzger, 1975; Marr, 1982).

1.6.1 Gestalt Theory

The term *Gestalt* (german: shape, form, figure), is defined here as the “essence or shape of an entity’s complete form”. The Gestalt concept goes back to the work of von Ehrenfels (von Ehrenfels, 1890), in the late 19th century. Gestalt theory² was later extended, amongst others, by Max Wertheimer, Kurt Koffka and Wolfgang Köhler between 1912 and 1920 (Ash, 1985). One of its important elements is the holistic realization, that “a whole is not simply the sum of its parts, but a synergistic ‘whole effect,’ or [precisely] gestalt” (Behrens, 2004; von Ehrenfels, 1890). In particular, related to motion perception, the theory notes that “the effect of apparent movement is generated not so much by its individual elements as by their dynamic interrelation” (Behrens, 2004; Wertheimer, 1912). Consequently, an entity’s properties cannot be derived directly from the properties of its constituent parts, since the whole is “more than just the sum of its parts”. The last, Aristotelic theorem was termed *Übersummativität* (over-summativity) by von Ehrenfels (von Ehrenfels, 1890).

²While the Gestalt theory actually involves a much wider field of psychology and beyond, it is important to note, that here we only refer to its psycho-physical elements of visual perception.

1.6.2 Gestalt principles

The (visual) Gestalt theory involves a short list of fundamental Gestalt principles: emergence, reification, multistability and invariance. These four principles are illustrated in figure 1.5. Note that Gestalt theory only describes these principles, but does not provide a complete physical explanation for them.

Emergence is the key phenomenon attributed to over-summativity. Indeed, the terms are sometimes used as synonyms. Emergence describes the – sometimes sudden and surprising – perception of objects in an image, that could not be spotted by independent very local analysis alone. In other words, “Emergence [...] refers to the arising of novel and coherent structures, patterns, and properties during the process of self-organization in complex systems” – in this case the gain of information on objects from low-level visual patterns (Goldstein, 1999). At first sight, the image shown in figure 1.5(a) is composed of an arbitrary random pattern of black spots. Only after global study, the picture of a Dalmatian dog *emerges*. The dog cannot be spotted based on individual analysis of its spots, i.e., we cannot identify the dog’s ear, nose or legs prior to the visual emergence of the whole dog. A similar observation holds for the horses pictured in the Pintos painting, in figure 1.5(b).

Reification refers to the constructive perceptive phenomenon, where the mind introduces objects into a scene in order to provide a plausible explanation for what is actually perceived. This is illustrated in figure 1.5(c).

Multistability is the tendency to switch back and forth, in ambiguous visual situations, between different plausible interpretations. Probably the most famous example is the Necker cube, illustrated in figure 1.5(d).

Invariance, finally, is the capacity of the mind to recognize and discern objects independently of (simple) transformations such as rotations, scale and perspective, as shown in figure 1.5(e).

1.6.3 Prägnanz and the Gestalt Laws

Another fundamental ingredient of Gestalt theory is the concept of *Prägnanz* (german: pithiness). This concept describes the fact, that the human mind automatically perceives objects and patterns as grouped according to a few basic rules. These rules are commonly referred to as “Gestalt laws” or “Principles of grouping”. The five core laws are in order: proximity, similarity, closure, good continuation, and common fate. A few other, contemporary grouping laws – such as symmetry, convexity, connectedness, and good form – are less commonly agreed upon.

The five core Gestalt laws are illustrated in figure 1.6. The law of **proximity** obviously predicts, that items located close to each other will likely be considered as a whole and will be grouped together. In the shown example 1.6(b), the spacing between circles is slightly smaller vertically than horizontally, and consequently the dots clearly appear arranged in columns. If the circles are shaded as in 1.6(c), then the law of **similarity** can overrule this im-

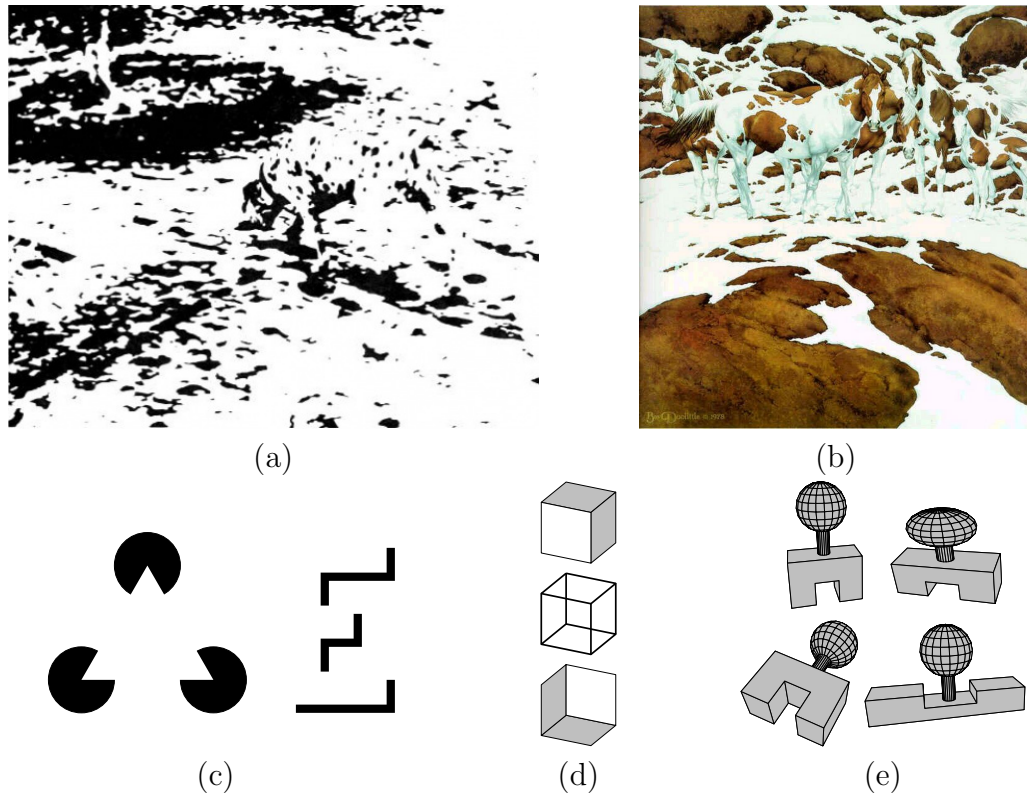


Figure 1.5: *Gestalt principles.* **(a)-(b)** Emergence. After some time, an intelligible object suddenly emerges out of the “unorganized, random” pattern. (a) Dalmatian dog. (b) Pintos, by Bev Doolittle, 1978. **(c)** Reification. The mind adds items to the scene in order to make the perceived visual information more plausible. Here: Kanizsa triangle and ‘E’. **(d)** Multistability. In ambiguous situations, one flips between different plausible interpretations of the visual experience. Here: the Necker cube. **(e)** Invariance. Objects are easily identified and discerned even under simple spatial transformations.

pression, and despite the column arrangement, the objects now appear grouped in rows. The law of **closure**, shown in 1.6(d), closes holes and bridges gaps between fragments of objects in order to build a whole, such as the circle and the well-known IBM-logo, also based on strong prior knowledge. The law of **good continuation** is the *Occam’s Razor* equivalent to Gestalt theory. It states, that we group and organize crossing line fragments in a way that lets each individual line be as smooth and regular as possible. Also, line fragments with similar properties (thickness, zigzag, straight or curved etc.) naturally group together. In our example in 1.6(e), one automatically perceives a straight diagonal crossing a curved line – while the lines could as well have 90-degree angles, by joining the two upper parts versus the two lower fragments, for example. The law of **common fate**, finally, lets us group items that are perceived to share the same fate, for example based on common movement, as sketched in




Figure 1.6: *Gestalt laws.* **(a)** Baseline: regularly spaced, identical circles evoke no particular grouping. **(b)** Law of proximity: due to reduced vertical spacing, the dots arrange in columns. **(c)** Law of similarity: horizontally common shadings groups the dots in rows. **(d)** Law of closure: the dots form a circle, and despite the horizontal interlacing the letters are clearly intelligible. **(e)** Law of good continuation: one tends to see a straight diagonal cutting a curved line. **(f)** Law of common fate: the two upward moving circles clearly separate from the interleaved downward moving ones.

1.6(f).

1.6.4 Motion coherence

A well-known study by Newsome and Paré in 1988 determined the lower limit of motion coherence required to evoke perception of global net motion (Newsome and Paré, 1988). In this study, monkeys were exposed to successions of random dot patterns on CRT screens, where each dot lives for $20 - 30 \mu s$, before it is replaced by a new dot, of which a controllable portion is replaced at a fixed displacement. This setup is illustrated in figure 1.7. If all successors are random, as in the ‘no correlation’ or 0%-coherent state, then no net motion can be perceived. The more dots, however, are replaced by correlated partners, the more the net motion becomes perceivable on a global level. It is to note that for low correlation levels, say around 10%, no long living streaks are apparent (only 1% of dots will have 2 correlated successors), so that motion can not be deduced from local cues alone. Monkeys had to decide whether the global net motion was upward or downward. The lower threshold of motion perception was around 1.8% motion coherence, below which the monkeys guessed entirely randomly, whereas motion detection performance was almost perfect around




Figure 1.7: *Motion perception experiment according to Newsome (Newsome and Paré, 1988).* **(a)** All dots are replaced by entirely randomly placed partners. The resulting apparent motion pattern involves no motion coherence. **(b)** 50% motion coherence. Half of the points are replaced by correlated partners, i.e., they apparently move coherently. **(c)** Full motion coherence: all dots are replaced by correlated successors. What is the minimal amount of coherence required for perception of a global motion pattern?

12% coherence (Newsome et al., 1989).

In order to better assess this amazing capability of the visual system to detect the relevant, coherent parts of motion at such high levels of noise, we revisited the Newsome study and performed a few tests in a self-experiment. We were able to reproduce parts of the study in humans and determined similar thresholds and characteristic curves. A few exemplary results are shown in figure 1.8.

1.7 A New Registration Framework

In this thesis we develop a new registration framework that tries to generalize existing approaches beyond current limitations and in particular addresses some shortcomings. Indeed, the additive variational framework, including Gaussian regularization in particular, comes with a few important drawbacks.

1.7.1 Gaussian Regularization Shortcomings

Firstly, it is well-known, that these Gaussian regularizers are not feature-preserving and that the obtained solutions are typically overly smooth. This is mainly due to the fact that with Gaussian regularizers, the only criterion defining the amount of information exchange, i.e., smoothing, is distance in image space. Indeed, one could say that Gaussian regularization ignores all Gestalt laws save the law of proximity.

It would be very reasonable, though, to restrain regularization to within distinct objects only. If objects are recognizable by boundaries in the reconstructed unknown data, then replacing the linear regularization by anisotropic regularization might help to do the job. Amongst other non-linear, anisotropic

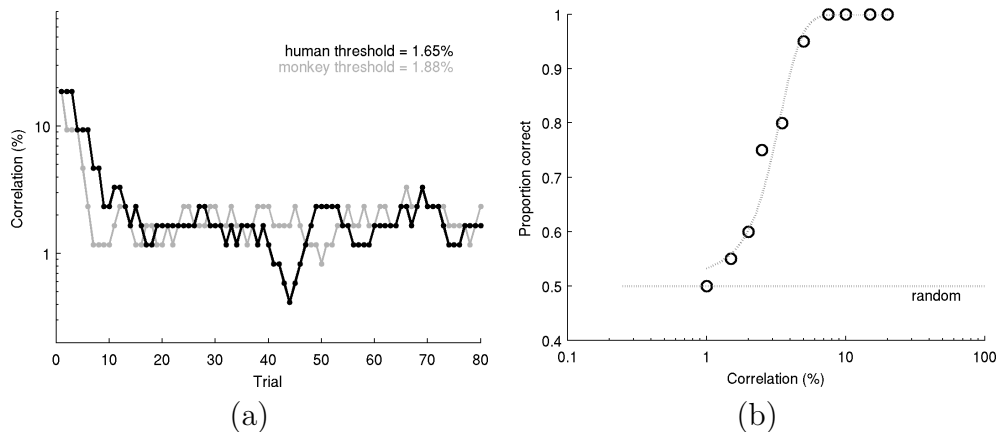


Figure 1.8: *Motion perception self-experiment.* **(a)** Determination of the human motion perception correlation threshold according to (Newsome and Paré, 1988): “The correlation value was well above threshold on the initial trial, and the probability of a decrease in correlation was 37% for each correct response. The correlation value increased following each incorrect response. [...] The threshold was calculated as the mean of the correlation values at each turning point (maxima and minima) during the last 60 trials.” The monkey-data of the original study is shown in gray. **(b)** Determination of the human motion perception correlation characteristic according to (Newsome et al., 1989). For different correlation values, direction detection was repeatedly tested ($n = 20$). Below around 1% motion coherence, the study subject guesses fully randomly, whereas above 10% correlation, the net motion is always perceived correctly.

DISCLAIMER: In this study, no PhD students were harmed and all experiments were conducted in accordance to all state and federal regulations. As liquid rewards were employed, coffee intake was controlled so as to maintain the student in a healthy but motivated state. Correct choices were rewarded with a drop of coffee.

diffusion schemes, the Rudin-Osher-Fatemi (ROF) or Total Variation (TV) model (Rudin et al., 1992), has therefore gained in importance.

Other approaches that are currently “en vogue” are bilateral filtering (Tomasi and Manduchi, 1998) and non-local regularization (Gilboa and Osher, 2008; Buades et al., 2008; Lou et al., 2009; Elmoataz et al., 2008; Peyré et al., 2008; Bougleux et al., 2011).

1.7.2 Drawbacks of Additive Regularization

Secondly, the additive energy model is not re-parametrization invariant. In image processing, the property of parametrization invariance is a very rare, but actually highly desirable property, so as in human vision. Indeed, there is no reason why the chosen parametrization of the image domain should influence the outcome of the process. And yet, many currently used image processing methods lack this important invariance property. While this is

absolutely acceptable in the current majority of cases, where images are classically parametrized by Cartesian coordinates, with the advent of wide-angle and catadioptric imaging, for example, non-Cartesian images are today gaining in importance. Such images are widely used in omnidirectional vision and robot navigation, for example, where ego-motion and position can be derived from a sequence of images, e.g. (Yagi et al., 1994; Gaspar et al., 2000; Bunschoten and Krose, 2003). The spherical feature maps of the cerebral cortex employed in this thesis obviously are non-Cartesian images as well.

1.7.3 Gestalt and the Beltrami Framework

In their seminal work (Sochen et al., 1998), Sochen, Kimmel and Malladi introduced the powerful Beltrami framework for image denoising and enhancement. This model is based on the Polyakov model (Polyakov, 1981) introduced in string theory for physics. The Polyakov model represents strings as harmonic maps in high-dimensional and curved spaces defined by Riemannian manifolds. Adopting this pure geometric point of view amounts to seeing objects such as images, shapes, or vector fields as geodesics or harmonic maps, much like a 2D topographic map corresponds to a three-dimensional surface in the real world.

The potential of this geometric framework lies in the general definition of the space-feature manifold and the choice of its metric. In particular, the metric can be chosen such that the Polyakov energy corresponds to an arbitrary interpolation between quadratic or total variation gradient penalty. The features are not restricted to scalar values but include vector features encountered in color, texture or multispectral image analysis (Kimmel et al., 2000). Similarly, the embedding is not limited to 2-dimensional image surfaces and generalizes naturally to n -dimensional manifolds associated to volumetric or time varying images or videos. Moreover, the choice of the metric enables the study of complex geometries inherent to scale-space methods (Bresson et al., 2006) and non-flat images generated e.g., by catadioptric or omnidirectional cameras (Bogdanova et al., 2007).

Depending on the exact embedding, the Beltrami framework can be implemented to incorporate several grouping principles from Gestalt theory. The laws of proximity and similarity are the most obvious ones. Embedding, e.g., the deformation field of the registration problem, allows mimicking the law of common fate. Closure and good continuation are closely related to TV-like regularization. It seems that the Beltrami framework is thus a very versatile and elegant way to reproduce in computer vision certain concepts known from human vision.

1.7.4 Weighted Beltrami

The second issue, absence of re-parametrization invariance, is more related to the additive structure of the variational model. Here, we note that the Beltrami framework can be equipped with a weighting function that allows

multiplicative coupling of the data term with the regularization term of the specific inverse problem. Indeed, the resulting energy functional has very interesting fundamental properties, such as geometric regularization interpolating between Gaussian and TV regularization, re-parametrization invariance, applicability to any Riemannian manifold and intrinsic automatic data-dependent modulation of the local regularization strength. The weighted Beltrami energy thus promises to be an interesting approach for the image registration problem at hand.

Moreover, research in related fields, such as compressed sensing (Donoho, 2006) and optimization theory (Glowinski and Le Tallec, 1989), has yielded very efficient optimization algorithms for the TV/L^1 case, which can partially be applied to the weighted Beltrami framework as well.

1.8 Organization of this thesis

We postulate that the weighted Beltrami framework represents an important step towards a unifying variational framework for geometric image processing, with a high degree of generality and a multitude of beneficial properties. Based on intuitions from Gestalt theory and given the particular role of regularization in motion estimation and image registration, we believe that this framework is an ideal fit for the cortical registration problem at hand. The central idea is (1) to use the Beltrami energy of the embedded unknown data as regularization term, and (2) to link data and regularization term through multiplication instead of addition.

1.8.1 Contributions

In this thesis, we develop a novel image registration framework, called Geodesic Active Fields, which is based on the weighted Beltrami energy, and we apply it to perform registration of cortical feature maps of the human brain.

The contributions of this thesis include:

- The formulation of an **appropriate anisotropic scale-space of cerebro-cortical feature maps** on the triangulated sphere, which allows later for hierarchical coarse-to-fine registration of the feature maps based on morphogenic concepts. The scale space is built upon the Beltrami framework.
- The formulation of **image registration as a weighted minimal surface problem**, within the weighted Beltrami framework. We provide the general setup and more detailed instantiations for different image geometries and modalities. The proposed Geodesic Active Fields (GAF) registration framework is able to deal with images defined in traditional Euclidean domains, but seamlessly extends to images in Riemannian spaces, such as curved manifolds or scale-spaces. The framework is

parametrization invariant, and the multiplicative coupling adaptively modulates the regularization strength based on the local alignment quality and noise level. We provide weighting functions for the common monomodal-Gaussian-noise case, as well as absolute error measure for monomodal-impulse-noise, and joint entropy for multimodal registration.

- **An efficient minimization scheme for the GAF framework**, called FastGAF, which allows optimizing the registration energy in time comparable to state-of-the-art registration schemes. We use splitting of the underlying deformation field into two constrained copies, and use an augmented Lagrangian approach to reformulate the constrained minimization as an unconstrained saddle point problem, efficiently solved using alternate minimization.
- The instantiation of the **FastGAF scheme on spherical meshes**, thus allowing for efficient registration of pairs of triangulated spherical features maps. This involves particular definitions, e.g., of the deformation field, its embedding and parallel transport, on the sphere.
- **A validation of the spherical registration scheme** in a leave-one-out cross-validation scheme, based on 39 cortical feature maps and their manual parcellations.
- A sketch of a possible inclusion of **diffeomorphic and inverse consistent deformation models**, with positive preliminary results.

1.8.2 Structure

The structure of this thesis is closely fit to the aforementioned contributions.

First, in chapter 2, we review the Beltrami framework in detail. We recall the general definitions, provide a few insightful examples, and give an extensive list of current applications. We complete with a quick overview of numerical schemes employed in Beltrami energy optimization so far.

For starters, chapter 3 then develops a scale space of cortical feature maps based on the Beltrami framework. This chapter also exposes the type of spherical maps and their structure, which are to be registered later on.

We introduce the Geodesic Active Fields (GAF) framework in chapter 4. This framework is a general, geometric framework for image registration, and here we present it in several show-case applications, e.g., for stereo vision, 2D Euclidean and multiscale registration, as well as registration on a spherical patch.

In order to overcome some numerical issues of the prototype implementation, and to make the optimization more efficient, a fast numerical scheme is presented in chapter 5. This scheme uses splitting and augmented Lagrangians so as to alternately minimize the data-term weighting function and the regularization term, respectively. We show, that the scheme compares favorably to state-of-the-art methods, both in terms of speed and registration quality.

In chapter 6 we formulate the spherical version of the Fast GAF scheme. Due to the specificities of the triangulated spherical surface, this involves a few particular (re-)definitions.

A validation study based on the pairwise registration of 39 subject brains is presented in chapter 7. We implement a Bayesian atlas fusion scheme which, from a set of pairwise registered parcellations, predicts the cortical parcellation of a subject hemisphere. This allows investigation of the registration performance in a leave-one-out cross-validation approach.

In the main body of the thesis, we only provide a very simple additive deformation model. Today, medical image registration tends to the use of diffeomorphisms, which restricts the space to physically more meaningful deformations. In chapter 8 we sketch, how a state-of-the-art diffeomorphic deformation model can be integrated into the GAF framework.

We conclude this thesis in chapter 9. We also provide potential follow-up work to further improve the proposed cortical registration scheme presented in this thesis. Moreover, it is shown how we believe that the insights and achievements of this thesis can be applied to a broader set of inverse problems in computer vision and image processing.

Finally, a few tools that we implemented to work with spherical images are detailed in the appendix A.

The Beltrami Framework

2

一分耕耘，一分收穫

“If one does not plow, there will be no harvest.”

Chinese proverb.

THE BELTRAMI FRAMEWORK was introduced in (Sochen et al., 1998) for low level vision, in particular image denoising and enhancement. This model is based on the Polyakov model (Polyakov, 1981) introduced in string theory for physics. While the Polyakov model represents strings as harmonic maps in high-dimensional and curved spaces defined by Riemannian manifolds, the Beltrami framework does similarly with images. In its simplest setting, this framework is not more complicated than an exercise known to every hiker in the mountains: understand a flat image as a topographical map of an actual 3D terrain, such as illustrated in figure 2.1.

Indeed, the Beltrami framework associates the spatial coordinates along with the features by defining an embedding, say $X : (x, y) \mapsto (x, y, I)$ for a 2D gray-scale image, as illustrated in figure 2.2. One now considers the image to be a 2D manifold embedded in 3D space. The Beltrami energy then measures the area of this non-flat manifold, which is a measure of image regularity. Minimizing the Beltrami energy amounts to regularizing (smoothing) the image, but in a very particular way. Namely, the Beltrami flow is proportional to the projection of the embedded manifold’s mean curvature vector on the feature component. Therefore, the diffusion flow is significantly less important at feature edges, and the smoothing is actually feature preserving.




Figure 2.1: *The Swiss view of the Beltrami framework.* A flat map corresponds to a higher-dimensional reality. Here: the intensity information (false-color) of the rectangular topographic map of the *Matterhorn* region translates directly into a three-dimensional terrain model. Real-world distances can be measured on the 2D map, by using an appropriate metric pulled-back from real-space. This is an easy exercise for every (non-GPS-based) hiker, and the Beltrami framework proposes a similar approach for image processing. (geodata ©swisstopo)

The Beltrami framework has a number of other interesting properties, such as parametrization invariance and the ability to work on any Riemannian domain (such as non-flat or multiscale images). Also, more than just gray-scale intensity can be embedded, and the range of possible applications is long, but we shall come to this in a moment.

The rest of this thesis will largely rely on the theory and concepts presented in this chapter. In the next paragraphs, we will first recall a few basic definitions from differential geometry, as well as review some scale-space theory and history. We will then present the Beltrami framework in more detail and provide an overview of a variety of its current applications. We will end this chapter with a short summary of numerical approaches that are used to optimize the Beltrami energy.

2.1 Preliminaries (1) – Definitions from Differential Geometry

In differential geometry, the first and second fundamental forms are quadratic forms on the tangent plane of a smooth n -dimensional hypersurface embedded in a $(n + 1)$ -dimensional Riemannian manifold, and are usually denoted by I and II, respectively. They were originally introduced by Gauss, who dealt with 2D parametric surfaces embedded in \mathbb{R}^3 .

2.1.1 First fundamental form I

First, define the surface as $\vec{r} = \vec{r}(u^1, u^2) = (x(u^1, u^2), y(u^1, u^2), z(u^1, u^2))^T$ in \mathbb{R}^3 . Let $\vec{r}_{,\alpha} = \frac{\partial \vec{r}}{\partial u^\alpha}$, $\alpha = 1, 2$ be the intrinsic tangent vectors. The second degree




Figure 2.2: *Beltrami framework in a nutshell.* A gray-level image is embedded in 3D according to $X : (x, y) \rightarrow (x, y, I)$. **(a)** Pixels across contours (red) are considered more distant than neighbors in flat regions or along contours (blue). In 3D, the shaded patches have area \sqrt{g} -times bigger than what they cover in 2D. The Polyakov energy computes the total hypersurface of the embedded manifold, which is a measure of the image regularity. **(b)** The minimizing flow is oriented along the normal-vector of the manifold and proportional to the mean curvature. Since only the features are allowed to move, the mean curvature vector \vec{H} is projected onto the respective feature dimension \vec{e}_I . The resulting flow is therefore significantly less important at feature contours, and Beltrami diffusion becomes feature-preserving.

polynomial

$$\begin{aligned} I(du^1, du^2) &= E(du^1)^2 + 2F du^1 du^2 + G(du^2)^2 \\ &= [du^1 du^2] \begin{bmatrix} E & F \\ F & G \end{bmatrix} \begin{bmatrix} du^1 \\ du^2 \end{bmatrix}, \end{aligned} \quad (2.1)$$

with $E = \vec{r}_{,1} \cdot \vec{r}_{,1}$, $F = \vec{r}_{,1} \cdot \vec{r}_{,2}$, $G = \vec{r}_{,2} \cdot \vec{r}_{,2}$ is called the first fundamental form, and the matrix $\begin{bmatrix} E & F \\ F & G \end{bmatrix}$ is called the metric tensor of the parametrization.

2.1.2 Second fundamental form II

Now, without loss of generality, suppose the surface be the graph of a twice continuously differentiable function $z = f(x, y)$. Let further the plane $z = 0$ be tangent to the surface at the origin. As a consequence, f and its first order partial derivatives with respect to x and y vanish at the origin. The Maclaurin expansion of f starts with quadratic terms, i.e.:

$$z = L \frac{x^2}{2} + Mxy + N \frac{y^2}{2} + O^3. \quad (2.2)$$

Accordingly, the second fundamental form at the origin is written

$$\begin{aligned} \text{II}(dx, dy) &= Ldx^2 + 2Mdx dy + Ndy^2 \\ &= [dxdy] \begin{bmatrix} L & M \\ M & N \end{bmatrix} \begin{bmatrix} dx \\ dy \end{bmatrix}, \end{aligned} \quad (2.3)$$

with $L = f_{xx}$, $M = f_{xy}$, $N = f_{yy}$ the second partial derivatives of f at the origin.

More generally, let again $\vec{r} = \vec{r}(u^1, u^2)$ be a parametrization of a surface in \mathbb{R}^3 . The normalized cross product of the partial derivatives $\vec{r}_{,\alpha}$ defines a field of unit normal vectors \vec{n} . If the coefficients $b_{\alpha\beta}$ are given by the projections of the second partial derivatives $\vec{r}_{,\alpha\beta} = \frac{\partial^2 \vec{r}}{\partial u^\alpha \partial u^\beta}$ on to the normal line \vec{n} , i.e.

$$b_{\alpha\beta} = \vec{n} \cdot \vec{r}_{,\alpha\beta}, \quad (2.4)$$

then the second fundamental form is usually written as

$$\text{II}(du^1, du^2) = b_{\alpha\beta} du^\alpha du^\beta, \quad (2.5)$$

where Einstein summation convention has been used.

2.1.3 On curvature

The mean curvature H of a surface is an extrinsic measure of curvature, derived from differential geometry, that locally describes the curvature of an embedded surface in its ambient space.

Let p be a point on the surface S and consider all curves formed by the intersection of S and any normal plane passing through the point p . Every such curve has an associated curvature, and of all curves at least one is characterized maximal (κ_1), and one as minimal (κ_2). These two curvatures are known as the principal curvatures of S , and they are given by the eigenvalues of the shape operator \mathcal{S} :

$$-\mathcal{S} = \begin{bmatrix} L & M \\ M & N \end{bmatrix} \begin{bmatrix} E & F \\ F & G \end{bmatrix}^{-1}. \quad (2.6)$$

The associated eigenvectors of this shape operator are the corresponding principal directions. While the product of the two principal curvatures, i.e. the determinant of \mathcal{S} , is called the Gaussian curvature, the mean curvature is defined as the average of the principal curvatures:

$$H = \frac{1}{2}(\kappa_1 + \kappa_2). \quad (2.7)$$

The mean curvature vector at p is then defined as

$$\vec{H} = H\vec{n}, \quad (2.8)$$

where \vec{n} is the unit normal to the surface S in p .

2.1.4 Higher dimensions and co-dimensions

In higher dimensions $n > 2$, the second fundamental form is linked with the shape operator \mathcal{S} of a hypersurface:

$$\text{II}(\vec{v}, \vec{w}) = \mathcal{S}(\vec{v}) \cdot \vec{w} = -\nabla_{\vec{v}} \vec{n} \cdot \vec{w} = \vec{n} \cdot \nabla_{\vec{v}} \vec{w}, \quad (2.9)$$

where $\nabla_{\vec{v}} \vec{w}$ is the covariant derivative of the ambient manifold.

The last description of the second fundamental form also generalizes to any arbitrary co-dimension, which means that the dimension of the ambient manifold might exceed the embedded surface by more than one, e.g. for a 2D surface embedded in \mathbb{R}^4 . In this case, II is a quadratic form on the tangent space with values in the normal bundle, and it can be defined by

$$\text{II}(\vec{v}, \vec{w}) = (\nabla_{\vec{v}} \vec{w})^\perp, \quad (2.10)$$

where $(\nabla_{\vec{v}} \vec{w})^\perp$ denotes the orthogonal projection of the covariant derivative onto the normal bundle.

2.1.5 First variation of area

The *first variation of area* formula relates the mean curvature of a hypersurface to the rate of change of its area, as it evolves in the outward normal direction. If $\Sigma(t)$ is a smooth family of oriented hypersurfaces in M , such that the velocity of each point is given by the outward unit normal at that point, then the first variation of area is given by

$$\frac{d}{dt} dA = H \cdot dA, \quad (2.11)$$

where dA is the area element on $\Sigma(t)$. This equation also reflects the fact, that a surface with zero mean curvature everywhere has minimal area. These are called minimal surfaces, and mean curvature flow minimizes the surface area.

2.2 Preliminaries (2) – On Scale and Space

The way we perceive the world depends on the scale of the aperture we use to measure it (Morse, 1994). This principle is trivially known as “missing the forest for the trees”, and has also particular validity in image processing: images are naturally composed of objects which are meaningful only at a given scale of observation (Marr and Hildreth, 1980; Koenderink, 1984).

As a consequence of the above, the scale at which one measures a certain property becomes an additional dimension of the imaging space, giving rise to Witkin’s patented notion of a *scale-space* (Witkin, 1983). He introduced the concept of artificially generating larger (coarser) scales of an image through low-pass filtering.

2.2.1 Linear Scale Space

The initial requirements for a scale space, i.e. the scale space axioms, were restrictive. These fundamental axioms for a scale-space obtained by convolution $L(x, y, t) = L(x, y) * g(x, y, t)$ with the smoothing kernel $g(x, y, t)$ include:

- linearity,
- shift invariance,
- semi-group structure, $g(x, y, t_1) * g(x, y, t_2) = g(x, y, t_1 + t_2)$, and cascade smoothing, $L(x, y, t_2) = g(x, y, t_2 - t_1) * L(x, y, t_1)$,
- existence of an infinitesimal generator A , i.e. $\exists A : \partial_t L(x, y, t) = (AL)(x, y, t)$,
- non-creation and non-enhancement of local extrema, i.e. $\partial_t L(x, y, t) \leq 0$ at spatial maxima and $\partial_t L(x, y, t) \geq 0$ at spatial minima,
- rotational symmetry, i.e. $g(x, y, t) = h(x^2 + y^2, t)$ for some h ,
- scale invariance, i.e. $\hat{g}(\omega_x, \omega_y, t) = \hat{h}(\frac{\omega_x}{\varphi(t)}, \frac{\omega_y}{\varphi(t)})$, for some $\varphi(t)$,
- positivity, $g(x, y, t) \geq 0$,
- normalization, $\int g(x, y, t) dx dy = 1$.

It was shown that the unique solution to these requirements is the normalized Gaussian kernel (Babaud et al., 1986)

$$g(x, y, t) = \frac{1}{2\pi t} e^{-(x^2 + y^2)/(2t)}, \quad (2.12)$$

where the variance σ^2 of the Gaussian is given by the scale parameter t . This kernel also has the nice property of being separable, i.e. $g(x, y, t) = g(x, t) \cdot g(y, t)$.

2.2.2 Relation to regularization

Now, consider an image restoration problem in the form of a Tikhonov regularized inverse problem (1.3). The observed image f is a noise-corrupted copy ($\Phi = \mathcal{I}$) of the unknown underlying image u , which we want to recover:

$$f = u + n \quad (2.13)$$

Assuming the noise is Gaussian and zero-mean, one would usually solve this using a variational approach:

$$u^* = \operatorname{argmin}_u \|f - u\|_2^2 \quad (2.14)$$

Since this problem is ill-posed, Tikhonov-like regularization is employed in order to render it well-posed:

$$\min_u \{ \|u - f\|_2^2 + \alpha \|\Gamma u\|_2^2 \}, \quad (2.15)$$

where we choose $\Gamma := D$ the first order differential operator, in order to enforce smoothness of the solution.

The Euler-Lagrange equation of this problem reads

$$\alpha \Delta u = u - f, \quad (2.16)$$

and with $u(0) = f$ as initial condition, the minimizer satisfies

$$\frac{u(\alpha) - u(0)}{\alpha} = \Delta u(\alpha), \quad (2.17)$$

which equates to an Euler-implicit time step of the heat flow:

$$\frac{\partial u(t)}{\partial t} = \Delta u(t). \quad (2.18)$$

The Green's function of this heat-flow problem is, again, uniquely given by the Gaussian kernel $g(x, t)$:

$$u(t) = (g * f)(x, t) \quad (2.19)$$

Thus, regularization, diffusion and scale-spaces are intimately related concepts (Florack et al., 2004; Scherzer and Weickert, 2000). Interestingly, the solution to a very local (heat) diffusion process is given by a global convolution with a kernel of infinite support.

2.2.3 Anisotropy

It is well-known, that the linear scale space is not feature preserving, and that solutions obtained using Gaussian regularizers are overly smooth at object boundaries. Therefore, one may define a more general energy functional that can be minimized through a diffusion process. Let

$$E = \frac{1}{2} \int C(\|\nabla u\|^2) dx \quad (2.20)$$

be an energy. Then the corresponding gradient descent equation is described by a diffusion PDE of the form:

$$\frac{\partial u}{\partial t} = -\frac{\partial E}{\partial u} = \nabla \cdot (c(\|\nabla u\|^2) \nabla u) \quad (2.21)$$

where $c(\cdot) = C'(\cdot)$ models the varying heat conductivity. In the above case of Tikhonov regularization, one has $C(\|\nabla u\|^2) = \|\nabla u\|^2$ and therefore $c = 1$ yields linear, isotropic diffusion.

Perona and Malik used anisotropic diffusion to build a scale space for edge-detection (Perona and Malik, 1990). They considered the following anisotropic (actually just inhomogeneous) conductivities:

$$c_{pm1}(\|\nabla u\|^2) = e^{-\left(\frac{\|\nabla u\|}{K}\right)^2}, \text{ and} \quad (2.22)$$

$$c_{pm1}(\|\nabla u\|^2) = \frac{1}{1 + \left(\frac{\|\nabla u\|}{K}\right)^2}. \quad (2.23)$$

To render the edge detector $\|\nabla u\|^2$ more robust against noise, Perona and Malik had to pre-filter the image with a pre-determined, non-adaptive Gaussian smoothing kernel.

Catté et al. integrated this ad-hoc filtering into the non-linear diffusion framework in a more elegant and strict way (Catte et al., 1992). They evaluate the conductivity on a continuously smoothed version of the image. To obtain a robust edge detector for the conductivity function, at any time t the image u is filtered with the derivative of a Gaussian kernel G_t of growing scale:

$$\frac{\partial u}{\partial t} = \nabla \cdot (c(\|DG_t * u\|^2)\nabla u) \quad (2.24)$$

On another note, choosing $C(\|\nabla u\|^2) = 2\sqrt{\|\nabla u\|^2}$ corresponds to the total variation energy or ROF model (Rudin et al., 1992; Rudin and Osher, 1994):

$$E = \int \|\nabla u\| dx \quad (2.25)$$

and conductivity $c(\|\nabla u\|^2) = \frac{1}{\sqrt{\|\nabla u\|^2}}$:

$$\frac{\partial u}{\partial t} = -\frac{\partial E}{\partial u} = \nabla \cdot \frac{\nabla u}{\|\nabla u\|}. \quad (2.26)$$

Alvarez tuned the TV flow equation, in order to obtain yet another anisotropic diffusion scale-space (Alvarez et al., 1992):

$$\frac{\partial u}{\partial t} = g(\|G * \nabla u\|) \|\nabla u\| \nabla \cdot \frac{\nabla u}{\|\nabla u\|}, \quad (2.27)$$

where G is a smoothing kernel, e.g., a Gaussian, and $G * \nabla u$ is a smoothed local estimate of the image gradient in presence of noise; $g(\cdot)$ is a non-increasing real function that vanishes at infinity. The tuning of the TV-flow by an edge-detecting function limits smoothing across edges and acts feature-preserving.

In an attempt to unify several nonlinear scale-spaces in a common formalism, and basing them on the well-understood linear scale-space, Florack et al. introduced a more strict definition of nonlinear scale-spaces through particular types of metric transforms (Florack et al., 1995). Florack's transforms are required to be diffeomorphisms in the plane, i.e., the underlying metrics

may well depend on the actual image but are constrained to be flat. It is, however, not necessary to explicitly compute the actual transform corresponding to the chosen metric. Unfortunately, however, with the restriction to flat metrics, this framework is not able to reproduce all of the previous anisotropic diffusion scale-spaces, s.a. the Perona-Malik model.

Beyond, other “ad-hoc” examples of anisotropic smoothing and regularization exist that are edge preserving or even edge enhancing, see e.g. (Weickert, 1997, 1999).

2.3 Beltrami Embedding, Geodesic Distance and Edge Detection

Related to Florack’s approach for non-linear scale-spaces, Sochen, Kimmel and Malladi introduced the Beltrami framework (Kimmel et al., 1997; Sochen et al., 1998). In the Beltrami framework, now, transforms are explicitly non-flat and existing non-linear scale-spaces are shown to be reproducible by this approach.

Let us consider a function $\phi(\mathbf{x}) \in W^{1,2} : \mathbb{R}^n \mapsto \mathbb{R}^p$, where $W^{1,2}$ denotes a Sobolev space, i.e. the L^2 norm of both the function and its first order derivative are finite:

$$\|\phi(\mathbf{x})\|_2 + \|\nabla\phi(\mathbf{x})\|_2 < \infty. \quad (2.28)$$

In simple terms, the function ϕ is required to be finite and differentiable. Note that here, the spatial parameter of ϕ is written \mathbf{x} rather than \vec{x} , since it may as well refer to non-Euclidean coordinates.

In this general form, $\mathbf{x} \in \mathbb{R}^n$ denotes the coordinates in n -dimensional space, and $\phi(\mathbf{x}) \in \mathbb{R}^p$ are the associated p -dimensional features. We thus have a mapping from space to features:

$$\phi : (\text{space}) \mapsto (\text{features}). \quad (2.29)$$

2.3.1 Embedding

Now, the Beltrami embedding associates the spatial coordinates along with the features. It defines a mapping $X : \mathbb{R}^n \mapsto \mathbb{R}^{n+p}$ as follows:

$$X : (\text{space}) \mapsto (\text{space}, \text{features}). \quad (2.30)$$

This amounts to seeing the function ϕ as a non-flat surface embedded in a higher dimensional space, much like a topographic map has a three-dimensional counterpart in the real world.

More generally, we have an n -dimensional image manifold Σ with coordinates $\sigma^{1\dots n}$, with associated p -dimensional features $\phi^{1\dots p}(\sigma^{1\dots n})$. This manifold can be embedded in an m -dimensional manifold M with coordinates $X^{1\dots m}$, with $m = n + p$. The embedding map $X : \Sigma \mapsto M$ is given by m functions of n variables:

$$X : (\sigma^1, \dots, \sigma^n) \mapsto (X^1(\sigma^1, \dots, \sigma^n), \dots, X^m(\sigma^1, \dots, \sigma^n)), \quad (2.31)$$

where the Jacobian needs to be of rank n and the map must be an injection to be an embedding.

To take the geometry of the manifolds into account, Σ and M are equipped with the metric tensors $g_{\mu\nu}$ and h_{ij} , and we get the Riemannian manifolds $(\Sigma, g_{\mu\nu})$ and (M, h_{ij}) , respectively.

Therefore, on M the line element is defined as

$$ds^2 = h_{ij}dX^i dX^j \quad i, j \in \{1, \dots, m\}, \quad (2.32)$$

where Einstein summation convention is used, i.e. identical indices that appear one up and one down are summed over.

Similarly, the metric $g_{\mu\nu}$ on Σ measures locally the distances at a point as follows:

$$ds^2 = g_{\mu\nu}d\sigma^\mu d\sigma^\nu \quad \mu, \nu \in \{1, \dots, n\}. \quad (2.33)$$

2.3.2 Induced Metric and Edge Detection

Naturally, the metric $g_{\mu\nu}$ is chosen as the induced metric, obtained by the *pullback*-relation: $g_{\mu\nu} = h_{ij}\partial_\mu X^i \partial_\nu X^j$.

Given $g_{\mu\nu}$, the (hyper-)area element on Σ is computed as

$$dA = \sqrt{g} \prod_{\mu=1}^n d\sigma^\mu = \sqrt{g} d^n \sigma \quad (2.34)$$

where g is the determinant of the metric tensor $g_{\mu\nu}$. The determinant g is clearly an edge detector function with respect to the embedded features. As such, it has been applied in several image segmentation applications based on active contours, e.g. (Sagiv et al., 2006; Derraz et al., 2009; Houhou et al., 2009).

2.3.3 Polyakov Energy

To measure the “weight” of the embedding $X : \Sigma \mapsto M$, Sochen et al. use the Polyakov action (Polyakov, 1981; Kimmel et al., 1997; Sochen et al., 1998):

$$S[X^i, g_{\mu\nu}, h_{ij}] = \int \sqrt{g} g^{\mu\nu} \partial_\mu X^i \partial_\nu X^j h_{ij}(X) d^n \sigma, \quad (2.35)$$

where, again, g is the determinant of the image metric, and $g^{\mu\nu}$ is the inverse, such that $g^{\mu\nu} g_{\nu\gamma} = \delta_\gamma^\mu$ (δ_γ^μ is the Kronecker delta).

Through calculation of the Euler-Lagrange equations, Sochen et al. suggest the following minimizing flow with respect to the embedding X^i (gradient descent):

$$\partial_t X^i = -\frac{1}{2\sqrt{g}} h^{il} \frac{\delta S}{\delta X^l} = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu X^i) + \Gamma_{jk}^i \partial_\mu X^j \partial_\nu X^k g^{\mu\nu} = \Delta_g X^i = H^i, \quad (2.36)$$

where Δ_g corresponds to the Laplace-Beltrami operator, the natural generalization of the Laplacian to Riemannian manifolds, equal to the i -th component of the mean curvature vector \vec{H} as defined in (2.8). The prefactor $-\frac{1}{2\sqrt{g}}h^{il}$ was introduced in order to make the resulting flow geometric (parametrization invariant), without changing the solutions of the Euler-Lagrange-equation, though. Further, Γ_{jk}^i are the Levi-Civita connection coefficients, i.e. the Christoffel symbols, with respect to the metric h_{ij} :

$$\Gamma_{jk}^i = \frac{1}{2}h^{il}(\partial_j h_{kl} + \partial_k h_{jl} - \partial_l h_{jk}). \quad (2.37)$$

If the metric h_{ij} is constant over the whole space, these coefficients are zero. The corresponding gradient descent equation then shortens to:

$$\partial_t X^i = -\frac{1}{2\sqrt{g}}h^{il}\frac{\delta S}{\delta X^l} = \frac{1}{\sqrt{g}}\partial_\mu(\sqrt{g}g^{\mu\nu}\partial_\nu X^i). \quad (2.38)$$

If, more specifically, the induced metric is used for $g_{\mu\nu}$, then the Polyakov energy reduces to the Euler functional:

$$S[X^i, g_{\mu\nu}, h_{ij}] = \int \sqrt{g}d^n\sigma, \quad (2.39)$$

which simply measures the hyperarea of the embedded manifold M . In the context of image processing, this energy is also often referred to as the *Beltrami energy*. The particular choice of the induced metric is well-understood in the 2D case, where the induced metric optimizes the Polyakov energy. In higher dimensions, however, a similar optimality has not been shown, yet.

2.3.4 Area Minimization, Anisotropic Diffusion and Feature Preservation

The PDE describing the area minimizing flow in (2.38) differs from the linear heat-flow PDE (2.18) only in the generalization of the Laplacian from Δ to Δ_g . This difference, however, is essential. Since the Beltrami-flow describes the isotropic heat-flow diffusion *on the embedded space-feature manifold*, diffusion becomes anisotropic in the original domain. As a consequence, the Beltrami flow is feature-preserving. Smoothing within almost homogeneous regions (in terms of embedded features) is faster than diffusion across feature contours.

Intuitively, we want to show this property from two different perspectives.

First, the general heat equation reads

$$\frac{\partial u}{\partial t} = \frac{1}{\rho}\nabla \cdot (\mathbf{C}\nabla\phi), \quad (2.40)$$

where for Beltrami diffusion the density can be identified as $\rho = \sqrt{g}$ and the conductivity tensor $\mathbf{C} = \sqrt{g}g^{\mu\nu}$, respectively. We see that based on this

density and conductivity, the heat flow is favored in homogeneous regions or along contours of ϕ .

From another perspective, the geometric embedding introduces the feature dimensions and therefore points which are close neighbors in the original space, but having different features, are considered farther apart than identical neighbors. Since diffusion is calculated based on the resulting geodesic distance between such neighbors, again we realize that diffusion is favored between pixels that are close and similar.

2.3.5 Examples

1D scalar ($n = p = 1$)

First, we consider the simplest case of a scalar function $\phi(x) \in \mathbb{R}$. The geometric embedding can be defined as

$$\begin{cases} X : (x) \mapsto (x, \phi(x)) \\ \{h_{ij}\} = \text{diag}(1, \beta^2) = \begin{bmatrix} 1 & 0 \\ 0 & \beta^2 \end{bmatrix} \\ \{g_{\mu\nu}\} = [1 + \beta^2 \phi_x^2] \\ g = 1 + \beta^2 \phi_x^2 \end{cases}, \quad (2.41)$$

where β is the aspect ratio between space x and feature $\phi(x)$. Considering the standard Beltrami energy $S = \int \sqrt{g} \, dx$, we can compute the minimizing flow analytically (Sochen et al., 2001):

$$\Delta_g \phi = \frac{\phi_{xx}}{(1 + \beta^2 \phi_x^2)^2}. \quad (2.42)$$

For $\beta \rightarrow 0$ we find classical Gaussian diffusion, otherwise the diffusion is limited at strong signal gradients.

2D scalar ($n = 2, p = 1$)

Let us now consider a more interesting 2D gray-scale image $\phi(x, y) \in \mathbb{R}$:

$$\begin{cases} X : (x, y) \mapsto (x, y, \phi(x, y)) \\ \{h_{ij}\} = \text{diag}(1, 1, \beta^2) \\ \{g_{\mu\nu}\} = \begin{bmatrix} 1 + \beta^2 \phi_x^2 & \beta^2 \phi_x \phi_y \\ \beta^2 \phi_x \phi_y & 1 + \beta^2 \phi_y^2 \end{bmatrix} \\ g = 1 + \beta^2 \|\nabla \phi\|^2 \end{cases}. \quad (2.43)$$

Here again, the analytic minimizing flow can be calculated easily:

$$\Delta_g \phi = \frac{(1 + \beta^2 \phi_y^2) \phi_{xx} - 2\beta^2 \phi_x \phi_y \phi_{xy} + (1 + \beta^2 \phi_x^2) \phi_{yy}}{(1 + \beta^2 \phi_x^2 + \beta^2 \phi_y^2)^2}. \quad (2.44)$$

For $\beta \rightarrow 0$ this 2D Beltrami operator clearly converges to the standard Laplacian. Indeed, for small x we can linearize the square root as follows:

$$\sqrt{1+x} \approx 1 + \frac{x}{2} \quad (2.45)$$

and therefore:

$$\int \sqrt{1 + \beta^2 \|\nabla \phi\|^2} dx \approx \int 1 + \frac{\beta^2}{2} \|\nabla \phi\|^2 dx. \quad (2.46)$$

In contrast, for $\beta \gg 1$ we obtain:

$$\lim_{\beta \rightarrow \infty} \frac{1}{\beta} \int \sqrt{1 + \beta^2 \|\nabla \phi\|^2} dx = \int \|\nabla \phi\| dx, \quad (2.47)$$

i.e. the Beltrami energy approximates the total variation norm (Rudin et al., 1992; Rudin and Osher, 1994).

2D vectorial ($n = 2, p > 1$)

The Beltrami framework can not only cope with scalar features to be embedded along with the space coordinates. Let us consider the case of a multispectral feature $(\phi_1(x, y), \dots, \phi_p(x, y))$ to be embedded:

$$\begin{cases} X : (x, y) \mapsto (x, y, \phi^1(x, y), \dots, \phi^p(x, y)) \\ \{h_{ij}\} = \text{diag}(1, 1, \beta_1^2, \dots, \beta_p^2) \\ \{g_{\mu\nu}\} = \begin{bmatrix} 1 + \sum_{i=1}^p \beta_i^2 (\phi_x^i)^2 & \sum_{i=1}^p \beta_i^2 \phi_x^i \phi_y^i \\ \sum_{i=1}^p \beta_i^2 \phi_x^i \phi_y^i & 1 + \sum_{i=1}^p \beta_i^2 (\phi_y^i)^2 \end{bmatrix} \\ g = 1 + \sum_{i=1}^p \beta_i^2 \|\nabla \phi^i\|^2 + \frac{1}{2} \sum_{i,j=1}^p \beta_i^2 \beta_j^2 \|\nabla \phi^i \times \nabla \phi^j\|^2 \end{cases} \quad (2.48)$$

The metric determinant g now includes a term containing the magnitude of the cross product between the respective feature gradient vectors, which represents an effective coupling between different channels of the embedded image. In the context of color image denoising, where the multispectral features are given by the different color channels, e.g. $(x, y) \mapsto (x, y, I^r(x, y), I^g(x, y), I^b(x, y))$, it was shown that this gradient alignment coupling corresponds well to the fundamental Lambertian image model (Kimmel et al., 2000).

2.3.6 Weighted Beltrami Energy

Weighted minimal surface problems have been known in image processing at least since the development of Geodesic Active Contours (GAC) (Caselles et al., 1997). Faugeras formulated 3D surface reconstruction from multiple views as a weighted minimal surface problem as well (Faugeras and Keriven, 1998). Without actually mentioning the Beltrami framework, Kimmel adapted this problem into a weighted Beltrami energy model for shape from stereo and auto-stereograms (Kimmel, 2002).

Sagiv used a Beltrami energy weighted by inverse data-fidelity to smooth orientation data of Gabor image descriptors (Sagiv et al., 2001), and Sarti et al. minimize the area of an embedded subject surface weighted by an edge detector of an underlying image (Sarti et al., 2002). In (Bresson et al., 2006), the Euler-Lagrange equations for the weighted Beltrami energy of arbitrarily dimensional manifold and features have been derived:

$$S_f = \int f \sqrt{g} d^n \sigma, \quad (2.49)$$

where $f = f(X^i, g_{\mu\nu}, h_{ij})$. This energy now describes the weighted hyperarea of the embedded manifold. The corresponding gradient descent equation is

$$\partial_t X^i = f H^i + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu X^i - \frac{m \cdot n}{2} \partial_k f h^{ki}, \quad (2.50)$$

where H^i corresponds to the original Beltrami-flow given in (2.38).

2.3.7 Parametrization invariance

As pointed out in the introduction of this thesis, images in computer vision and biomedical imaging are always a projection of an underlying physical reality. Any such projection can be parametrized in many different ways, depending upon the needs and particularities of a specific imaging technique or application. However, most of the state-of-the-art methods in the domain are designed to work with Cartesian parametrizations, and many have difficulties in generalizing to different image geometries, accruing e.g., from omnidirectional or wide-angle cameras. The Beltrami framework, however, thanks to its geometric foundation, is intrinsically parametrization invariant, viz., the Beltrami energy of an embedded image does not change with the parametrization of the image. Here, we provide a quick demonstration from a calculus point of view.

Parametrization invariance of the original and the weighted Beltrami energy can best be illustrated by actually performing a change of variables. Therefore, let $\mathbf{x} = (x_1, \dots, x_n)$ be a first parametrization of the image domain. Without loss of generality, we may obtain a re-parametrization into a second set of parameters $\mathbf{y} = (y_1, \dots, y_n)$ using a one-to-one, continuously differentiable mapping:

$$\mathbf{y} = (y_1(x_1, \dots, x_n), \dots, y_n(x_1, \dots, x_n)) = \mathbf{y}(\mathbf{x}). \quad (2.51)$$

Now, both metric tensors ${}^x g_{\mu\nu}$ and ${}^y g_{\alpha\beta}$ of the two embeddings are obtained as the respective induced metrics, through the pullback relation, from the *same* common metric h_{ij} of *the same* embedded manifold X^i :

$$\begin{cases} {}^x g_{\mu\nu} = h_{ij} \partial_\mu X^i \partial_\nu X^j \\ {}^y g_{\alpha\beta} = h_{ij} \partial_\alpha X^i \partial_\beta X^j \end{cases} \quad (2.52)$$

It is known from calculus, that a change of variables introduces the determinant of the Jacobian $D\mathbf{y}(\mathbf{x})$ of the mapping:

$$\int f(\mathbf{y}) \sqrt{{}^y g(\mathbf{y})} d\mathbf{y} = \int f(\mathbf{x}) \sqrt{{}^y g(\mathbf{x})} |\det(D\mathbf{y}(\mathbf{x}))| d\mathbf{x}. \quad (2.53)$$

If the energy term is parametrization invariant, this last expression needs to be equal to the weighted Polyakov energy parametrized by $\mathbf{x} = (x_1, \dots, x_n)$, i.e., we need to show that, for any f , we have:

$$\int f(\mathbf{x}) \sqrt{{}^y g(\mathbf{x})} |\det(D\mathbf{y}(\mathbf{x}))| d\mathbf{x} = \int f(\mathbf{x}) \sqrt{{}^x g(\mathbf{x})} d\mathbf{x}. \quad (2.54)$$

Therefore, one has to show that

$$\sqrt{{}^y g(\mathbf{x})} |\det(D\mathbf{y}(\mathbf{x}))| = \sqrt{{}^x g(\mathbf{x})}. \quad (2.55)$$

Using the chain rule, we have $\frac{\partial X^i}{\partial x_\mu} = \frac{\partial X^i}{\partial y_\alpha} \cdot \frac{\partial y_\alpha}{\partial x_\mu}$, where $\frac{\partial y_\alpha}{\partial x_\mu}$ stands for the Jacobian of the mapping. Thus

$${}^x g_{\mu\nu} = h_{ij} \partial_\alpha X^i \partial_\beta X^j \cdot \frac{\partial y_\alpha}{\partial x_\mu} \cdot \frac{\partial y_\beta}{\partial x_\nu} = {}^y g_{\alpha\beta} \cdot \frac{\partial y_\alpha}{\partial x_\mu} \cdot \frac{\partial y_\beta}{\partial x_\nu}, \quad (2.56)$$

and the determinant ${}^x g$ is effectively given by

$${}^x g = {}^y g \cdot (\det(D\mathbf{y}(\mathbf{x})))^2. \quad (2.57)$$

This amounts to equality (2.55) and therefore we have shown that the (weighted) Beltrami energy is indeed re-parametrization invariant.

2.4 Some Applications

In the next few paragraphs, we recall a few elementary applications of the Beltrami framework in approximately increasing order of complexity. Running the Polyakov energy minimization without any other constraints produces a scale-space. In lieu of the Gaussian gradient penalty, the framework can as well be employed as regularizer, by adding the Polyakov energy as regularization penalty to the original energy of the inverse problem. Other models can be built by tuning the original Polyakov energy in a more unexpected way, some examples of which will be presented below as well.

Note, that the Beltrami embedding was initially introduced to deal with gray-level and color images, volumes and movies (Kimmel et al., 2000, 1997; Sochen et al., 1998). Beyond, it is straightforward to embed textural features e.g., Gabor wavelets, the local structure tensor or patches (Sagiv et al., 2001, 2006; Houhou et al., 2009). Moreover, the Beltrami framework even applies to images whose features (pixel values) are tensors, useful for example in the context of Diffusion tensor image (DTI) regularization (Gur et al., 2009).

2.4.1 Image Denoising, Wiener Filtering, and Blind Deconvolution

First, let us consider again the image denoising problem, as stated in (2.13), and use its variational form (2.14). Instead of the linear, Gaussian regularizer, add the Polyakov energy of the embedded image $u(x, y)$ for anisotropic regularization:

$$X : (x, y) \mapsto (x, y, u) \quad (2.58)$$

such that the restoration is formulated as the following minimization problem:

$$\min_u \left\{ \int (u - u_0)^2 dx dy + \alpha \int \sqrt{g} dx dy \right\} \quad (2.59)$$

This is the most basic application of Beltrami regularization.

If an image is not only degraded by noise, but also affected by some spatial filter (motion blur, out-of-focus blur, sampling psf, etc.), one is faced with a deconvolution problem, as introduced in (1.4).

Beltrami-regularized anisotropic Wiener filtering was shown to outperform the standard Wiener filtering approach, thanks to its feature-preserving properties (Feigin and Sochen, 2009):

$$\min_u \left\{ \int (u * h - u_0)^2 dx dy + \alpha \int \sqrt{g} dx dy \right\} \quad (2.60)$$

If, in addition, the blurring kernel is unknown, it has to be estimated as well, a problem which is called blind deconvolution. Kaftory et al. proposed to solve for both the underlying image and the blurring kernel using two Polyakov regularizers (Kaftory et al., 2005). The joint minimization problem

$$\min_{u, h} \left\{ \int (u * h - u_0)^2 dx dy + \alpha_u \int \sqrt{g_u} dx dy + \alpha_h \int \sqrt{g_h} dx dy \right\}, \quad (2.61)$$

where g_u and g_h refer to the metric determinants corresponding to the embeddings of the image $(x, y) \mapsto (x, y, u)$ and the kernel $(x, y) \mapsto (x, y, h)$, respectively, is solved using alternate minimization of the following two sub-problems:

$$\min_u \left\{ \int (u * h - u_0)^2 dx dy + \alpha_u \int \sqrt{g_u} dx dy \right\} \quad (2.62)$$

$$\min_h \left\{ \int (u * h - u_0)^2 dx dy + \alpha_h \int \sqrt{g_h} dx dy \right\}. \quad (2.63)$$

2.4.2 Image Segmentation

In 2002, Sarti, Malladi and Sethian introduced a “A Geometric Model for Boundary Completion”, called “Subjective Surfaces”, based on the Beltrami

embedding and the weighted Polyakov energy (Sarti et al., 2002). They use an edge-detector function $g(x, y)$

$$g(x, y) = \frac{1}{(|\nabla G_\sigma * I|/\beta)^2} \quad (2.64)$$

as weighting function in the Polyakov energy of an embedded (subjective) surface ϕ :

$$\begin{cases} X : (x, y) \mapsto (x, y, \phi(x, y)) \\ S = \int g(x, y) \sqrt{1 + \|\nabla \phi\|^2} dx dy \end{cases} \quad (2.65)$$

Given appropriate initialization, this model evolves towards a quasi-piecewise constant surface, representing the different *Gestalts* (subjective surfaces) in the underlying image.

A more general approach for image segmentation was presented in (Bresson et al., 2006), where the minimizing flow was computed for arbitrary weighting functions, and used for multiscale active contours. The formalism is essentially the same as in (Sarti et al., 2002), except that the embedded surface is now considered to be the level-set-function (LSF) representation (Osher and Sethian, 1988) of a gradient-based active contour (Kass et al., 1988; Caselles et al., 1997). The multiscale embedding reads:

$$X : (x_1, \dots, x_n, \sigma) \mapsto (x_1, \dots, x_n, \sigma, \phi) \quad (2.66)$$

and the Multi-scale Active Contour (MAC) energy is obtained as

$$E_{MAC} = \int f(x_1, \dots, x_n, \sigma, \phi) \sqrt{1 + |\nabla \phi|^2 + \rho^2 \phi_\sigma^2} \frac{1}{c^{n+1} \rho} d\sigma \prod_{i=1}^n dx_i, \quad (2.67)$$

where c and ρ refer to the conductivity and mass density of the diffusion model generating the employed scale-space (Bresson et al., 2006). The weighting function is again chosen as an edge-detector function. Here, the weighting function acts as “data-term”, while the Polyakov energy itself only establishes regularity, within and across scale. Periodic reinitialization (Adalsteinsson and Sethian, 1995) is used to maintain the Level-set-function (LSF) close to a signed distance function (SDF).

A completely different approach was proposed by Estellers et al., where segmentation is obtained without actual edge detector (Estellers et al., 2011). Instead, features f^1, \dots, f^k of the image to be segmented are embedded along with the segmenting LSF ϕ , and the segmentation is driven by the gradient coupling property observed with non-scalar embeddings. The embedding map is defined as

$$X : (x^1, \dots, x^n) \mapsto (x^1, \dots, x^n, f^1, \dots, f^k, \phi) \quad (2.68)$$

and the metric on M is considered $[h_{i,j}] = \text{diag}(1, \dots, 1, \alpha_1^2, \dots, \alpha_k^2, \beta^2)$. The induced metric on Σ is then

$$g_{\mu\nu} = \delta_{\mu\nu} + \sum_i \alpha_i^2 f_\mu^i f_\nu^i + \beta^2 \phi_\mu \phi_\nu, \quad (2.69)$$

where $\delta_{\mu\nu}$ is the Kronecker delta. The determinant of the metric tensor is then computed as $g = g_{11}g_{22} - (g_{12})^2$ and the squared hyper-surface element corresponds to

$$g = 1 + \sum_i \alpha_i^2 |\nabla f^i|^2 + \beta^2 |\nabla \phi|^2 + \frac{1}{2} \sum_{i \neq j} \alpha_i^2 \alpha_j^2 \|\nabla f^i \times \nabla f^j\|^2 + \sum_i \alpha_i^2 \beta^2 \|\nabla f^i \times \nabla \phi\|^2. \quad (2.70)$$

We clearly see that the last term $\|\nabla f^i \times \nabla \phi\|$ measures the coupling of the features f^1, \dots, f^k with the segmentation function ϕ . Minimizing the Polyakov energy results in an alignment of the level-set-function gradients with the feature-gradients. The trade-off between gradient fidelity and levelset regularity is governed by the ratio β^2 versus α_i^2 , as can be seen in (2.70). Note that the embedded features can both be local features (e.g. intensity, patch, textural descriptors), or region-based features. In particular, the well-known Chan-Vese model (Chan and Vese, 2001) has successfully been integrated (Estellers et al., 2011).

2.4.3 Motion Estimation Regularization

As an extension of the pioneering work of Nagel and Enkelmann (Nagel and Enkelmann, 1986) and others, who had realized that isotropic Gaussian motion regularization is suboptimal, Ben-Ari et al. proposed to use the Beltrami framework as geometric regularizer, capable of generalizing several precedent models, for the stereo-vision case (Ben-Ari and Sochen, 2004, 2008), and for optical flow estimation (Ben-Ari and Sochen, 2009).

Their framework would embed the deformation vectors along with the image intensity:

$$X : (x, y) \mapsto (x, y, I, u, v) \quad (2.71)$$

The metric on the embedding manifold is chosen as $[h_{ij}] = \text{diag}(1, 1, \beta^2, \epsilon^2, \epsilon^2)$, where β^2 and ϵ^2 allow scaling the intensity and deformation components independently from the spatial features. The pulled-back metric tensor determinant, g , can now be derived in the usual way, and the optical flow (or registration) problem is formulated using the Polyakov energy as additive regularization term:

$$\min_{u,v} \left\{ \int (I(x, y) - I(x + u, y + v))^2 dx dy + \alpha \int \sqrt{g} dx dy \right\} \quad (2.72)$$

Based on different choices for β^2 and ϵ^2 , this model allows reducing deformation field regularization across edges in the image and/or the deformation field. Geometric regularization in presence of independent objects moving in an image (Gestalts) is clearly an advantage, since there are cases where boundaries in images – in terms of intensities or even texture – are good predictors of deformation field boundaries (Nagel and Enkelmann, 1986). Note that the extension to image features other than simple gray-value, such as parametric

texture descriptors or non-parametric patches or histograms, is straightforward in this framework.

2.4.4 Registration with Topological Defects

There are cases in medical image registration, where the constant brightness assumption does not hold. This typically occurs when lesions or tumors are present in one of the images only. One way to handle this issue, is to explicitly model these topological changes e.g., by masking, inpainting (Sdika and Pelletier, 2009) or explicit modeling (Bach Cuadra et al., 2002). Another way is to explicitly allow for intensity changes by introducing an additional deformation vector, in the feature direction (Barber et al., 2007).

Inspired by the work of Vemuri (Vemuri et al., 2003), Wyatt proposed to embed an image to be registered as $X : (x, y) \mapsto (x, y, I)$, and to use the normal vector \vec{U} of the resulting manifold as directional cue for deformation field evolution, i.e. along both spatial and intensity axes (Wyatt and Laurienti, 2006):

$$\frac{\partial X}{\partial t} = \phi \vec{U}, \quad (2.73)$$

where the speed is governed by the current image mismatch between the target image I_T and the evolving image $I(t)$, $\phi = I_T - I(t)$. Substituting ϕ and the definition of the normal vector

$$\vec{U} = \frac{1}{\sqrt{1 + \beta^2 \|\nabla I\|^2}} \begin{bmatrix} -\beta^2 I_x \\ -\beta^2 I_y \\ 1 \end{bmatrix}, \quad (2.74)$$

into (2.73) yields the following deformation field evolution $d\vec{V}$:

$$d\vec{V} = -\frac{I_T - I(t)}{\sqrt{1 + \beta^2 \|\nabla I\|^2}} \begin{bmatrix} -\beta^2 I_x \\ -\beta^2 I_y \\ 1 \end{bmatrix} \quad (2.75)$$

Note that repeated Gaussian filtering is required to smooth and regularize the evolution of the deformation field.

Later, this framework was refined, replacing the Vemuri-like evolution equation with a weighted Beltrami energy minimization (Wyatt et al., 2009; Li et al., 2011). The embedding under consideration remains unchanged, $X : (x, y) \mapsto (x, y, I)$, but the intensity mismatch is now introduced as a weighting function $\phi = (I_T - I(t))^2$, and one minimizes the weighted Beltrami energy

$$S = \int \phi \sqrt{g} dx dy \quad (2.76)$$

In this framework, the intensity scale β^2 is related to the permissiveness with respect to intensity modulations, and can locally be adapted based on the likelihood of topological change. Again, Gaussian smoothing is tacitly employed for regularization.

2.4.5 Non-standard Image Geometries

We close the section on Beltrami applications with a short list of more uncommon image geometries. Indeed, the Riemannian manifold theory, upon which the Beltrami framework is built, particularly lends itself to image geometries other than standard Euclidean space. It is not surprising, therefore, that various non-Cartesian applications have been proposed.

First, the framework has been formulated to work on implicitly defined manifolds (Sochen et al., 2003b,a,c). Shortly later, examples of the Beltrami flow on triangulated surfaces were presented (Lopez-Perez et al., 2004).

Equally based on triangulated surfaces, Spira et al. introduced a short-time kernel for Beltrami diffusion (Spira et al., 2007), and compared it against bilateral filtering (Tomasi and Manduchi, 1998), see also (Barash, 2002).

A first extension to more general graphs has been presented by Harrison, in order to smooth the parameters of a field of Gaussian processes (Harrison et al., 2007, 2008).

2.5 Numerical Schemes

In this section, we provide a brief overview of different numerical schemes and techniques that have been employed in order to minimize the (weighted) Beltrami energy in different contexts. We first introduce the fundamental implementations as explicit, semi-implicit or fully implicit schemes, before spending some time on different techniques employed to speed up the computations.

2.5.1 Explicit, semi-implicit and implicit

The direct way of implementing the Beltrami flow $\partial_t X^i = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu X^i)$ is obtained by using an Euler explicit forward scheme:

$$X_{k+1}^i = X_k^i + \partial_t X_k^i \cdot \tau, \quad (2.77)$$

where τ is the time-step of the PDE integration. After discretization in time, one obviously needs to discretize in space, which is most simply achieved by employing central differences, twice, as e.g., in (Ben-Ari and Sochen, 2009).

One may then rewrite the flow as the product of a matrix $L_k = L(X_k^i)$, whose coefficients are precisely given by the central differences scheme, and the current feature vector X_k^i :

$$\partial_t X_k^i = L_k X_k^i, \quad (2.78)$$

and therefore the Euler explicit forward scheme becomes:

$$X_{k+1}^i = X_k^i + \tau L_k X_k^i = (I + \tau L_k) X_k^i. \quad (2.79)$$

This scheme, however, is heavily limited by the CFL condition (Courant et al., 1928), and the time step τ has to be chosen carefully. An upper bound that ensures stability of such a scheme is provided in (Dascal and Sochen, 2005).

In order to make the integration scheme more robust, it was proposed to use a semi-implicit approach instead, e.g., (Weickert et al., 1998). Such a scheme is not identical, but shares the same first-order Taylor expansion. The fundamental idea is to make the flow component semi-implicit:

$$X_{k+1}^i = X_k^i + \tau L_k X_{k+1}^i = (I - \tau L_k)^{-1} X_k^i. \quad (2.80)$$

Such a semi-implicit scheme for the Beltrami flow is unconditionally stable, and the choice of τ is only limited by the tradeoff between efficiency and accuracy (Weickert et al., 1998). In particular thanks to the very sparse and regular structure of L , this problem can be solved relatively efficiently, as will be shown in the next section.

It is to note, that the fully implicit scheme, where the matrix $L_{k+1} = L(X_{k+1}^i)$ is employed instead, becomes non-linear and difficult to solve.

2.5.2 Operator splitting techniques

Indeed, the matrix L is very sparse, but due to the multiple dimensions of the diffusion, it is not nicely banded. In the one-dimensional case, however, the diffusion matrix L is simply tri-diagonal and the system can be solved efficiently (Thomas algorithm). In practice, many authors therefore reduce the multidimensional diffusion problem into a series (LOD) or average (AOS) of several independent 1D diffusion processes, that can be solved more easily (Weickert et al., 1998; Malladi and Ravve, 2002; Barash et al., 2003; Dascal et al., 2009).

A “locally one-dimensional” (LOD) scheme (Janenko, 1971), approximates the right-hand-term matrix by a product over all such matrices corresponding to one-dimensional diffusion. The corresponding semi-implicit scheme for a m -dimensional Beltrami diffusion problem then looks like

$$X_{k+1}^i = \left\{ \prod_{d=1}^m (I - \tau L_k^d)^{-1} \right\} X_k^i, \quad (2.81)$$

where L_k^d is the 1D diffusion matrix along the d -th dimension. This problem can now be solved by *subsequently* applying the Thomas algorithm m -times. Note that the order of the different 1D problems influences the solution of the overall problem.

In a more parallel and independent approach, the additive operator splitting scheme (AOS) was proposed (Lu et al., 1992). There, the different 1D diffusion processes are solved independently, e.g., using the Thomas algorithm, and then averaged over:

$$X_{k+1}^i = \frac{1}{m} \sum_{d=1}^m (I - m\tau L_k^d)^{-1} X_k^i, \quad (2.82)$$

This semi-implicit AOS scheme is first order accurate in time. Note that compared to LOD, here the time step of each individual diffusion is m -times larger.

2.5.3 Vector extrapolation

Vector extrapolation is a method to accelerate convergence of vector sequences (Jbilou and Sadok, 2000). The basic idea is to guess the limit of a vector sequence by looking at a series of samples. In the context of Beltrami-filtering the vector sequence samples correspond to subsequent states in the variational iteration of the (explicit) scheme, intended to converge in the limit to the solution of the problem (Dascal et al., 2007; Rosman et al., 2009).

Formally, let $\mathbf{x}_n, \dots, \mathbf{x}_{n+k}$ be a given sequence of N -dimensional column vector samples. Now we define the first and second forward difference vectors

$$\begin{aligned}\mathbf{u}_j &= \mathbf{x}_{j+1} - \mathbf{x}_j \\ \mathbf{w}_j &= \mathbf{u}_{j+1} - \mathbf{u}_j\end{aligned}\tag{2.83}$$

and build the $N \times (j+1)$ matrices

$$\begin{aligned}\mathbf{U}_j^{(n)} &= [\mathbf{u}_n; \mathbf{u}_{n+1}; \dots; \mathbf{u}_{n+j}] \\ \mathbf{W}_j^{(n)} &= [\mathbf{w}_n; \mathbf{w}_{n+1}; \dots; \mathbf{w}_{n+j}]\end{aligned}\tag{2.84}$$

Mainly, two different schemes are used to derive an estimation of the limit of the sequence: Minimal polynomial extrapolation (MPE) and reduced rank extrapolation (RRE).

Minimal polynomial extrapolation

In MPE, the limit of the vector sequence is estimated as

$$\mathbf{s}_{n,k} = \sum_{j=0}^k \gamma_j \mathbf{x}_{n+j},\tag{2.85}$$

where the coefficients γ_j are estimated as follows:

1. Obtain a least square solution of

$$\mathbf{U}_{k-1}^{(n)} \mathbf{c} = -\mathbf{u}_{n+k}\tag{2.86}$$

e.g. using Gram-Schmidt.

2. Denote $(c_0, c_1, \dots, c_{k-1})^T := \mathbf{c}$ and set $c_k = 1$. Now compute

$$\gamma_j = \frac{c_j}{\sum_{i=0}^k c_i}, \quad 0 \leq j \leq k,\tag{2.87}$$

assuming that $\sum_{i=0}^k c_i \neq 0$, otherwise an estimate of the limit does not exist.

Reduced rank extrapolation

With RRE, the approximation $\mathbf{s}_{n,k}$ of the limit is estimated differently:

$$\mathbf{s}_{n,k} = \mathbf{x}_n + \sum_{j=0}^{k-1} \xi_j \mathbf{x}_{n+j}, \quad (2.88)$$

where, this time, the coefficients ξ_j are determined by solving the overdetermined linear system

$$\mathbf{W}_{k-1}^{(n)} \xi = -\mathbf{u}_n, \quad (2.89)$$

with $\xi = (\xi_0, \xi_1, \dots, \xi_{k-1})$.

Vector extrapolation for Beltrami

For the Beltrami filtering application, the vector sequence is generated by repeated application of the explicit scheme. After a few such iterations from an initial guess, an extrapolation step is interleaved. The estimated limit $\mathbf{s}_{n,k}$ then serves as starting point for the next round of explicit iterations, and so on until convergence. Such a scheme has been called cycling (Sidi, 1991).

2.5.4 Decoupling schemes

An entirely different approach is proposed in (Rosman et al., 2010), based on a decoupling-scheme known in computer vision for quite some time (Cohen, 1996). A similar scheme has successfully been employed for efficient minimization in TV-regularized image denoising (Tai and Wu, 2009).

Note that, in contrast to the aforementioned numerical improvements, here the optimization of the Polyakov energy is directly modified, instead of just a speed up of its minimizing PDE flow.

The deconvolution problem using Beltrami regularization is originally defined similarly to (2.60), except that one includes multiple color channels, I^R , I^G and I^B :

$$\min_{I^i} \left\{ \int \left(\frac{\alpha}{2} \sum_{R,G,B} (HI^i - I_0^i)^2 + \sqrt{1 + \beta^2 \sum_{R,G,B} \|\nabla I^i\|^2 + \frac{\beta^4}{2} \sum_{R,G,B} \sum_{j \neq i} \|\nabla I^i \times \nabla I^j\|^2} \right) \right\}, \quad (2.90)$$

where H is the blurring kernel and I_0^i refers to the observed, deteriorated image channels. Similar to (Tai and Wu, 2009), the gradient is decoupled by introducing a new variable \mathbf{q}^i . The problem is now rewritten as a constrained minimization:

$$\min_{I^i, \mathbf{q}^i} \left\{ \int \left(\frac{\alpha}{2} \sum_{R,G,B} (HI^i - I_0^i)^2 + \sqrt{1 + \beta^2 \sum_{R,G,B} \|\mathbf{q}^i\|^2 + \frac{\beta^4}{2} \sum_{R,G,B} \sum_{j \neq i} \|\mathbf{q}^i \times \mathbf{q}^j\|^2} \right) \right\} \quad (2.91)$$

$$s.t. \quad \mathbf{q}^i = \nabla I^i.$$

Then, the constrained problem is rephrased as the following augmented Lagrangian:

$$\mathcal{L}(I^i, \mathbf{q}^i, \boldsymbol{\lambda}^i) = \int \left(\begin{aligned} & \frac{\alpha}{2} \sum_{R,G,B} (HI^i - I_0^i)^2 + \\ & \sqrt{1 + \beta^2 \sum_{R,G,B} \|\mathbf{q}^i\|^2 + \frac{\beta^4}{2} \sum_{R,G,B} \sum_{j \neq i} \|\mathbf{q}^i \times \mathbf{q}^j\|^2} \\ & + \sum_{R,G,B} \langle \boldsymbol{\lambda}^i, \mathbf{q}^i - \nabla I^i \rangle + \frac{r}{2} \sum_{R,G,B} \|\mathbf{q}^i - \nabla I^i\|^2 \end{aligned} \right) \quad (2.92)$$

where $\boldsymbol{\lambda}^i$ is the Lagrangian multiplier. This problem can be optimized efficiently using alternate minimization of the data-term and regularity-term.

3

Cortical Scale space

玉不琢不成器

“Jade must be chiseled before it can be considered a gem.”

Chinese proverb.

MUCH LIKE A TREE, the human brain exhibits a highly convoluted and irregular structure, with lots of complexity and variability: sulci and gyri vary a lot between subjects. On the other hand, high level structures of the brain – the “big picture” – are highly conserved, such as the two hemispheres, the lobes and main folds. A hierarchical representation of these structures is important for example in the context of intersubject registration: Considering the complexity of the cortical surface, directly involving local small-scale features would mislead the registration to be trapped in local minima. A robust method needs to rely on large-scale features, describing the main landmarks of the cortex, such as the main gyri or sulci. Subsequently, features are to be iteratively refined to drive the registration more locally and reach the desired precision.

Further, it is commonplace to represent the cortical surface as a spherical map, obtained through inflation of the original surface (Hebert et al., 1995; Fischl et al., 1999a). Available features of the mesh geometry are e.g. surface normals (Gu et al., 2004) or curvature (MacDonald et al., 2000), geodesic distances to gyral crown vertices (Robbins et al., 2004) or stereotaxic coordinates (Toga, 1999). The concept of spherical attribute images is illustrated in Fig. 3.1 using the example of mean curvature as feature.




Figure 3.1: *Creation of spherical attribute image.* (a) The cortical surface of a hemisphere is extracted and attributes mapped to its vertices. (b) Partially unfolded cortical sheet, textured with mean curvature of the original sheet. (c) Complete inflation yields the spherical feature map. Map data was equalized for improved visualization. (d) Nomenclature of vertices, neighbors and angles in a one-ring patch of the spherical mesh.

Here, we are interested in defining a scale-space that lends itself for a meaningful hierarchical representation of structures on cortical feature maps on the sphere¹. The scale-space is expected to produce “generic” brain images at coarse scale, adding more and more individual details at finer scales.

3.1 Introduction

The convolutions of the cortical sheet are very complex but by no means entirely random. There seems to exist a common “base-model” for the human cortex, which is then modulated on an individual basis. Here, we have neither the intention nor the qualification to provide a complete description of human cerebro-cortical morphogenesis. Nevertheless, we want to highlight a few results from corresponding research, since they motivate the particular choice of scale-space proposed in this chapter.

3.1.1 From gyrification to scale-space

The process of increased folding and fissuring of the cerebral cortex during gestation and, for a certain time, post-partem, is referred to as *gyrification*. The degree of gyrification can be assessed by an index either globally (Zilles

¹Parts of this chapter have been published in (Zosso and Thiran, 2009).

et al., 1988), or locally (Schaer et al., 2008). The exact and complete causality of the cortical folding is still largely unclear, but a few elements are commonly agreed upon. One central hypothesis behind cortical folding pathways is the influence of fiber tension within the white matter (Van Essen, 1997). This has also been underlined in a more recent study (Toro and Burnod, 2005). In parallel, it was realized that sulci exhibit different degrees of inter-subject stability or variability. In a monozygotic twin study, it was shown that (1) inter-subject variability is generally lower among monozygotic twins compared to between unrelated subjects, and (2) that among monozygotic twins the variability was particularly low for deep sulci, which embryologically develop first (Lohmann et al., 1999). This immediately suggests the existence of some underlying, stable sulcal ground-pattern, which is genetically coded at least to some extent, while the smaller structures are rather the result of subsequent emergence. Corresponding results were shown in (Lohmann et al., 2008).

It is self-evident, that in a cortical registration framework, the feature-map scale-space should respect this hierarchical pattern of variability to a great extent. Indeed, it was hypothesized that a geodesic diffusion scale-space of mean curvature of the cortical surface amounts to reversing the cortical folding process Cachia et al. (2003). In that work, the authors relate coarse-scale features with elementary cortical folds, called the sulcal roots, that seem to be particularly stable across individuals. Solely designed to yield landmarks for a primal sketch of folding patterns, however, their scale-space does not use a global spatial parametrization of the cortical surface and works with local planar parametrizations instead.

In our context, however, we are interested in a scale-space defined on the spherical mesh of the inflated cortical feature maps. To this end, we want to combine the advantages of a mean curvature scale-space, supposed to reveal stable sulcal roots at larger scales, with the benefits of a simple spherical parametrization. Further, we propose to generalize the linear diffusion by using the Beltrami framework that allows an additional degree of freedom with respect to the diffusion properties, and that opens the possibility of including multi-spectral features.

3.1.2 Beltrami flow on triangulated manifolds

Manifolds and embeddings need not necessarily be analytically defined through parametrizations and metrics. Instead, the image manifold can be sampled as discrete values at triangulated vertices, see figure 3.1. The Beltrami embedding is then simply obtained by enhancing the coordinates of each vertex by the corresponding feature components, optionally scaled by a scaling factor β to obtain the desired aspect ratio. The topology of the triangulation, however, remains unchanged.

In this context, different schemes for triangulated manifolds have been presented: Bajaj et al. base their approach on Loop's subdivision scheme to smooth both the surface and functions on that surface (Bajaj and Xu, 2003).

Lopez et al. estimate the Laplace-Beltrami operator on a triangulated manifold by discretizing the surface divergence of the gradient field of the function on that surface Σ (Lopez-Perez et al., 2004). This approach has the advantage of keeping the geometry of the mesh fix (i.e. only the geometry of Σ matters), but requires re-evaluation of the function gradient over that geometry at each iteration.

In the case of a triangulated manifold, however, it is not necessary to calculate the Beltrami flow through evaluation of the intrinsic Beltrami operator formulation. Instead, we propose to take advantage of the equivalence with the mean curvature vector as shown in (2.38), which can directly be estimated from the enhanced triangulated mesh M . The problem then reduces to finding a good and efficient curvature estimator.

Let us recall, that the *first variation of area* formula (2.11) relates the mean curvature of a hypersurface to the rate of change of its area, as it evolves in the outward normal direction.

3.1.3 Discretizations on 2D triangulations in 3D

Most of the mean curvature discretizations found in the literature deal with 2D manifolds embedded in a 3D space. An intuitive approach to estimate curvatures of a triangulated manifold consists in estimating the first and second fundamental forms presented in the previous chapter. Hamann fitted a bivariate polynomial locally to the vertex of interest and its first-order neighbors, from which the relevant coefficients can be derived (Hamann, 1993). This approach yields both principal curvatures κ_1 and κ_2 separately, so both the intrinsic Gaussian and the extrinsic mean curvature can be computed.

On the other hand, a major family of mean curvature-only estimators are of the following general form:

$$\vec{H}(\vec{x}_i) = \sum_{j \in N_1(i)} w_{ij}(\vec{x}_i - \vec{x}_j), \quad (3.1)$$

where $N_1(i)$ denotes the 1-ring neighbors of vertex i , i.e., vertices that are directly connected to i through a single edge, and where w_{ij} are some weights assigned to this edge. Several estimators of this class have been discussed in (Xu, 2004, 2006).

In a very simple approach, Taubin proposed the inverse of the number of 1-ring neighbors as a possible choice for the edge weight (Taubin, 1995): $w_{ij} = 1/|N_1(i)|$, where $|\cdot|$ denotes the cardinality of a set. The resulting estimator is also called the *umbrella operator*. A more complex proposition by the same author makes use of a positive function ϕ :

$$w_{ij}^{(1)} = \frac{\phi(\vec{x}_i, \vec{x}_j)}{\sum_{k \in N_1(i)} \phi(\vec{x}_i, \vec{x}_k)}, \quad (3.2)$$

where $\phi(\vec{x}_i, \vec{x}_j)$ can for example be the surface area of the two faces sharing the edge (\vec{x}_i, \vec{x}_j) , or some power of the edge-length.

Desbrun et al. proposed a more sophisticated weighting (Desbrun et al., 1999), derived using the first variation of area formula given in (2.11):

$$w_{ij}^{(2)} = \frac{\cot \alpha_j + \cot \beta_j}{4A}, \quad (3.3)$$

where α_j and β_j are the two angles opposite to the edge (\vec{x}_i, \vec{x}_j) in the two triangles sharing this edge, and A is the sum of the areas of the triangles having \vec{x}_i as a common vertex, i.e. the area of the 1-ring patch. The complete mean curvature vector estimator writes as follows:

$$\vec{H}(\vec{x}_i) = \frac{1}{4A} \sum_{j \in N_1(i)} (\cot \alpha_j + \cot \beta_j)(\vec{x}_i - \vec{x}_j). \quad (3.4)$$

Meyer et al's discretization differs only in the area term (Meyer et al., 2002): they choose the area of the Voronoi region around vertex \mathbf{x}_i , so that adjacent patches do not overlap². This difference is of minor importance in the present context and we will focus on (3.4).

3.1.4 Generalization to arbitrary co-dimension

The mean curvature estimator described above has been designed for 2D manifolds embedded in 3D. Interestingly, and Meyer et al. have mentioned it in (Meyer et al., 2002), it generalizes to surface embeddings in arbitrary dimensional space.

In their original derivation, Desbrun et al. used the cross product of two triangle edges to calculate the triangle area (Desbrun et al., 1999):

$$A_n = \frac{1}{2} \|(\vec{x}_i - \vec{x}_j) \times (\vec{x}_i - \vec{x}_{j+1})\|, \quad (3.5)$$

$$A_n^2 = \frac{1}{4} \epsilon_{abc} (\vec{x}_i - \vec{x}_j)_b (\vec{x}_i - \vec{x}_{j+1})_c \epsilon_{ade} (\vec{x}_i - \vec{x}_j)_d (\vec{x}_i - \vec{x}_{j+1})_e, \quad (3.6)$$

where Einstein summation notation was used, and where ϵ_{abc} is the permutation symbol. Fortunately, the tensor notation is exactly the same for the *wedge* product $(\vec{x}_i - \vec{x}_j) \wedge (\vec{x}_i - \vec{x}_{j+1})$, which is the generalization of the cross product to vector spaces different from \mathbb{R}^3 . Accordingly, the *cotangent* between two vectors \vec{u} and \vec{v} is simply generalized as

$$\cot(\vec{u}, \vec{v}) = \frac{\cos(\vec{u}, \vec{v})}{\sin(\vec{u}, \vec{v})} = \frac{\vec{u} \cdot \vec{v}}{\|\vec{u} \wedge \vec{v}\|}, \quad (3.7)$$

and (3.4) becomes perfectly valid for a 2D surface embedded in any higher dimensional space³.

²The Voronoi area is chosen for non-obtuse triangles only, and some specific care needs to be taken for obtuse triangles.

³(Meyer et al., 2002) generalize differently, without exploiting the general form of the original derivation. They avoid the cross product term in the triangle area formulation by replacing it with vector norm and inner-product terms, only, which is also interesting from an implementation point of view.

3.2 Proposed Beltrami scale space

In the system we propose, the curvature of M is estimated in each vertex according to (3.4), and only the feature components of this vector are used to update the mesh's (feature-) coordinates. The feature component of the curvature vector at vertex i writes:

$$[\vec{H}_i]_c = \frac{1}{4A} \sum_{j \in N_1(i)} (\cot \alpha_j + \cot \beta_j)(c_i - c_j), \quad (3.8)$$

viz. c_i the mean-curvature feature at node i , and where the patch area A and the cotangents $\cot \alpha_j$ and $\cot \beta_j$ are computed within the enhanced mesh M .

3.2.1 Comparison with Lopez' discretization

As a comparison, the discretization of the Beltrami operator of Lopez et al. (Lopez-Perez et al., 2004) reads as follows:

$$[\vec{H}_i]_c = \frac{1}{A' \sqrt{1 + \beta |\nabla_{SC}|^2}} \sum_{j \in N_1(i)} \frac{(\cot \alpha'_j + \cot \beta'_j)}{\sqrt{1 + \beta |\nabla_{SC}|^2}} (c_i - c_j), \quad (3.9)$$

where A' and the cotangents $\cot \alpha'_j$ and $\cot \beta'_j$ are computed within the original mesh Σ . The norm of the wedge product, and therefore the area term and the denominator of the cotangent, are simply scaled by $\sqrt{1 + \beta |\nabla_{SC}|^2}$ between original and enhanced mesh, which is reflected by the double appearance of this factor in (3.9) with respect to (3.8). The inner product however behaves a bit more complex, and its scaling is not accounted for. Let \mathbf{u}' and \mathbf{v}' be two edge-vectors in the original manifold (i.e. vectors with zero feature component), and be \mathbf{u} and \mathbf{v} the corresponding vectors in the enhanced mesh. These vectors are related as

$$\begin{aligned} \vec{u} &= \vec{u}' + \vec{n} \langle \nabla_{SC}, \vec{u}' \rangle, \\ \vec{v} &= \vec{v}' + \vec{n} \langle \nabla_{SC}, \vec{v}' \rangle, \end{aligned} \quad (3.10)$$

where ∇_{SC} is the direction of steepest feature ascent within the manifold, and \vec{n} is the unit normal, i.e. the feature basis vector. The inner product in the enhanced mesh now writes

$$\begin{aligned} \langle \vec{u}, \vec{v} \rangle &= \langle \vec{u}' + \vec{n} \langle \nabla_{SC}, \vec{u}' \rangle, \vec{v}' + \vec{n} \langle \nabla_{SC}, \vec{v}' \rangle \rangle \\ &= \langle \vec{u}', \vec{v}' \rangle + \langle \nabla_{SC}, \vec{u}' \rangle \langle \nabla_{SC}, \vec{v}' \rangle. \end{aligned} \quad (3.11)$$

The term $\langle \nabla_{SC}, \vec{u}' \rangle \langle \nabla_{SC}, \vec{v}' \rangle$ vanishes only if at least one of the two vectors \vec{u}' and \vec{v}' is orthogonal to the gradient. In the general case, however, this term represents an error in the estimation, and we expect (3.8) to perform more faithfully.

3.2.2 Numerical PDE solution

In matrix notation, the Beltrami flow of the feature c , using the weights of (3.3) is written as

$$\partial_t \mathbf{c} = W\mathbf{c}, \quad (W)_{ij} = \begin{cases} -w_{ij}, & i \neq j \\ \sum_{k \neq i} w_{ik}, & i = j \end{cases}, \quad (3.12)$$

where $\mathbf{c} = (c_1, \dots, c_N)$ denotes the feature vector over all vertices.

The resulting explicit, progressive Euler scheme for numerical solution reads

$$\mathbf{c}^{k+1} = (I - \tau W^k) \mathbf{c}^k, \quad (3.13)$$

where I is the identity matrix and τ is a small time step. However, this scheme is only conditionally stable, viz. provided that τ is sufficiently small.

On the other hand, the corresponding semi-implicit scheme is

$$(I + \tau W^k) \mathbf{c}^{k+1} = \mathbf{c}^k, \quad (3.14)$$

where the weights W are estimated in the current mesh. To find out if this scheme really creates a discrete scale-space, we check it against the criteria provided in (Weickert et al., 1998). It can be shown, that the inverse of the matrix term $(I - \tau W)$ in the above scheme satisfies both continuity, unit row sum, non-negativity, positive diagonal and irreducibility; but it is not symmetric. Therefore, the scheme is unconditionally stable and exhibits the desired smoothing effects. In the semi-implicit scheme, the time-step τ is only limited by the desired accuracy of the PDE solution. However, due to lack of symmetry, it does not guarantee average gray-level invariance, and the system does not necessarily converge to a constant steady state.

3.2.3 Histogram equalization

In the case of cortical mean-curvature maps, the identical sensors assumption does not necessarily hold. The maps are based on surfaces extracted from MR images, and one has to expect some variations at all levels of the pipeline. Also, the mean curvature map is biased by the size of the underlying brain. Nevertheless, one has strong priors on the correspondence of features between the maps: Sulci and gyri will always be represented by opposite local extrema of the mean curvature map.

The scale affects the intensity distributions of the cortical maps: The “non-creation and non-enhancement of local extrema”-property of the governing PDE causes a reduction of the image intensity dynamics over time. Luckily, normalization of the cortical maps easily corrects this bleaching and bias over scales. Equalization even allows for an improved contrast in the relevant domain, thanks to compression of the distribution tails.

Normalization of the mean curvature allows largely removing the bias between different maps and corrects bleaching over time:

$$c' = \frac{c - \mu(c)}{\sigma(c)}, \quad (3.15)$$

where c denotes the original mean curvature, c' its normalized version, μ and σ are the sample mean and standard deviation respectively. However, this normalization is a shift-and-scale operation only, and does not affect the structure of the intensity distribution. In particular, the outliers keep their extreme position.

One way to act against both bleaching, bias and outliers is histogram equalization. Be $p(c)$ the probability density function (PDF) associated with a map, called its histogram. Then, the cumulative distribution function (CDF) is calculated as:

$$F(c) = \int_{c_{min}}^c p(u)du, \quad (3.16)$$

where c_{min} denotes the lower bound of the value range. For histogram equalization, one is looking for the transfer function $c'' = T(c)$, that renders the intensity distribution uniform. For any given transfer function, the following equation holds:

$$p(c'')dc'' = p(c)dc, \quad (3.17)$$

i.e. the number of points mapped from c to c'' remains unchanged. A uniform distribution has constant PDF, which means one wants $p(c'') = \frac{1}{S}$ constant, where S represents the size of the output range. Rearranging yields

$$\frac{dc''}{dc} = S \cdot p(c), \quad (3.18)$$

and after integration on both sides we get

$$c'' = S \cdot \int_{c_{min}}^c p(u)du = S \cdot F(c), \quad (3.19)$$

which means that to obtain an equalized histogram, the feature is replaced by its cumulative distribution function, and eventually scaled to fit the desired output range. By default, we choose $S = 1$ and the output range is equal to $[0, 1]$.

To simplify the computation of the CDF, we take advantage of the quasi-normality of the mean curvature distribution. Maximum likelihood estimation yields the parameters of the closest Gaussian distribution, of which the CDF can efficiently be determined.

3.2.4 Scale space

The collection of all maps along the PDE integration forms the scale space, whose samples now have coordinates $c_i^{(k)}$, where i is the vertex index, denoting

its spatial location \vec{x}_i , and where k is the index along the time sampling, representing the scale. Notice, that two pairs of samples $(c_i^{k_1}, c_i^{k_2})$ and $(c_i^{k_3}, c_i^{k_4})$ are equidistant in scales, if their time sample indices share the same ratio, i.e. $\frac{k_1}{k_2} = \frac{k_3}{k_4}$. This is equivalent to the statement, that uniform scale sampling requires exponential sampling of the time axis, i.e. $\sigma_k = \sigma_0 \cdot q^k$, where σ_0 is the initial (finest) scale, and q is a multiplicative scale change factor (Morse, 1994).

3.3 Results

In the next few paragraphs, we briefly present and discuss the main results obtained with the proposed Beltrami scale-space on spherical triangulations, applied to cortical mean-curvature maps.

3.3.1 Equalization

First, in figure 3.2, we show the effect of histogram equalization on the cortical mean-curvature maps, projected on to the inflated cortical surface for more intuitive localization. The original feature maps are governed by long-tail outliers, which minimize the relative contrast for most of the map. This is greatly alleviated after equalization, where the tails are compressed and overall contrast is improved. The stratification of the intensity range into a normalized interval also reduces bleaching and fading-out of structures through scales. It therefore theoretically becomes possible, to even register cortical feature maps of different scales. This is potentially interesting when the original feature maps are extracted from subjects at different stages of gyrification, where less developed brains allegedly require less smoothing compared to a mature cortical convolution.

3.3.2 Scale-space

Fig. 3.3 illustrates the scale-space obtained for the right hemisphere of one subject. The equalized feature maps exhibit the desired simplification of cortical structures at coarser scales. The mean magnitude of the surface feature gradient is significantly reduced, and the gradient orientation gains in coherence. Both observations reflect the minimization of the gradient divergence by the Beltrami flow. The ragged outline of structures in the fine-scale map is reduced to abstract sketches at coarser scales, as is exemplarily illustrated by the evolution of the central sulcus in the thresholded image series (marked in red).

3.3.3 Choice of β

A comparison of two choices for the parameter β is provided in Fig. 3.4(a). While small β minimizes the L_2 norm and corresponds to Gaussian-like diffusion, high values approximate TV-norm minimization. The former provides smoother maps, but has tendency to overly merge nearby sulcal structures, whereas the latter results in more jagged, piecewise smooth maps, but keeps structures longer apart.

3.3.4 Morphogenic signification

By the two subject hemispheres in Fig. 3.4(b)&(c) we show, that the simplified large-scale maps of the proposed scale-space do not evolve arbitrarily, but remain anatomically meaningful. In deed, the sulcal regions identified by thresholding of the curvature map correspond well to the anatomical labels of the original map. This comes with the property, that topological differences between subjects, e.g. different interruption or fusion patterns of sulci are generally conserved in the simplified cortical maps, and will only disappear at extremely large scales. Therefore, if two brains with different topology are to be registered, for example, these differences are still to be coped with by the registration algorithm. Nevertheless, this difficult task is greatly alleviated by the important simplification of the structures thanks to the scale-space we present.

3.4 Conclusions

Based on the well-established Beltrami framework, we embedded a spherical feature map of the cortical surface in a higher-dimensional space. The embedded triangulated mesh evolved under Beltrami-flow, and the collection of all generated maps eventually forms a scale-space. Already at fine scales some smoothing is beneficial, as it reduces noise due to surface extraction artifacts and numerical error in the creation of the feature map. Our results further showed the desired simplification of highly convoluted structures into more conceptual sketches at coarser scales. Thus, this scale-space allows a hierarchical representation of structures on the surface of the cerebral cortex.

Computing the scale-space on a spherical inflation of the cortical feature map is beneficial for further processing of the images. The sphere is a well-defined surface and it is easy to define mappings between different subjects, which greatly facilitates the implementation of an image registration pipeline, for example, as will be seen in later chapters of this thesis. The Beltrami framework is capable of dealing with non-scalar features and thus gives direct access to multi-spectral feature maps. Moreover, in contrast to plain diffusion, the Beltrami framework provides an additional degree of freedom that controls the anisotropy. Indeed, some anisotropy is required to favor simplification




Figure 3.2: Histogram equalization of cortical mean curvature maps. **(a)** Original mean curvature map of the right hemisphere, histogram (bars) and cumulative density function (shaded area) of the right hemisphere. A few outliers create long tails and decrease contrast. The distribution is close to normal, as indicated by the superimposed PDF/CDF. **(b)** Map and histogram after equalization. For simplicity, the equalization is based on the cumulative density function estimated using the fitted normal, creating slight deviations from perfect uniformity. The extreme values are compressed and contrast is increased. The output range is automatically mapped to $[0, 1]$. Almost 5% of the original values clot at exactly zero, due to the artificial severing of the interhemispheric junction, thereby creating an important peak after equalization around 0.6. **(c)&(d)** Idem for a different subject.




Figure 3.3: *Scale space of cortical feature maps.* (a) & (b) The equalized spherical map and projection on the partially inflated surface, at scales k . (c) Median thresholding illustrates the simplification of structures. Main structures, e.g. the central sulcus (red), are well preserved in coarser scales. (d) & (e) Magnitude and orientation of the gradient of the equalized map. The map is smoothed, and gradient orientation coherence increased.




Figure 3.4: *Anisotropy parameter β^2 and label coherence.* (a) $\beta \ll 1$ leads to Gaussian-like diffusion, creating smooth maps, while the TV-norm-like $\beta \gg 1$ yields more jagged maps, but better isolates sulci. (b) Thresholded maps for two subjects at $\beta = 1$, $k = 32$. The smoothed maps remain anatomically meaningful through scales. (c) As comparison: the labels of main sulci at finest scale.

versus fusion. The outline of structures is smoothed “early” in the scale-space while fusion of nearby structures is delayed due to steep feature transitions.

Therefore, the simplified sketches well reflect the underlying cortical structures, and the main anatomical content of the maps is conserved. The “generic” brain images can be used in a cortical registration framework, where large-scale features are required to drive the registration process globally, before finer features can be aligned locally. The actual goodness of the proposed scale-space, however, will only be quantifiable after its integration into such an image processing application.

Recently, a similar diffusion scheme for cortical feature maps was proposed by (Joshi et al., 2009). The essential differences with respect to our method are, that (1) they work on p -harmonic parametrizations of the cortical surface (Joshi et al., 2004), instead of our triangulated spherical meshes. (2) They propose an isotropic and a Perona-Malik anisotropic diffusion scheme, instead of our Beltrami scale-space. Finally, (3) they employ a Crank-Nicolson second order numerical scheme in time, compared to the first-order semi-implicit scheme proposed in our approach. This choice offers improved accuracy at the cost of reduced robustness with respect to oscillations.

4

Geodesic Active Fields

“To get through the hardest journey we need take only one step at a time, but we must keep on stepping.”

Chinese proverb.

IN IMAGE REGISTRATION, one looks for the underlying deformation field that best maps one image onto another. This is a classic ill-posed inverse problem, which is usually solved by *adding* a regularization term.

Here, we propose a multiplicative coupling between the registration and the regularization term, which turns out to be equivalent to embedding the deformation field in a weighted minimal surface problem¹. Then, the deformation field is driven by a minimization flow toward a harmonic map corresponding to the solution of the registration problem. This proposed approach for registration shares close similarities with the well-known geodesic active contours model in image segmentation, where the segmentation term (the edge detector function) is coupled with the regularization term (the length functional) via multiplication as well. As a matter of fact, our proposed geometric model is actually the exact mathematical generalization to vector fields of the weighted length problem for curves and surfaces introduced by Caselles-Kimmel-Sapiro (Caselles et al., 1997).

The energy of the deformation field is measured with the Polyakov energy weighted by a suitable image distance, borrowed from standard registration models. We investigate three different weighting functions, the squared error and the approximated absolute error for monomodal images, and the local

¹Parts of this chapter have been published in (Zosso et al., 2011b).

joint entropy for multimodal images.

As compared to specialized state-of-the-art methods tailored for specific applications, our geometric framework involves important contributions. Firstly, our general formulation for registration works on any parametrizable, smooth and differentiable surface, including non-flat and multiscale images. In the latter case, multiscale images are registered at all scales simultaneously, and the relations between space and scale are intrinsically being accounted for. Secondly, this method is, to the best of our knowledge, the first re-parametrization invariant registration method introduced in the literature. Thirdly, the multiplicative coupling between the registration term, i.e. local image discrepancy, and the regularization term naturally results in a data-dependent tuning of the regularization strength. Finally, by choosing the metric on the deformation field one can freely interpolate between classic Gaussian and more interesting anisotropic, TV-like regularization.

4.1 Motivation

In many registration problems, separate objects (Gestalts) in the images are displaced and deformed independently. This is illustrated by a study on the individual movements of separate parts between slices of histological samples, where regularization has been delimited by explicitly modeled boundaries (Pitiot and Guimond, 2008). Other examples can be found in computer vision, where the optical flow often exhibits piece-wise constant or piece-wise smooth regions, with distinct boundaries (Sabatini et al., 2003). Geometric regularization based on the Beltrami embedding of the flow- or deformation field offers some nice advantages in this respect. Flow-driven regularization intrinsically allows for sharper transitions and isolated regions. Further, there are cases where boundaries in images – in terms of intensities or even texture – are good predictors of deformation field boundaries (Nagel and Enkelmann, 1986). In (Ben-Ari and Sochen, 2009) were proposed both, a purely flow-driven and a second, combined flow-intensity driven regularizer, based on the Beltrami framework. This additional intensity cue in the deformation field embedding increases the geodesic distance between independent homogeneous Gestalts and helps defining sharp deformation field boundaries between them.

On another note, we realize that data-dependent regularization has also become important when dealing with outlier pixels. In rigid registration, the influence of mismatching regions can be drastically reduced by cropping the image distance function, e.g., by using Tukey’s biweight instead of squared error as an instance of robust statistics (Reuter et al., 2010). In non-rigid registration, on the other hand, image distance metric and a regularization penalty are commonly incorporated into a single energy minimization model, a.k.a. variational model, e.g., (Hermosillo et al., 2002). The energy functionals are typically of the general form

$$E = E_{\text{data}} + r \cdot E_{\text{regularization}}, \quad r > 0. \quad (4.1)$$

where the balancing parameter r can have a severe impact on the registration result. Its choice is arbitrary and the optimal depends on several conditions. Beyond the simple choice of a constant, one can estimate a local measure of image data reliability to spatially adapt the strength of regularization (Tang et al., 2010), while in atlas-based registration this information can equally be derived from atlas statistics.

The goal of this work is to define a novel image registration scheme using a geometric approach. We couple the registration term and the regularization term locally, **by multiplication**. Hence, we embed the deformation field in a higher dimensional space and define a variational model using the **weighted** Polyakov energy. While the Polyakov energy itself only provides a regularity constraint – harmonic map –, the weighting allows driving the deformation field toward low image dissimilarity. This is in close analogy to geodesic active contours in image segmentation (Caselles et al., 1997), where the segmentation term, i.e. the edge detector function, is coupled with the regularizing length function through multiplication as well. Because our model actually represents a mathematical generalization to vector fields of the weighted length problem for curves and surfaces, we call this model *geodesic active fields* (GAF) for image registration.

As will become clearer in the next sections, the GAF framework has several appealing properties. One of these properties is that the multiplicative link between data and regularization term represents an automatic data-dependent modulation of the local regularization strength by the current alignment quality. The weighting function increases regularization in regions where low matching quality indicates missing confidence, e.g., due to a higher level of noise, whereas lower regularity is required in regions where a good fit can be provided.

The structure of remainder of this chapter is as follows. In section 4.2 we show how the weighted Polyakov framework can be used to define an abstract geometric image registration model. We then derive from this general image registration model several instances for stereo vision as well as flat and non-flat 2D image registration in section 4.3. Then, in section 4.4, we instantiate an extension of the framework to multiscale image registration. Section 4.5 studies different weighting functions. Finally, we show some illustrative, preliminary results obtained with our geodesic active fields framework in section 4.6 and we discuss our model in section 4.7.

4.2 Geodesic Active Fields

In this section we define the general evolution equation for the geodesic active fields for image registration. In contrast to the Beltrami framework for image denoising, we do not embed images, but the deformation field that relates the image pair to be registered. The deformation field is embedded as a mapping between the n -dimensional image domain and a m -dimensional space, where

$m > n$. This is achieved by letting the components of the deformation field become additional dimensions of the embedding space. We will then define metric tensors on the deformation field and the corresponding GAF energy to be minimized. The embedded deformation field manifold then evolves toward a *weighted* minimal surface, where the weighting function attracts it to a deformation field that brings the two images into registration.

In the most general form, we register a pair of n -dimensional images defined on a Riemannian domain Ω with coordinates $\mathbf{x} = (x_1, \dots, x_n)$. The deformation field acts along $p \leq n$ dimensions, i.e., $\mathbf{u} : \Omega \mapsto \mathbb{R}^p$, $\mathbf{u}(\mathbf{x}) = (u_1(\mathbf{x}), \dots, u_p(\mathbf{x}))$.

The embedding X and the metric tensors h_{ij} and $g_{\mu\nu}$ on the deformation field are chosen as follows:

$$\begin{cases} X : (x_1, \dots, x_n) \rightarrow (x_1, \dots, x_n, u_1, \dots, u_p) \\ h_{ij} \text{ is arbitrary} \\ g_{\mu\nu} = \partial_\mu X^i \partial_\nu X^j h_{ij}, \end{cases} \quad (4.2)$$

where x_1, \dots, x_n denote the spatial components of the image and u_1, \dots, u_p are the components of the dense deformation field. These equations are introduced in the weighted Polyakov functional (2.49) and its minimization flow (2.50), leading to the following general registration energy functional and the minimizing evolution flow of the geodesic active fields (GAF):

$$\begin{cases} E_{GAF} = \int f \sqrt{g} \prod_{i=1}^n dx_i \\ \partial_t u_i = f H^{n+i} + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu u_i - \frac{m \cdot n}{2} \partial_k f h^{k(n+i)}, \quad 1 \leq i \leq p, \end{cases} \quad (4.3)$$

where the weighting function $f = f(\mathbf{x}, \mathbf{u})$ is arbitrary for now, and will be defined in more detail in section 4.5.

The main contributions of this framework are:

1. The freedom to register images on any Riemannian manifold, i.e., on any smooth and parametrized surface. This will be developed further in sections 4.3 and 4.4.
2. The invariance under re-parametrization of the proposed energy, like the GAC energy (Caselles et al., 1997) for the segmentation problem.
3. The freedom to choose the metric h_{ij} in the embedding space to obtain different regularizing behavior, as known from the versatility of the Beltrami framework.
4. The intrinsic data-dependent modulation of the local regularization strength thanks to the multiplicative coupling.

In image registration, the property of parametrization invariance is a very rare, but actually highly desirable property. Indeed, there is no reason why the chosen parametrization of the image domain should influence the outcome




Figure 4.1: Omnidirectional *bunny* image obtained from a catadioptric system (Bogdanova et al., 2007). **(a)** Flat 2D raw image obtained from the camera, in Cartesian parametrization. **(b)** Panoramic representation, obtained by polar re-parametrization of the raw disc. **(c)** Mapping of the bunny on the parabolic manifold, $P^2 = (r \cos \phi, r \sin \phi, r^2)$, corresponding to the focal projection of the catadioptric system. Classical registration of images of this type with standard methods will produce different results for the disc and panorama parametrizations. The GAF energy uses a metric corresponding to the actual image geometry and is therefore invariant to the chosen parametrization.

of the registration process. And yet, many currently used image registration methods lack this important invariance property.

The relevance of the contributions one and two can be clarified with the example of catadioptric images illustrated in Fig. 4.1. Such images are widely used in omnidirectional vision and robot navigation, for example, where ego-motion and position can be derived from a sequence of images, e.g. (Yagi et al., 1994; Gaspar et al., 2000; Bunschoten and Krose, 2003). Because standard image registration methods ignore the paraboloid geometry of the actual image, they agnostically work on either one of the flat parametrized image versions. As can be clearly seen, there are important distortions between the raw, disc representation on the one hand, and the polar panorama view on the other hand, in Fig. 4.1(a) and (b), respectively. Obviously, a simple energy like mean squared error, employed in many standard methods, such as the popular Demons algorithm (Thirion, 1998), fails to be re-parametrization invariant on those images,

$$\int (\mathcal{M} - \mathcal{F})^2 dx dy \neq \int (\mathcal{M} - \mathcal{F})^2 d\phi dr, \quad (4.4)$$

and the registration result depends on the selected parametrization. In contrast, with the GAF energy, a metric tensor h_{ij} is derived from the actual image geometry, like the paraboloid in Fig. 4.1(c), and the registration result becomes independent of the chosen image representation.

4.3 Stereo Vision and Image Registration

In the previous section, we have defined the general, abstract energy of GAF and its corresponding gradient descent flow. In the following paragraphs, we instantiate this general concept for specific applications, namely stereo vision and 2D image registration in the Euclidean case, as well as image registration on non-flat manifolds. In other words, we will define specific image geometries and deformation field embeddings, derive the corresponding metric tensors, and thus concretize the GAF energy and its flow. The weighting function f , however, remains unspecified and will be described in detail only later, in section 4.5.

4.3.1 The general Euclidean case

Let us first consider the case of n -dimensional images defined on well-known Euclidean domains Ω with Cartesian coordinates $\vec{x} = (x_1, \dots, x_n)$. We look for a deformation field acting along $p \leq n$ dimensions.

The embedding X of the deformation field, and the corresponding metric tensors h_{ij} and $g_{\mu\nu}$ are chosen as follows:

$$\begin{cases} X : (x_1, \dots, x_n) \rightarrow (x_1, \dots, x_n, u_1, \dots, u_p) \\ h_{ij} = \text{diag}(\underbrace{1, \dots, 1}_n, \underbrace{\beta^2, \dots, \beta^2}_p) \\ g_{\mu\nu} = \partial_\mu X^i \partial_\nu X^j h_{ij} = \delta_{\mu\nu} + \beta^2 \sum_{i=1}^p \partial_\mu u_i \partial_\nu u_i, \end{cases} \quad (4.5)$$

where β is the scaling factor applied to the deformation field components to get the desired aspect ratio. In analogy to the Beltrami framework, this parameter interpolates between isotropic Gaussian, and anisotropic TV-like smoothing of the deformation field. Now, the general Euclidean registration energy functional and the minimizing evolution flow, obtained by plugging the above choice into (4.3), take the following form:

$$\begin{cases} E_{GAF} = \int f \sqrt{g} \prod_{i=1}^n dx_i \\ \partial_t u_i = f H^{n+i} + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu u_i - \frac{m \cdot n}{2\beta^2} f u_i, \quad 1 \leq i \leq p. \end{cases} \quad (4.6)$$

4.3.2 Stereo vision

Simply put, in stereo vision the depth information corresponding to a location is encoded as the lateral shift between its representation in two adjacent image acquisitions (Scharstein and Szeliski, 2002). The recovered depth information is used in e.g. satellite imaging or robot vision to reconstruct the observed scene. The lateral shift can be determined by registration of the two images, where only lateral deformation is allowed. That is, the deformation field $\vec{u}(\vec{x})$,

with $\vec{x} = (x, y)$, has only one component, u along the x -dimension. We choose the following embedding and metric tensors:

$$\begin{cases} X : (x, y) \rightarrow (x, y, u) \\ h_{ij} = \text{diag}(1, 1, \beta^2) \\ g_{\mu\nu} = \begin{bmatrix} 1 + \beta^2 u_x^2 & \beta^2 u_x u_y \\ \beta^2 u_x u_y & 1 + \beta^2 u_y^2 \end{bmatrix}, \quad g = 1 + \beta^2 |\nabla u|^2. \end{cases} \quad (4.7)$$

Introducing those equations into (4.3), we get the following energy functional and evolution equation:

$$\begin{cases} E_{GAF} = \int f \sqrt{1 + \beta^2 |\nabla u|^2} dx dy \\ \partial_t u = f H^u + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu u - \frac{3}{\beta^2} f u, \end{cases} \quad (4.8)$$

where H^u is simply the 3rd component of the mean curvature vector:

$$H^u = \frac{g_{11} u_{yy} - 2 \cdot g_{12} u_{xy} + g_{22} u_{xx}}{g^2}. \quad (4.9)$$

4.3.3 2-D image registration

In the case of registration, involving deformations along all image dimensions, one has $p = n$ and $m = 2n$. Here, as an example without loss of generality, we consider the registration of 2D images. The deformation field is described by $\vec{u} = (u, v)$, respectively along x and y of $\vec{x} = (x, y)$:

$$(u, v) : (x, y) \in \Omega \mapsto (u, v) = (u(x, y), v(x, y)) \in \mathbb{R}^2. \quad (4.10)$$

We choose the following embedding and metric tensors:

$$\begin{cases} X : (x, y) \rightarrow (x, y, u, v) \\ h_{ij} = \text{diag}(1, 1, \beta^2, \beta^2) \\ g_{\mu\nu} = \begin{bmatrix} 1 + \beta^2(u_x^2 + v_x^2) & \beta^2(u_x u_y + v_x v_y) \\ \beta^2(u_x u_y + v_x v_y) & 1 + \beta^2(u_y^2 + v_y^2) \end{bmatrix} \\ g = 1 + \beta^2(|\nabla u|^2 + |\nabla v|^2) + \beta^4 \|\nabla u \times \nabla v\|^2, \end{cases} \quad (4.11)$$

The expression of the determinant g has become quite cumbersome. The term $\beta^4 \|\nabla u \times \nabla v\|^2$ measures the misalignment of the gradients between different deformation field components (Kimmel et al., 2000). All these settings put into the general equations produce the following energy functional and minimizing flow:

$$\begin{cases} E_{GAF} = \int f \sqrt{1 + \beta^2(|\nabla u|^2 + |\nabla v|^2) + \beta^4 \|\nabla u \times \nabla v\|^2} dx dy \\ \partial_t u = f H^u + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu u - \frac{4}{\beta^2} f u \\ \partial_t v = f H^v + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu v - \frac{4}{\beta^2} f v. \end{cases} \quad (4.12)$$

4.3.4 Registration on non-flat manifolds

One of the main contributions of the proposed framework is that the image domain does not necessarily have to be Euclidean. Indeed, images to be registered can be defined on any Riemannian manifold, i.e., on any smooth and parametrized surface. In the Euclidean case, the spatial coordinates were directly given by the image domain parameters. In the non-Euclidean case, the spatial coordinates of the image are more complicated functions of the domain parameters instead.

To give a basic example, that will be illustrated in section 4.6, consider a spherical patch S described by two angles, θ and ϕ , on which the images are defined:

$$\begin{cases} S : (\theta, \phi) \in \Omega \subset \mathbb{R}^2 \rightarrow (x, y, z) \in \mathbb{R}^3 \\ \Omega = [\theta_{\min}, \theta_{\max}] \times [\phi_{\min}, \phi_{\max}] \\ 0 < \theta_{\min} < \theta_{\max} < \pi, \quad 0 \leq \phi_{\min} < \phi_{\max} < 2\pi \\ x(\theta, \phi) = \sin \theta \cos \phi \\ y(\theta, \phi) = \sin \theta \sin \phi \\ z(\theta, \phi) = \cos \theta \end{cases} \quad (4.13)$$

The induced metric on S is naturally given by $g_S = \text{diag}(1, \sin^2 \theta)$. Further, let the deformation field $(\vartheta(\theta, \phi), \varphi(\theta, \phi))$ act on the two angles describing the patch. This suggests the following embedding:

$$\begin{cases} X : (\theta, \phi) \in \Omega \subset \mathbb{R}^2 \rightarrow (\theta, \phi, \vartheta, \varphi) \in \mathbb{R}^4 \\ h_{ij} = \text{diag}(g_S, \beta^2 g_S) = \text{diag}(1, \sin^2 \theta, \beta^2, \beta^2 \sin^2 \theta) \end{cases} \quad (4.14)$$

where the metric tensor h_{ij} has been set by taking the induced metric g_S of the patch parametrization into account. The pullback relation yields the following metric tensor $g_{\mu\nu}$ in parameter space:

$$g_{\mu\nu} = \begin{bmatrix} 1 & 0 \\ 0 & \sin^2 \theta \end{bmatrix} + \beta^2 \begin{bmatrix} \vartheta_\theta^2 + \sin^2 \theta \varphi_\theta^2 & \vartheta_\theta \vartheta_\phi + \sin^2 \theta \varphi_\theta \varphi_\phi \\ \vartheta_\theta \vartheta_\phi + \sin^2 \theta \varphi_\theta \varphi_\phi & \vartheta_\phi^2 + \sin^2 \theta \varphi_\phi^2 \end{bmatrix}. \quad (4.15)$$

Given this metric tensor h_{ij} , the embedding space is not Euclidean anymore, and the computation of the mean curvature vector involves the Levi-Civita connection as in (2.36), to account for the Riemannian part.

For the spherical patch, only two relevant Christoffel symbols computed by (2.37) differ from zero:

$$\Gamma_{jk}^\varphi = \begin{cases} \cot \theta & j \hat{=} \varphi, k \hat{=} \theta, \\ \cot \theta & j \hat{=} \theta, k \hat{=} \varphi, \\ 0 & \text{otherwise,} \end{cases} \quad (4.16)$$

and $\Gamma_{jk}^\vartheta = 0 \ \forall j, k$, where, with some abuse of notation, $j \hat{=} \varphi$ denotes the one j corresponding to the parameter φ , and consequently $\Gamma_{jk}^\varphi = \Gamma_{jk}^i|_{i \hat{=} \varphi}$. This

gives the following evolution equations for the deformation field:

$$\begin{cases} \partial_t \vartheta = f H^\vartheta + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu \vartheta - \frac{4}{\beta^2} f \vartheta, \\ \partial_t \varphi = f H^\varphi + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu \varphi - \frac{4}{\beta^2 \sin^2 \theta} f \varphi, \\ H^\vartheta = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu \vartheta), \\ H^\varphi = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu \varphi) + 2 \cot \theta (g^{\theta\theta} \varphi_\theta + g^{\theta\phi} \varphi_\phi), \end{cases} \quad (4.17)$$

where $g^{\theta\theta} = g^{\mu\nu}|_{\mu=\theta, \nu=\theta}$, and $g^{\theta\phi} = g^{\mu\nu}|_{\mu=\theta, \nu=\phi}$.

4.4 Multiscale Image Registration

4.4.1 Motivation

The concept of scale spaces and multi-scale images has been defined already in section 2.2. In summary, it states that the scale at which one measures a certain property becomes an additional dimension of the imaging space. A scale-space is built by artificially generating larger (coarser) scales of an image through low-pass filtering (Witkin, 1983).

Scale-spaces have particular importance in the context of image registration. As an example, let us consider the human brain. It exhibits a highly convoluted and irregular structure, with high complexity and variability. Considering the complexity of the cortical surface, directly involving local small-scale features would mislead the registration to be trapped in bad local minima. A robust method needs to rely on large-scale features, describing the main landmarks of the cortex, such as the main gyri or sulci, while small-scale features drive the registration more locally to reach the desired precision.

The most intuitive and commonly used approach to multi-scale image registration consists of repeated, hierarchical registration at single scales – from coarse to fine. The result of one stage is used as initialization for the next finer scale. This pyramidal approach has reasonable computational load, but the link between scales is relatively weak, however, and unidirectional: information is only relayed from coarse to fine. Here, we propose a method of registering pairs of entire scale-spaces. All scales are registered simultaneously, thus allowing for bidirectional communication between scales.

The geometry of a large class of scale-spaces can be defined by a general metric tensor (Eberly, 1994; Bresson et al., 2006):

$$h_{ij} = \text{diag} \left(\underbrace{\frac{1}{c^2}, \dots, \frac{1}{c^2}}_{n \text{ times}}, \frac{1}{c^2 \rho^2} \right), \quad (4.18)$$

where the first n elements of the diagonal correspond to the spatial dimensions x_1, \dots, x_n , and the last element refers to the scale σ . c and ρ are two functions that represent the conductance and the density in the general model of heat




Figure 4.2: *Linear and Beltrami scale space.* (a)–(b) Linear and Beltrami scale-spaces of the Von Koch snowflake. The scale σ increases linearly from bottom to top of the image stack, thus constituting an additional image dimension. (c)–(d) Multiscale representation of a T1-weighted magnetic resonance image of a human brain, in a linear and Beltrami scale-space.

transfer. The spatial derivative within such a scale-space is now obtained as $c\nabla$, whereas the scale derivative is given by $\rho c\partial_\sigma$. The natural heat equation, that defines the scale-space, is:

$$\partial_\sigma I = \frac{1}{\rho} \nabla \cdot (c \nabla I). \quad (4.19)$$

Different choices for c and ρ yield different well-known scale-spaces. The linear scale-space, e.g., corresponds to $c = \sigma$ and $\rho = 1$: $\partial_\sigma I = \sigma \Delta I$. The Perona-Malik scale-space is reproduced with $\rho = 1$ and $c = \exp(-\alpha |\nabla I|^2)$, $\alpha > 0$ (Perona and Malik, 1990). The Beltrami flow of Sochen-Kimmel-Malladi requires $c = \rho = \frac{1}{\sqrt{1+\beta^2 |\nabla I|^2}}$ (Kimmel et al., 1997; Sochen et al., 1998). The linear and the Beltrami scale-space are illustrated at the example of the fractal image of a *Von Koch snowflake*, and a single slice of a T1-weighted brain MR image in Fig. 4.2.

4.4.2 Multiscale active deformation fields

Multiscale images have an additional image dimension: the scale σ . Along this scale-dimension, no deformation takes place. The multiscale deformation field is embedded as follows:

$$\begin{cases} X : (x_1, \dots, x_n, \sigma) \rightarrow (x_1, \dots, x_n, \sigma, u_1, \dots, u_n) \\ h_{ij} = \text{diag}(\frac{1}{c^2}, \dots, \frac{1}{c^2}, \frac{1}{c^2 \rho^2}, \frac{\beta^2}{c^2}, \dots, \frac{\beta^2}{c^2}) \\ g_{\mu\nu} = \partial_\mu X^i \partial_\nu X^j h_{ij}, \end{cases} \quad (4.20)$$

where the structure of the metric tensor h_{ij} is arbitrary, and inspired by (4.18).

Considering a linear scale-space, i.e., $c = \sigma$ and $\rho = 1$, the embedding thus looks like:

$$\begin{cases} X : (x_1, \dots, x_n, \sigma) \rightarrow (x_1, \dots, x_n, \sigma, u_1, \dots, u_n) \\ h_{ij} = \frac{1}{\sigma^2} \text{diag}(1, \dots, 1, 1, \beta^2, \dots, \beta^2) \\ g_{\mu\nu} = \partial_\mu X^i \partial_\nu X^j h_{ij} = \frac{1}{\sigma^2} \left(\delta_{\mu\nu} + \beta^2 \sum_{i=1}^n \partial_\mu u_i \partial_\nu u_i \right). \end{cases} \quad (4.21)$$

Again, as for non-flat image domains, the multiscale embedding is not Euclidean, and the Levi-Civita connection (2.37) is required to compute the complete mean curvature vector according to (2.36).

Note, that the deformation field $\vec{u} = \vec{u}(\vec{x}, \sigma)$ evolves at all scales *simultaneously*. At each scale, the deformation field is attracted by the corresponding data term, while coherence between scales is obtained thanks to the regularizing power of harmonic maps.

4.4.3 Multiscale 2-D image registration

In the case of 2D images to be registered, the only relevant non-zero Christoffel symbols computed as (2.37) are $\Gamma_{u\sigma}^u = \Gamma_{\sigma u}^u = \Gamma_{v\sigma}^v = \Gamma_{\sigma v}^v = -\frac{1}{\sigma}$. The evolution equations for both components (u, v) of the deformation field along (x, y) are

$$\begin{cases} \partial_t u = f H^u + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu u - \frac{15\sigma^2}{2\beta^2} f_u \\ \partial_t v = f H^v + \partial_k f g^{\mu\nu} \partial_\mu X^k \partial_\nu v - \frac{15\sigma^2}{2\beta^2} f_v \\ H^u = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu u) - \frac{2}{\sigma} \partial_\mu u g^{\mu\sigma} \\ H^v = \frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu v) - \frac{2}{\sigma} \partial_\mu v g^{\mu\sigma} \end{cases}, \quad (4.22)$$

where, with some abuse of notation, $g^{\mu\sigma}$ denotes the column ν of the inverse of the metric tensor $g^{\mu\nu}$ corresponding to the scale σ .

4.5 Weighting Function for the Registration Problem

The purpose of the weighting function f is to drive the deformation field toward minimal surfaces that bring the two images into registration. As such, the flow should stop when the deformed image perfectly matches the target image. Hence, the weighting function is naturally chosen to be an image distance metric, which approaches zero when the two images match.

4.5.1 Deformation model

The weighting function is the place, where the deformation field actually gets to act on the images. Therefore, it is crucial to define the particular deformation

model we want to use. First, we work with Euler coordinates. That is, for any pixel in the fixed image, the corresponding pixel is looked up in the moving image using a coordinate mapping. The corresponding location in the moving image will almost never fall on an exact pixel location and interpolation will be required.

Here, we use a very simple scheme, where the look-up is based on a shift by addition. The transform operator T is thus defined as:

$$T_0 : T_0 \vec{x} = \vec{x} + \vec{u}(\vec{x}), \quad (4.23)$$

where addition is implicitly understood only along the $p \leq n$ dimensions of the image that are deformed. Also, for simplicity we shall ignore any boundary issues and finite support.

This very basic deformation model embodies only a restricted set of properties. By definition, the displacement needs to be at least twice differentiable, otherwise the Riemannian manifold cannot be constructed and mean curvature cannot be computed. Other than that, no further guarantees exist: the deformation is not necessarily invertible as nothing explicitly prevents the Jacobian to become negative. Further, it is not enforced to be surjective (onto), and homeomorphism or even diffeomorphism are not guaranteed properties. It is important to realize, however, that this is a restriction of the employed deformation model and not of the GAF framework as a whole. More sophisticated deformation models can be used to obtain these properties².

Once the deformation model has been defined, corresponding fixed and moving image locations can be mapped, and the matching quality can be quantified using one of several distance metrics, of which we present some in the following paragraphs.

4.5.2 Squared error

If the images have been acquired using similar sensors, one can generally assume that the same entities are pictured at the same feature intensity in both images. An intuitive and simple choice for monomodal image registration subject to additive Gaussian noise is the squared error metric (Toga, 1999):

$$f^{(1)T_0}(\vec{x}, \vec{u}) = (\mathcal{M}(T_0 \vec{x}) - \mathcal{F}(\vec{x}))^2 = (\mathcal{M}(\vec{x} + \vec{u}(\vec{x})) - \mathcal{F}(\vec{x}))^2, \quad (4.24)$$

where \mathcal{F} and \mathcal{M} refer to the fix and moving images respectively. The evolution equation (4.3) includes the partial derivatives of the weighting function with

²Very recently, Vercauteren *et al.* introduced exponential map diffeomorphisms in the Demons framework (Vercauteren et al., 2009). There, at each iteration one looks for an infinitesimally small update $d\vec{s}$ to $\vec{s} = 0$, that is applied through composition of its exponential map with the existing diffeomorphic deformation. We have integrated this more complicated deformation field model into our GAF framework as well. See chapter 8 on diffeomorphisms.

respect to all components of the embedding. For the function given in (4.24), these are obtained as follows:

$$\begin{cases} f_{\vec{x}}^{(1)T_0} = 2 \cdot (\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x})) \cdot (J^T \nabla \mathcal{M}(T_0\vec{x}) - \nabla \mathcal{F}(\vec{x})) \\ f_{\vec{u}}^{(1)T_0} = 2 \cdot (\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x})) \cdot \nabla \mathcal{M}(T_0\vec{x}) \end{cases} \quad (4.25)$$

where $\nabla \mathcal{F}$ and $\nabla \mathcal{M}$ refer to the gradients of the fix and moving images, respectively, and where J^T denotes the transpose of the Jacobian of the deformed field:

$$J_{ij} = \delta_{ij} + \frac{\partial u_i}{\partial x_j}. \quad (4.26)$$

4.5.3 Local joint entropy

If images of different modality are to be registered, the above squared error metric is not a suitable distance metric anymore. Instead, mutual information is a commonly accepted similarity criterion in this case (Wells III et al., 1996; Viola and Wells III, 1997; Maes et al., 1997).

Mutual information is a global measure on the joint (p^{fm}) and marginal (p^f and p^m) histograms of the fixed and moving images:

$$MI = \sum_{i_1, i_2} p^{fm} \ln(p^{fm}) - \sum_{i_1} p^f \ln(p^f) - \sum_{i_2} p^m \ln(p^m), \quad (4.27)$$

where $p^{fm} = p^{fm}(i_1, i_2)$ etc. Let us assume, that the marginal entropies remain constant throughout the whole registration process, as they only depend on the fix and moving image separately. Maximizing mutual information is thus equal to minimizing the joint entropy.

The same joint entropy, i.e., the expectation of the negative logarithm of the joint probability, can also be computed in the image domain, instead of using the above histograms (Rogelj et al., 2003):

$$H^{fm} = - \sum_{i_1, i_2} p^{fm} \ln(p^{fm}) = \frac{1}{|\Omega|} \sum_{\vec{x}} -\ln(p^{fm}), \quad (4.28)$$

where $p^{fm} = p^{fm}(\mathcal{F}(\vec{x}), \mathcal{M}(\vec{x} + \vec{u}))$. The negative logarithm denotes the local joint entropy. This local joint entropy has a minimum value of 0 (if the joint probability matches 1), and is unbound positive. This provides us with a local measure that corresponds well to the weighting function criteria stated above.

Consequently, we define the following information theory based weighting function for multi-modal image registration:

$$f^{(2)T_0}(\vec{x}, \vec{u}) = -\ln(p^{fm}(\mathcal{F}(\vec{x}), \mathcal{M}(T_0\vec{x}))). \quad (4.29)$$

Using this weighting function, the goodness of a local alignment is measured by the frequency of similar intensity pairs in the rest of the image.

The partial derivatives along spatial components $f_{\vec{x}}^{(2)}$ are easily estimated numerically. The partial derivatives along deformation field components are obtained using the chain rule:

$$f_{\vec{u}}^{(2)T_0} = -\frac{p_m^{fm}(\mathcal{F}(\vec{x}), \mathcal{M}(T_0\vec{x}))}{p^{fm}(\mathcal{F}(\vec{x}), \mathcal{M}(T_0\vec{x}))} \cdot \nabla \mathcal{M}(T_0\vec{x}), \quad (4.30)$$

where $p_m^{fm}(i_1, i_2)$ is the partial derivative of the histogram along the dimension corresponding to the moving image.

4.5.4 Absolute error

For non-smooth deformation fields, e.g., observed in optical flow-based image registration, the L^1 -norm may perform better as data term (Zach et al., 2007). The L^1 -norm measures the absolute error between the two images,

$$f^{(3)T_0}(\vec{x}, \vec{u}) = |\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x})|, \quad (4.31)$$

and it can be approximated by a differentiable function:

$$f^{(3)T_0}(\vec{x}, \vec{u}) = \sqrt{(\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x}))^2 + \varepsilon^2}, \quad (4.32)$$

where $1 \gg \varepsilon > 0$. The partial derivatives of the approximated function are obtained easily:

$$\begin{cases} f_{\vec{x}}^{(3)T_0} = \frac{(\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x}))}{f^{(3)T_0}(\vec{x}, \vec{u})} \cdot (J^T \nabla \mathcal{M}(T_0\vec{x}) - \nabla \mathcal{F}(\vec{x})), \\ f_{\vec{u}}^{(3)T_0} = \frac{(\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x}))}{f^{(3)T_0}(\vec{x}, \vec{u})} \cdot \nabla \mathcal{M}(T_0\vec{x}). \end{cases} \quad (4.33)$$

4.5.5 Data term and regularization balancing

In practice, we found useful to extend the weighting function by a positive constant, to convey a minimal weight to regularization. This is required in two cases: first, a pixel pair might accidentally fit well and locally produce zero discrepancy. As a consequence without a minimal weight, regularization would not be able to release the trapped pixels from their local minima. On the other hand, minimal regularization weight is required by the aperture problem, otherwise displacement would not get propagated into matched, homogeneous regions (Stumpf, 1911; Todorovic, 1996). The general form of the weighting function is thus

$$f = 1 + \alpha f^{(i)}, \quad (4.34)$$

where $f^{(i)}$ is one of the image distance metrics specified above, and α is the balancing parameter, that scales the image discrepancy w.r.t. the constant minimal weight. This form represents a Polyakov energy functional, where the image distance metric corresponds to an additional penalty weighting. A big α will favor high data-fidelity, whereas a small value limits the modulating

impact of the image discrepancy and increases overall regularization. Note that choosing 1 as the minimum weight renders the weighting function in some way symmetric to the square root of the metric tensor determinant, which shares the same lower bound.

4.6 Results

We have implemented the geodesic active fields and ran it on several test problems. Here, the results are presented in order of task complexity. As for all forward schemes, the step length, and thus the speed of the registration, is heavily limited by the stability of the integration. The implementation was done using Matlab®(R2009a) on a standard 2.4GHz Intel®Core™2 Duo desktop machine, running a 64bit Fedora Core 11.

4.6.1 Mean curvature estimation

In (2.38), mean curvature is expressed as the anisotropic divergence of the coordinate gradient. Except for the simple stereo case, where an analytical expression of mean curvature was given, explicit expressions are cumbersome. Instead, we propose to estimate the mean curvature vector numerically, by using central differences twice.

In the 2D case of flat and non-flat images, this amounts to the same scheme as was already proposed by (Ben-Ari and Sochen, 2009) and of which numerical properties have been studied and discussed in (Dascal and Sochen, 2005). The numeric scheme for the mean curvature vector in the multiscale case is obtained in the same manner.

4.6.2 A few words on β and regularization

It might be useful to illustrate the influence of the scaling factor β on the deformation field smoothness. The analysis is easiest in the stereo case. For higher co-dimensions the analysis becomes more tedious and is beyond the scope of the present paper. We refer the reader to similar studies in the field of color and vector image denoising (Kimmel et al., 2000), and optical flow regularization (Ben-Ari and Sochen, 2009). To begin with, a pair of images is registered, where the one-directional deformation field u is initialized with a single local impulse. To study the impulse response of the regularization only, we wish $f = 1$ being constant, and set $\alpha = 0$. The so clutched GAF energy now corresponds exactly to the Beltrami framework for image denoising. Without surprise, after a few iterations, the deformation field has diffused, as illustrated in Fig. 4.3(a). Next, the deformation field is initialized with a unit step, disturbed by uniformly distributed, additive random noise, as shown in Fig. 4.3(b). In Fig. 4.3(c)–(e), the role of the parameter β becomes clear: The regularizer changes from Gaussian filtering for low β , to highly anisotropic,




Figure 4.3: *Regularization for different embedding scales β^2 .* **(a)** Mid-time response, i.e. after some 100 iterations, to a single impulse in the deformation field under constant data term $f = 1$. **(b)** initial unit-step deformation field with uniform additive noise, and **(c)–(f)** its smoothing by the regularizer at different β .

feature preserving TV-norm-like filtering at higher values. The actual choice of the parameter value depends on the available *a priori* knowledge on the deformation field regularity for a specific registration task. For computer vision applications such as motion detection and stereo vision, where entire image regions move as individual blocks (Gestalts), a higher β is preferable to allow for sharp deformation boundaries. Also think of the skulls in Fig. 1.1, where the rigid skull and mandible may be in a different relative pose in the image pair, whereas other applications would require more smooth transitions, thus motivating small β .

4.6.3 Application to stereo vision

An example of stereo vision depth recovery was performed as shown in Fig. 4.4. The image pair `tsukuba` is a well known test image, taken from the middlebury benchmark set for stereo vision. The registration was set up according to the embedding and evolution equation described in section 4.3.2 and using the absolute error weighting function (4.32). In our current implementation of which the goal is to illustrate the concept, the depth recovery result is fair, but does not yet achieve the quality of specifically tailored state-of-the-art stereo vision tools.




Figure 4.4: *Stereo vision depth recovery.* (a)–(b) The tsukuba test image for stereo vision and the image pair difference. (c) The recovered disparity map, (d) Ground truth.

4.6.4 Application to medical imaging

The third case deals with registration of a highly misaligned monomodal medical image pair. Two roughly corresponding axial slices through the T1-weighted MRI volume of different subjects are to be registered. The images have a resolution of 256×256 pixels. Registration is set up with the squared error weighting function. The slices are well aligned by registration, as illustrated in Fig. 4.5. Note, that the subtle differences in the folding pattern cannot effectively be compensated by the dense deformation field, i.e., the global outline of the skull and brain structures are aligned, but gyri and sulci remain largely individual.

Another case aims at registering a pair of multimodal medical images at resolution 317×317 . The first image is the same T1 brain slice as above. The second image now is a deformed slice in T2 weighting. For multimodal registration, we use the information theory based joint local entropy as weighting function (4.29), again in the diffeomorphic setting using exponential map compositions. At the fine resolution, the resampling of an entire image takes considerable time, as well as the computation of the joint histogram. Accordingly, the whole registration process takes around 3 minutes. Registration is widely successful, as is indicated by the before and after checkerboard and




Figure 4.5: 2D registration of brain monomodal MRI slices. **(a)–(b)** Fix and moving image. **(c)** The moving image warped by the recovered deformation field. **(d)** The estimated deformation field. **(e)–(f)** Intensity differences before and after registration.

overlay images provided in Fig. 4.6. Compared to the robust squared error weighting function, the local joint entropy is much more delicate with respect to the initial condition, but allows to register images of different modalities.

4.6.5 Registration on non-flat manifolds

To illustrate the model on a non-flat manifold, we have implemented the spherical patch described in (4.17). First, the purpose of the pullback relation is nicely illustrated in Fig. 4.7. We picture the impulse response that corresponds to the diffusion of a local non-zero spot in the deformation field without data term, i.e., α is simply set to 0, thus $f = 1$. On the spherical patch, the impulse response is isotropic and equal both close to the North pole and close to the equator. Isotropy on the spherical manifold requires a high degree of anisotropy in the rectangular parameter domain, as low- θ regions map denser on the sphere. This required anisotropy is directly obtained thanks to the pullback relation between the metrics h_{ij} on the patch and $g_{\mu\nu}$ in the parameter space. Further, the registration has been tested on an artificially deformed pair of topological maps of the Earth, see Fig. 4.8. The patch spans a good part of the northern hemisphere and some of the southern hemisphere of a globe, hence covering parts of both Americas, entire Europe, Africa, the Atlantic Ocean and of western Asia. Thus, the registration framework is shown to work on non-flat manifolds, such as the sphere.

4.6.6 Multiscale image registration

Finally, the multiscale registration case is tested on a pair of artificially deformed T1 brain images. Images are repeatedly lowpass filtered with a Gaussian to generate a linear diffusion scale-space. The multiscale image stacks prior to and after registration are shown in Fig. 4.9, as well as the corresponding intensity residues. Registration succeeds quite well, as illustrated by the almost entirely removed intensity errors.

4.7 Discussion and Conclusion

In this paper, we have presented a novel, purely geometric method, called geodesic active fields, to register images. The fundamental idea is to embed deformation fields in a weighted minimal surface energy and evolve the deformation field toward minimal surfaces, while being attracted by configurations that bring the images into registration. The process amounts to looking for an optimal hyper-contour in the space of all possible deformations in terms of image mismatch and deformation field regularity. This point of view reveals the close analogy to geodesic active contours in image segmentation (Caselles et al., 1997), that can be derived from a weighted Polyakov energy as well (Bresson et al., 2006), hence the name geodesic active fields (GAF).




Figure 4.6: *Multimodal 2D registration of brain MRI slices.* (a) The fix T1 image. (b) The artificially deformed T2 weighted image. (c)–(d) Red-green overlay of T1 and T2 images prior to and after registration. (e)–(f) The joint intensity histograms prior to and after registration. While initially, the histogram is widely spread, registration results in important histogram focalization.




Figure 4.7: *Diffusion on non-flat image domain.* **(a)** Diffusion takes place in parameter domain (θ, ϕ) , governed by the metric tensor $g_{\mu\nu}$. At low θ , diffusion is highly anisotropic. **(b)** Diffusion as seen on the embedded spherical patch. Both impulse responses look the same and are isotropic. This is obtained through the pullback relation that links the respective metrics h_{ij} on the patch and $g_{\mu\nu}$ in the parameter space.




Figure 4.8: *Registration of a spherical patch of the earth's topographical map.* **(a)** The fixed map with its approximate coastlines highlighted as white contours. **(b)** The moving image, with the coastlines of the fixed image superimposed. **(c)** The warped moving image after registration. The map fits well with the superimposed fixed coastlines, except at a few locations where small misregistration is observed (e.g. Red Sea). The colormap indicates height in meters above (below) sea level. **(d)** The fixed image in rectangular, flat parameter domain. **(e)–(f)** The intensity difference before and after registration.




Figure 4.9: *Multiscale registration of a brain MRI slice.* **(a)** The multiscale stack of the fixed image of a T1 brain slice. **(b)** The moving image of a T1 brain slice, obtained by synthetic deformation of the fixed image. **(c)** The intensity difference illustrates the misalignment at all scales. **(d)** The registered moving image. **(e)** The residue after registration is significantly reduced. Some misregistration is observed at the frontal parts of the skull.

4.7.1 Contributions

In contrast to classic approaches in variational methods, which make use of purely additive competition between data and regularization term, our method combines the two energy contributions in a multiplicative way. In fact, the data term is represented by a local image distance function, that acts as multiplicative weighting on the geometric regularization term, resulting in a weighted surface energy. We recall the main contributions of the proposed framework:

Registration of non-flat and multiscale images

We have derived the minimizing flow of this weighted minimal surface for different image registration configurations. First, the framework applies to standard Euclidean images, defined on Cartesian planes and volumes. Further, our proposed method also directly generalizes to images on Riemannian manifolds, such as non-flat image domains and various scale-spaces, and ultimately the combination of both. In true multiscale registration and in contrast to hierarchic multiresolution approaches, image pairs are registered at all scales simultaneously. Communication between different scales is bidirectionally achieved by the regularization term, smoothing the deformation field across scales. In this context, we contribute a framework which has the advantage over classical approaches of automatically taking the relation between space and scale into account. Useful applications of non-flat image registration can easily be

found in computer vision, e.g. motion detection or scene reconstruction from omnidirectional images.

Parametrization invariance

The second contribution of the proposed framework is the invariance of the registration result with respect to the parametrization chosen to describe the image domain. This result is also very intuitive, as by construction the employed energy measures the weighted hyperarea of the embedded deformation field, which is inherently independent of the parametrization that is used to describe this manifold.

Data-dependent, spatially-adaptive regularization

The multiplicative coupling of data-term and regularization intrinsically produces a data-dependent local modulation of the regularization strength. Naturally, one selects the one image discrepancy measure to be minimized that is the best estimate of alignment quality one can get. It is thus intuitive to let this same reliability estimate tune the local amount of regularization required. In practice, this might be particularly useful in medical image pairs that violate the premier assumption of actual existence of a one-to-one mapping between them, like a pair of images with and without lesions. In these instances, the adaptive regularization might help filling-in “the blanks” with a more regular deformation field than in the surrounding tissues that can be well aligned. It is also useful in images with regions of different noise levels. We thus require a smaller amount of global regularization, compared to classical additive schemes, where the non-adaptive regularization force always causes a bias off the optimal data position in the end result. Also, thanks to the multiplicative coupling, data-term and regularity compete very locally, in contrast to additive methods, where image distance metric and deformation field regularity compete as global measures on the whole image domain. Note, that the data-dependency of the regularization in GAF, based on the current local alignment of images, is different from (Tang et al., 2010), where regularization strength depends on individual image (gradient) intensities.

Geometric regularization.

In a similar context, the geometric nature of the regularization, in particular its freedom to choose the amount of anisotropy through the parameter β , can allow for sharper deformation field transitions than classical Gaussian regularization. This is needed in cases, where individual objects move or deform independently, and where deformation cues from separate objects should not overly interact. Also, TV-like regularization reduces the impact of deformation field outliers, as diffusion of the error is limited. Such outliers can occur at locations of actual image dissent, which can be observed for example with occlusions in stereo vision.

4.7.2 Conclusions

We would like to end this chapter with some concluding remarks:

Weighting functions

We have provided three instances of weighting function, namely squared error (4.24), absolute error (4.32) and local joint entropy (4.29). On the one hand, the absolute and squared error weighting functions minimize the global L^2 and L^1 -norm between the two images, and are suitable for monomodal image registration. The local joint entropy, on the other hand, maximizes the mutual information between images, and lends itself to multimodal image registration.

The parameters α and β

It is important to emphasize the role of the parameters α and β . First, β tunes the aspect ratio between the deformation field dimensions and the spatial dimensions in the embedding. In the simplest case of stereo matching, it has been shown that this allows interpolating between L^2 and L^1 -norm minimization of the deformation field gradient magnitudes, whereas interpretation is more difficult in the general case. Second, note that β only changes the nature of the regularization, but not its relative weight with respect to the data term, which is precisely the role of the balancing parameter α .

Preliminary results and limitations

In this chapter, we have only shown preliminary results, based on very simple discretized forward Euler schemes. These are results for illustrative purposes only, that cannot compete with tightly tailored, and specifically tuned state-of-the-art solutions to practical applications. As we focus on the theoretical and methodological aspects of our image registration framework, we did not develop efficient and accurate numerical schemes to challenge established state-of-the-art methods. Consequently we do not compare quantitatively with other registration methods.

The most stringent limitations of the current GAF implementations are numerical stability (mean curvature estimation) and computational complexity (small time steps). Consequently, the following chapter will, therefore, focus on bringing the GAF energy in a suitable form for more efficient numerical implementations, both in terms of speed, accuracy, and stability.

As mentioned, the embedding we propose for GAF corresponds to the flow-driven geometric regularizer proposed in (Ben-Ari and Sochen, 2009). The second, combined flow-intensity driven regularizer of that article is not exploited in the proposed GAF formulation, but inclusion is straightforward. We propose to go even one step further by embedding textural features rather than intensities, to address cases where Gestalts are defined by regions of homogeneous texture rather than flat intensity.

A Fast Scheme

5

“Denn Zeit ist Leben. Und das Leben wohnt im Herzen.”

Michael Ende, *Momo*.

IN THE PREVIOUS CHAPTER, we have introduced the concept of GAF along with some prototype results for purely illustrative purposes. There, we focused explicitly on the theoretical aspects of GAF and its conceptual strengths, e.g. re-parametrization invariance and applications on non-Euclidean images. The prototype implementations served as a proof of concept, while their main drawback was the relatively poor computational performance, compared to state-of-the-art methods, including Demons.

Here, we want to address these numerical and performance shortcomings and provide a numerical scheme that considerably improves the speed of the GAF energy minimization¹.

The basic idea is to use a splitting and an augmented Lagrangian methods borrowed from optimization theory, such as (Glowinski and Le Tallec, 1989; Lions and Mercier, 1979), which aims at solving some optimization problems efficiently through simpler sub-optimization problems. In our case, it consists in splitting the data term and the regularization term, which are coupled with the product operator. The splitting is then processed in an augmented Lagrangian approach to guarantee the equivalence with the original optimization problem (Nocedal and Wright, 2006). However, we observe that the splitting method is applied to a non-convex functional (product of the data term and the regularization term) which is thus not necessarily guaranteed to converge

¹Parts of this chapter have been published in (Zosso et al., 2011a).

to a global solution, even if all our experiments converged to satisfying solutions. The main interest of the proposed method is the splitting of the rather complex GAF problem into smaller sub-problems, for which fast resolution schemes exist.

The rest of this chapter is organized as follows. First, we will summarize a direct implementation. Then, in section 5.2 we provide a short introduction into the method of augmented Lagrangians (AL) for optimization with constraints. In section 5.3 we will make use of a splitting scheme and AL to get a fast minimization scheme for the GAF energy. We setup both, a stereo vision and a full 2D registration test case, in section 5.4. Results will be presented in section 5.5, and we will give a short conclusion and outlook in section 5.6.

5.1 Geodesic Active Fields – Direct Implementation

5.1.1 Discretization

As we work with sampled images and deformation fields, we adapt the notation of variables and the energy accordingly. First, let $\mathbf{x} \in \mathbb{R}^{N \times n}$ denote the matrix containing the coordinates of all samples, where N is the number of spatial samples, and where n is the dimension of the images. Pixel i is located at $\mathbf{x}_i \in \mathbb{R}^n$. Similarly, $\mathbf{u}_i = \mathbf{u}(\mathbf{x}_i) \in \mathbb{R}^p$ is the deformation vector at \mathbf{x}_i . For any $i \in [1, N]$ let us write $x_{i,j}$, $j \in [1, n]$ and $u_{i,q}$, $q \in [1, p]$ to denote a specific component of those vectors. Finally, the weighting function and the square root of the metric tensor are rewritten as column vectors, $F, G \in \mathbb{R}^N$, respectively:

$$F_i = f(\mathbf{x}_i, \mathbf{u}_i) \quad \forall i \in [1, N] \quad (5.1)$$

$$G_i = \sqrt{g(\mathbf{x}_i, \mathbf{u}_i)} \quad \forall i \in [1, N] \quad (5.2)$$

This allows rewriting the GAF energy in terms of a standard vectorial inner product:

$$E_{GAF} = \langle F, G \rangle = \sum_{i=1}^N F_i G_i = F^T G. \quad (5.3)$$

In summary, the notation changes as follows:

$$\begin{aligned} \mathbf{x} \in \Omega \subset \mathbb{R}^n &\rightarrow \mathbf{x} \in \mathbb{R}^{N \times n} \\ \mathbf{u}(\mathbf{x}) \in \mathbb{R}^p &\rightarrow \mathbf{u}(\mathbf{x}) \in \mathbb{R}^{N \times p} \\ f(\mathbf{x}, \mathbf{u}) \in \mathbb{R}_+ &\rightarrow F(\mathbf{x}, \mathbf{u}) \in \mathbb{R}_+^N \\ \sqrt{g(\mathbf{x}, \mathbf{u})} \in \mathbb{R}_+ &\rightarrow G(\mathbf{x}, \mathbf{u}) \in \mathbb{R}_+^N \\ E_{GAF}(\mathbf{u}) = \int_{\Omega} f \sqrt{g} d\mathbf{x} &\rightarrow E'_{GAF}(\mathbf{u}) = F^T G \end{aligned}$$

5.1.2 Direct implementation for Cartesian images

In the previous chapter, a direct implementation of the GAF minimizing flow (4.12) is given by the following explicit (forward) Euler scheme, for each location $i \in [1, N]$ and deformation field component $p \in [1, p]$:

$$u_{i,q}^{k+1} = u_{i,q}^k + \Delta t \frac{du_{i,q}^k}{dt}, \quad (5.4)$$

$$\begin{aligned} \frac{du_{i,q}^k}{dt} = & F_i H_{i,q} + \sum_{j=1}^{m=n+p} \left(D_j F_i \sum_{\mu,\nu=1}^n g_i^{\mu\nu} D_\mu X_i^j D_\nu u_{i,p}^k \right) \\ & - \frac{mn}{2\beta^2} D_{n+q} F_i \end{aligned} \quad (5.5)$$

where $H \in \mathbb{R}^{N \times p}$ are the deformation field components of the embedded manifold's mean curvature vector, and $H_{i,q} = H_{i,q}(\mathbf{u}^k) \in \mathbb{R}$ denotes the component corresponding to $u_{i,q}$ at pixel i . The computation of H will be detailed below. $g_i^{\mu\nu} \in \mathbb{R}_+$ is the component μ, ν of the metric tensor at pixel i . $D_j F_i \in \mathbb{R}$ is the j -th component of the manifold gradient of F at pixel i . The first order differentials in space, e.g. $D_j F_i$ or $D_j u_{i,p}$, $j \in [1, n]$, are estimated using central differences, whereas the gradients w.r.t. the deformation field components, $D_{n+q} F_i$, $q \in [1, p]$, can be calculated analytically as a function of the estimated image gradients.

5.1.3 Discretized Laplace-Beltrami operator

To estimate the Laplace-Beltrami operator on the deformation field, we use a second order central differences standard scheme, e.g. (Ben-Ari and Sochen, 2009). To simplify the notation, we shorten the expression of the anisotropy tensor:

$$\begin{pmatrix} a & b \\ b & c \end{pmatrix}_i = \sqrt{g_i} g_i^{\mu\nu}. \quad (5.6)$$

The scheme can now be written as a multiplication of a matrix $L \in \mathbb{R}^{N \times N}$ with the deformation field components $\mathbf{u} \in \mathbb{R}^{N \times p}$:

$$H = L\mathbf{u} \quad (5.7)$$

$$L_{i,j} = \frac{1}{\sqrt{g_i}} \cdot \begin{cases} -a(i) - c(i) \\ -\frac{1}{2}(a(i_{+,0}) + a(i_{-,0})) \\ -\frac{1}{2}(c(i_{0,+}) + c(i_{0,-})) & j = i \\ \frac{1}{2}(a(i) + a(j)) & j = i_{\pm 1,0} \\ \frac{1}{2}(c(i) + c(j)) & j = i_{0,\pm 2} \\ \frac{(\pm 1) \cdot (\pm 2) (b(i_{\pm 1,0}) + b(i_{0,\pm 2}))}{4} & j = i_{\pm 1,\pm 2} \end{cases} \quad (5.8)$$

where $i_{\pm,0} \mapsto (m \pm 1, n)$ denotes the index of the right/left-hand neighbor of pixel $i \mapsto (m, n)$, and $i_{0,\pm} \mapsto (m, n \pm 1)$, respectively. Further, \pm_1 versus \pm_2 are independent, whilst all \pm_1 are synchronous.

Based on the CFL condition (Courant et al., 1928), an upper bound on the time step that guarantees stability of such a scheme w.r.t. mean curvature vector flow, i.e. the first component of the complete GAF flow, is provided in (Dascal and Sochen, 2005). In our context, this bound is given by

$$\Delta t \leq \frac{h^2}{8\beta \cdot \max_{m,n} \left\{ \frac{a_{m\pm\frac{1}{2},n}}{\sqrt{g_{m,n}}}, \frac{c_{m,n\pm\frac{1}{2}}}{\sqrt{g_{m,n}}} \right\} \cdot \max_i \{F_i\}} \quad (5.9)$$

where $a_{m\pm\frac{1}{2},n} = \frac{1}{2}(a_{m,n} + a_{m\pm 1,n})$ and $c_{m,n\pm\frac{1}{2}} = \frac{1}{2}(c_{m,n} + c_{m,n\pm 1})$.

5.2 Lagrangian Multiplier

Lagrangian multiplier methods are a powerful and commonly used technique for constrained optimization. They are known to outperform ordinary penalty methods – where the constraint is only encouraged but not enforced (Bertsekas, 1976; Nocedal and Wright, 2006). Indeed, a combination of both ordinary penalty and Lagrangian multiplier terms in the energy functional leads to the so-called Augmented Lagrangian (AL), that exhibits better convergence properties (Glowinski and Le Tallec, 1989; Nocedal and Wright, 2006).

To motivate the AL method briefly, let us consider a simple multidimensional, constrained optimization problem. Let $J, K \in \mathbb{R}$ be real scalar functions of the vector \mathbf{x} . J is the function to be optimized, and K represents a linear equality constraint to be satisfied:

$$\min_{\mathbf{x}} J(\mathbf{x}) \quad s.t. \quad K(\mathbf{x}) = 0. \quad (5.10)$$

5.2.1 Classical Lagrange multiplier

First, introducing a Lagrange multiplier $\lambda \in \mathbb{R}$ allows integrating the constraint within the Lagrangian function L :

$$L(\mathbf{x}, \lambda) = J(\mathbf{x}) + \lambda K(\mathbf{x}). \quad (5.11)$$

Now, it is easy to see that the solutions of (5.10) are located at the stationary points of the Lagrangian, i.e. where $\nabla_{\mathbf{x},\lambda} L = \mathbf{0}$. On the one hand, $\partial L / \partial \lambda = 0$ ensures the constraints are actually verified – i.e. the inner product vanishes –, on the other hand then the derivatives of L with respect to \mathbf{x} need to be zero for J to be extremal.

5.2.2 Augmented Lagrangian

The augmented Lagrangian consists in the introduction of both a Lagrange multiplier to enforce the constraints, and an ordinary quadratic penalty term

that improves the convergence properties of some primal-dual minimization algorithms (Glowinski and Le Tallec, 1989):

$$AL(\mathbf{x}, \lambda, r) = J(\mathbf{x}) + \lambda K(\mathbf{x}) + \frac{r}{2} K(\mathbf{x})^2. \quad (5.12)$$

Given some initial guess for \mathbf{x}^0 and λ^0 , the AL can now be minimized iteratively in three steps:

$$\begin{cases} \mathbf{x}^{k+1} &= \operatorname{argmin}_{\mathbf{x}} AL(\mathbf{x}, \lambda^k, r^k), \\ \lambda^{k+1} &= \lambda^k + \rho^k \cdot K(\mathbf{x}^{k+1}), \\ r^{k+1} &\geq r^k, r^k \rightarrow \infty. \end{cases} \quad (5.13)$$

In (Glowinski and Le Tallec, 1989), it was shown under the assumption K being a linear constraint then (5.13) converges for $r^k = r$ being a positive constant and $0 < \rho^k \leq 2r$. Thus we may choose $\rho^k = r$. Recent TV-regularized image processing problems, such as (Goldstein and Osher, 2009; Tai and Wu, 2009), can be derived from such an AL scheme.

5.2.3 Dealing with multiple constraints

Let $K_i : \mathbb{R}^n \rightarrow \mathbb{R}$, $i = 1 \dots m$ represent m multiple equality constraints. The augmented Lagrangian method can incorporate them all, by defining m corresponding Lagrangian multipliers λ_i :

$$AL(\mathbf{x}, \boldsymbol{\lambda}, r) = J(\mathbf{x}) + \sum_{i=1}^m \lambda_i K_i(\mathbf{x}) + \frac{r}{2} \sum_{i=1}^m K_i(\mathbf{x})^2. \quad (5.14)$$

It is convenient to combine the m different constraint functions into a single, vector-valued constraint $\mathbf{K} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, and to define a vectorial Lagrangian $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_m) \in \mathbb{R}^m$. We may now rewrite the augmented Lagrangian incorporating multiple constraints as

$$AL(\mathbf{x}, \boldsymbol{\lambda}) = J(\mathbf{x}) + \boldsymbol{\lambda}^T \mathbf{K}(\mathbf{x}) + \frac{r^k}{2} |\mathbf{K}(\mathbf{x})|^2, \quad (5.15)$$

where $|\mathbf{K}(\mathbf{x})|^2 = \mathbf{K}^T \mathbf{K}$

Assuming \mathbf{K} being linear, the optimization scheme changes accordingly:

$$\begin{cases} \mathbf{x}^{k+1} &= \operatorname{argmin}_{\mathbf{x}} AL(\mathbf{x}, \boldsymbol{\lambda}^k, r), \\ \boldsymbol{\lambda}^{k+1} &= \boldsymbol{\lambda}^k + \rho \cdot \mathbf{K}(\mathbf{x}^{k+1}). \end{cases} \quad (5.16)$$

5.3 FastGAF Energy Minimization

The results of the previous chapter were based on a direct implementation of the flow (4.3), using a simple forward Euler scheme. Here, we speed up the optimization task with a splitting approach, that minimizes the weighting

function and the metric tensor term of the GAF energy separately, tightly coupled through an AL method.

Let us recall that $\mathbf{u} \in \mathbb{R}^{N \times p}$ is a matrix of size $N \times p$, where p is the co-dimension of the deformation field, i.e., the number of associated deformation field components.

5.3.1 Splitting

Firstly, we transform the unconstrained GAF energy minimization problem over one deformation field \mathbf{u} :

$$\min_{\mathbf{u}} \{E_{GAF} = F(\mathbf{u})^T G(\mathbf{u})\} \quad (5.17)$$

into an equivalent, constrained minimization problem on two coupled deformation fields \mathbf{u} and \mathbf{v} :

$$\min_{\mathbf{u}, \mathbf{v}} \{E_{GAF} = F(\mathbf{u})^T G(\mathbf{v})\} \quad s.t. \quad \mathbf{u} = \mathbf{v}. \quad (5.18)$$

5.3.2 Augmented Lagrangian

We can obtain an unconstrained minimization problem and guarantee to satisfy the linear constraint $\mathbf{u} = \mathbf{v}$ using the following AL scheme (Glowinski and Le Tallec, 1989; Nocedal and Wright, 2006):

$$\min_{\mathbf{u}, \mathbf{v}} \left\{ F(\mathbf{u})^T G(\mathbf{v}) + \langle \boldsymbol{\lambda}, \mathbf{u} - \mathbf{v} \rangle_{N,p} + \frac{r}{2} \|\mathbf{u} - \mathbf{v}\|_F^2 \right\}, \quad (5.19)$$

where $\boldsymbol{\lambda} \in \mathbb{R}^{N \times p}$ is the Lagrangian multiplier matrix, in the same space as \mathbf{u} and \mathbf{v} , $\langle \boldsymbol{\lambda}, \mathbf{u} - \mathbf{v} \rangle_{N,p} = \sum_{i=1}^N \sum_{d=1}^p \lambda_{i,d} (u_{i,d} - v_{i,d})$ is the scalar product between two matrices of the same dimension $N \times p$, and where $\|M\|_F^2 = \sum_{i,j} m_{i,j}^2$ is the square of the Matrix Frobenius norm. From now on, $r > 0$ is a positive constant; how to choose it will be discussed later in section 5.5.2.

Now, we apply algorithm (5.13) to solve the AL problem (5.19):

$$\begin{cases} (\mathbf{u}^{k+1}, \mathbf{v}^{k+1}) = \operatorname{argmin}_{\mathbf{u}, \mathbf{v}} \left\{ F(\mathbf{u})^T G(\mathbf{v}) \right. \\ \quad \left. + \langle \boldsymbol{\lambda}^k, \mathbf{u} - \mathbf{v} \rangle_{N,p} + \frac{r}{2} \|\mathbf{u} - \mathbf{v}\|_F^2 \right\}, \\ \lambda^{k+1} = \lambda^k + \rho(\mathbf{u}^{k+1} - \mathbf{v}^{k+1}), \end{cases} \quad (5.20)$$

Further, this minimization w.r.t. \mathbf{u} and \mathbf{v} can be carried out separately, and we get the following split optimization scheme:

$$\begin{cases} \mathbf{u}^{k+1} = \operatorname{argmin}_{\mathbf{u}} \left\{ F(\mathbf{u})^T G(\mathbf{v}^k) \right. \\ \quad \left. + \langle \boldsymbol{\lambda}^k, \mathbf{u} - \mathbf{v}^k \rangle_{N,p} + \frac{r}{2} \|\mathbf{u} - \mathbf{v}^k\|_F^2 \right\}, \\ \mathbf{v}^{k+1} = \operatorname{argmin}_{\mathbf{v}} \left\{ F(\mathbf{u}^{k+1})^T G(\mathbf{v}) \right. \\ \quad \left. + \langle \boldsymbol{\lambda}^k, \mathbf{u}^{k+1} - \mathbf{v} \rangle_{N,p} + \frac{r}{2} \|\mathbf{u}^{k+1} - \mathbf{v}\|_F^2 \right\}, \\ \lambda^{k+1} = \lambda^k + r(\mathbf{u}^{k+1} - \mathbf{v}^{k+1}), \end{cases} \quad (5.21)$$

In the next paragraphs, we will present how the two respective subminimization problems can be tackled efficiently.

5.3.3 Subminimization w.r.t. \mathbf{u}

The first subminimization problem in (5.21) deals with the optimization of the deformation field through the weighting function $F(\mathbf{u})$, i.e., the image distance function, while keeping the regularization term $G(\mathbf{v}^k)$ fix:

$$E_1(\mathbf{u}) = F(\mathbf{u})^T G(\mathbf{v}^k) + \langle \boldsymbol{\lambda}^k, \mathbf{u} - \mathbf{v}^k \rangle_{N,p} + \frac{r}{2} \|\mathbf{u} - \mathbf{v}^k\|_F^2. \quad (5.22)$$

To simplify things, we linearize F around the current estimate $F(\mathbf{u}^k)$. The Euler-Lagrange equation w.r.t. \mathbf{u} is then obtained as:

$$0 = \text{diag}(G(\mathbf{v}^k)) \cdot \frac{\partial F}{\partial \mathbf{u}}(\mathbf{u}^k) + \boldsymbol{\lambda}^k + r\mathbf{u} - r\mathbf{v}^k \quad (5.23)$$

$$\mathbf{u}^{k+1} = \mathbf{v}^k - \frac{1}{r} \left(\boldsymbol{\lambda}^k + \text{diag}(G(\mathbf{v}^k)) \cdot \frac{\partial F}{\partial \mathbf{u}}(\mathbf{u}^k) \right). \quad (5.24)$$

Stability of this step is limited by the constant r : a small r results in a wider step away from the current estimate, and consequently the linear approximation of the image might not be good enough. We present a solution in the next paragraph.

5.3.4 Balancing the computational complexity

The first order approximation requires small step sizes, i.e. big r . Further, as we will see below, the exact inverse solution of the second problem is computationally more challenging, but stable irrespective of the step size r . It would be interesting, both to balance the computational complexity between the two tasks, and to alleviate the step size restrictions in the first task. Here we propose a fixed-point scheme to address these limitations. First, let us introduce a virtual time t and define a corresponding gradient descent equation:

$$\frac{\partial \mathbf{u}}{\partial t} = - \frac{\partial E_1(\mathbf{u})}{\partial \mathbf{u}} \quad (5.25)$$

We discretize in time using a semi-implicit scheme:

$$\frac{\mathbf{u}^{n+1} - \mathbf{u}^n}{\tau} = - \text{diag}(G(\mathbf{v}^k)) \cdot \frac{\partial F}{\partial \mathbf{u}}(\mathbf{u}^n) - \boldsymbol{\lambda}^k - r\mathbf{u}^{n+1} + r\mathbf{v}^k \quad (5.26)$$

Using $\mathbf{u}^{n=0} = \mathbf{u}^k$ as initial condition, we iterate the scheme

$$\mathbf{u}^{n+1} = \frac{\mathbf{u}^n - \tau \text{diag}(G(\mathbf{v}^k)) \cdot \frac{\partial F}{\partial \mathbf{u}}(\mathbf{u}^n) - \tau \boldsymbol{\lambda}^k + \tau r \mathbf{v}^k}{1 + \tau r} \quad (5.27)$$

until convergence towards a fixed point such that $\mathbf{u}^{k+1} = \mathbf{u}^{n \rightarrow \infty}$. We choose $\tau = \frac{1}{Lr}$, $L \in \mathbb{N}$, where we shorten the step-size of the image disparity gradient by a factor L with respect to the direct solution (5.24). In practice, roughly $2L$ iterations are enough to achieve a satisfying convergence.

5.3.5 Subminimization w.r.t. \mathbf{v}

The second subminimization problem in (5.21) deals with the optimization of the deformation field through the regularization term $G(\mathbf{v})$, while, this time, keeping the weighting function term $F(\mathbf{u}^{k+1})$ fix:

$$E_2(\mathbf{v}) = F(\mathbf{u}^{k+1})^T G(\mathbf{v}) + \langle \boldsymbol{\lambda}^k, \mathbf{u}^{k+1} - \mathbf{v} \rangle_{N,p} + \frac{r}{2} \|\mathbf{u}^{k+1} - \mathbf{v}\|_F^2. \quad (5.28)$$

The derivative of the first term with respect to the deformation field \mathbf{v} can be discretized in analogy to the Laplace-Beltrami operator:

$$\frac{\partial F(\mathbf{u}^{k+1})^T G(\mathbf{v})}{\partial \mathbf{v}} = -\text{diag}(F(\mathbf{u}^{k+1})) \cdot W \mathbf{v} \quad (5.29)$$

where we introduce the Laplacian-like matrix $W \in \mathbb{R}^{N \times N}$ given by

$$W_{i,j} = 2\beta^2 \cdot \begin{cases} -a(i) - c(i) \\ \quad -\frac{1}{2}(a(i_{+,0}) + a(i_{-,0})) \\ \quad -\frac{1}{2}(c(i_{0,+}) + c(i_{0,-})) & j = i \\ \frac{1}{2}(a(i) + a(j)) & j = i_{\pm 1,0} \\ \frac{1}{2}(c(i) + c(j)) & j = i_{0,\pm 2} \\ \frac{(\pm 1) \cdot (\pm 2)(b(i_{\pm 1,0}) + b(i_{0,\pm 2}))}{4} & j = i_{\pm 1,\pm 2} \end{cases} \quad (5.30)$$

that differs from the Laplace-Beltrami discretization L only in its prefactor.

The complete Euler-Lagrange equation w.r.t. E_2 writes:

$$0 = -\text{diag}(F(\mathbf{u}^{k+1})) \cdot W \mathbf{v} - \boldsymbol{\lambda}^k - r(\mathbf{u}^{k+1} - \mathbf{v}) \quad (5.31)$$

$$\mathbf{v}^{k+1} = \left(I - \frac{1}{r} \text{diag}(F(\mathbf{u}^{k+1})) \cdot W \right)^{-1} \left(\mathbf{u}^{k+1} + \frac{1}{r} \boldsymbol{\lambda}^k \right). \quad (5.32)$$

Since we are only interested in the solution of this linear system, it is not necessary to explicitly compute the inverse of the right-hand-side matrix. Instead, more sophisticated solvers can exploit the sparsity of the system. We call this scheme the *exact inverse FastGAF*, in contrast to the approximate inverse FastGAF that will be introduced below. Note that this scheme highly resembles a semi-implicit Euler diffusion scheme, and therefore it can be considered unconditionally stable w.r.t. r (Weickert et al., 1998).

5.3.6 A Jacobi scheme for approximate inversion

Instead of an exact solution, which is computationally expensive, we propose to use an iterative Jacobi scheme as an approximated solution, see e.g., (Duncan and Lynch, 1991). Therefore, let us decompose the system matrix in a diagonal D and a remainder matrix R :

$$D + R = \left(I - \frac{1}{r} \text{diag}(F(\mathbf{u}^{k+1})) \cdot W \right), \quad (5.33)$$

where D and R can be given explicitly:

$$\begin{aligned}
 D_{ii} = & 1 + \frac{2\beta^2 F_i}{r} (a(i) + c(i)) \\
 & + \frac{\beta^2 F_i}{r} (a(i_{+,0}) + a(i_{-,0})) \\
 & + \frac{\beta^2 F_i}{r} (c(i_{0,+}) + c(i_{0,-}))
 \end{aligned} \tag{5.34}$$

$$R_{ij} = \begin{cases} 0 & j = i \\ -\frac{\beta^2 F_i}{r} (a(i) + a(j)) & j = i_{\pm 1,0} \\ -\frac{\beta^2 F_i}{r} (c(i) + c(j)) & j = i_{0,\pm 2} \\ -\frac{(\pm 1)(\pm 2)\beta^2 F_i (b(i_{\pm 1,0}) + b(i_{0,\pm 2}))}{2r} & j = i_{\pm 1,\pm 2} \end{cases} \tag{5.35}$$

Let us denote the right hand term $\mathbf{b} = (\mathbf{u}^{k+1} + \frac{1}{r}\lambda^k)$. Now, the Jacobi scheme approximates \mathbf{v}^{k+1} as a fixed point solution, with $\mathbf{v}^{k+\frac{0}{J}} = \mathbf{u}^{k+1}$ as possible initialization:

$$(D + R)\mathbf{v} = \mathbf{b} \tag{5.36}$$

$$\mathbf{v}^{k+\frac{j}{J}} = \frac{\mathbf{b} - R\mathbf{v}^{k+\frac{j-1}{J}}}{D}, \quad j = 1, \dots, J \tag{5.37}$$

where convergence can be guaranteed under two sufficient conditions:

- $D + R$ is strictly diagonally dominant, i.e. $\forall i : |D_i| > \sum_j |R_{ij}|$, or
- $D + R$ is definite positive, i.e. $\forall \mathbf{z} \neq \mathbf{0} : \mathbf{z}^T (D + R) \mathbf{z} > 0$.

In the present case, the first condition amounts to

$$\forall i : \quad r > \frac{\beta^2 F_i}{2} \sum_{\pm 1, \pm 2} (b(i_{\pm 1,0}) + b(i_{0,\pm 2})), \tag{5.38}$$

i.e. convergence of the Jacobi scheme is not unconditionally guaranteed under the first condition. However, there is an easy upper bound on the step length $\frac{1}{r}$, but the system gets worse conditioned the closer one is to the true (rough) solution. Meanwhile, the second condition, i.e., definite positivity of the Laplacian, is provided and the scheme is expected to converge.

5.3.7 Complete approximate-inverse FastGAF algorithm

The complete FastGAF energy minimization algorithm now summarizes as follows:

$$\left\{ \begin{array}{l} \mathbf{u}^{k+\frac{l+1}{2L}} = \frac{\mathbf{u}^{k+\frac{l}{2L}} - \frac{1}{Lr} \text{diag}(G(\mathbf{v}^k)) \cdot \frac{\partial F}{\partial \mathbf{u}}(\mathbf{u}^{k+\frac{l}{2L}}) - \lambda^k + r\mathbf{v}^k}{1+r} \\ \quad l = 0 \dots (2L-1) \\ \mathbf{v}^{k+\frac{j}{J}} = \begin{cases} \frac{\mathbf{u}^{k+1} + \frac{1}{r}\lambda^k - R\mathbf{u}^{k+1}}{D} & j = 1 \\ \frac{\mathbf{u}^{k+1} + \frac{1}{r}\lambda^k - R\mathbf{v}^{k+\frac{j-1}{J}}}{D} & j = 2 \dots J \end{cases} \\ \lambda^{k+1} = \lambda^k + r(\mathbf{u}^{k+1} - \mathbf{v}^{k+1}) \end{array} \right. \quad (5.39)$$

5.4 Experiments

We test the FastGAF scheme on two example applications: stereo vision disparity recovery, and 2D registration on a brain MRI slice. The results are compared to the direct implementation of the GAF minimizing flow, and the popular Demons algorithm, both in terms of registration error and time requirements. To illustrate the geometrical flexibility of the proposed FastGAF scheme, we show the results of a hemi-spherical toy-registration problem.

5.4.1 Multiresolution optimization and performance metrics

For all registration algorithms we employ the same coarse-to-fine multiresolution strategy. The images are low-pass filtered and sub-sampled at different resolutions and the minimization is executed in multiple stages. We select corresponding/matching parameters for all methods, in terms so of balancing α , metric β and step size r . We measure endpoint error EPE (in pixels) and percentage of wrong pixels $w1$ i.e., rate of pixels off by more than 1 pixel, as registration quality indicators. Each level stops after a pre-definite number of iterations K .

To measure the convergence speed of the different methods, we launch registration for a whole range of different numbers of iterations K , and plot the obtained respective quality measures against the wall clock time t required.

5.4.2 Stereo vision

An example of stereo vision depth recovery problem is shown in Fig. 5.1. The image pair **tsukuba** is a well known test image, taken from the middlebury benchmark set for stereo vision (Scharstein and Szeliski, 2002). We are aware of the existence of a broad variety of high performance algorithms for the stereo vision depth recovery problem. Neither Demons, nor our proposed method are

particularly well suited to compete with those state-of-the-art stereo vision methods. Here, we simply explore the stereo vision problem as an illustrative example due to its relatively simple embedding. The registration is set up according to the embedding described in (4.7) and using the absolute error weighting function (4.32).

Registration takes place at 5 scales and we use the L_1 image distance metric. Parameters are chosen as $\alpha = 5$, $\beta^2 = 10$. For the exact-inverse FastGAF scheme, we use $L = 100$ to balance the computational load between the two subminimization steps, whereas with the approximate-inverse scheme with $J = 4$ Jacobi iterations only, $L = 4$ seems more appropriate.

As shown in Fig. 5.1, the proposed FastGAF scheme outperforms the state-of-the-art Demons method.

5.4.3 2D registration

The second case deals with 2D registration of a highly misaligned monomodal medical image pair. An axial slice through a T1 MRI volume is heavily deformed by a given 2D deformation field. The initial average endpoint error is 7.3 pixels. The images have a resolution of 317×317 pixels and are both affected by 5% additive Gaussian noise. Registration is set up with the squared error weighting function (4.24). The image pair and initial error are illustrated in Fig. 5.2, along with the respective results of the Demons and FastGAF registration. We also compare against a very competitive optical-flow method by Sun et al., that uses a fast non-local regularization approach (Sun et al., 2010).

For the brain registration case we use 5 scales and employ the L_2 distance measure. Further parameters are $\alpha = 50$, $\beta^2 = 3$. With the exact inversion scheme, we use $L = 100$, whereas with approximate inversion we chose $L = 10$ for $J = 4$ Jacobi iterations.

Demons converges to a wrong-pixel rate of $w1 = 13.2\%$ within $t = 664$, corresponding to a final $EPE = 0.515$. The optical flow method converges considerably faster within $t = 26$ seconds, yielding $w1 = 13.3\%$ and $EPE = 0.547$. The proposed FastGAF scheme requires $t = 84$ to obtain $w1 = 6.8\%$, $EPE = 0.327$.

5.4.4 Hemi-spherical registration

One of the main strengths of the Geodesic Active Fields method for image registration, is its intrinsic ability to deal with non-Euclidean images. Here we show an example of a spherical image, parametrized through stereographic projection. While we use the well-known topography of the Earth as a toy example, realistic applications can be found in omnidirectional vision e.g., (Tosic et al., 2005). Indeed, in (Geyer and Daniilidis, 2000, 2001) it was shown, that a sensor image of a catadioptric camera – i.e. using spherical, paraboloid or hyperboloid mirror – is equivalent to a stereographic projection of the spherical plenoptic function.




Figure 5.1: *Depth recovery from stereo: FastGAF versus Demons.* (a)&(b) The tsukuba test image for stereo vision and the image pair difference. (c) Ground truth disparity map. (d) Demons registration. (e) Approximate-inverse FastGAF. (f)&(g) The location of wrong pixels (black) in the Demons and FastGAF result. Both methods primarily “fail” at occlusions.




Figure 5.2: 2D registration of brain MRI slice: FastGAF versus Demons. **(a)&(b)** Fix and moving image. **(c)** Intensity differences before registration. **(d)** Residue after registration with state-of-the-art Demons, **(e)** Sun et al.(Sun et al., 2010), **(f)** FastGAF. (All difference images amplified at equal gain.)

The metric of the stereographic projection is conformal to the regular Euclidean metric (Bogdanova et al., 2007). The choice of the metric h_{ij} of the embedding space is immediate:

$$h_{ij} = \frac{4}{1 + x^2 + y^2} \text{diag}(1, 1, \beta^2, \beta^2) \quad (5.40)$$

Three typical stereographic projections, polar, equatorial and oblique, are sketched in Fig. 5.3(a)–(c). Here we chose an oblique projection of the Earth’s topographical map for the sake of generality. The map is artificially deformed and successfully restored through registration, see Fig. 5.3(d)–(f).

5.5 Discussion

Qualitative results of the registration process using the proposed approximate-inverse FastGAF scheme and the state-of-the-art Demons method were shown in fig. 5.1 and fig. 5.2, respectively. Here, a more quantitative view of the respective algorithm performances is given in fig. 5.4 and fig. 5.5. In the next few paragraphs, we first discuss the performance differences between the four methods considered and highlight the advantages of the proposed scheme. Then, we have a closer look at the different roles of the parameters involved

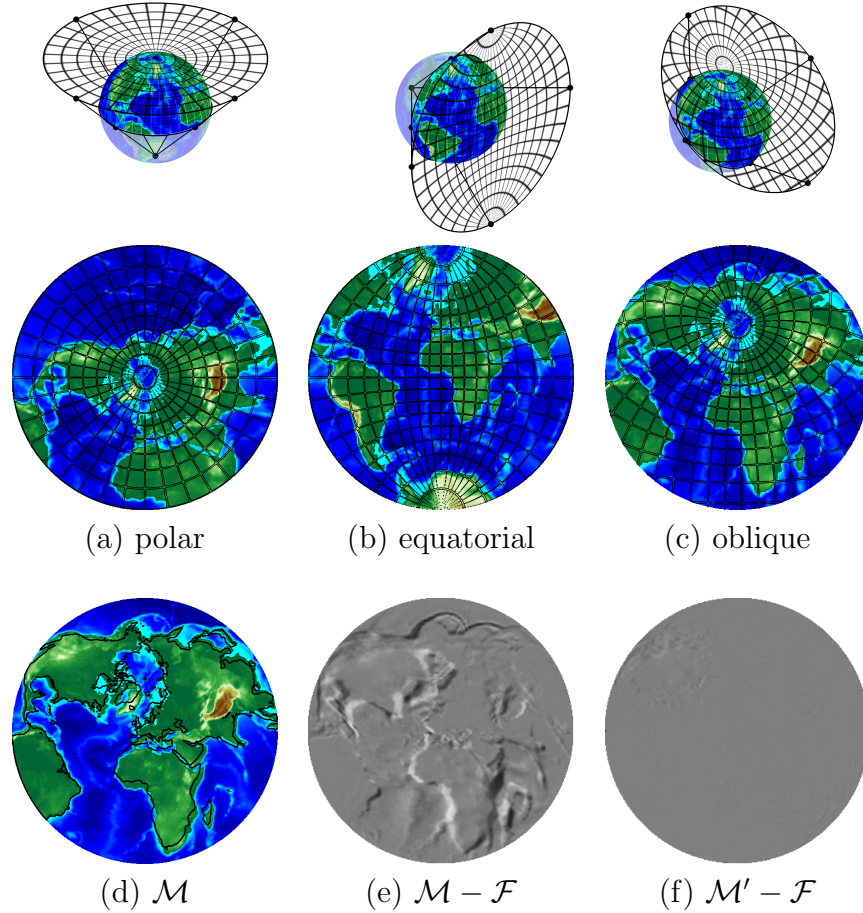


Figure 5.3: *Spherical registration through stereographic projections. (a)–(c)* Different stereographic projections of the sphere. **(d)** Chosen oblique projection and artificial deformation (original coastlines in black). **(e)&(f)** Intensity differences before and after registration.

in the FastGAF scheme. In particular, we highlight the importance of an augmented Lagrangian scheme for constraint internalization, compared to either pure Lagrangian or penalty deformation field coupling.

5.5.1 Algorithms and performance comparison

The four methods compared, i.e., the direct, baseline GAF implementation, the state-of-the-art Demons method, and the proposed accurate- and approximate-inverse FastGAF minimization, share a lot of common features, but have different parameters and stability requirements.

The main commonalities and differences can be resumed as follows:

- With Demons, the step size, i.e. the number of iterations required until convergence, is essentially determined by the desired smoothness of the deformation field. While the smoothing kernel g is fixed, preferably




Figure 5.4: Performance comparison of registration schemes for stereo depth recovery. Demons (red, circles), baseline GAF (green, triangles), exact-inverse FastGAF (cyan, plus), approximate-inverse FastGAF (blue, squares). **(a)** Percentage of wrong pixels (i.e. more than 1 pixel off). **(b)** Average end-point error (*EPE*) on the recovered disparity field.




Figure 5.5: Performance comparison of registration schemes for 2D image registration. Demons (red, circles), baseline GAF (green, triangles), exact-inverse FastGAF (cyan, plus), approximate-inverse FastGAF (blue, squares). **(a)** Percentage of wrong pixels w_1 (i.e. more than 1 pixel off). **(b)** Average end-point error (*EPE*) on the recovered disparity field.

Gaussian, the smoothness is entirely determined by the relative step-size of the data-term, governed by both r and $1/\sigma_i^2$. In this respect, σ_i^2 and σ_T^2 are actually redundant.

- The step size of the direct GAF implementation is heavily limited by the CFL-like stability criterion of the mean curvature flow component. Note that the criterion is depending both on the range of the involved metric, and the gradients of the deformation field. The expected, i.e. non-worst-case stability criterion is actually observed well less restrictive, and a β^2 dependency, rather than β^3 can be realistically assumed. Also, the effect of the central differences approximation of the advective middle term in the minimizing flow is diffusive. Additionally, in contrast to the two other schemes, the curvature flow does never act on its own, but always simultaneously with the advective and data term, which generally acts in favor of stability.
- The proposed accurate-inverse FastGAF scheme offers unconditional stability with respect to the curvature flow. The maximum progression rate is only limited by the stability of the data term, which is greatly enhanced thanks to the L -fold subdivision of the data step. The number of required iterations is therefore reduced importantly, at the expense of higher cost per iteration due to the sparse matrix inversion.
- The latter cost can be drastically reduced by employing an approximate inversion scheme, as proposed with the approximate-inverse FastGAF scheme.

Overall, the approximate-inverse FastGAF scheme can provide the various benefits of the geometric GAF framework for image registration in very attractive computational times, comparing favorably to state-of-the-art methods.

5.5.2 The roles of the parameters r , L and J

The approximate-inverse FastGAF scheme introduces three new parameters to the problem. Here we want to illustrate and discuss their respective role in the optimization framework from a performance-based point of view.

The first novel parameter introduced with the augmented Lagrangian scheme is the penalty weight r . The penalty term is supposed to stabilize the minimization problem. Indeed, it introduces a quadratic energy on the distance between the separate split deformation fields and acts as a “leash” between the two, its elasticity being governed by r . The bigger this parameter, the closer u and v are tied together and the less they can diverge at each iteration. This reduces the lag between u and v and is the only way to restrict the step size at each iteration. The step size limitation is critical with respect to the data term, as due to the complexity of the images under consideration, the linear approximation of the weighting function only holds within a close vicinity of

the current configuration. The impact of r on the registration performance is illustrated in fig. 5.6. It becomes clear, that below a certain level the optimization is unstable with respect to the data term, whereas above a certain threshold only the speed of the algorithm is affected. In particular, it becomes obvious that the penalty term itself is vital for stability, and a pure Lagrangian approach implies severe convergence issues. In contrast, the Lagrangian multiplier is not an absolute requirement for the split GAF iterations to converge, as illustrated in fig. 5.6b). However, as can be seen from the charts, it improves the convergence of the algorithm especially for low penalty weight and thus confirms to be of true interest, considering its little computational extra effort.

In order to overcome the small-steps requirement of the data-term, we propose to carry out several data-step optimizations within a fixed-point scheme before doing a more important smoothing step. The number of data-term optimizations per regularity-term optimization is governed by the parameter L . At r fixed, the image distance term is optimized more precisely, while by replacing $r \leftarrow r/L$, the leash length can be extended at the same image distance precision, thus reducing the number of smoothing steps required. This speed gain is illustrated in fig. 5.7. In particular with the exact-inverse FastGAF scheme, this parameter can be seen as a way of balancing the computational load between both optimization sub-tasks, to gain in efficiency.

In the approximate-inverse FastGAF scheme, finally, we replace the costly full inversion of the diffusion matrix by an approximating Jacobi scheme. There, the number of Jacobi iterations is given by J , controlling the desired precision of the approximate inverse. While the exact inverse has full support (full matrix), now the support of the smoothing stencil is of size $2J + 1$ in each dimension and the speedup is obviously considerable. The performance for different J can be seen in fig. 5.8. Close to the critical step-length r , the algorithm does not properly converge for small, in particular odd, J . Otherwise, the speedup comes at little loss of precision, which is reasonable since the smoothing is essentially operating in a closed loop.

5.6 Conclusions

In this chapter, we have presented a splitting scheme for GAF, based on the method of augmented Lagrangians. The optimization takes place in three alternating steps. In the first problem, the data-term is optimized using a fixed-point scheme derived from a semi-implicit gradient-descent discretization. The second task optimizes with respect to the smoothing-term, initially by exactly solving a sparse linear system, a method we call *exact-inverse FastGAF*. Then, we substitute with an approximate inversion within the semi-implicit smoothing step, using a few fast Jacobi iterations only, to obtain the *approximate-inverse FastGAF* scheme. The third step consists in the update of the Lagrangian multipliers by integration of the residuals between the split deformation field copies. This scheme results in a considerable speedup of




Figure 5.6: Influence of penalty term and Lagrangian multiplier on registration performance. **(a)** Smaller r increases the learning rate, i.e. the step length, at the cost of stability due to a too big data step (supra-critical step-length). Above a certain level, once stability achieved, increasing r only scales the computation time without gaining in precision (sub-critical step-length). **(b)** Compared to penalty-term only, the Lagrangian multiplier has a favorable impact. In particular around the critical step-length, the registration outcome is significantly improved, while farther away in the sub-critical zone, the impact is less pronounced.




Figure 5.7: Computational load balancing L . Subdividing the data-step allows reducing r by L while roughly keeping the same data-stability. This can provide speedups, as in effect a smoothing step (and Lagrangian update) is only performed after L data steps. Relatively small load balancing factors, here about $L = 10$ yields best performance. Beyond, e.g., at $L = 100$, the speed gain becomes negligible while the regularization step is performed too seldom and registration quality is decreased. **(a)** The speedup is important for computationally heavy smoothing task $J = 20$ (or exact inverse). **(b)** For lightweight smoothing, $J = 4$, the speed gains are less pronounced.

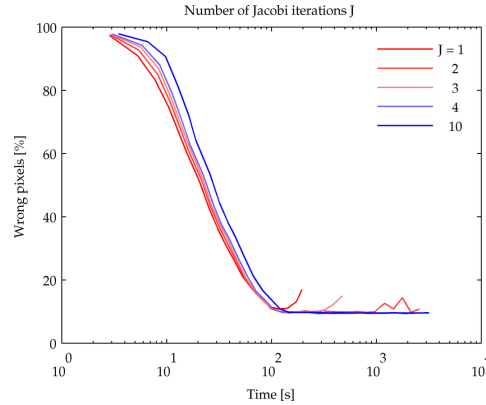


Figure 5.8: *Number of iterations J of the Jacobi approximate inversion scheme.* While too coarse approximations for $J = 1, 2, 3$ are insufficient in the present case, all even $J \geq 4$ yield similar, good registration quality, while scaling the computational time accordingly.

the registration process with respect to the baseline GAF, and also compares favorably to the Demons state-of-the-art registration method. While we describe promising solvers for those sub-problems, note that those can still be improved, which will be investigated in a future work. The Jacobi scheme already provides a nice speedup, but more recent and efficient methods e.g., additive operator splitting (AOS) may further improve performance (Weickert, 1997; Weickert et al., 1998; Malladi and Ravve, 2002; Dascal et al., 2009).

From a more fundamental perspective, geodesic active fields can be considered a generalization of the Demons method in several respects. First, the GAF framework is designed to work on Riemannian manifolds and is thus not restricted to Euclidean images. Although Demons can be generalized to non-flat images as well, e.g. (Yeo et al., 2010), on non-Cartesian grids the speed advantages of Gaussian convolution are lost. While Demons regularization is explicitly Gaussian – it penalizes the L_2 -norm of the deformation field gradient – Beltrami regularization offers a tunable interpolation between Gaussian L_2 and more anisotropic, TV-norm like L_1 regularization. Also, the GAF framework offers the advantage of being parametrization invariant. Finally, the GAF registration framework comes with no preferred image discrepancy measure, whereas Demons has a strong preference for the L_2 -norm on the image differences (SSD). With the approximate-inverse FastGAF scheme presented in this paper, we are able to achieve those relative advantages in very competitive computation times.

Fast Scheme on the Sphere

6

“What is mind but motion in the intellectual sphere?”

Oscar Wilde.

UNTIL NOW, the proposed Geodesic Active Fields registration framework was defined on general manifolds. The main motivation of this thesis, however, is to register cortical feature maps defined on spherical triangulations. Therefore, now we want to apply this framework, to images defined on triangulations of the sphere¹. First, we establish the necessary tools to define the local embedding and to calculate the minimizing flow. Then, we characterize the impulse response and the denoising properties of the regularizer, and illustrate the framework on synthetic images.

6.1 GAF on the sphere

Geodesic active fields embed the deformation field in a higher dimensional space and define a variational model using the weighted Polyakov energy. While the Polyakov energy itself only provides a regularity constraint, the weighting f drives the deformation field towards low image dissimilarity. This approach directly generalizes to non-Euclidean images, and thus automatically allows to work, e.g., with non-flat or multiscale images, that are smooth and parametrizable.

¹Parts of this chapter have been published in (Zosso and Thiran, 2010).




Figure 6.1: *Nomenclature and deformation on the sphere. (a) Deformation model. (b) One-ring patch on a spherical mesh and its local basis vectors. (c) Local coordinates of the one-ring neighbors within the tangential plane.*

Due to the “hairy ball”-theorem, however, it is known that no artifact-free, global parametrization of the whole sphere exists. Therefore the GAF framework cannot be applied directly to the whole spherical image. Instead, we decide to work in local coordinates, by defining a local coordinate chart for each vertex of the mesh. Using those local coordinate charts, the global deformation field can be embedded locally. This involves the (re-)definition of some entities for the spherical case.

6.1.1 Triangulated spherical images

We denote by $\mathcal{F}(\vec{x}), \mathcal{M}(\vec{x}) \in \mathbb{R}$, $\vec{x} \in S^2$, respectively, the fixed and moving spherical image. The images are sampled at a finite number of points, corresponding to the nodes of a triangular mesh of the spherical surface. By convention, we will mainly work with the mesh of the fixed image, composed of the N nodes at $\vec{x}_i \in S^2$, $i = 0 \dots (N - 1)$. The moving image will only be accessed through interpolation at arbitrary locations on S^2 .

For each node \vec{x}_i , every node \vec{x}_j that is directly connected through a common edge is called “1-ring neighbor”, the set of which is written $N_1(i)$.

6.1.2 Local Coordinates and Parallel Transport

Now, we define a local, orthonormal coordinate chart $E_i = [\vec{e}_1 \ \vec{e}_2 \ \vec{e}_3]$, $E_i \in SO(3)$, for each vertex \vec{x}_i :

$$\begin{aligned} \vec{e}_1 &= \frac{\vec{x}_j - \vec{x}_i}{\|\vec{x}_j - \vec{x}_i\|} \wedge \vec{n}_i \\ \vec{e}_2 &= \vec{n}_i \wedge \vec{e}_1 \end{aligned} \tag{6.1}$$

$$\vec{e}_3 = \vec{n}_i \tag{6.2}$$

where \vec{x}_j is one of the 1-ring neighbors, $j \in N_1(i)$, as illustrated in figure 6.1(b), and \vec{n}_i denotes the unit outward normal to the sphere at \vec{x}_i . The projection of

the coordinates of the neighbors into the newly defined local coordinate system of the tangent plane is then given by

$${}_i\vec{\theta}_j = E_i^T(\vec{x}_j - \vec{x}_i) \quad (6.3)$$

Note that the node \vec{x}_i has coordinates ${}_i\vec{\theta}_i = (0 \ 0 \ 0)^T$ in its own local coordinate chart, see figure 6.1(c).

To work with the tangent vectors of neighboring vertices, we report them to \vec{x}_i using parallel transport. The tangent vector \vec{t}_j is parallelly transported along the great circle defined by \vec{x}_i and \vec{x}_j . Let the axis $\vec{\omega}_{ij}$ of this great circle be

$$\vec{\omega}_{ij} = \frac{\vec{x}_j \wedge \vec{x}_i}{\|\vec{x}_j \wedge \vec{x}_i\|}. \quad (6.4)$$

The great circle is a geodesic, and therefore parallel transport has two main properties: components perpendicular to the geodesic remain perpendicular, and components along it remain parallel. Thus we construct for the parallelly transported ${}_i\vec{t}_j$:

$${}_i\vec{t}_j = \langle \vec{t}_j, \vec{\omega}_{ij} \rangle \vec{\omega}_{ij} + \langle \vec{t}_j, (\vec{\omega}_{ij} \wedge \vec{x}_j) \rangle (\vec{\omega}_{ij} \wedge \vec{x}_i) = \Omega_{ij} \vec{t}_j, \quad (6.5)$$

where the coefficients of the rotation matrix $\Omega_{ij} \in SO(3)$ can be identified as

$$[\Omega_{ij}]_{ab} = [\vec{\omega}_{ij}]_a [\vec{\omega}_{ij}]_b + [\vec{\omega}_{ij} \wedge \vec{x}_j]_a [\vec{\omega}_{ij} \wedge \vec{x}_i]_b. \quad (6.6)$$

Let (t_j^1, t_j^2) denote the tangential coefficients of the vector \vec{t}_j w.r.t. the basis E_j :

$$\vec{t}_j = E_j \begin{pmatrix} t_j^1 \\ t_j^2 \\ 0 \end{pmatrix} \quad (6.7)$$

and analogously $({}_i t_j^1, {}_i t_j^2)$ are the respective coefficients after local transport to \vec{x}_i and expressed w.r.t. the basis E_i :

$${}_i\vec{t}_j = E_i \begin{pmatrix} {}_i t_j^1 \\ {}_i t_j^2 \\ 0 \end{pmatrix} \quad (6.8)$$

Consequently, those coefficients are related through a 2D rotation matrix $\Lambda_{ij} \in SO(2)$:

$$\begin{pmatrix} {}_i t_j^1 \\ {}_i t_j^2 \end{pmatrix} = [{}_i \vec{e}_1 \ {}_i \vec{e}_2]^T \Omega_{ij} [{}_j \vec{e}_1 \ {}_j \vec{e}_2] \begin{pmatrix} t_j^1 \\ t_j^2 \end{pmatrix} = \Lambda_{ij} \begin{pmatrix} t_j^1 \\ t_j^2 \end{pmatrix}, \quad (6.9)$$

where the transformation is obtained as $\Lambda_{ij} = [{}_i \vec{e}_1 \ {}_i \vec{e}_2]^T \Omega_{ij} [{}_j \vec{e}_1 \ {}_j \vec{e}_2]$.

6.1.3 Deformation Field

Before we can set up the embedding, we need to define the deformation field model. We encode the individual displacement of each mesh vertex \vec{x}_i into \vec{x}_i' as a local tangent vector \vec{u}_i , as illustrated in figure 6.1(a). The actual displacement of the mesh node is on the great circle along this tangent vector, and the sine of the angle between \vec{x}_i and \vec{x}_i' is equal to the length of the tangent vector \vec{u}_i , as shown in figure 6.1(a). This choice, and the fact that all vertices reside on the unit sphere, provides a simple expression for \vec{u}_i using vector products (Yeo et al., 2008):

$$\vec{u}_i = -\vec{x}_i \wedge (\vec{x}_i \wedge \vec{x}_i'). \quad (6.10)$$

Conversely, the displaced vertex is given as follows:

$$\vec{x}_i' = \sqrt{1 - \|\vec{u}_i\|^2} \cdot \vec{x}_i + \vec{u}_i \quad (6.11)$$

The displacement vector \vec{u}_i at node \vec{x}_i can be described by its two tangent coefficients in the respective local basis. To parametrize the whole set of N deformation field vectors, we write their local coordinate chart coefficients as a column vector $\mathbf{u} \in \mathbb{R}^{2N}$, where

$$\vec{u}_i = E_i \begin{pmatrix} u_{2i} \\ u_{2i+1} \\ 0 \end{pmatrix}. \quad (6.12)$$

If we substitute this expression for \vec{u}_i into the definition (6.11) of the displaced vertex \vec{x}_i , we get

$$\vec{x}_i'(\mathbf{u}) = \sqrt{1 - u_{2i}^2 - u_{2i+1}^2} \cdot \vec{x}_i + [\vec{e}_1 \quad \vec{e}_2] \begin{pmatrix} u_{2i} \\ u_{2i+1} \end{pmatrix}, \quad (6.13)$$

since E_i is orthonormal.

6.1.4 Data term

An intuitive primer for monomodal image registration is the squared error metric (Toga, 1999), leading to:

$$f(\vec{x}, \vec{x}') = 1 + \frac{\alpha}{2} \cdot (\mathcal{M}(\vec{x}') - \mathcal{F}(\vec{x}))^2, \quad (6.14)$$

where \mathcal{F} and \mathcal{M} refer to the fix and moving images, respectively. The 1 provides a lower bound of local weight in the Polyakov energy, and α allows to balance the impact of the data-term with respect to the regularization. The data-term is evaluated at the nodes of the fixed image, and we write $F \in \mathbb{R}_+^N$ for the set of observations, where

$$F_i = f(\vec{x}_i, \vec{x}_i'). \quad (6.15)$$

The data-term clearly depends on the deformation field and its local parameters. Substituting these local coordinates coefficients yields:

$$F_i(\mathbf{u}) = 1 + \frac{\alpha}{2} \cdot \left(\mathcal{M} \left(\sqrt{1 - u_{2i}^2 - u_{2i+1}^2} \cdot \vec{x}_i + E_i \begin{pmatrix} u_{2i} \\ u_{2i+1} \\ 0 \end{pmatrix} \right) - \mathcal{F}(\vec{x}_i) \right)^2. \quad (6.16)$$

Now, the derivatives with respect to the local coefficients u_{2i} and u_{2i+1} are obtained using the chain rule:

$$\frac{\partial F_i}{\partial u_{2i}} = \alpha \cdot (\mathcal{M}(\vec{x}_i') - \mathcal{F}(\vec{x}_i)) \cdot \left\langle \frac{\partial \vec{x}_i'}{\partial u_{2i}}, \vec{\nabla} \mathcal{M}(\vec{x}_i') \right\rangle, \quad (6.17)$$

$$\frac{\partial F_i}{\partial u_{2i+1}} = \alpha \cdot (\mathcal{M}(\vec{x}_i') - \mathcal{F}(\vec{x}_i)) \cdot \left\langle \frac{\partial \vec{x}_i'}{\partial u_{2i+1}}, \vec{\nabla} \mathcal{M}(\vec{x}_i') \right\rangle, \quad (6.18)$$

where $\vec{\nabla} \mathcal{M}$ refers to the 3D gradient of the moving image.

The derivatives $\frac{\partial \vec{x}_i'}{\partial u_{2i}}$ and $\frac{\partial \vec{x}_i'}{\partial u_{2i+1}}$ in turn are calculated as:

$$\frac{\partial \vec{x}_i'}{\partial u_{2i}} = \frac{-\vec{x}_i \cdot u_{2i}}{\sqrt{1 - u_{2i}^2 - u_{2i+1}^2}} + {}_i\vec{e}_1 \quad (6.19)$$

$$\frac{\partial \vec{x}_i'}{\partial u_{2i+1}} = \frac{-\vec{x}_i \cdot u_{2i+1}}{\sqrt{1 - u_{2i}^2 - u_{2i+1}^2}} + {}_i\vec{e}_2. \quad (6.20)$$

For later use, we introduce the “forked” vector $\mathbb{F} \in \mathbb{R}^{2N}$:

$$\mathbb{F}_{2i}(\mathbf{u}) = \mathbb{F}_{2i+1}(\mathbf{u}) = F_i(\mathbf{u}) \quad (6.21)$$

and define its flow $\mathbb{F}' \in \mathbb{R}^{2N}$ as its relevant partial derivatives with respect to the deformation field

$$\mathbb{F}'_{2i}(\mathbf{u}) = \frac{\partial F_i}{\partial u_{2i}} \quad (6.22)$$

$$\mathbb{F}'_{2i+1}(\mathbf{u}) = \frac{\partial F_i}{\partial u_{2i+1}} \quad (6.23)$$

6.1.5 Deformation Field Embedding and Energy

In the geodesic active fields framework, the registration problem is solved by minimizing the Polyakov energy calculated over a single, global embedding. Since no such artifact-free embedding of the whole sphere exists, and because the Polyakov energy is parametrization invariant, we may sum up the energy contributions of a covering set of local embeddings. Here, we propose to parcel the spherical surface into the set of N Voronoi regions V_i around \vec{x}_i of the triangular mesh that covers it. We can then approximate the Voronoi-patch-wise Polyakov energy based on the local coordinates defined above.




Figure 6.2: *Neighborhood and Voronoi region on triangular meshes. (a) One-ring neighbors and angles opposite to an edge. (b) Voronoi region on a non-obtuse triangle. (reproduced from Meyer et al. (2002))*

More formally, since

$$S^2 = \bigcup_{i=0}^{N-1} V_i, \quad (6.24)$$

$$\forall i \neq j : V_i \cap V_j = \emptyset, \quad (6.25)$$

the global Polyakov energy can be computed as a sum of those N local contributions:

$$S = \int_{S^2} f \sqrt{g} d\theta = \sum_i \int_{V_i} f \sqrt{g} d\theta = \sum_i S_i, \quad (6.26)$$

where V_i denotes the Voronoi region around \vec{x}_i , and where each such patch can be parametrized differently. Thereby we instantiate N local embeddings X_i , one for each Voronoi patch.

On the local Voronoi region V_i , the Polyakov energy can be integrated using different techniques. We propose a quadrature that uses the terminology given in figure 6.2.

First, let ∂V_{ij} denote the Voronoi boundary between \vec{x}_i and \vec{x}_j , as well as its length. On that boundary, we assume the deformation field to be the mean of its samples at \vec{x}_i and \vec{x}_j . We further suppose the deformation field to vary linearly within the Voronoi sub-region V_{ij} , spanned by \vec{x}_i and ∂V_{ij} . Therefore, the deformation field components' gradients are constant on that same region. Further, these gradients are both oriented orthogonal to the boundary ∂V_{ij} . Finally, we consider the data-term to be constant over the whole Voronoi patch. The determinant of the metric tensor is thus simplified to

$$g_{ij}(\mathbf{u}) = 1 + \beta^2 \frac{(i u_{2j} - u_{2i})^2 + (i u_{2j+1} - u_{2i+1})^2}{\|\vec{x}_j - \vec{x}_i\|^2}, \quad (6.27)$$

where

$$\begin{pmatrix} i u_{2j} \\ i u_{2j+1} \end{pmatrix} = \Lambda_{ij} \begin{pmatrix} u_{2j} \\ u_{2j+1} \end{pmatrix} \quad (6.28)$$

are the deformation field coefficients of node \vec{x}_j parallelly transported to and expressed in the local basis of node \vec{x}_i . Now, the local Polyakov energy can be approximated by

$$S_i \approx \sum_{j \in N_1(i)} f(\vec{x}_i, \vec{x}_i') \sqrt{g_{ij}} V_{ij}. \quad (6.29)$$

Let us now compute V_{ij} . One has for the angles in the triangle PQR:

$$a + b + c = \pi/2, \quad (6.30)$$

and therefore $a = \pi/2 - \angle Q$ and $c = \pi/2 - \angle R$. Therefore we have

$$\partial V_{ij} = \frac{1}{2} (\cot \alpha_{ij} + \cot \beta_{ij}) \|\vec{x}_j - \vec{x}_i\| \quad (6.31)$$

and immediately

$$V_{ij} = \frac{1}{8} (\cot \alpha_{ij} + \cot \beta_{ij}) \|\vec{x}_j - \vec{x}_i\|^2 \quad (6.32)$$

The complete spherical GAF energy is thus approximated as:

$$E_{GAF} \approx \frac{1}{8} \sum_i f(\vec{x}_i, \vec{x}_i') \sum_{j \in N_1(i)} \sqrt{g_{ij}} (\cot \alpha_{ij} + \cot \beta_{ij}) \|\vec{x}_j - \vec{x}_i\|^2 \quad (6.33)$$

6.2 Minimizing Scheme

In the following paragraphs, we will first define a scheme to efficiently minimize the spherical GAF energy. We then specify the tools required to estimate image gradients and to perform point localization and mesh interpolation.

6.2.1 Discretization and Splitting

First, let us rewrite the approximated spherical GAF energy minimization in terms of an inner product:

$$\min_{\mathbf{u}} \{F(\mathbf{u})^T G(\mathbf{u})\} \quad (6.34)$$

where we introduce the vector $G \in \mathbb{R}^N$ as

$$G_i(\mathbf{u}) = \sum_{j \in N_1(i)} \sqrt{g_{ij}(\mathbf{u})} (\cot \alpha_{ij} + \cot \beta_{ij}) \|\vec{x}_j - \vec{x}_i\|^2 \quad (6.35)$$

Again, we fork the disparity term G to obtain $\mathbb{G} \in \mathbb{R}^{2N}$:

$$\mathbb{G}_{2i}(\mathbf{u}) = \mathbb{G}_{2i+1}(\mathbf{u}) = G_i(\mathbf{u}) \quad (6.36)$$

Now, as in the previous chapter, we transform this unconstrained energy minimization problem over one deformation field \mathbf{u} into an equivalent, constrained minimization problem on two coupled deformation fields \mathbf{u} and \mathbf{v} :

$$\min_{\mathbf{u}, \mathbf{v}} \{F(\mathbf{u})^T G(\mathbf{v})\} \quad s.t. \quad \mathbf{u} = \mathbf{v}. \quad (6.37)$$

We can obtain an unconstrained minimization problem and guarantee to satisfy the linear constraint $\mathbf{u} = \mathbf{v}$ using the following augmented Lagrangian scheme (Glowinski and Le Tallec, 1989; Nocedal and Wright, 2006):

$$\min_{\mathbf{u}, \mathbf{v}} \left\{ F(\mathbf{u})^T G(\mathbf{v}) + \langle \boldsymbol{\lambda}, \mathbf{u} - \mathbf{v} \rangle + \frac{r}{2} \|\mathbf{u} - \mathbf{v}\|^2 \right\}, \quad (6.38)$$

where $\boldsymbol{\lambda} \in \mathbb{R}^{2N}$ is the Lagrangian multiplier, and $r > 0$ is a constant.

This augmented Lagrangian problem (6.38) can now be iteratively solved in three separate steps:

$$\begin{cases} \mathbf{u}^{k+1} = \operatorname{argmin}_{\mathbf{u}} \left\{ F(\mathbf{u})^T G(\mathbf{v}^k) + \langle \boldsymbol{\lambda}^k, \mathbf{u} - \mathbf{v}^k \rangle + \frac{r}{2} \|\mathbf{u} - \mathbf{v}^k\|^2 \right\}, \\ \mathbf{v}^{k+1} = \operatorname{argmin}_{\mathbf{v}} \left\{ F(\mathbf{u}^{k+1})^T G(\mathbf{v}) + \langle \boldsymbol{\lambda}^k, \mathbf{u}^{k+1} - \mathbf{v} \rangle + \frac{r}{2} \|\mathbf{u}^{k+1} - \mathbf{v}\|^2 \right\}, \\ \lambda^{k+1} = \lambda^k + r(\mathbf{u}^{k+1} - \mathbf{v}^{k+1}), \end{cases} \quad (6.39)$$

6.2.2 Data-term optimization

The first sub-problem minimizes with respect to the deformation field \mathbf{u} and hence optimizes the data term F . We denote by E_1 the energy associated with the first sub-problem:

$$E_1 = F(\mathbf{u})^T G(\mathbf{v}^k) + \langle \boldsymbol{\lambda}^k, \mathbf{u} - \mathbf{v}^k \rangle + \frac{r}{2} \|\mathbf{u} - \mathbf{v}^k\|^2 \quad (6.40)$$

Let us introduce a virtual time t and define a corresponding gradient descent equation with respect to u_{2i} :

$$\frac{\partial u_{2i}}{\partial t} = - \frac{\partial E_1(\mathbf{u})}{\partial u_{2i}} \quad (6.41)$$

We discretize in time using a semi-implicit scheme with time-step τ :

$$\frac{u_{2i}^n - u_{2i}^{n+1}}{\tau} = \frac{\partial F_i(\mathbf{u}^n)}{\partial u_{2i}} G_i(\mathbf{v}^k) + \boldsymbol{\lambda}_{2i}^k + r(u_{2i}^{n+1} - v_{2i}^k) \quad (6.42)$$

$$\frac{u_{2i+1}^n - u_{2i+1}^{n+1}}{\tau} = \frac{\partial F_i(\mathbf{u}^n)}{\partial u_{2i+1}} G_i(\mathbf{v}^k) + \boldsymbol{\lambda}_{2i+1}^k + r(u_{2i+1}^{n+1} - v_{2i+1}^k) \quad (6.43)$$

Using $\mathbf{u}^{n=0} = \mathbf{u}^k$ as initial condition, we iterate the scheme

$$u_{2i}^{n+1} = \frac{u_{2i}^n - \tau \frac{\partial F_i(\mathbf{u}^n)}{\partial u_{2i}} G_i(\mathbf{v}^k) - \tau \boldsymbol{\lambda}_{2i}^k + \tau r v_{2i}^k}{1 + \tau r} \quad (6.44)$$

$$u_{2i+1}^{n+1} = \frac{u_{2i+1}^n - \tau \frac{\partial F_i(\mathbf{u}^n)}{\partial u_{2i+1}} G_i(\mathbf{v}^k) - \tau \boldsymbol{\lambda}_{2i+1}^k + \tau r v_{2i+1}^k}{1 + \tau r} \quad (6.45)$$

until convergence towards a fixed point such that $\mathbf{u}^{k+1} = \mathbf{u}^{n \rightarrow \infty}$. Using the forked vectors $\mathbb{F}(\mathbf{u})$ and $\mathbb{G}(\mathbf{v})$ we can rewrite the scheme as

$$\mathbf{u}^{n+1} = \frac{\mathbf{u}^n - \tau \mathbb{F}'(\mathbf{u}^n) \circ \mathbb{G}(\mathbf{v}^k) - \tau \boldsymbol{\lambda}^k + \tau r \mathbf{v}^k}{1 + \tau r} \quad (6.46)$$

where $\mathbf{a} \circ \mathbf{b} = \text{diag}(\mathbf{a})\mathbf{b}$ denotes component-wise multiplication of two vectors of same size.

In practice, we choose the time-step as $\tau = \frac{1}{Lr}$, $L \in \mathbb{N}$, and observe that roughly $2L$ iterations are enough to achieve a satisfying convergence.

6.2.3 Regularity-term optimization

The second minimization problems optimizes the deformation field \mathbf{v} with respect to the regularity-term G . The associated energy to be minimized is E_2 :

$$E_2 = F(\mathbf{u}^{k+1})^T G(\mathbf{v}) + \langle \boldsymbol{\lambda}^k, \mathbf{u}^{k+1} - \mathbf{v} \rangle + \frac{r}{2} \|\mathbf{u}^{k+1} - \mathbf{v}\|^2 \quad (6.47)$$

The Euler-Lagrange equations for the E_2 energy with respect to the deformation field \mathbf{v} are easily deduced:

$$0 = F_i(\mathbf{u}^{k+1}) \frac{\partial G_i(\mathbf{v})}{\partial v_{2i}} - \lambda_{2i}^k - r(u_{2i}^{k+1} - v_{2i}) \quad (6.48)$$

$$0 = F_i(\mathbf{u}^{k+1}) \frac{\partial G_i(\mathbf{v})}{\partial v_{2i+1}} - \lambda_{2i+1}^k - r(u_{2i+1}^{k+1} - v_{2i+1}) \quad (6.49)$$

We establish the derivative of the regularity term G :

$$\frac{\partial G_i(\mathbf{v})}{\partial v_{2i}} = \sum_{j \in N_1(i)} \frac{\partial g_{ij}(\mathbf{v})}{\partial v_{2i}} \cdot \frac{\cot \alpha_{ij} + \cot \beta_{ij}}{2\sqrt{g_{ij}(\mathbf{v})}} \|\vec{x}_j - \vec{x}_i\|^2 \quad (6.50)$$

where further:

$$\frac{\partial g_{ij}(\mathbf{v})}{\partial v_{2i}} = \frac{-2\beta^2}{\|\vec{x}_j - \vec{x}_i\|^2} ({}_i v_{2j} - v_{2i}) \quad (6.51)$$

Thus:

$$\frac{\partial G_i(\mathbf{v})}{\partial v_{2i}} = - \sum_{j \in N_1(i)} \frac{\beta^2 (\cot \alpha_{ij} + \cot \beta_{ij})}{\sqrt{g_{ij}(\mathbf{v})}} ({}_i v_{2j} - v_{2i}) \quad (6.52)$$

and by analogy:

$$\frac{\partial G_i(\mathbf{v})}{\partial v_{2i+1}} = - \sum_{j \in N_1(i)} \frac{\beta^2 (\cot \alpha_{ij} + \cot \beta_{ij})}{\sqrt{g_{ij}(\mathbf{v})}} ({}_i v_{2j+1} - v_{2i+1}) \quad (6.53)$$

Let us recall how the neighboring deformation field parameters are parallelly transported to the central node:

$${}_i v_{2j} = \Lambda_{ij}^{(11)} v_{2j} + \Lambda_{ij}^{(12)} v_{2j+1} \quad (6.54)$$

$${}_i v_{2j+1} = \Lambda_{ij}^{(21)} v_{2j} + \Lambda_{ij}^{(22)} v_{2j+1} \quad (6.55)$$

We rewrite the Euler-Lagrange equations in matrix form:

$$\left(I_{2N} + \frac{\beta^2}{r} \text{diag}(\mathbb{F}(\mathbf{u}^{k+1})) \cdot W(\mathbf{v}^{k+1}) \right) \mathbf{v}^{k+1} = \frac{1}{r} \boldsymbol{\lambda}^k + \mathbf{u}^{k+1} \quad (6.56)$$

where I_{2N} is the $2N \times 2N$ identity matrix and where the coefficients of the matrix $W(\mathbf{v}) \in \mathbb{R}^{2N \times 2N}$ are identified as:

$$W_{ab} = \begin{cases} -w_{ij}(\mathbf{v})\Lambda_{ij}^{(11)} & j \in N_1(i) \wedge a = 2i \wedge b = 2j \\ -w_{ij}(\mathbf{v})\Lambda_{ij}^{(12)} & j \in N_1(i) \wedge a = 2i \wedge b = 2j + 1 \\ -w_{ij}(\mathbf{v})\Lambda_{ij}^{(21)} & j \in N_1(i) \wedge a = 2i + 1 \wedge b = 2j \\ -w_{ij}(\mathbf{v})\Lambda_{ij}^{(22)} & j \in N_1(i) \wedge a = 2i + 1 \wedge b = 2j + 1 \\ \sum_{j \in N_1(i)} w_{ij}(\mathbf{v}) & (a = b = 2i) \vee (a = b = 2i + 1) \\ 0 & \text{otherwise} \end{cases}, \quad (6.57)$$

where $\forall j \in N_1(i)$ the coefficients $w_{ij}(\mathbf{v})$ are given by:

$$w_{ij}(\mathbf{v}) = \frac{\cot \alpha_{ij} + \cot \beta_{ij}}{\sqrt{g_{ij}(\mathbf{v})}} \quad (6.58)$$

Since the coefficients w_{ij} vary with the deformation field \mathbf{v} , the fully implicit system (6.56) is non-linear and difficult to solve. Therefore we approximate $W(\mathbf{v}^{k+1})$ by $W(\mathbf{v}^k)$ and rather optimize the linear, semi-implicit scheme:

$$\left(I_{2N} + \frac{\beta^2}{r} \text{diag}(\mathbb{F}(\mathbf{u}^{k+1})) \cdot W(\mathbf{v}^k) \right) \mathbf{v}^{k+1} = \frac{1}{r} \boldsymbol{\lambda}^k + \mathbf{u}^{k+1} \quad (6.59)$$

6.2.4 Spherical tools

A few tools required to process data on spherical meshes, namely gradient estimation, point location and interpolation, are presented in the appendix.

6.3 Results

We will, in the next chapter, use this spherical registration framework to build an automatic parcellation scheme of the human cerebral cortex, based on spherical feature maps and associated manual parcellations of gyral regions. Here, we simply demonstrate the validity of the chosen approach by presenting the regularizer behavior (impulse response and denoising) of the GAF framework on the sphere, and provide a simple example of registration based on a synthetic image pair.

6.3.1 Impulse response and denoising

The regularization of the GAF framework is characterized by its impulse response and denoising properties. The impulse response is obtained by disconnecting the data-term, i.e., by setting $\alpha = 0$ and thus the weighting function to




Figure 6.3: *Deformation field regularization on the sphere.* **(a)-(b)** Impulse response on a coarse and fine sphere mesh. **(c)-(d)** Denoising for small and large values of β . Small values correspond to Gaussian-like diffusion, whereas large values approximate TV-diffusion.

$f = 1$. The diffusion of an initial single peak is illustrated in figure 6.3(a)-(b). Regularization is largely independent of the sampling density, as the field looks almost the same after two runs on a fine and a coarse sphere mesh. Denoising for two different β is illustrated in figure 6.3(c)-(d). Smaller β produces Gaussian-like diffusion, whereas higher β corresponds to feature-preserving TV-norm minimization.

6.3.2 Synthetic image pair

To illustrate the complete registration framework, a pair of synthetic images has been registered as shown in figure 6.4. The parameters were chosen as follows:

- squared error weighting function
- $\alpha = 2.5$ (weighting function balance)
- $\beta = 0.05$ (aspect ratio of deformation field embedding)
- $r = 40$ (augmented Lagrangian penalty)
- $\rho = 0.1r$ (Lagrangian update step-size)
- $L = 2$ (data-step length, resulting in 6 fixed-point scheme iterations)
- $J = 50$ (number of Jacobi iteration in regularization-step)
- linear interpolation

Convergence is achieved in less than 100 iterations. Registration is successful, as the resampled moving image $\mathcal{M}(\vec{x}')$ nicely matches the fixed image $\mathcal{F}(\vec{x})$, under reasonable deformation $\vec{x} \rightarrow \vec{x}'$.




Figure 6.4: *Registration of a synthetic image pair on the sphere. (a)-(b) Fix and moving image. (c) Warped moving image after successful registration. (d) Deformed mesh.*

6.4 Discussion and Conclusions

We have introduced the geodesic active fields framework for image registration to applications, where images are defined on spherical meshes. We have shown that the proposed method works as expected under different aspects. The regularizer has the desired impulse response irrespective of the image sampling density, the anisotropy of the diffusion is tunable by the parameter β , and the framework performs correctly on a set of synthetic test images.

There are, finally, a few particularities of the spherical case, that we want to quickly discuss in the next paragraphs.

6.4.1 Trivial solutions and rigid registration

It is to note that, in the spherical case, the all-zero deformation field is the only one with trivial unweighted Polyakov energy. This is in contrast to planar image parametrizations with Neumann boundary conditions, where all constant-valued deformation fields are equally regular. Thus, the boundary-free mesh of the sphere allows only the null solution as steady state in absence of a data-term. While in the plane, the Beltrami regularizer is transparent to global (translational) deformation fields, this is clearly not the case on the sphere, where the most global deformation is described by a rotation around a single axis, which already incurs a cost in terms of Beltrami energy.

If such a global rotation between two spherical images is to be compensated for, one should consider pre-registering the image pair globally, e.g., using a method of moments or PCA, or by optimizing a global parametric variational model.

6.4.2 Matrix inversion

Due to the varying connectivity of the vertices in the mesh, and since their ordering is entirely arbitrary and irregular, the diffusion matrix W is sparse but not regularly structured. In particular it does not exhibit a tri-diagonal banded

structure at all. In MATLAB, it takes almost as much time to construct the sparse matrix, as is afterwards required to actually solve the linear system of equations using the Jacobi algorithm described earlier.

6.4.3 Data interpolation

The most important computational load, however, is spent on data interpolation. As is outlined in the appendix, spherical warping and resampling requires to locate the containing triangle on the moving image mesh for each query point of the fixed mesh. Even if, thanks to good initialization and sophisticated guessing, the correct triangle is tested as first guess, determining the exact position of the query point within the triangle in terms of barycentric coordinates requires resolution of a small linear system. This is less likely when the mesh resolution increases, since the same displacement of a query point with respect to the previous look-up covers a greater number of triangles to be checked.

Based on the calculated barycentric coordinates, both the moving image feature and its gradient need to be interpolated at that specific location. Although we have implemented cubic Hermitian interpolation, we stick to simple linear interpolation due to computation time considerations. Even so, since the moving image needs to be resampled several times (here: 6 times) within the fixed-point scheme, the data-term optimization takes roughly 4 times longer than 50 iterations of the Jacobi scheme for approximated matrix inversion.

In conclusion, we realize that an important speed-up could be obtained by re-implementing the spherical warping in a faster way, e.g., using C and MEX for use with all other MATLAB routines.

Mapping the Human Brain

7

三人成虎

“Three men make a tiger.”

Chinese proverb.

IN THE INTRODUCTION TO THIS THESIS, we mentioned the utility of cortical registration to the establishment of more and more precise functional localizations and mapping of the human cortex, in particular within statistical group studies. Before the proposed FastGAF scheme on the sphere can actually be used in such a context, we should first provide some validation. We will do so by inverting the data-flow: we use existing cortical label maps to evaluate the performance of our registration scheme: In this chapter, we use the previously defined spherical registration scheme to build an automated parcellation system of the human cortex.

Based on an ensemble of 39 different subject brains, each of which was manually labeled into 35 gyral regions per hemisphere, we propose to test the influence of pairwise registration on the performance of parcellation-by-label-fusion. Therefore, we setup a leave-one-out cross-validation, where the information of 38 brains is projected onto the remaining query-brain, without and after spherical registration, and the predicted parcellation is then compared to the available manual “ground-truth”. We can show, that the proposed registration scheme has a small, but significantly positive influence on the performance of this simple automatic parcellation scheme.

7.1 Background and Concept

Mainly, two different classes of automatic labeling and parcellation systems for the human cortex have been developed so far. One class achieves sulcal basin identification based on learned shape and neighborhood statistics, e.g. (Lohmann and von Cramon, 2000; Rivière et al., 2002; Tao et al., 2002; Tu et al., 2007). The other class is characterized by subject-to-template registration and subsequent local classification, e.g. (Fischl et al., 2004; Desikan et al., 2006).

7.1.1 Proposed parcellation scheme

Here, we follow the latter parcellation scheme, i.e., we want to assign a label to each bit of a subject cortex based on fusion of label information available for a certain number of reference brains. We will then want to illustrate the improvement of parcellation thanks to the spherical registration. Since the focus of this thesis is on image registration rather than atlas fusion and classification schemes, we will largely rely on existing classification and label fusion schemes. In essence, here we substitute the registration scheme proposed in (Fischl et al., 2004) with our own Fast Geodesic Active Fields framework, and try reproducing their sophisticated classifier as close as it gets.

7.1.2 Leave-one-out cross-validation

The parcellation-validation framework takes the following “leave-one-out” structure: First, all feature maps of the same (right) hemisphere in the data set are Beltrami-low-pass filtered, as described in chapter 3, so as to construct an anisotropic scale space for each hemisphere. Then, the resulting maps of all subjects are mutually registered using the “Fast Geodesic Active Fields on the Sphere” framework as seen in chapter 6, in a brute-force pairwise registration scheme. Based on these individual mappings, for each subject the available feature and label maps of every other brain in the set can be projected on to the local brain. At this point, for every node of a subject’s feature map, there is a list of registered feature and label maps issued of all other brains in the data set, based on which the local label of the subject brain can be “predicted”. We then compare those predictions to the available manual ground truth for each individual brain and obtain an error measure (Dice similarity coefficient, DSC). The same label prediction can be made without prior registration of subject brains, i.e., simply assuming identity transforms between the different subjects’ maps. Comparing the hereby resulting error measure to the complete framework, including registration, will highlight the (hypothetically beneficial) impact of registration on the parcellation scheme.

7.2 Materials and Methods

In the next paragraphs, we will describe in more detail the available brain data, the registration scheme and parameters used for pairwise registration, the exact classification scheme employed for Bayesian label fusion onto the subject brain surface, as well as the specific implementation using shell scripts for job scheduling and the progress monitoring dashboard.

7.2.1 Brain Data

The set of 39 cortical maps used in this study are part of the “Anonymous Buckner” data set, which was already used in (Desikan et al., 2006), and which was made publicly available by B.T. Thomas Yeo¹. As stated in their data disclaimer, these data are a small subset of subjects originating from Washington University, collected by Randy Buckner and colleagues.

Surfaces and features

Only the surfaces and labels are provided so as to preserve the anonymity of the original subjects. Indeed, from the original MR image volumes, the cortical white-matter-gray-matter and the gray-matter-CSF (pial) surface was extracted, separately for each hemisphere using FreeSurfer (Dale et al., 1999). In addition, smoothed, inflated and spherically parametrized versions are included as well (Fischl et al., 1999a).

Moreover, the data set also contains feature maps, such as the mean curvature of the white matter surface and of the inflated surface, as well as average convexity (sulcal depth, (Fischl et al., 1999a)).

Manual label maps

Finally, manual labeling of each hemisphere into 35 gyral units by Raul Desikan (Desikan et al., 2006) constitutes the most valuable part of the data set. The labels are respectively shown on the pial, the inflated, and the spherical surface in figure 7.1. These manual labels will serve as atlas information in the training set, and will simultaneously represent the ground-truth information that allows some performance measurement of the registration-classification scheme within a leave-one-out cross-validation scheme. A indicative list of the different regions is provided in table 7.1. For an exact definition of the how the regions were delineated, however, we refer to the original paper (Desikan et al., 2006).

¹see http://people.csail.mit.edu/ythomas/code_release/AnonBuckner39.tar.gz at the time of this thesis

lateral

medial




Figure 7.1: *Manual parcellation scheme in the Buckner cerebral cortex data set.* The delineation was done manually in a sulcal approach, i.e., by “tracing from the depth of one sulcus to another, thus incorporating the gyrus within” (Desikan et al., 2006). The parcellation is shown both from a lateral and medial view, as well as on **(a)** the pial surface, **(b)** the inflated surface, **(c)** the sphere. The pial view allows for easiest identification of the gyral regions, while the inflated surface reveals otherwise hidden buried structures, such as the banks of the superior temporal sulcus or the pericalcarine fissure. The spherical surface is the configuration which is actually used for registration and classification.

Table 7.1: *List of parcellated structures.* Reproduced from (Yeo et al., 2010), according to the definitions in (Desikan et al., 2006).

1. Unlabeled	13. Lateral Orbito Frontal	25. Pre-Central Gyrus
2. Banks of Sup. Temporal S.	14. Lingual	26. Pre-Cuneus
3. Caudal Anterior Cingulate	15. Medial Orbito Frontal	27. Rostral Anterior Cingulate
4. Caudal Middle Frontal G.	16. Middle Temporal Gyrus	28. Rostral Middle Frontal
5. Corpus Callosum	17. Parahippocampal	29. Superior Frontal Gyrus
6. Cuneus	18. Paracentral	30. Superior Parietal Complex
7. Entorhinal	19. Pars Opercularis	31. Superior Temporal Gyrus
8. Fusiform Gyrus	20. Pars Orbitalis	32. Supramarginal
9. Inferior Parietal Complex	21. Pars Triangularis	33. Frontal Pole
10. Inferior Temporal Gyrus	22. Pericalcerine fissure	34. Temporal Pole
11. Isthmus Cingulate	23. Post-Central Gyrus	35. Transverse Temporal
12. Lateral Occipital	24. Posterior Cingulate	

7.2.2 Registration scheme

Given the complexity of the cortical feature maps, the high computational cost of fine scaled spherical resampling, and the relatively important deformations to compensate, we propose to register the cortical feature maps in a hierarchical multiscale approach, both in terms of feature scale and mesh resolution.

Multiscale registration

The most intuitive and commonly used approach to hierarchical multi-scale registration is to perform repeated registration at single scales, from coarse to fine. The result of one stage is used as initialization for the next finer scale. In contrast to simultaneous registration at multiple scales, this sequential approach has reasonable computational load. However, the link between scales is relatively weak, and unidirectional: information is only passed from coarse to fine. Only the information of the next coarser scale is exploited, and as pure initialization it is not necessarily considered in any further iterations.

Downsampling

At each scale of the multiresolution pyramid, appropriate scales of both feature maps to be registered are resampled on a “regular” spherical mesh, obtained as subdivision of an Icosahedron, see section A.1 in the appendix. The subdivision depth is increasing at each level, and the resampled maps become finer at each stage. This resampling allows for considerable speed-up, in particular at coarse level, since the number of vertices is drastically reduced from more than 100’000 down to just a few thousand vertices. The resampled meshes are then registered using a fixed number of iterations of the spherical FastGAF scheme. At the first, coarsest scale, the deformation fields are initialized by the identity transform (zero deformation), while at finer scales the result of the previous scale is up-sampled through linear interpolation. Also, we store the deformation fields at the end of each stage in separate files for later use.

Registration parameters

Based on just a very few trial runs on a single pair of subject feature maps, we chose the following parameters for the Beltrami scale-space generation:

- $\beta = 0.01$ (embedding aspect ratio)
- $\tau = 0.01$ (PDE integration time step)

And for the FastGAF registration framework, we chose:

- $\alpha = 0.75$ (data versus regularization balance)
- $\beta = 1$ (embedding aspect ratio)
- $r = 50$ (weight of the Augmented Lagrangian penalty term)
- $L = 4$ (semi-implicit data-step subdivisions)
- 12 data-step fixed-point scheme iterations
- 100 iterations per scale, enough to converge at each scale
- linear interpolation
- absolute error image distance weighting function (L^1)
- 4 different scales (Icosahedron subdivisions ranging from 162 to 10'242 vertices), corresponding to feature map scales (time points) 30, 22, 14 and 6, see figure 7.2

7.2.3 Classification Scheme

After pairwise registration of a subject to all other maps in the data-sets, we use the established spatial relationships to project both the feature maps and manual parcellations into the respective subject space. This data allows constructing an elaborate automatic parcellation scheme for the subject brain. Here, we use a Bayesian approach, as suggested in (Fischl et al., 2004), generalized in (Sabuncu et al., 2010) and applied in (Desikan et al., 2006; Yeo et al., 2010).

Bayesian approach

The central idea is to formulate the conditional probability of a parcellation P given a surface model S in terms of the conditional likelihood of S given P , and the prior probability of the parcellation:

$$p(P|S) \propto p(S|P)p(P) \quad (7.1)$$




Figure 7.2: *Spherical multiresolution scheme for faster registration.* **(a)–(d)** The full resolution cortical feature maps at different scales in the Beltrami scale-space, from coarse to fine. **(e)–(h)** The same feature maps, sub-sampled at the nodes of increasingly dense spherical meshes, obtained from icosahedron subdivisions of different depth. Reducing the sampling density allows for an important speedup of the registration task, without important loss of map detail. (Here we chose gray-scale colormap in order to avoid artifacts that are actually due to bad color interpolation in the rendering.)

In order to determine a maximum *a posteriori* (MAP) estimate of the parcellation P , we thus need to define both the likelihood model $p(S|P)$, where S is the observed geometry of the subject surface, and a prior model for the parcellation, $p(P)$.

Surface model

First, we look for a suitable surface model, i.e. the conditional PDF of the geometry given a parcellation. If we assume the noise at each vertex to be independent from all other vertices, then we may rewrite the global likelihood of the observed geometry as a product of the local probabilities at each single vertex:

$$p(S|P) = \prod_i p(S(x_i) | P(x_i)) \quad (7.2)$$

where $S(x)$ denotes the local observation of the surface geometry at vertex x . In particular, we consider a local feature vector composed of the equalized cortical feature maps of mean curvature, at different scales (here: 3, 13, and 23) as illustrated in figure 7.3. This is in contrast to the initial description of this classifier in (Fischl et al., 2004), where the geometry was described respectively by the average convexity (Fischl et al., 1999a) and mean curvature at a single scale only.

Now, at each vertex and for each class (label) in the atlas, the geometrical feature vector probability density function is modeled as a Gaussian distribution, of which the mean vector and covariance matrix are estimated directly from the available data:

$$p(S(x_i)|P(x_i) = c) = N(S(x_i)|\mu_c(x_i), \Sigma_c(x_i)) \quad (7.3)$$

Parcellation prior

Second, we need to define a prior on the parcellation P . Here, it is assumed only that the spatial distribution of labels can be approximated by a non-stationary Markov random field (MRF). In contrast, the initial model in (Fischl et al., 2004) is also anisotropic. Formally, this reduced the dependency of the local label from all other labels in the map to the first order neighbors $x_j \in N_1(x_i)$ only:

$$p(P(x_i)|P(x_j \neq x_i)) = p(P(x_i)|\{P(x_j)|x_j \in N_1(x_i)\}) \quad (7.4)$$

This restriction permits the probability of the entire parcellation to be rewritten in terms of neighborhood or clique potentials, using the Hammersley-Clifford theorem and by introducing Gibbs distributions (Fischl et al., 2004). Therefore:

$$p(P) = \prod_i p(P(x_i)|\{P(x_j)|x_j \in N_1(x_i)\}) \quad (7.5)$$

and, using Bayes' rule:

$$p(P) \propto \prod_i p(P(x_i)) p(\{P(x_j)|x_j \in N_1(x_i)\}|P(x_i)) \quad (7.6)$$

where $p(P(x_i))$ is directly estimated from the number of occurrences of a label at a given vertex in all available reference data sets.

To render this problem computationally tractable, we further assume that only the first order conditional dependence matters, i.e., the likelihood of a label given its neighbors can be expressed as the product of some pairwise probabilities. This leads to:

$$p(P) \propto \prod_i p(P(x_i)) \prod_{j \in N_1(x_i)} p(P(x_j)|P(x_i)) \quad (7.7)$$




Figure 7.3: *Multiple scales of cortical feature used as geometry descriptor.* The different Beltrami scales of the cortical feature map, that are used as feature vectors to describe the local geometry for automatic parcellation of the cerebral cortex.

Here, in contrast to (Fischl et al., 2004), we do not exploit the spatial relationship of the respective neighbors, and simply propose

$$p(P(x_j)|P(x_i)) \propto \exp(\delta(P(x_j), P(x_i)) \ln(\lambda)) = \lambda^{\delta(P(x_j), P(x_i))} \quad (7.8)$$

where δ denotes Kronecker delta, and $\lambda \geq 1$ is a constant that controls clique parcellation smoothness. As a result, the probability of label c at x_i will be proportional to λ^{f_c} , where f_c denotes the number of occurrences of the same label c among the neighbors $x_j \in N_1(x_i)$. Clearly, at $\lambda = 1$, the neighboring occurrences have no impact, whereas $\lambda \rightarrow \infty$ yields “winner-takes-it-all”.

MAP estimate

As in (Fischl et al., 2004), we employ the iterated conditional modes (ICM) algorithm (Besag, 1986) in order to make the MAP estimation of the proposed MRF computationally tractable. This involves initialization of the parcellation with the MAP estimate assuming $\lambda = 1$. Then, the parcellation is sequentially updated at each location by computing the label $P(x_i)$ that maximizes the conditional posterior probability $p(P(x_i)|\{P(x_j)|x_j \in N_1(x_i)\}, S(x_i))$:

$$\begin{aligned} P(x_i) &= \operatorname{argmax}_c p(P(x_i) = c | \{P(x_j)|x_j \in N_1(x_i)\}, S(x_i)) \\ &= p(S(x_i)|P(x_i) = c) p(P(x_i) = c) \prod_{j \in N_1(x_i)} p(P(x_j)|P(x_i) = c) \end{aligned} \quad (7.9)$$

This expression differs from the one given in (Fischl et al., 2004), where instead of $\prod_{j \in N_1(x_i)}$ one maximizes over $\sum_{j \in N_1(x_i)}$, the reason of which is currently under investigation. The impact of this difference on the outcome of the parcellation, however, is only minor.

It is to note, also, that in contrast to (Fischl et al., 2004), we do not perform any manual post-processing of the label maps.

7.2.4 Scripting and Process Monitoring

The data set contains 39 brains, of which we will arbitrarily consider the right hemisphere only. In order to achieve full-blown pairwise registration,

each subject has to be mapped to all 38 peers, and we have to perform 1'482 complete spherical registrations. Given the aforementioned registration setup and the current implementation in 64-bit MATLAB[®] R2009a, with Fedora 13 running on an Intel[®] Core[™] 2 at 2.40 *GHz* and with 8 *GB* of RAM installed, the registration of a single pair of cortical maps takes roughly one hour. Almost 95% of this time are spent on spherical interpolation, in particular for the actual warping of the moving image at each iteration. Considering that 2 jobs can run in parallel on the dual core computer, already the registration only is expected to take around 30 days of computation time!

In order to protect these lengthy computations against computer failures and data loss, we implemented a small stateless shell script daemon, which defines the different computational tasks for each subject or pair of subjects (mesh pre-processing, scale-space generation, registration, label and feature warping, classification, etc.) as different jobs. These jobs are continuously processed by launching the corresponding MATLAB script as soon as all required precursor data and free CPU are available. Upon OS failure, power outage or program crash, the computations could then easily be resumed with only little loss of data and computation time, and without important human intervention. Also, the number of simultaneous tasks can easily be scaled according to the number of CPU available on the respective computer, and generalization to clusters and grid is possible without important effort. Finally, the status of all the different jobs is automatically reported within a single HTML-file, which allows easy monitoring of the jobs and progress over the Web.

7.3 Results and Discussion

Having run the pairwise registration and subsequent leave-one-out automatic parcellation of the 39 right hemispheres in the data set, we will now provide and discuss the obtained results. Three main aspects are being evaluated: first, we want to get an idea of the success of the pairwise registration on its own. We do so by computing the pairwise overlap of region labels prior to and after registration. Second, we will try to characterize the label fusion scheme. In particular, we want to compare against naive majority voting, and we illustrate the impact of the MRF smoothing term. Finally, we analyze the parcellation performance of the overall scheme, including both registration and automatic label fusion, to get an idea of the impact of pairwise registration using the proposed FastGAF method on the outcome of the subsequent parcellation.

7.3.1 Measuring label consensus: Dice's similarity coefficient

In order to quantify the overlap, viz. consensus, between two label maps, projected on the same surface, we use the Dice's similarity coefficient (DSC)

Table 7.2: *Pairwise labelmap agreement prior to and after registration.* Accuracy is measured in “percent correct”, with nodes weighted by surface area, corresponding to DSC. Without registration versus using the proposed spherical FastGAF framework.

	Dice similarity coefficient (DSC) [%]			
	w/o registration		FastGAF	
	mean	median	mean	median
Labelmap agreement	61.0145	63.0689	62.3963	64.5976

(Dice, 1945). For any structure i , the respective DSC is computed as

$$S_i(l_1, l_2) = \frac{2 \int_{S^2} \delta(l_1, i) \delta(l_2, i) dA}{\int_{S^2} \delta(l_1, i) + \delta(l_2, i) dA}, \quad (7.10)$$

i.e., the ratio between twice the area of the sphere where both label maps share the label i , over the total area of label i on both maps. This measure ranges from 0 to 1, where 0 corresponds to completely disjoint labelization, and 1 to complete consensus w.r.t. the label i . The consolidated DSC for the entire set of labels can be constructed as a weighted average of the individual S_i :

$$S(l_1, l_2) = \frac{\sum_i S_i(l_1, l_2) \int_{S^2} \delta(l_1, i) + \delta(l_2, i) dA}{\sum_i \int_{S^2} \delta(l_1, i) + \delta(l_2, i) dA} = \frac{\int_{S^2} \delta(l_1, l_2) dA}{\int_{S^2} dA} \quad (7.11)$$

which is the portion of the spherical area where both label maps are in agreement.

7.3.2 Registration

In a first step, we want to assess the improvement of feature map alignment due to the pairwise registration. Therefore, we compare the labelmap agreement between pairs of subjects prior to and after registration. The key figures are given in table 7.2, and illustrated in more detail in the boxplots provided in figure 7.4.

In essence, we show that the improvement in feature map alignment due to pairwise registration may be small in terms of mean DSC improvement (increase from 61.01% to 62.40%) but *consistent and significant*.

A closer look at the subject-wise DSC distributions in the boxplots of figure 7.4 reveals that there are a number of subjects, whose brain seem to be consistently different from the rest of the data set. These brains (subject 5, 10, 14, 17, 23, 32 and 33) exhibit a clearly lower DSC than the sample average, both before and after registration. Unsurprisingly, the worst subject (32), will also perform weakest in the automatic parcellation, as we will see soon below.




Figure 7.4: *DSC of pairwise labelmap agreement, with and without registration.* The subjects are sorted by id in data set, and are numbered below. **(a)** In each group, the gray box corresponds to the result without registration, and the black box shows the DSC after FastGAF registration. **(b)** Pairwise DSC difference shows small but significant (pairwise t-test, $\alpha = 0.01$) improvement due to registration for all structures.

7.3.3 Label fusion scheme

Before we can look at the outcome of the parcellation cross-validation, we want to characterize and discuss the proposed label fusion scheme. Therefore we apply the proposed automatic parcellation scheme to one subject, given the labels of all others. In particular, we want to show the difference between a simple majority voting, the proposed Bayesian scheme without spatial smoothing ($\lambda = 1$), and the full scheme with $\lambda > 1$. Note that majority voting (MV) is equivalent to uniform conditional probabilities $p(P|S)$ (no feature dependency) and $p(P(x_i)|P(x_j))$ (no smoothing) in the Bayesian model.

The different obtained parcellations are illustrated in figure 7.5. Note that majority voting intrinsically produces rather smooth maps, while the Bayesian approach without smoothing yields label maps with lots of small “perturbations”. These perturbations, however, are efficiently reduced by the spatial smoothness prior.

Both schemes, MV and Bayesian, produce labelmaps that are similarly different from the actual manual parcellation. The feature-based surface model $p(P|S)$ can improve the parcellation to some extent, locally. But the labelmaps are consistently shifted off the true map, nonetheless.

The spatial smoothing based on different choices for the parameter λ is illustrated in figure 7.6. As expected, the ICM algorithm converges rapidly, and after around 12 iterations, the labelization only oscillates slightly. The impact of the actual choice of $\lambda > 1$ is rather minor. Choosing $\lambda = 100$ clearly produces smoother maps than $\lambda = 1.25$, but the difference is not dramatic. Some small label-islands remain even when choosing λ that large. This is, why (Fischl et al., 2004) performs manual correction of the automatic parcellations.

7.3.4 Parcellation Cross-validation

The overall results of the automatic registration-parcellation leave-one-out cross-validation scheme are reported in table 7.3. In particular, we provide the mean and median DSC of the parcellation with and without prior registration by the proposed FastGAF algorithm. For comparison, we also show the corresponding numbers for the simple majority voting and the unsmoothed Bayesian model.

To illustrate the variability of quality of the obtained parcellations, we also give the numbers for the worst and best subjects, respectively. As expected, the worst subject is sub32, which already had worst pairwise label agreement before label fusion. At the other end of the spectrum, the best parcellated brain, sub09, already showed best pairwise label agreement.

The numbers show, that on average the smoothed Bayesian classifier performs best. Using the proposed Bayesian classifier, roughly 75% of the spherical surface are correctly labeled. Further, we clearly see that FastGAF registration consistently improves this figure by ca. 1.3%.

To get a better intuition of where mis-parcellation typically occurs, we

project the manual and automatic parcellation of an averagely performing subject on to the pial and inflated surface, see figure 7.7. Further, we highlight the mis-parcellated regions on the ground-truth labelmap, as well as on the actual feature map. This clearly reveals, that most of the error is committed at the boundaries only of the different gyral regions. A look on the feature map suggests, that this happens in particular at regions that are not particularly well-defined in terms of surface geometry, e.g., within highly variable sulcal regions, or around the artificially dissected corpus callosum.

The same illustrations are produced for the extreme subjects 09 and 32, in figure 7.8. In the worst case, the mis-parcellation is clearly due to an outlier cortex: The overall shape as well as the cortical folding pattern are clearly particular with respect to the rest of the brains in the data-set. The parcellation of subject 09, however, is almost perfect, with almost 90% correctly labeled surface. The remaining mis-parcellation is extremely concentrated at the region boundaries, and at this point become almost arbitrary with respect to manual delineation.

Finally, we want to shed some light on the parcellation quality of the 35 individual gyral regions. Therefore, the structure-wise DSC was computed for all parcellations, and the obtained data are shown in the boxplots of figure 7.9. This illustration clearly shows, that the parcellation quality improves with bigger structures, which is a natural consequence of parcellation error being concentrated around boundaries. Bigger structures not only are better parcellated, the DSC variability among test brains is also drastically smaller. Comparing the DSC values with and without registration, we can identify again a systematic improvement thanks to registration. In the boxplot of pairwise differences, we find that most structures are significantly improved, whereas only few structures suffer by registration, and all of the deterioration is non-significant. We identify a few structures that perform particularly bad: This includes the frontal pole (33), the temporal pole (34), the corpus callosum (5), the entorhinal (7) and the pericalcarine fissure (22). All of these structures are either ill-defined (e.g., only in relation to their surrounding structures and not by specific geometry), exhibit important natural intersubject variability, or suffer from feature map design artifacts, such as the dissected corpus callosum.




Figure 7.5: *Visual evaluation of different label fusion schemes. (a)* Underlying cortical feature map. **(b)** Manual delineated “ground truth” parcellation. **(c)** Labels fused using pure majority voting. **(d)** Result of proposed Bayesian classifier, without spatial smoothing ($\lambda = 1$). **(e)** Spatial smoothing with $\lambda = 100$ in the spatial prior for multi-atlas fusion.




Figure 7.6: *Effect of spatial smoothing for varying λ .* We show the estimated parcellations after ICM iterations 2, 4, 6, and 20, for different $\lambda > 1$ ranging from 1.25 to 100. ICM converges quickly, and beyond ca. 13 iterations, the labelmap has converged to just tiny oscillations. Already small λ successfully remove noise from the labelmaps, and larger values just straighten the borders somewhat more. A few isolated label-patches remain, however, that should be removed manually.

Table 7.3: *Parcellation accuracy for different atlas fusion schemes.* Accuracy is measured in “percent correct” (overall DSC), with nodes weighted by surface area. Without registration versus using the proposed spherical FastGAF framework. In addition, we give the accuracy of the worst (sub32) and the best (sub09) individual parcellation.

Fusion Method	Dice similarity coefficient (DSC) [%]			
	w/o registration		FastGAF	
	mean	median	mean	median
Majority voting	71.0723	74.9535	72.4938	76.5458
Bayes $\lambda = 1$	73.2996	74.7249	74.6315	75.9541
Bayes $\lambda = 100$	73.9480	75.4433	75.2338	76.7319
	sub32	sub09	sub32	sub09
Majority voting	45.2556	87.8660	46.6625	88.2811
Bayes $\lambda = 1$	49.7565	86.1914	51.5298	86.4088
Bayes $\lambda = 100$	49.9253	86.9561	51.8386	87.2013




Figure 7.7: *Manual and automatic parcellation for an average subject.* First and second row show manual (“ground truth”) and automatic parcellation results on the pial and inflated cerebro-cortical surface, respectively, from lateral and medial view. Most mis-parcellation occurs at boundaries. Due to the absence of any manual post-processing of the automatic parcellation results, a few, clearly identifiable and easily correctable isolated mis-parcellations occur.




Figure 7.8: *Manual vs. automatic parcellation for best and worst subject.* Cf. caption 7.7. As is clearly visible, the worst parcellated subject 09 exhibits a particular brain shape and an atypical folding pattern, difficult to parcellate. The parcellation of the best subject, 32, is close to perfect: The remaining mis-parcellation is very marginal, and close to arbitrary.




Figure 7.9: Dice similarity coefficients for individual structures, with and without registration. The structures are sorted by increasing mean area, and are numbered below. **(a)** In each group, the gray box corresponds to the result without registration, and the black box shows the DSC after FastGAF registration. The solid red curve indicates the relative cortical area of a structure, and the dotted line shows cumulative area. **(b)** Pairwise DSC difference shows significant (pairwise t-test, $\alpha = 0.01$) improvement due to registration for most structures (black), gray boxes indicate non-significant improvement, red highlights non-significant deterioration. As mismatches mainly occur close to structure borders, small structures automatically exhibit weaker performance and higher variance. It is to note, that some of the worst performing structures are ill-defined on the feature maps that we use, such as the frontal pole (33), the temporal pole (34), the corpus callosum (5), the entorhinal (7) or the pericalcarine fissure (22).

7.4 Conclusions and Outlook

In this chapter, we have proposed to build an automatic cortical parcellation scheme, based on pair-wise registration of a subject map against all maps in the atlas. In a second step, the registered feature and label-maps in the atlas are used within a Bayesian classifier, of which the MAP estimate is solved for using ICM.

Based on a data-set of 39 subject brain surfaces, each of which composed of feature maps and manual delineations into 35 gyral regions on each hemisphere, we were able to test the proposed parcellation scheme in a leave-one-out cross-validation.

Using the results of this cross-validation, we were able to show:

1. pairwise registration of a subject brain against the atlas brains improves the pairwise labelmap agreement consistently and significantly.
2. the proposed classifier outperforms simple majority voting,
3. estimating the MAP solution using ICM is efficient and converges fast,
4. most importantly, registering the spherical feature maps using the proposed FastGAF algorithm improves the parcellation outcome consistently and significantly.

7.4.1 Comparison to the state-of-the-art

Although the improvement thanks to spherical registration is consistent and significant, we have to admit that the impact is surprisingly small. Indeed, interestingly, in (Yeo et al., 2010) a high parcellation quality of around 88% is claimed as *average* performance for the same data set in a similar cross-validation scheme, both using the Freesurfer registration scheme as well as spherical Demons. Nonetheless, their figure is extremely difficult to interpret against our algorithm, since several poorly performing structures are excluded from their measure, and it is unclear how the original classifier borrowed from (Fischl et al., 2004) was exactly built and improved. Moreover, they allow for manual post-processing of the automatically established labelmaps, which we explicitly don't.

7.4.2 Potential improvements of the registration scheme

As already pointed out in the conclusions of the spherical FastGAF algorithm at the end of chapter 6, rigid rotations of the whole sphere do not come free of cost with the proposed Beltrami regularizer. Therefore, one should definitely consider pre-registering the feature-map pairs using rigid deformation (rotation only), or with low-dimensional parametric deformation.

Also, we want to clearly state that the chosen set of registration parameters was only coarsely determined based on a few trial runs on a single pair

of images, in order to protect against overfitting. To some extent, and within limits of computational complexity, the parameters could be further optimized, by maximizing the performance of the parcellation scheme on the right hemispheres, and subsequent testing on the left hemispheres, for example.

Another weakness is the brute-force pairwise registration of all brains in the data-set. As mentioned earlier, here this involves 1'482 registrations. An alternative approach consists in the iterative construction of an average or statistical atlas brain, to which subjects are subsequently registered. The strength of this approach is also its main weakness: averaging removes individual modulations and subjects are registered to the common base-pattern. For the same reason, individual features of the subject may get lost and not find a suitable counterpart. But the very idea of the Bayesian approach for label-fusion is precisely to match a subject brain to all individuals in the atlas set, and to weight the label-voting based on how good two labelmaps can be matched, thus intrinsically allowing for multiple brain models in the atlas.

7.4.3 Potential improvements of the label-fusion

Although not the focus of this thesis, we would like to comment on a few possible improvements of the classification scheme.

Spatial prior

Intuitively, a lot of performance difference between the proposed method described in this chapter, and the reference results from literature, is due to the classifier. In particular, we admit that the spatial prior is probably overly simple. Since the labels are based on gyral regions, transitions between different labels are expected to be more likely in sulcal regions. This prior on the transition probability based on feature value is currently not being accounted for.

ICM for MAP estimate

The important advantage of ICM over other MRF solvers, such as simulated annealing (Geman and Geman, 1984), is its extremely simple structure and fast convergence. However, it has been shown that ICM corresponds to instant freezing from a simulated annealing point of view (Besag, 1986). ICM therefore has tendency to fix the underlying label-map too quickly. It would be interesting to see, whether slower cooling could provide slower fixation and smoother parcellations. In particular, the ability of simulated annealing to escape from local extrema might help reducing the occurrence of isolated label patches, that otherwise have to be removed manually.

Label definition

Finally, since the region boundaries of the current 35 gyral regions are mainly located at sulcal cavities of the feature maps, there is not much contrast between the geometry vectors on either side. Both the registration is imprecise, and the surface model $p(P|S)$ is rather unspecific at exactly these locations. This obviously does not help in getting sharp and precise label boundaries. If, in contrast, one had chosen separate regions for sulci and gyri, label boundaries would instead occur at the edges of the feature map, with much better feature discrimination between both sides of the label boundary.

Diffeomorphic Deformations

8

騎驢找馬

“Riding a mule while looking for a horse.”

Chinese proverb.

IN BIOMEDICAL IMAGE REGISTRATION, diffeomorphic deformation fields are widely accepted as a powerful class of deformations (Christensen et al., 1996; Joshi and Miller, 2000; Beg et al., 2005; Rueckert et al., 2006; Ashburner, 2007; Vercauteren et al., 2009; Yeo et al., 2008). By definition, a diffeomorphism is an invertible deformation that maps one image to another, such that both the deformation and its inverse are smooth (differentiable). A diffeomorphic deformation preserves the topology of the image and prevents it from folding, even at large deformations. Moreover, invertibility of the deformation field is an essential property when it comes to rendering the registration process inverse consistent or even symmetric (Rogelj and Kovacic, 2006). On the other hand, inverse consistent image registration tries to optimize the forward and inverse deformation fields simultaneously, to reduce the impact of the choice of subject and target image roles. Symmetric methods enforce equivalence between registration directions by construction.

In this chapter, we sketch how the GAF framework can be extended to diffeomorphic and inverse-consistent deformation model. The extended framework will be illustrated on simple toy examples.

8.1 Background

Different approaches exist, that try to obtain a diffeomorphic deformation field. Models, which only penalize non-positive determinants of the transform Jacobians, may well guarantee deformations to be free of microfolds locally, but they have no control, however, over macroscopic folding and are not necessarily diffeomorphic. Fluid-like deformation models are shown to prevent folding completely, but require complex computations. The use of exponential maps results in truly diffeomorphic deformations. In parametric image registration, diffeomorphism is most of the time only implicitly achieved, through constraints on model parameters or specific regularization terms (Narayanan et al., 2005).

8.1.1 Exponential maps

The most intuitive way of incorporating exponential maps in diffeomorphic image registration is inspired by the works of Joshi and Miller, Beg and Ashburner (Joshi and Miller, 2000; Beg et al., 2005; Ashburner, 2007). While Joshi and Miller optimize a regularized, time-varying velocity field, which transports points in space, Beg and Ashburner use a single, static velocity field. Let \vec{u} describe the static velocity field, and the transform be obtained as its exponential map, e.g., $T\vec{x} = \exp(\vec{u})$. The deformation field is directly guaranteed to be diffeomorphic and invertible – the inverse is obtained as the exponential map of the negative velocity field. But in return, a local change of the velocity field does not only affect the displacement of a single pixel, but all pixels of which the trajectories pass through the support of that change. Consequently, the expression of the metric flow, i.e., the partial derivatives of the image metric with respect to the velocity field, gets cumbersome (Ashburner, 2007). Further, it is suspected that the class of diffeomorphisms that can actually be generated by a static velocity field is quite limited.

8.1.2 Diffeomorphic Demons

Very recently, Vercauteren *et al.* introduced exponential map diffeomorphisms in the Demons framework (Vercauteren et al., 2009). There, at each iteration one looks for an infinitesimally small update $d\vec{s}$ to $\vec{s} = 0$, that is applied through composition of its exponential map with the existing diffeomorphic deformation:

$$T_1 : T_1\vec{x} = (\vec{x} + \vec{u}) \circ \exp(\vec{s}), \quad (8.1)$$

where composition of two diffeomorphisms produces another diffeomorphism, within the limits of discrete numerical computations. At each iteration, the deformation field \vec{u} is updated as:

$$\begin{cases} \vec{s} \leftarrow \vec{0} + dt \cdot d\vec{s} = dt \cdot d\vec{s} \\ \vec{x} + \vec{u} \leftarrow (\vec{x} + \vec{u}(\vec{x})) \circ \exp(\vec{s}), \end{cases} \quad (8.2)$$

where dt denotes the time step. This scheme has the advantage of keeping the impact of the update very local. Also, successive composition of diffeomorphisms is commonly believed to generate a wider space of possible deformations, than exponential maps of static velocity fields.

8.1.3 Deformation model in GAF

In the GAF framework, the deformation field is embedded as a mapping between the n -dimensional image domain and a m -dimensional space, where $m > n$. This is achieved by letting the components of the deformation field become additional dimensions of the embedding space. The embedded manifold then evolves towards a weighted minimal surface, while being attracted by a deformation field that brings the two images into registration. Further, we dealt with a deformation field $\vec{u}(\vec{x}) : \vec{x} \in \Omega \subset \mathbb{R}^n \rightarrow (u_1, \dots, u_p) \in \mathbb{R}^p$ that describes the spatial displacement along $p \leq n$ dimensions of an n -dimensional image of support Ω by addition:

$$T_0 : T_0\vec{x} = \vec{x} + \vec{u}(\vec{x}), \quad (8.3)$$

where addition is implicitly understood only along the $p \leq n$ dimensions of the image that are actually deformed.

Once corresponding fixed and moving image locations can be mapped, the matching quality can be quantified using one of several distance metrics. As an example for monomodal image registration subject to additive Gaussian noise let us consider the squared error metric (Toga, 1999):

$$f^{T_0}(\vec{x}, \vec{u}) = (\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x}))^2 = (\mathcal{M}(\vec{x} + \vec{u}(\vec{x})) - \mathcal{F}(\vec{x}))^2, \quad (8.4)$$

where \mathcal{F} and \mathcal{M} refer to the fix and moving images respectively.

Here, we propose to integrate the *demons-like* deformation scheme T_1 into the geodesic active fields model. In a first step, we simply adopt the new deformation model and compute the corresponding evolution equations of the GAF energy. In a second step, we take profit of the availability at low cost of an estimate of the inverse deformation field, to extend the weighting function to be symmetric. This means that both, the matching of the fixed image with the resampled moving image, and the matching of the inversely warped fixed image with the moving image, are considered in the weighting function, thus removing some asymmetry in the registration process.

The structure of the remainder of this chapter is as follows. First, section 8.2 presents the integration of exponential mapping into the registration scheme to provide diffeomorphic deformation fields. In section 8.3 we extend the weighting function by inclusion of the inverse deformation data term to make the registration inverse consistent. Finally, we show some illustrative, preliminary results obtained with our diffeomorphic geodesic active fields framework in section 8.4 and we discuss our model in section 8.5.

8.2 Diffeomorphic Deformation Model

Before we can tackle the integration of diffeomorphisms into the GAF deformation model, we shall briefly recall how diffeomorphisms are related to exponential maps.

8.2.1 Exponential map

Diffeomorphisms do not form a vector space with respect to addition: addition of diffeomorphisms does not necessarily produce a diffeomorphism, and the results may not be invertible. Instead, diffeomorphisms can be smoothly composed and inverted (Vercauteren et al., 2009). Formally, diffeomorphic deformation fields on the image manifold form a Lie group \mathcal{G} . The group element is a smooth manifold, and the associated operation is composition. The neutral element is given by identity I , and an inverse of the diffeomorphism exists by definition.

To this Lie group, a corresponding Lie algebra \mathfrak{g} can be associated. The underlying vector space of this Lie algebra is the tangent space of \mathcal{G} around the neutral element Id . The Lie algebra and the Lie group are linked through the exponential map (Vercauteren et al., 2009):

$$\exp : \mathfrak{g} \rightarrow \mathcal{G}, \quad \exp(X) = \gamma(1), \quad (8.5)$$

where $\gamma : \mathbb{R} \rightarrow \mathcal{G}$ is the unique one-parameter subgroup of \mathcal{G} whose tangent vector at identity is given by X . The exponential map can be constructed as the integral curve

$$\frac{d}{dt}\gamma(t) = X \circ \gamma(t), \quad (8.6)$$

where X is a stationary vector field, and with initial condition $\gamma(0) = I$. Physically interpreted, the diffeomorphic deformation is the movement of particles governed by a stationary velocity field X , after unit time. The exponential map is a smooth mapping from a neighborhood of 0 in \mathfrak{g} to a neighborhood of I in \mathcal{G} . Note that being a one-parameter subgroup implies $\gamma(2t) = \gamma(t) \circ \gamma(t)$, and more generally $\exp(tX) = \gamma(t)$ (Bossa et al., 2008). In particular, this yields

$$\exp(X) = \exp(2^{-N}X)^{2^N}, \quad (8.7)$$

which leads to the scaling and squaring algorithm for fast vector field exponential computation (Arsigny et al., 2006; Bossa et al., 2008):

1. **Scaling:** Divide X by a factor 2^N , so that $2^{-N}X$ is sufficiently close to zero (for example $N = 8$)
2. **Exponentiation:** Perform an explicit first order integration using the Taylor expansion of the exponential, $\exp(2^{-N}X) \approx I + 2^{-N}X$.
3. **Squaring:** N recursive squarings, $\exp(2^{-n+1}X) = \exp(2^{-n}X) \circ \exp(2^{-n}X)$.




Figure 8.1: Convergence of exponential maps with increasing scaling and squaring depths N . The blue arrows indicate the velocity field. The thick red lines indicate the individual trajectories of a set of vertices during squaring, while the thin red lines show the front propagation. $N = 0$ corresponds to the basic, non-diffeomorphic model and can introduce heavy folding. Indeed, the whole set of points is mangled onto a very small region. For N increasing, the exponential map approximation becomes rapidly more accurate.

Composition of sampled deformations is obtained using recursive interpolation.

Fig. 8.1 shows the convergence towards diffeomorphic maps as a function of the scaling and squaring depth N . The difference between warping a mesh with a vector field and its exponential map is illustrated in Fig. 8.2.

8.2.2 Diffeomorphic weighting functions

First, we want to identify the update $d\vec{s}$ in (8.2) with the minimizing flow $d\vec{u}$ corresponding to the formulation in (4.3). We have:

$$\begin{cases} \exp(\vec{s})|_{\vec{s}=0} = I, \\ \frac{\partial}{\partial \vec{s}}|_{\vec{s}=0} \exp(\vec{s}) = I. \end{cases} \quad (8.8)$$

Therefore, we obtain the following equality:

$$d\vec{u} = J^T \cdot d\vec{s}. \quad (8.9)$$

where J^T refers to the transpose of the transform Jacobian.

Given the compositive deformation model T_1 , the squared error weighting function writes:

$$f^{T_1}(\vec{x}, \vec{u}) = (\mathcal{M}(\vec{x} + \vec{u}) \circ \exp(\vec{s}) - \mathcal{F}(\vec{x}))^2. \quad (8.10)$$




Figure 8.2: *Warping a mesh with a vector field and its exponential. (a)* A square mesh warped by a vector field. *(b)* Warping by its exponential map. The velocity field is considered zero outside the square, therefore the mesh cannot flow beyond that border. The exponential map produces a diffeomorphic deformation of the mesh, whereas direct deformation by the vector field introduces folding.

The spatial gradient of the weighting function, $\nabla f^{(1)T_1}$, is most easily obtained numerically through finite differences of the available samples. The partial derivatives with respect to the update are:

$$f_{\vec{s}}^{T_1}(\vec{x}, \vec{u})|_{\vec{s}=0} = 2(\mathcal{M}(\vec{x} + \vec{u}) \circ \exp(\vec{s}) - \mathcal{F}(\vec{x})) \quad (8.11)$$

$$\cdot J^T \nabla \mathcal{M}(\vec{x} + \vec{u}) \cdot \partial_{\vec{s}} \exp(\vec{s}), \quad (8.12)$$

and, more interestingly, we get the metric flow with respect to the deformation field \vec{u} :

$$f_{\vec{u}}^{T_1} = 2(\mathcal{M}(T_0 \vec{x}) - \mathcal{F}(\vec{x})) \cdot \nabla \mathcal{M}(T_0 \vec{x}), \quad (8.13)$$

which is exactly equal to the non-diffeomorphic, additive flow (4.25).

The diffeomorphic counterparts of, e.g., the absolute error and local joint entropy weighting function are calculated analogously.

From an implementational point of view, there is only one difference between the additive, non-diffeomorphic T_0 , and the compositive diffeomorphic T_1 deformation model: in the latter, the update is obtained as composition with the exponential map of the affinely transformed flow (4.3):

$$\vec{u} \leftarrow (\vec{x} + \vec{u}(\vec{x})) \circ \exp(dt \cdot J^{-T} d\vec{u}) - \vec{x}. \quad (8.14)$$

8.2.3 Inversibility

The main motivation for diffeomorphic deformations is the existence of a smooth inverse. In this section, we shall compute the inverse deformation along with the regular update of \vec{u} .

At the end of every iteration $k \geq 0$, the forward transform is composed with the exponential map of the update. It is therefore also natural to compose the inverse of the update with the current inverse transform, written $(\vec{x} +$

$\vec{p}^k(\vec{x})$). The inverse of the update, in turn, is easily obtained by computing the exponential map of the negative velocity field, and we get:

$$\begin{cases} \vec{x} + \vec{u}^{k+1} \leftarrow (\vec{x} + \vec{u}^k(\vec{x})) \circ \exp(\vec{s}), \\ \vec{x} + \vec{p}^{k+1} \leftarrow \exp(-\vec{s}) \circ (\vec{x} + \vec{p}^k(\vec{x})). \end{cases} \quad (8.15)$$

The proof using mathematical induction is straightforward. At any iteration $k + 1$, invertibility requires:

$$\begin{cases} (\vec{x} + \vec{u}^k(\vec{x})) \circ \exp(\vec{s}) \circ \exp(-\vec{s}) \circ (\vec{x} + \vec{p}^k(\vec{x})) = I, \\ \exp(-\vec{s}) \circ (\vec{x} + \vec{p}^k(\vec{x})) \circ (\vec{x} + \vec{u}^k(\vec{x})) \circ \exp(\vec{s}) = I, \end{cases} \quad (8.16)$$

which is obviously true, iff $(\vec{x} + \vec{u}^k)^{-1} = (\vec{x} + \vec{p}^k)$. Indeed, at the first iteration, $\vec{u}^0 = \vec{0} = \vec{p}^0$, and thus $(\vec{x} + \vec{u}^0)^{-1} = I = (\vec{x} + \vec{p}^0)$.

8.3 Inverse Consistent Registration Model

An inverse consistent registration model measures both the quality of the forward and inverse deformation field. This has the advantage of removing some asymmetry between the moving and fixed image by introducing further mirror terms. Let us again consider the squared error weighting function, now involving the inverse update $\exp(-\vec{s})$ and the current estimate of the inverse deformation field \vec{p} :

$$\begin{aligned} f^{T_1*}(\vec{x}, \vec{u}) &= \frac{1}{2} \left((\mathcal{M}(\vec{x} + \vec{u}) \circ \exp(\vec{s}) - \mathcal{F}(\vec{x}))^2 \right. \\ &\quad \left. + (\mathcal{M}(\vec{x}) - \mathcal{F}(\exp(-\vec{s})) \circ (\vec{x} + \vec{p}))^2 \right). \end{aligned} \quad (8.17)$$

The partial derivatives with respect to the update \vec{s} are:

$$\begin{aligned} f_{\vec{s}}^{T_1*}(\vec{x}, \vec{u})|_{\vec{s}=0} &= (\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x})) \cdot J^T \nabla \mathcal{M}(T_0\vec{x}) \\ &\quad + (\mathcal{M}(\vec{x}) - \mathcal{F}(T_0^{-1}\vec{x})) \cdot \nabla \mathcal{F}(T_0^{-1}\vec{x}). \end{aligned} \quad (8.18)$$

Substituting (8.9) into this, the partial derivatives of the weighting function with respect to the deformation field \vec{u} , finally, write:

$$f_{\vec{u}}^{T_1*} = (\mathcal{M}(T_0\vec{x}) - \mathcal{F}(\vec{x})) \cdot \nabla \mathcal{M}(T_0\vec{x}) \quad (8.19)$$

$$+ (\mathcal{M}(\vec{x}) - \mathcal{F}(T_0^{-1}\vec{x})) \cdot J^{-T} \nabla \mathcal{F}(T_0^{-1}\vec{x}). \quad (8.20)$$

8.4 Results

We have implemented the diffeomorphic and inverse consistent geodesic active fields and ran them on a toy test problem. All single-scale problems have been solved in a pyramidal multi-resolution approach (Modersitzki, 2004). Images are low-pass filtered and diadically downsampled in several stages, to

allow fast coarse-to-fine registration. The flow is integrated using a simple forward Euler scheme. As for all forward schemes, the step length, and thus the speed of the registration, is heavily limited by the stability of the integration. The implementation was done using Matlab®(R2009a) on a standard 2.4GHz Intel®Core™2 Duo desktop machine, running a 64bit Fedora Core 11.

A nice illustrative example is the registration of a star to a disc. These two binary images have been created at resolution 350×350 . For comparison, both the additive squared error weighting function (4.24) and the diffeomorphic squared error weighting function (8.10) are applied. The setup, the warped images and the resulting deformation fields are all shown in figure 8.3. The warped images of both direct and diffeomorphic deformation field are of comparable quality, that could be further improved by registering at higher resolution deformation fields. For the example we provide, the main difference between diffeomorphic and non-diffeomorphic registration is the inverse deformation field that is automatically obtained within the diffeomorphic setup. Indeed, we show that composition of the forward transform T_1 with its co-evolving inverse T_1^{-1} produces a map very close to identity, as expected.

8.5 Discussion and Conclusion

In this chapter, we have presented a sketch of how diffeomorphic deformations can be used within the geodesic active fields framework. By means of exponential maps we have defined deformations that are diffeomorphic, i.e., that keep the warping free of folds. Indeed, at each iteration, the deformation field is updated through composition with a small diffeomorphic field, obtained as the exponential map of the computed flow, keeping the whole deformation field diffeomorphic. Simultaneously, the exponential map of the reverse flow, i.e., the inverse update, is composed with the current estimate of the inverse deformation to update the inverse diffeomorphism.

In image registration, existence and computation of the inverse deformation is as important as obtaining the forward transform, because this allows mapping points between fixed and moving image in both directions. We have shown within the limits of precision of numerical computations, that composition of the forward and inverse diffeomorphism actually matches identity.

In a second step, the incorporation of the inverse deformation into the weighting function has enabled us to symmetrize the data term. This greatly reduces the impact of the role attribution between fixed and moving image, as the image mismatch is optimized in both directions simultaneously. Indeed, in the non-symmetric setting, only the edges of the moving image result in deformation forces, whereas in the symmetric weighting function both images contribute equally.

Despite this symmetric data term, the GAF model does not provide complete symmetry, as the GAF functional still measures the weighted surface of just the forward deformation field, only. Had we used the LDDMM version




Figure 8.3: *Diffeomorphic registration of a star to a disc.* **(a)–(c)** The fixed star and the moving disc, and their difference. **(d)–(f)** The computed non-diffeomorphic deformation field, T_0 , the warped disc and the remaining intensity difference. **(g)–(i)** Deformation field T_1 , warped disc and remaining intensity difference in the diffeomorphic setting. **(j)** Inverse transformation T_1^{-1} in the diffeomorphic setup. **(k)** Composition of the forward and backward transform is very close to identity: $T_1 \circ T_1^{-1} \approx I$. **(l)** The star warped with the inverse transform maps close to the disc. In this case, both the non-diffeomorphic and diffeomorphic deformation are free of folds. The precision of the registration, i.e. the remaining shape differences, are mainly due to the limited resolution we chose for the deformation field.




Figure 8.4: *Additive and diffeomorphic deformation on the sphere.* Deformation field for **(a)** $N = 0$ (basic additive model) and **(b)** $N = 10$ (accurate diffeomorphic model). While the basic model involves a lot of crossing lines, corresponding to important folding, the diffeomorphic model produces smooth stream-lines as point trajectories.

of diffeomorphic deformations, i.e., if the deformation field was described and regularized by its logarithm, the static velocity field \vec{u} , then complete symmetry could be obtained at the prize of much more complex flow computations. Indeed, the inverse of such a deformation field is described by the negative of its logarithm, which obviously has equal hyperarea.

As future work, it will now be interesting to minimize the inverse-consistent, symmetric diffeomorphic deformation model within the fast scheme, and, more importantly, on the sphere. The difference between standard and diffeomorphic deformation fields on the sphere is illustrated in figure 8.4. However, since computing the exponential map of a velocity field involves several recursive compositions through interpolation (deformation field warping), the speed of spherical data interpolation will quickly become an important bottleneck.

Conclusions

9

“Success is not final, failure is not fatal: it is the courage to continue that counts.”

Winston Churchill.

THE GOAL OF THIS THESIS was to develop a method to register human cerebro-cortical feature maps on the sphere. First, prior to registration, the spherical maps are anisotropically low-pass filtered using a Beltrami scale-space, so as to allow for morphogenically meaningful hierarchical registration. Then, based on a few concepts from Gestalt theory, and the central hypothesis that Beltrami-regularization of the inverse problem would allow reproducing some of the grouping phenomena observed in human vision, we have developed a novel geometric framework for image registration in general, and on the sphere in particular. In order to speed up the computation, a fast scheme has been implemented. Finally, the spherical registration has been employed and proven useful in a leave-one-out cross-validation study of automatic cortical feature map parcellation, based on a data set consisting of 39 pairs of left and right hemispheres, each of which being manually parcellated into 35 gyral regions.

9.1 Summary and Main Contributions

The main contributions of this thesis can be summarized as follows:

9.1.1 Scale-space of cortical feature maps

Given the highly convoluted structure of the cerebro-cortical sheet, it is today commonly accepted to work with spherical representations of the extracted actual cortical surface. Cortical folding patterns vary a lot between subjects. The degree of inter-subject variability, however, changes hierarchically with the “importance” of the cortical fold in question. It was shown in different studies that some structures are particularly stable between subjects, while smaller structures exhibit more variability. It was also suggested, that the degree of variability is related to brain morphogenesis, and that a scale-space of cortical mean curvature well reproduces the involved hierarchy.

In this thesis, we have therefore chosen to pre-process the cortical feature maps into an anisotropic scale-space, based on the Beltrami framework. We propose to simply include the data of the map, associated to a particular node of the spherical triangulated mesh, as additional coordinate of the embedding space. Mesh regularity is then measured using the Beltrami energy, of which the minimizer is efficiently computed using mean curvature vector estimators, well-known from computer graphics. Our results indicate that the simplification in coarser scales of the generated scale-space are appropriate simplifications of the initial feature map. The pattern of sulci and gyri is successfully simplified from the detailed initial map into coarser and more abstract sketches, while keeping the important main structures well-conserved.

9.1.2 Geodesic Active Fields

The proposed Geodesic Active Fields (GAF) framework for image registration is built around the Beltrami framework. We propose to use an embedding of the unknown deformation field and the associated Beltrami energy as regularizer in the inverse problem of image registration. Despite the increased computational complexity, geometric regularization has several appealing advantages over standard Gaussian regularization. In particular, the flow-driven regularization increases separation between different Gestalts in the observed images, i.e., regularization is reduced across Gestalt boundaries, defined by sharp transitions in deformation space.

In addition, we replace the traditional additive coupling between data-term (the image mismatch) and regularity term (Beltrami energy) by multiplication. Indeed, we use the local image-mismatch as a weighting function in the Beltrami energy, and the registration problem is actually redefined as a weighted minimal surface problem. Again, this choice makes the optimization more complex, but involves some very interesting properties. First, the resulting weighted Beltrami energy functional is invariant under re-parametrization, i.e., choosing one or another parametrization of the images (Cartesian, polar, ...) does not change the registration problem. Second, thanks to the generality of the Beltrami framework, the registration method can be applied to images defined on any Riemannian manifold, and is not restricted to the flat

Euclidean domain. In particular, we show examples of patches on curved surfaces and multiscale images. Finally, the multiplicative coupling between data and regularization involves an automatic, intrinsic modulation of the regularization strength depending upon the local quality of alignment, and/or noise level. Indeed, in the initial phase of registration, the data-term will generally perform bad, and regularity is required to be stronger. However, later in the optimization, images are better aligned, regularization becomes less strict and local fine-tuning may start. The same observation holds for spatially varying noise level. Low-noise regions can be aligned in more detail compared to regions that are more severely affected – or regions where a one-to-one correspondence cannot be established due to “topological changes”, such as lesions.

9.1.3 Fast Scheme

With the FastGAF scheme, we have shown an efficient optimization scheme for the weighted Beltrami approach, that can outperform current state-of-the-art approaches both in terms of registration quality, as well as computation time. The fast scheme relies on a splitting approach, where data-term (the weighting function) and regularity term (the “Beltrami term”) are optimized over separate deformation fields, that are, however, constrained to be equal. This constraint is addressed using an Augmented Lagrangian method, and the resulting problem can be optimized efficiently in three different sub-tasks. Currently, we show some promising solutions for these respective sub-problems: A fixed-point scheme based on semi-implicit gradient descent for the data-term, and approximate inverse based on a few Jacobi-iterations on a semi-implicit scheme for the regularity term problem. For sure though, even faster schemes may still be implemented, based, e.g., on operator splitting such as AOS (implemented but not reported in this thesis).

Since the problem is non-linear and non-convex, and because of the multiplicative relation between the split terms, no guarantee on convergence exists so far for this optimization method. Nonetheless, we have experienced good stability and convergence for sufficiently large – but finite and actually reasonably small – penalty weight. This is a clear advantage over pure penalty methods, where an infinitely large penalty weight is required in order to really enforce the (equality) constraint.

9.1.4 Registration on the Sphere

In order to register the spherical feature maps obtained after full inflation of the cortical surface mesh, we adapted the GAF framework to the sphere. This is somewhat particular, since the full sphere does not comply with the Riemannian requirements of the GAF method: No single, global and smooth parametrization can describe the complete spherical surface without singularities. To address this problem, we embed the spherical deformation field patch-wise, by discretizing the global deformation field hyperarea over each

Voronoi region separately. We introduce local coordinate charts and local coordinates around each vertex. The resulting discretized registration energy can be efficiently optimized by employing a similar fast minimization scheme as in the previous, “flat” case.

The registration scheme is shown to work successfully on a basic artificial image pair. It is to note that currently, most of the computational complexity of the spherical registration problem that we propose actually stems from spherical data interpolation; repeated warping and resampling of the entire moving image is required at each minimization iteration.

9.1.5 Automatic cerebro-cortical parcellation

In order to assess the spherical registration framework, and to provide a practical application, we have implemented a simple automatic parcellation scheme based on a set of available atlas brain surfaces together with their manually delineated gyral labelizations. The parcellation scheme consists of pairwise registration of a subject hemisphere to all available atlas hemispheres, individually. After registration, the atlas feature maps and labelizations are projected into subject space, where a Bayesian classifier can be trained. This classifier combines (1) the *a priori* probability of observing a specific label at a vertex, estimated from the observed frequency in the atlas data, (2) the likelihood of observing a specific label given the local subject surface geometry and the class-specific surface geometries learned from the atlas brains, and (3) the first-order Markovian likelihood of observing a specific label at a vertex given the estimations of its immediate mesh neighbors, as a measure of parcellation smoothness. The MAP solution of this classifier is very efficiently estimated using iterated conditional modes (ICM), in just a few iterations.

In a leave-one-out cross-validation study based on this automatic parcellation procedure, we were able to show that (1) the proposed Bayesian classifier outperforms popular majority voting atlas fusion. Further, thanks to the registration scheme proposed in this thesis, we obtain (2) significant improvement of the pairwise label agreement before classification, and subsequently (3) significant improvement of the parcellation after atlas fusion.

9.1.6 Diffeomorphic Deformations

As the small cherry on top, we sketch how recent, more sophisticated deformation models can easily be integrated into our proposed GAF framework. Indeed, we provide a deformation model and update rule that imitates the diffeomorphic demons method. Further, we show how the framework can then easily be made more inverse consistent, by including a symmetric image matching term based on the now-available inverse deformation. The data-symmetry reduces the registration differences due to the arbitrary role attribution of fixed and moving image.

Nevertheless, we note that the proposed data-symmetric diffeomorphic deformation model is still not entirely symmetric, since the regularization term only acts on the forward deformation field. However, the same restriction holds for many state-of-the-art models as well.

9.2 Future Work

The conclusions of the different parts constituting this thesis have revealed directions of future research, that may lead to further improvement of the proposed methods.

9.2.1 Short term

In particular, we think of the following “immediate” propositions:

- Today, we have not yet fully exploited the spatially adaptive regularization strength of the GAF framework. In particular, it would be interesting to test the framework on image pairs exhibiting spatially varying noise levels and/or topological defects, such as lesions present in only one of the images. Comparison with state-of-the-art methods, such as Demons, should then allow to prove this alleged strength of the proposed GAF algorithm.
- In order to overcome some shortcomings of the current spherical registration model, it would be important to include the diffeomorphic deformation model in the spherical scheme.
- In the automatic parcellation scheme, rigid pre-registration of the spherical feature maps might help improving the pairwise inter-subject agreement between labelmaps, and thus the outcome of the parcellation after Bayesian atlas fusion¹. Also, the registration parameters of the proposed spherical registration model can be optimized to improve the performance. Here, we the full-blown pairwise registration of all subjects was run only once, also in order to protect against overfitting.
- We also point out, that the current Bayesian atlas fusion scheme can still be improved. In particular, we suggest including better spatial and smoothness priors, e.g., since we know that boundaries between gyral regions will favorably be located at sulcal cusps (as opposite to gyral ridges).

¹Actually, in the course of this thesis, we have also developed a parametric, multiscale registration framework on the sphere, theoretically capable of rigid registration, but we decided not to include it in this document due to missing experimental validation and to keep the thesis at reasonable scope.

- The proposed GAF model only embeds the deformation field components along with the spatial coordinates, and the resulting regularizer is purely flow-driven. It is straightforward, however, to obtain a combined flow-intensity-driven regularizer, by including intensity or other features of the underlying image pair. This is in full accordance with the idea, that edges in the images are valid predictors for edges in the deformation field. Ultimately, since the Beltrami energy involves a coupling term between different embedded feature channels, this will allow resigning from a data-term weighting function altogether. Indeed, the mutual alignment of both image gradients, as well as flow-intensity driven regularization of the unknown deformation field, are all obtained by minimizing this extended hypersurface, corresponding to a model that we may call “Harmonic Active Fields” in analogy to the similarly defined harmonic active contours.

9.2.2 And beyond

Finally, we would like to end this thesis with a more “long distance perspective” of interesting research directions, unfolding from results achieved and insights gained during the last few years while working intensely in this domain.

Indeed, multiplicative regularization, with TV as regularization penalty, was already proposed in the context of inverse problems in wave propagation and scattering, (Abubakar and van den Berg, 2001), and shown to offer some benefit in image deblurring (Abubakar et al., 2004). However, only very recently, this work has raised some interest in the computer vision community (Orozco Rodríguez, 2011; Lam, 2011). Also, in contrast to the deformation field regularization proposed in this thesis, there data-term and regularizer are multiplied on a global level, i.e., both terms compete not locally, but globally.

Now, we propose a slightly different, more local approach. The main interest of multiplicative regularization is the adaptive modulation of the required regularization strength by the current *local* “regression quality”, i.e., the local error on the data-term. Depending upon the specific inverse problem, this error is constituted of *local* mismatch, misalignment, or noise level. Now, based on the general intuition on geometric regularization, we believe that the weighted Beltrami framework has great potential as generalized regularizing functional, and propose the following general approach, as a substitute for the classical Tikhonov regularization approach:

Weighted minimal surface problem: The unknown data u is embedded and the Polyakov energy penalizes its variations to obtain regularity. The weighting function includes the local contribution to the global data-term $\|\Phi u - f\|_p^p$, which we write $|\Phi u - f|_p^p(x)$ with some abuse of notation. In case of L^1 denoising, this term might be given by $|u - f|(x)$. The weighting

then drives the minimization towards the solution of the inverse problem:

$$\min_u \left\{ E = \int (1 + \alpha |\Phi u - f|_p^p(x)) \sqrt{g(u)} dx \right\} \quad (9.1)$$

Note, however, that determining the local data term contribution $|\Phi u - f|_p^p(x)$ may not be straightforward, depending on the specific inverse problem.

Since Beltrami regularization is used instead of TV, and because the multiplication is on a local basis, the resulting energy functional inherits very interesting properties already observed with GAF: (1) it maintains re-parametrization invariance, (2) it is applicable to any Riemannian manifold and not just specific Cartesian images, and (3) it seamlessly generalizes to regularization of vectorial inverse problems, such as deformations or color.

We are confident that, much like the FastGAF scheme, this approach will likely benefit from research in related fields, such as compressed sensing (Donoho, 2006) and optimization theory (Glowinski and Le Tallec, 1989), where very efficient optimization algorithms have been developed.

Also, we will focus further research on how the idea of data-term weighting can be extended to more recent regularization concepts, such as non-local regularization (Gilboa and Osher, 2008; Buades et al., 2008; Lou et al., 2009; Elmoataz et al., 2008; Peyré et al., 2008; Bougleux et al., 2011).

Some Spherical Tools

A

AS ILLUSTRATED in figure 3.1, the cortical maps are sampled at the vertices of a triangular surface mesh. In order to work with these samples, several tools need to be defined for the mesh and the function defined on it. Particular entities that are of interest are the gradient of the function defined on the mesh, i.e. its first order Taylor expansion in any vertex, as well as the interpolation of the function and its gradient at any other surface point.

One can distinguish two major approaches in defining those tools. The first one consists of expanding the mesh and the function using (differentiable) basis functions in order to perform analytical resolution of the aforementioned linear operations. The second approach exploits the available topological information of the triangulation to perform discretization of the desired operations in a more direct fashion. While the former way is more rigorous from a mathematical point of view, the latter has important benefits with respect to implementation and computation time, what makes it the approach of our choice.

A.1 Spherical Mesh Generation

A perfectly regular distribution of an arbitrary number n of nodes on the spherical surface is not easy to determine. If regularity is defined in terms of a minimal total potential energy I_n based on relative distances between the n points, the problem is called the *Fekete problem*. One choice to define the potential energy of an n -tuple of points $w_n = \{\vec{x}_1, \dots, \vec{x}_n | \vec{x}_i \in S^2\}$ uses Riesz's

energies:

$$I_n = \sum_{i \neq j} \frac{1}{\|\vec{x}_i - \vec{x}_j\|^s}, \quad (\text{A.1})$$

where in the general case $s > 0$. The choice of $s = 1$ leads to the standard Coulomb potential used to model electrons repelling each other (the *Thomson problem* (Thomson, 1904)). The limit $s \rightarrow \infty$ corresponds to the packaging problem, where the minimal angle between adjacent points needs to be maximized (the *Tammes problem* (Tammes, 1930)).

A.1.1 Regular polyhedra

Three regular, triangular polyhedra are solutions to the Thomson problem, corresponding to $n = 4, 6, 12$: the tetrahedron (4 vertices, 4 faces, 6 edges), the octahedron (6 vertices, 8 faces, 12 edges), and the icosahedron (12 vertices, 20 faces, 30 edges). The vertices of the isocahedron of edge-length 2 are

$$(x, y, z)_{12} = \{(0, \pm 1, \pm \phi), (\pm 1, \pm \phi, 0), (\pm \phi, 0, \pm 1)\},$$

where $\phi = (1 + \sqrt{5})/2$ is the golden ratio. The unit sphere octahedron vertices are given simply by

$$(x, y, z)_8 = \{(\pm 1, 0, 0), (0, \pm 1, 0), (0, 0, \pm 1)\},$$

and the tetrahedron vertices can be found in the corners of a cube:

$$(x, y, z)_4 = \{(1, 1, 1), (-1, -1, 1), (-1, 1, -1), (1, -1, -1)\}.$$

For arbitrary n , the Riesz energy criteria define a highly non-linear optimization problem. The resulting point sets are called *minimum energy* (ME) points. Other regularity criteria maximize the determinant of an arbitrary spherical polynomial basis. Such point distributions are called extremal or *maximum determinant* (MD) points (Womersley and Sloan, 2001).

A.1.2 Subdivision

A computationally interesting approach to generate highly regular point distributions on the sphere is recursive triangular subdivision of the above polyhedra. Each edge is split in two, and new edges connect the additional vertices to replace each triangle by four sub-triangles. Consequently, the triangular subdivision creates vertices with 6 neighbors each. The icosahedron has vertices with 5 neighbors each, and the resulting subdivision triangulation is therefore of highest regularity. In contrast, the tetrahedron with only 3 neighbors per vertex creates important artifacts, also with respect to regularity of the subdivision triangles.

At each iteration, the additional vertices are projected onto the spherical surface. An illustration of the subdivision meshes at different levels for each

of the three considered polyhedra is shown in figure A.1. In total, the number of faces and edges is multiplied by 4 at each iteration, while the number of vertices is given by the sum of vertices and edges in the previous iteration:

$$f_{i+1} = 4f_i \quad (\text{A.2})$$

$$e_{i+1} = 4e_i \quad (\text{A.3})$$

$$v_{i+1} = v_i + e_i, \quad (\text{A.4})$$

where f_i , e_i and v_i denote the number of faces, edges and vertices in the mesh at iteration i , respectively. Moreover, these three numbers obey Euler's formula at all time:

$$\chi = v - e + f = 2, \quad (\text{A.5})$$

where χ is called the Euler characteristic. On top of that, in triangular meshes the number of edges and faces is directly calculated by the number of vertices as

$$e = 3v - 6 \quad (\text{A.6})$$

$$f = 2v - 4. \quad (\text{A.7})$$

Therefore, the number of vertices distributed on the sphere evolves as

$$v_{i+1} = 4v_i - 6, \quad (\text{A.8})$$

and the set of possible numbers of vertices distributed on the sphere by a triangular subdivision scheme is spanned by a geometric series of ratio 4:

$$v_i = 4^i(v_0 - 2) + 2, \quad (\text{A.9})$$

where the coefficient v_0 is the number of vertices in the initial polyhedron. At first sight, this set of possible node configurations might seem too restrictive. The geometric series, however, corresponds well to the exponential sampling of equidistant scales.

A.2 Mesh smoothing

Some of the meshes produced by freesurfer are not perfectly inflated to the sphere. Although all vertices are eventually projected on the spherical surface, the mesh contains some remaining microfolds. These folds are not visible at macroscopic scale and probably have no impact. At small scale, they are characterized by vertices being located within neighboring or slightly distant triangles. Three layers of the mesh cover the same spherical region. This not only undermines the spherical triangle walking algorithm presented later in this appendix – the middle layer is oriented oppositely to the spherical tangent plane – it also means that the spherical parametrization of the cortical surface is not bijective. Three cortical points map to the same spherical location.




Figure A.1: Polyhedra subdivision for regular point distribution on the sphere. The integer v under the subdivision mesh indicates the number of vertices distributed on the sphere. The shaded patch corresponds to a single face in the initial polyhedron. **Top:** Tetrahedron and 4 subsequent subdivision levels. The meshes suffer from heavy artifacts and uneven triangle size. **Center:** Octahedron and 4 subsequent subdivision levels. The subdivided meshes have 6 vertices that are connected to 4 neighbors only. **Bottom:** Icosahedron and 4 subsequent subdivision levels. The icosahedron subdivisions contains 12 vertices that have only 5 direct neighbors. Among the three polyhedron subdivision meshes, the icosahedron presents highest regularity.

The unmet bijectivity condition represents a major problem. To overcome it, we propose to perform mesh smoothing. Implicit fairing using Laplacian diffusion does a good job (Desbrun et al., 1999).

The Laplacian matrix of the mesh is built as follows:

$$L_{ij} = \begin{cases} \frac{1}{m_i} & j \in N_1(i) \\ -1 & i = j \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.10})$$

This very simple Laplacian is also known as the umbrella operator. It drives a vertex i to the center of gravity of its surrounding first-order neighbors $j \in N_1(i)$.




Figure A.2: *Detail of the original and smoothed mesh, using implicit Laplacian smoothing.* The colored triangles initially contain vertices of neighboring faces and thus indicate the presence of microfolds. Coloring is conserved in the right-hand image, although no more folds are detected. Both images represent the same angle-of-view.

The diffusion equation of the vertex coordinates $\mathbf{x} \in \mathbb{R}^{N \times 3}$ is written as

$$\frac{\partial \mathbf{x}}{\partial t} = \lambda L \mathbf{x}, \quad (\text{A.11})$$

where λ is the diffusion coefficient. We apply time-shifted evaluation to formulate an implicit integration scheme, also known as backward Euler scheme:

$$(I - \lambda dt L) \mathbf{x}^{n+1} = \mathbf{x}^n, \quad (\text{A.12})$$

where the exponent refers to the time index of the coordinates. After each integration, the vertices must be reprojected on the spherical surface. The parameter λdt governs the speed and precision of the integration scheme. In contrast to the simpler forward Euler scheme, the implicit scheme is unconditionally stable, but requires resolution of a linear system. As the Laplacian matrix is highly sparse, this is nevertheless achieved in reasonable time. The results are illustrated in figure A.2.

A.3 Gradient estimation

At several occasions, we need to estimate the gradient of a function on a spherical triangulated mesh. Here we present a solution we chose in our implementation.

A.3.1 Definitions

Let the surface be defined by $\vec{r} = \vec{r}(u^1, u^2) = (x(u^1, u^2), y(u^1, u^2), z(u^1, u^2))^T$ in \mathbb{R}^3 ¹. Let $\vec{r}_{,\alpha} = \frac{\partial \vec{r}}{\partial u^\alpha}$, $\alpha = 1, 2$ be the intrinsic tangent vectors, defining the

¹The gradient of the attached function over the 2D manifold is an intrinsic measure, i.e., the dimension and structure of the embedding space does not matter.

tangent plane. Further, we define a data function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, that maps a feature to each point of the 2D manifold ². The gradient of the function with respect to the parametrization is then defined as

$$\nabla_S f = \left(\frac{\partial f}{\partial u^1}, \frac{\partial f}{\partial u^2} \right), \quad (\text{A.13})$$

and within the embedding space, this is written

$$\nabla f = [\vec{r}_{,1}, \vec{r}_{,2}] \left(\frac{\partial f}{\partial u^1}, \frac{\partial f}{\partial u^2} \right)^T = \vec{r}_{,\alpha} f_{,\alpha}, \quad (\text{A.14})$$

where in the second equality $f_{,\alpha} = \frac{\partial f}{\partial u^\alpha}$, and Einstein summation was used.

A.3.2 Gradient estimation on triangulated maps

Unfortunately, the data function is only sampled at the vertices. These samples can be interpolated using differentiable basis functions to analytically compute its gradient at the vertices. Nevertheless, the data values attached to the vertices are usually interpolated linearly over the surface, which is not differentiable at the vertices, and estimators need to be used.

Again, one can use spatial averaging of the gradient information around the vertex (Lopez-Perez et al., 2004). Alternatively, the gradient can be fitted using least squares approximation (Frank and Lang, 2000). As one is interested in the gradient *within* the tangent plane of the manifold, it is useful to enforce this by design. Therefore we will extend the latter approach to this end through projection of the data into the local plane, prior to gradient fitting.

First, we construct a local coordinate chart and determine the local coordinates of the neighboring vertices, as defined in (6.2) and (6.3).

A first order Taylor expansion of the feature map f around \vec{x}_i , using these local coordinates $\vec{\theta}_i$ now yields:

$$f(\vec{x}_j) = f(\vec{x}_i) + \left\langle \vec{\theta}_j, \nabla_S f \right\rangle + \epsilon. \quad (\text{A.15})$$

To simplify the notation, we define

$$f'_j = \frac{f(\vec{x}_j) - f(\vec{x}_i)}{\|\vec{\theta}_j\|_2}, \quad (\text{A.16})$$

where $\|\cdot\|_2$ denotes the L_2 -norm, and rewrite equation (A.15) accordingly:

$$f'_j = \left\langle \frac{\vec{\theta}_j}{\|\vec{\theta}_j\|_2}, \nabla_S f \right\rangle + \epsilon. \quad (\text{A.17})$$

²The extension to multidimensional features is straightforward and omitted for simplicity




Figure A.3: *Gradient estimation on the cortical map.* **(a)** The mean-curvature map is projected on the partially inflated surface. **(b)** The norm of the estimated gradient, and **(c)** its azimuth.

Now we look for an optimal $\nabla_S f$ that minimizes ϵ in the least square sense. We define the following matrix

$$A_{kj} = \frac{[\vec{\theta}_j]_k}{\|\vec{\theta}_j\|_2}, \quad (\text{A.18})$$

where $k = 1, 2$ denotes the components of the local coordinates vector $\vec{\theta}_j$. The minimization problem and its direct solution can now be written in matrix notation as

$$\begin{aligned} \widehat{\nabla_S f} &= \operatorname{argmin}_{\vec{g}} (\|A\vec{g} - \vec{y}\|_2) \\ &= (A^T A)^{-1} A^T \vec{y}. \end{aligned} \quad (\text{A.19})$$

In figure A.3 we show an example of estimated gradient of a cortical feature map, projected onto the inflated surface for easier understanding.

A.4 Spherical point location and data interpolation

In order to perform image or deformation field interpolation, resampling and warping on the spherical surface, we need to locate a query point (e.g. vertex of the fixed mesh) within the correct triangle in the target mesh (e.g. moving image). Point location is one of the basic problems in computational geometry. Let S be a (triangular) mesh and p the query point, then point location means finding the polygon (triangle) T in S containing p . Point location is a key step for the interpolation of a data function, but is also used in real time classification.

A.4.1 Triangle containment

In order to determine whether a point p be contained by the triangle T , it is useful to define the barycentric coordinates of p with respect to T . Let the triangle be defined by its three vertices \vec{v}_i , \vec{v}_j and \vec{v}_k . These three vectors form

a basis of \mathbb{R}^3 , and any point p with Cartesian coordinates \vec{p} may be represented as a weighted sum of them, i.e.

$$\vec{p} = {}_T\lambda_1\vec{v}_i + {}_T\lambda_2\vec{v}_j + {}_T\lambda_3\vec{v}_k, \quad (\text{A.20})$$

where ${}_T\lambda_1$, ${}_T\lambda_2$ and ${}_T\lambda_3$ are the barycentric coordinates of p with respect to T . If the point p lies in the same plane as the triangle T , then the barycentric coordinates are subjected to the constraint

$${}_T\lambda_1 + {}_T\lambda_2 + {}_T\lambda_3 = 1. \quad (\text{A.21})$$

Further, if p lies within the triangle T , all barycentric coordinates are non-negative. Substituting ${}_T\lambda_3 = 1 - {}_T\lambda_1 - {}_T\lambda_2$ into equation (A.20) and rearranging yields

$$\vec{p} = \vec{v}_k + {}_T\lambda_1(\vec{v}_i - \vec{v}_k) + {}_T\lambda_2(\vec{v}_j - \vec{v}_k), \quad (\text{A.22})$$

where \vec{p} is required to be in the triangular plane. p is now contained by T if both ${}_T\lambda_1$ and ${}_T\lambda_2$ are non-negative, and if ${}_T\lambda_1 + {}_T\lambda_2 \leq 1$.

In the spherical case, both the query point p and the vertices of the triangles lie on the spherical hull, but the triangle edges are cords of the sphere. Consequently, the query point will generally not lie within the triangular plane. In this case, triangle containment is understood as the radial beam from the sphere center through the query point intersecting the triangular plane within the triangle. Equation (A.22) is extended by an additional degree of freedom along the radial beam, defined by the vector $(\vec{p} - \vec{c})$:

$$\vec{p} = \vec{v}_k + {}_T\lambda_1(\vec{v}_i - \vec{v}_k) + {}_T\lambda_2(\vec{v}_j - \vec{v}_k) + \epsilon(\vec{p} - \vec{c}), \quad (\text{A.23})$$

where the value of ϵ is not of actual interest. The introduction of this additional degree of freedom leads to some ambiguity as triangles of opposite sides of the sphere might now concurrently apply as containing triangles. In such a case, the candidate triangle exhibiting smaller $|\epsilon|$ should be selected.

A.4.2 Classical point location schemes

A naive solution to point location is exhaustive search through the whole set of mesh-cells. Such an algorithm is expected to perform in $O(n)$. The first efficient point location algorithm in the plane was proposed by Dobkin and Lipton (Dobkin and Lipton, 1976). They put parallel lines L through any vertex of a planar mesh, and any pair of neighboring lines L_i and L_{i+1} forms a band. The band itself is structured into subregions by a set of non-intersecting mesh edges. For a new point p , both the band and the subregion within the band can be determined in $O(\log n)$, using binary search. However, the bands and subregions need to be pre-calculated, requiring a considerable amount of space $O(n^2 \log n)$ and time $O(n^2)$. Kirkpatrick proposed a point location algorithm making use of hierarchical subdivision of the strictly triangular mesh cells, that locates in $O(\log n)$ as well, but requires only $O(n)$ space and $O(n)$ time for

preparation (Kirkpatrick, 1983). They recursively embed several mesh cells in one or more parent cells up to a certain height during preparation, while for point location the algorithm descends this hierarchy, looking for the respective daughter cell containing the query point p . Other authors have provided more practical algorithms and implementations with similar performance (Sarnak and Tarjan, 1986; Edelsbrunner et al., 1986).

In contrast to the aforementioned algorithms, that all are based on a *divide and conquer* strategy, another class is formed by so-called walk-through algorithms. The general complexity of these algorithms is $O(\sqrt{n})$, without any extra preparation except having direct access to the neighboring triangles. Starting from a random triangle, these algorithms walk through the mesh until they find the triangle containing the point p . Several authors provided different decisions on how to progress through the mesh: Mücke et al. perform a segment intersection test (Mücke et al., 1996). A line segment between the starting point and the query point is drawn, and one walks only over triangles that are intersected by this line segment. Earlier, Guibas and Stolfi proposed the *Counter Clock Wise Search* (Guibas and Stolfi, 1985). They compute the determinant of an edge with respect to the query point and compare this to the determinant of the vertex opposite to that edge – if the two determinants are of opposite sign, then the edge is crossed. In case both remaining edges of the triangle could be crossed, there is ambiguity, and the path begins to meander. Sundareswara et al. resolve this ambiguity with the use of barycentric coordinates to obtain a straighter path (Sundareswara and Schrater, 2003). This algorithm computes the barycentric coordinates of the query point with respect to the current triangle, and crosses the edge corresponding to the two highest coordinates, until all barycentric coordinates are positive. Wu et al. port this point location algorithm to the sphere, and add a pre-processing stage that provides a good starting triangle for each query point p (Wu et al., 2005). This is basically a combination of a divide-and-conquer algorithm to determine the start triangle, and of a walk-through algorithm for final point location.

A.4.3 Spherical triangle walk

In the present context, spherical interpolation will take place in 2 well-defined situations: During warping, to interpolate the fixed map at the vertices of the moving map, and during velocity field exponentiation, defining the deformation field in a diffeomorphic setting. In both cases, a good initial guess for the starting triangle can be obtained easily, and spherical triangle walk is a promising point location and interpolation method.

For warping, as the deformation does not change drastically between iterations, the triangle hit of the previous iteration can be used as a good guess for triangle walk initialization. At the very first iteration, such a “history” is not available, but nice guesses can easily be approximated by (randomly) looking for corresponding triangles no farther than a certain threshold.

To locate a point p on the sphere and starting at a given triangle T , we compute the local coordinates ${}_T\lambda_1$ and ${}_T\lambda_2$ of p w.r.t. T according to equation (A.23). If p is not contained in T , the coordinates can be signed in four different ways, and the next triangle is selected accordingly. If both coordinates are positive, the next triangle selected shares the edge $\vec{v}_i\vec{v}_j$, opposite to the vertex \vec{v}_k . If only one coordinate is positive, the corresponding edge is crossed, i.e. the next triangle will be the adjacent face, which is opposite to the negative vertex. Finally, if both coordinates are negative, the query point lies “behind” \vec{v}_k , and we select a triangle linked to \vec{v}_k , but that does not share any of the vertices \vec{v}_i or \vec{v}_j . Triangle walk proceeds until the containing triangle is found, usually within a reasonable amount of time. See figure A.4 for a long walk on the sphere.

A major condition for this triangle walk scheme to work, is having a sufficiently convex mesh. In deed, if the mesh is too much convoluted, the triangle plane and the “mean” tangent plane are oppositely oriented. Thus, the algorithm will walk in the wrong direction, and it can get stuck indefinitely. An easy but only partially successful solution to this problem consist in performing short random walks as soon as the normal walk gets stuck. Alternatively, the mesh is processed to remove undesirable folds, as seen below. An instance of a triangle walk getting lost in a microfold, and its success in a smoothed mesh, are shown in figure A.4.

A.4.4 Vertex binning

In cases, where the mesh surface is not regular, i.e. not “sufficiently” spheroid, and where smoothing is not possible, we propose a different approach to point location. It is based on a divide-and-conquer preparation step. But in contrast to the previously mentioned methods, we index faces according to the spatial location of their vertices only. This is based on the fundamental assumption that if the query point p is in the immediate vicinity of a vertex v , then there is a high chance that the containing triangle shares this vertex. For a sufficiently regular triangulation, i.e. non-obtuse triangles, this is generally true ³.

Our method consists in binning the vertices as a preparation step. Once the vertices are binned, they are replaced by the corresponding faces. The same face can therefore be assigned to three different bins. A query point can directly be mapped into a corresponding bin as well. According to the fundamental assumption mentioned above, we expect to find the containing triangle in the same bin, and the all respective faces of the bin are checked for containment according to equation (A.23). If no triangle within the selected bin contains the query point, an appropriate fallback strategy needs to be defined. The expected complexity of the algorithm does not depend on the number of faces or vertices of the mesh, i.e. $O(1)$, as the size of the bins can be

³If a 1-ring patch consists of only non-obtuse triangles, then it entirely covers the Voronoi region corresponding to its central vertex. All points having the central vertex as its nearest neighbor then automatically lie within one of the 1-ring triangles.




Figure A.4: *Triangle walk for spherical point location.* **(a)** Global view of a single trajectory over a relatively large walking distance. Point location is initialized at the top left end of the trajectory, the target point to be located is marked with a green circle, at the lower right end of the trajectory. **(b)** Density plot of 250 trajectories between the same start and target triangles. When both containment coordinates are negative, the algorithm chooses randomly among the triangles that do not share any of the vertices \vec{v}_i or \vec{v}_j , providing some randomness w.r.t. to the chosen trajectory. All observed trajectories cover between 161 and 207 triangles, the mode being at 192. **(c)** Close-up of a trajectory trapped in a microfold of a different mesh. The point to be located is the barycenter of the triangle shaded in grey. **(d)** After mesh smoothing, the microfold has disappeared. The dense region is easily traversed, and the target is reached within few iterations.

chosen arbitrarily small. The worst-case performance depends on the fallback strategy.

A simple approach is to bin the vertices of the mesh according to their spherical coordinates, θ and ϕ . This requires $O(n)$ both in space and time. If no triangle in the bin contains the query point, full blown search is performed over the whole mesh. This θ - ϕ binning is easily implementable, but has important drawbacks. The main problem resides in the fact, that a triangle can span part of a bin, even without having one of its vertices in this bin. This is particularly the case in polar regions, where the bins get extremely narrow. In case of failure, the fallback strategy is not very interesting, as it directly leads to $O(n)$.

A.4.5 Octree binning

A more sophisticated approach makes use of hierarchical binning of the vertices, to allow for a smoother fallback strategy. We propose to bin the vertices according to a octree division of the 3D space embedding the spherical mesh. Also, coordinate transforms are thus not necessary anymore, as the 3D coordinates of the vertices can directly be used for binning.

A global cube enveloping all the vertices of the mesh is divided into 8 identical octant cubes. Each of these cubes is recursively split in the same way, until a certain stop criterion is met. The set of all cubes at all levels form the nodes of an octree. Each cube in the tree has either eight children or no children. All cubes, that have no children, are called terminal nodes, or leaves of the tree. The volumes spanned by the terminal nodes form a partition of the initial volume, i.e. each point within the initial cube is contained in exactly one terminal cube. Each terminal node writes its label into a cubic lookup-table, typically a 3D image of resolution corresponding to the dimensions of the smallest terminal node.

We choose to stop the recursive splitting, if the number of vertices contained in a cube falls below a certain threshold. Each node of the octree is aware of the vertices it contains, and it knows its parent cube. Again, once the vertices are binned, they are replaced by the corresponding faces. The binning requires a total of $O(n \log n)$ space, and can be achieved in $O(n \log n)$ time.

Thanks to the hierarchical structure of the octree binning, one disposes of a gradual fallback strategy, with full blown search necessary only at the very last resort. Although worst-case search performance is still of $O(n)$, the average performance with respect to θ - ϕ binning is expected to improve significantly, thanks to the intermediate fallback steps.

A.4.6 Interpolation

Once the containing triangle of a query point has been determined according to equation (A.23), we use linear interpolation to estimate the data function at the query point. The barycentric coordinates, i.e. the coefficients ${}_T\lambda_1$ and




Figure A.5: *Different aspects of octree vertex binning.* (a) The regular bin-pattern on the sphere, obtained using θ - ϕ binning, as opposed to (b) irregular binning pattern resulting from spatial octree splitting. (c) 3D visualization of the octree terminal nodes, containing the spherical vertices. Empty cubes are only drawn in the posterior half. Cube color corresponds to subdivision depth, transparency indicates lower density cubes. (d) Terminal node cell-size histogram. The threshold below which a cube is not further split has been chosen at 40. (e) Frequency plot of the different node levels, showing that not many empty cubes are generated as overhead. 7 levels were sufficient to bin the 130k vertices in bins smaller than 40. (f) The triangular spherical mesh has been resampled at 1.6M vertices of a regular θ - ϕ spherical mesh. More than 99% of these points could be located within the leaf cube, and only one instance required full blown search.

${}_T\lambda_2$, are used to interpolate the data function from the three vertex-samples as follows:

$$\hat{f}(\vec{p}) = {}_T\lambda_1 f(\vec{v}_i) + {}_T\lambda_2 f(\vec{v}_j) + (1 - {}_T\lambda_1 - {}_T\lambda_2) f(\vec{v}_k). \quad (\text{A.24})$$

Bibliography

- Abubakar, A. and van den Berg, P. M. (2001). Total variation as a multiplicative constraint for solving inverse problems. *IEEE Transactions on Image Processing*, 10(9):1384–1392.
- Abubakar, A., van den Berg, P. M., Habashy, T. M., and Braunisch, H. (2004). A multiplicative regularization approach for deblurring problems. *IEEE Transactions on Image Processing*, 13(11):1524–1532.
- Adalsteinsson, D. and Sethian, J. A. (1995). A Fast Level Set Method for Propagating Interfaces. *Journal of Computational Physics*, 118(2):277–269.
- Adelson, E. H. and Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. *Nature*, 300:523–525.
- Alvarez, L., Lions, P.-L., and Morel, J.-M. (1992). Image Selective Smoothing and Edge Detection by Nonlinear Diffusion. II. *SIAM Journal on Numerical Analysis*, 29(3):845–866.
- Angenent, S., Haker, S., Tannenbaum, A., and Kikinis, R. (1999). On the Laplace-Beltrami operator and brain surface flattening. *IEEE Transactions on Medical Imaging*, 18(8):700–711.
- Ardekani, B. a., Guckemus, S., Bachman, A., Hoptman, M. J., Wojtaszek, M., and Nierenberg, J. (2005). Quantitative comparison of algorithms for inter-subject registration of 3D volumetric brain MRI scans. *Journal of Neuroscience Methods*, 142(1):67–76.
- Arsigny, V., Commowick, O., Pennec, X., and Ayache, N. (2006). A Log-Euclidean Framework for Statistics on Diffeomorphisms. In *MICCAI*, pages 924–931.
- Ash, M. G. (1985). Die experimentelle Psychologie an den deutschsprachigen Universitäten von der Wilhelminischen Zeit bis zum Nationalsozialismus. In Ash, M. G. and Geuter, U., editors, *Geschichte der deutschen Psychologie im 20. Jahrhundert*, page 62. Westdeutscher Verlag.

- Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *Neuroimage*, 38(1):95–113.
- Ashburner, J. and Friston, K. J. (1999). Nonlinear Spatial Normalization Using Basis Functions. *Human Brain Mapping*, 7:254–266.
- Audette, M. A., Ferrie, F. P., and Peters, T. M. (2000). An algorithmic overview of surface registration techniques for medical imaging. *Med. Image Anal.*, 4(3):201–217.
- Babaud, J., Witkin, A. P., Baudin, M., and Duda, R. O. (1986). Uniqueness of the Gaussian kernel for scale-space filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(1):26–33.
- Bach Cuadra, M., Gomez, J., Hagmann, P., Pollo, C., Villemure, J.-G., Dawant, B. M., and Thiran, J.-P. (2002). Atlas-based segmentation of pathological brains using a model of tumor growth. In *MICCAI 2002*, pages 380–387. Springer.
- Bajaj, C. L. and Xu, G. (2003). Anisotropic Diffusion of Surfaces and Functions on Surfaces. *ACM Transactions on Graphics*, 22(1):4–32.
- Bakircioglu, M. M. (1999). *Landmark matching on brain surfaces via large deformation diffeomorphisms on the sphere*, volume 3661. SPIE.
- Bar, L., Sochen, N., and Kiryati, N. (2006). Semi-blind image restoration via Mumford-Shah regularization. *IEEE Transactions on Image Processing*, 15(2):483–493.
- Barash, D. (2002). Fundamental relationship between bilateral filtering, adaptive smoothing, and the nonlinear diffusion equation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(6):844–847.
- Barash, D., Schlick, T., Israeli, M., and Kimmel, R. (2003). Multiplicative Operator Splittings in Nonlinear Diffusion: From Spatial Splitting to Multiple Timesteps. *J. Math. Imaging Vis.*, 19(1):33–48.
- Barber, D. C., Oubel, E., Frangi, a. F., and Hose, D. R. (2007). Efficient computational fluid dynamics mesh generation by image registration. *Medical image analysis*, 11(6):648–62.
- Beg, M. F., Miller, M. I., Trouvé, A., and Younes, L. (2005). Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms. *Int. J. Comput. Vis.*, 61(2):139–157.
- Behrens, R. R. (2004). Art, Design and Gestalt Theory.
- Ben-Ari, R. and Sochen, N. (2004). Non-isotropic regularization of the correspondence space in stereo-vision. In *ICPR 2004*, volume 4, pages 293–296.

- Ben-Ari, R. and Sochen, N. (2008). A geometric approach for regularization of the data term in stereo-vision. *J. Math. Imaging Vis.*, 31(1):17–33.
- Ben-Ari, R. and Sochen, N. (2009). A geometric framework and a new criterion in optical flow modeling. *J. Math. Imaging Vis.*, 33(2):178–194.
- Bertero, M., Poggio, T. A., and Torre, V. (1988). Ill-Posed Problems in Early Vision. In *Proceedings of the IEEE*, volume 76, pages 869–889.
- Bertsekas, D. P. (1976). Multiplier methods: A survey. *Automatica*, 12(2):133–145.
- Besag, J. (1986). On the statistical analysis of dirty pictures (with discussion). *Journal of the Royal Statistical Society. Series B*, 45(3):259–302.
- Bogdanova, I., Bresson, X., Thiran, J.-P., and Vanderghenst, P. (2007). Scale space analysis and active contours for omnidirectional images. *IEEE Transactions on Image Processing*, 16(7):1888–1901.
- Bossa, M., Zacur, E., and Olmos, S. (2008). Algorithms for computing the group exponential of diffeomorphisms: Performance evaluation. In *CVPRW*, pages 1–8. IEEE.
- Bougleux, S., Peyré, G., and Cohen, L. (2011). Non-local regularization of inverse problems. *Inverse Problems and Imaging*, 5(2):511–530.
- Bresson, X., Vanderghenst, P., and Thiran, J.-P. (2006). Multiscale Active Contours. *Int. J. Comput. Vis.*, 70(3):197–211.
- Brown, L. G. (1992). A survey of image registration techniques. *ACM Comput. Surveys*, 24(4):325–376.
- Buades, A., Coll, B., and Morel, J.-M. (2008). Nonlocal Image and Movie Denoising. *Int. J. Comput. Vis.*, 76:123–139.
- Bunschoten, R. and Krose, B. (2003). Robust scene reconstruction from an omnidirectional vision system. *IEEE Transactions on Robotics and Automation*, 19(2):351–357.
- Cachia, A., Mangin, J.-F., Riviere, D., Kherif, F., Boddaert, N., Andrade, A., Papadopoulos-Orfanos, D., Poline, J.-B., Bloch, I., Zilbovicius, M., Sonigo, P., Brunelle, F., and Regis, J. (2003). A primal sketch of the cortex mean curvature: a morphogenesis based approach to study the variability of the folding patterns. *IEEE Transactions on Medical Imaging*, 22(6):754–765.
- Cammoun, L., Gigandet, X., Sporns, O., Thiran, J., Deppen, P., Krieger, E., Maeder, P., Meuli, R., Hagmann, P., Bovet, P., and Do, K. (2009). Connectome alterations in schizophrenia. *NeuroImage*, 47:S157.

- Caselles, V., Kimmel, R., and Sapiro, G. (1997). Geodesic Active Contours. *Int. J. Comput. Vis.*, 22(1):61–79.
- Catte, F., Lions, P.-L., Morel, J.-M., and Coll, T. (1992). Image Selective Smoothing and Edge Detection by Nonlinear Diffusion. *SIAM Journal on Numerical Analysis*, 29(1):182.
- Chan, T. F. and Vese, L. A. (2001). Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277.
- Christensen, G. E., Joshi, S. C., and Miller, M. I. (1997). Volumetric transformation of brain anatomy. *IEEE Transactions on Medical Imaging*, 16(6):864–77.
- Christensen, G. E., Rabbitt, R. D., and Miller, M. I. (1996). Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(10):1435–1447.
- Chui, H., Rambo, J., Duncan, J., Schultz, R., and Rangarajan, A. (1999). Registration of Cortical Anatomic Structures via Robust 3D Point Matching. *Lecture Notes in Computer Science*, 1613.
- Cohen, L. D. (1996). Auxiliary Variables and Two-Step Iterative Algorithms in Computer Vision Problems. *J. Math. Imaging Vis.*, 6(1):59–83.
- Collins, D. L., Le Goualher, G., and Evans, A. C. (1998). Non-linear Cerebral Registration with Sulcal Constraints. In *MICCAI'98*.
- Courant, R., Friedrichs, K., and Lewy, H. (1928). Über die partiellen Differenzengleichungen der mathematischen Physik. *Mathematische Annalen*, 100(1):32–74.
- Dale, A. M., Fischl, B., and Sereno, M. I. (1999). Cortical surface-based analysis - I: Segmentation and surface reconstruction. *Neuroimage*, 9(2):179–194.
- Dascal, L., Rosman, G., and Kimmel, R. (2007). Efficient Beltrami filtering of color images via vector extrapolation. In *SSVM'07: Proceedings of the 1st international conference on Scale space and variational methods in computer vision*, pages 92–103. Springer-Verlag.
- Dascal, L., Rosman, G., Tai, X.-C., and Kimmel, R. (2009). On Semi-implicit Splitting Schemes for the Beltrami Color Flow. In *Scale Space and Variational Methods in Computer Vision*, volume 5567 of *Lecture Notes in Computer Science*, pages 259–270. Springer Berlin / Heidelberg.
- Dascal, L. and Sochen, N. A. (2005). A Maximum Principle for Beltrami Color Flow. *SIAM J. Appl. Math.*, 65(5):1615–1632.

- Derraz, F., Taleb-Ahmed, A., Peyrodie, L., Pinti, A., Chikh, A., and Bereksi-Reguig, F. (2009). Active Contours Based Battachryya Gradient Flow for Texture Segmentation. In *CISP '09*, pages 1–6.
- Desbrun, M., Meyer, M., Schröder, P., and Barr, A. H. (1999). Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. ACM SIGGRAPH '99*, pages 317–324.
- Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., and Killiany, R. J. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31(3):968–80.
- Dice, L. R. (1945). Measures of the Amount of Ecologic Association Between Species. *Ecology*, 26(3):297–302.
- Dobkin, D. and Lipton, R. J. (1976). Multidimensional Searching Problems. *SIAM J. Comput.*, 5(2):181–186.
- Donoho, D. L. (2006). Compressed Sensing. *IEEE Transaction on Information Theory*, 52(4):1289–1306.
- Duncan, D. B. and Lynch, M. A. M. (1991). Jacobi Iteration in Implicit Difference Schemes for the Wave Equation. *SIAM J. Numer. Anal.*, 28(6):1661–1679.
- Eberly, D. H. (1994). A differential geometric approach to anisotropic diffusion. In ter Haar Romeny, B. M., editor, *Geometry-Driven Diffusion in Computer Vision*, pages 371–392. Kluwer.
- Edelsbrunner, H., Guibas, L. J., and Stolfi, J. (1986). Optimal Point Location in a Monotone Subdivision. *SIAM J. Comput.*, 15(2):317–340.
- Elmoataz, A., Lezoray, O., and Bogleux, S. (2008). Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing. *IEEE Transactions on Image Processing*, 17(7):1047–1060.
- Estellers, V., Zosso, D., Bresson, X., and Thiran, J.-P. (2011). Harmonic active contours for multichannel image segmentation. In *18th IEEE International Conference on Image Processing ICIP2011*.
- Faugeras, O. and Keriven, R. (1998). Variational principles, surface evolution, PDE's, level set methods, and the stereo problem. *IEEE Transactions on Image Processing*, 7(3):336–344.
- Feigin, M. and Sochen, N. (2009). Anisotropic Regularization for Inverse Problems with Application to the Wiener Filter with Gaussian and Impulse Noise. In *Scale Space and Variational Methods in Computer Vision*, volume 5567

- of *Lecture Notes in Computer Science*, pages 319–330. Springer Berlin / Heidelberg.
- Fischl, B., Sereno, M. I., and Dale, A. M. (1999a). Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage*, 9(2):195–207.
- Fischl, B., Sereno, M. I., Tootell, R. B., and Dale, A. M. (1999b). High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.*, 8(4):272–284.
- Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Ségonne, F., Salat, D. H., Busa, E., Seidman, L. J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., and Dale, A. M. (2004). Automatically Parcellating the Human Cerebral Cortex. *Cerebral Cortex*, 14(1):11–22.
- Florack, L., Duits, R., and Bierkens, J. (2004). Tikhonov regularization versus scale space: a new result. In *International Conference on Image Processing*, volume 1, pages 271–274.
- Florack, L. M. J., Salden, A. H., ter Haar Romeny, B. M., Koenderink, J. J., and Viergever, M. A. (1995). Nonlinear scale-space. *Image and Vision Computing*, 13(4):279–294.
- Forsberg, D., Andersson, M., and Knutsson, H. (2010). *Adaptive anisotropic regularization of deformation fields for non-rigid registration using the morphon framework*. Number x. IEEE.
- Frank, K. and Lang, U. (2000). Gradient and curvature approximation in data dependent surface simplification. *Comput. Visual Sci.*, 2:221–228.
- Gaspar, J., Winters, N., and Santos-Victor, J. (2000). Vision-based navigation and environmental representations with an omnidirectional camera. *IEEE Transactions on Robotics and Automation*, 16(6):890–898.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *Journal of Applied Statistics*, 20(6):721–741.
- Geyer, C. and Daniilidis, K. (2000). A Unifying Theory for Central Panoramic Systems and Practical Implications. In *Computer Vision ECCV 2000*, volume 1843 of *Lecture Notes in Computer Science*, pages 445–461. Springer Berlin / Heidelberg.
- Geyer, C. and Daniilidis, K. (2001). Catadioptric Projective Geometry. *Int. J. Comput. Vis.*, 45(3):223–243.
- Gholipour, A., Kehtarnavaz, N., Briggs, R., Devous, M., and Gopinath, K. (2007). Brain functional localization: a survey of image registration techniques. *IEEE transactions on medical imaging*, 26(4):427–51.

- Gilboa, G. and Osher, S. (2008). Nonlocal Operators with Applications to Image Processing. *Multiscale Modeling & Simulation*, 7(3):1005–1028.
- Gizzi, M. S., Katz, E., Schumer, R. a., and Movshon, J. a. (1990). Selectivity for Orientation and Direction of Motion of Single Neurons in Cat Striate and Extrastriate Visual Cortex. *Journal of Neurophysiology*, 63(6):1529–1543.
- Glowinski, R. and Le Tallec, P. (1989). *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia.
- Goldstein, J. (1999). Emergence as a Construct: History and Issues. *Emergence*, 1(1):49–72.
- Goldstein, T. and Osher, S. (2009). The Split Bregman Method for L1-Regularized Problems. *SIAM J. Imaging Sci.*, 2(2):323–343.
- Gonzalez, R. C. and Woods, R. E. (1992). *Digital Image Processing*. Addison Wesley.
- Gray, H. (1918). *Anatomy of the Human Body*. Lea & Febiger, Philadelphia, 20th edition.
- Gu, X., Wang, Y., Chan, T. F., Thompson, P. M., and Yau, S. T. (2004). Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, 23(8):949–958.
- Guibas, L. and Stolfi, J. (1985). Primitives for the manipulation of general subdivisions and the computation of Voronoi. *ACM Trans. Graph.*, 4(2):74–123.
- Gur, Y., Pasternak, O., and Sochen, N. (2009). Fast GL(n)-Invariant Framework for Tensors Regularization. *Int. J. Comput. Vis.*, 85(3):211–222.
- Hadamard, J. (1902). Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin*, 13:49–52.
- Hagmann, P., Cammoun, L., Gigandet, X., Gerhard, S., Ellen Grant, P., Wedeen, V., Meuli, R., Thiran, J.-P., Honey, C. J., and Sporns, O. (2010). MR connectomics: Principles and challenges. *Journal of neuroscience methods*, 194(1):45–34.
- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., and Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS biology*, 6(7):e159.
- Hagmann, P., Cammoun, L., Martuzzi, R., Maeder, P., Clarke, S., Thiran, J. P., and Meuli, R. (2006a). Hand preference and sex shape the architecture of language networks. *Hum. Brain Mapp.*, 27(10):828–835.

- Hagmann, P., Jonasson, L., Maeder, P., Thiran, J.-P., Wedeen, V. J., and Meuli, R. (2006b). Understanding diffusion MR imaging techniques: from scalar diffusion-weighted imaging to diffusion tensor imaging and beyond. *Radiographics : a review publication of the Radiological Society of North America, Inc*, 26 Suppl 1(suppl 1):S205–23.
- Hagmann, P., Thiran, J. P., Jonasson, L., Vandergheynst, P., Clarke, S., Maeder, P., and Meuli, R. (2003). DTI mapping of human brain connectivity: statistical fibre tracking and virtual dissection. *Neuroimage*, 19(3):545–554.
- Hamann, B. (1993). Curvature approximation for triangulated surfaces. *Computing*, 8:139–153.
- Harrison, L. M., Penny, W., Ashburner, J., Trujillo-Barreto, N., and Friston, K. J. (2007). Diffusion-based spatial priors for imaging. *NeuroImage*, 38(4):677–695.
- Harrison, L. M., Penny, W., Daunizeau, J., and Friston, K. J. (2008). Diffusion-based spatial priors for functional magnetic resonance images. *NeuroImage*, 41(2):408–23.
- Hebert, M., Ikeuchi, K., and Delingette, H. (1995). A Spherical Representation for Recognition of Free-form Surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(7):681–690.
- Hellier, P., Barillot, C., Corouge, I., Gibaud, B., Le Goualher, G., Collins, D. L., Evans, A., Malandain, G., Ayache, N., Christensen, G. E., and Johnsson, H. J. (2003). Retrospective evaluation of intersubject brain registration. *IEEE Transactions on Medical Imaging*, 22(9):1120–1130.
- Hermosillo, G., Chefd’hotel, C., and Faugeras, O. (2002). Variational Methods for Multimodal Image Matching. *International Journal of Computer Vision*, 50(3):329–343.
- Hildreth, E. C. (1984). Computations underlying the measurement of visual motion. *Artificial Intelligence*, 23(3):309–354.
- Hildreth, E. C. and Koch, C. (1987). The analysis of visual motion: from computational theory to neuronal mechanisms. *Annual review of neuroscience*, 10:477–533.
- Honey, C. J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J. P., Meuli, R., and Hagmann, P. (2009). Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences of the United States of America*, 106(6):2035–40.

- Hong, H. and Park, I. K. (2010). Single Image Motion Deblurring Using Anisotropic Regularization. *IEEE International Conference on Image Processing ICIP 2010*, pages 1149–1152.
- Horn, B. K. P. and Schunck, B. G. (1981). Determining optical flow. *Artificial Intelligence*, 17:185–203.
- Houhou, N., Thiran, J.-P., and Bresson, X. (2009). Fast Texture Segmentation Based on Semi-Local Region Descriptor and Active Contour. *Numer. Math. Theor. Meth. Appl.*, 2(4):445–468.
- Janenko, N. (1971). *The method of fractional steps*. Springer.
- Jbilou, K. and Sadok, H. (2000). Vector extrapolation methods. Applications and numerical comparison. *Journal of Computational and Applied Mathematics*, 122:149–165.
- Joshi, A. A., Shattuck, D. W., Thompson, P. M., and Leahy, R. M. (2004). Cortical Surface Parameterization by p-Harmonic Energy Minimization. In *ISBI 2004*, volume 1, pages 428–431.
- Joshi, A. A., Shattuck, D. W., Thompson, P. M., and Leahy, R. M. (2009). A parameterization-based numerical method for isotropic and anisotropic diffusion smoothing on non-flat surfaces. *IEEE Transactions on Image Processing*, 18(6):1358–1365.
- Joshi, S. C. and Miller, M. I. (2000). Landmark matching via large deformation diffeomorphisms. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, 9(8):1357–70.
- Kaftory, R., Sochen, N. A., and Zeevi, Y. Y. (2005). Variational Blind Deconvolution of Multi-Channel Images. *Int. J. Imaging Syst. Technol.*, 15:56–63.
- Kass, M., Witkin, A., and Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331.
- Kimmel, R. (2002). 3D Shape Reconstruction from Autostereograms and Stereo. *Journal of Visual Communication and Image Representation*, 13:324–333.
- Kimmel, R., Malladi, R., and Sochen, N. (2000). Images as Embedded Maps and Minimal Surfaces: Movies, Color, Texture, and Volumetric Medical Images. *Int. J. Comput. Vis.*, 39(2):111–129.
- Kimmel, R., Sochen, N., and Malladi, R. (1997). From high energy physics to low level vision. *Scale-Space Theory in Computer Vision*, 1252:236–247.
- Kirkpatrick, D. (1983). Optimal Search in Planar Subdivisions. *SIAM J. Comput.*, 12(1):28–35.

- Koenderink, J. J. (1984). The structure of images. *Biol. Cybern.*, 50(5):363–370.
- Kybic, J. and Unser, M. (2003). Fast parametric elastic image registration. *IEEE Transactions on Image Processing*, 12(11):1427–1442.
- Lam, M. C. H. (2011). Edge-Preserving Multiplicative Regularization with Weighted Total Variation.
- Lancaster, J. L., Glass, T. G., Lankipalli, B. R., Downs, H., Mayberg, H., and Fox, P. T. (1995). A modality-independent approach to spatial normalization of tomographic images of the human brain. *Human Brain Mapping*, 3(3):209–223.
- Lester, H. and Arridge, S. (1999). A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32(1):129–149.
- Li, X., Long, X., Wyatt, C. L., and ADNI (2011). Registration of Images with Topological Change via Riemannian Embedding. In *ISBI 2011*, pages 1247–1252.
- Lions, P. L. and Mercier, B. (1979). Splitting Algorithms for the Sum of Two Nonlinear Operators. *SIAM J. Numer. Anal.*, 16(6):964–979.
- Liu, T., Shen, D., and Davatzikos, C. (2004). Deformable registration of cortical structures via hybrid volumetric and surface warping. *Neuroimage*, 22(4):1790–1801.
- Lohmann, G. and von Cramon, D. Y. (2000). Automatic labelling of the human cortical surface using sulcal basins. *Medical image analysis*, 4(3):179–88.
- Lohmann, G., von Cramon, D. Y., and Colchester, A. C. F. (2008). Deep sulcal landmarks provide an organizing framework for human cortical folding. *Cerebral cortex*, 18(6):1415–1420.
- Lohmann, G., von Cramon, D. Y., and Steinmetz, H. (1999). Sulcal Variability of Twins. *Cerebral Cortex*, 9(7):754–763.
- Lopez-Perez, L., Deriche, R., and Sochen, N. (2004). The Beltrami Flow over Triangulated Manifolds. In *CVAMIA and MMBIA Workshop at ECCV 2004*, pages 135–144.
- Lou, Y., Zhang, X., Osher, S., and Bertozzi, A. (2009). Image Recovery via Nonlocal Operators. *Journal of Scientific Computing*, 42(2):185–197.
- Lu, T., Neittaanmaki, P., and Tai, X.-C. (1992). A parallel splitting up method for partial differential equations and its application to Navier-Stokes equations. *RAIRO Mathematical Modelling and Numerical Analysis*, 26(6):673–708.

- Lucas, B. D. and Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. In *International Joint Conference on Artificial Intelligence*, pages 674–679.
- Lynall, M.-E., Bassett, D. S., Kerwin, R., McKenna, P. J., Kitzbichler, M., Muller, U., and Bullmore, E. (2010). Functional Connectivity and Brain Networks in Schizophrenia. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(28):9477–9487.
- MacDonald, D., Kabani, N., Avis, D., and Evans, A. C. (2000). Automated 3D extraction of inner and outer surfaces of cerebral cortex from MRI. *Neuroimage*, 12(3):340–356.
- Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., and Suetens, P. (1997). Multimodality image registration by maximization of mutual information. *IEEE transactions on medical imaging*, 16(2):187–98.
- Maintz, J. B. A. and Viergever, M. A. (1998). A survey of medical image registration. *Med. Image Anal.*, 2(1):1–36.
- Malladi, R. and Ravve, I. (2002). Fast Difference Schemes for Edge Enhancing Beltrami Flow. In *ECCV 2002*, volume 2350 of *Lecture Notes in Computer Science*, pages 343–357.
- Marr, D. (1982). *Vision*. W. H. Freeman and Company.
- Marr, D. and Hildreth, E. (1980). Theory of Edge Detection. *Proc. Roy. Soc. Lon. Series B, Biol. Sci.*, 207(1167):187–217.
- Mattes, D. (2001). *Nonrigid multimodality image registration*, volume 4322. SPIE.
- Metzger, W. (1975). *Gesetze des Sehens*. W. Kramer & Co., Frankfurt am Main, 3rd edition.
- Meyer, M., Desbrun, M., Schröder, P., and Barr, A. H. (2002). Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III*, pages 35–57. Springer-Verlag.
- Miga, M. I., Sinha, T. K., Cash, D. M., Galloway, R. L., and Weil, R. J. (2003). Cortical surface registration for image-guided neurosurgery using laser-range scanning. *IEEE Transactions on Medical Imaging*, 22(8):973–985.
- Modersitzki, J. (2004). *Numerical Methods for Image Registration*. Oxford Univ. Press.
- Morozov, V. A. (1975). Linear and nonlinear ill-posed problems. *Journal of Mathematical Sciences*, II(6):706–736.

- Morse, B. S. (1994). *Computation of Object Cores from Grey-Level Images*. PhD thesis, University of North Carolina, Chapel Hill.
- Movshon, J. A., Adelson, E. H., Gizzi, M. S., and Newsome, W. T. (1985). The Analysis of Moving Visual Patterns. *PONTIFICIAE ACADEMIAE SCIENTIARVM SCRIPTA VARIA*, 54:117–151.
- Mücke, E. P., Saias, I., and Zhu, B. (1996). Fast randomized point location without preprocessing in two- and three-dimensional Delaunay triangulations. In *SCG '96: Proceedings of the twelfth annual symposium on Computational geometry*, pages 274–283.
- Nagel, H.-H. and Enkelmann, W. (1986). An Investigation of Smoothness Constraints for the Estimation of Displacement Vector Fields from Image Sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(5):565–593.
- Narayanan, R., Fessler, J. A., Park, H., and Meyer, C. R. (2005). Diffeomorphic Nonlinear Transformations: A Local Parametric Approach for Image Registration. In *Information Processing in Medical Imaging*, pages 174–185.
- Newsome, W. T., Britten, K. H., and Movshon, J. A. (1989). Neuronal correlates of a perceptual decision. *Nature*, 341(6237):52–54.
- Newsome, W. T. and Paré, E. B. (1988). A Selective Impairment of Motion Perception Following Lesions of the Middle Temporal Visual Area (MT). *The Journal of Neuroscience*, 8(6):2201–2211.
- Nieuwenhuys, R., Voogd, J., and Van Huijzen, C. (2008). *The Human Central Nervous System*. Springer-Verlag, Berlin Heidelberg New York, 4th edition.
- Nocedal, J. and Wright, S. J. (2006). *Numerical optimization*. Springer, Berlin, 2nd edition.
- Nordström, N. K. (1990). Biased anisotropic diffusion: a unified regularization and diffusion approach to edge detection. *Image and Vision Computing*, 8(4):318–327.
- Orozco Rodríguez, J. A. (2011). *Regularization Methods for Inverse Problems*. PhD thesis, University of Minnesota.
- Osher, S. and Sethian, J. A. (1988). Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49.
- Pei, Y. C., Hsiao, S. S., and Bensmaia, S. J. (2008). The tactile integration of local motion cues is analogous to its visual counterpart. *Proceedings of the National Academy of Sciences of the United States of America*, 105(23):8130–5.

- Perona, P. and Malik, J. (1990). Scale-Space and Edge Detection Using Anisotropic Diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639.
- Peyré, G., Bougleux, S., and Cohen, L. D. (2008). Non-local Regularization of Inverse Problems. In *ECCV 2008*, pages 57–68.
- Pitiot, A. and Guimond, A. (2008). Geometrical regularization of displacement fields for histological image registration. *Med. Image Anal.*, 12(1):16–25.
- Pluim, J. P. W., Maintz, J. B. A., and Viergever, M. A. (2003). Mutual-information-based registration of medical images: a survey. *IEEE transactions on medical imaging*, 22(8):986–1004.
- Pock, T., Urschler, M., Zach, C., Beichel, R., and Bischof, H. (2007). A Duality Based Algorithm for TV-L1-Optical-Flow Image Registration. In *MICCAI 2007*, pages 511–518.
- Poggio, T., Torre, V., and Koch, C. (1985). Computational vision and regularization theory. *Nature*, 317(6035):314–319.
- Polyakov, A. M. (1981). Quantum geometry of bosonic strings. *Phys. Lett. B*, 103(3):207–210.
- Pratt, W. K. (2001). *Digital Image Processing, PIKS inside*. John Wiley.
- Reuter, M., Rosas, H. D., and Fischl, B. (2010). Highly accurate inverse consistent registration: A robust approach. *Neuroimage*, 53(4):1181–1196.
- Rivière, D., Mangin, J.-F., Papadopoulos-Orfanos, D., Martinez, J.-M., Frouin, V., and Régis, J. (2002). Automatic recognition of cortical sulci of the human brain using a congregation of neural networks. *Medical image analysis*, 6(2):77–92.
- Robbins, S., Evans, A. C., Collins, D. L., and Whitesides, S. (2004). Tuning and comparing spatial normalization methods. *Med. Image Anal.*, 8(3):311–323.
- Rogelj, P. and Kovacic, S. (2006). Symmetric image registration. *Med. Image Anal.*, 10(3):484–493.
- Rogelj, P., Kovacic, S., and Gee, J. C. (2003). Point similarity measures for non-rigid registration of multi-modal data. *Computer Vision and Image Understanding*, 92(1):112–140.
- Rosman, G., Dascal, L., Sidi, A., and Kimmel, R. (2009). Efficient Beltrami Image Filtering via Vector Extrapolation Methods. *SIAM J. Imaging Sci.*, 2(3):858–878.

- Rosman, G., Tai, X.-C., Dascal, L., and Kimmel, R. (2010). Augmented Lagrangian for Polyakov Action Minimization in Color Images. In *ICNAAM 2010: International Conference of Numerical Analysis and Applied Mathematics 2010*, volume 1281, pages 1018–1021.
- Rudin, L. and Osher, S. (1994). Total Variation based Image Restoration with Free Local Constraints. *Proc. IEEE Int. Conf. Image Process.*, 1:31–35.
- Rudin, L., Osher, S., and Fatemi, E. (1992). Nonlinear Total Variation based Noise Removal Algorithms. *Phys. D*, 60:259–268.
- Rueckert, D., Aljabar, P., Heckemann, R. A., Hajnal, J. V., and Hammers, A. (2006). Diffeomorphic Registration Using B-Splines. In *MICCAI*, volume 4191 of *Lecture Notes in Computer Science*, pages 702–709, Berlin, Heidelberg. Springer.
- Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L., Leach, M. O., and Hawkes, D. J. (1999). Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721.
- Sabatini, S., Solari, F., and Bisio, G. (2003). Lattice Models for Context-Driven Regularization in Motion Perception. In *Neural Nets*, volume 2859 of *Lecture Notes in Computer Science*, pages 35–42. Springer Berlin / Heidelberg.
- Sabuncu, M., Yeo, B. T., Van Leemput, K., Fischl, B., and Golland, P. (2010). A Generative Model for Image Segmentation Based on Label Fusion. *IEEE transactions on medical imaging*, 29(10):1714–1729.
- Sagiv, C., Sochen, N. A., and Zeevi, Y. Y. (2001). Gabor Feature Space Diffusion via the Minimal Weighted Area Method. In *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 621–635.
- Sagiv, C., Sochen, N. A., and Zeevi, Y. Y. (2006). Integrated active contours for texture segmentation. *IEEE Transactions on Image Processing*, 15(6):1633–1646.
- Sarnak, N. and Tarjan, R. E. (1986). Planar point location using persistent search trees. *Commun. ACM*, 29(7):669–679.
- Sarti, A., Malladi, R., and Sethian, J. A. (2002). Subjective Surfaces : A Geometric Model for Boundary Completion. *International Journal of Computer Vision*, 46(3):201–221.
- Schaer, M., Cuadra, M. B., Tamarit, L., Lazeyras, F., Eliez, S., and Thiran, J.-P. (2008). A Surface-Based Approach to Quantify Local Cortical Gyrfication. *IEEE Transactions on Medical Imaging*, 27(2):161–170.

- Scharstein, D. and Szeliski, R. (2002). A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comput. Vis.*, 47(1):7–42.
- Scherzer, O. and Weickert, J. (2000). Relations Between Regularization and Diffusion Filtering. *Journal of Mathematical Imaging and Vision*, 12:43–63.
- Sdika, M. and Pelletier, D. (2009). Nonrigid Registration of Multiple Sclerosis Brain Images Using Lesion Inpainting for Morphometry or Lesion Mapping. *Human Brain Mapping*, 30(4):1060–1067.
- Sidi, A. (1991). Efficient implementation of minimal polynomial and reduced rank extrapolation methods. *Journal of Computational and Applied Mathematics*, 36(3):305–337.
- Sochen, N., Deriche, R., and Lopez-Perez, L. (2003a). The Beltrami flow over implicit manifolds. In *9th IEEE ICCV*, pages 832–839.
- Sochen, N., Deriche, R., and Lopez-Perez, L. (2003b). The Beltrami Flow over Manifolds. Technical Report 4897, INRIA.
- Sochen, N., Deriche, R., and Lopez-Perez, L. (2003c). Variational Beltrami flows over manifolds. In *IEEE ICIP*, volume 1, pages 861–864.
- Sochen, N., Kimmel, R., and Bruckstein, A. M. (2001). Diffusions and confusions in signal and image processing. *Journal of Mathematical Imaging and Vision*, 14(3):195–209.
- Sochen, N., Kimmel, R., and Malladi, R. (1998). A general framework for low level vision. *IEEE Transactions on Image Processing*, 7(3):310–318.
- Spira, A., Kimmel, R., and Sochen, N. (2007). A Short-Time Beltrami Kernel for Smoothing Images and Manifolds. *IEEE Transactions on Image Processing*, 16(6):1628–36.
- Stumpf, P. (1911). Über die Abhängigkeit der visuellen Bewegungsempfindung und ihres negativen Nachbildes von den Reizvorgängen auf der Netzhaut. *Zeitschrift für Psychologie*, 59:321–330.
- Sun, D., Roth, S., and Black, M. J. (2010). Secrets of Optical Flow Estimation and Their Principles. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 2432–2439.
- Sundareswara, R. and Schrater, P. R. (2003). Extensible Point Location Algorithm. In *GMAG*, pages 84–89.
- Tai, X.-C. and Wu, C. (2009). Augmented Lagrangian Method, Dual Methods and Split Bregman Iteration for ROF Model. In *Scale Space and Variational Methods in Computer Vision*, pages 502–513.

- Tammes, P. M. L. (1930). On the origin of number and arrangement of places of exit on the surface of pollen grains. *Recueil des Travaux Botanique Neerlandais*, 27:1–84.
- Tang, L., Hamarneh, G., and Abugharbieh, R. (2010). Reliability-Driven, Spatially-Adaptive Regularization for Deformable Registration. In Fischer, B., Dawant, B., and Lorenz, C., editors, *WBIR 2010*, volume 6204 of *Lecture Notes in Computer Science*, pages 173–185.
- Tao, X., Prince, J. L., and Davatzikos, C. (2002). Using a statistical shape model to extract sulcal curves on the outer cortex of the human brain. *IEEE transactions on medical imaging*, 21(5):513–24.
- Taubin, G. (1995). A signal processing approach to fair surface design. In *Proc. ACM SIGGRAPH '95*, pages 351–358.
- Thévenaz, P. and Unser, M. (2000). Optimization of mutual information for multiresolution image registration. *IEEE Transactions on Image Processing*, 9(12):2083–2099.
- Thirion, J.-P. (1998). Image matching as a diffusion process: an analogy with Maxwell’s demons. *Medical Image Analysis*, 2(3):243–260.
- Thompson, D. W. (1917). *On Growth and Form*. Cambridge Univ. Press, Cambridge.
- Thompson, P. and Toga, A. W. (1996). A Surface-Based Technique for Warping Three-Dimensional Images of the Brain. *IEEE Transactions on Medical Imaging*, 15(4):402–417.
- Thompson, P. and Toga, A. W. (1999). Anatomically-Driven Strategies for High-Dimensional Brain Image Warping and Pathology Detection. In Toga, A. W., editor, *Brain Warping*. Academic Press.
- Thomson, J. J. (1904). On the Structure of the Atom: an Investigation of the Stability and Periods of Oscillation of a number of Corpuscles arranged at equal intervals around the Circumference of a Circle; with Application of the Results to the Theory of Atomic Structure. *Philosophical Magazine Series 6*, 7(39):237–265.
- Tichonov, A. N. (1963). Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics*, 4:1035–1038.
- Todorovic, D. (1996). A gem from the past: Pleikart Stumpf’s (1911) anticipation of the aperture problem, Reichardt detectors, and perceived motion loss at equiluminance. *Perception*, 25(10):1235–1242.
- Toga, A. W. (1999). *Brain Warping*. Academic Press.

- Toga, A. W. and Thompson, P. M. (2001). The role of image registration in brain mapping. *Image and vision computing*, 19(1-2):3–24.
- Tomasi, C. and Manduchi, R. (1998). Bilateral filtering for gray and color images. In *ICCV 1998*, pages 839–846. Narosa Publishing House.
- Toro, R. and Burnod, Y. (2005). A morphogenetic model for the development of cortical convolutions. *Cerebral cortex*, 15(12):1900–1913.
- Tosic, I., Bogdanova, I., Frossard, P., and Vanderghelynst, P. (2005). Multiresolution motion estimation for omnidirectional images. In *Proceedings of 13th EUSIPCO*.
- Tosun, D., Rettmann, M. E., and Prince, J. L. (2004). Mapping techniques for aligning sulci across multiple brains. *Med. Image Anal.*, 8(3):295–309.
- Tu, Z., Zheng, S., Yuille, A. L., Reiss, A. L., Dutton, R. a., Lee, A. D., Galaburda, A. M., Dinov, I., Thompson, P. M., and Toga, A. W. (2007). Automated extraction of the cortical sulci based on a supervised learning approach. *IEEE transactions on medical imaging*, 26(4):541–52.
- Vaillant, M. and Davatzikos, C. (1997). Finding parametric representations of the cortical sulci using an active contour model. *Med. Image Anal.*, 1(4):295–315.
- Van Essen, D. C. (1997). A tension-based theory of morphogenesis and compact wiring in the central nervous system. *Nature*, 385(6614):313–8.
- Van Essen, D. C. (1998). Functional and structural mapping of human cerebral cortex: Solutions are in the surfaces. *Proceedings of the National Academy of Sciences*, 95(3):788–795.
- Van Essen, D. C., Drury, H. A., Joshi, S., and Miller, M. I. (2000). Functional and structural mapping of human cerebral cortex: solutions are in the surfaces. *Adv. Neurol.*, 84:23–34.
- Vemuri, B. C., Ye, J., Chen, Y., and Leonard, C. M. (2003). Image registration via level-set motion: Applications to atlas-based segmentation. *Medical Image Analysis*, 7:1–20.
- Vercauteren, T., Pennec, X., Perchant, A., and Ayache, N. (2009). Diffeomorphic demons: Efficient non-parametric image registration. *Neuroimage*, 45(1, Supplement 1):S61–S72.
- Viceic, D. (2007). *Postlesional Plasticity in central auditory processing: MRI and fMRI image registration and statistical signal processing*. PhD thesis, EPFL.

- Viceic, D., Fornari, E., Thiran, J. P., Maeder, P. P., Meuli, R., Adriani, M., and Clarke, S. (2006). Human auditory belt areas specialized in sound recognition: a functional magnetic resonance imaging study. *Neuroreport*, 17(16):1659–1662.
- Viola, P. and Wells III, W. M. (1997). Alignment by Maximization of Mutual Information. *Int. J. Comput. Vis.*, 24(2):137–154.
- von Ehrenfels, C. (1890). Über ‘Gestaltqualitäten’. *Vierteljahrsschrift für wissenschaftliche Philosophie*, 14:249–292.
- Weickert, J. (1997). A review of nonlinear diffusion filtering. In *Scale-space theory in computer vision*, volume 1252 of *Lecture Notes in Computer Science*, pages 3–27, Berlin, Germany. Springer-Verlag.
- Weickert, J. (1999). Coherence-Enhancing Diffusion Filtering. *International Journal of Computer Vision*, 31(2):111–127.
- Weickert, J. and Schnörr, C. (2001). A Theoretical Framework for Convex Regularizers in PDE-Based Computation of Image Motion. *International Journal of Computer Vision*, 45(3):245–264.
- Weickert, J., ter Haar Romeny, B. M., and Viergever, M. A. (1998). Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Transactions on Image Processing*, 7(3):398–410.
- Wells III, W. M., Viola, P., Atsumi, H., Nakajima, S., and Kikinis, R. (1996). Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–51.
- Wertheimer, M. (1912). Experimentelle Studien über das Sehen von Bewegung. *Zeitschrift für Psychologie*, 61(1):161–265.
- West, J., Fitzpatrick, J., Wang, M., Dawant, B., Maurer Jr, C., Kessler, R., Maciunas, R., Barillot, C., Lemoine, D., Collignon, A., and Al, E. (1997). Comparison and evaluation of retrospective intermodality brain image registration techniques. *Journal of Computer Assisted Tomography*, 21(4):554.
- Witkin, A. P. (1983). Scale-space filtering. *Int. Joint Conf. AI*, 2:329–332.
- Womersley, R. S. and Sloan, I. H. (2001). How Good can Polynomial Interpolation on the Sphere be? *Advances in Computational Mathematics*, 14:195–226.
- Wu, Y., He, Y., and Tian, H. (2005). A spherical point location algorithm based on barycentric coordinates. In *Lecture Notes in Computer Science*, volume 3482, pages 1099–1108.

- Wyatt, C. L. and Laurienti, P. J. (2006). Nonrigid registration of images with different topologies using embedded maps. In *IEEE Engineering in Medicine and Biology Conference*, volume 1, pages 4823–4827.
- Wyatt, C. L., Li, X., and Gong, X. (2009). A Framework for Registration of Images with Varying Topology using Embedded Maps : Reimannian Embedding Spaces. Technical Report c, Virginia Polytechnic Institute and State University.
- Xu, G. (2004). Convergent Discrete Laplace-Beltrami Operators over Triangular Surfaces. In *Proceedings of the Geometric Modeling and Processing 2004 (GMP'04)*.
- Xu, G. (2006). Discrete Laplace-Beltrami operator on sphere and optimal spherical triangulations. *International Journal of Computational Geometry and Applications*, 16(1):75–93.
- Yagi, Y., Kawato, S., and Tsuji, S. (1994). Real-time omnidirectional image sensor (COPIS) for vision-guided navigation. *IEEE Transactions on Robotics and Automation*, 10(1):11–22.
- Yeo, B., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., and Golland, P. (2008). Spherical Demons: Fast Surface Registration. In *MICCAI 2008*, pages 745–753.
- Yeo, B. T. T., Sabuncu, M. R., Vercauteren, T., Ayache, N., Fischl, B., and Golland, P. (2010). Spherical Demons: Fast Diffeomorphic Landmark-Free Surface Registration. *IEEE Transactions on Medical Imaging*, 29(3):650–668.
- You, Y.-L. and Kaveh, M. (1999). Blind Image Restoration by Anisotropic Regularization. *IEEE Transactions on Image Processing*, 8(3):396–407.
- Zach, C., Pock, T., and Bischof, H. (2007). A Duality Based Approach for Realtime TV-L1 Optical Flow. In *Pattern Recognition*, volume 4713 of *Lecture Notes in Computer Science*, pages 214–223, Berlin, Heidelberg. Springer.
- Zilles, K., Armstrong, E., Schleicher, A., and Kretschmann, H.-J. (1988). The human pattern of gyrification in the cerebral cortex. *Anatomy and Embryology*, 179(2):173–179.
- Zitová, B., Flusser, J., and Zitova, B. (2003). Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000.
- Zosso, D., Bresson, X., and Thiran, J.-P. (2011a). Fast Geodesic Active Fields - A Splitting Scheme for Geometric Image Registration. *IEEE Transactions on Image Processing*.

- Zosso, D., Bresson, X., and Thiran, J.-P. (2011b). Geodesic Active Fields - A Geometric Framework for Image Registration. *IEEE Transactions on Image Processing*, 20(5):1300–1312.
- Zosso, D., Noirhomme, Q., Davare, M., Macq, B., Olivier, E., Thiran, J.-P., and De Craene, M. (2006). Normalization of Transcranial Magnetic Stimulation Points by Means of Atlas Registration. In *EUSIPCO 2006*.
- Zosso, D. and Thiran, J.-P. (2009). A scale-space of cortical feature maps. *IEEE Signal Processing Letters*, 16(10):873–876.
- Zosso, D. and Thiran, J.-P. (2010). Geodesic Active Fields on the Sphere. In *International Conference on Pattern Recognition, ICPR2010*.

Curriculum Vitae

Dominique ZOSSO

Born 23.05.1983 in Bern, Switzerland (28 years old)

Common law marriage, 2 children

Contact Information

Professional	EPFL STI IEL LTS5 Station 11, ELD 232 CH-1015 Lausanne	Private	Route de Chavannes 68 CH-1007 Lausanne
--------------	--	---------	---

phone:	+41 21 693 46 22	phone:	+41 21 624 62 51
--------	------------------	--------	------------------

fax:	+41 21 693 76 00	mobile:	+41 79 427 48 80
------	------------------	---------	------------------

email: dominique.zosso@epfl.ch

web: <http://people.epfl.ch/zosso>

Education and Employments

2007 - 2011 **PhD** (Dr. ès sciences)
École Polytechnique Fédérale de Lausanne (EPFL), Switzerland.
Signal Processing Laboratory (LTS5), Prof. J.-Ph. Thiran.

2006 - 2007 **Research Assistant**
Biozentrum University of Basel, Switzerland. Structural Bioinformatics Group and Swiss Institute of Bioinformatics (SIB), Prof. T. Schwede & Dr. M. Podvinec.

2001 - 2006 **MSc. Electrical and Electronics Eng.** (Ing. él. dipl. EPF)
École Polytechnique Fédérale de Lausanne (EPFL), Switzerland.

2005 - 2006 **Master thesis, Erasmus exchange**
Université catholique de Louvain, Louvain-la-Neuve, Belgium.
Telecommunications and Remote Sensing Laboratory (TELE), Prof. B. Macq & Dr. M. Decraene.

1997 - 2001 **Matura**
Mathematisch Naturwissenschaftliches Gymnasium Bern-Kirchenfeld, Bern, Switzerland.

Achievements and Awards

- | | |
|------|---|
| 2008 | 8th IEEE EMBS International Summer School on Biomedical Imaging, Berder, France |
| 2006 | Chavannes Municipality Award
(“Prix de la Commune de Chavannes”) |
| 2001 | Meyer Award for scientific excellence
(“Meyer-Preis”) |
| 2001 | Award of the High School Association
(“Preis des Gymnervereins”) |

Funding

- | | |
|-------------|---|
| 2011 - 2013 | Fellowship for prospective researchers (PostDoc at UCLA)
Swiss National Science Foundation (SNF) |
| 2007 - 2011 | PhD student fellowship
National Competence Center in Biomedical Imaging (NCCBI) |

Talks and Seminars

- | | |
|------|--|
| 2010 | NCCBI Meeting, EPFL, Lausanne Switzerland |
| 2010 | Math Dept. (Prof. L. Vese), UCLA, Los Angeles, CA, USA |
| 2010 | Radiology Dept. (Prof. S. Warfield), Children’s Hospital, Boston, MA, USA |
| 2010 | Martinos Center for Biomedical Imaging (Prof. B. Fischl), MGH, Boston, MA, USA |

Publications

10+ research papers in peer-reviewed journals and conference proceedings. See separate full list of publications.

Reviewer for: IEEE Transactions on Image Processing, IEEE Signal Processing Letters, IEEE Journal of Display Technology, Computer Vision and Image Understanding, and the IEEE conferences ICIP, ISBI and EMBC.

Transferable and Other Skills

Tutoring: formal training and practical experience as undergraduate and graduate Teaching Assistant in various EPFL courses, including “Image Processing” and “Advanced Image Processing and Analysis”.

Programming and modelling languages: C/C++, Matlab, Java, Pascal, Perl, PHP, mySQL.

Military training: Conscript Sergeant (squad leader and platoon sergeant) in a reconnaissance battalion, Head of Company HQ, responsible for various technical and tactical training.

List of Publications

Journals

- 2011 **D. Zosso**, X. Bresson, J.-Ph. Thiran, *Fast Geodesic Active Fields for Image Registration based on Splitting and Augmented Lagrangian Approaches*, IEEE Transactions on Image Processing (submitted).
- . V. Estellers, **D. Zosso**, R. Lai, J.-Ph. Thiran, S. Osher, X. Bresson, *An Efficient Algorithm for Level Set Method Preserving Distance Function*, IEEE Transactions on Image Processing (submitted).
- . **D. Zosso**, X. Bresson, J.-Ph. Thiran, *Geodesic Active Fields – A Geometric Framework for Image Registration*, IEEE Transactions on Image Processing 20(5):1300-1312, 2011.
- 2009 **D. Zosso**, J.-Ph. Thiran, *A Scale-Space of Cortical Feature Maps*, IEEE Signal Processing Letters 16(10):873-876, 2009.
- 2007 **D. Zosso**, M. Bach Cuadra, J.-Ph. Thiran, *Direct Fourier Tomographic Reconstruction Image-To-Image Filter*, The Insight Journal, vol. July-December 2007.

Conference Proceedings

- 2011 V. Estellers, **D. Zosso**, X. Bresson, J.-Ph. Thiran, *Harmonic Active Contours for multichannel image segmentation*, 18th IEEE International Conference on Image Processing ICIP2011, accepted.
- 2010 **D. Zosso**, X. Bresson, J.-Ph. Thiran, *Geodesic Active Fields – A Geometric Framework for Image Registration*, 8th International Conference of Numerical Analysis and Applied Mathematics ICNAAM2010, Rhodes, Greece, p.1027-1030.
- . **D. Zosso**, J.-Ph. Thiran, *Geodesic Active Fields on the Sphere*, 20th International Conference on Pattern Recognition ICPR2010, Istanbul, Turkey, p.4484-4487.

- 2008 D. Ferrario, M. Bach Cuadra, M. Schaer, N. Houhou, **D. Zosso** et al., *Brain Surface Segmentation of Magnetic Resonance Images of the Fetus*, 16th European Signal Processing Conference EUSIPCO2008, Lausanne, Switzerland.
- . **D. Zosso**, B. Le Calennec, M. Bach Cuadra, K. Aminian et al., *Bi-planar 2D-to-3D Registration in Fourier Domain for Stereoscopic X-Ray Motion Tracking*, SPIE Medical Imaging 2008, San Diego, CA, USA, vol. 6914.
- 2007 **D. Zosso**, K. Arnold, T. Schwede, M. Podvinec, *SWISS-TANDEM: A Web-based Workspace for MS/MS Protein Identification on PC Grids*, 20th IEEE International Symposium on Computer-Based Medical Systems CBMS'07, Maribor, Slovenia, p.406-411.
- . **D. Zosso**, M. Podvinec, M. Mller, R. Aebersold et al., *Tandem Mass Spectrometry Protein Identification on a PC Grid*, 5th HealthGrid Conference, Geneva, Switzerland. In: Studies in Health Technology and Informatics, vol. 126, p.3-12. (Paper Award Finalist).
- . **D. Zosso**, M. Podvinec, K. Arnold, T. Schwede, *SWISS-Tandem: A web-based workspace for MS/MS protein identification on a PC-grid*, Joint Basel Computational Biology Conference [BC]2 and USGEB2007, Basel, Switzerland.
- 2006 **D. Zosso**, Q. Noirhomme, M. Davare, B. Macq et al., *Normalization of Transcranial Magnetic Stimulation Points by Means of Atlas Registration*, 14th European Signal Processing Conference EUSIPCO2006, Florence, Italy.

Unpublished / In Preparation

Original Research Articles

D. Zosso, X. Bresson, J.-Ph. Thiran, *Fast Geodesic Active Fields on the Sphere for Registration of Cortical Feature Maps*.

D. Zosso, X. Bresson, J.-Ph. Thiran, *A Symmetric and Diffeomorphic Deformation Model for the Geodesic Active Fields Framework*.

Invited Reviews

D. Zosso, N. Sochen, R. Kimmel, J.-Ph. Thiran, X. Bresson, *Polyakov Variational Model: A Unified Geometric Framework for Image Processing*, Numerical Mathematics: Theory, Methods and Applications.