



NTNU – Trondheim
Norwegian University of
Science and Technology

Anisotropic Total Variation Based Image Restoration Using Graph Cuts

Bjørn Rustad

Master of Science in Physics and Mathematics

Submission date: February 2015

Supervisor: Markus Grasmair, MATH

Norwegian University of Science and Technology
Department of Mathematical Sciences

Abstract

In this thesis we consider a particular kind of edge-enhancing image restoration method based on total variation. We want to address the fact that the total variation method in some cases leads to contrast loss in thin structures. To reduce the contrast loss a directional dependence is introduced through an anisotropy tensor. The tensor controls the regularization applied based on the position in the image and the direction of the gradient. It is constructed using edge information extracted from the noisy image. We optimize the resulting functional using a graph cut framework; a discretization which is made possible by a coarea and a Cauchy–Crofton formula. In the end we perform numerical studies, experiment with the parameters and discuss the results.

Sammendrag

I denne masteroppgaven ser vi på en spesifikk kant-bevarende støyfjerningsalgoritme basert på «total variation». Vi tar for oss at «total variation» i noen tilfeller fører til tap av kontrast i detaljer og tynne strukturer. For å redusere kontrast-tapet introduserer vi en retningsavhengig anisotropitensor. Denne tensoren kontrollerer støyfjerningen basert på posisjonen i bildet, og retningen til gradienten i punktet. Den blir konstruert basert på kant-informasjon fra det opprinnelige støyete bildet. Vi minimerer den resulterende funksjonale i et graf-kutt-rammeverk, som er gjort mulig ved hjelp av en coarea- og en Cauchy–Crofton-likning. Vi avslutter med en numerisk studie, eksperimentering med parametrene og diskusjon av resultatene.

Preface

This master thesis concludes my study at the Applied Physics and Mathematics Master's degree program with specialization in Industrial Mathematics at the Norwegian University of Science and Technology (NTNU).

I would like to thank my supervisor Markus Grasmair at the Department of Mathematical Sciences for invaluable help and discussion throughout my work with my project and this thesis.

Finally I would like to thank my family for their support, and Mats, Lars, Kine, Hager, Edvard and Henrik for productive discussions around the coffee pot.

Bjørn Rustad, February 8, 2015.

Contents

1	Introduction	1
2	Methods in image restoration	3
2.1	Diffusion filtering	4
2.2	Total variation	6
3	Continuous formulation	11
3.1	Anisotropic total variation	11
3.2	Well-posedness	16
3.3	Anisotropic coarea formula	19
3.4	Anisotropic Cauchy–Crofton formula	24
4	Discrete formulation	32
4.1	Discretization	32
4.2	Graph cut approach	41
5	Maximum flow	47
5.1	Flow graphs	47
5.2	Augmenting path algorithms	49
5.3	Other algorithms	50
5.4	Push–relabel algorithm	51
5.5	Boykov–Kolmogorov algorithm	60
6	Results	66
6.1	Tensor parameters	66
6.2	Neighborhood stencils	70
6.3	Restoration	73
7	Discussion and conclusion	76

Bibliography	79
List of Figures	83
List of Tables	85
List of Symbols	87
A C++ implementation	89
A.1 main.cpp	90
A.2 image.hpp	92
A.3 image.cpp	93
A.4 anisotropy.hpp	96
A.5 anisotropy.cpp	96
A.6 graph.hpp	99
A.7 graph.cpp	101
A.8 selectionrule.hpp	109
A.9 selectionrule.cpp	110
A.10 neighborhood.hpp	111

Introduction

Image processing is becoming an increasingly important part of our modern computerized world. Tasks previously only performed by humans, like detecting edges, recognizing textures and inferring shapes and motions can now be performed algorithmically. The background of these methods spans several fields, including psychology and biology for the study of human vision, statistics and analysis for the mathematical background, and computer science for their implementation and performance analysis.

Image restoration methods are concerned with trying to remove noise or recover otherwise degraded images. Possible noise can result from the physical nature of light traveling to your sensor, dust on your lens, and many other sources. Therefore numerous different approaches to denoising exist, each having their own strengths and weaknesses. Some of these are introduced in Chapter 2, and one of the main challenges they all face is the recovery of edges.

A method well known for recovering edges is the total variation method, as the total variation does not favor smooth gradients over edges. I gave an overview of this method in my project work [1], where I used a graph cut framework to obtain a numerical solution. The method consists of trying to reduce the total variation of the image, while still staying “close” to the original.

A problem with the total variation method is that contrast is often lost, especially in fine details and thin structures. In this thesis we try to alleviate this. We extend the method by introducing an anisotropy tensor into the total variation, thus making it directionally dependent. This means we can control the regularization applied to the image based on position and direction. The main idea is then to reduce the regularization applied across edges in the image, while we still regularize along them.

The variational problem we obtain is a convex minimization problem, and many optimization approaches exist. We choose to discretize in such a way that

we can apply the same graph cut framework used in my project work [1]. Through the coarea formula, the functional is decomposed into a sequence of minimization problems, one for each level of the image. These separate level problems are then transformed and discretized further using an anisotropic Cauchy–Crofton formula that we develop. Similar formulas have been presented before in other contexts.

A nice property of this numerical approach is that we can prove that the graph cut framework finds an exact global minimizer of the discrete functional. Additionally we verify that the discrete functional is consistent with the continuous one.

We present and implement two maximum flow algorithms that allow us to find minimum cuts corresponding to minimizers of the discrete functionals. The push-relabel algorithm is considered to be the fastest and most versatile for general graphs, while the Boykov–Kolmogorov algorithm is specially tailored for the type of graphs we find in these kinds of imaging applications. We describe every part of the method in detail such that it can be easily implemented by the reader. In addition, a C++ implementation is attached.

In the end we present numerical results that show how the different parameters affect the restoration, and we look into and explain some artifacts caused by approximations in the discretization. Further we look at how the introduction of the anisotropy in certain cases amend some of the weaknesses of the total variation method. We particularly look at how contrast loss is reduced in images containing thin structures such as fingerprints.

Methods in image restoration

There are numerous methods in image restoration, but we do not have time nor space to discuss them all. In short overview, which is an extension of the one given in my project [1], we will focus on the methods related to the anisotropic total variation method considered later in this thesis. See [2] and [3] for more background on image processing in general.

In this chapter, and also in the rest of the thesis we will assume that we are given an image $f : \Omega \rightarrow \mathbb{R}$ where Ω is a rectangular, open domain. Because of limitations in the numerical method used, the codomain is \mathbb{R} and we are thus restricted to monochrome, or grayscale images. Such images are produced in large numbers by for example ultrasound, X-ray and MRI machines.

The space in which the image f resides in will vary, but since we are looking at image restoration methods, we assume that it includes some kind of noise. Depending on the application and how the image is obtained, one might construct different models describing different types of noise.

We will assume that the given image f is a combination of an underlying, actual image u^* , and some noise δ . The simplest model is additive noise where the assumption is that $f = u^* + \delta$. There is also multiplicative noise where $f = u^* \cdot \delta$. An other much seen noise type is salt and pepper noise, which is when black and white pixels randomly appear in the image.

These are only models, and in the real world the noise might be more complex, and even come from a combination of sources. Depending on the application, the goal might not even be to recover u^* , but rather to obtain an output which fulfills certain smoothness or regularity properties. In any case, we will continue denoting the noisy input image f and use u for the output image in the description of the restoration methods.

2.1 Diffusion filtering

Diffusion filtering is a broad group of filtering and restoration methods based on physical diffusion processes. The basic idea is to take the noisy image as the initial value of some diffusion process, and then let it evolve for some time. The best known method is probably the Gaussian filter or Gaussian blur, in which one convolves the image with the Gaussian function

$$K_\sigma(x, y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right). \quad (2.1)$$

In the discrete setting where the image consists of a grid of pixels, the Gaussian blur amounts to calculating each pixel in the output image as a weighted average of its neighboring pixels in the input image.

The Gaussian function happens to be the fundamental solution of the heat equation $\partial_t u = \Delta u$. Convolution with $K_\sigma(x, y)$ with the original image f is therefore equivalent to solving the heat equation with f as initial value, until some time $T > 0$ depending on σ . Boundary conditions have to be specified of course, and one common choice is to symmetrically extend the image in x and y directions, which corresponds to zero flux boundary conditions.

By basic Fourier analysis it is possible to show that the Gaussian filter is a low-pass filter which attenuates high frequencies. Further theory can be found in Weickert's book on anisotropic diffusion [4].

The main concern with the Gaussian filter is that it will, in addition to smoothing out possible noise, remove details from the image, which motivates the next set of methods, where the amount of diffusion can vary between different parts of the image.

2.1.1 Non-linear diffusion filtering

In the theory of the heat equation one can introduce a *thermal diffusivity* α such that the equation becomes

$$\begin{cases} \partial_t u = \operatorname{div}(\alpha(\nabla u)\nabla u), \\ u|_{t=0} = f. \end{cases} \quad (2.2)$$

The thermal diffusivity $\alpha(\nabla u) = \alpha(x, \nabla u)$ is material dependent, and can also vary throughout the object. It specifies how well heat travels through the specific point in the object. We can make use of this in the image restoration context by specifying different diffusivity in different parts of the image, in an effort to reduce noise without losing image detail. Optimally, we would like there to be a lot of

diffusion in smooth parts of the image, and not so much in areas with a lot of details.

One much-studied non-linear diffusion equation is the Perona–Malik equation

$$\partial_t u = \operatorname{div} \left(\frac{\nabla u}{1 + \frac{|\nabla u|^2}{\lambda^2}} \right). \quad (2.3)$$

The thermal diffusivity $\alpha(\nabla u) = (1 + |\nabla u|^2/\lambda^2)^{-1}$ varies from 1 in smooth areas to 0 as the norm of the gradient $|\nabla u|$ grows.

This particular form of the thermal diffusivity has been shown to be related to how brightness is perceived by the human visual system. The model has some theoretical problems related to well-posedness, for more information see [4].

A different kind of non-linear diffusion model is the total variation flow which can be formulated as

$$\partial_t u = \operatorname{div} \frac{\nabla u}{|\nabla u|}, \quad (2.4)$$

where the diffusivity has a similar effect of reducing the diffusion in areas of high variation. As the name suggests, this model can be related to the variational total variation formulation presented later. One forward Euler time-step in the solution of this partial differential equation corresponds to the Euler–Lagrange equation of the variational formulation.

Note that we follow Weickert’s terminology when it comes to the distinction between non-linear and anisotropic diffusion methods. The Perona–Malik equation, and other diffusion equations with non-homogenous diffusivities, are often by others called anisotropic, as the diffusivity depends on the location. We will name these methods non-linear and spare the anisotropy term for the “real” anisotropic methods. These are methods where the diffusivity is a tensor, and thus both location and direction dependent.

2.1.2 Anisotropic diffusion

The diffusivity is made directionally dependent by introducing a diffusion *tensor* $A(u)$ such that the initial boundary value problem becomes

$$\begin{cases} \partial_t u = \operatorname{div} (A(u)\nabla u) & \text{on } \Omega \times (0, \infty), \\ u|_{t=0} = f & \text{on } \Omega, \\ A(u)\nabla u \cdot \nu = 0 & \text{on } \partial\Omega \times (0, \infty), \end{cases} \quad (2.5)$$

where ν is the outer normal of Ω . The tensor $A(u)$ is constructed such as to diminish the effect of ∇u across what we believe to be edges in the image. This

way, there will also be less diffusion through these edges. Weickert [4] suggests constructing $A(u)$ based on the edge estimator ∇u_σ where

$$u_\sigma := K_\sigma * \tilde{u} \quad (2.6)$$

and \tilde{u} is an extension of u from Ω to \mathbb{R}^2 made by symmetrically extending u across the boundary of Ω . Assuming we are at an edge in the image, the direction of ∇u_σ should be perpendicular to the edge, while its magnitude will provide information on the steepness of the edge.

To extract this information, and also to identify features on a larger scale, the *structure tensor* is introduced

$$S_\rho(x) := K_\rho * (\nabla u_\sigma \otimes \nabla u_\sigma), \quad (2.7)$$

where the convolution with the Gaussian function K_ρ is done component-wise. The anisotropy tensor $A(u)$ can then be constructed based on the eigenvectors and eigenvalues of $S_\rho(x)$. The structure tensor and its properties will be discussed further when we introduce our anisotropic total variation functional.

Assuming some smoothness, symmetry and uniform positive definiteness on $A(u)$ one can prove well-posedness, regularity and an extremum principle of the problem (2.5) as done in [4].

However, even if the diffusivity tensor was introduced to reduce the amount of smoothing across edges, the solution of (2.5) will still be infinitely differentiable [4], i.e. $u(T) \in C^\infty(\Omega)$ for $T > 0$. Thus there are no real discontinuities, and no real edges in the solution.

Further, the anisotropic diffusion may introduce structure based on noise, when there really was no structure to begin with. This is a problem we aim to avoid in our anisotropic total variation method.

2.2 Total variation

Total variation was initially introduced to the field of image restoration by Rudin, Osher and Fatemi in [5] and is usually formulated as a minimization problem

$$\min_{u \in L^p(\Omega)} F(u), \quad (2.8)$$

$$F(u) = \underbrace{\int_{\Omega} |u - f|^p dx}_{\text{fidelity term}} + \beta \underbrace{\int_{\Omega} |\nabla u| dx}_{\text{regularization term}},$$

where p is normally taken to be 1 or 2. The fidelity term penalizes images u that are far from the original image f . The regularization term is the total variation

of the image, and minimizing it will reduce the variation and thus regularize the image. The β parameter controls the strength of the regularization. Note that $u = f$ is a minimizer of the fidelity term, while a constant image $u = c$ is a minimizer of the regularization term.

As this restoration method is the one which will be extended later in this thesis, we will look a little bit more deeply into the background and the numerical methods relating to it.

Since we do not only want to consider differentiable images $u \in C^1(\Omega)$ for which the gradient exists, we introduce the total variation using the distributional derivative.

Definition 2.1 (Total variation). *Given a function $u \in L^1(\Omega)$, the total variation of u , often written $\int_{\Omega} |Du| dx$, where the D is the gradient taken in the distributional sense, is*

$$\text{TV}(u) = \int_{\Omega} |Du| dx = \sup \left\{ \int_{\Omega} u \cdot \text{div } \varphi dx : \varphi \in C_c^{\infty}(\Omega, \mathbb{R}^2), \|\varphi\|_{L^{\infty}(\Omega)} \leq 1 \right\}. \quad (2.9)$$

The test functions φ are taken from $C_c^{\infty}(\Omega, \mathbb{R}^2)$, the space of smooth functions from Ω to \mathbb{R}^2 with compact support in Ω .

Note that since Ω is open and bounded, the test functions φ vanish on the boundary of Ω . Thus no variation is measured at the boundary.

As we are searching for an image with low total variation, it is useful to introduce the space of functions of bounded variation.

Definition 2.2 (Functions of bounded variation). *The space of functions of bounded variation $\text{BV}(\Omega)$ is the space of functions $u \in L^1(\Omega)$ for which the total variation is finite, i.e.,*

$$\text{BV}(\Omega) = \{u \in L^1(\Omega) : \text{TV}(u) < \infty\}. \quad (2.10)$$

Our optimization problem has thus become

$$\min_{u \in \text{BV}(\Omega)} \int_{\Omega} |u - v|^p dx + \beta \text{TV}(u). \quad (2.11)$$

As with any restoration method, the total variation method has its strengths and weaknesses. Its main strength is its ability to recover edges in the input image. The total variation of a section only takes the absolute change into account, and does not favor gradual changes like the diffusion methods.

There is also a theoretical result stating that the set of edges in the solution u is contained in the set of edges in the original image f , thus no new edges are created [6]. However, in the presence of noise, the method may introduce or rather




Figure 2.1: Although the original gradient was smooth, the total variation method manages to find structure in the noise, and create edges in the restored image.




Figure 2.2: A fingerprint heavily regularized using the total variation method. The originally white and black ridges have been brought closer in value, to reduce the total variation.

“find” new edges that were not in the original image, since flat sections of zero variation are encouraged by the functional. This effect is called the *stair-casing effect*, and can be seen in Figure 2.1 where a noisy gradient has been restored using the total variation method.

Fine details, thin objects and corners may suffer from contrast loss since bringing them closer to their surroundings reduces the total variation. An example of this is shown in Figure 2.2, where a not particularly noisy fingerprint image has been strongly regularized. The original black and white levels have been brought closer to yield a lower total variation in the regularized image.

2.2.1 Numerical methods

See [7] for an overview of some of the numerical methods relating to total variation image restoration. Amongst others it describes some dual and primal-dual methods, as well as the graph cut approach we take in this thesis.

Graph cut approach

Using graph cuts is the approach we will be taking later when considering the anisotropic total variation regularization, and it is therefore valuable to briefly look into how graph cuts are used in the case of regular total variation.

A graph cut is a set of edges that when removed will separate the graph into two disconnected parts. A minimum cut is a cut such that the sum of the weight of the edges in the cut is minimal. It has been shown that for some discrete functionals, it is possible to construct graphs for which the minimum cuts correspond to minimizers of the functional.

In the discrete setting our image consists of pixels, and is represented by a function $u : \mathcal{G} \rightarrow \mathcal{P}$ where \mathcal{G} is a regular grid over Ω , and $\mathcal{P} = \{0, \dots, L-1\}$ is the discrete set of pixel values, or *levels*. We denote the value in pixel x as $u(x) = u_x$.

For an image u and a level λ we denote the *level set* by $\{u > \lambda\}$, defined as the set $\{x \in \Omega : u_x > \lambda\}$. The thresholded image u^λ , an indicator function, is then defined as

$$u^\lambda = \chi_{u > \lambda}. \quad (2.12)$$

Here, χ_E signifies the characteristic function of the set E , the function which is equal to one in every point in E , and zero elsewhere.

The idea of the graph cut approach is to decompose the minimization problem into one minimization problem for each level of the image, and then solve them separately before combining the results.

Through careful manipulation of the continuous functional in (2.11) it is possible to obtain a discrete functional decomposed as a sum over all the level values on the form

$$F(u) = \sum_{\lambda=0}^{L-2} \sum_x F_\lambda^x(u_x^\lambda) + \beta \sum_{\lambda=0}^{L-2} \sum_{(x,y)} F^{x,y}(u_x^\lambda, u_y^\lambda) =: \sum_{\lambda=0}^{L-2} F_\lambda(u^\lambda) \quad (2.13)$$

where the sum over (x, y) is over all pixel pairs (x, y) in a neighbor relation, i.e. pixels that are “close” to each other. The actual form of the functional, and the steps to construct it will be presented later.

The graph cut we find will for each level λ give us the thresholded image u^λ , and they can then be combined to form the complete image u .

When constructing the graph used to find the thresholded image u^λ , we have two special vertices, one representing the set $\{u > \lambda\}$, and one which represents

the set $\{u \leq \lambda\}$. The pixels are then connected to these vertices with a weight representing how strongly they are related to the corresponding set. This weight will be based on the value of F_λ^x .

Additionally there are connections between pixels in a neighborhood relation, representing the energy $F^{x,y}$. Thus when finding a cut, we partition the pixels into the sets $\{u > \lambda\}$ and $\{u \leq \lambda\}$. And if in addition the cut is minimal, we know that the edges cut have minimal weight, and can prove that the u^λ found minimizes the functional in (2.13).

Continuous formulation

In the previous chapter we saw that there are many different approaches to the image restoration problem, all with their own strengths and weaknesses. The method considered in this thesis is an anisotropic total variation formulation, and the aim is to keep the strengths of the anisotropic diffusion and total variation methods, while eliminating some of their weaknesses.

This chapter will be devoted to the continuous formulation of the method. We will look at the functional we want to minimize and its different forms, and briefly discuss its well-posedness. Through the anisotropic coarea formula, the anisotropic total variation is rewritten as an integral of the perimeter of all the level sets of the image.

Following that, the anisotropic Cauchy–Crofton formula is introduced to make it feasible to calculate the perimeter of these level sets. All of this leads up to the discretization of our functional in the next chapter.

3.1 Anisotropic total variation

The method considered will build on the total variation regularization method of Section 2.2. From anisotropic diffusion in Section 2.1.2 we borrow the idea of making the regularization in each point directionally dependent. We introduce the anisotropic total variation

$$\mathrm{TV}_A(u) = \int_{\Omega} \sqrt{\nabla u(x)^T A(x) \nabla u(x)} \, dx \quad (3.1)$$

for all $u \in C^1(\Omega)$. We assume here that $A(x)$ is continuous and positive definite, and we will later need the eigenvalues of $A(x)$ to be uniformly bounded below and above. If $A(x)$ is the identity matrix we get the regular total variation found in

(2.8). When minimizing the regular total variation, we will also try to reduce the variation over known edges in the image. This can lead to unwanted contrast loss, especially in fine details. By controlling $A(x)$ such that the contribution of $\nabla u(x)$ is reduced across known edges, we hope to retain the regularization properties of the original method while reducing this contrast loss. If the variation across an edge is “ignored” by the functional, there is no gain in reducing the height of the edge as before.

Note that $u(x)$ and $A(x)$ are always dependent on the position in the image x , but we will sometimes drop the x , when the intended meaning is clear.

As we will not always be working with differentiable images, we extend the definition of the total variation functional. Being symmetric positive definite, the matrix A can be factored into two symmetric matrices as $A = A^{1/2}A^{1/2}$. We can then write

$$\begin{aligned}
\mathrm{TV}_A(u) &= \int_{\Omega} |A^{1/2}\nabla u| \, dx \\
&= \sup_{|\xi(x)| \leq 1} \int_{\Omega} (A^{1/2}\nabla u)^T \xi \, dx \\
&= \sup_{|\xi(x)| \leq 1} \int_{\Omega} \nabla u \cdot A^{1/2}\xi \, dx \\
&= \sup_{|\xi(x)| \leq 1} \int_{\Omega} u \operatorname{div}(A^{1/2}\xi) \, dx \\
&= \sup_{\eta^T A^{-1}\eta \leq 1} \int_{\Omega} u \operatorname{div} \eta \, dx,
\end{aligned} \tag{3.2}$$

where ξ and $\eta = A^{1/2}\xi$ are in $C_c^\infty(\Omega, \mathbb{R}^2)$, the space of smooth vector fields with compact support. In the following we define the norms $\|\xi\|_A = \sup_x (\xi^T A \xi)^{1/2}$ and $\|\eta\|_A^* = \sup_x (\eta^T A^{-1}\eta)^{1/2}$, and with that we present the formal definition of the anisotropic total variation.

Definition 3.1 (Anisotropic total variation). *For a function $u \in L^2(\Omega)$ and a continuous symmetric positive definite tensor $A : \Omega \rightarrow \mathbb{R}^{2 \times 2}$ we define the anisotropic total variation*

$$\mathrm{TV}_A(u) = \sup \left\{ \int_{\Omega} u \operatorname{div} \xi \, dx : \xi \in C_c^\infty(\Omega, \mathbb{R}^2), \|\xi\|_A^* \leq 1 \right\}. \tag{3.3}$$

With this extended definition, we have arrived at a minimization problem where we seek to find a minimizer of the functional

$$F(u) = \int_{\Omega} (u - f)^2 \, dx + \beta \mathrm{TV}_A(u). \tag{3.4}$$




Figure 3.1: A noisy fingerprint on the left, and the largest eigenvalue of the structure tensor is $|\nabla f_\sigma(x)|^2$ on the right, which—as we can see—functions as an edge detector.

Similar functionals have been considered in [8] and [9]. The question is now how to construct the anisotropy tensor $A(x)$ to get the improvements we hope for, and how the introduction of the tensor affects our numerical solution method.

3.1.1 Anisotropy tensor

There are many possible choices for the anisotropy tensor $A(x)$. Our constraints are that we have assumed it to be continuous and symmetric positive definite, and we have some wishes for its properties. We would first and foremost like it to down-weight ∇u in (3.1) across true edges, while maintaining normal regularization properties in smooth sections.

By true edges we mean that that we do not want the tensor to be sensitive to noise in the image, and thus find edges where there are none, so we somehow want to be sure about the edges we find.

Edges can be found in many different ways, but as suggested by Weickert in his book on Anisotropic Diffusion [4], and briefly mentioned in Section 2.1.2, a good starting point is the *edge detector* ∇f_σ . The image is smoothed by a Gaussian filter as described in Section 2.1: $f_\sigma = K_\sigma * \tilde{f}$, where \tilde{f} is the symmetric extension of the initial image f in \mathbb{R}^2 . The smoothing parameter σ is called the *noise scale*, and it controls the scale at which details are considered to be noise.

As seen in Figure 3.1, the edge detector is fine for detecting edges, but it can not give us information about larger structures, like corners and textures,

which is why we introduce the *structure tensor* $S_\rho(x)$. First consider the tensor $S_0(x) = \nabla f_\sigma(x) \otimes \nabla f_\sigma(x)$. It is symmetric positive semi-definite, and obviously contains no more information than the edge detector itself. Its eigenvalues are $\lambda_1 = |\nabla f_\sigma(x)|^2$ and $\lambda_2 = 0$ with corresponding eigenvectors v_1 and v_2 parallel and perpendicular to $\nabla f_\sigma(x)$ respectively.

To detect features in a neighborhood around the point x , such as corners, curved edges and coherent structures we introduce the component-wise convolution with K_ρ such that

$$S_\rho(x) := K_\rho * (\nabla f_\sigma(x) \otimes \nabla f_\sigma(x))(x). \quad (3.5)$$

The parameter ρ , called the *integration scale*, controls the size of the neighborhood which affects the structure tensor. Thus it defines the size of the structures we want our anisotropy tensor to be sensitive to.

The smoothed tensor $S_\rho(x)$ can easily be verified to be symmetric positive semi-definite, just like $S_0(x)$. In addition, when $\rho > 0$, the elements of S_ρ are smooth maps from Ω to \mathbb{R} .

We order the two real eigenvalues such that $\lambda_1 \geq \lambda_2$ and denote the corresponding eigenvectors v_1 and v_2 . From the characteristic polynomial of $S_\rho(x) = \begin{pmatrix} s_{11} & s_{12} \\ s_{12} & s_{22} \end{pmatrix}$ we obtain a closed form expression for the eigenvalues

$$\lambda = \frac{1}{2} \left(s_{11} + s_{22} \pm \sqrt{(s_{11} - s_{22})^2 + 4s_{12}^2} \right). \quad (3.6)$$

The vector v_1 will then indicate the direction of most variation in the neighborhood. An edge will give $\lambda_1 \gg \lambda_2 \approx 0$, while smooth areas will give $\lambda_1 \approx \lambda_2 \approx 0$. In corners we have variation in the direction of v_1 but also perpendicular to v_1 , so we will have $\lambda_1 \approx \lambda_2 \gg 0$. Thus the quantity $(\lambda_1 - \lambda_2)^2$ will be large around edges and small in smooth or non-coherent areas.

To extract this information from the structure tensor, we decompose it as

$$S_\rho(x) = U(x)\Lambda(x)U(x)^T, \quad (3.7)$$

where

$$\Lambda(x) = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \quad (3.8)$$

has the eigenvalues $\lambda_1 \geq \lambda_2$ on its diagonal, while $U(x)$ is a rotation matrix and has the eigenvectors of $S_\rho(x)$ as its columns. From this we construct a new matrix $A(x) = U(x)\Sigma(x)U(x)^T$ where

$$\Sigma(x) = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix}. \quad (3.9)$$

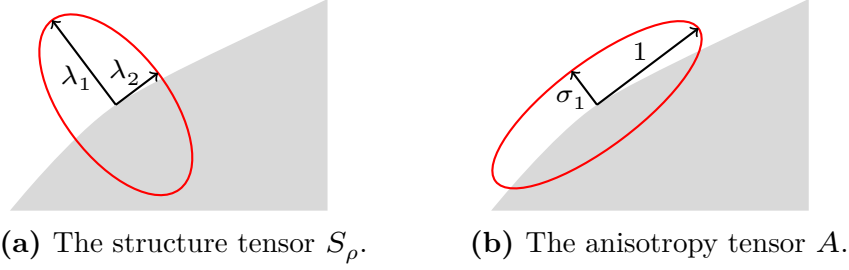


Figure 3.2: An edge with the structure and anisotropy tensors visualized using their eigenvectors and eigenvalues.

and for σ_1 and σ_2 we choose

$$\begin{aligned}\sigma_1 &= \left(1 + \frac{(\lambda_1 - \lambda_2)^2}{\omega^2}\right)^{-1}, \\ \sigma_2 &= 1.\end{aligned}\tag{3.10}$$

Thus the eigenvectors of $A(x)$ and $S_\rho(x)$ are equal, while the eigenvalues are different. A visualization of the two tensors can be seen in Figure 3.2 where the two tensors are shown at an edge in the image.

In smooth areas, $\sigma_1 \approx 1$ and $A(x)$ will be close to the identity matrix. At or around edges, σ_1 , which corresponds to the eigenvector perpendicular to the edge, will be small.

Around corners $A(x)$ will be close to the identity matrix, which gives regularization similar to smooth areas. This is one possible down-side of this tensor choice, as rounded corners may occur.

The parameter ω controls the amount of anisotropy in the method, such that if it is very large we are left with the identity matrix and our method becomes the regular total variation method. Note also that changing the parameter ω implicitly affects the amount of regularization applied. For an image u , decreasing ω will, all else being equal, decrease the lowest eigenvalue of $A(x)$ and in turn decrease the anisotropic total variation $\text{TV}_A(u)$.

For the case where $\lambda_1 = \lambda_2$, the $U(x)$ in our decomposition is not well-defined. This is not a problem however, since $\Sigma(x)$ will be the identity matrix, so any orthogonal matrix will suffice for $U(x)$.

Note that the eigenvalues of S_ρ are continuous, and so are the eigenvectors (ignoring their sign) except possibly when $\lambda_1 = \lambda_2$. Thus A is also continuous except possibly in these points. When $\lambda_1 = \lambda_2$ however, the eigenvalues σ_1 and σ_2 of A will both be 1, and A is the identity matrix. Thus we can argue that if $S_\rho(x) \rightarrow \lambda I$ then $A(x) \rightarrow I$ and A is continuous in all of Ω .

See [10] for a different tensor construction, made to enhance flow structures in the image, relevant in for example fingerprint analysis.

3.2 Well-posedness

The theory of existence and uniqueness for these kinds of variational methods is a minefield of more or less subtle problems. Even if we restrict ourselves to a nice space such as $L^2(\Omega)$ we will at some point run into problems. The discussion here is not meant to give the most rigorous background, but rather an overview of what needs to be shown. Some problems will be worked around, while others will be skipped with a reference to further theory.

The basic things we ask of our functional

$$F(u) = \int_{\Omega} (u - f)^2 + \beta \text{TV}_A(u) \quad (3.11)$$

to have a *well-posed* problem are lower semi-continuity and coercivity for existence, and convexity for uniqueness. We restrict ourself to $L^2(\Omega)$ which makes sense with our fidelity term, assuming that $f \in L^2(\Omega)$.

We consider the weak topology, as it will allow us to arrive at an existence result relatively easily. We say that a sequence f_n in $L^2(\Omega)$ converges weakly to f if

$$\lim_{n \rightarrow \infty} \int_{\Omega} f_n \xi \, dx = \int_{\Omega} f \xi \, dx \quad (3.12)$$

for all $\xi \in L^2(\Omega)$ and we write $f_n \rightharpoonup f$. A weakly convergent sequence is a sequence that converges in the weak topology.

3.2.1 Convexity

We start with convexity as it is the easiest to show. Being quadratic, the fidelity term of our functional

$$\int_{\Omega} (u - f)^2 \, dx \quad (3.13)$$

is obviously strictly convex. This can be shown by expanding and rearranging the strict convexity condition

$$\int_{\Omega} (\lambda u_1 + (1 - \lambda)u_2 - f)^2 \, dx < \lambda \int_{\Omega} (u_1 - f)^2 \, dx + (1 - \lambda) \int_{\Omega} (u_2 - f)^2 \, dx \quad (3.14)$$

to obtain that it is equivalent to

$$-\lambda(1 - \lambda) \int_{\Omega} (u_1 - u_2)^2 \, dx < 0 \quad (3.15)$$




Figure 3.3: A lower semi-continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$ can have discontinuities, but for a convergent sequence $x_k \rightarrow x$ we always have $f(x) \leq \liminf_{k \rightarrow \infty} f(x_k)$.

which is true for $0 < \lambda < 1$ and $u_1 \neq u_2$.

The anisotropic total variation

$$\mathrm{TV}_A(u) = \sup_{\|\xi\|_A^* \leq 1} \int_{\Omega} u \operatorname{div} \xi \, dx \quad (3.16)$$

can be thought of as—and has the properties of—a semi-norm, and is therefore convex. The sum of the fidelity and regularization terms is thus strictly convex, which, given the existence of a minimizer, implies uniqueness.

3.2.2 Coercivity

Coercivity relates to how the functional behaves when the norm of the image u tends to infinity. What we need in order to conclude with existence is weak sequential coercivity. Thus we need all level sets $F^\alpha = \{u \in L^2(\Omega) : F(u) \leq \alpha\}$ to be *weakly sequentially pre-compact*, meaning that all sequences in the set contain a subsequence weakly converging to an element of the closure of the set.

It is obvious from the fidelity term that for some fixed $f \in L^2(\Omega)$, if $\|u\|_{L^2} \rightarrow \infty$ then $F(u) \rightarrow \infty$. This implies that all the level sets F^α are bounded. Since $L^2(\Omega)$ is a Hilbert space all bounded sequences contain a weakly convergent subsequence. Thus all the level sets F^α are weakly sequentially pre-compact.

3.2.3 Lower semi-continuity

The lower semi-continuity is the most tricky part, and this is where we will take some shortcuts. Lower semi-continuity for a functional F at a point u means that at points u_ϵ close to u , the functional takes values either close to or above $F(u)$. More specifically, for every sequence u_k converging to u , we have $F(u) \leq \liminf_k F(u_k)$. For a function $f : \mathbb{R} \rightarrow \mathbb{R}$ this can be visualized as in Figure 3.3.

Since our space $L^2(\Omega)$ is of infinite dimensions things become a little problematic here. The problem lies in the fact that a functional which is continuous with respect to sequences is not necessarily continuous with respect to the underlying topology. In other words, in these spaces, there can be a difference between sequential continuity and topological continuity. Topological continuity implies sequential continuity, but the converse does not hold. One way to get around this would be to consider topological *nets*, an extension of sequences, but for simplicity, and because it might not add much to the understanding of the restoration method, we will stick to proving sequential lower semi-continuity and referring to further theory. For further reading on the theory of sequential versus topological continuity see for example Megginson's book on Banach space theory [11].

The mapping $u \mapsto \int_{\Omega} u \xi dx$ is weakly continuous for all $\xi \in L^2(\Omega)$. Note that when we write weakly continuous it is not a weaker version of continuity, but rather continuity in the weak topology, and the same goes for weak lower semi-continuity.

Before arguing that our own functional is sequentially weakly lower semi-continuous, we present a needed result.

Lemma 3.2. *Assume that the functional $F : L^2(\Omega) \rightarrow \mathbb{R}$ is defined by*

$$F = \sup_i F_i \tag{3.17}$$

where all the F_i are sequentially weakly lower semi-continuous, then F is sequentially weakly lower semi-continuous, meaning that for any sequence $u_k \rightharpoonup u$ we have $F(u) \leq \liminf_k F(u_k)$.

Proof. For any sequence $u_k \rightharpoonup u$ in $L^2(\Omega)$ we have

$$F(u) = \sup_i F_i(u) \leq \sup_i \liminf_{k \rightarrow \infty} F_i(u_k) \tag{3.18}$$

from the sequential weak lower semi-continuity of F_i . Using that $\liminf_{k \rightarrow \infty} u_k = \sup_k \inf_{l \geq k} u_l$, we obtain

$$\begin{aligned} F(u) &\leq \sup_i \sup_k \inf_{l \geq k} F_i(u_l) \\ &= \sup_k \sup_i \inf_{l \geq k} F_i(u_l) \\ &\leq \sup_k \inf_{l \geq k} \sup_i F_i(u_l) \\ &= \liminf_{k \rightarrow \infty} F(u_k) \end{aligned} \tag{3.19}$$

which proves that F is sequentially weakly lower semi-continuous. \square

In our functional in (3.4), we first consider the fidelity term, and rewrite it as a supremum

$$\int_{\Omega} (u - f)^2 dx = \sup \left\{ \int_{\Omega} (u - f) \xi dx : \xi \in L^2(\Omega), |\xi(x)| \leq |u(x) - f(x)| \right\} \quad (3.20)$$

As the map $u \mapsto \int_{\Omega} (u - v) \xi dx$ is continuous in the weak topology, the fidelity term is thus a supremum of weakly continuous functionals, and is thus by Lemma 3.2 sequentially lower semi-continuous.

For the regularization term the approach is similar. With our extended definition from (3.3), we have

$$\mathrm{TV}_A(u) = \sup \left\{ \int_{\Omega} u \operatorname{div} \xi dx : \xi \in C_c^\infty(\Omega, \mathbb{R}^2), \|\xi\|_A^* \leq 1 \right\} \quad (3.21)$$

This is again a supremum of weakly continuous functionals. Thus the regularization term is by Lemma 3.2 also sequentially weakly lower semi-continuous.

The sum of the two terms is trivially sequentially weakly lower semi-continuous functional since

$$\begin{aligned} F_1(u) + F_2(u) &\leq \liminf_{k \rightarrow \infty} F_1(u_k) + \liminf_{k \rightarrow \infty} F_2(u_k) \\ &= \lim_{k \rightarrow \infty} \left(\inf_{l \geq k} F_1(u_l) + \inf_{l \geq k} F_2(u_l) \right) \\ &\leq \liminf_{k \rightarrow \infty} (F_1(u_k) + F_2(u_k)), \end{aligned} \quad (3.22)$$

and thus our functional is sequentially weakly lower semi-continuous.

The usual ways of going from coercivity and lower semi-continuity to existence do not work in infinite dimensions. But with sequential coercivity and sequential lower semi-continuity in the weak topology we can conclude that we have existence from [12, Theorem 5.1].

3.3 Anisotropic coarea formula

The anisotropic coarea formula we present here will allow us to write the anisotropic total variation as an integral over the levels of the image. For a similar presentation of the regular coarea formula for all $f \in \mathrm{BV}(\Omega)$ see [13].

First we define the thresholded image at level s .

Definition 3.3 (Thresholded image). *The thresholded image at level s is the function*

$$u^s(x) = \begin{cases} 1 & \text{if } u(x) > s, \\ 0 & \text{otherwise.} \end{cases} \quad (3.23)$$

This will be used throughout the rest of the thesis. Note that given the thresholded image for every level, we are able to reconstruct the image as

$$u(x) = \sup \{s : u^s(x) = 1\}. \quad (3.24)$$

The thresholded image definition also allows us to write a non-negative image $u \geq 0$ as an integral over all the layers

$$u(x) = \int_0^\infty u^s(x) ds. \quad (3.25)$$

Note that (3.25) only holds for non-negative images, which complicates the proof of the anisotropic coarea formula a little.

Theorem 3.4 (Anisotropic coarea formula). *Given an image $u \in \text{BV}(\Omega)$, the anisotropic total variation can be written as an integral over all the levels*

$$\text{TV}_A(u) = \int_{-\infty}^\infty \text{TV}_A(u^s) ds. \quad (3.26)$$

For the proof we will avoid measure theory and follow a proof given in [9], but first we will present a necessary result from measure theory.

Theorem 3.5 (Lebesgue's Dominated Convergence theorem). *Let $\{f_n\}$ be a sequence of real-valued measurable functions on a space S with measure $d\mu$ which converges almost everywhere to a real-valued measurable function f . If there exists an integrable function g such that $|f_n| \leq g$ for all n , then f is integrable and*

$$\lim_{n \rightarrow \infty} \int_S f_n d\mu = \int_S f d\mu. \quad (3.27)$$

For a proof and further background on measure theory and Lebesgue integration theory see for example [14].

Proof of the anisotropic coarea formula. Assume that $u \in C^1(\Omega) \cap \text{BV}(\Omega)$. The extension to all functions $u \in \text{BV}(\Omega)$ will not be considered here, but for the case of regular total variation see [15, Theorem 5.3.3].

Proof of upper bound. Assume that $u \geq 0$ such that the integral representation in (3.25) holds, then inserting (3.25) into the extended total variation definition in (3.3) gives

$$\begin{aligned} \text{TV}_A(u) &= \sup_{\|\xi\|_A^* \leq 1} \int_\Omega \left(\int_0^\infty u^s ds \right) \text{div} \xi dx = \sup_{\|\xi\|_A^* \leq 1} \int_\Omega \int_0^\infty u^s \text{div} \xi ds dx \\ &\leq \int_0^\infty \left(\sup_{\|\xi\|_A^* \leq 1} \int_\Omega u^s \text{div} \xi dx \right) ds = \int_0^\infty \text{TV}_A(u^s) ds. \end{aligned} \quad (3.28)$$

For $u \leq 0$ we use that $\text{TV}_A(-v) = \text{TV}_A(v)$ and that $\text{TV}_A(c + v) = \text{TV}_A(v)$ for any constant c . Note that $-u \geq 0$ and that its thresholded image $(-u)^s$ will be exactly the opposite of u^{-s} , that is $(-u)^s = 1 - u^{-s}$. This allows us to show that

$$\begin{aligned} \text{TV}_A(u) &= \text{TV}_A(-u) \leq \int_0^\infty \text{TV}_A((-u)^r) dr = \int_0^\infty \text{TV}_A(1 - u^{-r}) dr \\ &= \int_0^\infty \text{TV}_A(u^{-r}) dr = \int_{-\infty}^0 \text{TV}_A(u^s) ds. \end{aligned} \quad (3.29)$$

Following from the supremum definition of the anisotropic total variation in (3.3), we obtain the inequality

$$\begin{aligned} \text{TV}_A(u_1 + u_2) &= \sup_{\|\xi\|_A^* \leq 1} \int_{\Omega} (u_1 + u_2) \operatorname{div} \xi \, dx \\ &\leq \sup_{\|\xi\|_A^* \leq 1} \int_{\Omega} u_1 \operatorname{div} \xi \, dx + \sup_{\|\xi\|_A^* \leq 1} \int_{\Omega} u_2 \operatorname{div} \xi \, dx \\ &= \text{TV}_A(u_1) + \text{TV}_A(u_2). \end{aligned} \quad (3.30)$$

Next, we write a general u as a difference of two positive functions $u = u_+ - u_-$ where $u_+ = \max\{u, 0\}$ and $u_- = -\min\{u, 0\}$. Inserting (3.28) and (3.29) into (3.30) we obtain

$$\begin{aligned} \text{TV}_A(u) &\leq \text{TV}_A(u_-) + \text{TV}_A(u_+) = \text{TV}_A(-u_-) + \text{TV}_A(u_+) \\ &\leq \int_{-\infty}^0 \text{TV}_A((-u_-)^s) ds + \int_0^\infty \text{TV}_A(u_+^s) ds \\ &= \int_{-\infty}^0 \text{TV}_A(u^s) ds + \int_0^\infty \text{TV}_A(u^s) ds = \int_{-\infty}^\infty \text{TV}_A(u^s) ds. \end{aligned} \quad (3.31)$$

Note that u_+ and u_- will not be differentiable everywhere, but we did not use the differentiability of u in this part of the proof.

Proof of lower bound. Define the function

$$m(t) = \int_{\{x \in \Omega : u(x) \leq t\}} \|\nabla u\|_A \, dx, \quad (3.32)$$

and note that $m(\infty) = \text{TV}_A(u)$ and $m(-\infty) = 0$. Since $m(t)$ is non-decreasing with t , we can apply the existence theorems of Lebesgue [16, Thm. 17.12, 18.14] to conclude that $m'(t)$ exists almost everywhere and that the following inequality holds:

$$\int_{-\infty}^\infty m'(t) \, dt \leq m(\infty) - m(-\infty) = \text{TV}_A(u). \quad (3.33)$$




Figure 3.4: Visualization of the cut-off function $\eta_r(t)$ and its derivative.

Next, fix an $s \in \mathbb{R}$ and define the cut-off function

$$\eta_r(t) = \begin{cases} 0 & \text{if } t < s, \\ (t-s)/r & \text{if } s \leq t < s+r, \\ 1 & \text{if } t \geq s+r, \end{cases} \quad \eta'_r(t) = \begin{cases} 0 & \text{if } t < s, \\ 1 & \text{if } s < t < s+r, \\ 0 & \text{if } t > s+r, \end{cases} \quad (3.34)$$

visualized in Figure 3.4. By composing the function η_r with our image u and using Green's formula, for example from [8, Corollary 9.32] we obtain

$$\int_{\Omega} -\eta_r(u) \operatorname{div} \xi \, dx = \int_{\Omega} \eta'_r(u) \nabla u \cdot \xi \, dx = \frac{1}{r} \int_{\{s < u < s+r\}} \nabla u \cdot \xi \, dx, \quad (3.35)$$

for all vector fields $\xi \in C_c^\infty(\Omega, \mathbb{R}^2)$. The measure of $\{x : u(x) = \lambda \text{ and } \nabla u(x) \neq 0\}$ is zero for all λ following from [13, Corollary I, Section 3.1.2], thus we can ignore the sets $\{u = s\}$ and $\{u = s+r\}$ in the integrals. Assuming that $\|\xi\|_A^* \leq 1$ we obtain from (3.32) and (3.35) that

$$\begin{aligned} \frac{m(s+r) - m(s)}{r} &= \frac{1}{r} \int_{\{s < u \leq s+r\}} \|\nabla u\|_A \, dx \\ &\geq \frac{1}{r} \int_{\{s < u \leq s+r\}} \nabla u \cdot \xi \, dx = \int_{\Omega} -\eta_r(u) \operatorname{div} \xi \, dx. \end{aligned} \quad (3.36)$$

As the limit when $r \rightarrow 0$ of the left-hand side exists almost everywhere, suppose it exists at $s \in \mathbb{R}$. The integrand on the right-hand side $-\eta_r(u) \operatorname{div} \xi$ approaches $-u^s \operatorname{div} \xi$ pointwise almost everywhere. We also have the bound $|\eta_r(u) \operatorname{div} \xi| \leq |u^s \operatorname{div} \xi|$ and know that $\xi \in C_c^\infty(\Omega, \mathbb{R}^2)$ is bounded following from the extreme value theorem. Thus we can apply Lebesgue's dominated convergence theorem, giving that $|u^s \operatorname{div} \xi|$ is integrable and

$$\int_{\Omega} -\eta_r(u) \operatorname{div} \xi \, dx \rightarrow \int_{\Omega} -u^s \operatorname{div} \xi \, dx \quad (3.37)$$

From (3.36) we then obtain

$$m'(s) \geq - \int_{\Omega} u^s \operatorname{div} \xi \, dx. \quad (3.38)$$

As this holds for any $\|\xi\|_A^* \leq 1$, we get from the extended total variation definition in (3.3) that $m'(s) \geq \operatorname{TV}_A(u^s)$ almost everywhere and conclude using (3.33) that

$$\operatorname{TV}_A(u) \geq \int_{-\infty}^{\infty} m'(t) \, dt \geq \int_{-\infty}^{\infty} \operatorname{TV}_A(u^s) \, ds. \quad (3.39)$$

Combining the upper and lower bounds just proved, we have equality. \square

This coarea formula is our first step in transforming the anisotropic total variation into an easily discretizable expression. It allows us to consider each level λ separately when calculating the anisotropic total variation.

The anisotropic total variation of the thresholded images occurring in the anisotropic coarea formula is very much related to the size of the boundary of the level set, as the only variation in a characteristic function occurs at the boundary of the set. This is why we introduce the following definition of the anisotropic set perimeter.

Definition 3.6 (The anisotropic set perimeter). *Given an anisotropy tensor A the anisotropic perimeter of a set U in Ω is defined as*

$$\operatorname{Per}_A(U; \Omega) = \operatorname{TV}_A(\chi_U). \quad (3.40)$$

The anisotropic set perimeter is not like the regular set perimeter and does not measure the length of the boundary of the set, but it can for sufficiently nice level sets be calculated in the following way

$$\begin{aligned} \operatorname{Per}_A(\{u > s\}; \Omega) &= \operatorname{TV}_A(u^s) \\ &= \sup_{\|\xi\|_A^* \leq 1} \int_{\Omega} u^s \operatorname{div} \xi \, dx \\ &= \sup_{\|\xi\|_A^* \leq 1} \int_{\{u > s\}} \operatorname{div} \xi \, dx \\ &= \sup_{\|\xi\|_A^* \leq 1} \int_{\partial\{u > s\}} \nu_s \cdot \xi \, dt \\ &= \sup_{\|\eta\| \leq 1} \int_{\partial\{u > s\}} \nu_s \cdot A^{1/2} \eta \, dt \\ &= \int_{\partial\{u > s\}} \sqrt{\nu_s A \nu_s} \, dt. \end{aligned} \quad (3.41)$$




Figure 3.5: The blue line is parametrized by the angle ϕ and the distance from the origin to the line ρ , or alternatively, the pair (ν, ρ) .

Here, ν_s is the unit exterior normal of the level set $\{u > s\}$. Note that because of the compact support of ξ in Definition 3.1, the parts of the boundary of U that overlap with the boundary of Ω will not be included in the perimeter.

Exterior normals and perimeters of level sets of any function $u \in \text{BV}(\Omega)$ will not be considered here, but can for the isotropic case be found in for example [15, Section 5.4 and 5.5].

Using the anisotropic coarea formula and inserting the anisotropic perimeter definition we transform the anisotropic total variation and are left with the problem of minimizing the following functional

$$F(u) = \int_{\Omega} (u - f)^2 dx + \beta \int_{-\infty}^{\infty} \text{Per}_A(\{u > \lambda\}; \Omega) d\lambda. \quad (3.42)$$

The transformation is motivated by our upcoming anisotropic Cauchy–Crofton integration formula, and the discretization, where an approximation of the perimeter will be computed using a graph cut machinery.

3.4 Anisotropic Cauchy–Crofton formula

In the fields of integral geometry and geometric measure theory there are a number of interesting integral formulas. Several of them fall in a category often referred to as *Cauchy–Crofton style formulas*, and give ways to measure geometric objects using the set of all lines in the plane. The formulas presented here will give a way to measure the length of a curve by counting the times it intersects lines in the set of all lines. The first formula will be for the isotropic case, and we will use it to prove the anisotropic formula following it.

We write \mathcal{L} for the set of all lines in the plane, and parametrize them as shown in Figure 3.5. A line is parametrized by the angle $\phi \in [0, 2\pi)$ of the normal going to the origin, and the distance $\rho \in [0, \infty)$ from origin to the line. Sometimes it is more convenient to consider a unit vector ν giving the direction of the line instead of the angle parameter ϕ . We denote a line by $\ell_{\phi, \rho} = \ell_{\nu, \rho}$ where ν is a unit vector along the line, i.e. $\nu = (-\sin \phi, \cos \phi)^T$. By defining the measure on this set $d\mathcal{L} = d\phi d\rho$ we are ready to introduce the Cauchy–Crofton formula. Note that the measure $d\mathcal{L}$ is invariant under rotations.

Theorem 3.7 (The Euclidean Cauchy–Crofton formula). *Given a differentiable curve C in \mathbb{R}^2 , the length of this curve $|C|$ is related to the set of lines \mathcal{L} as follows*

$$\int_{\mathcal{L}} \#(\ell_{\phi, \rho} \cap C) d\mathcal{L}(\ell_{\phi, \rho}) = 2|C|, \quad (3.43)$$

where $\#(\ell_{\phi, \rho} \cap C)$ is the number of times the line $\ell_{\phi, \rho}$ intersects the curve C .

Proof. See [17, Theorem 3, Section 1-7]. □

If our space is equipped with a metric tensor $M(x)$ such that the inner product of two vectors a and b in a point x is calculated as $\langle a, b \rangle_M = \langle a, M(x)b \rangle$, then the length of a curve γ parametrized by some parameter t becomes

$$|\gamma|_M = \int_{\gamma} \sqrt{\langle \dot{\gamma}, M(\gamma(t)) \dot{\gamma} \rangle} dt. \quad (3.44)$$

We will now present and prove a Cauchy–Crofton formula in this case where our domain is equipped with a metric tensor in each point. This elegant formula is very useful when we later will discretize our perimeter calculation. The set of lines \mathcal{L} is then discretized in a reasonable way, and the length of the curve C can be approximated by a sum over all these lines.

Theorem 3.8 (The anisotropic Cauchy–Crofton formula). *Assume that our space Ω is equipped with a continuous positive definite metric tensor $M(x)$, whose eigenvalues are bounded by $0 < k \leq \lambda_2 \leq \lambda_1 \leq K < \infty$ for all $x \in \Omega$. The Cauchy–Crofton formula for a differentiable curve C of finite length then becomes*

$$|C|_M = \int_{\mathcal{L}} \sum_{x \in \ell_{\nu, \rho} \cap C} \frac{\det M(x)}{2(\nu^T \cdot M(x) \cdot \nu)^{3/2}} d\mathcal{L}(\ell_{\nu, \rho}). \quad (3.45)$$

Proof of the anisotropic Cauchy–Crofton formula. Assume first that our space is equipped with a constant metric tensor M . The length of a curve in this space

can be calculated by transforming the curve and applying the Euclidean Cauchy–Crofton formula

$$|C|_M = \int_C \sqrt{\langle \dot{C}, M \dot{C} \rangle} dt = \int_C \sqrt{\langle M^{1/2} \dot{C}, M^{1/2} \dot{C} \rangle} = |M^{1/2} C| \quad (3.46)$$

$$= \int_{\mathcal{L}} \#(\ell_{\phi, \rho} \cap M^{1/2} C) d\mathcal{L}(\ell_{\phi, \rho}) \quad (3.47)$$

$$= \int_{\mathcal{L}} \#(M^{-1/2} \ell_{\phi, \rho} \cap C) d\mathcal{L}(\ell_{\phi, \rho}) \quad (3.48)$$

$$= \int_{\mathcal{L}} \#(m_{\phi, \rho} \cap C) |J_M(\ell_{\phi, \rho})| d\mathcal{L}(m_{\phi, \rho}). \quad (3.49)$$

Here $J_M(\ell_{\phi, \rho})$ is the Jacobian of the coordinate transformation $F : \mathcal{L} \rightarrow \mathcal{L}$, which maps $\ell_{\phi, \rho} \mapsto M^{1/2} \ell_{\phi, \rho}$.

We will now compute the Jacobian $J_M(\ell_{\phi, \rho})$. As $M \in \mathbb{R}^{2 \times 2}$ is symmetric, so is $M^{1/2}$, and it admits a decomposition $M^{1/2} = U \Sigma U^T$ where the components correspond to the following coordinate transformations

$$U(\ell_{\nu, \rho}) = \ell_{\phi + \xi, \rho} = \ell_{U\nu, \rho} \quad (3.50)$$

$$U^T(\ell_{\nu, \rho}) = \ell_{\phi - \xi, \rho} = \ell_{U^T\nu, \rho} \quad (3.51)$$

$$\Sigma = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} = \begin{pmatrix} \sqrt{\lambda_1} & 0 \\ 0 & \sqrt{\lambda_2} \end{pmatrix} \quad (3.52)$$

As U and U^T correspond to rotations and our measure \mathcal{L} is invariant under rotations, U and U^T do not have direct contributions to the Jacobian. They do however affect the input angle of the operator Λ such that $J_M(\ell_{\phi, \rho}) = J_{\Sigma^2}(U^T \ell_{\phi, \rho})$. Thus we will now compute $J_{\Sigma^2}(\ell_{\phi, \rho})$. Given a line

$$\ell_{\phi, \rho} = \begin{pmatrix} \rho \cdot \cos \phi \\ \rho \cdot \sin \phi \end{pmatrix} + \mathbb{R} \begin{pmatrix} -\sin \phi \\ \cos \phi \end{pmatrix}, \quad (3.53)$$

the operator Σ transforms it into

$$\Sigma \ell_{\phi, \rho} = \begin{pmatrix} \sigma_1 \rho \cdot \cos \phi \\ \sigma_2 \rho \cdot \sin \phi \end{pmatrix} + \mathbb{R} \begin{pmatrix} -\sigma_1 \sin \phi \\ \sigma_2 \cos \phi \end{pmatrix}, \quad (3.54)$$

which equals the line $\ell_{\theta, \eta}$ with

$$\theta = \arctan \left(\frac{\sigma_1}{\sigma_2} \tan \phi \right) \quad (3.55)$$

$$\eta = \left\langle \begin{pmatrix} \sigma_1 \rho \cdot \cos \phi \\ \sigma_2 \rho \cdot \sin \phi \end{pmatrix}, \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix} \right\rangle = \sigma_1 \rho \cdot \cos \phi \cdot \cos \theta + \sigma_2 \rho \cdot \sin \phi \cdot \sin \theta. \quad (3.56)$$

As $\partial_\rho \theta = 0$, the Jacobian becomes $|J_{\Sigma^2}(\ell_{\phi,\rho})| = \partial_\phi \theta \cdot \partial_\rho \eta$. Differentiation yields

$$\partial_\phi \theta = \frac{\frac{\sigma_1}{\sigma_2} \sec^2 \phi}{1 + \frac{\sigma_1^2}{\sigma_2^2} \tan^2 \phi} = \frac{\sigma_1 \sigma_2}{\sigma_1^2 \sin^2 \phi + \sigma_2^2 \cos^2 \phi}, \quad (3.57)$$

$$\partial_\rho \eta = \sigma_1 \cos \phi \cdot \cos \theta + \sigma_2 \sin \phi \cdot \sin \theta. \quad (3.58)$$

In the expression for $\partial_\rho \eta$ we insert θ from (3.55) and use that $\sin(\arctan(x)) = x/\sqrt{1+x^2}$ and that $\cos(\arctan(x)) = 1/\sqrt{1+x^2}$ to obtain

$$\partial_\rho \eta = \frac{\sigma_1 \cos \phi + \sigma_2 \sin \phi \frac{\sigma_1}{\sigma_2} \tan \phi}{\sqrt{1 + \frac{\sigma_1^2}{\sigma_2^2} \tan^2 \phi}} = \frac{\sigma_1 \sigma_2}{\sqrt{\sigma_1^2 \sin^2 \phi + \sigma_2^2 \cos^2 \phi}}. \quad (3.59)$$

If $\nu = (\nu_x, \nu_y)^T$ is a unit vector along the line $\ell_{\phi,\rho} = \ell_{\nu,\rho}$ then

$$|J_{\Sigma^2}(\ell_{\nu,\rho})| = \frac{\sigma_1^2 \sigma_2^2}{(\sigma_1^2 \sin^2 \phi + \sigma_2^2 \cos^2 \phi)^{3/2}} = \frac{\sigma_1^2 \sigma_2^2}{(\sigma_1^2 \nu_x^2 + \sigma_2^2 \nu_y^2)^{3/2}} = \frac{\det \Sigma^2}{(\nu^T \cdot \Sigma^2 \cdot \nu)^{3/2}}. \quad (3.60)$$

We are interested in the Jacobian of the whole transformation $J_{\Sigma^2}(U^T \ell_{\nu,\rho})$, so all that is left to do is insert $U^T \ell_{\nu,\rho}$ to obtain

$$|J_M(\ell_{\nu,\rho})| = |J_{\Sigma^2}(U^T \ell_{\nu,\rho})| = \frac{\det M}{(\nu^T U \cdot \Sigma^2 \cdot U^T \nu)^{3/2}} = \frac{\det M}{(\nu^T \cdot M \cdot \nu)^{3/2}} \quad (3.61)$$

We have now proved that for a constant metric tensor M , the length of the differentiable curve C with regards to this tensor can be calculated as

$$|C|_M = \int_C \sqrt{\langle \dot{C}, M \dot{C} \rangle} dt = \int_{\mathcal{L}} \#(\ell_{\nu,\rho} \cap C) \frac{\det M}{(\nu^T \cdot M \cdot \nu)^{3/2}} d\mathcal{L}(\ell_{\nu,\rho}). \quad (3.62)$$

Further we argue that the similar formula in (3.45) holds for a non-constant but continuous metric tensor $M(x)$. By partitioning the domain into disjoint sets U_i such that $\Omega = \cup_i U_i$, we make a piecewise constant approximation $M_\pi(x)$ such that if $x \in U_i$ then $M_\pi(x) = M(x_i)$ for some fixed $x_i \in U_i$. We then approximate (3.62) by

$$|C|_{M_\pi} = \sum_i \int_{\mathcal{L}} \#(\ell_{\nu,\rho} \cap C \cap U_i) w_i(\nu) d\mathcal{L}(\ell_{\nu,\rho}) \quad (3.63)$$

where w_i is the weight-function used in the set U_i , that is,

$$w_i(\nu) = \frac{\det M(x_i)}{(\nu^T \cdot M(x_i) \cdot \nu)^{3/2}}. \quad (3.64)$$

We further simplify the approximation by introducing the global weight-function $w_\pi(\nu, x)$ which is equal to $w_i(\nu)$ when $x \in U_i$. It can be written as

$$w_\pi(\nu, x) = \frac{\det M_\pi(x)}{(\nu^T \cdot M_\pi(x) \cdot \nu)^{3/2}}. \quad (3.65)$$

Using this weight in (3.63) we can get rid of the sum over the partition i and form a sum of all intersection point of C and the line $\ell_{\nu, \rho}$ currently being integrated over. The approximation becomes

$$\begin{aligned} |C|_{M_\pi} &= \sum_i \int_{\mathcal{L}} \sum_{x \in \ell_{\nu, \rho} \cap C \cap U_i} w_\pi(\nu, x) d\mathcal{L}(\ell_{\nu, \rho}) \\ &= \int_{\mathcal{L}} \sum_{x \in \ell_{\nu, \rho} \cap C} w_\pi(\nu, x) d\mathcal{L}(\ell_{\nu, \rho}). \end{aligned} \quad (3.66)$$

Now it only remains to show that the left- and right-hand side of (3.66) converges to the left- and right-hand side of (3.45).

As our partition π is refined, the weight $w_\pi(x)$ converges pointwise to the continuously varying weight

$$w(\nu, x) = \frac{\det M(x)}{(\nu^T \cdot M(x) \cdot \nu)^{3/2}} \quad (3.67)$$

found in (3.45).

Recall from (3.44) that the left-hand side is calculated as

$$|C|_{M_\pi} = \int_C |\dot{C}(t)|_{M_\pi} dt = \int_C \sqrt{\dot{C}(t)^T M_\pi(C(t)) \dot{C}(t)} dt. \quad (3.68)$$

We know that $M_\pi(x)$ converges pointwise to $M(x)$, and thus $|\dot{C}(t)|_{M_\pi}$ converges pointwise to $|\dot{C}(t)|_M$. We have assumed bounds on the eigenvalues of $M(x)$ such that, according to the Rayleigh principle

$$K \geq \lambda_1 = \max_{\xi} \frac{\xi^T M_\pi(x) \xi}{\xi^T \xi} \quad (3.69)$$

and therefore we have the bound

$$\xi^T M_\pi(x) \xi \leq K \|\xi\|^2, \quad \forall \xi. \quad (3.70)$$

Thus the integrand of (3.68) is bounded by $g(t) = (K \cdot \dot{C}(t)^T \dot{C}(t))^{1/2}$. We know that $g(t)$ is integrable as its integral is exactly $\sqrt{K} |C|$ and we have assumed that

the curve is of finite length. This means we can apply Lebesgue’s dominated convergence theorem to see that $|C|_{M_\pi} \rightarrow |C|_M$.

We apply the same theorem to show that the right-hand side of (3.66) converges. Recall the definition of w_π in (3.65). The numerator is equal to $\sigma_1^2 \sigma_2^2 = \lambda_1 \lambda_2$ and is by assumption bounded from above by K^2 .

Next we need to bound $\nu^T M_\pi(x) \nu$ away from zero. According to the Rayleigh principle

$$\lambda_2 = \min_{\|\xi\|=1} \xi^T M_\pi(x) \xi \quad (3.71)$$

and thus $\nu^T M_\pi(x) \nu \geq \lambda_2 \geq k$. The weight function w_π is then bounded such that

$$\sum_{x \in \ell_{\nu,\rho} \cap C} w_\pi(\nu, x) \leq \sum_{x \in \ell_{\nu,\rho} \cap C} \frac{K^2}{k^{3/2}} = \frac{K^2}{k^{3/2}} \cdot \#(\ell_{\nu,\rho} \cap C) =: g(\ell_{\nu,\rho}). \quad (3.72)$$

This is integrable following from the Euclidean Cauchy–Crofton formula in Theorem 3.7 and the fact that we assumed C to be of finite length:

$$\int_{\mathcal{L}} g(\ell_{\nu,\rho}) d\mathcal{L}(\ell_{\nu,\rho}) = \frac{K^2}{k^{3/2}} |C| < \infty. \quad (3.73)$$

Thus we can apply the dominated convergence theorem again and conclude that

$$\int_{\mathcal{L}} \sum_{x \in \ell_{\nu,\rho} \cap C} w_\pi(\nu, x) d\mathcal{L}(\ell_{\nu,\rho}) \rightarrow \int_{\mathcal{L}} \sum_{x \in \ell_{\nu,\rho} \cap C} w(\nu, x) d\mathcal{L}(\ell_{\nu,\rho}) \quad (3.74)$$

which—as both sides of the equality in (3.66) have been shown to converge—leaves us with what we wanted to prove

$$|C|_M = \int_{\mathcal{L}} \sum_{x \in \ell_{\nu,\rho} \cap C} \frac{\det M(x)}{2(\nu^T \cdot M(x) \cdot \nu)^{3/2}} d\mathcal{L}(\ell_{\nu,\rho}). \quad (3.75)$$

□

With the anisotropic coarea formula in Theorem 3.4 we have a way to calculate the anisotropic total variation by integrating the anisotropic perimeter of each level set of the image. We will now see how the anisotropic Cauchy–Crofton formula can help us calculate the perimeters of the level sets. In the Euclidean case, which here would amount to setting the anisotropy tensor A equal to the identity matrix I , the perimeter coincides nicely with the length of the boundary curve, assuming some regularity for the boundary. In the general case we need to be more careful. As can be seen in (3.41), the anisotropic perimeter is calculated by integrating the norm of the *normal vector* around the boundary, while the anisotropic curve

length in (3.44) is the integral of the norm of the *tangent vector* of the curve. Thus a 90° rotation separates the two.

If P is a 90° rotation matrix we have

$$\begin{aligned} \text{Per}_A(U; \Omega) &= \int_{\partial U} \sqrt{\langle \nu_{\partial U}, A(x) \nu_{\partial U} \rangle} dt \\ &= \int_{\partial U} \sqrt{\langle P \nu_{\partial U}, PA(x) P^T P \nu_{\partial U} \rangle} dt. \end{aligned} \quad (3.76)$$

We simplify the equation by defining the metric tensor $M(x) = PA(x)P^T$ and letting $\gamma = \partial U \cap \Omega$ be an arclength parametrization of the boundary of U that does not overlap with the boundary of Ω

$$\text{Per}_A(U; \Omega) = \int_{\gamma} \sqrt{\langle \dot{\gamma}, M(x) \dot{\gamma} \rangle} dt. \quad (3.77)$$

Now we make sure that all the assumptions of the anisotropic Cauchy–Crofton formula in Theorem 3.8 are fulfilled so that it can be applied to the curve length integral we have constructed in (3.77).

The structure tensor is constructed as described in Section 3.1.1

$$S_{\rho}(x) = (K_{\rho} * (\nabla f_{\sigma} \otimes \nabla f_{\sigma}))(x). \quad (3.78)$$

Because of the convolutions with the Gaussian function, this is a smooth map from $\bar{\Omega}$ to $\mathbb{R}^{2 \times 2}$. As we can see in (3.6), the eigenvalues depend continuously on the coefficients of the elements in the structure tensor $S_{\rho}(x)$. The extreme value theorem states that a continuous real-valued function on a nonempty compact space is bounded above. Thus the eigenvalues $\lambda_1 \geq \lambda_2$ of $S_{\rho}(x)$ are bounded above. Moreover, by the construction in (3.10), there exists uniform bound k such that the smallest eigenvalue σ_1 of the anisotropy tensor $A(x)$ is bounded away from zero, as

$$\sigma_1 = \left(1 + \frac{(\lambda_1 - \lambda_2)^2}{\omega^2}\right)^{-1} \geq \left(1 + \frac{\lambda_1^2}{\omega^2}\right)^{-1} \geq k > 0. \quad (3.79)$$

Hence our metric tensor $M(x) = PA(x)P^T$ is continuous and positive definite with bounded eigenvalues $k \leq \sigma_1 \leq \sigma_2 \leq K = 1$ and thus the curve length calculation in (3.77) fulfills all the assumptions of the anisotropic Cauchy–Crofton formula in Theorem 3.8. Hence we can apply the formula to calculate the perimeter in (3.77) as

$$\text{Per}_A(U; \Omega) = \int_{\mathcal{L}} \sum_{x \in \ell_{\nu, \rho} \cap \gamma} \frac{\det M(x)}{2(\nu^T \cdot M(x) \cdot \nu)^{3/2}} d\mathcal{L}(\ell_{\nu, \rho}) ds, \quad (3.80)$$

where $\gamma = \partial U \cap \Omega$. Note that P does not affect the determinant, i.e. $\det A = \det PAP^T = \det M$, and from our decomposition in (3.10) we see that the transformation $PAP^T \rightarrow M$ actually amounts to switching the two eigenvalues σ_1 and σ_2 in Σ .

This concludes the treatment of the continuous problem. We have seen how the anisotropic coarea formula in Theorem 3.4 allows us to calculate the anisotropic total variation as an integral of the perimeter of all the level sets. Through the anisotropic Cauchy–Crofton formula in Theorem 3.8 these perimeters are calculated by an integral over the set of all lines. We are then left with the functional

$$F(u) = \int_{\Omega} (u - f)^2 + \beta \text{TV}_A(u), \quad (3.81)$$

where

$$\text{TV}_A(u) = \int_{-\infty}^{\infty} \int_{\mathcal{L}} \sum_{x \in \ell_{\nu, \rho} \cap \gamma_s} \frac{\det M(x)}{2(\nu^T \cdot M(x) \cdot \nu)^{3/2}} d\mathcal{L}(\ell_{\nu, \rho}) ds, \quad (3.82)$$

and $\gamma_s = \partial\{u > s\} \cap \Omega$. Within the restrictions that these theorems put on the tensor $M(x)$, we have chosen a construction where one eigenvalue is always 1, while the other varies from 1 in smooth areas towards 0 around edges, with the corresponding eigenvector perpendicular to the edge.

Discrete formulation

The whole transformation from the initial functional in (3.4), through the anisotropic coarea formula and the Cauchy–Crofton formula was motivated by the discrete formulation which will be described here. After discretizing the functional, we will see how a graph cut approach can be used to find a global minimizer in polynomial time.

4.1 Discretization

Assume that our discrete images are given on a uniform grid \mathcal{G} , where each grid point is called a pixel. The image is a function giving each pixel a value in the set of levels $\mathcal{P} = \{0, \dots, L-1\}$, such that $u : \mathcal{G} \rightarrow \mathcal{P}$. This is a reasonable assumption for digital grayscale images. Thus, when discretizing the functional in (3.81), we have to consider that our images now have both discrete domain and co-domain.

The integrals in (3.81) will be approximated by discrete sums. First the fidelity term is discretized without too much trouble, while with the regularization term, there is more choice as to how to discretize the set of lines \mathcal{L} . In the end we will verify that our discretization is consistent with the continuous functional.

4.1.1 Fidelity term

Since it is not affected by our introduction of the anisotropy tensor, the fidelity term can be discretized as in my project work [1]. For some pixel position $x \in \mathcal{G}$ and some level value $k \in \mathcal{P}$, we define the following function

$$N_x(k) = |k - f_x|^2 \tag{4.1}$$

which is the value of the fidelity term if we were to give u_x a value of k . This allows us write

$$\int_{\Omega} |u - f|^2 dx \approx \sum_{x \in \mathcal{G}} |u_x - f_x|^2 \Delta x = \sum_{x \in \mathcal{G}} N_x(u_x) \Delta x. \quad (4.2)$$

The reason we introduce the function $N_x(k)$ is that we want to apply the following decomposition formula, which holds for any function $F(k)$ taking values $k \in \mathcal{P}$:

$$\begin{aligned} F(k) &= \sum_{\lambda=0}^{k-1} (F(\lambda+1) - F(\lambda)) + F(0) \\ &= \sum_{\lambda=0}^{L-2} (F(\lambda+1) - F(\lambda)) I(\lambda < k) + F(0), \end{aligned} \quad (4.3)$$

where $I(x)$ is the indicator function that takes the value 1 if x is true, and 0 if x is false. Since $I(\lambda < u_x) = u_x^\lambda$ we rewrite (4.2) and obtain

$$\sum_{x \in \mathcal{G}} |u_x - f_x|^2 = \sum_{x \in \mathcal{G}} N_x(u_x) = \sum_{\lambda=0}^{L-2} \sum_{x \in \mathcal{G}} (N_x(\lambda+1) - N_x(\lambda)) u_x^\lambda + N_x(0). \quad (4.4)$$

As our domain is discretized uniformly, we drop the constant Δx , and absorb it into our parameter β of (3.81). Note that since our image takes values in $\mathcal{P} = \{0, \dots, L-1\}$, the thresholded image u^{L-1} is equal to zero everywhere.

4.1.2 Regularization term

Discretizing the regularization term is more challenging. We introduce the discrete levels to get

$$\int_{-\infty}^{\infty} \text{Per}_A(\{u > \lambda\}; \Omega) d\lambda \approx \sum_{\lambda=0}^{L-2} \text{Per}_A(\{u > \lambda\}; \Omega) \Delta\lambda. \quad (4.5)$$

As with the Δx difference, we can absorb the $\Delta\lambda$ difference into the β parameter of (3.81). The perimeter is then calculated using a discretized version of the Cauchy–Crofton formula introduced in Theorem 3.8. Again, we stop the sum at $L-2$ since the level set $\{u > L-1\}$ is empty and has zero perimeter.

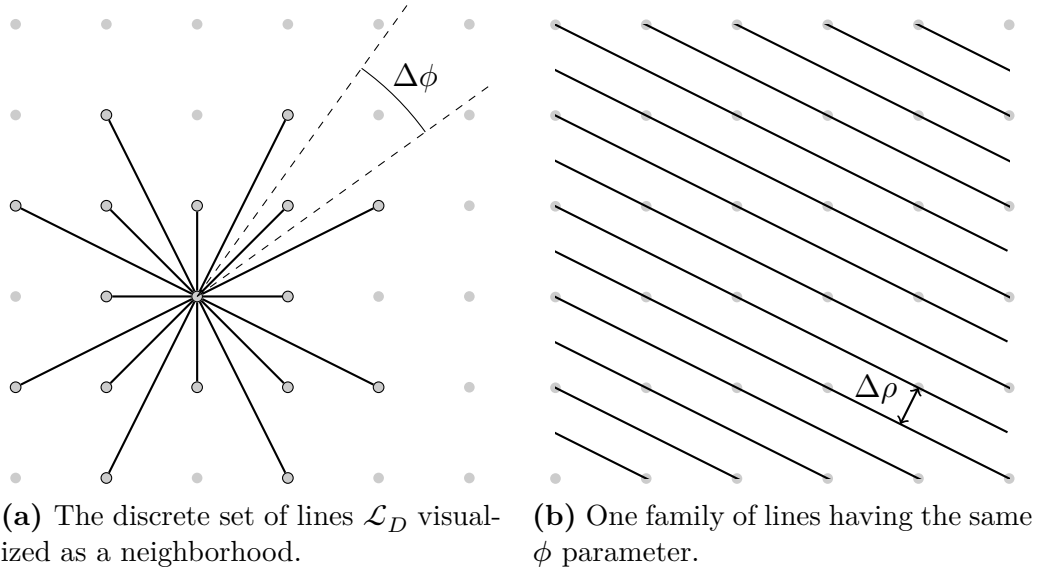


Figure 4.1: The set of lines \mathcal{L} is discretized to \mathcal{L}_D where each line belongs to a family given by ϕ , the angle parameter.

Discrete anisotropic Cauchy–Crofton formula

By approximating the integral Theorem 3.8 by a discrete sum we obtain the approximation

$$\begin{aligned}
 |C|_M &= \int_{\mathcal{L}} \sum_{x \in \ell_{\nu, \rho} \cap C} \frac{\det M(x)}{2(\nu^T \cdot M(x) \cdot \nu)^{3/2}} d\mathcal{L}(\ell_{\nu, \rho}) \\
 &\approx \sum_{\ell_{\nu, \rho} \in \mathcal{L}_D} \sum_{x \in \ell_{\nu, \rho} \cap C} \frac{\det M(x)}{2(\nu^T \cdot M(x) \cdot \nu)^{3/2}} \Delta \ell_{\nu, \rho} \\
 &= \sum_{\nu} \sum_{\rho} \sum_{x \in \ell_{\nu, \rho} \cap C} \frac{\det M(x)}{2(\nu^T \cdot M(x) \cdot \nu)^{3/2}} \Delta \rho \Delta \nu.
 \end{aligned} \tag{4.6}$$

The set of lines \mathcal{L} has been discretized into the set \mathcal{L}_D . Note that C is still a differentiable curve, not yet discretized. Being a difference in the ρ parameter of our line discretization in Figure 3.5, the difference $\Delta \rho$ represents the distance from one line to the next in a line family as shown in Figure 4.1b, and thus depends on the angle ϕ considered. The difference $\Delta \phi$ is taken to be the average of the distance to the two neighboring line families as shown in Figure 4.1a, and thus also depends on ϕ .

The choice of our discrete set of lines \mathcal{L}_D is important, as it will decide the accuracy of our approximation. We need some sensible restrictions on the set




Figure 4.2: Here our intersection approximation would not be correct, as only the intersection with edge b is counted in (4.15), even though the curve intersects edge a twice.

\mathcal{L}_D to simplify the further discussion. All lines intersect at least two grid points, and from the periodicity of our grid they thus intersect an infinite number of grid points. This puts some restrictions on the angles we can choose. For each angle included, we include all possible lines of that family, meaning there are no grid points without a line of that family intersecting it.

The set of lines can then be represented by the neighborhood of a pixel as shown in Figure 4.1a. We write $\mathcal{N}(x)$ for the neighborhood of grid point x . Extending the edges shown in Figure 4.1a gives all lines going through the point considered. Figure 4.1b shows all lines of a given family, i.e. lines having the same angle parameter ϕ .

Thus not only have we discretized the set of lines, but each line is made up of edges going from one grid point to the next. We will denote such an edge by e or e_{ab} when its endpoints are $a, b \in \mathcal{G}$. Thus we rewrite the discretization of (4.6), and sum over all the edges in the discretization \mathcal{L}_D to obtain

$$|C|_M \approx \sum_e \sum_{x \in e \cap C} \frac{\det M(x) \|e\|^3}{2(e^T \cdot M(x) \cdot e)^{3/2}} \Delta\phi \Delta\rho. \quad (4.7)$$

This is beginning to look like something we can calculate. One difficulty is finding the intersections $e \cap C$. The exact calculations of these points will not fit into our graph cut framework later, and thus for an edge e we will only consider the question of “did e cross C or not?” This amounts to checking whether the terminals of e lie on different sides of the curve C . This approximation is exact for zero or one intersection points, but will, as we see in Figure 4.2, be wrong when we have more.

The second difficulty is that in the discrete setting, we will only have an approximation of the metric tensor $M(x)$ for each point $x \in \mathcal{G}$, and it is thus not available for arbitrary intersection points in Ω . For an intersection of edge e we will utilize the average of the tensor in the two endpoints of the edge. Thus for an




Figure 4.3: A visual argument showing that $\delta^2 = \Delta\rho \|e\|$. If extended to the whole plane, there will be the same amount of blue squares as red rectangles as each grid point is the upper left corner of both a blue and a red rectangle. Thus their areas must be equal.

intersection point x somewhere on the edge e_{ab} , we approximate the metric tensor by

$$M(x) \approx M(e_{ab}) = \frac{M(a) + M(b)}{2}, \quad (4.8)$$

the component-wise average of the tensors in the two end points of the edge. Recall that we have already done some spatial smoothing of the structure tensor in (3.5) corresponding to the *integration scale* ρ , and thus we expect the tensors $M(a)$ and $M(b)$ to be similar for edges e of reasonably short length.

We also remark that using the Rayleigh principle, it is easy to conclude that the eigenvalues of the tensor approximation $M(e_{ab})$ are bounded below and above by the smallest and largest eigenvalues of $M(a)$ and $M(b)$.

We have now almost arrived at our final curve length approximation, but we need a way to calculate the inter-line distance $\Delta\rho$ which will be provided by the following lemma.

Lemma 4.1. *For each family of lines given by an angle parameter ϕ in the uniform grid of size δ we have the relation*

$$\delta^2 = \|e\| \Delta\rho. \quad (4.9)$$

Proof. Consider a line ℓ intersecting the point in the grid given by the indices $(p, q) \in \mathbb{Z}^2$. The distance $\Delta\rho$ from this line ℓ to the neighboring lines can then be calculated as a minimum over the distance to all other grid points.

The lines are split into edges $e = (\delta s, \delta t)^T$ where $s, t \in \mathbb{Z}$ are coprime such that e does not intersect any other grid points than its two endpoints.

We then calculate the minimal distance to a grid point not on the line ℓ as

$$\begin{aligned}\Delta\rho &= \min_{(p',q') \in \mathcal{G} \setminus \ell} \left\{ \left\langle \delta[p-p', q-q'], \frac{e^\perp}{\|e^\perp\|} \right\rangle \right\} \\ &= \min_{(p',q') \in \mathcal{G} \setminus \ell} \left\{ \delta^2 \cdot \frac{t(p-p') - s(q-q')}{\|e\|} \right\}.\end{aligned}\quad (4.10)$$

Since s and t are coprime, there exists $a, b \in \mathbb{Z}$ such that $at - bs = 1$, and since the $\Delta\rho$ cannot be zero, we obtain

$$\Delta\rho = \frac{\delta^2}{\|e\|}.\quad (4.11)$$

A visual argument for the same result can be seen in Figure 4.3. \square

Inserting $\Delta\rho = \delta^2/\|e\|$ and the tensor approximation of (4.8) into the curve length approximation of (4.7) we obtain

$$|C|_M \approx \sum_{e \cap C} \frac{\det M(e) \|e\|^2 \delta^2 \Delta\phi}{2 (e^T \cdot M(e) \cdot e)^{3/2}},\quad (4.12)$$

where the sum is over all edges crossing the curve.

The curve length we initially wanted to calculate was the level set perimeter $\text{Per}_A(\{u > \lambda\}; \Omega)$ in (4.5). To find edges that crosses this boundary curve, we identify the edges that have one terminal inside the level set, and the other outside. Thus we rewrite the sum over $e \cap C$ such that

$$\text{Per}_A(\{u > \lambda\}; \Omega) \approx \sum_{e_{ab}} |u_a^\lambda - u_b^\lambda| \frac{\det M(e_{ab}) \|e_{ab}\|^2 \delta^2 \Delta\phi}{2 (e_{ab}^T \cdot M(e_{ab}) \cdot e_{ab})^{3/2}}.\quad (4.13)$$

The absolute value $|u_a^\lambda - u_b^\lambda|$ is 1 if one of a and b lie inside the level set and the other lies outside, and 0 otherwise. In other words the absolute value is one if e_{ab} crosses the perimeter of $\{u > \lambda\}$ an odd number of times, and zero otherwise.

Thus we have arrived at our final discretization, which takes the form

$$F(u) = \sum_\lambda \sum_x F_x^\lambda(u_x^\lambda) + \beta \sum_\lambda \sum_{(x,y)} F_{x,y}^\lambda(u_x^\lambda, u_y^\lambda) =: F^\lambda(u^\lambda),\quad (4.14)$$

$$F_x^\lambda(u_x^\lambda) = (N_x(\lambda + 1) - N_x(\lambda)) \cdot u_x^\lambda,$$

$$F_{x,y}^\lambda(u_x^\lambda, u_y^\lambda) = |u_x^\lambda - u_y^\lambda| \frac{\det M(e_{xy}) \|e_{xy}\|^2 \delta^2 \Delta\phi}{2 (e_{xy}^T \cdot M(e_{xy}) \cdot e_{xy})^{3/2}}.\quad (4.15)$$

Recall that $N_x(\lambda) = |\lambda - f_x|^2$.

If we minimize F_λ to obtain u^λ for each level separately, it is obvious that we will also minimize the sum over all F_λ . However, it is not guaranteed that the obtained thresholded images u^λ can be combined to make an output image u . They were defined as $u^\lambda = \chi_{u > \lambda}$, so we need them to be monotonically decreasing in increasing level values, i.e.

$$u_x^\lambda \geq u_x^\mu, \quad \forall \lambda \leq \mu, \quad \forall x \in \mathcal{G}. \quad (4.16)$$

Later we will present two graph cut algorithms that find the thresholded images minimizing each level, *while guaranteeing that they meet this requirement*.

Consistency

Consistency relates to whether a solution to the continuous problem fits in the discretized equation, in other words, whether the discretized equation approximates the continuous one.

It is obvious that the discretization of the fidelity term in (4.2) is consistent. The sum is a midpoint rule approximation of the integral. As the grid is refined and $\delta \rightarrow 0$ the sum will converge to the integral.

For the regularization term we will argue that for a differentiable curve C , the discretization of our domain Ω and the set of lines \mathcal{L} gives a discrete Cauchy–Crofton formula that is consistent with the continuous one. We will show that for an increasingly refined discrete domain \mathcal{G} , there exists a choice for \mathcal{L}_D that leads to a consistent Cauchy–Crofton formula. For convenience we will use a neighborhood representation of \mathcal{L}_D similar to the one in Figure 4.1a.

If we consider the edges e of each family separately, the curve length approximation in (4.12) can be written

$$\begin{aligned} |C|_M &= \int_\nu \int_\rho \sum_{x \in \ell_{\nu,\rho} \cap C} \frac{\det M(x)}{2 (\nu^T \cdot M(x) \cdot \nu)^{3/2}} d\rho d\nu \\ &\approx \sum_\nu \sum_\rho \sum_{e_{\nu,\rho} \cap C} \frac{\det M(e_{\nu,\rho}) \|e_{\nu,\rho}\|^3}{2 (e_{\nu,\rho}^T \cdot M(e_{\nu,\rho}) \cdot e_{\nu,\rho})^{3/2}} \Delta\rho \Delta\nu. \end{aligned} \quad (4.17)$$

As described in the construction of this formula, there are four main approximations used. Firstly there is the fact that we do not consider the actual intersection points, but only whether an edge crosses the curve or not. Secondly we have the tensor which is averaged as in (4.8). And then we have the discretizations of our two line parameters ν and ρ .

It is intuitive that if $\sup \|e\| \rightarrow 0$, the number of times the differentiable curve C can cross a given edge decreases. We will not prove convergence, but rather assume that the special cases where it might not work, are negligible.




Figure 4.4: The discretization in the ρ dimension can be regarded as a midpoint rule approximation of the integral, since the difference $\Delta\rho$ the same for all lines in one line family.




Figure 4.5: We showed that the maximal angle difference $\Delta\phi_k$ goes to zero. The discretization in the ϕ dimension can be viewed as a rectangle approximation rule of the integral, as the summand is evaluated at ϕ_k , somewhere inside the interval $\Delta\phi_k$.

Further, if $\sup \|e\| \rightarrow 0$ it is obvious that the tensor average in (4.8) converges to the tensor in the intersection point.

Consider now the discretization in ρ . For each ϕ parameter, our discretization in the ρ dimension can be regarded as a midpoint rule as shown in Figure 4.4. Thus if $\sup \Delta\rho \rightarrow 0$, this part of the discretization is consistent.

The discretization in the ϕ dimension can also be regarded as a version of the *rectangle method*, although not the midpoint rule. As shown in Figure 4.5, the circle is split into intervals

$$\left[\frac{\phi_{k-1} + \phi_k}{2}, \frac{\phi_k + \phi_{k+1}}{2} \right] \quad (4.18)$$

of length $\Delta\phi_k = (\phi_{k+1} + \phi_{k-1})/2$. The summand is evaluated at ϕ_k , somewhere inside the interval. Thus if $\sup \Delta\phi_k \rightarrow 0$, this discretization is also consistent.

To show that all these properties can be fulfilled, we look at a particular neighborhood stencil construction. Consider a square centered around a grid point with




Figure 4.6: To show that we have a consistent discretization of the Cauchy–Crofton integral formula, we construct a discrete set of lines \mathcal{L}_D such that the length of the edges $\|e\|$, the angle differences $\Delta\phi$ (here a and b) and the distance between lines $\Delta\rho$ goes to zero as $\delta \rightarrow 0$.

side lengths $\sqrt{\delta}$ as shown in Figure 4.6. As δ goes to zero, the size of this square will go to zero. Inside this square we can fit a square of $n^2 = \lfloor 1/\sqrt{\delta} \rfloor^2$ grid points. This means that the number of grid points along the outer edge of this square n goes to infinity.

We include all grid points inside the square in our neighborhood, except for multiple points that lie on the same line from the origin. If two or more grid points lie on the same line, we include only the one closest to the origin. This implies that for each grid point along the outer edge of this square, we include in our neighborhood a grid point having the same angle ϕ to the x -axis.

This construction can be seen in Figure 4.6 for $n = 5$. The maximal angle between two lines $\phi_k - \phi_{k-1}$ will always be when one of ϕ_k and ϕ_{k-1} is horizontal or vertical, shown in Figure 4.6 as the angle a . Thus the largest $\Delta\phi_k$ will be when $\phi_k = m\pi/2$ for $m \in \mathbb{Z}$, so around vertical and horizontal edges. The supremum can then be calculated to be

$$\sup \Delta\phi_k = 2 \cdot \sup \frac{\phi_{k+1} - \phi_k}{2} = \arctan \frac{1/n}{n/2} = \arctan \frac{2}{n^2} \rightarrow 0. \quad (4.19)$$

Further we see that the edge length will be bounded by half of the diagonal of the square such that

$$\|e\| \leq \sqrt{\delta}/2 \rightarrow 0. \quad (4.20)$$

And finally we know from Lemma 4.1 that for each line family $\delta^2 = \Delta\rho \|e\|$ and $\|e\| \geq \delta$. Thus for the inter-line distance $\Delta\rho_k$ we have

$$\sup \Delta\rho_k = \sup \frac{\delta^2}{\|e\|} \leq \frac{\delta^2}{\delta} = \delta \rightarrow 0. \quad (4.21)$$

Hence the approximation has been shown to be equivalent to well-known, and consistent integral approximations, where the summand converges to the integrand, and the differences $\Delta\phi$ and $\Delta\rho$ go to zero. Thus the perimeter approximation in (4.5) is consistent with the continuous formulation in Theorem 3.8.

Note that as we will work with digital images with fixed resolutions, we do not really have the chance to refine our discretization. We do however have to take these things into account when creating our neighborhood stencil, to make sure that we get a reasonable approximation of the perimeter lengths.

4.2 Graph cut approach

The discretization we arrived at in (4.15) can be minimized using graph cuts. For each level λ , a minimum graph cut is found to produce the corresponding level set $\{u > \lambda\}$. These are then combined to form the final restored image u .

In this section we will look at how these graphs are constructed such that their minimum cuts correspond to the minimizers of the functional F^λ . The description is taken with some small adjustments from my project work [1], and is included here for completeness. An implementation of the described approach can be found in Appendix A.

4.2.1 Graphs

Using the notation of [18] we will denote a directed graph as $G = (V, E)$ where V is a finite set of vertices, and E is a binary relation on V . If $(u, v) \in E$ we say that there is an edge from u to v in the graph G .

We introduce the non-negative capacity function $c : V \times V \rightarrow [0, \infty)$. Only edges $(u, v) \in E$ can have a positive capacity $c(u, v) = q > 0$ and it means that it is possible to send a *flow* of maximum q units from u to v . For convenience we will let $c(u, v) = 0$ for any pair $(u, v) \notin E$, and we do not allow self-loops in our graph. When a directed graph G is equipped with capacity function c , one might call it a capacitated graph or a flow network, but as all our graphs will be capacitated from this point, we will just call them graphs and we write $G = (V, E, c)$.

There are two special vertices in the graph, the source s and the sink t . Contrary to other vertices, which can neither produce nor receive excess flow, the source can produce and the sink can receive an unlimited amount of flow. The most basic problem in graph flow theory is the question of how much flow it is possible to send through the graph from the source to the sink.

What we seek in our final graph is a minimum s - t -cut, a “minimal” line through the graph that cuts a set of edges and divides the vertex set in two, separating the source from the sink.

Definition 4.2 (*s-t-cut*). Given a graph $G = (V, E, c)$, an *s-t-cut* (S, T) of G is a partition of V into S and $T = V - S$ such that $s \in S$ and $t \in T$. The capacity of the cut is

$$c(S, T) = \sum_{u \in S} \sum_{v \in T} c(u, v), \quad (4.22)$$

and a minimum *s-t-cut* is a cut whose capacity is minimum over all *s-t-cuts*.

Note that there might exist several minimum *s-t-cuts* in a graph, resulting in different partitions of V . This is why we need to verify later that the cuts we obtain result in stackable thresholded images.

4.2.2 Graph representable functionals

The next step is to find a way to construct a graph such that we can minimize the functional in (4.15) by finding a minimum *s-t-cut*. We will do this by creating small and simple graphs representing the separate summands of the functional. For these small graphs it will be easy to verify that the minimal cut also minimizes the corresponding part of the functional, and they can then be merged giving a graph representing the complete functional in (4.15).

First we need to establish the definition of a graph representable function, presented by Kolmogorov and Zabih in [19].

Definition 4.3 (Graph representable functions). A function $\mathcal{E}(x_1, \dots, x_n)$ of n binary variables is graph-representable if there exists a graph $G = (V, E, c)$ with terminals s and t , and a subset of vertices $V_0 = \{v_1, \dots, v_n\} \subseteq V - \{s, t\}$ such that, for any configuration $(x_1, \dots, x_n) \in \{0, 1\}^n$, the value of the energy $\mathcal{E}(x_1, \dots, x_n)$ is equal to a constant plus the cost of the minimum *s-t-cut* among all cuts $C = (S, T)$ where $x_i = 0 \Leftrightarrow v_i \in S$ and $x_i = 1 \Leftrightarrow v_i \in T$, $\forall 1 \leq i \leq n$.

From this definition we see that if we have a graph-representable function \mathcal{E} it is possible to find an exact global minimum of \mathcal{E} by finding a minimal *s-t-cut* in a graph representing \mathcal{E} .

Furthermore Kolmogorov and Zabih present an important result concerning what kinds of functions are graph-representable.

Theorem 4.4 (Identification of graph representable functions). Given an energy function \mathcal{E} of n binary variables of the form

$$\mathcal{E}(x_1, \dots, x_n) = \sum_i \mathcal{E}^i(x_i) + \sum_{i < j} \mathcal{E}^{i,j}(x_i, x_j), \quad (4.23)$$

then \mathcal{E} is graph representable if and only if each term $\mathcal{E}^{i,j}$ satisfies the inequality

$$\mathcal{E}^{i,j}(0, 0) + \mathcal{E}^{i,j}(1, 1) \leq \mathcal{E}^{i,j}(0, 1) + \mathcal{E}^{i,j}(1, 0). \quad (4.24)$$

This theorem will allow us to verify that our functional actually is graph representable.

Finally, the following theorem, proved by Kolmogorov and Zabih in [19], will be crucial in our graph construction.

Theorem 4.5 (Additivity). *The sum of a finite number of graph-representable functions*

$$\mathcal{E}(x_1, \dots, x_n) = \sum_k \mathcal{E}^k(x_1, \dots, x_n), \quad (4.25)$$

each represented by a graph $G^k = (V^k, E^k, c^k)$, is graph-representable by $G = (V, E, c)$ where $V = \cup_k V^k$, $E = \cup_k E^k$ and $c(u, v) = \sum_k c^k(u, v)$.

It allows us to construct small graphs representing the different summands of our functional (4.15), before adding them together to create a final graph representing complete functional.

Note that when we apply this theorem later, we will assume that all the summands of (4.15) have the whole picture as their domain. It is unproblematic to extend $F_\lambda^x(u_x^\lambda)$ and $F^{x,y}(u_x^\lambda, u_y^\lambda)$ such that they take all the pixels their argument and then ignore all pixels except the ones they actually depend on.



(a) The graph when $F_\lambda^x(1) > 0$, with constant equal to 0. (b) Graph when $F_\lambda^x(1) < 0$, with constant equal to $-F_\lambda^x(1)$.

Figure 4.7: The graph construction for the fidelity term $F_\lambda^x(u_x^\lambda)$. See Table 5.1 for an overview of the different possible cuts, and on why this construction works.

4.2.3 Graph construction

We will construct a graph in such a way that if a vertex u_x^λ ends up on the source side of the cut we set $u_x^\lambda = 1$, and if it ends up on the sink side we set $u_x^\lambda = 0$, as

Table 4.1: Each row represents one of the two possible values of $u_x^\lambda \in \{0, 1\}$. The functional $F_\lambda^x(u_x^\lambda)$ and minimum cut obtaining this configuration is shown. The last two columns show the capacities of the cut for each of the two graph constructions in Figure 4.7. We verify that for each of the two graph constructions, the cut capacities are equal to the functional value, plus a constant.

u_x^λ	$F_\lambda^x(u_x^\lambda)$	Min. cut (S, T)	Graph (a) cut cap.	Graph (b) cut cap.
0	0	$(\{s\}, \{u_x^\lambda, t\})$	0	$-F_\lambda^x(1)$
1	$F_\lambda^x(1)$	$(\{s, u_x^\lambda\}, \{t\})$	$F_\lambda^x(1)$	0

in Definition 4.3. This is an arbitrary choice, but still something we have to keep in mind through the rest of the section.

We will now consider the two kinds of summands in the discrete functional (4.15). The fidelity term coming from our aim to keep the output image close to the original image, and the regularization term coming from our aim to minimize the total variation.

Fidelity term

The fidelity term of our discrete functional (4.15) simplifies to

$$F_\lambda^x(0) = 0 \quad (4.26)$$

$$F_\lambda^x(1) = N_x(\lambda + 1) - N_x(\lambda) \quad (4.27)$$

where $F_\lambda^x(1)$ might be positive or negative depending on λ and the pixel value f_x .

Figure 4.7 shows how graphs can be constructed to represent this part of functional. The construction differs depending on whether $F_\lambda^x(1)$ is positive or negative. Table 5.1 shows how the cuts correspond to the values of u_x^λ and we can easily verify that the constructed graph actually represents the fidelity term.

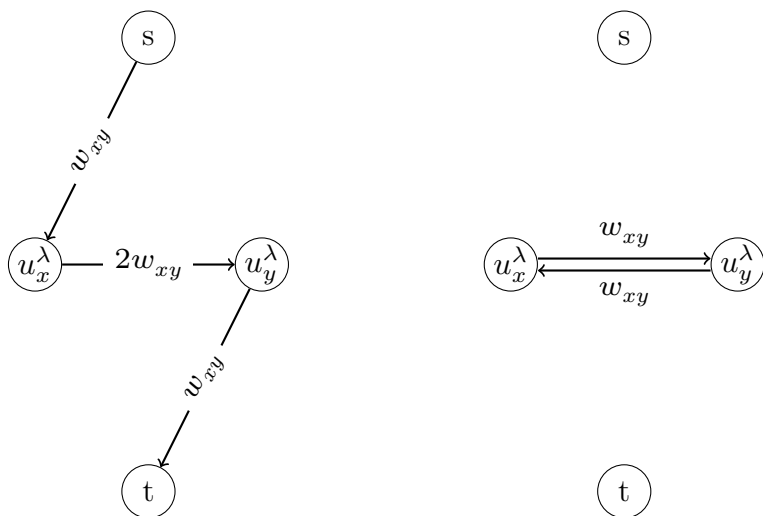
Regularization term

For our regularization term in (4.15) on the form

$$F^{x,y}(u_x^\lambda, u_y^\lambda) = w_{xy} |u_x^\lambda - u_y^\lambda| \quad (4.28)$$

we have

$$\begin{aligned} F^{x,y}(0, 0) &= 0, \\ F^{x,y}(0, 1) &= w_{xy}, \\ F^{x,y}(1, 0) &= w_{xy}, \\ F^{x,y}(1, 1) &= 0, \end{aligned} \quad (4.29)$$



(a) Representing $F^{x,y}(u_x^\lambda, u_y^\lambda)$ with a constant term of w_{xy} . (b) Representing $F^{x,y}(u_x^\lambda, u_y^\lambda)$ with a constant term of 0.

Figure 4.8: Two alternative ways of constructing a graph representing the energy term $F^{x,y}(u_x^\lambda, u_y^\lambda)$. See Table 4.2 for an overview of the different possible configurations of $(u_x^\lambda, u_y^\lambda)$ and how they correspond to cuts through the graph.

and by Theorem 4.4 our functional is graph representable. In [19], Kolmogorov and Zabih presents a way to construct a graph for any graph representable function on the form shown in Theorem 4.4. Since the energies in (4.29) are especially simple, the construction and presentation is simplified.

Figure 4.8 shows two different ways of how a graph can be constructed to represent the regularization term. See Table 4.2 for an overview of how the values of $(u_x^\lambda, u_y^\lambda)$ corresponds to cuts in the graph.

Figure 4.9 shows a visualization of how the final graph might look with all its edges. The source will have a lot of outgoing edges, one for each pixel, while the sink has one incoming edge from each pixel. The vertices corresponding to the pixels are only connected to the source, the sink, and their neighboring pixels, so their edge degree is much smaller than for s and t .

Table 4.2: An overview of the possible configurations of the variables in the term $F^{x,y}(u_x^\lambda, u_y^\lambda)$. For each configuration the corresponding functional value and the cut yielding this configuration is shown. The last two columns show the capacities of the cut in the two alternative graph constructions shown in Figure 4.8. We verify that for each of the two graph constructions, the cut capacities are equal to the functional value, plus a constant.

$(u_x^\lambda, u_y^\lambda)$	$F^{x,y}(u_x^\lambda, u_y^\lambda)$	Min. cut (S, T)	Alt. (a) cut cap.	Alt. (b) cut cap.
(0, 0)	0	$(\{s\}, \{u_x^\lambda, u_y^\lambda, t\})$	w_{xy}	0
(0, 1)	w_{xy}	$(\{s, u_y^\lambda\}, \{u_x^\lambda, t\})$	$2w_{xy}$	w_{xy}
(1, 0)	w_{xy}	$(\{s, u_x^\lambda\}, \{u_y^\lambda, t\})$	$2w_{xy}$	w_{xy}
(1, 1)	0	$(\{s, u_x^\lambda, u_y^\lambda\}, \{t\})$	w_{xy}	0




Figure 4.9: A visualization of the final graph, with edges between neighboring pixels, and edges connecting the source and sink to the rest of the graph. The edge capacities are not visualized, and of course some of them might be zero while others might be very high.

Maximum flow

In the previous section we have seen how finding the minimum cut of carefully constructed graph can give us the thresholded image minimizing the functional for one level value λ . We will in this and the next section see how such a minimum cut can be found by sending flow through the graph and trying to identify the “bottleneck”. This chapter, except for the description of the Boykov–Kolmogorov algorithm is taken with some adjustments from my project work [1] and is included here for completeness. Implementations of the push-relabel and Boykov–Kolmogorov algorithms can be found in Appendix A.

5.1 Flow graphs

We have already introduced capacities, and briefly mentioned the notion of flow as something limited by the capacity. In other words, flow is something we can send through our graph, but the capacity limits how much we can send along each edge. It is useful to imagine a water supply graph with pipes of different sizes and a water source and sink.

Formally, we introduce the *flow* as the function $f : V \times V \rightarrow [0, \infty)$. This function keeps count of how much flow we are sending through each edge of our graph and must satisfy the following two constraints

Capacity constraint: For all $u, v \in V$, $0 \leq f(u, v) \leq c(u, v)$, i.e., for every pair of vertices, the flow is less than or equal to the capacity.

Flow conservation: For all $u \in V - \{s, t\}$

$$\sum_{v \in V} f(v, u) = \sum_{v \in V} f(u, v), \quad (5.1)$$

i.e., for any vertex except the source and the sink, the flow into the vertex must be equal to the flow out of the vertex.

Note that we have defined f with all pairs of vertices as its domain, even though it is only non-zero on edges $(u, v) \in E$. This makes it easier to write sums as in the flow conservation constraint. We say that an edge (u, v) is *saturated* if $f(u, v) = c(u, v)$.

We define $|f|$ as the net amount of flow from the source to the sink in the graph. Because of the flow conservation constraint, this can be calculated as the net amount of flow going out of the source

$$|f| = \sum_{v \in V} f(s, v) - \sum_{u \in V} f(u, s). \quad (5.2)$$

Furthermore we denote the *net flow* across an s - t -cut $C = (S, T)$ as

$$f(S, T) = \sum_{u \in S} \sum_{v \in T} f(u, v) - \sum_{v \in T} \sum_{u \in S} f(v, u). \quad (5.3)$$

Note how this definition differs from the capacity of a cut $c(S, T)$ in Definition 4.2. While the capacity of a cut represents how much flow it is maximally possible to send from S to T , the net flow across a cut represents the net amount of flow going across the cut, counting negatively the flow that goes back from T to S .

For any s - t -cut we have that

$$|f| = f(S, T). \quad (5.4)$$

This is quite intuitive given the flow conservation constraint, and a full chain of arguments can be found in [18].

Following from (5.3) and (5.4), we find that

$$\begin{aligned} |f| &= \sum_{u \in S} \sum_{v \in T} f(u, v) - \sum_{v \in T} \sum_{u \in S} f(v, u) \\ &\leq \sum_{u \in S} \sum_{v \in T} f(u, v) \\ &\leq \sum_{u \in S} \sum_{v \in T} c(u, v) \\ &\leq c(S, T) \end{aligned} \quad (5.5)$$

for any s - t -cut $C = (S, T)$. A very central result in graph flow theory called the max-flow min-cut theorem will be presented later. It states that the inequality of (5.5) becomes an equality when f is a maximum flow for some cut $C = (S, T)$, and that all such cuts are minimum cuts.

But how does this help us? If we know the maximum flow value, and we have an s - t -cut with capacity equal to the maximum flow, we actually have a minimum cut. The question is then, how do we find a maximum flow, and how do we find a minimum cut?

5.2 Augmenting path algorithms

The family of augmenting flow algorithms represent a popular approach to the maximum flow problem. The idea is simply to look for paths from the source to the sink that can carry additional flow, so-called augmenting paths, and then send the maximum possible amount of flow along such a path. When no such path exists anymore, no more flow can be sent from the source to the sink, and a maximum flow has been reached.

5.2.1 Residual graph

When further discussing approaches to solving the maximum flow problem we will need the notion of a residual graph $G_f = (V_f, E_f, c_f)$, which is derived from the original graph G and contains the edges along which it is possible to send additional flow. This means that E_f contains the edges (u, v) from E where $f(u, v) < c(u, v)$. But that is not all; an important realization is that it is also possible to push flow *back* along an edge where the flow is already positive. In other words, sending flow from v to u by canceling some or all of the flow that is already going from u to v .

Thus the capacity function c_f of our residual graph becomes

$$c_f(u, v) = \begin{cases} c(u, v) - f(u, v) & \text{if } (u, v) \in E, \\ f(v, u) & \text{if } (v, u) \in E, \\ 0 & \text{otherwise.} \end{cases} \quad (5.6)$$

The vertices V_f of G_f are the same as the original graph G , while the edges E_f are taken to be all pairs of vertices (u, v) with $c_f(u, v) > 0$. Since we can in E_f at most have all the original edges, and their reversals, we have $|E_f| \leq 2|E|$.

Note that there is ambiguity in the definition of the residual graph in the case where the original graph contains anti-parallel edges. One could avoid this by defining $c_f(u, v) = f(v, u) + c(u, v) - f(u, v)$ instead, or as they do in [18], disallow anti-parallel edges altogether. In any case it is not something we will have to think about in the implementation, since we will not actually construct the residual graph.

With the residual graph defined, we are ready to formally present the max-flow min-cut theorem.

Theorem 5.1 (Max-flow min-cut theorem). *If f is a flow in a graph $G = (V, E, c)$ with source s and sink t , then the following statements are equivalent:*

1. f is a maximum flow in G .
2. The residual graph G_f contains no augmenting paths.

3. $|f| = c(S, T)$ for some cut (S, T) of G .

See [18] for a proof, and remark that because of the inequality in (5.5), the cut in Statement 3 is a minimum cut. The theorem does not tell us how to find such a cut, and there are multiple ways. One possibility is to take S to be all vertices reachable from the source in the residual graph and $T = V - S$.

Figure 5.1 shows a simple graph which already has five units flowing from s to t . The marked path is a possible augmenting path, and note that it follows an edge in E in the reverse direction, made possible by the construction of the residual graph just described.

5.2.2 Ford-Fulkerson

The Ford-Fulkerson algorithm is the most basic augmenting path algorithm, which can be extended to more advanced algorithms. It is stated in pseudocode in Algorithm 1, and the idea is to augment the flow along paths from s to t until it is no longer possible.

There are different ways to find augmenting paths, and a common choice is to do a breadth-first search from the source until the sink is found, as this will yield the shortest possible augmenting path. This version of the algorithm is called Edmonds-Karp and has a running time of $O(|V||E|^2)$. See for example [18] for a description of the breadth-first search, and a formal proof of the running time of the Edmonds-Karp algorithm.

5.3 Other algorithms

There are many different maximum flow algorithms that fall into the augmenting path category, although we will see a different approach in the next section.




Figure 5.1: A graph with the flow and capacity of each edge shown as flow/capacity. The marked path is a valid augmenting path from s to t , and by increasing the flow along it, we will push flow back from v to u .

Algorithm 1 The Ford-Fulkerson max-flow algorithm

```

function FORD-FULKERSON( $G, s, t$ )
  while there exists a path  $p$  from  $s$  to  $t$  in the residual graph  $G_f$  do
     $\alpha \leftarrow \min\{c_f(u, v) : (u, v) \in p\}$ 
    for all  $(u, v) \in p$  do
      if  $(u, v) \in E$  then
         $f(u, v) += \Delta f$ 
      else
         $f(v, u) -= \Delta f$  ▷ Push flow back
      end if
    end for
  end while
end function

```

The algorithm of Dinitz, originally published in 1970, later improved on, and described by the original author in [20], is a variant of the augmenting path algorithm. It maintains a distance labeling $d(u)$ of the vertices $u \in V$ in the graph, where $d(u)$ is the shortest path from the s to u in the residual graph. This can be computed with a simple breadth-first search. The next step is to construct a *blocking flow* f' , using only edges in $E'_f = \{(u, v) \in E_f : d(u) + 1 = d(v)\}$. The blocking flow is such that if we augment the flow f by f' , there is no longer any paths from s to t following edges in E'_f . After the blocking flow f' has been found and added to f , the distance labels are recalculated, and the label of the sink will be increased by at least one.

Boykov and Kolmogorov present a variant of the augmenting path algorithm in [21], specialized for the kinds of graphs occurring in graphical applications. It will be presented in Section 5.5

5.4 Push–relabel algorithm

The push-relabel algorithm is a different approach to the maximum flow problem, presented by Goldberg and Tarjan in [22]. Unlike the augmenting flow algorithms, it does not maintain a valid flow f in the graph at all times, but still obtains a valid maximum flow when the algorithm terminates.

5.4.1 Preflow

Instead of maintaining a valid flow, we introduce the concept of a *preflow* by relaxing the flow conservation constraint from earlier. We allow positive excess in the vertices and the flow conservation constraint from before then becomes

Preflow conservation: For all $u \in V - \{s, t\}$

$$\sum_{v \in V} f(v, u) \geq \sum_{v \in V} f(u, v), \quad (5.7)$$

i.e., for any vertex except the source and the sink, the flow into the vertex must be greater or equal to the flow out of the vertex.

As in most of the cited push-relabel literature, we define $N = |V|$, and for all vertices $u \in V$ we define the excess

$$e(u) = \sum_{v \in V} f(v, u) - \sum_{v \in V} f(u, v), \quad (5.8)$$

which represents the amount of flow which disappears in vertex u . Equivalent to the preflow conservation constraint is stating that $e(u) \geq 0$ for all vertices $u \in V - \{s, t\}$.

In addition to the flow, we maintain a height map $d : V \rightarrow \mathbb{N}$ that satisfies $d(s) = 0$, and for every edge (u, v) in the residual graph, i.e. every edge with $c_f(u, v) > 0$, we require that $d(u) \leq d(v) + 1$. For all vertices u , the label $d(u)$ will be a lower bound on the length from u to t in G_f which is why it is also often called a distance labeling.

A vertex u is *active* if $u \in V - \{s, t\}$, it has positive excess ($e(u) > 0$) and $d(u) < N$. These are the vertices we want to operate on to increase the preflow.

5.4.2 Basic operations

The algorithm performs two basic operations, the *push* and *relabel* operations, while always maintaining a valid preflow f and a valid distance labeling d .

The push procedure

The push procedure moves excess flow from an active vertex along an edge $(u, v) \in E_f$ for which $d(u) = d(v) + 1$, i.e. to a vertex with a smaller distance label. We call such edges *admissible*. See Algorithm 2 for a pseudocode implementation of the push operation.

Assuming that f is a valid preflow, it is easy to verify that the preflow f and labeling d remain valid after running the push procedure on some admissible edge (u, v) .

The relabel procedure

The relabel procedure is our tool for changing the distance labels of the vertices. It changes the label of a vertex to the greatest possible value, which is one more

Algorithm 2 The push procedure of the Push-Relabel algorithm.

```

function PUSH( $u, v$ )
   $\Delta f \leftarrow \min(c_f(u, v), e(u))$ 
  if  $(u, v) \in E$  then
     $f(u, v) += \Delta f$ 
  else
     $f(v, u) -= \Delta f$  ▷ Push flow back
  end if ▷ Excess  $e(u)$  and  $e(v)$  will also change
end function

```

than the lowest label among its neighbors in the residual graph. See Algorithm 3 for a pseudocode implementation.

Algorithm 3 The relabel procedure of the Push-Relabel algorithm

```

function RELABEL( $u$ )
  if there is a  $v \in V$  such that  $(u, v) \in E_f$  then
     $d(u) \leftarrow \min\{d(v), \forall v \in V : (u, v) \in E_f\} + 1$ 
  else
     $d(u) \leftarrow N$  ▷  $u$  becomes inactive
  end if
end function

```

If d was a valid labeling before running the relabel procedure, then we still have $d(u) \leq d(v) + 1$ for all neighbors v of u in the residual graph, and d remains a valid labeling. The capacity constraint and preflow constraint remain satisfied assuming they were satisfied before the procedure was started.

5.4.3 Final algorithm

These basic procedures are then applied to active vertices and admissible edges until we obtain our minimum cut. We will see later that when there are no more active vertices, we can extract the minimum cut from the graph.

In the first phase of the algorithm we initialize a valid preflow and distance labeling by saturating all edges out of the source s , and then setting its distance label $d(s) = N$. We then apply the push and relabel procedures where applicable until there are no more active vertices and a maximum preflow is obtained.

A vertex u can only be successfully relabeled to obtain a new label if the outgoing edges of u in the residual graph have changed since the previous relabeling. This is why the push and relabel procedures often are combined into a *discharge* procedure as shown in Algorithm 4. When it is run on an active vertex u , we

push as much as possible of the excess flow to other vertices before the vertex is relabeled.

Algorithm 4 The discharge procedure of the Push-Relabel algorithm

```

function DISCHARGE( $u$ )
  for all  $v \in V$  such that  $(u, v) \in E_f$  do
    if  $c_f(u, v) > 0$  and  $d(u) = d(v) + 1$  then PUSH( $u, v$ )
    end if
  end for
  if  $e(u) > 0$  then RELABEL( $u$ )
  end if
end function

```

In the second phase of the algorithm this maximum preflow is turned into a maximum flow by returning excess flow which did not reach the sink from inside the graph back to the source. We can skip this part of the algorithm, as it is possible to identify a minimum cut as soon as the first phase is finished, and the following theorem allows us to do that.

Theorem 5.2 (Cut identification). *Given a graph $G = (V, E, c)$, assume that the first phase of the push-relabel algorithm has terminated so that no more active vertices remain. Then there exists a $k \in \{1, \dots, N - 1\}$ such that there is no vertex with label k . For every such k the vertex sets $S = \{u \in V : d(u) > k\}$ and $T = \{u \in V : d(u) < k\}$ define a minimum cut $C = (S, T)$ in G .*

Proof. There are N vertices, the source has label N and the sink has label 0 , and the $N - 2$ remaining vertices can not occupy all the $N - 1$ labels in $\{1, \dots, N - 1\}$, so there must exist an k as described.

There can be no edge $(u, v) \in E_f$ such that $u \in S$ and $v \in T$, as this would imply $k \leq d(u) - 1 \leq d(v) < k$. From the construction of E_f we now know that all edges in E from S to T are saturated, and all edges from T to S carry no flow. This implies the capacity of the cut is equal to the flow through the cut, i.e. $c(S, T) = f(S, T)$.

Since the first phase of the algorithm has terminated, there can be no active vertices, and therefore no excess in T , except for the sink. If all flow excess in vertices in S is returned to the source, we can apply the max-flow min-cut theorem to conclude that $C = (S, T)$ is a minimum s - t -cut, since the cut capacity is equal to the flow. \square

We will see later that with the gap relabeling heuristic, there will always be a gap at label $k = N - 1$ such that we can construct our cut by taking $S = \{u \in V : d(u) \geq N\}$.

Note that the vertices in S are vertices earlier described as being on the source side of the cut, as no additional flow can go from these vertices to the sink.

5.4.4 Complexity

In their original article [22], Goldberg and Tarjan analyze the complexity of the push-relabel algorithm by considering the maximum number of basic operations we can possibly do before the algorithm terminates.

The number of relabelings is in $O(|V|^2)$ since every time the procedure is applicable to a vertex u , the label $d(u)$ increases by at least one.

The number of saturating pushes is in $O(|V| |E|)$. When a push along (u, v) is saturating, the label of v has to increase with at least 2 before a push can saturate the same edge (in the opposite direction). Since the number of relabelings of a single vertex is bounded by $|V|$, and we have $|E|$ edges, this gives the stated number of saturating pushes.

The number of non-saturating pushes is the most complicated to bound, and will also make up the asymptotic running time of the algorithm. The idea is to define

$$\phi = \sum_{u \text{ active}} d(u), \quad (5.9)$$

and look at how much this number changes throughout the algorithm. It starts at zero and ends at zero. Every non-saturating push from u to v makes ϕ decrease by at least one since it makes u inactive (but might activate v). The total increase in ϕ due to relabelings is less than $|V|^2$. A saturating push from u to v increases ϕ by at most $|V|$, since v might become active.

Even if ϕ is always increased by relabelings and saturating pushes, we can bound the number of non-saturating pushes by

$$|V|^2 + |V| \underbrace{c |V| |E|}_{\#(\text{saturating pushes})} \quad (5.10)$$

which means that in the general case, the algorithm has a complexity of $O(|E| |V|^2)$.

5.4.5 Vertex selection rules

Until now we have stated that the discharge procedure is run on active vertices until there are no more active vertices left. The choice of the order in which to discharge these active vertices remain, and multiple possibilities exist.

The «First In, First Out» (FIFO) approach is to always maintain a queue of active vertices. When the vertex from the beginning of the queue is discharged, other vertices might become active, and these are added at the end of the queue.

The original article of Goldberg and Tarjan [22] contains a proof that this selection rule gives a complexity of $O(|V|^3)$, which can be very good if you have a dense graph.

Using the highest level selection rule one always discharges the active vertex with the largest distance label. Goldberg and Tarjan state that this rule also gives a complexity of $O(|V|^3)$ while this bound is improved to $O(|V|^2 \sqrt{|E|})$ in an article by Cheriyan and Maheshwari [23].

These are complexity bounds, and the actual running time of the algorithm, which can only be determined by implementing it and running it, varies a lot with the structure of the input graph.

Cherkassky and Goldberg describe the algorithm along with different selection rules, heuristics and their implementation in [24].

5.4.6 Heuristics

Different heuristics exist that can speed up the algorithm considerably. Being heuristics, they are not guaranteed to work, and might perform differently on different kinds of graphs. The most used heuristics are the gap and global relabeling heuristics, both aiming to reduce the total number of relabeling steps.

The gap relabeling heuristic aims to find a label k as in Theorem 5.2 such that no vertex has that label. From vertices u with $d(u) > k$, there are no unsaturated edges going to vertices with smaller distance labels, so no more flow can ever find its way from these vertices to the sink. These vertices are therefore given the label N and never considered again as they will never become active. Algorithm 5 shows a pseudocode representation of what is done once a gap k is found.

Algorithm 5 The gap procedure of the Push-Relabel algorithm

```

function GAP( $k$ )
  for all  $u \in V$  such that  $d(u) > k$  do
     $d(u) \leftarrow N$ 
  end for
end function

```

The preflow and capacity constraints are still valid after the gap relabeling procedure, as only the labels d are changed. For the labeling d one has to verify that $d(u) \leq d(v) + 1$ for every edge $(u, v) \in E_f$ in the residual network. If none, both, or only v is relabeled, this is trivial. It is not possible that only u would be relabeled, as this would imply that $d(u) \geq d(v) + 2$ which is not a valid initial labeling.

When running the push-relabel algorithm with the gap heuristic, we can be sure that there will never be a vertex u with label $d(u) = N - 1$ at the end of




Figure 5.2: A sketch of how the fidelity energy term $F_\lambda^x(1)$ in (4.15) increases monotonically with λ .

the algorithm, i.e. we know that there will always be a gap at label $N - 1$. This can be seen using the same reasoning as in Theorem 5.2, because if there was a vertex with label $N - 1$, there would only be $N - 3$ vertices possibly having labels in $\{1, \dots, N - 2\}$, so a gap must exist somewhere in that interval. When using the gap relabeling heuristic, such a gap can not exist, so we can conclude that there is no vertex with label $N - 1$.

Using Theorem 5.2 we can then conclude that the sets $S = \{u \in V : d(u) \geq N\}$ and $T = V - S$ form a minimum cut of the graph.

5.4.7 Parametric push-relabel algorithm

Now we have an algorithm for finding a minimum s - t -cut in a graph, so let's return to the graph constructed in Section 4.2.3. For every level $\lambda \in \{0, \dots, L - 1\}$ we want to find a minimum s - t -cut which gives us the thresholded image u^λ . These can then hopefully be stacked together to form the final image u .

Graph reuse

Solving L separate minimum cut problems seems like a lot of work, but when using the push-relabel algorithm we will, if we do things in the right order, be able to reuse the graph when going from one label to the next.

Going back to the graph representations in Figure 4.7 and Figure 4.8 we know that only the capacity of edges from sub-graphs representing the fidelity term depend on our level parameter λ . From the expression in (4.26), visualized in Figure 5.2, we see that the energy term $F_\lambda^x(1)$ increases monotonically with increasing λ parameter. Let $u, v \in V - \{s, t\}$. Since the edges in Figure 4.7 now are the only ones depending on λ , the following is true for *decreasing* values of λ

Edges from s to u : As seen in Figure 4.7b the capacity of these edges will increase monotonically with decreasing λ parameter.

Edges from u to v : These edges do not depend on λ and will remain unchanged.

Edges from v to t : As seen in Figure 4.7a the capacity of these edges will decrease monotonically with decreasing λ parameter.

After running the push-relabel algorithm for $\lambda = k$, we are left with a graph $G = (V, E, c)$, a preflow f and a labeling d . To obtain the graph for $\lambda = k - 1$ we have to change the capacity of two different kinds of edges, and this is done in the following way to keep the capacity, preflow and labeling constraints satisfied.

Edges from s to u : The capacity $c(s, u)$ is increased, and the flow is set to be equal to the capacity $f(s, u) = c(s, u)$. The vertex u might have an increased excess $e(u)$, which might in turn make it active.

Edges from v to t : The capacity $c(v, t)$ is decreased. If it is decreased to a value below the current flow value, we set $f(v, t) = c(v, t)$ which will decrease the excess of the sink t , and increase the excess of v .

None of these actions will create new edges in the residual graph, and we do not change the labeling d , so the labeling constraints are also satisfied in the new graph.

Through this procedure we have easily created the graph for $\lambda = k - 1$, and the distance labels remain the same. As the distance labels always increase monotonically, we have a head start compared to the case where we reset the flow and labels.

Output image construction

We mentioned already in Chapter 4 that in order to be able to construct our output image u , the thresholded images u^λ would have to stack one on top of the other. Because of the reuse of the distance labels between the iterations of the push-relabel algorithm, we can guarantee that this is possible.

Consider two subsequent runs of the push-relabel algorithm, for labels λ and $\lambda - 1$ ending with distance labels d^λ and $d^{\lambda-1}$ respectively. We already know that the distance labels d are monotonically increasing. This means that the set $S = \{u \in V : d(u) \geq N\}$ is increasing in size, more precisely, we have the inclusion

$$\{u \in V : d^\lambda(u) \geq N\} \subseteq \{u \in V : d^{\lambda-1}(u) \geq N\}. \quad (5.11)$$

For a pixel $x \in S$ we will set $u_x^\lambda = 1$, which together with the previous inclusion property implies

$$u_x^\lambda \leq u_x^{\lambda-1} \quad (5.12)$$

for all $x \in \mathcal{G}$. Being equivalent with the inequality in (4.16), this means our algorithm produces stackable thresholded images u^λ .

We then construct our output image u by giving each pixel the value

$$u_x = \min\{\lambda \in \{0, \dots, L-1\} : u_x^\lambda = 0\}. \quad (5.13)$$

This marks the end of the description of the implemented algorithm, but we will further discuss some possible improvements.

5.4.8 Divide and conquer

The possibility of re-using the graph between separate level is a very nice property of the push-relabel algorithm, but there are further room for improvements. Consider one pixel x with value u_x , and imagine we only wanted to find the value of this pixel. One could go through all pixel values $\lambda \in (L-1, \dots, 0)$, and see when u_x^λ changes from 1 to 0, just as we do for all the pixels in the algorithm above. Ignoring graph re-use this would have us solve $O(L)$ maximum flow problems.

Improving on this we could employ the idea of binary search to find the value of u_x in only $O(\log_2 L)$ time. After finding one cut, we know whether u_x is above or below the current λ value, and by choosing λ as the midpoint of the current possible range of u_x , we can cut the search space in half for each iteration of the algorithm.

We can extend this idea to the problem of finding all pixel values. Instead of running the algorithm for successively decreasing values of λ , we choose some λ in the middle of the range $\{0, \dots, L-1\}$. The cut we obtain consists of two sets $S = \{u \in V : d(u) \geq N\}$ and $T = V - S$. We know that no more flow can be sent from S to T , even if we decrease the value of λ and adjust the capacities accordingly.

The idea is now that we have halved the possible λ interval for *all* pixels. We continue by considering the two sets S and T separately, and applying the algorithm recursively, at each time halving the λ interval until we have the value of every pixel.

Combining the divide and conquer approach with the parametric push-relabel algorithm is a bit problematic. For all pixels $x \in T$ we know that $u_x \leq \lambda$, and we can reuse the graph when decreasing λ . However for pixels $x \in S$, we seek to find $u_x > \lambda$, meaning we have to increase λ which does not allow graph re-use.

Goldfarb and Yin [25] have found that the divide and conquer approach only yields improved performance when using the L^2 norm in the fidelity term.

See [26], [27] and [25] for more information.

5.5 Boykov–Kolmogorov algorithm

Algorithm 6 The Boykov–Kolmogorov maximum flow algorithm

```

function BOYKOVKOLMOGOROV( $G, s, t$ )
   $A := \{s, t\}, O := \emptyset, S = \{s\}, T = \{t\}$ 
   $s.\text{color} = S, t.\text{color} = T$ 
  while True do
     $e \leftarrow \text{GROW}(G, A)$ 
    if not  $e$  then
      break
    end if
    AUGMENT( $G, s, t, e, O$ )
    ADOPT( $G, s, t, O$ )
  end while
end function

```

A maximum flow algorithm specialized for the type of graphs found in imaging applications is described by Boykov and Kolmogorov in [21]. Their algorithm is an augmenting path algorithm where the paths are found using trees that keep track of possible partial paths.

The idea of the algorithm is to always maintain the structure of two trees, one rooted in the source and one rooted in the sink, which are built up using only edges that can carry additional flow. The algorithm consists of different phases, one being the *grow* phase, where the trees are grown by adding additional edges. When an unsaturated edge connecting the two trees are found, the *augment* phase starts. In this phase, an augmenting path through the trees is recovered, and the flow is augmented. Some of the edges in the trees may then become saturated, and can no longer be a part of their tree. The vertices they connect to the tree become *orphans* which in the *adopt* phase are adopted back into their trees, or alternatively freed from their tree connection.

There is one tree denoted S with the source s as its root, and another tree denoted T with the sink t as its root. These trees are disjoint, and all edges in the tree S can carry flow towards the leaves of the tree, while all edges in T can carry flow towards the sink t .

A vertex can either be part of one of these trees, or be a free vertex, and we write

$$u.\text{color} = \begin{cases} S & \text{if } u \in S \\ T & \text{if } u \in T \\ \text{Free} & \text{otherwise.} \end{cases} \quad (5.14)$$

The vertices in the trees S and T are either active or passive, and we write $u \in A$ if u is active. The active vertices are those at the boundary of the tree, which can possibly connect to other vertices to grow the tree. The passive vertices are internal in the tree, and edges to their neighbors are either a part of the tree, or completely saturated.

The main loop of the algorithm is as shown in Algorithm 6. In the grow procedure, the trees are grown from their active vertices until a new path is found. The flow is then augmented along this path in the augment procedure, and orphan vertices might be created. These orphan vertices are then either joined back into their respective trees, or become free, in the adopt procedure.

Algorithm 7 The grow procedure of the Boykov–Kolmogorov maximum flow algorithm

```

function GROW( $G, A$ )
  while  $|A| \neq 0$  do
     $u \leftarrow$  one node from  $A$ 
    for all  $v$  such that  $\text{TREECAP}(u, v) > 0$  do
      if  $v.\text{color} = \text{Free}$  then
         $v.\text{color} \leftarrow u.\text{color}$ 
         $v.\text{parent} \leftarrow u$ 
         $A := A \cup \{v\}$ 
      else if  $v.\text{color} \neq u.\text{color}$  then
        return  $e := (u, v)$ 
      end if
    end for
    Remove  $u$  from  $A$ 
  end while
  return False
end function

```

The grow procedure is shown in Algorithm 7. It goes through the set of active vertices A to try to expand the trees S and T . When considering an active vertex $u \in S$, we want to grow by finding vertices v such that $c_f(u, v) > 0$, while when considering an active vertex $u \in T$, we want to grow by finding vertices v such that $c_f(v, u) > 0$. This is why the TREECAP is introduced which is defined as

$$\text{TREECAP}(u, v) = \begin{cases} c_f(u, v) & \text{if } u \in S, \\ c_f(v, u) & \text{if } u \in T. \end{cases} \quad (5.15)$$

Thus if a non-saturated edge is found from an active vertex u to a free vertex v , then v is added to the tree of u . The tree structure is stored by keeping a parent attribute in each non-free node, such that for example $v.\text{parent} = u$.

If a non-saturated edge is found from an active vertex u to a vertex v in the other tree, the two trees connect, and we can return the edge e since we have an augmenting path. Note that if this happens, the vertex u is still active. It only becomes passive when all neighbors are considered without yielding an augmenting path.

Algorithm 8 The augment procedure of the Boykov–Kolmogorov maximum flow algorithm

```

function AUGMENT( $G, s, t, e, O$ )
     $p \leftarrow$  path from  $s$  to  $t$  through  $e$                                  $\triangleright$  through the trees  $S$  and  $T$ 
     $\Delta f \leftarrow \min\{c_f(u, v) : (u, v) \in p\}$ 
    for all  $(u, v) \in p$  do
         $f(u, v) += \Delta f$ 
         $f(v, u) -= \Delta f$ 
        if  $c_f(u, v) = 0$  then
            if  $u.\text{color} = v.\text{color} = S$  then
                 $v.\text{parent} = \text{Null}$ 
                 $O := O \cup \{v\}$ 
            else if  $u.\text{color} = v.\text{color} = T$  then
                 $u.\text{parent} = \text{Null}$ 
                 $O := O \cup \{u\}$ 
            end if
        end if
    end for
end function

```

The augment procedure finds the augmenting path going through the tree S , the edge e and the tree T . The maximal possible flow is then sent along this path. At least one edge will then become saturated. If a saturated edge occurs in the trees S or T , the edge terminal farthest from the tree root is marked as an orphan. Note that even if a whole subtree is disconnected from the main tree, only the root of this subtree is marked as an orphan. The adopt procedure will take care of either reconnecting, or freeing all the vertices in the disconnected subtree.

The adopt procedure processes all the vertices in the set of orphans O . These vertices are either single vertices, or roots of disconnected subtrees. For an orphaned vertex u we look through its neighbors to find a possible parent vertex v . It should belong to the same tree as u , and also satisfy $\text{TREECAP}(v, u) > 0$. In addition, v should be connected to one of the tree roots s or t . This keeps us from connecting to other orphaned vertices or subtrees, and is checked by the procedure $\text{TREEORIGIN}(v)$, which follows the parent information until reaching either s , t or an orphaned vertex.

Algorithm 9 The adopt procedure of the Boykov–Kolmogorov maximum flow algorithm

```

function ADOPT( $G, s, t, O$ )
  while  $|O| \neq 0$  do
     $u \leftarrow$  one node from  $O$ 
    Remove  $u$  from  $O$ 
    found  $\leftarrow$  False
    for all  $v$  such that  $\text{TREECAP}(v, u) > 0$  do
      if  $u.\text{color} \neq v.\text{color}$  then
        continue
      end if
      orig  $\leftarrow$   $\text{TREEORIGIN}(v)$ 
      if orig  $\neq s$  and orig  $\neq t$  then
        continue
      end if
      found  $\leftarrow$  True
       $u.\text{parent} \leftarrow v$ 
      break
    end for
    if found  $\neq$  True then
      for all  $v$  such that  $u.\text{color} = v.\text{color}$  do
        if  $\text{TREECAP}(v, u) > 0$  then
           $A := A \cup \{v\}$ 
        end if
        if  $v.\text{parent} = u$  then
           $O := O \cup \{v\}$ 
           $v.\text{parent} \leftarrow \text{Null}$ 
        end if
         $u.\text{color} \leftarrow \text{Free}$ 
         $A := A - \{u\}$ 
      end for
    end if
  end while
end function

```




Figure 5.3: A modified version of the subgraph corresponding to the fidelity term $F_\lambda^x(u_x^\lambda)$ that has positive and non-decreasing edges with decreasing λ .

If a potential parent v of u is *not* found, u becomes a free vertex. All vertices that had u as its parent are orphaned, and are thus treated by the adopt procedure later. Vertices that are in one of the trees S or T , and have a non-saturated edge to this newly freed vertex u , become active.

When no more active vertices remain, there are no more possible augmenting paths, and the algorithm terminates with a maximum flow. A proof of correctness can be found in Kolmogorov's PhD thesis [28].

5.5.1 Graph reuse

As with the push-relabel algorithm described earlier, it is not necessary to completely restart the Boykov–Kolmogorov algorithm for every level of the image, but we can reuse the graph and the trees S and T in successive runs.

Since this algorithm does not deal with the relaxed preflow concept, the updates from one level to the next have to be done in a different way. We can no longer decrease the capacity of edges, as this could break the flow conservation constraint. However, the graph construction in Section 4.2.2 did allow for the addition of an arbitrary constant to all the edges of the sub-graph. Thus the construction in Figure 5.3 is also valid. We have added the non-negative constant $\max\{F_{L-1}^x(1), 0\}$ to all the edges of the graph in Figure 4.7b. As λ goes from $L-1$ to 0 both edges stay non-negative. The edge from u_x^λ to t is non-decreasing with decreasing λ parameter.

This construction allows us to update the capacities of the edges, while retaining the flow and the trees S and T . A nice property that follows is that the partition $(S, V - S)$ is after each run a minimum cut since the tree S has been

Table 5.1: Each row represents one of the two possible values of $u_x^\lambda \in \{0, 1\}$. The functional $F_\lambda^x(u_x^\lambda)$ and minimum cut obtaining this configuration is shown. The last column shows the capacities of the cut for the graph construction in Figure 5.3. We verify that the cut capacities are equal to the functional value, plus a constant.

u_x^λ	$F_\lambda^x(u_x^\lambda)$	Min. cut (S, T)	Graph cut cap.
0	0	$(\{s\}, \{u_x^\lambda, t\})$	$\max\{F_{L-1}^x(1), 0\} - F_\lambda^x(1)$
1	$F_\lambda^x(1)$	$(\{s, u_x^\lambda\}, \{t\})$	$\max\{F_{L-1}^x(1), 0\}$

grown as much as possible. Further, no vertex already in the tree S will leave S when the trees are updated and algorithm is run for lower λ values. This can be seen from the fact that if u is in S and the algorithm has terminated, all paths from s to t going through u have a saturated edge somewhere after u . For all these paths, the capacity will only change for edges before u , as we only change the capacity of edges (s, v) . Thus these saturated edges will stay saturated, and u will stay in S .

5.5.2 Performance improvements

There are several open choices in the implementation of the algorithm, for example the order in which active vertices are processed. As recommended in [21], we have implemented a «First-In-First-Out» queue. This ensures that at least the first augmenting path found is a shortest path, although later the distance information is lost in the adoption stage.

In the adoption stage it is possible, and perhaps preferable to seek a possible parent that is closest to the root of the tree, and adopt that vertex as parent, instead of the first one found.

Because of the particular graph construction, all vertices except s and t are connected directly to s and t . Thus there are as many two-edged paths from s to t as there are pixels in the image. When increasing the capacity of edges (v, t) , a quick sweep over these two-edged paths to send any possible flow may speed up the algorithm.

Results

In the previous chapters we have carefully constructed a method always trying to argue how the choices we make can have a positive impact on the image restoration results. The anisotropic total variation is the main advance from my project work [1], and we would like to see how the introduction of the anisotropy affects the performance of the restoration algorithm.

6.1 Tensor parameters

The anisotropy was introduced into the total variation in order to lessen the regularization applied across what we know, or at least are pretty sure to be edges in the image. Before considering the anisotropy tensor construction described earlier, we will look at how a simple pre-described tensor affects the regularization. Imagine a tensor which is a diagonal matrix $A = \text{diag}(1, \epsilon)$, where $\epsilon \ll 1$ is small. This would result in us down-weighting the size of ∇u in the y -direction, and thus regularization mostly in the x -direction.

The results of this experiment can be seen in Figure 6.1, where a noisy picture of a circle has been restored in two different ways, first with a uniform anisotropy tensor $A = \text{diag}(1, \epsilon)$, and then with the tensor described in Section 3.1.1. We see in Figure 6.1b that the uniform tensor gives a strong smoothing in the x -direction, but no apparent smoothing in the y -direction.

It is useful to look at intermediate results to better understand what happens and how the parameters affect our end result. In Figure 6.2 we show the different stages of our image restoration algorithm. We start out with a section of the much-used Lena test image in Figure 6.2a. The first step in the construction of the structure tensor is to apply a Gaussian blur with parameter σ , and the result is shown in Figure 6.2b. The blurring is done so that the edge detector ∇f_σ ,

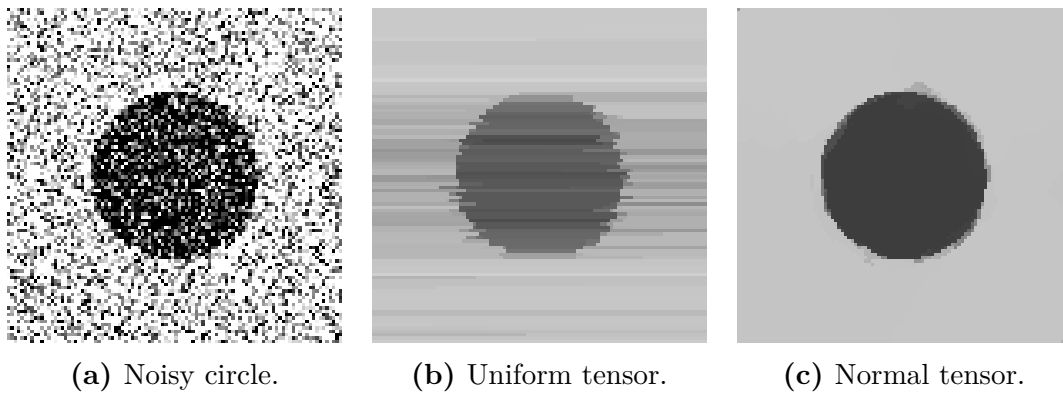


Figure 6.1: A noisy circle first restored using a uniform tensor, resulting mostly in smoothing in the x -direction. Then restored using the tensor described in Section 3.1.1.

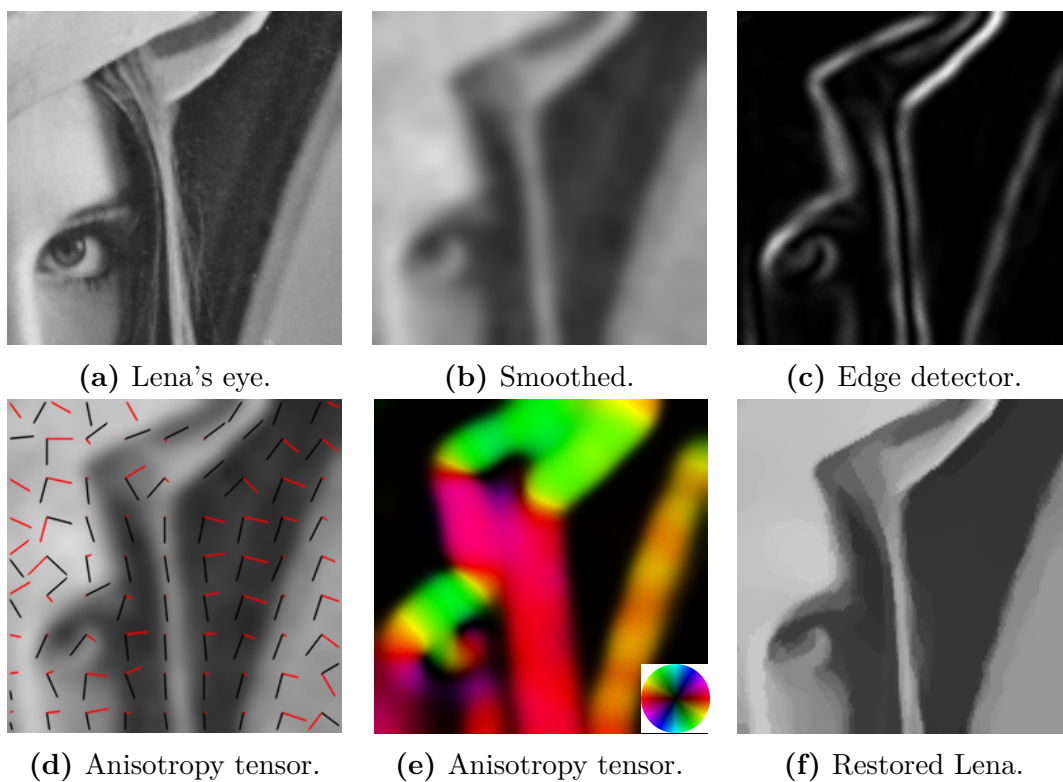


Figure 6.2: The different stages of the restoration algorithm, showing the original image, the smoothed image, the edge detector, two visualizations of the anisotropy tensor, and finally the restored image.




Figure 6.3: Color wheel used for tensor visualization.

visualized as $|\nabla f_\sigma|^2$ in Figure 6.2c, is not too sensitive to noise in the image. The structure tensor is then constructed using ∇f_σ and smoothed according to the integration scale ρ . The resulting anisotropy tensor is visualized in Figure 6.2d. In a selection of points, the eigenvectors of the tensor have been drawn. The length of the vectors have been scaled by the corresponding eigenvalue.

Another way of visualizing the tensor is shown in Figure 6.2e. The brightness and color in each point is decided by the size of the smallest eigenvalue and the direction of its corresponding eigenvector using the color wheel in Figure 6.3. The direction decides the color, while the brightness, which is the radius in the color wheel, is set to $1/\sigma_1 - 1$, where σ_1 is the smallest eigenvalue. Thus the stronger the anisotropy, the brighter the color, while we expect uniform areas in the original image to be black in the tensor visualization.

Finally, Figure 6.2f shows the restored image, and we see that it has been heavily regularized. How the anisotropy affected the regularization is not obvious however.

To look further into the effects of the parameters in our tensor construction we have a constructed a zebra pattern of increasing width as shown in Figure 6.4. The anisotropy introduced should in theory help reduce contrast loss in this situation, by reducing the regularization applied in the x -direction across the edges. There is however the question of how the different scales σ and ρ affect the regularization.

In Figure 6.4b the noise scale is increased such that the edge detector, and thus the anisotropy tensor, considers the finest lines to be noise rather than details. Thus this left-most part is regularized to a smooth gray area.

The integration scale ρ of (3.5) controls the size of the structures we want to detect with our anisotropy tensor. In Figure 6.4c a high value for ρ means that the structure found in the inner square is almost completely ignored in favor of the larger, more coherent structure around it.

Another visualization of the anisotropy tensor can be seen for a fingerprint image in Figure 6.5. We see that the gradual changes in the direction of the

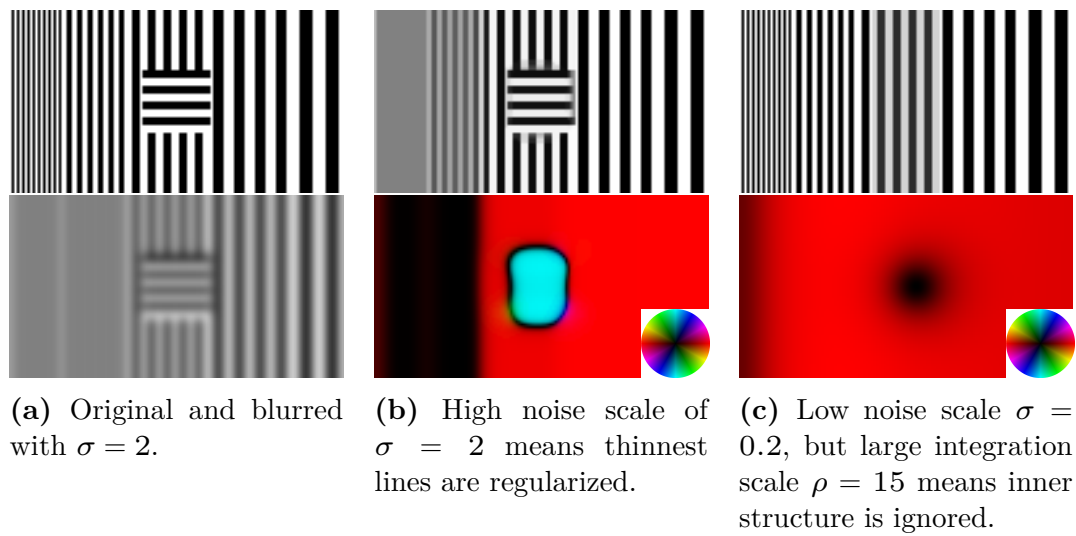


Figure 6.4: An example constructed to show the effects of the parameters σ and ρ in the anisotropy tensor. The noise scale σ controls the smoothing done before edge detection. The integration scale ρ controls the size of the structures considered by the tensor. Parameters: $|\mathcal{N}| = 8$, $\beta = 10000$, $\omega = 100$.




Figure 6.5: A noisy fingerprint with the anisotropy tensor visualized for different integration scales ρ . Parameters: $|\mathcal{N}| = 32$, $\sigma = 3$, $\beta = 15000$, $\omega = 150$.

edges are captured by the tensor. Even the singularities in the fingerprint can be identified by finding the dark spots, as the gradient of the image in these areas do not point in one single direction.

6.2 Neighborhood stencils

The neighborhoods were introduced as a way to describe the discrete set of lines \mathcal{L}_D in the discretization of the regularization term in Section 4.1.2. And as discussed, we want the stencil to have many short edges, such that the angular differences $\Delta\phi$, the inter-line distances $\Delta\rho$ and the edge lengths e are “small.”

Figure 6.6 shows three images of different shapes that are heavily regularized using different neighborhood stencils. In these results we see some of the artifacts that may occur because of this particular discretization method. It is clear that the stencil of size 8 prefers horizontal, vertical and 45° perimeters, such that the circle in Figure 6.6a is shaped like an octagon when regularized in Figure 6.6b. However the stencil of size 72 manages to keep the circular shape.

In Figure 6.6d, we see that the stencil of size 8 actually favors octagon-like shapes, as the restored shape is the same octagon, just with some contrast loss.

The tilted square in Figure 6.6g, hints that a stencil of size 4 neighborhood favors horizontal and vertical edges only. While the larger stencil of size 72 retains the square shape but rounds the corners some.

As mentioned before, there is a discretization error which relates to the length of the edges in the neighborhood. We approximated the number of times an edge e_{ab} crosses level set boundary by $|u_a^\lambda - u_b^\lambda|$, an approximation that becomes worse for long edges. Thus a larger neighborhood is not always better, even if it will reduce the artifacts discussed above. In Figure 6.7, a noisy image of Lena has been restored using two neighborhood stencils, and there are obvious differences. For the stencil of size 72, the restored image still contains some pixel-sized noise. An explanation can be found in Figure 6.8. We see why the length of a one-pixel curve is underestimated by the Cauchy–Crofton formula when some of the edges are long. Because many edges e_{ab} cross the curve cross while $|u_a^\lambda - u_b^\lambda| = 0$, thus these edges are ignored completely in our perimeter approximation.

An additional demonstration that this problem mostly relates to small sized noise is shown in Table 6.1. The table shows how our discrete Cauchy–Crofton formula approximates the circumference of circles of different radii. Note that the circumference approximated is that of an actual continuous circle $u : \mathbb{R}^2 \rightarrow \{0, 1\}$ and not a discrete representation. We see that the perimeter of the smallest circle is grossly underestimated by the larger neighborhoods. Thus the perimeter of one-pixel noise will be underestimated, and in turn the contribution to the total variation by one-pixel noise will be smaller for large neighborhoods.




Figure 6.6: Different test images restored without anisotropy, different neighborhoods and a large restoration parameter β . We see how different neighborhood stencils introduce different artifacts.




Figure 6.7: Noisy Lena restored using different neighborhood stencils. Note how the large neighborhood introduces some pixel size artifacts.




Figure 6.8: A selection of edges contributing to the wrong approximation of the length of a single pixel's perimeter. The edges cross the perimeter twice, but we count zero crossings in our approximation. This leads to residual pixel noise in the restored image, when large neighborhood stencils are used.

Table 6.1: The circumference of circles of different radii r measured by the discretized Cauchy–Crofton formula in (4.15), using different neighborhood stencils.

$ \mathcal{N} $	$r = 0.5$	$r = 1.5$	$r = 5.5$	$r = 50.5$
4	3.14	9.42	34.56	317.3
8	2.68	10.27	33.94	317.5
16	2.08	9.97	34.59	316.8
32	1.63	9.24	34.44	317.2
48	1.40	8.40	34.29	317.3
72	1.21	7.45	33.95	317.2
$2\pi r$	3.14	9.42	34.56	317.3

6.3 Restoration

We have claimed that the anisotropic total variation will reduce some of the contrast loss one can encounter with regular total variation regularization. In Figure 6.9 the noisy fingerprint from Figure 6.5a has been restored three times with different parameters. The first using regular total variation and $\beta = 15000$, then the second with $\beta = 15000$ and anisotropy $\omega = 150$. This gives an obvious contrast enhancement, but it is however not obvious what this tells us about the quality of the anisotropic algorithm. As previously mentioned, increasing the anisotropy (decreasing ω) means decreasing the total amount of regularization applied, and less contrast loss is an expected outcome of decreasing the regularization, disregarding the anisotropy.

Thus in the last image we have also used $\omega = 150$ but the restoration amount β is increased such that the amount of noise removed $\|u - f\|_{L^2}$ is approximately equal to the noise removed when using regular total variation in Figure 6.9a. By visual inspection, the last image seems to have somewhat higher contrast. This can be confirmed by calculating the standard deviation of the images, one of many possible contrast measures, which gives 22.9, 37.9 and 27.2 respectively.

To further compare the three restored images, a single row has been extracted from the three images and is shown against each other in Figure 6.10. Note that although the results in Figure 6.9c look promising compared to Figure 6.9a, some details in the singularity in the upper left is actually lost. The metric tensor is approximately the identity matrix there, and thus an increased restoration parameter β leads to increased smoothing compared to the regular total variation of Figure 6.9a.

So far we have just inspected the restored image visually to judge the results of the method. Depending on the application, we might have different wishes for the

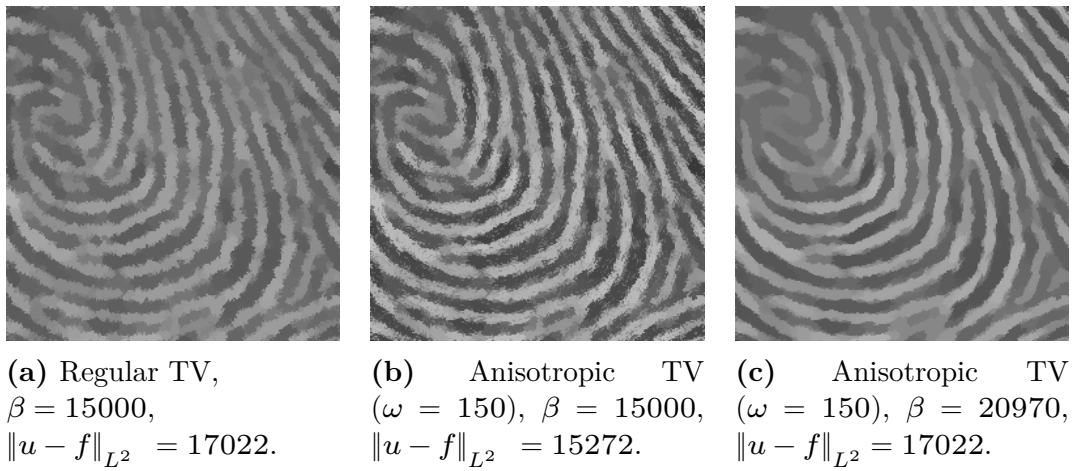


Figure 6.9: A noisy fingerprint restored using both isotropic and anisotropic total variation to see how the contrast loss compares. Parameters: $|\mathcal{N}| = 32$, $\sigma = 3$, $\rho = 10$.

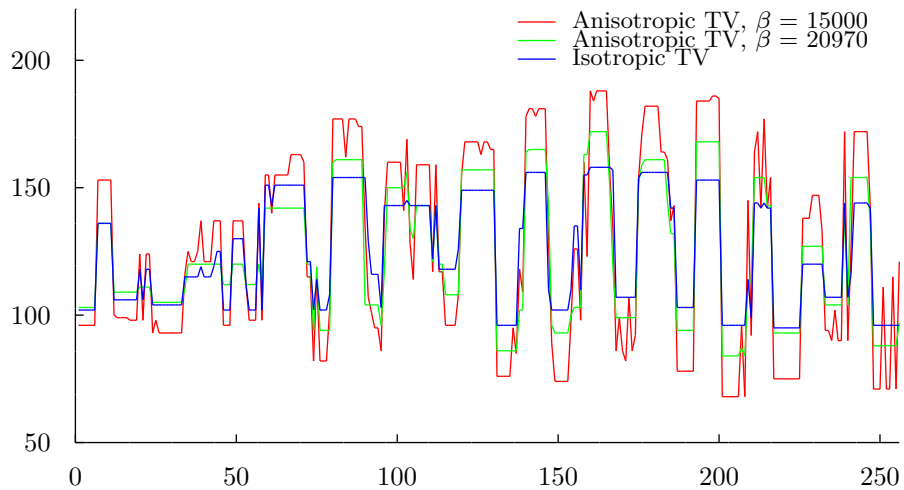


Figure 6.10: A one-dimensional slice through the restored fingerprint images of Figure 6.9, showing a difference in the contrast.

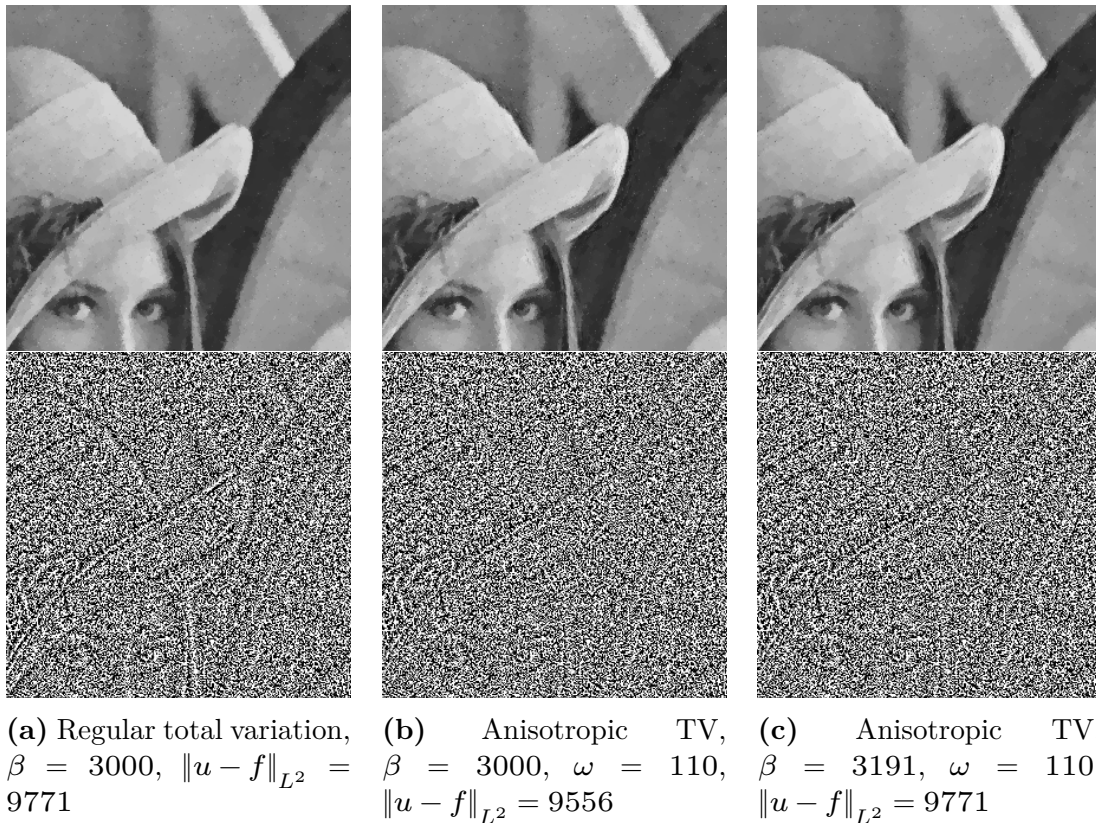


Figure 6.11: Lena with additive Gaussian noise ($\sigma = 30$) shown restored with different parameters, with corresponding method noise below. A slight improvement in the amount of detail in the method noise is seen from (a) to (c).

results. In many cases we just want to remove the noise δ in $f = u^* + \delta$ where the captured image f is assumed to consist of an actual image u^* and an additional noise term δ . The noise removed by the restoration method, $\hat{\delta} = f - u^*$ is called the method noise. Any assumptions of properties we have on δ , we optimally want $\hat{\delta}$ to fulfill as well. Thus if we assumed independently distributed Gaussian noise, we want to remove that kind of noise. An indication of problems can in that case be if the method noise contains a lot of the details of the image.

Figure 6.11 shows the results after restoring a noisy Lena picture with different parameters, and their method noise. We see that the method noise of Figure 6.11a shows slightly more details than the anisotropic counterpart in Figure 6.11c.

Discussion and conclusion

In this thesis we have seen how the total variation restoration method can be extended using an anisotropy tensor, and how this fits into the discretization and graph cut framework used in my project work [1]. The anisotropy tensor was introduced in hopes of reducing the amount of regularization applied across edges in the image, and by that preventing contrast loss. It was constructed based on the structure tensor which contains local information about edge direction and steepness.

In order to arrive at a discrete functional that could be minimized using the same graph cut framework as in my project work [1], we transformed the continuous functional using an anisotropic coarea and Cauchy–Crofton formula, both described in detail. The resulting discrete functional was shown—under some restrictions—to be consistent with the continuous functional.

The successive graph cuts, each giving a minimizer for a functional of a level in the image, were found using maximum flow algorithms. The push-relabel algorithm is known to have good all-round performance, while the Boykov–Kolmogorov algorithm is tailored for the kinds of graphs appearing in imaging applications. They were both described and implemented.

A goal was to give readers a good understanding of the inner workings of the method, as well as all details necessary for a working implementation. Thus we put effort into describing all steps going from the initial continuous formulation, to the discretization, and finally the maximum flow algorithms.

One part of the thesis consists of studies of how the different parameters affect the performance of the algorithm. We will not give any unified conclusion as to whether this method is “better” or “worse” than the regular total variation method it is based on, or other methods. In different applications the input images have different properties, and one might also have different hopes for, and restrictions on the output image. But we have seen how the anisotropy tensor affects the

restoration, and that it can have positive effects on contrast loss as well as method noise.

Further work is possible in the study of the continuous problem, its well-foundedness, and also the anisotropic coarea and Cauchy–Crofton formulas, which can be studied on a measure-theoretic foundation. Also, the construction of the anisotropy tensor offers choices that can be explored further.

Regarding the discretization, the choice of neighborhood stencil also allows for further discussion, as approximation error can be traded for algorithm performance. One particular possibility would be the application of non-uniform stencils with varying stencil size depending on the local level of detail. This could give better performance without sacrificing too much solution accuracy.

Bibliography

- [1] Bjørn Rustad. Total variation based image restoration using graph cuts. 2014.
- [2] Bernd Jähne. *Digital Image Processing*. Springer-Verlag, Berlin Heidelberg, 2005.
- [3] Gilles Aubert and Pierre Kornprobst. *Mathematical problems in image processing*, volume 147 of *Applied Mathematical Sciences*. Springer, New York, second edition, 2006.
- [4] Joachim Weickert. *Anisotropic diffusion in image processing*. European Consortium for Mathematics in Industry. B. G. Teubner, Stuttgart, 1998.
- [5] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1–4):259–268, 1992.
- [6] V. Caselles, A. Chambolle, and M. Novaga. Total variation in imaging. In Otmar Scherzer, editor, *Handbook of Mathematical Methods in Imaging*, pages 1016–1057. Springer, New York, 2011.
- [7] Raymond Chan, Tony Chan, and Andy Yip. Numerical methods and applications in total variation image restoration. In *Handbook of Mathematical Methods in Imaging*, pages 1059–1094. Springer, 2011.
- [8] Markus Grasmair and Frank Lenzen. Anisotropic total variation filtering. *Appl. Math. Optim.*, 62(3):323–339, 2010.
- [9] C. Olsson, M. Byrod, N.C. Overgaard, and F. Kahl. Extending continuous cuts: Anisotropic metrics and expansion moves. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 405–412, Sept 2009.

- [10] Joachim Weickert. Coherence-enhancing diffusion filtering. *International Journal of Computer Vision*, 31(2-3):111–127, 1999.
- [11] Robert E. Megginson. *An introduction to Banach space theory*, volume 183 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998.
- [12] Otmar Scherzer, Markus Grasmair, Harald Grossauer, Markus Haltmeier, and Frank Lenzen. *Variational methods in imaging*, volume 167 of *Applied Mathematical Sciences*. Springer, New York, 2009.
- [13] Lawrence C. Evans and Ronald F. Gariepy. *Measure theory and fine properties of functions*. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL, 1992.
- [14] Robert G Bartle. *The Elements of Integration and Lebesgue Measure*. John Wiley & Sons, 1995.
- [15] William P. Ziemer. *Weakly differentiable functions*, volume 120 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1989.
- [16] Edwin Hewitt and Karl Stromberg. *Real and abstract analysis. A modern treatment of the theory of functions of a real variable*. Springer-Verlag, New York, 1965.
- [17] Manfredo P. do Carmo. *Differential geometry of curves and surfaces*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1976.
- [18] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, third edition, 2009.
- [19] Vladimir Kolmogorov and Ramin Zabih. What energy functions can be minimized via graph cuts? In *Computer Vision—ECCV 2002*, pages 65–81. Springer, 2002.
- [20] Yefim Dinitz. Dinitz’ algorithm: the original version and Even’s version. In *Theoretical computer science*, volume 3895 of *Lecture Notes in Comput. Sci.*, pages 218–240. Springer, Berlin, 2006.
- [21] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(9):1124–1137, 2004.
- [22] Andrew V. Goldberg and Robert E. Tarjan. A new approach to the maximum-flow problem. *J. Assoc. Comput. Mach.*, 35(4):921–940, 1988.

- [23] Joseph Cheriyan and S. N. Maheshwari. Analysis of preflow push algorithms for maximum network flow. *SIAM Journal on Computing*, 18(6):1057–1086, 1989.
- [24] B. V. Cherkassky and A. V. Goldberg. On implementing the push-relabel method for the maximum flow problem. *Algorithmica*, 19(4):390–410, 1997.
- [25] Donald Goldfarb and Wotao Yin. Parametric maximum flow algorithms for fast total variation minimization. *SIAM Journal on Scientific Computing*, 31(5):3712–3743, 2009.
- [26] Giorgio Gallo, Michael D. Grigoriadis, and Robert E. Tarjan. A fast parametric maximum flow algorithm and applications. *SIAM J. Comput.*, 18(1):30–55, 1989.
- [27] Dorit S. Hochbaum. An efficient algorithm for image segmentation, Markov random fields and related problems. *J. ACM*, 48(4):686–701 (electronic), 2001.
- [28] Vladimir Kolmogorov. *Graph Based Algorithms for Scene Reconstruction from Two or More Views*. PhD thesis, Cornell University, September 2003.
- [29] Bjørn Rustad. C++ implementation. <https://github.com/burk/image-restoration>.
- [30] G. Bradski. Opencv library. *Dr. Dobb's Journal of Software Tools*, 2000.

List of Figures

2.1	The stair-casing effect of regular total variation restoration	8
2.2	Heavily regularized fingerprint showing contrast loss	8
3.1	The edge detector $ \nabla f_\sigma(x) $	13
3.2	Eigenvectors and eigenvalues of the structure and anisotropy tensors	15
3.3	Lower semi-continuous function f	17
3.4	The cut-off function $\eta_r(t)$ and its derivative	22
3.5	Line parametrization by angle and distance	24
4.1	The discrete set of lines \mathcal{L}_D	34
4.2	Visualization of the approximation error relating to long edges . . .	35
4.3	A visual argument for $\delta^2 = \Delta\rho \ e\ $	36
4.4	Consistency argument for the line distance ρ	39
4.5	Consistency argument for the angle discretization	39
4.6	Consistent neighborhood construction	40
4.7	Graph construction for the fidelity term F_λ^x	43
4.8	Graph construction for the regularization term $F^{x,y}$	45
4.9	The structure of the final graph construction	46
5.1	Augmenting flow example graph	50
5.2	Fidelity energy $F_\lambda^x(1)$ as a function of λ	57
5.3	Modified fidelity subgraph for the Boykov–Kolmogorov algorithm .	64
6.1	Restoration using a uniform tensor	67
6.2	The different stages of the restoration algorithm visualized	67
6.3	Color wheel	68
6.4	Effects of the σ and ρ parameters on restoration	69
6.5	Anisotropy tensor visualized for a fingerprint test image	69
6.6	Shapes restored to show artifacts introduced by neighborhood choice	71

6.7	Noisy Lena restored using different neighborhood stencils	72
6.8	Approximation error when calculating the perimeter of one pixel . .	72
6.9	Comparison of contrast loss in regular and anisotropic total variation	74
6.10	1D slice through restored fingerprints showing contrast differences .	74
6.11	Lena restored, shown with method noise	75

List of Tables

4.1	Graph construction for $F_\lambda^x(u_x^\lambda)$	44
4.2	Graph construction for $F^{x,y}(u_x^\lambda, u_y^\lambda)$	46
5.1	Graph construction for $F_\lambda^x(u_x^\lambda)$ in the Boykov–Kolmogorov algorithm	65
6.1	Circle circumferences estimated by the Cauchy–Crofton formula . .	73

List of Symbols

$ C $	The length of the curve C , page 25
$ C _M$	The length of the curve C using the metric tensor M , page 25
$\alpha(\nabla u)$	Scalar thermal diffusivity, page 4
β	Regularization parameter, page 6
$BV(\Omega)$	Functions of bounded variation in Ω , page 7
δ	Noise function, page 3
$\ell_{\nu,\rho}$	Line given by tangent vector ν and distance to origin ρ , page 25
$\ell_{\phi,\rho}$	Line given by angle of normal ϕ and distance to origin ρ , page 25
χ_E	Characteristic function of the set E , page 9
$\ \eta\ _A^*$	The norm $\sup_x (\eta^T A^{-1} \eta)^{1/2}$, page 12
$\ \xi\ _A$	The norm $\sup_x (\xi^T A \xi)^{1/2}$, page 12
\mathcal{G}	Regular grid of pixels over Ω , page 9
\mathcal{L}	The set of all straight lines in the plane, page 25
$\mathcal{N}(x)$	Neighborhood of pixel x , page 35
\mathcal{P}	Discrete set of pixel values, or levels, page 9
∇f_σ	Edge detector, page 13
Ω	Rectangular, open domain, page 3

ω	Anisotropy parameter, page 15
$\text{Per}_A(U; \Omega)$	Anisotropic perimeter of set U using anisotropy tensor A , page 23
\tilde{u}	Symmetric extension of u from Ω to \mathbb{R}^2 , page 6
$\text{TV}(u)$	Total variation of u , page 7
$\text{TV}_A(u)$	Anisotropic total variation of u given anisotropy tensor A , page 12
$A(u)$	Thermal diffusivity tensor, page 5
$A(x)$	Anisotropy tensor, page 11
$c(u, v)$	Capacity function $c : V \times V \rightarrow [0, \infty)$, page 41
$C^\infty(\Omega)$	Space of infinitely differentiable functions from Ω to \mathbb{R} ., page 6
$C_c^\infty(\Omega, \mathbb{R}^2)$	Space of smooth functions from Ω to \mathbb{R}^2 with compact support in Ω , page 7
$d(u)$	Height map or distance labeling $d : V \rightarrow \mathbb{N}$, page 52
$e(u)$	Excess in vertex u , page 52
f	Original, noisy image, page 3
$G = (V, E, c)$	Graph given by the vertices V , edges E , and capacity function c , page 41
$J_M(\ell_{\phi, \rho})$	Jacobian of the coordinate transformation induced by the metric tensor M , page 26
$K_\sigma(x, y)$	The Gaussian kernel, page 4
$L^p(\Omega)$	Real functions f on Ω for which $\int_\Omega f ^p < \infty$, page 6
$M(x)$	The metric tensor, page 25
$N_x(k)$	Helper function $N_x(k) = k - f_x ^2$, page 33
s	The source vertex, page 41
$S_\rho(x)$	Structure tensor, page 6
t	The sink vertex, page 41
u^*	Usually unknown, actual image without noise, page 3
u^λ	Thresholded image at level λ , page 9

Appendix **A**

C++ implementation

A C++ implementation of the method described is included here and can also be found online at [29]. It uses the open computer vision library OpenCV [30] to load and save image files and contains compilation and usage instructions. The implementation has been tested on an installation of the Ubuntu Linux distribution, but it should in theory be portable to other platforms supported by OpenCV. A rudimentary graphics interface has also been made, to make it easier to play with the parameters of the algorithm.

Both the push-relabel and the Boykov–Kolmogorov algorithms have been implemented. Although an effort has been made to improve the performance of both implementations, they are not ment to beat the fastest. The focus has rather been on clarity and understanding.

Note that when implementing maximum flow algorithms it is not a good idea, memory- and performance-wise, to actually construct the residual graph G_f . Instead, every time we update the flow $f(u, v)$ we set the flow in the opposite direction to its negative value $f(v, u) = -f(u, v)$. Then we can at any time, consider the value $c(u, v) - f(u, v)$ in the place of the residual capacity $c_f(u, v)$.

For the gap relabeling heuristic of the push-relabel algorithm, we need to have a easy way of finding when a gap occurs. This is done by keeping track of how many vertices exist with each label.

When the capacities have been updated in the Boykov–Kolmogorov algorithm, flow is sent along all two-edge paths such that they do not have to be considered by the main loop of the algorithm.

A.1 main.cpp

Main function of the command line executable. Here we read and parse command line parameters.

```

1  #include <iostream>
2  #include <cstdlib>
3  #include <cmath>
4  #include <cstdio>
5  #include <unistd.h>
6  #include <opencv2/opencv.hpp>
7  #include "graph.hpp"
8  #include "selectionrule.hpp"
9  #include "neighborhood.hpp"
10 #include "image.hpp"
11 #include "anisotropy.hpp"
12
13 using namespace std;
14 using namespace cv;
15
16 int main(int argc, char *argv[])
17 {
18     int p = 2;
19     double beta = 10;
20
21     int neighbors = 8;
22
23     double sigma = 10.0;
24     double rho = 10.0;
25     double gamma = 10000.0;
26
27     int c;
28
29     /* Read command line parameters beta and p. */
30     while ((c = getopt(argc, argv, "b:p:r:s:g:n:fh")) != -1) {
31         switch (c)
32         {
33             case 'p':
34                 p = atoi(optarg);
35                 break;
36             case 'b':
37                 beta = atof(optarg);
38                 break;
39             case 'g':
40                 gamma = atof(optarg);
41                 break;
42             case 'r':
43                 rho = atof(optarg);
44                 break;
45             case 's':
46                 sigma = atof(optarg);
47                 break;
48             case 'n':
49                 neighbors = atoi(optarg);
50                 break;
51             case '?':
52                 if (optopt == 'p' || optopt == 'b' || optopt == 'g'
53                     || optopt == 'n' || optopt == 'r'
54                     || optopt == 's') {
55                     fprintf(stderr, "Option -%c requires an argument.\n",
56                             optopt);
57                 }
58                 else if (isprint(optopt)) {
59                     fprintf(stderr, "Unknown option `-%c'.\n", optopt);

```

```

60         }
61         else {
62             fprintf(stderr, "Unknown option character '\\x%x'.\n",
63                     optopt);
64         }
65         return 1;
66     default:
67         exit(1);
68     }
69 }
70
71 /*
72  * Non-option arguments are now in argv from index optind
73  * to index argc-1
74  */
75 Mat image;
76 image = imread(argv[optind], CV_LOAD_IMAGE_GRAYSCALE);
77
78 if (!image.data) {
79     cout << "Loading image failed" << endl;
80     return -1;
81 }
82
83 cout << "Using gamma = " << gamma << endl;
84 cout << "Using rho = " << rho << endl;
85 cout << "Using sigma = " << sigma << endl;
86
87 Mat_<Tensor> tensors = Mat_<Tensor>::zeros(image.rows, image.cols);
88 Mat blur, edge, structure, color;
89 createAnisotropyTensor(tensors, image, sigma, rho, gamma,
90                       blur, edge, structure, color);
91 imwrite(argv[optind + 1], blur);
92 imwrite(argv[optind + 2], edge);
93 imwrite(argv[optind + 3], structure);
94 imwrite(argv[optind + 4], color);
95
96 /*
97  * Network only handles integer edges, so we increase the scale a bit.
98  */
99 int a;
100 int b;
101 a = 100;
102 b = beta;
103
104 /*
105  * Specify the neighbors of a pixel.
106  */
107 cout << "Creating size " << neighbors << " neighborhood." << endl;
108 Neighborhood neigh;
109 if (neighbors >= 4) {
110     neigh.add( 1, 0, b * 1.0);
111     neigh.add( 0, 1, b * 1.0);
112     neigh.add(-1, 0, b * 1.0);
113     neigh.add( 0,-1, b * 1.0);
114 }
115
116 if (neighbors >= 8) {
117     neigh.add( 1, 1, b * 1.0/sqrt(2.0));
118     neigh.add(-1, 1, b * 1.0/sqrt(2.0));
119     neigh.add( 1,-1, b * 1.0/sqrt(2.0));
120     neigh.add(-1,-1, b * 1.0/sqrt(2.0));
121 }
122
123 if (neighbors >= 16) {
124     neigh.add8(1, 2, 1.0);

```

```

125     }
126
127     if (neighbors >= 32) {
128         neigh.add8(3, 1, 1.0);
129         neigh.add8(3, 2, 1.0);
130     }
131
132     if (neighbors >= 48) {
133         neigh.add8(1, 4, 1.0);
134         neigh.add8(3, 4, 1.0);
135     }
136
137     if (neighbors >= 72) {
138         neigh.add8(1, 5, 1.0);
139         neigh.add8(2, 5, 1.0);
140         neigh.add8(3, 5, 1.0);
141     }
142
143     cout << "Neighborhood: " << endl;
144     neigh.setupAngles();
145     for (Neighborhood::iterator it = neigh.begin(); it != neigh.end(); ++it) {
146         cout << it->x << ", " << it->y << ": " << it->dt * 180 / M_PI << endl;
147     }
148
149     Mat out = image.clone();
150     restoreAnisotropicTV(image, out, tensors, neigh, a, b, p);
151
152     cout << "Writing output to " << argv[optind + 5] << endl;
153     imwrite(argv[optind + 5], out);
154
155     return 0;
156 }

```

A.2 image.hpp

Calculates and sets up graph weights based on the input image and parameters.

```

1  #pragma once
2
3  #include <vector>
4  #include <set>
5  #include <opencv2/opencv.hpp>
6  #include "graph.hpp"
7  #include "selectionrule.hpp"
8  #include "neighborhood.hpp"
9  #include "anisotropy.hpp"
10
11 void createEdgesAnisotropic(
12     FlowGraph& network,
13     Neighborhood& neigh,
14     int beta,
15     const cv::Mat_<Tensor>& tensors
16 );
17 void setupSourceSink(FlowGraph& network, cv::Mat& in, int alpha, int label, int p);
18 void setupSource(FlowGraph& network, cv::Mat& in, int alpha, int label, int p);
19 void setupSink(FlowGraph& network, cv::Mat& in, int alpha, int label, int p);
20
21 void restoreAnisotropicTV(
22     cv::Mat& in,
23     cv::Mat& out,

```

```

24         cv::Mat_<Tensor>& tensors,
25         Neighborhood& neigh,
26         int alpha, int beta, int p
27     );

```

A.3 image.cpp

```

1  #include <iostream>
2  #include <fstream>
3  #include <algorithm>
4  #include <set>
5  #include <iterator>
6  #include <cmath>
7  #include <opencv2/opencv.hpp>
8  #include "image.hpp"
9  #include "neighborhood.hpp"
10 #include "selectionrule.hpp"
11 #include "anisotropy.hpp"
12
13 using namespace std;
14 using namespace cv;
15
16 int f(int u, int v, int p) {
17     return p == 2 ? (u - v) * (u - v) : abs(u - v);
18 }
19
20 /* Fidelity energy term. */
21 int Ei(int label, int pix, int u, int p) {
22     return (f(label+1, pix, p) - f(label, pix, p)) * u;
23 }
24
25 void createEdgesAnisotropic(
26     FlowGraph& network,
27     Neighborhood& neigh,
28     int beta,
29     const Mat_<Tensor>& tensors
30 ) {
31
32     int rows = tensors.rows;
33     int cols = tensors.cols;
34     int pixels = rows * cols;
35
36     /*
37      * Add sink edges first, so that the first push in discharge
38      * will go towards the sink. The capacities are set up in
39      * setupSourceSink.
40     */
41     for (int i = 0; i < pixels; ++i) {
42         network.addEdge(i, network.getSink(), 0);
43     }
44
45     /*
46      * Create internal edges, which do not depend on the current
47      * level.
48     */
49     for (int i = 0; i < rows; ++i) {
50         for (int j = 0; j < cols; ++j) {
51             Neighborhood::iterator it;
52
53             for (it = neigh.begin(); it != neigh.end(); ++it) {
54
55                 /* Only add edges for right half of the neighborhood. */

```

```

56         if (it->x < 0)
57             continue;
58         if (it->x == 0 && it->y > 0)
59             continue;
60
61         int x = j + it->x;
62         int y = i + it->y;
63
64         if (x >= 0 && x < cols && y >= 0 && y < rows) {
65
66             Mat ee = (Mat_<double>(2, 1) << it->x, it->y);
67             Mat M1 = Mat(tensors(i, j));
68             Mat M2 = Mat(tensors(y, x));
69
70             Mat M3 = (M1 + M2) / 2.0;
71
72             double w = beta
73                 * norm(ee) * norm(ee)
74                 * determinant(M3)
75                 * it->dt
76                 / pow(ee.dot(M3 * ee), 3.0 / 2.0);
77
78             network.addDoubleEdge(
79                 i*cols + j,
80                 y*cols + x,
81                 w
82             );
83         }
84     }
85 }
86
87
88 /*
89 * Add edges from the source. Capacities are set up in
90 * setupSourceSink.
91 */
92 for (int i = 0; i < pixels; ++i) {
93     network.addEdge(network.getSource(), i, 0);
94 }
95 }
96
97 /*
98 * Change the capacities of the edges connecting the source
99 * and the sink to the rest of the network, as these edges
100 * are dependent on the current level.
101 */
102 void setupSourceSink(FlowGraph& network, Mat& in, int alpha, int label, int p) {
103     std::vector<int> s_caps(in.rows * in.cols);
104     std::vector<int> t_caps(in.rows * in.cols);
105
106     for (int j = 0; j < in.rows; ++j) {
107         for (int i = 0; i < in.cols; ++i) {
108             int e1 = Ei(label, in.at<uchar>(j, i), 1, p);
109
110             if (0 < e1) {
111                 t_caps[j*in.cols + i] += e1 - 0;
112             }
113             else {
114                 s_caps[j*in.cols + i] += 0 - e1;
115             }
116         }
117     }
118
119     for (size_t i = 0; i < s_caps.size(); ++i) {
120         network.changeCapacity(network.getSource(), i, alpha * s_caps[i]);

```



```

121     }
122
123     for (size_t i = 0; i < t_caps.size(); ++i) {
124         network.changeCapacity(i, network.getSink(), alpha * t_caps[i]);
125     }
126 }
127
128 /* Update edges connected to the source. Used by Boykov--Kolmogorov. */
129 void setupSource(FlowGraph& network, Mat& in, int alpha, int label, int p) {
130     std::vector<int> s_caps(in.rows * in.cols);
131
132     for (int j = 0; j < in.rows; ++j) {
133         for (int i = 0; i < in.cols; ++i) {
134             int elinit = Ei(255, in.at<uchar>(j, i), 1, p);
135             int e1 = Ei(label, in.at<uchar>(j, i), 1, p);
136
137             s_caps[j*in.cols + i] += max(elinit, 0) - e1;
138         }
139     }
140
141     for (size_t i = 0; i < s_caps.size(); ++i) {
142         if (s_caps[i] != 0)
143             network.changeCapacity(network.getSource(), i, alpha * s_caps[i]);
144     }
145 }
146
147 /* Update edges connected to the sink. Used by Boykov--Kolmogorov. */
148 void setupSink(FlowGraph& network, Mat& in, int alpha, int label, int p) {
149     std::vector<int> t_caps(in.rows * in.cols);
150
151     for (int j = 0; j < in.rows; ++j) {
152         for (int i = 0; i < in.cols; ++i) {
153             int elinit = Ei(255, in.at<uchar>(j, i), 1, p);
154
155             t_caps[j*in.cols + i] += max(elinit, 0);
156         }
157     }
158
159     for (size_t i = 0; i < t_caps.size(); ++i) {
160         if (t_caps[i] != 0)
161             network.changeCapacity(i, network.getSink(), alpha * t_caps[i]);
162     }
163 }
164
165 /* Restore image. */
166 void restoreAnisotropicTV(
167     Mat& in,
168     Mat& out,
169     Mat_<Tensor>& tensors,
170     Neighborhood& neigh,
171     int alpha, int beta, int p
172 ) {
173
174     int rows = in.rows;
175     int cols = in.cols;
176     int pixels = rows * cols;
177     int source = pixels;
178     int sink = pixels + 1;
179
180     FIFORule frule(pixels + 2);
181     SelectionRule& rule = frule;
182     FlowGraph network(rows * cols + 2, source, sink, rule);
183
184     createEdgesAnisotropic(network, neigh, beta, tensors);
185

```

```

186 #ifdef BOYKOV_KOLMOGOROV
187     setupSink(network, in, alpha, 255, p);
188 #endif
189
190     for (int label = 255; label >= 0; --label) {
191         cout << "Label: " << label << endl;
192
193 #ifdef PUSH_RELABEL
194         setupSourceSink(network, in, alpha, label, p);
195         network.minCutPushReLabel(source, sink);
196 #else
197         setupSource(network, in, alpha, label, p);
198         network.minCutBK(source, sink);
199 #endif
200
201         /* Use the cut to update the output image. */
202         for (int j = 0; j < rows; ++j) {
203             for (int i = 0; i < cols; ++i) {
204                 if (!network.cut[j*cols + i])
205                     out.at<uchar>(j, i) = label;
206             }
207         }
208     }
209 }

```

A.4 anisotropy.hpp

Construction of anisotropy tensor.

```

1  #pragma once
2
3  #include <opencv2/opencv.hpp>
4
5  typedef cv::Matx<double, 2, 2> Tensor;
6
7  void createAnisotropyTensor(
8      cv::Mat_<Tensor>& tensors,
9      cv::Mat& in,
10     double sigma,
11     double rho,
12     double gamma,
13     cv::Mat& blur,
14     cv::Mat& edge,
15     cv::Mat& structure,
16     cv::Mat& color
17 );
18
19 void createUniformAnisotropyTensor(cv::Mat_<Tensor>& tensors, cv::Mat& in, double gamma);

```

A.5 anisotropy.cpp

```

1  #include <opencv2/opencv.hpp>
2  #include <cstdlib>
3  #include <cstdio>
4  #include "anisotropy.hpp"
5
6  using namespace cv;
7  using namespace std;
8

```

```

9 void createAnisotropyTensor(
10     Mat_<Tensor>& tensors,
11     Mat& in,
12     double sigma,
13     double rho,
14     double gamma,
15     cv::Mat& blur,
16     cv::Mat& edge,
17     cv::Mat& structure,
18     cv::Mat& color
19 ) {
20
21     Mat grad;
22
23     /* Apply blur by sigma. */
24     GaussianBlur(in, blur, Size(0,0), sigma, 0, BORDER_REFLECT);
25
26     Mat grad_x, grad_y;
27     Mat kernel;
28
29     Point anchor = Point(-1, -1);
30
31     kernel = Mat::zeros(1, 3, CV_16S);
32
33     kernel.at<short>(0, 0) = -1;
34     kernel.at<short>(0, 1) = 0;
35     kernel.at<short>(0, 2) = 1;
36
37     /* Calculate gradient in x and y. */
38     filter2D(blur, grad_x, CV_64F, kernel, anchor, 0, BORDER_REFLECT);
39     transpose(kernel, kernel);
40     filter2D(blur, grad_y, CV_64F, kernel, anchor, 0, BORDER_REFLECT);
41
42     Mat x_sq, y_sq, xy;
43     Mat x_sqr, y_sqr, xyr;
44
45     /* Element of outer product. */
46     x_sq = grad_x.mul(grad_x);
47     y_sq = grad_y.mul(grad_y);
48     xy   = grad_x.mul(grad_y);
49
50     /* Integration scale (rho) smoothing. */
51     GaussianBlur(x_sq, x_sqr, Size(0,0), rho, 0, BORDER_REFLECT);
52     GaussianBlur(y_sq, y_sqr, Size(0,0), rho, 0, BORDER_REFLECT);
53     GaussianBlur(xy , xyr , Size(0,0), rho, 0, BORDER_REFLECT);
54
55     Mat evec, eval;
56
57     Mat h = Mat::zeros(in.size(), CV_64F);
58     Mat s = Mat::zeros(in.size(), CV_64F);
59     Mat v = Mat::zeros(in.size(), CV_64F);
60     Mat hsv;
61     vector<Mat> channels;
62
63     edge.create(in.size(), CV_64F);
64
65     for (int i = 0; i < in.rows; ++i) {
66         for (int j = 0; j < in.cols; ++j) {
67             /* Structure tensor. */
68             Tensor b = Tensor(
69                 x_sqr.at<double>(i,j),
70                 xyr.at<double>(i,j),
71                 xyr.at<double>(i,j),
72                 y_sqr.at<double>(i,j)
73             );

```

```

74
75      /* Structure tensor without rho smoothing, for the edge detector. */
76      Tensor c = Tensor(
77          x_sq.at<double>(i,j),
78          xy.at<double>(i,j),
79          xy.at<double>(i,j),
80          y_sq.at<double>(i,j)
81      );
82
83      /* Returns the eigenvectors as row vectors! */
84      eigen(b, eval, evec);
85
86      double s1 = eval.at<double>(0);
87      double s2 = eval.at<double>(1);
88
89      if (s2 > s1)
90          fprintf(stderr, "OOPS: Wrong ordering of eigenvalues\n");
91
92      double l1 = 1.0;
93      double l2 = 1.0 / (1.0 + (s1 - s2) * (s1 - s2) / (gamma*gamma));
94
95      Mat eval2 = eval.clone();
96      eval2.at<double>(0) = l1;
97      eval2.at<double>(1) = l2;
98      tensors(i, j) = Tensor(Mat(evec.t() * Mat::diag(eval2) * evec));
99
100     h.at<double>(i, j) = fmod(atan2(evec.at<double>(1), evec.at<double>(0))
101         * 180.0 / M_PI + 180.0, 180.0);
102     s.at<double>(i, j) = 0;
103     v.at<double>(i, j) = 1.0/(1.0 + l2);
104
105     /* Returns the eigenvectors as row vectors! */
106     eigen(c, eval, evec);
107
108     s1 = eval.at<double>(0);
109     s2 = eval.at<double>(1);
110
111     edge.at<double>(i, j) = s1;
112 }
113 }
114 normalize(edge, edge, 0, 255, NORM_MINMAX, CV_8U);
115
116 /* Create tensor visualization, double the size. */
117 structure.create(in.size(), CV_64F);
118 GaussianBlur(in, structure, Size(0,0), sigma, 0, BORDER_REFLECT);
119 cvtColor(structure, structure, CV_GRAY2BGR);
120 resize(structure, structure, Size(0, 0), 2, 2);
121 for (int i = 0; i < in.rows; i += 15) {
122     for (int j = 0; j < in.cols; j += 15) {
123         Tensor b = tensors(i, j);
124
125         eigen(b, eval, evec);
126
127         double s1 = eval.at<double>(0);
128         double s2 = eval.at<double>(1);
129
130         Point2f p1(evec.row(0));
131         Point2f p2(evec.row(1));
132         line(structure, 2*Point(j, i), 2*Point2f(j, i) + 20 * s2 * p1,
133             CV_RGB(255,0,0), 1.2, CV_AA);
134         line(structure, 2*Point(j, i), 2*Point2f(j, i) + 20 * s1 * p2,
135             CV_RGB(0,0,0), 1.2, CV_AA);
136     }
137 }
138

```

```

139     Mat ho, so, vo;
140     h.convertTo(ho, CV_8U);
141     normalize(s, so, 255, 255, NORM_MINMAX, CV_8U);
142     normalize(v, vo, 0, 255, NORM_MINMAX, CV_8U);
143
144     channels.push_back(ho);
145     channels.push_back(so);
146     channels.push_back(vo);
147     merge(channels, hsv);
148
149     cvtColor(hsv, color, CV_HSV2BGR);
150 }
151
152 /* Create a uniform tensor, for testing. */
153 void createUniformAnisotropyTensor(Mat_<Tensor>& tensors, Mat& in, double gamma) {
154     Mat eval, eval2;
155     for (int i = 0; i < in.rows; ++i) {
156         for (int j = 0; j < in.cols; ++j) {
157             Tensor b = Tensor(1, 1, 1, 1);
158
159             /* Returns the eigenvectors as row vectors! */
160             eigen(b, eval, eval2);
161
162             double l1 = 1.0 / gamma;
163             double l2 = 1.0;
164
165             Mat eval2 = eval.clone();
166             eval2.at<double>(0) = l1;
167             eval2.at<double>(1) = l2;
168             tensors(i, j) = Tensor(Mat::diag(eval2));
169         }
170     }
171 }

```

A.6 graph.hpp

Graph class with the maximum flow algorithms.

```

1  #pragma once
2
3  #include <iostream>
4  #include <queue>
5  #include <vector>
6  #include <list>
7  #include <set>
8  #include "selectionrule.hpp"
9
10 enum Color { FREE = 0, SOURCE, SINK };
11
12 class Edge {
13 private:
14
15 public:
16     int from, to;
17     int cap;
18     int flow;
19     int index;
20
21     Edge(int f, int t, int c, int i) :
22         from(f), to(t), cap(c), flow(0), index(i) {}
23 };
24

```

```

25 class Vertex {
26 private:
27
28 public:
29     std::vector<Edge> e;
30     Color c;
31     bool active;
32     Edge *p;
33     int height;
34     int excess;
35     int si;
36     int ti;
37
38     Vertex(Color c, bool a) :
39         c(c), active(a), p(NULL), height(0), excess(0), si(0), ti(0) {}
40
41     Vertex() :
42         c(FREE), active(false), p(NULL), height(0), excess(0), si(0), ti(0) {}
43 };
44
45 class FlowGraph {
46 private:
47     int N;
48     int source, sink;
49     std::vector<Vertex > G;
50     std::vector<int> count;
51     SelectionRule& rule;
52
53     /* BK stuff */
54     std::queue<int> bkq;
55     std::queue<int> orphans;
56
57     int lastGrowVertex;
58     size_t lastIndex;
59 public:
60     std::vector<char> cut;
61
62     FlowGraph(int N, int source, int sink, SelectionRule& rule) :
63         N(N),
64         source(source),
65         sink(sink),
66         G(N),
67         count(N+1),
68         rule(rule),
69         lastGrowVertex(-1),
70         cut(N) {
71
72     #ifdef BOYKOV_KOLMOGOROV
73         bkq.push(source);
74         bkq.push(sink);
75         G[source].c = SOURCE;
76         G[sink].c = SINK;
77         G[source].active = true;
78         G[sink].active = true;
79     #endif
80     }
81
82     int getSource() const { return source; }
83     int getSink() const { return sink; }
84     void addEdge(int from, int to, int cap);
85     void addDoubleEdge(int from, int to, int cap);
86     void changeCapacity(int from, int index, int cap);
87     void resetFlow();
88     void resetHeights();
89

```

```

90     void push(Edge &e);
91     void push(Edge &e, int f);
92     void relabel(int u);
93     void gap(int h);
94     void discharge(int u);
95     void minCutPushRelabel(int source, int sink);
96
97     void minCutBK(int source, int sink);
98     int augment(Edge *e);
99     int treeCap(const Edge& e, Color col) const;
100    int treeOrigin(int u, int &len) const;
101    void adopt();
102    Edge *grow();
103 };

```

A.7 graph.cpp

```

1  #include <iostream>
2  #include <queue>
3  #include <stack>
4  #include <vector>
5  #include <algorithm>
6  #include <cassert>
7  #include <cstdlib>
8
9  #include "graph.hpp"
10
11 using namespace std;
12
13 /* Add an edge from one vertex to another. */
14 void FlowGraph::addEdge(int from, int to, int cap) {
15     G[from].e.push_back(Edge(from, to, cap, G[to].e.size()));
16     if (from == to) G[from].e.back().index++;
17     int index = G[from].e.size() - 1;
18     G[to].e.push_back(Edge(to, from, 0, index));
19
20     if (from == source)
21         G[to].si = index;
22
23     if (to == sink)
24         G[from].ti = index;
25 }
26
27 /*
28  * Add an edge and at the same time an antiparallel edge
29  * with the same capacity.
30 */
31 void FlowGraph::addDoubleEdge(int from, int to, int cap) {
32     G[from].e.push_back(Edge(from, to, cap, G[to].e.size()));
33     G[to].e.push_back(Edge(to, from, cap, G[from].e.size() - 1));
34 }
35
36 #ifndef PUSH_RELABEL
37
38 /*
39  * Change the capacity of an edge. Need the from-vertex and
40  * the index of the edge in its edge list (returned from addEdge.
41 */
42 void FlowGraph::changeCapacity(int from, int to, int cap) {
43     int index;
44     if (from == source) {
45         index = G[to].si;

```

```

46     } else if (to == sink) {
47         index = G[from].ti;
48     } else {
49         exit(1);
50     }
51
52     int diff = G[from].e[index].flow - cap;
53
54     G[from].e[index].cap = cap;
55
56     /* Check if we need to reduce the flow. */
57     if (diff > 0) {
58         G[from].excess += diff;
59         G[to].excess -= diff;
60         G[from].e[index].flow = cap;
61         G[to].e[G[from].e[index].index].flow = -cap;
62         rule.add(from, G[from].height, G[from].excess);
63     }
64 }
65
66 #else
67
68 /*
69  * Change the capacity of an edge. Need the from-vertex and
70  * the index of the edge in its edge list (returned from addEdge)
71  */
72 void FlowGraph::changeCapacity(int from, int to, int cap) {
73     int index;
74     if (from == source) {
75         index = G[to].si;
76     } else if (to == sink) {
77         index = G[from].ti;
78     } else {
79         exit(1);
80     }
81
82     /* Nodes in the S set can not send any more flow anyways. */
83     if (from == source && G[to].c == SOURCE)
84         return;
85
86     G[from].e[index].cap = cap;
87
88     if (from != source)
89         return;
90
91     /* Push flow along two-edged path. */
92     int si = G[to].si;
93     int ti = G[to].ti;
94
95     Edge *sv, *vt;
96     sv = &G[source].e[si];
97     vt = &G[to].e[ti];
98
99     int rs = sv->cap - sv->flow;
100    int rt = vt->cap - vt->flow;
101
102    if (rs > 0 && rt > 0) {
103        int m = min(rs, rt);
104        push(*sv, m);
105        push(*vt, m);
106
107        if (m == rs && G[to].p == sv) {
108            G[to].p = NULL;
109            orphans.push(to);
110        }

```



```

111         if (m == rt && G[to].p == vt) {
112             G[to].p = NULL;
113             orphans.push(to);
114         }
115     }
116 }
117
118 /* Activate vertices if the new edge was not saturated. */
119 if (G[from].e[si].flow != G[from].e[si].cap) {
120     if (!G[from].active) {
121         bkq.push(from);
122         G[from].active = true;
123     }
124     if (!G[to].active) {
125         bkq.push(to);
126         G[to].active = true;
127     }
128 }
129 }
130
131 #endif
132
133 /* Reset all flow and excess. */
134 void FlowGraph::resetFlow() {
135     for (size_t i = 0; i < G.size(); ++i) {
136         for (size_t j = 0; j < G[i].e.size(); ++j) {
137             G[i].e[j].flow = 0;
138         }
139         G[i].excess = 0;
140     }
141 }
142
143 /* Reset all distance labels. */
144 void FlowGraph::resetHeights() {
145     for (size_t i = 0; i < G.size(); ++i) {
146         G[i].height = 0;
147     }
148     fill(count.begin(), count.end(), 0);
149 }
150
151 /* Push along an edge. */
152 void FlowGraph::push(Edge &e) {
153     int flow = min(e.cap - e.flow, G[e.from].excess);
154     G[e.from].excess -= flow;
155     G[e.to].excess += flow;
156     e.flow += flow;
157     G[e.to].e[e.index].flow -= flow;
158
159     rule.add(e.to, G[e.to].height, G[e.to].excess);
160 }
161
162 /* Push given flow along an edge. */
163 void FlowGraph::push(Edge &e, int f) {
164     e.flow += f;
165     G[e.to].e[e.index].flow -= f;
166 }
167
168 /* Relabel a vertex. */
169 void FlowGraph::relabel(int u) {
170     count[G[u].height]--;
171     G[u].height = 2*N;
172
173     for (size_t i = 0; i < G[u].e.size(); ++i) {
174         if (G[u].e[i].cap > G[u].e[i].flow) {
175             G[u].height = min(G[u].height, G[G[u].e[i].to].height + 1);

```

```

176         }
177     }
178
179     if (G[u].height >= N) {
180         G[u].height = N;
181     }
182     else {
183         count[G[u].height]++;
184         rule.add(u, G[u].height, G[u].excess);
185     }
186 }
187
188 /* Relabel all vertices over the gap h to label N. */
189 void FlowGraph::gap(int h) {
190     for (size_t i = 0; i < G.size(); ++i) {
191         if (G[i].height < h) continue;
192         if (G[i].height >= N) continue;
193
194         rule.deactivate(i);
195
196         count[G[i].height]--;
197         G[i].height = N;
198     }
199
200     rule.gap(h);
201 }
202
203 /* Discharge a vertex. */
204 void FlowGraph::discharge(int u) {
205     size_t i;
206     for (i = 0; i < G[u].e.size() && G[u].excess > 0; ++i) {
207         if (G[u].e[i].cap > G[u].e[i].flow
208             && G[u].height == G[G[u].e[i].to].height + 1) {
209             push(G[u].e[i]);
210         }
211     }
212
213     if (G[u].excess > 0) {
214         /* Check if a gap will appear. */
215         if (count[G[u].height] == 1)
216             gap(G[u].height);
217         else
218             relabel(u);
219     }
220 }
221
222 /* Run the push-relabel algorithm to find the min-cut. */
223 void FlowGraph::minCutPushRelabel(int source, int sink) {
224     G[source].height = N;
225
226     rule.activate(source);
227     rule.activate(sink);
228
229     for (size_t i = 0; i < G[source].e.size(); ++i) {
230         G[source].excess = G[source].e[i].cap;
231         push(G[source].e[i]);
232     }
233     G[source].excess = 0;
234
235     int c = 0;
236     /* Loop over active nodes using selection rule. */
237     while (!rule.empty()) {
238         c++;
239         int u = rule.next();
240         discharge(u);

```

```

241     }
242
243     /* Output the cut based on vertex heights. */
244     for (size_t i = 0; i < cut.size(); ++i) {
245         cut[i] = G[i].height >= N;
246     }
247 }
248
249 /* The capacity of the edge, in the direction given by the tree. */
250 int FlowGraph::treeCap(const Edge& e, Color col) const {
251     if (col == SOURCE)
252         return e.cap - e.flow;
253     else if (col == SINK)
254         return G[e.to].e[e.index].cap - G[e.to].e[e.index].flow;
255     else
256         return 0;
257 }
258
259 /* Try to grow the trees S and T from the active vertices. */
260 Edge *FlowGraph::grow() {
261     while (!bkq.empty()) {
262         int p = bkq.front();
263         if (!G[p].active) {
264             bkq.pop();
265             continue;
266         }
267
268         size_t i = 0;
269
270         /*
271          * If we're growing from the same vertex as before,
272          * continue with the same index, if not restart at 0.
273         */
274         if (lastGrowVertex == p)
275             i = lastIndex;
276
277         lastGrowVertex = p;
278
279         for (; i < G[p].e.size(); ++i) {
280             lastIndex = i;
281             Edge *e = &G[p].e[i];
282             int q = e->to;
283
284             if (G[p].c == G[q].c)
285                 continue;
286
287             if (treeCap(*e, G[p].c) <= 0)
288                 continue;
289
290             if (G[q].c == FREE) {
291                 /* Found a free vertex, add it. */
292                 G[q].c = G[p].c;
293
294                 int len;
295                 if (G[p].c == SOURCE) {
296                     G[q].p = e;
297                     assert(treeOrigin(p, len) == source);
298                 } else if (G[p].c == SINK) {
299                     G[q].p = &G[q].e[e->index];
300                     assert(treeOrigin(p, len) == sink);
301                 } else {
302                     cout << G[p].c << endl;
303                     exit(1);
304                 }
305                 (void)len;

```

```

306
307         G[q].active = 1;
308         bkq.push(q);
309     }
310     else if (G[q].c != G[p].c) {
311         /* The trees meet! */
312         if (G[p].c == SOURCE) {
313             return e;
314         }
315         else if (G[p].c == SINK) {
316             return &G[q].e[e->index];
317         }
318         else {
319             exit(1);
320         }
321     }
322     return NULL;
323 }
324 }
325
326     bkq.pop();
327     G[p].active = 0;
328 }
329
330 /* Path is empty */
331 return NULL;
332 }
333
334 /* Augment along path given by the edge e, and implicitly by the trees. */
335 int FlowGraph::augment(Edge* e) {
336     int m = e->cap - e->flow;
337
338     /* Find maximum flow we can send. */
339     Edge *cur = e;
340     while (cur != NULL) {
341         m = min(m, cur->cap - cur->flow);
342         cur = G[cur->from].p;
343     }
344
345     cur = e;
346     while (cur != NULL) {
347         m = min(m, cur->cap - cur->flow);
348         cur = G[cur->to].p;
349     }
350
351     cur = e;
352     bool back = true;
353     int len = 0;
354     /* Loop through path and update flow. */
355     while (cur != NULL) {
356         /* If saturated, we must orphanize. */
357         if (cur->cap - cur->flow == m) {
358             int u = cur->from;
359             int v = cur->to;
360
361             if (G[u].c == SOURCE && G[v].c == SOURCE) {
362                 if (v != source && v != sink) {
363                     orphans.push(v);
364                     G[v].p = NULL;
365                 }
366             }
367             if (G[u].c == SINK && G[v].c == SINK) {
368                 if (u != source && u != sink) {
369                     orphans.push(u);
370                     G[u].p = NULL;

```

```

371         }
372     }
373 }
374 len++;
375 push(*cur, m);
376
377 /*
378  * If we reach the source, we must start again
379  * in e, and go towards the sink.
380  */
381 if (back) {
382     cur = G[cur->from].p;
383     if (cur == NULL) {
384         back = false;
385         cur = G[e->to].p;
386     }
387 } else {
388     cur = G[cur->to].p;
389 }
390 }
391 return len;
392 }
393
394 /* Find the origin of vertex u. */
395 int FlowGraph::treeOrigin(int u, int &len) const {
396     int cur = u;
397     len = 0;
398
399     if (G[cur].c == SOURCE) {
400         while (G[cur].p != NULL) {
401             cur = G[cur].p->from;
402             len++;
403         }
404     } else if (G[cur].c == SINK) {
405         while (G[cur].p != NULL) {
406             cur = G[cur].p->to;
407             len++;
408         }
409     } else {
410         exit(1);
411     }
412
413     return cur;
414 }
415
416 /* Adopt orphans. */
417 void FlowGraph::adopt() {
418     while (orphans.size() > 0) {
419         int u = orphans.front();
420         orphans.pop();
421
422         assert(G[u].c != FREE);
423
424         int minlen = 1000000000;
425         int minidx = -1;
426         /* Aim to find parent close to the root of the tree. */
427         for (size_t i = 0; i < G[u].e.size(); ++i) {
428             int v = G[u].e[i].to;
429
430             if (G[u].c != G[v].c)
431                 continue;
432
433             if (treeCap(G[v].e[G[u].e[i].index], G[u].c) <= 0)
434                 continue;
435

```

```

436     int len;
437     int origin = treeOrigin(v, len);
438     if (origin != source && origin != sink)
439         continue;
440
441     if (len < minlen) {
442         minlen = len;
443         minidx = i;
444     }
445     if (minlen <= 2) break;
446     /* Found a possible parent */
447 }
448
449 bool found = false;
450 if (minidx != -1) {
451     int i = minidx;
452     int v = G[u].e[i].to;
453     int len;
454     int origin = treeOrigin(v, len);
455
456     if (origin == source) {
457         G[u].p = &G[v].e[G[u].e[i].index];
458         found = true;
459     } else if (origin == sink) {
460         G[u].p = &G[u].e[i];
461         found = true;
462     } else {
463         exit(1);
464     }
465 }
466
467 /* If not found, free vertex, and orphanize possible children. */
468 if (!found) {
469     for (size_t i = 0; i < G[u].e.size(); ++i) {
470         int v = G[u].e[i].to;
471
472         if (G[u].c != G[v].c)
473             continue;
474
475         if (treeCap(G[v].e[G[u].e[i].index], G[v].c) > 0) {
476             G[v].active = true;
477             bkq.push(v);
478         }
479
480         if (v == source || v == sink)
481             continue;
482
483         if (G[v].p
484             && (G[v].p->to == u
485                || G[v].p->from == u)) {
486             orphans.push(v);
487             G[v].p = NULL;
488         }
489     }
490
491     G[u].c = FREE;
492
493     G[u].active = false;
494     /* We might still have u in the queue */
495 }
496 }
497 }
498
499 void FlowGraph::minCutBK(int source, int sink) {
500     lastGrowVertex = -1;

```

```

501     adopt();
502
503     int numpaths = 0;
504     double totlen = 0;
505     while (true) {
506         Edge *e;
507         e = grow();
508
509         if (e == NULL) {
510             /* Empty path. */
511             break;
512         }
513
514         totlen += augment(e);
515         numpaths++;
516         adopt();
517     }
518
519     cout << "Avg length: " << double(totlen) / double(numpaths) << endl;
520
521     int size1 = 0, size2 = 0;
522     for (size_t i = 0; i < cut.size(); ++i) {
523         if (G[i].c == SOURCE) size1++;
524         else if (G[i].c == SINK) size2++;
525         cut[i] = G[i].c == SOURCE;
526     }
527     cout << "Inbetweeners: " << cut.size() - size1 - size2 << endl;
528 }

```

A.8 selectionrule.hpp

Highest level and FIFO selection rules for the push-relabel algorithm.

```

1  #pragma once
2
3  #include <iostream>
4  #include <queue>
5  #include <vector>
6  #include <exception>
7
8  class EmptyQueueException : public std::exception {
9      virtual const char* what() const throw()
10     {
11         return "Empty Queue Exception";
12     }
13 };
14
15 class SelectionRule {
16
17     protected:
18         int N;
19
20     private:
21         std::vector<char> active;
22
23     public:
24         virtual int next(void) = 0;
25         virtual void add(int u, int height, int excess) = 0;
26         virtual bool empty(void) = 0;
27         virtual void gap(int h) = 0;
28
29         void activate(int u) { active[u] = 1; }

```

```

30     void deactivate(int u) { active[u] = 0; }
31     bool isActive(int u) { return active[u]; }
32
33     SelectionRule(int N) : N(N), active(N) {}
34
35     virtual ~SelectionRule() {}
36 };
37
38 class HighestLevelRule : virtual public SelectionRule {
39 private:
40     int highest;
41     std::vector<std::queue<int> > hq;
42
43     void updateHighest(void);
44
45 public:
46     virtual int next(void);
47     virtual void add(int u, int height, int excess);
48     virtual bool empty(void);
49     virtual void gap(int h);
50
51     HighestLevelRule(int N) : SelectionRule(N), highest(-1), hq(N) {}
52 };
53
54 class FIFORule : virtual public SelectionRule {
55 private:
56     std::queue<int> q;
57
58 public:
59     virtual int next(void);
60     virtual void add(int u, int height, int excess);
61     virtual bool empty(void);
62     virtual void gap(int h);
63
64     FIFORule(int N) : SelectionRule(N) {}
65 };
66
67

```

A.9 selectionrule.cpp

```

1  #include "selectionrule.hpp"
2
3  using namespace std;
4
5  void HighestLevelRule::gap(int h) {
6      highest = h - 1;
7      updateHighest();
8  }
9
10 void HighestLevelRule::updateHighest(void) {
11     int s = highest;
12     highest = -1;
13
14     for (int i = s; i >= 0; --i) {
15         if (hq[i].size() > 0) {
16             highest = i;
17             break;
18         }
19     }
20 }
21

```



```

22 bool HighestLevelRule::empty(void) {
23     return highest < 0;
24 }
25
26 int HighestLevelRule::next(void) {
27     if (empty()) throw EmptyQueueException();
28
29     int u = hq[highest].front();
30     hq[highest].pop();
31     deactivate(u);
32
33     updateHighest();
34
35     return u;
36 }
37
38 void HighestLevelRule::add(int u, int height, int excess) {
39     if (isActive(u)) return;
40     if (height >= N) return;
41     if (excess == 0) return;
42
43     hq[height].push(u);
44     if (height > highest) highest = height;
45 }
46
47 void FIFORule::gap(int h) {
48 }
49
50 int FIFORule::next(void) {
51     int u;
52
53     if (!q.empty()) {
54         u = q.front();
55         q.pop();
56         deactivate(u);
57     }
58     else {
59         throw EmptyQueueException();
60     }
61
62     return u;
63 }
64
65 void FIFORule::add(int u, int height, int excess) {
66     if (isActive(u)) return;
67     if (height >= N) return;
68     if (excess == 0) return;
69
70     activate(u);
71     q.push(u);
72 }
73
74 bool FIFORule::empty(void) {
75     return q.empty();
76 }

```

A.10 neighborhood.hpp

Neighborhood class that also calculates angular differences.

```

1  #pragma once
2

```

```

3  #include <iostream>
4  #include <vector>
5  #include <cmath>
6
7  class Coord {
8  public:
9      int x;
10     int y;
11     int w;
12     mutable double dt;
13
14     Coord() : x(0), y(0), w(0), dt(0.0) {}
15     Coord(int x, int y, int w) : x(x), y(y), w(w), dt(0.0) {}
16
17     double angle() {
18         return atan2(y, x);
19     }
20 };
21
22 class CoordCompare {
23 public:
24     bool operator()(const Coord& lhs, const Coord& rhs) {
25         double p1 = copysign(1.0 - lhs.x/(fabs(lhs.x) + fabs(lhs.y)), lhs.y);
26         double p2 = copysign(1.0 - rhs.x/(fabs(rhs.x) + fabs(rhs.y)), rhs.y);
27
28         return p1 < p2;
29     }
30 };
31
32 class Neighborhood {
33 private:
34     std::set<Coord, CoordCompare> v;
35
36 public:
37     void add(int x, int y, int w) {
38         v.insert(Coord(x, y, w));
39     }
40
41     void add8(int x, int y, int w) {
42         v.insert(Coord(x, y, w));
43         v.insert(Coord(-x, y, w));
44         v.insert(Coord(x, -y, w));
45         v.insert(Coord(-x, -y, w));
46         v.insert(Coord(y, x, w));
47         v.insert(Coord(-y, x, w));
48         v.insert(Coord(y, -x, w));
49         v.insert(Coord(-y, -x, w));
50     }
51
52     std::set<Coord, CoordCompare>::iterator begin() { return v.begin(); }
53     std::set<Coord, CoordCompare>::iterator end() { return v.end(); }
54     std::set<Coord, CoordCompare>::reverse_iterator rbegin() { return v.rbegin(); }
55     std::set<Coord, CoordCompare>::reverse_iterator rend() { return v.rend(); }
56     size_t size() { return v.size(); }
57
58     typedef std::set<Coord, CoordCompare>::iterator iterator;
59
60     void setupAngles() {
61         for (Neighborhood::iterator it = this->begin();
62              it != this->end();
63              ++it) {
64             Coord prev;
65             Coord next;
66
67             if (it == this->begin()) {

```

```
68         prev = *(this->rbegin());
69     } else {
70         prev = *(--it);
71         it++;
72     }
73
74     if (++it == this->end()) {
75         next = *(this->begin());
76     } else {
77         next = *it;
78     }
79     it--;
80
81     it->dt = next.angle() - prev.angle();
82     while (it->dt > 2.0 * M_PI)
83         it->dt -= 2.0 * M_PI;
84     while (it->dt < 0.0)
85         it->dt += 2.0 * M_PI;
86     it->dt /= 2.0;
87     }
88 }
89 };
```