



HAL
open science

Efficient computational tools for the statistical analysis of shape and asymmetry of 3D point sets

Benoît Combès

► **To cite this version:**

Benoît Combès. Efficient computational tools for the statistical analysis of shape and asymmetry of 3D point sets. Signal and Image Processing. Université Rennes 1, 2010. English. NNT: . tel-00684994

HAL Id: tel-00684994

<https://theses.hal.science/tel-00684994>

Submitted on 3 Apr 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

N° d'ordre: 4185

THÈSE

présentée

devant l'Université de Rennes 1

pour obtenir

le grade de : DOCTEUR DE L'UNIVERSITÉ DE RENNES 1
Mention TRAITEMENT DU SIGNAL ET DES IMAGES

par

Benoît COMBÈS

Équipe d'accueil : Visages - IRISA

École Doctorale : Matisse

Composante universitaire : IFSIC

Titre de la thèse :

*Efficient computational tools for the statistical analysis
of shape and asymmetry
of 3D point sets*

À soutenir le jeudi 18 novembre 2010 devant la commission d'examen

M. :	Président	DU JURY	Président
MM. :	Grégoire	MALANDAIN	Rapporteurs
	Jean-François	MANGIN	
MM. :	Etienne	MÉMIN	Examineurs
	Antoine	BALZEAU	
	Neil	ROBERTS	
	Sylvain	PRIMA	

Table des matières

Table des matières	1
Introduction	7
I Estimating bilateral symmetries of 3D point sets	19
1 Introduction	21
1.1 Introduction	21
1.2 Notations & basic formulas	22
2 Bibliography	25
2.1 Approaches based on points distribution	25
2.1.1 Basics	25
2.1.2 Generalisation	26
2.1.3 Advantages and limitations	27
2.2 Approaches based on points superimposition	27
2.2.1 Basics	27
2.2.2 Optimal \tilde{P} with known correspondents (\tilde{x}'_i)	29
2.2.3 Minimisation strategies	30
2.2.4 Generalisation	32
2.2.5 Advantages and limitations	33
2.3 Approaches based on planes clustering	34
2.3.1 Basics	34
2.3.2 Clustering	34
2.3.3 Improvements	35
2.3.4 Advantages and limitations	36
2.4 Other methods	36
2.4.1 EGI	36
2.5 Conclusions	37
3 Symmetry plane estimation	39
3.1 Coarse detection of plausible reflections	39
3.1.1 Subsampling using keypoints	39

3.1.2	Matching keypoints and weighing pairs	40
3.1.3	Building, modifying and pruning the affinity matrix	41
3.1.3.1	Enforcing global constraint on C	41
3.1.3.2	Pruning C	41
3.1.4	Clustering mirror symmetries	42
3.1.4.1	Solution 1 : Euclidean distance on parameters	42
3.1.4.2	Solution 2 : Riemannian distance on angular components	42
3.1.4.3	Solution 3 : Point-to-plane distance and distance between angles	44
3.1.5	Checking the mirror symmetries	45
3.1.6	Results and evaluation	46
3.1.6.1	Parameters	46
3.1.6.2	Some observations	46
3.1.6.3	Some results	46
3.2	Accurate & robust estimation of a reflection	48
3.2.1	A fast, robust and accurate algorithm	48
3.2.1.1	Algorithm	48
3.2.1.2	Rejection of outliers	49
3.2.1.3	Multiscale scheme & estimation of σ	50
3.2.1.4	Coarse-to-fine approach	51
3.2.1.5	Algorithm	51
3.2.2	Symmetrisation of A	51
3.2.3	Results and evaluation	52
3.2.3.1	Parameters & algorithm	52
3.2.3.2	Evaluation on symmetrical data	53
3.2.3.3	Evaluation on asymmetrical data	53
3.2.3.4	Results on real data	55
4	Symmetry surface estimation	57
4.1	Extension to symmetry surface	57
4.1.1	Formulation	57
4.1.2	Initialisation	59
4.1.3	Measuring the curvature of C	60
4.1.4	Some results	60
5	Conclusion	61
5.1	Contributions	61
5.2	Limits & Perspectives	61
5.2.1	Coarse estimation	61
5.2.2	Accurate and robust estimation	61
5.2.3	Symmetry surface	62

II Non linear registration of point sets	63
6 Introduction	65
7 Bibliography	67
7.1 Introduction	67
7.2 Mixture models and point set registration	68
7.2.1 Data attachment term & Minimisation	69
7.2.1.1 KL-BASED METHODS : Kullback-Leibler divergence between a mixture of Gaussians and a mixture of Diracs : the ML approach	69
7.2.1.2 KL-BASED METHODS : Jensen-Shannon divergence between GMMs	71
7.2.1.3 L_2 -BASED METHODS : L_2 distance between 2 GMMs	73
7.2.1.4 L_2 -BASED METHODS : Correlation between GMMs :	75
7.2.2 Regularisation	76
7.2.2.1 The thin plate splines	76
7.2.2.2 The coherent point drift (CPD)	77
7.2.2.3 Diffeomorphisms	78
7.2.3 Some additions	80
7.2.3.1 Landmarks	80
7.2.3.2 Adding higher order information	80
7.2.4 A fundamental minimisation tool for GMMs : the fast kernel com- putation	81
7.2.4.1 Fast Gauss transform	81
7.2.4.2 Kd-tree subdivision of the space	81
7.3 Other approaches :	82
7.3.1 Spectral based correspondences (and extensions)	82
7.3.2 Surface as a level set function	82
7.3.3 Diffeomorphic matching of distributions	83
7.3.3.1 Matching of Dirac distributions	83
7.3.3.2 Currents	84
8 Methods	87
8.1 Non linear registration as a clustering problem	88
8.1.1 General formulation	88
8.1.2 Robustifying the estimator and reducing the computational burden	90
8.1.3 Energetic interpretation	91
8.2 Symmetrising the matching process	92
8.2.1 Minimisation with respect to T	93
8.3 Adding priors	95
8.3.1 General idea	95
8.3.2 Designing π	95
8.3.2.1 Designing π using labels	95

8.3.2.2	Designing π using descriptors	96
8.3.3	Implementation	96
8.3.4	Initialising T using π only	97
8.4	Symmetrising the estimation of T : A framework for a consistent & symmetric EM-ICP	100
8.4.1	Cost function	100
8.4.2	Minimisation	101
8.5	Efficient approximation techniques	103
8.5.1	Local models	103
8.5.1.1	Locally affine	103
8.5.1.2	Local translations	104
8.5.1.3	Discussion	105
8.5.2	Efficient global models	106
8.5.2.1	Formulation	106
8.5.2.2	Choosing a kernel	107
8.5.2.3	Efficient choices	109
8.5.2.4	M-step in a nutshell	109
8.5.2.5	Note on vectorial kernels	110
9	Evaluation	111
9.1	Descriptions	111
9.1.1	Descriptions of the algorithms and of the parameters	111
9.1.2	Description of the data	111
9.1.2.1	Thalamus	111
9.1.2.2	Caudate nuclei	111
9.1.2.3	Ventricles	111
9.1.2.4	Osseous labyrinths	112
9.1.2.5	Endocasts	112
9.1.3	Description of the experiments	112
9.1.3.1	Experiment 1 : MSE on simulated data	112
9.1.3.2	Experiment 2 : Landmark-to-landmark distance on real data	113
9.2	Results	114
9.2.1	Experiment 1	114
9.2.2	Experiment 2	114
9.2.3	Additional measurements	114
9.3	Summary	115
10	Conclusion	117
10.1	Contributions	117
10.2	Perspectives	118

III Applications	121
11 Tools for group studies & Application	123
11.1 Statistical Shape Analysis	123
11.1.1 Introduction & Brief state-of-the-art	123
11.1.2 Mean shape EM-ICP and extensions	124
11.1.2.1 Notations	124
11.1.2.2 The mean shape EM-ICP	124
11.1.2.3 First interpretation : the Fréchet-like mean	126
11.1.2.4 Second interpretation : Procrustean mean shape on un- labelled point sets	127
11.1.3 Virtual correspondence & projection	129
11.1.4 Comparison of populations	130
11.2 Quantification of asymmetries over/between populations	131
11.2.1 Computation of individual asymmetry maps	131
11.2.2 Computation of a mean cortical shape & projection of the asym- metry mappings	131
11.2.3 Computation of the symmetry surface	132
12 Asymmetries	133
12.1 Note on asymmetry maps :	133
12.2 Illustrating our pipeline : study on Situs inversus	134
12.2.1 Subjects and image acquisition	135
12.2.2 Preprocessing	135
12.2.3 Mean, standard deviation and significant asymmetries in controls	135
12.2.4 Computation of regional asymmetries	137
12.2.5 Individual asymmetry on SI subjects	137
12.2.6 Significant differences of asymmetry between controls and SI sub- jects	140
12.3 Results	142
12.3.1 Controls	142
12.3.2 Situs inversus	142
13 Conclusion	145
13.1 Contributions	145
13.2 Discussions	145
13.2.1 Mean shape EM-ICP strategies	145
13.2.2 Pipeline for asymmetry quantification	146
13.2.2.1 Alternative solutions	146
14 General conclusion	149
14.1 Ongoing studies on brain asymmetry	150
14.2 Other ongoing applications	151
14.3 List of publications	153

14.3.1	International conferences	153
14.3.2	International conferences : under submission	154
IV	Appendix	155
A	Closed-form solution for symmetry plane estimation	157
A.0.3	Principal axes	157
A.0.3.1	Fitting a plane in point set	157
A.0.3.2	Scatter matrix and inertia matrix	157
A.0.4	Optimal solution for symmetry plane estimation with known matchings	158
B	The EM algorithm	161
B.1	Formulation générale	161
B.2	L'EM et les GMMs	162
B.3	Discussion	163
	Bibliographie	177
	Table des figures	179

Introduction

Anatomie commune et variations anatomiques

L'idée d'exprimer l'ensemble des représentants d'une population donnée (espèce ou sous-groupe d'une espèce) comme la combinaison d'une anatomie commune et de variations autour de celle-ci est la clef de voûte de nombreux travaux méthodologiques, dont cette thèse fait partie, visant à l'élaboration de modèles statistiques de l'anatomie. Cette idée n'est pas nouvelle et trouve son origine dans certains développements théoriques de la définition des espèces et de leur histoire, dont voici quelques illustrations.

Au XVIIIème siècle, les naturalistes Carl von Linné et Georges-Louis Leclerc de Buffon définissent les espèces comme des groupes d'individus partageant des caractéristiques morphologiques communes et pouvant se reproduire. Le critère d'interfécondité ne peut pas toujours être vérifié : c'est le cas pour les fossiles, les organismes asexués ou pour des espèces difficiles à observer. L'étude de la variabilité morphologique d'une espèce devient alors un exercice primordial pour la définition des espèces et l'affectation d'un individu à une espèce.

Au XIXème siècle, la théorie de l'évolution développée par Charles Darwin et Alfred Wallace donne une interprétation et une perspective nouvelles à cette vision : l'anatomie commune dérive de celle d'un ancêtre commun et la variabilité résulte de la transmission d'innovations de génération en génération. De surcroît, ces innovations sont soumises à une sélection dite naturelle : les traits qui favorisent la survie et la reproduction voient leurs fréquences s'accroître d'une génération à l'autre.

Cependant, pour Darwin, l'hérédité se fait sur la base des caractères acquis et il faut attendre la fin du XIXème siècle pour que le darwinisme adopte le postulat du mutationnisme, d'après lequel les mutations génétiques (variations brusques et héréditaires) sont les matériaux de l'évolution sur lesquels s'exerce la sélection naturelle.

À cette idée de sélection de caractères survenant aléatoirement, s'est ajoutée celle introduite par D'Arcy Thompson selon laquelle les innovations possibles étaient contraintes par la structure des tissus biologiques et des forces auxquelles ils sont soumis. Ainsi, la palette des innovations possibles au sein d'une structure ne serait pas aussi large que celle que générerait un pur processus aléatoire, mais serait contenue par des contraintes physiques et biologiques.

Quels que soient les enjeux théoriques, l'idée suggérée est de considérer l'ensemble des anatomies définissant une espèce, ou un sous groupe de celle-ci, comme un ensemble de variations "contrôlées" autour d'une base commune.




FIG. 1 – **A gauche : Classification des formes de feuilles** tirés de “Hortus Cliffortianus” de Carl Von Linné. **A droite : Les pinsons de Darwin** L’isolation des Pinson dans les îles Galápagos produit environ une douzaine de nouvelles espèces. Les différences anatomiques les plus importantes entre ces espèces se trouvent dans la taille et la forme de leurs becs. Ces variations s’expliquent par les différences de nourriture de chaque espèce. Leurs comportements tout comme leurs chants sont différents.

La symétrie comme caractéristique anatomique d’intérêt

Les propriétés de symétrie associées à un organisme constituent une caractéristique anatomique particulièrement étudiée et une grande partie de cette thèse trouve son application dans l’étude des asymétries bilatérales de structures anatomiques.

La symétrie des organismes

La majeure partie des êtres vivants possède l’étonnante caractéristique d’être anatomiquement organisée autour d’une symétrie globale. En fait, les organisations selon des symétries radiales (*i.e.* par rapport à un axe) et bilatérales (*i.e.* par rapport à un plan) apparaissent comme des stratégies adoptées par les organismes pour optimiser leur relation à l’environnement. Les organisations selon une symétrie radiale concernent, par exemple, les organismes marins (comme le plancton) pour lesquels la possibilité de mouvement est limitée. Dans ce cas, afin d’optimiser leurs contacts avec l’environnement, les fonctions sensorielles voire métaboliques sont réparties régulièrement autour d’un axe de symétrie (la perte de la symétrie haut-bas étant due à la pesanteur). L’organisation du corps selon une symétrie bilatérale nous est plus familière et coïncide avec la faculté de se mouvoir. Cette symétrie est alors interprétée comme la perte de la symétrie antéro-postérieure présente chez les organismes à symétrie radiale due à la nécessité de concentrer les fonctions motrices à l’arrière et les fonctions sensorielles à l’avant du corps.




FIG. 2 – Quelques exemples des symétries du monde vivant

Fluctuation autour de la symétrie parfaite : les asymétries fluctuantes

Cependant cette symétrie n'est jamais parfaite et s'observe le plus souvent sous une forme approximative. On nomme asymétries fluctuantes les écarts à une symétrie idéalement parfaite. Celles-ci sont généralement interprétées comme des imperfections du développement de l'organisme. En particulier, la quantification de l'asymétrie fluctuante au sein d'une population est une mesure pertinente de la stabilité des individus dans des conditions données : en situation de stress environnementaux (malnutrition, parasites, *etc*) ou génétiques (mutations, consanguinité, *etc*), l'asymétrie fluctuante augmente. La symétrie serait aussi un critère de choix lors de la sélection sexuelle : les mâles les plus symétriques sont préférés par les femelles pour l'accouplement. Ainsi, ces dernières fournissent à leur descendance les traits d'un individu a priori plus apte à se développer normalement dans son environnement.

Asymétries intentionnelles : les asymétries directionnelles et les anti-symétries

Les asymétries précédentes évoquaient un écart à une symétrie supposée parfaite dans le cas idéal. Cependant, il est relativement commun d'observer des asymétries que l'on pourrait qualifier d'intentionnelles : dans le cas idéal, le trait est asymétrique. Ainsi, les asymétries directionnelles désignent des traits bilatéraux qui se développent de manière plus importante d'un côté que de l'autre. Dans la plupart des cas, ce type d'asymétrie est prédéterminé par le génotype de l'individu. Cette asymétrie concerne

un grand nombre de nos organes : le poumon droit a 3 lobes alors que le gauche en a 2, le rein gauche est plus élevé que le droit ou encore, comme nous allons le voir par la suite, les deux hémisphères du cerveau humain ne sont pas identiques.

L'antisymétrie concerne quant à elle les traits bilatéraux qui se développent de façon prédominante d'un côté de l'organisme et ce, aussi souvent à droite qu'à gauche. En général, aucun côté n'est favorisé *a priori* et l'action de l'environnement est déterminante (c'est le cas de la spécialisation des pinces chez le homard).

Pour résumer, on peut représenter les trois types de symétrie à l'échelle d'une population par des courbes où :

- l'asymétrie fluctuante est une cloche centrée en 0,
- l'asymétrie directionnelle est une cloche centrée en une valeur non nulle,
- l'antisymétrie est la somme de deux cloches centrées de part et d'autre de 0.

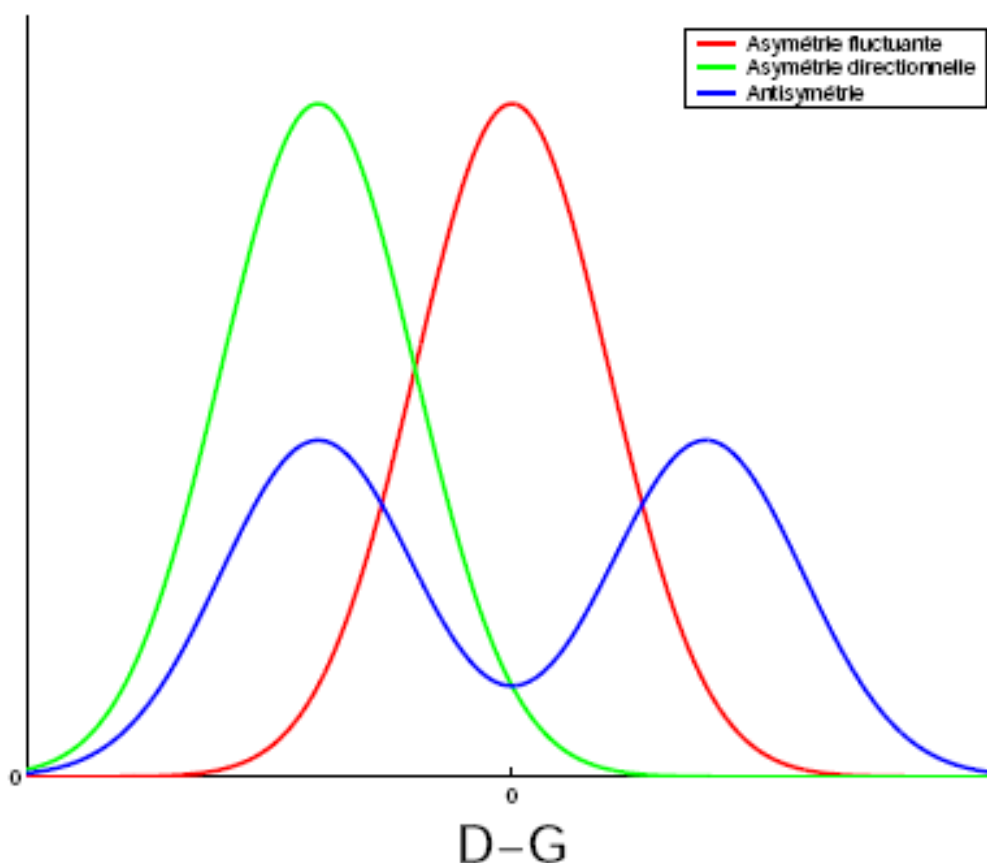


FIG. 3 – Représentation graphique des trois types de symétrie à l'échelle d'une population

Des asymétries intrigantes : le cerveau humain

Avec ses 2 hémisphères dans lesquels chaque structure (lobes, ventricules, *etc.*) possède un homologue dans l'hémisphère opposé, le cerveau apparaît comme globalement anatomiquement symétrique (Figure 4). Cependant, il possède de nombreuses asymétries.

Asymétries fonctionnelles

Tout d'abord, d'un point de vue fonctionnel, les différentes aires de part et d'autre de la fissure interhémisphérique sont associées à des fonctions différentes [LeM99]. L'aire de Broca constitue un exemple intéressant d'asymétrie fonctionnelle : cette aire liée à la production du langage articulé est dans la plupart des cas présente dans l'hémisphère gauche du cerveau. Plus généralement, on associe à l'hémisphère gauche des tâches analytiques (parole, lecture, écriture) et à l'hémisphère droit des tâches plus pratiques comme l'appréhension de l'espace et des formes (même si la réalisation d'une tâche nécessite souvent l'usage des 2 hémisphères). Cette prépondérance du rôle de l'un ou l'autre des hémisphères cérébraux pour une tâche donnée est connue sous le nom de dominance cérébrale et on désigne parfois l'hémisphère lié au langage comme l'hémisphère dominant.

Il est aussi intéressant de noter la troublante asymétrie qui lie préférence manuelle et latéralité des zones du langage : chez 99% des droitiers (environ 89% de la population), la fonction du langage se situe dans l'hémisphère gauche. Cependant, chez les gauchers et les ambidextres, la situation est plus ambiguë puisque chez 56% d'entre eux, les aires associées au langage sont situées à gauche et chez les 44% restants, elles se situent dans l'hémisphère droit ou sont réparties entre les 2 hémisphères. Ainsi, dominance cérébrale et préférence manuelle seraient liées. À ceci, le genre vient ajouter son influence : il a été noté que les femmes victimes de lésions cérébrales dans l'hémisphère gauche comme dans l'hémisphère droit du cerveau étaient moins sujettes à l'aphasie (trouble du langage affectant l'expression ou la compréhension du langage parlé ou écrit) que les hommes [McG80]. Cette observation suggère ainsi une latéralisation plus importante du cerveau des hommes que de celui des femmes.

Asymétries anatomiques

En plus de ces asymétries des répartitions fonctionnelles, le cerveau possède plusieurs asymétries anatomiques notables. Ces asymétries sont observées depuis plus de 150 ans, tout d'abord via des études post-mortem puis via la neuroimagerie. Les asymétries cérébrales les plus décrites dans la littérature sont probablement celles du planum temporale [GL68], qui est généralement plus large à gauche qu'à droite, ou encore de la scissure de Sylvius [RMH76], qui est plus longue à gauche qu'à droite et dont la terminaison suggère un tracé plus vertical à droite qu'à gauche chez la plupart des gens. En plus de ces asymétries locales, le cerveau humain a été décrit comme étant sujet à une torsion globale, dans le sens anti-horaire et selon un axe vertical (c'est-à-dire parallèle au corps humain) [WLMW82]. La magnitude de cette torsion a été quantifiée et

interprétée de plusieurs façons : i) comme une différence de taille ou de volume des 2 hémisphères [BWB⁺94] ou ii) comme une avancée antéro postérieure du lobe frontal droit par rapport au gauche et du lobe occipital gauche par rapport au droit [VB62], ou encore iii) comme une déviation vers la gauche (resp. droite) de la fissure interhémisphérique au niveau des lobes frontaux (resp. occipitaux).

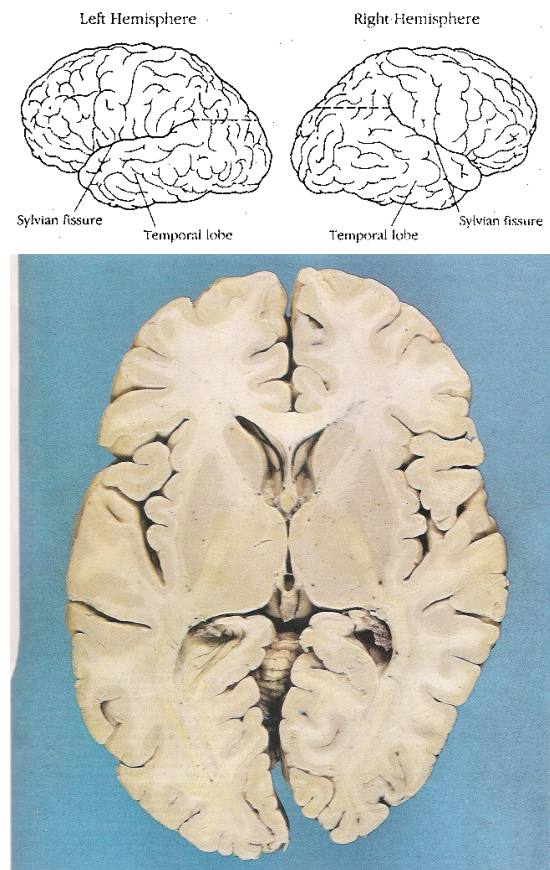


FIG. 4 – Illustration de l'asymétrie de la fissure de Sylvius et de la torsion globale du cerveau

Tiré de <http://www.benbest.com/science/anatmind/anatmd6.html>

Des questions sans réponses

À l'heure actuelle, on ne sait pas quelle est la prévalence de l'effet de torsion du cerveau dans une population donnée. Cependant, il est la clé de voûte de plusieurs théories. En effet, pour certains, il est considéré comme le substrat de la dominance du langage chez les humains. Il a été de surcroît observé comme absent chez les patients schizophrènes (chez lesquels il a été suggéré que la dominance cérébrale n'est pas développée) [BWB⁺94, Cro04b] et chez les grands singes [Cro04a]. Cependant, chacune de

ces considérations est très controversée [BMP⁺05, MBR⁺03, GHBB98, CH01] et pour l’heure, il n’existe pas de consensus.

Toutes ces observations étrangement liées, parfois contradictoires, font des asymétries cérébrales un sujet d’étude important.

Quantifier les variations : un challenge aux enjeux multiples

Qu’il s’agisse de l’anatomie globale ou d’une de ses caractéristiques comme l’asymétrie, l’existence d’une base commune et de “lois” régissant les variations d’une population autour de cette base évoque inmanquablement la question de l’identification de chacun de ces acteurs.

Modèles descriptifs déterministes et modèles génératifs stochastiques

Quand les causes physiques et biologiques de la variabilité observée sont bien comprises, il est parfois possible de construire un modèle de type génératif (et déterministe). Ainsi, la somme des variabilités est expliquée par des paramètres physiques et biologiques régissant les interactions entre les acteurs. Cependant ces interactions sont souvent compliquées, nombreuses, s’effectuent sur des échelles (de temps et d’espace) très différentes et sont parfois mal comprises. Ainsi, dans de nombreux cas et à partir des connaissances dont on dispose actuellement, un tel modèle ne peut pas être construit.

Si on ne peut expliquer, on peut observer et reproduire : l’approche alternative consiste alors à construire un modèle génératif (et aléatoire). Pour cela, on considère une population d’observations comme un ensemble de réalisations d’un processus aléatoire dont on va chercher à estimer les lois. A partir de ces lois, on peut alors identifier la base commune et les lois de variations autour de celle-ci. Dans cette thèse, nous nous intéressons à la construction de tels modèles.

Enjeux des modèles génératifs

Les applications liées à la construction de modèles génératifs sont nombreuses. Tout d’abord, ils permettent de caractériser de manière compacte, voire interprétable, les variations observées au sein d’une population de structures anatomiques. Ensuite, les comparaisons de modèles de populations différentes peuvent mener à l’identification de caractéristiques discriminantes entre ces populations. Ainsi, on peut comparer l’anatomie de populations de patients à celle de populations de sujets sains afin de mieux comprendre le développement des pathologies, ou encore proposer une aide au diagnostic en cherchant à savoir à quelle population un individu donné semble appartenir. Ce type d’application est enfin très utile en paléanthropologie où l’une des questions récurrentes consiste à déterminer à quelle espèce appartient tout nouveau fossile découvert, voire à définir une nouvelle espèce.

Construction de modèles génératifs : amers et morphométrie

La conception de modèles génératifs a tout d'abord été traitée comme un problème de statistiques multivariées sur un ensemble finie mesures faites sur les données (distances, angles, âge, *etc*). Une première révolution a été opérée par le développement de statistiques dans l'espace des formes (méthode de superposition procrustéenne) par Kendall [DM98] et de l'analyse à base de modèles de plaques minces (analyse des efforts mis en jeu pour passer d'une forme à une autre) par Bookstein [Boo89]. Ces méthodes, encore largement utilisées, permettent d'analyser directement la forme des structures, représentées par des ensembles d'amers anatomiques identifiés de manière unique sur chaque sujet. Au-delà des enjeux méthodologiques présentés par de telles approches, celles-ci souffrent de limitations liées à la nature même des données traitées. En effet, les structures étudiées ne possèdent en général qu'un nombre très limité d'amers anatomiques. L'ensemble de ces amers ne permet alors qu'une représentation très schématique des structures étudiées et leur analyse n'offre qu'une vision limitée des variations en jeu. De surcroît, la capture d'amers est une opération laborieuse.

Plus récemment, le développement des méthodes de vision par ordinateur et la démocratisation des systèmes d'acquisition numérique ont stimulé le développement de méthodes basées non plus sur un ensemble d'amers mais sur une description complète de la structure étudiée. Dans la plupart des cas, cette description est donnée par l'intermédiaire d'une image 3D provenant d'acquisitions CT ou IRM. De nombreuses méthodes permettent d'établir des modèles de formes à partir de ces données expérimentales. Cependant, les données brutes offrent une description de l'ensemble des structures présentes dans le champ d'acquisition et ne permettent pas de se concentrer sur une structure donnée. De plus, elles sont soumises à plusieurs types d'artefacts (bruit, inhomogénéité) susceptibles de gêner les méthodes d'estimation de modèles. Une technique alternative consiste à segmenter la structure d'intérêt puis à considérer uniquement la surface délimitant son contour. Plusieurs méthodes ont été proposées dans la littérature pour le traitement de telles données et nous proposons dans cette thèse de poursuivre cette approche. Cependant, il est important de noter que ces opérations d'extraction, qu'elles soient faites automatiquement ou manuellement, impliquent une perte vis-à-vis des données expérimentales. De plus, certains algorithmes de segmentation utilisent des *a priori* de forme et les données extraites par de telles méthodes sont susceptibles de biaiser la construction d'un modèle statistique.

Positionnement et organisation du document

L'objectif premier de cette thèse est de proposer un ensemble d'outils permettant de quantifier et de comparer statistiquement l'asymétrie globale au sein d'une population et entre deux populations de structures représentées par des nuages de points. Plusieurs outils ont été proposés dans la littérature [ÓLD⁺07, CDD07, CKDD08, LKAD⁺09] pour la réalisation d'une telle tâche et on peut distinguer 2 approches principales : les approches basées sur le calcul et l'analyse des déformations liant la structure d'intérêt à son image symétrique et les approches basées sur l'extraction et l'analyse de caractéris-

tiques de plus haut niveau (forme globale, épaisseur corticale, *etc*). Ces deux approches adoptent des points de vue différents et sont complémentaires.

Cependant, tous ces travaux reposent sur l'intégration d'outils de recalage dans une chaîne de traitement pour la quantification des asymétries. Ainsi, la continuité méthodologique n'est pas assurée et il est parfois difficile d'interpréter les résultats de ces méthodes. Dans ce sens, il est intéressant de noter que le plan de symétrie n'est jamais estimé par ces méthodes mais est calculé à partir de recalage rigide ou d'alignement manuel. De plus, certaines d'entre elles [CDD07, CKDD08, LKAD⁺09] ont de forts *a priori* sur les données (topologie, échantillonnage, *etc*) et sont dédiées à des structures particulières.

Nous proposons ici un ensemble d'outils alternatifs permettant la quantification et la comparaison des asymétries au sein et entre populations de manière rapide, robuste et sans *a priori* sur les données. En particulier, ces outils permettent le calcul de cartes statistiques d'asymétries directionnelles locales et l'analyse de la déviation d'une "surface" de symétrie par rapport à une symétrie parfaite.

Plus précisément, nos choix nous ont amenés au développement de méthodes (au sein d'un même cadre probabiliste) pour i) l'estimation robuste d'un plan de symétrie, ii) la définition de surfaces permettant d'étendre la notion de symétrie par rapport à un plan et d'offrir une nouvelle approche pour quantifier l'écart d'un objet à une symétrie parfaite, iii) le recalage non-linéaire et iv) la création de modèles statistiques de formes. Pour chacun de ces points, nous nous sommes efforcés de proposer des méthodes robustes ne supposant pas d'*a priori* sur les données (échantillonnage, topologie, *etc.*) et permettant de gérer des données de tailles importantes en un temps raisonnable. La suite de cette thèse est décomposée de la manière suivante :

Partie 1 : Estimation de symétrie bilatérale dans des nuages de points 3D

Chapitre 1 : Dans ce chapitre, nous présentons brièvement les enjeux méthodologiques et applicatifs liés à l'estimation de la symétrie bilatérale dans des nuages de points 3D.

Chapitre 2 : Dans ce chapitre, nous proposons une taxonomie des méthodes existantes d'estimation de symétrie bilatérale. Nous proposons 3 grandes approches et montrons comment une grande partie des méthodes de la littérature s'insère dans ces approches.

Chapitre 3 : Forts de cette taxonomie, nous développons dans le 3eme chapitre une méthode pour l'estimation de toutes les symétries bilatérales dans un objet représenté par un nuage de points 3D. Cette méthode est composée de 2 sous-méthodes. La première, basée sur des outils de clustering, permet d'estimer de manière grossière tous les plans de symétrie d'un objet. La seconde permet d'estimer de manière précise et robuste chacun de ces plans de symétrie étant donné une initialisation. Elle est basée sur un estimateur robuste de type maximum de vraisemblance qui est estimé par un algorithme de type Expectation-Maximisation. Les résultats obtenus par chacun de ces deux algorithmes sont illustrés sur des données synthétiques et réelles. Nous montrons

en particulier que nos outils permettent d'estimer le plan de symétrie d'un objet de manière très précise même quand la part d'asymétrie est importante.

Chapitre 4 : Dans ce chapitre, nous montrons comment étendre la notion de symétrie par rapport à un plan à celle de symétrie par rapport à une surface, et comment estimer de telles surfaces sur des objets 3D. Nous proposons quelques applications possibles de cette surface de symétrie.

Chapitre 5 : Dans ce chapitre, nous analysons les performances de nos méthodes et la pertinence de nos choix et proposons quelques pistes pour de travaux futurs.

Partie 2 : Recalage non-linéaire de nuages de points 3D

Chapitre 6 : Dans ce chapitre, nous présentons brièvement les enjeux méthodologiques et applicatifs liés au calcul des déformations superposant 2 nuages de points

Chapitre 7 : Le chapitre 7 est dédié à un état de l'art des méthodes de recalage non-linéaire de nuages de points et de surfaces. De même qu'au chapitre 2, nous nous efforçons de mettre en évidence les hypothèses de base souvent cachées sur lesquelles reposent nombre d'algorithmes de la littérature, permettant une appréhension plus objective de leurs différences et de leurs points communs. Nous montrons en particulier qu'une technique particulièrement fructueuse consiste à considérer le problème du recalage comme un problème de minimisation d'une distance entre deux densités de probabilité par rapport à une fonction paramétrisant une des densités. Ainsi, les nuages de points sont représentés sous formes de mélanges de densités de probabilités et l'on va chercher la transformation déformant "au mieux" un mélange vers l'autre. Nous montrons que l'algorithme EM-ICP appartient à cette classe de méthodes et mène à des solutions efficaces dans le cas du recalage rigide.

Chapitre 8 : Malgré ses capacités à traiter le problème du recalage rigide de manière efficace, nous observons que l'EM-ICP est asymétrique et qu'une extension au cas non rigide peut s'avérer très consommatrice en temps machine et est susceptible de converger vers une solution inappropriée si une bonne initialisation n'est pas fournie. Dans le chapitre 8, nous proposons des solutions pour chacune de ces limites. Pour cela, nous enrichissons le critère original et proposons des solutions algorithmiques efficaces pour chacune de ces modifications. En particulier, nous proposons :

- de symétriser le critère,
- d'introduire des probabilités a priori sur les mises en correspondance entre les points des deux nuages,
- d'utiliser deux nouveaux opérateurs de régularisation de champs de déformations très efficaces. Le premier est basé sur l'affectation locale de modèles de transformations simples (affine ou translation) qui sont ensuite régularisés à travers l'espace, le second est basé sur l'analyse fréquentielle du champ de déformations et la théorie des noyaux reproduisants.

A la fin de ce chapitre, nous sommes en mesure de proposer un algorithme permettant le recalage de nuages de points de tailles importantes (10000 points) de manière

efficace (3 minutes sur un PC standard).

Chapitre 9 : Dans ce chapitre, nous évaluons et illustrons les performances de notre algorithme et les avantages relatifs des différentes améliorations proposées sur des données synthétiques et réelles.

Chapitre 10 : Dans ce chapitre, nous analysons les performances de nos méthodes et la pertinence de nos choix et proposons quelques pistes pour des travaux futurs.

Partie 3 : Outils pour l'analyse de groupe & Applications

Chapitre 11 : L'objectif premier de cette thèse est de proposer un ensemble d'outils permettant de caractériser et de comparer les asymétries présentes au sein et entre des populations de structures anatomiques. Dans le chapitre 11, nous montrons comment nos choix méthodologiques mènent à des solutions à la fois simples et pertinentes pour la construction de modèles statistiques de formes et l'analyse de groupe. En particulier, nous montrons comment calculer des cartes d'asymétries individuelles et comment construire une forme moyenne et projeter l'ensemble des cartes d'asymétries individuelles sur celle-ci.

Chapitre 12 : Dans ce chapitre, nous illustrons nos outils par une étude comparant les asymétries corticales de 2 sujets Situs Inversus (condition congénitale dans laquelle les principaux viscères et organes sont inversés dans une position en miroir par rapport à leur situation normale) et celles d'une population contrôle. Cette étude utilise l'ensemble des outils développés dans cette thèse (recalage non linéaire, estimation de plan et de surface de symétrie et estimation de forme moyenne)

Chapitre 13 : Dans ce chapitre, nous discutons la pertinence de notre approche pour la quantification d'asymétries.

Chapitre 14 : Dans ce chapitre, nous concluons ce document et illustrons l'apport de notre travail par la description de plusieurs études en cours impliquant de nombreux collaborateurs venus de disciplines différentes.

Chaque partie peut être lue de manière indépendante.

Première partie

Estimating bilateral symmetries of 3D point sets

Chapitre 1

Introduction

1.1 Introduction

Most natural or manufactured objects exhibit some form of symmetry. In mathematical terms, it means that they remain invariant or almost invariant under certain transformations, typically : rotations (with respect to a point or a line) or reflections (with respect to a point, a line or a plane). More specifically, the set of symmetries of a 3D object forms a group in which each element can either be a rotation, a reflection or a composition of them.

In this work, we focus our interest on reflection (*i.e.* bilateral) symmetry in 3D point sets. In computer vision, it is often crucial to use these symmetry features for common image processing tasks such as recognition, denoising, registration, segmentation, *etc.* Most initial works considered the estimation of reflections in 3D point sets in an ideal case where objects under study exhibit single or multiple perfect symmetries [Ata85, WWV85]. These algorithms convert the estimation of reflections in 3D point sets into a 1D pattern matching problem for which efficient solutions are known [Knu98]. However, in practice most real objects exhibit approximate (as opposed to perfect) and even partial (*i.e.* that concern only a small subpart of the object) symmetries. Moreover, an object can exhibit several approximate global and partial symmetries. For about 20 years, the computer vision community has proposed several methods to estimate multiple global/partial approximate symmetries in 3D point sets.

In the following, we first propose a taxonomy including a major part of existing methods for the estimation of reflections in 3D point sets (Chapter 2). This taxonomy elucidates some links between the different methods of the literature and allows an analysis of the relative advantages and drawbacks of the different strategies. Based on this taxonomy, we propose a method to coarsely estimate all the “meaningful” reflections of a 3D point set (Section 3.1). Then, we propose a method that, given an initialisation, estimates a reflection in a 3D point set in a robust and accurate way (Section 3.2). More precisely, we consider that at least a small subset of the object is not affected by asymmetries and we define the approximate symmetry plane of the object as the one minimising the average sum of squared distances between the points belonging to

this highly symmetrical part. The combination of both approaches allows an automatic, accurate and robust estimation of all the meaningful reflections in 3D point sets. Then we extrapolate the notion of symmetry about a plane to that of symmetry about a differentiable 3D surface C (Section 4.1) and propose an algorithm to estimate such a surface in 3D point sets.

1.2 Notations & basic formulas

- We note $X = \{x_1, \dots, x_N\}$ a 3D point set and $x_G = \frac{1}{N} \sum_{x_i \in X} x_i$ its centroid.
- We consider P as a 3D plane and $S_P : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ as the reflection with respect to P . Each plane P can be uniquely parametrised as $P = (\alpha, \beta, d)$ where $\alpha \in [-\pi, \pi]$, $\beta \in [-\pi/2, \pi/2]$ and $d \in \mathbb{R}^+$ represent respectively the azimuth angle, the latitude angle and the distance from P to the origin (Fig. 1.1). An equivalent but more convenient parametrisation for calculus is $P = (n, d)$ where n is a unit vector orthogonal to P (Fig. 1.1). The normal n is linked to α and β by the following relationship :

$$n = (n_x, n_y, n_z) = (\cos(\beta) \cos(\alpha), \sin(\beta) \cos(\alpha), \sin(\alpha))$$

\Leftrightarrow

$$\alpha = \arcsin\left(\frac{n_z}{\|n\|}\right), \beta = \begin{cases} 0 & \text{if } n_x = n_y = 0 \\ \arctan\left(\frac{n_z}{n_x}\right) & \text{if } n_x \text{ and } n_y > 0 \\ 180 + \arctan\left(\frac{n_y}{n_x}\right) & \text{else} \end{cases}$$

- The distance from a point x to a plane $P = (n, d)$ is given by :

$$d(x, P) = |n^T x - d|. \quad (1.1)$$

- The image of point x through reflection with respect to $P = (n, d)$ is given by :

$$S_P(x) = (I_3 - 2nn^T)x + 2dn, \quad (1.2)$$

where I_3 is the 3×3 identity matrix.




FIG. 1.1 – Parametrisation of a plane in 3D.

Chapitre 2

Bibliography

In this chapter, we aim at designing a taxonomy of existing methods for the estimation of (multiple/partial) approximate bilateral symmetries of 3D points. We identify 3 main approaches :

- approaches based on points distribution (Section 2.1),
- approaches based on points superimposition (Section 2.2),
- approaches based on planes clustering (Section 2.3),

and discuss their characteristics. In practice, almost all methods proposed in the literature are a combination of these 3 approaches. We briefly list other methods that do not fit in this taxonomy in Section 2.4.

2.1 Approaches based on points distribution

2.1.1 Basics

The general view of this approach is to define the optimal plane P as “**the plane that best passes through the point set X** ”. Thus, P can be simply considered as the plane minimising \mathcal{E}_1 :

$$\begin{aligned}\mathcal{E}_1(P) &= \sum_{x_i \in X} d(x_i, P)^2 \\ &= \sum_{x_i \in X} |n^T x_i - d|^2.\end{aligned}$$




FIG. 2.1 – 2D illustration of “the plane P that best passes through the point set X ”

Note that when working with meshes, a simple improvement of this approach consists in working with sums of integrals over triangles instead of sums over vertices. This way, the approach is less dependent on the sampling of X [VSR01].

The partial derivatives of \mathcal{E}_1 with respect to d and n must vanish at the minimum. The first condition implies that d must be equal to $n^T x_G$. Thus, given that $\|n\| = 1$, n must minimise :

$$\sum_{x_i \in X} |n^T (x_i - x_G)|^2 = \sum_{x_i \in X} \|x_i - x_G\|^2 - \sum_{x_i \in X} \|n \times (x_i - x_G)\|^2, \quad (2.1)$$

\times being the cross-product. The first term at right of the equal sign is a constant with respect to n . By an analogy with mechanics, considering that each point of X has a mass equal to one, the second term is simply the second-order moment of X with respect to the axis aligned with n and passing through x_G . In other words, the plane minimising \mathcal{E}_1 is orthogonal to the principal axis of X having the highest moment of inertia (sometimes called the “figure axis”). The intuitive definition of \mathcal{E}_1 is thus very satisfying as it can be shown that any perfect symmetry plane of an object is orthogonal to one of its principal axes, which suggests that all the (local) optima of \mathcal{E}_1 can define a relevant approximate symmetry plane of X . This suggests a simple algorithm to find all the approximate reflections of X :

- compute the 3 eigenvalues and related eigenvectors of the scatter matrix of X (related to the inertia matrix in a straightforward way).
- check the 3 planes orthogonal to these eigenvectors and passing through x_G for symmetry (computing a distance - to be defined - between X and its mirrored version).

One can see that this algorithm implicitly assumes that X has less than 3 distinct symmetry planes. If this is not the case, then at least 2 of the eigenvalues of the scatter matrix are equal. This causes \mathcal{E}_1 to exhibit an infinity of optima (any n belonging to the eigenspace associated to these equal eigenvalues), that do not necessarily correspond to valid symmetry planes of X . As an example, a cuboid (resp. a cube) has only 2 (resp. 1) distinct eigenvalues, but 5 (resp. 9) symmetry planes. Even worse, a cylinder (resp. a sphere) has only 2 (resp. 1) distinct eigenvalues, but an infinite number of symmetry planes.

2.1.2 Generalisation

To deal with this problem when two of the eigenvalues are equal, Minovic et al. [MIK93] proposed a solution based on the use of complex moments to study the 2D section of X passing through x_G and orthogonal to the marked axis (*i.e.* the one associated to the single eigenvalue). Unfortunately, this solution does not work when the 3 eigenvalues are equal.

More recently, Martinet et al. [MSHS06] proposed to cope with this problem by generalising the notion of moments with respect to an axis to higher orders. These generalised moment functions write :

$$\mathcal{E}_1^{(2p)}(P) = \sum_{x_i \in X} d(x_i, n)^{2p} = \sum_{x_i \in X} \|n \times (x_i - x_G)\|^{2p}. \quad (2.2)$$

They show that any perfect symmetry plane of X is an optimum of $\mathcal{E}_1^{(2p)} \forall p \in \mathbb{N}$, which extends the previously stated relationship between symmetry planes and principal axes. Looking for planes $(n, n^T x_G)$ that optimise all the even order moments allows to be much more selective compared to just considering the second order moment : lots of planes optimising $\mathcal{E}_1 = \mathcal{E}_1^{(2)}$ do not optimise $\mathcal{E}_1^{(2p)}$ ($p > 1$) and thus can be easily discarded. In practice, Martinet et al. optimise $\mathcal{E}_1^{(2p)}$ for a small number of values p (typically 1,2,3). In absence of a closed-form solution for $p > 1$, this is done by i) finding P where $\nabla \mathcal{E}_1^{(2p)}$ is close to 0 by an exhaustive search on a discrete adaptive sampling of the directions n and ii) performing a steepest descent algorithm on $\|\nabla \mathcal{E}_1^{(2p)}\|$ for each candidate find in i). The whole process is implemented efficiently using spherical harmonics. Normals n that simultaneously reduce $\|\nabla \mathcal{E}_1^{(2p)}\|$ to zero for the selected orders $2p$ are considered as interesting directions for reflection. Then for each such extracted direction n , the corresponding plane P is checked by computing a distance between $S_P(X)$ and X . Interestingly, this approach allows to identify all the symmetries of X (i.e. simple, rotational or spherical symmetries and their improper forms, including reflections).

2.1.3 Advantages and limitations

The key property of this approach is to avoid the search of bilateral correspondences between points of X by just considering the distribution of the point locations with respect to the plane. This results in simple and effective methods that exhibit good results in presence of an almost symmetrical object. It explains why it has been used in several works to detect and identify different types of approximate [MIK93, CV08] and partial [PeS⁺06] symmetries of 3D objects. However it performs very badly in case of large deviation from perfect symmetry. Furthermore, making this approach robust is a difficult issue as large residuals of \mathcal{E}_1 or higher order moments do not necessary correspond to outliers but must often to points far from the plane. However, some proposition consisting in simply robustifying \mathcal{E}_1 either with M-estimators [PeS⁺06] or least median of squares estimators [LR09] have exhibited good results.

2.2 Approaches based on points superimposition

2.2.1 Basics

Zabrodsky *et al.* [ZPA92, ZPA95] defined what they call the symmetry distance as **“a quantifier of the minimum effort required to turn a shape X into a mirror symmetric shape”**. The symmetry plane of this perfectly symmetrical shape X^* can then be considered as the approximate symmetry plane of X . Typically, this “effort”

is quantified as the sum of squared displacements from X to X^* . The approximate symmetry plane is then designed as a minimiser of :

$$\mathcal{E}_2(P) = \min_{(x_i^*) \in \mathbb{R}^3} \sum_{x_i \in X} \|x_i - x_i^*\|^2$$

with $\forall x_i^* \in X^*, \exists x_j^* \in X^*$ such that $x_i^* = S_P(x_j^*)$

FIG. 2.2 – 2D illustration of “the plane P for which the minimum effort is required to turn X into a mirror symmetric point set with respect to plane P ”. The circles represent (x_i^*) , the crosses represent (x_i) .

One can first derive the optimal displaced points (x_i^*) . The constraint of symmetry for X^* imposes that for all $x_j \in X$, there exists a unique point $x_{\pi(j)}$ in X such that $S_P(x_j^*) = x_{\pi(j)}$ ($\pi : [1, \dots, N] \rightarrow [1, \dots, N]$). Thus, solving for the optimal (x_i^*) implicitly induces a pairing between points of X . As a result, for all j such that $\pi(j) \neq j$, one can write :

$$\sum_{x_i \in X} \|x_i - x_i^*\|^2 = \|x_j - x_j^*\|^2 + \|x_{\pi(j)} - S_P(x_j^*)\|^2 + C, \quad (2.3)$$

where C does not depend on x_j^* . Differentiating this expression with respect to x_j^* gives :

$$\tilde{x}_j^* = \frac{x_{\pi(j)} + S_P(x_j)}{2}. \quad (2.4)$$

When considering the case where $\pi(j) = j$, we find the same expression for the optimal \tilde{x}_j^* . In this case, it implies that $x_j^* \in P$. Using the optimal values (\tilde{x}_i^*) , $\mathcal{E}_2(P)$ can then be simply reformulated as :

$$\mathcal{E}_2(P) = \min_{\pi} \sum_{x_i \in X} \|x_{\pi(i)} - S_P(x_i)\|^2 / 4, \quad (2.5)$$

where $\forall i, j, \pi(i) = j \Leftrightarrow \pi(j) = i$ (*i.e.* the pairing permutation function π is both bijective and involutive).

In other words, the plane P minimising $\mathcal{E}_2(P)$ can be viewed as “**the plane that minimises the left-right differences of X** ” and $\mathcal{E}_2(P)$ can be reinterpreted as :

$$\mathcal{E}_2(P) = \min_{x'_i \in X} \sum_{x_i \in X} \|S_P(x_i) - x'_i\|^2/4$$

where $\forall i, j, x'_i = x_j \Leftrightarrow x'_j = x_i$.




FIG. 2.3 – 2D illustration of “the plane P that minimises the left-right differences of X ”.

With this formulation, minimising \mathcal{E}_2 appears as an interleaved problem that consists in finding both i) the optimal left-right matchings (x_i, x'_i) (with $\forall i, j, x'_i = x_j \Leftrightarrow x'_j = x_i$) and ii) the optimal symmetry S_P best superposing (x'_i) on (x_i) .

In the following, we first discuss closed-form solutions for the problem of computing the optimal approximate symmetry plane when the left-right matchings are known. Then we describe the three main strategies that have been proposed in the literature to minimise $\mathcal{E}_2(P)$.

2.2.2 Optimal \tilde{P} with known correspondents (\tilde{x}'_i)

If one assumes that the optimal set of correspondents $\tilde{x}'_i \in X$ is known, two main approaches have been proposed to find the optimal plane P that best superposes (\tilde{x}'_i) on (x_i) .

• Estimation using an intermediate rigid-body transformation :

A first, intuitive approach consists in choosing an arbitrary plane K and rewriting this minimisation problem as :

$$\tilde{P} = \arg \min_{S_P} \sum_{x_i \in X} \|\tilde{x}'_i - S_P \circ S_K \circ S_K(x_i)\|^2$$

Then one can note that :

- if $K \parallel P$ then $S_P \circ S_K$ is a translation
- if $K \not\parallel P$ then $S_P \circ S_K$ is a rotation

In other words, for any plane K , $T = S_P \circ S_K$ is a rigid-body transformation. Thus, an intuitive approach to compute \tilde{P} is to fix a plane K and to find the rigid-body transformation \tilde{T} best superposing X and $S_K(X)$:

$$\tilde{T} = \arg \min_T \sum_{x_i \in X} \|\tilde{x}'_i - T \circ S_K(x_i)\|^2$$

Different closed-form solutions exist for this last problem using, for instance, the

unit quaternions [Hor87, FH86] or the singular value decomposition [AHB87, HHN88, Ume91].

Unfortunately, this approach suffers from a major flaw : the optimal transformation \tilde{T} , composed with S_K , does not necessarily define a proper reflection, and thus the optimal plane \tilde{P} cannot be computed directly using $S_{\tilde{P}} = \tilde{T} \circ S_K$. For instance, if \tilde{T} is a pure translation that is not perpendicular to K , then $\tilde{T} \circ S_K$ is not a reflection. Several heuristic approaches have been proposed in the literature to estimate \tilde{P} from \tilde{T} [WPW03, ZRF⁺06, MCP⁺06, ZPA95] but to our knowledge, none of these is guaranteed to yield a plane \tilde{P} really minimising $\mathcal{E}_2(P)$ (when correspondents (\tilde{x}'_i) are fixed).

• **Direct estimation of a reflection :**

Recently and independently, we and several other authors [CHW⁺08a, MGP07, FEdlF⁺08, MBM00] have proposed a closed-form solution for the former minimisation problem without resorting to an intermediate couple (K, \tilde{T}) . It gives, in a general form (the proof is provided in Appendix A) :

For a given square matrix A of size $N \times N$ (independent of P), the plane $P = (d, n)$ that minimises

$$\sum_{x_i \in X} \sum_{x_j \in X} A_{i,j} \|x_j - S_P(x_i)\|^2 \text{ with } \forall i, j \ A_{i,j} \geq 0 \quad (2.6)$$

is characterised by :

- n colinear with the eigenvector corresponding to the smallest eigenvalue of the 3×3 matrix B , where

$$B = \sum_{(x_i, x_j) \in X^2} A_{i,j} [(x_i - g_1 + x_j - g_2)(x_i - g_1 + x_j - g_2)^T$$

$$- (x_i - x_j)(x_i - x_j)^T]$$

$$\text{and } \begin{cases} g_1 = \frac{1}{\sum_i \sum_j A_{i,j}} \sum_{(x_i, x_j) \in X^2} A_{i,j} x_i \\ g_2 = \frac{1}{\sum_i \sum_j A_{i,j}} \sum_{(x_i, x_j) \in X^2} A_{i,j} x_j \end{cases} .$$

- $d = \frac{1}{2}(g_1 + g_2)^T n$.

Notice that this formula is linked to that obtained by Horn [Hor87] and Faugeras and Hebert [FH86] for rigid-body registration of point clouds. There is an analogy between the parameters (n, d) (unit normal vector, distance to the origin) and their parameters (q, t) (unit quaternion, translation component).

2.2.3 Minimisation strategies

In practice, both optimal correspondents (\tilde{x}'_i) and optimal plane \tilde{P} are unknown and one has to estimate both parameters. In the literature three main strategies have been

proposed to achieve this estimation.

Brute force minimisation : search over matchings (x_i, x'_i)

The set of admissible solutions for the matchings (x_i, x'_i) is the set of involutive permutations of (x_i) . It contains $\sum_{k=0}^{\lfloor N/2 \rfloor} \frac{N!}{(N-2k)! 2^k k!}$ elements [Knu98]. Zabrodsky [ZPA95] simply proposed to tackle the problem using an exhaustive search within this set. For each set of generated matchings (x_i, x'_i) , one can compute the corresponding optimal plane (Section 2.2.2) and the one providing the lowest criterion will be the optimum. Unfortunately, the combinatorial explosion limits this approach to small point sets.

Brute force minimisation : search over planes P

Podolak and colleagues [PSG⁺06] proposed to minimise \mathcal{E}_2 (more precisely a normalised version of \mathcal{E}_2) by sampling the space of \mathbb{R}^3 planes in order to test a finite number of them. For that, they proposed to represent the point set X by its Gaussian distance transform in order to avoid the matching process when having to compute \mathcal{E}_2 with a known plane. In practice, the computation of $\mathcal{E}_2(P)$ for each of the tested planes is approximated via an importance sampling strategy.

This minimisation strategy relies on the conversion from surface to volume data which can be problematic in terms of memory/time requirements. Moreover, due to the sampling scheme this strategy is only able to give a coarse estimate of local minima of \mathcal{E}_2 and needs a refinement of the solutions.

Alternated minimisation

Another strategy, if one has a prior knowledge of the plane P , is to perform the left-right matchings based on this plane, recompute the optimal plane based on these matchings, and iterate. This yields to the following algorithm :

$$\begin{array}{l} \text{Step i) } (\tilde{x}'_i) = \arg \min_{x'_i \in X} \sum_{x_i} \|x'_i - S_{\tilde{P}}(x_i)\|^2 \\ \text{with } \forall i, x'_i = x_j \Leftrightarrow x'_j = x_i. \\ \text{Step ii) } \tilde{P} = \arg \min_P \sum_{x_i} \|\tilde{x}'_i - S_P(x_i)\|^2 \end{array}$$

This algorithm monotonically converges to a local minimum of \mathcal{E}_2 . Obviously, as this method exhibits a local convergence, it assumes that a good initialisation of P is available.

Step ii) is the problem described in Section 2.2.2. To our knowledge, there is no efficient solution for Step i) and an exhaustive search is needed. In practice, this step can be solved removing the constraint that $\forall i, j, x'_i = x_j \Leftrightarrow x'_j = x_i$. In this condition, each \tilde{x}'_i can be computed independently :

$$\forall i, \tilde{x}'_i = \arg \min_{x'_i \in X} \|x'_i - S_{\tilde{P}}(x_i)\|^2 \quad (2.7)$$

Step i) then consists in solving N closest point problems that can be efficiently dealt with using (for instance) a kd -tree [Ben75]. Removing the constraint is interesting for two reasons. First it allows an efficient optimisation of the modified criterion \mathcal{E}_2 . Second,

in many cases, the constraint is irrelevant as it implicitly assumes that there are as many points on each side of the plane and that the sampling of X is homogeneous. As these conditions are generally not met, this constraint can lead to inappropriate results.

Note that, when removing the constraint, the proposed algorithm is nothing else but the well-known iterative closest point algorithm [BM92] where the unknown transformation is a reflection instead of a rigid-body transformation. This ICP strategy can be viewed as a method to minimise a constraint-free version of the original criterion \mathcal{E}_2 proposed by Zabrodsky *et al.* This strategy has been used intensively in the context of symmetry plane estimation [CHW⁺08a, MGP07, FEdIF⁺08, WPW03, ZRF⁺06, MCP⁺06, PSG⁺06].

2.2.4 Generalisation

A more general formulation

In the previous subsections, we proposed a comprehensive view for the original criterion proposed by Zabrodsky *et al.* In this section, we consider a more generic solution that consists in considering a space \mathcal{F} representing the set of the 3D shapes. If one considers the shape under study as $f \in \mathcal{F}$, following the approach of Zabrodsky, one can propose the following criterion :

$$\mathcal{E}'_2(P) = \min_{f^* \in \text{Sym}(P)} \|f - f^*\|_{\mathcal{F}},$$

where $\text{Sym}(P) = \{f \in \mathcal{F} | S_P(f) = f\}$ and $\|\cdot\|_{\mathcal{F}}$ is a norm on \mathcal{F} . This last criterion can be simply rewritten as :

$$\mathcal{E}'_2(P) = \|f - \Pi_{\text{Sym}(P)}(f)\|_{\mathcal{F}}, \quad (2.8)$$

where $\Pi_{\text{Sym}(P)}(f)$ is the projection of f onto the subspace $\text{Sym}(P)$. Then, following Kazhdan and colleagues [BCD⁺03], one can show that :

$$\mathcal{E}'_2(P) = \|f - (f + S_P(f))/2\|_{\mathcal{F}}^2 = \|f - S_P(f)\|_{\mathcal{F}}^2/4. \quad (2.9)$$

This again shows the equivalency between finding the minimum effort to turn a shape into a perfectly symmetrical shape (Eq. 2.8) and minimising the residual distance between f and $S_P(f)$ (Eq. 2.9).

Choice of f

One has to propose an efficient representation f for the point set under study. The method of Podolak and colleagues [PSG⁺06] described previously simply consists in designing f as the Gaussian Euclidean distance transform of X (which is a continuous function from \mathbb{R}^3 to \mathbb{R}^+). It could be interesting to try other models such as mixture models.

Choice of $\|\cdot\|_{\mathcal{F}}$

In the previous sections, the norm $\|\cdot\|_{\mathcal{F}}$ was chosen to be a classical 2-norm. However, the choice of this norm is crucial as it defines the notion of approximate reflection. Different norms leads to different interpretations and thus to different results.

- A first alternative to the 2-norm is to consider a robust cost function. This is particularly relevant if one considers the sum of point-to-point squared distances proposed by Zabrodsky. Indeed in this case, one can simply introduce a robust function ρ instead of the classical 2-norm, leading to a M-estimator of the form :

$$\mathcal{E}_2''(P) = \min_{x'_i \in X} \sum_{x_i \in X} \rho(\|S_P(x_i) - x'_i\|)/4,$$

In this case, the estimation of P will be performed by focusing on points belonging to highly symmetrical parts of X . This solution is pragmatic and convenient as, by choosing an adequate function ρ (e.g. a Leclerc function), efficient and simple algorithmic solutions are provided [CHW⁺08b]. Note that $\sum \rho(\|\cdot\|)$ does not define a norm and thus we cannot necessarily be interpreted as a symmetry distance as defined by Zabrodsky *et al.*

- Another choice is to design a norm that consists of a measure of stretch and bend between the object and its symmetrical image. Samir and colleagues [SSDK09] proposed to implement this idea rewriting Pb 2.9 by using an arbitrary plane K and optimising an intermediate rigid-body transformation R (similarly to what is done in Sec. 2.2.2) :

$$\mathcal{E}_2'''(R) = \arg \min_R \|f - R \circ S_K(f)\|_{\mathcal{F}} \quad (2.10)$$

Then one has to describe F and $\|\cdot - \cdot\|_{\mathcal{F}}$ and to provide minimisation strategies for Criterion 2.10. In practice, there exists no convenient framework allowing to directly design \mathcal{F} as the set of 3D point sets or of 3D surfaces. However, if one considers \mathcal{F} as the set of 3D close curves, the framework of elastic deformations [Jos07] provides efficient solutions to compute $\|\cdot - \cdot\|_{\mathcal{F}}$ and R (and thus an optimal symmetry plane P). Then, Samir *et al.* proposed to deal with surfaces by representing them as an indexed collection of concentric curves. The minimisation over R is performed other each of the generated curves independently and the resulting transformations are averaged to compute a global rigid-body transformation R .

Although this view provides an interesting alternative to the classical 2-norm, the approach is limited and is likely to be inadequate for many applications. First, in essence this method is not robust to outliers. Moreover, the representation proposed for 3D surfaces is not straightforward and is debatable as the use of the set of indexed curves imposes i) a constraint on the correspondences (*i.e.* the lines having the same index are necessarily matched with each other) and ii) to build a set of concentric closed curves on the mesh, which clearly limits the method to a certain class of concentric and well-balanced structures.

2.2.5 Advantages and limitations

Superimposition based methods rely on a pragmatic formulation of the expected reflections. More specifically, these methods depend on the expression of the departure from a point (or the complete object representation) to its symmetrical correspondent.

As many expressions (*e.g.* based on M-estimators or elastic deformations) can be proposed, the framework is very adaptable. The brute force optimisation of the underlying criterion is of high complexity (mainly because of the underlying pairing process). However, two efficient optimisation strategies have been proposed. The first one is based on a discretisation of the space of the planes and on an importance sampling strategy. It is not initialisation dependent but suffers from problems due to memory usage and of inaccuracy. The second one, based on the ICP algorithm, is more accurate and computationally attractive but is initialisation dependent. In practice, these two methods can be used successively to obtain fine estimates of all the significant reflections of a 3D object.

2.3 Approaches based on planes clustering

2.3.1 Basics

The general view of this approach is to define the optimal plane P as **“the plane with respect to which the most points are mirror-symmetrical”**. This simply writes as :

$$\mathcal{E}_3(P) = \text{card}(x_i \in X \text{ such that } S_P(x_i) \in X)$$

FIG. 2.4 – 2D illustration of "the plane with respect to which the most points are mirror-symmetrical".

For each pair of points (x_i, x_j) of X^2 ($i \neq j$), there is a unique plane $P_{i,j}$ such that $x_i = S_{P_{i,j}}(x_j)$ (and reciprocally). The method consists in accumulating local evidences of reflection by generating all such planes $P_{i,j}$ considering all possible pairs of points in X (there are $(\text{card}(X)^2 - \text{card}(X))/2$ of them). Then reflections with the most occurrences are considered as significant reflections of X .

In practice, when a surface is only approximately symmetrical, it is likely that each plane $P_{i,j}$ will occur only once. In this case, however, potential approximate symmetry planes will tend to form clusters in the space of \mathbb{R}^3 planes. The approach thus consists in finding significant clusters in this space and estimating a representative for each of these clusters.

2.3.2 Clustering

The key idea under this clustering process is to i) model the distributions followed by the $P_{i,j}$ s and ii) extract significant local maxima of this distribution. In practice, two

major (and related [Dah09]) methods can be used.

- The first method, based on the Hough transform, [Lea93] performs the estimation of the optimal planes P by considering a discrete 3-dimensional histogram defined on the space of variables defining the planes. For each one of the $P_{i,j}$ s, this histogram is incremented. Then its local maxima allow to estimate the parameters of the planes. Several ways to increment the discrete 3-dimensional histogram have been proposed in the literature. Given a sample $P_{i,j}$, the simplest one consists in incrementing by one the value of the bin representing $P_{i,j}$. This method is computationally very efficient but of limited accuracy. Indeed, the discretisation steps have to be large enough to group close planes but small enough to allow accuracy.

- The second method is based on the kernel density estimation method. It consists in approximating a continuous density function (*i.e.* Parzen window estimate) from the discrete sample ($P_{i,j}$) using a kernel function. Using an exponential kernel, it writes :

$$f(P|P_{i,j}) \propto \sum_{x_i} \sum_{x_j} \exp(\mathcal{D}(P_{i,j}, P)^2 / (2h^2)),$$

where \mathcal{D} is a distance function in the space of the planes of \mathbb{R}^3 . Then significant local maxima (*i.e.* modes) of f are considered as plausible symmetry planes of X . In practice, the Mean Shift algorithm [CM02] allows to find all the modes of f . The parameter h^2 can be viewed as the expected variance of locations in the sample belonging to a same mode. By building a continuous function f from the $P_{i,j}$ s, the problem due to the discretisation process of the Hough-based approach is avoided. However, the value of h can have a critical impact on the results and its choice is decisive [CM02]. Finally, note that this method is computationally more demanding than the Hough transform.

For each of these two clustering methods, some of the extracted reflections are likely to have no physical meaning. As a consequence, each one of them are tested with respect to the complete dataset X .

2.3.3 Improvements

In practice, two strategies have been implemented to limit the search space and improve the estimation of P :

- the first strategy consists in working only on a small subset of points instead of the complete set X . The considered subset of points is chosen to describe the object sufficiently well.
- the second consists in testing the compatibility between each pair of points (x_i, x_j) using differential properties of the surface around them. If the test fails, $P_{i,j}$ is considered as not plausible and is removed from the estimation. Note that in order to perform this comparison efficiently, the subsampling strategy described above must select points having non ambiguous differential properties.

In practice, Mitra and colleagues [MGP07] implement these ideas by i) discarding points having almost equal maximal and minimal curvatures and ii) considering $P_{i,j}$ as not plausible if curvatures of points x_i and x_j are too different from each other.

2.3.4 Advantages and limitations

In essence, the methods based on plane clustering are designed to find all the significant reflections of an object. The use of efficient subsampling and testing strategies allows efficient estimations but also implies a lack of accuracy. Moreover, the underlying clustering methods are parameters sensitive (especially the variance h^2 of the Mean shift approach) and thus are not well-suited to design a robust definition of the bilateral symmetries of a 3D object.

2.4 Other methods

2.4.1 EGI

This method is based on the computation of an orientation histogram built using the normals generated from the data, called the extended Gaussian image (EGI). This histogram is considered as a discrete sphere in which each facet contains the number of occurrences of a given orientation on the surface of interest [SS97]. More recently, Pan et al. [PWQW06] have proposed to weigh each facet with the inverse of the Gaussian curvature of the point, allowing to give less importance to areas where there is an important variation of the normal. This EGI has the advantage of being an adequate representation of the feature under study. Indeed, considering that the histogram is the discrete representation of a continuous domain, it can be properly smoothed, allowing to provide a certain robustness to noise. If an object is symmetric, then so is its EGI. Thus, the optimal plane is searched around principal axes one by one by maximising a correlation measure between the EGI and its reflection with respect to this plane (using exhaustive search).

Limitations :

This method is quite simple as it allows to turn the problem of estimating the symmetry plane of any 3D point set to that of finding the symmetry plane of an EGI. However, this method is quite restrictive. Indeed, first of all it is restricted to convex structures in which the notion of neighbourhood (needed to compute the normal or the curvature) is clearly defined. Moreover, it is tessellation-dependent. Normals and curvatures are quantities quite much more sensitive to noise than the point locations. Finally, this approach implicitly assumes that the object under study has only weak asymmetries : any strong deviation from perfect symmetry (either normal, pathological, or caused by image artifacts such as occlusions) is likely to bias the estimation. Moreover, some implementation choices are debatable : in the refinement process, the estimation of n is uncorrelated with the estimation of d which is clearly inappropriate. The technique of minimisation used is an exhaustive search within a discrete set of planes that is inevitably made on discrete domain, which clearly implies a loss of accuracy.

2.5 Conclusions

This taxonomy allowed to elucidate the nature of the different symmetry plane estimators. In particular, we showed that the principal axes analysis is a special case of the generalised moment analysis recently designed by Martinet *et al.* and that a variant of the well-known ICP algorithm minimises a "relaxed" version of the symmetry distance designed by Zabrodsky *et al.* Moreover, one notices that each of the 3 approaches we identified has proper characteristics :

- Distribution based methods are quite simple, do not require initialisation and are computationally efficient but likely to fail when dealing with large deviation from (multiple) perfect global symmetries.
- Superimposition based methods are particularly flexible and adaptive. Moreover they can be associated to efficient and accurate algorithmic solutions when one has a prior knowledge on the solution.
- Clustering based methods do not require initialisation and are particularly suited to the estimation of all the significant reflections of an object but are subject to inaccuracy and to a strong dependency on the minimisation parameters.

Chapitre 3

Symmetry plane estimation

Based on the previous taxonomy, we develop a method to estimate all the partial and approximate reflections of a 3D point set X in a robust and accurate way. More specifically, this method is composed of two sub-methods. The first one (Section 3.1) is based on a new method of planes clustering and allows to estimate coarsely all approximate and partial symmetries of X without any prior on the solution. Using these planes as initialisation, the second one (Section 3.2) is based on a new superimposition method and relies on a robust ML estimator for the optimal symmetry planes. This estimator is computed using the EM algorithm initialised with each of the previously identified meaningful symmetries and leads to an accurate and robust estimate of P . In this context, the local convergence of the EM is beneficial as it allows to estimate several potential symmetries from different initialisations.

3.1 Coarse detection of plausible reflections

3.1.1 Subsampling using keypoints

Based on the curvature information, we devise a keypoint detector, which allows to focus on meaningful subsets of points while drastically reducing the computation time of the overall algorithm. The principal curvatures [DC76] (called κ_1 and κ_2) are invariant to isometries (and thus to reflections), and thus they have the potential to be used to design such a detector, that must be i) distinctive, *i.e.* detecting points very different from their neighbourhood and ii) relatively insensitive to noise and small distortions of the point set (*e.g.* due to natural deviation from perfect symmetry or imperfect digitisation). Both characteristics are achieved as follows :

▷ First, we discard umbilic points (points for which $|\kappa_1| = |\kappa_2|$) from the computation. For these points (and more generally for points with $|\kappa_1| \approx |\kappa_2|$), directions of maximal and minimal curvature are ambiguous and not well defined. Thus, we only keep points such that the ratio $|\kappa_1/\kappa_2|$ is higher than a threshold κ_m . In particular, flat parts of the surface are not very informative and contain many such umbilic points.

▷ Second, we keep points with high maximal curvature that mainly correspond to salient features of the surface. In practice we select points whose $|\kappa_1|$ is higher than

a threshold κ_M . To account for the fact that some regions of the surface are likely to exhibit many less keypoints than others, and then could be unfairly excluded for the overall symmetry detection process, we adapt the previous selection procedure by i) cutting the surface into contiguous patches of similar size (we use a size of 40) ii) setting a local threshold κ_M for each patch. The figure 3.3 shows the location of the keypoints on various objects.

To achieve robustness to noise and to small distortions, the curvature properties of a given point x are computed from a quadratic surface that is approximated in the neighbourhood of x [GI04].




FIG. 3.1 – Location of keypoints on various objects.

3.1.2 Matching keypoints and weighing pairs

We recall that for two different keypoints x_i and x_j there exists a single reflection $S_{P_{i,j}}$ such that :


$$x_i = S_{P_{i,j}}(x_j) \text{ (and reciprocally } x_j = S_{P_{i,j}}(x_i))$$

The plane $P_{i,j}$ can be simply tested for the hypothesis that x_i and x_j are mirror images of each other. The principal curvatures at points x_i and x_j are invariant to isometries, and the Darboux frame is not, but both can be used to weigh $P_{i,j}$ as a potential symmetry plane of the surface. This is done by checking the compatibility between the augmented Darboux frame at x_i and the flipped augmented Darboux frame at x_j . In practice, to cope with the ambiguities on the orientation of the Darboux frames, the weight $c_{i,j}$ can be computed as (normals n and principal directions e_1 and e_2 are considered to be unit vectors, and \wedge is the cross-product) :

where $0 \geq \beta \geq 1$ weighs the relative influence of the compatibility of the augmented Darboux frame over the compatibility of the curvatures for the computation of $c_{i,j}$.

Note that $c_{i,j} = c_{j,i}$. In addition, $c_{i,i}$ is set to an arbitrary value c_m to take into account that the point x_i can be a midline structure (and thus has no bilateral counterpart).

$$\begin{aligned}
c_{i,j} = & (1 - \beta) (|\kappa_1^{x_i}| - |\kappa_1^{x_j}| + |\kappa_2^{x_i}| - |\kappa_2^{x_j}|) \\
& + \beta (\arcsin |n^{x_i} \wedge S_{P_{i,j}}(n^{x_j})| \\
& + \arcsin |e_1^{x_i} \wedge S_{P_{i,j}}(e_1^{x_j})| \\
& + \arcsin |e_2^{x_i} \wedge S_{P_{i,j}}(e_2^{x_j})|)
\end{aligned}$$

FIG. 3.2 – 2D illustration of the computation of $c_{i,j}$.

3.1.3 Building, modifying and pruning the affinity matrix

The previous cost values $c_{i,j}$ are turned into affinity measures by setting $C_{i,j} \propto \exp(\frac{-c_{i,j}}{2\sigma^2})$ where σ^2 represents the expected variance of the noise on $c_{i,j}$: in essence, the higher σ , the flatter $C_{i,j}$. The matrix $(C_{i,j})$ encodes the probabilities that two keypoints are mirror images of each other. In the following, we call N the size of the matrix C .

3.1.3.1 Enforcing global constraint on C

In the previous section, we built a matrix c (and C) considering each pair of points independently. However, in some applications, we have a knowledge about the global behaviour of the matches and thus on the structure of the matrix. Adding constraints on the structure of the matrix can be considered as the following problem : build a matrix as closed (in a sense to be defined) as possible from the symmetric matrix C and respecting some properties encoding the global constraint.

Consider the case when we are looking for a single symmetry plane (like in face recognition/medical imaging analysis) ; it can be interesting to specify that each point has a single counterpart (one-to-one matching) and consequently to impose the matrix C to be doubly stochastic. This way, the matrix will be different from the one proposed and will take into account potential incoherency between lines and columns probability. Several methods [GHR98, BCT07, Sin64] can be used for such a task and their respective merits have to be evaluated.

3.1.3.2 Pruning C

In terms of computational time, it can be very constraining to consider the $N \times N$ matchings whereas some of them have a very low $C_{i,j}$ value. Thus, we choose to cancel some coefficients of the matrix via :

- either a best one rule :
$$\begin{aligned}
\forall i, j \quad C_{i,j} = 1 & \text{ if } j = \arg \max_k C_{i,k} \\
C_{i,j} & = 0 \text{ else}
\end{aligned}$$
- or a threshold rule :
$$\begin{aligned}
\forall i, j \quad C_{i,j} = 1 & \text{ if } C_{i,j} \geq \delta \\
C_{i,j} & = 0 \text{ else}
\end{aligned}$$

Note that after this step, $C_{i,j}$ is a binary matrix.

3.1.4 Clustering mirror symmetries

In order to find clusters of $P_{i,j}$ s, we first build a Parzen window estimate from the discrete samples $P_{i,j}$ using a kernel function and then compute its modes. When choosing an exponential kernel function, the Parzen window estimate of the probability density function of the $(P_{i,j})$ is given by :

$$p(P|C) \propto \sum_{x_i} \sum_{x_j} C_{i,j} \cdot \exp(-\mathcal{D}(P_{i,j}, P)^2/(2h^2)), \quad (3.1)$$

where h^2 can be viewed as the expected variance of locations in the sample belonging to a same mode. The first issue consists in properly designing the distance $\mathcal{D}(P_{i,j}, P)$. The second issue consists in finding local maxima of the resulting pdf p . In fact, both issues are related as the optimisation method used to find local maxima of p results from the expression of \mathcal{D} . In the following, we investigate several propositions for \mathcal{D} and for the subsequent optimisation strategy. None of these solutions is completely satisfying.

3.1.4.1 Solution 1 : Euclidean distance on parameters

In previous works, Mitra *et al.* [MGP07] proposed to consider the set of planes $\in \mathbb{R}^3$ as a 3-dimensional Euclidean space for which each dimension represents one of the three parameters α , β and d . By considering a different kernel for each dimension, the formula 3.1 can be rewritten as :

$$p(P = (\alpha, \beta, d)|C) \propto \sum_{x_i} \sum_{x_j} C_{i,j} \cdot \exp(-|\alpha_{i,j} - \alpha|^2/(2h_1^2)) \cdot \exp(-|\beta_{i,j} - \beta|^2/(2h_2^2)) \cdot \exp(-|d_{i,j} - d|^2/(2h_3^2)). \quad (3.2)$$

The bandwidth parameters h_1 , h_2 and h_3 can be distinct and one can find all modes of the pdf p using the classical mean-shift procedure [FH75]. We call the corresponding algorithm MeanShift1.

This solution seems natural and simple. However the considered distance between parameters does not provide a suitable distance between planes. This is particularly enlightening to consider the singular case where $\beta \rightarrow 0$: the normal n tends to point to the north pole of the unit sphere. In this case, all the planes $(\alpha, \beta \rightarrow 0, d)$ ($\forall \alpha$, and for a fixed d) will be very close to each other whereas \mathcal{D} will change quadratically with α . More generally, this distance depends on the coordinate system frame in which X is expressed.

3.1.4.2 Solution 2 : Riemannian distance on angular components

In fact, the set of planes of \mathbb{R}^3 can be seen as the Cartesian product of the unit 3-dimensional sphere (called S^3) and of a 1-dimensional Euclidean space : $P \in S^3 \times \mathbb{R}$. Each $P_{i,j}$ is represented in the joint space by the pair $(n_{i,j}, d_{i,j}) \in S^2 \times \mathbb{R}$. Note that, as we consider $d \in \mathbb{R}$ and no longer in \mathbb{R}^+ , two different parameter sets represent a single plane. This choice is needed for theoretical purpose and in practice, we will ensure a

single parametrisation of the planes by systematically imposing $d > 0$. An appropriate distance for each of these two spaces is :

$$\mathcal{D}(n_1, n_2) = \text{acos}(n_1^T n_2) \text{ and } \mathcal{D}(d_1, d_2) = |d_1 - d_2|.$$

The two resulting spaces are Riemannian manifolds (more precisely, a Lie group and an Euclidean space) for which exponential and logarithmic maps can be computed. The former are straightforward for the Euclidean part \mathbb{R} . For S^3 they are given by [BF98] :

$$\begin{aligned} \log_{n_2}(n_1) &= \text{acos}(n_1^T n_2) \cdot (n_1 - n_2 \cdot (n_1^T n_2)) / \|(n_1 - n_2 \cdot (n_1^T n_2))\|, \\ \exp_{n_2}(n_1) &= n_2 \cos(\|n_1\|) + n_1 \sin(\|n_1\|) / \|n_1\|. \end{aligned}$$

By considering two different kernels for the angular (n) and linear (d) components of the joint domain, the Equation 3.1 can be rewritten as :

$$p(P = (n, d) | C) \propto \sum_i \sum_j C_{i,j} \cdot \exp(-\text{acos}(n^T n_{i,j})^2 / (2h_1^2)) \cdot \exp(-|d - d_{i,j}|^2 / (2h_2^2)). \quad (3.3)$$

The bandwidth parameters h_1 and h_2 are distinct for the angular (n) and linear (d) components of the joint domain. The clustering procedure consists in estimating local maxima of p using the non-linear mean-shift procedure [SM09] that entirely relies on the operators \mathcal{D} , \log and \exp described above.

Algo MeanShift2: Riemannian mean-shift

```

∀ i, j let ni,j be the normal associated to plane Pi,j,
    let di,j be the distance to origin associated to plane Pi,j
for all Ck,l == 1
    let n = nk,l
    let d = dk,l
do
    ntemp =  $\frac{\sum_{i,j} C_{i,j} \exp(-\text{acos}(n^T n_{i,j})^2 / h_1^2) \cdot \exp(-|d - d_{i,j}|^2 / h_2^2) \cdot \log_n(n_{i,j})}{\sum_{i,j} C_{i,j} \exp(-\text{acos}(n^T n_{i,j})^2 / h_1^2) \cdot \exp(-|d - d_{i,j}|^2 / h_2^2)}$ 
    dtemp =  $\frac{\sum_{i,j} C_{i,j} \exp(-\text{acos}(n^T n_{i,j})^2 / h_1^2) \cdot \exp(-|d - d_{i,j}|^2 / h_2^2) \cdot d_{i,j}}{\sum_{i,j} C_{i,j} \exp(-\text{acos}(n^T n_{i,j})^2 / h_1^2) \cdot \exp(-|d - d_{i,j}|^2 / h_2^2)}$ 
    n = exp(ntemp)
    d = dtemp
while convergence is not achieved
retain (n, d) as a local mode
append close modes

```

Notice that as \mathcal{D} is not an Euclidean distance, p does not define a true kernel density [SM09].

3.1.4.3 Solution 3 : Point-to-plane distance and distance between angles

The Euclidean distance between the distance to origin d is problematic. First contrary to the distance between the normals, one can easily see that such a distance depends on the frame in which we express the studied planes, which is clearly inappropriate. Second, it has no significance when trying to compute a distance between secant planes. Indeed, as an illustration, the planes P_2' and P_2'' (and each plane parallel to them) have no reason to be at a different distance from P_1 but it is easy to see that $|d_{P_2''} - d_{P_1}| < |d_{P_2'} - d_{P_1}|$ (Fig. 3.3).




FIG. 3.3 – 2D illustration of the two problems related to the use of the parameter d for clustering.

However it is necessary to be able to distinguish planes having a small angular distance (*i.e.* planes that are almost parallel) but being distinct (*i.e.* not superposed) from each other. For that, we consider that an optimal plane P has to be close to the midpoints $\frac{x_i+x_j}{2}$ corresponding to the planes $P_{i,j}$ converging to the mode P . Thus, following the approach recently proposed by Dahyot [Dah09], one can express the pdf as :

$$p(P = (n, d)|C) \propto \sum_i \sum_j C_{i,j} \exp\left(-\frac{d(\frac{x_i+x_j}{2}, P)^2}{2h_1^2}\right) \quad (3.4)$$

$$p(P = (n, d)|C) \propto \sum_i \sum_j C_{i,j} \exp\left(-\frac{(d - (\frac{x_i+x_j}{2})^T n)^2}{2h_1^2}\right) \quad (3.5)$$

Moreover, we incorporate a measure of similarity between the normals of the planes :

$$p(P = (n, d)|C) \propto \sum_i \sum_j C_{i,j} \exp\left(-\frac{(d - (\frac{x_i+x_j}{2})^T n)^2}{2h_1^2}\right) \exp\left(-\frac{\text{acos}(n_{i,j}^T n)^2}{2h_2^2}\right) \quad (3.6)$$

However, due the non linear nature of the last term, the resulting optimisation would be difficult. Thus, we choose to replace the term by $|\text{acos}(n_{i,j}^T n)|^2$ by the quadratic cost $(1 - n_{i,j}^T n)^2$, leading to :

$$p(P = (n, d)|C) \propto \sum_i \sum_j C_{i,j} \exp\left(-\frac{(d - (\frac{x_i+x_j}{2})^T n)^2}{2h_1^2}\right) \exp\left(-\frac{(1 - n_{i,j}^T n)^2}{2h_2^2}\right) \quad (3.7)$$

The two underlying distances do not depend on the frame in which we express the planes. Then one can compute the derivatives of p with respect to d and n . This results in the following mean-shift algorithm :

Algo MeanShift3: Point to plane distance and angles between normals

$\forall i, j$ let $n_{i,j}$ be the normal associated to plane $P_{i,j}$,
 let $m_{i,j}$ be the midpoint associated to plane $P_{i,j}$,
 let $d_{i,j}$ be the distance to origin associated to plane $P_{i,j}$.
 for all $C_{k,l} == 1$
 let $n = n_{k,l}$
 let $d = d_{k,l}$
 let $m = m_{k,l}$
 do
 $\forall i, j$ $R_{i,j}(d, n) = \exp\left(-\frac{(d-m_{i,j}^T n)^2}{2h_1^2}\right) \exp\left(-\frac{(1-(n_{i,j}^T n))^2}{2h_2^2}\right)$
 $d = \frac{\sum_{i,j} C_{i,j} \frac{m_{i,j}^T n}{h_1^2} R_{i,j}(d, n)}{\sum_{i,j} C_{i,j} R_{i,j}(d, n)}$
 build the 3×3 matrix A such that $A_{k,l} = \sum_{i,j} C_{i,j} \left(\frac{m_{i,j}^k m_{i,j}^l}{h_1^2} + \frac{n_{i,j}^k n_{i,j}^l}{h_2^2} \right) R_{i,j}(d, n)$
 build the $3D$ vector b such that $b = \sum_{i,j} C_{i,j} \frac{m_{i,j}}{h_1^2} R_{i,j}(d, n)$
 $n = A^{-1}b$
 while convergence is not achieved
 retain (n, d) as a local mode
 append close modes

Note that, to the best of our knowledge, it is not possible to combine both the Riemannian approach of subsection 3.1.4.2 for n and the point to plane distance of the present subsection.

3.1.5 Checking the mirror symmetries

Significant modes are now considered as plausible reflections. However, some of these can have no physical meaning and it is important to test each plausible model with respect to the complete dataset X . For this purpose, we consider all pairs of keypoints x_i, x_j associated to a local maxima P during the mean-shift procedure. For each of these pairs, we test (recursively) if points neighbouring x_i are superposed with the symmetrical image with respect to the plane P of points neighbouring x_j . We compute an inlier rate from this test. If this rate is greater than a threshold $nbmin$, P is considered as a symmetry plane of X . This writes more formally as :

for all pairs of points (x_i, x_j) associated to cluster P

i) let N_{x_i}, N_{x_j} be respectively the set of points neighbouring x_i and x_j

ii) for all points x_k of N_{x_i}

if $\exists x_l \in N_{x_j}$ such that $\|S_P(x_k) - x_l\| \leq \delta$ increment $nbIn$

and add N_{x_k} to N_{x_i} and N_{y_l} to N_{x_j}

else increment $nbOut$

if $\frac{nbIn}{nbOut} \geq nbmin$, P is a symmetry plane of X

3.1.6 Results and evaluation

3.1.6.1 Parameters

The different parameters used in the previous sections are empirically set to :

- $\kappa_m = 0.9$,
- $\kappa_M = 0.9$ times the maximal curvature within the patch,
- $\sigma = 0.1$ times the extent of X ,
- $c_m = 0.1$,
- $\delta = 0.05$ times the extent of X ,
- $nbmin = 0.8$.

3.1.6.2 Some observations

MeanShift1 (Section 3.1.4.1) and MeanShift2 (Section 3.1.4.2) provide frame-dependent results. This is an undesirable property and makes these algorithms very difficult to evaluate, as the obtained results can be very different from each other by simply moving rigidly the object under study. By contrast, MeanShift3 does not depend on the frame in which we express the object under study and provides stable results. In the following, we use MeanShift3. In practice, it always provides quite good estimate of the different symmetries of the object under study. As previously mentioned, the bandwidth parameter has a strong influence on the results. In practice, we did not manage to set this parameter to a single value for all the object under study. In practice, we manually choose the parameter for a given structure and keep this choice for point sets corresponding to the same class objects (brain, ventricle, chair, *etc*). For a more automated approach, bandwidth selection methods could be investigated [CRM01].

As we design this algorithm to provide a coarse estimation of all the bilateral symmetries of an object, we do not investigate its accuracy.

3.1.6.3 Some results

We illustrate MeanShift3 by displaying some results obtained in Figure 3.4.




FIG. 3.4 – Some results obtained on real data for coarse estimation of the plausible meaningful reflections of an object

3.2 Accurate & robust estimation of a reflection

The previous algorithm (Section 3.1) allows to detect all the meaningful reflections of a 3D object in an efficient way. However in many cases, symmetry has to be estimated very precisely and in a robust way (*e.g.* for face recognition or preoperative facial surgery). In these conditions, we have to define properly what is the approximate symmetry plane of an object or of a part of it and to provide a good way to find this approximate symmetry plane in an accurate way. If we assume that we have a first estimate of the expected plane (for example provided by the clustering-based method we devised in the previous section), a good strategy for this purpose is the widely-used ICP-formalism (Section 2.2.3). We propose to extend and improve this formalism below to make it more accurate and more robust.

3.2.1 A fast, robust and accurate algorithm

As shown in Section 2.2.1, one can define the approximate symmetry planes of an object X as the significant local minima of :

$$\mathcal{E}_2(P) = \min_{x'_i \in X} \sum_{x_i \in X} \|S_P(x_i) - x'_i\|^2 \quad (3.8)$$

If one has a prior knowledge of the plane P , the considered local minima can be reached efficiently using an ICP-like algorithm. However, as the classical ICP algorithm, which is essentially a least squares minimisation, this ICP-like algorithm is not robust to outliers. Classical techniques used in the case of rigid registration can be adapted to deal with this issue [CSK05, KKM03]. However, another critical problem is that as for the classical ICP algorithm, \mathcal{E}_2 is based on discrete, one-to-one matches between the two sides of the object under study, relying on the closest point principle. In practice, this causes the criterion \mathcal{E}_2 to exhibit lots of undesirable local minima close to more relevant minima, which are prone to make the ICP-like algorithm (which has a monotonic, local convergence) fail. To deal with this limitation, in the context of rigid-body registration of point clouds, Rangarajan and colleagues proposed to define a new criterion using multiple, weighted matches [RMP⁺96]. They showed this criterion to be well-behaved and smoother, with a reduced number of local minima, than \mathcal{E}_2 . A simple way to implement this idea, and to introduce such fuzzy matches between points, is to cast the minimisation problem in a probabilistic framework, and then to use classical techniques of statistical inference to estimate the unknown transformation.

Following Chui & Rangarajan [CR03] and Granger & Pennec [GP02], who introduced these techniques in the context of rigid-body registration of point clouds, we propose an adaptation of the criterion \mathcal{E}_2 , based on mixture models, to compute an approximate symmetry plane of an object given an initialisation.

3.2.1.1 Algorithm

We define the *approximate symmetry plane* P of a point set X as the one best superposing X and its reflection $S_P(X)$ about P . For that, we consider X as a noised

version of $S_P(X)$. This allows to consider each point x_i of X as the realisation of a random variable whose distribution is a mixture model composed of $\text{card}(X)$ Gaussian laws $(\mathcal{N}(S_P(x_k), \sigma^2 I))_{x_k \in X}$. This leads to the following likelihood function :

$$L(P) = \prod_{x_i \in X} \sum_{x_j \in X} \pi_{i,j} p_j(x_i; P)$$

where $\pi_{i,j}$ are prior probabilities ($\forall i, \sum_j \pi_{i,j} = 1$) conveying the probability that the point x_i comes from the distribution $p_j(\cdot; P) = \mathcal{N}(S_P(x_j), \sigma^2 I)$ without knowing anything else.

The Expectation-Maximisation algorithm [DLR77a] (EM) can be used to maximise the likelihood function L and yields the very simple following iterative scheme (see Appendix B for more details about the derivation) :

Step 0 : Initialise \tilde{P}

E-step : $\forall i, j \tilde{A}_{i,j} = \frac{\pi_{i,j} \exp(-\|x_j - S_{\tilde{P}}(x_i)\|^2 / (2\sigma^2))}{\sum_k \pi_{i,k} \exp(-\|x_k - S_{\tilde{P}}(x_i)\|^2 / (2\sigma^2))}$

M-step : $\tilde{P} = \arg \min_P \sum_{(x_i, x_j) \in X^2} \tilde{A}_{i,j} \|x_j - S_P(x_i)\|^2$

Step 3 : if \tilde{P} has changed go to E-step else finish

This scheme can be seen as the successive estimation of a *match matrix* A describing the fuzzy point-to-point correspondences between X and $S_P(X)$ (E-step), and of the unknown reflection plane P (M-step). The noise variance σ^2 is not estimated and can be seen as a scale parameter, which controls the fuzziness of the match matrix A . Typically, it can be given a high initial value and then decreased in a multiscale framework for optimal performances.

The M-step can be solved efficiently using the general theorem we stated in Section 2.2.2. The ICP-like algorithm we presented in Section 2.2.3 can actually be seen as a simplified version of this EM algorithm. The point x_j with the highest value $A_{i,j}$ is the closest point of $S_P(x_i)$ in X , so the E-step is very similar to Step i). If the matrix A is made binary, by keeping only this closest point into account, then the M-step is identical to Step ii).

3.2.1.2 Rejection of outliers

The E-step involves the computation of the distance between each point $S_{\tilde{P}}(x_i)$, $x_i \in X$ and all the points x_j of X . Actually, when the point $S_{\tilde{P}}(x_i)$ is far away from X (which can happen for instance when x_i has no satisfying counterpart on the other side of the object), the values $A_{i,\cdot}$ can be not negligible because of the normalisation term (see E-step) and thus influence negatively the criterion to maximise. A simple way to tackle this problem is to reject the points x_j in X that are farther from $S_{\tilde{P}}(x_i)$ than a predefined cut-off distance, which amounts to give them a null weight $\tilde{A}_{i,j}$. This last process is equivalent to consider p_j as a truncated Gaussian density function instead of a classical Gaussian law. With this slight modification one can easily show that the algorithm still converges to a local maximum of the modified likelihood function. In practice, the threshold must depend on σ : there is a high uncertainty on the matches when σ is

large, while there is increased confidence on these matches when σ decreases. Thus we choose to set this threshold at the value $l\sigma$. The E-step then becomes :

-
- i) **Initialise** \tilde{A} as the null matrix
 - ii) $\forall x_i$ in X
 - $L_i = \{x_j \in X \text{ such that } \|x_j - S_{\tilde{P}}(x_i)\| < l\sigma\}$
 - $\forall x_j \in L_i, \tilde{A}_{i,j} = \exp(-\|x_j - S_{\tilde{P}}(x_i)\|^2/2\sigma^2)$
- Normalise** the i^{th} line of \tilde{A}
-

The set L_i is built by a search in a kd -tree space-partitioning built at the beginning of the algorithm. Note that this modification allows A to be sparse and reduces drastically both the computational and memory usage burden.

Notice that the subsequent algorithm can be seen as the iterative minimisation over A and P of the following criterion :

$$\sum_{i,j} A_{i,j} \rho_{l\sigma}(\|x_j - S_P(x_i)\|^2) + 2\sigma^2 \sum_{i,j} A_{i,j} \log(A_{i,j})$$

with $\forall i, \sum_j A_{i,j} = 1$ and where $\rho_\delta : r \mapsto r$ if $r < \delta$ and δ else (with $\delta > 0$).

In other words, it amounts to define the approximate symmetry plane as the one minimising the average sum of squared distances between the points on one side of the object and their reflection about the plane, but only among the pairs of points for which these distances are lower than a predefined threshold $l\sigma$.

3.2.1.3 Multiscale scheme & estimation of σ

The convergence of this new algorithm is very dependent on the parameter σ . A small σ allows to be very selective on the pairs of points to be significantly taken into account for the estimation of P . Thus, if the algorithm is given a good initial plane, a small σ allows good accuracy. If this is not the case, there is little chance that two mismatched points will ever be matched properly during the iterations. On the contrary, a large value of σ allows these symmetrical, but initially mismatched points, to have growing influence during the iterations. Based on this, to allow both accuracy and robustness, we follow a multiscale approach, by running successive EM algorithms with decreasing σ values. In this scheme, σ acts as a scale parameters, and allows to progressively refine the plane estimation. In practice, the initial scale factor σ_0 is successively divided by a constant value $f > 1$ until it is equal to, or lower than a predefined final value σ_f .

The initial σ_0 can be initialised thanks to the output of the plane clustering method. For that we consider σ_0 as the Hausdorff distance between the two sets N_{x_i} and $S_P(N_{x_j})$ computed in Section 3.1.5.

Note that this approach is slightly different from that of Granger & Pennec, who update the scale factor at each iteration of the EM algorithm. Experimentally, we observed that letting the EM algorithms converge and choosing a higher f value led to better results in our case, in terms of both robustness and accuracy, without significantly increasing the computational time. This multiscale scheme can be seen as a deterministic annealing procedure where σ is analogous to the temperature parameter.

3.2.1.4 Coarse-to-fine approach

At the beginning of this multiscale scheme, large values of σ only lead to a gross estimation of the unknown plane. Consequently, it is useless to take the entire point set X into account at these stages. As a result, we propose a coarse-to-fine approach, where X is decimated at large σ values, and then refined progressively when σ decreases. We use the same technique as Granger & Pennec to compute the successive decimated clouds, named X' , in which the points x_i are chosen. At a given σ value, we iteratively merge each point of X included in spheres of radius $k\sigma$ (where k is a constant to define). Each of these spheres finally contains N_m points (with $\sum_m N_m = N = \text{card}(X)$), which are then replaced by their centroids. The set of centroids, each one being given a weight N_m (later used in the estimation of P), constitutes the decimated cloud used at scale σ . Note that the original set X is entirely taken into account when $k\sigma$ is lower than or equal to the minimal resolution of X .

Note that Granger & Pennec [GP02] have experimentally shown (in case of rigid-body registration) that the transformation yielding the global maximum likelihood is different from that globally minimising the ICP criterion. Tsin & Kanade [TK04] have actually shown that in case of two identical, aligned point clouds, the identity transformation is usually not a stationary point of the likelihood function. The same can be said about our problem, where we look for a reflection best superposing the cloud with itself. However, the bias between the correct solution and the ML estimate is virtually null when σ is very small, as at the end of the multiscale scheme.

3.2.1.5 Algorithm

Finally, we have the algorithm :

Algo Sym1: multiscale robust EM
Initialise \tilde{P} and σ Repeat Decimation : $X' =$ decimation of X with a radius $k\sigma$ Repeat E-step : see Section 3.2.1.2 M-step : $\tilde{P} = \arg \min_P \sum_{x_i \in X'} \sum_{x_j \in X} N_i \tilde{A}_{i,j} \ x_j - S_P(x_i)\ ^2$ Until \tilde{P} does not change Decrease scale factor $\sigma = \max(\sigma/f, \sigma_f)$ Until \tilde{P} does not change and $\sigma \leq \sigma_f$

3.2.2 Symmetrisation of A

As previously mentioned, $A_{i,j}$ can be seen as a measure of whether the point x_i is the bilateral counterpart of x_j or not. Following this observation, it would be natural to impose that $A_{i,j} = A_{j,i}$, $\forall(i, j)$. However, there is no simple way to impose such a constraint on our large sparse matrix A efficiently.

However, one can use the fact that the EM algorithm described in Section 3.2.1.2 can be seen as the iterative minimisation over A and P of the following criterion :

$$\sum_{i,j} A_{i,j} \rho_{l\sigma}(\|x_j - S_P(x_i)\|^2) + 2\sigma^2 \sum_{i,j} A_{i,j} \log(A_{i,j})$$

with *forall* i , $\sum_j A_{i,j} = 1$. We propose to restate the symmetry plane estimation as :

$$\begin{aligned} \tilde{P} = \arg \min_{P,A} & \sum_{i,j} A_{i,j} \rho_{l\sigma}(\|x_j - S_P(x_i)\|^2) + 2\sigma^2 \sum_{i,j} A_{i,j} \log(A_{i,j}) \\ & + \sum_{i,j} B_{i,j} \rho_{l\sigma}(\|x_j - S_P(x_i)\|^2) + 2\sigma^2 \sum_{i,j} B_{i,j} \log(B_{i,j}), \end{aligned}$$

with $\forall i$, $\sum_j A_{i,j} = 1$ and $\forall j$, $\sum_i B_{i,j} = 1$. This criterion is equal to :

$$\sum_{i,j} (A_{i,j} + B_{i,j}) \rho_{l\sigma}(\|x_j - S_P(x_i)\|^2) + 2\sigma^2 \sum_{i,j} A_{i,j} \log(A_{i,j}) + 2\sigma^2 \sum_{i,j} B_{i,j} \log(B_{i,j}),$$

with $\forall i$, $\sum_j A_{i,j} = 1$ and $\forall j$, $\sum_i B_{i,j} = 1$. One can easily show that the matrix $(A_{i,j} + B_{i,j})$ is symmetric.

This new criterion can be minimised by the following algorithm :

Algo Sym2: symmetric multiscale robust EM

Initialise \tilde{P} and σ

Repeat

Repeat

E-step : compute $A_{i,j}$ and $B_{i,j}$

M-step : $\tilde{P} = \arg \min_P \sum_{x_i \in X, x_j \in X} (\tilde{A}_{i,j} + \tilde{B}_{i,j}) \|x_j - S_P(x_i)\|^2$

Until \tilde{P} does not change

Decrease scale factor $\sigma = \max(\sigma/f, \sigma_f)$

Until \tilde{P} does not change and $\sigma \leq \sigma_f$

This last algorithm implies that A and B are squared matrices. In particular, it suggests that we no longer use decimate versions of X . Thus to limit the computational burden, we apply this symmetric constraint only when σ (and thus the cut-off distance $l \times \sigma$) becomes small.

3.2.3 Results and evaluation

3.2.3.1 Parameters & algorithm

The different parameters for both Sym1 and Sym2 presented in the previous sections are empirically set to :

- σ_0 estimated from the clustering based method,

- threshold for outliers rejection $l = 3$,
- sphere radius for subsampling $k = 1$.

We compare the proposed methods (Sym1 and Sym2) with two methods based on a robustified ICP-like algorithm that we do not present in this document (see [CHW⁺08b] for more details). These two methods rely on replacing the classical least squares estimator either with a Leclerc estimator (leading to what we call a mICP) or a trimmed estimator (leading to what we call a trICP).

3.2.3.2 Evaluation on symmetrical data

In this section, we investigate the accuracy and capture range of Sym1 and Sym2 on perfectly symmetrical data that we generate by i) positioning a symmetry plane manually, ii) replacing all the points located on the right side of the plane by the mirror image of the points located on the left side. We work on a face (100k points, bounding box of about 200mm \times 200mm \times 200mm), a ventricle (10k points, bounding box of about 80mm \times 35mm \times 50mm) and a virtual endocast (100k points, bounding box of about 95mm \times 115mm \times 85mm). To apply the same process on all data, we rescale them to the size of the face data. We then apply angular offsets between 0 and 40 degrees and linear offsets between 0 and 60 mm to the ground truth symmetry plane and use it to initialise our algorithm. After convergence, we compute the angular and linear errors (called respectively θ and τ) of the estimated plane compared to the ground truth solution. For large linear offsets (below 60 mm) and large angular offsets (below 31 degrees), both Sym1 and Sym2 always converge to a plane for which θ and τ are lower than 10^{-15} . In practice, we did not observe differences of performance between Sym1 and Sym2.

3.2.3.3 Evaluation on asymmetrical data

In this section, we evaluate the robustness and accuracy of Sym1 and Sym2 on asymmetrical data. For this purpose, we add artefacts to the perfectly symmetrical face, ventricle and endocast point sets of the previous section whose symmetry plane is still considered as the ground truth. We generate artefacts as follows :

- ▷ Noise is modeled as isotropic Gaussian white, with variance δ^2 and is added to each point.
- ▷ Occlusions are generated by removing a given quantity of adjacent points. We generate between 1 and 4 areas of occlusion this way on each data.
- ▷ Asymmetries are generated by randomly choosing a point x_D close to the surface and deforming each point x of the cloud to a new position x' according to :

$$\triangleright x' = x + K \times G_v(x - x_D) \frac{\overrightarrow{x - x_D}}{\|\overrightarrow{x - x_D}\|}$$

where G_v is a 3D non-normalised Gaussian function of variance v^2 and K is the deformation strength. We generate on a given data between 3 and 8 randomly located asymmetric areas this way.

	mICP	trICP	Sym1	Sym2
max τ	0.61	5.14	0.12	0.11
max θ	1.2	2.82	0.15	0.13
mean (θ, τ)	(0.12,0.65)	(0.76,1.14)	(0.06,0.04)	(0.06,0.04)
var (θ, τ)	(0.02,0.01)	(0.75,0.16)	(10^{-3} ,0.01)	(10^{-3} ,0.01)

TAB. 3.1 – Statistics on (θ, τ) for mICP, tICP [CHW⁺08b], Sym1 initialised with tICP and Sym2 initialised with tICP.

By randomly combining these artefacts, we generate a set of 150 surfaces with varying levels of artefacts (for the face, the rescaled ventricle and the rescaled endocast). The parameters are chosen such that $(K, v^2) \in [0, 20] \times [0, 25]$, 0 to 20% of outliers and a noise of variance $\delta^2 = 0.3$. Examples of data built this way with the resulting estimated planes are shown on Fig 3.5.



FIG. 3.5 – Estimation of the symmetry plane on ground truth, asymmetrical data. Estimation of the plane by Sym2 on simulated data.

Tab. 3.1 and 3.2 show statistics over the 3×150 experiments for mICP, trICP, Sym1 and Sym2. We observe that over the 150 experiments, Sym1 and Sym2 have a very low mean and maximum errors for both θ and τ , with very low variance. These values are higher for trICP and mICP, but we also demonstrate the relevance of using trICP to initialise Sym1 in terms of computational time.

To assess the breakdown point of Sym1 and Sym2, we increase drastically the size, number of the deformations and of the occlusions and we observe that our estimator always leads to a good estimate for \tilde{P} (with an error less than 0.5 mm and 0.5 degrees) when the number of inliers ¹ is higher than 30%. Beyond this breakdown point, the error begins to increase until it reaches errors of 10 degrees and 8 mm. Figure 3.6 shows the average errors (obtained from a set of 200 experiments performed on the endocast data) with percentage of outliers going from 0 to 95 %. Note that Sym1 gives similar results.

¹We consider a point x as an inlier if there exists a point x_j of X such that $\|S_{\hat{P}}(x_j) - x_j\|^2 < l \times \sigma_{final}$, where \hat{P} is the ground truth plane.




FIG. 3.6 – Mean angular and linear errors with respect to the percentage of outlying points in the data.

	mICP	tICP	Sym2	tICP + Sym2
mean t (100k points)	80s	30s	600s	160s
max t (100k points)	120s	80s	1360s	380s
mean t (10k points)	18s	6s	40s	12s
max t (10k points)	36s	12s	88s	35s

TAB. 3.2 – Statistics on run time t for mICP, tICP, Sym2 and Sym2 initialised with tICP performed on a standard PC with an Intel Core Duo T7700 at 2.4GHz with 2GB Ram.

3.2.3.4 Results on real data

In Fig. 3.7, we display the symmetry plane obtained with Sym2 for a woman’s face with a large area of missing points on the left cheek, the Stanford bunny, a chair with a missing leg, a chimpanzee endocast, a human brain cortex and a human brain ventricle.

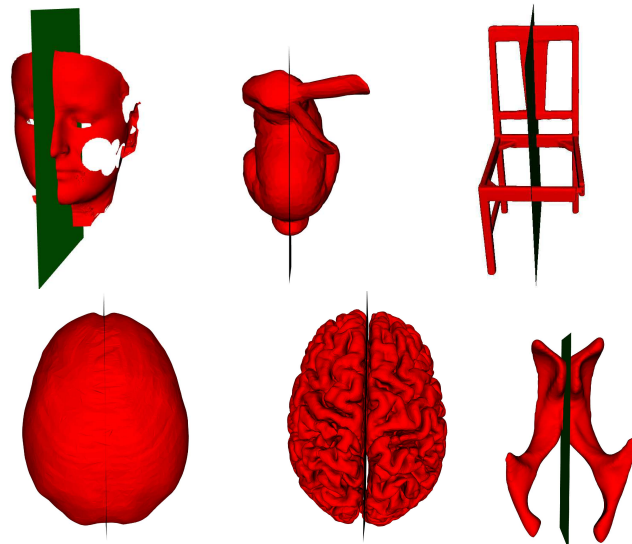


FIG. 3.7 – Estimation of the symmetry plane using Sym2 on different structures.

Chapitre 4

Symmetry surface estimation

In this section, we extrapolate the notion of symmetry about a plane to that of symmetry about a differentiable 3D surface C . Such a symmetry provides an alternative suitable description of the departure from a perfect symmetry and can be used for example to quantify asymmetries or to retrodeform a bilateral object subject to a global deformation (such as a human fossil skull with taphonomic deformations).

4.1 Extension to symmetry surface

4.1.1 Formulation

To extrapolate the notion of symmetry about a plane to that of symmetry about a differentiable 3D surface C , we define the reflection of a point x_i about a surface C as the reflection of this point about its projection on this surface (Fig. 4.1). Following from the projection theorem¹, it can be readily demonstrated that the plane i) orthogonal to the line linking x_i and its projection and ii) passing through this projection is tangent to the surface C . This observation allows to generalise the notion of symmetry with respect to a plane by considering locally the symmetry about a surface C as a classical symmetry about a plane (see Figure 4.1).

This *symmetry surface* is modeled as a smooth parametric function (we simply choose a low-order bivariate polynomial function $x^w = C(x^u, x^v)$, noting x^u, x^v, x^w the three spatial coordinates of point x). Then, following the minimisation sketch proposed in the previous section, the problem of finding the best "symmetry surface" C of X can be written as :

$$\tilde{C} = \arg \min_{C,A} \sum_{x_i, x_j \in X^3} A_{i,j} \rho_{l\sigma}(\|x_j - S_C(x_i)\|^2) + 2\sigma^2 \sum_{i,j} A_{i,j} \log(A_{i,j})$$

with $\forall i$

¹Let \mathcal{H} be a Hilbert space and \mathcal{M} a closed subspace of \mathcal{H} . Corresponding to any vector x in \mathcal{H} , there is a unique vector m_0 in \mathcal{M} such that $\|x - m_0\| \leq \|x - m\|$ for all $m \in \mathcal{M}$. Furthermore, a necessary and sufficient condition that $m_0 \in \mathcal{M}$ is the unique minimising vector is that $x - m_0$ is orthogonal to \mathcal{M} .




FIG. 4.1 – **A curved symmetry surface** : 2D illustration of the curved symmetry surface.

, $\sum_j A_{i,j} = 1$ and $\forall i, j A_{i,j} \geq 0$ where $S_C(x_i)$ defines the local symmetry transformation for x_i in X . We propose to minimise the previous criterion using a strategy similar to that proposed in the previous section :

Initialise \tilde{C} and σ

Repeat

Decimation : $X' =$ decimation of X with a radius $k\sigma$

Repeat

E-step : $\tilde{A} = \arg \min_A \sum_{x_i \in X'} \sum_{x_j \in X} N_i A_{i,j} \|x_j - S_{\tilde{C}}(x_i)\|^2 + 2\sigma^2 \sum_{i,j} A_{i,j} \log(A_{i,j})$

M-step : $\tilde{C} = \arg \min_C \sum_{x_i \in X'} \sum_{x_j \in X} N_i \tilde{A}_{i,j} \|x_j - S_C(x_i)\|^2$

Until \tilde{C} does not change

Decrease scale factor $\sigma = \max(\sigma/f, \sigma_f)$

Until \tilde{C} does not change and $\sigma \leq \sigma_f$

For the E-step, we propose an optimal solution :

E-step :

Initialise \tilde{A} as the null matrix

$\forall x_i$ in X'

Project x_i on \tilde{C} at $p_{\tilde{C}}(x_i)$.

Compute $S_C(x_i) = p_{\tilde{C}}(x_i) + (p_{\tilde{C}}(x_i) - x_i)$

$L_i = \{x_j \in X' \text{ such that } \|x_j - S_{\tilde{C}}(x_i)\| < l\sigma\}$

$\forall x_j \in L_i, \tilde{A}_{i,j} = \exp(-\|x_j - S_{\tilde{C}}(x_i)\|^2/2\sigma^2)$

Normalise the i^{th} line of \tilde{A}

To the best of our knowledge, the M-Step has no simple solution. We choose to approximate it by the following procedure :

M-step : \tilde{C} is the surface best fitting the mid-points $\tilde{A}_{i,j}(x_i + x_j)/2$:

$$\tilde{C} = \arg \min_C \sum_{i,j} \tilde{A}_{i,j} \|C((x_i^u + x_j^u)/2, (x_i^v + x_j^v)/2) - (x_i^w + x_j^w)/2\|^2$$




FIG. 4.2 – **A curved symmetry surface** : 2D illustration of the approximate solution proposed for the M-step. \tilde{C} is the surface best fitting the mid-points computed from the correspondences (x_i, y_i) .

This choice is illustrated in Figure 4.2. This algorithm does not ensure to converge to a solution of the problem. However, in order to ensure its good behaviour : i) we first model C with a second degree polynomial and then increase its order through the iterations and ii) we propose to add landmark points belonging to the midline structure to the criterion to constrain the solution appropriately.

Note on large bending : Notice that our polynomial model has two strong limitations that in practice are only problematic when the object under study exhibits a symmetry with respect to a surface C with a large global bending.

First, C is modelled as a polynomial. This imposes that for each value (x^u, x^v) there exists a unique x^w . In practice, this limits C to the modeling of surfaces subject to a relatively small global bending.

Secondly, we designed the reflection of x_i with respect to C by considering the symmetry with respect to the plane orthogonal to the line linking x_i and its projection $p_C(x_i)$ on C . However, in practice this choice can be seen as quite arbitrary as there are potentially several points c of C such that their tangent is orthogonal to the line linking c and x_i . In particular, this aspect implies that our transformation is not ensured to respect the equality $S_C(S_C(x)) = x$.

In practice, it is easy to observe that for a small deviation of the surface from a perfect plane, the reflection we have defined is the most appropriate and that, as a result, our choice is justified.

4.1.2 Initialisation

In practice, this algorithm is quite sensitive to the provided initial value for C . We propose the following method to provide a relevant initial function C :

- First, we compute the approximate symmetry plane P of the point set X (using Sym1).

- we parcellate the object into E pairs of patches $(X^1, \dots, X^e, \dots, X^E)$ that are face-to-face with respect to the previously estimated plane.
- Then, for each of these subsets X^e , we estimate an approximate symmetry plane P^e (using Sym2).
- Finally, we approximate a surface C from the set of the computed planes in a robust manner (parametrising P as $P(x, y) = z$) :

$$C^{init} = \sum_{c \in 1, \dots, C} \sum_{(x_i, y_i) \in P^c} \rho(\|C(x_i, y_i) - P^c(x_i, y_i)\|)$$

where ρ is a Leclerc function allowing a robust fitting of the surface C .

4.1.3 Measuring the curvature of C

Once C is estimated, studying its curvature can be of interest. The different directional curvatures of C are simply given by its Hessian matrix.

4.1.4 Some results




FIG. 4.3 – **Two applications of the symmetry surface estimation** : From left to right : i) a cortical surface, ii) the symmetry surface estimated on the cortex (for a better visualisation of the results, we display each hemisphere with a different colour), iii) colour representation of the absolute maximal curvature values along the antero-posterior axis (hot colours indicates high curvatures), iv) symmetry surface estimated on the Tautavel skull and v) we straighten the Tautavel skull according to the estimated symmetry surface [Ege86]

Chapitre 5

Conclusion

5.1 Contributions

In Chapter 3, we proposed a comprehensive overview of the methods for the estimation of symmetry planes in 3D point sets. Then, in Chapter 4, we designed two original methods allowing respectively the coarse estimation of all significant reflections in a 3D point set and the robust and accurate estimation of a symmetry plane given an initialisation. We showed that the resulting estimates of the symmetry plane are accurate even when the number of outliers is large. Then, we proposed to define the symmetry with respect to a smooth surface and proposed an algorithm to compute this surface of symmetry on 3D point sets.

5.2 Limits & Perspectives

5.2.1 Coarse estimation

Although the method we investigated (MeanShift3) for the coarse estimation of all the symmetry planes of a 3D point sets gives good results, this study is still quite preliminary and further evaluations are needed to investigate the properties of the different proposed algorithms. Moreover, we think that the use of methods to select the bandwidth parameters are needed to propose a more automated tool.

5.2.2 Accurate and robust estimation

In our framework, we indirectly define the approximate symmetry plane as the one minimising the average sum of squared distances between the points on one side of the object and their reflection about the plane, but only among the pairs of points for which these distances are lower than a predefined threshold $l\sigma$. In practice, the parameter σ decreases throughout the iterations until it reaches a small predefined value σ_f . We think that the estimation of the parameter σ_f on the basis of a trade-off between a maximal size for the symmetrical parts and a minimal value for the criterion will be an interesting improvement.

An important limitation of our approach is that it assumes that at least a small fraction of the data has not been contaminated significantly by asymmetries. As a result, this approach is likely to fail when dealing, for example, with symmetrical objects to which a global large scale deformation has been applied. For these objects, there does not remain any symmetrical part from the original symmetry. In this case, an appropriate approach would be to define, as proposed by Samir and colleagues, the optimal symmetry plane \tilde{P} as the one minimising the deformation between X and $S_{\tilde{P}}(X)$. We discuss such a possibility in Chapter 13.1.

5.2.3 Symmetry surface

Although in general, the algorithm we proposed for the estimation of a symmetry surface C provides good results, it is likely to provide surfaces C exhibiting local oscillations. To avoid such a behaviour, we can add a prior $P(C)$ on the ML problem to enforce C to be smooth. We performed preliminary studies designing C as a Fourier polynomial and $P(C) = \exp(-\sum_i (b/i!) \|C^i\|^2)$ (where C^i is the i th derivatives of C with respect to the spatial coordinates and where b is a parameter weighing the relative penalisation of the derivatives of C) [dFFG93]. Secondly, more appropriate solutions for the M-step have to be investigated. One could for example try to investigate solutions based on gradient-based minimisation of the original M-step.

Deuxième partie

Non linear registration of point sets

Chapitre 6

Introduction

Registration generally refers to the process of computing and applying a geometrical transformation to a first dataset to superpose it to a second data set, so as to make the homologous structures in both sets coincide. The need for automated registration methods is common to many fields such as computer vision, medical imaging, biometrics, quality control, *etc.* In practice, a registration method implicitly assumes the choice of i) a basis information to represent the structures to register (*e.g.* images, point sets, combination of both), ii) a model to explicit the nature of the expected transformations, movements and iii) a metric to specify what alignment means. The literature abounds with methods that are specific instances of this general view and many specificities such as *inverse consistent* registration (the deformation computed from a structure A to a structure B is the inverse of that computed from B to A) or *group-wise* registration (one wants to align together structures without specifying an a priori template that would bias the alignment) have emerged.

In this work, we assume that the data under study are represented as sets of 3D points representing the outline surfaces of the structures to register. Chapter 7 is a comprehensive review of non-linear registration methods involving such data. In this chapter, it appears that probabilistic methods based on the alignment of mixture models representing the point sets instead of the point sets themselves is a particularly fruitful approach. In particular, this view allows to consider the registration problem as a problem of estimation of unknown parameters linking two mixture models corrupted by noise and outliers. Interestingly, one can show that the EM-ICP algorithm [GP02] belongs to this wide class of approaches and leads to a simple and powerful formulation when dealing with rigid-body deformations. Moreover, as we will see in Part III, this view allows to introduce probabilities of correspondences between the points belonging to the two sets to register and leads to interesting properties to lead group-wise studies (*e.g.* template building). However, to the best of our knowledge, only little effort has been made to improve this framework and to adapt it to handle non-linear deformations efficiently. In particular, one observes that such a framework is inherently asymmetric, time consuming and of limited capture range. In Chapter 8, we tackle each of these flaws by modifying the original EM-ICP criterion and by designing subsequent efficient

algorithmic solutions to optimise it. This results in an algorithm allowing to register large data sets efficiently whereas most state-of-the-art methods are devised to deal with small data sets (a few hundreds). Finally, in Chapter 9 we evaluate our contributions and compare them with state-of-the-art methods.

Chapitre 7

Bibliography

7.1 Introduction

In this chapter, we present an overview of methods allowing the non-linear registration of 3D point sets structured as a mesh or not. More specifically, we focus on methods having no strong priors on the structures to register (*e.g.* topological constraints [YSV⁺10, SHMS07, ZHM07, Rob04], tessellation constraints [Rob04]) or on the expected deformations (*e.g.* articulated motion [MHK⁺08], isometric deformations [HAWG08]). Moreover, considering the huge number of methods for point set registration, we focus on the methods that we think are the most representative. We subdivide these methods into 4 main classes : methods based on mixture models representations of point sets, methods based on modal matrix representations of point sets, methods based on level set representations of point sets and methods based on Schwartz distribution representations of point sets. As we will see in the following, the methods based on mixtures modelling make up the major part of the algorithms proposed in the literature and will be studied with more details than others.

In Section 7.2, we study methods based on the modelling of point sets as mixture models. An important number of registration algorithms can be seen as specific instances of this wide class of methods and we aim to provide a comprehensive view where we explicit the sometimes hidden i) mixture models, ii) similarity measure, iii) expression of a transformation model T (generally relying on a regulariser $L(T)$) allowing to "deform" a mixture and iv) subsequent minimisation.

In Section 7.3.1, we describe registration methods based on the analysis of modal matrices of the point sets to register. In Section 7.3.2, we describe registration methods based on the registration of level set representations of the point sets instead of the surfaces themselves. In Section 7.3.3, we describe registration methods based on the matching of Schwartz distributions.

7.2 Mixture models and point set registration

Considering the point set registration as a problem of parameters estimation where the point sets are subject to noise and outliers, it seems natural to use probability density functions (pdf) to model each point set. More pragmatically, such a modelling allows to alleviate the complicated problems of defining correspondences in point sets (as one-to-one matching does not always exist) by indirectly introducing probabilities of matching leading to a more appropriate/pragmatic definition of the superimposition of two point sets. Moreover it allows to benefit from the numerous well-grounded estimators/minimisation schemes developed in the information theory literature.




FIG. 7.1 – **Point sets as Gaussian Mixture Models (GMMs)** : Mid-sagittal outlines of corpus callosa shapes represented as probability density functions, from [WVR06]

Let $X = \{x_1, \dots, x_N\}$ and $Y = \{y_1, \dots, y_M\}$ be two point sets. Let T be a transformation. Let f_T and g be the mixture models having the points of $T(X)$ and Y as centroids :

$$f_T(z) = 1/N \sum_i p_f(T(x_i) - z)$$

and

$$g(z) = 1/M \sum_j p_g(y_j - z)$$

p_f and p_g being two pdfs. In the following these pdf will be either considered as isotropic Gaussian (called $\psi(\cdot; \mu, \sigma^2 I)$) or Dirac (called $\delta(\cdot)$).

In this section, we consider the registration problem as a problem of minimisation of divergence between f_T and g with respect to T . Different divergences, modelling for p_f/p_g (Section 7.2.1) and for the transformation T have been investigated in the literature (Section 7.2.2) and are summed up in the following.

7.2.1 Data attachment term & Minimisation

7.2.1.1 KL-BASED METHODS : Kullback-Leibler divergence between a mixture of Gaussians and a mixture of Diracs : the ML approach

The Kullback-Leibler (KL) divergence between two pdf writes :

$$KL(g||f_T) = \int_{\mathbb{R}^3} g(z) \log \left(\frac{g(z)}{f_T(z)} \right) dz \quad (7.1)$$

The KL divergence can be interpreted as the “coding penalty” associated with selecting a distribution f_T to approximate a true distribution g [Cov06]. It is not symmetric and does not respect the triangle inequality (thus it is not a distance strictly speaking).

One can show that minimising the KL divergence between a mixture of N Gaussians f_T and a mixture of M Diracs $g(z) = \frac{1}{M} \sum_j \delta(y_j - z)$ is equivalent to solving the ML problem where the points of Y are considered as draws of a mixture of N Gaussians with centroids $T(X)$.

Derivation :

$$KL(g||f_T) = \int g(z) \log(g(z)) dz - \int g(z) \log(f_T(z)) dz, \quad (7.2)$$

the first term being a constant with respect to T , one can write (specifying the nature of f_T and g)

$$\arg \min_T KL(g||f_T) = - \arg \min_T \int g(z) \log(f_T(z)) dz, \quad (7.3)$$

$$= \arg \max_T \int \sum_j \delta(y_j - z) \log(f_T(z)) dz, \quad (7.4)$$

$$= \arg \max_T \sum_j \log(f_T(y_j)), \quad (7.5)$$

$$= \arg \max_T \sum_j \log \left(\sum_i \psi(y_j; T(x_i), \sigma^2 I) \right), \quad (7.6)$$

$$(7.7)$$

which corresponds to a maximum likelihood estimator of T (see Section 8.1 for an extended analysis of the ML formulation).

In this formulation, the mixtures f_T and g are not modelled the same way. This aspect has a strong interpretation in the clustering context (see Section 8.1 for more details). This asymmetric formulation can be undesirable in the context of atlas learning or unbiased registration.

Optimisation : The Expectation Maximisation (EM) algorithm [DLR77b] can be used to solve this ML problem (see Section 8.1 for more details). Contrary to gradient-based optimisation algorithms (such as the quasi-Newton algorithm), the EM algorithm does not need additional parameters specific to the optimisation process and achieves a monotonic convergence. However, this convergence is linear and the algorithm is likely to need many iterations to converge.

Robustification : The estimator obtained by the minimisation of the KL divergence is known not to be robust [BHHJ98, Sco99]. By extension the last ML estimator has the same flaw. To alleviate this problem, one can consider the M-estimation that has been designed as an extension of ML estimation (M-estimation stands for maximum likelihood like estimation). Indeed, the previous ML problem writes $\arg \max_T \sum_j \log(f_T(y_j))$ (Eq. 7.5) ; extending this approach the M-estimation proposes to consider the following problem : $\arg \max_T \sum_j \rho(y_j; T)$ where ρ is an arbitrary function. In essence, the M-estimation consists in proposing suitable functions ρ . Equivalently, with a probabilistic viewpoint, it consists in designing a new density function f_T .

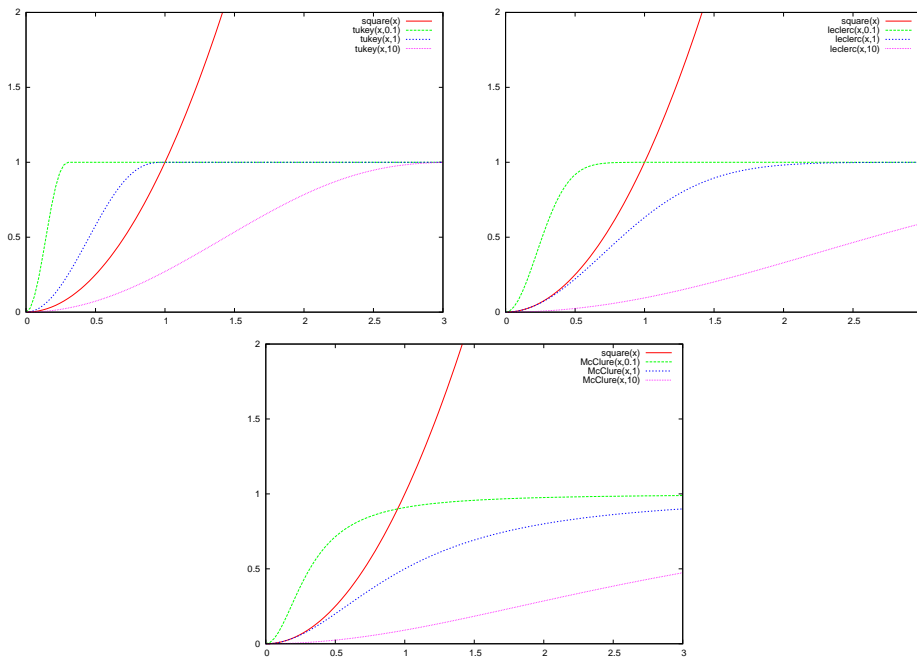


FIG. 7.2 – Some classical functions ρ used in M-estimation $Tukey(x, b) = \min(1 - (1 - x^2/b)^3, 1.0)$, $Leclerc(x, b) = 1 - \exp(-x^2/b)$ and $McClure(x, b) = x^2/(b + x^2)$.

For a certain class of functions ρ , the literature provides i) approaches to tune their robustness and ii) generic algorithmic solutions (the iteratively reweighted least squares [GR92]). The design of the function ρ generally implies the choice of i) a global shape generally indicating the way that $\rho(x)$ increases with x and ii) a scale parameter allowing to adapt ρ (that is generally not a linear function) to the parameters (y_j, T) . This scale

parameter must (named b on Figure 7.2) be carefully chosen as it can have a critical impact on the results.

Figure 7.2 references some commonly used functions ρ .

Applications : The effectiveness of the couple ML/EM has motivated the development of several specialisations of the framework. The method has been first derived to deal with rigid transformations [RCM⁺97] [Wel97], and then in a robust and efficient way [GP02]. More recently, different authors [CR00b, CP09, MS09, HFY⁺10] have extended the framework to non-linear transformations using different deformation models. The addition of priors on matchings has been investigated in [CP09] and some works have proposed methods to estimate (locally) the variance parameters of the Gaussian involved in f_T [HFY⁺10, MS09].

While these methods share a common viewpoint, their implementation choices lead to different characteristics mainly in terms of robustness and of computational and memory burden.

The Iterative Closest Point (ICP) algorithm : the Classification Maximum Likelihood (CML) approach : The well-known Iterative Closest Point algorithm [BM92] can be seen as an algorithm maximising a classification ML (CML) criterion (see Section 8.1 for more details). The iterative closest point algorithm has been intensively used to register point sets and the works using and extending the ICP are numerous. Particularly interesting works include that of Fitzgibbon [Fit03] (showing that the performance of the ICP is comparable to that of a Levenberg-Marquardt algorithm on the ICP criterion) and Pottman [PHYH06] (elucidating links between ICP and gradient descent minimisation). Moreover, the key extensions of the ICP can be summed up by the paper of Chui and colleagues [CR00b] (softassign, non linear transformation model and Sinkhorn normalisation) and that of Rusinkiewicz and colleagues [RL07] (robustification, weighting and subsampling strategies).

7.2.1.2 KL-BASED METHODS : Jensen-Shannon divergence between GMMs

As previously mentioned, the KL divergence is not symmetric. This asymmetry is likely to introduce an undesirable bias in the registration process. To cope with this limitation, Wang and colleagues [WVRE08] have proposed to use the Jensen-Shannon (JS) divergence (also termed total divergence to the average) that appears as a symmetrised KL divergence. It writes :

$$JS(f_T||g) = \frac{1}{2}KL(f_T||m_T) + \frac{1}{2}KL(g||m_T), \quad (7.8)$$

where $m_T = \frac{1}{2}(f_T + g)$. It can equivalently be written as :

$$JS(f_T||g) = H(f_T + g) - (H(f_T) + H(g)). \quad (7.9)$$

where H is the Shannon entropy. It shares many of the properties of the KL divergence (in particular, its subsequent estimator is not robust). However, it is theoretically more relevant than the KL divergence as its square root defines a proper distance.

Derivation : In the case where one considers two GMMs, no closed-form solution allows to easily express neither $H(f_T)$ or $JS(f_T||g)$ and one needs to approximate it by generating a finite number Q of samples s_q^f ($q \in [1, \dots, Q]$) from f_T . Then using a Monte Carlo approximation (see [Rob06], p. 315), one can write :

$$H(f_T) = - \int_{\mathbb{R}^3} f_T(z) \log(f_T(z)) dz \approx - \frac{1}{Q} \sum_{q=1}^Q \log(f_T(s_q^f)) \quad (7.10)$$

$$= - \frac{1}{Q} \sum_{q=1}^Q \log \left(\frac{1}{N} \sum_i \psi(s_q^f; T(x_i), \sigma^2 I) \right). \quad (7.11)$$

In the same way :

$$H(g) = - \frac{1}{Q} \sum_{q=1}^Q \log \left(\frac{1}{M} \sum_j \psi(s_q^g; y_j, \sigma^2 I) \right). \quad (7.12)$$

Then, one can simply write the JS divergence between f_T and g as :

$$JS(f_T||g) = \frac{1}{Q} \sum_{q=1}^Q \log \left(\frac{1}{N} \sum_i \psi(s_q^f; T(x_i), \sigma^2 I) \right) + \frac{1}{Q} \sum_{q=1}^Q \log \left(\frac{1}{M} \sum_j \psi(s_q^g; y_j, \sigma^2 I) \right) \quad (7.13)$$

$$- \frac{1}{Q} \sum_{q=1}^Q \log \left(\frac{1}{N} \sum_i \psi(s_q^f; T(x_i), \sigma^2 I) + \frac{1}{M} \sum_j \psi(s_q^g; y_j, \sigma^2 I) \right) \quad (7.14)$$

Optimisation : The derivatives of the cost function with respect to T can be analytically computed and the gradient descent algorithm (redrawing the s_q^f s and s_q^g throughout the iterations) can be used to minimise $JS(f_T||g)$.

Applications : The Jensen-Shannon divergence can be extended to compare an arbitrary number of distributions and has been proposed to properly define group-wise registration [WVR06] :

$$JS(f_1 || \dots || f_n) = \sum_i H(f_i) - H\left(\sum_i f_i\right). \quad (7.15)$$

Extension : the CDF JS divergence Wang *et al.* [WVR06] proposed to investigate solutions based on the cumulative distribution function of the spatial point distribution instead of their pdf. More precisely, they study the cumulative residual entropy (that has been developed first for grey level image registration). Let F_i be the cumulative distribution associated to the pdf f_i , the cumulative residual entropy (CRE) is defined as :

$$CRE(f_i) = - \int_{\mathbb{R}^{3+}} (1 - F_i(\lambda)) \log(1 - F_i(\lambda)) d\lambda$$

Relying on integral form (cumulative distribution function) unlike pdf (a derivative form), it defines a regular measure that is quite immune to noise and outliers. Then, following the JS divergence (Eq. 7.15), the CDF JS divergence is defined as :

$$JS_{cdf}(f_1 || \dots || f_n) = CRE(F_{\sum_i f_i}) - \sum_i CRE(F_i),$$

where $F_{\sum_i f_i}$ is the cumulative distribution function associated to the pdf $\frac{1}{n} \sum_i f_i$. Then designing f_i as a pdf depending from the parameter T_i , the function JS_{cdf} and its derivatives with respect to each transformation T_i can be estimated using a Parzen window estimate of $CRE(X)$ and the overall criterion can be minimised with respect to T_i using the conjugate gradient method.

Note that, Wang and colleagues recently extended their work to another divergence based on the Havrda-Charvat entropy [WMH10]. This work shares many properties with the CDF JS divergence.

7.2.1.3 L_2 -BASED METHODS : L_2 distance between 2 GMMs

The natural inability of the KL based estimator to cope with outliers have stimulated the use of divergence leading to more robust estimators for T . In particular, Jian and colleagues proposed to use the L_2 distance between two pdf [BHHJ98] :

$$L_2E(f||g) = \int_{\mathbb{R}^3} (f_T(z) - g(z))^2 dz. \quad (7.16)$$

The estimator obtained the minimisation of this distance is known to be robust and it can be shown that it belongs to the class of M-estimators [BHHJ98]. On the other hand it is asymptotically less efficient than KL divergence *i.e.* when the number of samples tends to infinity the variance of its estimator is larger than the one of KL. Its robustness properties are convenient as the scale and the shape parameter of the estimator are implicitly (as opposed to classical M-estimators). On the other hand, it implies that one has no control on its behaviour (as opposed to classical M-estimators).

Derivation :

$$\begin{aligned} \arg \min_T L_2E(f_T||g) &= \arg \min_T \int f_T(z)^2 dz - 2 \int g(z) f_T(z) dz + \int g(z)^2 dz, \\ &= \arg \min_T \int f_T(z)^2 dz - 2 \int g(z) f_T(z) dz. \end{aligned} \quad (7.17)$$

In case of two GMMs, it gives :

$$\begin{aligned}
\arg \min_T L_2 E(f_T || g) &= \arg \min_T \int \left[\sum_i \psi(z; T(x_i), \sigma^2) \right]^2 dz \\
&\quad - 2 \int \left(\sum_j \psi(z; y_j, \sigma^2) \right) \left(\sum_i \psi(z; T(x_i), \sigma^2) \right) dz \\
&= \arg \min_T \int \sum_i \psi(z; T(x_i), \sigma^2)^2 dz \\
&\quad + 2 \int \sum_{i_1 \neq i_2} \psi(z; T(x_{i_1}), \sigma^2) \psi(z; T(x_{i_2}), \sigma^2) dz \\
&\quad - 2 \int \left(\sum_j \psi(z; y_j, \sigma^2) \right) \left(\sum_i \psi(z; T(x_i), \sigma^2) \right) dz \quad (7.18)
\end{aligned}$$

Using the fact that $\int \psi(z; \mu_1, \sigma_1^2) \psi(z; \mu_2, \sigma_2^2) dz = \psi(0; \mu_1 - \mu_2, \sigma_1^2 + \sigma_2^2)$, one can finally write :

$$\arg \min_T L_2 E(f_T || g) = 2 \sum_{i_1 \neq i_2} \psi(0; T(x_{i_1}) - T(x_{i_2}), 2\sigma^2) \quad (7.19)$$

$$- 2 \sum_j \sum_i \psi(0; y_j - T(x_i), 2\sigma^2). \quad (7.20)$$

Optimisation : The derivatives of the cost function with respect to T can be explicitly computed and a quasi-Newton method can be used to minimise $L_2 E(f_T || g)$.

Links between KL and L_2 : The density power divergence has been defined by Basu and colleagues [BHHJ98] as :

$$d_\alpha(f, g) = \int \left[\frac{1}{\alpha} f^{1+\alpha}(x) - \left(1 + \frac{1}{\alpha} f(x) g^\alpha(x) \right) + g^{1+\alpha}(x) \right] dx, \quad (7.21)$$

where $\alpha \geq 0$ can be seen as a parameter characterising the trade-off between robustness and asymptotic efficiency [BHHJ98]. L_2 corresponds to $\alpha = 1$ and KL is considered as the asymptotic case where $\alpha = 0$. These two cases are widely used as they are the only ones to exhibit closed-form expressions when dealing with mixture of Diracs/Gaussians.

Applications : To the best of our knowledge, Jian and colleagues [JV05] were the first to use the L_2 distance between GMMs in the context of rigid and non-rigid (thin plate spline local warps) point sets registration. Their results illustrate the natural robustness of $L_2 E$.

More recently work of Roy and colleagues [RGR07] proposed to use the L_2 distance in a similar context. The main difference with the work of Jian and colleagues is that the

registration is not performed directly on the point sets but on intermediate Gaussian mixtures that are first fitted to the data X and Y (using an EM algorithm). Moreover, the authors proposed to estimate the variance of the noise σ^2 characterising the mismatch between the GMMs.

Generalized L_2 distance : Recently, the L_2 distance between two pdf has been extended to define a measure between an arbitrary number of pdf [WVS09] :

$$L_2E(f_1||\dots||f_n) = \sum_{i=1}^n L_2E(f_i, \frac{1}{n} \sum_j f_j).$$

When considering GMM, this distance still leads to a closed-form expression.

7.2.1.4 L_2 -BASED METHODS : Correlation between GMMs :

The correlation between two GMMs is defined as

$$C(f_T, g) = \int_{\mathbb{R}^3} f_T(z)g(z)dz. \quad (7.22)$$

Derivation 1 : One can show that when dealing with isometric transformations, it is equivalent to maximise $C(f_T, g)$ and to minimise the $L_2E(f_T, g)$ with respect to T . Indeed, as previously mentioned :

$$\arg \min_T L_2E(f_T||g) = \arg \min_T \int f_T(z)^2 dz - 2 \int g(z)f_T(z)dz. \quad (7.23)$$

if one assumes that T is an isometric transformation then the first term vanishes

$$\begin{aligned} \arg \min_T L_2E(f_T||g) &= \arg \min_T - \int g(z)f_T(z)dz \\ &= \arg \max_T C(f_T, g). \end{aligned}$$

As a result, one can extrapolate the properties (efficiency, robustness, minimisation strategies) of L_2 to correlation C . This result is only true when dealing with isometric transformations. To our knowledge this point sets kernel correlation framework has never been used to perform non-rigid registration.

Without explicitly formulating the problem in terms of mixture models, Tsin and Kanade [TK04] were the first to propose to align two point sets by maximising the (kernel) correlation between them.

Derivation 2 : More recently, Sandhu and colleagues [SDT09] proposed to assume that the modes of each of the densities composing the mixtures are far apart from each other. Under this assumption, one can consider only a component of a mixture x_i for each y_j when computing the correlation [SDT09] :

$$\tilde{T} = \arg \max_T \sum_j \max_i \int \psi(z; y_j)\psi(z; T(x_i))dz, \quad (7.24)$$

this interestingly leads to (with the assumption that T is an isometric transformation) :

$$\tilde{T} = \arg \max_T \sum_j \max_i \exp(-\frac{1}{4} \|y_j - T(x_i)\|^2). \quad (7.25)$$

This last criterion can interestingly be viewed as a robust version of the criterion minimised by the ICP algorithm (in fact it is exactly an ICP criterion where the classical square function is replaced by a Leclerc function [Hub81]). The main contribution of this work consists in showing how to use particle filtering methods to optimise this criterion. Their algorithm leads to particularly efficient solutions in case of partial matchings.

7.2.2 Regularisation

In the previous section, we focused on the divergences/distances that can be used as data attachment term in the registration context without specifying T . In practice, T can be either a parametric transformation (such as a rigid-body or a polynomial transformation) possibly penalised or can define a dense field $T(x) = Ax + t(x)$ where A is a global linear deformation and t is regularised by an operator L :

$$\min_T d(T(X), Y) + \alpha L(T) \quad (7.26)$$

where $\alpha > 0$ weighs the influence of the regularisation L over the data-attachment d . In the following, we present transformation T of this form.

Moreover, we focus on the solution of the following problem (known as approximation problem) :

$$\min_T \sum_i w_i \mathcal{C}(\|c_i - T(v_i)\|) + \alpha L(t) \quad (7.27)$$

where $w_i \in \mathbb{R}^+$, c_i and $v_i \in \mathbb{R}^3$ and $\mathcal{C} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a cost function. One can see this problem as a general form of registration where the correspondences (c_i, v_i) are known.

7.2.2.1 The thin plate splines

The thin plate spline (TPS) deformations model the physical bending of a thin sheet of metal. This model has been first introduced in the registration community by Bookstein [Boo89] and has been used intensively to interpolate/approximate deformations from manually labeled point sets.

Noting $T(x) = (T_u(x), T_v(x), T_w(x))$, this regulariser writes :

$$L(T) = L(T_u) + L(T_v) + L(T_w) \quad (7.28)$$

where

$$L(T_u) = \int_{\mathbb{R}^3} \frac{\partial^2 T_u}{\partial u^2}(x) + \frac{\partial^2 T_u}{\partial v^2}(x) + \frac{\partial^2 T_u}{\partial w^2}(x) + 2 \frac{\partial^2 T_u}{\partial u \partial v}(x) + 2 \frac{\partial^2 T_u}{\partial v \partial w}(x) + 2 \frac{\partial^2 T_u}{\partial u \partial w}(x) dx \quad (7.29)$$

and similarly for $L(T_v)$ and $L(T_w)$.

The advantage of this model is clearly to offer physical justification. Notice that Cachier and Ayache have studied vectorial TPS deformations *i.e.* the components f_u , f_v and f_w are not penalized independently but depend from each other [CA04]. An illustration of the interest of such a deformation is given in Figure 7.3.




FIG. 7.3 – **Scalar versus vectorial TPS deformations** One of four points initially forming a square is forced to move downwards. With scalar TPS (left), vertical lines remain vertical and straight. With vectorial TPS (right), the displacement field involves both vertical and horizontal components. From [CA04]

Optimisation The solution of the corresponding approximation problem has the form :

$$T(x) = A.x + t + \sum_{i=0, \dots, N} w_i \cdot k_{TPS}(\|x - x_i\|) \quad (7.30)$$

where :

- A is a 3×3 matrix representing the affine component of T .
- t is a $3D$ translation vector.
- (w_i) with $w_i \in \mathbb{R}^3$ characterises the local warps of T .
- $k_{TPS}(x) = x^3$ is the radial basis function characterising TPS.

The value of $L(T)$ for given parameters (w_i) is called the bending energy and is given by $\sum_{i,j} w_i^T w_j k_{TPS}(\|x_i - x_j\|)$.

When \mathcal{C} is classical square function ($\mathcal{C} : x \mapsto x^2$), the approximation problem (Eq. 7.27) has a closed-form solution consisting in a QR decomposition and the solving of linear systems.

Applications : For the last 10 years, the TPS model has been used extensively in the non-linear registration literature (*e.g.* [CR00a, TC07, WMH10, WVS09, RGR07, JV05]).

7.2.2.2 The coherent point drift (CPD)

This model comes from the motion coherence theory (MCT) developed by Yuille and Grzywacz [YG89]. The regulariser corresponding to the CPD writes :

$$L(T) = L(T_u) + L(T_v) + L(T_w), \quad (7.31)$$

where

$$L(T_u) = \sum_{n=0}^{\infty} \frac{(\beta)^{2n}}{n!2^n} (D^n(T_u))^2, \quad (7.32)$$

where $D^{2n}T_u = \Delta^{2n}T_u$ and $D^{2n+1}T_u = \nabla\Delta^{2n}T_u$ (where Δ is the Laplacian operator and ∇ the gradient operator). $L(T_v)$ and $L(T_w)$ have similar expressions. $\beta > 0$ is a parameter allowing to weigh the increasing of penalisation $\frac{(\beta)^{2n}}{n!2^n}$ when the order of the derivatives D^n increases.

Optimisation The solution of the corresponding approximation problem (Eq. 7.27) has the form :

$$T(x) = \sum_{i=0, \dots, N} w_i k_{CPD}(\|x - x_i\|) \quad (7.33)$$

where $k_{CPD}(x) = \exp(-x^2/(2\beta))$. The value of $L(T)$ for given parameters (w_i) is given by $\sum_{i,j} w_i^T w_j k_{CPD}(\|x_i - x_j\|)$. When \mathcal{C} is a classical square function, the approximation problem (Eq. 7.27) is a linear system consisting in the solving of a linear system.

The advantages of this model are i) to introduce the parameter β allowing to tune the regularisation and thus to have a better control on the solution, ii) to offer theoretical justifications [YG89].

Relationship between TPS and CPD : It appears that both TPS and CPD regularisers lead to analogous expressions for T . In fact, this kind of approaches can be generalised using the Reproducing Kernel Hilbert Space (RKHS) theory. More results about this generalisation will be given in Section 8.5.

Applications : To the best of our knowledge, there are only a few works using operators coming from the motion coherence theory in context of the point set registration [MS09].

7.2.2.3 Diffeomorphisms

A diffeomorphism T is defined as a C^∞ function, for which T^{-1} exists and is C^∞ . Such a property is of great interest. In particular, it ensures that the deformation will be one-to-one, and that overlaps are not possible.

There are several ways to generate diffeomorphic transformations. A particularly powerful one consists in considering a time dependent transformation $T(x, t)$ that is generated as the solution of a flow transport equation involving a time dependent velocity $v(x, t)$ that is subject to a smoothness cost. The final transformation T is given by the integration of the velocity over the time t :

$$T(x) = T(x, 1) = x + \int_0^1 v(T(x, t), t) dt \quad (7.34)$$

The transformations described in the previous sections do not allow to ensure that T is a diffeomorphic transformation. One of the main practical differences between the classical formalism and the diffeomorphic formalism is that the “classical” transformation generate deformations parametrised as $T(x) = x + t(x)$ (where t is a 3D vector). This can be seen as a limit case of Eq. 7.34 where the integrated trajectory $\int_0^1 v(T(x, t), t) dt$ is a straight line (*i.e.* a simple 3D vector). When deformations are small this assumption can be considered as correct (this is the so-called “small deformation paradigm”). When the deformations become larger, it is more realistic to consider the deformation as a time dependent velocity describing 3D curves (this is the so-called “large deformation paradigm”).

Optimisation Thus, one considers T as a time dependent function $T(x, t)$ (with $t \in [0, 1]$) that is generated as the solution of the transport equation [JM00, CY01] :

$$\frac{\partial T(x, t)}{\partial t} = v(T(x, t), t) \quad (7.35)$$

with $T(x, 0) = x$ and $v(\cdot, t)$ subject to a smoothness cost $\|Lv(\cdot, t)\|^2$ (where L is a regulariser). This is an intricate problem linking v/T by Equations 7.34 and 7.35, we note T as T^v to recall (when necessary) this dependency. Following this notation, the approximation problem 7.26 takes the form :

$$(T, v) = \arg \min_{T, v} d(T^v(X), Y) + \int_0^1 Lv(\cdot, t) dt. \quad (7.36)$$

When \mathcal{C} is a classical square function, the solution exhibits a general form for v that can be introduced into the last equation that then depends only on $T(\cdot, t)$ evaluated at points $x_i \in X$ (and on its first order derivatives). Then, in order to be able to minimise it, one discretises the time interval $[0, 1]$ as $\cup_k [t_k, t_{k+1}]$ and for each interval $[t_k, t_{k+1}]$, the velocity is assumed to be constant (thus the derivatives of $T(\cdot, t)$ can be computed very easily) and the approximation problem 7.36 becomes :

$$T(x_i, t_k)_{\forall i, k} = \arg \min_{T(x_i, t_k)_{\forall i, k}} \left[\frac{1}{t_k - t_{k-1}} \sum_k \sum_i (T^v(v_i, t_k) - T^v(v_i, t_{k-1}))^T \right]. \quad (7.37)$$

$$\left(\int_{t_{k-1}}^{t_k} (K_{T, t})^{-1} dt \right) (T^v(v_i, t_k) - T^v(v_i, t_{k-1})) + \sum_i [c_i - T^v(v_i, 1)]^2. \quad (7.38)$$

where $K_{T, t}$ is a $3 \text{ card}(X) \times 3 \text{ card}(X)$ matrix defined by a kernel that characterises the operator L . The derivatives with respect to $T(\cdot, t_k)$ of the underlying criterion can be explicitly computed and the optimal $\hat{T}(\cdot, t_k)$ can be estimated using a gradient descent algorithm on each point x_i and at each time step t_k .

Applications

To the best of our knowledge, Guo and colleagues [GR09] are the only authors to propose to use such transformations in the GMM based point set registration context. Indeed, although the model exhibits desirable properties, it is very time/memory consuming and thus likely to be limited to the registration of structures that can be represented by only a small number of points.

7.2.3 Some additions

7.2.3.1 Landmarks

Imposing the transformation T to superpose pairs of landmarks (that have been manually identified) in addition to the classical data attachment term allows to introduce *a priori* knowledge into the automatic registration process. As a result, it provides a way to constrain the transformation and to lead to a realistic/desirable solution. In practice, this just consists in adding the following term to the criterion :

$$d_L(T) = w_L \sum_{(l_k^X, l_k^Y) \in L} \|l_k^Y - T(l_k^X)\|^2, \quad (7.39)$$

where $w_L > 0$ weighs the influence of the landmark term over the divergence term and $L = ((l_k^X, l_k^Y))$ is the set of matched landmarks. In essence, this term is simple and close to classical form for $d(T(X, Y))$ (with the main difference that the correspondences are known) and thus classical optimisation schemes (such as EM or gradient descent) can easily be extended to support this new term [ARV07, CP09].

7.2.3.2 Adding higher order information

Most methods just consider superimposition of point locations to build a data attachment term. When the point sets are tessellated, higher order information such as normals and curvatures can be derived and used in the data attachment term. The main interest of such an adding is to increase the capture range of the algorithm and to provide more realistic solutions.

Such approaches assume that one can compute how the normals evolve when the surface is deformed. Two main methods related with the GMM framework (in an ICP-like approach) have been proposed to deal with this problem.

Feldmar and Ayache [FA96] showed how to include the normals and curvatures by considering how these quantities are modified in the special case of locally affine deformations. They devised a criterion combining point coordinates, normals and curvatures and locally affine regularisation and showed how to minimise it in an iterative manner similarly to the original ICP algorithm. However, solving for the optimal transformation given correspondences in this framework is a nonlinear problem which makes the overall scheme computationally expensive and difficult to implement in case of large surfaces.

More recently Munch et al. [MCP10] integrated the work of Kambamethu et al. into a ICP-like process. Kambamethu and colleagues proposed a simple formula showing

how the unit normal changes when a surface undergoes a small deformation [KGH93, KGHL03, LK03]. This formula is quite specific as it involves a differentiable orthogonal parametrization of the surfaces and of the deformation field. However, under these strong conditions, it results in the addition of a linear terms that can be efficiently integrated in classical minimisation scheme. The main drawback of this solution is the need to compute explicitly, at each step of the minimisation scheme, a differentiable orthogonal parametrization for both X and T which can be very time consuming.

7.2.4 A fundamental minimisation tool for GMMs : the fast kernel computation

7.2.4.1 Fast Gauss transform

The discrete Gauss Transform (GT) is defined as :

$$G(y_j) = \sum_{x_i \in X} w_i \exp(-\|y_j - x_i\|^2 / 2\sigma^2), \quad (7.40)$$

where $w_i \in \mathbb{R}^+$, $\sigma^2 \in \mathbb{R}^+$ and y_j and $x_i \in \mathbb{R}^3$

The computation of cost functions involving Gaussian mixture models or of their gradient generally leads to a large number of evaluations of the GT at different point y_j (generally derived from a matrix-vector product). Its direct computation at $\text{card}(Y)$ points y_j needs $\mathcal{O}(\text{card}(X) \times \text{card}(Y))$ operations.

The fast Gauss transform [GS91] allows to approximate the $\text{card}(Y)$ computations in $\mathcal{O}(\text{card}(X) + \text{card}(Y))$ operations (with a constant factor depending on the desired precision). This method relies on the replacing of the complete sums on $x_i \in X$ with smaller sums, computed from series expansion of the exponential function.

7.2.4.2 Kd-tree subdivision of the space

When using a truncated Gaussian function the discrete Gauss Transform becomes :

$$G(y_j) = \sum_{x_i \in X} w_i \rho_\delta(\|y_j - x_i\|^2) \exp(-\|y_j - x_i\|^2 / 2\sigma^2), \quad (7.41)$$

where $\rho_\delta : x \mapsto 1$ if $x < \delta$ and 0 else.

When δ is small each $G(y_j)$ can be computed very efficiently using a *kd*-tree [Ben75]. On the other hand, when δ becomes very large, *kd*-tree performs very poorly whereas the FGT performs very efficiently. The main advantage of the *kd*-tree is its ability to support non linear metrics efficiently. As an example, suppose that we design a function c between points of Y and X such that $c(x_i, y_j) = 0$ if the normals of x_i and y_j are close to each other and $c(x_i, y_j) = \infty$ else :

$$G(y_j) = \sum_i w_i \rho_\delta(\|y_j - x_i\|^2 + c(x_i, y_j)) \exp(-(\|y_j - x_i\|^2 + c(x_i, y_j)) / 2\sigma^2) \quad (7.42)$$

In this condition, $G(y_j)$ can be computed efficiently by limiting the evaluation of $c(.,.)$ to a small subset of X :

compute $\mathcal{S} = \{x_i \in X \text{ such that } \|x_i - y_j\|^2 < \delta\}$ (using a kd-tree)
 $G(y_j) = 0$
for all $x_i \in \mathcal{S}$
 if $(\|y_j - x_i\|^2 + c(x_i, y_j) < \delta)$
 then $G(y_j) = G(y_j) + \exp(-(\|y_j - x_i\|^2 + c(x_i, y_j))/(2\sigma^2))$

On the contrary, the FGT does not provide any solution for this problem without evaluating $c(x_i, y_j)$ for all $x_i \in X$.

7.3 Other approaches :

7.3.1 Spectral based correspondences (and extensions)

The spectral method initially proposed by Shapiro and Brady [SB92] consists in building a modal matrix for each point set. The modal matrix for a point set $X = (x_i)$ is computed by i) building a $\text{card}(X) \times \text{card}(X)$ symmetric proximity matrix $G_{i,j} = \exp(-\|x_i - x_j\|/(2\sigma^2))$, ii) performing a SVD decomposition of $G = VDVT^T$ where D contains the (positive) eigenvalues of G in a decreasing order. The matrix V is the modal matrix. Note that G (thus V) is invariant under rigid-body transformation. Each row of V represents one of the original points of X whereas each column entry measures how the original points are distributed among the different eigen-modes of G . Once, one has computed the modal matrices V_X and V_Y of X and Y , the strategy consists in considering these measures as a shape descriptor almost invariant with respect to the expected transformation. As a result, the point-to-point matching is done by comparing rows of both modal matrices. This can be done either by a best one rule [SB92] or by building correspondence probabilities [CH00] (when both point sets have not the same number of points the larger modal matrix is truncated).

This technique can be embedded in an iterative scheme involving the optimisation of a transformation (such as an affine [CH00] or a TPS [TC07] transformation) with known correspondences and the spectral correspondences between points (with updated positions of the point set).

Note that other strategies sharing a common view with the one we present in this section have been proposed in the literature (*e.g.* [LH05, SLH91]).

Advantages and drawbacks : This kind of strategies has the advantage to be quite simple, fast and efficient in case of small structures subject to transformation relatively closed to an affine transformation.

7.3.2 Surface as a level set function

In this subsection, X and Y are considered to be structured as surfaces (or as meshes). Instead of registering directly the two surfaces of interest X and Y , Albrecht and colleagues [LAV07, DLAV07] proposed to represent them as the zero level of a

signed distance 3D function I :

$$I(x) = \begin{cases} \text{dist}(x, X) & \text{if } x \in \textit{outside}(X) \\ 0 & \text{if } x \in X \\ -\text{dist}(x, X) & \text{if } x \in \textit{inside}(X) \end{cases} \quad (7.43)$$

These 3D functions are represented on 3D grids and their registration consists in registering both 3D grids together. In other words, the original problem can be treated as a 3D image registration problem.

In essence, any registration algorithm can be used in this context. Albrecht and colleagues propose to use the Thirion's demons algorithm and modifies its cost function in order to consider the mean curvature of I as a supplemental source of information leading the registration process.

Advantages and drawbacks : This method is quite convenient as it enjoys the numerous algorithmic solutions proposed in the 3D image registration literature. The drawback of such a strategy is clearly to need the conversion of the surfaces to 3D volumes. Indeed, for identical data resolution, the memory usage (and possibly the computational requirement) needed for registration much larger for 3D images than for meshes.

7.3.3 Diffeomorphic matching of distributions

7.3.3.1 Matching of Dirac distributions

The framework of diffeomorphic matching of distributions has been recently developed by Glaunés and colleagues in a series of papers [GTY04, VG05].

Regularisation Following the approach of Section 7.2.2.3, one generates diffeomorphisms as the solution of the flow equation :

$$\frac{\partial T(x, t)}{\partial t} = v(T(x, t), t), \quad (7.44)$$

with $T(x, 0) = x$ and $t \in [0, 1]$.

One builds a Hilbert space \mathcal{V} containing the elements $v(., t)$ and that is equipped with a norm $|\cdot|_{\mathcal{V}}$. One then defines a group $G_{\mathcal{V}} = \{T^v(., 1) \text{ such that } |v(x, \cdot)|_{\mathcal{V}} \in L_1\}$ and one equips it with a right-invariant geodesic distance

$d(T_1, T_2) = \inf(\int_0^1 \|v(., t)\|_{\mathcal{V}} dt \text{ such that } T^v(., 1) \circ T_1(., 1) = T_2(., 1)$. The group $G_{\mathcal{V}}$ contains only diffeomorphisms and will be the space in which one searches our transformation T . $d(T_1, I_d)$ is used as a regulariser (I_d being the identity transformation).

Data attachment term The authors show that if one considers the case of a weighted sum of Diracs (which is a discrete Schwartz distribution) localised at the points of X : $\nu = \sum_i a_i \delta(x_i)$ then $T^v \nu = \sum_i a_i \delta(T^v(x_i))$.

In order to compare two such distributions, the authors notice that all differences of distributions are contained in the dual of an Hilbert space containing continuous bounded functions on \mathbb{R}^3 . One notes \mathcal{I} such a space and \mathcal{I}^* its dual. Thus one can consider $\|v - \mu\|_{\mathcal{I}^*}$ as a data attachment term. (where $\|\nu\|_{\mathcal{I}^*} = \sup(\int f d\nu, f \in \mathcal{I}, |f|_{\mathcal{I}} = 1)$ defines a norm for \mathcal{I}^*). Under few assumptions on \mathcal{I} and $|\cdot|_{\mathcal{I}}$ and by noting $T^v \mu = \sum_i a_i \delta_{T^v(x_i)}$, $\nu = \sum_j b_j \delta_{y_j}$ and $T^v \mu - \nu = \sum_i c_i \delta_{z_i^v}$, the authors show that one can write $\|T^v \mu - \nu\|_{\mathcal{I}^*} = \sum_{i,j} c_i c_j k_{\mathcal{I}}(z_i^v, z_j^v)$ (z_i^v depends on v/T) where $k_{\mathcal{I}}$ is a simple kernel function (characterising the space \mathcal{I}).

Global functional This leads to the following problem :

$$\tilde{v} = \arg \min_v d(T^v, I) + 1/\sigma^2 \|T^v \mu - \nu\|_{\mathcal{I}^*} \quad (7.45)$$

$$= \arg \min_v \int_0^1 \int_x \|v(x, t)\|_{\mathcal{V}}^2 dt + 1/\sigma^2 \|T^v \mu - \nu\|_{\mathcal{I}^*} \quad (7.46)$$

Moreover, one can show that the optimal v has the following form (when trajectories $x_i(t) = T^v(x_i, t)$ are fixed) : $\tilde{v}(x, t) = \sum k_{\mathcal{V}}(x_i(t), x) \alpha_i(t)$. where $k_{\mathcal{V}}$ is a kernel characterising the space \mathcal{V} . In other words, the problem 7.46 writes :

$$\tilde{v} = \arg \min_{\alpha_i(t), x_i(t)} \int_0^1 \sum_{i,j} \alpha_j(t)^T k_{\mathcal{V}}(x_i(t), x_j(t)) \alpha_i(t) dt + 1/\sigma^2 \sum_{i,j} c_i c_j k_{\mathcal{I}}(z_i^v, z_j^v). \quad (7.47)$$

The derivatives of the underlying cost function with respect to $\alpha_i(t)$ can then be explicitly computed and the functional can be minimised with respect to $\alpha_i(t)$ and $x_i(t)$ using a gradient descent algorithm.

7.3.3.2 Currents

The previous was dedicated to the modeling of set of points. More recent works have shown how to embed more distributions of more elaborated structures called m-vectors : 1-vectors are segments and 2-vectors are 3D frames. In particular, a 3D mesh can be seen as a distribution of 2-vectors (also called 2d-currents) and Vaillant and Glaunes [VG05] show how to consider current instead of classical discrete Schwartz distributions in the previous formalism. In essence, the problem has the same form than previously but μ and ν now represent distributions of 2-vectors of \mathbb{R}^3 . Then, the problem consists in modeling efficiently μ, ν and the action of T^v *i.e.* characterising the space \mathcal{I}^* .

Advantages and drawbacks : These well grounded methods are very interesting as they allow to model the complete surface representation *i.e.* points and normals and guarantee the transformation to be diffeomorphic. On the other hand, their computational burden can limit their applications. However, recent works of Durrleman and colleagues [DPTA08] proposed very efficient minimisation schemes. Another limitation, concerns the fact the space of currents contains shapes but also other elements that are

Other approaches :

85

not shape. This can be quite problematic when for example trying to compute the mean shape of a population.

Chapitre 8

Methods

In this chapter, we propose an efficient and robust algorithm for the non-linear registration of point sets.

First, in Section 8.1 we present the EM-ICP framework by showing that a simple probabilistic modelling of the registration problem allows to solve it by using standard clustering techniques. In particular, the unknown non-linear transformation best superposing the two point sets can then be estimated according to the maximum *a posteriori* principle (MAP) and using (typically) the EM algorithm. This optimisation then boils down to a simple iterative estimation of fuzzy point-to-point correspondences (E-step) (encoded in what is often termed the *match matrix*) and of the non-linear transformation (M-step) in turn.

Then we progressively improve this initial algorithm :

- In Section 8.2, we propose to tackle problems due to the asymmetric nature of the MAP principle. For this purpose, we notice that the derived EM algorithm can be seen as an iterative alternated minimisation (over the match matrix and the unknown transformation) of an (energetic) criterion and we propose to modify it. In particular, we explicitly add a second match matrix that is column stochastic in addition to the first match matrix that is row stochastic. This modification only changes the E-step and improves the matching process.
- In Section 8.3, we show how to specify priors on the match matrices with only minor changes to the optimisation algorithm. These priors based on local and global shape descriptors are likely to significantly improve the capture range and the convergence speed of the algorithm.
- In Section 8.4, we propose to symmetrise the estimation of transformation by jointly computing the forward and the backward deformation fields linking the two point sets and enforcing them to be compatible with each other.
- In Section 8.5, we devise two alternative efficient solutions to regularise the deformation field. The first one is based on the regularisation of simple local models (translation or affine) attributed to each point of the set and regularised over the space. The second one stands on the Reproducing Kernel Hilbert Space (RKHS) theory and on the Fourier analysis and consists in building efficient regularisers

leading to closed-form solutions based on sparse linear algebra. Both of these two regularisers are much less computational and memory demanding than the classical Thin Plate Spline regulariser.

8.1 Non linear registration as a clustering problem

8.1.1 General formulation

The problem of interest in this chapter is to find the transformation T best superposing two free-form surfaces X and Y (represented by point clouds or meshes). A convenient probabilistic viewpoint on this classical problem is to consider the surface Y as a noised version of $T(X)$. Hypothesizing an isotropic Gaussian noise on data X and Y , a simple way to formulate this viewpoint is to assume that each sample y_j has been drawn independently from any one of $\text{card}(X)$ possible 3-variate normal distributions with means $T(x_k)$ and covariance matrices $\sigma^2 I$ (with σ unknown).

This way, the registration problem becomes a *clustering* problem, whose challenge is i) to find the *label* of each point y_j , *i.e.* the one out of $\text{card}(X)$ possible distributions from which y_j has been drawn, and ii) to estimate the parameters of these $\text{card}(X)$ distributions. The connection between registration and clustering becomes clear when one realises that i) actually amounts to match each point y_j in Y with a point x_k in X , while ii) simply consists in computing T given these matches. This viewpoint is extremely fruitful, as it allows one to refer to classical clustering techniques and especially the maximum likelihood principle to solve the registration problem. Three different paradigms have been especially followed in this context [Mar75]. Let us introduce some notations first :

$$\begin{aligned} \forall k \in 1 \dots \text{card}(X), \psi_k(\cdot; T) &= \mathcal{N}(T(x_k), \sigma^2 I) \\ \forall j \in 1 \dots \text{card}(Y), \forall k \in 1 \dots \text{card}(X), z_{jk} &= 1 \text{ iff } y_j \text{ comes from } \psi_k(\cdot; T) \end{aligned}$$

In the Classification Maximum Likelihood (CML) approach, T is considered as a fixed unknown parameter and one tries to find the indicator variables z_{jk} and the transformation T so as to maximise the criterion CL [SS71] :

$$CL = \prod_{y_j \in Y} \prod_{x_k \in X} [\psi_k(y_j; T)]^{z_{jk}} \quad (8.1)$$

The problem is typically solved by the Classification EM (CEM) algorithm [Gil92], which can be shown to find an at least local maximum of the criterion CL and proceeds as follows, in an iterative way, starting from an initial value \tilde{T} :

$$\begin{aligned} \mathbf{EC\text{-}step} : \forall j, \tilde{z}_{jk} &= 1 \text{ iff } k \text{ maximises } \psi_k(y_j; \tilde{T}) \\ \mathbf{M\text{-}step} : \tilde{T} &= \arg \min_T \sum_{j,k} \tilde{z}_{jk} \|y_j - T(x_k)\|^2 \end{aligned}$$

In other words, the Expectation-Classification (EC) step consists in matching each point y_j of Y with the closest point in $\tilde{T}(X)$, while the Maximisation (M) step consists in computing the transformation best superposing these pairs of matched points. When

T describes a rigid-body transformation, this algorithm is nothing else than the popular Iterated Closest Point algorithm [BM92].

In the Maximisation Likelihood (ML) approach, the indicator values z_{jk} are no longer considered as unknown quantities to estimate, but rather as hidden/unobservable variables of the problem. This is actually a drastic and fundamental change of viewpoint, as the focus is no longer on assigning each y_j to one of the distributions ψ_k but rather on estimating the parameters of the Gaussian mixture made of these distributions. If we involve priors π_{jk} on the indicator variables ($\forall j, k, 0 < \pi_{jk} < 1$, and $\forall j, \sum_k \pi_{jk} = 1$), the likelihood then simply writes [Day69] :

$$L = \prod_{y_j \in Y} \sum_{x_k \in X} \pi_{jk} \psi_k(y_j; T) \quad (8.2)$$

In essence, the prior π_{jk} conveys the probability that the point y_j comes from the distribution ψ_k without knowing anything else. The criterion L can be maximised by using the popular EM algorithm, which converges to an at least local maximum of the likelihood [DLR77b]. If we consider the priors π_{jk} as known beforehand and if we introduce the notation A_{jk} as the posterior probability of the hidden indicator variable z_{jk} to be equal to 1, the EM algorithm writes (see appendix B for a derivation of the algorithm) :

$$\begin{array}{l} \text{E-step : } \tilde{A}_{jk} = \frac{\pi_{jk} \exp[-\|y_j - \tilde{T}(x_k)\|^2 / (2\sigma^2)]}{\sum_i \pi_{ji} \exp[-\|y_j - \tilde{T}(x_i)\|^2 / (2\sigma^2)]} \\ \text{M-step : } \tilde{T} = \arg \min_T \sum_{j,k} \tilde{A}_{jk} \|y_j - T(x_k)\|^2 \end{array}$$

The parameter σ , which is not estimated in this framework, acts as a scale parameter. The problem of estimating σ (or even a different σ_k for each of the model ψ_k) is not investigated in this work. It can be given an initial value and decreased throughout the iterations for improved performances. When T describes a rigid-body transformation and priors π_{jk} are chosen to be uniform, this algorithm is nothing else than the EM-ICP algorithm [GP02].

In the Maximum A Posteriori (MAP) approach, instead of simply considering T as a fixed unknown parameter of the pdf, one can consider it as a random variable on which priors (acting as regularisers on T) can be easily specified. Then, the ML estimation can be easily turned into a maximum *a posteriori* (MAP) problem with only slight modifications to the optimisation scheme, as shown by Green [Gre90]. Different choices can be proposed for T and associated priors in this context can be proposed. If $p(T)$ is a prior of the form $p(T) \propto \exp(-\alpha L(T))$ then the optimal transformation can be found using the MAP principle (also termed penalised ML) :

$$L = \prod_{y_j \in Y} \sum_{x_k \in X} \pi_{jk} \psi_k(y_j; T) p(T) \quad (8.3)$$

The EM algorithm then writes :

$$\begin{aligned} \mathbf{E}\text{-step} : \tilde{A}_{jk} &= \frac{\pi_{jk} \exp[-\|y_j - \tilde{T}(x_k)\|^2 / (2\sigma^2)]}{\sum_i \pi_{ji} \exp[-\|y_j - \tilde{T}(x_i)\|^2 / (2\sigma^2)]} \\ \mathbf{M}\text{-step} : \tilde{T} &= \arg \min_T \sum_{j,k} \tilde{A}_{jk} \|y_j - T(x_k)\|^2 + \alpha L(T) \end{aligned}$$

Note that one can consider a similar adaptation for the CML approach.

Interpretation & Extensions :

Intuitively, the ML approach is a *fuzzy* version of the CML. It appears clearly from the iterative formulas of both algorithms that the classification likelihood is an “all-or-nothing” version of the likelihood, leading to a “bumpier” and harder-to-maximise criterion, something that is well known by those who are familiar with the ICP algorithm. Note that the ML formulation followed by the EM algorithm leads to the same iterative formulas that would have resulted from the addition of a barrier function on the indicator variables in the ICP criterion [CR00a].

8.1.2 Robustifying the estimator and reducing the computational burden

In practice, this algorithm suffers from outliers *i.e.* points of X having no homologous in Y . To alleviate this problem, one can consider $\psi_k(\cdot; T)$ as a truncated Gaussian probability density function (pdf) with cut-off distance $\delta > 0$. This modification has beneficial effects on the properties of the algorithm. Indeed, it allows : i) to reduce drastically the computational burden of the E-step (by the use of a *kd-tree* [Ben75]), ii) to reduce the impact of points of Y having no correct homologous in X in the estimation of the optimal T *i.e.* to improve the robustness of the criterion iii) to increase the convergence speed of the overall algorithm. The following algorithm can be shown to converge to an at least local maximum of the new (truncated) criterion :

Algo Reg1: Robust EM-ICP

E-step : initialise $A = (A_{jk})$ as the null matrix

$\forall x_k \in X$;

$S = \{y_j \in Y \text{ such that } \|y_j - \tilde{T}(x_k)\|^2 < \delta\}$ (using a *kd-tree*)

$\forall y_j \in S$; $A_{jk} = \exp(-(\|y_j - \tilde{T}(x_k)\|^2 / (2\sigma^2)))$

$\forall y_j \notin S$; A_{jk} is left equal to 0.

$\forall y_j \in Y$;

if $\sum_i A_{ji} \neq 0$

$\tilde{p}_j = 1$

$\forall x_k \in X$, $\tilde{A}_{jk} = A_{jk} / \sum_i A_{ji}$ (normalisation)

else

$\tilde{p}_j = 0$

M-step : $\tilde{T} = \arg \min_T \sum_{j,k} \tilde{p}_j \tilde{A}_{jk} \|y_j - T(x_k)\|^2 + \alpha L(T)$

This introduces a set (p_j) of binary variables such that $\forall j$, p_j is null if point y_j is considered as an outlier (thus has no correspondent) and is equal to one else.

8.1.3 Energetic interpretation

It can be shown that the former algorithm is an iterative alternated minimisation (over A and T) of the following criterion :

$$\mathcal{E}1(X, T(Y), A) = \sum_{j,k} A_{jk} \rho_{\delta}(\|y_j - T(x_k)\|^2) + 2\sigma^2 \sum_{j,k} A_{jk} \log(A_{jk}) + 2\sigma^2 \alpha L(T) \quad (8.4)$$

with $\forall j, \sum_k A_{jk} = 1$ and where $\rho_{\delta} : r \mapsto r$ if $r < \delta$ and δ else (δ corresponds to the cut-off distance of the truncated pdf of Section 8.1.2). To simplify the notation, we introduce $\alpha' = 2\sigma^2 \alpha$. Note that the ρ function could be advantageously replaced by a smoother function (leading to a more classical M-estimation scheme). However, our function has the advantage to remain interpretable in terms of pdf (a truncated Gaussian function) and thus allows to conserve the probabilistic interpretation of the different parameters (and particularly of the match matrix A).

It can be given an energetic interpretation of this criterion whose 3 terms respectively represent :

- a data-attachment term (where the classical quadratic cost function is replaced by a more robust cost function),
- a barrier function allowing to control the fuzziness of A (the higher σ^2 , the greater the fuzziness); in practice, this term convexifies the criterion,
- a regularisation term.

8.2 Symmetrising the matching process

A particularly undesirable property of the last formulation is the asymmetric constraint $\sum_k A_{jk} = 1, \forall j$ (*i.e.* A is row stochastic). In practice, on the basis of the maximum a posteriori probability principle, for a given match matrix A , the correspondent in X of a point y_j is given by x_c where $c = \arg \max_c A_{jc}$. This leads to many-to-one matchings between points of X and Y . In particular, there is no direct constraint to enforce a point of X to have a correspondent in Y . This makes the algorithm unable to achieve good matchings in some specific configurations. It is particularly enlightening to consider the case when surfaces are far from each other (example on Fig 8.1).




FIG. 8.1 – **Effect of the asymmetric normalisation constraint on A .** The "correct" matchings are $(x_i, y_i)_{i=1\dots 5}$. One considers T as the identity. From left to right : i) two point sets X and Y , ii) distances involved in the computation of A_{11} when registering Y on X . iii) distances involved in the computation of A_{11} when registering X on Y . As an example, if one considers $\sigma = 1$, then in case ii) $A_{11} = 0.01$ and points x_1 and y_1 have only little chance to be matched at the end of the overall process whereas in case iii) $A_{11} = 0.45$.

To alleviate this problem, Rangarajan and colleagues [RCM⁺97] have proposed to impose the matrix A to be doubly stochastic (*i.e.* $\forall k, \sum_j A_{jk} = 1$ and $\forall j, \sum_k A_{jk} = 1$) instead of simply stochastic. With this new constraint on A , the E-step has no longer a simple solution. As a result, they approximate the optimal solution for A by performing a Sinkhorn normalisation [Sin64] on the original (*i.e.* non-normalised) matrix A at the end of the E-step. However, this method is not applicable to matrices having null entries and thus cannot be applied when using a truncated Gaussian pdf. In practice, this limits its application to small data sets.

As an alternative, we propose to modify the energetic criterion $\mathcal{E}1$ minimised by the EM algorithm (Eq. 8.4) and to introduce a new match matrix B in this criterion

that we impose to be column stochastic :

$$\begin{aligned} \mathcal{E}2(X, T(Y), A, B) &= \sum_{j,k} A_{jk} \rho_\delta(\|y_j - T(x_k)\|^2) + 2\sigma^2 \sum_{j,k} A_{jk} \log(A_{jk}) \quad (8.5) \\ &+ \sum_{j,k} B_{jk} \rho_\delta(\|y_j - T(x_k)\|^2) + 2\sigma^2 \sum_{j,k} B_{jk} \log(B_{jk}) \\ &+ \alpha' L(T), \end{aligned}$$

with $\forall j, \sum_k A_{jk} = 1$ and $\forall k, \sum_j B_{jk} = 1$.

Notice that if we consider the matrix $C = (\frac{A_{jk} + B_{jk}}{\text{card}(X) + \text{card}(Y)})$, one can state that $\forall j, k$ $C_{j,k} \geq 0$ and $\sum_{j,k} C_{j,k} = 1$.

8.2.1 Minimisation with respect to T

While the minimisation of $\mathcal{E}2(X, T(Y), A, B)$ with respect to A and B when T is fixed (E-Step) is straightforward and of low complexity (using the improvement proposed in Section 8.1.2), its minimisation with respect to T (M-Step) generally consists in solving a numerical system of size proportional to $\text{card}(Y) \times \text{card}(X)$. When dealing with large point sets X and Y , this can have severe effects on the computational and memory requirements. In this section, we propose to reduce the complexity of this minimisation from $\text{card}(Y) \times \text{card}(X)$ to $\text{card}(X)$.

For this purpose, we now consider that T is represented as the initial position plus a displacement field : $T(x_k) = x_k + t(x_k)$ and L is a regulariser on t . First, the M-Step consists of :

$$\text{M-step : } \tilde{t} = \arg \min_t \sum_{j,k} (\tilde{p}_j \tilde{A}_{jk} + \tilde{q}_k \tilde{B}_{jk}) \|y_j - x_k - t(x_k)\|^2 + \alpha' L(t)$$

where, analogously to (p_j) , (q_k) are the binary variables associated to the matrix B .

The derivative of $\mathcal{E}_M = \sum_{j,k} (\tilde{p}_j \tilde{A}_{jk} + \tilde{q}_k \tilde{B}_{jk}) \|y_j - x_k - t(x_k)\|^2 + \alpha' L(t)$ with respect to $t(x_k)$ is :

$$\begin{aligned} \frac{\partial \mathcal{E}_M}{\partial t(x_k)} &= -2 \left(\sum_j \tilde{p}_j \tilde{A}_{jk} y_j + \sum_j \tilde{q}_k \tilde{B}_{jk} y_j \right) \\ &+ 2 \left(\left(\sum_j \tilde{p}_j \tilde{A}_{jk} + \sum_j \tilde{q}_k \tilde{B}_{jk} \right) (x_k + t(x_k)) \right) + \alpha' \frac{\partial L(t(x_k))}{\partial t(x_k)} \end{aligned}$$

Calling $\tilde{B}_{.k} = \tilde{q}_k \sum_j \tilde{B}_{jk}$ and $\tilde{A}_{.k} = \sum_j \tilde{p}_j \tilde{A}_{jk}$, this gives :

$$\frac{\partial \mathcal{E}_M}{\partial t(x_k)} = -2 \sum_j (\tilde{p}_j \tilde{A}_{jk} + \tilde{q}_k \tilde{B}_{jk}) y_j + 2(\tilde{A}_{.k} + \tilde{B}_{.k})(x_k + t(x_k)) + \alpha' \frac{\partial L(t(x_k))}{\partial t(x_k)}$$

which is the derivative with respect to $t(x_k)$ of :

$$\sum_k (\tilde{A}_{.k} + \tilde{B}_{.k}) \left\| \frac{\sum_j (\tilde{p}_j \tilde{A}_{jk} + \tilde{q}_k \tilde{B}_{jk}) y_j}{\tilde{A}_{.k} + \tilde{B}_{.k}} - x_k - t(x_k) \right\|^2 + \alpha' L(t) \quad (8.6)$$

Thus, it is equivalent to minimise criteria 10.1 and 8.6 with respect to t . However, the first problem consists of solving a system of size $\mathcal{O}(\text{card}(X) \times \text{card}(Y))$ whereas the second consists of solving a system of size $\mathcal{O}(\text{card}(X))$. Then the overall algorithm to minimise Eq.10.1 can be expressed as :

Algo Reg2: Symmetric robust EM-ICP

E-step :

compute $\tilde{A}; \tilde{B}; (\tilde{p}_j); (\tilde{q}_k)$

for all k ,

compute $\tilde{A}_{.k} + \tilde{B}_{.k}$,

compute $\tilde{y}_k = \sum_j (\tilde{p}_j \tilde{A}_{jk} + \tilde{q}_k \tilde{B}_{jk}) y_j / (\tilde{A}_{.k} + \tilde{B}_{.k})$

M-step : solve the approximation problem :

$\arg \min_T \sum_k (\tilde{A}_{.k} + \tilde{B}_{.k}) \|\tilde{y}_k - x_k - t(x_k)\|^2 + \alpha' L(t)$

8.3 Adding priors

8.3.1 General idea

The computation of matrices A and B is essentially based on the spatial proximity between the points $T(x_k)$ and y_j . This is unsatisfactory, for two reasons. First, this distance is highly conditioned by the previous estimation of T , which in turn depends on the previous estimation of $A_{jk} + B_{jk}$ and so on. This chicken-and-egg problem limits the capture range of the algorithm, which is likely to converge to a bad solution if no good initial T is given. Second, in many applications it is difficult to design a physical model T capturing the expected deformation between two structures. Thus, the global maximiser of the ML criterion is likely not to be realistic.

Some efforts have been made to include richer information in the matching process in addition to the classical spatial proximity between points, *e.g.* based on the similarity of the normals at points x_k and y_j . Such approaches assume that one can compute how the normals evolve when the surface is deformed. This generally results in adding nonlinear terms into the initial criterion and leads to an intractable M-Step.

On the other hand, a more generic and simple method consists in specifying a prior probability π_{jk} between the points x_k and y_j to be matched, that we suppose to be independent of 1) the spatial proximity between the points of the two surfaces and 2) the unknown transformation T . By specifying relevant priors π_{jk} , we introduce additional information on matches independent of the transformation that thus allows to compute reliable posteriors even for a bad initial estimate of T .

8.3.2 Designing π

In practice, we choose to design $\pi = (\pi_{jk})$ such that $\pi_{jk} \propto \exp(-\beta c(y_j, x_k))$ where $c : X \times Y \rightarrow \mathbb{R}^+$ conveys the cost of matching points y_j and x_k , independently of T . The parameter $\beta > 0$ weighs the influence of π_{jk} over $\|y_j - T(x_k)\|$ during the E-step.

The equivalent energetic criterion is :

$$\mathcal{E}\mathcal{Z}(X, T(Y), A, B) = \sum_{j,k} (A_{jk} + B_{jk}) \rho_\delta(\|y_j - T(x_k)\|^2 + \beta c(y_j, x_k)) \quad (8.7)$$

$$+ 2\sigma^2 \sum_{j,k} A_{jk} \log(A_{jk}) + 2\sigma^2 \sum_{j,k} B_{jk} \log(B_{jk}) + \alpha' L(T) \quad (8.8)$$

Depending on the information to encode (continuous value or label), we propose two approaches to build c .

8.3.2.1 Designing π using labels

The cost function c can be computed via the comparison between labels on points (cortical sulci/gyri). We define : $c(y_j, x_k) = 0$ if points j and k have compatible labels and $c(y_j, x_k) = \textit{penalty} > 0$ else. In particular, this view allows to use pair of landmarks

in the registration process ($c(y_j, x_k) = 0$ if points j and k correspond to the same landmark and $c(y_j, x_k) = \textit{penalty} > 0$ else). One can also extract the crest lines from both meshes [GWM01] as they constitute salient features. Each point is given a label depending on whether it belongs to a crest line or not. Then, we define $c_{crest}(y_j, x_k) = 0$ if y_j and x_k have the same label and $c(y_j, x_k) = \textit{penalty}$ else.

8.3.2.2 Designing π using descriptors

c can be computed via the comparison between continuous values (or vectors) $d(x)$ describing the surface around the considered points. To account for potential inaccuracies on $d(\cdot)$, we define the measure as : $c^d(y_j, x_k) = 0$ if $\|d(y_j) - d(x_k)\| < \tau$ and $c^d(y_j, x_k) = \textit{penalty} > 0$ else.

Then we choose $d(\cdot)$ among the local/global shape descriptions designed in the literature. In our case, we are interested into descriptors :

- invariant to a certain class of transformations
- robust to noise
- robust to small distortions

Among them, we choose to use :

- the shape index $d(x) = sh(x)$ [KvD92] that describes the local shape irrespective of the scale and that is invariant to similarities. To achieve robustness to noise and small distortions, we compute it by fitting a quadratic surface in the neighborhood of the considered point. The fitting is performed by i) approximating a unit normal at point x , ii) defining a local coordinate system (where x is at $(0; 0; 0)$ and the unit normal lies along the positive z-axis) and iii) fitting a quadratic surface of the type $f(u; v) = au^2 + buv + cv^2$ in the least-squares sense using the neighbors of x . The shape index can then be expressed as a function of a , b and c : $sh(x) = -2/\pi \arctan\left(\frac{2(a(x)+c(x))+2b(x)^2}{2|a(x)-c(x)|}\right)$.
- the curvedness $d(x) = cu(x)$ [KvD92] that specifies the amount of curvature and that is invariant to rigid-body transformations. We compute it using the same techniques that we used for shape index : $cu(x) = \sqrt{\frac{(a(x)+b(x)+|a(x)-b(x)|+c(x))^2+(a(x)+b(x)-|a(x)-b(x)|+c(x))^2}{2}}$.
- the (normalised) total geodesic distance $d(x) = tgd(x)$ [AFK07] that is invariant to isometries in the shape space (including non-elastic deformations). This distance is defined as $tgd(x) = \sum_j \frac{d_G(x, x_j)}{\max_j \sum_k d_G(x_j, x_k)}$, where $d_G(x_j, x_k)$ is the geodesic distance between x_j and x_k . It is computed efficiently using a graph representation of the (tessellated) point set and the Dijkstra's algorithm.

8.3.3 Implementation

In practice, adding non uniform priors π only changes the way the matrices A and B are computed (E-step). We propose the following efficient algorithm for this purpose :




FIG. 8.2 – Mapping of descriptor values : From left to right : curvedness, shape index and total geodesic distance on two lateral ventricles. Homologous anatomical landmarks yield qualitatively the same descriptor values.

Algo Reg3: Symmetric robust EM-ICP with priors

E-step :

initialise A and B to the null matrix

$\forall x_k \in X$;

$\mathcal{S} = \{y_j \in Y \text{ such that } \|y_j - \tilde{T}(x_k)\|^2 < \delta\}$ (using a kd -tree)

$\forall y_j \in \mathcal{S}$;

if $\|y_j - \tilde{T}(x_k)\|^2 / (2\sigma^2) + \beta c(y_j, x_k) \leq \delta$

$A_{jk} = \exp(-(\|y_j - \tilde{T}(x_k)\|^2 / (2\sigma^2) + \beta c(y_j, x_k)))$

$\forall y_j \notin \mathcal{S}$; A_{jk} is left equal to 0.

$B = A$

normalise A in rows and B in columns and compute (\tilde{p}_j) and (\tilde{q}_k) (see Section 8.1)

for all x_k , compute $\tilde{A}_{.k} + \tilde{B}_{.k}$ and \tilde{y}_k (see Section 8.2.1)

compute $\tilde{y}_k = \sum_j (\tilde{p}_j \tilde{A}_{jk} + \tilde{q}_k \tilde{B}_{jk}) y_j / (\tilde{A}_{.k} + \tilde{B}_{.k})$

M-step : solve the approximation problem :

$\arg \min_t \sum_k (\tilde{A}_{.k} + \tilde{B}_{.k}) \|\tilde{y}_k - x_k - t(x_k)\|^2 + \alpha' L(t)$

8.3.4 Initialising T using π only

By considering $\beta \rightarrow \infty$, we assume that the spatial proximity term $\|y_j - T(x_k)\|^2$ is not reliable. In this asymptotic case, the algorithm allowing to minimise the criterion $\mathcal{E3}$ can be written as :

E-Step :

compute the matrix $A + B$ without considering the proximity term $\|y_j - T(x_k)\|^2$

M-Step :

compute $\tilde{T} = \arg \min \sum_{j,k} (A_{jk} + B_{jk}) \rho_\delta(\|y_j - T(x_k)\|^2) + \alpha' L(T)$

Note that this algorithm is not iterative as the E-Step does not depend on T . Moreover, it does not depend on an initial value for T or A and is invariant to the relative position of X and Y and thus, assuming that no correct initialisation is provided for T , this algorithm is a good candidate to estimate an initial transformation T . Typically, we

want to estimate an initial rigid-body transformation, similarity or affine transformation between X and Y . In practice, if reliable priors were provided, this last algorithm would be able to compute a good estimate for such a transformation model. However, the function c is not discriminative enough to be able to compute reliable matrices A and B . As a result during the M-step, the estimation is “contaminated” by a very large number of outliers. Whereas the ρ function provides sufficient robustness properties when the number of outliers is small with respect to the number of inliers, it is not able to provide correct results in this context.

To tackle this lack of robustness, we restate our problem minimisation of residuals into a problem of mode seeking in the residuals space that has been demonstrated to tolerate more than 50% of outliers and that exhibits very efficient minimisation strategies [WMH10] :

$$\mathcal{E}_{Init}(X, T(Y)) = - \sum_{j,k} U_\delta(c(x_k, y_j)) U_\epsilon(\|y_j - T(x_k)\|^2) + \alpha' L(T) \quad (8.9)$$

where $U_\delta : x \mapsto 1$ if $x < \delta$ and 0 else is a uniform kernel and $\epsilon > 0$. The value of this criterion can be interpreted as the number of pairs $(y_j, T(x_k))$ among the pairs satisfying $c(x_i, y_j) < \delta$ such that $\|y_j - T(x_k)\|^2 \leq \epsilon$.

One can minimise the criterion \mathcal{E}_{Init} using the following algorithm :

Algo Reg4: Initialisation T with priors

- Step i)** compute the set $S = \{(x_k, y_j) \in X \times Y \text{ such that } c(x_k, y_j) < \delta\}$
Step ii) compute $\tilde{T} = \arg \min n(T, S) + \alpha' L(T)$
-

where $n(T, S)$ is the number of matchings $(x_k, y_j) \in S$ within an error tolerance (*i.e.* $\|y_j - T(x_k)\|^2 < \epsilon$).

Roughly speaking Step ii) is nothing else but a RANSAC procedure [FB81]. In other words, we use all matchings having a prior lesser than δ (Step i)) and we perform a RANSAC procedure to estimate T on this set of matchings (Step ii)). Moreover, to avoid degenerate solutions and to improve our estimator, we add the supplementary constraint that the matchings are one-to-one *i.e.* each point x_k and each point y_j can be matched at most one time when computing $n(T, S)$. Note that other kernels than the uniform kernel U can be used [WMH10] and are able to achieve better robustness. In practice the uniform kernel is sufficient for our applications (and allows a simpler and more efficient minimisation than other kernels).




FIG. 8.3 – **Example of initialisation on two ventricles.** Left : Initialisation with center of mass and principal axes superimposition. Right : Affine transformation estimation with our method

8.4 Symmetrising the estimation of T : A framework for a consistent & symmetric EM-ICP

Whereas we tried to reduce the inherent asymmetry of the EM-ICP framework in Section 8.2, it seems difficult to obtain *inverse consistent* registration (*i.e.* the registration of one set to the other provides the inverse transformation of that obtained when switching the two sets) in this EM-ICP framework. Such a property is very desirable, especially when building an atlas from a set of anatomical structures. First trials focused on the estimation of the transformation : Joshi & Miller [JM00] showed how to build a diffeomorphism between two sets of matched points, but without guaranteeing inverse consistency. In parallel, Johnson & Christensen proposed a method towards this goal, but using thin-plate splines where invertibility is not ensured [JC02]. These solutions are theoretically attractive but computationally redhibitory in case of large point sets, which limits their use to simple anatomical structures.

In this section, we propose a symmetric formulation of the registration problem extending the criteria developed in previous section, that allows to jointly compute the forward and the backward deformation fields linking the two point sets.

8.4.1 Cost function

Let T^X and T^Y be respectively the backward and forward unknown transformations superposing X and Y . Let A^X (resp. A^Y) be the match matrix describing the correspondences between $T^X(X)$ and Y (resp. $T^Y(Y)$ and X). Following the EM-ICP framework in its energetic formulation, we consider the matrices A^X and A^Y as hidden variables of the problem and we design our cost function as :

$$\begin{aligned} \mathcal{E}_4(T^X, T^Y, A^X, A^Y) = & \mathcal{E}_1(Y, T^X(X), A^X) + \mathcal{E}_1(X, T^Y(Y), A^Y) \\ & + \gamma \mathcal{E}_c(T^X \circ T^Y, I) + \gamma \mathcal{E}_c(T^Y \circ T^X, I) \end{aligned} \quad (8.10)$$

where :

- \mathcal{E}_1 is the criterion corresponding to the EM-ICP described in the previous sections. It can be equivalently replaced by criteria \mathcal{E}_2 or \mathcal{E}_3 .
- $\mathcal{E}_c(T^Y \circ T^X, I)$ is a consistency term that measures the discrepancy between transformations T^X and T^Y . Without this term, estimations of T^X and T^Y would be completely independent. This new term couples them and forces them to be compatible with each other. We design it as :

$$\mathcal{E}_c(T^Y \circ T^X, I) = \sum_{x_j \in X} \|T^Y \circ T^X(x_j) - x_j\|^2.$$

It is easy to observe that $\mathcal{E}(T^X \circ T^Y, I)$ is null for $T^X = (T^Y)^{-1}$. However, for any minimum of criterion (8.10), there is no insurance that T^X and T^Y are invertible

and that $\mathcal{E}(T^X \circ T^Y, I)$ is null. Similarly, we design

$$\mathcal{E}_c(T^X \circ T^Y, I) \sum_{y_j \in Y} \|T^X \circ T^Y(y_j) - y_j\|^2.$$

– $\gamma > 0$ is a parameter weighting the influence of \mathcal{E}_c .

8.4.2 Minimisation

The criterion $\mathcal{E}4$ can be optimised by an iterative two-step algorithm that consists in minimisation over $A^X; A^Y$ and $T^X; T^Y$ in turn :

Algo Reg5: Symmetric consistent EM-ICP

Init \tilde{T}^X and \tilde{T}^Y as the identity function

E-step : $\tilde{A}^X; \tilde{A}^Y = \arg \min_{A^X, A^Y} \mathcal{E}4(\tilde{T}^X, \tilde{T}^Y, A^X, A^Y)$

M-step $\tilde{T}^X; \tilde{T}^Y = \arg \min_{T^X, T^Y} \mathcal{E}4(T^X, T^Y, \tilde{A}^X, \tilde{A}^Y)$

The E-step can be solved by computing independantly A^X and A^Y similarly to what is done in Section 8.1.2. For the M-step, we propose to optimise iteratively the criterion with respect to each one of the two unknowns T^X and T^Y :

M-step.init estimate the optimal \tilde{T}^X and \tilde{T}^Y dropping the consistency term \mathcal{E}_c
M-step.a) $T^X = \tilde{T}^X$ and $T^Y = \tilde{T}^Y$
M-step.b) $\tilde{T}^X = \arg \min_{T^X} \mathcal{E}4(T^X, T^Y, \tilde{A}^X, \tilde{A}^Y)$
M-step.c) $\tilde{T}^Y = \arg \min_{T^Y} \mathcal{E}4(T^X, T^Y, \tilde{A}^X, \tilde{A}^Y)$

If one notices that $\mathcal{E}1(Y, T^X(X), A^X)$ is composed of a data attachment term $\mathcal{E}_d(Y, T^X(X), A^X)$ and a regularisation term $\mathcal{E}_r(T^X)$, this algorithm intuitively consists in alternatively estimating T^X as a compromise between data attachment ($\mathcal{E}_d(Y, T^X(X), A^X)$), regularisation ($\mathcal{E}_r(T^X)$) and consistency with T^Y ($\mathcal{E}_c(T^Y \circ T^X, I)$) and T^Y as a compromise between the three other symmetric terms. By specifying the form of the transformation T^X as the initial position plus a displacement field : $T^X(x_j) = x_j + t^X(x_j)$ and R is a regulariser on t^X (similarly for T^Y and t^Y). Then M-step can be written as :

init estimate the optimal \tilde{T}^X and \tilde{T}^Y dropping the consistency term \mathcal{E}_c
M-step.a) $T^X = \tilde{T}^X$ and $T^Y = \tilde{T}^Y$
M-step.b) $\tilde{t}^X = \arg \min_{t^X} \sum_{i,j} A_{i,j}^X \|y_i - x_j - t^X(x_j)\|^2 + \alpha' L(t^X) + \gamma \sum_j \|\tilde{t}^Y(x_j + t^X(x_j)) + t^X(x_j)\|^2 + \gamma \sum_i \|t^X(y_i + \tilde{t}^Y(y_i)) + \tilde{t}^Y(y_i)\|^2$
M-step.c) $\tilde{t}^Y = \arg \min_{t^Y} \sum_{i,j} A_{i,j}^Y \|x_i - y_j - t^Y(y_j)\|^2 + \alpha' L(t^Y) + \gamma \sum_j \|\tilde{t}^X(y_j + t^Y(y_j)) + t^Y(y_j)\|^2 + \gamma \sum_i \|t^Y(x_i + \tilde{t}^X(x_i)) + \tilde{t}^X(x_i)\|^2$

Figure 8.4 illustrates the principle of the chosen compatibility term with the previous parametrisation.

Due to the terms $\sum_j \|\tilde{t}^Y(x_j + t^X(x_j)) + t^X(x_j)\|^2$ in M-Step.a) and $\sum_i \|\tilde{t}^X(y_i + t^Y(y_i)) + t^Y(y_i)\|^2$ in M-Step.b), which are somewhat redundant with their symmetric




FIG. 8.4 – **Illustration of the interest of the compatibility term** : From left to right i) Point x is deformed towards Y by the deformation $t(x)^X$; ii) Point y is deformed towards X by the deformation $\tilde{t}^Y(y)$; iii) the compatibility term imposes that $t^X(y + \tilde{t}^Y(y))$ should be equal to $-\tilde{t}^Y(y)$; iv) The resulting deformation field \tilde{t}^X tends to be a compromise (in the sense of the regularisation) between data attachment and compatibility.

counterparts, the two problems are very intricate and thus difficult to solve. Thus, similarly to what is done by Chui *et al.* [CRZL04] in a related context, we drop them, which allows to reformulate M-step.b) and M-step.c) as two independent approximation problems :

$$\begin{aligned} \text{M-step.b)} \quad \tilde{t}^X &= \arg \min_{t^X} \sum_{i,j} A_{i,j}^X \|y_i - x_j - t^X(x_j)\|^2 \\ &\quad + \alpha' L(t^X) + \gamma \sum_i \|t^X(y_i + \tilde{t}^Y(y_i)) + \tilde{t}^Y(y_i)\|^2 \\ \text{M-step.c)} \quad \tilde{t}^Y &= \arg \min_{t^Y} \sum_{i,j} A_{i,j}^Y \|x_i - y_j - t^Y(y_j)\|^2 \\ &\quad + \alpha' L(t^Y) + \gamma \sum_j \|t^Y(x_j + \tilde{t}^X(x_j)) + \tilde{t}^X(x_j)\|^2 \end{aligned}$$

In practice, only a few iterations are necessary to significantly decrease the criterion.

8.5 Efficient approximation techniques

Solving the M-step is highly conditioned by the choice of the regulariser L . Several choices have been proposed in the literature and the most commonly used is probably the Thin Plate Spline (TPS) regulariser. It has the main advantages to exhibit a closed-form solution for the approximation problem and to be justified by a physical interpretation. However, this choice also implies computational and memory usage requirements that limit its application to data of small size. In this section, we focus on building a tractable (in terms of minimisation strategy) and powerful (in terms of reliability of the model) regulariser L .

The requirements of the M-step are different for the algorithms Reg1, Reg2, Reg3 and Reg4 developed in the previous sections. More specifically :

- Reg1 has no particular specification on L and T ,
- Reg2 and Reg3 imply that the transformation is represented as the initial position plus a displacement field : $T(x_k) = x_k + t_k$,
- Reg4 implies that the transformation is represented as the initial position plus a displacement field : $T(x_k) = x_k + t_k$ and that this displacement field can be approximated and computed at points that do not belong to the point set.

We propose different regularisers, some of them cannot be applied in all of the previously developed algorithms. To alleviate complicated notations, we restate our problem in the following general form (that we call approximation problem) :

$$\tilde{f} = \arg \min_f \sum_j p_j \|c_j - f(v_j)\|^2 + \alpha' L(f), \quad (8.11)$$

with $\forall j, p_j \in \mathbb{R}^+, c_j \in \mathbb{R}^3, v_j \in \mathbb{R}^3$ and $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$.

8.5.1 Local models

In this section, we devise regularisers based on the decomposition of f into several local transformations. For that purpose, we assign a given transformation model at each point of the point set and we regularise the transformations over the space. This solution implies that the points v_j s are structured as a mesh. In practice, such models provide very efficient optimisation strategies.

8.5.1.1 Locally affine

This way, one can assign an affine transformation that we call T_j to each point v_j of the mesh and spatial coherency is ensured by a regularisation on the T_j over the space (M_j and t_j being the linear and translation components of T_j) :

$$L(f = (T_j)_{j=1, \dots, \text{card}(V)}) = \sum_{(k1, k2) \in C_2} w_{k1, k2} ((1 - \eta) \|M_{k1} - M_{k2}\|_F^2 + \eta \|t_{k1} - t_{k2}\|^2)$$

where $\|\cdot\|_F$ is the Frobenius norm, C_2 the set containing the indices of points of V that are neighbours, $0 < \eta < 1$ weighs the influence of the terms and $w_{k1, k2} > 0$ are

coefficients weighting the relative influence on the regularisation between the different neighbour points to a same point (typically $w_{k1,k2}$ will increase with the spatial distance between v_{k1} and v_{k2}).

Then solving the corresponding approximation problem (Pb 8.13) is performed using the Iterated Coordinate Descent (ICD) algorithm, which consists in considering sequentially each $T_j \in T$ and maximising the criterion 8.13 for T_j while keeping the other transformations fixed. This algorithm converges monotonically to a (at least local) maximum of the criterion, and boils down to :

Algo Approx1: M-step with locally affine deformation model

M-step :

Step 0 : Initialise the \tilde{T}_j

Step 1 : $\forall j \tilde{T}_j = \arg \min_{T_j, t_j} p_j \|c_j - M_j v_j - t_j\|^2$
 $+ \alpha' \sum_{k \in \mathcal{V}_j} w_{k,j} \left((1 - \eta) \|M_j - \tilde{M}_k\|_F^2 + \eta \|t_j - \tilde{t}_k\|^2 \right)$

Step 2 : If the \tilde{T}_j s have changed go to Step 1 else finish

where \mathcal{V}_j contains the indices of neighbours of point c_j in X . It can be shown that Step 1 has a unique solution that is given by the following closed form solution :

$$\left\{ \begin{array}{l} \tilde{M}_j = \left[\frac{\alpha' \eta p_j}{p_j + \alpha' \eta \sum_k w_{k,j}} (v_j \sum_k w_{k,j} - \sum_{k \in \mathcal{V}_j} w_{k,j} \tilde{t}_k) v_j^T + \alpha' (1 - \eta) \sum_{k \in \mathcal{V}_j} w_{k,j} \tilde{M}_k \right] \\ \left[\frac{\alpha' \eta p_j \sum_k w_{k,j}}{p_j + \alpha' \eta \sum_k w_{k,j}} v_j v_j^T + \alpha' (1 - \eta) \sum_k w_{k,j} I_3 \right]^{-1} \\ \tilde{t}_j = \frac{1}{p_j + \alpha' \eta \sum_k w_{k,j}} \left(p_j (c_j - \tilde{M}_j v_j) + \alpha' \eta \sum_{k \in \mathcal{V}_j} w_{k,j} \tilde{t}_k \right) \end{array} \right.$$

8.5.1.2 Local translations

A similar approach consists in decomposing t in a set of linear displacements and to design L as :

$$L(f = (t_j)_{j=1, \dots, \text{card}(V)}) = \sum_{(k1,k2) \in C_2} w_{k1,k2} \|t_{k1} - t_{k2}\|^2$$

leading to

Algo Approx2: M-step with local translation deformation model

M-step :

Step 0 : Initialise \tilde{t}_j

Step 1 : $\forall j \tilde{t}_j = \arg \min_{t_j} p_j \|c_j - v_j - t_j\|^2 + \alpha' \sum_{k \in \mathcal{V}_j} w_{k,j} \|t_j - \tilde{t}_k\|^2$

Step 2 : If the \tilde{t}_j have changed go to Step 1 else finish

where \mathcal{V}_j contains the indices of neighbours of point v_j in V . It can be shown that Step 1 has a unique solution that is given by the following closed form solution :

$$\tilde{t}_j = \frac{1}{p_j + \alpha' \eta \sum_k w_{k,j}} \left(p_j(c_j - v_j) + \alpha' \eta \sum_{k \in \mathcal{V}_j} w_{k,j} \tilde{t}_k \right)$$

8.5.1.3 Discussion

In practice, these two models provide very efficient optimisation strategies. However, each of these models suffers from a lack of theoretical justifications.

- The locally affine model can lead to degenerate solutions as the distance between matrices (based on the Frobenius norm) does not necessarily provide a suitable distance between transformations. A better formulation would directly rely on comprehensive parameters describing the transformation (*i.e.* rotation angles, stretch and scale factors). However such a formulation would result in intractable optimisation strategies.

- The local translation model is more difficult to interpret. However, an interesting comparison with the first order regularisation (*e.g.* [TC07, M'02]) can be done. The first order regulariser L is defined as :

$$L(f = (f^1, f^2, f^3)^T) = \int |\nabla f_1(v)|^2 + |\nabla f_2(v)|^2 + |\nabla f_3(v)|^2 dv \quad (8.12)$$

where $f(v) = (f_1(v), f_2(v), f_3(v))^T$ and ∇ is the gradient operator. In the image grid, such a regulariser can be discretised as :

$$L(f = (t_j)_{j=1, \dots, \text{card}(V)}) = \sum_{(k1, k2) \in C_2} \|t_{k1} - t_{k2}\|^2$$

where C_2 represents the set of cliques of order 2 of the image/volume data. This form appears identical to the one we provided in the previous section. Interestingly, performing Euler Lagrange optimality condition on $L(f)$ gives :

$$\Delta f_1 = 0; \Delta f_2 = 0; \Delta f_3 = 0$$

that can be stated as the convergence at an infinite time of the following system of equations :

$$\frac{\partial f_1}{\partial t} = \text{div}(\nabla f_1); \frac{\partial f_2}{\partial t} = \text{div}(\nabla f_2); \frac{\partial f_3}{\partial t} = \text{div}(\nabla f_3),$$

that corresponds to a heat diffusion process on f and shows a relationship between the first order regulariser and an isotropic Gaussian smoothing.

This analogy helps understanding our regularisation that in practice provides good results. However, it is limited by the fact that the cliques defined in the voxel domain correspond to a well-designed finite model of the first order derivatives whereas this is not the case for the vertices (in the general case).

In the context of medical image registration, we find that the locally translational model provides the more realistic results. However, in the case of articulated motion, the locally affine model provides more adapted results.

8.5.2 Efficient global models

In this section, we devise a new solution for the M-step based on the Reproducing Kernel Hilbert Space (RKHS) theory and the Fourier analysis. The advantage of this solution over the local models is to offer a well designed theoretical framework in which the regulariser is defined as a function of the frequencies of f (allowing to tune efficiently the regulariser) and produces a solution f defined all over the space (allowing to interpolate efficiently the transformation). This framework leads to a closed-form solution for the M-step that can be tackled very efficiently using sparse linear algebra. Moreover, the regulariser L depends on the choice of a kernel that can be easily modified to fit the applications. As we will see in the evaluation section, the regularisers proposed in this section lead to more realistic solutions but are less efficient (in terms of computational time and memory usage) than local models.

8.5.2.1 Formulation

In this framework, we consider that the transformation is represented as the initial position plus a displacement field and thus recast our problem as :

$$\tilde{f} = \arg \min_f \sum_j p_j \|c_j - (v_j + f(v_j))\|^2 + \alpha' L(f), \quad (8.13)$$

with $\forall j, c_j \in \mathbb{R}^+, v_j \in \mathbb{R}^3$ and $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$.

We consider our problem in a space of admissible solutions \mathcal{H} that we span using a positive definite kernel (pdk) $k^1 : \mathcal{H} = \{f | f(\cdot) = \sum_{i=0}^{\infty} k(q_i, \cdot) w_i, w_i \in \mathbb{R}^3, q_i \in \Omega; \|f\|_{\mathcal{H}} < \infty\}^C$ where S^C denotes the completion of the set S and where $\Omega \subset \mathbb{R}^3$. This space is endowed with the inner product $\langle f, h \rangle_{\mathcal{H}} = \sum_{i,j=0}^{\infty} w_i^T k(q_i, q_j) w_j$. The space \mathcal{H} is a Hilbert space with reproducing kernel k (or more compactly a RKHS) [Wen05]. Then we assume that $f \in \mathcal{H}$ and define our regulariser $L(f)$ as $\|f\|_{\mathcal{H}}$:

$$\tilde{f} = \arg \min_{f \in \mathcal{H}} \sum_j p_j \|c_j - (v_j + f(v_j))\|^2 + \beta \|f\|_{\mathcal{H}}. \quad (8.14)$$

One of the key advantage of RKHS is that one can show [SHS01] that the values taken by the solution \tilde{f} at the points $v_1, \dots, v_j, \dots, v_M$ can be expressed as $\tilde{f}(v_j) = \sum_{i=1}^M k(v_i, v_j) w_i$ and then formulate the last minimisation problem as :

$$(\tilde{w}) = \arg \min_{(w)} \sum_{j=1..M} p_j \|c_j - (v_j + \sum_{i=1..M} k(v_j, v_i) w_i)\|^2 + \beta \sum_{i,j=1..M} w_i^T k(v_j, v_i) w_j$$

Vanishing the derivatives gives a linear system whose solution can be expressed in a closed-form as :

$$W = (D(P)K + \beta I)^{-1} D(P)[C - V],$$

¹More generally, we could use a vectorial positive definite kernel \mathbf{k} (in our case, $\mathbf{k}(\cdot, \cdot)$ would be a 3×3 matrix). By simply considering a scalar pdk, we indirectly restrict our study to vectorial pdk \mathbf{k} of the form $\mathbf{k}(\cdot, \cdot) = k(\cdot, \cdot)I$. However, note that all the results of this section can be extended to vectorial pdks.

where $V = [v_1, \dots, v_M]^T$, $C = [c_1, \dots, c_M]^T$, $W = [w_1, \dots, w_M]^T$, $K = (k(v_i, v_j))_{i,j}$ is the M by M matrix associated to kernel k and $D(P)$ is the diagonal matrix formed by the p_j values.

Now, the challenge is to choose a kernel corresponding to a relevant regulariser.

8.5.2.2 Choosing a kernel

In order to design a suitable k , one can use an interesting relationship with Fourier-based stabilisers. Let $\forall f$ integrable,

$$L(f) = L(f_1) + L(f_2) + L(f_3)$$

where for $i = 1, 2, 3$,

$$L(f_i) = \frac{1}{(2\pi)^3} \int_{-\infty}^{\infty} |f_i^*(\omega)|^2 \phi^*(\|\omega\|/b)^{-1} d\omega,$$

where $*$ is the Fourier transform operator, $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is an integrable function and b is a real positive rescaling factor. Let $\mathcal{F} = \{f : \mathbb{R}^3 \rightarrow \mathbb{R}^3 | L(f) < \infty\}$. Interestingly, one can state that if the function $(q_i, q_j) \rightarrow \phi(\|q_i - q_j\|)$ is a pdk then \mathcal{F} is a RKHS whose reproducing kernel is given by $k(q_i, q_j) = b \times \phi(b \times \|q_i - q_j\|)$ and such that $\forall f \in \mathcal{F}$, $\|f\|_{\mathcal{F}} = L(f)$ [Wen05, Sid04]. This dual view is convenient as it allows to design a wide variety of efficient regularisers directly into the Fourier domain.

In order to design an efficient regulariser, we have to choose ϕ^{*-1} as a high-pass filter. This way, high frequencies of the deformation will be drastically penalised whereas low frequencies will only be penalised a little. Generally, $\phi_{[0,\infty]}^*$ is a monotonically decreasing function and the most important element that characterises its influence on the regularisation is the way it decreases that indicates the amount of penalisation with respect to frequencies. Particularly, the frequencies for which $\phi^*(\|\omega\|/b)$ is null are forbidden. The two parameters β and b allow to handle the regularisation properties : β is a quantitative parameter (it indicates the amount of smoothness) whereas b is more qualitative (in a way, it defines what the term "smoothness" means). More precisely, b can be seen as a parameter contracting (resp. dilating) the kernel function $\phi^*(.)$ and thus decreasing (resp. increasing) the range of admissible frequencies. Figure 8.5 illustrates results of approximation of 2D noise deformation fields with different parameters β and b for a given kernel.




FIG. 8.5 – Effect of parameters β and b on the approximation of a noisy field.

Finally, note that an undesirable effect of this approach is to penalise the null frequency *i.e.* the average of the field (as $1/\phi^*(0)$ is not null). One removes this penalisation by simply ensuring that the deformation field has the same norm before and after regularisation. Figure 8.5 shows the influence of b and β when approximating a noisy field when choosing ϕ as the Wu kernel [Wen05].

Why using frequencies ? Designing L as a function penalising T in terms of its spatial frequencies can be of great interest. Roughly speaking, high frequencies of T concern details and local changes of the deformation field whereas its low frequencies concern the global aspects of the deformation. Thus, tuning the parameter b allows to introduce a multiscale approach by first trying to capture a global deformation and then, if needed, local deviations from this global deformation. This view allows to adapt our algorithm to different applications needing either a fine registration (*e.g.* automatic labeling of substructures) or a more global registration (*e.g.* statistical shape analysis) of two structures. Figure 8.6 illustrates the influence of different kernels and different scale parameters b on the regularisers. Figure 8.7 illustrates that modifying β allows to characterise different scales of deformations linking two structures.

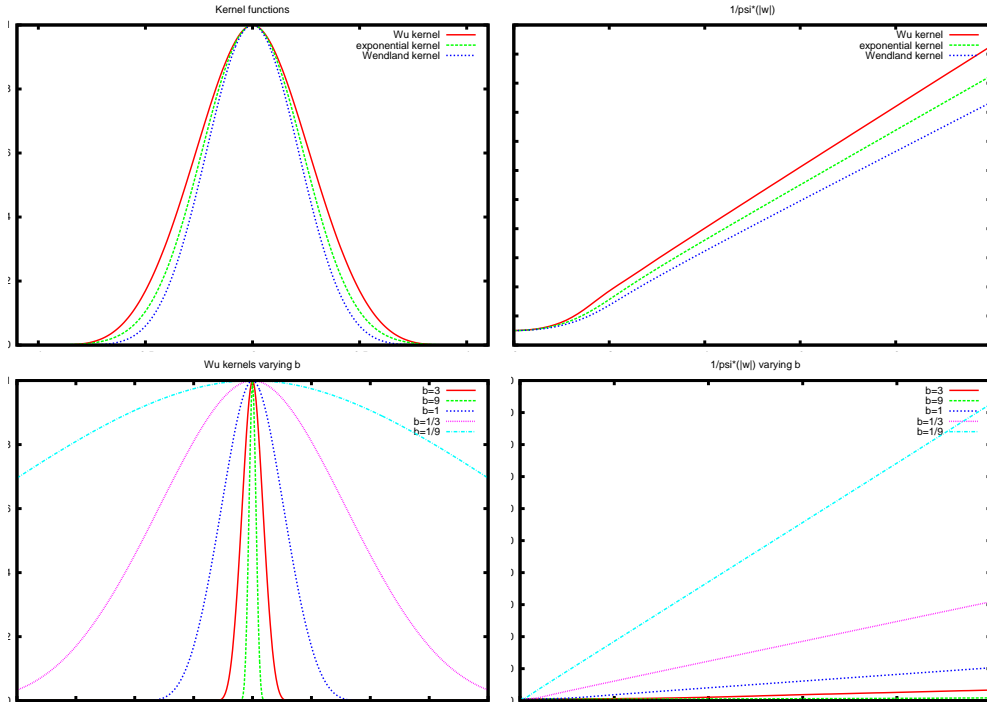


FIG. 8.6 – **Different positive definite kernels ϕ and their associated value $\phi^*(\|\omega\|/b)^{-1}$.** From left to right and top to bottom : i) 2D plot of three kernels ($b=1$), ii) 2D plot of the value of $\phi^*(\|\cdot\|)^{-1}$ for the three above-mentioned kernels, iii) 2D plot of $\phi(\cdot \times b)$ for the Wu kernel with different b , iv) 2D plot of $\phi^*(\|\omega\|/b)^{-1}$ for the Wu kernel with different b .

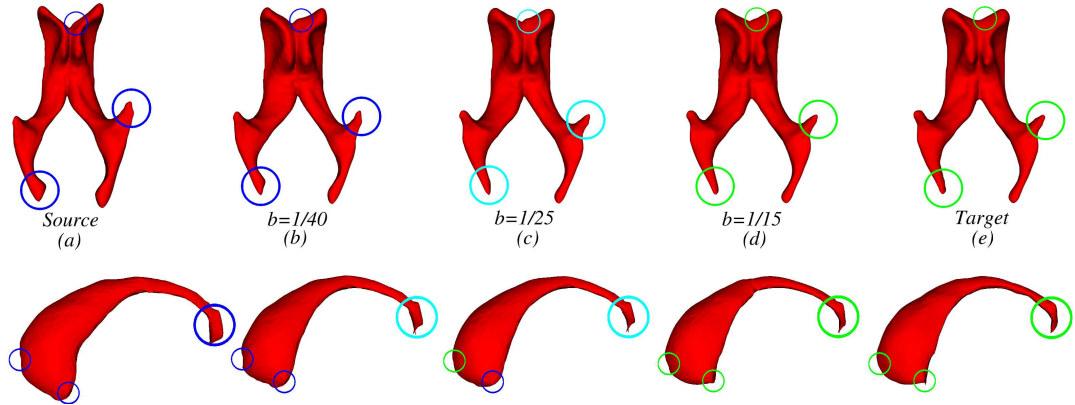


FIG. 8.7 – **Influence of b on the registration.** **First line :** We register the source ventricle (a) on the target ventricle (e) by increasing b throughout the iterations (without modifying the other parameters). Intermediate registration results for a given b are represented between the source and the target (b), (c) and (d). For a small b , the source is only globally transformed towards the target and small patterns are left unchanged (dark blue circles). When b increases, small patterns tend to fit the target (green circles). Light blue circles indicate intermediate configurations of the patterns. **Second line :** Same description that for the first line but with caudate nuclei.

8.5.2.3 Efficient choices

Although we propose a closed-form solution for the approximation problem, it consists in solving a $M \times M$ system. This can be problematic when M increases (in terms of memory usage and of computational time). Suppose that we choose a compactly supported pdk (*i.e.* for $\forall x, \forall y$ such that $\|y\| > r; k(x, y) = 0$), then i) $D(P)K + \beta I$ is a sparse matrix that can be computed using a kd -tree and ii) computing W consists in solving a sparse system. Some interesting compactly supported pdk corresponding to low-pass filters have been proposed in the literature (such as Wendland, Wu or Buhmann functions). Moreover, techniques to generate a wide variety of them have been proposed [Wen05]. Alternatively one can use a highly decreasing function and approximate it by zeroing all its values over a given threshold. We experimentally find the compact support kernel of Wu ($\phi_{2,3}$) as the one providing the best results and we will use it in the following. Note that the top-right plot in Figure 8.7 shows that the Wu kernel penalises faster high frequencies than the exponential or Wendland kernels (for a given b). More experiments will be of interest to evaluate the respective added value of the different possible kernels.

8.5.2.4 M-step in a nutshell

Algo Approx3: M-step with RKHS

M-step :for each $v_i \in V$ $\mathcal{S}_i = \{v_k \in V \text{ such that } \|v_k - v_i\|^2 < b\}$ (using a *kd*-tree)for each v_k in \mathcal{S}_i

$$K(i, k) = b \times \phi(b \times \|v_k - v_i\|) \times \tilde{p}_i$$

$$K(i, i) = K(i, i) + \alpha'$$

preconditioning on K (using sparse algebra)solve $KW^1 = D(P)[V^1 - C^1]$ (using sparse algebra)solve $KW^2 = D(P)[V^2 - C^2]$ (using sparse algebra)solve $KW^3 = D(P)[V^3 - C^3]$ (using sparse algebra)

where V^1 , V^2 and V^3 are respectively the vectors extracted from the first, second and third columns of matrix V (the same for C and W).

8.5.2.5 Note on vectorial kernels

By building a scalar kernel k we consider each component of the deformation field independently ($L(f) = L(f_1) + L(f_2) + L(f_3)$). In fact, all the previous theorems stay true when designing vectorial positive definite kernels. In our case, the evaluation $k(., .)$ will take its values in the set of the 3 by 3 matrix, the generated space will be such that $\mathcal{H} = \{f | f(.) = \sum_{i=0}^{\infty} k(q_i, .)w_i, w_i \in \mathbb{R}^3, q_i \in \mathbb{R}^3; \|f\|_{\mathcal{H}} < \infty\}^C$ and the dot product writes $\langle f, h \rangle_{\mathcal{H}} = \sum_{i,j=0}^{\infty} w_i^T k(q_i, q_j)w_j$. This way, the corresponding Fourier-based stabiliser writes :

$$L(f) = \frac{1}{(2\pi)^3} \int_{\mathbb{R}^3} f^*(\omega) \phi^*(\omega)^{-1} f^*(\omega) d\omega$$

With this formulation the cross frequencies (and indirectly cross derivatives) can be penalised . In our work, we restrict our study to scalar kernels.

Chapitre 9

Evaluation

9.1 Descriptions

9.1.1 Descriptions of the algorithms and of the parameters

We use the following notations : **Algo-Approx**.


With **Algo** \in {Algo1 (simple EM-ICP), Algo2 (symmetric EM-ICP), Algo3 (symmetric EM-ICP + priors), Algo4 (symmetric EM-ICP + priors + consistency)} and **Approx** \in {Approx1 (Local translation), Approx2 (RKHS)}.

We also investigate the performances of a state-of-the-art method the TPS-L2[JV05] using the code provided by the authors on the web. Note that we try to evaluate both the TPS-RPM [CR00a] and CPD [MSCP07] algorithms on our data but none of them has provided results after 3 hours (even for the smallest data).

9.1.2 Description of the data


9.1.2.1 Thalamus

We use 13 right thalamus segmented with an in-house software from 3T T1-weighted MRI of healthy subjects. The data contain about 1.000 points.




9.1.2.2 Caudate nuclei

We use 22 Caudate nuclei segmented with an in-house software from 3T T1-weighted MRI of healthy subjects. The data contain about 2.000 points.




9.1.2.3 Ventricles

We use 3 ventricles segmented with itknap (itknap.org) from 3T T1-weighted MRI of healthy subjects. The data contain about 10.000 points.



9.1.2.4 Osseous labyrinths

We use 10 Osseous labyrinths (5 right and 5 left symmetrised) segmented with Amira (amira.com) from CT of modern Homo sapiens. The data and contain about 20.000 points.



9.1.2.5 Endocasts

We use 19 brain endocasts segmented with itknap (itknap.org) from CT of great apes. The data contain about 20.000 points.



9.1.3 Description of the experiments

9.1.3.1 Experiment 1 : MSE on simulated data

In Experiment 1, we evaluate the ability of our different algorithms to recover global and local smooth deformations. This allows to evaluate both the effectiveness of our different minimisation strategies and the ability of our two deformation models to fit other deformation models widely used in the medical imaging context. For this purpose, we choose one surface in each of the 6 previously described datasets and deform these surfaces 100 times using :

- randomly generated non-linear transformations using the widely used TPS or CPD deformations (see Section 7.2.2 for more details),
- randomly generated local smooth deformations of the form : $t(x) = x + K \times G_v(x - x_c)n_x$, where x_c is the centre of the deformations (randomly chosen on the surfaces), n_x is the normal vector at point x , G_v is a 3D non-normalised Gaussian function of variance v^2 and K is the deformation strength.

Then we register the original and deformed surfaces and compute the overall residual distance between the known corresponding points between the surfaces and the overall

Barron error between the transformation field¹. These two error measures are then averaged over the 100 simulations for each of the surfaces.




FIG. 9.1 – **Illustration of the generation of the ground truth data.** From left to right : i) original data, ii) we generate a random local deformation (the resulting distance between the corresponding points is mapped), iii) we generate a random global deformation (TPS) (the resulting distance between the corresponding points is mapped) and iv) superimposition of both the original (red) and deformed (green) data.




FIG. 9.2 – **Examples of pairs of ground truth data :** From left to right, thalamus, ventricles, osseous labyrinths and endocasts

9.1.3.2 Experiment 2 : Landmark-to-landmark distance on real data

In Experiment 2, we investigate the ability of our algorithm to register real data. For this purpose, we manually select landmarks on each of 2 datasets (ventricles and osseous labyrinths). Then we choose one of the surfaces in each of the two datasets to be the template, and register all other surfaces to this template. We evaluate the mean residual errors on the landmarks and we average these errors over all subjects for each of the two datasets.




FIG. 9.3 – Position of landmarks on two of our data

	algo1-reg1	algo2-reg1	algo3-reg1	algo1-reg2	algo2-reg2	algo3-reg2	algo4-reg2	TPS-L2
thalamus	1.02/1.96	0.21/1.02	0.23/0.76	0.82/1.45	0.12/0.67	0.16/0.57	0.17/0.65	0.14/0.57
caudate	1.82/2.72	0.63/1.89	0.70/1.21	1.34/1.78	0.31/1.21	0.35/0.97	0.32/1.19	0.33/1.05
ventricle	1.94/3.94	0.79/1.91	0.87/1.13	1.43/2.99	0.64/1.66	0.65/1.56	0.66/1.53	-
osseous	0.29/0.87	0.17/0.63	0.21/0.44	0.24/0.86	0.11/0.52	0.09/0.45	0.14/0.49	-
endocast	3.23/6.76	2.78/4.77	3.09/4.63	2.99/6.11	2.01/3.61	2.11/2.92	2.12/3.01	-

TAB. 9.1 – Statistics (mean/max) on the end point error

9.2 Results

9.2.1 Experiment 1

The Barron error on deformation field gives similar results expect that the difference of accuracy between reg1 and reg2 is more accentuated (reg2 decreases the error from about 44% with respect to reg1 with the same algo and on a given data). The conclusions on these results are given in Section 9.3.

9.2.2 Experiment 2

	algo1-reg1	algo2-reg1	algo3-reg1	algo1-reg2	algo2-reg2	algo3-reg2	algo4-reg2	TPS-L2
ventricle	2.41/3.01	1.43/2.03	1.37/1.98	1.24/2.43	1.05/1.53	1.06/1.22	1.05/1.34	-
osseous	0.91/1.87	0.61/1.21	0.59/1.01	0.64/1.64	0.35/0.56	0.34/0.57	0.34/0.56	-

TAB. 9.2 – Statistics (mean/max) on landmark-to-landmark distance

The conclusions on these results are given in Section 9.3.

9.2.3 Additional measurements

Run-time

¹the Barron error between two vectors t_1 and t_2 is defined as $\arccos\left(\frac{t_1^T t_2}{\|t_1\| \|t_2\|}\right)$

	algo1-reg1	algo2-reg1	algo3-reg1	algo1-reg2	algo2-reg2	algo3-reg2	algo4-reg2	TPS-L2
thalamus	0m06s	0m16s	0m19s	0m14s	0m19s	0m21s	0m37s	5m12s
caudate	0m13s	0m22s	0m24s	0m32s	0m40s	1m00s	2m20s	5m37s
ventricle	0m45s	1m02s	1m12s	4m04s	4m50s	5m24s	10m10s	fail
endocast	1m05s	1m23s	1m29s	8m12s	9m04s	12m17s	23m01s	fail

TAB. 9.3 – Mean execution-time on a standard personal computer needed by the different algorithms for the different structures

Each program is compiled using the same compiler and the same compilation options. Table 9.3 indicates mean run-time of the different algorithms for the different structures.

The conclusions on these results are given in Section 9.3.

9.3 Summary

We summarise our results as follows :

- All our algorithms allow to register our data whereas none of the 3 state-of-the-art methods we investigated was able to register point sets larger than 5000 points.
- Algo3-reg2, Algo3-reg2 and Algo4-reg2 provide results comparable to the state of the art TPS-L2 (on the small data set).
- Fourier-stabiliser based regularisation leads to lower errors (especially for the Barron error) than local models. However, local models are computationally more attractive than Fourier-stabiliser.
- For any regulariser, Algo-2 (symmetric EM-ICP) importantly improves the results.
- Adding priors on the match matrix (Algo-3) does not improve the mean results (except for the osseous labyrinth, that are complicated and curved structures) but it decreases the value maximum error over the experiments.
- Algo-4 (Symmetric EM-ICP + consistency) is significantly computationally more demanding than all other algorithms and does not improve the registration in average.

Further evaluation is needed and will be of great interest. In particular, we are currently working on evaluating the influence of b for different applications.

Chapitre 10

Conclusion

10.1 Contributions

In this part, we first proposed a review of registration methods based on the modeling of point sets as mixture models, aiming to provide a comprehensive view where we explicated the sometimes hidden i) mixture models, ii) similarity measure, iii) expression of a transformation model T (generally relying on a regulariser $L(T)$) allowing to "deform" a mixture and iv) subsequent minimisation.

Then, we showed that a simple probabilistic modelling of the registration problem allows to solve it by using standard clustering techniques. In particular, the unknown non-linear transformation best superposing the two point sets can then be estimated according to the maximum *a posteriori* principle (MAP) and using (typically) the EM algorithm. This optimisation then boils down to a simple iterative estimation of fuzzy point-to-point correspondences (E-step) (encoded in what is often termed the *match matrix*) and of the non-linear transformation (M-step) in turn. The resulting strategy is limited by three main aspects. Firstly, the derivation of the MAP principle leads to an asymmetric formulation. In particular, in this framework, the match matrix arises as a row stochastic matrix (leading to many-to-one matchings). This asymmetric formulation makes the algorithm unable to achieve a good matching in specific cases. Secondly, both the E-step and the M-step are very time and memory consuming when dealing with large data. Finally, the overall iterative scheme exhibits a monotonic convergence that leads to a local maximum of the likelihood function and thus can provide a bad estimate of the transformation when a bad initialisation is used. We proposed efficient solutions for each of these problems.

In Section 8.2, we proposed to tackle problems due to the asymmetric nature of the MAP principle. For this purpose, we noticed that the derived EM algorithm can be seen as an iterative alternated minimisation (over the match matrix and the unknown transformation) of an (energetic) criterion and proposed to modify it. In particular, we explicitly added a second match matrix that is column stochastic. This modification only changes the E-step and improves the matching process. The modification of the energetic framework alters the interpretation of the probabilistic framework and we do not know

whether it is still possible to interpret the criterion in the probabilistic framework. In practice, we observed that this new constraint can be seen as an optimisation tool. Indeed, the final value of the original EM-ICP criterion :

$$\mathcal{E}2(X, T(Y), A, B) = \sum_{j,k} A_{jk} \rho_{\delta}(\|y_j - T(x_k)\|^2) + 2\sigma^2 \sum_{j,k} A_{jk} \log(A_{jk}) + \alpha' L(T),$$

is generally higher when using the classical EM-ICP algorithm than when using the algorithm minimising the new criterion. Thus, we adopt the following arrangement : we first register the point sets using the column and line normalisation constraint and, at the end of the convergence we make a few iterations removing the column constraint. It allows us to benefit the minimisation ability of the symmetrised constraint and to keep a probabilistic interpretation for A (*a posteriori* probability of correspondences between points).

In Section 8.3, we showed how to specify priors on the match matrix with only minor changes to the optimisation algorithm. These priors based on local and global shape descriptors allow to significantly improve the capture range and the convergence speed of the algorithm.

In Section 8.4, we proposed a symmetric consistent formulation of the registration problem that allows to jointly compute the forward and the backward deformation fields linking the two point sets which are enforced to be compatible with each other.

In Section 8.5, we devised two alternative efficient solutions for the M-step. The first one is based on the regularisation of simple local models (translation or affine) attributed to each point of the set and regularised over the space. The second one stands on the Reproducing Kernel Hilbert Space (RKHS) theory and on the Fourier analysis and consists in building efficient regularisers leading to closed-form solutions based on sparse linear algebra. Both of these methods lead to accurate results when dealing with deformation generated from classical models. The local models are particularly relevant to deal with large data in a short time. The RKHS-based model seems to lead to more accurate solutions but is more time and memory demanding which can be limiting in presence of very large data sets.

As shown in Section 9, the combination of these improvements leads for the first time to an efficient EM-based algorithm for the non-linear registration of large point sets (with the exception of the recent work of Myronenko and colleagues [MS09] that we did not yet investigate).

Finally, note that as we will see in Part III, our probabilistic modelling leads to efficient and simple solutions to deal with group-wise study : this fact is important to judge our methodological choices.

10.2 Perspectives

We think that more evaluations and comparisons are needed. First, one can go closer to real situations when generating ground truth data by for example modifying the sampling of the point sets. Moreover our database of real data was quite small and limits our evaluations on real data.

In the following we address some possible improvements that could be investigated in further works.

Fitting the data : Our methodology relies on modelling point sets as mixture of pdf. An interesting work would consist in trying to fit mixture models directly on the point sets (or on the surface) by estimating both mean and variance of a set of Gaussian laws. Then the registration could be done on the fitted models. We think that such an approach will provide less sampling dependent results.

Data-attachment : In other works (not presented in this thesis) [MCP10], we investigated solutions to use normal in the classical data-attachment term. This framework is still computationally intractable because of the need of computing, at each iteration of the algorithm, an orthogonal parametrisation of the transformed surface. However, by designing efficient numerical tools for such a computation, this approach could improve the performance of our algorithm by increasing its capture range.

Deformation : As previously shown, Fourier-based regularisers offer a simple and suitable way to compute the deformations linking the point sets at different scales. The design of an elaborated and well-grounded strategy to compute/analyse these different deformations will be of great interest.

Moreover, the design of efficient diffeomorphic splines approximation would be helpful to make possible the computation of diffeomorphic transformations in our formalism.

Robustness : We proposed to use a very simple robust estimator (truncated square function) whereas the literature abounds with robust estimators. Replacing our estimator with a more elaborated M-estimator would be benefic. However, we chose a truncated square function as it corresponds to a truncated Gaussian law in the clustering framework and thus allows to preserve a probabilistic interpretation of the variables (which will be of great interest in the following).

Optimisation : It would be quite interesting to maximise the likelihood function directly using a gradient-based method instead of an EM algorithm and to investigate the superiority (or the complementarity) of one of the methods.

Troisième partie

Applications

Chapitre 11

Tools for group studies & Application

Registration methods have a huge number of applications in medical imaging *e.g* neuro-navigation, segmentation, building of statistical atlases, *etc*. Among them, statistical shape analysis is of great interest to characterise and compare anatomical structures over and between populations in an automatic and comprehensive way. In this chapter, we develop pipelines for statistical shape analysis (Section 11.1) and quantification of asymmetries over and between populations of anatomical structures represented by point sets optionally structured as a mesh (Section 11.2). These pipelines rely on the tools we proposed in the previous parts of this thesis.

11.1 Statistical Shape Analysis

11.1.1 Introduction & Brief state-of-the-art




FIG. 11.1 – Mean shape and first mode of variation ($\pm 2 \sqrt{\lambda_1}$) on 40 human hands (described by labelled landmarks) computed with a classical Procrustes analysis (From [SG02]) and on 10 osseous labyrinths (described by free-form surfaces) computed with an automatic method described in this chapter (Section 11.1.2.4).

The challenge underlying statistical shape analysis consists in building a statistical model describing the behaviour of the observed structures in terms of probability distributions. More specifically, this analysis generally consists in i) defining and computing a mean shape M from the observations (the shape that is the closest, in a sense to be

defined, to all shapes of the population) and ii) characterising the main variations of the shapes around M . When each shape is represented by a set of labelled landmarks, a standard statistical model is the point distribution model [CTCG95] classically based on the Procrustean framework [DM98] and a principal component analysis. When it is not the case, *i.e.* when the shapes are represented by implicit surfaces, meshes or simple point sets, alternative strategies have been proposed. Among them one can quote :

- Methods based on the modelling of shapes by currents. The space of currents is linear and performing statistics over it is straightforward [DPTA09].
- Methods based on the minimum description length (MDL) theory [HWWpM05]. The key principle of these methods is to consider that the simplest explains better : one looks for the most compact model (*i.e.* involving as few parameters as possible) that best explains the population. The correspondences between the shapes and the model are computed as a trade-off between a maximum compactness of the model and a minimum error between the data generated by the model and the data under study.
- Methods based on the definition of the mean shape as the minimiser of a registration-like criterion involving Gaussian mixtures modeling of the surfaces [CRZL04, HPE⁺08b, HEP⁺09, WVRE08].

Among this last class of methods, the Mean Shape EM-ICP (MS-EM-ICP) recently developed by Hufnagel and colleagues [HPE⁺08b] is based on the original EM-ICP developed by Granger and Pennec [GP02]. As a result, it shares many properties with the registration method we presented in part 2. In the following, we first present the original mean shape EM-ICP algorithm and study its limitations (Section 11.1.2.2). Then, we propose two extensions of the MS-EM-ICP aiming at tackling these limitations (Section 11.1.2.3 and Section 11.1.2.4).

11.1.2 Mean shape EM-ICP and extensions

11.1.2.1 Notations

Let :

- $X^c = \{x_1^c, \dots, x_{n^c}^c\}$ be a point set representing a shape under study, where $x_i^c \in \mathbb{R}^3$ is a 3D point. Note that each point set does not necessarily contain the same number of points.
- $X = \{X^1, \dots, X^c, \dots, X^C\}$ be the set of C shapes defining a population under study.
- $M = \{m_1, \dots, m_N\}$ be the mean shape corresponding to the population X . We consider the number of points N (arbitrary) fixed.

11.1.2.2 The mean shape EM-ICP

Following the work of Hufnagel *et al.* [HPE⁺08b], we propose to define the mean shape \tilde{M} as the one that maximises a likelihood function modelling \tilde{M} as noised versions

of $T^c(X^c)$ ($\forall c$)¹. This likelihood function, very close to the one presented for the registration problem in Section 8.1, can be maximised using an EM algorithm which can be seen as an alternated minimisation of a criterion involving 3 main unknown interdependent parameters $T^c/A^c/M$. This criterion writes :

$$\tilde{M} = \arg \min_{M, A^c, T^c} \sum_{c=1}^C \left[\sum_{x_i^c \in X^c} \sum_{m_j \in M} A_{i,j}^c \|T^c(x_i^c) - m_j\|^2 + 2\sigma^2 \sum_{x_i^c \in X^c, m_j \in M} A_{i,j}^c \log(A_{i,j}^c) \right],$$

with $\forall c, \forall i, \sum_j A_{i,j}^c = 1$,

where :

- for all c , $A^c = (A_{i,j}^c)$ is the match matrix encoding the *a posteriori* probabilities of correspondence between points of X^c and of M ,
- T^c is either a rigid-body or an affine transformation superposing X^c on M .

In this framework, one can interpret the mean shape as the one for which the “fuzzy residuals” $\sum_{x_i^c \in X^c} \sum_{\tilde{m}_j \in \tilde{M}} \tilde{A}_{i,j}^c \|\tilde{T}^c(x_i^c) - \tilde{m}_j\|^2$ subject to optimal transformations \tilde{T}^c and optimal match matrices \tilde{A}^c is minimal. This criterion is minimised until local convergence by the following 3-step iterative algorithm (that is an EM algorithm) :

Init : \tilde{M} is one of the shape from X . $\forall c, \tilde{T}^c = Id$

Step 1 : update \tilde{A}^c
for each c , for all i, j , $\tilde{A}_{i,j}^c = \frac{\exp(-\|\tilde{T}^c(x_i^c) - \tilde{m}_j\|^2 / (2\sigma^2))}{\sum_k \exp(-\|\tilde{T}^c(x_i^c) - \tilde{m}_k\|^2 / (2\sigma^2))}$

Step 2 : update \tilde{T}^c
for each c , $\tilde{T}^c = \arg \min_{T^c} \sum_i \sum_j \tilde{A}_{i,j}^c \|T^c(x_i^c) - \tilde{m}_j\|^2$

Step 3 : update \tilde{M}
for all j , $\tilde{m}_j = \frac{1}{\sum_c \sum_i \tilde{A}_{i,j}^c} \sum_c \sum_i \tilde{A}_{i,j}^c \tilde{T}^c(x_i^c) = \frac{1}{C} \sum_c \sum_i \tilde{A}_{i,j}^c \tilde{T}^c(x_i^c)$

Although this method seems simple and natural, it exhibits some limitations. In particular, the transformations T^c are modeled as linear, similarity or rigid-body transformations. As a result, the criterion is likely to provide a poor estimate of the correspondences between point sets and may fail to characterise complicated structures. An illustration of this flaw is given in the first row of Figure 11.2. In the following, we propose to improve the previous definition of the mean shape by considering the transformations T^c ($c \in [1, \dots, C]$) as non-linear transformations subject to regularisations $L(T^c)$. Interestingly, depending on the way we interpret the original MS-EM-ICP criterion, this improvement can lead to two different paradigms. These two paradigms are equivalent in the case described by Hufnagel and colleagues (when T^c is a simple

¹In fact, Hufnagel and colleagues define $T^c(M)$ as a noised version of X^c , $\forall c$: this is the forward scheme. Theoretically speaking, the forward scheme is better than the backward scheme (that we propose to use). In practice, we chose to adopt the backward scheme because it simplifies considerably the notations and the interpretation but it is possible to adapt all the statements of this part to fit the forward view.

linear transformation) but quite different when dealing with non-linear deformations. In the two following paragraphs, we discuss these two paradigms that lead respectively to design \tilde{M} as a Fréchet-like mean or as a Procrustean-like mean.

11.1.2.3 First interpretation : the Fréchet-like mean

A first view simply consists in rewriting the last criterion as :

$$\tilde{M} = \arg \min_M \sum_c d(M, X_c)^2$$

where :

$$d(M, X_c)^2 = \min_{A^c, T^c} \sum_{x_i \in X^c} \sum_{m_j \in M} A_{i,j}^c \|T^c(x_i^c) - m_j\|^2 + 2\sigma^2 \sum_{i,j} A_{i,j}^c \log(A_{i,j}^c).$$

If $d(.,.)$ defined a proper distance, this would be the definition of the Fréchet mean of X for d . Then the natural solution to handle non-linear transformations consists in replacing the linear transformations by a non-linear transformations subjects to regularisations L :

$$d(M, X_c)^2 = \min_{A^c, T^c} \sum_{x_i \in X^c} \sum_{m_j \in M} A_{i,j}^c \|T^c(x_i^c) - m_j\|^2 + 2\sigma^2 \sum_{i,j} A_{i,j}^c \log(A_{i,j}^c) + \alpha' L(T^c).$$

where α' weighs the importance of the regularisation over the data attachment term. Roughly speaking, this criterion defines the mean shape \tilde{M} as the shape that minimises both the “mismatches” $\sum_{x_i \in X^c} \sum_{m_j \in M} A_{i,j}^c \|T^c(x_i^c) - m_j\|^2$ and the deformation energies $\alpha' L(T^c)$ between each X^c and M . The choice of the parameter α' and of L is crucial as it has a strong impact on the definition of the mean shape and thus on the interpretation of the subsequent statistical model. A large α' tends to favour regularisation over mismatching whereas a small α' penalises it. However, note that $\alpha' L(T^c)$ has to increase sufficiently when T^c moves away from non-linear solutions. As an illustration, if $\alpha' \rightarrow 0$ then for all M , there exists T^c such that $\sum_{x_i \in X^c} \sum_{m_j \in M} A_{i,j}^c \|T^c(x_i^c) - m_j\|^2 = 0$ and $\alpha' L(\tilde{T}^c) \rightarrow 0$. As a result, the optimal mean shape \tilde{M} can take an arbitrary value, which is clearly inappropriate.

Moreover in some applications, it is meaningless to capture in the model simple differences of global orientations of the shape X^c (as they are most of the time completely arbitrary). Then, the regulariser L has not to penalise such differences and thus has to be invariant to the rigid-body (or rigid-body plus scale) component of the deformations.

Similarly to the original MS-EM-ICP, this criterion can then be minimised by the following algorithm :

Step 1 :	update \tilde{A}^c
	for each c , for all i, j , $\tilde{A}_{i,j}^c = \frac{\exp(-\ \tilde{T}^c(x_i^c) - \tilde{m}_j\ ^2 / (2\sigma^2))}{\sum_k \exp(-\ \tilde{T}^c(x_k^c) - \tilde{m}_j\ ^2 / (2\sigma^2))}$
Step 2 :	update \tilde{T}^c :
	for each c , $\tilde{T}^c = \sum_i \sum_j \arg \min_{T^c} \tilde{A}_{i,j}^c \ T^c(x_i^c) - \tilde{m}_j\ ^2 + \alpha' L(T^c)$
Step 3 :	update \tilde{M} :
	for all j , $\tilde{m}_j = \frac{1}{C} \sum_c \sum_i \tilde{A}_{i,j}^c \tilde{T}^c(x_i^c)$

Notice that, by using our contributions, one can improve this method by symmetrising the matching process (Section 8.2), imposing symmetric consistency of the transformation (Section 8.4), adding landmarks/descriptors/labels (Section 8.3).

11.1.2.4 Second interpretation : Procrustean mean shape on unlabelled point sets

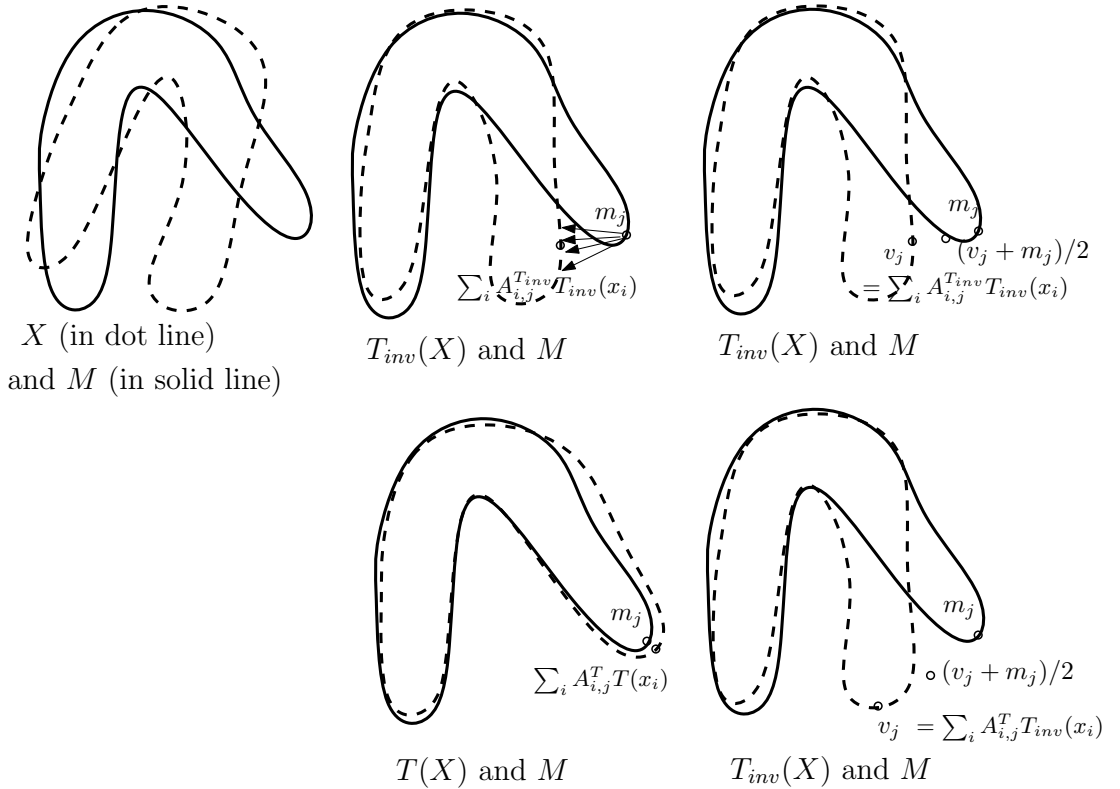


FIG. 11.2 – **First row : Illustration of the classical MS-EM-ICP.** From left to right : i) the point set X and M , ii) $T_{inv}(X)$ and M (T_{inv} is a rigid-body transformation) and the resulting matching $\sum_i A_{i,j} T_{inv}(x_i)$ and iii) the resulting correspondences $v_j = \sum_i A_{i,j} T_{inv}(x_i)$ and mean $(v_j + m_j)/2$, **Second row : Illustration of the improved MS-EM-ICP.** From left to right : i) $T(X)$ on M (T is a non-linear smooth transformation) and the resulting matching $\sum_i A_{i,j} T(x_i)$ and ii) the resulting correspondences $v_j = \sum_i A_{i,j} T(x_i)$ and mean $(v_j + m_j)/2$.

Procrustean mean is a Fréchet mean for the space of labelled point sets. More specifically, it is defined on a shape space \mathcal{S} [DM98]. In this framework, the shape of an object is defined as the geometrical information that remains when we filter out translation, rotation (and optionally scale). Then \mathcal{S} is defined as the set of all possible shape of an object. More formally, \mathcal{S} is defined as the orbit space of the non-coincident point set configurations under the action of the rigid-body (or similarity) transformation.

It is quite different from the previously developed mean as the departure from a shape to another is then considered as the sum of the squared distances between each correspondent point (which are assumed to be known) after an optimal rigid alignment of the shapes (assuming that M and V are set of labelled points) :

$$d(M, V)^2 = \arg \min_R \sum_i \|R(v_i) - m_i\|^2. \quad (11.1)$$

In other words, the shape is what remains after registration (*i.e.* the residuals in the shape space) and the departures from the mean are computed only point-to-point distances (whereas the value of bending energy $L(T^c)$ is taken into account in the design of the departure to the mean in the Fréchet-like mean). By identification, in the original MS-EM-ICP, the transformation underlying this registration is modeled by T^c and the residuals $\sum_i A_{i,j} \|\tilde{T}^c(x_i^c) - m_j\|^2$ defines the departure from X^c to M .

Following this view, the original MS-EM-ICP can be interpreted as a two-step algorithm that consists in : i) finding correspondences in X^c of the points $\tilde{m}_j : \tilde{v}_j^c = \sum_i \tilde{A}_{i,j} \tilde{T}^c(x_i^c)$ (after an optimal rigid alignment T^c) and ii) computing the subsequent mean : $\tilde{m}_j = \frac{1}{C} \sum_c \tilde{v}_j^c$, and iii) iterating :

Step 1 : update correspondences (\tilde{v}_j^c) (after an optimal rigid alignment \tilde{T}^c)
compute \tilde{T}^c and \tilde{A}^c
$\tilde{v}_j^c = \sum_i \tilde{A}_{i,j} \tilde{T}^c(x_i^c)$
Step 2 : update \tilde{M} :
for each j , $\tilde{m}_j = \frac{1}{C} \sum_c \tilde{v}_j^c$

This simple algorithm describes an automatic way to compute a classical Procrustean mean shape. For a better understanding of this process, one can regard it as :

$$\tilde{M} = \arg \min_M \sum_c d(M, V^c(X, M))^2, \quad (11.2)$$

where M and $V^c(., .)$ are labelled point sets and $d(., .)$ is the Procrustean distance described in Eq. 11.1. Each labelled point set $V^c(X, M)$ ($c \in [1, \dots, N]$) is the result of a registration process of X^c on M . Then the previous algorithm can be described as i) the estimation of V^c knowing the population X^c and M and ii) the estimation of M knowing V^c .

As previously claimed, the previous algorithm is likely to fail to characterise complicated structures. Thus, we propose to improve the estimation of the correspondences v_j^c by introducing a non-linear deformation linking M and X^c into the criterion. This new term will improve the estimation of the matrices A^c and thus the subsequent estimation of the correspondences (v_j^c). However, to preserve the spirit of the Procrustean mean, the non-linear deformation has to affect only the estimation of the matrices A^c and not to change the nature of the shape space (that is defined as an orbit space under the action of a simple similarity transformation). This view results in the following algorithm :

Step 1 : update correspondences (\tilde{v}_i^c) (after an optimal rigid alignment \tilde{T}^c) :
 $\forall c$, compute \tilde{T}^c and \tilde{A}^c (using a non-linear registration)
 $\forall c, \forall j$, compute $v_j^c = \sum_i \tilde{A}_{i,j}^c \tilde{T}_{inv}^c(x_i^c)$

Step 2 : update \tilde{M} :
for each j , $\tilde{m}_j = \frac{1}{C} \sum_c \tilde{v}_j^c$

where T_{inv}^c is the rigid component of T^c . An illustration of this strategy is given in the second row of Figure 11.2.

This algorithm is designed as an extension of the MS-EM-ICP and we did not find any proper criterion it minimises. However, we observe that the value $\sum_c d(M, V^c(X, M))^2$ decreases throughout the iterations and that the obtained mean shape fits well the data X .

Conclusion :

When modeling T as a linear transformation (without penalisation), the two developed algorithms (Fréchet-like mean and automatic Procrustean mean) coincide and are equivalent with the algorithm proposed by Hufnagel and colleagues. When using non-linear transformations, the two algorithms describe two different points of view. As for now, we have not investigated the respective added values of both frameworks to the estimation of a statistical model but we plan to investigate their generalisation abilities (indicating how closely the shape model matches an unknown observation) and specificity (indicating how the obtained shape model describe shapes similar to the ones in the training set X).

In the following, we choose to use the automatic Procrustean mean to perform mean shape estimation and statistical shape analysis because of i) its simplicity of interpretation as opposed to the Fréchet-like mean (largely depending on L/α') and ii) the easiness of the underlying statistical analysis.

11.1.3 Virtual correspondence & projection

As previously mentioned, the match matrix allows to compute the “virtual” correspondent of each point of the mean shape in each shape of the population under study using the *a posteriori* probability of matchings encoded in $A_{i,j}^c$ ($\forall c$) : $v_j^c = \sum_i A_{i,j}^c T_{inv}^c x_i^c$. In the following, we assume that the shapes X^c are initially rigidly registered together and thus that $T_{inv}^c = Id$ and $v_j^c = \sum_i A_{i,j}^c x_i^c$.

Another interesting use of this *a posteriori* probability concerns the projection of pointwise individual mapping on the mean shape M . Consider that each point x_i^c of X^c contains a scalar information (*e.g.* cortical thickness, local asymmetry, curvature, ...) that we call $s_i^{X^c}$. Then, the projection of $S^{X^c} = (s_i^{X^c})$ on M is given by $s_j^{X^c \rightarrow M} = \sum_i A_{i,j}^c s_i^{X^c}$. As $\forall j, \sum_i A_{i,j}^c = 1$, the “interpolated” value at m_j is a weighted mean of the values at different points X . This last property will be useful for the projection of individual asymmetry maps on a reference template.

11.1.4 Comparison of populations

Once a mean shape M and the non-linear transformation best superposing M on each X^c are computed, several techniques allow the comparison between populations of shapes using either the different bending energies [LMS⁺08] or displacement fields [PLNS04], and t-tests or permutation tests.

11.2 Quantification of asymmetries over/between populations

In this section, based on the previously described tools, we develop a pipeline for the quantification of asymmetries over or between populations. We consider that all the structures can be and are oriented in a common orthogonal frame consisting of three axes that we call anterior-posterior, left-right and head-foot.

11.2.1 Computation of individual asymmetry maps

Let $X^c \in X$ be a point set representing an anatomical structure under study. Its *individual asymmetry maps* are computed as follows :

1. computation of the *approximate symmetry plane* P of X^c using the algorithms described in Part 1.
2. computation of the *asymmetry field* as the deformation field best superposing X^c and $S_P(X^c)$ using the algorithms described in Part 2.
3. computation of the 3 *individual asymmetry maps* by projecting each pointwise vector of the asymmetry field on the 3 coordinate axes. This allows to differentiate the anterior-posterior, left-right and head-foot components of the asymmetry field. This leads to 3 different scalar *asymmetry maps* for each point set X^c . We call these asymmetry maps $S^{X^c,AP}$, $S^{X^c,LR}$ and $S^{X^c,HF}$.

11.2.2 Computation of a mean cortical shape & projection of the asymmetry mappings

We consider C point sets X^1, \dots, X^C representing the C structures under study. Their individual asymmetry maps $S^{X^c,AP}$, $S^{X^c,LR}$ and $S^{X^c,HF}$ have to be put in a common geometry to be compared. For this purpose, we use the method described above to iteratively compute :

- the point set M representing the mean shape,
- the match matrices $\{A^1, \dots, A^C\}$ (describing the fuzzy point-to-point correspondences between the point sets X^c and M),
- the transformations T^c (best superposing the point sets X^c and M).

Once the mean point set M and the fuzzy match matrices A^C are computed, we project each individual (scalar) asymmetry maps $S^{X^c,AP}$, $S^{X^c,LR}$ and $S^{X^c,HF}$ on M which provides the *normalised individual asymmetry maps* $S^{X_i \rightarrow AP}$, $S^{X_i \rightarrow M,LR}$ and $S^{X_i \rightarrow M,HF}$. As pointed out in Section 11.1.3, the mapping at point $m_j \in M$ was defined as $s_j^{X^c \rightarrow M,AP} = \sum_i A_{ij}^c s_i^{X^c,AP}$ (and similarly for $s^{X^c \rightarrow M,LR}$ and $s^{X^c \rightarrow M,HF}$). The whole process of computing 3 asymmetry maps for one simplified cortex and projecting these maps on the mean shape is described in Fig. 11.3.

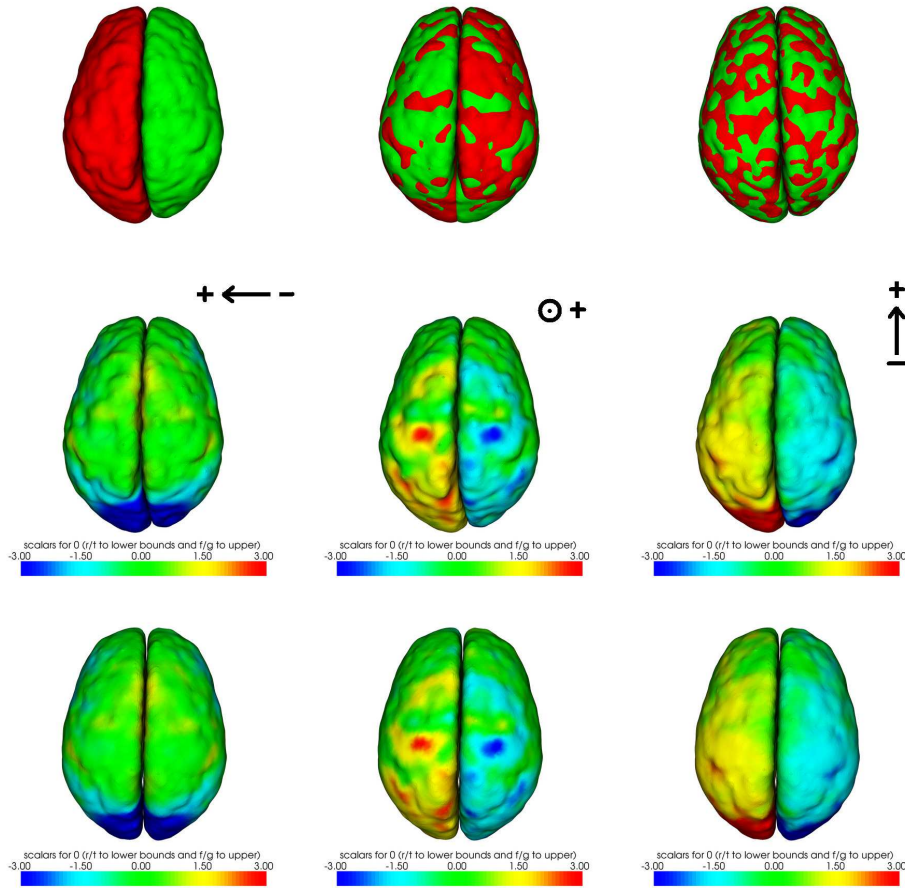


FIG. 11.3 – **Mapping asymmetries on one subject** : All the rows are described from left to right. First row : simplified point set X - superposition of $S_P(X)$ on X - superposition of $T \circ S_P(X)$ on X (the patchy appearance shows the residual errors after nonlinear registration of X and $S_P(X)$). Second row : left-right (width), head-foot and posterior-anterior (protrusion) components of the asymmetry field. Third row : projection of the 3 asymmetry maps of X (second row) on the mean point set M .

11.2.3 Computation of the symmetry surface

When dealing with cortical surfaces, we can analyse the deviations of the inter-hemispheric fissure by estimating the symmetry surfaces of the point sets (using the method developed in Section 4.1). As opposed to the mid-sagittal plane, which is only used as an intermediate result to compute the asymmetry field, this surface offers an insight into the configuration of the inter-hemispheric fissure, and can be analysed as such.

Chapitre 12

Asymmetries

In this chapter, we propose to illustrate our previously developed pipeline for the quantification of asymmetries over and between populations by performing a study comparing cortical asymmetries of Situs Inversus subjects with control subjects.

12.1 Note on asymmetry maps :

The interpretation of the directional asymmetry maps we defined in the previous chapter can be quite misleading. To alleviate some bad interpretations of the results proposed in this chapter let us first consider the two following points.




FIG. 12.1 – How to read asymmetry maps.

Ambiguous colour maps : Figures 12.1.a and 12.1.b show how to read asymmetry maps. Let X be an object in black, its approximate symmetry plane P and its symmetrical image $S_P(X)$ (in red), we compute the deformation superposing X on $S_P(X)$ (black arrows). The left-right component of the deformation field (Fig 12.1.a) points towards the same direction in the left and in the right side of the object. Consequently, the color map (indicating the left-right component of the asymmetry field) on the two

sides of the object will be roughly symmetrical. However, the deformation indicates that the left side extends more laterally than the right side.

For the components of the deformation field that are not orthogonal to P (*i.e.* antero-posterior Fig 12.1.b and head foot), the deformation field corresponding to analogous structures points towards opposite directions. In our example(Fig 12.1.b) , the top-left (where the deformation points towards the **positive** direction) and bottom-left (where the deformation points towards the **negative** direction) area both protrude more than their counterpart.

To sum up, one has to look carefully at the direction of the axis to interpret the results : hot (resp. cold) colors do not necessarily imply expansion (resp. shrinking).

Asymmetric maps : Figure 12.1.c illustrates why asymmetry maps are likely to look asymmetric. Let be an object X (in black) and its symmetrical image $S_P(X)$ (in red), registering X on $S_P(X)$ and registering $S_P(X)$ and X is similar and theoretically should lead to equivalent results. However, the origin of the vectors composing the displacement field are different as the object is not symmetrical and thus the resulting asymmetry maps have no reason to be symmetric. In other words, the computed information (*i.e.* the deformation field) is symmetric but its representation is not.

12.2 Illustrating our pipeline : study on Situs inversus

This work is the result of a collaboration with Neil Roberts and David N. Kennedy.

Subjects with situs inversus (SI) have the potential to provide unique clues into the developmental mechanisms underlying the torque and its relationship with hemispheric dominance for language. Situs inversus is a very rare condition (*i.e.* affecting 1 in 10,000 live births) where all the visceral organs are on the opposite side of the body to where they would be expected and as if they were reflected in a mirror. General health, however, is not generally affected and thus clinical imaging studies are very scarce [Woo86, CGH⁺93, KOT⁺99, IHF⁺10].

In short, Kennedy and colleagues and Ihara and colleagues [KOT⁺99, IHF⁺10] each studied a different three subjects with situs inversus (3 very strong right-handers for Kennedy data and 1 very weak right-handedness, 1 very weak left-hander and 1 very strong right-hander for Ihara data). Both groups report that brain torque is reversed in all three subjects with situs inversus and both groups suggest that leftward anatomical asymmetry of language structures *i.e.* planum temporale (2 out of 3 in [KOT⁺99] and 3 out of 3 in [IHF⁺10]) and IFG (3 out of 3 in [IHF⁺10]) is present in situs inversus, but differ in that Kennedy suggests left hemisphere functional dominance for language using fMRI [KOT⁺99] in all 3 subjects and Ihara suggests right hemisphere functional dominance for language in 2 out of 3 subjects using MEG [IHF⁺10]. Therefore, structurally there seems to be preserved local asymmetries of language-related structures in situs inversus but reversed torque and maybe also reversed functional dominance for language.

To our knowledge, the relationship between the brain torque (measured in terms of the protrusion and/or lateral extension of the frontal and occipital lobes) and de-

viations in the inter-hemispheric plane has not been previously studied. Our goal was to perform such a study using our tools on two right-handed male subjects with situs inversus in comparison with a group of eleven age-matched right-handed male control subjects. The two subjects with situs inversus have already been studied by Kennedy and colleagues [KOT⁺99] and we aim to replicate and extend their observations, by using a more homogeneous control population in terms of age, handedness and sex.

12.2.1 Subjects and image acquisition

Data were available for two right-handed male subjects with situs inversus [KOT⁺99]. These two subjects are denoted as subject 1170 and 1175 in the following. Both were aged 33 years. The control subjects are eleven right-handed male volunteers, ranging in age from 24 to 42 years (mean 28 ± 4.8). A $256 \times 256 \times 182$ 3D T1-weighted image of 1mm isotropic resolution was acquired for each control subject on a Philips Achieva 3T System (Philips Medical Systems, Best, The Netherlands). The series of images covered the whole brain and were acquired without any gap in the axial plane parallel to the anterior-posterior commissure line.

12.2.2 Preprocessing

For each subject’s MRI, the following pipeline was applied. We first segmented grey matter and separated each hemisphere (`surfer.nmr.mgh.harvard.edu`). The mesh was smoothed to remove inter-subject variability arising from the gyri and sulci. For this purpose, we applied the following process on each mask (each hemisphere separately). We applied a closing (*i.e.* a dilation followed by an erosion) of the corresponding binary mask and repeated this process k times. In practice, the larger k , the “smoother” the cortex. Then for each such produced mask, we extracted a mesh using marching cubes (`vtk.org`). Using $k = 10$, we obtained visually satisfying results, giving meshes of about 150k points (Fig. 11.3, first row).

12.2.3 Mean, standard deviation and significant asymmetries in controls

We had a collection of n ($\times 3$) scalar normalised individual asymmetry maps $S^{X_1 \rightarrow M}, \dots, S^{X_n \rightarrow M}$ projected on a common mean mesh M . For each point of M and each of the 3 components of the asymmetry, we computed the mean asymmetry and its standard deviation over the 11 subjects and we performed a pointwise t-test with the null hypothesis H_0 : “There is a perfect symmetry”. We corrected the obtained p -values for multiple comparisons by i) fixing a suprathreshold $\alpha = 0.05$, ii) building the permutation distribution of the maximal suprathreshold cluster size (by shuffling the p -values over the mesh and computing the size of the largest cluster of points with a p -value greater than α , using 20,000 permutations) iii) setting the critical suprathreshold cluster size as the $[\alpha \times 20,000] + 1^{\text{th}}$ largest value over the sampling distribution and iv) removing the clusters having a size smaller than the critical suprathreshold cluster size from the statistical map [NH02]. The results are displayed on Figures 12.2 and 12.3.




FIG. 12.2 – Mean for each of the 3 components of the asymmetry field (11 subjects). Each row shows one of the components (from top to bottom : left-right, head-foot, posterior-anterior) with different views.

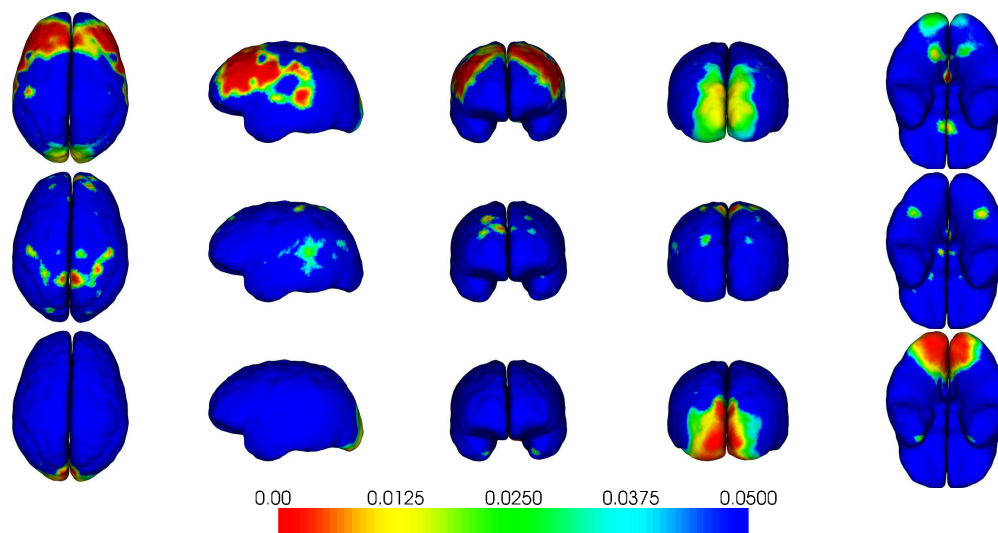


FIG. 12.3 – p-value maps for each of the 3 components of the asymmetry field (11 subjects). Each row shows one of the components (from top to bottom : left-right, head-foot, posterior-anterior) with different views. The null hypothesis is that there is no asymmetry (t-test, corrected for multiple comparisons using suprathreshold cluster size tests [NH02]).

12.2.4 Computation of regional asymmetries

For an easier interpretation of the results, we integrated the pointwise asymmetry values over the frontal and occipital lobes of each of the subjects. Both lobes were extracted following affine registration of the LONI LPBS40/SPM5 atlas [SMA⁺08] to each subject of each of the subjects. Then for each subject, we computed the average left-right ("width") and the posterior-anterior ("protrusion") components for both the frontal (F) and occipital (O) lobes. To be even more synthetic, for each subject we noted the lobe (L or R, *i.e.* left or right) that appears *wider* and that *protrudes more* than its counterpart. When the difference in width or protrusion between the hemispheres was low ($< 10^{-1}$ mm), we noted 0 instead of L or R. For both the width (L-R), protrusion (P-A) components of the asymmetry field over the frontal and occipital lobes), we performed a t-test with the null hypothesis H_0 : "There is a perfect symmetry". This integrated analysis was performed both on controls and SI subjects. In the same way, we averaged the mean curvature ("deviation") of the symmetry surface for both lobes. L (resp. R) indicates that the left (resp. right) hemisphere deviates towards the right (resp. left). For both lobes, we performed a t-test with the null hypothesis H_0 : "There is no deviation of the inter-hemispheric fissure". This integrated analysis was performed both on controls and SI subjects. The results are displayed on Tab. 12.1.

	1	2	3	4	5	6	7	8	9	10	11	p
O protrusion	L	L	L	L	R	L	L	L	L	L	0	p<0.05
O width	L	L	0	L	R	L	L	L	L	L	L	p<0.05
O curvature	L	L	0	L	R	L	L	L	L	L	L	p<0.05
F protrusion	R	R	R	R	L	R	L	R	L	L	0	p>0.05
F width	R	L	0	R	R	0	L	0	L	R	L	p>0.05
F curvature	R	R	R	0	R	0	0	R	R	R	0	p>0.05

TAB. 12.1 – **Protrusion/width interhemispheric differences and curvature of the symmetry surface.** We display these differences for the 11 subjects for occipital (O) and frontal (F) lobe. More details in the core of the text.

12.2.5 Individual asymmetry on SI subjects

The asymmetry maps are shown in Figure 12.4 and Figure 12.5 and the integrated values in Table 12.3.2.

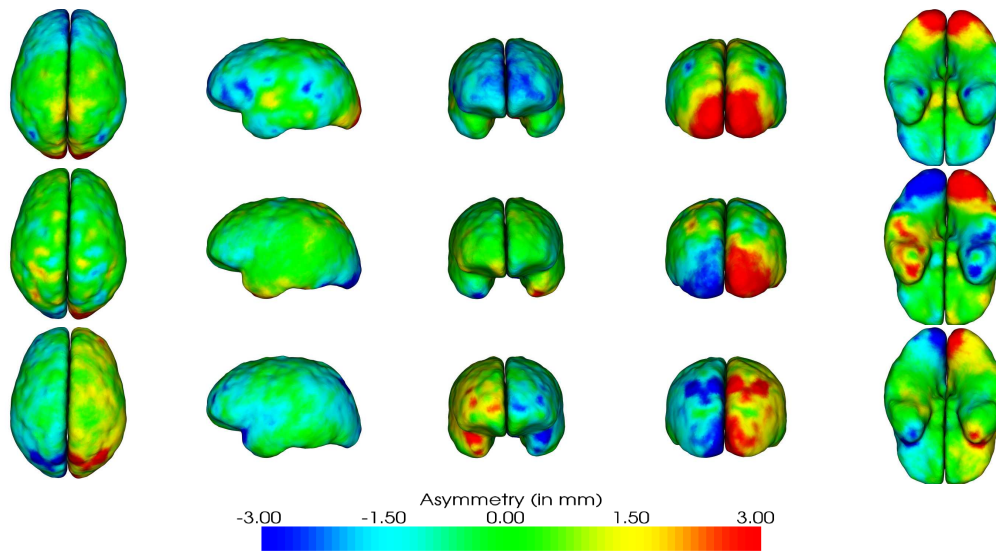


FIG. 12.4 – Asymmetry map for each of the 3 components of 1170.

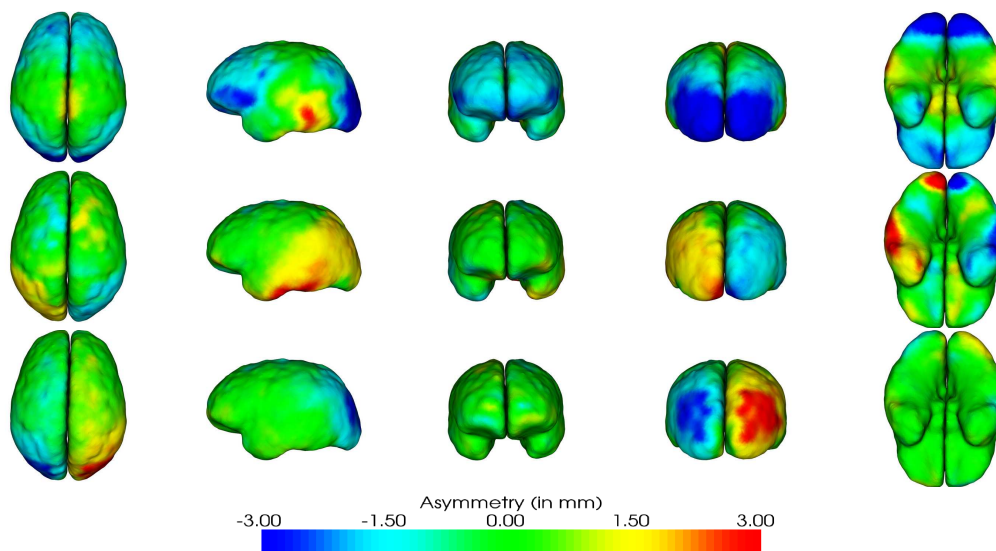


FIG. 12.5 – Asymmetry map for each of the 3 components of 1175.

-	1170	1175	p
O protrusion	R	R	$p < 0.05$
O width	L	R	$p > 0.05$
O curvature	L	R	$p > 0.05$
F protrusion	L	L	$p < 0.05$
F width	L	L	$p < 0.05$
F curvature	L	L	$p < 0.05$

TAB. 12.2 – Integrated values for the two SI subjects and p-value (t-test) with H_0 : SI subjects belong to the control population.

12.2.6 Significant differences of asymmetry between controls and SI subjects

We performed two tests to compare the control and the SI populations. First, we simply proposed to compute a pointwise p -value using a t-test comparing independently each SI subject with the control population (H_0 : “The given SI subject belongs to the control population”). The p -values were then corrected for multiple comparisons using the same technique as in Section 12.2.3. The computed p -values are displayed in Figures 12.6 and 12.7.

Then we performed a pointwise permutation test (H_0 : both SI subjects belong to the control population) using the mean difference as a statistic. The p -value at each point is simply the proportion of the permutation distribution of the statistic greater than or equal to the observed statistic [Goo00]. The p -values were then corrected for multiple comparisons using the same technique as in 12.2.3. The computed p -values are displayed in Figures 12.8.

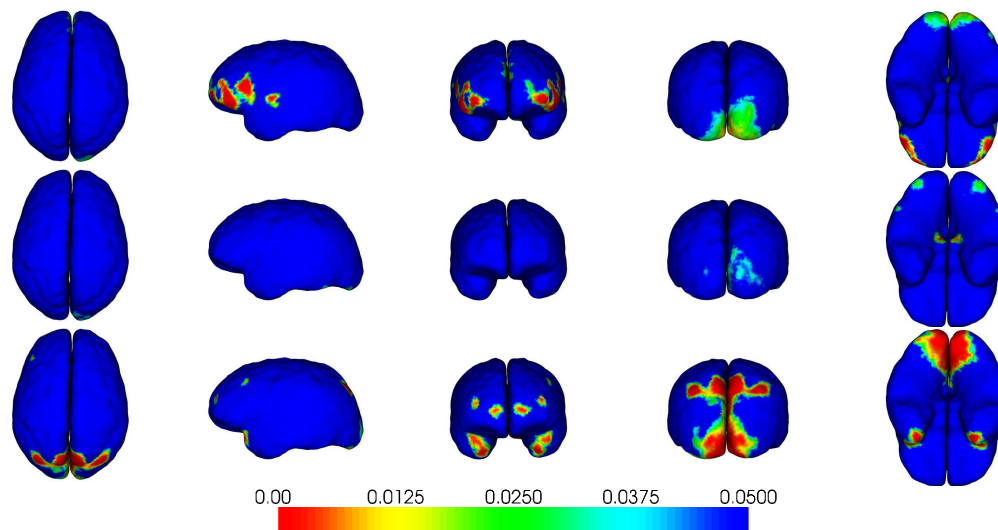


FIG. 12.6 – p -value of (corrected) t-test with H_0 : 1170 belongs to the control population.

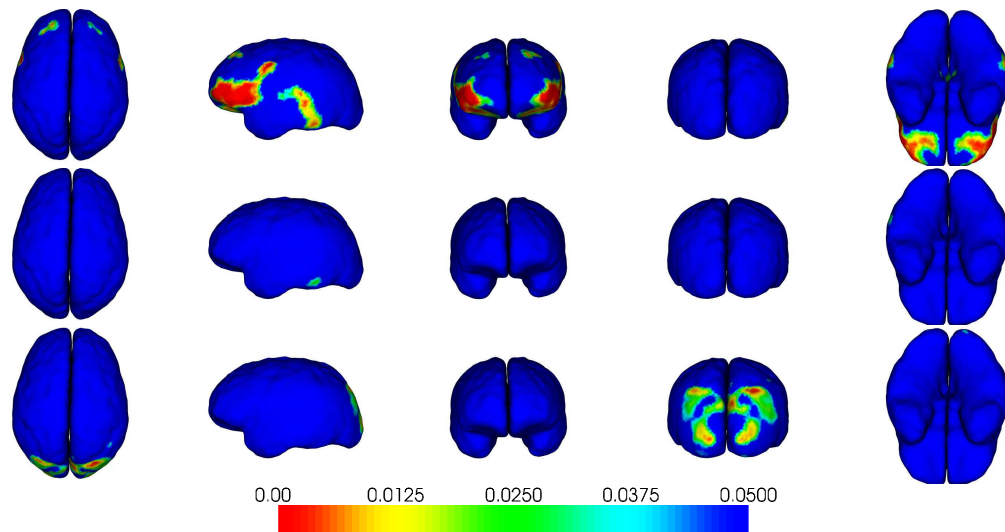


FIG. 12.7 – p-value of (corrected) t-test with H_0 : 1175 belongs to the control population.

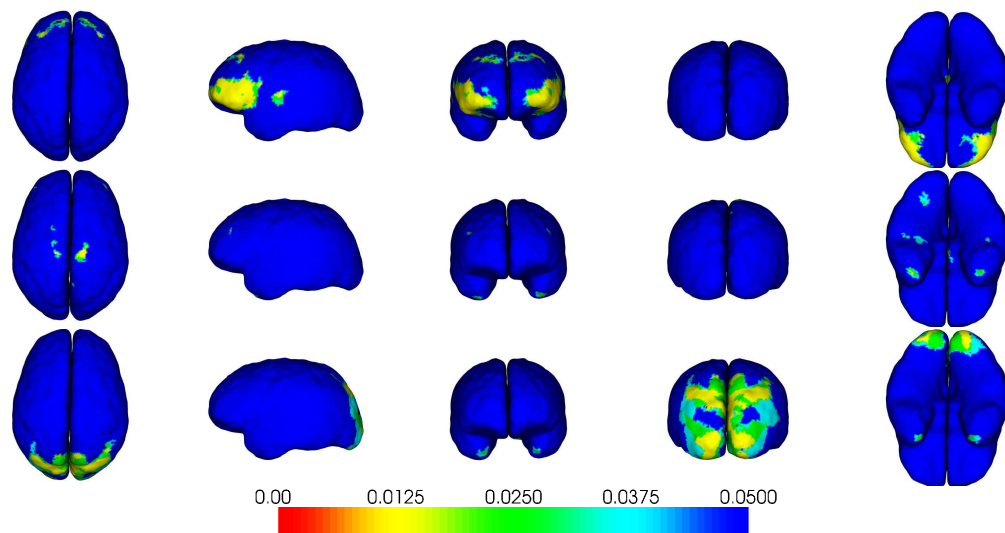


FIG. 12.8 – p-value of permutation test with H_0 : SI population and control population are identical.

12.3 Results

12.3.1 Controls

Occipital lobes on controls : At the population level, and at the chosen significance level ($p=0.05$, corrected for multiple comparisons) the left occipital lobe appears to be wider (or bending towards the other side) and protrudes more posteriorly than the right occipital lobe (Fig. 12.2, first and third rows). These results are corroborated by the averaging of the asymmetry values (width, protrusion, bending) over the occipital lobes, where only one of the subjects does not seem to fit the general pattern (Table 12.1).

Frontal lobes on controls : The tip of the frontal lobes do not appear as significantly asymmetrical at the chosen level ($p=0.05$) for any of the 3 components. This could be due to the high variability in frontal lobes, which yields low t statistics and thus high p -values. However, when we average the 3 asymmetry components in the frontal lobes and look at the results at an individual level (Table 12.1), we see that the right frontal lobe is wider than the left in only 4 out of 11 subjects and protrudes more anteriorly in only 6 out of 11 subjects, which corroborates the results of the statistical maps. However, the fissure seems to bend towards the right in 7 out of 11 subjects, and never seems to bend towards the left.

Parietal lobes on controls : At last and interestingly, we also find an area in the parietal lobe with a small significant "vertical" (head-foot) asymmetry : the right side seems to be "higher" compared to its counterpart (Fig. 12.2, second row).

Summary : Overall, these results do not support the existence of a counterclockwise torque at the population level, although the left occipital lobe seems to be significantly wider than the opposite side, while bending towards it and protruding more posteriorly. Intriguingly, we recover a significant left-right asymmetry on top of the frontal lobes (Fig. 12.2, first row), in the same direction as in the occipital lobes, that is, it seems bigger on the left side. This area is located close to the Broca's area.

12.3.2 Situs inversus

Right occipital lobes protrude more anteriorly than the left in the two subjects, while it is the opposite in the frontal lobes. The left frontal lobes are also wider than the right counterpart and bend towards it. On the contrary, width and curvature measurements are different between the two subjects in the occipital lobes (Table). This is corroborated by Fig. 12.4 and 12.5, where one can also see that the head-foot components are opposite in the occipital lobes between the two subjects (Table , Figures 12.4 and 12.5).

Comparison : The two SI subjects appear significantly different from the control subjects in the frontal lobes for the left-right component and in the occipital lobes

for the antero-posterior component (Fig 12.8). Interestingly, both SI and controls show significant, but opposite, protrusion asymmetries in the occipital lobes. This observation is our most striking finding and the occipital protrusion appears to be a significant discriminative pattern between the two populations. This results are preliminary and more interpretations are needed to have a real understanding of our findings. More subjects are to be included too.

Chapitre 13

Conclusion

13.1 Contributions

The design of tools for the quantification of asymmetries over and between populations of anatomical structures represented by point sets was an important goal of this PhD work and in this sense, our methodological choices find new justifications. Indeed, they lead to particularly simple and powerful solutions for group-wise studies. In particular, the initial probabilistic consideration allows to define the mean shape of a given population and the projection of individual point-wise information on the mean shape in a simple way. This application is particularly related to our probabilistic framework. Moreover, one shows that the method to compute a symmetry surface that we design in Part I provides an alternative suitable description of the departure from a perfect symmetry. We also illustrated our methodology on a real application.

13.2 Discussions

In Section 13.2.1 we address some questions and remarks concerning the proposed MS-EM-ICP algorithms and in Section 13.2.2 we address questions and remarks concerning our asymmetry analysis pipeline.

13.2.1 Mean shape EM-ICP strategies

It is clear that a deeper analysis and evaluation of the two proposed methods (Fréchet-like mean and Procrustean-like mean) has to be performed. However, the first results are very encouraging. Moreover, notice that the results obtained by the original MS-EM-ICP has been validated in a series of papers [HPE⁺08b, HPE⁺08a, HPE⁺07].

At that time, it is difficult to assess the relative added value of each one of the improvements we proposed and to determine which one will be better suited for a given application. However, one can notice that the mean (and thus the underlying statistical model) computed from the Fréchet-like mean algorithm strongly depends on the design of the regulariser L . Generally speaking, there exists no model to characterise

the deformation of a given organ and the design of L can be very problematic. By contrast, the Procrustean-like mean does not depend directly on L but on the quality of the correspondences obtained during the registration process that we think to be more independent of the design of L . For this reason, this last method is less dependent on ad hoc, and sometimes quite arbitrary, choices.

13.2.2 Pipeline for asymmetry quantification

13.2.2.1 Alternative solutions

We chose to adopt a quite natural strategy for the mapping of asymmetries on population of shapes. This choice is debatable and several other choices are conceivable. In the following, we list existing or possible alternative approaches and discuss them.

a) Joint symmetry plane and asymmetry field computation :

In our pipeline, we considered the estimation of the approximate symmetry plane and the estimation of asymmetry maps as 2 successive independent problems. However, this view can lead to contradiction. As an example, nothing can avoid that some points considered as corresponding to a highly symmetrical part of the shape under study during the estimation of the approximate symmetry plane are then considered as highly asymmetric (i.e. the points are associated to large displacement vectors) during the estimation of the asymmetry field. We propose to tackle this flaw using a criterion of the form :

$$\tilde{P} = \arg \min_{P,T} \sum_{i,j} A_{i,j} \|x_i - T(S_P(x_j))\|^2 + \alpha' L(T) + 2\sigma^2 \sum_{i,j} A_{i,j} \log(A_{i,j}),$$

with this solution, one estimates both the asymmetry field and the symmetry plane in a single task. Moreover, if one considers α' small (in this case, the residuals will tend to be null), one can interpret \tilde{P} as the plane that minimises the deformation between X and its symmetrical image. This provides an alternative definition of the symmetry plane. We tried to design such a solution but the minimisation leads to poor estimate of both the symmetry plane and the asymmetry field.

b) Asymmetry quantification as a registration problem :

Recently, Ólafsdóttir *et al.* [ÓLD⁺07] proposed to quantify asymmetries of 3D surfaces by registering a template, perfectly symmetrical surface, to the surface of interest in a non-linear way. The difference between the two displacement vectors needed to map bilateral points of the template with points of the surface is used to quantify the local asymmetry at these points. This approach is quite nice as it allows to compute in a single task both the asymmetry field and its projection to a template. However, such an approach can lead to substantial errors as small errors during the registration process are likely to importantly bias the asymmetry quantification. Moreover, this technique is not based on a well grounded definition of the asymmetry and appears as quite tricky.

As a result, the obtained asymmetry maps might be quite difficult to interpret. An illustration of this flaw is given in Figure 13.1.




FIG. 13.1 – From left to right : a deformation field is applied to a perfectly symmetrical face ; estimated asymmetry map by our strategy ; normalised asymmetry map on the template ; template-based asymmetry map as estimated by the method of Ólafsdóttir [ÓLD⁺07]. Both strategies are implemented using the same registration algorithm.

c) Asymmetry quantification as a classical shape analysis problem :

Asymmetry analysis is a shape analysis problem and could be tackled as such. If most of the subjects of a population exhibit a significant common asymmetry pattern, then it is likely that the mean shape will also exhibit it and that their asymmetries can be visualised via a classical analysis (*e.g.* modes of variations). However to our opinion, two main flaws limit this approach. First, the statistical shape modelling is a more difficult problem than the one consisting in quantifying asymmetries over a population. This is mainly due to the fact that it is generally more difficult to explain the variability between two different structures than the variability between a structure and its symmetrical image. Secondly, shape analysis does not focus on asymmetry and it is likely that the visualisation and the interpretation of the asymmetry would be hidden by the numerous variations among the population.

Chapitre 14

General conclusion

In this work, we presented a pipeline allowing the quantification of asymmetries between and within populations of anatomical structures. It is composed of a deformation-based approach and of the computation and the study of symmetry surfaces. An application of this pipeline is illustrated in the study comparing situs inversus and control subjects (Chapter 12). This pipeline relies on a set of methods that have been exposed independently in this work. However, all these methods have been developed in a unified framework relying on a MAP definition of the problem (optionally modified and improved) :

Let $F : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be a transformation best superposing (in a very general sense) two point sets X and Y . If we consider the points of X as the means of a Gaussian mixture model (GMM), the points in Y as independent samples of this GMM, and $\exp(-\beta L(F))$ a prior on F , then the MAP estimate of F is :

$$\tilde{F} = \arg \max_F \left[\prod_{y_j \in Y} \sum_{x_k \in X} \pi_{jk} p_k(y_j; F) \right] \exp(-\beta L(F))$$

where the $p_k(\cdot; T) = \mathcal{N}(F(x_k), \sigma^2 I)$ are Gaussian pdfs and the π_{jk} s ($\forall j, k, 0 < \pi_{jk} < 1$ and $\forall j, \sum_k \pi_{jk} = 1$) are the mixture proportions. In essence, π_{jk} conveys the probability that the point y_j in Y is matched with the point x_k in X without knowing anything else.

The methodological contributions and perspectives of this PhD were mentioned in the conclusion chapters finishing each part of this document. To our opinion, our main contributions have consisted in providing pragmatic and robust tools for the non-linear registration of large point sets and the accurate quantification and comparison of asymmetries over and between populations. These needs does not consist of an arbitrary choice but corresponded to a real demand met since the beginning of my PhD. The several ongoing studies (that are listed below) with various scientists going from paleoanthropologists to psychiatrists illustrate this aspect.

14.1 Ongoing studies on brain asymmetry

SI vs control cortical asymmetries (in collaboration with Neil Roberts and David N. Kennedy) :

Subjects with situs inversus (SI) have the potential to provide unique clues into the developmental mechanisms underlying the torque and its relationship with hemispheric dominance for language. We presented first results concerning this study in Chapter 12.

Bonobo vs Humans cortical asymmetries (in collaboration with Neil Roberts) :

Typical patterns of anatomical/functional asymmetries have long been thought to be typical of human beings (*Homo sapiens*). There has been much debate recently on this puzzling question, and some have claimed that these asymmetry patterns are actually shared by other great apes. Answering this key question could have important implications in the understanding of language origins and in what makes the humans unique. In this study, we quantify and compare the global asymmetry of 10 humans and 10 chimpanzees (acquired on the same MR system).

Brain-endocasts comparison (in collaboration with Marc Fournier, Neil Roberts and José Braga) :

Endocranial casts are commonly used to infer the shape, size, asymmetry or overall organization of the brain of fossil species. However, the inner surface of the skull is separated from the brain by the meninges, and thus endocasts are only an indirect representation of the brain. Few studies seek to assess the faithfulness of this representation, upon which is based much of our knowledge on the evolution of the brain of hominids. The aim of this work is to automatically map in 3D the distance between the brain and the inner surface of the skull in humans and chimpanzees to see how some morphological characteristics of the endocast (and in particular its asymmetries) relate to those of the brain in these two populations. MR images of humans and chimpanzees are used to automatically extract their brain and virtual endocast. Then point-to-point distances are evaluated between the two surfaces. An automatic computation of a mean endocranial shape within each population provides an average distance map whose spatial variation is investigated. Finally bilateral asymmetries of the brain and endocast are also computed, mapped and averaged on the mean endocranial shape of each population. Results show that the local distance between the brain and the virtual endocast strongly varies between anatomical regions, but symmetrically with respect to the mid-sagittal plane. It follows that global asymmetries of the brain and of the endocast are very similar in these two populations. First results are given in Figure 14.1.

Chimpanzees Vs Bonobos Vs ancestors endocranial asymmetries (in collaboration with Francis Thackeray, José Braga, Antoine Balzeau and Emmanuel Gillissen) :

The study of endocranial asymmetries of hominids is a central topic in paleoneurology. However, our knowledge about the emergence of these asymmetries during human evolution is still limited. This is partly due to the fact that, so far, the 3D asymme-




FIG. 14.1 – Distance color map pattern between the endocast and brain hull shown on the average endocast mesh. (a) Top view. (b) Bottom view. (c) Front view. (d) Right side view. (e) Rear view. (f) Left side view. The linear color map scale is in millimeters and the distance values are shown on the color bar.

tries have been mostly analysed using landmarks and semi-landmarks based methods. Such methods are limited as they only provide a partial description of the anatomy and thus of the possible asymmetries. One key problem before assessing the evolution of patterns of asymmetry in hominids is the identification of confounding factors such as intra-specific variability as well as age and sex influences. We study virtual endocasts of 60 *Pan paniscus* and 59 *Pan troglodytes* of different dental age and sex. 3D statistical analyzes are led to assess significantly asymmetrical areas on the endocasts within each population, and a comparison is made between the two populations. Several fossil hominin endocasts (such as Australopithecine species and fossil *Homo* species) will be then reexamined in light of the previously estimated variability with these two extant *Pan* species.

We display here some preliminary results on famous fossiles :

- Cro-Magnon¹ endocranial surface was segmented using AtrCore from a CT scan and contains about 650k points (Fig.14.2).
- Mrs Ples² endocranial surface was segmented from a CT scan $515 \times 512 \times 998$ of resolution $0.348 \times 0.348 \times 0.2$ mm using Amira and contains about 159k (Fig.14.3).

14.2 Other ongoing applications

Osseous labyrinths (in collaboration with José Braga and Didier Descouens) :

The bony labyrinth consists of three parts (the two vestibular sacs, the three semicircular canals and the cochlea) and houses two functional systems. The vestibular system provides one way of motion detection in a three-dimensional space. The cochlea is specialised for sound detection. These two functional systems represent relatively autonomous modules. The close anatomical relationship between the bony labyrinth and the corresponding receptor endorgans provides an opportunity to study osteological spe-

¹ *Homo sapiens* dated from about 30 thousands of years discovered at Les Eyzies (France) in 1868 by Louis Lartet

² *Australopithecus africanus* dated from 2.6 - 2.8 million of years and discovered at Sterkfontein (South Africa) in 1947 by Robert Broom and John T. Robinson

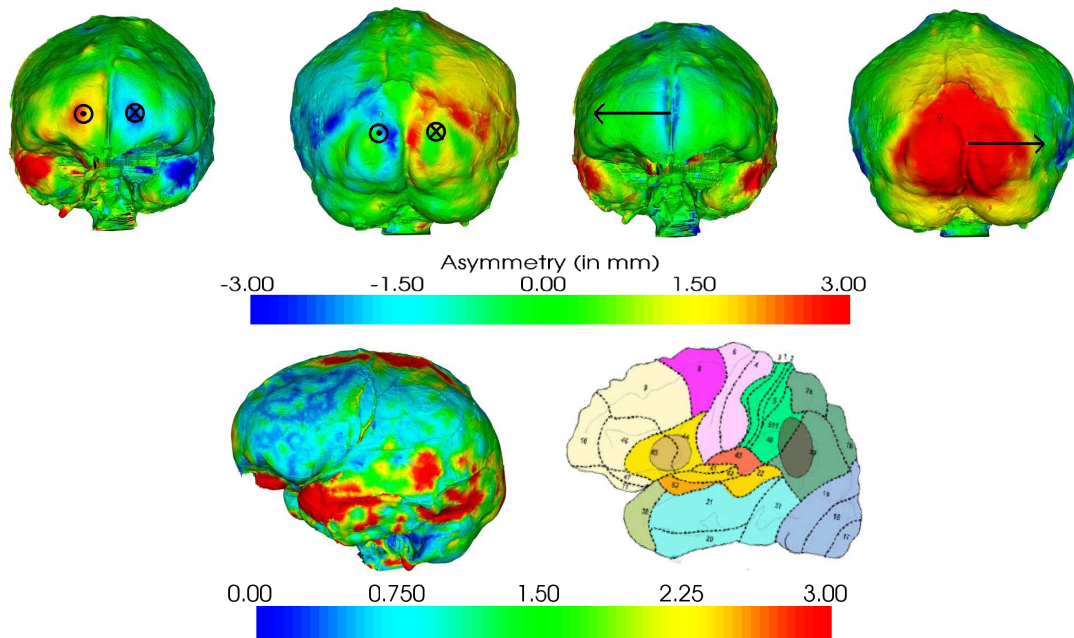


FIG. 14.2 – Global torque pattern and local asymmetries on CM1.

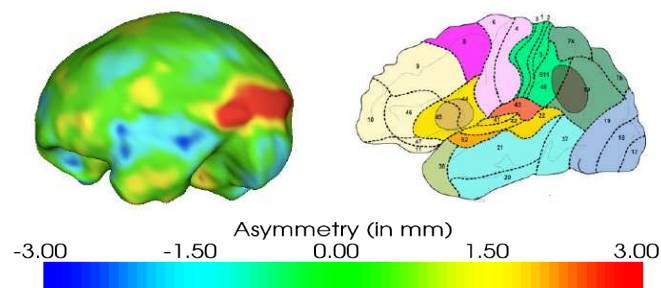


FIG. 14.3 – Asymmetry map on Mrs Ples. Areas 17-18-19-39-40, the left side protrudes more than the right. Areas 41 to 45, the right side protrudes more than the left

cimens (including fossils). So far, this has been done mainly by using linear, surface and volume measurements separately for the semicircular canals and the cochlea. This study is motivated by the fact that the investigation of the three-dimensional (3D) anatomical variation of the bony labyrinth in extant species represents a prerequisite for the interpretation of their fossil closest relatives. This prerequisite has not been fulfilled yet due to (i) the complexity of the labyrinth 3D geometry; (ii) the lack of high resolution data; (iii) the lack of appropriate and expert-independent comparative methods. We propose to use 3D geometrical models of 40 bony labyrinths reconstructed from micro-CT scans of extant humans, chimpanzees (both common and pygmy) and baboons to investigate some new morphological features which could be used to assess some motion or hearing

parameters of fossil taxa. Even if these assessments remain speculative, they allow us to investigate the co-evolution of the two functional systems of the inner ear in extant or extinct higher primates. First results are given in Figure 14.4.




FIG. 14.4 – Mean shape and first mode of variation ($\pm 2 \sqrt{\lambda_1}$) and on 10 osseous labyrinths : the first mode can be interpreted as a size change of the canals.

Neuronavigation (in collaboration with Pierre Hellier and Charles Garraud) :

An engineer is currently integrating our non linear-registration methods into a neuronavigation system. As possible patents are considered, we cannot give more details about this work.

14.3 List of publications

14.3.1 International conferences

1. **Benoît Combès**, Sylvain Prima. An efficient EM-ICP algorithm for symmetric consistent non-linear registration of point sets. In 13th International Conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI'2010, Tianzi Jiang, A. Colchester, J. Duncan (eds.), Lecture Notes in Computer Science, Pékin, Chine, Septembre 2010.

2. **Benoît Combès**, José Braga, Francis Thackeray, Sylvain Prima. 3D automated quantification of asymmetries on fossil endocasts. In 79th Annual Meeting of the American Association of Physical Anthropologists (AAPA), Albuquerque, État-Unis, Avril 2010.

3. Daniel Münch, **Benoît Combès**, Sylvain Prima. A modified ICP algorithm for normal-guided surface registration. In Proceedings of SPIE Medical Imaging 2010 : Image Processing, Benoit M. Dawant, David R. Haynor (eds.), San Diego, État-Unis, Février 2010.

4. **Benoît Combès**, Sylvain Prima. Setting priors and enforcing constraints on matches for nonlinear registration of meshes. In 12th International Conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI'2009, G.-Z. Yang, D.J. Hawkes, D. Rueckert, J.A. Noble, C.J. Taylor (eds.), Lecture Notes in Computer Science, Pages 175-183, Londres, Royaume-Uni, Septembre 2009.

5. **Benoît Combès**, Sylvain Prima. Prior affinity measures on matches for ICP-like

nonlinear registration of free-form surfaces. In 6th IEEE International Symposium on Biomedical Imaging : From Nano to Macro, ISBI'2009, Pages 370-373, Boston, États-Unis, Juin 2009.

6. **Benoît Combès**, Sylvain Prima. New algorithms to map asymmetries of 3D surfaces. In 11th International Conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI'2008, D.N. Metaxas, L. Axel, G. Fichtinger, G. Székely (eds.), Lecture Notes in Computer Science, Pages 17-25, New York, États-Unis, Septembre 2008.

7. **Benoît Combès**, Robin Hennessy, John Waddington, Neil Roberts, Sylvain Prima. Automatic symmetry plane estimation of bilateral objects in point clouds. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR'2008, Anchorage, États-Unis, Juin 2008.

8. **Benoît Combès**, Robin Hennessy, John Waddington, Neil Roberts, Sylvain Prima. An algorithm to map asymmetries of bilateral objects in point clouds. In 5th IEEE International Symposium on Biomedical Imaging : From Nano to Macro, ISBI'2008, Pages 1139-1142, Paris, France, Mai 2008.

14.3.2 International conferences : under submission

Benoît Combès, Marc Fournier, José Braga, Gérard subsol, Antoine Balzeau, Emmanuel Gilissen, Francis Thackeray, Sylvain Prima Automatic, landmark-free quantification of 3D endocranial asymmetries in extant and fossil species : new insights into paleoneurology. In 80th Annual Meeting of the American Association of Physical Anthropologists (AAPA), In Submission

José Braga, Marc Fournier, **Benoît Combès**, Didier Descouens, Gérard Subsol, Sylvain Prima. Evolutionary perspectives of high resolution three-dimensional comparisons of bony labyrinths in humans, chimpanzees and baboons. In 80th Annual Meeting of the American Association of Physical Anthropologists (AAPA), In Submission

Marc Fournier, **Benoît Combès**, José Braga, Neil Roberts, William Hopkins and Sylvain Prima Mapping the distance between the brain and the inner surface of the skull in humans and chimpanzees In 80th Annual Meeting of the American Association of Physical Anthropologists (AAPA), In Submission

Quatrième partie

Appendix

Annexe A

Closed-form solution for symmetry plane estimation

A.0.3 Principal axes

A.0.3.1 Fitting a plane in point set

$$\begin{aligned}\mathcal{E}_1(P) &= \sum_x d(x_i, P) = \sum_i |n^T x_i - d|^2 \\ \frac{\partial \mathcal{E}}{\partial d} &= -2 \sum (n^T x - d) \\ \frac{\partial \mathcal{E}}{\partial d} = 0 &\Rightarrow \sum n^T x = \text{card}(X)d \\ \frac{\partial \mathcal{E}}{\partial d} = 0 &\Rightarrow d = n^T x_G\end{aligned}$$

Setting this new expression for d we have at the optimum :

$$\mathcal{E}_1(P) = \sum_i (n^T x_i - d)^2 = \sum_i (n^T x_i - n^T x_G)^2 = \sum_i (n^T (x_i - x_G))^2 = n^T \sum_i ((x_i - x_G)(x_i - x_G)^T) n$$

As a result optimal n are among the eigenvectors of matrix $\sum_i ((x_i - x_G)(x_i - x_G)^T)$ (called the scatter matrix).

A.0.3.2 Scatter matrix and inertia matrix

The scatter matrix S is linked with the inertia matrix J by the expression :

$$S = J - \left(\sum_i x_i^2 + y_i^2 + z_i^2 \right) . I_3,$$

Let \tilde{v} and $\tilde{\lambda}$ be respectively an eigenvector and an eigenvalue of J . Then

$$(J - \left(\sum_i x_i^2 + y_i^2 + z_i^2 \right) . I_3) \tilde{v} = J . \tilde{v} - \left(\sum_i x_i^2 + y_i^2 + z_i^2 \right) \tilde{v}$$

$$= (\tilde{\lambda} + \sum_i (x_i^2 + y_i^2 + z_i^2))v.$$

As a result, \tilde{v} is an eigenvector of $\sum_i ((x_i - x_g)(x_i - x_g)^T)$ with eigenvalue $\tilde{\lambda} + \sum_i (x_i^2 + y_i^2 + z_i^2)$.

A.0.4 Optimal solution for symmetry plane estimation with known matchings

Considering (1), the minimisation of (2) with respect to P can be written :

$$\tilde{P} = (\tilde{d}, \tilde{n}) = \arg \min_{(d, n)} \sum_i w_i \|y_i - (I_3 - 2nn^T)x_i - 2dn\|^2 \text{ with the constraint } \|n\| = 1$$

The Lagrangian $\mathcal{L}(d, n, \lambda)$ associated to the previous constrained minimisation problem is :

$$\mathcal{L}(d, n, \lambda) = \left(\sum_i w_i \|y_i - (I_3 - 2nn^T)x_i - 2dn\|^2 \right) - \lambda (\|n\|^2 - 1)$$

For the optimal \tilde{d} , one necessarily has $\frac{\partial \mathcal{L}}{\partial d} = 0$, which writes :

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial d} = 0 &\Leftrightarrow \sum_i w_i [-2n^T (y_i - (I_3 - 2nn^T)x_i - 2dn)] = 0 \\ &\Leftrightarrow \sum_i w_i [n^T (y_i + x_i) - 2d] = 0 \\ &\Leftrightarrow d = \frac{1}{2 \sum_i w_i} \sum_i w_i (x_i + y_i)^T n \\ &\Leftrightarrow d = \frac{1}{2} (x_g + y_g)^T n \end{aligned}$$

noting

$$x_g = \frac{1}{\sum_i w_i} \sum_i w_i x_i \text{ and } y_g = \frac{1}{\sum_i w_i} \sum_i w_i y_i.$$

Substituting this expression of d in the Lagrangian, we obtain :

$$\begin{aligned} \mathcal{L}(d, n, \lambda) &= \left[\sum_i w_i \| (I_3 - 2nn^T)x_i + n^T (x_g + y_g)n - y_i \|^2 \right] - \lambda (\|n\|^2 - 1) \\ &= \left[\sum_i w_i \| (I_3 - nn^T)x_i - nn^T x_i + nn^T x_g + nn^T y_g - nn^T y_i + nn^T y_i - y_i \|^2 \right] \\ &\quad - \lambda (\|n\|^2 - 1) \end{aligned}$$

because $\forall(u, v) \in \mathbb{R}^3 \times \mathbb{R}^3$, $v^T u v = v v^T u$; then we can rewrite :

$$\mathcal{L}(d, n, \lambda) = \left[\sum_i w_i \| (I_3 - nn^T) (x_i - y_i) - nn^T (x_i - x_g + y_i - y_g) \|^2 \right] - \lambda(\|n\|^2 - 1)$$

By developing the sum, we get :

$$\begin{aligned} \mathcal{L}(d, n, \lambda) &= \left[\sum_i w_i \| (I_3 - nn^T) (x_i - y_i) \|^2 \right] + \left[\sum_i w_i \| nn^T (x_i - x_g + y_i - y_g) \|^2 \right] \\ &\quad - \left[2 \sum_i w_i (x_i - y_i)^T (I_3 - nn^T) nn^T (x_i - x_g + y_i - y_g) \right] - \lambda(\|n\|^2 - 1) \end{aligned}$$

Considering that $(I_3 - nn^T) nn^T = 0$, we get :

$$\mathcal{L}(d, n, \lambda) = \left[\sum_i w_i \| (I_3 - nn^T) (x_i - y_i) \|^2 \right] + \left[\sum_i w_i \| nn^T (x_i - x_g + y_i - y_g) \|^2 \right] + 0 - \lambda(\|n\|^2 - 1)$$

By developing the two sums, we get :

$$\begin{aligned} \mathcal{L}(d, n, \lambda) &= \left[\sum_i w_i (x_i - y_i)^T (I_3 - nn^T) (I_3 - nn^T) (x_i - y_i) \right] \\ &\quad + n^T \left[\sum_i w_i (x_i - x_g + y_i - y_g) (x_i - x_g + y_i - y_g)^T \right] n - \lambda(\|n\|^2 - 1) \end{aligned}$$

considering that $(I_3 - nn^T)(I_3 - nn^T) = I_3 - nn^T$, that leads to :

$$\begin{aligned} \mathcal{L}(d, n, \lambda) &= \sum_i w_i (x_i - y_i)^2 - n^T \left[\sum_i w_i (x_i - y_i) (x_i - y_i)^T \right] n \\ &\quad + n^T \left[\sum_i w_i (x_i - x_g + y_i - y_g) (x_i - x_g + y_i - y_g)^T \right] n - \lambda(n^T n - 1) \end{aligned}$$

Finally, we get :

$$\mathcal{L}(d, n, \lambda) = \sum_i w_i (x_i - y_i)^2 - n^T (A - \lambda I_3) n + \lambda$$

with

$$A = \sum_i w_i \left[(x_i - x_g + y_i - y_g)(x_i - x_g + y_i - y_g)^T - (x_i - y_i)(x_i - y_i)^T \right]$$

Consequently, the optimal n is such that :

$$\frac{\partial \mathcal{L}}{\partial n} = 0 \Leftrightarrow (A - \lambda I_3)n = 0$$

That is to say that n is an eigenvector of the symmetric matrix A . This yields :

$$\mathcal{L}(d, n, \lambda) = \sum_i w_i (x_i - y_i)^2 + \lambda \text{ at the optimum} \quad (\text{A.1})$$

The first term of this expression does not depend on n and the second term is minimal when λ is minimal. In other words, the optimal n is a unit eigenvector of A corresponding to the smallest eigenvalue λ of A .

Annexe B

The EM algorithm

B.1 Formulation générale

On se place dans le cadre d'un problème "aux données incomplètes" : on observe une v.a Y qui est en relation avec une v.a X non observée mais qui ne la détermine pas complètement ¹. La distribution suivie par la v.a X est fonction d'un paramètre θ inconnu et est modélisée par $p(X|\theta)$.

L'objectif de l'algorithme EM est de calculer le paramètre θ expliquant au mieux la distribution observée de Y , au sens du maximum de vraisemblance :

$$\tilde{\theta} = \arg \max_{\theta} L(\theta) = \arg \max_{\theta} \log(p(Y|\theta))$$

Comme la vraisemblance $\log(p(Y|\theta))$ n'est pas directement accessible, on l'exprime par marginalisation des états cachés (en considérant que X prend sa valeur parmi un nombre fini d'états) :

$$\tilde{\theta} = \arg \max_{\theta} L(\theta) = \arg \max_{\theta} \sum_X \log(P(X, Y|\theta))$$

L'idée de l'algorithme EM consiste à construire une borne inférieure de choix pour L . Dans ce but, on réécrit

$$L(\theta) = \log \left(\sum_X \frac{q(X)P(X, Y|\theta)}{q(X)} \right),$$

où $q(X)$ est une distribution arbitraire des données cachées X . Puis, on remarque (en utilisant l'inégalité de Jensen) que quelque soit q :

$$L(\theta) = \log \left(\sum_X \frac{q(X)p(Y, X|\theta)}{q(X)} \right) \geq \sum_X q(X) \log \left(\frac{p(Y, X|\theta)}{q(X)} \right) = G(\theta, q(X)).$$

¹ Y est donc la forme incomplète de l'observation X

Ainsi, pour une valeur de θ donnée, la vraisemblance est minorée par la fonction G quelque soit q . Cependant, on ne cherche pas simplement à borner L mais à la borner “au plus près”, c’est à dire que l’on va chercher la distribution q qui maximise $G(\theta, q)$ (pour un θ donné) ; la recherche de cette distribution est l’étape E. Cette étape ne modifie pas du tout la valeur de la vraisemblance (simplement car celle-ci ne dépend pas de q). Une fois la distribution q optimisée, la recherche du paramètre θ maximisant $G(\theta, q)$ constitue l’étape M. La succession de ces étapes compose l’algorithme EM :

tant que $\tilde{\theta}$ n’est pas stable
Etape E : $\tilde{q} = \arg \max_q G(\tilde{\theta}, q(W))$
Etape M : $\tilde{\theta} = \arg \max_{\theta} G(\theta, \tilde{q}(X))$

Plus précisément, on peut montrer que

- $G(\theta, q(X)) = -KL(q(X)||p(X|Y, \theta)) + L(\theta)$; ainsi l’étape E est optimale quand $KL(q(X)||p(X|Y, \theta)) = 0$, c’est à dire quand $q(X) = p(X|Y, \theta)$. On remarque que dans cette condition $G(\theta, q(X)) = L(\theta)$.
- $G(\theta, q(X)) = \sum_X q(X) \log(p(Y, X|\theta)) - \sum_X q(X) \log(X)$; ainsi l’étape M est optimale quand $\tilde{\theta} = \arg \max_{\theta} \sum_Y \tilde{q}(X) \log(p(Y, X|\theta))$.

Par suite, on peut réécrire l’algorithme précédent :

tant que $\tilde{\theta}$ n’est pas stable
Etape E : $\tilde{q} = p(X|Y, \theta)$
Etape M : $\tilde{\theta} = \arg \max_{\theta} \sum_X \tilde{q}(X) \log(p(Y, X|\theta))$

Cette séquence de minimisation possède 3 caractéristiques qui en font un algorithme de choix pour la maximisation de la vraisemblance :

- elle augmente de manière monotone (et même linéaire) la valeur de $L(\theta)$,
- elle ne nécessite aucun paramètre spécifique à la procédure d’optimisation (contrairement à la plupart des algorithmes basés gradient),
- pour un grand nombre de problèmes classiques, les étapes E et M ont des solutions simples, voir analytiques.

B.2 L’EM et les GMMs

On suppose que notre pdf $p(y|\theta)$ est construite selon un mélange de K pdf :

$$p(y|\theta) = \sum_{k=1}^K p(y|k, \theta)P(k|\theta),$$

où $p(y|k, \theta) \sim \mathcal{N}(f_{\theta}(\mu_k), f_{\theta}(\Sigma_k))$ représente la pdf de la k ème Gaussienne dont les moyennes et les variances dépendent d’une fonction de θ et où $P(k|\theta)$ représente la probabilité qu’une donnée choisie au hasard ait été générée par la k ème composante du mélange. L’état caché x_i caractérise l’appartenance de l’échantillon y_i à la k ème Gaussienne ; on note $p(x_i = k|y_i, \theta)$ la probabilité que l’échantillon x_i provienne de la k ème Gaussienne sachant la valeur du paramètre θ et l’observation y_i .

Il reste alors à suivre le schéma proposé dans la section précédente :

- L'étape E s'écrit simplement :

$$\begin{aligned} p(x_i = k|y_i, \theta) &= \frac{p(y_i|x_i = k, \theta)p(x_i = k|\theta)}{p(x_i|\theta)} = \frac{p(y_i|x_i = k, \theta)p(x_i = k|\theta)}{\sum_k p(y_i|x_i = k, \theta)p(x_i = k|\theta)} \\ &= \frac{\exp((y_i - f_\theta(\mu_k))^T f_\theta(\Sigma_k)^{-1}(y_i - f_\theta(\mu_k)))p(x_i = k|\theta)}{\sum_k \exp((y_i - f_\theta(\mu_k))^T f_\theta(\Sigma_k)^{-1}(y_i - f_\theta(\mu_k)))p(x_i = k|\theta)} \end{aligned}$$

- L'étape M s'écrit :

$$\begin{aligned} \tilde{\theta} &= \arg \max_{\theta} \sum_i \sum_k q(x_i) \log(p(y_i|x_i = k, \theta)p(y_i = k|\theta)) \\ &= \arg \max_{\theta} \sum_i \sum_k q(x_i) \log(p(y_i|x_i = k, \theta)) + \sum_i \sum_k q(x_i) \log(p(y_i = k|\theta)) \\ &= \arg \max_{\theta} \sum_i \sum_k q(x_i) ((y_i - f_\theta(\mu_k))^T f_\theta(\Sigma_k)^{-1}(y_i - f_\theta(\mu_k))) + \sum_i \sum_k q(x_i) \log(p(y_i = k|\theta)) \end{aligned}$$

Considérant cette forme générale, les spécialisations décrites aux chapitres 3 et 8 ne posent aucune difficulté.

B.3 Discussion

Au-delà du point de vue purement théorique, l'algorithme EM se révèle en pratique rapide et converge vers un maximum proche du maximum global [Ran98]. Toutefois, malgré ces bonnes dispositions, dans certains cas cet algorithme s'avère plus lent et moins efficace que des méthodes de minimisation directe (telles que le gradient conjugué) de la fonction de vraisemblance [Ran98]. En pratique, il est difficile de statuer à l'avance sur les performances relatives des différentes méthodes de maximisation mais il est à noter que la convergence de l'EM va dépendre de l'allure de la fonction $G(\theta, q)$, alors que celle des méthodes basées gradient dépendra de l'allure de la vraisemblance L .

Pour finir, notons que comme la convergence de l'EM dépend grandement de la valeur initiale de θ qui lui est fournie, celle-ci devra être choisie avec soin.

Bibliographie

- [AFK07] D. Aouada, S. Feng, and H. Krim. Statistical Analysis of the Global Geodesic Function for 3D Object Classification. In *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, volume 1, pages I-645–I-648, 2007.
- [AHB87] K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5) :698–700, 1987.
- [ARV07] B. Amberg, S. Romdhani, and T. Vetter. Optimal Step Nonrigid ICP Algorithms for Surface Registration. In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1–8, 2007.
- [Ata85] M. J. Atallah. On symmetry detection. *IEEE Trans. Comput.*, 34(7) :663–666, 1985.
- [BCD⁺03] Michael Kazhdan Bernard, Bernard Chazelle, David Dobkin, Thomas Funkhouser, and Szymon Rusinkiewicz. A reflective symmetry descriptor for 3D models. *Algorithmica*, 38 :201–225, 2003.
- [BCT07] Zheng-Jian Bai, Delin Chu, and Roger C. E. Tan. Computing the nearest doubly stochastic matrix with a prescribed entry. *SIAM J. Sci. Comput.*, 29(2) :635–655, 2007.
- [Ben75] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9) :509–517, 1975.
- [BF98] Samuel R. Buss and Jay P. Fillmore. Spherical averages and applications to spherical splines and interpolation. *ACM Transactions on Graphics*, 20 :95–126, 1998.
- [BHHJ98] Ayanendranath Basu, Ian R. Harris, Nils L. Hjort, and M. C. Jones. Robust and efficient estimation by minimising a density power divergence. *Biometrika*, 85(3) :549–559, 1998.
- [BM92] P.J. Besl and N.D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2) :239–256, December 1992.
- [BMP⁺05] Thomas R Barrick, Clare E Mackay, Sylvain Prima, Frederik Maes, Dirk Vandermeulen, Timothy J Crow, and Neil Roberts. Automatic analysis of cerebral asymmetry : an exploratory study of the relationship between

- brain torque and planum temporale asymmetry. *Neuroimage*, 24(3) :678–691, February 2005.
- [Boo89] Fred L. Bookstein. Principal warps : Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(6) :567–585, June 1989.
- [BWB⁺94] R. M. Bilder, H. Wu, B. Bogerts, G. Degreef, M. Ashtari, J. M. Alvir, P. J. Snyder, and J. A. Lieberman. Absence of regional hemispheric volume asymmetries in first-episode schizophrenia. *Am J Psychiatry*, 151(10) :1437–1447, October 1994.
- [CA04] Pascal Cachier and Nicholas Ayache. Isotropic energies, filters and splines for vector field regularization. *Journal of Mathematical Imaging and Vision*, 20(3) :251–265, May 2004.
- [CDD07] M.K. Chung, K.M Dalton, and R.J. Davidson. Unified cortical asymmetry analysis in autism via weighted-spharm. In *Human Brain Mapping conference*, 2007.
- [CGH⁺93] L. Cohen, C. Gény, O. Hermine, F. Gray, and J. D. Degos. Crossed aphasia with visceral situs inversus. *Ann Neurol*, 33(2) :215–218, February 1993.
- [CH00] Marco Carcassoni and Edwin R. Hancock. Spectral correspondence for deformed point-set matching. In *AMDO*, pages 120–132, 2000.
- [CH01] C. Cantalupo and W. D. Hopkins. Asymmetric broca’s area in great apes. *Nature*, 414(6863) :505, November 2001.
- [CHW⁺08a] Benoît Combès, Robin Hennessy, John Waddington, Neil Roberts, and Sylvain Prima. Automatic symmetry plane estimation of bilateral objects in point clouds. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR’2008*, Anchorage, USA, June 2008.
- [CHW⁺08b] Benoît Combès, Robin Hennessy, John Waddington, Neil Roberts, and Sylvain Prima. An algorithm to map asymmetries of bilateral objects in point clouds. In *5th IEEE International Symposium on Biomedical Imaging : From Nano to Macro, ISBI’2008*, pages 1139–1142, Paris, France, May 2008.
- [CKDD08] M.K. Chung, D.J. Kelley, K.M. Dalton, and R.J. Davidson. Quantifying cortical surface asymmetry via logistic discriminant analysis. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW ’08. IEEE Computer Society Conference on*, pages 1–8, June 2008.
- [CM02] D. Comaniciu and P. Meer. Mean shift : A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24 :603–619, 2002.
- [Cov06] Thomas M. Cover and Joy A. Thomas *Elements of Information Theory*. Wiley, 2006.

- [CP09] Benoît Combès and Sylvain Prima. Setting priors and enforcing constraints on matches for nonlinear registration of meshes. In G.-Z. Yang, D.J. Hawkes, D. Rueckert, J.A. Noble, and C.J. Taylor, editors, *12th International Conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI'2009*, Lecture Notes in Computer Science, pages 175–183, London, UK, September 2009. Springer.
- [CR00a] Haili Chui and Anand Rangarajan. A new algorithm for non-rigid point matching. *cvpr*, 2044 ;2052, 2000.
- [CR00b] Haili Chui and Anand Rangarajan. A feature registration framework using mixture models. *Mathematical Methods in Biomedical Image Analysis, IEEE Workshop on*, 182 :190, 2000.
- [CR03] Haili Chui and Anand Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2-3) :114–141, 2003.
- [CRM01] Dorin Comaniciu, Visvanathan Ramesh, and Peter Meer. The variable bandwidth mean shift and data-driven scale selection. In *in Proc. 8th Intl. Conf. on Computer Vision*, pages 438–445, 2001.
- [Cro04a] T. Crow. Directional asymmetry is the key to the origin of modern Homo sapiens (the Broca-Annett axiom) : A reply to Rogers’ review of The Speciation of Modern Homo Sapiens. *Laterality*, 9(2) :233–242, April 2004.
- [Cro04b] T.J. Crow. Cerebral asymmetry and the lateralization of language : core deficits in schizophrenia as pointers to the genetic predisposition. *Current Opinion in Psychiatry*, 17(2) :97–106, March 2004.
- [CRZL04] Haili Chui, Anand Rangarajan, Jie Zhang, and Christiana Morison Leonard. Unsupervised learning of an atlas from unlabeled point-sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26 :160–172, 2004.
- [CSK05] D. Chetverikov, D. Stepanov, and P. Krsek. Robust Euclidean alignment of 3D point sets : the trimmed iterative closest point algorithm. *Image Vision Computing*, 23(3) :299–309, 2005.
- [CTCG95] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models—their training and application. *Comput. Vis. Image Underst.*, 61(1) :38–59, 1995.
- [CV08] M. Chaouch and A. Verroust-Blondet. A novel method for alignment of 3D models. In *Shape Modeling and Applications, 2008. SMI 2008. IEEE International Conference on*, pages 187–195, 2008.
- [CY01] V. Camion and L. Younes. Geodesic interpolating splines. In *EMMCVPR*, pages 513–527, 2001.
- [Dah09] R. Dahyot. Mean-shift for statistical hough transform. Internal report 2009.

- [Day69] N. E. Day. Estimating the Components of a Mixture of Normal Distributions. *Biometrika*, 56(3) :463–474, 1969.
- [DC76] Manfredo P. Do-Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, February 1976.
- [dFFG93] D. de Falco, M. Frontini, and L. Gotusso. A unifying approach to the regularization of fourier polynomials. *Numerical Algorithms*, 5 :419–424, 1993. 10.1007/BF02109422.
- [DLAV07] Andreas Dedner, Marcel Lüthi, Thomas Albrecht, and Thomas Vetter. *Curvature Guided Level Set Registration Using Adaptive Finite Elements*, pages 527–536. 2007.
- [DLR77a] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the em algorithm. *J. Royal Statistical Society, Series B*, 39(1) :1–38, 1977.
- [DLR77b] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1) :1–38, 1977.
- [DM98] I.L. Dryden and K.V. Mardia. *Statistical Shape Analysis*. John Wiley Sons, 1998.
- [DPTA08] Stanley Durrleman, Xavier Pennec, Alain Trouvé, and Nicholas Ayache. Sparse approximation of currents for statistics on curves and surfaces. In *MICCAI '08 : Proceedings of the 11th International Conference on Medical Image Computing and Computer-Assisted Intervention, Part II*, pages 390–398, Berlin, Heidelberg, 2008. Springer-Verlag.
- [DPTA09] Stanley Durrleman, Xavier Pennec, Alain Trouvé, and Nicholas Ayache. Statistical models of sets of curves and surfaces based on currents. *Medical Image Analysis*, 13(5) :793 – 808, 2009. Includes Special Section on the 12th International Conference on Medical Imaging and Computer Assisted Intervention.
- [Ege86] E.H. Egelman. An algorithm for straightening images of curved filamentous structures. *Ultramicroscopy*, 19(4) :367 – 373, 1986.
- [FA96] Jacques Feldmar and Nicholas Ayache. Rigid, affine and locally affine registration of free-form surfaces. *Int. J. Comput. Vision*, 18(2) :99–119, 1996.
- [FB81] Martin A. Fischler and Robert C. Bolles. Random sample consensus : a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6) :381–395, June 1981.
- [FEdlF+08] Lorenz Fieten, Jörg Eschweiler, Matías de la Fuente, Sascha Gravius, and Klaus Radermacher. Automatic extraction of the mid-sagittal plane using an icp variant. volume 6918, page 69180L. SPIE, 2008.
- [FH75] K. Fukunaga and L. Hostetler. The estimation of the gradient of a density function, with applications in pattern recognition. *Information Theory, IEEE Transactions on*, 21(1) :32–40, 1975.

- [FH86] O.D. Faugeras and M. Hebert. The representation, recognition, and locating of 3-d objects. *Int. J. Rob. Res.*, 5(3) :27–52, 1986.
- [Fit03] Andrew W. Fitzgibbon. Robust registration of 2d and 3d point sets. *Image and Vision Computing*, 21(13-14) :1145 – 1153, 2003. British Machine Vision Computing 2001.
- [GHBB98] P. J. Gannon, R. L. Holloway, D. C. Broadfield, and A. R. Braun. Asymmetry of chimpanzee planum temporale : humanlike pattern of Wernicke’s brain language area homolog. *Science*, 279(5348) :220–222, January 1998.
- [GHR98] William Glunt, Thomas L. Hayden, and Robert Reams. The nearest doubly stochastic matrix to a real matrix with the same first moment. *Numerical Linear Algebra with Applications.*, 5(6) :475–482, 1998.
- [GI04] Jack Goldfeather and Victoria Interrante. A novel cubic-order algorithm for approximating principal direction vectors. *ACM Trans. Graph.*, 23(1) :45–63, 2004.
- [Gil92] Gilles Celeux and Gérard Govaert. A classification EM algorithm for clustering and two stochastic versions. *Computational Statistics and Data Analysis*, 14(3) :315–332, 1992.
- [GL68] N. Geschwind and W. Levitsky. Human brain : left-right asymmetries in temporal speech region. *Science*, 161(837) :186–167, July 1968.
- [Goo00] Phillip Good. *Permutation Tests : A Practical Guide to Resampling Methods for Testing Hypotheses*. Springer, 2nd edition, 2000.
- [GP02] S. Granger and X. Pennec. Multi-scale EM-ICP : A fast and robust approach for surface registration. volume 2353 of *Lecture Notes in Computer Science*, pages 418–432, Copenhagen, Denmark, May 2002. Springer.
- [GR92] Donald Geman and George Reynolds. Constrained restoration and the recovery of discontinuities. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(3) :367–383, 1992.
- [GR09] Hongyu Guo and Anand Rangarajan. Diffeomorphic point matching with applications in biomedical image registration. *International Journal of Tomography and Statistics*, 15(F10), 2009.
- [Gre90] Peter J. Green. On Use of the EM for Penalized Likelihood Estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, 52(3) :443–452, 1990.
- [GS91] Leslie Greengard and John Strain. The fast gauss transform. *SIAM J. Sci. Stat. Comput.*, 12(1) :79–94, 1991.
- [GTY04] Joan Glaunes, Alain Trouvé, and Laurent Younes. Diffeomorphic matching of distributions : A new approach for unlabelled point-sets and sub-manifolds matching. In *In CVPR*, pages 712–718, 2004.
- [GWM01] Stefan Gumhold, Xinlong Wang, and Rob Macleod. Feature Extraction from Point Clouds. In *In Proceedings of the 10 th International Meshing Roundtable*, pages 293–305, 2001.

- [HAWG08] Qi-Xing Huang, Bart Adams, Martin Wicke, and Leonidas J. Guibas. Non-rigid registration under isometric deformations. *Comput. Graph. Forum*, 27(5) :1449–1457, 2008.
- [HEP⁺09] Heike Hufnagel, Jan Ehrhardt, Xavier Pennec, Nicholas Ayache, and Heinz Handels. Computation of a probabilistic statistical shape model in a maximum-a-posteriori framework. *Methods of Information in Medicine*, 48(4) :314–319, 2009.
- [HFY⁺10] R Horaud, F Forbes, M Yguel, G Dewaele, and J Zhang. Rigid and articulated point registration with expectation conditional maximization. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99) :1–8, 2010.
- [HHN88] B.K.P. Horn, H.M. Hilden, and S. Negahdaripour. Closed-form solution of absolute orientation using orthonormal matrices. *Journal of the Optical Society of America A*, 5(7) :1127–1135, July 1988.
- [Hor87] B. Horn. Closed-form solution of absolute orientation using unit quaternions. *J of the Optical Society of America*, 4 :629–642, 1987.
- [HPE⁺07] Heike Hufnagel, Xavier Pennec, Jan Ehrhardt, Heinz Handels, and Nicholas Ayache. Point-based statistical shape models with probabilistic correspondences and affine em-icp. In Springer, editor, *Bildverarbeitung fuer die Medizin 2007*, Informatik aktuell, pages 434–438. Springer, March 2007. Third prize in category best scientific work.
- [HPE⁺08a] H. Hufnagel, X. Pennec, J. Ehrhardt, N. Ayache, and H. Handels. Comparison of statistical shape models built on correspondence probabilities and one-to-one correspondences. In *Proc. SPIE Symposium on Medical Imaging '08*, volume 6914 of *SPIE Conference Series*, San Diego, USA, February 2008.
- [HPE⁺08b] H. Hufnagel, X. Pennec, J. Ehrhardt, N. Ayache, and H. Handels. Generation of a Statistical Shape Model with Probabilistic Point Correspondences and EM-ICP. *International Journal for Computer Assisted Radiology and Surgery (IJCARs)*, 2(5) :265–273, March 2008.
- [Hub81] Peter J. Huber. *Robust statistics / Peter J. Huber*. Wiley, New York :, 1981.
- [HWWpM05] Tobias Heimann, Ivo Wolf, Tomos Williams, and Hans peter Meinzer. 3d active shape models using gradient descent optimization of description length. In *in Proc. IPMI*, pages 566–577. Springer, 2005.
- [IHF⁺10] Aya Ihara, Masayuki Hirata, Norio Fujimaki, Tetsu Goto, Yuka Umekawa, Norihiko Fujita, Yasushi Terazono, Ayumu Matani, Qiang Wei, Toshiki Yoshimine, Shiro Yorifuji, and Tsutomu Murata. Neuroimaging study on brain asymmetries in situs inversus totalis. *J Neurol Sci*, 288(1-2) :72–78, January 2010.

- [JC02] H.J. Johnson and G.E. Christensen. Consistent landmark and intensity-based image registration. *IEEE Transactions on Medical Imaging*, 21(5) :450–461, May 2002.
- [JM00] S.C. Joshi and M.I. Miller. Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing*, 9(8) :1357–1370, August 2000.
- [Jos07] Eric Klassen Shantanu H. Joshi. A novel representation for riemannian analysis of elastic curves in \mathbb{R}^n . *CVPR*, 1-7, June 2007.
- [JV05] Bing Jian and B.C. Vemuri. A robust algorithm for point set registration using mixture of gaussians. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 2, pages 1246–1251 Vol.2, 17-21 Oct. 2005.
- [KGH93] C. Kambhamettu, DB Goldgof, and M. He. On a study of invariant features in nonrigid transformations. In *Qualitative Vision, 1993., Proceedings of IEEE Workshop on*, pages 118–127, 1993.
- [KGHL03] C. Kambhamettu, D. Goldgof, M. He, and P. Laskov. 3d nonrigid motion analysis under small deformations. *Image and Vision Computing*, 21(3) :229–245, 2003.
- [KKM03] Shunichi Kaneko, Tomonori Kondo, and Atsushi Miyamoto. Robust matching of 3d contours using iterative closest point algorithm improved by M-estimation. *Pattern Recognition*, 36(9) :2041–2047, 2003.
- [Knu98] Donald E. Knuth. *The art of computer programming, volume 3 : (2nd ed.) sorting and searching*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1998.
- [KOT+99] D. N. Kennedy, K. M. O’Craven, B. S. Ticho, A. M. Goldstein, N. Makris, and J. W. Henson. Structural and functional brain asymmetries in human situs inversus totalis. *Neurology*, 53(6) :1260–1265, October 1999.
- [KvD92] Jan J. Koenderink and Andrea J. van Doorn. Surface shape and curvature scales. *Image and Vision Computing*, 10(8) :557–564, October 1992.
- [LAV07] Marcel Lüthi, Thomas Albrecht, and Thomas Vetter. Curvature guided surface registration using level sets, 2007.
- [Lea93] V. F. Leavers. Which hough transform? *CVGIP : Image Underst.*, 58(2) :250–264, 1993.
- [LeM99] M. LeMay. Functional and anatomical asymmetries of the human brain. *Eur J Neurol*, 6(1) :79–85, January 1999.
- [LH05] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. *Computer Vision, IEEE International Conference on*, 2 :1482–1489, 2005.
- [LK03] P. Laskov and C. Kambhamettu. Curvature-based algorithms for nonrigid motion and correspondence estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10) :1349–1354, 2003.

- [LKAD⁺09] Oliver C. Lyttelton, Sherif Karama, Yasser Ad-Dab'bagh, Robert J. Zatorre, Felix Carbonell, Keith Worsley, and Alan C. Evans. Positional and surface area asymmetry of the human cerebral cortex. *NeuroImage*, 46(4) :895 – 903, 2009.
- [LMS⁺08] Xinyang Liu, Washington Mio, Yonggang Shi, Ivo Dinov, Xiuwen Liu, Natasha Leporé, Franco Leporé, Madeleine Fortin, Patrice Voss, Maryse Lassonde, and Paul M. Thompson. Models of normal variation and local contrasts in hippocampal anatomy. In *MICCAI '08 : Proceedings of the 11th International Conference on Medical Image Computing and Computer-Assisted Intervention, Part II*, pages 407–415, Berlin, Heidelberg, 2008. Springer-Verlag.
- [LR09] Yu-Shen Liu and Karthik Ramani. Robust principal axes determination for point-based shapes using least median of squares. *Computer-Aided Design*, 41(4) :293–305, April 2009.
- [M'02] Mémin, E. and Pérez, P. Hierarchical estimation and segmentation of dense motion fields. *Int. Journal of Computer Vision*, 46(2) :129–155, February 2002.
- [Mar75] F. H. C. Marriott. Separating Mixtures of Normal Distributions. *Biometrics*, 31(3) :767–769, 1975.
- [MBM00] K. V. Mardia, F. L. Bookstein, and I. J. Moreton. Statistical assessment of bilateral symmetry of shapes. *Biometrika*, 87(1) :285–300, 2000.
- [MBR⁺03] C.E. Mackay, T.R. Barrick, N. Roberts, L.E. DeLisi, F. Maes, D. Vandermeulen, and T.J. Crow. Application of a new image analysis technique to study brain asymmetry in schizophrenia. *Psychiatry Research : Neuroimaging*, 124(1) :25–35, September 2003.
- [McG80] McGlone. Sex differences in human brain asymmetry : A critical survey. 1980.
- [MCP⁺06] E. De Momi, J. Chapuis, I. Pappas, G. Ferrigno, W. Hallermann, A. Schramm, and M. Caversaccio. Automatic extraction of the mid-facial plane for cranio-maxillofacial surgery planning. *Int J Oral Maxillofac Surg*, 35(7) :636–642, July 2006.
- [MCP10] Daniel Munch, Benoît Combès, and Sylvain Prima. A modified ICP algorithm for normal-guided surface registration . In *Proceedings of SPIE Medical Imaging 2010 : Image Processing*, San Diego, USA, February 2010.
- [MGP07] N. J. Mitra, L. Guibas, and M. Pauly. Symmetrization. In *ACM Transactions on Graphics*, volume 26, pages #63, 1–8, 2007.
- [MHK⁺08] Diana Mateus, Radu P. Horaud, David Knossow, Fabio Cuzzolin, and Edmond Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2008.

- [MIK93] P. Minovic, S. Ishikawa, and K. Kato. Symmetry identification of a 3-D object represented by octree. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 15(5) :507–514, 1993.
- [MS09] Andriy Myronenko and Xubo Song. Point-set registration : Coherent point drift. *PAMI*, May 2009.
- [MSCP07] Andriy Myronenko, Xubo Song, and Miguel Carreira-Perpinan. Non-rigid point set registration : Coherent point drift. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems (NIPS) 19*, pages 1009–1016, Cambridge, MA, 2007. MIT Press.
- [MSHS06] Aurélien Martinet, Cyril Soler, Nicolas Holzschuch, and François X. Sillion. Accurate detection of symmetries in 3D shapes. *ACM Trans. Graph.*, 25(2) :439–464, 2006.
- [NH02] T.E. Nichols and A.P. Holmes. Nonparametric permutation tests for functional neuroimaging : A primer with examples. *Human Brain Mapping*, 15 :1–25, 2002.
- [ÓLD⁺07] H. Ólafsdóttir, S. Lanche, T.A. Darvann, N.V. Hermann, R. Larsen, B.K. Ersbll, E. Oubel, A.F. Frangi, P. Larsen, C.A. Perlyn, G.M. Morriss-Kay, and S. Kreiborg. A point-wise quantification of asymmetry using deformation fields : application to the study of the Crouzon mouse model. In N. Ayache, S. Ourselin, and A. Maeder, editors, *10th International Conference on Medical Image Computing and Computer-Assisted Intervention, MICCAI'2007*, volume 4792 of *lncs*, pages 452–459, Brisbane, Australia, October 2007. Springer.
- [PeS⁺06] Konrad Polthier, Alla Sheffer. Folding meshes : Hierarchical mesh segmentation based on planar symmetry. 2006.
- [PHYH06] Helmut Pottmann, Qi-Xing Huang, Yong-Liang Yang, and Shi-Min Hu. Geometry and convergence analysis of algorithms for registration of 3D shapes. *International Journal of Computer Vision*, 67(3) :277–296, May 2006.
- [PLNS04] D. Pantazis, R.M. Leahy, T.E. Nichols, and M. Styner. Statistical surface-based morphometry using a nonparametric approach. In *Biomedical Imaging : Nano to Macro, 2004. IEEE International Symposium on*, pages 1283 – 1286 Vol. 2, 15-18 2004.
- [PSG⁺06] Joshua Podolak, Philip Shilane, Aleksey Golovinskiy, Szymon Rusinkiewicz, and Thomas Funkhouser. A planar-reflective symmetry transform for 3d shapes. *ACM SIGGRAPH*, 25 :549—559, 2006.
- [PWQW06] Gang Pan, Yueming Wang, Yipeng Qi, and Zhaohui Wu. Finding symmetry plane of 3D face shape. In *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on*, volume 3, pages 1143–1146, 2006.
- [Ran98] Anand Rangarajan. Everything you always wanted to know about the EM, Internal report 1998.

- [RCM⁺97] A. Rangarajan, H. Chui, E. Mjolsness, S. Pappu, L. Davachi, P. Goldman-Rakic, and J. Duncan. A robust point-matching algorithm for autoradiograph alignment. *Med Image Anal*, 1(4) :379–398, September 1997.
- [RGR07] Arunabha S Roy, Ajay Gopinath, and Anand Rangarajan. Deformable density matching for 3d non-rigid registration of shapes. *Miccai*, 10(Pt 1) :942–9, 2007.
- [RL07] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. In *Proceedings of the Third Intl. Conf. on 3D Digital Imaging and Modeling*, pages 145–152, 2007.
- [RMH76] A. B. Rubens, M. W. Mahowald, and J. T. Hutton. Asymmetry of the lateral (sylvian) fissures in man. *Neurology*, 26(7) :620–624, July 1976.
- [RMP⁺96] Anand Rangarajan, Eric Mjolsness, Suguna Pappu, Lila Davachi, Patricia S. Goldman-Rakic, and James S. Duncan. A robust point matching algorithm for autoradiograph alignment. pages 277–286, 1996.
- [Rob04] Steven M. Robbins. *Anatomical standardization of the human brain in euclidean 3-space and on the cortical 2-manifold*. PhD thesis, McGill University, 2004.
- [Rob06] Robert, Christian P *Le choix bayésien*. Springer, 2006.
- [SB92] Larry S Shapiro and J Michael Brady. Feature-based correspondence : an eigenvector approach. *Image and Vision Computing*, 10(5) :283 – 288, 1992. BMVC 1991.
- [Sco99] D.W. Scott. Remarks on fitting and interpreting mixture models. *Computing Science and Statistics*, 31 :104–109, 1999.
- [SDT09] Romeil Sandhu, Samuel Dambreville, and Allen Tannenbaum. Point set registration via particle filtering and stochastic dynamics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 99, 2009.
- [SG02] Mikkel B. Stegmann and David Delgado Gomez. A brief introduction to statistical shape analysis. Technical report, University of Denmark, DTU, 2002.
- [SHMS07] Li Shen, Heng Huang, Fillia Makedon, and Andrew J. Saykin. Efficient registration of 3d spharm surfaces. *Computer and Robot Vision, Canadian Conference*, 0 :81–88, 2007.
- [SHS01] Bernhard Schölkopf, Ralf Herbrich, and Alex J. Smola. A generalized representer theorem. In *In Proceedings of the Annual Conference on Computational Learning Theory*, pages 416–426, 2001.
- [Sid04] T. Sidlofov. Existence and uniqueness of minimization problems with Fourier based stabilizers. In *Compstat*, Prague, 2004.
- [Sin64] Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35(2) :876–879, 1964.

- [SLH91] G. L. Scott and H. C. Longuet-Higgins. An algorithm for associating the features of two images. *Proceedings. Biological sciences / The Royal Society*, 244(1309) :21–26, April 1991.
- [SM09] Raghav Subbarao and Peter Meer. Nonlinear mean shift over riemannian manifolds. *Int. J. Comput. Vision*, 84(1) :1–20, 2009.
- [SMA⁺08] David W. Shattuck, Mubeena Mirza, Vitria Adisetiyo, Cornelius Hojat-kashani, Georges Salamon, Katherine L. Narr, Russell A. Poldrack, Robert M. Bilder, and Arthur W. Toga. Construction of a 3d probabilistic atlas of human cortical structures. *NeuroImage*, 39(3) :1064 – 1080, 2008.
- [SS71] A. J. Scott and M. J. Symons. Clustering Methods Based on Likelihood Ratio Criteria. *Biometrics*, 27(2) :387–397, 1971.
- [SS97] Changming Sun and J. Sherrah. 3D symmetry detection using the extended gaussian image. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 19(2) :164–168, 1997.
- [SSDK09] Chafik Samir, Anuj Srivastava, Mohamed Daoudi, and Sebastian Kurttek. On analyzing symmetry of objects using elastic deformations. In *VISSAPP (1)*, pages 194–200, 2009.
- [TC07] Tommy Tang and Albert Chung. Non-rigid image registration using graph-cuts. In *Medical Image Computing and Computer-Assisted Intervention*, volume 4791 of *Lecture Notes in Computer Science*, pages 916–924, 2007.
- [TK04] Y. Tsin and T. Kanade. A correlation-based approach to robust point set registration. pages Vol III : 558–569, 2004.
- [Ume91] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(4) :376–380, 1991.
- [VB62] G. Von Bonin. Anatomical asymmetries of the cerebral hemispheres. In V.B. Mountcastle, editor, *Interhemispheric Relations and Cerebral Dominance*, pages 1–6. Johns Hopkins Press, Baltimore, USA, 1962.
- [VG05] Marc Vaillant and Joan Glaunès. Surface matching via currents. In *Proceedings of Information Processing in Medical Imaging (IPMI 2005)*, number 3565 in *Lecture Notes in Computer Science*, pages 381–392, 2005.
- [VSR01] D.V. Vranic, D. Saupe, and J. Richter. Tools for 3D-object retrieval : Karhunen-Loeve transform and spherical harmonics. In *Multimedia Signal Processing, 2001 IEEE Fourth Workshop on*, pages 293–298, 2001.
- [Wel97] William M. Wells, III. Statistical approaches to feature-based object recognition. *Int. J. Comput. Vision*, 21(1-2) :63–98, 1997.
- [Wen05] H Wendland. *Scattered Data Approximation*. Cambridge Monographs on Applied and Computational Mathematics, 2005.

- [WLMW82] D. R. Weinberger, D. J. Luchins, J. Morihisa, and R. J. Wyatt. Asymmetrical volumes of the right and left frontal and occipital regions of the human brain. *Ann Neurol*, 11(1) :97–100, January 1982.
- [WMH10] Hanzi Wang, Daniel Mirotu, and Gregory D. Hager. A generalized kernel consensus-based robust estimator. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32 :178–184, 2010.
- [Woo86] R. P. Woods. Brain asymmetries in situs inversus. A case report and review of the literature. *Arch Neurol*, 43(10) :1083–1084, October 1986.
- [WPW03] Y. Wu, G. Pan, and Z. Wu. Face authentication based on multiple profiles extracted from range data. *Audio- and Video-Based Biometric Person Authentication*, pages 515–522, 2003.
- [WVR06] Fei Wang, Baba C. Vemuri, and Anand Rangarajan. Groupwise point pattern registration using a novel CDF-based Jensen-Shannon divergence. In *CVPR*, pages 1283–1288, 2006.
- [WVRE08] Fei Wang, Baba C. Vemuri, Anand Rangarajan, and Stephan J. Eisenchenk. Simultaneous nonrigid registration of multiple point sets and atlas construction. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(11) :2011–2022, 2008.
- [WVS09] Fei Wang, Baba Vemuri, and Tanveer Syeda-Mahmood. Generalized L2-Divergence and Its Application to Shape Alignment. *IPMI* pages 227–238, 2009.
- [WWV85] J.D. Wolter, T.C. Woo, and R.A. Volz. Optimal algorithms for symmetry detection in two and three dimensions. 1 :37–48, 1985.
- [YG89] Alan L. Yuille and Norberto M. Grzywacz. A mathematical analysis of the motion coherence theory. *International Journal of Computer Vision*, 3(2) :155–175, June 1989.
- [YSV⁺10] B.T.T. Yeo, M.R. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, and P. Golland. Spherical demons : Fast diffeomorphic landmark-free surface registration. *Medical Imaging, IEEE Transactions on*, 29(3) :650–668, march 2010.
- [ZHM07] Guangyu Zou, Jing Hua, and Otto Muzik. Non-rigid Surface Registration Using Spherical Thin-Plate Splines. *MICCAI 2007*, pages 367–374.
- [ZPA92] H. Zabrodsky, S. Peleg, and D. Avnir. A measure of symmetry based on shape similarity. In *Computer Vision and Pattern Recognition, 1992. Proceedings CVPR '92., 1992 IEEE Computer Society Conference on*, pages 703–706, 1992.
- [ZPA95] H. Zabrodsky, S. Peleg, and D. Avnir. Symmetry as a continuous feature. *IEEE PAMI*, 17 :1154–1166, 1995.
- [ZRF⁺06] Liyan Zhang, Anshuman Razdan, Gerald E. Farin, John Femiani, Myung-Soo Bae, and Charles Lockwood. *3D face authentication and recognition*

based on bilateral symmetry analysis. The Visual Computer, 22(1) :43-55, 2006.

Table des figures

1	A gauche : Classification des formes de feuilles tirés de “Hortus Cliffortianus” de Carl Von Linné. A droite : Les pinsons de Darwin L’isolation des Pinson dans les îles Galápagos produit environ une douzaine de nouvelles espèces. Les différences anatomiques les plus importantes entre ces espèces se trouvent dans la taille et la forme de leurs becs. Ces variations s’expliquent par les différences de nourriture de chaque espèce. Leurs comportements tout comme leurs chants sont différents. . . .	8
2	Quelques exemples des symétries du monde vivant	9
3	Représentation graphique des trois types de symétrie à l’échelle d’une population	10
4	Illustration de l’asymétrie de la fissure de Sylvius et de la torsion globale du cerveau	12
1.1	Parametrisation of a plane in 3D.	23
2.1	2D illustration of “the plane P that best passes through the point set X ”	25
2.2	2D illustration of “the plane P for which the minimum effort is required to turn X into a mirror symmetric point set with respect to plane P ”. The circles represent (x_i) , the crosses represent (x_i^*)	28
2.3	2D illustration of “the plane P that minimises the left-right differences of X ”.	29
2.4	2D illustration of "the plane with respect to which the most points are mirror-symmetrical".	34
3.1	Location of keypoints on various objects.	40
3.2	2D illustration of the computation of $c_{i,j}$	41
3.3	2D illustration of the two problems related to the use of the parameter d for clustering.	44
3.4	Some results obtained on real data for coarse estimation of the plausible meaningful reflections of an object	47
3.5	Estimation of the symmetry plane on ground truth, asymmetrical data. Estimation of the plane by Sym2 on simulated data. . . .	54
3.6	Mean angular and linear errors with respect to the percentage of outlying points in the data.	55
3.7	Estimation of the symmetry plane using Sym2 on different structures.	56

4.1	A curved symmetry surface : 2D illustration of the curved symmetry surface.	58
4.2	A curved symmetry surface : 2D illustration of the approximate solution proposed for the M-step. \tilde{C} is the surface best fitting the mid-points computed from the correspondences (x_i, y_i)	59
4.3	Two applications of the symmetry surface estimation : From left to right : i) a cortical surface, ii) the symmetry surface estimated on the cortex (for a better visualisation of the results, we display each hemisphere with a different colour), iii) colour representation of the absolute maximal curvature values along the antero-posterior axis (hot colours indicates high curvatures), iv) symmetry surface estimated on the Tautavel skull and v) we straighten the Tautavel skull according to the estimated symmetry surface [Ege86]	60
7.1	Point sets as Gaussian Mixture Models (GMMs) : Mid-sagittal outlines of corpus callosa shapes represented as probability density functions, from [WVR06]	68
7.2	Some classical functions ρ used in M-estimation $Tukey(x, b) = \min(1 - (1 - x^2/b)^3, 1.0)$, $Leclerc(x, b) = 1 - \exp(-x^2/b)$ and $McClure(x, b) = x^2/(b + x^2)$	70
7.3	Scalar versus vectorial TPS deformations One of four points initially forming a square is forced to move downwards. With scalar TPS (left), vertical lines remain vertical and straight. With vectorial TPS (right), the displacement field involves both vertical and horizontal components. From [CA04]	77
8.1	Effect of the asymmetric normalisation constraint on A . The "correct" matchings are $(x_i, y_i)_{i=1..5}$. One considers T as the identity. From left to right : i) two point sets X and Y , ii) distances involved in the computation of A_{11} when registering Y on X . iii) distances involved in the computation of A_{11} when registering X on Y . As an example, if one considers $\sigma = 1$, then in case ii) $A_{11} = 0.01$ and points x_1 and y_1 have only little chance to be matched at the end of the overall process whereas in case iii) $A_{11} = 0.45$	92
8.2	Mapping of descriptor values : From left to right : curvedness, shape index and total geodesic distance on two lateral ventricles. Homologous anatomical landmarks yield qualitatively the same descriptor values. . .	97
8.3	Example of initialisation on two ventricles . Left : Initialisation with center of mass and principal axes superimposition. Right : Affine transformation estimation with our method	99

8.4	Illustration of the interest of the compatibility term : From left to right i) Point x is deformed towards Y by the deformation $t(x)^X$; ii) Point y is deformed towards X by the deformation $\tilde{t}^Y(y)$; iii) the compatibility term imposes that $t^X(y + \tilde{t}^Y(y))$ should be equal to $-\tilde{t}^Y(y)$; iv) The resulting deformation field \tilde{t}^X tends to be a compromise (in the sense of the regularisation) between data attachment and compatibility.	102
8.5	Effect of parameters β and b on the approximation of a noisy field.	107
8.6	Different positive definite kernels ϕ and their associated value $\phi^*(\ \omega\ /b)^{-1}$. From left to right and top to bottom : i) 2D plot of three kernels ($b=1$), ii) 2D plot of the value of $\phi^*(\ \cdot\)^{-1}$ for the three above-mentioned kernels, iii) 2D plot of $\phi(\cdot \times b)$ for the Wu kernel with different b , iv) 2D plot of $\phi^*(\ \omega\ /b)^{-1}$ for the Wu kernel with different b	108
8.7	Influence of b on the registration. First line : We register the source ventricle (a) on the target ventricle (e) by increasing b throughout the iterations (without modifying the other parameters). Intermediate registration results for a given b are represented between the source and the target (b), (c) and (d). For a small b , the source is only globally transformed towards the target and small patterns are left unchanged (dark blue circles). When b increases, small patterns tend to fit the target (green circles). Light blue circles indicate intermediate configurations of the patterns. Second line : Same description that for the first line but with caudate nuclei.	109
9.1	Illustration of the generation of the ground truth data. From left to right : i) original data, ii) we generate a random local deformation (the resulting distance between the corresponding points is mapped), iii) we generate a random global deformation (TPS) (the resulting distance between the corresponding points is mapped) and iv) superimposition of both the original (red) and deformed (green) data.	113
9.2	Examples of pairs of ground truth data : From left to right, thalamus, ventricles, osseous labyrinths and endocasts	113
9.3	Position of landmarks on two of our data	114
11.1	Mean shape and first mode of variation ($\pm 2 \sqrt{\lambda_1}$) on 40 human hands (described by labelled landmarks) computed with a classical Procrustes analysis (From [SG02]) and on 10 osseous labyrinths (described by free-form surfaces) computed with an automatic method described in this chapter (Section 11.1.2.4).	123

11.2	First row : Illustration of the classical MS-EM-ICP. From left to right : i) the point set X and M , ii) $T_{inv}(X)$ and M (T_{inv} is a rigid-body transformation) and the resulting matching $\sum_i A_{i,j}T_{inv}(x_i)$ and iii) the resulting correspondences $v_j = \sum_i A_{i,j}T_{inv}(x_i)$ and mean $(v_j + m_j)/2$, Second row : Illustration of the improved MS-EM-ICP. From left to right : i) $T(X)$ on M (T is a non-linear smooth transformation) and the resulting matching $\sum_i A_{i,j}T(x_i)$ and ii) the resulting correspondences $v_j = \sum_i A_{i,j}T_{inv}(x_i)$ and mean $(v_j + m_j)/2$	127
11.3	Mapping asymmetries on one subject : All the rows are described from left to right. First row : simplified point set X - superposition of $S_P(X)$ on X - superposition of $T \circ S_P(X)$ on X (the patchy appearance shows the residual errors after nonlinear registration of X and $S_P(X)$). Second row : left-right (width), head-foot and posterior-anterior (protrusion) components of the asymmetry field. Third row : projection of the 3 asymmetry maps of X (second row) on the mean point set M	132
12.1	How to read asymmetry maps.	133
12.2	Mean for each of the 3 components of the asymmetry field (11 subjects). Each row shows one of the components (from top to bottom : left-right, head-foot, posterior-anterior) with different views.	136
12.3	p-value maps for each of the 3 components of the asymmetry field (11 subjects). Each row shows one of the components (from top to bottom : left-right, head-foot, posterior-anterior) with different views. The null hypothesis is that there is no asymmetry (t-test, corrected for multiple comparisons using suprathreshold cluster size tests [NH02]).	136
12.4	Asymmetry map for each of the 3 components of 1170.	138
12.5	Asymmetry map for each of the 3 components of 1175.	138
12.6	p-value of (corrected) t-test with H_0 : 1170 belongs to the control population.	140
12.7	p-value of (corrected) t-test with H_0 : 1175 belongs to the control population.	141
12.8	p-value of permutation test with H_0 : SI population and control population are identical.	141
13.1	From left to right : a deformation field is applied to a perfectly symmetrical face ; estimated asymmetry map by our strategy ; normalised asymmetry map on the template ; template-based asymmetry map as estimated by the method of Ólafsdóttir [ÓLD ⁺ 07]. Both strategies are implemented using the same registration algorithm.	147

14.1	Distance color map pattern between the endocast and brain hull shown on the average endocast mesh. (a) Top view. (b) Bottom view. (c) Front view. (d) Right side view. (e) Rear view. (f) Left side view. The linear color map scale is in millimeters and the distance values are shown on the color bar.	151
14.2	Global torque pattern and local asymmetries on CM1.	152
14.3	Asymmetry map on Mrs Ples. Areas 17-18-19-39-40, the left side protudes more than the right. Areas 41 to 45, the right side protudes more than the left	152
14.4	Mean shape and first mode of variation ($\pm 2 \sqrt{\lambda_1}$) and on 10 osseous labyrinths : the first mode can be interpreted as a size change of the canals.	153

Résumé

Dans cette thèse, nous proposons un ensemble d'outils permettant de quantifier et de comparer statistiquement l'asymétrie au sein d'une population ou entre deux populations de structures représentées par des nuages de points 3D. Ces outils sont définis dans un cadre cohérent qui repose sur une modélisation des nuages de points sous forme de mélanges de densités de probabilité. Ces densités dépendent de paramètres décrivant par exemple une réflexion ou encore une déformation tridimensionnelle. L'estimation des paramètres superposant au mieux les mélanges est effectuée par des variantes d'algorithmes de type Espérance-Maximisation. En particulier, nous proposons des algorithmes pour i) l'estimation d'un plan de symétrie d'un nuage de points, ii) l'estimation d'une surface de symétrie d'un nuage de points, iii) l'estimation d'une déformation non-linéaire superposant au mieux 2 nuages de points et iv) la création de modèles statistiques de formes à partir d'un ensemble de nuages de points. Nous montrons ensuite comment combiner ces méthodes pour générer des cartes statistiques d'asymétrie locale (moyenne, variance, asymétries significatives, différences significatives, *etc*) ainsi que des mesures globales caractérisant l'asymétrie des populations étudiées. Enfin, nous appliquons nos outils pour comparer les asymétries corticales de 2 sujets Situs Inversus (maladie congénitale dans laquelle les principaux organes viscéraux sont situés en miroir par rapport à leur situation normale) avec celles d'une population contrôle.

Abstract

In this thesis, we propose a set of tools for the statistical quantification and comparison of asymmetry within a population or between two populations of structures represented by 3D point sets. These tools are designed in a coherent framework relying on a modeling of the point sets as mixtures of probability density functions depending on parameters describing for example a reflection or a 3D deformation. The estimation of the parameters best superposing the mixtures is performed with variants of the Expectation-Maximisation algorithm. Particularly, we propose algorithms allowing i) the estimation of a symmetry plane of a point set, ii) the estimation of a symmetry surface of a point set, iii) the estimation of a non-linear deformation best superposing 2 point sets and iv) the building of statistical shape models from a set of point sets. Then, we show how to combine these methods to build statistical asymmetry maps (mean, standard deviation, significant asymmetries, significant differences, *etc*) and global measurements characterising the asymmetry of the populations under study. Finally, we use our tools to compare the cortical asymmetries of 2 Situs Inversus subjects (congenital condition in which the major visceral organs are mirrored from their normal positions) with those of a control population.