PEOPLES DEMOCRATIC REPUBLIC OF ALGERIA
MINISTRY OF HIGHER EDUCATION AND SCIENTIFIC RESEARCH

Mohamed Boudiaf University of M'Sila
Faculty of Mathematics and Computer Science
Department of Mathematics

Master memory

Field: Mathematics and Computer Science
Department : Mathematics
Option : Partial Differential Equations and Applications

Theme

Chan–Vese Model for Image Segmentation: Theory and Numerical Implementation

Presented by:
Boularas Yasmina

Publicly defended on : September 24, 2026

In front of the jury composed of:

BASTI Bilal	MCA ,	M'sila University	**Chairperson**
CHOUDER Rafaa	MCA ,	M'sila University	**Supervisor**
BENHAMIDOUCHE Noureddine	Prof ,	M'sila University	**Co-Supervisor**
BOUNAB Noura	MCA ,	M'sila University	**Examiner**

University year 2025/2026

Contents

General Introduction 1

1 *Preliminaries and Basic Definitions* 4

1.1 Introduction to Image Processing . . . 4

1.1.1 The Image Society . . . 4

1.1.2 Mathematical Definition of an Image . . . 5

1.1.3 What is Image Processing ? . . . 6

1.2 Image Derivatives, Gradients and Edge Detection . . . 8

1.2.1 Classical Edge Detection Operators: An Overview . . . 10

1.2.2 Limits of Purely Gradient-Based Segmentation . . . 10

1.2.3 Region-Based versus Edge-Based Segmentation . . . 11

1.3 Elements of Differential Geometry: Curvature . . . 11

1.4 Partial Differential Equations . . . 15

1.5 Level Set Method . . . 18

2 *From the Mumford–Shah Model to the Chan–Vese Model* 21

2.1 Image Segmentation as a Variational Problem . . . 21

2.2 The Mumford–Shah Functional . . . 22

2.3 The Space of Functions of Bounded Variation . . . 24

2.4 The Chan–Vese Simplification: The Piecewise-Constant, Two-Phase Case . . . 26

2.5 Comparison with Classical Active Contour Models . . . 29

3 *Mathematical Formulation of the Chan–Vese Model* 33

3.1 Elements of the Calculus of Variations . . . 33

3.2 Level Set Representation of the Chan–Vese Energy . . . 36

3.3 Optimal Values of the Region Constants . . . 37

3.4 First Variation and the Euler–Lagrange Equation . . . 38

3.5 The Chan–Vese Evolution Equation . . . 40

3.6 Regularization of the Heaviside and Dirac Functions . . . 41

3.7 Boundary Condition . . . 42

3.8 Role of the Model Parameters 42

4 *Numerical Implementation and Experiments* **45**

4.1 Finite Difference Discretization of the Evolution Equation 45

4.2 Numerical Stability and Energy Monitoring 47

4.3 Boundary Condition 48

4.4 Reinitialization of the Level Set Function 48

4.5 The Complete Algorithm 50

4.6 Experimental Results 51

General Conclusion **58**

Bibliography **60**

كلية الرياضيات والإعلام
Faculté des Mathématiques et de l'Informatique

Dedication

In the name of God,the most gracious,the most merciful.

*I dedicate this work to my dear father **Abdelkarim**, who taught me the value of effort, patience, and honesty, and whose sacrifices have brought me to this day.*

*To my beloved mother **Noura**, the heart of our home, whose tenderness and endless prayers have carried me through my hardest moments.*

*To my brothers especially my brother **Islam**, the light of my eyes and sisters, my safe place and my greatest support. To my family and dear friends, who shared with me the laughter, the stress, and the long nights.*

To my supervisor Mr. Rafaa Chouder, my Co-supervisor Mr. Benhamidouche Noureddine and all my teachers, for their guidance and trust. And to everyone who believed in me, even when I doubted myself. With all my love and gratitude, I dedicate this humble work to you..

Acknowledgments

First of all, I would like to thank **"Allah"** the all-powerful, for having given us strength and patience since born in this life.

I am grateful to my supervisor, Mr. **Rafaa Chouder** and my Co-supervisor Mr. **Noureddine Benhamidouche** for their accurate guidance, support, and suggestions that made for this work. Their valuable advice carried me through every stage of writing my memory. It is truly impossible to adequately express my appreciation to them in mere words.

I extend my respect and gratitude to the members of the Jury, Dr. **Bilal Basti** and Dr. **Noura Bounab**, for graciously agreeing to examine and evaluate my work. Their insightful remarks and suggestions will undoubtedly contribute to enhancing the quality of this memory.

My family, especially my parents, played an indispensable role in the completion of this work. Their unwavering support and encouragement have been invaluable.

Last but not least, I would like to express my gratitude to all my teachers and everyone who has directly or indirectly contributed to the realization of this work.

Faculté des Mathématiques et de l'Informatique

Abstract

Image segmentation is a fundamental task in image processing and computer vision. This thesis presents a mathematical study and numerical implementation of the two-phase Chan–Vese model, a variational active contour model based on regional statistics rather than explicit edge detection.

The mathematical formulation is developed within the level set framework, starting from the Mumford–Shah model and using the calculus of variations to derive the evolution equation. Regularized Heaviside and Dirac delta functions are introduced for numerical treatment, with region-wise constants updated during the evolution.

For the numerical solution, the nonlinear partial differential equation is discretized using finite differences on a uniform grid and an explicit forward Euler scheme. The procedure incorporates the curvature term, boundary conditions, time-step selection, energy monitoring, and periodic Godunov-type reinitialization. The algorithm is implemented in MATLAB and tested on synthetic and real images. Results are evaluated using qualitative visual assessment, the Dice coefficient, and discrete energy evolution. The study also discusses limitations of the piecewise-constant two-phase formulation.

Keywords: Image segmentation; Chan–Vese model; active contours; level set method; Mumford–Shah model; calculus of variations; finite differences.

ملخص

تُعدّ تجزئة الصور من المهام الأساسية في معالجة الصور والرؤية الحاسوبية، وتهدف إلى تقسيم الصورة إلى مناطق ذات معنى. تقدم هذه المذكرة دراسة رياضية وتنفيذًا عدديًا لنموذج Chan–Vese ثنائي الطور، وهو نموذج تغايري للمنحنيات النشطة يعتمد على الإحصاءات الخاصة بالمناطق بدلًا من الكشف الصريح عن الحواف.

تُطوَّر الصياغة الرياضية ضمن إطار مجموعة المستوى، انطلاقًا من نموذج Mumford–Shah وباستخدام حساب التغايرات لاشتقاق معادلة التطور. كما يتم إدخال دالتي Heaviside و Dirac المنتظمتين للمعالجة العددية، مع تحديث الثوابت الخاصة بالمناطق أثناء التطور.

وللحل العددي، تُجزَّأ معادلة التفاضل الجزئية غير الخطية بالفروق المحدودة على شبكة منتظمة ومخطط Euler الصريح. ويتضمن الإجراء حد الانحناء، والشروط الحدية، واختيار الخطوة الزمنية، ومراقبة الطاقة، وإعادة التهيئة الدورية من نوع Godunov . وقد تم تنفيذ الخوارزمية باستخدام MATLAB واختبارها على صور اصطناعية وحقيقية، وتقييم النتائج بصريًا وباستخدام معامل Dice وتطور الطاقة المتقطعة. كما تناقش الدراسة أهم حدود الصياغة ثنائية الطور ذات القيم الثابتة على مستوى المناطق.

الكلمات المفتاحية : تجزئة الصور؛ نموذج Chan–Vese ؛ المنحنيات النشطة؛ طريقة مجموعة المستوى؛ نموذج Mumford–Shah ؛ حساب التغايرات؛ الفروق المحدودة.

General Introduction

Image segmentation is one of the fundamental tasks in image processing and computer vision. Its main objective is to partition an image into meaningful regions corresponding, for example, to objects, structures, or homogeneous areas. Accurate segmentation is an essential preliminary step in many applications, including medical image analysis, object recognition, computer vision, remote sensing, and biological image processing. Despite its apparent simplicity, image segmentation remains a challenging problem, because real images may contain noise, weak or discontinuous boundaries, intensity variations, and objects with complex geometries.

Among the mathematical approaches developed for image segmentation, variational and partial differential equation (PDE)-based methods have received considerable attention. These approaches formulate segmentation as an optimization or evolution problem in which a suitable energy functional measures the quality of a proposed partition. The resulting mathematical formulation provides a natural framework for incorporating geometric properties of contours, statistical information from image regions, and regularization mechanisms that improve the stability of the segmentation.

A major milestone in variational image segmentation is the Mumford–Shah model [9], which provides a general framework for simultaneously approximating an image by a piecewise-smooth function and identifying its discontinuity set. Although mathematically powerful, the original model leads to a difficult free-discontinuity problem. This motivated the development of simplified formulations adapted to particular segmentation tasks. In this context, the Chan–Vese model [6] provides a region-based active contour formulation derived from a piecewise-constant approximation of the Mumford–Shah functional.

Unlike classical edge-based active contour models, the Chan–Vese model does not require the image gradient to be explicitly used to stop the evolving contour at an edge. Instead, it relies primarily on statistical information describing the intensities inside and outside the evolving contour. This feature allows the model to handle boundaries that may be weak, blurred, or not clearly defined by local image gradients. At the same time, the length regularization term introduces a geometric constraint that controls the complexity of the evolving contour.

The mathematical formulation of the Chan–Vese model is naturally expressed using the level set method. Rather than representing the contour explicitly as a parametric curve, the contour is rep-

resented implicitly as the zero level set of a higher-dimensional function. This formulation makes it possible to handle changes in topology during the evolution, such as the merging or splitting of contour components. The associated energy functional combines a length term, an optional area term, and two region-based fidelity terms. The corresponding Euler–Lagrange equation leads to a nonlinear evolution equation involving the curvature of the level set function and the difference between the squared deviations of the image intensities from the estimated regional means.

From a computational point of view, solving the Chan–Vese evolution equation requires suitable numerical approximations of the spatial derivatives, curvature term, regularized Heaviside and Dirac delta functions, as well as an appropriate time discretization. The numerical implementation must also take into account boundary conditions and the possible need for reinitialization of the level set function. These aspects are important because the quality and stability of the numerical solution directly influence the resulting segmentation.

The main objective of this work is therefore to study the Chan–Vese model from both mathematical and computational perspectives, with particular emphasis on its formulation as a level set variational model and its numerical implementation. The work begins with the mathematical concepts required to understand the model, including image derivatives, geometric notions, curvature, and level set functions. It then examines the connection between the Mumford–Shah framework and the Chan–Vese model, highlighting the piecewise-constant, two-phase formulation, its variational interpretation, and the functional-analytic setting – the space of functions of bounded variation – in which existence of a minimizer can be established. This is followed by a recollection of the calculus-of-variations tools needed to pass rigorously from an energy functional to a partial differential equation, and by the full derivation of the Chan–Vese evolution equation within the level set framework.

Finally, a finite-difference numerical scheme is developed for the discretization of the level set evolution equation. The spatial derivatives and curvature are approximated on a regular Cartesian grid, while an explicit Euler scheme is used for the temporal evolution. Numerical considerations concerning the time-step restriction, boundary treatment, energy evaluation, convergence, and level set reinitialization are also discussed. A MATLAB implementation is used to perform numerical experiments on representative images in order to illustrate the behavior of the model and examine the influence of its principal parameters.

The remainder of this work is organized into four chapters. Chapter 1 introduces the mathematical and image-processing concepts required throughout the study: image representation, gradients and edge detection, the region-based point of view, curvature, partial differential equations, and the level set method. Chapter 2 presents the transition from the Mumford–Shah model to the Chan–Vese formulation; it introduces, immediately after the Mumford–Shah functional, the

space of functions of bounded variation needed to make its length-regularization term rigorous, and uses this framework both to derive the Chan–Vese energy and to establish existence of a minimizer, before comparing the model with other active contour approaches. Chapter 3 opens by recalling the calculus-of-variations tools – functionals, first variation, the Euler–Lagrange equation – needed to pass rigorously from an energy to a PDE, then develops the mathematical formulation of the Chan–Vese model using the level set framework and derives its associated evolution equation. Chapter 4 focuses on the numerical implementation, including spatial and temporal discretization, curvature approximation, reinitialization, convergence considerations, and numerical experiments.

Chapter 1

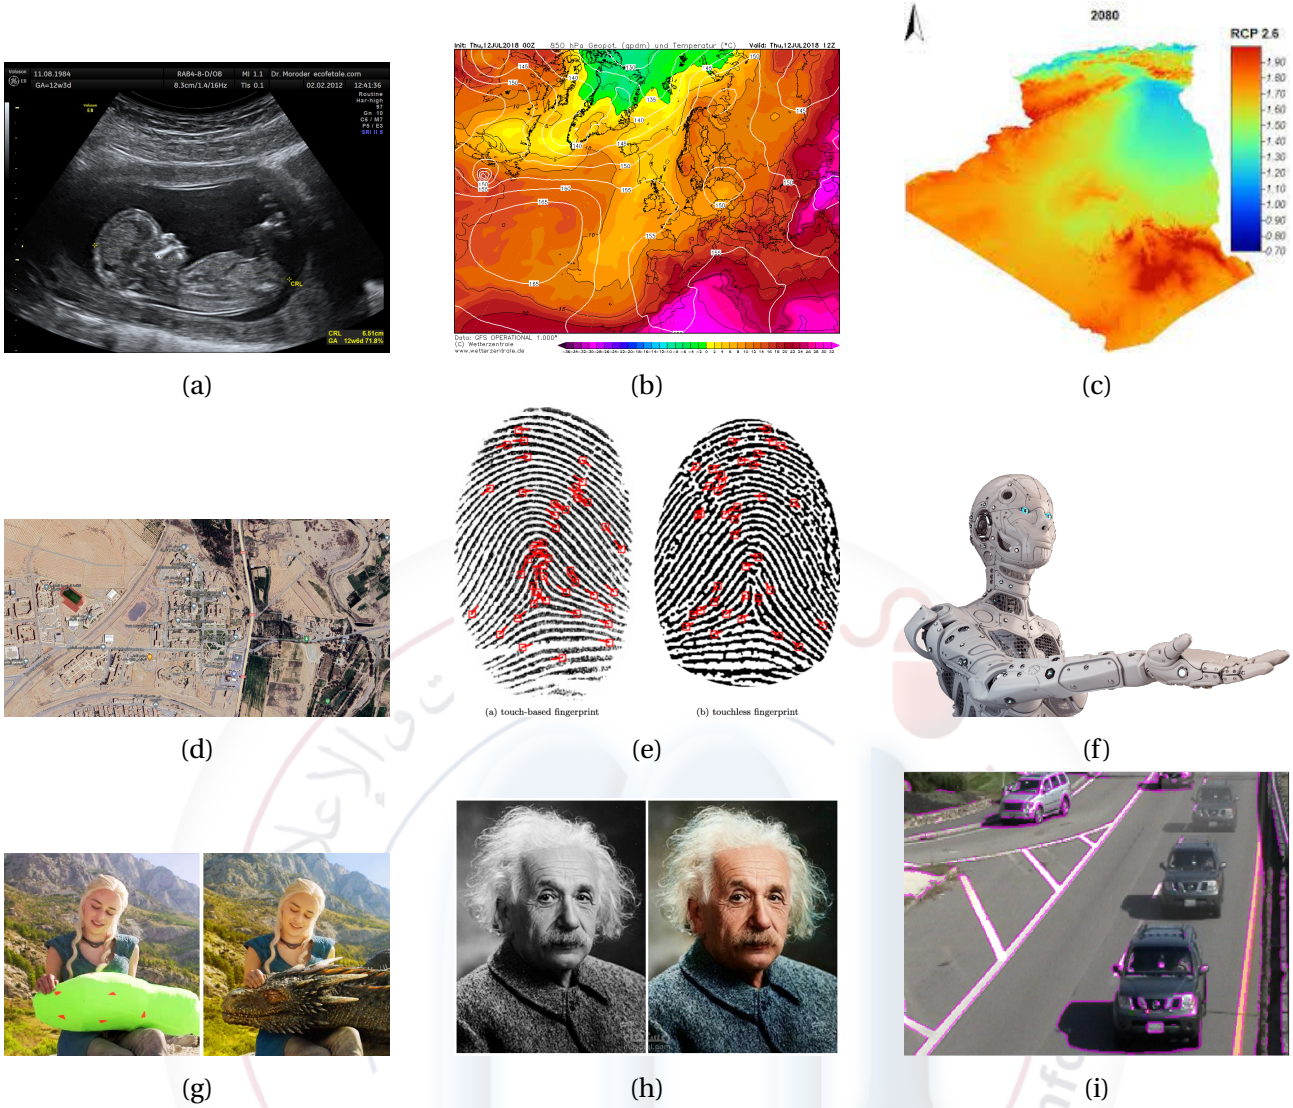
Preliminaries and Basic Definitions

The goal of this chapter is to introduce the fundamental concepts of image processing and to present the mathematical tools that will be used throughout the subsequent chapters. It provides the necessary background for the study of the Chan–Vese segmentation model by first presenting basic notions of image representation, gradients, and edge detection, together with the alternative region-based point of view on which Chan–Vese relies. It then introduces essential elements of differential geometry, with particular emphasis on curvature and geometric flows that govern contour evolution. Finally, it discusses partial differential equations and the level set method, which together constitute the main analytical and computational framework used to derive and implement the Chan–Vese model in the following chapters. Two further tools that are used repeatedly afterwards – the calculus of variations and the space of functions of bounded variation – are introduced later, directly where they are first put to work: the former opens Chapter 3, and the latter is placed in Chapter 2, immediately after the Mumford–Shah functional whose regularization term it makes rigorous.

1.1 Introduction to Image Processing

1.1.1 The Image Society

Images are a central means of communication in modern life, valued for their compactness and immediacy: they capture moments and events, and can be analysed, enhanced, or combined with other information. Image processing methods are used across almost every area of science and technology – medical imaging, satellite and aerial imaging, weather forecasting, fingerprint analysis, robotics, quality control, and multimedia and video processing, among others (see Figure 1.1).



(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

Figure 1.1: Some applications where image processing is concerned.

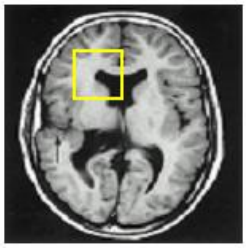
1.1.2 Mathematical Definition of an Image

In image processing, a continuous Gray-scale image is considered that is a two-dimensional function

$$\begin{aligned} u \quad &: \quad \mathbb{R}^2 \to \mathbb{R} \\ &\quad (x,y) \to u(x,y) \end{aligned}$$

The variable (x,y), for digital images, actually belongs to $\mathbb{N}^2$, it is a pixel of the image. $u(x,y)$ represents the luminous intensity, or the Gray level of the image at the pixel (x,y).

Digital images are most commonly presented as a matrix of scalars for gray-scale images or vectors for colour images. Digital Gray scale images on the other hand are sampled and quantized. Sampling is the discretization of the image domain. Here an image consists of Gray values of a rectangular point grid $\{u_{i,j} \setminus \ i = 0,...,N-1 \text{ and } j = 0,...,M-1\}$. A grid point (i,j) is called pixel. Here N and M are the width and height of the image in pixels, respectively. $u_{i,j}$ denotes the Gray



125	131	147	171	173	154	145	128	123
121	124	142	163	169	161	147	135	128
135	131	139	148	155	163	155	144	138
150	150	141	146	150	155	158	157	158
193	189	171	161	162	167	176	182	188
197	207	204	199	198	199	203	205	206
183	202	215	222	219	217	217	214	208
181	198	210	216	218	217	217	37	22
180	197	209	213	216	212	212	91	25
191	201	211	215	214	208	208	21	23
206	208	210	215	215	211	212	16	14

Figure 1.2: Gray level digital image

value of pixel (i, j). (See Figure 1.2).

For a colour image, it is not enough to know its Gray level for each pixel: it is necessary to know the intensity of each of the three channels of the fundamental colours, the red R, the green G, and the blue B. A colour image can then be defined as a vector function, (See Figure 1.3),

$$\begin{aligned} u \quad &: \quad \mathbb{R}^2 \to \mathbb{R}^3 \\ &\quad (x, y) \to u(x, y) = (R(x, y); G(x, y); B(x, y)). \end{aligned}$$

1.1.3 What is Image Processing ?

Image processing is a branch of computer science and applied mathematics concerned with the analysis, enhancement, and manipulation of digital images. Its primary objective is to transform images into a form that is either more suitable for human interpretation or more appropriate for subsequent automated analysis.

Closely related to image processing is the field of computer vision, whose goal is to extract meaningful information from visual data and to emulate, as much as possible, the capabilities of the human visual system. Computer vision involves higher-level tasks such as image segmentation, object recognition, three-dimensional reconstruction, motion analysis, and scene understanding. To accomplish these tasks, computer vision systems often rely on image processing techniques combined with machine learning, variational methods, optimization techniques, and geometric approaches.

In contrast, image processing mainly focuses on improving the quality of images or extracting relevant low-level information from them. Typical image processing tasks include:

- **Feature extraction:** Detecting important image structures such as edges, corners, ridges, and texture patterns, which play a crucial role in visual perception and subsequent analysis.



$RED ==$

124	120	116	116	114	120	125	100	105	128	119
122	117	115	114	112	112	116	89	95	120	113
118	115	113	113	112	111	114	87	93	119	112
120	117	115	116	114	110	113	85	92	117	109
155	155	149	148	152	153	147	141	139	143	147
178	178	173	172	175	180	176	172	170	173	174

$GREEN ==$

145	141	139	139	137	139	142	114	118	141	136
141	138	135	137	135	132	135	106	109	134	132
137	134	133	133	132	136	137	107	112	138	135
139	136	135	136	134	138	138	108	112	137	134
147	147	141	140	144	144	138	132	132	136	140
166	166	161	160	163	167	163	159	160	163	166

$BLUE ==$

150	146	145	145	143	146	150	123	127	150	144
147	143	142	143	141	139	142	114	118	143	139
143	140	140	140	139	141	143	114	119	145	141
145	142	142	143	141	142	143	114	119	144	139
136	136	130	129	133	135	129	123	122	126	130
154	154	149	148	151	158	154	150	150	153	155

Figure 1.3: Color image

- **Image segmentation:** Partitioning an image into homogeneous regions with similar characteristics, while accurately locating the boundaries between different objects or regions.
- **Image deblurring:** Recovering a sharp image from a blurred observation by compensating for image degradation caused by motion, defocus, or other acquisition imperfections.
- **Image denoising:** Removing unwanted noise while preserving significant image details such as edges and fine structures.
- **Image enhancement:** Improving visual quality through contrast adjustment, histogram equalization, sharpening, and other enhancement techniques.

Image processing constitutes the foundation of many modern applications, including medical imaging, remote sensing, surveillance systems, industrial inspection, autonomous vehicles, and multimedia technologies.
Among all image processing tasks, image segmentation occupies a central position because it provides the first step toward object extraction and scene understanding. The Chan–Vese model studied in this dissertation belongs to this class of segmentation techniques. Unlike edge-based approaches, it relies on a variational, region-based formulation combined with the level set method, and it is derived from the more general Mumford–Shah functional, as will be detailed in Chapter 2.

1.2 Image Derivatives, Gradients and Edge Detection

The detection of object boundaries is a fundamental task in image processing and computer vision. A large family of segmentation methods, known as *edge-based methods*, relies on the observation that significant intensity variations often correspond to object edges. Consequently, differential operators play a central role in extracting boundary information from images, and classical active contour models such as Snakes and Geodesic Active Contours are built directly on this idea. The gradient is one of the most important mathematical tools for measuring local intensity variations. It provides both the magnitude and direction of the strongest intensity change and therefore serves as the basis for many edge detection techniques. We review these classical notions here for two reasons: first, because they form part of the common mathematical background of PDE-based segmentation, and second, because they will later be used in Chapter 4 as a basis of comparison with the region-based Chan–Vese model, which – as shown in Section 1.2.3 below – does not itself rely on the image gradient.

Image Derivatives

Let $I : \Omega \subset \mathbb{R}^2 \to \mathbb{R}$ be a grayscale image defined on the domain Ω.

The first-order partial derivatives of the image are

$$I_x = \frac{\partial I}{\partial x}, \qquad I_y = \frac{\partial I}{\partial y}.$$

These derivatives measure the rate of intensity variation along the horizontal and vertical directions, respectively.

For a discrete image, derivatives are approximated numerically. The forward difference approximations are given by

$$I_x(i,j) \approx I(i+1,j) - I(i,j),$$

and

$$I_y(i,j) \approx I(i,j+1) - I(i,j).$$

More accurate approximations are obtained using centered differences:

$$I_x(i,j) \approx \frac{I(i+1,j) - I(i-1,j)}{2},$$

$$I_y(i,j) \approx \frac{I(i,j+1) - I(i,j-1)}{2}.$$

These approximations constitute the basis of many numerical edge detection operators, and the same finite-difference philosophy will be used in Chapter 4 to discretize the Chan–Vese evolution equation.

Gradient Magnitude and Orientation

The image gradient is defined as

$$(I_x, I_y).$$

The gradient vector points in the direction of the maximum increase of image intensity, and its magnitude

$$\sqrt{I_x^2 + I_y^2}$$

measures the strength of local intensity variations, while its orientation

$$\tan^{-1}\left(\frac{I_y}{I_x}\right)$$

indicates the direction of steepest intensity increase. Regions with nearly constant intensity correspond to small gradient values, whereas edges correspond to large gradient magnitudes: since image edges are characterized by abrupt intensity transitions, they are generally associated with local maxima of the gradient magnitude. Consider for instance a one-dimensional intensity profile $I(x)$. In a homogeneous region, $I(x)$ = constant, and therefore $\frac{dI}{dx} = 0$, whereas at an edge location the intensity changes rapidly, so that $\left|\frac{dI}{dx}\right| \gg 0$. This observation naturally extends to two-dimensional images through the gradient norm $|\nabla I|$.

In practice, image edges correspond to discontinuities or sharp transitions in the image intensity function. These discontinuities may arise from changes in object reflectance, illumination, texture, or depth. Detecting such transitions is a key step in edge-based segmentation and object extraction algorithms.

Second-Order Derivatives and the Laplacian

Second-order derivatives provide complementary information about image structures. The Laplacian operator is defined by

$$\Delta I = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2},$$

and measures the local curvature of the image intensity surface. Zero-crossings of the Laplacian of a smoothed image often indicate the presence of edges, an idea used in the Laplacian-of-Gaussian edge detector [8].

1.2.1 Classical Edge Detection Operators: An Overview

Several classical numerical operators have been proposed to estimate image gradients on a discrete grid, typically as small convolution masks applied to the image: the Roberts operator [14], computationally cheap but noise-sensitive; the Prewitt [13] and Sobel [15] operators, which incorporate local averaging for better robustness to noise; and, most notably, the Canny detector [4], which combines Gaussian smoothing, gradient computation, non-maximum suppression, and hysteresis thresholding to achieve good edge localization with a single response per edge. Since differentiation amplifies noise, all of these operators are typically preceded by Gaussian smoothing, $G_\sigma * I$, whose scale parameter σ trades off detail preservation against noise reduction. We recall these operators here only for completeness, as a point of comparison in Chapter 4; none of them is used by the Chan–Vese model itself.

1.2.2 Limits of Purely Gradient-Based Segmentation

Edge-based methods such as Snakes and Geodesic Active Contours use image gradients to build an edge indicator function, typically of the form

$$g(|\nabla I_\sigma|) = \frac{1}{1+|\nabla I_\sigma|^2}, \qquad \text{or} \qquad g(|\nabla I_\sigma|) = e^{-|\nabla I_\sigma|^2},$$

which satisfies $0 < g \leq 1$, tends to 1 in homogeneous regions ($|\nabla I_\sigma| \approx 0$), and tends to 0 near strong edges. The evolving contour is then attracted toward locations where g is small, i.e. toward regions of large gradient.

While effective for images with sharp, well-contrasted boundaries, this strategy has two well-known limitations: it depends on a pre-smoothing scale σ that must be tuned, and it performs poorly

when boundaries are weak, blurred, or corrupted by noise, since $|\nabla I_\sigma|$ may never become large enough along the true boundary. These limitations motivate an alternative, *region-based* formulation of the segmentation problem, which is the point of view adopted by Mumford–Shah and, subsequently, by Chan–Vese.

1.2.3 Region-Based versus Edge-Based Segmentation

Segmentation methods can broadly be classified into two families.

Edge-based methods search directly for the object boundary by following locations of high gradient magnitude, as described above. The final contour is entirely determined by local gradient information along its own trajectory.

Region-based methods, by contrast, look for a partition of the image domain Ω into subregions that are homogeneous with respect to some statistical descriptor of intensity – typically the mean, but possibly also the variance or texture. Rather than asking "where is the gradient large?", a region-based model asks "which partition of Ω best explains the observed image as a collection of nearly constant pieces?". The boundary between regions is then obtained indirectly, as the interface that separates statistically different areas, and not as a curve chasing large gradient values. This distinction is essential for the present work: the Chan–Vese model, which will be derived in Chapter 2 as a simplified case of the Mumford–Shah functional, is a *region-based* model. It measures the fit of a candidate partition using the mean intensities inside and outside the evolving contour, and does not use $|\nabla I|$ or any edge indicator function $g(|\nabla I|)$ in its energy. As a consequence, it remains applicable to images with smooth or noisy boundaries where purely gradient-based methods tend to fail, at the cost of being restricted (in its basic two-phase form) to images that are reasonably well approximated by piecewise-constant regions.

1.3 Elements of Differential Geometry: Curvature

In this section we recall some basic definitions and properties of differential geometry and focus on the notion of curvature. Curvature has different meanings depending on the "object" that we are considering. We shall define the notion of curvatures for parametrized curves (for snakes, for instance) and for images that can be represented by their isovalues.

Parametrized Curves

Let $x(p) = \big(x_1(p), x_2(p)\big)$ be a regular planar oriented curve in $\mathbb{R}^2$, $0 < p < 1$. We denote

$$
\begin{aligned}
&T(p) = x'(p) = \big(x_1'(p), x_2'(p)\big), \text{ the tangent vector at } x(p),\\
&N(p) = \big(-x_2'(p), x_1'(p)\big), \text{ the normal vector at } x(p),\\
&s(p) = \int_0^p \sqrt{\big(x_1'(r)\big)^2 + \big(x_2'(r)\big)^2}\, dr, \text{ the curvilinear abscissa (or arc length).}
\end{aligned}
$$

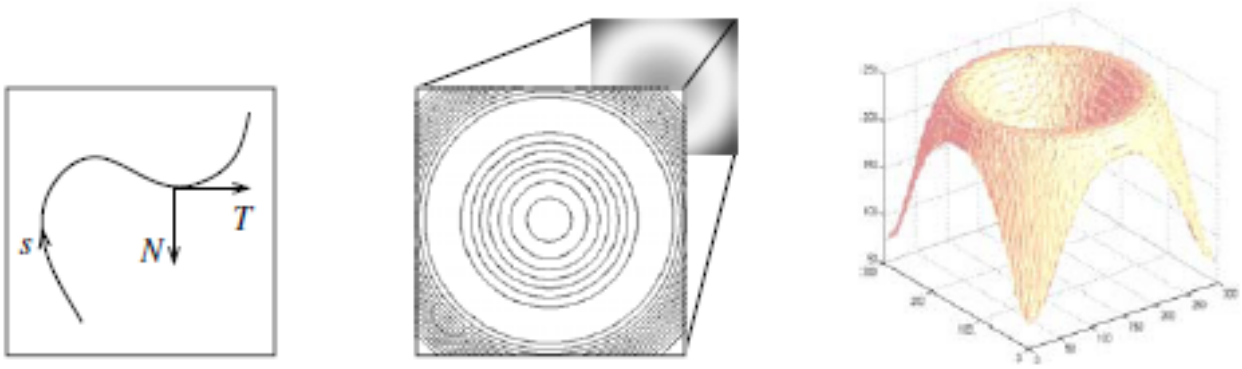


Figure 1.4: Different "objects" encountered in image analysis: parametrized curves and curves as level sets of an image.

If $x(p)$ is regular, we can parametrize it by s, and then $T(s) = \frac{dx}{ds}(s)$ is such that $|T(s)| = 1$. The curvature vector is $\frac{dT}{ds}$, and since $T = \frac{dx}{ds}$, we have

$$\frac{dT}{ds} = \frac{d^2x}{ds^2}.$$

One can show that the curvature vector is collinear to $\frac{N(s)}{|N(s)|}$, i.e.,

$$\frac{dT}{ds}(s) = \kappa(s)\frac{N(s)}{|N(s)|},$$

where $\kappa(s)$ is the curvature and $\frac{1}{|\kappa(s)|}$ is the radius of curvature.
For any parametrization we have

$$\kappa(p) = \frac{x_1'(p)x_2''(p) - x_2'(p)x_1''(p)}{\left((x_1'(p))^2 + (x_2'(p))^2\right)^{3/2}}. \tag{1.1}$$

Moreover,

$$\frac{1}{|x'(p)|}\frac{\partial}{\partial p}\left(\frac{x'(p)}{|x'(p)|}\right) = \kappa(p)\left(\frac{N(p)}{|N(p)|}\right). \tag{1.2}$$

Curves as Level Sets of a Function u

Let us now consider the case where $x(s)$, parametrized by its curvilinear abscissa, is the k-level set of a function

$$u : \mathbb{R}^2 \to \mathbb{R},$$

that is,

$$x(s) = \{(x_1(s), x_2(s)) \; ; \; u(x_1(s), x_2(s)) = k\}.$$

By differentiating the equality $u(x_1(s), x_2(s)) = k$ with respect to s, we obtain

$$x_1'(s)\, u_{x_1} + x_2'(s)\, u_{x_2} = 0. \tag{1.3}$$

where u_{x_1} stands for $\frac{\partial u}{\partial x_1}\big(x_1(s), x_2(s)\big)$. Therefore, the vectors $\big(x_1'(s), x_2'(s)\big)$ and $(-u_{x_2}, u_{x_1})$ are collinear. For some λ we have

$$\begin{cases} x_1'(s) = -\lambda u_{x_2}, \\ x_2'(s) = \lambda u_{x_1}, \end{cases} \tag{1.4}$$

so the vectors (u_{x_1}, u_{x_2}) and $(-u_{x_2}, u_{x_1})$ are respectively normal and tangent to the curve $x(s)$.

If we differentiate again (1.3) with respect to s, we obtain

$$(x_1'(s))^2 u_{x_1x_1} + (x_2'(s))^2 u_{x_2x_2} + 2x_1'(s)x_2'(s)u_{x_1x_2} + x_1''(s)u_{x_1} + x_2''(s)u_{x_2} = 0.$$

Using (1.4), this becomes

$$\lambda^2\big((u_{x_1})^2 u_{x_2x_2} + (u_{x_2})^2 u_{x_1x_1} - 2u_{x_1}u_{x_2}u_{x_1x_2}\big) + \frac{1}{\lambda}\big(x_1''(s)x_2'(s) - x_2''(s)x_1'(s)\big) = 0.$$

Since $|x'(s)| = 1$, we obtain from (1.4) $\lambda^2 = \frac{1}{|\nabla u|^2}$.

Therefore, using (1.1), we deduce the expression of the curvature (of course we assume $|\nabla u(x)| \neq 0$):

$$\kappa = \frac{(u_{x_1})^2 u_{x_2x_2} + (u_{x_2})^2 u_{x_1x_1} - 2u_{x_1}u_{x_2}u_{x_1x_2}}{\big((u_{x_1})^2 + (u_{x_2})^2\big)^{3/2}}. \tag{1.5}$$

The curvature can be expressed as

$$\kappa = \operatorname{div}\left(\frac{\nabla u}{|\nabla u|}\right). \tag{1.6}$$

This formula shows that the curvature of a level set can be computed directly from the image function without an explicit parametrization of the curve. This property is fundamental in level set methods, and formula (1.6) is precisely the expression that will appear inside the regularization (length) term of the Chan–Vese energy once it is written in level set form in Chapter 3.

This expression reveals an important connection between differential geometry and partial differential equations. The curvature is obtained as the divergence of the unit normal vector field, making it possible to compute geometric quantities directly from image derivatives. This observation is at the heart of geometric active contours and level set methods, including Chan–Vese.

Curvature Flow

A curve evolution can generally be decomposed into a tangential and a normal component,

$$\frac{\partial C}{\partial t} = V_T T + V_N N.$$

The tangential component only changes the parametrization of the curve and does not modify its geometric shape. Therefore, from a geometric point of view, only the normal component influences the evolution of the interface. For this reason, geometric active contour models – including Chan–Vese – are usually written in the form

$$\frac{\partial C}{\partial t} = V_N N.$$

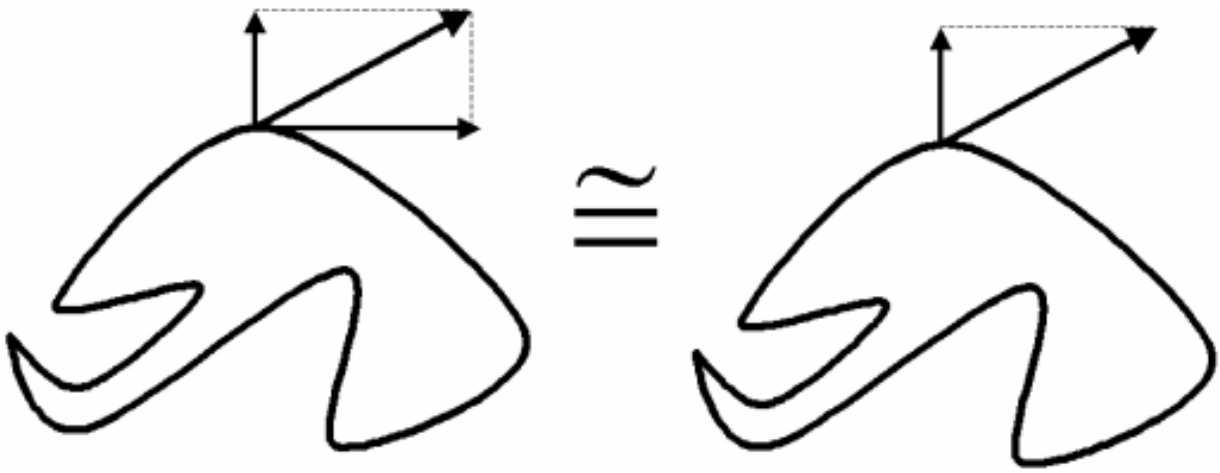


Figure 1.5: Most general form of a curve evolution flow. The velocity is decomposed into its tangential and normal components, the former not affecting the geometry of the flow, just its parameterization.

Figure 1.5 illustrates this decomposition.

Among all possible geometric flows, one of the most fundamental is obtained when the normal velocity is chosen equal to the curvature,

$$V_N = \kappa.$$

The evolution law then becomes

$$\frac{\partial C}{\partial t} = \kappa N.$$

This motion is known as curvature flow. Under curvature flow, highly curved regions move faster than flatter regions, producing a smoothing effect on the contour.

Consequently, irregularities and noise are progressively removed while the overall shape is preserved.

Curvature flow plays a central role in geometric image processing: it will reappear in Chapter 3 as the regularizing part of the Chan–Vese evolution equation, where it is weighted by a length parameter μ and combined with region-based terms that attract the contour towards the actual object boundary rather than shrinking it indefinitely.

Geometric Quantities in the Level Set Framework

In the level set method, a curve is represented implicitly as the zero level set of a function $\phi(x, y, t)$. The contour is defined by

$$C(t) = \{(x, y) \,:\, \phi(x, y, t) = 0\}.$$

Within this framework, the unit normal vector becomes

$$N = \frac{\nabla \phi}{|\nabla \phi|},$$

and the curvature can be expressed as

$$\kappa = \operatorname{div}\left(\frac{\nabla\phi}{|\nabla\phi|}\right).$$

These formulas are of fundamental importance because they allow geometric quantities to be computed directly from the level set function without explicit parameterization of the contour.
This section introduced the main geometric concepts required for the study of active contour models, including tangent and normal vectors, curvature, curve evolution, and curvature-driven flows. These notions constitute the mathematical foundation of geometric deformable models. In Chapter 3, they will be combined with the region-based, gradient-free energy of Mumford–Shah and Chan–Vese to derive the evolution equation studied in this dissertation.

1.4 Partial Differential Equations

Partial Differential Equations (PDEs) constitute one of the most powerful mathematical tools for modeling dynamic phenomena involving several independent variables. They arise naturally in numerous scientific disciplines, including fluid mechanics, heat transfer, elasticity, electromagnetism, and image processing.
Over the last decades, PDE-based techniques have become increasingly important in computer vision and image analysis. Many fundamental tasks such as image denoising, enhancement, restoration, segmentation, and contour detection can be formulated as the solution of suitable partial differential equations.
In particular, both edge-based active contour models and region-based models such as Chan–Vese rely on geometric and variational PDEs to describe the evolution of object boundaries. Therefore, a basic understanding of partial differential equations is essential before introducing the Chan–Vese model in Chapter 2.

Definition of a Partial Differential Equation

A partial differential equation is an equation involving an unknown function and its partial derivatives with respect to several independent variables.
Let $u = u(x_1, x_2, \ldots, x_n)$ be an unknown function defined on a domain $\Omega \subset \mathbb{R}^n$. A general PDE can be written as

$$F\left(x, u, \nabla u, \nabla^2 u, \ldots\right) = 0,$$

where

- $x = (x_1, \ldots, x_n)$ denotes the spatial variables,

- ∇u denotes the gradient of u,
- $\nabla^2 u$ denotes the Hessian matrix,
- F is a given function.

When time is involved, the unknown function is written as $u = u(x,t)$, and the PDE typically describes the temporal evolution of a physical quantity.

Diffusion and Nonlinear Diffusion

The simplest diffusion model is the heat equation $u_t = \Delta u$, i.e., in two dimensions, $u_t = u_{xx} + u_{yy}$. In image processing, the intensity plays the role of temperature: the equation smooths intensity variations and reduces noise, and its maximum principle guarantees that no new extrema are created. However, isotropic diffusion of this type also blurs edges indiscriminately.
More generally, one may consider $u_t = \operatorname{div}(D\nabla u)$ for a diffusion coefficient D; taking D constant recovers the heat equation, while letting D depend on image features yields *nonlinear diffusion*, which can preserve edges. A classical instance is

$$u_t = \operatorname{div}\big(g(|\nabla u|)\nabla u\big), \qquad g(s) = \frac{1}{1+\left(\frac{s}{K}\right)^2},$$

where g is a decreasing function [12]: diffusion remains strong where $|\nabla u| \approx 0$ (homogeneous regions) and is reduced where $|\nabla u| \gg 0$ (edges). This diffusivity-based logic underlies many gradient-based PDE methods; it is recalled here only for contrast, since, as shown in Section 1.2.3, the Chan–Vese functional studied in this dissertation follows a different, region-based logic and does not itself involve any diffusivity function $g(|\nabla u|)$.

Hamilton–Jacobi Equations

Another important family of PDEs is the Hamilton–Jacobi equations. Their general form is

$$u_t + H(\nabla u) = 0,$$

where H is called the Hamiltonian. These equations naturally arise in front propagation problems and geometric curve evolution.
A particularly important example is

$$u_t + F|\nabla u| = 0,$$

where F denotes a propagation speed. This equation will reappear in the level set formulation of moving interfaces, and, in Chapter 3, in the level set formulation of the Chan–Vese evolution.

Geometric Partial Differential Equations

Geometric PDEs describe the evolution of shapes and interfaces according to intrinsic geometric quantities.
Consider a family of curves $C(q,t)$. The evolution law can be written as

$$\frac{\partial C}{\partial t} = FN,$$

where

- F is the normal velocity,
- N is the unit normal vector.

If $F = \kappa$, one obtains the curvature flow

$$\frac{\partial C}{\partial t} = \kappa N.$$

Curvature flow smooths the contour by eliminating small irregularities.
Geometric PDEs of this type constitute the mathematical basis of active contour models in general, and their combination with region-based terms will yield the Chan–Vese evolution equation in Chapter 3.

PDEs in Image Segmentation

PDEs provide a natural framework for image segmentation.
They offer several advantages:

- mathematical rigor;
- robustness to noise;
- smooth contour evolution;
- incorporation of geometric constraints;
- compatibility with variational methods.

In active contour models, image boundaries are extracted by evolving a curve according to a PDE that combines image information – edge-based or region-based – with geometric regularization.
The resulting evolution equation generally contains two competing mechanisms:

1. a data-fitting force driving the contour toward the object boundary (based on the gradient in edge-based models, or on regional intensity statistics in region-based models such as Chan–Vese);

2. a smoothing force based on curvature.

The balance between these terms determines the quality of the final segmentation.

1.5 Level Set Method

The level set method is a powerful numerical and mathematical framework for representing and tracking evolving interfaces. It was introduced by Osher and Sethian [10] in the late 1980s and has since become a fundamental tool in computational geometry, fluid dynamics, and image processing.

In the context of image segmentation, the level set method provides a flexible way to represent evolving contours without requiring explicit parameterization. This allows natural handling of complex deformations such as splitting and merging of shapes, which are difficult to manage in classical parametric active contour models. This flexibility is precisely what makes the level set method the natural computational framework chosen by Chan and Vese [6] to implement their region-based segmentation model, as detailed in Chapters 2 and 3.

Implicit Representation of Curves

Let a moving curve be denoted by

$$C(t) \subset \mathbb{R}^2.$$

Instead of representing the curve explicitly, the level set method introduces an auxiliary scalar function

$$\phi : \Omega \times [0, +\infty) \to \mathbb{R},$$

such that the curve is represented as the zero level set:

$$C(t) = \{(x, y) \in \Omega \, : \, \phi(x, y, t) = 0\}.$$

In this formulation:

- $\phi(x, y, t) > 0$ inside the curve,
- $\phi(x, y, t) < 0$ outside the curve,
- $\phi(x, y, t) = 0$ defines the interface.

This implicit representation eliminates the need to explicitly track points along the curve. In the Chan–Vese model, the sign of ϕ is used directly to distinguish the two regions "inside" and "outside" the contour over which the regional intensity averages are computed.

Signed Distance Function

A common and convenient initialization of the level set function is the signed distance function defined by

$$\phi(x,y,0) = \begin{cases} d(x,y,C_0) & \text{inside the curve,} \\ -d(x,y,C_0) & \text{outside the curve,} \end{cases}$$

where $d(x,y,C_0)$ denotes the Euclidean distance from the point (x,y) to the initial contour C_0.
The signed distance function satisfies the important property

$$|\nabla\phi| = 1.$$

This property improves numerical stability and is often preserved through reinitialization procedures, a technique that will be used in Chapter 4 to stabilize the numerical scheme for the Chan–Vese equation.

Geometric Quantities in Level Set Form

One of the main advantages of the level set framework is that geometric quantities can be expressed directly in terms of ϕ.
The unit normal vector to the interface is given by

$$N = \frac{\nabla\phi}{|\nabla\phi|}.$$

The curvature of the level set is expressed as

$$\kappa = \operatorname{div}\left(\frac{\nabla\phi}{|\nabla\phi|}\right).$$

These expressions allow curvature-driven motions to be implemented without explicit curve parameterization, and they are exactly the terms used to write the length-regularization part of the Chan–Vese energy in level set form.

Evolution of the Level Set Function

Consider a moving curve $C(t)$ evolving with normal velocity F:

$$\frac{\partial C}{\partial t} = FN.$$

In the level set formulation, this evolution is translated into a PDE for ϕ:

$$\phi_t + F|\nabla\phi| = 0.$$

This equation is known as the level set equation or Hamilton–Jacobi type equation.
It describes how the implicit function ϕ evolves so that its zero level set moves with velocity F in the normal direction. In Chapter 3, F will be replaced by the specific region-based and curvature terms of the Chan–Vese model, yielding its full evolution equation.

Interpretation of the Level Set Equation

The term $F|\nabla\phi|$ ensures that motion occurs in the normal direction only.
The sign of F determines the direction of motion:

- $F > 0$ leads to outward propagation,
- $F < 0$ leads to inward contraction.

This formulation is independent of parameterization and remains valid even when the topology of the curve changes – a property that allows the Chan–Vese contour to naturally split or merge as it converges toward several objects at once.

Advantages of the Level Set Method

The level set method offers several important advantages:

- **Topological flexibility:** splitting and merging of curves occur naturally.
- **Numerical stability:** geometric quantities are computed from ϕ.
- **High dimensional extension:** the method extends easily to 3D surfaces.
- **Implicit formulation:** no explicit parameterization is required.

These properties make the level set method particularly well suited to region-based segmentation problems, and in particular to the Chan–Vese model.

In summary, this chapter has introduced the two complementary points of view on image segmentation – edge-based and region-based – together with the differential-geometric tools (curvature, curve evolution), the partial differential equations, and the level set framework. In the next chapter, these concepts will be combined with the Mumford–Shah energy functional, and complemented by the space $BV(\Omega)$ needed to give the length term a rigorous meaning, in order to derive the Chan–Vese model and its associated evolution equation.

Chapter 2

From the Mumford–Shah Model to the Chan–Vese Model

Chapter 1 introduced the two main paradigms of image segmentation – edge-based and region-based – together with the mathematical tools (curvature, PDEs, level sets) that are common to both. The purpose of the present chapter is to put these tools to work and to derive, step by step, the model that is the subject of this dissertation.

We proceed in four stages. We first present the Mumford–Shah functional [9], the general variational formulation of image segmentation from which almost every modern region-based model descends, and we explain why it is extremely difficult to minimize directly. We then introduce the functional-analytic framework – the space $BV(\Omega)$ of functions of bounded variation and the associated perimeter functional – that is needed to give the length-regularization term of Mumford–Shah a rigorous meaning even for irregular contours. We show how Chan and Vese [6], by restricting attention to the simplest non-trivial case of the Mumford–Shah functional – a two-phase, piecewise-constant image – obtained a model that is both mathematically tractable and numerically implementable, using exactly the region-based logic introduced in Section 1.2.3, and we use the BV framework to justify that a minimizer of this simplified energy actually exists. Finally, we compare the resulting "active contours without edges" model with the classical, edge-based active contour models (Snakes and Geodesic Active Contours), in order to make precise in what sense Chan–Vese constitutes a genuinely different approach to segmentation, and not merely a variant of the classical ones.

The full variational derivation of the Chan–Vese energy in level set form, together with its associated Euler–Lagrange equation, is postponed to Chapter 3, which opens with the calculus-of-variations tools required for that derivation; the present chapter focuses on the modeling ideas and on the passage from Mumford–Shah to Chan–Vese.

2.1 Image Segmentation as a Variational Problem

Before introducing the Mumford–Shah functional itself, it is useful to state precisely what an "ideal" segmentation should achieve, in variational terms.

Let $\Omega \subset \mathbb{R}^2$ be the image domain and let $u_0 : \Omega \to \mathbb{R}$ be the observed (possibly noisy) grayscale

image. Intuitively, we look for:

- a partition of Ω into a finite number of regions $\Omega_1, \dots, \Omega_n$, separated by a set of curves K, inside which the image intensity varies smoothly (or is constant), and
- across which the intensity is allowed to jump, these jumps corresponding to the actual object boundaries.

This is a joint problem: we must simultaneously recover a smoothed (or simplified) version u of the image *and* the discontinuity set K that separates the regions – K being precisely the segmentation we are looking for. Neither u nor K is known in advance, and each depends on the other: knowing K makes it easy to compute an optimal smooth u on each side, while knowing a good candidate u suggests where K should be placed (namely, along the discontinuities of u). This chicken-and-egg structure is what makes the problem genuinely variational, and it is the structure that Mumford and Shah proposed to formalize directly as the minimization of a single energy functional over the pair (u, K).

2.2 The Mumford–Shah Functional

Formulation

Mumford and Shah [9] proposed to formalize the joint smoothing/segmentation problem described above as the minimization, over all pairs (u, K), of the energy functional

$$F_{MS}(u, K) = \mu \int_{\Omega} (u - u_0)^2 \, dx + \int_{\Omega \setminus K} |\nabla u|^2 \, dx + \nu \, \mathrm{Length}(K), \tag{2.1}$$

where:

- $u_0 : \Omega \to \mathbb{R}$ is the given (observed) image;
- $u : \Omega \to \mathbb{R}$ is the unknown, piecewise-smooth approximation of u_0;
- $K \subset \Omega$ is the unknown discontinuity set (a union of curves), across which u is allowed to jump, and along which u is not required to be differentiable;
- $\mu, \nu > 0$ are fixed weighting parameters.

Each term of (2.1) has a clear interpretation:

- the **fidelity term** $\mu \int_{\Omega} (u - u_0)^2 dx$ penalizes the distance between the reconstructed image u and the observed data u_0, and prevents the trivial minimizer $u = 0$;

- the **smoothness term** $\int_{\Omega\setminus K}|\nabla u|^2\,dx$ penalizes variations of u *away from* K, forcing u to be smooth inside each region, exactly as in the diffusion-based regularization recalled in Section 1.4;

- the **regularization term** ν Length(K) penalizes the total length of the discontinuity set, and prevents K from being an arbitrarily long or fragmented curve that would simply interpolate the noise in u_0. Giving a rigorous meaning to this length term for an arbitrary, possibly irregular, set K is precisely the purpose of the space $BV(\Omega)$ introduced in Section 2.3 below.

A minimizing pair (u^*, K^*) thus realizes a compromise: u^* is as close as possible to u_0, is smooth away from K^*, and K^* is as short as possible while still separating the regions where a genuine jump in intensity is needed to keep the fidelity term small.

Why the Mumford–Shah Functional is Difficult to Minimize

Despite its conceptual clarity, functional (2.1) poses several serious mathematical and numerical difficulties.

The unknown K is a set, not a function. Unlike u, the discontinuity set K does not live in a linear (vector) space: it is an abstract closed subset of Ω, generally a union of curves of unknown number, length, and topology. There is, a priori, no natural notion of "small perturbation" of K that would allow one to differentiate the K-dependence of F_{MS} directly.

The problem is a free-discontinuity problem. The domain of integration $\Omega\setminus K$ in the second term itself depends on the unknown K, and the natural function space for u (functions that are smooth away from an unknown, lower-dimensional discontinuity set) does not coincide with any classical Sobolev space $W^{1,2}(\Omega)$. This class of problems, called *free-discontinuity problems*, requires substantially more advanced tools (the space $SBV(\Omega)$ of *special* functions of bounded variation, and the theory of rectifiable sets) than what is introduced in this dissertation, and a full existence theory is beyond its scope.

Direct numerical minimization is impractical. Even setting aside the theoretical difficulties, minimizing F_{MS} numerically over pairs (u, K) is very costly: K must be represented and updated (its length computed, its topology allowed to change), while u must be recomputed on the correspondingly updated domain $\Omega\setminus K$ at every step. Early numerical approaches (e.g. elliptic approximations such as that of Ambrosio and Tortorelli [1], which replace K by an auxiliary smooth "edge function" and let a small parameter tend to zero) alleviate some of these difficulties, but remain considerably more expensive than what is required for many practical segmentation tasks, particularly when only a two-region segmentation is sought.

These difficulties motivate looking for a *simplified* version of the Mumford–Shah functional, one

that is restrictive enough to be tractable both theoretically and numerically, yet expressive enough to solve a large and practically important class of segmentation problems. This is exactly the route taken by Chan and Vese, once the length term itself is placed on a rigorous footing – the task of the next section.

2.3 The Space of Functions of Bounded Variation

The energy functionals used in image segmentation, and in particular the Mumford–Shah functional above and the Chan–Vese functional derived from it, involve a regularization term that penalizes the length of the segmenting curve, or equivalently the perimeter of the region it encloses. To give a rigorous meaning to this term – one that remains well defined even for irregular, non-smooth contours, which are unavoidable when working with level set functions and noisy images – it is necessary to work in the space of functions of *bounded variation*, denoted $BV(\Omega)$, rather than restricting attention to smooth or continuously differentiable functions.

Definition

Let $\Omega \subset \mathbb{R}^2$ be a bounded open set. A function $u \in L^1(\Omega)$ is said to be of *bounded variation* if its distributional gradient Du is a (finite) Radon measure, i.e. if the quantity

$$\int_\Omega |Du| := \sup\left\{ \int_\Omega u \, \operatorname{div}(\xi)\, dx \,:\, \xi \in C_c^1(\Omega;\mathbb{R}^2),\ \|\xi\|_{L^\infty(\Omega)} \le 1 \right\} \tag{2.2}$$

is finite. This quantity is called the *total variation* of u on Ω, and the space

$$BV(\Omega) = \left\{ u \in L^1(\Omega) \,:\, \int_\Omega |Du| < +\infty \right\}$$

is called the space of functions of bounded variation on Ω. Equipped with the norm $\|u\|_{BV(\Omega)} = \|u\|_{L^1(\Omega)} + \int_\Omega |Du|$, $BV(\Omega)$ is a Banach space.

Why $BV(\Omega)$ Rather than $W^{1,1}(\Omega)$?

If u is smooth, definition (2.2) reduces, after an integration by parts, to the familiar expression

$$\int_\Omega |Du| = \int_\Omega |\nabla u|\, dx,$$

so that $BV(\Omega)$ can be seen as an extension of the Sobolev space $W^{1,1}(\Omega)$. The crucial difference is that $BV(\Omega)$ also contains functions with *jump discontinuities* across hypersurfaces (curves, in $\mathbb{R}^2$), such as the indicator function of a region with a rectifiable boundary – functions that do not belong to $W^{1,1}(\Omega)$. This is essential for image segmentation, since the object we ultimately wish to reconstruct or approximate is typically a piecewise-constant image, i.e. an image with genuine jumps across the segmenting curves, not a smooth function.

Sets of Finite Perimeter and the Perimeter Functional

Let $E \subset \Omega$ be a measurable set and let χ_E denote its indicator (characteristic) function. The set E is said to have *finite perimeter* in Ω if $\chi_E \in BV(\Omega)$, and its *perimeter* is then defined as the total variation of χ_E:

$$\mathrm{Per}(E;\Omega) := \int_\Omega |D\chi_E|. \tag{2.3}$$

When the boundary $\partial E \cap \Omega$ is a sufficiently regular curve C (e.g. piecewise C^1), this abstract definition coincides with the classical, intuitive notion of perimeter:

$$\mathrm{Per}(E;\Omega) = \mathrm{Length}(C) = \int_C ds,$$

where ds denotes the arc-length element. The advantage of definition (2.3) over the classical arc-length integral is that it remains meaningful for sets E whose boundary is not smooth, or whose boundary is only known implicitly through a level set function ϕ – exactly the situation encountered when $E = \{\phi > 0\}$ in the Chan–Vese model.

Perimeter in Level Set Form

For the Chan–Vese model, the region $E = \{\phi > 0\}$ is described implicitly through a level set function ϕ (see Section 1.5). Using the (regularized) Heaviside function H, characterized informally by $H(s) = 1$ for $s > 0$ and $H(s) = 0$ for $s < 0$, one can write $\chi_E = H(\phi)$, and the perimeter of E can then be expressed, formally, as

$$\mathrm{Per}(\{\phi > 0\};\Omega) = \int_\Omega |\nabla H(\phi)|\, dx = \int_\Omega \delta(\phi)\,|\nabla \phi|\, dx,$$

where $\delta = H'$ is the Dirac measure. This last expression is precisely the length-regularization term that will appear in the Chan–Vese energy functional below, and its regularized (smoothed) version, together with a discussion of H and δ, will be treated in detail in Chapter 4 when the model is discretized numerically.

Relevance to Segmentation

Working within $BV(\Omega)$ and using the perimeter functional (2.3) as a regularizer serves two purposes that are essential for a rigorous treatment of the Chan–Vese model:

- it allows piecewise-constant approximations of the image – the natural output of a region-based segmentation – to be legitimate competitors in the minimization problem, since indicator functions of sufficiently regular sets belong to $BV(\Omega)$;
- it penalizes irregular, oscillating, or overly fragmented contours (which would have infinite or very large perimeter), thereby ensuring that the minimization problem is well-posed and

that minimizing sequences admit a convergent subsequence, by the compactness properties of $BV(\Omega)$.

These two properties are what make it possible to prove existence of a minimizer for the Mumford–Shah and Chan–Vese functionals, as used below when the Chan–Vese energy is introduced.

2.4 The Chan–Vese Simplification: The Piecewise-Constant, Two-Phase Case

Restricting u to be Piecewise Constant

Chan and Vese [6] observed that a very large number of practical segmentation tasks do not actually require the full generality of Mumford–Shah: it is often enough to model the image as consisting of (at most) two regions of *constant* intensity, an object and a background, separated by a single curve C. This corresponds to setting, in (2.1), u to be piecewise constant rather than merely piecewise smooth:

$$u(x) = \begin{cases} c_1 & \text{if } x \text{ is inside } C, \\ c_2 & \text{if } x \text{ is outside } C, \end{cases}$$

for two unknown constants $c_1, c_2 \in \mathbb{R}$, with $K = C$ a single closed curve. This is sometimes called the *minimal partition problem*, or the *reduced (cartoon) Mumford–Shah problem*, and it is the simplest non-trivial special case of (2.1).

Under this restriction, u is constant on each side of C, so that $\nabla u = 0$ almost everywhere in $\Omega \setminus C$, and the smoothness term $\int_{\Omega\setminus K} |\nabla u|^2 dx$ in (2.1) vanishes identically. The Mumford–Shah functional therefore collapses, in this restricted setting, to

$$F(c_1, c_2, C) = \mu \,\mathrm{Length}(C) + \lambda_1 \int_{\mathrm{inside}(C)} (u_0 - c_1)^2 \, dx + \lambda_2 \int_{\mathrm{outside}(C)} (u_0 - c_2)^2 \, dx, \tag{2.4}$$

where $\mu, \lambda_1, \lambda_2 > 0$ are fixed weighting parameters, and $\mathrm{Length}(C)$ denotes the length – or, equivalently, by Section 2.3, the perimeter – of the curve C. Chan and Vese further introduced an optional area term $\nu \cdot \mathrm{Area}(\mathrm{inside}(C))$, with $\nu \in \mathbb{R}$, which can accelerate the evolution and control whether the curve tends to shrink or expand, without changing the nature of the model.

This is precisely the "Active Contours Without Edges" functional. Two remarks are worth making immediately, both of which connect directly back to Chapter 1:

- **Only two ingredients remain.** A regularization term (the length of C, exactly as in the general curvature-flow discussion of Section 1.3) and two *region-based* fidelity terms (the fit of a constant c_i to the data u_0 inside/outside C). There is no smoothness term, because u is piecewise constant by construction.

- **No image gradient appears anywhere in** (2.4)**.** The fit of each region is judged only through its mean intensity error, not through $|\nabla u_0|$. This is exactly the region-based logic anticipated in Section 1.2.3: the model asks "does a constant c_1 (resp. c_2) explain the data well inside (resp. outside) C?" rather than "is the gradient large along C?".

Optimal Constants for a Fixed Curve

A useful and instructive observation is that, for a *fixed* curve C, the optimal constants c_1, c_2 minimizing (2.4) can be computed explicitly. Since C is fixed, Length(C) is a constant with respect to c_1, c_2, and minimizing (2.4) over c_1 and c_2 separately amounts to two independent least-squares problems,

$$\min_{c_1} \int_{\text{inside}(C)} (u_0 - c_1)^2 dx, \qquad \min_{c_2} \int_{\text{outside}(C)} (u_0 - c_2)^2 dx,$$

whose (unique) solutions are simply the *average intensities* of u_0 inside and outside C:

$$c_1(C) = \text{average}\big(u_0\big) \text{ inside } C, \qquad c_2(C) = \text{average}\big(u_0\big) \text{ outside } C. \tag{2.5}$$

This is a first, elementary illustration of a general Euler–Lagrange-type principle (here applied to the two real variables c_1, c_2 rather than to a function or a curve, and made precise in general form at the start of Chapter 3): setting the partial derivatives of F with respect to c_1 and c_2 to zero yields exactly (2.5). In practice, this means that at every stage of the curve evolution described in Chapter 3, the constants c_1, c_2 are simply recomputed as the current mean intensities inside and outside the evolving contour – they are not independent unknowns to be evolved by a separate PDE.

Remaining Unknown: the Curve C

Once c_1, c_2 are expressed in terms of C via (2.5), the energy (2.4) becomes a functional of C alone,

$$F(C) = F\big(c_1(C), c_2(C), C\big),$$

and the segmentation problem reduces to finding the curve C that minimizes $F(C)$. This is now a curve-functional minimization problem, to which the general calculus-of-variations machinery recalled at the start of Chapter 3 applies directly: computing its first variation with respect to C and setting it to zero will yield an Euler–Lagrange equation for C, of the geometric form $\partial C/\partial t = V_N N$, whose normal velocity V_N combines the curvature of C (from the length term) with the two regional fidelity terms. This derivation, together with its reformulation using the level set function ϕ of Section 1.5, is carried out in full in Chapter 3.

Existence of a Minimizer

Before turning to the numerical treatment of the model, it is important to know that the minimization problem

$$\min_{c_1,c_2,C} F(c_1,c_2,C)$$

associated with the energy (2.4) is well posed, in the sense that a minimizer actually exists. We state this fact without giving a full proof, which would require additional technical tools beyond the scope of this dissertation; the interested reader is referred to the original article of Chan and Vese [6] and to the general existence theory for the Mumford–Shah functional developed in [9].
The existence argument follows the classical *direct method* of the calculus of variations, and relies precisely on the tools introduced in Section 2.3:

- one enlarges the class of admissible curves C to the class of sets $E \subset \Omega$ of *finite perimeter*, and rewrites the energy (2.4) as a functional of the pair (c_1, c_2, χ_E), using $\mathrm{Per}(E;\Omega)$ in place of $\mathrm{Length}(C)$;
- one considers a minimizing sequence $(c_1^n, c_2^n, \chi_{E_n})$, i.e. a sequence along which F decreases toward its infimum;
- because $\mathrm{Length}(C)$ (equivalently $\mathrm{Per}(E_n;\Omega)$) is controlled along this sequence, the compactness theorem for sets of finite perimeter in $BV(\Omega)$ guarantees that a subsequence of χ_{E_n} converges (in the L^1 sense) to the indicator function χ_{E^*} of some set E^* of finite perimeter, while c_1^n, c_2^n, being bounded, converge (up to a subsequence) to some c_1^*, c_2^*;
- finally, one checks that F is *lower semi-continuous* with respect to this mode of convergence, which allows one to pass to the limit in the inequality $F(c_1^n, c_2^n, E_n) \leq F(c_1, c_2, E)$ for every competitor (c_1, c_2, E), and to conclude that (c_1^*, c_2^*, E^*) is indeed a minimizer.

This argument is exactly why $BV(\Omega)$ and the perimeter functional, rather than a smooth parametrization of C, were introduced in Section 2.3: they provide a class of admissible curves that is simultaneously rich enough to contain any reasonable segmenting curve, and compact enough for a minimizing sequence to have a convergent subsequence – the two properties on which the direct method rests.

Scope and Limitations of the Two-Phase Model

The piecewise-constant, two-phase restriction is what makes the model both theoretically tractable, in the sense made precise above, and numerically implementable with a single level set function,

as will be shown in Chapters 3 and 4. This tractability, however, comes at a price: the model assumes that the image can be reasonably well approximated by exactly two regions of (nearly) constant intensity. Images containing more than two significantly different intensity levels, or genuine texture (rather than piecewise-constant regions), are not well captured by (2.4) and require extensions of the model. Chan and Vese themselves proposed such a *multiphase* extension, using several level set functions simultaneously to represent more than two regions [16]; this extension lies outside the scope of the present dissertation, which restricts attention throughout to the original two-phase model, but it is mentioned again briefly in the general conclusion as a natural direction for future work.

Summary: From Mumford–Shah to Chan–Vese

Table 2.1 summarizes, term by term, how the restriction to a piecewise-constant, two-phase image transforms the general Mumford–Shah functional (2.1) into the Chan–Vese energy (2.4).

Ingredient	**Mumford–Shah**	**Chan–Vese (two-phase)**
Unknown image u	piecewise *smooth*, arbitrary number of regions	piecewise *constant*, exactly two regions (c_1, c_2)
Smoothness term	$\int_{\Omega\setminus K} \lvert\nabla u\rvert^2 dx$	vanishes identically ($\nabla u = 0$ on each region)
Fidelity term	$\mu \int_\Omega (u-u_0)^2 dx$	$\lambda_1 \int_{\text{in}} (u_0-c_1)^2 dx + \lambda_2 \int_{\text{out}} (u_0-c_2)^2 dx$
Regularization	ν Length(K), K of arbitrary topology	μ Length(C), C a single (or few) closed curve(s)
Unknowns	the pair (u, K)	the triple (c_1, c_2, C)
Dependence on $\lvert\nabla u_0\rvert$	none (region-based, general case)	none (region-based, two-phase case)
Numerical difficulty	high (free-discontinuity problem)	moderate (single level set function, Chapters 3–4)

Table 2.1: From the general Mumford–Shah functional to the Chan–Vese two-phase energy.

2.5 Comparison with Classical Active Contour Models

To situate the Chan–Vese model precisely within the landscape of active contour methods, and to justify the claims made informally in Chapter 1 (Section 1.2.3), we now compare it with the two classical families of edge-based active contours: parametric Snakes and Geodesic Active Contours.

Snakes (Parametric Active Contours)

The original active contour model, introduced by Kass, Witkin and Terzopoulos [7] under the name *Snakes*, represents the contour explicitly as a parametrized curve $C(p) = (x(p), y(p))$, $p \in [0,1]$, and defines its energy as

$$E_{\text{snake}}(C) = \underbrace{\int_0^1 \left(\alpha\,|C'(p)|^2 + \beta\,|C''(p)|^2\right) dp}_{\text{internal energy}} - \underbrace{\int_0^1 |\nabla u_0(C(p))|^2\, dp}_{\text{external (image) energy}}, \tag{2.6}$$

where $\alpha, \beta > 0$ control the curve's elasticity and rigidity. The internal energy keeps the curve smooth (penalizing stretching and bending), while the external energy attracts the curve toward locations of large image gradient, i.e. toward edges.

This model has two well-known limitations, both of which motivated later developments:

- **Explicit parametrization.** Because C is represented by a finite set of parametrized points, the model cannot naturally handle changes of topology (splitting into several curves, or merging of curves) without additional, non-variational bookkeeping.
- **Strong dependence on initialization and on edge strength.** The external energy in (2.6) depends on $|\nabla u_0|$ evaluated only along the current curve, so the snake must be initialized close to the true boundary, and it may fail entirely if the boundary is weak, blurred, or corrupted by noise – exactly the gradient-based limitation already discussed in Section 1.2.2.

Geodesic Active Contours

Geodesic Active Contours (GAC), introduced by Caselles, Kimmel and Sapiro [5], reformulate the snake model geometrically and implicitly, using the level set framework of Section 1.5. The curve is no longer parametrized explicitly; instead, it is represented as the zero level set of a function ϕ, and it evolves according to

$$\frac{\partial \phi}{\partial t} = g(|\nabla u_0|)\big(\kappa + \nu\big)|\nabla \phi| + \nabla g(|\nabla u_0|) \cdot \nabla \phi, \tag{2.7}$$

where $\kappa = \operatorname{div}(\nabla\phi/|\nabla\phi|)$ is the curvature introduced in Section 1.3, ν is a constant driving force, and $g(|\nabla u_0|)$ is the edge-stopping (edge indicator) function of Section 1.2.2, typically $g(s) = 1/(1 + |\nabla u_{0,\sigma}|^2)$.

The passage to the level set framework solves the topology problem of Snakes: because ϕ is defined on the whole domain Ω rather than only along the curve, its zero level set can split or merge freely as ϕ evolves, exactly as discussed in Section 1.5. This is a genuine improvement over parametric Snakes.

However, GAC remains, at its core, an **edge-based** model: the factor $g(|\nabla u_0|)$ in (2.7) is precisely the edge indicator function discussed critically in Section 1.2.2, and it plays exactly the same role as the external energy of Snakes – it slows the evolution of ϕ near large gradients, and lets it pass through freely (with speed governed only by $\kappa + \nu$) wherever $|\nabla u_0|$ is small. As a consequence, GAC inherits the same fundamental limitation identified in Section 1.2.2: if the object boundary is weak, blurred, or lost in noise, $g(|\nabla u_0|)$ never becomes small enough along that boundary, and the evolving contour fails to stop there, typically leaking through the gap.

Where Chan–Vese Departs from Both Models

The Chan–Vese model shares with GAC the implicit, level-set-based representation of the curve – and therefore the same topological flexibility – but departs radically from both Snakes and GAC in the nature of the force driving the contour. Comparing (2.4) with (2.6) and (2.7) makes the distinction precise:

- **No edge indicator function.** Neither $g(|\nabla u_0|)$ nor any other function of the image gradient appears in the Chan–Vese energy (2.4). The curve is not stopped by high gradients; it is attracted toward the boundary that best separates two regions of (nearly) constant intensity, as measured by the fidelity terms of (2.4).

- **Global, not local, data term.** The fidelity terms of (2.4) are integrals over the *whole* inside and outside of C, not evaluations restricted to a neighborhood of the curve as in (2.6) and (2.7). The curve therefore "feels" the global statistics of each region, which is what allows it to detect boundaries with very weak or even vanishing local gradient, as long as the two regions differ, on average, in intensity.

- **Reduced sensitivity to initialization.** Because the driving force is global rather than restricted to a neighborhood of C, the Chan–Vese contour can, in principle, be initialized far from the true boundary (even as a small circle placed anywhere in the image) and still converge toward it, whereas Snakes in particular require an initialization close to the target boundary.

These properties come, as already noted above, at the cost of the piecewise-constant assumption: Chan–Vese is the natural choice when the objects and background are reasonably homogeneous in intensity, whereas GAC (and, more generally, edge-based methods) may still be preferable when boundaries are sharp and well contrasted but regions are strongly textured or non-uniform in intensity.

Conclusion of the Chapter

This chapter has traced the path from the general Mumford–Shah functional (2.1) to the specific two-phase, piecewise-constant energy (2.4) proposed by Chan and Vese, introducing along the way the space $BV(\Omega)$ and the perimeter functional needed to give the length-regularization term a rigorous meaning and to establish existence of a minimizer. This simplification was shown, term by term, to preserve exactly the region-based, gradient-free philosophy introduced in Chapter 1, while discarding the smoothness term that makes the general Mumford–Shah problem so difficult to handle. The comparison with Snakes and Geodesic Active Contours has clarified in what precise sense Chan–Vese is not merely another level-set active contour model, but a model driven by fundamentally different, global and region-based, information.

What remains to be done is the rigorous variational derivation of the Chan–Vese model: rewriting the energy (2.4) in level set form using the function ϕ and the regularized Heaviside and Dirac functions, computing its first variation, and deriving the associated Euler–Lagrange equation and gradient-descent evolution PDE for ϕ. This is the object of Chapter 3, which opens with the calculus-of-variations tools that make this passage rigorous.

Chapter 3

Mathematical Formulation of the Chan–Vese Model

Chapter 2 introduced the Chan–Vese model [6] as a region-based variational model derived from a piecewise-constant approximation of the Mumford–Shah functional: it seeks a contour that separates an image into two phases according to their region statistics rather than relying explicitly on image gradients.

The purpose of the present chapter is to formulate this model rigorously in the level set framework. We begin by recalling the general calculus-of-variations tools – functionals, first variation, and the Euler–Lagrange equation – that make rigorous the passage from an energy functional to a partial differential equation; these tools are used repeatedly throughout the rest of the chapter. We then express the Chan–Vese energy functional using the level set function of Section 1.5 and introduce the Heaviside and Dirac functions; derive the optimal values of the constants c_1 and c_2; compute the first variation of the energy and obtain the associated Euler–Lagrange equation together with its natural boundary condition; derive the gradient-flow evolution equation and its normal-velocity interpretation; and conclude with a detailed discussion of the role of every parameter of the model. Equation (3.36) below is the equation that will be discretized and implemented numerically in Chapter 4.

3.1 Elements of the Calculus of Variations

The Chan–Vese model, like the Mumford–Shah functional from which it is derived (Chapter 2), is not postulated directly as a partial differential equation: it is obtained as the minimizer of an energy functional. We therefore recall here the basic notions of the calculus of variations that make this passage, from an energy to a PDE, rigorous, before applying them to the Chan–Vese energy in the sections that follow.

Functionals

Let $\Omega \subset \mathbb{R}^2$ be a bounded open set and let $\mathscr{A}$ be a set of admissible functions (or admissible curves) defined on Ω. A *functional* is a mapping

$$E : \mathscr{A} \longrightarrow \mathbb{R}, \qquad u \longmapsto E(u),$$

that associates a real number to each admissible function u, rather than a number to each point as an ordinary function does. In image processing, $E(u)$ typically measures how well a candidate function u (an image, a contour, or a level set function ϕ) explains the observed data, penalized by some regularity requirement.

A generic example, which covers most of the functionals encountered in this dissertation, has the form

$$E(u) = \int_{\Omega} L\big(x, u(x), \nabla u(x)\big)\, dx, \tag{3.1}$$

where L is a given function, called the *Lagrangian* of the problem, depending on the position x, the value $u(x)$, and the gradient $\nabla u(x)$.

First Variation and the Euler–Lagrange Equation

The central problem of the calculus of variations is to characterize the minimizers of a functional such as (3.1). The key tool for this is the notion of *first variation*.

Let u be a (smooth) minimizer of E over $\mathscr{A}$, and let h be an admissible perturbation (a test function, vanishing on $\partial\Omega$ if u is fixed on the boundary). For $\varepsilon \in \mathbb{R}$ small, consider the real-valued function

$$j(\varepsilon) = E(u + \varepsilon h).$$

Since u minimizes E, the function j attains its minimum at $\varepsilon = 0$, and therefore

$$j'(0) = 0.$$

The quantity $j'(0)$, denoted $\delta E(u; h)$, is called the *first variation* of E at u in the direction h. Computing it for the functional (3.1) using the chain rule gives

$$\delta E(u; h) = \int_{\Omega} (L_u\, h + L_{\nabla u} \cdot \nabla h)\, dx,$$

where L_u and $L_{\nabla u}$ denote the partial derivatives of L with respect to u and ∇u respectively. Integrating the second term by parts (and using $h = 0$ on $\partial\Omega$ to discard the boundary term) yields

$$\delta E(u; h) = \int_{\Omega} \big(L_u - \operatorname{div}\big(L_{\nabla u}\big)\big)\, h\, dx.$$

Since this must vanish for every admissible perturbation h, the fundamental lemma of the calculus of variations implies that the bracketed term itself must vanish. This gives the *Euler–Lagrange equation* associated with E:

$$L_u - \operatorname{div}\big(L_{\nabla u}\big) = 0 \qquad \text{in } \Omega, \tag{3.2}$$

together with natural boundary conditions on $\partial\Omega$ arising from the discarded boundary terms. Equation (3.2) is precisely the variational derivative of E, denoted $\dfrac{\delta E}{\delta u} = 0$: any smooth minimizer of E must satisfy it.

Gradient Descent and Artificial Time

Solving the Euler–Lagrange equation (3.2) directly is often difficult. A standard strategy is to embed the static equation into a dynamic process by introducing an artificial time variable t and letting u evolve by the associated *gradient descent* (or gradient flow):

$$\frac{\partial u}{\partial t} = -\frac{\delta E}{\delta u} = \operatorname{div}\big(L_{\nabla u}\big) - L_u. \tag{3.3}$$

Starting from a suitable initial condition $u(x,0) = u_0(x)$, the solution $u(x,t)$ decreases the energy E at every time step,

$$\frac{d}{dt} E\big(u(\cdot,t)\big) = -\int_\Omega \left|\frac{\delta E}{\delta u}\right|^2 dx \le 0,$$

and converges, under suitable conditions, to a steady state u_∞ satisfying $\dfrac{\delta E}{\delta u}(u_\infty) = 0$, i.e. a critical point of E. This is exactly the strategy that will be used later in this chapter to pass from the Chan–Vese energy functional to its evolution equation, and in Chapter 4 to discretize and solve that equation numerically.

Case of a Functional Depending on a Curve

When the unknown is a curve C rather than a function u – as in active contour models – the same principle applies, with h now representing an admissible perturbation of the curve (a vector field along C). The first variation of a curve functional $E(C)$ again takes the form of an integral along C of a quantity multiplying the perturbation, and setting this quantity to zero yields the Euler–Lagrange equation of the curve functional. Written as a gradient descent, this equation takes the geometric form already introduced in Section 1.3,

$$\frac{\partial C}{\partial t} = -\nabla E(C) = V_N N,$$

where V_N is precisely the Euler–Lagrange expression associated with $E(C)$, projected onto the normal direction N. This is the variational mechanism that produces the curvature term κN whenever

$E(C)$ contains the length (perimeter) of C – made rigorous, for irregular C, by the $BV(\Omega)$ framework of Section 2.3 – and it is the same mechanism that will produce, later in this chapter, the full Chan–Vese evolution equation once the length term is combined with the region-based data-fitting terms.

3.2 Level Set Representation of the Chan–Vese Energy

The Level Set Function and the Two Regions

Let $\Omega \subset \mathbb{R}^2$ denote the image domain and let $u_0 : \Omega \to \mathbb{R}$ be the given image. Following Section 1.5, the evolving contour C is represented implicitly by a level set function $\phi : \Omega \to \mathbb{R}$ such that

$$C = \{x \in \Omega : \phi(x) = 0\}. \tag{3.4}$$

Throughout this chapter we adopt the sign convention

$$\begin{cases} \phi(x) > 0, & x \text{ is inside } C, \\ \phi(x) < 0, & x \text{ is outside } C, \\ \phi(x) = 0, & x \in C, \end{cases} \tag{3.5}$$

so that the inside and outside regions are represented implicitly by

$$\Omega_{\text{in}} = \{x \in \Omega : \phi(x) > 0\}, \qquad \Omega_{\text{out}} = \{x \in \Omega : \phi(x) < 0\}. \tag{3.6}$$

This single function ϕ now encodes, implicitly, everything that was previously described through the explicit curve C: its interior, its exterior, and the curve itself, with no explicit parametrization required and no restriction on its topology, exactly as discussed in Section 1.5. The contour can deform, split, merge, or disappear during the evolution of ϕ.

Heaviside and Dirac Delta Functions

Let H denote the Heaviside function,

$$H(z) = \begin{cases} 1, & z > 0, \\ 0, & z < 0. \end{cases} \tag{3.7}$$

Under the sign convention adopted above, the characteristic function of the inside region is represented by

$$\chi_{\Omega_{\text{in}}}(x) = H(\phi(x)), \qquad \chi_{\Omega_{\text{out}}}(x) = 1 - H(\phi(x)). \tag{3.8}$$

The Dirac delta distribution is formally related to the Heaviside function by $\delta(z) = H'(z)$. Using the coarea formula, as already anticipated in Section 2.3, the length of C can be written in level set form as

$$\text{Length}(C) = \int_\Omega \delta(\phi)|\nabla\phi|\, dx, \tag{3.9}$$

and the area of the inside region as

$$\text{Area}(\Omega_{\text{in}}) = \int_{\Omega} H(\phi)\, dx. \tag{3.10}$$

The Energy in Level Set Form

Substituting these expressions into the parametric Chan–Vese energy of Chapter 2, the functional becomes

$$\begin{aligned} F(c_1, c_2, \phi) = \mu \int_{\Omega} \delta(\phi)|\nabla\phi|\, dx + \nu \int_{\Omega} H(\phi)\, dx \\ + \lambda_1 \int_{\Omega} (u_0 - c_1)^2 H(\phi)\, dx + \lambda_2 \int_{\Omega} (u_0 - c_2)^2 \big(1 - H(\phi)\big)\, dx, \end{aligned} \tag{3.11}$$

where $\mu, \nu, \lambda_1, \lambda_2$ are the model parameters: the first term penalizes the length of the contour, the second controls the area enclosed by the contour, and the last two terms measure the discrepancy between the image and the representative values c_1, c_2 in the two phases. The parameter ν is optional, in the sense that the original Chan–Vese formulation also considers the case $\nu = 0$; when $\nu \neq 0$, the area term introduces a bias toward expansion or contraction of the interior region.

Every region-dependent quantity of Chapter 2 has thus been rewritten as an integral over the fixed domain Ω, with the unknown region encoded entirely through the sign of ϕ inside the (now fixed) integrands $H(\phi)$ and $\delta(\phi)$ – exactly what makes the level set framework so convenient here, since the domain of integration no longer moves with the unknown, unlike $\Omega \setminus K$ in the original Mumford–Shah functional (Section 2.2).

The corresponding minimization problem is

$$\min_{c_1, c_2, \phi} F(c_1, c_2, \phi). \tag{3.12}$$

For a fixed ϕ, the values c_1 and c_2 are obtained by minimizing the fidelity terms; this is considered next.

3.3 Optimal Values of the Region Constants

For a fixed level set function ϕ, the optimal values of c_1 and c_2 are obtained by differentiating the energy with respect to these constants, exactly as in Section 2.4. Considering c_1 first,

$$\frac{\partial F}{\partial c_1} = 2\lambda_1 \int_{\Omega} (c_1 - u_0)\, H(\phi)\, dx. \tag{3.13}$$

Setting this derivative equal to zero gives $\int_{\Omega} (c_1 - u_0)\, H(\phi)\, dx = 0$, so that, provided $\int_{\Omega} H(\phi)\, dx > 0$,

$$c_1(\phi) = \frac{\displaystyle\int_{\Omega} u_0\, H(\phi)\, dx}{\displaystyle\int_{\Omega} H(\phi)\, dx}. \tag{3.14}$$

Similarly, differentiation with respect to c_2 gives $\dfrac{\partial F}{\partial c_2} = 2\lambda_2 \displaystyle\int_\Omega (c_2 - u_0)\big(1 - H(\phi)\big)\,dx$, so that, provided $\int_\Omega \big(1 - H(\phi)\big)\,dx > 0$,

$$c_2(\phi) = \frac{\displaystyle\int_\Omega u_0\,\big(1 - H(\phi)\big)\,dx}{\displaystyle\int_\Omega \big(1 - H(\phi)\big)\,dx}. \tag{3.15}$$

Thus c_1 and c_2 are simply the mean image intensities in the inside and outside phases, respectively – the level set counterpart of (2.5) from Chapter 2. As before, $c_1(\phi)$ and $c_2(\phi)$ are not independent unknowns: these expressions are recomputed at every stage of the numerical evolution described in Chapter 4, allowing the region statistics to adapt automatically as the contour changes.

3.4 First Variation and the Euler–Lagrange Equation

Setting Up the Variation

We now derive the evolution equation associated with the Chan–Vese functional, following the general method of Section 3.1. Consider a perturbation of the level set function of the form $\phi_\eta = \phi + \eta h$, where h is a sufficiently smooth test function and η is a small real parameter, and define

$$j(\eta) = F(c_1, c_2, \phi + \eta h). \tag{3.16}$$

The first variation of the functional is $\delta F(\phi; h) = \dfrac{d}{d\eta} j(\eta)\Big|_{\eta=0}$. We compute the contributions of the four terms of the energy separately.

Variation of the Length Term

Consider $F_{\text{length}}(\phi) = \mu \displaystyle\int_\Omega \delta(\phi)|\nabla\phi|\,dx$. For $\phi_\eta = \phi + \eta h$,

$$\frac{d}{d\eta}\Big|_{\eta=0} F_{\text{length}}(\phi_\eta) = \mu \int_\Omega \left[\delta'(\phi)|\nabla\phi|\,h + \delta(\phi)\frac{\nabla\phi\cdot\nabla h}{|\nabla\phi|}\right] dx. \tag{3.17}$$

To simplify the second term, integration by parts, exactly as in the general derivation of Section 3.1, gives

$$\int_\Omega \delta(\phi)\frac{\nabla\phi\cdot\nabla h}{|\nabla\phi|}\,dx = -\int_\Omega h\,\operatorname{div}\!\left(\delta(\phi)\frac{\nabla\phi}{|\nabla\phi|}\right) dx + \int_{\partial\Omega} h\,\delta(\phi)\frac{\partial_n\phi}{|\nabla\phi|}\,ds, \tag{3.18}$$

where $\partial_n\phi = \nabla\phi\cdot n$ and n is the outward unit normal to $\partial\Omega$. Using the identity

$$\operatorname{div}\!\left(\delta(\phi)\frac{\nabla\phi}{|\nabla\phi|}\right) = \delta'(\phi)|\nabla\phi| + \delta(\phi)\operatorname{div}\!\left(\frac{\nabla\phi}{|\nabla\phi|}\right), \tag{3.19}$$

the terms containing $\delta'(\phi)|\nabla\phi|$ in (3.17) and (3.18) cancel exactly. Introducing the curvature of the zero level set, as in Section 1.3 and Section 1.5,

$$\kappa = \operatorname{div}\!\left(\frac{\nabla\phi}{|\nabla\phi|}\right), \tag{3.20}$$

the first variation of the length term reduces to

$$\delta F_{\text{length}}(\phi; h) = -\mu \int_\Omega \delta(\phi)\,\kappa\,h\,dx + \mu \int_{\partial\Omega} \delta(\phi)\frac{\partial_n\phi}{|\nabla\phi|}\,h\,ds. \tag{3.21}$$

Variation of the Area Term

The area contribution is $F_{\text{area}}(\phi) = \nu \int_\Omega H(\phi)\,dx$. Since $H'(\phi) = \delta(\phi)$, its first variation is immediate:

$$\delta F_{\text{area}}(\phi; h) = \nu \int_\Omega \delta(\phi)\, h\, dx. \tag{3.22}$$

Variation of the Fidelity Terms

For the inside region, $F_1(\phi) = \lambda_1 \int_\Omega (u_0 - c_1)^2 H(\phi)\,dx$, so that

$$\delta F_1(\phi; h) = \lambda_1 \int_\Omega (u_0 - c_1)^2\, \delta(\phi)\, h\, dx. \tag{3.23}$$

Similarly, for the outside region, $F_2(\phi) = \lambda_2 \int_\Omega (u_0 - c_2)^2 \big(1 - H(\phi)\big)\,dx$, so that

$$\delta F_2(\phi; h) = -\lambda_2 \int_\Omega (u_0 - c_2)^2\, \delta(\phi)\, h\, dx. \tag{3.24}$$

(As in Section 2.4, the dependence of c_1, c_2 themselves on ϕ contributes no extra term here, by the same envelope-type argument used in Chapter 2.)

Total First Variation and the Euler–Lagrange Equation

Collecting the four contributions (3.21)–(3.24) yields

$$\begin{aligned} \delta F(\phi; h) = &\int_\Omega \delta(\phi)\Big[-\mu\kappa + \nu + \lambda_1 (u_0 - c_1)^2 - \lambda_2 (u_0 - c_2)^2 \Big] h\, dx \\ &+ \mu \int_{\partial\Omega} \delta(\phi) \frac{\partial_n \phi}{|\nabla \phi|}\, h\, ds. \end{aligned} \tag{3.25}$$

Requiring $\delta F(\phi; h) = 0$ for every admissible h, first supported in the interior of Ω and then supported near $\partial\Omega$, exactly as in the abstract derivation (3.2) of Section 3.1, shows that, under a boundary condition making the boundary integral vanish, the functional derivative is formally given by

$$\frac{\delta F}{\delta \phi} = \delta(\phi) \left[-\mu\kappa + \nu + \lambda_1 (u_0 - c_1)^2 - \lambda_2 (u_0 - c_2)^2 \right], \tag{3.26}$$

and that a stationary point of the energy satisfies the Euler–Lagrange equation

$$\delta(\phi) \left[-\mu\kappa + \nu + \lambda_1 (u_0 - c_1)^2 - \lambda_2 (u_0 - c_2)^2 \right] = 0 \qquad \text{in } \Omega. \tag{3.27}$$

It is important to distinguish this stationary condition from the evolution equation, obtained next by introducing a gradient flow that decreases the energy.

3.5 The Chan–Vese Evolution Equation

Gradient Descent in Artificial Time

As explained in general terms in Section 3.1, the gradient-descent evolution associated with the energy is defined by embedding (3.27) into an artificial time variable t, $\phi = \phi(x,t)$, through

$$\frac{\partial \phi}{\partial t} = -\frac{\delta F}{\delta \phi}. \tag{3.28}$$

Substituting (3.26) into (3.28) gives the Chan–Vese *evolution equation*:

$$\frac{\partial \phi}{\partial t} = \delta(\phi)\left[\mu \operatorname{div}\left(\frac{\nabla \phi}{|\nabla \phi|}\right) - \nu - \lambda_1 (u_0 - c_1)^2 + \lambda_2 (u_0 - c_2)^2\right] \qquad \text{in } \Omega \times (0, +\infty), \tag{3.29}$$

with $c_1 = c_1(\phi(\cdot,t))$, $c_2 = c_2(\phi(\cdot,t))$ recomputed from (3.14)–(3.15) at every time t, and an initial condition $\phi(x,0) = \phi_0(x)$ specifying the initial contour (see Chapter 4). The three terms inside the brackets have distinct roles: the curvature term $\mu \operatorname{div}(\nabla\phi/|\nabla\phi|)$ regularizes the contour and tends to reduce its length, exactly as the curvature flow of Section 1.3; the area term $-\nu$ introduces an expansion or contraction bias; and the two fidelity terms drive the contour toward a partition for which the inside and outside regions are well represented by the constants c_1 and c_2.

Normal Velocity Interpretation

Equation (3.29) can also be interpreted through the normal velocity of the evolving zero level set, in the spirit of the general level set evolution $\phi_t + F|\nabla\phi| = 0$ of Section 1.5. Let $C(t)$ be the evolving zero level set, defined by $\phi(x(t),t) = 0$. Differentiating with respect to time gives

$$\nabla\phi \cdot \frac{dx}{dt} + \phi_t = 0. \tag{3.30}$$

If $n = \nabla\phi/|\nabla\phi|$ denotes the unit normal, the normal velocity of $C(t)$ satisfies

$$V_N = -\frac{\phi_t}{|\nabla\phi|}. \tag{3.31}$$

The geometric motion of the zero level set is therefore governed by the balance between curvature, area, and fidelity forces described above. At a formal steady state ($\phi_t = 0$), the bracketed expression in (3.29) vanishes on the zero level set, yielding

$$\mu\kappa = \nu + \lambda_1 (u_0 - c_1)^2 - \lambda_2 (u_0 - c_2)^2, \tag{3.32}$$

a balance between contour regularization, area bias, and region-based data fidelity – the discrete, numerical counterpart of which will be used in Chapter 4 to define a stopping criterion for the iterative scheme. No image gradient $|\nabla u_0|$ and no edge-stopping function $g(|\nabla u_0|)$ appear anywhere in (3.29): the entire driving force is the pair of regional averages c_1, c_2, exactly as promised by the region-based philosophy of Section 1.2.3 and confirmed by the comparison with Snakes and Geodesic Active Contours in Section 2.5.

3.6 Regularization of the Heaviside and Dirac Functions

The Heaviside function and the Dirac delta distribution are discontinuous, respectively singular, and are therefore not directly suitable for numerical implementation, as already anticipated in Section 2.3. Following the regularization commonly used in numerical implementations of the Chan–Vese model [6], we define

$$H_\varepsilon(z) = \frac{1}{2}\left(1 + \frac{2}{\pi}\arctan\left(\frac{z}{\varepsilon}\right)\right), \qquad \varepsilon > 0, \tag{3.33}$$

whose derivative gives the corresponding regularized Dirac function

$$\delta_\varepsilon(z) = H'_\varepsilon(z) = \frac{1}{\pi}\cdot\frac{\varepsilon}{\varepsilon^2 + z^2}. \tag{3.34}$$

The parameter ε controls the width of the transition between the two phases: as ε decreases, the regularized Heaviside function approaches the discontinuous Heaviside function away from the transition region. Two properties of this regularization are essential for what follows, and both will be used again in Chapter 4: δ_ε is smooth and strictly positive everywhere, so the regularized model assigns a smooth transition of width $O(\varepsilon)$ across $\{\phi = 0\}$; and because $\delta_\varepsilon(z) > 0$ everywhere – not only near $z = 0$, as a compactly supported mollification of δ would give – the evolution equation below is influenced by the whole domain Ω at every step, not only by a narrow band around the current curve. This is what allows the Chan–Vese contour, in practice, to detect boundaries even when it is initialized far from them, as already anticipated in Section 2.5.

Using these regularizations, the numerical energy is

$$\begin{aligned} F_\varepsilon(c_1, c_2, \phi) = \mu \int_\Omega \delta_\varepsilon(\phi)|\nabla\phi|\,dx + \nu \int_\Omega H_\varepsilon(\phi)\,dx \\ + \lambda_1 \int_\Omega (u_0 - c_1)^2 H_\varepsilon(\phi)\,dx + \lambda_2 \int_\Omega (u_0 - c_2)^2\big(1 - H_\varepsilon(\phi)\big)\,dx, \end{aligned} \tag{3.35}$$

and the corresponding regularized Chan–Vese evolution equation, which is the equation actually discretized and solved numerically in Chapter 4, is

$$\frac{\partial\phi}{\partial t} = \delta_\varepsilon(\phi)\left[\mu\,\mathrm{div}\left(\frac{\nabla\phi}{|\nabla\phi|}\right) - \nu - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2\right]. \tag{3.36}$$

In a numerical implementation, the singularity associated with $|\nabla\phi|$ in the curvature term is also avoided by introducing a small positive regularization parameter when evaluating it. For example, one may use

$$|\nabla\phi|_\eta = \sqrt{|\nabla\phi|^2 + \eta^2}, \qquad \eta > 0, \tag{3.37}$$

and approximate the curvature by

$$\kappa_\eta = \mathrm{div}\left(\frac{\nabla\phi}{\sqrt{|\nabla\phi|^2 + \eta^2}}\right). \tag{3.38}$$

This additional regularization is a purely numerical device, used to avoid division by zero when the gradient of the level set function becomes very small; it does not change the conceptual structure of the Chan–Vese model, and will be revisited in Chapter 4 together with the finite-difference discretization of (3.36).

3.7 Boundary Condition

The integration by parts performed in the first variation, equation (3.18), produces the boundary term $\mu \int_{\partial\Omega} \delta(\phi) \frac{\partial_n \phi}{|\nabla \phi|} h \, ds$. A natural boundary condition associated with the variational formulation is therefore

$$\delta(\phi) \frac{\partial_n \phi}{|\nabla \phi|} = 0 \qquad \text{on } \partial\Omega. \tag{3.39}$$

In numerical implementations, this condition is commonly enforced through the stronger, homogeneous Neumann condition

$$\frac{\partial \phi}{\partial n} = 0 \qquad \text{on } \partial\Omega, \tag{3.40}$$

which prevents artificial flux of the level set function through the image boundary and is convenient for the finite-difference discretization used in Chapter 4.

3.8 Role of the Model Parameters

The behavior of the Chan–Vese model depends strongly on the parameters $\mu, \nu, \lambda_1, \lambda_2$, and ε, and a correct understanding of their role is essential for the experiments of Chapter 4.

μ: Length Regularization Parameter

The parameter $\mu \geq 0$ controls the contribution of the contour length term, exactly as in the general curvature-flow discussion of Section 1.3. Increasing μ places greater emphasis on short and regular contours and therefore generally produces smoother segmentations, at the risk, taken to an extreme, of preventing the contour from detecting small or thin objects, or from splitting to capture several disconnected objects, since each additional connected component of C increases $\text{Length}(C)$. Conversely, a smaller value of μ allows the contour to adapt more closely to image structures, although it may also make the segmentation more sensitive to noise and small-scale variations.

ν: Area Parameter

The parameter ν controls the contribution of the area term $\nu \int_{\Omega} H(\phi) \, dx$. With the sign convention adopted here, the corresponding force in the gradient-flow equation is $-\nu$, so the sign of ν

influences the tendency of the contour to expand or contract: $\nu > 0$ biases the contour toward shrinking, while $\nu < 0$ biases it toward expanding from a small initial curve. The case $\nu = 0$ removes the explicit area bias altogether; this is a common and often sufficient choice in practice, and is used as the default setting in Chapter 4.

λ_1 and λ_2: Fidelity Parameters

The parameters λ_1 and λ_2 determine the relative importance of the two fidelity terms $\lambda_1 \int_\Omega (u_0 - c_1)^2 H(\phi)\, dx$ and $\lambda_2 \int_\Omega (u_0 - c_2)^2 (1 - H(\phi))\, dx$. When $\lambda_1 = \lambda_2$ (the most common choice, used throughout Chapter 4 unless stated otherwise), the two phases are treated symmetrically; different values allow the relative importance of the inside and outside fitting terms to be adjusted, for instance when one region is known to be more homogeneous than the other. In all cases, increasing λ_1, λ_2 relative to μ shifts the balance of (3.36) toward the data terms and away from the curvature term, i.e. toward a more detailed, less regularized contour; decreasing them has the opposite effect.

ε: Heaviside Regularization Parameter

The parameter $\varepsilon > 0$ controls the smooth approximation of the Heaviside function and its derivative introduced in Section 3.6. A smaller value provides a sharper transition around the zero level set, closer to the ideal (discontinuous) model, whereas a larger value produces a smoother transition and a more global influence of δ_ε over Ω. In numerical experiments, ε must be chosen consistently with the spatial discretization and the numerical approximation of the curvature, a point taken up again in Chapter 4.

Conclusion of the Chapter

This chapter developed the mathematical formulation of the two-phase Chan–Vese model within the level set framework. Starting from the variational formulation, the first variation and gradient descent procedure led to the Euler–Lagrange equation and the corresponding level set evolution equation, together with the natural homogeneous Neumann boundary condition. The region constants $c_1(\phi)$ and $c_2(\phi)$ were identified as the mean intensities of the inside and outside regions, while the regularization of the Heaviside and Dirac functions yielded a smooth evolution equation suitable for numerical treatment. The roles of the model parameters μ, ν, λ_1, λ_2, and ε were also discussed.

The resulting regularized evolution equation, together with its boundary condition and the update rules for c_1 and c_2, provides the continuous model to be approximated numerically. Chapter 4

therefore focuses on its finite-difference discretization, the numerical treatment of the curvature term, the reinitialization of ϕ, the choice of the time step, and the implementation and experimental evaluation of the resulting algorithm.

Chapter 4

Numerical Implementation and Experiments

Chapter 3 formulated the Chan–Vese model within the level set framework and derived the regularized evolution equation

$$\frac{\partial \phi}{\partial t} = \delta_\varepsilon(\phi)\left[\mu\,\mathrm{div!}\left(\frac{\nabla\phi}{|\nabla\phi|}\right) - \nu - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2\right], \tag{4.1}$$

together with the homogeneous Neumann boundary condition and the expressions for the region constants $c_1(\phi)$ and $c_2(\phi)$. Equation (4.1) is a nonlinear partial differential equation and, for a general image, does not admit a closed-form solution. The purpose of this chapter is therefore to construct a practical numerical approximation of the continuous model, describe its implementation, and evaluate its behavior through numerical experiments.

The chapter is organized as follows. Section 1 discretizes (4.1) on a uniform pixel grid using finite differences, following the finite-difference principles introduced in Section 1.2. Section 2 discusses the numerical stability of the resulting explicit scheme and introduces the discrete energy as a diagnostic tool. Section 3 specifies the discrete form of the boundary condition. Section 4 describes the reinitialization procedure used to keep ϕ close to a signed distance function, as discussed in Section 1.5. Section 5 combines these ingredients into a complete numerical algorithm. Finally, Section 6 presents the experimental results and analyzes the numerical behavior of the Chan–Vese model in relation to the theoretical properties established in Chapters 2 and 3.

4.1 Finite Difference Discretization of the Evolution Equation

The Discrete Grid

Let the image u_0 be sampled, as in Section 1.1.2, on a uniform grid of $N \times M$ pixels, and denote by $u_{0,i,j}$ the intensity at pixel (i, j) and by $\phi^n_{i,j}$ the value of the level set function at the same pixel and at discrete time $t_n = n\Delta t$, $n = 0, 1, 2, \ldots$. The spatial grid sizes h_x, h_y are taken equal, $h_x = h_y = h$, and, as is natural for digital images, $h = 1$ in the experiments. The discrete level set function is thus represented by the matrix $\phi^n = (\phi^n_{i,j})_{1\le i\le N, 1\le j\le M}$, whose zero level set is the numerical approximation of the evolving contour at time t_n.

Finite Difference Operators

Following Section 1.2, for an interior grid point the forward and backward differences are

$$D_x^+\phi_{i,j} = \frac{\phi_{i+1,j} - \phi_{i,j}}{h}, \qquad D_x^-\phi_{i,j} = \frac{\phi_{i,j} - \phi_{i-1,j}}{h},$$

and similarly $D_y^\pm\phi_{i,j}$, while the central differences are

$$D_x^0\phi_{i,j} = \frac{\phi_{i+1,j} - \phi_{i-1,j}}{2h}, \qquad D_y^0\phi_{i,j} = \frac{\phi_{i,j+1} - \phi_{i,j-1}}{2h}, \tag{4.2}$$

which provide a natural approximation of the gradient, $\nabla\phi_{i,j} \approx (D_x^0\phi_{i,j}, D_y^0\phi_{i,j})$.

Regularized Gradient Magnitude

The curvature operator contains the normalized gradient $\nabla\phi/|\nabla\phi|$, whose direct evaluation is problematic where $|\nabla\phi|$ is very small. To avoid division by zero, we use, as anticipated in Section 3.6, the regularized magnitude

$$|\nabla\phi_{i,j}|_\eta = \sqrt{\left(D_x^0\phi_{i,j}\right)^2 + \left(D_y^0\phi_{i,j}\right)^2 + \eta^2}, \qquad \eta > 0 \text{ small.} \tag{4.3}$$

This regularization is purely numerical, introduced to avoid division by zero in the discrete curvature; it should be distinguished from the regularization of the Heaviside and Dirac functions controlled by ε.

Discretization of the Curvature Term

The curvature term $\kappa = \operatorname{div}(\nabla\phi/|\nabla\phi|)$ of Section 1.3 and Section 3.6 is discretized by first computing the normalized gradient at the interfaces between neighboring pixels and then taking its discrete divergence, preserving the structure "curvature = divergence of normalized gradient". For the horizontal flux at $(i+\frac{1}{2}, j)$,

$$q^x_{i+\frac{1}{2},j} = \frac{D_x^+\phi_{i,j}}{\sqrt{\left(D_x^+\phi_{i,j}\right)^2 + \left(D_y^0\phi_{i+\frac{1}{2},j}\right)^2 + \eta^2}}, \qquad D_y^0\phi_{i+\frac{1}{2},j} = \frac{1}{2}\left(D_y^0\phi_{i,j} + D_y^0\phi_{i+1,j}\right), \tag{4.4}$$

and analogously for $q^x_{i-\frac{1}{2},j}$ (using D_x^-) and for the vertical fluxes $q^y_{i,j+\frac{1}{2}}$, $q^y_{i,j-\frac{1}{2}}$ (using $D_y^\pm$ and averaged D_x^0). The discrete curvature at (i,j) is then the discrete divergence of these four half-point fluxes,

$$\kappa_{i,j} = \frac{q^x_{i+\frac{1}{2},j} - q^x_{i-\frac{1}{2},j}}{h} + \frac{q^y_{i,j+\frac{1}{2}} - q^y_{i,j-\frac{1}{2}}}{h} \approx \operatorname{div}\left(\frac{\nabla\phi}{|\nabla\phi|_\eta}\right)_{i,j}, \tag{4.5}$$

the discrete counterpart of the regularized curvature κ_η of (3.38). This discretization is local – $\kappa_{i,j}$ depends only on a neighborhood of (i,j) – which makes it straightforward to implement.

Discretization of the Region Constants

For every time step, the constants c_1, c_2 are recomputed from the current ϕ, using the regularized Heaviside function H_ε of Section 3.6:

$$c_1^n = \frac{\sum_{i,j} u_{0,i,j}\, H_\varepsilon(\phi_{i,j}^n)}{\sum_{i,j} H_\varepsilon(\phi_{i,j}^n) + \zeta}, \qquad c_2^n = \frac{\sum_{i,j} u_{0,i,j}\big[1 - H_\varepsilon(\phi_{i,j}^n)\big]}{\sum_{i,j}\big[1 - H_\varepsilon(\phi_{i,j}^n)\big] + \zeta}, \tag{4.6}$$

the discrete counterparts of (3.14)–(3.15), where $\zeta > 0$ is a small positive constant added to the denominators to avoid numerical problems in degenerate situations (in normal segmentations, where both phases have non-negligible area, this safeguard has practically no effect).

Time Discretization

Discretizing the time derivative of (4.1) by the forward Euler approximation $\partial\phi/\partial t \approx (\phi_{i,j}^{n+1} - \phi_{i,j}^n)/\Delta t$, and evaluating the regularized Dirac function $\delta_\varepsilon(\phi_{i,j}^n) = \frac{1}{\pi}\cdot\frac{\varepsilon}{\varepsilon^2 + (\phi_{i,j}^n)^2}$ and the discrete curvature (4.5) at the current iteration, the explicit update reads

$$\phi_{i,j}^{n+1} = \phi_{i,j}^n + \Delta t\, \delta_\varepsilon(\phi_{i,j}^n)\Big[\mu\kappa_{i,j}^n - \nu - \lambda_1(u_{0,i,j} - c_1^n)^2 + \lambda_2(u_{0,i,j} - c_2^n)^2\Big]. \tag{4.7}$$

All quantities on the right-hand side are evaluated at iteration n, so the scheme is explicit. The algorithm has an important feedback structure: the current contour determines the phase means c_1^n, c_2^n, these means determine the fidelity force, and the resulting force changes the contour at the next time step.

4.2 Numerical Stability and Energy Monitoring

The CFL Condition

The explicit scheme (4.7) requires an appropriate choice of Δt. The main restrictive term is the curvature contribution, which has a diffusive character: recall from Section 1.4 that, for the explicit heat equation $u_t = D\Delta u$, stability requires $\Delta t \le h^2/(4D)$. The Chan–Vese equation is nonlinear, however, and its effective diffusion coefficient depends on both $\delta_\varepsilon(\phi)$ and the geometry of the level set, so the heat-equation condition cannot be interpreted as a rigorous stability theorem for the full nonlinear scheme. Nevertheless, since $0 < \delta_\varepsilon(\phi) \le \delta_\varepsilon(0) = 1/(\pi\varepsilon)$, a useful conservative CFL-type guideline based on the curvature coefficient is

$$\Delta t \le \frac{h^2}{4\mu\,\delta_\varepsilon(0)} = \frac{\pi\varepsilon h^2}{4\mu}. \tag{4.8}$$

This condition is used as a practical time-step restriction in the present implementation; it should not be interpreted as a proof that the complete nonlinear discrete energy decreases monotonically

for every parameter configuration. It does, however, capture the expected qualitative dependence on the numerical parameters: for fixed h, μ, decreasing ε increases $\delta_\varepsilon(0)$ and forces a smaller admissible time step, and increasing μ has the same effect. The fidelity terms also contribute to the explicit update and may become important for particular choices of λ_1, λ_2, the intensity range, and the initialization; in practice it is therefore preferable to use a time step smaller than the nominal bound (4.8) when necessary, and to monitor the numerical evolution as described next.

Energy Monitoring

The continuous Chan–Vese model is obtained from the gradient-flow formulation of Section 3.5, for which the energy is expected to decrease along a smooth continuous evolution under appropriate assumptions; for the discrete explicit scheme, however, monotone energy decrease is not automatically guaranteed by the CFL-type condition (4.8). The discrete energy is therefore monitored numerically, at every iteration n, as

$$\begin{aligned} F_\varepsilon^n = \mu \sum_{i,j} \delta_\varepsilon(\phi_{i,j}^n)\, |\nabla \phi_{i,j}^n|_\eta + \nu \sum_{i,j} H_\varepsilon(\phi_{i,j}^n) \\ + \lambda_1 \sum_{i,j} (u_{0,i,j} - c_1^n)^2 H_\varepsilon(\phi_{i,j}^n) + \lambda_2 \sum_{i,j} (u_{0,i,j} - c_2^n)^2 \big[1 - H_\varepsilon(\phi_{i,j}^n)\big], \end{aligned} \tag{4.9}$$

(with the appropriate quadrature factor h^2 included when $h \neq 1$). Monitoring F_ε^n allows one to determine experimentally whether the selected parameters produce a stable, energy-decreasing evolution; if oscillations or increases in the energy are observed, the time step can be reduced.

4.3 Boundary Condition

The continuous formulation of Chapter 3 leads to the natural boundary condition $\delta(\phi)\, \partial_n \phi / |\nabla \phi| = 0$ on $\partial\Omega$, equation (3.39). In the numerical implementation, we impose the stronger homogeneous Neumann condition

$$\frac{\partial \phi}{\partial n} = 0 \qquad \text{on } \partial\Omega, \tag{4.10}$$

approximated, for a finite-difference implementation, by a mirror (zero-normal-gradient) condition – e.g. $\phi_{0,j} = \phi_{1,j}$ at the left boundary, and analogously at the remaining boundaries. This prevents artificial normal variation of the level set function through the image boundary.

4.4 Reinitialization of the Level Set Function

Motivation

As recalled in Section 1.5, the level set method does not require ϕ to be a signed distance function in the continuous formulation; however, maintaining a reasonable gradient magnitude is advantageous numerically, since the curvature operator contains the normalized gradient $\nabla\phi/|\nabla\phi|$.

During repeated numerical updates, ϕ may depart from its initial signed-distance structure, which can reduce the quality of the finite-difference approximation and make the evolution less robust. A reinitialization procedure is therefore applied periodically; it is a purely numerical technique for the level set representation, not an additional term in the Chan–Vese energy.

The Reinitialization Equation

Let ϕ_0 denote the level set function immediately before reinitialization. An auxiliary function $\psi(x,\tau)$, initialized by $\psi(x,0)=\phi_0(x)$, is evolved in a fictitious time τ by

$$\frac{\partial \psi}{\partial \tau} = S_\eta(\phi_0)\big(1-|\nabla\psi|\big), \qquad S_\eta(\phi_0)=\frac{\phi_0}{\sqrt{\phi_0^2+\eta^2}}, \tag{4.11}$$

where S_η is a smoothed approximation of the sign function. This is a Hamilton–Jacobi equation of the type introduced in Section 1.4; its desired steady state satisfies $|\nabla\psi|\approx 1$ away from the zero level set, while $S_\eta(\phi_0)$ ensures that the zero level set of ψ coincides, at every τ, with that of ϕ_0. In a discrete computation this property is only approximate, and depends on the pseudo-time step, the number of iterations, and the spatial discretization.

Upwind Discretization

Because (4.11) is a Hamilton–Jacobi equation, an upwind (Godunov) discretization is preferable to a central-difference approximation. Writing $a_{i,j}=D_x^-\psi_{i,j}$, $b_{i,j}=D_x^+\psi_{i,j}$, $c_{i,j}=D_y^-\psi_{i,j}$, $d_{i,j}=D_y^+\psi_{i,j}$, the Godunov approximation of $|\nabla\psi|$ is

$$G_{i,j}^{+}=\sqrt{\max(a_{i,j},0)^2+\min(b_{i,j},0)^2+\max(c_{i,j},0)^2+\min(d_{i,j},0)^2} \tag{4.12}$$

where $S_\eta(\phi_{0,i,j})\geq 0$, and

$$G_{i,j}^{-}=\sqrt{\min(a_{i,j},0)^2+\max(b_{i,j},0)^2+\min(c_{i,j},0)^2+\max(d_{i,j},0)^2} \tag{4.13}$$

where $S_\eta(\phi_{0,i,j})<0$, so that $G_{i,j}=G_{i,j}^{+}$ or $G_{i,j}^{-}$ according to the sign of $S_\eta(\phi_{0,i,j})$. The pseudo-time update is then

$$\psi_{i,j}^{m+1}=\psi_{i,j}^{m}+\Delta\tau\, S_\eta(\phi_{0,i,j})\big(1-G_{i,j}^{m}\big), \tag{4.14}$$

with $\Delta\tau$ chosen small enough for the upwind scheme, and reinitialization stopped after a prescribed number of pseudo-time iterations or by a suitable convergence criterion.

Frequency of Reinitialization

Reinitializing at every step of the main evolution would be needlessly expensive and could introduce unnecessary numerical perturbations. In the present implementation, ϕ is reinitialized

every K main iterations, with $K = 10$ or $K = 20$ used as typical initial choices; the numerical experiments should verify that reinitialization improves the conditioning of ϕ without producing an appreciable displacement of the zero level set.

4.5 The Complete Algorithm

Combining the ingredients of Sections 1–4, the full numerical procedure is summarized in Algorithm 4.1.

Algorithm 4.1: Numerical Chan–Vese Segmentation

1. **Input:** grayscale image u_0; parameters $\mu, \nu, \lambda_1, \lambda_2, \varepsilon, \eta$; time step Δt satisfying (4.8); maximum number of iterations $N_{\max}$; reinitialization period K; stopping tolerance tol.
2. **Initialization:** construct ϕ^0 (preferably a signed distance function) whose zero level set represents the initial contour C_0.
3. **For** $n = 0, 1, \ldots, N_{\max} - 1$:
   (a) Compute $H_\varepsilon(\phi^n)$ and the region constants c_1^n, c_2^n using (4.6).
   (b) Compute the discrete normalized-gradient fluxes and the curvature κ^n using (4.4)–(4.5).
   (c) Compute $\delta_\varepsilon(\phi^n)$ and update ϕ^{n+1} using the explicit scheme (4.7).
   (d) Enforce the Neumann boundary condition (4.10) on ϕ^{n+1}.
   (e) If the reinitialization period has been reached, reinitialize ϕ^{n+1} using (4.14).
   (f) Compute $E_\phi^n = \|\phi^{n+1} - \phi^n\|_2 / (\|\phi^n\|_2 + \zeta)$; if $E_\phi^n <$ tol, stop.
4. **Output:** the final level set function ϕ, the zero level set $C = \{x : \phi(x) = 0\}$, and the final constants c_1, c_2.

The algorithm emphasizes the alternating structure of the numerical procedure: the region statistics are updated from the current segmentation, while the segmentation is subsequently updated using the resulting data-fitting forces.

Implementation Notes

The algorithm above is implemented in MATLAB, whose matrix operations are well suited to a finite-difference scheme of this kind. The image and the level set function are stored as matrices of the same dimensions, and the implementation is organized directly around the steps of the algorithm: initialization; computation of $H_\varepsilon(\phi)$ and $\delta_\varepsilon(\phi)$; computation of c_1, c_2; computation of the normalized gradient and curvature (4.4)–(4.5); the explicit update (4.7); treatment of the boundary; optional periodic reinitialization; and evaluation of the stopping criterion and the discrete energy (4.9). The principal routine can be represented schematically as

$\phi^{n+1} = \mathcal{T}(\phi^n, u_0; \mu, \nu, \lambda_1, \lambda_2, \varepsilon, \eta, \Delta t)$, and must consistently use the sign convention $\phi > 0$ inside the contour, $\phi < 0$ outside, since reversing it without simultaneously changing the Heaviside and fidelity terms would reverse the interpretation of the two phases.

4.6 Experimental Results

The experiments presented in this section are designed to investigate the numerical behavior of the two-phase Chan–Vese model and to illustrate, through representative examples, the main properties discussed in Chapters 2 and 3. In particular, the experiments examine the evolution of the active contour, the influence of image characteristics such as noise and contrast, the detection of objects with different geometries, the ability to handle multiple objects, and the applicability of the model to real grayscale images.

Two main categories of images are considered: synthetic images and real images. Synthetic images are used first because they provide controlled test cases in which the geometry and intensity characteristics of the objects can be prescribed. They therefore provide a convenient setting for examining the behavior of the numerical implementation. The synthetic examples include clean two-level images, noisy images, low-contrast images, objects with smooth contours, and images containing several disjoint objects. These examples allow the evolution of the contour and the effect of the region-based data term to be observed under different conditions.

For noisy images, additive Gaussian noise may be introduced according to

$$u_0 = u_{\text{clean}} + \sigma_n \xi,$$

where ξ denotes a zero-mean, unit-variance Gaussian field and σ_n controls the noise level. This formulation provides a simple way of examining the behavior of the model when the observed image is a perturbed version of an underlying clean image. Low-contrast examples are used to investigate cases in which the intensity difference between the object and its surrounding region is relatively small.

The second category consists of real grayscale images, including medical, ultrasound, natural, and benchmark-type images. These examples are intended to illustrate the practical behavior of the model on images that are less idealized than the synthetic test cases. In accordance with the two-phase piecewise-constant formulation, the selected examples are characterized by regions that can be reasonably distinguished according to their average intensities. For each segmentation result, the corresponding piecewise-constant approximation is also displayed. This approximation provides a direct visualization of the two-region representation associated with the Chan–Vese model.

Synthetic Images

We first consider synthetic images in order to examine the evolution of the contour under controlled conditions. Figure 4.1 shows the segmentation of a clean synthetic image. The initial contour is displayed together with the contours obtained after 50 and 100 iterations.

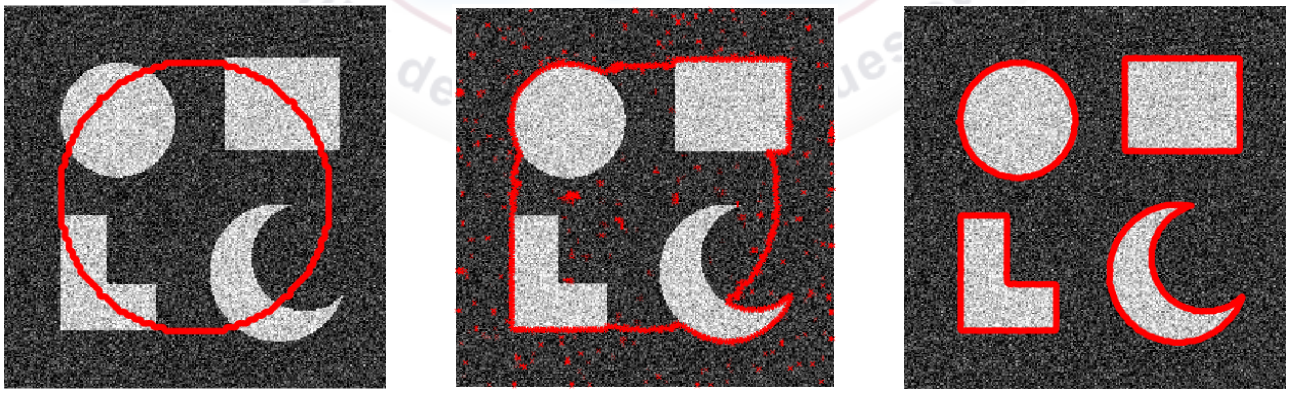


(a) (b) (c)

Figure 4.1: Chan–Vese segmentation of a clean synthetic image. Left: u_0 with the initial contour. Middle: the contour after 50 iterations. Right: the contour after 100 iterations.

As shown in Figure 4.1, the contour evolves from its initial position toward the boundary separating the two intensity regions. The successive contours illustrate the numerical evolution of the level set representation and show how the segmentation develops as the iterations proceed. For a clean two-level image, the distinction between the two regions is sufficiently clear for the region-based data term to guide the contour toward the object boundary.

The effect of noise is illustrated in Figure 4.2. In this experiment, the input image is obtained by adding Gaussian noise to the synthetic image.



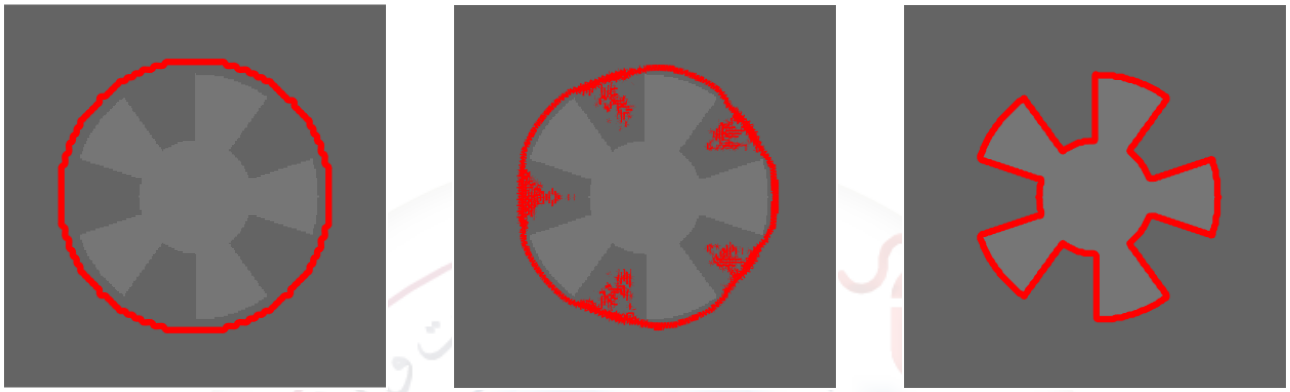
(a) (b) (c)

Figure 4.2: Chan–Vese segmentation of a noisy synthetic image. Left: u_0 with the initial contour. Middle: the contour after 50 iterations. Right: the contour after 100 iterations.

Figure 4.2 illustrates that the segmentation is not based exclusively on local gradient information. Instead, the Chan–Vese model uses the mean intensity of the regions inside and outside the evolv-

ing contour through the data-fitting term. Consequently, the contour can still evolve toward the main object boundary even when the image contains intensity fluctuations caused by noise. The result also illustrates the role of the region-based formulation in situations where a purely edge-based criterion may be affected by small-scale variations in the image.

A different situation arises when the contrast between the object and its surrounding region is reduced. Figure 4.3 presents a low-contrast synthetic example.

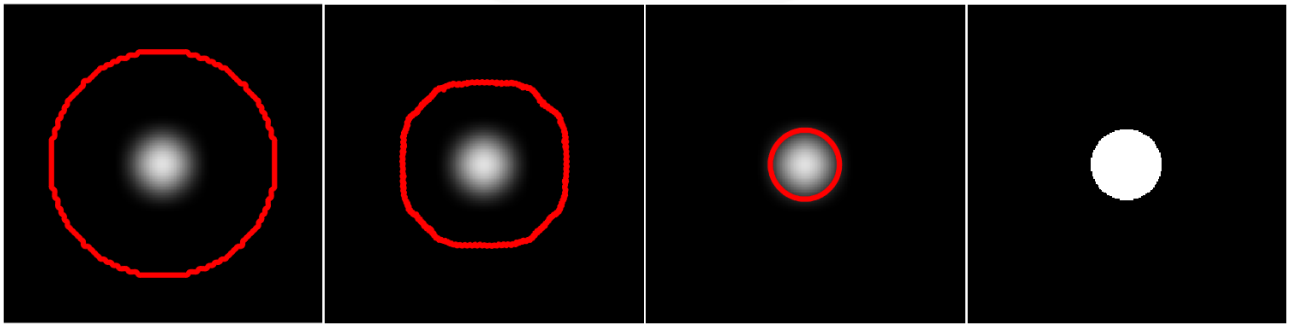


(a) (b) (c)

Figure 4.3: Chan–Vese segmentation of a low-contrast synthetic image. Left: u_0 with the initial contour. Middle: the contour after 50 iterations. Right: the contour after 100 iterations.

The result in Figure 4.3 illustrates the behavior of the model when the intensity difference between the two regions is less pronounced. Although the boundary may not correspond to a strong edge, the region-based formulation can use the difference between the average intensities of the two phases to drive the contour. This example therefore illustrates an important distinction between the Chan–Vese model and classical edge-based active contours.

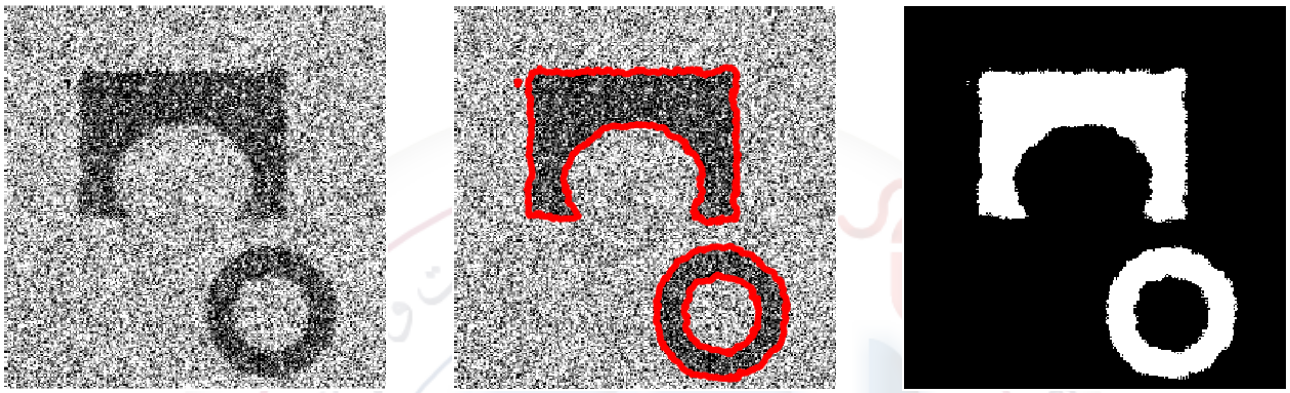
The evolution of the contour for an object with a smooth boundary is considered next. Figure 4.4 presents several stages of the segmentation together with the resulting piecewise-constant approximation.



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

Figure 4.4: Chan–Vese segmentation of an object with a smooth contour. Left: u_0 with the initial contour. Middle-left: the contour after 80 iterations. Middle-right: the contour after 150 iterations. Right: the corresponding piecewise-constant approximation.

The successive contours in Figure 4.4 provide a visual representation of the evolution of the level set function toward the final segmentation. The last image shows the corresponding piecewise-constant approximation obtained from the final partition of the image domain. In particular, the approximation assigns one representative intensity to each of the two regions, thereby illustrating directly the piecewise-constant assumption on which the two-phase Chan–Vese model is based.
The ability to handle several objects is illustrated in Figure 4.5, where multiple objects with different shapes are present in a noisy image.



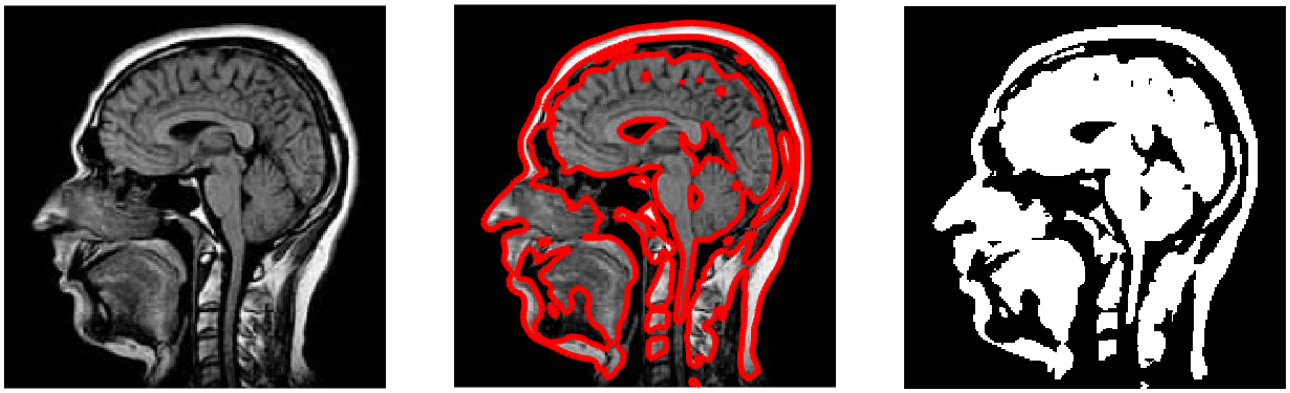
(a) (b) (c)

Figure 4.5: Detection of multiple objects in a noisy synthetic image. Left: u_0. Middle: the final Chan–Vese segmentation. Right: the corresponding piecewise-constant approximation.

Figure 4.5 illustrates the ability of the level set representation to describe several objects within a single computational domain. The final segmentation identifies the regions represented by the evolving contour, while the corresponding piecewise-constant approximation displays the two intensity levels associated with the resulting partition. This example also illustrates the topological flexibility of the level set formulation, since the evolving zero level set can represent multiple components without explicitly parameterizing each contour separately.
Taken together, the synthetic experiments demonstrate several characteristic features of the Chan–Vese model. The contour evolves according to region statistics rather than relying exclusively on local edge information, allowing the model to be applied to clean, noisy, and low-contrast images. The experiments also illustrate the flexibility of the level set formulation with respect to the geometry and number of objects.

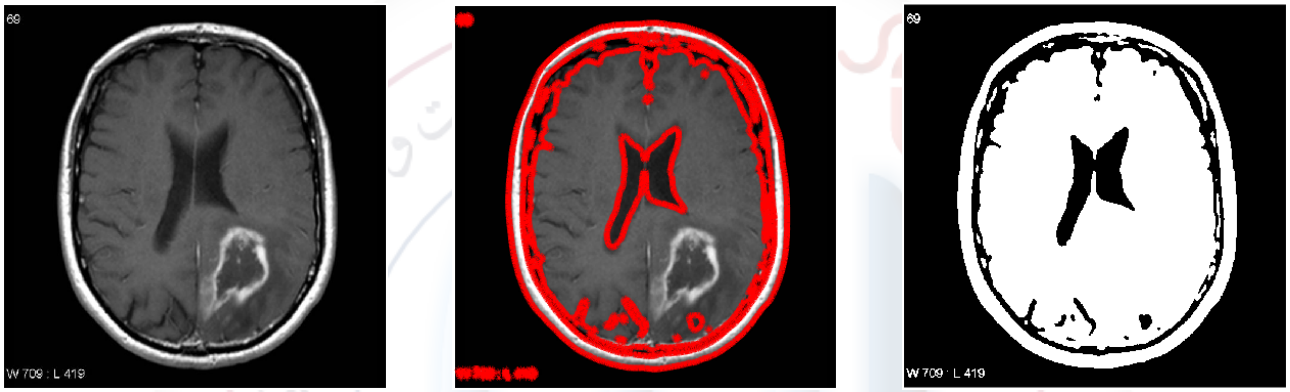
Real Images

We next consider real grayscale images in order to examine the behavior of the numerical implementation on less idealized data. In each example, the original image, the final Chan–Vese segmentation, and the corresponding piecewise-constant approximation are displayed.
The first two examples are medical images, as shown in Figures 4.6 and 4.7.



(a) (b) (c)

Figure 4.6: Chan–Vese segmentation of a medical image. Left: u_0. Middle: the final segmentation. Right: the corresponding piecewise-constant approximation.



(a) (b) (c)

Figure 4.7: Chan–Vese segmentation of a second medical image. Left: u_0. Middle: the final segmentation. Right: the corresponding piecewise-constant approximation.

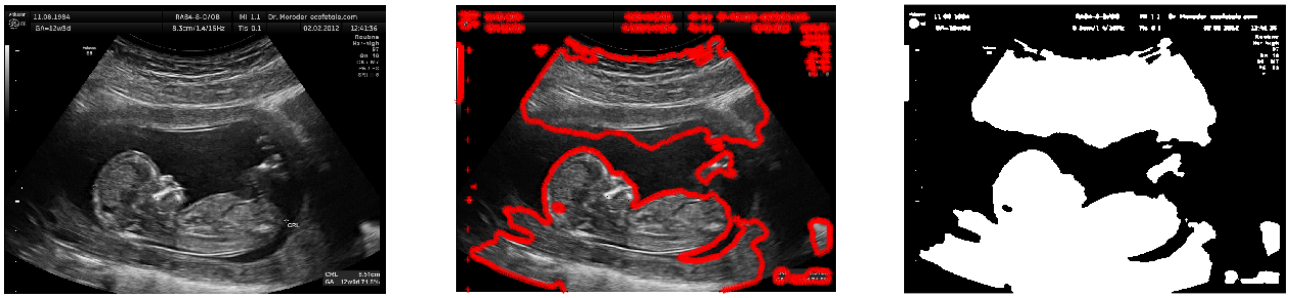
The medical-image experiments show how the region-based model can be applied to images whose structures are not represented by ideal geometric shapes. The corresponding piecewise-constant approximations provide a simplified representation of the detected regions and make it possible to visualize the intensity partition produced by the model.

An ultrasound image is considered next in order to examine the behavior of the model on another type of medical data.

Figure 4.8 illustrates the application of the model to ultrasound data. As in the previous examples, the final partition is determined by the region-based criterion, and the piecewise-constant approximation represents the average intensity assigned to each of the two phases.

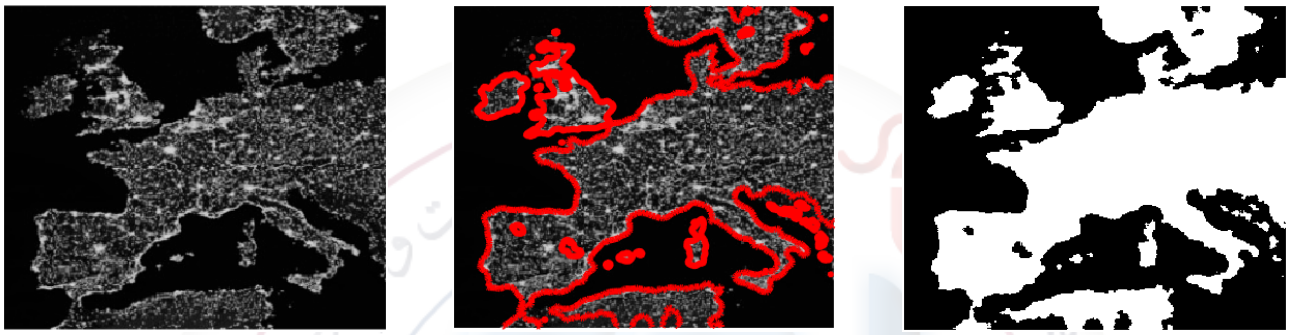
The next example is the Europe night-lights image, which provides a different type of real grayscale data.

The result in Figure 4.9 further demonstrates the use of the two-phase model on a real image with spatially varying intensity. The segmentation separates the image into two regions according to the intensity statistics represented in the Chan–Vese energy, while the piecewise-constant approx-



(a) (b) (c)

Figure 4.8: Chan–Vese segmentation of an ultrasound image. Left: u_0. Middle: the final segmentation. Right: the corresponding piecewise-constant approximation.

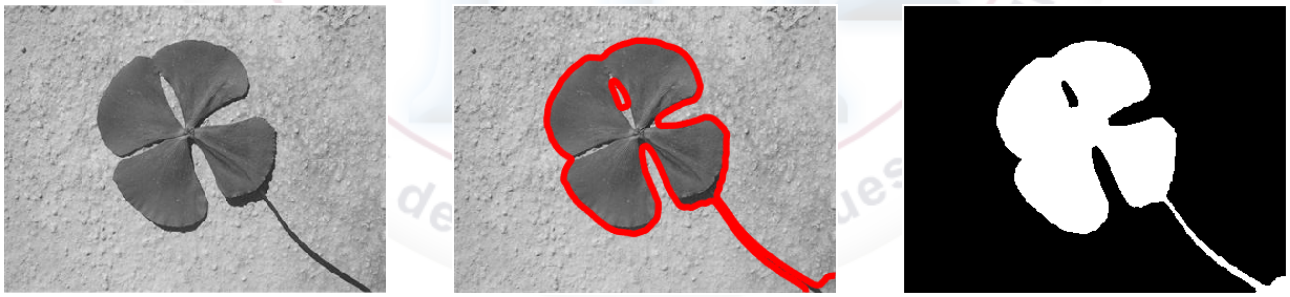


(a) (b) (c)

Figure 4.9: Chan–Vese segmentation of the Europe night-lights image. Left: u_0. Middle: the final segmentation. Right: the corresponding piecewise-constant approximation.

imation gives a simplified representation of the resulting partition.

Finally, two additional natural-image examples are considered in Figures 4.10 and 4.11.

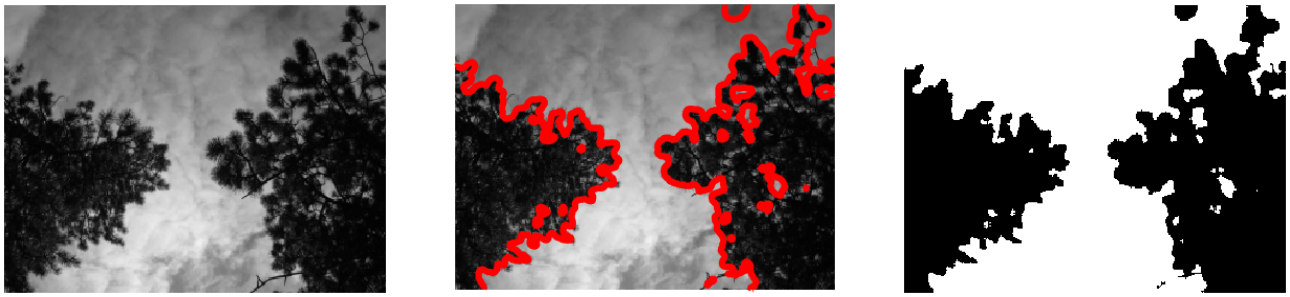


(a) (b) (c)

Figure 4.10: Chan–Vese segmentation of a natural image. Left: u_0. Middle: the final segmentation. Right: the corresponding piecewise-constant approximation.

The natural-image examples demonstrate that the numerical implementation can also be applied to images with more complex structures and intensity distributions. The resulting segmentations should, however, be interpreted in accordance with the assumptions of the two-phase piecewise-constant model: the method is primarily driven by differences in regional mean intensities rather than by detailed texture or purely local edge information.

Overall, the real-image experiments complement the synthetic tests by demonstrating the behav-



(a) (b) (c)

Figure 4.11: Chan–Vese segmentation of a natural image. Left: u_0. Middle: the final segmentation. Right: the corresponding piecewise-constant approximation.

ior of the Chan–Vese model on practical image data. Across the examples, the final segmentation and the associated piecewise-constant approximation provide two complementary views of the result: the former emphasizes the detected partition of the image domain, while the latter illustrates the intensity representation associated with the two phases of the model.

These experiments therefore provide a numerical illustration of the theoretical formulation developed in the preceding chapters. In particular, they show how the level set representation enables the evolution of the contour without an explicit parametrization of its geometry, while the region-based data term allows the segmentation to be driven by differences in the average intensities of the two phases.

Conclusion of the Chapter

This chapter presented the numerical implementation of the two-phase Chan–Vese model. The evolution equation was discretized using finite differences, with the curvature term approximated through the divergence of a normalized gradient. Numerical regularization, a practical CFL-type time-step restriction, and a discrete energy diagnostic were also introduced. The homogeneous Neumann boundary condition and periodic level-set reinitialization were incorporated to improve the numerical evolution.

The resulting scheme was applied to synthetic and real grayscale images, including clean, noisy, low-contrast, multiple-object, medical, ultrasound, and natural images. The obtained contours and piecewise-constant approximations illustrate the behavior of the model under different image conditions and complement the mathematical analysis developed in the previous chapters.

General Conclusion

This thesis has been devoted to the mathematical study and numerical implementation of the Chan–Vese model for image segmentation, from its variational origins to a practical numerical algorithm.

Chapter 1 introduced the mathematical background required for this study: image derivatives and edges, curvature, level set functions, partial differential equations, and the distinction between edge-based and region-based segmentation. Chapter 2 presented the Mumford–Shah functional as the general variational framework for image segmentation, introduced the space $BV(\Omega)$ of functions of bounded variation and the perimeter functional needed to give its length-regularization term a rigorous meaning, and derived the Chan–Vese model as its two-phase, piecewise-constant specialization: rather than detecting edges explicitly, the model seeks a contour separating the image into two regions represented by their mean intensities c_1 and c_2. The same BV framework was used to establish existence of a minimizer for this simplified energy, and the chapter closed with a comparison of Chan–Vese against the classical Snakes and Geodesic Active Contours models.

Chapter 3 first recalled the calculus-of-variations tools – functionals, first variation, the Euler–Lagrange equation – needed to pass rigorously from an energy to a PDE, then reformulated the Chan–Vese model in the level set framework, representing the contour implicitly through a function ϕ, and introduced the regularized Heaviside and Dirac functions needed for a numerical treatment. Computing the first variation of the resulting energy yielded the Chan–Vese evolution equation

$$\frac{\partial \phi}{\partial t} = \delta_\varepsilon(\phi)\left[\mu \operatorname{div}\left(\frac{\nabla \phi}{|\nabla \phi|}\right) - \nu - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2\right], \tag{4.15}$$

in which the contour evolves under the joint action of a geometric regularization force and region-based fidelity forces, with c_1, c_2 recomputed at every stage as the optimal constant representatives of the two phases.

Chapter 4 focused on the numerical implementation. The evolution equation was discretized using finite differences and an explicit time-stepping scheme, together with appropriate boundary treatment and periodic reinitialization of the level set function. The resulting algorithm was implemented in MATLAB and tested on synthetic and real grayscale images, illustrating its behavior for different geometries, noise levels, contrasts, and image types.

Overall, this work provides a complete mathematical and computational treatment of the classical two-phase Chan–Vese model, combining geometric regularization, region-based image statistics, and numerical approximation within a single, coherent framework for image segmentation.

Bibliography

[1] L. Ambrosio and V. M. Tortorelli. Approximation of functionals depending on jumps by elliptic functionals via Γ-convergence. Communications on Pure and Applied Mathematics, 43(8):999-1036, 1990.

[2] G. Aubert and P. Kornprobst. Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations. Springer, 2006.

[3] G. Aubert and L. Blanc-Féraud. Some remarks on the equivalence between 2D and 3D classical snakes and geodesic active contours. The International Journal of Computer Vision, 34(1):19-28, 1999.

[4] J. Canny. A computational approach to edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-8(6):679-698, 1986.

[5] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. The International Journal of Computer Vision, 22(1):61-79, 1997.

[6] T. F. Chan and L. A. Vese. Active contours without edges. IEEE Transactions on Image Processing, 10(2):266-277, 2001.

[7] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. The International Journal of Computer Vision, 1(4):321-331, 1988.

[8] D. Marr and E. Hildreth. Theory of edge detection. Proceedings of the Royal Society of London. Series B, Biological Sciences, 207(1167):187-217, 1980.

[9] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. Communications on Pure and Applied Mathematics, 42(5):577-685, 1989.

[10] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. Journal of Computational Physics, 79(1):12-49, 1988.

[11] G. Sapiro. Geometric Partial Differential Equations and Image Analysis. Cambridge University Press, 2001.

[12] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. IEEE Transactions on Pattern Analysis and Machine Intelligence, 12(7):629-639, 1990.

[13] J. M. S. Prewitt. Object enhancement and extraction. In B. S. Lipkin and A. Rosenfeld, editors, Picture Processing and Psychopictorics, pages 75-149. Academic Press, New York, 1970.

[14] L. G. Roberts. Machine Perception of Three-Dimensional Solids. PhD thesis, Massachusetts Institute of Technology, 1963.

[15] I. Sobel and G. Feldman. A 3x3 isotropic gradient operator for image processing. Presented at the Stanford Artificial Intelligence Project (SAIL), 1968.

[16] L. A. Vese and T. F. Chan. A multiphase level set framework for image segmentation using the Mumford and Shah model. The International Journal of Computer Vision, 50(3):271-293, 2002.

Faculté des Mathématiques et de l'Informatique