

Fortran Based Chan-Vese Segmentation For Noisy MRI Images

Bachelor Thesis

Author A.J Brito Lima *
Supervisors Prof. dr. ir. M. B. van Gijzen
 Ir. K. Slepova
 Dr. ir. W. T. van Horsen
Faculty EEMCS
 Delft University of Technology

July 1, 2024

Laymen Summary

To diagnose diseases using low-field MRI devices, automatic image recognition techniques are required to detect anomalies in a given image. A critical component of this process is image segmentation, which involves dividing the image into coherent regions with similar features to detect the anomalies. Low quality images make segmentation difficult by the existence of noise in the images. The Chan-Vese model is a segmentation technique to segment images into two regions. This report aims to segment the low quality images obtained by the low-field MRI devices with use of the Chan-Vese model. This model is found to be particularly well-suited for handling the challenges posed by low-quality images which the low-field MRI devices produce. Future research should investigate if it is possible to further segment the region with the objects, such that each object has a separate region.

Summary

An important component of image processing is image segmentation, which involves dividing the image into coherent regions with similar features. This makes it possible to apply automatic image recognition techniques. This is very useful in diagnosing diseases in medical image analyses, for example when analysing MRI images. Low-cost mobile MRI scanners tend to have lower resolution and higher levels of noise compared to the standard scanners. This noise and distortion makes image segmentation in the images obtained by those low-cost scanners very challenging.

In this thesis, the use of the Chan-Vese model to segment the low-field MRI images is explored. The Chan-Vese algorithm does not rely on edge detection, but rather segments images based on the intensity variations between regions. This makes it that the model has less trouble with image recognition in lower image quality images than other contour models. The model divides the image in to two regions. A region with the background of the image and a region containing the objects of the image. The research question that this thesis is trying to answer is: can the Chan-Vese model be used effectively to segment the low quality images provided by the low-cost mobile MRI scanners?

The process of finding the optimal segmentation using the Chan-Vese model can be formulated as solving a nonlinear diffusion partial differential equation (PDE). This PDE captures the evolution of a level set function, which implicitly represents the segmentation boundary. The goal of the model is to find the optimal segmentation boundary. This will be done with the use of the computer language Fortran. For the time integration, a semi-implicit scheme will be implemented using Additive Operator Splitting (AOS). AOS is a numerical technique that simplifies the computation of PDEs. It involves breaking down the original PDE into a series of simpler sub-problems that can be solved sequentially. Each sub-problem corresponds to a part of the discretization matrix. To solve the AOS discretization, the Thomas Algorithm will be used. The Thomas Algorithm is a solver for tridiagonal systems that uses LU decomposition.

From the results, it is found that the Chan-Vese segmentation technique can be effective in finding the right segmentation boundary to segment the low-field MRI images. This model proves to be reliable, as the segmentations derived from the provided MRI scans accurately differentiate the scanned objects from the background and any surrounding noise. During testing, it was observed that the model is very sensitive to changes in the parameters. Even slight adjustments to these parameters can have a large impact on the outcome of the segmentation. Future research should investigate if it is possible to further segment the region with the objects, such that each object has a separate region.

Contents

1	Introduction	4
2	Theory	6
2.1	Simplified Mumford-Shah model	6
2.2	Level Set Formulation	7
2.3	Minimization	9
3	Implementation	13
3.1	Spatial discretization	13
3.2	Time integration	17
3.3	Additive Operator splitting	19
4	Results	22
4.1	Chosen parameters	22
4.2	Heaviside function	23
4.3	Noise and Smoothing	25
4.4	Choice of the area penalty	28
4.5	Choice of λ	31
5	Discussion	33
5.1	Disadvantages of gradient descent	33
5.2	Parameter Sensitivity	33
5.3	Choice of length penalty	33
5.4	Two regions segmentation	34
5.5	Reinitialization	34
6	Conclusion	35
	Bibliography	37
	Appendix	39
A	Original MRI images	39

Introduction

Since the first use of X-ray imaging the importance of medical imaging has grown significantly [1]. After X-ray imaging, other imaging methods arose, such as Positron Emission Tomography, Computer Tomography and Magnetic Resonance Imaging. Magnetic Resonance Imaging (MRI) is a non-invasive imaging technology that often used for disease detection, diagnosis and treatment monitoring. It produces detailed images of almost everything in the human body, including the organs, muscles and blood vessels.

The MRI devices are not accessible in 66% of the world [2]. In most developing countries, the utilization of MRI is limited partly due to the high cost of the systems and a lack of infrastructure. To help combat this issue, the TU Delft, Leiden University Medical Center, Penn State and Mbarara University of Science and Technology have developed a low-cost, portable, low-field MRI devices [3, 4, 5, 6]. MRI devices contain magnets. The magnetic field the magnets produces causes the protons in the body to align. These low-field MRI devices uses magnets that have a magnetic strength of less than 1.5 Tesla [6], whereas high-end MRI scanners uses magnets that have a magnetic strength of more than 1.5 Tesla. By replacing the costly magnets with weaker alternatives, the price of the scanners reduces significantly. While this results in significant cost savings, it also brings the challenge of lower image quality, which can affect diagnostic accuracy. The ongoing challenge is to develop low-field scanners that strike an optimal balance between affordability and diagnostic utility, ensuring that the images produced are good enough to be used effectively in medical diagnosis.

Hydrocephalus is prevalent in infants is between one and thirty-two per ten thousand births [7]. It is a condition in which too much cerebrospinal fluid accumulates in the cavities of the brain. Many of the patients with this condition remain untreated. A complication of hydrocephalus can be brain damage. Without proper treatment the condition can have fatal consequences. In children, hydrocephalus has a mortality rate of 0 to 3 % depending on the duration of the follow-up [8]. Hydrocephalus can be treated if it is properly diagnosed and MRI is widely used to diagnose this condition.

Low-field MRI devices have a lower image quality, as the images obtained from those scanners contain noise and distortions. The noise and distortion makes image segmentation in the images obtained by those low-cost scanners very challenging. Most image segmentation methods typically identify the boundary of a region based on the intensity difference between neighbouring pixels. This approach relies on clear and distinct edges

Chapter 1. Introduction

to segment the regions of interest. However, this method may struggle when the boundaries of the object to be segmented are not well-defined due to noise.

The Chan-Vese model is a method for image segmentation in scenarios where the boundary of the objects whose boundaries are not necessarily defined by gradient [9]. This can be the case because the image is smoothed or the edges are not well-defined due to noise. The method is based on level set techniques and active contours without edges. It was introduced by Tony Chan and Luminita Vese in 2001. The Chan-Vese method is an image segmentation method that segments the image into two regions based on the intensity variations between regions.

This thesis aims to evaluate the model's ability to provide accurate and reliable segmentations, thereby assessing its suitability to segment the MRI images obtained from a low-field MRI scanner. The research question that will be answered is: can the Chan-Vese model be used effectively to segment the low quality images provided by the low-cost mobile MRI scanners?

The Chan-Vese model must successfully separate the scanned objects from the background and also distinguishing the objects from the noise typically found in images produced by low-field MRI scanners.

In Chapter 2, the primary focus will be on the theory behind the Chan-Vese model. This chapter will include the theoretical background, the derivation of the method and the optimization process. In chapter 3, numerical implementation of the algorithm will be described. Chapter 4 will present the results obtained from applying the segmentation method to various images.

Theory

In this chapter, the theory behind the Chan-Vese model will be explored [9]. The first part of this chapter contains the derivation of the Chan-Vese functional. In the second part the optimization of finding the best segmentation is discussed. The Chan-Vese algorithm is a segmentation technique used in image processing that divides an image into two regions: one inside a contour and one outside it. The goal of the two region segmentation is to partition the image into two regions with distinct average intensities. The algorithm seeks to find a contour that minimizes a functional, leading to an optimal separation of the image into these two regions. For the use of the Chan-Vese method only greyscale images will be considered, as the MRI scans that will be tested are all in greyscale.

2.1 Simplified Mumford-Shah model

Let u_0 denote a greyscale image. The goal is to find a curve $C(s) : [0, 1] \rightarrow \mathbb{R}^2$ that fit closely the boundaries of the objects in the image. Using this curve, the image can be segmented in two pieces. One which will contain the objects in the image, that part is found inside the curve. The other part will contain the background and is found outside the curve. Starting with an arbitrary curve, a narrow band around the curve will be evolved subject to constraints. The goal is to move the curve towards the boundaries of the objects to obtain the optimal curve to segment the image into objects and background.

The optimal curve will be founded by minimization of a functional, based on the Mumford-Shah functional [10]:

$$\arg \min_{u, C} F(u, C) = \mu \text{Length}(C) + \lambda \int_{\Omega} (u_0(x) - u(x))^2 dx + \int_{\Omega \setminus C} |\nabla u(x)|^2 dx. \quad (2.1)$$

In this functional, the domain of the image is defined as Ω , the curve that will be evolved is defined as C and u is a piecewise-smooth function. The functional consist of three main terms, each serving a distinct purpose. The first term acts as a length penalty, ensuring the smoothness of the curve C . This term minimizes the total length of the curve, favouring smoother curves over highly irregular ones.

The second term encourages the function u to closely approximate the original image u_0 . This fidelity term measures the discrepancy between the function u and the original image u_0 . It ensures that the segmented regions closely match the actual image content.

The third term ensures that the function u itself remains smooth. This term prevents abrupt changes or discontinuities within the function and promotes gradual transitions within the segmented regions.

The Chan-Vese model is a particular case of this functional. In the Chan-Vese model, u is considered a piecewise constant function that can attain only two values (2.2), typically representing the object and the background in the image.

$$u(x) = \begin{cases} c_1, & \text{if } x \text{ is inside } C, \\ c_2, & \text{if } x \text{ is outside } C. \end{cases} \quad (2.2)$$

The curve C is a variable curve that defines the boundary between these two regions. The terms c_1 and c_2 represent the average values of u_0 inside and outside the curve C , respectively. Additionally, the Chan-Vese model includes a term that penalizes the enclosed area of the curve C . This area term regulates the size of the segmented regions. The functional of Chan-Vese method then becomes the functional:

$$\begin{aligned} \arg \min_{c_1, c_2, C} F(c_1, c_2, C) &= \mu \text{Length}(C) + \nu \text{Area}(\text{inside}(C)) \\ &+ \lambda_1 \int_{\text{inside}(C)} |u_0(x, y) - c_1|^2 dx dy + \lambda_2 \int_{\text{outside}(C)} |u_0(x, y) - c_2|^2 dx dy. \end{aligned} \quad (2.3)$$

The terms in functional impose the following restrictions when evolving the curve. The first term keeps the curve smooth by penalizing the length of the curve. The second term penalizes the enclosed area of the curve, controlling the size of the curve. Lastly, the third and fourth terms penalize the difference between the function u and the greyscale image u_0 . The pixel values in the regions must not deviate much from the mean value c_1 and c_2 of their respective regions. By finding the curve that minimizes this functional, the best curve for segmenting the image u_0 into two distinct regions is determined. This optimal curve effectively separates the object from the background, achieving a precise image segmentation.

2.2 Level Set Formulation

For this minimization problem, it is necessary to minimize over all possible curves C . This can be achieved using the level set formulation introduced by Osher and Sethian in 1988 [11]. The advantage of the level set formulation is that it allows for the minimization involving curves without the need to parameterize them explicitly. Instead of working directly with the curves, the level set method represents the curves C as the zero level set of an Lipschitz function $\phi : \Omega \rightarrow \mathbb{R}$. Furthermore, the sign of ϕ allows us to distinguish between the inside and outside of the curve C . Specifically, the regions where ϕ is positive correspond to the inside of C , and the regions where ϕ is negative correspond to the outside.

$$\begin{cases} C = \{\phi(x, y) = 0 \mid (x, y) \in \Omega\} \\ \text{inside}(C) = \{\phi(x, y) > 0 \mid (x, y) \in \Omega\}, \\ \text{outside}(C) = \{\phi(x, y) < 0 \mid (x, y) \in \Omega\}. \end{cases} \quad (2.4)$$

It is important to note that the level set representation for a given curve C is not unique. If the zero level set of ϕ represents a given curve C , then every other Lipschitz function $\tilde{\phi}$ where every point has the same sign as in ϕ i.e $\text{sign}(\phi(x, y)) = \text{sign}(\tilde{\phi}(x, y))$ for every $(x, y) \in \Omega$ will also represent the same curve C . The sign function is defined as:

$$\text{sign}(x) = \begin{cases} -1, & \text{if } x < 0, \\ 0, & \text{if } x = 0, \\ 1, & \text{if } x > 0. \end{cases} \quad (2.5)$$

To minimize the functional (2.3), the variable C will be replaced by the level set function ϕ . The terms of the functional (2.3) will be integrated over the whole domain. To define the functional in terms of ϕ , the Heaviside function H and the Dirac delta measure δ_0 are going to be needed. These functions are defined respectively by (2.6) and (2.7).

$$H(z) = \begin{cases} 1, & \text{if } z \geq 0 \\ 0, & \text{if } z < 0 \end{cases} \quad (2.6)$$

$$\delta_0(z) = \frac{d}{dz} H(z) \quad (2.7)$$

Since the Dirac delta measure as defined in (2.7) does not exist at $z = 0$, slightly regularized versions of the Heaviside function will be considered in the implementation. This regularization makes the function differentiable around $z = 0$. Because of this the δ_0 function is defined at $z = 0$. For the implementation, the regularizations in (2.8) and (2.9) will be considered. The regularizations take on slightly different values around $z = 0$ as can be seen in the graph from the paper "Active Contours Without Edges" [9]. To define the inside and outside of the curve, the Heaviside function H is used. The outside of the curve can be found where the Heaviside function is zero.




Figure 2.1: The two different regularizations of the Heaviside function.

$$H(z) = \begin{cases} 1, & \text{if } z > \varepsilon \\ 0, & \text{if } z < -\varepsilon \\ \frac{1}{2} \left[1 + \frac{z}{\varepsilon} + \frac{1}{\pi} \sin \left(\frac{\pi z}{\varepsilon} \right) \right], & \text{if } |z| \leq \varepsilon \end{cases} \quad (2.8)$$

$$H_2(z) = \frac{1}{2} \left(1 + \frac{2}{\pi} \arctan \left(\frac{z}{\varepsilon} \right) \right) \quad (2.9)$$

Using the Heaviside function and the fact that ϕ is Lipschitz continuous, the length of the curve can be expressed as the integral of $|\nabla\phi|$ along the curve [12]. The area enclosed by C can be determined by integrating $H(\phi)$ in the domain Ω . Therefore, with use of the Heaviside function, the terms of the Chan-Vese functional (2.3) can be rewritten as:

$$\begin{aligned} \text{Length}(C) &= \int_{\Omega} \delta_0(\phi(x, y)) |\nabla\phi(x, y)| dx dy, \\ \text{Area}(C) &= \int_{\Omega} H(\phi(x, y)) dx dy, \\ \int_{\phi>0} |u_0(x, y) - c_1|^2 dx dy &= \int_{\Omega} |u_0(x, y) - c_1|^2 H(\phi(x, y)) dx dy, \\ \int_{\phi<0} |u_0(x, y) - c_2|^2 dx dy &= \int_{\Omega} |u_0(x, y) - c_2|^2 (1 - H(\phi(x, y))) dx dy. \end{aligned} \quad (2.10)$$

The terms c_1 and c_2 can be defined as functions of ϕ using the level set formulation. As long as the curve C has a non-empty interior and exterior in Ω , the value of c_1 and c_2 can be expressed respectively as (2.11) and (2.12).

$$c_1(\phi) = \frac{\int_{\Omega} u_0(x, y) H(\phi(x, y)) dx dy}{\int_{\Omega} H(\phi(x, y)) dx dy} \quad (2.11)$$

$$c_2(\phi) = \frac{\int_{\Omega} u_0(x, y) (1 - H(\phi(x, y))) dx dy}{\int_{\Omega} (1 - H(\phi(x, y))) dx dy} \quad (2.12)$$

All the terms in (2.10) can be combined into one integral, as they are all integrated over the same area. This will result in the level set formulation for the Chan-Vese functional (2.3).

$$\begin{aligned} \arg \min_{c_1, c_2, \phi} F(c_1, c_2, \phi) &= \int_{\Omega} \mu \delta_0(\phi(x, y)) |\nabla\phi(x, y)| + \nu H(\phi(x, y)) \\ &\quad + \lambda_1 |u_0(x, y) - c_1|^2 H(\phi(x, y)) + \lambda_2 |u_0(x, y) - c_2|^2 (1 - H(\phi(x, y))) dx dy \end{aligned} \quad (2.13)$$

2.3 Minimization

Now that the level set formulation of the functional is defined, the next step is to find the function ϕ that minimizes the functional. Minimizing the functional (2.13), will be done

using the Euler–Lagrange equation [13]. For the minimization, the c_1 and c_2 are fixed. The terms inside the integral of (2.13) form the Lagrangian functional $G(\phi)$:

$$G(\phi) = \mu\delta_0(\phi) |\nabla\phi| + \nu H(\phi) + \lambda_1 |u_0 - c_1|^2 H(\phi) + \lambda_2 |u_0 - c_2|^2 (1 - H(\phi)). \quad (2.14)$$

Suppose that the functional (2.13) finds its minimal value for the function $u(x, y)$. Let $f(x, y)$ be a smooth function that equals zero on the boundary of the domain Ω , and let ε be a variable. The solution u will be perturbed by adding εf to that solution. Define the function $\Phi(\varepsilon)$ as $\Phi(\varepsilon) = F(u + \varepsilon f)$. The total derivative of the function Φ with respect to ε will be calculated at the point $\varepsilon = 0$.

$$\frac{d}{d\varepsilon}\Phi(\varepsilon)\Big|_{\varepsilon=0} = \int_{\Omega} \frac{d}{d\varepsilon}G(u + \varepsilon f)\Big|_{\varepsilon=0} dx dy. \quad (2.15)$$

This process involves differentiating the Lagrangian function $G(\phi)$ in two parts. First for the length penalty term inside $G(\phi)$ and then for the other terms inside $G(\phi)$. It will be assumed that $\frac{d\phi}{dn} = 0$ on the boundary of the image.

$$\begin{aligned} & \int_{\Omega} \frac{d}{d\varepsilon} \left[\mu\delta_0(u) |\nabla u + \varepsilon f| \right] \Big|_{\varepsilon=0} dx dy = \\ & \int_{\Omega} \frac{d}{d\varepsilon} \left[\mu\delta_0(u + \varepsilon f) |\nabla(u + \varepsilon f)| \right] \Big|_{\varepsilon=0} + \left[\mu\delta_0(u + \varepsilon f) \frac{d}{d\varepsilon} |\nabla u + \varepsilon f| \right] \Big|_{\varepsilon=0} dx dy = \\ & \int_{\Omega} \mu\delta'_0(u) f |\nabla u| + \mu\delta_0(u) \frac{\frac{d}{d\varepsilon} \left[\left(\frac{d}{dx}(u + \varepsilon f) \right)^2 + \left(\frac{d}{dy}(u + \varepsilon f) \right)^2 \right] \Big|_{\varepsilon=0}}{2 |\nabla u|} dx dy = \\ & \int_{\Omega} \mu\delta'_0(u) f |\nabla u| + \mu\delta_0(u) \frac{\nabla u \cdot \nabla f}{|\nabla u|} dx dy \end{aligned} \quad (2.16)$$

The product rule of divergence states:

$$\nabla \cdot (g_1 g_2) = g_1 \nabla \cdot g_2 + g_2 \cdot \nabla g_1. \quad (2.17)$$

Let $g_2 = \frac{\nabla u}{|\nabla u|}$ and $g_1 = \mu\delta_0(u) f$. This gives that:

$$\begin{aligned} \nabla \cdot \left(\mu\delta_0(u) f \frac{\nabla u}{|\nabla u|} \right) &= \mu\delta_0(u) f \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) + \frac{\nabla u}{|\nabla u|} \cdot \nabla (\mu\delta_0(u) f) \\ &= \mu\delta_0(u) f \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) + \frac{\nabla u}{|\nabla u|} \cdot (\mu\delta'_0(u) f \nabla u + \mu\delta_0(u) \nabla f), \\ &= \mu\delta_0(u) f \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) + \mu\delta'_0(u) f \frac{|\nabla u|^2}{|\nabla u|} + \mu\delta_0(u) \frac{\nabla u \cdot \nabla f}{|\nabla u|}. \\ &= \mu\delta_0(u) f \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) + \mu\delta'_0(u) f |\nabla u| + \mu\delta_0(u) \frac{\nabla u \cdot \nabla f}{|\nabla u|}. \end{aligned} \quad (2.18)$$

Therefore it holds that:

$$\mu\delta_0(u)\frac{\nabla u \cdot \nabla f}{|\nabla u|} = \nabla \cdot \left(\mu\delta_0(u)f\frac{\nabla u}{|\nabla u|} \right) - \mu\delta_0(u)f\nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) - \mu\delta'_0(u)f|\nabla u|. \quad (2.19)$$

Substituting this result back into the integral gives the following:

$$\int_{\Omega} \frac{d}{d\varepsilon} \left[\mu\delta_0(u) |\nabla u + \varepsilon f| \right] \Big|_{\varepsilon=0} = \int_{\Omega} \nabla \cdot \left(\mu\delta_0(u)f\frac{\nabla u}{|\nabla u|} \right) - \mu\delta_0(u)f\nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) dx dy. \quad (2.20)$$

According to the Divergence Theorem, we have:

$$\int_{\Omega} \nabla \cdot \left(\mu\delta_0(u)f\frac{\nabla u}{|\nabla u|} \right) dx dy = \int_{\partial\Omega} \left(\mu\delta_0(u)f\frac{\nabla u}{|\nabla u|} \right) \cdot ndS. \quad (2.21)$$

f was chosen to be an arbitrary smooth function that equals zero on the boundary of the domain. Since $f(x, y)$ vanishes on the boundary, the integrand in the boundary integral is zero wherever it is evaluated. This leaves us with:

$$\int_{\Omega} \frac{d}{d\varepsilon} \left[\mu\delta_0(u) |\nabla u + \varepsilon f| \right] \Big|_{\varepsilon=0} = \int_{\Omega} -\mu\delta_0(u)f\nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) dx dy. \quad (2.22)$$

For the other terms, the following derivation is found:

$$\begin{aligned} \int_{\Omega} \frac{d}{d\varepsilon} \Big|_{\varepsilon=0} \left[\nu H(u+\varepsilon f) + \lambda_1 |u_0(x, y) - c_1|^2 H(u+\varepsilon f) + \lambda_2 |u_0(x, y) - c_2|^2 (1 - H(u + \varepsilon f)) \right] = \\ \int_{\Omega} \delta_0(u)f (\nu + \lambda_1 |u_0(x, y) - c_1|^2 - \lambda_2 |u_0(x, y) - c_2|^2). \end{aligned} \quad (2.23)$$

Combining the two parts again gives the condition (2.24) on u . Since the functional attains a minimum value at u , any perturbation of u increases u . This makes it so that the function Φ reaches an extreme value at $\varepsilon = 0$. This means that the derivative in that point should be equal to zero.

$$\begin{aligned} \int_{\Omega} \frac{d}{d\varepsilon} \Big|_{\varepsilon=0} G(u+\varepsilon f) dx dy = \int_{\Omega} \delta_0(u)f \left(-\mu\nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) + \nu + \lambda_1 |u_0(x, y) - c_1|^2 \right) \\ - \delta_0(u)f (\lambda_2 |u_0(x, y) - c_2|^2) dx dy = 0 \end{aligned} \quad (2.24)$$

Since this condition holds for all smooth functions f that vanish on the boundary of Ω , it follows from Dubois-Reymond's lemma [13] that the derivative of G with respect to ε should be zero. This implies that if a solution u exists and the boundary condition that $\frac{\partial u}{\partial \hat{n}} = 0$ is satisfied, then this solution must fulfil the differential equation (2.25). The existence of the solution has been proved in [10].

$$\begin{cases} \delta_{\varepsilon}(u) \left(\mu\nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) - \nu - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) = 0 \\ \frac{\delta_0(u)f}{|\nabla u|} \frac{\partial u}{\partial \hat{n}} = 0 \end{cases} \quad (2.25)$$

To find this solution, the differential equation (2.25) needs to be solved. One effective method for solving this type of equation is gradient descent. Due to the use of gradient descent, a local minima will be found for the functional. To implement gradient descent an artificial time $t \geq 0$ will be introduced into the function ϕ . Rather than solving the differential equation (2.25) directly, the stationary solution of the following differential equation will be found:

$$\frac{\partial \phi}{\partial t} = \delta_\varepsilon(\phi(t, x, y)) \left(\mu \nabla \cdot \left(\frac{\nabla \phi(t, x, y)}{|\nabla \phi(t, x, y)|} \right) - v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) = 0. \quad (2.26)$$

The process begins by defining the initial curve at $t = 0$ as $\phi_0(x, y)$. This gives the initial condition $\phi_0(0, x, y) = \phi_0(x, y)$. The problem changes to solving the partial differential equation (2.27) with the specified initial and boundary conditions.

$$\begin{cases} \frac{\partial \phi}{\partial t} = \delta_\varepsilon(\phi) \left(\mu \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) - v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) = 0 & \text{in } (0, \infty) \times \Omega, \\ \phi(0, x, y) = \phi_0(x, y) & \text{in } \Omega, \\ \frac{\delta_\varepsilon \phi}{|\nabla \phi|} \frac{\partial \phi}{\partial \hat{n}} = 0 & \text{on } \partial \Omega. \end{cases} \quad (2.27)$$

Implementation

Solving Chan-Vese minimization problem (2.27) will be done numerically with the use of Fortran. Fortran is a powerful tool for scientific and numerical computing and is known for its high performance in numerical computations. The language has been optimized for mathematical operations and array handling. In this chapter, the numerical implementation of the Chan-Vese segmentation will be discussed.

3.1 Spatial discretization

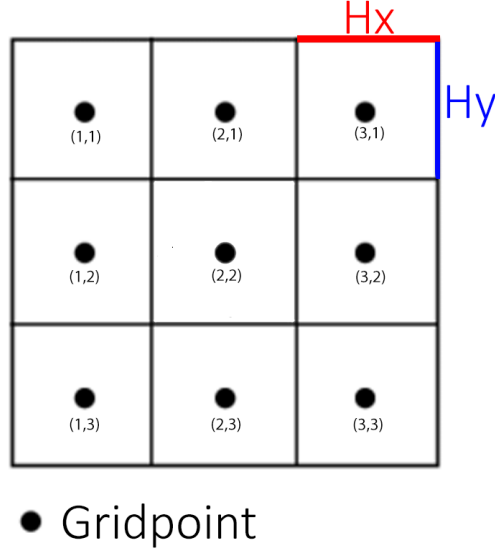
To numerically solve the Chan-Vese minimization problem (2.27), the partial differential equation needs to be first discretized using the finite difference method (FDM). FDM approximates solutions to differential equations by discretizing the problem's domain into a finite grid of points containing the image values. Derivatives are approximated using differences between function values at these discrete points.

$$\begin{cases} \delta_\varepsilon(\phi) \left(\mu \nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) - v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) = 0 \text{ in } \Omega, \\ \phi(0, x, y) = \phi_0(x, y) \\ \frac{\delta_\varepsilon}{|\nabla \phi|} \frac{\partial \phi}{\partial \hat{n}} = 0 \text{ on } \partial \Omega. \end{cases} \quad (3.1)$$

First, a grid will be defined. It is assumed that the greyscale image u_0 is on a rectangular, equidistant cell-centred grid $\Omega = 0, \dots, M \times 0, \dots, N$, where M denotes the amount of pixels in the horizontal direction and N denotes the number of pixels in vertical direction of the image. The pixels are defined to be horizontally and vertically 1 unit apart, $h_x = h_y = 1$. As they are both chosen to be 1, the term h will be used to refer to h_x and h_y . The grid points are concentrated at the centre of the pixels. The grid points are denoted as follows:

$$\phi_{i,j} = \phi\left(i - \frac{1}{2}h_x, j - \frac{1}{2}h_y\right) \quad i \in \{1, \dots, N\}, \quad j \in \{1, \dots, M\}. \quad (3.2)$$

To find the spatial discretization of equation (3.1), the the term $\nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right)$ need to be discretized. The derivatives $\frac{\partial}{\partial x} \left(\frac{\phi}{|\nabla \phi|} \right)$ and $\frac{\partial}{\partial y} \left(\frac{\phi}{|\nabla \phi|} \right)$ will be approximated using central

Figure 3.1: Control volumes of a 3×3 grid.

differences with a step size of $\frac{h}{2}$ around $\phi_{i,j}$.

$$\nabla \cdot \left(\frac{\nabla \phi}{|\nabla \phi|} \right) = \frac{\partial}{\partial x} \left(\frac{\partial}{\partial x} \left(\frac{\phi}{|\nabla \phi|} \right) \right) + \frac{\partial}{\partial y} \left(\frac{\partial}{\partial y} \left(\frac{\phi}{|\nabla \phi|} \right) \right). \quad (3.3)$$

$$\frac{\partial}{\partial x} \left(\frac{\partial}{\partial x} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \right) = \frac{\frac{\partial}{\partial x} \left(\frac{\nabla \phi_{i+\frac{1}{2},j}}{|\nabla \phi_{i+\frac{1}{2},j}|} \right) - \frac{\partial}{\partial x} \left(\frac{\nabla \phi_{i-\frac{1}{2},j}}{|\nabla \phi_{i-\frac{1}{2},j}|} \right)}{h} + \mathcal{O}(h^2) \quad (3.4)$$

$$\frac{\partial}{\partial y} \left(\frac{\partial}{\partial y} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \right) = \frac{\frac{\partial}{\partial y} \left(\frac{\nabla \phi_{i,j+\frac{1}{2}}}{|\nabla \phi_{i,j+\frac{1}{2}}|} \right) - \frac{\partial}{\partial y} \left(\frac{\nabla \phi_{i,j-\frac{1}{2}}}{|\nabla \phi_{i,j-\frac{1}{2}}|} \right)}{h} + \mathcal{O}(h^2) \quad (3.5)$$

Central differences will again be used to discretize $\frac{\partial}{\partial x} (\nabla \phi)$ and $\frac{\partial}{\partial y} (\nabla \phi)$ with a step size of $\frac{h}{2}$. By applying this discretization to the equation in (3.1), the equation in (3.10) is found. Note that this equation has an error of order $\mathcal{O}(h^2)$.

$$\frac{\partial}{\partial x} \left(\frac{\phi_{i+\frac{1}{2},j}}{|\nabla \phi_{i+\frac{1}{2},j}|} \right) = \frac{1}{h} \frac{\phi_{i+1,j} - \phi_{i,j}}{|\nabla \phi_{i+\frac{1}{2},j}|} + \mathcal{O}(h^2), \quad (3.6)$$

$$\frac{\partial}{\partial x} \left(\frac{\phi_{i-\frac{1}{2},j}}{|\nabla \phi_{i-\frac{1}{2},j}|} \right) = \frac{1}{h} \frac{\phi_{i,j} - \phi_{i-1,j}}{|\nabla \phi_{i-\frac{1}{2},j}|} + \mathcal{O}(h^2), \quad (3.7)$$

$$\frac{\partial}{\partial y} \left(\frac{\phi_{i,j+\frac{1}{2}}}{|\nabla \phi_{i,j+\frac{1}{2}}|} \right) = \frac{1}{h} \frac{\phi_{i,j+1} - \phi_{i,j}}{|\nabla \phi_{i,j+\frac{1}{2}}|} + \mathcal{O}(h^2), \quad (3.8)$$

$$\frac{\partial}{\partial y} \left(\frac{\phi_{i,j-\frac{1}{2}}}{|\nabla \phi_{i,j-\frac{1}{2}}|} \right) = \frac{1}{h} \frac{\phi_{i,j} - \phi_{i,j-1}}{|\nabla \phi_{i,j-\frac{1}{2}}|} + \mathcal{O}(h^2). \quad (3.9)$$

$$\begin{aligned} & \frac{\mu}{h^2} \frac{\phi_{i,j+1} - \phi_{i,j}}{|\nabla \phi_{i,j+\frac{1}{2}}|} \delta_\varepsilon(\phi_{i,j}) + \frac{\mu}{h^2} \frac{\phi_{i+1,j} - \phi_{i,j}}{|\nabla \phi_{i+\frac{1}{2},j}|} \delta_\varepsilon(\phi_{i,j}) - \frac{\mu}{h^2} \frac{\phi_{i,j} - \phi_{i,j-1}}{|\nabla \phi_{i,j-\frac{1}{2}}|} \delta_\varepsilon(\phi_{i,j}) - \frac{\mu}{h^2} \frac{\phi_{i,j} - \phi_{i-1,j}}{|\nabla \phi_{i-\frac{1}{2},j}|} \delta_\varepsilon(\phi_{i,j}) \\ & \delta_\varepsilon(\phi_{i,j}) \left(-v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) + O(h^2) = 0 \quad i \in \{1, \dots, N\}, \quad j \in \{1, \dots, M\} \end{aligned} \quad (3.10)$$

Discretizing the Neumann boundary condition around the boundaries: $\phi_{i,M+\frac{1}{2}}$, $\phi_{N+\frac{1}{2},j}$, $\phi_{\frac{1}{2},j}$ and $\phi_{i,\frac{1}{2}}$ using central differences with a step size of $\frac{h}{2}$ leads to the equations in (3.11). With the introduction of the virtual points $\phi_{i,M+1}$, $\phi_{N+1,j}$, $\phi_{0,j}$ and $\phi_{i,0}$, with $i \in \{1, \dots, N\}$, $j \in \{1, \dots, M\}$, it is found that on the boundary certain terms of the equation (3.10) are equal to zero. This is either because $\delta_\varepsilon(\phi_{i,j}) = 0$ on that part of the boundary or by the identities detailed in (3.12).

$$\begin{aligned} & \frac{\delta_\varepsilon(\phi_{i,j})}{|\nabla \phi_{\frac{1}{2},j}|} \frac{\phi_{1,j} - \phi_{0,j}}{h} + O(h^2) = 0, \\ & \frac{\delta_\varepsilon(\phi_{i,j})}{|\nabla \phi_{N+\frac{1}{2},j}|} \frac{\phi_{N+1,j} - \phi_{N,j}}{h} + O(h^2) = 0, \\ & \frac{\delta_\varepsilon(\phi_{i,j})}{|\nabla \phi_{i,\frac{1}{2}}|} \frac{\phi_{i,1} - \phi_{i,0}}{h} + O(h^2) = 0, \\ & \frac{\delta_\varepsilon(\phi_{i,j})}{|\nabla \phi_{i,M+\frac{1}{2}}|} \frac{\phi_{i,N+1} - \phi_{i,M}}{h} + O(h^2) = 0. \end{aligned} \quad (3.11)$$

$$\begin{cases} \phi_{1,j} = \phi_{0,j}, & \text{if } \delta_\varepsilon(\phi_{1,j}) \neq 0 \\ \phi_{N,j} = \phi_{N+1,j}, & \text{if } \delta_\varepsilon(\phi_{N,j}) \neq 0 \\ \phi_{i,1} = \phi_{i,0}, & \text{if } \delta_\varepsilon(\phi_{i,1}) \neq 0 \\ \phi_{i,M} = \phi_{i,M+1}, & \text{if } \delta_\varepsilon(\phi_{i,M}) \neq 0 \end{cases} \quad (3.12)$$

On the boundary on the south, the equation reduces to:

$$\begin{aligned} & -\mu \frac{\phi_{i+1,M} - \phi_{i,M}}{|\nabla \phi_{i+\frac{1}{2},M}|} \delta_\varepsilon(\phi_{i,M}) + \mu \frac{\phi_{i,M} - \phi_{i,M-1}}{|\nabla \phi_{i,M-\frac{1}{2}}|} \delta_\varepsilon(\phi_{i,M}) + \mu \frac{\phi_{i,M} - \phi_{i-1,M}}{|\nabla \phi_{i-\frac{1}{2},M}|} \delta_\varepsilon(\phi_{i,M}) = \\ & h^2 \delta_\varepsilon(\phi_{i,M}) \left(-v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) + O(h^3) \quad i \in \{1, \dots, N\} \end{aligned} \quad (3.13)$$

On the boundary on the east, the equation reduces to:

$$\begin{aligned} & -\mu \frac{\phi_{N,j+1} - \phi_{N,j}}{|\nabla \phi_{N,j+\frac{1}{2}}|} \delta_\varepsilon(\phi_{N,j}) + \mu \frac{\phi_{N,j} - \phi_{N,j-1}}{|\nabla \phi_{N,j-\frac{1}{2}}|} \delta_\varepsilon(\phi_{N,j}) + \mu \frac{\phi_{N,j} - \phi_{N-1,j}}{|\nabla \phi_{N-\frac{1}{2},j}|} \delta_\varepsilon(\phi_{N,j}) = \\ & h^2 \delta_\varepsilon(\phi_{N,j}) \left(-v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) + O(h^3) \quad j \in \{1, \dots, M\} \end{aligned} \quad (3.14)$$

On the boundary on the north, the equation reduces to:

$$\begin{aligned} & -\mu \frac{\phi_{i,1} - \phi_{i,0}}{|\nabla \phi_{i,\frac{1}{2}}|} \delta_\varepsilon(\phi_{i,0}) - \mu \frac{\phi_{i+1,0} - \phi_{i,0}}{|\nabla \phi_{i+\frac{1}{2},0}|} \delta_\varepsilon(\phi_{i,0}) + \mu \frac{\phi_{i,0} - \phi_{i-1,0}}{|\nabla \phi_{i-\frac{1}{2},0}|} \delta_\varepsilon(\phi_{i,0}) = \\ & h^2 \delta_\varepsilon(\phi_{i,0}) \left(-v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) + O(h^3) \quad i \in \{1, \dots, N\} \end{aligned} \quad (3.15)$$

On the boundary on the west, the equation reduces to:

$$\begin{aligned}
& -\mu \frac{\phi_{0,j+1} - \phi_{0,j}}{|\nabla \phi_{0,j+\frac{1}{2}}|} \delta_\varepsilon(\phi_{0,j}) - \mu \frac{\phi_{1,j} - \phi_{0,j}}{|\nabla \phi_{\frac{1}{2},j}|} \delta_\varepsilon(\phi_{0,j}) + \mu \frac{\phi_{0,j} - \phi_{0,j-1}}{|\nabla \phi_{0,j-\frac{1}{2}}|} \delta_\varepsilon(\phi_{0,j}) = \\
& h^2 \delta_\varepsilon(\phi_{0,j}) \left(-v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) + O(h^3) \quad j \in \{1, \dots, M\}
\end{aligned} \tag{3.16}$$

The magnitude of the gradient vector will be discretized using central differences. As these magnitudes are not known in the points $(i \pm \frac{1}{2}, j)$, and $(i, j \pm \frac{1}{2})$, the value of those coefficients will be determined by taking the average of neighbouring grid points. The magnitude of the gradient is regularized by including a small parameter ρ in the denominator to prevent division by zero. In the implementation, ρ is taken as 10^{-6} .

$$\frac{1}{|\phi_{i,j}|} = \frac{1}{\sqrt{\left(\frac{\phi_{i+1,j} - \phi_{i-1,j}}{h}\right)^2 + \left(\frac{\phi_{i,j+1} - \phi_{i,j-1}}{h}\right)^2 + \rho}} \tag{3.17}$$

$$\begin{aligned}
\frac{1}{|\phi_{i+\frac{1}{2},j}|} &= \frac{1}{2} \left(\frac{1}{|\phi_{i,j}|} + \frac{1}{|\phi_{i+1,j}|} \right) \\
\frac{1}{|\phi_{i-\frac{1}{2},j}|} &= \frac{1}{2} \left(\frac{1}{|\phi_{i,j}|} + \frac{1}{|\phi_{i-1,j}|} \right) \\
\frac{1}{|\phi_{i,j+\frac{1}{2}}|} &= \frac{1}{2} \left(\frac{1}{|\phi_{i,j}|} + \frac{1}{|\phi_{i,j+1}|} \right) \\
\frac{1}{|\phi_{i,j-\frac{1}{2}}|} &= \frac{1}{2} \left(\frac{1}{|\phi_{i,j}|} + \frac{1}{|\phi_{i,j-1}|} \right)
\end{aligned} \tag{3.18}$$

The terms c_1 and c_2 can be numerically approximated as found in (3.19):

$$\begin{aligned}
c_1(\phi) &= \frac{\sum_{i=1}^N \sum_{j=1}^N u_{0,i,j} H(\phi_{i,j})}{\sum_{i=1}^N \sum_{j=1}^N H(\phi_{0,i,j})}, \\
c_2(\phi) &= \frac{\sum_{i=1}^N \sum_{j=1}^N u_{0,i,j} (1 - H(\phi_{i,j}))}{\sum_{i=1}^N \sum_{j=1}^N (1 - H(\phi_{0,i,j}))}.
\end{aligned} \tag{3.19}$$

With this discretization, the partial differential equation found in (3.1) can be rewritten using the non-linear operator $A(w)$ and $f(w)$.

$$\begin{cases} \mathbf{A}(w)w = \mathbf{f}(w) \\ \phi(0, x, y) = \phi_0(x, y) \\ \frac{\delta_\varepsilon}{|\nabla \phi|} \frac{\partial \phi}{\partial \bar{n}} = 0 \text{ on } \partial\Omega. \end{cases} \tag{3.20}$$

To write the equation in this form, the two-dimensional grid of the pixels in the image has to be converted in such a way that it becomes a vector. For this implementation the pixels will be ordered in a horizontal organization. The numbering of the pixels will start in the top left corner. A grid point $\phi_{i,j}$ will be found in the one-dimensional vector w at the place $w_{i-N(j-1)}$. Let w_k be the value of the grid point k , then w_{k+1} is the pixel right

of w_k and w_{k+N} is the pixel below w_k . The operator $\mathbf{A}(w)$ is a $NM \times NM$ and has the structure:

$$\mathbf{A}(w) = \begin{cases} \mathbf{A}_{i,i+1}(w) = \frac{\mu\delta_\varepsilon(w_i)}{2h^2} \left(\frac{1}{|w_i|} + \frac{1}{|w_{i+1}|} \right), & \text{if } i \bmod (N) \neq 1 \\ \mathbf{A}_{i,i-1}(w) = \frac{\mu\delta_\varepsilon(w_i)}{2h^2} \left(\frac{1}{|w_i|} + \frac{1}{|w_{i-1}|} \right), & \text{if } i \bmod (N) \neq 0 \\ \mathbf{A}_{i,i+M}(w) = \frac{\mu\delta_\varepsilon(w_i)}{2h^2} \left(\frac{1}{|w_i|} + \frac{1}{|w_{i+N}|} \right), & \text{if } i \leq M(N-1) \\ \mathbf{A}_{i,i-M}(w) = \frac{\mu\delta_\varepsilon(w_i)}{2h^2} \left(\frac{1}{|w_i|} + \frac{1}{|w_{i-N}|} \right), & \text{if } i > N \\ \mathbf{A}_{i,i}(w) = -\mathbf{A}_{i,i+1} - \mathbf{A}_{i,i-1} - \mathbf{A}_{i,i+M} - \mathbf{A}_{i,i-M}, \\ \mathbf{A}_{i,j}(w) = 0, & \text{in all other entries.} \end{cases} \quad (3.21)$$

where $i \in \{1, \dots, NM\}$. The operator $\mathbf{f}(w)$ is of length NM where each entry is defined as:

$$\mathbf{f}_i(w) = \delta_\varepsilon(w_i) \left(-v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right). \quad (3.22)$$

3.2 Time integration

To perform time integration for the partial differential equation (PDE) given by equation (2.27), various methods can be considered, ranging from explicit methods like Explicit Euler to implicit methods such as Implicit Euler. Explicit methods might initially seem ideal due to their lower computational effort compared to implicit methods. However, for this specific equation, the stability of explicit methods poses a significant concern. For instance, with the Forward Euler method, a stable timestep is not always guaranteed. The Forward Euler method achieves stability for real eigenvalues if the inequality (3.23) holds for all eigenvalues λ_{ij} .

$$-2 \leq \lambda \Delta t \leq 0 \quad (3.23)$$

By applying the Gershgorin circle theorem, it can be determined that the eigenvalues λ of the operator A , as defined in equation (3.21), are bounded within the bounds specified (3.24).

$$\left\{ \begin{array}{l}
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right)| \leq \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right), \text{ if } i = N, j = 0 \\
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} \right)| \leq \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} \right), \text{ if } i = N, j = M \\
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right)| \leq \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right), \text{ if } i = 0, j = 0 \\
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} \right)| \leq \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} \right), \text{ if } i = 0, j = M \\
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right)| \leq \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right), \text{ if } i = 0 \\
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right)| \leq \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right), \text{ if } j = M \\
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right)| \leq \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right), \text{ if } i = N \\
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} \right)| \leq \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} \right), \text{ if } j = 0 \\
|\lambda + \frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right)| \leq \\
\frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right), \text{ otherwise}
\end{array} \right. \quad (3.24)$$

Given that $|\nabla\phi|$ is non-negative, the Gershgorin circle theorem leads to the following bounds for the eigenvalues λ :

$$\min_{i,j} \left[-\frac{2\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right) \right] \leq \lambda \leq 0. \quad (3.25)$$

The upper bound for λ is non-positive and the time steps Δt are positive. This leads to that the inequality $\lambda\Delta t \leq 0$ is always satisfied. For the lower bound the following inequality must hold:

$$-2 \leq \min_{i,j} \left[-\frac{2\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right) \right] \Delta t. \quad (3.26)$$

This means that for the time step dt to be stable, the following must hold:

$$\min_{i,j} \left[\frac{\mu\delta_\varepsilon(\phi_{i,j})}{h^2} \left(\frac{1}{|\nabla\phi_{i+\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j-\frac{1}{2}}|} + \frac{1}{|\nabla\phi_{i-\frac{1}{2},j}|} + \frac{1}{|\nabla\phi_{i,j+\frac{1}{2}}|} \right) \right] \Delta t \leq 1. \quad (3.27)$$

This indicates that the maximum stable timestep is dependent on the magnitude of the gradients of ϕ . The points of ϕ that have the gradients smallest gradients dictate the stable time step that can be chosen. The magnitude of the gradient (3.17) is regularized by ρ and ρ is taken as 10^{-6} . The fact that the value of the magnitude is at least ρ implies the existence of a stable time step. In the worst case scenario where there is a point on the curve with $|\nabla\phi_{i,j}| = \rho$, Δt can only have values of the magnitude $\frac{\rho}{\mu}$. This means that the stable time step might be small, resulting in a scenario where convergence requires a substantial number of iterations.

3.3 Additive Operator splitting

As explicit methods are not ideal for this problem, the choice was made to use the semi-implicit Additive Operator Splitting (AOS) method [14]. AOS is a numerical technique used to solve partial differential equations by decomposing the discretization matrix in parts. The AOS scheme has several advantages. The semi-implicit formulation is stable for all time steps and as such it allows for bigger time steps. Furthermore, the algorithm can be executed in parallel along the x and y directions. It was found in the spatial discretization that for each point ϕ , the derivative with respect to the artificial time t is taken as:

$$\begin{aligned} \frac{d\phi_{i,j}}{dt} = & \mu \frac{\phi_{i,j+1} - \phi_{i,j}}{|\nabla\phi_{i,j+\frac{1}{2}}|} \delta_\varepsilon(\phi_{i,j}) + \mu \frac{\phi_{i+1,j} - \phi_{i,j}}{|\nabla\phi_{i+\frac{1}{2},j}|} \delta_\varepsilon(\phi_{i,j}) - \mu \frac{\phi_{i,j} - \phi_{i,j-1}}{|\nabla\phi_{i,j-\frac{1}{2}}|} \delta_\varepsilon(\phi_{i,j}) - \mu \frac{\phi_{i,j} - \phi_{i-1,j}}{|\nabla\phi_{i-\frac{1}{2},j}|} \delta_\varepsilon(\phi_{i,j}) \\ & + h^2 \delta_\varepsilon(\phi_{i,j}) \left(-v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 \right) + O(h^3) \quad i \in \{1, \dots, N\}, \quad j \in \{1, \dots, M\}. \end{aligned} \quad (3.28)$$

To simplify the differential equation, two terms will be introduced: $a_{1,i,j}$ and $a_{2,i,j}$:

$$a_{1,i,j} = \frac{1}{2} \left(\frac{1}{|\phi_{i,j}|} + \frac{1}{|\phi_{i+1,j}|} \right) \quad \text{and} \quad a_{2,i,j} = \frac{1}{2} \left(\frac{1}{|\phi_{i,j}|} + \frac{1}{|\phi_{i-1,j}|} \right). \quad (3.29)$$

Let \mathbf{f} again be equal to $-v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2$. With the use of the terms $a_{1,i,j}$ and $a_{2,i,j}$, the time step discretization for the point inside the boundary reduces to (3.30), where the superscript of ϕ denotes the time step.

$$\begin{aligned} \frac{\phi_{i,j}^{k+1} - \phi_{i,j}^k}{dt} + O(dt) = & \frac{\mu}{h^2} a_{1,i,j} (\phi_{i+1,j}^{k+1} - \phi_{i,j}^{k+1}) - \frac{\mu}{h^2} a_{1,i-1,j} (\phi_{i,j}^{k+1} - \phi_{i-1,j}^{k+1}) \\ & + \frac{\mu}{h^2} a_{2,i,j} (\phi_{i,j+1}^{k+1} - \phi_{i,j}^{k+1}) - \frac{\mu}{h^2} a_{2,i,j-1} (\phi_{i,j}^{k+1} - \phi_{i,j-1}^{k+1}) \\ & - v - \lambda_1(u_0 - c_1)^2 + \lambda_2(u_0 - c_2)^2 + O(h^3). \end{aligned} \quad (3.30)$$

With the timestep discretization, the partial differential equation (2.27) can be written as:

$$\begin{aligned} \frac{w^{k+1} - w^k}{dt} = & \sum_{j=1}^2 \mathbf{A}_j(w^k) w^{k+1} + \mathbf{f}(w^k), \\ w^{k+1} = & [I - dt \sum_{j=1}^2 \mathbf{A}_j(w^k)]^{-1} (w^k + dt \mathbf{f}(w^k)). \end{aligned} \quad (3.31)$$

For this, three non-linear operators $\mathbf{A}_1(w)$ (3.32), $\mathbf{A}_2(w)$ (3.33) and $\mathbf{f}(w)$ are used. $\mathbf{f}(w)$ is defined as in (3.22).

$$\mathbf{A}_1(w) = \begin{cases} \mathbf{A}_{1,i,i+1}(w) = \frac{\mu \delta_\varepsilon(w_i)}{2h^2} \left(\frac{1}{|w_i|} + \frac{1}{|w_{i+1}|} \right), & \text{if } i \bmod (N) \neq 1 \\ \mathbf{A}_{1,i,i-1}(w) = \frac{\mu \delta_\varepsilon(w_i)}{2h^2} \left(\frac{1}{|w_i|} + \frac{1}{|w_{i-1}|} \right), & \text{if } i \bmod (N) \neq 0 \\ \mathbf{A}_{1,i,i}(w) = -\mathbf{A}_{1,i,i+1} - \mathbf{A}_{1,i,i-1} \\ \mathbf{A}_{1,i,j}(w) = 0, & \text{in all other entries} \end{cases} \quad (3.32)$$

$$\mathbf{A}_2(w) = \begin{cases} \mathbf{A}_{2_{i,i+1}}(w) = \frac{\mu\delta_\varepsilon(w_i)}{2h^2} \left(\frac{1}{|w_i|} + \frac{1}{|w_{i+M}|} \right), & \text{if } i \leq M(N-1) \\ \mathbf{A}_{2_{i,i-1}}(w) = \frac{\mu\delta_\varepsilon(w_i)}{2h^2} \left(\frac{1}{|w_i|} + \frac{1}{|w_{i-N}|} \right), & \text{if } i > N \\ \mathbf{A}_{2_{i,i}}(w) = -\mathbf{A}_{2_{i,i+1}} - \mathbf{A}_{2_{i,i-1}} \\ \mathbf{A}_{2_{i,j}}(w) = 0, & \text{in all other entries} \end{cases} \quad (3.33)$$

Calculating the inverse in (3.31) would lead to a large computation effort. In order to address this problem, a modification of this semi-implicit scheme is considered, the additive operator splitting (AOS) scheme. The semi-implicit scheme and the AOS scheme have the same first-order Taylor expansions in dt and both schemes are consistent to the original equation [14]. This will result in the scheme:

$$w^{k+1} = \frac{1}{2} \sum_{j=1}^2 [I - 2dt\mathbf{A}_j(w^k)]^{-1} (w^k + dt\mathbf{f}(w^k)). \quad (3.34)$$

The operators:

$$B_m = [I - 2dt\mathbf{A}_m(w^k)], \quad (3.35)$$

describe the one-dimensional diffusion processes along the x and y axes. It will be shown that they can be inverted in an efficient and stable way by the Thomas algorithm [14]. First it has to be shown that they are invertible. To do this it will be shown that the operators are strictly diagonally dominant. The operators are strictly diagonally dominant if:

$$|B_{m_{i,i}}| > \sum_{j \neq i} |B_{m_{i,j}}| \quad \forall i \in \{1, \dots, NM\}. \quad (3.36)$$

The value of $B_{m_{i,i}}$ is:

$$B_{m_{i,i}} = 1 + 2dt (\mathbf{A}_{m_{i,i+1}} + \mathbf{A}_{m_{i,i-1}}) \quad (3.37)$$

The terms in $B_{m_{i,i}}$ are all non-negative, thus $B_{m_{i,i}}$ has always a value of at least one. The value of $|B_{m_{i,j}}|$ is:

$$|B_{m_{i,j}}| = 2dt\mathbf{A}_{m_{i,j}}(w^k). \quad (3.38)$$

Only for $j \in \{i-1, i+1\}$ does $\mathbf{A}_{m_{i,j}}$ have a non-zero value. From this it can be found that the values on the diagonal are 1 more than the sum of the values of the other entry in the row. For that reason, the inequality in 3.36 holds. This means that the operators are strictly diagonally dominant.

Now that it is shown that the tridiagonal systems are invertible, the Thomas algorithm [15] can be used to invert the matrices. The benefits of the Thomas algorithm are that it is stable for every strictly diagonally dominant system matrix and that the algorithm is very efficient. In the algorithm only three diagonals of B_j are used for computation. The algorithm only requires $5MN - 4$ multiplications/divisions and $3MN - 3$ subtraction for each matrix B_j . Thus, the algorithm is solved in linear time [14].

To solve the linear system (3.39) for w^{k+1} , the Thomas algorithm takes of three steps. First the LU decomposition of B_j has to be found. The matrix will be decomposed into the product of a lower triangular L matrix and an upper triangular matrix U . After that the system $Ly = w^k + dt\mathbf{f}$ is solved for y and at last $Uw^{k+1} = y$ is solved for w^{k+1} . This will be done for both the x and y direction and by taking the average of the two found results for w^{k+1} the w for the next time step is found.

$$B_j(w^k)w^{k+1} = w^k + dt\mathbf{f}(w^k) \quad (3.39)$$

Results

4.1 Chosen parameters

The Chan-Vese segmentation method will be tested on a variety of images. MRI images of a papaya and a bell pepper were obtained by Ksenia Slepova in 2023 at the MRI laboratory of MUST (Mbarara University of Science and Technology) using the low-cost MRI scanner. Additionally, MRI images of a corn and a melon were created by Andy Ellison for the collection of Magnetic Resonance Imaging of Foods, and the MRI image of a brain was sourced from the article "Brain imaging with portable low-field MRI" [16]. These images are provided in Appendix A.

In these experiments, we will examine the impact of the choice of the regularization of the Heaviside function, the effect of the image quality on the segmentation and the influence of the parameters on the segmentation. The effect of the image quality is assessed by applying the model to images that contain noise and to images that have been filtered. Unless specified otherwise, the parameters used in the experiments can be found in table (4.1).

Table 4.1: The standard parameters used in the experiments

Standard parameters	
Parameter	Value
μ	0.5
ν	0
λ_1	1
λ_2	1
dt	0.01
Initial contour	$\sin\left(\frac{x\pi}{5}\right)\sin\left(\frac{y\pi}{5}\right)$

The AOS method has to stop when the stationary solution is found. To do that, the method is terminated when the L^2 difference between w^{k+1} and w^k is below a certain tolerance level. For this implementation the tolerance level is chosen to be 10^{-6} . The images in the experiments below contain a red contour. This contour indicates where the segmentation boundary gained from the model.


4.2 Heaviside function

For the terms c_1 and c_2 found in (2.11) and (2.12) a Heaviside function is used in the computation of the value of those terms. Both regularizations of the Heaviside functions defined earlier (2.8) and (2.9) will be considered for the computation of the terms c_1 and c_2 . For convenience, they are stated again in (4.1) and (4.2). In H_1 an ε level of 10^{-4} is taken, while for H_2 the value of ε is chosen to be 1. The ε in H_1 was chosen to be significantly smaller to see if a much sharper regularization of the Heaviside function leads to different results. These regularizations indeed lead to different results when applied to the same images.

$$H_1(z) = \begin{cases} 1, & \text{if } z > \varepsilon \\ 0, & \text{if } z < -\varepsilon \\ \frac{1}{2} \left[1 + \frac{z}{\varepsilon} + \frac{1}{\pi} \sin \left(\frac{\pi z}{\varepsilon} \right) \right], & \text{if } |z| \leq \varepsilon \end{cases} \quad (4.1)$$

$$H_2(z) = \frac{1}{2} \left(1 + \frac{2}{\pi} \arctan \left(\frac{z}{\varepsilon} \right) \right) \quad (4.2)$$


An MRI image of a corn can be seen in figures (4.1a) and (4.1b). The H_1 regularization segments the corn kernels and the light areas of the corncob from the background with some artefacts on the peel. The H_2 regularization also contains the peel in addition to the corn kernels and the corncob inside the curve. The darker part of the corncob is now included in the contour in contrast to the H_1 regularization. Depending on if the peel of the corn is of importance, one of H_1 or H_2 is preferred from the other in this case.



(a) The curve found with the H_1 regularization in 142 iterations. (b) The curve found with the H_2 regularization in 180 iterations.

Figure 4.1: MRI scan of a corn.


An MRI image of a melon can be seen in the figures (4.2a) and (4.2b). The H_1 regularization segments the seeds from the melon from the pulp of the melon and the background with some artefacts on the left side of the melon. The H_2 regularization tries to segment the pulp and seeds together from the background. It fails to do so at the upper and lower part of the melon. On the ends of the melon, the pulp starts to become darker, hence the segmentation takes those parts as background in the segmentation. Once again, the H_1 and H_2 regularizations serve different purposes based on what is wanted to be segmented.



(a) The curve found with the H_1 regularization in 144 iterations. (b) The curve found with the H_2 regularization in 181 iterations.

Figure 4.2: MRI scan of a melon.

On the MRI images of a brain in (4.3a) and (4.3b) it can be seen that the H_2 regularization perform better at segmenting the inner darker areas from the lighter areas of the brain. Both regularization methods fail to accurately segment the darker parts of the brain from the lighter parts. The difference in intensity between the brain and the background is large. Because of that the whole brain is considered as the object that the model is trying to segment. With just a slight increase in brightness of the background, the segmentation starts to differentiate more between the lighter and darker area of the nut as seen in (4.4a) and (4.4b).




(a) The curve found with the H_1 regularization in 141 iterations. (b) The curve found with the H_2 regularization in 179 iterations.

Figure 4.3: MRI scan of a brain.

The H_1 regularization is significantly more impacted by the contrast difference between the background and the bright objects within the image. This sensitivity results in the H_1 regularization attempting to segment the bright white areas from the rest of the image. The model creates a contour around these high-contrast regions. The H_2 regularization is more lenient to allow light gray areas inside the contour.

Based on the expected contrast between the background and the brightest objects, one regularization can be preferred above the other. If the goal is to segment regions with high contrast, the H_1 regularization might be preferred due to its sensitivity to sharp differences in intensity. H_2 can be used for less pronounced contrast differences.

Furthermore, both regularizations can be effectively utilized to segment different el-



(a) The curve found with the H_1 regularization for the brain with brighter background. (b) The curve found with the H_2 regularization for the brain with brighter background.

Figure 4.4: MRI scan of a brain with a brighter background.

ements within the same set of images. They can also potentially be used in tandem to segment the gray areas of the image.

4.3 Noise and Smoothing


The Chan-Vese model, unlike gradient-based contour models, can still produce good results even when dealing with blurry or noisy images [9]. This robustness stems from the approach to image segmentation, which does not rely on image gradients, but instead utilizes region-based information, making it less sensitive to the sharpness of object boundaries.

Smoothing is often used to remove noise while preserving the structure of the objects. In the case of smoothed images, the boundaries of objects are not sharply defined. When an image is smoothed, the gradient information becomes less distinct. This often leads to poor segmentation.

Images taken with the low-field scanner of a papaya and a bell pepper were used to test the effectiveness of the model. The images obtained from the scan contain noise. Smoothed version of the images were also tested. The images were filtered with the use of the Nordstrom model [17]. The images were tested on both regularizations of the Heaviside functions.


In the noisy images of the papaya (4.5b) and of the bell pepper (4.6b), the H_2 regularization did not perform well. The presence of noise points in the images led to their inclusion within the contours, which was not intended. This resulted into a bad segmentation, where the model mistakenly identified noise artefacts as part of the object boundaries. In the bell pepper image it even failed to capture the upper and bottom part of the boundary of the pepper. The H_1 regularization demonstrated a better ability to exclude noise points when forming a contour around the objects, as seen in the image of the papaya (4.5a) and the image of the bell pepper (4.6a). This regularization effectively ignored the noise in the image of the bell pepper. In the image of the papaya a line of noise was included in the middle of the image. This line of noise probably was included as the noise points in that line are significantly brighter than the other noise points on the image. The H_1 regularization resulted in more accurate contours that closely fit to the true boundaries of

the objects, as barely any noise points were included inside the contour.



(a) The curve found with the H_1 regularization in 139 iterations. (b) The curve found with the H_2 regularization in 181 iterations.

Figure 4.5: MRI scan of a papaya.



(a) The curve found with the H_1 regularization in 139 iterations. (b) The curve found with the H_2 regularization in 180 iterations.

Figure 4.6: MRI scan of a bell pepper.

Using the segmentation boundaries found for the noisy images of the bell pepper and the papaya, the objects can now be separated from the background. The images will be divided in two parts, one containing the background and one containing the objects. The pixels that are not included in the given image will be set to have a value of 1. This will make the pixels white. This will result in the segmentations found in (4.7) and (4.8).

The model was also evaluated on images containing artificially added Gaussian white noise. Specifically, the model was tested on images of corn and melon, with Gaussian white noise imposed onto the original images. Despite the presence of this noise, the segmentations produced by the Chan-Vese model remained largely consistent with those obtained from noise-free images. The biggest difference can be found in the segmentation of the MRI image of the melon with the H_2 regularization (4.10b). In contrast to the segmentation of the noise-free image (4.2b), the curve found for the noisy image is less smooth at the top and bottom of the melon.

This consistency can be attributed to the nature of Gaussian white noise. It affects the entire image uniformly, thus the mean intensity values change uniformly across all regions. Due to this the relative differences in the intensity are preserved. The Chan-Vese




Figure 4.7: Segmentation of image of the bell pepper with the H_1 regularization. From left to right: the original image, the segmented background and the bell pepper.





Figure 4.8: Segmentation of image of the papaya with the H_1 regularization. From left to right: the original image, the segmented background and the papaya.

model effectively handles the added noise, maintaining its ability to accurately segment the images into distinct regions.


For the filtered (smoothed) versions of the images, the H_1 regularization again outperformed the H_2 regularization. In the inner boundary H_1 performs much better at following the star shaped boundary than H_2 . The H_2 regularization also contained some noise points on the right side. The H_2 regularization fails to find the contour around the bell pepper in (4.12b).

For both regularizations it can be seen that in the filtered image of the bell pepper found in (4.12a) and (4.12b), the model fails to segment the blank space inside the bell peppers into the region of the background. In the noisy image of the bell pepper (4.6a) the H_1 regularization succeeds in doing so.




(a) H_1 , Gaussian white noise with a standard deviation of 0.1. (b) H_2 , Gaussian white noise with a standard deviation of 0.1.

Figure 4.9: Filtered MRI scan of the corn with added white noise.



(a) H_1 , Gaussian white noise with a standard deviation of 0.1. (b) H_2 , Gaussian white noise with a standard deviation of 0.1.

Figure 4.10: Filtered MRI scan of the melon with added white noise.




(a) The curve found with the H_1 regularization in 139 iterations. (b) The curve found with the H_2 regularization in 180 iterations.

Figure 4.11: Filtered MRI scan of a papaya.

4.4 Choice of the area penalty

In the previous section, it was observed that the H_2 regularization tended to include noise points within the curve, and specifically for the MRI scan of the bell pepper, the H_2 regularization failed to accurately enclose the bell pepper within the curve. To address




(a) The curve found with the H_1 regularization in 138 iterations. (b) The curve found with the H_2 regularization in 180 iterations.

Figure 4.12: Filtered MRI scan of the bell pepper.

these issues, an area penalty can be introduced. By applying a positive area penalty, the model incentivizes the curve to minimize its area. This discourages the curve from expanding unnecessarily to include noise points.

This effect is evident in the images referenced in figures (4.13a) and (4.14a). In these examples, the majority of noise points are effectively excluded from the curve. Notably, in the bell pepper image (4.14a), the curve even excludes empty spaces within the bell pepper, which was not achieved with the H_2 regularization alone. In the filtered image of the papaya (4.13b), the curve now excludes the noise points to the right of the papaya, demonstrating an improvement over the curve obtained without the area penalty (4.11b).


The biggest improvement is seen in the filtered image of the bell pepper. Previously, the H_2 regularization failed to wrap the curve around the bell pepper's boundary (4.12b). However, with the addition of the area penalty, the segmentation not only outlines the bell pepper's boundary more precisely but also excludes the empty space within the bell pepper. This result surpasses the segmentation achieved with H_1 regularization without the length penalty.



(a) MRI scan of papaya with $\nu = 14 \cdot 10^{-13}$ (b) Filtered MRI scan of a papaya with $\nu = 14 \cdot 10^{-13}$.

Figure 4.13: Segmentation with a non-zero area penalty on the MRI scan of the papaya.

The area penalty can also be employed to replicate the results obtained from H_1 regularization on MRI scans of the melon and corn. The H_2 regularization originally included the peel of the corn and the pulp of the melon within their respective curves. By applying




(a) MRI scan of a bell pepper with $\nu = 50 \cdot 10^{-13}$. (b) Filtered MRI scan of a bell pepper with $\nu = 54 \cdot 10^{-13}$.

Figure 4.14: Segmentation with a non-zero area penalty on the MRI scan of the bell pepper.

an area penalty, these parts can be excluded because the penalty discourages the curve from expanding to include them, thus keeping the area smaller.

The results of applying the area penalty are illustrated in figures (4.15a) and (4.13b). By attempting to replicate the effects of the length penalty, we managed to achieve results comparable to those obtained with the H_1 regularization. This approach even succeeded in excluding some unnecessary points from within the contour, enhancing the precision of the segmentation. However, there are some limitations. In the segmentation of the corn image, only a few points of the corncob are included within the contour, which indicates that the curve does not fully contain the corncob in the object region. Similarly, in the segmentation of the melon image, certain spots of the seeds are not included within the curve.



(a) MRI scan of a corn, $\nu = 4 \cdot 10^{-12}$. (b) MRI scan of a melon, $\nu = 4 \cdot 10^{-12}$.

Figure 4.15: Imitating the results of the H_1 regularization using the H_2 regularization on the model containing a non-zero area penalty.

In the MRI image of the brain (4.3b), the segmentation without an area penalty failed to accurately segment the lighter parts of the brain from the darker parts. With the introduction of an area penalty this segmentation can be achieved as seen in (4.16).

The inclusion of an area penalty can significantly improve segmentation results. By incorporating an area penalty, the model is discouraged from expanding the contour to include noise points, leading to better segmentations.




Figure 4.16: MRI image of the brain with H_2 regularization and $\nu = 5 \cdot 10^{-12}$. From left to right: the original image, the segmented background, the lighter part of the brain.

4.5 Choice of λ

The choice of the λ parameter plays a big part in the segmentation process, allowing for fine-tuning of how much the pixels inside and outside the curve can deviate from the mean of their respective areas. This parameter gives control over the uniformity of the pixel values within and outside the segmented region. By adjusting λ , one can dictate the model's range of pixel values inside the contour. For example, increasing the value of λ_1 slightly directs the model to favour curves around regions with more uniform intensity. This means that the model becomes more inclined to create segments around objects that have consistent intensity patterns. Consider the example of the MRI scan of the corn (4.17). By increasing λ_1 , the model segments the image such that only the peel and the small parts of the boundary of the corncob are included within the curve. This is because these parts of the corn have relatively uniform intensity. On the other hand, the corn kernels, which are less uniform in intensity, are excluded from the segmented area.

A similar effect can be observed in the image of the melon (4.18). In this case, the seeds are not included in the segmentation. The seeds are excluded as they do not exhibit uniform intensity patterns. The segmentation model, prefers to outline regions that are more homogenous in terms of intensity.



Figure 4.17: H_2 regularization with $\lambda_1 = 1 + 53 \cdot 10^{-12}$. From left to right: the original image, the segmented background and the peel of the corn.




Figure 4.18: H_2 regularization with $\lambda_1 = 1 + 53 \cdot 10^{-12}$. From left to right: the original image, the segmented background and the found pulp of the melon.

Discussion

5.1 Disadvantages of gradient descent

One of the major drawbacks of the chosen method for optimizing the curve in Chan-Vese segmentation is its reliance on gradient descent, which frequently converges to a local minimum rather than the global minimum. This issue becomes particularly problematic if the equation has multiple local minima. In such scenarios, the outcome produced by gradient descent heavily dependent on the initial contour selected for the model. Therefore, making an appropriate initial guess is essential to increase the likelihood of finding a more optimal solution. Without a good initial guess, the model may settle into a suboptimal local minimum, potentially leading to less effective results.

5.2 Parameter Sensitivity

The model relies on multiple parameters that must be adjusted for each unique application. Incorrect parameter settings can significantly degrade the quality of segmentation results. During testing, it was observed that even slight adjustments to these parameters can have a large impact on the outcome of the segmentation. For instance, the difference between the parameters λ_1 and λ_2 plays a critical role. If this difference is too large the model may fail to find a curve around the zero level of the ϕ function. If the area penalty is set excessively high, the model also fails to find a curve around the zero level of the ϕ function. This zero level contour is crucial as it defines the boundary between the regions inside and outside the segmented objects. In the experiments tiny differences between λ_1 and λ_2 were usable. By a difference of more than 10^{-10} it was not able to find a curve. The same holds for ν as an area penalty of more that 10^{-11} was unusable.

5.3 Choice of length penalty

The length penalty in the Chan-Vese segmentation model ensures the smoothness of the curve that splits the segmented regions. This penalty term regularizes the boundary, preventing it from becoming overly complex by favouring shorter and smoother curves.

When investigating the impact of varying the length penalty parameter ϕ , it was observed that changes in ϕ did not affect the smoothness of the boundary. Regardless of the value of ϕ , the resulting boundary maintained a consistent level of smoothness. The effect of adjusting the parameter ϕ was found on the convergence rate of the model. Specifically, larger values of ϕ led to the stationary solution being reached in fewer iterations. Further research should be done on the effect of the length penalty.

5.4 Two regions segmentation

The Chan-Vese model is specifically designed for segmenting objects within an image into two distinct regions. This means that it cannot distinguish between different objects within the same segmented region. As a result, differentiating between various objects within that region is not possible using the standard Chan-Vese model.

To effectively segment images containing multiple objects, the traditional single-level set method must be modified to incorporate a multiphase level set method. The multiphase level set method, as introduced by L.A. Vese and T.F. Chan in 2002 [18], involves the introduction of additional curves. These extra curves enable the segmentation process to separate the image into more than two regions, allowing for the differentiation of multiple objects within the image. By employing these additional curves, the multiphase level set method extends the capabilities of the original model, making it suitable for more complex segmentation tasks where distinguishing between several distinct objects is necessary. For future applications, it is recommended to implement the multiphase level set method.

5.5 Reinitialization

The level set function can be reinitialized after every couple iterations by replacing ϕ with the signed distance function to its zero-level curve. When working with level sets and Dirac delta functions, a standard procedure is to reinitialize to ϕ [9]. This change does not modify the boundary itself, but this prevents the level set function to become too flat. It comes with the downside that it can prevent the curve from growing [9]. Reinitialization can be done with the evolution equation (5.1) as proposed in [19]:

$$\psi_\tau = \text{sign}(\phi(t))(1 - |\psi|)\psi(0, x, y) = \phi(t, x, y). \quad (5.1)$$

$\phi(t, x, y)$ is our solution at the time t . The new ϕ will become the ψ that is obtained at the stationary solution of the partial differential equation. The results found in the experiments deemed to be sufficient to produce good segmentations and for that reason it was chosen to not include reinitialization. Reinitialization would lead to more computations and would risk preventing the curve from growing. For future applications, reinitialization could be used in the model.

Conclusion

The purpose of this report was twofold. The first was to implement the Chan-Vese model in Fortran. The second was to determine whether the Chan-Vese model can successfully separate the objects from the background in the low quality images produced by low-field MRI scanners.

With appropriate parameter settings, the segmentation process successfully minimizes the inclusion of unnecessary noise points within the region of the object. The ability of the model to find precise boundaries in smoothed images shows its capability to handle variations in image quality.

The choice of regularization for the Heaviside function influences the results of the Chan-Vese segmentation technique. Different regularizations of the Heaviside function lead to the inclusion of different parts of the image as objects. This means that the boundaries of the segmented regions vary depending on the regularization method used. For instance, a sharper regularization of the Heaviside function can result in more defined and distinct boundaries. This approach is beneficial when there is a need to clearly separate objects with sharp edges from the background or when there is a significant contrast between the objects and their surroundings. A smoother regularization of the Heaviside function might lead to more gradual transitions between the foreground and background. By choosing the right regularization method, we can make the segmentation process fit the specific details of the image and the goals of the analysis better.

Adjusting the area penalty (ν parameter) allows for fine-tuning the segmentation process by controlling the size of the curve. This parameter plays a crucial role in determining how large or small the segmented region will be. When the initial segmentation includes more parts of the objects than needed, modifying the area penalty becomes particularly useful. By increasing the area penalty, the algorithm is discouraged from creating larger curves, thus preventing the inclusion of unnecessary sections within the contour. This helps in achieving a more precise segmentation by focusing only on the relevant parts of the objects.

The λ parameters are crucial when segmenting objects with uniform colours. By increasing the λ values for the inside or outside of the contour, the segmentation process can be adjusted to better encapsulate uniform objects. If an object has a uniform colour, increasing the λ_1 parameter helps the contour to more accurately follow the boundaries of the uniformly coloured object. Similarly, increasing the λ_2 parameter can help in sit-

Chapter 6. Conclusion

uations where the background is uniform. This adjustment makes the algorithm more attuned to the uniformity of the pixel values outside the contour, helping to exclude areas that do not belong to the object. This fine-tuning ensures that the segmented area reflects the uniform characteristics of the objects being analysed.

The remaining task is to answer the research question: Can the Chan-Vese model be effectively used to segment the low-quality images produced by the low-cost mobile MRI scanner?. The Chan-Vese segmentation technique can be effective in finding the right segmentation boundary to segment the low-field MRI images. This method proves to be reliable, as the segmentations derived from the provided MRI scans accurately differentiate the scanned objects from the background and any surrounding noise. This distinction is crucial in medical diagnostics, where precise identification of anatomical structures is necessary.

Summarizing, the Chan-Vese model segments the low quality images well, with appropriate choice of parameters. The model is limited in dividing the images into two regions. Future research should investigate if it is possible to further segment the region with the objects, such that each object has a separate region.

Bibliography

- [1] Elizabeth A Krupinski and Yulei Jiang. “Anniversary paper: evaluation of medical imaging systems”. In: *Medical physics* 35.2 (2008), pp. 645–659.
- [2] World Health Organization et al. “Global atlas of medical devices”. In: (2017).
- [3] Merel de Leeuw den Bouter, Martin B. van Gijzen, and Rob Remis”. “Conjugate gradient variants for ℓ p-regularized image reconstruction in low-field MRI”. In: *SN Applied Sciences* 1.12 (2019), p. 1736.
- [4] T O’Reilly, WM Teeuwisse, and AG Webb. “Three-dimensional MRI in a homogeneous 27 cm diameter bore Halbach array magnet”. In: *Journal of Magnetic Resonance* 307 (2019), p. 106578.
- [5] Johnes Obungoloch et al. “Design of a sustainable prepolarizing magnetic resonance imaging system for infant hydrocephalus”. In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 31 (2018), pp. 665–676.
- [6] Merel de Leeuw den Bouter et al. “Description of a low-field MRI scanner based on permanent magnets”. In: *CVCS 2020 Colour and Visual Computing Symposium 2020 Proceedings of the 10th Colour and Visual Computing Symposium 2020 (CVCS 2020) Gjøvik, Norway, and Virtual, September 16-17, 2020*. Vol. 2688. 2020.
- [7] M De Koleva and O Jesus. “Hydrocephalus”. In: *Hydrocephalus. In StatPearls. Treasure Island* (2021).
- [8] Matthieu Vinchon, Harold Rekate, and Abhaya V Kulkarni. “Pediatric hydrocephalus outcomes: a review”. In: *Fluids and Barriers of the CNS* 9 (2012), pp. 1–10.
- [9] Tony F Chan and Luminita A Vese. “Active contours without edges”. In: *IEEE Transactions on image processing* 10.2 (2001), pp. 266–277.
- [10] David Bryant Mumford and Jayant Shah. “Optimal approximations by piecewise smooth functions and associated variational problems”. In: *Communications on pure and applied mathematics* (1989).
- [11] Stanley Osher and James A Sethian. “Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations”. In: *Journal of computational physics* 79.1 (1988), pp. 12–49.
- [12] LawrenceCraig Evans. *Measure theory and fine properties of functions*. Routledge, 2018.

- [13] J.J.I.M. van Kan, A. Segal, and F. Vermolen. “Numerical methods in scientific computing”. In: (2023).
- [14] J. Weickert, B.M.T.H. Romeny, and M.A. Viergever. “Efficient and reliable schemes for nonlinear diffusion filtering”. In: *IEEE transactions on image processing* 7.3 (1998), pp. 398–410.
- [15] Yao Zhang et al. “A hybrid method for solving tridiagonal systems on the GPU”. In: *GPU Computing Gems Jade Edition*. Elsevier, 2012, pp. 117–132.
- [16] W Taylor Kimberly et al. “Brain imaging with portable low-field MRI”. In: *Nature Reviews Bioengineering* 1.9 (2023), pp. 617–630.
- [17] K Niklas Nordström. “Biased anisotropic diffusion: a unified regularization and diffusion approach to edge detection”. In: *Image and vision computing* 8.4 (1990), pp. 318–327.
- [18] Luminita A Vese and Tony F Chan. “A multiphase level set framework for image segmentation using the Mumford and Shah model”. In: *International journal of computer vision* 50 (2002), pp. 271–293.
- [19] Mark Sussman, Peter Smereka, and Stanley Osher. “A level set approach for computing solutions to incompressible two-phase flow”. In: *Journal of Computational physics* 114.1 (1994), pp. 146–159.

Original MRI images




Figure .1: MRI image of a corn by Andy Ellison




Figure .2: MRI image of a melon by Andy Ellison



Figure .3: MRI image of a brain




Figure .4: MRI image of a bell pepper by Ksenia Slepova




Figure .5: MRI image of a papaya by Ksenia Slepova