Morphology, Geometry and Statistics in non-standard imaging

Emmanuel Chevallier

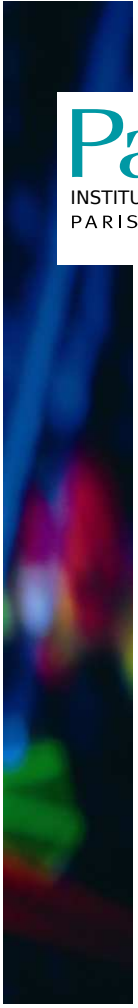

ParisTech
INSTITUT DES SCIENCES ET TECHNOLOGIES
PARIS INSTITUTE OF TECHNOLOGY

MINES ParisTech

École doctorale n°432 : Sciences des Métiers de l'Ingénieur

Doctorat ParisTech

T H È S E

pour obtenir le grade de docteur délivré par

l'École nationale supérieure des mines de Paris

Spécialité « Morphologie Mathématique »

présentée et soutenue publiquement par

Emmanuel CHEVALLIER

18 novembre 2015

Morphologie, Géométrie et Statistiques en imagerie non-standard

Morphology, Geometry and Statistics in non-standard imaging

Directeur de thèse : **Jesús ANGULO**

Jury

M. Fernand MEYER, Directeur de recherche, Mines ParisTech	Président
M. Michel BERTHIER, Professeur, Université de La Rochelle	Rapporteur
Mme. Isabelle BLOCH, Professeur, Telecom ParisTech	Rapporteur
Mme. Gabriele STEIDL, Professeur, Technische Universität Kaiserslautern	Rapporteur
M. Frédéric BARBARESCO, Senior Scientist, THALES	Examinateur
M. Jasper van de GRONDE, Post-doctorant, Groningen University	Examinateur
M. Christian RONSE, Professeur, Université de Strasbourg	Examinateur

T H È S E

MINES ParisTech
CMM - Centre de Morphologie Mathématique, Mathématiques et Systèmes
35, rue Saint-Honoré, 77305 Fontainebleau, France

Résumé, Abstract

Résumé

Le traitement d'images numériques a suivi l'évolution de l'électronique et de l'informatique. Il est maintenant courant de manipuler des images à valeur non pas dans $\{0, 1\}$, mais dans des variétés ou des distributions de probabilités. C'est le cas par exemple des images couleurs où de l'imagerie du tenseur de diffusion (DTI). Chaque type d'image possède ses propres structures algèbriques, topologiques et géometriques. Ainsi, les techniques existantes de traitement d'image doivent être adaptés lorsqu'elles sont appliquées à de nouvelles modalités d'imagerie. Lorsque l'on maniuple de nouveaux types d'espaces de valeurs, les précédant opérateurs peuvent rarement être utilisés tel quel. Même si les notions sous-jacentes ont encore un sens, un travail doit être mené afin de les exprimer dans le nouveau contexte.

Cette thèse est composée de deux parties indépendantes. La première, « Morphologie matématiques pour les images non standards », concerne l'extension de la morphologie mathématique à des cas particuliers où l'espace des valeurs de l'image ne possède pas de structure d'ordre canonique. Le chapitre 2 formalise et démontre le problème de l'irregularité des ordres totaux dans les espaces métriques. Le résultat principal de ce chapitre montre qu'étant donné un ordre total dans un espace vectoriel multidimensionnel, il existe toujours des images à valeur dans cet espace tel que les dilatations et les érosions morphologiques soient irrégulières et incohérentes. Le chapitre 3 est une tentative d'extension de la morphologie mathématique aux images à valeur dans un ensemble de labels non ordonnés.

La deuxième partie de la thèse, « Estimation de densités de probabilités dans les espaces de Riemann » concerne l'adaptation des techniques classiques d'estimation de densités non paramétriques à certaines variétés Riemanniennes. Le chapitre 5 est un travail sur les histogrammes d'images couleur dans le cadre de métriques perceptuelles. L'idée principale de ce chapitre consiste à calculer les histogrammes suivant une approximation euclidienne local de la métrique perceptuelle, et non une approximation globale comme dans les espaces perceptuels standards. Le chapitre 6 est une étude sur l'estimation de densité lorsque les données sont des lois Gaussiennes. Différentes techniques y sont analysées. Le résultat principal est l'expression de noyaux pour la métrique de Wasserstein.

Mots clés:
Morphologie Mathématique, Ordres Totaux, Traitement d'Image, Estimation de Densités, Géometrie Riemannienne, Métrique de Fisher, Métrique de Wasserstein, Histogrammes d'Images

Abstract

Digital image processing has followed the evolution of electronic and computer science. It is now current to deal with images valued not in $\{0,1\}$ or in gray-scale, but in manifolds or probability distributions. This is for instance the case for color images or in diffusion tensor imaging (DTI). Each kind of images has its own algebraic, topological and geometric properties. Thus, existing image processing techniques have to be adapted when applied to new imaging modalities. When dealing with new kind of value spaces, former operators can rarely be used as they are. Even if the underlying notion has still a meaning, a work must be carried out in order to express it in the new context.

The thesis is composed of two independent parts. The first one, "Mathematical morphology on non-standard images", concerns the extension of mathematical morphology to specific cases where the value space of the image does not have a canonical order structure. Chapter 2 formalizes and demonstrates the irregularity issue of total orders in metric spaces. The main results states that for any total order in a multidimensional vector space, there are images for which the morphological dilations and erosions are irregular and inconsistent. Chapter 3 is an attempt to generalize morphology to images valued in a set of unordered labels.

The second part "Probability density estimation on Riemannian spaces" concerns the adaptation of standard density estimation techniques to specific Riemannian manifolds. Chapter 5 is a work on color image histograms under perceptual metrics. The main idea of this chapter consists in computing histograms using local Euclidean approximations of the perceptual metric, and not a global Euclidean approximation as in standard perceptual color spaces. Chapter 6 addresses the problem of non parametric density estimation when data lay in spaces of Gaussian laws. Different techniques are studied, an expression of kernels is provided for the Wasserstein metric.

Key words:
Mathematical Morphology, Total Orders, Image Processing, Density Estimation, Riemanninan Geometry, Fisher Metric, Wasserstein Metric, Image Histograms

Acknowledgments

I am convinced I have been very lucky to be supervised by Jesùs Angulo, he has been an ideal director. Thanks to his right balance between guidance and freedom, the three last years have been an exciting scientific experience.

I would like to thank Jasper van de Gronde, for the multiple discussions and constructive suggestions during our respective academic visits. Jasper is always ready to spend time to help his colleagues.

I would like to thank Daniel Bennequin, Jean-Philippe Anker, Frédéric Barbaresco and Ivar Farup for their help on various problems.

I also would like to thank the referees for their help in the improvement of the manuscript.

At last, I would like to thank my family, my friends from Malakableau and all the members of the Center for Mathematical Morphology. Thanks to them, the last three years have been particularly memorable.

Contents

1. Introduction

Résumé

Ce chapitre présente les motivations et les thèmes abordés dans cette thèse. L'augmentation progressive de la mémoire disponible et l'évolution des dispositifs physiques d'acquisition ont permis l'apparition d'images de plus en plus complexes. Les techniques existantes de traitement doivent être adaptées à chaque nouvelle modalité d'image. Deux problématiques sont abordées dans cette thèse. La première concerne la morphologie mathématique, dans des cas particuliers où l'espace des valeurs de l'image ne dispose pas intrinsèquement d'une structure de treillis. La seconde concerne l'estimation de densité dans des variétés Riemanniennes. Ce problème apparaît notamment lors de la construction d'histogrammes d'images à valeurs dans des variétés Riemanniennes.

1.1. Motivation: Non standard images

Since its origins, digital signal processing follows the evolution of electronic technology. In the sixties, with the development of computers, image processing has become a discipline in itself. For technological reasons, first computers could only handle images composed of two distinct levels, 0 and 1. Thus, the first techniques of image processing including mathematical morphology where developed for binary images. Since the sixties, the development of sensors and computers has widened digital image processing to numerous new forms of images. The nature of images evolves on two levels. On the one hand, the spatial support has evolved from a regular grid of pixels to an arbitrary graph. On the other hand, the value contained in each pixel or node of the graph takes more and more complex forms. The increase of available memory has first enabled the transition from binary to gray-scale images. The memory is now sufficient to describe a significant amount of manifolds such as colors or probability distributions. Diffusion tensor imaging (DTI) is a typical case of non standard images. A physical device measures the motion of water molecules and models it by a centered 3-dimensional Gaussian distribution. The set of 3-dimensional Gaussian distributions are a 6-dimensional manifold, naturally represented by the 3×3 symmetric positive definite matrices. Another case of images valued in distributions arise from the popularization of superpixels. In order to simplify images, neighbor pixels can be gathered according to a given criterion, to form a set of superpixels. Usually a single value is chosen to represent a superpixel. However, it is sometimes interesting to preserve more information and to assign a distribution of values to each superpixel instead of a single value.

Each kind of images has its own algebraic, topological and geometric properties. These properties depend on the graph of pixels, on the value space, and on the nature of the observed objects. In this thesis, we leave aside the problem of the graph of pixels and focus on the value space and its structure. Usually, an image processing operator depends on the structure of the image. Some of them have a certain generality. For instance, the

mean and the median filters are defined on a large variety of spaces. However the definition of the filters depends on the value space of the image: the mean and the median, originally defined using addition, division and an order, have been reformulated in metric spaces. When dealing with new kind of value spaces, former operators can rarely be used as they are. Even if the underlying notion has still a meaning, a work must be carried out in order to express it in the new context.

1.2. Statistics in non-Euclidean spaces

The adaptation of standard quantities of statistic to new spaces is an important problem. For instance, the expected value was first defined using the algebraic structure of $\mathbb{R}$. This definition is naturally extended to random phenomena valued in $\mathbb{R}^n$. Consider a distribution μ, taking values on the surface of the earth S, for instance the distribution of the population in a country. The notion of mean or expected value has an intuitive meaning but can not be defined from algebraic operations. One has to find the counterpart of the previous notion based on the only intrinsic structure of S: the notion of distance $d(x \in S, x' \in S)$. Without entering uniqueness problems, the following quantities

$$E^S = \text{argmin}_x \int_{x' \in S} d(x, x')^2 d\mu,$$

$$E^S_{emp}(n) = \text{argmin}_x \sum_i d(x, x_i)^2.$$

correspond to the original expected value and empirical expected value in $\mathbb{R}$ when replacing the distance on S by the distance on $\mathbb{R}$. Furthermore, they verify similar convergence properties ($E^S_{emp}(n) \to E^S$) for many standard distances d. A mean formalized according to this definition is called a Fréchet mean. The situation is similar for the space of colors. The set of colors has intrinsically no algebraic structure, but posses intrinsic metrics. Thus a mean between colors should be a Fréchet mean.

In this thesis we are not specifically interested in the notion of mean but in the notion of density of a distribution. Unlike the mean where the definition must be adapted, the definition of the density is not a problem as long as the value space has a reference measure. However the estimation of a probability density from a set of samples strongly depends on the context. Consider a space $\mathcal{X}$ and two metrics d_1 and d_2 on $\mathcal{X}$. The most straightforward way of estimating a density is to use a tilling of the space. Estimating a density consists in counting a number of samples in each bin of the tilling and dividing by their surface. Finding good tilling is a difficult problem that has to be studied for each case $(\mathcal{X}, d_i)$. The same holds for every density estimation techniques. Chapter 6 studies deeply density estimation on spaces of Gaussian and multivariate Gaussian distributions.

1.3. Mathematical morphology for non standard images

Similarly to statistics, it is possible to observe how mathematical morphology have progressed with the evolution of the value spaces of images. Mathematical morphology is one of the first sound theories of digital image processing. When dealing with binary images, three neighborhood operators naturally arise: replacing the value of a pixel by the majority, the minimal or the maximal value in the neighborhood. Replacing the value of a pixel

by the majority value in a neighborhood is called a median filter. The two other operators called erosion, ε, and dilation, δ, respectively form the famous couple of operators at the origin of mathematical morphology. This pair of operators present remarkable stability properties. For a suitable choice of neighborhood, the operator formed by the composition of an erosion and a dilation is stable after one iteration, that is to say $\varepsilon \circ \delta = \varepsilon \circ \delta \circ \varepsilon \circ \delta$. By symmetry, or by duality in the morphological vocabulary, the same holds for $\delta \circ \varepsilon$. The composition $\varepsilon \circ \delta$ is called a closing and $\delta \circ \varepsilon$ an opening. This stability property, also called idempotency, together with the scale-space semigroups of openings and closings, are at the heart of morphological image filtering. The erosion and the dilation can be extended to grey-scale images by taking the minimal and maximal value in the neighborhood according to the standard order on reals. This generalization preserves the fundamental properties of stability. The theory is now generalized in terms of complete lattice theory, a lattice being a partially ordered set where couples of elements have a unique supremum and infimum. Pairs of erosion and dilation verifying a pseudo-inverse property called "adjunction" generate idempotent compositions. In this framework, the value space of the image has to be endowed with a complete lattice structure in order to use morphological operators. This is natural for grey-scale images but often problematic for images valued in multidimensional spaces. Colors can be represented in a three dimensional vector space. Standard lattice structures used to process colors all have important drawbacks. Partial orders tend to introduce false colors while total orders are not compatible with the Euclidean distance as Chapter 2 shows.

The idea of formulating the "best" total order for a given image is the object of Chapter 2.

Several parallel works to this thesis aim at defining morphological operators without the notion of complete lattice. Chapter 3 proposes a morphological framework not based on lattices, for images valued in a set of independent categories. Approach introduced by Gronde and Roerdink (2015) proposes to drop the associative law of the lattice structure. Unlike the proposition of chapter 3 their framework enables to work in a continuous space, however the idempotence of openings and closings no longer holds. Meanwhile, work by Carlinet and Géraud (2015) build a tree of shapes on multivariate images without requiring an order structure.

1.4. Thesis organization and main results

This thesis was originally focused on mathematical morphology. In the chapter "Unsupervised Morphology" of the PhD thesis Velasco-Forero (2012), the author proposes to build an order on the value space of an image from the distribution of the values using a notion of statistical depth. In theory, there is no need to compute a density associated to the point cloud of values in order to compute the statistical depth. Indeed, the statistical depth is defined directly from the point cloud of values. However estimating the density of the distribution of the values has an interest in term of computation time of the statistical depth. Trying to apply this idea to order DTI data, we encountered the problem of density estimation of symmetric positive definite matrices. This problem took a growing importance and gradually led to an autonomous part of the thesis on density estimation on Riemannian spaces. Meanwhile our development of mathematical morphology changed direction and the original idea of ordering DTI data according to a statistical depth has been abandoned.

The thesis is thus composed of two independent parts. The first part named "Mathe-

matical morphology on non-standard images" is composed of two chapters:

- Chapter 2 formalizes and demonstrates the irregularity issue of total orders in metric spaces, its main result being the lemma 2.3.1. This problem has been known for a while in the mathematical morphology community, see pioneering work by Chanussot and Lambert (1998) or more recently, by Flórez-Revuelta (2005). The problem is reformulated so as to better take into account the needs of mathematical morphology. Our original formal result states that any total order on a multidimensional space possesses a form of irregularity that generates inconsistent and irregular erosions and dilations. A method is proposed to build total orders adapted to each image, by minimizing the irregularity of the order.

- n-ary images refer in the thesis to images valued in a set of n elements or categories, with no specific mathematical structure. Each value is called a label. n-ary images appear when the studied objects have equivalent semantic levels. Consider for instance a traffic image with blue, red and yellow cars. Blue, red and yellow represent objects of the same semantic level. In this situation, distances and orders on colors present an arbitrary character: saying that blue is closer to yellow than red or saying that blue<yellow<red has no meaning in terms of the studied objects. This situation often appears in pixelwise classified images. It is difficult to address data that are intrinsically unordered using standard morphological operators. Chapter 3 is an attempt to generalize morphology on binary images to discrete n-ary images. The previous considerations already have motivated several works, see Busch and Eberle (1995) and Ronse and Agnus (2005), but led to different directions. The main originality of this work is the introduction of multiple erosions and dilations. The main results of the chapter are,
  - The definition of an erosion ε_i and a dilation δ_i per label;
  - A proof of the idempotence of $\gamma_i = \varepsilon_i \circ \delta_i$ and $\varphi_i = \delta_i \circ \varepsilon_i$. Recall that this results is standard in the binary case by adjunction theory;
  - A proof of semigroup $\varepsilon_{i,R_1} \circ \varepsilon_{i,R_2} = \varepsilon_{i,R_1+R_2}$. In other words, an erosion with a ball of radius R_1 composed by a erosion with a ball of radius R_2 is equivalent to an erosion with a ball of radius $R_1 + R_2$. The result also holds for dilation but is directly deducted from the binary case. This enables to express the n-ary operations as flow propagations or operator composition;
  - The operator $F = (\varphi_1 \circ ... \circ \varphi_n)$, involving different labels, is no longer idempotent. However, $F_k = (\varphi_1 \circ ... \circ \varphi_n)^k$ is stable when k goes to infinity. F is a good candidate to filter out image objects with shape/shape criteria;
  - A first extension to the continuous n-ary case.

The second part entitled "Probability density estimation on Riemannian spaces" is composed of three chapters:

- Chapter 4 introduces a background on probability and geometry for a typical reader from image processing community.

- Chapter 5 addresses the issue of image histogram computation. A formalization of the notion of image histogram in term of density estimation is first proposed. The second part of the chapter is a study on perceptual color image histograms. The perceptual

metric on colors is assumed to be Riemannian. Local scalar products are provided by standard ellipse or ellipsoid datasets, such as the MacAdam ellipses. Usually perceptual histograms are built according to the Euclidean structure of perceptual spaces such as the Lab space. The use of local Euclidean approximations enables to follow more precisely the Riemannian metric without increasing the computational complexity.

- Chapter 6 is a guide to non parametric probability density estimation on spaces of multivariate Gaussian distributions. In other words it provides methods to estimate a density when the set of samples $\{x_i\}$ is a set of Gaussian distributions. This situation is more and more common in signal and image processing. The case of Gaussian distributions also includes the space of positive definite matrices and the hyperbolic space of dimension two. The construction of the estimated density depends on the choice of a metric. The metrics studied in this chapter are the main standard metrics on Gaussian distributions: the Fisher metric, the Wasserstein metric and Euclidean metrics on parameters. The main results are,
  - An expression of kernels for kernel density estimation on multivariate Gaussian distributions when the space is endowed with the Wasserstein metric;
  - An expression of kernels for kernel density estimation on multivariate centered Gaussian distributions when the space is endowed with the Fisher metric. This result is not entirely original, the different factors composing the kernel are already present in the specialized literature, but are here systematically developed;
  - The study of partial quantities (mean, eigenvalues and rotation) under the different metrics.

Part I.

Mathematical morphology on non standard images

2. The irregularity issue of total orders on metric spaces and its consequences for mathematical morphology

Résumé

Nous nous intéressons dans ce chapitre au problème de l'irrégularité des ordres totaux dans les espaces métriques pour la morphologie mathématique. Nous donnons tout d'abord une formulation rigoureuse du problème. Une nouvelle approche est proposée pour y faire face, en adaptant l'ordre à l'image à traiter. Étant donné une image et un ordre total, nous définissons une fonction de coût évaluant l'importance des conflits lors des opérateurs morphologiques. L'ordre proposé est alors une minimisation de cette fonction de coût. Un des atouts de cette méthode est sa généralité : le seul ingrédient nécessaire à la construction de l'ordre est un graphe des distances entre les valeurs présentes dans l'image. L'ordre peut donc être calculé pour toute les images à valeur dans des espaces métriques où l'expression des distances est connue. Nous présentons quelques résultats sur des images couleurs, sur l'imagerie du tenseur de diffusion (DTI), et sur des images à valeur dans le demi plan supérieur de Poincaré.

Ce chapitre est principalement adapté de Chevallier and Angulo (2015).

Abstract

We address in this chapter the problem of irregularity of total orders in metric spaces and its implications for mathematical morphology. We first give a rigorous formulation of the problem. Then, a new approach is proposed to tackle the issue by adapting the order to the image to be processed. Given an image and a total order, we define a cost that evaluates the importance of the conflict for morphological processing. The proposed order is then built as a minimization of this cost function. One of the strength of the proposed framework is its generality: the only ingredient required to build the total order is the graph of distances between values of the image. The adapted order can be computed for any image valued in a metric space where the distance is explicitly known. We present results for color images, diffusion tensor images (DTI) and images valued in the hyperbolic upper half-plane.

This chapter is mainly adapted from Chevallier and Angulo (2015).

2.1. Introduction

Since its emergence in the sixties, mathematical morphology has become one of the major theory of nonlinear image processing. Originally used for binary images in Matheron (1974), the theory has followed the technical evolution of computer science which has

enabled the manipulation of more and more complex images, see Serra (1982). The set theory was sufficient to study binary images. Later, the emergence of gray-scale images required the introduction of the notion of order. From the works by Serra (1988) and Heijmans and Ronse (1990), the theory of mathematical morphology is now fully based on the lattice theory. Value spaces endowed with a total order form the most comfortable framework for morphological processing. However, if it is natural to endow gray-scale images with a total order, the task is more difficult when the pixel values do not have a one-dimensional structure, see Serra (1993) for a pioneering discussion on the topic. Indeed, we show that the information contained in a total order is too weak to completely represent the value space. In many situations, the use of a partial order such as a product order on a vector space is preferred. Using a product order is equivalent to processing components independently. Then this order structure becomes natural but some information about the geometry of the original value space is lost. Both choices present a loss of information.

The interest of product orderings in mathematical morphology is still a matter of recent research, see for instance Gronde and Roerdink (2013) and Burgeth and Kleefeld (2013), mainly by considering the geometric and invariant properties of the underlying space.

Other recent works on partial ordering for morphological operators on vector images were motivated by the need of taking into account some prior information to order vectors: either by learning the order from training samples in Velasco-Forero and Angulo (2011) or by building the order according to the outlier distribution in Velasco-Forero and Angulo (2012). Therefore, partial ordering can become image adaptive and consequently leads to more relevant morphological operators.

On the other hand, the study of total orders remains mainly limited to lexicographic orders, see Hanbury and Serra (2001); Angulo (2007, 2010); Aptoula and Lefèvre (2007). Since the emergence of multivariate images, very few papers have addressed the problem of total orders in a general way. Previous approaches of total ordering mainly focused on building regular orders on hypercubes see Chanussot and Lambert (1998) and Flórez-Revuelta (2005). Beyond the theoretical result, the two ideas of the present work are the following. Firstly, given an image we restrict the value space to the values that are actually present in the image. This idea is already present in Lezoray et al. (2008) where they order the values present in the image according to a Hamiltonian path. Although the motivation is different, the idea is also present in Ledda and Philips (2005). Secondly, we take into account the locations of the values in the image. Given an image, these considerations enable to find more regular orders. One can therefore consider our approach as an image adapted total ordering. We formulate this task as an optimization problem which cannot be solved using classical optimization techniques. Thus, we introduce a hierarchical recursive algorithm aiming at finding an approximated solution.

The chapter is organized as follows.

- Section 2.2 provides a background on notations and basic notions.
- A discussion on existing ordering strategies on vector spaces is given in Section 2.3. In particular, a strong theoretical result points out the limitation of total orderings in terms of discontinuity of vector morphological operators.
- Section 2.4 gives the motivation for introducing an order adapted to a given image. A cost function is introduced to measure the quality of a total order regarding the topological conflict.

- The recursive algorithm developed to minimize the cost function is fully described in Section 2.5.
- A short discussion on the invariance properties of the adapted order is given in Section 2.6.
- Section 2.7 presents some results of morphological image processing using the image adapted total order. We present results for color images, diffusion tensor images (DTI) and images valued in the hyperbolic upper half-plane.
- Conclusions and perspectives close the chapter in Section 2.8.

2.2. Notations and recalls

We set here a few notations and recall elementary operators of mathematical morphology, see for instance Serra (1982); Soille (2004). Let us consider an image I as a function:

$$I : \begin{cases} \Omega \to \mathcal{V} \\ p \mapsto I(p) \end{cases}$$

where Ω is the support space of pixels p: typically $\Omega \subset \mathbb{Z}^2$ or $\mathbb{Z}^3$ for discrete images. The pixel values of the image belong to the space $\mathcal{V}$. Typically we have $\mathcal{V} \subset \mathbb{R}$ for grey-scale images, $\mathcal{V} \subset \mathbb{R}^n$ for multivariate vector images, or $\mathcal{V} \subset \mathcal{M}$ for manifold valued images. We address here images where $\mathcal{V}$ is any metric space. Points in $\mathcal{V}$ will generally be called colors. We denote by $I(\Omega) \subset \mathcal{V}$ the set of colors of $\mathcal{V}$ present in the image $I(p)$.

Unlike linear processing mainly based on linear convolution (i.e., weighted average), mathematical morphology is based on sup and inf-convolution. The numerous choices of convolution kernels lead to a large variety of processing. Thus, the two basic operators of mathematical morphology are the erosion and the dilation of an image $I(p)$, $I : \Omega \to \mathbb{R}$, by $B \subset \Omega$ given respectively by

$$\varepsilon_B(I)(p) = \inf_{q \in B(p)} \{I(q)\}, \tag{2.1}$$

$$\delta_B(I)(p) = \sup_{q \in \check{B}(p)} \{I(q)\}, \tag{2.2}$$

where the set B defines the structuring element (the equivalent of the convolution kernel), $\check{B}$ is the transpose of B (i.e., symmetric set with respect to the origin), and $B(p)$ defines the neighborhood of p according to the shape of B. Note that here we only focus on flat structuring elements. Other morphological filters, such as the opening $\gamma_B(I)$ and closing $\varphi_B(I)$, are obtained by composition of dilation and erosion; i.e,

$$\gamma_B(I) = \delta_B\left(\varepsilon_B(I)\right), \tag{2.3}$$

$$\varphi_B(I) = \varepsilon_B\left(\delta_B(I)\right). \tag{2.4}$$

More evolved filters and transforms are obtained from composition of openings/closings. Another nonlinear operator particularly useful in image denoising, also based on ordering, is the median filter:

$$m_B(I)(p) = \operatorname*{med}_{q \in B(p)} \{I(q)\}. \tag{2.5}$$

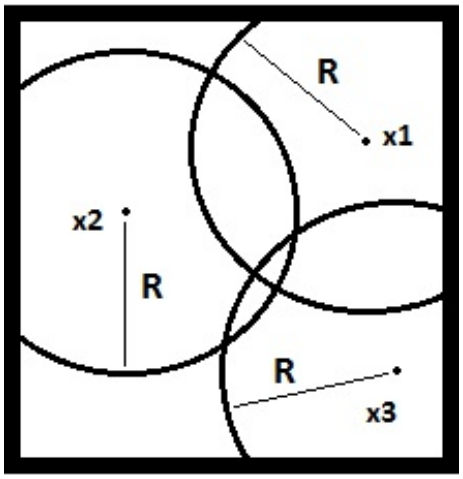

Figure 2.1.: The metric space X is here a rectangle of the Euclidean plan

2.3. Existing total orders

The problem of total ordering for multivariate images is a relatively well known problem in mathematical morphology. The essential difficulty is that the topology induced by a total order on a multidimensional space cannot reproduce the natural topology of the vector space, see discussion in Chanussot (1998). Arising as a milestone limitation, we have the following theoretical result.

Lemma 2.3.1 *Let (X, d) be a metric space endowed with a total order $\leq$. Suppose that there exist a positive real number R and three points $x_1, x_2, x_3 \in X$ such that $d(x_i, x_j) > R$ for $i \neq j$ and that the complementary $B^C(x_i, R)$ of each ball $B(x_i, R)$ is connected, as in Fig. 2.1. Then for all $r > 0$, there exist three points a, b and c in X such that*

$$\begin{cases} a \leq b \leq c, \\ d(a, b) \geq R, \\ d(a, c) \leq r. \end{cases}$$

Proof We can assume that $x_1 < x_2 < x_3$. We argue by contradiction and assume that there exists $r > 0$ such that for all $a, b, c \in X$ one at least of the three above conditions doesn't hold. The idea of the proof is to show that under this assumption, we have either $\forall y \notin B(x_2, R), y < x_2$ or $\forall y \notin B(x_2, R), y > x_2$. This would contradict the existence of the triple (x_1, x_2, x_3) with $x_1 < x_2 < x_3$. By our assumption, for any triple (a, b, c) of points in X, the three inequalities $a \leq b$, $d(a, b) \geq R$ and $d(a, c) \leq r$ imply $c < b$. Consider the set E of points a in $B^C(x_2, R)$ such that $a \leq x_2$. Recall that a topological space B is connected iff the only open and closed subsets A of B are $A = \emptyset$ and $A = B$. We use this definition with $B = B^C(x_2, R)$ and $A = E$ to exhibit a contradiction. The considered topology is now the induced topology on $B^C(x_2, R)$. Making use of our assumption with the triple (a, b, c), where a is in E, $b = x_2$ and $c \in B(a, r) \cap B^C(x_2, R)$, we see that for all a in E, $B(a, r) \cap B^C(x_2, R) \subset E$. Since $B(a, r) \cap B^C(x_2, R)$ is a ball of $B^C(x_2, R)$, E is an open subset of $B^C(x_2, R)$. Again by our assumption, if a point $c \in B^C(x_2, R)$ is not in E then $B(c, r) \cap B^C(x_2, R)$ cannot contain a point $a \in E$, otherwise the point c which is in

$B(a,r)$, would be in E. It follows that $B^C(x_2,R)\setminus E$ is also an open subset of $B^C(x_2,R)$, hence E is an open and closed subset of $B^C(x_2,R)$. Since $B^C(x_2,R)$ is connected, the open and closed subsets E and $B^C(x_2,R)\setminus E$ cannot be both nonempty, a contradiction with the existence of the points x_1 and x_3.

The order is going back and forth between point a, point b and point c. Note that in the particular case where $X=\mathbb{R}^{n>1}$, this result is not a corollary of the fact that there is no continuous bijection between $\mathbb{R}^{n>1}$ and $\mathbb{R}$. Indeed, not every order on $\mathbb{R}^{n>1}$ can be represented by a bijection on $\mathbb{R}$ (this would contradict the least-upper-bound property). This lemma tells us that for any total order in $\mathbb{R}^n$, functions supremum of two points, $\mathbb{R}^n\times\mathbb{R}^n\to\mathbb{R}^n:(x,y)\mapsto\sup(x,y)$ and infimum $\inf(x,y)$ present high irregularities with respect to the Euclidean metric, i.e. for any $r,R>0$ there exists $x_0,x_1,y_0,y_1\in\mathbb{R}^n$ such that $||(x_0,y_0)-(x_1,y_1)||<r$ and $R<||\sup(x_0,y_0)-\sup(x_1,y_1)||$. Let I be an image composed of two objects of distinct colors with smooth boundaries. The first object is composed of two values x_0 and x_1 with $d(x_0,x_1)<r$. Points x_0 and x_1 are two versions of a color $\mathcal{C}_x$. The second object is composed by an unique value y with $R<d(x_0,y)$, representing color $\mathcal{C}_y$. Assume that $\sup(x_0,y)=x_0$ and $\sup(x_1,y)=y$. Regarding the object, the morphological dilation and erosion are inconsistent and might introduce aliasing on boundaries. The lemma is valid in any metric space and has strong negative implications. Namely, given a total order, it is always possible to find an image where the erosion and dilation introduce important irregularities. An illustration of this phenomenon is given by the following toy example depicted in Fig. 2.2. The RGB image is composed of 3 different colors. Two close black ($\mathcal{C}_x$) represented by $a=(0,0,0)$ and $b=(1,0,0)$, and a blue ($\mathcal{C}_y$) represented by $c=(0,0,255)$. According to the lexicographic ordering on coordinates, we are precisely in the situation described in the lemma. Fig. 2.2 shows us the result of a dilation using a 3×3 square as structuring element.

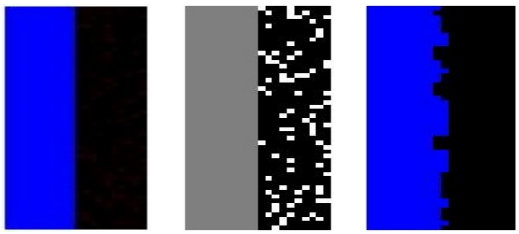

(a) (b) (b)

Figure 2.2.: (a) Original image, (b) grey-scale representation of the lexicographic order, (c) lexicographic dilation.

Through the study of space filling curves, see Sagan (1994), the work Chanussot and Lambert (1998) proposes several total orders on $\mathbb{Z}^n$ that preserve as far as possible the notion of neighbourhood. For each point of the discretized multidimensional vector space, the spatial neighbourhood and the neighbourhood in the chain of the order are compared, as shown in Fig. 2.3 extracted from Chanussot and Lambert (1998).

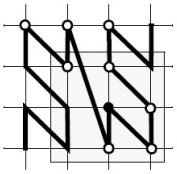

Figure 2.3.: Quantitative evaluation of the topological distortion (figure extracted from Chanussot and Lambert (1998)). Five white dots are included in the gray square, the weight assigned to the black point is 5/8. These weights are averaged over all the points of the region of interest, as the RGB cube for color images.

The average comparison between the spatial neighbourhood and the neighbourhood in the chain of the order gives a measure of the preservation of the neighbourhoods. This measure enables us to compare every total orders on $\mathbb{Z}^n$. Chanussot and Lambert (1998) shows that the neighborhoods are significantly better preserved by the Peano curve and the bit-mixing paradigm than by the lexicographic order. Despite the fact that the Peano curve gives slightly better results than the bit-mixing strategy paradigm, authors of Chanussot and Lambert (1998) chose not to use the Peano curve because this order does not preserve the vectorial structure; i.e., in two dimensions, the point $(0, 8)$ can be greater than the point $(15, 15)$.

In the present work we chose to focus the study only on topological considerations, without taking into account the algebraic structure of the value space. Note that the algebraic structure is often present by default without link with the physical quantity.

2.4. An order adapted to a given image

2.4.1. Motivation

As we have discussed above, it is not possible to create a total order that preserves neighbourhood on a multidimensional space. The philosophy of Chanussot and Lambert (1998) and Flórez-Revuelta (2005) is to try to minimize the difference between spatial neighbourhoods and neighbourhoods in the space of order. However it is possible to push this idea further. Even for the best total order in the sense of the measure defined in Fig. 2.3, our lemma tells us that the processing of a particular image can give highly irregular results. As a consequence, it might be more interesting to look for the best order being given an image, instead of looking for the best order in general. Indeed, restricting the evaluation of a total order to a particular image, largely enhance the potential quality of the order. An order on a multidimensional space can present important jumps that might not affect the processing of a particular image.

Firstly, given a specific image, the evaluation of a total order only gives importance to colors that lay in the image. Introduction of other (false) colors is due to the use of partial ordering. With total ordering, the values are chosen among the vectors of the input and as such there is no new colors. Indeed, no other colors are introduced by flat erosion or dilation. Decreasing the size of the set to be ordered can lead to more regular orders. More precisely, it becomes easier to find orders that avoid having two points close in the color space and far in the order. Secondly, problematic triplets of values described in the lemma can be tolerated to the extend that values do not occur in same zones of the image. Indeed,

if two values never occur simultaneously in a structuring element, there relative position in the order will not affect the processing.

As explained in Chanussot and Lambert (1998), it is clear that if points are close in the color space, they should remain close in the order. However, the reverse is not always required. Let us consider a binary image represented on the real line, where black and white are not represented by 0 and 1 but by 0 and 10. This situation presents two points close in the order chain and distant regarding the metric of the color space. However, this does not introduce any irregularity in the morphological operators. Thus the cost defined in Fig. 2.3 penalizes situations that are not problematic for mathematical morphology.

In order to transpose this topological intuition of *closeness both in value space and in spatial space should imply closeness in the order*, we introduce a cost function to be minimized by the total ordering.

Both works Ledda and Philips (2005) and Lezoray et al. (2008) proposed to restrict the order to values present in the image. Authors of Ledda and Philips (2005) propose to order values according to their importance in the histogram of the image. While being a natural idea, it does not rely on any geometrical consideration on the value space. Authors of Lezoray et al. (2008) propose to build local orders based on optimal Hamiltonian path. They do not order the entire set of values of the image, and the associated dilation and erosion are not strictly speaking morphological operators. However the present work is in the continuation of their idea. We try here to build an order on the full value set present in the image, optimal with respect to a topological criteria.

2.4.2. Cost function

Given an image, we would like to define a cost function that measures the quality of the order regarding rank based operators. We first need to quantify the criteria of the previous section. We define the following notion of co-occurrence of values a and b:

$$C_I(a,b) = Card\{p \in \Omega \mid \exists q \in B(p),\ \ (I(p), I(q)) \in \{(a,b),(b,a)\}\}. \tag{2.6}$$

The computation of co-occurrences involves that one has fixed a typical size/shape of the structuring element B which will be used in subsequent processing. The motivation for introducing a notion of co-occurrence in the cost function is the following: if two colors are not present in the same zones of the image, they cannot produce aliasing. Let a and c be two close points in $\mathcal{V}$ according to a given metric distance $d(a,c)$, such that $a < c$. a and c are two versions of a first color $\mathcal{C}_1$. Let b be a third point far from a and c. b represents a second color $\mathcal{C}_2$, distinct from $\mathcal{C}_1$. If $C_I(a,b) = 0$ or $C_I(c,b) = 0$, no irregularity is created by the triplet a, b, c regardless of the chosen order. Indeed in the image the color $\mathcal{C}_2$ is in presence of at most one version of the color $\mathcal{C}_1$. However if $C_I(a,b)$ and $C_I(c,b)$ are both significant, the color $\mathcal{C}_2$ is neighboring of the two versions of $\mathcal{C}_1$. It is then important that $a < c < b$ or $b < a < c$. Thus the quality of a total order can be measured by evaluating to what extend the following property is respected:

$$\left\{\begin{array}{c} a < c \\ C_I(a,b) \text{ and } C_I(c,b) \text{ significant} \\ d(a,c) \text{ small} \\ d(a,b) \text{ significant} \end{array}\right. \tag{2.7}$$

$$\Longrightarrow a < c < b \text{ or } b < a < c$$

Adjectives "significant" and "small" induce an increasing or decreasing behavior of the cost function with respect to the different quantities. Let us endow $I(\Omega)$ with a total order $\leq$. We can define the following quantity:

$$P_I(\leq) = \sum_{\substack{a,b,c \in I(\Omega), \\ a<c<b}} g(a,b,c) \cdot (C_I(c,a) \wedge C_I(c,b)), \tag{2.8}$$

where $\wedge$ stands for min, with,

$$g(a,b,c) = f\left((d(c,a) \wedge d(c,b)), d(a,b)\right),$$

and $f(\cdot,\cdot)$ an increasing function according to the first variable and decreasing according to the second.

Given an image $I(p)$, this adapted cost function is more tolerant for some specific orders than the cost function defined in Chanussot and Lambert (1998). The cost function P_I has been designed to represent as well as possible what is expected of an order. One of its main advantage is to take into account the above condition, see Eq. (2.7), weaker than the one required in Chanussot and Lambert (1998), see Fig. 2.3. However, as a standard image often contains more than ten thousand different colors, this cost function presents the serious drawback of not being computable. Thus, given two orders, it is difficult to compare them using this cost function. Nevertheless, it is possible to try to minimize this cost function using a recursive procedure, without computing globally the cost for the full set of points.

2.5. Minimisation of the cost function

2.5.1. Overview of the algorithm

The idea is to divide the set $I(\Omega)$ into a collection $\mathcal{C}$ of clusters and to compute an optimal order on $\mathcal{C}$ considering each cluster as a single point. Then each cluster is ordered individually. The orders on individual clusters are merged according to the order on $\mathcal{C}$ to obtain an order on $I(\Omega)$. The point of the clustering is to make this operation possible by reducing the number of parameters of the minimization. Here are the main steps of the algorithm.

- Perform a clustering of the data $I(\Omega)$ in a given number n of clusters: $\{Cluster_i\}_{1\leq i\leq n}$.
- Order the n clusters according to $C(i,j)$, $D(i,j)$, and the corresponding cost function P where
  - $D(i,j)$ represents the minimum distances between clusters i and j.
  - $C(i,j)$ represents the co-occurrence of clusters i and j in different structuring elements.
- Perform the same procedure recursively on each cluster.
- After a given threshold, stop the recursion. For each remaining cluster, order all its points according to a criterion based on distances to previous and next clusters.

- Merge the orders: for x in $Cluster_i$ and y in $Cluster_j$, $x < y$ if and only if $Cluster_i < Cluster_j$, or $i = j$ and $x < y$ in $Cluster_i$.

In this short description we did not mention a significant source of complications. Indeed, when the recursive procedure is applied to $Cluster_i$, one has to take into account neighbour clusters. Indeed, at the first step of the recursion, nothing has to be taken into account except the considered set of colors $I(\Omega)$. However, unlike the set $I(\Omega)$, $Cluster_i$ can no longer be considered as isolated from the rest of the color values. If there exist colors c_k in $Cluster_i$ and colors c_l in $Cluster_j$ such that $d(c_k, c_l)$ is small, then it is not possible to order $Cluster_i$ without taking into account $Cluster_j$. To order $Cluster_i$, one needs to know the set of its neighbour clusters and their relative ordering. These are the main ingredients of the *recursive algorithm*. The algorithm uses five functions.

- A function `Neighbour`
  `input`: a list S_1 of clusters and a $Cluster_{index}$ in S_1
  `output`: a sublist $S_{neighbour}$ of N_{neighb} elements of S_1

  The elements of $S_{neighbour}$ are selected from the list S_1 according to their co-occurrence with $Cluster_{index}$. The co-occurrence between the clusters are computed from the image $\tilde{I}$ where $\tilde{I}(p) = i$ if and only if $I(p) \in Cluster_i$. We impose $Cluster_{index} \in S_{neighbour}$. N_{neighb} is an arbitrary parameter of the algorithm, see the result section.

- A function `IndexCutting`
  `input`: a list S_1 of clusters and a $Cluster_{index}$ in S_1
  `output`: a list $S_{Cluster_{index}}$ of N_{Clust} sub-clusters
  of $Cluster_{index}$. N_{Clust} is an other arbitrary constant of the algorithm, see the result section.

  Call the function `Neighbour` to replace S_1 by the shorter list $S_{neighbour}$. For each cluster $Cluster_i$ in $S_{neighbour}$, find the point c_i in $Cluster_{index}$ that minimizes the distance to $Cluster_i$. If $i = index$, let c_i be the barycentre of $Cluster_{index}$. The set $\{c_i\}$ is completed with random points in $Cluster_{index}$ to reach a minimum number of points. Perform a clustering of $Cluster_{index}$ using a k-means algorithm initialized with the $\{c_i\}$. Return $S_{Cluster_{index}}$ the new set of clusters.

- A function `MainOrder` (To be used when the recursion depth is $\leq$ threshold)
  `input`: a list S_1 of clusters with an order $<_1$, a $Cluster_{index}$ in S_1 and a list $S_{Cluster_{index}}$ of sub-clusters of $Cluster_{index}$
  `output`: an order $<_2$, on
  the list $S_2 = S_1 \cup S_{Cluster_{index}} \setminus Cluster_{index}$.

  Perform a minimization of P on S_2, using the co-occurrences of clusters $C(Cluster_i, Cluster_k)$ and the distance $D(Cluster_i, Cluster_k)$, such that the new total order $<_2$ is compatible with the initial order $<_1$. If $Cluster_i$ and $Cluster_j$ are in S_1,

  $$(Cluster_i <_1 Cluster_j) \Rightarrow (Cluster_i <_2 Cluster_j)$$

  Furthermore, for $Cluster_i$ in $S_{Cluster_{index}}$ and $Cluster_j$ in $S'_1 = S_1 \setminus Cluster_{index}$,

  $$Cluster_i <_2 Cluster_j \Leftrightarrow Cluster_{index} <_1 Cluster_j$$

  The situation is summarized in Fig. 2.4.

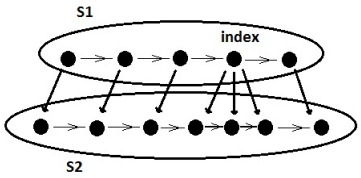

Figure 2.4.: Construction of $<_2$ (see description of the algorithm).

- A function `SimpleOrder` (To be used when the recursion depth is > threshold)
  `input`: a list S_1 of clusters with an order and a $Cluster_{index}$ in S_1
  `output`: an order on $Cluster_{index}$.

  Select in S_1 the nearest neighbour of $Cluster_{index}$ greater than $Cluster_{index}$ and the nearest neighbour lower than $Cluster_{index}$. Order elements of $Cluster_{index}$ according to their distance to the two clusters selected during the previous step.

- A function `FuncRecursive`

```
FuncRecursive;
input  : An ordered list of clusters (S_1, <_1) , an Index Cluster Cluster_index, a
         depth D
output: An order on Cluster_index given by a list
if D > threshold then
 | SimpleOrder (Cluster_index)
else
 | S'_1 ← Neighbour(S_1, Cluster_index);
 | S_index ← IndexCutting(Cluster_index);
 | S_2 ← S'_1 ∪ S_index \ Cluster_index;
 | (<_2) ← MainOrder(S'_1, Cluster_index, S_index);
 | listorder ← empty list;
 | for all the sons in S_index taken increasingly for <_2 do
 |  | listorder ← Concat (listorder, FuncRecursive (S_2, son, D + 1))
 | end
end
```

The main program is simply the function `FuncRecursive` called with the list S_1 reduced to $I(\Omega)$ with
$Cluster_{index} = I(\Omega)$ and $D = 0$.

Fig. 2.5 presents the different steps on a simple example.

The non-deterministic aspect of the clustering step induces that two successive minimizations of the cost function do not necessarily produce identical orders. Furthermore, the order defined on $\varepsilon(I)(\Omega)$ is not necessarily induced by the order defined on $I(\Omega)$. Thus, in order to preserve standards properties of mathematical morphology such as the idempotence of composed dilation and erosion, $\varepsilon \circ \delta$, one has to use the same order for all operators on the complete processing chain.

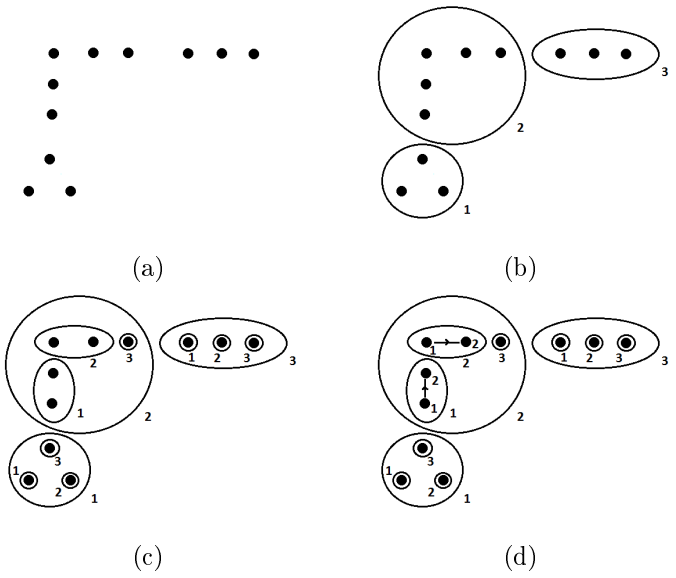

Figure 2.5.: This example shows the successive steps of the algorithm launched with *threshold* = 2. (a) Initial point cloud; (b) *depth* = 1 ; (c) *depth* = 2; (d) *depth* = *threshold* : points in each cluster are ordered according to their distance to neigbhour clusters.

2.5.2. Optimization over the permutation space

We note that, according to this algorithm the set to be ordered is no longer the set of all colors present in the image, but the set of clusters S at each level of recursion. The cost function P can be calculated if the cardinal n of S is reasonable. If the set S does not exceed $n = 9$ elements, it is conceivable to calculate the cost P of each permutation and select the minimum. However, this procedure requires an important computational time. An alternative solution consists in optimizing the cost function iteratively. It is well known that from a permutation it is possible to reach any other permutation by composing transpositions. Given a permutation, we select the transposition that minimizes the cost function, and we repeat the process until we fall in a minimum. As there is no proof that the cost function is convex, the minimum might only be a local minimum. However, under $n = 9$ elements it is possible to compare the result with the global exploration. In all tested situations with $n = 9$ elements, the optimization by transposition does not fall in a local minimum.

2.5.3. Computational cost

Given an order, the computation of the cost function P involves a sum over

$$\binom{n}{3} = \frac{n(n-1)(n-2)}{6},$$

terms where n is the number of points. If N_{Clust} is the number of clusters produced by the function `IndexCutting`, the minimization of P is performed in $(N_{Clust})^{depth}$ times. Our Matlab code launched on a standard computer gives $1,25.10^{-2}s$ for a single computation of P and 1.8 for its minimization over a set of 10 points (the second figure is an average

over 50 runs). The choice of the clustering algorithm might induce significant variation in the global computational cost. Using a k-means algorithm, $N_{clust} = 10$, $N_{neighb} = 15$ and $depth = 2$, our Matlab code requires 8 minutes to process an image of size 256×256 pixels composed of 40000 different colors.

2.6. Order invariances

The cost function $P(I)$ does not depend on the coordinates of colors but only on their mutual distances and their co-occurrences. As the notion of co-occurrences remains unchanged under bijective transformations, the cost $P(I)$ is invariant under any isometric transformation. Note that this is not the case for the lexicographic and the bit-mixing order. However, the choice of a particular function $f(\cdot, \cdot)$ has an influence on the class of invariances. For instance, the function

$$f(min_{i,j}, d(c_i, c_j)) = \frac{min_{i,j}}{d(c_i, c_j)}, \tag{2.9}$$

where $min_{i,j} = d(c, c_i) \wedge d(c, c_j)$ also provides invariance of $P(I)$ under homothetic transformations. We can also note that for the function

$$f(min_{i,j}, d(c_i, c_j)) = \frac{min_{i,j}^{\alpha}}{d(c_i, c_j)^{\beta}}, \tag{2.10}$$

homothetic transformations simply result in the multiplication of $P(I)$ by a positive constant. Consequently, as the notion of minimum is invariant under increasing transformation, the minimization of $P(I)$ should remain relatively stable. If T is an isometric (or homothetic) transformation of the value space $\mathcal{V}$, and ϕ a morphological operator $\{\Omega, \mathcal{V}\} \rightarrow \{\Omega, \mathcal{V}\}$, then T and ϕ commute for any image I, i.e.,

$$\phi(T(I)) = T(\phi(I)).$$

2.7. Results of morphological image processing

For each image, the minimization is launched with the following parameters:

- Function f:

  $$f(x, y) = x \cdot G(y), \tag{2.11}$$

  where $G(.)$ is a gate function with linear decrease, see Fig. 2.6. The parameter of the gate G is set to $1/10$ of the diameter of the value space $\mathcal{V}$.

  The cost P becomes:

  $$P = \sum_{a < c < b} (d(a, c) \wedge d(a, b)) \times G(d(a, b)) \times (C_I(c, a) \wedge C_I(c, b)). \tag{2.12}$$

- At each level of the recursion, the set S is divided into $n \leq 10$ clusters, assuming that its size allows it.
- The number of neighbors selected by the function `Neighbour` is set to 15 at most, depending on the number of existing clusters. During the `MainOrder` function the cost P is evaluated on set of size 25 at most.
- The recursion is stopped when $depth = 2$.

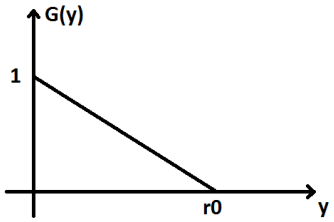

Figure 2.6.: The function $f(x, y)$ used in the experiment is of the form $x \cdot G(y)$. The parameter r_0 of the gate G is related to the parameter r of the Lemma. Triplet of values $a < b < c$ such that $d(a, c) \geq r_0$ has no influence in the cost function.

2.7.1. 2D point sets

Before going into morphological operations on images, we present here a comparison of the lexicographic, the bit-mixing, and the adapted order on a few sets of points in a square of $\mathbb{R}^2$, without introducing the notion co-occurrences. Fig. 2.7 presents an example of ordering of 20 points, according to the lexicographic order, the bit-mixing order, and the adapted order. The adapted order is directly minimized over the set of permutations, without going through the clustering step.

Over a set of 200 points, the cost P can still be computed, which enables the comparison between the different orders. The influence of the number of clusters, the depth of the recursion and the spatial configuration is studied in Fig. 2.8. Each experiment focuses on a specific geometric pattern. Points are randomly drawn around the geometric patterns. The image adapted order is computed for a recursion depth varying from 1 to 4. At each depth, the adapted order and its cost are computed for a number of clusters varying on odd numbers from 2 to 12. For each configuration, the costs of the different total orders are averaged over 20 runs. The second column of Fig. 2.8 presents the costs depending on the depth. At each depth, the cost is averaged over the choice of the number of clusters. The third column represents the evolution of the cost according to the number of clusters at $depth = 2$. Fig. 2.8 enables to draw several conclusions. Except exceptional situations, the recursive minimization provides better results than the lexicographic and the bit-mixing orders. It succeeds in adapting to each situation. The decrease of the cost function with the depth lays mainly during the two first steps of the recursion. This fact motivated the choice of $depth = 2$. Fig. 2.8 illustrates that the quality of the different orders strongly depends of the configuration of the point cloud. Indeed, although the lexicographic order seems to be the less interesting order, it is possible to find configurations where it performs better than the bit-mixing and the adapted order. The behavior of the adapted order with respect to the number of clusters is more difficult to handle that the influence of the depth. For rectilinear distributions, increasing the number of cuts of the point set seems to slightly decrease the quality of the adapted order. When the geometrical configuration has an intrinsic number of clusters, overestimating the number of clusters does not guarantee to obtain optimal results. In case of uniform distribution of the points, the cost seems to be globally decreasing with the number of clusters. The absence of assumption on the spatial configuration pushes to chose a significant number of clusters. In the next experiments on images, 10 clusters has been judged as the best compromise between accuracy and computation time.

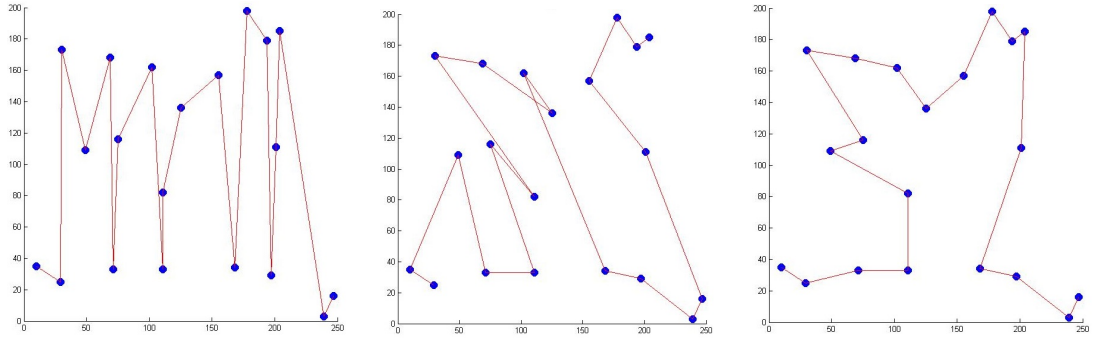

(a) (b) (c)

Figure 2.7.: A set of 20 points is ordered according to the lexicographic order in (a), the bit-mixing order in (b), and the adapted order in (c).

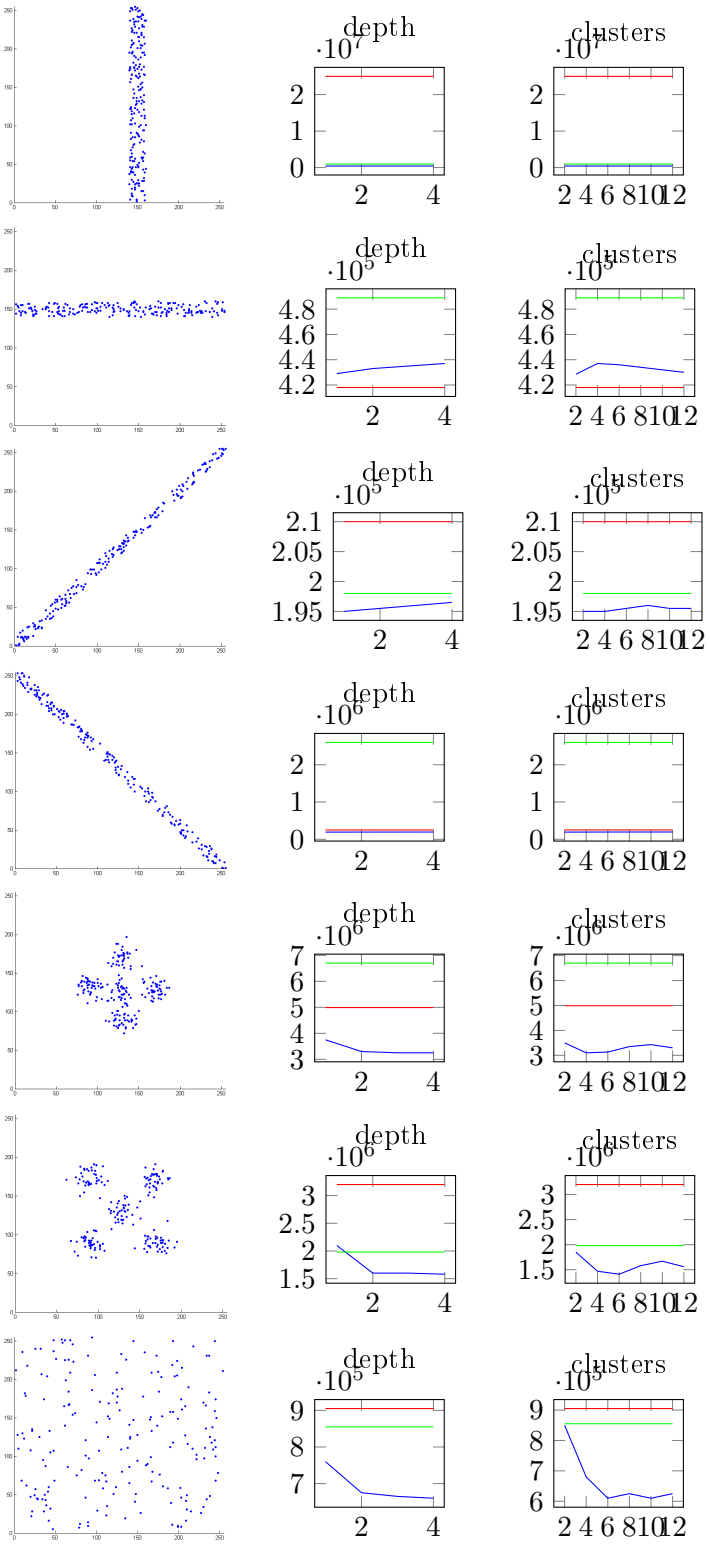

Figure 2.8.: Each row corresponds to a spatial distribution. The first column represents a typical point cloud. The second column represents the average dependence of the depth of the recursion. The third column represents the influence of the number of clusters at $depth = 2$. In the different plots, the lexicographic order is in red, the bitmixing in green and the adapted order in blue.

2.7.2. Color Imaging

Unfortunately unlike the previous 2D point sets, the set of colors in a standard image is too large to enable the computation of the cost of a total order. It is then not possible to numerically evaluate the performance of the optimization. However, in the following examples the dependence on the order of the irregularities produced by the morphological processing is visually appreciable.

We present results of morphological color processing obtained for two different RGB images. The first one is a microscopic blood vessel from a fluorescence microscope, the second one is a natural color image. For both of them P is minimized by the recursive algorithm discussed in previous sections. The distance $d(.,.)$ between colors is the Euclidean distance of the RGB color space.

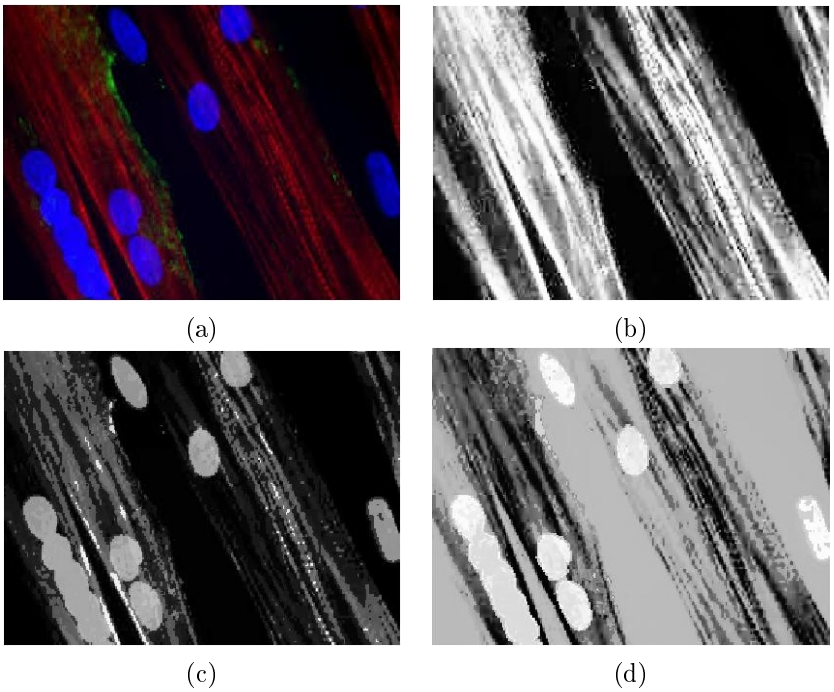

Figure 2.9.: Projection of the total order on the image support: (a) original RGB image $I(p)$; (b) lexicographic order (R $\rightarrow$ G $\rightarrow$ B); (c) bit-mixing order; (d) our image adapted total order.

Fig. 2.9 represents the RGB color image together with projections of the total orders on the image support for the three studied total orders. As we can see, the blue spots are totally invisible to the lexicographic ordering. Note that the lexicographic ordering starts with red, and ends with blue. The bit-mixing paradigm and the image adapted total order give different results but are both coherent with the original image.

Using each total order, we can compute morphological color operators. Fig. 2.10 gives the corresponding openings and closings using as structuring element a square of 7×7 pixels. As expected, the lexicographic ordering produces important aliasing around the blue spots. The bit-mixing paradigm and the image adapted total order give different results but both preserve the regularity of boundaries.

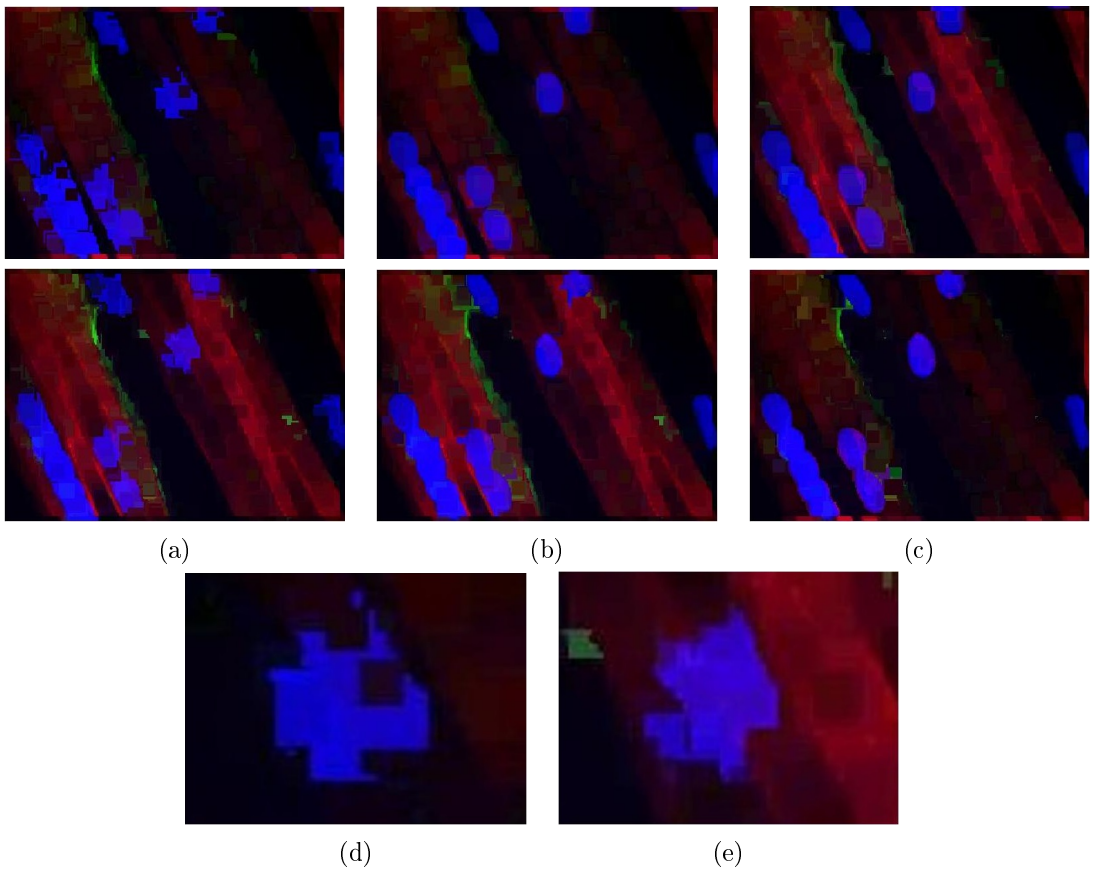

Figure 2.10.: Morphological processing of image from Fig. 2.9: openings (top row) $\gamma_B(I)$ and closings (middle row) $\varphi_B(I)$ using lexicographic order in (a), bit-mixing order in (b) and our image adapted total order in (c). Structuring element is B is a square of 7×7 pixels. Zoom on aliasing introduced by the lexicographic order in (d) and (e).

The same study is performed on the second example. Fig. 2.11 provides the original RGB image and the image representation of the three orders. Unlike the previous examples, lexicographic order is able to distinguish all the interesting objects of the image. Furthermore, it seems to give an order smoother than the bit-miximg paradigm and the image adapted total order. As we can observe on Fig. 2.12, which gives the result of median filtering, the regularity of the grey-scale projection of the lexicographic order is only "an illusion". Opening and closing operators of size 9×9 are depicted in Fig. 2.13 and Fig. 2.14. Both lexicographic and bit-mixing order present aliasing on blue and yellow boundaries. The adapted order presents a slight aliasing in the background, which is not present with lexicographic and bit-mixing orders. However on this example, our image adapted total order is the only one of the three orders that provides satisfying results in terms of regularity for opening and closing.

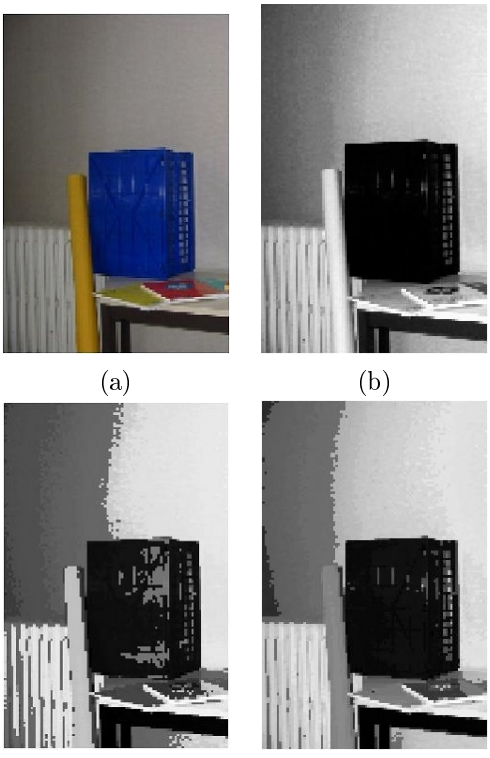

(a) (b)

(c) (d)

Figure 2.11.: Projection of the total order on the image support: (a) original RGB image $I(p)$; (b) lexicographic order; (c) bit-mixing order; (d) our image adapted total order.

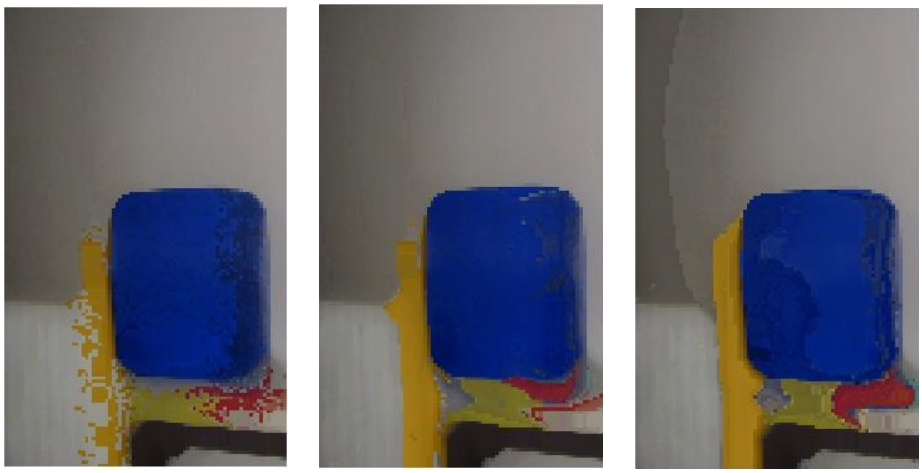

(a) (b) (c)

Figure 2.12.: Median filtering Fig. 2.11 using lexicographic order in (a), bit-mixing order in (b) and our image adapted total order in (c).

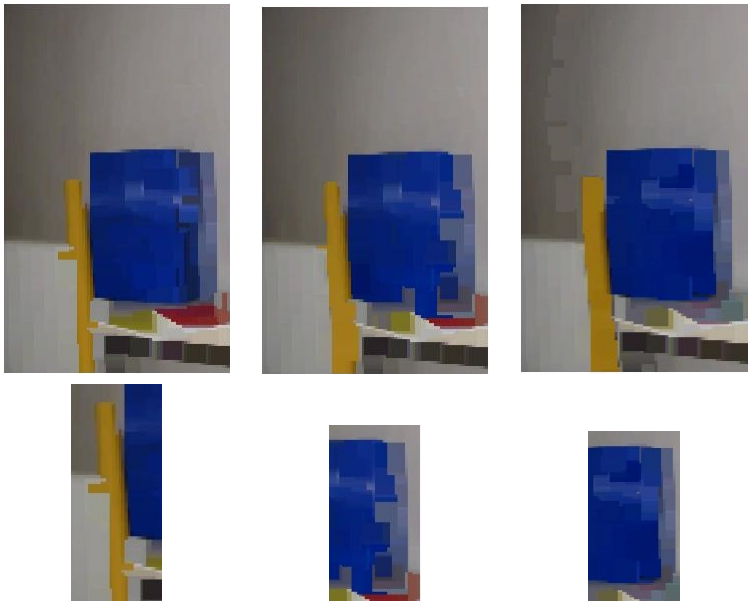

(a) (b) (c)

Figure 2.13.: Closings of image from Fig. 2.11, using lexicographic order in (a), bit-mixing order in (b) and our image adapted total order in (c). The structuring element is a square of size 9×9. Top row: closings, bottom row: zoom on main aliasing zones.

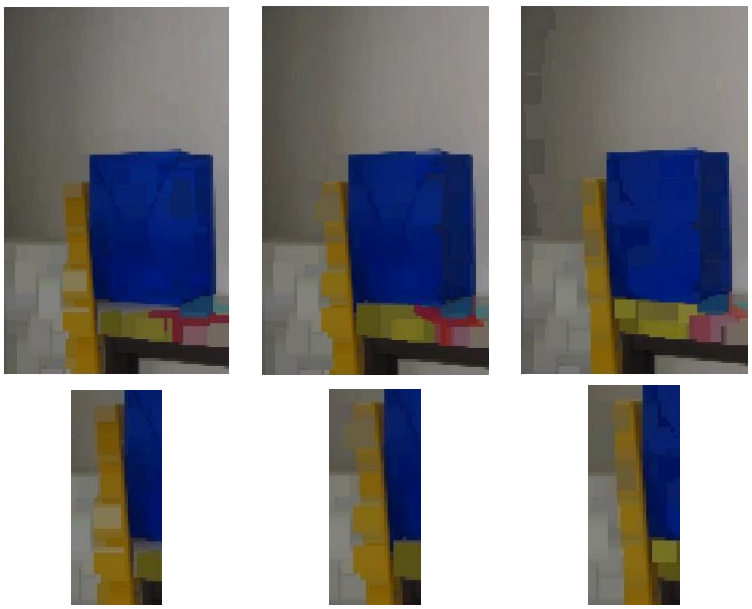

(a) (b) (c)

Figure 2.14.: Openings of image from Fig. 2.11, using lexicographic order in (a), bit-mixing order in (b) and our image adapted total order in (c). The structuring element is a square of size 9×9. Top row: openings, bottom row: zoom on main aliasing zones.

2.7.3. Diffusion Tensor Imaging (DTI)

We recall that in DTI, each pixel of the image contains a 3×3 symmetric positive definite matrix, that is a point in the space SPD(3). A matrix of SPD(3) can be represented as an ellipsoid, thus DTI image can be visualized using ellipsoids, see the example in Fig. 2.15.

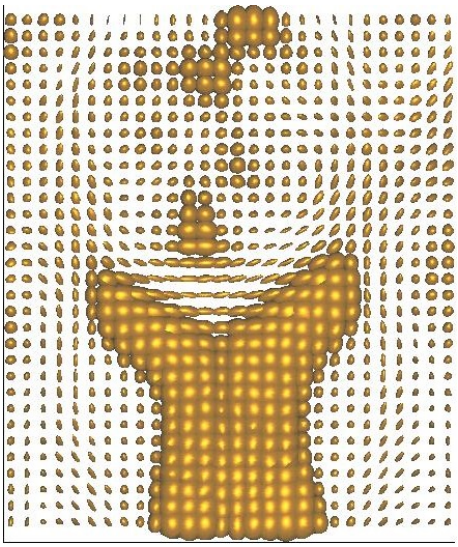

Figure 2.15.: Example of DTI image.

Morphological processing of DTI data often requires a transition to uni-dimensional data, see for instance the approach introduced in Rittner and de Alencar Lotufo (2009). The use of Loewner order was introduced in Burgeth and Kleefeld (2013), so as to preserve the matrix information. The Loewner order presents the important drawback of not inducing a lattice. Accepting to lose standard properties of mathematical morphology, it was proposed in Burgeth and Kleefeld (2013) a reasonable definition of the supremum and the infimum which enables the construction of a dilation and an erosion. The proposed supremum and infimum are continuous however the Loewner order is only a partial order. Gronde and Roerdink (2015) recently proposed a new morphological framework to address DTI data. However, this new framework is also based on partial order relationships. Similarly to the color section, we restrict the comparison of the irregularities to total orders.

As the matrix is symmetric, it can be parametrized by the 6 upper-diagonal coefficients, see Fig. 2.16.

C1	C2	C3
	C4	C5
		C6

Figure 2.16.: Standard parametrization of a symmetric matrix.

This parametrization of SPD(3) induces a lexicographic and a bit-mixing order, named respectively $LEX1$ and $BMIX1$ in what follows. Another approach consists in representing SPD(n) matrices with rotations and eigenvalues. A symmetric matrix can be

diagonalized in an orthonormal basis. Then each SPD(3) matrix can be represented by 3 eigenvalues and a rotation matrix. Using any angular representation of the rotation matrix, the SPD(3) matrix can be represented by 3 eigenvalues and 3 angles. Let us consider such a parametrization where the eigenvalues are sorted decreasingly. The associated lexicographic order is named $LEX2$ and the bit-mixing order $BMIX2$.

Using our framework, the image adapted total order is calculated according to two metrics:

- the metric associated with the Frobenius scalar product, i.e., $\langle A, B\rangle = \mathrm{tr}(AB^t)$;
- the Log-Euclidean metric introduced in Arsigny et al. (2006), i.e., $d(A, B) = \| \log(A) - \log(B)\|$.

In every processing example, the structuring element B is a square of 5×5 pixels. Fig. 2.17 and Fig. 2.18 show respectively the results of openings and closings for the different total orders. As in standard morphological processing, opening and closing removes small objects respectively from the image foreground and background.

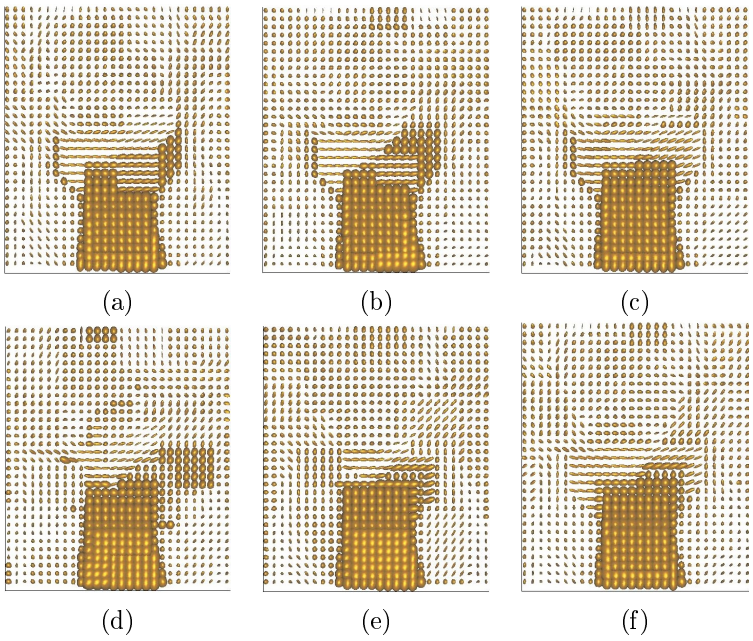

Figure 2.17.: Morphological opening of a DTI image (original image in Fig. 2.15): using $LEX1$ in (a), $LEX2$ in (b), $BMIX1$ in (c), $BMIX2$ in (d), image adapted total order using Frobenius norm in (e), and using Log-Euclidean norm in (f).

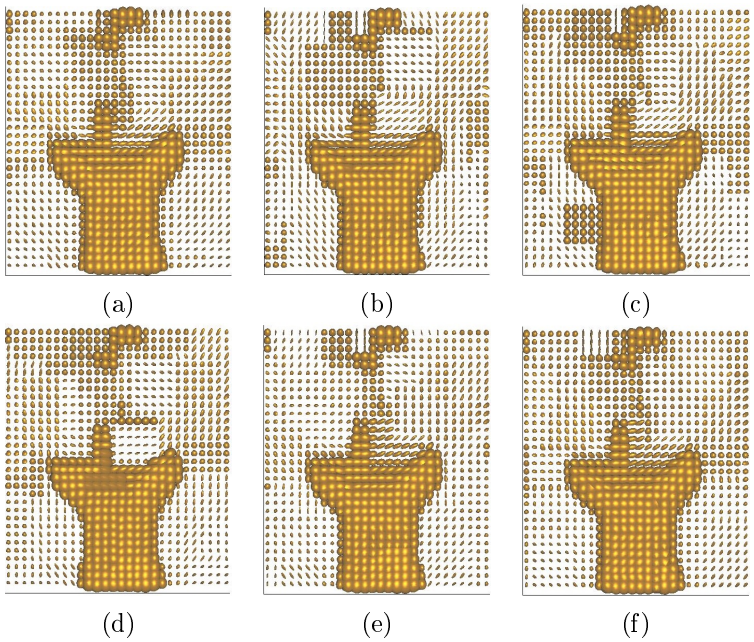

Figure 2.18.: Morphological closing of a DTI image (original image in Fig. 2.15): using $LEX1$ in (a), $LEX2$ in (b), $BMIX1$ in (c), $BMIX2$ in (d), image adapted total order using Frobenius norm in (e), and using Log-Euclidean norm in (f).

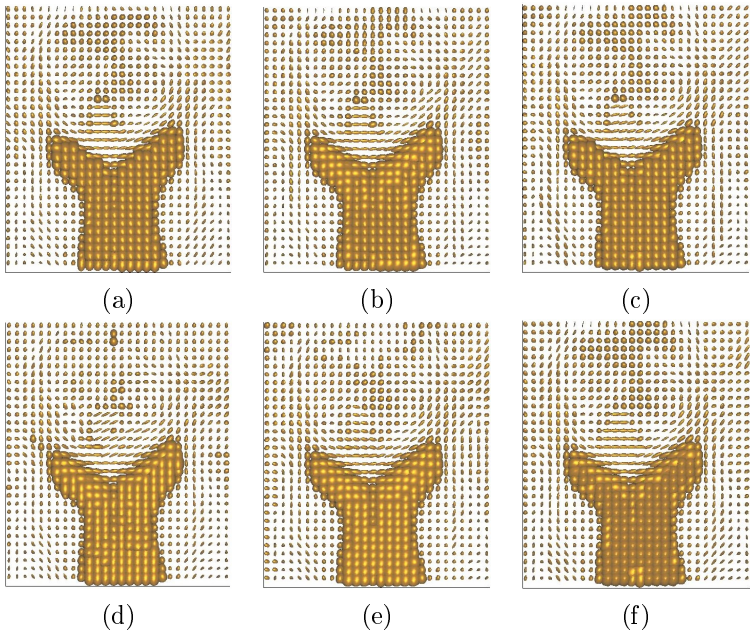

Figure 2.19.: Morphological reconstruction of a DTI image (reference image in Fig. 2.15 and marker images are the openings in Fig. 2.17): using $LEX1$ in (a), $LEX2$ in (b), $BMIX1$ in (c), $BMIX2$ in (d), image adapted total order using Frobenius norm in (e), and using Log-Euclidean norm in (f).

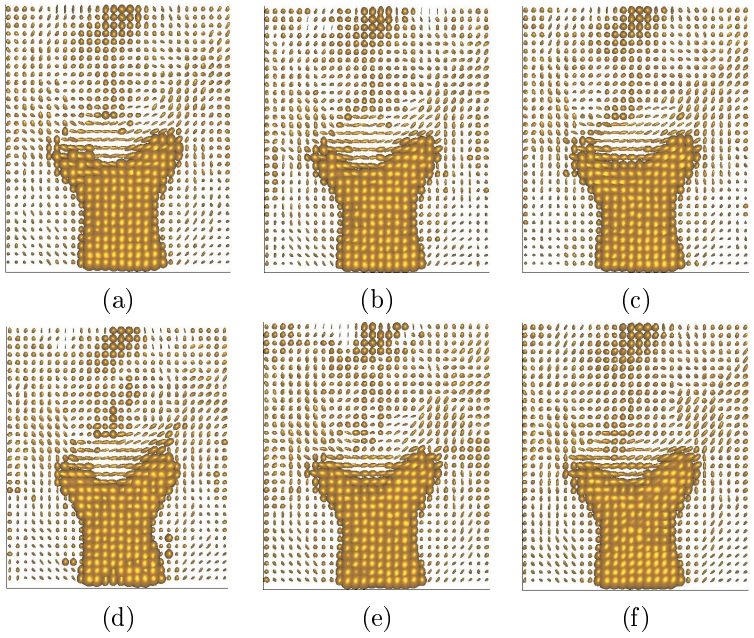

Figure 2.20.: Median filtering of a DTI image (reference image in Fig. 2.15 and marker images are the openings in Fig. 2.17): using $LEX1$ in (a), $LEX2$ in (b), $BMIX1$ in (c), $BMIX2$ in (d), image adapted total order using Frobenius norm in (e), and using Log-Euclidean norm in (f).

Using the openings from Fig. 2.17 as markers, the geodesic reconstruction Soille (2004) has also been computed. The results are depicted in Fig. 2.19 and as expected, this operator based on geodesic dilations will recover the contours of the objects which have not been suppressed by the opening.

Finally, the comparative result of median filtering is presented in Fig. 2.20.

Visually, the image adapted total orders give results at least as good as other orders, and better ones in several situations. The choice of the Euclidean or log-Euclidean metric does not seem to introduce important changes.

2.7.4. Total order in the Poincaré upper half plane

As discussed in previous sections, the proposed order is based only on the notion of distance, independently from the algebraic structure. This framework suits perfectly to the case of images whose pixels lies in a Riemannian manifold, where the distance is known. We present here a situation where the image is valued in such a manifold.

Mathematical morphology for images valued in uni-dimensional Gaussian laws has already been studied in Angulo and Velasco-Forero (2014). One of the most common distances on Gaussians is the distance induced by the Riemannian metric called the Fisher metric. This is a simple example where the value space is not a vector space. In the case of one dimensional Gaussian parametrized by their means μ and their standard deviation σ, the distance between two Gaussians G_1 and G_2 is given by:

$$d(G_1, G_2) \quad = \quad \cosh^{-1}(1 + \frac{\frac{1}{2}(\mu_1 - \mu_2)^2 + (\sigma_1 - \sigma_2)^2}{2\sigma_1\sigma_2}).$$

In the mean/standard deviation half-plane, the shortest paths of the Fisher metric are half ellipses centered on the μ-axis. The mean/standard deviation half-plane endowed with the Fisher metric is a Poincaré upper half plane of curvature $-\frac{1}{2}$.

The studied example is a time lapse sequence of grey-scale images from a retina. At each pixel, we dispose of 20 successive acquisitions. By assuming a Gaussian distribution on the successive acquisition, we obtain a Gaussian valued image represented in Fig. 2.21. The observation of the distribution of Gaussians in the upper half plane presented in Fig. 2.21(c) strongly suggests to endow the space of Gaussian laws with the Fisher metric. The proposed solution to obtain total orders is currently the only total order framework that takes the geometry of the space into account.

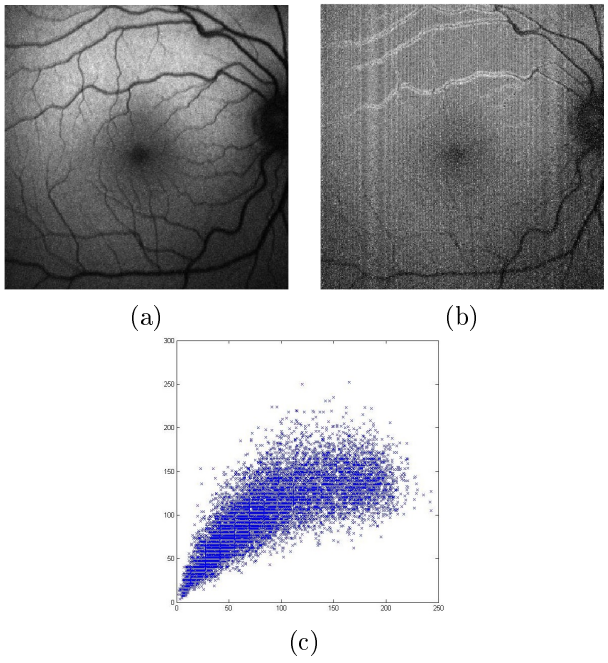

Figure 2.21.: Time lapse sequence of images from the retina: (a) mean image, (b) standard deviation image, (c) corresponding point cloud in the uni-dimensional Gaussian laws space.

Note that to minimize P in this case study, one has now to perform a clustering on the hyperbolic upper half plane. We choose to use the model centroids proposed in Galperin (1993) which enables a simple implementation of the k-means algorithm. The adapted order is computed with respect to the Euclidean and the Fisher metric. Figures 2.22- 2.25 show that the Riemannian framework might lead to better results. The central black spot of the image, one of the anatomical elements of the retina, is for instance easier extracted using Riemannian framework, see Fig. 2.24 and 2.25.

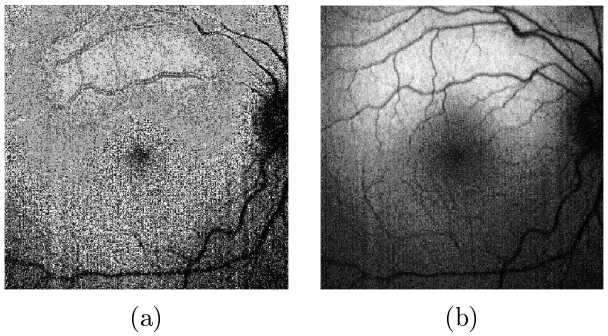

Figure 2.22.: Visualization of order map: (a) total order projection using Euclidean metric, (b) total order projection using Fisher metric.

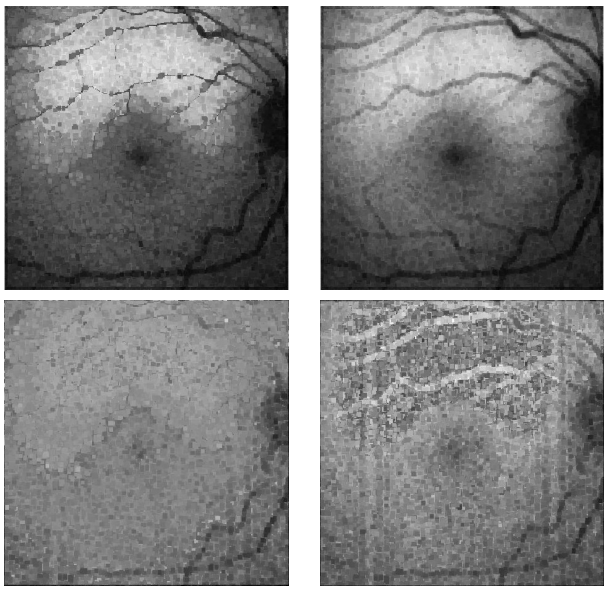

(a) (b)

Figure 2.23.: Top row: mean image, Bottom row: standard deviation image. (a) Closing in the Euclidean framework, (b) closing in the Riemannian framework.

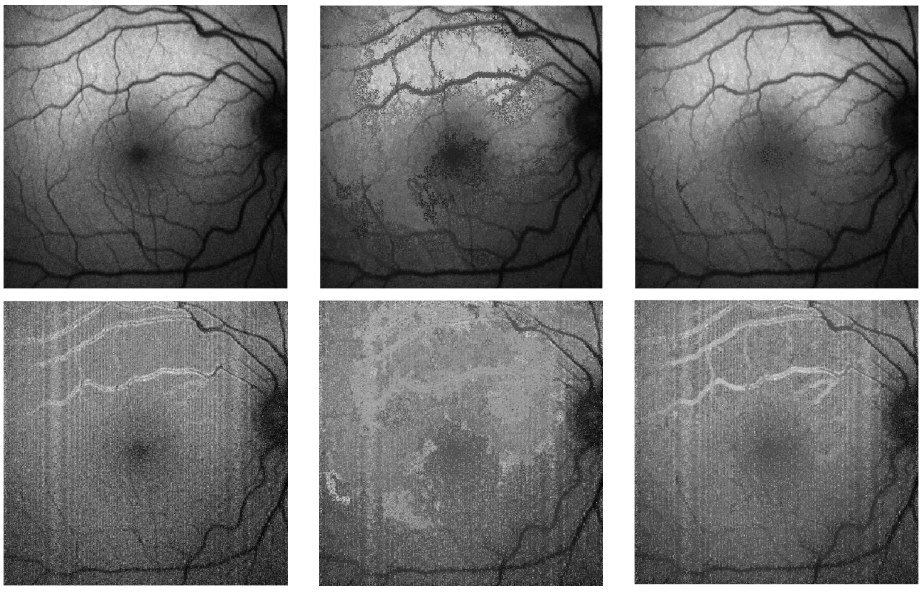

(a) (b) (c)

Figure 2.24.: Top row: mean image, Bottom row: standard deviation image. (a) Original image, (b) closing by reconstruction in the Euclidean framework, (c) closing by reconstruction in the Riemannian framework.

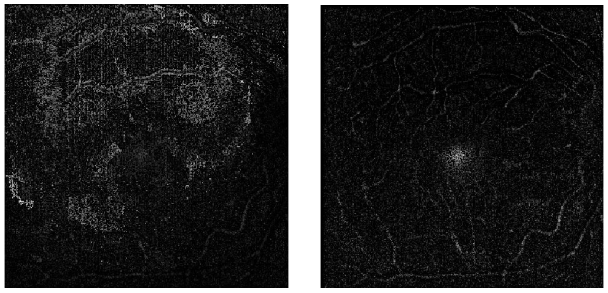

(a) (b)

Figure 2.25.: Residue image of the closing by reconstruction working on the Euclidean framework in (a), and on the Riemannian framework in (b).

2.8. Conclusions and Perspectives

To our knowledge, this work is the first study that rigorously formulates and demonstrates the irregularity issue of total orders in metric spaces. Given a total order, the introduced Lemma tells us there are always images where morphological operators introduce irregularities and aliasing. Given an image and a total order, we exhibit the problematic triplets of values. The identification of problematic triplets leads us to a cost function measuring the quality of a total order regarding the back and forth problem. Due to the restriction of the problem to a specific image and to a precise identification of the problematic situations, we know that the order minimizing the cost function produces less irregularities than the existing propositions, namely Chanussot and Lambert (1998); Flórez-Revuelta (2005).

As the cost function is not explicitly computable due to the size of the set to be ordered, we propose a recursive minimization procedure based on successive clustering of the set to order, so as to find a total order adapted to the image. Future work might be led on boundaries regularity measurement so as to provide a quantitative measure of the quality of the minimization. A reasonable lead would be to evaluate the irregularity of the level-sets channel-wise and merge the results.

A strength of the proposed framework is its generality. Indeed, for a large majority of images, the value space is equipped with a metric. The image adapted order can thus be computed for almost any image. Note that for a given type of images, the adapted metric is often open to debate, see for instance Ledoux et al. (2013).

We have shown the interest of our method on some examples, however the bit-mixing total order proposed in Chanussot and Lambert (1998) remains an interesting solution in the Euclidean case. Indeed our current minimization procedure does not yet provide average results significantly better than the bit-mixing, while the latter is independent from the image and requires no pre-processing.

The minimization of the cost function is subject to several potential improvements, on its accuracy on the one hand, and on the computation speed on the other hand. The use of techniques such as evolutionary algorithms might for instance enable to significantly enhance the number of clusters at each step of the recursion, see Flórez-Revuelta (2015). Thus our future research will be essentially focused on the minimization of the cost function.

3. N-ary morphology

Résumé

La morphologie mathématique pour les images binaires peut être entièrement d'écrite dans le cadre de la théorie des ensembles. Cependant, cette théorie n'est pas suffisante pour d'écrire la morphologie des images en niveau de gris. Ce type d'images requiert l'introduction de la notion d'ordre. Plus généralement, la morphologie mathématique est maintenant d'écrite à l'aide de la notion de treillis. Depuis quelques décennies, plusieurs travaux visent à appliquer des opérateurs morphologiques aux images multivariées telles que les images couleurs. Ces travaux ce basent principalement sur la notion d'ordre vectoriels. Cependant, aucune de ces tentatives n'a donné de résulats entièrement satisfaisants. Au lieu de viser directement le cas multivarié, nous proposons ici une extension de la morphologie binaire à une situation intermédiaire : les images à valeur dans un ensemble fini de labels non ordonnés. Nous proposons ensuite une extension à un cas continu.

Ce chapitre est principalement extrait de Chevallier et al. (2015a).

Abstract

Mathematical morphology on binary images can be fully described by set theory. However, it is not sufficient to formulate mathematical morphology for grey scale images. This type of images requires the introduction of the notion of partial order on grey levels, together with the definition of sup and inf operators. More generally, mathematical morphology is now described within the context of the lattice theory. For a few decades, attempts are made to use mathematical morphology on multivariate images, such as color images, mainly based on the notion of vector order. However, none of these attempts has given fully satisfying results. Instead of aiming directly at the multivariate case we propose first an extension of binary mathematical morphology to an intermediary situation: images composed of a finite number of independent unordered labels. We propose then a second extension to a continuous case.

This chapter is mainly adapted from Chevallier et al. (2015a).

3.1. Introduction

The philosophy of the previous chapter was to adapt the data to the theory of mathematical morphology. The aim was to put an interesting structure on data, a total order in this case, that would enable the use of morphological operators. However in the multivariate case, addressing data only through an order, partial or total, is often not fully compatible with the original structure. Product orders for instance are equivalent to marginal processing. Lemma 2.3.1 shows that total orders on $\mathbb{R}^{n>1}$ are never compatible with distance information in morphological processing.

Instead of adapting the structure of the data to the mathematical morphology theory, one could try to adapt the theory to the data. This chapter do not aim directly at the full multivariate case. We propose to adapt the morphological framework firstly to images valued in n independent categories and secondly to images valued in mixtures of this categories.

Let us start by quoting Serra (1982). "Objects in space generally have three dimensions, which are reduced to two dimensions in a photograph or on the retina. In this projection the luminance of the point located along a line oriented directly away from the viewer are not summed, because most physical objects are not translucent to light rays, in the way they would be to X ray but are opaque. Consequently, any object that is seen hides those that are placed beyond it with respect to the viewer: this self-evident property is a basic one. In fact it serves as a starting point for mathematical morphology, since, whenever we wish to describe quantitatively phenomena in this domain, a set-theoretic approach must be used.". Here is one of the main differences between (nonlinear) image processing and sound processing. In sound processing, audio sources of the same environment are summed, whereas opaque objects of a 3D scene hide themselves. In image processing we might not want to mix colors of neighbor objects, as does the linear convolution, because this mix might have no meaning in terms of objects. Thus this key idea can be reformulated by saying that mathematical morphology is based on the reduction and the extension of the surface of the different objects over their neighbors. That leads naturally to the two basic morphological operators in binary images.

Binary images. In such images, there are only two kinds of objects: black or white objects. Two dual and adjoint operators have been defined: the erosion and the dilation. The erosion extends the black objects over the white objects, the dilation extend the white objects over the black objects. Formally, a binary image can be seen as a support set Ω, and X a subset of Ω. Set X represents the white points of the image and the complement $\complement X = X^c$ the black ones. We assume here that Ω is a subset of $\mathbb{R}^n$, which is generally the case. Let B be a subset of Ω called the structuring element. We assume that Ω has a translation operation. The erosion $\varepsilon_B(X)$ and the dilation $\delta_B(X)$ of X according to a structuring element are defined as follows, see Matheron (1974):

$$\varepsilon_B(X) = \bigcap_{y\in B} X_{-y} = \{p \in \Omega \,:\, B_p \subset X\} = \left\{x \,:\, \forall p \in \check{B}, x \in X_p\right\}, \tag{3.1}$$

$$\delta_B(X) = \bigcup_{y\in B} X_y = \{x+y \,:\, x \in X, y \in B\} = \left\{p \in \Omega \,:\, X \cap \check{B}_p \neq \emptyset\right\}, \tag{3.2}$$

where $\check{X} = \{-x \,:\, x \in X\}$ is the transpose of X (or symmetrical set with respect to the origin O) and $X_p = \{x+p \,:\, x \in X\}$ the translate of X by p. We note that erosion and dilation correspond respectively to the Minkowski substraction and addition of set X by B; i.e., $\varepsilon_B(X) = X \ominus B$ and $\delta_B(X) = X \oplus B$. For the sake of simplicity, we limit the rest of our notation to symmetric structuring elements: $B = \check{B}$.

Grey-scale images. With the appearance of grey-scale images, mathematical morphology was reformulated in terms of inf and sup convolution where the kernel is the structuring element B, see Serra (1982). That corresponds to the so-called flat case, where the kernel is just a "shape". It is also possible to use structuring functions, but this is out

of the scope of this work. An image is now considered as a function I defined as

$$I : \begin{cases} \Omega \to V \\ p \mapsto I(p) \end{cases}$$

where V is the set of grey-levels, which can be generally assumed as a subset of the real line $V \subset \mathbb{R}$. Grey-scale flat erosion and dilation of I by structuring element B are now defined as follows:

$$\varepsilon_B(I)(p) = \inf_{q \in B_p} \{f(q)\}, \tag{3.3}$$

$$\delta_B(I)(p) = \sup_{q \in B_p} \{f(q)\}. \tag{3.4}$$

In this classical framework, each grey-level is not fully considered as an independent label (i.e., a different category) but simply as an intermediary level between black and white. This point of view is actually justified when interesting objects of the images are local extrema. A mathematical way to see this formalism is to note that flat erosion and dilation commute with processing upper level set by upper level set, that is operators (3.3) and (3.6) are respectively equivalent to

$$\varepsilon_B(I)(p) = \sup \left\{h \in \mathbb{R} \, : \, p \in \left(X_h^+(I) \ominus B\right)\right\}, \tag{3.5}$$

$$\delta_B(I)(p) = \sup \left\{h \in \mathbb{R} \, : \, p \in \left(X_h^+(I) \oplus B\right)\right\}, \tag{3.6}$$

where $X_h^+(I)$ denotes the upper level set at height h of function I; i.e.,

$$X_h^+(I) = \{p \in \Omega \, : \, I(p) \geq h\}.$$

(a) (b) (c)

Figure 3.1.: Grey-level morphological processing: (a) original image I, (b) closing $\varphi_B(I)$, (c) opening $\gamma_B(I)$.

Let us see what happens in the situation depicted in Fig. 3.1. It corresponds to process a rather simple grey-level by a closing and an opening. We recall that the closing of I by B is the composition of a dilation followed by a erosion; i.e., $\varphi_B(I) = \varepsilon_B(\delta_B(I))$. The opening is just the dual operator; i.e., $\gamma_B(I) = \delta_B(\varepsilon_B(I))$. Closing (resp. opening) is appropriate to remove dark (resp. bright) structures smaller than the structuring element B. This behavior is based on the fact that the dilation "reduces" dark structures by B while the erosion "restores" the dark structures which are still present.

In the current example, it is not possible to remove the central grey spot using erosion and dilation with B larger than the spot size. This grey spot is not considered as an interesting object in itself but simply as an intermediary value between the black object and the white object. If this assumption is often coherent, this is not always the case.

Let consider the grey-scale image in Fig. 3.2(a). In this image, each grey level has the same semantic level: each represents a different component, sometimes called a phase. However, in the morphological grey-scale framework, the grey is processed as an intermediary level. It is possible to replace each grey level by a color, see Fig. 3.2(b). We would like then to process both images using the same approach.

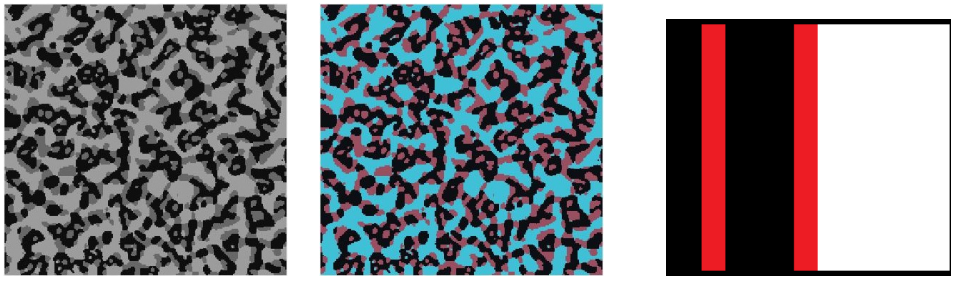

(a) (b) (c)

Figure 3.2.: (a) A three independent grey-level image and (b) the same image where the grey values has been replaced by colors. (c) Example of multivariate (color) image.

Multivariate images. For multivariate images, no canonical framework has yet appeared. Most processing consist in endowing the structure with a partial order relationship. The structure has to be a complete lattice in order to define erosion and dilation in terms of inf and sup. The notion of order induces the notion of intermediary level as in the precedent framework. It is rather intuitive to consider grey as an intermediary between black and white. The situation is different for colors: using the product ordering in the RGB space red is an intermediary color between black and white. This assumption often has no meaning in term of image processing. In an image, the color red usually has the same semantic level as black or white: the color red represents a class of objects, it is not a transition between black objects and white objects. If the phenomenon already exists in the grey-scale framework, it turns out now to be a significant problem, see example in Fig. 3.2(c). As the red usually has a real signification in terms of a particular class of objects, it is very natural to try to remove the red spot, which is not possible using generic classical morphology. The more the image has a complex semantic structure, such as a color image, the more it is difficult to find a lattice structure which makes every interesting object an extremum.

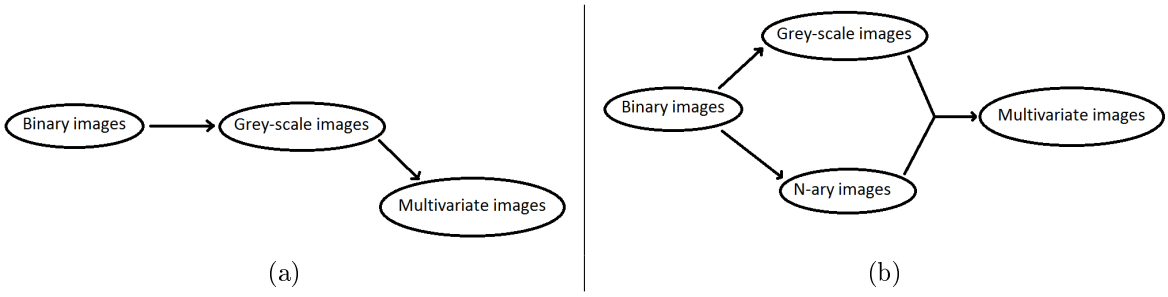

Figure 3.3.: From binary morphology to multivariate morphology: (a) classical framework, (b) need of the intermediate n-ary morphology.

Aim and chapter organization. The historical way of extending the mathematical morphology theory from binary to multivariate images is summarized in Fig.3.3(a). However, the gap between grey-scale and multivariate is much more significant than the gap between binary and grey-scale. As discussed previously, the grey-scale structure only

enables to define intermediary levels between two references. This structure is obviously too weak to describe a multivariate information. Avoiding the notion of order becomes one of the goals of recent developments of mathematical morphology. Authors of Gronde and Roerdink (2015) propose to lighten the lattice structure, while authors of Carlinet and Géraud (2015) build a tree of shapes on multivariate images without requiring a total orders. Although very different from these propositions, our work is based on the same underlying idea.

Before extending mathematical morphology to multivariate images, we might first want to define a coherent approach for mathematical morphology with n independent unordered labels, without considering them as intermediary levels. This is the aim of the first part of the chapter. The difference between the frameworks can be interpreted in term of a change of metric on the value space:

- grey-scale framework: $\forall(i,j),\ \ d(label_i, label_j) = |i-j|$;
- $n-$ary framework: $\forall(i,j),\ \ d(label_i, label_j) = 1$.

This work is not the first to consider the problems of classical mathematical morphology for images composed of independents labels. Authors of Busch and Eberle (1995) and Ronse and Agnus (2005) have very similar motivations but the development we propose is different. In contrast to operators proposed in Busch and Eberle (1995), Ronse and Agnus (2005) or labeled openings from Hanbury and Serra (2001), we are interested in filling gaps left by anti-extensive operators. We note that the theory of morphological operators for partitions introduced in Ronse (2013) and hierarchies of partitions in Meyer (2011) are not compatible with our framework.

The second part of the chapter aims to extend the morphology defined in the first part on independent (discrete) labels to a continuous case. This part of the theory is less elaborated than the discrete case but provides already the key ingredients for future developments. The continuous setting involves to work with images valued in the $(n-1)$-simplex. The approach adopted here will be based on a fuzzy representation and limited initially to the edges of the simplex, i.e., only two labels (or categories) contributes at each pixel. We should point out that our fuzzy framework is not related to the theory of fuzzy morphology as developed for instance by Deng and Heijmans (2002); Bloch (2009). In parallel to the approach considered here for the $(n-1)$-simplex, another recent work by Franchi and Angulo (2015) has considered the extension of morphological operators using partial orders on this space.

The rest of the chapter is organized as follows.

- A proposition of n-ary morphological operators and a study of their theoretical properties is done in Sections 3.2 and 3.3.
- Some applications to image filtering are discussed in Section 3.4.
- Section 3.5 proposes a continuous extension of n-ary morphology for images arising from fuzzy classifications.
- Section 3.6 concludes and closes the chapter.

3.2. n-ary morphological operators

Let us come back to the key idea of mathematical morphology is to reduce and extend objects over their neighbors. In the case of binary images, two operations were introduced: the erosion extends the black over the white and the dilation extends the white over the black. In a general way, we would like to allow to reduce and extend the surface of each label of object. This makes four theoretic operations in the binary case, reduced to two in practice due to the coincidence of certain operations: reducing the black is the same as extending the white and conversely. This duality is one of the basic principle of binary morphology.

3.2.1. Dilation and erosion of label i

Let I be an n-ary image defined as

$$I : \begin{cases} \Omega \to \{1, 2, \cdots, n\} \\ p \mapsto I(p) \end{cases}$$

In the $n-$ary case, it seems natural to try to introduce the corresponding pair $(\varepsilon_i, \delta_i)$ of operators for each label i. Erosion ε_i is the operator that reduces the surface of the objets of label i, and dilation δ_i the operator that extends the label i. Above $n > 2$, we unfortunately lose the duality between operations, such that the number of elementary operators is then equal to $2n$. Let us formulate more precisely these operators.

The *dilation of label i on image I by structuring element B* presents no difficulty:

$$\delta_i(I; B)(x) = \begin{cases} I(x) & \text{if } \forall p \in B_x, I(p) \neq i \\ i & \text{if } \exists p \in B_x, I(p) = i \end{cases} \tag{3.7}$$

Operator $\delta_i(I; B)$ implies that if the label i is not present in the structuring element B_x, $I(x)$ is not modified. Thus $\delta_i(I; B)$ extends objects of label i over their neighbors. The case of the erosion presents more theoretical difficulties. Indeed, if we want to reduce the objects of label i, we need to decide how to fill the gaps after the reduction, see problem in Fig. 3.4.

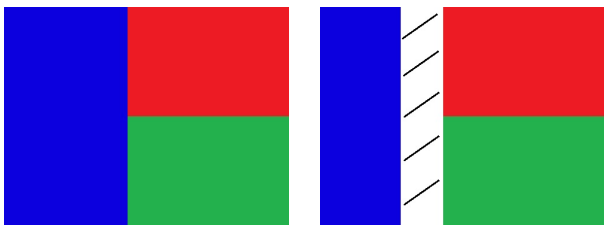

(a) (b)

Figure 3.4.: (a) Original image, (b) image after applying the "blue erosion": one needs to decide how to fill the gap.

Let us first define the erosion for pixels where there are no ambiguities. Thus the *erosion of label i on image I by structuring element B* is given by

$$\varepsilon_i(I; B)(x) = \begin{cases} I(x) & \text{if } I(x) \neq i \\ i & \text{if } \forall p \in B_x, I(p) = i \\ \theta(x, I) & \text{otherwise} \end{cases} \tag{3.8}$$

We will address later definition of $\theta(x, I)$. Sections 2.2 and 2.3 are independent of θ. Although the image is a partition of Ω the proposed framework differs from Ronse and Agnus (2005) and Ronse (2013).

3.2.2. Opening and closing of label i

Once the dilation and erosion have been defined, we can introduce by composition of these two operators the *opening and the closing on I by B of label i* respectively as

$$\gamma_i(I;B) = \delta_i \circ \varepsilon_i = \delta_i\left(\varepsilon_i(I;B);B\right), \tag{3.9}$$
$$\varphi_i(I;B) = \varepsilon_i \circ \delta_i = \varepsilon_i\left(\delta_i(I;B);B\right). \tag{3.10}$$

Let us set a few notations used in the following. If ϕ is an operator, let ϕ^k be $\phi \circ .... \circ \phi$ the iteration of ϕ, k times. Let $\phi_{|A}$ be the restriction of ϕ to the subset A. Let us set $E_i^I = I^{-1}(i)$. To simplify, $\mathbf{1}_{E_i^I}$ will be noted $\mathbf{1}_i^I$.

Due to the fact that operators (3.7) and (3.8) are not adjoint, classical theory of opening and closing by adjunction, see Serra (1988); Heijmans (1994), does not apply to our framework. Hence, we need to explore which properties still hold and in particular, if operators (3.9) and (3.10) are idempotent. Indeed, we have such a property of stability.

Proposition 3.2.1 *Opening and closing of label i are idempotent operators, i.e.,*

$$\gamma_i(I;B) = \gamma_i^2(I;B),$$
$$\varphi_i(I;B) = \varphi_i^2(I;B).$$

Proof Since the binary opening is idempotent, one has $E_i^{\gamma_i(I)} = E_i^{\gamma_i^2(I)}$. Furthermore we have that $E_j^{\gamma_i(I)} \subset E_j^{\gamma_i^2(I)}$, for all $j \neq i$. Since sets $(E_i)_i$ from a partition of the support space, necessarily $E_j^{\gamma_i(I)} = E_j^{\gamma_i^2(I)}$, $\forall j$. Indeed, if all the elements of a partition are extensive, then they all remain stable. Then $\gamma_i = \gamma_i^2$.

Properties cannot directly be transported by duality, as in binary morphology, however the property remains true for the closing. We first show the binary property $\varepsilon\delta\varepsilon = \varepsilon$. The binary erosion and opening can be written as

$$\varepsilon_B(X) = \cup_{B_x \subset X}\{x\}, \quad \text{and} \quad \gamma_B(X) = \cup_{B_x \subset X} B_x.$$

Then $\varepsilon(\gamma(X)) = \cup_{B_x \subset \gamma(X)}\{x\}$. Since $\{B_x \subset \gamma(X)\} = \{B_x \subset \cup_{B_x \subset X} B_x\} = \{B_x \subset X\}$, then $\varepsilon(\gamma(X)) = \varepsilon(X)$. Thus, $\varepsilon\delta\varepsilon = \varepsilon$ and by duality, $\delta\varepsilon\delta = \delta$. Then $E_i^{\delta_i} = E_i^{\delta_i\varepsilon_i\delta_i}$. It can be shown that $E_j^{\delta_i} \subset E_j^{\delta_i\varepsilon_i\delta_i}$ for all $j \neq i$. Using the same reasoning as in the proof for the opening, we have that for all j, $E_j^{\delta_i} = E_j^{\delta_i\varepsilon_i\delta_i}$. In other words, $\delta_i\varepsilon_i\delta_i = \delta_i$. Thus $\varepsilon_i\delta_i\varepsilon_i\delta_i = \varepsilon_i\delta_i$, or equivalently $\varphi_i = \varphi_i^2$.

3.2.3. Composed n-ary filters

We can now try to define label filters from the openings and the closings of label i. In binary morphology, the simplest filters are of the following form: $\gamma \circ \phi$ and $\phi \circ \gamma$. In the n-ary framework, with $n = 2$, they can be rewritten as

$$\gamma_1 \circ \gamma_2 = \phi_2 \circ \phi_1, \quad \text{and} \quad \gamma_2 \circ \gamma_1 = \phi_1 \circ \phi_2.$$

(a) (b) (c)

Figure 3.5.: Opening and closing on a 3-ary image: (a) original image I, (b) opening of red color $\gamma_{red}(I;B)$, (c) closing of black color $\varphi_{black}(I;B)$.

The opening removes peaks smaller than the structuring element and the closing removes holes, which are dual notions in binary morphology. However, peaks and holes are no longer a dual notion in n-ary morphology with $n > 2$. Fig. 3.5 illustrates the difference between openings and closings on a 3-ary image: three colors, black, white and red. The structuring element is a square whose size is half of the width of the red line. Removing the red line using ϕ_{black} requires a structuring element twice bigger than with γ_{red}.

As a good candidate to filter out small object of a labeled image I, independently of the label of the objects, we introduce the operator ψ, named *composed n-ary filter by structuring element* B, defined as

$$\psi(I;B) = \gamma_n(I;B) \circ \gamma_{n-1}(I;B) \circ \cdots \circ \gamma_1(I;B). \tag{3.11}$$

Unfortunately on the contrary to $\gamma \circ \phi$ in binary morphology, ψ is generally not idempotent. Worst, the sequence ψ^k do not necessarily converge. However we still have a stability property for relevant objects. Let us be more precise.

Proposition 3.2.2 *Let Ω be a finite set. Given a structuring element B, the interior with respect to B of the composed n-ary filter $\psi(I;B)$ converges for any image I, i.e.,*

$$\forall i, \varepsilon(E_i^{\psi^k}) \text{ converges .}$$

Proof Since $\varepsilon = \varepsilon \circ \delta\varepsilon$, $\forall i$, $\varepsilon(E_i^{\psi^k}) = \varepsilon(E_i^{\gamma_i \circ \psi^k})$. Furthermore, since $\varepsilon(E_i^{\psi^k}) \subset \varepsilon(E_i^{\gamma_j \circ \psi^k})$, we have that $\forall i, \varepsilon(E_i^{\psi^k}) \subset \varepsilon(E_i^{\psi^{k+1}})$. Since Ω is a finite set, $\varepsilon(E_i^{\psi^k})$ converges.

This property ensures that the variations between ψ^k and ψ^{k+1} do not affect the interior of objects and is only limited to boundaries. Nevertheless, as we shown in section 4, ψ^k is almost always stable after a few iterations.

3.2.4. n-ary geodesic reconstruction

The binary reconstruction can be transposed in the n-ary framework as follows. The proposition of reconstruction is similar to the one proposed in Ronse and Agnus (2005). Given two labeled images R and M, for each label i,

- Perform a binarisation of the reference R and the marker M between i and $\complement i$, which correspond respectively to binary images X_i and Y_i.
- Compute $\gamma^{\text{rec}}(X_i;Y_i)$, that is the binary geodesic reconstruction of the marker in the reference.

Then, the *n-ary geodesic reconstruction of the reference R by the marker M* is given by

$$\gamma^{\text{rec}}(R;M)(x) = \begin{cases} i & \text{if } x \in \gamma^{\text{rec}}(X_i;Y_i) \\ M(x) & \text{if } \forall i,\ x \notin \gamma^{\text{rec}}(X_i;Y_i) \end{cases}$$

Fig.3.6 illustrates the difference between classical geodesic reconstruction and the proposed n-ary reconstruction. For the classical reconstruction, the 3-label image is simply viewed as a grey-scale image.

(a) (b) (c) (d)

Figure 3.6.: Geodesic reconstruction of a 3-ary image: (a) reference image R, (b) marker image M, (c) classical grey-scale reconstruction, (d) n-ary reconstruction $\gamma^{\text{rec}}(R;M)$.

The aim of this definition is to symmetrize labels. In Fig.3.6 (d), the grey object is considered as an object in itself. The proposed reconstruction is a connected operator in the sense of Salembier and Serra (1995)

3.2.5. n-ary morphological gradient

The classical symmetric morphological gradient of a binary or grey-scale image I is defined as the difference between the dilation and the erosion:

$$grad(I)(x) = \delta_B(I)(x) - \varepsilon_B(I)(x),$$

where B is typically the unit ball. This concept can be extended to an $n-$ary image I by the following formulation of the *n-ary morphological gradient*:

$$grad(I)(x) = \left(\mathbf{1}_{\delta_1(I;B)}(x), \mathbf{1}_{\delta_2(I;B)}(x), \cdots, \mathbf{1}_{\delta_n(I;B)}(x)\right). \tag{3.12}$$

where $\mathbf{1}_{\delta_i(I;B)}(x)$ stands for $\mathbf{1}_{E_i^{\delta_i(I;B)}}(x)$ and $grad(I)(x)$ is an element of $\{0,1\}^n$. As the standard gradient, a large part of its information is contained in its norm. Here, $\|grad(I)(x)\|_1 = \sum_i \mathbf{1}_{\delta_i(I)}(x)$ tells us how many different labels lays in the neighbourhood of each point x.

Fig. 3.7 shows the differences between the classical morphological gradient and the n-ary morphological gradient of a labeled image. We note that for the classical color gradient, the involved dilation/erosion are computed using a lexicographic (total) ordering. Both gradients have similar aspect. However, their respective grey levels have different meanings. In the classical framework, the intensity represents a difference of position in the order, whereas in the n-ary framework, the intensity represents the local variety of labels.

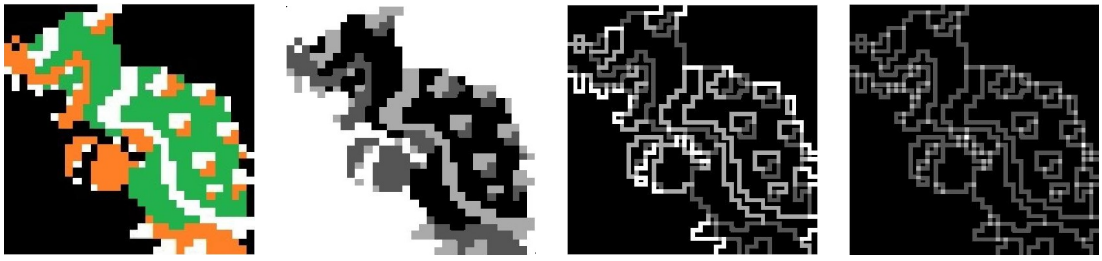

(a) (b) (c) (d)

Figure 3.7.: Computing the gradient of a color image: (a) original n-ary image, (b) projection of a total color order on a grey-scale image, (c) classical morphological gradient based on order represented by (b), (d) L^1 norm of the n-ary morphological gradient $\|grad(I)(x)\|_1$.

3.3. On the choice of an erosion of label i

Before any application, we need to come back to the erosion problem. More precisely, we need to define a consistent rule to fill the space created by the erosion operation.

First of all, we note that the definition in Eq. (3.8) of the erosion ε_i of label i does not indicate how to behave on the following set:

$$A = \{x \mid I(x) = i \text{ and } \exists p \in B_x \text{ such that } I(p) \neq i\}.$$

For points $x \in A$ we have to decide by which label to replace label i and therefore to define ε_i on A, i.e.,

$$\varepsilon_i(I;B)(x) = \begin{cases} I(x) & \text{if } I(x) \neq i \\ i & \text{if } \forall p \in B_x, I(p) = i \\ ? & \text{if } x \in A \end{cases} \tag{3.13}$$

Many alternatives are possible. Two criteria have to be taken into account: (i) the direct coherence in terms of image processing, and (ii) the number of morphological properties verified by the erosion, such as $\varepsilon_i(I;kB) = \varepsilon_i^k(I;B)$ where $kB = \{kx \mid x \in B\}$ (i.e., homothetic of size k). Let us consider in particular the three following rules for $x \in A$:

1. **Fixed-label erosion:** Erosion always fills the gaps with label 1 (or any other fixed label):

$$\varepsilon_i(I;B)(x) = 1. \tag{3.14}$$

2. **Majority-based erosion:** Erosion takes the value of the major label different from i in the structuring element B:

$$\varepsilon_i(I;B)(x) = min(\arg\max_{j \neq i} (\text{Card}\,\{p \in B_x | I(p) = j\})). \tag{3.15}$$

3. **Distance-based erosion:** Erosion replaces label i by the closest label on the support space Ω:

$$\varepsilon_i(I;B)(x) = min(\arg\min_{j \neq i} d_x^j). \tag{3.16}$$

where $d_x^j = \inf\{\|x - p\|_\Omega \mid p \in \Omega,\ I(p) = j\}$.

The majority-based erosion (3.15) and distance-based erosion (3.16) are initially not defined in case of equality. Hence the apparition of the *min*. Obviously, fixed-label erosion (3.14) satisfies $\varepsilon_i(I; kB) = \varepsilon_i^k(I; B)$, but is not coherent in terms of image processing.

(a) (b) (c) (d) (e)

Figure 3.8.: Comparison of erosions of black color: (a) original image I, (b) erosion $\varepsilon_{black}(I; B)$ using majority-based formulation (3.15), (c) erosion $\varepsilon_{black}(I; B)$ using distance-based formulation (3.16). Let l be the width of the red line, the structuring element B is a square whose size is now between l and $2l$. (d) iterated black erosion $\varepsilon_{black}^2(I; B)$ using majority-based formulation, (e) erosion $\varepsilon_{black}(I; 2B)$ using majority-based formulation.

Majority-based erosion and distance-based erosion both look potentially interesting in terms of image processing; however, as shown in the basic example of Fig. 3.8, majority-based erosion (to compare Fig. 3.8(b) to Fig. 3.8(c)) can produce unexpected results. The same example, now Fig. 3.8(d) and Fig. 3.8(e), shows that majority-based erosion do not satisfies $\varepsilon_i(I; kB) = \varepsilon_i^k(I; B)$.

Let us formalize the iterative behavior of the distance-based erosion by the following result on isotropic structuring elements.

Proposition 3.3.1 *Let (Ω, d) be a compact geodesic space and I a n-ary image on Ω. For any $R_1, R_2 > 0$ the distance-based erosion of label i satisfies:*

$$\varepsilon_i(I; B_{R_1+R_2}) = \varepsilon_i(I; B_{R_2}) \circ \varepsilon_i(I; B_{R_1}).$$

where B_R is the open ball of radius R.

Proof For the sake of notation, let $X = E_i^I$ and $X_j = E_j^I$, for $j \neq i$. Let $X' = \{x \in X | d(x, \complement X) \geq R_1\}$, X' is the binary eroded of X. Let $proj_X(a) = \{b | b \in \overline{\complement X}, d(a, b) = d(a, \complement X)\}$ with $\overline{A}$ the closure of A. Let $\mathcal{P} = \{X_j\}$ and for $a \in \Omega$, $I_a^{\mathcal{P}} = \{j | \exists b \in proj_X(a) \cap \overline{X_j}\}$. Let $X_j' = \{a \notin X' | proj_X(a) \cap \overline{X_j} \neq \emptyset\}$ and $\mathcal{P}' = \{X_j'\}$. Using the property of the *min* function, $min(A \cup B) = min(\{min(A), min(B)\})$, it can be shown that $I_a^{\mathcal{P}} = I_a^{\mathcal{P}'}$ for all $a \in X'$ implies proposition 3.3.1.

- Let $j \in I_a^{\mathcal{P}}$. Thus, $\exists b \in proj_X(a) \cap \overline{X_j}$. Let γ be a geodesic joining a and b with $\gamma(0) = b$ and $\gamma(d(a, b)) = a$. Let $c = \gamma(R_1)$. For all $x \in \gamma([0, R_1[)$, $x \notin X'$ and $b \in proj_X(x)$. Hence $x \in X_j'$ and $c \in \overline{X_j'}$. The triangle inequality gives $c \in proj_{X'}(a)$. Indeed, assuming that $c \notin proj_{X'}(a)$ easily leads to $b \notin b \in proj_X(x)$. Thus, $j \in I_a^{\mathcal{P}'}$.
- Let $j \in I_a^{\mathcal{P}'}$. Thus, $\exists c \in proj_{X'}(a) \cap \overline{X_j'}$. By definition of the closure, $\exists c_n \to c, c_n \in X_j'$. $c_n \in X_j' \Rightarrow (\exists d_n \in proj_X(c_n), d_n \in \overline{X_j})$. The compact assumption enables us to consider that $d_n \to d$ (it is at least valid for a sub-sequence). $d \in \overline{X_j}$. By continuity, $d(d, c) = lim(d_n, c_n) = lim(\complement X, c_n) = d(\complement X, c)$. Thus, $d \in proj_X(c)$. Proposition 3.3.2 tells us that $d \in proj_X(a)$. Hence $j \in I_a^{\mathcal{P}}$.

Thus, $I_a^{\mathcal{P}} = I_a^{\mathcal{P}'}$.

Proposition 3.3.2 *Using the notation introduced in the demonstration of proposition 3.3.1:*

$$\forall a \in X', \forall c \in proj_{X'}(a), \forall d \in proj_X(c), d \in proj_X(a).$$

Proof Let $a \in X', c \in proj_{X'}(a), d \in proj_X(c)$.

- Since X' is closed, $\complement X'$ is open and the existence of geodesics implies that $c \notin \overline{\complement X'}$. Hence $c \in X'$ and $d(c, \complement X) \geq R_1$. Since c is a projection on $\complement X'$, we have $c \in \overline{\complement X'}$, $d(c, \complement X) \leq R_1$. Hence $d(c, \complement X) = R_1$ and $d(c, d) = R_1$.
- Let $b \in proj_X(a)$ and γ a geodesic such that $\gamma(0) = a$ and $\gamma(d(a, b)) = b$. Let $c' = \gamma(sup\{t | \gamma(t) \in X'\})$. Since X' is closed, $c' \in X'$. We have $d(c', \complement X) \geq R_1$. Since $c' \in \overline{\complement X'}$, $d(a, c') \geq d(a, \complement X') = d(a, c)$. Thus, $d(a, b) = d(a, c') + d(c', b) \geq d(a, c) + d(c', \overline{\complement X}) = d(a, c) + d(c', \complement X) \geq d(a, c) + R_1 = d(a, c) + d(c, d) \geq d(a, d)$. Hence $d \in proj_X(a)$.

Note that in a compact subset of a vector space, since proposition 3.3.1 is valid for any norm, the property holds for any convex structuring element. Note also that result 3.3.1 is based on a metric definition of the erosion. For vector spaces norms, this erosion is identical to the translation based erosion. Proposition 3.3.1 strongly links the distance-based erosion to a wave-front propagation. Without entering the details, we present the general idea of wave-front propagation and the relation with n-ary morphology. Let the curve $C_0(s) = (x_0(s), y_0(s))$ be the contour of a planar shape S. Let $C(s, \lambda)$ be the curve defined by the following partial differential equation:

$$\begin{cases} \frac{\partial C(s,t)}{\partial t} = \alpha \overrightarrow{N}(s, t) \\ C(s, 0) = C_0(s) \end{cases} \tag{3.17}$$

where N is the normal vector at $C(s, t)$. Authors of Arehart et al. (1993) and Sapiro et al. (1993) point out that $C(s, \lambda)$ corresponds to the dilated or the eroded, depending on the choice $\alpha = 1$ or $\alpha = -1$, of S by a ball of radius λ. The underlying idea is that a dilation (resp. erosion) can be decomposed into a succession of infinitesimal dilations (resp. erosion). The situation is slightly more complicated in the n-ary case, see Fig.3.9. The propagation of each labels $j \neq i$ in the gap left by label i during ε_i is up to a certain point described by a proper wave-front propagation following Eq.(3.17), see Fig.3.9(b). However note that in Fig.3.9(b), if *Front* 1 follows Eq.(3.17), it is not the case of *Front* 2. *Front* 2 is elongated following the trace left by the point p.

According to this discussion, in all what follows we adopt the distance-based erosion.

3.4. Applications to image filtering

In the first case study, depicted in Fig. 3.10 and Fig. 3.11, we consider the behavior and interest of the composed n-ary filter $\psi(I; B)$. The aim is to filter out objects smaller than the structuring element B of the 4-ary labeled image Fig. 3.10(a). Results in Fig. 3.10(b) and (c) are respectively the classical color operator $\gamma_B \circ \varphi_B$ and $\varphi_B \circ \gamma_B$ obtained by using total order of Fig.3.7(b). Then, we compare in Fig. 3.11 four alternative composed n-ary filter $\psi(I; B)$, which correspond to different permutations of the composition of openings

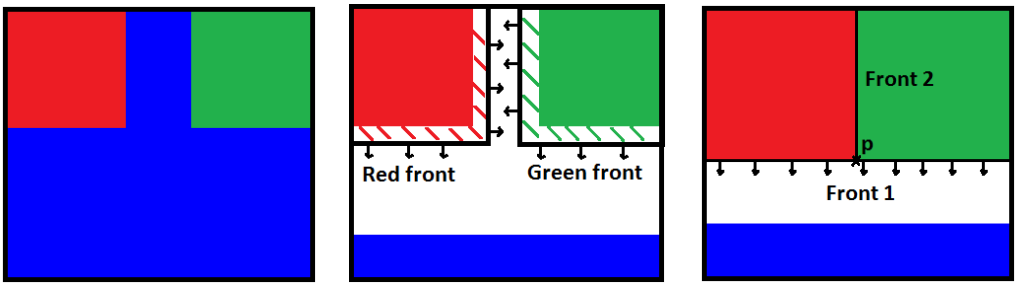

(a) (b) (c)

Figure 3.9.: (a) Original labeled image I, (b) first propagation regime (c) second propagation regime.

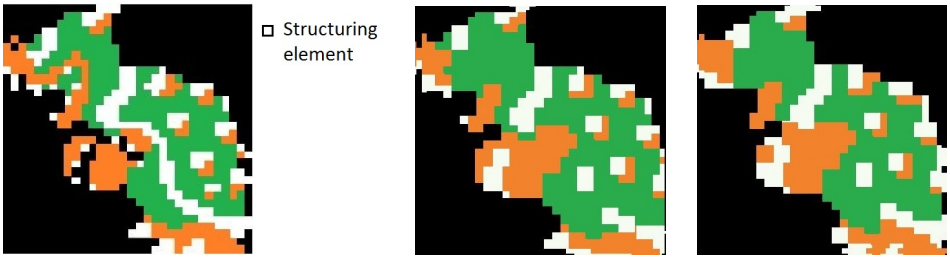

(a) (b) (c)

Figure 3.10.: Morphological filtering of a 4-ary image: (a) original color image I and structuring element B, (b) color operator $\gamma_B \circ \varphi_B$ using a color order, (c) color operator $\varphi_B \circ \gamma_B$ using same order.

of label i. On the one hand, note that in all the cases, the iterated filter converges rather fast to a stable (idempotent) result and that the difference between the first iteration and the final result are rather similar. On the other hand, the different permutations produce different results, however in all the cases, the small objects seems better removed than in the case of the classical color order operators.

The second example, given in Fig. 3.12 attempts to regularize the 3-ary image, by removing small objects without deforming the contours of the remaining objects. More precisely, Fig. 3.12(a) represents the electron microscopy image of a ceramic eutectic, with three different phases after segmentation. The filtering process is composed of two steps: morphological size filter followed by geodesic reconstruction. We compare the results of filtering of the color image according to two pipelines: (i) color total order framework, Fig. 3.12(b), where the filter is an opening by reconstruction composed with a closing by reconstruction; (ii) n-ary framework, Fig. 3.12(c), where the marker is a n-ary filter $\psi(I;B)$ followed by a n-ary geodesic reconstruction. In the case of the color ordering, black and blue are extreme color whereas red is the intermediary color. As we can observe, using the order-based approach all red objects that lay between black and blue objects are not extracted. Both n-ary and color total order frameworks give the same results for the blue grains. This corresponds to the fact that, in the color order, blue is an extreme color whereas red is an intermediary. Therefore the n-ary framework provides a more symmetric processing of all the labels.

The last example is a classification image from the brain. Fig. 3.13 (a) is a result of

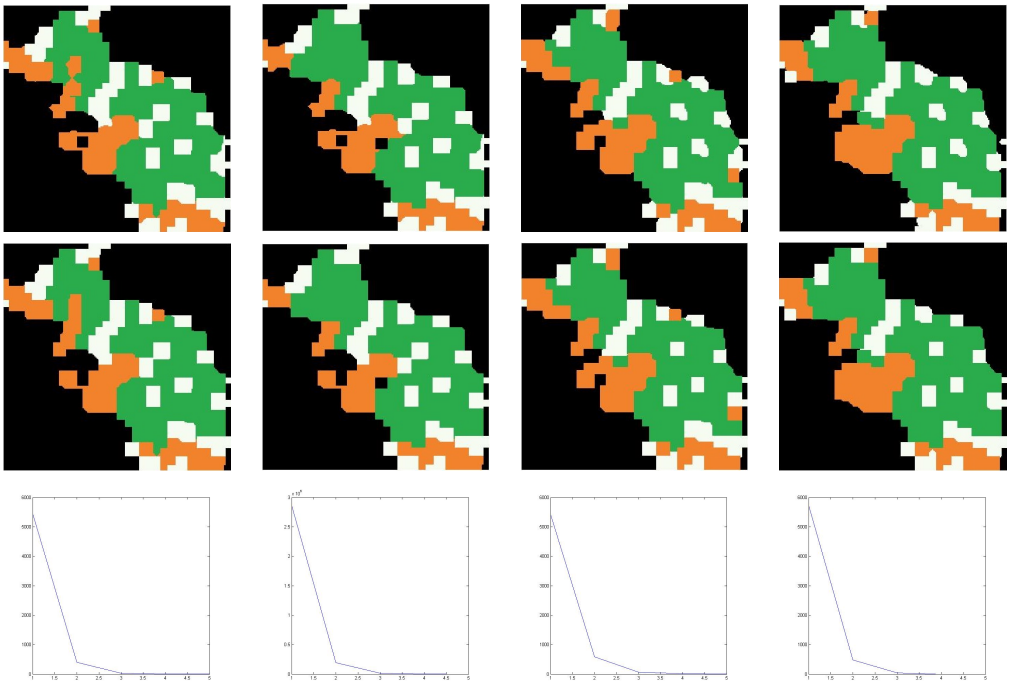

(a) (b) (c) (d)

Figure 3.11.: n-ary processing of Fig. 3.10 (a). first row: composed n-ary filter $\psi(I;B)$, second row: iterated composed n-ary filter $\psi^k(I;B)$ until convergence, third row: convergence speed w.r.t. to k. Column (a) $\psi(I;B) = \gamma_1 \circ \gamma_2 \circ \gamma_3 \circ \gamma_4$, (b) $\psi(I;B) = \gamma_2 \circ \gamma_1 \circ \gamma_3 \circ \gamma_4$, (c) $\psi(I;B) = \gamma_3 \circ \gamma_4 \circ \gamma_1 \circ \gamma_2$, (d) $\psi(I;B) = \gamma_4 \circ \gamma_3 \circ \gamma_2 \circ \gamma_1$.

a classification where the red represents the grey matter, the green represents the white matter and the blue represents the cerebrospinal fluid. The processing is same as for the second example. The miss-classified white matter around the brain and some miss-classified grey matter spots around the cerebrospinal fluid are successfully removed by the n-ary framework, whereas they remain after the classical processing.

3.5. A first extension to a fuzzy case

A natural candidate space to generalize n-ary morphology is the $(n-1)$-simplex, that is to say the space of discrete probability distributions on n elements. Instead of assigning one and only one label per pixel, we assign now a discrete probability distribution representing a local mixture of classes.

In this setting, an image is a function of the form:

$$I : \begin{cases} \Omega \to \{(a_1, \cdots, a_n),\ a_i \geq 0,\ \sum_i a_i = 1\} \\ p \mapsto I(p) \end{cases}$$

The aim of this section is to adapt to the continuous $n-$ary setting, the definitions of the dilation and erosion per label as in (3.7) and (3.8).

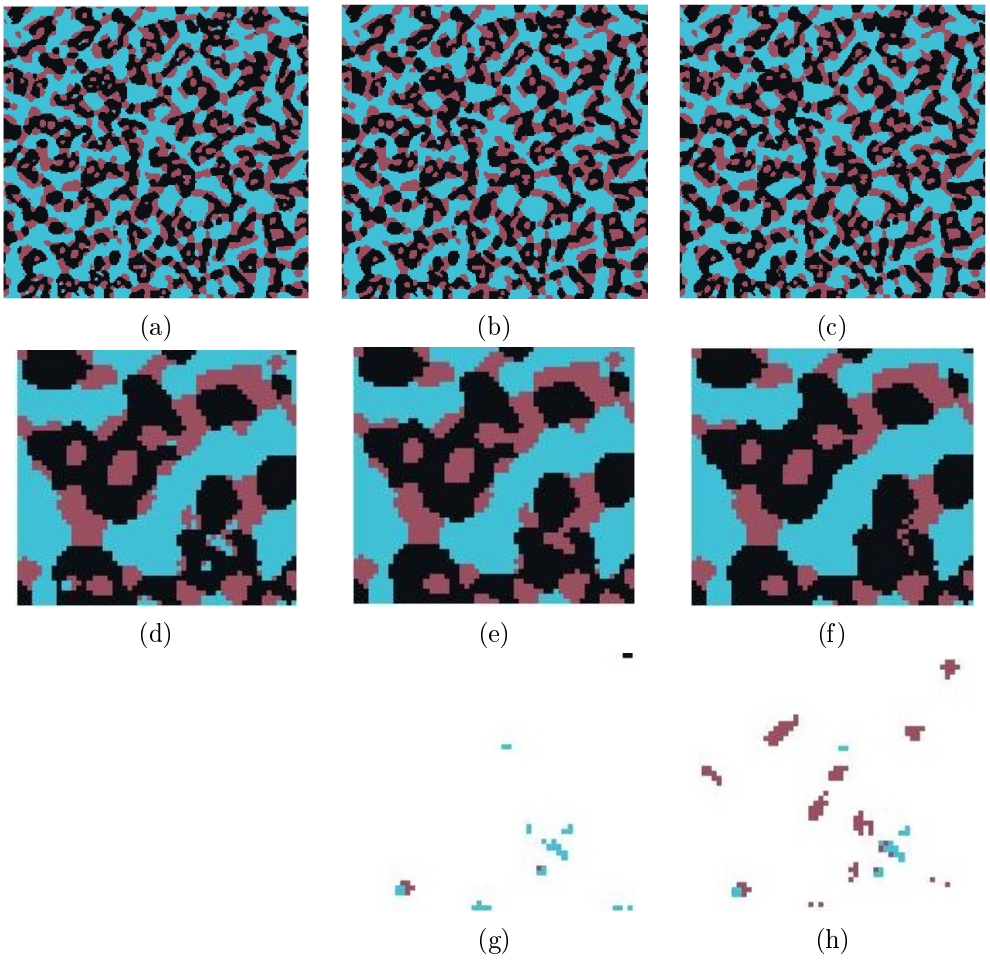

Figure 3.12.: Image size-regularization: (a) original image I, (b) classical order-based filtering, (c) n-ary based filtering, (d) zoom-in (a), (e) zoom-in (b), (f) zoom-in (c) , (g) residue between (d) and (e), (h) residue between (d) and (f). See the text for details.

As we have pointed out in the introduction, the approach developed here has no direct link with the classical grey-scale fuzzy morphology, see for instance Deng and Heijmans (2002); Bloch (2009) as relevant works on the latter topic.

3.5.1. Basic operators and their properties

Similarly to (discrete) n-ary case the dilation is more straightforward than the erosion. Let $I(p)_i$ be the i-th element of $I(p)$, and I_i be the real valued image associated to the i-th component. A natural way to dilate the component i is to perform a grey-scale dilation of the component i, followed by a rescaling of the other components in order to belong to the $(n-1)$-simplex:

$$\delta_i(I;B)(x)_j = \begin{cases} \delta(I_i;B)(x) & \text{if } j = i \\ \frac{1-\delta_i(I;B)(x)_i}{1-I(x)_i}.I(x)_j & \text{if } j \neq i \end{cases} \tag{3.18}$$

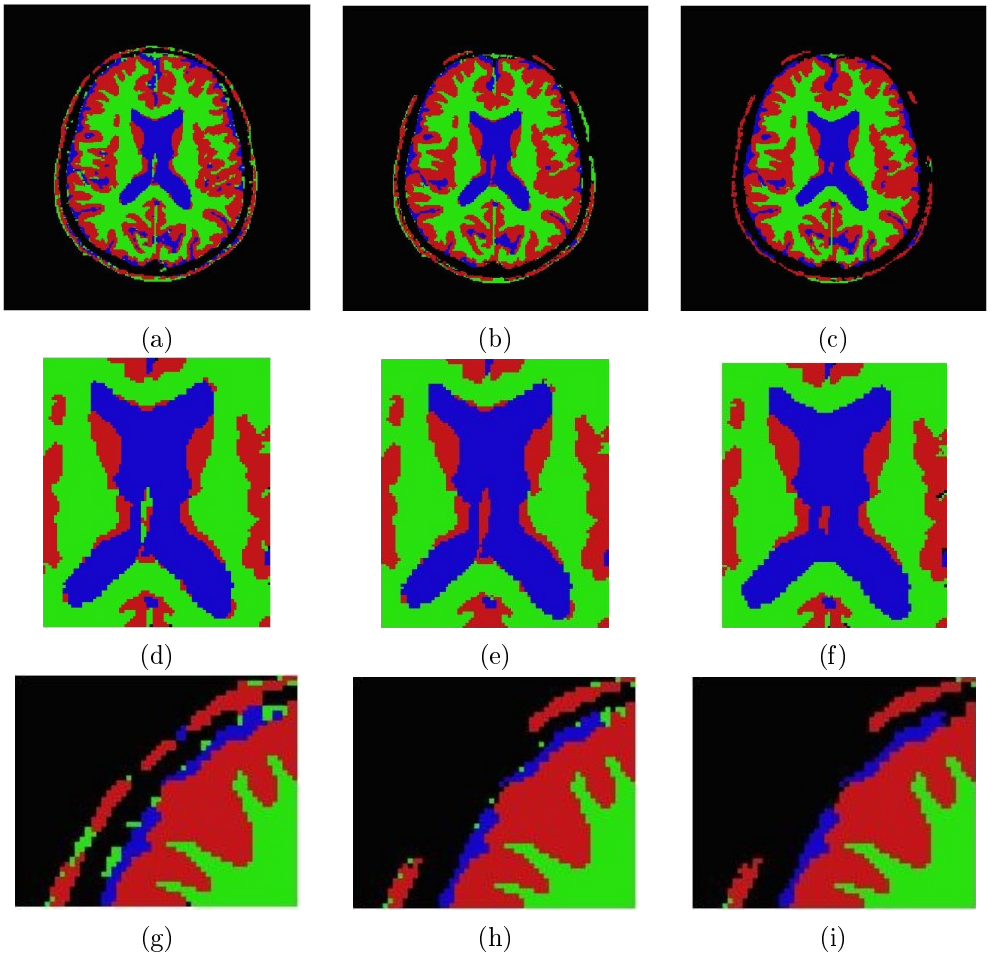

Figure 3.13.: Image size-regularization: (a) original image I, (b) classical order-based filtering, (c) n-ary based filtering, (d) zoom-in (a), (e) zoom-in (b), (f) zoom-in (c), (g) zoom-in (a), (h) zoom-in (b),(i) zoom-in (c). See the text for details.

with

$$\frac{1-\delta_i(I;B)(x)_i}{1-I(x)_i}=1, \quad \text{when} \quad \delta_i(I;B)(x)_i = I(x)_i = 1.$$

One could be tempted to define the erosion the same way, that is to erode the component i and rescaling the other components. In the n-ary case this is equivalent to leaving the gap of Fig.3.4(b) empty. The surface of the label i is reduced, but not replaced by a specific label, which is not entirely satisfying neither consistent with the discrete case. Furthermore, contrary to the dilation the rescaling is impossible when $I(x)_j = 0$ for all $j \neq i$.

The amount to be filled is $I(x)_i - \varepsilon(I_i;B)(x)$, where $\varepsilon(I_i;B)$ stands for the classical erosion of the partial domain defined function I_i. One has to decide with which label or set of labels to complete the distribution. Unlike in section 3.2.2, the stability properties of composed operators of the form $\varphi_i = \varepsilon_i \circ \delta_i$ or $\gamma_i = \delta_i \circ \varepsilon_i$ strongly depends on that choice.

Following the philosophy from Section 3.3, the components used to fill the gaps should

be chosen among components present in the local structuring element B_x. The difficulty is that many of the possible choices do not provide idempotence of φ_i and γ_i.

Consider for instance the following erosion rule. A grey-scale erosion is performed on component i. At x, let $j \neq i$ be the index of the component with the highest value in B_x. When j is not uniquely determined, j is chosen as small as possible. $I(x)_i - \varepsilon(I_i;B)(x)$ is added to the component j. More formally, the erosion can be written as:

$$\begin{cases} \varepsilon_i(I;B)(x)_i = \varepsilon(I_i,B)(x) \\ \\ \varepsilon_i(I;B)(x)_j = I_j(x) + I_i(x) - \varepsilon(I_i,B)(x) \ \text{ for } \ j = min(\arg\min_{k\neq i}(\delta(I_k,B))) \\ \\ \varepsilon_i(I;B)(x)_k = I_k(x) \ \text{ for } \ k \neq i,j \end{cases} \tag{3.19}$$

This rule is an interesting candidate of erosion ε_i, however it does not provide idempotency of φ_i and γ_i. Indeed we did not find a satisfying rule providing the stability of φ_i and γ_i without restricting the generality of the framework. The first restriction consist in working only on the edges of the simplex, that is to say with distribution composed of two components at most, i.e., only two values of the mixture $(a_1,\cdots,a_n)$ can be different from 0. If this set of distribution can not longer deal with "triple points", it is still enough to describe transitions between two objects. The second restriction makes the filling rule depending on the location on the image. Consider the following rule. A grey-scale erosion of the component i is performed. When at x, i is the majority component and when there exists a component $j \neq i$ majority at $p \in B_x$, the value of the component j becomes $1-\varepsilon(I_i,B)(x)$ and all other components are set to 0. If i is not the majority component at x, or if there is no other majority component in B_x, the remaining component are rescaled. Formally, the restricted erosion is given by

$$\begin{cases} \varepsilon_i(I;B)(x)_i = \varepsilon(I_i,B)(x) \\ \\ \text{if } I_i(x) > 0.5 \ \text{ and } \ \max_{k\neq i}(\delta(I_k,B)) > 0.5 \\ \qquad \varepsilon_i(I;B)(x)_j = 1 - \varepsilon(I_i,B)(x) \ \text{ for } \ j = \min(\arg\max_{k\neq i}(\delta(I_k,B))) \\ \qquad \varepsilon_i(I;B)(x)_{k\neq i,j} = 0 \\ else, \\ \qquad \varepsilon_i(I;B)(x)_{k\neq i} = \frac{1-\varepsilon_i(I;B)(x)_i}{1-I(x)_i}.I(x)_k \end{cases} \tag{3.20}$$

As mentioned in Section 3.2.3 and illustrated Fig. 3.5, the opening operator $\gamma_i = \delta_i \circ \varepsilon_i$ tends to be more interesting than the closing in the n-ary framework. The intuitive considerations about the openings and closings are confirmed by the fact that morphological properties are easier to recover for openings than for closings. It was the case for the n-ary filter introduced in Eq.(3.11), and it is now the case for the stability properties.

Proposition 3.5.1 *In the case of the n-ary continuous dilation of component i as in* (3.18) *and the restricted n-ary continuous erosion of component i, we have the following stability of the n-ary continuous opening:*

$$\gamma_i^2(I;B) \ = \ \gamma_i^3(I;B).$$

Proof Let $I_j^{\varepsilon_i} = \varepsilon_i(I;B)_j$. The grey-scale morphology gives $I_i^{\gamma_i} = I_i^{\gamma_i^2}$. Let us first show that for all j, if $I_j^{\gamma_i}(x) > 0.5$ then $I_j^{\gamma_i}(x) = I_j^{\gamma_i^2}(x)$. The result is clear for $j = i$. For $j \neq i$,

since $I_i^{\gamma_i}(x)$ is not the majority component of $I^{\gamma_i}(x)$, the j-th component is simply rescaled after ε_i. Thus $I_j^{\varepsilon_i \circ \gamma_i}(x) \geq I_j^{\gamma_i}(x) > 0$. Since δ_i affects the components different from i only by a rescaling, and that $I_i^{\delta_i \circ \varepsilon_i \circ \gamma_i}(x) = I_i^{\gamma_i}(x)$, we have that $I_j^{\delta_i \circ \varepsilon_i \circ \gamma_i}(x) = I_j^{\gamma_i}(x)$. We will now show that for $I_j^{\gamma_i}(x) \leq 0.5$, $I_j^{\gamma_i^2}(x) = I_j^{\gamma_i^3}(x)$. $I_j^{\varepsilon_i \circ \gamma_i}(x)$ is determined by the majority components of $I^{\gamma_i}(p)$ for $p \in B_x$. Since the majority components are stable between I^{γ_i} and $I^{\gamma_i^2}$, we have that $I_j^{\varepsilon_i \circ \gamma_i}(x) = I_j^{\varepsilon_i \circ \gamma_i^2}(x)$. Finally since δ_i only introduce rescaling on component different from i, $I_j^{\delta_i \circ \varepsilon_i \circ \gamma_i}(x) = I_j^{\delta_i \circ \varepsilon_i \circ \gamma_i^2}(x)$. Hence the result.

3.5.2. Experimental results

Let us consider the example of Fig.3.15(a), which corresponds to a grey-scale image of a triphasic material. Each phase is characterized by a range of grey level. On a semantic level, each phase has an equivalent level. It is then arbitrary to order them according to their grey level. A fuzzy extension of the n-ary framework is a natural framework to process these data. The restriction to a two-classes mixture is particularly adapted to images originally acquired in grey-scale. The two-classes mixture is built following the "fuzzy classification" of grey-scale described in the diagram Fig.3.14 for $n = 3$ phases. Values a, b and c were chosen manually for the experiment but the procedure can easily be automated.

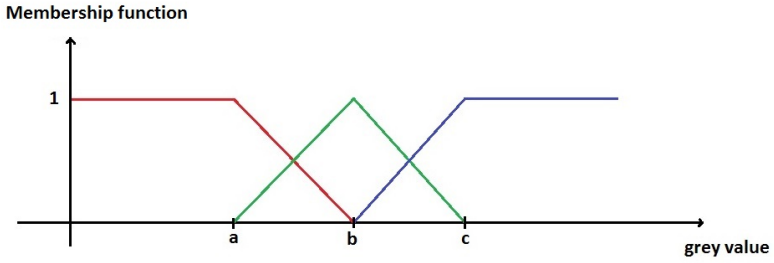

Figure 3.14.: Fuzzy classification of the grey levels

Fig.3.15(b) is a RGB representation of this fuzzy classification. The information contained in Fig.3.15(a) and (b) only differs from the flat zones of Fig.3.14. In Fig.3.15(a) and (b), intermediate grey levels appear at the transition between the darker phase and the lighter phase. These grey level are not representative of a phase but are transition artifacts. Classical grey-scale morphological operators cannot directly address these artifacts. Using the continuous n-ary framework, zones associated to these artifacts easily erased by an n-ary opening of the associated component, see Fig.3.15(c). Note that it is in general not possible to come back to grey-scale image from the fuzzy image.

The initial image of the second example is a color one, see Fig.3.16(a). We start by considering exclusively the Hue component, which is given in Fig.3.16(b). Then, a 3-components fuzzy classification is performed on the Hue channel, see Fig.3.16(c). The Hue channel classification and representation is achieved similarly to the previous example

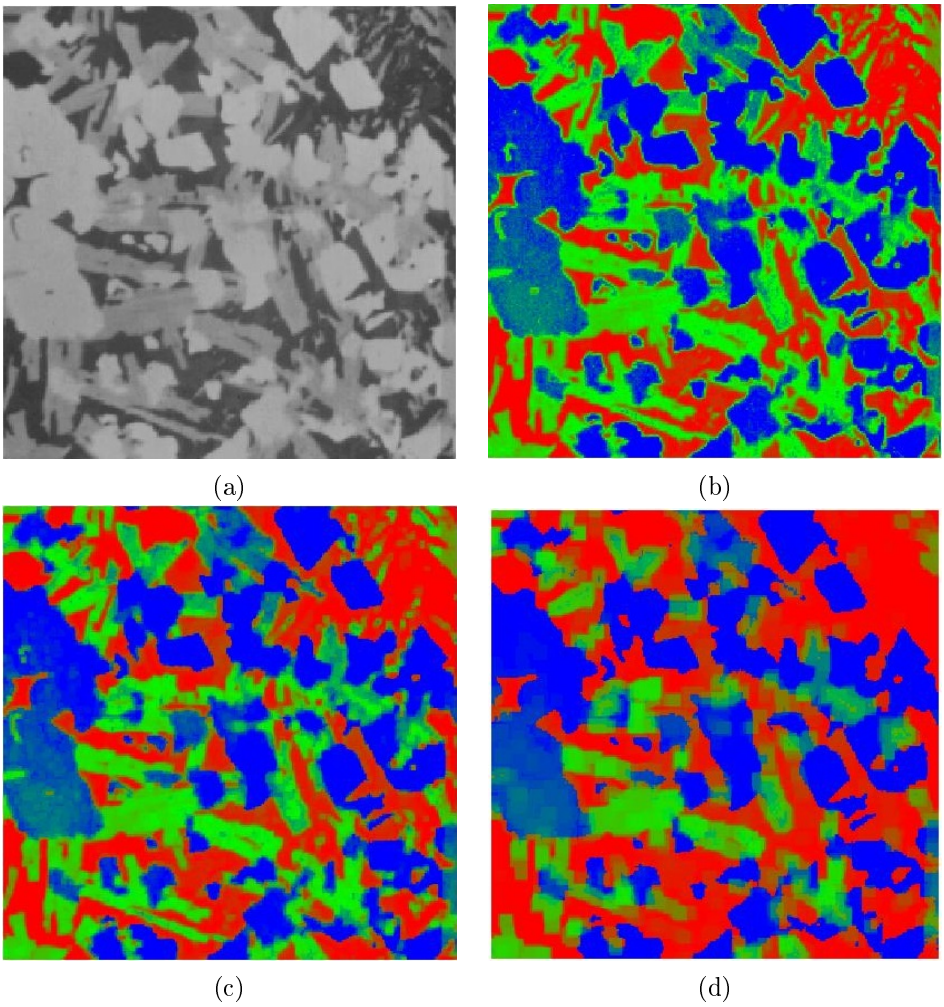

Figure 3.15.: Example of continuous n-ary filtering on grayscale image: (a) triphasic material grey-scale image I, (b) fuzzy classification of grey-scale image represented in the RGB space I_{fuzzy} , (c) fuzzy n-ary opening of green component $\gamma_{green}(I_{fuzzy})$ by a 3×3 square, (d) fuzzy n-ary opening of green component $\gamma_{green}(I_{fuzzy})$ by a 7×7 square.

except that data lay on a circle instead of a line. The fuzzy n-ary filter $\gamma_1 \circ \gamma_2 \circ \gamma_3$ is then applied to Fig.3.16(c) and the result is depicted in Fig.3.16(d).

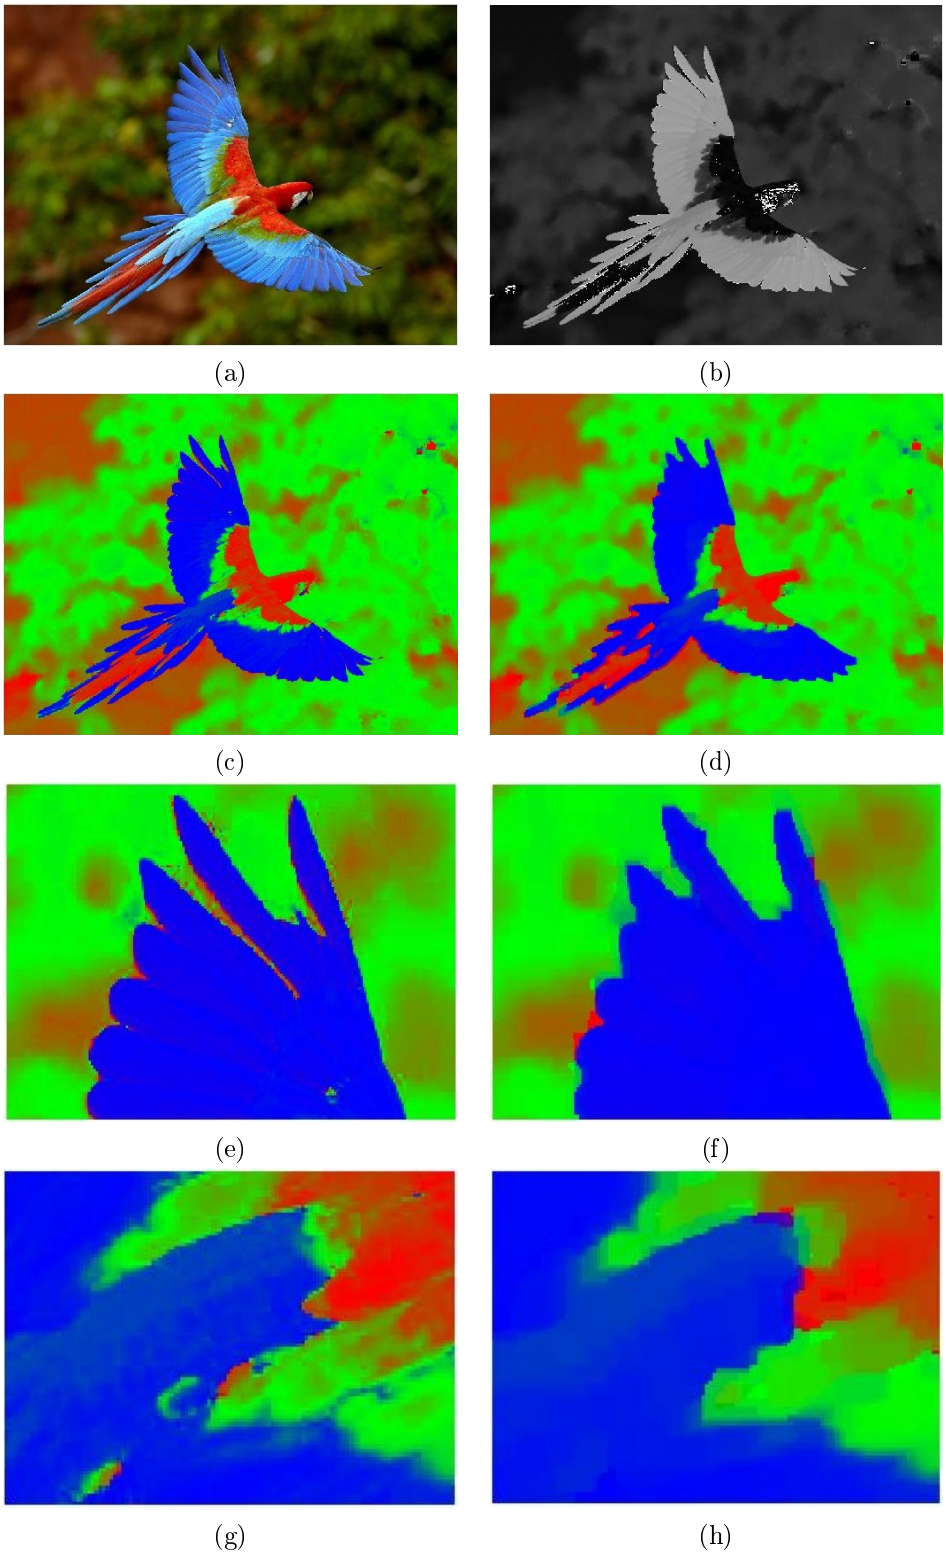

Figure 3.16.: Example of continuous n-ary filtering on hue image: (a) original color image I_{RGB}, (b) Hue channel of I_{RGB}, (c) fuzzy classification from the Hue channel I_{fuzzy} represented in the RGB space, (d) fuzzy n-ary filter of I_{fuzzy}, (e) zoom-in (c), (f) zoom-in (d), (g) zoom-in (c), (h) zoom-in (d).

3.6. Conclusions and perspectives

We proposed here an approach to extend mathematical morphology to images composed of n independent labels (or classes). The approach presents two key particularities: first, the number of elementary operators is increased; second, the absence of an indeterminate class. This absence of indeterminate label transforms the notion of erosion the problem into a problem of gap filling.

In our framework, classical theory of morphological adjunction does not apply and consequently, exploring properties of composed operators requires specific considerations. Hence, we proved that some of the elementary properties of standard morphological operators are preserved, such as the idempotence of openings and closings per label. Despite its quasi experimental validity, the main lost property is the granulometric semigroup. Beyond the mathematical properties, one of the natural consequences of this n-ary framework is the definition of a new reconstruction operator. The main application of the proposed operators is the filtering of image objects according to their size/shape, the presented examples demonstrate the relevance of the n-ary operators.

Fuzzy classification is a natural extension of the discrete n-ary framework to the continuous case, by working on the edges of the $(n-1)$-simplex. Even if theoretical properties are more difficult to recover in continuous cases, the result obtained in practice are promising. Our future research will focus on the improvement of the theoretical framework for the continuous cases, including in particular a relevant definition of geodesic reconstruction, as well as on the study of more practical case-studies and applications.

Part II.

Probability density estimation on Riemannian spaces

4. Theoretical notions and notations of probability and differential geometry

Résumé

Ce court chapitre introduit des notions théoriques et les notations nécéssaires à la compréhension de la deuxième partie de la thèse. Nous avons choisi de ne pas introduire la notion de variété. En effet, mis à part les espaces de rotations, tout les espaces présents dans la suite de l'étude peuvent être décrits en une unique carte: ils sont tous homéomorphes à $\mathbb{R}^n$ pour un certain n. Il n'y a donc pas besoin d'introduire une définition abstraite de plan tangent à partir de classes d'équivalences de courbes. Pour plus de détails sur la théorie de la mesure, voir Tao (2011), sur la géometrie Riemannnienne voir Berger (2003), et sur l'analyse harmonique dans les espaces symétriques, voir Helgason (1993).

Abstract

This short chapter introduces theoretical notions and notations necessary for a good understanding of the second part of the thesis. We chose not to introduce the notion of manifold. Indeed, expect spaces of rotations, every space present in the following study can by described by a unique chart. In other words, they are homeomorphic to $\mathbb{R}^n$. There is then no need to introduce an abstract definition of the tangent space based on equivalence classes. The reader interested on a deeper insight into measure theory is referred to Tao (2011), for more details on Riemannian geometry to Berger (2003) and for a reference on analysis on symmetric spaces to Helgason (1993).

4.1. Basics on measure theory and probability

Let $\Omega = \cup\{\omega\}$ be a set of outcomes. Let $\mathcal{P}(\Omega)$ be the set of subsets of Ω and $\mathcal{A}$ be a subset of $\mathcal{P}(\Omega)$. An element of $\mathcal{P}(A)$ is called an event. $\mathcal{A}$ represents the subset of $\mathcal{P}(\Omega)$ where it is possible to quantify the probability. A probability measure is a function μ_Ω that goes from $\mathcal{A}$ to $\mathbb{R}^+$ which associates a probability to each event. Measure μ_Ω and set $\mathcal{A}$ must verify several properties. First, it is imposed that:

$$\Omega \in \mathcal{A},\ \mu_\Omega(\Omega) = 1.$$

Let A, B be in $\mathcal{P}(\Omega)$ such as $A \cap B = \emptyset$. If it is possible to quantify the probability of A and B, that is to say $A, B \in \mathcal{A}$ then it is possible to quantify the probability of the union:

$$\mu_\Omega(A \cup B) = \mu_\Omega(A) + \mu_\Omega(B).$$

As a consequence $A \cup B \in \mathcal{A}$. The property is extended to any family for disjoint countable sets $(A_i)_{i \in I \subset \mathbb{N}}$:

$$\cup_i A_i \in \mathcal{A},$$

$\mu_\Omega(\cup_i A_i) = \Sigma_i \mu_\Omega(A_i)$.

The last property required of $\mathcal{A}$ is:

$\forall A \in \mathcal{A},\ \complement A \in \mathcal{A},$

where $\complement A$ is the complement of A. Under these assumptions, $\mathcal{A}$ is called a σ-algebra and μ a probability measure on Ω. In a more general way, a measure μ is a probability measure without the constraint $\mu(\Omega) = 1$. In a probability space, the measure is used to quantify the probability of an event. And in a more general way, the measure is used to formalize the notion of area and volume.

We can now define the notion of random variable. Let V be a space with its own σ-algebra. Let $\mathcal{A}_\Omega$ and $\mathcal{A}_V$ be respectively the σ-algebras of Ω and V. A random variable X is a measurable function from Ω to V. A measurable function is a function such that:

$A \in \mathcal{A}_V \Rightarrow X^{-1}(A) \in \mathcal{A}_\Omega.$

The measurable hypothesis implies that the variable X can be seen as a transport of the measure of probability μ from the space Ω to the space V. Indeed, the function μ_X on $\mathcal{A}_V$ defined by:

$\mu_X(A) = \mu_\Omega(X^{-1}(A)),$

is a probability measure on V.

The concept of density can now be introduced. Let E be a set and $\mathcal{A}_E$ a σ-algebra of E. Let μ_1 and μ_2 be two measures on E. A measure is a way of measuring volumes. Consequently, measures play a central role in the notion of integral: the integral of an indicator function of a set A is equal to its volume, that is to say to its measure. The integral of an ordinary function is defined as the limit of a weighted sum of integral of indicator functions. The notion of integral depends then on the choice of a reference measure. We note $\int f d\mu$ the integral of the function f with respect to the measure μ. We say that $f = \frac{d\mu_1}{d\mu_2}$ is the density of the measure μ_1 with respect to the measure μ_2 if we have

$$\forall A \in \mathcal{A}_E,\ \mu_1(A) = \int_A f d\mu_2.$$

Let μ_V be a reference measure on V. For some specific random variable X, the measure μ_X has a density f with respect to the reference measure of V. Then f is called the density of the random variable X. In a finite dimensional vector space there is a unique measure, up to a scaling factor, invariant under translations. The translation invariant measure that normalizes the unit hyper-cube is called the Lebesgue measure. In the case where V is a vector space the reference measure μ_V is often the Lebesgue measure. Fig. 4.1 illustrates a change of reference measure.

Let E and F be sets equipped with σ-algebras, and $f : E \to F$ an application . Any measure μ defined on E can be transformed into a measure $f^*(\mu)$ on F by:

$$f^*(\mu)(A) = \mu(f^{-1}(A)). \tag{4.1}$$

The smallest σ-algebra of $E \times F$ that contains all the products
$U \times V, (U, V) \in \mathcal{A}_E \times \mathcal{A}_F$ is called the product σ-algebra of $E \times F$. A product measure μ on $E \times F$ is a measure such that there exist two measures μ_E and μ_F respectively on E and F with:

$\forall (A, B) \in \mathcal{A}_E \times \mathcal{A}_F, \mu(A \times B) = \mu_E(A)\mu_F(B).$

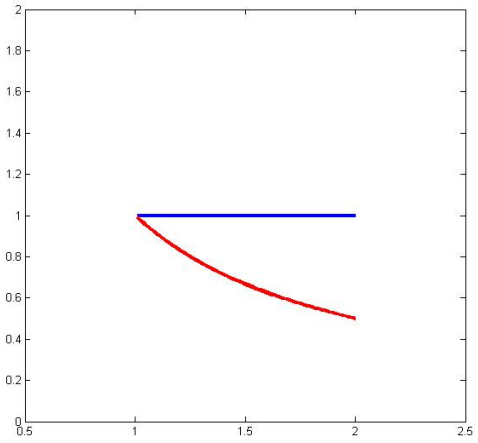

Figure 4.1.: Let X be a random variable valued in $[1, 2]$. The blue curve is the density of μ_X with respect to the Lebesgue measure of $\mathbb{R}$ restricted to $[1, 2]$. The red curve is the density of μ_X with respect to the measure $\mu([a, b]) = log(b) - log(a)$.

4.2. Basics of Riemaniann geometry

Let $\mathcal{M}$ be a topological space, homeomorphic to an open subset of $\mathbb{R}^n$. An homeomorphism is a continuous bijection whose converse is also continuous. Let ϕ be an homeomorphism from an open subset $U_\phi \subset \mathbb{R}^n$ to $\mathcal{M}$. ϕ is referred to as a parametrization of $\mathcal{M}$. A Riemannian metric is a smooth field of scalar product on U_ϕ. In other words, a Riemannian metric associates a positive definite matrix $G(x)$ to each point $x \in U_\phi$ depending smoothly on the point, see Fig. 4.2.

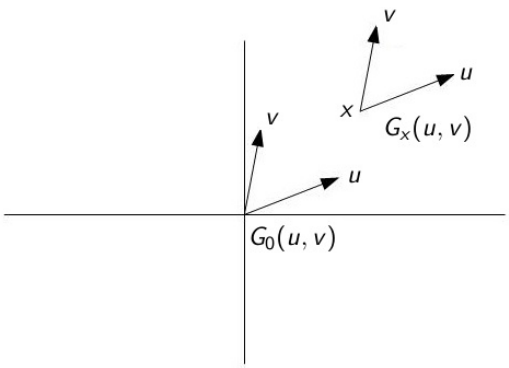

Figure 4.2.: Riemannian space.

A smooth path is a map $\gamma : [a, b] \to \mathcal{M}$ such that $\phi^{-1} \circ \gamma$ is continuously smooth. Let γ be such a path. The Riemannian metric induces a notion of length on smooth paths as follows:

$$L(\gamma) = \int_a^b \sqrt{\langle (\phi^{-1} \circ \gamma)'(t), (\phi^{-1} \circ \gamma)'(t) \rangle_{(\phi^{-1} \circ \gamma)(t)}} dt, \tag{4.2}$$

where $\langle \cdot, \cdot \rangle_{(\phi^{-1} \circ \gamma)(t)}$ is the scalar product attached to the point $(\phi^{-1} \circ \gamma)(t)$. The notion of shortest path between two points induces a distance on $\mathcal{M}$. The distance $d(p, q)$ associated with the Riemannian metric is given by

$$d(p, q) = inf_{\gamma}\{L(\gamma)\}, \tag{4.3}$$

where the infimum is taken over all the smooth paths from p to q. A path realizing this minimum is called a geodesic path. Geodesic paths can be seen as straight segments on $\mathcal{M}$. Geodesics are paths which are locally shortest paths.

Given $p \in \mathcal{M}$ with $\phi(x) = p$, the set of vectors $u \in \mathbb{R}^n$ attached to x is noted $T_p\mathcal{M}$ and is called the tangent space at p. Unlike the case of manifolds where the tangent space is defined as an equivalence class isomorphic to $\mathbb{R}^n$, we simply have here $T_p\mathcal{M} = \mathbb{R}^n$.

It can be shown that given $a < 0$, $0 < b$, vector $u \in \mathbb{R}^n$, there is only one geodesic $\gamma, [a, b] \to \mathcal{M}$ such that $\gamma(0) = p$ with tangent vector $(\phi^{-1} \circ \gamma)'(\gamma^{-1}(p)) = u$ such that $||(\phi^{-1} \circ \gamma)'(t)||_{(\phi^{-1} \circ \gamma)(t)} = ||u||_x$ for all $t \in]a, b[$. It can be then shown that there exists a unique geodesic such that $\gamma(0) = p$, $(\phi^{-1} \circ \gamma)'(\gamma^{-1}(p)) = u$ and that its domain can not be extended. This geodesic is noted $geod_p(u)$. The exponential map, see Fig. 4.3, is then defined by

$$exp_p : \left\{ \begin{array}{ccc} T_p\mathcal{M} & \to & \mathcal{M} \\ u & \mapsto & geod_p\left(\frac{1}{||u||_x} u\right)(||u||_x) \end{array} \right.$$

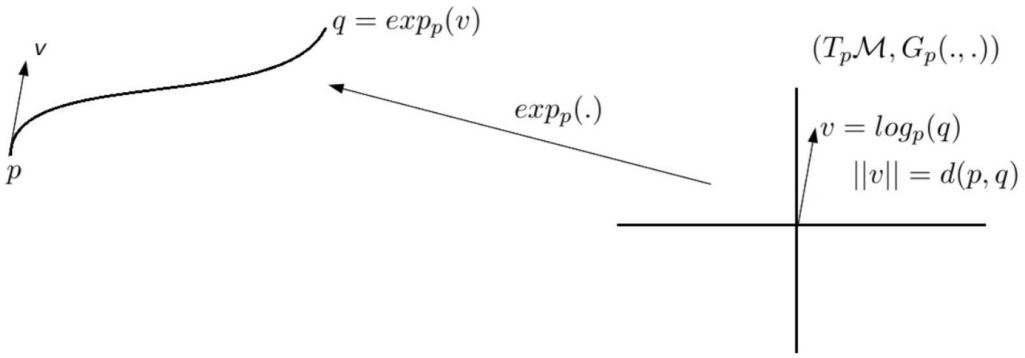

Figure 4.3.: Exponential map.

Note that this exponential has generally no link with the classical exponential application. However, they happen to coincide in specific cases.

The radius of injectivity r_p^{inj} of exp_p is the largest r such that exp_p restricted to the centered ball of radius r is injective. Inside this ball, $log_p(.)$ denotes the inverse of $exp_p(.)$.

Let $\mathcal{S}$ be a subset of $\mathcal{M}$ homeomorphic to an open subset of $\mathbb{R}^m$ with $m \leq n$. The Riemannian metric on $\mathcal{M}$ naturally induces a Riemannian metric on $\mathcal{S}$. The metric on $\mathcal{S}$ is called the induced metric. In the expression of the geodesic distance $d_{\mathcal{S}}(p, q) = inf_{\gamma}\{L(\gamma)\}$, the infimum is then taken over all the paths γ that stay in $\mathcal{S}$.

A Riemannian metric also induces a measure. The matrix of the metric $G(x)$ is symmetric and can be written $G(x) = AA^t$ where A is a $n \times n$ inversible matrix and represents

a local linear change of coordinates that induces the new scalar product $G(x)$. The modification of a unitary volume associated with the linear change of coordinates is given by

$$det(A) = \sqrt{det(G(x))}. \tag{4.4}$$

The density of the Riemannian measure with respect to the Lebesgue measure of U_ϕ is given by $x \in U_\phi \mapsto \sqrt{det(G(x))}$.

In many situations, given a metric, it is impossible to obtain an analytic expression of the distance between two points of $\mathcal{M}$. An isometry of $\mathcal{M}$ is bijection $h : \mathcal{M} \to \mathcal{M}$ such that $d(p_1, p_2) = d(h(p_1), h(p_2))$. Using the parametrization ϕ an isometry is indifferently seen as an application $h : \mathcal{M} \to \mathcal{M}$ or $h : U_\phi \to U_\phi$. The set $ISO(\mathcal{M}, G)$ of isometries is a group: i) the composition of two isometries is an isometry and since an isometry is a bijection, ii) every isometry has an inverse which is also an isometry. Many properties of the geometry induced by the metric can be deduced from the study of $ISO(\mathcal{M}, G)$. A group action from a group A on a set B is an application $act : (A, B) \to B$ compatible with the left-multiplication of A:

$$\forall a_1, a_2 \in A,\ b \in B,\ act(a_2, act(a_1, b)) = act(a_2 a_1, b).$$

The orbit of an element $b \in B$ is defined as:

$$Orb(b_0) = \{b \in B\ |\exists a \in A,\ act(a, b_0) = b\}.$$

An action is said transitive when $Orb(b) = B$. In other words, from a point b in B, it is possible to reach every point of B through the action act. Note that if $Orb(b) = B$ for one b in B, it is also true for all the points in B.

Group of isometries $ISO(\mathcal{M}, G)$ naturally acts on $\mathcal{M}$:

$$act(iso, p) = iso(p),\ iso \in ISO,\ p \in \mathcal{M}.$$

Riemannian manifold $\mathcal{M}$ is said to be homogeneous if the action of $ISO(\mathcal{M}, G)$ is transitive, i.e., given two points of $p_1, p_2 \in \mathcal{M}$, there is always an isometry iso such that $p_2 = iso(p_1)$. In a homogeneous space, all the points are similar in term of geometry. Vector spaces and spheres are examples of homogeneous spaces. However an open disc is not a homogeneous space: if p and q are two points such that $d(0, p)$ is different from $d(0, q)$, there is no isometry iso_{disc} of the disc such that $q = iso_{disc}(p)$.

Let p be a point in $\mathcal{M}$. The group $\mathcal{K}_p$ of isotropy of p is the subgroup of $ISO(\mathcal{M}, G)$ that fixes p, i.e.,

$$\mathcal{K}_p = \{k \in ISO(\mathcal{M}, G) | k(p) = p\}.$$

$\mathcal{M}$ is homogeneous if and only if $\mathcal{M}$ is in bijection with the quotient space $ISO(\mathcal{M}, G)/\mathcal{K}_p$.

A space $\mathcal{M}$ is said to be symmetric if for all $p \in \mathcal{M}$ there exists an element $k \in \mathcal{K}_p$ such that:

$$dk = -id,$$

where dk is the differential of k. In other words, reversing the geodesics starting from a point p is an isometry: central symmetries are isometries. It can be shown symmetric spaces are homogeneous.

Space $\mathcal{M}$ is said to be isotropic if for all $p \in \mathcal{M}$ and for all couples of vectors $u, v \in \mathbb{R}^n$, there is an element k of the group of isotropy of p such that:

$$dk(u) = v.$$

At each point p, an isotropic space $\mathcal{M}$ "looks the same" in every direction. Very few spaces are isotropic. However, several spaces like torus are locally isotropic. It can be shown that isotropic spaces are symmetric and therefore homogeneous.

4.3. Convolution on Riemannian spaces

We introduce now some basic ingredients of analysis on Riemannian spaces. The notion of convolution $f * g$ of two functions $\mathcal{M} \to \mathbb{R}$ is highly related to the notion of homogeneity. In $\mathbb{R}^n$ the convolution is defined as follows:

$$(f *_{\mathbb{R}^n} g)(p) = \int_{\mathbb{R}^n} f(x)g(p-x)dx, \quad p \in \mathbb{R}^n. \tag{4.5}$$

Thus, the convolution of a function f by a kernel g consists in the integral of the translated kernel in each point x of the support space, weighted by $f(x)$. The group law $+$ of $\mathbb{R}^n$ is an isometry that enables to transport the kernel in the whole space.

If the space is not homogeneous, the transport of the kernel might be impossible without modification of its shape.

Examples. Let g_0 be the triangular kernel centered in 0 defined on the open unit disc of $\mathbb{R}$ by $g(x) = 1 - |x|$, $x \in [-1, 1]$. Kernel g cannot be naturally transported at any other point of the disc. On the other hand if g_1 is defined on the unit circle of $\mathbb{R}^2$ by $g_1(e^{i\pi\theta}) = 1 - |\theta|$, $\theta \in [-1, 1]$ can easily transported to any point $p = e^{i\pi\theta_0}$ of the circle: $g_p(e^{i\pi\theta}) = g_1(e^{i\pi(\theta-\theta_0)})$.

In an homogeneous space, it is possible to transport a kernel from a reference point to any other point by an isometry. However, there are often several isometries that send a point p_{ref} on a point q. If $iso \in ISO(\mathcal{M}, G)$ is such that $iso(p_{ref}) = q$, then for any $k \in \mathcal{K}_{p_{ref}}$, $iso(k(p_{ref})) = q$. Since there are several potential candidates to transport the kernel from p_{ref} to q, the convolution kernel g is required to be invariant under the isotropy group of the reference point. Then the transport of the kernel from p_{ref} to q is independent of the choice of the isometry. Let $iso_{q,p_{ref}}$ be any isometry that sends q on p_{ref} such that $g_q(p) = g_{p_{ref}}(iso_{q,p_{ref}}(p))$. Thus in an homogeneous space $\mathcal{M}$, the convolution of a function $f \in L^2$ and a $\mathcal{K}_{p_{ref}}$-invariant kernel $g_{p_{ref}} \in L^2$ is naturally defined by:

$$(f *_{\mathcal{M}} g)(p) = \int_{q \in \mathcal{M}} f(q)g_q(p)d\mu, \tag{4.6}$$

where μ is the Riemannian measure. Note that the convolutions $*_{\mathcal{M}}$, and the standard convolution on $\mathbb{R}^n$ slightly differ. Indeed, on $\mathbb{R}^n$ the isometry used to send p on q is always the translation of vector $q - p$ so that the classical convolution on $\mathbb{R}^n$ does not require an isotropic kernel.

In $\mathbb{R}^n$ the Laplacian of a function f is defined as follows:

$$\Delta_{\mathbb{R}^n} f = div(gradf) = \sum \partial_i^2 f.$$

It is possible to adapt the definition of the gradient and of the divergence operators on $(\mathcal{M}, G)$. The expression of the so-called Laplace-Beltrami operator becomes:

$$\Delta_{\mathcal{M}} f(x) = \frac{1}{\sqrt{det(G(x))}} \sum_{i,j} \frac{\partial}{\partial x_i}(\sqrt{det(G(x))} g^{i,j} \frac{\partial f}{\partial x_j}),$$

where $g^{i,j}$ are the coefficient of $G(x)^{-1}$.

In $\mathbb{R}^n$, the functions involved in the Fourier transform are eigenfunctions of the Laplacian operator:

$$\Delta_{\mathbb{R}^n} e^{i\langle\omega,.\rangle} = ||\omega||^2 e^{i\langle\omega,.\rangle}, \omega, x \in \mathbb{R}^n.$$

The generalization of the Fourier transform on symmetric spaces is called the Fourier-Helgason transform. As trigonometric exponentials, the Helgason functions $e_\omega(x)$ are eigenfunctions of the Laplace-Beltrami operator. The theoretical construction of Helgason functions $e_\omega(x)$ (generalization of plane waves) can be found in Helgason (2006). Explicit expression of these functions are provided when necessary in next chapters. The transform and the inverse transform are of the form:

$$\mathcal{H}(f)(\omega) = \int_x f(x)\overline{e_\omega(x)} d\mu_1(x), \tag{4.7}$$

$$f(x) = \int_\omega \mathcal{H}(f)(\omega) e_\omega(x) d\mu_2(\omega). \tag{4.8}$$

One of the major interests of this transform is that it preserves properties of the Euclidean Fourier transform regarding convolution. Let f be a real function and g be a $\mathcal{K}_{p_{ref}}$-invariant kernel. We have for instance the following property

$$\mathcal{H}(f *_{\mathcal{M}} g) = \mathcal{H}(f) \times \mathcal{H}(g).$$

5. Image histograms

Résumé

L'histogramme est l'un des outils les plus simples et les plus importants du traitement d'image. Il est notament utile dans des problèmes de rehaussement de contraste, de segmentation, de recherche d'image par contenu, etc. La méthode classique de construction de l'histogramme d'une image consiste à découper l'espace des valeurs en cases régulières et à compter le nombre de pixels tombant dans chaque case. La densité obtenue présente en générale d'importantes discontinuités. On préfère donc parfois utiliser d'autres techniques d'estimation de densité tel que la méthode des noyaux. La section 5.2 est une discussion générale sur la notion d'histogramme d'image et d'estimation de densité. La section 5.3 est un travail sur les histogrammes d'image couleurs dans le cadre des métriques perceptuelles. Cette section est principalement extraite de Chevallier et al. (2015b).

5.1. Introduction

The histogram computation of a scalar image consists in counting the number of pixels at each different intensity value. This extremely simple approximation to the univariate density distribution of image intensities is one of the most important image processing tools to address problems such as contrast enhancement (by histogram linear stretching or using advanced approaches, see Sapiro et al. (1993)); image segmentation (by 1D clustering); texture processing, see Portilla and Simoncelli (2000); image retrieval, see Gong et al. (1996), etc. The standard way of computing a histogram is to cut the value space into regular bins and to count the number of pixels that fall into each bin. However the obtained histogram presents important discontinuities. One thus prefers sometimes to use kernel methods or other density estimation techniques. This chapter starts by a general discussion on image histograms, in particular Section 5.2 formalizes the notion in terms of density estimation. Section 5.3 presents a work on perceptual color histograms extracted from conference paper Chevallier et al. (2015b).

5.2. Image histogram and density estimation

In this section we set the theoretical links and notation between image histograms and density estimation. Let us consider an image I as the map:

$$I : \left\{ \begin{array}{ccc} \Omega & \rightarrow & V \\ p & \mapsto & I(p) \end{array} \right.$$

We have for instance $V = \mathbb{R}$ for grey-scale images or $V = \mathbb{R}^n$ for multispectral images. Domain Ω is the support space of pixels/voxels, typically a subset of $\mathbb{R}^2$ or $\mathbb{R}^3$ such as a rectangle or a parallelepiped.

We endow Ω with a measure μ_Ω that reflects its geometry: since the spatial distances are usually evaluated using Euclidean distance, this measure is typically the Lebesgue measure. Indeed the prior importance of an area is often proportional to its Lebesgue measure. Measure μ_Ω is generally finite, that is to say $\mu_\Omega(\Omega) < \infty$.

The application I transports the measure μ_Ω on V. The image measure on V is noted $I^*(\mu_\Omega)$, see Eq. (4.1). The measure $I^*(\mu_\Omega)$ represents the distribution of values. It often carries a useful synthetic information of the image I. Thus, one might be interested in the study of this distribution. We have access to a finite number of evaluations of I in points $p \in \Omega$ called pixels, from which we try to deduce information on $I^*(\mu_\Omega)$. The value space V is usually endowed with a reference measure μ_V. The standard strategy is to build a density f with respect to μ_V from the set of observations $\{I(p), p \text{ a pixel }\}$ such that the associated measure approximates $I^*(\mu_\Omega)$. Note that in many situations $I^*(\mu_\Omega)$ does not have a density with respect to μ_V. This is for instance the case for color images $I : [0,1]^2 \to \mathbb{R}^3$ when μ_Ω and μ_V are Lebesgue measures. Then the density f does not represent exactly the measure $I^*(\mu_\Omega)$.

We assume that pixels p are uniformly distributed with respect to μ_Ω. Roughly, for $A \subset \Omega$,

$$\frac{1}{N} \sum_p \mathbf{1}_A(p) \approx \mu_\Omega(A),$$

where N is the number of pixel and $\mathbf{1}_A$ the indicator function of A. Then, if μ_Ω is normalized, the set of pixels can be considered as a set of N independent realizations of a random variable of law μ_Ω. It follows directly that the set $\{I(p), p \text{ a pixel }\}$ can be considered as a set of N independent realizations of a random variable of law $I^*(\mu_\Omega)$. The study of $I^*(\mu_\Omega)$ is thus often interpreted as a problem of probability density estimation. This is for instance the case of color histograms developed in Section 5.3. However, note that this way of addressing the problem do not take advantage of the specificity of the situation.

Let X be a random variable from (Ω, μ_Ω) to (V, μ_V). Assume known a set of independent realizations $X(\omega_i)$, $\{\omega_i \in \Omega\}$. The estimation of the density of $X^*(\mu_\Omega)$ with respect to μ_V is a very difficult problem since no assumption is made on $X(\omega \notin \{\omega_i\})$. The amount of information contained in $\{X(\omega_i)\}$ about $X^*(\mu_\Omega)$ is extremely weak, on might even say null. Thus the estimated density cannot be determined with confidence, but can only be "guessed". When building an histogram, the choice of the bin size is an essential parameter and there is no precise way of determining if two peaks should be gathered in the same bin.

However, the situation in image processing is different. Unlike the standard case of random variables, the space Ω often has a metric structure and the function I has regularity properties with respect to this metric. Thus, knowing the values $\{I(p_i)\}$ brings information on $I(p \notin \{p_i\})$. The set $\{I(p_i)\}$ carries then a significant amount of information about $I^*(\mu_\Omega)$. Assume for instance a rule of interpolation of $I(p \notin \{p_i\})$ from $\{I(p_i)\}$. The density f of $I^*(\mu_\Omega)$ with respect to μ_V can then be exactly computed using formulas such as the coarea formula, see Morgan (2009).

We first give an almost rigorous formula for $I : [0,1]^2 \mapsto \mathbb{R}$, where spaces are endowed with Lebesgue measures. When possible let $\gamma_h(t)$ be a unit speed parametrization of

$I^{-1}(h)$. Then we have

$$f(h) = \frac{dI^*(Leb_{[0,1]^2})}{dLeb_{\mathbb{R}}}(h) = \int \frac{1}{||\nabla I(\gamma_h(t))||_2} dt.$$

For a function $\phi, \mathbb{R}^n \mapsto \mathbb{R}^k$, the k-dimensional Jacobian determinant $J_k(x)$ is defined by the maximal volume of the image of a unit k-dimensional cube by the differential of ϕ. The general formula for $I, [0,1]^n \mapsto \mathbb{R}^k$ with $k < n$ is,

$$f(h) = \frac{dI^*(Leb_{[0,1]^n})}{dLeb_{\mathbb{R}^k}}(h) = \int_{I^{-1}(h)} \frac{1}{J_k(x)} d\mathcal{H}^{n-k}(I^{-1}(h)),$$

where $\mathcal{H}^{n-k}$ refers to the $(n-k)$-Hausdorff measure, see Morgan (2009).

Note that an interpolation of $I : [0,1]^2 \mapsto \mathbb{R}$ by pieces of linear functions leads to an elementary computation of the density estimator f. The previous problem of the size of bins no longer exists, the shape of the computed density f is intrinsic to the interpolated image. While they did not follow exactly the same line of thoughts, Authors of Duffy et al. (2013) ended up at the same conclusions. Their work still remains quite undisclosed in the image processing community. Fig. 5.1 presents histograms of a 16 bits gray-scale image computed using different methods.

5.3. Perceptual color histograms

One sometimes considers the color space $\mathcal{C}$ as a part of a three dimensional Euclidean space $\mathbb{R}^3$. Under this assumption, the histogram of a color image can be built in the same way as for gray-scale images. However, the distances induced on colors by the human perceptual system cannot be represented by a Euclidean space structure. Observation showed that the perceptual relation between colors is better represented in the framework of Riemannian manifolds.

The local metrics of the Riemannian structure are experimentally measured by a set of ellipses, such as the MacAdam ellipses introduced in Macadam (1942), BFD-P by Luo and Rigg (1986) and RIT-DuPont by Berns et al. (1991). This Riemannian structure makes the construction of the histogram difficult. On the one hand, except rare situations, there are no regular tilings of the space. On the other hand, kernel methods have been generalized to Riemannian manifolds in Pelletier (2005), but requires the knowledge of the geodesics.

We propose here an approach that takes into account the Riemannian structure of $\mathcal{C}$ while keeping the computation in the Euclidean framework. Thus we propose a way of building histograms for images $I : [0,1]^2 \to \mathcal{C}$ that respects the perceptual distances better than histograms built in Euclidean spaces, and without increasing the computation time.

5.3.1. Perceptual metric on colors

Already Riemann used colour as an illustration of the applicability of his geometry in Riemann (1867), and concrete examples of such colour geometries were developed by Helmholtz (1891), Schrödinger (1920) and Stiles (1946). A scalar product is entirely determined by its unit ball. Expressed in vector coordinates, the associated unit ball takes the form of an ellipse in two dimensions or of an ellipsoid in three dimensions. Thus, the Riemannian metric is given by a field of ellipses or ellipsoids.

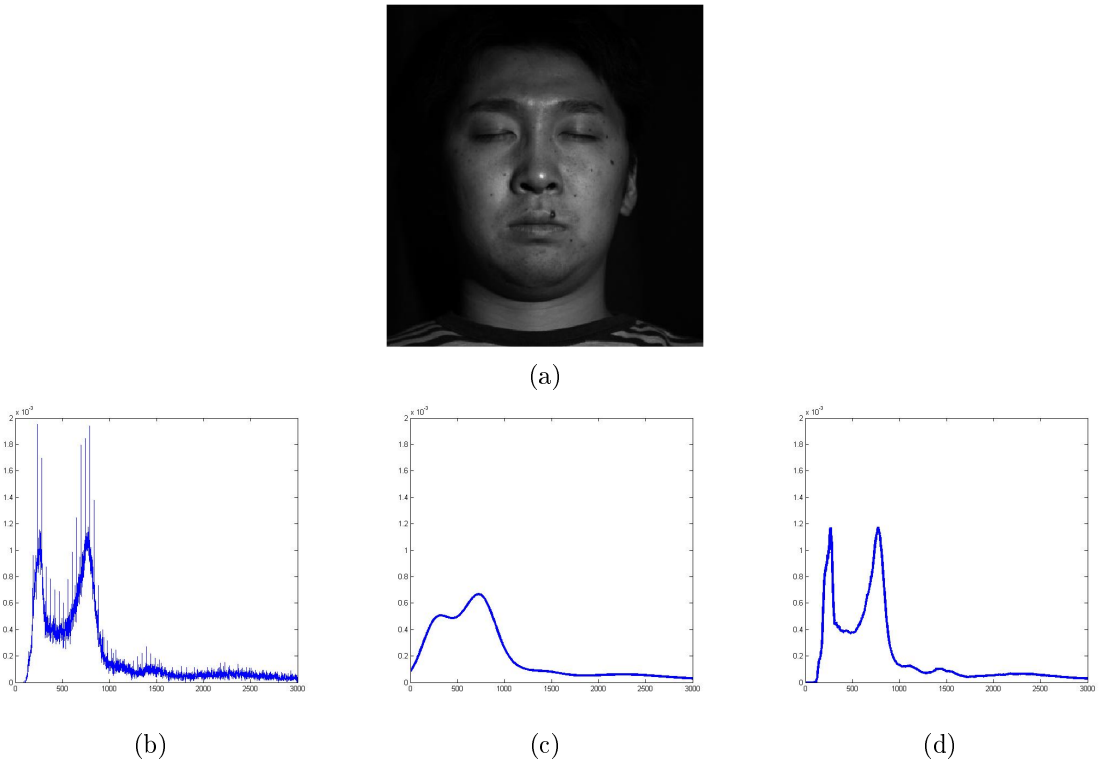

(a)

(b) (c) (d)

Figure 5.1.: Image (a) is a 16 bits gray-scale image. Graph (b) represents the number of pixels associated with each of the 2^{16} levels. This is often the quantity used in image processing. Graph (c) is a kernel density estimate of the distribution of the gray values. The kernel bandwidth is determined by a cross-validation technique. Graph (d) is the density computed using the coarea formula. The computation is achieved using an interpolation of the image by pieces of linear functions.

5.3.1.1. Ellipses, local metric

The first experimental determination of the field of ellipses describing the Riemannian metric of the colour space was performed by Macadam (1942). The experiment consisted of about 25 000 colour matches with one observer, and the ellipses were derived from the covariance matrices of the repeated observations. Later, it has become common practice to denote ellipses obtained in this manner as JND (just noticeable difference) ellipses or ellipsoids.

Later, another type of experiment has become more commonplace. Pairs of colours that are barely perceptually different, are presented to the observer, who is given the task to estimate the magnitude of the perceptual distance using a set of standard pairs. Ellipses, ellipsoids and metrics obtained in this way are normally denoted supra-threshold ellipses. Examples of supra-threshold color difference based data include BFD-P, see Luo and Rigg (1986) and RIT-DuPont, see Berns et al. (1991).

5.3.1.2. Global model

Data sets of measurements provide information on distances through the local metric or through distances between specific colors. A global model provides an analytic expression of the distance between two arbitrary colors. The closest the proposed expression is to the Riemannian perceptual distances, the better the model is. The more conventional procedure for going from a tristimulus space to a space closer linked to a perceptual homogeneous space typically includes the following steps.

- First, apply a linear transform in the tristimulus space such that the base gets close to the cone fundamentals of the retina.
- Secondly, perform a non-linear compression of the coordinates (e.g., logarithmic or cubic root) in order to mimic the non-linear response of the human visual system.
- Finally, perform a linear transformation of the resulting coordinates in order to correspond better to the perceptual attributes of color.

Typically, the first coordinate is a weighted sum of the coordinates and represent a lightness correlate, whereas the two other coordinates are weighted differences, and represent color opponent channels such as, e.g., red–green and blue–yellow.

In order to identify the different parameters of the various transforms, different optimisation criteria are used. In the CIELAB colour space, see Robertson (1977), the parameters were optimised so that the lightness should correspond to perceived lightness, and that the Euclidean metric in the resulting space should correspond to perceptual colour differences. For the IPT colour space, see F Ebner (1998), the parameters were optimised in order to achieve a constant perceived hue along straight radial lines in cylindrical coordinates. It is furthermore reasonably well established that in such perceptual spaces, the Euclidean metric is not the one best corresponding to the perceived colour differences, and other models have been proposed, see, e.g., Luo et al. (2001) and Farup (2014). In the hyperbolic models proposed in Farup (2014), histograms can be computed using adapted kernels, see next Chapter.

5.3.2. Kernel density estimation

There are various ways of addressing the problem of probability density estimation. In the Euclidean context the most popular techniques are mainly the histograms, the kernels, and the characteristic function density estimator. The characteristic function density estimator consists in the estimation of the Fourier transform or series of the density. In what follows, we chose to focus on the kernel method. Recall that the kernel method in the Euclidean case has the following expression for the density estimator:

$$f(x) = \frac{1}{k} \sum_{p_i \in \{pixels\}} \frac{1}{\lambda^n} K\left(\frac{||x - I(p_i)||}{\lambda}\right), \tag{5.1}$$

where λ is a scaling parameter, n the dimension of the space, k the number of pixels, and $K : \mathbb{R}_+ \to \mathbb{R}_+$ a map which obeys the following properties:
- $\int_{\mathbb{R}^n} K(||x||)dx = 1$,
- $\int_{\mathbb{R}^n} xK(||x||)dx = 0$,
- $sup(K(x)) = K(0)$.

In the present work, we assume a supplementary condition of bounded support $K(x > 1) = 0$.

5.3.3. Local Euclidean approximations

In general Riemannian manifolds, computing the distance between two arbitrary points given the metric field is a difficult problem. Indeed, finding the distance is a minimization problem over a set of paths. However, for two close points, the local metric provides a satisfying approximation of the Riemannian distance. A probability density measures the ratio between the probability of an infinitesimal volume element and its volume. It is thus a local notion. The central idea of our contribution is to take advantage of the fact that histograms mainly involves local phenomena. Since in a Riemannian manifold the computation of an histogram does not involve computation of long geodesics, the need of a global model that provides distances between every pairs of colors is of lower importance than in most applications.

Fig. 5.2(a) shows a set of ellipses in the projective ab plane of the CIELAB color space. Let us assume that these ellipses represent the local perceptual metric. Let c be a point where the metric has been measured through the ellipse $\mathcal{E}_c$. In a neighborhood of c, computing distances using the metric measured at c is a better approximation of the perceptual distance than using the canonical Euclidean distance of the ab plane. At a point p where the metric is originally unknown, a metric interpolated from the neighbor points c_i has all the odds of being more relevant than the canonical Euclidean metric of the map, see Fig. 5.2(c).

Let $d_R(p, q)$ be the perceptual distance between color p and color q. $d_R(p, q)$ is the Riemannian distance induced by the field of ellipses. Let $||p-q||$ be the distance associated with the canonical scalar product of the ab plane, and $||p-q||_c$ be the distance associated with the scalar product induced by the ellipse $\mathcal{E}_c$. Let $B(c, R)$ and $B_c(c, R)$ be the respective balls of center c and radius R. The previous discussion can be formalized as follows. It can be shown that:

$$lim_{x\to c}\frac{||x-c||_c}{d_R(x,c)} = 1,$$

while if $||.||_c \neq ||.||$, the equality case being exceptional, thus

$$lim_{x\to c}\frac{||x-c||}{d_R(x,c)} \neq 1.$$

Therefore for such a c there exists $A > 0$ such that,

$$\forall R > 0, \exists x \in B(c, R), A < \left|\frac{||x-c||}{d_R(x,c)} - 1\right|. \tag{5.2}$$

On the other hand there exists a real positive number $R_c = R_{c,A}$ such that,

$$\forall x \in B(c, R_c), \left|\frac{||x-c||_c}{d_R(x,c)} - 1\right| < A. \tag{5.3}$$

We have

$$sup_{B(c,R_c)}\left(\left|\frac{||x-c||_c}{d_R(x,c)} - 1\right|\right) < A < sup_{B(c,R_c)}\left(\left|\frac{||x-c||}{d_R(x,c)} - 1\right|\right).$$

Thus for $x \in B(c, R_c)$, $||x-c||_c$ is preferred to $||x-c||$. Consider a kernel K and a scaling parameter λ such that

$$\lambda \leq R_c \text{ and } B_c(c, \lambda) \subset B(c, R_c).$$

For $x \in B(c, R_c)$, $K\left(\frac{||x-c||_c}{\lambda}\right)$ is preferred to $K\left(\frac{||x-c||}{\lambda}\right)$. For $x \notin B(c, R_c)$, $K\left(\frac{||x-c||_c}{\lambda}\right) = K\left(\frac{||x-c||}{\lambda}\right) = 0$. Therefore, under these assumptions on the scaling parameter λ, the histogram estimator

$$f(x) = \frac{1}{k} \sum_{p_i \in \{pixels\}} \frac{1}{\lambda^n} K\left(\frac{||x - I(p_i)||_{I(p_i)}}{\lambda}\right), \tag{5.4}$$

is preferred to the classical histogram. We think that the hypothesis on λ is reasonable in practice, its validation is a subject of further research. Note that the higher the resolution of the image is, the smaller λ is and then the more the hypothesis becomes reasonable.

5.3.3.1. Metric interpolation and Euclidean approximation

Let $\mathcal{M}$ be topological space, homeomorphic to an open subset of $\mathbb{R}^n$. Let ϕ be an homeomorphism from $U_\phi \subset \mathbb{R}^n$ to $\mathcal{M}$. A set of scalar products G_{c_i} is given for a set of points $\{c_i\} \in \mathcal{M}$. We consider here the problem of interpolation of the field of metrics. Let F_1 and F_2 be two smooth metric fields that coincide with the observed ellipses at the points $\{c_i\}$. Despite the intuition, if no assumption is made on ϕ regarding the Riemannian distance, there are no criteria that enables to prefer F_1 or F_2. The problem of metric tensor interpolation is thus a difficult problem. In this work, we adopt an elementary solution. Ellipses are represented in the projective ab plane. A Delaunay triangulation with respect to the canonical Euclidean metric of the plane is performed on the set $\{c_i\}$, Fig. 5.2(b). At a point p in the triangle $c_ic_jc_k$ the parameters of the interpolated ellipse $\mathcal{E}_p$ are linearly interpolated between the parameters of $\mathcal{E}_{c_i}$, $\mathcal{E}_{c_j}$, and $\mathcal{E}_{c_k}$ with respect to the barycentric coordinates of p. If p does not belong to one of the triangles of the Delaunay triangulation, we set $\mathcal{E}_p = \mathcal{E}_q$ where q is the projection of p on the convex hull of the set of centers.

5.3.3.2. Experimental results

The RIT-DuPont dataset from Berns et al. (1991) shows that the perceptual metric is dependent of the luminance. Nevertheless, for visualization purpose we choose to abandon the luminance information in order to work with two dimensional data. The MacAdam ellipses were measured at a fixed luminance, in the CIE chromaticity diagram. The ellipses are transported in the $L = 40$ plane of the CIELAB space. Forgetting the luminance coordinate, one obtains then a transport of the MacAdam ellipses in an ab plane. Remind that the proposed framework is independent of the dimension and can be used in three dimensional spaces with standard datasets of ellipsoids.

Fig. 5.3 represents the density of the Riemannian measure with respect to the Lebesgue density of the plane. Recall that the expression of the density is given by $\sqrt{det(G)}$ where G is the metric tensor derived from the ellipse.

Let us consider the color image given in Fig. 6.2.2(a) and the three alternative color histograms compared in Fig. 6.2.2. Histograms (b) and (c) aim a studying the density f with respect to the perceptual Riemannian volume measure. The main difference between

(c) and (d) is that in (d) the shape of the kernel follows the Riemannian metric. The density with respect to the Euclidean measure is visibly different from the histogram with respect to the Riemannian measure. The amplitude of the upper spot, representing white colors, is significantly decreased when using the Riemannian measure. Perceptually, this results from the fact that the eyes have an higher sensitivity around white than around blue.

5.3.4. Conclusion and perspectives on perceptual color histograms

Given a set of ellipses representing the perceptual metric on colors, we proposed an approach for histogram computation that takes into account the Riemannian structure of the perceptual metric without introducing supplementary computational complexity. Indeed, the step of ellipses interpolation only has to be achieved once and does thus not introduce computational complexity. The relevance of the approach is conditioned by the relevance of the set of perceptual ellipses and the quality of the interpolation. The deep problem of metric tensor interpolation has been partially left aside and will be subject of future research. The second topic of our future research will be on the convergence of the proposed histogram to the density of the underlying random variable with respect to the interpolated Riemannian metric.

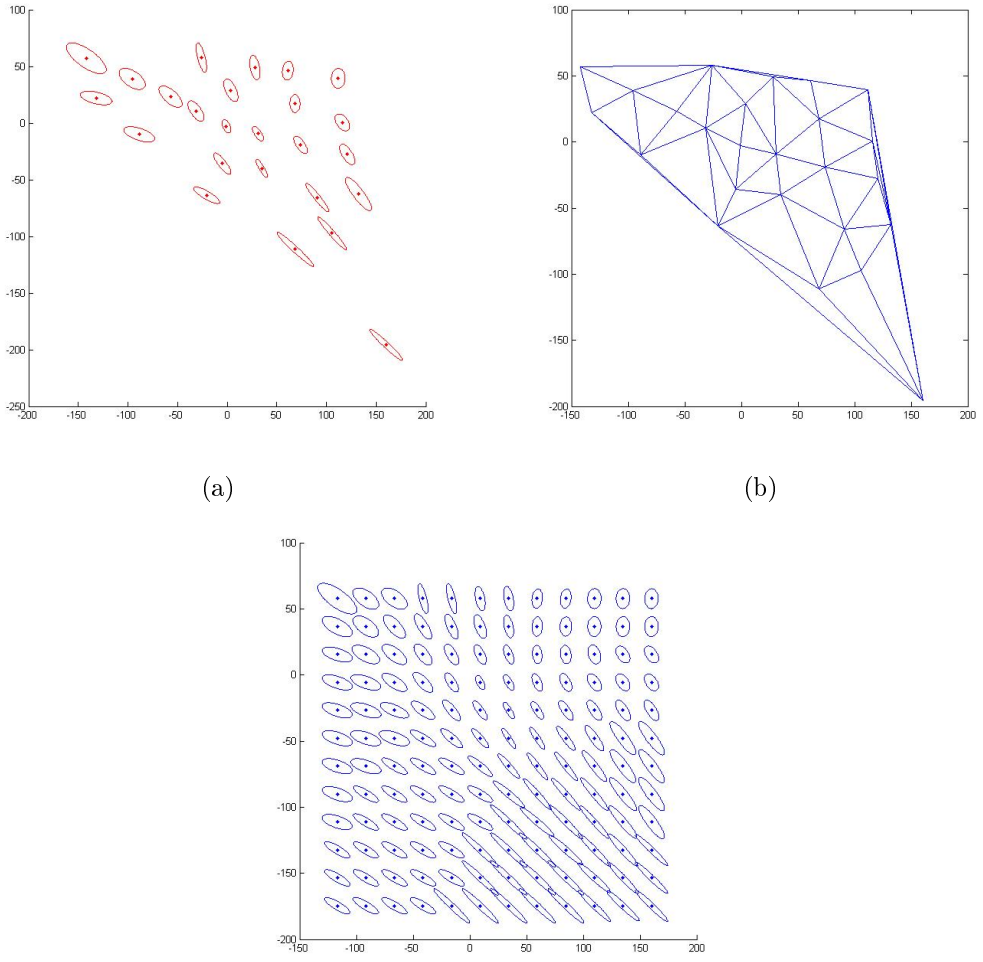

(a) (b)

(c)

Figure 5.2.: (a) MacAdam ellipses transported in the projective *ab* plane, (b) Delaunay triangulation of the space from center of MacAdam ellipses, (c) interpolated ellipses.

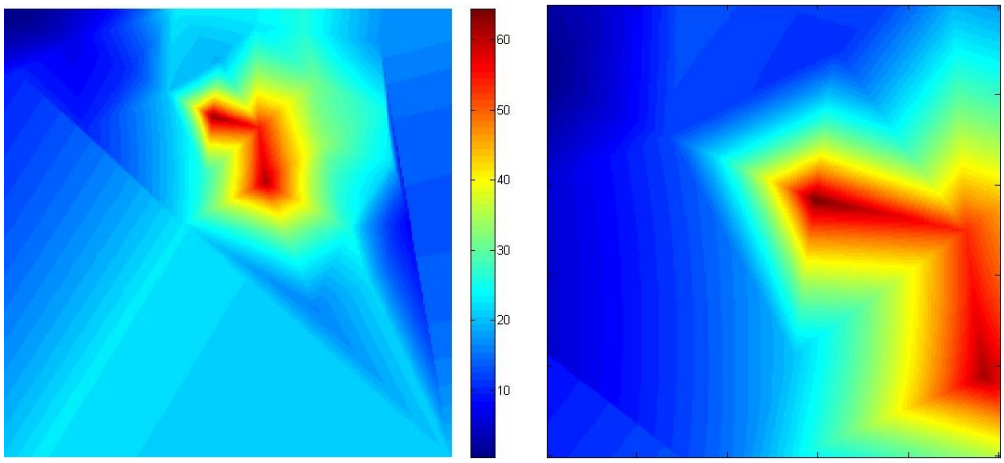

(a) (b)

Figure 5.3.: (a) Local density change induced by the interpolated ellipses in *ab* plane, (b) Zoom adapted to colors present in the color image 6.2.2(a).

(a)

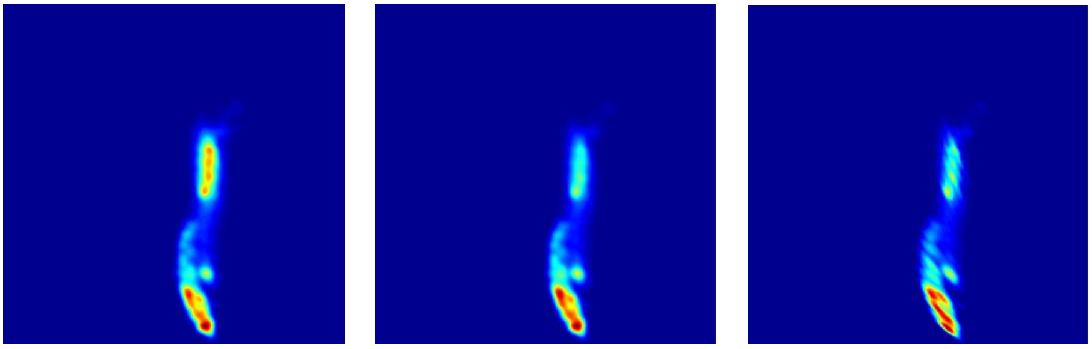

(b) (c) (d)

Figure 5.4.: Example of histogram of image (a) computed using: (b) the canonical Euclidean metric of the *ab* projective plane, (c) the canonical metric followed by a division by the local density of the perceptual metric, (d) the formula (5.4).

6. Non-parametric probability density estimation on spaces of Gaussian laws

Résumé

Le traitement du signal et de l'image manipulent des données se trouvant dans des espaces de plus en plus variés. Chaque type de données a ses propres structures algébriques et géométriques. Les lois Gaussiennes et les matrices symétriques définies positives sont de plus en plus présentes en traitement du signal et de l'image, d'où l'interêt de l'étude de la géométrie de ces espaces. Ce chapitre porte sur l'estimation de densités de probabilités dans le cas de variables aléatoires à valeurs dans les lois gaussiennes. L'étude se place principalement dans le contexte des géometries induites par les métriques de Fisher et de Wasserstein. Pour chaque des métriques et dans chaque espace de lois, différentes techniques d'estimation sont étudiées.

6.1. Introduction

Signal and image processing are dealing nowadays with data laying in more and more various spaces. Each type of data posses its own geometric and algebraic structures. Data laying in spaces of Gaussian laws or symmetric positive definite matrices $\mathrm{SPD}(n)$ are more and more present in signal and image processing, hence the importance of the study of the geometry of these spaces.

The most common geometries of Gaussian laws are the Fisher information metric and the Wasserstein metric. The Fisher metric on multivariate centered Gaussians happens to be similar to the affine invariant metric on $\mathrm{SPD}(n)$, that is to say the metric invariant under the action of invertible matrices $Gl(n)$, i.e.,

$$G \cdot \Sigma \mapsto G\Sigma G^t.$$

Furthermore the Fisher metric on univariate Gaussians is a hyperbolic metric of dimension two. Thus understanding the Fisher metric has consequences beyond Gaussian laws. The Wasserstein metric is another name for the earth mover's distance. Let S be a set of probability distributions on E. A fundamental difference between the Fisher metric and the Wasserstein metric on S, is that the Fisher metric do not relay on a geometry of E while the Wasserstein metric does. Indeed, the Wasserstein metric depends on the geometry of E through a transportation cost. The dependence on the geometry of E constitutes one of the main specificities and interests of the Wasserstein metric.

Median and means are fundamental quantities in signal and image processing. Several studies already address their definition and computation for the Fisher and the Wasserstein metrics, see for instance Pennec (2004); Yang (2009); Agueh and Carlier (2011); Arnaudon et al. (2013); Bini and Iannazzo (2013). Although being a secondary problem, density estimation is an important tool of signal and image processing which is still little studied

on spaces of Gaussian laws. It is in particular useful, for instance, for segmenting point clouds or for Bayesian classification. The study of distributions on SPD(n) is subject to recent studies. In Said et al. (2015), it is proposed for instance a generalization of the Gaussian distribution on SPD(n). To our knowledge, the existing literature dedicated to density estimation is mainly restricted to Huckemann et al. (2010), Peter T. Kim and Donald St. P. Richards (2011) and Asta (2014). These previous works have focussed on techniques derived from the orthogonal series density estimation for the Fisher metric in the case of centered multivariate and univariate Gaussian laws.

This chapter is a practical guide to density estimation on spaces on Gaussian laws and can be summarized as follows.

- Main techniques of density estimation on Riemannian manifolds are reviewed in Section 6.2.
- Application to each specific situation of Gaussian law spaces are discussed in Sections 6.4-6.13.
- The computation of kernels for the Fisher and Wasserstein metric are provided. The expression of the different factors constituting these kernels in the case of the Fisher metric are already known in the literature, see Terras (1988); Wijsman (1990).
- The contributions of Sections 6.6.2 and 6.7.2 consist in presenting a computation relying on a low level Riemannian geometry and Lie groups, and providing a "ready-to-use" expression of the kernels. To our knowledge, the expression of kernels for the Wasserstein metric is new.
- Being an important problem for practical applications, the density estimation of partial quantities such as mean eigenvalues and rotation is discussed in Section 6.13.

6.2. Non parametric density estimation on Riemannian spaces

Let Ω be a space, endowed with a σ-algebra and a probability measure p. Let X be a random variable $\Omega \to \mathcal{M}$ with $\mathcal{M}$ homeomorphic to U an open subset of $\mathbb{R}^n$. The homeomorphism is noted ϕ. $\mathcal{M}$ is equipped with a Riemannian metric G. The Riemannian measure is called vol and the measure on $\mathcal{M}$ induced by X is noted μ_X. We assume that μ_X has a density, noted f, with respect to vol, and that the support of X is a compact set noted $Supp$. Let $(x_1, .., x_k) \in \mathcal{M}^k$ be a set of draws of X.

The Dirac measure in point a is defined as:

$$\delta_a(U) : \begin{cases} 1 \text{ if } a \in U \\ 0 \text{ if } a \in \complement U \end{cases}$$

Let $\mu_k = \frac{1}{k}\sum_i \delta_{x_i}$ denotes the empirical measure of the set of draws. This section presents the main techniques of estimation of f from the set of draws $(x_1, .., x_k)$. The estimated density in x in $\mathcal{M}$ is noted $\hat{f}_k(x) = \hat{f}(x, x_1, ..., x_k)$. Observe that $\hat{f}_k(x)$ can be seen as a random variable. The relevance of density estimation technique depends on several aspects. When the space allows it, the estimation technique should not privilege specific directions or locations. This results in an isotropy and a homogeneity condition. In the kernel method for instance, a kernel function K_{x_i} is placed at each observation x_i.

Firstly, in order to treat directions equally, the function K_{x_i} should be invariant by the isotropy group of x_i. Secondly, for an other observation x_j, functions K_{x_i} and K_{x_j} should be similar up to the isometries that send x_i on x_j. These considerations strongly depend on the geometry of the space: if the space is not homogeneous and the isotropy group is empty, these indifference principles have no meaning. The convergence of the different estimation techniques is widely studied. Results were first obtained in the Euclidean case, and are gradually extended to the probability densities on manifold, see Hendriks (1990); Pelletier (2005); Huckemann et al. (2010); Asta (2014). The last relevant aspect, is computational. Each estimation technique has its own computational framework, which presents pro and cons given the different applications. For instance, the estimation by orthogonal series presents an initial pre-processing, but provides a fast evaluation of the estimated density in compact manifolds.

6.2.1. Estimation from the Euclidean structure of the parametrization

Assume the term $\sqrt{det(G(x))}$ known for all $x \in U$. According to Chapter 4,

$$\frac{dvol}{dLeb}(x) = \sqrt{det(G(x))}.$$

If we have an estimator $\hat{f}^{Eucl}$ of the density with respect to the Lebesgue measure of U, noted Leb_U, it is possible to obtain an estimator of the density with respect to the Riemannian measure vol:

$$\hat{f} = \hat{f}^{Eucl}\frac{dLeb_U}{dvol} = \frac{1}{\sqrt{det(G(x))}}\hat{f}^{Eucl}. \tag{6.1}$$

The estimation $\hat{f}$ is a probability density with respect to the Riemannian measure vol. However the estimation does not respect homogeneity and isotropy considerations. If $\hat{f}^{Eucl}$ is constructed from an Euclidean kernel K, the "Riemannian shape" of K will differ from x_i to x_j in an uncontrolled way.

6.2.2. Histograms

We provide here an L^2 convergence rate of the histogram when the random variable X is valued on a Riemannian space, as formulated in Section 4.2, and whose support is compact. We chose not to introduce the notion of manifold in the thesis. The result is however valid on Riemannian manifolds. Even if this results might already be known, we have not been able to find it in the literature. The closest result we are aware of in can be found in Geffroy (1974), where the author provides a condition on the uniform convergence of the histogram when X is valued in a metric space, without assumptions on its support.

Let $(A_j^k)_{0\leq j\leq J_k}$ be a partition of U. Let α_j be the number of draws that lays in A_j^k, i.e., $\alpha_j = Card\{x_i \in A_j^k\}$. The histogram of the draws (x_i) is the function

$$\hat{f}_k = \frac{1}{k}\sum_j \frac{1}{vol(A_j^k)}\alpha_j \mathbf{1}_{A_j^k}, \tag{6.2}$$

where $\mathbf{1}_E$ is the indicator function of E. $\hat{f}_k$ is an estimator of the density f.

We suppose that:
i. The density f is α-Lipschitz with respect to the Riemannian metric;
ii. There exists a constant C such that, for all k and all i,j, one has

$$\frac{vol(A_j^k)}{vol(A_i^k)} \leq C,$$

and C will be called the homogeneity factor;
iii. $\beta_k = max_j(R_j^k)$ goes to zero when k goes to infinity, where $2R_j^k$ is the diameter of A_j^k.

Estimator $\hat{f}_k(x)$ is seen as a random variable of the set of draws. The quantity

$$e = \mathbb{E}\left(\int_{Supp}(f(x) - \hat{f}_k(x))^2 dx\right) = \int_{Supp}\mathbb{E}((f(x) - \hat{f}_k(x))^2)dx,$$

is the mean square error of the estimation. It can be developed as:

$$\begin{aligned} e &= \int_{Supp}(f(x) - \mathbb{E}(\hat{f}_k(x))^2 dx + \int_{Supp}\mathbb{E}((\mathbb{E}(\hat{f}_k(x)) - \hat{f}_k(x))^2) \\ &= \int_{Supp}(f(x) - \mathbb{E}(\hat{f}_k(x)))^2 + \int_{Supp} Var(\hat{f}_k(x))dx \\ &= e_1 + e_2. \end{aligned}$$

For x in A_j^k the random variable $kvol(A_j^k)\hat{f}_k(x) = \alpha_j$ has a binomial distribution $\mathcal{B}(k, \mu_X(A_j^k))$, hence

$$\mathbb{E}(\hat{f}_k(x)) = \frac{\mu_X(A_j^k)}{vol(A_j^k)},$$

$$Var(\hat{f}_k(x)) = \frac{k\mu_X(A_j^k)(1 - \mu_X(A_j^k))}{(kvol(A_j^k))^2} \leq \frac{\mu_X(A_j^k)}{kvol(A_j^k)^2}.$$

It follows that

$$e_2 \leq \sum_j \int_{A_j^k} \frac{\mu_X(A_j^k)}{kvol(A_j^k)^2} dx = \sum_j vol(A_j^k)\frac{\mu_X(A_j^k)}{kvol(A_j^k)^2}.$$

Call N_k the number of bins and $\mathcal{V}$ the measure of the support $vol(Supp)$. By the assumption (ii), $CN_k vol(A_j^k) \geq \mathcal{V}$, therefore

$$e_2 \leq \sum_j \frac{\mu_X(A_j^k)}{k\frac{\mathcal{V}}{CN_k}} = \frac{CN_k}{k\mathcal{V}}.$$

Note that $\sum_j \mu_X(A_j^k) = 1$ since μ_X is a probability.

Assuming that the A_j^k are connected for each j, there exists y_j^k in A_j^k such that

$$E(\hat{f}_k(x)) = \frac{\mu_X(A_j^k)}{vol(A_j^k)} = \frac{1}{vol(A_j^k)}\int_{A_j^k} f(x)dx = f(y_j^k).$$

We have

$$e_1 = \sum_j \int_{A_j^k} (f(x) - \mathbb{E}(\hat{f}_k(x)))^2 dx = \sum_j \int_{A_j^k} (f(x) - f(y_j^k)^2 dx,$$

and since f is α-Lipschitz we obtain

$$\begin{aligned} e_1 &\leq \sum_j (\alpha.R_j^k)^2 vol(A_j^k) \\ &\leq (\alpha\beta_k)^2 \mathcal{V}. \end{aligned}$$

Finally we have

$$e \leq \frac{CN_k}{k\mathcal{V}} + (\alpha\beta_k)^2 \mathcal{V}.$$

Assume now that there exist positive real numbers C_1,C_2 such that we have

$$\forall j,\ C_1 (R_j^k)^n \leq vol(A_j^k) \leq C_2 (R_j^k)^n.$$

When the radius R goes to zero, the volume of a ball B_R becomes proportional to $\mathbb{R}^n$, see Gray (1974). Thus the existence of C_1,C_2 is an assumption on the isotropy of the partition. It follows that for all i, j

$$\frac{C_1 (R_j^k)^n}{C_2 (R_i^k)^n} \leq C,$$

which implies that $min_j(R_j^k) \geq \left(\frac{C_1}{CC_2}\right)^{1/n} \beta_k = \gamma^{1/n}\beta_k$. Hence

$$N_k C_1 \gamma \beta_k^n \leq \mathcal{V},$$

and

$$e \leq \frac{C}{k\mathcal{V}} \times \frac{\mathcal{V}}{C_1\gamma\beta_k^n} + (\alpha\beta_k)^2\mathcal{V} = \frac{A}{k\beta_k^n} + B\beta_k^2 = F(\beta_k),$$

where A and B are two constants depending on the geometry and on the density f. Let β_k be the value that minimize the right hand side of the above inequality. We have:

$$\beta_k = \left(\frac{n.A}{2Bk}\right)^{\frac{1}{n+2}},$$

and therefore the minimum of F is:

$$min(F) = D.k^{\frac{1}{n+1}}.$$

The obtained convergence rates are similar to the rates of Euclidean regular multidimensional histograms. The constant D depends on the homogeneity and the isotropy of the partition through C, C_1 and C_2. In practice, finding a partition with satisfying factors of homogeneity and isotropy can be difficult. It can be noted that since the bins are fixed,

even with ideal factors, $C = 1$, $C_1 = C_2$, the estimation can not be exactly invariant to isometries.

Compared to other methods, the main advantage of histograms is often the low algorithmic complexity and the simplicity of use. The different costs that have to be taken into account are: the construction of the partition and the cost $c(x)$ of the search of the bin A_j^k such that $x \in A_j^k$. Function $c(x)$ depends on the geometry of the space. The second technical aspect that has to be taken into account is the nature of the representation of the partition: list, graph, matrix, etc. Indeed for further processing, the histogram is often viewed as an image, its structure plays then an important role.

6.2.3. Kernel

This part is mainly based on the work by Pelletier (2005). Let $K : \mathbb{R}_+ \to \mathbb{R}_+$ be a map which verifies the following properties:
i) $\int_{\mathbb{R}^d} K(||x||)dx = 1$,
ii) $\int_{\mathbb{R}^d} xK(||x||)dx = 0$,
iii) $K(||x|| > 1) = 0$,
iv) $sup(K(x)) = K(0)$.

Let $p \in \mathcal{M}$ and $x \in U$ with $\phi(x) = p$. In the general case, given a point $p \in \mathcal{M}$, exp_p defines an injective application only on a neighborhood of $0 \in T_p\mathcal{M}$. However in every studied situations of the present work, the parametrization induced by exp_p is injective on the whole space. Recall that when $T_p\mathcal{M}$ is endowed with the scalar product $G(x)$, Euclidean distances to p corresponds to Riemannian distance to p. The Lebesgue measure of $T_p\mathcal{M}$ is noted Leb_p. The function θ_p defined by:

$$\theta_p : q \mapsto \theta_p(q) = \frac{dvol}{dexp^*(Leb_p)}(q), \tag{6.3}$$

is the density of the Riemannian measure of $\mathcal{M}$ with respect to the Lebesgue measure Leb_p after the identification of $\mathcal{M}$ and $T_p\mathcal{M}$ induced by exp_p, see Fig.6.1.

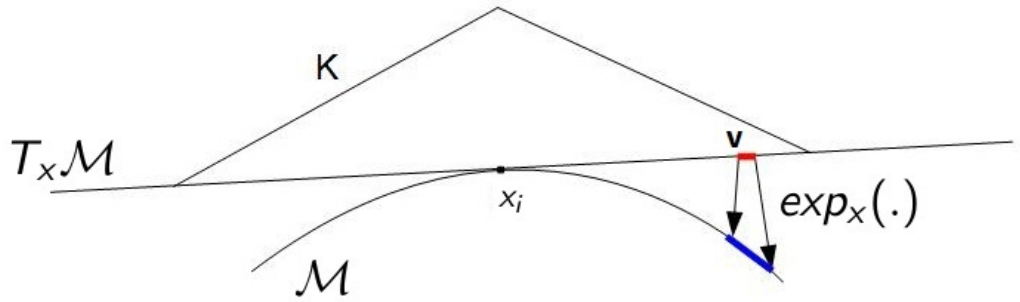

Figure 6.1.: Volume change θ_{x_i} induced by the exponential map

Given K and a positive radius r, the estimator of f proposed by Pelletier (2005) is

defined by:

$$\hat{f}_k = \frac{1}{k}\sum_i \frac{1}{r^n}\frac{1}{\theta_{x_i}(x)} K\left(\frac{d(x,x_i)}{r}\right). \tag{6.4}$$

The corrective factor $\theta_{x_i}(x)^{-1}$ is necessary since the kernel K originally integrates to one with respect to the Lebesque measure, while we want it now to integrate to one with respect to the Riemannian measure. It can be noted that this estimator is the usual kernel estimator is the case of Euclidean space. Under reasonable assumptions on the true density f, the shape of the kernel does not have a significant impact on the quality of the estimation in the Euclidean context, see for instance Silverman (1986). Simulations summarized in Fig. 6.4 experimentally confirms the result in a specific Riemannian situation.

Let $r^{inj} = \sup_{p\in\mathcal{M}}(r_p^{inj})$. Let $\delta = \sup\mathcal{K}$ the supremum of the sectional curvature in $\mathcal{M}$. For a definition of the sectional curvature, see. Carmo (1992). Then, Pelletier (2005) provides the following result.

Theorem 6.2.1 (Pelletier (2005)) *Let x_i be an arbitrary point of $\mathcal{M}$. Let μ be a measure on $\mathcal{M}$ whose density with respect to the measure vol is*

$$\frac{1}{r^n}\frac{1}{\theta_{x_i}(.)} K\left(\frac{d(.,x_i)}{r}\right).$$

For $r < \min\{\frac{r^{inj}}{2}, \frac{\pi}{4\sqrt{\delta}}\}$ when $\delta > 0$ or $r < \frac{r^{inj}}{2}$ when $\delta \leq 0$, x_i is an intrinsic mean of μ.

Theorem. 6.2.1 ensures the kernel put on x_i is "centered" on x_i.

Despite convergence rates provided by Pelletier (2005) are obtained for compact manifolds without boundaries, they remain valid for non-compact manifolds when the support of X is compact. Indeed the double manifold, see Munkres (1966), enable to see a compact manifold with boundaries as a submanifold of a compact manifold without boundaries. Assume f two-times differentiable with bounded second covariant derivative. See Parker (2015) for a definition of covariante derivative. Pelletier (2005) also provides the following convergence rates.

Theorem 6.2.2 (Pelletier (2005)) *For r satisfying conditions of theorem. 6.2.1, there exist a constant C_f such that:*

$$\int_{Supp} \mathbb{E}[(f(x) - \hat{f}_k(x))^2] d\mu \leq C_f\left(\frac{1}{kr^n} + r^4\right),$$

then for r equivalent to $k^{-\frac{1}{n+4}}$:

$$\int_{Supp} \mathbb{E}[(f(x) - \hat{f}_k(x))^2] d\mu = O(k^{-\frac{4}{n+4}}).$$

For the expression of the constants see details in Pelletier (2005).

Given a reference point $p_{ref} \in \mathcal{M}$, let

$$\tilde{K}(q) = \frac{1}{kr^n}\frac{1}{\theta_{p_{ref}}(q)} K(d(p_{ref}, q)).$$

Lemma 6.2.3 *If ϕ is an isometry of $\mathcal{M}$, then $\theta_p(q) = \theta_{\phi(p)}(\phi(q))$.*

Recall that μ_k is the empirical measure of the set of draws.

Theorem 6.2.4 *If $\mathcal{M}$ is homogeneous, then $\hat{f}_k = \mu_k * \tilde{K}$.*

Proof Since $d(x_i, .)$ and $\theta_p(.)$ are invariant by the isotropic group of p_{ref}, also is $\tilde{K}$. Thus the convolution is well defined. Using notations of Chapter 4, one has

$$\begin{aligned}
(\delta_i * \tilde{K})(p) &= \int \delta_i(q)\tilde{K}_q(p)d\mu \\
&= \tilde{K}_{x_i}(p) \\
&= \tilde{K}(iso_{x_i,p_{ref}}(p)) \\
&= \frac{1}{kr^n}(\theta_{p_{ref}}(iso_{x_i,p_{ref}}(p)))^{-1}K(d(p_{ref}, iso_{x_i,p_{ref}}(p))) \\
&= \frac{1}{kr^n}\frac{1}{\theta_{x_i}(p)}K(d(x_i,p)).
\end{aligned}$$

Hence,

$$\begin{aligned}
\mu_k * \tilde{K} &= \sum_i \delta_i * \tilde{K} \\
&= \frac{1}{kr^n}\sum_i (\theta_{x_i}(.))^{-1}K(d(x_i,.)) \\
&= \hat{f}_k.
\end{aligned}$$

Therefore, when $\mathcal{M}$ is homogeneous, the density estimation proposed by Pelletier (2005) is a convolution. Each location and direction are processed as similarly as possible. As a result, it can be easily shown that for an isometry α we have:

$$\hat{f}_k(x, x_1, ..., x_k) = \hat{f}_k(\alpha(x), \alpha(x_1), ..., \alpha(x_k)).$$

It must be noted that the kernel method requires the existence of explicit expression of distances and of the function θ_p, which renders it unusable in several situations. The kernel density estimation presents better convergence rates than histograms. The only computational aspect that has to be taken into account is the cost of evaluation of the density at a point x in $\mathcal{M}$. The computational cost of $\hat{f}_k(x)$ is in $O(k)$. The constant mainly depends on the computational cost of distances. Note that when the number of evaluations is important, histograms present lower computational cost.

The k-nearest neighbor is an interesting variant of the standard kernel density estimation. It consists in making the parameter r depend on x_i, by setting r as the distance to the k-th nearest neighbor of x_i. This estimator has been studied on Riemannian manifolds in Henry et al. (2011).

6.2.4. Orthogonal series

Instead of directly studying the density f, the estimation is made from the estimation of the scalar product between f and a set of orthonormal functions $\{e_j\}$. Recall that for g and h in $L^2(\mathcal{M})$:

$$\langle g, h\rangle = \int_{\mathcal{M}} g\overline{h} dvol,$$

where $\overline{x}$ denotes the complex conjugate of x. We have that:

$$\langle f, e_i\rangle = \int f\overline{e_j} dvol = \mathbb{E}\left(\overline{e_j}(X)\right).$$

Thus $\langle f,e_j\rangle$ can be approximated by an estimation of the expectation, i.e.,

$$\mathbb{E}\left(e_j(X)\right) \approx \frac{1}{k}\sum_{j=1}^{k} e_j\left(x_i\right)$$

Or in other words:

$$\int e_j f dvol \approx \int e_j d\mu_k.$$

Now, given an integer N, the orthogonal series estimator is defined as:

$$\hat{f} = \frac{1}{k}\sum_{j=-N}^{N}\left[\sum_{i=1}^{k}\overline{e_j}\left(x_i\right)\right] e_j. \tag{6.5}$$

The formula becomes an integral when the number of functions e_j is uncountable. First the base has to be ordered such that the norm of the rest of the decomposition, i.e., $\sum_{|j|>N}\langle f, e_j\rangle e_j$, decreases as fast as possible for regular functions. Second, in order to process locations and directions indifferently, the basis functions must present regularity properties regarding the metric. Eigenfunctions of the Laplacian operator are very good candidates. Indeed the Fourier transform on $\mathbb{R}^n$, or the Fourier–Helgason transform on symmetric spaces, are highly related to the convolution by isotropic kernel. When the e_j functions are Fourier functions, the estimation technique is also called the characteristic function method.

When the space is compact, the spectrum of the Laplacian operator is discrete. Under the compact assumption, let $\{e_j\}$ be an orthonormal Hilbert basis given by the eigenfunctions of the Laplacian. Given $T > 0$ we define:

$$F_T = \{e_j | \Delta e_j = \lambda e_j \text{ with } \lambda < T\}.$$

Using this orthogonal series, the density estimator becomes:

$$\hat{f} = \frac{1}{k}\sum_{e_j\in F_T}\left[\sum_{i=1}^{k} e_j\left(x_i\right)\right] e_j.$$

Hendriks (1990) provides the following convergence rates for density estimators based on eigenfunctions of the Laplacian.

Theorem 6.2.5 (Hendriks (1990)) *Suppose f is s times differentiable with square integrable derivatives. Let $T_0 > 0$. There are constants A (depending on $\mathcal{M}$) and B (depending on f) such that for $T \geq T_0$,*

$$\int_{Supp} \mathbb{E}(f(x) - \hat{f}_k(x))^2 dx \leq A\frac{T^{n/2}}{k} + BT^{-s},$$

and for a suitable decrease of T,

$$\int_{Supp} \mathbb{E}(f(x) - \hat{f}_k(x))^2 dx \leq O(k^{\frac{-2s}{2s+n}}).$$

Thus, the convergence rate for the orthogonal series is similar to the rate of the kernel estimation.

The Fourier series on a parallelepiped of $\mathbb{R}^n$ are a particular case of decomposition on eigenfunctions of the Laplacian.

When the space is not compact, the estimator takes an integral form. In the case of a continuous spectrum,

$$F_T = \{e_s | \Delta e_s = \lambda_s e_s(y),\ \lambda_s < T, s \in \mathcal{S}\},$$

for some uncountable set $\mathcal{S}$. Then, in this uncountable case, one has

$$\hat{f} = \frac{1}{k}\int_{F_T} \left[\sum_{i=1}^{k} e_s(x_i)\right] e_s d_* s,$$

where d_*s is the spectral measure on $\mathcal{S}$ and in the case of symmetric spaces is the Harish-Chandra $c-$function, see Helgason (2001).

The explicit expressions of F_T and of the estimators are provided when necessary in the following sections. For the space of positive definite matrices equipped with the metric invariant by the action of $Gl(n)$ on $\mathrm{SPD}(n)$,$(G, M) \mapsto GMG^t$, Peter T. Kim and Donald St. P. Richards (2011) provided convergence rates of the estimator based on the Fourier–Helgason transform. Helgason functions are eigenvalues of the Laplacian operator, i.e.,

$$\mathcal{H}(\Delta f)(w) = \lambda_\omega \mathcal{H}(f)(w).$$

Then, $\Delta^{\sigma/2}$ is defined as the operator verifying:

$$\mathcal{H}(\Delta^{\sigma/2} f)(w) = \lambda_\omega^{\sigma/2} \mathcal{H}(f)(w).$$

Let $H_\sigma(Q) = \{f \in C^\infty, ||\Delta^{\sigma/2} f|| < Q\}$. We have the following result.

Theorem 6.2.6 (Peter T. Kim and Donald St. P. Richards (2011)) *If $f \in H_\sigma(Q)$ with $2\sigma > dim\,\mathrm{SPD}(n)$*

$$\int_{Supp} E(f(x) - \hat{f}_k(x))^2 = O(k^{-2\sigma/(2\sigma + dim\,\mathrm{SPD}(n))}).$$

The hypothesis on σ is slightly different but highly related to the differentiability hypothesis of the kernel method. Similar rates on the hyperbolic space of dimension two can be found here Huckemann et al. (2010).

The Fourier–Helgason method on SPD(n) and on the hyperbolic space are convolutions:

$$\hat{f}_k = \mu_k * \mathcal{H}^{-1}(\mathbf{1}_{F_T}).$$

Recall that the Fourier–Helgason of the convolution of a function by a kernel is a product of transforms when the kernel is invariant by the isotropy group. This can be verified on expressions in following sections, by showing that $\mathcal{H}^{-1}(\mathbf{1}_{F_T})$ is invariant by the isotropy group $SO(n)$ for SPD(n) and $SO(2)$ for the hyperbolic space. Thus, as for the kernel method, each location and direction are processed as similarly as possible. The main drawback of this method in the non-compact case is its computational cost: the estimation requires the computation of an integral.

6.2.5. Curse of dimensionality

Observe that the convergence rates deteriorate with the dimension in $k^{a/(b+n)}$ where a and b are constants and n is the dimension of the space. To reach the same precision, the number of draws (x_i) has to increase exponentially with the dimension. Recall that the dimension of the space of centered n-dimensional Gaussian laws is of dimension $\frac{1}{2}(n^2+n)$ and that the dimension of n-dimensional Gaussian laws is of dimension $\frac{1}{2}(n^2+3n)$. Thus the quality of the density estimation quickly deteriorates with the dimension of the underlying space. Fortunately, as shows section. 6.4.1, classical metrics on spaces of Gaussian laws allow to reduce the dimension by separating the different components of the Gaussian law (mean, eigenvalues and rotation). Dimension can also be reduced using principal geodesic analysis on manifolds, see Fletcher and Joshi (2004).

6.3. Classical metrics on Gaussian laws

6.3.1. The Fisher metric

We consider here a set $\mathcal{M}$ of probability measures on a measure space (X, σ_X, μ) homeomorphic to an open subset of $\mathbb{R}^n$. Let $\theta = (\theta_1, ..., \theta_n) \in U \subset \mathbb{R}^n$ be an homeomorphic parametrization of $\mathcal{M}$. The distributions are assumed to have a density $f(., \theta)$. The Fisher matrix $G(\theta)$ associated to the parametrization θ is defined as follows, see Amari et al. (2007):

$$g(\theta)_{i,j} = E\left[\frac{\partial ln(f(x,\theta))}{\partial \theta_i}\frac{\partial ln(f(x,\theta))}{\partial \theta_j}\right] = \int_x \frac{\partial ln(f(x,\theta))}{\partial \theta_i}\frac{\partial ln(f(x,\theta))}{\partial \theta_j} f(x,\theta)dx. \quad (6.6)$$

The matrix $G(\theta)$ is symmetric positive definite. The matrices $\{G(\theta), \theta \in U\}$ induce a smooth field of positive definite matrices on U, that is to say a Riemannian metric. It can be shown that the distances and the geodesic paths induced on $\mathcal{M}$ by this Riemannian metric do not depend on the choice of the parametrization. It is interesting to note that the Fisher metric do not rely on a metric on X.

Let P and Q be two probability measures. The relative entropy is defined as follows:

$$D(P,Q) = \int_x ln\left(\frac{dP}{dQ}(x)\right) dP.$$

Function $D(.,.)$ is not symmetric and is therefore not a distance. The relative entropy is also called the Kullback–Leibler divergence. The Kullback–Leibler divergence measures an informational difference between P and Q.

Let $D(\theta, \theta_0)$ be the relative entropy between the laws parametrized by θ and θ_0. It can be shown that:

$$\frac{\partial D(\theta_0, \theta)}{\partial \theta_i}(\theta_0) = 0,$$

and

$$\frac{\partial^2 D(\theta_0, \theta)}{\partial \theta_i \partial \theta_j}(\theta_0) = g(\theta_0)_{i,j}.$$

Thus, the Fisher matrix is also the Hessian of the Kullback–Leibler divergence.

6.3.2. The Wasserstein metric/ Earth mover's distance

The Wasserstein metric on probability distributions is the metric of optimal transport as founded by Monge (1781), see Villani (2009) for a modern and complete mathematical overview of the problem. If the probability distributions are seen as earth heaps, the Wasserstein distance between two distributions is the minimum cost that is needed to transform the first heap into the second one. The cost of transport of earth depends of the amount of earth and the traveled distance. Due to this interpretation, the Wasserstein metric is also called the earth mover's distance.

More formally, let X and Y be two random variables valued in a space S endowed with a σ-algebra $\mathcal{A}_S$. Let μ_X and μ_Y be the measures induced on S by X and Y. Let $\Gamma(\mu_X, \mu_Y)$ be the set of measures on $S \times S$ such that μ_X and μ_Y are the first and the second marginals, i.e., for $\gamma \in \Gamma(\mu_X, \mu_Y)$, one has

$$\gamma(A, S) = \mu_X(A), \quad \gamma(S, B) = \mu_X(B),$$

for all $A, B \in \mathcal{A}_S$. Thus $\Gamma(\mu_X, \mu_Y)$ represents the set of possible transports of the measure μ_X on the measure μ_Y. The p-Wasserstein distance between X and Y is then defined by:

$$d(X, Y) = \left[inf_{\gamma \in \Gamma(\mu_X, \mu_Y)} \int_{S \times S} d(a, b)^p d\gamma \right]^{1/p}. \tag{6.7}$$

For most families probability distributions, the Wasserstein distance has no explicit expression. Fortunately, the expression of the 2-Wasserstein distance between Gaussian laws has a simple form. Given two Gaussian distributions $\mathcal{N}_1$, $\mathcal{N}_2$ of mean m_1, m_2 and covariances Σ_1, Σ_2 the distance becomes:

$$d(\mathcal{N}_1, \mathcal{N}_2)^2 = |m_1 - m_2|^2 + tr(\Sigma_1) + tr(\Sigma_2) - 2tr(\sqrt{\Sigma_2^{1/2} \Sigma_1 \Sigma_2^{1/2}}).$$

Takatsu (2011) has showed that the 2-Wasserstein distance on Gaussian laws is a Riemannian metric distance. We restrict the present study to the case of the 2-Wasserstein distance.

6.4. The space of multivariate Gaussian laws

In this section we study the space of multivariate Gaussian laws under two classical Riemannian metrics: the Fisher metric induced by the Fisher matrix, and the Wasserstein metric induced by the optimal transport. A Gaussian law is determined by its mean vector m and its covariance matrix Σ and is noted $\mathcal{N}_{m,\Sigma}$. It is well known that the expression of its density $f_{m,\Sigma}$ is written as:

$$f_{m,\Sigma}(x) = \frac{1}{(2\pi)^{N/2} det(\Sigma)^{1/2}} e^{\frac{-1}{2}(x-m)^t \Sigma^{-1}(x-m)}.$$

The space of n-dimensional Gaussian laws, noted $\mathcal{G}(n)$, can be identified with the product between the vector space $\mathbb{R}^n$ and the space of positive definite matrices $\mathrm{SPD}(n)$. The space $\mathrm{SPD}(n)$ is an open cone included in the vector space of symmetrical matrices $Sym(n)$. Let $\{e_{m_i}\}_{i<n}$ be the canonical basis of $\mathbb{R}^n$ and $\{E_{i,j}\}$ be the canonical basis of $M(n)$. Let

$$E'_{i,i} = E_{i,i}, \quad \text{and} \quad E'_{i,j} = \frac{1}{\sqrt{2}}(E_{i,j} + E_{j,i}), \quad \text{if} i < j.$$

Using this parametrization, $\left\{E'_{i,j}\right\}_{i\leq j}$ forms a basis of the space of symmetric matrices. The space $Sym(n)$ is identified to $\mathbb{R}^{\frac{n(n+1)}{2}}$ through the coordinates in the basis $\left\{E'_{i,j}\right\}_{i\leq j}$. The union $\{x_l\} = \{e_{m_i}\}_{i<n} \cup \left\{E'_{j,k}\right\}_{j\leq k}$ is a basis of $\mathbb{R}^n \times Sym(n)$.

Let consider

$$\phi : \left\{ \begin{array}{ccc} \mathbb{R}^n \times \mathrm{SPD}(n) & \to & \mathcal{G}(n) \\ (m,\Sigma) & \mapsto & \mathcal{N}_{m,\Sigma} \end{array} \right.$$

Since $\mathbb{R}^n \times \mathrm{SPD}(n)$ is an open subset of the vector space $\mathbb{R}^n \times Sym(n)$ the situation is identical to the one described in Chapter 4.

The diagonalization $\Sigma = RDR^t$ of Σ induces another parametrization of $\mathcal{G}(n)$:

$$\psi : \left\{ \begin{array}{ccc} \mathbb{R}^n \times \mathbb{R}^n_+ \times SO(n) & \to & \mathcal{G}(n) \\ (m,\lambda,R) & \mapsto & \mathcal{N}_{m,(RD_\lambda R^t)} \end{array} \right.$$

where $D_\lambda = \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix}$. In a certain way, (λ, R) are similar to the polar coordinates, with the difference that R is not necessary of determinant 1 in polar coordinates. The parametrization ψ is not injective, i.e., given $\Sigma \in \mathrm{SPD}(n)$ there exists several $(\lambda, R) \in \mathbb{R}^n_+ \times SO(n)$ such that $\Sigma = RD_\lambda R^t$. If this is not a problem in most of theoretical works, it is when one comes to applications. Indeed it is important to always use the same representation of each object.

Let $\tilde{\mathcal{G}}(n)$ be the set of Gaussian laws $\mathcal{N}_{m,\Sigma}$ such that Σ has distinct eigenvalues. Note that the difference between $\tilde{\mathcal{G}}(n)$ and $\mathcal{G}(n)$ has a null measure for standard measures. Let define

$$E = \{\lambda \in \mathbb{R}^n_+ \,|\, \forall 0 \leq i < j \leq n,\ \lambda_i < \lambda_j\},$$

and

$$H = \{D_\lambda \,|\, \lambda \in \{+1,-1\}^n,\ det(D_\lambda) = 1\}.$$

Note that H is a group. The quotient $SO(n)/H = \{RH, R \in SO(n)\}$ can be identified to the fundamental domain $\mathcal{D} \subset SO(n)$ where for $R \in \mathcal{D}$, the maximal element with the smallest index element of each column $1 \leq j < n$ is positive. In other words

$$D = \{R \in SO(n) \,|\, \forall j < n, R_{\min\{\text{argmax}_i\{R_{i,j}\}\},j} > 0\}.$$

Then, parametrization

$$\tilde{\psi} : \left\{ \begin{array}{ccc} \mathbb{R}^n \times E \times \mathcal{D} & \to & \tilde{\mathcal{G}}(n) \\ (m, \lambda, R) & \mapsto & \mathcal{N}_{m,(RD_\lambda R^t)} \end{array} \right.$$

is an homeomorphism.

Let α be the change of parametrization $\alpha = \phi^{-1} \circ \psi$:

$$\alpha : \left\{ \begin{array}{ccc} \mathbb{R}^n \times \mathbb{R}^n_+ \times SO(n) & \to & \mathbb{R}^n \times \text{SPD}(n) \\ (m, \lambda, R) & \mapsto & (m, RD_\lambda R^t) \end{array} \right.$$

On any locally compact group there exists, up to a scaling factor, a unique measure invariant under the group law. Such a measure is called a Haar measure. The Lebesgue measure is a Haar measure. Since the set of rotations is a compact group, there exist a Haar measure $Haar_{SO(n)}$ on rotations.

Let μ_ϕ be the Lebesgue measure on $U_\phi = \mathbb{R}^n \times \text{SPD}(n)$, given by

$$\mu_\phi = Leb_{\mathbb{R}^n} \times Leb_{\text{SPD}(n)}.$$

Let μ_ψ be the product measure between the Lebesgue measure of $\mathbb{R}^n \times \mathbb{R}^n_+$ and the Haar measure of $SO(n)$, i.e.,

$$\mu_\psi = Leb_{\mathbb{R}^n} \times Leb_{\mathbb{R}^n_+} \times Haar_{SO(n)}.$$

We are now interested in the local volume change induced by the application α, that is to say the ratio between the volume of an infinitesimal volume element in $\mathbb{R}^n \times \mathbb{R}^n_+ \times SO(n)$ and its image in $\mathbb{R}^n \times \text{SPD}(n)$. Let us introduce

$$R_0.(m, \lambda, R) = (m, \lambda, R_0 R) \quad \text{and} \quad (m, \lambda, R).R_0 = (m, \lambda, RR_0).$$

Let $A \subset \mathbb{R}^n \times \mathbb{R}^n_+ \times SO(n)$. Since a matrix $R \in SO(n)$ is of determinant 1, one has $\mu_\phi(\alpha(A)) = \mu_\phi(R\alpha(A)R^t)$ for any $R \in SO(n)$. Since $R\alpha(A)R^t = \alpha(R.A)$, $\mu_\phi(\alpha(A)) = \mu_\phi(\alpha(R.A))$. By definition of the Haar measure, $\mu_\psi(A) = \mu_\psi(R.A)$. Thus the ratio between the volumes of infinitesimal volume elements are independent of the rotation.

Let $\mu_{\tilde{\psi}}$ be the restriction of μ_ψ to $\mathbb{R}^n \times E \times \mathcal{D}$. As defined in Section 4.1, $\alpha^*(\mu_{\tilde{\psi}})$ is the transport of $\mu_{\tilde{\psi}}$ on $\mathbb{R}^n \times \text{SPD}(n)$ by α. Thus

$$\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, \Sigma)$$

denotes the density of μ_ϕ with respect to $\alpha^*(\mu_{\tilde{\psi}})$ at (m, Σ). For $\Sigma = RD_\lambda R^t$ with $D_\lambda \in E, R \in \mathcal{D}$, we abuse notation and write $\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, D_\lambda, R)$ for $\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, \Sigma)$. Following the previous discussion, $\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, D_\lambda, R)$ is independent of the rotation, thus

$$\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, D_\lambda, R) = \frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, D_\lambda, I)$$

.

The matrix exponential $e^{(.)}$ defines an homeomorphism from a neighborhood of the null matrix in the space of anti-symmetric matrices and a neighborhood of I in $\mathcal{D} \in SO(n)$. Let $\hat{\mu_\psi}$ be the Lebesgue measure on $\mathbb{R}^n \times \mathbb{R}^n_+ \times ASym$, where $ASym$ is the set of anti-symmetric matrices. The computation of $\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, D, I)$ is made through the following diagram:

$$\begin{array}{ccc} (\mathbb{R}^n \times \mathbb{R}^n_+ \times ASym(n), \hat{\mu_\psi}) & & \\ \Big\downarrow {\scriptstyle Id_{2n} \times e^{(.)}} & \searrow^{\hat{\alpha}} & \\ (\mathbb{R}^n \times \mathbb{R}^n_+ \times \mathcal{D}, \mu_{\tilde{\psi}}) & \xrightarrow[\alpha]{} & (\mathbb{R}^n \times \mathrm{SPD}(n), \mu_\phi) \end{array}$$

where $\tilde{\alpha}(m, \lambda, A) = (m, e^A D_\lambda (e^A)^t)$. The local volume change induced by $\hat{\alpha}$ at $(m, \lambda, 0) \in \mathbb{R}^n \times \mathbb{R}^n_+ \times ASym(n)$ is a product between local volume change of $Id_{2n} \times e^{(.)}$ and α. Thus, knowing the volume change induced by $\hat{\alpha}$ and $Id_{2n} \times e^{(.)}$ at $(m, \lambda, 0)$ is enough to determined the volume change induced by α. We can normalize the Haar measure on rotations such that $Id_{2n} \times e^{(.)}$ does not introduce local at $(m, \lambda, 0)$. Then the volume change induced by $\hat{\alpha}$ and α are equal and given by the Jacobian determinant of $\hat{\alpha}$.

Let $A_{i,j} = \frac{1}{\sqrt{2}}(E_{i,j} - E_{j,i})$. The partial derivatives of $\hat{\alpha}$ at $(m, \lambda, 0)$ in basis

$$\mathcal{B} = \{e_{m_i}\} \cup \{e_{\lambda_i}\} \cup \{A_{i,j}\} \text{ and } \mathcal{B}' = \{e_{m_i}\} \cup \left\{E'_{i,j}\right\},$$

can easily be computed and are given by

$$\frac{\partial \hat{\alpha}}{\partial e_{m_i}}(m, \lambda, 0) = e_{m_i},$$

$$\frac{\partial \hat{\alpha}}{\partial e_{\lambda_i}}(m, \lambda, 0) = E'_{i,i},$$

$$\frac{\partial \hat{\alpha}}{\partial A_{i,j}}(m, \lambda, 0) = (\lambda_j - \lambda_i) E'_{i,j}. \tag{6.8}$$

Since the $E'_{i,j}$ form an orthonormal system, we have:

$$\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, D_\lambda, I) = \prod_{i<j\leq n} |\lambda_i - \lambda_j|, \tag{6.9}$$

and thus using independency of rotation:

$$\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}(m, D_\lambda, R) = \prod_{i<j\leq n} |\lambda_i - \lambda_j|, \tag{6.10}$$

when eigenvalues are distinct. This result can be found in Example 8.7, page 158, of Wijsman (1990).

6.4.1. Separation of mean, eigenvalues, and rotation

This section shows that in many classical cases, measures on multivariate Gaussian laws can be decomposed as a product between a measure on the mean space, a measure on the eigenvalue space, and a measure on the rotation space. Observe that rotations naturally act on $\mathcal{G}(n)$, i.e.,

$$R_0.\mathcal{N}_{m,\lambda,R_1} = \mathcal{N}_{m,\lambda,R_0R_1}.$$

Theorem 6.4.1 *Let μ be a measure on $\tilde{\mathcal{G}}(n)$, invariant under the action of rotations. Measure μ is equivalent to a product measure,*

$$\mu = \mu_{mean} \times \mu_\lambda \times \mu_R,$$

where μ_{mean} is a measure on $\mathbb{R}^n$, μ_λ is a measure on E and μ_R is a measure on $SO(n)$.

Proof This is a direct consequence of Theorem 8.2, page 155, of Wijsman (1990).

In most situations, measures μ on $\mathcal{G}(n)$ present the desired invariances, and $\mu(\mathcal{G}(n) \setminus \tilde{\mathcal{G}}(n)) = 0$. Writing μ as a product measure enables to define densities for the mean, the eigenvalues and the rotations separately, which is not always possible, see Section 6.13. This is in particular the case for the measures induced by the Fisher metric, the Wasserstein metric, and the Lebesgue measure of $\mathbb{R}^n \times \mathrm{SPD}(n)$.

6.4.2. The Fisher metric

Unlike multivariate centered and univariate Gaussian laws, the space of multivariate Gaussian laws under the Fisher metric is not a symmetric space. As in the large majority of the Riemannian manifolds, there is no explicit expression of distances. We only have access to the expression of the metric. An attempt was made by authors of Lovrić et al. (2000) to symmetrize the space.

Let u and v be vectors of $\mathbb{R}^n \bigoplus Sym(n)$. Let index m and Σ denote the components associated to $\mathbb{R}^n$ and $Sym(n)$ respectively.

The scalar product of the Fisher metric at (m, Σ) can be rewritten, see Skovgaard (1984), as

$$< u, v >^{Fisher}_{m,\Sigma} = u_m^t \Sigma^{-1} v_m + \frac{1}{2} tr(\Sigma^{-1} u_\Sigma \Sigma^{-1} v_\Sigma). \tag{6.11}$$

Let $G^{Fisher}_{\mathcal{N}_{m,\Sigma}}$ be the matrix of the scalar product in the canonical basis of $\mathbb{R}^n \bigoplus Sym(n)$. Let vol_{Fisher} be the measure associated with the Riemannian metric. The density of the Fisher volume measure with respect to the Lebesgue measure of $\mathbb{R}^n \times \mathrm{SPD}(n)$ is given by

$$\frac{dvol_{Fisher}}{d\mu_\phi} = \sqrt{det(G^{Fisher}_{\mathcal{N}_{m,\Sigma}})}.$$

The invariance of the Fisher metric under the action of rotations implies the invariance of the measure, i.e.,

$$\sqrt{det(G^{Fisher}_{\mathcal{N}_{m,RDR^t}})} = \sqrt{det(G^{Fisher}_{\mathcal{N}_{m,D}})}.$$

Thus the volume change only need to be computed for the case diagonal matrices. Let $D \in \mathrm{SPD}(n)$ be a diagonal matrix of eigenvalues λ_i. Using Expression (6.11), we can now

calculate the different scalar product between the vectors of the canonical basis. More precisely,

$\forall i, j, k$:

$$< e_{m_i}, E'_{j,k} >^{Fisher}_{\mathcal{N}_{m,D}} = 0,$$

$\forall i, j$:

$$< e_{m_i}, e_{m_j} >^{Fisher}_{\mathcal{N}_{m,D}} = \delta_{i,j}\lambda_i^{-1},$$

where $\delta_{i,j} = 1$ if $i = j$ and $\delta_{i,j} = 0$ if $i \neq j$. In addition,

$\forall i, j$:

$$\langle E'_{i,i}, E'_{j,j} \rangle^{Fisher}_{\mathcal{N}_{m,D}} = \delta_{i,j}\lambda_i^{-2},$$

$\forall (i \neq j, k \neq l)$:

$$\langle E'_{i,j}, E'_{k,l} \rangle^{Fisher}_{\mathcal{N}_{m,D}} = \frac{1}{2}\delta_{i,k}\delta_{j,l}(\lambda_i\lambda_j)^{-1},$$

$\forall j \neq k$:

$$\langle E'_{i,i}, E'_{j,k} \rangle^{Fisher}_{\mathcal{N}_{m,D}} = 0.$$

Thus, one can write

$$G^{Fisher}_{\mathcal{N}_{m,D}} = \begin{array}{c} \begin{array}{ccccccc} & e_{m_i} & & E'_{i,i} & & E'_{i,j} & \end{array} \\ \begin{pmatrix} \frac{1}{\lambda_1} & & & & & & \\ & \ddots & & & & & \\ & & \frac{1}{\lambda_2} & & & & \\ & & & \frac{1}{\lambda_1^2} & & & \\ & & & & \ddots & & \\ & & & & & \frac{1}{\lambda_n^2} & \\ & & & & & & \frac{1}{2\lambda_i\lambda_j} \end{pmatrix} \end{array}$$

and now, from $G^{Fisher}_{\mathcal{N}_{m,D}}$, we have

$$\frac{dvol_{Fisher}}{d\mu_\phi} = \sqrt{det(G^{Fisher}_{\mathcal{N}_{m,\Sigma}})} = \frac{1}{2^{\frac{n(n-1)}{2}}}\sqrt{\prod_i \frac{1}{\lambda_i^{n+2}}} = \frac{1}{2^{\frac{n(n-1)}{2}}} det(\Sigma)^{-(n+2)/2}. \quad (6.12)$$

As announced in section 6.4.1, the measure induced by the Fisher metric is a product between a measure on the mean, a measure on eigenvalues and a measure on rotations, i.e.,

$$\begin{aligned} dvol_{Fisher} &= \sqrt{det(G^{Fisher}_{\mathcal{N}_{m,\Sigma}})}\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}d\mu_{\tilde{\psi}} \\ &= \frac{1}{2^{\frac{n(n-1)}{2}}}\prod_{i<j\leq n}|\lambda_i - \lambda_j|\sqrt{\prod_i \frac{1}{\lambda_i^{n+2}}}\, dLeb_{\mathbb{R}^n}\, dLeb_{\mathbb{R}^n_+}\, dHaar_{SO(n)}. \end{aligned} \quad (6.13)$$

6.4.3. The Wasserstein metric

The 2-Wasserstein metric is a product metric between the space of means and the space of covariance matrices:

$$d_{W_2}(\mathcal{N}_{m_1,\Sigma_1},\mathcal{N}_{m_2,\Sigma_2})^2 = d_m(m_1,m_2)^2 + d_\Sigma(\Sigma_1,\Sigma_2)^2,$$

with

$$d_m(m_1,m_2) = |m_1 - m_2|,$$

$$d_\Sigma(\Sigma_1,\Sigma_2) = tr(\Sigma_1) + tr(\Sigma_2) - 2tr(\sqrt{\Sigma_1^{1/2}\Sigma_2\Sigma_1^{1/2}}).$$

Takatsu (2011) proved that this distance is induced by a Riemannian metric. Let vol_W be the measure associated with such Riemannian metric.

6.4.3.1. The Riemannian volume element

The calculation of the volume measure is similar to the calculation of the volume measure associated with the Fisher metric. Let u and v be vectors of $\mathbb{R}^n \bigoplus Sym(n)$. Let index m and Σ denote the components associated to $\mathbb{R}^n$ and $Sym(n)$ respectively. The scalar product at (m,Σ) according to Takatsu (2011) is given by

$$< u, v >^W_{m,\Sigma} = u_m^t v_m + tr(\tilde{u}_\Sigma \Sigma \tilde{v}_\Sigma),$$

with

$$\begin{aligned} u_\Sigma &= \Sigma\tilde{u}_\Sigma + \tilde{u}_\Sigma\Sigma, \\ v_\Sigma &= \Sigma\tilde{v}_\Sigma + \tilde{v}_\Sigma\Sigma. \end{aligned} \tag{6.14}$$

Let $G^W_{\mathcal{N}_{m,\Sigma}}$ be the matrix of the scalar product in the canonical basis of $\mathbb{R}^n \bigoplus Sym(n)$. The density of the Wasserstein volume measure with respect to the Lebesgue measure of $\mathbb{R}^n \times \mathrm{SPD}(n)$ is given by

$$\frac{dvol_W}{d\mu_\phi} = \sqrt{det(G^W_{\mathcal{N}_{m,\Sigma}})}.$$

The invariance of the Wasserstein metric under the action of rotations implies the invariance of the measure, i.e.,

$$\sqrt{det(G^W_{\mathcal{N}_{m,RDR^t}})} = \sqrt{det(G^W_{\mathcal{N}_{m,D}})}.$$

Thus the volume change only need to be computed for the case diagonal matrices. Let $D \in \mathrm{SPD}(n)$ be a diagonal matrix of eigenvalues λ_i. Equations. (6.14) can now be solved. Since,

$$\begin{aligned} 2\lambda_i E'_{i,i} &= DE'_{i,i} + E'_{i,i}D, \\ (\lambda_i + \lambda_j)E'_{i,j} &= DE'_{i,j} + E'_{i,j}D, \end{aligned}$$

equations (6.14) are equivalent to

$$\begin{aligned} u_\Sigma &= P(\tilde{u}_\Sigma), \\ v_\Sigma &= P(\tilde{v}_\Sigma), \end{aligned}$$

where P is a linear operator of $Sym(n)$ whose matrix form is

$$P: \begin{array}{c} \\ E'_{i,i} \\ \\ \\ E'_{i,j} \end{array} \begin{pmatrix} 2\lambda_1 & & & & \\ & \ddots & & & \\ & & 2\lambda_n & & \\ & & & (\lambda_i+\lambda_j) & \end{pmatrix}$$

The different terms of the metric are given by:

$\forall(i,j \leq k)$:

$$\langle e_{m_i}, E'_{j,k}\rangle^W_{\mathcal{N}_{m,D}} = 0,$$

$\forall(i,j)$:

$$\langle e_{m_i}, e_{m_j}\rangle^W_{\mathcal{N}_{m,D}} = 1,$$

$\forall(i,j)$:

$$\langle E'_{i,i}, E'_{j,j}\rangle^W_{\mathcal{N}_{m,D}} = \delta_{i,j}\frac{1}{4\lambda_i},$$

where again $\delta_{i,j} = 1$ if $i = j$ and $\delta_{i,j} = 0$ if $i \neq j$,

$\forall(i,j \neq k)$:

$$\langle E'_{i,i}, E'_{j,k}\rangle^W_{\mathcal{N}_{m,D}} = 0,$$

$\forall i < j, k < l$:

$$\langle E'_{i,j}, E'_{k,l}\rangle^W_{\mathcal{N}_{m,D}} = \delta_{i,k}\delta_{j,l}\frac{1}{2(\lambda_i+\lambda_j)}.$$

Thus, we obtain

$$G^W_{\mathcal{N}_{m,D}} = \begin{array}{c} \\ e_{m_i} \\ \\ \\ E'_{i,i} \\ \\ \\ E'_{i,j} \end{array} \begin{pmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & 1 & & & & \\ & & & \frac{1}{4\lambda_1} & & & \\ & & & & \ddots & & \\ & & & & & \frac{1}{4\lambda_n} & \\ & & & & & & \frac{1}{2(\lambda_i+\lambda_j)} \end{pmatrix}$$

As for the Fisher measure, we compute

$$\frac{dvol_W}{d\mu_\phi} = \sqrt{det(G^W_{\mathcal{N}_{m,\Sigma}})} = \frac{1}{2^{\frac{1}{2}n^2+\frac{3}{2}n}}\sqrt{\prod_i \frac{1}{\lambda_i}\prod_{i<j}\frac{1}{(\lambda_i+\lambda_j)}}, \tag{6.15}$$

and the measure induced by Wasserstein metric is written as

$$\begin{aligned} dvol_W &= \sqrt{det(G^W_{\mathcal{N}_{m,\Sigma}})}\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})}d\mu_\psi \\ &= \frac{1}{2^{\frac{1}{2}n^2+\frac{3}{2}n}}\prod_i \frac{1}{\sqrt{\lambda_i}}\prod_{i<j}\frac{|\lambda_i-\lambda_j|}{\sqrt{(\lambda_i+\lambda_j)}}\, dLeb_{\mathbb{R}^n}\, dLeb_{\mathbb{R}^n_+}\, dHaar_{SO(n)}. \end{aligned} \tag{6.16}$$

6.4.3.2. The volume change of the exponential map

We are interested here in the computation of the following quantity:

$$\theta_{\Sigma_1} : \Sigma_2 \mapsto \theta_{\Sigma_1}(\Sigma_2) = \frac{dvol_W}{dexp^*(Leb_{\Sigma_1})}(\Sigma_2),$$

that is to say the density of the Riemannian measure of $\mathcal{G}(n)$ with respect to the Lebesgue measure of the parametrization of $\mathcal{G}(n)$ induced by exp_{Σ_1}. As we have stated, Wasserstein metric is a product between the Euclidean metric on $\mathbb{R}^n$ and a Riemannian metric on $\mathrm{SPD}(n)$. The Euclidean part can be omitted in this section, since it does not affect the volume change.

It has been shown in Takatsu (2011) that the L^2-Wasserstein distance on centered Gaussian measure is induced by a Riemannian metric on $\mathrm{SPD}(n)$. Furthermore, the application

$$\Pi : \left\{ \begin{array}{ccc} Gl_n & \to & \mathrm{SPD}(n) \\ G & \mapsto & GG^t \end{array} \right.$$

is a Riemannian submersion when GL_n is endowed with the scalar product $\langle A, B\rangle = tr(AB^t)$ and $\mathrm{SPD}(n)$ is embedded with the Riemannian metric. For $A \in Gl(n)$ the kernel of $d\Pi$ is called the vertical space, and its orthogonal the horizontal space. For $\Sigma = RDR^t \in \mathrm{SPD}(n)$, let $\Sigma^{1/2} = RD_{(\sqrt{\lambda_i})}R^t$, such that $\Pi(\Sigma^{1/2}) = \Sigma$. The horizontal space at $\Sigma^{1/2}$, see Takatsu (2011), is given by:

$$\mathcal{H}_{\Sigma^{1/2}} = \{X\Sigma^{1/2}, X \in Sym\}.$$

$\mathcal{H}_{\Sigma^{1/2}}$ can be identified with $T_\Sigma M$ via $d\Pi$. Let $\Pi_A(Z) = \Pi(A + Z)$. For simplicity reasons, the exponential map will be expressed in $\mathcal{H}_{\Sigma^{1/2}}$. For $Z \in \mathcal{H}_{\Sigma^{1/2}}$ classical results on Riemannian submersions give that the image by Π of the geodesic $\Sigma^{1/2} + tZ$ is the geodesic $exp_\Sigma(tZ)$. Hence, one has

$$exp_\Sigma(Z) = \Pi_{\Sigma^{1/2}}(Z).$$

Fig. 6.2 sums up the situation. Let $\tilde{\Pi}_A$ be the restriction of Π_A to $\mathcal{H}_A$. The density of interest is a product of two factors, i.e.,

$$\frac{dvol_W}{dexp^*(Leb_{\mathcal{H}_{\Sigma^{1/2}}})} = \frac{dvol_W}{dLeb_{\mathrm{SPD}(n)}} \frac{dLeb_{\mathrm{SPD}(n)}}{dexp^*(Leb_{\mathcal{H}_{\Sigma^{1/2}}})},$$

where $\frac{dvol_W}{dLeb_{\mathrm{SPD}(n)}}$ was computed in Eq. (6.15). The computation of $\frac{dLeb_{\mathrm{SPD}(n)}}{dexp^*(Leb_{\mathcal{H}_{\Sigma^{1/2}}})}$ can be made through the Jacobian of the exponential application. The computation of the differential of $\tilde{\Pi}_A$ gives:

$$\begin{aligned} d\tilde{\Pi}_A(XA, X \in Sym(n))(YA, Y \in Sym(n)) &= (A + XA)(YA)^t + YA(A + XA)^t \\ &= AA^tY + XAA^tY + YAA^t + YAA^tX \\ &= AA^tY + YAA^t + XAA^tY + YAA^tX. \end{aligned}$$

The invariance of the Wasserstein metric under the action of rotations enables to restrict the computation of θ_{Σ_1} to θ_D with D a diagonal matrix. So,

$$\begin{aligned} d\tilde{\Pi}_{D^{1/2}}(XD^{1/2})(YD^{1/2}) &= DY + YD + XDY + (XDY)^t \\ &= (I + X)DY + Y((I + X)D)^t. \end{aligned}$$

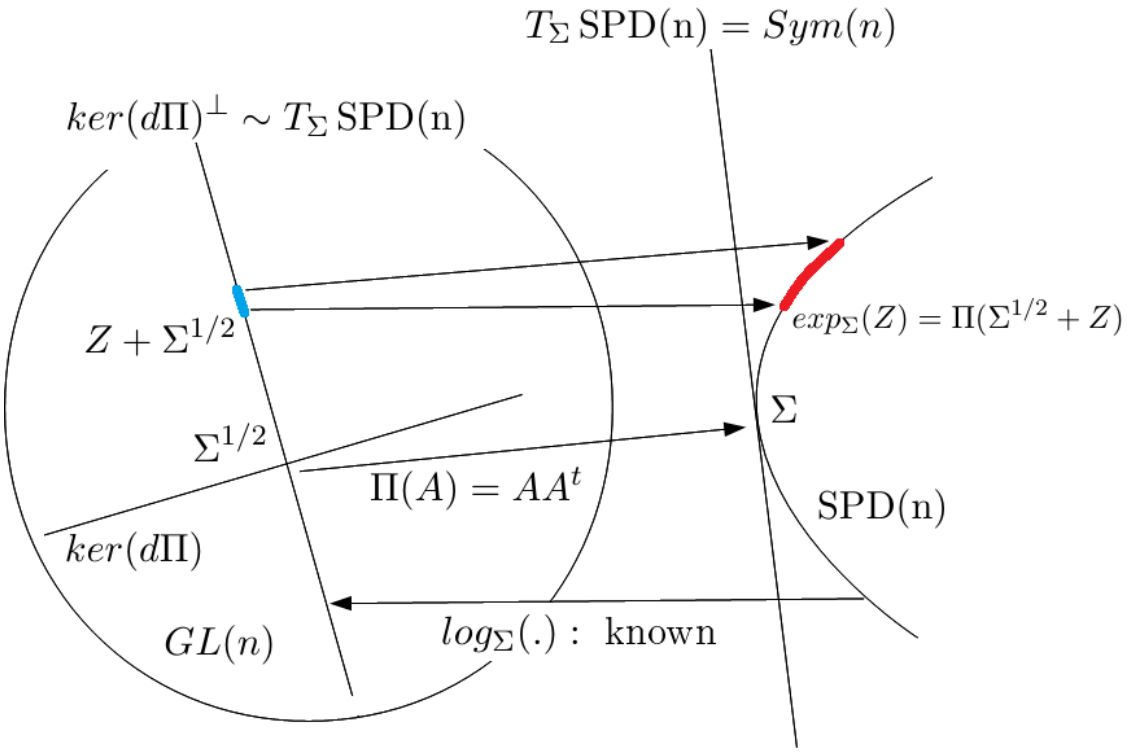

Figure 6.2.: The application Π is a Riemannian submersion for the flat metric on $Gl(n)$ and the Wasserstien metric on $\mathrm{SPD}(n)$. The volume change between the blue and the red area appears in the expression of kernels.

The following result enables the computation of the volume change of the application $Y \mapsto d\tilde{\Pi}_{D^{1/2}}(XD^{1/2})(YD^{1/2})$.

Theorem 6.4.2 *Let A be a square matrix of order n and*

$$S_A : \left\{ \begin{array}{ccc} Sym(n) & \to & Sym(n) \\ Y & \mapsto & YA + A^tY \end{array} \right.$$

Then

$$det(S_A) = \prod_{i \leqslant j} (\lambda_i + \lambda_j) = 2^n det(A) \prod_{i<j} (\lambda_i + \lambda_j),$$

where the λ_i are the complex eigenvalues of A. Furthermore, $det(S_A)$ depends only on the coefficients of the characteristic polynomial of A

$$\chi_A(\lambda) = det(\lambda I - A) = \prod_{i=1}^{n} (\lambda - \lambda_i) = \lambda^n + \prod_{k=1}^{n} (-1)^k \sigma_{n-k} \lambda^k.$$

1. *If $n = 2$, $det(S_A) = 4(detA)(traceA)$;*

2. *If* $n = 3$, $det(S_A) = 2^3 detA(\sigma_1\sigma_2 - \sigma_3)$;
3. *If* $n = 4$, $det(S_A) = 2^4 detA(\sigma_1\sigma_2 - \sigma_3 - \sigma_4\sigma_1^2 - \sigma_3^2)$;
4. *If* $n \geq 4$, $det(S_A)$ *can be obtained with the help of standard libraries of mathematical programming languages using the decomposition on the elementary symmetric polynomials.*

Proof 1. Suppose A is diagonalisable and $(X_i)_{1\leq i\leq n}$ is a basis of eigenvectors of A. Then

$$S_A(X_iX_j^t + X_jX_i^t) = (\lambda_i + \lambda_j)(X_iX_j^t + X_jX_i^t),$$

and $(X_iX_j^t + X_jX_i^t)_{1\leqslant i\leqslant j\leqslant n}$ is a basis of $Sym(n)$. Consequently

$$det(S_A) = \prod_{i\leqslant j}(\lambda_i + \lambda_j) = 2^n det(A)\prod_{i<j}(\lambda_i + \lambda_j)$$

2. $det(S_A)$ is a symmetric polynomial of the λ_i. The fundamental theorem of symmetric polynomial states that any symmetric polynomial can by expressed as a polynomial function of the elementary symmetric polynomials. Since the coefficients of the characteristic polynomial of A are the elementary symmetric polynomials in λ_i, $det(S_A)$ is a polynomial function of the coefficient of the characteristic polynomial of A.
3. In the general case, any matrix A can be seen as a limit of diagonalizable matrices A_n. By continuity we have $det(S_A) = \lim det(S_{A_n})$ and the coefficients of the polynomial χ_A are the limits of the coefficients of the polynomials χ_{A_n}. Therefore the formula can be extended to any matrix.

The basis $\mathcal{B} = (E'_{i\leq j}D^{1/2})$ of $\mathcal{H}_{D^{1/2}}$ is orthogonal for the scalar product $\langle A, B\rangle = tr(AB^t)$, with $\left\langle E'_{ii}D^{1/2}, E'_{ii}D^{1/2}\right\rangle = \lambda_i$ and $\left\langle E'_{i<j}D^{1/2}, E'_{i<j}D^{1/2}\right\rangle = \frac{\lambda_i+\lambda_j}{2}$. Thus in an orthonormal basis of $\mathcal{H}_{D^{1/2}}$, we have

$$det(\mathcal{B}) = \sqrt{\prod\lambda_i\prod_{i<j}\frac{\lambda_i + \lambda_j}{2}}.$$

Let consider the map

$$f : \left\{\begin{array}{ccc} \mathcal{H}_{D^{1/2}} & \to & Sym(n) \\ YD^{1/2} & \mapsto & Y \end{array}\right.$$

Expressed in orthonormal basis, $det(f) = det(\mathcal{B})^{-1}$. Since $d\tilde{\Pi}_{D^{1/2}}(XD^{1/2}) = S_{(I+X)D}\circ f$, the determinant of $d\tilde{\Pi}_{D^{1/2}}$ expressed in orthonormal basis is given by

$$det(d\tilde{\Pi}_{D^{1/2}}) = det(S_{(I+X)D})det(f) = \frac{det(S_{(I+X)D})}{det(\mathcal{B})}.$$

Hence the expression of the volume change is given by,

$$\frac{dLeb_{\mathrm{SPD}(n)}}{dexp^*(Leb_{\mathcal{H}_{\Sigma^{1/2}}})} = det(d\tilde{\Pi}_{D^{1/2}}) = \frac{det(S_{(I+X)D})}{det(\mathcal{B})}. \tag{6.17}$$

Given two matrices $\Sigma_0, \Sigma_1 \in \text{SPD}(n)$, Takatsu (2011) provides the expression of the vector $Z \in \mathcal{H}_{\Sigma_0^{1/2}}$:

$$Z = (\Sigma_1^{1/2}(\Sigma_1^{1/2}\Sigma_0\Sigma_1^{1/2})^{-1/2}\Sigma_1^{1/2} - I)\Sigma_0^{1/2},$$

such that $\Pi_{\Sigma_0^{1/2}}(Z) = \Sigma_1$. Thus,

$$\begin{aligned}\theta_D(\Sigma) &= \frac{dvol_W}{dLeb_{\text{SPD}(n)}} \frac{dLeb_{\text{SPD}(n)}}{dexp^*(Leb_{\mathcal{H}_{D^{1/2}}})} \\ &= \frac{1}{2^{n(n+1)}} \frac{det(S_A)}{\sqrt{\prod \lambda_i^\Sigma \prod_{i<j}(\lambda_i^\Sigma + \lambda_j^\Sigma)}\sqrt{\prod \lambda_i^D \prod_{i<j}(\lambda_i^D + \lambda_j^D)}},\end{aligned}$$

where (λ_i^D) and (λ_i^Σ) are the eigenvalues of D and Σ respectively, and where

$$A = \Sigma^{1/2}(\Sigma^{1/2}D\Sigma^{1/2})^{-1/2}\Sigma^{1/2}D^{1/2}.$$

We recall that A corresponds to the matrix $(I + X)D$ of Eq. (6.17).

Finally, using the invariance of the Wasserstein metric under the action of rotations, we have for $\Sigma_0 = RDR^t$:

$$\theta_{\Sigma_0}(\Sigma_1) = \frac{1}{2^{n(n+1)}} \frac{det(S_A)}{\sqrt{\prod \lambda_i^{\Sigma_1} \prod_{i<j}(\lambda_i^{\Sigma_1} + \lambda_j^{\Sigma_1})}\sqrt{\prod \lambda_i^{\Sigma_0} \prod_{i<j}(\lambda_i^{\Sigma_0} + \lambda_j^{\Sigma_0})}} \tag{6.18}$$

where $(\lambda_i^{\Sigma_0})$ and $(\lambda_i^{\Sigma_1})$ are the eigenvalues of Σ_0 and $R^t\Sigma_1 R$ respectively,

$$A = (R^t\Sigma_1 R)^{1/2}((R^t\Sigma_1 R)^{1/2}D(R^t\Sigma_1 R)^{1/2})^{-1/2}(R^t\Sigma_1 R)^{1/2}D^{1/2},$$

and $det(S_A)$ is given in Theorem 6.4.2.

6.5. Density estimation on multivariate Gaussian laws

Let Ω be a space endowed of with a σ-algebra and a probability measure p. Let X be a random variable $\Omega \to \mathcal{G}(n)$. We assume that the support of X is a compact set, noted *Supp*. The space of Gaussian laws $\mathcal{G}(n)$ is equipped successively with the Fisher and the Wasserstein metric. The Riemannian measure is called *vol* and the measure on $\mathcal{G}(n)$ induced by X is noted μ_X. We assume that μ_X has a density, noted f, with respect to *vol*, and that the support of X is a compact set noted *Supp*. Let $(x_1, .., x_k) \in \mathcal{M}^k$ be a set of draw of X and $\mu_k = \frac{1}{k}\sum_i \delta_{x_i}$ denotes the empirical measure of the set of draws. To each point $x \in \mathcal{G}(n)$ quantities m_x and Σ_x representing the associated mean and covariance are associated.

6.5.1. The Fisher metric

The expression of distances is unknown, there are no known interesting tilings, and there is no expression of the eigenfunctions of the Laplace operator. Thus there is no way of estimating densities based on a fully Riemannian approach. The only option here is to

perform a density estimation in an Euclidean context, and to adapt the result to the Riemannian measure. If X is seen as a random variable in $\mathrm{SPD}(n)$ the estimation gives:

$$\hat{f}^{Fisher} = \hat{f}^{Eucl}\frac{d\mu_\phi}{dvol} = 2^{\frac{n(n-1)}{2}}\sqrt{\prod_i \lambda_i^{n+2}}\hat{f}^{Eucl}, \tag{6.19}$$

where $\hat{f}^{Eucl}$ is an estimation of the density of X with respect to the measure induced by the scalar product $\langle A, B\rangle = tr(AB^t)$. However the estimation does not respect homogeneity and isotropy considerations.

6.5.2. The Wasserstein metric

As with the Fisher metric, there are no known interesting tiling, and there is no expression of the eigenfunctions of the Laplace operator. However it is possible to obtain expressions of kernels. Given a scaling parameter $r > 0$, the estimator becomes:

$$\hat{f}_k^r(x) = \frac{1}{k}\sum_i \frac{1}{r^n}\frac{1}{\theta_{\Sigma_{x_i}}(\Sigma_x)}K\left(\frac{d_{W_2}(x, x_i)}{r}\right), \tag{6.20}$$

where $\theta_{\Sigma_{x_i}}(\Sigma_x)$ was computed in Eq. (6.18).

6.6. The space of multivariate centred Gaussian laws

We consider specifically the study of multivariate Gaussian laws, after adding the property $m = 0$. The set of multivariate centered Gaussian laws is noted $\mathcal{G}_{m=0}(n)$. In $\mathcal{G}_{m=0}(n)$, a Gaussian laws is determined by its covariance matrix Σ only. It elements are noted $\mathcal{N}_\Sigma$. The expression of its density f_Σ is:

$$f_{m,\Sigma}(x) = \frac{1}{(2\pi)^{N/2}det(\Sigma)^{1/2}}e^{\frac{-1}{2}(x)^t\Sigma^{-1}(x)}.$$

Everything written in section 6.4 remains valid after removal of the mean. The measures μ_ϕ and μ_ψ become respectively:

$$\mu_\phi = Leb_{\mathrm{SPD}(n)},$$

$$\mu_\psi = Leb_{\mathbb{R}_+^n} \times Haar_{SO(n)},$$

and $\mathcal{G}_{m=0}(n)$ is identified with $\mathrm{SPD}(n)$, such that $\mathcal{N}_\Sigma$ and Σ are used indifferently.

The different expressions associated to the Wasserstein metric computed in section 6.4.3 do not change for the case of centered Gaussian laws. Furthermore, space $\mathcal{G}(n)$ does not gain any property after imposing $m = 0$ in the Wasserstein metric. Thus this section focuses on the Fisher metric.

The Fisher metric on $\mathcal{G}_{m=0}(n)$ happens to be the metric induced by the Fisher metric on $\mathcal{G}(n)$. On the contrary of $\mathcal{G}(n)$, $\mathcal{G}_{m=0}(n)$ is symmetric and is then easier to study. Let $GL(n)$ be the set of invertible matrices of size n. For any G in $GL(n)$, the application $\Sigma \mapsto G\Sigma G^t$ is an isometry of $\mathcal{G}_{m=0}(n)$. This application corresponds to a coordinate change for covariance matrices. The action of $GL(n)$ on $\mathcal{G}_{m=0}(n)$ induced by this application is

transitive. Indeed, for any $\Sigma \in \mathrm{SPD}(n)$, let $\Sigma^{1/2} \in GL(n)$. The orbit of the identity I is $\mathrm{SPD}(n)$: $\Sigma = \Sigma^{1/2} I \Sigma^{1/2}$. The stabilizator of I are the orthogonal matrices. Thus $\mathcal{G}_{m=0}(n)$ can be identified to $GL(n)/O(n)$.

There is an explicit expression of the distance between two laws:

$$d_{Fisher}(\mathcal{N}_{\Sigma_1}, \mathcal{N}_{\Sigma_2}) = ||log(\Sigma_1^{-1/2} \Sigma_2 \Sigma_1^{-1/2})||,$$

where $||.||$ is the norm associated with the scalar product $tr(AB^t)$.

6.6.1. The Riemannian volume element

By adapting the calculations done in section 6.4.2, we obtain that the metric $G^{Fisher}_{\mathcal{N}_D}$ expressed in the canonical basis of $\mathrm{SPD}(n)$ is given by

$$G^{Fisher}_{\mathcal{N}_D} = \begin{array}{c} \begin{matrix} \quad E'_{i,i} & \quad E'_{i,j} \end{matrix} \\ \begin{pmatrix} \frac{1}{\lambda_1^2} & & & \\ & \ddots & & \\ & & \frac{1}{\lambda_n^2} & \\ & & & \frac{1}{2\lambda_i\lambda_j} \end{pmatrix} \end{array}$$

Therefore, using the classical machinery,

$$\frac{dvol_{Fisher}}{d\mu_\phi} = \sqrt{det(G^{Fisher}_{\mathcal{N}_\Sigma})} = \frac{1}{2^{\frac{n(n-1)}{2}}} \sqrt{\prod_i \frac{1}{\lambda_i^{n+1}}} = \frac{1}{2^{\frac{n(n-1)}{2}}} det(\Sigma)^{-(n+1)/2}. \quad (6.21)$$

This expression can be found in Terras (1988).

Furthermore, the Riemannian volume element is now

$$\begin{aligned} dvol_{Fisher} &= \sqrt{det(G^{Fisher}_{\mathcal{N}_\Sigma})} \frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})} d\mu_{\tilde{\psi}} \\ &= \frac{1}{2^{\frac{n(n-1)}{2}}} \prod_{i<j\leq n} |\lambda_i - \lambda_j| \sqrt{\prod_i \frac{1}{\lambda_i^{n+1}}}\, dLeb_{\mathbb{R}^n_+}\, dHaar_{SO(n)}. \end{aligned} \quad (6.22)$$

6.6.2. The volume change of the exponential map

We are interested here in the computation of the following quantity:

$$\theta_{\Sigma_1} : \Sigma_2 \mapsto \theta_{\Sigma_1}(\Sigma_2) = \frac{dvol_{Fisher}}{dexp^*_{\Sigma_1}(Leb_{Sym(n)})}(\Sigma_2),$$

that is to say the density of the Riemannian measure of $\mathcal{G}_{m=0}(n)$ with respect to the Lebesgue measure of the parametrization induced by exp_Σ. Since $\mathrm{SPD}(n)$ is homogeneous, the volume change can be computed at the identity matrix I. For the affine invariant metric we have that exp_I is the classical exponential application on matrices, see Lenglet et al. (2006). The density is split into two terms:

$$\frac{dvol_{Fisher}}{dexp^*(Leb_{Sym})} = \frac{dvol_{Fisher}}{dLeb_{\mathrm{SPD}(n)}} \frac{dLeb_{\mathrm{SPD}(n)}}{dexp^*(Leb_{Sym(n)})},$$

where $\frac{dvol_{Fisher}}{dLeb_{\mathrm{SPD}(n)}}$ was already computed in Eq. (6.21). Recall the function defined in section 6.4:

$$\alpha : \begin{cases} \mathbb{R}^n \times SO(n) \to Sym(n) \\ (\lambda, R) \mapsto RD_\lambda R^t \end{cases}$$

Let $e\tilde{x}p$ be defined as:

$$e\tilde{x}p : \begin{cases} \mathbb{R}^n \times SO(n) \to \mathbb{R}^n_+ \times SO(n) \\ (\lambda, R) \mapsto \left(e^{\lambda_i}, R\right) \end{cases}$$

The computation of $\frac{dLeb_{\mathrm{SPD}(n)}}{dexp^*(Leb_{Sym(n)})}$ is made using the following commutation diagram,

$$\begin{array}{ccc} (\mathbb{R}^n \times Q, \mu_{\tilde{\psi}} = Leb(\mathbb{R}^n) \times Harr_{SO(n)}) & \xrightarrow{e\tilde{x}p} & (\mathbb{R}^n \times Q, \mu_{\tilde{\psi}}) \\ \Big\downarrow \alpha & & \Big\downarrow \alpha \\ (Sym(n), \mu_\phi = Leb(Sym)) & \xrightarrow[exp]{} & (\mathrm{SPD}(n), \mu_\phi) \end{array}$$

that is to say $exp \circ \alpha = \alpha \circ e\tilde{x}p$. For the sake of simplicity, μ_ϕ is here extended to $Sym(n)$. Knowing the volume changes induced by α and $e\tilde{x}p$ enables the computation of the volume change induced by exp. The local volume change induced by α was computed in section 6.4 and gives:

$$\frac{d\mu_\phi}{d\alpha^*(\mu_{\tilde{\psi}})} = \prod_{i<j} |\lambda_j - \lambda_i|.$$

We address now the volume change $\frac{dexp^*(\mu_{\tilde{\psi}})}{d\mu_{\tilde{\psi}}}$ induced by $e\tilde{x}p$, where $exp^*(\mu_{\tilde{\psi}})$ denotes the image measure of $\mu_{\tilde{\psi}}$ by $e\tilde{x}p$.

Let $F : \mathbb{R}^n \times Q \to \mathbb{R}, F(\lambda, R) = f(\lambda)g(R)$ such that

$$\begin{aligned} \int_{\mathbb{R}^n \times Q} F(\lambda, R) dexp^*(\mu_{\tilde{\psi}}) &= \int_{\mathbb{R}^n \times Q} F(e\tilde{x}p(\lambda, R)) d\mu_{\tilde{\psi}} \\ &= \int_{\mathbb{R}^n \times Q} F(e^\lambda, R) d\mu_{\tilde{\psi}} \\ &= \int_{\mathbb{R}^n} f(e^\lambda) d\lambda \int_Q g(R) dHaar_{SO(n)} \\ &= \int_{\mathbb{R}^n} f(\Lambda) \prod_{i=1}^n \frac{1}{\Lambda_i} d\Lambda \int_Q g(R) dHaar_{SO(n)} \\ &= \int_{\mathbb{R}^n \times Q} F(\Lambda, R) \prod_{i=1}^n \frac{1}{\Lambda_i} d\mu_{\tilde{\psi}}. \end{aligned}$$

Since the set of separable functions $F \in L^1$ is dense in L^1, we have that:

$$\frac{dexp^*(\mu_{\tilde{\psi}})}{d\mu_{\tilde{\psi}}}(\lambda, R) = \prod \frac{1}{\lambda_i} \quad \text{and} \quad \frac{d\mu_{\tilde{\psi}}}{dexp^*(\mu_{\tilde{\psi}})}(\lambda, R) = \prod \lambda_i.$$

Using

$$\frac{d\mu_\phi}{dexp^*(\mu_\phi)}(RD_{e^\lambda}R^t) \frac{d\mu_\phi}{d\mu_{\tilde{\psi}}}(\lambda, R) = \frac{d\mu_\phi}{d\mu_{\tilde{\psi}}}(e^\lambda, R) \frac{d\mu_{\tilde{\psi}}}{dexp^*(\mu_{\tilde{\psi}})}(e^\lambda, R),$$

we have that

$$\frac{d\mu_\phi}{dexp^*(\mu_\phi)}(RD_{e^\lambda}R^t)\prod_{i<j}|\lambda_i-\lambda_j| = \prod_{i<j}|e^{\lambda_i}-e^{\lambda_j}|e^{\sum_i \lambda_i},$$

and when eigenvalues are distinct,

$$\frac{d\mu_\phi}{dexp^*(\mu_\phi)}(RD_{e^\lambda}R^t) = \frac{dLeb_{\mathrm{SPD}(n)}}{dexp^*(Leb_{Sym(n)})}(RD_{e^\lambda}R^t) = \prod_{i<j}\frac{|e^{\lambda_i}-e^{\lambda_j}|}{|\lambda_i-\lambda_j|}e^{\sum_i \lambda_i}. \quad (6.23)$$

The continuity of the volume change enables to extend the expression everywhere. Let $log(\Sigma) = RD_{(log(\lambda_i))}R^t$, where again $\Sigma = RDR^t \in \mathrm{SPD}(n)$. The application log is the inverse the exponential application on $\mathrm{SPD}(n)$. Combining Eq. (6.23) and Eq. (6.21), we obtain for $\Sigma \in \mathrm{SPD}(n)$:

$$\begin{aligned}
\theta_I(\Sigma) &= \frac{1}{2^{n(n-1)/2}}\sqrt{\frac{1}{\prod e^{(n+1)\lambda_i}}}\prod\frac{|e^{\lambda_j}-e^{\lambda_i}|}{|\lambda_j-\lambda_i|}\prod e^{\lambda_i} \\
&= \frac{1}{2^{n(n-1)/2}}\prod\frac{|e^{\lambda_j}-e^{\lambda_i}|}{|\lambda_j-\lambda_i|}\prod e^{(1-\frac{n+1}{2})\lambda_i} \\
&= \frac{1}{2^{n(n-1)/2}}\prod\frac{1}{|\lambda_i-\lambda_j|}\prod 2e^{\frac{\lambda_i+\lambda_j}{2}}\frac{|e^{\frac{\lambda_i-\lambda_j}{2}}-e^{\frac{\lambda_j-\lambda_i}{2}}|}{2}\prod e^{(1-\frac{n+1}{2})\lambda_i} \\
&= \frac{2^{n(n-1)/2}}{2^{n(n-1)/2}}\prod\frac{1}{|\lambda_i-\lambda_j|}\prod sinh\left(\frac{|\lambda_i-\lambda_j|}{2}\right)\prod e^{\frac{\lambda_i+\lambda_j}{2}}\prod e^{(1-\frac{n+1}{2})\lambda_i} \\
&= \prod\frac{1}{|\lambda_i-\lambda_j|}\prod sinh\left(\frac{|\lambda_i-\lambda_j|}{2}\right)\prod e^{\frac{(n-1)}{2}\lambda_i}\prod e^{(1-\frac{n+1}{2})\lambda_i} \\
&= \prod\frac{sinh\left(\frac{|\lambda_i-\lambda_j|}{2}\right)}{|\lambda_i-\lambda_j|},
\end{aligned}$$

where the λ_i are the eigenvalues of $log(\Sigma)$. This expression is not new since the factor in $sinh(.)$ is known to be the Riemannian volume element in polar coordinates. Let $\Sigma^{1/2} = RD_{(\sqrt{\lambda_i})}R^t$. Using the invariance of the metric we obtain finally,

$$\theta_{\Sigma_1}(\Sigma_2) = 2^{n(n-1)/2}\prod\frac{sinh\left(\frac{|\lambda_i-\lambda_j|}{2}\right)}{|\lambda_i-\lambda_j|}, \quad (6.24)$$

where the λ_i are the eigenvalues of $log(\Sigma_1^{1/2}\Sigma_2\Sigma_1^{1/2})$.

6.6.3. Fourier–Helgason transform

The $\mathcal{G}_{m=0}(n)$ is symmetric and there is a Fourier–Helgason transform defined on it. This paragraph is mainly extracted from Peter T. Kim and Donald St. P. Richards (2011). Let

$$p_s : \left\{\begin{array}{ccc} \mathrm{SPD}(n) & \to & \mathbb{C} \\ \Sigma & \mapsto & \prod_i |\Sigma|_i^{s_i} \end{array}\right. \quad (6.25)$$

where $|\Sigma|_i$ is the determinant of Σ after removing row and column i.

For $s \in \mathbb{C}^n$ and $K \in SO(n)$, the Helgason functions (plane waves) have the following expression:

$$e_{s,R} : \left\{ \begin{array}{ccc} \mathrm{SPD}(n) & \to & \mathbb{C} \\ \Sigma & \mapsto & p_s(R\Sigma R^t) \end{array} \right. \tag{6.26}$$

Let C_c^∞ be the set of infinitely differentiable functions with compact support. The Fourier–Helgason transform of $f \in C_c^\infty$ is given by:

$$\mathcal{H}(f)(s,R) = \int_{\Sigma \in \mathrm{SPD}(n)} f(\Sigma)\overline{e_{s,R}(\Sigma)} dvol_{Fisher}, \tag{6.27}$$

where $\overline{x}$ denotes the complex conjugate of x.

Let $\rho = (\frac{1}{2}, \cdots, \frac{1}{2}, \frac{1}{4}(1-n)) \in \mathbb{R}^n$ and

$$\mathbb{C}^n(\rho) = \{s \in \mathbb{C}^n | Re(s_i) = -\rho_i\}. \tag{6.28}$$

For $a, b \in \mathbb{C}$ with $Re(a), Re(b) > 0$, let

$$B(a,b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)},$$

where Γ is the gamma function. For $s \in \mathbb{C}^n$ let

$$c_m(s) = \prod_{1 \le i \le j \le n-1} \frac{B(\frac{1}{2}, s_i + \cdots + s_j + \frac{1}{2}(j-i+1))}{B(\frac{1}{2}, \frac{1}{2}(j-i+1))},$$

and

$$\omega_n = \frac{\prod_{j=1}^n \Gamma(j/2)}{(2i\pi)^n \pi^{n(n+1)/4} n!}.$$

Using this notation, the measure μ_s on $\mathbb{C}^n$ is defined as

$$d\mu_s = \omega_n |c_n(s)|^{-2} ds_1 \cdots ds_n.$$

For $f \in C_c^\infty$, Terras (1988) and Peter T. Kim and Donald St. P. Richards (2011) provided the inversion formula,

$$f(\Sigma) = \int_{s \in \mathbb{C}^n(\rho)} \int_{k \in SO(n)/H} \mathcal{H}(f)(s,R) e_{s,R}(\Sigma) d\bar{k} d\mu_s,$$

where $d\bar{k}$ stands for the normalized factorization of the Haar measure of $SO(n)$ on $SO(n)/H$. Following the discussion of section 6.4.1, the expression can be rewritten as

$$f(\Sigma) = \int_{s \in \mathbb{C}^n(\rho)} \int_{R \in \mathcal{D}} \mathcal{H}(f)(s,R) e_{s,R}(\Sigma) dHaar_{SO(n)} d\mu_s, \tag{6.29}$$

where $Haar_{SO(n)}$ is taken such that $Haar_{SO(n)}(\mathcal{D}) = 1$.

6.7. Density estimation on multivariate centred Gaussian laws

Let X be a random variable in $\Omega \to \mathcal{G}_{m=0}(n)$. Notations of this section are similar to those of section 6.2. For the reasons explained in the previous section, we are here only interested in the Fisher metric. The case of the Wasserstein metric is the same studied in section 6.5.2.

6.7.1. Histograms

Let (A_j) be a partition of $\mathcal{G}_{m=0}(n)$. Let α_j be the number of draws that lay in A_j, $\alpha_j = Card\{x_i \in A_j\}$. The histogram of the draws (x_i) is the function

$$\hat{f}_k = \frac{1}{k}\sum_j \frac{1}{vol(A_j)}\alpha_j \mathbf{1}_{A_j},$$

where $\mathbf{1}_E$ is the indicator function of E. Thus $\hat{f}_k$ is an estimator of the density f. In order to respect the criterion mentioned in section 6.2, bins should be isometric and verify some isotropy properties. In the Euclidean context, square bins present an interesting compromise between isotropy and computational simplicity. The main drawback of this method is the difficulty to obtain a regular partition of the space. A lead is to see $\mathcal{G}_{m=0}(n)$ as the imaginary part of the Siegel upper half space. See Zelikin (2000) for an introduction to the Siegel upper half space. Regular polygonal tiling on the hyperbolic space, which can be extended to the Siegel space, are widely studied and are good candidates for the histogram bins. However the absence of homothetic transforms makes the adaptation of the estimation to the scaling factor difficult.

6.7.2. Kernel density estimation

The explicit expression of the change of volume $\theta_p(q)$ induced by the exponential map enables us to use the kernel density estimation. Given a scaling parameter $r > 0$, the estimation becomes:

$$\hat{f}_k^r = \frac{1}{k}\sum_i \frac{1}{r^n}\frac{1}{\theta_{\Sigma_{x_i}}(\Sigma_x)} K\left(\frac{d_{Fisher}(x, x_i)}{r}\right), \tag{6.30}$$

where $\theta_{\Sigma_{x_i}}(\Sigma_x)$ is given in Eq. (6.24). Thus, we have the following expression for the Pelletier kernel density estimator

$$\hat{f}_k^r = \frac{1}{k}\sum_i \frac{2^{n(n-1)/2}}{r^n}\prod_{p<q}\frac{sinh(\frac{|\lambda_p-\lambda_q|}{2})}{|\lambda_p-\lambda_q|} K\left(\frac{||log(\Sigma_{x_i}^{-1/2}\Sigma_x\Sigma_{x_i}^{-1/2})||}{r}\right), \tag{6.31}$$

where the $\lambda_.$ are the eigenvalues of $log(\Sigma_{x_i}^{-1/2}\Sigma_x\Sigma_{x_i}^{-1/2})$.

6.7.3. Orthogonal series density estimation

As discussed in previous section, the symmetry of the space gives rise to a Fourier–Helgason transform and its inverse. The empirical transform is given by:

$$\mathcal{H}_{emp}(f)(s, R) = \sum_i f(\Sigma_{x_i})\overline{e_{s,R}(\Sigma_{x_i})},$$

where $e_{s,R}$ is the plane wave defined in Eq. (6.26).

For $s \in \mathbb{C}^n(\rho)$, see Eq. (6.28), let

$$r_j = s_j + s_{j+1} + \cdots + s_n + \frac{1}{4}(n-2j+1) \text{ and } \lambda_s = -(r_1^2 + r_2^2 + \cdots + r_n^2 - \frac{1}{48}n(n^2-1)),$$

then

$$\Delta e_{s,R}(\Sigma) = \lambda_s e_{s,R}(\Sigma).$$

Let $\mathbb{C}^n(\rho, T) = \{s \in \mathbb{C}^n(\rho), \lambda_s < T\}$. For a cut-off parameter $T > 0$, the associated density estimator is then given by:

$$\hat{f}_k^T(\Sigma) = \int_{s \in \mathbb{C}^n(\rho,T)} \int_{R \in Q} \mathcal{H}_{emp}(f)(s, R) e_{s,R}(\Sigma) dHaar_{SO(n)} d\mu_s. \tag{6.32}$$

Let $\mathbf{1}_{F_T} = \mathbb{C}^n(\rho, T) \times SO(n)$. It can be verified that $\mathcal{H}^{-1}(\mathbf{1}_{F_T})$ is invariant under the action of $SO(n)$. Thus, one has

$$\hat{f}_k^T = \mu_k * \mathcal{H}^{-1}(\mathbf{1}_{F_T}).$$

In other words, the estimation do not privilege specific directions or locations. The main drawback of this method is its computational cost: one has to evaluate integral (6.32). The integral present in the transform is due to the non-compactness of the underlying space $\mathrm{SPD}(n)$. In $\mathbb{R}^n$, the assumption that the random variable lays in a parallelepiped P, that is to say $Supp \subset P$, enables the use of Fourier series instead of the Fourier transform. Then the estimation involves only a finite sum. Unfortunately, there are no compact domains $D \subset \mathrm{SPD}(n)$ where the eigenfunctions of the Laplacian are known.

6.8. The space of Gaussian laws with fixed rotation

Given a rotation matrix R, we study here the specificity of the space of Gaussian laws of the form $\mathcal{N}_{(m,RDR^t)}$ where D is diagonal matrix, under the Wasserstein metric. We choose here to represent the Gaussian law $\mathcal{N}_{(m,RDR^t)}$ by its mean $m \in \mathbb{R}^n$ and a set of standard deviation $\sigma \in \mathbb{R}^n$, with σ_i^2 an eigenvalue of D. In the (m, σ) parametrization, the expression of the metric found in (6.4.3.1) takes the following form:

$$G^{Wasserstein}_{\mathcal{N}_{m,\sigma}} = \begin{array}{c} \begin{matrix} \quad & e_{m_i} & & e_{\sigma_i} & \quad \end{matrix} \\ \begin{pmatrix} 1 & & & & \\ & \ddots & & & \\ & & 1 & & \\ & & & 1 & \\ & & & & \ddots \\ & & & & & 1 \end{pmatrix} \end{array}$$

Thus in the (m, σ) parametrization, the metric is Euclidean. The density estimation can then be achieved using standard techniques.

6.9. The space of Gaussian laws with fixed mean and rotation

Given a mean m and a rotation R, we study here the specificity of the space of Gaussian laws of the form $\mathcal{G}_{m,R}(n) = \{\mathcal{N}_{(m,RDR^t)}, \text{with } D \text{ diagonal}\}$, under the Fisher metric. The metric computed in (6.4.2) expressed in the basis $\{e_{\lambda_i}\}$ gives:

$$G^{Fisher}_{\mathcal{N}_{m,RD_\lambda R^t}} = \begin{array}{c} \begin{matrix} \quad & e_{\lambda_i} & \quad \end{matrix} \\ \begin{pmatrix} \frac{1}{\lambda_1^2} & & \\ & \cdot & \\ & & \frac{1}{\lambda_n^2} \end{pmatrix} \end{array}$$

After the change of coordinates $\gamma_i = log(\lambda_i)$, the metric becomes:

$$G^{Fisher}_{\mathcal{N}_{m,RD_\lambda R^t}} = \begin{pmatrix} e_{\gamma_i} & & \\ 1 & & \\ & . & \\ & & 1 \end{pmatrix}$$

Thus in the $\gamma = log(\lambda)$ parametrization, the metric is Euclidean. The density estimation can then be achieved using standard techniques.

6.10. The space of Gaussian laws with fixed covariance

Given a covariance matrix Σ, we study now the specificity of the space of Gaussian laws of the form $\mathcal{N}_{(m,\Sigma)}$, under the Fisher metric. Note that the mean under the Wasserstein metric follows the canonical Euclidean scalar product. From Eq. (6.11) in $\mathcal{G}(n)$:

$$< u, v >^{Fisher}_{m,\Sigma} = u_m^t \Sigma^{-1} v_m + \frac{1}{2} tr(\Sigma^{-1} u_\Sigma \Sigma^{-1} v_\Sigma),$$

and we have that, at a fixed Σ, it is given by

$$< u_m, v_m >^{Fisher}_{m} = u_m^t \Sigma^{-1} v_m.$$

The metric on the mean is Euclidean, of scalar product Σ^{-1}. The density estimation can then be achieved using standard techniques.

6.11. The space of univariate Gaussian laws

This section addresses the case of $\mathcal{G}(n = 1)$. Space $\mathcal{G}(n = 1)$ is identified with the half plane $\mathbb{R} \times \mathbb{R}_+$. That means that each normal law $\mathcal{N}$ is described by its mean $m \in \mathbb{R}$ and its (positive) standard deviation $\sigma \in \mathbb{R}_+$. Note that, as in section 6.4.3, there is a slight change of convention with respect to the previous sections since the equivalent of the covariance matrix Σ would be the square of the standard deviation σ^2. The study is focused on the Fisher metric since the case of the Wasserstein metric has been analyzed in section 6.8.

6.11.1. Fisher metric and the Poincaré upper half plane

The Poincaré upper half plane of curvature R is a half plane $\{x \in \mathbb{R}, y \in \mathbb{R}_+\}$ endowed with the following Riemannian metric:

$$G_{(x,y)} = a^2 \begin{pmatrix} 1/y^2 & 0 \\ 0 & 1/y^2 \end{pmatrix}, \tag{6.33}$$

with $R = -\frac{1}{a^2}$. The Poincaré upper half plane is a model of hyperbolic geometry, see Cannon et al. (1997). The distance between (x_1, y_1) and (x_2, y_2) is given by:

$$d((x_1, y_1), (x_2, y_2)) = a \cosh^{-1}\left(1 + \frac{(x_2 - x_1)^2 + (y_2 - y_1)^2}{2y_1 y_2}\right). \tag{6.34}$$

The Poincaré upper half plane is isotropic: each location and directions are equivalent. After having considered the change of convention between Σ and σ, the matrix of the

Fisher metric on $\mathcal{G}(n=1)$ can be deduced from Eq. (6.4.2), i.e.,

$$G^{Fisher}_{(m,\sigma)} = \begin{pmatrix} 1/\sigma^2 & 0 \\ 0 & 2/\sigma^2 \end{pmatrix} \tag{6.35}$$

Consider now the application $s : (m,\sigma) \mapsto (x,y) = (m, \sqrt{2}\sigma)$. Solving the equation:

$$(ds)^t . G^{Fisher}_{(x,y)} . ds = G^{Fisher}_{(m,\sigma)} \tag{6.36}$$

where ds is the Jacobian matrix of s, the expression of the metric is obtained in the new parametrization:

$$G^{Fisher}_{(x,y)} = 2 \begin{pmatrix} 1/y^2 & 0 \\ 0 & 1/y^2 \end{pmatrix} \tag{6.37}$$

Thus, under the Fisher metric, the space $\mathcal{G}(n=1)$ is isometric to a Poincaré upper half plane of curvature $-\frac{1}{2}$. Let $\mathbb{H}$ denotes just the Poincaré upper half plane of curvature $-\frac{1}{2}$.

6.11.2. The Riemannian volume element

In the (x,y) parametrization, the Riemannian volume measure *vol* has the following expression:

$$\frac{dvol}{dLeb_{(x,y)}} = \sqrt{det(G^{Fisher}_{(x,y)})} = \frac{2}{y^2}. \tag{6.38}$$

6.11.3. The volume change of the exponential map

Given a reference point p, any point in polar coordinates (r,α) of the hyperbolic space is defined as the point $q_{(r,\alpha)}$ at distance r of p on the geodesic with initial direction $\alpha \in \mathbb{S}^1$. Since the hyperbolic space is isotropic, the expression of the length element in polar coordinates depends only on r, see Anker and Ostellari (2003); Grigoryan (2009). Expressed in polar coordinates the hyperbolic metric of curvature -1 is:

$$\mathbf{g} = dr^2 + \sinh(r)^2 \mathbf{g}_{\mathbb{S}^1}.$$

For a hyperbolic space of curvature $-\frac{1}{2}$, the expression becomes:

$$\mathbf{g} = 2(dr^2 + \sinh(r)^2 \mathbf{g}_{\mathbb{S}^1}).$$

In the basis $\left\{ e_r = \frac{dp_{(r,\alpha)}}{dr}, e_\alpha = \frac{dp_{(r,\alpha)}}{rdr} \right\}$, the metric takes the following form:

$$G_{(r,\alpha)} = 2 \begin{pmatrix} 1 & 0 \\ 0 & \sinh(r)^2 \frac{1}{r^2} \end{pmatrix} \tag{6.39}$$

Thus, if we consider *vol* as a measure on the parametrization (r,α), one has

$$\frac{dvol}{dLeb_{(r,\alpha)}}(r,\alpha) = \sqrt{det(G_{(r,\alpha)})} = 2\frac{\sinh(r)}{r}, \tag{6.40}$$

where $r = d(p,q)$. The polar coordinates are a polar expression of the exponential map at p. Indeed the exponential map at p can be defined as $exp_p(r(cos(\alpha), sin(\alpha))) = q_{(r,\alpha)}$.

Therefore, considering now *vol* as a measure on the Riemannian space, the volume change of the exponential is just:

$$\frac{dvol}{dexp^*(Leb_p)}(q_{(r,\alpha)}) = 2\frac{\sinh(r)}{r}. \tag{6.41}$$

We obtain then

$$\theta_p(q) = 2\frac{\sinh(d(p,q))}{d(p,q)}, \tag{6.42}$$

where θ is the function used in Eq. (6.3).

6.11.3.1. Fourier–Helgason transform

We have said that $\mathcal{G}(n=1)$ is symmetric, therefore there exists a Fourier–Helgason transform. This paragraph is mainly extracted from Huckemann et al. (2010). Space $\mathbb{H}$ is identified with the complex upper half plane: $(1,0)=1$ and $(0,1)=i$. For $(s,k)\in\mathbb{C}\times SO(2)$, the Helgason functions (plane waves) have in the case the following expression:

$$e_{s,k_u} : \left\{ \begin{array}{ccc} \mathbb{H} & \rightarrow & \mathbb{C} \\ z & \mapsto & Im(k_u.z)^s = e^{Re(s)\log(Im(k_u.z))}e^{iIm(s)\log(Im(k_u.z))} \end{array} \right. \tag{6.43}$$

where $Re(.)$, $Im(.)$ denotes the real and imaginary part,

$$k_u = \begin{pmatrix} \cos(u) & \sin(u) \\ -\sin(u) & \cos(u) \end{pmatrix},$$

and

$$k_u.z = \frac{z\cos(u)+\sin(u)}{\cos(u)-z\sin(u)}.$$

We have that

$$\Delta e_{s,k_u} = s(s-1)e_{s,k_u}$$

, see Huckemann et al. (2010). Let $f\in C_c^\infty(\mathbb{H})$, where $C_c^\infty(\mathbb{H})$ denotes the set of infinitely differentiable with compact support. For $(s,k)\in\mathbb{C}\times SO(2)$, the Fourier–Helgason transform of f is given by:

$$\mathcal{FH}(f)(s,k_u) = \int_{\mathbb{H}} f(z)\overline{e_{s,k_u}(z)}dvol. \tag{6.44}$$

The inverse transform is given by:

$$f(z) = \int_{t\in\mathbb{R}}\int_{u=0}^{u=2\pi} \mathcal{FH}(f)(\frac{1}{2}+it,k_u)e_{\frac{1}{2}+it,k_u}(z)\frac{1}{8\pi^2}t\tanh(\pi t)dudt. \tag{6.45}$$

6.12. Density estimation on univariate Gaussian laws under the Fisher metric

Let X be a random variable in $\Omega\rightarrow\mathcal{G}(n=1)$. Notations of this section are similar to notation of section 6.2, except for the set of draw which now is $(z_1,...,z_k)$, with $z_i=(m_{z_i},\sigma_{z_i})$. For the reasons explained in the previous section we are here only interested in the Fisher metric.

6.12.1. Histograms

The problem has already been partially addressed in section 6.7.1. Recent works on Voronoï diagrams in the hyperbolic space Nielsen and Nock (2010) might provide alternative solutions through the construction of adaptive bins.

6.12.2. Kernel

The explicit expression of the change of volume $\theta_p(q)$ induced by the exponential map enables us to use the kernel density estimation. Given a scaling parameter $r > 0$, the estimation becomes:

$$\hat{f}_k^r = \frac{1}{k}\sum_i \frac{1}{r^n}\frac{1}{\theta_{z_i}(z)} K\left(\frac{d_{Fisher}(z,z_i)}{r}\right), \tag{6.46}$$

where $\theta_{z_i}(z)$ is given in Eq. (6.42). Then, we have the close expression:

$$\hat{f}_k^r = \frac{1}{k}\sum_i \frac{1}{r^n}\frac{d_{Fisher}(z,z_i)}{2\sinh(d_{Fisher}(z,z_i))} K\left(\frac{d_{Fisher}(z,z_i)}{r}\right), \tag{6.47}$$

where here $d_{Fisher}(z,z_i) = \sqrt{2}\cosh^{-1}\left(1+\frac{(m_z-m_{z_i})^2+2(\sigma_z-\sigma_{z_i})^2}{2\sigma_z\sigma_{z_i}}\right)$.

6.12.3. Orthogonal series

The symmetry of the space provides a Helgason transform and its inverse. The empirical transform is given by:

$$\mathcal{H}_{emp}(f)(s,k_u) = \sum_i f(z_i)\overline{e_{s,k}(z_i)},$$

where e_{s,k_u} is the plane wave defined in 6.43.

For a cut-off parameter $T > 0$, the associated density estimator is then given by:

$$\hat{f}_k^T(z) = \int_{|t|<T}\int_{u=0}^{u=2\pi} \mathcal{H}_{emp}(f)(\frac{1}{2}+it,k_u)e_{\frac{1}{2}+it,k_u}(z)\frac{1}{8\pi^2}t\tanh(\pi t)dudt. \tag{6.48}$$

Let $\mathbf{1}_{F_T} = [-T,T]\times SO(2)$. It can be verified that $\mathcal{H}^{-1}(\mathbf{1}_{F_T})$ is invariant under the action of $SO(2)$. Thus,

$$\hat{f}_k^T = D_k * \mathcal{H}^{-1}(\mathbf{1}_{F_T}),$$

and, consequently the estimation do not privilege specific directions or locations.

The main drawback of this method is its computational cost. Indeed the situation is similar to that considered in section 6.7.3.

6.13. Partial quantities: mean, eigenvalues and rotation

This section addresses the study of the standard partial quantities that are the mean, the eigenvalues and the rotation. On the one hand, depending on the application, the interesting information is sometimes carried by only one or two partial quantities. On the

other hand, as mentioned in section 6.2.5 the curse of dimensionality pushes us to reduce the dimension of the studied objects.

We start the discussion with an example. Let X be a random variable valued in $\mathbb{R}^2$ of density f. Consider the distribution of the radius of the polar coordinates (r, θ). Firstly, the Lebesgue measure of $\mathbb{R}^2$ can be written as a product $dLeb = rd\theta dr$. Thus the Lebesgue measure of $\mathbb{R}^2$ induces a measure on radiuses rdr. The density computed with respect to rdr can be interpreted as an average of f over a slice of constant radius. Secondly, the metric on $\mathbb{R}^2$ induces a natural metric on the space of radius seen as the quotient space $\mathbb{R} \sim \mathbb{R}^2/\theta$. The natural quotient metric is the Euclidean metric on $\mathbb{R}$, the associated measure being the Lebesgue measure dr. Thus this example shows that there is not a unique way of addressing the density estimation of partial quantities.

Remind the context of section 6.2. Let Ω be a space, endowed of with a σ-algebra and a probability measure p. Space $\mathcal{G}(n)$ is equipped with a Riemannian metric G, the associated Riemannian measure is called vol. Let X be a random variable, $X, \Omega \mapsto \mathcal{G}(n)$. The measure on $\mathcal{G}(n)$ induced by X is noted μ_X. We assume that μ_X has a density, noted f, with respect to vol. $\tilde{\mathcal{G}}(n)$ can be identified to the product $\mathbb{R}^n \times E \times \mathcal{D}$, see section 6.4. Recall that $vol(\mathcal{G}(n) \setminus \tilde{\mathcal{G}}(n)) = 0$. Under this identification, the measure vol is described by a product $vol = \mu_{mean} \times \mu_\lambda \times \mu_R$, see section 6.4.1. Let X_{mean}, X_λ and X_R be the random variables naturally induced by the identification between $\tilde{\mathcal{G}}(n)$ and $\mathbb{R}^n \times E \times \mathcal{D}$. Note that X_R is not always properly defined. This problem can be neglected given that $vol(\mathcal{G}(n) \setminus \tilde{\mathcal{G}}(n)) = 0$ and that the law of X has a density with respect to vol. Let $\mu_{X_{mean}}$, μ_{X_λ}, μ_{X_R} be the associated measures.

6.13.1. Average over slices

Measure $\mu_{X_{mean}}(A)$ is the probability that $m_x \in A$ when x follows μ_X. Thus,

$$\mu_{X_{mean}}(A) = \mu_X(A \times E \times \mathcal{D}).$$

Let f_{mean} be the density of $\mu_{X_{mean}}$ with respect to μ_{mean}. Since

$$\mu_X(A \times E \times \mathcal{D}) = \int_A \int_E \int_{\mathcal{D}} f d\mu_R d\mu_\lambda d\mu_{mean},$$

we have,

$$f_{mean}(m) = \int_E \int_{\mathcal{D}} f(m, \lambda, R) d\mu_R d\mu_\lambda.$$

Let $\alpha = (\mu_\lambda(E)\mu_R(\mathcal{D}))^{-1}$. Quantity αf_{mean} can be interpreted as an average value of the original density f over a slice $E \times \mathcal{D}$. The same hold for μ_{X_λ} and μ_{X_R}. Let $\mu^{Eucl}_{\cdot}$ be the measure induced by the canonical Euclidean metric of $\mathbb{R}^n \times Sym(n)$. Up to a scaling

factor, we recall here the expression of the different measures:

$$\begin{aligned}
\mu_{mean}^{Eucl} &= Leb,\\
d\mu_{\lambda}^{Eucl} &= \prod_{i<j\leq n} |\lambda_i - \lambda_j| dLeb(\lambda),\\
\mu_{R}^{Eucl} &= Haar_{SO(n)},\\
\mu_{mean}^{Fisher} &= Leb,\\
\mu_{R}^{Fisher} &= Haar_{SO(n)},\\
\mu_{mean}^{Wasserstein} &= Leb,\\
d\mu_{\lambda}^{Wasserstein} &= \prod_{i} \frac{1}{\sqrt{\lambda_i}} \prod_{i<j} \frac{|\lambda_i - \lambda_j|}{\sqrt{(\lambda_i + \lambda_j)}} dLeb,\\
\mu_{R}^{Wasserstein} &= Haar_{SO(n)},
\end{aligned}$$

where $Haar_{SO(n)}$ is restricted to $\mathcal{D}$. The expression of μ_{λ}^{Fisher} varies between $\mathcal{G}(n)$ and $\mathcal{G}_{m=0}(n)$. When $\mathcal{G}(n)$ is the underlying space, we have:

$$d\mu_{\lambda}^{Fisher} = \prod_{i<j\leq n} |\lambda_i - \lambda_j| \sqrt{\prod_i \frac{1}{\lambda_i^{n+2}}} dLeb, \tag{6.49}$$

and when $\mathcal{G}_{m=0}(n)$ is the underlying space the expression becomes:

$$d\mu_{\lambda}^{Fisher} = \prod_{i<j\leq n} |\lambda_i - \lambda_j| \sqrt{\prod_i \frac{1}{\lambda_i^{n+1}}} dLeb \tag{6.50}$$

Fig. 6.3 shows several cases of density change for $n = 2$.

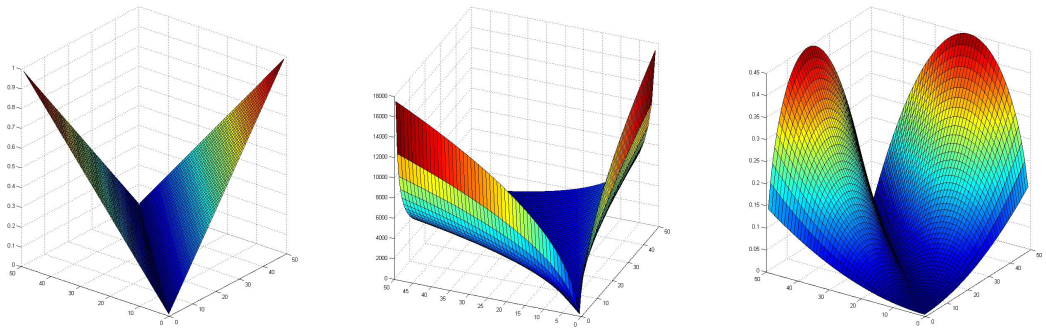

(a) (b) (c)

Figure 6.3.: For $n = 2$, (a),(b) and (c) are visualizations of $d\mu_{\lambda}^{Eucl}$, $d\mu_{\lambda}^{Wass}$ and $d\mu_{\lambda}^{Fisher}$ (for $\mathcal{G}_{m=0}(n)$) respectively.

The question of the underlying metric associated to these measures has not yet been clarified. In the absence of underlying metric, the notions of homogeneity and isotropy disappear. Thus the estimation can be achieved using the approach described section 6.2.1. Despite the absence of clear argument, the Lebesgue measure and the Haar measure push to the use of an underlying translation invariant metric.

6.13.2. Quotient metric on partial quantities

The space associated to each partial quantity can be seen as a quotient of $\mathcal{G}(n)$ by the rest of the partial quantities. The quotient $\mathcal{G}(n)/(\mathbb{R}_+^n \times SO(n))$ denotes the identification elements with all those which differ only in the eigenvalues or the rotation. Thus $\mathcal{G}(n)/(\mathbb{R}_+^n \times SO(n))$ can be interpreted as the space of means. In several cases, the Riemannian metric on $\mathcal{G}(n)$ induces a canonical metric on such quotient space. This metric induces a measure on the partial quantity, which enable the estimation of density of the associated random variable $X_{(.)}$.

6.13.2.1. Mean

For the Wasserstein metric and the Euclidean metric of $\mathbb{R}^n \times Sym(n)$, the quotient space $\mathbb{R}^n \sim \mathcal{G}(n)/(\mathbb{R}_+^n \times SO(n))$ inherits naturally of the canonical Euclidean metric. Indeed, each couple (λ, R) induces the same metric on m. The density estimation can thus be achieved using standard techniques. Note that the induced Lebesgue measure is similar to μ_{mean}.

Recall the expression of the scalar product of the Fisher metric (6.11):

$$< u, v >_{m,\Sigma}^{Fisher} = u_m^t \Sigma^{-1} v_m + \frac{1}{2} tr(\Sigma^{-1} u_\Sigma \Sigma^{-1} v_\Sigma).$$

Each covariance matrix Σ induces a metric on the space of means. However, the metric varies when Σ varies. Thus the Fisher metric does not induces a canonical metric on $\mathbb{R}^n \sim \mathcal{G}(n)/(\mathbb{R}_+^n \times SO(n))$.

6.13.2.2. Eigenvalues

For $\lambda \in \mathbb{R}^n$, let D_λ be the associated diagonal matrix. Fixing a mean m and a rotation R induces a metric on eigenvalues:

$$d_{m,R}(\lambda_1, \lambda_2) = d(\mathcal{N}_{m,D_{\lambda_1},R}, \mathcal{N}_{m,D_{\lambda_2},R}). \tag{6.51}$$

The group of rotations naturally acts on Gaussian laws by $R_a \cdot_{Left} \mathcal{N}_{m,\lambda,R_b} = \mathcal{N}_{m,\lambda,R_a R_b}$. Furthermore for the Euclidean metric of $\mathbb{R}^n \times Sym(n)$, the Fisher and the Wasserstein metric, we have that for all $R \in SO(n)$:

$$d(\mathcal{N}_{m_1,\lambda_1,R_1}, \mathcal{N}_{m_2,\lambda_2,R_2}) = d(R \cdot_{Left} \mathcal{N}_{m,\lambda,R_1}, R \cdot_{Left} \mathcal{N}_{m,\lambda,R_1}). \tag{6.52}$$

Thus since $d_{m,R}$ is independent of m and R, the three metrics induce a canonical metric on the quotient $\mathbb{R}_+^n \sim \mathcal{G}(n)/(\mathbb{R}^n \times SO(n))$.

On the one hand according to section 6.8, the Wasserstein metric on eigenvalues at fixed rotation is the canonical Euclidean metric after the change of coordinates $\sigma = \sqrt{\lambda}$. On the other hand, according to section 6.9 the Fisher metric on eigenvalues at fixed mean and rotation is the canonical Euclidean metric after the change of coordinates $\gamma = \log(\lambda)$. Finally, it is easy to see that the quotient metric induced by the Euclidean metric of $\mathbb{R}^n \times Sym(n)$, is the Euclidean metric on $\mathbb{R}_+$. Note that the associated measures differs from μ_λ.

6.13.2.3. Rotation

Given a Gaussian law $\mathcal{N}_{m,\lambda,R}$, the Euclidean metric of $\mathbb{R}^n \times Sym(n)$, the Fisher and the Wasserstein metric induce a left-invariant metric on Q:

$$d_{\mathcal{N}_{m,\lambda,R}}(R_1 \in Q, R_2 \in Q) = d(R_1 \cdot_{Left} \mathcal{N}_{m,\lambda,R}, R_2 \cdot_{Left} \mathcal{N}_{m,\lambda,R}). \tag{6.53}$$

However according to the expressions of the metrics in the basis $(E'_{i\leq j})$ and the differential of $\tilde{\alpha}$, see Eq. 6.8, it can be verified that the left-invariante metric is dependent of the choice of $\mathcal{N}_{m,\lambda,R}$. Thus neither the Euclidean metric of $\mathbb{R}^n \times Sym(n)$, nor the Fisher nor the Wasserstein metric induce a canonical quotient metric on $SO(n) \sim \mathcal{G}(n)/(\mathbb{R}^n \times \mathbb{R}^n_+)$.

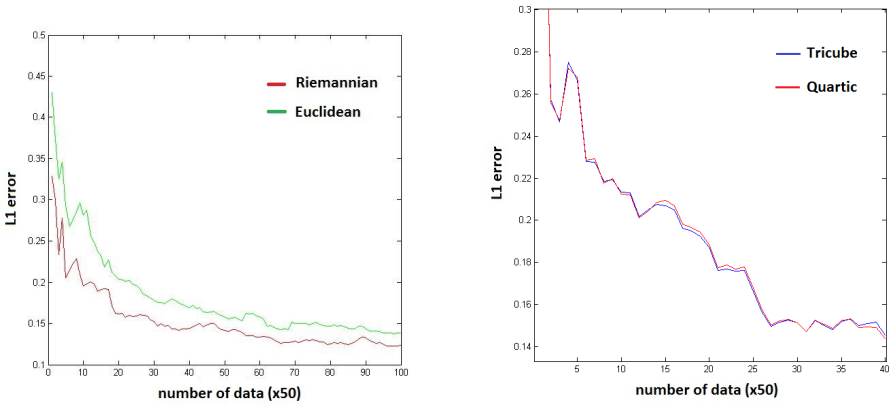

(a) (b)

Figure 6.4.: (a): Convergence of density estimations using an Euclidean and a Riemannian Kernel, in a Riemannian context, (b): the density is estimated following section 6.2.3 using two standard kernels of the plane, $K(||x||) = \frac{3}{\pi}(1-||x||^2)^2\mathbf{1}_{||x||<1}$ and $K(x) = \frac{220}{81\pi}(1-||x||^3)^3\mathbf{1}_{||x||<1}$.

6.14. Experimental section

Every standard density estimation technique involves a scaling parameter. This scaling factor controls the influence of the observation x_i on the estimated density at x, depending on the distance between x and x_i. In the experiments, the scaling factor has been chosen following the framework proposed in Duin (1976): a cross validation of the likelihood of the estimator.

This section start with an illustration, see Fig. 6.4(a), of the importance of the underlying metric on the density estimation method. Points are drawn in the Poincaré upper half plane according to the following density:

$$f(z) = \sqrt{d_{Fisher}(z,(0,1))^2 - 1}\mathbf{1}_{d_{Fisher}(z,(0,1))<1}.$$

From the draws, firstly, the density is estimated using an Euclidean kernel method, followed by an adaptation to the Riemannian measure, as described in section 6.2.1. Secondly, the density is estimated using a Riemannian kernel, see section 6.2.3. The base kernel is a quartic kernel $K(||x||) = \frac{3}{\pi}(1-||x||^2)^2\mathbf{1}_{||x||<1}$. Fig. 6.4(a) shows the convergence of the two estimations to the true density. The second experiment, see Fig. 6.4(b), illustrates the limited impact of the shape of the kernel K.

6.14.1. Histograms of multiple grey-scale image acquisition

The studied example is a time lapse sequence of grey-scale images from a retina. At each pixel, we have 20 successive fast acquisitions. By assuming a Gaussian distribution on the time series, we obtain a Gaussian valued image represented in Fig. 6.5.

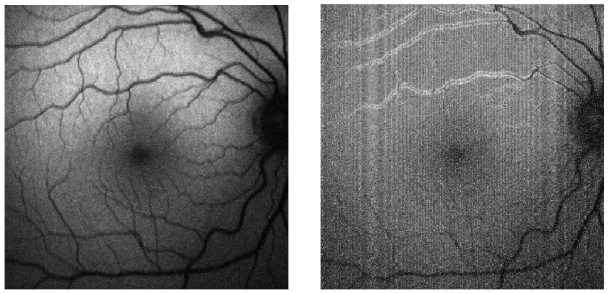

(a) (b)

Figure 6.5.: Time lapse sequence of images from the retina: (a) mean image, (b) standard deviation image.

Each pixel of the image contains an univariate Gaussian law. We recall that for univariate Gaussian laws, the Wasserstein metric is the Euclidean metric of the (m, σ) plane and that the Fisher metric is the Poincaré metric. Histograms of Gaussian-valued image Fig. 6.5 are computed in Fig. 6.6 with respect the Fisher and the Wasserstein metric, using the appropriate kernel density estimation.

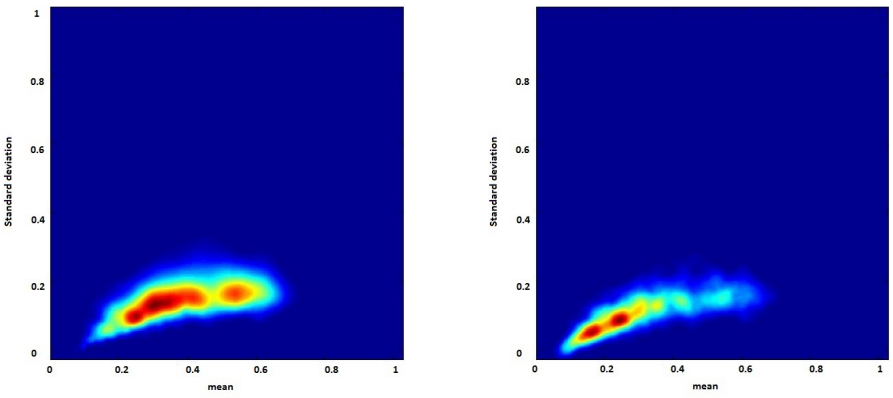

(a) (b)

Figure 6.6.: Histogram of Gaussian-valued image Fig. 6.5 using the Fisher metric in (a) and the Wasserstein metric in (b).

6.14.2. Density estimation in structure tensor images

The structure tensor image from a grayscale image is a field of symmetric positive definite matrices. Despite the fact that these matrices are not directly covariance matrices of Gaussian laws, they can be studied under the Fisher and the Wasserstein metric of $\mathcal{G}_{m=0}(n)$, typically $n = 2$ for 2D images and $n = 3$ for 3D ones. Fig. 6.7(b)-(g) shows the densities of the two eigenvalues $\lambda_1 \leq \lambda_2$ of the structure tensor field computed from Fig. 6.7(a). For each metric, the densities of eigenvalues are computed following considerations of section 6.13. For the induced measures μ_λ^{Eucl}, μ_λ^{Fisher} and $\mu_\lambda^{Wasserstein}$, the densities are firstly estimated using an Euclidean kernel density estimation, followed by an adaptation to the reference measures, given in section 6.13.1. According to section 6.13.2.2, each quotient

metric is Euclidean in the right parametrization. In the adapted parametrizations, the densities are obtained using an Euclidean kernel density estimation.

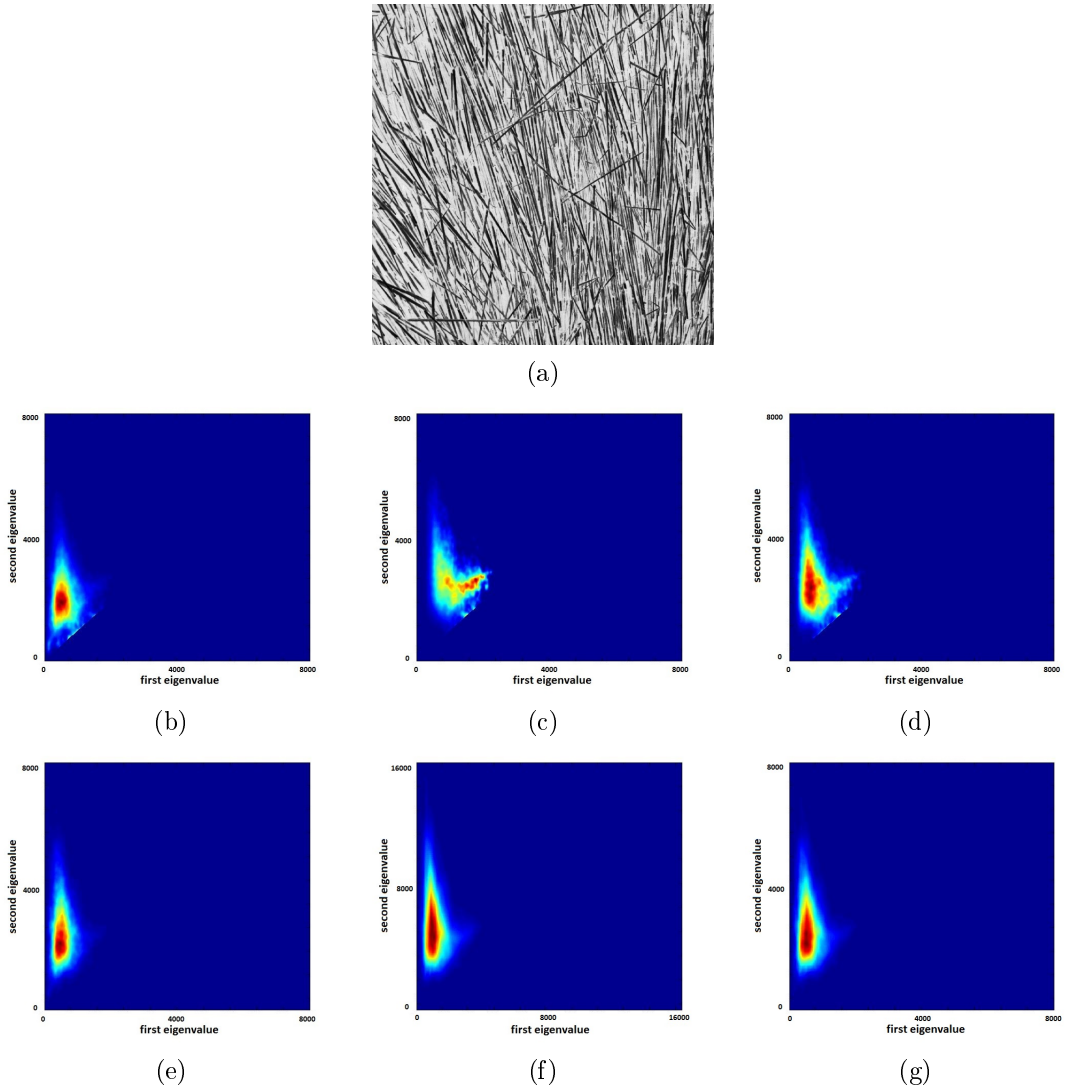

(a)

(b) (c) (d)

(e) (f) (g)

Figure 6.7.: Densities estimated from the structure tensor field computed from image (a). Top row, densities computed from the induces measures μ_{λ}: (b) Euclidean metric of $\mathbb{R}^n \times Sym(n)$, (c) Fisher metric, (d) Wasserstein metric. Bottom row, densities computed using the quotient metrics: (e) Euclidean metric of $\mathbb{R}^n \times Sym(n)$, (f) Fisher metric (note the change of scale), (f) Wasserstein metric.

Given a set of samples, the watershed transform of the complement of its associated density provides a non parametric clustering Soille (2004), similar to that of mean-shift algorithm. When the set of samples are the values taken by an image, the clustering of the samples can be interpreted as a segmentation of the image. Fig. 6.8 presents an example of texture segmentation according to the watershed transform of the density of the structure tensor field using measures.

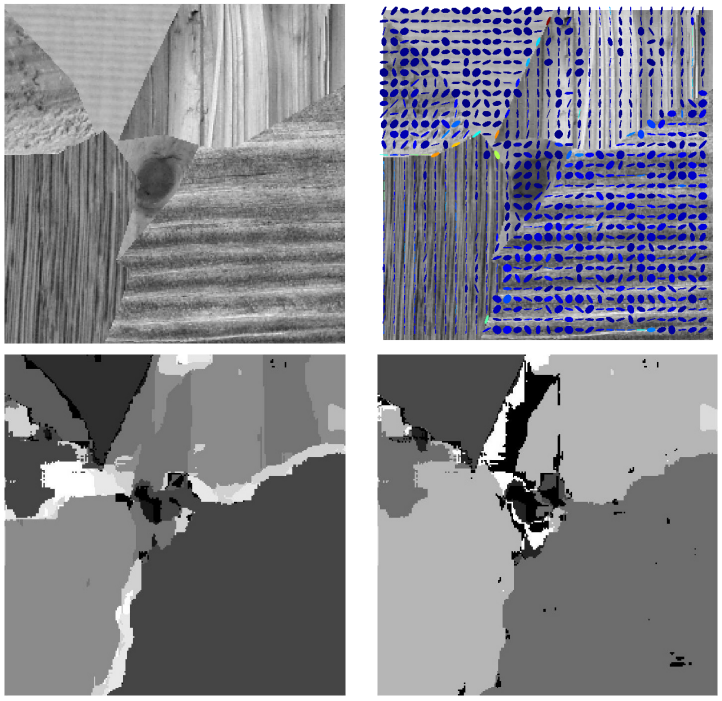

Figure 6.8.: Texture image segmentation: (a) and (b) are respectively a texture image and its associated structure tensor field. The density of the eigenvalues of the structure tensors is estimated with respect to the two measures discussed in section 6.13 induced by the Wasserstein metric: (c) segmentation with respect to the Lebesgue measure on eigenvalues, and (d) segmentation with respect to μ_λ^{Wass}.

6.14.3. Density estimation in Diffusion Tensor Imaging

Diffusion tensor imaging (DTI) data represent an average motion of water molecules in each voxel of the image, modeled by a three-dimensional centered Gaussian law. Gaussian laws are therefore parametrized by their covariance matrices. Thus, the resulting image is valued in SPD(3). Fig. 6.9(a) is a visual representation of a covariance matrix field on a two dimensional slice of a brain. Fig. 6.9(c) is the corresponding fractional anisotropy. Fig. 6.9(b) is a segmentation of the brain based on the distribution of the largest eigenvalue of each matrix.

The limiting factor for the use of density estimation is the poor spatial resolution. Indeed, the spatial resolution is not yet sufficient to enable the estimation of densities on local spatial neighborhoods.

6.14.4. Radar density estimation on the Poincaré disk

In this section the random variable is not an image I but a set of coefficients derived from a radar signal. In particular, the draws (x_i) are not explicitly Gaussian laws, but falls into the Poincaré disk, an analog of the Poincaré upper half plane.

Let us discuss briefly how radar data are related to hyperbolic space via reflection coefficients, for more details see Arnaudon et al. (2013). Each radar cell is a complex vector $z = (z_0, \cdots, z_{n-1})$ considered as a realization of a centered stationary Gaussian process

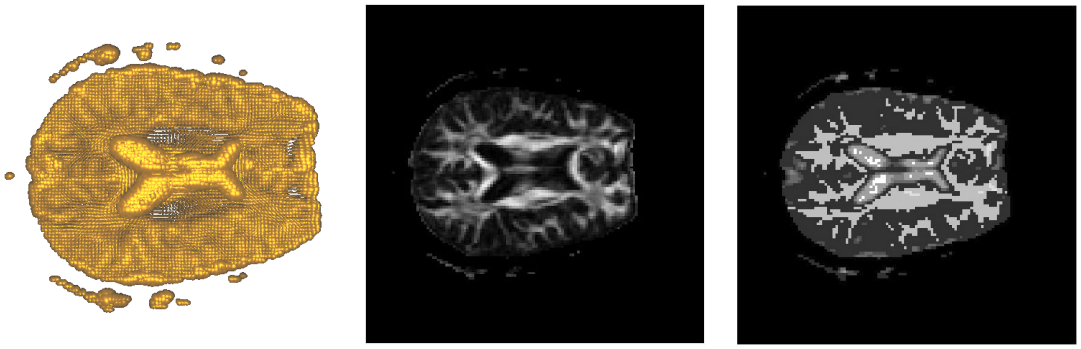

(a) (b) (c)

Figure 6.9.: DTI segmentation by watershed partitioning of its estimated density (histogram): (a) original SPD(3)-valued image, (b) its map of fractional anisotropy. The density of the largest eigenvalues of each matrix is evaluated against the Lebesgue measure of $\mathbb{R}^n$ and (c) is a segmentation corresponding to the watershed transform of the complement of the this density.

$Z = (Z_0, \cdots, Z_{n-1})$ of covariance matrix $R_n = \mathbb{E}[ZZ^*]$. The matrix R_n has a Toeplitz structure. For $1 \leq k \leq l \leq n-1$, the k-th order autoregressive estimate of Z_l is given by $\hat{Z}_l = -\sum_{j=1}^{k} a_j^{(k)} Z_{l-j}$, where the autoregressive coefficients $a_1^{(k)} \cdots a_k^{(k)}$ are chosen such that the mean squared error $\mathbb{E}(|Z_l - \hat{Z}_l|^2)$ is minimized. In practice, reflection coefficients are estimated by regularized Burg algorithm BARBARESCO and Barbaresco (1996). The last autoregressive coefficient $a_k^{(k)}$ is called the k-th reflection coefficient, denoted by μ_k and which has the property $|\mu_k| < 1$. The coefficient for $k = 0$ corresponds to the power, denoted $P_0 \in \mathbb{R}_+^*$. The reflection coefficients induce a (diffeomorphic) map φ between the Toeplitz Hermitian positive definite (HPD) matrices of order n, $\mathcal{T}^n$, and reflection coefficients:

$$\varphi : \ \mathcal{T}^n \to \mathbb{R}_+^* \times \mathbb{D}^{n-1}, \ \ R_n \mapsto (P_0, \mu_1, \cdots, \mu_{n-1})$$

where $\mathbb{D} = \{\zeta \in \mathbb{C} \ : \ |\zeta| < 1\}$ is the open unit disk of the complex plane. Diffeomorphism φ is very closely related to theorems of Trench (1964).

The Riemannian geometry of the space of reflection coefficients has been explored in Barbaresco (2013) through the Hessian of Kähler potential. The metric appropriate metric on $\mathcal{T}^n$ is

$$ds^2 = n\frac{dP_0^2}{P_0^2} + \sum_{k=1}^{n-1}(n-k)\frac{|d\mu_k|^2}{(1-|\mu_k|^2)^2}. \tag{6.54}$$

According to the metric (6.54) the space $\mathcal{T}^n$ can be seen as a product of the Riemannian manifold $\left(\mathbb{R}_+^*, ds_0^2\right)$, with $ds_0^2 = n(dP_0^2/P_0^2)$ (logarithmic metric multiplied by n), and $(n-1)$ copies of $\left(\mathbb{D}, ds_k^2\right)_{1 \leq k \leq n-1}$, with $ds_k^2 = (n-k)ds_{\mathbb{D}}^2$. $\left(\mathbb{R}_+^* \times \mathbb{D}^{n-1}, ds^2\right)$ is a Cartan–Hadamard manifold whose sectional curvatures are bounded, i.e., $-4 \leq K \leq 0$. This metric is related to information geometry and divergence functions, see discussion in Barbaresco (2014). From the product metric, closed forms of the Riemannian distance, arc-length parameterized geodesic, etc. can be obtained, see Barbaresco (2013); Arnaudon et al.

(2013). $\mathbb{D}$ endowed with the metric ds is called the Poincaré disk. The application

$$f(z) = \frac{z - i}{z + i},$$

is an isometry between the Poincaré upper half plan and the Poincaré disk. Thus the Poincaré upper half plan and the Poincaré disk are isometric. Fig. 6.10 presents the estimations of the marginal densities coefficients μ_k.

Data used in the experimental tests are radar observations from THALES X-band Radar, recorded during 2014 field trials campaign at Toulouse Blagnac Airport for European FP7 UFO study (Ultra-Fast wind sensOrs for wake-vortex hazards mitigation). Data are representative of Turbulent atmosphere monitored by radar in rainy conditions. Fig. 6.10 and 6.11 illustrate the density estimation of the six reflection coefficients on the Poincaré unit disk. For each coefficient the dataset is composed of 120 draws.

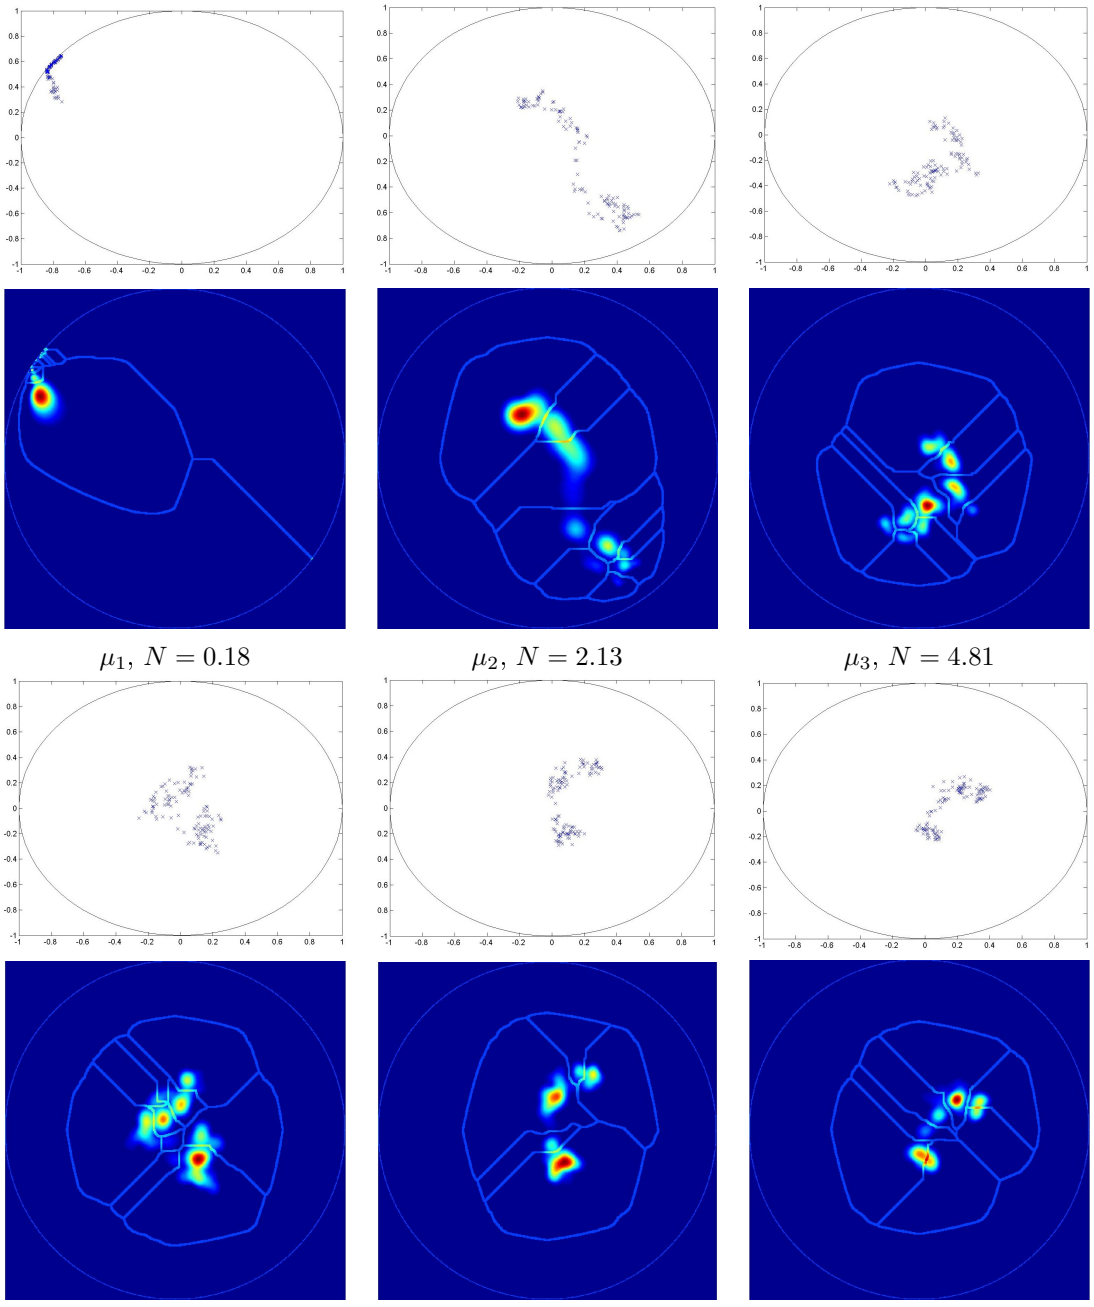

μ_1, $N = 0.18$ μ_2, $N = 2.13$ μ_3, $N = 4.81$

μ_4, $N = 3.83$ μ_5, $N = 5.26$ μ_6, $N = 6.99$

Figure 6.10.: Kernel density estimation of the distribution of the 6 first coefficients μ_k under rainy conditions, with $K(x) = \frac{3}{\pi}(1-x^2)^2\mathbf{1}_{x<1}$. For visualization clarity, densities are normalized by $N = \sup_{z\in\mathbb{D}}\{\hat{f}(z)\}$. Contours correspond to the watershed transform of the complement of estimated densities.

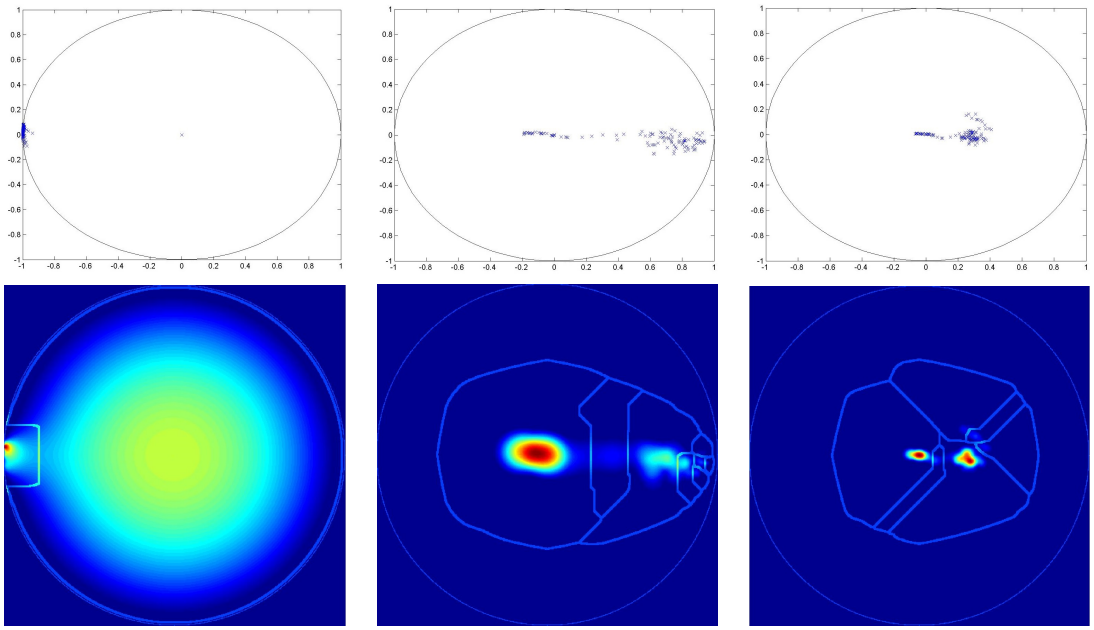

μ_1, $N = 0.007$ μ_2, $N = 1.61$ μ_3, $N = 14.86$

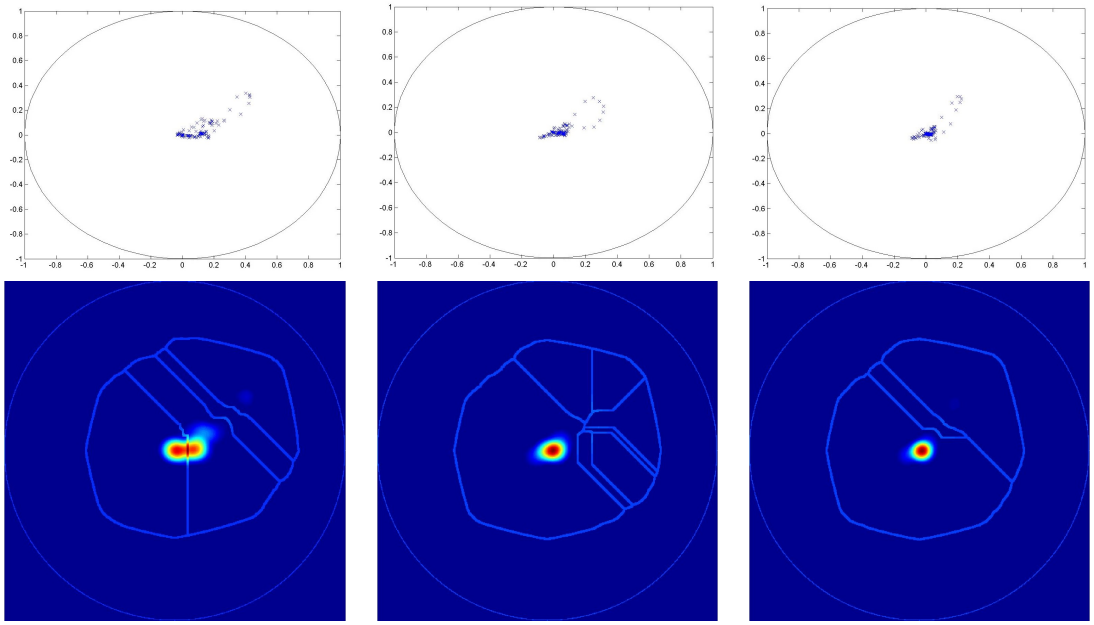

μ_4, $N = 8.82$ μ_5, $N = 17.31$ μ_6, $N = 23.29$

Figure 6.11.: Kernel density estimation of the distribution of the 6 first coefficients μ_k in normal conditions, with $K(x) = \frac{3}{\pi}(1-x^2)^2\mathbf{1}_{x<1}$. For visualization clarity, densities are normalized by $N = \sup_{z\in\mathbb{D}}\{\hat{f}(z)\}$. Contours correspond to the watershed transform of the complement of estimated densities.

6.15. Conclusions and perspectives

The whole theory and corresponding practical solutions to density estimation have been analyzed for standard spaces underlying Gaussian laws.

Under each space and each metric, four types of density estimation are considered. The first one consists in classical estimations in the Euclidean context of the parameters, multiplied by a density ratio to obtain a density with respect to the desired measure. This type of density estimation is possible in each studied space and metric, but does not take into account the geometry of the space: if two metrics induce the same volume measure then the estimated density does not depend on the choice of one the metrics. The three other types of estimation, histograms, kernels and orthogonal series rely more directly on the geometry. On the one hand, despite their low computational cost, finding good tiling make the use of histograms difficult. On the other hand, finding an orthogonal basis of functions adapted to orthogonal series density estimation is also a difficult problem. In the few situations where the eigenfunctions of the Laplacian operator are known, the estimator presents a high computational complexity.

The kernel density estimation seems to be the most adapted in most cases. It presents a reasonable computational complexity, and expressions of kernels are ready to use in each studied situations, except the space of multivariate Gaussian laws under the Fisher metric.

To our knowledge, the most original contribution of this work are the expression of kernels for the case of the Wasserstein metric, and the study of partial quantities. The latter being particularly useful in practical cases. The results of the chapter is summed up in tables 6.1-6.4.

	multivariate	multivariate centered	univariate
adaptation from the Lebesgue measure	Eq. 6.12, Eq. 6.13	Eq. 6.21, Eq. 6.22	Eq. 6.38
histograms	-	Siegel tiling	hyperbolic tiling
orthogonal series	-	Eq. 6.32	Eq. 6.48
kernels	-	Eq. 6.31	Eq. 6.47

Table 6.1.: Density estimation on Gaussian laws under the Fisher metric.

	mean	eigenvalues	rotations (Q)
quotient metric	-	Euclidean on $\gamma = \log(\lambda)$	-
product measure	Lebesgue	Eq. 6.49 or Eq. 6.50	$Haar_{SO(n)}$

Table 6.2.: Structures on partial quantities under the Fisher metric

Future work will focus mainly on the application of our result to different applications in image and signal processing. Densities are useful objects in segmentation and classification. The mean shift algorithm is a standard approach that search maxima of a density associated with a point cloud and then segment it into modes or classes. This algorithm, proposed by Fukunaga and Hostetler (1975) has been widely studied in $\mathbb{R}^n$. It is mainly used in

	multivariate	multivariate centered	univariate
adaptation from the Lebesgue measure	Eq. 6.15, Eq. 6.16	Eq. 6.15, Eq. 6.16	Euclidean on $\sigma = \sqrt{\lambda}$
histograms	-	-	Euclidean on $\sigma = \sqrt{\lambda}$
orthogonal series	-	-	Euclidean on $\sigma = \sqrt{\lambda}$
kernels	Eq. 6.20	Eq. 6.20	Euclidean on $\sigma = \sqrt{\lambda}$

Table 6.3.: Density estimation on Gaussian laws under the Wasserstein metric

	mean	eigenvalues	rotations (Q)
quotient metric	Euclidean	Euclidean on $\sigma = \sqrt{\lambda}$	-
product measure	Lebesgue	$\frac{d\mu_\lambda^{Wass}}{dLeb(\lambda)} = \prod_i \frac{1}{\sqrt{\lambda_i}} \prod_{i<j} \frac{\lvert\lambda_i - \lambda_j\rvert}{\sqrt{(\lambda_i + \lambda_j)}}$	$Haar_{SO(n)}$

Table 6.4.: Structures on partial quantities under the Wasserstein metric

point cloud segmentation, but also in tracking and smoothing. Hence, point clouds can be segmented according to the local maxima of an associated density using the mean shift algorithm. It presents for instance the advantage of not requiring assumptions on the number of clusters. The corresponding arbitrary quantity is the scaling factor of the kernels, but unlike the number of clusters, the scaling factor can have a physical interpretation. Consider the embedding of a signal $f : E \to F$ in the product space $E \times F$. The segmentation of the point cloud in $E \times F$ produces edge-preserving smoothing of the signal. The mean shift algorithm has been studied on Riemannian manifolds in Subbarao and Meer (2009a). However in order to gain generality authors chose to work with "pseudo" kernels in the sense that they do not take into account the volume change between the exponential map and the manifold. Indeed, this term is generally unknown. The study can thus be resumed for spaces Gaussian laws, except multivariate Gaussians under the Fisher metric. Especially since this thesis focuses on mathematical morphology and on density estimation, it is important to point out the similarity between the mean shift algorithm and the watershed transform of the density. Indeed the mean shift can be interpreted as a steepest ascent on the density.

Probability densities are also important tools in Bayesian classification. The classification decision consists in maximizing the probability of a class given an observation. This requires a priori information on classes and probability densities knowing the class. Our future research on this topic will study the influence of the choice of the metric on classification results. This type of classification has been addressed on SPD(n) in a recent work Said et al. (2015), where they model conditional densities by (parametric) Gaussian distributions on SPD(n). We remind that our approach here allows us a non-parametric density estimation on SPD(n).

Part III.

Conclusion

7. Conclusions and perspectives

Résumé

Ce chapitre résume les différents résultats de la thèse et présente les perspectives associées à ces résultats.

Le lemme 2.3.1 du chapitre 2 formalise les limites de la qualité des ordres totaux dans les espaces métriques, en terme de morphologie mathématique. A chaque ordre total, il est possible de trouver des images pour lesquelles les érosions et les dilatations sont irrégulières et incohérentes. Il est possible de construire une fonction de coût mesurant la qualité d'un ordre total. La procédure de minimisation de cette fonction de coût proposé dans le chapitre 2 permet de construire un ordre adapté à une image donnée. Les résultats de la procédure de minimisation sont intéressants mais sont encore sans doute sujets à améliorations. Les limites des ordres totaux ont conduit à s'orienter vers un autre type de données : les images à valeur dans un ensemble de labels non ordonnés. Un nouveau cadre morphologique est proposé pour traiter ces images. Une partie importante des propriétés des opérateurs binaire est conservée dans ce nouveau cadre. Les propriétés théoriques étant maintenant en majeure partie déterminées, les travaux futurs se concentreront sur les applications.

La formalisation de la notion d'histogramme d'image en terme d'estimation de densité suggère l'utilisation de la formule de la coaire. Cette idée à en fait déjà été explorée dans Duffy et al. (2013) dans le cas des images $I, \mathbb{R}^p \to \mathbb{R}^{q \leq p}$. Un travail peut-être envisagé sur le cas des images $I, \mathbb{R}^p \to \mathbb{R}^{q > p}$.

L'étude des histogrammes d'images couleurs dans le cadre des métriques perceptuelles propose de se baser non pas sur une approximation euclidienne globale de la métrique mais seulement sur une approximation locale. Cela permet d'être plus fidèle à la métrique perceptuelle, sans augmenter la complexité des calculs. Ce travail soulève le problème important de l'interpolation d'une métrique riemannienne.

Le chapitre 6 passe en revue plusieurs techniques d'estimation de densités de probabilités dans le cas où les données sont des lois Gaussiennes. Parmis les méthodes étudiées, la méthode des noyaux semble la plus adapté à la plupart des situations. Les travaux futurs se concentreront sur les applications pratiques.

7.1. Conclusions

7.1.1. Part I

Our reflections on mathematical morphology led to the following results. Lemma 2.3.1 of chapter 2 formalizes a limitation of the quality of total orders on multidimensional spaces for mathematical morphology: given a total order, there are always images where morphological operators will introduce irregularities and aliasing on processed images. In this thesis, the notion of quality of orders and operators is restricted to topological considerations. Given an image it is possible to asses the quality of a total order regarding its

impact on morphological operators. This quality measure is represented by a cost function introduced in Chapter 2. Thus given an image it is natural to search orders optimizing this quality measure. The main drawback of the cost function is its computational complexity. A recursive minimization procedure aiming at minimizing the cost has been proposed. Results in various contexts have been presented. The orders obtained by our minimization procedure produce interesting morphological erosions and dilations. However, orders have to be computed for each images and the minimization of the cost is a time demanding operation. Thus, the minimization framework does not yet an undeniable interest in front of orders such as the bit-mixing. Indeed the bit-mixing order presents interesting topological properties while being easy to evaluate and independent from the image.

The idea of n-ary morphology introduced in Chapter 3 has emerged as a circumvention of the theoretical limits of total orders: a structure based on only two operators can not properly take into account the geometry of a multidimensional space. On the real line each element can be approached by only two sides. In multidimensional spaces, the topological incompatibility exhibited by Lemma 2.3.1 arises from the increase of directions by which a point can be approached. This corresponds to the connectivity assumption on the complementary of balls. Hence the idea of introducing multiple pairs of erosions and dilations appears naturally. Each pair would then be associated with a pole of the space, each pole playing an equivalent role in the studied image. This underlying idea led to a formalization of n-ary morphology. Recall that n-ary images refer here to images valued in a finite set where no mathematical structure is assumed. In comparison with other works addressing the same type of images, the two main specificities of the present paradigm are the introduction of multiple operators and the absence of an "indeterminate" label or category. Several elementary properties of the binary framework are recovered. An extension to a continuous case has been proposed. Unfortunately theoretical properties are not recovered as naturally as in the n-ary framework. However, operators produce promising results in practice.

7.1.2. Part II

The formalization of image histograms in terms of density estimation suggests to address the estimation using the coarea formula. This approach has already been proposed recently in Duffy et al. (2013), originally aiming at visualization applications.

Perceptual color spaces like the *Lab* space are built such that their natural Euclidean metric is closer to the perceptual metric than the Euclidean metric of other spaces such as the *RGB* space. Since density estimation mainly involves local computations, the Riemannian assumption on the perceptual metric on colors suggest to locally adapt computations using the local scalar product. Indeed, as shown in Chapter 5 this enables to be more faithful to the Riemannian metric without requiring additional computation cost. The analyzed example shows the impact of the local adaptation.

Non-parametric probability density estimation has been discussed in Chapter 6 for several standard spaces of Gaussian laws, under different metrics. For each case, several density estimation techniques have been discussed and compared. Expressions of kernels for kernel density estimation, and functions involved in orthogonal series density estimation are provided when possible. The most original contributions of this work consists in the expression of kernels for the Wasserstein metric, and in the study of the structure of partial quantities (mean, eigenvalues and rotation) under different metrics. Partial quantities play an important role since studied phenomena often depend on only one or two parameters

of the Gaussian laws.

7.2. Perspectives

7.2.1. Part I

Our contributions and results on mathematical mathematical morphology lead to the following perspectives:

- The minimization procedure of the measure of irregularities of total order is not yet fully convincing. However, we are confident in the fact that the optimization can be significantly improved. The use of techniques such as evolutionary algorithms might for instance enable to significantly enhance the number of clusters at each step of the recursion, see Flórez-Revuelta (2015). This is currently an ongoing work.
- After having examined the transposition of properties of binary morphology to n-ary morphology, our future works will focus on applications. The main considered areas are material science images and pixelwise classified images. The proposed extension of n-ary morphology to a continuous case is still a preliminary work. Before addressing applications, future researches will continue investigating new continuous operators, in particular openings by reconstruction and levelings.

7.2.2. Part II

Ideas explored in Chapter 5 on images histograms provides several avenues for research:

- While more intuitive for functions $f : \mathbb{R}^p \to \mathbb{R}^{q \leq p}$ the coarea formula can also be formulated of functions $f : \mathbb{R}^p \to \mathbb{R}^{q > p}$. The existing work Duffy et al. (2013) only addresses the case of real valued functions. A work can thus be carried on so as to extend their proposition to color and multispectral images.
- The preliminary work on perceptual color histograms raises firstly the problem of Riemannian metric interpolation. The main difficulty is that despite the primary intuition, there is no intrinsic criteria that enables to evaluate the quality of an interpolation. The quality evaluation of an interpolation necessarily depends on the parametrization. A first study must then be carried out in order to understand how to determine in which parametrization the interpolation should be performed. We did not find - at least, not yet - a clear answer to this question in the literature. Secondly, it would be interesting to study the convergence of the proposed density to the true density of the distribution of colors with respect to the interpolated Riemannian metric.

Chapter 6 describes and gives strategies to estimate densities on spaces of Gaussian laws. Our future researches will focus on the potential uses of the estimated densities.

- The mean shift algorithm proposed by Fukunaga and Hostetler (1975) can be adapted to spaces of Gaussian laws using the expressions of kernels computed in Chapter 6. Given a point cloud, the mean-shift algorithm is a standard mode seeking technique. This technique is mainly used for point cloud segmentation tracking and smoothing Subbarao and Meer (2009b). When used for segmentation purposes, the mean-shift algorithm is very similar to a watershed transform of the estimated density of

the point cloud Paris and Durand (2007). The theoretical properties of the algorithm on spaces of Gaussian laws and its applications form an interesting subject of future studies.

- Density estimation is also an important tool in classification. This is the subject of recent studies on symmetric positive definite matrices using parametric density under the affine-invariant metric. Recall that the affine-invariant metric is similar to the Fisher metric on multivariate centered Gaussians. In the recent work Said et al. (2015), it is proposed to model the distribution of each classes by a parametric model and to take classification decision according to the estimated probabilities. The expression of kernels for the Fisher and the Wasserstein metric enables to conduct similar studies using non-parametric density estimation.

Publications related to the thesis

Chevallier, E., Angulo, J., Feb. 2014a. The discontinuity issue of total orders on metric spaces and its consequences for mathematical morphology. Accepted for publication in Journal of Mathematical Imaging and Vision.
URL https://hal.archives-ouvertes.fr/hal-00948232/document

Chevallier, E., Angulo, J., Oct. 2014b. Image adapted total ordering for mathematical morphology on multivariate images. In: 2014 IEEE International Conference on Image Processing (ICIP). pp. 2943–2947.

Chevallier, E., Barbaresco, F., Angulo, J., Feb. 2015a. Probability density estimation on the hyperbolic space applied to radar processing. Accepted at GSI 2015.
URL https://hal.archives-ouvertes.fr/hal-01121090/document

Chevallier, E., Chevallier, A., Angulo, J., Aug. 2014. Computing Histogram of Tensor Images Using Orthogonal Series Density Estimation and Riemannian Metrics. In: 2014 22nd International Conference on Pattern Recognition (ICPR). pp. 900–905.

Chevallier, E., Chevallier, A., Angulo, J., 2015b. N-ary Mathematical Morphology. In: Benediktsson, J. A., Chanussot, J., Najman, L., Talbot, H. (Eds.), Mathematical Morphology and Its Applications to Signal and Image Processing. No. 9082 in Lecture Notes in Computer Science. Springer International Publishing, pp. 339–350.
URL http://link.springer.com/chapter/10.1007/978-3-319-18720-4_29

Chevallier, E., Farup, I., Angulo, J., Feb. 2015c. Histograms of images valued in the manifold of colours endowed with perceptual metrics. Accepted at GSI 2015.
URL https://hal.archives-ouvertes.fr/hal-01121213/document

An extended version of Chevallier et al. (2015b) has been submitted to the journal "Mathematical Morphology - Theory and Applications". The chapter "Non-parametric probability density estimation on spaces of Gaussian laws" is currently being transformed into a journal article.

Bibliography

Agueh, M., Carlier, G., Jan. 2011. Barycenters in the Wasserstein Space. SIAM Journal on Mathematical Analysis 43 (2), 904–924.
URL http://epubs.siam.org/doi/abs/10.1137/100805741

Amari, S.-i., Nagaoka, H., Harada, D., 2007. Methods of information geometry, nachdr. Edition. No. 191 in Translations of mathematical monographs. American Math. Soc. [u.a.], Providence, RI.

Angulo, J., Jul. 2007. Morphological colour operators in totally ordered lattices based on distances: Application to image filtering, enhancement and analysis. Computer Vision and Image Understanding 107 (1–2), 56–73.
URL http://www.sciencedirect.com/science/article/pii/S1077314206002165

Angulo, J., Jan. 2010. Geometric algebra colour image representations and derived total orderings for morphological operators – Part I: Colour quaternions. Journal of Visual Communication and Image Representation 21 (1), 33–48.
URL http://www.sciencedirect.com/science/article/pii/S1047320309001473

Angulo, J., Velasco-Forero, S., 2014. Morphological Processing of Univariate Gaussian Distribution-Valued Images Based on Poincaré Upper-Half Plane Representation. In: Nielsen, F. (Ed.), Geometric Theory of Information. Signals and Communication Technology. Springer International Publishing, pp. 331–366.
URL http://link.springer.com/chapter/10.1007/978-3-319-05317-2_12

Anker, J.-P., Ostellari, P., 2003. The Heat Kernel on Noncompact Symmetric Spaces. Amer. Math. Soc. Transl. Ser. 2.

Aptoula, E., Lefèvre, S., Nov. 2007. A Comparative Study on Multivariate Mathematical Morphology. Pattern Recogn. 40 (11), 2914–2929.
URL http://dx.doi.org/10.1016/j.patcog.2007.02.004

Arehart, A., Vincent, L., Kimia, B., May 1993. Mathematical morphology: The Hamilton-Jacobi connection. In: , Fourth International Conference on Computer Vision, 1993. Proceedings. pp. 215–219.

Arnaudon, M., Barbaresco, F., Yang, L., Aug. 2013. Riemannian Medians and Means With Applications to Radar Signal Processing. IEEE Journal of Selected Topics in Signal Processing 7 (4), 595–604.

Arsigny, V., Fillard, P., Pennec, X., Ayache, N., Aug. 2006. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. Magnetic Resonance in Medicine 56 (2), 411–421.

Asta, D. M., Nov. 2014. Kernel Density Estimation on Symmetric Spaces. arXiv:1411.4040 [math, stat]ArXiv: 1411.4040.
URL http://arxiv.org/abs/1411.4040

BARBARESCO, Barbaresco, F., 1996. Super Resolution Spectrum Analysis Regularization: Burg, Capon and Ago-antagonistic Algorithms. EUSIPCO-96, 2005–2008.

Barbaresco, F., 2013. Information Geometry of Covariance Matrix: Cartan-Siegel Homogeneous Bounded Domains, Mostow/Berger Fibration and Fréchet Median. In: Nielsen, F., Bhatia, R. (Eds.), Matrix Information Geometry. Springer Berlin Heidelberg, pp. 199–255.
URL http://link.springer.com/chapter/10.1007/978-3-642-30232-9_9

Barbaresco, F., Aug. 2014. Koszul Information Geometry and Souriau Geometric Temperature/Capacity of Lie Group Thermodynamics. Entropy 16 (8), 4521–4565.
URL http://www.mdpi.com/1099-4300/16/8/4521

Berger, M., 2003. A Panoramic View of Riemannian Geometry. Springer Berlin Heidelberg, Berlin, Heidelberg.
URL http://link.springer.com/10.1007/978-3-642-18245-7

Berns, R. S., Alman, D. H., Reniff, L., Snyder, G. D., Balonon-Rosen, M. R., Oct. 1991. Visual determination of suprathreshold color-difference tolerances using probit analysis. Color Research & Application 16 (5), 297–316.
URL http://onlinelibrary.wiley.com/doi/10.1002/col.5080160505/abstract

Bini, D. A., Iannazzo, B., Feb. 2013. Computing the Karcher mean of symmetric positive definite matrices. Linear Algebra and its Applications 438 (4), 1700–1710.
URL http://www.sciencedirect.com/science/article/pii/S0024379511006616

Bloch, I., 2009. Duality vs. adjunction for fuzzy mathematical morphology and general form of fuzzy erosions and dilations. Fuzzy Sets and Systems (13), 1858–1867.

Burgeth, B., Kleefeld, A., 2013. Morphology for Color Images via Loewner Order for Matrix Fields. In: Hendriks, C. L. L., Borgefors, G., Strand, R. (Eds.), Mathematical Morphology and Its Applications to Signal and Image Processing. No. 7883 in Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 243–254.
URL http://link.springer.com/chapter/10.1007/978-3-642-38294-9_21

Busch, C., Eberle, M., Aug. 1995. Morphological Operations for Color-Coded Images. Computer Graphics Forum 14 (3), 193–204.
URL http://onlinelibrary.wiley.com/doi/10.1111/j.1467-8659.1995.cgf143_0193.x/abstract

Cannon, J. W., Floyd, W. J., Kenyon, R., Parry, W. R., 1997. Hyperbolic Geometry.

Carlinet, E., Géraud, T., 2015. A Color Tree of Shapes with Illustrations on Filtering, Simplification, and Segmentation. In: Benediktsson, J. A., Chanussot, J., Najman, L., Talbot, H. (Eds.), Mathematical Morphology and Its Applications to Signal and Image Processing. No. 9082 in Lecture Notes in Computer Science. Springer International Publishing, pp. 363–374.
URL http://link.springer.com/chapter/10.1007/978-3-319-18720-4_31

Carmo, M. P. d., 1992. Riemannian Geometry. Birkhäuser.

Chanussot, J., Jan. 1998. Approches vectorielles ou marginales pour le traitement d'images multi-composantes. Chambéry.
URL http://www.theses.fr/1998CHAMS025

Chanussot, J., Lambert, P., 1998. Total Ordering Based on Space Filling Curves for Multivalued Morphology. In: Proceedings of the Fourth International Symposium on Mathematical Morphology and Its Applications to Image and Signal Processing. ISMM '98. Kluwer Academic Publishers, Norwell, MA, USA, pp. 51–58.
URL http://dl.acm.org/citation.cfm?id=295095.295110

Chevallier, E., Angulo, J., 2015. The irregularity issue of total orders on metric spaces and its consequences for mathematical morphology. Journal of Mathematical Imaging and Vision.

Chevallier, E., Chevallier, A., Angulo, J., 2015a. N-ary Mathematical Morphology. In: Benediktsson, J. A., Chanussot, J., Najman, L., Talbot, H. (Eds.), Mathematical Morphology and Its Applications to Signal and Image Processing. No. 9082 in Lecture Notes in Computer Science. Springer International Publishing, pp. 339–350.
URL http://link.springer.com/chapter/10.1007/978-3-319-18720-4_29

Chevallier, E., Farup, I., Angulo, J., Feb. 2015b. Histograms of images valued in the manifold of colours endowed with perceptual metrics. Accepted at GSI 2015.
URL https://hal.archives-ouvertes.fr/hal-01121213/document

Deng, T.-Q., Heijmans, H. J. A. M., Mar. 2002. Grey-Scale Morphology Based on Fuzzy Logic. Journal of Mathematical Imaging and Vision 16 (2), 155–171.
URL http://link.springer.com/article/10.1023/A%3A1013999431844

Duffy, B., Carr, H., Moller, T., Feb. 2013. Integrating Isosurface Statistics and Histograms. IEEE Transactions on Visualization and Computer Graphics 19 (2), 263–277.

Duin, R. P. W., 1976. On the Choice of Smoothing Parameters for Parzen Estimators of Probability Density Functions. IEEE Transactions on Computers 25 (11), 1175–1179.

F Ebner, M. D. F., 1998. Development and Testing of a Color Space (IPT) with Improved Hue Uniformity. Final Program and Proceedings - IS and T/SID Color Imaging Conference, 8–13.

Farup, I., May 2014. Hyperbolic geometry for colour metrics. Optics Express 22 (10), 12369–12378.

Fletcher, P. T., Joshi, S., 2004. Principal Geodesic Analysis on Symmetric Spaces: Statistics of Diffusion Tensors. In: Sonka, M., Kakadiaris, I. A., Kybic, J. (Eds.), Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis. No. 3117 in Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 87–98.
URL http://link.springer.com/chapter/10.1007/978-3-540-27816-0_8

Flórez-Revuelta, F., 2005. Ordering of the RGB Space with a Growing Self-organizing Network. Application to Color Mathematical Morphology. In: Duch, W., Kacprzyk, J., Oja, E., Zadrożny, S. (Eds.), Artificial Neural Networks: Biological Inspirations – ICANN

2005. No. 3696 in Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 385–390.
URL http://link.springer.com/chapter/10.1007/11550822_60

Flórez-Revuelta, F., Apr. 2015. Topology-Preserving Ordering of the RGB Space with an Evolutionary Algorithm. In: Mora, A. M., Squillero, G. (Eds.), Applications of Evolutionary Computation. No. 9028 in Lecture Notes in Computer Science. Springer International Publishing, pp. 517–528.
URL http://link.springer.com/chapter/10.1007/978-3-319-16549-3_42

Franchi, G., Angulo, J., 2015. Ordering on the Probability Simplex of Endmembers for Hyperspectral Morphological Image Processing. In: Benediktsson, J. A., Chanussot, J., Najman, L., Talbot, H. (Eds.), Mathematical Morphology and Its Applications to Signal and Image Processing. No. 9082 in Lecture Notes in Computer Science. Springer International Publishing, pp. 410–421.
URL http://link.springer.com/chapter/10.1007/978-3-319-18720-4_35

Fukunaga, K., Hostetler, L., Jan. 1975. The estimation of the gradient of a density function, with applications in pattern recognition. IEEE Transactions on Information Theory 21 (1), 32–40.

Galperin, G. A., May 1993. A concept of the mass center of a system of material points in the constant curvature spaces. Communications in Mathematical Physics 154 (1), 63–84.
URL http://link.springer.com/article/10.1007/BF02096832

Geffroy, J., 1974. Sur l'estimation d'une densité dans un espace métrique.

Gong, Y., Chuan, C. H., Xiaoyi, G., Mar. 1996. Image indexing and retrieval based on color histograms. Multimedia Tools and Applications 2 (2), 133–156.
URL http://link.springer.com/10.1007/BF00672252

Gray, A., Apr. 1974. The volume of a small geodesic ball of a Riemannian manifold. The Michigan Mathematical Journal 20 (4), 329–344.
URL http://projecteuclid.org/euclid.mmj/1029001150

Grigoryan, A., 2009. Heat Kernel and Analysis on Manifolds. American Mathematical Soc.

Gronde, J. J. v. d., Roerdink, J. B. T. M., 2013. Group-Invariant Frames for Colour Morphology. In: Hendriks, C. L. L., Borgefors, G., Strand, R. (Eds.), Mathematical Morphology and Its Applications to Signal and Image Processing. No. 7883 in Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 267–278.
URL http://link.springer.com/chapter/10.1007/978-3-642-38294-9_23

Gronde, J. J. v. d., Roerdink, J. B. T. M., 2015. Sponges for Generalized Morphology. In: Benediktsson, J. A., Chanussot, J., Najman, L., Talbot, H. (Eds.), Mathematical Morphology and Its Applications to Signal and Image Processing. No. 9082 in Lecture Notes in Computer Science. Springer International Publishing, pp. 351–362.
URL http://link.springer.com/chapter/10.1007/978-3-319-18720-4_30

Hanbury, A., Serra, J., Dec. 2001. Morphological operators on the unit circle. IEEE Transactions on Image Processing 10 (12), 1842–1850.

Heijmans, H., Ronse, C., Jun. 1990. The algebraic basis of mathematical morphology I. Dilations and erosions. Computer Vision, Graphics, and Image Processing 50 (3), 245–295.
URL http://linkinghub.elsevier.com/retrieve/pii/0734189X90901480

Heijmans, H. J. A. M., 1994. Morphological image operators. Academic Press.

Helgason, S., 1993. Geometric Analysis on Symmetric Spaces. American Mathematical Soc.

Helgason, S., Jun. 2001. Differential Geometry, Lie Groups, and Symmetric Spaces. American Mathematical Society, Providence, R.I.

Helgason, S., 2006. Non-Euclidean Analysis. In: Prékopa, A., Molnár, E. (Eds.), Non-Euclidean Geometries. No. 581 in Mathematics and Its Applications. Springer US, pp. 367–384.
URL http://link.springer.com/chapter/10.1007/0-387-29555-0_18

Helmholtz, H., 1891. Versuch einer erweiterten Anwendung des Fechnerschen Gesetzes im farbensystem. Z. Psychol. Physiol. Sinnesorg., 1–30.

Hendriks, H., 1990. Nonparametric Estimation of a Probability Density on a Riemannian Manifold Using Fourier Expansions. The Annals of Statistics 18 (2).

Henry, G., Muñoz, A., Rodriguez, D., Jun. 2011. k-Nearest neighbor density estimation on Riemannian Manifolds. arXiv:1106.4763 [math, stat]ArXiv: 1106.4763.
URL http://arxiv.org/abs/1106.4763

Huckemann, S. F., Kim, P. T., Koo, J.-Y., Munk, A., Aug. 2010. Möbius deconvolution on the hyperbolic plane with application to impedance density estimation. The Annals of Statistics 38 (4), 2465–2498.
URL http://projecteuclid.org/euclid.aos/1278861254

Ledda, R., Philips, W., 2005. W.: Majority Ordering for Colour Mathematical Morphology. In: In: Proceedings of the XIIIth European Signal Processing Conference.

Ledoux, A., Richard, N., Capelle-Laize, A.-S., Sep. 2013. How to specify or identify the most accurate multispectral distance function for mathematical morphology? In: Colour and Visual Computing Symposium (CVCS), 2013. pp. 1–7.

Lenglet, C., Rousson, M., Deriche, R., Faugeras, O., Oct. 2006. Statistics on the Manifold of Multivariate Normal Distributions: Theory and Application to Diffusion Tensor MRI Processing. Journal of Mathematical Imaging and Vision 25 (3), 423–444.
URL http://link.springer.com/10.1007/s10851-006-6897-z

Lezoray, O., Meurie, C., Elmoataz, A., Jul. 2008. Graph-based ordering scheme for color image filtering. International Journal of Image and Graphics 08 (03), 473–493.
URL http://www.worldscientific.com/doi/abs/10.1142/S0219467808003192

Lovrić, M., Min-Oo, M., Ruh, E. A., Jul. 2000. Multivariate Normal Distributions Parametrized as a Riemannian Symmetric Space. Journal of Multivariate Analysis 74 (1), 36–48.
URL http://www.sciencedirect.com/science/article/pii/S0047259X99918535

Luo, M. R., Cui, G., Rigg, B., Oct. 2001. The development of the CIE 2000 colour-difference formula: CIEDE2000. Color Research & Application 26 (5), 340–350.
URL http://onlinelibrary.wiley.com/doi/10.1002/col.1049/abstract

Luo, M. R., Rigg, B., Mar. 1986. Chromaticity-discrimination ellipses for surface colours. Color Research & Application 11 (1), 25–42.
URL http://onlinelibrary.wiley.com/doi/10.1002/col.5080110107/abstract

Macadam, D. L., May 1942. Visual Sensitivities to Color Differences in Daylight. Journal of the Optical Society of America 32 (5), 247.
URL https://www.osapublishing.org/josa/abstract.cfm?uri=josa-32-5-247

Matheron, G., 1974. Random sets and integral geometry. Wiley series in probability and mathematical statistics. Wiley, New York.

Meyer, F., 2011. Adjunctions on the lattice of hierarchies.

Monge, G., 1781. Mémoire sur la théorie des déblais et des remblais. De l'Imprimerie Royale.

Morgan, F., 2009. Geometric measure theory: a beginner's guide, 4th Edition. Academic Press/Elsevier, Amsterdam ; Burlington, MA.

Munkres, J. R., 1966. Elementary Differential Topology. Princeton University Press.

Nielsen, F., Nock, R., 2010. Hyperbolic Voronoi Diagrams Made Easy. In: 2013 13th International Conference on Computational Science and Its Applications. Vol. 0. IEEE Computer Society, Los Alamitos, CA, USA, pp. 74–80.

Paris, S., Durand, F., Jun. 2007. A Topological Approach to Hierarchical Segmentation using Mean Shift. In: IEEE Conference on Computer Vision and Pattern Recognition, 2007. CVPR '07. pp. 1–8.

Parker, T. H., Jan. 2015. Geometry Primer.

Pelletier, B., 2005. Kernel density estimation on Riemannian manifolds. Statistics & Probability Letters (3), 297–304.

Pennec, X., Jan. 2004. Probabilities and Statistics on Riemannian Manifolds : A Geometric approach. report.
URL https://hal.inria.fr/inria-00071490/document

Peter T. Kim, Donald St. P. Richards, Jan. 2011. Deconvolution Density Estimation on the Space of Positive Definite Symmetric Matrices. In: Nonparametric Statistics and Mixture Models. WORLD SCIENTIFIC, pp. 147–168.
URL http://www.worldscientific.com/doi/abs/10.1142/9789814340564_0010

Portilla, J., Simoncelli, E. P., Oct. 2000. A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients. International Journal of Computer Vision 40 (1), 49–70.
URL http://link.springer.com/article/10.1023/A%3A1026553619983

Riemann, B., 1867. Über die Hypothesen, welche der Geometrie zu Grunde liegen. Abh. Kgl. Ges. Wiss. Gött.

Rittner, L., de Alencar Lotufo, R., 2009. Segmentation of DTI based on tensorial morphological gradient. Vol. 7259. pp. 72591E–72591E–12.
URL http://dx.doi.org/10.1117/12.811754

Robertson, A. R., Mar. 1977. The CIE 1976 Color-Difference Formulae. Color Research & Application 2 (1), 7–11.
URL http://onlinelibrary.wiley.com/doi/10.1002/j.1520-6378.1977.tb00104.x/abstract

Ronse, C., Aug. 2013. Ordering Partial Partitions for Image Segmentation and Filtering: Merging, Creating and Inflating Blocks. Journal of Mathematical Imaging and Vision 49 (1), 202–233.
URL http://link.springer.com/article/10.1007/s10851-013-0455-2

Ronse, C., Agnus, V., May 2005. Morphology on Label Images: Flat-Type Operators and Connections. Journal of Mathematical Imaging and Vision 22 (2-3), 283–307.
URL http://link.springer.com/article/10.1007/s10851-005-4895-1

Sagan, H., 1994. Space-filling curves. Universitext. Springer-Verlag, New York.

Said, S., Bombrun, L., Berthoumieu, Y., Manton, J., Jul. 2015. Riemannian Gaussian Distributions on the Space of Symmetric Positive Definite Matrices. arXiv:1507.01760 [math, stat]ArXiv: 1507.01760.
URL http://arxiv.org/abs/1507.01760

Salembier, P., Serra, J., 1995. Flat zones filtering, connected operators, and filters by reconstruction. IEEE transactions on image processing: a publication of the IEEE Signal Processing Society 4 (8), 1153–1160.

Sapiro, G., Kimmel, R., Shaked, D., Kimia, B. B., Bruckstein, A. M., Sep. 1993. Implementing continuous-scale morphology via curve evolution. Pattern Recognition 26 (9), 1363–1372.
URL http://www.sciencedirect.com/science/article/pii/003132039390142J

Schrödinger, E., Jan. 1920. Grundlinien einer Theorie der Farbenmetrik im Tagessehen. Annalen der Physik 368 (21), 397–426.
URL http://onlinelibrary.wiley.com/doi/10.1002/andp.19203682102/abstract

Serra, J. (Ed.), Feb. 1988. Image Analysis and Mathematical Morphology, Vol. 2: Theoretical Advances, 1st Edition. Academic Press, London.

Serra, J. C., 1993. Anamorphoses and function lattices, 2–11.

Serra, J. P., 1982. Image analysis and mathematical morphology. Academic Press, London ; New York.

Silverman, B. W., Apr. 1986. Density Estimation for Statistics and Data Analysis. Chapman and Hall, Boca Raton.

Skovgaard, L. T., Jan. 1984. A Riemannian Geometry of the Multivariate Normal Model. Scandinavian Journal of Statistics 11 (4), 211–223.
URL http://www.jstor.org/stable/4615960

Soille, P., 2004. Morphological Image Analysis. Springer Berlin Heidelberg, Berlin, Heidelberg.
URL http://link.springer.com/10.1007/978-3-662-05088-0

Stiles, W. S., Jan. 1946. A modified Helmholtz line-element in brightness-colour space. Proceedings of the Physical Society 58 (1), 41.
URL http://iopscience.iop.org/0959-5309/58/1/305

Subbarao, R., Meer, P., Mar. 2009a. Nonlinear Mean Shift over Riemannian Manifolds. International Journal of Computer Vision 84 (1), 1–20.
URL http://link.springer.com/article/10.1007/s11263-008-0195-8

Subbarao, R., Meer, P., Mar. 2009b. Nonlinear Mean Shift over Riemannian Manifolds. International Journal of Computer Vision 84 (1), 1–20.
URL http://link.springer.com/article/10.1007/s11263-008-0195-8

Takatsu, A., 2011. Wasserstein geometry of Gaussian measures. Osaka Journal of Mathematics - OSAKA J MATH 48 (2011).

Tao, T., 2011. An Introduction to Measure Theory. American Mathematical Soc.

Terras, A., 1988. Harmonic Analysis on Symmetric Spaces and Applications II. Springer New York, New York, NY.
URL http://link.springer.com/10.1007/978-1-4612-3820-1

Trench, W. F., Sep. 1964. An Algorithm for the Inversion of Finite Toeplitz Matrices. Journal of the Society for Industrial and Applied Mathematics 12 (3), 515–522.
URL http://www.jstor.org/stable/2946327

Velasco-Forero, S., Jun. 2012. Topics in mathematical morphology for multivariate images. phdthesis, Ecole Nationale Supérieure des Mines de Paris.
URL https://pastel.archives-ouvertes.fr/pastel-00820581/document

Velasco-Forero, S., Angulo, J., Nov. 2011. Supervised Ordering in : Application to Morphological Processing of Hyperspectral Images. IEEE Transactions on Image Processing 20 (11), 3301–3308.

Velasco-Forero, S., Angulo, J., Nov. 2012. Random Projection Depth for Multivariate Mathematical Morphology. IEEE Journal of Selected Topics in Signal Processing 6 (7), 753–763.

Villani, C., 2009. Optimal Transport. Vol. 338 of Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, Berlin, Heidelberg.
URL http://link.springer.com/10.1007/978-3-540-71050-9

Wijsman, R. A., 1990. Invariant measures on groups and their use in statistics. No. v. 14 in Lecture notes-monograph series. Institute of Mathematical Statistics, Hayward, Calif.

Yang, L., 2009. Riemannian Median and Its Estimation. LMS Journal of Computation and Mathematics 13.

Zelikin, M. I., 2000. Control Theory and Optimization I Homogeneous Spaces and the Riccati Equation in the Calculus of Variations. Springer Berlin Heidelberg, Berlin, Heidelberg.
URL http://dx.doi.org/10.1007/978-3-662-04136-9

Morphology, Géométrie et Statistiques en imagerie non-standard.

Résumé : Le traitement d'images numériques a suivi l'évolution de l'électronique et de l'informatique. Il est maintenant courant de manipuler des images à valeur non pas dans $\{0, 1\}$, mais dans des variétés ou des distributions de probabilités. C'est le cas par exemple des images couleurs où de l'imagerie du tenseur de diffusion (DTI). Chaque type d'image possède ses propres structures algébriques, topologiques et géométriques. Ainsi, les techniques existantes de traitement d'image doivent être adaptés lorsqu'elles sont appliquées à de nouvelles modalités d'imagerie. Lorsque l'on manipule de nouveaux types d'espaces de valeurs, les précédents opérateurs peuvent rarement être utilisés tel quel. Même si les notions sous-jacentes ont encore un sens, un travail doit être mené afin de les exprimer dans le nouveau contexte.

Cette thèse est composée de deux parties indépendantes. La première, « Morphologie matématiques pour les images non standards », concerne l'extension de la morphologie mathématique à des cas particuliers où l'espace des valeurs de l'image ne possède pas de structure d'ordre canonique. Le chapitre 2 formalise et démontre le problème de l'irregularité des ordres totaux dans les espaces métriques. Le résultat principal de ce chapitre montre qu'étant donné un ordre total dans un espace vectoriel multidimensionnel, il existe toujours des images à valeur dans cet espace tel que les dilatations et les érosions morphologiques soient irrégulières et incohérentes. Le chapitre 3 est une tentative d'extension de la morphologie mathématique aux images à valeur dans un ensemble de labels non ordonnés.

La deuxième partie de la thèse, « Estimation de densités de probabilités dans les espaces de Riemann » concerne l'adaptation des techniques classiques d'estimation de densités non paramétriques à certaines variétés Riemanniennes. Le chapitre 5 est un travail sur les histogrammes d'images couleur dans le cadre de métriques perceptuelles. L'idée principale de ce chapitre consiste à calculer les histogrammes suivant une approximation euclidienne local de la métrique perceptuelle, et non une approximation globale comme dans les espaces perceptuels standards. Le chapitre 6 est une étude sur l'estimation de densité lorsque les données sont des lois Gaussiennes. Différentes techniques y sont analysées. Le résultat principal est l'expression de noyaux pour la métrique de Wasserstein.

Mots clés : Morphologie Mathématique, Ordres Totaux, Traitement d'Image, Estimation de Densités, Géometrie Riemannienne, Métrique de Fisher, Métrique de Wasserstein, Histogrammes d'Images.

Morphology, Geometry and Statistics in non-standard imaging.

Abstract: Digital image processing has followed the evolution of electronic and computer science. It is now current to deal with images valued not in $\{0, 1\}$ or in gray-scale, but in manifolds or probability distributions. This is for instance the case for color images or in diffusion tensor imaging (DTI). Each kind of images has its own algebraic, topological and geometric properties. Thus, existing image processing techniques have to be adapted when applied to new imaging modalities. When dealing with new kind of value spaces, former operators can rarely be used as they are. Even if the underlying notion has still a meaning, a work must be carried out in order to express it in the new context.

The thesis is composed of two independent parts. The first one, "Mathematical morphology on non-standard images", concerns the extension of mathematical morphology to specific cases where the value space of the image does not have a canonical order structure. Chapter 2 formalizes and demonstrates the irregularity issue of total orders in metric spaces. The main results states that for any total order in a multidimensional vector space, there are images for which the morphological dilations and erosions are irregular and inconsistent. Chapter 3 is an attempt to generalize morphology to images valued in a set of unordered labels.

The second part "Probability density estimation on Riemannian spaces" concerns the adaptation of standard density estimation techniques to specific Riemannian manifolds. Chapter 5 is a work on color image histograms under perceptual metrics. The main idea of this chapter consists in computing histograms using local Euclidean approximations of the perceptual metric, and not a global Euclidean approximation as in standard perceptual color spaces. Chapter 6 addresses the problem of non parametric density estimation when data lay in spaces of Gaussian laws. Different techniques are studied, an expression of kernels is provided for the Wasserstein metric.

Keywords: Mathematical Morphology, Total Orders, Image Processing, Density Estimation, Riemanninan Geometry, Fisher Metric, Wasserstein Metric, Image Histograms.

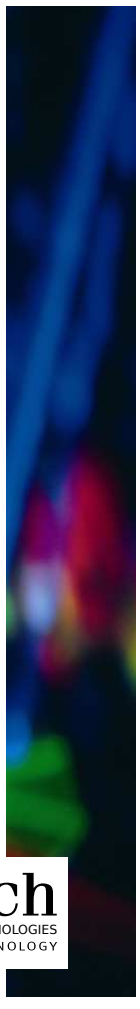

MINES ParisTech

ParisTech
INSTITUT DES SCIENCES ET TECHNOLOGIES
PARIS INSTITUTE OF TECHNOLOGY