Diffeomorphic image registration taking topological differences into account : metamorphosis on brain MRI containing glioblastomas

Anton Francois

Université Paris Cité

ipnp

LABORATOIRE MAP5

UNIVERSITÉ PARIS-CITÉ

École doctorale de Science Mathématiques de Paris Centre (ED 386)

Laboratoire d'accueil : Laboratoire de Mathématiques Appliquées de Paris 5 (MAP5), UMR 8145 CNRS

Diffeomorphic image registration taking topological differences into account.

Metamorphosis on brain MRI containing Glioblastomas

THÈSE DE DOCTORAT MATHÉMATIQUES

Spécialité : Mathématiques Appliquées

Anton FRANÇOIS

Encadré par
Joan A. GLAUNÈS and Pietro GORI

Présentée et soutenue publiquement le : 23 mai 2023

Après avis des rapporteurs : FRANÇOIS ROUSSEAU (IMT Atlantique)
JOHN ASHBURNER (University College London)

Jury de soutenance :

FRANÇOIS ROUSSEAU	(IMT Atlantique)	Rapporteur
JOHN ASHBURNER	(University College London)	Rapporteur
JOAN GLAUNÈS	(Université Paris Cité)	Directeur de thèse
PIETRO GORI	(Télécom Paris)	Membre invité
IRÉNE KALTENMARK	(Univesité Paris Cité)	Examinatrice
ALAIN TROUVÉ	(ENS-Paris-Saclay)	Examinateur

NNT :

> “La neuvième forme tue à coup sûr le chameau. Elle blesse à mort le lion. Mais l’enfant que tu sauras peut-être devenir pourrait lui survivre. Ces trois métamorphoses peuvent être les étapes d’une vie, d’un amour, d’une quête...
> Penses-y quand tu seras sur le bord du monde,...”

Alain Damasio, *La Horde du contrevent*

Université Paris Cité
Campus Saint-Germain
45 rue des Saints-Pères
75006 Paris FRANCE

Remerciements

J'ai eu la chance d'avoir Joan Glaunès et Pietro Gori comme directeurs de thèse qui m'ont fait bénéficié de leurs points de vue et forces respectives. Merci pour les diverses relectures et retours sur mon travail. Merci à toi Joan pour ta patience et ta disponibilité pendant nos entrevues. Merci aussi d'avoir toujours pris le temps de discuter lorsque nous nous croisions dans les couloirs. Pietro, je te remercie car tu as su me poser les bonnes questions pour me positionner au sein de la communauté scientifique et ses multiples usages.

Je remercie tout particulièrement John Ashburner et François Rousseau d'avoir rapporté ma thèse.

Ces presque quatre ans ont aussi été parsemés de nombreuses discussions et d'échanges inspirants. Je ne peux citer ici que les plus récurrents: Julie Delon, Remi Abergel et bien sûr Joan Glaunès qui m'ont tous trois accueilli au MAP5 en m'acceptant en stage. Barbara Gris qui a fait partie de mon commité de mi-parcours et qui a toujours su être bienveillante et à l'écoute. J'aimerais aussi remercier Clément Debacker et toute l'équipe de l'ima-brain à l'IPNP qui m'ont accueilli et fourni un ordinateur doté d'une jolie puissance de calcul. Je dois enfin mentionner Raphaël Tinarrage, qui est avant tout mon ami mais qui est aussi devenu mon collègue pendant nos babillages segmentatif.

J'ai eu la chance de faire mon doctorat au MAP5 où la communauté est soudée, bienveillante et joyeuse. Chacun.e d'entre eux participe uniquement à ce vivant patchwork et ce sera une nouvelle fois difficile d'en faire une liste exhaustive. Dès mon arrivée, en tant que stagiaire puis doctorant, j'y ai été accueilli par les "anciens"; Claire qui émane de la bienveillance continûment, Julianna qui pétille et qui a trouvé le temps de me donner quelques cours de brésilien, Vincent et ses réponses toujours exhaustives, Pierre et sa volonté de transmettre, Arthur et sa largesse comfortable, Noura et nos longues et passionnées discussions. Je ne suis pas arrivé au MAP5 seul. Ceux de ma génération ont aussi su se révéler exceptionnels; Pierre-Louis mon "frère de thèse" en qui j'ai trouvé un ami sincère, Florian qui a réussi à finir sa thèse le premier avec brio - ça a été un plaisir de partager les responsabilités du séminaire des Doctorants ainsi que les droits d'admin du discord avec toi -, Antoine M. acolyte avec qui je passe la ligne d'arrivée et, pour finir, Remi L. et son ineffable sens de l'humour chaotique. Puis viennent ceux que j'appellerai toujours les "nouveaux" même si pour certains, ils sont déjà presque à la fin de leur aventure ! Zoé qui dès qu'elle se laisse déconcentrer émane une gentillesse pure, Guillaume à qui je passe le flambeau et qui avec son esprit capable, j'en suis sûr, le transformera en brasier, Loïc et nos discussions absurdes et hautes en couleurs qui m'ont toujours apporté le sourire, Alexander grâce à qui je me suis enfin laissé aller à faire des étirements et poiriers divers dans le bureau, Antoine S. mon voisin de bureau complice, Martin, nouveau compagnon danseur, Medhi qui m'a sorti faire des tours en courant. Mais aussi Yen, Safa, Eloi, Herb, Diala, Sonia, Ivan, Adrien, Remi B., Michel, Ariane, Mariem, Chabane, Charlie, Keanu ...

Je n'aurais pas traversé ces trois ans sans mes proches. J'ai la chance d'avoir une famille qui me soutient. Ma mère et les pauses dans son jardin fleuri, les soirées de travail chez mon père, ainsi que ma retraite de redaction chez ma tante Monika. Je pense aussi bien sûr à Carmen, Hadrien, Jola, Feliks et Thomas.

Pour ma famille de coeur, ceux avec qui je ne partage pas le sang mais plutôt les rires et les larmes. Vous le savez, je vous aime et je ne pourrai pas l'exprimer assez ici. Capucine, malgré notre récente rencontre, tu as éclairé et rendu ma fin de thèse si légère à vivre. Je suis chanceux de t'avoir rencontrée. Nolwen, tu m'as accompagné pendant les années

d'inconscience en début de thèse. Tu es une personne sur laquelle j'aime pouvoir compter et je t'en suis reconnaissant. Salomé, tu as su trouver une place de choix dans ma thèse; grâce à tes illustrations celle-ci brille de mille feux ! Il va sans dire que je n'aurais pas pu le faire sans toi ! Raph, mon ami de toujours, en nous rencontrant nous nous sommes promis de faire une thèse. Voilà chose faite! Merci pour tant de choses. Maé et Saïd, vous êtes mes piliers sur lesquels je sais que je peux compter. Camille Souffron et Pierre Roux je garde en mémoire nos discussions folles et passionnantes qui ne manquent jamais de m'inspirer. Je n'oublie pas les danseureuses qui ont été une famille et mon ressourcement ! En particulier: Florian, Rosalia, Nathasha, Blandine, Espérence et d'autres que je n'aurai pas la place de citer.

A la mémoire des grains de café sacrifiés, ce manuscrit en a dépendu.

Contents

List of Figures

0

Notation & Lecture guide

Notations

Mathematical notation

Spaces

- Ω is the definition domain of the image. In general, images are seen as functions defined on a subset of $\mathbb{R}^d$, where $d \in \{2, 3\}$ is the dimension. Depending on the context, Ω can be seen as the image set of pixels or voxels. H, W, D are respectively the number of pixels in height, width and depth. (*i.e.:* $\Omega = H \times W$ in 2D and $\Omega = H \times W \times D$ for 3D pictures as the ones acquired from MRI.)
- $\mathcal{I}$ is the space of images $I : \Omega \to \mathbb{R}$.
- G is the group of deformation. By default, G is set to be Diff the group of $\mathscr{C}^1(\Omega, \Omega)$ diffeomorphisms.
- H is any Hilbert space, and L^2 is the space of 2 times integrable functions.

Operators:

- $v \cdot w$: where $v, w : \Omega \mapsto \mathbb{R}^d$ and $\cdot$ is the **Euclidean scalar product**.
- $g \triangleright m$: where $g \in G$ a group and $m \in M$ a Manifold, $g \triangleright m$ is the **action** of g on m
- $v \diamond f$: where v is a vector fields, and f a differentiable function, $\diamond$ is a **differential operator**.

PRODUCTS:

- We will note $\langle \bullet, \bullet \rangle_A$ the **scalar products** with the corresponding **norms** $\|a\|_A^2 = \langle a, a \rangle_A$ where A is a functional space
- We will note $(\bullet|\bullet)$ the **dual product of an Hilbert space** H given by the Riesz representation theorem, such as for every continuous linear functional $g \in H'$, there exists a unique vector $f_g \in H$ written as $g(x) = \langle x, f_g \rangle_H = (f_g|x)$ for all $x \in H$.

DERIVATION:

- We will use $\partial_x f = \frac{\partial f}{\partial x}$ for the **the partial derivative in the direction** x to lighten notations, where f is a differentiable function. Depending on the context we can use one or the other.
- We will use $\dot{f} = \partial_t f = \frac{\partial f}{\partial t}$ for the **time derivative** to lighten notations, where f is a time differentiable function. Depending on the context we can use one or the other.

LETTERS:

- $\mathcal{L}$: Lagrangian (*c.f.,* Section 2.3.1)
- $\mathscr{L}$: Lie derivative (*c.f.,* Section 2.1)
- $\mathbb{L}, \mathbb{K}$: differential operator for RKHS, L, K corresponding kernels (*c.f.,* Section 2.2.3.c)

DIFFEOMORPHISMS:

- $\mathbf{v}$ is a temporal vector field, $\mathbf{v} = (v_t)_{t\in[0,1]}$.
- $\varphi^{\mathbf{v}}_{t,s}, 0 \leq t < s \leq 1$ is the diffeomorphism integrated from $\mathbf{v}$, from t to s.
- $\varphi^{\mathbf{v}}_t \doteq \varphi^{\mathbf{v}}_{0,t}$. When only one value is passed, the integration is implied as starting from 0. Similarly, $\varphi^{\mathbf{v}} \doteq \varphi^{\mathbf{v}}_1$ is the complete forward integration from 0 to 1.

You can find more details on the definition in Section 2.2.3.b.

In text

- `functions` and `Classes` are always written with this font and the source code can be found in one of the Github repositories Demeter_metamorphosis or gliomaSegmentation_TDA
- Math using the 'mathtt' font is the discretisation analogue of the same continuous function. For example, the temporal image is $I : [0, 1] \to \mathbb{R}$ and its discretisation with N step is $\mathtt{I} : [\![\mathtt{0}, \mathtt{N} - \mathtt{1}]\!] \to [\mathtt{0}, \mathtt{1}]$ (the image is normalised to have values in between 0 and 1).

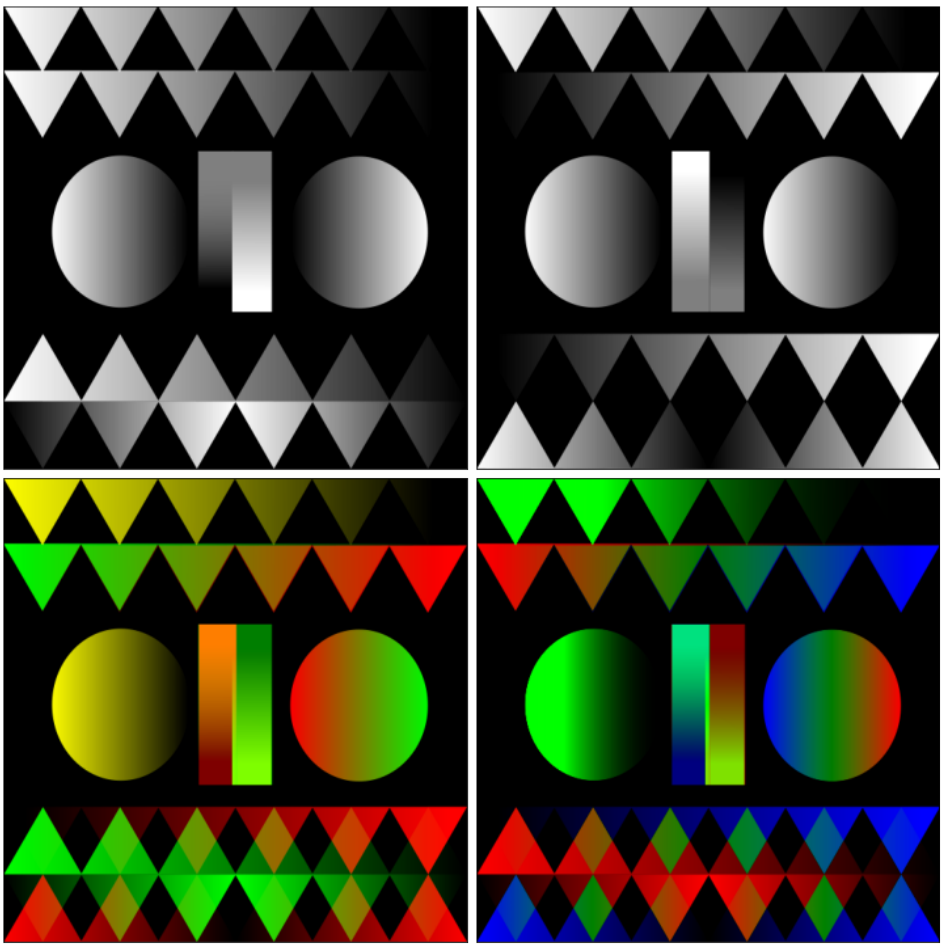

Figure 1: Image comparison color code Top- Two images to compare. Bottom- (left) 'old' colour code used in Chapter 3, (right) 'new' colour code used in Chapter 2.

Colour Code

Numerous techniques have been developed for comparing images, among which one of the most efficient is the method of displaying two images alternately. However, due to the limitations of the pdf format, this technique cannot be used for image comparison. In the case of grayscale images, it is possible to utilize the RGB format for comparative analysis of images. At the outset of my thesis, I employed a rudimentary technique which involved assigning each image to the first (red) and second (green) colour channels. The resulting images are depicted in the lower-right corner of Figure 1, with yellow representing similar pixel values and green and red representing diverging pixel values. However, I observed confusion among the audience while presenting my results, which prompted me to devote more effort to image comparison. This led me to develop an improved technique that results in images similar to the one depicted in the lower-left corner of Figure 1, where similar pixel values are displayed in green and diverging ones in red and blue, with slightly enhanced values for low and non-zero values. It should be noted that the figures presented in Chapter 3 were

created using the old convention of assigning red and green channels to the first and second images, respectively. However, in Chapter 2, the improved technique of displaying similar pixel values in green and diverging ones in red and blue was employed. Unfortunately, due to time constraints, it was not feasible to redo all the experiments with the new convention. As a result, a mix of both conventions may be present in the figures presented throughout this thesis. Nonetheless, the comparative analysis of images remains valid and informative for the purposes of this research.

A vulgarisation comic

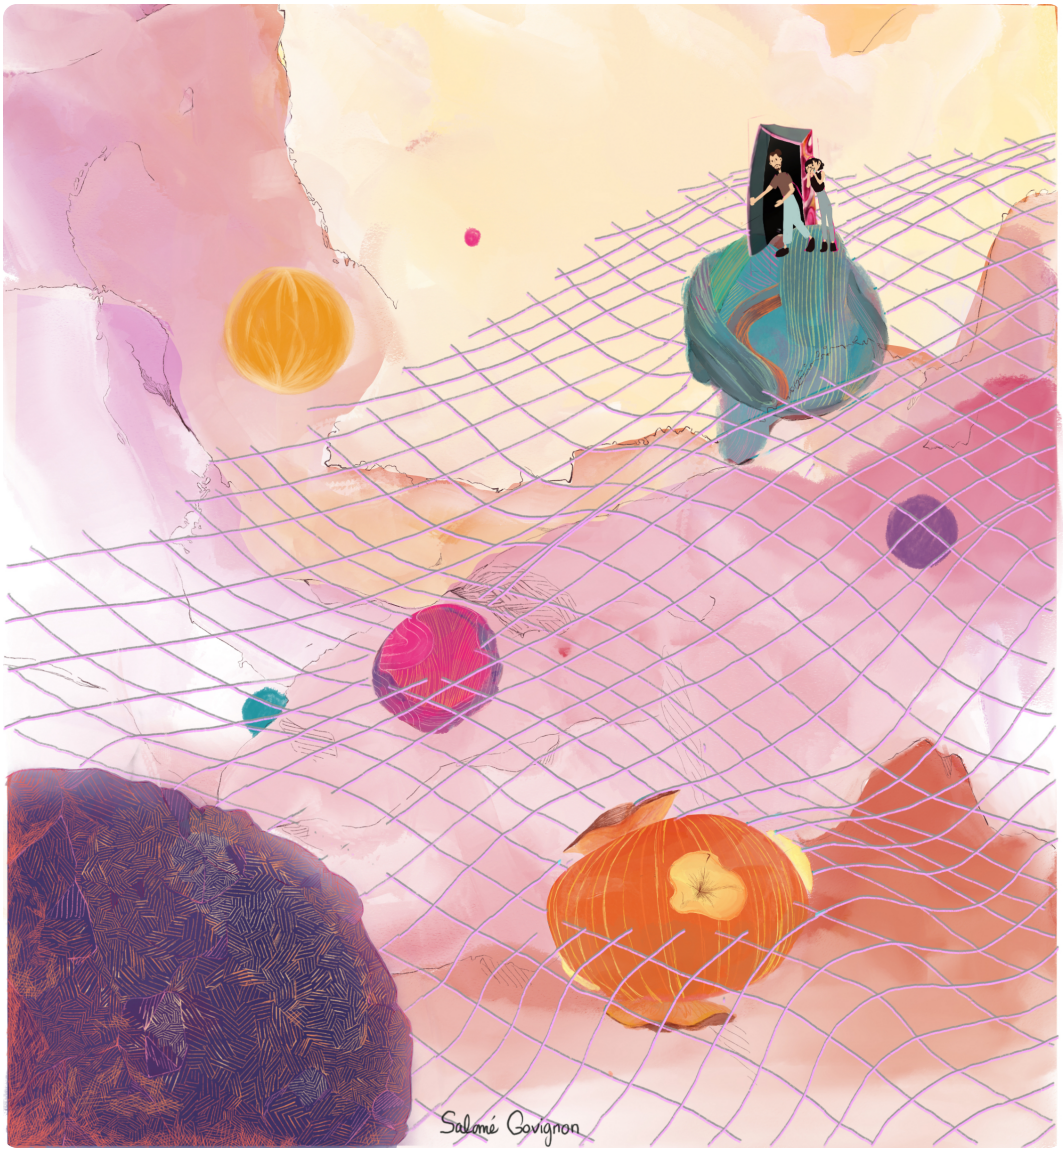

The primary duty of a researcher is to generate knowledge that is verifiable and reproducible. However, it is also important for the researcher to communicate their research findings to the wider public. As someone with a passion for comic books, I decided to try my hand at making my research more accessible by creating two short stories that summarise key ideas of my thesis. You can find these stories in sections 1.7 and 2.6. I was fortunate to have the help of **Salomé Govignon**, an illustrator and friend, who lent a touch of magic to my work through her delicate drawings. Her contribution was invaluable in bringing the two stories to life, and I am grateful for her assistance.

1

Introduction - fr

Ce projet de doctorat se trouve à l'intersection de plusieurs disciplines. Nous appliquons nos travaux à la recherche médicale sur les glioblastomes, qui sont un genre de tumeurs du cerveau. À partir de données radiologiques, nous voulons inférer une forme moyenne en fonction de leurs positions dans le cerveau. Pour cela, nous utilisons les mathématiques appliquées, plus particulièrement la théorie des espaces de formes, domaine qui est lui-même à l'intersection de plusieurs champs des mathématiques et de la physique. Il se base sur la géométrie différentielle mais possède des liens forts avec la dynamique des fluides, tout en dépendant beaucoup de techniques d'optimisation. Dans cette introduction, nous prenons le temps de détailler la pertinence de l'application et introduisons les concepts d'espaces de formes ainsi que du recalage.

1.1 Motivation médicale et aperçu des données

1.1.1 À propos des Glioblastomes

GÉNÉRALITÉS SUR LES TUMEURS. Une tumeur est une masse de cellules anormales qui se développe dans le corps. Ces cellules anormales se multiplient de manière incontrôlée, formant une masse qui peut envahir les tissus et les organes environnants. Les tumeurs peuvent être *bénignes*, ce qui signifie qu'elles ne sont pas cancéreuses et ne se propagent pas à d'autres parties du corps. Si une tumeur n'est pas bénigne, on dit qu'elle est *maligne*. Les premières tumeurs peuvent gêner la fonction corporelle en redirigeant les nutriments à leur seul bénéfice, privant potentiellement d'autres tissus essentiels. Les cancers (avec une emphase sur le pluriel) peuvent prendre de nombreuses formes et sont extrêmement hétérogènes en termes de forme, de composition et de symptômes. En effet, leur origine est souvent due à une mutation malheureuse survenue pendant la division cellulaire, altérant leurs fonctions normale.

Malgré cette variabilité, on peut distinguer plusieurs parties d'une tumeur, notamment la principale *masse tumorale*, qui est composée de cellules anormales, et le tissu environnant, qui peut être normal ou peut contenir un mélange de cellules normales et anormales. La tumeur peut également avoir une couche de cellules appelée capsule tumorale, qui entoure la principale masse tumorale et la sépare du tissu sain. La tumeur peut également avoir des vaisseaux sanguins et lymphatiques, qui peuvent l'aider à se développer et à se propager à d'autres parties du corps. À l'intérieur de la masse tumorale, une *nécrose* peut se produire. La nécrose est un type de mort cellulaire qui se produit lorsque les cellules sont endommagées ou blessées au-delà de toute réparation. Contrairement à d'autres formes de mort cellulaire, comme l'apoptose, la nécrose est un processus chaotique et incontrôlé qui peut être nuisible pour le corps. La nécrose peut se produire dans n'importe quel type de tissu, mais elle est le plus souvent associée aux dommages tissulaires causés par une blessure ou une infection. Lorsque la nécrose se produit, les cellules affectées gonflent et éclatent, causant des *œdèmes* en libérant leur contenu dans les tissus environnants, ce qui peut causer une inflammation et des dommages aux cellules voisines. Dans les cas graves, la nécrose peut conduire à la mort des tissus ou des organes.

Aux frontières de la tumeur cancéreuse, les cellules sont plus actives et peuvent se propager dans les tissus environnants de manière diffuse, dans ce cas, on parle de *cancer infiltrant*. Le mélange de cellules saines et cancéreuses à la frontière rend ce type de cancer difficile à traiter chirurgicalement. Lorsque ces tumeurs atteignent les vaisseaux sanguins ou le système lymphatique, elles peuvent également se propager à d'autres parties du corps, ce

qui les rend plus agressives et plus difficiles à contrôler. Le traitement du cancer infiltrant implique souvent une combinaison de chirurgie, de radiothérapie et de chimiothérapie pour tuer les cellules cancéreuses et les empêcher de se propager.

PLUS SUR LES GLIOMES Il n'y a que deux types de cellules dans le système nerveux central (SNC) : les *neurones* qui sont considérés comme le support de la cognition et les cellules *gliales* qui sont toutes les autres cellules du SNC qui répondent à tous les besoins des neurones. Les neurones ont une grande variété de formes et de fonctions, par exemple les astrocytes fournissent des nutriments, régulent le flux sanguin et protègent les neurones des anticorps.

Les tumeurs cérébrales sont le plus souvent des *gliomes*, un terme général utilisé pour décrire les tumeurs cérébrales primaires. Les tumeurs neuronales sont très rares et sont surtout bénignes. En effet, les neurones ont un processus de reproduction complexe très différent des autres cellules. Les gliomes prennent leur origine dans les cellules gliales et leurs noms spécifiques proviennent du nom de leur cellule (*e.g., : astrocytome*) [Hanif et al., 2017]. Le glioblastome multiforme est le type d'astrocytome primaire le plus grave et le plus fréquent. Il représente plus de 60% de toutes les tumeurs cérébrales chez l'adulte [Rock et al., 2012]. Leurs causes ne sont pas bien comprises, l'exposition à des doses élevées de rayonnement ionisant étant le seul facteur de risque confirmé [Hanif et al., 2017]. Le traitement souvent recommandé est une chirurgie précise qui peut améliorer la qualité de vie du patient en réduisant la charge tumorale, et de fait aide à contrôler les convulsions ou même inverser les déficits neurologiques.

La localisation de la tumeur est un paramètre clé dans les soins des patients atteints de glioblastome car elle est corrélée aux caractéristiques démographiques ainsi qu'aux symptômes. De fait elle affecte le pronostic du patient et influence la prise en charge chirurgicale [Roux and et al., 2019].

De précédentes études sur la pathogenèse ont montré que l'emplacement le plus fréquent pour les GB sont les hémisphères cérébraux. 95% des GB apparaissent dans la région sus-tentorielle (partie supérieure de l'encéphale), tandis que seulement quelques-uns dans le cervelet, le tronc cérébral et la moelle épinière [Nakada et al., 2011]. À une échelle macroscopique, les GB sont assez hétérogènes en forme. S'ils sont irrégulièrement formés, ils se développent généralement dans la substance blanche. Il a été démontré que selon le lobe où la tumeur apparaît, les symptômes varient. Par exemple, les patients atteints de GB situé dans la région du lobe temporal montrent souvent des problèmes auditifs et visuels, tandis que ceux qui ont une tumeur dans le lobe frontal peuvent présenter un changement de personnalité [Hanif et al., 2017]. De plus, la distribution de l'œdème/nécrose entraîne différents effets secondaires chez le patient : une augmentation progressive de la taille de la tumeur et une augmentation de l'œdème entourant la tumeur entraînent un déplacement du contenu intracrânien, ce qui provoque des maux de tête. Une tumeur localisée dans des sites tels que le cortex éloquent, le tronc cérébral ou le ganglion basal ne peut pas être opérée et ces patients ont généralement un pronostic plus sombre [Mrugala, 2013].

La classification de l'organisation mondiale de la santé (OMS) des glioblastomes est un système utilisé pour classer et décrire différents types de glioblastome [Villa et al., 2018; Berger et al., 2022]. Cette classification est basée sur l'apparence microscopique des cellules cancéreuses, ainsi que sur les caractéristiques génétiques et moléculaires de la tumeur. C'est un outil important pour les cliniciens et les chercheurs, car elle permet de caractériser les caractéristiques spécifiques d'une tumeur particulière et de guider les décisions de traitement.

Le grade OMS des glioblastomes est généralement déterminé en analysant un échantillon de tissu prélevé lors d'une biopsie ou d'une intervention chirurgicale. Cependant, l'IRM peut également être utilisée pour identifier certaines caractéristiques du glioblastome qui peuvent être indicatives de son grade OMS. Les glioblastomes de haut grade ont tendance à être plus grands, avoir une forme irrégulière et diffuse, tout en étant très intenses sur les images d'IRM avec contraste, être associés à des zones de nécrose et infiltrer le tissu cérébral environnant. Il est important de noter que l'IRM seule n'est pas suffisante pour classer avec précision un glioblastome. Un échantillon de tissu doit être analysé par un pathologiste pour déterminer définitivement le grade OMS de la tumeur.

Le premier diagnostic est toujours basé sur la visualisation de la tumeur et est réalisé à l'aide de techniques d'imagerie radiologique, la plus courante étant l'IRM.

1.1.2 Imagerie par Résonance Magnétique pour la visualisation de cerveaux

Une machine d'imagerie par résonance magnétique (IRM) utilise des champs magnétiques puissants et des ondes radio pour créer des images détaillées de l'intérieur du corps. Le patient est allongé sur une table qui est déplacée dans la machine IRM, qui contient un grand et puissant aimant. L'aimant aligne les particules atomiques dans le corps, et des ondes radio sont utilisées pour décaler ces particules de leur alignement initial. Lorsque les particules reviennent à leur alignement initial, elles émettent un signal qui est détecté par la machine IRM et utilisé pour créer une image de l'intérieur du corps. Les images produites par une machine IRM peuvent être utilisées pour diagnostiquer une large gamme de conditions médicales.

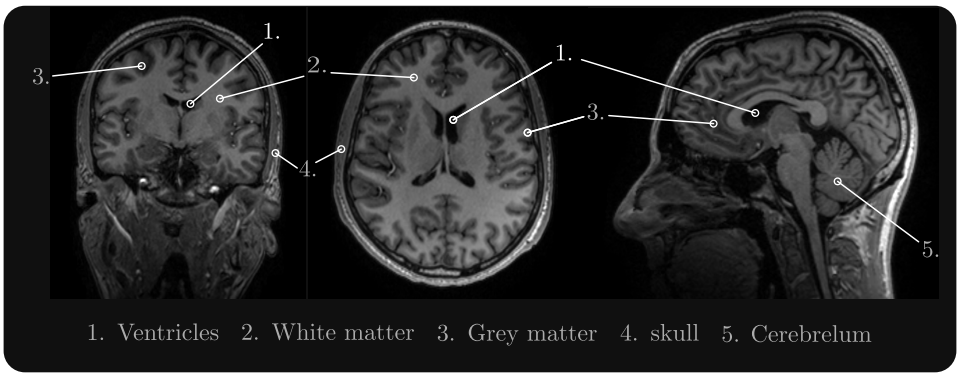

Figure 1.1: Glossaire succinct de l'anatomie du cerveau.

La figure 1.1 illustre une IRM T1 d'un cerveau sain avec le crâne. Le crâne est généralement retiré lors du pré-traitement, cependant, des résidus peuvent subsister. Les tissus cérébraux peuvent être divisés en quatre régions :

1. Les **ventricules** sont des chambres dans le cerveau, remplies de liquide céphalo-rachidien (LCR). Il y a quatre ventricules dans le cerveau, deux dans le cerveau (les ventricules latéraux gauche et droit) et deux dans le tronc cérébral (le troisième et le quatrième ventricule). Dans cette thèse, nous nous référerons principalement aux ventricules latéraux gauche et droit. La fonction principale des ventricules est de produire et de faire circuler le liquide céphalo-rachidien, un liquide clair et aqueux qui entoure

et amortit le cerveau et la moelle épinière. Il remplit plusieurs fonctions importantes, notamment la protection du cerveau et de la moelle épinière contre les lésions mécaniques, le maintien de l'équilibre correct des nutriments et des déchets dans le SNC, ainsi que la régulation de la pression à l'intérieur du crâne.

2. La **matière blanche** correspond aux grandes zones plus claires de la figure 1.1. Elle est composée de longs axones, entourés d'une substance grasse constituée de cellules gliales appelées myéline. La matière blanche doit son nom à l'apparence blanche de la gaine de myéline, qui contraste avec l'apparence grise des corps cellulaires des neurones qui constituent la matière grise. Notez que ces tons ne sont pas conservés par l'IRM (voir le paragraphe suivant).

3. La **matière grise** est principalement composée de corps cellulaires de neurones et de synapses et est principalement responsable du traitement et de l'interprétation des informations.

4. Le **cervelet** est une région plus petite située à l'arrière de la tête sous le cerveau. Il est constitué de deux hémisphères et est relié au tronc cérébral et à la moelle épinière. La partie extérieure contient des neurones et la partie intérieure communique avec le cortex cérébral. Il est responsable de la coordination des mouvements, de l'équilibre et de la posture. Il reçoit des entrées des systèmes sensoriels du corps, y compris les yeux, les oreilles et les muscles. Il utilise ces informations pour effectuer des ajustements de mouvement afin de maintenir l'équilibre et la précision. De nouvelles études explorent les rôles du cervelet dans la pensée, les émotions et le comportement social, ainsi que son possible rôle dans l'addiction, l'autisme et la schizophrénie.

Plusieurs séquences sont généralement nécessaires pour évaluer de manière adéquate un tissu, et la combinaison de séquences est appelée protocole d'IRM. En particulier pour le diagnostique des gliomes, les radiologues réalisent au moins quatre modalités d'IRM aux patients suspectés, qui aident à distinguer différents organes [Villanueva-Meyer et al., 2017; Ginsberg et al., 1998] :

- **T1** ou pondéré en T1, rend les graisses brillantes et l'eau sombre. En particulier, le liquide céphalorachidien dans les ventricules apparaît plus sombre. Sur la figure 1.1, on peut voir une première IRM pondérée en T1.

- **T1ce** ou pondéré en T1 avec injection de produit de contraste, est une IRM où le patient reçoit une injection d'agents intraveineux. Ces agents augmentent le contraste des tissus des zones où ils ont fui la barrière hémato-encéphalique vers les tissus interstitiels. Dans les gliomes diffus, l'augmentation de contraste est positivement corrélée au grade de la tumeur, bien que certains gliomes de haut grade et certains gliomes de bas grade puissent ne montrer aucune ou une amélioration minimale.

- **T2** ou pondérées en T2, chaque tissu a une valeur T2 intrinsèque, mais des facteurs externes (comme l'inhomogénéité du champ magnétique) peuvent diminuer le temps de relaxation T2, diminuant ainsi la valeur. Les liquides sont lumineux et les graisses sont intermédiaires.

- **T2-FLAIR** ou inversion récupérée pondérée en T2 des fluides est un balayage T2 avec un processus qui supprime le signal du liquide céphalo-rachidien dans les images qui en résultent. Les tissus cérébraux sur les images FLAIR apparaissent similaires

aux images pondérées en T2 avec de la matière grise plus lumineuse que de la matière blanche. Ainsi, contrairement à la pondération en T1, le liquide céphalo-rachidien et les ventricules apparaissent sombres. Le glioblastome induit des hyperintensités T2 corticales cérébrales. FLAIR a une intensité élevée dans l'œdème péri-tumoral (vasogénique et infiltrant) ainsi que dans les lésion de la matière blanche et la gliose. À noter cependant que dans de nombreux gliomes, l'anomalie du signal hyper-intense T2/FLAIR peut ne pas être dissociable de la lésion tumorale principale.

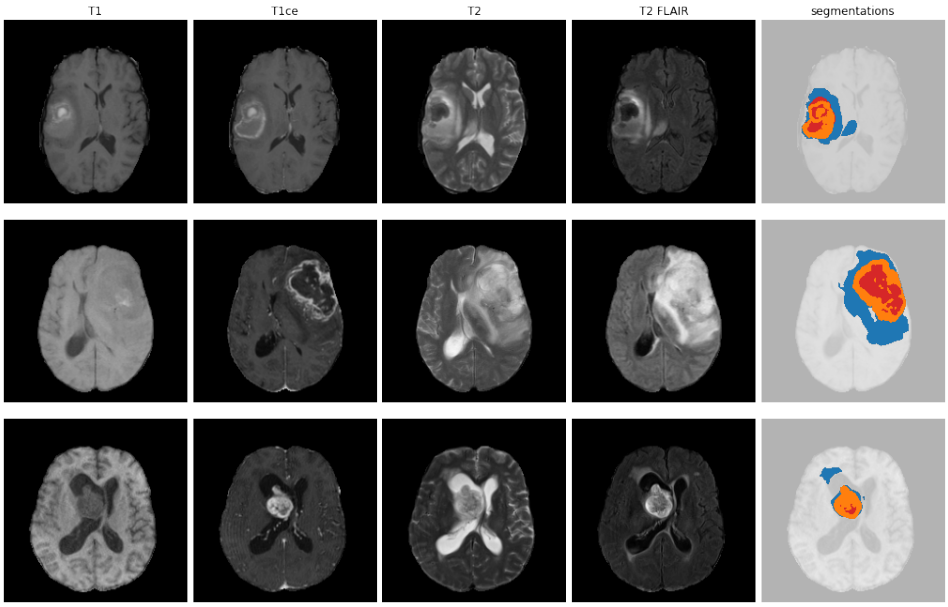

Figure 1.2: Exemples de segmentation de glioblastomes : Chaque ligne contient des coupes d'images IRM dans quatre modalités différentes pour le même patient à la même date. On peut comparer les différences entre les modalités d'image et la variation de couleur inter-sujets au sein des mêmes modalités. Légende de la segmentation : rouge - Noyau nécrotique (TC) ; orange - Tumeur active (ET) ; bleu - œdèmes (WT).

Nous avons vu que chaque modalité d'imagerie est utile pour visualiser des structures cérébrales spécifiques. Cependant, il est important de noter que la qualité et l'apparence des IRM peuvent varier considérablement en raison d'un certain nombre de facteurs, notamment la machine IRM spécifique utilisée et les conditions d'acquisition. De plus, la pathologie spécifique observée peut également avoir un impact sur l'apparence des images. Par conséquent, il est courant que les images IRM de la même modalité obtenues à partir de différents patients présentent des différences significatives, voir Figure 1.2. Dans la dernière colonne, on peut voir un exemple de segmentation du glioblastome. Une segmentation en trois classes a été acceptée [Baid and et al., 2021]. Cette figure montre une segmentation manuelle effectuée par des experts. Les annotations comprennent la tumeur en prise de contraste avec le gadolinium (ET), le tissu envahi, l'œdéme péri-lésionnel (WT) et le cœur tumoral nécrotique (TC). Notez tout de même que WT veut dire 'Whole tumour' ou "toute la tumeur" en anglais, comme pour le challenge BraTS2021. Ainsi, dans ce document nous allons l'utiliser pour désigner les 'tissus tumoraux qui ne sont pas la nécrose'.

1.2 Des Atlas pour cartographier le cerveau

Avec l'IRM, nous disposons d'un outil pour visualiser des cerveaux individuels, cependant, pour comparer les images prises par des radiologues à une référence, un cerveau représentatif est nécessaire. En général, les *atlas* sont un outil pour étudier la forme des objets. Quelques définitions de ce terme coexistent en fonction de leurs domaines respectifs. Il se trouve que dans cette thèse, nous sommes à l'intersection d'au moins trois d'entre eux. Nous donnons ici des définitions courtes et nous en détaillerons deux dans les sous-sections suivantes :

- **Les Atlas de l'espace de formes** est l'une des notions principales de cette thèse et sera expliquée plus en détail dans la section 1.2.2. Un espace de formes est une représentation mathématique des formes, où chaque point dans l'espace représente une forme différente. Dans un espace de formes, la distance entre deux points reflète à quel point les formes sont similaires. En étudiant les formes dans un espace de formes, les chercheurs peuvent analyser et comparer les formes, et comprendre comment elles changent et évoluent au fil du temps. La théorie des espaces de formes est un outil important dans des domaines tels que l'imagerie médicale et la biomécanique. Dans cette thèse, à l'exception de cette section, l'espace de formes fait référence à l'*espace de formes difféomorphique*. Il s'agit d'un type spécifique d'espace de formes qui est utilisé pour représenter des formes qui sont continûment déformables l'une dans l'autre, en préservant leur topologie. Cela signifie que les formes sont transformées de manière à ce que leur structure sous-jacente soit préservée et que la transformation soit inversible. Dans l'espace de formes difféomorphique, les formes sont représentées par des points, et les transformations difféomorphes correspondent à des chemins entre les points. Une fois que cet espace est construit, on peut estimer une forme moyenne appelée *template*.
- En biologie, les atlas font souvent référence aux **Atlas statistiques**. Un atlas statistique est une collection de cartes ou d'images qui montrent la distribution d'une variable ou d'une caractéristique particulière dans une population définie. Ces cartes peuvent être utilisées pour visualiser et analyser des données dans un contexte spatial, permettant aux chercheurs de voir des motifs et des tendances qui ne seraient pas apparents à partir de données tabulaires seules. Les atlas statistiques sont souvent utilisés dans le domaine médical pour visualiser et comparer les données liées à la santé et à la maladie. Par exemple, un atlas statistique du cerveau pourrait montrer la distribution des structures cérébrales ou de la fonction cérébrale dans une population, ou la prévalence d'un trouble cérébral particulier dans différentes régions du cerveau. Nous détaillerons cette méthode dans la sous-section 1.2.1.
- En mathématiques, **l'Atlas topologique** est utilisé pour la description d'objets abstraits comme les variétés. Il est composé d'individus *cartes* décrivant localement la variété. Les cartes juxtaposées doivent décrire leurs régions de chevauchement de manière équivalente. Bien que ce soit une notion importante en géométrie différentielle, nous n'aurons pas besoin de travailler directement avec cette notion.

Si le traitement d'un patient est toujours à traiter au cas par cas, il est nécessaire d'avoir une représentation d'un sujet sain. Les atlas peuvent servir à comparer des sujets sains et malades. Par exemple, Le modèle cérébral de l'Institut neurologique de Montréal (MNI) est un système de référence standard utilisé dans la recherche en neuroimagerie pour aligner et comparer des images cérébrales de différents individus Fonov et al. [2009, 2011]. Il est basé

sur un ensemble d'images cérébrales qui ont été moyennées pour créer un cerveau "moyen", qui représente l'anatomie moyenne d'un grand groupe de personnes.

Le modèle de cerveau MNI est créé en alignant un grand nombre d'images cérébrales à l'aide d'un logiciel spécialisé, puis en moyennant les images pour créer un cerveau composite qui représente l'anatomie moyenne du groupe. Ce modèle est ensuite utilisé comme système de référence pour aligner d'autres images cérébrales, ce qui permet aux chercheurs de comparer l'anatomie et la fonction cérébrales entre différents individus. Le modèle de cerveau MNI est largement utilisé dans le domaine de l'imagerie cérébrale et est devenu un système de référence standard pour de nombreux chercheurs. Il est particulièrement utile pour aligner des images cérébrales provenant de différentes modalités d'imagerie, et pour comparer des images cérébrales provenant de différentes populations ou groupes. Dans la littérature, il est souvent appelé l'Atlas MNI de manière trompeuse. En effet, il peut être obtenu à partir d'Atlas d'espace de forme (mais n'en est qu'une partie, voir la section 1.2.2) et est souvent utilisé comme arrière-plan dans les Atlas statistiques.

1.2.1 Les Atlas Statistiques pour l'études du positionnement des Glioblastomes

Les glioblastomes peuvent apparaître dans n'importe quelle partie du cerveau, mais ils sont le plus souvent situés dans les hémisphères cérébraux, qui sont les moitiés gauche et droite du cerveau. Comme nous l'avons mentionné précédemment, l'emplacement du glioblastome peut influencer les symptômes ressentis par un patient, ainsi que les options chirurgicales et de traitement disponibles. Par exemple, un glioblastome situé dans les hémisphères cérébraux peut causer des symptômes tels que des maux de tête, des nausées, des vomissements et des convulsions, tandis qu'un glioblastome dans le tronc cérébral peut causer des symptômes tels que des difficultés d'équilibre, une faiblesse ou des changements dans la parole ou la vision [Bilello et al., 2016; Parisot et al., 2016]. Simpson et al. [1993] n'a observé aucune différence de survie pour les différentes tailles de tumeurs et a constaté que les patients atteints de tumeurs du lobe frontal survivaient plus longtemps que ceux ayant des lésions du lobe temporal ou pariétal, concluant que la localisation est un indicateur crucial du pronostic.

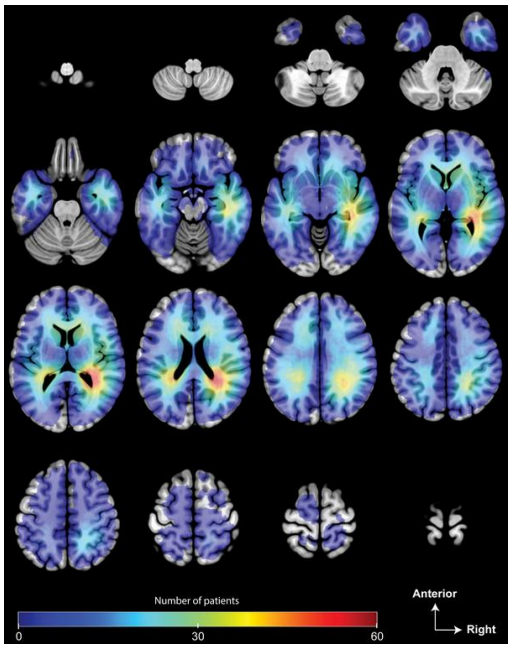

Figure 1.3: Emplacement et fréquence des glioblastomes de type sauvage pour l'isocitrate déshydrogénase (IDH) (n = 392). La carte de fréquence de couleur illustre l'emplacement et le nombre de patients atteints de glioblastome de type sauvage pour l'IDH. Les images sont affichées selon la convention d'affichage neurologique. Figure issue de Roux and et al. [2019].

Pour obtenir de tels résultats, les neurochirurgiens doivent étudier l'emplacement moyen des glioblastomes et donc construire un atlas statistique : une carte où l'arrière-plan est un modèle cérébral moyen et sur laquelle

chaque voxel est associé à sa fréquence d'apparition de gliome [De Witt Hamer et al., 2013; Bilello et al., 2016; Parisot et al., 2016; Roux and et al., 2019; Sagberg et al., 2019]. En d'autres termes, pour chaque région, ils portent la probabilité estimée d'apparition de la tumeur.

Nous allons prendre l'exemple de l'étude récente de Roux and et al. [2019] pour expliquer la méthode de construction d'un atlas statistique. Les auteurs ont construit un tel atlas pour aider au pronostic des patients en fonction de l'âge, du sexe et d'autres données cliniques. La méthode de construction peut être divisée en trois étapes :

1. Tout d'abord, un expert a manuellement segmenté les composantes de nécrose et de rehaussement des glioblastomes à l'aide de l'interface graphique utilisateur *Multi-image Analysis GUI Mango*, qui permet une visualisation et une transformation simples des images en 3D. **2.** Ensuite, ils ont utilisé le modèle MNI comme référence de cerveau sain pour enregistrer chaque image de patient malade dessus. Pour ce faire, ils ont utilisé des méthodes de recalage difféomorphique Ripollés et al. [2012] basée sur l'algèbre de Lie exponentielle (voir Section 2.1.2) et le masquage de coût. **3.** Avant de construire l'atlas, la segmentation et l'enregistrement ont été vérifiés par un autre expert formé, et les auteurs affirment qu'aucune correction n'a été nécessaire. Enfin, toutes les segmentations ont été superposées dans toutes les régions d'intérêt, obtenant ainsi une carte de fréquence 3D basée sur différents paramètres.

Les chercheurs ont découvert que les glioblastomes IDH de type sauvage étaient le plus souvent situés dans la substance blanche sous-corticale des zones sous-ventriculaires (*i.e :* fine couche de cellules située près des ventricules latéraux) dans les deux hémisphères du cerveau. Il y avait des différences dans l'emplacement du glioblastome en fonction de la zone sous-ventriculaire concernée. Chez les patients atteints de glioblastomes près des ventricules, les tumeurs étaient plus susceptibles de se trouver dans la corne antérieure, l'atrium droit et gauche, et la corne temporale droite et gauche des ventricules latéraux. Ils ont également établi des liens entre les symptômes observés et l'emplacement. Par exemple, ils affirment que, parmi les 81 patients chez qui le glioblastome a été détecté en raison de symptômes de pression intracrânienne élevée, il était plus probable que les glioblastomes soient situés dans les lobes frontaux droits.

L'emplacement est donc un facteur très important. Cependant, des analyses ultérieures ont montré que la relation entre l'emplacement de la tumeur et l'incidence des convulsions préopératoires dépend du grade de malignité du gliome [Pallud et al., 2016]. De plus, les médecins cherchent également à établir des liens avec la forme de la tumeur mais manquent d'un test facile à utiliser pour leur hypothèse. Par exemple, les cellules tumorales infiltrantes ont tendance à suivre les vaisseaux sanguins et sont plus susceptibles de se propager à d'autres régions du cerveau. Pour ce faire, on pourrait construire un atlas de l'espace de forme ("Shape-space Atlas").

1.2.2 Introduction aux espaces de formes

La notion d'espace de formes n'est pas très intuitive pour le lecteur non initié. Dans cette section, on commence par expliquer deux espaces de formes classiques et élémentaires pour introduire l'espace de formes difféomorphique, qui est notre objet d'intérêt principal. Si vous êtes plus généralement intéressés par les espaces de formes, vous pouvez vous référer à Bauer et al. [2014].

1.2.2.a L'espace de Kendall

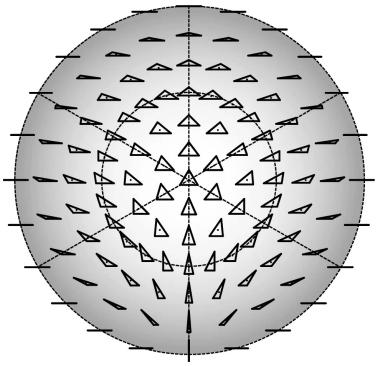

Figure 1.4: L'espace de forme de Kendall pour les triangles en deux dimensions. Le diagramme montre la vue depuis le "pôle nord" correspondant à un triangle équilatéral (au centre) sur un hémisphère. Le cercle pointillé extérieur correspond à l'équateur et contient les triangles colinéaires, où les trois sommets se trouvent sur une ligne droite. Les six méridiens représentés par des lignes pointillées droites contiennent des triangles isocèles. Figure de Klingenberg [2020]

L'espace de Kendall est un espace mathématique utilisé pour représenter une forme comme une collection de points et toutes ses variations naturelles possibles. Il est nommé d'après le mathématicien David G. Kendall, qui a introduit le concept dans son travail sur l'analyse de la forme [Kendall, 1984; Kendall et al., 2009]. Dans cet espace, chaque élément est une forme.

Par exemple, Klingenberg [2020] a réalisé une étude très intéressante sur la forme des ailes de la drosophile. Dans cet article, il donne l'exemple de l'espace de formes pour les triangles qui peut être représenté sur une sphère, donc en deux dimensions (voir Figure 1.4). Plus précisément, il s'agit du quotient de l'espace des configurations de trois points dans $\mathbb{R}^2$ (donc avec une dimension de base de 6) par les similarités de $\mathbb{R}^2$ (un groupe avec une dimension de 4 : 2 paramètres de translation, un de rotation, un d'homothétie), ce qui explique la dimension de 2. Deux points voisins représentent deux formes ou triangles différents qui sont proches, au sens où l'un est égal à l'autre par une petite déformation. Pour des ensembles de points plus complexes, les formes sont réparties sur une hypersphère, fournissant une représentation des formes pratique. Cela fournit un cadre pour effectuer des manœuvres statistiques telles que l'estimation des formes moyennes ou la caractérisation de la variation des formes autour de ces moyennes. Cela peut être fondamental pour les applications biologiques de morphométrie géométrique.

Cependant, l'espace de Kendall convient uniquement à la comparaison d'un ensemble de points par l'analyse de Procruste et ne prend pas en compte l'arrière-plan. Il n'est donc pas adapté à nos applications d'imagerie.

1.2.2.b Espaces de formes à contour actifs

Un espace de formes plus adapté aux images pourrait être l'espace de formes à contours actifs [Cootes et al., 2000], qui est basé sur des modèles de contours actifs (parfois appelés "snakes") [Kass et al., 1988]. Les modèles de contours actifs sont des méthodes visant à délimiter les contours d'un objet dans une image en ajustant une courbe donnée à ses bords.

(a) Exemples de visages annotés de landmarks (points de repères)

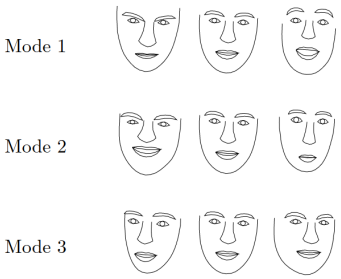

(b) Variations des trois paramètres principaux d'un modèle de visage, d'une variance de ± 3.

Figure 1.5: Espaces de formes à contours actifs appliqués aux visages. Figures issue de Cootes et al. [2000]

Ces méthodes ne résolvent pas tout le problème de la détection de contours dans les images, car la méthode nécessite la connaissance de la forme de contour souhaitée au préalable.

L'espace de formes à contours actifs est appris à partir d'une collection de courbes obtenues par l'application d'une méthode de contour actif, puis les axes sont déduits à l'aide d'une analyse en composantes principales (PCA). Tout d'abord, un ensemble de repères connectés définis par un utilisateur est augmenté par une méthode de contour actif, puis grossièrement aligné dans un cadre de coordonnées commun (voir Figure 1.5a). Ensuite, une réduction dimensionnelle est appliquée à l'ensemble des repères en utilisant une PCA. Dans ce processus, on récupère les axes principaux et donc une représentation de l'espace de forme. Une fois que le modèle est appris, on peut représenter la variation de forme sur ses différents axes ou modes (voir Figure 1.5b). Ces 'modes' expliquent la variation globale due aux changements de posture 3D, qui provoquent le mouvement de tous les points de repère les uns par rapport aux autres. Les modes moins significatifs provoquent des changements plus petits et plus locaux.

Cependant, ils sont dérivés directement des statistiques d'un ensemble d'apprentissage et ne séparent pas toujours de manière évidente la variation de forme. Même s'il est considéré comme un 'espace de formes d'image' dans la littérature, c'est un espace de formes pour les repères avec une métrique L^2 utilisant des techniques de contour actif pour trouver les points. De plus, ces modèles ne sont pas évidents à généraliser en 3D car une courbe deviendrait une surface, nécessitant l'utilisation de la triangulation. Plus important encore, l'arrière-plan n'est pas pris en compte au cours du processus, ne permettant pas d'avoir une carte de déplacement directe entre les points de cet espace.

Nous devrions également noter le modèle d'apparence active (active appearance model en anglais) [Edwards et al., 1998] qui correspond à un modèle statistique de la forme et de l'apparence d'un objet dans une nouvelle image en prenant en compte l'arrière-plan. Un ensemble d'images avec des repères est utilisé pendant la phase de formation pour construire le modèle. L'algorithme utilise des techniques des moindres carrés pour faire correspondre rapidement deux nouvelles images.

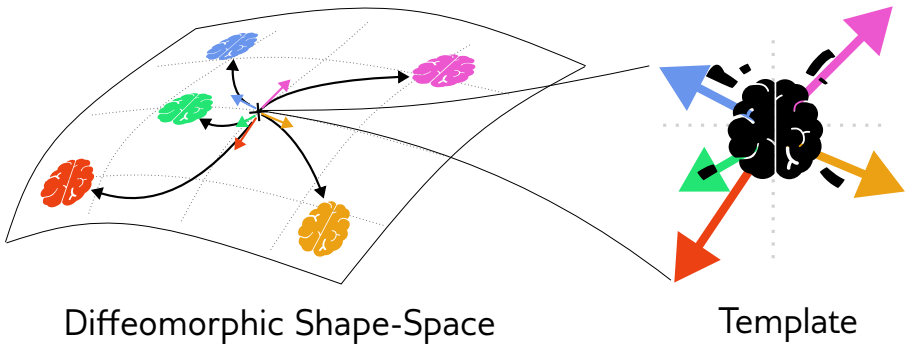

Figure 1.6: Vue schématique d'un Espace de Forme difféomorphique. Chaque cerveau représente un point de données, la surface est la variété de Riemann sur laquelle se trouvent les déformations (flèches noires). Les déformations sont entièrement décrites par leur vitesse initiale (flèches colorées). En trouvant le centre de cet espace, on peut déduire le modèle de référence (cerveau noir).

1.2.2.c Espace de Formes Difféomorphique

L'espace de formes difféomorphique trouve ses racines dans le travail de D'Arcy-Thompson [Thompson, 1942] et a été développé ultérieurement dans la théorie des motifs de Grenander [Grenander, 1993]. La spécificité de ce cadre est que la déformation met en correspondence la forme observée à la forme modèle, appelée *template*. En effet, ces espaces sont construits à partir de la structure induite par des groupes de difféomorphismes agissant sur les formes par déformation. Le *template* peut être choisi fixe ou être estimé, dans les deux cas, il prends le rôle de centroïde de l'espace. De plus, nous supposons que toutes les formes dans cet espace peuvent être assorties de transformations bijectives. En effet, théoriquement, on défini des difféomorphismes sur tout le domaine, ce sont donc des fonctions continues qui cartographient non seulement les objets mais également l'arrière-plan ou l'espace ambiant.

Ainsi, le modèle est extrêmement polyvalent, les modèles récents de templates déformables se concentrant spécifiquement sur les déformations représentées par des difféomorphismes agissant sur des repères, des courbes, des surfaces, des images ou même un mélange de ces objets. Plus précisément, l'espace des difféomorphismes induit une métrique riemannienne à droite invariante sur l'espace des formes de l'espace ambiant, donnant un sens de distance entre les objets. [Dryden and Mardia, 2016; Miller and Younes, 2001; Trouvé and Younes, 2003; Charon, 2013]

Cet espace des formes possède plusieurs propriétés intéressantes. En raison de sa nature de variété riemannienne de dimension infinie, la déformation minimale reliant deux points de l'espace peut être vue comme une géodésique, et la quantité de déformation sert de métrique riemannienne sur l'espace. Cela permet l'application de diverses opérations statistiques, telles que le calcul de moyennes et la réalisation d'interpolations entre formes en trouvant un chemin entre leurs points correspondants dans l'espace.

Par construction, l'espace des formes difféomorphique est similaire à une variété de Banach. Une variété de Banach est une généralisation d'un espace de Banach. Il s'agit d'un espace de dimension infinie avec un atlas topologique. Ainsi, c'est un espace topologique dans lequel chaque voisinage d'un point est homéomorphe à un ensemble ouvert. Plus formellement, cela signifie qu'il est localement "lisse et se comporte bien".

On peut construire un *template* à partir d'un ensemble de données donné en trouvant

la forme moyenne. Elle est obtenue par le biais de la moyenne de Karcher, ou du centre de masse riemannien de l'espace, qui est analogue au centre de masse en dimension finie. En particulier, l'atlas MNI est obtenu par des processus similaires et peut être appelé un template.

Il est aussi théoriquement possible de classifier des points appartenant à un tel espace. Il va sans dire que réaliser une classification dans un espace de formes difféomorphiques n'est pas une tâche évidente. Malgré cela, des recherches ont déjà été menées. Par exemple, Srivastava et al. [2005] ont présenté une telle approche dans un espace de formes pour les courbes. Ils ont opté pour une approche stochastique pour regrouper une collection de formes en k clusters en minimisant la variance totale "à l'intérieur du cluster". De manière intéressante, ils organisent les formes dans une structure d'arbre de telle sorte que les formes affichent une résolution croissante à mesure que l'on descend dans l'arbre, avec une stratégie ascendante.

Dans la suite, lorsque nous nous référerons à l'espace de formes, il sera toujours sous-entendu que nous faisons référence à l'espace de formes difféomorphique.

1.2.3 Motivation originale : Un nouveau genre d'atlas

La motivation initiale de mon projet de thèse était de construire un nouveau genre d'atlas pour représenter le glioblastome. Cependant, son principe doit être précisément défini et de nombreuses questions doivent être résolues. Cette section présente quelques idées spéculatives, et bien que cette thèse ne couvre qu'une partie de celles-ci, elles peuvent esquisser un objectif de recherche à long terme.

Ce nouvel atlas combinerait les avantages des atlas statistiques et des atlas difféomorphiques, satisfaisant les objectifs du premier tout en incorporant des informations du second. On pourrait construire un espace de formes difféomorphique pour étudier les formes du glioblastome en elles-mêmes. Cependant, cela nécessiterait d'étudier les formes déconnectées du cerveau, sans tenir compte de la localisation des tumeurs dans le cerveau, qui est une information cruciale comme discuté précédemment dans la Section 1.2.1.

Déterminer l'emplacement exact de l'origine de la tumeur peut être difficile en raison du déplacement causé par l'effet de masse de la tumeur et de la variabilité inhérente de l'anatomie cérébrale. Avec une méthode disponible qui fournit de telles informations, il serait possible de regrouper les formes de gliomes en fonction de leur origine (*e.g. :* frontal, temporal, etc.) ou de la structure anatomique avec laquelle elles sont en contact (*e.g. :* ventricules, cervelet, etc.). De plus, nous pourrions calculer des atlas difféomorphiques en parcelles définis, par exemple, les lobes. L'analyse serait alors effectuée indépendamment dans chaque parcelle où l'on pourrait étudier des cartes de fréquence sur le point d'origine des tumeurs et en même temps utiliser les diffeomorphismes pour étudier l'évolution et la forme finale des tumeurs qui naissent dans cette parcelle. Cela améliorerait considérablement les limites des atlas statistiques, qui agrègent simplement des voxels contenant des tumeurs et ne peuvent pas différencier la source du cancer.

Pour obtenir des informations de localisation, une technique d'enregistrement qui correspond à un cerveau sain avec un cerveau cancéreux, en tenant compte des différences topologiques, peut être utilisée. Cependant, à notre connaissance, il n'existe actuellement aucune technique de recalage capable de réaliser une mise en correspondance efficace et robuste à de telles variations topologiques et de construire un espace de formes difféomorphique. C'est le sujet principal de cette thèse.

1.3 Une brève revue des techniques de recalage

Comme nous l'avons vu dans les parties précédentes, les médecins ont besoin de techniques de recalage anatomique précises à des fins de diagnostic ou de recherche. En pratique, cela est réalisé en mettant en correspondance les images voxel par voxel avec une transformation géométrique anatomiquement plausible.

Ces mises en correspondance sont généralement modélisées comme des difféomorphismes, car elles permettent de créer une déformation réaliste voxel-à-voxel sans modifier la topologie de l'image source. Il existe une vaste littérature traitant de ce sujet. Certains auteurs ont proposé d'utiliser des champs de vecteurs stationnaires, en utilisant l'exponentielle de champ de vecteurs de l'algèbre de Lie [Arsigny et al., 2006; Ashburner, 2007; Lorenzi et al., 2013] ou plus récemment, des méthodes basées sur l'apprentissage profond [Rohé et al., 2017; Yang et al., 2017; Balakrishnan et al., 2019; Mansilla et al., 2020; Niethammer et al., 2019; Mok and Chung, 2020]. Nous n'utiliserons pas ces méthodes, cependant, une brève présentation de l'exponentielle d'algèbre de Lie peut être trouvée dans la partie 2.1.2. Par exemple, Roux and et al. [2019] recale des images en utilisant DARTEL, qui est basé sur les exponentielles de Lie.

D'autres auteurs ont utilisé les *Large Diffeomorphic Deformation Metric Mapping*(LDDMM) qui utilisent des champs de vecteurs variant dans le temps pour définir une métrique riemannienne invariante à droite sur le groupe des difféomorphismes. En particulier, cette métrique peut être utilisée pour construire un espace de formes difféomorphiques, fournissant des notions utiles de géodésiques, de plus courts chemins et de distance entre les images [Avants et al., 2008; Beg et al., 2005; Younes, 2019; Zhang and Fletcher, 2018]. Dans cette méthode, le champ de déformation est représenté comme un flux de champs de vecteurs, et le problème de trouver le champ de déformation est réduit à trouver une courbe ou un chemin dans un certain espace qui représente le flux de champs de vecteurs. Il faut intégrer le long de la courbe en utilisant des schémas empruntés à la dynamique des fluides. Une partie significative des performances de LDDMM vient du choix du schéma. Bien que ce ne soit pas un nouvel algorithme, des travaux récents continuent d'étudier son implémentation, visant à améliorer à la fois la qualité de l'appariement et la vitesse/complexité [Hernandez, 2018; Mang et al., 2019; Brunn et al., 2021b; Hernandez, 2021].

Pour trouver la géodésique, l'algorithme LDDMM utilise une approche de minimisation d'énergie. La fonctionnelle d'énergie se compose de deux parties : un terme de régularisation, qui encourage le champ de déformation à être régulier en conservant le moment, et un terme de similarité, qui mesure la similarité entre l'image transformée et l'image cible. Le terme de régularisation est généralement mis en œuvre en utilisant un terme de régularisation, tel qu'un flou gaussien, et le terme de similarité est généralement mis en œuvre en utilisant une mesure telle que la différence quadratique moyenne entre l'image transformée et l'image cible.

Le recalage optimal peut être vu comme un chemin optimal qui est une géodésique. Une ressemblance théorique avec la dynamique des fluides confère à LDDMM plusieurs propriétés physiques avantageuses. En particulier, des théorèmes de conservation assurent que certaines quantités, comme le moment, sont conservées le long des chemins géodésiques. La définition du moment dans le cadre de LDDMM est difficile à introduire sans équations. Nous explorerons ce concept dans les sections 2.2.3.c et 2.3.3.

Une approche pour implémenter LDDMM consiste à utiliser le tir géodesique (*geodesic shooting*), où un champ de vitesse est propagé le long de la courbe et le champ de vecteur

initial est mis à jour de manière itérative. Nous introduisons scrupuleusement tous les détails dans la section 2.2.

1.4 Contribution de ma thèse

1.4.1 Schémas semi-Lagrangiens pour l'implémentation de LDDMM et des Métamorphoses

Même si de nombreuses publications ont présenté des résultats avec l'algorithme des Métamorphoses, contrairement à LDDMM, aucune implémentation publique n'était disponible. La première contribution de cette thèse est l'implémentation des Métamorphoses et de LDDMM pour des images en 2D et 3D.

Pour calculer les Métamorphoses, il faut intégrer un système PDE géodésique en utilisant un schéma numérique à partir d'une condition initiale. L'enregistrement est réalisé via le "shooting" en mettant à jour cette condition initiale ou résiduelle. C'est un défi car les schémas classiques nécessitent beaucoup d'itérations pour respecter les conditions CFL (CourantFriedrichsLewy) et donc converger, ce qui rend l'optimisation subséquente longue. J'ai choisi d'utiliser un schéma semi-lagrangien sur les images et les résidus pour pouvoir garder le nombre d'itérations raisonnable et donc réduire le temps d'optimisation.

L'implémentation est orientée objet, ce qui permet une grande polyvalence et des modifications faciles. Nous l'avons développée en utilisant `PyTorch` et tirons parti de l'accélération GPU pour accélérer les calculs.

Le code est entièrement disponible sur GitHub pour que tout le monde puisse l'utiliser et contribuer à l'adresse suivante : `https://github.com/antonfrancois/Demeter_metamorphosis`.

POUR EN SAVOIR PLUS RENDEZ-VOUS DANS LE CHAPITRE 2, PARTIE 2.5

1.4.2 Métamorphose contrainte : un modèle étendu

Si la Métamorphose permet des recalages de très bonne qualité, la dissociation entre les changements géométriques et d'intensité n'est pas unique et dépend fortement des hyperparamètres définis par l'utilisateur. Cela rend l'interprétation des résultats difficile, limitant son utilisation clinique. Par exemple, pour aligner un modèle sain sur une image avec une tumeur, on pourrait s'attendre à ce que la méthode ajoute des intensités uniquement pour créer de nouvelles structures (*e.g.,* des tumeurs) ou pour compenser les changements d'intensité dus à la pathologie (*e.g.,* l'œdème). Toutes les autres structures devraient être correctement alignées uniquement par les déformations. Cependant, en fonction des hyperparamètres, l'algorithme peut décider de tenir compte des différences morphologiques (*e.g.,* l'effet de masse des tumeurs) en modifiant l'apparence plutôt qu'en appliquant des déformations. Cette limitation vient principalement du fait que les changements d'intensité peuvent théoriquement être appliqués dans tout le domaine de l'image. Cependant, dans de nombreuses applications cliniques, on dispose généralement de connaissances préalables sur la position des variations topologiques entre une image saine et une image pathologique (*e.g.,* la position de la tumeur et de l'œdème).

Dans cette partie de la thèse, nous abordons les difficultés associées à l'algorithme des Métamorphoses en proposant un cadre pour incorporer des connaissances préalables dans

le modèle. Notre approche consiste à modifier le problème d'enregistrement de manière à préserver la capacité de déduire le système d'EDP qui sous-tend l'algorithme des Métamorphoses. La théorie de la Métamorphose est assez délicate à gérer, mais dans le cas des images, le principe est de définir un problème de recalage, potentiellement inspiré d'un problème de contrôle optimal [Miller et al., 2015], et de déduire un système d'EDP à partir de celui-ci. Nous proposons un ajustement de la Métamorphose qui permet d'ajouter des contraintes sur le problème d'enregistrement en concordance avec certains *priors* préalables. Nous appelons ce cadre Métamorphose contrainte. Nous présentons deux types spécifiques de *priors* préalables qui peuvent être incorporées dans le modèle : (1) un masque de croissance généré à partir d'une segmentation donnée, et (2) un champ qui guide la déformation dans une direction souhaitée.

Notre implémentation de l'algorithme des Métamorphoses contraintes repose sur notre implémentation précédente des Métamorphoses et est conçue pour être suffisamment flexible pour permettre à tout utilisateur d'incorporer ses propres contraintes supplémentaires. Nous démontrons l'efficacité de notre approche à travers des expériences sur des glioblastomes, en utilisant des ensembles de données BraTS, en comparant avec des méthodes de pointe.

POUR EN SAVOIR PLUS RENDEZ-VOUS DANS LE CHAPITRE 3,
CADRE MATHÉMATIQUE : SECTION 3.2.3
VALIDATION SUR UN EXEMPLE JOUET : SECTION 3.2
VALIDATION SUR DES DONNÉES RÉELLES : SECTION 3.3

1.4.3 Segmentation utilisant l'Analyse de Données Topologiques (TDA)

En collaboration avec Raphaël Tinarrage, nous avons développé un outil de segmentation de tumeurs utilisant l'Analyse de Données Topologiques (TDA). Nous cherchons à détecter des composants caractéristiques dans les modalités FLAIR et T1ce en prenant des hypothèses sur leurs topologies. En utilisant la TDA, nous sommes capables de sélectionner et de trouver des composants connectés et des cycles qui composent un gliome sans l'utilisation d'outils statistiques. De plus, en automatisant le processus de segmentation, nous acquérons une compréhension plus approfondie des caractéristiques des glioblastomes en termes de leur emplacement, leur forme et leur intensité dans les images IRM. En effet, en ayant une méthode de segmentation robuste, nous posons les bases pour construire un atlas difféomorphique des glioblastomes. Le code est entièrement disponible sur GitHub pour que chacun puisse l'utiliser et y contribuer : `https://github.com/antonfrancois/gliomaSegmentation_TDA`.

POUR EN SAVOIR PLUS RENDEZ-VOUS DANS LE CHAPITRE 4

1.4.4 Outil de visualisation

Les algorithmes LDDMM et des Métamorphoses produisent une séquence d'images ou une image temporelle, qui peut être interprétée comme une représentation visuelle du déplacement de l'image source vers l'image cible. Cependant, en raison des défis liés à la manipulation d'images 3D et du manque de plugins Python facilement disponibles pour visualiser à la fois des images temporelles et des déformations simultanément lors du débogage, nous avons développé un outil qui résout ces problèmes.

Notre outil, implémenté à l'aide de la bibliothèque `vedo` [Musy et al., 2023], permet d'afficher:

- une seule image temporelle 3D
- la comparaison de deux images temporelles en les affichant côte à côte et en les superposant d'un simple clic.
- le résultat d'une optimisation LDDMM, Métamorphose ou Métamorphose Contrainte (CM), comprenant des informations utiles telles que les paramètres utilisés ou le nombre de pixels où la déformation a déchiré l'arrière-plan (si la déformation n'est pas difféomorphique)
- En plus du recalage, on peut également visualiser la déformation sous la forme d'une collection de vecteurs 3D.

On peut voir une démonstration dans cette vidéo: `https://www.youtube.com/watch?v=-qzQ1B6DdSg` et retrouver le code sur le GitHub antonfrancois/Demeter_metamorphosis.

1.5 Organisation des chapitres

CHAPITRE 2: couvre toutes les connaissances de base sur LDDMM et les Métamorphoses, en commençant par un petit détour donnant des éléments d'algèbre de Lie qui sont utiles pour comprendre les concepts clés de la géométrie différentielle. Nous en profitons pour expliquer brièvement l'exponentielle de Lie comme outil pour le recalage. Ensuite, nous donnons une démonstration scrupuleuse de LDDMM, introduisant lentement les nombreuses propriétés. Nous continuons en expliquant en détail les propriétés de conservation de la quantité de mouvement et comment cela nous permet d'utiliser le tir géodésique. Après avoir digéré toute cette théorie, nous sommes prêts à comprendre le cadre Métamorphique et ses dérivations géodésiques avec facilité. Nous terminons le chapitre en expliquant notre mise en œuvre des deux méthodes susmentionnées en utilisant des schémas semi-lagrangiens (voir Contribution 1.4.1) et en donnant quelques premiers résultats.

CHAPITRE 3: introduit le cadre des Métamorphoses contraintes (CM) (voir Contribution 1.4.2). Après la formulation mathématique, nous présentons l'exemple jouet (voir Contribution 1.4.4) que nous avons utilisé pour le débogage et pour une première validation. Ensuite, nous présentons les performances des CM sur des séquences d'IRM sur deux applications différentes. Tout d'abord, dans la continuité de la motivation originale de cette thèse, nous le présentons comme un outil pour l'enregistrement d'un modèle de cerveau sain sur un cerveau avec un glioblastome. Ensuite, nous avons profité du défi BraTSReg 2022 [Baheti and et al., 2021] pour tester CM sur l'alignement de chirurgies de suivi.

CHAPITRE 4: présente une méthode de segmentation de glioblastomes en utilisant TDA. Nous commençons par faire une brève revue des méthodes de segmentation de glioblastomes en utilisant l'apprentissage automatique. Ensuite, nous couvrons quelques articles utilisant TDA pour l'analyse d'IRM. Ensuite, dans la sous-section 4.2, nous faisons une brève introduction aux différents concepts de TDA utiles pour l'analyse d'IRM avec TDA. Nous sommes maintenant prêts à présenter notre algorithme de segmentation dans la sous-section

4.3. Enfin, nous présentons nos résultats et discutons des améliorations que nous voulons implémenter, dans la sous-section 4.4.

Conclusion Générale & Perspectives : nous élaborons davantage sur le nouveau type d'atlas présenté dans la sous-section 1.2.3. Nous expliquons comment le travail que nous avons présenté dans cette thèse peut servir à développer un tel atlas.

1.6 Articles publiés et Préprints

1. A. François, P. Gori, and J. Glaunès. Metamorphic image registration using a semi-lagrangian scheme. In *SEE GSI*, 2021

2. A. François, M. Maillard, C. Oppenheim, J. Pallud, I. Bloch, P. Gori, and J. Glaunès. Weighted metamorphosis for registration of images with different topologies. In *WBIR*, pages 8–17, 2022

3. M. Maillard, A. François, J. Glaunès, I. Bloch, and P. Gori. A deep residual learning implementation of metamorphosis. In *IEEE ISBI*, 2022

1

Introduction - en

This Dissertation lies at the intersection of several disciplines. We apply our work to medical research on glioblastomas (brain tumours). From radiological data, we aim to infer an average shape based on its position in the brain. To do this, we use applied mathematics, specifically the theory of shape spaces. This field is itself at the intersection of several areas of mathematics and physics. It is based on differential geometry but has strong links with fluid dynamics while relying heavily on optimisation techniques. In this introduction, we take the time to detail the relevance of the application and introduce the concepts of shape spaces and registration.

1.1 Medical motivation & data overlook

1.1.1 About Glioblastomas

GENERALITIES ABOUT TUMOURS. A tumour is a mass of abnormal cells that grows in the body. These abnormal cells grow and divide in an uncontrolled way, forming a mass that can invade surrounding tissues and organs. Tumours can be *benign*, which means they are not cancerous and do not spread to other parts of the body. If a tumour is not benign, we say it is *malignant*. The former type of tumours can impede body functions by redirecting nutrients to their only benefit, potentially depriving other essential tissues. Cancers (with an emphasis on the plural) can take many forms and are extremely heterogeneous in shape, composition and symptoms. Indeed their origin is often due to an unfortunate mutation that occurred during cell division, altering their normal function.

Despite the variability, one can distinguish several parts of a tumour, including the main *tumour mass*, which is composed of the abnormal cells, and the surrounding tissue, which may be normal or may contain a mix of normal and abnormal cells. The tumour may also have a layer of cells called the tumour capsule, which surrounds the main tumour mass and separates it from healthy tissues. The tumour may also have blood and lymph vessels, which can help it grow and spread to other parts of the body. Inside the tumour mass, *necrosis* can happen. Necrosis is a type of cell death that occurs when cells are damaged or injured beyond repair. Unlike other forms of cell death, such as apoptosis, necrosis is a chaotic and uncontrolled process that can be harmful to the body. Necrosis can occur in any type of tissue but is most commonly associated with tissue damage caused by injury or infection. When necrosis occurs, the affected cells swell and rupture causing *Oedemas* by releasing their contents into the surrounding tissues. This can cause inflammation and damage to nearby cells. In severe cases, necrosis can lead to tissue or organ death.

At the cancerous tumour borders, cells are more active and may spread into surrounding tissues in a diffusive way. In this case, it is called *Infiltrative cancer*. Healthy and cancerous cells being mixed together at the frontier, make this type of cancer difficult to treat surgically. When these tumours reach blood vessels or the lymphatic system, they can also spread to other parts of the body, which can make them more aggressive and harder to control. Treatment for infiltrative cancer often involves a combination of surgery, radiation therapy, and chemotherapy to kill the cancer cells and prevent them from spreading.

MORE ON GLIOMAS There are only two types of cells in the central nervous system (CNS): *neurons* are channels for cognition and *Glial* cells are the other CNS cells that fulfil all neuron needs. The former has a high variety of forms and functions. For example, Astrocytes provide nutrients, and regulate blood flow, while protecting neurons from antibodies. Brain tumours are most of the time *Glioma*, which is a generic term describing primary brain tumours. Neuronal tumours are very rare and mostly benign. Glioma roots in glial cells and their specific names come from their cell's name (*e.g.,* astrocytoma) [Hanif et al., 2017]. Glioblastoma (GB) multiform is the most malignant and frequently occurring type of primary astrocytoma. It accounts for more than 60% of all brain tumours in adults [Rock et al., 2012]. Their causes are not well understood, with exposure to high-dose ionizing radiation as the only confirmed risk factor [Hanif et al., 2017]. The often recommended treatment is precise surgery which can help improve the patient's quality of life by reducing the tumour burden, controlling seizures or even reversing neurological deficits.

Tumour location is a key parameter in the care of patients with Glioblastoma because it correlates with demographic characteristics, symptoms, surgical management, delivery of subsequent oncologic treatments, and, ultimately affects the patient's prognosis. [Roux and et al., 2019]. Previous pathogenesis research has shown that the most frequent location for GB is the cerebral hemispheres. 95% of GB arise in the supratentorial region (upper part), while only a few in the cerebellum, brainstem and spinal cord [Nakada et al., 2011]. At a macroscopic scale, GBs are quite heterogeneous in form and irregularly shaped but usually arise in white matter. It has been shown that depending on the lobe where the tumour arises, symptoms vary. For example, patients with a GB located in the temporal lobe area often show hearing and visual problems, while those who have one in the frontal lobe might demonstrate personality change [Hanif et al., 2017]. Also, the distribution of oedema/necrosis leads to different secondary effects in the patient: a gradual increase in tumour size and increased oedema surrounding the tumour leads to a shift in intracranial contents, resulting in headaches. A tumour that resides in sites like the eloquent cortex brain stem or basal ganglia cannot go through surgery and these patients usually have worse prognoses [Mrugala, 2013].

The World Health Organisation (WHO) classification of glioblastomas is a system used to classify and describe different types of glioblastoma [Villa et al., 2018; Berger et al., 2022]. This classification is based on the microscopic appearance of the cancer cells, as well as on the genetic and molecular characteristics of the tumour. It is an important tool for clinicians and researchers, as it helps to characterise the specific features of a particular tumour and to guide treatment decisions.

The WHO grade of glioblastoma is typically determined by analysing a tissue sample taken during a biopsy or surgery. However, magnetic resonance imaging (MRI) can also be used to identify certain characteristics of glioblastoma that may be indicative of its WHO grade. High-grade glioblastomas tend to be larger, have an irregular, diffuse shape, show strong enhancement on contrast-enhanced MRI scans, be associated with areas of necrosis, and infiltrate surrounding brain tissue. It is important to note that MRI alone is not sufficient for accurately grading a glioblastoma. A tissue sample must be analysed by a pathologist to definitively determine the WHO grade of the tumour.

The first diagnosis is always based on tumour visualisation and is made using radiological imaging techniques, the most common being MRI.

1.1.2 Magnetic Resonance Imaging for brain visualisation

A magnetic resonance imaging (MRI) machine uses strong magnetic fields and radio waves to create detailed images of the inside of the body. The patient lies on a table that is moved into the MRI machine, which contains a large, powerful magnet. The magnet aligns the hydrogen atomic particles in the body, and radio waves are used to knock these particles out of alignment. As the particles return to their original alignment, they emit a radio-frequency signal that is detected by the MRI machine and used to create an image of the inside of the body. The images produced by an MRI machine can be used to diagnose a wide range of medical conditions.

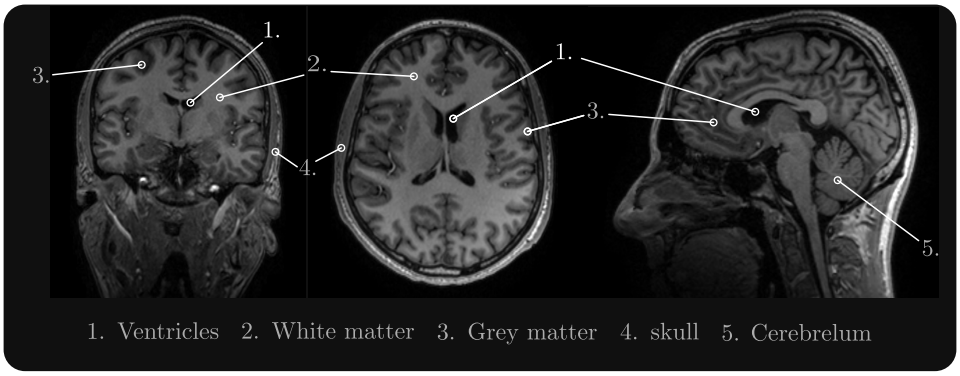

Figure 1.1: A short glossary of brain organs visible in MRIs, (T1)

Figure 1.1 illustrates a T1 MRI of a healthy brain, including the skull. The skull is usually removed during preprocessing, however, leftovers may remain. Brain tissues can be split into four regions:

1. **Ventricles** are chambers in the brain, filled with cerebrospinal fluid (CSF). There are four ventricles in the brain, two in the cerebrum (the left and right lateral ventricles) and two in the brainstem (the third and fourth ventricles). In this thesis, we will mostly refer to the left and right ones. The primary function of the ventricles is to produce and circulate cerebrospinal fluid, a clear, watery fluid that surrounds and cushions the brain and spinal cord. It provides several important functions, including protecting the brain and spinal cord from mechanical injury, maintaining the correct balance of nutrients and waste products in the CNS as well as regulating the pressure inside the skull.

2. The **White Matter** corresponds to the large lighter areas in Figure 1.1. It is composed of long axons, surrounded by a fatty substance made of glial cells called myelin. White matter is so named because of the white appearance of the myelin sheath, which contrasts with the grey appearance of the nerve cell bodies that make up Gray matter. Note that these tones are not conserved by MRI (see next paragraph).

3. **Gray Matter** is mostly composed of neurons somas and synapses and is primarily responsible for processing and interpreting information.

4. **Cerebellum** is a smaller region located at the back of the head under the cerebrum. It is made up of two hemispheres and is connected to the brainstem and the spinal

cord. The outer portion contains neurons, and the inner area communicates with the cerebral cortex. It is responsible for coordinating movement, balance, and posture. It receives input from sensory systems in the body, including the eyes, ears, and muscles, It uses this information to make adjustments to movement in order to maintain balance and accuracy. New studies are exploring the cerebellums roles in thought, emotions and social behaviour, as well as its possible involvement in addiction, autism and schizophrenia.

Multiple sequences are usually needed to adequately evaluate a tissue, and the combination of sequences is referred to as an MRI protocol. In particular, for Glioma diagnosis, most suspected patients undergo at least four MRIs modalities, that help distinguish different organs [Villanueva-Meyer et al., 2017; Ginsberg et al., 1998]:

- **T1** or T1-weighted make fats appear bright and water dark. In particular, the CSF in the ventricles appears darker. In Figure 1.1, one can see a first T1 MRI scan.
- **T1ce** or T1-weighted contrast-enhanced is an MRI, where the patient got injected with intravenous agents, that increase tissue contrast by accentuating areas where contrast agents have leaked out of the bloodbrain barrier into the interstitial tissues. Within diffuse gliomas, contrast enhancement is positively correlated with tumour grade, although a few high-grade gliomas and some lower-grade gliomas may show no or minimal enhancement.
- On **T2** or T2-weighted images, each tissue has an inherent T2 value, but external factors (such as magnetic field inhomogeneity) can decrease the T2 relaxation time, thus diminishing the value. Fluids are bright and fats are intermediate.
- **T2-FLAIR** or T2-fluid-attenuated inversion recovery is a T2 scan with a process that removes the signal from the cerebrospinal fluid in the resulting images. Brain tissues on FLAIR images appear similar to T2 weighted images with grey matter brighter than white matter. Thus, unlike T2, CSF and ventricles appear dark. Glioblastoma induces Cerebral cortical T2 hyper-intensities. FLAIR has a high intensity in peritumoral oedema (vasogenic and infiltrative), non-enhancing tumour, white matter injury and gliosis. Vasogenic oedema represents a reactive increase in extracellular water due to leakage of plasma fluid. In many gliomas, the T2/FLAIR hyper-intense signal abnormality may be indistinguishable from the primary mass lesion.

We have seen that each imaging modality is useful in visualising specific brain structures. However, it is important to note that the quality and appearance of MRI images may vary greatly due to a number of factors, including the specific MRI machine being used and the acquisition conditions. Additionally, the specific pathology being observed can also impact the appearance of the images. As a result, it is common for MRI images of the same modality obtained from different patients to exhibit significant differences, see Figure 1.2. In the last column, one can see an example of glioblastoma segmentation. A segmentation in three classes is commonly used for glioma: the necrotic tumour core (TC), the GD-enhancing tumour (ET) and the peritumoral edematous/invaded tissue (WT) [Baid and et al., 2021]. Note that WT stands for 'Whole tumour' as for the BraTS2021 challenge, in this dissertation we use it to refer to 'Tumoural tissues that are not the tumoural core'. This figure displays manual segmentation done by experts.

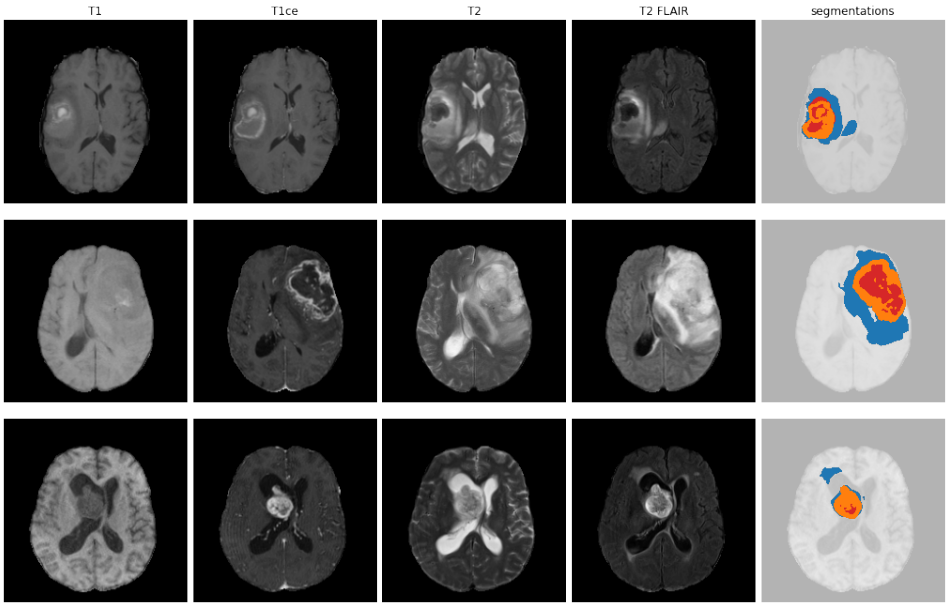

Figure 1.2: GB segmentation examples: Each row contains MR Images slices in four modalities of the same patient at the same time. One can see both the tendencies between different image modalities and the inter-subject colour variation within the same modalities. Segmentation legends: Red - Necrotic Core (TC); Orange - Enhancing Tumour (ET); Blue - Oedemas (WT).

1.2 Atlases to map the brain

With MRI, we have a tool for visualising individual brains, however, to compare images taken by radiologists to a reference, a representative brain is needed. In general *Atlases* are a tool for object shape studies. A few definitions of this term co-exist depending on their respective fields. It happens that in this thesis we are at the intersection of at least three of them. We give here short definitions and we will detail two of them in the following subsections :

- **Shape-space Atlases** is one of the main notions for this thesis and will be explained with more details in Section 1.2.2. A shape space is a mathematical representation of shapes, where each point in the space represents a different shape. In a shape space, the distance between two points reflects how similar the shapes are. By studying shapes in a shape space, researchers can analyse and compare shapes, and understand how they change and evolve over time. The theory of shape spaces is an important tool in fields such as computer vision, medical imaging, and biomechanics, among others. In this dissertation, with the exception of Section 1.2.2, Shape-space refers to the *diffeomorphic shape-space*. It is a specific type of shape space that is used to represent shapes that are continuously deformable into one another, preserving the topology of the shapes. This means that the shapes are transformed in a way such that their underlying structure is preserved and the transformation invertible. In the diffeomorphic shape space, shapes are represented as points, and diffeomorphic transformations correspond to paths between points. Once this space is constructed, one can estimate an average shape called *template*.

- In biology and medicine, atlases often refer to **Statistical Atlases.** A statistical atlas can be seen as a frequency map that shows the distribution of a particular variable or feature within a defined population. These maps can be used to visualise and analyse data in a spatial context, allowing researchers to see patterns and trends that may not be apparent from tabular data alone. Statistical atlases are often used in the medical field to visualise and compare data related to health and disease. For example, a statistical atlas of the brain might show the distribution of brain structures or brain function across a population, or the prevalence of a particular brain disorder in different regions of the brain. We will detail this concept further in Subsection 1.2.1.
- In mathematics, **the topological Atlas** is used for the description of abstract objects like manifolds. It is composed of individuals *charts* describing locally the manifold. Juxtaposed charts must 'describe' their overlapping regions in 'equivalent' ways. While this is an important notion in differential geometry, we will not need to work directly with this notion.

If sick patients were always to be treated individually, one would need a representation of a healthy subject. Atlases can serve the purpose of comparing healthy and ill subjects. The Montreal Neurological Institute (MNI) brain template is a standard reference system used in neuroimaging research to align and compare brain images from different individuals Fonov et al. [2009, 2011]. It is based on a set of brain images that have been averaged together to create a "mean" brain, which represents the average anatomy of a large group of people.

The MNI brain template is created by aligning a large number of brain images using specialised software, and then averaging the images together to create a composite brain that represents the average anatomy of the group. This template is then used as a reference system to align other brain images, allowing researchers to compare brain anatomy and function across different individuals. The MNI brain template is widely used in the field of neuroimaging and has become a standard reference system for many researchers. It is particularly useful for aligning brain images from different imaging modalities, such as MRI and PET, and for comparing brain images from different populations or groups. In the literature, it is often called the MNI Atlas in a misleading way. Indeed, it can be obtained from shape-space Atlases (but are only a part of it, see section 1.2.2) and are often used as background in Statistical Atlases.

1.2.1 Statistical Atlases for studying Glioblastoma localisation

Glioblastomas can occur in any part of the brain, but they are most commonly found in the cerebral hemispheres, which are the left and right halves of the brain. As we have mentioned earlier, the location of glioblastoma can influence the symptoms a patient experiences, as well as the surgical and treatment options that are available. For example, a glioblastoma located in the cerebral hemispheres may cause symptoms such as headache, nausea, vomiting, and seizures, while a glioblastoma in the brain stem may cause symptoms such as difficulty with balance, weakness, or changes in speech or vision [Bilello et al., 2016; Parisot et al., 2016]. Simpson et al. [1993] saw no difference in survival for the different tumour sizes and asses that patients with frontal lobe tumours survived longer than those with temporal or parietal lobe lesions, concluding that localisation is a crucial prognosis indicator.

To give such results, Neurosurgeons need to study average Glioblastomas location and therefore build a statistical atlas: A map where the background is an averaged brain template

and on which each voxel is associated with its glioma apparition frequency [De Witt Hamer et al., 2013; Bilello et al., 2016; Parisot et al., 2016; Roux and et al., 2019; Sagberg et al., 2019]. In other words, they carry for each region the tumour apparition estimated probability.

We will use the recent paper of Roux and et al. [2019] as an example of the method to build Statistical Atlases. The authors have built such an atlas to help the patient prognosis in function of age, sex and other clinical data. The construction method can be split into three steps:

1. An expert first manually segmented the enhancing and necrotic components of glioblastomas with the help of the Multi-image Analysis GUI Mango, which allows making simple 3D image visualisation and transformation.
2. Then they use the MNI template as a healthy reference to register each ill image on it. To do so they used Diffeomorphic Anatomic Registration methods Ripollés et al. [2012] using Lie Algebra Exponential (see Section 2.1.2) and cost Masking.
3. Before constructing the atlas, segmentation and registrations were checked by another trained expert, and the authors claim that no correction needed to be made. Finally, all segmentation where superimposed in all regions of interest, obtaining a 3D frequency map based on different parameters.

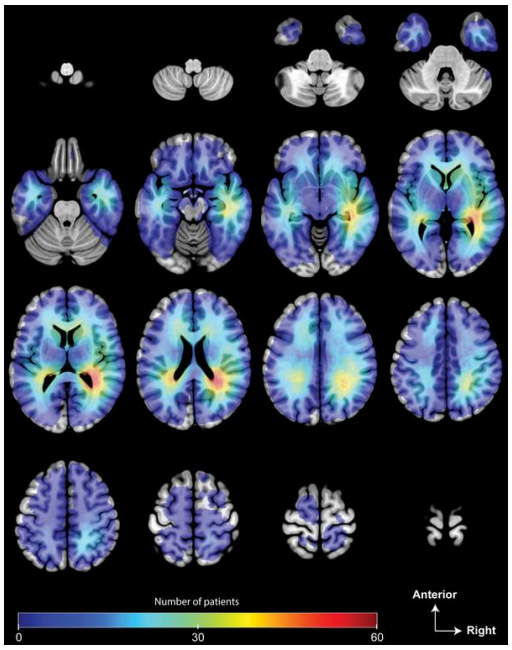

Figure 1.3: Location and frequency of isocitrate dehydrogenase (IDH) wild-type glioblastoma (n = 392). The colour frequency map illustrates the location of and the number of patients with IDH wild-type glioblastoma. Images are displayed in the neurologic display convention. Figure from Roux and et al. [2019]

The researchers found that IDH wild-type glioblastomas were most commonly located in the subcortical white matter of the subventricular zones (*i.e:* Thin layer of cells located near the lateral ventricles) in both hemispheres of the brain. There were differences in glioblastoma location depending on whether the tumour involved the subventricular zone. In patients with glioblastomas near ventricles, tumours were more likely to be located in the anterior horn, right and left atrium, and right and left temporal horn of the lateral ventricles. They also made links between the observed symptoms and the location. For example, they assert that, among the 81 patients in whom the glioblastoma was detected due to symptoms of elevated intracranial pressure, there was a greater likelihood of the glioblastomas being situated in the right frontal lobes.

So location is indeed a very important factor. However, previous analysis has shown that the relationship between tumour location and preoperative seizure incidence depends on the glioma grade of malignancy [Pallud et al., 2016]. In addition, physicians are also interested to seek the links with the tumour shape but lack an easy-to-use test for their hypothesis. For example, infiltrative tumour cells tend to follow blood vessels and are more susceptible to spreading to other brain regions. To do so, one could build a Shape-space Atlas.

1.2.2 Shape-spaces comprehensive introduction

The notion of shape-space is not a very intuitive one for the unfamiliar reader. In this section, I start by explaining two classical an elementary shape spaces to introduce the diffeomorphic space-space, which is the focus of our interest. If you are interested more generally in shape spaces, I advise you to read the review made by Bauer et al. [2014].

1.2.2.a Kendall's space

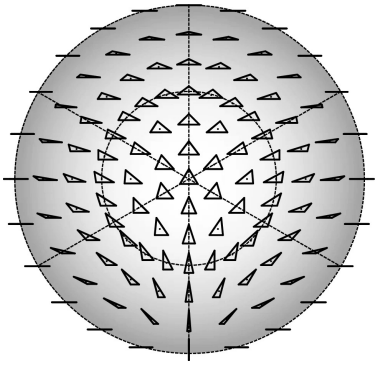

Figure 1.4: Kendalls shape space for triangles in two dimensions. The diagram shows the view from the north pole corresponding to an equilateral triangle (centre) onto one hemisphere. The outer dashed circle corresponds to the equator and contains the collinear triangles, where all three vertices lie on a straight line. The six meridians shown as straight dashed lines contain isosceles triangles. Figure from Klingenberg [2020]

Kendall's space is a mathematical space that is used to represent a shape as a collection of points and all its possible natural variations. It is named after the mathematician David G. Kendall, who introduced the concept in his work on shape analysis [Kendall, 1984; Kendall et al., 2009]. In this space, each element is a shape.

For example, Klingenberg [2020] made a very interesting study of the drosophila wing shape. In this article, he gives the case of the triangle shape space which can be represented on a sphere, therefore, is two-dimensional (see Figure 1.4). More precisely, it is the quotient of the space of configurations of three points in $\mathbb{R}^2$ (thus with a base dimension of 6) by the similarities of $\mathbb{R}^2$ (a group with a dimension of 4: 2 parameters of translation, one of rotation, one of homothety), which explains the dimension of 2. Two neighbouring points represent two different shapes, or triangles, that are close, in the sense that one is equal to the other by only a small deformation. For more complex point sets, the shapes are spread on a hyper-sphere, providing a handy representation of the shapes to work with. it provides a framework to perform statistical manoeuvres such as estimating average shapes or characterising variation of shapes around those averages. This can be fundamental for geometric morphometric biological applications.

However, the Kendall space is suitable to compare a set of landmarks only by Procrustes analysis and does not take the background into account. It is therefore not suitable for our imaging applications.

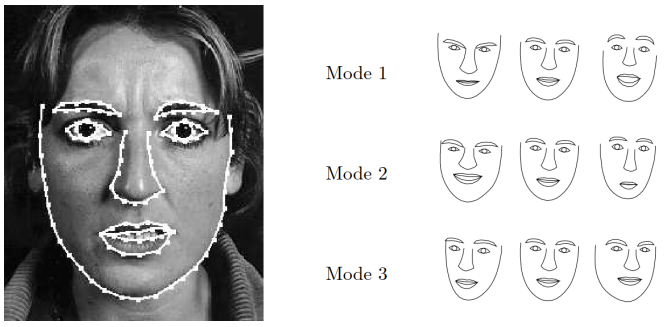

(a) Example face image annotated with landmarks

(b) Effect of varying each of first three face model shape parameters in turn between ± 3 s.d.

Figure 1.5: Active shape space of faces. Figures from Cootes et al. [2000]

1.2.2.b Active Shape Space

A more suitable shape space for images could be the active Shape Space [Cootes et al., 2000], which is based on Active contour models (sometimes called snakes) [Kass et al., 1988]. Active contour models are methods aiming at outlining object boundaries in an image by fitting a given curve to its edges. These methods do not solve the entire problem of finding contours in images, since the method requires knowledge of the desired contour shape beforehand. The Active Shape space is learned using a collection of such obtained curves by applying Principal Component Analysis (PCA). First, a set of connected landmarks defined by a user is augmented by an Active contour method and then roughly aligned in a common coordinate frame (see Figure 1.5a). Then a dimensional reduction is applied to the set of landmarks using a PCA. In the process, one retrieves the main axes and therefore a shape space representation. Once the model is learned, one can represent shape variation on its different axis or modes (see Figure 1.5b). These modes explain global variation due to 3D pose changes, which cause movement of all the landmark points relative to one another. Less significant modes cause smaller, more local changes.

However, they are derived directly from the statistics of a training set and will not always separate shape variation in an obvious manner. Even if it is considered as an 'image shape space' by the literature, it is a shape space for landmarks with a L^2 metric utilising active contour techniques to find the points. In addition, these models are not obvious to generalise to 3D as a curve would become a surface, needing the usage of triangulation. More importantly, the background is not taken into account during the process, not allowing to have a direct displacement map between points of this space.

We should note also the active appearance model [Edwards et al., 1998] which matches a statistical model of an object's shape and appearance to a new image taking the background into account. A set of images with landmarks is used during the training phase to build the model. The algorithm uses least squares techniques to quickly match two new images.

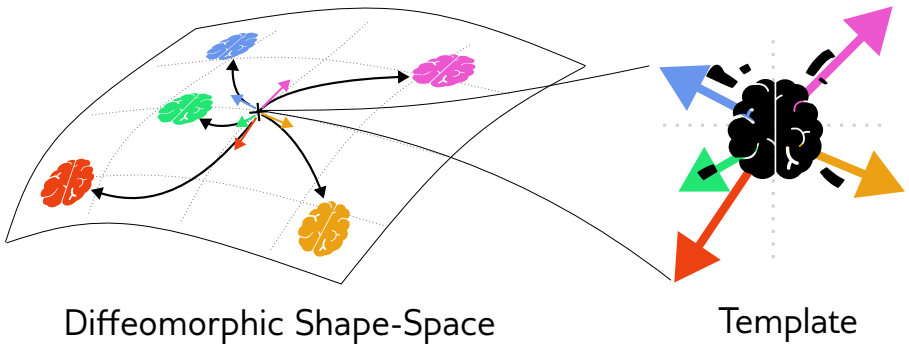

Figure 1.6: Schematic view of a Diffeomorphic-Shape space. Each brain represents a data point, the surface is the Riemannian manifold on which lies the deformations (black arrows). The deformations are fully described by their initial velocity (coloured arrows). By finding the centroid of this space, one can deduce the template (black brain).

1.2.2.c Diffeomorphic Shape Space

The diffeomorphic shape-space has roots in D'Arcy-Thompson's work [Thompson, 1942] and was further developed in Grenander's Pattern Theory [Grenander, 1993]. The specificity of this framework is that the deformation maps a reference shape, called a *template*, to the observed shape. The template can be chosen fixed or estimated, either way, it serves the purpose of the space's centroid. In addition, we take the assumption that all forms in this space can be matched by bijective transformations. Indeed, these spaces are constructed from the structure induced by groups of diffeomorphisms acting on shapes through deformation. Theoretically, these diffeomorphisms being defined on the whole domain are thus continuous functions and map not only the objects but also the background, or the ambient space.

Thus the model is extremely versatile, recent models of deformable templates focus specifically on deformations represented by diffeomorphisms acting on landmarks, curves, surfaces, images or even a mixture of these objects. More precisely, the space of diffeomorphisms induces a right-invariant Riemannian metric on the space of shapes of the ambient space, giving a sense of distance between objects. [Dryden and Mardia, 2016; Miller and Younes, 2001; Trouvé and Younes, 2003; Charon, 2013]

This shape space possesses a number of interesting properties. As a result of its nature as an infinite-dimensional Riemannian manifold, the minimal deformation connecting two points in the space can be viewed as a geodesic, and the amount of deformation serves as a Riemannian metric on the space. This allows for the application of various statistical and mathematical operations, such as computing means and performing interpolation between shapes by finding a path between their corresponding points in the space. By construction, the diffeomorphic shape-space is similar to a Banach manifold. A Banach manifold is a generalisation of a Banach space. It is an infinite-dimensional space with a topological atlas. Thus is it a topological space in which each point neighbourhood is homeomorphic to an open set. More informally speaking, it means that it is locally 'smooth and well behaved'.

One can build a *template* from a given data set by finding the average shape. It is obtained through the Karcher mean, or the Riemannian Centre of Mass of the space, which is analogous to the centre of mass in finite dimension. In particular, the MNI atlas is obtained through similar processes and can be called a template [Fonov et al., 2009, 2011].

It is also theoretically possible to cluster points within a diffeomorphic shape space. However, it goes without saying that it is not an obvious task as well. Despite this, research has been already conducted. For example, Srivastava et al. [2005] presented such an approach in a shape space for curves. They opted for a stochastic approach to cluster a collection of shapes in k clusters minimising the total "within-cluster" variance. Interestingly, they organise shapes in a tree structure such that shapes display increasing resolution as we move down the tree, with a bottom-up strategy.

In the latter when we will refer to shape-space, it would always be assumed that we refer to the shape space of diffeomorphism.

1.2.3 Original motivation: A new atlas genre

The original motivation for my PhD project was to build a new atlas genre for glioblastoma representation. However, its principle itself has to be precisely set and a lot of questions need to be answered. This section presents some speculative ideas, and if this dissertation covers only a part of those, they can sketch a long time research objective.

This new atlas would combine the advantages of both statistical and diffeomorphic atlases, satisfying the goals of the former while incorporating information from the latter. One could build a diffeomorphic shape space to study glioblastoma shapes by themselves. However, it would require studying the shapes disconnected from the brain, not accounting for the tumours' localisation within the brain, which is crucial information as previously discussed in Section 1.2.1.

Determining the exact location of the tumour origin can be challenging due to the natural growth of the tumour and the displacement of the surrounding tissue caused by the tumour mass effect and the inherent variability of brain anatomy. With an available method that provides such information, it would be possible to cluster glioma shapes based on their origin (*e.g.,* frontal, temporal, etc. lobes) or the anatomical structure they are in contact with (*e.g.,* ventricles, lobes, cerebellum, etc.). In addition, we could compute diffeomorphic atlases into parcels/clusters defined, for instance, beforehand as the lobes. The analysis would then be carried out independently in each parcel where one could investigate frequency maps about the origin point of the tumours and at the same time use diffeomorphisms to study the evolution and the final shape of the tumours that are born in that parcel. This would significantly enhance the limitations of statistical atlases, which merely aggregate voxels containing tumours and cannot differentiate the original location of cancer.

To obtain localisation information, a registration technique that matches a healthy brain with a cancerous brain, taking into account topological differences, can be used. However, to the best of our knowledge, there is currently no registration technique that can perform efficient registration robust to such topological variations and also construct a diffeomorphic shape space. This is the main topic of this dissertation.

1.3 A short review of Registration techniques

As we have seen in the previous sections, for diagnosis or research purposes, physicians need accurate anatomical registrations. In practice, this is achieved by mapping images voxel-wise with a plausible anatomical transformation.

These mappings are usually modelled as diffeomorphisms, as they allow for the creation of a realistic one-to-one deformation without modifying the topology of the source image. There exists a vast literature dealing with this subject. Some authors proposed to use stationary vectors fields, using the Lie algebra vector field exponential [Arsigny et al., 2006; Ashburner, 2007; Lorenzi et al., 2013], or, more recently, Deep-Learning based methods [Rohé et al., 2017; Yang et al., 2017; Balakrishnan et al., 2019; Mansilla et al., 2020; Niethammer et al., 2019; Mok and Chung, 2020]. We will not use these methods, however, a short presentation of the Lie-Algebra exponential can be found in Section 2.1.2. For example, Roux and et al. [2019] register images using DARTEL, which is based on Lie-Exponentials.

Other authors used the Large Diffeomorphic Deformation Metric Mapping (LDDMM) that utilises time-varying vector fields to define a right-invariant Riemannian metric on the group of diffeomorphisms. Among other advantages, this metric can be used to build a diffeomorphic shape space, providing useful notions of geodesics, shortest paths and distances between images [Avants et al., 2008; Beg et al., 2005; Younes, 2019; Zhang and Fletcher, 2018]. In this method, the deformation field is represented as a flow of vectors fields , and the problem of finding the deformation field is reduced to finding a curve or a path in a certain space that represents the flow of vector fields. One needs to integrate along the curve using schemes borrowed from fluid dynamics. A significant part of LDDMM performances comes from the scheme choice. While it is not a new algorithm, recent works continue to study its implementation, aiming to improve both the matching quality and speed/complexity [Hernandez, 2018, 2021; Mang et al., 2019; Brunn et al., 2021b].

To find the geodesic, the LDDMM algorithm uses an energy minimisation approach. The energy functional consists of two parts: a smoothness term, which encourages the deformation field to be smooth by conserving the momentum, and a similarity term, which measures the similarity between the transformed image and the target image. The smoothness term is typically implemented using a regularisation term, such as a Gaussian blurring, and the similarity term is typically implemented using a similarity measure, such as the mean squared difference between the transformed image and the target image. Optimal registration can be seen as an optimal path which is a geodesic. The connection with fluid dynamics endows LDDMM with several advantageous physical properties. In particular, conservation theorems ensure that certain quantities, like momentum, are conserved along geodesic paths. The momentum definition within the LDDMM framework is challenging to introduce without equations. We will explore this concept in Sections 2.2.3.c and 2.3.3. One approach to implement LDDMM is through the implementation of geodesic shooting, whereby a velocity field is propagated along the curve and the initial vector field is iteratively updated. We introduce scrupulously all details in Section 2.2.

LDDMM has been applied to many data and in various goals. For example, Khan et al. [2008] and later Kutten et al. [2016] used LDDMM to propagate brain segmentations along with annotations. However, clinical or morphometric studies often include an alignment step between a healthy template (or atlas) and images with lesions, alterations or pathologies, like white matter multiple sclerosis or tumour. In such applications, source and target images show a different topology, thus making the use of diffeomorphisms problematic, which are by definition one-to-one mappings. Several solutions have been proposed in order to take into account such topological variations. One of the first methods was Cost-Function Masking [Brett et al., 2001], where authors simply excluded the lesions from the image similarity cost. It is versatile and easy to implement, but it does not give good results when working with big lesions. Sdika and Pelletier [2009] proposed an inpainting method that only works on small lesions. Niethammer and et al. [2011] proposed Geometric Metamorphosis, which

combines two deformations to align pathological images which need to have the same topology. Another strategy, when working with brain images with tumours, is to use biophysical models Gooya et al. [2011a]; Scheufele and et al. [2019] to mimic the growth of a tumour into a healthy image and then perform the registration (see for instance GLISTR [Ali et al., 2012]). However, this solution is slow, computationally heavy, specific to a particular kind of tumour and needs many different imaging modalities. Other works proposed to solve this problem using Deep-Learning techniques [Bône et al., 2020; Han et al., 2020b; Maillard et al., 2022; Shu and et al, 2018]. However, these methods strongly depend on the data set and on the modality they have been trained on, and might not correctly disentangle shape and appearance changes.

The Metamorphic framework [Holm et al., 2009; Trouvé and Younes, 2005; Younes, 2019] can be seen as a relaxed version of LDDMM in which residual time-varying intensity variations are added to the diffeomorphic flow, therefore allowing for topological changes. The image evolution is not only modelled by deformation, we allow adding intensity at each time for every voxel, making topological changes possible. We add intensity through the momentum (the one from the LDDMM formulation). In Metamorphosis, the definition of momentum is set as the difference between the ideal image evolution and the deformation induced by the vector fields. In other words, between the partial derivative in time of the image and the image transport by the field. Thus we also call it the *residual*. One can read Section 2.4 for more details.

1.4 My thesis's Contribution

1.4.1 Semi-Lagrangian schemes for LDDMM and Metamorphosis implementation

Even if a lot of publications displayed results with the Metamorphosis algorithm, unlike LDDMM, no public implementation was available. The first contribution of this thesis is the implementation of both Metamorphosis and LDDMM for images in 2D and 3D. To compute Metamorphosis, one must integrate over a geodesic PDE system using a numerical scheme from an initial condition. The registration is done via shooting by updating this initial condition or residual. This is challenging because classical schemes require a lot of time steps, to respect the CourantFriedrichsLewy (CFL) conditions and therefore to converge. Making the subsequent optimisation long to converge as well. I choose to use a semi-Lagrangian scheme on both images and residual to be able to keep the number of times reasonable and thus the optimisation time.

The implementation is object-oriented, which allows for versatility in usage and easy modifications. We developed it using `PyTorch` and makes use of GPU acceleration for speed gain.

The code is fully available on GitHub for anyone to use and contribute to at `https://github.com/antonfrancois/Demeter_metamorphosis`.

More About in Chapter 2, Section2.5

1.4.2 Constrained Metamorphosis: An extended model

If Metamorphosis leads to very good registrations, the disentanglement between geometric and intensity changes is not unique and it highly depends on user-defined hyper-parameters. This makes interpretation of the results hard, thus hampering its clinical usage. For instance, in order to align a healthy template to an image with a tumour, one would expect that the method adds intensities only to create new structures (*i.e.:* tumours) or to compensate for intensity changes due to the pathology (*i.e.:* oedema). All other structures should be correctly aligned solely by the deformations. However, depending on the hyper-parameters, the algorithm might decide to account for morphological differences (*i.e.:* the mass effect of tumours) by changing the appearance rather than applying deformations. This limitation mainly comes from the fact that the additive intensity changes can theoretically be applied all over the image domain. However, in many clinical applications, one usually has prior knowledge about the position of the topological variations between a healthy image and a pathological one (*e.g.,* tumour and oedema position).

In this part of the thesis, we address the difficulties associated with the Metamorphosis algorithm by proposing a framework for incorporating prior knowledge into the model. Our approach involves modifying the registration problem in a way that preserves the ability to deduce the system of PDEs that underlies the Metamorphosis algorithm. The theory of Metamorphosis is quite delicate to handle, but in the case of images the principle is to set a registration problem, potentially inspired from an optimal control one [Miller et al., 2015], and deduce a system of PDE from it. We propose an adjustment to Metamorphosis that allows adding constraints on the registration problem by also matching some given priors. We call this framework Constrained Metamorphosis. We present two specific types of priors that can be incorporated into the model: (1) a growing mask generated from a given segmentation, and (2) a field that guides the deformation in a desired direction.

Our implementation of the Constrained Metamorphosis algorithm builds upon our previous implementation of Metamorphosis and is designed to be flexible enough to allow any user to incorporate its own additional priors. We demonstrate the effectiveness of our approach through experiments on glioblastomas, using BraTS datasets, comparing with state-of-the-art methods.

More About in Chapter 3,
Mathematical framework: Section 3.2.3
Validation on toy-Example: Section 3.2
Validation on real data: Section 3.3

1.4.3 Segmentation using TDA

In collaboration with Raphaël Tinarrage, we developed a tumour segmentation tool using Topological Data Analysis (TDA). We aim to detect characteristic components within the FLAIR and T1ce modalities by taking assumptions on their topologies. Using TDA we are able to select and find connected components and cycles that compose a glioma without the usage of any statistical tool. Additionally, by automating the segmentation process, we gain insight into the location, shape, and intensity of glioblastomas in MR images, which gives us a deeper understanding of glioblastoma characteristics. In fact, by having a robust segmentation method, we lay the foundation for building a diffeomorphic glioblastoma atlas.

The code is fully available on GitHub for anyone to use and contribute to at: `https://github.com/antonfrancois/gliomaSegmentation_TDA`.

More About in Chapter 4

1.4.4 Toy-example construction

I proposed a 2D toy example to mimic the registration of cancer growing inside an organ pushing the surrounding by "mass effect". The process is divided into three steps: first, a deformation is applied to the initial image using a bspline field generated from randomly chosen control points; second, a random deformation is generated around a small ball to mimic the growth of necrosis, using a bspline field again. The process is repeated a number of times, and the resulting deformation is applied to the image from the first step. Finally, we add intensity changes around cancer to mimic Oedema infiltration. The toy example can be customised by randomly selecting four parameters: the control matrix for the background deformation, the initialisation point for cancer, and the number and amplitude of the bspline field for the cancer growth deformation.

More About in Chapter 3, Section 3.2

1.4.5 Visualisation tool

Both the LDDMM and Metamorphosis algorithms produce a sequence of images or a temporal image, that can be interpreted as a visual representation of the displacement from the source to the target images. However, due to the challenges associated with handling 3D images and the lack of readily available Python plugins for visualising both temporal images and deformations simultaneously during debugging, we developed a tool that addresses these issues.

Our tool, which is implemented using the `vedo` library [Musy et al., 2023], enables to display

- a single 3D temporal image
- the comparison of two temporal images by displaying them side by side and overlaying them just with a click.
- the result of a Metamorphosis or Constrain metamorphosis (CM) optimisation, including useful information like parameters that were used or the number of pixels where the deformation tore the background (if the deformation is not diffeomorphic)
- In addition to the registration, one can also visualise the deformation as a collection of 3D vectors.

One can see a demonstration in this video: `https://www.youtube.com/watch?v=-qzQ1B6DdSg` and find its code on GitHub at antonfrancois/Demeter_metamorphosis.

1.5 Chapter organisation

Chapter 2 covers all the Background knowledge about LDDMM and Metamorphosis, starting with a small hiatus giving elements of Lie algebra which are helpful to understand core concepts of differential geometry. I take this opportunity to explain briefly Lie Exponential as a tool for registration. Then I give a cautious demonstration of LDDMM, slowly introducing the numerous properties. I continue by expatiating about the momentum conservation properties and how it allows us to use geodesic Shooting. Having digested all this theory, one is ready to understand the Metamorphic framework and its geodesic derivations with ease. I finish the Chapter by explaining my implementation of both aforesaid methods using semi-Lagrangian schemes (see Contribution 1.4.1) and giving some first results.

Chapter 3: introduces the Constrained Metamorphosis (CM) framework (see Contribution 1.4.2). After the mathematical formulation, I present the toy example (see Contribution 1.4.4) I used for debugging and making a first validation. Then I present CM performances on MRI sequences on two different applications. First, in continuation of this thesis's original motivation, we present it as a tool for registering a healthy brain template on a brain with glioblastoma. Second, we took the BraTSReg 2022 Challenge [Baheti and et al., 2021]] as an opportunity to test CM on follow-up surgeries alignment.

Chapter 4: presents a glioblastoma segmentation method using TDA. We first make a short review of methods doing Glioblastoma segmentation using Machine Learning. Then we cover some papers using TDA for MRI analysis. Then, in Subsection 4.2, we make a short introduction to the different TDA concepts useful for analysing MRI with TDA. We are now ready to present our segmentation algorithm in Subsection 4.3. Finally, we present our results and discuss the improvement we want to implement, in Subsection 4.4.

General Conclusion & Perspectives : Finally we elaborate further on the new kind of atlas presented in Subsection 1.2.3. We elaborate more on how the work we presented in this thesis can serve to develop such an atlas.

1.6 Published Articles and Preprints

1. A. François, P. Gori, and J. Glaunès. Metamorphic image registration using a semi-lagrangian scheme. In *SEE GSI*, 2021

2. A. François, M. Maillard, C. Oppenheim, J. Pallud, I. Bloch, P. Gori, and J. Glaunès. Weighted metamorphosis for registration of images with different topologies. In *WBIR*, pages 8–17, 2022

3. M. Maillard, A. François, J. Glaunès, I. Bloch, and P. Gori. A deep residual learning implementation of metamorphosis. In *IEEE ISBI*, 2022

1.7 Comic: Diffeomorphic Shape Space

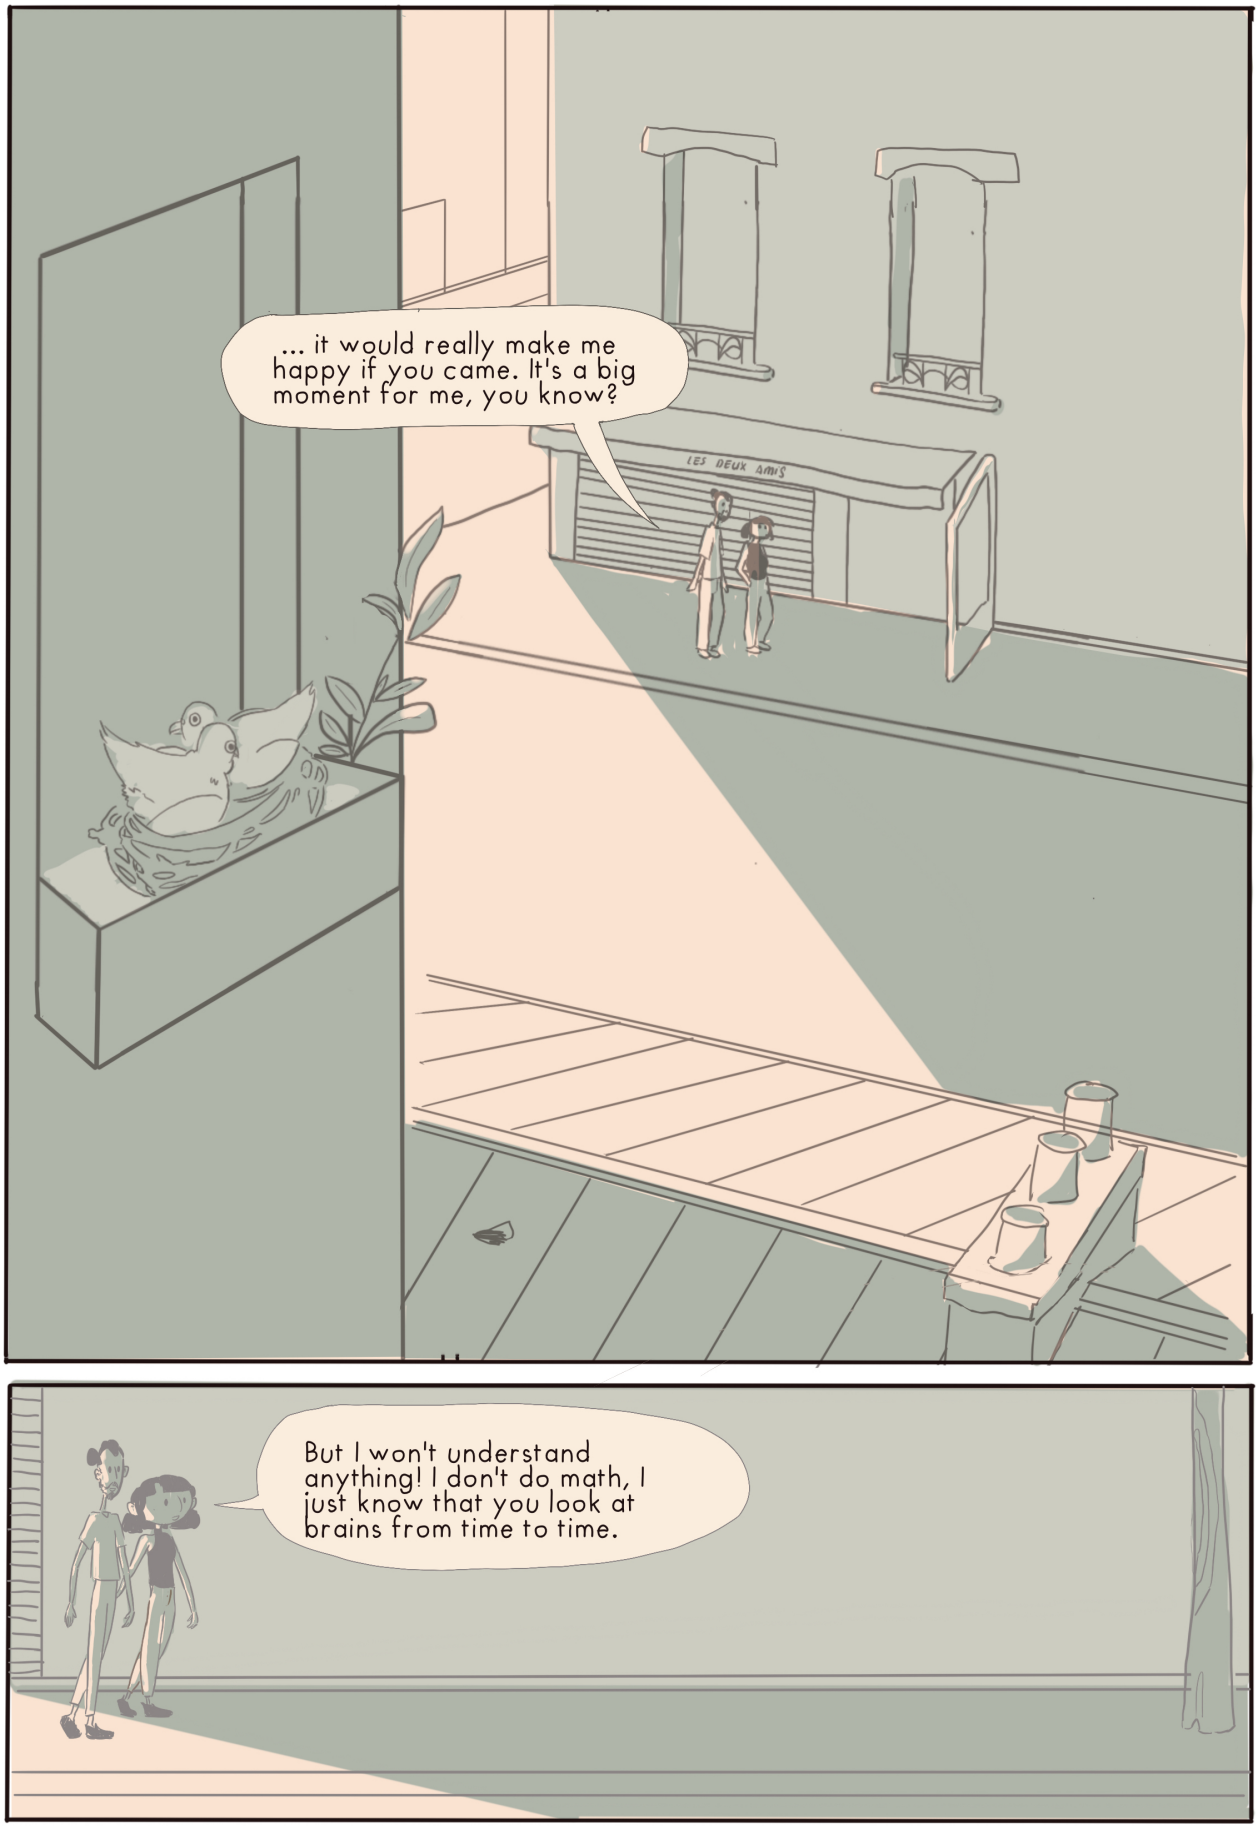
... it would really make me happy if you came. It's a big moment for me, you know?
LES DEUX AMIS
But I won't understand anything! I don't do math, I just know that you look at brains from time to time.

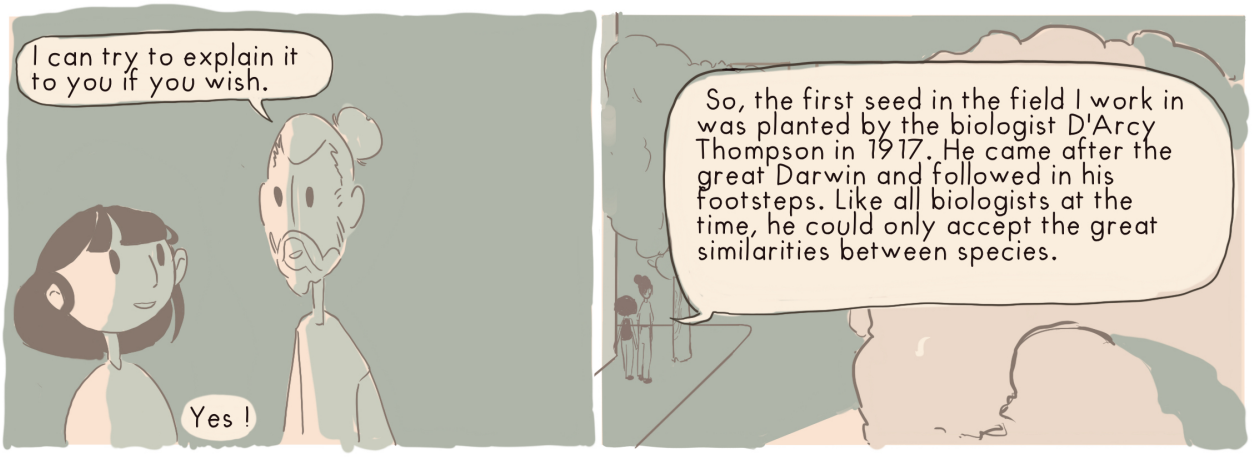

He noticed and created a first method for comparing species by comparing the deformation necessary to align them perfectly. Of course, he didn't have a computer at his disposal and he didn't limit himself to linear deformations, meaning: zoom, rotation, and shearing

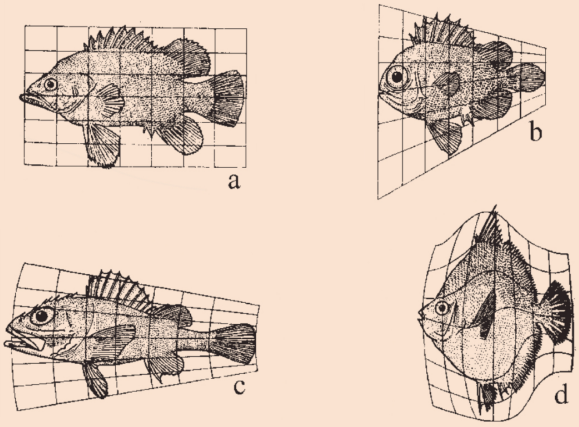

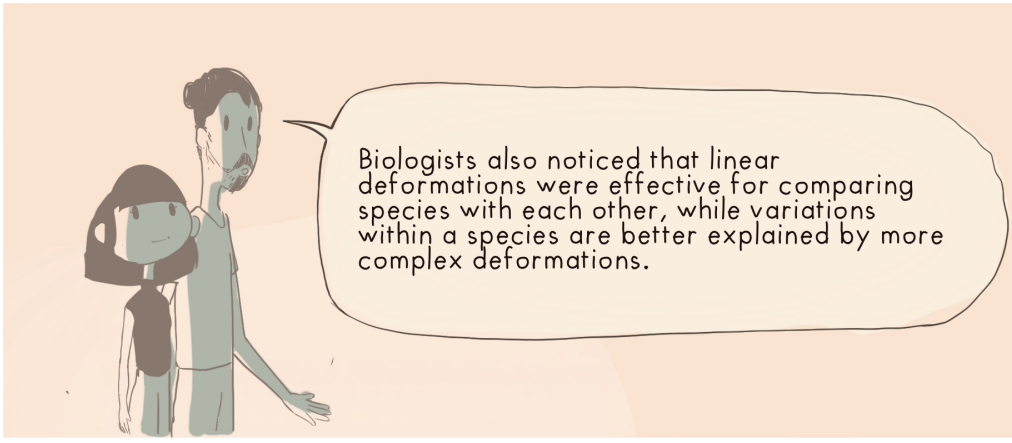

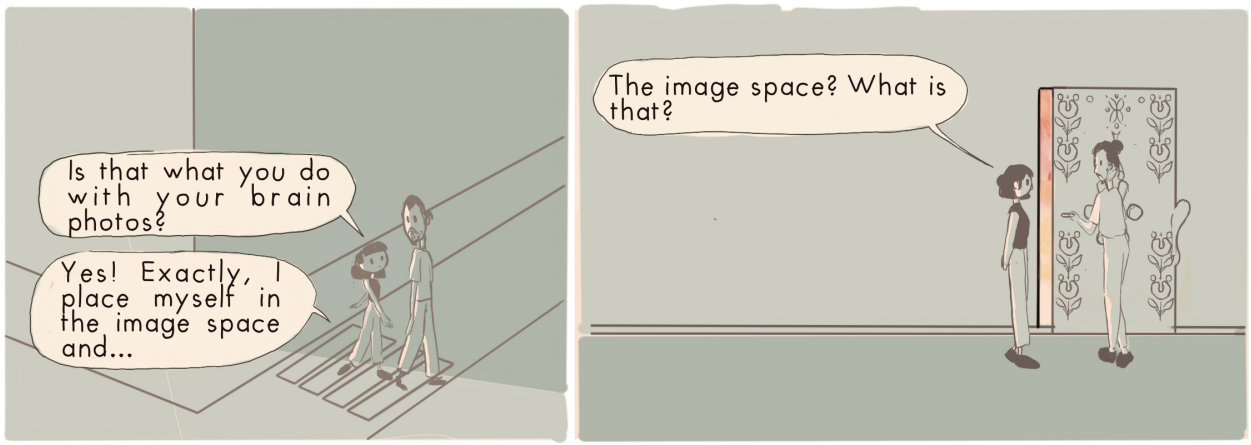


It's a high-dimensional space where every point is a possible image variation. For example, for a 100x100 pixel image, there are 10,000 dimensions.


Yes, it's mind-boggling. And if you randomly pick an image, it's very likely to look like this:


I come on, let's try to find a brain.


We'll never get there, it's too tiring...
AAAA
HAAA
HAAAAAA


Outch !
No, my name is Ulf, Grenander
Wow, how do you manage to move around so easily?
Thanks to our geodesic surfboards!
What's that? And what are those lines that appeared?
A group of diffeomorphic deformations!
MICHEL MILLER
As you saw, it's difficult to move around in the space of images, which is of far too high dimension. Since we are interested only in images of organs that have the same overall structure, we know that we can reach one image by deforming another.
Starting from an image, we will only reach a portion of the space of images, but it is precisely those that interest us! Remarkably, this subset of the space is a manifold, that is, a surface on which we know how to move! Following a deformation is like following a path in the space of images!
Our surfboards are precisely for that, they follow a straight line on this space that we call geodesic. And the best part about all of this is that they measure the distance they have traveled based on the "amount" of deformation, which allows us to compare the images with each other.
Great, do you know where we can find these geodesic surfs ?
Haha, we're in the world of math, everything is in your head, just think about it really hard!


That's right! By the way, if you want to get out of here, you'll find an exit in this image.
Thank you for the information!
But won't we get lost?
Good question! But no, don't worry, this method allows us to build the space of shapes, also known as an "Atlas".
From a series of images, we construct an image, the "Template", which is at the center of the generated surface. That is, at the point where the sum of distances to reach each image is the shortest.
An average image is not enough to make an Atlas, is it?
Yes, you are right. We find the Template, but we also store all the deformation paths from this point to all the other images.
Well, it's time to get out of here. The image seems to be in that direction.


C
C
C
C
C

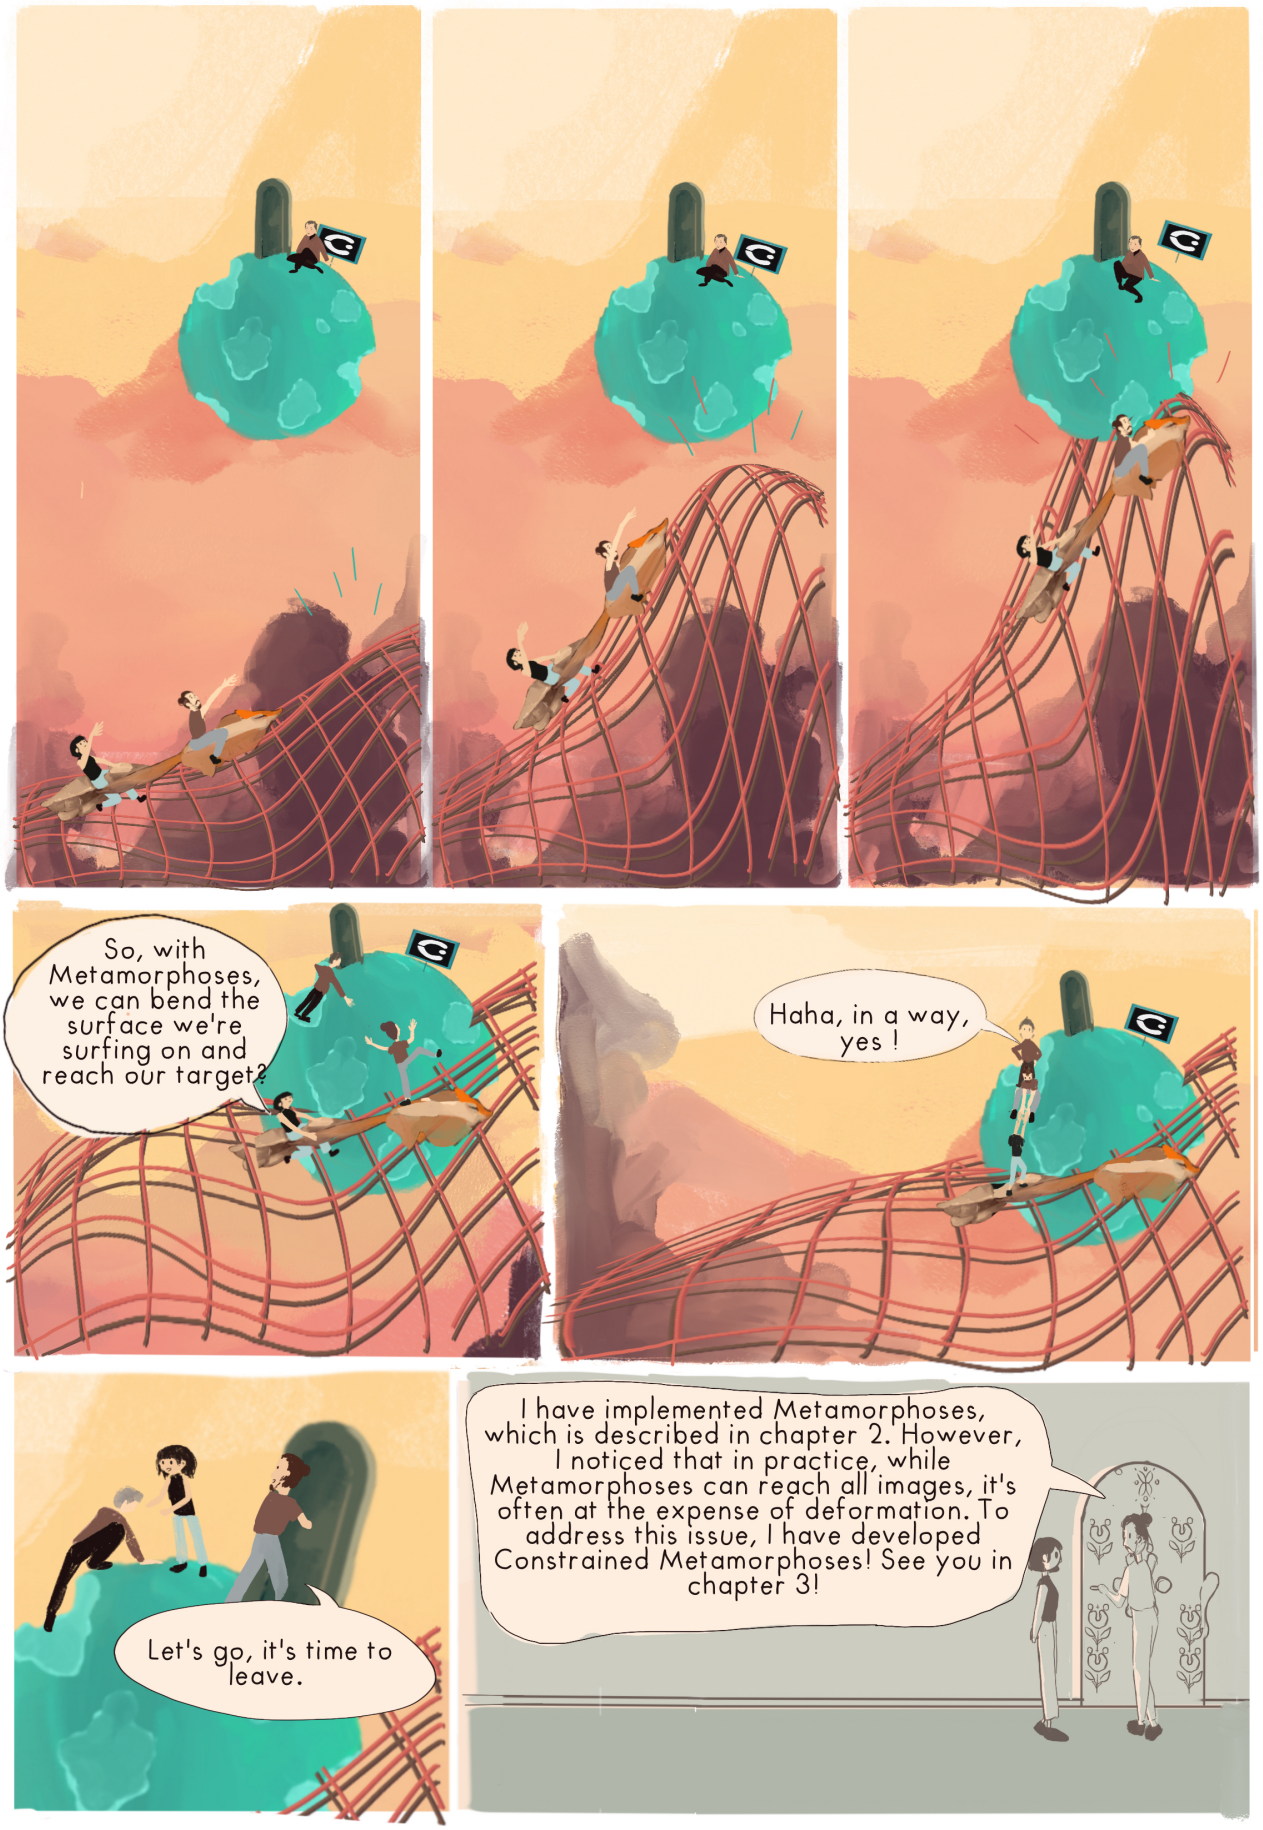
So, with Metamorphoses, we can bend the surface we're surfing on and reach our target?
Haha, in a way, yes !
Let's go, it's time to leave.
I have implemented Metamorphoses, which is described in chapter 2. However, I noticed that in practice, while Metamorphoses can reach all images, it's often at the expense of deformation. To address this issue, I have developed Constrained Metamorphoses! See you in chapter 3!

2

Background : LDDMM & Metamorphoses

In this chapter, we will study the LDDMM and the Metamorphosis frameworks which are the main theories on which is based most of my work. To give some first definitions and cover some concepts we start with a short introduction to Lie groups. At the section end, one will be able to see a way to produce a diffeomorphic map from one manifold to another. Armed with this knowledge, we explain the LDDMM framework. Rather than defining it shortly, I decided to construct it slowly giving some insight into most of its many properties. Then we introduce the Metamorphosis framework in the context of images and present our implementation that was partially presented in François et al. [2021].

2.1 A brief background on Lie groups and algebras

If you read historical papers about Computational Anatomy, you will often encounter Lie algebras, leading to surprisingly convenient properties. Especially for someone like me, three years ago, that had no clue what it was about. In this thesis, deformations will be in the group of diffeomorphisms, which is infinitely dimensional and not smooth enough to be a Lie group. Therefore, I could avoid talking about the Lie theory. However, many concepts we will cover, share similarities with Lie groups. I still decided to go through some notions as their understanding helps grasp the shape space theory. Cédric Villani said, talking about Grothendieck's work, that a way of solving complex mathematical problems is to lift them

to a more general formulation until it is miraculously solved. I found that it is the case here. This section's aim is not to give a proper introduction to Lie algebras, but rather to present some interesting properties and consequences that I find helpful for comprehension.

2.1.1 General Knowledge

Most of this section is a short agglomerate of notions and properties from the book "Introduction to Smooth Manifolds" by John Lee [2012]. I won't formally cover the definitions of smooth maps, smooth manifold, and topological atlas, nor give proof here. However, I strongly advise the unfamiliar reader to take some time manipulating these notions, which are very well introduced in the cited book.

2.1.1.a Lie Groups

Lie groups are a natural tool for modelling transformations under which an object is invariant, endowed with the group operation of composition, called the symmetry group. An example could be the Lie group of $d \times d$ matrices of 3D rotations, modelling the position of a rigid body with a fixed centre of mass. It is a group under matrix multiplication.

Definition 2.1 (Lie group)**.** *A* Lie group *is a smooth manifold G (without boundary) that is also a group, with the property that the multiplication map j and inversion map i are both smooth, where j and i are given by:*

$$\begin{aligned} j : G \times G &\to G & \qquad i : G &\to G \\ j(g,h) &= gh & \qquad i(g) &= g^{-1}. \end{aligned} \tag{2.1}$$

A Lie group is, in particular, a ***topological group*** *(a topological space with a group structure such that the multiplication and inversion maps are continuous).*

Lie groups also involve group actions, which are external composition laws of the group on a set. If G is a group and M is a set, a *left action* of G on M is a map

$$\begin{aligned} G \times M &\to M \\ (g, m) &\mapsto g \triangleright m \end{aligned} \tag{2.2}$$

that satisfies the usual group action properties:

$$\begin{aligned} \forall m \in M, \quad & e \triangleright m = m \\ \forall (g, g') \in G \times G, \forall m \in M \quad & g' \triangleright (g \triangleright m) = (g'g) \triangleright m. \end{aligned} \tag{2.3}$$

The group action will be written $\triangleright$ in this manuscript. The usual notation is $g \cdot m$, as for the scalar product, and sometimes the derivation operator as well. As I find this notation mix confusing, we will use the dot only for the Euclidean scalar product.

A *right action* is defined in a similar way as a map $M \times G \to M$ and appropriate composition laws are such that:

$$\begin{aligned} \forall m \in M, \quad & e \triangleright m = m \\ \forall (g, g') \in G \times G, \forall m \in M \quad & (m \triangleright g) \triangleright g' = m \triangleright (gg'). \end{aligned} \tag{2.4}$$

A right action can always be converted to a left action by the trick of defining $g \triangleright m$ to be $m \triangleright g^{-1}$, and a left action can similarly be converted to a right action. This will be used all along the thesis for composing vector fields and images (see Section 2.5.1.a). We usually focus our attention on left actions, because their group law has the property that group elements multiplication corresponds to map composition.

There are many useful properties for Lie groups to study, like isometry groups, transitive orbits, and free action, but only the notion of orbit will be used in the manuscript.

Definition 2.2 (orbit)**.** *For each $m \in M$, the orbit of m, denoted by $G \triangleright m$, is the set of all images of m under the action by elements of G:*

$$G \triangleright m = \{g \triangleright m : g \in G\} \tag{2.5}$$

In other terms the orbit can be seen as the collection of objects attainable from m using elements of G.

2.1.1.b Vector Fields

A student in mathematics or physics has encountered vector fields in his study of multivariate calculus. In that setting, a vector field on an open subset $\Omega \subset \mathbb{R}^n$ is simply a continuous map from Ω to $\mathbb{R}^n$, which can be visualized as attaching an 'arrow' to each point of Ω. We can transpose this idea to smooth manifolds, where we attach to each point m an arrow part of the tangent space at m. We can see the vector field as a map $M \to TM$.

Definition 2.3 (Vector fields on a smooth manifold)**.** *Let be M a smooth manifold, a vector field v on M is a continuous section of the projection $\pi : TM \to M$, where TM is the tangent bundle of M, such that for all $m \in M$:*

$$\pi \circ v_m = m. \tag{2.6}$$

We will denote $\mathcal{V}(M)$ the set of all smooth vector fields on M

We note v_m the value of v at m instead of $v(m)$ to be consistent with the notation for the tangent bundle elements. In other terms, a vector field is such that $v_m \in T_m M$ for each $m \in M$.

As for the vector space case, one should visualise smooth vector fields on manifolds as arrows attached to each point of M, chosen to be tangent to M, and varying continuously from point to point. Even if we set a different definition of a vector field, one can still see a vector field as an order-1 linear differential operator and write $v = \sum_{i=1}^{d} v^i(x) \frac{\partial}{\partial x_i}$. The main difference is that at each point m one has to use the specific local coordinate system $\{x_i\}$ defined by $T_m M$. We will redefine the vector field in a vector space in Definition 2.10.

In general, it is not obvious to differentiate functions with vector fields or even vector fields together. Indeed, if $F : M \to N$ is a smooth map and v a vector field on M, then for each point $m \in M$; we obtain a vector $dF_m(v_m) \in T_{F(m)}N$ by applying the differential of F to v_p. However, this does not in general define a vector field on N. Indeed if F is only injective, one can not decide which vector to assign to a point without any antecedent. If F is only surjective, then there might be points without any vector (see Figure 2.1). Thus we define the relation in between two vector fields defined on different manifolds by F using its derivative as it is a map such that $dF : TM \to TN$.

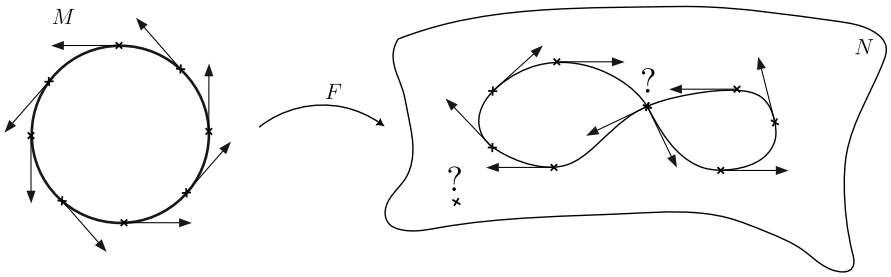

Figure 2.1: The differential of a vector field is not always a vector field

Definition 2.4 (F-related vector fields)**.** *Let M and N two manifolds with a vector field v (resp. w) defined on M (resp. N). Let suppose that it exists $F : M \mapsto N$ such that for each $m \in M$, $dF_m(v_m) = w_{F(m)}$. We say that v and w are F-related.*

In particular, if $F : M \to N$ is a diffeomorphism, for every v defined on M there is a unique smooth vector field on N that is F-related to v (see Proposition 8.19 in Lee [2012]). Conversely, if F is a diffeomorphism, we are sure that objects keep their topology (*e.g.,* in Figure 2.1 the circle can't be mapped to a lemniscate).

A natural object generated by smooth vector fields on smooth manifolds are *integral curves*: smooth curves within the manifold having a velocity given by the vector field at each point.

Consider v is a smooth vector field* (*i.e.:* $v \in \mathcal{V}(M)$) on M. An integral curve is defined as a smooth curve $\gamma : [0,1] \to M$ for all $t \in [0,1]^\dagger$ passing at $m \in M$ at a time t_0 and such that:

$$\gamma(t_0) = m; \qquad \dot{\gamma}(t) = v_{\gamma(t)}, \forall t \in [0,1]. \tag{2.7}$$

This means that at every time $t \in [0,1]$ the integral curve has a velocity equal to a vector in $T_{\gamma(t)}M$ and $\dot{\gamma}(t)$ being the derivative of γ at t is also the "direction in which γ points".

We are now equipped to talk about flows, which one can see as a collection of integral curves on M associated with a vector field v. For each $t \in \mathbb{R}$, we can define a map $\Phi(t) : M \to M$ by taking all points $m \in M$ and letting them follow the directions given by v on their respective integral curves. We will use indistinctively the notations $\Phi_t(x) = \Phi(t,x)$ depending on the context. Each map Φ_t "slides" the manifold along the integral curves. The Translation lemma implies that $t \mapsto \Phi_{t+s}(m)$ is an integral curve of v starting at $m' = \Phi_s(m)$ and thus $\Phi_t \circ \Phi_s(m') = \Phi_{t+s}(m)$. It is time to define the global flow.

Definition 2.5 (global flow)**.** *A global flow on M is a continuous left $\mathbb{R}$-action on M; that is, a continuous map $\Phi : \mathbb{R} \times M \to M$ satisfying the following properties for all $s, t \in \mathbb{R}$ and $m \in M$:*

$$\Phi(t, \Phi(s,m)) = \Phi(t+s, m), \qquad \Phi(0,m) = m \tag{2.8}$$

Remarks:

*Here smooth implies $\mathscr{C}^\infty$, however in a less restricted domain such as differential geometry, one could use $v \in \mathscr{C}^{r-1}, r \geq 2$ and M being itself of class $\mathscr{C}^r$

†We could have defined the time interval to be any interval $\mathcal{T}$

1. The global flow is also sometimes called a *one-parameter group action* to emphasize the fact that as left $\mathbb{R}$-action on M it depends mostly on time. They were introduced to define infinitesimal transformation along the flow.
2. As we restricted the time interval to $[0,1]$ we could have defined $\Phi : [0,1] \times M \to M$
3. As is the case for any continuous group action, each map $\Phi_t : M \to M$ is a homeomorphism, and if the flow is smooth, Φ_t is a diffeomorphism

2.1.1.c Lie algebra & Lie brackets

Let v and w two smooth vector fields on M. The Lie bracket is an operator that assign two vector fields defined on M in a way to obtain a vector field as well, and is noted $[v,w]$. The key intuition being that $[v,w]$ is the derivative of w along the flow generated by v. Recalling that vector fields are derivative operators, to construct such operation, it would have been natural to try directly applying vw to a function. Indeed, given a smooth function $f : M \mapsto \mathbb{R}$, vf is still a smooth function (proof in Lee [2012]). It follows that we can apply w to this smooth function and still obtain a smooth function $w(vf)$. However the map $f \mapsto wvf$ does not follow the product rule for derivation as the Schwartz theorem can not be generalized to this case (*i.e.:* the second derivatives $w(vf)$ and $v(wf)$ might differ). There is few equivalent ways to define the Lie derivative and its brackets, we will see the vector fields as derivations and start by introducing the Lie derivative noted $\mathscr{L}_v$.

We saw that, each $v \in \mathcal{V}(M)$ is a map such that $M \to TM$ and is a differential operator acting on smooth functions f. Accordingly, we define, $v(f)$ to be another function whose value at a point m is the directional derivative of f at m in the direction $v(m)$. In general, the commutator $\delta_1 \circ \delta_2 - \delta_2 \circ \delta_1$ of any two derivations δ_1 and δ_2 is again a derivation, where $\circ$ denotes composition of operators. Getting back on the Lie derivative for vectors, it defines a derivative in the sense that it is linear and verify the Leibniz formula. Also $\mathscr{L}_v\mathscr{L}_w - \mathscr{L}_w\mathscr{L}_v$ is a derivative on M and identify with the definition of a Lie Bracket.

Definition 2.6 (Lie Bracket)**.** *We define the Lie Bracket operator* $[v,w] : \mathcal{V}(M) \times \mathcal{V}(M) \mapsto \mathcal{V}(M)$ *as*

$$[v,w]f = vwf - wvf.$$

satisfying:

- *bi-linearity*
- *anti-symmetry, i.e.:* $[v,w] = -[w,v]$.
- *Jacobi identity, i.e.:* $[u,[v,w]] + [w,[u,v]] + [v,[w,u]] = 0$ *with* $u,v,w \in \mathscr{C}^\infty$

We have the property that the Lie bracket of any pair of smooth vector fields is still a smooth vector field. With this definition, the derivative composition rule is respected and the Lie derivative is a differential operator Ślebodziński [1931] noted $\diamond$

$$\mathscr{L}_v f(m) = df(m)\left.\frac{d}{dt}\right|_{t=0} \Phi_t(m) = \mathrm{d}f(m)(v(m)) = v \diamond f. \tag{2.9}$$

If we put a Riemannian structure on M, then we can write similarly to Equation 2.37:

$$\mathscr{L}_v f(m) = v \diamond f = v \cdot \nabla f(m) = \sum_{i=1}^{d} v^i(m) \frac{\partial f}{\partial x_i}(m). \tag{2.10}$$

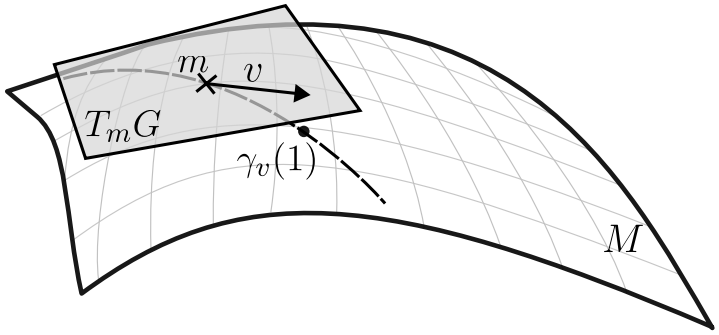

Figure 2.2: Vector field exponential visualisation

with $v^i(m)$ being the i-th coordinate of the vector at m.

Note that in general the Lie derivative is defined for a differential form α of degree p but in this thesis, we only study cases where $p \leq 1$, $p > 1$ begin the tensor case.

Definition 2.7 (Lie Algebra over $\mathbb{R}^d$)**.** *A* Lie algebra *over $\mathbb{R}^d$ is a vector space $\mathfrak{g}$ together with the binary operation $[\cdot,\cdot] : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$.*

To sum up, a Lie group both a geometrical object: a manifold (*e.g.,* A circle) with an algebraic structure, meaning that for all operations $g, g' \in G$, will stay in G, $g \cdot g' \in G$. They provide a way to express the concept of a continuous family of symmetries for geometric objects. By differentiating the Lie group action, you get a Lie algebra action, which is a linearisation of the group action. As a linear object, a Lie algebra is often a lot easier to work with than working directly with the corresponding Lie group. For example, the space of all smooth vector fields on a smooth manifold M is a Lie algebra under the Lie bracket. Another important example: if G is a finite-dimensional Lie group, the set of all smooth left-invariant vector fields on G is a Lie subalgebra of the smooth vector field space. Therefore, G is a Lie algebra. This last example case is called *the* Lie algebra of G, usually noted Lie(G). In addition, Lie(G) is finite-dimensional, and in fact, has the same dimension as G itself.

2.1.2 Vector field exponential

Let us study, as an example, the vector field exponential, which is a way to construct a diffeomorphism from any vector field v. The idea here is that we build the group G around v. We assume that the vector field v is part of the tangent space of an abstract manifold G at identity (*i.e.:* $v \in T_{\mathrm{Id}}G$). Then by following the integral curve associated with v.

Definition 2.8 (Exponential map)**.** *Let be v a vector field. The exponential map is a map from the Lie algebra of G*

$$\begin{aligned} \exp : T_{\mathrm{Id}}G &\to M \\ \exp(v) &= \gamma(1) \end{aligned}$$

where $\gamma : [0,1] \to M$ is the unique integral curve of M whose tangent vector at the identity is equal to v.

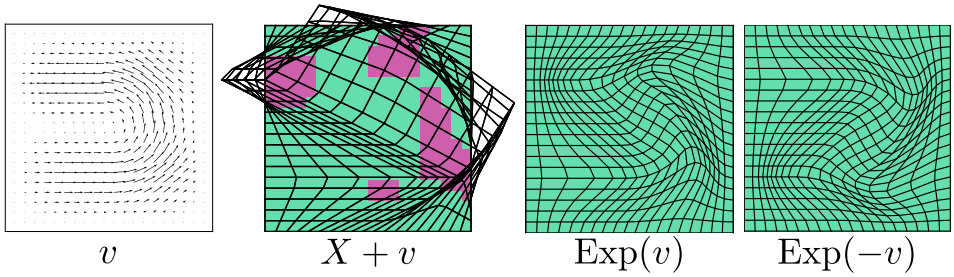

Figure 2.3: Exponential map effects on a vector field. Grids show deformations, the red and green scares behind indicate the sign of the jacobian determinant. Red presence indicates that the deformation is not diffeomorphic.

In the particular case were $M \subset \mathbb{R}^d$ we can compute the exponential as described in Ashburner [2007]; Vercauteren et al. [2009], where a fast algorithm is proposed. The idea is to take any vector v and scale it down by a factor l such that $\mathrm{Id} + lv$ is a bijective deformation. To obtain a defomation with the same amplitude as $\mathrm{Id} + v$, one can solve the ordinary differential equation:

$$\dot{\varphi}_t = (lv) \circ \varphi_t, \quad l \leq 1. \tag{2.11}$$

Indeed the iterative composition of bijective transformations returns a bijective one. The Lie exponential also have the properties that for $t, s \in T$ and $t + s \in \mathcal{T}, \exp((t + s)v) = \exp(tv) \circ \exp(sv)$. Thus for any integer k

$$\exp(v) = \exp(k^{-1}v)^k \tag{2.12}$$

where the power operation is the composition of spatial transformations. This leads to the algorithm and Code 2.1. The main idea behind is that starting from a first approximation of $\varphi_0 = \exp(2^{-N}\mathbf{v}) \simeq \mathrm{Id} + 2^{-N}\mathbf{v}$ and then iterating for $k < N \in \mathbb{N}, \varphi_k = \varphi_{k-1} \circ \varphi_{k-1}$.

Source Code 2.1: Fast exponential integrator.

```
def fast_exp_integrator(in_vectField,N):
        """ Compute Exp(in_vectField) """
        # First order approximation exp(2**(-N) * v) ~ Id + 2**(-N) * v
        in_vectField *= 1/(2**N)
        grid_def = id_grid + sign*in_vectField

        # Field Composition divided 2**N times
        tmp_grid = torch.zeros(grid_def.shape)
        for n in range(N):
                field = tb.grid2im(grid_def-id_grid)
                tmp_grid = grid_def.clone()
                interp_vectField = grid_sample(field,tmp_grid)

        grid_def += tb.im2grid(interp_vectField)

        return grid_def
```

We add a small feature in order to choose automatically the scaling needed. If we scale a displacement field by 2^N we will need to compose it N time. Arguably, on one side the higher N is, best the approximation is, on the other side, the higher N is the more we apply successive interpolation and introduce computational errors. In any case, for speed optimisation reasons, it is interesting to find the smallest N possible for v to induce a one-to-one transformation We implement our methods with `PyTorch` where the interpolation function `grid_sample` require a deformation grid defined on a $\Omega = [-1, 1]^d$ domain. So we will use this convention in all the following.

Proposition 2.1. *For any smooth vector field v defined on a grid $[-1, 1]^2$ with $H \times W$, the related transformation $\varphi^v(x) = x + 2^{-N}v(x)$ is bijective if*

$$N = \max\left(1, \left\lceil \frac{\log(\max_x(\|v(x)\|_\infty)) - \log(\min(1/H, 1/W))}{\log(2)} \right\rceil\right) \tag{2.13}$$

Proof. A sufficient criterion for φ^v to be bijective, v being smooth, is

$$\max_x \left\|2^{-N}v(x)\right\|_\infty \leq \min(1/H, 1/W) \tag{2.14}$$

$$\Leftrightarrow N \geq \frac{\log(\max_x(\|v(x)\|_\infty)) - \log(\min(1/H, 1/W))}{\log(2)} \tag{2.15}$$

□

Remark 1: *This proposition was written for $d = 2$ but it is straightforward to adapt to any other dimension.*

Remark 2: *In Vercauteren et al. [2009], authors advise to choose N such that $\|2^{-N}v(m)\|_\infty \leq 0.5$ but it is for a grid $[0, H] \times [0, W]$ with step one.*

We demonstrate the fast vector field exponential in Figure 2.3 where we took a vector field v generated by splines. We show that the basic deformation, which consists in adding v to a regular grid X, is not diffeomorphic (to say the least). Then we demonstrate that $\exp(v)$ has a positive determinant of Jacobian at every pixel $x \in X$. One can verify the interesting property that $\exp(v) \circ \exp(-v) = \mathrm{Id}$.

The Lie Exponential is a suitable method for generating diffeomorphic deformations from a single vector field and has been used for registration for DARTEL Ashburner [2007] or DEMON-related algorithms for example Vercauteren et al. [2009]; Pennec et al. [1999]; Mansi et al. [2010]. The previous methods propose to deduce a diffeomorphism from a single vector field, in the rest of the manuscript, we aim to register objects using more complex deformations generated by time-dependent vector fields. It is reasonable as it is way less computationally and memory intensive. Also, most of the modern research dealing with the task of registration also estimates such generated deformations using machine learning [Rohé et al., 2017; Yang et al., 2017; Balakrishnan et al., 2019; Mansilla et al., 2020; Niethammer et al., 2019; Mok and Chung, 2020].

2.1.3 Finite versus Infinite-dimensional Lie Groups

One can also define a Lie group from an infinite-dimensional Manifold. If most definitions remain the same, the big difference is that they are not locally compact. This can be problematic, indeed some results in finite dimensions are not true in general in infinite dimensions [Schmid, 2010], for example:

- There is no Implicit function theorem or inverse function theorem in infinite dimensions. Thus it is not always possible to represent an intricate function through a collection of simpler analytic ones.
- If G is a finite-dimensional Lie group, the exponential map $\exp : \mathfrak{g} \to G$ is a local diffeomorphism from a neighbourhood of zero in $\mathfrak{g}$ onto a neighbourhood of the identity in G; hence exp defines canonical coordinates on the Lie group G. As we cannot ensure local compactness, this is not true in infinite dimensions.
- If $f_1, f_2 : G_1 \to G_2$ are smooth Lie group homomorphism between finite-dimensional Lie groups, then f is smooth. This is not true in infinite dimensions.
- If $f : G \to H$ is a continuous group homomorphism between finite-dimensional Lie groups, the f is smooth. It is not always the case in infinite dimensions.
- If $\mathfrak{g}$ is any finite-dimensional Lie algebra, then there exists a connected finite-dimensional Lie group G with $\mathfrak{g}$ as its Lie algebra; that is, $\mathfrak{g} \approx T_{\mathrm{Id}}G$. This is not true in finite dimensions.

Here is a classical example of an infinite Lie group with a Lie algebra. Let M be a finite-dimensional manifold, $\mathscr{C}^\infty(M)$ the smooth function on M along with the group operation is a Lie group[‡]. We can set the Lie algebra to be $\mathfrak{g} = T_{\mathrm{Id}}\mathscr{C}^\infty(M) \approx \mathscr{C}^\infty$ with the trivial bracket $[v, w] = 0$ and $\exp = \mathrm{Id}$. Moreover, If we complete these spaces in the $\mathscr{C}^k$-norm, $k < \infty$, then the completed space is a Banach-Lie group. Alternatively, if we complete the space with the H^s-Sobolev norm with $s > (1/2)\dim(M)$ then it is a Hilbert-Lie group [Schmid, 2010].

Obviously, the diffeomorphism group is a very important example of an infinite dimensional Lie group, while not being a Banach Lie group. We will give more details about this group in Sections 2.2.2 and 2.3.2.

2.2 LDDMM framework construction

In this Section, I will try to make a gentle and exhaustive Large Diffeomorphic Metric Mapping (LDDMM) introduction. The method is not only a registration technique but also a powerful tool to compare objects embedded in complex data. We already did an equation-free introduction in Chapter 1. I assume that you have a general idea of LDDMM and are thirsty to understand how can a single theory could have so many properties.

I will start my presentation with the core concept of LDDMM: the ability to compare images by a flow of diffeomorphisms through a Riemannian distance, the amount of deformation giving us the searched metric. The method, introduced for the first time three decades ago by Grenander [1993], had some time to mature. Therefore, just like water finding the straightest path to the ground, the Shape-space theory found shortcuts leading to very convenient properties and making proofs more straightforward. However, the study of the earliest papers and careful construction of the framework help us understand these properties' origins.

[‡]more precisely an Abelian Frechet Lie group

This long Section will be organised as follow: I start in Subsection 2.2.1 by summarizing the article Miller and Younes [2001], which was at that time the aggregate of major breakthroughs in the field. This study will give the foundation to construct a distance on the space of image $\mathcal{I}$ only with the help of a group of one-to-one transformations G.

Then in Subsection 2.2.2, I give some algebraic arguments on why the group of diffeomorphisms is a very good candidate for G.

I follow in Subsection 2.2.3 by parametrizing these diffeomorphisms with the help of temporal vector fields, conveniently chosen among a Hilbert space named V. V is first introduced generally with the properties we wish it had. A short hiatus about how we integrate temporal vector fields in practice is also given, as it may help to manipulate vector fields and diffeomorphisms to understand later parts. Finally, I give some consequences of the Hilbert space construction on its scalar products and give links with Eulerian and Lagrangian mechanics in physics.

Eventually we define the scalar product of V and therefore the aforesaid norm in Section 2.2.4. We introduce the Reproducing Kernel Hilbert Spaces (RKHS) that are not only perfect for our use case but are making proofs easier, give a sense of scale and come with efficient implementations. In the last part, I take the liberty to take a step away from LDDMM and introduce its multi-scale counterpart: LDDKBM.

Finally, in Section 2.2.5 I finish by presenting some results using LDDMM on images, that I use to make a first discussion. I do not detail the implementation in this Section, as is it detailed later, we use Metamorphosis in a particular case to compute LDDMM.

2.2.1 Historic geodesics construction, first properties

In this subsection, we will introduce core concepts of LDDMM, in the study of the well-written paper *Group Action, Homeomorphisms, and Matching: A general Framework* by Miller and Younes [2001]. This work lies on the foundations held conjointly by Grenander [1993]; Dupuis et al. [1998b] and Trouvé [1995].

Later contributions bringing opportune insight gave some results more directly. For example, the contribution of D. Holm made the bridge with fluid dynamics and Lagrangian mechanics (see Subsection 2.3.1). We can also mention the parametrisation of the group of transformations using Reproducing Kernels Hilbert Spaces (see Subsection 2.2.4.b). We could also choose to present LDDMM through optimal control as it is done in Miller et al. [2015, 2006]; Younes [2019]. However, while I was reading about LDDMM, many properties seemed magical to me, and going through the old literature helped me sort all of this out.

We aiming at registering images I and I' part of a functional space $\mathbb{I}$ maps $\Omega \subset \mathbb{R}^d \to [0, a], a < \infty$, where Ω is an open set. We aim at constructing a distance on images using the group of bijective transformations $G : \Omega \to \Omega$. This construction will allow us to compare only images that are attainable by an element of G. Thus for a given $I \in \mathbb{I}$ we define $\mathcal{I}_I \subseteq \mathbb{I}$ the orbit of I by G (*i.e.:* $\mathcal{I}_I \doteq G \triangleright I$). However, if $I' \in \mathcal{I}_I$, there exit $g \in G$ such that $g \triangleright I = I'$ and thus $\mathcal{I}_I = \mathcal{I}_{I'}$. Conversely, for a given G, there exist two images $I, I' \in \mathbb{I}$ such as $I' \notin \mathcal{I}_I$. Thus we have already seen that a method based on bijective deformations will restrain us to a certain area of the space $\mathbb{I}$, which depending on the application, can be a desired feature. In the following we define $\mathcal{I}$ to be equal to some $\mathcal{I}_I$ for an unspecified image I and by writing $I, I' \in \mathcal{I}$ we, therefore, assume that they belong to the same orbit.

We will start by setting all the properties a distance on $\mathcal{I}$, $d_{\mathcal{I}}$, must have to answer

our registration problem. We will follow by constructing this distance on G the space of deformations before extending it to the expanded space $\mathcal{S} = G \times \mathcal{I}$ to compare a pair of images by the amount of deformation.

2.2.1.a Distance between registered objects

For any $I, I' \in \mathcal{I}$, we aim to build a distance $d_{\mathcal{I}}$ to evaluate the amount of deformation to match I on I'. Let $\mathcal{S} = G \times \mathcal{I}$ be the set of *deformed images* (or registered objects in Miller and Younes [2001]) and G acts on $\mathcal{S}$ through the operation $h \triangleright (g, I) = (hg, h \triangleright I)$. We will define a distance on images $d_{\mathcal{I}}$ as the set distance $D_{\mathcal{S}}$ between the orbits (Id, I) and (Id, I') under the action of g.

$$\begin{aligned} d_{\mathcal{I}}(I, I') &= D_{\mathcal{S}}(G \triangleright (\mathrm{Id}, I), G \triangleright (\mathrm{Id}, I')) \\ &= \inf\{d_{\mathcal{S}}(a, a') : a = (g, g \triangleright I); a' = (g', g' \triangleright I'); g, g' \in G\} \end{aligned} \tag{2.16}$$

In Equation 2.16, one can see that the set distance $D_{\mathcal{S}}$ is defined using the distance on $\mathcal{S}$, $d_{\mathcal{S}}$, that remains theoretical for now. We will carefully craft it from the distance of G in subsection 2.2.1.b. In the previous definition, the functional $d_{\mathcal{I}}$, is close to being a distance, one needs only to ensure it respects the triangular identity. Consequently, we need one more assumption: the distance $d_{\mathcal{S}}$ must be left-invariant to the group action, namely for all $h, g, g' \in G$, all $I, I' \in \mathcal{I}$

$$d_{\mathcal{S}}(h \triangleright (g, I), h \triangleright (g', I')) = d_{\mathcal{S}}((g, I), (g', I')). \tag{2.17}$$

In an informal way, the left invariance states that the distance d takes the shortest path g from I to I' and any detour h won't affect the result. That leads to the author's first claim:

Proposition 2.2 (Miller and Younes [2001] : Proposition 1)**.** *Let $d_{\mathcal{S}}$ be a distance on $\mathcal{S}$ which is invariant by the left action of G, then the function $d_{\mathcal{I}}$, defined on $\mathcal{I} \times \mathcal{I}$ by*

$$d_{\mathcal{I}}(I, I') = \inf\{d_{\mathcal{S}}((g, g \triangleright I), (g', g' \triangleright I')) : g, g' \in G\} \tag{2.18}$$

is symmetrical, satisfies the triangular inequality and is such that $d_{\mathcal{I}}(I, I) = 0$ for all I.

We can remark that it is direct that one can write Equation 2.18 as $d_{\mathcal{I}}(I, I') = \inf\{D_{\mathcal{S}}((\mathrm{Id}, I), (g, g \triangleright I')) : g \in G\}$ since d is invariant by left composition (*i.e.:* $d_{\mathcal{S}}((g, g \triangleright I), (g', g' \triangleright I')) = d_{\mathcal{S}}((g^{-1}g, g^{-1}g \triangleright I), (g^{-1}g', g^{-1}g' \triangleright I') = d_{\mathcal{S}}((\mathrm{Id}, I), (g'', g'' \triangleright I'))$). It means that a pair of images are compared using only one transformation g. This leads to parametrisations helping distinguish two ways of comparing objects :

- **Subject to subject comparison** (or Homogenous) where we match directly the two objects $I, I' \in \mathcal{I}$ with the best elements $g, g' \in G$ which register I and I' to the same object such that $g^{-1} \triangleright I = g'^{-1} \triangleright I'$ again its specific distance could be:

$$d_{\mathrm{hom}}(I, I') = \inf\{d_{\mathcal{S}}((\mathrm{Id}, I), (g, g \triangleright I')) : g \in G, g \triangleright I = I'\}$$

- **Template based comparison** with a reference representative object $\tilde{I}$ in $\mathcal{I}$ (*i.e.:* a template) we compute the shortest paths g, g' which register I and I' to the template such that $g \triangleright \tilde{I} = I$ and $g' \triangleright \tilde{I} = I'$ and so the distance could be written

$$d_{\mathrm{temp}}(I, I') = \inf\{d_{\mathcal{S}}((g, g \triangleright I), (g', g' \triangleright I')) : g, g' \in G, g^{-1} \triangleright I = g'^{-1} \triangleright I' = \tilde{I}\}$$

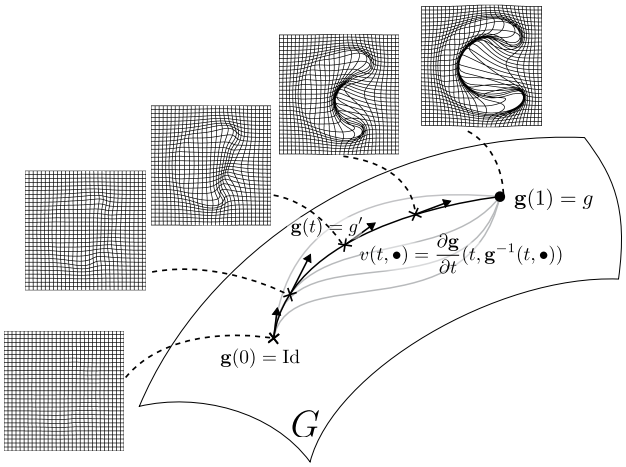

Figure 2.4: How the group G is also a Manifold. In particular, we can define paths $\mathbf{g} \in G$, the shortest being a geodesic.

These two points of view are common in shape space analysis and show the method's versatility. We can compare a pair of images directly or we can compare them having the context of a dataset. For now, we only underlined the properties such distance must obey, in the next paragraph, we will study the first parametrisation.

2.2.1.b Geodesic distance construction.

GEODESIC DISTANCE ON G: Let G be a group of transformations and let's assume it is a smooth group everywhere. G also acts on itself through the composition. Let's consider $\mathbf{g} : [0,1] \to G$ a smooth path for which the time derivative is defined $\frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t}$ for all $t \in [0,1]$ by setting:

$$\left.\frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t}\right|_{t=t_0} = \frac{\partial \mathbf{g}}{\partial t}(t_0). \tag{2.19}$$

Each element $g \in G$ is in correspondence with the tangent space $T_g G$ which contains all the vectors $\left.\frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t}\right|_{t=t_0}$ for all paths going though g at time t_0, formally we can write:

$$\left\{ \left.\frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t}\right|_{t=t_0} : \mathbf{g} : [0,1] \to G; \mathbf{g}(t_0) = g \right\} \tag{2.20}$$

A translation or a left action on $\mathbf{g}$ by $h \in G$ being given by the mapping

$$h \triangleright \mathbf{g} : \Big(\big(t \mapsto g(x)\big) \mapsto \big(t \mapsto g \circ h(x)\big) \Big) \tag{2.21}$$

and illustrated in Figure 2.5. The translated left action derivative is thus the invertible linear mapping: $\forall t \in [0,1]$

$$\begin{aligned} \frac{\mathrm{d}h \triangleright \mathbf{g}}{\mathrm{d}t} : \Omega &\to T_{g(t)\circ h}G \\ x &\mapsto \frac{\partial \mathbf{g}(t)}{\partial t}(h(x)) \end{aligned} \tag{2.22}$$

The path being transported, the derivatives are now part of other tangent spaces of G: we have for all $t \in [0,1]$, $\partial_t \mathbf{g}(t) \in T_{\mathbf{g}(t)}G$ and $\partial_t(h \triangleright \mathbf{g}(t)) \in T_{\mathbf{g}(t)\circ h}G$.

Each point g of G being associated with its tangent space, itself being a vector space, there exists a corresponding norm that we call $\| \bullet \|_g$. However, we don't necessarily know them. To overcome this difficulty, we can define them all just by defining the norm for the identity deformation, and retrieve the others by translation: The collection of left-invariant norms $\| \bullet \|_g, g \in G$ can be fully determined by $\| \bullet \|_{\mathrm{Id}}$. If we define $\| \bullet \|_{\mathrm{Id}}$, for all differentiable paths $\mathbf{g}$ with $\mathbf{g}(t_0) = \mathrm{Id}$ such that

$$\left\| \frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t}(t_0) \right\|_{\mathrm{Id}} = N_G \left(\frac{\partial \mathbf{g}}{\partial t}(t_0, \bullet) \right) \tag{2.23}$$

where N_G is a functional norm on the space of functions $v : \Omega \mapsto \mathbb{R}^d$. By left translation invariance, we deduce the norms for any path $\mathbf{g}$ from N_G; if $\mathbf{g}(s) = g$,

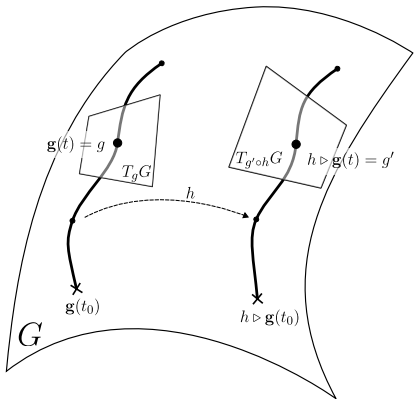

Figure 2.5: How the group G is also a Manifold. In particular, we can define paths $\mathbf{g} \in G$, the shortest being a geodesic.

$$\left\| \frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t}(t_0) \right\|_g = N_G \left(\frac{\partial \mathbf{g}}{\partial t}(t_0, \mathbf{g}^{-1}(t, \bullet)) \right). \tag{2.24}$$

Even if we will give a formal parametrisation of the temporal vector field in Definition 2.10, one can see arising the vector field essence here by setting

$$v(t,x) = \frac{\partial \mathbf{g}}{\partial t}(t, \mathbf{g}^{-1}(t,x)), \tag{2.25}$$

where v carries the information of an infinitesimal displacement among the deformation $\mathbf{g}$ at a time t (see Figure 2.4).

Knowing every norm on the path $\mathbf{g}$ let us compute the total energy of the deformation, using a Riemannian metric integrating over the norms such that

$$E(\mathbf{g}) = \int_0^1 \left\| \frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t} \right\|^2_{\mathbf{g}(t)} dt = \int_0^1 N_G \left(v(t, \bullet) \right)^2 dt \tag{2.26}$$

thus the associated geodesic distance

$$d_G(g, g') = \inf \left\{ \sqrt{E(\mathbf{g})} : \mathbf{g}(0) = g, \mathbf{g}(1) = g' \right\} \tag{2.27}$$

is left invariant for the action of G on itself.

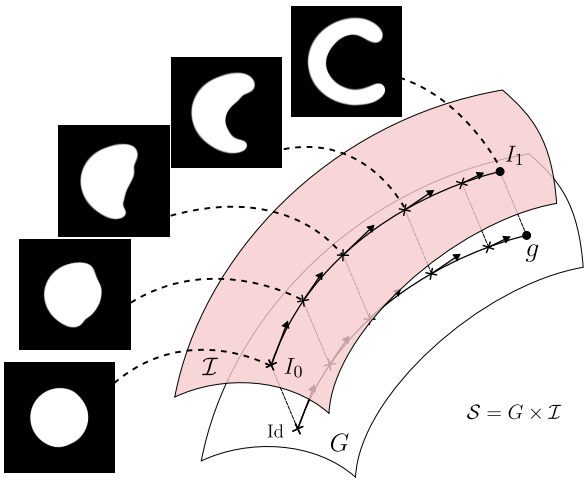

Figure 2.6: A path on $\mathcal{I}$ generated by the corresponding path on G

DISTANCE ON $\mathcal{S}$ In this paragraph, we will extend the previous result to the product space $\mathcal{S} = G \times \mathcal{I}$. The reasoning is the same, with the addition of the smooth space of images $\mathcal{I}$ which elements are functions $I : \Omega \to \mathbb{R}$ on which G is acting. Like for Equations 2.19 and 2.22 we define the time derivative of a path $\mathbf{s} : [0,1] \to \mathcal{S}$ and its translation $h \triangleright \mathbf{s}$, with $h \in G$:

$$\frac{\mathrm{d}\mathbf{s}}{dt} = \left(\frac{\partial \mathbf{g}}{\partial t}, \frac{\partial \mathbf{I}}{\partial t}\right); \qquad \frac{\mathrm{d}h \triangleright \mathbf{s}}{dt} = \left[(t,x) \mapsto \left(\frac{\partial \mathbf{g}}{\partial t}(t, h(x)), \frac{\partial \mathbf{I}}{\partial t}(t, h(x))\right)\right] \tag{2.28}$$

We define the collection of norms $\| \bullet \|_{g,I}$, for $(g, I) \in \mathcal{S}$ from the case $g = \mathrm{Id}$

$$\left\| \frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t}(t_0) \right\|^2_{\mathrm{Id},I} = N_G \left(\frac{\partial \mathbf{g}}{\partial t}(t_0, \bullet), I \right)^2 + N_{\mathcal{I}} \left(\frac{\partial \mathbf{I}}{\partial t}(t_0, \bullet), I \right)^2 \tag{2.29}$$

for the chosen functional norms N_G and $N_{\mathcal{I}}$, the former being defined for functions $v : \Omega \to \mathbb{R}^d$ and the later for function $I : \Omega \to \mathbb{R}$. Assuming that $\mathbf{I} : [0,1] \to \mathcal{I}$ is the image path such that $\mathbf{I}(t) = \mathbf{g}(t) \triangleright I$ for any $I \in \mathcal{I}$ and path $\mathbf{g} \in G$; we obtain the norm for any $g \in G$ applying the same pullback as in Equation 2.24.

$$\left\| \frac{\mathrm{d}\mathbf{g}}{\mathrm{d}t}(t_0) \right\|^2_{g,I} = N_G \left(\frac{\partial \mathbf{g}}{\partial t}(t_0, \mathbf{g}^{-1}(t, \bullet)), \mathbf{I}(\mathbf{g}^{-1}(t, \bullet)) \right)^2 + N_{\mathcal{I}} \left(\frac{\partial \mathbf{I}}{\partial t}(t_0, \mathbf{g}^{-1}(t, \bullet)), \mathbf{I}(\mathbf{g}^{-1}(t, \bullet)) \right)^2 \tag{2.30}$$

We set v as in Equation 2.25 and its *alter ego* for images; the infinitesimal image variation by the deformation $\mathbf{g}$ at time t: $\breve{I} = I(t, \mathbf{g}^{-1}(t, x))$. Assuming $\breve{I}$ is differentiable both in time and space, we have the identity:

$$\frac{\partial \mathbf{I}}{\partial t}(t,x) = \frac{\partial \breve{\mathbf{I}}}{\partial t}(t, \mathbf{g}(x)) + \frac{\partial \breve{\mathbf{I}}}{\partial x}(t, \mathbf{g}(x)) \cdot \frac{\partial \mathbf{g}}{\partial t}(t,x) \tag{2.31}$$

$$\Leftrightarrow \frac{\partial \mathbf{I}}{\partial t}(t, \mathbf{g}^{-1}(t,x)) = \frac{\partial \breve{\mathbf{I}}}{\partial t}(t,x) + \frac{\partial \breve{\mathbf{I}}}{\partial x}(t,x) \cdot v(t,x) \tag{2.32}$$

I will comment on it after the conclusion: The energy of the path on $\mathcal{S}$ is given by the following theorem, illustrated in Figure 2.6. The previous discussion holds for proof.

Theorem 2.1. *(Miller and Younes [2001]) Let $(N_G(\bullet, I), I \in \mathcal{I})$ and $(N(\bullet, I), I \in \mathcal{I})$ be two collections of norms. Associate to the temporal vector field $t \mapsto \mathbf{v}(t, \bullet)$ and the path $t \mapsto \breve{\mathbf{I}}(t, \bullet)$, where $\mathbf{v}(t, x) \in \mathbb{R}^d$ and $\breve{\mathbf{I}}(t, x) \in \mathbb{R}$, the energy*

$$E(\mathbf{v}, \breve{\mathbf{I}}) = \int_0^1 N_G\big(\mathbf{v}(t, \bullet); \breve{\mathbf{I}}(t, \bullet)\big)^2 dt + \int_0^1 N_{\mathcal{I}} \left(\frac{\partial \breve{\mathbf{I}}}{\partial t}(t, \bullet) + \frac{\partial \breve{\mathbf{I}}}{\partial x}(t, \bullet) \cdot v(t, \bullet); \breve{\mathbf{I}}(t, \bullet) \right)^2 dt. \tag{2.33}$$

Then, the function

$$d_{\mathcal{I}}(I, I') = \inf \left\{ \sqrt{E(v, \breve{\mathbf{I}})} : v, \breve{\mathbf{I}}, \breve{\mathbf{I}}(0, \bullet) = I, \breve{\mathbf{I}}(1, \bullet) = I \right\} \tag{2.34}$$

is a distance on $\mathcal{I}$.

The Equation 2.32 is the Lie derivative of $\breve{\mathbf{I}}$ in the direction of the vector field v. This description is obviously very general, the results being applicable to any kind of manifold. A famous toy example is the space of a landmarks triplet that actually happens to live on a sphere surface [Kendall et al., 2009; Klingenberg, 2020]. In this thesis application, we consider dense images defined on a 3D Euclidean subspace. One could argue that we do not need such elaborated definitions. However, we consider Euclidean space as the identity deformation that we will transport, each voxel carrying its intensity along the flow. Therefore this Euclidean space must be considered like a 3D manifold, and G is another of higher dimension, where each point is a deformed version of it (Figure 2.4). Also if $\breve{\mathbf{I}}$ is part of $\mathcal{I}$, it is not on the path $\mathbf{I}$ as the former is more the infinitesimal variation happening in between two close points of the last. This is why, we handled the partial derivative of $\mathbf{I}$ in time with caution, as the notion will come back often through the manuscript.

One can see that theorem 2.1 is a generalisation of LDDMM as given in Trouvé [1995]; Dupuis et al. [1998b] where given I and $I' \in \mathcal{I}$, one minimise:

$$\int_0^1 \|v(t, \bullet)\|^2 dt + \int_\Omega \big(I(x) - I' \circ \mathbf{g}^{-1}(1, x)\big)^2 dx \tag{2.35}$$

While commenting on their results Miller and Younes [2001] referred to this kind of matching as viscous matching in opposition to elastic one, more common at the beginning of the last millennium. They were right, as we will see, the geodesic formulation we will end up with is alike to the one of a viscous non-diffusive fluid (*c.f.:* Theorem 2.9). As we will see later, this theorem paves the road for Metamorphosis as we detail in Section 2.4.

We have seen that thanks to the group structure we pass from a hard functional problem of deforming images to a geometrical problem, giving us a sense of distance in between functions I through geodesics. Moreover, in a very general setting where there is no implied norms for every point of a manifold (*e.g.,*$\mathcal{I}$) we can deduce all norms on the path from G. The following sections will focus on the parametrisation of the space G and methods to register the following geodesics. Trouvé [1995] set regularity constraints to rigorously define a group G that shares many properties of finite-dimensional Lie groups. We will present a slightly weaker definition of G in the next section.

2.2.2 Diffeomorphisms from group theory point of view

The diffeomorphism group Diff, composed of the infinite-dimensional one-to-one smooth transformation, can be considered a good candidate for G. As we will see in this subsection,

we can put easily a group structure on Diff. We will give the general definition of diffeomorphism and we will wait for Sections 2.2.3 and 2.2.4.b to give the complete characterisation used for this thesis application.

Definition 2.9 (Diff(M, N))**.** *Given two manifolds M and N, a differentiable map $f : M \to N$ is called a* diffeomorphism *if it is a bijection and its inverse $f^{-1} : N \to M$ is differentiable as well.*

In particular, we will note Diff(M) *the group of diffeomorphisms from M to itself.*

The study of links between geometrical objects and an algebraical structure is quite common. For example, a compact topological space U is characterized by the real-valued continuous function algebra $\mathscr{C}(U)$. Moreover, if U and U' are compact spaces, every morphism $\mathscr{C}(U) \to \mathscr{C}(U')$ is a map such that $f \mapsto f \circ w$, where $w : U' \to U$ is continuous. In other words, we can retrieve U as the set of the morphisms $\mathscr{C}(U) \to \mathbb{R}$. These results can be extended to the group of diffeomorphisms, with the cost of more technical proofs. In particular, and it is important for us, a smooth manifold is characterized by its group of diffeomorphisms [Filipkiewicz, 1982; Mourtada, 2015]. Furthermore, if M is a compact smooth manifold, it is possible to define a distance $d : \text{Diff}(M) \times \text{Diff}(M) \to [0, \infty)$ such that the topology induced by d is $\mathscr{C}^\infty$. In fact, Diff(M) has a natural smooth structure which makes it an infinite-dimensional Lie group [Mourtada, 2015]. The Lie algebra of Diff(M) is then the vector fields algebra on M and its exponential application is the flow: $\exp(tv) = \Phi_v^t$. In addition, if we combine the results presented in Section 2.2.1, along with the fact that if we set a diffeomorphism group on any Manifold, the obit of a point is the manifold itself. We have the property, in particular: $\text{Diff} \triangleright \text{Id} = \text{Diff}$. In other words, we know that every diffeomorphism from the group is attainable from the identity deformation. [Milnor, 1965]

In this thesis, applications will focus on images I, from a manifold $\mathcal{I}$ where its elements are functions such that $I : \Omega \subset \mathbb{R}^d \to \mathbb{R}$. We will also call this manifold the *Morphological space*. Typically we will choose images as functions from $L^2(\Omega, \mathbb{R})$ [Beg et al., 2005; Younes, 2019; Vialard et al., 2011] or even with bounded variation (BV) [Vialard and Santambrogio, 2009]. Therefore we are not within the framework of Lie theory anymore. However, close enough to make analogies. Here, The algebraic model is that any image $I \in \mathcal{I}$ is carried by its coordinate system represented by a diffeomorphic transformation $\varphi \in \text{Diff}(\mathcal{I})$. φ interact with an image according to (φ, I), denoted algebraically as $(\varphi, I) \mapsto \varphi \triangleright I$. Thus, an elementary element exists, under the name of the identity transformation Id and elements of Diff are stable, as dictated by Equations 2.3. In particular, the second equation would be: $\varphi' \triangleright (\varphi \triangleright I) = (\varphi' \circ \varphi) \triangleright I$. We take as an assumption that the morphological space orbit of I is homogenous under the actions of Diff elements, leading that for all $I, J \in \mathcal{I}$ there exists $\varphi \in \text{Diff}(\mathcal{I})$ such that $J = \varphi \triangleright I$. In other words, we chose $\mathcal{I}$ to be composed only of elements attainable by a diffeomorphism. For images, it means that the topology is preserved, in terms of structure (see Figure 2.7). Indeed elements of $\varphi \in$ Diff when applied to objects generate a symmetries group.

Because we will work in setting when image coordinates Ω are subsets of the Euclidean space $\mathbb{R}^d$, the group left action under the flow translate as the inverse function composition:

$$(\varphi, I) \mapsto \varphi \triangleright I = I \circ \varphi^{-1}. \tag{2.36}$$

As both $\mathcal{I}$ and Diff are smooth manifolds, any pair of points at their surface can be joined by a path, subset of an integral curve[§]. Thus, we can deduce a Riemannian distance on Diff

[§]One could define flows on Diff with the 'sliding conditions' of Equation 2.8 holding.

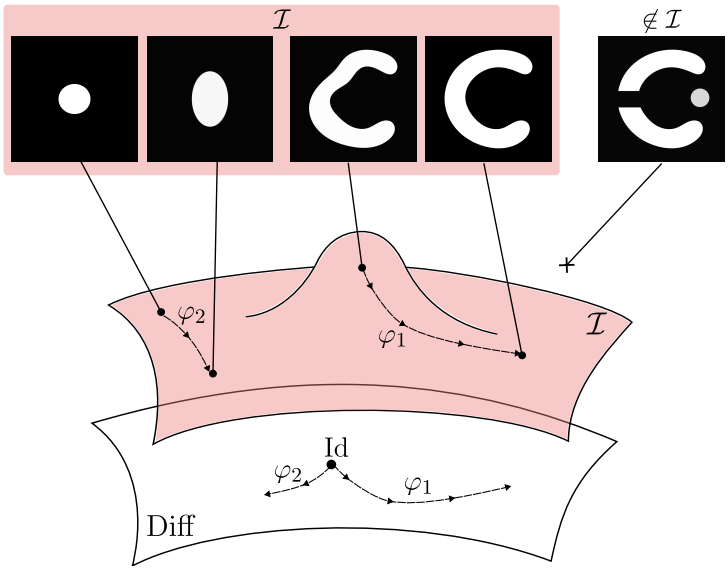

Figure 2.7: Schematic view of the shape space. We can move in the functional space of images with elements of Diff. The orbit $\mathrm{Diff} \triangleright I$ generate $\mathcal{I}$, the subspace of images with the same topology.

noted d_{Diff} for any diffeomorphisms $\varphi, \varphi' \in \mathrm{Diff}$ that also implies a distance on $\mathcal{I}$ noted $d_{\mathcal{I}}$. (*c.f.:* Sections 2.3.1 and 2.2.1). From the definition of Morphological space, we know that for any pair of images $I, J \in \mathcal{I}$, it exists $\varphi \in \mathrm{Diff}$ such that $\varphi \triangleright I = J$. Furthermore, because $\mathcal{I}$ is smooth, if a pair of images is close on $\mathcal{I}$ for $d_{\mathcal{I}}$, φ will be close to the identity as well for the associated distance on Diff. Formally, if I and J are such that $d_{\mathcal{I}}(I, J) < \varepsilon_1$ for any $\varepsilon_1 \in \mathbb{R}^{+*}$ then it exist a diffeomorphism φ matching I on J and a small constant $\varepsilon_2 \in \mathbb{R}^{+*}$ such that $d_{\mathrm{Diff}}(\mathrm{Id}, \varphi) < \varepsilon_2$.

We saw in this section a first sketch of the Shape space, and that Diff is a good candidate for modelling transformation as it can be seen algebraically as a group and results from Section 2.2.1.a holds. However one can argue that we still have not rigorously parametrised Diff. Your patience will be rewarded in the next section.

2.2.3 Temporal vector fields to parametrise large diffeomorphic deformations

2.2.3.a A Hilbert space for smooth vector fields

Let S and T be two images from the image space $\mathcal{I}$ defined on the same domain Ω. We want to register the source image S on the target T. We have already seen vector fields arising from Equation 2.25, here we will give a more formal definition.

Definition 2.10 (Displacement Vector field)**.** *We call a displacement vector field any map* $v = (v^1(x), \cdots, v^d(x)) : \Omega \mapsto \mathbb{R}^d$ *belonging to* $\mathscr{C}^1_0(\Omega, \mathbb{R}^d)$*, which are the* C^1 *fields on* Ω *vanishing at infinity along with their derivatives.*

As a remark, and as we have seen before in Section 2.1.1.b, we can view v as a *derivation*

operator on $\mathscr{C}^1_0$ defined for all $f \in \mathscr{C}^1_0(\Omega)$ as:

$$v \diamond f(x) = \sum_{i=1}^{d} v^i(x) \frac{\partial f}{\partial x_i}(x). \tag{2.37}$$

So v is identified to the derivation operator : $v = \sum_{i=1}^{n} v^i(x) \partial_{x_i}(x)$.

Note: We see here that this definition coincides with its generalised form on Manifolds seen in Section 2.1. For this thesis applications, we will only work with vector fields defined on Euclidean vectors spaces; but understanding the more general form on manifolds will help visualise core concepts of LDDMM and Metamorphosis.

This definition underlines the fact that the derivation operator $I \mapsto v \diamond I$ performs an infinitesimal action of the vector field on an image.

To compute a diffeomorphism from a single vector field is not an obvious task, and the approximation $\Phi(x) = x + v(x)$ holds only if the vector field $v : \Omega \mapsto \mathbb{R}^d$ is small enough. One can see a non-working example in Figure 2.3. Thus integrating a temporal vector field to a diffeomorphic deformation must be done with caution. We are interested in a subclass of diffeomorphisms obtained form the ordinary differential equation for all $t \in [0, 1]$

$$\dot{\varphi}_t = v_t \circ \varphi_t; \quad \varphi_0 = \mathrm{Id} \tag{2.38}$$

We will denote by $\mathbf{v}$ a *time dependent vector field* which is an ordered collection of vector fields $(v_t)_{t\in[0,1]}$. A discrete approximation with N elements can be defined as the N-uplet $\mathbf{v} = (\mathtt{v}_{\mathtt{t}_0=0}, \mathtt{v}_{\mathtt{t}_1}, \cdots, \mathtt{v}_{\mathtt{t}_{N-1}=1})$. During the dissertation, we will use symbols in `Typewriter` fonts for discretised versions of the continuous variables we are working with.

We will denote the deformation given by integrating over $\mathbf{v}$ by $\varphi^{\mathbf{v}}$ (see Section 2.2.3.b for more details). We want to make sure that the transformation $\varphi^{\mathbf{v}}$, such that $T = \varphi^{\mathbf{v}} \circ S$ or equivalently $S \circ (\varphi^{\mathbf{v}})^{-1} = T$, is a diffeomorphism and $\varphi^{\mathbf{v}}$ transports the point $x \in \Omega$ to a realistic position. (see Section 2.5.1.a for more detail about image transport implementation).

Definition 2.11 (admissible vector space V)**.** *The vector space V of vector fields $v : \Omega \to \mathbb{R}^d$ is said admissible if:*

- *V is a Hilbert space and its norm and scalar product are denoted by $\|\cdot\|_V$ $\langle \cdot, \cdot \rangle_V$.*
- *$(V, \|\cdot\|^2_V)$ continuously injects itself in $(\mathscr{C}^1_0(\Omega, \mathbb{R}^d), \|\cdot\|_{1,\infty})$ the C^1 fields on Ω vanishing at infinity along with their derivatives :*

$$V \hookrightarrow \mathscr{C}^1_0(\Omega, \mathbb{R}^d), \tag{2.39}$$

  meaning that its exists $k \in \mathbb{R}^+$ such that for all $v \in V$, $\|v\|_{1,\infty} \doteq \|v\|_\infty + \|dv\|_\infty \leq k\|v\|_V$.

The study of V is important as the elements of V need to be sufficiently smooth to produce a flow of diffeomorphisms [Dupuis et al., 1998a]. This leads to a first characterisation of the set of admissible transformations:

$$\mathrm{Diff}_V \doteq \left\{ \varphi^{\mathbf{v}} : \int_0^1 \|v_t\|^2_V \, dt < \infty \right\} \tag{2.40}$$

We are now equipped to translate the results of Section 2.2.1 for $G = \mathrm{Diff}$.

Theorem 2.2. *Let V be a space of admissible vector fields and Diff_V the related group of diffeomorphisms. For all $\phi \in \mathrm{Diff}_V$ and $\mathbf{v} \in L^2([0,1], V)$,*

$$d_{\mathrm{Diff}}(\mathrm{Id}, \phi) \doteq \inf \left\{ \left(\int_0^1 \|v_t\|_V^2 \, dt \right)^{\frac{1}{2}} : \mathbf{v} \in L^2([0,1], V), \quad \varphi_1^{\mathbf{v}} = \phi \right\} \tag{2.41}$$

is a left-invariant metric on Diff_V and $(\mathrm{Diff}_V, d_{\mathrm{Diff}})$ is a complete metric space.

This theorem by A.Trouvé is proven in this form in Theorem 7.17 of Younes [2019] book. Authors often refer to $E(v) = \int_0^1 \|v_t\|_V^2 \, dt$ as the *energy* of a deformation, as it gives a measure of the total displacement on the ambient space. It corresponds to the length of a path on G given by a Riemannian distance of Diff_V and can be interpreted as a *geodesic distance* on the group.

By coupling theorems 2.1 and 2.2 we are able to compare images in the space $\mathcal{I}$ using the deformation given by d_{Diff}:

$$d_{\mathcal{I}}(S,T) = \inf \left\{ \left(\int_0^1 \|v_t\|_V^2 \, dt \right)^{\frac{1}{2}} : \mathbf{v} \in L^2([0,1], V), \; S \circ \varphi^{\mathbf{v}} = T, \right\} \tag{2.42}$$

We are now almost ready to measure the distance between objects using a diffeomorphism flow.

2.2.3.b Hiatus: Temporal vector field integration in practice.

Before giving more theoretical results on the scalar product deduced from the V norm, I would like to give some details on the implementation of diffeomorphisms.

From the definition given by Equation 2.40 we can deduce a formula to get the deformation $\varphi_t^{\mathbf{v}}$ at any given time t. If we use the ODE notation from Equation 2.38 the flow at any given time is written

$$\varphi_t^{\mathbf{v}} = \mathrm{Id} + \int_0^t v_s \circ \varphi_s^{\mathbf{v}} ds, \quad v_s \in \mathbf{v}. \tag{2.43}$$

where $\varphi_t(x)$ is the transported position of x at time t. The map $t \to \varphi_t^{\mathbf{v}}$ is unique for a given $\mathbf{v}$ and differentiable for any time and voxel [Glaunès, 2005]. It is also possible to start the integration midway, with $0 \leq t_1 < t_2 \leq 1$:

$$\varphi_{t_1,t_2}^{\mathbf{v}} = \mathrm{Id} + \int_{t_1}^{t_2} v_s \circ \varphi_s^{\mathbf{v}} ds, \quad v_s \in \mathbf{v}. \tag{2.44}$$

It is immediate that for $s < r < t, \varphi_{s,r}^{\mathbf{v}} \circ \varphi_{r,t}^{\mathbf{v}} = \varphi_{s,t}^{\mathbf{v}}$.

If we integrate the vector fields backwards, through the relation:

$$\begin{cases} \dot{\phi}_t & = (-v_{1-t}) \circ \phi_t \\ \phi_1 & = \mathrm{Id} \end{cases} \qquad \forall t \in [0,1], v_t \in \mathbf{v} \tag{2.45}$$

We can rewrite the system such that

$$\varphi_{t_2,t_1}^{\mathbf{v}} = \mathrm{Id} + \int_{t_1}^{t_2} -v_{t_2-s} \circ \varphi_{t_2,t_2-s}^{\mathbf{v}} ds, \quad v_s \in \mathbf{v} \tag{2.46}$$

and in particular

$$\varphi^{\mathbf{v}}_{1,0} = \mathrm{Id} + \int_0^1 -v_{1-s} \circ \varphi^{\mathbf{v}}_{1,t_2-s} ds \tag{2.47}$$

One can verify that $\varphi^{\mathbf{v}}_t \circ (\varphi^{\mathbf{v}}_t)^{-1} = \mathrm{Id}$ for all $t \in [0,1]$ and in particular that $\varphi^{\mathbf{v}} \circ (\varphi^{\mathbf{v}})^{-1} = \mathrm{Id}$.

In practice, for two vector fields v and w the operation $v \circ (\mathrm{Id} + w)$ is made by interpolation. One can read the implementation principle to obtain both $\varphi^{\mathbf{v}}$ and $(\varphi^{\mathbf{v}})^{-1}$ in Code 2.2.

Source Code 2.2: Temporal field integrator implementing the construction of $\varphi^{\mathbf{v}}$ and $(\varphi^{\mathbf{v}})^{-1}$. You can find a full tutorial within the file `vector_field_to_flow_man.ipynb`

```
def temporal_field_integrator_forward(vectField):
        grid_def = id_grid + vectField

        # change the vector field to image convention shape [B,C,...]
        in_vectField_im = tb.grid2im(vectField)

        for t in range(1,in_vectField.shape[0]):
                tmp_grid = grid_def.detach()
                # interpolate the field on the deformation
                interp_vectField = grid_sample(in_vectField_im[t,...]
                                        ,tmp_grid)

                # change back to vector field to fields shape conventions
                grid_def += tb.im2grid(interp_vectField)

        return grid_def

v  = # a temporal vector field of shape [T,...,D]
# compute the deformation phi_v
phi_v = temporal_field_integrator_forward(v)

# compute the inverse deformation of phi_v
v_inverse  = - v.flip(0) # put the field in reverse order from 1 to 0.
inv_phi_v = temporal_field_integrator_forward(v_inverse)
```

2.2.3.c Energy and momenta: Eulerian and Lagrangian formulation

We have just defined a Hilbert space and used its associated norm. We will provide more abstract details regarding its scalar product modelling. However, we will wait until Section 2.2.4.b to parameterise it using the theory of Reproducing Kernel Hilbert Spaces (RKHS). For now, we will provide some general insights. Let v and w be vector fields, elements of the Hilbert space V. We define the linear form $w \mapsto \langle v, w \rangle_V$ as $\mathbb{L}v$. Actually, this identity is simply a translation of what has just been defined: we have called $\mathbb{L}v$ the linear form $w \mapsto \langle v, w \rangle_V$, which is mathematically expressed as $\mathbb{L}v(w) \doteq \langle v, w \rangle_V$, or equivalently, using the notation of dual brackets, $(\mathbb{L}v|w) \doteq \langle v, w \rangle_V$. From this definition, we know that $\mathbb{L} : V \to V^*$ maps V to its dual and can be a generalised function. We will note its inverse $\mathbb{K} : V^* \to V$. In the latter, we will be in the case where the generalised functions can

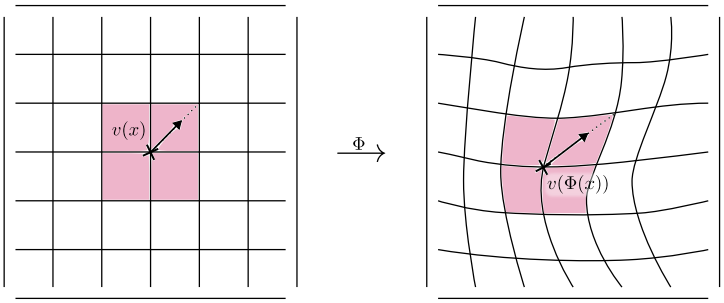

Figure 2.8: Adjoint action representation

be identified as real one: For example one can define $(\mathbb{L}v, v)$ with a function $\Psi : \Omega \mapsto \mathbb{R}^d$ integrable such that: $(\mathbb{L}v, v) = \int_\Omega \Psi(x) \cdot v(x)dx$.

As we will see in Subsection 2.2.3.d, one will need to choose the operator $\mathbb{L}$ for practical utilisation. Historically, $\mathbb{L}$ was set as a differential operator [Trouve et al., 1995; Beg et al., 2005]. Nowadays we choose the operator which will define the space V.

We saw earlier that the diffeomorphism $\varphi^{\mathbf{v}} \in \mathrm{Diff}_V$, with $\mathbf{v} \in L^2([0,1], V)$ model the displacement of particles $x \in \Omega$ over time. From a physical point of view, the vector $v_t(x)$ gives their velocity, it is called the *Eulerian velocity*. It is common to define the kinetic energy of the system at a given time by $E(v_t) = \frac{1}{2}\|v_t\|_V^2$, the total energy spent during the deformation is then given by integrating $E(v_t)$ over time [Miller et al., 2006; Beg et al., 2005; Younes, 2019] such that

$$E(\mathbf{v}) = \int_0^1 \|v_t\|_V^2 \, dt; \quad v_t \in \mathbf{v}. \tag{2.48}$$

Note that fluid dynamics mechanics describes kinetic energy as the sum of particle energies in a similar way. Again, in analogy with a standard mechanical system, the linear form $\mathbb{L}\mathbf{v}$ is interpreted as the *global momentum* of the system. To summarise, if v_t is the Eulerian velocity field at time t, the momentum at time t is given by $\mathbb{L}v_t$ and will be called the momentum in *Eulerian* coordinates. It is used to describe mechanical quantities in the current configuration at each time.

The complementary point view is called *Lagrangian*: It is when one describes quantities from the initial configuration following individual pixels' position through time. From Equation 2.38 one can see that a diffeomorphism is made by transporting its elements in the velocity direction given by v_t. On the other way around, given a deformation φ we can retrieve vector fields through the relation:

$$\tilde{v}_t(x) = \frac{\mathrm{d}}{\mathrm{d}s}\left(\varphi_t^{-1}(\varphi_{t+s}(x))\right)\Big|_{s=0} = (d\varphi_t)^{-1}(v_t \circ \varphi_t)(x). \tag{2.49}$$

As for Equation 2.25, it is just the translation of the fact that v_t acts infinitesimally on φ. This led to the adjoint definition:

Definition 2.12 (Adjoint action)**.** *Let φ be a diffeomorphism of Ω and $v \in V$. We denote $\mathrm{Ad}_\varphi v$ the vector field on Ω defined by*

$$\mathrm{Ad}_\varphi v : v \mapsto (\mathrm{d}\varphi)v \circ \varphi^{-1} \tag{2.50}$$

Remark: *In a finite case and with sufficient smoothness criteria to get Lie groups, the adjoint defines a fundamental Lie group operation and is called the adjoint action of G on its Lie algebra.*

From Definition 2.12 we have also the relation $v_t = \mathrm{Ad}_\Phi v\ \tilde{v}_t$. The new vector field under the adjoint action has to be interpreted as the transformation of v under the deformation generated by φ as illustrated in Figure 2.8. Note that the direction of v has moved relative to the background but stayed still relative to the moving grid. In fact, one can see it as a change of coordinates, where one wants to use the new coordinate system given by the flow. Or from the particle viewpoint, one replace x by $x' = \varphi(x)$, and the new particle moves with velocity:

$$\partial_t x' = \mathrm{d}\varphi(x)\partial_t x = \mathrm{d}\varphi(x)v_t(x) = (\mathrm{d}\varphi v_t)\circ\varphi^{-1}(x') \tag{2.51}$$

We retrieve the definition of adjoint. We will continue this discussion in Section 2.3.2 giving alternative interpretations from an algebraic viewpoint in the case of infinite Lie groups.

However, in practice, we do not necessarily know the expression of the dual adjoint. Thus We have to be careful in choosing $\mathbb{L}$ to be able to implement the update of the momentum. From this discussion, we see that it is not obvious to choose a good scalar product V and we will see in the next sub-section a very convenient tool to do so.

2.2.3.d Construction of the admissible vector space from the operator

In this Section, we will choose an operator L and construct an admissible vector space. This is a theoretic discussion and will only be useful for Section 2.3.2 where we discuss about momentum conservation within a special setting. Thus if one is mainly interested in the practical parts, one can skip this paragraph and read the next one.

We will define it through an operator $L : \mathscr{C}^\infty_c(\Omega, \mathbb{R}^d) \to L^2(\Omega, \mathbb{R}^d)$, with $\mathscr{C}^\infty_c(\Omega, \mathbb{R}^d)$ being the space of C^∞ maps with compact support. Its inverse is a smoothing operator. In this Section we will denote $\mathscr{C}^p(\Omega, \mathbb{R}^d)$ by $\mathscr{C}^p$ and $L^2(\Omega, \mathbb{R}^d)$ by L^2, to ease notation. We now choose L to be symmetric. We also set L as strongly monotonic, which means that there exists a constant c such that for all $v \in \mathscr{C}^\infty_c, \langle v, Lv\rangle_{L^2} \geq c\langle v, v\rangle_{L^2}$. We thus have positive definiteness.

We define the norm and scalar product associated to L as

$$\|v\|^2_L = \langle v, v\rangle_L \doteq \langle Lv, v\rangle_{L^2} = \int_\Omega Lv \cdot v\ dx, \quad v \in \mathscr{C}^\infty_c. \tag{2.52}$$

Where $\langle \bullet, \bullet\rangle_{L^2}$ is the scalar product of the square-integrable functions Hilbert space $L^2(\Omega, \mathbb{R}^d)$.

It can be done by identifying $\mathbb{L}$ the duality operator and L with Friedrich's extension theorem [Zeidler, 2012]. The reader may refer to [Younes [2019], Section 8.2] for proof.

Theorem 2.3 (Friedrich's extension)**.** *The inner product*

$$\langle \bullet, \bullet\rangle_L = \langle L\bullet, \bullet\rangle_{L^2} : \mathscr{C}^\infty_c \times \mathscr{C}^\infty_c \to \mathbb{R} \tag{2.53}$$

can be extended to an inner product

$$\langle \bullet, \bullet\rangle_V : V \times V \to \mathbb{R} \tag{2.54}$$

where V is a dense subspace of L^2 with respect to its norm, and such that $\mathscr{C}^\infty_c$ is a dense subspace of V with respect to $\|\bullet\|^2_V$. The operator L can also be extended to the duality operator of V, $\mathbb{L} : V \to V^*$. The extensions have the properties:

1. $(V, \| \bullet \|_V)$ *is continuously embedded in* $(L^2, \| \bullet \|_{L^2})$
2. *for all* $v, w \in \mathscr{C}_c^\infty$, $\langle v, w \rangle_V = \langle v, w \rangle_L = \langle Lv, w \rangle_{L^2} = (\mathbb{L}v | w)$
3. V *is a Hilbert space with respect to* $\langle \bullet, \bullet \rangle_V$.

This proves the completeness of V with respect to its norm.

In addition, the *duality paradox* states that if V is a subset of a Hilbert space H, then $V \subset H = H^* \subset V^*$. $\mathbb{L}$ being an extension of L and $\mathscr{C}_c^\infty \subset V$, we can see L as an operator with values in the dual space $(L^2)^*$. Moreover, the Fiedrich extension a with a lot of interesting properties. For example, Theorem 8.13 in Younes [2019] states that the Friedrich operator is bijective and self-adjoint and its inverse $\mathbb{K} = \mathbb{L}^{-1}$ as well. We do not details these points and refer to Younes's book.

We are interested in the construction of V which is given by the next theorem which characterises a space V using an orthonormal sequence ϕ_n in $\mathscr{L}^2(\Omega, \mathbb{R}^d)$.

Theorem 2.4 (Younes [2019],Theorem 8.15)**.** *Assume that* Ω *is bounded, and* $L : \mathscr{C}_c^\infty \to L^2$ *is symmetric and satisfies*

$$\langle Lv, v \rangle_{L^2} \leq c \|v\|_\infty^2, \quad \forall v \in \mathscr{C}_c^\infty(\Omega, \mathbb{R}^d) \tag{2.55}$$

for some constant $c > 0$*. Then the space* V *associated to* L *via the Friedrich extension is continuously embedded in* $\mathscr{C}^0(\Omega, \mathbb{R}^d)$ *and there exists an orthonormal basis,* (ϕ_n) *in* $L^2(\Omega, \mathbb{R}^d)$ *and a decreasing sequence of positive numbers* (ρ_n)*, which tends to* 0 *such that*

$$V = \left\{ v \in L^2(\Omega, \mathbb{R}^d) : \sum_{n=1}^{\infty} \frac{1}{\rho_n} \langle v, \phi_n \rangle_{L^2}^2 < \infty \right\}. \tag{2.56}$$

Moreover,

$$\mathbb{L}v = \sum_{n=1}^{\infty} \frac{1}{\rho_n} \langle v, \phi_n \rangle_{L^2} \phi_n, \tag{2.57}$$

whenever $v \in V$ *satisfies*

$$\sum_{n=1}^{\infty} \left(\frac{\langle v, \phi_n \rangle_{L^2}}{\rho_n} \right)^2 < \infty \tag{2.58}$$

With this theorem, we defined V and then identify the properties of its operator L, then we defined an operator $\mathbb{L}$ and show that the vector space deduced can be expressed through an orthonormal basis. However, this result does not match yet with the Definition 2.11 as we do not have a $\mathscr{C}^1$ embedding. We can have stronger regularity with Laplacian operators with Sobolev injections, which we do not detail here.

2.2.4 Construction of the admissible vector space using RKHS theory

In the previous sections we have defined V as a Hilbert space and then covered properties of deformation deduced from such space, then we saw a construction going the other way around, by defining the admissible vector space from an operator in Paragraph 2.2.3.d. Now we will use the Reproducing Kernels Hilbert Spaces (RKHS) theory to build a practical admissible vector space from a reproducing kernel.

2.2.4.a RKHS General properties

This section's objective is to give the theoretical tools for building the vector space V introduced in Section 2.2.3. We will present here some results of(RKHS). In the scalar case they were introduced by Aronszajn [1950] for functional analysis. However, we need to study the theory extension to vector-valued functions and refer to Glaunès [2005]; Younes [2019]; Charon [2013].

Definition 2.13 (Reproducing Kernel Hilbert Space (RKHS))**.** *Let $(H, \|\cdot\|_H)$ be a Hilbert space of functions on a set $\mathcal{A}$ taking values in $\mathbb{R}^d$. H is a RKHS if the functional $\delta_x^\alpha : f \in H \mapsto f(x)\cdot\alpha, \forall\alpha\in\mathbb{R}^d, \forall x\in\mathcal{A}$ is a continuous linear form on H.*

In other words, H is an RKHS if and only if all functionals δ_x^α are in the dual H^*. From this, using the Riesz representation theorem, we know that for all $x\in\mathcal{A}$ and $\alpha\in\mathbb{R}^d$, there exists an element $\mathbb{K}\delta_x^\alpha \in H$ such that

$$\forall f\in H, f(x)\cdot\alpha = \delta_x^\alpha(f) = \langle \mathbb{K}\delta_x^\alpha, f\rangle_H \tag{2.59}$$

This result using a Hilbertian scalar product $\langle\bullet,\bullet\rangle_H$, underlines the symmetric positive structure of the reproducing kernel and leads to the following: We can write $\mathbb{K}\delta_x^\alpha = K(x,\bullet)\alpha$ where $\mathbb{K}\delta_x^\alpha\in H$. We will call such functions $K(x,\bullet)\alpha$ fundamental functions of the RKHS. With Equation 2.59 we can combine two of theses functions, for all $x,y\in\mathcal{A}$ and $\alpha,\beta\in\mathbb{R}$:

$$\langle K(x,\bullet)\alpha, K(y,\bullet)\beta\rangle_H = \beta\cdot K(x,y)\alpha = \beta\cdot K(y,x)\alpha. \tag{2.60}$$

from which we can deduce that $K(x,y) = K(y,x)^T$.

A central result of the RKHS theory holds in this theorem from Aronszajn [1950].

Theorem 2.5. *To any positive kernel K that maps $\mathcal{A}\times\mathcal{A}\mapsto\mathscr{M}_d(\mathbb{R})$, where $\mathscr{M}_d(\mathbb{R})$ is the space of squared matrices of size d, there corresponds a* unique *RKHS H of functions from $\mathcal{A}$ to $\mathbb{R}^d$ whose reproducing kernel is K.*

We will use this result, defining a reproducing kernel to automatically construct a Hilbert space as a consequence. One can note by re-arranging the definitions that for any linear form $\nu\in H^*, x\in A$ and $\alpha\in\mathbb{R}^d$, we have the following identifications:

$$\mathbb{K}\nu(x)\cdot\alpha = (\delta_x^\alpha|\mathbb{K}\nu) = \langle\mathbb{K}\delta_x^\alpha, \mathbb{K}\nu\rangle_H = (\nu|\mathbb{K}\delta_x^\alpha) = (\nu|K(x,\bullet)\alpha) \tag{2.61}$$

2.2.4.b From reproducing kernels to an admissible vector space.

In the previous sections, we defined V and then identified the properties of its operator L, then we defined an operator $\mathbb{L}$ and show that the vector space deduced can be expressed through an orthonormal basis. There exists an alternative way to define V by first defining a kernel K.

We now use the reproducing kernels to generate a Hilbert space and we are interested in the inherited properties. In this subsection, we will work with kernel defined on $\mathcal{A}=\Omega=\mathbb{R}^d$. Note that we could choose Ω as an open set of $\mathbb{R}^d$, only the discussion about translation would be more difficult.

With the following proposition, proved in the dissertation of Glaunès [2005], we will use the natural regularity vector space properties of RKHS for the construction of V.

Proposition 2.3. *Let K be a $2p$-continuously differentiable positive kernel. Let us also assume that all the derivatives of order $l \leq p$ are bounded and vanish at infinity. Then the corresponding reproducing kernel is continuously embedded into $\mathscr{C}_0^p(\mathbb{R}^d, \mathbb{R}^d)$.*

This property tells us that the Hilbert space H generated by K is a good candidate for V, since $H \hookrightarrow \mathscr{C}_0^p(\mathbb{R}^d, \mathbb{R}^d)$. In particular, if we set $p = 1$, we respect the settings given in Definition 2.11. Moreover, as it is covered and proved in Glaunès [2005] and Younes [2019], the reproducing kernels come with some very convenient properties. We need a metric independent of object position or rotation. Fortunately, scalar reproducing kernels that are function of the euclidean distance, such that $K(x, y) = h(\|x - y\|) \cdot \mathrm{Id}_{\mathbb{R}^d}$, $h : \mathbb{R} \mapsto \mathbb{R}$ are invariant by translation and isometry. Translation invariance for vector fields means that for a translation θ of $\mathbb{R}^d$, the map $v \mapsto v \circ \theta$ is isometric in V. The isometry invariance means that any global rotation of the vector field (as part of the action of the orthonormal group) keeps its V norm unchanged.

We could study the whole class of isometry invariant kernels. However, it would not be necessary as in this dissertation, we will study only simple kernels such as Gaussian kernels.

Definition 2.14 (Gaussian reproducing kernel)**.** *For all $x, y \in \mathbb{R}^d$, $\sigma \in \mathbb{R}^+$*

$$K_\sigma(x, y) = \exp\left(-\frac{|x - y|^2}{2\sigma^2}\right) \cdot \mathrm{Id}_{\mathbb{R}^d} \tag{2.62}$$

We will now study the convolution operator $K\star$ and its associated bilinear form, which is for all $v, w \in L^2(\Omega, \mathbb{R}^d)$

$$\langle K \star v, w \rangle_{L^2} = \int_\Omega \int_\Omega \Big(K(x, y)v(x)\Big) \cdot w(y) dy dx \tag{2.63}$$

With the previous discussions, it can be shown that the operator $(K\star)$ is self-adjoint, positive definite. We can see that $\mathbb{K}$ is an extension of $(K\star)$ with the help of Equation 2.61.

Finally, we have a complete parametrisation of the V-norm. In Section 2.2.3.d, where we have constructed V from an operator, one could directly compute the V norm $\|v\|_V^2 = \langle Lv, v \rangle_{L^2}$ if v is regular enough. In this section, we cannot directly work with the V norm as we define $\mathbb{K}$ only and we do not necessarily know the expression of its inverse. However, if we have a $v \in V$ and we know $\psi = \mathbb{L}v \in V^*$, then we can compute the V norm:

$$\|v\|_V^2 = \langle \psi, v \rangle_{L^2} = \langle \psi, K \star \psi \rangle_{L^2}. \tag{2.64}$$

Note that it is only the case if $\psi \in L^2 \subset V^*$. This kind of computation with a Gaussian convolution is more stable and it will be critical from the numerical experiments. If $K = K_\sigma$ is a Gaussian kernel, L could be computed with a deconvolution or within the Fourier space, however, it is rather impractical.

The form given by Equation 2.62 allows us to set a scale at which two points interact with each other. The smaller σ is, the more the matching will generate localised deformation. Reversely a 'big' σ will lead to overall smoothed deformations, possibly missing details. In this context, tuning well the hyper-parameter σ is a critical point (see Section 2.5.2.c for visualisation). However, one can avoid fine-tuning using multi-scale kernels as demonstrated in the next Section.

In this section, we have seen a powerful tool for the parametrisation of admissible vector spaces. In addition, the deduced norm on V is invariant by translation and rotation, while scalable through the parameter σ. You can find more details on the implementation of reproducing kernels in Section 2.5.2.a

2.2.4.c Multi-scale resolution

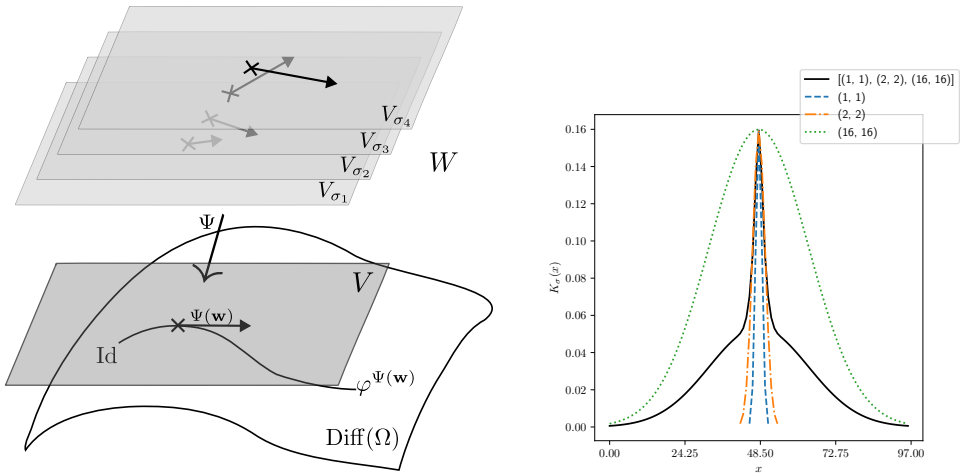

(a) Schematic view of LDDKBM

(b) Comparison of Gaussian and related multi-scale kernel.

Figure 2.9: Multi-scale kernel visualisation. (b) Each line is a slice of a 2D kernel. The dotted lines are Gaussian kernels and the solid one is multi-scale.

The Gaussian reproducing kernel as defined in Definition 2.14 has the limitation to be parametrised by σ. As it is often the case with parametric methods, the choice of an appropriate scale can be challenging and leads to compromises between regularity and registration quality. In particular, this phenomenon arises commonly with images having local constant intensity. Sommer et al. [2011] came with LDDKBM as a multi-scale kernel bundle that we will detail in this section.

Let us recall that for LDDMM the integration over the V norms induces a Riemannian metric (see Section 2.2.3 and 2.2.4.b. Sommer et al. [2011] extend the tangent space V to the family $W = \{V_\sigma\}_{\sigma\in\mathfrak{S}}$ of Hilbert spaces V_σ with $\mathfrak{S}$ being a set of scalars. Each $\sigma \in \mathfrak{S}$ is the scale parameter of V_σ. We now look at a registration among paths within the vector bundle $G_V \times W$, denoted the multi-scale kernel bundle. W inherits the vector space structure from each V_σ along with a pre-Hilbertian norm:

Definition 2.15 (W norm)**.** *Let $w = \{w^\sigma : \Omega \to \mathbb{R}^d, \sigma \in \mathfrak{S}\} \in W$ be a set of vector fields, such that each $w^\sigma \in V_\sigma$ is an admissible vector space. We define the W norm as the sum of the V_σ norm.*

$$\|w\|_W^2 = \sum_{\sigma\in\mathfrak{S}} \|w^\sigma\|_{V_\sigma}^2 \tag{2.65}$$

Note: If $\mathfrak{S}$ is a continuous set, one can write Equation 2.65 as an integral.

Similarly to the temporal vector fields, $\mathbf{w} = (w_t)_{t\in[0,1]}$ is a collection of elements of W. So the energy giving the length of the path $\phi^{\mathbf{w}}$ for the exact registration problem $I_1 = S\circ\varphi^{\mathbf{w}} = T$ is given by

$$E^W(\mathbf{w}) = \int_0^1 \|w_s\|_W^2 ds \tag{2.66}$$

As it is, the W space at a given location contains a different vector for each σ, it is necessary to define a map $\Psi : W \to V$ to be able to integrate the flow equation. We use the mean, defining $\Psi(w) = \frac{1}{\#\mathfrak{S}}\sum_{\sigma\in\mathfrak{S}} w_\sigma$ where $\#\mathfrak{S}$ is the cardinal of $\mathfrak{S}$. Then we obtain the deformation $\varphi^{\Psi(\mathbf{w})}$ using Equation 2.43.

We keep most properties of the RKHS using the multi-scale bundle. However, we will see in Section 2.3 some LDDMM geodesics conservation properties. In particular, a geodesic path on V keeps the norms $\|v_t\|_V$ constant in time. For LDDKBM, the momentum is conserved along optimal paths in $G_V \times W$ though $\|w_t\|_W$ is not constant. This occurs because the new energy is not directly related to a metric in the Riemannian sense. You can find more detail relative to the implementation in Section 2.5.2.b.

2.2.5 Illustration example & conclusion

In this section, we have introduced the LDDMM framework and we are ready to start its study. Before making a summary of the properties we encountered as a matter of conclusion, I would like to show you some LDDMM registration examples, implemented with our library Demeter. It uses the same principles as the Metamorphosis implementation. Therefore we will wait for Section 2.5 to detail it.

For illustration purposes, we reproduce the famous example of D'Arcy Thomson registering fishes. To perform the registration all images were converted to grayscale, and then we applied the deformation to each channel to obtain the coloured image. You can find a similar example in Joan Glaunès' course in the landmark case. One can note that in our case we do not need to provide or find key points to match. Each column corresponds to one registration, if one compares fish shapes, one can see that the matching is good while preserving small details like if the mouth is open or not. If in this figure, colour is for aesthetics, it helps also to see the body deformation.

In Section 1.2.2, we introduced the concept of diffeomorphic shape-space. The registering methods under the name LDDMM are the backbone of this concept. We have seen how the group structure of diffeomorphisms can match two objects, in particular images. Emphasising early on the property that these methods must register objects having the same topology (see Figure 2.7 for an example). Indeed, the theoretical object $\mathcal{I}_I$ defined in the section beginning can be built from a collection of representative images. We choose to parametrise diffeomorphisms using flows of temporal vector fields chosen from the Hilbert space V. This gave us, first, a definition of a Riemannian distance for images by integrating those fields and second, a vector space to work with, representing accurately images. As a Hilbert space distance is induced from its inner product, we started to define it abstractly with the help of the differential operator L such that $\langle v, w\rangle_V = \langle Lv, w\rangle_{L^2}$. With the help of the RKHS theory, we choose a kernel such that for images $K = L^{-1}$ is a Gaussian convolution of parameter σ, giving an image-relative sense of scale. We also studied its multi-scaled counterpart. To conclude, the LDDMM seems to have numerous advantages in theory. In the next section we will study a strategy to find optimal paths.

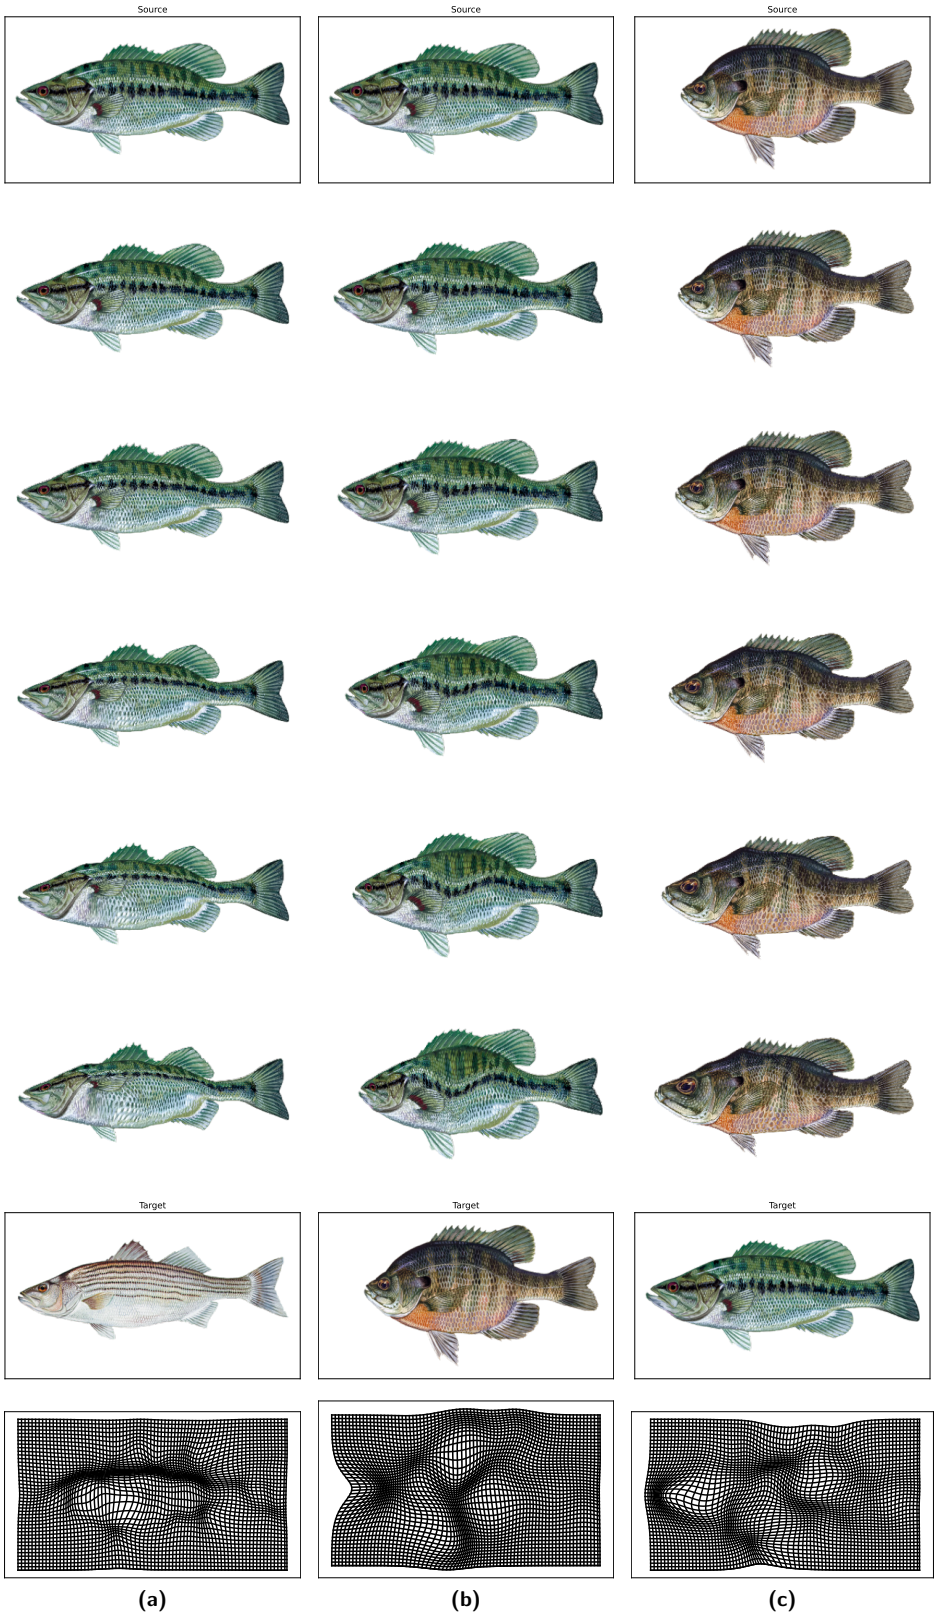

Figure 2.10: Each column shows an LDDMM registration example. Source images (top) were registered on targets (bottom) in grayscale, along with the found deformation grid.
Images from: https://www.wildlifedepartment.com/fishing/resources/species

2.3 From Motion equations to geodesic Shooting

In the previous section, we built the LDDMM framework. We know how a pair of images can be registered and we even have the ability to deduce a distance between these two. However, in practice, we did not talk about an optimisation strategy yet. We know that solutions exist, which is nice, but we lack the ability to find them. This problem will be addressed in this section. There are several ways to register objects using LDDMM. Historically, three main strategies were employed. At first, the problem was solved in the space of deformations but it is not only a high dimensional problem to solve, for some object's trajectory, there may exist no $v \in V$ to transport the object to the target, thus during numerical optimisation, resulting in diverging minimisation costs. Then, some were also deducing a gradient from the Hamiltonian form using the control formalism with the Pontrayagin principle, as it is reviewed in Miller et al. [2015]. Finally, one can use *shooting methods*, which we will detail in the following and which is one of the main focuses of this dissertation.

In numerical analysis, shooting methods aim to solve boundary value problems by reducing them to an initial value problem. For example, in ballistic, the full trajectory of a cannon ball can be deduced from the initial velocity, with the knowledge of some physical properties (like gravitation or air density). It is done by the study of the canon ball momentum over time. This section aims to find momentum conservation properties along a geodesic path in the shape-space.

Fortunately, the shape-space theory has deep connections with physics, especially fluid dynamics. In this chapter, I will go over some notions that are widely known by physicists and translate them in a way a Mathematician would have written about.

The section is organised as follows: We start by giving some elemental notions of Lagrangian mechanics, namely, the Euler-Lagrange equations and the Noether theorem. Then we make a discussion around the momentum conservation for infinite Lie groups. Finally, we introduce a geodesic shooting method (see Section 2.5 for implementation details). If one wants to skip a theoretical discussion on momentum conservation for infinite dimensional Lie groups, one can go directly to Subsection 2.3.3 where we set the geodesic shooting strategy for our space of interest.

2.3.1 Lagrangian mechanics Simple case

Physicists like referring to optimisations problems as energies to minimise, following the least action principles. From a physical viewpoint, the Lagrangian $\mathcal{L}$ is defined as the difference between kinetic energy and potential energy. The former is relative to the speed and the second is about the position. For example, in the study of a moving object, the Lagrangian allows focusing on the system energy rather than the force, the function summarising the dynamics of the entire system. Indeed, let m_0 and m_1 two points on a smooth manifold M and $\Gamma(m_0, m_1)$ the set of smooth paths $\gamma : [0, 1] \to M$ starting at $\gamma(0) = m_0$ and ending at $\gamma(1) = m_1$. They call E the action functional describing the amount of change in a physical system over time, defined such that

$$\begin{aligned} E : \Gamma(m_0, m_1) &\to \mathbb{R} \\ E(\gamma) &= \int_0^1 \mathcal{L}(t, \gamma(t), \dot{\gamma}(t))dt. \end{aligned} \tag{2.67}$$

Note that this notion, while having similar names is not to be mistaken with the group action. One can remark that if the Lagrangian is defined as a norm functional of $\dot{\gamma}$, E is a Riemannian distance. In this case, minimal paths on M are geodesics.

One can be interested in the characterisation of paths for which E is a stationary point, meaning that the derivative of E is null. The Euler-Lagrange equations give such results. To intuit their meaning we can think of the fact that for a given differentiable function, one can find its extrema by studying the vanishing point of its derivative. The Euler-Lagrange equation gives a condition for which a function is stationary according to Hamilton's principle. Before giving the theorem in a more general form, I will give it in the case where objects live in a finite dimensional vector space, as this simpler case is easier to intuit. A proof in a more general context has been proposed by Courant and Hilbert [1954].

Theorem 2.6 (Euler-Lagrange equations - finite dimensional case)**.** *Let $M \subset \mathbb{R}^d$ be a smooth manifold and E being defined from Equation 2.67 the Lagrangian $\mathcal{L} : [0,1]\times M\times\mathbb{R}^d \to \mathbb{R}$ a smooth function two times continuously differentiable. A path $\gamma \in \Gamma(m_0, m_1)$, noting γ^i the element of its i-th component, is a stationary point of the action functional E if and only if*

$$\frac{\partial \mathcal{L}}{\partial \gamma_i}(t, \gamma(t), \dot{\gamma}(t)) - \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial \mathcal{L}}{\partial \dot{\gamma}_i}(t, \gamma(t), \dot{\gamma}(t)) = 0, \quad \forall i \in \{1, \dots, d\}. \tag{2.68}$$

Proof. We aim to find the minimal path $\tilde{\gamma} \in \Gamma(m_0, m_1)$ for the metric given by E. We will call $\tilde{\gamma}_\varepsilon$ a perturbation of $\tilde{\gamma}$ by a differentiable function ν satisfying $\nu(0) = \nu(1) = 0$, such that for all time: $\tilde{\gamma}_\varepsilon(t) = \tilde{\gamma}(t) + \varepsilon\nu(t)$. We define

$$E_\varepsilon = \int_0^1 \mathcal{L}(t, \tilde{\gamma}_\varepsilon(t), \dot{\tilde{\gamma}}_\varepsilon(t))dt = \int_0^1 \mathcal{L}_\varepsilon dt. \tag{2.69}$$

The total derivative of E_ε with respect to ε is

$$\frac{\mathrm{d}E_\varepsilon}{\mathrm{d}\varepsilon} = \int_0^1 \frac{\mathrm{d}}{\mathrm{d}\varepsilon}\mathcal{L}_\varepsilon dt. \tag{2.70}$$

Now focusing on the total derivative of $\mathcal{L}_\varepsilon$,recalling that t does not depend from ε:

$$\begin{aligned}\frac{\mathrm{d}\mathcal{L}_\varepsilon}{\mathrm{d}\varepsilon} &= \frac{\mathrm{d}t}{\mathrm{d}\varepsilon}\frac{\partial \mathcal{L}_\varepsilon}{\partial t} + \frac{\mathrm{d}\gamma_\varepsilon}{\mathrm{d}\varepsilon}\frac{\partial \mathcal{L}_\varepsilon}{\partial \gamma_\varepsilon} + \frac{\mathrm{d}\dot{\gamma}_\varepsilon}{\mathrm{d}\varepsilon}\frac{\partial \mathcal{L}_\varepsilon}{\partial \dot{\gamma}_\varepsilon} \\ &= \nu(t)\frac{\partial \mathcal{L}_\varepsilon}{\partial \gamma_\varepsilon} + \dot{\nu}(t)\frac{\partial \mathcal{L}_\varepsilon}{\partial \dot{\gamma}_\varepsilon}.\end{aligned}$$

Then

$$\frac{\mathrm{d}E_\varepsilon}{\mathrm{d}\varepsilon} = \int_0^1 \left[\nu(t)\frac{\partial \mathcal{L}_\varepsilon}{\partial \gamma_\varepsilon} + \dot{\nu}(t)\frac{\partial \mathcal{L}_\varepsilon}{\partial \dot{\gamma}_\varepsilon}\right] dt. \tag{2.71}$$

When $\varepsilon = 0$ we know that E_ε must be minimal and

$$\left.\frac{\mathrm{d}E_\varepsilon}{\mathrm{d}\varepsilon}\right|_{\varepsilon=0} = \int_0^1 \left[\nu(t)\frac{\partial \mathcal{L}}{\partial \gamma} + \dot{\nu}(t)\frac{\partial \mathcal{L}}{\partial \dot{\gamma}}\right] dt = 0. \tag{2.72}$$

We use integration by part to lose the time derivative of ν

$$\int_0^1 \left[\nu(t)\frac{\partial \mathcal{L}}{\partial \gamma} + \dot{\nu}(t)\frac{\partial \mathcal{L}}{\partial \dot{\gamma}}\right] dt = \int_0^1 \left[\frac{\partial \mathcal{L}}{\partial \gamma} - \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial \mathcal{L}}{\partial \dot{\gamma}}\right]\nu(t)dt + \left[\nu(t)\frac{\partial \mathcal{L}}{\partial \dot{\gamma}}\right]_0^1 = 0 \tag{2.73}$$

Recalling that $\nu(0) = \nu(1) = 0$ the last term at the border is zero and we have.

$$\int_0^1 \left(\frac{\partial \mathcal{L}}{\partial \gamma} - \frac{\mathrm{d}}{\mathrm{d}t} \frac{\partial \mathcal{L}}{\partial \dot{\gamma}} \right) \nu(t) dt = 0 \tag{2.74}$$

By the fundamental lemma of calculus, we deduce the result :

$$\frac{\partial \mathcal{L}}{\partial \gamma} - \frac{\mathrm{d}}{\mathrm{d}t} \frac{\partial \mathcal{L}}{\partial \dot{\gamma}} = 0 \tag{2.68}$$

□

Remark 1: In Hamiltonian physics we call the *conjugate momentum* the quantity p_i:

$$p_i \dot{=} \frac{\partial \mathcal{L}}{\partial \dot{\gamma}_i} \tag{2.75}$$

with $\dot{\gamma}_i$ being the generalised velocity.

Noether's theorem [Noether, 1918] is handy to study Lagrangian invariances of a system by given coordinate transformations. It allows studying a physical problem regarding its symmetries, translations and evolution in time. The Noether theorem is quite general, we only give intuition here in a simple case. It states, in physical terms, that: "*If a system has a continuous symmetry property, then there are corresponding quantities whose values are conserved in time.*" Let us translate the above sentence into a more mathematical formulation.

The word *symmetry* here refer to a diffeomorphism α of M such that the Lagrangian is invariant (*i.e.:* $\mathcal{L} \circ d\alpha = \mathcal{L}$, with $d\alpha = (\alpha, \frac{\partial \alpha}{\partial \gamma})$. An *infinitesimal symmetry* of $\mathcal{L}$ however, is a vector field v such that the one parameter sub-group $t \mapsto \alpha_t^v$ is a symmetry of $\mathcal{L}$, in other words, it a one-dimensional Lie group of transformations. It is the expression of a law of a physical quantity conservation and is usually expressed as a continuity equation. Thus a group of transformations letting $\mathcal{L}$ invariant is invariant for its integral $E = \int \mathcal{L}$ as well.

Theorem 2.7 (Noether)**.** *Let $\mathcal{L} : [0,1] \times M \times \mathbb{R}^d \to \mathbb{R}$ be a Lagrangian and v a infinitesimal symmetry of $\mathcal{L}$. Then the function $f : [0,1] \times M \times \mathbb{R}^d$ defined such that*

$$f(t, \gamma, \dot{\gamma}) = \frac{\partial \mathcal{L}}{\partial \dot{\gamma}}(t, \gamma, \dot{\gamma}) v(\gamma) \tag{2.76}$$

is an integral of the Euler-Lagrange equations of $\mathcal{L}$.

Proof. (elements) The proof of this theorem formulation can be made by deriving with respect to s, at $s = 0$ the equation:

$$\mathcal{L}\left(t, \alpha_s^v(\gamma), \frac{\partial \alpha_s^v}{\partial \gamma}(\gamma, \dot{\gamma}) \right) = \mathcal{L}(t, \gamma, \dot{\gamma}) \tag{2.77}$$

and choosing a path γ that satisfy the Euler-Lagrange equations. □

Example: *For every infinitesimal transformation of a path $\dot{\gamma} = \alpha(\gamma)$ such that the action integral E is stationary for a given path γ, then the quantity*

$$I = \frac{\partial \mathcal{L}}{\partial \dot{\gamma}} \frac{\partial \gamma(s)}{\partial s} \tag{2.78}$$

is conserved.

To recapitulate, with the Theorem 2.6, by computing the Euler-Lagrange equations, we have the assurance that the resulting equations will parameterise a geodesic equation for a Riemannian Manifold. If any geodesic in between two points, is not always a minimal path, it is easier to look among the set of geodesics rather than in the wider set of paths. Finally, Noether's Theorem will give us the tools to check the momentum invariance, allowing us to fully set up a geodesic integration from the initial momenta.

In this paragraph discussion, we have chosen convenient assumptions for not having to care about localised coordinate changes within a manifold. However, many generalisations exist. We will now come back to our problem of interest by studying the paper of Miller et al. [2006], which gives a collection of such results for different cases.

2.3.2 Lagrangian for infinite dimensional Lie algebras

We will now study the Euler-Lagrange equation for momentum conservation in LDDMM, where the group of diffeomorphism is a Lie group. Miller et al. [2006] gave elements of proof for the Euler-Lagrange equations using the adjoint for infinite Lie groups. This 'result' is mostly educational and must be taken as a basis to build intuition. It is delicate to present the result here because there is no way to validate all the assumptions. We have seen in Section 2.1.3 some basic differences between infinite and finite-dimensional Lie groups. To recall the LDDMM construction from Section 2.2, we construct G from the Hilbert space V, itself defined from an operator $L : V \to V^*$.

If one wanted to make G a Lie group, such that a Lie algebra is defined on it. Indeed, the infinite group of diffeomorphisms G defined earlier is not a Lie group. On simpler finite Lie groups of deformations, like the group of Matrices, a Lie algebra is well defined. However, it is more complicated for infinite groups. A first guess would be to change slightly the diffeomorphism group of Definition 2.11 to make it a Lie group, making sure a Lie algebra is well defined. Let $G = \mathrm{Diff}^\infty(\Omega)$ be the group of $\mathscr{C}^\infty$ diffeomorphisms (*i.e.:* one to one maps $\varphi : \Omega \to \Omega$ such that φ and φ^{-1} are $\mathscr{C}^\infty$). This group admit a Lie group structure and the corresponding Lie algebra $\mathfrak{g} = T_{\mathrm{Id}}\mathrm{Diff}^\infty(\Omega) \approx \mathrm{Vec}^\infty(\Omega)$, where Vec^∞ is a space of smooth vector fields (*i.e.:* $\mathrm{Vec}^\infty \hookrightarrow \mathscr{C}^\infty(\Omega, \mathbb{R}^d)$) [Schmid, 2010]. Taking weaker smoothness assumptions, like making the vector field space $\mathscr{C}^k$ or H^s with finite k and s, will lose the derivative conservation of the Lie brackets. We can define a scalar product on the Lie algebra $\mathfrak{g}$, that we denote $\langle \bullet, \bullet \rangle_{\mathfrak{g}}$ and its norm $\| \bullet \|_{\mathfrak{g}}$. It is in a way related to the V norm, in fact, we define it with the help of an operator $L : \mathfrak{g} \to \mathfrak{g}^*$ such that $\langle v, w \rangle_{\mathfrak{g}} = (Lv|w)$ as we did before. However, $\mathfrak{g}$ is not a Hilbert space, losing some important properties of LDDMM. A second guess would be to choose $\mathfrak{g}$ as a Hilbert space and find a way to define a Lie algebra structure on it. A good candidate, a smooth and infinitely dimensional Hilbert space, could be the Sobolev group H^s, defining $L = (\mathrm{Id} - \Delta)^s$. However, there is no way to define a Lie algebra structure on it, despite the advancements made in this recent paper by Bruno et al. [2019]. Arguably, despite the problem not being well posed, the formal discussion helps to figure out the desired dynamics.

In this context, we have an alternative interpretation of the adjoint (see Definition 2.12) from an algebraic viewpoint. The adjoint coincides with the Lie bracket. Precisely the adjoint action Ad of G on its Lie algebra $\mathfrak{g}$ and the associated adjoint action ad of $\mathfrak{g}$ on itself

are given with their dual operator Ad^* and ad^*, by

$$\begin{aligned} \mathrm{Ad}_\varphi w &= (\mathrm{d}\varphi) w \circ \varphi^{-1} & \mathrm{ad}_v w &= [v, w] = (\mathrm{d}v)w - (\mathrm{d}w)v \\ (\mathrm{Ad}^*_\varphi f | w) &= (f | \mathrm{Ad}_\varphi w) & (\mathrm{ad}^*_v f | w) &= (f | \mathrm{ad}_v w) \end{aligned} \tag{2.79}$$

with $\varphi \in G, w \in \mathfrak{g}, f \in \mathfrak{g}^*$. In our case we have:

$$(Lv_t | v_t) = (Lv_t | \mathrm{Ad}_{\varphi_t^\mathbf{v}} \tilde{v}_t) = (\mathrm{Ad}^*_{\varphi_t^\mathbf{v}} Lv_t | \tilde{v}_t) \tag{2.80}$$

We retrieve the momentum in Lagrangian coordinates within the expression $\mathrm{Ad}^*_{\varphi_t^\mathbf{v}} Lv_t$

We are now ready to give the Euler-Lagrange equations for the Infinite Lie group. In the next equation, one can remark that there is a sign difference with Theorem 2.6 equation. According to the authors, it is due to the switch from a left-invariance to a right-invariance in the diffeomorphism case.

Proposition 2.4 (Euler-Lagrange equation - Infinite Lie Group, Miller et al. [2006])**.** *The Euler-Lagrange equation for the kinetic energy is given by*

$$\frac{\mathrm{d}Lv}{\mathrm{d}t} + \mathrm{ad}^*_v(Lv) = 0. \tag{2.81}$$

In the case where $Lv \in H$, it is a function and one has

$$\mathrm{ad}^*_v = \nabla \cdot ((Lv)v) + \mathrm{d}v^* Lv, \tag{2.82}$$

where $\nabla\cdot$ is the divergence.

Proof. The principle of the proof is very similar to the one of Theorem 2.6, only some technical parts vary.

Let $((t, \varepsilon) \mapsto \varphi_t^{\overset{\varepsilon}{\rightsquigarrow}})$ be a collection of path such that $\varphi_t^{\overset{\circ}{\rightsquigarrow}} = \varphi_t^\mathbf{v}$ is a geodesic and for any $\varepsilon > 0, \varphi_t^{\overset{\varepsilon}{\rightsquigarrow}}$ is a 'perturbed' path by ε. As elements of Diff^∞ are generated by vector fields, we say that their variation is given such that $\frac{\partial v_t^{\overset{\varepsilon}{\rightsquigarrow}}}{\partial \varepsilon} = h_t^{\overset{\varepsilon}{\rightsquigarrow}}$. Finally, we have

$$\frac{\partial \varphi_t^{\overset{\varepsilon}{\rightsquigarrow}}}{\partial t} = v_t^{\overset{\varepsilon}{\rightsquigarrow}} \circ \varphi_t^{\overset{\varepsilon}{\rightsquigarrow}} \text{ and } \frac{\partial \varphi_t^{\overset{\varepsilon}{\rightsquigarrow}}}{\partial \varepsilon} = \nu_t^{\overset{\varepsilon}{\rightsquigarrow}} \circ \varphi_t^{\overset{\varepsilon}{\rightsquigarrow}} \tag{2.83}$$

where $\nu_t^{\overset{\varepsilon}{\rightsquigarrow}}$ is the perturbation of the minimising diffeomorphism with $\nu_0^{\overset{\varepsilon}{\rightsquigarrow}} = \nu_1^{\overset{\varepsilon}{\rightsquigarrow}} = 0$. We will in this section work with higher smoothness assumptions than we will have in our application.

The first part of the proof is to get an expression of $h_t^{\overset{\varepsilon}{\rightsquigarrow}}$ isolated from the rest. To do so we can use the Clairaut identity (*i.e.:* symmetries of the second derivatives) by deriving Equation 2.83. Reminding that G is a Lie group with the group operation being the composition (*i.e.:* In Equation 2.3 $j(f, g) = f \circ g$), one can note that $v_t^{\overset{\varepsilon}{\rightsquigarrow}} \triangleright \varphi_t^{\overset{\varepsilon}{\rightsquigarrow}} = v_t^{\overset{\varepsilon}{\rightsquigarrow}} \circ \varphi_t^{\overset{\varepsilon}{\rightsquigarrow}}$, the derivation of this quantity must be done using the product rule:

$$\begin{aligned} \frac{\partial^2 \varphi}{\partial \varepsilon \partial t} &= \frac{\partial^2 \varphi}{\partial t \partial \varepsilon} \\ \Leftrightarrow h_t \circ \varphi_t + \mathrm{d}_{\varphi_t} v_t \nu_t \circ \varphi_t &= \frac{\partial \nu_t}{\partial t} \circ \varphi + \mathrm{d}_{\varphi_t} \nu_t v_t \circ \varphi_t \\ \Leftrightarrow h_t &= \frac{\partial \nu_t}{\partial t} + \mathrm{d}\nu_t v_t - \mathrm{d}v_t \nu_t \\ \Leftrightarrow h_t &= \frac{\partial \nu_t}{\partial t} + [\nu_t, v_t] = \frac{\partial \nu_t}{\partial t} - [v_t, \nu_t] \end{aligned} \tag{2.84}$$

This last equation makes sense: the variation of the vector field by ε is the difference between the infinitesimal variation of the path perturbation ν and the evaluation of v through the flow induced by ν.

Then we study the variation induced by ε on the paths through the integral of the norms:

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\varepsilon}\int_0^1 \|v_t^{\overset{\varepsilon}{\rightsquigarrow}}\|_{\mathfrak{g}}^2 dt &= 2\int_0^1 \langle v_t, h_t\rangle_{\mathfrak{g}} dt \\ &= 2\int_0^1 \left\langle v_t, \frac{\partial \nu_t}{\partial t} - [v_t, \nu_t]\right\rangle_{\mathfrak{g}} dt \\ &= 2\int_0^1 \left(Lv_t \middle| \frac{\partial \nu_t}{\partial t}\right) - 2\int_0^1 (Lv_t | \mathrm{ad}_v \nu_t)\end{aligned} \tag{2.85}$$

Since φ_t is extremal, this expression vanishes for all ν and by integration by part we get:

$$\frac{\mathrm{d}Lv}{\mathrm{d}t} + \mathrm{ad}_v^*(Lv) = 0 \tag{2.81}$$

□

In the same article Miller, Trouvé and Younes give conditions for the momenta conservation. I will give some results without the full justification and kindly advise the reader to refer to their article. You will find many well-justified arguments on different variable conservation properties. In particular, the authors focus on inexact image matching. They give the minimisation problem:

$$E_{\mathrm{LDDMM}}(\mathbf{v}) = \|S\circ(\varphi^{\mathbf{v}})^{-1} - T\|_{L^2}^2 + \lambda\int_0^1 \|v_t\|_V^2\,; \qquad \lambda\in\mathbb{R}, \tag{2.86}$$

and state that the optimal solution should satisfy at each time t,

$$Lv_t = -\frac{1}{\lambda}|\mathrm{d}\varphi_{0,t}^{\mathbf{v}}|(I_t - T\circ\varphi_{1,t}^{\mathbf{v}})\nabla I_t \tag{2.87}$$

with $I_t = S\circ(\varphi_{0,t}^{\mathbf{v}})^{-1}$ and $|\mathrm{d}\varphi|$ being the Jacobian of φ (a result also detailed in Beg et al. [2005]). In the case where $\varphi^{\mathbf{v}}$ perfectly match S on T, $I_t = S\circ(\varphi_{0,t}^{\mathbf{v}})^{-1} = T\circ\varphi_{1,t}^{\mathbf{v}}$ and $Lv_1 = 0$. We, therefore, have the momentum expression at each time and in particular at the origin:

$$Lv_0 = -\frac{1}{\lambda}|\mathrm{d}\varphi^{\mathbf{v}}|(S\circ(\varphi^{\mathbf{v}})^{-1} - T)\nabla S \tag{2.88}$$

These expressions can be also understood through the concept of *momentum maps* as explained by Bruveris et al. [2011]. The importance of the momentum map in geometric mechanics is due to Noethers theorem. Noethers theorem states that the generalised momentum L is a constant of motion for the system under consideration when its Hamiltonian is invariant under the action of G on $\mathfrak{g}$. This theorem enables one to turn symmetries of the Hamiltonian into conservation laws.

This paves the way to control the registration from the initial step only. They then give a dynamical system that controls a geodesic path on G. Setting $Lv_t = z_t\nabla I_t$, Miller et al. [2006] state that an image transported along a geodesic path is given by

$$\begin{cases} v_t &= L^{-1}(z_t\nabla I_t) \\ \dot{z}_t &+ (dz_t)v_t + \mathrm{div}(v_t)z_t = 0 \\ \dot{I}_t &+ v_t\cdot\nabla I_t = 0 \end{cases}. \tag{2.89}$$

This result can also be found in Beg et al. [2005] for LDDMM and it is a limit case from the proof in the geodesic transport equations of Metamorphosis introduced in Trouvé and Younes [2005]. I will give the proof of the Equation 2.89 result in the image-Metamorphic with $G = \mathrm{Diff}_V$ (and not Diff^∞) case. Surprisingly, the Metamorphosis case is easier to understand than the LDDMM one.

2.3.3 Geodesic shooting

In this section, we will work again with the assumptions we need for image registration. We will reuse the notions and notations introduced in Section 2.2. Generally speaking, we have seen that geodesic derived from the Euler-Lagrange equations principle of least action, which states that the true path taken by a physical system between two points is the one that minimizes the action, a function that describes the energy of the system. In the case of LDDMM, the action describes the difference between the two images being registered and the geodesic equations ensure that the deformation fields generated by the algorithm minimise this difference while conserving momentum. At this point, one can already control the whole deformation by integrating the geodesic system of PDE 2.89. Recall that if we did not prove the geodesic existence for $G = \mathrm{Diff}_V$ yet, we will see it in the next section as it is a particular case of the image-Metamorphic geodesic equations of Theorem 2.9.

We now built a shooting method to find the optimal deformation. Conceptually it is close to the *adjoint method* which is a numerical method, used among others, for fluid dynamics problems. This method, covered in the work of Céa [1986], uses the Lagrangian function to drastically reduces the systems dependence on the number of control parameters. Thus rather than finding the shortest path between images, we are looking among geodesic paths starting at S the one that ends closer to T. This is done via an optimisation process with data attachment and regularisation terms. The first term of Equation 2.86 ensures we get closer to T and we choose to use the L2 norm between images which is the SSD. The regularisation term, consisting in integrating the V norm, comes from the initial exact matching formulation.

$$E_{\mathrm{LDDMM}}(\mathbf{v}) = \|S \circ (\varphi^{\mathbf{v}})^{-1} - T\|_{L^2}^2 + \lambda \int_0^1 \|v_t\|_V^2 \, dt; \qquad \lambda \in \mathbb{R}, \tag{2.86}$$

As we already discussed, $S \circ (\varphi^{\mathbf{v}})^{-1}$ is integrated through PDE numerical schemes and can be deduced entirely from the initial conditions. Note that following the construction of Subsection 2.2.4.b, we do not know the expression of L. However, for a given geodesic, we can evaluate the value of the norm of its vector fields at every time through the relation given by Equation 2.64.

We can further simplify the cost 2.86 with the help of the next theorem.

Theorem 2.8. *(Glaunès [2005], Thm. 6) For any $\phi, \psi \in \mathrm{Diff}_V$ it exists $\mathbf{v} \in L^2([0,1], V)$ such that $\varphi^{\mathbf{v}} = \psi \circ \phi^{-1}$ and*

$$d_{\mathrm{Diff}_V}(\phi, \psi) = \sqrt{\int_0^1 \|v_t\|_V^2 \, dt}. \tag{2.90}$$

In addition, for such a temporal vector field $\mathbf{v}$, the norm $\|v_t\|_V$ is constant for almost every t. (i.e.: $\|v_0\|_V = \|v_t\|_V$, for almost every $t \in [0,1]$).

This theorem tells us that for any pair of diffeomorphisms in Diff_V we can find one temporal vector field $\mathbf{v}$ to compare them, knowing that the V norms of $\mathbf{v}$ will stay constant over time. Thus, with Theorem 2.8 we can remove the integration of the V-norms obtaining:

$$E_{\mathrm{LDDMM}}(\mathbf{v}) = \|S \circ (\varphi^{\mathbf{v}})^{-1} - T\|_{L^2}^2 + \lambda' \|v_0\|_V^2 \, dt; \qquad \lambda' \in \mathbb{R}. \tag{2.86.b}$$

The cost functions 2.86 and 2.86.b have thus equivalent solutions! In practice, for the implementation, we can set $\lambda' = \mathtt{n}\lambda$, where $\mathtt{n}$ is the number of integration steps. Note, however, that setting $\lambda = 0$, theoretically does not change the expected solution. This will be used in the following, indeed we saw in Section 2.2 that one can parameterise a diffeomorphism with elements from the Hilbert space V, with Theorem 2.8 a given diffeomorphism $\varphi^{\mathbf{v}}$ following the geodesic equations 2.89 is fully parameterised by the initial vector field $v_0 \in \mathbf{v}$. More precisely, by the initial momentum z_0, as v_0 is deduced from z_0 such that $v_0 = L^{-1}(z_0 \nabla S) = K \star (z_0 \nabla S)$.

Having built the LDDMM framework and opted for an optimisation strategy, we are now ready to introduce the Metamorphosis framework. We will cover the LDDMM implementation details in Section 2.5, as we utilise geodesic shooting for both methods in a combined framework.

2.4 Metamorphosis

The Metamorphic framework lies on the foundations of LDDMM we explored in Section 2.2 and 2.3.3. Metamorphosis's main application is to solve the attainability issue with diffeomorphisms in image matching. Indeed, pairs of images are, in general, not related by a diffeomorphism. Conceptually, for images, it supplements the transformation by allowing intensity changes to the diffeomorphic deformation. Thus changing the infinitesimal action. Theoretically, a lot of work has been made to ensure the compatibility of this extended framework with the LDDMM one [Trouvé and Younes, 2005; Trouvé and Younes, 2005; Holm et al., 2009].

In LDDMM, registrations were seen as paths on the group of deformations which was itself a Banach manifold, where deformations were used to define a Riemannian metric, making it a Riemannian manifold as well. Metamorphosis is a generalisation of this concept. Again let's follow its construction presented by Trouvé and Younes [2005] and take strong assumptions on the deformation smoothness. In this section introduction, we will work with M a Riemannian manifold containing visual objects (*e.g.,* images) and set $G = \mathrm{Diff}^{\infty}(M)$ and its Lie algebra $\mathfrak{g}$, properly introduced in Section 2.3.2. $\mathfrak{g}$ is a space of vector field and has the same role as V. Then, in Subsection 2.4.1, we will compute the Euler-Lagrange equation with more convenient smoothness assumptions and pose Metamorphosis as a shooting optimisation problem.

Let's assume that a classical norm is defined on M, noted $|\bullet|_M$, we aim to construct a new Riemannian metric on M through a new norm, noted $\|\bullet\|_M$. For example, if M is a space of images, $|\bullet|_M$ can be the L^2 norm.

Definition 2.16 (Metamorphosis)**.** *Let M be a Riemannian manifold. A* Metamorphosis *on M is defined as a pair of curves $\mathbf{m} = (m_t)_{t\in[0,1]} = ((g_t, z_t))_{t\in[0,1]}$ respectively on G and M, with $g_0 = \mathrm{Id}$*

With g_t being the deformation part (*e.g.,* $G = \mathrm{Diff}$) and δ_t the template evolution. When z_t is constant, the Metamorphosis is a pure deformation. Metamorphoses, by the evolution

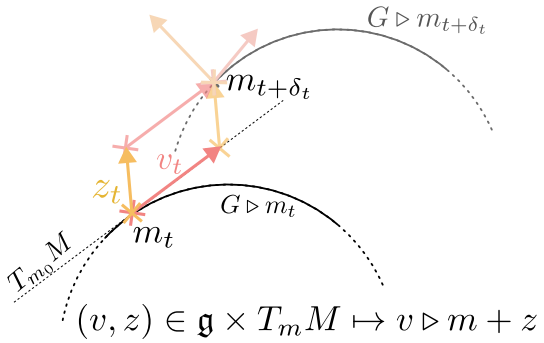

Figure 2.11: Metamorphic infinitesimal maps. The transformation is done via a succession of diffeomorphic deformations generated from $v \in \mathfrak{g}$ and of intensity changes z over a time step δ_t.

of their image, provide a convenient representation of combinations of a group action and of a variation on M. An infinitesimal deformation presents such a decomposition and is defined as a map:

$$\begin{aligned} \Phi_M : (\mathfrak{g}, T_m M) &\to T_m M \\ (v, z) &\mapsto v \diamond m + z. \end{aligned} \tag{2.91}$$

One can see an illustration in Figure 2.11. Similarly to LDDMM, a temporal vector field $\mathbf{v} = (v_t)_{t\in[0,1]}$ can be used to generate a diffeomorphic deformation $\varphi^{\mathbf{v}}$.

We define the new norm on M as the sum of the $\mathfrak{g}$ norm and the previous norm on M, with $\rho > 0$:

$$\|\eta\|_M^2 = \inf \left\{ \|v\|_{\mathfrak{g}} + \rho |z|_M^2 : \eta = \Phi_M(v, z) \right\} \tag{2.92}$$

$\mathfrak{g}$ is a Hilbert space and thus $\mathfrak{g} \times T_m M$ have a Hilbert structure. $\|\eta\|_M^2$ is the norm of the linear projection of $(0, \eta)$ on $\Phi_M^{-1}(0)$. Metamorphosis can be used to define a new Riemannian metric on M through the energy of a curve $\mathbf{m}$, such that

$$\begin{aligned} E(\mathbf{m}) &= \int_0^1 \left\| \frac{\mathrm{d}m_t}{\mathrm{d}t} \right\|_M^2 dt \\ &= \inf_{(t \mapsto v_t \in \mathfrak{g})} \left(\int_0^1 \|v_t\|_{\mathfrak{g}}^2 dt + \rho \int_0^1 \left| \frac{\mathrm{d}m_t}{\mathrm{d}t} - v_t \diamond m \right|_M^2 dt \right). \end{aligned} \tag{2.93}$$

The distance between two elements m and m' in M can therefore be computed by minimising $U(v_t, m_t) = E(m_t)$ over all curves $\mathbf{m}$ with boundary condition $m = m_0$ and $m' = m_1$ (exact matching).

Finally, by computing the Euler-Lagrange equation and with the Noether theorem, we can find conservation properties as described in Section 2.3. We will now explicit the geodesic equations for Metamorphosis with vector fields in V.

2.4.1 Theoretical background and proof

We now return to our case of interest for images and $G = \mathrm{Diff}_V$ in concordance with Sections 2.2.3 and 2.2.4. As we just saw, Metamorphoses join additive intensity changes with the deformations. The goal of Metamorphosis is to register an image I to J using variational

methods with an intensity additive term $z_t \in L^1([0,1],\mathbb{R})$. An infinitesimal step for image evolution can be defined as :

$$\partial_t I_t = v_t \diamond I_t + \mu z_t = -v_t \cdot \nabla I_t + \mu z_t, \quad \text{s.t. } I_0 = I \quad \mu \in \mathbb{R}^+. \tag{2.94}$$

One can control the amount of deformation vs photometric changes by varying the hyper-parameter $\mu \in \mathbb{R}^+$. As described by Trouvé & Younes in Trouvé and Younes [2005]; Younes [2019], the $\{z_t\}$ have to be the 'leftovers' of the transport of I_t by v_t toward the exact registration, as it can be seen by rewriting Equation 2.94 as $z_t = \frac{1}{\mu}(\partial_t I_t - v_t \cdot \nabla I_t)$. From this formulation, z is called the *residual*. However, it is also the momentum we have encountered in previous sections.

In order to find the optimal $(v_t)_{t\in[0,1]}$ and $(z_t)_{t\in[0,1]}$, one can minimise the exact matching functional Younes [2019]; Holm et al. [2009]; Richardson and Younes [2016] using Equation 2.94:

$$E_M(I,v) = \int_0^1 \|v_t\|_V^2 + \rho \|z_t\|_{L^2}^2 \, dt, \quad \text{s.t. } I_1 = J, I_0 = I; \quad \rho \in \mathbb{R}, \tag{2.95}$$

Theorem 2.9. *The geodesic equations deduced from E_M are :*

$$\begin{cases} v_t &= -\frac{\rho}{\mu} K \star (z_t \nabla I_t) \\ \partial_t z_t &= - \quad \nabla \cdot (z_t v_t) \\ \partial_t I_t &= -v_t \cdot \nabla I_t + \mu z_t \end{cases} \tag{2.96}$$

$\nabla \cdot (zv) = \mathrm{div}(zv)$ is the divergence of the field v times z at each pixel, K_σ is the chosen translation invariant RKHS kernel and $\star$ is the convolution. In practice, K is often a Gaussian blurring kernel Miller et al. [2006]; Vialard et al. [2011].

Proof. This proof can be found in a more general formulation in Younes [2019]; Holm et al. [2009].

Recalling from Equation 2.94 that $\|z_t\|_{L^2}^2 = \frac{1}{\mu^2}\left\|\dot{I} + v_t \cdot \nabla I_t\right\|_{L^2}^2$ and that the RKHS norm V is $\|v_t\|_V^2 =< (K\star)^{-1}(v_t), v_t >_{L^2}$ where $(K\star)^{-1}$ is an abstract differential operator. In order to compute the geodesics we will use the Euler-Lagrange equations.

Let begin by the variation regarding $v \in L^2([0,1],V)$ and consider the functional

$$F(t,v,\dot{v}) = E_M(I,\bullet)$$

We first compute the partial derivative of F with respect to the second argument v.

$$\begin{aligned} D_v F \cdot h &= 2\int_0^1 \langle v_t, h_t\rangle_V dt + 2\rho \int_0^1 \langle z_t, \frac{1}{\mu}\nabla I_t h_t\rangle_{L^2} dt \\ &= \int_0^1 \left\langle 2\big((K\star)^{-1} v_t + \frac{\rho}{\mu} z_t \nabla I_t\big), h_t \right\rangle_{L^2} dt \end{aligned}$$

by identification we have $\nabla_v F = 2((K\star)^{-1}v_t + \frac{\rho}{\mu} z \nabla I_t)$. Using a similar method we see that $\nabla_{\dot{v}} F = 0$. The Euler Lagrange equation being $\nabla_v F - \frac{d}{dt}\nabla_{\dot{v}} F = 0$ and k being linear it yields that

$$v = -\frac{\rho}{\mu} K \star (z \nabla I) \tag{2.96.a}$$

We now make the variation with respect to I. Let consider the equation g such that $G(t, I, \dot{I}) = E_M(\bullet, v)$. We begin by calculating the partial derivative relative to I.

$$\begin{aligned}
D_I G \cdot h &= 2\rho \int_0^1 \left\langle \frac{1}{\mu}(\dot{I}_t + \nabla I_t v_t), \frac{1}{\mu}\nabla h_t v_t \right\rangle_{L^2} dt \\
&= 2\rho \int_0^1 \left\langle z_t, \frac{1}{\mu}\nabla h_t v_t \right\rangle_{L^2} dt \\
&= \frac{2\rho}{\mu} \int_0^1 \int_\Omega z_t(x) \sum_{i=1}^d \frac{\partial h_t}{\partial x_i}(x) v_t^i(x) dx dt \\
\text{by part integration} \qquad &= -\frac{2\rho}{\mu} \int_0^1 \int_\Omega \sum_{i=1}^d h_t(x) \frac{\partial (z_t v_t^i)}{\partial x_i}(x) dx dt \\
&= \int_0^1 \left\langle -\frac{2\rho}{\mu} \nabla \cdot (z_t v_t), h \right\rangle_{L_2} dt.
\end{aligned}$$

and so $\nabla_I G = -\frac{2\rho}{\mu}\text{div}(z_t v_t)$. Similarly we can prove that $\nabla_{\dot{I}} G = \frac{2\rho}{\mu} z$. As the Euler-Lagrange equation is $\nabla_I G - \frac{d}{dt}\nabla_{\dot{I}} G = 0$ it yield our last equation :

$$\dot{z}_t = -\text{div}(z_t v_t) \tag{2.96.b}$$

By assembling equations 2.94, 2.96.a and 2.96.b we obtain the desired system. □

Note : The choice of μ and ρ varies in the literature. Richardson and Younes [2016]; Younes [2019] chose $\mu = 1$ and $\rho = 1/\sigma^2$, with σ being any real value (We use their notation here, it has no link with the Gaussian kernel). We tested many parameter configurations and found out that it could be interesting to be able to tweak them. We will detail it later.

The last line of Equations 2.96 is the advection term, simulating the movement of non-diffusive material. The second (continuity) equation is a conservative form which ensures that the amount of deformation is preserved on the whole domain over time. Thus, given the initial conditions of the system, $I = I_0$ and z_0, one can integrate in time the system of Equations 2.96 to obtain I_1. Note that v_0 can be computed from z_0, making z the only unknown. Furthermore, one can notice that the energy in Equation2.97 is conserved (*i.e.:* constant along the geodesic paths) and therefore the time integrals may be replaced by the norms at time 0.

Here, we propose to solve metamorphosis as an inexact matching problem. This allows us to have a unifying cost function (*i.e.:* Hamiltonian) for both LDDMM and metamorphosis:

$$H(z_0) = \frac{1}{2}\|I_1 - T\|_{L_2}^2 + \lambda\Big[\|v_0\|_V^2 + \rho\|z_0\|_{L_2}^2\Big] \tag{2.97}$$

with $\|v_0\|_V^2 = \langle z_0 \nabla S, K \star (z_0 \nabla S)\rangle$. The hyper-parameters λ and ρ define the amount of total regularisation and intensity changes respectively. Please notice that when $\mu = 0$, we retrieve the LDDMM cost function (Equation 2.86.b).

2.5 Implementation

The implementation of Metamorphosis is based on two main parts, imbricated one onto the other: First, one needs to resolve the minimisation problem given by Equation 2.97. Second, at every iteration, one needs to compute I_1 and thus define an accurate numerical scheme for Equations 2.96 integration. From now on, when we will write *iteration*, we will refer to an optimisation step and when we write *time step*, we will refer to a numerical scheme increment.

In this Section, we will mostly focus on the numerical scheme implementation. As a cautious both quality and efficiency of the integration scheme will be critical factors. The former is obviously important for the image's exploitability. The latter, as at each iteration the whole integration will need to be redone. Parts of this Section's work have been published and presented in François et al. [2021]. The code is fully available on GitHub at https://github.com/antonfrancois/Demeter_metamorphosis for anyone to use and contribute to.

In the historical implementation, a great deal was made on finding the deformation and energy gradients [Beg et al., 2005; Richardson and Younes, 2016]. For our implementation, we will totally skip this part by using the auto-differentiation of `PyTorch`. The fact that it will make the implementation easier will allow us to test variations of the original cost and deduced geodesic equations, as we will see in Chapter 3. Assuming we implemented a function `f(S)` that returns the integrated image, such that `f(S)` is differentiable in `PyTorch` standards (*i.e.:* There exist the corresponding back-propagation function). Therefore, `f(S)` must be implemented with `PyTorch` functions if we are not ready to provide a back-probation for the other ones. Good optimisation schemes are proposed by `Torch` and we use some of them directly.

The program core is organised into two classes or classes that inherit from them.

1.The abstract class `Geodesic_integrator` contains all the tool for integration, and its child `Metamorphosis_path` that make the Metamorphic or LDDMM geodesic integration. Most of the implementation details we cover in this Section are embedded within those classes.

2. The abstract class `Optimize_geodesicShooting` use a child of `Geodesic_integrator` to perform the Optimisation Scheme. The class built to work with `Metamorphosis_path` is called `Optimize_metamorphosis`.

We start the presentation by comparing Eulerian and Lagrangian schemes for the transport of an image by a vector field in Subsection 2.5.1. Then as a matter of compromise, we will present the semi-Lagrangian schemes that preserve the better of both worlds. Then, in Subsection 2.5.2 we cover the implementation of a Reproducing Kernel, introduced in Subsection 2.2.4. We finish by presenting an alternative numerical scheme ensuring images to stay sharp in Section 2.5.3. Each subsection deals with elements that required careful choices and planning. Rather than making a long general discussion at the end of the chapter, we wind up each subsection with results and a discussion related. As a conclusion to this Chapter, we will discuss if Metamorphosis is a suitable method for registering brains with glioblastomas.

2.5.1 Numerical scheme for the geodesic equation transport

In this subsection, we aim at solving the image evolution equation given by

$$\partial_t I_t + v_t \cdot \nabla I_t - \mu z_t = 0 \tag{2.96.c}$$

2.5.1.a Applying deformation to images

Most of the PDE numerical schemes aim to solve a relaxed version where only the transport takes place, by setting $\mu = 0$. For pedagogical purposes, we will discuss this case first and will deal with the case $\mu > 0$ from Subsection 2.5.1.b. We need to choose a numerical scheme with good numerical stability and not too computationally intensive. First one needs to make the classical choice between Eulerian and Lagrangian formulation that can be summarised by the discussion:

Let be $\alpha(x,t)$ a scalar function giving the density of a given fluid in time and space, and $u(x,t)$ a function giving the said fluid inputs in the domain. One can write the evolution of α in Lagrangian form:

$$\frac{\mathrm{d}\alpha}{\mathrm{d}t} = u, \tag{2.98}$$

or in Eulerian form as

$$\frac{\partial\alpha}{\partial t} + v\frac{\partial\alpha}{\partial x} = u, \tag{2.99}$$

with the velocity being defined such that $\frac{\mathrm{d}x}{\mathrm{d}t} = v$. This transport equation when $u = 0$ is called the *advection equation*. Obviously, this equivalence of formulation comes from the definition of the total derivative (*i.e.:* $\frac{\mathrm{d}}{\mathrm{d}t} = \frac{\partial}{\partial t} + \frac{\mathrm{d}x}{\mathrm{d}t}\frac{\partial}{\partial x}$). For a more intuitive explanation, the reader can refer to the comic of Section 2.6.

As one can see Equation 2.96.c is already in Eulerian form. Let's start by this study:

THE EULERIAN VIEWPOINT tracks the image intensity changes on a fixed grid. This formulation is without a doubt the easiest to understand and implement. Indeed recalling that with an infinitesimal displacement vector field v the infinitesimal transformation step is $-v \cdot \nabla I_t$ and with the definition 2.10 we have the identity:

$$v \triangleright I_t = \sum_{i=1}^{d} -v^i(x)\frac{\partial I_t}{\partial x_i}(x) = -v \cdot \nabla I_t \tag{2.100}$$

The implementation is straightforward and can be found in 2.3, however, one should be careful using it. Indeed the choice of a time step small enough is crucial. Arguably one could determine δ_t using the Courant-Friedrichs-Lewy (CFL) condition, which gives a necessary condition for PDE solvers convergence. It consist in controlling that a value C is bounded by another $C_{\max}$:

$$0 \leq C \doteq \delta_t \left(\max_{x\in\Omega} \sum_{i=1}^{d} \frac{|v^i(x)|}{\delta_{x^i}} \right) \leq C_{\max} \tag{2.101}$$

where δ_t is the time step, δ_{x^i} is the spatial step in between two voxel in the i-th direction, and $C_{\max}$ is a constant chosen by the user. The spatial step varies depending on the

convention we are in: in pixel conventions, $\delta_{x^i} = 1$ and in the `PyTorch` grid convention $\delta_{x^i} = \frac{2}{H-1}$, where H is the number of voxels in the i-th direction. In this form, the CFL condition does not depend on the smoothness of the image. One way to circonvoluate this is to set $C_{\max} = 1/\max_{x\in\Omega} \|\nabla I_t\|_\infty$. As the values of I are set to be in $[0,1]$ thus $\partial_{x_i} I_t(x) \leq 1, \forall t \in [0,1], \forall x \in \Omega$, we can set arbitrarily but reasonably $C_{\max} = 1$. Therefore we can deduce a suitable δ_t:

$$\begin{aligned} C = \delta_t \max_{x\in\Omega} \sum_{i=1}^{d} \frac{|v^i(x)|}{\delta_{x^i}} \leq \frac{\delta_t}{\sum_{i=1}^{d} \delta_{x^i}} \max_{x\in\Omega} d\|v^i(x)\|_\infty \leq 1 \\ \delta_t \leq \frac{\sum_{i=1}^{d} \delta_{x^i}}{\max_{x\in\Omega} d\|v^i(x)\|_\infty} \end{aligned} \tag{2.101}$$

where $\|v^i(x)\|_\infty = \max_{i\leq d} |v^i(x)|$. For example, if the image is 300×300 pixels wide and the displacement field v is of 100 pixels in one direction, then we must choose $\delta_t \leq \frac{1}{200}$. So to integrate I over a fixed vector field v, one must iterate Equation 2.100 more than 200 times (`n_step = 200` in code 2.3:l 20). In practice, however, we often need even more iterations.

Source Code 2.3: Eulerian image displacement

```
1  def _image_Eulerian_integrator_(self,image,vector_field,t_max,n_step):
2          """ image integrator using an Eulerian scheme
3
4          :param image: (tensor array) of shape [1,1,H,W] or [1,1,D,H,W]
5          :param vector_field: (tensor array) of shape [1,H,W,2] or [1,D,H,W,3]
6          :param t_max: (float) the integration will be made on [0,t_max]
7          :param n_step: (int) number of time steps in between [0,t_max]
8
9          :return: (tensor array) of shape [1,1,H,W] integrated with vector_field
10         """
11         dt = t_max/n_step
12         for t in torch.linspace(0,t_max,n_step):
13                 grad_I = tb.spacialGradient(image,dx_convention='pixel')
14                 grad_I_scalar_v = (grad_I[0]*tb.grid2im(vector_field)).sum(dim=1)
15                 image = image - grad_I_scalar_v * dt
16         return image
17
18 # Transport of an image respecting the amplitude of v
19 n_step = ... # must respect the CFL condition.
20 deformed_image = _image_Eulerian_integrator_(image,v,1,n_step)
21
22 # Transport of an image in the direction of v[t] part of a temporal vector field V.
23 n_step = len(V)
24 deformed_image =  _image_Eulerian_integrator_(image,v[t],1/n_step,1)
```

One could argue that using more complex Eulerian integration schemes would be better, like the Runge-Kutta methods. However, if they are more stable they would require to access values in between the grids, and therefore would need to be interpolated. It is possible but would be much slower than the Lagrangian methods discussed in the next paragraph.

In Metamorphosis $\mu > 0$ in Equation 2.96.c, the image update would be realised through the relation :

```
image =  _image_Eulerian_integrator_(image,v,1/n_step,1) + mu * residual
```

with `v` and `residual` recomputed at each iteration, making the process computationally intense (see also Result&Discussion paragraph of Subsection 2.5.1.b).

LAGRANGIAN VIEW POINT The implementation of an image displacement by a deformation ϕ is straightforward and stable as long as ϕ is diffeomorphic. It is a sufficient condition. For almost any vector fields[¶] we can always make it diffeomorphic using the Lie Exponential (see Section 2.1.2). For a temporal field $\mathbf{v}$, one can find $\varphi^{\mathbf{v}}$ by temporal vector fields integration, as detailed in Section 2.2.3.b.

Even if we often consider images as functions defined on an open set, their nature is discrete and lives on a grid of pixels. Note that we use the form $I \circ (\varphi^{\mathbf{v}})^{-1} = J$ in image registration rather than the direct form used for landmarks of surfaces form which would be $\varphi^{\mathbf{v}} \circ S = T$. Indeed this last form is convenient to send a point in space to another location, but in the image case, we need to catch the pixel's value at a given location. By applying $\varphi^{\mathbf{v}}$ to S we take the intensities of the pixels at their new location, and then we interpolate the image on a regular grid. Ashburner and Friston [2011] We use the interpolations provided by `PyTorch` in the function `grid_sample`. In 3D, only the tri-linear interpolation was implemented, thus for coherence, we use the bi-linear one in 2D.

In practice, it is realised by the line

```
deformed_image = F.grid_sample(image,inverse_deformation)
```

In `PyTorch` only the tri-linear interpolation is implemented for volumetric deformation.

With equation 2.43 we have for all $t \in [0, 1]$, $\varphi_t^{\mathbf{v}}$ a diffeomorphic deformation. The image deformation computation is simply done by evaluating the images at the position given by $\varphi^{\mathbf{v}}$, using cubic interpolation. A higher order of interpolation could have been preferred, however, there were not implemented for 3D images. Moreover, they would have been slower.

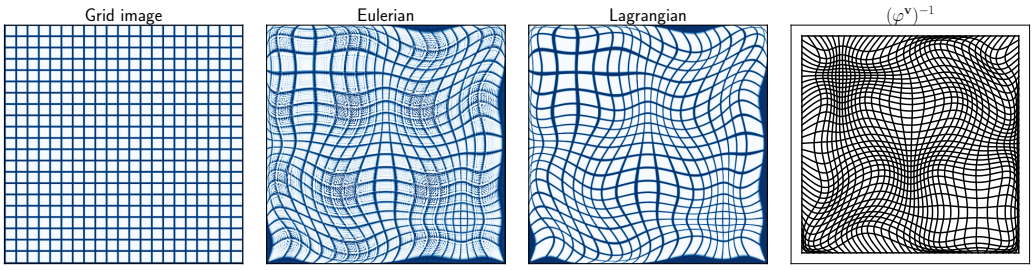

Figure 2.12: Image of a grid deformed using Lagrangian or Eulerian methods. Both images in the centre were displaced by a deformation $(\varphi^{\mathbf{v}})^{-1}$ with an Eulerian scheme (right) and a Lagrangian one (left). See the jupyter notebook: defomation_adventure.ipynb.

RESULTS AND DISCUSSION: Both points of view allow us to make similar deformation on images. The Lagrangian schemes are computationally stable and provide good tolerance to

[¶]We assume that we do not use ill cases constructed as a counter-example.

numerical-accuracy-related errors. Historically Lagrangian methods were avoided because of the numerous interpolations needed for particle tracking, which made the computation too costly [Beg et al., 2005]. Eulerian methods are conceptually easier to implement but as they depends on the image gradient they are very sensitive to image smoothness. The more there are sharp edges in the picture the more instabilities will generate abnormalities as seen in Figure 2.12.b (white stripes). To correct this issue, one would need to decrease the time step, leading to very large computational times. Today Lagrangian methods are way faster and more accurate and recent works started to drop Eulerian forms.

It is illustrated in Figure 2.12 where we deform a grid 700×700 image using both methods. The one shown through Eulerian integration was obtained for 500 iterations when the CFL conditions would require 225937 iterations and takes 3000x more time than the Lagrangian method. The Eulerian panel shows an example of the instabilities the method can create. You can check the Jupyter notebook, linked in the figure, for the full experience. To apply the Lagrangian scheme to the image, we needed to compute a deformation from the vector field. It was done using the fast-Exponential map (see Section 2.1.2) and required 6 interpolations (*i.e.:* $\exp(v) = \exp(2^{-6}v)^6$).

Integration of the geodesics is a crucial computational point for both LDDMM and Metamorphosis. In the case of image registration using LDDMM, Beg et al. [2005] initially described a method based on gradient descent which could not retrieve exact geodesics, as shown in Vialard et al. [2011]. An actual shooting method was then proposed in Vialard et al. [2011] for LDDMM-based registration of images. To the best of our knowledge, the only shooting method proposed in the literature for image metamorphosis is the one proposed in Richardson and Younes [2016]. It is based on a Lagrangian frame of reference and therefore it is not well suited for large images showing complicated deformations, as could be the case when registering healthy templates to patients with large tumours. Here, we propose to use a semi-Lagrangian scheme.

2.5.1.b From Eulerian to semi-Lagrangian formulation

We have seen two concurrent points of view for implementing numerical integration of ODE and PDE. We will now discuss a numerical scheme for the geodesic system of equation 2.96 integration. These paragraphs are the takeaways from my first paper François et al. [2021]. If in the previous paragraph, Lagrangian schemes were faster and more accurate, they work only for pure deformation. Thus these schemes seem not suitable for Metamorphosis as its image variations are modelled with transport and intensity changes at each time step. With this condition, Eulerian schemes seem to be the most natural candidate for flow integration. However, they are way too slow to be a viable choice. Hopefully, semi-Lagrangian schemes are a good compromise.

The main idea of semi-Lagrangian schemes is to stop evaluating the values of a function α (from Equations 2.98 and 2.99) on a regular fixed grid as we would do with Eulerian schemes. But rather deduce a new grid at each time step, with the help of the flow. How do we proceed? At each time step, for each voxel x, we integrate backwards in time by a δ_t to get the value of the voxel that will end up at position x.

$$\mathtt{I[t,x]} \approx \mathtt{I[t-1, x - v[t,x]]} \tag{2.102}$$

Obviously, $\mathtt{x - v[t,x]}$ will rarely end up on a regular grid and `I[t-1, x -v[t,x]]` must be obtained through interpolation. Unlike the pure Lagrangian procedure, this approach keeps

the fluid parcels evenly distributed throughout the flow and facilitates the computation of spatial derivatives via finite differences. [Durran, 2013]

In our case, we compute the deformation of a grid corresponding to a small displacement $\mathrm{Id} - \delta t\, v_t$, and then interpolate the values of the image I_t on the grid. This can be summarised by $I_{t+\delta t} \approx I_t \circ (id - \delta t\, v_t)$. Semi-Lagrangian schemes are more stable and don't need as many iterations. Too many iterations would blur the images due to the successive bilinear or trilinear (in 3D) interpolations.

We will now explain step by step the numerical scheme used for integration over the PDE system of Theorem 2.9.

$$\begin{cases} v_t &= -\frac{\rho}{\mu} K \star (z_t \nabla I_t) \\ \partial_t z_t &= - \nabla \cdot (z_t v_t) \\ \partial_t I_t &= -v_t \cdot \nabla I_t + \mu z_t \end{cases} \tag{2.96}$$

Let's reformulate Equation 2.96 in a semi-Lagrangian formulation, As one can see, computing the vector field can be done straight forward from the equation as no temporal derivative is involved. You can find all this implementation in the class `Metamorphosis_Path` within the file `metamorphosis.py`. In this tutorial, we integrate with `n_step` and it is set by the user.

Source Code 2.4: `_update_field_()` is a method that can be found in the class `Metamorphic_path` within the file `metamorphosis.py`

```
def _update_field_(image, residual):
    grad_image = tb.spacialGradient(image)
    field = tb.im2grid(kernelOperator( # see next section for the kernelOperator Implementation
        (-(residuals.unsqueeze(2) * grad_image).sum(dim=1))
                                      ))
    field *= get_field_cst_mult()
    return field
field = _update_field_(image,residual)
inverse_deformation = id_grid - field/n_step
```

In semi-Lagrangian formulation both $\dot{I}$ and $\dot{z}$ deformations are computed solving the characteristic equation:

$$dx(t) = v(t, x(t)), x(0) = \mathrm{Id} \tag{2.103}$$

We can now deal with the advection part. From the Eulerian formulation for the closed domain $\tau \times \Omega$, we can write:

$$\partial_t I_t + v_t \cdot \nabla I_t - \mu z_t = \partial_t I(t,x) + \sum_{i=1}^{d} \partial_{x_i} I(t,x) v_{x_i}(t, x_i) - \mu z(t,x) = 0. \tag{2.96.c}$$

where we use for convenience the notations $v_t = v(t,x) = (v_{x_1}(t,x), \cdots, v_{x_d}(t,x)), t \in \tau, x \in \Omega$ [Efremov et al., 2014]. The implementation is straightforward:

```
updated_image = F.grid_sample(image,inverse_deformation) + mu * residual[t]/n_step
```

Using semi-Lagrangian schemes for image evolution has been suggested by Beg et al. [2005] but not implemented. However, to our knowledge using it on the momentum is an

original idea. Thus, we can also rewrite the continuity equation as:

$$\begin{aligned}\partial_t z_t + \nabla \cdot (z_t v_t) &= \partial_t z_t + \sum_{i=1}^{d} \partial_{x_i}(z(t,x) \times v_{x_i}(t,x)) = 0 \\ &= \partial_t z_t + v_t \cdot \nabla z_t + (\nabla \cdot v_t) z_t = 0\end{aligned} \tag{2.96.b}$$

We can translate this equation to code such that:

Source Code 2.5: `_update_residuals_semiLagrangian_` is a method that can be found in the class `Metamorphic_path` within the file `metamorphosis.py`

```
def _update_residuals_semiLagrangian_(deformation):
    div_v_times_z = residuals * tb.Field_divergence()(field)[0,0]
    updated_residual = tb.imgDeform(residuals,
                                    inverse_deformation,
                                    clamp=False # important to keep negative values.
                                )
                    - div_v_times_z/n_step
    return updated_residual
```

Therefore in practice, Equations 2.96(b,c) transport are evaluated by interpolation. Note that by using a semi-Lagrangian scheme for residuals we avoid computing a discrete approximation of $\nabla \cdot (z_t v_t)$, but still need to compute an approximation of $\nabla \cdot v_t$. However, the momentum z_t, like the image I_t, is potentially non-smooth, while v_t is smooth due to its expression through the convolution operator K.

RESULTS AND DISCUSSION We will compare three presented computational options to integrate over the geodesics: 1- the Eulerian scheme, 2- the semi-Lagrangian approach and 3- a combination of the two, where we use the semi-Lagrangian scheme for the advection and the Eulerian scheme for the residuals (as it is suggested in Beg et al. [2005]).

In Figure 2.13, we can observe the lack of stability of Eulerian methods compared to the semi-Lagrangian ones. Even if the chosen time step is rather small, the Eulerian scheme produces ripples (in purple in the residuals) and the integration fails (see the estimated deformation). On the contrary, semi-Lagrangian schemes converge to a better deformation with a higher time step. It should also be noted that the full semi-Lagrangian scheme (advection and continuity equations) is perfectly stable without showing ripples, as it is instead the case for the advection-only semi-Lagrangian scheme.

It is clear that semi-Lagrangian schemes are a better solution than pure Eulerian ones. Note that we did not test Runge-Kutta methods, because of their huge computational complexity. However, a recent publication that I found while writing this dissertation proposed a semi-Lagrangian Runge-Kutta integration for LDDMM [Hernandez, 2021] using the Runge-Kutta algorithm described in Guo [2013]. Nor paper published their implementation, it would be interesting to re-implement their work for comparison. Apart from our integration methods, the paper of M. Hernandez has a major difference from ours: They implemented the gradient update with forward and backward passes following their Band-Limited implementation [Hernandez, 2018] when we use the auto-differentiation from PyTorch. In addition, the recent works of Brunn et al. [2021a,b] optimised LDDMM for GPU usage written in C++ and Cuda. According to their claims, they provided the fastest implementation to date. Another lead could come from recent advancements in the computation of

fluid-dynamics schemes, which went through huge improvement with the push for the media industries, being always faster and more realistic. For example, McNamara et al. [2004] work shows already impressive results. The main difficulty in these methods would be to compute the method's backward gradients.

Speed and computation efficiency was not our main focus for our implementation, and we did not want to implement LDDMM only. Indeed we utilised the similarities of both LDDMM and Metamorphosis geodesic equations to craft a unified integration. Moreover, we had in mind the possibility of adapting the framework (this is the topic of Chapter 3), and thus focus our effort on an object-oriented, easily customisable implementation. This is where the auto-differentiation of `PyTorch` came necessary. Indeed, for a given registration cost, one needs only to derive the Euler-Lagrange equations and the gradient is computed automatically through the chain rule. To sum up, since the publication of our implementation [François et al., 2021], faster codes for LDDMM were published, however, ours remain the only one available to our knowledge for Metamorphosis. In addition, it was written in Python with the idea to conduct experiments and easily bring modifications to the framework.

2.5.2 Kernel Implementation details

2.5.2.a Gaussian kernel

We use the definition of the V norm defined from its inverse operator $\mathbb{K}$ which is itself the extension of the convolution kernel $(K_\sigma \star)$ as seen in Subsection 2.2.4.b. As we have seen in this section, we have a trick to compute the V norm. It remains to compute $\mathbb{K}_\sigma$ as is it directly involved in Equation 2.96.a.

The Gaussian kernel of Definition 2.14 convoluted to a field $v \in V$ gives

$$K_\sigma \star v(x) = \int_\Omega K_\sigma(x,y)v(y)dy \tag{2.104}$$

The choice of the kernel size in pixels has a significant impact on the computation time. Because the kernel is Gaussian only neighbouring voxels are significant. Thus we can cut the convolution mask where the values tend to be zero. We choose to set its size to be $\lceil 6 \times \sigma \rceil$ and at least 7 pixel wide: (so the value of the Gaussian kernel is about $\exp(3/\sigma) \sim 10^{-4}$ at its edges). The convolution will be performed very often in the implementation, and increasing its efficiency has a significant impact on the total computation time. The direct convolution is the only convolution provided by `PyTorch`. As for big kernels, it is way faster to perform the convolution in the Fourier space, it was made possible from `PyTorch` 1.7 where a differentiable `fft` was induced. One can see benchmark results on CPU[‖] in Figure 2.14. A clear version of the implementation can be found in Code 2.6

Source Code 2.6: Sample of code for Gaussian kernel. Here is presented the essence of the computation done in the class `GaussianRKHS` within the file `reproducing_kernels.py`. Relying on the functions `get_Gaussian_kernel` and `convolution_filter` that must be implemented with care.

```
1 def Gaussian_convolution(sigma,input):
2     big_odd = lambda val : max(6,int(val*6)) + (1 - max(6,int(val*6)) %2)
```

[‖]At the time I implemented it, the gain in speed was also effective in GPU. However more efficient torch convolution has been released for GPU since. My implementation chooses automatically the best filter.

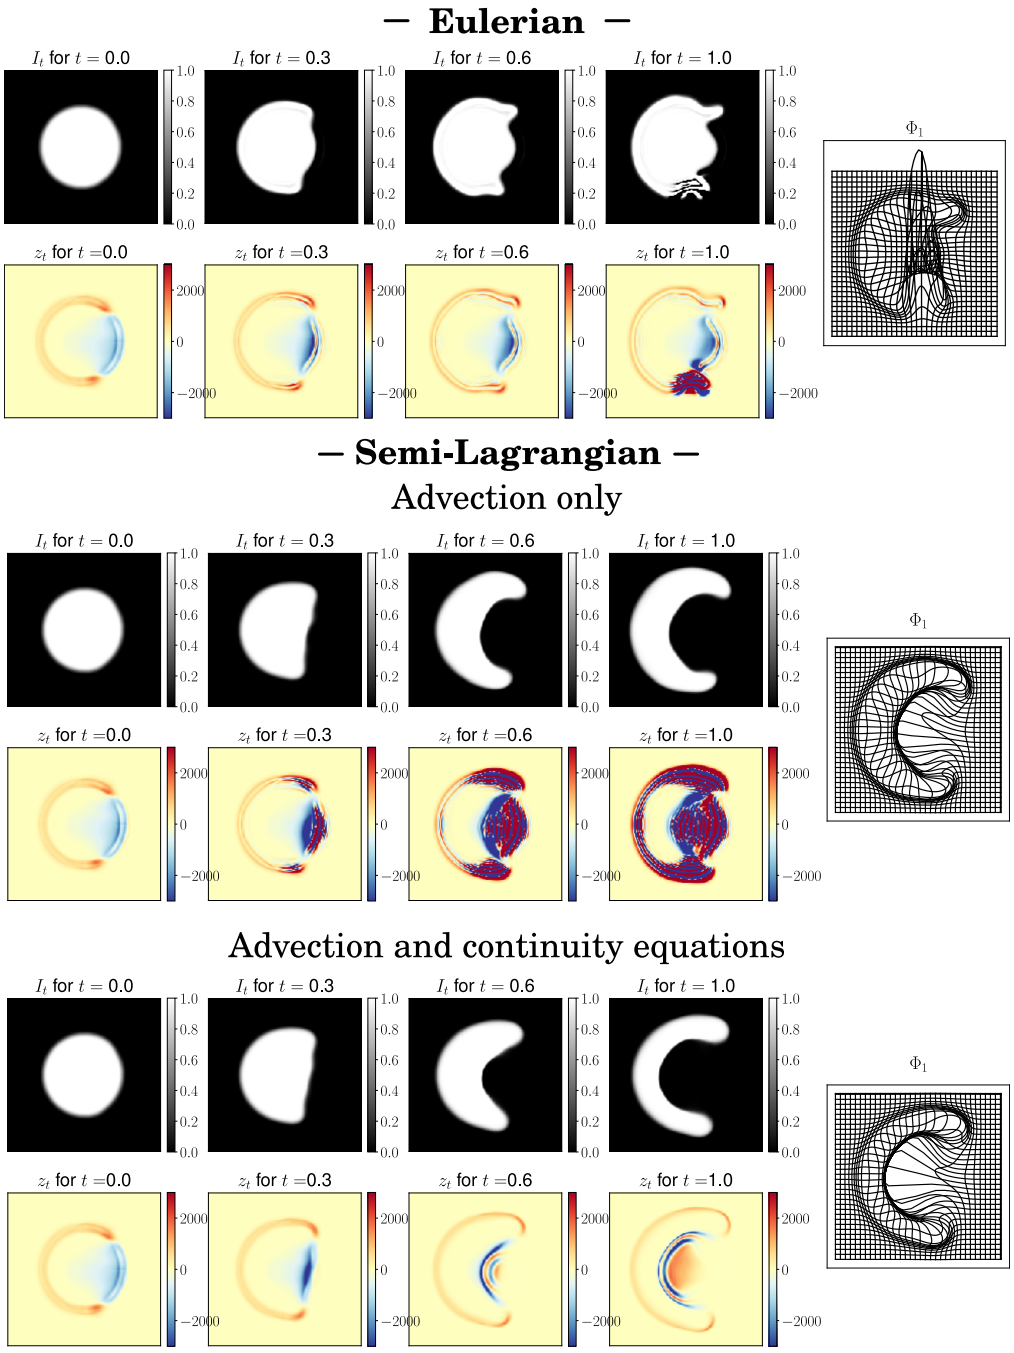

Figure 2.13: Comparison between the stability of the 3 geodesic shootings schemes proposed for LDDMM In each row, we show four intermediary shooting steps with the same initial z_0 and RKHS for v. The black and white pictures are the images, below the corresponding z. The deformations grids on the right are obtained by integrating over all $v_t, \forall t \in [0, 1]$. The shooting was performed using a z_0 obtained from LDDMM optimisation towards a 'C' picture ($\mu = 0$). The Eulerian and semi-Lagrangian schemes have a time step of $1/38$ and $1/20$ respectively.

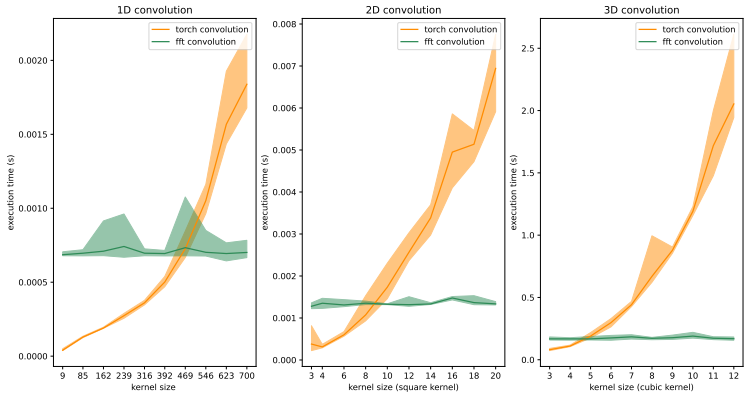

Figure 2.14: Direct versus Fourier convolution benchmark. For each kernel sizes 10 convolutions on a random signal and kernel had been performed. In 2D (resp. 3D) a kernel size of 6 is a 6x6 matrix (resp. 6x6x6). The solid line is the mean computation time, and the filled values are the minimal and maximal computation time in seconds. The benchmark has been realised on my laptop with a Processor: Intel® Core™ i9-9980HK CPU @ 2.40GHz × 16

```
    kernel_size = tuple([big_odd(s) for s in sigma]) # select the good kernel_size
    kernel = get_Gaussian_kernel(kernel_size,sigma) # 2d or 3d depending on the input shape
    return convolution_filter(input,kernel,border_type)
```

2.5.2.b Multi-scale Gaussian kernel

Let's recall the take-away of Section 2.2.4.c: One can sum up the V norms at different scales and take advantage of the combined momenta information. For the method implementation, Sommer et al. [2011] advise applying each kernel K_σ to the corresponding subspace of W and only then computing $\varphi^{\Psi(\mathbf{w})}$. Specifically they precise that building a kernel by summing Gaussians of different scales changes only the shape of the kernel and does not allow different momentum at different scales. This is true for an implementation of the W norm. However, in our case, we just need to compute the updated vector field v from Equation 2.96.a with the multi-scale kernel. Assuming v is fixed, following their advice would mean computing a convolution for each $\sigma \in \mathfrak{S}$ independently, at every time t, at every iteration and then averaging it such that

$$\frac{1}{\#\mathfrak{S}} \sum_{\sigma \in \mathfrak{S}} (K_\sigma \star v) \tag{2.105}$$

However, We say that it is not necessary. Let pose $K_{\mathfrak{S}}$ as the average of the Gaussian kernels of parameter σ, for a fixed $v \in V$

$$K_{\mathfrak{S}} = \frac{1}{\#\mathfrak{S}} \sum_{\sigma \in \mathfrak{S}} K_\sigma, \tag{2.106}$$

it is direct that

$$\frac{1}{\#\mathfrak{S}} \sum_{\sigma \in \mathfrak{S}} (K_\sigma \star v) = K_{\mathfrak{S}} \star v. \tag{2.107}$$

because we perform the convolution on the same vector field v. This allows us to compute the kernel $K_{\mathfrak{S}}$ once before the optimisation starts and perform only one convolution to compute $K_{\mathfrak{S}} \star v$. (see Figure 2.9b). One can find a plain interpretation in Code 2.7 with the kernel computation from other Gaussian kernels.

Source Code 2.7: Sample of code for a multi Gaussian kernel Here is presented the essence of the computation done in the class `multi_GaussianRKHS` within the file `reproducing_kernels.py`.

```
def multiScale_Gaussian_kernel(sigma_list,input):
    """
    sigma: list of tuples of `dim' length
        ex: [(3,3,3),(5,4,3)], code for two 3d kernels with sigma 3 in each direction
        for the first and with variable sigma in the second.
    """
    ## Make the kernel
    # For each sigma, compute the maximum kernel size for each dimension
    _ks = []
    for sigma in list_sigmas:
        big_odd = lambda val : max(6,int(val*6)) + (1 - max(6,int(val*6)) %2)
        kernel_size = tuple([big_odd(s) for s in sigma])
        _ks.append(kernel_size)
    # Get the max of each dimension.
    _dim = len(kernel_size)
    kernel_size = tuple([
        max([s[i] for s in _ks ]) for i in range(_dim)
                ])
    list_sigma = list_sigmas

    # Make a mean of all kernels for each filtered pixel
    kernel = torch.cat(
        [ prod(sigma)*get_Gaussian_kernel(kernel_size,sigma) for sigma in list_sigmas ]
            ).sum(dim=0)
    kernel /= len(list_sigmas)

    ## Apply the filter
    return self.filter(input,kernel)

```

2.5.2.c Choice of scale, LDDMM vs LDDKBM

Before looking at more natural data, I would like to compare the Gaussian kernel from Definition 2.14 with the multi-scale one from Section 2.2.4.c. The comparison will be done using toy examples from Figure 2.15. The target image (b) has been constructed to be very challenging to register from the source image (a). In the examples, all images are 300×300 pixel wide. On the panel (c) one can see the source and target images concatenated together in an RGB image. Indeed the deformation needs to go through the narrow passes (red arrows, circle) and then expand the source image to fill the pointy shapes (blue arrows, triangle) while keeping the transformation diffeomorphic. In itself, the pointy shapes would be hard to match from an object with a smooth outline, because the diffeomorphism nature tends

(a) Source **(b)** Target **(c)** Source vs Target

Figure 2.15: Source and Target image of Figure. (c) The comparison of images is made by superposing two images in different channels such that matching parts appear green and parts that are intense in only one image appear blue or red. Arrows of (b) are explained in the text below. 2.16

to preserve features like edges. For example, registering a circle targeting a square will end up as a rounded square.

In Figure 2.16 we can find registration examples at different scales σ. The three first columns show the registration integration at times .1, .5, and 1 respectively, and the two middle ones show the registered image superposed on the target along with the corresponding deformation. Finally, the last column shows the shape of each kernel, which is originally 2D. One can notice their respective size according to the discussion in Section 2.2.4.b. The four top rows show four examples using classical mono-scale LDDMM, and the four bottom ones show examples of multi-scale LDDKBM as described in Section 2.2.4.c. The analysis of the mono-scale results indicates that σ is a critical parameter to tune. Indeed, we see that choosing a σ too small prevents the momentum and the deformation is minimal; choosing a σ too big brings to a less accurate deformation. The best registration displayed according to the SSD was obtained for $\sigma = 10$. However, the three top registration are not diffeomorphic and are therefore not admissible: the red dots behind the deformation grid indicate a negative determinant of the deformation Jacobian. For $\sigma = 20$ the deformation is diffeomorphic but the registration is quite poor.

The last four rows show that by choosing a set of sigmas about in the good scale, we obtain good results with little only variations on the obtained matchings. In the right column, one can see the different Gaussian kernel shapes used with the kernel applied. The black solid line is the average of the other. Sommer et al. [2011] advised against using a unique kernel as a mixture of others. However, we found that because of our implementation, it is a reasonable choice as we detail in Section 2.5.2.b. For now, we can observe that a kernel like the one $\sigma = \{1, 10, 20\}$ takes high importance on details thanks to the centre pointy shape while smoothing the surroundings with its tail.

2.5.3 Sharp integration

This work is inspired by an idea of Matthis Maillard.

As we saw in the last paragraph, the semi-Lagrangian integration scheme allows us to register images with fewer integration steps, however, the successive interpolation errors make the output image blurry. We will discuss here a method to use the semi-Lagrangian approach while keeping the output image sharp. In methods registering using pure deformation like LDDMM, the solution is simply to compose the fields to an identity grid and

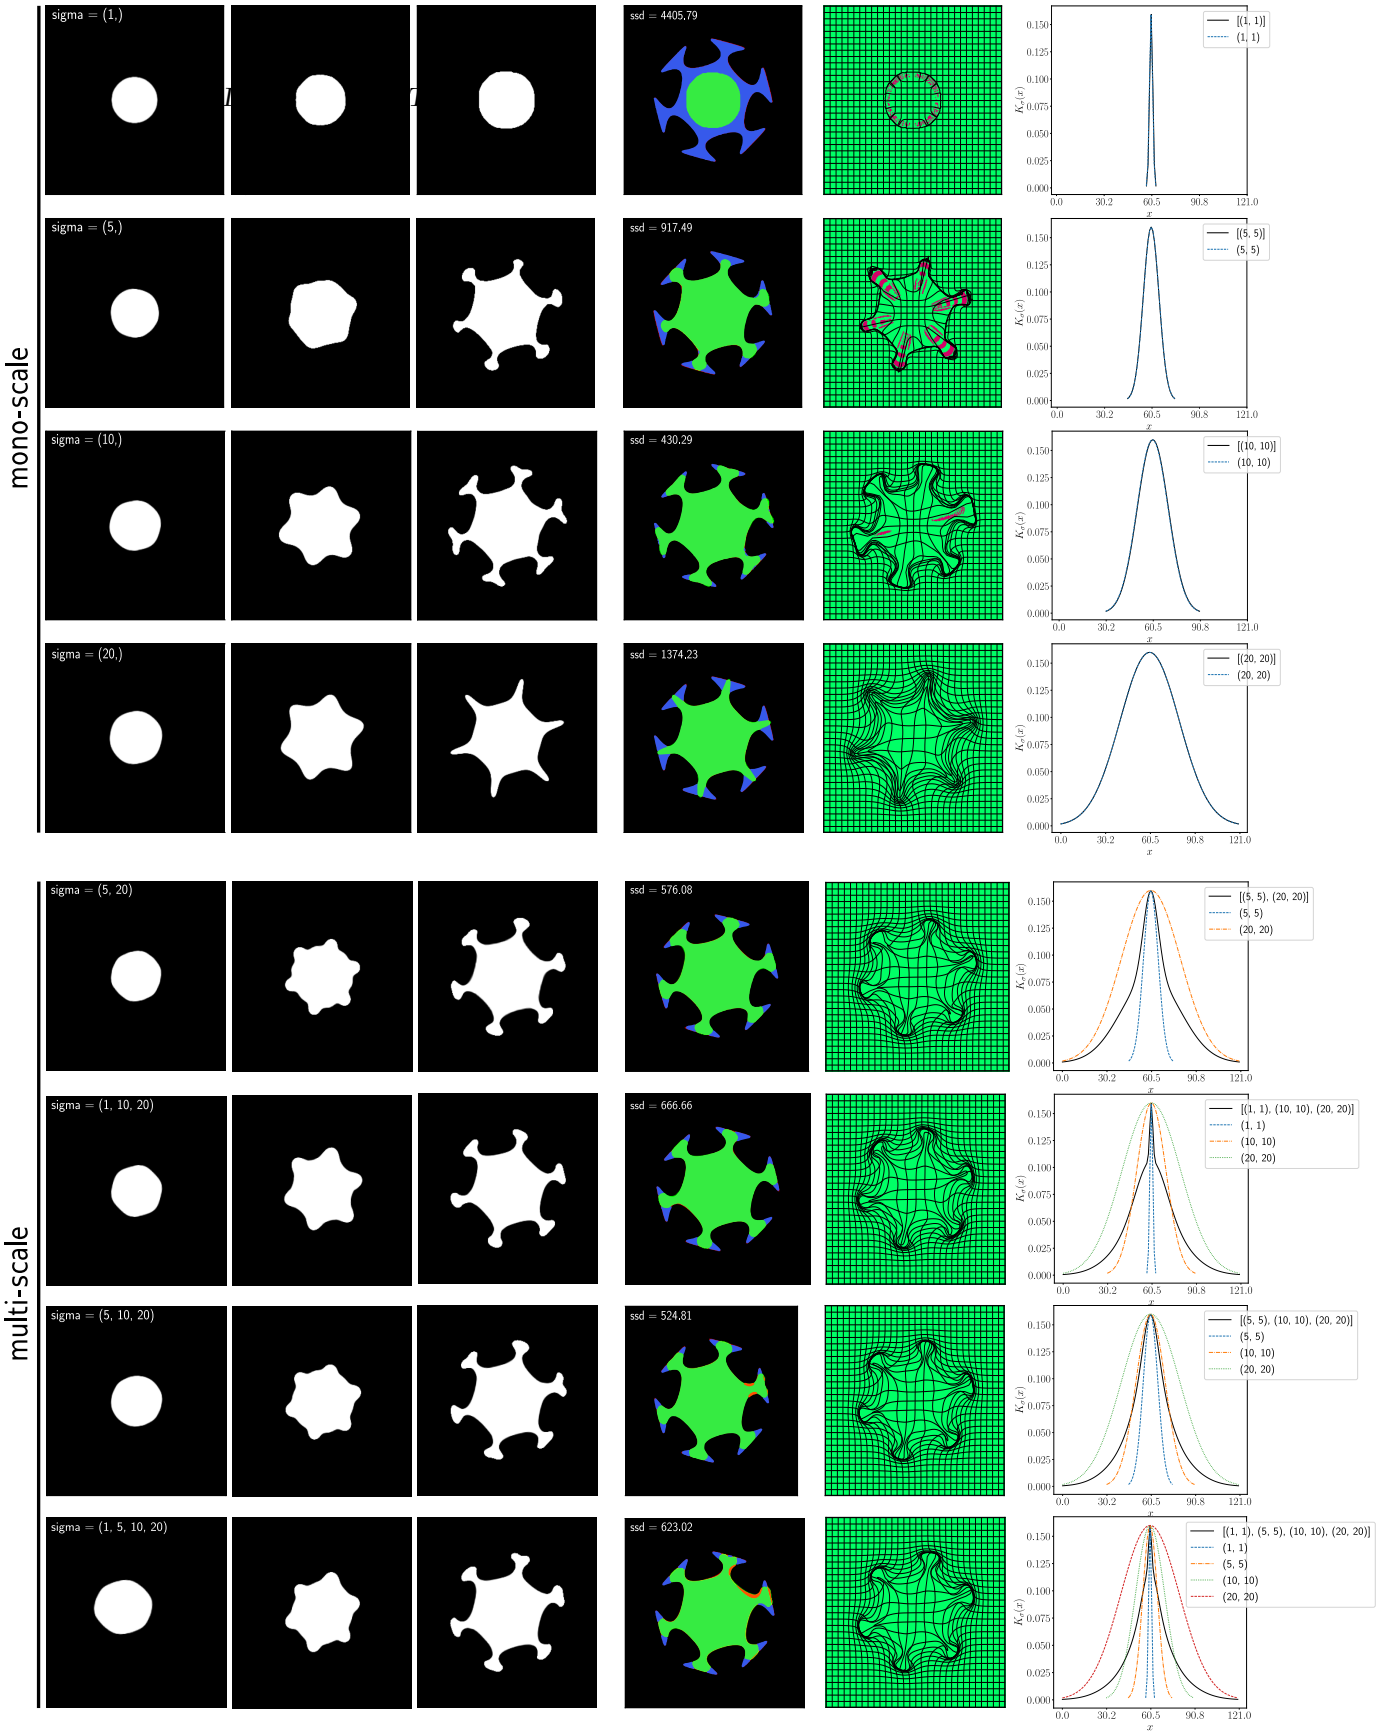

Figure 2.16: Registering details depending on the Gaussian kernel scale. In order to fit the big figure, the caption detail can be found within the text.

apply the subsequent deformation to the image, interpolating it only once. We first generate a deformation flow φ_1 such that $\dot{\varphi}_t = v_t \circ \Phi_t$, $(v_t)_{t\in[0,1]} \in V$ being a temporal vector field and thus :

$$\varphi^{\mathbf{v}} = \mathrm{Id} + \int_0^1 v_t \circ \varphi_t^{\mathbf{v}} dt. \tag{2.108}$$

Then, we want to be obtain a sharp image by simply composing $\varphi^{\mathbf{v}}$ to the source image S. As one can expect, it is not as simple for Metamorphosis. The image integration is given by the ODE: $\dot{I} = -v_t \cdot \nabla I_t + \mu z_t$, at each time, intensity is added (or subtracted) by z_t that will be later transported by v_l, with $t < l \leq 1$. A continuous formulation can be found in Trouvé and Younes [2005](Theorem 2) give some insight on the image integration:

$$I_t = I_0 \circ (\varphi_{0,t}^{\mathbf{v}})^{-1} + \mu \int_0^t z_s \circ \varphi_{s,t}^{\mathbf{v}} ds \tag{2.109}$$

If we want to build an integration scheme based on semi-Lagrangian formulation keeping the image sharpness we need to construct carefully the deformations of the residuals. We will provide some definitions before detailing the "sharp integrator".

Definition 2.17 (discrete forward deformation)**.** *Let be* $n, k \in \mathbb{N}, 0 \leq k \leq n \leq N$

$$\varphi_{n+1}^k = (\mathrm{Id} - \mathrm{v}_n) \circ (\mathrm{Id} - \mathrm{v}_{n-1}) \circ \cdots \circ (\mathrm{Id} - \mathrm{v}_{k+1}) \circ (\mathrm{Id} - \mathrm{v}_k)$$

with the convention : $\varphi_n^n = \mathrm{Id}$ *and* $\varphi_n^{n+1} = (\mathrm{Id} + \mathrm{v}_n)$

With this notation we have an immediate property :

$$\varphi_{n+1}^k = \varphi_{n+1}^n \circ \varphi_n^k. \tag{2.110}$$

We will use this to avoid re-doing many time the same deformation computation. We can also define the backward deformation

Definition 2.18 (discrete backward deformation)**.** *Let be* $n, k \in \mathbb{N}, 0 \leq k \leq n+1 = N$

$$\varphi_k^n = (\mathrm{Id} + \mathrm{v}_k) \circ (\mathrm{Id} + \mathrm{v}_{k+1}) \circ \cdots \circ (\mathrm{Id} + \mathrm{v}_{n+1}) \tag{2.111}$$

with the convention : $\varphi_{n+1}^{n+1} = \mathrm{Id}$ *and* $\varphi_{n+1}^n = (\mathrm{Id} + \mathrm{v}_n)$

Preparing the implementation, Figure 2.17 helps to understand the operation order. In each coloured box, one must first compute the one on the inside by applying the previously computed residuals before summing them. To do so we need to keep track of all the deformation as well. We will call the lower triangular matrix such that

$$\begin{aligned}\Phi &= \begin{pmatrix} \varphi_1^0 & & & & \\ \varphi_2^0 & \varphi_2^1 & & & \\ \varphi_3^0 & \varphi_3^1 & \varphi_3^2 & & \\ \vdots & & & \ddots & \\ \varphi_N^0 & \cdots & \varphi_N^k & \cdots & \varphi_N^{N-1} \end{pmatrix} \\ &= \begin{pmatrix} \varphi_1^0 & & & & & \\ \varphi_2^1 \circ \varphi_1^0 & & \varphi_2^1 & & & \\ \varphi_3^2 \circ \varphi_2^0 & & \varphi_3^2 \circ \varphi_2^1 & \varphi_3^2 & & \\ \vdots & & & & \ddots & \\ \varphi_N^{N-1} \circ \varphi_{N-1}^0 & \cdots & & \varphi_N^{N-1} \circ \varphi_{N-1}^k & \cdots & \varphi_N^{N-1} \end{pmatrix}. \end{aligned} \tag{2.112}$$

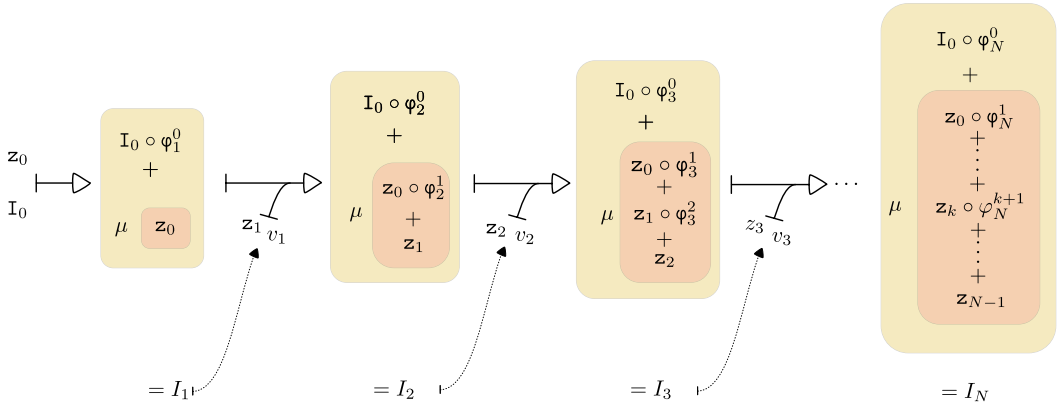

Figure 2.17: Sharp Metamorphosis Integration scheme

At this point, one can already see an algorithmic trade-off has to be made between memory and computation. To save computation time, one can compute the k-th row of Φ can be done by computing φ_k^{k-1} and composing it with the $(k-1)$-th row (see Equation 2.112). Alternatively, one can recompute every deformation. When the residual is computed, one can simply add it to the transported image.

We are now ready to tackle its implementation which can be found in Code 2.8. The main function is `_step_sharp_semiLagrangian_`, the computation of the Matrix 2.112 is computed at each time step within function `_update_sharp_intermediary_field_` and the deformations of each residual are computed within `_compute_sharp_intermediary_residuals_`.

Source Code 2.8: Integration step of the sharp Metamorphosis scheme class `Metamorphosis_path` within the file `metamorphosis.py`. Note that this implementation use `_update_field_` and `_update_residual_semiLagrangian_` that were defined in Code 2.4 and 2.5 respectively.

```
def _update_sharp_intermediary_field_(self):
    self._phis[self._i][self._i] = self.id_grid - self.field/self.n_step
    if self._i > 0:
        for k,phi in enumerate(self._phis[self._i - 1]):
            self._phis[self._i][k] = phi + tb.compose_fields(-self.field/self.n_step,phi)

def _compute_sharp_intermediary_residuals_(self):
    resi_cumul = torch.zeros(self.residuals.shape,device=device)
    for k,phi in enumerate(self._phis[self._i][1:]):
        resi_cumul += tb.imgDeform(self.residuals_stock[k][None],
                                   phi,
                                   clamp=False)
    resi_cumul = resi_cumul + self.residuals
    return resi_cumul

def _step_sharp_semiLagrangian(step):
    self._update_field_()
    self._update_sharp_intermediary_field_()

```

```
    resi_cumul = self._compute_sharp_intermediary_residuals_()

    # free the memory
    if self._i > 0: self._phis[self._i - 1] = None

    self._update_image_semiLagrangian_(source,resi_cumul,
                                       sharp=True)
    # actual code in _update_image_semiLagrangian_()
    self.image = tb.imgDeform(self.source,self._phis[self._i][0])
    self.image += (resi_cumul*self.mu)/self.n_step


    self._update_residuals_semiLagrangian_(self._phis[self._i][self._i])

    return (self.image, self.field, self.residuals)
```

On a general note, one can remark that the method to recover the intensity addition for each voxel in the source image can be deduced from these pieces of code. In other words, let $x_i \in \Omega$ be a voxel, at time t it is transported to the position $\varphi^{\mathbf{v}}_{0,t}(x_i)$. Let $J_t(x_i)$ be the cumulative intensity of x_i at time t starting from zero, defined such that:

$$J_t(x_i) = \mu \int_0^t z_s \circ \varphi^{\mathbf{v}}_{t,s}(x_i) ds. \tag{2.113}$$

Note that $\varphi^{\mathbf{v}}_{t,s}$ is the backward transformation as $s \leq t$. To find the total intensity changes added at the shooting end, we need to compute the collections of deformations: $\{\mathrm{Id}, \varphi_1^0, \varphi_2^1, \cdots, \varphi_0^N\}$ before composing them to the saved residuals $\{\mathtt{z}_0, \mathtt{z}_1, \cdots, \mathtt{z}_N\}$ and summing the result. This resembles the computations done in the method `_compute_sharp_intermediary_residuals_`

RESULTS AND DISCUSSION The sharp integration scheme is compared with the classical one in Figure 2.18. In this subsection, the 'classical scheme' refers to the non-sharp one. The images to match are difficult, the big circle must be morphed into the three-horned one while making one disk appear on the right of the shape, and a disk disappears at the centre. From a global perspective, the two integration schemes show similar performances. However, both produce visual defects. The classical scheme becomes increasingly blurry with successive interpolations (see zoom D), while the sharp scheme creates local instabilities (see zoom A). However, for the sharp scheme, these instabilities are effectively erased by the addition of intensities. In reality, these instabilities appear in both methods, but successive interpolations erase them. We have not determined with certainty why they appear, but they only occur when the deformation is significant and the number of integration steps is too low. If we focus on the appearance/disappearance of the disks, both must be explained by changes in intensities. However, we observe two scenarios. The disappearing sphere disappears gradually without generating deformation, while the other appears from right to left. This is due to the fact that the shape of the centre undergoes a local deformation to the left at that location. The addition of intensity thus follows the movement (see zoom B-D). Additionally, in this example, the disk that appears to the right of the shape causes strong growth in deformation only in the sharp image (see panel a). The advantage that

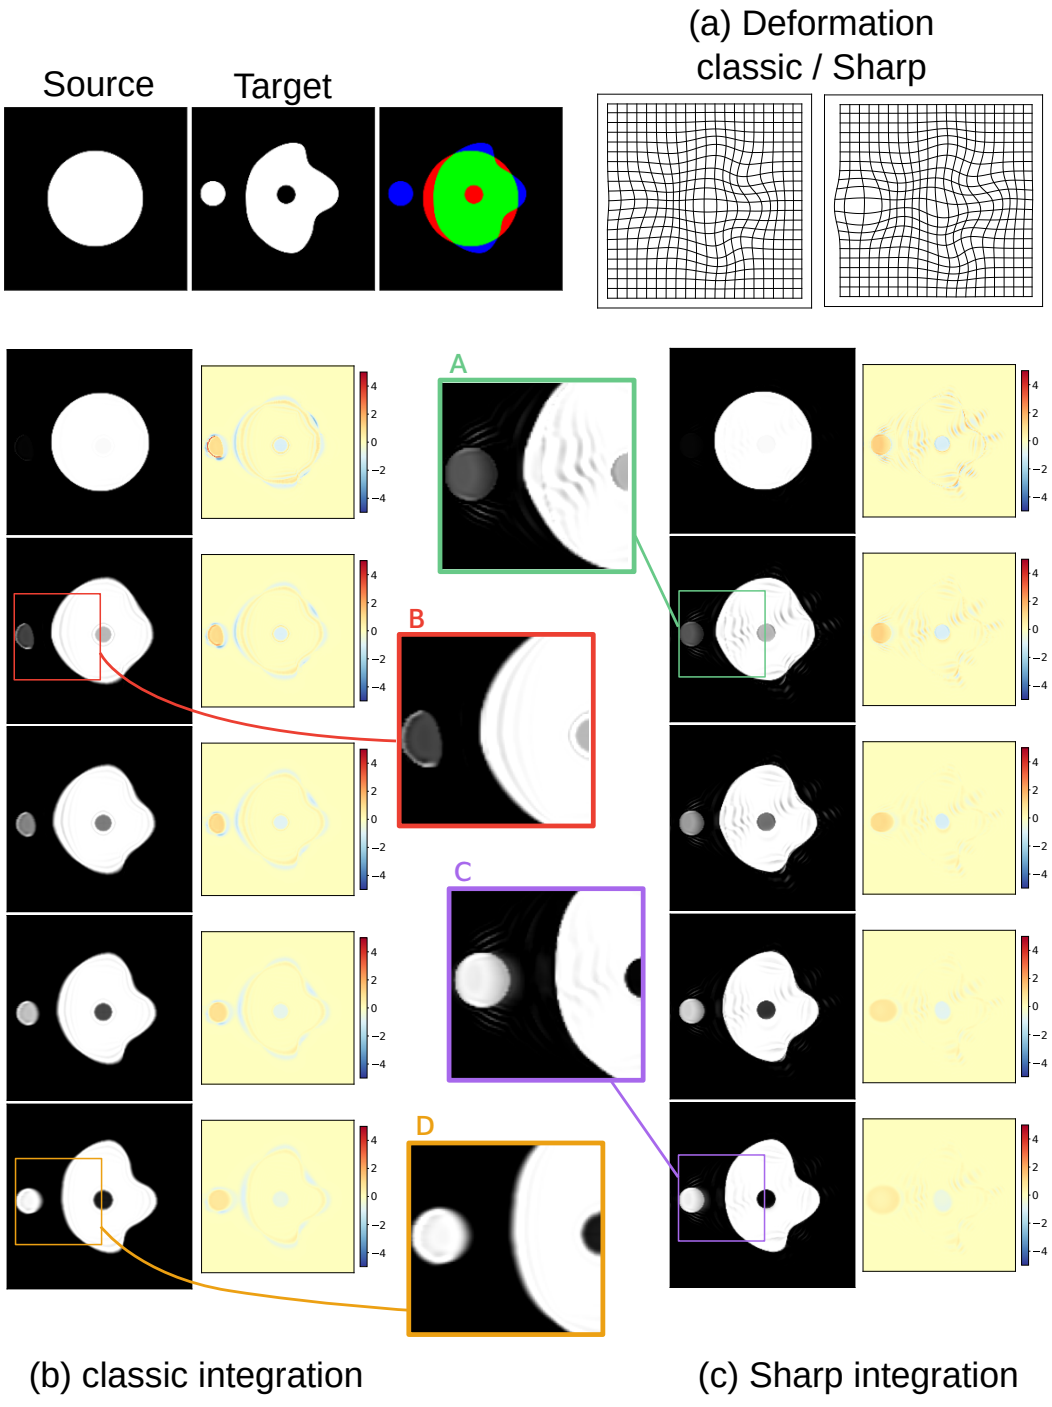

Figure 2.18: Metamorphosis classical vs Sharp integration. Both integration were conducted with 50 time steps and $\mu = 1, \rho = 10, \sigma = 15$.

the classical integration scheme retains over the sharp scheme is that it is significantly less complex and faster to execute.

2.5.4 Metamorphosis - General Results & Discussion

If, in theory, Metamorphosis can always archive a perfect registration at least using photometric changes, the geodesic varies with the hyper-parameters. Therefore, changing the image integration, and the deformation as a whole. One can see it in Figure 2.19, where the matching problem is quite challenging for several reasons. First, the 'ideal' matching to perform is unclear. Arguably, one would want to transform the big starting circle to the 'C' shape, rather than re-scaling it to the small circle. However, the cut into the 'C', which can be seen as a local occlusion in the capture, should not be taken into account. Second, as can be seen in the Source/Target comparison, the large circle overlap all structures in the target image. Due to the use of a L^2 norm, the algorithm could be tempted to match the three structures at the same time. Third, the displacement to estimate is very large compared to the image size (about half of its width), making the integration difficult and increasing the number of time steps required.

From the study of Figure 2.19 alone, we can learn a lot about the behaviour of Metamorphosis. We break these intuitions into points:

- **Residual entanglement:** By studying residual columns (in colour, left) one can visualise the momentum value over time. In column (a), it is clear that most of the registration is done by intensity additions, with the negative (blues) and positive (orange) areas of the residual. On the other two (b, c), with a higher ρ the positive (resp. negative) part of the residual will generate a deformation in the direction (resp. opposite direction) of the image gradient. At each time t and at each pixel x, a given value $z(t, x)$ encodes both deformation and intensity changes.

- ρ **variation:** In the geodesics equations 2.96 a higher ρ will promote higher amplitude vector fields and consequently a larger deformation. It can be seen as a way to control the entanglement discussed in the previous point. We clearly see that if all Metamorphosis integration ends up being very similar to the target, deformations differ greatly. Thus depending on ρ the amount of intensities changes varies as well. Because of this dependence, it is hard to analyse the output of metamorphoses. If one wants to analyse those quantities, for example assessing the number of topology changes, one must take ρ into account. One can see that any of the Metamorphosis's results we present retrieve the full 'C' with deformation. However by increasing ρ ever more, one will get results closer to LDDMM.

- μ **variation to get LDDMM** μ control the amount of intensity added at each time step. It works together with ρ. The actual important value to tune is their ratio $p = \rho/\mu$, which I call p the power magnitude of the field. Thus, on the contrary of what I said in François et al. [2022], I advise setting $\mu = 1$ when one wants to do Metamorphosis. If one wants to match with LDDMM, the framework allows it and it suffices to set $\mu = \rho = 0$.

- **Semi-Lagrangian scheme makes deformations time step independent** for a fixed ρ. If we decrease the number of time steps with Eulerian schemes, the process may diverge, not producing any results. With semi-Lagrangian schemes, they converge

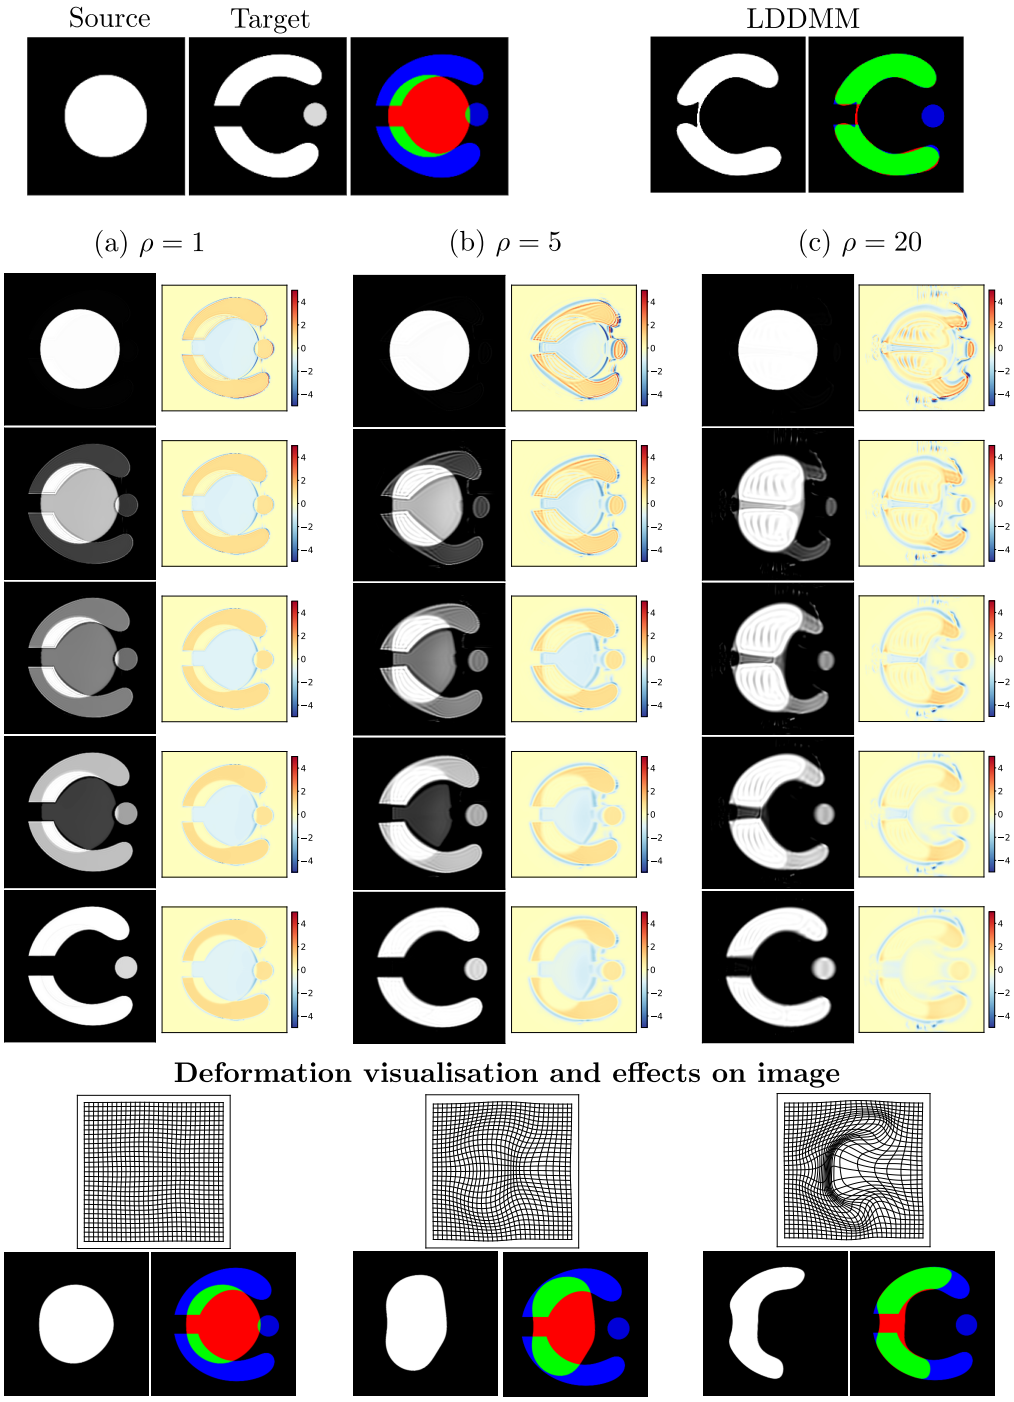

Figure 2.19: Influence of ρ on the Metamorphic integration ($\mu = 1$). On each panels,(left) Image integration I_t, top: $t = .1$,bottom $t = 1$. (right) Residuals z_t, (bottom) Integrated deformation. The coloured image at the top is the comparison of source vs target images. It is made by superposing the two images in different channels such that matching parts appear green and parts that are intense in only one image appear blue or red.

most of the time and produce similar deformations for different time steps. One can set the number of steps down to 3 still retrieving the deformation, making the computation very fast. However, it produces some instabilities when adding intensities and generates artefacts on images. It is still useful for parameter tuning.

- **Automatic hyper-parameter choice.** In contemplating the potential avenues for future research, it is posited that the development of a methodology that employs a semi-Newton scheme, Bayesian inference or machine learning algorithms to automatically identify optimal parameters would be highly desirable. While currently more of a wishful comment than a concrete perspective, the advent of such a technique will make Metamorphosis significantly easier to use.

- **Have you tried using AI ?** It is a good and legit question! We have presented in Maillard et al. [2022] a deep residual learning (ResNet-based) implementation of Metamorphosis that drastically reduces the computational time at inference. Indeed, if the training is very time-consuming, this type of method is much faster at inference time than classical Metamorphosis which needs to optimize a function for every pair of images. The registration principle is inspired by Metamorphosis, and images still follow PDEs similar to 2.96.a and 2.96.c. However, the divergence term of the geodesic equation is replaced by a ResNet, estimating a new unknown high-dimensional function over which we have no control. On one hand, it is very efficient to estimate deformation. On the other hand, as we do not follow Euler-Lagrange equations, using it to build a shape-space is questionable.

- **Discussing intensity addition through semi-Lagrangian schemes.** As mentioned earlier semi-Lagrangian schemes increased the PDE integration stability, making it usable. However, it is still not perfect. We aim to model the transport of an image I through the relation $\dot{I}_t + v_t \cdot \nabla I_t = u_t$ where $u_t : \Omega \times [0,1] \to \mathbb{R}$, making it an equation of transport with an external source as introduced in Section 2.5.1.a. Solving this scheme numerically is challenging depending on the control one has on u, as discussed there is room for improvement. After the optimisation is completed, assuming that z_t is known for all t, it is possible to recover the intensity addition for each voxel in the source image using a process similar to the one presented for the sharp integrator in subsection 2.5.3. However, its interpretation remains unclear and it shares the same pro and cons as the said sharp integrator.

- **Does Metamorphosis gives a better result than LDDMM?** The short answer is ... it depends on the objective. It is the case if one aims at having a good and smooth photometric and geometric deformation linking two images. Indeed, Metamorphosis allows going beyond the image orbit by the action of Diff. However, if one aims at analysing the deformation it is less clear. Indeed both method can fail in different ways: In figure 2.19 LDDMM match the circle with the 'C' acknowledging the cut without a topology break when Metamorphosis register by deformation and compensate the rest with intensity changes. For intuition, one can see the intensity change as the reverse projection of the infinitesimal intensity changes δ to the infinitesimal deformation by g (see Figure 2.11).

The effectiveness of the Metamorphosis method in the registration of 3D brain MRI remains to be seen. Results shown in Figure 2.20 demonstrate that the deformations produced by Metamorphosis seem to make sense in minimising the sum of squared differences. However, the application of LDDMM results in the expansion of the ventricles and grey matter

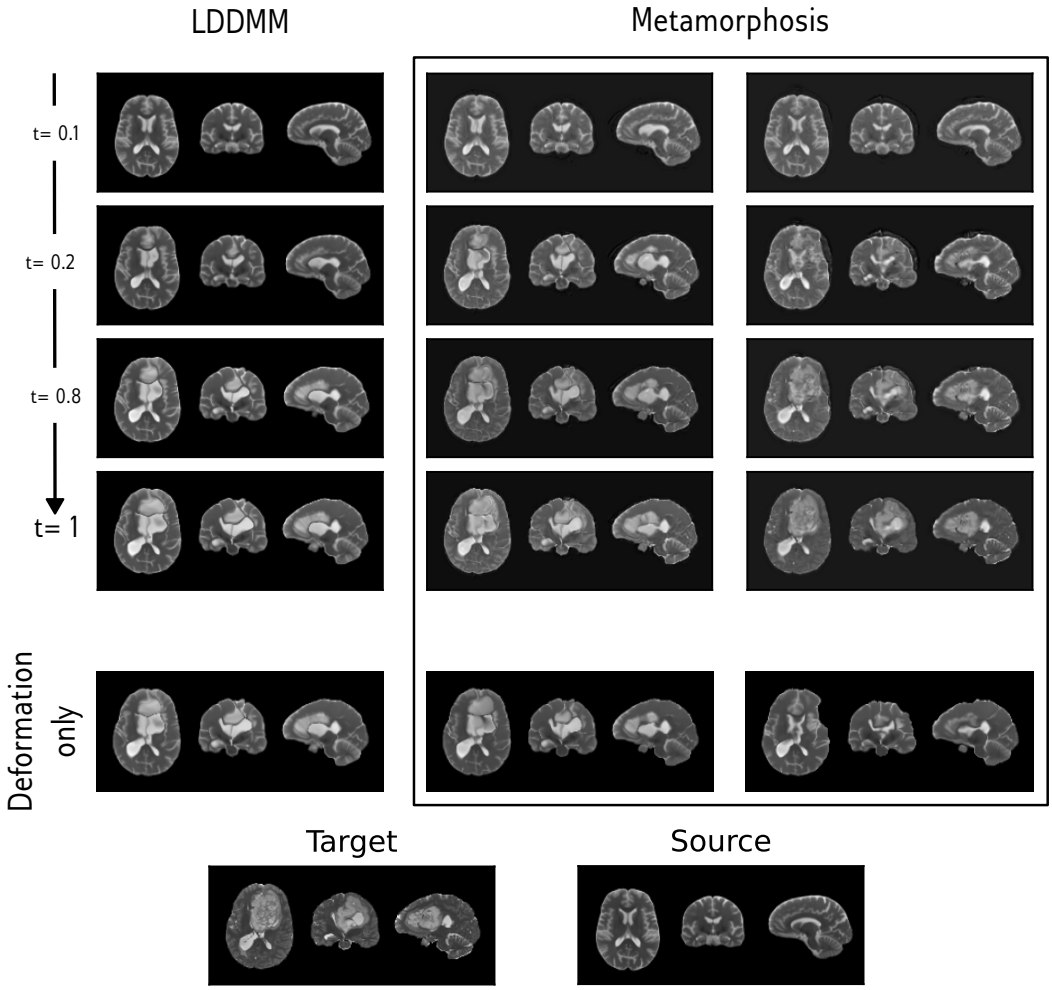

Figure 2.20: "Failed" registration LDDMM & Metamorphosis on brains with Glioblastomas

folds in the frontal lobes. The utilisation of cost-function masking (CFM) would produce less aberrant deformation at the cost of not producing any displacement within the mask, particularly in the front ventricles. On the other hand, the application of Metamorphosis results in intense changes in the tumour and a non-realistic deformation that squishes a portion of the brain. In short, the final result resembles the target, but the deformation produced is not anatomically plausible. Neither LDDMM nor Metamorphosis prove to be suitable tools for the registration of the pathological brain with tumors in its current form. We extend the framework with Constrained Metamorphoses, which will be introduced in the subsequent chapter as a potential solution to the aforementioned problems.

2.6 Comic: Semi-Lagrangian schemes

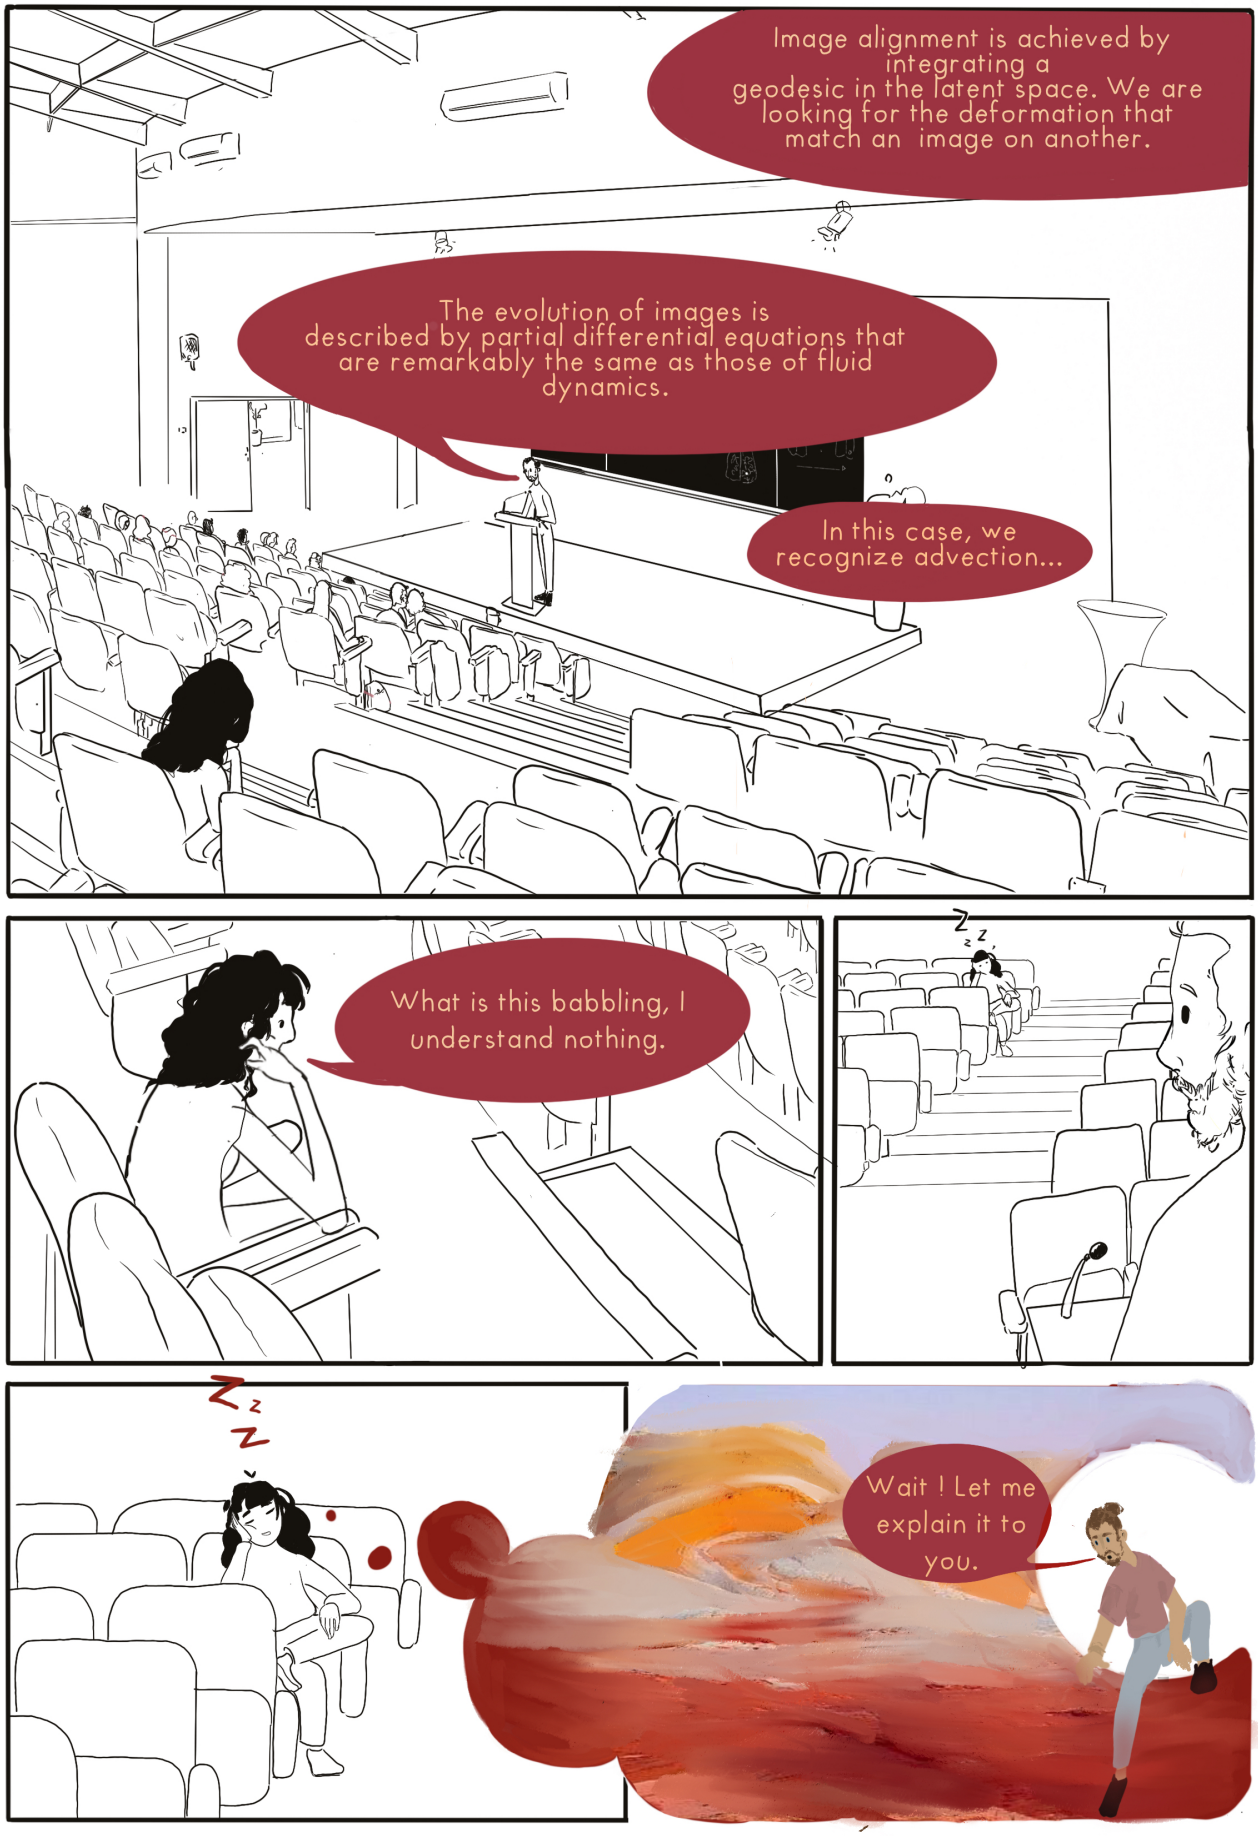
Image alignment is achieved by integrating a geodesic in the latent space. We are looking for the deformation that match an image on another.
The evolution of images is described by partial differential equations that are remarkably the same as those of fluid dynamics.
In this case, we recognize advection...
What is this babbling, I understand nothing.
Wait ! Let me explain it to you.


You may try, but I do not make any promises to understand.
You will see, It is not that hard.
I was aiming at implementing Metamorphoses, I had at disposition the formula of a straight line between an image pair.
I have noticed, as other before me, that these equations bear a strong resemblance to the equations of fluid dynamics.
If these equations describe the evolution of the fluid, it must be possible to perform this operation using a computer in order to represent it.
It seams difficult
Indeed
Thankfully, a whole branch of physics tries to answer that. But I know nothing about it...
Let's see Lagrange, maybe he will give us an idea.

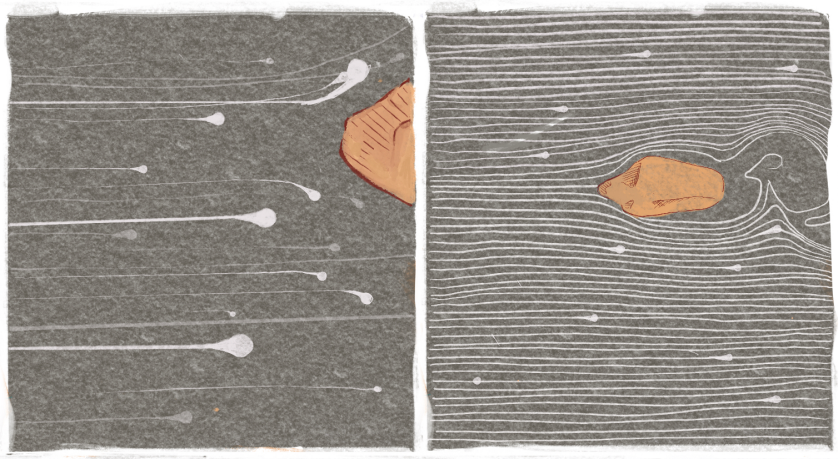

A fluid can be represented by a collection of particles, in water it would be water molecules.

In the liquid state, the flow of the fluid is continuous and regular, a mathematician could say differentiable. This means that locally, the particles are going in the same general direction. To reformulate it: if we take a particle at random at a time t, its neighboring molecule will have a similar trajectory.

This property means that to describe the flow as a whole, it's sufficient to follow a few well-chosen particles!


This seems like a tiring technique and in practice
it wouldn't be useful for me. The particles are likely to move too far away from each other, or worse, leave the image...
Oh, look, there's Euler, he'll surely have another idea for studying flows.
134,135 ...
Excuse me sir, but why are you counting?
Oh ?! ... Hello youngsters, I didn't hear you coming.
To describe the flow, unlike Lagrange, I prefer to describe the flow by the quantity of particles that cross a given place.
And I don't like to run
And once I know the quantity of particles in my basket
, I only need to count the quantity of particles that enter and leave.
This variation is the derivative.
Wow, it's perfect for describing the evolution of an image over time. Each basket here is a pixel, and since I know the amount of variation in the number of particles in a basket based on its neighbors, I can describe the evolution of my entire image.
Very well, but I do not see the link with your thesis.
Our goal is to do an image registration. That is to find the vectors that describe the flow, which will in turn carry our image. Hopefully it will end up looking as much as possible like the target image.

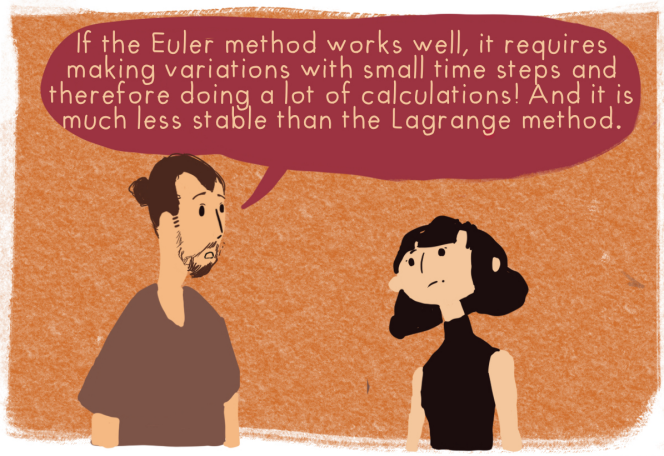

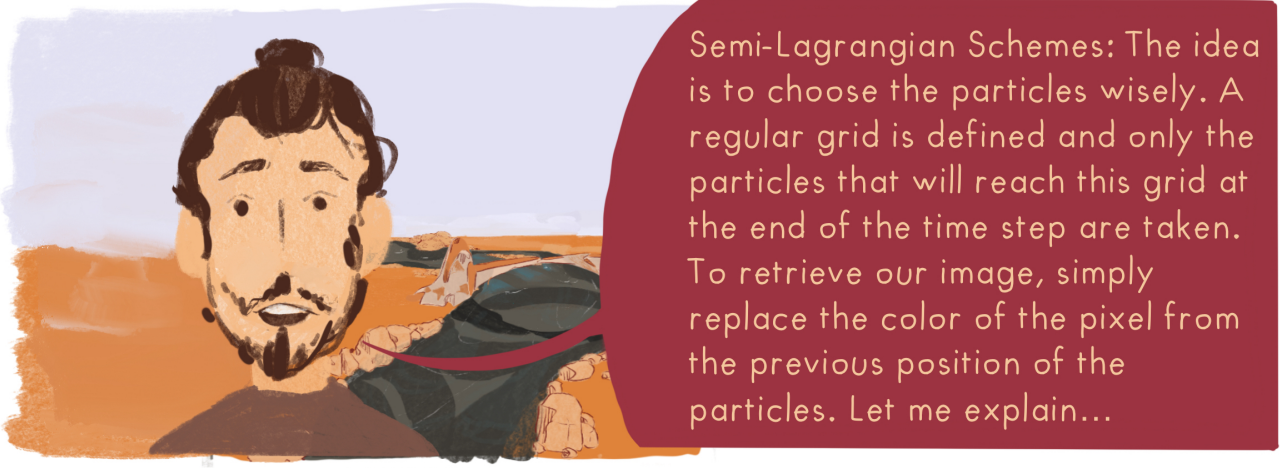

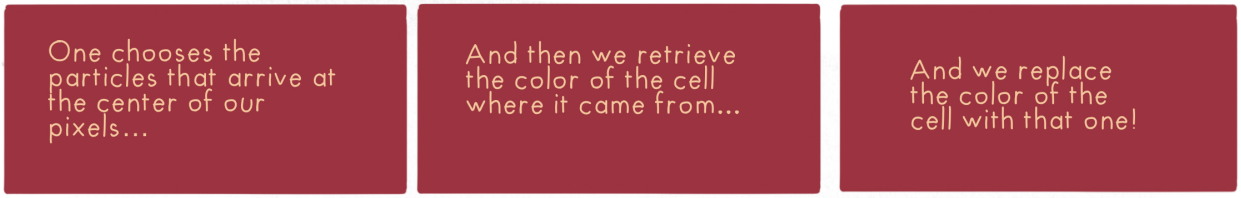

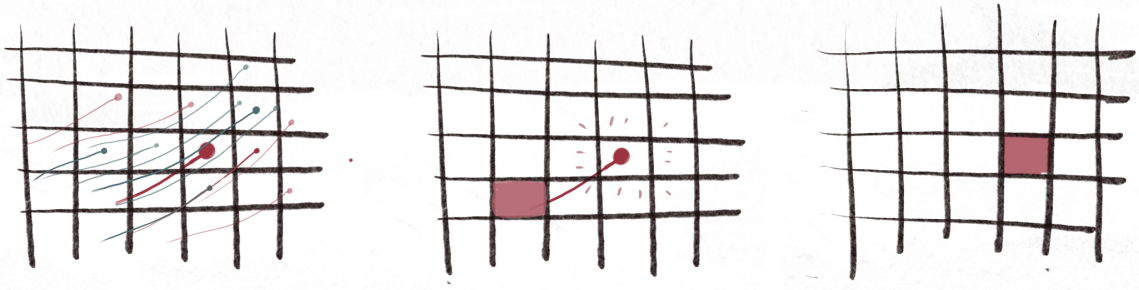

Finally, we can see the evolution of such a numerical scheme

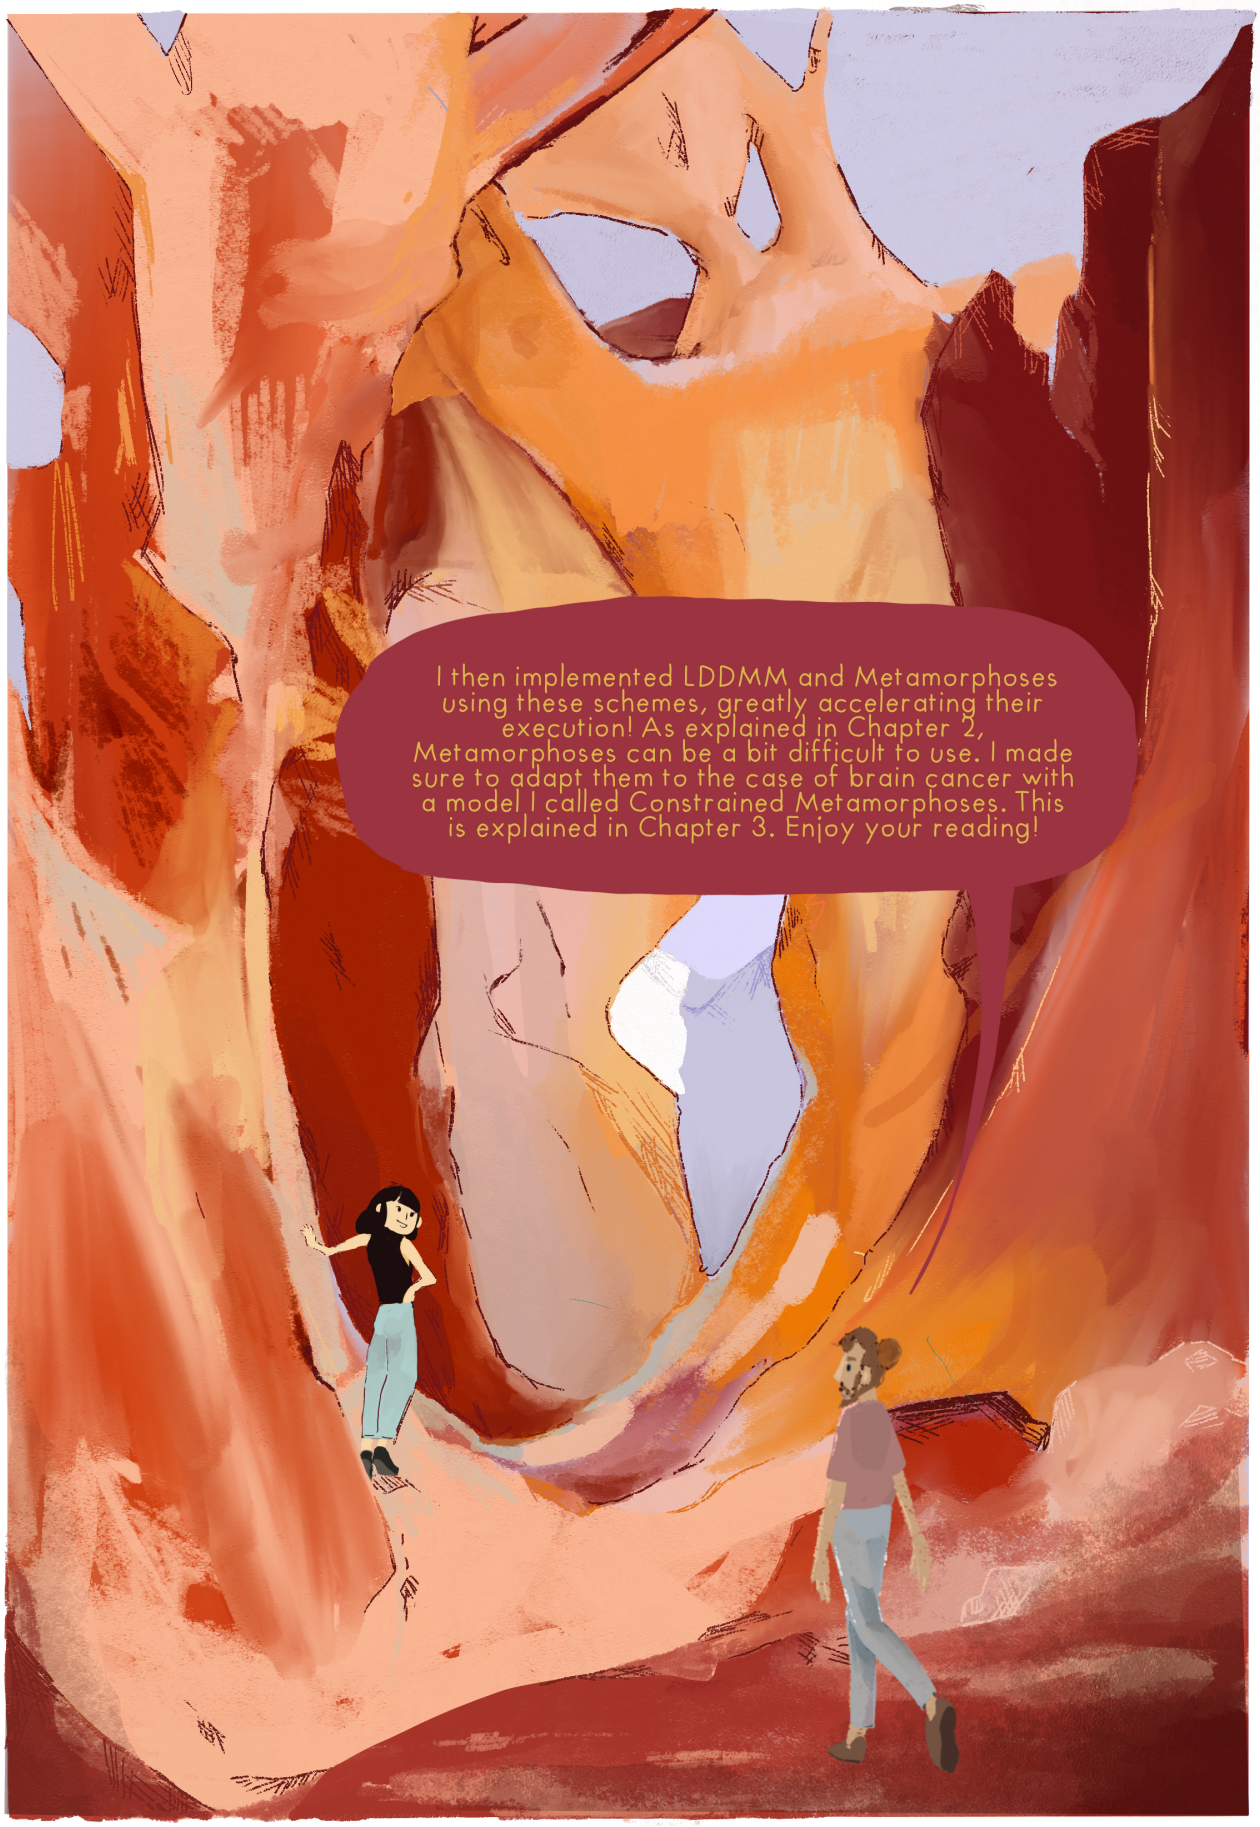
I then implemented LDDMM and Metamorphoses using these schemes, greatly accelerating their execution! As explained in Chapter 2, Metamorphoses can be a bit difficult to use. I made sure to adapt them to the case of brain cancer with a model I called Constrained Metamorphoses. This is explained in Chapter 3. Enjoy your reading!

3

Constrained & Weighted Metamorphosis

In this chapter, we will cover this dissertation's main contributions. Considering the practical problems encountered while using the classical Metamorphosis algorithm, aiming at registering MRI with glioblastomas (also see Section 3.1.1). Acknowledging the biological nature of the data, we should not explain topological change solely by intensity apparition. In medical image analysis, one often comes with a lot of prior information, as we know the expected morphological structures of the organs and their underlying mechanisms. We will focus on the specific case of registering cancerous brains to healthy ones. In other terms, a matching problem where the topology between images differs, with in addition a 'mass effect' induced by the appearing tumour.

This chapter aims to present an extension of the Metamorphic framework, where one can constrain deformations to behave as dictated by a given model derived from prior information. In our case, we will constrain the localisation of topological addition (*i.e.:* intensity changes) and locally the direction of the flow, roughly modelling the tumour growth via LDDMM. We will use a glioblastoma segmentation as prior, and deduce a 'growing' flow by registering a small ball to the segmentation. We detail our strategy in Section 3.1. The objective of the Constrained Metamorphosis (CM) framework is to register a pair of images taking into account those priors. It can be seen as solving a constrained optimisation problem. In short, we design new energies by leveraging prior anatomical and biological information from which geodesic equations are deduced (i.e., Euler-Lagrange equations). Similarly to Metamorphosis, we have a system of PDE and we will solve the registration problem via geodesic shooting.

The formulation of CM is general and can be adapted to different cases. In this dissertation, we focus on the aforesaid registration problem. We thus present two new metrics: the first allows the addition of intensity only in a given mask, while the second makes the deduced flow match a given one (that mimics prior biological information).

It is interesting to notice that Holm, Trouvé and Younes already used the Constrained Metamorphosis name in Holm et al. [2009]. However, differently from here, in their work, this term was related to a simple addition of a fixed constraint to the geodesic equations. It was a mere theoretical proposition and they did not thoroughly analyse it with practical experiments. Instead, the model we propose here can be seen as a more generic generalisation of Metamorphosis and we decided to stick to this self-explanatory name. In addition, this work has been built to be versatile. Different user cases and applications could lead to other types of constraints.

In Section 3.1 we present the Constrained Metamorphosis framework theoretical part. Starting by defining the metrics and justifying their expression. Then, we give the formulation of a geodesic equation in the most general form and adapt it to some interesting cases with respect to the goal of this Thesis. To conclude the section, we detail the usage of the method using our Python library Demeter.
In Section 3.2, we illustrate the method through a toy-Example. Finally, in Section 3.3 we apply and test the method on real MRI data.

3.1 The Growing Constrained Metamorphosis framework

3.1.1 From Metamorphosis limitations for Pathological brain to prior utilisation

Before getting to the mathematical formulation, let's be more precise on how we intend to use the Constrained Metamorphosis framework for brain registration. Starting by analysing the limitation of a classical Metamorphosis.

In the previous chapter, we tested the Metamorphic registration method on a couple of MR images: a healthy brain template and a pathological image with a brain glioblastoma (GB). However, as shown in Figure 2.20, the results were disappointing. Specifically, the method failed to register the anatomical areas that had been displaced by the GB mass effect.

From a biological perspective, there are at least three factors that may explain this limitation:

- Tumour growth may have moved tissues away from each other to the point where they no longer overlap. As we use sum of squared differences (SSD) as the data attachment term, the algorithm does not attempt to match them.
- While the tumour develops, it reroutes nutrients that were previously destined for healthy cells. This creates attrition in the surrounding tissues, and neighbouring cells may die, resulting in the disappearance of tissue parts visible on the MRI scan at the macroscopic scale.
- The tumour not only grows and pushes the surrounding tissues but also infiltrates them, inducing intensity changes with only minimal displacement.

To correctly match or compensate for the apparition of such changes, Metamorphosis must create displacement and generate intensity changes at the same time. If in theory, that is exactly its purpose, as we saw in Chapter 3 conclusion, that the amount of the deformation *vs.* intensity changes were controlled by hyper-parameters. Thus, this trade-off was applied to the whole image when it would have been desirable to have different local hyper-parameter ratios on the image.

Here, we want to leverage the fact that anatomical images are very structured and embedded with many *a-priori* information, especially about organ spatial organisation. The main idea of this chapter is to find a way to guide registrations with this prior information. We developed it having in mind a growing structure (*i.e.*, the tumour) that lies within another one (*i.e.*, the brain) and therefore it might need adjustments for other applications.

Our goal is to help the registration to correctly locate the topological additions and at the same time to follow the direction of the mass effect.

We decided to select the two main structures of glioblastoma, namely the necrotic core only (ET and TC in Figure 1.2) and the whole tumour, including Oedemas infiltration (WT). One can note that in this case, the former mask is included in the later. We want to restrict intensity addition to the lesion only, assuming that healthy tissues must be matched by deformation only, thus LDDMM. We also aim to simulate the tumour growth via a very simple model. The exact details will be further described for each specific case presented in the Chapter but the general principle remains. We select the necrotic mask, initialise a small ball within, and make it grow into the mask. We choose to make the ball grow via an LDDMM registration, however, one could use a biophysical model instead. The LDDMM registration gives us a displacement field that we will later utilise to guide the registration of the whole brain.

Obviously, we assume that we can detect these *a-priori* within a given image, in our case the tumour. Segmentation masks are a convenient way to apply specific computation locally. They can be manually drawn or, for large data sets, it could be worth it to develop automatic methods. In this chapter, we use a deep-learning-based method, as recent improvements made automatic segmentation highly reliable [Liu et al., 2021; Luu and Park, 2022; Yuan, 2022]. One can also refer to Chapter 4 to read about our attempt to GB segmentation using TDA. In the following section, we will start by defining metrics used to ensure that the final registration will efficiently respect our priors. Then, we will see how these metrics affect the geodesics.

3.1.2 Defining some metrics

In this section, we aim to define a Riemannian metric on cancerous images, inspired by the one constructed from Metamorphosis in Section 2.4. This sub-section will focus on the construction of a norm (as a variation of Equation 2.92) that will allow us to take *priors* as described in previous section into account.

We aim at registering two brain images, where one at least contains a glioblastoma, with a Necrotic and Oedema part. Here, we aim to construct the metrics $N_{\mathscr{O}}(v, P, w)$ and $N_{\mathscr{W}}(z, Q, f)$. $N_{\mathscr{O}}(v, P, w)$ will *orient* v using a given temporal vector field w at locations given by a mask P of the Oedema. $N_{\mathscr{W}}(z, Q, f)$ will *weight* the intensity additions given by z by adding local constraints defined by a function f and a mask Q defining the necrosis location. Images infinitesimal updates will be given by the relation

$$\delta_t I_t = v_t \triangleright I_t + \mu Q_t z_t = -v_t \cdot \nabla I_t + \mu Q_t z_t, \quad \mu \in \mathbb{R}^+. \tag{3.1}$$

with $v \cdot \nabla I$ being the Euclidean scalar product between the field and the image gradient. Also, the Riemannian metric is

$$E_C(I, v) = \int_0^1 \|v_t\|_W^2 + \gamma N_{\mathscr{O}}(v_t, P_t, w_t) + N_{\mathscr{W}}(z_t, Q_t, f)dt. \tag{3.2}$$

We search for the minimal path such as $I_0 = S$, I_1 is found with Equation 3.1. In this section, we will always refer to the W space and norm (see Section 2.2.4.c) as we usually want to do multi-scale registration, however, one could just interchange it with the V space if one wants to do so, as it can be seen as a particular case of W.

3.1.2.a Weighted norm: $N_{\mathscr{W}}(z, Q, f)$

The norm presented here was introduced first in my conference paper "Weighted Metamorphosis for registration of images with different topologies" (François et al. [2022]). I present it here in a slightly more general form.

We saw in the previous chapter that even if Metamorphosis can lead to interesting registrations, the disentanglement between geometric and intensity changes is not unique and it highly depends on user-defined hyper-parameters. This makes the interpretation of the results hard, thus hampering its clinical usage. For instance, to align a healthy template to an image with a tumour, one would expect that the method adds intensities only to create new structures (*i.e.:* tumours) or to compensate for intensity changes due to the pathology (*i.e.:* oedema). All other structures should be correctly aligned solely by the deformations. However, depending on the hyper-parameters, the algorithm might decide to account for morphological differences (*i.e.:* mass effect of tumours) by changing the appearance rather than applying deformations. This limitation mainly comes from the fact that the additive intensity changes can theoretically be applied all over the image domain. However, in many clinical applications, one usually has prior knowledge about the position of the topological variations between a healthy image and a pathological one (*e.g.:* tumour and oedema position).

To prevent the addition of topological additions outside of a given mask M taking values in between 0 and 1 possibly rough. We define the following metric:

Definition 3.1 (Weighted semi-norm)**.** *First, let F and f be functions taking values in $[0,1]$ and mapping to $\mathbb{R}$, such that*

$$F(u) = uf(u). \tag{3.3}$$

In addition, F must be two times derivable and $f(u) \neq 0$ for every $u \in [0,1]$.

Let be $M : \Omega \mapsto [0,1]$ a mask. We define the generalised weighted norm on residuals such as :

$$\|z\|^2_{F(M)} = \left\|\sqrt{F(M)}z\right\|^2_{L^2} = \langle z, F(M)z\rangle_{L^2} = \langle z, Mf(M)z\rangle_{L^2} \tag{3.4}$$

As the reader may note, this definition is quite general and f remains to be defined. By experience, we found that choosing a very simple definition for f gives already good results. For example in François et al. [2022], the function f was set to be a constant $f(M) = \rho, \rho \in \mathbb{R}^+$. The assumption that the function f must be positive everywhere is technical and comes to avoid a division by zero in the proof of Theorem 3.1. The definition of Equation 3.3 has been chosen to avoid adding intensity outside of the mask by having the property that for all $x \in \Omega, M(x) = 0 \implies F(M(x)) = 0$. The other way around, however, is not true (*i.e.:* $F(M(x)) = 0 \not\implies M(x) = 0$).

As one can remark, $\| \bullet \|_{F(M)}$ is only a semi-norm, indeed $z = 0$ is not the only solution to the equation $\|z\|_{F(M)} = 0$. Having in mind the Metamorphic Riemannian metric construction from Section 2.4, one could wonder if having a semi-norm instead of a norm will suffice. In short, yes. As we will see in the incoming proof of Theorem 3.1, we care more about having a scalar product defined everywhere.

3.1.2.b Oriented norm: $N_{\mathscr{O}}(v, P, w)$

In this subsection, we aim at defining a metric that will constrain the estimated deformation to follow a given direction. One can take as an example the work of Risser et al. [2013]. The authors register lung CT scans and want to impose deformation with non-smooth discontinuous sliding conditions near the rib cage while keeping the invertibility of the deformation. They propose an adapted version of the log-Demons algorithm where a regularised term has been added to constrain the field orientation directly. We propose here a norm in the same fashion. In the following, we will assume that we have a known temporal vector field $w = (w_t)_{t\in[0,1]}$ having some interesting property that we want to recover at given locations given by a mask $P = (P_t)_{t\in[0,1]}$. In our case, as detailed in this section introduction, w would be a vector field modelling the cancer growth and P the mask of the growing tumour.

A DETOUR BY DROPPED IDEAS Before getting to the norm crafted by our means, I find it interesting to detail some ideas we explored and why we did not retain them. I hope that this short paragraph about negative results will give some insights into our final choice.

Our first direction of research was to minimise the angle θ in between $v(t, x)$ and $w(t, x)$ for a given time and location. We can work with the definition of the cosine

$$\cos\theta = \frac{v(t,x)\cdot w(t,x)}{\|v(t,x)\|_{L^2}\|w(t,x)\|_{L^2}} \tag{3.5}$$

integrating over the pixels and shifting the value of the function to zero when the vectors are aligned.

$$\begin{aligned} N_{\mathscr{O}\cos} &: V^2 \to [0,1] \\ N_{\mathscr{O}\cos}(v_t, w_t) &= \frac{1}{2} - \frac{1}{2|\Omega|}\int_{x\in\Omega} \frac{v_t \cdot w_t}{\|v_t\|_2\|w_t\|_2} dx. \end{aligned}$$

If this form seems to be a good candidate, it will prevent us to have a clean form of the vector field formula in the geodesic equations. Indeed, looking at the proof of theorem 3.1 at equations 3.18 it would be hard to set v_t apart as we must keep the term $K^{-1}v_t$ from the derivative of the V norm.

Our second thought, still keeping the ambition to restrain the angle only, was to utilise outward pointing unitary vectors from the mask P at every time. A big advantage of this solution would allow the bypass of the vector field w usage. It was inspired by the Green theorem and a very practical lemma for RKHS theory.

Let $\vec{n}_x$ be a unitary vector field pointing toward the desired direction. For example, taking the case presented for the Weighted Norm (Section 3.1.2.a) where we have a growing mask, $x \in \Omega \mapsto \vec{n}_x(t)$ could be the vector field pointing toward the exterior of the mask P at its border and zero elsewhere at a given time t. In practice, the border locations x can be retrieved using the spatial gradient on a thresholded mask. An interesting lemma states that we can make the scalar product V from the Euclidean one.

Lemma 3.1. *Let $\vec{n}$ and v two vectors field mapping $\Omega \mapsto \mathbb{R}^d$, it exists a kernel K and a scalar product V such that*

$$\vec{n}_x \cdot v(x) = \langle K(x, \bullet)\vec{n}_x,\ v(x)\rangle_V \tag{3.6}$$

Proof. As $\vec{n}_x \cdot v(x)$ is a continuous linear form for each x in Ω: $\psi_x : v(x) \mapsto \vec{n}_x \cdot v(x) \in \mathbb{R}$, where $\psi_x \in V^*$ is the dual of V. Using the Riesz theorem we have $\forall \psi \in V^*, \exists \hat{\psi} \in V$ such as $\psi(v) = \langle \hat{\psi}, v \rangle_V$

$$\langle K(x, \bullet)\vec{n}_x,\ v(x) \rangle_V = \langle K(x, \bullet)K^{-1} \star \vec{n}_x,\ v(x) \rangle_{L^2} = \vec{n}_x \cdot v(x) \tag{3.7}$$

□

This seems interesting at first glance. For example, using lemma 3.1 we can rewrite the projection of the difference of difference two vector fields v and w on $\vec{n}_x$ such as

$$\begin{aligned} \int_{x\in\Omega} \langle v_t(x) - w_t(x), \vec{n}_x \rangle_2 dx &= \int_{x\in\Omega} \langle K(x, \bullet)\vec{n}_x, v_t - w_t \rangle_V dx \\ &= \left\langle \int_{x\in\Omega} K(x, \bullet)\vec{n}_x dx, v_t - w_t \right\rangle_V . \end{aligned} \tag{3.8}$$

The integral inside the brackets is made of known objects and therefore can be preprocessed. However, for any $v \in V$, one can remark that the differential in the direction v does not depends of v:

$$D_v \left[\int_{x\in\Omega} \langle v_t(x), \vec{n}_x \rangle_2 dx \right] = \int_{x\in\Omega} \vec{n}_x dx. \tag{3.9}$$

The fact that neither v nor w were appearing in this differential, thus won't be part of the geodesics, was bothering us. We did not try to implement it due to time constraints. Instead, we chose a much simpler and quite efficient solution, that will be described in the following.

PROPOSED ORIENTED NORM Let's assume we have a temporal vector field w_t as a prior for our matching problem, in which we are confident that it matches or at least goes in the right direction near some voxels x, let's say given by a mask P_t. We need to set a cost being small when the researched field v_t is close to w_t near some voxels x given by a temporal mask $P_t(x)$. (more details in section 2.1)

Thereby, we minimise the differences of v_t by w_t inside the norm W (see Definition 2.15)

$$N_{\mathscr{O}}(v_t) = \|P_t(v_t - w_t)\|_W^2 \tag{3.10}$$

It ensures v_t follows the field given by w_t without displaying rough turns at the edges of P_t. As we will see in the future theorem proof, it is essential to use a W norm here rather than another (*e.g.:* At first glance a L^2 norm could have been a good candidate).

3.1.3 The Constrain Metamorphosis Geodesics

3.1.3.a Theorem and proof

Theorem 3.1 (Constrained Metamorphosis)**.**

Let be $(P_t)_{t\in[0,1]}, (Q_t)_{t\in[0,1]} \in L^2([0,1], \mathscr{C}(\Omega, [0,1]))$ *two smooth temporal masks and* $(w_t)_{t\in[0,1]} \in L^2([0,1], V)$ *a vector field. We consider* P, Q *and* w *as known constraints. The minimal path in between two images* $S = I_0$ *and* $T = I_1$ *with the evolution :*

$$\partial_t I_t = -v_t \cdot \nabla I_t + \mu Q_t z_t, \quad \mu \in \mathbb{R}^+ \tag{3.11}$$

and the metric

$$E_C(I, v) = \int_0^1 \|v_t\|_W^2 + \gamma \|P_t(v_t - w_t)\|_W^2 + \left\| \sqrt{F(Q_t)} z_t \right\|_{L^2}^2 dt \tag{3.12}$$

is given by the following set of geodesic equations :

$$\begin{cases} (L + \gamma P_t L P_t) v_t &= -\frac{1}{\mu} K \star (f(Q_t) z_t \nabla I_t) + P_t L P_t w_t \\ \partial_t z_t &= -f(Q_t)^{-1} \left[\nabla \cdot (f(Q_t) z_t v_t) + \partial_t f(Q_t) \right] \\ \partial_t I_t &= -v_t \cdot \nabla I_t + \mu Q_t z_t \end{cases} \tag{3.13}$$

Proof. We will find the geodesic equations 3.21 from the image evolution 3.11 and the cost functional E_C by doing the variations according to images I and vector fields v successively for solving their Euler-Lagrange Equations. To do so we define two Lagrangian:

$$\mathcal{L}_I(t, I, \dot{I}) = E_C(\bullet, v) \text{ and } \mathcal{L}_v(t, v, \dot{v}) = E_C(I, \bullet) \tag{3.14}$$

We start by computing the fields variations h with respect to v by finding $\mathcal{L}_v$'s gradient. As $\|v\|_W^2 = \langle Lv, v \rangle_{L^2} = \int_{\sigma \in \mathfrak{S}} \|v\|_{V_\sigma}^2 d\sigma = \sum_{\sigma \in \mathfrak{S}} \left\langle K_\sigma^{-1} v, v \right\rangle_{L^2}$, its differential is $D_v \|v\|_W^2 .h = \sum_{\sigma \in \mathfrak{S}} \left\langle K_\sigma^{-1} v, h \right\rangle_{L^2}$. (see Section 2.2.4.c and 2.5.2.b)

For the generalised weighted norm on residual, we first need to note that we can re-write Equation 3.11 such as

$$\partial_t I_t = v_t \cdot \nabla I_t - \mu \frac{F(Q_t)}{f(Q_t)} z_t \tag{3.15}$$

leading to

$$\|z\|_{F(Q)}^2 = \left\langle z, \frac{1}{\mu} f(Q)(\partial_t I - v \cdot \nabla I) \right\rangle_{L^2} \tag{3.16}$$

thus

$$D_v \mathcal{L}_v \cdot h = \int_0^1 \left\langle L v_t + \gamma P_t L (P_t(v_t - w_t)) + \frac{1}{\mu} f(Q_t) z_t \nabla I_t, h_t \right\rangle_{L^2} dt. \tag{3.17}$$

One can remark that $D_{\dot{v}} \mathcal{L}_v = 0$ so its Euler-Lagrange equations are :

$$\begin{aligned} & \nabla_v \mathcal{L}_v - \partial_t \nabla_{\dot{v}} \mathcal{L}_v = 0 \\ \Leftrightarrow \quad & L v_t + \gamma P_t L (P_t(v_t - w_t)) = -\frac{1}{\mu} f(Q_t) z_t \nabla I_t \\ \Leftrightarrow \quad & (L + \gamma P_t L P_t) v_t - P_t L P_t w_t = -\frac{1}{\mu} f(Q_t) z_t \nabla I_t \\ \Leftrightarrow \quad & (L + \gamma P_t L P_t) v_t = -\frac{1}{\mu} f(Q_t) z_t \nabla I_t + P_t L P_t w_t \end{aligned} \tag{3.18}$$

using the linearity of the operator L and $K^{-1} = L$ being its inverse operator.

We can now compute the variations of $\mathcal{L}_I$ with respect to I and $\dot{I}$. One can check that the differential of $\|\bullet\|_V^2$ by I and $\dot{I}$ is zero and using Equation 3.16 find:

$$\begin{aligned} D_I\mathcal{L}_I \cdot h &= 2\int_0^1 \left\langle z_t, \frac{1}{\mu} f(Q_t)\nabla h_t \cdot v_t \right\rangle_{L^2} dt \quad \text{by integration by part} \\ &= \int_0^1 \left\langle -\frac{2}{\mu}\nabla \cdot (f(Q_t)z_t v_t), h_t \right\rangle_{L^2} dt \end{aligned} \tag{3.19}$$

Similarly $D_{\dot{I}}\mathcal{L}_I \cdot h = \int_0^1 \langle \frac{2}{\mu} f(Q_t)z_t, h_t\rangle_{L^2} dt$. We can now tackle the Euler-Lagrange equations:

$$\begin{aligned} & \nabla_I \mathcal{L}_I - \partial_t \nabla_{\dot{I}} \mathcal{L}_I = 0 \\ \Leftrightarrow \quad & -\frac{2}{\mu}\nabla \cdot (f(Q_t)z_t v_t) - \partial_t \left[\frac{2}{\mu} f(Q_t) z_t\right] = 0 \\ \Leftrightarrow \quad & f(Q_t)\partial_t z_t = -\nabla \cdot (f(Q_t)z_t v_t) + z_t \partial_t f(Q_t) \\ \Leftrightarrow \quad & \partial_t z_t = -\frac{1}{f(Q_t)} \left[\nabla \cdot (f(Q_t)z_t v_t) + z_t \partial_t f(Q_t)\right] \end{aligned} \tag{3.20}$$

Finally, the combination of Equations 3.11, 3.18 and 3.20 describe the searched geodesic variations. □

With this theorem, we have shown that the Riemannian metric E_C admit a minimising path formulation for the Euler-Lagrange equations. Thus, among the paths matching two images, we can choose the minimising one by solving the exact matching problem under constraints of Equation 3.2. As one can note, the proof and reasoning are very similar to the one we have presented in Section 2.4.

As one can notice, it this version of the vector field evolution (Equations 3.18) one apply the differential operator P_tL. It is problematic because, as detailed in Section 2.2.2.c, we define L through its inverse K, a Gaussian smoothing kernel or a combination of those. It results that we don't know the explicit expression of L and are not able to compute P_tLP_t nor inverse P_tL. To overcome this difficulty, we will approximate $P_tLP_t \approx LP_t$. This approximation is theoretically challenging because:

- If P_t is a binary mask, one might argue that if
  - in all the open domain where $P_t(x) = 1$, one has $P_tLP_tv = LPtv = Lv$
  - in all the open domain where $Pt(x) = 0$, one has $P_tLP_tv = LP_tv = 0$
  - at the boundary, the mask is discontinuous, which implies that LP_tv is not defined in the classical sense, and becomes a distribution (with enough smoothness assumption on v). Hence P_tLP_tv does not equal LP_tv at this boundary.
- Conversely, if P_t is a smooth mask taking values between 0 and 1, then LP_tv is well defined even in the classical sense. However, it remains that P_tLP_tv is not equal to LP_tv for every x such that $0 < P_t(x) < 1$.

Nevertheless, this approximation holds reasonably well over the majority of the domain, at least in the examples tested in this document. Thus, the geodesic equations 3.13 of Theorem 3.1 can be re-written as follows:

Proposition 3.1. *If we assume that $P_t L P_t = L P_t$ then the geodesic equations 3.13 can be re-written such that:*

$$\left\{\begin{array}{rl} v_t & = -\frac{1}{\mu(1+\gamma P_t)} K \star \left(f(Q_t) z_t \nabla I_t\right) + \frac{\gamma P_t}{1+\gamma P_t} w_t \\ \partial_t z_t & = -f(Q_t)^{-1} \left[\nabla \cdot \left(f(Q_t) z_t v_t\right) + \partial_t f(Q_t)\right] \\ \partial_t I_t & = -v_t \cdot \nabla I_t + \mu Q_t z_t \end{array}\right. \tag{3.21}$$

Proof. This set of equations can be deduced by rewriting the Equivalences 3.18 within Theorem 3.1:

$$\begin{aligned} & \nabla_v \mathcal{L}_v - \partial_t \nabla_{\dot{v}} \mathcal{L}_v = 0 \\ \Leftrightarrow \quad & L v_t + \gamma P_t L (P_t (v_t - w_t)) = -\frac{1}{\mu} f(Q_t) z_t \nabla I_t \\ \Leftrightarrow \quad & L v_t + \gamma L (P_t (v_t - w_t)) = -\frac{1}{\mu} f(Q_t) z_t \nabla I_t \\ \Leftrightarrow \quad & (1+\gamma P_t) v_t = -\frac{1}{\mu} K \star \left(f(Q_t) z_t \nabla I_t\right) + P_t L P_t w_t \end{aligned} \tag{3.22}$$

by linearity of the operator L and with $K = L^{-1}$, with L being a differential operator. □

In the rest of the document, we will always make this assumption and refer to Equations 3.21 when writing Constrained Metamorphosis Geodesic. Accordingly, in Section 3.1.3.b we will detail some particular cases using Proposition 3.1.

In practice we solve the problem in a non-exact minimisation approach, using geodesic shooting as we previously covered for Metamorphosis in Section 2.4.1. We minimise the cost:

$$H(z_0) = \mathrm{Sim}(I_1, T) + \lambda \left[\|v_0\|_W^2 + \gamma \|P_0(v_0 - w_0)\|_W^2 + \left\| \sqrt{F(Q_0)} z_0 \right\|_{L^2}^2 \right] \tag{3.23}$$

In theory, the Sim function could be chosen and changed without changing the geodesics nature. However, it will have an impact on convergence speed, and potentially on the direction of approach of the optimal path (*e.g:* approaching the target by the 'top' rather than the 'left'). If one problem admits more than one minimising geodesic to match two images, it could result in different matching strategies. By default, I use the L^2 norm on images which is the Ssd. Note that we can set $\lambda = 0$ and still retrieve good registration as long as we follow the geodesics. Before getting to the practical implementation of the CM framework, let us list more convenient particular cases of Theorem 3.1 taking into account Proposition 3.1.

3.1.3.b Particular cases and remarks

We will now give some particular cases of Theorem 3.1 taking into account Proposition 3.1 that can be used depending on the needs. We set as a convention that when we talk about a method (*e.g.:* Weighted Metamorphosis) we will set f constant (*e.g:* $f(M) = \rho$) unless we specify otherwise.

WEIGHTED METAMORPHOSIS (WM) is the case when one does not try to orient the field and is obtained by setting $\gamma = 0$. If we choose f in its simplest form: $f(M) = \rho$ (*i.e:* $F(M) = \rho M$) then we obtain the Weighted Metamorphosis we described in François et al. [2022]. Indeed when $\gamma = 0$ and setting f as a constant function applying the Theorem 3.1 one gets the geodesic set of equations:

$$\begin{cases} v_t &= -\frac{\rho}{\mu} K \star (z_t \nabla I_t) \\ \partial_t z_t &= -\nabla \cdot (z_t v_t) \\ \partial_t I_t &= -\nabla I_t \cdot v_t + \mu Q_t z_t \end{cases} \tag{3.24}$$

Choosing f as a constant function makes the formula very close to the Metamorphosis case as it removes its time derivatives from the system, and therefore it is easier to understand and implement.

One could find making f more complex interesting by, for example, enhancing the flow at mask borders by defining:

$$F(M) = M(1 + \rho|\nabla M|) \tag{3.25}$$

This slightly more advanced function allows giving a factor of multiplication at the border of the mask. If $\rho > 0$ the vector field will be stronger at its edge.

Let's check if the function induces well-defined geodesics. As $f(M_t) = 1 + \rho|\nabla M_t|$ we have $f(M(t,x)) = 1$ when the mask is constant on a compact (*i.e.:* $\exists \varepsilon > 0, \forall x \in \mathscr{B}(x, \varepsilon)$ such as $\nabla M_t(x) = 0$), and $\nabla M_t < 1$ so $F(M_t) < 1 + \rho$ therefore f is well defined.

There is more work for computing the time derivative of f. As

$$\begin{aligned} \partial_t \left|\nabla M_t(x)\right| = \partial_t \left(\sum_{i=1}^{d} \left(\partial_{x_i} M(t,x)\right)^2 \right)^{\frac{1}{2}} &= \frac{\sum_{i=1}^{d} \partial_{x_i} M(t,x) \partial_t \partial_{x_i} M(t,x)}{\left|\nabla M_t(x)\right|} \\ &= \frac{\nabla M(t,x) \nabla \partial_t M(t,x)}{\left|\nabla M_t(x)\right|} \end{aligned}$$

we have $\partial_t f = \rho \frac{\nabla M(t,x) \nabla \partial_t M(t,x)}{\left|\nabla M_t(x)\right|}$. Note that $\partial_t M_t = \nabla M_t \cdot w_t$ might already been computed during the making of M_t and so we can approximate $\partial_t M_t = \frac{1}{\delta t}(M_{t+\delta t} - M_t)$ Finally, we obtain the geodesic equations :

$$\begin{cases} v_t &= -K \star (\frac{1}{\mu}(1 + \rho|\nabla M_t|) z_t \nabla M_t) \\ \partial_t z_t &= -\frac{1}{1+\rho|\nabla M_t|} \left[\nabla \cdot ((1 + \rho|\nabla M_t|) z_t v_t) + \frac{\rho}{\mu} \frac{\nabla M(t,x) \nabla \partial_t M(t,x)}{\left|\nabla M_t(x)\right|} z_t \right] \\ \partial_t I_t &= -\nabla I_t \cdot v_t + \mu M_t z_t \end{cases} \tag{3.26}$$

We have an example of function f for which geodesics exists. It demonstrates how someone could introduce a biological model or tweak the framework further for their specific applications. However, in practice, this system requires choosing the integration carefully. Moreover, the complexity of the formula will probably make the computation slow.

ORIENTED METAMORPHOSIS (OM) are obtained when the function $f(M) = \rho$ and the mask Q_t is set to one everywhere. Namely, it is the case of classical Metamorphosis with the oriented norm. Then the geodesics are

$$\begin{cases} v_t &= -\frac{\rho}{\mu(1+\gamma P_t)} K \star (z_t \nabla I_t) + \frac{\gamma P_t}{1+\gamma P_t} w_t \\ \partial_t z_t &= -\nabla \cdot (z_t v_t) \\ \partial_t I_t &= -\nabla I_t \cdot v_t + \mu z_t \end{cases} \tag{3.27}$$

One can observe that the addition of the oriented norm impacts the vector field geodesic equation only. With this formulation we can easily balance w's influence on v, indeed $\frac{\gamma}{1+\gamma} \underset{\gamma \to +\infty}{\longrightarrow} 1$, converge quickly and stays between 0 and 1 for $\gamma \geq 0$. If one set $\gamma \approx p = \mu/\rho$ then both members of v have the same importance.

COMBINATION WITH A BIOPHYSICAL MODEL Gooya et al. [2011b, 2012] implemented a joint model for registering and Segmenting glioblastomas through a biophysical model using a PDE system. Assuming it can be derived, the Constrained Metamorphosis framework allows adding such constraints within the geodesic equations via a combination of the norms N and $N_{\not W}$. We do not give the details here and refer to future publications.

CONSTRAINED METAMORPHOSIS PARAMETER TUNING HINTS: Lets assume than $f(M) = \rho$, using Theorem 3.1 we have the geodesics equations:

$$\begin{cases} v_t &= -\frac{\rho}{\mu(1+\gamma P_t)} K \star (z_t \nabla I_t) + \frac{\gamma P_t}{1+\gamma P_t} w_t \\ \partial_t z_t &= -\nabla \cdot (z_t v_t) \\ \partial_t I_t &= -\nabla I_t \cdot v_t + \mu Q_t z_t \end{cases} \tag{3.28}$$

One can remark that there are at least three constants that are delicate to tune: μ, ρ and γ. By observing the systems of Equations 3.24 and 3.28 first line, we can deduce that what matters is the quotient $p = \frac{\rho}{\mu}$ which help scale the vector field. Once μ set, by default to 1, it is easier to tune p than ρ. Then having chosen p in CM we want to set the amount of influence of the field w_t on v_t, which is controlled by γ. As both terms in v_t are divided by $(1 + \gamma P_t)$, it is more convenient to set $\gamma = k \times \rho$ and I advice to choose $k \in [0, 1]$

3.1.4 How to use Constrained Metamorphosis with Demeter

3.1.4.a An implementation easy to use

Just as for Metamorphosis, the code is divided into two classes. A first one, `Constrained_meta_path` inheriting from `Geodesic_integrator` implementing the integration of the geodesic equations 3.28. It takes as parameter for initialisation a class inheriting from `Residual_norm_function`, which implements the squared norm $\|z\|^2_{F(M)}$ (see Equation 3.4). By default we pass the class `Residual_norm_identity` which sets $f = \text{Id}$. A second one, `Constrained_Optim` inheriting from `Optimize_geodesicShooting` implements the geodesic shooting algorithm, minimising Equation 3.23. It takes as argument a class inheriting from `Data_cost` and implements a similarity function in between I_1 and T. In this thesis, I didn't have time to study properly the effects of different data metrics, but having it in mind, I made it easy to change Sim in the code. The default one is `Ssd`. The code displayed in Code 3.1 is a wrapper function API executing the optimisation.

Source Code 3.1: Constrained Metamorphosis usage

```
1 def constrained_metamorphosis(source,target,residual,
2         rf_method,mu,rho,mask_w,
3         mp_orienting,gamma,mask_o,
4         sigma,cost_cst,sharp,
5         n_iter,grad_coef):
```

```
        mask = mp_orienting.image_stock.to(source.device)
        orienting_field = mp_orienting.field_stock.to(source.device)

        if rf_method == 'identity':
                rf_method = Residual_norm_identity(mask,mu,rho)
        elif rf_method == 'borderBoost':
                rf_method = Residual_norm_borderBoost(mask,mu,rho)
        else:
                raise ValueError(f"rf_method must be 'identity' or 'borderBoost'")

        residual.requires_grad = True

        mp_constr = Constrained_meta_path(orienting_mask=mask,
                orienting_field=orienting_field,
                residual_function=rf_method,
                mu=mu,rho=rho,gamma=gamma,
                sigma_v=sigma,
                sharp=sharp,
                # no n_step, it is defined from mask.shape[0]
        )
        mr_constr = Constrained_Optim(source,target,mp_constr,
                cost_cst=cost_cst,
                optimizer_method='adadelta')
        mr_constr.forward(residual,n_iter=n_iter,grad_coef=grad_coef)
        return mr_constr
```

3.1.4.b Easy to tweak: Reversed Constrained Metamorphosis model definition and implementation

In this section, we aim to demonstrate how we can use both the theoretical framework and our implementation to test different ideas. As a matter of example, we aim to change the direction of the registration, deforming the image with lesions to match healthy areas. It is a common practice as it works well for retrieving plausible deformations Han et al. [2020a]; Maillard et al. [2022]. For Metamorphosis, however, it is less clear how to retrieve the pixel intensity difference.

Let's first define the model step by step mathematically. I deliberately swap the images to still have T the image with the tumour and S the healthy one even if now the source image is T and the target S. N will be the necrosis segmentation and Q the Oedema one. The objective here is to use WM for registration using the mask Q which will follow along with the deformation and make vanish the necrosis mask. This is obtained by re-expressing the Equation 3.12 as

$$E_C(I, v) = \int_0^1 \|v_t\|_W^2 + \left\| \sqrt{Q \circ (\varphi^{\mathbf{v}})^{-1}} z_t \right\|_{L^2}^2 dt \tag{3.29}$$

with $\varphi_t^{\mathbf{v}}$ being the deformation at a given time t integrated with Equation 2.43. Using

Theorem and Proposition 3.1 we obtain the geodesic equations to integrate over:

$$\begin{cases} v_t &= -\frac{\rho}{\mu} K \star (z_t \nabla I_t) \\ \partial_t z_t &= -\nabla \cdot (z_t v_t) \\ \partial_t I_t &= -\nabla I_t \cdot v_t + \mu (Q \circ (\varphi^{\mathbf{v}})^{-1}) z_t \end{cases} \tag{3.30}$$

Finally and as seen before, we solve the problem with geodesic shooting minimising the cost:

$$J_r(z_0) = \|I_1 - S\|_{L^2}^2 + \kappa \left\|N \circ \Phi^{-1}\right\|_{L^2}^2 + \lambda \left[\|v_0\|_W^2 + \rho \|z_0\|_{Q \circ \Phi_1^{-1}}^2\right]. \tag{3.31}$$

I_1 and Φ being deduced from the PDE system 3.30 and with $\kappa, \lambda, \rho \in \mathbb{R}^+$. We added the term $\kappa \left\|N \circ \Phi^{-1}\right\|_{L^2}^2$ to the data cost with the idea to help the algorithm to make vanish the necrosis by deformation. Note that following the terminology introduced in 3.1.3.b, we perform a kind of Weighted Metamorphosis as we do not constrain directly the v.

With our implementation of the Constrained Metamorphic framework, it is quite easy to test model variations. In this example (Code 3.2) we have set a new data cost and custom residual norm. Starting with the data cost, we write the class `Reduce_field_ssd` (line 5) that inherits from the abstract class `DataCost` and we combine the default SSD cost with the `Ssd` class. Then we cannot use the default `Residual_norm_function` as we need to get all the deformations of oedema masks while we compute the vectors fields v. At this point, we need to implement Equations 3.3 and 3.15 within the method `F_div_f` (line 28) which by default returns mu times the mask. At each integration steps it will compute the intermediary deformation of the mask before intensity additions. Finally, the function `execute_reduce` performs the whole optimisation.

Source Code 3.2: Example of a custom Constrained Metamorphosis

```
1  import metamorphosis as mt
2
3  class Reduce_field_ssd(mt.DataCost):
4          """ norm || N \circ Phi^{-1} ||^2"""
5          def __init__(self,target,mask_to_reduce,kappa = 1):
6                  super(Reduce_field_ssd, self).__init__()
7                  self.ssd = mt.Ssd(target)
8                  self.mask_to_reduce = mask_to_reduce
9                  self.mult = kappa/prod(target.shape[2:])
10
11         def __call__(self):
12                 mask_deform = tb.imgDeform(self.mask_to_reduce.cpu(),
13                 self.optimizer.mp.get_deformator(),
14                 dx_convention='pixel')
15                 return self.ssd() + self.mult * (mask_deform**2).sum()
16
17 class Residual_identity_picker(mt.Residual_norm_function):
18         """ norm ||\sqrt{ Q \circ \diffi} z_t||^2"""
19         def __init__(self,seg_necrosis,n_step,mu=1,rho=None):
20                 super(Residual_identity_picker, self).__init__(seg_necrosis,mu,rho)
21                 self.temp_int = vff.FieldIntegrator(method='temporal',save=False)
22                 self.n_step = n_step
23
```

```
        def f(self,t): # mandatory: f is abstract in Residual_norm function
                return self.rho

        def F_div_f(self,t):
                """ perform the operation F(M)/f(M) """
                deform_mask = self.mask.clone()
                if t > 0:
                        field = self.geo_int.field_stock[:t] / self.geo_int.n_step
                        deformator = self.temp_int(field,forward=False)
                        deform_mask = tb.imgDeform(self.mask,deformator,
                        dx_convention='pixel')

                return deform_mask * self.mu

        def dt_F(self,t):
                return 0

        def execute_reduce(residual,lamb,mu,rho,kappa,sigma,n_step,n_iter):
                residual.requires_grad = True

                reduce_dTerm = Reduce_field_ssd(target,mask_necrosis,kappa)

                rf = Residual_identity_picker(mask_oedema,n_step,mu,rho)
                mp_reduce = mt.Constrained_meta_path(residual_function=rf,
                n_step=n_step,
                sigma_v=sigma,
                sharp=True
                )
                mr_reduce = mt.Constrained_Optim(source,target,mp_reduce,
                cost_cst=lamb,
                data_term= reduce_dTerm
                )
                mr_reduce.forward(residual,n_iter=n_iter,grad_coef=1)

                return mr_reduce
```

Results analysis obtained with this model can be found in Section 3.2.4

3.2 Toy-Examples: Construction & Experiments

Before dealing with real data, namely brain 3D MRI scans, we will demonstrate our methods on 2D toy-Examples. It was interesting for us to take some time to develop good toy examples for our registration tasks, in order to better control the characteristics of source and target images, to reduce computational time (fewer pixels) while debugging and to ease visualisation and interpretation of results. Simply taking MRI slices would have not been practical due to the 3D nature of the displacement and to the fact that images are highly complex. Using a controlled environment allowed us to better understand the proposed method and the role of its hyper-parameters.

3.2.1 Weighted Metamorphosis - The growing mask

In this section, we will present our work published in François et al. [2022] where we introduced Weighted Metamorphosis for the first time. In the next section, we will present more refined toy examples. However, the study presented in this section justifies the need of having a temporal growing mask. As one can retrieve from Section 3.1.2.a and Theorem and Proposition 3.1 we minimise the cost:

$$J_{\mathrm{WM}}(z_0) = \|I_1 - T\|_{L_2}^2 + \lambda \left[\|v_0\|_V^2 + \rho \|z_0\|_{M_0}^2 \right], \quad \lambda \in \mathbb{R}^+, I_0 = S \tag{3.32}$$

where $\|z_0\|_{M_0}^2$ is the identity weighted norm (*i.e.:* $\|z_t\|_{M_t}^2 = \langle z_t, M_t z_t \rangle_{L^2}$) and I_1 is obtained by integrating the geodesic equations 3.24.

The definition of the weight function $M_t : [0,1] \times \Omega \to [0,1]$ is quite generic and could be used to register any kind of topological/appearance differences. Here, we restrict to brain tumour images and propose to use an evolving segmentation mask as a weight function. We assume that we already have the binary segmentation mask B of the tumour (comprising both oedema and necrosis) in the pathological image and that healthy and pathological images are rigidly registered so that B can be rigidly moved onto the healthy image. Our goal is to obtain an evolving mask M_t that somehow mimics the tumour growth in the healthy image starting from a smoothed small ball in the centre of the tumour (M_0) and smoothly expanding it towards B. We generate M_t by computing the LDDMM registration between M_0 and B. Please note that here one could use an actual biophysical model (Gooya et al. [2011a]; Scheufele and et al. [2019]) instead of the proposed simplistic approximation based on LDDMM. However, it would require prior knowledge, correct initialisation and more than one imaging modality. The main idea is to smoothly and slowly regularise the transformation so that the algorithm first modifies the appearance only in a small portion of the image, trying to align the surrounding structure only with deformations. In this way, the algorithm tries to align all structures with shape changes adding/removing intensity only when necessary. This should prevent the algorithm from changing the appearance instead of applying deformations (*i.e.:* better disentanglement) and avoid wrong overlapping between new structures (e.g. tumour) and healthy ones.

Figure 3.1 presents qualitative results. We can observe the differences in the geodesic image evolution for LDDMM, Metamorphosis and Weighted Metamorphosis with a constant and evolving mask. First, LDDMM cannot correctly align all grey ovals and Metamorphosis results in an image very similar to the target. However, most of the differences are accounted for with intensity changes rather than deformations. By contrast, when using the proposed evolving mask (fourth row), the algorithm initially adds a small quantity of intensity in the middle of the image and then produces a deformation that enlarges it and correctly pushes away the four grey ovals. In the third row, a constant mask ($M_t = M_1, \forall t \in [0,1]$) is applied. One can observe that, in this case, the bottom and left ovals overlap with the created central triangle and therefore pure deformations cannot correctly match both triangle and ovals.

In all methods, the registration was done with the same field smoothness regularisation σ and integration steps. Please note that the four grey ovals at the border are not correctly matched with LDDMM and, to a lesser extent, also with our method. Two factors can explain this. First, the L2-norm data term since these shapes do not overlap between the initial source and target images and therefore the optimiser cannot match them. Secondly, the construction of the Figure 3.1 template was made using Inkscape and as I found out later the images could not be matched by a single diffeomorphic deformation. Indeed the

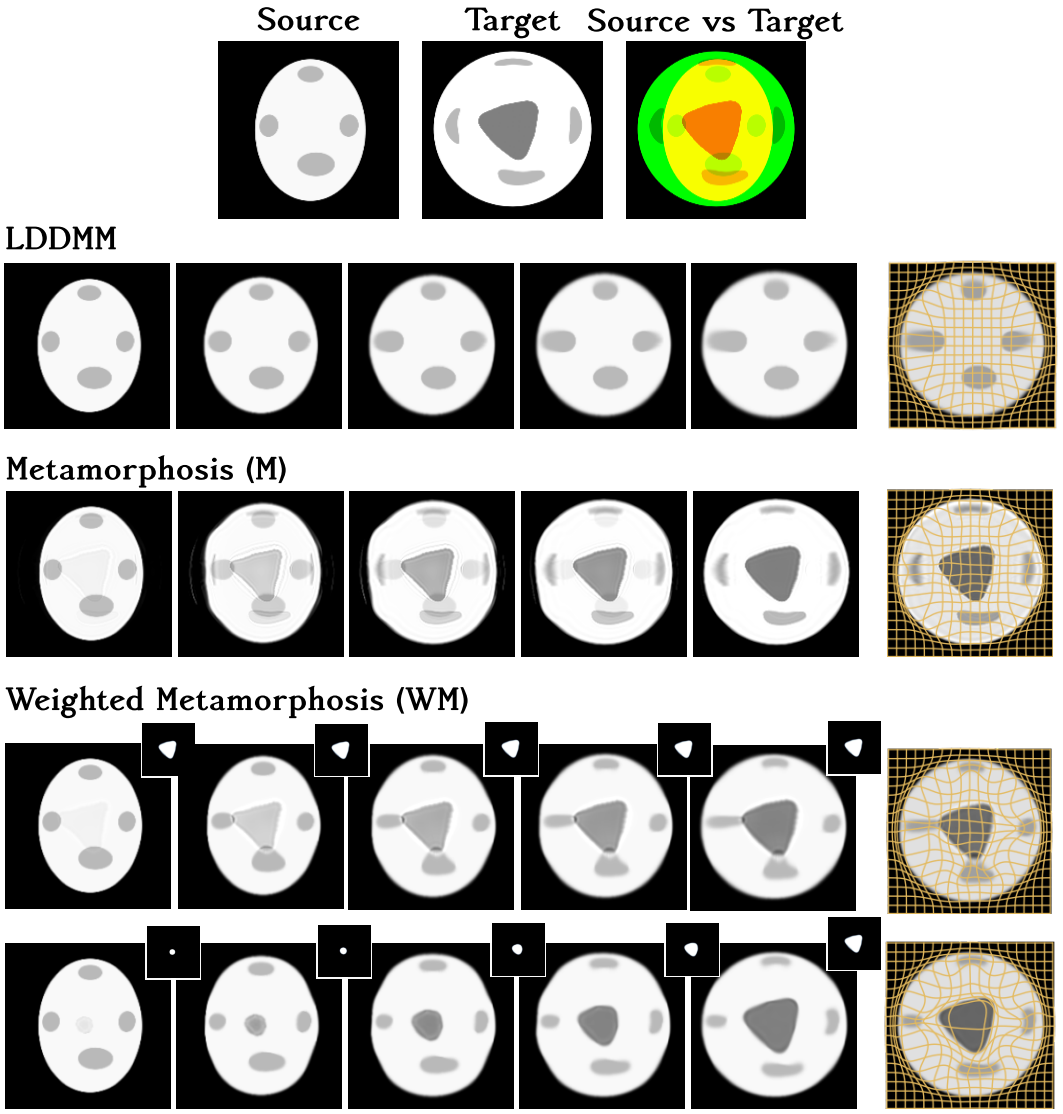

Figure 3.1: Comparison between LDDMM, Metamorphosis and WM. Image registration toy example. Differently from the Source image (S), the Target image (T) has a big central triangle that has grown "pushing" the surroundings ovals. Note that the bottom and left ovals in S overlap with the triangle in T. The two last rows show our method using a constant and time-evolving mask. The used mask is displayed in the top right corner of each image.

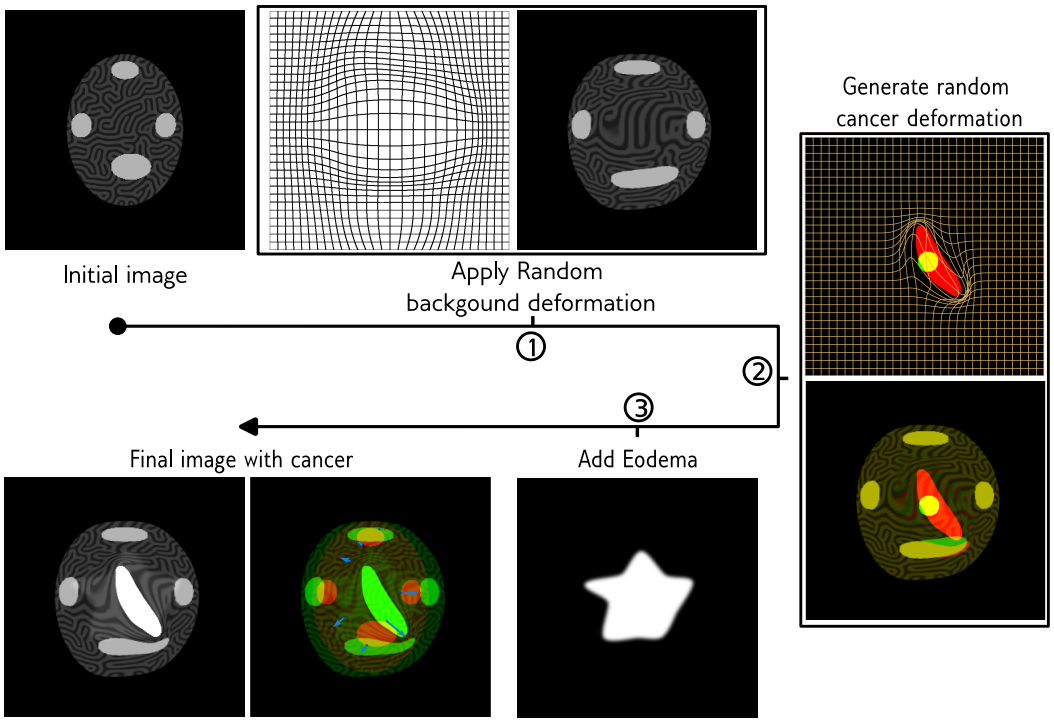

Figure 3.2: Construction of a mimicking Glioblastoma toy-Example, see Figure 1 for a color code explanation.

background deformation of the oval on the circle and the displaced organ were not really compatible. In more technical terms, the target was not in the deformation orbit of the image, this was corrected in the next section. The takeaway message of this section was the need for a growing mask to help generate the 'mass effect' deformation.

3.2.2 The construction of a refined Toy-Example

Shortly after the paper presented in the previous section submission (*i.e.:* François et al. [2022]), I understood why for some couple of brain MRIs the algorithm was failing. To do so I needed the intuition on why it was working on the toy example of Section 3.2.1. By letting the intensity appear only in the middle of the 'brain', we create an intensity variation and thus an edge. Recalling that the field is obtained from an equation such as $v_t = K \star (z_t \nabla I_t)$, a potential field pushing away other structures will arise. However, on real data, if the image is textured at the tumour location, the edges could be badly oriented thus preventing the expected mass effect push. Besides this toy example mimic only the necrosis growth, forgetting about the infiltrative intensity changes. The need for a more refined toy example was pressing.

The previous toy example construction was made using the drawing vector graphic tool Inkscape. This time the brain background will be very textured, presenting both registration problems of necrosis and oedema apparition. The construction process is illustrated in Figure 3.2. It is decomposed into three phases :

STEP 1: We apply to the initial image a deformation obtained from a bspline field obtained from a randomly chosen control point.

STEP 2: From a small ball, we generate a random deformation mimicking a necrosis growth using a bspline field again with random control points. As explained in Code 3.3, we first generate a random smooth field everywhere on Ω, then restrain this field around the necrosis smoothed edges. We apply the deformation deduced from this field to an image and repeat the process. We want the field to be random but keep some control over it. A first-time, having a general global deformation progressively allowing finer deformation. It is done on line 8 by increasing the number of control points after a few iterations (variables `p` and `q`). After having generated the deformation series, we integrate over the vector fields to have a single deformation that we can apply to the Source image of step 1.

STEP 3: Eventually, we add the 'Oedema infiltration' by adding `.1` to the image at a location given by a predefined mask.

To sum up, four parameters can be chosen randomly. **1.** The background deformation control matrix, needs to have strong deformation at the centre only. **2.** The cancer initialisation point, needs to be placed around the centre of the Oedema mask. **3-4.** The deformation amount for cancer growth is controlled by the increments number and the amplitude of the bspline field.

Source Code 3.3: Necrosis growth construction for toy-example

```
1  sigma = 14
2  smooth = GaussianBlur2d(kernel_size, sigma)
3
4  n_iter = 30
5  field_stock = torch.zeros((n_iter,)+size+(2,))
6  necrosis = ini_ball.clone() # Initialisation of the tumour as a small ball
7  for i in range(n_iter):
8          p = q = max(4, i//4)
9
10         # Building a field
11         cm = (torch.rand((2,p,q))*2-1)*50 # Random control points
12         field = field2D_bspline(cm,size,degree=(2,2),dim_stack=-1)[None]*10
13
14         # restraining the field to the edges of the necrosis
15         grad_necrosis = tb.spacialGradient_2d(cancer).sum(dim=2)[0]
16         field = smooth(grad_cancer * field)
17         field_stock[i] = field[0] # save
18
19         # incremental necrosis growth
20         deform = reg_grid - field
21         necrosis = tb.imgDeform(necrosis,deform,dx_convention='pixel')
22
23 # FieldIntegrator is Explained in Chapter 1
24 int_deformator = FieldIntegrator(method='temporal')(field_stock,forward=False)
25 int_deformation = FieldIntegrator(method='temporal')(field_stock,forward=True)
26
```

```
27  cancer_sharp = tb.imgDeform(ini_ball,int_deformator,dx_convention='pixel')
```

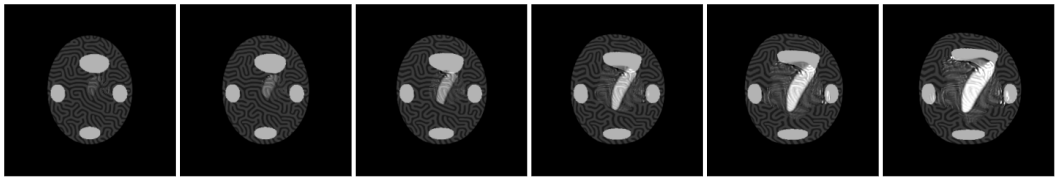

(a) Registration evolution

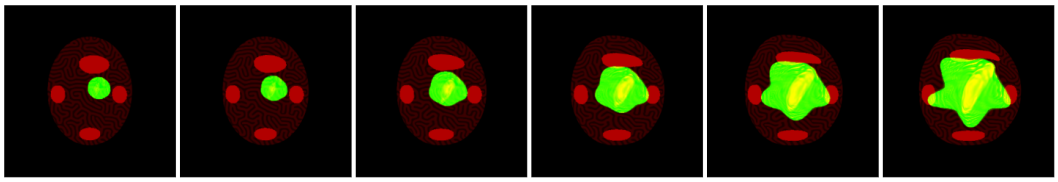

(b) Deformation only

Figure 3.3: Weighted Metamorphosis fails on the refined ToyExample `mu=1,rho=10,gamma=0` [add circle ?]

As we have seen, we have made a procedure to build a toy example data set, with the perspective to perform a benchmark. However, in this dissertation, we study the registration of only the image pair of Figure 3.2. One advantage of the textures added compared to the white background toy example, is that we can easily track the movement with the background as well. We use two quantitative metrics to evaluate the registration quality that focuses on the deformation:

- **Ovals dice score:** On both images, there are four ovals, representing organs, that get deformed. In Figure 3.2, one can see the actual displacement. We can retrieve a ground truth mask of these ovals by applying the spline deformation to the oval mask of the initial image. We can compare an estimated deformation with the DICE score of both estimated and ground truth deformed oval masks.
- **Landmarks:** With the help of the textured background, we set manually landmarks, that will be used to validate the deformations retrieved with the proposed methods. Landmark distance works well for validating images and will be also used later on real data. One can argue that we could also directly compare our estimated deformations with the bspline generated one. This is true and it will definitely be the case when benchmarking on a toy example dataset.

With this more realistic toy example, we see that Weighted Metamorphosis mismatch images near the tumour, as it can be seen in Figures 3.3 and 3.8 (right column). Note that the procedure to construct the mask is the same as the one for the Oedema mask in Section 3.2.3.a. On the image showing deformation only, we can clearly see that the displacement induced by the 'tumour growth' has been explained by intensity changes. Note that there are a lot of ripples in the intensities addition, it will be discussed later in the next section. This is the kind of behaviour that has motivated the development of the Oriented norm (discussed in Section 3.1.2.b) and leads to the next section.

3.2.3 Constrained Metamorphosis

3.2.3.a Building priors

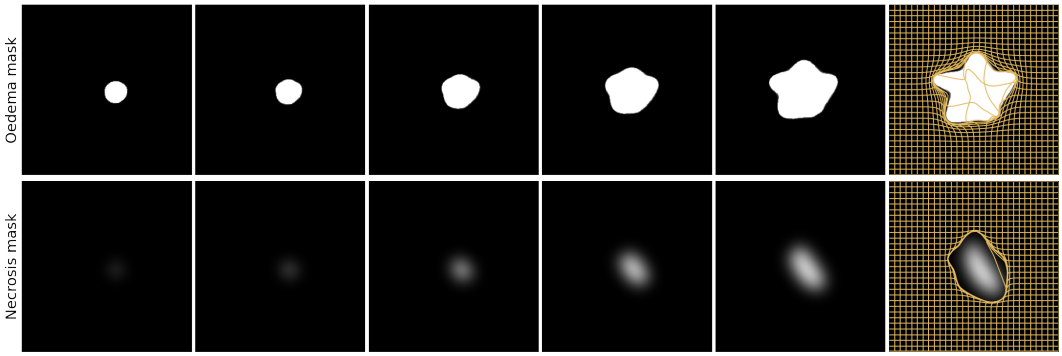

Figure 3.4: Visualisation of the masks P_t, Q_t and the deformation induced by w_t Each row shows the masks evolution through time. In the last column, we can see superposed the deformations generated by w_t (up) and $w_t \times P_t$ (below).

Before executing CM, it is necessary to carefully compute constraining priors. We need two growing masks and a temporal vector field. A mask Q_t with the information of the Oedema location, another one P_t for the necrosis and a vector field w_t transporting the mask information. We construct these three objects in one LDDMM registration detailed in Algorithm 3.4. First, we build a target image using the Oedema and Necrosis segmentations, setting the pixels at `.5` and `1` at respectively the Oedema and necrosis location else `0`. Secondly, we do the same for a source image with small balls at each segmentation barycentre. Finally, we register the source on the target, getting a temporal field W_t and image M_t. M_t is then used for our mask of interest, taking values at each time close to `1` as the necrosis mask P_t and close to `.5` as the oedema mask Q_t. We then smooth the mask P_t only as we want to use the field w_t near the necrosis but keep the fields smooth near its edges. One can see the masks in Figure 3.4.

Source Code 3.4: P_t, Q_t and w_t construction

```
# make target image
val_o,val_n = .5,1  # Oedema and necrosis pixel values
segs = torch.zeros(seg_necrosis.shape)
segs[seg_oedeme > 0] = val_o
segs[seg_necrosis > 0] = val_n

# make source image
ini_ball_n,_ = tb.make_ball_at_shape_center(seg_necrosis,force_radius=5)
ini_ball_o,_ = tb.make_ball_at_shape_center(seg_necrosis,force_radius=18)
ini_ball_on = torch.zeros(ini_ball_o.shape)
ini_ball_on[ini_ball_o > 0] = val_o
ini_ball_on[ini_ball_n > 0] = val_n

# LDDMM registration
n_steps = 10
```

```
residuals = 0 # short-cut for setting residuals as a zero matrix
sigma = [3,7,20]  # multi-scale registration
mr_mask =lddmm(ini_ball_on.to(device),segs,residuals,
        sigma=sigma,cost_cst=.0001,integration_steps=n_steps,
        n_iter=2500,grad_coef=50)

# Get mask and field
sig = 14
smooth = kornia.filters.GaussianBlur2d((sig*6+1,sig*6+1), (sig, sig))

# Oedema mask
mask_o = mr_mask.mp.image_stock.clone() # We want to have values close to val_0
mask_o[mask_o > val_o - .02 ] = 1       # to avoid unwanted holes in the mask.

# Necrosis mask
mask_n = mr_mask.mp.image_stock.clone()
mask_n[mask_n <= val_o + .05] = 0
mask_n[mask_n > 0] = 1
mask_n = smooth(mask_n)

w = mr_mask.field_stock
```

3.2.3.b Results

Having built priors, we can register our toy examples. Using Theorem and Proposition 3.1 for retrieving geodesic equation and minimising the cost:

$$J_{\mathrm{CM}}(z_0) = \|I_1 - T\|_{L_2}^2 + \lambda \left[\|v_0\|_W^2 + \gamma \|P_0(v_0 - w_0)\|_W^2 + \rho \|z_0\|_Q^2 \right], \quad \lambda \in \mathbb{R}^+, I_0 = S \tag{3.33}$$

gets us our registered images and matching deformation.

As one can see in Figures 3.6(a-c) the deformation deduced by CM gives good results qualitatively and quantitatively (see Validation metrics in Section 3.2.2). One can compare LDDMM, M, WM and CM outputs in Figure 3.7, and find dice and landmarks scores for each method. It is clear that CM surpass all others methods retrieving the closest deformation to the ground truth. Only talking about deformations: (far-left) LDDMM try to explain the necrosis apparition by deforming the upper organ (red arrow),

(centre-left) pure Metamorphosis archive to explain the intensities changes at the expense of the amount of deformation. Note that depending on the parameter ρ, the retrieved deformation could be the same as LDDMM (ρ too high) or the identity (ρ close to zero). (centre-right) Weighted Metamorphosis improves greatly the result while having trouble detaching the necrosis from the organ. Indeed during the mask's growth, the top organ overlaps the mask and is therefore explained by intensity additions, failing to recover this deformation locally. (far right) Finally, CM manage to register correctly all organs and gets the closest score for landmark distance (see Figures 3.6(a-c)). The landmark that failed to be matched is the one in the background textures, which is quite hard to match because of its high local similarities. In addition, in Figure 3.6c one can see that the background was not paired very precisely near the necrosis, which is expected as the necrosis removed some background material and induce strong deformations.

If we examine intensity changes in quality, some may be disappointed. Indeed the entanglement problem of Metamorphosis discussed in Chapter 2 is still present here, only restrained to the mask support. A balance between deformation and topological addition is still to be found. One can observe ripples appearing, (orange arrows in Figure 3.6a and 3.7). It is unclear if it is caused by the residuals z being used both for generating the vector fields v_t and the intensities additions or because of integration instabilities. Indeed, especially in these images, the background textures have a lot of variations that need to be displaced, we find these lines as the blue/red edges in the momentum (see Figure 3.5). Unfortunately, they end up added to the final image. This is a structural flaw of Metamorphosis that needs more work to be addressed. ;

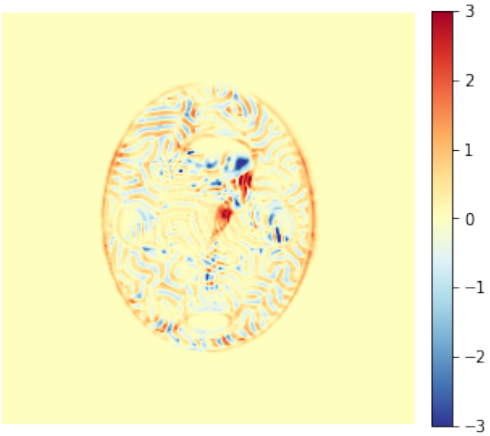

Figure 3.5: Constrained metamorphosis initial momentum z_0

3.2.4 Reversed Constrained Metamorphosis on toy-Example

We present here results from the variation of the 'Reversed' Weighted Metamorphosis (RWM) model defined in Section 3.1.4.b. As it is mentioned, we aim to transform the cancerous brain into the healthy one. In this model, to constrain the intensity addition, we simply rearrange the mask Q_t from Section 3.2.3.a in reverse order (*i.e.:* Q_{1-t}).

The registration result of RWM on the toy-Examples is displayed in Figure 3.9. By analysing panel (a) one can see that the tumour has been pushed to a fine line with a rather satisfying registration at first glance. Moreover one can clearly see that the background got very close to the right colour, despite the texture making it difficult. We measure the dice score and the landmark distance by applying the inverse deformation on the initial image, giving very good quantitative results as can be seen in the images on the right. Qualitatively, we can ponder these results as the deformation is both very strong and seems too localised. The estimated matching pushed the 'cancer' to a small line about at the location of the ball as initialised in step 2 of Figure 3.2. In the surroundings of the tumour, within the oedema part of the registration has been done by intensity changes. By a close inspection of panel (c) central image texture, one can see that the deformation gets quite close to the source image, which is impressive given the large deformation and the data occlusion.

The vanishing of the tumour is coherent with the data term $\kappa \left\| N \circ \Phi^{-1} \right\|_{L^2}^2$ we added, which intends to make the necrosis mask disappear. However, in Figure 3.10 one can see that changing the importance of κ has little impact on the registration. Worse, as the number of iterations to gets below a given value ('nifc' in Figure 3.10) indicates to us, the higher κ was, the longer one had to wait for convergence. Therefore, both ruling out this hypothesis and challenging the decision to add a new data term.

In conclusion, the model has generated a deformation that effectively addresses the difficult registration problem. However, there remains uncertainty about how to recover the

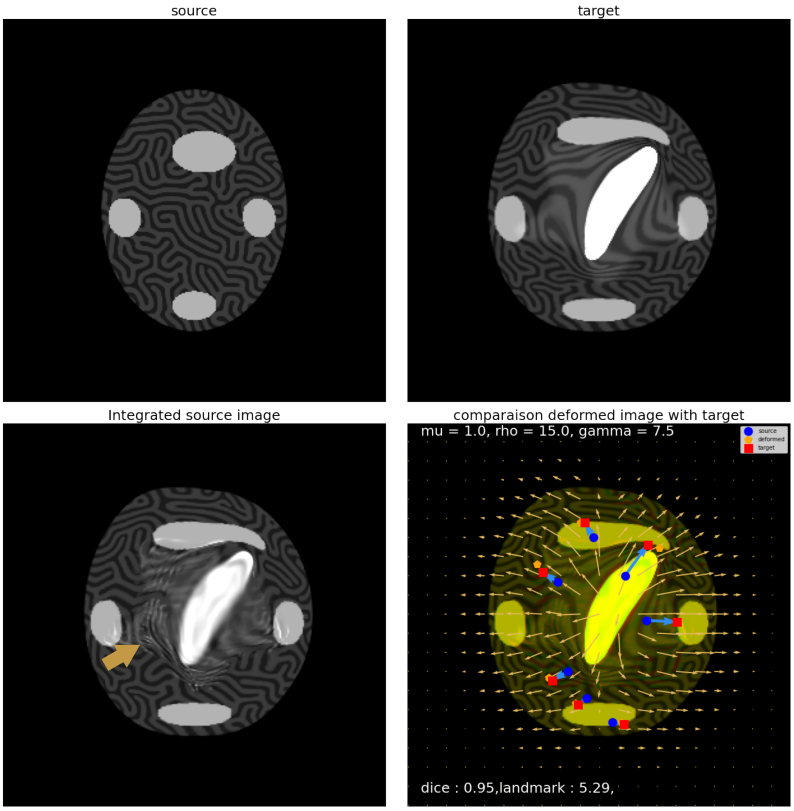

(a) Registration summary

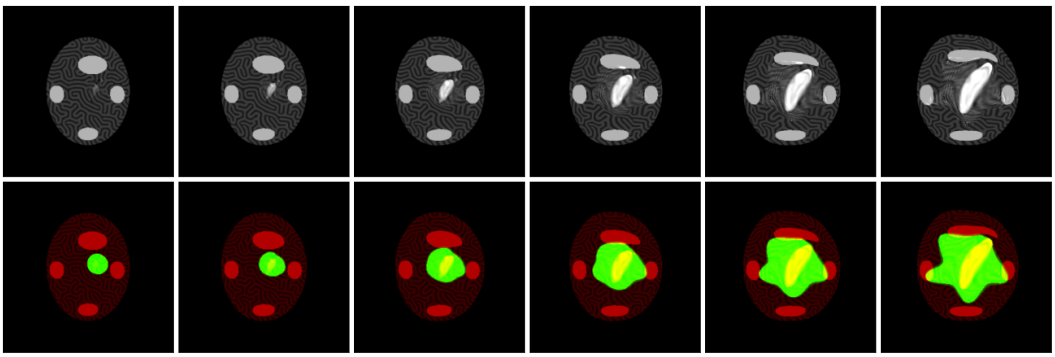

(b) Registation evolution (top) with mask superposition (bottom)

(c) Deformation only

Figure 3.6: Constrained Metamorphosis on ToyExample. See Figure 1 for an image comparison colour code explanation.

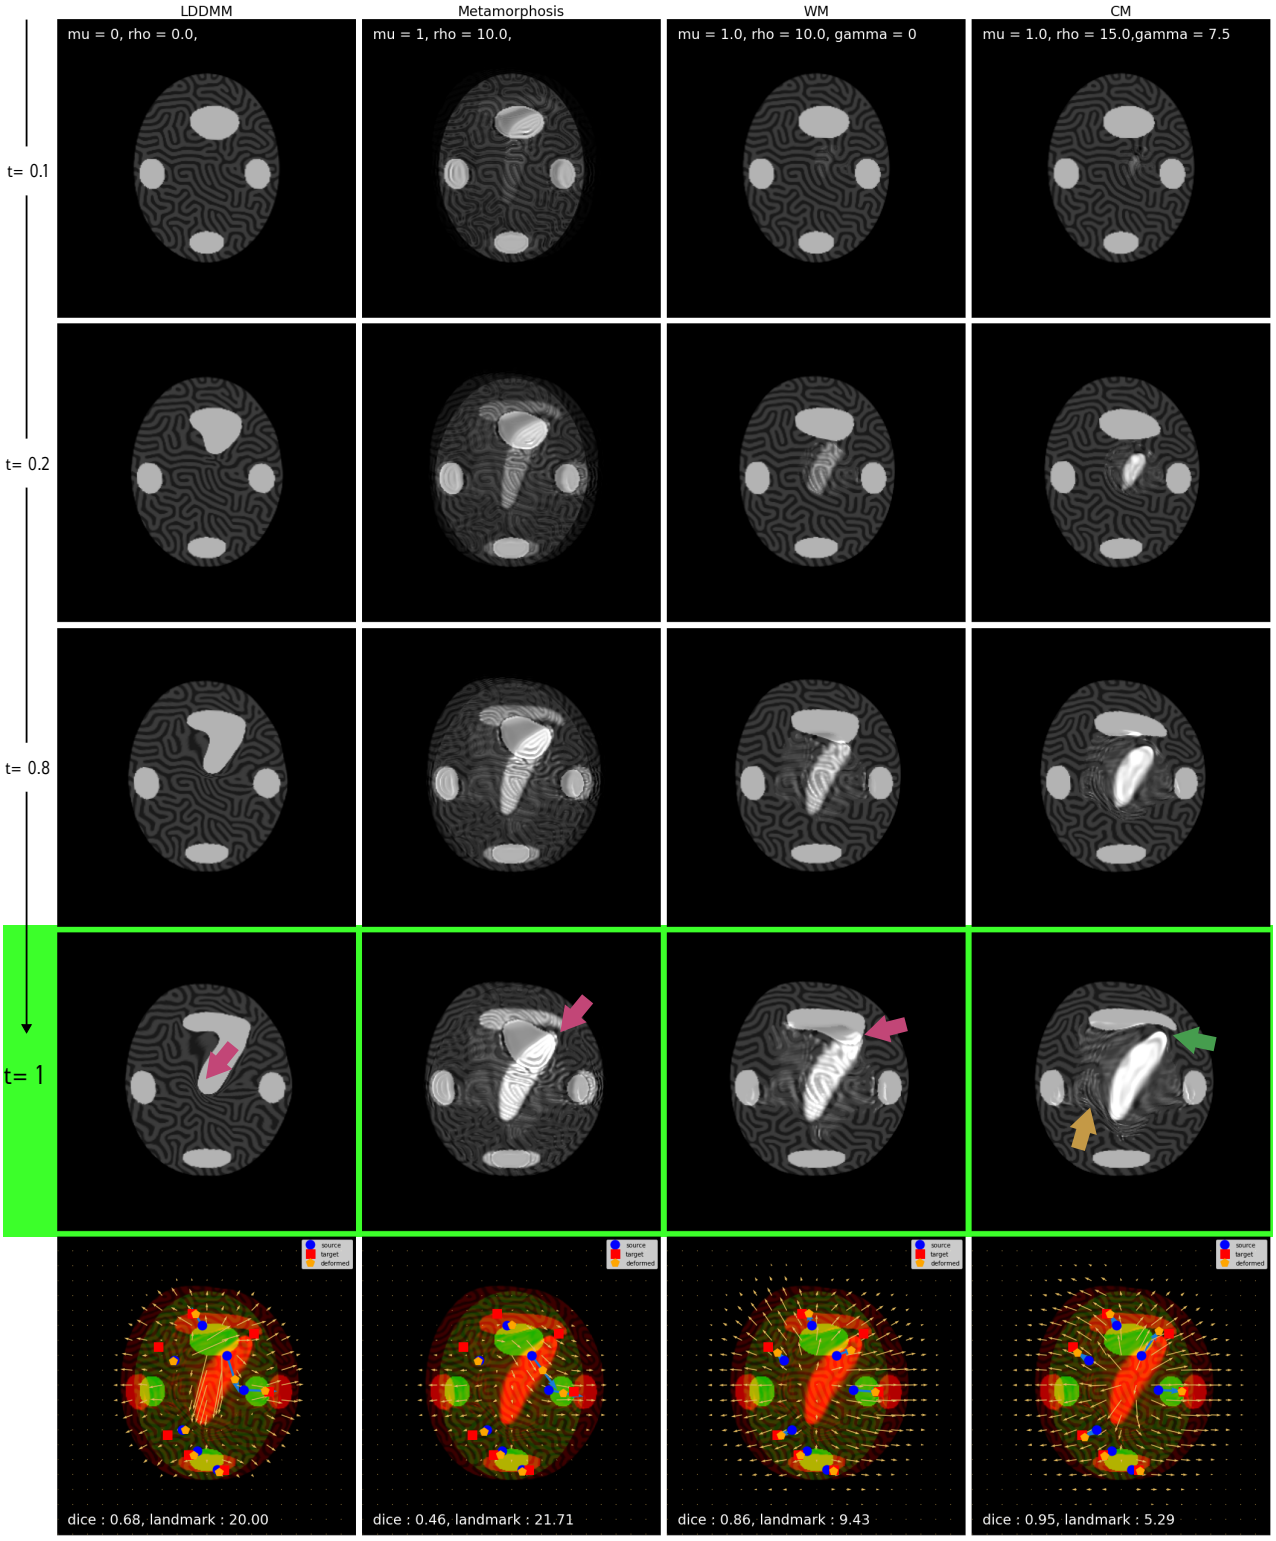

Figure 3.7: Comparison of Metamorphosis-related methods on the forward direction. Each column represents an image integration. The last row is a summary of the registration. See Figure 1 for an image comparison colour code explanation.

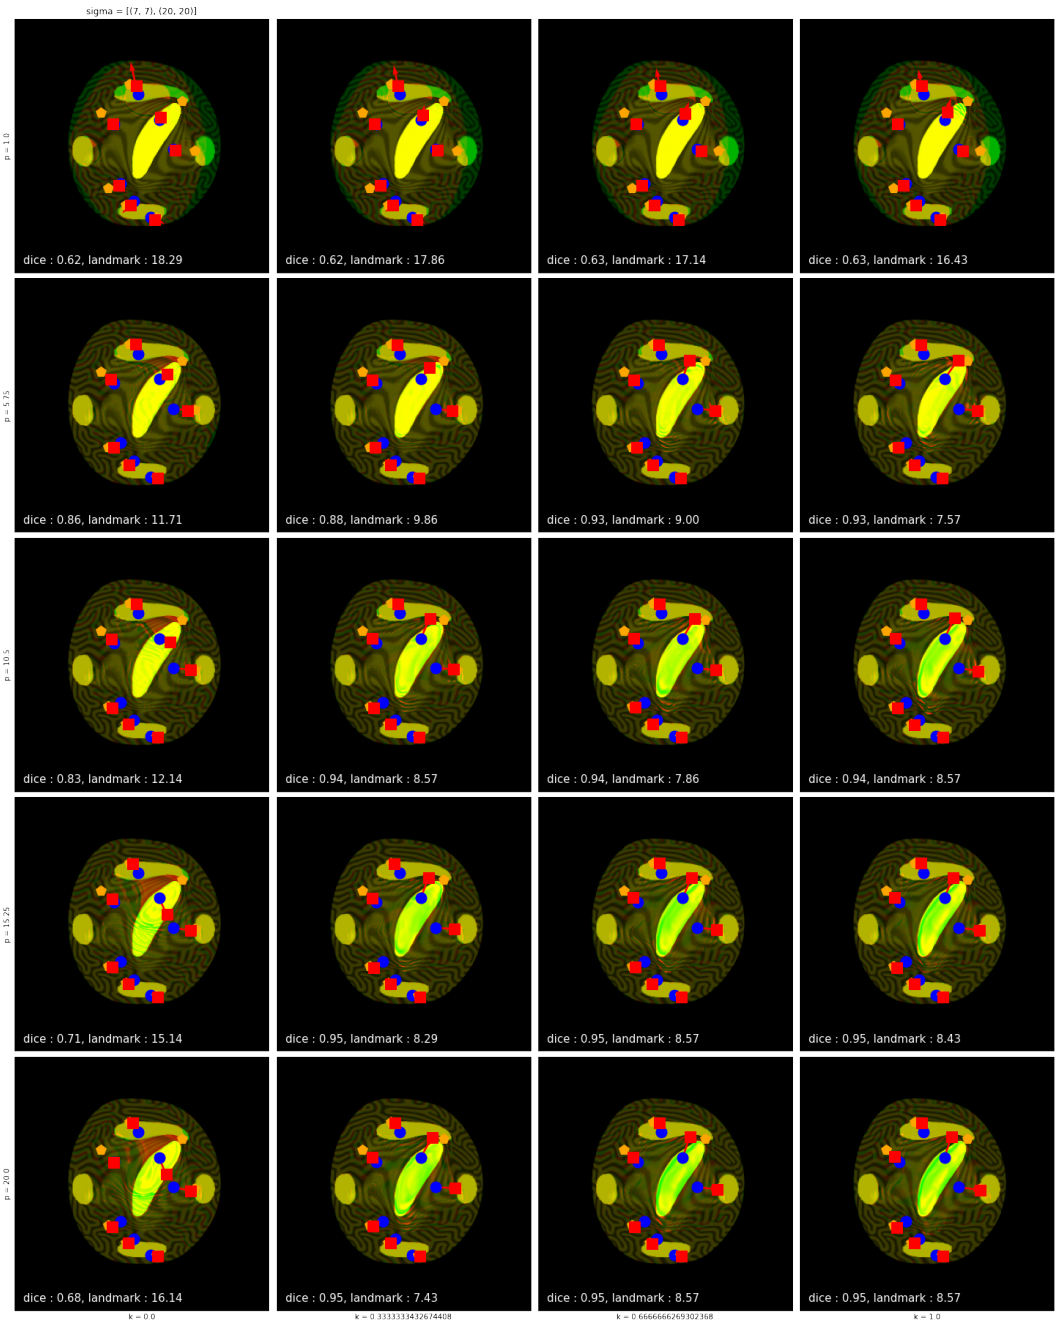

Figure 3.8: Constrained Metamorphosis sensibility to parameters. See Figure 1 for an image comparison colour code explanation.

information that was removed through intensity subtraction in order to re-incorporate it into the final image.

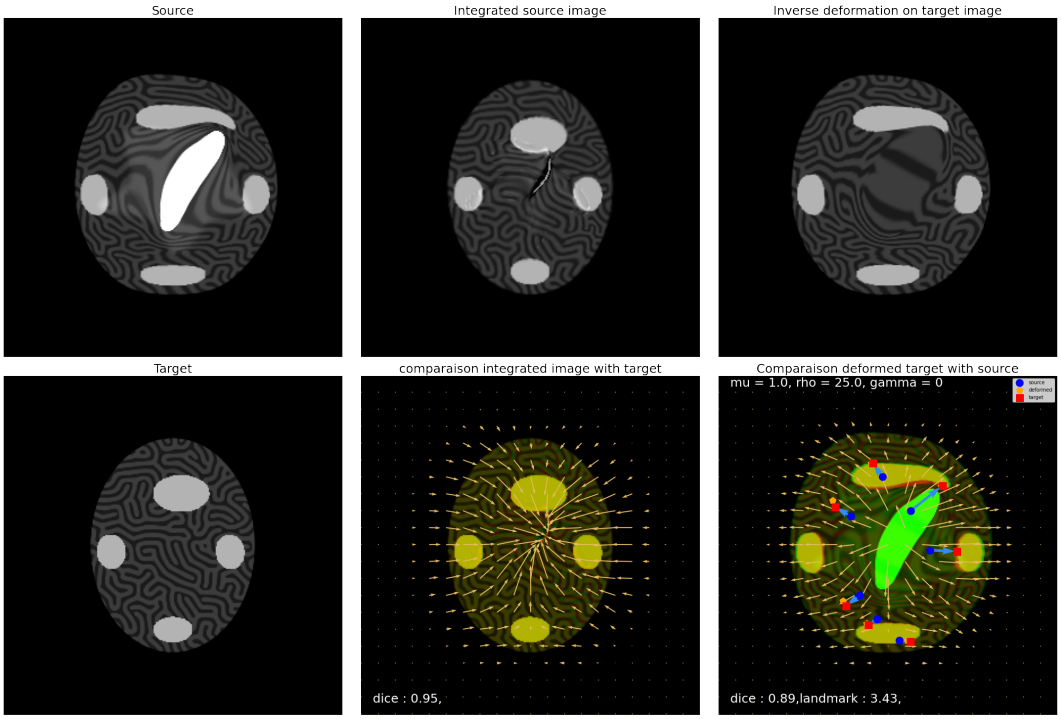

(a) Registration deformation (centre) and inverse deformation (right) summary

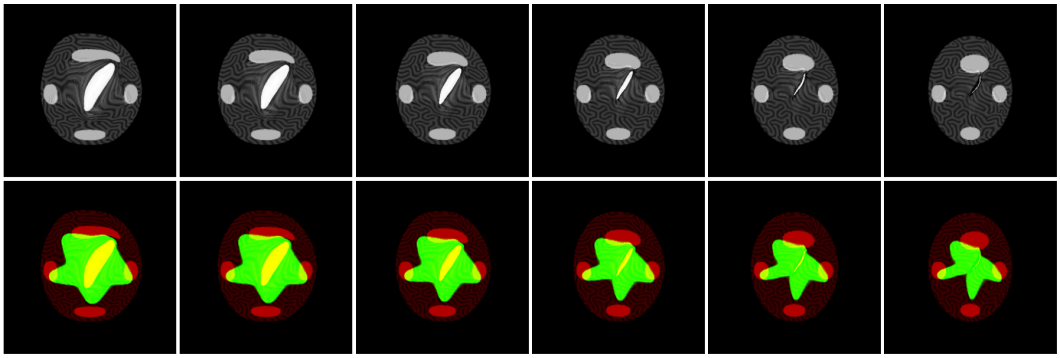

(b) Registration evolution (top) with mask superposition (bottom)

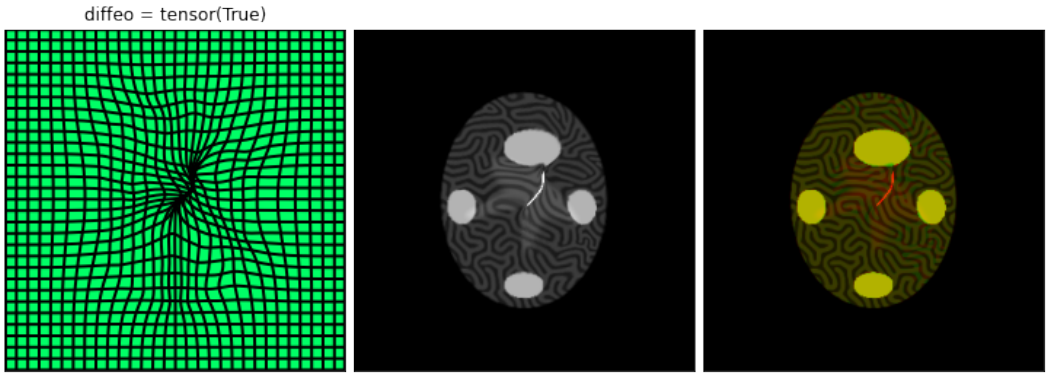

(c) Deformation only

Figure 3.9: Reversed Weighted Metamorphosis on toy example

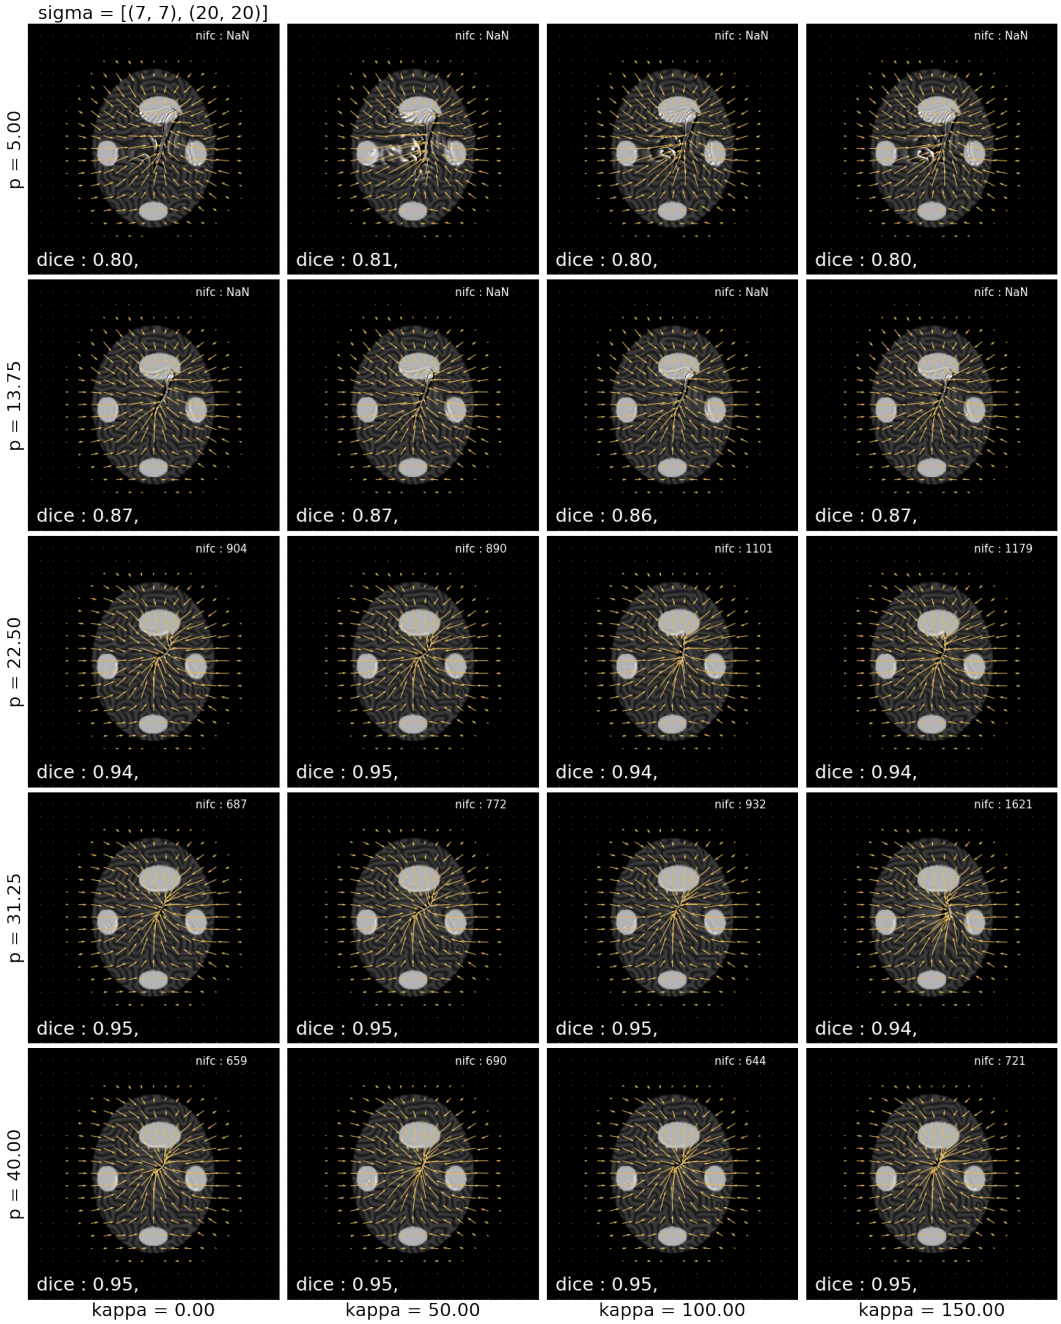

Figure 3.10: Reverse Weighted Metamorphosis sensibility to parameters. nifc: in each image right corner corresponds to the first iteration number where the ssd was below 50. If the value is NaN, the optimisation never reached a value smaller than 50.

3.3 Registering Cancerous brains

In this section, we present the results of applying the Constrained Metamorphosis framework, to real 3D MRI brain images with tumours, demonstrating its efficacy in tackling this challenging registration problem. First, we will discuss the process of matching a healthy brain template to a cancerous one, which is similar to the toy-example problem presented in the previous section. In the second part of this section, we will investigate the use of our framework to register cancerous brain MRIs with their follow-up surgery analogues. While the framework was not initially designed for this purpose, we were able to take advantage of the BraTSReg2022 Challenge to obtain relevant data and validation tools, enabling us to further evaluate the effectiveness of our method.

3.3.1 From template to cancerous

3.3.1.a BraTS2021 Dataset and sri template

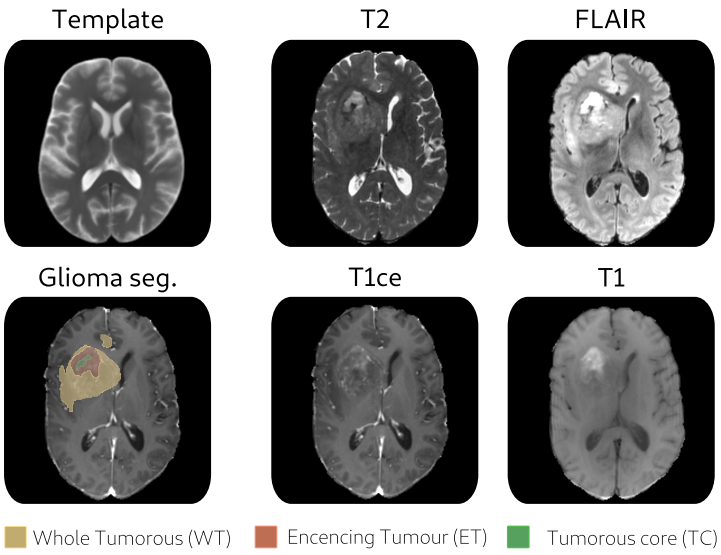

Figure 3.11: Visualisation of BraTS2021

The *center for Biomedical Image Computing & Analytics* (CBICA) of the Perelman school of medicine run the Brain Tumour Segmentation (BraTS) challenge for ten years. The BraTS2021 challenge goal [Baid and et al., 2021; Menze and et al., 2015; Bakas et al., 2017] was to segment efficiently intrinsically heterogeneous brain glioblastoma sub-regions in mpMRI scans. The sub-regions considered for evaluation are the 'enhancing tumour' (ET), the 'tumorous core' (TC), and the 'whole tumour' (WT) (see Figure 3.11). The ET part is described by areas that show hyper-intensity in T1ce when compared to T1, but also when compared to 'healthy' white matter in T1ce. The TC describes the bulk of the tumour, which is what is typically resected. The TC entails the ET, as well as the necrotic parts of the tumour. The appearance of the necrosis is typically hypo-intense in T1ce when compared to T1. The WT describes the complete extent of the disease, as it entails the

Table 3.1: Quantitative evaluation for different registration methods. Results were computed on a test set of 50 2D 240x240 images from the BraTS 2021 dataset. - $(*)$ SSD for CFM is computed over the domain outside the mask.

Method	LDDMM	Meta.	WM (ours)	MAE	Voxelm.	CFM
SSD (final)	223 ± 51	$\mathbf{36 \pm 9}$	65 ± 71	497 ± 108	166.71 ± 37	$49^{*} \pm 28$
SSD (def.)	-	112 ± 21	$\mathbf{102 \pm 76}$	865 ± 172	-	-
Dice score	68.6 ± 11.9	74.1 ± 9.3	$\mathbf{77.2 \pm 10.1}$	60.6 ± 8.79	66.8 ± 10	45.0 ± 13.5

tumorous core and the peritumoral edematous/invaded tissue, which is typically depicted by a hyper-intense signal in FLAIR. We must note the rare exception of the astrocytomas (IDH-mutant, 1p19q non-codeleted) that present a T2/FLAIR mismatch [Deguchi et al., 2020]. We approximate the necrosis part for building masks as the union of the Tumorous core (CT) and the Enhanced part (ET). The Oedema is covered by the rest (WT).

For the challenge, they provided a set of 1600 MRI with four modalities: Native (T1), contrasted and enhanced (T1ce), T2-weighted(T2) and T2-fluid Attenuated Inversion Recovery (FLAIR). It came along with a physician-made manual segmentation. Answering the challenge is out of this chapter's scope and it will be discussed in chapter 4, we decided to use BraTS2021 segmentations for building our prior mask in a similar fashion than discussed in section 3.2.2

This dataset provides a lot of test images and we still have to choose a source moving image. We decided to use the SRI24 Altas [Torsten et al., 2010] as a normal human brain anatomy standard reference. The main reason is the fact all images from BraTS2021 were registered to this template and resampled to 1mm^3 as part of their prepossessing pipeline as stated in Bakas and et al. [2018]. Furthermore, the SRI template is well suited for label propagation and spatial normalisation, providing very sharp images for a template.

3.3.1.b Weighted Metamorphosis on 2D MRI slices

As for Section 3.2.1, Our first attempt to solve this problem was published in François et al. [2022]. For evaluation, we used T1 MR images from the BraTS 2021 dataset [Baid and et al., 2021; Menze and et al., 2015]. At this time we had some issues with the registrations of 3D images, so for each patient, selected the same slice for 50 patients resising them to 240x240 and making sure that a tumour was present. We then proceeded to register the healthy brain template SRI24 (Torsten et al. [2010]) to each of the selected slices (see Fig.3.12 for two examples). To evaluate the quality of the alignment we used three different measures in Table 3.1:

1. the Sum of Squared Differences (SSD) (*i.e.:* L2-norm) between the target (T) and the transformed source (S) images. This is a natural choice as it is used in the cost function.

2. the SSD between T and the deformed S without considering intensity changes. This is necessary since Metamorphoses could do a perfect matching without using deformations but only intensity changes.

3. A Dice score between the segmentations of the ventricles in the deformed S and T. The ventricles were manually segmented. All methods should correctly align the ventricles using solely pure deformations since these regions are (theoretically) not infiltrated by the tumour (*i.e.:*, no intensity modifications) and they can only be displaced by the tumour mass effect.

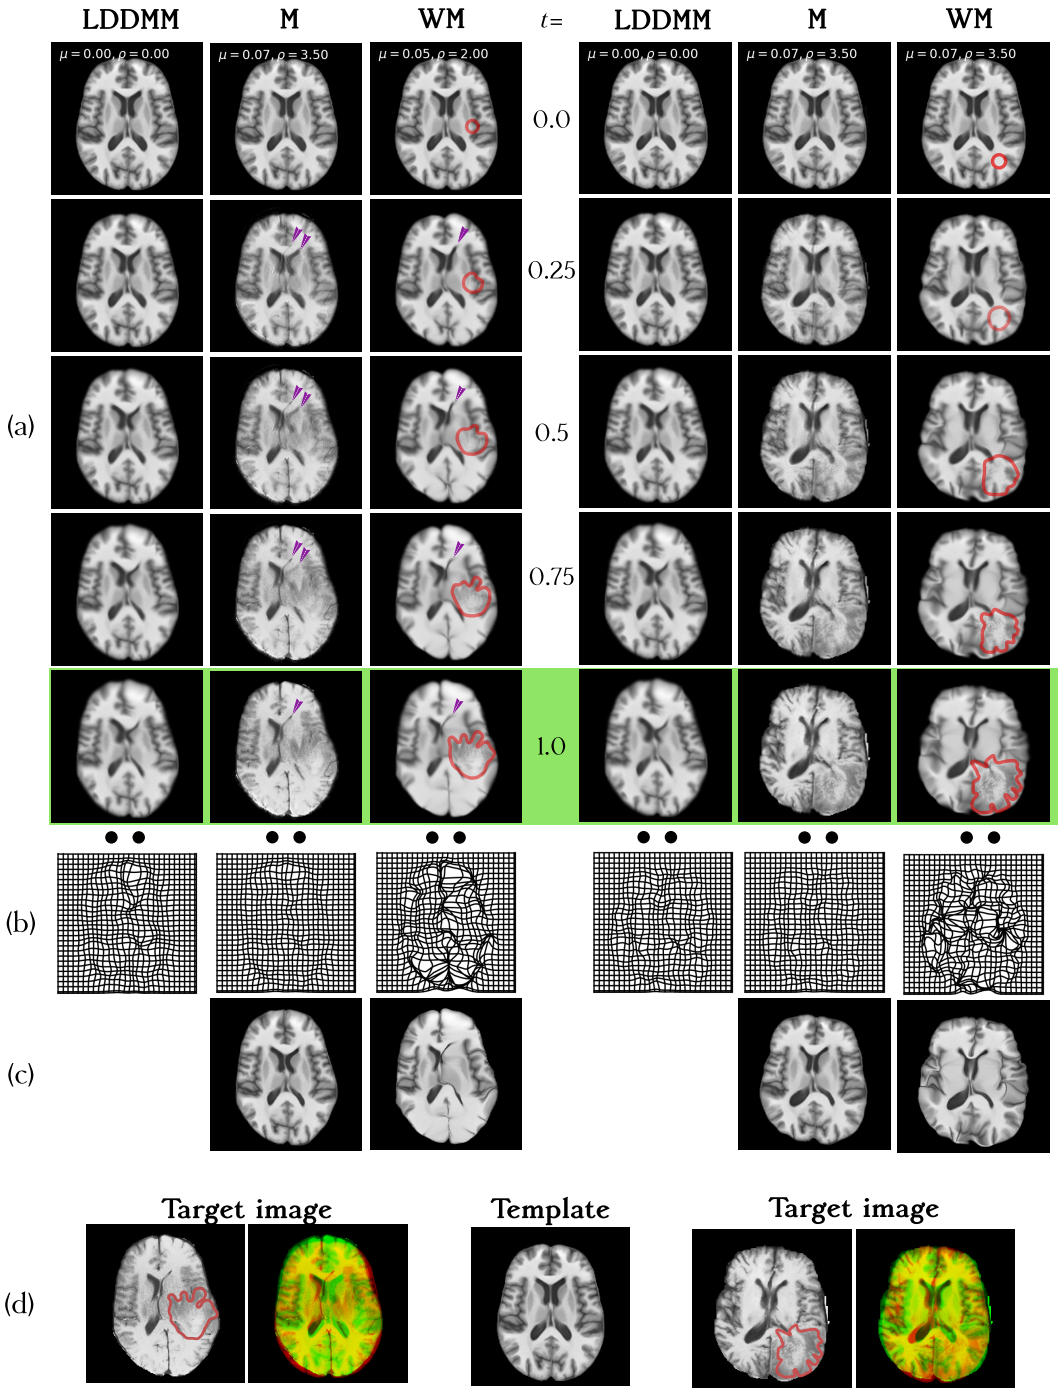

Figure 3.12: Registrations on MRI brain slices presenting brain tumours. Two examples from BraTS 2021 database Baid and et al. [2021]; Menze and et al. [2015]. Comparison of geodesic shooting for LDDMM, Metamorphosis (M) and Weighted Metamorphosis (WM). (a&d) On the target images and the geodesic integration, the temporal mask is indicated by the red outline. The final result of each integration can be seen in the green outlined row. (b) The deformation grids retrieved from each method and (c) the template image deformed without intensity additions for each concerned method. Purple arrows in columns 2 and 3 in the top right part of each image show the evolution of one ventricle through registration. While M makes the ventricle disappear and reappear, WM coherently displaces the structure. (d) Target images with the segmentation outlined in red; the coloured image is its superposition with the source. see `animations in GitHub in notebook : brains_weightedMetamorphosis.ipynb`

We compared WM with LDDMM [Beg et al., 2005], Metamorphosis [Trouvé and Younes, 2005], using the implementation of François et al. [2021], Metamorphic Auto-Encoder (MAE) [Bône et al., 2020], Voxelmorph [Balakrishnan et al., 2019] and Cost Function Masking (CFM) [Brett et al., 2001] (see Table 3.1). Please note that we did not include other deep-learning methods, such as Han et al. [2020b]; Maillard et al. [2022], since they only work the other way around, namely they can only register images with brain tumours to healthy templates.

As expected, Metamorphosis got the best score for SSD (final) as it is the closest to an exact matching method. However, WM outperformed all methods in terms of Dice score obtaining a very low SSD (both final and deformation-only). This means that our method correctly aligned the ventricles, using only the deformation, and at the same time it added intensity only where needed to globally match the two images (*i.e.*, good disentanglement between shape and appearance).

In this section, we discussed results for 2D images, which must be seen as proof of concept. Indeed even if we did our best to select slices in both source and target images that were displaying the same organs, the deformation induced by the cancer is likely to have components perpendicular to the plane, making the problem impossible to solve. From now on, all discussed results will be in 3D. However, the theory and the codes can be applied to both.

3.3.1.c Constrain Metamorphosis prior information pipeline

More effort is required in constructing priors for Constrained Metamorphosis compared to WM. In this section, we will elaborate on the pipeline employed for mask and field construction, as outlined in Code 3.5. Similar to Section 3.2.3, the objective is to generate two temporal masks: the first mask Q_t emulates the whole tumour including the Oedema while the second mask P_t represents the necrosis. One can incorporate intensity changes within both masks and the field will be oriented around the necrosis one only. Accurate localisation of the 'cancer' initialisation is crucial when working with real data. The procedure for constructing the masks can be divided into five steps.

STEP 1: We select the intersection of the white matter within the template and of the Glioma segmentation in the target (see line 8). Indeed tumours typically originate in the white matter and do not infiltrate the ventricles [Hanif et al., 2017]. This intersection makes sense as the tumour will grow but not change its localisation. furthermore, as mentioned in Section 3.3.1.a, all brains where rigidly aligned beforehand.

STEP 2: Instead of searching for the intersection barycentre or centroid, we determine the *Pole of inaccessibility* to deal with the non-compact twisted shape that often results after Step 1. The Pole of inaccessibility is a concept from geography that refers to the farthest internal point in a closed surface. Although exact algorithms exist for polygons [Garcia-Castellanos and Lombardo, 2007], we opt for a faster and more image-related approach. To approximate the Pole of inaccessibility, we erode the segmentation until only a small seed remains (lines 10-20), and compute its centre as the mean in each dimension.

STEP 3: The seeds returned by step 2 being quite small, it would be hard to find a diffeomorphism to the segmentations. The initial necrosis mask P_t is obtained by dilating it. Dilating about the same number of times we needed to erode is a good amount to get an initialisation of reasonable size.

STEP 4: Then we compute the oedema initialisation mask Q_t, by placing a ball at the seed centre, making sure to stay in the tumour mask and to remove all intersections with the ventricles on the source image.

STEP 5: Finally, we obtain the temporal masks by stacking the previously discussed initialisation and registering them with the glioma segmentations using LDDMM.

```
# 1./ get all segmentations
seg_oedeme = seg_tumour > 0
seg_necrosis = seg_tumour > .5

whiteMatter_seg = get_template_whiteMatter()
ventricule_seg = get_template_vesicule_seg()

intersection = seg_necrosis * whiteMatter_seg

# 2./ erode the intersection until necrosis seed is small enough
eroded = intersection
ini_ero_sum = eroded.sum()
count = 0
while eroded.sum()  > max(ini_ero_sum//100,100):
        eroded = binary_erosion(eroded,iterations=1)
        count += 1

# find center of seed
indexes = (eroded == 1).nonzero()
center = tuple([int(array.mean()) for array in reversed(indexes)])

# 3./ dilate the seed to have a starting image of reasonable size
dilated = binary_dilation(eroded,iterations=max(min(count,5),1))

# 4./ construct eadema initialisation
ini_ball_o,_ = make_ball_at_shape_center(seg_oedeme,force_center=center)

# clean the inital ball by removing ventricule intersection
ini_ball_o[ventricule_seg[None,None] == 1] = 0
# and removing iniball outside of the target oedema
ini_ball_o[seg_tumour == 0] =0

# 5./ Stack initialisation images
ini_ball_on = torch.zeros(seg_tumour.shape)
ini_ball_on[ini_ball_o > 0] = val_o
ini_ball_on[dilated] = val_n

```

```
# and register them.
sigma = [1,3]
mr_mask = mt.lddmm(source = ini_ball_on,
        target = seg_tumour,
        sigma=[1,3],
        cost_cst=.0001,
        integration_steps=10,
        n_iter=500,
        grad_coef=1,
        sharp=True)
```

Source Code 3.5: P_t, Q_t and w_t construction for brats2021

3.3.1.d Results: Metamorphosis Methods comparison on 3D MRI

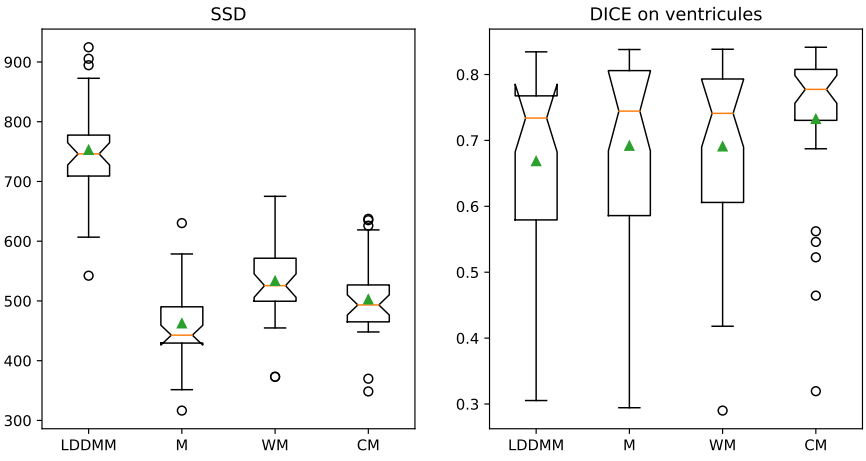

Figure 3.13: Quantitative results on brats2021. Notched box plots for 35 3D MRI. The orange line is the median and the green triangle represents the mean. The notches represent the 95% confidence interval around the median. The flipped appearance on the LDDMM DICE box is expected behaviour and means that the confidence interval is greater than the upper quartile.

As before we used SSD and a DICE score on ventricle segmentation to perform quantitative registration quality validation. As we didn't hear of an easy and validated way to segment ventricles on heavily deformed brains, I went through the tedious task of segmenting 35 subjects manually*. I want to note, however, that with the results of Chapter 4, we might be able to segment ventricles as well. For this registration task, we used T2 MRI images of size `120x120x78`. We chose the parameter balance $\mu = 1$ and $\rho = 5$ to advantage deformation over intensity changes (see Section 3.1.3.b).

In Figure 3.13 one can see that the three Metamorphosis based methods outperform LDDMM for the SSD. The best being pure Metamorphosis as expected, as it can add

*I have a thought to my medical colleagues that perform these kinds of tasks routinely on larger scales.

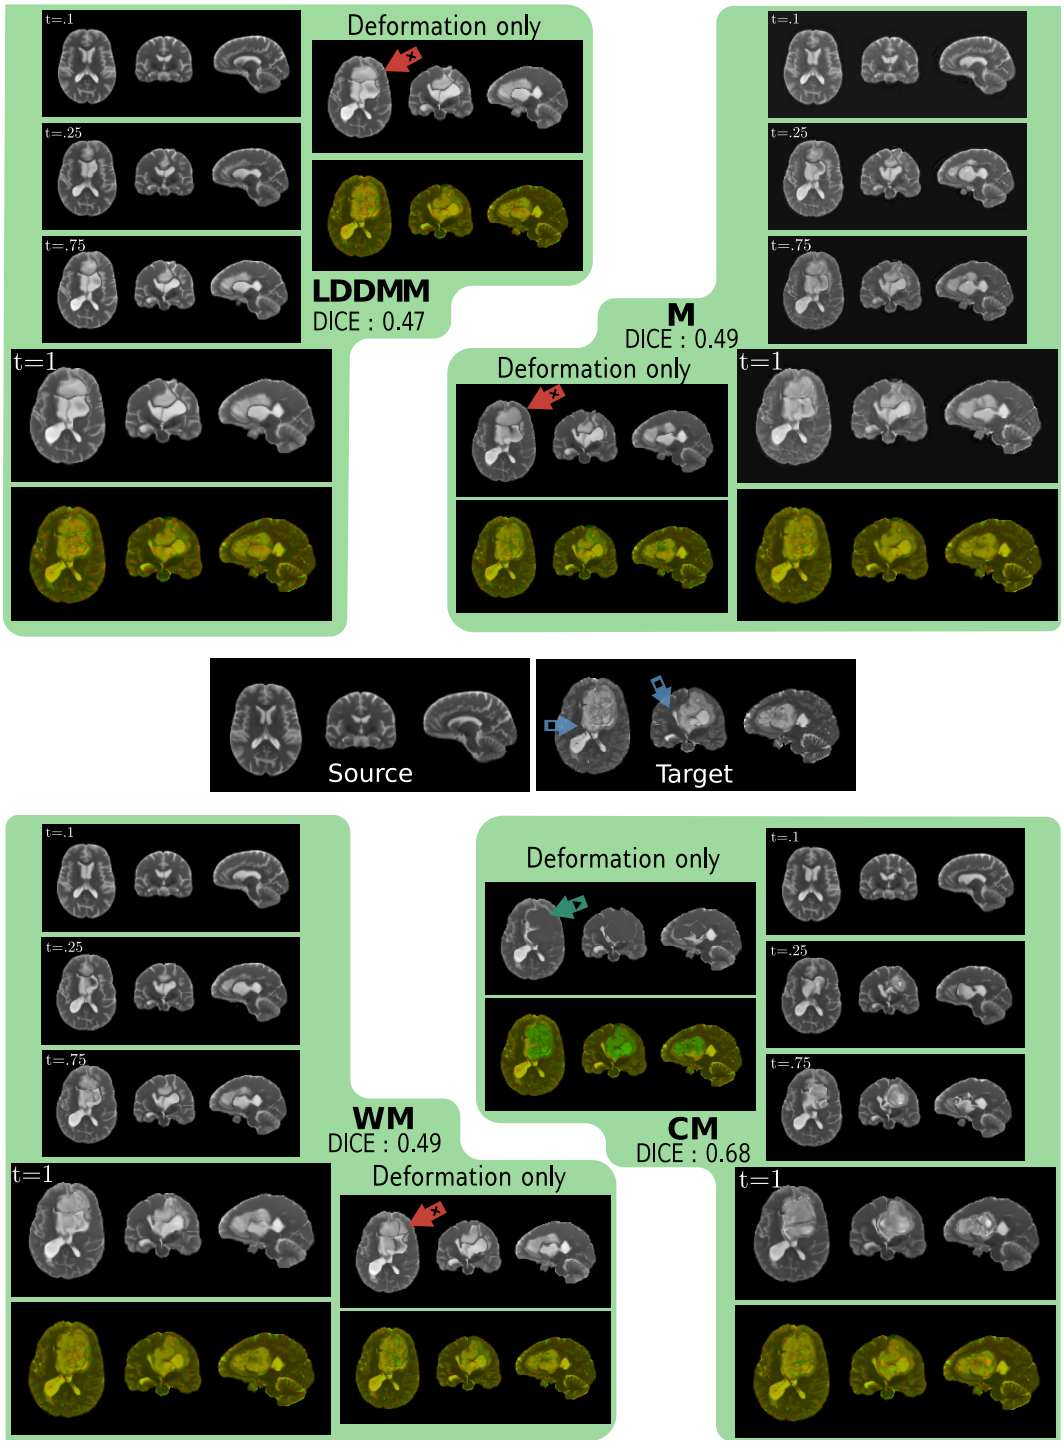

Figure 3.14: Visualisation of the different methods on BraTS2021 Green circulated areas display information for respectively LDDMM, M, WM and CM. Each one constrains the image evolution, marked with the time it is in white. The last one is bigger and compared with the target below. On the side, one can see the source image is solely deformed by all methods. Old colour code used, see Figure 1 for an image comparison colour code explanation.

intensity changes everywhere. However, WM and CM are close in terms of SSD scores, which means that significant intensity changes occurred in the lesion region. Talking only about deformation, one can see that CM outperform all three other methods on the ventricle DICE score, having only a few outliers. It is known that even LDDMM match well brain with small lesions, with CM we widened the use cases.

In Figure 3.14 we show a typical result where CM improve the matching significantly compared to other methods. The target image here displays a huge glioma, inducing high pressure into the patient's brain and compressing the temporal part of the ventricles, making them almost invisible (blue-squared arrows). It is a hard task to retrieve the deformation induced by the tumour. LDDMM, M and WM chose wrongly to deform the parietal lobe folds and the ventricles to match the tumour (red-cross arrows). On contrary, CM oriented by the a-priori field matched the glioma with the white matter region (green-triangle arrow), which is closer to an anatomically plausible match. We are not sure and can not verify that this deformation is the actual one, but it is at least a step in the good direction.

In Figure 3.15 we study the worst result obtained. None of the methods tested managed to register this image, with a DICE score of about 0.3 each. The thin red arrows show the ventricle displacement matching with the tumour. The white dotted line on the brain compared with the temporal mask helps to see that the necrosis mask is growing only toward the skull. Thus the final displacement field follows the general displacement of the temporal mask field. The algorithm fails because, in this example, the initialisation of the cancer ball was misplaced at the edge of the necrosis.

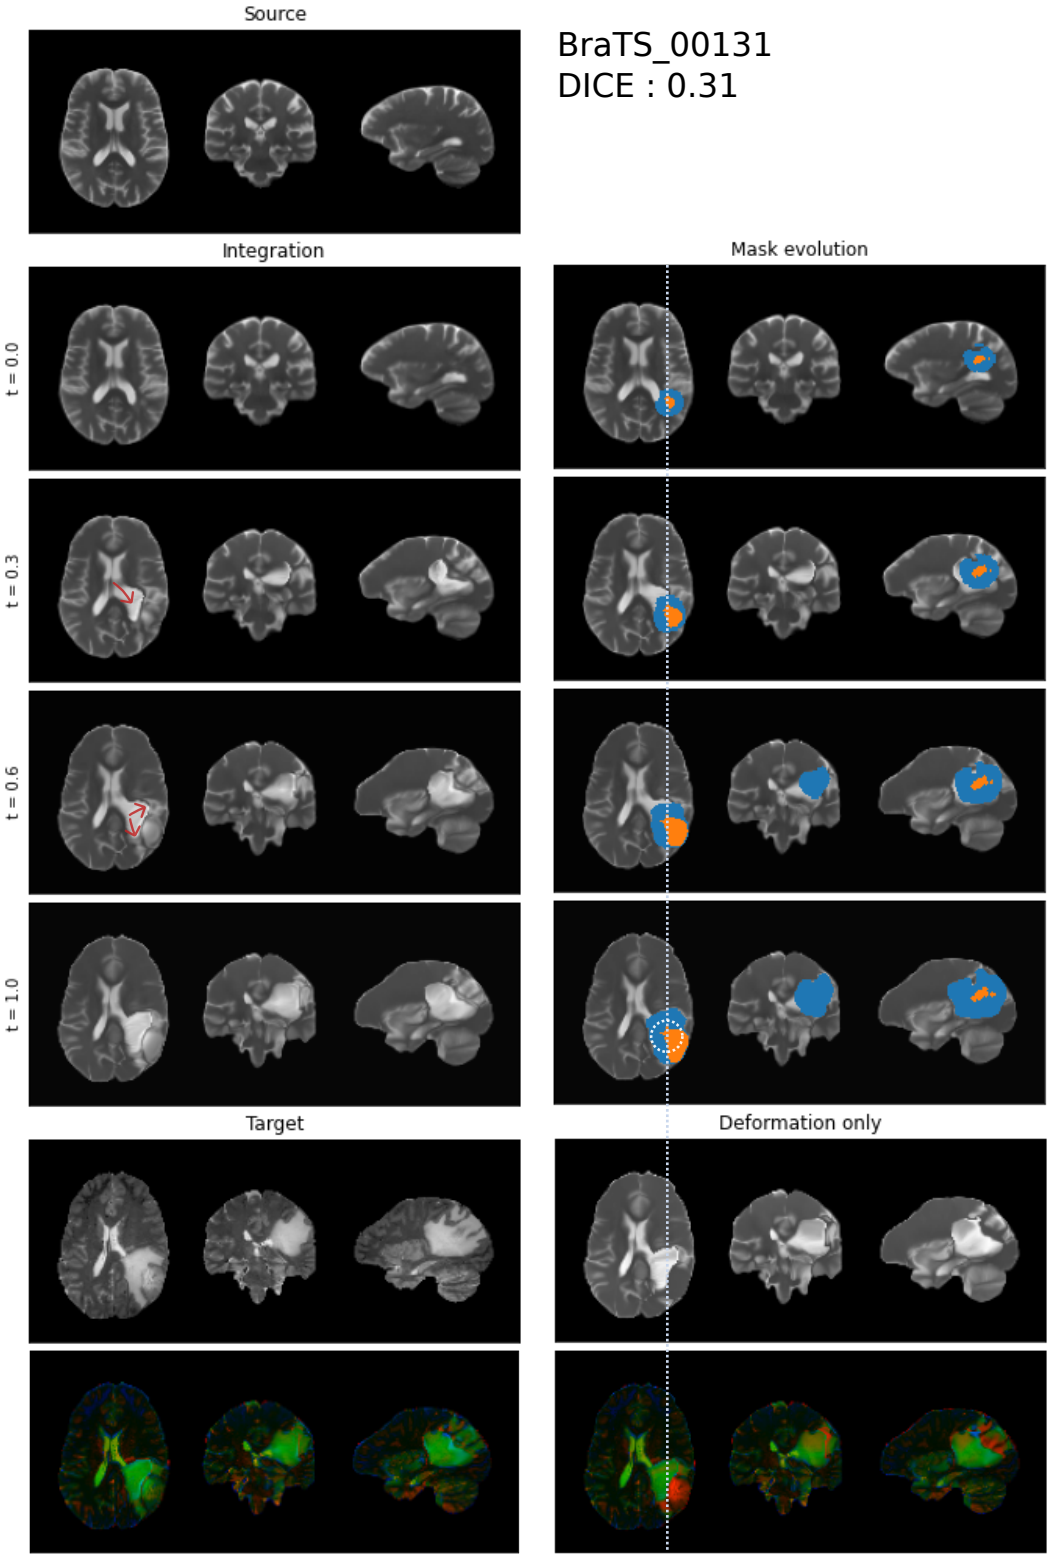

Figure 3.15: Outlier example of failed registration on the image BraTS_00131. New colour code used, see Figure 1 for an image comparison colour code explanation.

3.3.2 An other application: follow up surgery

3.3.2.a BraTSReg2022 Dataset

The BraTSReg2022 challenge organised by Baheti and et al. [2021] was a registration task aiming at matching MRI showing glioma with a follow-up surgery of the same patient. As the event organisers have pointed out, this registration task is hard due to all changes induced by the glioma growth as previously discussed in Section 1.1.2. However, these features were already part of the model in CM and the registration task is easier than the one we set up for BraTS2021. Indeed, because the pair of images are from the same patient, the source and target images showed the same overall structures outside the cancerous lesion (e.g.: grey matter folds,...) and only small displacement. On the contrary, the data set shows a lot of variation in the type of resection and how much the surrounding tissues filled the resulting hole.

The data set is made of 140+20 pairs of images (training + validation sets). Four modalities where provided: Native (T1), contrasted and enhanced (T1ce), T2-weighted(T2) and T2-fluid Attenuated Inversion Recovery (FLAIR) of size $240 \times 240 \times 155$. The data was gathered from different institutions and then pre-processed and curated for each scans size and resolution, according to a common anatomical template.

For assessing the quality of the registration, landmarks were provided for each patient, on both images for the Training set and on the follow-up operation for the Validation one. Meaning that challengers were expected to match from the follow-up to pre-operative brain and provided their estimated deformed landmark. As both images were acquired from the same patient it was possible to set these landmarks at recognisable anatomical shapes of the brain, blood vessels bifurcation and others. Note that even if a deformation perfectly matches landmarks, one is not sure to have found the true 'displacement' on the whole background.

In Figure 3.16 one can see two brain examples to be matched. In the left image, the tumour has been removed and the tissues resected well (green arrow). However, some Oedema was left over leaving a whitish region (blue arrow). On the right, the brain shows a hole on the side, having the same colour as the ventricles in pre-operative brains. On the follow-up ones, the hole is about the same size, with a darker colour, just as the Oedema.

3.3.2.b CM on the validation set

Our algorithm can be summarised in three parts :

1. Segmenting tumour, with Oedema/Necrosis distinction.
2. Registering the segmentations to get temporal masks using metamorphosis.
3. Applying Constrained Metamorphosis with the obtained temporal mask.

STEP 1: The segmentation was done by applying a U-net to both source and target images. It was trained on the BraTS2021 dataset [Baid and et al., 2021; Menze and et al., 2015] which contains annotated glioma segmentations along with four modalities. One could argue that those follow-up operative images should not be segmented by such a U-net, ... and they would be right. While testing out methods on the Training set, we observed that some

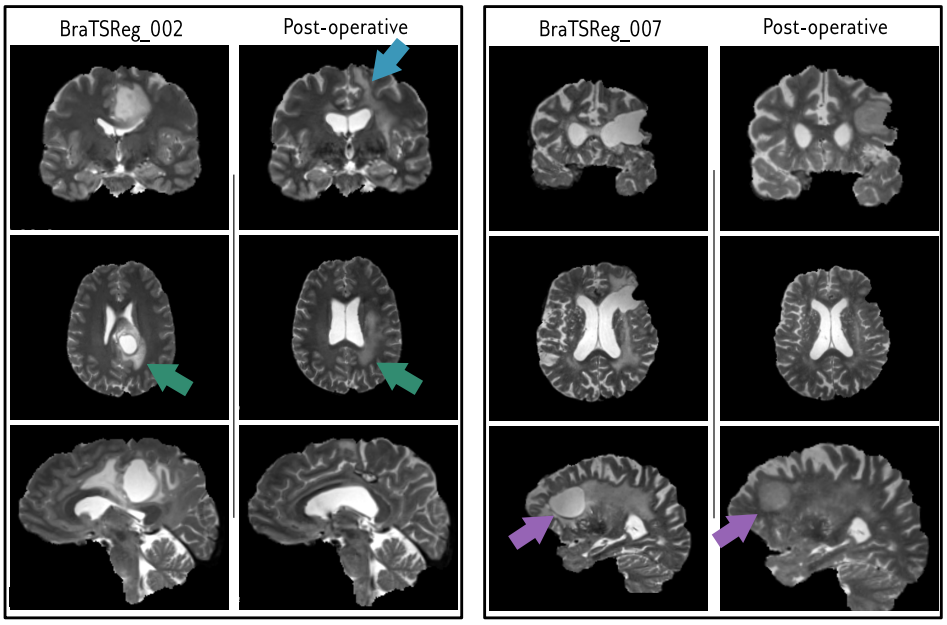

Figure 3.16: BraTSReg2022 data visualisation

bad results could be explained by wrong segmentations, which motivated the attempt to segment glioma using TDA (*c.f.:* Chapter 4). We selected the segmentations that were close to reality in a sub-training set.

STEP 2: We register one segmentation to another using Metamorphosis, setting the necrosis parts to 1 and the Oedema one to .5 (see Fig. 3.18, top rows). Getting this first registration is crucial as it is the one that will constrain the actual registration. We used Metamorphosis rather than LDDMM as sometimes new disconnected areas appeared which should not be matched by deformation.

STEP 3: Finally, we extract information from the registration, namely the deformation w_t and deduce two masks. The first one P_t represents necrotic parts only (i.e.: voxels value close to 1) that will be used to orient the final deformation. The second one Q_t will take the value '1' for the whole tumour and will allow intensities changes. Finally, we register images using Constrained Metamorphosis.

In Figure 3.17 one can see quantitative results showing deformed landmark distance to their target. we see that CM gives better results than LDDMM and M for large lesions, while reducing the score variance, making it more predictable. We observe qualitatively good registrations for the three methods for small cancers as it was usually the case for classical methods (*e.g.:* Brett et al. [2001]; Sdika and Pelletier [2009]). However, the model presented is still mono-modal for the main part, the registrations of step three being computed only using the T1ce modality. Only the segmentation of step one was realised with all available modalities. Adapting the CM model to use all modalities for the final registration is possible and may improve the results.

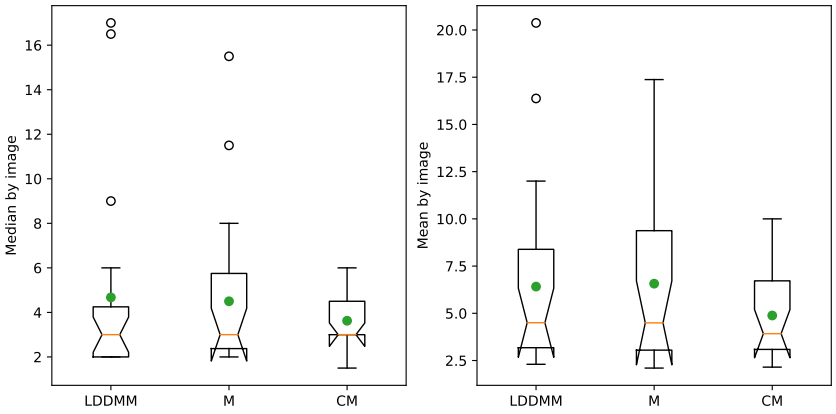

Figure 3.17: Landmark distance after registration on the validation dataset for LDDMM, Metamorphosis (M) and Constrained Metamorphosis (CM). (right) Median of individual landmarks distance. (left) The mean of individual landmarks distance. The orange line is the median and the green dot represents the mean. The y-axis measure unit is the pixel.

We are confident that our results can be further improved with more time. In particular, the mask construction would need more caution. As, because of the classical problem of entanglement in Metamorphosis François et al. [2022] it was hard to set the right parameters for all mask construction, some needing more deformation and some more intensity changes. This impact was high because of the strong dependence of CM on the priors. Furthermore, because of memory limitation on the GPU we had at disposition, we proceeded to make the registrations on images half the size of the ones provided (*i.e.:* `120x120x77`) and then re-interpolated the final image and deformation to the original size which may have reduced the accuracy of the registrations.

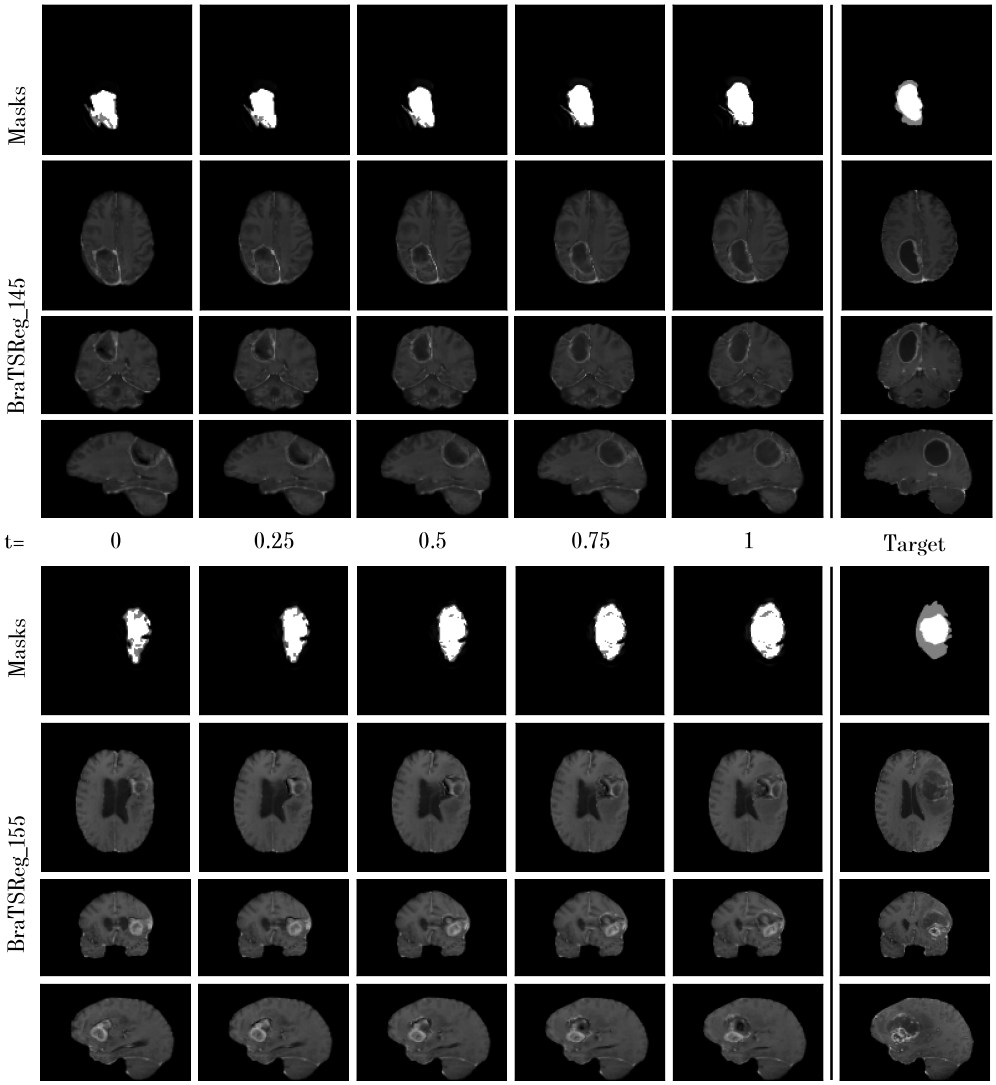

Figure 3.18: Qualitative result with Constrained Metamorphosis. Two examples from the validation data-set are shown. On first row we can see the temporal mask described in 1. Below one can see the Axial, coronal and Sagittal view of the registered brain. The last column is the Target image.

3.4 Conclusion

In this chapter, we have addressed the issues of inconsistent deformation observed in other methods, such as M and LDDMM. In the three conducted experiments, we have shown that Constrained Metamorphosis can improve registration outcomes while improving its deformation exploitability. Obviously, provided that appropriate priors are used. Thus it is critical creating accurate masks to achieve good results. With Metamorphosis we relied on the optimisation process to obtain both the ideal displacement and intensity changes from the images. While Constrained Metamorphosis (CM) offers the benefit of constraining the registration process with a model, it may also inherit from its limitations. In other words, the reliability of CM is dependent on the accuracy of the priors it follows. We have introduced parameters (*i.e.:* μ, γ) making CM more or less dependent on its model, which can be tuned in accordance with the confidence we have with it. However, they may be hard to choose for a whole data set. In more general terms, the need to tune parameters with CM can be considered a drawback, as is the case with any method with hyperparameters. The implementation of CM has been made by adapting the implementation of M, thus the discussion in Section 2.5 is also applicable in this specific case.

In this dissertation, we present mono-modal registrations only, however, we saw in Section 1.1.2 that each modality shows different tissues and it turns out it is hard to infer every tissue position from only one. We could also use cross-modality techniques aiming at augmenting an MRI by synthesizing images from one modality to another [Ge et al., 2019; Azad et al., 2022]. For example, a T1-weighted MRI can be translated to a T2-weighted MRI using a GAN, and the resulting images can be used to train a segmentation algorithm designed for T2-weighted images. Another lead is to extend the Metamorphic model to the multi-modal case, treating each modality as a different channel and deducing a common field for each. This is a work in progress in collaboration with Guillaume Sérieys from the MAP5 laboratory at *Université de Paris-Cité*.

Studies of longitudinal data sets have shown that during growth or ageing phenomena, the observed organisms are subject to transformations over time that is no longer diffeomorphic. One reason might be the gradual creation of new material. The evolution of the shape can then be described by the joint action of a deformation process and a material apparition process [Kaltenmark, 2016]. The work of Kaltenmark and Trouvé [2017] can serve as inspiration, as it seeks to strike a balance between accurately replicating growth dynamics and keeping diffeomorphisms simple for the purpose of comparing subjects.

4

TDA for glioma segmentation

Anatomical segmentation in MRI refers to the process of identifying and separating different structures within an MRI scan of the body. This process can be performed by a computer algorithm or a human operator using specialised software. The algorithm or operator segments the scan into different regions of interest, in our case within the brain, based on differences in image intensity, shape and size.

Accurate segmentation of glioblastomas is important for several reasons. Firstly, it enables medical professionals to make informed treatment decisions, such as the choice of surgical intervention or radiation therapy, by providing a clear understanding of the size and location of the tumour. Secondly, it is a valuable tool for monitoring disease progression and evaluating the effectiveness of treatment over time. Finally, glioblastoma segmentation is essential for the development of computer-aided diagnosis systems, which have the potential to significantly improve the accuracy and efficiency of many medical imaging algorithms. Among them, one can think of the need for prior data in Constrained Metamorphosis as we have seen in Chapter 3.

In this chapter, we present a method utilising tools from Topological Data Analysis (TDA) in order to perform glioblastoma segmentation, offering several advantages over traditional machine learning approaches. Unlike machine learning methods, which typically require large annotated data sets to train the model and may be prone to overfitting, our method is train-free and can be easily adapted to different data sets and segmentation needs. Additionally, TDA provides a more interpretable and stable framework for segmentation, as it leverages topological features to perform the segmentation in a series of simple and well-defined steps. These benefits, combined with the expertise of Raphaël Tinarrage from the FGV EMAp laboratory in Rio de Janeiro, make our approach an attractive option for glioblastoma segmentation. It should be noted that this project was completed without the direct involvement of my supervisors.

While working on the Constrained Metamorphosis (CM) framework for glioblastomas it seems logical to think about building this segmentation method. Partly because one could construct the CM prior mask (*i.e.:* P_t, see Section 3.1.2 for general definition and 3.3.1.c for a practical case.) before or during the registration optimisation. Additionally, by automating the segmentation process, we gain insight into the location, shape, and intensity of glioblastomas in MR images, which gives us a deeper understanding of glioblastoma characteristics. In fact, by having a robust segmentation method, we lay the foundation for building a diffeomorphic glioblastoma atlas. This atlas could then be used as a reference for future segmentation tasks and help to standardise the analysis of glioblastomas across different studies and institutions. Thus, the development of this method is not only beneficial for the Constrained Metamorphosis framework but also has the potential to make a significant impact in the field of medical imaging analysis. At the time of my thesis submission, this is an ongoing project and we present here our initial results.

In this chapter, we will first make a short review of methods doing glioblastoma segmentation using Machine Learning. Then we will cover some papers using TDA for MRI analysis. Then, in Subsection 4.2, we will make a short introduction to the different TDA concepts useful for analysing MRI with TDA. We will then be ready to present our segmentation algorithm in Subsection 4.3. Finally, we will present our results and discuss the improvement we want to implement in Subsection 4.4. The code for this project is fully available on GitHub at `https://github.com/antonfrancois/gliomaSegmentation_TDA` for anyone to use and contribute to.

4.1 Related works

IA-SEGMENTATION METHODS FOR GLIOMA: Glioma segmentation is a widely studied subject, and a multitude of methods have been developed, many of which are based on machine learning. Among these methods, at least two prominent categories can be identified: those based on *generative models* and those based on *discriminative models*. The former category of methods requires prior information, such as the shape or appearance of the tumour, and generally converges quickly. The latter category employs techniques such as random forests, support vector machines, conditional random fields and last but not least deep learning which typically exhibits higher accuracy [Zhang et al. [2021]; Zhao and Jia [2015]; Rucco et al. [2020]].

Let us cover a few representative articles on this topic. Islam et al. [2020] proposed an efficient multilevel segmentation method in four steps: They start with a common preprocessing/filtering phase, then cluster the MRIs using a k-mean, continue with an optimal threshold and watershed segmentation technique followed by a morphological operation to separate the tumour using CNN. Zhang et al. [2021] make boundaries of the tumour clear by a simple fusion of FLAIR and T2 images. Then they train dense 2D-CNN using novel architectures and loss function. Also, Liu et al. [2021] propose to solve the inter-class ambiguity problem in segmentation.

Last but not least, the eight selected teams of the BraTS2021 challenge (see Section 3.3.1.a) used deep-learning techniques for segmentation, the most efficient ones being the ones using nn-UNets architectures. The winning team had dice scores of 0.8835, 0.8878 and 0.9319 for the enhancing tumour (ET), the tumour core (TC), and the whole tumour (WT), respectively. The other teams got very similar scores (inter-study variance being of order 10^{-2}) [Luu and Park [2022]; Yuan [2022]; Futrega et al. [2022]; Siddiquee and Myronenko [2021]; Ma and Chen [2022]; Kotowski et al. [2022]; Ren et al. [2022]; Jia et al. [2022]].

TDA FOR MRI ANALYSIS: TDA is a field at the intersection of computational geometry, algebraic topology and data analysis. It aims at capturing relevant geometric and topological information from datasets. Since its emergence in the 2000s, it has been applied to a wide range of problems, from medicine, physics, computer vision and machine learning, among others, as listed in Oudot [2017]. The reader may consult Carlsson [2009] and Chazal and Michel [2021] for an introduction to TDA. In this subsection, we review how TDA has been used in the context of segmentation and analysis of MRI.

A first application of this theory consists in designing segmentation by *constraining their topology*. This idea has been used in Clough et al. [2019], in the context of segmentation from Cardiac Magnetic Resonance data. The authors train a CNN, using the usual loss computed from the DICE score, to segment parts of the image. In addition to this loss, they add a *topological loss*, calculated using TDA. Based on 'prior topological knowledge', the segmentation is constrained to be close to a pre-defined shape. For instance, in the problem of segmenting the myocardium of the left ventricle of the heart, the authors use the knowledge that the myocardium is ring-shaped. Similarly, they applied this idea to the problem of placenta segmentation, knowing that it forms one connected component with no holes [Clough et al., 2020]. This method was later extended to the problem of multi-class segmentation [Byrne et al., 2021].

A similar approach is used in Qaiser et al. [2016, 2019], in the context of segmentation of tumours of colorectal cancer, based on Hematoxylin and Eosin stained slides images. The

method consists in first cutting the image into patches, selecting relevant patches with the help of a CNN, and then determining which of these patches exhibit a tumour using TDA. The topology of the patches presents two characteristic profiles: ones that contain less connected components and holes (corresponding to infected tissues) and ones that contain more holes and connected components (corresponding to healthy tissues). These characteristics are computed using persistent homology and collected in the *persistent homology profile* curve. The differences between these topologies are explained by the fact that, in infected tissues, nuclei tend to have atypical characteristics, irregular shape and size.

From a second point of view, TDA can also be used as a *feature for other machine learning algorithms*. In this context, no topological prior is known, and TDA is seen as an exploratory tool. For instance, in Crawford et al. [2020], the authors compute the *Smooth Euler Characteristic Transform* of MRI of tumours. Using these features as the input of statistical models, they obtain an accurate prediction of clinical outcomes. In the same vein, Saadat-Yazdi et al. [2021] compute the *Betti curves* of brain MRI, and show that they can be used as a predictor of the presence for Alzheimer's disease. As a last example, TDA has been used in Rucco et al. [2020] for the analysis of glioblastoma using FLAIR modality MRIs. After selecting two-dimensional slices, the authors compute various topological features, such as *Euler characteristics*, *persistent entropy* and *generator entropy*. These features are given to a machine learning classifier, allowing personal diagnosis and comparison of the tumour pre- and post-treatment.

We stress that persistent homology is not only defined for images, but also for different structures, such as graphs or point clouds. An example of such a construction is described in Lee et al. [2012], where the authors use PET scans to build a graph whose nodes are 'regions of interest' of the brain, and whose edges are selected according to the correlation of the FDG-PET emissions of the regions. It is shown that the persistent homology of these graphs allows to distinguish patients with attention-deficit hyperactivity disorder, autism disorder and control subjects.

Last, it must be mentioned that applications of TDA also inspire *theoretical developments*. An example is given by the introduction of Decorated Merge Trees in Curry et al. [2022]. This notion, which can be seen as a refinement of the persistence diagrams, allows to track the evolution of the topology of images, but also their localisation. They exemplify their method on segmentation of glioblastoma from MRIs.

4.2 TDA for images

The aim of this section is to introduce *persistent homology*, the most popular technique of TDA. It is a theoretical framework that allows inferring the *homology groups* of a dataset [Edelsbrunner et al., 2000; Zomorodian and Carlsson, 2005; Niyogi et al., 2008]. In Subsection 4.2.1, we will briefly introduce some notions of homology, without getting too much into the details. We refer the reader to Hatcher [2000] for a thorough presentation of homology theory and algebraic topology.

Among the several variations of persistent homology, *cubical persistent homology* is particularly adapted to deal with images. The construction is based on the notion of *filtration*, which we will present in Subsection 4.2.2. From such a filtration, one defines a *persistence diagram*, the main object of TDA, that we will introduce in Subsection 4.2.3.

4.2.1 Singular homology

Abstractly, homology is a *functor* from the category of topological spaces to the category of vector spaces. In particular, it is an operator that transforms topological spaces into vector spaces. This way, we are able to convert topological problems into linear problems, which we hope are easier to solve.

Many theories of homology exist, and we will present here a particular one: the singular homology with coefficients in $\mathbb{Z}/2\mathbb{Z}$. Here, by $\mathbb{Z}/2\mathbb{Z}$, we mean the finite field with two elements, also denoted $(\{0,1\},+,\times)$. In what follows, and when there is no ambiguity, we may say 'homology' instead of 'singular homology over $\mathbb{Z}/2\mathbb{Z}$'.

Homology allows to associate to any topological space X a sequence of $\mathbb{Z}/2\mathbb{Z}$-vector spaces $H_0(X)$, $H_1(X)$, $H_2(X)$, ..., called *singular homology groups*. The construction of these vector spaces is based on a combination of topology and algebra that we do not present here. Instead, we will explain roughly what they represent in the next paragraphs.

It is worth noting that linear algebra over the field $\mathbb{Z}/2\mathbb{Z}$ is similar to that over $\mathbb{R}$: the notions of independent vectors, spanning sets and bases are the same. In particular, each vector space over $\mathbb{Z}/2\mathbb{Z}$ admits a dimension (potentially infinite). If such a vector space has a finite dimension equal to d, then it is isomorphic to the product vector space $(\mathbb{Z}/2\mathbb{Z})^d$.

We now turn to the interpretation of homology groups. Let $i \in \mathbb{N}$, and consider the i^{th} homology group $H_i(X)$. If it is finitely generated, which we will suppose, then it admits a finite dimension, denoted $d(i)$. This integer $d(i)$ is also called the i^{th} *Betti number* of X. It carries topological information about X:

- $d(0)$, the dimension of $H_0(X)$, is equal to the number of connected components of X.
- $d(1)$, the dimension of $H_1(X)$, is equal to the number of 'independent loops' in X. For instance, the circle has $H_1(X) = \mathbb{Z}/2\mathbb{Z}$, and the torus $H_1(X) = (\mathbb{Z}/2\mathbb{Z})^2$ (see Figure 4.1)
- $d(2)$, the dimension of $H_2(X)$, is equal to the number of 'independent voids' in X. For instance, the sphere and the torus have $H_2(X) = \mathbb{Z}/2\mathbb{Z}$.

Except for dimension 0, this list should not be treated as a formal mathematical result, but only as heuristic interpretations.

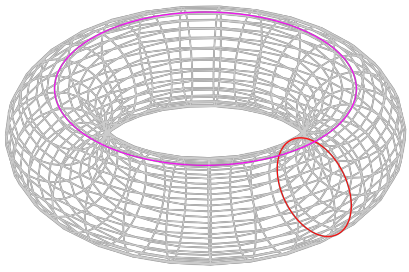

Figure 4.1: The torus is made of one connected component, hence $H_0(X) = \mathbb{Z}/2\mathbb{Z}$. Besides, it contains two independent loops, hence $H_1(X) = (\mathbb{Z}/2\mathbb{Z})^2$. Last, it has one void, so $H_2(X) = \mathbb{Z}/2\mathbb{Z}$.

In order to use homology in the context of images, one has to transform images into topological spaces. This is easily done when the image is *binary*, that is when it only has

black and white pixels. Let $I : \Omega \to \{0, 1\}$ be a binary image with domain Ω, and consider the collection $I^{-1}(\{1\}) \subset \Omega$ of its black pixels. Seen as a subset of Ω, this collection can be seen as a topological space. Hence homology is well-defined.

As an example, we represent in Figure 4.2 a 2-dimensional image, whose homology groups are $H_0(X) = (\mathbb{Z}/2\mathbb{Z})^3$ and $H_1(X) = (\mathbb{Z}/2\mathbb{Z})^2$.

When X is a topological space obtained from a binary image, as described in the previous paragraph, another theory of homology may be used: *cubical homology*. In the same way, that singular homology transforms topological spaces into vector spaces, cubical homology transforms binary images into vector spaces. Moreover, if the topological space comes from a binary image, then their singular and cubical homology groups are equal. The advantage of cubical homology is that it can be directly computed: the algorithm boils down to Gaussian elimination. This has been implemented in several Python libraries, such as `giotto-tda`, `cubical Ripser` or `Gudhi` [Tauzin et al., 2021; Kaji et al., 2020; Maria et al., 2014].

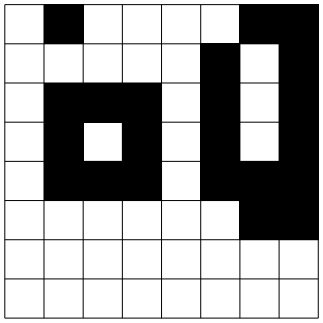

Figure 4.2: A 2d binary image, consisting of three connected components and two 'holes'.

4.2.2 Filtrations

The aim of this work is to analyse brain MRIs. These are 3-dimensional greyscale images. They are not binary images, and there is no obvious way to convert them canonically into topological spaces. Therefore, homology theory cannot be used directly. In order to circumvent this issue, we can follow the pipeline of persistent homology. The idea consists in building not one but a collection of topological spaces, called a *filtration*. More precisely, a filtration of images is an increasing family of binary images, indexed by a parameter $t \in [0, 1]$. Such a filtration will be denoted $\{I^t \mid t \in [0, 1]\}$.

Let $I \colon \Omega \to [0, 1]$ denote a greyscale image. There exist various popular filtrations that can be defined from I, such as the height, radial and density filtrations, implemented in GIOTTO-TDA, or the sublevel and superlevel sets filtrations, implemented in CUBICAL RIPSER and GUDHI [Tauzin et al., 2021; Kaji et al., 2020; Maria et al., 2014]. In this work, we will focus on these two last filtrations.

For any $t \in [0, 1]$, let us define I^t as the set of pixels with intensity lower or equal to t. The set I^t is a topological space—a union of cubes—, and we have the relation $I^s \subset I^t$ for any $s, t \in [0, 1]$ such that $s \leq t$. In other words, the family $\{I^t \mid t \in [0, 1]\}$ is an increasing sequence of binary images. It is called the *sublevel sets filtration*. Note that, since the image has values in $[0, 1]$, the parameter t parses all the possible values a pixel can take, and the last image, I^1, is equal to the whole domain Ω. An example is given in Figure 4.3a, where I is a brain MRI, the T2 modality of the SRI template [Torsten et al., 2010]. The SRI template has already been introduced in Subsection 3.3.1.a and is represented in Figure 3.11. Similarly, if we define I^t as the set of pixels with intensity greater or equal to t, then we obtain the *superlevel sets filtration* of I. It is represented in Figure 4.3b.

Within a filtration, one can understand the parameter t as a temporal value. The more t increases, the more pixels are added to the image. In the case of the sublevel sets filtration, pixels of low intensity are added first. In Figure 4.3a, one sees that the first pixels added

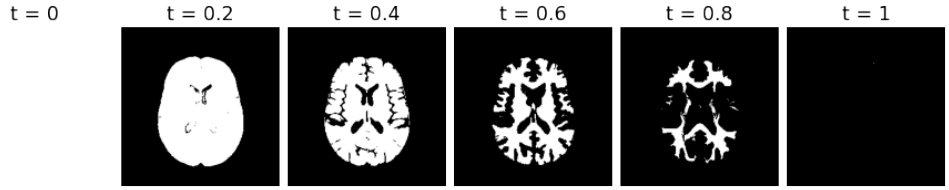

(a) Sublevel sets filtration

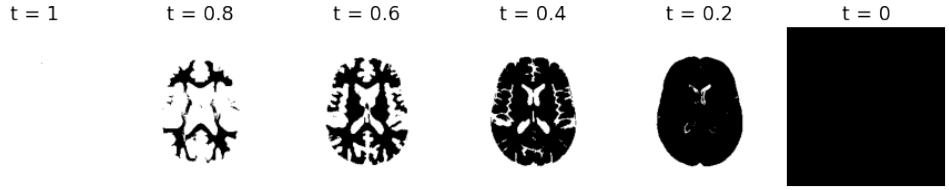

(b) Superlevel sets filtration

Figure 4.3: Examples of sublevel (a) and superlevel (b) sets filtrations on the SRI brain template MRI.

are the ones in the background (filtration value $t = 0$).

On the contrary, in the superlevel sets filtration, pixels of high intensity are added first. In figure 4.3b, we see that this corresponds to the ventricles and the grey matter. In other words, the most luminous parts of the image appear the earliest in filtration. This idea will be used in Subsection 4.3.2 when we will propose a segmentation method for glioblastoma. Indeed, on a FLAIR modality MRI, glioblastoma tends to be represented by the pixels of the highest intensity. That is to say, we expect the tumour to be the first element to appear in the filtration. An example is given in Figure 4.4, for a FLAIR modality MRI of a brain presenting a tumour.

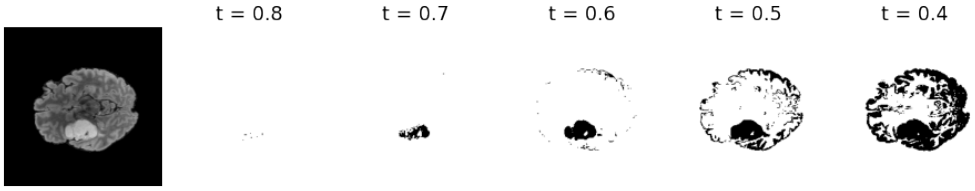

Figure 4.4: Superlevel sets filtration on a FLAIR modality MRI of a brain with tumour.

4.2.3 Persistence diagrams

We can now define the main objects of TDA: the persistence modules and their corresponding persistence diagrams. These objects allow summarising the information contained in a filtration, seen from a topological point of view.

Let $\{I^t \mid t \in [0,1]\}$ be a filtration, built from the image I, as defined in the previous subsection. The filtration is a family of greyscale images, hence we can apply cubical (or simplicial) homology to each of the I^t's. Let us fix a $i \in \mathbb{N}$, and apply the i^{th} homology functor to the filtration. This yields a family of vector spaces $\{H_i(I^t) \mid t \in [0,1]\}$. This structure is called a *persistence module*. It is a purely algebraic object, which gathers the

homology groups of I at various scales.

As an example, let us consider the superlevel sets filtration. The persistence module $\{H_0(I^t) \mid t \in [0,1]\}$ gathers the number of connected components of all the superlevel sets of I, as represented in Figure 4.5. As one can read, at the beginning of the filtration, the images consist of many connected components, therefore the corresponding homology groups $H_0(I^t)$ have large dimension. At some point in the filtration, all the components merged together, and only one big component exists.

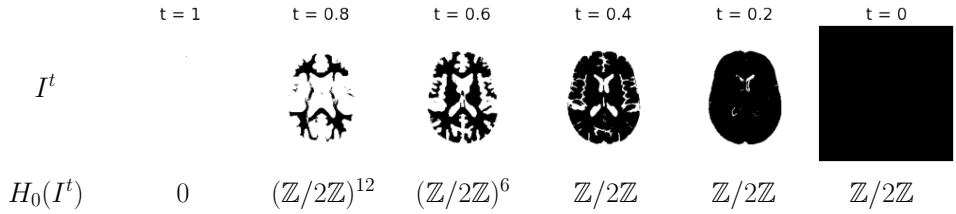

Figure 4.5: A superlevel sets filtration, and the corresponding H_0-persistence module

Actually, in a persistence module, we can extract more information than just the homology groups $H_0(I^t)$. Using the inclusion maps $I^s \hookrightarrow I^t$ for $s \leq t$, we can 'track' the evolution of the homological features. That is, we are able to tell whether a cycle of $H_i(I^s)$ is still alive in $H_i(I^t)$. The interval on which a feature exists is called its *persistence*. In practice, one interprets cycles of large persistence as relevant features of the dataset, while cycles of low persistence are considered as noise. The formal definition of this idea is based on the property of *functoriality* of the homology, which we won't explain here.

The persistence of all the cycles is recorded in the *persistence diagram* of the persistence module. It is a set of points $\mathcal{P}$ of the form $p = (t_b, t_d)$, with $t_b \leq t_d$. A point of the persistence diagram is interpreted as a homological feature, born at time t_b and dead at time t_d.

In the case of images, to each point $p = (t_b, t_d)$ of the persistence diagram corresponds a *birth* pixel p_b, which gives birth to a cycle (a new connected component, a new H_1 cycle, *etc.*), and a *death* pixel p_d, which kills the cycle (merge the component to another one, fill the H_1 cycle, *etc.*). In particular, in the H_0 persistence diagram, there always is a point that dies at infinity. It represents the connected component consisting of all pixels.

In Figure 4.6, one can see the persistence diagrams of the SRI template, for the two filtrations defined above: superlevel and sublevel sets. As shown in Garin et al. [2020], there exists a duality between these two diagrams. That is, one can be transformed into the other, via an explicit process. As a consequence, we can choose only one of the filtrations, and still have the same information. In what follows, we will choose to work with the superlevel sets filtration, since it is easier to interpret.

4.2.4 Persistent homology of the SRI template

As presented in the previous subsection, the persistence diagram of the filtration of an image gathers the evolution of its homology, while the thresholding parameter t varies. Each point $p = (t_b, t_d)$ of the diagram represents a cycle, born at t_b and killed at t_d. The quantity $t_d - t_b$, called the *persistence* of the cycle, is understood as the duration of the cycle. An important interpretation is that, given persistence diagram, points with high persistence

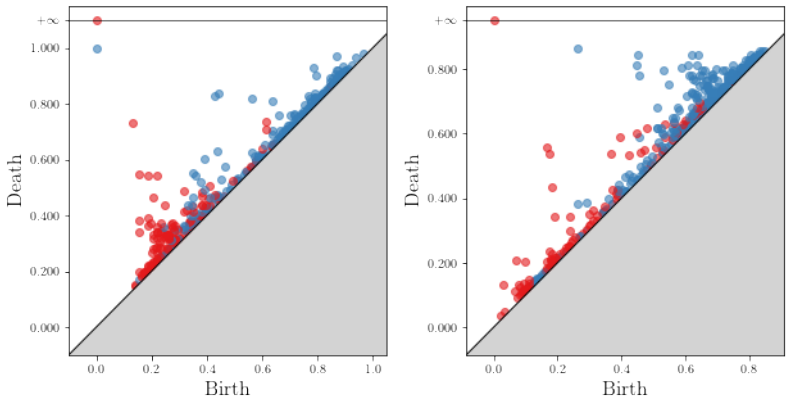

Figure 4.6: The persistence diagram of the sublevel sets filtration of the SRI template (left), and the diagram of its superlevel sets filtration (right). H_0 is represented in red, and H_1 in blue.

correspond to relevant topological features of the image, while points with small persistence can be considered as 'topological noise'. Visually, points with high persistence are those far away from the diagonal.

As an illustration, let us consider the SRI template. We choose a slice of the image and compute the persistence of its superlevel sets filtration. We stress that, in the following subsections, we will consider the persistent homology of the full 3-dimensional image, and not only of its slices. We choose here to analyse a 2-dimensional slice for visualisation ease.

We first study its persistence diagram in H_0, that is, the evolution of its connected component. As shown in Figure 4.7, three red points seem significantly distant from the diagonal. They correspond to three connected components, that evolve independently, without merging with others. For each of these three persistent cycles, we consider the points $p = (t_b, t_d)$ of the diagram, extract the pixel of birth, and plot the connected component of the pixel at time t_d. It represents the connected component, just before it merges with another one. As we can see in the figure, these three main components are part of the white matter, disconnected in this slice.

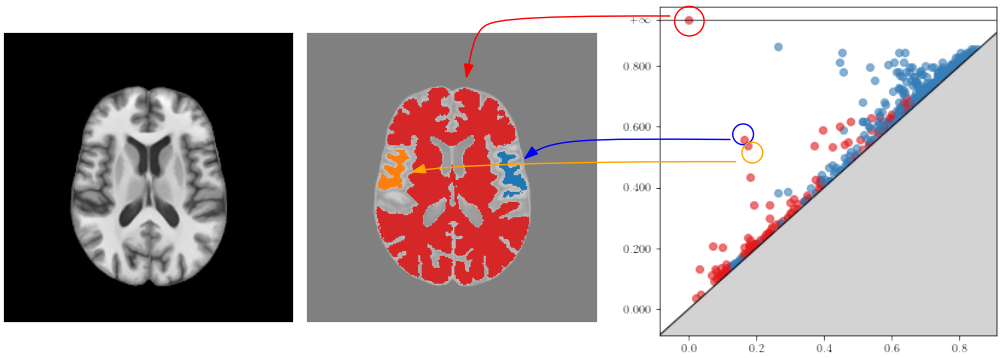

Figure 4.7: Slice of the SRI template (left), the persistence diagram of its superlevel sets filtration (right), and its most persistent H_0-cycles (middle).

We now turn to the analysis of the persistence diagram in H_1, that is, the holes. As we can see on the persistence diagram, four points seem away from the diagonal, with one point particularly off. In order to represent the corresponding holes, we circled them in Figure 4.8. We see that they correspond to the lateral and third ventricles (again disconnected in this slice).

It is worth mentioning that a particular difficulty arises when considering homology groups H_i with $i \geq 1$: the cycles are not uniquely localised. That is, several sets of pixels may represent the same cycle. In the figure, our circles are arbitrary representatives of these cycles. The problem of finding relevant representatives of persistent cycles, for instance, cycles with minimal length or volume, is an active topic of research in TDA, but we will not be concerned with this problem in the following [Escolar and Hiraoka, 2016; Obayashi, 2018].

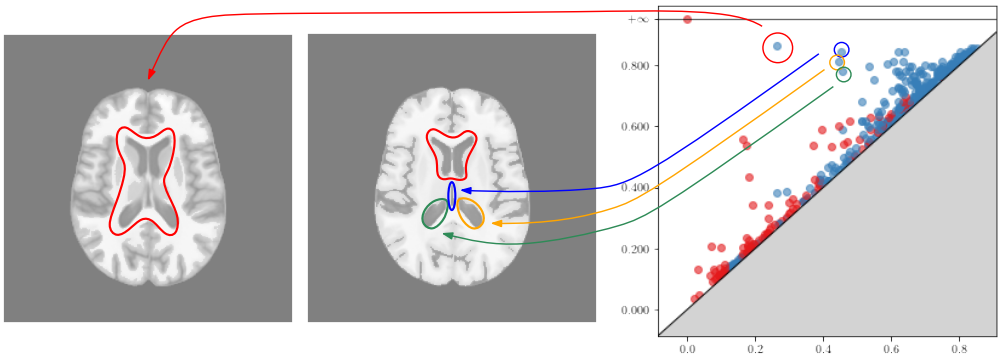

Figure 4.8: Slice of the SRI template (left and middle) with the most persistent H_1-cycles, and the persistence diagram of its superlevel sets filtration.

4.3 A segmentation method based on TDA

4.3.1 Formulation of the problem and description of the method

The problem we consider can be formulated as follows: given an MRI of a brain with glioblastoma, we wish to segment the tumour into its different components. More precisely, we have access to two modalities of the MRI, FLAIR and T1ce, as introduced in Chapter 1.1.2. The three tumoural components are: Enhancing Tumour (ET), Tumorous Core (TC) and oedema (WT: whole tumour), as described in Sections 1.1.2 and 3.3.1.a. All along this section, we will use as running examples a few MRIs from the collection BRATS2021.

NOTATIONS: Let I_{FLAIR} and $I_{\text{T1ce}} \colon \Omega \to [0,1]$ denote respectively the FLAIR and T1ce modalities. The three components of the segmentation are denoted X_{ET}, X_{TC} and X_{WT}. They are subsets of Ω. By the *global segmentation*, or *union*, we refer to the subset $X = X_{\text{ET}} \cup X_{\text{TC}} \cup X_{\text{WT}}$.

MODEL: The algorithm we propose is based on TDA, hence it is worth giving the topological insights which support the method. We make the assumption that the segmentation of the tumour, as given by the specialists, satisfies the following hypotheses:

1. The global segmentation X is made of one connected component and appears significantly more luminous than the other voxels in the modality FLAIR.
2. The component X_{ET} is homotopy equivalent to a sphere and appears significantly luminous in the modality T1ce.
3. The component X_{TC} is in the interior of the spherical shape X_{ET}.
4. The component X_{WT} is in the exterior of the spherical shape X_{ET}.

This model, although rather simplistic, can be observed on several MRIs, as shown in the examples of Figure 4.9. In Subsection 4.4.1, we will quantify, via the formulation of a more formal model, how many images fall under these assumptions.

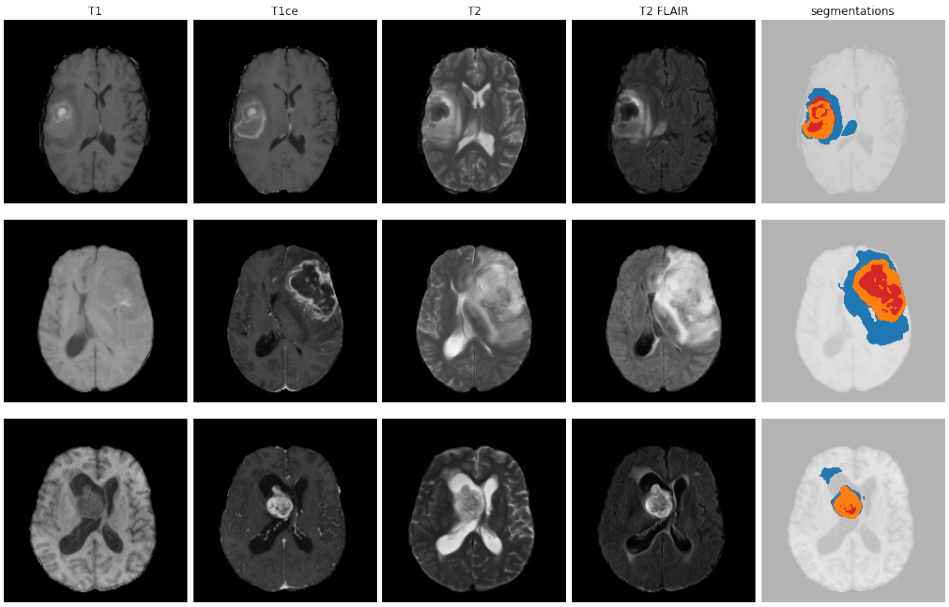

Figure 4.9: Glioblastomas examples in four modalities Each row contains MR Images slices in four modalities of the same patient at the same time. One can see both the tendencies between different image modalities and the inter-subject colour variation within the same modalities. Segmentation legends: Red - Necrotic Core (TC); Orange - Enhancing Tumour (ET); Blue - Oedemas (WT).

GLOBAL IDEA: Having in mind the model, we present a three-step method utilising the FLAIR and T1ce modality for glioma segmentation.

Step 1. Tumour selection: Following Hypothesis 1, the FLAIR modality shows a hyperintense area corresponding to the tumoural activity. This step consists in selecting this whole subset of voxels. We consider the superlevel set filtration of the image, starting at $t = 1$ and select a threshold at which a connected component in the filtration I^t is equal to

the tumour. The choice of this threshold is based on an analysis of the number of voxels contained in each connected components.

Having selected a connected component, we propose two refinement procedures. First, we fill the holes of the estimated segmentation that may appear because the necrosis is often darker on FLAIR images. Secondly, we improve the segmentation by checking whether the estimated tumour is matching the image edges. These procedures are detailed in paragraphs 'Refinement' of Subsection 4.3.2.

Step 2. Identification of ET: Following Hypothesis 2, the ET component of the tumour forms a sphere in the T1ce image. Hence, we estimate ET as the brightest H_2 feature in the persistent homology of this image. More precisely, we compute the persistent homology of the image restricted to the whole segmentation X obtained above, since we seek ET in X.

Step 3. Identification of TC and WT: In order to obtain the segmentation of TC and WT, we consider the global segmentation X, from which we remove ET. This binary image can be partitioned into connected components. We identify the components of TC as those in the interior of ET and the components of WT as those in its exterior (as indicates by Hypotheses 3 and 4).

4.3.2 Step 1: Segmentation of the whole tumour

FIRST ESTIMATION: In this step, our goal is to select the largest hyper-intense region present in the FLAIR image. To achieve this, we analyse the number of voxels in the filtration process over time, starting from $t = 1$ and moving to $t = 0$. While the number of voxels increases steadily, we anticipate a sharp increase in the number of voxels when the voxels corresponding to the white and grey matter are included. Therefore, we select the value of t just before this sharp increase. To do this, we examine the derivative of the number of voxels with respect to t, and choose the first derivative value that exceeds a predefined threshold. This process is similar to a binarisation algorithm, such as Otsu's method. Once we have selected the optimal value of t, we then identify the largest connected component, which corresponds to the region with the most voxels.

The core element of this step is therefore the suggestion of an optimal t implemented in the function `suggest_t` that one can find in Code 4.1, itself split into three parts:

1. We first slightly blur the image. For any $t \in [0, 1]$, we define I^t_* as the set of voxels forming the biggest connected component. This notation is chosen because for all $t, I^t_* \subset I^t$.

2. We study the number of voxels of I^t_*, denoted $t \mapsto \#I^t_*$, and its derivative $t \mapsto \mathrm{d}\#I^t_*$, computed by finite differences. The map $t \mapsto \#I^t_*$ is monotonically increasing while t decreases. We expect it to behave roughly like a sigmoid function, having a sudden rise at the moment a big amount of medium-valued voxels are added (*c.f.:* Figure 4.10 plots). In practice, The derivative is computed by finite difference. In this form, because of the image intensity range or the chosen time step, the choice of `dt_thershold` over a database is not obvious. We solve the problem with the normalisation found in code 4.1, l.17-18.

3. Varying t from 1 to 0 we search for the first significant jump of the filtration by selecting the first t such that $\mathrm{d}\#I^t >$`dt_thershold`.

Source Code 4.1: suggest_t

```
def suggest_t(img,pos=None, N= 25,dt_threshold=1):
        """

        :param img: Image
        :param pos: Position of the hightest pixel.
        :param N: Number of `t` to test
        :param dt_threshold: First gap threshold.
        :return:
        """
        tmax = img.max()  if pos is None else img[pos[0],pos[1],pos[2]]
        t_list = np.linspace(0.01,tmax,N)
        filtr = np.zeros(t_list.shape)
        for i,t in enumerate(t_list):
                cc = get_highest_connectedComponent(img,t)

        filtr[i] = (cc==1).sum()
        filtr_dt = (filtr[:-1] - filtr[1:])*(N/(tmax-0.01))
        filtr_dt_norm = len(filtr_dt)*filtr_dt/filtr_dt.sum()

        best_t = -1
        index = N-2
        while best_t < 0:
                dt = filtr_dt_norm[index]
                if dt > dt_threshold:
                        best_t = t_list[index+2]
                        break
                else:
                        index -= 1

        return best_t,filtr,filtr_dt,t_list
```

In Figure 4.10 one can observe the `suggest_t` procedure, comparing the filtration cardinal and its derivative with the t selected, alongside the corresponding brightest components. On panel 4.10.a the selected component is very close to the ground truth. On panel 4.10.b, however, one can see that the process should have selected a slightly higher filter (t too low) to fit better to the true segmentation.

REFINEMENT - STICK TO EDGES: In the previous paragraph, we saw that t depends on a threshold (*c.f.:* in Code 4.1). It happens that t is chosen a bit off, and the segmentation 'leaks' on non-cancerous tissues. To correct this issue, one can slide t to make the edges of the segmentation match the edges of the brain. By edges, we refer to boundaries between different objects or areas of an image that have a significant change in intensity or colour.

We observed that a simple 3D Sobel filter is not very handy on MRI images for finding edges. Therefore, we use the procedure described in Rorden et al. [2022], based on the second derivative (difference of Gaussian) of the image to generate edges.

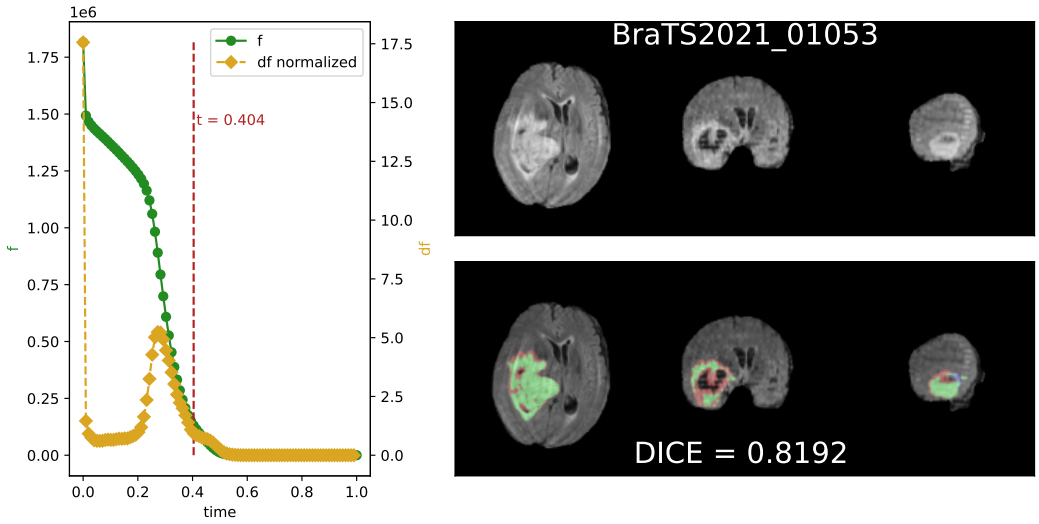

Figure 4.10: Automatic t selection by the function `suggest_t`. (left) plots with $I^t(\tilde{x})$ in green and its derivative in yellow, their axes are respectively located on the left and right sides. (centre) Brain slice to segment. (right) Superposition of the estimated segmentation with the ground truth. Green areas correspond to a correct match, Red areas are where our segmentation selected unwanted pixels, and Blue (none) are the ground truth not selected.

In order to evaluate the Edge match, we create a metric by utilizing a DICE score to measure the degree of overlap. Let I represent the image, and let X denote the estimated segmentation. Let us also assume that there exists a function E which takes an image as input and produces a binary output image with ones representing only the edges of the original image. Then, we define edgeDice as follows:

$$\text{edgeDice}(I, X) = \text{DICE}(E(I) \times E(X), E(X)) \tag{4.1}$$

Next, we evaluate the `edgeDice` metric by testing values within a narrow range of the suggested t value. This approach ensures that the metric is not influenced by other edges, such as those belonging to other organs or the skull, which could be present outside of the narrow range.

Note that if we found this method promising, this technique does not improve the results systematically yet. Indeed the identification of relevant edges turned out to be challenging to handle on 3D MRI and needs to be fine-tuned for the data set.

Source Code 4.2: edgesEnhancer

```
def respect_edges(img_flair,true_seg,t_suggested,pos,
                  n_test = 20,ovrlap=.1, plot= False,ax= None,edgeDice=None):
    edgesDICE_stock = np.zeros(n_test)

    t_list = np.linspace(t_suggested - ovrlap, t_suggested + ovrlap,n_test)

    best_ed = -1
    for i,t in enumerate(t_list):
        # find the component at a given position
```

```
        seg_union = sTDA.getConnectedComponent(img_flair,pos,t)
        # compute edgeDice
        edgesDICE_stock[i]= edgeDice(img_flair,seg_union)

    # select the new best t
    new_t_idx = np.argmax(edgesDICE_stock)
    new_t = t_list[new_t_idx]

    return new_t,best_seg,ax
```

REFINEMENT - CLOSING HOLES: Given a segmentation of the global tumour X obtained by the method described above, we perform a post-processing step. It consists in filling the holes in X, following Hypothesis 1 of our model. This is done by computing the connected components of the binary image $\Omega \backslash X$, identifying the components that are not the background , and adding them to X. More precisely, we apply the function `ndimage.measurements.label` of `scipy` to the binary image $\Omega \setminus X$. As an output, each voxel is associated an integer, corresponding to the index of the connected component it belongs to. We identify the connected component of the background as the one containing the most voxels. An explicit implementation is given in Code 4.3.

Source Code 4.3: filling_holes

```
    ' Fill the holes of seg_union '
    # Define the complementary segmentation
    seg_remaining = (seg_union>0)*1-1
    seg_remaining[seg_remaining<0]=1

    # Classify the complementary segmentation into connected components
    labels = skimage.measure.label(seg_remaining, background=0)
    components = [(labels==i)*1 for i in range(1,np.max(labels)+1)]

    if len(components)>0:
        # Identify the non-background componnents
        components_cardinal = [np.sum(component) for component in components]
        cardinalmax = max(components_cardinal)
        remainingcomponents = np.sum([components[i] for i in range(len(components))
                                      if components_cardinal[i]<cardinalmax],0)

        # Add the non-background components to seg_union
        seg_union[remainingcomponents>0] = 1
```

4.3.3 Step 2: Identification of ET

In the previous step, we estimated the global segmentation $X \subset \Omega$, based on the image I_{FLAIR}. Note that, according to our model, we can decompose X as $X_{\text{ET}} \cup X_{\text{TC}} \cup X_{\text{WT}}$, although we do not know the components at this point. In this second step, we will use the image I_{T1ce}, and the previously estimated segmentation X, to obtain X_{ET}.

According to the model, the enhancing tumour X_{ET} is the boundary of the tumour and is highly intense in the image. Hence, using a superlevel sets filtration, we should see in the H_2 persistence diagram a persistent cycle, representing the sphere formed by the boundary of the tumour. More precisely, we compute the persistent homology of the superlevel sets filtration on the image I_{T1ce} restricted to X. This is represented in Figure 4.11a, in the case of our running example (see Figure 4.9). One green point appears particularly far away from the diagonal: it represents the persistent cycle we are looking for. In order to automatise this procedure, we follow two steps:

1. We compute the persistence diagram of the superlevel sets filtration of the image I_{T1ce} restricted to X, and identify the H_2-point of highest persistence.
2. We extract the pixels that form this cycle as follows. Let t_b be the birth value of this persistent cycle, and x the corresponding pixel that gave birth to it. We then define X_{ET} as the connected component of the pixel x in the binary image $I^{t_b}_{\text{T1ce}}$.

The resulting segmentation can be seen on Figure 4.11b. An implementation is given in Code 4.4.

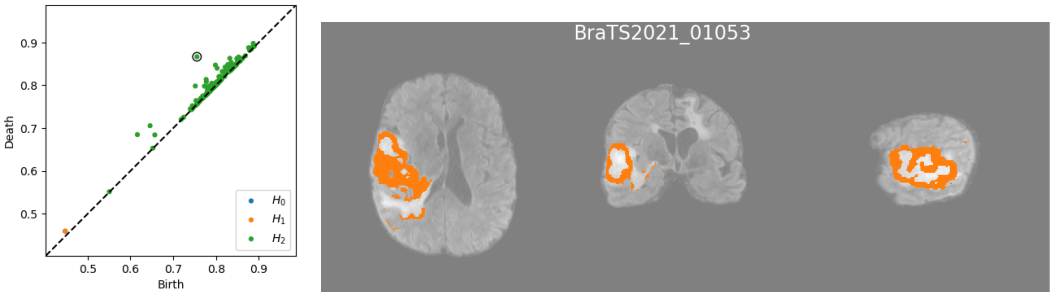

(a) **(b)**

Figure 4.11: Persistence diagram of I_{T1ce} restricted to X, with the most persistent point circle (a) and the resulting segmentation of X_{ET} (b).

Source Code 4.4: Step 2: Estimation of X_{ET}

```
    def GetConnectedComponent(img, pos, t):
        '''
        Get the connected component of the voxel pos = (x,y,z) at time t.
        The output is a binary image.
        Background value of img must be 0 (as conventional).
        '''
        imt = (img>=t)*1
        if imt[pos[0],pos[1],pos[2]]==0:
            raise ValueError('The voxel pos is not active at time t.')
        labels = skimage.measure.label(imt, background=0)
        labeltumor = labels[pos[0],pos[1],pos[2]]
        imtumor = (labels == labeltumor)*1
        return imtumor

    # Compute persistence diagram with cripser
```

```
    seg_union_t1ce = img_t1ce*seg_union
    barcode = cripser.computePH(1-seg_union_t1ce,maxdim=3)

    # Gather the non-infinite points
    H2 = [list(bar[1::]) for bar in barcode if bar[0]==2 and bar[2]<1]

    # Sort the points and select the most persistent one
    H2 = [bar for _,bar in sorted(zip([bar[1]-bar[0] for bar in H2],H2))[::-1]]
    point = H2[0]

    # Extract contour H2
    pos = np.array(point[2:5]).astype(int)
    t = point[0]+0.0001
    seg_contour = GetConnectedComponent(seg_union_t1ce, pos, 1-t)
```

4.3.4 Step 3: Identification of TC and WT

At this point, we have estimated the segmentation of the global tumour, X, and the enhancing tumour X_{ET}. We now identify the components X_{TC} and X_{WT}. This step does not depends on the initial MRI, but only on the subsets X and X_{ET}.

Following our model, TC corresponds to the part of the tumour that lies *inside* ET, and WT to the part that lies *outside* ET (and still within X). In order to identify these parts, we apply the following procedure.

1. We consider the subset $X \setminus X_{\mathrm{ET}}$, the complementary of ET in the segmentation, and compute its connected components.
2. Note that we may have more than two connected components. For each of them, we compute whether it is 'inside' or 'outside' the contour, according to the following rule: it is 'inside' if most of its neighbouring pixels (obtained by dilatation) belong to X_{ET}.
3. Connected components that are 'inside' are added to X_{TC}, and those that are 'outside' are added to X_{WT}.

As a convention, we save the segmentation as an image $\Omega \to \{0, 1, 2, 4\}$, where 0 represents the background, 1 represents TC, 2 represents WT and 4 represents ET. The resulting segmentation can be seen in Figure 4.12. An implementation is given in Code 4.5.

Source Code 4.5: Step 3: Identification of X_{TC} and X_{WT}

```
    # Identify connected components of the complement with skimage.measure.label
    seg_union_nocontour = seg_union - seg_contour
    labels = skimage.measure.label(seg_union_nocontour, background=0)
    components = [(labels==i)*1 for i in range(1,np.max(labels))]

    # Classify components: TC or WT
    seg_final = seg_contour*4 # define seg_final
    for component in components:
```

Figure 4.12: Segmentation after step 3. WT is represented blue, ET in orange and TC in red.

```
        componentdilated = scipy.ndimage.binary_dilation(component,iterations=1)
        componentcontour = componentdilated - component
        meanvalue = np.mean(seg_contour[np.where(componentcontour>0)])
        if meanvalue>=1/2:
            # if most of the boundary points are in ET, then the component is in TC
            seg_final[component>0] = 1
        else:
            # otherwise, it is in TC
            seg_final[component>0] = 2
```

4.4 Results & Discussion

4.4.1 Validation on BRATS2021

In order to quantify the performance of the algorithm, we applied it to the whole collection of MRIs of BRATS 2021, consisting of 1250 images. For each image, we computed the DICE score for each segmentation class (union, TC, WT and ET), between the estimation of our algorithm and the segmentations provided by the specialists. The results are gathered in Figure 4.13. Some examples of segmentations with a high score are represented in Figure 4.15, and with a low score in Figure 4.16.

By inspecting Figure 4.13, We observe that the results seem correct for the whole segmentation (median approximately 0.8) but relatively poor for TC, WT and ET (median approximately 0.2, 0.2 and 0.4). One can take a closer look at the images where the algorithm performs poorly in Figure 4.16, which reveals low scores are mainly attained by images that do not satisfy the model we introduced in Subsection 4.3.1. We observed that either the enhancing tumour does not surround the necrosis, or it does but form a perforated sphere. In both cases, the algorithm cannot partition the domain into the interior and exterior of ET, leading subsequently to an incorrect estimation of the other components of the segmentation.

In order to evaluate the importance of the model, we restricted the analysis to only a subset of images, the ones that satisfy Hypothesis 3 of the model described in Subsection 4.3.1. More precisely, we consider the images such that applying a binary dilatation of X_{TC}

yields new pixels of which at least half are included in the image. We computed that 298 out of 1250 images satisfy this assumption. We computed the DICE scores of the segmentations we obtained on these images and gathered the results in Figure 4.14. One sees that, in this case, segmentations are of higher quality: the median score for the whole segmentation, TC, WT and ET are approximately and respectively 0.8, 0.5, 0.6 and 0.7.

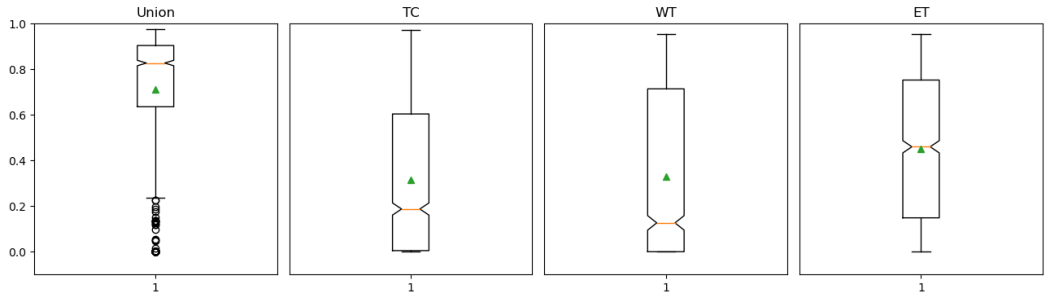

Figure 4.13: For the 1250 MRI of BRATS 2021, we plot the Dice scores of the union segmentation, the TC, WT and ET.

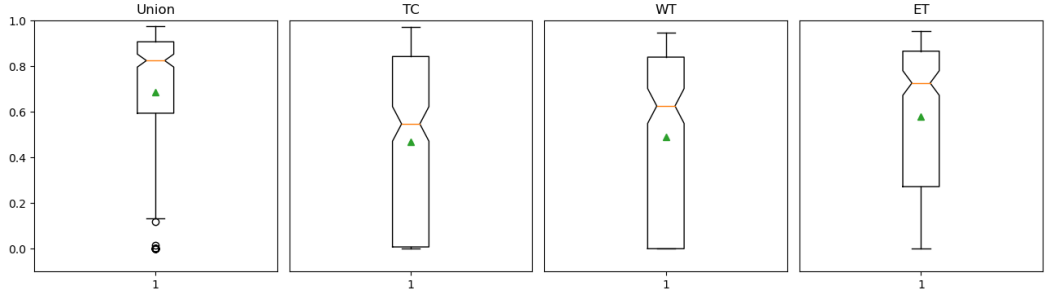

Figure 4.14: Over the 1250 MRI of BRATS 2021, the model is verified on 298 images. We plot, for these images, the Dice scores of the union segmentation, the TC, WT and ET.

4.4.2 Preprocessing

One approach we have taken for normalising our images involves simple techniques such as scaling the pixel values to lie within the range of 0 and 1, normalising the image patches to a common mean and variance, and using various equalisation techniques. Currently, we are utilising an affine 0-1 normalisation strategy by adding the minimum pixel value to all pixels and then dividing by the maximum value. However, we believe that we can enhance our results with a more sophisticated normalisation pipeline. After all, most of the literature on machine learning places significant emphasis on the preprocessing of data, as it is said to improve results considerably.

One could use superpixel methods to divide an image into smaller, homogeneous regions that are more easily analysed than individual pixels. Superpixel methods work by aggregating similar pixels into larger groups, typically based on colour, texture, or other image features. This process reduces the number of elements in the image, making it easier to perform image processing tasks while preserving the structure of the original image that can be

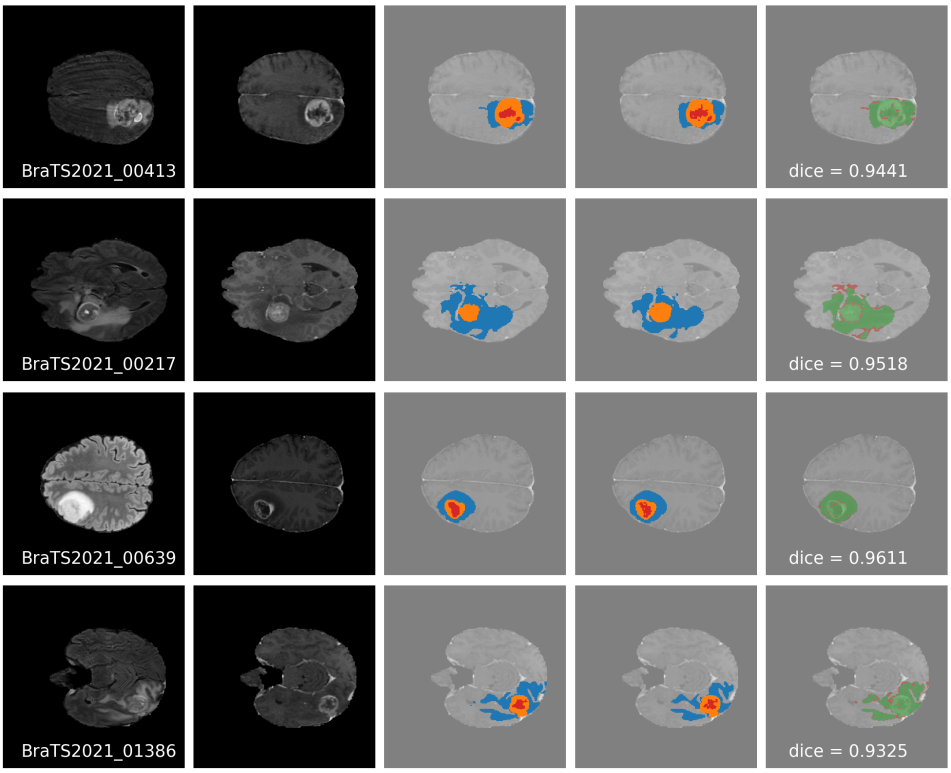

Figure 4.15: Results of our algorithm, in cases where the model is valid. On each row, from left to right: an image of modality FLAIR, of modality T1CE, the segmentation we obtained, the segmentation provided by the specialists and the comparison of both segmentation green being well matched and red otherwise.

used for further analysis and processing. For example, Shivhare and Kumar [2020] proposed a new method for brain tumour detection by combining features from Fluid-Attenuated Inversion Recovery (FLAIR) MRI scans with the graph-based manifold ranking algorithm. The algorithm has three main steps: After converting their image into superpixels, the rank of each superpixel or node is computed based on its affinity against certain selected nodes as the background prior. Finally, the relevance of each node with the background prior is computed and represented as a tumour map. However, the TDA principle is to avoid choosing a specific threshold before treatment. We put aside this idea as we find it irrelevant to our method.

Alternatively, we could study SPM's strategy that uses a Bayesian generative modelling perspective representing the probability density of the data in the most accurate but parsimonious way possible [Ashburner, 2009]. Indeed, their non-deterministic approaches are still exploitable in combination with TDA. Their segmentation algorithms discriminate tissue intensity by fitting a mixture of Gaussian models to their intensity ranges, also taking into account the tissue expected proportion within the brain.

As we presented in Section 1.1.2 each modality holds different information. A natural idea is to highlight targeted information by combining different modalities. There exist different approaches more or less involved to implement. A simple idea was proposed by Zhang et al. [2021] which simply consists of the algebraic superposition of FLAIR and T2 images:

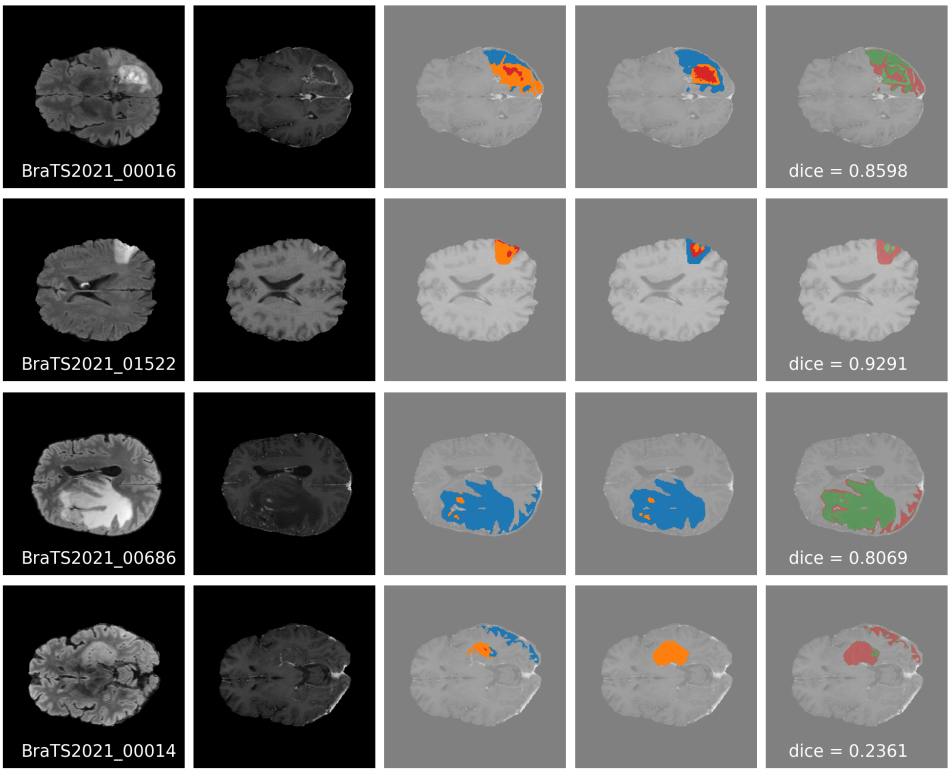

Figure 4.16: Results of our algorithm, in cases where the model is NOT valid. On each row, from left to right: an image of modality FLAIR, of modality T1CE, the segmentation we obtained, the segmentation provided by the specialists and the comparison of both segmentation green being well matched and red otherwise.

$(\texttt{FLAIR} + \texttt{T2})/(\texttt{FLAIR} \times \texttt{T2})$. We tried this very idea, correcting their equation slightly for our purposes. However, if it works for certain examples, it remains a significant number of images for which the improvement is unclear or worsened. To our understanding, it might indeed improve the performance of a neural network, having more information to infer from, yet for our simple technique, it does not seem to be consistent enough. More sophisticated methods exist, like the one detailed in Rajinikanth et al. [2018] based on Discrete Wavelet Transform PCA. We did not have the time to test such a method yet.

4.4.3 Refining the model with a preliminary topology analysis

In the previous section, we saw that only a quarter of the BraTS2021 dataset respects our model assumptions. We must revise the model, still focusing on the topology and keeping it simple.

Let us briefly revisit a research avenue that we previously abandoned: Our model's first assumption posits that the tumour X is the most luminous connected component in the FLAIR image. Once we have identified the optimal threshold using the `select_t` method, we must select the connected component that corresponds to the tumour. Initially, we attempted to accomplish this by selecting the connected component that contained the brightest pixel. In most cases, it was indeed selecting the right component, however, it

may fail if the FLAIR modality is not hyper-intense. This only occurs for astrocytomas (IDH-mutant, 1p19q non-codeleted) that present a T2/FLAIR mismatch [Deguchi et al., 2020]: On T2 weighted images, these tumours have extensive areas of fairly homogeneous and strikingly high signal. On FLAIR, instead, the majority of these areas become relatively hypointense in signal due to incomplete suppression. The occurrence of those mutants is unclear in the literature we reviewed. This is why we opted for selecting the component containing the most voxels, as it usually solves the problem. Nevertheless, it still happens on rare occurrences, that we select the grey matter. We hope that the correct connected component will be selected by at least one of those strategies.

To this day, the most significant challenge related to our model remains the ability to differentiate between the various sections of tumours. One example of this is when the tumour comes in contact with the skull, causing the ET component to have the same topology as a bowl rather than being homotopic to a sphere, as intended. This situation is illustrated in Figure 4.9. A potential solution to this issue would be to incorporate the skull mask during the search for the H_2 component.

More generally, we plan to define other simple models that would cover the variety of topologies of glioblastoma. Such models could be found by analysing the topology of the whole ground truth segmentations, and by identifying different glioma shape profiles. Segmentations come in the form of piece-wise constant images, eliminating all noise-related problems. Therefore, their persistence diagram should be easy to analyse and possibly to cluster. From this analysis, we hope to identify problematic cases and their occurrences, at least in the BraTS2021 dataset.

4.4.4 Perspective: Comparing with a known shape

We have seen that our method can fail for two reasons. First during the segmentation of the tumour union, depending if we select the brightest or the biggest component, we might take the wrong one (see discussion about IDH-mutant in Section 4.3.1). We hope, however, that the correct component within brain MRI will be chosen by at least one of these strategies. Second, while annotating the tumour, the simplistic topological model we use leads to incorrect segmentation.

For now, we lack the ability to segregate bad segmentations from good ones. This is why we need a representative glioblastoma to compare our segmentation with. In this thesis introduction, we referred to the diffeomorphic shape space theory as a great tool for such problems. For example, the works of Khan et al. [2008] and later Kutten et al. [2016] utilise a diffeomorphic shape space to segment and propagate labels from a brain template to segment unknown brains. The main difference between our methods is that in their case all tissues are to be segmented and are always structured in the same way. When in our case we have to segregate the tissues and determine the localisation of the tumour.

However, in a speculative approach, we can give a general approach outline. First, we would need to build the shape space. Using the BraTS2021 database, we can estimate a glioblastoma template. As it would be advised for more traditional statistical methods, we can split the database into two parts. One to construct the template and the other will be used to ensure that we can properly segment tumours that were not used to construct the shape space. At this point, we believe that one could determine if an estimated segmentation is realistic or not by computing its distance from the shape space. Then, cluster this space according to their topological signatures, getting a collection of models. Finally, we perform

a segmentation following each method and choose the most realistic.

This outline arise many questions: How to perform the clustering in such a space? How will we choose the models? Will the Riemannian distance of the space be enough to determine if a given segmentation is realistic? While this idea requires a lot of research, it fits well in this dissertation object as that is coherent with the idea of making tumourous brain atlases as discussed in section 1.2.3. Moreover, from this perspective, building a shape-space of glioblastoma would have been a first step in any case.

4.4.5 Conclusion

We have presented a simple yet promising methodology to segment glioblastomas based on TDA and simple imaging techniques. We have seen that the segmentation problem can be split into two sub-problems. First locating correctly the tumour within the MRI and second dividing the tumour into three annotated classes. We have shown an overall good performance on the union segmentation. Besides we do not put too much effort into a pre- or post-processing pipeline yet. However, splitting the tumour into classes turned out to be more difficult than expected, for plural reasons. First, we have shown that our topological assumption was verified only on a quarter of the BraTS data set. We can also incriminate biological and image normalisation irregularities. If image normalisation can be addressed with better preprocessing, biological variability must be addressed by designing and selecting new strategies. In our case, studying the topological variations across the dataset may give us the insight to design new segmentation pipelines. Then it will remain to select the correct pipeline, which will probably require to use of statistical methods. One could use machine learning, or with more delight, build a glioblastoma atlas and use traditional statistical techniques.

General Conclusion & Perspectives

At the outset of this thesis, as outlined in Section 1.2.3, we identified the need for a novel atlas that combined both Diffeomorphic and Statistical Atlases. Implying the need for reliable registration tools that can both 'explain' a displacement and be used to generate a diffeomorphic shape-space. The choice of the Metamorphic framework to register MRI data with different topologies was natural. The work presented in this thesis has advanced our practical understanding of these methods and may lead to promising results in the development of accurate and clinically useful atlases. In Chapter 2, we demonstrated our implementation of Metamorphosis and its initial results, which showed limitations in registering Glioblastomas due to momentum entanglement. In Chapter 3, we proposed a solution to this problem by incorporating a model to guide the Metamorphic registration. Our results showed a significant improvement in performance with the use of a simplistic model. In Chapter 4, we introduced the first iteration of a Glioma segmentation tool using TDA, which among other things, provided insights into the topology of Glioblastomas. We discovered that the simplistic model we had been using may not be adequate and the work on TDA segmentation may help us to define a more suitable model to use with CM. Based on the insights gained, we will now expand on potential research directions for this novel atlas. Just before that, let me review some works on Atlas construction necessary to understand the upcoming perspectives.

SHORT AND PRACTICAL REVIEW OF ATLAS CONSTRUCTION To begin, let us review some of the mechanisms involved in Diffeomorphic Shape-Spaces. We have presented in Section 2.2.1.a a theoretical foundation of such spaces, however, we did not cover a practical implementation method. Let be $x_i \in M$ a collection of data points on a general metric space M. The template $\hat{x}$ is defined as the Fréchet mean of the shape-space manifold M, via the distance d_{Diff} (see Section 2.2.3) such that:

$$\hat{x} = \underset{x}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^{N} d_{\text{Diff}}(x, x_i)^2 \tag{4.2}$$

Many variations of the template and shape space construction algorithms have been proposed, that can be split into Template-based and Template-free approaches.

Template-based approaches choose an initial reference image and all subjects are registered onto it [Vaillant et al., 2004; Lorenzen et al., 2005; Fonov et al., 2009, 2011; Legouhy et al., 2019]. However, this can introduce a bias towards the first reference image. To mitigate this, the inverse average transformation from the registration is used to compensate for this

bias. The process is then iterated, with the average model of the previous iteration used as the new reference image. By repeating this process, an unbiased atlas can be generated. For example, Fonov et al. [2009, 2011] built the MNI brain template that is one of the most used brain references today. Note that they do not use LDDMM but rather a log-exp-deformation-based method for computational constraints. Also, Lorenzen et al. [2005] claim constructing an unbiased atlas construction on images, by estimating a template by co-registering all subjects. They used LDDMM, but only on 2D slices on a database of 14 brains.

Template-free approaches, however, avoid the bias by not relying on an initial reference [Lorenzen et al., 2005; Glaunes and Joshi, 2006]. Two main strategies have been proposed: a pairwise and a groupwise method. Pairwise methods require registering each subject towards all the others. From these pairwise registrations, so-called unbiased subjects can be created that are then averaged. These procedures are long and time-consuming, which is aggravated by the fact that an additional data supply requires redoing the whole shape-space construction process.
Iteratives approaches aims to solve this problem [Cury et al., 2014; Legouhy et al., 2022]. These approaches assume that the template should be the centroid of the Riemannian space, and they use an iterative centroid scheme to approach it. The process begins by matching two images and selecting a point along the geodesic, which is considered the initial template. The centroid is then updated by matching it to another image in the dataset and selecting a point along the geodesic to be the new template. This process is repeated until all images have been processed. This approach can save time compared to the traditional method of constructing shape spaces for each additional data supply. Legouhy et al. [2022] propose a generalisation of the iterative diffeomorphic centroid approach to brain image atlasing, also capable of merging two atlases. Furthermore, they show that the ordering has only a slight influence on the result, compared to a template-based approach. Assuming we have constructed a shape-space, the interpolation of two shapes, denoted as A and B, is accomplished by computing the average of their initial momenta. Specifically, the initial momenta that register the template to both shapes, z_0^A and z_0^B, are averaged, and the resulting average momentum is used to compute the deformation from the template to the interpolated shape. Thus yielding the final interpolated shape.

A Metamorphic Atlas ? When registering images with different topologies using LDDMM, it is not realistic to expect a perfect match between the images. Instead, the result will be a projection of the target image onto the orbit of the image by the group of deformations. Thus, if one builds a diffeomorphic atlas from such images, these projections would be used as representatives rather than the actual targets. This introduces a source of error in tools that aim to explain shape variations. One could argue that this also happens when registering images with Metamorphosis. Indeed, as one combines intensity changes with deformation, the deformation part is the projection on the orbit as well (*c.f.:* Figure 2.20). However, the principle of Metamorphosis itself lies on explaining differences that cannot be explained by deformation. Furthermore, intensities changes can be analysed as discussed in this point of Chapter 2 conclusion.

Building this new kind of atlas by substituting LDDMM with Metamorphosis should be possible. Indeed we saw that as long as the momentum is conserved, one can control a path from its initial velocity, thus allowing us to adapt the atlas constructions methods reviewed in the previous paragraphs. We have to note, however, that as Metamorphosis completes the registration with intensity changes. Thus the resulting space will no longer be a shape-space

strictly speaking. For a given set of hyper-parameters, we could learn to interpret those and gain insight into the whole atlas. *It remains to study and determine how to interpret the variations between points within such space.*

A way to start answering this question is to ponder the meaning of a metamorphic registration. One might assume that two brains with tumours have the same topology, by stating that they have the same number of structures, roughly speaking a brain with a hole within. However, from a deformation perspective, this is not always the case. For instance, if we take two brains, one with a tumour in the temporal lobe and the other in the frontal lobe, a matching connecting voxels of the same nature only, would translate the tumour from one lobe to the other. In a brain, this would require many aberrant deformations. The most plausible solution would be to make one tumour disappear and the other vanishes. This is likely what would happen with Metamorphoses, aside from the entanglement issues between deformation and intensity. However, the implications of these topology changes in an atlas are not clear. *Will the resulting atlas be a brain with multiple small tumour seeds? Or a brain similar to a healthy brain?* To the best of our knowledge, these question has not been studied yet.

It would be interesting to make preliminary experiments, comparing LDDMM- and Metamorphosis-based Atlases, on simple images, such as the 'C' with a dot one (see Figure 2.20). We have shown that CM helps to improve the displacement relevance compared to the Metamorphosis one, in particular when mass effect is in play. Building an atlas with CM should not be too different from one built using M. Thus, in a second time we could study a CM-based Atlas. Before considering real data, we advise studying the toy-example data set we introduced in Section 3.2, to be able to validate the method within a controlled setting.

LOCALISED GLIOMA ATLAS After reading the previous paragraph, some may conclude that creating an atlas with metamorphoses poses risks, given our current lack of knowledge about how to interpret such atlases and extract useful information from them.

To avoid answering the questions above, like the problem of glioma translation from one lobe to another, we propose an alternative approach: constructing an atlas of glioblastomas for each lobe. Our method involves registering tumorous brains onto healthy ones using CM, such that the tumour disappears at specific voxels. We then collect the glioblastomas that have been pushed into each lobe, using a segmentation method to isolate the tumour from the rest of the brain. With these glioblastomas, we construct a shape space for each lobe using LDDMM.

In the perspective section of Chapter 4 (see Section 4.4.3), it was suggested that studying the topological variations of glioblastoma could be useful. One way to accomplish this would be to link the topological analysis with the analysis of the shape spaces. Specifically, we could study the variability of the tumour shape across different patients and how it relates to the location of the tumour within the brain. By combining this information with the Glioma topology analysis, it may be possible to identify specific patterns of tumour shape that are associated with different clinical outcomes or treatment responses. Furthermore, exploring the connection between topological analysis and shape space analysis could yield additional insights. For example, clustering techniques could be used to group together glioblastomas with similar topological features or similar shapes. This could potentially provide new insights into the underlying mechanisms of glioblastoma growth and the relationship between structure and function in the brain.

Bibliography

G. Ali, P. Kilian, B. Michel, C. Luigi, B. George, M. Elias, and D. Christos. GLISTR: Glioma image segmentation and registration. *IEEE Transactions on Medical Imaging*, 31: 1941–54, 2012.

N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950. URL http://dx.doi.org/10.2307/1990404.

V. Arsigny, O. Commowick, X. Pennec, and N. Ayache. A Log-Euclidean Framework for Statistics on Diffeomorphisms. In *MICCAI*, Lecture Notes in Computer Science, pages 924–931. 2006.

J. Ashburner. A fast diffeomorphic image registration algorithm. *NeuroImage*, 38(1):95–113, 2007.

J. Ashburner. Computational anatomy with the spm software. *Magnetic resonance imaging*, 27(8):1163–1174, 2009.

J. Ashburner and K. J. Friston. Diffeomorphic registration using geodesic shooting and gauss–newton optimisation. *NeuroImage*, 55(3):954–967, 2011.

B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, 12(1), 2008.

R. Azad, N. Khosravi, M. Dehghanmanshadi, J. Cohen-Adad, and D. Merhof. Medical image segmentation on mri images with missing modalities: A review. *arXiv preprint arXiv:2203.06217*, 2022.

B. Baheti and et al. The brain tumor sequence registration challenge: Establishing correspondence between pre-operative and follow-up mri scans of diffuse glioma patients, 2021.

U. Baid and et al. The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv:2107.02314*, 2021.

S. Bakas and et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge, 2018. URL https://arxiv.org/abs/1811.02629.

S. Bakas, H. Akbari, A. Sotiras, M. Bilello, M. Rozycki, J. S. Kirby, J. B. Freymann, K. Farahani, and C. Davatzikos. Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Scientific Data*, 4(1), 2017.

G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca. VoxelMorph: A Learning Framework for Deformable Medical Image Registration. *IEEE Transactions on Medical Imaging*, 38(8):1788–1800, 2019.

M. Bauer, M. Bruveris, and P. W. Michor. Overview of the geometries of shape spaces and diffeomorphism groups. *Journal of Mathematical Imaging and Vision*, 50(1):60–97, 2014.

M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157, 2005.

T. R. Berger, P. Y. Wen, M. Lang-Orsini, and U. N. Chukwueke. World health organization 2021 classification of central nervous system tumors and implications for therapy for adult-type gliomas: a review. *JAMA oncology*, 2022.

M. Bilello, H. Akbari, X. Da, J. M. Pisapia, S. Mohan, R. L. Wolf, D. M. ORourke, M. Martinez-Lage, and C. Davatzikos. Population-based mri atlases of spatial distribution are specific to patient and tumor characteristics in glioblastoma. *NeuroImage: Clinical*, 12:34–40, 2016.

M. Brett, A. Leff, C. Rorden, and J. Ashburner. Spatial Normalization of Brain Images with Focal Lesions Using Cost Function Masking. *NeuroImage*, 14(2):486–500, 2001.

M. Brunn, N. Himthani, G. Biros, M. Mehl, and A. Mang. Claire: constrained large deformation diffeomorphic image registration on parallel computing architectures. *Journal of open source software*, 6(61), 2021a.

M. Brunn, N. Himthani, G. Biros, M. Mehl, and A. Mang. Fast gpu 3d diffeomorphic image registration. *Journal of Parallel and Distributed Computing*, 149:149–162, 2021b.

T. Bruno, M. M. Peloso, A. Tabacco, and M. Vallarino. Sobolev spaces on lie groups: embedding theorems and algebra properties. *Journal of Functional Analysis*, 276(10): 3014–3050, 2019.

M. Bruveris, F. Gay-Balmaz, D. D. Holm, and T. S. Ratiu. The momentum map representation of images. *Journal of nonlinear science*, 21(1):115–150, 2011.

N. Byrne, J. R. Clough, G. Montana, and A. P. King. A persistent homology-based topological loss function for multi-class cnn segmentation of cardiac mri. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 3–13. Springer, 2021.

A. Bône, P. Vernhet, O. Colliot, and S. Durrleman. Learning Joint Shape and Appearance Representations with Metamorphic Auto-Encoders. In *MICCAI*, volume 12261, pages 202–211. 2020.

G. Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2): 255–308, 2009.

J. Céa. Conception optimale ou identification de formes, calcul rapide de la dérivée directionnelle de la fonction coût. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique*, 20(3):371–402, 1986.

N. Charon. *Analysis of geometric and functional shapes with extensions of currents : applications to registration and atlas estimation.* Theses, École normale supérieure de Cachan - ENS Cachan, Nov. 2013.

F. Chazal and B. Michel. An introduction to Topological Data Analysis: fundamental and practical aspects for data scientists. *Frontiers in Artificial Intelligence*, 4, 2021.

J. Clough, N. Byrne, I. Oksuz, V. A. Zimmer, J. A. Schnabel, and A. King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

J. R. Clough, I. Oksuz, N. Byrne, J. A. Schnabel, and A. P. King. Explicit topological priors for deep-learning based image segmentation using persistent homology. In *International Conference on Information Processing in Medical Imaging*, pages 16–28. Springer, 2019.

T. Cootes, E. Baldock, and J. Graham. An introduction to active shape models. *Image processing and analysis*, 328:223–248, 2000.

R. Courant and D. Hilbert. Methods of mathematical physics. *Bulletin of the American Mathematical Society*, 60:578–579, 1954.

L. Crawford, A. Monod, A. X. Chen, S. Mukherjee, and R. Rabadán. Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis. *Journal of the American Statistical Association*, 115(531):1139–1150, 2020.

J. Curry, H. Hang, W. Mio, T. Needham, and O. B. Okutan. Decorated merge trees for persistent topology. *Journal of Applied and Computational Topology*, pages 1–58, 2022.

C. Cury, J. A. Glaunès, and O. Colliot. Diffeomorphic iterative centroid methods for template estimation on large datasets. *Geometric Theory of Information*, pages 273–299, 2014.

P. C. De Witt Hamer, E. J. Hendriks, E. Mandonnet, F. Barkhof, A. H. Zwinderman, and H. Duffau. Resection probability maps for quality assessment of glioma surgery without brain location bias. *PloS one*, 8(9):e73353, 2013.

S. Deguchi, T. Oishi, K. Mitsuya, Y. Kakuda, M. Endo, T. Sugino, and N. Hayashi. Clinicopathological analysis of t2-flair mismatch sign in lower-grade gliomas. *Scientific reports*, 10(1):1–6, 2020.

I. L. Dryden and K. V. Mardia. *Statistical shape analysis: with applications in R*, volume 995. John Wiley & Sons, 2016.

P. Dupuis, U. Grenander, and M. I. Miller. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of Applied Mathematics*, 56(3):587–600, 1998a.

P. Dupuis, U. Grenander, and M. I. Miller. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of applied mathematics*, pages 587–600, 1998b.

D. R. Durran. *Numerical methods for wave equations in geophysical fluid dynamics*, volume 32. Springer Science & Business Media, 2013.

H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. In *Proceedings 41st annual symposium on foundations of computer science*, pages 454–463. IEEE, 2000.

G. J. Edwards, C. J. Taylor, and T. F. Cootes. Interpreting face images using active appearance models. In *Proceedings Third IEEE International Conference on Automatic Face and Gesture Recognition*, pages 300–305. IEEE, 1998.

A. Efremov, E. Karepova, V. Shaydurov, and A. Vyatkin. A computational realization of a semi-lagrangian method for solving the advection equation. *JAM*, 2014:1–12, 2014.

E. G. Escolar and Y. Hiraoka. Optimal cycles for persistent homology via linear programming. In *Optimization in the Real World: Toward Solving Real-World Optimization Problems*, pages 79–96. Springer, 2016.

R. P. Filipkiewicz. Isomorphisms between diffeomorphism groups. *Ergodic Theory and Dynamical Systems*, 2(2):159171, 1982. doi: 10.1017/S0143385700001486.

V. Fonov, A. C. Evans, K. Botteron, C. R. Almli, R. C. McKinstry, D. L. Collins, B. D. C. Group, et al. Unbiased average age-appropriate atlases for pediatric studies. *Neuroimage*, 54(1):313–327, 2011.

V. S. Fonov, A. C. Evans, R. C. McKinstry, C. R. Almli, and D. Collins. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage*, (47):S102, 2009.

A. François, P. Gori, and J. Glaunès. Metamorphic image registration using a semi-lagrangian scheme. In *SEE GSI*, 2021.

A. François, M. Maillard, C. Oppenheim, J. Pallud, I. Bloch, P. Gori, and J. Glaunès. Weighted metamorphosis for registration of images with different topologies. In *WBIR*, pages 8–17, 2022.

M. Futrega, A. Milesi, M. Marcinkiewicz, and P. Ribalta. Optimized u-net for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 15–29. Springer, 2022.

D. Garcia-Castellanos and U. Lombardo. Poles of inaccessibility: A calculation algorithm for the remotest places on earth. *Scottish Geographical Journal*, 123(3):227–233, 2007. doi: 10.1080/14702540801897809.

A. Garin, T. Heiss, K. Maggs, B. Bleile, and V. Robins. Duality in persistent homology of images. *arXiv preprint arXiv:2005.04597*, 2020.

C. Ge, I. Y.-H. Gu, A. S. Jakola, and J. Yang. Cross-modality augmentation of brain mr images using a novel pairwise generative adversarial network for enhanced glioma classification. In *2019 IEEE international conference on image processing (ICIP)*, pages 559–563. IEEE, 2019.

L. E. Ginsberg, G. N. Fuller, M. Hashmi, N. E. Leeds, and D. F. Schomer. The significance of lack of mr contrast enhancement of supratentorial brain tumors in adults: histopathological evaluation of a series. *Surgical neurology*, 49(4):436–440, 1998.

J. A. Glaunes and S. Joshi. Template estimation form unlabeled point set data and surfaces for computational anatomy. In *1st MICCAI workshop on mathematical foundations of computational anatomy: geometrical, statistical and registration methods for modeling biological shape variability*, 2006.

J. A. Glaunès. *Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l'anatomie numérique*. PhD thesis, 2005. Thèse de doctorat dirigée par Younes, Laurent et Trouvé, Alain Mathématiques Paris 13 2005.

A. Gooya, G. Biros, and C. Davatzikos. Deformable registration of glioma images using em algorithm and diffusion reaction modeling. *IEEE Transactions on Medical Imaging*, 30 (2):375–390, 2011a.

A. Gooya, K. M. Pohl, M. Bilello, G. Biros, and C. Davatzikos. Joint segmentation and deformable registration of brain scans guided by a tumor growth model. In *Medical image computing and computer-assisted intervention: MICCAI... International Conference on Medical Image Computing and Computer-Assisted Intervention*, volume 14, page 532. NIH Public Access, 2011b.

A. Gooya, K. M. Pohl, M. Bilello, L. Cirillo, G. Biros, E. R. Melhem, and C. Davatzikos. Glistr: glioma image segmentation and registration. *IEEE transactions on medical imaging*, 31(10):1941–1954, 2012.

U. Grenander. General pattern theory: A mathematical study of regular structures. 1993.

D. X. Guo. A semi-lagrangian runge-kutta method for time-dependent partial differential equations. *J. Appl. Anal. Comput*, 3(3):251–263, 2013.

X. Han, Z. Shen, Z. Xu, S. Bakas, H. Akbari, M. Bilello, C. Davatzikos, and M. Niethammer. A deep network for joint registration and reconstruction of images with pathologies. In *Machine Learning in Medical Imaging: 11th International Workshop, MLMI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings 11*, pages 342–352. Springer, 2020a.

X. Han, Z. Shen, Z. Xu, S. Bakas, H. Akbari, M. Bilello, C. Davatzikos, and M. Niethammer. A Deep Network for Joint Registration and Reconstruction of Images with Pathologies. In *International Workshop on Machine Learning in Medical Imaging*, pages 342–352, 2020b.

F. Hanif, K. Muzaffar, K. Perveen, S. M. Malhi, and S. U. Simjee. Glioblastoma multiforme: A review of its epidemiology and pathogenesis through clinical presentation and treatment. *Asian Pac. J. Cancer Prev.*, 18(1):3–9, Jan. 2017.

A. Hatcher. *Algebraic topology*. Cambridge Univ. Press, Cambridge, 2000. URL `https://cds.cern.ch/record/478079`.

M. Hernandez. Band-limited stokes large deformation diffeomorphic metric mapping. *IEEE Journal of Biomedical and Health Informatics*, 23(1):362–373, 2018.

M. Hernandez. Combining the band-limited parameterization and semi-lagrangian runge–kutta integration for efficient pde-constrained lddmm. *Journal of Mathematical Imaging and Vision*, 63(5):555–579, 2021.

D. D. Holm, A. Trouvé, and L. Younes. The euler-poincaré theory of metamorphosis. *Quarterly of Applied Mathematics*, 67(4):661–685, 2009.

R. Islam, S. Imran, M. Ashikuzzaman, and M. M. A. Khan. Detection and classification of brain tumor based on multilevel segmentation with convolutional neural network. *Journal of Biomedical Science and Engineering*, 13(4):45–53, 2020.

H. Jia, C. Bai, W. Cai, H. Huang, and Y. Xia. Hnf-netv2 for brain tumor segmentation using multi-modal mr imaging. *arXiv preprint arXiv:2202.05268*, 2022.

S. Kaji, T. Sudo, and K. Ahara. Cubical ripser: Software for computing persistent homology of image and volume data. *arXiv preprint arXiv:2005.12692*, 2020.

I. Kaltenmark. *Geometrical Growth Models for Computational Anatomy*. PhD thesis, Université Paris Saclay (COmUE), 2016.

I. Kaltenmark and A. Trouvé. Estimation of a Growth Development with Partial Diffeomorphic Mappings. working paper or preprint, Dec. 2017. URL https://hal.science/hal-01656670.

M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International journal of computer vision*, 1(4):321–331, 1988.

D. G. Kendall. Shape manifolds, procrustean metrics, and complex projective spaces. *Bulletin of the London mathematical society*, 16(2):81–121, 1984.

D. G. Kendall, D. Barden, T. K. Carne, and H. Le. *Shape and shape theory*. John Wiley & Sons, 2009.

A. R. Khan, L. Wang, and M. F. Beg. Freesurfer-initiated fully-automated subcortical brain segmentation in mri using large deformation diffeomorphic metric mapping. *Neuroimage*, 41(3):735–746, 2008.

C. P. Klingenberg. Walking on kendalls shape space: understanding shape spaces and their coordinate systems. *Evolutionary Biology*, 47(4):334–352, 2020.

K. Kotowski, S. Adamski, B. Machura, L. Zarudzki, and J. Nalepa. Coupling nnu-nets with expert knowledge for accurate brain tumor segmentation from mri. In *International MICCAI Brainlesion Workshop*, pages 197–209. Springer, 2022.

K. S. Kutten, J. T. Vogelstein, N. Charon, L. Ye, K. Deisseroth, and M. I. Miller. Deformably registering and annotating whole clarity brains to an atlas via masked lddmm. In *Optics, Photonics and Digital Technologies for Imaging Applications IV*, volume 9896, pages 282–290. SPIE, 2016.

H. Lee, H. Kang, M. K. Chung, B.-N. Kim, and D. S. Lee. Persistent brain network homology from the perspective of dendrogram. *IEEE transactions on medical imaging*, 31(12): 2267–2277, 2012.

J. Lee. *Introduction to Smooth Manifolds*. Springer, 2012.

A. Legouhy, O. Commowick, F. Rousseau, and C. Barillot. Unbiased longitudinal brain atlas creation using robust linear registration and log-euclidean framework for diffeomorphisms. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1038–1041. IEEE, 2019.

A. Legouhy, F. Rousseau, C. Barillot, and O. Commowick. An iterative centroid approach for diffeomorphic online atlasing. *IEEE Transactions on Medical Imaging*, 41(9):2521–2531, 2022.

Z. Liu, L. Tong, L. Chen, F. Zhou, Z. Jiang, Q. Zhang, Y. Wang, C. Shan, L. Li, and H. Zhou. Canet: Context aware network for brain glioma segmentation. *IEEE Transactions on Medical Imaging*, 40(7):1763–1777, 2021.

P. Lorenzen, B. C. Davis, and S. Joshi. Unbiased atlas formation via large deformations metric mapping. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005: 8th International Conference, Palm Springs, CA, USA, October 26-29, 2005, Proceedings, Part II 8*, pages 411–418. Springer, 2005.

M. Lorenzi, N. Ayache, G. B. Frisoni, and X. Pennec. LCC-Demons: A robust and accurate symmetric diffeomorphic registration algorithm. *NeuroImage*, 81:470–483, 2013.

H. M. Luu and S.-H. Park. Extending nn-unet for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 173–186. Springer, 2022.

J. Ma and J. Chen. Nnunet with region-based training and loss ensembles for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 421–430. Springer, 2022.

M. Maillard, A. François, J. Glaunès, I. Bloch, and P. Gori. A deep residual learning implementation of metamorphosis. In *IEEE ISBI*, 2022.

A. Mang, A. Gholami, C. Davatzikos, and G. Biros. Claire: A distributed-memory solver for constrained large deformation diffeomorphic image registration. *SIAM Journal on Scientific Computing*, 41(5):C548–C584, 2019.

T. Mansi, X. Pennec, M. Sermesant, H. Delingette, and N. Ayache. LogDemons revisited: Consistent regularisation and incompressibility constraint for soft tissue tracking in medical images. In *MICCAI*, volume 6362, pages 652–659. 2010.

L. Mansilla, D. H. Milone, and E. Ferrante. Learning deformable registration of medical images with anatomical constraints. *Neural Networks*, 124:269–279, 2020.

C. Maria, J.-D. Boissonnat, M. Glisse, and M. Yvinec. The gudhi library: Simplicial complexes and persistent homology. In *International congress on mathematical software*, pages 167–174. Springer, 2014.

A. McNamara, A. Treuille, Z. Popović, and J. Stam. Fluid control using the adjoint method. *ACM Transactions On Graphics (TOG)*, 23(3):449–456, 2004.

B. H. Menze and et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging*, 34(10):1993–2024, 2015.

M. Miller and L. Younes. Group actions, homeomorphisms, and matching: A general framework. *International Journal of Computer Vision*, 41:61–84, 01 2001. doi: 10.1023/A:1011161132514.

M. I. Miller, A. Trouvé, and L. Younes. Geodesic shooting for computational anatomy. *Journal of Mathematical Imaging and Vision*, 24(2):209–228, 2006.

M. I. Miller, A. Trouvé, and L. Younes. Hamiltonian systems and optimal control in computational anatomy: 100 years since d'arcy thompson. *Annual review of biomedical engineering*, 17:447–509, 2015.

J. Milnor. *Topology from the Differentiable Viewpoint*. University Press of Virginia, 1965.

T. C. W. Mok and A. C. S. Chung. Large Deformation Diffeomorphic Image Registration with Laplacian Pyramid Networks. In *MICCAI*, 2020.

J. Mourtada. Structure algébrique des groupes de diéomorphismes. page 16, 2015.

M. M. Mrugala. Advances and challenges in the treatment of glioblastoma: a clinicians perspective. *Discovery medicine*, 15(83):221–230, 2013.

M. Musy, G. Jacquenot, G. Dalmasso, J. Lee, R. de Bruin, neoglez, J. Soltwedel, A. Pollack, F. Claudi, C. Badger, A. Sol, Z.-Q. Zhou, B. Sullivan, B. Lerner, D. Hrisca, D. Volpatto, Evan, F. Matzkin, mkerrinrapid, N. Schlömer, RichardScottOZ, RobinEnjalbert, X. Lu, and O. Schneider. marcomusy/vedo: 2023.4.3, Jan. 2023. URL https://doi.org/10.5281/zenodo.7545806.

M. Nakada, D. Kita, T. Watanabe, Y. Hayashi, L. Teng, I. V. Pyko, and J.-I. Hamada. Aberrant signaling pathways in glioma. *Cancers*, 3(3):3242–3278, 2011.

M. Niethammer and et al. Geometric Metamorphosis. In *MICCAI*, volume LNCS 6892, pages 639–646, 2011.

M. Niethammer, R. Kwitt, and F.-X. Vialard. Metric Learning for Image Registration. In *CVPR*, pages 8455–8464, 2019.

P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008.

E. Noether. Invariante variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1918:235–257, 1918. URL http://eudml.org/doc/59024.

I. Obayashi. Volume-optimal cycle: Tightest representative cycle of a generator in persistent homology. *SIAM Journal on Applied Algebra and Geometry*, 2(4):508–534, 2018.

S. Y. Oudot. Persistence theory: from quiver representations to data analysis. volume 209. American Mathematical Soc., 2017.

J. Pallud, A. Roux, and M. Zanello. Relationship between tumour location and preoperative seizure incidence depends on glioma grade of malignancy. *Epileptic Disorders: International Epilepsy Journal with Videotape*, 18(1):107–108, 2016.

S. Parisot, A. Darlix, C. Baumann, S. Zouaoui, Y. Yordanova, M. Blonski, V. Rigau, S. Chemouny, L. Taillandier, L. Bauchet, et al. A probabilistic atlas of diffuse who grade ii glioma locations in the brain. *PloS one*, 11(1):e0144200, 2016.

X. Pennec, P. Cachier, and N. Ayache. Understanding the demons algorithm: 3d non-rigid registration by gradient descent. In *MICCAI*, volume 1679, pages 597–605. 1999.

T. Qaiser, K. Sirinukunwattana, K. Nakane, Y.-W. Tsang, D. Epstein, and N. Rajpoot. Persistent homology for fast tumor segmentation in whole slide histology images. *Procedia Computer Science*, 90:119–124, 2016.

T. Qaiser, Y.-W. Tsang, D. Taniyama, N. Sakamoto, K. Nakane, D. Epstein, and N. Rajpoot. Fast and accurate tumor segmentation of histology images using persistent homology and deep convolutional features. *Medical image analysis*, 55:1–14, 2019.

V. Rajinikanth, S. C. Satapathy, N. Dey, and R. Vijayarajan. Dwt-pca image fusion technique to improve segmentation accuracy in brain tumor analysis. In *Microelectronics, Electromagnetics and Telecommunications: Proceedings of ICMEET 2017*, pages 453–462. Springer, 2018.

J. Ren, W. Zhang, N. An, Q. Hu, Y. Zhang, and Y. Zhou. Ensemble outperforms single models in brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 451–462. Springer, 2022.

C. L. Richardson and L. Younes. Metamorphosis of images in reproducing kernel Hilbert spaces. *Adv Comput Math*, 42(3):573–603, 2016.

P. Ripollés, J. Marco-Pallarés, R. de Diego-Balaguer, J. Miró, M. Falip, M. Juncadella, F. Rubio, and A. Rodriguez-Fornells. Analysis of automated methods for spatial normalization of lesioned brains. *Neuroimage*, 60(2):1296–1306, 2012.

L. Risser, F.-X. Vialard, H. Y. Baluwala, and J. A. Schnabel. Piecewise-diffeomorphic image registration: application to the motion estimation between 3D CT lung images with sliding conditions. *Med. Image Anal.*, 17(2):182–193, feb 2013.

K. Rock, O. McArdle, P. Forde, M. Dunne, D. Fitzpatrick, B. O'Neill, and C. Faul. A clinical review of treatment outcomes in glioblastoma multiformethe validation in a non-trial population of the results of a randomised phase iii clinical trial: has a more radical approach improved survival? *The British journal of radiology*, 85(1017):e729–e733, 2012.

M.-M. Rohé, M. Datar, T. Heimann, M. Sermesant, and X. Pennec. SVF-Net: Learning Deformable Image Registration Using Shape Matching. In *MICCAI*, page 266, 2017.

C. Rorden, R. Newman-Norlund, C. Drake, D. R. Glen, J. Fridriksson, T. Hanayik, and P. A. Taylor. Improving 3d edge detection for visual inspection of mri coregistration and alignment. *bioRxiv*, pages 2022–09, 2022.

A. Roux and et al. MRI atlas of IDH wild-type supratentorial glioblastoma: Probabilistic maps of phenotype, management, and outcomes. *Radiology*, 293(3):633–643, Dec. 2019.

M. Rucco, G. Viticchi, and L. Falsetti. Towards personalized diagnosis of glioblastoma in fluid-attenuated inversion recovery (flair) by topological interpretable machine learning. *Mathematics*, 8(5):770, 2020.

A. Saadat-Yazdi, R. Andreeva, and R. Sarkar. Topological detection of alzheimers disease using betti curves. In *Interpretability of Machine Intelligence in Medical Image Computing, and Topological Data Analysis and Its Applications for Medical Data*, pages 119–128. Springer, 2021.

L. M. Sagberg, D. H. Iversen, E. H. Fyllingen, A. S. Jakola, I. Reinertsen, and O. Solheim. Brain atlas for assessing the impact of tumor location on perioperative quality of life in patients with high-grade glioma: A prospective population-based cohort study. *NeuroImage: Clinical*, 21:101658, 2019.

K. Scheufele and et al. Coupling brain-tumor biophysical models and diffeomorphic image registration. *Comput Methods Appl Mech Eng*, 347:533–567, 2019.

R. Schmid. Infinite-dimensional lie groups and algebras in mathematical physics. *Advances in Mathematical Physics*, 2010, 2010.

M. Sdika and D. Pelletier. Nonrigid registration of multiple sclerosis brain images using lesion inpainting for morphometry or lesion mapping. *Human Brain Mapping*, 30(4): 1060–1067, 2009.

S. N. Shivhare and N. Kumar. Brain tumor detection using manifold ranking in flair mri. In *Proceedings of ICETIT 2019: Emerging Trends in Information Technology*, pages 292–305. Springer, 2020.

Z. Shu and et al. Deforming Autoencoders: Unsupervised Disentangling of Shape and Appearance. *ECCV*, 2018.

M. M. R. Siddiquee and A. Myronenko. Redundancy reduction in semantic segmentation of 3d brain tumor mris. *arXiv preprint arXiv:2111.00742*, 2021.

J. Simpson, J. Horton, C. Scott, W. Curran, P. Rubin, J. Fischbach, S. Isaacson, M. Rotman, S. Asbell, J. Nelson, et al. Influence of location and extent of surgical resection on survival of patients with glioblastoma multiforme: results of three consecutive radiation therapy oncology group (rtog) clinical trials. *International Journal of Radiation Oncology* Biology* Physics*, 26(2):239–244, 1993.

W. Ślebodziński. Sur les équations de hamilton. *Bull Acad Roy Belg*, 17:864–870, 1931.

S. Sommer, M. Nielsen, F. Lauze, and X. Pennec. A multi-scale kernel bundle for lddmm: Towards sparse deformation description across space and scales. In G. Székely and H. K. Hahn, editors, *IPMI*, pages 624–635, 2011.

A. Srivastava, S. H. Joshi, W. Mio, and X. Liu. Statistical shape analysis: Clustering, learning, and testing. *IEEE Transactions on pattern analysis and machine intelligence*, 27(4):590–602, 2005.

G. Tauzin, U. Lupo, L. Tunstall, J. B. Pérez, M. Caorsi, A. M. Medina-Mardones, A. Dassatti, and K. Hess. giotto-tda:: A topological data analysis toolkit for machine learning and data exploration. *J. Mach. Learn. Res.*, 22(39):1–6, 2021.

D. W. Thompson. *On growth and form*, volume 2. Cambridge university press Cambridge, 1942.

R. Torsten, N. M. Zahr, E. V. Sullivan, and A. Pfefferbaum. The SRI24 multichannel atlas of normal adult human brain structure. *Human brain mapping*, 31(5):798–819, 2010.

A. Trouvé and L. Younes. Matching deformable objects. mise en correspondance dobjets déformables matching deformable objects mise en correspondance d'objets déformables. *International Information and Engineering Technology Association (IIETA)*, 2003.

A. Trouvé and L. Younes. Local geometry of deformable templates. *SIAM Journal on Mathematical Analysis*, 37(1):17–59, 2005.

A. Trouvé and L. Younes. Metamorphoses through lie group action. *Foundations of computational mathematics*, 5(2):173–198, 2005.

A. Trouve, A. Trouv, and A. T. E. Diffeomorphisms groups and pattern matching in image analysis, 1995.

A. Trouvé. An infinite dimensional group approach for physics based models in pattern recognition. *International Journal of Computer Vision - IJCV*, 01 1995.

M. Vaillant, M. I. Miller, L. Younes, and A. Trouvé. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23:S161–S169, 2004.

T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. Diffeomorphic demons: Efficient non-parametric image registration. 45(1):S61–S72, 2009.

F.-X. Vialard and F. Santambrogio. Extension to BV functions of the large deformation diffeomorphisms matching approach. 347(1):27–32, 2009. ISSN 1631-073X. doi: 10.1016/j.crma.2008.11.019.

F.-X. Vialard, L. Risser, D. Rueckert, and C. J. Cotter. Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. *International Journal of Computer Vision*, 97(2):229–241, 2011.

C. Villa, C. Miquel, D. Mosses, M. Bernier, and A. L. Di Stefano. The 2016 world health organization classification of tumours of the central nervous system. *La Presse Médicale*, 47(11-12):e187–e200, 2018.

J. E. Villanueva-Meyer, M. C. Mabray, and S. Cha. Current clinical brain tumor imaging. *Neurosurgery*, 81(3):397–415, 2017.

X. Yang, R. Kwitt, M. Styner, and M. Niethammer. Quicksilver: Fast Predictive Image Registration - a Deep Learning Approach. *NeuroImage*, 158:378–396, 2017.

L. Younes. *Shapes and Diffeomorphisms*. Springer Berlin Heidelberg, 2019.

Y. Yuan. Evaluating scale attention network for automatic brain tumor segmentation with large multi-parametric mri database. In *International MICCAI Brainlesion Workshop*, pages 42–53. Springer, 2022.

E. Zeidler. *Applied functional analysis: applications to mathematical physics*, volume 108. Springer Science & Business Media, 2012.

M. Zhang and P. T. Fletcher. Fast diffeomorphic image registration via Fourier-approximated Lie algebras. *International Journal of Computer Vision*, 127(1):61–73, 2018.

X. Zhang, Y. Hu, W. Chen, G. Huang, and S. Nie. 3d brain glioma segmentation in mri through integrating multiple densely connected 2d convolutional neural networks. *Journal of Zhejiang University-SCIENCE B*, 22(6):462–475, 2021.

L. Zhao and K. Jia. Deep feature learning with discrimination mechanism for brain tumor segmentation and diagnosis. In *2015 international conference on intelligent information hiding and multimedia signal processing (IIH-MSP)*, pages 306–309. IEEE, 2015.

A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

Diffeomorphic image registration taking topological differences into account
Metamorphosis on brain MRI containing Glioblatomas

Abstract: This thesis addresses the problem of registering images with different topologies with diffeomorphic deformation. We focus on the case of medical images of glioblastomas, a type of brain tumour.
Firstly, we implemented both Metamorphosis and LDDMM for images in 2D and 3D. Our implementation is object-oriented and developed using `PyTorch`, allowing for versatility in usage and easy modifications. We also used a semi-Lagrangian scheme on both images and residual.The implementation is GPU-accelerated, and we demonstrate the effectiveness of our approach through experiments on glioblastomas using BraTS datasets.
Secondly, we address the difficulties associated with the Metamorphosis algorithm by proposing a framework for incorporating prior knowledge into the model, called Constrained Metamorphosis. The framework allows for adding constraints on the registration problem by also matching given *priors*. We present two specific types of *priors* that can be incorporated into the model: a growing mask generated from a given segmentation and a field that guides the deformation in a desired direction. We demonstrate the effectiveness of our approach through experiments on glioblastomas using BraTS datasets, comparing with state-of-the-art methods.
Finally, we developed a tumour segmentation tool using Topological Data Analysis (TDA) to detect characteristic components within the FLAIR and T1ce modalities.

Keywords: Diffeomorphic registration, LDDMM, Metamorphosis, Glioblastomas.

Recalage difféomorphique d'images respectant les différences topologiques
Metamorphoses sur des IRM de cerveaux atteints de Glioblastomes

Résumé : Cette thèse aborde le problème du recalage d'images ayant des topologies différentes avec une déformation diffeomorphique. Nous nous concentrons sur le cas des images médicales de glioblastomes, un type de tumeur cérébrale.
Tout d'abord, nous avons implémenté à la fois les Metamorphoses et LDDMM pour des images en 2D et 3D. Notre implémentation est orientée objet et développée à l'aide de `PyTorch`, permettant une grande versatilité d'utilisation et des modifications faciles. Nous avons également utilisé un schéma semi-lagrangien sur les images et les résidus.L'implémentation est accélérée par GPU, et nous démontrons l'efficacité de notre approche à travers des expériences sur des glioblastomes en utilisant les données BraTS.
Dans un second temps, nous abordons les difficultés pratiques associées aux Métamorphoses en proposant un cadre pour incorporer des connaissances préalables dans le modèle, appelé Métamorphoses Contraintes. Le cadre permet d'ajouter des contraintes sur le problème de recalage en utilisant également des *a-priori*. Nous présentons deux types spécifiques de *prior* qui peuvent être incorporées dans le modèle : un masque de croissance généré à partir d'une segmentation donnée et un champ qui guide la déformation dans une direction souhaitée. Nous démontrons l'efficacité de notre approche à travers des expériences sur des glioblastomes en utilisant des ensembles de données BraTS, en comparant avec des méthodes de pointe.
Enfin, nous avons développé un outil de segmentation de tumeurs utilisant l'analyse de données topologiques (TDA) pour détecter des composants caractéristiques dans les modalités FLAIR et T1ce.

Mots-Clefs : Recalage Diffeomorphique, LDDMM, Metamorphoses, Glioblastomes.