Image processing for materials applications

Bruno Figliuzzi

INSTITUT POLYTECHNIQUE DE PARIS

Habilitation à diriger des recherches

Image processing for materials applications

Habilitation à diriger des recherches de l'Institut Polytechnique de Paris
Spécialité : Informatique, Données et IA

HDR présentée et soutenue à Paris, le 30 novembre 2023, par

BRUNO FIGLIUZZI

Composition du Jury :

Isabelle Bloch Professeur, Sorbonne Université	Présidente
Saïd Moussaoui Professeur, Centrale Nantes	Rapporteur
Nicolas Passat Professeur, Université de Reims Champagne-Ardenne	Rapporteur
Claudia Redenbach Professeur, Université de Kaiserslautern	Rapporteuse
Jocelyn Chanussot Professeur, Grenoble INP	Examinateur

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Modern materials are characterized by complex geometries at the microscopic scale. Quantifying the influence of the microstructure on the macroscopic physical and mechanical properties of a material has therefore become a major research topic in materials science. Furthermore, the exploitation of experimental images of microstructures and the increasing use of new imaging modalities like hyper-spectral imaging open interesting research perspectives in the field of image processing. I present in this Habilitation manuscript the research work that I have been conducting since 2014 at the Center for Mathematical Morphology at Mines Paris, which revolves around these questions and in particular around the following three axes:

1. The development of stochastic geometry models to characterize and simulate the geometries of microstructures,
2. The development of segmentation algorithms dedicated to material microstructures,
3. The use of hyper-spectral imaging for characterizing the surface of materials.

Previously, during my PhD thesis and my post-doctorate, I have also worked on research topics in fluid mechanics such as the flow of thin liquid films [Figliuzzi et al., 2012a,b], capillary rise [Figliuzzi and Buie, 2013; Joung et al., 2014] or non-linear electrophoresis [Figliuzzi et al., 2014, 2016a].

1.1 Morphological models and stochastic geometry

The development of morphological models aimed at describing the microstructure of heterogeneous materials is a well-established line of research at the Center for Mathematical Morphology (CMM), which offers strong potential for exploitation in the industrial sector. Stochastic geometry is a relatively mature field scientifically, but the development of innovative heterogeneous materials constantly supports the reflection on new geometric models of microstructures.

A significant part of my research work has focused on the topic of the geometry of microstructures, through several collaborations with industrials in particular. Between 2013 and 2016, I worked on a research project carried out in direct partnership with the

Japanese tire company Yokohama Rubber Co. on the modeling of the microstructure and of the visco-elastic response of a multi-scale material used to manufacture pneumatic materials. In 2015, I obtained with François Willot (CMM) a funding from the Air Force Office of Scientific Research to develop microstructure models for a copper film deposited by *coldspray*. This study combined experimental and simulation aspects. The research project was conducted in collaboration with Vincent Bortolussi as part of his PhD thesis work, conducted at the center for materials engineering of Mines Paris under the supervision of Prof. Michel Jeandin. Coldspray is a deposition process that works by projecting particles at very high speed onto a substrate at room temperature. In the case of the study, this process is used to deposit copper particles on the fuselage of an aircraft to provide protection in the event of a lightning strike. Our objective was to investigate the influence of the geometry of the substrate microstructure on the macroscopic properties of electrical conduction. We describe the development of this model, which is particularly representative of the researches that I conducted on the topic of morphological models, in chapter 2 of the manuscript.

More recently, I have been working with L'Oréal R&I on the development of new morphological models to simulate the microstructure of colloidal media or the distribution of melanosomes in the epidermis. Most of this work could not be published due to confidentiality reasons. On a scientific level, this work allowed me nevertheless to develop a method for optimizing the parameters of these models based on the MCMC [Figliuzzi et al., 2021] algorithm. I present this approach in chapter 3 of the manuscript.

1.2 Supervised segmentation

The experimental campaigns carried out in materials science often lead to the obtaining of substantial image databases, the exploitation of which must be automated as much as possible. Currently, the most efficient segmentation methods use supervised learning algorithms based on convolutional neural networks. In the case of images of microstructures, it is however rare to have enough manually segmented images to train these architectures. Techniques such as transfer learning are also insufficient to address the lack of training data, as microstructure images are often very distinct from the natural images typically used to train most existing convolutional network architectures. The development of segmentation methods using only a limited amount of data therefore remains an important axis of research, which conditions the applicability of the most efficient segmentation methods to actual industrial problems.

My main contribution to this axis of research is probably the PhD thesis of Kaïwen Chang between 2016 and 2019, during which we developed an algorithm based on the Eikonal equation to generate a superpixel segmentation of a given image based on color and texture information [Chang and Figliuzzi, 2019, 2020]. The superpixel segmentation can be thought as a first step toward the complete segmentation of the image and can be progressively merged until the obtaining of an actual segmentation of the image. By relying on carefully handcrafted features directly computed on the spatial supports defined by the superpixels, it is possible to learn a similarity measure between adjacent superpixels with a limited number of training examples. The problem of segmenting the image can thus be reformulated as a clustering problem on the region adjacency graph associated to the superpixel segmentation. To perform the clustering, we relied

on a generalization of the Eikonal equation to graph structures. I describe this research in chapter 5 of the manuscript.

In the case of relatively homogeneous images, I also sought to explore a solution to the absence of training data consisting in developing morphological models allowing the joint generation of a base of simulated images and of a ground truth in order to train a supervised segmentation algorithm. In 2020, as part of David Paulovics' internship, supervised in collaboration with F. Blanc (Nice Physics Institute), we successfully used this approach to automate the segmentation of images of suspended particles. This research is described in chapter 6.

1.3 Hyperspectral imaging

The last axis of research I am working on is hyperspectral imaging and in particular its use for the chemical characterization of surfaces. The principle of this imaging modality is to record the intensity of the light received by the sensor over a wide range of wavelengths, usually in the infrared domain. Due to the low spatial resolution of hyperspectral sensors, the observed spectra generally correspond to a combination, linear or not, of the spectra of the different elementary materials present in the observation area. The identification of these elementary spectra largely remains an open problem. In particular, the potential variability of the elementary spectra within a given observation scene, the case of highly mixed observations and non-linear mixing models remain very difficult to deal with at present.

As part of Tarek Zenati's thesis, conducted in partnership with the steel company ArcelorMittal, we developed an algorithm to detect and characterize the presence and thickness of oxide layers on steel surfaces. From an industrial point of view, it is indeed important to prevent the formation of such layers, which can significantly impact the wettability properties of steel sheets. Our approach is based on a parsimonious regression approach originally introduced for the processing of satellite images [Iordache et al., 2011], which makes it possible to bypass the issue of the strong non-linearity of the optical interactions leading to the formation of the hyperspectral measurements that we are processing [Zenati et al., 2022b]. This work is presented in the final chapter of the manuscript.

Part I

Morphological models and stochastic geometry

Understanding the impact of the microstructure on the overall physical and mechanical properties of a material has recently emerged as a prominent research area [Capasso, 2003; Jeulin, 2021; Ohser and Schladitz, 2009; Redenbach, 2009; Redenbach and Vecchio, 2011; Torquato, 2002]. One intriguing aspect of this research area involves generating random microstructures that accurately replicate the geometric features of a material [Altendorf and Jeulin, 2011; Jeulin and Moreaud, 2006; Liebscher and Redenbach, 2013; Moreaud et al., 2012; Peyrega et al., 2011]. These simulated microstructures can subsequently serve as a foundation for exploring the physical and mechanical properties of heterogeneous materials through extensive numerical simulations. This approach is particularly valuable for understanding the influence of microstructure on the macroscopic properties of the material and holds significant potential for practical applications in the industrial sector.

I describe in this part of the manuscript selected research topics that I conducted on the simulation of materials microstructures and on morphological models more generally. In chapter 2, after recalling the basic theory of morphological models, I present a research study that aimed at developing microstructure models for a copper coating deposited using the coldspray process, conducted in collaboration with Vincent Bortolussi as part of his PhD thesis at the Center for Materials engineering of Mines Paris. In this study, our objective was to examine how the geometry of the substrate microstructure affects electrical conduction at the macroscopic level. I conclude chapter 2 by presenting a research work establishing a link between the Eikonal equation and a class of tessellation models, which allows to generate random tessellations with irregular boundaries. I discuss then in chapter 3 a novel approach for the parameterization of morphological models, which aims to determine a set of parameters for the model that allows him to accurately reproduce an experimental microstructure.

Chapter 2

Multi-scale models of random microstructures

Most materials used in contemporary life and industry are heterogeneous and exhibit a complex internal microstructure. The microstructure is therefore a key feature of the global material that largely determines most of its physical properties at the macroscopic level [Torquato, 2002]. It can be interesting to generate random microstructures that reproduce accurately the geometrical characteristics of the original material, as the simulated microstructures can in turn serve as a basis to investigate the physical or mechanical properties of heterogeneous materials through extensive numerical simulations. This approach is in particular interesting to better understand the influence of the microstructure on the physical properties of the material at the macroscopic scale.

We begin in this chapter by presenting basic notions on morphological models in section 2.1, before detailing in section 2.2 research work that I have carried out on the microstructures of coatings deposited by coldspray process in collaboration with researchers from the Center for Material engineering at Mines Paris [Bortolussi et al., 2018, 2020]. This research subject is particularly representative of several studies that I have conducted on the simulation of material microstructures [Figliuzzi et al., 2016a; Gasnier et al., 2015]. We present finally in section 2.3 a research work that establishes a link between the resolution of the Eikonal equation on a domain and the generation of Voronoï and Johnson-Mehl tessellations [Figliuzzi, 2019]. This works introduces an original way to construct tessellations of space with irregular ("rough") boundaries.

2.1 Basics of morphological models

Mathematical morphology and stochastic geometry provide efficient tools for both analysis and simulation of heterogeneous microstructures. Mathematical morphology is a theory for the analysis and processing of geometrical structures based on set theory and topology [Soille, 2003]. It finds most of its applications in the field of image processing and random structures simulation. The basic idea behind mathematical morphology is to analyze a set A of some topological space E by probing it with a compact set K, classicaly referred to as *structuring element*.

The basic operators of mathematical morphology are dilation and erosion, which

are defined in the following manner. Let A be a subset of $\mathbb{R}^d$. The dilated of the set A by the structuring element K is the set

$$A \oplus K = \{x \in \mathbb{R}^d | K_x \cap A \neq \emptyset\}, \tag{2.1}$$

where K_x is the translated of the compact K at $x \in E$. Similarly, the eroded of the set A by the structuring element K is the set

$$A \ominus K = \{x \in \mathbb{R}^d | K_x \subset A\}. \tag{2.2}$$

Dilation and erosion are dual operators with respect to the complement, in the sense that dilating the set A by the structuring element K is equivalent to erode A^c by K.

Dilation and erosion can be seen as the fundamental bricks of mathematical morphology, from which derive almost all other operators. For instance, by combining erosion and dilation, we can define two new morphological operators. Let A, B be subsets of E. The *closing* A^B and the *opening* A_B of the set A by B are defined as follows:

$$A^B = (A \oplus \check{B}) \ominus B, \tag{2.3}$$

and

$$A_B = (A \ominus \check{B}) \oplus B, \tag{2.4}$$

where $\check{B} = \{x \in \mathbb{R}^d | -x \in B\}$ is the symmetric set of B.

2.1.1 Random set theory

Let us assume in this section that we dispose of a set of experimental images of some materials microstructure. Once segmented, the images of the microstructure still carry a huge amount of information. As a consequence, we need to develop mathematical techniques to capture the main features of the complex geometry under scrutiny. A natural way to describe inclusions in a matrix is for instance to consider a set A, representing the included particles, and its complementary set A^c, representing the matrix. The study of such models falls into the scope of stochastic geometry, a mathematical discipline which aims at providing a systematic description of random spatial patterns. The development of this discipline is intimately correlated to the one of mathematical morphology.

Let's go back to our previous example of inclusions in a matrix. As mentioned previously, to study the obtained sets, the idea behind mathematical morphology is to define a structuring element K (e.g. a point, or a disk with some given diameter) and to use it to probe the subsets A and A^c. The simplest relations that one can build to study A and A^c are the following ones: at a given point of the porous media, either the structuring element K hits the set A or it is disjoint from it. At each point, the knowledge of wheter or not some structuring element K hits a set A is sufficient to completely characterize A.

The deterministic approach exposed here proves however untractable in practice. Obviously, processing each one of the microstructures constituting heterogeneous media is generally out of scope, since this would involve a considerable amount of data. Nevertheless, considering the limited amount of information available on the material,

this issue can be entirely alleviated by relying on a statistical approach [Matheron, 1975]. From this perspective, instead of determining at each point if some structuring element intersects A, we try to estimate the *probability* that the considered structuring element intersects A. A benefit of this approach is that we can determine statistical laws on experimental samples and generalize these laws to larger portions of the same material as long as the samples are statistically relevant.

In this framework, the set A representing the solid phase becomes completely characterized by the functional $T(K)$ defined for all structuring element K by

$$T(K) = P(A \cap K \neq \emptyset) = 1 - P(K \cap A^c). \tag{2.5}$$

$T(K)$ is called the *Choquet capacity* of the random closed set A [Choquet, 1953]. Interestingly, the Choquet capacity is closely related to dilation and erosion operators. For all compact set $K \subset \mathbb{R}^d$, we have indeed

$$T(K) = P(K \cap A \neq \emptyset) = P(x \in A \oplus \check{K}) \tag{2.6}$$

Through the Choquet capacity, mathematical morphology provides a solid mathematical framework to investigate the microstructure geometry. Since each compact set K brings its own information of the studied set A, the choice of structuring element allows one to conduct very specific statistical measurements on the random set A. For instance, if one chooses K to be a single point, the choquet capacity yieds

$$T(\{x\}) = P(\{x\} \cap A \neq \emptyset) = P(x \in A), \tag{2.7}$$

which is the *spatial law* of the set A. Similarly, if one chooses K to be the set $\{x, x+h\}$, the choquet capacity allows to calculate the *covariance* of the random closed set.

$$T(\{x, x+h\}) = P(x \in A, x+h \in A). \tag{2.8}$$

The covariance of the set A at a given point x and for a distance h is the probability that x and $x+h$ both belong to A. The covariance C provides useful information about the spatial arrangement of the random set A. In particular, it accounts for the presence of several scales in the studied set or for periodicity.

2.1.2 Morphological models

As stated previously, one is often interested in developing stochastic models of the microstructure that reproduce accurately the geometrical characteristics of the original material.

Poisson point processes The basic ingredients of stochastic geometry models are random point processes. A random point process $\mathcal{P}$ is a collection of random points of the space $\mathbb{R}^d$. A particular role is played by Poisson point processes. Let $\theta > 0$ be a positive real number. A Poisson point process on $\mathbb{R}^d$ is a point process such that the number $N(K)$ of points contained in any region K of $\mathbb{R}^d$ is a Poisson random variable with *intensity* θ:

$$P(N(K) = n) = \frac{\theta(K)^n}{n!} \exp(-\theta(K)). \tag{2.9}$$

By adopting the point of view of mathematical morphology, we can generate a Poisson point process in a domain Ω of $\mathbb{R}^d$ by exploring Ω with a compact structuring element K and implanting a random number of points following the probability law (2.9) at each location.

Boolean model Most classical models of stochastic geometry rely on Poisson point processes. An archetypal model is for instance the Boolean model. The Boolean model is a grain model which is obtained by implanting random primary grains A' on the N germs $(x_n)_{1\leq n\leq N}$ of a Poisson points process $\mathcal{P}$ with intensity θ. Note that primary grains can possibly overlap. The resulting set A is

$$A = \bigcup_{x_n \in \mathcal{P}} A'_{x_n}, \tag{2.10}$$

where A'_{x_n} denotes the translated of the primary grain A' at point x_n:

$$A'_{x_n} = \{x_n + y, \quad y \in A'\}. \tag{2.11}$$

Any shape can be used for the grain A', including convex, non-convex or even non connected sets.

Interestingly, the Boolean model is tractable analytically. For instance, one can easily prove that the number $N(K)$ of primary grains intersected by any compact region K follows a Poisson distribution of parameter $\theta\bar{\mu}(\check{A}' \oplus K)$:

$$P(N = n) = \frac{\theta^n \bar{\mu}(\check{A}' \oplus K)^n}{n!} \exp(-\theta\bar{\mu}(\check{A}' \oplus K)). \tag{2.12}$$

In this expression, $\bar{\mu}(\check{A}' \oplus K)$ denotes the average Lebesgue measure (the average surface in $\mathbb{R}^2$ or the average volume in $\mathbb{R}^3$) of a primary grain A' dilated by the compact set K. This result guarantees that the number of primary grains in any bounded window remains almost surely finite, and enables us to calculate the Choquet capacity of the Boolean model

$$T(K) = 1 - \exp(-\theta\bar{\mu}(\check{A}' \oplus K)). \tag{2.13}$$

Knowing an analytical expression for the Choquet capacity, we can easily determine the covariance of the Boolean model that yields

$$C(h) = P(x \in A, x + h \in A) = 1 - \exp(-\theta\bar{\mu}(\check{A} \oplus l_h)), \tag{2.14}$$

where l_h denotes the structuring element constituted by both points 0 and h, h being some vector of $\mathbb{R}^d$.

Random tessellations A *tessellation* or *mosaic* is a division of the $d-$dimensional Euclidean space $\mathbb{R}^d$ into polyhedra. Such geometrical patterns can be observed in many natural situations, as shown in figure 2.1. Random tessellation models have therefore been extensively used in physics, materials science and chemistry.

Definition 1. *A tessellation in $\mathbb{R}^d$ is a countable system $\mathcal{T}$ of subsets satisfying the following conditions:*

- *$\mathcal{T} \in \mathcal{F}_{lf}(\mathbb{R}^d)$, meaning that $\mathcal{T}$ is a locally finite system of nonempty closed sets.*
- *The sets $K \in \mathcal{T}$ are compact, convex and have interior points.*
- *The sets of $\mathcal{T}$ cover the space,*

$$\cup_{K\in\mathcal{T}} K = \mathbb{R}^d \tag{2.15}$$

- *If K and K' are two sets of $\mathcal{T}$ then their interiors are disjoint.*

We denote by $\mathbb{T}$ the set of all tessellations.

The faces of a cell $\mathcal{C}$ of the tessellation are the intersections of $\mathcal{C}$ with its supporting hyperplanes. A $k-$face is a face of dimension k. Among all possible $k-$faces, the $0-$faces, or vertices, and the $1-$faces, or edges, are of particular interest. The $k-1$ dimensional faces of a $k-$dimensional polytope will be referred to as its facets.

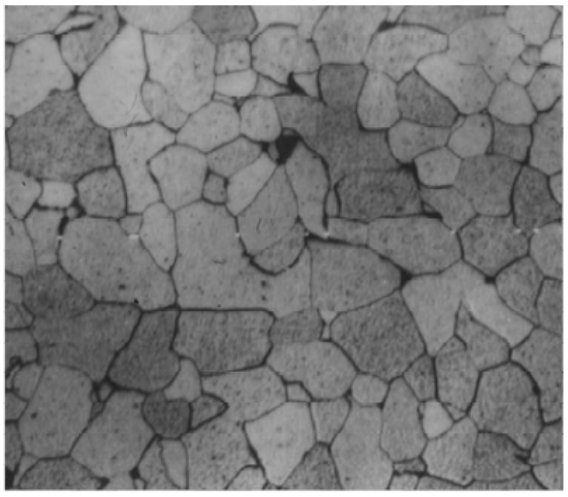

Figure 2.1: Steel polycristal microstructures are naturally organized in the form of a tessellation.

In this manuscript, we will particularly focus on one specific tessellation model, the Poisson-Voronoï model. This model has been studied extensively and is a classical model in stochastic geometry. Let Ω denote a given volume in $\mathbb{R}^3$. A Voronoï tessellation is a tessellation built from a Poisson point process $\mathcal{P}$ in the space $\mathbb{R}^3$. Every point x of $\mathbb{R}^3$ is associated to the class $\mathcal{C}_i$ containing all points of $\mathbb{R}^3$ closer from the point x_i of $\mathcal{P}$ than from any other point of $\mathcal{P}$. Hence, the classes $(C_i)_{i=1,\ldots,N}$ are defined by

$$C_i = \left\{ y \in \mathbb{R}^3, \forall j \neq i, \|x_i - y\| \leq \|x_j - y\| \right\}. \tag{2.16}$$

In practice, Voronoï tessellations are characterized by one single parameter, namely the intensity of the underlying point process. A general description of Poisson-Voronoï tessellations in $\mathbb{R}^d$ can be found in Møller [Møller, 1989, 1992, 1994].

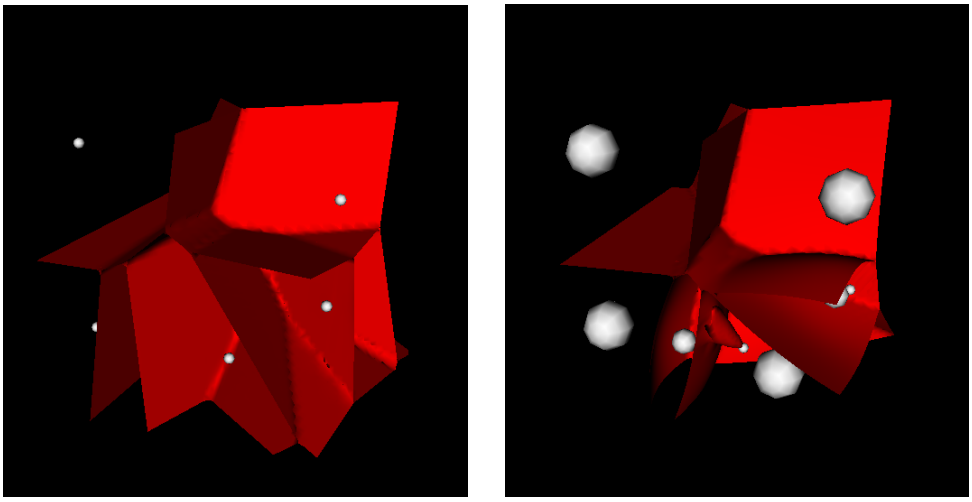

Figure 2.2: Voronoï (left) and Johnson-Mehl (right) tessellations in $\mathbb{R}^3$ generated with the software VtkSim [Faessel, 2016]. The centers of the tessellation cells are represented in grey. The first implanted germs are represented with larger radii for the Johnson-Mehl tessellation.

The last tessellation model that we briefly describe here is the Johnson-Mehl model. Johnson-Mehl tessellations can be seen as a sequential version of the Voronoï model, where the Poisson points are implanted sequentially with time. All classes grow then isotropically with the same rate, and the growth of crystal boundaries is stopped when they meet. All Poisson points falling in an existing crystal are removed. From a mathematical perspective, a Johnson-Mehl tessellation is constructed from a sequential Poisson point process where the points $(x_i)_{i=1,\dots,N}$ are implanted sequentially at a time $(t_i)_{i=1,\dots,N}$. The classes $(C_i)_{i=1,\dots,N}$ corresponding to the points $(x_i)_{i=1,\dots,N}$ are defined by

$$C_i = \left\{ y \in \mathbb{R}^3, \forall j \neq i, t_i + \frac{\|x_i - y\|}{v} \leq t_j + \frac{\|x_j - y\|}{v} \right\}. \tag{2.17}$$

Note that when all times are set to zero, we recover the classical Poisson-Voronoï tessellation model. The Johnson-Mehl tessellation model was introduced by Johnson and Mehl to describe crystallization processes [Avrami, 1939; Gilbert, 1962; Johnson and Mehl, 1939]. The paper [Møller, 1992] of Møller provides a unified exposition of Random Johnson-Mehl tessellations.

2.2 Microstructure modelling with morphological models

We present in this section a study dedicated to the development of a morphological model used for studying the electrical characteristics of a bi-phased coating produced through thermal spraying [Amsellem et al., 2008; Beauvais et al., 2008; Delloro et al., 2014]. This coating is specifically designed to shield recent commercial aircraft bodies from lightning strikes by providing a conductive surface that can effectively dissipate electrical charges. To our knowledge, this study represents one of the earliest efforts to model the electrical properties of coldspray deposits based on their microstructure. This research work was carried out in collaboration with Vincent Bortolussi (PhD student at Centre des Matériaux, Mines Paris), Michel Jeandin (Centre des matériaux, Mines

Paris), François Willot and Matthieu Faessel (Centre de morphologie mathématique, Mines Paris) and was partially funded by a grant from Air Force Office of Scientific Research (AFOSR). This study led to several publications related to the morphological model itself [Bortolussi et al., 2018] and to its exploitation to study the electric conductivity of coldspray coatings [Bortolussi et al., 2020].

2.2.1 Context

Modern commercial aircrafts are constructed using carbon fiber-reinforced polymers, which provide excellent mechanical properties while being significantly lighter than traditional aluminum alloys. However, the polymer matrix in these materials is generally highly electrically insulating, which poses security concerns in the event of lightning strikes. Advanced aerospace composites are typically fabricated using a matrix of PEEK (Poly-Ether-Ether-Ketone), a thermoplastic polymer that offers good mechanical and thermal properties while being an excellent insulator. To address the need for electrical charge dissipation in the event of lightning strikes, a layer of copper mesh is typically applied to the composite body. Copper is an easily machinable and corrosion-resistant electrical conductor with low electrical resistivity.

2.2.2 Experimental data

To simplify the manufacturing and assembly of copper meshing, a new coating method was developed by Vincent Bortolussi that relies on copper powder being thermally sprayed onto composite parts. Specifically, an adherent and electrically conductive layer is achieved by using a powder mixture containing 80% volumetric of spherical copper powder (10-35 microns) and 20% of irregular PEEK particles (35-65 microns). The mixture is applied using the cold-gas dynamic spraying or "cold spray" process.

In our work, we relied on microscopic images of the coating to study the resulting microstructure. Prior to observation, the cold spray coatings underwent cross-sectioning and polishing. To do this, the coating samples were cut in two directions: along the spraying path and orthogonal to it. Unfortunately, cutting and polishing caused debonding of the copper particles due to poor mechanical anchorage in the matrix. This resulted in dark holes at the surface. It should be noted that manual polishing can have a significant impact on the debonding phenomenon. To counteract this, the samples were metallized with a layer of Gold-Palladium that was only 2 nm thick, using a Cressington sputter coater. This was a crucial step as the layer modified the color of the PEEK matrix, greatly enhancing the color gradient between phases. To observe the cross-sections, we used a Leica optical microscope at $\times 20$ magnification in bright field. The microscope had a resolution of 0.2428 microns per pixel. We selected this observation scale to obtain a representative fraction of copper while also highlighting PEEK interstices.

The microstructure of the cold spray coating can be observed in Fig. 2.3, where yellow copper particles are embedded in a grey PEEK matrix that contains dark footprints. The matrix is made up of irregular PEEK particles that become highly deformed upon high speed impact, resulting in a dense structure with no visible pores at this scale. The copper particles deform only upon impact with each other, resulting in limited plas-

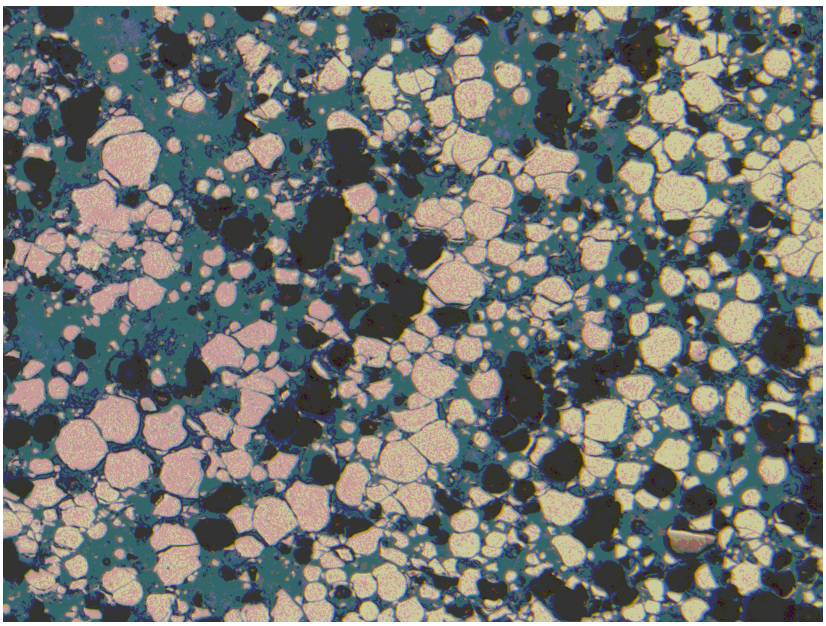

Figure 2.3: Optical microscope cross-section of the coating microstructure with debonded particles in black (2560 $\times$ 1920 pixels representing a surface area of 620 $\times$ 476 μm^2) [Bortolussi et al., 2018].

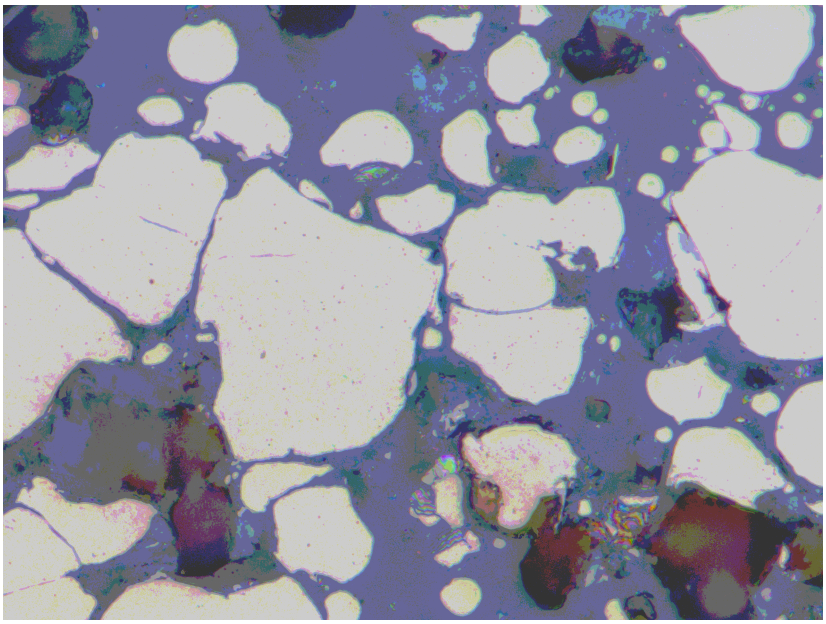

Figure 2.4: Magnified optical microscope cross-section of the microstructure of the coating (2560 $\times$ 1920 pixels representing a surface area of 248 $\times$ 186 μm^2) [Bortolussi et al., 2018].

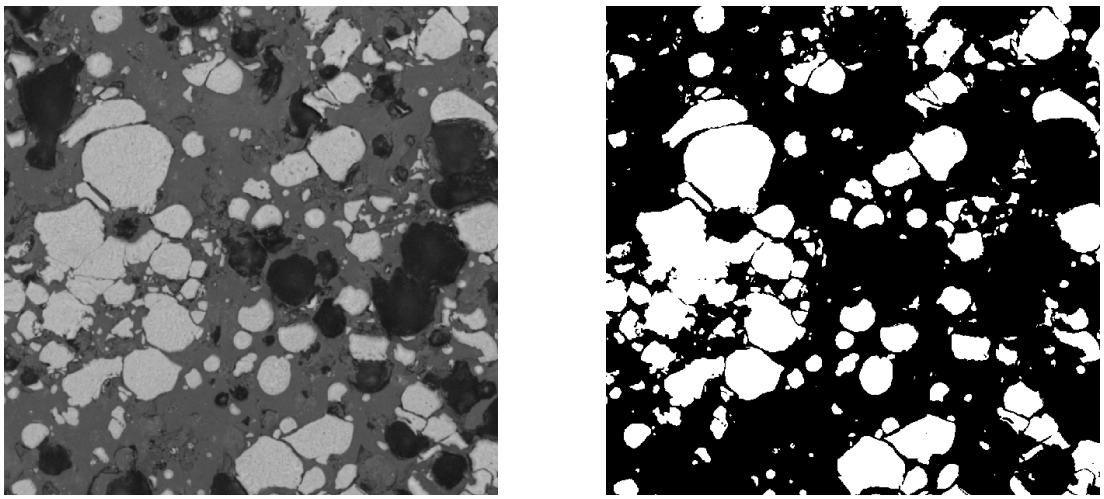

(a) (b)

Figure 2.5: Original image (left) and corresponding segmentation mask (right) [Bortolussi et al., 2018].

tic strain and forming a network of copper clusters. In typical cold spraying of metal particles onto a metal substrate, chemical bonding and inter-diffusion occur at the interfaces between particles and substrate. However, in our case, optical observation at a larger scale highlights thin PEEK layers are lying between deformed copper particles, as shown in Fig. 2.4. These interstices prevent direct contact between copper particles and are filled with electrically insulating PEEK, allegedly increasing resistivity and thus lowering the conductivity of the coating. Fig. 2.3 is an optical microscope cross-section of the coating microstructure with debonded particles appearing in black. The image comprises 2560 $\times$ 1920 pixels representing a surface area of 620 $\times$ 476 μm^2. Fig. 2.4 is a magnified optical microscope cross-section of the coating's microstructure, with the image comprising 2560 $\times$ 1920 pixels representing a surface area of 248 $\times$ 186 μm^2.

The first step of our study consisted in developing a segmentation algorithm of the experimental images in order to identify the copper particles within the microstructures as well as the interstices separating these particles. We refer the reader seeking additional details on the segmentation algorithm to the original article [Bortolussi et al., 2018]. Segmentation results are displayed in Fig. 2.5. Overall, we performed the segmentation of $K = 13$ large images of the microstructure, and we used these images to compute experimental covariance and granulometry curves that characterize the coating microstructure.

2.2.3 Morphological model

We present in this section the multi-scale morphological model [Jeulin, 2012] developed for describing the microstructure geometry. To construct the model, we make the assumption that the microstructure consists of copper spheres embedded in a PEEK matrix, separated by thin PEEK interstices. The covariance, copper fraction, granulometry, and interstice thickness are measured on 2D slices of the coating. Therefore, the parameters of the 3D model need to be inferred from this 2D information. The final

two-scale model is based on a two-step simulation process:

1. The first step of the simulation process corresponds to the generation of a Boolean model of spheres with intensity θ. The radii of the spheres follow a Gamma distribution law with parameters λ and a. This first scale of the model intends to represent the set of copper particles, which form aggregates due to interpenetration.

2. The second step involves interstice implantation based on a modified Johnson-Mehl tessellation. This two-scale model provides a representation of the 3D microstructure of the coating based on the 2D information obtained from the slices.

Aggregates simulation We assume that the copper particles aggregates can be represented by a Boolean model of spheres and that the distribution of the radii of the spheres is described by a Gamma law. The probability density function of the Gamma law is given by

$$p(r, \lambda, a) = \frac{r^{a-1}}{\Gamma(a)\lambda^a} \exp\left(-\frac{r}{\lambda}\right), \tag{2.18}$$

where Γ denotes the Gamma function. The average radius of the typical sphere is $k\lambda$. Its variance is $k\lambda^2$. The average surface of the typical grain is

$$S_v = \int_0^{+\infty} \frac{4\pi r^{a+1}}{\Gamma(a)\lambda^a} \exp\left(-\frac{r}{\lambda}\right) \mathrm{d}r = 4\pi\lambda^2 a(a+1). \tag{2.19}$$

Similarly, its average volume is

$$V_v = \int_0^{+\infty} \frac{4\pi r^{a+2}}{3\Gamma(a)\lambda^a} \exp\left(-\frac{r}{\lambda}\right) \mathrm{d}r = \frac{4\pi}{3}\pi\lambda^3 a(a+1)(a+2). \tag{2.20}$$

To determine the 3D parameters of the model using 2D measurements, we use the stereological formulae

$$\theta_3 V_v = \theta_2 \bar{A}, \quad \theta_3 S_v = \frac{4}{\pi}\theta_2 \bar{L}, \tag{2.21}$$

where $\bar{A}$ is the mean area of the sliced spheres and $\bar{L}$ their perimeter. We need to relate the 2D measurements to the parameters of the Boolean model. To that end, we rely on Miles' formulae [Chiu et al., 2013; Miles, 1972; Schneider and Weil, 2008]

$$A_a = 1 - e^{(-\theta_2 \bar{A})}, \quad L_a = \theta_2 \bar{L}(1 - A_a), \tag{2.22}$$

where A_a is the mean surface fraction of copper on segmented images and L_a is the mean perimeter of the copper phase on segmented images divided by the total surface. Using Miles' formulae in conjunction with stereological formulae [Chiu et al., 2013; Schneider and Weil, 2008], we find, for the Boolean model

$$\mathcal{A}_a = 1 - \exp(-\theta V_v), \quad \mathcal{L}_a = \theta\frac{\pi}{4} S_V \exp(-\theta V_v). \tag{2.23}$$

Overall, there are three unknowns in the model, namely the intensity θ of the Boolean model and the parameters a and λ of the Gamma distribution. Hence, we can express all parameters as functions of a. Using (2.23), we find

$$\theta = -\frac{3}{4\pi\lambda^3 a(a+1)(a+2)} \ln(1 - \mathcal{A}_a), \tag{2.24}$$

and

$$\lambda = -\frac{3\pi}{4(a+2)\mathcal{L}_a}(1-\mathcal{A}_a)\ln(1-\mathcal{A}_a). \tag{2.25}$$

To determine the parameters of the stochastic model, we rely on a maximum likehood approach to find the parameters that minimize the least-square distance between the covariance of the simulated microstructure and the covariance that is measured on the available experimental dataset. However $\mathcal{L}_a$ is highly influenced by interconnection between particles. As many particles remain in contact due to interstices segmentation, computing θ only from λ and a provides a more robust algorithm.

Interstices implantation To simulate the interstices between the particles of the same aggregate, we use a random Johnson-Mehl tessellations restricted to each aggregate, or connected component of the first scale of the microstructure. For each connected aggregate $\mathcal{A}$ of the first scale of the simulation, we consider n germs and we simulate a random number ζ_n according to an exponential law with some mean k. The cells of the tessellation are then defined by

$$C_i = \left\{ y \in \mathcal{A}, \forall j \neq i, t_i + \frac{\|x_i - y\|_G}{v} + \zeta_i \leq t_j + \frac{\|x_j - y\|_G}{v} \right\}. \tag{2.26}$$

In this relation, $\|\cdot\|_G$ denotes the geodesic distance with respect to the realization of the Boolean model. With this definition, we note that some points of the aggregates do not belong to any class of the Johnson-Mehl tessellation. We consider that these points form the interstices between the grains of the microstructure.

A significant question remains, which is how to select the initial germs of the tessellation and the germination times. While selecting the germs, our aim is to preserve the geometrical shape of the grains of the microstructure. Hence, we seek to set the germs in the center of connected components to simulate a granulary microstructure. To the end, we rely on the *h-maxima* [Serra, 1982] of the distance function to generate the germs. The *h-maxima* of the distance function form connected components. For each component, we select its barycentre to be the location of a germ. The threshold for the *h-maxima* is selected after an optimization procedure that aims at minimizing the distance between the granulometries. For each germ n, we denote by d_n the value of the Euclidean distance function at the location of the germ. The germination time associated to germ n is defined to be

$$t_n = (\max_m d_m) - d_n. \tag{2.27}$$

With this choice of germination time, the cells border of bigger spheres are kept close from the edges of the spheres and we are able to preserve the geometrical shape of the grains constituting the microstructure.

2.2.4 Maximum likehood for parameters estimation

The parameters estimation follows a two-step process. For estimating the parameters a and λ of the first scale of the model, we rely on a Nelder-Mead procedure [Nelder and Mead, 1965]. At each iteration, the parameters a and λ are fixed and θ is computed using Eq. (2.24). Then, we generate a 1000×1000 pixels slice, extracted from a larger

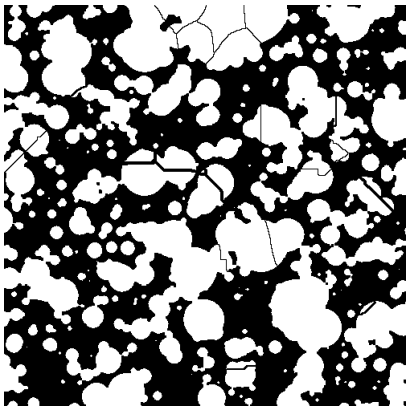

Figure 2.6: 500×500 pixels slice of a microstructure simulated with the optimized parameters. This slice is taken from a larger 3D volume of size 500×500×500 pixels. A pixel corresponds to a resolution of 0.2428 microns [Bortolussi et al., 2018].

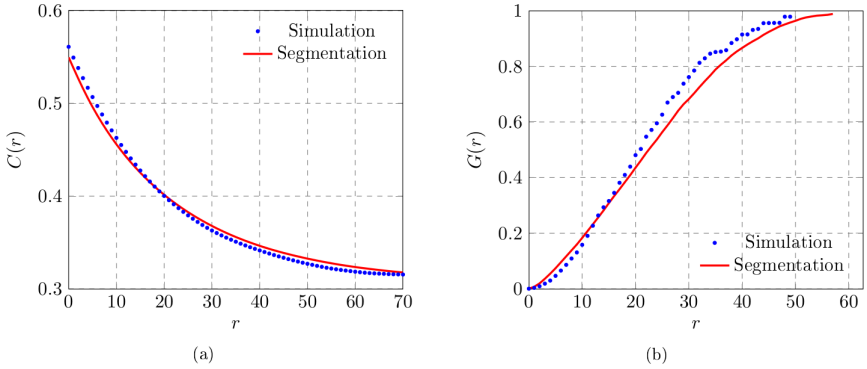

Figure 2.7: Simulated and experimental microstructures comparisons between covariances (a) and granulometries (b) [Bortolussi et al., 2018].

Aa	θ	λ	a	k	h
55.3 %	4.0×10^{-3} µm^{-3}	2.3 µm	0.8	1.2 µm	8.68 µm

Table 2.1: Model parameters

3D domain to account for border effects. The spheres are implanted within this given volume, requiring it to be large enough to reach a representative copper fraction. Finally, we measure the covariance associated to the generated slice of the material, which is compared to the experimental one. The parameters a and λ are then updated to minimize the least square distance between both covariances. It takes less than 40 iterations to fit the experimental covariance.

For the second scale of the model, we rely on a similar procedure, this time involving the granulometry measurements, to estimate the threshold value h of the h-maxima distance that is used to select the germ location and the parameter of the exponential law describing interstice thickness. Note that to allow a better fit of the covariances we slightly increase A_a to counterbalance copper surface fraction decreasing due to interstices implantation. The usual value is 57% instead of 55.3%.

2.2.5 Results and discussion

A 2D slice of the simulated microstructure obtained with the parameters estimated with the maximum likehood approach is represented in Fig. 2.6. The parameters are reported in Table 2.1, with A_a the measured surface fraction of copper on segmented images, θ the intensity of the poisson point process, λ and a the parameters of the Gamma distribution of the spheres radii and h the h-maxima of the distance function used for selecting the germs of the tesselation. The covariance and the granulometry of the simulated and of the experimental microstructures are shown in Fig. 2.7. Covariances and granulometries taken from simulated and segmented images show very good agreement, assessing the validity of the model.

The model that we developed makes it possible to reproduce accurately the microstructure of coldspray coatings containing copper particles. The model was subsequently used to optimize the microstructure of the coating with the aim of increasing its electric conductivity. This ultimately led to the manufacturing of coldspray coating with high electric conductivity, as described in [Bortolussi et al., 2020].

2.3 Eikonal based tessellations

We present in this section a research work that establishes a link between the Johnson-Mehl and Voronoï tessellation models and the Eikonal equation on a domain. A direct incentive for this study was the development of efficient algorithms for generating the tessellations used to simulate the interstices of the coldspray microstructure considered in section 2.2. We base the writing of this section on the article [Figliuzzi, 2019] published in 2019 in Image Analysis & Stereology and related to this research study. The analogy established between the Voronoï tessellation and the Eikonal equation has two direct applications: the generation of images of the tessellation in an efficient manner,

and the generation of tessellations with rough boundaries.

2.3.1 Eikonal equation

The Eikonal equation describes the propagation of a wave on an open domain Ω through the non-linear partial differential equation:

$$\forall x \in \Omega, \quad \nabla t(x) = \frac{1}{u(x)}. \tag{2.28}$$

In Eq. (2.28), ∇ denotes the gradient operator, and $t(x)$ and $u(x)$ represent the first arrival time of the wave at point x and the velocity of the wave at location x, respectively.

In order to compute the image of a Johnson-Mehl tessellation on a domain Ω, let us consider the specific version of the Eikonal equation. Given a Poisson point process $\mathcal{P}$ with realizations $(x_i)_{i=1,...,N}$ and corresponding germination times $(\tau_i)_{i=1,...,N}$, we consider the problem:

$$\begin{cases} \nabla t(x) = \dfrac{1}{u(x)} \\ t(x_i) = \tau_i, \forall i = 1, .., N. \end{cases} \tag{2.29}$$

If the velocity $u(x)$ is constant, then we can use the following proposition (see [Sethian, 1996]):

Proposition 1. *On the domain Ω, the solution t of problem (2.29) satisfies, for all $x \in \Omega$,*

$$t(x) = \inf_{i=1,...,N} (t_i + ||x - x_i||^2). \tag{2.30}$$

Proposition 1 states that solving the Eikonal equation is an efficient way to compute the Johnson-Mehl distance on a grid of points, as it precisely matches the distance used for constructing the Johnson-Mehl tessellation. Efficient algorithms are available for this task, including the fast marching algorithm described in the next section [Dejnozková and Dokládal, 2003; Sethian, 1996].

2.3.2 Fast marching algorithm

The fast marching (FM) algorithm compute the arrival times in the domain iteratively by following the propagation of the waves coming from the germs. Here, we present a slightly adapted version of the fast marching algorithm that keeps track of the labels of the initial germs during the propagation, in order to facilitate the generation of the Johnson-Mehl tessellation. We restrict ourselves to the $2D$ framework to keep notations simple. We initialize the algorithm as follows:

1. We initialize the map of *arrival times* by setting the arrival at each voxel x_i equal to $t_i = \infty$, except if the considered voxel contains a germ from $\mathcal{P}$. In this case, the arrival time at x_i is set to be the germination time τ_j.

2. We initialize a map of *labels* that keeps tracks of the wave that has arrived at each location first. We affect the label j to the voxels x_i containing one of the germs $(g_j)_{1\leq j\leq N}$ are associated the label j. We affect the label 0 to the other voxels.

3. Finally, we group all voxels containing one of the germ in a set referred to as the *narrow band*.

At each iteration, we extract the voxel (X, Y) of the narrow band with smallest arrival time and we affect it to the *frozen set*. Next, we compute the arrival times of the wave at all neighboring voxels and we add these voxels to the narrow band. At each neighbor location (x, y) of (X, Y), we compute the arrival time by solving Eikonal equation

$$||\nabla T||(x, y)u(x, y) = 1. \tag{2.31}$$

The main difficulty at this point is to propose a discretization for the gradient term. Following the work [Sethian, 1996], we discretize Eikonal equation in the following manner

$$\begin{aligned}(t_{x,y} - \min(t_{x,y+1}, t_{x,y-1}))^2 + \\ (t_{x,y} - \min(t_{x+1,y}, t_{x-1,y}))^2 = \frac{1}{u(x, y)^2}\end{aligned} \tag{2.32}$$

Equation (2.32) has two distinct solutions. To respect the consistency of the scheme, the arrival time must be higher than the time $t(X, Y)$ at the selected point in the narrow band. Therefore, we select for the arrival time $t(x, y)$ the largest solution of (2.32).

Once the arrival time $t(x, y)$ has been computed, we can encounter two distinct situations:

- When the neighbor point (x, y) is in the narrow band, it means that it has already been affected an arrival time $t^{old}(x, y)$. If $t(x, y) < t_{old}(x, y)$, then we affect the arrival time $t(x, y)$ to (x, y) as well as the label of point (X, Y). On the contrary, if $t(x, y) > t_{old}(x, y)$, the label and the arrival time at (x, y) remain unchanged.
- When the neighbor point (x, y) is not in the narrow band, we affect to it the arrival time $t(x, y)$ as well as the label of (X, Y), and we add it to the narrow band.

At each iteration of the algorithm, it is necessary to extract the element of the narrow band with the smallest arrival time. To sort the elements of the narrow band in an efficient manner, we use a binary heap to store them. We refer the reader interested by more details on the fast marching algorithm implementation to the original articles [Malladi et al., 1995; Sethian, 1996, 1999].

2.3.3 Applications

We describe in this section two potential applications of the proposed approach for generating a tessellation: the first application describes a computationally efficient method for generating an image of a Voronoï or Johnson-Mehl tessellation from a set of germs. The second application is related to the generation of tessellation with rough boundaries that it is made possible by the use of Eikonal equation.

Generation of images from vectorial simulations

For materials engineering applications, it is often necessary to simulate images of a generated Voronoï or Johnson-Mehl tessellation. In practice, the tessellation is entirely characterized by the location of its germs and by the corresponding implantation times in the case of the Johnson-Mehl tessellation. The image generation is achieved by discretizing the domain Ω and by evaluating the implicit functions associated with the tessellation's cells at each voxel location. However, this approach can be computationally expensive, particularly when high resolution is required. If we denote the number of voxels used to discretize Ω as P, the average number of Poisson points in $\mathcal{P}$ is proportional to P. As a result, the image generation algorithm's complexity is $O(P^2)$, which can represent a significant amount of computation when using this method for a vectorial simulation. By contrast, using the approach based on the Eikonal equation yields a complexity in $O(N \log N)$, which is particularly appealing for large values of N.

Tessellation with rough boundaries

We have considered so far a constant velocity field u in the Eikonal equation (2.29) for constructing the Johnson-Mehl tessellation. However, when using the Eikonal equation to build the mosaic, it is possible to use any velocity fields on the domain Ω, as long as $u(x) > 0$ for all $x \in \Omega$. This observation opens the way to the development of tessellation models that cannot be directly obtained with classical approaches including the Voronoï and the Johnson-Mehl tessellations described previously. A potential application is for instance the generation of tessellations of space with rough boundaries between adjacent cells.

Let $\mathcal{P}$ be the realization of a marked point process with intensity θ on an open domain Ω, for which the marks are drawn according to the uniform distribution $\mathcal{U}([0, L])$ on an interval $[0, L] \subset \mathbb{R}_+$. The points in $\mathcal{P}$ as well as their respective marks allow to compute a Johnson-Mehl tessellation on the domain Ω by solving problem (2.29). We construct a random velocity field u on domain Ω by

1. computing a Voronoï tessellation of Ω based upon a Poisson point process $\mathcal{P}_v$ with intensity θ_v, and

2. setting a random velocity drawn according to some distribution $\mathcal{D}$ in each cell of the Voronoï tessellation.

By solving the problem

$$\begin{cases} \nabla t(x) = \dfrac{1}{u(x)} \\ t(x_i) = \tau_i, \forall i = 1, .., N \end{cases} \tag{2.33}$$

we obtain a tessellation of the domain Ω ressembling the Johnson-Mehl tessellation generated by solving (2.29), but with rough boundaries between adjacent cells. It is worth noting that the roughness observed at the boundaries is closely linked to the characteristic length of the fluctuations in the velocity field. This characteristic length is dependent on the intensity of the Voronoï tessellation, denoted by θ_v, which is used

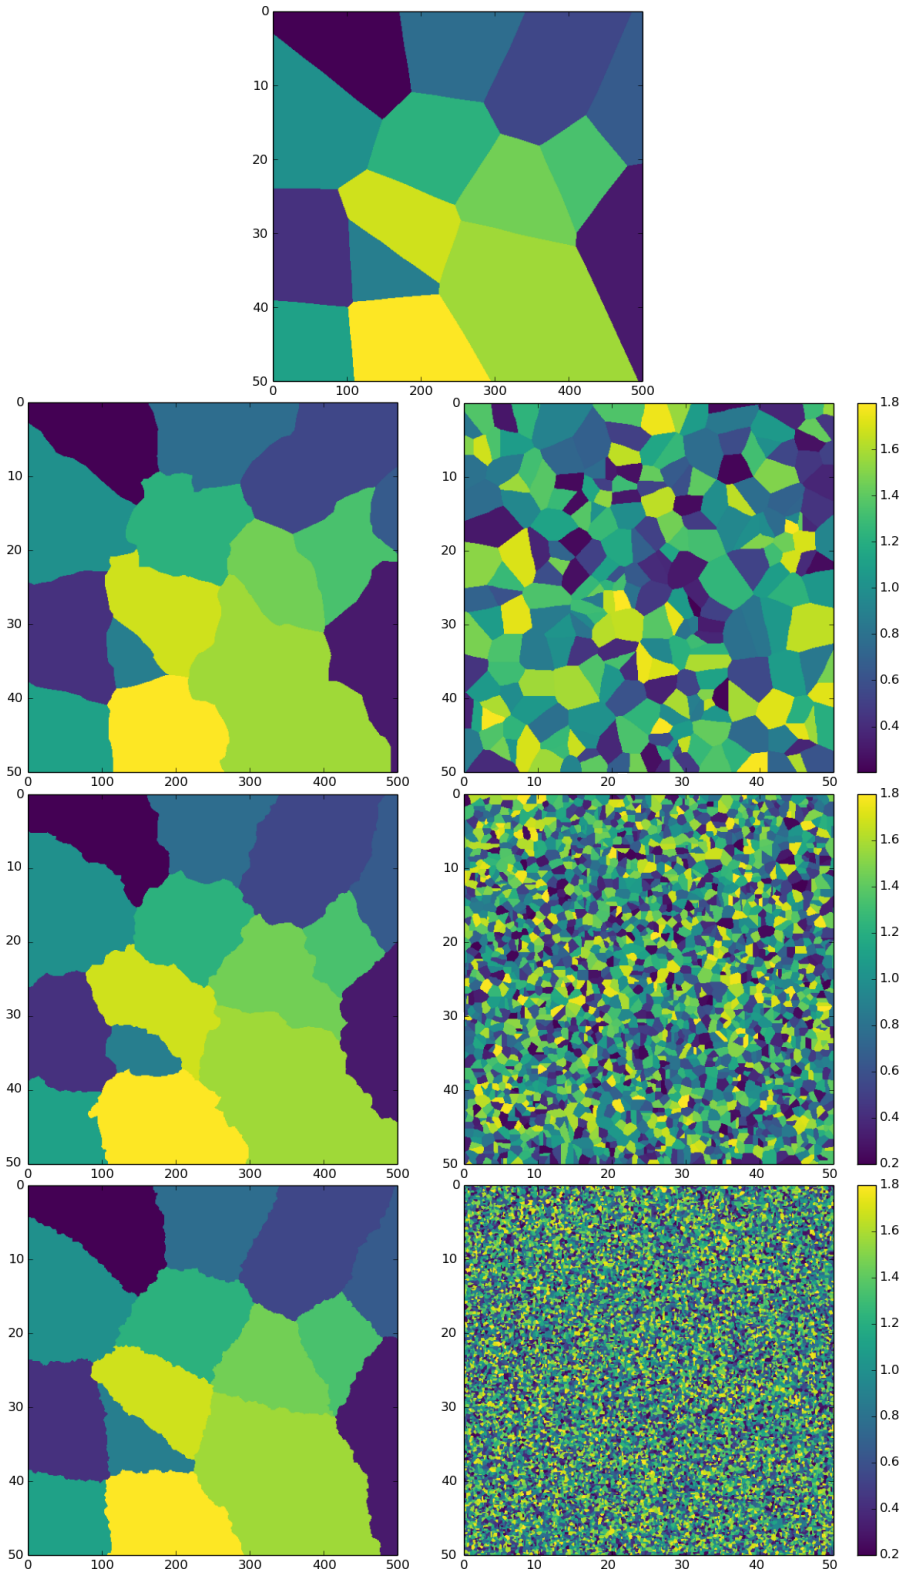

Figure 2.8: Rough Voronoï tessellations (left) and corresponding local velocity field (right). The tessellations were generated with the same set of germs $\mathcal{P}$ on a 50×50 domain, discretized on a 500×500 regular grid. The local velocity fields are realizations of Voronoï tessellations with respective intensities $\theta = 0$, $\theta = 0.1$, $\theta = 1$ and $\theta = 10$ [Figliuzzi, 2019].

to compute the velocity field. A larger value of θ_v results in a greater number of smaller Voronoï cells, leading to velocity variations over shorter distances. As a result, the generated Johnson-Mehl tessellation displays boundaries with a high concentration of roughness patterns with relatively small amplitudes. In contrast, Johnson-Mehl tessellations generated from velocity fields constructed with a smaller value of θ_v exhibit boundaries with a lower density of roughness patterns with larger amplitudes. In the original article [Figliuzzi, 2019], we developed a multi-scale tortuosity descriptor, which is able to characterize the tortuosity of the boundary between two adjacent cells by iteratively computing the tortuosity of smoothened versions of the boundary. This descriptor allowed us to establish a direct correlation between the value of θ_v and the tortuosity of the boundaries at distinct scales. We refer the reader interested by additional details to the original article [Figliuzzi, 2019].

To illustrate our algorithm, we present in figure 2.8 simulations of Voronoï tessellations with rough boundaries computed with different values of velocity parameter θ_v. In this figure, we generated four distinct Voronoï tessellations on a domain with size 50×50 discretized on a 500×500 regular grid. To generate the tessellation, we started from the same realization of a Poisson point process $\mathcal{P}$ with intensity $\theta = 5 \times 10^{-2}$. We generated a random velocity field for each one of the tessellations by first constructing a Voronoï tessellation of the domain, and then selecting a random velocity drawn from the uniform distribution $\mathcal{U}$ on the interval $[0.2, 1.8]$ for each cell of the tessellation. The characteristic length λ_v of the velocity variation can be related to the intensity θ_v of the point process used to construct the germs of the tessellation through the relationship

$$\lambda_v \simeq \sqrt{\frac{1}{\theta_v}}.$$

We note that when λ_v is small, it leads to the obtaining of small roughness patterns at the cell boundaries. Conversely, large values of λ_v lead to larger sizes of roughness pattern. The proposed method therefore allows to control, at least qualitatively, the amplitude and the size of the roughness patterns observed at the boundaries.

2.4 Conclusion

I discussed in this chapter several morphological models used to describe the microstructure of cold spray coatings or to generate tessellations of space with rough boundaries. These models share the characteristic of being described by a small number of parameters that are fixed in order to match statistical features estimated from realizations of the model with features computed from experimental images of the microstructure. In the next chapter of the manuscript, we present a general methodology for selecting these parameters based on the Monte Carlo Markov Chains (MCMC) algorithm.

2.5 Related works

Microstructure models for rubber materials Between 2013 and 2016, I worked on a research project carried out in direct partnership with the Japanese tire company Yokohama Rubber Co. on the modeling of the microstructure and of the visco-elastic response of a material used to manufacture pneumatic materials. This project was conducted in collaboration with Dominique Jeulin, Franccois Willot and Matthieu Faessel (Center for Mathematical Morphology, Mines Paris).

Rubber are classically constituted by a polymer matrix filled by black carbon particles. The physical properties of rubber are therefore strongly influenced by the volume fraction and spatial distribution of fillers [El Moumen et al., 2013, 2014; Jean et al., 2011]. In particular, when nano-sized fillers form extended percolating networks, they have a significant impact on the macroscopic mechanical properties of rubbers. To predict the mechanical properties of rubber materials, we combined the use of morphological models that describe the multi-scale microstructure of rubber with efficient Fourier-based numerical algorithms that allow to compute the effective visco-elastic properties of the material based on an accurate description of its microstructure. The morphological models were developed from experimental transmission microscope images of the microstructure and take into account the non-uniform distribution of fillers observed within these images.

In this study, we proposed various types of multi-scale models to describe the observed microstructure. Specifically, aggregates of black carbon particles were positioned outside of an exclusion polymer, simulated through two distinct approaches. The first approach utilized a Boolean model of spheres, while the second model employed a mosaic model constructed from a Johnson-Mehl tessellation. The aggregates were then characterized as the intersection between a Boolean model of spheres and the complementary region of the exclusion polymer. We simulated the individual carbon black particles within the aggregates by using a Cox-Boolean model of spheres. The multi-scale model was parameterized based on the experimental covariance and on the cumulative granulometry measured on the experimental images.

The morphological model was exploited by François Willot along with a novel Fourier-based algorithm to estimate the viscoelastic properties of linear heterogeneous media under harmonic conditions, which demonstrated that the proposed numerical method is effective in calculating the viscoelastic response of microstructures containing rubbers and fillers. This work was published in 2016 in *Technische Mechanik* [Figliuzzi et al., 2016a] and presented at the *European Conference on Constitutive Models for Rubbers X* in 2017 [Koishi et al., 2017].

Microstructure models for pressed explosive materials In [Gasnier et al., 2015], we developed a numerical approach to homogenization specifically designed for pressed explosives based on an energetic molecular crystal referred to as TATB (1,3,5-triamino-2,4,6-trinitrobenzene). Our approach combined virtual microstructure modeling and Fourier-based computations to compute effective thermoelastic properties for the material. In particular, we used a Johnson-Mehl tessellation model with anisotropic grain growth to simulate the microstructure of the material. Extensive Fourier computations were conducted on microstructures generated with the model to determine effective thermoelastic properties for the heterogeneous material, which demonstrated excellent agreement with experimental data, particularly in relation to the overall thermal expansion coefficient, and adhere to established bounds for polycrystalline anisotropic thermoelasticity.

2.6 Related publications

- B. Figliuzzi. Eikonal-based models of random tessellations. *Image Analysis and Stereology*, 38(1), 15-23, 2019.

- V. Bortolussi, B. Figliuzzi, F. Willot, M. Faessel, M. Jeandin. Morphological modeling of cold spray coatings. *Image Analysis and Stereology*, 37(2), 145-158, 2018.
- M. Koishi, N. Kowatari, B. Figliuzzi, M. Faessel, F. Willot, D. Jeulin. Computational material design of filled rubbers using multi-objective design exploration. In *Constitutive Models for Rubber X* (pp. 467-473). CRC Press, 2017.
- V. Bortolussi, F. Borit, A. Chesnaud, M. Jeandin, M. Faessel, B. Figliuzzi, G. Surdon. Cold spray of metal-polymer composite coatings onto carbon fiber-reinforced polymer (CFRP). In *International Thermal Spray Conference*, 2016.
- B. Figliuzzi, D. Jeulin, M. Faessel, F. Willot, M. Koishi, N. Kowatari. Modelling the microstructure and the viscoelastic behaviour of carbon black filled rubber materials from 3D simulations. *Technische Mechanik*, 32(1-2), 22-46, 2016.
- J. B. Gasnier, F. Willot, H. Trumel, B. Figliuzzi, D. Jeulin, M. Biessy. A Fourier-based numerical homogenization tool for an explosive material. *Matériaux et Techniques*, 103(3), 308, 2015.

Chapter 3

Morphological models parameterization

We presented in the previous chapter different morphological models allowing to reproduce the microstructure of materials from experimental imaging data. Typically, these morphological models rely on a set of parameters Θ that must be carefully selected to align with experimental observations. One common approach to determine these parameters is by minimizing an objective function, often representing the difference between measurements obtained from simulations and actual experimental data. Numerous methods have been proposed in literature for this purpose.

A classic technique involves minimizing the objective function using a gradient descent algorithm. Gradient descent algorithms are iterative optimization algorithms of the first order that aim to locate local minima of a differentiable function by progressively advancing along steps in the direction of the negative gradient in the parameter space. These algorithms are highly effective for minimizing strongly convex functions, but applying them to the problem at hand presents several challenges. Firstly, for morphological models, computing the gradient of the objective function is often analytically intractable and necessitates the use of finite difference schemes to evaluate the gradient. This process requires computing at least $D + 1$ evaluations of the objective function, where D represents the dimension of the parameter space. Moreover, the inherent statistical variability of morphological models often results in noisy gradient estimates. Secondly, the objective function is often non-convex, implying that a gradient descent algorithm will converge only to a local minimum of the objective function.

Alternatives to the gradient descent algorithm exist in the literature. Notably, in [Wang et al., 2015, 2017, 2018] and [Figliuzzi et al., 2016a], the authors suggest employing the Nelder-Mead algorithm for the optimization, an heuristic approach used to find the minimum or maximum of an objective function in a multidimensional space. It is a direct search method that relies solely on iterative evaluations of the objective function, making it suitable for cases where analytical computation of the gradient is not feasible.

In this chapter of the manuscript, I present researches that I conducted on the relatively overlooked question of morphological models parameterization. This study was partly inspired by issues encountered in the context of a research project with

L'Oréal R&I, and was published in [Figliuzzi et al., 2021]. In this study, we introduce an approach based on a Bayesian formulation of the problem. This allows us to come up with a posterior distribution model for the parameters of the microstructure model conditioned on the experimental observations. To generate samples from the posterior distribution, we employ a Monte Carlo Markov Chains (MCMC) algorithm. This approach offers several advantages compared to other methods. Specifically, the MCMC algorithm produces samples that accurately represent the actual posterior distribution, enabling the detection of parameter correlations within the model and quantification of their respective influences.

3.1 Mathematical model

3.1.1 Problem statement

Let us assume that we dispose of a set of previously segmented experimental images of the microstructure of a material, referred to as the observations $\mathcal{O}$. We can then characterize the geometry of the microstructure by computing a number of statistical features including the covariance, the granulometry or the granulometry of the complementary image. The covariance of a random set A was defined in chapter 2. It is the function C_A defined on $\mathbb{R}^d \times \mathbb{R}^d$ by

$$C_A(x, x+h) = P(x \in A, x+h \in A), \tag{3.1}$$

where h is some vector of $\mathbb{R}^d$. For a stationary random set, the covariance is a function of the distance h only:

$$C_A(x, x+h) = C_A(h). \tag{3.2}$$

The granulometry is an other statistical characteristics of the microstructure, constructed as follows. Let K be a convex set. We consider the family of structuring elements $(K_\lambda, \lambda > 0)$, where $K_\lambda = \lambda K$. The operator defined for all closed set $A \subset \mathbb{R}^d$ by

$$\Phi_\lambda(A) = (A \ominus \check{K_\lambda}) \oplus K_\lambda, \tag{3.3}$$

where $\ominus$ is the morphological erosion and $\oplus$ the morphological dilation, is a granulometry. The granulometry by openings of a random set A provides a characterization of the size distribution of the elements of A. The random set A can be characterized by recording, for distinct values of λ, the volume fraction of the residual set $A \backslash \Phi_\lambda(A)$. Granulometry measurements can be computed on the random set A, as well as on its complementary set A^c.

In the remainder of this chapter, we will denote by $m_\mathcal{O}$ and m_Θ the statistical measurements conducted on the observations and the statistical measurements conducted on simulated microstructure, respectively. The parameters Θ of the model are typically selected to minimize the difference between $m_\mathcal{O}$ and m_Θ. The determination of the parameters of morphological model can therefore be formulated as the following constrained optimization problem:

$$\hat{\Theta} \quad := \quad \arg\min_{\Theta} \|m_\mathcal{O} - m_\Theta\|_2^2 \text{ subject to } \Theta \in \mathcal{D}. \tag{3.4}$$

In (3.4), $\mathcal{D}$ denotes the set of admissible parameters for the model.

3.1.2 Bayesian formulation

Starting from measurements obtained from a morphological model and from experimental observations of the studied microstructure, we use a Bayesian approach to determine the parameters of the model. The likelihood that the observations were obtained with the set of parameters Θ can be defined as follows:

$$p(\mathcal{O}|\Theta) = \frac{1}{\sqrt{2\pi|\Lambda|}} \exp((m_{\mathcal{O}} - m_{\Theta})^T \Lambda^{-1} (m_{\mathcal{O}} - m_{\Theta})), \tag{3.5}$$

In Eq. (3.5), we usually consider a diagonal covariance matrix Λ and we assign a weight to the parameters corresponding to their respective importance. The arbitrary choice of a Gaussian distribution is motivated by its simplicity, but other distributions could potentially be considered for the likelihood.

We resort to physical considerations to define a prior distribution on the parameters Θ. Usually, we are able to define a range of variation for all parameters of the model based on their physical validity:

$$\Theta_{\min} \preceq \Theta \preceq \Theta_{\max}, \tag{3.6}$$

where $\preceq$ is a component-wise inequality. We simply consider a uniform distribution between the lower and the upper bound for each parameter as prior distribution:

$$p(\Theta) := \mathcal{U}_{\Theta_{\min}, \Theta_{\max}}. \tag{3.7}$$

Thanks to Bayes formula, we obtain an expression for the posterior distribution of the parameters Θ knowing the observations:

$$p(\Theta|\mathcal{O})p(\mathcal{O}) = p(\mathcal{O}|\Theta)p(\Theta). \tag{3.8}$$

Since the (unknown) probability law for the observation is independent from Θ, the posterior distribution is proportional to the product of the likelihood and of the prior distribution:

$$p(\Theta|\mathcal{O}) \propto p(\mathcal{O}|\Theta)p(\Theta). \tag{3.9}$$

The problem of the model parameterization becomes the determination of a set of parameters maximizing the posterior probability $p(\Theta|\mathcal{O})$:

$$\hat{\Theta} = \text{argmax}_{\Theta} p(\Theta|\mathcal{O}). \tag{3.10}$$

3.1.3 Sampling from the posterior distribution

A question now arises, which is how to solve the optimization problem from Eq. (3.10) in practice? An obvious solution could be to use a gradient ascent algorithm to find local maxima in the parameters space. In this section, we propose a distinct approach consisting in using a Monte Carlo Markov Chains (MCMC) algorithm to generate samples from the posterior distribution [Andrieu et al., 2003; Robert et al., 1999]. This approach offers several benefits compared to gradient ascent. Computing the gradient of the measurements with respect to the parameters is often challenging due to two

primary factors. Firstly, there is usually no readily available analytical formula to calculate the gradient. Secondly, the function m_Θ is a stochastic function dependent on the parameter set Θ, making it difficult to estimate the gradient using a finite difference scheme. Moreover, the MCMC approach possesses an additional advantage as it enables the generation of samples that accurately represent the posterior distribution. This capability allows for the capture of noteworthy characteristics of the microstructure, including parameter correlations and characteristic ranges of parameter variation. The algorithm that we proposed works as follows:

- *Initialization:* At initialization, we first generate a set of parameters from the prior distribution:
$$\Theta_1 \sim \mathcal{U}_{\Theta_{\min},\Theta_{\max}}. \tag{3.11}$$
- *n-th iteration:* At each iteration, we repeat the following steps:
  1. *Parameters sampling* we generate a set of parameters $\hat{\Theta}$ from the current state Θ_n according to a proposal distribution q:
  $$\hat{\Theta} \sim q(\hat{\Theta}|\Theta_n) \tag{3.12}$$
  2. *Simulation* We compute a simulation of the morphological model with this set of parameters, and perform measurements on the computed microstructure. The measurements $m_{\hat{\Theta}}$ are used to compute the posterior distribution, up to the constant factor $p(\mathcal{O})$:
  $$p(\hat{\Theta}|\mathcal{O}) \propto p(\mathcal{O}|\hat{\Theta})p(\hat{\Theta}). \tag{3.13}$$
  3. *Accept/reject step* We compute the so-called Hastings ratio r, defined by:
  $$r = \min\left(1, \frac{p(\hat{\Theta}|\mathcal{O})q(\Theta_n|\hat{\Theta})}{p(\Theta_n|\mathcal{O})q(\hat{\Theta}|\Theta_n)}\right) \tag{3.14}$$
  Noteworthy, the Hastings ratio r is independent of the quantity $p(\mathcal{O})$. The new set of parameters $\hat{\Theta}$ is accepted or rejected with probability r:
  $$\Theta_{n+1} = \begin{cases} \hat{\Theta} & \text{if } u < r \\ \Theta_n & \text{otherwise} \end{cases} \tag{3.15}$$
  where u is a random variable sampled from the uniform distribution $\mathcal{U}(0,1)$.

3.2 Numerical experiments

3.2.1 Experiments

To illustrate our methodology, we consider two examples of microstructures generated with morphological models. Our objective is to recover the parameterization of these models from statistical measurements conducted on realizations of the models. To that end, for each example, we generate a set of $N = 10$ random microstructures in a domain of size 30 by 30 in $\mathbb{R}^2$.

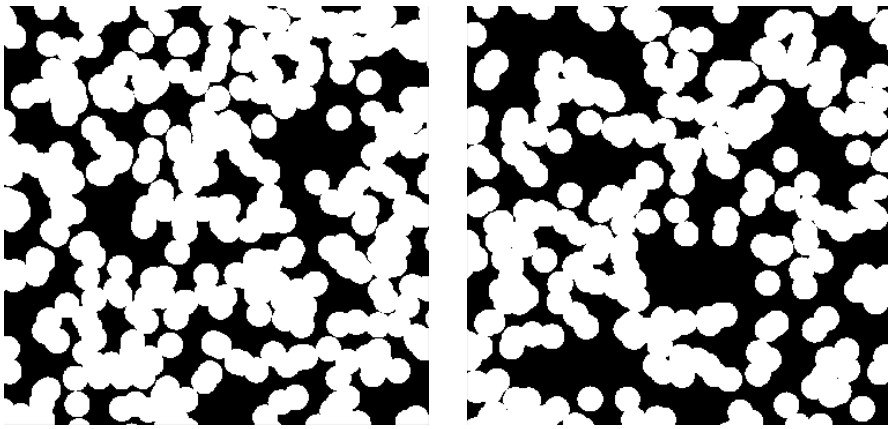

Figure 3.1: Sample of the microstructure generated with the Boolean model of disks (left) and with the Cox-Boolean model (right) [Figliuzzi et al., 2021].

- The first model corresponds to a Boolean model of disks with constant radii $R = 0.5$ and with intensity $\theta = 0.45$.
- The second model corresponds to a Cox-Boolean model of disks corresponding to a two scales microstructure, with the first scale consisting in exclusion zones that remain empty of any inclusions, modeled by a Boolean model of disk with constant radii $R_e = 1$ and intensity $\theta_e = 0.1$. The second scale simulates inclusions located in the complementary of the exclusion zones. The inclusions are described by a Boolean model of disks with intensity $\theta = 0.45$ and radii sampled according to a normal distribution with mean $R = 0.5$ and standard deviation 0.1.

Realizations of both morphological models are displayed in Fig. 3.1. To obtain a set of measurements $m_{\mathcal{O}}$ for each morphological model, we averaged the covariance, the granulometry and the granulometry of the complementary as computed on 10 model realizations.

The upper and lower bounds for the parameters used to define the prior distribution are indicated in table 3.1 for the Boolean model and in table 3.2 for the Cox-Boolean model.

We consider the same diagonal covariance matrix Λ in the likelihood function (3.5) for both morphological models. To put more emphasis on the correlations at the smallest scales, we employ the following expression for the diagonal coefficients of Λ:

$$\Lambda_{n,n} = \lambda\left(1 + \frac{1}{n}\right)^{-1}, \tag{3.16}$$

where $\lambda = 0.1$. Next, we initialize the Metropolis-Hastings algorithm by sampling a first set of parameters $\Theta_1 := (\theta_1, R_1)$ from the prior distribution. At each step of the algorithm, the proposal distribution for each parameter is the truncated normal distribution:

$$q(\theta|\theta_n) = \mathcal{N}_{tr}(\theta_n, \sigma_\theta, \theta_{min}, \theta_{max}), \tag{3.17}$$

where θ_{min} (resp. θ_{max}) is the minimal (resp. maximal) possible value of the parameter

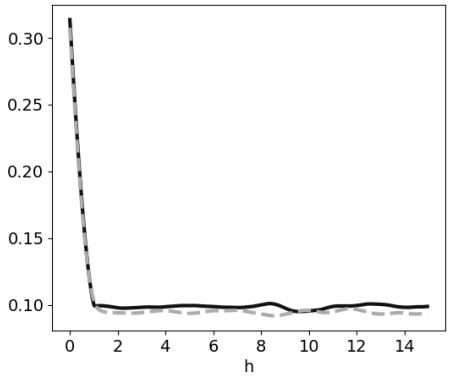

(a) Covariance

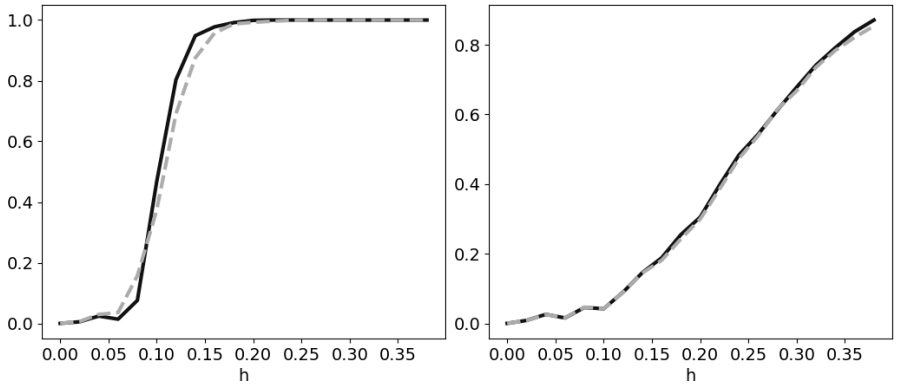

(b) Granulometry measurements of the microstructure (left) and of the complementary (right)

Figure 3.2: Comparison between the experimental measurements and the measurements obtained with simulations conducted with the MAP parameters identified by the MCMC algorithm for the Boolean model of disks [Figliuzzi et al., 2021].

as established for defining the prior distribution, and the standard deviation σ is set to:

$$\sigma_\theta = \frac{\theta_{max} - \theta_{min}}{\Delta}. \tag{3.18}$$

We fix $\Delta = 20$ to get an acceptation rate of 0.3 in the algorithm.

We stop the algorithm after 500 iterations, and we discard the first 100 parameters sampled by the algorithm, as they remain highly dependent on the initial conditions. Finally, we estimate the set of optimal parameters for the microstructure by considering the maximal a posteriori (MAP) value found for the sampled parameters:

$$\hat{\Theta} := \arg \max_{n=100,\dots,500} p(\Theta_n|\mathcal{O}). \tag{3.19}$$

Results and discussion

The optimal set of parameters $\hat{\Theta}$ found by the Metropolis-Hastings algorithm are displayed in Tab. 3.1 for the Boolean model of disks, along with the standard deviation of the sampled parameters, and in Tab. 3.2 for the Cox-Boolean model. The standard deviation provides valuable information about the sensitivity of the parameters: if a

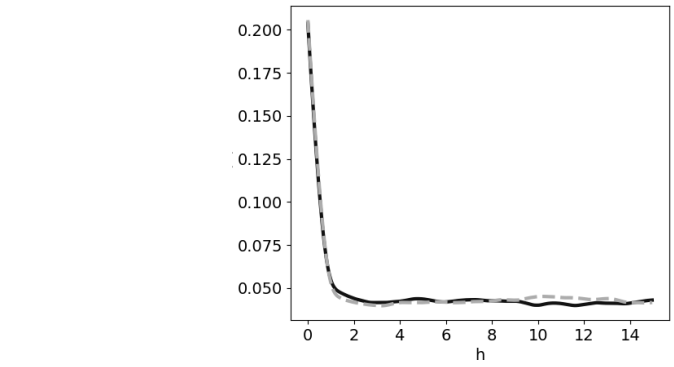

(a) Covariance

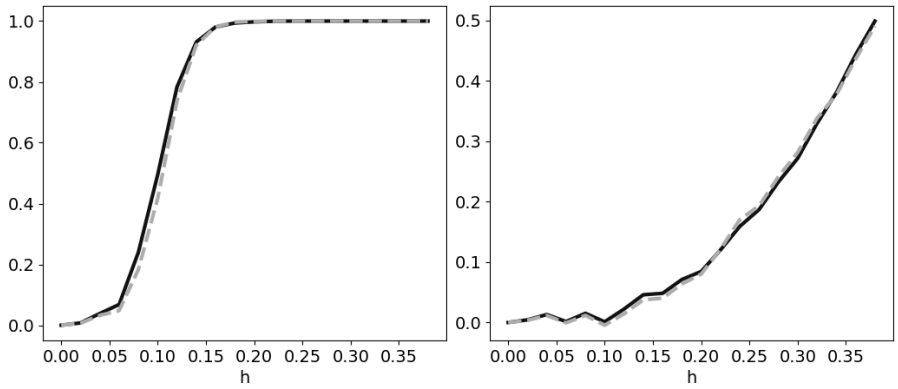

(b) Granulometry measurements of the microstructure (left) and of the complementary (right)

Figure 3.3: Comparison between the experimental measurements and the measurements obtained with simulations conducted with the MAP parameters identified by the MCMC algorithm for the Cox-Boolean model of disks [Figliuzzi et al., 2021].

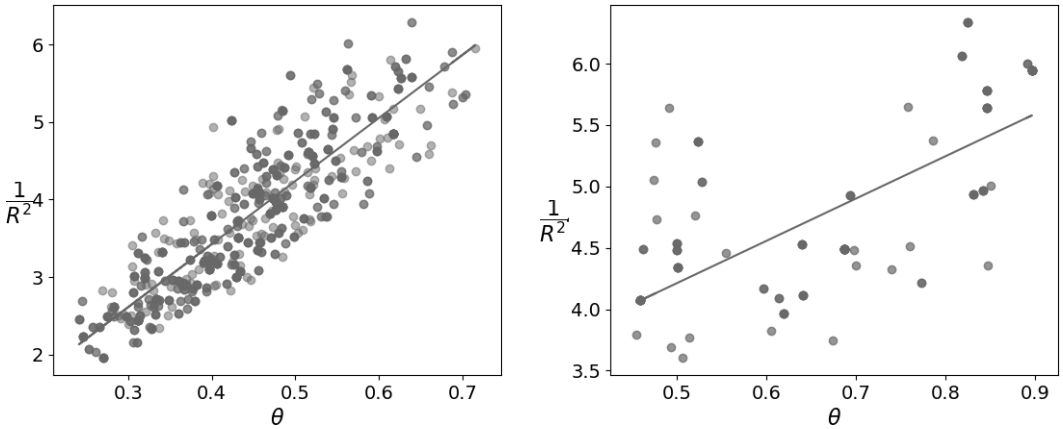

Figure 3.4: Scatterplot of the parameters $1/R^2$ and θ sampled by the Metropolis-Hastings algorithm for the Boolean model of disks (left) and for the Cox-Boolean model (right). The color intensity of each point corresponds to the frequency of the corresponding parameters within the samples [Figliuzzi et al., 2021].

	θ $[L^{-1}]$	R $[L]$
Target	0.45	0.5
$\hat{\Theta}$	0.44 ± 0.05	0.52 ± 0.025
Min. Val.	0.2	0.15
Max. val.	0.9	1.05

Table 3.1: Optimal parameters found by the Metropolis-Hastings algorithm for the Boolean model of disks.

	θ_e $[L^{-1}]$	R_e $[L]$	θ $[L^{-1}]$	R
Target	0.1	1.	0.45	0.5
$\hat{\Theta}$	0.13	0.85	0.45	0.5
Std. Dev.	0.03	0.25	0.05	0.04
Min. Val.	0.05	0.5	0.2	0.15
Max. val.	0.2	2.	0.9	1.05

Table 3.2: Optimal parameters found by the Metropolis-Hastings algorithm for the Cox-Boolean model of disks.

parameter has experienced only minor variations throughout the iterations of the algorithm, it implies that when this parameter deviates from its optimal value, there is a substantial deviation from the targeted measurements. In other words, even small changes in the parameter can have a significant impact on the agreement between the predicted measurements and the desired outcome. Overall, we can note that there is a good agreement between the parameters identified by the algorithm and the ones corresponding to the experimental microstructure. A comparison between the measurements as conducted on the target microstructures simulated with both morphological models and as obtained with the optimal set of parameters $\hat{\Theta}$ is displayed in Fig. 3.2 and 3.3. Again, we note a fairly good agreement between the measurements.

Fig. 3.4 displays a scatterplot of the parameters sampled by the Metropolis-Hastings algorithm for both morphological models. Interestingly, we note that the parameters θ and R sampled by the algorithm are strongly correlated in the case of the Boolean model of disks. We performed a linear regression between the parameters θ and $1/R^2$, to find a coefficient of determination equal to 0.81. For a Boolean model of disks, the volume fraction v is known to be given by

$$v = 1 - \exp(\pi R^2 \theta). \tag{3.20}$$

This leads to the linear relationship

$$\frac{1}{R^2} = -\frac{\pi}{\log(1-v)}\theta. \tag{3.21}$$

The covariance measurements are highly sensitive to the volume fraction. Hence, by construction, the Metropolis-Hastings will usually select a set of parameters that allows to obtain a volume fraction similar to the one measured on the experimental samples. As evidenced in Fig. 3.4 , the parameters are indeed distributed around a straight line. For the Cox-Boolean model, we note that the parameters θ and R sampled by the algorithm remain correlated, but the correlation is less obvious than for the Boolean

model. A linear regression performed between the parameters θ and $1/R^2$ yields a coefficient of determination equal to 0.51. Due to the presence of exclusion zones in the microstructure, Eq. (3.21) is not valid anymore, which explains this reduced correlation.

3.3 Experiments

In this section, we demonstrate the application of our proposed optimization approach to investigate an experimental microstructure. The microstructure under examination consists of pigments embedded in a matrix, forming a colloid film. These pigments possess an elongated ellipsoidal shape similar to that of a needle. To capture the microstructure for analysis, a dataset of experimental images was acquired using transmission electron microscopy (TEM) imaging. Figure 3.6 provides an illustrative example of one such TEM image. The analysis of the experimental TEM images reveals a noticeable tendency of the pigments to aggregate within the microstructure, resulting in the formation of clusters. These clusters exhibit a characteristic size that significantly exceeds the individual pigment size.

To describe the microstructure, we use a two-scale model as follows:

- The first scale of the simulation describes the pigments aggregates through a Boolean model of inclusion spheres characterized by two parameters, namely the intensity θ_i of the Poisson point $\mathcal{P}$ process used to construct the Boolean model and the radii R_i of the spheres.
- The second scale of the model simulates the pigments within the aggregates using a Boolean model of ellipsoids whose largest semi-axis follows a normal law with mean L and standard deviation σ_L, and whose smallest semi-axes follow a normal law with mean l and standard deviation σ_l. A hardcore distance h is introduced in the underlying point process. Ellipsoids are implanted at locations sampled from a Poisson point process with intensity θ restricted to the set constituted by the spheres constitutive of the first scale of the model. Ellipsoids belonging to the same inclusion sphere share the same orientation, sampled uniformly from Euler's angles.

The description of the microstructure model requires eight parameters. These parameters include the intensity θ_i of the Poisson point process for the aggregates, their radius R_i, the lengths of the semi-axes L and l of the pigments, along with their corresponding standard variations σ_L and σ_l, the intensity θ of the point process used for the pigments and the hardcore distance h.

To determine the model parameters, we employ the Bayesian approach explained in the preceding section. We initiate the Metropolis-Hastings algorithm by setting λ to 1 during the simulation, and we terminate the algorithm after 200 iterations. Similar to before, we estimate the optimal parameter set for the microstructure by selecting the sample that exhibits the highest posterior probability among all the sampled parameters. Throughout the MCMC run, we generate TEM images of the microstructure by first creating a 3D sample of the model, then extracting a thick slice from the 3D volume that matches the thickness of the experimental slices used for obtaining the TEM images, and finally projecting the thick slice onto a single plane. An illustrative example of the microstructure generated using this approach is presented in Figure 3.6.

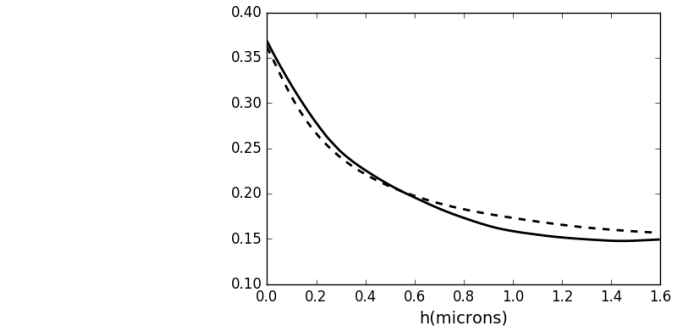

(a) Covariance

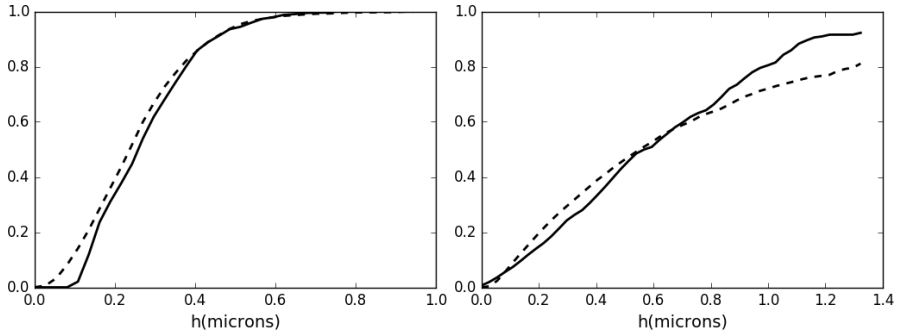

(b) Granulometry measurements

Figure 3.5: Comparison between the covariance and the granulometries as obtained on the experimental TEM images (dashed lines) and on the simulated TEM images for the "needles" microstructure [Figliuzzi et al., 2021].

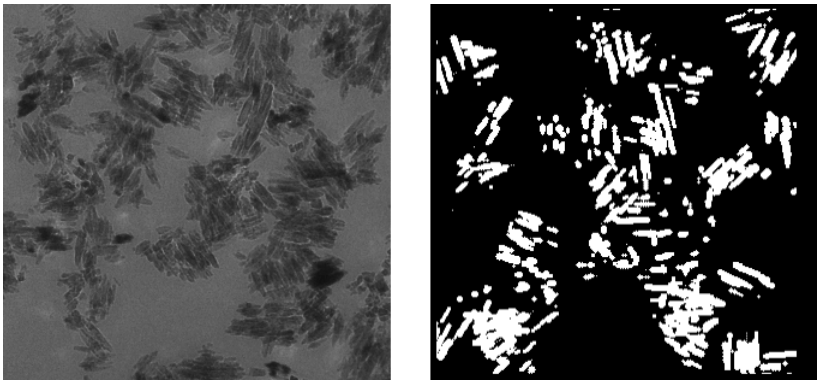

Figure 3.6: Needles microstructure: experimental TEM image (left) and binary image simulated with the morphological model using the optimal set of parameters identified by the Bayesian approach [Figliuzzi et al., 2021].

θ_i (µm^{-3})	R_i (µm)	θ (µm^{-3})	L (µm)	l (µm)	h (µm)
0.56	0.54	45.	0.56 ± 0.08	0.07 ± 0.01	0.2

Table 3.3: Optimal parameters for the "needles" microstructure.

Covariance and granulometry measurements are subsequently performed on both the experimental and simulated TEM images.

The optimal set of parameters $\hat{\Theta}$ determined by the Metropolis-Hastings algorithm for the morphological model is shown in Table 3.3. A comparison is presented in Figure 3.5 between the measurements conducted on the experimental microstructure and those obtained using the optimal set of parameters. We observe that the optimal parameter set enables us to achieve a good agreement between the experimental measurements and the measurements performed on simulations of the morphological models.

3.4 Conclusion

In this chapter, we presented a Bayesian approach for determining the parameters of morphological models of microstructure using measurements obtained from experimental microstructure images. Through various examples, we demonstrated that this approach effectively identifies the optimal parameters for different morphological models and detects potential correlations among the model parameters. Our proposed methodology offers several advantages compared to other parameter determination approaches. Specifically, it eliminates the need for computing the gradient of the model with respect to the parameters and provides a sequence of parameter samples, enabling the quantification of the parameters sensitivity and the identification of potential correlations between them.

3.5 Related publications

- B. Figliuzzi, A. Montaux-Lambert, F. Willot, G. Naudin, P. Dupuis, B. Querleux, E. Huguet. A Bayesian Approach to Morphological Models Characterization. *Image Analysis and Stereology*, 40(3), 171-180, 2021.

Part II
Image segmentation

In materials science, experimental campaigns can generate large databases of images that require automatic processing for being efficiently exploited. Supervised learning algorithms based on convolutional neural networks constitute the state-of-the-art approaches for performing image segmentation. However, training these architectures depends upon the availability of images segmented manually, which are often not available in sufficient quantity. This challenge is further compounded by the fact that images from materials science experiments are markedly different from the natural images that are typically used for training convolutional neural networks. Consequently, transfer learning techniques alone are insufficient to address the lack of training data. As a result, the development of segmentation methods that can effectively handle limited amounts of data remains a crucial area of research. The applicability of the most effective segmentation methods to actual industrial problems is dependent on overcoming this challenge. This can be done in two manners: first, by pursuing the development of segmentation techniques that require minimal training data; second, by relying on data augmentation and image generation techniques to obtain sufficiently large training datasets for the task at hand. This part of the manuscript presents my main contributions to these two lines of research.

My main contribution to the development of segmentation techniques is related to the work conducted during the PhD thesis of Kaïwen Chang between 2016 and 2019, which I co-directed along with Jesús Angulo at the Centre for Mathematical Morphology [Chang, 2019]. During the course of this PhD thesis, we developed an algorithm based on the Eikonal equation to generate the superpixel segmentation of a given image based on color and texture information [Chang and Figliuzzi, 2019, 2020]. Any superpixel segmentation can be represented by a graph whose nodes represent the superpixels of the image and whose edges represent the adjacency relations between superpixels. Using different features extracted from each superpixel, we showed that it was possible to learn a distance characterizing the dissimilarity between neighboring superpixels with a limited number of images. We finally relied on a generalization of the Eikonal equation to graph structures to merge the superpixels of the image into an actual segmentation. I describe these researches in chapters 4 and 5. In chapter 6, I present a study conducted in collaboration with the Institut de Physique de Nice, which explores the second line of research mentioned previously, and aims to segment images in a supervised way from a dataset of images synthesized with a morphological model.

Chapter 4

Eikonal-based superpixels

Superpixel algorithms are a class of techniques that seek to divide an image into smaller regions consisting of similar pixels, providing more meaningful and analyzable representations than raw pixels. One major advantage of superpixels is that they significantly reduce the amount of calculations required for further processing, as their number is much lower than that of the original pixels. Moreover, due to their homogeneous nature, superpixels form subregions of the image that are particularly relevant for feature computation. As a result, superpixel algorithms are frequently used as a pre-processing step in various applications such as object classification, image segmentation, or depth estimation, as demonstrated by studies like [Zitnick and Kang, 2007] and [Fulkerson et al., 2009].

A significant amount of literature in computer vision and image analysis is dedicated to the topic of superpixels [Achanta et al., 2012; Li and Chen, 2015; Stutz et al., 2018]. Graph based methods rely on a graph representation of the image. These methods include the normalized cut algorithm (NC) of Shi and Malik [Malik et al., 2001; Shi and Malik, 2000] or the efficient graph-based segmentation algorithm (GS) of Felzenswalb and Huttenlocher [Felzenszwalb and Huttenlocher, 2004]. By contrast, clustering based methods proceed by iteratively refining clusters of pixels until some convergence criterion is met. These methods notably include mean shift [Comaniciu and Meer, 2002], watershed [Beucher and Lantuéjoul, 1979; Beucher and Meyer, 1992; Figliuzzi et al., 2017; Meyer and Beucher, 1990; Vincent and Soille, 1991], turbopixel [Levinshtein et al., 2009] or waterpixel [Cettour-Janet et al., 2019; Machairas et al., 2015] algorithms, respectively.

This chapter presents in a synthetic manner the work conducted during the PhD thesis of Kaïwen Chang on the development of a novel superpixel algorithm [Chang, 2019]. Our objective was to use superpixels as a first segmentation step to perform the complete segmentation of an image by region merging techniques. We noticed during this work that the quality of the resulting segmentation was strongly related to the quality of the initial superpixel segmentation. Hence, we dedicated part of our efforts to the development of a clustering-based algorithm for generating a superpixel partition, which will be referred to as fast-marching based superpixels (FMS) algorithm in the remainder of this manuscript. In this algorithm, following an idea originally introduced for the Eikonal-based region growing for efficient clustering algorithm (ERGC) of Buyssens *et*

al. [Buyssens et al., 2014a,b,c], we used the Eikonal equation to compute the superpixel partition of the image. This work resulted in two scientific publications [Chang and Figliuzzi, 2019, 2020], which we summarize in the present chapter.

4.1 Algorithm

The FMS algorithm operates by likening the growth of regions on an image to waves propagating through a non-uniform medium, where the growth rate depends on the local color and texture. The wave propagation is governed by the stationary Eikonal equation. Hence, the task of constructing the superpixel partition essentially involves solving the Eikonal equation on the image domain using a velocity field that is determined by the local color and texture. A similar approach was introduced by Buyssens *et al.* in 2014 [Buyssens et al., 2014a,b,c], but there are notable differences between the two methods, such as the expression for the local velocity as a function of the image content, the use of texture features, and the region update mechanism during propagation.

4.1.1 The Eikonal equation on a bounded domain

We recall in this paragraph basic notions on the stationary Eikonal equation, already discussed in section 2.3 of this manuscript. Let Ω denote some bounded domain in $\mathbb{R}^2$. In what follows, we consider a wave front emerging from a set of points on the domain and propagating at a velocity $u := u(x)$ specified at every location $x \in \Omega$. Let us denote by $x \to T(x)$ the function associating with each point $x \in \Omega$ the first arrival time of the propagation front. It can be shown [Sethian, 1996] that T is solution of the so-called stationary Eikonal equation

$$\|\nabla T(x)\| = \frac{1}{u(x)}, \quad \forall x \in \Omega. \tag{4.1}$$

The stationary Eikonal equation (4.1) must be complemented by boundary conditions specified on $\partial\Omega$. One usually considers a function g defined on the boundary $\partial\Omega$ so that $T(x) = g(x)$ for all $x \in \partial\Omega$. Generally, the function g is taken to be identically 0. Hence, the stationary Eikonal equation becomes

$$\begin{cases} ||\nabla T(x)|| = \dfrac{1}{u(x)}, & \forall x \in \Omega \\ T(x) = 0, & \forall x \in \partial\Omega. \end{cases} \tag{4.2}$$

For all x in Ω, the solution $T(x)$ of Eq. (4.2) can be interpreted as the minimal time required to travel from x to the domain boundary $\partial\Omega$. In other words, the Eikonal equation allows us to compute the shortest distance between any point x of the domain Ω and the boundary $\partial\Omega$. Efficient numerical methods can be found in the literature to solve the Eikonal equation on a domain. Among these methods, the fast marching algorithm, originally introduced by Sethian *et al.*, ranks among the most popular. It works by iteratively following the wavefront propagation and computing the first arrival time step by step [Sethian, 1996].

4.1.2 Region growing

We describe in this section the practical implementation of the superpixel generation algorithm. Let us first introduce some notations. A pixel in image $\mathcal{I}$ is denoted by p and its coordinates in the image by (x, y). We denote by $\mathbf{C}(p)$ the color at pixel p in the CIELAB color space. Similarly, we denote by $\mathbf{T}(p)$ a vector of features characterizing the local texture at pixel p. The proposed algorithm is equivalent to solving the stationary Eikonal equation on the image domain with the fast marching algorithm for a velocity field depending upon the local content of the image.

Let K denote the requested number of superpixels. We initialize the algorithm by selecting K seeds on a regular grid. To avoid placing a seed on a boundary, following [Achanta et al., 2012], we place the seed at the local minimum of the gradient in a 3×3 neighborhood of the nodes of the grid.

For each seed s_i, we define a velocity field $u_i(p)$, which depends on the color and on the texture of both the seed and the location $p := (x, y)$ in the image. The labels of the seeds are then gradually propagated from the labeled pixels to the unlabeled pixels according to the local velocities $(u_i(p))_{i=1,..,K}$. The propagation is described by the Eikonal equation

$$\begin{cases} ||\nabla T(p)|| = \dfrac{1}{u(p)}, & \forall p \in \mathcal{I} \\ T(p) = 0, & \forall p \in \partial\mathcal{I}. \end{cases} \tag{4.3}$$

In this expression, $u(p)$ is the local velocity at pixel p, $\partial\mathcal{I}$ corresponds to the subset of the image $\mathcal{I}$ constituted by the seeds $\{s_1, s_2, ..., s_K\}$ and $T(p)$ is the minimal traveling time from $\partial\mathcal{I}$ to p. Once seeds have been selected, we solve Eq. (4.3) using the fast marching algorithm, whose cost is known to be in $O(N \log N)$ [Sethian, 1996, 1999], N being the number of pixels in the image.

4.1.3 Local velocity model

We can use the algorithm described in the previous paragraph with any non-negative velocity model $u(p)$. One of the main difficulties related to superpixel segmentation is that certain regions of the image are extremely textured. In these regions, the colorimetric distance between two pixels presumes in no way that these two pixels belong to the same cluster or not. To account for texture, we considered a velocity model incorporating both local color and texture information. This velocity model is based on a distance D between the image pixels and the seeds defined by

$$D(p, s_i) = w_0 \|\mathbf{C}_p - \mathbf{C}_i\|_2 + w_1 d(\mathbf{T}_p, \mathbf{T}_i), \tag{4.4}$$

where $\mathbf{T}_p$ is a vector of features characterizing the local texture at pixel p, $\mathbf{T}_i$ is the corresponding vector of features for the $i-$th cluster center s_i, $d(\mathbf{T}_p, \mathbf{T}_i)$ a texture distance between the pixel p and the $i-$th cluster center, and w_0 and w_1 are positive coefficients weighting color and texture contributions, respectively.

Several approaches can be implemented to obtain the features vector $\boldsymbol{T}_p$. In our study, we constructed a texton map based on some texture classifier and we associated

with each pixel in the image an histogram of textons computed in a local window of fixed size. We could then define a texture distance between two pixels of the image by considering the χ^2 distance between the textons histograms associated with both pixels.

The local velocity $u_i(p)$ associated with the propagation of region i is computed by relying on the exponential kernel

$$u_i(p) = \exp(-w_0\|\boldsymbol{C}_p - \boldsymbol{C}_i\|_2 - w_1 d(\boldsymbol{T}_p, \boldsymbol{T}_i)). \tag{4.5}$$

In the FMS algorithm, the velocity is at its maximum and equals 1 when both the pixel and the cluster center have identical color and texture characteristics. The weighting parameters w_0 and w_1 that are used to compute the color and texture distances need to be chosen carefully. Increasing these parameters leads to better adherence to the image boundaries because the propagation velocity becomes more sensitive to variations in color and texture. However, using high values of w_0 and w_1 (i.e. > 8) can cause the local velocity to become very low. This, in turn, can result in the creation of many small, isolated superpixels that correspond to low-contrast regions of the image, which can degrade the compactness of the final superpixels partition. Therefore, the choice of parameters involves a trade-off between boundary adherence performance and topological considerations.

4.1.4 Refinement

Images often contain regions of interest that vary greatly in scale, making it challenging to select seeds for superpixel segmentation using a uniform grid. To achieve better boundary adherence in the superpixel segmentation, we employed a refinement strategy that utilizes the map of arrival times.

The map $T(p)$ obtained at the end of the superpixel generation algorithm contains information about the arrival time at each pixel p, and can be used to refine the superpixel partition by adding additional seeds. When the arrival time at a pixel p within a region $\mathcal{R}_i$ associated with seed s_i is high, it indicates that the region boundary propagating from s_i has passed through pixels that are highly dissimilar to s_i. Therefore, the arrival times in each region $\mathcal{R}_i$ can be used as a criterion to determine if that region should be split further. A similar refinement strategy is employed in [Buyssens et al., 2014b], where the distance map is used to generate new superpixels for refinement at the end of the over-segmentation process.

The refinement is conducted as follows. Let us denote by $\mathcal{B}$ the set of pixels belonging to the superpixel boundaries and by $\delta(\mathcal{B})$ the dilated set of $\mathcal{B}$ by a disk of radius 2 pixels. Then, the maximal arrival time in region $\mathcal{R}_i$ is defined to be

$$t_i = \max_{p \in \mathcal{R}_i \cap \delta(\mathcal{B})^c} T(p), \tag{4.6}$$

where $\delta(\mathcal{B})^c$ is the complementary set of $\delta(\mathcal{B})$. To further refine the superpixel segmentation, we select the k regions with highest t_i and we add seeds at the corresponding locations before re-propagating, k being a parameter fixed by the user. We exclude the pixels that belong to the region $\delta(\mathcal{B})$ to avoid implanting a seed directly on a boundary. This procedure is repeated iteratively until the desired number of superpixels is obtained and significantly improves the boundary adherence of the resulting superpixel partition.

It must however be noted that the refinement strategy usually reduces the compactness of the obtained superpixel segmentation by increasing the density of superpixels in selected regions of the image.

4.1.5 Comparison with other algorithms

It is interesting to point out the main differences between the FMS and the ERGC algorithms [Buyssens et al., 2014a,b,c]. ERGC is a clustering-based algorithm that computes superpixel partitions by relying upon the Eikonal equation. The most significant difference between both algorithms is related to the velocity models that are employed in both cases. For ERGC, the local velocity field is simply given by:

$$u_i(p) = \frac{1}{||\boldsymbol{C}_p - \boldsymbol{C}_i||_2}. \tag{4.7}$$

In particular, the velocity is very high when the pixel and the seed have similar color characteristics. Another important difference is that ERGC only considers local color information to compute the local velocity, while FMS is also able to incorporate texture information. FMS is also very close to an algorithm referred to as Iterative Spanning Forest (ISF) [Vargas-Muñoz et al., 2019], which constructs superpixels by iteratively merging pixels with the aim of maximizing intra-clusters color and spatial distance proximity. The maximization is conducted through a greedy procedure. Even if ISF cannot be directly interpreted in terms of waves propagating according to the stationary Eikonal equation, the greedy procedure for merging superpixels is similar to the fast marching algorithm used to compute the superpixel partition in FMS and ERGC.

Overall, the main differences between the proposed approach and both ISF and ERGC are the incorporation of texture information and the choice of the velocity model to compute the clustering, as well as the refinement procedure and the choice of initial seeds, that are not implemented for ISF and ERGC, respectively.

4.2 Experiments and discussion

We evaluate in this section the performance of the FMS algorithm and compare it with several state of the art superpixel algorithms including SLIC [Achanta et al., 2012], ERGC [Buyssens et al., 2014a,b,c] and ISF [Vargas-Muñoz et al., 2019]. These algorithms share the characteristics of generating superpixels through clustering procedures, either by relying on a K-means algorithm, like SLIC, or on iterative agglomeration procedures, like ERGC, ISF and FMS. To perform the evaluation, we use the Berkeley Segmentation Dataset 500 (BSDS500) [Arbelaez et al., 2011; Martin et al., 2001]. This dataset contains 500 images and provides several ground truth manual segmentations for each image. We point out that we also presented in [Chang, 2019; Chang and Figliuzzi, 2020] an evaluation of the performance of FMS on a database of images recombining texture patches extracted from images representing stripes, constructed by Giraud *et al.* with the intent to evaluate a superpixel algorithm adapted to highly textured images [Giraud et al., 2019].

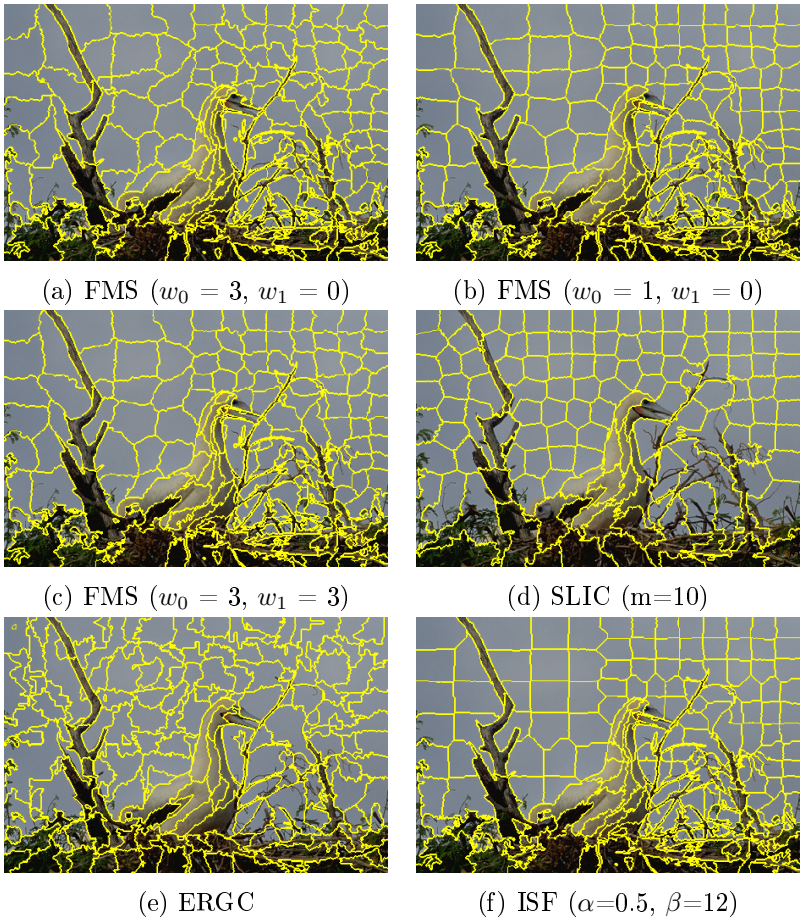

(a) FMS ($w_0 = 3$, $w_1 = 0$) (b) FMS ($w_0 = 1$, $w_1 = 0$)

(c) FMS ($w_0 = 3$, $w_1 = 3$) (d) SLIC (m=10)

(e) ERGC (f) ISF (α=0.5, β=12)

Figure 4.1: Illustration of the superpixel segmentation on an image of the BSDS500, for $K = 200$ superpixels [Chang and Figliuzzi, 2020].

4.2.1 Experiments on the BSD500 dataset

To evaluate the performance of the algorithms, we considered the following metrics: boundary recall, undersegmentation error, compactness and contour density. These metrics are classic in the field of image segmentation and we refer the reader seeking more details to the articles [Chang and Figliuzzi, 2019, 2020; Schick et al., 2012].

Figure 4.2 displays the recall, undersegmentation error, compactness and density metrics of FMS, SLIC, ISF and ERGC averaged over the 500 images of the BSDS500. These metrics are represented for a number K of superpixels ranging from 100 to 600. We also display in Fig. 4.1 superpixel partitions obtained with these algorithms on an image of the BSDS500. In our results, we present several versions of the FMS algorithm: two versions using only color information, and a third one incorporating texture information through an additional texton channel computed with a bank of Gabor filters [Jain and Farrokhnia, 1991]. We selected the parameter $w_0 = 3$ for the color term in the velocity expression (4.5) by relying on a grid search to maximize boundary adherence while keeping a reasonable compactness for the superpixel partition. To illustrate the trade-off between compactness and boundary recall controlled by the parameter

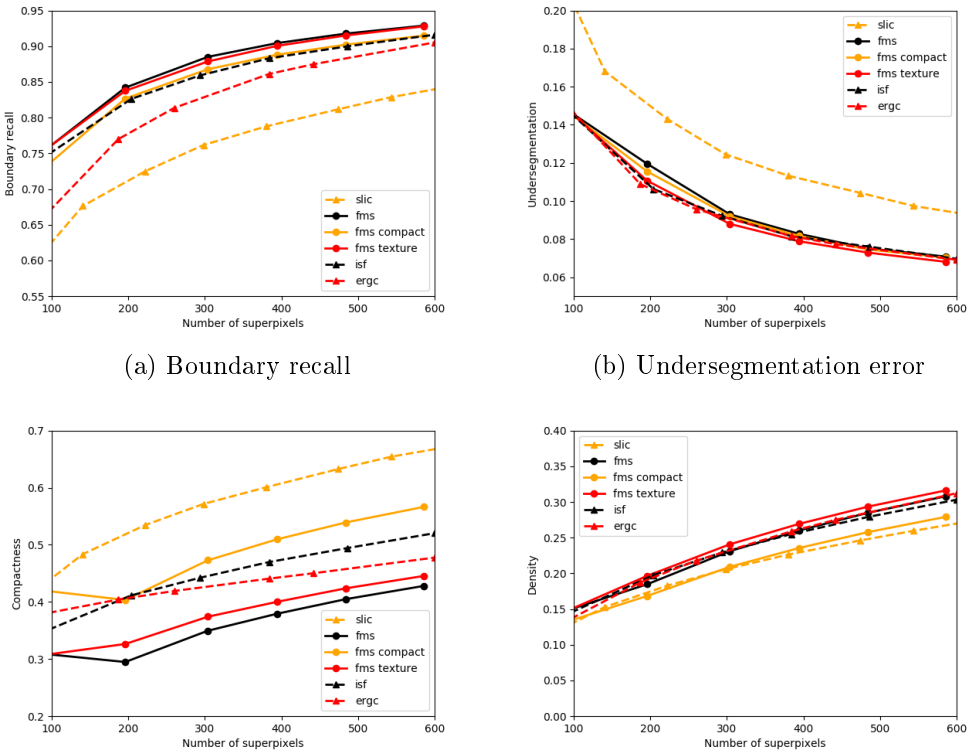

(a) Boundary recall

(b) Undersegmentation error

(c) Compactness

(d) Density

Figure 4.2: Comparison between SLIC, ERGC, IFS and FMS algorithms on BSDS500 [Chang and Figliuzzi, 2020].

w_0, we also displayed the metrics as obtained with the set of parameters $w_0 = 1$ and $w_1 = 0$. Finally, we computed the superpixel partition using the parameters $w_0 = 3$ and $w_1 = 3$ to evaluate the texture influence. Additional experiments aiming at characterizing the influence of the refinement and of the seeds initialization are also presented in the original article [Chang and Figliuzzi, 2019, 2020] related to the present work.

In terms of recall, we can note that the FMS algorithm yields better results than SLIC and ERGC algorithms, and performs slightly better than ISF when $w_0 = 3$. However, the boundary recall should be considered along with compactness, and we can note that the compactness of the FMS algorithm is significantly lower than the one of ISF when $w_0 = 3$. Interestingly, when specifying $w_0 = 1$ and therefore favoring the construction of compact superpixels, we are able to obtain a similar boundary recall than the one reached by ISF, but with higher compactness. The gain in terms of superpixel compactness is also clearly visible in Fig. 4.1. The compact FMS algorithm also exhibits a smaller density when compared to all other algorithms except SLIC. The performances of SLIC appear to be significantly below the ones of the other algorithms. Regarding the compactness, we can note that the compactness decreases between $K = 100$ and $K = 200$ for the FMS algorithm. This is caused by the fact that no refinement steps are applied when computing the partition with $K = 100$ superpixels. Due to the refinement procedure, the size of the superpixels constructed with the FMS algorithm is locally

adapted depending on the content of the image. Therefore, the refinement is performed at the expense of the superpixels uniformity: the refinement yields superpixels that are heterogeneous in scale and size. We can finally note that in terms of undersegmentation, all considered superpixel algorithms are relatively similar.

4.3 Conclusion

During Kaïwen Chang's PhD thesis research [Chang, 2019], we developed a fast-marching based algorithm for generating superpixel partitions of images, building upon an idea initially used in the ERGC algorithm by Buyssens *et al.* [Buyssens et al., 2014a,b,c]. The significant contribution of this study was the introduction of a new expression for the velocity term, which allowed for texture information to be incorporated into the computation of the superpixel partition. In addition, different strategies were also proposed to refine the segmentation.

We evaluted the FMS algorithm on the Berkeley Segmentation Database 500 and we found that it constructs superpixel partitions with slightly higher recall and similar undersegmentation error compared to similar superpixel algorithms such as ERGC, ISF, and SLIC. Moreover, we demonstrated that the inclusion of texture information improved the compactness of the partition without compromising boundary recall.

One potential extension of this work could involve incorporating gradient information into the local velocity model to better account for contours and image discontinuities. It would also be interesting to attempt to automatically estimate optimal parameters, such as w_0 and w_1, based on the processed image, a strategy that has been successfully employed in several other superpixel algorithms [Achanta and Susstrunk, 2017; Giraud et al., 2019]. Additionally, developing a multiscale version of the algorithm, where superpixels are computed and refined at different scales using a pyramid representation of the image, could be useful. This multiscale approach could generate a hierarchy of superpixel partitions and a corresponding saliency map, which could be utilized as a criterion for superpixel merging.

4.4 Related works

In addition to this work on the generation of superpixels, we also worked during Kaiwen Chang's PhD thesis on a segmentation algorithm based on the watershed algorithm [Beucher and Meyer, 1992; Vincent and Soille, 1991]. In general, applying a watershed algorithm directly on the image gradient results in an over-segmented image. To overcome this problem, markers are classically employed to indicate the approximate locations of the objects to be segmented [Meyer and Beucher, 1990]. The primary challenge in marker-controlled segmentation is therefore to determine appropriate markers locations. Our algorithm relies on a multi-scale representation [Burt and Adelson, 1987] of the image to select relevant markers at different scales, which enables the obtaining of a multi-scale segmentation of the image. Our approach uses the discrete decimated wavelet transform to obtain successive approximations of the image [Mallat, 1999]. For each approximation, the minima of the gradient image are propagated back to the original image space and selected as markers for the watershed transform. Interestingly, this technique establishes a hierarchical structure for the detected contours, since the most important contours in the image are usually persistent through the subsequent

coarse approximations in the decomposition. This work was presented at the 13th International Symposium on Mathematical Morphology (ISMM 13th) [Figliuzzi et al., 2017].

4.5 Related publications

- K. Chang and B. Figliuzzi. Fast marching based superpixels. *Mathematical Morphology-Theory and Applications*, 4(1), 127-142, 2020.
- K. Chang and B. Figliuzzi. Fast marching based superpixels generation. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing* (pp. 350-361). Springer, July 2019.
- B. Figliuzzi, K. Chang, M. Faessel. Hierarchical segmentation based upon multiresolution approximations and the watershed transform. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing* (pp. 185-195). Springer, 2017.

Chapter 5

Eikonal-based region merging

In chapter 4, we presented a superpixel algorithm using the Eikonal equation to compute the superpixel partition of an image. We go one step further in this chapter and present a generalization of this algorithm to the framework of graphs, which allows to merge image regions into an actual segmentation of the image. The chapter is organized as follows: in section 5.1, we introduce a generalization of the Eikonal equation to the setting of graphs. In section 5.2, we explain how the Eikonal equation can be leveraged to perform region merging in the image. We present finally in section 5.3 the results obtained with our algorithm on the Berkeley Segmentation Dataset and compare these results to the ones obtained with a classical algorithm for performing graph clustering: the normalized cut [Shi and Malik, 2000]. The research works presented in this chapter was conducted during the PhD thesis of Kaïwen Chang between 2016 an 2019 [Chang, 2019].

5.1 Eikonal equation on an undirected graph

We describe in this section how to extend the continuous Eikonal equation to the setting of graphs. A graph is a structure employed to depict a collection of objects, wherein certain objects are connected. The objects are denoted by *vertices* or *nodes*, while the connections between pairs of vertices are represented by *edges*.

5.1.1 Eikonal equation

In what follows, we consider an undirected, path-connected graph $\mathcal{G} := (V, E)$. We assume that each edge (i, j) in E connecting vertices V_i and V_j carries a weight w_{ij} that describes the similarity between nodes V_i and V_j, and that $t : V \to \mathbb{R}$ is some function defined on the set V of all vertices of the graph. To keep notations simple, we note t_u the value of the function t at vertice u.

Definition 2. *Let v be some vertex of the graph. We denote by $\mathcal{N}_v$ the set of all neighbor vertices of v i.e. the set of all vertices that are connected to v. For all vertices u in $\mathcal{N}_v$, the morphological derivative of t at v with respect to u is:*

$$Dt(u, v) := w_{uv}(t_u - t_v)^+, \tag{5.1}$$

where the quantity $(t_u - t_v)^+$ is defined by $(t_u - t_v)^+ = \max(0, t_u - t_v)$.

Definition 2 allows us to define the gradient of t in the following manner:

Definition 3. *The gradient of t at vertex v is the vector*

$$\nabla t(v) := (Dt(u,v))_{u\in\mathcal{N}_v}. \tag{5.2}$$

Based on definitions 2 and 3, we can propose a formulation of the Eikonal equation adapted to graph structures by using an analogy with the continuous setting. In the continuous setting, the Eikonal equation relates the L^p norm of the local gradient to the local velocity u in the open domain Ω:

$$u(x)\|\nabla t(x)\|_p = 1, \quad \forall x \in \Omega. \tag{5.3}$$

For a graph structure, the equivalent formulation is therefore

$$u(v)\|\nabla t(v)\|_p = 1, \quad \forall v \in V, \tag{5.4}$$

where $u(v)$ denotes a local velocity associated to node v.

Here, we focus on the case where $p = \infty$. In this case, the Eikonal equation is given at each vertex v of $\mathcal{G}$ by

$$u(v) \max_{u\in\mathcal{N}_v} w_{uv}(t_u - t_v)_+ = 1. \tag{5.5}$$

The choices of $p = \infty$ and of the morphological derivative in Eq. (2) are not accidental and expression (5.5) arises naturally if we consider a wave propagating on the entire graph $\mathcal{G}$ and arriving at a given vertice $v \in \mathcal{G}$. For all neighbor vertices $u \in \mathcal{N}_v$, we have necessarily

$$t_v \leq t_u + \frac{1}{w_{uv}}, \tag{5.6}$$

and there exists a particular neighbor $\hat{u}$ of v such that

$$t_u = t_{\hat{u}} + \frac{1}{w_{\hat{u}v}}. \tag{5.7}$$

If we put these requirements together, we obtain the single equation

$$\|\nabla t(v)\| := \max_{u\in\mathcal{N}_v} w_{uv}(t_v - t_u)^+. \tag{5.8}$$

It is finally straightforward to define boundary conditions for the Eikonal equation, by simply selecting a subset $\partial\mathcal{G} = (v_i)_{1\leq i\leq k}$ of k vertices in V and specifying that $\forall i = 1, ..., k$, $t_{v_i} = 0$.

5.1.2 Eikonal equation and shortest path distance between two vertices

Definition 4 (Shortest path distance between two vertices)**.** *Let us denote by $\mathcal{P}(u,v)$ the set of all paths in $\mathcal{G}$ that connect vertices u and v. Since $\mathcal{G}$ is path-connected, $\mathcal{P}(u,v)$ contains at least one element. A path p in $\mathcal{G}(V,E)$ is a collection of edges in E. The distance $D(p)$ associated to path p is simply the sum of the weights of the edges constitutive of p. The shortest path distance between two vertices u and v in $\mathcal{G}$ is:*

$$\hat{D}(u,v) = \min_{p\in\mathcal{P}(u,v)} D(p) \tag{5.9}$$

Building upon this definition, we can define the distance function between any vertex v of the graph and the subset $\partial\mathcal{G} = (v_i)_{1\leq i\leq k}$ to be

$$d(v, \partial\mathcal{G}) = \inf_{i=1,...,k} \hat{D}(v, v_i). \tag{5.10}$$

Proposition 2 relates the distance function to the solution of the Eikonal equation:

Proposition 2. *Let $\mathcal{G}$ be an undirected, weighted graph. Then, the gradient of the distance function $d(\cdot, \partial\mathcal{G})$ satisfies the Eikonal equation*

$$||\nabla t(v)|| = 1$$

with boundary conditions $t(v_i) = 0, \forall i = 1, ..., k$.

5.1.3 Fast marching algorithm on graphs

We present in this section the generalization of the fast marching algorithm already encountered in chapters 2 and 4 to undirected graph structures. To that end, let us consider an undirected, path-connected graph $\mathcal{G} := (V, E)$ such that each edge (i, j) in E carries a weight w_{ij}. The fast marching algorithm seeks to determine the solution of the Eikonal equation. However, instead of iteratively solving this equation for each vertex of $\mathcal{G}$ until convergence, it works by following the front propagation within the graph to compute the arrival times. During the procedure, the vertices of $\mathcal{G}$ are divided into three distinct subsets:

- The *frozen* set groups all vertices already reached by the propagation front.
- Vertices adjacent to the frozen points but not reached by the front yet are grouped in the *narrow band*.
- The remaining vertices constitute a subset referred to as the *far away* set.

Initialization

1. We affect the arrival time $t = 0$ to all vertices in $\partial\mathcal{G}$ and we add them to the narrow band.
2. We label all other vertices as *far away* and we affect them the arrival time ∞. The frozen set is initially empty.
3. In order to keep track of the shortest paths between each vertex in $\mathcal{G}$ and the boundary $\partial\mathcal{G}$, we affect the label i to each vertex in $\{v_i, i = 1, ..., k\}$. All other vertices are labelled 0.

Iteration At each iteration, we extract the vertex v of the narrow band with the smallest arrival time and we label it as *frozen*. Next, we compute the arrival times for each neighbor u of v not belonging to the frozen set, by solving equation (5.5) and by considering that the arrival times at the neighbor nodes w of u are t_w if $w \in \mathcal{C}_i$ and ∞ otherwise, where $\mathcal{C}_i$ is the subset of $\mathcal{G}$ containing the points reached by the front that emerged from the i-th vertex in $\partial\mathcal{G}$. We assume obviously that $v \in \mathcal{C}_i$. Once the arrival time t of a neighbor point u has been computed, two situations can be encountered:

- When u is in the narrow band, it has already been affected an arrival time and it is affected to one of the subsets $\mathcal{C}_j, j = 1, ..., k$. If the new arrival time is smaller than the current one, the arrival time is updated and the vertex u is affected to the subset $\mathcal{C}_i$.
- When the neighbor vertex u is in the *far away* set, we add it to the narrow band with the computed arrival time and to the subset $\mathcal{C}_i$.

Stopping condition The fast marching algorithm stops when the narrow band is empty.

5.2 Application to superpixels merging

We describe in this section an algorithm that allows to perform the clustering of a similarity graph based on the resolution of the Eikonal equation in order to perform the segmentation of color images. A similar idea was proposed in 2014 in [Buyssens et al., 2014c].

Let I be a color image. We denote by $I(p, q)$ the color of the pixel located at position (p, q) in the image. In what follows, we assume that I is represented as the union of N disjoint superpixels $(\mathcal{S}_i)_{i=1,...,N}$:

$$I = \cup_{1 \leq i \leq N} \mathcal{S}_i, \tag{5.11}$$

with $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ if $i \neq j$.

Region adjacency graph We can associate to the superpixel partition of I a graph referred to as its region adjacency graph $\mathcal{G}$. This region adjacency graph is a representation of the image I as an undirected graph, whose vertices $(V_i)_{i=1,...,N}$ are associated with the superpixels $(\mathcal{S}_i)_{i=1,...,N}$. Two vertices V_i and V_j are linked by an edge of the graph if and only if the corresponding superpixels $\mathcal{S}_i$ and $\mathcal{S}_j$ share a common boundary in the image. In the following, we will adopt the notation $\mathcal{G} := (V, E)$ when referring to the region adjacency graph, where V is the set of all vertices (superpixels) in $\mathcal{G}$ and E the set of all edges.

We can specify a weight for each edge in E by defining a function $w : E \to [0, 1]$ which associates to the edge e_{ij} joining vertices v_i and v_j a quantity $w_{ij} \in [0, 1]$, interpreted as a dissimilarity measure between superpixels $\mathcal{S}_i$ and $\mathcal{S}_j$. Several approaches have been considered in the literature to compute the dissimilarity weights w_{ij}, including the color distance between $\mathcal{S}_i$ and $\mathcal{S}_j$ or the strength of the gradient at the boundary. Here, we will assume that, for each pair $(\mathcal{S}_i, \mathcal{S}_j)$ of adjacent superpixels, we were able to estimate the probability p_{ij} that $\mathcal{S}_i$ and $\mathcal{S}_j$ belong to the same segment of the image. It is clear that the quantity

$$w_{ij} = \exp(-p_{ij}), \tag{5.12}$$

then defines a dissimilarity measure between vertices V_i and V_j of $\mathcal{G}$.

Graph clustering According to proposition 2, it is possible to partition a graph $\mathcal{G}$ into $K \geq 1$ subgraphs by relying on the Eikonal equation. To that end, we start by selecting K vertices $(v_k)_{1 \leq k \leq K}$ of $\mathcal{G}$. Then, solving the Eikonal equation on $\mathcal{G}$ with

boundary conditions set to be $t(v_k) = 0, \forall k = 1, ..., K$ allows to compute the distance $w(v, \partial\mathcal{G})$ of the shortest path linking each vertex $v \in V$ to the closest vertex in the subset $\partial\mathcal{G}) := (v_k)_{1 \leq k \leq K}$. Since the graph $\mathcal{G}$ is path-connected, for $k = 1, ..., K$, the subsets

$$\mathcal{C}_k = \{v \in V, w(v, v_k) = w(v, \partial\mathcal{G})\} \tag{5.13}$$

constitute a partition of $\mathcal{G}$ into K connected subgraphs. We can use this approach to compute a partition of the region adjacency graph associated to a superpixel segmentation and therefore coarsen the segmentation.

Algorithm The merging algorithm that we proposed works by iteratively solving the Eikonal equation on the region adjacency graph of the superpixel segmentation $\mathcal{S}$ of I and adapting the boundary conditions. The algorithm start with an initial segmentation $\mathcal{S}^0$ containing N superpixels. Typical values for N are in the order of 500 to 800 superpixels. Our objective is to significantly reduce the number of segments in the image to a value around $50 - 100$. The superpixel merging is conducted as follows:

1. To initialize the algorithm, K vertices $(v_k)_{1 \leq k \leq K}$ are chosen randomly in the region adjacency graph.
2. The Eikonal equation is solved for the region adjacency graph with boundary conditions $t(v_k) = 0, \forall k = 1, ..., K$. After this step, the graph is clustered into K separated subgraphs $(\mathcal{G}_k)_{k=1,...,K}$. For $k = 1, ..., K$, we denote by V_k and E_k the set of the vertices and of the edges of $\mathcal{G}_i$, respectively.
3. For each subgraph $(\mathcal{G}_k)_{k=1,...,K}$, we search for the edge e_k in E_k with maximal weight w_k. We denote by $n_{0,k}$ and $n_{1,k}$ the vertices in V_k linked by e_k. Then, we select the subgraph $\mathcal{G}_j$ whose maximal internal weight is the highest and we add the nodes $n_{0,j}$ and $n_{1,j}$ to the boundary conditions. This step allows to refine the previously obtained segmentation.
4. We solve the Eikonal equation for the region adjacency graph with the updated boundary conditions $t(v_k) = 0, \forall k = 1, ..., K + 2$, where $v_{k+1} = n_{0,j}$ and $v_{k+2} = n_{1,j}$.
5. We iterate between steps 3 and 4 until some stopping criterion is met.

Two distinct stopping criteria can be used in the algorithm. A first stopping criterion consists in stopping the algorithm iterations when a specified number of segments are obtained, the advantage of this approach being that it enables to control the number of segments obtained in the final segmentation. However, this approach can potentially yield a segmentation with segments still containing highly dissimilar superpixels. A second stopping criterion consists in specifying a probability threshold t and in iterating between steps 3 and 4 until no subgraph contains weights higher than this threshold.

5.3 Results and discussion

We present in this section results obtained on the Berkeley Segmentation Dataset (BSDS500) with the proposed merging approach. Starting from a superpixel partition with roughly 500 superpixels, our objective is to reduce the number of superpixels

to around $60-80$ by merging them. The oversegmentation that results from this process contains a "reasonable" number of segments and can serve as a solid foundation for implementing a classification algorithm depending on higher-level features to accomplish the segmentation.

5.3.1 Experiments

Our proposed merging method is agnostic to the choice of superpixel generation algorithm. In our experiments, we started from an image partition containing $K = 500$ superpixels computed with the fast marching based algorithm described in Chapter 4, and we used this initial partition to construct a region adjacency graph (RAG) on the image. The merging procedure was completed in two steps. The first step allows to reduce the number of superpixels from $N = 500$ to $N = 100 - 200$. After the first step, the weights of the RAG are updated and a second merging procedure is conducted in order to obtain an oversegmentation with $N = 50 - 100$ regions.

Dissimilarity measure

The edges of the RAG must carry a dissimilarity measure constructed through Eq. (5.12) from the probability that the adjacent regions associated with the edge belong to the same segment. To compute the dissimilarity measure, we learned the similarity measure between regions by using a regression algorithm, which takes as input features extracted from the pair of adjacent superpixels and returns a score that can be interpreted as the probability of merging each pair of adjacent regions. These features incorporate in particular various color and texture distances between the adjacent superpixels and information on the gradient strength at the boundary separating the superpixels. We refer the reader interested by additional details on the selected features to the manuscript [Chang, 2019].

To train the classifier, we use Berkeley Segmentation Dataset 500 (BSDS500) [Martin et al., 2001] in the following manner. For each training image in the BSDS500, we dispose of K distinct manual segmentation. Let $\mathcal{T}^k := (T_i^k)_{i=1,\dots K_k}$ be the $k-$th segmentation mask in the ground truth, which contains K_k segments. We can associate each superpixel $\mathcal{S} := (S_i)_{i=1,\dots N}$ to the region $\mathcal{T}^k = (T_i^k)_{i=1,\dots K_k}$ that it intersects the most:

$$A_i^k = \arg \max_{j=1,\dots,K_n} \mathcal{A}(S_i \cap \mathcal{T}_j^k), \tag{5.14}$$

where $\mathcal{A}(S_i \cap \mathcal{T}_j^k)$ is the area of the intersection and A_i^k denotes the index of the region associated with S_i in the k-th manual segmentation. For a given edge (S_i, S_j) of the RAG, we define the target similarity measure $\hat{w}_{ij}$ to be:

$$\hat{w}_{ij} = \exp\left(-\frac{1}{K}\sum_{k=1}^{K} 1_{\{A_i^k = A_j^k\}}\right).$$

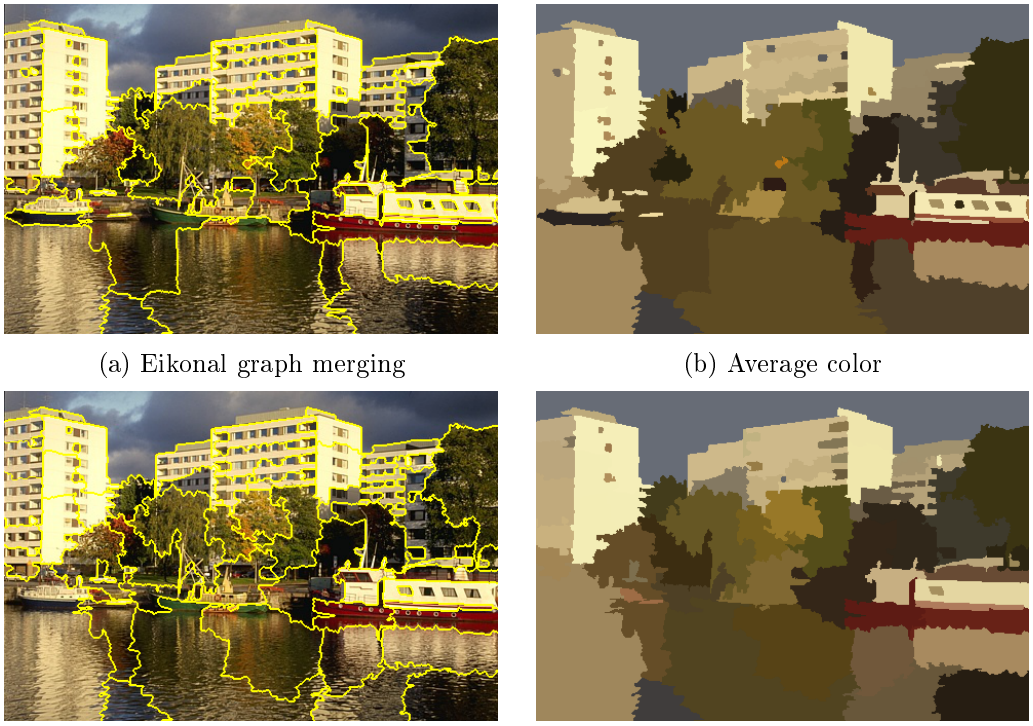

(a) Eikonal graph merging (b) Average color

(c) Ncut (d) Average color

Figure 5.1: Segmentation results obtained with the Eikonal and normalized cut algorithms. The boundary recall are 0.64 and 0.63, respectively and the boundary precision are 0.33 and 0.29. The final partition comprises 70 segments [Chang, 2019].

5.3.2 Results and discussion

To evaluate the performance of our algorithm, we compared it to a classical approach for performing graph clustering, the normalized cut algorithm (Ncut) [Shi and Malik, 2000]. In addition to classical metrics including boundary recall and precision, we also report the results in terms of the segmentation covering metrics. Segmentation covering provides a measures of the average matching between a segmentation and a given ground truth. It is defined by

$$SC(S, S_g) = \sum_{s_i \in S} \frac{|s_i|}{|\mathcal{P}|} \max_{s_j \in S_g} \frac{|s_i \cap s_j|}{|s_i \cup s_j|}. \tag{5.15}$$

In this expression, S and S_g are the segmentation and the ground truth segmentation, respectively. $\mathcal{P}$ is the set of pixels, so that $|\mathcal{P}|$ corresponds to the total number of pixels in the image. The segmentation covering metric computes, for each segment $s_i \in S$, the area of the largest intersection over union with the corresponding ground truth segment s_j. The calculation is weighted by the number of pixels within segment s_i and normalized by the total number of pixels in the image. The maximum possible value 1 can be achieved when S is identical with S_g.

The results obtained with our merging approach are provided in Tab. 5.1. We display one example of obtained segmentation in Fig. 5.1. Overall, we can note that the

results obtained with the Eikonal equation on graphs are better than the ones obtained with the normalized cut. In particular, the Eikonal based approach brings significant improvement over the normalized cut algorithm with respect to the boundary precision and to the segmentation covering metrics. In our experiment, we opted for the second stopping criterion for the Eikonal algorithm, which consists in specifying a probability threshold t and in iterating between the last two steps of the algorithm until no cluster contains weights higher than t. A drawback of this approach is that it makes it difficult to control the number of segments obtained at the end of the procedure. Hence, to facilitate the comparison between the Eikonal and Ncut approaches, for each image, we computed the normalized cut segmentation with the same number of segments as the one yielded by the Eikonal algorithm.

	Eikonal	NCut
Boundary recall	0.76 ± 0.09	0.74 ± 0.08
Boundary precision	0.36 ± 0.12	0.31 ± 0.11
Seg. covering	0.32 ± 0.12	0.23 ± 0.09
Number of segments	66 ± 10	63 ± 9

Table 5.1: Results of the Eikonal and of the normalized cut algorithm on the test images of the BSD.

It is interesting to discuss in greater detail the stopping criterion used for Eikonal algorithm, in particular the choice of the threshold t. This threshold can indeed be fixed in an adaptative manner depending on the set of weights as observed on the adjacency graph. In our experiments, we followed the following procedure to select the threshold value:

1. Sort the edges increasingly according to their weight.

2. Arbitrarily select a proportion of edges considered to be actual contours, and set the corresponding weight as threshold.

Another interesting thing to notice is that in spite of the refinement step, the obtained segmentation depends on the initial choice of germs. Due to the initial random selection of germs, certain areas of the image may become artificially over-segmented. To address this issue and minimize the impact of the initial seed selection, a straightforward post-processing operation can be performed at the conclusion of the algorithm. This operation involves extracting the highest merging probability observed at the boundaries between adjacent regions, thereby establishing a dissimilarity measure between these regions. Subsequently, the pairs of regions are sorted in ascending order based on their dissimilarity values. As the algorithm progresses, pairs of adjacent regions are evaluated, and if their dissimilarity falls below a specified threshold for the refinement step, they are merged together. It is important to note that each region can only be merged once during this process to prevent the creation of adjacency graph regions with edges exceeding the threshold weight. By implementing this post-processing step, the number of clustered regions can be effectively reduced, while the overall quality of the segmentation is only minimally affected. This post-processing was applied in our ex-

periment, which explains why the number of segments obtained with the Eikonal based algorithm and the normalized cut are slightly different in Tab. 5.1.

5.4 Conclusion

In this chapter, we presented a novel algorithm that performs region merging on the superpixel segmentation of an image. This algorithm is based upon a region adjacency graph representation of the superpixel partition. Each edge in the graph corresponds to a pair of adjacent superpixels and carries a weight accounting for the dissimilarity between these superpixels. The algorithm uses a generalization of the Eikonal equation to the framework of graphs to perform the graph clustering, an idea originally introduced by [Buyssens et al., 2014c]. During Kaiwen Chang's PhD thesis, our main contribution was to propose to learn the weights of the graph based on annotated data. We evaluated the performance of our Eikonal-based approach by comparing it to a classical graph clustering algorithm, namely the normalized cut [Shi and Malik, 2000]. Our results demonstrated improvements over the classical algorithm, particularly in terms of segmentation covering and boundary precision metrics.

The initial goal of the algorithm was to process images from experiments in materials science. For these images, it is difficult to obtain annotated training data, and therefore it is crucial to develop segmentation algorithms requiring a low number of data to be trained. So far, our merging algorithm has only been applied to perform the segmentation of natural images taken from the Berkeley Segmentation Dataset. Hence, an obvious short term perspective of the work presented in this chapter would be to evaluate it on images obtained during physics experiments. We expect these images to be more homogeneous in terms of content than natural images, so that they might require a lower amount of training example to be segmented using a supervised approach.

Chapter 6

Supervised segmentation from synthesized data

As pointed out in the introduction of this part of the manuscript dedicated to image segmentation, state-of-the-art image segmentation techniques currently rely on supervised learning algorithms including convolutional neural networks. However, segmenting images obtained during physics experiments with these algorithms requires to train them on manually segmented images that are often not available. In addition, the images that we want to study are significantly different from natural images, therefore making transfer learning techniques unsuitable to overcome the lack of training data. As a consequence, developing effective segmentation methods that can handle limited annotated data is a critical research area. I describe in this chapter research work conducted in collaboration with David Paulovics (Student, Institut de Physique de Nice), Dr. Frédéric Blanc (Researcher, Institut de Physique de Nice) and Théo Dumont (Student, Mines Paris, PSL University) that seek to train a neural network for performing the segmentation of images obtained during rheology experiments based on a dataset of training images entirely synthesized by a morphological model.

6.1 Context

Image processing techniques are crucial for interpreting the outcomes of rheometry experiments on non-Brownian suspensions. The two most significant quantities for characterizing suspension properties are the concentration fields and the viscosity. Traditionally, these are determined by recording images of the suspension particles during flow at regular time intervals. Figure 6.1 illustrates the principle of a recording device, where a flat laser sheet illuminates a transparent suspension and excites the fluorescence of a dye dissolved in the liquid. A camera perpendicular to the laser plane captures the fluorescent light. An image obtained from this device is displayed in Figure 6.1. In this image, the spherical particles of the suspension appear as black disks.

The measurement of the concentration field relies on the detection and segmentation of the particles present in the image [D'Ambrosio et al., 2021; Snook et al., 2016]. Currently, the algorithms that are utilized for this specific task are relatively "classical" image processing algorithms, that often rely on mathematical morphology tools [Blanc et al., 2013; D'Ambrosio et al., 2021; Dijksman et al., 2017; Dougherty, 1992; Kimme

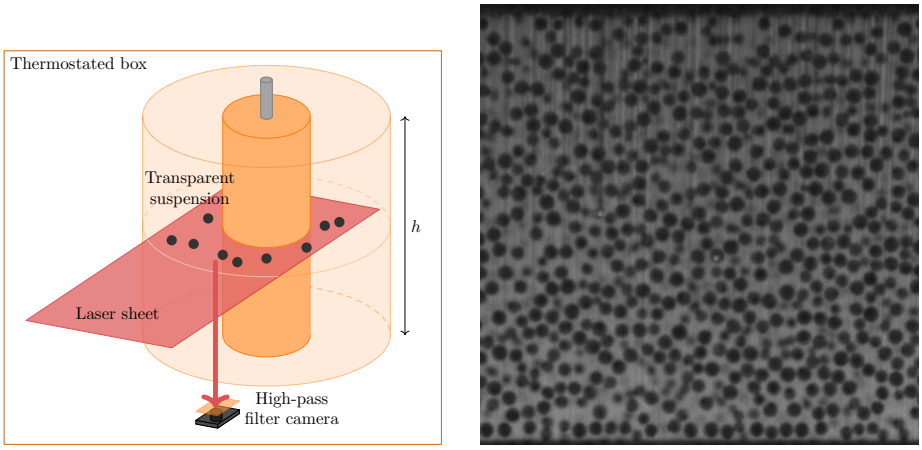

Figure 6.1: Left: schematic view of a recording device [Blanc et al., 2013]. Right: experimental image of a granular suspension from [D'Ambrosio et al., 2021]. The boundaries of the flow cell containing the suspension are visible at the top and at the bottom of the image.

et al., 1975; Snook et al., 2016]. It is usually necessary to properly parameterize these algorithms and to adjust the parameterization depending on the image being processed, which makes the use of these algorithms relatively difficult in practice.

To overcome these difficulties, we developed an image processing algorithm based on a convolutional network [Chen et al., 2017], which brings an advantage over traditional image processing techniques by alleviating the need of updating the algorithm parameterization for each novel image. However, like all supervised learning algorithms, convolutional networks require a dataset of annotated experimental images to be trained. Constructing the ground truth is a time consuming task, which renders the use of supervised algorithms difficult in a lot of problems related to physical applications. In addition, the annotation process can be error prone. The main originality of our approach is that we entirely trained the network on a series of synthetic images generated with morphological models rather than on images of real experiments.

The lack of annotated data has long been identified as a critical issue that prevents the use of state-of-the-art supervised algorithms in many image processing problems. In particular, annotating images obtained during physical experiments is often expensive, which triggers interest in alternative methods where ground truth images are generated in a synthetic manner. The development of such methods is increasingly being studied in the literature [Barisin et al., 2022; Jahanian et al., 2021; Nagy et al., 2022; Ravuri and Vinyals, 2019]. In [Besnier et al., 2020], a generative adversarial network is for instance used to generate a dataset of images similar to those of ImageNet. These generated images are then used to train a classification network. In [Baradad et al., 2021], the authors investigate image generation models that produce images from simple random processes. These generated images are subsequently used as training data for a visual representation learner. Finally, in [Barisin et al., 2022], morphological models are used to generate synthetic 3D images of cracks in concrete that are used to train a supervised segmentation model.

6.2 Segmentation algorithm

6.2.1 Network architecture

To perform the image segmentation, we use the Context Aggregation Network (CAN) introduced in [Yu and Koltun, 2016]. This network is entirely composed of convolutional layers, making it adaptable to any size of input image. Its main particularity is that it gradually aggregates contextual information without losing resolution through the use of dilated convolutions whose field of view increases exponentially over the successive network layers. This exponential growth yields global information aggregation with a very compact structure [Chen et al., 2017; Yu and Koltun, 2016].

The CAN architecture is composed of a set of basis layers $\{L^{(s)}\}_{1\leq s\leq \ell}$. We modify the output of the original network so that it is composed of an image with two channels corresponding to a segmentation mask M for the granular suspension particles and of an image C used to locate the centers of the particles, respectively.

The detailed architecture of the network is presented in Tab. 6.1. Each block $L^{(s)}$ for $s \in [\![2, \ell-2]\!]$ is made of a 3×3 *dilated convolution* with kernel $K^{(s)}$ and dilation parameter $r^{(s)} = 2^{s-1}$, followed by an *adaptive batch normalization* layer $\Psi^{(s)}$ [Chen et al., 2017] and a *leaky rectifier linear unit* (leaky ReLU) non-linear activation function Φ. The depth d of all hidden convolutional layers is kept fixed in the CAN architecture.

Layer $L^{(s)}$ for $s=$		1	2	3	4	5	6	7
Input channels		3	24	24	24	24	24	24
Output channels		24	24	24	24	24	24	2
Conv.	kernel size	3×3	3×3	3×3	3×3	3×3	3×3	1×1
	dilation $r^{(s)}$	1	2	4	8	16	1	1
	padding	1	2	4	8	16	1	0
Adaptive BN		✓	✓	✓	✓	✓	✓	✓
Leaky ReLU		✓	✓	✓	✓	✓	✓	—
Number of parameters		722	5258	5258	5258	5258	5258	50

Table 6.1: Architecture of the Context Aggregation Network (CAN). The total number of trainable parameters for this architecture is 27162 [Paulovics et al., 2023].

The penultimate layer of the network is a classic convolution layer with a filter with size 3×3. The final layer is a 1×1 convolution used to perform dimension reduction. The neural network produces a segmentation mask M and an image C with bi-dimensional Gaussian functions placed at locations corresponding to the centers of the detected particles. We obtain a labeled image of the detected particles by applying a watershed algorithm [Vincent and Soille, 1991] to the segmentation mask M, previously thresholded at the value $1/2$, with the local maxima of C selected as markers.

6.2.2 Generation of synthetic training images

As mentioned previously, the use of convolutional neural networks can be challenging due to the substantial amount of annotated data required for training. Manual annotation is especially arduous when dealing with large quantities of particles in each image,

which can number in the thousands.

To address this issue, we utilize synthetic images created using a morphological model to train the neural network. This approach enables us to acquire training images with corresponding ground truth information without the need for manual annotation of a subset of experimental images. However, generating synthetic images that closely resemble the experimental images is crucial to ensure that the trained neural network architecture has good generalization properties.

Our approach consists in generating gray level images encoded on 8 bits through the use of random morphological models. The image generation proceeds in several subsequent steps:

- **Step 1.** We start by specifying the dimension $w \times h$ of the synthetic image and we build a mask specifying the location of the wall of the flow cell at the image borders. We assign distinct gray levels to the mask and to the interior to obtain an intensity image denoted $\bar{I}$.
- **Step 2.** The experimental images exhibit quasi-periodic stripes patterns. To simulate these patterns, we perturb the intensity at each pixel location $[x, y]$ in the image according to the relationship:

$$\hat{I}_1[x, y] = \bar{I} + \sum_{i=1}^{2} A_i \cos(2\pi f_i \phi(x, y)) \,. \tag{6.1}$$

  In this equation, the amplitudes A_1 and A_2 and the frequencies f_1 and f_2 are specified randomly for each generated image from uniform distributions on specified intervals. The quantities $\phi(x, y)$ defined at each location are independent random variables drawn from a normal distribution with mean x and standard deviation σ and are used instead of the coordinate x in order to add randomness to the geometry of the patterns.
- **Step 3.** We use a Boolean model of disks to simulate a mask for the particles. We recall (see Chap 2.1) that the Boolean model is a grain model obtained by implanting independent random primary grains G' on the germs $\{x_k\}$ of a Poisson point process with intensity θ. The resulting set $\mathcal{G}$ is

$$\mathcal{G} = \bigcup_{x_k \in \mathcal{P}} G'_{x_k} \,, \tag{6.2}$$

  where G'_{x_k} denotes the translated of the primary grain G' at point x_k. In general, the grains of a Boolean model can overlap. To avoid this, we add the grains of the Boolean model sequentially. When a grain intersects a grain which is already present, we simply remove it from the simulation. The primary grains used to construct the model are random disks whose radii are drawn according to a normal distribution with mean $\bar{R}$ and standard deviation σ_R specified for each image. In practice, we draw $\bar{R}$ and σ_R from pre-defined uniform distribution for each generated image. A gray level is finally selected independently for each particle according to an uniform law on a specified interval. The gray level background is set equal to 255. This results in the obtaining of a particle image $\hat{P}$. The synthetic

image is updated by taking the minimum value between the background image $\hat{I}_1$ and the particles image $\hat{P}$:

$$\hat{I}_2[x,y] = \min\{\hat{I}_1[x,y], \hat{P}[x,y]\}\,. \tag{6.3}$$

The particle image $\hat{P}$ is used to generate a binary mask image $\hat{M}$ indicating the presence of the particles in the generated image. In addition, we create an image $\hat{C}$ recording the centers $(x_i, y_i)_{1\leq i\leq N}$ of the N implanted suspension particles by setting:

$$\hat{C}[x,y] = \sum_{i=1}^{N} \frac{1}{2\pi s^2} \exp\left(-\frac{(x-x_i)^2+(y-y_i)^2}{2s^2}\right), \tag{6.4}$$

where s is the size of the selected Gaussian kernel. In this image, each particle is identified by a normalized bi-dimensional Gaussian function. $\hat{M}$ and $\hat{C}$ constitute the ground truth images associated with the synthetic image.

- **Step 4.** To complete the image generation, we add blur to the synthetic image by convolving it with a Gaussian kernel G with standard deviation set equal to 3 pixels, as well as white noise. The synthetic image is therefore described by:

$$\hat{I}[x,y] = \max\{0, (\hat{I}_2 * G)[x,y] + W[x,y]\}\,, \tag{6.5}$$

where the quantities $\{W[x,y]\}_{1\leq x\leq w,\, 1\leq y\leq h}$ are independent centered Gaussian random variables with specified standard deviation.

All the parameters involved in the description of the model can be adapted depending on the set of experimental images under scrutiny. We display in Fig. 6.2 a synthetic image of the suspension constructed with the aforementioned procedure. We remark that synthetic images are visually very close to the suspension images obtained in the experiments.

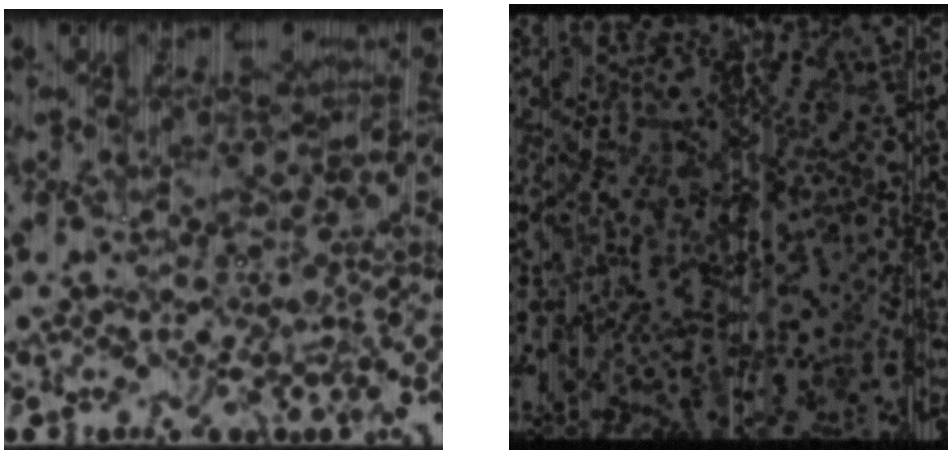

Figure 6.2: Left: experimental image of the suspension. Right: synthetic image of the suspension constructed with our procedure [Paulovics et al., 2023].

6.2.3 Training of the neural network

To train the neural network architecture, we generated a training set and a validation set containing respectively 2240 and 360 synthetic images along with their corresponding ground truth images. We used the Euclidean distance between the output of the network and the ground truth images as loss function to train the algorithm, therefore formulating the segmentation as a regression problem. We relied on data augmentation techniques [Shorten and Khoshgoftaar, 2019] to improve the robustness of the network : the network was fed with random crops of the training images with randomly distorted gray level histogram. To train the neural network, we used the Adam optimizer with a learning rate initially set to 0.1 and a batch size of 8, and we divided the learning rate by a factor of 2 every 50 epochs. We fixed the maximal number of epochs to 400, and retained the weights of the epoch that led to the minimal error on the validation set.

		CAN				**K-means**			
Image	Particles	Recall	Prec.	D[px]	IoU	Recall	Prec.	D[px]	IoU
#1	1339	0.949	0.978	0.68	0.784	0.827	0.954	1.25	0.626
#2	1329	0.944	0.977	0.67	0.771	0.839	0.963	1.46	0.581
#3	964	0.926	0.983	0.72	0.752	0.817	0.962	2.05	0.487
#4	1051	0.947	0.996	0.55	0.869	0.808	0.948	1.75	0.666
#5	487	0.961	0.998	0.5	0.883	0.879	0.949	1.49	0.671
Avg	1034	0.945	0.986	0.62	0.812	0.834	0.955	1.6	0.606

		Otsu thresholding				**Adaptive thresholding**			
Image	Particles	Recall	Prec.	D[px]	IoU	Recall	Prec.	D[px]	IoU
#1	1339	0.819	0.953	1.24	0.639	0.845	0.956	1.2	0.654
#2	1329	0.833	0.967	1.46	0.591	0.875	0.965	1.25	0.612
#3	964	0.812	0.958	1.92	0.5	0.898	0.964	1.3	0.585
#4	1051	0.808	0.948	1.75	0.666	0.866	0.921	1.78	0.662
#5	487	0.830	0.967	1.43	0.707	0.899	0.946	1.52	0.697
Avg	1034	0.821	0.959	1.56	0.621	0.877	0.95	1.2	0.66

Table 6.2: Segmentation metrics for the Context Aggregation Network (CAN), K-means, Otsu and adaptive thresholding algorithms.

6.3 Results and discussion

6.3.1 Evaluation dataset

In order to quantitatively evaluate the results and investigate the generalization capability of the algorithm to real experimental images, we performed manual annotation on 5 experimental images. Each image was annotated by labeling all suspension particles with a disk, providing the center and radius of each particle. While 5 images may seem relatively small in number, these images are large in size and contain a significant number of particles, as shown in Tab. 6.2. Consequently, the detection results were tested against a substantial number of particles, ensuring the statistical validity of the findings. Moreover, we deliberately included a low-quality image in the experimental

dataset, featuring a prominent illumination gradient and noticeable blurriness. This specific image is depicted in Fig. 6.3. On average, each image in the dataset contained 1037 particles. Manual annotation of an image typically required one to two hours, illustrating the significant time investment involved. This highlights the practical challenge of manually annotating an entire set of images for training a convolutional network architecture, which is often unfeasible in many applications.

6.3.2 Detection metrics

In order to evaluate the segmentation results, we developed an approach that allows to establish a one-to-one correspondence between the particles detected by the algorithm (the "detections") and the particles present in the ground truth. We refer the reader interested by more details on this approach to the original preprint [Paulovics et al., 2023].

Once the correspondence established between the particles and the ground truth, we determined the number `fp` of false detections by counting the number of detections not associated with any particle. Similarly, we determined the number `fn` of undetected particles in the ground truth by counting the number of particles in the ground truth left unassociated. The number `tp` of correct detections is to the number of correspondences established between the particles of the ground truth and the detections. The ability of the algorithm to properly detect the suspension particles is described in terms of precision and recall, defined by

$$\text{Recall} = \frac{\texttt{tp}}{\texttt{tp} + \texttt{fn}}, \quad \text{Precision} = \frac{\texttt{tp}}{\texttt{tp} + \texttt{fp}}. \tag{6.6}$$

For all correct detections, we computed different metrics that characterize the quality of the segmentation including the distance $D = \|c_p - \hat{c}_q\|_2$ between the center of the detection and the actual center of the particle as annotated in the ground truth or the intersection over union (IoU) of the particle and the detection, defined by

$$\text{IoU}(\mathcal{P}_p, \mathcal{D}_q) = \frac{\mathcal{P}_p \cap \mathcal{D}_q}{\mathcal{P}_p \cup \mathcal{D}_q}. \tag{6.7}$$

6.3.3 Results and discussion

We compare the results of the convolutional network architecture trained on synthetic data to results obtained with traditional algorithms, including Otsu thresholding [Otsu, 1979], adaptive thresholding [Gonzalez and Woods, 2007] and K-means segmentation [Bishop, 2006]. The results for the CAN neural network, the K-mean segmentation and the Otsu and adaptive thresholding algorithms for the segmentation metrics (precision P, recall R and average IoU between the particle and the detection mask) are reported in Fig. 6.5 and Tab. 6.2, where we also report the average distance D (in pixels) between the centers of the particles and the centroid of their corresponding detection masks.

In general, it can be observed that the convolutional neural network (CNN) performs significantly better than traditional methods in terms of segmentation quality for all

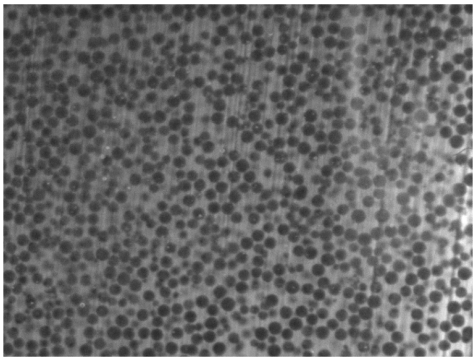

(a) Original image

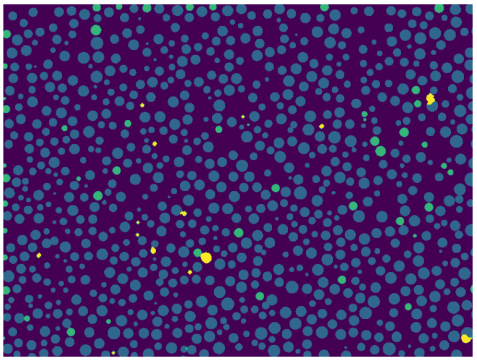

(b) CAN segmentation

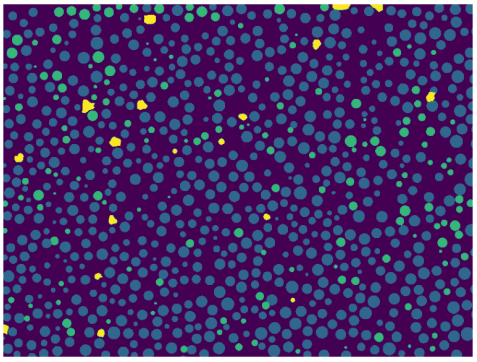

(c) K-means segmentation

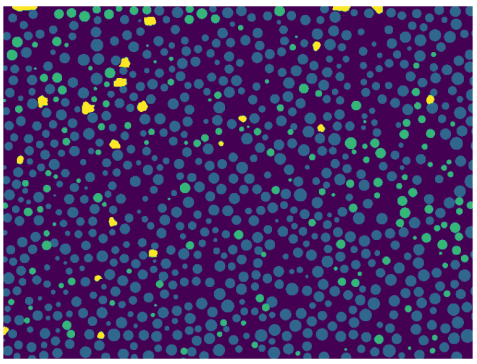

(d) Otsu thresholding

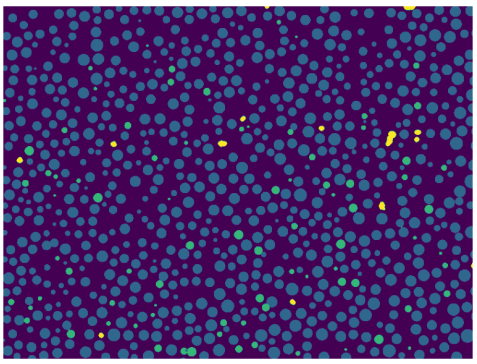

(e) Adaptive thresholding

Figure 6.3: Segmentation results obtained for image #3 (a) with the CAN network (b), K-means (c), Otsu thresholding (d) and adaptive thresholding (e). Correct detections (`tp`) are displayed in blue, false positives (`fp`) in yellow and false negatives (`fn`) in green [Paulovics et al., 2023].

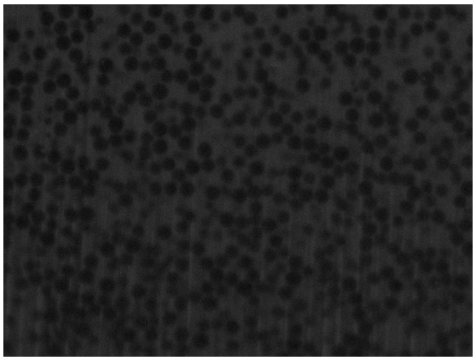

(a) Original image

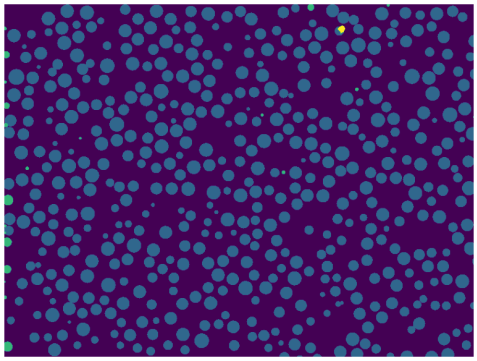

(b) CAN segmentation

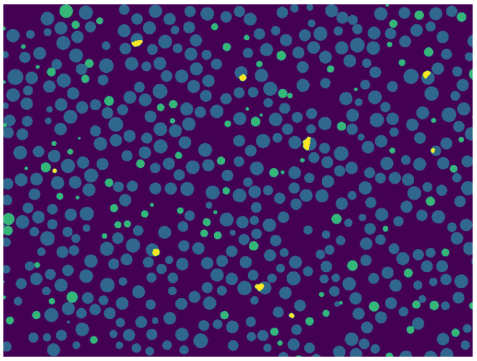

(c) K-means segmentation

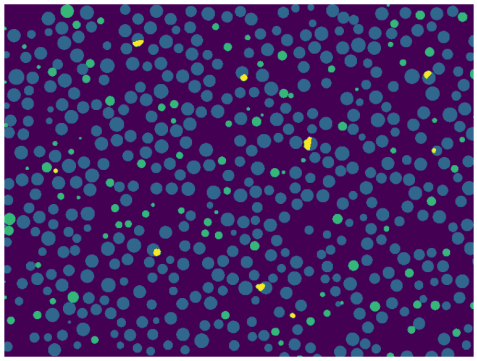

(d) Otsu thresholding

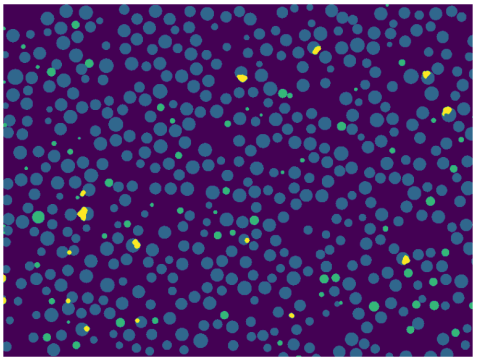

(e) Adaptive thresholding

Figure 6.4: Segmentation results obtained for image #4 (a) with the CAN network (b), K-means (c), Otsu thresholding (d) and adaptive thresholding (e). Correct detections (`tp`) are displayed in blue, false positives (`fp`) in yellow and false negatives (`fn`) in green [Paulovics et al., 2023].

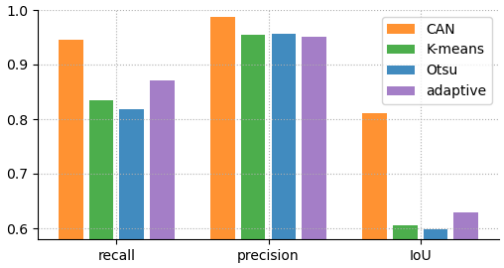

Figure 6.5: Average over the 5 images of Tab. 6.2 of the segmentation metrics (recall, precision, IoU) for the CAN, K-means, Otsu and adaptive thresholding algorithms [Paulovics et al., 2023].

evaluation metrics. However, the algorithm shows slightly lower performance for image #3 in the test dataset, which can be attributed to its lower quality compared to the other images, making segmentation more challenging. Interestingly, all proposed approaches exhibit relatively similar precision results, with values systematically exceeding 0.95. It is mostly for the recall metric that the CNN significantly outperforms other methods. Furthermore, there is a substantial improvement in particle center localization, with an average below one pixel for test images using the CNN architecture. Additionally, the intersection over union (IoU) metric shows significantly better results than conventional algorithms. Notably, adaptive thresholding outperforms K-means or Otsu thresholding, emphasizing the importance of local threshold adaptation rather than relying on global image information.

To illustrate the performance of the compared algorithms, segmentation examples are presented in Figures 6.3 and 6.4. Blue represents correctly detected particles from the ground truth, green indicates false negatives, and yellow highlights false positives superimposed on the ground truth image. It is worth noting that false positives and negatives often occur for particles with ambiguous or inaccurate manual segmentation. In practice, we can therefore consider that the CNN architecture can provide segmentation of similar quality to manual segmentation.

6.4 Conclusion and perspectives

In the study described in this chapter, we introduced a novel approach for segmenting experimental images of a suspension by training a convolutional neural network on synthetic images generated with a morphological model. We demonstrated in particular that the CNN exhibits good generalization properties and outperforms traditional segmentation algorithms when applied to real images of the suspension. From a broader perspective, efficient image processing techniques are essential for making the most of images collected during physical experiments. This study highlights the value of using morphological models to generate reliable training samples in situations where annotated images are unavailable, therefore enabling the use of current state-of-the-art supervised approaches in image segmentation.

6.5 Related work

In addition to this work on the segmentation of images from rheology experiments, I also worked as part of Robin Alais PhD thesis on the development of a convolutional neural network architecture that evaluates the quality of retinal images by assessing the visibility of the macular region. The algorithm deems an image to be of acceptable quality if the macular region is completely visible and within the field of view. In addition, the method can pinpoint the location of the fovea with a maximal error of 0.34 mm for acceptable images. The algorithm is based on a lightweight fully-convolutional network, which is several thousand times smaller than state-of-the-art networks classically used for this specific task, and achieves near-human performance for assessing macula visibility and fovea localization. The main advantage of the method is that it can easily be integrated into portable retinographs, reduce the number of ungradable images and therefore save both patient and physician time. It is a significant step towards automating the screening process for retinal pathologies, including diabetic retinopathy, which constitutes a major global healthcare concern. This work was published in the journal *Biomedical Signal Processing and Control* [Alais et al., 2020].

6.6 Related publications

- D. Paulovics, B. Figliuzzi, T. Dumont and F. Blanc. A supervised algorithm entirely trained on a synthetic dataset to segment granular suspension images. *Preprint*, 2023.
- R. Alais, P. Dokládal, A. Erginay, B. Figliuzzi, E. Decencière. Fast macula detection and application to retinal image quality assessment. *Biomedical Signal Processing and Control*, 55, 101567, 2020.

Part III
Hyperspectral imaging

Hyperspectral imaging is an image modality that captures and processes images at many different wavelengths, allowing for the analysis of objects or scenes in great detail. A hyperspectral sensor collects information in the form of a set of spatial images where each image corresponds to a narrow band of wavelengths of the electromagnetic spectrum. Its spatial resolution is a key characteristic of the sensor: when pixels are too large, the measured spectra is usually a combination of multiple elementary materials present in the observation scene. In this situation, sophisticated post-processing methods have to be applied to identify these materials. The development of these methods has led to the emergence of an active field of research referred to as hyperspectral unmixing (HU). HU algorithms aim at identifying the *spectral signatures* or *endmembers* of the elementary materials and at quantifying their relative contributions or *abundances* to the hypersensor measurements.

A classical approach in HU is to assume that the measured spectra are linear combinations of the endmembers present in the observed surface weighted by their respective abundances. The linear mixing model has a straightforward physical interpretation: it corresponds to situations where each incoming light ray only interacts with a single endmember before reaching the sensor. A large corpus of methods based upon geometrical or statistical approaches have been considered to solve the problem of linear unmixing, which is now relatively well understood [Bioucas-Dias et al., 2012]. Still, a number of problems remain open:

- *Endmembers variability.* The spectral signatures of the elementary materials often exhibit intrinsic variability within the observations. The issue of endmembers variability is still an active ongoing research topic in HU. Two approaches are classically employed in the literature for dealing with endmembers variability: modeling endmembers as sets, or modeling endmembers as statistical distributions [Zare and Ho, 2013].
- *Highly mixed observations.* When the mixing is linear, the hyperspectral observations are embedded in a low-dimensional simplex. Geometrical methods try to identify this simplex by assuming the presence of observations either at the vertices or at the boundary of the simplex. These assumptions are however not satisfied when the observations are highly mixed. Statistical approaches have notably been proposed in the literature to deal specifically with this issue [Bioucas-Dias et al., 2012].

- *Nonlinear mixing models.* In situations where the incoming electromagnetic wave interacts with more than one material present at the scene, the mixing model is usually nonlinear. The development of HU algorithms adapted to nonlinear mixing remains an active topic of research [Heylen et al., 2014].

Hyperspectral imaging is traditionally associated with remote sensing applications, but is increasingly used in novel applications including surface characterization of chemical imaging. The optical interactions leading to the hyperspectral observations in these applications are very different to the ones classically occurring in remote sensing, which foster the development of dedicated signal processing methods to process them. From a broader perspective, the development of hyperspectral sensors for applications in chemistry and materials science raises novel questions on a theoretical level: the mixing models encountered in these applications are often strongly non-linear, the endmember variability is enhanced and the observations are often highly mixed. This topic is the last axis of research that I wish to explore in the coming years, in connection with the thematic of image processing for the materials engineering. My main contributions to this research field to day is the PhD thesis of Tarek Zenati, which I co-supervised along with Jesús Angulo (Center for Mathematical Morphology, Mines Paris, PSL University) between 2019 and 2023. During this PhD thesis, we worked on the development of algorithmic methods enabling the detection and the chemical characterization of selective oxides that can form on steel sheet surfaces during their manufacturing [Zenati et al., 2022b,c]. The PhD thesis was funded by ArcelorMittal Research through the CIFRE program. I summarize in chapter 7 the methodology that we developed.

Chapter 7

Oxide detection using hyperspectral imaging

During their production, a number of alloying elements including manganese (Mn), aluminum (Al) or silicon (Si) are usually incorporated into steel sheets. These additives can subsequently cause the formation of selective oxides, which can significantly impact the surface properties of the produced sheets. The wettability of the surface can for instance be altered by the presence of oxides, which is a significant concern when additional coatings have to be deposited on the steel to protect it from oxidation. Preventing the formation of selective oxides at the steel surface is therefore essential, which explains why the development of image processing techniques aimed at detecting their formation has been an active topic of research over the past few years in the steel industry [Ferté, 2014; Ham et al., 2016; Hui Ham et al., 2016; Masci et al., 2012, 2013; Zenati et al., 2022a,b,c].

During the course of the PhD thesis of Tarek Zenati, we investigated the possibility of using hyperspectral observations to perform the detection of selective oxides at the surface of steel sheets. One of the main issue is that due to the limited spatial resolution of hyperspectral sensors, the hyperspectral observations measured at each surface location are often highly mixed and potentially incorporate the contribution of several oxide species. The inability to distinguish between these oxides is a barrier to the development of surface analysis techniques based on hyperspectral observations.

Our primary contribution during the PhD was to propose a novel approach for detecting the formation of oxide mixtures on steel surfaces and for estimating the thickness of the corresponding oxidized layers. Our approach involves the prior identification of all potential mixtures of oxides that can form on the surface and the computation of a comprehensive dictionary of all associated spectra. This allowed us to bypass the issue of the nonlinear mixing model by reformulating the problem of oxide determination as a sparse regression problem.

The outline of this chapter is as follows. In section 7.1, we present the optical model used to describe the formation of the hyperspectral measurements conducted on the steel. We introduce next the sparse regression based approach in section 7.2. Finally, we summarize in section 7.3 some experimental results obtained during the PhD study.

7.1 Optical model

We describe in this section the optical model used to describe the formation of the hyperspectral observations. This model relies on the assumption that the substrate is covered by homogeneous layers constituted of a mixture of several oxides.

Permittivity models A major issue related to the modeling is that it requires a precise knowledge of the permittivity of the oxides that can form at the surface for an extended range of wavelengths. During their formation, the oxide layers are intrinsically heterogeneous and contain several chemical species. As a consequence, the permittivity of the resulting heterogeneous material depends on the permittivities of the pure oxides, on their crystallinity and on the microstructure geometry all together. A potential solution to account for the presence of heterogeneous oxides could be to compute an effective permittivity from the permittivities of the elementary oxides from the mixture by using a homogenization model. This approach is however difficult to follow in practice, for two main reasons: 1. it is difficult to fully account for the complexity of the microstructure geometry with a simple homogenization model and 2. computing the exact permittivity of the pure oxides is in itself a difficult task.

To overcome these difficulties, we used experiments conducted on oxidized samples in the PhD study [Hui Ham, 2018] to determine the permittivities of elementary mixtures of oxides that can form at the surface. To account for the potentially large range of chemical compositions of the oxide layers that form at the surface, we assumed that these layers contain several elementary mixtures. To compute the effective permittivity of the resulting heterogeneous oxide, we used the self-consistent model [Tolstoy et al., 2003], for which the effective permittivity ϵ_{SC} is defined implicitly through the equation

$$\sum_{i=1}^{n} p_i \frac{\epsilon_i - \epsilon_{SC}}{\epsilon_i + 2\epsilon_{SC}} = 0, \text{ with } \sum_{i=1}^{n} p_i = 1, \tag{7.1}$$

where n is the number of elementary oxides, and p_i and ϵ_i the volume fraction and the permittivity of the i-th elementary oxide, respectively.

Reflection on multiple layers The model describing the formation of the hyperspectral observations relies on a simplified representation of the surface geometry, described as a stack of homogeneous oxide layers parallel to the surface. We denote by xy the plane parallel to the surface and by z the orthogonal direction. According to the superposition principle, we can decompose the electric field E onto a field E_f that propagates in the forward direction and a field E_b that propagates in the backward direction [Tolstoy et al., 2003]:

$$E = E_f e^{ik_z z} + E_b e^{-ik_z z}. \tag{7.2}$$

In Eq. (7.2), k_z denotes the component of the wave vector along direction z. Let us denote by z_k the location of the interface separating layers k and $k+1$. At each layer interface, according to Young-Fresnel equations, we have

$$\begin{pmatrix} E_f(z_k^-) \\ E_b(z_k^-) \end{pmatrix} = \underbrace{\frac{1}{t_{k,k+1}} \begin{pmatrix} 1 & r_{k,k+1} \\ r_{k,k+1} & 1 \end{pmatrix}}_{T_{k,k+1}} \begin{pmatrix} E_f(z_k^+) \\ E_b(z_k^+) \end{pmatrix}, \tag{7.3}$$

where $t_{k,k+1}$ and $r_{k,k+1}$ are the Fresnel coefficients for transmission and reflection, respectively. These coefficients can be directly linked to the complex permittivity ϵ of the oxides constituting the layers and to the incidence angle ϕ_k at the interface. Let us denote by ξ_k the generalized complex index of refraction, defined by

$$\xi_k = n_k \cos\phi_k, \tag{7.4}$$

where n_k is the complex optical index of the oxide in the k-th layer. Then, when the electromagnetic wave is s-polarized, the Fresnel coefficients are given by

$$r^{(s)}_{k,k+1} = \frac{\xi_{k+1} - \xi_k}{\xi_{k+1} + \xi_k}, \quad t^{(s)}_{k,k+1} = \frac{2\xi_k}{\xi_{k+1} + \xi_k}. \tag{7.5}$$

Similarly, when the electromagnetic wave is p-polarized, the Fresnel coefficients are

$$r^{(p)}_{k,k+1} = \frac{\epsilon_k\xi_{k+1} - \epsilon_{k+1}\xi_k}{\epsilon_k\xi_{k+1} + \epsilon_{k+1}\xi_k}, \quad t^{(p)}_{k,k+1} = \frac{2n_k n_{k+1}\xi_k}{n_k^2\xi_{k+1} + n_{k+1}^2\xi_k}. \tag{7.6}$$

The propagation of the electric fields within the k-th layer is in turn described by

$$\begin{pmatrix} E_f(z_k^+) \\ E_b(z_k^+) \end{pmatrix} = \underbrace{\begin{pmatrix} e^{ik_z h_k} & 0 \\ 0 & e^{-ik_z h_k} \end{pmatrix}}_{T_k(h_k)} \begin{pmatrix} E_f(z_{k+1}^-) \\ E_b(z_{k+1}^-) \end{pmatrix}, \tag{7.7}$$

where $h_k = z_{k+1} - z_k$ is the thickness of the k-th layer. A global equation describing the field propagation for the multi-layered geometry can be obtained by computing the product between all transmission matrices. For a geometry with K layers, we have, since there is no electric field propagating backward in the steel,

$$\begin{pmatrix} E_i(0^+) \\ E_r(0^+) \end{pmatrix} = \underbrace{\prod_{k=1}^{K} T_{k-1,k}T_k(h_k)T_{k,k+1}}_{T(h_1,\ldots,h_K)} \begin{pmatrix} E_t(z_{K+1}) \\ 0 \end{pmatrix}, \tag{7.8}$$

where E_i, E_r and E_t correspond to the incident, reflected and transmitted fields, respectively. In Eq. (7.8), the indices $k = 0$ and $k = K + 1$ are associated to the air and to the steel substrate. In the presence of a single layer of oxide on a steel substrate, the system is described by:

$$\begin{pmatrix} E_i(0^+) \\ E_r(0^+) \end{pmatrix} = \underbrace{T_{0,1}T_1(h)T_{1,2}}_{T(h)} \begin{pmatrix} E_t(h^-) \\ 0 \end{pmatrix}, \tag{7.9}$$

where h is the layer thickness. The global transmission and reflection coefficients t and r for the entire system can easily be computed from the coefficients of $T(h)$ as a function of the layer thickness:

$$r = \frac{[T(h)]_{21}}{[T(h)]_{11}}, \qquad t = \frac{1}{[T(h)]_{11}}. \tag{7.10}$$

It is finally straightforward to compute the reflectance, defined as

$$R(h) = |r(h)|^2. \tag{7.11}$$

7.2 Sparse unmixing approach

We describe in this section our methodology to invert the optical model described in section 7.1. The problem can be summarized as follows: given a hyperspectral observation, is it possible to identify the chemical composition and the thickness of the oxide layers that led to this observation? In what follows, we make the assumption that the measured reflectance spectra result from a linear combination of a small number of elementary spectra $\vec{S}_i$ corresponding to a layer of thickness h_i of some mixture i of k oxides with respective volume fractions $\vec{v_i} := (v_i^{(1)}, \ldots, v_i^{(k)})$:

$$\vec{S}(\lambda) = \sum_{i=1}^{n} \alpha_i \vec{S}_i(h_i, \vec{v_i}; \lambda). \tag{7.12}$$

In expression (7.12), λ corresponds to the wavelength and the coefficients $(\alpha_i)_{i=1,\ldots,n}$ are the abundances of the different mixtures of oxides and can be considered as effective surface fractions for the oxidized layers. This model can be interpreted as an approximation of the optical interactions occurring in the geometrical setting described in Fig. 7.1, where several mixtures of oxide with distinct thicknesses are present in the observation region.

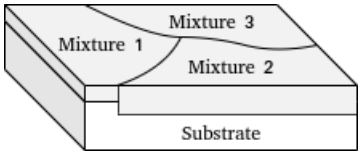

Figure 7.1: Multiple layers of oxides on a steel substrat [Zenati et al., 2023].

A possible approach for characterizing the presence of an oxidized layer is to calculate, for each mixture of oxide potentially present at the surface, the reflectance spectra corresponding to different reference thicknesses and different volume fractions for the elementary oxides present in the layer. The pre-calculated spectra can then be stored in a dictionary D, which takes the form of a matrix of size $L \times N$. The columns of D stack the calculated spectra, L designating here the number of wavelengths. The problem of determining the elementary reflectance spectra is then reformulated as the search for a linear combination involving a small number of elementary spectra from the dictionary D that allows to reconstruct the observed spectrum $\vec{S}$ with good accuracy.

Let $\vec{S}$ be an experimental reflectance observation. Our aim is to find a sparse vector $\vec{\alpha}$ such that $\vec{S} = D\vec{\alpha}$ in order to represent $\vec{S}$ as a linear combination of a small number of pre-calculated oxide spectra taken from the dictionary. From a mathematical perspective, this can be formulated as the following optimization problem:

$$\hat{\vec{\alpha}} = \underset{\vec{\alpha} \in \mathbb{R}^N}{\arg\min} \|\vec{S} - D\vec{\alpha}\|^2 + \gamma \|\vec{\alpha}\|_0. \tag{7.13}$$

In formulation (7.13), the quantity $\vec{\alpha}$ corresponds to the coefficients of a vector selecting the elements of the dictionary D used for the reconstruction of the observed spectrum $\vec{S}$. $\|\vec{\alpha}\|_0$ refers to the number of non-zero components of the vector $\vec{\alpha}$:

$$\|\vec{\alpha}\|_0 = \#\{i \in [\![1, N]\!], \alpha_i \neq 0\}. \tag{7.14}$$

Formulation (7.13) can be seen as the search for a compromise between a vector of coefficients $\vec{\alpha}$ that properly reconstructs the reflectance $\vec{S}$ through the minimization of the term $\|\vec{S} - D\vec{\alpha}\|^2$, while remaining parsimonious through the minimization of the term $\|\vec{\alpha}\|_0$. The positive parameter γ establishes a trade-off between the parsimony of $\vec{\alpha}$ and the accuracy $\|\vec{S} - D\vec{\alpha}\|^2$ of the reconstruction: a high value of γ leads to parsimonious solutions to the detriment of the reconstruction accuracy. In practice, problem (7.13) is difficult to solve due to its non-convexity. To obtain an exact solution, it is necessary to consider all possible combinations of parameters, which leads to a combinatorial complexity. Different algorithms are nevertheless available in the literature to obtain approximate solutions to the problem, including the *orthogonal matching pursuit* (OMP) [Bourguignon et al., 2015; Mallat and Zhang, 1993]. We discuss in the next sections two alternative reformulations of problem (7.13) that present the advantage of being computationally tractable.

7.2.1 Path LASSO approach

A first alternative to formulation (7.13) is to penalize the problem by the ℓ^1 norm instead of the ℓ^0 norm, leading to the new optimization problem:

$$\hat{\vec{\alpha}} = \underset{\vec{\alpha} \in \mathbb{R}^N}{\arg\min} \|\vec{S} - D\vec{\alpha}\|^2 + \gamma \|\vec{\alpha}\|_1, \tag{7.15}$$

where the norm l^1 of $\vec{\alpha}$ is given by $\|\vec{\alpha}\|_1 = \sum_{i=1}^{N} |\alpha_i|$. The main advantage of this second formulation, known under the name of *Least Absolute Shrinkage and Selection Operator* (LASSO [Tibshirani, 1996]) is that (7.15) is a convex optimization problem. Efficient algorithms are therefore available to solve it [Boyd et al., 2004; Bubeck et al., 2015; Efron et al., 2004].

The Lasso formulation is only relevant when applied to normalized observations. All spectra are therefore subtracted their mean and divided by their standard deviation:

$$\vec{S}(\lambda) \leftarrow \frac{\vec{S}(\lambda) - \bar{S}}{\sigma_S}, \tag{7.16}$$

where

$$\bar{S} = \frac{1}{L} \sum_{i=1}^{L} \vec{S}(\lambda_i), \qquad \sigma_S^2 = \frac{1}{L} \sum_{i=1}^{L} (\vec{S}(\lambda_i) - \bar{S})^2. \tag{7.17}$$

In practical situations, fixing the value of the parameter γ that controls the trade-off between the parsimony of the solution and the reconstruction accuracy is not straightforward. The most obvious solution is to fix the value of γ *a priori*. An alternative is to start from a high value of γ and to solve problem (7.15) for decreasing values of the parameter γ, an approach referred to as *path Lasso* in the literature. Starting with a high value of γ yields a highly parsimonious $\hat{\vec{\alpha}}(\gamma)$ solution. By progressively reducing the value of γ, we obtain solutions that are less sparse but that lead to higher reconstruction accuracy. We retain the solution satisfying

$$\sum_{i=1}^{N} \hat{\vec{\alpha}}(\gamma)_i \simeq 1. \tag{7.18}$$

The fact that the coordinate i of the solution $\hat{\vec{\alpha}}(\gamma)$ is non-zero indicates that it was necessary to select the oxide corresponding to the column i of the dictionary D to reconstruct the observed spectrum. As the reconstruction problem is penalized by the ℓ^1 norm, only a small number of coefficients are non-zero, and only a small number of oxides are involved in the reconstruction of the spectrum. The criterion (7.18) is relatively easy to interpret: when it is satisfied, the quantity $D\hat{\vec{\alpha}}(\gamma)$ can be seen as a weighted average of different elements of the dictionary.

7.2.2 Group LASSO approach

A second alternative to formulation (7.13) is to rely on the group Lasso algorithm. The group Lasso algorithm, introduced in [Yuan and Lin, 2006], allows to perform a sparse selection of groups of variables. In the problem that we consider, the hyperspectral observation in the dictionary D can naturally be divided in G distinct groups $(\mathcal{G}_1, \ldots, \mathcal{G}_G)$ by selecting all spectra obtained with a given mixture of oxides, regardless of the thickness of the oxide layer and of the volume fraction of the elementary oxides constituting the mixture. The group Lasso formulation of the oxide detection problem amounts to solve the following optimization problem:

$$\hat{\vec{\alpha}} = \underset{\vec{\alpha}\in\mathbb{R}^G}{\arg\min} \|\vec{S} - \sum_{g=1}^{G} D\vec{\alpha}_g\|^2 + \gamma \sum_{g=1}^{G} \sqrt{p_g}\|\vec{\alpha}_g\|_2. \tag{7.19}$$

In expression (7.19), the quantity $\vec{\alpha}_g$ corresponds to the restriction of the vector $\vec{\alpha}$ to the entries of the dictionary D that belong to the g-th group

$$\forall g \in [\![1, G]\!], \quad \vec{\alpha}_g[i] = \begin{cases} \vec{\alpha}[i] & \text{if } i \in \mathcal{G}_g \\ 0 & \text{otherwise,} \end{cases} \tag{7.20}$$

where the quantities $(p_g)_{g\in[\![1,G]\!]}$ denote the cardinal of each group. We can note that when each group contains only one variable, the formulation (7.19) reduces to the Lasso (7.15). The main interest of the group Lasso formulation is that it allows to impose the sparsity on the selection of the oxide mixtures that can be present at the surface, regardless of their thicknesses and of the relative contribution of the elementary oxides to the mixture. The group Lasso algorithm therefore tries to recover the hyperspectral observation with a limited number of oxides mixtures, but does not impose sparsity within the spectra corresponding to similar oxide mixtures. It is possible to impose a sparse selection of the spectra within a given group by considering the sparse group Lasso formulation, which consists in adding a ℓ^1 penalty to the vector $\vec{\alpha}$. The sparse group Lasso formulation reads

$$\hat{\vec{\alpha}} = \underset{\vec{\alpha}\in\mathbb{R}^G}{\arg\min} \|\vec{S} - \sum_{g=1}^{G} D\vec{\alpha}_g\|^2 + \gamma_1 \sum_{g=1}^{G} \sqrt{p_g}\|\vec{\alpha}_g\|_2 + \gamma_2\|\vec{\alpha}\|_1 \tag{7.21}$$

and can be used to perform the spectra identification. Its main difference with the path Lasso algorithm is that it imposes sparsity on both the selection of the oxides mixtures that can intervene in the observation recovery through the group penalty and on the selection of spectra within a given group through the ℓ^1 penalty.

7.2.3 Algorithm

Before applying the proposed methodology, a prerequisite is to pre-compute, for each mixture of oxides potentially appearing at the substrate surface, the reflectance spectra corresponding to different reference thicknesses and different chemical compositions, before stacking all pre-computed spectra in a dictionary D. The estimation and characterization of the oxides present on the observed reflectance spectrum are then carried out in two steps:

1. The first step of the algorithm consists in identifying a small number of layer of oxides in the dictionary that allows to reconstruct the spectral observation. This can be done either with the path Lasso formulation or the group Lasso formulation. The purpose of this step is to identify a subset of spectra present in the observation.

2. In the second step of the algorithm, we use the subset $\mathcal{S}$ of spectra from D identified by the path or group Lasso algorithm to perform the recovery of the hyperspectral observation $\vec{S}$. This is done by considering the least-square minimization problem

$$\hat{\vec{\alpha}}(\gamma) = \arg\min_{\vec{\alpha}} \|\vec{S} - D_{\mathcal{S}}\vec{\alpha}\|_2^2 \quad \text{s.t} \quad \vec{\alpha} \succeq 0, \quad \sum_{i=1}^{\text{Card}(\mathcal{S})} \vec{\alpha}_i = 1. \tag{7.22}$$

As a general rule, nothing prevents the path or group Lasso algorithms from selecting columns of D corresponding to the same mixture of oxides, but for different thicknesses and chemical compositions. We therefore estimate the thickness and the chemical composition of each mixture of oxides effectively selected by calculating their average, weighted by the value of the coefficients of $\hat{\vec{\alpha}}(\gamma)$. More precisely, for a given oxide $\mathcal{O}$, let us denote by $\hat{\vec{\alpha}}(\gamma, \mathcal{O})$ the vector such that:

$$\forall i \in [\![1, N]\!], \quad \hat{\alpha}(\gamma, \mathcal{O})_i = \begin{cases} \hat{\alpha}(\gamma)_i \text{ if the i-th column refers to oxide } \mathcal{O} \\ 0 \text{ otherwise.} \end{cases} \tag{7.23}$$

Let $h \in \mathbb{R}^N$ be the vector containing, for all $n \in [\![1, N]\!]$ the thicknesses of the pre-calculated spectra in the columns of D. For a given mixture of oxides $\mathcal{O}$, we estimate the thickness of the oxide layer to be:

$$h(\mathcal{O}) = \frac{h^T \hat{\vec{\alpha}}(\gamma, \mathcal{O})}{\|\hat{\alpha}(\gamma, \mathcal{O})\|_1} = \frac{1}{\|\hat{\vec{\alpha}}(\gamma, \mathcal{O})\|_1} \sum_{i=1}^{N} h_i \hat{\vec{\alpha}}(\gamma, \mathcal{O})_i. \tag{7.24}$$

7.3 Experiments

7.3.1 Numerical experiments

We present and discuss in this section numerical experiments conducted to evaluate the proposed algorithm. In these experiments, we use the optical properties of $K = 5$

distinct oxides, including aluminum oxide (Al_2O_3), silicon dioxide (SiO_2), manganese oxide (MnO) and mixtures of manganese and aluminum oxides (Mn-Al-O, Mn*-Al-O) with different stochiometries. The permittivities of the MnO, Mn-Al-O and Mn*-Al-O oxides were obtained using the approach described in section 7.1, while the permittivities of Al_2O_3 and SiO_2 were directly obtained from the literature [Brun et al., 2013; Meneses et al., 2014]. Based on the optical properties of the oxides, we construct a dictionary D for 8 reference thicknesses ranging from 20 to 300 nm. For each pair of oxide and thickness, we compute the corresponding reflectance for 186 wavelengths ranging uniformly from 6.5 to 12.6 microns. We incorporate the reflectance of the substrate in the dictionary. The shape of the dictionary D is therefore 41×186.

In our experiment, we randomly select k oxides, with k selected randomly between 1 and 3, and k random thicknesses $(h_1, \ldots, h_k)$ between 10 nm and 350 nm according to a uniform law. Surface fractions $(\alpha_1, \ldots, \alpha_k)$ are drawn for the selected oxides according to a uniform Dirichlet distribution with dimension k. Based upon model (7.12), we compute then the optical response of the corresponding geometry. This procedure allows us to simulate the obtaining of experimental reflectance spectra. Next, using the dictionary D, we rely on the algorithm described in section 7.2 to estimate the proportion of each one of the considered oxides at the surface and their respective thicknesses. The estimation is then compared to the actual geometric setting. The results of the experiments are summarized in table 7.1. We obtain a precision of 0.91 and a recall of 0.92 when multiple oxides are present at the surface with the path Lasso algorithm, and a precision of 0.92 and a recall of 0.89 with the group Lasso algorithm.

We display in Fig. 7.3 the results of the estimation on two generated samples. The first sample corresponds to a mixture of Al_2O_3 (5.7%, estimation: 5.8%) and SiO_2 (94.3%, estimation: 94.2%) with respective thicknesses 234 nm (estimation: 200 nm) and 215 nm (estimation: 214 nm). The second sample corresponds to a mixture of MnO (31.3%) and Mn*-Al-O (68.7%) with respective thicknesses 179 nm and 216 nm. However, our estimation yields a mixture of substrate (10.3%) and Mn*-Al-O (89.7%) with thickness 182 nm. In this case, the proposed approach fails to properly estimate the correct proportion of oxides and even to identify one of the actual oxides present at the surface. However, as can be observed in figure 7.3, the spectrum is still reconstructed with high accuracy. This example illustrates one intrinsic difficulty related to our approach: the difficulty arises when the elements of the dictionary are highly correlated, meaning that they share similar features. In such cases, multiple combinations of spectra from the dictionary can result in the same reconstruction, which can lead to inaccuracies in estimating the correct volume fraction of oxides in the observed reflectance spectrum. The ℓ^1 penalization may not be sufficient to overcome this difficulty and ensure the accuracy of the reconstruction. This highlights the importance of carefully selecting a suitable dictionary and regularization technique when applying this approach.

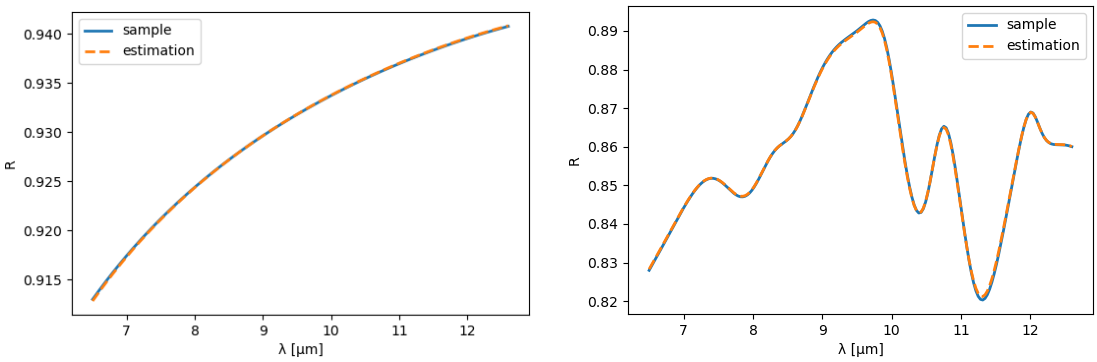

Figure 7.2: Spectrum recovery in the presence of a single layer of oxide (left: SiO_2, 120 nm; right: Mn-Al-O, 183 nm) [Zenati et al., 2023].

	Path Lasso		Group Lasso	
	Single oxide	Multiple oxides	Single oxide	Multiple oxides
Recall	1	0.917	0.992	0.885
Precision	0.956	0.910	0.954	0.924
Thickness error	2.6 nm	10.7 nm	2.6 nm	6.6 nm

Table 7.1: Detection metrics

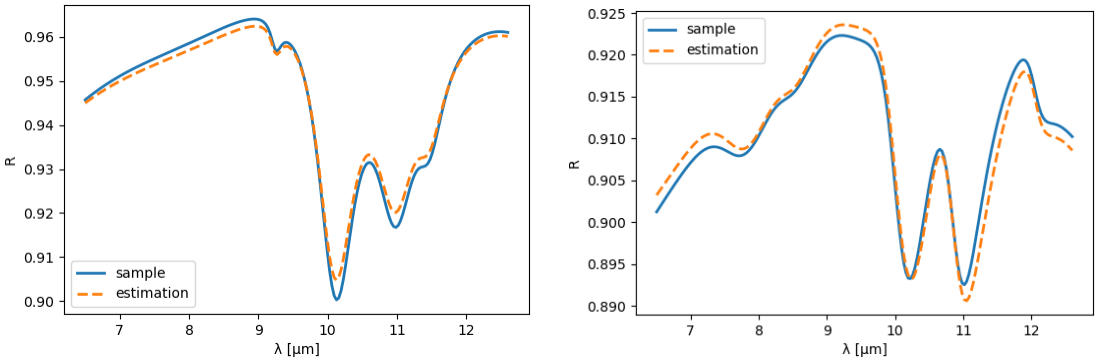

Figure 7.3: Reflectance recovery in the presence of a multiple layers of oxides (left: mixture of SiO_2 and Mn-Al-O; right: mixture of MnO and Mn-Al-O*) [Zenati et al., 2023].

7.3.2 Experiments

Samples preparation and measurements

We summarize in this section experiments conducted on steel samples to validate the algorithm. During these experiments, we subjected steel sheets samples to two distinct thermal cycles, referred to as HP1180 and DP980, to control the formation of distinct species of oxides at their surface, while varying the value of the dew point to obtain different structures of oxides at the surface.

X-ray photoelectron spectrometry (XPS) and Glow-discharge optical emission spectroscopy (GD-OES) analysis were conducted to characterize the surface composition. The XPS analysis enabled us to obtain quantitative information on the atomic species directly located at the surface of the samples. The GD-OES analysis, conducted by progressively eroding the surface of the samples, provided us with information on the atomic species across the depth of the samples. The experiments allowed us to obtain $K = 5$ oxidized samples identified by the references 19 and 22 for the HP1180 thermal cycles (with respective dew points of $-40°$C and $-20°$C) and 5, 9 and 11 for the DP980 thermal cycle (with respective dew points of $-50°$C, $-35°$C and $-5°$C). Overall, we can notice that the thermal cycle HP1180 favors the obtaining of a mixture of Mn_2SiO_4 and SiO_2 oxides referred to as Mn-Si-O mixture in the following, while the thermal cycle DP980 favors the growth of manganese oxides (MnO) at the surface. The high resolution spectra from XPS analysis confirm the presence of these oxides. Scanning electron microscopy (SEM) images of cross-sections of the samples were also registered during the experiments to provide a qualitative view of geometry of the oxide layers.

Results and discussion

The patterns of oxides in the experimental samples are complex and cannot be accurately described by the simple geometric model presented in section 7.1. This model assumes that the oxides form layers with uniform thicknesses on the surface, which only partially approximates the actual geometry. Furthermore, comparing the estimates obtained from our approach to the results of X-ray photoelectron spectroscopy (XPS) measurements is challenging because XPS only characterizes the presence of oxides over a few nanometers, making it difficult to quantitatively compare the results. Instead, our goal is to qualitatively verify that the oxides detected on the samples are also present in XPS and glow discharge optical emission spectroscopy (GD-OES) analyses and that the geometry matches the observations made using scanning electron microscopy (SEM) at different depths and at the surface. For instance, SEM images can reveal the presence of oxide covering layers, which should also be detected by our estimation. It is also possible to qualitatively compare the estimations obtained on different samples. Samples 19 and 22 are made of the same steel grade but have undergone different thermal cycles, resulting in a significant amount of silica oxides detected by XPS, while other samples should contain more manganese oxides.

Our proposed approach requires the construction of a dictionary that stores reference reflectance spectra pre-computed with the optical model. To construct the dictionary, we assume that the oxides that can form at the steel surfaces can be of three distinct

natures: 1. a mixture of substrate, SiO_2 and Mn-Si-O, 2. a mixture of substrate, Mn-AlO and Al_2O_3, and 3. a mixture of substrate and MnO. We assume that these mixtures can be observed with different thicknesses and volume fractions for each one of the individual oxide components and we use the self-consistent approximation to compute the effective permittivity of the material. We compute the dictionary of spectra used for the inversion procedure by calculating the optical response of a layer of each considered mixture of oxides for a given range of 8 thicknesses and different combinations of volume fractions. This leads to a dictionary containing $K = 246$ spectra. We present and discuss below the estimation results obtained for one experimental sample for each thermal cycle. We refer the reader to the original article [Zenati et al., 2023] for a complete overview of our results.

HP1180 thermal cycle We present here the estimation results obtained for one of the samples obtained through the thermal cycle HP1180 (sample 19). The estimation conducted from the hyperspectral measurements indicates the presence of a layer constituted of a mixture of steel ($v_f = 0.27$), Mn-Si-O ($v_f = 0.69$) and SiO_2 ($v_f = 0.04$). The thickness of the layer is estimated to be 139 nm, and its relative abundance is $\alpha = 0.89$, which indicates that the oxide layer almost covers the entire sample. Traces of Al_2O_3 are also detected at the surface, with a relative abundance of 0.11. These estimations are consistent with observations of the sample conducted on SEM images, which confirm the presence of a covering oxide layer with a thickness ranging from 100 to 150 nm. They are also in good agreement with the XPS measurements conducted at the surface and the GD-OES measurements conducted across the depth of sample, which indicates the presence of aluminum oxides near the surface.

	Mixture 1		Mixture 2			Mixture 3		
	Sub.	MnO	Sub.	Mn-Si-O	SiO2	Sub.	Mn-Al-O	Al2O3
v_f	-	-	0.27	0.69	0.04	0	0	1
α	-		0.89			0.11		
h	-		139 nm			190 nm		

Table 7.2: Estimation results obtained for sample 19.

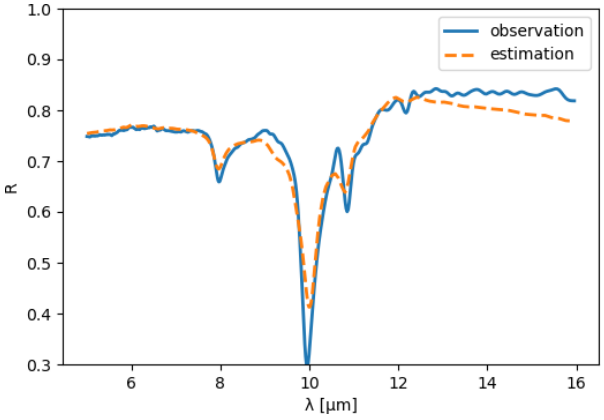

Figure 7.4: Recovery of the observations for sample 19 [Zenati et al., 2023].

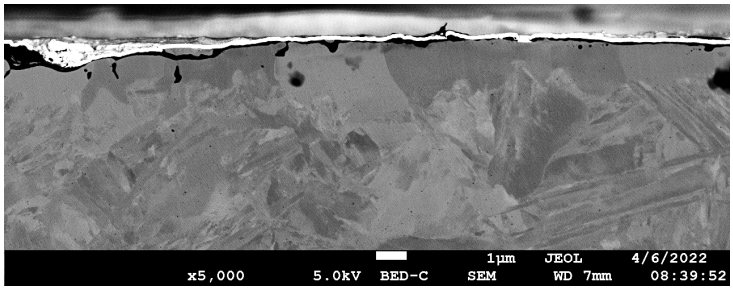

Figure 7.5: Scanning electron microscopy image of a slice of sample 19. The covering layer of oxides appears in black on the image [Zenati et al., 2023].

DP980 thermal cycle We present here the estimation results obtained for one of the samples obtained through the thermal cycle DP980 (sample 5). The estimation from the hyperspectral measurements indicates that the main layer present at the surface is a mixture between steel and MnO, with respective volume fractions 0.37 and 0.63. A mixture of steel, Mn-Si-O and SiO_2 is also detected by the algorithm, as well as a layer of aluminum oxides. The presence of these oxides is confirmed by the XPS analysis. In comparison with samples 19 and 22, the additives used for sample 5 contain more manganese, which is well detected by the algorithm. However, the observations conducted on the SEM images seem to indicate that the proportion of steel at the surface is under-estimated by our analysis. We attribute this under-estimation to the fact that there is a high correlation between the characteristic spectra observed for MnO layers and for the substrate, which leads the estimation algorithm to significantly over-estimate the presence of MnO oxides to the detriment of steel.

	Mixture 1		Mixture 2			Mixture 3		
	Sub.	MnO	Sub.	Mn-Si-O	SiO2	Sub.	Mn-Al-O	Al2O3
v_f	0.37	0.63	0.13	0.40	0.47	0.14	0.43	0.43
α	0.44		0.32			0.24		
h	182 nm		183 nm			87 nm		

Table 7.3: Estimation results obtained for sample 5.

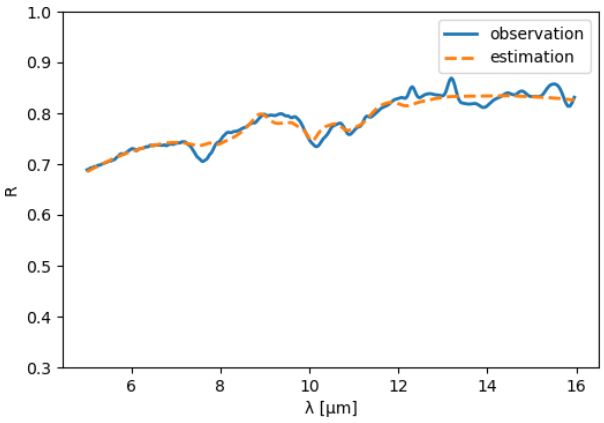

Figure 7.6: Recovery of the observations for sample 5 [Zenati et al., 2023].

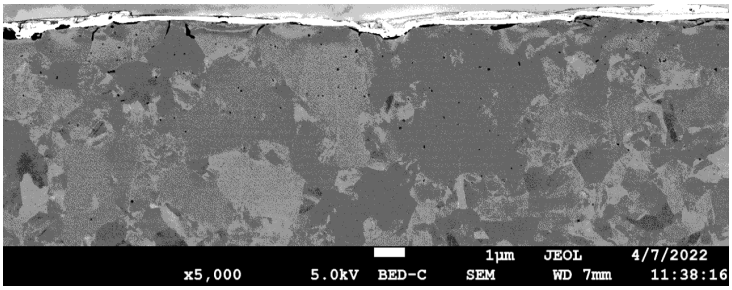

Figure 7.7: Scanning electron microscopy image of a slice of sample 5 [Zenati et al., 2023].

Overall, we note that there is a good agreement between the estimations obtained with the proposed approach and the observations that are conducted on the experimental samples. We note in particular that the algorithm succeeds in detecting the presence of a covering oxide layer for sample 19, which is the only one to be completely covered with oxides on the surface. The estimated thickness is also consistent with the observations made on the sample. We can also note that the estimates made on sample 5, and more generally for all samples associated with the DP980 thermal cycle show a relative preponderance of manganese oxides compared to silicon oxides. This is again in agreement with the experimental observations conducted on the samples. Finally, it should be noted that the estimated reflectance curves are close to the curves measured experimentally for all the studied samples. The algorithm has nevertheless important limitations due to the strong correlations that can occur between some of the reflectance spectra present in the dictionary. The strong correlation that exists at the wavelengths considered between the reflectance of the manganese oxides and that of the steel substrate explains for instance the over-estimation of the proportion of manganese oxides MnO for the samples with the DP980 thermal cycle.

7.4 Conclusions and perspectives

In this chapter, we presented a novel methodology that utilizes hyperspectral measurements to detect selective oxide surface layers on steel sheets. Our approach involves pre-computing optical reflectance associated with heterogeneous oxide layers of distinct thickness and chemical composition to create a reference spectrum dictionary. This enables us to formulate the highly nonlinear unmixing problem as a sparse regression problem, which can be solved using various approaches, including the path and group Lasso algorithms. The group Lasso algorithm is particularly suitable for the dictionary's intrinsic structure, which contains groups of heterogeneous oxides with the same elementary components but different thicknesses and stochiometries. Our numerical experiments demonstrate that our approach can accurately estimate the presence of selective oxides with high recall and precision. However, we also identified a major limitation of our method: when spectra in the dictionary are strongly correlated, different oxide combinations can potentially produce similar reflectance spectra. To address this issue, incorporating spatial information to enforce consistency between estimates at adjacent locations could be helpful. Total variation based algorithms have been proposed for this purpose in the literature [Chouzenoux et al., 2013; Iordache et al., 2012;

Moussaoui et al., 2012], which we will explore in future research. Finally, we applied our estimation method to experimental samples previously characterized using XPS analysis and scanning electron microscopy. The results were qualitatively consistent with the XPS and SEM observations, providing further evidence for the effectiveness of our approach.

7.5 Related work on hyperspectral imaging

In addition to Tarek Zenati's PhD thesis work, I had the opportunity to contribute to the field of hyperspectral unmixing within the European project M3S. This project aimed to develop a multimodal microscopy platform to improve the diagnosis and prognosis of chronic lymphocytic leukemia. As part of the platform, Raman spectroscopy imaging was incorporated.

In this study, our objective was to design a Bayesian algorithm for extracting endmembers and estimating their abundances in situations where the observations were highly mixed but prior information on the abundances was available. Bayesian approaches provide a robust mathematical framework for incorporating prior knowledge into problem descriptions. We extended the Bayesian algorithm initially described in [Dobigeon et al., 2009] to handle scenarios where informative priors were readily accessible for the endmember abundances.

Specifically, we introduced bound parameters that allowed us to incorporate qualitative prior information on the abundances, which could be estimated using a simulated annealing algorithm. The development of this method was primarily motivated by the challenge of unmixing Raman spectra obtained from B-lymphocyte cells. Due to the limited spatial resolution of the acquisition, the Raman spectra often contained contributions from the nucleus, cytoplasm, and cell membrane. However, it was known that the contributions of the membrane and cytoplasm to the measured spectra were limited, with abundances typically ranging around 0.2 or 0.3 at most. Therefore, their respective spectra were never observed as distinct entities within the available hyperspectral observations. This research study resulted in a conference publication [Figliuzzi et al., 2016b].

7.6 Related publications

- T. Zenati, B. Figliuzzi and S. H. Ham. Surface oxides characterization based on hyperspectral observations . *Chemometrics and Intelligent Laboratory Systems*, in review, 2023.
- T. Zenati, B. Figliuzzi and S. H. Ham. Characterization of surface oxides from hyperspectral measurements. In *Workshop on Hyperspectral Image and Signal Processing : Evolution in Remote Sensing (WHISPERS)* (pp. 1-5). IEEE. 2022.
- T. Zenati, B. Figliuzzi and S. H. Ham. A supervised approach for the detection of surface oxides from hyperspectral measurements. In *Workshop on Hyperspectral Image and Signal Processing : Evolution in Remote Sensing (WHISPERS)* (pp. 1-5). 2022.
- T. Zenati, B. Figliuzzi and S. H. Ham. Surface Oxide Detection and Characterization Using Sparse Unmixing on Hyperspectral Images. In *International Con-*

ference on Image Analysis and Processing (pp. 291-302). Springer. 2022.

- B. Figliuzzi, S. Velasco-Forero, M. Bilodeau, J. Angulo. A bayesian approach to linear unmixing in the presence of highly mixed spectra. In *International Conference on Advanced Concepts for Intelligent Vision Systems* (pp. 263-274). Springer, 2016.

Chapter 8

Conclusion and perspectives

I have presented in this manuscript research work that I have carried out since 2013 on image processing and its applications, in particular in the field of materials science. To conclude the manuscript, I would like to provide some perspectives on the different research topics addressed in this work.

Morphological models The first research axis detailed in this manuscript concerns the development of morphological models allowing to reproduce the microstructure of heterogeneous materials. My research perspectives on the subject are essentially applicative. In the short term, as part of several research projects carried out in collaboration with an industrial partner, we are seeking to develop a morphological model making it possible to reproduce the geometric distribution of melanosomes in the different layers of the epidermis. Coupled with optical simulation tools, this model aims to study the influence of a certain number of biological parameters of the skin, such as the density of melanosomes or their distribution in the form of clusters, on the optical properties of the skin. We also intend to adapt the MCMC approach presented in chapter 3 in order to determine, for given optical properties, which biological parameters of the skin allow for the obtaining of these properties.

Image segmentation My main research focus in the field of image segmentation is to advance the development of supervised approaches that utilize synthesized datasets for segmenting sets of similar images obtained during physics experiments. These images are often available in large quantities, but annotating them is a time-consuming task. Consequently, using supervised algorithms for image segmentation is generally not feasible. As discussed previously, one potential solution to address this issue is to jointly generate synthesized images that closely resemble the experimental ones and ground truth annotations. This approach is challenging in general, but it could be viable for the specific types of images that we aim to synthesize, such as suspension or microstructure images. These images are indeed characterized by relatively simple geometries. There is a substantial body of literature accurately describing material microstructures using morphological models. Hence, leveraging morphological models could enable the generation of synthesized training datasets, allowing us to train supervised segmentation algorithms for various practical problems, as demonstrated in the research work presented in the manuscript. It remains however to be determined to

what extend the approach presented in Chap. 6 can be utilized for other types of image datasets.

In the short term, I am exploring this approach in the field of medical imaging as part of Pablo Jeken Rico's thesis, which I co-supervise in collaboration with Elie Hachem (Centre de mise en forme des matériaux, Mines Paris, PSL University). The primary objective of this research project is to develop a segmentation model for three-dimensional MRI images of aneurysms located at the base of the brain. The limited availability of annotated images poses a significant challenge for developing the model. To overcome this challenge, we plan to investigate the generation of synthetic MRI images using morphological models, similar to the approach presented in Chap. 6 for segmenting images from rheology experiments. By artificially increasing the size of the available training data through synthetic image generation, we hope to be able to subsequently apply advanced supervised segmentation techniques to these images.

Looking ahead, in the long term, I aim to explore the combined use of morphological models and generative models, such as generative adversarial networks (GANs) [Goodfellow et al., 2014], to facilitate the synthesis of training images. For example, it may be of interest to employ a GAN on top of a morphological model to generate realistic noise and blur in the synthesized images.

Hyperspectral imaging The last axis of research described in this manuscript revolves around hyperspectral imaging. As emphasized previously, hyperspectral imaging has traditionally been linked to remote sensing applications, but its usage is expanding into new domains, such as surface characterization and chemical imaging. The optical phenomena that lead to the hyperspectral observations in these applications differ significantly from those encountered in traditional remote sensing, necessitating the development of specialized signal processing methods. Interestingly, these developments intersect current research topics that remain largely open and generate significant interest in the hyperspectral unmixing literature: highly mixed observations, nonlinear mixing models and endmembers variability. To address these issues, I would like to explore two distinct approaches.

1. *Optimal transport based hyperspectral unmixing* A classical approach to hyperspectral unmixing is to formulate the problem within the framework of non-negative matrix factorization (NMF). In that case, the problem boils down to determining a factorization of the hyperspectral observation matrix into two non-negative matrices corresponding to the endmembers and to their respective abundances, respectively. The main issue with NMF is that the factorization is not unique, which triggers the need for some kind of regularization. A common approach to perform the regularization is to exploit the geometry of the observations, which are known to lie in a low dimensional simplex in the observations space. Here, I would like to explore an alternative NMF-based formulation of blind hyperspectral unmixing, where the regularization term is the optimal transport distance between some discrete probability distribution directly related to the abundances and an arbitrary probability distribution [Genevay et al., 2016, 2018; Peyré et al., 2019]. By enforcing the algorithm to identify endmembers whose associated abundances follow some pre-specified distribution, this kind of regularization could be useful to deal with highly mixed observations.

2. *Variational Autoencoders* Let us assume that we dispose of a set $(\vec{x}_i)_{1\leq i\leq n}$ of n hyperspectral observations in $\mathbb{R}^d$ resulting from the mixing of k distinct endmembers characterized by some degree of variability. The observations $(\vec{x}_i)_{1\leq i\leq n}$ can be interpreted as random realizations of a conditional probability distribution $p(\vec{x}_i|S_i,\vec{a}_i)$ describing the mixing model, where $S_i = (\vec{s}_1^i,\ldots,\vec{s}_k^i) \in \mathbb{R}^{d\times k}$ and $\vec{a}_i = (a_i^1,\ldots,a_i^k)^T \in \mathbb{R}^k$ are the endmembers matrix and the abundances vector that led to the i-th hyperspectral observation, respectively. The quantities S_i and $\vec{a}_i$ are latent variables of the model and have to be estimated for each observation $\vec{x}_i$. To account for endmembers variability, we can define a prior distribution on the endmembers matrices $(S_i)_{1\leq i\leq n}$ by considering for instance that for $i = 1,\ldots,n$, the endmembers $\vec{s}_k^i \in \mathbb{R}^d$ are realizations of a multivariate Gaussian distribution with mean $\vec{\mu}_k$ and specified covariance matrix. The task of hyperspectral unmixing can be summarized as the estimation of 1) the parameters $(\vec{\mu}_1,\ldots\vec{\mu}_k) \in \mathbb{R}^{d\times k}$ characterizing the prior endmembers distribution and 2) of the latent variables $(\vec{s}_1^i,\ldots\vec{s}_k^i) \in \mathbb{R}^{d\times k}$ and $\vec{a}_i \in \mathbb{R}^k$ associated with the observation $(\vec{x}_i)$, for all $i = 1,\ldots,n$.

   My objective is to explore the use of a variational autoencoder (VAE) architectures to perform these estimations [Kingma and Welling, 2013; Kingma et al., 2019]. VAEs are constituted of two independent parameterized models: an *encoder* or *recognition model*, and a *decoder* or *generative model*. Here, we will use the generative model leading to the observations as decoder. This model is parameterized by the quantities $(\vec{\mu}_1,\ldots,\vec{\mu}_k)$. For the encoder g, we will consider a probability distribution parameterized by the output of a neural network f with a set of parameters ϕ, which will be trained to approximate the conditional distribution of the latent variables $(S_i)_{1\leq i\leq n}$ and $(\vec{a}_i)_{1\leq i\leq n}$ knowing the observations. The interest of this approach is 1) that it incorporates the mixing model directly into the decoder, which enables to consider nonlinear models and 2) that it models the endmembers by statistical distributions, which allows to account for endmembers variability [Zare and Ho, 2013].

Bibliography

Bibliography

R. Achanta and S. Susstrunk. Superpixels and Polygons Using Simple Non-iterative Clustering. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4895–4904, Honolulu, HI, July 2017. IEEE. ISBN 978-1-5386-0457-1.

R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence*, 34(11):2274–2282, 2012.

R. Alais, P. Dokládal, A. Erginay, B. Figliuzzi, and E. Decencière. Fast macula detection and application to retinal image quality assessment. *Biomedical Signal Processing and Control*, 55:101567, 2020.

H. Altendorf and D. Jeulin. Random-walk-based stochastic modeling of three-dimensional fiber systems. *Physical Review E*, 83(4):041804, 2011.

O. Amsellem, K. Madi, F. Borit, D. Jeulin, V. Guipont, M. Jeandin, E. Boller, and F. Pauchet. Two-dimensional (2d) and three-dimensional (3d) analyses of plasma-sprayed alumina microstructures for finite-element simulation of young's modulus. *Journal of Materials Science*, 43(12):4091–4098, 2008.

C. Andrieu, N. De Freitas, A. Doucet, and M. I. Jordan. An introduction to mcmc for machine learning. *Machine learning*, 50:5–43, 2003.

P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):898–916, May 2011.

M. Avrami. Kinetics of phase change. I. general theory. *The Journal of Chemical Physics*, 7(12):1103–1112, 1939.

M. Baradad, J. Wulff, T. Wang, P. Isola, and A. Torralba. Learning to see by looking at noise. *Advances in Neural Information Processing Systems*, 34, 2021.

T. Barisin, C. Jung, F. Müsebeck, C. Redenbach, and K. Schladitz. Methods for segmenting cracks in 3d images of concrete: A comparison based on semi-synthetic images. *Pattern Recognition*, 129:108747, 2022.

S. Beauvais, V. Guipont, M. Jeandin, D. Jeulin, A. Robisson, and R. Saenger. Study of the porosity in plasma-sprayed alumina through an innovative three-dimensional

simulation of the coating buildup. *Metallurgical and Materials Transactions A*, 39 (11):2711–2724, 2008.

V. Besnier, H. Jain, A. Bursuc, M. Cord, and P. Pérez. This dataset does not exist: training models from generated images. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2020.

S. Beucher and C. Lantuéjoul. Use of watersheds in contour detection. In *International workshop on image processing, real-time edge and motion detection*, 1979.

S. Beucher and F. Meyer. The morphological approach to segmentation: the watershed transformation. *Optical Engineering-New York-Marcel Dekker Incorporated-*, 34:433–433, 1992.

J. M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, and J. Chanussot. Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE journal of selected topics in applied earth observations and remote sensing*, 5(2):354–379, 2012.

C. M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.

F. Blanc, E. Lemaire, A. Meunier, and F. Peters. Microstructure in sheared non-brownian concentrated suspensions. *Journal of rheology*, 57(1):273–292, 2013.

V. Bortolussi, B. Figliuzzi, F. Willot, M. Faessel, and M. Jeandin. Morphological modeling of cold spray coatings. *Image Analysis & Stereology*, 37(2):145–158, 2018.

V. Bortolussi, B. Figliuzzi, F. Willot, M. Faessel, and M. Jeandin. Electrical conductivity of metal–polymer cold spray composite coatings onto carbon fiber-reinforced polymer. *Journal of Thermal Spray Technology*, 29:642–656, 2020.

S. Bourguignon, J. Ninin, H. Carfantan, and M. Mongeau. Exact sparse approximation problems via mixed-integer programming: Formulations and computational performance. *IEEE Transactions on Signal Processing*, 64(6):1405–1419, 2015.

S. Boyd, S. P. Boyd, and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

J.-F. Brun, L. Del Campo, D. De Sousa Meneses, and P. Echegut. Infrared optical properties of α-alumina with the approach to melting: γ-like tetrahedral structure and small polaron conduction. *Journal of Applied Physics*, 114(22):223501, 2013.

S. Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.

P. J. Burt and E. H. Adelson. The laplacian pyramid as a compact image code. In *Readings in computer vision*, pages 671–679. Elsevier, 1987.

P. Buyssens, I. Gardin, and S. Ruan. Eikonal based region growing for superpixels generation: Application to semi-supervised real time organ segmentation in CT images. *IRBM*, 35(1):20–26, Dec. 2014a.

P. Buyssens, I. Gardin, S. Ruan, and A. Elmoataz. Eikonal-based region growing for efficient clustering. *Image and Vision Computing*, 32(12):1045–1054, Dec. 2014b. ISSN 0262-8856.

P. Buyssens, M. Toutain, A. Elmoataz, and O. Lézoray. Eikonal-based vertices growing and iterative seeding for efficient graph-based segmentation. In *IEEE International Conference on Image Processing (ICIP 2014)*, page 5 pp., Paris, France, Oct. 2014c.

V. Capasso. *Mathematical Modelling for Polymer Processing: Polymerization, Crystallization, Manufacturing*, volume 2. Springer, 2003.

P. Cettour-Janet, C. Cazorla, V. Machairas, Q. Delannoy, N. Bednarek, F. c. o. Rousseau, E. Decencière, and N. Passat. Watervoxels. *Image Processing On Line IPOL*, 9:317–328, 2019.

K. Chang. *Machine learning for image segmentation.* PhD thesis, Université Paris sciences et lettres, 2019.

K. Chang and B. Figliuzzi. Fast marching based superpixels generation. In *Mathematical Morphology and Its Applications to Signal and Image Processing: 14th International Symposium, ISMM 2019, Saarbrücken, Germany, July 8-10, 2019, Proceedings 14*, pages 350–361. Springer, 2019.

K. Chang and B. Figliuzzi. Fast marching based superpixels. *Mathematical Morphology-Theory and Applications*, 4(1):127–142, 2020.

Q. Chen, J. Xu, and V. Koltun. Fast image processing with fully-convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2497–2506, 2017.

S. N. Chiu, D. Stoyan, W. S. Kendall, and J. Mecke. *Stochastic geometry and its applications.* John Wiley & Sons, 2013.

G. Choquet. Theory of capacities. *Annales de l'institut Fourier*, 5(131–295):54, 1953.

E. Chouzenoux, S. Moussaoui, M. Legendre, and J. Idier. Algorithme primal-dual de points intérieurs pour l'estimation pénalisée des cartes d'abondances en imagerie hyperspectrale. *Traitement du signal*, 30(1-2):35–59, 2013.

D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 24(5): 603–619, 2002.

E. D'Ambrosio, F. Blanc, and E. Lemaire. Viscous resuspension of non-brownian particles: determination of the concentration profiles and particle normal stresses. *Journal of Fluid Mechanics*, 911, 2021.

E. Dejnozková and P. Dokládal. A parallel algorithm for solving the eikonal equation. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03).*, volume 3, pages III–325. IEEE, 2003.

F. Delloro, M. Faessel, H. Proudhon, D. Jeulin, M. Jeandin, E. Meillot, and L. Bianchi. A morphological approach to the modeling of the cold spray process. In *ITSC 2014*, pages 221–226. DVS, 2014.

J. A. Dijksman, N. Brodu, and R. P. Behringer. Refractive index matched scanning and detection of soft particles. *Review of Scientific Instruments*, 88(5):051807, 2017.

N. Dobigeon, S. Moussaoui, M. Coulon, J.-Y. Tourneret, and A. O. Hero. Joint bayesian endmember extraction and linear unmixing for hyperspectral imagery. *IEEE Transactions on Signal Processing*, 57(11):4355–4368, 2009.

E. Dougherty. *Mathematical morphology in image processing*. CRC press, 1992.

B. Efron, T. Hastie, I. Johnstone, R. Tibshirani, et al. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.

A. El Moumen, T. Kanit, A. Imad, and H. El Minor. Effect of overlapping inclusions on effective elastic properties of composites. *Mechanics Research Communications*, 53:24–30, 2013.

A. El Moumen, A. Imad, T. Kanit, E. Hilali, and H. El Minor. A multiscale approach and microstructure design of the elastic composite behavior reinforced with natural particles. *Composites Part B: Engineering*, 66:247–254, 2014.

M. Faessel. Vtksim software, 2016.

P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *International journal of computer vision*, 59(2):167–181, 2004.

M. Ferté. *Étude et analyse de couches minces par techniques multi-spectroscopiques pour une application sur une ligne de galvanisation*. PhD thesis, Université de Lorraine, 2014.

B. Figliuzzi. Eikonal-based models of random tessellations. *Image Analysis & Stereology*, 38(1):15–23, 2019.

B. Figliuzzi and C. Buie. Rise in optimized capillary channels. *Journal of Fluid Mechanics*, 731:142–161, 2013.

B. Figliuzzi, D. Jeulin, A. Lemaître, G. Fricout, J.-J. Piezanowski, and P. Manneville. Numerical simulation of thin paint film flow. *Journal of Mathematics in Industry*, 2 (1):1–20, 2012a.

B. Figliuzzi, D. Jeulin, A. Lemaitre, G. Fricout, J.-J. Piezanowski, and P. Manneville. Rheology of thin films from flow observations. *Experiments in fluids*, 53:1289–1299, 2012b.

B. Figliuzzi, W. H. R. Chan, J. Moran, and C. R. Buie. Nonlinear electrophoresis of ideally polarizable particles. *Physics of Fluids*, 26(10), 2014.

B. Figliuzzi, W. Chan, C. Buie, and J. Moran. Nonlinear electrophoresis in the presence of dielectric decrement. *Physical Review E*, 94(2):023115, 2016a.

B. Figliuzzi, S. Velasco-Forero, M. Bilodeau, and J. Angulo. A bayesian approach to linear unmixing in the presence of highly mixed spectra. In *International Conference on Advanced Concepts for Intelligent Vision Systems*, pages 263–274. Springer, 2016b.

B. Figliuzzi, K. Chang, and M. Faessel. Hierarchical segmentation based upon multi-resolution approximations and the watershed transform. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, pages 185–195. Springer, 2017.

B. Figliuzzi, A. Montaux-Lambert, F. Willot, G. Naudin, P. Dupuis, B. Querleux, and E. Huguet. A bayesian approach to morphological models characterization. *Image Analysis & Stereology*, 40(3):171–180, 2021.

B. Fulkerson, A. Vedaldi, and S. Soatto. Class segmentation and object localization with superpixel neighborhoods. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 670–677. IEEE, 2009.

J.-B. Gasnier, F. Willot, H. Trumel, B. Figliuzzi, D. Jeulin, and M. Biessy. A fourier-based numerical homogenization tool for an explosive material. *Matériaux & Techniques*, 103(3):308, 2015.

A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. *Advances in neural information processing systems*, 29, 2016.

A. Genevay, G. Peyré, and M. Cuturi. Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. PMLR, 2018.

E. N. Gilbert. Random subdivisions of space into crystals. *The Annals of mathematical statistics*, pages 958–972, 1962.

R. Giraud, V.-T. Ta, N. Papadakis, and Y. Berthoumieu. Texture-aware superpixel segmentation. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 1465–1469. IEEE, 2019.

R. C. Gonzalez and R. E. Woods. *Digital Image Processing (3rd Edition)*. Pearson, 2007.

I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

S. H. Ham, M. Ferie, C. Carteret, G. Fricout, J. Angulo, and F. Capon. Hyperspectral imaging as an analytical tool for thin single and multilayer oxides characterization: A laboratory study. In *2016 8th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, pages 1–5. IEEE, 2016.

R. Heylen, M. Parente, and P. Gader. A review of nonlinear hyperspectral unmixing methods. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(6):1844–1868, 2014.

S. Hui Ham. *Propriétés physico-chimiques des oxydes de surface et analyse des données de l'imagerie hyperspectrale. Physicochemical properties of surface oxides and data analysis of hyperspectral imaging.* PhD thesis, Université de Lorraine, 2018.

S. Hui Ham, M. Ferté, G. Fricout, L. Depalo, and C. Carteret. In-situ spectral emissivity measurement of alloy steels during annealing in controlled atmosphere. *QIRT 2016*, pages 315–321, 2016.

M.-D. Iordache, J. M. Bioucas-Dias, and A. Plaza. Sparse unmixing of hyperspectral data. *IEEE Transactions on Geoscience and Remote Sensing*, 49(6):2014–2039, 2011.

M.-D. Iordache, J. M. Bioucas-Dias, and A. Plaza. Total variation spatial regularization for sparse hyperspectral unmixing. *IEEE Transactions on Geoscience and Remote Sensing*, 50(11):4484–4502, 2012.

A. Jahanian, X. Puig, Y. Tian, and P. Isola. Generative models as a data source for multiview representation learning. *arXiv preprint arXiv:2106.05258*, 2021.

A. K. Jain and F. Farrokhnia. Unsupervised texture segmentation using gabor filters. *Pattern recognition*, 24(12):1167–1186, 1991.

A. Jean, D. Jeulin, S. Forest, S. Cantournet, and F. N'Guyen. A multiscale microstructure model of carbon black distribution in rubber. *Journal of microscopy*, 241(3): 243–260, 2011.

D. Jeulin. Morphology and effective properties of multi-scale random sets: A review. *Comptes Rendus Mécanique*, 340(4):219–229, 2012.

D. Jeulin. *Morphological models of random structures.* Springer, 2021.

D. Jeulin and M. Moreaud. Percolation of multi-scale fiber aggregates. In *S4G, 6th International Conference Stereology, Spatial Statistics and Stochastic Geometry. Prague*, pages 26–29, 2006.

W. A. Johnson and R. F. Mehl. Reaction kinetics in processes of nucleation and growth. *Trans. Aime*, 135(8):396–415, 1939.

Y. S. Joung, B. M. Figliuzzi, and C. R. Buie. Design of capillary flows with functionally graded porous titanium oxide films fabricated by anodization instability. *Journal of colloid and interface science*, 423:143–150, 2014.

C. Kimme, D. Ballard, and J. Sklansky. Finding circles by an array of accumulators. *Communications of the ACM*, 1975.

D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

D. P. Kingma, M. Welling, et al. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019.

M. Koishi, N. Kowatari, B. Figliuzzi, M. Faessel, F. Willot, and D. Jeulin. Computational material design of filled rubbers using multi-objective design exploration. In *Constitutive Models for Rubber X*, pages 467–473. CRC Press, 2017.

A. Levinshtein, A. Stere, K. N. Kutulakos, D. J. Fleet, S. J. Dickinson, and K. Siddiqi. Turbopixels: Fast superpixels using geometric flows. *IEEE transactions on pattern analysis and machine intelligence*, 31(12):2290–2297, 2009.

Z. Li and J. Chen. Superpixel segmentation using Linear Spectral Clustering. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1356–1363, June 2015.

A. Liebscher and C. Redenbach. Statistical analysis of the local strut thickness of open cell foams. *Image Analysis & Stereology*, 32(1):1–12, 2013.

V. Machairas, M. Faessel, D. Cárdenas-Peña, T. Chabardes, T. Walter, and E. Decencière. Waterpixels. *IEEE Transactions on Image Processing*, 24(11):3707–3716, 2015.

J. Malik, S. Belongie, T. Leung, and J. Shi. Contour and texture analysis for image segmentation. *International journal of computer vision*, 43(1):7–27, 2001.

R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE transactions on pattern analysis and machine intelligence*, 17(2):158–175, 1995.

S. Mallat. *A wavelet tour of signal processing*. Elsevier, 1999.

S. G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on signal processing*, 41(12):3397–3415, 1993.

D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001.

J. Masci, U. Meier, D. Ciresan, J. Schmidhuber, and G. Fricout. Steel defect classification with max-pooling convolutional neural networks. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6. IEEE, 2012.

J. Masci, U. Meier, G. Fricout, and J. Schmidhuber. Multi-scale pyramidal pooling network for generic steel defect classification. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2013.

G. Matheron. *Random sets and integral geometry*, volume 1. Wiley New York, 1975.

D. D. S. Meneses, M. Eckes, L. Del Campo, and P. Echegut. Phase transformations of crystalline SiO_2 versus dynamic disorder between room temperature and liquid state. *Journal of Physics: Condensed Matter*, 26(25):255402, 2014.

F. Meyer and S. Beucher. Morphological segmentation. *Journal of visual communication and image representation*, 1(1):21–46, 1990.

R. E. Miles. The random division of space. *Advances in applied probability*, pages 243–266, 1972.

J. Møller. Random tessellations in $\mathbb{R}^d$. *Advances in Applied Probability*, 21(1):37–73, 1989.

J. Møller. Random johnson-mehl tessellations. *Advances in applied probability*, 24(4): 814–844, 1992.

J. Møller. *Lectures on random Voronoi tessellations*. Springer, 1994.

M. Moreaud, D. Jeulin, V. Morard, and R. Revel. Tem image analysis and modelling: application to boehmite nanoparticles. *Journal of microscopy*, 245(2):186–199, 2012.

S. Moussaoui, J. Idier, and E. Chouzenoux. Primal dual interior point optimization for penalized least squares estimation of abundance maps in hyperspectral imaging. In *2012 4th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, pages 1–4. IEEE, 2012.

M. Nagy, L. Istrate, M. Simtinică, S. Travadel, and P. Blanc. Automatic detection of marine litter: A general framework to leverage synthetic data. *Remote Sensing*, 14 (23):6102, 2022.

J. A. Nelder and R. Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.

J. Ohser and K. Schladitz. *3D images of materials structures: processing and analysis*. John Wiley & Sons, 2009.

N. Otsu. A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics*, 9(1):62–66, 1979.

D. Paulovics, B. Figliuzzi, T. Dumont, and F. Blanc. A supervised algorithm entirely trained on a synthetic dataset to segment granular suspension images, 2023.

G. Peyré, M. Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

C. Peyrega, D. Jeulin, C. Delisée, and J. Malvestio. 3d morphological modelling of a random fibrous network. *Image Analysis & Stereology*, 28(3):129–141, 2011.

S. Ravuri and O. Vinyals. Classification accuracy score for conditional generative models. *Advances in neural information processing systems*, 32, 2019.

C. Redenbach. Microstructure models for cellular materials. *Computational Materials Science*, 44(4):1397–1407, 2009.

C. Redenbach and I. Vecchio. Statistical analysis and stochastic modelling of fibre composites. *Composites Science and Technology*, 71(2):107–112, 2011.

C. P. Robert, G. Casella, and G. Casella. *Monte Carlo statistical methods*, volume 2. Springer, 1999.

A. Schick, M. Fischer, and R. Stiefelhagen. Measuring and evaluating the compactness of superpixels. In *Proceedings of the 21st international conference on pattern recognition (ICPR2012)*, pages 930–934. IEEE, 2012.

R. Schneider and W. Weil. *Stochastic and integral geometry*. Springer, 2008.

J. Serra. *Image analysis and mathematical morphology*. Academic press, 1982.

J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.

J. A. Sethian. Fast marching methods. *SIAM review*, 41(2):199–235, 1999.

J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000.

C. Shorten and T. M. Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019.

B. Snook, J. E. Butler, and É. Guazzelli. Dynamics of shear-induced migration of spherical particles in oscillatory pipe flow. *Journal of Fluid Mechanics*, 786:128–153, 2016.

P. Soille. *Morphological image analysis: principles and applications*. Springer-Verlag New York, Inc., 2003.

D. Stutz, A. Hermans, and B. Leibe. Superpixels: An evaluation of the state-of-the-art. *Computer Vision and Image Understanding*, 166:1–27, 2018.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

V. P. Tolstoy, I. Chernyshova, and V. A. Skryshevsky. *Handbook of infrared spectroscopy of ultrathin films*. John Wiley & Sons, 2003.

S. Torquato. *Random heterogeneous materials: microstructure and macroscopic properties*, volume 16. Springer, 2002.

J. E. Vargas-Muñoz, A. S. Chowdhury, E. B. Alexandre, F. L. Galvão, P. A. V. Miranda, and A. X. Falcão. An iterative spanning forest framework for superpixel segmentation. *IEEE Transactions on Image Processing*, 28(7):3477–3489, 2019.

L. Vincent and P. Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 13(06):583–598, 1991.

H. Wang, A. Pietrasanta, D. Jeulin, F. Willot, M. Faessel, L. Sorbier, and M. Moreaud. Modelling mesoporous alumina microstructure with 3d random models of platelets. *Journal of microscopy*, 260(3):287–301, 2015.

H. Wang, F. Willot, M. Moreaud, M. Rivallan, L. Sorbier, and D. Jeulin. Numerical simulation of hindered diffusion in γ-alumina catalyst supports. *Oil & Gas Science and Technology–Revue d'IFP Energies nouvelles*, 72(2):8, 2017.

H. Wang, D. Jeulin, F. Willot, L. Sorbier, and M. Moreaud. Modelling of the microstructure of mesoporous alumina constrained by morphological simulation of nitrogen porosimetry. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 553:378–396, 2018.

F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. In *International Conference on Learning Representations (ICLR)*, May 2016.

M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68 (1):49–67, 2006.

A. Zare and K. Ho. Endmember variability in hyperspectral analysis: Addressing spectral variability during spectral unmixing. *IEEE Signal Processing Magazine*, 31(1): 95–104, 2013.

T. Zenati, B. Figliuzzi, and S. Ham. A supervised approach for the detection of surface oxides from hyperspectral measurements. In *2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, pages 1–5. IEEE, 2022a.

T. Zenati, B. Figliuzzi, and S. Hui Ham. Surface oxide detection and characterization using sparse unmixing on hyperspectral images (in press). In *1st Int. Workshop on intelligent Systems in human and artificial perception (ISHAPE)*, 2022b.

T. Zenati, B. Figliuzzi, and S. Hui Ham. A supervised approach for the detection of surface oxides from hyperspectral measurements. In *2022 12th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*. IEEE, 2022c.

T. Zenati, F. B, and S. Ham. Endmember variability in hyperspectral analysis: Addressing spectral variability during spectral unmixing. *Chemometrics and Intelligent Laboratory Systems (in review)*, 2023.

C. L. Zitnick and S. B. Kang. Stereo for image-based rendering using image oversegmentation. *International Journal of Computer Vision*, 75(1):49–65, 2007.

Appendix A

Curriculum Vitae

A.1 Research experience

2016 - Now: Mines Paris, Centre for Mathematical Morphology
Maître assistant
Research and teaching activities in image processing and mathematical morphology

2014 - 2016: Mines Paris, Centre for Mathematical Morphology
Assistant Professor in Tenure track
Research and teaching activities in image processing and mathematical morphology

2012 - 2013: Massachusetts Institute of Technology (Cambridge, USA)
Shapiro Postdoctoral Fellow
Postdoctoral researcher in the research group of Professor Cullen R. Buie
Research and teaching activities in fluid mechanics

2008 - 2011: ArcelorMittal (Maizières-lès-Metz, France)
Research engineer
Study of the flow of a thin fluid film on a rough surface (CIFRE PhD Thesis)

2008: AREVA NP (La Défense, France)
Internship
Study of the MCMC calculation scheme of a nuclear reactor

2007: TRANSVER GmbH (Munich, Germany)
Internship
Development of a probabilistic traffic prediction algorithm in a road network

A.2 Education

2008 - 2011: Mines Paris - PhD Thesis
CIFRE PhD thesis in Mathematical Morphology in partnership with ArcelorMittal Research under the supervision of Prof. Dominique Jeulin
Title: Modeling of the roughness and the appearance of painted steel surfaces

2005 - 2008: Mines Paris - Civil engineer
Specialization in nuclear engineering

A.3 Main collaborative research projects

A.3.1 Direct collaborations with industrials

2021 - 2022: L'Oréal R&I Development of a morphological model for studying the influence of the repartition of melanosomes in the epidermis on the optical properties of the skin.

2019 - 2020: Fareco Fayat Development of a people counting algorithm in a vehicle for carpooling certification. Project carried out in collaboration with Nicolas Petit (Center for Automatic Control and Systems, Mines Paris).

2016 - 2019: L'Oréal R&I Development of a morphological model to describe the microstructure of a colloidal medium in order to enhance its optical properties.

2013 - 2016: Yokohama Rubber Co. Development of a random morphological model aimed at simulating the microstructure of a pneumatic material in order to improve its mechanical properties.

A.3.2 Collaborative research

2021 - Now: Participation in the Cure Project Development of an aneurysm segmentation algorithm as part of the ERC Consolidator Cure grant awarded to Prof. Elie Hachem (Center for Materials Forming, Mines Paris)

2015 - 2016: Participation in the European Project M3S Development of

a multimodal microscopy platform to improve the diagnosis and prognosis of chronic lymphocytic leukemia

2015 - 2016: Co-responsible of the Coldspray Project Development of a morphological model to describe the microstructure of a substrate deposited by coldspray process. Collaboration with Vincent Bortolussi, Michel Jeandin (Center for Materials, Mines Paris), and François Willot (CMM, Mines Paris). Funding from the Air Force Office for Scientific Research (AFOSR)

2015 - 2016: Responsible of the Electrophoresis Project Development of a process for depositing gradient surfaces using non-linear electrophoresis. Collaboration with Alain Thorel and Anthony Chesnaud (Center for Materials, Mines Paris), funded through the Carnot M.I.N.E.S institute's call for additive manufacturing projects.

A.4 Students supervision

A.4.1 Post-doctoral supervision

2017 - 2018: Jihane Belhadj Co-advisor *Bayesian seismic tomography based on compressed sensing techniques.*

A.4.2 PhD students supervision

2022 - Now: Hana Yahia *Detection of the behavior of drilling tools by machine learning: Application to detection of events and optimization of performances* CIFRE PhD in partnership with the company Drillscan. Co-advisor with Thomas Romary, Laurent Gerbaud and Florent Di Meglio (Mines Paris).

2021 - Now: Pablo Jeken Rico *Deep reinforcement learning for the control of intracranial aneurysms and the optimization of stents.* PSL Cofund fellowship.
Co-advisor with Elie Hachem and Aurélien Larcher (Mines Paris).

2020 - 2023: Tarek Zenati *Characterization of the physico-chemical properties of steel surfaces by hyperspectral imaging.* CIFRE PhD in partnership with the company ArcelorMittal Research.
Co-advisor with Jesus Angulo (Mines Paris)

2016 - 2019: Kaïwen Chang *Machine learning for image segmentation*
Co-advisor with Jesus Angulo (Mines Paris)

2020 - 2023: Robin Alais *Image segmentation Problems and Contributions to*

Mathematical Morphology.
Co-advisor with Etienne Decencière and Petr Dokladal (Mines Paris)

A.4.3 Internships supervision

2020: David Paulovics, Nice University *Image processing with machine learning algorithms*
Co-advisor with Frédéric Blanc (Nice Physics Institute)

2016: Xiaojun Pan, Huazhong University of Science & Technology, China *Numerical simulation of thin film flow on horizontal substrate*

2015: Jian Chen, Huazhong University of Science & Technology, China *Research on the electrophoretic deposition of Al_2O_3 particles with AC and DC fields*

2013: Wai Hong Ronald Chan, Massachusetts Institute of Technology, USA) *Non linear electrophoresis of ideally polarizable particles*

A.5 Professional service

- Secretary of the International Society for Stereology and Image Analysis (ISSIA)
- Organizer of the annual seminar Journée ISS France (2016 - Now)
- Reviewer for the journals Physics of fluids, Journal of fluid mechanics, Pattern recognition Letters, Image Analysis & Stereology
- Examiner for the PhD thesis *"A novel way to measure the rheology of liquid"* of Abdulrahman Jasim Al-Behadili, University of Canterbury, NZ (2018).
- Examiner for the PhD thesis *"Physicochemical properties of surface oxides and data analysis of hyperspectral imaging."* of Grace Shu Hui Ham, Université de Lorraine, France (2018).

Appendix B

Teaching

B.1 Civil engineering program, Mines Paris, France

2020 - 2023: Signal processing 1st year course of the civil engineering program at Mines Paris. This course is an introduction to Fourier analysis and its application to signal processing. I am currently co-responsible of the course with Florent Di Meglio (Mines Paris).

2022 - 2023: Data science 1st year course of the civil engineering program at Mines Paris. This course is an introduction to statistics and machine learning. I was responsible of this course in 2022 and 2023 as a replacement for Chloé-Agathe Azencott (Mines Paris).

2019 - 2023: Option Ingénierie digitale des systémes complexes: Image processing 3rd year course of the civil engineering program at Mines Paris. This course is part of the specialization *ingénierie digitale des systémes complexes* offered to final-year students. It takes the form of a week-long course focused on image processing with applications in materials engineering. I am currently co-responsible of this course with Henry Proudhon (Mines Paris).

2021 - 2023: Option Mareva 3rd year course of the civil engineering program at Mines Paris. This course is part of the specialization *Mareva* offered to final-year students. I am giving a lecture during the course on articial and convolutional neural networks and on object detection. The course responsible is Beatriz Marcotegui (Mines Paris).

2016 - 2019: Integral calculus 1st year course of the civil engineering program at Mines Paris. I was participating to the teaching of the practical sessions. The course responsible was Lionel Rosier (Mines Paris).

2014 - 2017: Thermal-fluid engineering 2nd year course of the civil engineering program at Mines Paris. I was participating to the teaching of the

practical sessions. The course responsibles were Dominique Marchio and François Cauneau (Mines Paris).

2013 - 2024: Physics of random media 2nd year course of the civil engineering program at Mines Paris. I am teaching the aspects of the course related to mathematical morphology and morphological models. The course responsible is François Willot (Mines Paris).

B.2 Masters program, Mines Paris, France

2021 - 2023: Computer vision, DMS Master Course from the *Master Design des Matériaux et Surfaces* (DMS). I am giving a lecture on convolutional networks for objects detection. The course responsible is Henry Proudhon (Mines Paris).

2017 - 2022: Heterogeneous materials, DMS Master Course from the *Master Design des Matériaux et Surfaces* (DMS). I was teaching the aspects of the course related to mathematical morphology and morphological models. The course responsible is Sébastien Joannes (Mines Paris).

B.3 Undergraduate courses, Massachusetts Institute of Technology, Cambridge, USA

2012 - 2013: Thermal-fluid Engineering II 4th year course of the undergraduate program of the Mechanical Engineering Department at Massachusetts Institute of Technology. I was teaching some lectures and the practical sessions of the course. The course responsibles were Cullen Buie (Spring 2012, Spring 2013) Pierre Lermusiaux (Spring 2012) and Rohit Karnik (Fall 2012).

Appendix C

List of publications

C.1 Journal articles

2023 T. Zenati, B. Figliuzzi and S. H. Ham. Surface oxides characterization based on hyperspectral observations . *Chemometrics and Intelligent Laboratory Systems*, accepted, 2023.

2021 B. Figliuzzi, A. Montaux-Lambert, F. Willot, G. Naudin, P. Dupuis, B. Querleux, E. Huguet. A Bayesian Approach to Morphological Models Characterization. *Image Analysis and Stereology*, 40(3), 171-180, 2021.

2020 K. Chang, B. Figliuzzi. Fast marching based superpixels. *Mathematical Morphology-Theory and Applications*, 4(1), 127-142, 2020.

V. Bortolussi, B. Figliuzzi, F. Willot, M. Faessel, M. Jeandin. Electrical conductivity of metal-polymer cold spray composite coatings onto carbon fiber-reinforced polymer. *Journal of Thermal Spray Technology*, 29(4), 642-656, 2020.

R. Alais, P. Dokládal, A. Erginay, B. Figliuzzi, E. Decencière. Fast macula detection and application to retinal image quality assessment. *Biomedical Signal Processing and Control*, 55, 101567, 2020.

2019 B. Figliuzzi. Eikonal-based models of random tessellations. *Image Analysis and Stereology*, 38(1), 15-23, 2019.

2018 E. Decenciere, R. Alais, P. Dokladal, B. Figliuzzi, A. Erginay. Efficient estimation of eye fundus color image quality with convolutional neural networks. *Investigative Ophthalmology & Visual Science*, 59(9), 1735-1735, 2018.

J. Belhadj, T. Romary, A. Gesret, M. Noble, B. Figliuzzi. New parameterizations for Bayesian seismic tomography. *Inverse Problems*, 34(6), 065007, 2018.

V. Bortolussi, B. Figliuzzi, F. Willot, M. Faessel, M. Jeandin. Morphological modeling of cold spray coatings. *Image Analysis and Stereology*, 37(2), 145-158, 2018.

2016 B. Figliuzzi, W. H. R. Chan, C. R. Buie, J. L. Moran. Nonlinear electrophoresis in the presence of dielectric decrement. *Physical Review E*, 94(2), 023115, 2016.

B. Figliuzzi, D. Jeulin, M. Faessel, F. Willot, M. Koishi, N. Kowatari. Modelling the microstructure and the viscoelastic behaviour of carbon black filled rubber materials from 3D simulations. *Technische Mechanik*, 32(1-2), 22-46, 2016.

2015 J. B. Gasnier, F. Willot, H. Trumel, B. Figliuzzi, D. Jeulin, M. Biessy. A Fourier-based numerical homogenization tool for an explosive material. *Matériaux et Techniques*, 103(3), 308, 2015.

2014 B. Figliuzzi, W. H. R. Chan, J. L. Moran, C. R. Buie. Nonlinear electrophoresis of ideally polarizable particles. *Physics of Fluids*, 26(10), 102002, 2014.

Y. S. Joung, B. Figliuzzi, C. R. Buie. Design of capillary flows with functionally graded porous titanium oxide films fabricated by anodization instability. *Journal of colloid and interface science*, 423, 143-150, 2014.

2013 B. Figliuzzi, C. R. Buie. Rise in optimized capillary channels. *Journal of Fluid Mechanics* 731, 142-161, 2013.

2012 B. Figliuzzi, D. Jeulin, A. Lemaitre, G. Fricout, J.J. Piezanowski, P. Manneville. Rheology of thin films from flow observations.*Experiments in Fluids* 1-11, 2012.

B. Figliuzzi, D. Jeulin, A. Lemaitre, G. Fricout, J.J. Piezanowski, P. Manneville. Numerical Simulation of Thin Paint Film Flow. *Journal Of Mathematics In Industry* 2(1), 1-20, 2012.

C.2 Conference proceedings

2022 T. Zenati, B. Figliuzzi and S. H. Ham. Characterization of surface oxides from hyperspectral measurements. In *Workshop on Hyperspectral Image and Signal Processing : Evolution in Remote Sensing (WHISPERS)* (pp. 1-5). IEEE. 2022.

T. Zenati, B. Figliuzzi and S. H. Ham. A supervised approach for the detection of surface oxides from hyperspectral measurements. In *Workshop on Hyperspectral Image and Signal Processing : Evolution in Remote Sensing (WHISPERS)* (pp. 1-5). 2022.

T. Zenati, B. Figliuzzi and S. H. Ham. Surface Oxide Detection and Characterization Using Sparse Unmixing on Hyperspectral Images. In *International Conference on Image Analysis and Processing* (pp. 291-302). Springer. 2022.

2019 K. Chang and B. Figliuzzi. Fast marching based superpixels generation. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing* (pp. 350-361). Springer, July 2019.

2017 M. Koishi, N. Kowatari, B. Figliuzzi, M. Faessel, F. Willot, D. Jeulin. Computational material design of filled rubbers using multi-objective design exploration. In *Constitutive Models for Rubber X* (pp. 467-473). CRC Press, 2017.

B. Figliuzzi, K. Chang, M. Faessel. Hierarchical segmentation based upon multi-resolution approximations and the watershed transform. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing* (pp. 185-195). Springer, 2017.

2016 B. Figliuzzi, S. Velasco-Forero, M. Bilodeau, J. Angulo. A bayesian approach to linear unmixing in the presence of highly mixed spectra. In *International Conference on Advanced Concepts for Intelligent Vision Systems* (pp. 263-274). Springer, 2016.

V. Bortolussi, F. Borit, A. Chesnaud, M. Jeandin, M. Faessel, B. Figliuzzi, G. Surdon. Cold spray of metal-polymer composite coatings onto carbon fiber-reinforced polymer (CFRP). In *International Thermal Spray Conference*, 2016.

C.3 Conferences

2015 B. Figliuzzi, W. H. R. Chan, C. R. Buie and J. Moran. Dielectric decrement effects in electrokinetics. *Bulletin of the American Physical Society*, 60, Boston, 2015.

J. B. Gasnier, B. Figliuzzi, M. Faessel, F. Willot, D. Jeulin, H. Trumel. 3D Morphological modeling of a polycrystaline microstructure with non-convex, anisotropic grains. In *14th International Congress for Stereology and Image Analysis (ICSIA)*, 2015.

2014 J. L. Moran, W. H. R. Chan, C. R. Buie, B. Figliuzzi. Dielectric Decrement Effects on Nonlinear Electrophoresis of Ideally Polarizable Particles. *Bulletin of the American Physical Society*, 59, 2014.

2013 Y. S. Joung, B. Figliuzzi, C. R. Buie. Design of Capillary Flows with Spatially Graded Porous Films. In *APS Meeting Abstracts*, 2013.

2012 B. Figliuzzi, C. R. Buie. Wicking flow in optimized capillary channels. *Bulletin of the American Physical Society*, 57, San Diego, Nov. 2012.

B. Figliuzzi, D. Jeulin, A. Lemaitre, G. Fricout. In situ determination of the local rheological parameters of a fluid by inverse method. *2nd International Conference on Material Modelling*, Paris, Sept. 2011.

B. Figliuzzi, D. Jeulin, G. Fricout, J.J. Piezanowski. Characterization of surface roughness by means of wavelet analysis. *13^{th} International Conference on Metrology and Properties of Engineering Surfaces*, London, April 2011.

2011 B. Figliuzzi, D. Jeulin, A. Lemaitre, G. Fricout. In situ determination of the local rheological parameters of a fluid by inverse method. *2nd International Conference on Material Modelling*, Paris, Sept. 2011.

B. Figliuzzi, D. Jeulin, G. Fricout, J.J. Piezanowski. Characterization of surface roughness by means of wavelet analysis. *13^{th} International Conference on Metrology and Properties of Engineering Surfaces*, London, April 2011.