Total variation-based image restoration for a single-sensor digital colour camera

by

James D. Trotter

Thesis

for the degree of

Master of Science
in
Applied Mathematics
Computational Science

Faculty of Mathematics and Natural Sciences
University of Oslo

December 2012

Abstract

Modern digital colour cameras are faced with a number of challenges in producing high-quality images, including noisy sensor measurements and chromatic aberration due to dispersion in the optics. In addition, most digital colour cameras use a single sensor combined with a set of colour filters to capture red, green and blue wavelengths of light at different spatial locations in a mosaic-like pattern. Hence, some form of interpolation, often called demosaicking, is required to produce a full colour image. These image restoration tasks are formulated as ill-posed inverse problems and solved through a regularisation inspired by the total variation (TV) denoising algorithm of Rudin, Osher and Fatemi. This leads to convex variational problems and edge-preserving image restorations. To solve these problems, an efficient primal-dual algorithm from convex analysis is adopted. In addition to some standard image restoration problems, we apply these methods to chromatic aberration and demosaicking. A TV-based demosaicking model is developed based on a decomposition of the image into luminance and chrominance components which are then regularised separately. The proposed method demonstrates improved results for demosaicking a set of standard test images.

Acknowledgements

First of all, I am indebted to my supervisors, Knut Mørken, Lars Aurdal, Jan Tore Korneliussen, and Per Ove Husøy. For sharing your knowledge, offering advice, and for every one of our helpful discussions, thank you.

So many others have also offered their help along the way. Among them I would like to thank Prof. Keigo Hirakawa for taking the time to listen and for offering his valuable insight and perspective. I would also like to thank Siddhartha Mishra for help, inspiration and for being a great teacher.

I am grateful to my fellow students who have never turned me down when I have asked questions. In particular, I want to mention Magne Nordaas, Mikkel Brudvik Sanderud, Andreas Stusvik Haug, Håkon Østerbø, and Kristina Dahl.

I am thankful to those who have otherwise helped me with this thesis by reading through it and giving corrections and advice.

Last, but certainly not least. Thank you, Une.

James David Trotter,

December 2012

Contents

List of Symbols

$\mathbb{R}$	The set of real numbers
$\bar{\mathbb{R}}$	The set of extended real numbers $\mathbb{R} \cup \{+\infty, -\infty\}$
$\mathbb{R}^n$	n-dimensional Euclidean space
$\lvert\cdot\rvert$	The Euclidean norm, $\lvert x\rvert = \sqrt{\sum_{k=1}^n x_k^2}$
$x \cdot y$	The Euclidean inner product, $x \cdot y = \sum_{k=1}^n x_k y_k$
χ_Ω	The characteristic function of the set $\Omega \subset \mathbb{R}^n$
δ_C	The indicator function of the convex set $C \subset X$
f_{x_j}	The j-th partial derivative of f, $\partial f / \partial x_j$
∇f	The gradient of f, $\nabla f(x) = (f_{x_1}(x), \ldots, f_{x_n}(x))$
$\nabla^2 f$	The Hessian of f, $\nabla^2 f(x) = \left(f_{x_i x_j}(x)\right)_{i,j=1}^n$
Δf	The Laplacian of f, $\Delta f(x) = \sum_{i=1}^n \left(f_{x_i x_i}(x)\right)^2$
$\operatorname{div} \boldsymbol{f}$	The divergence of $\boldsymbol{f}(x) = \left(f^1(x), \ldots, f^n(x)\right)$, $\operatorname{div} \boldsymbol{f}(x) = \sum_{i=1}^n f_{x_i}^i(x)$
$TV(f)$	The total variation of f, $TV(f) = \int_\Omega \lvert\nabla f(x)\rvert \, dx$
X	Usually denotes a topological vector space
X^*	The dual space of continuous linear functionals on X
$\lVert\cdot\rVert_X$	Norm on the space X
$(\cdot,\cdot)_X$	Inner product on the space X
$\langle\cdot,\cdot\rangle_{X,V}$	Pairing of two topological vector spaces X and V
$C^k(\Omega)$	Space of k-times continuously differentiable functions
$C^\infty(\Omega)$	Space of infinitely differentiable, or smooth functions
$C_c(\Omega)$	Space of continuous functions with compact support
$L^p(\Omega)$	Space of p-integrable functions, $1 \leq p < +\infty$, with respect to the Lebesgue measure
$L^\infty(\Omega)$	Space of essentially bounded functions, i.e., bounded except on a set of measure zero with respect to the Lebesgue measure
$W^{k,p}(\Omega)$	Sobolev space of k-times weakly differentiable functions in $L^p(\Omega)$
$BV(\Omega)$	Space of functions of bounded variation

CHAPTER 1

Introduction

The use of computers for restoring degraded images was pioneered in the field of astronomical imaging in the 1960s and 70s. The Nobel Prize-winning invention of the charge-coupled device (CCD) in 1969 by Willard Boyle and George E. Smith was soon adopted by astronomers, thereby producing vast amounts of images. This data was obtained in an environment that is far from ideal due to, for example, atmospheric disturbances and ever-present noise. To extract meaningful information from these observations, the original, undegraded data had to be recovered. Thus, the CCD technology became a driving force in the development of digital image restoration techniques.

Recording an image relies on the imperfect measurements of some sensor device, such as a CCD or complementary metal-oxide-semiconductor (CMOS) image sensor. In addition to the noise inherent in the sensing process, there is blur and distortion due to movement, aberrations due to atmospheric effects, and loss of information during the transfer of data. Every imaging application has to deal with various such effects that diminish the quality of the images which it produces. Therefore, image restoration is regarded as a fundamental operation in image processing and computer vision where it is responsible for providing good-quality data, either for viewing or further processing.

1.1 Image restoration problems

In this thesis, we consider some image degradations that arise in most digital cameras. These degradations include, of course, noise, but also blurring effects due to the camera lens. Moreover, we address two additional problems that are specific to digital colour cameras. The first is known as *chromatic aberration* and occurs because of dispersion in the optics of an imaging device. That is, the refractive index of the camera lens material varies with the wavelength of light. This results in a relative shift of the different colours and the appearance of incorrect, coloured fringes surrounding objects (see Figure 1.1).

The second problem arises when employing a *colour filter array (CFA)* and a single image sensor to obtain three-channel colour images in a cost-effective way. This arrangement assigns every pixel of the image to a single colour

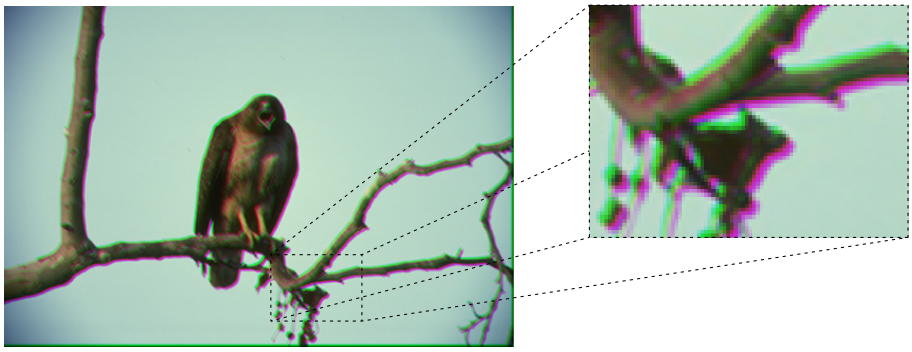

Figure 1.1: *False colour fringes. The figure shows an example of how chromatic aberration can lead to* green–purple fringing *in a colour image. Since different wavelengths of light travel along different paths in the optical system of an imaging device, the colour channels are shifted relative to one another. This results in false colours being recorded at the boundaries between bright and darker regions.*

channel in a mosaic-like pattern. Therefore, a process known as *demosaicking* (or CFA interpolation) is required to estimate the missing colour values at every pixel. This process is illustrated in Figure 1.2. Demosaicking is a very special kind of interpolation problem where particular care must be taken to avoid artefacts in the final image. Although it is an intentional result of the camera design, we will see that it is convenient to regard the effects of the CFA as a kind of degradation of the original image. These image degradations are discussed in further detail in Chapter 2.

For our purposes, image restoration can be considered in terms of the three following problems:

- Noise removal (or denoising)
- Blur removal (or deblur)
- Inpainting

The problem of sensor noise obviously belongs to the first category, whereas blur and chromatic aberration fall under the second. Indeed, we explain in Chapter 2 that chromatic aberration can be considered as a kind of "blur" where the colour channels of the image are blurred by varying amounts.

The third problem, inpainting, refers to restoring an image when some parts of it are missing. This may be due to, for example, bits of the image being lost during the transfer of data, or scratches in an old photographic film. Roughly speaking, an image is defined on some domain Ω, and the image data happens to be missing within some subdomain $\Gamma \subset \Omega$, sometimes called the *hole* or the *inpainting domain*, see Figure 1.3. The inpainting problem is to fill in the missing image values corresponding to the inpainting domain Γ in such a way that the restored image is meaningful and visually pleasing. We argue that demosaicking can also be regarded as an inpainting problem where the values of some colour channels are missing at every pixel.

The above three problems belong to the general class of *inverse problems*. An inverse problem is defined (rather vaguely) as the inverse operation to some forward problem. In the case of image restoration, the forward problem

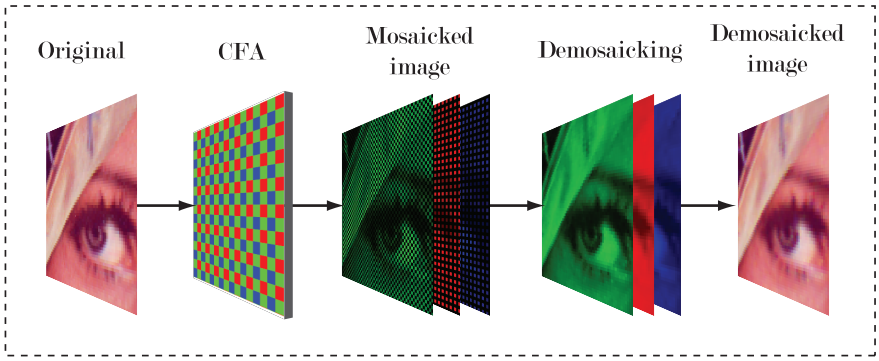

Figure 1.2: *Demosaicking. In a single-sensor digital colour camera a colour filter array (CFA) assigns to every pixel only one colour channel. The colour channels in the resulting mosaicked image are effectively subsampled versions of the colour channels in the original image. Therefore, it is necessary to perform an interpolation step, known as demosaicking, to estimate the missing pixel values in each colour channel.*

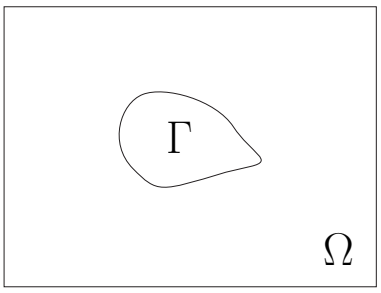

(a) Inpainting domain

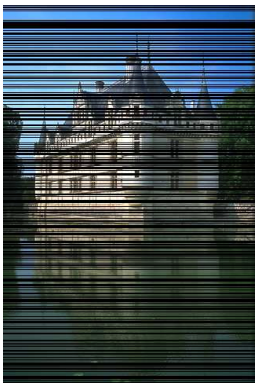

(b) An inpainting problem

Figure 1.3: *Inpainting. Figure (a) shows an abstract representation of the domain Ω on which the image is defined and the inpainting domain $\Gamma \subset \Omega$ where there is missing information. The process of inpainting refers to filling in the missing parts of the image, i.e., those that lie within the inpainting domain Γ. Figure (b) is an example of an image for which some information is missing. More precisely, about 60% of the lines in the image are missing. The inpainting domain then corresponds to the missing lines, and an inpainting is required to fill in the missing data.*

is to describe the degraded image in terms of the original. Naturally, the inverse is then to recover the original, knowing only the degraded image. The theory surrounding inverse problems is extensive and has applications to, e.g., parameter estimation, geophysics, inverse scattering theory, and, of course, image restoration, see, e.g., [46], [74].

There are a large number of methods for solving inverse problems, including many that apply specifically to image restoration. We mention a few here:

- Filtering (see, e.g., [35])
- Stochastic modelling (e.g., [47])
- Wavelet-based methods and multi-resolution analysis (e.g., [27])
- PDE-based methods (e.g., [3])
- Variational methods and regularisation (e.g., [5])

These approaches are also related in various ways, see [19]. Here we choose to focus on variational methods, and particularly those based on regularisation. This is mainly because variational methods provide a general and flexible framework for dealing with the variety of image restoration problems that we are interested in. We proceed to describe these methods in a general context. Then we present the particular models that are to be the central objects of study in this thesis.

1.2 Variational methods for image restoration

Our approach is to study denoising, deblurring and inpainting within a single variational framework. Moreover, we aim to adapt this approach to deal with chromatic aberration and demosaicking. Variational methods are well known in the image processing literature and have been applied to diverse problems such as segmentation, motion estimation, optical flow and multiple view reconstruction, to name a few, see, e.g., [17]. These methods are based on minimising some cost functional over a set of feasible solutions. A classical example from image restoration is the Wiener filter (e.g., [35]) which is designed to give the denoised or deblurred solution with the least mean square error compared to the original image. Another example is the Mumford–Shah functional for image segmentation [56].

Roughly speaking, the cost functional encodes information about which images are deemed to be a superior restoration compared to others. Although the choice of cost functional can be quite flexible, we limit our discussion to a specific family of cost functionals that lead to so-called *total variation-based methods*. These methods began with the image denoising algorithm in [65], which has become known as the *Rudin–Osher–Fatemi (ROF) model*, named after the authors. The ROF model is based on minimising the *total variation (TV)*, which is another name for the L^1-norm of the gradient. This approach has been quite successful for some image denoising and deblurring problems (see, e.g., [6]), mainly because the edges of the original image are restored sharply. The main characteristic of the model is that it penalises small-scale variations that are associated with noise while still retaining the edges of the image. This is in contrast to a number of other variational methods that tend to blur the edges of the image.

Other variational methods, e.g., for inpainting [72], have also incorporated TV for regularisation, and they have indeed proved to be successful for image restoration problems, at least for the following two reasons. Firstly, they are edge-preserving in the sense that the edges of the original image can often be recovered sharply, even in the presence of noise and blur. Secondly, the associated variational problems are convex, which implies that a globally optimal solution can be computed with reasonable efficiency.

The remainder of this section is devoted to introducing the concept of regularisation, and to present the ROF model and some related total variation-based methods. First, we acquaint ourselves with some notation. In the context of variational image processing, it is useful to consider an image as a function of continuous spatial variables, rather than a discrete set of pixel values. The number of colour channels remains, however, fixed. In the following we will use bold notation $\boldsymbol{f}$ to denote a function that takes values in $\mathbb{R}^n$ for $n \geq 2$, and the components of $\boldsymbol{f}$ are $\boldsymbol{f}(x) = \big(f^1(x), \ldots, f^n(x)\big)$.

Let Ω be an open and bounded subset of $\mathbb{R}^m$. We will refer to Ω as the image domain. An image is then represented by a function $\boldsymbol{f} : \Omega \to \mathbb{R}^n$. Usually, we take $m = 2$ and let Ω be some rectangular domain, for example, the unit square $\Omega = (0, 1)^2$. For a digital colour image, there are typically $n = 3$ colour channels, each representing the intensity of red, green or blue wavelengths of light. We note, however, that nothing in particular is gained from restricting ourselves to a specific choice of m and n, and in some cases values other than those we have mentioned become interesting.

Let X and Y be two vector spaces whose elements are functions from Ω to $\mathbb{R}^n$. Suppose that an image $\boldsymbol{u} \in X$ is subjected to the following transformation, which we call the *forward image model*

$$\boldsymbol{f} = A\boldsymbol{u} + \boldsymbol{\eta}. \tag{1.1}$$

Here $\boldsymbol{\eta} \in Y$ denotes the noise represented by some stochastic process, and $A : X \to Y$ is a linear operator, for example, the blur or colour fringing induced by chromatic aberration, or the sub-sampling by the CFA. Finally, $\boldsymbol{f} \in Y$ denotes the image that is observed, and which represents a degraded version of the original.

Generally, recovering the original image $\boldsymbol{u}$ is *ill-posed* in the sense that there is no unique solution. This is due to the following facts. First, the operator A might not be invertible (e.g., if it is a compact integral operator, which is usually the case for deblurring problems, as we will see in Chapter 2.) Second, we do not know much about the noise $\boldsymbol{\eta}$, other than perhaps a few statistics such as its mean and variance. Consequently, we cannot obtain $\boldsymbol{u}$ directly by solving (1.1).

To obtain a well-posed problem, it is necessary to introduce additional constraints that favour certain solutions over others. One way to accomplish this is to consider the regularised problem

$$\min_{\boldsymbol{u} \in X} \ \varphi(\boldsymbol{u}) + \psi(A\boldsymbol{u} - \boldsymbol{f}), \tag{1.2}$$

where $\varphi : X \to \mathbb{R}$ and $\psi : Y \to \mathbb{R}$ are positive convex functionals. (The importance of convexity will be discussed further in Chapter 3.)

The term on the right in (1.2) is called the *data fidelity term*. This is because it forces the solution $\boldsymbol{u}$ to remain close to the observed data. The choice of

the functional ψ usually depends on the type of noise that is affecting the observations in (1.1). The remaining term is called the *regularisation term* and φ is called the *regularisation functional*. This part represents the additional constraints on the solution, as mentioned above. Traditionally, φ is chosen so that smooth or regular solutions are preferred, hence the term regularisation. In a certain sense the regularisation term can be likened to specifying a prior model for the solution because the functional φ usually restricts the solution to lie in some particular subspace of X.

Let us consider a concrete example. Suppose that we are dealing with scalar images so that in (1.1) the functions u and f and the noise η take values in $\mathbb{R}$. Let $|\cdot|$ denote the Euclidean norm in $\mathbb{R}^n$. Recall that $L^p(\Omega)$ for $1 \leq p < \infty$ is the space of p-integrable functions on Ω with respect to the Lebesgue measure that are identified upon agreeing almost everywhere, i.e., except on sets of measure zero. This is a complete vector space with respect to the norm

$$\|g\|_{L^p(\Omega)} := \left(\int_\Omega |g(x)|^p \, dx \right)^{1/p}.$$

Also, the space $L^\infty(\Omega)$ consists of essentially bounded functions, i.e., bounded except on a set of measure zero. It is a complete normed vector space with the essential sup-norm

$$\|g\|_{L^\infty(\Omega)} := \operatorname{ess\,sup} g := \inf \{a \in \mathbb{R} \mid \mu(\{x \mid g(x) > a\}) = 0\},$$

where μ is the Lebesgue measure. Moreover, in the special case $p = 2$, we have that $L^2(\Omega)$ is a Hilbert space equipped with the inner product

$$(f, g)_{L^2(\Omega)} := \int_\Omega f(x) g(x) \, dx.$$

Usually, we will drop the subscript and only write $(\cdot, \cdot)$ for the inner product when no confusion can arise.

Now, take $X = Y = L^2(\Omega)$ in the regularised problem (1.2). Usually, we choose the data fidelity term to give a least-squares approximation. That is, we choose ψ to be the mapping

$$\psi(u) = \frac{\lambda}{2} \|Au - f\|^2_{L^2(\Omega)},$$

where $\lambda > 0$ is a parameter. Suppose also that the regularisation functional is $\varphi(u) = \|Tu\|^2_{L^2(\Omega)}$, where T is some linear operator on $L^2(\Omega)$, e.g., the identity operator or a differential operator. Then we obtain the classical method known as Tikhonov regularisation [75],

$$\min_{u \in L^2(\Omega)} \|Tu\|^2_{L^2(\Omega)} + \frac{\lambda}{2} \|Au - f\|^2_{L^2(\Omega)}. \tag{1.3}$$

Here the parameter λ balances the regularity of the solution against the fit to the data. That is, for smaller values of λ the solution is more regular, whereas larger values of λ force the solution to be closer to the observed data f. By setting $T = I$, where I is the identity operator, one obtains the minimum norm solution. Another choice is to use the Laplacian Δu or the gradient ∇u, thereby penalising variations in the solution u.

In the discrete case, the method of Tikhonov regularisation reduces to solving a linear system of equations, and it is therefore relatively fast to compute the solution. Unfortunately, this regularisation is inappropriate for some image restoration problems because it imposes *too much* regularity. By overly smoothing the solution, the edges of the image are lost and appear to be blurred.

To remedy this problem, Rudin, Osher and Fatemi [65] proposed instead to use the regularisation functional $\varphi(u) = \|\nabla u\|_{L^1(\Omega)}$, that is, the L^1-norm of the gradient. This functional is also called the *total variation*. The proposed regularisation leads to the *Rudin–Osher–Fatemi (ROF) model*

$$\min_{u \in L^2(\Omega)} \|\nabla u\|_{L^1(\Omega)} + \frac{\lambda}{2} \|Au - f\|^2_{L^2(\Omega)} \,. \tag{1.4}$$

Roughly speaking, the L^1-norm does not so much penalise variations that are associated with edges in the image. It can be shown, under some fairly weak assumptions on the operator A, that the solution to the ROF model lies in the space $BV(\Omega)$ of functions of bounded variation. This is a function space in which certain types of discontinuities are allowed, and it is therefore well suited for edge-preserving image restoration. We will consider bounded variation functions more closely in Chapter 3.

Another way to motivate the ROF model is the following. In the discrete counterpart of the ROF model, images are represented by vectors in some finite-dimensional vector space, rather than functions. Moreover, the total variation is replaced by the vector ℓ^1-norm of the gradient (or, more precisely, a finite difference approximation thereof). It turns out that the solution to this discretised ROF model has a nearly sparse gradient [28]. That is, the gradient of the restored image is a vector consisting mainly of zeros. This in turn means that the image will tend to be made up of piecewise constant regions, and the boundaries between these regions are marked by sharp edges.

One could attempt to find the sparsest possible solution by minimising the vector ℓ^0-"norm"[1], i.e., the number of non-zero elements of a vector. However, this is a combinatorial optimisation problem and it is known to be NP-hard. By instead using the ℓ^1-norm, we can find a *nearly* sparse solution, but the minimisation problem becomes convex. Consequently, the solution may be computed feasibly by using the methods of convex analysis. This principle has also been applied to other problems in sparse recovery such as basis pursuit, see, e.g., [22].

Since we are primarily concerned with restoration problems for colour images, we also consider the extension of the ROF model (1.4) to vector-valued functions. Although it is possible to define this generalisation in more than one way, see, e.g., [9, 33], we follow the work of Bresson and Chan in [11], and consider a so-called *vectorial ROF (VROF) model*. This method will be applied to restore images suffering from chromatic aberration.

In addition, Shen and Chan [72] used the ROF model as the basis for their local non-texture inpainting method. Let Ω denote the image domain and let f be an image defined on Ω. Suppose that there is an inpainting region $\Gamma \subset \Omega$ wherein the data $f\big|_{\Gamma}$, i.e., the restriction of f to Γ, has been lost. The TV

[1] The quotation marks are used because this is not actually a norm in the usual sense.

inpainting model is then the minimisation problem

$$\min_{u \in L^2(\Omega)} \|\nabla u\|_{L^1(\Omega)} + \frac{\lambda}{2} \|u - f\|^2_{L^2(\Omega \setminus \Gamma)} . \tag{1.5}$$

This is precisely the ROF model, except that the data fidelity term has been modified to reflect that we have no data in the inpainting domain Γ.

Finally, we present a new TV-based model for demosaicking. This is based on an idea of Alleysson *et al.* in [2], namely that the demosaicking of a colour image can be greatly improved by decomposing the image into a luminance component and a chrominance component. The luminance represents the spatial information in the image, whereas the chrominance represents colour information. The latter tends to vary much less than the luminance, and this heuristic is incorporated into our TV demosaicking model. This leads to improved demosaicking results compared to other regularisations. We will return to discuss these TV-based methods in further detail in Chapter 4.

In the meanwhile, we provide the following short summary. We are concerned with image restoration problems that are relevant to single-sensor digital colour cameras. In particular, this includes chromatic aberration and demosaicking. The focus is on a variational approach, more specifically the total variation-based method of the ROF model and some variants of it. The ROF model is already known to perform well for image denoising and deblurring, and so it makes for a good starting point. Our goal is to demonstrate that a TV-based approach can also be used to remove chromatic aberration, and to solve the demosaicking problem.

1.3 The contributions of this thesis

Part of this thesis is dedicated to giving an introduction to total variation-based image restoration methods. In order to achieve this, we have included a review of several such existing methods in Chapter 4, including the ROF model, the vectorial ROF model for colour image restoration, and the TV inpainting model. This part does not represent any new work in itself, and the applications to image denoising and deblurring in Chapter 6 are fairly well known in the variational image processing literature.

On the other hand, we make an attempt at tying the above mentioned image restoration problems to some quite concrete applications in digital colour imaging. So, in Chapter 2, we expend some effort to explain where these colour image restoration problems come from.

The main contribution is the TV-based demosaicking model that is presented in Section 4.4. This is related to the algorithm of Alleysson *et al.* [2] which demonstrates that the demosaicking is greatly improved by treating the image as a sum of luminance and chrominance components. We adapt this approach to the framework of TV-based regularisation by applying the regularisation to the luminance and chrominance components of the image.

To numerically compute the solution to the TV demosaicking model, we apply a recent numerical algorithm that was introduced by Chambolle and Pock in [17] for minimising the sum of two convex functionals. This applies directly to the regularisation problem (1.2). The algorithm is based on the primal-dual

formulation of the given variational problem. For the TV demosaicking model, the corresponding primal-dual algorithm is given in Section 5.3.

The demosaicking problem has previously been approached from a regularisation point of view. For example, Menon and Calvagno [54] and Condat [24] both propose a quadratic Tikhonov regularisation. The regularisation in [24] is also performed in a luminance-chrominance basis. The use of TV-based regularisation for demosaicking was also considered by Saito and Komatsu [68–70]. However, their method employs only a simple channel-by-channel TV regularisation. Moreover, the numerical method considered is a simple gradient descent that relies on a smooth approximation to the total variation and is known to be very inefficient.

1.4 The organisation of this thesis

The rest of this thesis is organised as follows. In the next chapter we explain the fundamentals of a digital CCD/CMOS colour imaging device and how the images produced by such a device become degraded. Chapter 3 provides the mathematical background necessary for understanding the variational methods studied in later chapters. This includes a brief introduction to convex analysis and also some results regarding the total variation and functions of bounded variation. In Chapter 4 we present a number of TV-based regularisation methods. This includes the ROF model for denoising and deblurring grayscale images, which is then generalised to the vectorial ROF model for colour images. In addition, we present a TV-based model for the demosaicking problem. We provide a proof of the existence of a solution to this problem using a standard argument from the calculus of variations. Numerical methods for the TV-based regularisation methods are then discussed in Chapter 5, and a primal-dual algorithm for the TV demosaicking model is derived. In Chapter 6 we apply these methods to a number of image degradation problems for both grayscale and colour images, including noise removal, deblur, demosaicking and removal of chromatic aberration effects. Finally, we end with some concluding remarks in Chapter 7 and offer some ideas that can be pursued in future work.

CHAPTER 2

Image degradation in a single-sensor digital colour camera

To develop mathematically sound methods for image restoration, it is important to understand how an image becomes degraded. Therefore, we begin by giving a short explanation of the main components in a single-sensor CCD or CMOS colour imaging device. Once we understand the basic operation of such a system, we can, in the next section, identify the various sources of image degradation and model their behaviour.

2.1 The digital CCD/CMOS imaging pipeline

The purpose of a digital imaging system is to obtain a digital representation of the spatial distribution of radiance (Watt/solid angle/unit area) in a scene. This is achieved by measuring the amount of light incident on an image sensor and quantising the values to be represented by a given number of bits. Figure 2.1 gives a schematic representation of the imaging process including the main stages that we will now describe. The following presentation is based on the articles by Healey and Kondepudy [40] and Ramanath *et al.* [61]. Janesick [45] provides a thorough survey of the CCD technology, whereas CMOS image sensors are covered in Nakamura [58]. Concepts relating to colour imaging and colour perception can be found in more detail in the books by Reinhard [62] and Sharma [71]. Finally, an excellent survey on demosaicking is given by Gunturk *et al.* in [38] or by Menon and Calvagno in [55].

A brief tour of the digital camera

The *scene irradiance* (Watt/unit area) is the light which is emitted or reflected by objects in a scene, travels in the direction of the imaging device and is accumulated on the 2-D image plane where the image sensor lies. Similarly, the

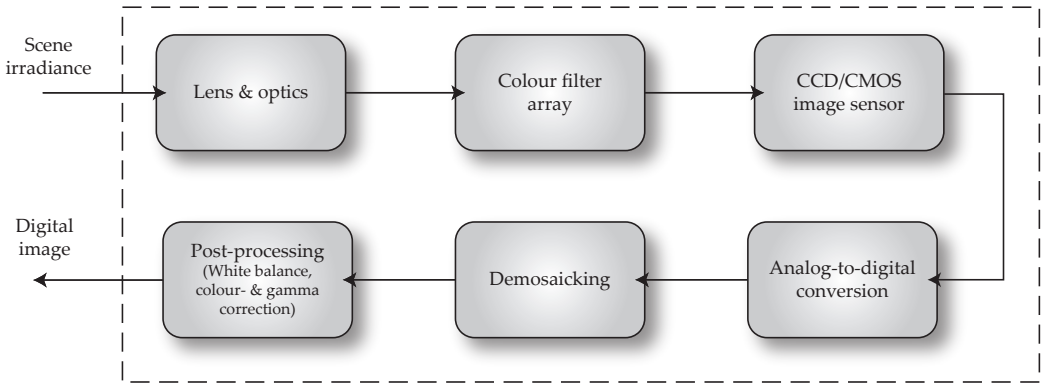

Figure 2.1: *The digital imaging pipeline. The diagram is a simplified representation of the main components and processing stages in a single-sensor digital CCD/CMOS imaging device. The scene irradiance denotes the light which is focused onto the image plane by the lens and optics. The colour filter array (CFA) ensures that each photo-detector in the image sensor records wavelengths corresponding to a single colour, e.g. red, green or blue, to which that pixel is assigned. The sensor itself measures the number of incident photons on a 2-D array of photo-detectors, producing a signal that is quantised by the analog-to-digital converter. A process known as demosaicking is then necessary to estimate missing colours at every pixel due to the CFA. Further post-processing, (e.g. white balance, colour- and gamma correction), is also necessary to produce a final digital image that is suitable for viewing.*

term *scene spectral irradiance* (Watt/unit area/wavelength) is used to denote the light incident on the image plane whenever we need to distinguish it by wavelength. Once the light reaches the imaging device, it is first led through the optical system. This includes a lens and focus control for focusing light onto the image sensor, an optical low-pass filter to eliminate high frequency spatial variations that can cause aliasing artefacts, and a spectrally selective filter that reflects or absorbs unwanted wavelengths, such as infrared radiation. The aperture is the adjustable hole through which the incoming light passes. In a still camera, a mechanical shutter is used to control the amount of time the sensor is exposed to incident light.

The image sensor uses a 2-D array of photo-detectors, known as photo-sites or potential wells, to measure the scene irradiance. At each photo-site in a CCD, photons strike the silicon wafer and produce an electron-hole pair. The charge generated by this process is proportional to the number of incident photons. The photo-detectors in a CMOS sensor consist of a photo-diode and a readout circuit. The photo-diodes are known as active-pixel sensors (APS) because the generated charge is amplified within the photo-detector. This is not the case in a CCD where the pixels are said to be passive. The signal produced at each photo-site is obtained by integrating over time the charge generated by incident photons.

Consider the transformation of the scene spectral irradiance into the discrete set of pixel values and colour channels that we will call the *image spectral irradiance*. Let $\Omega \subset \mathbb{R}^2$ denote the image plane and suppose it is partitioned

uniformly into M-by-N sub-domains $\Omega = \bigcup_{i,j}^{M,N} \Omega_{i,j}$, where each sub-domain corresponds to the spatial extent of a single photo-detector or pixel. Let $R(x, \lambda)$ denote the scene spectral irradiance (Watts/unit area/wavelength) incident on the image plane. Assume that these values are recorded for P colour channels $L = (L^1, \ldots, L^P)$, and each channel consists of M-by-N pixels, $L^k \in \mathbb{R}^{MN}$ for $k = 1, \ldots, P$. Then the image spectral irradiance of the (i, j)-th pixel of the k-th colour channel is given by

$$L_{i,j}^k = T \iint\limits_{\Omega_{i,j} \times \mathbb{R}} R(x, \lambda) S_{i,j}^k(x, \lambda) \, dx \, d\lambda. \tag{2.1}$$

Here T is the integration time (in seconds) and $S_{i,j}^k(x, \lambda)$ is the spatio-spectral response (electrons/Joule) of the photo-site, i.e., the ratio of electrons collected per incident light energy. Usually, this can be written as the product $S_{i,j}^k(x, \lambda) = S_{i,j}(x) q^k(\lambda)$ of a spatial response $S_{i,j}(x)$ and a spectral component $q^k(\lambda)$, sometimes called the quantum efficiency or spectral sensitivity associated with the k-th colour channel.

After an exposure, the charge packets at each photo-site are read out and quantised by an analog-to-digital converter (ADC). The CCD gets its name from a process known as charge coupling that is used to repeatedly shift electrons to adjacent photo-sites to be read out sequentially. The digital signal produced by the ADC is a set of image intensity or brightness values that represent the scene irradiance. There is a non-linear relationship between the image brightness and the image irradiance which is described by the *camera response function* (see [36], [37]). This is an important quantity for calibrating imaging devices, particularly in applications where the irradiance values are required to be in proper radiometric units.

In addition, the image brightness values are adjusted, so that when they are rendered on a display, the reproduced colours better correspond with our perception and how they were perceived when they were captured. This is a post-processing stage that consists of several steps, including *colour correction*, *white balance*, and *gamma correction*. For further details, see, e.g., [71].

Digital colour imaging

Initially, the photo-detectors in the image sensor have no way of distinguishing between photons of different wavelengths. However, to obtain multi-channel data, and colour images in particular, it is necessary to separately measure the scene irradiance for a given number of wavelengths. For colour images it is conventional to record three colour channels corresponding to red, green and blue light. This is done to exploit or emulate the trichromacy of the human visual system (see, e.g., [62]) where there are three types of colour-sensing cells, known as short (S), medium (M) and long (L) cones, which primarily respond to wavelengths of 420nm, 530nm and 560nm, respectively. All visible colours are experienced through different combinations of stimuli of the three types of cone cells.

One way of obtaining multi-channel images is to take multiple exposures of the same scene, each with different spectral filters that allow only certain wavelengths to pass. Unfortunately, this requires the scene to be static and unchanged between exposures. This is rarely the case. A second alternative is

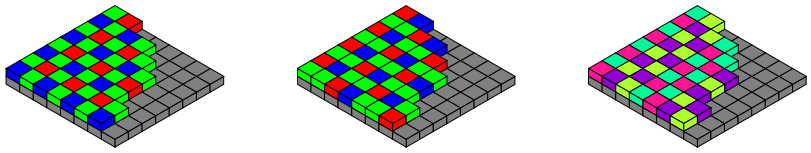

(a) Bayer pattern (b) Lukac pattern (c) Hirakawa–Wolfe pattern

Figure 2.2: *Colour filter array designs. A colour filter array consists of a spectral filter for each photo-detector in the image sensor, thus assigning a single colour to each pixel. Figure 2.2(a) shows the most commonly used configuration, known as the Bayer CFA. Figure 2.2(b) shows a pattern proposed by Lukac and Plataniotis in [51] together with a framework for CFA design. The configuration in Figure 2.2(c) uses panchromatic elements, i.e., combinations of red, green and blue filters, and was obtained by Hirakawa and Wolfe in [43] as an optimal pattern based on a set of optimisation criteria. (Bayer pattern image courtesy of Colin M.L. Burnett under the GNU Free Documentation License (GFDL)).*

to use a single exposure, but multiple image sensors, one for each channel. For digital colour cameras, this is known as "3CCD" or "three-chip" technology [59], and it is realised by using a trichroic prism that splits the light into separate red, green and blue beams, directing each to a separate sensor. However, image sensors are very expensive components and this approach still requires a careful calibration and registration to match the images captured by each sensor.

A third, more cost-effective solution can be achieved using only a single sensor and a *colour filter array (CFA)*, a mosaic-like pattern of spectral filters, one for each photo-detector. See Figure 2.2. In this way, multiple colour channels can be acquired by a single sensor during a single exposure. Although the drawback is a reduced resolution for each colour channel, since each pixel can only belong to one channel, this design is so much cheaper compared to the three-sensor alternative that it is widely used in digital cameras.

The most commonly used mosaic pattern for a CFA is the Bayer pattern [8], shown in Figure 2.2(a). It consists of odd rows alternating between green and red pixels and even rows alternating between blue and green pixels. It is usually argued [38] that the green component most closely corresponds to the human perception of luminance, and it is therefore assigned twice the number of pixels. This seems reasonable because the human visual system has less acuity to chrominance variations (see, e.g., [2]), which are usually identified with the red and blue components. Figure 2.2 also shows two of the many other CFA configurations which have been proposed. See [43] and [51] for others.

As mentioned in Chapter 1, employing a CFA in a single-sensor camera introduces an additional processing stage in the pipeline. Since each photo-detector has sensed only one colour, a process known as *demosaicking* (or CFA interpolation) is required to synthesise the sub-sampled data into a full-colour image. For example, consider the Bayer pattern (Figure 2.2(a)). The following steps are necessary to reconstruct a full-colour image. The green colour values must be estimated at the red and blue pixels, the red colour values must be estimated at the green and blue pixels, and, finally, the blue colour values must be estimated at the green and red pixels. The resulting image is then

represented by red, green and blue image brightness values at every pixel.

The literature on the demosaicking problem is comprehensive, and a great number of methods, spanning the whole repertoire of image processing, have been proposed to solve this problem. We will not provide a review here, but instead we refer to the surveys by Gunturk *et al.* [38] and Menon and Calvagno [55]. In Chapter 4, we will return to discuss demosaicking in the context of total variation-based regularisation methods.

2.2 Modelling image degradation

In this section, we consider images captured by a single-sensor digital colour camera and discuss some of the ways in which they become degraded during this process. This includes sensor noise, blur and chromatic aberration due to the optical lens system, and a loss of information as a result of employing a colour filter array. These degradations give rise to the image restoration problems which we will study later, and they can all be represented as part of the forward image model.

Recall that the forward image model describes an observed image $\boldsymbol{f}$ in terms of the undegraded original $\boldsymbol{u}$ through the relation

$$\boldsymbol{f} = A\boldsymbol{u} + \boldsymbol{\eta}, \tag{2.2}$$

where A is a linear operator and $\boldsymbol{\eta}$ represents noise. In the following, we describe in further detail the nature of the degradation operator and the noise.

Sensor noise

The presence of noise in the image sensor results in random perturbations of the measured image irradiance values. To describe this randomness, we first make a few definitions (see, e.g., [60]). In this section, we let Ω denote a sample space consisting of all possible outcomes of a given random experiment. Let I be a non-empty index set (e.g., $I = [0, 1]$), then a collection $\{X_i\}_{i \in I}$ of random variables is called a *stochastic process*. If the underlying index set I is not an interval, but instead a subset of $\mathbb{R}^m$, $m \geq 2$, then the above collection of random variables is called a *random field*.

Now, let $U \subset \mathbb{R}^m$ denote the image domain and let $\boldsymbol{f}$ be a vector-valued image with values in $\mathbb{R}^n$. The noise, denoted by $\boldsymbol{\eta}$, is a random field. Roughly speaking, this can be regarded as assigning a random variable $\boldsymbol{\eta}_x$ to each point $x \in U$ in the image domain. We classify this stochastic process as follows. First, it may be *signal-independent*, or additive, in which case its distribution does not depend on the image. Otherwise, we say that it is signal-dependent. Although this dependence can be of a quite general sort, we usually take this to mean that the variance of the noise is proportional to the image $\boldsymbol{f}$. Second, if a stochastic process $\boldsymbol{\eta}$ follows a Gaussian distribution, it is said to be *white* (in the strong sense) if each random variable $\boldsymbol{\eta}_x$ for $x \in U$ is independently and identically distributed with zero mean and equal variance.

The study of noise in image sensors has been covered by a number of works, e.g., the books by Janesick [45] and Nakamura [58]. The following presentation is chiefly based on the article by Liu *et al.* [48].

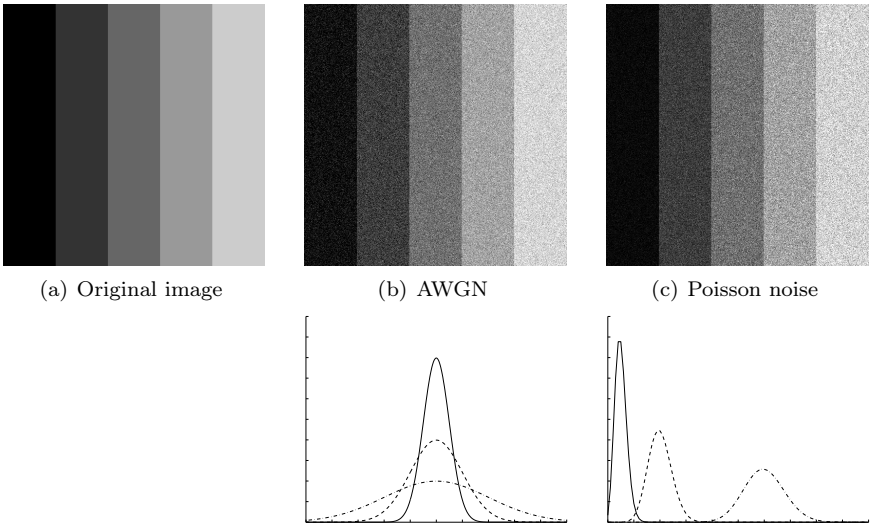

(a) Original image (b) AWGN (c) Poisson noise

(d) Gaussian distribution (e) Poisson distribution

Figure 2.3: *Noise models. Figures 2.3(b) and 2.3(c) show two different types of noise that apply to digital CCD/CMOS cameras. These are additive white Gaussian noise (AWGN) and signal-dependent Poisson-distributed noise, respectively. They appear quite similar, except that the variance of the Poisson noise is proportional to the image intensity. Below are a few realisations of the Gaussian and Poisson distributions that govern these noise models.*

For our purposes, the most important sources of noise are *dark current noise* and *shot noise*. In addition, each photo-detector suffers from *fixed-pattern noise*, and, during the analog-to-digital conversion, the image is subjected to *quantisation noise*. The fixed-pattern noise is caused by individual differences between photo-sites that are introduced during the manufacturing process. For example, the photo-sites may differ in size or quantum efficiency, i.e., the ratio of photons that are converted into electrons. Since these effects are mostly deterministic, they can be measured and more or less eliminated. The quantisation noise is associated with the process of analog-to-digital conversion. This defines the minimum noise level that the imaging device can achieve, but it is usually so small that it can be disregarded entirely.

Dark current is the generation of electrons from thermal energy in the silicon wafer of the sensor. These electrons are not related to any incoming photons, and are therefore called "dark" electrons. The number of "dark" electrons collected at each photo-site is proportional to the integration time and follows a Gaussian distribution. This effect is also highly dependent on temperature. Even though it can be reduced by cooling the device, it accounts for a significant portion of the noise in practical situations. The dark current noise experienced by a photo-detector is regarded as independent of that of other photo-detectors. Hence, the appropriate noise model is an *additive white Gaussian noise (AWGN)*, and it is illustrated in Figure 2.3(b).

Shot noise is the term for random fluctuations in the number of incoming photons due to the particle nature of light. It follows a Poisson distribution with a variance proportional to the image irradiance. Therefore, brighter areas

of the image will contain larger perturbations from the true pixel values. This is shown in Figure 2.3(c). Unlike dark current noise, shot noise is a fundamental limitation that cannot be eliminated by physical means. Although it is mainly in low-light conditions that shot noise dominates (in terms of signal-to-noise-ratio) over other noise sources, it also becomes a serious concern when the size of the photo-sites shrinks to accomodate more elements and a higher resolution [41].

In conclusion, the noise in a real CCD or CMOS image sensor can be modelled by two components. The first is an additive white Gaussian noise, and the second is a signal-dependent Poisson-distributed noise. These are usually considered separately, but in some cases, e.g., [50], they are modelled simultaneously as a mixed Poisson–Gaussian process. Under some circumstances, especially when the number of events is large (i.e., a large number of incoming photons), the Poisson noise is considered to be sufficiently close to a Gaussian-distributed noise, and the AWGN assumption is a good approximation. Yet another possibility is to transform the image in a manner so that the transformed noise becomes Gaussian-distributed, e.g., by using the Anscombe transform [52]. This approach is quite common for denoising images with Poisson-distributed noise. With this reasoning, we will henceforth consider mainly the Gaussian-distributed noise component.

We note also that there is a practical aspect of employing real CCD and CMOS noise models in that certain parameters, such as the mean and variance of the noise, must be somehow estimated. We leave this as a purely practical consideration and refer the reader to Liu *et al.* [48].

Blur and chromatic aberration

The optical system of an imaging device is designed to project a 3-D scene onto a 2-D image domain. In practice, this projection also introduces distortions and aberrations in the acquired image, and objects appear to be warped by the shape of the lens.

It is well known in optics (see, e.g., [10]) that any finite aperture causes diffraction. That is, the fact that the light passes through a small hole causes the photons to interact and create interference patterns. If one should attempt to capture an image of an ideal point-light source (a unit impulse), one would instead observe that the light is spread out over a larger area of the image plane. If the photo-detectors and aperture are small enough, this will appear as a blur covering several pixels. In addition, the use of optical low-pass filters and imperfections in the lens gives rise to further blurring effects. Finally, if it is large enough, the relative motion of the camera and the object being imaged will also induce a similar effect, known as motion blur. These effects are illustrated in Figure 2.4.

The blurring of an image can be modelled in terms of an integral operator acting on the scene irradiance values. Recall that Ω denotes the image domain. For simplicity, consider a scalar image $u : \Omega \to \mathbb{R}$ representing the scene irradiance. Let $f : \Omega \to \mathbb{R}$ denote the same irradiance values after they have passed through the optical system and have been subjected to blur. Then u and f are related by the operator equation $f = Au$, where A is the continuous linear

(a) Original *cameraman* image (b) Gaussian blur (c) Motion blur

Figure 2.4: *Blur in images. The original image is shown in 2.4(a). The image in 2.4(b) has been subject to a Gaussian blur which can sometimes be used as a simplified model for the blur caused by a camera lens. In 2.4(c), the blur is caused by simulating the camera shaking during the exposure.*

integral operator defined by

$$f(x) = (Au)(x) := \int_{\Omega} K(x, y)u(y)\, dy. \tag{2.3}$$

The kernel $K(x, y)$ is known as the *point-spread function (PSF)*, and it can be interpreted as follows. Suppose we position a point-light source such that, ideally, it were only to excite the point $y \in \Omega$. The function $K(x, y)$ measures the actual excitation under the given optical system at every point $x \in \Omega$.

Notice that in the discrete case whenever $u, f \in \mathbb{R}^N$, then (2.3) reduces to a matrix-vector multiplication $f = Au$, where A is an N-by-N matrix.

In the special case when the blur is spatially invariant, i.e., it is the same over the entire image, then the PSF can be written as a function of a single variable $K(x, y) = K(x - y)$. Equation (2.3) then reduces to a convolution $f = K * u$. This case is particularly interesting because it is very efficient to compute. Numerical algorithms often exploit the fact that in the discrete case the operator A can be represented by a Toeplitz matrix which is diagonalised by the Fourier transform.

For colour images the matter becomes slightly more complicated. Another well known fact in optics (again, see [10]) is that the index of refraction of the lens material varies with the wavelength of light. This mean that photons of different wavelengths corresponding to different colours will travel along different paths in the imaging device even though they originate from the same point. This phenomenon is known as dispersion, and it is the source of two kinds of chromatic aberrations in colour imaging devices. The first is known as *lateral* chromatic aberration and manifests as a geometric shift between the different colour channels, see Figure 2.5(a). The second is a *longitudinal* (or axial) aberration and is the result of the different wavelengths focusing at different points along the optical axis, see Figure 2.5(b). A typical artefact in colour images that appears as a consequence of chromatic aberration is *colour fringing*, and an example is shown in Figure 1.1 in Chapter 1.

Chromatic aberration can be modelled by a linear integral operator in the same way as the blur we considered above. Consider the multi-channel images

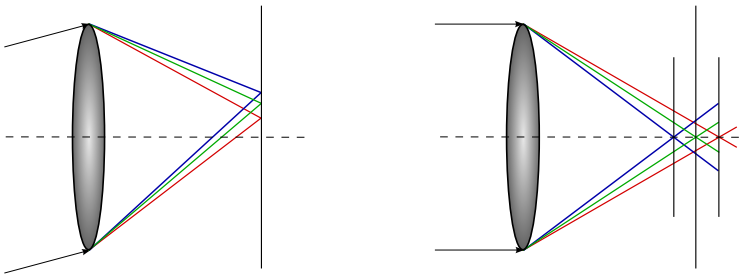

(a) Lateral aberration (b) Longitudinal aberration

Figure 2.5: *Chromatic aberration. The index of refraction in the lens material varies with wavelength, thereby causing light of different colours to* (i) *appear shifted relative to one another, as shown in 2.5(a), and* (ii) *focus at different distances along the optical axis, as shown in 2.5(b).*

$\boldsymbol{u}$ and $\boldsymbol{f}$ taking values in $\mathbb{R}^n$. Once more, $\boldsymbol{u}$ represents the scene irradiance, whereas $\boldsymbol{f}$ represents the irradiance after it has passed through the optics. Then they are related by the equation $\boldsymbol{f} = A\boldsymbol{u}$, and the integral operator A is defined by

$$\boldsymbol{f}(x) = (A\boldsymbol{u})(x) := \int_\Omega K(x,y)\boldsymbol{u}(y)\, dy. \tag{2.4}$$

The difference from (2.3) is that now the kernel $K(x,y)$ is an n-by-n matrix for each $x, y \in \Omega$. Furthermore, we interpret the integral of a vector-valued function $\boldsymbol{g}(x) = \big(g^1(x), \ldots, g^n(x)\big)$ to mean the vector of component-wise integrals,

$$\int_\Omega \boldsymbol{g}(x)\, dx := \begin{pmatrix} \int_\Omega g^1(x)\, dx \\ \vdots \\ \int_\Omega g^n(x)\, dx \end{pmatrix}.$$

The matrix $K(x,y) = (K_{i,j}(x,y))_{i,j=1}^n$ describes dependencies between the different colour channels. More precisely, the (i,j)-th component represents a cross-channel influence between the i-th and j-th colour channels. However, when dealing with chromatic aberration, we normally consider $K(x,y)$ to be a diagonal matrix of n independent PSFs. Hence, there is no "leakage" between colour channels.

In practice, we are required to estimate the PSF of a given imaging device, should we want to compensate for the effects of blur and chromatic aberration. In certain cases, a parametric model can be used. For example, it is quite comon to assume that the PSF is a Gaussian function. In contrast, Delbracio *et al.* [26] recently suggested a procedure for estimating the PSF of a real-world camera with sub-pixel accuracy and without being constrained to a parametric representation. Moreover, their procedure is formulated as a well-posed minimisation problem. It is sometimes possible to perform an image restoration without prior knowledge of the PSF. The idea is then to simultaneously estimate the PSF and the restored image as they are both subject to uncertainty. This is referred to as a *blind* image restoration problem, see, e.g., [79] for further details.

Colour filter array

Next, we turn to modelling the effects of the colour filter array. Let $\boldsymbol{u} : \Omega \to \mathbb{R}^n$ denote the image irradiance. Then, at each point $x \in \Omega$, consider the multiplication of the image irradiance $\boldsymbol{u}(x)$ with a very simple n-by-n colour filter matrix $C(x)$ that discards the values of certain colour channels. The resulting mosaicked image irradiance values are $\boldsymbol{f} = A\boldsymbol{u}$, where A is given by

$$\boldsymbol{f}(x) = (A\boldsymbol{u})(x) := C(x)\boldsymbol{u}(x). \tag{2.5}$$

As an example, take the Bayer pattern shown in Figure 2.2(a). The following description is easily adapted to other configurations. Let $\Omega = \bigcup_{i,j}^{M,N} \Omega_{i,j}$ denote the image domain and its partitioning into M-by-N subdomains, each corresponding to the area of a pixel. For an ideal Bayer CFA, the colour filter matrix is diagonal, $C(x) = \operatorname{diag}(C_R(x), C_G(x), C_B(x))$, where

$$\begin{aligned}
C_R(x) &= \begin{cases} 1, & x \in \Omega_{i,j},\ i \text{ is odd and } j \text{ is even,} \\ 0, & \text{otherwise.} \end{cases} \\
C_G(x) &= \begin{cases} 1, & x \in \Omega_{i,j},\ i \text{ and } j \text{ are both odd or both even,} \\ 0, & \text{otherwise.} \end{cases} \\
C_B(x) &= \begin{cases} 1, & x \in \Omega_{i,j},\ i \text{ is even and } j \text{ is odd,} \\ 0, & \text{otherwise.} \end{cases}
\end{aligned} \tag{2.6}$$

Compare this with the Bayer pattern in Figure 2.2(a).

In practice, the spectral filters of the CFA can exhibit flaws that sometimes transmit photons of unwanted wavelengths. Moreover, photons and electrons can "leak" into neighbouring photo-detectors, an effect known as *cross-talk* [41]. This cannot be accounted for by the simple point-wise multiplication of irradiance values with a colour filter matrix. Instead, we write the mosaicking operator A as an integral operator, exactly as in (2.4), the kernel of which is denoted by $C(x, y)$. As in the case with chromatic aberration, the kernel $C(x, y)$ must be somehow estimated or constructed based on knowledge of the CFA, the responses of the individual spectral filters, and the cross-talk effects.

2.3 Summary

In this chapter we have described the basic operation of a single-sensor CCD or CMOS digital colour camera, and we have considered how the images produced by such a device become degraded during the process. The measurements performed by the image sensor are perturbed mainly by dark current noise and shot noise. The former is additive white Gaussian-distributed, whereas the latter is Poisson-distributed with a variance proportional to the measured irradiance values. Next, there are optical effects that cause the final image to suffer from blur and chromatic aberration, both of which are modelled by a linear integral operator whose kernel is the point-spread function of the optical system. Finally, the effects of the CFA can also be described in terms of a linear operator, which in the ideal case, i.e., when there is no cross-talk, takes the very simple form of a pointwise multiplication by a colour filter matrix.

We proceed to discuss how these effects might be undone, or at least compensated for by a post-processing of the digital image, and we begin, in the next chapter, with some of the necessary mathematical background.

CHAPTER 3

Mathematical preliminaries

In this chapter we give a brief review of some mathematical concepts that are central to understanding the variational image restoration methods in Chapter 4. First, we briefly survey some relevant definitions and theorems from convex analysis. This is the framework in which convex optimisation problems, such as the ROF model, are best understood. The central result here is one which provides sufficient conditions for the existence of a minimiser of a convex function. Second, we introduce the notion of total variation and bounded variation functions. This is tied directly to the regularisation of our inverse imaging problems. The total variation is precisely the regularisation functional that is used in the ROF model, and bounded variation functions serve to characterise the solutions to the regularised inverse problems. This material will be used to prove the existence of a unique solution to the ROF model and some related TV-based models in Chapter 4.

3.1 Basic convex analysis

Convexity is a very desirable property in variational problems. If a given problem is known to be convex, then, generally, a globally optimal solution, if it exists, can be computed, and it can be done within a reasonable amount of time. Fortunately, the variational problems that we will consider, namely the ROF model and related problems, are convex. More generally, the field of convex analysis has applications to linear and nonlinear programming, as well as stochastic programming, optimal control theory and partial differential equations [64].

The following material is chiefly based on Rockafellar [63], [64]. It is assumed that the reader has some familiarity with real and functional analysis and also some basic definitions from topology. See, for instance, Rudin [66], [67] for the former and Munkres [57] for the latter.

First, we introduce some notation. Throughout this chapter, let X denote a vector space over $\mathbb{R}$. Usually, this can be thought of as representing $\mathbb{R}^n$, but it will also be necessary to consider more general infinite-dimensional vector spaces whose elements are functions. The usual Euclidean inner product of

$x, y \in \mathbb{R}^n$ is denoted by $x \cdot y$, and the corresponding norm $x \mapsto \sqrt{x \cdot x}$ is denoted by $|x|$. Whenever a vector space X is provided with an inner product, we denote this by the bilinear form $(\cdot, \cdot) : X \times X \to \mathbb{R}$. If X is a normed space, its norm is written $\|\cdot\|_X$, or sometimes just $\|\cdot\|$, if there can be no confusion.

Moreover, X is said to be a *topological vector space* whenever it is endowed with a topology such that the operations of vector addition and scalar-vector multiplication,

$$(x, y) \mapsto x + y, \qquad (\lambda, x) \mapsto \lambda x,$$

are continuous. A *locally convex space* is a vector space with a topology that is generated by a fundamental system of convex neighbourhoods relative to the origin, (see [67]). All the usual topological vector spaces encountered in analysis are locally convex. In particular normed spaces are locally convex.

It is sometimes convenient to allow a function to take on the values $\pm\infty$. Therefore, we define the set of *extended real values* by $\bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty, -\infty\}$. Furthermore, we augment the usual arithmetic operations with the following definitions:

$$\begin{aligned}
&\alpha + \infty = \infty + \alpha = +\infty && \text{for} \quad -\infty < \alpha \leq \infty, \\
&\alpha - \infty = -\infty + \alpha = -\infty && \text{for} \quad -\infty \leq \alpha < \infty, \\
&\alpha \cdot (\pm\infty) = (\pm\infty) \cdot \alpha = \pm\infty && \text{for} \quad 0 < \alpha \leq \infty, \\
&\alpha \cdot (\pm\infty) = (\pm\infty) \cdot \alpha = \mp\infty && \text{for} \quad -\infty \leq \alpha < 0, \\
&0 \cdot (\pm\infty) = (\pm\infty) \cdot 0 = 0.
\end{aligned}$$

Notice that the expressions $\infty - \infty$ and $-\infty + \infty$ remain undefined and should therefore be avoided.

Convex sets

A convex set is one which contains every line segment connecting a pair of points in the set.

Definition 3.1 (Convex set)**.** *A set $C \subset X$ is said to be* convex *if*

$$\alpha x + (1 - \alpha) y \in C \tag{3.1}$$

for every $x, y \in C$ and $0 \leq \alpha \leq 1$.

A typical example of convex sets are those associated with an *affine hyperplane* $\mathcal{H}$, which, in $\mathbb{R}^n$, is given by the set $\mathcal{H} = \{x \in \mathbb{R}^n \mid a \cdot x = \alpha\}$ for some $\alpha \in \mathbb{R}$ and a non-zero normal vector $a \in \mathbb{R}^n$. More generally, in a vector space X, $\mathcal{H}$ is determined by a linear equation $\mathcal{H} = \{x \in X \mid \ell(x) = \alpha\}$, where $\alpha \in \mathbb{R}$ and $\ell : X \to \mathbb{R}$ is a non-zero linear functional on X. An affine hyperplane separates the space X into two convex sets, known as the *open half-spaces*, given by

$$\{x \in X \mid \ell(x) < \alpha\} \quad \text{and} \quad \{x \in X \mid \ell(x) > \alpha\}.$$

Similarly, the same affine hyperplane divides X into the *closed half-spaces*

$$\{x \in X \mid \ell(x) \leq \alpha\} \quad \text{and} \quad \{x \in X \mid \ell(x) \geq \alpha\},$$

which are also convex.

Theorem 3.2. *The intersection of an arbitrary collection of convex sets is convex.*

Proof. Let $\{C_i\}_{i\in I}$ be a collection of convex sets, $0 \leq \alpha \leq 1$ and $x, y \in \cap_{i\in I} C_i$. For every $i \in I$ we have $x, y \in C_i$, and, from convexity, the point $\alpha x + (1-\alpha)y$ lies in C_i. Therefore it also lies in the intersection $\cap_{i\in I} C_i$. □

The *convex hull* of an arbitrary set $A \subset X$, denoted by $\operatorname{co} A$, is the smallest convex set that contains A. Of course, since convex sets are closed under intersection, $\operatorname{co} A$ is the intersection of all convex convex sets that contain A.

Convex functions

The central objects of study in convex analysis are convex functions.

Definition 3.3 (Convex function)**.** *Let $C \subset X$ be a non-empty convex set. A function $f : C \to \mathbb{R}$ is said to be* convex *on C if*

$$f(\alpha x + (1-\alpha)y) \leq \alpha f(x) + (1-\alpha) f(y), \tag{3.2}$$

for every $x, y \in C$ and $0 < \alpha < 1$. Moreover, f is said to be strictly convex *if the above holds with a strict inequality whenever $x \neq y$.*

A (strictly) concave function is one whose negative is (strictly) convex.

Ultimately, the goal is to solve a minimisation problem on the following form: Given a convex set $C \subset X$ and a convex *objective function* $f : C \to \bar{\mathbb{R}}$, find $x \in C$ such that

$$f(x) = \inf_{y\in C} f(y), \tag{3.3}$$

or, in other words, $f(x) \leq f(y)$, for all $y \in C$. Every x that satisfies this criterion is termed a *solution* to the minimisation problem, and we write,

$$f(x) = \min_{y\in C} f(y) \quad \text{and} \quad x = \arg\min_{y\in C} f(y).$$

It is often more convenient to consider an objective function defined on the whole vector space X, rather than restricted to a subset C. This is done by extending $f : C \to \mathbb{R}$ with the value $+\infty$ outside of the set C

$$\bar{f}(x) = \begin{cases} f(x), & \text{if } x \in C, \\ +\infty, & \text{otherwise.} \end{cases}$$

Notice that minimising f over C is equivalent to minimising $\bar{f}$ over the whole space X. Furthermore, $\bar{f}$ is convex on X if and only if *(i)* C is a convex set, and *(ii)* f is convex on C. From now on, without loss of generality, it is sufficient to consider the case where $C = X$.

The *effective domain* of $f : X \to \bar{\mathbb{R}}$ is the set

$$\operatorname{dom} f = \{x \mid f(x) < +\infty\}. \tag{3.4}$$

Also, f is said to be *proper* if it is nowhere equal to $-\infty$ and not identically equal to $+\infty$.

There are, in fact, equivalent ways of defining the notion of a convex function. We mention here the relation between a convex function and its *epigraph*, that is, the set of points lying "on or above the graph" of f

$$\text{epi } f = \{(x, r) \mid x \in X \text{ and } r \geq f(x)\}. \tag{3.5}$$

This is described in the following proposition.

Proposition 3.4 (§4 in [63]). *A function $f : X \to \bar{\mathbb{R}}$ is convex if and only if its epigraph* epi f *is a convex subset of $X \times \mathbb{R}$.*

Proof. For $x, y \in X$, the points $(x, f(x))$ and $(y, f(y))$ lie in epi f. If $0 < \alpha < 1$, then $(\alpha x+(1-\alpha)y, \alpha f(x)+(1-\alpha)f(y))$ is in epi f if and only if f is convex. □

For differentiable and twice differentiable functions, convexity can be related to the first and second derivatives. For simplicity, consider an exteded real-valued function f defined on $\mathbb{R}^n$. Recall that the gradient of f is the vector of partial derivatives $\nabla f(x) = (f_{x_1}(x), \ldots, f_{x_n}(x))$, whereas the Hessian of f is the matrix of second partial derivatives $\nabla^2 f(x) = (f_{x_i x_j}(x))_{i,j=1}^n$.

Proposition 3.5 (§4 in [63]). *Let f be a real-valued function on $\mathbb{R}^n$.*

(i) If f is differentiable, then f is convex if and only if for every $x, y \in \mathbb{R}^n$,

$$f(y) \geq f(x) + \nabla f(x) \cdot (y - x). \tag{3.6}$$

(ii) If f is twice differentiable, then f is convex if and only if the Hessian $\nabla^2 f(x)$ is positive semi-definite for every $x \in \mathbb{R}^n$. That is, for every $x, y \in \mathbb{R}^n$,

$$y \cdot \left(\nabla^2 f(x)\, y\right) \geq 0. \tag{3.7}$$

Proof. Suppose that f is differentiable and satisfies the first-order condition (3.6). If we let $z = \alpha x + (1 - \alpha)y$, we have the following two inequalities

$$\begin{aligned} f(x) &\geq f(z) + \nabla f(z) \cdot (x - z) \\ f(y) &\geq f(z) + \nabla f(z) \cdot (y - z). \end{aligned}$$

A combination of the two inequalities gives

$$\alpha f(x) + (1 - \alpha) f(y) \geq f(z) + \nabla f(z) \cdot (\alpha(x - z) + (1 - \alpha)(y - z)) = f(z).$$

Suppose f is convex and differentiable and $0 \leq t \leq 1$, then

$$f(x + t(y - x)) \leq f(x) + t(f(y) - f(x)).$$

Dividing both sides by t and rearranging, we get

$$f(y) \geq f(x) + \frac{f(x + t(y - x)) - f(x)}{t}.$$

Letting $t \to 0$, the term on the right becomes a derivative with respect to t at $t = 0$ of the function $g(t) = f(x + t(y - x)) - f(x)$. The derivative is $g'(t) = \nabla f(x + t(y - x)) \cdot (y - x)$, and setting $t = 0$ gives the first-order condition (3.6).

The second-order condition (3.7) now follows from Taylor's theorem. □

This can be generalised further to hold in an infinite-dimensional vector space endowed with an inner product. Then the gradient and Hessian correspond to the first and second-order Fréchet derivatives, respectively.

The following proposition provides another useful way of identifying convex functions. Recall that a function $f : X \to \bar{\mathbb{R}}$ is said to be *subadditive* if for every $x, y \in X$ we have $f(x+y) \leq f(x) + f(y)$. Also, f is said to be *positively one-homogeneous* if $f(\alpha x) = \alpha f(x)$ for $\alpha > 0$ and $x \in X$.

Proposition 3.6. *Let $f \colon X \to \bar{\mathbb{R}}$ be subadditive and positively one-homogeneous, then f is convex.*

Proof. Let $x, y \in X$ and $0 < \alpha < 1$, then

$$f(\alpha x + (1-\alpha)y) \leq f(\alpha x) + f((1-\alpha)y) = \alpha f(x) + (1-\alpha) f(y). \qquad \square$$

More convex functions can be constructed through certain operations on functions which are already known to be convex. We mention a few examples here, and note that more are considered, e.g., in [63].

Proposition 3.7 (§5 in [63])**.**

(i) If f_1 and f_2 are proper convex functions, then $f_1 + f_2$ is convex.

(ii) If f is convex and $\lambda \geq 0$, then the function $(\lambda f)(x) = \lambda(f(x))$ is convex.

(iii) Let $f : X \to \bar{\mathbb{R}}$ be proper and convex, and suppose that φ is a convex function on $\mathbb{R}$ which is non-decreasing. Then $h(x) = \varphi(f(x))$ is convex on $\mathbb{R}^n$.

(iv) Let $\{f_i\}_{i \in I}$ be a non-empty family of convex functions. Then the pointwise supremum

$$g(x) := \sup_{i \in I} f_i(x)$$

is also convex.

Proof. These results all follow directly from the definition of convexity (3.2). We will show (iii), and note that the other are derived similarly. Let $x, y \in X$, and $0 < \alpha < 1$, then

$$\begin{aligned} h\left(\alpha x + (1-\alpha)y\right) &= \varphi\left(f\left(\alpha x + (1-\alpha)y\right)\right) \\ &\leq \varphi\left(\alpha f(x) + (1-\alpha) f(y)\right) \\ &\leq \alpha h(x) + (1-\alpha) h(y). \end{aligned}$$

$\square$

The following are some important examples of convex functions.

Example 3.8 (Some convex functions)**.**

(a) The Euclidean norm $|\cdot| : \mathbb{R}^n \to \mathbb{R}$, defined by $x \mapsto \sqrt{x \cdot x}$, is convex. This follows from Proposition 3.6, since $|\cdot|$ satisfies the triangle inequality (subadditivity) and is positively one-homogeneous

$$|x+y| \leq |x| + |y|, \qquad |\lambda x| = \lambda |x| \quad \text{for} \quad x, y \in \mathbb{R}^n,\ \lambda \geq 0.$$

In fact, Proposition 3.6 shows that every norm and semi-norm is convex.

(b) The function $f(x) = x^p$ defined on $\mathbb{R}$ is convex for $x \geq 0$ and $1 \leq p < \infty$, or $x > 0$ and $-\infty < p \leq 0$. It is concave whenever $x \geq 0$ and $0 \leq p \leq 1$. This follows from Proposition 3.5 and considering the sign of the second derivative $f''(x) = p(p-1)x^{p-2}$.

(c) The *indicator function* of a convex set C is defined by

$$\delta_C(x) := \begin{cases} 0 & \text{if } x \in C, \\ +\infty & \text{otherwise.} \end{cases} \tag{3.8}$$

The epigraph of $\delta_C(\cdot)$ is the set $C \times [0, +\infty)$, which is convex. Hence, Proposition 3.4 implies that $\delta_C(\cdot)$ is convex.

The following example will be useful later in Chapter 4. More specifically, it relates to the strict convexity of the data fidelity functional in our variational image restoration problems.

Example 3.9 (Some strictly convex functions)**.** Let X be a normed vector space with norm $\|\cdot\|$.

(a) The square of the norm $x \mapsto \|x\|^2$ is strictly convex. To see this, let $x, y \in X$ such that $x \neq y$ and $0 < \alpha < 1$. Then

$$\begin{aligned} \|\alpha x + (1-\alpha)y\|^2 &\leq \alpha^2 \|x\|^2 + (1-\alpha)^2 \|y\|^2 \\ &< \alpha \|x\|^2 + (1-\alpha) \|y\|^2 . \end{aligned}$$

(b) Let $z \in X$, then the mapping $x \mapsto \|x - z\|^2$ is also strictly convex,

$$\begin{aligned} \|\alpha x + (1-\alpha)y - z\|^2 &= \|\alpha(x-z) + (1-\alpha)(y-z)\|^2 \\ &< \alpha \|x - z\|^2 + (1-\alpha) \|y - z\|^2 . \end{aligned}$$

(c) Lastly, let $A : X \to X$ be a linear map. Provided that A is injective, i.e., $Ax \neq Ay$ for $x \neq y$, the function $x \mapsto \|Ax\|^2$ is strictly convex,

$$\|A(\alpha x + (1-\alpha)y)\|^2 = \|\alpha Ax + (1-\alpha)Ay\|^2 < \alpha \|Ax\|^2 + (1-\alpha) \|Ay\|^2 .$$

Existence of minimisers

We now outline the general strategy for showing the existence of a minimiser of a convex function. This relies on a standard argument from the calculus of variations. In Chapter 4 this is applied to the ROF model and other TV-based image restoration problems. More generally, necessary and sufficient conditions for the existence of minimisers is also related to the subdifferential calculus for convex functions, see [64].

One of the most important properties of convex functions, seen from the point of view of optimisation, is the fact that a local minimiser is also a global minimiser.

Theorem 3.10. *Let $f : X \to \bar{\mathbb{R}}$ be convex. Suppose that $x \in X$ is a local minimiser of f, i.e., for some $\epsilon > 0$, we have $f(x) \leq f(y)$ for every $y \in X$ such that $\|y - x\| \leq \epsilon$. Then $f(x) \leq f(y)$ for every $y \in X$.*

Proof. Let $y \in X$ and consider the convex combination $z = tx + (1-t)y$ for $0 < t < 1$. Then, for sufficiently large t, we have $\|z - x\| \leq \epsilon$, and, consequently

$$f(x) \leq f(tx + (1-t)y) \leq tf(x) + (1-t)f(y).$$

By rearranging, this shows that $f(x) \leq f(y)$. □

We begin with some definitions that are required to state the sufficient conditions for the existence of a minimiser of a convex function. The first is a topological property of functions that can be considered entirely separate from convexity. Let X denote a locally convex space.

Definition 3.11 (Lower semi-continuity)**.** *A function $f : X \to \bar{\mathbb{R}}$ is said to be* lower semi-continuous *on X if for every $x \in X$,*

$$f(x) \leq \liminf_{k\to\infty} f(x_k)$$

for every sequence $\{x_k\}_{k=1}^{\infty}$ such that $x_k \to x$ as $k \to \infty$.

The following proposition relates the lower semi-continuity of a function to its epigraph.

Proposition 3.12 (§7 in [63])**.** *A function $f : X \to \bar{\mathbb{R}}$ is lower semi-continuous if and only if its epigraph is closed.*

Proof. Suppose f is lower semi-continuous and let $\{(x_k, t_k)\}_{k=1}^{\infty} \subset \operatorname{epi} f$ be a sequence converging to (x, t) as $k \to \infty$. From the lower semi-continuity and the definition of the epigraph, we have

$$f(x) \leq \liminf_{k\to\infty} f(x_k) \leq \liminf_{k\to\infty} t_k = t.$$

Since $(x, t) \in \operatorname{epi} f$, the epigraph of f is closed.

Suppose epi f is closed and consider the sequence $\{(x_k, f(x_k))\}_{k=1}^{\infty} \subset \operatorname{epi} f$. Since epi f is closed, the sequence has a limit (x, t) in epi f. Then,

$$f(x) \leq t = \liminf_{k\to\infty} f(x_k).$$ □

For every function $f : X \to \bar{\mathbb{R}}$, there exists a greatest lower semi-continuous function $h \leq f$. The function h is called the *lower semi-continuous hull* of f, denoted $\operatorname{lsc} f$. Furthermore, $\operatorname{lsc} f$ coincides with the function whose epigraph is the closure in $X \times \mathbb{R}$ of the epigraph of f.

The following, closely related definition is useful in order to avoid some uninteresting and degenerate cases.

Definition 3.13 (Closure)**.** *Let $f : X \to \bar{\mathbb{R}}$. The* closure $\operatorname{cl} f$ *of f is the lower semi-continuous hull of f, unless f attains the value $-\infty$ at some point, in which case the closure is identically equal to $-\infty$. Furthermore, f is said to be* closed *if $\operatorname{cl} f = f$.*

Thus, a proper convex function is closed if and only if it is lower semi-continuous. Further, every closed convex function f is the pointwise supremum of the collection of all affine functions h such that $h \leq f$ (see Theorem 12.1 in [63]). Geometrically, this comes from the fact that epi f is the intersection

of the closed half-spaces in $X \times \mathbb{R}$ that contain it. These closed half-spaces are just the epigraphs of the affine functions $h \leq f$, and their intersection is the epigraph of the pointwise supremum of such functions h.

The next definition describes another sufficient condition that is necessary for the existence of a minimiser. Let X, Y denote normed vector spaces.

Definition 3.14 (Coercivity)**.** *Suppose that $f : X \to \bar{\mathbb{R}}$ and $\operatorname{dom} f \subset Y$. We say that f is* coercive *on Y if there exist constants $C_1 > 0$ and $C_2 \geq 0$ such that*

$$f(x) \geq C_1 \|x\|_Y - C_2, \quad \forall x \in \operatorname{dom} f. \tag{3.9}$$

Finally, the following theorem states the sufficient conditions for the existence and the uniqueness of a minimiser.

Theorem 3.15 (Existence of minimiser)**.** *Suppose that $f : X \to \bar{\mathbb{R}}$ is a proper convex function and $\operatorname{dom} f \subset Y$. Then f has a minimiser provided that the following conditions hold*

(i) f is lower semi-continuous,

(ii) f is coercive on Y,

(iii) Bounded sets in Y are relatively sequentially compact in X.

The solution is unique if f is strictly convex.

Proof.

1. Define $m := \inf_{x \in X} f(x)$. Take a minimising sequence $\{x_k\}_{k=1}^{\infty} \subset Y$ such that $f(x_k) \to m$ as $k \to \infty$.

2. Observe that $\{f(x_k)\}$ is a convergent sequence in $\mathbb{R}$, hence it is bounded. Together with the coercivity condition, this implies that $\{x_k\}$ is bounded uniformly in Y.

3. From the relative sequential compactness of bounded sets in Y, there exists $x \in Y$ and a subsequence $\{x_{k_j}\} \subset \{x_k\}$ such that $x_{k_j} \to x$ in X.

4. From the lower semi-continuity of f, we have that

$$f(x) \leq \liminf_{j \to \infty} f(x_{k_j}) = m.$$

   Moreover, since $m \leq f(x)$ by definition then $f(x) = m$. Thus, x is a minimiser of f.

5. To show uniqueness, let x, y be two minimisers of f such that $x \neq y$. From the strict convexity of f we have

$$f\left(\frac{x+y}{2}\right) < \frac{1}{2} f(x) + \frac{1}{2} f(y) = m.$$

   This contradicts the assumption that x and y are two different minimisers of f. □

Usually, the above theorem is used in the context of a reflexive Banach space X supplied with the weak topology. Then, from the Banach-Alaoglu theorem, the closed unit ball in X is weakly sequentially compact. However, the particular function space that arises in connection with the total variation-based methods in Chapter 4, is not reflexive. Although, as we will see in Section 3.2, it does have a compactness property that allows the above theorem to be utilised.

3.2 Functions of bounded variation

Now we turn to the second half of the mathematical background material that is presented in this chapter. In this section, we introduce the total variation. This is the regularisation functional that is used in the ROF model and other image restoration problems in Chapter 4. Moreover, it is associated with a certain Banach space of functions of bounded variation. This function space has previously been studied in connection with minimal surfaces [32], and the generalised solutions of other partial differential equations such as hyperbolic conservation laws [44]. It is necessary for us to study this particular function space in order to characterise the solutions to the ROF model.

Most of the results in this section may seem fairly technical, but the aim is to establish the sufficient conditions required in Theorem 3.15 for minimisation problems involving the total variation functional. For a more complete reference, see Evans and Gariepy [31], upon which most of this material is based. Some of the standard facts from Sobolev space theory that are used in the proofs can be found, e.g., in Evans [30].

Let Ω denote an open subset of $\mathbb{R}^n$. Recall that the Euclidean norm in $\mathbb{R}^n$ is denoted by $|\cdot|$.

Definition 3.16 (Total variation for differentiable functions)**.** *Let f be a real-valued continuously differentiable function on Ω. The* total variation (TV) *of f is defined by*

$$TV(f) := \int_\Omega |\nabla f(x)| \, dx. \tag{3.10}$$

For a differentiable function $f \in C^1(\Omega)$ the total variation is exactly the norm of the gradient ∇f in $L^1(\Omega; \mathbb{R}^n)$. The above definition holds also for functions in the Sobolev space $W^{1,1}(\Omega)$ if ∇f is interpreted in terms of the weak partial derivatives of f. Even more generally, e.g., if f is only locally integrable, then the gradient must be interpreted in the sense of distributions (and it can be shown that the gradient of a function with finite TV is a vector-valued Radon measure, see [31]).

It is possible to derive an equivalent definition of the total variation that holds also for integrable functions without mentioning distributions.

Definition 3.17 (Total variation)**.** *For $f \in L^1(\Omega)$ the* total variation (TV) *of f is defined by*

$$TV(f) := \sup_{\boldsymbol{\varphi} \in \mathcal{V}} -\int_\Omega f \operatorname{div} \boldsymbol{\varphi}, \tag{3.11}$$

where $\mathcal{V}$ is the set

$$\mathcal{V} := \left\{ \boldsymbol{\varphi} \in C_c^1(\Omega; \mathbb{R}^n), |\boldsymbol{\varphi}(x)| \leq 1 \text{ for every } x \in \Omega \right\}. \tag{3.12}$$

The following proposition shows that this is in fact a reasonable generalisation of the total variation.

Proposition 3.18. *Let $\Omega \subset \mathbb{R}^n$ be open and bounded and $f \in C^1(\Omega)$, then*

$$\int_\Omega |\nabla f| = \sup_{\boldsymbol{\varphi} \in \mathcal{V}} -\int_\Omega f \operatorname{div} \boldsymbol{\varphi}.$$

Proof. First, observe that from integration by parts and the Cauchy-Schwarz inequality we have

$$\begin{aligned}\sup_{\boldsymbol{\varphi} \in \mathcal{V}} -\int_\Omega f(x) \operatorname{div} \boldsymbol{\varphi}(x)\, dx &= \sup_{\boldsymbol{\varphi} \in \mathcal{V}} \int_\Omega \nabla f(x) \cdot \boldsymbol{\varphi}(x)\, dx \\ &\leq \sup_{\boldsymbol{\varphi} \in \mathcal{V}} \int_\Omega |\nabla f(x)|\, |\boldsymbol{\varphi}(x)|\, dx \\ &= \int_\Omega |\nabla f(x)|\, dx.\end{aligned}$$

Second, we have that

$$\int_\Omega |\nabla f| = \int_\Omega \nabla f \cdot \boldsymbol{\nu}, \qquad \boldsymbol{\nu}(x) := \begin{cases} \frac{\nabla f(x)}{|\nabla f(x)|}, & \text{if } \nabla f(x) \neq 0, \\ 0, & \text{otherwise.} \end{cases}.$$

The function $\boldsymbol{\nu}(x)$ is the unit normal to the level lines of f whenever the gradient of f is non-zero. Observe that $\|\boldsymbol{\nu}(x)\|_{L^\infty} \leq 1$. Moreover, because Ω is bounded it follows from Hölder's inequality that $\boldsymbol{\nu} \in L^p(\Omega; \mathbb{R}^n)$ for $1 \leq p \leq \infty$. Finally, recall that the set $C_c^1(\Omega; \mathbb{R}^n)$ is dense in $L^p(\Omega; \mathbb{R}^n)$, so that

$$\int_\Omega |\nabla f| = \int_\Omega \nabla f \cdot \boldsymbol{\nu} \leq \sup_{\boldsymbol{\varphi} \in \mathcal{V}} \int_\Omega \nabla f \cdot \boldsymbol{\varphi} = \sup_{\boldsymbol{\varphi} \in \mathcal{V}} -\int_\Omega f \operatorname{div} \boldsymbol{\varphi}. \qquad \square$$

Later, we will be working with minimisation problems involving the total variation. In that context, it is appropriate to regard $TV(\cdot)$ as an extended real-valued functional on $L^p(\Omega)$ for some $1 \leq p \leq \infty$. This makes sense whenever Ω is bounded, because it follows from Hölder's inequality that $L^p(\Omega) \subset L^1(\Omega)$ for $p > 1$. With this in mind, we summarise some important properties of this functional in the next result.

Theorem 3.19. *Let $\Omega \subset \mathbb{R}^n$ be open and bounded. Then $TV(\cdot)$ is a proper, closed convex function on $L^p(\Omega)$ for $1 \leq p \leq \infty$.*

Proof. First, $TV(\cdot)$ is non-negative, so it is nowhere equal to $-\infty$, and not identically equal to $+\infty$ (take, e.g., $f = 0$). Hence, it is proper.

Second, $TV(\cdot)$ is subadditive,

$$\begin{aligned}TV(f+g) &= \sup_{\boldsymbol{\varphi}} -\int_\Omega (f+g) \operatorname{div} \boldsymbol{\varphi}\, dx \\ &\leq \sup_{\boldsymbol{\varphi}} -\int_\Omega f \operatorname{div} \boldsymbol{\varphi}\, dx + \sup_{\boldsymbol{\varphi}} -\int_\Omega g \operatorname{div} \boldsymbol{\varphi}\, dx \\ &= TV(f) + TV(g).\end{aligned}$$

In addition, it is positively one-homogeneous, i.e., for $\alpha > 0$,

$$TV\left(\alpha f\right) = \sup_{\boldsymbol{\varphi}} - \int_\Omega \alpha f \operatorname{div} \boldsymbol{\varphi}\, dx = \alpha \sup_{\boldsymbol{\varphi}} - \int_\Omega f \operatorname{div} \boldsymbol{\varphi}\, dx = \alpha TV\left(f\right).$$

Then it follows from Proposition 3.6 that $TV(\cdot)$ is convex.

To see that $TV(\cdot)$ is lower semi-continuous, let $\{f_k\}_{k=1}^\infty$ be a sequence such that $f_k \to f$ in $L^p(\Omega)$. From Fatou's lemma we have

$$TV(f) = \sup_{\boldsymbol{\varphi}\in\mathcal{V}} - \int_\Omega f \operatorname{div} \boldsymbol{\varphi}\, dx \leq \liminf_{k\to\infty} \sup_{\boldsymbol{\varphi}\in\mathcal{V}} - \int_\Omega f_k \operatorname{div} \boldsymbol{\varphi}\, dx. \qquad \square$$

The effective domain of the TV functional is a subspace of $L^1(\Omega)$. This defines a Banach space whenever it is equipped with an appropriate norm.

Definition 3.20 (Bounded variation)**.** *We say that $f \in L^1(\Omega)$ has* bounded variation *provided that its total variation is finite. The space of functions of bounded variation is*

$$BV(\Omega) := \left\{ f \in L^1(\Omega) \mid TV\left(f\right) < +\infty \right\}.$$

This is a Banach space equipped with the norm

$$\|f\|_{BV(\Omega)} := \|f\|_{L^1(\Omega)} + TV\left(f\right). \tag{3.13}$$

It is easy to see that $BV(\Omega)$ is a vector subspace of $L^1(\Omega)$ and that (3.13) is in fact a norm. Completeness follows from the fact that $L^1(\Omega)$ is complete and the lower semi-continuity of $TV(\cdot)$. Notice also that for $f \in W^{1,1}(\Omega)$, the above norm is exactly the norm on $W^{1,1}(\Omega)$.

The remainder of this section is devoted to studying properties of the space $BV(\Omega)$. In particular, we show a certain compactness property that is required in Theorem 3.15. We also include a Poincaré inequality that is later used to show a coercivity condition on functionals involving $TV(\cdot)$.

Approximation and compactness

Since bounded variation functions can be hard to study directly, most of their properties are instead developed through approximations by smooth functions. These approximations are possible due to the following result.

Theorem 3.21 (Approximation by smooth functions [31])**.** *Let $f \in BV(\Omega)$. There exist functions $\{f_k\}_{k=1}^\infty \subset BV(\Omega) \cap C^\infty(\Omega)$ such that as $k \to \infty$*

(i) $f_k \to f$ in $L^1(\Omega)$, and

(ii) $TV\left(f_k\right) \to TV\left(f\right)$.

Proof. See Section 5.2 in [31]. $\square$

We now consider the relation between functions of bounded variation and certain other L^p-spaces. These results are extensions of standard facts from Sobolev space theory.

Recall that the Gagliardo–Nirenberg–Sobolev inequality (see [30]) states that for $1 \leq p < n$ the Sobolev space $W^{1,p}(\mathbb{R}^n)$ is continuously embedded in

$L^{p^*}(\mathbb{R}^n)$, where $p^* = np/(n-p)$ is the Sobolev conjugate of p. In the case $p = 1$, this can be extended to show that $BV(\mathbb{R}^n)$ is continuously embedded in $L^{n/(n-1)}(\mathbb{R}^n)$, see Section 5.6 in [31]. The case $n = 1$ is special in that every function in $W^{1,1}(I)$ for some interval $I = (a, b)$ is absolutely continuous, and therefore also in $L^\infty(I)$.

We will show a similar result, but for bounded domains $\Omega \subset \mathbb{R}^n$. This requires some additional conditions on the boundary $\partial\Omega$. Roughly speaking, we say that $\partial\Omega$ is C^1 if it is locally the graph of a continuously differentiable function, see [30].

Theorem 3.22 (Embedding [31])**.** *Let $\Omega \subset \mathbb{R}^n$ be open and bounded with C^1 boundary $\partial\Omega$. The space $BV(\Omega)$ is continuously embedded in $L^q(\Omega)$ for $1 \le q \le n/(n-1)$.*

Proof. From Theorem 3.21, let $\{f_k\}_{k=1}^\infty \subset BV(\Omega) \cap C^\infty(\Omega)$ be a sequence such that $f_k \to f$ in $L^1(\Omega)$ and $TV(f_k) \to TV(f)$. Observe that $f_k \in W^{1,1}(\Omega)$. As a consequence of the Gagliardo–Nirenberg–Sobolev inequality, there exists $C > 0$ such that

$$\|f_k\|_{L^{\frac{n}{n-1}}(\Omega)} \le C \|f_k\|_{W^{1,1}(\Omega)} = C \|f_k\|_{BV(\Omega)}.$$

Taking limits as $k \to \infty$ and using Fatou's lemma on the left-hand side gives

$$\left(\int_\Omega |f|^{\frac{n}{n-1}}\right)^{\frac{n-1}{n}} \le \liminf_{k\to\infty} \|f_k\|_{L^{\frac{n}{n-1}}(\Omega)} \le \lim_{k\to\infty} C \|f_k\|_{BV(\Omega)} = C \|f\|_{BV(\Omega)}.$$

Whenever Ω is bounded, it follows from Hölder's inequality that $L^{n/(n-1)}(\Omega)$ is continuously embedded in $L^q(\Omega)$ for $1 \le q \le n/(n-1)$. This completes the proof. □

The next result shows that, except for the case $p = n/(n-1)$, the above embedding is, in fact, compact. This is based on the Rellich–Kondrachov compactness theorem which states that the Sobolev space $W^{1,p}(\Omega)$ is compactly embedded in $L^q(\Omega)$ for $1 \le q < p^*$.

Theorem 3.23 (Compactness [31])**.** *Let $\Omega \subset \mathbb{R}^n$ be open and bounded with C^1 boundary $\partial\Omega$. Then the embedding $BV(\Omega) \subset L^q(\Omega)$ for $1 \le q < n/(n-1)$ is compact.*

Proof. Let $\{f_k\}_{k=1}^\infty$ be a sequence uniformly bounded in $BV(\Omega)$. That is,

$$\sup_k \|f_k\|_{BV(\Omega)} < \infty.$$

Using Theorem 3.21, construct an approximating sequence $\{g_k\}_{k=1}^\infty \subset C^\infty(\Omega)$ such that

$$\|f_k - g_k\|_{L^1(\Omega)} \le 1/k,$$
$$\sup_k TV(g_k) < \infty.$$

Observe that the sequence $\{g_k\}_{k=1}^{\infty}$ is bounded in $W^{1,1}(\Omega)$,

$$\begin{aligned}\|g_k\|_{W^{1,1}(\Omega)} = \|g_k\|_{BV(\Omega)} &\leq \|f_k\|_{BV(\Omega)} + \|f_k - g_k\|_{BV(\Omega)} \\ &\leq \|f_k\|_{BV(\Omega)} + \|f_k - g_k\|_{L^1(\Omega)} + TV(f_k) + TV(g_k) \\ &< \infty.\end{aligned}$$

By the Rellich–Kondrachov compactness theorem, there exists $f \in L^1(\Omega)$ and a subsequence $\left\{g_{k_j}\right\}_{j=1}^{\infty}$ such that $g_{k_j} \to f$ in $L^1(\Omega)$ as $j \to \infty$. Moreover, from the lower semi-continuity of $TV(\cdot)$ it follows that $f \in BV(\Omega)$.

Now, considering an appropriate subsequence of $\{f_k\}_{k=1}^{\infty}$, we find that

$$\left\|f_{k_j} - f\right\|_{L^1(\Omega)} \leq \left\|f_{k_j} - g_{k_j}\right\|_{L^1(\Omega)} + \left\|g_{k_j} - f\right\|_{L^1(\Omega)} \to 0$$

as $j \to \infty$. This shows that the embedding $BV(\Omega) \subset L^1(\Omega)$ is compact.

Finally, let $1 \leq q < n/(n-1)$. From the embedding $BV(\Omega) \subset L^{n/(n-1)}(\Omega)$, we have

$$\left\|f_{k_j} - f\right\|_{L^{\frac{n}{n-1}}(\Omega)} \leq \left\|f_{k_j} - f\right\|_{BV(\Omega)} \leq \left\|f_{k_j}\right\|_{BV(\Omega)} + \|f\|_{BV(\Omega)} < \infty.$$

Recall the interpolation inequality

$$\left\|f_{k_j} - f\right\|_{L^q(\Omega)} \leq \left\|f_{k_j} - f\right\|_{L^1(\Omega)}^{\theta} \left\|f_{k_j} - f\right\|_{L^{\frac{n}{n-1}}(\Omega)}^{1-\theta},$$

where $0 < \theta < 1$ and $1/q = \theta + (1-\theta)(n-1)/n$. Combining these two facts proves the theorem. □

The third result we include in this section is a Poincaré–Wirtinger inequality for $BV(\Omega)$. Let $\bar{f}$ denote the average value of f, i.e., $\bar{f} := \frac{1}{|\Omega|}\int_\Omega f(x)\,dx$. The proof follows precisely the same argument as in Theorem 3.22.

Theorem 3.24 (Poincaré–Wirtinger inequality [31])**.** *Let Ω be a bounded, open and connected subset of $\mathbb{R}^n$ with a C^1 boundary $\partial\Omega$. Suppose that $f \in BV(\Omega)$. Then there is a constant $C > 0$ such that*

$$\left\|f - \bar{f}\right\|_{L^{\frac{n}{n-1}}(\Omega)} \leq C\, TV\,(f)\,.$$

Proof. Let $f \in BV(\Omega)$ and take a sequence $\{f_k\}_{k=1}^{\infty} \subset BV(\Omega) \cap C^\infty(\Omega)$ such that $f_k \to f$ in $L^1(\Omega)$ and $TV\,(f_k) \to TV\,(f)$ as $k \to \infty$. Clearly, $f_k \in W^{1,1}(\Omega)$ and the Poincaré–Wirtinger inequality [30] states that

$$\left\|f_k - \bar{f}_k\right\|_{L^{\frac{n}{n-1}}(\Omega)} \leq C\,\|\nabla f_k\|_{L^1(\Omega;\mathbb{R}^n)} = C\,TV\,(f_k)\,,$$

for some $C > 0$. Then, from Fatou's lemma we have

$$\begin{aligned}\left\|f - \bar{f}\right\|_{L^{\frac{n}{n-1}}(\Omega)} &\leq \liminf_{k\to\infty} \left\|f_k - \bar{f}_k\right\|_{L^{\frac{n}{n-1}}(\Omega)} \\ &\leq \lim_{k\to\infty} C\,TV(f_k) \\ &= C\,TV(f).\end{aligned}$$

□

The compactness that was established in Theorem 3.23 does not hold for the embedding $BV(\Omega) \subset L^{n/(n-1)}(\Omega)$. However, using the Poincaré–Wirtinger inequality, we can show that this embedding is compact in the weak topology.

Theorem 3.25 (Weak compactness [1])**.** *Let Ω be a bounded, open and connected subset of $\mathbb{R}^n$ for $n \geq 2$ with a C^1 boundary $\partial\Omega$. Then the embedding $BV(\Omega) \subset L^{n/(n-1)}(\Omega)$ is weakly compact.*

Proof. Let $\{f_k\}_{k=1}^{\infty}$ be a bounded sequence in $BV(\Omega)$. From the Poincaré–Wirtinger inequality, the sequence $\{f_k\}$ is also bounded in $L^{n/(n-1)}(\Omega)$, as follows

$$\begin{aligned}
\|f_k\|_{L^{\frac{n}{n-1}}(\Omega)} &\leq \left\|f_k - \bar{f}_k\right\|_{L^{\frac{n}{n-1}}(\Omega)} + \left\|\chi_\Omega \bar{f}_k\right\|_{L^{\frac{n}{n-1}}(\Omega)} \\
&\leq |f_k|\,|\Omega|^{(n-1)/n} + C\,TV(f_k) \\
&\leq \|f_k\|_{L^1(\Omega)}\,|\Omega|^{(n-1)/n-1} + C\,\|f_k\|_{BV(\Omega)} \\
&\leq \left(|\Omega|^{-1/n} + C\right)\|f_k\|_{BV(\Omega)}\,.
\end{aligned}$$

Since the space $L^{n/(n-1)}(\Omega)$ is reflexive whenever $n \geq 2$, it follows from the Banach–Alaoglu theorem that there is a subsequence $\left\{f_{k_j}\right\}_{j=1}^{\infty}$ converging weakly to $f \in L^{n/(n-1)}(\Omega)$. □

3.3 Summary

In this chapter we have introduced two of the main tools that will be used to study the variational image restoration methods in Chapter 4. The first of these is convex analysis, which, roughly speaking, is concerned with optimisation problems where the objective function is convex. This is used in the next chapter to answer questions regarding the existence and uniqueness of solutions to the variational models for image restoration. The second important concept that has been introduced is the total variation, a proper, closed and convex functional that is central to the above mentioned variational methods.

CHAPTER 4

Total variation-based methods

The restoration of degraded images is, generally, an ill-posed inverse problem. The approach taken here is to instead consider the regularised problem, thereby imposing additional constraints which, hopefully, turn the problem into a well-posed one. In this chapter, we present the ROF model [65], a regularisation based on the total variation (TV) from Section 3.2. This particular regularisation favours piecewise constant solutions, and, therefore, it is often able to recover the edges of the original image.

First, we present the ROF model for restoring a blurred and noisy grayscale image. We provide a proof of the existence and uniqueness of the solution using the material from Chapter 3, and we give some simple examples that illustrate the properties of this model. Next, we consider the generalisation of this model to the restoration of colour images. As a particular special case of this, we regard a variational model by Shen and Chan [72] for image inpainting. Finally, we present an original TV-based model for the demosaicking problem. This is based on decomposing the image into luminance and chrominance components and applying a TV regularisation to each of them.

4.1 The Rudin–Osher–Fatemi model

For convenience, we recall the forward image model (1.1) from Chapter 1. We presently assume that we are working with scalar images defined on an open and bounded set $\Omega \subset \mathbb{R}^n$ with a C^1 boundary $\partial\Omega$.

Consider an image $u \in L^p(\Omega)$, for $1 \leq p \leq \infty$, that is subject to the following transformation

$$f = Au + \eta. \tag{4.1}$$

The observed image $f \in L^2(\Omega)$ is the result of the original image u undergoing a linear transformation $A : L^p(\Omega) \to L^2(\Omega)$ and being corrupted by noise η. The exact nature of the noise and the linear transformation A were discussed for various cases in the context of digital colour cameras in Chapter 2. Here we

assume that the noise is additive white Gaussian-distributed with zero mean and known variance. For the pure image denoising problem, the operator A is just the identity operator. Otherwise, it is usually an integral operator, e.g., a convolution by some point-spread function.

The *Rudin–Osher–Fatemi (ROF) model* is to solve the minimisation problem

$$\min_{u \in L^p(\Omega)} \left\{ F(u) := TV(u) + \frac{\lambda}{2} \|Au - f\|^2_{L^2(\Omega)} \right\}. \tag{4.2}$$

The functional $F : L^p(\Omega) \to \bar{\mathbb{R}}$ is called the *ROF functional*. The regularisation term is exactly the total variation $TV(\cdot)$ defined in Section 3.2, and the effective domain of F is $BV(\Omega)$. The data fidelity functional is a least-squares term where the parameter $\lambda > 0$ balances the regularity of the solution against the fit to the data. Usually, this parameter can be related to the variance of the noise (see [16]).

Existence and uniqueness

In this section, we consider the existence and uniqueness of the solution to the ROF model. These results were shown by Acar and Vogel in [1] and Chambolle and Lions in [16]. Our approach closely follows [1], and uses the material from Chapter 3.

First, we note some properties of the ROF functional.

Lemma 4.1. *The ROF functional F, defined in* (4.2)*, is*

(i) proper, closed and convex, and

(ii) strictly convex if A is injective.

Proof. Clearly, F is proper because it is non-negative and not identically equal to $+\infty$ (take, e.g., $u = 0$). Moreover, TV and the square of the L^2-norm are both lower semi-continuous (a consequence of Theorem 3.19 and Fatou's lemma, respectively), and so F is lower semi-continuous. Similarly, TV and the data fidelity term are both convex. In fact, the latter is strictly convex provided that A is injective, see Example 3.9. □

In view of Theorem 3.15, a coercivity condition is required on the ROF functional F. Such a condition is provided by the following lemma. The proof is due to Acar and Vogel [1].

Lemma 4.2 (Coercivity [1])**.** *Let $1 \leq p \leq n/(n-1)$. Then the ROF functional $F : L^p(\Omega) \to \bar{\mathbb{R}}$ is coercive on $BV(\Omega)$ provided that the operator A does not annihilate constant functions, that is, $A\chi_\Omega \neq 0$.*

Proof.

1. Every $u \in BV(\Omega)$ can be decomposed uniquely as a sum

$$u = v + w, \quad \text{where} \quad w = \left(\frac{1}{|\Omega|} \int_\Omega u \, dx \right) \chi_\Omega \quad \text{and} \quad \int_\Omega v \, dx = 0.$$

First, notice that $TV(u) = TV(v)$. From this decomposition, we have

$$\begin{aligned}\|u\|_{BV(\Omega)} &= \|u\|_{L^1(\Omega)} + TV(u) \\ &= \|v+w\|_{L^1(\Omega)} + TV(v+w) \\ &\leq \|v\|_{L^1(\Omega)} + \|w\|_{L^1(\Omega)} + TV(v).\end{aligned}$$

From Hölder's inequality and the Poincaré inequality (Theorem 3.24), there is a constant $C_1 > 0$ such that,

$$\|v\|_{L^p(\Omega)} \leq |\Omega|^{\frac{1}{p}-\frac{n-1}{n}} \|v\|_{L^{\frac{n}{n-1}}(\Omega)} \leq C_1\, TV(v).$$

Hence, we have

$$\|u\|_{BV(\Omega)} \leq \|w\|_{L^1(\Omega)} + \left(C_1 |\Omega|^{1-1/p} + 1\right) TV(v). \qquad (\star)$$

2. Consider the data fidelity term of the ROF functional. From the decomposition $u = v + w$, we find that

$$\begin{aligned}\|Au - f\|^2_{L^2(\Omega)} &= \|Av - f + Aw\|^2_{L^2(\Omega)} \\ &\geq \left(\|Av - f\|_{L^2(\Omega)} - \|Aw\|_{L^2(\Omega)}\right)^2 \\ &= \|Av - f\|^2_{L^2(\Omega)} + \|Aw\|_{L^2(\Omega)} \left(\|Aw\|_{L^2(\Omega)} - 2\|Av - f\|_{L^2(\Omega)}\right) \\ &\geq \|Aw\|_{L^2(\Omega)} \left(\|Aw\|_{L^2(\Omega)} - 2\|Av - f\|_{L^2(\Omega)}\right). \qquad (\star\star)\end{aligned}$$

Recall that the operator A is continuous, and thus the operator norm

$$\|A\| = \sup_{\|u\|_{L^p(\Omega)} \leq 1} \|Au\|_{L^2(\Omega)}$$

is bounded. Therefore,

$$\begin{aligned}\|Av - f\|_{L^2(\Omega)} &\leq \|Av\|_{L^2(\Omega)} + \|f\|_{L^2(\Omega)} \\ &\leq \|A\|\, \|v\|_{L^p(\Omega)} + \|f\|_{L^2(\Omega)} \\ &\leq \|A\|\, C_1\, TV(v) + \|f\|_{L^2(\Omega)}.\end{aligned}$$

Also, since w is a constant function and $Aw \neq 0$, there exists a constant $C_2 > 0$ such that

$$\|Aw\|_{L^2(\Omega)} = C_2 \|w\|_{L^1(\Omega)}.$$

Combining these facts and inserting into $(\star\star)$, we see that

$$\begin{aligned}\|Au - f\|^2_{L^2(\Omega)} &\geq \|Aw\|_{L^2(\Omega)} \left(\|Aw\|_{L^2(\Omega)} - 2\|Av - f\|_{L^2(\Omega)}\right) \\ &\geq C_2 \|w\|_{L^1(\Omega)} \left(C_2 \|w\|_{L^1(\Omega)} - 2\|A\|\, C_1\, TV(v) - 2\|f\|_{L^2(\Omega)}\right).\end{aligned}$$

3. Thus, we have the following bound on the ROF functional

$$\begin{aligned}F(u) &= TV(u) + \frac{\lambda}{2}\|Au - f\|^2_{L^2(\Omega)} \\ &\geq TV(v) + \frac{\lambda}{2} C_2 \|w\|_{L^1(\Omega)} \left(C_2 \|w\|_{L^1(\Omega)} - 2\|A\|\, C_1\, TV(v) - 2\|f\|_{L^2(\Omega)}\right).\end{aligned}$$

Suppose that $\left(C_2 \|w\|_{L^1(\Omega)} - 2\|A\| C_1 TV(v) - 2\|f\|_{L^2(\Omega)}\right) \geq 1$, then

$$F(u) \geq \frac{\lambda}{2} C_2 \|w\|_{L^1(\Omega)}.$$

Finally, from $(\star)$, it follows that

$$\begin{aligned}\|u\|_{BV(\Omega)} &\leq \|w\|_{L^1(\Omega)} + (C_1 |\Omega|^{1-1/p} + 1) TV(v) \\ &\leq \left(\frac{2}{\lambda C_2} + C_1 |\Omega|^{1-1/p} + 1\right) F(u).\end{aligned}$$

Otherwise, we have that

$$\|w\|_{L^1(\Omega)} < \frac{1 + 2\|A\| C_1 TV(v) + 2\|f\|_{L^2(\Omega)}}{C_2},$$

and thereby that

$$\|u\|_{BV(\Omega)} - \frac{1 + 2\|f\|_{L^2(\Omega)}}{C_2} \leq \left(\frac{2\|A\| C_1}{C_2} + C_1 |\Omega|^{1-1/p} + 1\right) F(u).$$

Hence, $\|u\|_{BV(\Omega)} \to \infty$ implies that $F(u) \to \infty$, which shows that F is coercive. □

Finally, we invoke Theorem 3.15 to prove the existence and uniqueness of the solution to the ROF model.

Theorem 4.3 (Existence-uniqueness)**.** *Let $\Omega \subset \mathbb{R}^n$ be an open and bounded domain with a C^1 boundary $\partial\Omega$. Suppose that $1 \leq p \leq n/(n-1)$. Then there exists a solution to the ROF model* (4.2)*, provided that the operator A does not annihilate constant functions. Moreover, the solution is unique if A is injective.*

Proof. First, from Lemma 4.1, the ROF functional F is a proper, closed convex function on $L^p(\Omega)$, strictly convex if A is injective. Second, $\operatorname{dom} F = BV(\Omega)$, and F is coercive on $BV(\Omega)$ by Lemma 4.2. From Theorem 3.23, bounded sets in $BV(\Omega)$ are relatively compact in $L^p(\Omega)$ whenever $1 \leq p < n/(n-1)$.

For the case $p = n/(n-1)$, Theorem 3.25 states that bounded sets in $BV(\Omega)$ are compact in $L^{n/(n-1)}(\Omega)$ with respect to the weak topology on $L^{n/(n-1)}(\Omega)$. From Mazur's lemma [30], a closed convex subset of a Banach space X is weakly closed, i.e., closed with respect to the weak topology on X. Since epi F is closed and convex, it is also weakly closed. In other words, F is lower semi-continuous with respect to the weak topology on $L^{n/(n-1)}(\Omega)$.

The result now follows from Theorem 3.15. □

Properties of the solution

The ROF model is particularly effective for restoring "blocky" or piecewise constant images. Although explicit solutions are, in general, not available, Strong and Chan [73] have given the exact solution in a few simple cases, i.e., for radially symmetric, piecewise constant functions. Their result shows that the ROF model tends to preserve the locations of edges and that it can have quite local effects, despite the fact that (4.2) is a global problem.

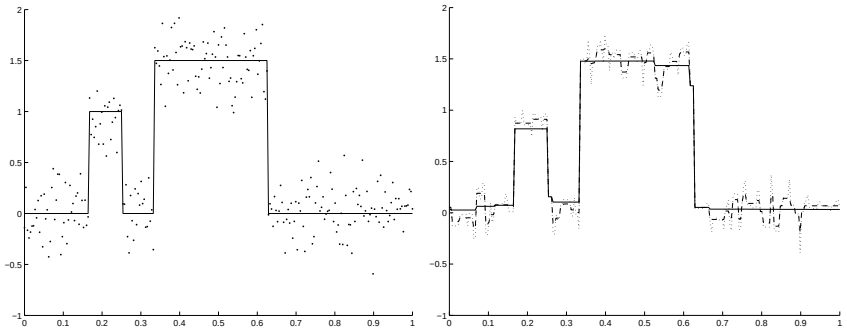

(a) Original and noisy signal (b) TV restoration for different λ

Figure 4.1: *TV regularisation. The figure displays the restoration of a noisy signal with the ROF model* (4.2)*, and the effect of various choices for the regularisation parameter λ. Figure (a) shows a piecewise constant signal and some measurements corrupted by Gaussian noise. Figure (b) shows the reconstruction of the noisy signal with $\lambda = 1$ (solid line), $\lambda = 5$ (dashed line) and $\lambda = 10$ (dotted line). Smaller values of λ smooth out oscillations and merge small-scale features.*

Essentially, the regularisation works by reducing the contrast between individual features in an image. The change in intensity of a piecewise constant region is inversely proportional to the regularisation parameter λ and the scale of the feature. Noise, which may be regarded as very small-scale features, is therefore effectively removed while larger features are preserved.

Example 4.4 (The regularisation parameter λ)**.** Consider the function on the unit interval defined by

$$u = \chi_{[1/6,1/4]} + \frac{3}{2}\chi_{[1/3,5/8]},$$

where χ_I is the indicator function of the set $I \subset \mathbb{R}$, i.e., $\chi_I(x) = 1$ if $x \in I$ and $\chi_I(x) = 0$ otherwise.

We sample the function at $N = 257$ uniformly spaced points and add some Gaussian noise. Denote the original signal by $u \in \mathbb{R}^N$ and the noisy signal by $f \in \mathbb{R}^N$. These are shown in Figure 4.1(a). The amount of noise added corresponds to a signal-to-noise ratio (SNR) of 10. The SNR is given by

$$SNR := \frac{\sqrt{\sum_{j=1}^N (u_j - \bar{u})^2}}{\sqrt{\sum_{j=1}^N (f_j - u_j)^2}}, \qquad \bar{u} = \frac{1}{N}\sum_{j=1}^N u_j.$$

Figure 4.1(b) shows the result of varying the regularisation parameter λ in the ROF model (4.2). The rate at which small-scale features are merged into larger-scale features is inversely proportional to λ.

The next example is a comparison of the ROF model with the classical Tikhonov regularisation.

Example 4.5 (TV vs. H^1-regularisation)**.** Consider the same noisy 1-D signal as in Example 4.4, shown in Figure 4.1(a).

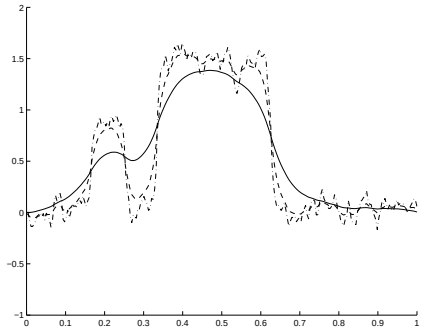

Figure 4.2: *Tikhonov regularisation. The figure shows the restoration of the noisy 1-D signal shown in Figure 4.1(a) using the H^1 semi-norm $\|\nabla u\|^2_{L^2(\Omega)}$ as the regularisation functional. The restoration is shown for $\lambda = 1$ (dash-dotted line), $\lambda = 0.1$ (dashed line) and $\lambda = 0.01$ (solid line).*

First, we replace $TV(\cdot)$ in (4.2) with a different choice of regularisation functional. A natural choice in the context of Tikhonov regularisation is the H^1 semi-norm $\|\nabla u\|^2_{L^2(\Omega)}$, in which case the minimisation is performed over the Sobolev space $H^1(\Omega)$. The restoration of the noisy signal using the H^1 semi-norm is shown in Figure 4.2.

There is a striking difference between the performance of the H^1 semi-norm and the TV functional, at least for the above piecewise constant signal. Generally, the H^1 semi-norm cannot restore edges as well as TV can, because it simply penalises too much the gradient associated with an edge. Another explanation is the following. For the H^1 semi-norm, the solution is characterised by the Euler–Lagrange equation

$$(\lambda I - 2\Delta)u = \lambda f.$$

Notice that the solution depends on the Laplacian Δu which is known to have strong isotropic smoothing properties. For more details, see [5].

See [76] for further comparisons of the ROF model and other regularisation functionals.

Numerous examples of the ROF model applied to 2-D images will be given in Chapter 6. We conclude this section with a final example, illustrating a limitation of the ROF model whenever it is used to recover smooth functions.

Example 4.6 (Staircase effect)**.** Consider the signal shown in Figure 4.3(a). The original signal contains some smooth regions that are not particularly well restored by the ROF model, as can be seen in Figure 4.3(b). The restoration here was computed with $\lambda = 5$ after adding Gaussian white noise with a standard deviation of $\sigma = 0.05$. The result shows the formation of a number of plateaus rather than the original smooth regions. This is known as the "staircase" effect, and it shows that TV regularisation favours piecewise constant solutions.

A number of improvements to the ROF model have been suggested to avoid the staircase effect, including minimisation of the second-order total variation, i.e., the total variation of the gradient. See [16, 18] for more.

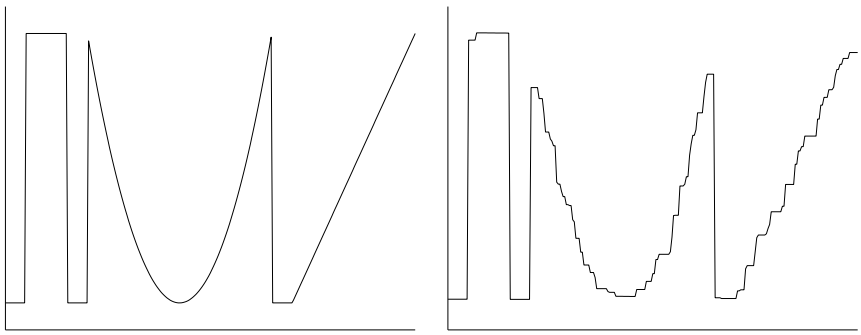

(a) Original (b) TV restoration

Figure 4.3: *The staircase effect. A small amount of noise was added to the piecewise smooth signal on the left before attempting to recover it with the ROF model. The result on the right shows the formation of many small piecewise constant steps rather than the smooth or affine transitions of the original signal. This illustrates a shortcoming of the ROF model which is known as the "staircase effect".*

4.2 The vectorial ROF model

In this section, we present the generalisation of the ROF model to vector-valued functions. This is necessary in order to consider restoration problems for colour and other multi-channel images. Note that the total variation does not extend uniquely to vector-valued functions. In fact, several alternatives have been suggested in the literature, for example, by Blomgren and Chan [9] and Goldluecke *et al.* [33]. However, we will restrict ourselves to the definition which is given by Bresson and Chan in [11] which is, perhaps, the most commonly used.

First, let the image domain $\Omega \subset \mathbb{R}^m$ be an open and bounded set with a C^1 boundary $\partial\Omega$. We now consider an image to have $n \geq 2$ colour channels. That is, an image $\boldsymbol{f} = (f^1, \ldots, f^n)$ is a function defined on Ω taking values in $\mathbb{R}^n$. Once more, let $1 \leq p \leq \infty$, and consider the forward image model for the observed image $\boldsymbol{f} \in L^2(\Omega; \mathbb{R}^n)$ given by

$$\boldsymbol{f} = A\boldsymbol{u} + \boldsymbol{\eta}. \tag{4.3}$$

The original image is $\boldsymbol{u} \in L^p(\Omega; \mathbb{R}^n)$, and $A : L^p(\Omega; \mathbb{R}^n) \to L^2(\Omega; \mathbb{R}^n)$ is a linear operator. The noise $\boldsymbol{\eta} \in L^2(\Omega; \mathbb{R}^n)$ is assumed to be additive white Gaussian-distributed with zero mean and known variance.

Suppose that the operator A does not introduce any dependencies between the channels. That is, if

$$A\boldsymbol{u} = \left(A_1 u^1, \ldots, A_n u^n\right),$$

where $A_1, \ldots, A_n$ are linear operators from $L^p(\Omega)$ to $L^2(\Omega)$. Then it is possible to solve the above ill-posed inverse problem by computing the solution to the ROF model (4.2) separately for each colour channel. Unfortunately, this approach does not take into account any potential correlations between the colour channels. This provides poor results for many colour image restoration problems, because images very often exhibit rather strong cross-channel correlations.

An alternative is to apply some sort of regularisation in which the coupling between the colour channels is taken into account. First, we present some more notation. Recall that the gradient matrix of $\boldsymbol{f}$ is given by

$$\nabla \boldsymbol{f} := \begin{pmatrix} (\nabla f^1)^T \\ (\nabla f^2)^T \\ \vdots \\ (\nabla f^n)^T \end{pmatrix} = \begin{pmatrix} f^1_{x_1} & f^1_{x_2} & \cdots & f^1_{x_m} \\ f^2_{x_1} & f^2_{x_2} & \cdots & f^2_{x_m} \\ \vdots & \vdots & \ddots & \vdots \\ f^n_{x_1} & f^n_{x_2} & \cdots & f^n_{x_m} \end{pmatrix}.$$

Similarly, for $\boldsymbol{\varphi} = (\boldsymbol{\varphi}^1, \dots, \boldsymbol{\varphi}^n) : \Omega \to \mathbb{R}^{n\times m}$, where each component is given by $\boldsymbol{\varphi}^i = (\varphi^{i,1}, \dots, \varphi^{i,m}) : \Omega \to \mathbb{R}^m$ for $i = 1, \dots, n$, the divergence of $\boldsymbol{\varphi}$ is defined by

$$\operatorname{div} \boldsymbol{\varphi} := \left(\operatorname{div} \boldsymbol{\varphi}^1, \operatorname{div} \boldsymbol{\varphi}^2, \dots, \operatorname{div} \boldsymbol{\varphi}^n\right).$$

For a matrix $A \in \mathbb{R}^{M,N}$ the Frobenius norm $\|\cdot\|_F$ is given by

$$\|A\|_F := \sqrt{\left(\sum_{i,j=1}^{M,N} a_{i,j}^2\right)}.$$

The following definition extends TV to vector-valued functions.

Definition 4.7 (Vectorial TV)**.** *For $\boldsymbol{f} \in L^1(\Omega; \mathbb{R}^n)$ the* vectorial TV (VTV) *of f is defined by*

$$TV(\boldsymbol{f}) := \sup_{\boldsymbol{\varphi} \in \mathcal{V}_n} -\int_\Omega \boldsymbol{f} \cdot \operatorname{div} \boldsymbol{\varphi}, \tag{4.4}$$

where $\mathcal{V}_n$ is the set

$$\mathcal{V}_n := \left\{ \boldsymbol{\varphi} \in C_c^1(\Omega; \mathbb{R}^{n\times m}) \;\middle|\; \|\boldsymbol{\varphi}(x)\|_F \le 1 \text{ for every } x \in \Omega \right\}.$$

Notice that for $n = 1$ the vectorial TV reduces to the scalar definition (3.11). Also, for differentiable functions $\boldsymbol{f} \in C^1(\Omega; \mathbb{R}^n)$, the vectorial TV reduces to

$$TV(\boldsymbol{f}) = \int_\Omega \|\nabla \boldsymbol{f}\|_F.$$

Next, we summarise some properties of the vectorial TV. The proof is exactly the same as in Theorem 3.19.

Theorem 4.8. *Let $\Omega \subset \mathbb{R}^m$ be open and bounded. Then $TV(\cdot)$ is a proper, closed convex function on $L^p(\Omega; \mathbb{R}^n)$ for $1 \le p \le \infty$.*

The effective domain of the vectorial TV defines a space $BV(\Omega; \mathbb{R}^n)$ of vector-valued functions of bounded variation.

Definition 4.9 (Bounded variation for vector-valued functions)**.** *We say that $\boldsymbol{f} \in L^1(\Omega; \mathbb{R}^n)$ has* bounded variation *provided that its vectorial total variation is finite. The space of vector-valued functions of bounded variation is*

$$BV(\Omega; \mathbb{R}^n) := \left\{ \boldsymbol{f} \in L^1(\Omega; \mathbb{R}^n) \mid TV(\boldsymbol{f}) < +\infty \right\}.$$

This is a Banach space equipped with the norm

$$\|\boldsymbol{f}\|_{BV(\Omega;\mathbb{R}^n)} := \|\boldsymbol{f}\|_{L^1(\Omega;\mathbb{R}^n)} + TV(\boldsymbol{f}). \tag{4.5}$$

Since there is usually no distinction between scalar and vector-valued functions in Sobolev space theory (see [30]), the results for $BV(\Omega)$ in Section 3.2 extend effortlessly to the vector-valued case. For further details, see [11].

Now, the *vectorial Rudin–Osher–Fatemi (VROF) model* is to solve the following minimisation problem

$$\min_{\boldsymbol{u}\in L^p(\Omega;\mathbb{R}^n)} \left\{ F(\boldsymbol{u}) := TV(\boldsymbol{u}) + \frac{\lambda}{2} \left\| A\boldsymbol{u} - \boldsymbol{f} \right\|^2_{L^2(\Omega;\mathbb{R}^n)} \right\}, \tag{4.6}$$

where the functional $TV : L^p(\Omega;\mathbb{R}^n) \to \bar{\mathbb{R}}$ is the vectorial TV.

The existence and uniqueness of the solution can be shown under the same assumptions as for the ROF model. A proof in the case that $A = I$ and $f \in L^\infty(\Omega;\mathbb{R}^n)$ was given by Bresson and Chan in [11].

Theorem 4.10 (Existence-uniqueness)**.** *Let $\Omega \subset \mathbb{R}^m$ be an open and bounded domain with a C^1 boundary $\partial\Omega$. Suppose that $1 \leq p \leq m/(m-1)$. Then there exists a solution to the VROF model* (4.6)*, provided that the operator A does not annihilate constant functions. Moreover, the solution is unique if A is injective.*

Proof. The coercivity in Lemma 4.2 is extends trivially to the VROF functional and the space $BV(\Omega;\mathbb{R}^n)$. Then the proof is the same as for Theorem 4.3. □

4.3 TV inpainting

In this section, we present a TV-based model for image inpainting that was proposed by Shen and Chan in [72].

One might as well work with colour images, so we assume the same notation as in the previous section. In particular, the image domain $\Omega \subset \mathbb{R}^m$ is an open and bounded set with a C^1 boundary $\partial\Omega$, and the number of colour channels is $n \geq 2$. In addition, let $\Gamma \subset \Omega$ be the inpainting domain where the image data is missing. In this case, fix $1 \leq p \leq \infty$, and the forward image model that relates the degraded image $\boldsymbol{f} : L^2(\Omega;\mathbb{R}^n)$ to the original image $\boldsymbol{u} : L^p(\Omega;\mathbb{R}^n)$ is given by

$$\boldsymbol{f}(x) = \boldsymbol{\eta}(x) + \begin{cases} \boldsymbol{u}(x) & \text{if } x \in \Omega \setminus \Gamma, \\ 0 & \text{otherwise.} \end{cases} \tag{4.7}$$

Recall that $\boldsymbol{\eta} \in L^2(\Omega;\mathbb{R}^n)$ is an additive white Gaussian noise. The above can be seen as a particular case of the forward image model (4.3), where the operator A is defined by

$$(A\boldsymbol{u})(x) := \begin{cases} \boldsymbol{u}(x) & \text{if } x \in \Omega \setminus \Gamma, \\ 0 & \text{otherwise.} \end{cases}$$

The TV inpainting model for recovering the original image is then

$$\min_{\boldsymbol{u}\in L^p(\Omega;\mathbb{R}^n)} \left\{ F(\boldsymbol{u}) := TV(\boldsymbol{u}) + \frac{\lambda}{2} \left\| \boldsymbol{u} - \boldsymbol{f} \right\|^2_{L^2(\Omega\setminus\Gamma;\mathbb{R}^n)} \right\}. \tag{4.8}$$

Once more, the functional $TV : L^p(\Omega;\mathbb{R}^n) \to \bar{\mathbb{R}}$ is the vectorial TV. Notice that the L^2-norm of the data fidelity term is only evaluated outside the inpainting domain Γ.

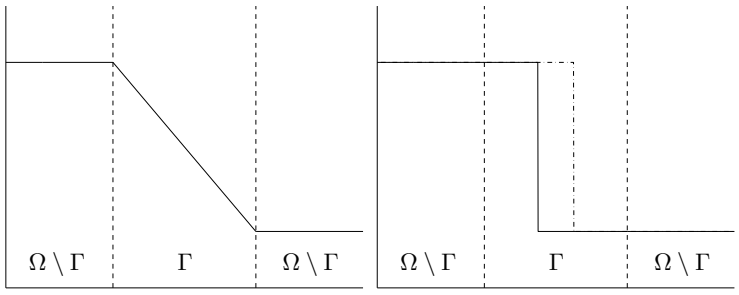

(a) Harmonic inpainting (b) TV inpainting

Figure 4.4: *The solution to the TV-based inpainting model is not necessarily unique. The figure illustrates a simple 1-D example of inpainting the missing region $\Gamma \subset \Omega$. On the left is the solution according to the harmonic inpainting model, i.e., a smooth interpolation between the boundary values. The figure on the right shows two possible solutions for the TV-based model which both have the same total variation.*

If we assume that $\boldsymbol{f}(x) = 0$ for $x \in \Gamma$, this merely becomes a special case of the VROF model (4.6). Hence, the existence of a solution to the TV inpainting model (4.8) follows from Theorem 4.10. On the other hand, it is quite easy to see that the operator A is no longer injective, and thus the data fidelity functional is not strictly convex. Therefore, unlike the VROF model, we can no longer guarantee that the solution is unique. This is most easily illustrated in the 1-D case, see Figure 4.4.

If there is no noise present in the image, we may wish to force the solution to agree exactly with the data outside the inpainting region. This can be seen as letting $\lambda \to \infty$ in (4.8). The inpainting problem can then be written

$$\min_{u \in L^p(\Omega;\mathbb{R}^n)} TV(\boldsymbol{u}) \quad \text{such that} \quad \boldsymbol{u}|_{\Omega\setminus\Gamma} = \boldsymbol{f}\big|_{\Omega\setminus\Gamma}. \tag{4.9}$$

For comparison, we will also use the harmonic inpainting model

$$\min_{\boldsymbol{u} \in L^2(\Omega;\mathbb{R}^n)} \|\nabla \boldsymbol{u}\|^2_{L^2(\Omega;\mathbb{R}^n)} + \frac{\lambda}{2} \|\boldsymbol{u} - \boldsymbol{f}\|^2_{L^2(\Omega\setminus\Gamma;\mathbb{R}^n)}. \tag{4.10}$$

This is based on a quadratic regularisation term such as that in Tikhonov regularisation. The harmonic inpainting model usually fills in missing pixels so as to give a smooth solution, see [72].

The following example demonstrates the difference between the harmonic and TV inpainting models.

Example 4.11 (Harmonic vs. TV inpainting)**.** Consider inpainting a step edge that has been partially occluded as shown in Figure 4.5(a). Now, compare the restoration of this edge with the harmonic and TV inpainting models as shown in Figure 4.5(b) and Figure 4.5(c). The harmonic inpainting produces fills the gap with a smooth transition. On the other hand, the TV inpainting is able to recover the edge sharply.

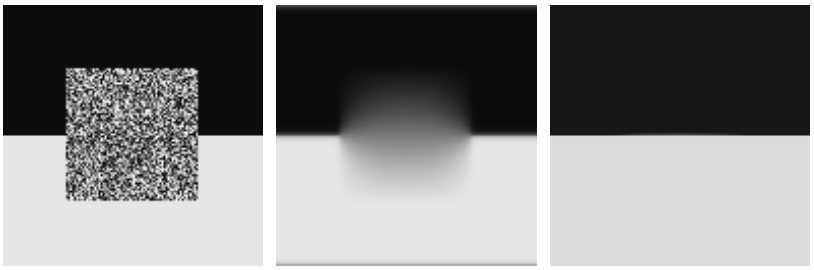

(a) An occluded step edge (b) Harmonic inpainting (c) TV inpainting

Figure 4.5: *Harmonic vs. TV inpainting. Figure (a) represents a horizontal edge between the black and white regions which has become occluded. The occlusion is represented by the noisy region in the centre of the image. Figures (b) and (c) show the result of inpainting the missing pixels with the harmonic inpainting model* (4.10) *and TV inpainting model* (4.8)*, respectively. Notice that the former produces a smooth fill-in of the occluded domain, whereas the latter recovers the horizontal edge.*

4.4 A TV-based model for demosaicking

Finally, we present a variational method for the demosaicking problem, i.e., the restoration of a full colour image that has been subsampled by a colour filter array (CFA). Recall that a CFA is a set of colour-selective filters that assigns each pixel in the image to a single colour channel. For a description of how this problem arises in a single-sensor digital colour camera, see Chapter 2.

Although this problem may be approached by inpainting each colour channel separately, or as a special case of the vectorial ROF model (4.6), we will see from the numerical experiments in Chapter 6 that the results then become quite disappointing. Instead, inspired by the work of Alleysson *et al.* [2], our method is based on decomposing the image into a luminance and a chrominance component and applying a TV regularisation to each of these components separately. From a regularisation point of view, this is a change of basis which better describes what we expect the restored image to look like.

The situation is the following. Once more, $\Omega \subset \mathbb{R}^m$ is an open and bounded set with a C^1 boundary $\partial\Omega$. The original RGB colour image $\boldsymbol{u} : \Omega \to \mathbb{R}^3$ consists of a red, green and blue colour channel $\boldsymbol{u} = (u^R, u^G, u^B)$. Each of these is subsampled according to some CFA. The mosaicked image $\boldsymbol{f}$ and the original image $\boldsymbol{u}$ are related through the forward image model (4.3), where the operator A is given by

$$(A\boldsymbol{u})(x) := C(x)\boldsymbol{u}(x). \tag{4.11}$$

Here $C(x) := \operatorname{diag}(C_R(x), C_G(x), C_B(x))$ is a 3-by-3 diagonal matrix. For each $x \in \Omega$, this describes the proportion of each colour that is transmitted by the CFA.

For the sake of simplicity, we will only consider the standard Bayer CFA (see Figure 2.2(a)). Then the matrix $C(x)$ was given by (2.6) in Section 2.2. Many demosaicking algorithms are tailored to this particular configuration. However, note that variational methods, such as the one presented here, work for arbitrary CFA patterns. In fact, it has been argued that this is the only way to compare the performance of various proposed CFA designs [24].

(a) Original (b) Luminance (c) Chrominance

Figure 4.6: *Luminance-chrominance decomposition. The image in (a) is decomposed into a luminance component carrying the spatial information in the image, shown in (b), and a chromatic component representing the colour information, shown in (c). In this example, the luminance is the mean of the three colour channels, i.e.,* $p_R = p_G = p_B = 1/3$.

Perhaps the simplest approach to demosaicking is to treat each colour channel separately and perform a bilinear interpolation to recover the missing pixel values. Unfortunately, this method has a tendency to blur edges and introduce false colours due to aliasing. This might be improved by applying the TV inpainting model (4.8) to each colour channel instead of a bilinear interpolation. However, any correlation between the different colour channels is still not taken into account. Yet another alternative is to use the vectorial ROF model (4.6) which introduces at least some dependency between the colour channels.

Another important approach, proposed by Alleysson *et al.* [2], comes from regarding the subsampled CFA image as the sum of an achromatic luminance component and a chromatic component. The former contains the spatial information of the image, whereas the latter represents the chromatic information. A key principle is that the human visual system is more sensitive to variations in luminance compared to chrominance. Moreover, for demosaicking with the Bayer CFA, it is shown in [2] that the luminance can be estimated with full spatial resolution, and only the chromatic component is subject to subsampling by the CFA. In other words, every pixel of the degraded image $\boldsymbol{f}$ carries some luminance information, and only the chrominance requires that we perform some sort of interpolation.

The decomposition of an image $\boldsymbol{u}$ into a luminance component $\Phi : \Omega \to \mathbb{R}$ and three opponent chrominatic components $\boldsymbol{\Psi} = (\Psi^1, \Psi^2, \Psi^3) : \Omega \to \mathbb{R}^3$ may be written

$$\boldsymbol{u} = \Phi + \boldsymbol{\Psi}. \tag{4.12}$$

The luminance component is defined by taking a convex combination of the colour channels

$$\Phi := \sum_{i \in \{R,G,B\}} p_i u^i, \tag{4.13}$$

where $p_i > 0$ is the weight associated with the i-th colour channel such that $\sum_i p_i = 1$. The chrominance is obtained by subtracting the luminance from the original image

$$\Psi^i := (1 - p_i)u^i - \sum_{j \neq i} p_j u^j. \tag{4.14}$$

An example of such a decomposition is shown in Figure 4.6.

It was observed by Alleysson *et al.* [2], that a decomposition similar to (4.12) holds for the mosaicked image $\boldsymbol{f}$,

$$\boldsymbol{f} = \Phi + A\boldsymbol{\Psi}, \tag{4.15}$$

where A is the CFA subsampling operator (4.11). The method in [2] estimates the luminance and subsampled chrominance components from (4.15) by linear filtering. Next, the missing chrominance information is interpolated, and the restored image is obtained as a sum of the luminance and interpolated chrominance components. This algorithm greatly outperforms most demosaicking methods that work in the usual RGB basis.

We now consider this idea in the setting of TV-based regularisation. This is achieved by decomposing the image into a luminance and a chrominance component, as in (4.12), and then applying a TV regularisation to each of them. Hence, we propose the following *TV demosaicking model*, which is to solve the minimisation problem

$$\min_{\boldsymbol{u}\in L^p(\Omega;\mathbb{R}^n)} \left\{ F(\boldsymbol{u}) := TV(\Phi) + \mu\, VTV(\boldsymbol{\Psi}) + \frac{\lambda}{2}\, \|A\boldsymbol{u} - \boldsymbol{f}\|^2_{L^2(\Omega;\mathbb{R}^n)} \right\}. \tag{4.16}$$

Here the functional $TV : L^p(\Omega) \to \bar{\mathbb{R}}$ is the usual scalar total variation, and $VTV : L^p(\Omega;\mathbb{R}^n) \to \bar{\mathbb{R}}$ denotes the vectorial TV, see Definition 4.7.

There are two parameters $\lambda > 0$ and $\mu > 0$ in the TV demosaicking model (4.16). As before, λ is a weight associated with the data fidelity term, and thus determines how close the solution should remain to the observed data. There are, however, two regularisation terms, and the weighting between them is determined by the parameter μ. Intuitively, larger values of μ ought to result in a more regular chrominance component, and, conversely, smaller values of μ should produce a more regular luminance component. The relationship between these parameters is explored further through numerical experiments in Chapter 6.

We end with a result that states the existence of a solution for the above TV demosaicking model.

Theorem 4.12 (Existence of solution)**.** *Let $\Omega \subset \mathbb{R}^m$ be an open and bounded domain with a C^1 boundary $\partial\Omega$. Suppose that $1 \le p \le m/(m-1)$. Then there exists a solution to the TV demosaicking model* (4.16)*, where A is an operator on the form given by* (4.11)*.*

Proof. The result follows from the exact same argument as for the ROF model in Theorem 4.3 and the VROF model in Theorem 4.10, provided that we can show that the functional F in (4.16) is coercive on $BV(\Omega;\mathbb{R}^n)$.

First, observe that for $\boldsymbol{u} = (u^1, \dots, u^n)$, we have

$$VTV(u) \le \sum_{i=1}^{n} TV(u^i).$$

This follows from the definition of the vectorial TV (4.4), that is,

$$\begin{aligned}\sup_{\boldsymbol{\varphi}\in\mathcal{V}_n} -\int_\Omega \boldsymbol{u}\cdot\operatorname{div}\boldsymbol{\varphi} &= \sup_{(\boldsymbol{\varphi}^1,\dots,\boldsymbol{\varphi}^n)\in\mathcal{V}_n} -\int_\Omega \sum_{i=1}^n u^i \operatorname{div}\boldsymbol{\varphi}^i \\ &\le \sum_{i=1}^n \sup_{\boldsymbol{\varphi}^i\in\mathcal{V}_1} -\int_\Omega u^i \operatorname{div}\boldsymbol{\varphi}^i \\ &= \sum_{i=1}^n TV(u^i).\end{aligned}$$

Now, since $\boldsymbol{u}$ can be decomposed as $\boldsymbol{u} = \Phi + \boldsymbol{\Psi}$, we have

$$\begin{aligned}VTV(\boldsymbol{u}) &= VTV(\Phi+\boldsymbol{\Psi}) \\ &\le n\,TV(\Phi) + VTV(\boldsymbol{\Psi}) \\ &\le C_1\left(TV(\Phi) + \mu\,VTV(\boldsymbol{\Psi})\right),\end{aligned}$$

where $C_1 := n\max\{1, 1/\mu\}$.

Finally, recall that the functional in the VROF model (4.6) is coercive on $BV(\Omega;\mathbb{R}^n)$, i.e., there exists $C_2 > 0$ such that

$$\begin{aligned}\|\boldsymbol{u}\|_{BV(\Omega;\mathbb{R}^n)} &\le C_2\left(VTV(\boldsymbol{u}) + \frac{\lambda}{2}\|A\boldsymbol{u}-\boldsymbol{f}\|^2_{L^2(\Omega;\mathbb{R}^n)}\right) \\ &\le C_1C_2\left(TV(\Psi) + \mu\,VTV(\boldsymbol{\Phi}) + \frac{\lambda}{2}\|A\boldsymbol{u}-\boldsymbol{f}\|^2_{L^2(\Omega;\mathbb{R}^n)}\right) \\ &= C\,F(\boldsymbol{u}),\end{aligned}$$

where $C := C_1C_2$. This shows that F is coercive on $BV(\Omega;\mathbb{R}^n)$. □

4.5 Summary

In this chapter we first introduced the ROF model for the denoising and deblurring of grayscale images. This problem possesses a unique solution under some assumptions on the operator A that models the blur in the forward image model (4.1). The proof of this relies on a standard method from the calculus of variations and much of the material from Chapter 3. Next, we have illustrated how the ROF model favours piecewise constant solutions, in contrast to the usual Tikhonov regularisation where the solution is much smoother and cannot represent edges sharply.

The vectorial ROF model is the generalisation of this to colour images. Rather than performing an image restoration by applying the ROF model separately to each colour channel, the VROF model is able to take into account some coupling between the colour channels. The TV inpainting model in Section 4.3, may be derived as a particular case of the VROF model. This provides an edge-preserving method for inpainting missing pixels in an image.

Finally, based on the approach of Alleysson *et al.* [2], we considered the decomposition of a colour image into a luminance and a chrominance component. Then, we proposed a TV-based model for demosaicking by applying a TV regularisation to the luminance and chrominance components separately.

CHAPTER 5

Numerical methods for TV-based regularisation

In this chapter, we consider numerical methods for computing the solution to the variational problems that were presented in Chapter 4. Some of these methods rely on certain tools from convex analysis, whereby a convex optimisation problem can be associated with another convex optimisation problem, known as the dual problem. From this duality follows the equivalence of mimising the primal objective function and maximising a certain dual objective function. The numerical algorithms we consider are divided into primal methods, dual methods and primal-dual methods, depending on whether they seek the optimal solution via the primal or dual objective function, or both. There are also many other numerical methods for the ROF model which we do not consider here, for example, those based on split Bregman iteration [34], graph cuts [25], or the augmented Lagrangian method [77].

For the sake of clarity in the following presentation, we will focus first and foremost on the ROF model in Section 5.1 and Section 5.2. However, the primal-dual methods in Section 5.3 are naturally considered in a general context that accomodates all of our TV-based regularisation problems.

Implementing these methods requires that we work in a discrete setting. In this, we consider only 2-D images defined on a rectangular grid. Thus, an image of M-by-N pixels becomes a vector $u \in \mathbb{R}^{MN}$, and $u_{i,j}$ denotes the value of the (i,j)-th pixel. Let X denote $\mathbb{R}^{MN}$ and also let $Y = X \times X$. Then the discretisation of the gradient operator $\nabla : X \to Y$ using forward differences and Neumann boundary conditions is

$$(\nabla u)_{i,j} = \big((\nabla u)^1_{i,j}, (\nabla u)^2_{i,j}\big)$$

where

$$(\nabla u)^1_{i,j} = \begin{cases} u_{i+1,j} - u_{i,j} & \text{if } i < M, \\ 0 & \text{if } i = M, \end{cases}$$
$$(\nabla u)^2_{i,j} = \begin{cases} u_{i,j+1} - u_{i,j} & \text{if } j < N, \\ 0 & \text{if } j = N. \end{cases}$$

The divergence operator is defined as $\operatorname{div} = -\nabla^*$, where $\nabla^* : Y \to X$ is the adjoint of the gradient, i.e., $(\nabla u, v) = (u, \nabla^* v)$ for every $u \in X$, $v \in Y$. It can easily be verified that for $\boldsymbol{p} = (p^1, p^2) \in Y$, we have

$$(\operatorname{div} \boldsymbol{p})_{i,j} = \begin{cases} p^1_{i,j} - p^1_{i-1,j} & \text{if } 1 < i < M, \\ p^1_{i,j} & \text{if } i = 1, \\ -p^1_{i-1,j} & \text{if } i = M, \end{cases} + \begin{cases} p^2_{i,j} - p^2_{i,j-1} & \text{if } 1 < j < N, \\ p^2_{i,j} & \text{if } j = 1, \\ -p^2_{i,j-1} & \text{if } j = N. \end{cases}$$

Using the above notation, the discretisation of the ROF model is then

$$\min_{u \in X} \|\nabla u\|_1 + \frac{\lambda}{2} \|Au - f\|_2^2 . \tag{5.1}$$

Here $\|u\|_2^2 = \langle u, u \rangle_X$ is the squared ℓ^2-norm, and $\|u\|_1 := \sum_{i,j} |u_{i,j}|$ is the ℓ^1-norm.

It is possible to show, by introducing a step size h into the finite difference operators and using the notion of Γ-convergence, that the solution to the discretisation (5.1) converges to the true solution of the ROF model (4.2) as the step size h becomes small, see [14] and the references therein.

For colour images, we take $X = \mathbb{R}^{MNP}$ where P is the number of colour channels. The gradient of a colour image $u = (u^1, \dots, u^P) \in X$ is defined componentwise

$$\nabla u := \begin{pmatrix} \nabla u^1 \\ \vdots \\ \nabla u^P \end{pmatrix},$$

and a similar definition holds for the divergence of $\boldsymbol{p} \in Y = X \times X$.

The discretisation of the VROF model (4.6) is

$$\min_{u \in X} \sum_{i,j} \sqrt{\sum_k \left|\nabla u^k_{i,j}\right|^2} + \frac{\lambda}{2} \|Au - f\|_2^2 . \tag{5.2}$$

For the TV inpainting (4.8), the data fidelity term is altered as follows

$$\min_{u \in X} \sum_{i,j} \sqrt{\sum_k \left|\nabla u^k_{i,j}\right|^2} + \frac{\lambda}{2} \sum_{(i,j,k) \in \Omega \setminus \Gamma} (u^k_{i,j} - f^k_{i,j})^2 . \tag{5.3}$$

Here we define $\Omega := \{(i,j,k) \mid 1 \le i \le M, 1 \le j \le N, 1 \le k \le P\}$, and the inpainting domain $\Gamma \subset \Omega$ is the set of indices for which the image data is missing.

Finally, consider the TV demosaicking model from Section 4.4. In the discrete case, the operator A in (4.11) is a diagonal matrix. The luminance and chrominance components of the image $u \in X$ are denoted by $\Phi \in \mathbb{R}^{MN}$ and $\Psi \in X$, respectively. Then u may be decomposed into the sum

$$u = \Phi + \Psi, \qquad \Phi := \sum_{k=1}^{P} \alpha_k u^k,$$

where $\alpha_k > 0$ and $\sum_k \alpha_k = 1$.

As with the TV inpainting model above, one can define the set $\Gamma \subset \Omega$ of indices where the (i,j)-th pixel of the k-th colour channel vanishes due to the colour filter array. Then the discrete counterpart of the TV demosaicking model becomes

$$\min_{u \in X} \|\nabla \Phi\|_1 + \sum_{i,j} \sqrt{\sum_k \left|\nabla \Psi_{i,j}^k\right|^2} + \frac{\lambda}{2} \sum_{(i,j,k) \in \Omega \setminus \Gamma} (u_{i,j}^k - f_{i,j}^k)^2. \tag{5.4}$$

5.1 Primal methods

Numerical methods that are based on directly minimising the ROF model (5.1) are referred to as *primal methods*. However, this approach is problematic because the total variation $TV(f)$ is not differentiable whenever $|\nabla f| = 0$. It was suggested by Acar and Vogel [1] and Vogel and Oman [76] to consider instead a slightly perturbed functional

$$TV_\beta(f) = \int_\Omega \sqrt{|\nabla f|^2 + \beta^2}\, dx \tag{5.5}$$

for $\beta \geq 0$. Note that this functional is differentiable for positive values of β, and for $\beta = 0$ it reduces to the usual TV for differentiable functions (a generalisation to integrable functions can also be made, see [1]). Furthermore, the functional TV_β is also convex and lower semi-continuous. In [1] it is shown that using this as a regularisation functional yields a well-posed problem under the same assumptions as for the ROF model. In addition, the solution converges to the solution of the ROF model as $\beta \to 0$. This is illustrated in Figure 5.1. Although it results in a worse approximation, the benefit of choosing β large in (5.5) is that numerical methods tend to become more robust. See, e.g., [76] and [20]. Each of the primal methods discussed here will depend on this smoothing to give stable results.

Euler–Lagrange equations

The original denoising algorithm presented by Rudin, Osher and Fatemi in [65] is based on solving the Euler–Lagrange equation for the ROF model. This can be derived by computing the gradient of the ROF functional. The gradient is given by

$$F'(u) = -\operatorname{div}\left(\frac{\nabla u}{\sqrt{|\nabla u|^2 + \beta^2}}\right) + \lambda A^*(Au - f). \tag{5.6}$$

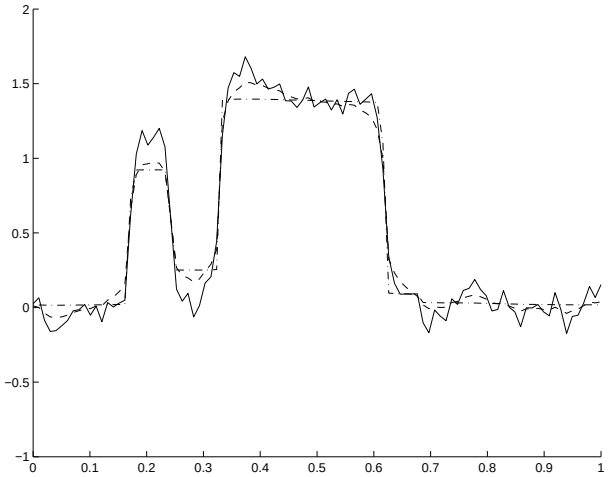

Figure 5.1: *Smooth approximation to the TV functional. The figure illustrates the effect of various choices of the parameter β for the functional TV_β in* (5.5)*. The solid line corresponds to $\beta = 1$, the dashed line corresponds to $\beta = 0.1$ and the dash-dotted line corresponds to $\beta = 0.01$. As β decreases, the solution approaches the true minimum of the ROF model.*

A minimiser of the ROF model is then a weak solution to the Euler–Lagrange equation

$$\begin{cases} F'(u) = 0 & \text{in } \Omega, \\ \dfrac{\partial u}{\partial \nu} = 0 & \text{on } \partial\Omega. \end{cases} \tag{5.7}$$

This is a non-linear elliptic PDE, and it can be difficult to solve directly. Rudin, Osher and Fatemi [65] proposed an artificial time marching scheme by letting u be a function also of a time variable $t > 0$ and driving the parabolic PDE

$$\begin{cases} \dfrac{\partial u}{\partial t} = F'_x(u) & \text{in } \Omega, \\ \dfrac{\partial u}{\partial \nu} = 0 & \text{on } \partial\Omega. \end{cases}$$

to a steady state as $t \to \infty$. The point is that the parabolic problem can be solved iteratively with an explicit finite difference scheme. It was also noted in [65], that upon choosing λ appropriately, this becomes the gradient projection method for non-linear programming.

In [76], Vogel and Oman discuss a gradient descent algorithm which is exactly the above method whenever λ is held fixed. Given some initial guess $u^{(0)}$, a descent direction $d^{(k)}$ is computed, and the following iteration is performed

$$u^{(k+1)} = u^{(k)} + \alpha^* d^{(k)}. \tag{5.8}$$

The step size α^* can either be fixed, or an optimal value

$$\alpha^* = \arg\min_{\alpha > 0} F(u^{(k)} + \alpha d^{(k)}), \tag{5.9}$$

may be computed with a line search or trust region algorithm.

For the method of steepest descent, the descent direction is the negative gradient $d^{(k)} = -F'(u^{(k)})$. Due to restrictions on the step size at each iteration, this method shows very poor convergence. This can be somewhat improved through a preconditioning such as that proposed by Marquina and Osher in [53].

Higher-order methods

Higher-order methods, such as Newton's method, can be formulated by choosing alternative descent directions. However, these methods involve higher-order derivatives of the minimising functional. The Hessian of F is the mapping $u \mapsto H(u)$, where $H(u)$ is the functional given by

$$(H(u))(v) = -\operatorname{div}\left(\frac{1}{\sqrt{|\nabla u|^2 + \beta^2}}\left(1 - \frac{\nabla u(\nabla u)^T}{|\nabla u|^2 + \beta^2}\right)\nabla v\right) + \lambda A^* A v. \quad (5.10)$$

Now, we obtain Newton's method by choosing the descent direction

$$d^{(k)} = -\left[H(u^{(k)})\right]^{-1} F'(u^{(k)}),$$

assuming that the Hessian is invertible at $u^{(k)} \in X$.

Vogel and Oman [76] also propose a quasi-Newton method which they call "lagged diffusivity fixed-point iteration." Although this method avoids explicitly evaluating the Hessian, it still relies on solving a linear system for each iteration.

Another higher-order method with improved convergence rate and domain of convergence was proposed by Chan, Golub and Mulet [20]. Since most of the problems with convergence are mainly due to the singularity of the gradient (5.6), their idea was to introduce the dual variable

$$\boldsymbol{p} = \frac{\nabla u}{\sqrt{|\nabla u|^2 + \beta^2}}.$$

Despite being non-linear and non-differentiable for $\beta = 0$, this expression is usually smooth. This is because $\boldsymbol{p}$ can in fact be interpreted as the normal to the level lines of the image.

The introduction of the dual variable $\boldsymbol{p}$ lets us rewrite the Euler–Lagrange equations (5.7) as a system of non-linear equations

$$\begin{cases} -\operatorname{div}\boldsymbol{p} + \lambda A^*(Au - f) = 0, \\ \boldsymbol{p}\sqrt{|\nabla u|^2 + \beta^2} - \nabla u = 0. \end{cases} \quad (5.11)$$

As before, the second derivatives can be computed and a Newton's method can be derived for this system, see [20] for the details.

Generally, the convergence of these higher-order methods requires very few iterations, especially compared to the gradient descent. However, each iteration requires solving a very large system of linear equations which consumes both time and memory. In particular, for $u \in \mathbb{R}^N$ the expression $\nabla u(\nabla u)^T$ in the Hessian (5.10) gives rise to a dense matrix with N^2 elements. Even for

moderately sized images, say 512×512 pixels, the storage requirements become prohibitively large.

Moreover, the smoothness of the functional TV_β heavily impacts the convergence of these iterative methods. Whereas we usually expect quadratic convergence from Newton's method, this only occurs for relatively large values of β and after a large number of iterations. When β is small, the step size α^* determined in the line search (5.9) also becomes small, restricting the size of each iteration. The lagged diffusivity method appears to be less severely affected by the parameter β and also appears to converge even without a globalisation strategy such as the line search step, see [76]. The Chan–Golub–Mulet primal-dual method also behaves somewhat better for small β, see [20].

In [21] Chan, Chan and Zhou consider a continuation procedure for Newton's method. The parameter β is initially chosen to be large and the corresponding solution is used as an initial guess for the next iteration where β is gradually reduced. The resulting algorithm converges for arbitrary β and results on the domain of convergence were proved using the Newton–Kantorovich theorem.

5.2 Dual methods

So far, the numerical methods that we have considered all rely on solving some smooth approximation to the ROF model. The dual methods that we now present are based on a quite different approach which does not require any such smoothing. This method was introduced for the ROF model by Chambolle in [14], and extended to the VROF model by Bresson and Chan in [11]. Unfortunately, it depends on a certain simplification that occurs only for the pure denoising problem, i.e., whenever $A = I$ in the ROF model (5.1). However, this obstacle is dealt with by the methods considered in the next section.

It is well known in convex analysis that minimising a given convex function is, under certain conditions, equivalent to maximising another convex function. These are referred to as the primal and dual objective functions, respectively. We will not attempt to cover the general theory of convex duality, instead we refer to [64] for more. It will only be necessary to present the dual problem associated with the ROF model (5.1) in the finite-dimensional case.

Define the *Lagrangian function* $K : X \times Y \to \bar{\mathbb{R}}$ by

$$K(u, \boldsymbol{p}) := -(u, \operatorname{div} \boldsymbol{p}) + \frac{\lambda}{2} \|Au - f\|_2^2 + \delta_{\mathcal{V}}(\boldsymbol{p}). \tag{5.12}$$

Here $\delta_{\mathcal{V}}(\cdot)$ is the indicator function of the closed convex set

$$\mathcal{V} := \{\boldsymbol{p} \in Y \mid |\boldsymbol{p}_{i,j}| \leq 1 \text{ for } 1 \leq i \leq M, 1 \leq j \leq N\}.$$

Notice that the ROF model (5.1) can now be written as the saddle-point problem

$$\inf_{u \in X} \sup_{\boldsymbol{p} \in Y} K(u, \boldsymbol{p}). \tag{5.13}$$

From convex duality, the sup and inf may be exchanged if there exists a saddle-point of K, i.e., a point $(\bar{u}, \bar{\boldsymbol{p}}) \in X \times Y$ such that

$$K(\bar{u}, \boldsymbol{p}) \leq K(\bar{u}, \bar{\boldsymbol{p}}) \leq K(u, \bar{\boldsymbol{p}}) \qquad \forall\, x \in X, y \in Y.$$

Moreover, a sufficient condition for the existence of such a saddle-point is that the mapping $u \mapsto K(u, \cdot)$ is convex and lower semi-continuous and $\boldsymbol{p} \mapsto K(\cdot, \boldsymbol{p})$ is concave and upper semi-continuous, see Example 13 in [64].

The dual problem is obtained by exchanging the sup and inf as follows

$$\sup_{\boldsymbol{p} \in Y} \inf_{u \in X} K(u, \boldsymbol{p}). \tag{5.14}$$

Notice that for a fixed $\boldsymbol{p}$, the Lagrangian $K(u, \boldsymbol{p})$ is quadratic and differentiable in u. Thus, it can easily be verified that the infimum with respect to u is attained whenever

$$-\operatorname{div} \boldsymbol{p} + \lambda A^*(Au - f) = 0, \tag{5.15}$$

where A^* is the adjoint of A, i.e., $(Au, v) = (u, A^*v)$ for every $u, v \in X$.

For the pure denoising problem, i.e., whenever $A = I$, we can explicitly solve for u in (5.15) and get

$$u = f + \frac{1}{\lambda} \operatorname{div} \boldsymbol{p}. \tag{5.16}$$

Inserting this into (5.14), the dual problem becomes the constrained maximisation problem

$$\sup_{\boldsymbol{p} \in \mathcal{V}} \left\{ -(f, \operatorname{div} \boldsymbol{p}) - \frac{1}{2\lambda} \|\operatorname{div} \boldsymbol{p}\|^2_{L^2(\Omega)} \right\}. \tag{5.17}$$

The corresponding Euler–Lagrange equations give the optimality condition

$$-\nabla \left(\frac{1}{\lambda} \operatorname{div} \boldsymbol{p} + f \right) + \alpha \boldsymbol{p} = 0, \tag{5.18}$$

where $\alpha(x) \geq 0$ is the Lagrange multiplier for the constraint $|\boldsymbol{p}| \leq 1$. The key observation by Chambolle [14] is the following: If $|\boldsymbol{p}| < 1$, then the Lagrange multiplier is $\alpha = 0$, and, otherwise, it is $\alpha = \left| \nabla(\frac{1}{\lambda} \operatorname{div} \boldsymbol{p} + f) \right|$. In either case, the second of the two expressions for α must hold.

Chambolle then proposes to solve (5.18) with the following semi-implicit gradient descent algorithm. Choose $\boldsymbol{p}^0$ and $\tau > 0$ and compute the following iteration for $n = 1, 2, \ldots$

$$\boldsymbol{p}^{n+1} = \frac{\boldsymbol{p}^n + \tau \nabla \left(\frac{1}{\lambda} \operatorname{div} \boldsymbol{p}^n + f \right)}{1 + \tau \left| \nabla \left(\frac{1}{\lambda} \operatorname{div} \boldsymbol{p}^n + f \right) \right|}. \tag{5.19}$$

In [14], the above iteration was shown to converge if $\tau \leq 1/8$. Thereafter, the solution to the ROF model may be obtained from the relation (5.16).

In [15], Chambolle proposes instead to use a simple gradient projection method for the dual problem. This leads to the iteration

$$\boldsymbol{p}^{n+1} = \frac{\boldsymbol{p}^n + \tau \nabla \left(\frac{1}{\lambda} \operatorname{div} \boldsymbol{p}^n + f \right)}{\max\left(1, \left| \boldsymbol{p}^n + \tau \nabla \left(\frac{1}{\lambda} \operatorname{div} \boldsymbol{p}^n + f \right) \right| \right)}. \tag{5.20}$$

It was noted by Aujol in [7] that this is actually an instance of a general duality-based method known as the Bermùdez–Moreno algorithm. From this, Aujol shows that the scheme converges provided that $\tau < 1/4$.

5.3 Primal-dual methods

The final class of algorithms that we will now consider are designed to exploit both the primal and dual formulations of the ROF model. In the general case, these methods are known as proximal forward-backward splitting [23], and they are studied in the context of variational image processing by Chambolle and Pock in [17]. Unlike the dual methods in the previous section, the following primal-dual approach is general enough to deal with a range of variational image processing problems, including the TV-based models for colour image restoration that were given in Chapter 4.

Therefore, following Chambolle and Pock in [17], we adopt a slightly more general point of view. That is, let X and Y be two finite-dimensional real vector spaces, and consider the problem of minimising

$$\min_{u \in X} \; F(Ku) + G(u). \tag{5.21}$$

Here G is a convex function on X, F is convex on Y, and $K : X \to Y$ is a linear operator.

Notice that the ROF model (5.1) is a special case of the above problem where $F(\cdot) = \|\cdot\|_1$, $G(\cdot) = \|A \cdot - f\|_2^2$, and $K = \nabla$. However, this general formulation also encompasses a number of other interesting problems, including basis pursuit and other ℓ^1-regularised problems, see [17].

Before we can present the primal-dual algorithm for solving the minimisation problem (5.21), we recall two important definitions from convex duality theory (see, e.g., [63], [64]). The first lies at the heart of duality theory for convex functions.

Definition 5.1 (Convex conjugate)**.** *Let $f : \mathbb{R}^n \to \bar{\mathbb{R}}$. The* convex conjugate *of f is the function $f^* : \mathbb{R}^n \to \bar{\mathbb{R}}$ defined by*

$$f^*(y) := \sup_{x \in X} \left\{ (x, y) - f(x) \right\}. \tag{5.22}$$

Roughly speaking, the convex conjugate f^* describes the convex hull of epi f in terms of its so-called supporting hyperplanes. These are the hyperplanes containing epi f such that at the point $(x, f(x))$ lies on the hyperplane for some $x \in X$.

The conjugacy operation $f \mapsto f^*$, also called the *Fenchel transform*, defines a one-to-one correspondence on the set of proper closed convex functions. In other words, two proper closed convex functions $f : \mathbb{R}^n \to \bar{\mathbb{R}}$ and $g : \mathbb{R}^n \to \bar{\mathbb{R}}$ are said to be in *duality* if

$$f = g^* \quad \text{and} \quad g = f^*.$$

Second, the primal-dual algorithms for (5.21) rely heavily on the following concept from convex analysis.

Definition 5.2 (Proximal mapping)**.** *Let $f : \mathbb{R}^n \to \bar{\mathbb{R}}$ be convex. The* proximal mapping *of f is defined by*

$$x = \operatorname{prox}_{\tau f}(\bar{x}) := \arg\min_x \left\{ \frac{\|x - \bar{x}\|^2}{2\tau} + f(x) \right\}. \tag{5.23}$$

This is, in a sense, a generalisation of the notion of projection. In fact, if $\delta_C(x)$ is the indicator function of a convex set C, then $\operatorname{prox}_{\tau\delta_C}(\cdot)$ is precisely the mapping that projects $x \in X$ onto C. Using the subdifferential calculus for convex functions, it is possible to show that this mapping is well defined for any $x \in \mathbb{R}^n$ whenever f is closed and convex. Moreover, the following relation, which is known as Moreau's identity, provides a way of computing the proximal mapping of the convex conjugate f^*,

$$x = \operatorname{prox}_{\tau f}(x) + \tau \operatorname{prox}_{1/\tau f^*}(x/\tau). \tag{5.24}$$

Now, we return to the convex minimisation problem (5.21). We claim that it is equivalent to the following saddle-point problem

$$\min_{u \in X} \max_{p \in Y} \langle Ku, p \rangle + G(u) - F^*(p), \tag{5.25}$$

and the corresponding dual problem

$$\max_{p \in Y} -(G^*(-K^*p) + F^*(p)). \tag{5.26}$$

We do not derive this here, but refer instead to Example 11 in [64].

Chambolle and Pock [17] define the following primal-dual algorithm for finding a saddle-point of (5.25). Choose $\tau, \sigma > 0$, $\theta \in [0, 1]$, an initial guess $(u^0, p^0) \in X \times Y$, and set $\bar{u}^0 = u^0$. Then, for $n = 0, 1, 2, \ldots$ perform the following iteration

$$\begin{cases} p^{n+1} = \operatorname{prox}_{\sigma F^*}(p^n + \sigma K \bar{u}^n) \\ u^{n+1} = \operatorname{prox}_{\tau G}(u^n - \tau K^* p^{n+1}) \\ \bar{u}^{n+1} = u^{n+1} + \theta(u^{n+1} - u^n) \end{cases} \tag{5.27}$$

Each iteration of this algorithm consists of three steps. Roughly speaking, the first two correspond to gradient descent steps in the dual and primal variables, respectively. For $\theta = 1$, the third step is a simple linear extrapolation based on the previous iterates. A convergence analysis for this case is given in [17]. For the choice $\theta = 0$, the algorithm (5.27) is, in fact, the Arrow–Hurwicz method [4], a classical algorithm from convex analysis. Chambolle and Pock also discuss the relation to several other algorithms.

The key idea behind the above algorithm is that the proximal mapping is usually inexpensive to evaluate. This means that each iteration in itself requires relatively few computations. Moreover, by updating both the primal and dual variables, the convergence is faster than the pure primal or pure dual methods. Chambolle and Pock [17] also discuss various ways in which to choose the step sizes τ and σ in (5.27) to accelerate the convergence. However, this assumes that either F or G has some additional regularity, e.g., that one of the functionals is uniformly convex. We will, however, not pursue this further, in particular, because the inpainting and demosaicking models in Section 4.3 and Section 4.4 do not have this regularity.

Proximal mappings

What remains is to describe the proximal mappings for the functionals involved in the discrete TV-based regularisation methods in the beginning of

this chapter. For the ROF and VROF models, the operator K in (5.25) is $K = \nabla$. We consider first the regularisation terms. For the ROF model, we have $F(\cdot) = \|\cdot\|_1$, and the proximal mapping is given by

$$\operatorname{prox}_{\tau F}(u) = \max\left(0, 1 - \frac{\tau}{|u|}\right) u.$$

Using the Moreau's identity (5.24), the proximal mapping of the dual F^* is the projection onto the unit ball in the dual norm

$$\operatorname{prox}_{\sigma F^*}(\boldsymbol{p}) = \frac{\boldsymbol{p}}{\max(1, |\boldsymbol{p}|)}. \tag{5.28}$$

For the VROF model, this is replaced by

$$\operatorname{prox}_{\sigma F^*}(\boldsymbol{p}) = \frac{\boldsymbol{p}}{\max\left(1, \sqrt{\sum_{k=1}^{P} |\boldsymbol{p}_k|^2}\right)}. \tag{5.29}$$

For the data fidelity functional, the proximal mapping is given by

$$\operatorname{prox}_{\tau G}(u) = (I + \tau\lambda A^* A)^{-1}(u + \tau\lambda A^* f). \tag{5.30}$$

Although this is not necessarily simple to compute, we are only interested in a few special cases. First, for the pure denoising problem $A = I$, this reduces to

$$\operatorname{prox}_{\tau G}(u) = \frac{u + \tau\lambda f}{1 + \tau\lambda}.$$

Second, there is the case when the operator A is a convolution with some point-spread function h, i.e., $Au = h * u$. If the convolution is computed with periodic boundary conditions, it is well known that the convolution operator is diagonalised by the Fourier transform. Therefore, the solution to the equation (5.30) can be written

$$\operatorname{prox}_{\tau G}(u) = \mathcal{F}^{-1}\left(\frac{\mathcal{F}(u) + \tau\lambda\mathcal{F}(h)^*\mathcal{F}(f)}{1 + \tau\lambda\mathcal{F}(h)^2}\right). \tag{5.31}$$

By using the FFT to compute the forward and inverse Fourier transforms, this calculation can be made very fast.

Finally, for the data fidelity functional used in the image inpainting problem (4.8), the proximal mapping becomes

$$\operatorname{prox}_{\tau G}(u) = \begin{cases} u^k_{i,j} & \text{if } (i,j,k) \in \Gamma, \\ \left(u^k_{i,j} + \tau\lambda f^k_{i,j}\right) / (1 + \tau\lambda) & \text{otherwise.} \end{cases}$$

The primal-dual algorithm for the TV demosaicking model

The TV demosaicking model (5.4) requires some special treatment, so we will explicitly describe the primal-dual algorithm for this model. Let $X = \mathbb{R}^{MNP}$, $U = \mathbb{R}^{MN}$, and define $Y = X \times X$ and $V = U \times U$. Here the dualisation is performed with respect to both the luminance $\Phi \in U$ and the chrominance $\Psi \in X$, as explained in the following.

Let $K : X \to V \times Y$ denote the mapping of the image u to the gradient of the luminance and chrominance components, i.e.,

$$K : u \mapsto (\nabla\Phi, \nabla\Psi).$$

If we define the mappings $L : u \mapsto \Phi$ and $C : u \mapsto \Psi$, we may alternatively write

$$K = \begin{pmatrix} \nabla \circ L \\ \nabla \circ C \end{pmatrix}, \qquad K^* = -\begin{pmatrix} L^* \circ \operatorname{div} & C^* \circ \operatorname{div} \end{pmatrix}.$$

Next, for $u \in X$, $\boldsymbol{v} \in V$ and $\boldsymbol{w} \in Y$, we define

$$F(\boldsymbol{v}, \boldsymbol{w}) := \|\boldsymbol{v}\|_1 + \mu \sum_{i,j} \sqrt{\sum_k \left|\boldsymbol{w}_{i,j}^k\right|^2}, \tag{5.32}$$

$$G(u) := \frac{\lambda}{2} \sum_{(i,j,k)\in\Omega\setminus\Gamma} (u_{i,j}^k - f_{i,j}^k)^2. \tag{5.33}$$

With these definitions, the TV demosaicking model (5.4) is now on the form of the minimisation problem (5.21). Moreover, this is equivalent to the primal-dual formulation (5.25), where the convex conjugate of F is

$$F^*(\boldsymbol{p}, \boldsymbol{q}) = \delta_{\mathcal{V}}(\boldsymbol{p}, \boldsymbol{q}), \quad \mathcal{V} := \{(\boldsymbol{p}, \boldsymbol{q}) \in V \times Y \mid |\boldsymbol{p}| \le 1, |\boldsymbol{q}| \le 1\}.$$

The proximal mappings of F^* and G are then given by

$$\operatorname{prox}_{\sigma F^*}(\boldsymbol{p}, \boldsymbol{q}) = \left(\frac{\boldsymbol{p}}{\max(1, |\boldsymbol{p}|)}, \frac{\boldsymbol{q}}{\max\left(1, \sqrt{\sum_{k=1}^P |\boldsymbol{q}_k|^2}\right)} \right) \tag{5.34}$$

$$\operatorname{prox}_{\tau G}(u) = \begin{cases} u_{i,j}^k & \text{if } (i,j,k) \in \Gamma, \\ \left(u_{i,j}^k + \tau\lambda f_{i,j}^k\right) / (1 + \tau\lambda) & \text{otherwise.} \end{cases} \tag{5.35}$$

Thus, we may now apply the primal-dual algorithm (5.27) for the TV demosaicking model.

5.4 Comparison of numerical methods

We end this chapter by comparing some of the numerical methods that we have described. For this purpose, we consider two different numerical experiments that use the ROF model for denoising a 2-D grayscale image. The point here is not so much the actual result of the denoising, but rather to consider the convergence of the different methods.

To begin with, we illustrate the effect of the smoothing parameter β in (5.5) on the convergence of the primal gradient descent method. Consider the problem of denoising the grayscale image in Figure 6.1(a) to which an additive white Gaussian noise of standard deviation $\sigma = 0.1$ was added.

During each iteration, the step size α in (5.8) is chosen by a backtracking line search method to ensure a sufficient decrease in the primal functional. The details of such a procedure can be found in [20]. The initial guess is set to be the observed image $u^{(0)} = f$, and three different choices for smoothing parameter

Method	$\epsilon = 10^{-4}$	$\epsilon = 10^{-5}$	$\epsilon = 10^{-6}$
Chambolle	61 (2.82s)	995 (30.72s)	9297 (241.09s)
Bermudez–Moreno	19 (0.92s)	262 (8.12s)	1979 (39.89)
Arrow–Hurwicz	21 (1.01s)	36 (1.32s)	56 (1.44)
Chambolle–Pock	21 (1.00s)	36 (1.23s)	56 (1.43)

Table 5.1: *Convergence rate for dual and primal-dual methods for the ROF model for denoising a 256×256 grayscale image.*

were tested, $\beta = 1$, 0.01, and 0.001. The norm of the gradient functional (5.6) is shown in Figure 5.2(a). Notice that for smaller values of β, the line search quickly restricts the step size. Thus, the convergence becomes very slow as β is decreased.

Generally, the higher-order primal methods converge in much fewer iterations than the gradient descent. This is illustrated in Figure 5.2(b) for Newton's method. A similar fast convergence also occurs for the primal-dual Newton's method, but it is somewhat less sensitive to the choice of β, see [20]. However, these schemes soon become too costly for larger images because of the need to solve a linear system during each iteration.

As we will now illustrate, the dual and primal-dual schemes do not have these problems. They do not require us to solve a linear system, and the convergence rate is generally faster than the primal gradient descent.

Our experiment is similar to that in [17], but much simplified. Consider a grayscale image of 256×256 pixels to which is added white Gaussian noise with a standard deviation of $\sigma = 0.1$. For the sake of comparison, a ground truth solution is computed with $\lambda = 16$ using the Chambolle–Pock algorithm running 100 000 iterations. Next, we compute the solution using the various dual and primal-dual schemes. The iterations are stopped once the error relative to the ground truth solution falls below some tolerance. The algorithms were tested with the following settings:

- Chambolle: $\tau = 1/8$.
- Bermùdez–Moreno: $\tau = 1/4$.
- Arrow–Hurwicz: $\tau = 0.01$, $\sigma = 12.5$, $\theta = 0$.
- Chambolle–Pock: $\tau = 0.01$, $\sigma = 12.5$, $\theta = 1$.

The results are presented in Table 5.1.

The results clearly indicates that the primal-dual methods converge in much fewer iterations than the pure dual methods, especially whenever a higher accuracy is desired. Observe that the Arrow–Hurwicz and Chambolle–Pock methods seem to perform almost identically. This indicates that for the current problem, the additional extrapolation step in (5.27) has little impact on the convergence of the primal-dual method. However, we have only used fixed step sizes in our trials. Chambolle and Pock [17] consider different ways of accelerating the convergence by selecting the step sizes appropriately during each iteration. With this sort of acceleration, they observed that the Arrow–Hurwicz method is generally the fastest.

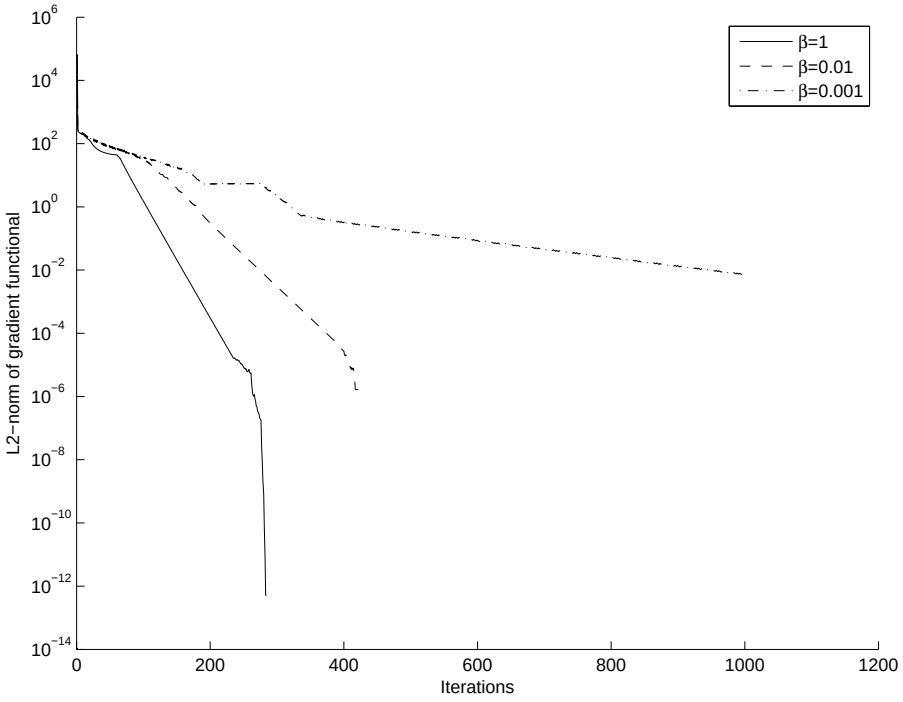

(a) Gradient descent

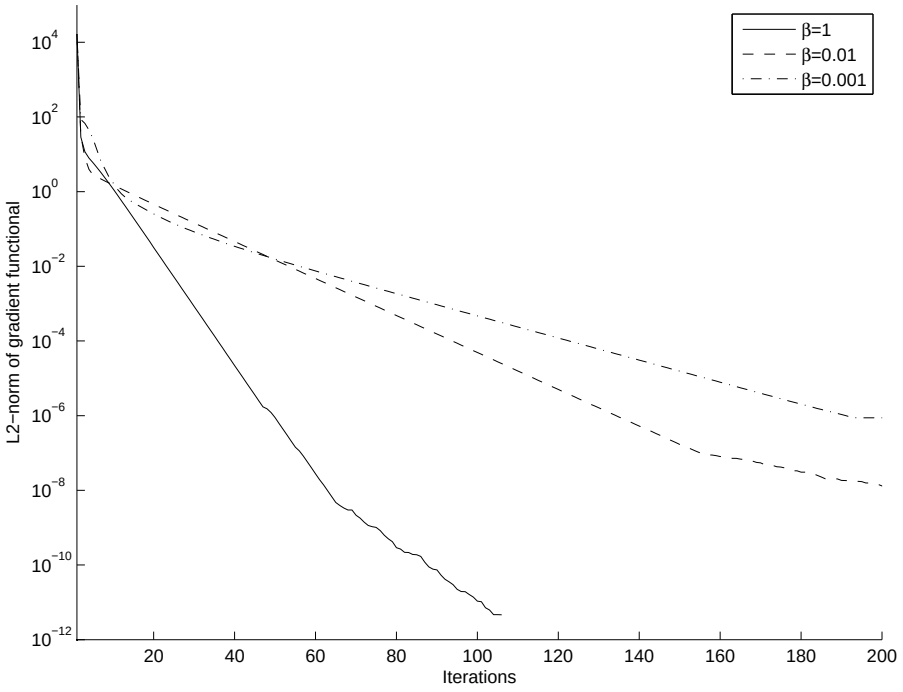

(b) Newton's method

Figure 5.2: *Convergence of primal methods.*

5.5 Summary

In this chapter we have presented the discretisation of the TV-based regularisation models from Chapter 4, and some numerical algorithms for solving them. Primal methods attempt to directly minimise the relevant cost functional using familiar algorithms from numerical optimisation, such as gradient descent or Newton's method. However, due to the ℓ^1-norm that appears in the regularisation term, the primal problem is non-differentiable. This means that we can only solve an approximate problem by smoothing the original functional. Generally, the convergence of the primal gradient descent is very slow. On the other hand, dual methods do not require any smoothing because the dual problem is differentiable. For the image denoising problem, Chambolle's algorithm [14] is a gradient projection in the dual variable, which gives a much faster convergence than the primal gradient descent. Finally, we have considered the primal-dual methods of Chambolle and Pock [17] that exploit the saddle-point formulation of the given problem. During each iteration, both the primal and dual variables are updated, yielding even faster convergence than the pure primal or pure dual approach. Some numerical evidence was also supplied to support this. In addition, this method is flexible enough to incorporate a variety of regularisation and data fidelity terms. Based on this general framework, the primal-dual method for the TV demosaicking model (5.4) was described in particular.

CHAPTER 6

Applications to image restoration problems

In this chapter we present and discuss the results of a number of image restoration experiments performed with the ROF model (4.2), the vectorial ROF model (4.6), the TV inpainting model (4.8), and also the TV demosaicking model (4.16). We begin with some standard denoising, deblurring and inpainting problems which have been demonstrated many times in the literature. However, include them here as stepping stones on the way to other image restoration tasks, and also because it illustrates the capacity and flexibility of TV-based regularisation. Next, we consider the problems of chromatic aberration and demosaicking for single-sensor digital colour cameras. The former can be treated as a deblurring problem where the point-spread function is different for each colour channel, whereas the latter is solved with the TV demosaicking model from Section 4.4.

In order for the results in this chapter to be reproducible, the Matlab code that is used in the following numerical experiments has been made available at `http://folk.uio.no/jamesdt/`.

The images used in this chapter are standard test images. Most of these are available from the USC-SIPI database (`http://sipi.usc.edu/database/`). In addition, a set of 24 colour images from the Kodak image database are used in the evaluation of demosaicking algorithms. These images are 512×768 pixels, and they depict a variety of scenes, as shown in Figure 6.6. Note that in the following experiments, the image intensities are scaled to lie in the range $[0, 1]$.

Comparing image restoration algorithms is a difficult subject. Sometimes it is possible to compare a given restoration to the original image, whenever that is available. However, this evaluation is usually subjective. The best image restoration is the one which is most visually pleasing. On the other hand, there exist a number of objective criteria that attempt to quantify how good a given restoration is compared to the original image. We will review some of these below. It is important to note that such criteria are only a convenient substitute for subjective evaluation, and, therefore, they should ideally reflect the same findings as with a subjective evaluation.

This difficulty is highlighted in the context of demosaicking by Longére *et al.* in [49]. Here the authors describe an experimental setup in which participants compared and evaluated images demosaicked by various methods. The result was a ranking of those demosaicking algorithms from best to worst. However, it was observed, perhaps surprisingly, that this ordering disagreed with that predicted by the standard S-CIELAB colour image metric [80]. Still, objective performance criteria can provide useful information and they are very commonly used in the literature. They must not be used blindly, but should instead be balanced by some form of visual inspection. Therefore, we have attempted to supplement our results in terms of objective performance measures with images that demonstrate the effects of a given method.

Now, let us present some objective performance criteria. Let $f \in \mathbb{R}^{MN}$ denote a grayscale image of M-by-N pixels, and let $\tilde{f} \in \mathbb{R}^{MN}$ be an approximation to f. The mean-square error (MSE) is given by

$$MSE := \frac{1}{MN} \sum_{i,j=1}^{M,N} (f_{i,j} - \tilde{f}_{i,j})^2. \tag{6.1}$$

The peak signal-to-noise (PSNR) ratio is a related logarithmic measure

$$PSNR := 10 \log_{10} \left(\frac{1}{\frac{1}{MN} \sum_{i,j=1}^{M,N} (f_{i,j} - \tilde{f}_{i,j})^2} \right). \tag{6.2}$$

This can be extended to P-channel colour images $f, \tilde{f} \in \mathbb{R}^{MNP}$. This is known as the colour peak signal-to-noise ratio (CPSNR) [2]

$$CPSNR := 10 \log_{10} \left(\frac{1}{\frac{1}{MNP} \sum_{i,j,k=1}^{M,N,P} (f_{i,j}^k - \tilde{f}_{i,j}^k)^2} \right). \tag{6.3}$$

6.1 Denoising and deblurring

The examples in this section are considered to be fairly standard in the variational image processing literature and they have been demonstrated many times for the ROF and VROF models.

Our first example is one of denoising a grayscale image for which the ROF model is known to produce very good results. Consider the image shown in Figure 6.1(a). An additive white Gaussian noise with a standard deviation of $\sigma = 0.5$ is added, producing the noisy image shown in Figure 6.1(b). Note that such a large amount of noise, the standard deviation corresponding to half of the maximum image intensity, represents a quite extreme case. The image restored by the ROF model with $\lambda = 1$ is shown in Figure 6.1(c). The results are quite pleasing since the noise has essentially been removed and the edges of the objects are still quite prominent.

Next, we consider the same experiment, but performed instead on the selected natural images shown in Figure 6.2. To each image, additive white Gaussian noise with a standard deviation of $\sigma = 0.1$ is added. Then the image is restored using the ROF model with the regularisation parameter $\lambda = 15$ which was determined experimentally. Once more, the ROF model succeeds in removing noise while recovering the edges of the image. However, since the

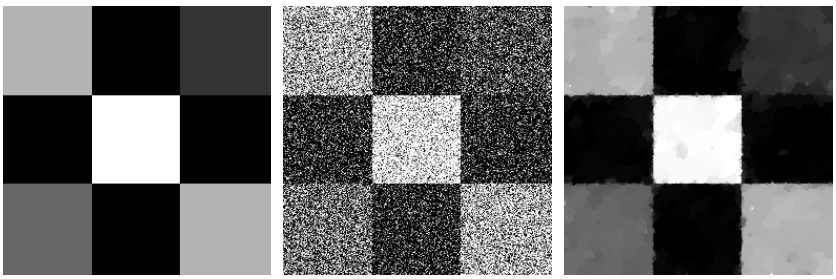

(a) (b) (c)

Figure 6.1: *Denoising a synthetic image. To the image on the left there was added a white Gaussian noise of standard deviation $\sigma = 0.5$. The result of denoising using the ROF model with $\lambda = 2$ is shown on the right.*

ROF model inherently focuses on restoring the geometric parts of the images, some of the finer details are lost in the process of removing the noise.

A final example of denoising a colour image is shown in Figure 6.3. Here the noise is additive white Gaussian-distributed with a standard deviation of $\sigma = 0.2$. Two results are shown. The first is of denoising by applying the ROF model to each channel individually, and the second comes from solving the VROF model where the colour channels are treated jointly. The latter is slighly better in terms of PSNR. Although the visible difference is not very large, the VROF model is slightly better at restoring some finer details and texture, e.g., the wall to the left in the image.

Next, we evaluate the ROF model for deblurring a grayscale image. For this, we follow the numerical experiments in [77] and [78]. This involves comparing the ROF model to built-in functions from the MATLAB Image Processing Toolbox. These are "`deconvreg`", "`deconvwnr`", and "`deconvlucy`". The algorithms implemented by these functions are well known in the image processing literature, see, e.g., [35]. To give a short description, `deconvreg` solves the Tikhonov regularisation problem (see eq. (1.3) in Section 1.2) with the Laplacian operator for regularisation term, `deconvwnr` computes the Wiener-filter restoration that minimises the mean-square error between the estimated image and the original, and `deconlucy` implements the Richardson-Lucy allgorithm, a constrained maximum likelihood problem for Poisson-distributed noise. Further details may be found in [78] and the references therein.

Consider the 512×512 grayscale image in Figure 6.4(a). The image was blurred with a PSF that is a Gaussian low-pass filter

$$K(x, y) = e^{-(x^2+y^2)/(2c^2)}, \tag{6.4}$$

with $c = 11$, and a square support that is 21 pixels wide. A Gaussian white noise with a standard deviation of $\sigma = 10^{-3}$ was added to the blurred image, and the result can be seen in Figure 6.4(b). The results of the image deblurring are shown next in Figure 6.4. The `deconvwnr` result exhibits quite a large amount of noise, whereas the `deconvlucy` and `deconvreg` show some "ringing" artefacts. For the ROF model, the regularisation parameter was chosen to be $\lambda = 5 \times 10^4$. In terms of both PSNR and visual quality, the ROF model shows

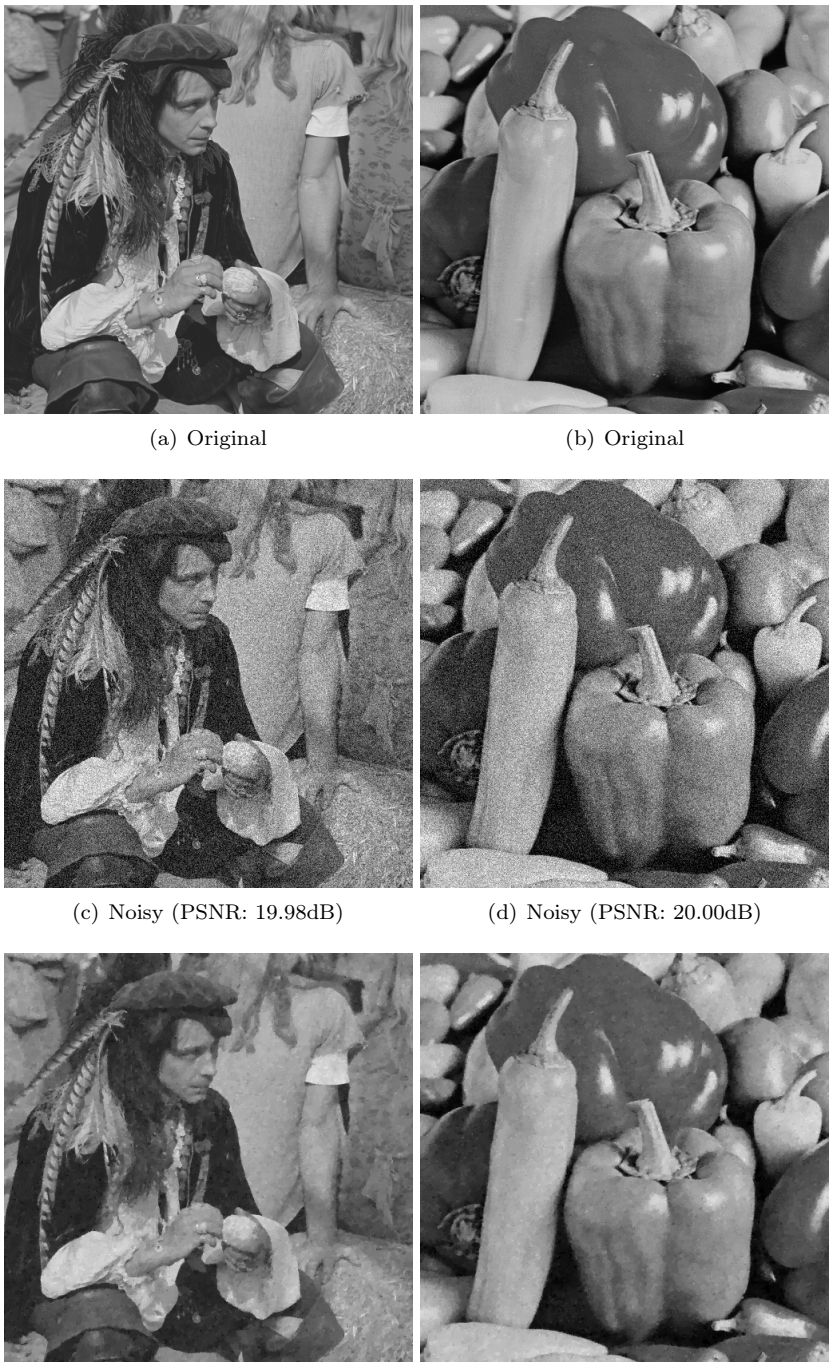

(a) Original (b) Original

(c) Noisy (PSNR: 19.98dB) (d) Noisy (PSNR: 20.00dB)

(e) Restoration (PSNR: 28.34dB) (f) Restoration (PSNR: 29.57dB)

Figure 6.2: *Denoising natural images. The first row displays the original grayscale images and the middle row shows the noisy images with additive white Gaussian noise with a standard deviation of $\sigma = 0.1$. The final row shows the restored images using the ROF model with a regularisation parameter of $\lambda = 15$. Below is the PSNR that was achieved by the restoration.*

(a) Original (b) Noisy (13.97 dB)

(c) Channel-by-channel ROF (23.73 dB) (d) VROF model (24.11 dB)

Figure 6.3: *Denoising colour images. The first row displays the original and noisy colour image. The noise is additive white Gaussian with a standard deviation of* $\sigma = 0.2$*. Below are the respective denoising results that come from applying the ROF model separately to each colour channel with* $\lambda = 6$*, and from applying the VROF model with* $\lambda = 4.1$ *to denoise the colour channels jointly. The PSNR is displayed below each restored image.*

the best performance. As in our previous examples, the restoration produced by the ROF model is able to represent discontinuities in the image.

6.2 Chromatic aberration

In this section, we consider the restoration of an image suffering from chromatic aberration. This may be treated as a colour image deblurring problem. Consider the colour image shown in Figure 6.5(a). To simulate the effect of axial chromatic aberration (see Figure 2.5(b)), the three colour channels are blurred by varying amounts. The PSF applied to each channel is a uniform disk with a radius of 10 pixels for the red and blue channels, and only 2 pixels for the green channel. Thereafter, additive white Gaussian noise was added with standard deviations $\sigma_1 = 0.001$ and $\sigma_2 = 0.01$ for two different examples. The degraded images, shown in Figure 6.5(b) and Figure 6.5(c), clearly display the effects of colour fringing.

Here we have only treated each colour channel independently and applied the ROF model to denoise and deblur each of them. The regularisation parameters for the two cases were chosen to be $\lambda = 50\,000$ and $\lambda = 1000$. The full restorations are shown in Figure 6.5(d) and Figure 6.5(e). The effects of the colour fringing are reversed, and since the amount of noise was not too severe, even much of the texture is recovered. This is true also for the noisier case, although a small amount of colour fringing can still be observed.

6.3 Demosaicking

In this section, we consider the demosaicking problem in the context of the variational methods we have studied. Recall that a single-sensor digital colour camera employs a colour filter array (CFA) to assign each pixel to one of a number of possible colour channels. Although variational methods are able to deal with any CFA configuration, we restrict our attention to the standard Bayer CFA (see Figure 2.2(a)). Two different regularisation approaches are compared. The first is to treat each colour channel independently and attempt to inpaint the missing pixels. This leads to unsatisfactory results that motivates the need for a better regularisation. Hence, the second approach uses the TV demosaicking model from Section 4.4 which is based on applying the regularisation in the luminance-chrominance basis.

First, let us consider a simple experiment that consists of inpainting a vertical edge that has been subsampled according to the green colour channel of the Bayer CFA. Figure 6.7(a) shows the subsampled image where every second pixel is missing. As we have already demonstrated in Example 4.11, the TV inpainting model (4.8) and the harmonic inpainting model (4.10) behave quite differently in this kind of situation. We expect the former to fill in the gaps in a smooth manner, whereas the TV inpainting favours a piecewise constant solution. The result of the inpaintings is shown in Figure 6.7(b) and Figure 6.7(c). The TV inpainting recovers the edge exactly, while the harmonic inpainting produces a zipper-like effect along the edge with alternating light and dark pixels.

In the red colour channel of the Bayer CFA, this effect is even more pronounced, see Figure 6.7. Here, we can also see that the TV inpainting does

(a) Original (b) Blurred and noisy (PSNR: 20.97dB)

(c) `deconvwnr` (PSNR: 22.73dB) (d) `deconvreg` (PSNR: 28.53dB)

(e) `deconvlucy` (PSNR: 22.92dB) (f) ROF model (PSNR: 30.16dB)

Figure 6.4: *Image deblurring. The original image (a) is blurred by a Gaussian PSF* (6.4) *with* $c = 11$*, and white Gaussian noise with a standard deviation* $\sigma = 10^{-3}$ *is added to produce the image in (b). The subsequent images show the restoration with (c)* `deconvwnr`*, (d)* `deconvreg`*, (e)* `deconvlucy`*, and (f) the ROF model.*

(a) Original

(b) $\sigma = 0.001$ (PSNR: 23.13dB)

(c) $\sigma = 0.01$ (PSNR: 23.04dB)

(d) Restoration (PSNR: 32.74dB)

(e) Restoration (PSNR: 27.67dB)

Figure 6.5: *Chromatic aberration. To simulate the effects of axial chromatic aberration, the red, green and blue colour channels of the original image (a) are blurred by varying amounts. The two images in the middle row display the blurred image with different amounts of noise added. The bottom row shows the resulting restoration from applying the ROF model to each colour channel separately.*

Figure 6.6: *The Kodak image database. This is a standard database of 24 colour images, often used to evaluate demosaicking algorithms. The images are referred to as no. 1–24, numbered from left to right and top to bottom.*

not precisely recover the edge in the red channel either. Instead, there is a transition across the edge which is one pixel wide.

It is clear that if either of these methods is used to naively inpaint the colour channels separately, the resulting demosaicking suffers some quite severe artefacts in the form of false colours and zipper effect. This can be improved by applying the VROF model (4.6) so that the colour channels are restored jointly. An example of this approach is displayed in Figure 6.8. Unfortunately, this method is not particularly successful, and the resulting demosaicking shows false colours especially near thin structures or in textured regions. This is seen clearly in Figure 6.8(b).

The luminance-chrominance decomposition of Alleysson *et al.* [2] is an improvement over many of the demosaicking algorithms that work in the usual RGB basis. This is also the motivation behind our TV-based demosaicking model (4.16) that applies the regularisation in the luminance-chrominance basis. Using the primal-dual algorithm from Section 5.3, we now consider the performance of this new regularisation.

Note that the algorithm in [2] is based on designing linear filters for estimating the luminance and chrominance components of the image. In practice, the optimal filter will actually depend on the underlying image. More precisely,

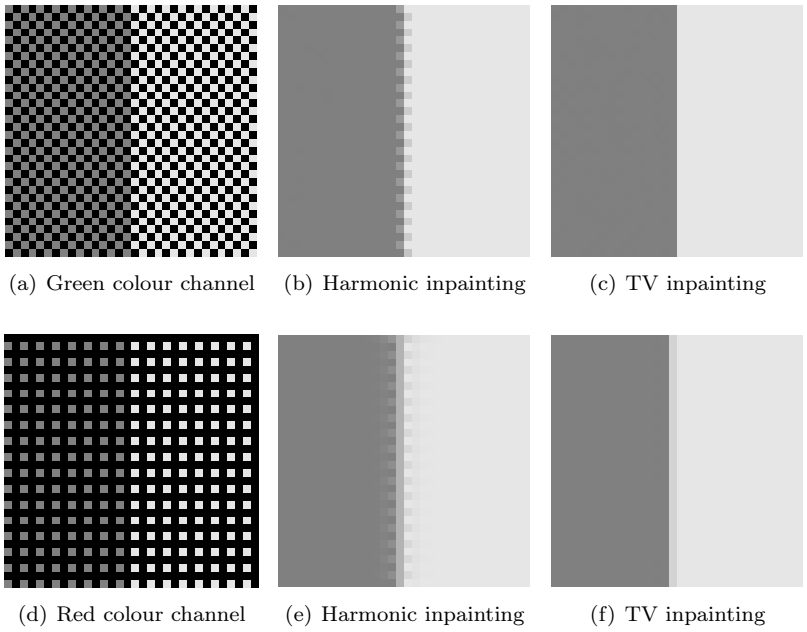

(a) Green colour channel (b) Harmonic inpainting (c) TV inpainting

(d) Red colour channel (e) Harmonic inpainting (f) TV inpainting

Figure 6.7: *Inpainting a vertical edge in the green and red colour channels of the Bayer CFA. The harmonic inpainting model* (4.10) *results in a zipper-like effect near edges. Although the TV inpainting model* (4.8) *represents an improvement, it is not able to entirely recover the edge in the red colour channel. This illustrates the problem of demosaicking by inpainting each colour channel separately.*

it will depend on exactly how smooth the luminance and chrominance components are. Therefore, it is necessary to strike a balance and select filters that perform well on average for a set of images. On the other hand, for the regularisation approach this can be more explicitly controlled by the parameters of the regularisation model. That is, the parameter μ in the TV demosaicking model (4.16) determines how much regularisation is applied to the luminance relative to the chrominance. For example, if an image happens to have a more regular chrominance component, μ should be chosen larger. The point is that these parameters are much easier to adapt to the underlying image compared to the design of a linear filter, as in the approach of Alleysson *et al.* in [2].

Now, we compare the method in [2] to our proposed TV demosaicking model. The entire set of 24 images from the Kodak database were subsampled according to the Bayer CFA and subsequently demosaicked with the TV demosaicking model. For this, the parameter λ that weights the data fidelity term was set to $\lambda = +\infty$. This corresponds to exactly enforcing the data fidelity constraint which is appropriate in the noiseless case. The chrominance regularity parameter was determined experimentally for each image. The resulting CPSNR values are shown in Table 6.1. In the same table, the results from demosaicking by bilinear interpolation and the method in [2] are also shown (these results were obtained from [29]).

For most of the images, the TV demosaicking model represents an improvement in CPSNR over the method of Alleysson *et al.*. Although the average

(a) VROF model

(b) Zoom-in

Figure 6.8: *Demosaicking with the VROF model. The image in (a) displays a demosaicking of image no. 5 from the Kodak test set. The demosaicking was performed by solving VROF model* (4.6) *with the corresponding operator A that models the effects of the Bayer CFA. The image in (b) is a zoomed-in region that clearly shows the appearance of false colours due to the demosaicking. This is particularly apparent near textures or thin structures such as the spokes of the bikes.*

Image no.	Bilinear interpolation	Alleysson *et al.* [2]	TV-based demosaicking
01	26.35	36.47	39.30
02	33.11	37.72	38.15
03	34.53	40.03	41.41
04	33.76	38.99	39.06
05	26.72	35.92	35.94
06	27.86	37.45	38.65
07	33.53	39.91	39.89
08	23.61	32.69	35.05
09	32.51	40.07	40.86
10	32.51	40.99	40.30
11	29.22	38.02	37.97
12	33.53	41.02	42.35
13	23.92	34.94	35.90
14	29.28	33.78	35.05
15	32.55	38.25	37.90
16	31.38	40.62	42.42
17	32.15	41.14	40.05
18	28.05	36.68	35.72
19	28.11	37.56	38.87
20	31.65	39.71	40.02
21	28.55	37.88	38.81
22	30.46	36.96	36.83
23	35.20	39.87	40.75
24	26.71	34.96	33.26
Average	30.22	37.99	38.52

Table 6.1: *Demosaicking results. The table displays the CPSNR of the demosaicked images from the Kodak test set for three methods (i) bilinear interpolation of each colour channel, (ii) the frequency domain luminance-chrominance estimation method of Alleysson et al. [2], and (iii) the proposed TV demosaicking model from Section 4.4.*

improvement is only about 0.5dB, for some images, e.g., no. 1 and no. 8, the improvement is more significant. In addition, a zoomed-in region of the demosaicking of image no. 3 from the Kodak test set is shown in Figure 6.9. This shows that the TV demosaicking model does not suffer from the same zipper-like artefacts that appear with the method [2], especially near the boundaries of objects where there is a transition in colour.

Still, the main advantage of the regularisation-based approach is that it can be used to simultaneously perform denoising and demosaicking. This is very important in practice for a real single-sensor digital colour camera. It has been argued, e.g., by Hirakawa and Parks [42], that considering denoising and demosaicking simultaneously has some benefits. An obvious benefit, at least for the the regularisation methods used here, is that no additional computational effort is required to perform this denoising. Though perhaps more important is the fact that the process of demosaicking changes the characteristics of the noise, thereby complicating the analysis and any subsequent denoising algorithm.

With this in mind, we provide a final example of demosaicking in the pres-

(a) Frequency domain method of Alleysson *et al.* [2]

(b) TV demosaicking model

Figure 6.9: *Demosaicking artefacts. The images display a zoomed-in region of the demosaicking of image no. 3 from the Kodak test set. The top image displays the result from the frequency domain luminance-chrominance estimation method of Alleysson et al. [2], and the bottom is the proposed TV demosaicking model. The top image exhibits a certain zipper-like effect that occurs with the method in [2] especially near colour edges. The proposed TV demosaicking model does not suffer from the same artefact.*

ence of noise. For this, we take a single image from the Kodak test set, subsampled according to the Bayer CFA and add Gaussian white noise of standard deviation $\sigma = 0.02$. The demosaicking is computed as above, but we have now set the data fidelity parameter to $\lambda = 1000$. The chrominance regularity parameter was set to $\mu = 1.46$. The noisy CFA image and the demosaicking are shown in Figure 6.10. The CPSNR of the resulting restoration is 34.70dB. We compare this to the same experiment performed by Menon and Calvagno in [54] with a different regularisation approach based on a quadratic regularisation term. For the same image, Menon and Calvagno reported a CPSNR of 32.95dB. Although this is only demonstrated for a single image, we might expect the TV regularisation to give a similar improvement in other cases.

6.4 Summary

In this chapter, we have covered a number of applications of TV-based regularisation methods to image restoration. Although using the ROF and VROF models for denoising and deblurring is fairly standard in the literature, nonetheless, they illustrate the power and versatility of these methods. These examples all illustrate the main characteristic of the TV-based regularisation, that is to say, its ability to restore edges in the degraded images. More importantly, we have considered two image restoration problems directly related to digital colour cameras. The first is an example of deblurring a colour image suffering from axial chromatic aberration. The results shown here are more a proof-of-concept that illustrate the feasibility of the TV regularisation approach. Second, we considered the TV demosaicking model from Section 4.4, and demonstrated improved results compared to the approach of Alleysson *et al.* [**?**].

(a) Noisy CFA image ($\sigma = 0.02$)

(b) TV-based demosaicking (CPSNR: 34.70dB)

Figure 6.10: *Demosaicking in the presence of noise. The image in (a) shows image no. 1 in the Kodak test set, subsampled according to the Bayer CFA and with additive white Gaussian noise of $\sigma = 0.02$. Below is the demosaicking of this image with the proposed TV demosaicking model.*

CHAPTER 7

Conclusion

Variational methods, based on minimising some cost or energy functional over a set of feasible solutions, represent a rather flexible approach to solving image restoration problems. By regularising the given problem, one can incorporate prior information about the solution. In a sense, this allows one to describe in rough terms what the restored image should look like.

The total variation-based methods in this thesis are an example of this approach. In this case, the regularisation is designed to favour piecewise constant solutions, and the resulting image restoration methods tend to restore both large-scale features and edges in the image. This is desirable because edges are particularly important visual cues. In the discrete case, such a regularisation is achieved by finding a solution with a nearly sparse gradient. This is in turn related to minimising the ℓ^1-norm of the gradient.

The numerical experiments in the previous chapter demonstrated this for a number of image restoration problems, including denoising, deblurring, chromatic aberration and demosaicking. The treatment of chromatic aberration was more a proof-of-concept to demonstrate that the TV-based regularisation approach also is able to deal with this effect.

In the case of demosaicking, the usual regularisation provided by the vectorial ROF model gave unsatisfactory results. In particular, the demosaicked images displayed false colours in textured regions or near thin structures. The proposed TV demosaicking model, based on the regularisation of the luminance and chrominance components of the image, resulted in an improved restoration in this case. Also, a certain zipper-like effect appearing near chromatic edges with the method of Alleysson *et al.* [2] was also suppressed with the proposed method.

In practice, images captured by a digital colour camera are always corrupted by noise. Therefore, the regularisation-based approach offers an additional advantage since it is already designed to account for this noise. The numerical experiments in the previous chapter show that the TV-based regularisation is successful also for deblurring, chromatic aberration and demosaicking in the presence of noise.

Future work

There is certainly more to be said with respect to comparing the proposed TV demosaicking model with other demosaicking methods. In particular, it would be interesting to further compare it to similar regularisation-based methods, especially for demosaicking in the presence of noise.

Also, the treatment of chromatic aberration only considered axial chromatic aberration. Although the case of lateral chromatic aberration could in theory be treated in the same framework, this cannot be realistically modelled as a spatially invariant convolution. Therefore, the operator in the forward image model (1.1) that is associated with this degradation does not have the same special structure as a convolution operator. In general, the matrix representation of such an operator consumes a lot of memory and can be difficult to manage. Even worse is the fact that one cannot exploit efficient FFT-based algorithms, such as for computing the proximal mapping (5.31) in the primal-dual method for deblurring in Chapter 5. However, in special cases, the degradation operator might have some other structure that can be exploited. Or, one might consider approximating such an operator by subdividing the image into a number of sub-images and applying convolutions with different kernels to each sub-image.

Although we have limited our discussion to regularisation based on the total variation functional, there are, of course, other possibilities. In particular, a number of transforms, including curvelets [13] and shearlets [39], have been introduced to provide sparse descriptions of images. It is easy to imagine that a regularisation based on one of these transforms is likely to give improved results. This has already been demonstrated for some image restoration problems, such as denoising and inpainting, via so-called frame-based methods [12]. This could be extended in particular to demosaicking and compared to the luminance-chrominance regularisation in this thesis.

Bibliography

[1] R. Acar and C.R. Vogel, *Analysis of bounded variation penalty methods for ill-posed problems*, Inverse Problems **10** (1994), no. 6, 1217–1229.

[2] D. Alleysson, S. Süsstrunk, and J. Hérault, *Linear demosaicing inspired by the human visual system*, Image Processing, IEEE Transactions on **14** (2005), no. 4, 439–449.

[3] L. Alvarez, F. Guichard, P.-L. Lions, and J.-M. Morel, *Axioms and fundamental equations of image processing*, Archive for Rational Mechanics and Analysis **123** (1993), 199–257.

[4] K.J. Arrow, L. Hurwicz, and H. Uzawa, *Studies in linear and non-linear programming*, vol. 2, Stanford University Press, 1958.

[5] G. Aubert and P. Kornprobst, *Mathematical problems in image processing: Partial differential equations and the calculus of variations*, Applied Mathematical Sciences, vol. 147, Springer, 2006.

[6] G. Aubert and L. Vese, *A variational method in image recovery*, SIAM Journal on Numerical Analysis **34** (1997), no. 5, 1948–1979.

[7] J.-F. Aujol, *Some first-order algorithms for total variation based image restoration*, Journal of Mathematical Imaging and Vision **34** (2009), no. 3, 307–327.

[8] B.E. Bayer, *Color imaging array*, U.S. Patent 3,971,065, 1976.

[9] P. Blomgren and T.F. Chan, *Color TV: Total variation methods for restoration of vector-valued images*, Image Processing, IEEE Transactions on **7** (1998), no. 3, 304–309.

[10] M. Born and E. Wolf, *Principles of optics*, sixth ed., Pergamon Press, 1980.

[11] X. Bresson and T.F. Chan, *Fast dual minimization of the vectorial total variation norm and applications to color image processing*, Inverse problems and Imaging **2** (2008), no. 4, 455–484.

[12] J.F. Cai, R.H. Chan, and Z. Shen, *A framelet-based image inpainting algorithm*, Applied and Computational Harmonic Analysis **24** (2008), no. 2, 131–149.

[13] E.J. Candes and D.L. Donoho, *Curvelets: A surprisingly effective nonadaptive representation for objects with edges*, Tech. report, Stanford University, Ca, Department of Statistics, 2000.

[14] A. Chambolle, *An algorithm for total variation minimisation and applications*, Journal of Mathematical Imaging and Vision **20** (2004), no. 1–2, 89–97.

[15] ———, *Total variation minimization and a class of binary MRF models*, 2005, pp. 136–152.

[16] A. Chambolle and P.-L. Lions, *Image recovery via total variation minimization and related problems*, Numerische Mathematik **76** (1997), no. 2, 167–188.

[17] A. Chambolle and T. Pock, *A first-order primal-dual algorithm for convex problems with applications to imaging*, Journal of Mathematical Imaging and Vision **40** (2011), no. 1, 120–145.

[18] T. Chan, A. Marquina, and P. Mulet, *High-order total variation-based image restoration*, SIAM J. Sci. Comput. **22** (2000), no. 2, 503–516.

[19] T. Chan and J. Shen, *Image processing and analysis: Variational, PDE, wavelet, and stochastic methods*, Society for Industrial and Applied Mathematics, 2005.

[20] T.F. Chan, G.H. Golub, and P. Mulet, *A nonlinear primal-dual method for total variation-based image restoration*, SIAM J. Sci. Comput. **20** (1999), no. 6, 1964–1977.

[21] T.F. Chan, H.M. Zhou, and R.H. Chan, *Continuation method for total variation denoising problems*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 2563, 1995, pp. 314–325.

[22] S.S. Chen, D.L. Donoho, and M.A. Saunders, *Atomic decomposition by basis pursuit*, SIAM Journal on Scientific Computing **20** (1998), no. 1, 33–61.

[23] P.L. Combettes and V.R. Wajs, *Signal recovery by proximal forward-backward splitting*, Multiscale Modeling & Simulation **4** (2005), no. 4, 1168–1200.

[24] L. Condat, *A generic variational approach for demosaicking from an arbitrary color filter array*, Image Processing (ICIP), 2009 16th IEEE International Conference on, 2009, pp. 1625–1628.

[25] J. Darbon and M. Sigelle, *Image restoration with discrete constrained total variation part i: Fast and exact optimization*, Journal of Mathematical Imaging and Vision **26** (2006), no. 3, 261–276.

[26] Mauricio Delbracio, Pablo Musé, Andrés Almansa, and Jean-Michel Morel, *The non-parametric sub-pixel local point-spread function estimation is a well posed problem*, International Journal of Computer Vision **96** (2012), no. 2, 175–194.

[27] D.L. Donoho, *De-noising by soft-thresholding*, Information Theory, IEEE Transactions on **41** (1995), no. 3, 613–627.

[28] ———, *For most large underdetermined systems of linear equations the minimal ℓ^1-norm solution is also the sparsest solution*, Communications on pure and applied mathematics **59** (2006), no. 6, 797–829.

[29] E. Dubois, *Frequency-domain methods for demosaicking of bayer-sampled color images*, Signal Processing Letters, IEEE **12** (2005), no. 12, 847–850.

[30] L.C. Evans, *Partial differential equations*, second ed., Graduate Studies in Mathematics, vol. 19, American Mathematical Society, 2010.

[31] L.C. Evans and R.F. Gariepy, *Measure theory and fine properties of functions*, CRC, 1992.

[32] E. Giusti, *Minimal surfaces and functions of bounded variation*, Monographs in Mathematics, vol. 80, Birkhauser, 1984.

[33] B. Goldluecke, E. Strekalovskiy, and D. Cremers, *The natural vectorial variation which arises from geometric measure theory*, SIAM J. Imag. Sciences **5** (2012), no. 2, 537–563.

[34] T. Goldstein and S. Osher, *The split Bregman method for L1-regularized problems*, SIAM Journal on Imaging Sciences **2** (2009), no. 2, 323–343.

[35] R.C. Gonzalez and R.E. Woods, *Digital image processing*, third ed., Pearson Education, 2008.

[36] M.D. Grossberg and S.K. Nayar, *Determining the camera response from images: What is knowable?*, Pattern Analysis and Machine Intelligence, IEEE Transactions on **25** (2003), no. 11, 1455–1467.

[37] ———, *Modeling the space of camera response functions*, Pattern Analysis and Machine Intelligence, IEEE Transactions on **26** (2004), no. 10, 1272–1282.

[38] B.K. Gunturk, J. Glotzbach, Y. Altunbasak, R.W. Schafer, and R.M. Mersereau, *Demosaicking: Color filter array interpolation*, Signal Processing Magazine, IEEE **22** (2005), no. 1, 44–54.

[39] K. Guo and D. Labate, *Optimally sparse multidimensional representation using shearlets*, SIAM journal on mathematical analysis **39** (2007), no. 1, 293–318.

[40] G.E. Healey and R. Kondepudy, *Radiometric CCD camera calibration and noise estimation*, Pattern Analysis and Machine Intelligence, IEEE Transactions on **16** (1994), no. 3, 267–276.

[41] K. Hirakawa, *Cross-talk explained*, Image Processing, 15th IEEE International Conference on, 2008, pp. 677–680.

[42] K. Hirakawa and T.W. Parks, *Joint demosaicing and denoising*, Image Processing, IEEE Transactions on **15** (2006), no. 8, 2146–2157.

[43] K. Hirakawa and P.J. Wolfe, *Spatio-spectral color filter array design for optimal image recovery*, Image Processing, IEEE Transactions on **17** (2008), no. 10, 1876–1890.

[44] H. Holden and N.H. Risebro, *Front tracking for hyperbolic conservation laws*, Applied Mathematical Sciences, vol. 152, Springer, 2011.

[45] J.R. Janesick, *Scientific charge-coupled devices*, SPIE Press, 2001.

[46] A. Kirsch, *An introduction to the mathematical theory of inverse problems*, second ed., Springer, 2011.

[47] S.Z. Li, *Markov random field modeling in image analysis*, Springer, 2009.

[48] C. Liu, R. Szeliski, S.B. Kang, C. Zitnick, and W. Freeman, *Automatic estimation and removal of noise from a single image*, Pattern Analysis and Machine Intelligence, IEEE Transactions on **30** (2008), no. 2, 299–314.

[49] P. Longére, X. Zhang, P.B. Delahunt, and D.H. Brainard, *Perceptual assessment of demosaicing algorithm performance*, Proceedings of the IEEE **90** (2002), no. 1, 123–132.

[50] F. Luisier, T. Blu, and M. Unser, *Image denoising in mixed Poisson–Gaussian noise*, Image Processing, IEEE Transactions on **20** (2011), no. 3, 696–708.

[51] R. Lukac and K.N. Plataniotis, *Color filter arrays: Design and performance analysis*, Consumer Electronics, IEEE Transactions on **51** (2005), no. 4, 1260–1267.

[52] M. Makitalo and A. Foi, *Optimal inversion of the Anscombe transformation in low-count Poisson image denoising*, Image Processing, IEEE Transactions on **20** (2011), no. 1, 99–109.

[53] A. Marquina and S. Osher, *Explicit algorithms for a new time dependent model based on level set motion for nonlinear deblurring and noise removal*, SIAM Journal on Scientific Computing **22** (2000), no. 2, 387–405.

[54] D. Menon and G. Calvagno, *Regularization approaches to demosaicking*, Image Processing, IEEE Transactions on **18** (2009), no. 10, 2209–2220.

[55] ———, *Color image demosaicking: An overview*, Signal Processing: Image Communication **26** (2011), 518–533.

[56] D. Mumford and J. Shah, *Optimal approximations by piecewise smooth functions and associated variational problems*, Communications on pure and applied mathematics **42** (1989), no. 5, 577–685.

[57] J.R. Munkres, *Topology*, second ed., Pearson, 1999.

[58] J. Nakamura (ed.), *Image sensors and signal processing for digital still cameras*, CRC Press, 2006.

[59] K.A. Parulski, *Color filters and processing alternatives for one-chip cameras*, Electron Devices, IEEE Transactions on **32** (1985), no. 8, 1381–1389.

[60] M.A. Pinsky and S. Karlin, *An introduction to stochastic modelling*, fourth ed., Academic Press, 2011.

[61] R. Ramanath, W.E. Snyder, Y. Yoo, and M.S. Drew, *Color image processing pipeline*, Signal Processing Magazine, IEEE **22** (2005), no. 1, 34–43.

[62] E. Reinhard, E.A. Khan, A.O. Akyz, and G.M. Johnson, *Color imaging: Fundamentals and applications*, A.K. Peters, Ltd., 2008.

[63] R.T. Rockafellar, *Convex analysis*, Princeton University Press, 1970.

[64] ———, *Conjugate duality and optimization*, Society for Industrial and Applied Mathematics, 1974.

[65] L.I. Rudin, S. Osher, and E. Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D: Nonlinear Phenomena **60** (1992), no. 1–4, 259–268.

[66] W. Rudin, *Principles of mathematical analysis*, third ed., McGraw–Hill, 1976.

[67] ———, *Functional analysis*, second ed., McGraw–Hill, 1991.

[68] T. Saito and T. Komatsu, *Sharpening-demosaicing method with a total-variation-based super-resolution technique*, Digital Photography, vol. 5678, Proc. SPIE, no. 177, 2005, pp. 1801–1812.

[69] ———, *Spatially adaptive super-resolution sharpening-demosaicing for a single solid-state color image sensor*, Digital Photography II (N. Sampat, J.M. DiCarlo, and R.A. Martin, eds.), Proc. SPIE, vol. 6069, 2006.

[70] ———, *Demosaicing method using the extended color total-variation regularization*, Digital Photography IV (J.M. DiCarlo and B.G. Rodricks, eds.), Proc. SPIE, vol. 6817, 2008.

[71] G. Sharma (ed.), *Digital color imaging handbook*, CRC Press, 2002.

[72] J. Shen and T.F. Chan, *Mathematical models for local nontexture inpaintings*, SIAM Journal on Applied Mathematics **62** (2002), no. 3, 1019–1043.

[73] D. Strong and T. Chan, *Edge-preserving and scale-dependent properties of total variation regularization*, Inverse Problems **19** (2003), no. 6, S165–S187.

[74] A. Tarantola, *Inverse problem theory and methods for model parameter estimation*, Society for Industrial and Applied Mathematics, 2005.

[75] A.N. Tikhonov and V.Y. Arsenin, *Solutions of ill-posed problems*, V.H. Winston and Sons, 1977.

[76] C.R. Vogel and M.E. Oman, *Iterative methods for total variation denoising*, SIAM Journal on Scientific Computing **17** (1996), no. 1, 227–238.

[77] C. Wu and X.C. Tai, *Augmented Lagrangian method, dual methods, and split Bregman iteration for ROF, vectorial TV, and high order models*, SIAM Journal on Imaging Sciences **3** (2010), no. 3, 300–339.

[78] J. Yang, W. Yin, Y. Zhang, and Y. Wang, *A fast algorithm for edge-preserving variational multichannel image restoration*, SIAM Journal on Imaging Sciences **2** (2009), no. 2, 569–592.

[79] Y.-L. You and M. Kaveh, *A regularization approach to joint blur identification and image restoration*, Image Processing, IEEE Transactions on **5** (2002), no. 3, 416–428.

[80] X. Zhang and B.A. Wandell, *Color image fidelity metrics evaluated using image distortion maps*, Signal processing **70** (1998), no. 3, 201–214.