Physically-Based Image Reconstruction

Dissertation zur Erlangung des Grades des Doktors der Ingenieurwissenschaften der Naturwissenschaftlich-Technischen Fakultäten der Universität des Saarlandes

vorgelegt von

Nico Persch

Saarbrücken, 2016

Mathematische Bildverarbeitungsgruppe
Fakultät für Mathematik und Informatik,
Universität des Saarlandes, 66041 Saarbrücken

Tag des Kolloquiums
26.07.2016

Dekan
Prof. Dr. Frank-Olaf Schreyer

Prüfungsausschuss

Prof. Dr. Matthias Hein (Vorsitzender)
Universität des Saarlandes, Deutschland

Prof. Dr. Joachim Weickert (1. Gutachter)
Universität des Saarlandes, Deutschland

Prof. Dr. Martin Welk (2. Gutachter)
Universität Hall/Tyrol, Österreich

Dr. Pascal Peter (akademischer Mitarbeiter)
Universität des Saarlandes, Deutschland

Copyright

Kurzzusammenfassung

Die Rekonstruktion gestörter oder verlorengegangener Daten ist eine der wesentlichen Herausforderungen der Bildverarbeitung und des Maschinensehens. In dieser Arbeit konzentrieren wir uns auf den physikalischen Bildgebungsprozess und nähern ihn mittels sogenannter Vorwärtsoperatoren an. Wir betrachten die Umkehrung dieser mathematisch stichhaltigen Formulierungen als das Ziel von Rekonstruktion. Wir gehen diese Aufgabe mit Variationsverfahren an, wobei wir unsere Methoden auf die spezifischen physikalischen Grenzen und Schwächen der verschiedenen Bildgebungsverfahren ausrichten. Der erste Teil dieser Arbeit betrifft die Bildverarbeitung. Wir stellen eine weiterentwickelte Rekonstruktionsmethode für 3-D Konfokal und STED (*stimulated emission depletion*) Mikroskopiebilder vor. Hierzu vereinigen wir Bildentrauschung, Entfaltung und anisotropes Einfüllen. Der zweite Teil betrifft maschinelles Sehen: Wir stellen eine neue Depth-from-Defocus Methode vor und entwerfen einen neuen Vorwärtsoperator, der wichtige physikalische Eigenschaften wahrt. Unser Operator passt gut in ein variationelles Gerüst. Zudem zeigen wir die Vorteile etlicher weitergehender Konzepte, wie eine gemeinsame Depth-from-Defocus- und Entrauschungsmethode, sowie Robustifizierungsstrategien. Außerdem zeigen wir die Vorteile des multiplikativen Euler-Lagrange Formalismus gegenüber dem additiven. Synthetische und reale Experimente in den Hauptkapiteln bestätigen die Anwendbarkeit und das Vermögen unserer Methoden.

Short Abstract

The reconstruction of perturbed or lost data is one of the fundamental challenges in image processing and computer vision. In this work, we focus on the physical imaging process and approximate it in terms of so-called forward operators. We consider the inversion of these mathematically sound formulations to be the goal of reconstruction. We approach this task with variational techniques where we tailor our methods to the specific physical limitations and weaknesses of different imaging processes. The first part of this work is related to image processing. We propose an advanced reconstruction method for 3-D confocal and *stimulated emission depletion* (STED) microscopy imagery. To this end, we unify image denoising, deconvolution and anisotropic inpainting. The second part is related to computer vision: We propose a novel depth-from-defocus method and design a novel forward operator that preserves important physical properties. Our operator fits well into a variational framework. Moreover, we illustrate the benefits of a number of advanced concepts such as a joint depth-from-defocus and denoising approach as well as robustification strategies. Besides, we show the advantages of the multiplicative Euler-Lagrange formalism compared to the additive one. Synthetic and real-world experiments within the main chapters confirm the applicability and the performance of our methods.

Zusammenfassung

Die Rekonstruktion gestörter oder verlorengegangener Information ist eine der wesentlichen Herausforderungen der Bildverarbeitung und des Maschinensehens. Der Bildgebungsprozess, der die 3-D Welt in ein 2-D Bild projiziert, ist ein gutes Beispiel für solch einen Informationsverlust. In dieser Arbeit entwickeln wir Modelle physikalischer Bildgebungsprozesse mittels mathematisch stichhaltiger Formulierungen – sogenannter Vorwärtsoperatoren. Solch ein Operator beschreibt formell, wie die echte Welt in ein Bild abgebildet wird. Dementsprechend präsentieren wir in dieser Dissertation verschiedene Verwirklichungen von Vorwärtsoperatoren und besprechen ihre Vor- und Nachteile. In dieser Arbeit verstehen wir unter dem Begriff der Rekonstruktion die Umkehrung des Vorwärtsoperators. Wir zeigen, wie man diese Aufgabe mit Variationsmodellen angeht und erklären, wie unsere Methoden auf die spezifischen physikalischen Grenzen und Schwächen der betrachteten Bildgebungsverfahren ausgerichtet werden können. In diesem Zusammenhang machen wir deutlich, auf welche Art und Weise der physikalische Prozess der Diffusion eine wichtige Rolle spielt. Wir wenden unsere Rekonstruktionsideen in zwei verschiedenen Bereichen an: Im Rahmen der Bildverarbeitung stellen wir eine weiterentwickelte Rekonstruktionsmethode vor, die speziell für 3-D Konfokal- und STED (*stimulated emission depletion*) Mikroskopie maßgeschneidert ist. Unsere Methode vereinigt Bildentrauschung, Entfaltung und Interpolation in einem gemeinsamen Ansatz. Dies ermöglicht uns, die speziellen Schwächen dieser Mikroskope zu bewältigen: Defokussierte Unschärfe, Poisson Rauschen und geringe axiale Auflösung. Demzufolge schlagen wir die Kombination von (i) Richardson-Lucy Entfaltung, welche besonders geeignet ist bezüglich des Rauschmodells, (ii) Bild-Restaurierung und (iii) anisotropem Einfüllen, welches speziell für die Verbesserung von länglichen Zellstrukturen konstruiert ist, in einem einzigen Modell vor. Im Rahmen des Maschinensehens schlagen wir eine neue Depth-from-Defocus Variationsmethode vor. Wir besprechen verschiedene Bildentstehungsmodelle und entwerfen einen Vorwärtsoperator, der wichtige physikalische Eigenschaften wahrt. Unser neuer Vorwärtsoperator kommt dem Dünne-Linsen Kamera-Modell nahe und passt gut in ein variationelles Gerüst. Zudem zeigen wir die Vorteile etlicher weitergehender Konzepte: Verwenden der vollen Information eines Mehrkanal-Signals, eine gemeinsame Depth-from-Defocus und Entrauschungsmethode sowie Robustifizierungs-Strategien. All diese Konzepte ermöglichen es uns, die Rekonstruktionsqualität zu verbessern. Ein weiterer wichtiger Beitrag dieser Dissertation ist die Veranschaulichung der Vorteile des multiplikativen Euler-Lagrange Formalismus in Hinblick auf die Minimierung der vorkommenden Variationsfunktionale. Dieser ist der gebräuchlichen

additiven Variante in zweierlei Hinsicht überlegen: Erstens erlaubt er uns, die Lösung auf den plausiblen positiven Bereich zu begrenzen. Zweitens ermöglicht er uns, ein semi-impliziteres Gradientenverfahren zu entwickeln, welches einen höheren Stabilitätsbereich aufweist. Synthetische und reale Experimente in den Hauptkapiteln bestätigen die Anwendbarkeit und das Vermögen unserer Methoden.

Abstract

The reconstruction of perturbed or lost information is one of the fundamental challenges in image processing and computer vision. The imaging process that projects the 3-D real world into a 2-D image is a good example for such a loss of information. In this work, we develop models of physical imaging processes in terms of mathematically sound formulations – so-called forward operators. Such an operator describes formally how the real world is mapped to an image. Thus, we present different realisations of forward operators in this dissertation and discuss their advantages and shortcomings. In this work, we understand reconstruction as the inversion of the forward operator. We show how to approach this task with variational models, and we explain how to tailor our methods to the specific physical limitations and weaknesses of the considered imaging processes. In that respect, we make obvious in which sense the physical process of diffusion plays an important role. We apply our reconstruction ideas in two different fields: In the context of image processing, we propose an advanced reconstruction method that is specifically tailored towards 3-D confocal and *stimulated emission depletion* (STED) microscopy. Our method unifies image denoising, deconvolution and interpolation in one joint approach. This allows us to handle the typical weaknesses of these microscopes: Out-of-focus blur, Poisson noise, and low axial resolution. Hence, we propose the combination of (i) Richardson-Lucy deconvolution which is especially suited for this noise model, (ii) image restoration, and (iii) anisotropic inpainting which is designed especially for the enhancement of elongated cell structures, in one single scheme. In the context of computer vision, we propose a novel variational depth-from-defocus method. We discuss different image formation models and design a forward operator that preserves important physical properties. Our novel forward operator approximates the thin lens camera model and fits well into a variational framework. Moreover, we illustrate the benefit of a number of advanced concepts: using the full information of a multi-channel signal, a joint depth-from-defocus and denoising approach, as well as robustification strategies. All these concepts allow us to improve the reconstruction results. Another important contribution of this dissertation is the demonstration of the advantages of the multiplicative Euler-Lagrange formalism regarding the minimisation of the occurring variational functionals. It is superior over the common additive one in two aspects: First, it allows us to constrain the solution to the plausible positive range. Second, it allows us to develop a more semi-implicit gradient descent scheme which exhibits a higher stability range. Synthetic and real-world experiments in the main chapters confirm the applicability and the performance of our methods.

Acknowledgements

First of all, I would like to express my gratitude to Prof. Dr. Joachim Weickert giving me this PhD position in the Mathematical Image Analysis (MIA) group, for offering me this interesting field of research, and for being my supervisor. Furthermore, I want to thank Prof. Dr. Martin Welk for the willingness and effort involved in acting as the external reviewer of my thesis.

My thanks are addressed to the German Federal Ministry for Education and Research (BMBF) who funded the first part of my work through the project 'Intracellular Transport of Nanoparticles: 3D Imaging and Stochastic Modelling'. The second part of my research has been funded by the *Deutsche Forschungsgemeinschaft (DFG)* through a *Gottfried Wilhelm Leibniz Prize* for Prof. Dr. Joachim Weickert and the Cluster of Excellence *Multimodal Computing and Interaction*. I would also like to express my sincere thanks to them.

It is a very important desire for me to express my very deep gratitude to my friends Oliver Demetz, David Hafner, Sebastian Hoffmann, Christopher Schroers, and Simon Setzer for the friendship, the support, and the scientific discussions. Besides that, I also want to extend my thanks for the spare time, the sporting activities and the lark together. I want to thank Mrs Ellen Wintringer for her help concerning organisational stuff, the friendly confabs and for regularly putting my dirty cup into the dishwasher.

Moreover, I would like to thank Prof. Dr. Andrés Bruhn, Prof. Dr. Martin Welk, and Prof. Dr. Michael Breuss for their great support and helpful guidance. I thank my current and former colleagues of the MIA group: Sven Grewenig, Pascal Gwosdek, Kai Uwe Hagenburg, Markus Mainberger, Pascal Peter, Martin Schmidt, and Laurent Hoeltgen for the fruitful discussions and the very friendly atmosphere. Further thanks are addressed to our actual system administrator Peter Franke as well as his predecessor Marcus Hargarter.

For having the understanding that I did not have much time to spend with her, and for the moments together, I thank my girlfriend Janine Schweitzer.

Finally, I want to thank my parents Anne and Karl-Heinz. They always motivated and supported me, making it possible for me to go this way. Indeed, thank you for all the things you have done for me.

Contents

Chapter 1

Introduction

The pursuit of improving imaging techniques has always been a competition against or with physics. Already in the Middle Ages, the Arabian scientist Ibn al-Haytham, also known as Alhazen, analysed the magnifying property of spherical glasses – lenses [Twyman, 1988; King, 2011]. His studies '*Book of Optics*' about vision, the physical phenomena of reflection and refraction are considered to be significant contributions in the field of early optics and to have inspired the development of future imaging techniques. About 600 years after Alhazen, Francesco Maria Grimaldi observed deviations from the ray model in the propagation of light [Hall, 1990]. The wave characteristic of light began to attract attention. Representatives of this wave theory include Pierre Ango [1682], Robert Hooke [1665] and Ignace-Gaston Pardies, see e.g. [Hall, 1990] and [Dijksterhuis, 2006]. The assumption of a wave-like character of light was later corroborated by Christiaan Huygens [1690] [Ziggelaar, 1980]. Eventually, Albert Einstein [1905] postulated the theory that light consists of energy quanta, so-called photons.

Once the principles of geometric optics were understood, one recognised that correct arrangements of several lenses allowed for the development of new optical instruments such as the telescope and the microscope. However, due to the wave character of light, the maximum resolution was considered for a long time to be physically limited by the diffraction limit as described by Ernst Abbe [1873]. Nowadays, about thousand years after the studies of Alhazen, we have knowledge about electromagnetism and quantum physics. Exploiting these theories has allowed for the development of scientific instruments such as *electron microscopes* [Knoll and Ruska, 1932] and *scanning tunnelling microscopes* (STM) [Binnig and Rohrer, 1983]. Moreover, the work of Stefan Hell has revealed a way to bypass the diffraction limit even working with visible light [Hell and Wichmann, 1994]. Although, the achievable resolution of all these techniques and the quality of acquired data

has become very impressive, one is now confronted with perturbations and weaknesses arising, for instance, from aberration of electromagnetic lenses [Goldstein et al., 2012], physical properties of the tip of the STM [Dongmo et al., 1996] or quantum phenomena such as *shot* or *Poisson* noise [Pawley, 2006; Conn, 2012].

So, in every age, researchers have been confronted with technological and physical limits and have developed clever concepts in order to increase the performance of state-of-the-art imaging devices. However, even if technological limits are pushed further and further, and some physical limits were bypassed, new ones were revealed. Hence, it seems that imaging techniques will never become perfect, which implies that perturbations will always be present in the acquired data. The different kind of degradations involved thereby depend strongly on the individual image acquisition technique. For instance, the degradations that appear in ultra-sound imaging are completely different from those appearing in methods working with visible light, other electromagnetic waves, or particles such as electrons. But even only among light microscopy based methods, one is confronted with different perturbations depending on the exact type of microscope and observed specimen. Thus, no matter how much effort is spent on the development of an image acquisition method, the captured information does not completely coincide with the real world.

However, another way to improve the quality of acquired data is post-processing: Image restoration and enhancement methods may help where acquisition methods reach their limits. Hence, these strategies, which are as old as image processing itself, are still a current topic. The ongoing goal is to further increase the gain of exploitable information and to bring it closer to the real world.

In addition to that, especially with regard to computer vision, another objective addresses the interpretation of the acquired data. Here, the aim is to analyse and interpret the captured scene based on one or several images automatically. This covers aspects such as automatic segmentation, the reconstruction of depth information and the recognition and interpretation of individual patterns, structures and objects.

1.1 Common Challenges in Image Reconstruction

Image processing and computer vision are typically confronted with certain classes of degradations and limitations. Depending on the individual image

Figure 1.1: Comparison of artificially created Gaussian distributed noise and Poisson distributed noise. **(a) Top:** Gaussian noise ($\mu = 0$) with increasing standard deviation σ_{noise} from left to right. **(b) Bottom:** Poisson noise with decreasing number photons per intensity value from left to right.

acquisition technique and capturing settings, often the captured information not only suffers from a single type of perturbation but from a combination of several types. To formulate deteriorations mathematically, we interpret an n-dimensional grey value image as a continuous function $f : \Omega_n \to \mathbb{R}$, where $\Omega_n \subset \mathbb{R}^n$ denotes the image domain. Furthermore, we denote by $\boldsymbol{x} = (x_1, \ldots, x_n)^T \in \Omega_n$ a location in the image domain. In the following, we discuss common degradations and limitations.

Noise. The phenomenon of noise constitutes a random change of signal values that might only be modelled stochastically. Most typical reasons for noise have a physical or chemical origin and usually occur directly in the photo sensor during image acquisition or in electronic circuits during signal processing. In this dissertation, we focus on two different kinds of noise: *additive white Gaussian noise* and *Poisson noise*. Figure 1.1 compares the degradation by artificially created Gaussian and Poisson distributed noise respectively.

Gaussian distributed noise mainly arises by signal amplification or within the image sensor and is, e.g. based on the thermal movement of charge carriers [Van Etten, 2006; Bovik, 2009]. Additive white Gaussian noise can be modelled as

$$f(\boldsymbol{x}) = g(\boldsymbol{x}) + \eta(\boldsymbol{x}) \,, \quad \boldsymbol{x} \in \Omega_n \,, \tag{1.1}$$

Figure 1.2: Hidden or lost signal information. **(a) Left:** Corrupted image. **(b) Right:** Mask indicating region D where information is known (white) and where information is missing (black).

where $g : \Omega_n \to \mathbb{R}_+$ denotes the original undisturbed signal, $f : \Omega_n \to \mathbb{R}$ the observed one and η the noise function. The noise function η follows a Gaussian distribution, i.e.

$$p(s) = \frac{1}{\sigma_{\text{noise}}\sqrt{2\pi}} \exp\left(-\frac{(s-\mu)^2}{2\sigma_{\text{noise}}^2}\right) \tag{1.2}$$

with mean μ and standard deviation σ_{noise} (see e.g. [Bovik, 2009]). While Gaussian noise is independent of the underlying intensity values, *shot* or *Poisson* noise does depend on them [Rangayyan, 2004]. Hence, it can generally be described by

$$f(\boldsymbol{x}) = \eta(g(\boldsymbol{x})) \ , \quad \boldsymbol{x} \in \Omega_n \ . \tag{1.3}$$

This type of noise is based on the particle nature of light and charge carriers. Thus, it occurs particularly in low light intensity imagery [Rangayyan, 2004]. For this reason, we take a closer look at Poisson noise in Section 3.2.2.

Typical denoising approaches are based on eliminating high frequencies, e.g. with the help of smoothing low-pass filters. This strategy is based on the fact that noise usually acts as a high frequency perturbation and it is based on the assumption that the original undisturbed image is piecewise smooth, i.e. it contains more low frequency components.

Partially missing information. Here, the information is only available on a subset $D \subset \Omega_n$ of the whole image domain Ω_n. In the remaining parts $\Omega_n \backslash D$, it is missing. We assume that D and Ω_n are known and that D at least partially surrounds $\Omega_n \backslash D$. If this is not the case, we would end up in an extrapolation task at the boundaries of Ω_n. Such a degradation is illustrated in Figure 1.2 and might originate from defects in the image sensor, hidden image information (due to overlapping objects, text characters

Figure 1.3: Image convolution. *From left to right:* **(a)** Input image (512×512 pixels). **(b)** Blurred with a 2-D Gaussian with $\sigma = 3.0$ (inserted in the upper left corner). **(c)** Ditto with $\sigma = 6.0$. **(d)** $\sigma = 9.0$.

or overpaintings) or dropouts in a transmission.

We will approach this problem in Section 2.3 with the help of suitable inpainting or interpolation strategies.

Blur. This phenomenon describes a blending or an averaging of neighbouring information. Common types of blur are *out-of-focus* and *motion blur*. Out-of-focus blur occurs if the light rays of an object point are not bundled to a single image point by the lens system. In contrast, motion blur occurs if the camera moves relative to the object during exposure time. Further reasons for a blurred acquisition can be atmospheric turbulences as well as imperfections of optical systems and diffraction phenomena. In this dissertation, we consider out-of-focus blur. Since blur constitutes a weighted average over some neighbourhood, the image formation can be modelled as

$$f(\boldsymbol{x}) = (H \circledast g) := \int_{\mathbb{R}^n} H(\boldsymbol{x}, \boldsymbol{y}) \cdot g(\boldsymbol{y}) \, \mathrm{d}\boldsymbol{y} \,, \tag{1.4}$$

where the spatially variant *point-spread function* (PSF) $H : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}_{0+}$ describes the weighting. If the blur does not change between different locations, i.e. if it is spatially invariant, it can be expressed with the help of the mathematical convolution operation:

$$f(\boldsymbol{x}) = (h * g)(\boldsymbol{x}) := \int_{\mathbb{R}^n} h(\boldsymbol{x} - \boldsymbol{y}) \cdot g(\boldsymbol{y}) \, \mathrm{d}\boldsymbol{y} \,, \tag{1.5}$$

where $h : \mathbb{R}^n \to \mathbb{R}_{0+}$ denotes a spatially invariant PSF. In the discrete setting, summation replaces integration. Figure 1.3 shows the effect of convolution with different 2-D Gaussian kernels with $\mu = 0$, i.e.:

$$h(\boldsymbol{x}) = \mathcal{K}_\sigma(\boldsymbol{x}) := \frac{1}{2\pi\sigma^2} \cdot e^{-\frac{-x_1^2 - x_2^2}{2\sigma^2}} \,. \tag{1.6}$$

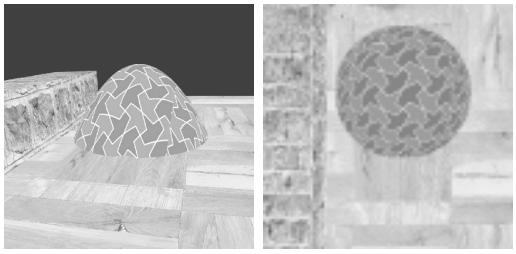

Figure 1.4: Loss of depth information. **(a) Left:** 3-D scene to be captured. **(b) Right:** 2-D image recorded by a pinhole camera placed above the scene.

To reconstruct the original sharp information, the blurring must be inverted. Such deblurring or deconvolution (if the PSF is spatially invariant) techniques are typically based on a (stabilised) inversion of the convolution operator as, e.g. proposed by Wiener [1949], variational methods as, e.g. described by Osher and Rudin [1994], Marquina and Osher [1999], Chan and Wong [1998], and You and Kaveh [1996a], or Bayesian-based schemes as suggested independently of each other by Richardson [1972] and Lucy [1974].

Loss of depth information. This constitutes rather a physical limitation than some kind of degradation and is – as its name implies – inherently coupled with 2-D photography. Generally, 2-D photography can be seen as a projective mapping of the 3-D world onto a 2-D image plane. This projection maps all world points lying on the same optical ray to the same point on the sensor. Due to the fact that such a projective mapping acts independently of the object distance to the camera, information about depth is lost. Hence, regarding Figure 1.4(b) one can only have at most an intuition of the depth profile of the scene shown in Figure 1.4(a). This intuition is based on the enormous past experiences of the human visual system. A reconstruction of 3-D information – also called *depth-map* or *topography* – based upon the data of the 2-D image 1.4(b) is not straightforwardly possible.

The reconstruction of depth information from a single 2-D image constitutes a very challenging task. The research on *shape-from-shading* methods has addressed this problem for more than 30 years. As the name implies, here the shading of objects is taken as a cue for the depth reconstruction [Horn and Brooks, 1989; Zhang et al., 1999; Vogel et al., 2009]. However the severe model assumptions limit the practical use of such approaches. Köser et al. [2011] propose a method that is applicable for the special case that the imaged scene contains symmetric structures. If at least two images varying in their perspective exist, one typically uses *stereo methods* [Alvarez et al., 2002; Slesareva et al., 2005] where the work of Marr and Poggio [1976] marks one

of the first approaches in this field of research. A quite different approach, namely, *depth-from-defocus* uses the amount of out-of-focus blur as a cue to reconstruct the depth. This idea is discussed in Chapter 4.

1.2 Scope of this Thesis

In this dissertation, our intention is to provide insights into the development of advanced reconstruction frameworks by means of incorporating essential physical properties. The aim is a specific tailoring of the reconstruction methods towards the individual systematic peculiarities and limitations of a considered imaging process. Confronted with conjunctions of the above-mentioned degradations and limitations, we illustrate the way of posing the reconstruction issue as the inversion of the associated imaging model. Obviously, the reconstruction performance depends fundamentally on the ability to approximate the physical imaging process. For this purpose, a large part of this work is devoted to exactly this task. Each main chapter starts with an analysis of the respective image acquisition technique, investigates reasons for degradations and limitations, and tries to formulate the imaging process mathematically. To this end, we use the notion of *forward operators* which in turn brings us to the question of what one should be aware of when modelling them.

The considered reconstruction tasks constitute ill-posed problems. Already with a small change in the input data, one may end up in a completely different reconstruction result. Further, the solution may be not unique. For this reason, the second main contribution of this thesis is an adequate and stable inversion strategy. Here, we focus on *variational methods* [Gelfand and Fomin, 2000; Aubert and Kornprobst, 2006]. They determine the solution as a minimiser of the discrepancy between the forward process and the captured data. The ill-posedness can be counteracted by following established regularisation schemes and by restricting the solution to only plausible values. However, not every forward operator is well suited for variational methods. Therefore, one main issue is finding a compromise between a good imitation of the physical imaging process and the feasibility of its inversion. The fact that not each detail of the image acquisition can be incorporated in turn raises the demand of advanced variational concepts such as robustification ideas.

1.3 Our Contributions

In this work, we suggest two novel reconstruction approaches: In Chapter 3 we propose an advanced method for the joint treatment of typical perturbations arising during 3-D cell recordings in *confocal laser scanning microscopy* (CLSM) [Minsky, 1988] and *stimulated emission depletion* (STED) microscopy [Hell and Wichmann, 1994]. These microscopy techniques typically suffer from Poisson noise, blur and a relative low axial resolution. This chapter is based on our journal publication [Persch et al., 2013] which in turns extends the conference publication of Elhayek et al. [2011]. Besides revisiting and discussing the ideas of the latter paper in a very detailed way, we extend this model in three aspects: First, we replace the isotropic regularisation and inpainting term of Elhayek et al. [2011] by an anisotropic operator. We explain how to exploit the physical process of diffusion in order to reach a smoothing behaviour that is tailored especially towards the elongated structures of the cell filament. For this purpose, we replace the scalar-valued diffusivity by a tensor-valued quantity, the so-called diffusion tensor [Weickert, 1998]. In this way, we are able to steer the diffusion process along the single fibres of the cell filament. Besides regularisation and inpainting, this strategy improves the handling of an incomplete fluorescence labelling. Second, a novel semi-implicit numerical scheme is proposed. It is more robust and – for the isotropic case – at the same time faster compared to its predecessor [Elhayek et al., 2011]. More precisely, our scheme allows larger amounts of regularisation and reaches a significantly faster convergence behaviour. Third, we perform extensive numerical experiments by evaluating our approach on real-world CLSM as well as STED images. A comparison with competing methods in the literature validates the suitability of our modifications.

Chapter 4 is based on our conference publication [Persch et al., 2014], as well as our submitted technical report [Persch et al., 2015]. Here, we address the depth-from-defocus problem and propose a novel approach that incorporates important physical properties. For many existing methods, the reconstruction of strong depth changes constitutes a main challenge. While a common remedy is given by imposing local equifocal assumptions – local patches of constant depth – we go a quite different way. Our idea is based on the modelling of a novel forward operator. It simulates the depth-of-field like the thin lens model very accurately. This is realised by additionally preserving a maximum-minimum principle w.r.t. the unknown image intensities. Moreover, we show how to embed this operator in a variational framework and derive its minimality conditions. For this purpose, we advocate the use of the *multiplicative* Euler-Lagrange formalism. This way, the solution can

be constrained to the physically plausible range and the ill-posedness of the problem can be mitigated. Furthermore, it allows us to follow a more stable and more efficient semi-implicit, gradient descent strategy similar to that presented in Chapter 3. Besides that, we explain how to handle multi-channel focal stacks, analyse the impact of robustification and propose a novel joint variational denoising and depth-from-defocus approach. We illustrate the achieved improvements by synthetic and real-world experiments. In that context, we also make abundantly clear the difference between depth-from-defocus and standard 3-D deconvolution.

1.4 Our Methodology

1.4.1 Forward Operator

As we interpret reconstruction as the inversion of the imaging process, we first need to formulate the image formation along with its degradations mathematically. This is exactly the task of the forward operator $\mathcal{F}$. It serves as an approximative mathematically sound formulation in order to describe the physical image formation process. Given the original undisturbed information (of an object to be captured) as argument, the outcome of the forward operator should resemble the result of the real imaging system. However, in our case it is only half the story because the second main requirement is that the forward operator has to be easily invertible. More precisely, it should be applicable within a variational framework. Therefore, the main issue in designing a suitable forward operator is finding a compromise between physical accuracy and computational feasibility. This leads us to a discussion about physical properties as well as specific imaging characteristics that should be incorporated or could be neglected, and others that must be guaranteed. Such investigations form the basis within each of our reconstruction frameworks and have to be done individually:

A forward operator that reasonably approximates the image formation process of CLSM and STED microscopy should imitate *Poisson* distributed noise and 3-D blur. A proper simplification of this is the assumption of spatially invariant blur. In this case, blurring can then be described by means of mathematical convolution of the sharp information with a 3-D point-spread function (PSF). The PSF describes the redistribution of a point light source and can be estimated physically within a separate processing step.

Regarding the depth-from-defocus problem, a suitable forward operator has to simulate the local out-of-focus blur based on the local depth information. Here, the energy distribution at strong depth changes constitutes a

severe challenge.

1.4.2 Variational Methods

Once we have understood the physical imaging process and found a way to formulate it reasonably by a mathematical forward operator $\mathcal{F}$, we need a suitable strategy to go the inverse direction. More precisely, we want to estimate the arguments to be applied to $\mathcal{F}$ approximating the captured data. As already mentioned, in contrast to the forward direction, however, such inverse problems are typically ill-posed and require a more elaborate technique. Variational methods offer an elegant way for exactly such tasks [Bertero et al., 1988; Gelfand and Fomin, 2000; Aubert and Kornprobst, 2006] and have found their way into a broad area of image processing and computer vision applications such as image restoration [Whittaker, 1923; Tikhonov, 1963; Rudin et al., 1992], optical flow [Horn and Schunck, 1981; Brox et al., 2004], image segmentation [Mumford and Shah, 1989; Ambrosio and Tortorelli, 1992; Morel and Solimini, 1994; Chan and Vese, 2001], and deconvolution techniques [Marquina and Osher, 1999; Chan and Wong, 1998; You and Kaveh, 1996a]. The principles of variational methods consist in determining the solution as a minimiser of an appropriate energy formulation. It is also possible to couple different problems in one joint functional. By sophisticated regularisation techniques, variational methods are able to handle the problem of ill-posedness. Moreover, we can modify the regularisation strategy such that a desired smoothing behaviour can be achieved.

To explain the principle of variational methods, let us now consider a general energy functional E of the form

$$E(u_1, \dots, u_p) = \underbrace{E_{\mathsf{D}}(f, u_1, \dots, u_p, \mathcal{F})}_{\text{data term}} + \alpha \cdot \underbrace{E_{\mathsf{S}}(u_1, \dots, u_p)}_{\text{smoothness term}} . \tag{1.7}$$

This functional is composed of two terms as well as a parameter α that balances both. The first term E_{D} is the so-called *data term*. It penalises the discrepancy between the captured data f and its approximation provided by the forward operator $\mathcal{F}$. The aim is to find the p functions $(u_1, \dots, u_p)$ of $\mathcal{F}$ such that the discrepancy attains its minimum. Within the data term, we consider only the part of the forward operation that acts deterministically. The remaining stochastic part as well as perturbations that are not regarded by the forward operator, e.g. calibration or detector problems are treated as noise. In order to cope with such noise, several discrepancy measures and penalisation strategies for the data term are proposed in the literature. While quadratic penalisation is especially suited for Gaussian distributed noise,

Csiszár's information divergence (I-divergence) [Csiszár, 1991] is suitable for Poisson noise [Steidl and Teuber, 2010; Welk, 2015]. Besides, robust data terms as already proposed in [Zervakis et al., 1995; Bar et al., 2005; Bruhn et al., 2005; Welk, 2010, 2015] tackle small deviations in the imaging model [Huber, 2004]. These deviations may be originated by the above-mentioned noise, i.e. complex degradations that cannot be simulated exactly by the forward operator.

Considering the data term on its own, the minimiser may be not unique. To handle this problem, the *smoothness term* denoted by E_{S} demands (piecewise) smoothness of the solution as a second constraint. This is usually enforced by penalising first or higher order derivatives of the unknown. In this way, variational methods profit from a filling in effect caused by regularisation at locations where no (in the case of inpainting) or not enough information is available.

In order to find a suitable minimiser of such a variational functional and by that a solution of the corresponding inverse problem, let us now introduce the *multiplicative Euler-Lagrange* formalism.

1.4.3 Multiplicative Euler-Lagrange Formalism

In the last section, we have discussed the variational principles using a general energy functional (see Equation 1.7). To obtain a suitable minimiser of such a functional, the *additive Euler-Lagrange* formalism constitutes a very popular technique and is usually the first choice. Less known than this classical approach is its *multiplicative* counterpart, which presents itself as a very promising alternative especially for our work.

Both, additive and multiplicative Euler-Lagrange variants require that the minimising functions fulfil the associated Euler-Lagrange equations whose demand is a vanishing variational gradient [Gelfand and Fomin, 2000]. To derive the variational gradient, the classical approach considers an additive perturbation of the functional with a test function. Here, the sought minimiser is not constrained to any specific range. Instead, in the multiplicative formalism, the perturbation is done multiplicatively. It behaves like a substitution with an exponential function and thus it constraints the solution to the positive range [Welk, 2010, 2015]. Furthermore, the associated Euler-Lagrange equations are eventually supplemented by a multiplication with the unknown function. Therefore, in searching for a suitable minimiser, we can pursue more semi-implicit gradient descent schemes offering both higher stability and efficiency. Moreover, it enables us to interpret the *Richardson-Lucy* (RL) deconvolution method [Richardson, 1972; Lucy, 1974] as an iterative

minimisation scheme for a specific functional. This allows us to combine the advantages of the variational calculus with those of RL deconvolution. By its adaptation to Poisson statistics, RL deconvolution is especially suited for low intensity imagery such as CLSM and STED microscopy.

1.5 Organisation of this Thesis

This thesis is organised as follows: In Chapter 2, we revisit some established basic reconstruction methods. The first part of this chapter starts with recapitulating the ideas behind image diffusion filtering in Section 2.1. We consider the physical process of diffusion and explain its powerfulness for image denoising. Isotropic as well as anisotropic diffusion filters are discussed and compared. We illustrate their continuous modelling as well as the way to their discrete counterparts. After that, we revisit variational image restoration and the classical additive Euler-Lagrange formalism in Section 2.2. The resulting *partial differential equation* (PDE) brings us to the approach of PDE-based inpainting in Section 2.3. Combining both challenges in one variational functional is the topic of Section 2.4, where we consider the idea of joint inpainting and restoration. The second part of this chapter is devoted to deconvolution methods and their comparison. To make the reader familiar with such methods, we briefly present the idea behind *inverse filtering* as well as the *Wiener filter* [Wiener, 1949] in Section 2.5.1. To this end, we also make a brief excursion to the *Fourier domain* (see e.g. [Gasquet et al., 1998; Bracewell, 1999]). After that, we describe the idea of the statistics-based Richardson-Lucy (RL) deconvolution method [Richardson, 1972; Lucy, 1974] in Section 2.5.2. How variational methods can be used for image deconvolution is the topic of Section 2.5.3.

Chapter 3 addresses our novel cell reconstruction approach. We start by giving some insights about confocal laser scanning microscopy (CLSM) in Section 3.1.1 and stimulated emission depletion microscopy (STED) in Section 3.1.2. Besides explaining the functional principles, we discuss their physical limitations as well as resulting typical perturbations. Since blur constitutes one of the main problems, in Section 3.1.3 we describe how the PSF of an optical system can be estimated. In Section 3.2.1 we design a suitable forward operator simulating such 3-D low light intensity techniques and show the way to a statistically justified data term in Section 3.2.2. The mathematical concepts behind the *multiplicative Euler-Lagrange* formalism are shown in Section 3.2.3. They also allow us to give the relation between RL deconvolution and Csiszár's information divergence. Section 3.2.4 revisits

the idea of Welk [2010] supplementing RL deconvolution to a robust and regularised approach. Its extension to a joint inpainting and deconvolution approach presented by Elhayek et al. [2011] is discussed in Section 3.2.5. After that, we present our fibre enhancement approach in Section 3.3 and propose a novel, fast, and stabilised iteration scheme in the following Section 3.4. This chapter is completed by discretising our models and performing evaluations on real-world microscopy data sets.

In Chapter 4, we turn to the field of computer vision, and present a novel variational depth-from-defocus approach. For this purpose, we discuss some appropriate image formation models in Section 4.1 and show the way from the thin lens camera model to our novel normalised forward operator as its suitable approximation. After that, we address our variational inversion strategies in Section 4.2. There, we model a suitable functional and derive the associated variational derivatives w.r.t. both variants, additive and multiplicative Euler-Lagrange formalism. In Section 4.2.3 we supplement our model with the multi-channel case and propose a robust variant in the following Section 4.2.4. To handle noisy focal-stacks, Section 4.3 presents a joint denoising and depth-from-defocus approach. We discretise the developed concepts in Section 4.4 and give experimental comparisons using real and synthetic data in Section 4.5.

Chapter 2

Basic Concepts

In this chapter, we want to revisit some established image processing techniques that counteract typical degradations such as noise and blur as well as partial loss of information (cf. Section 1.1). Our intention is to familiarise the reader with these reliable concepts as they form the foundations for each of our proposed reconstruction frameworks.

Denoising while preserving important signal features such as image edges is one of the main challenges in signal processing. To this end, we first revisit diffusion filtering. Inspired by the physical process of diffusion [Fourier, 1822; Graham, 1829; Fick, 1855; Carslaw and Jaeger, 1959; Cussler, 1997], we illustrate its adaptation to an edge preserving smoothing filter [Fritsch, 1992; Perona and Malik, 1987; Weickert, 1997]. Furthermore, in order to incorporate even directional information that will be necessary in view of our cell reconstruction approach of Chapter 3, we explain the idea behind anisotropic diffusion filtering [Weickert, 1996, 1998]. Apart from the continuous modelling of those techniques, we also consider their discrete formulations. On the one hand, this represents a necessary step before the application of these ideas to digital images. On the other hand, a discrete formulation of the problem allows one to consult numerical methods [Morton and Mayers, 2005; Evans et al., 1999; Mitchell and Griffiths, 1980] which in turn are necessary to solve even more challenging PDEs occurring in later chapters.

To obtain more insights into variational methods, we proceed with the basic ideas behind variational image restoration [Bertero et al., 1988; Demoment, 1989; Rudin et al., 1992; Schnörr, 1994; Charbonnier et al., 1994] and show how to formulate the denoising and restoration tasks within a suitable energy functional. The problem then becomes the search for a suitable minimiser to which end we derive the corresponding Euler-Lagrange (EL) equation. In this chapter, we only consider classical additive EL formalism [Gelfand and Fomin, 2000] because our focus here is to introduce the basic

concept.

Apart from image denoising, the estimation of missing information is a second elementary issue. This not only includes the reconstruction of lost information due to, e.g. transmission gaps or defects of the sensor, in the discrete setting, this also comprises the task of interpolation, i.e. increasing the resolution. Due to physical limitations, in three-dimensional optical microscopy, the axial resolution (z-direction) is always below the lateral one (x and y direction). With image interpolation and inpainting strategy, one is able to remove this discrepancy [Lehmann et al., 1999]. To this end, we first consider the strategy of PDE-based inpainting [Masnou and Morel, 1998; Bertalmío et al., 2000; Chan and Shen, 2001; Galić et al., 2005; Roussos and Maragos, 2007]. Next, in order to also handle noisy data, we revisit joint inpainting and image restoration approaches. Such approaches are proposed by Chan and Shen [2002] or by Weickert and Welk [2006] in the context of tensor-valued images.

We complete this chapter by revisiting several sophisticated deconvolution techniques. Therefore, we also briefly explain the *convolution theorem* where we make a brief excursion to the *Fourier spectrum* [Gasquet et al., 1998; Bracewell, 1999]. On the one hand, this will be necessary to understand *inverse filtering* and *Wiener filtering* [Wiener, 1949]. On the other hand, following the convolution theorem ameliorates the runtime as we are dealing with large spatially invariant blurring kernels in our frameworks.
After that, we illustrate the idea behind the iterative deconvolution technique of Richardson [1972] and Lucy [1974]. The last section of this chapter is devoted to variational deconvolution [Osher and Rudin, 1994; Marquina and Osher, 1999] and a direct comparison of the proposed deconvolution strategies.

2.1 Diffusion

Diffusion is a natural, mass-preserving physical process that equilibrates spatial changes in concentrations such as particles in a fluid or in a gas. To formulate this process mathematically, let us consider an n-dimensional closed system $\Omega_n \subset \mathbb{R}^n$ where $u : \Omega_n \times \mathbb{R}_{0,+} \to \mathbb{R}_{0,+}$ describes the concentration at each location $\boldsymbol{x} := (x_1, \ldots, x_n)^\top \in \Omega_n$ at evolution time $t \in \mathbb{R}_{0,+}$. If the concentration is not constant in space, a flux $\boldsymbol{j} \in \mathbb{R}^n$ that is proportional to the negative concentration gradient to equilibrate spatial variations occurs (*Fick's law*) [Fick, 1855]:

$$\boldsymbol{j} = -\boldsymbol{D} \cdot \boldsymbol{\nabla} u \,. \tag{2.1}$$

Table 2.1: Mathematical notations and definitions for 2-D and 3-D image data sets.

Symbol	2-D	3-D
Ω	$\Omega_2 \subset \mathbb{R}^2$	$\Omega_3 \subset \mathbb{R}^3$
u	$u : \Omega_2 \times \mathbb{R}_{0,+} \to \mathbb{R}_{0,+}$	$u : \Omega_3 \times \mathbb{R}_{0,+} \to \mathbb{R}_{0,+}$
$\boldsymbol{j}$	$(j_1, j_2)^\top$	$(j_1, j_2, j_3)^\top$
$\boldsymbol{D}$	$\boldsymbol{D} \in \mathbb{R}^{2\times 2}$	$\boldsymbol{D} \in \mathbb{R}^{3\times 3}$
$\boldsymbol{\nabla}$	$\boldsymbol{\nabla}_2 \cdot := (\partial_{x_1}\cdot, \partial_{x_2}\cdot)^\top$	$\boldsymbol{\nabla}_3 \cdot := (\partial_{x_1}\cdot, \partial_{x_2}\cdot, \partial_{x_3}\cdot)^\top$
div	$\mathrm{div}(\boldsymbol{j}) := \partial_{x_1} j_1 + \partial_{x_2} j_2$	$\mathrm{div}(\boldsymbol{j}) := \partial_{x_1} j_1 + \partial_{x_2} j_2 + \partial_{x_3} j_3$
Δ	$\Delta_2 \cdot := \partial_{x_1x_1} \cdot + \partial_{x_2x_2} \cdot$	$\Delta_3 \cdot := \partial_{x_1x_1} \cdot + \partial_{x_2x_2} \cdot + \partial_{x_3x_3} \cdot$

Here $\boldsymbol{\nabla} := (\partial_{x_1}, \ldots, \partial_{x_n})^\top$ denotes the n-dimensional gradient operator and ∂_{x_i} the partial derivative w.r.t. the i-th dimension. The *diffusion tensor* $\boldsymbol{D}$ can be seen as a descriptor of the mobility of particles and depends on a lot of different factors such as the temperature or the medium within the diffusion process occurs (see e.g. [Cussler, 1997]). In the anisotropic case, where the mobility of particles may vary for different directions, $\boldsymbol{D}$ is a positive definite symmetric matrix of size $n \times n$. In the isotropic case where no directional dependency exists $\boldsymbol{j}$ and $\boldsymbol{\nabla} u$ are parallel and $\boldsymbol{D}$ turns into scalar-valued diffusivity, usually denoted by g.

Due to the fact that mass is preserved, one has to follow the *continuity equation*. The temporal change in concentration $\partial_t u$ thus can be formulated by the *partial differential equation* (PDE)

$$\partial_t u = \mathrm{div}(\boldsymbol{D} \cdot \boldsymbol{\nabla} u) \;, \tag{2.2}$$

which is the so-called diffusion or *heat equation* [Fick, 1855; Cussler, 1997; Weickert, 1998]. To utilise the physical process of diffusion for image processing, one regards images as smooth (differentiable) functions $f : \Omega_n \to \mathbb{R}_+$ given on an image domain Ω_n. The diffusion process (Equation (2.2)) can then be transferred to image processing by interpreting grey values $f(\boldsymbol{x})$ as concentration values with the initial setting $u(\boldsymbol{x}, 0) := f(\boldsymbol{x})$. Since we are mainly dealing with 2-D and 3-D data sets in this thesis, Table 2.1 provides an overview of the mathematical notations.

Nowadays, diffusion filtering is a very popular method in image processing and computer vision. It can be found in a broad spectrum of applications such as the well-known edge preserving isotropic denoising approach of

Figure 2.1: Image denoising with homogenous diffusion. (2.3). *From left to right:* **(a)** Input image (Gaussian noise, $\sigma_{\text{noise}} = 45$). **(b)** After a diffusion time of $t = 4.5$. **(c)** Diffusion time $t = 18$. **(d)** Diffusion time $t = 40.5$.

Perona and Malik [1987], in image inpainting as well as image compression methods [Galić et al., 2005; Mainberger and Weickert, 2009; Galić et al., 2008; Hoffmann et al., 2013]. Additionally, it provides the foundation to understand and classify most of the advanced variational regularisation techniques [Scherzer and Weickert, 2000].

The diffusion process, and thus its smoothing behaviour, strongly depends on the diffusion tensor. It can describe an isotropic or anisotropic process, it can be variant or invariant w.r.t. the location and/or the evolution time. By adapting the diffusion tensor to the local image structure, one can create edge preserving and edge enhancing effects. In the forthcoming sections, we are going to briefly discuss different diffusion techniques and we are going to analyse their smoothing properties.

2.1.1 Isotropic Diffusion

Homogeneous Diffusion. Strictly speaking, *homogeneous diffusion* comprises diffusion processes having a spatially invariant diffusion tensor such that its entries do not change between different locations. Most often, however, one interprets this designation as the simplest case wherein the diffusion tensor reduces to the identity matrix $\boldsymbol{D} := \boldsymbol{I}$ (equal to $g = 1$ due to its isotropy). If we talk about homogeneous diffusion, we follow the last interpretation and refer to

$$\partial_t u = \Delta u \;, \tag{2.3}$$

where $\Delta u := \operatorname{div}(\boldsymbol{I} \cdot \boldsymbol{\nabla} u) = \sum_{i=1}^{n} \partial_{x_i} \partial_{x_i} u$ denotes the Laplacian of u [Iijima, 1959]. To obtain the solution of the PDE above, one possibility is to exploit the equivalence between homogeneous diffusion and Gaussian convolution [Cannon, 1984; Gonzalez-Velasco, 1996]. Accordingly, if f denotes the initial

image, the intensity distribution at evolving time $t > 0$ can be calculated by

$$u(\boldsymbol{x}, t) = (\mathcal{K}_{\sqrt{2t}} * f)(\boldsymbol{x}) \;, \tag{2.4}$$

where $\mathcal{K}_\sigma$ denotes a Gaussian (for the 2-D case cf. Equation (1.6)) with zero mean and standard deviation σ. The only parameter in this diffusion model is the evolution time t and the standard deviation σ, respectively. Figure 2.1 demonstrates the behaviour of homogeneous diffusion for different evolution times t applied to a noisy input image. The noise is Gaussian distributed with $\sigma_{\text{noise}} = 45$. While the equivalence with Gaussian convolution holds for homogeneous diffusion, we need more advanced strategies to even solve more challenging PDEs. Moreover, until now, our formulation has been done in a continuous setting, but digital images only provide sampled discrete data. For these reasons, let us investigate in later sections how such PDEs can be solved numerically.

Inhomogeneous Diffusion. Due to a spatially constant diffusivity, homogeneous diffusion can not take into account any local structural information. This results in a uniform blur over the whole image domain, and thus besides a denoising effect, in blurred and dislocated edges. For the human visual system, however, edges are one of the most significant image features. They are very important for our perception and make it easy for us to separate and identify different objects. For this reason, we are interested in filters providing good denoising capabilities but at the same time preserving edges: The idea behind *inhomogeneous diffusion* is to locally control the strength of the diffusion process by adapting its diffusivity g to the underlying image structure. The gradient magnitude of the image usually serves a suitable indicator of the local structure and can be taken either from the initial image f, so-called *linear inhomogeneous diffusion*, as proposed by Fritsch [1992], or the evolving image u as proposed by Perona and Malik [1987] (*nonlinear inhomogeneous diffusion*). In this thesis, we only regard the nonlinear case which leads to a more sophisticated approach offering better edge localisations. Making the diffusivity dependent on the gradient magnitude of the evolving image, the diffusion process can be formulated as

$$\partial_t u = \operatorname{div}(g(|\boldsymbol{\nabla} u|^2) \cdot \boldsymbol{\nabla} u) \;, \tag{2.5}$$

where the differentiable diffusivity function $g(s^2) > 0$ decreases with increasing s and should be close to 0 for $s \to \infty$ in order to damp the diffusion process near image discontinuities. While for Perona-Malik [Perona and Malik, 1987], Charbonnier [Charbonnier et al., 1997], and Weickert diffusivities

[Weickert, 1998] it holds additionally that $g(0) = 1$, total variation (TV) diffusivity [Rudin et al., 1992] becomes unbounded, or in the case of regularised TV [Feng and Prohl, 2002] bounded by $1/\epsilon$ with some small stabilising $\epsilon > 0$. Besides that, the Perona-Malik and Weickert diffusivities are designed in a way that they do not only offer an edge preserving, but also an edge enhancing effect, which is controllable by the so-called contrast parameter λ.

To handle strong noise as well as to avoid staircasing artefacts, one often takes the gradient magnitude of a presmoothed version u_σ of the evolving image u. This way, Equation (2.5) turns to

$$\partial_t u = \operatorname{div}(g(|\boldsymbol{\nabla} u_\sigma|^2) \cdot \boldsymbol{\nabla} u) \;, \tag{2.6}$$

where $u_\sigma := \mathcal{K}_\sigma * u$. One can show that such a presmoothing step turns the ill-posed problem of Perona-Malik diffusion to a well-posed one, referred as *regularised* Perona-Malik diffusion.

Discretisation. To apply isotropic diffusion to digital 2-D images sampled on a rectangular regular grid of $N := N_{x_1} \times N_{x_2}$ pixels, let $u_{i,j}^k$ approximate u at pixel (i, j) and evolving time $t = k \cdot \tau$. Here, τ denotes the time stepsize. If we further assume a cell size of $\boldsymbol{h}_2 := (h_{x_1}, h_{x_2})^\top$, and follow a finite differences scheme [Mitchell and Griffiths, 1980; Morton and Mayers, 2005], we can approximate Equation (2.6) in each pixel by

$$\begin{aligned}\frac{u_{i,j}^{k+1} - u_{i,j}^k}{\tau} = &\left(\frac{g_{i+1,j} + g_{i,j}}{2} \cdot \frac{u_{i+1,j}^k + u_{i,j}^k}{h_{x_1}^2} - \frac{g_{i,j} + g_{i-1,j}}{2} \cdot \frac{u_{i,j}^k + u_{i-1,j}^k}{h_{x_1}^2}\right) \\ &+ \left(\frac{g_{i,j+1} + g_{i,j}}{2} \cdot \frac{u_{i,j+1}^k + u_{i,j}^k}{h_{x_2}^2} - \frac{g_{i,j} + g_{i,j-1}}{2} \cdot \frac{u_{i,j}^k + u_{i,j-1}^k}{h_{x_2}^2}\right) \;,\end{aligned} \tag{2.7}$$

where $g_{i,j}$ approximates g at pixel (i, j) in a so-called *lagged diffusivity* manner (Kačanov-method) [Fučik et al., 1973; Chan and Mulet, 1999; Vogel, 2002], i.e. g is applied to u at the old iteration step k:

$$g_{i,j} := g\left(\left|[\boldsymbol{\nabla} u_\sigma]_{i,j}^k\right|^2\right) \;. \tag{2.8}$$

The notation $[\boldsymbol{\nabla} u_\sigma]_{i,j}^k = ([\partial_{x_1} u_\sigma]_{i,j}^k \;,\; [\partial_{x_2} u_\sigma]_{i,j}^k)^\top$ describes the approximation of the 2-D gradient operator. The first order derivatives can be obtained

with the help of central differences:

$$[\partial_{x_1} u]_{i,j}^k := \frac{u_{i+1,j}^k - u_{i-1,j}^k}{2h_{x_1}} ,$$
$$[\partial_{x_2} u]_{i,j}^k := \frac{u_{i,j+1}^k - u_{i,j-1}^k}{2h_{x_2}} . \tag{2.9}$$

Since there is no information available outside the image domain Ω, we follow *homogeneous Neumann boundary conditions* by mirroring the image at its boundaries $\partial\Omega$.

To formulate Equation (2.7) in a more compact way, we now change from a double-index notation to a single-index one (e.g. row-major ordering) and rearrange discrete 2-D signals $u_{i,j}$ in vectors $\boldsymbol{u} \in \mathbb{R}^N$. Further, let us define the symmetric matrix $\boldsymbol{A}(\boldsymbol{u}^k) \in \mathbb{R}^{N \times N}$ by its entries

$$a_{n,m} := \begin{cases} \frac{g_n + g_m}{2h_{x_\ell}^2} & (m \in \mathcal{N}_{x_\ell}(n)) , \\ -\sum_{\ell=1}^{2} \sum_{m \in \mathcal{N}_{x_\ell}(n)} \frac{g_n + g_m}{2h_{x_\ell}^2} & (m = n) , \\ 0 & (else) , \end{cases} \tag{2.10}$$

where $\mathcal{N}_{x_\ell}(n)$ denotes the neighbouring pixels in the ℓ-th direction of pixel n. Then, Equation (2.7) can be expressed with the help of a matrix-vector multiplication

$$\frac{\boldsymbol{u}^{k+1} - \boldsymbol{u}^k}{\tau} = \boldsymbol{A}(\boldsymbol{u}^k)\boldsymbol{u}^k . \tag{2.11}$$

Solving for $\boldsymbol{u}^{k+1}$ results in an *explicit* iteration scheme

$$\boldsymbol{u}^{k+1} = (\boldsymbol{I} + \tau \cdot \boldsymbol{A}(\boldsymbol{u}^k))\ \boldsymbol{u}^k , \tag{2.12}$$

with $\boldsymbol{u}^0 = \boldsymbol{f}$. Setting the diffusivity function $g_{i,j} := 1$, of course, results in a discretisation scheme for homogeneous diffusion. While this explicit approach can be implemented in a straightforward way, it suffers from the stability condition [Weickert, 1998]

$$\tau < \frac{1}{\max_n |a_{n,n}|} , \tag{2.13}$$

and thus requires a relatively large number of iterations. Such a restriction can be circumvented by going a slightly different way. It is based on considering the factor $\boldsymbol{u}$ in the right-hand side of Equation (2.11) not at the old

Figure 2.2: Image denoising by nonlinear isotropic diffusion using Charbonnier diffusivity [Charbonnier et al., 1997] in Equation (2.17). *From left to right:* **(a)** Input image. **(b)** Output with $\sigma = 0.4$, $t = 18$ and $\lambda = 1$. **(c)** Ditto with $\lambda = 2$. **(d)** Ditto with $t = 40.5$ and $\lambda = 2$.

time step k but at the new one $k+1$:

$$\frac{\boldsymbol{u}^{k+1} - \boldsymbol{u}^k}{\tau} = \boldsymbol{A}(\boldsymbol{u}^k)\boldsymbol{u}^{k+1} \,. \tag{2.14}$$

Eventually, solving for $\boldsymbol{u}^{k+1}$ leads to

$$\boldsymbol{u}^{k+1} = (\boldsymbol{I} - \tau \cdot \boldsymbol{A}(\boldsymbol{u}^k))^{-1}\, \boldsymbol{u}^k \,. \tag{2.15}$$

This so-called *semi-implicit* iteration scheme no longer suffers from the stability condition (2.13) [Weickert, 1998]. However, it requires to solve a linear system of equations in each iteration.

In later sections of this thesis, we demonstrate how we can take benefit from such semi-implicit ideas as they may help us to also solve very unstable PDEs.

Multi-Channel Images. Now, let us explain how to extend isotropic diffusion filtering for its application to multi-channel images (e.g. colour images). So, let $\boldsymbol{f} = (f_c)_{c \in \mathcal{C}}$ be an image with channel index set $\mathcal{C}$ (e.g. RGB). For homogeneous diffusion there is no need to couple different channels, since the diffusivity is constant and no information exchange between different channels is necessary. Hence, if we denote by $\boldsymbol{u} = (u_c)_{c \in \mathcal{C}}$ the evolving image, the homogeneous diffusion process can be performed channel-wise:

$$\begin{aligned} \partial_t u_c &= \Delta u_c \,, \quad \forall c \in \mathcal{C} \,, \\ \boldsymbol{u}(\boldsymbol{x}, 0) &= \boldsymbol{f}(\boldsymbol{x}) \,. \end{aligned} \tag{2.16}$$

In contrast, inhomogeneous diffusion needs an indicator that combines the structural information of all channels. This can be provided by taking

the sum over the squared gradient magnitude of each channel as proposed by Gerig et al. [1992]. Using it as an argument for the diffusivity function g, the diffusion equation reads

$$\partial_t u_c = \operatorname{div}\Big(g\big(\sum_{c\in\mathcal{C}} |\boldsymbol{\nabla} u_{\sigma,c}|^2\big) \cdot \boldsymbol{\nabla} u_c \Big) \quad \forall c \in \mathcal{C} \; . \tag{2.17}$$

Figure 2.2 shows the behaviour of nonlinear isotropic diffusion with Charbonnier diffusivity [Charbonnier et al., 1997].

2.1.2 Anisotropic Diffusion

In the last section, we have discussed the idea of isotropic inhomogeneous diffusion filtering. By adapting the scalar-valued diffusivity to the underlying (evolving) image structure, the diffusivity of the diffusion process is locally regulated such that blurring proceeds within flat regions but is attenuated at image discontinuities. Although it provides in this way an edge preserving denoising filter, it suffers from the limitation to not incorporate any directional information. Reducing the diffusivity in all directions equally at discontinuities may preserve image edges, but it may also preserve noise at those locations. Further, particularly regarding our cell reconstruction problem of Chapter 3, we need a regularisation technique that can be steered along coherent elongated image structures. For this purpose, let us now briefly discuss the nonlinear anisotropic diffusion scheme of Weickert [1998] for the two-dimensional case. Instead of a scalar-valued diffusivity, here, a symmetric, positive definite diffusion tensor $\boldsymbol{D} \in \mathbb{R}^{2\times 2}$ modulates the flux. While it is sufficient to use an edge detector such as the gradient magnitude in the isotropic case, for the anisotropic case, directional information about the underlying image structure is also required. This can be achieved by analysing the *structure tensor* $\boldsymbol{J} \in \mathbb{R}^{2\times 2}$ of Förstner and Gülch [1987], which is a symmetric, positive semi-definite matrix. For a multi-channel image $\boldsymbol{u} = (u_c)_{c\in\mathcal{C}}$, we follow Di Zenzo [1986] and Weickert [1999a] and consider the *joint structure tensor*

$$\boldsymbol{J}_\rho(\boldsymbol{u}_\sigma) := \sum_{c\in\mathcal{C}} \mathcal{K}_\rho * \left(\boldsymbol{\nabla} u_{\sigma,c} \cdot \boldsymbol{\nabla} u_{\sigma,c}^\top \right) \; , \tag{2.18}$$

where the two involved convolution operations act differently: While the inner one $\boldsymbol{u}_\sigma := \mathcal{K}_\sigma * \boldsymbol{u}$ offers robustness w.r.t. noise, the outer convolution (that is applied to all entries of the tensor) averages the directional information over some neighbourhood whose size is described by the standard

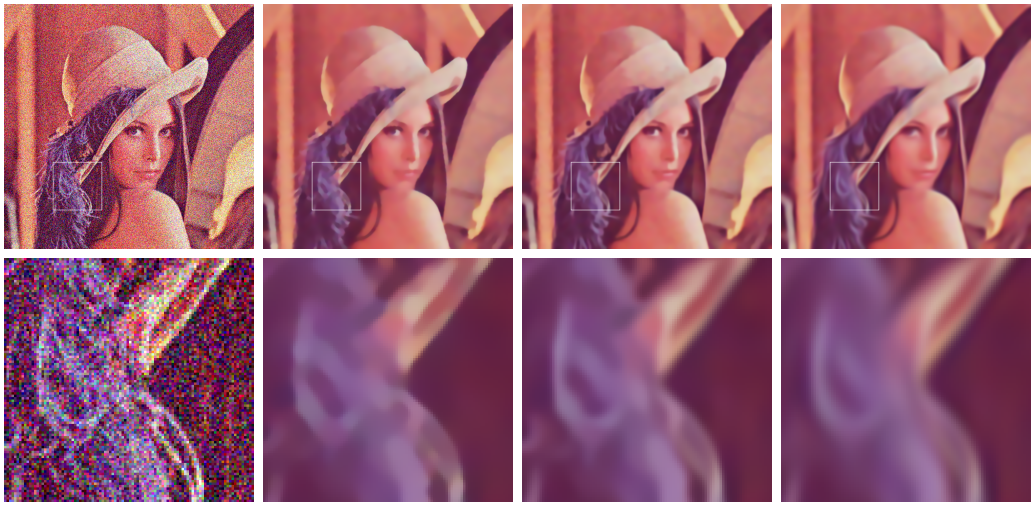

Figure 2.3: Comparison of isotropic (2.17) versus anisotropic diffusion (2.21). The white rectangle in the top images indicates the origin of the zoom-in shown in the bottom image. *From left to right:* **(a) Column 1:** Input image **(b) Column 2:** Isotropic with $\sigma = 0.4$, $t = 40.5$ and $\lambda = 2$. **(c) Column 3:** Anisotropc with $\sigma = 0.4$, $\lambda = 2$, $\rho = 1.0$, $t = 24$. **(d) Column 4:** $\rho = 4.0$.

deviation ρ. Eventually, an eigenvalue decomposition of the structure tensor

$$\boldsymbol{J}_\rho(u_\sigma) = \begin{pmatrix} \boldsymbol{v}_1 & \boldsymbol{v}_2 \end{pmatrix} \begin{pmatrix} \mu_1 & 0 \\ 0 & \mu_2 \end{pmatrix} \begin{pmatrix} \boldsymbol{v}_1^\top \\ \boldsymbol{v}_2^\top \end{pmatrix} \tag{2.19}$$

reveals the averaged (w.r.t. the integration scale ρ) direction v_1 of the highest contrast and v_2 ($v_1 \perp v_2$) of the lowest contrast. Accordingly, the respective eigenvalue μ_1 and μ_2 describes the contrast in the corresponding direction. Adopting the eigenvectors but penalising each of the eigenvalues in a separate way

$$\boldsymbol{D}(\boldsymbol{J}_\rho(u_\sigma)) := \begin{pmatrix} \boldsymbol{v}_1 & \boldsymbol{v}_2 \end{pmatrix} \begin{pmatrix} g_1(\mu_1) & 0 \\ 0 & g_2(\mu_2) \end{pmatrix} \begin{pmatrix} \boldsymbol{v}_1^\top \\ \boldsymbol{v}_2^\top \end{pmatrix} , \tag{2.20}$$

the diffusion tensor can be seen as a function of the structure tensor. Finally, we obtain an anisotropic diffusion equation by replacing the diffusivity function in Equation (2.6) by the diffusion tensor above:

$$\partial_t u = \operatorname{div}(\boldsymbol{D}(\boldsymbol{J}_\rho(u_\sigma)) \cdot \boldsymbol{\nabla} u) . \tag{2.21}$$

Such an anisotropic diffusion process now allows us to perform a different smoothing w.r.t. the directions $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$. For example, choosing $g_1(s^2) :=$

Table 2.2: Discretisation scheme to approximate Equation (2.23).

$\partial_{x_1}(a\ \partial_{x_1}u) \approx [\partial_{x_1}(a\ \partial_{x_1}u)]_{i,j} = \left(\frac{a_{i+1,j}+a_{i,j}}{2} \frac{u_{i+1,j}-u_{i,j}}{h_{x_1}^2} - \frac{a_{i,j}+a_{i-1,j}}{2} \frac{u_{i,j}-u_{i-1,j}}{h_{x_1}^2} \right)$
$\partial_{x_1}(b\ \partial_{x_2}u) \approx [\partial_{x_1}(b\ \partial_{x_2}u)]_{i,j} = \left(b_{i+1,j} \frac{u_{i+1,j+1}-u_{i+1,j-1}}{4h_{x_1}h_{x_2}} - b_{i-1,j} \frac{u_{i-1,j+1}-u_{i-1,j-1}}{4h_{x_1}h_{x_2}} \right)$
$\partial_{x_2}(b\ \partial_{x_1}u) \approx [\partial_{x_2}(b\ \partial_{x_1}u)]_{i,j} = \left(b_{i,j+1} \frac{u_{i+1,j+1}-u_{i-1,j+1}}{4h_{x_1}h_{x_2}} - b_{i,j-1} \frac{u_{i+1,j-1}-u_{i-1,j-1}}{4h_{x_1}h_{x_2}} \right)$
$\partial_{x_2}(c\ \partial_{x_2}u) \approx [\partial_{x_2}(c\ \partial_{x_2}u)]_{i,j} = \left(\frac{c_{i+1,j}+c_{i,j}}{2} \frac{u_{i+1,j}-u_{i,j}}{h_{x_2}^2} - \frac{c_{i,j}+c_{i,j-1}}{2} \frac{u_{i,j}-u_{i,j-1}}{h_{x_2}^2} \right)$

$1/\sqrt{1+s^2/\lambda^2}$ and $g_2(s^2) := 1$ results in edge preserving Charbonnier penalisation [Charbonnier et al., 1997] along the direction of the highest contrast, i.e. across edges, combined with homogeneous diffusion along the direction of the lowest contrast, i.e. along coherent structures, respectively. In Figure 2.3, we compare the smoothing behaviour of isotropic and anisotropic diffusion.

2.1.3 Discretisation

In the meantime, several approaches for a suitable discretisation scheme of the anisotropic diffusion model exist in the literature [Weickert, 1998, 1999b; Cottet and El Ayyadi, 1998; Weickert et al., 2013]. In that respect, one of the main challenges is the preservation of scale-space properties of the continuous anisotropic diffusion model even in the discrete setting. However, in this basic chapter we want to show only the basic principles of anisotropic diffusion. Hence, let us consider the standard discretisation of Weickert [1999b]. Thus, assuming a diffusion tensor $\boldsymbol{D}$ of the form

$$\boldsymbol{D} = \begin{pmatrix} a & b \\ b & c \end{pmatrix} , \tag{2.22}$$

the right-hand side of Equation (2.21) turns to

$$\operatorname{div} \begin{pmatrix} a\ \partial_{x_1}u + b\ \partial_{x_2}u \\ b\ \partial_{x_1}u + c\ \partial_{x_2}u \end{pmatrix} = \partial_{x_1}(a\ \partial_{x_1}u) + \partial_{x_1}(b\ \partial_{x_2}u) + \partial_{x_2}(b\ \partial_{x_1}u) + \partial_{x_2}(c\ \partial_{x_2}u) , \tag{2.23}$$

where each of the summands can be approximated by means of central differences (cf. Table 2.2). Even though the dependency of the diffusion tensor on u_σ is not denoted in the equation above, one should keep in mind that the entries of $\boldsymbol{D}$ depend on the evolving image. Due to the mixed derivatives

and the involved diagonal neighbours, we refrain from an explicit formulation of the entries of the matrix $\boldsymbol{A}$ and make use of a stencil notation (grey annotations indicates the pixel indices).

	$-\frac{b_{i-1,j}+b_{i,j+1}}{4h_{x_1}h_{x_2}}$	$\frac{c_{i,j+1}+c_{i,j}}{2h_{x_2}^2}$	$\frac{b_{i+1,j}+b_{i,j+1}}{4h_{x_1}h_{x_2}}$
	$\frac{a_{i-1,j}+a_{i,j}}{2h_{x_1}^2}$	$-\left(\frac{a_{i-1,j}+2a_{i,j}+a_{i+1,j}}{2h_{x_1}^2} + \frac{c_{i,j-1}+2c_{i,j}+c_{i,j+1}}{2h_{x_2}^2}\right)$	$\frac{a_{i+1,j}+a_{i,j}}{2h_{x_1}^2}$
$j \uparrow$	$\frac{b_{i-1,j}+b_{i,j-1}}{4h_{x_1}h_{x_2}}$	$\frac{c_{i,j-1}+c_{i,j}}{2h_{x_2}^2}$	$-\frac{b_{i+1,j}+b_{i,j-1}}{4h_{x_1}h_{x_2}}$

$i \rightarrow$

Figure 2.4: Variational image restoration (cf. Equation (2.24)). *From left to right:* **(a)** Input image. **(b)** With Whittaker-Tikhonov penaliser [Whittaker, 1923; Tikhonov, 1963] ($\alpha = 30$). **(c)** With Charbonnier penaliser ($\lambda = 2$, $\alpha = 20$) [Charbonnier et al., 1997]. **(d)** Ditto ($\alpha = 30$).

2.2 Image Restoration

In Section 1.4.2 we have already briefly introduced the principles of variational methods. As already mentioned, they pose the respective reconstruction problem as the minimisation of a suitable energy functional. While we have considered a general expression so far, let us now become more concrete and illustrate how such a technique can be used for image denoising. Furthermore, let us explain the concept of the classical additive Euler-Lagrange formalism in order to find a suitable minimiser of the considered functional.

2.2.1 Variational Model

According to the noise model of Equation (1.1), let us assume that $f : \Omega_n \to \mathbb{R}$ describes a noisy version of the original undisturbed signal g. Then, we are interested in a denoised signal u that approximates g. This approximation should at least fulfil two requirements. On the one hand, it should, of course, resemble the input data as exactly as possible. On the other hand, as noise usually acts as a high-frequency perturbation, the approximation should be (piecewise) smooth. Accordingly, an energy functional requiring these two aspects to its minimiser u can be modelled as

$$E(u) := \underbrace{\int_{\Omega_n} (u - f)^2 \, \mathrm{d}\boldsymbol{x}}_{\text{data term}} + \alpha \cdot \underbrace{\int_{\Omega_n} \Psi(|\boldsymbol{\nabla} u|^2) \, \mathrm{d}\boldsymbol{x}}_{\text{smoothness term}} , \tag{2.24}$$

where α weights between accuracy and smoothness, $\boldsymbol{\nabla} := (\partial_{x_1}, \ldots, \partial_{x_n})^\top$ denotes the gradient operator, and $\Psi : \mathbb{R}_{0,+} \to \mathbb{R}_+$ is a positive increasing function. The data term enforces the required similarity between u and f (this implies a forward operator that is just the identity operator). It is known that

a quadratic penaliser $r_{1,f}(u) := (u-f)^2$ (acting in a least square sense) of deviations is especially appropriate w.r.t. additive white Gaussian noise [Steidl and Teuber, 2010; Welk, 2015]. In the smoothness term, depending on the respective choice of Ψ, different smoothing behaviours can be achieved. Typical choices are Whittaker-Tikhonov penaliser $\Psi(s^2) := s^2$ [Whittaker, 1923; Tikhonov, 1963], Charbonnier penaliser $\Psi(s^2) := 2\lambda^2\sqrt{1 + s^2/\lambda^2}$ [Charbonnier et al., 1997] or total variation (TV) [Rudin et al., 1992] regularisation $\Psi(s^2) := |s|$.

To find a suitable minimiser of such a functional, let us explain the classical additive Euler-Lagrange formalism in the next section.

2.2.2 Classical Euler-Lagrange Formalism

After we have modelled a variational energy functional, we need a way to find a suitable minimiser. As we can see, for the n-dimensional case, Equation (2.24) is of the form

$$E(u) = \int_{\Omega_n} F(x_1, \ldots, x_n, u, u_{x_1}, \ldots, u_{x_n}) \, \mathrm{d}\boldsymbol{x} \; , \tag{2.25}$$

where u_{x_i} describes the first order partial derivative of u w.r.t. the i-th direction. From variational calculus, one knows that a minimiser necessarily has to fulfil the Euler-Lagrange equation

$$\frac{\delta E}{\delta u} = 0 \; , \tag{2.26}$$

as well as the corresponding boundary conditions [Gelfand and Fomin, 2000]. To derive the variational gradient (or functional derivative) $\frac{\delta E}{\delta u}$, one commonly follows the additive Euler-Lagrange formalism. It requires for each differentiable perturbation function v

$$\left\langle \frac{\delta E}{\delta u}, v \right\rangle \overset{!}{=} \frac{\partial}{\partial \epsilon} E(u + \epsilon \, v)\Big|_{\epsilon=0} \; , \tag{2.27}$$

where $\langle a(\boldsymbol{x}), b(\boldsymbol{x}) \rangle = \int a(\boldsymbol{x}) \cdot b(\boldsymbol{x}) \, \mathrm{d}\boldsymbol{x}$ denotes the standard inner product. The right-hand side of Equation (2.27) can be interpreted as a directional derivative in the direction of v. In the classical Euler-Lagrange formalism,

this perturbation acts additively and yields for the functional (2.25)

$$
\begin{aligned}
&\frac{\partial}{\partial\epsilon}E(u+\epsilon\, v)\Big|_{\epsilon=0} \\
&= \frac{\partial}{\partial\epsilon}\int_{\Omega_n} F(x_1,\ldots,x_n,u+\epsilon v,u_{x_1}+\epsilon v_{x_1},\ldots,u_{x_n}+\epsilon v_{x_n})\ \mathrm{d}\boldsymbol{x}\Big|_{\epsilon=0} \\
&\overset{*}{=} \int_{\Omega_n}\Big(F_u(x_1,\ldots,x_n,u+\epsilon v,u_{x_1}+\epsilon v_{x_1},\ldots,u_{x_n}+\epsilon v_{x_n})\cdot v \\
&+ \sum_{i=1}^{n} F_{u_{x_i}}(x_1,\ldots,x_n,u+\epsilon v,u_{x_1}+\epsilon v_{x_1},\ldots,u_{x_n}+\epsilon v_{x_n})\cdot v_{x_i}\Big)\mathrm{d}\boldsymbol{x}\Big|_{\epsilon=0} \\
&= \int_{\Omega_n}\Big(F_u(x_1,\ldots,x_n,u,u_{x_1},\ldots,u_{x_n})\cdot v \\
&+ \sum_{i=1}^{n} F_{u_{x_i}}(x_1,\ldots,x_n,u,u_{x_1},\ldots,u_{x_n})\cdot v_{x_i}\Big)\mathrm{d}\boldsymbol{x} \\
&= \int_{\Omega_n}\Big(F_u\cdot v+\sum_{i=1}^{n} F_{u_{x_i}}\cdot v_{x_i}\Big)\mathrm{d}\boldsymbol{x}\ . \qquad (2.28)
\end{aligned}
$$

The second row (*) is obtained according to the chain rule. For a better readability, we do not write the arguments from the last row on. If we further define the vector $\boldsymbol{F}_{\boldsymbol{\nabla}u} := (F_{u_{x_1}},\ldots,F_{u_{x_n}})^\top$, we can write the equation above as:

$$
\frac{\partial}{\partial\epsilon}E(u+\epsilon\, v)\Big|_{\epsilon=0} = \int_{\Omega_n}\Big(F_u\cdot v+\boldsymbol{F}_{\boldsymbol{\nabla}u}^\top\,\boldsymbol{\nabla}v\Big)\ \mathrm{d}\boldsymbol{x}\ . \qquad (2.29)
$$

By applying integration by parts for the higher dimensional case to the second summand, we can now reformulate Equation (2.27) to

$$
\begin{aligned}
\Big\langle\frac{\delta E}{\delta u},v\Big\rangle &\overset{!}{=} \int_{\Omega_n}F_u\cdot v\ \mathrm{d}\boldsymbol{x}-\int_{\Omega_n}\mathrm{div}\,(\boldsymbol{F}_{\boldsymbol{\nabla}u})\cdot v\ \mathrm{d}\boldsymbol{x} \\
&\quad+\int_{\partial\Omega_n}\Big(\boldsymbol{\eta}^\top\boldsymbol{F}_{\boldsymbol{\nabla}u}\Big)\cdot v\ \mathrm{d}\boldsymbol{x}\ ,\quad\forall v \\
&= \int_{\Omega_n}\left(F_u-\mathrm{div}\,(\boldsymbol{F}_{\boldsymbol{\nabla}u})\right)\cdot v\ \mathrm{d}\boldsymbol{x} \\
&\quad+\int_{\partial\Omega_n}\Big(\boldsymbol{\eta}^\top\boldsymbol{F}_{\boldsymbol{\nabla}u}\Big)\cdot v\ \mathrm{d}\boldsymbol{x}\ ,\quad\forall v \qquad (2.30)
\end{aligned}
$$

where $\partial\Omega_n$ denotes the image boundary with its unit normal vector $\boldsymbol{\eta}$. Since this requirement has to be fulfilled for all admissible perturbation functions v, it has to hold also for all those v that vanish at $\partial\Omega_n$, i.e. $v(\boldsymbol{x}) = 0$ for $\boldsymbol{x} \in \partial\Omega_n$. This case in turn leads us to the variational gradient

$$\left\langle \frac{\delta E}{\delta u}, v \right\rangle \stackrel{!}{=} \int_{\Omega_n} \left(F_u - \operatorname{div}\left(\boldsymbol{F}_{\boldsymbol{\nabla} u}\right)\right) \cdot v \, \mathrm{d}\boldsymbol{x}$$

$$\iff \qquad \frac{\delta E}{\delta u} = F_u - \operatorname{div}\left(\, \boldsymbol{F}_{\boldsymbol{\nabla} u} \,\right) \, . \tag{2.31}$$

Not imposing the restriction above to the set of perturbation functions i.e. also allowing v that do not vanish at $\partial\Omega_n$, and embedding the variational gradient of Equation (2.31) as the first argument of the scalar product (2.30), we obtain

$$\begin{aligned} \left\langle F_u - \operatorname{div}\left(\boldsymbol{F}_{\boldsymbol{\nabla} u}\right), v \right\rangle &\stackrel{!}{=} \int_{\Omega_n} \left(F_u - \operatorname{div}\left(\boldsymbol{F}_{\boldsymbol{\nabla} u}\right)\right) \cdot v \, \mathrm{d}\boldsymbol{x} \\ &\quad + \int_{\partial\Omega_n} \left(\boldsymbol{\eta}^\top \boldsymbol{F}_{\boldsymbol{\nabla} u}\right) \cdot v \, \mathrm{d}\boldsymbol{x} \, , \quad \forall v \\ \iff \qquad 0 &\stackrel{!}{=} \int_{\partial\Omega_n} \left(\boldsymbol{\eta}^\top \boldsymbol{F}_{\boldsymbol{\nabla} u}\right) \cdot v \, \mathrm{d}\boldsymbol{x} \, , \quad \forall v \, , \end{aligned} \tag{2.32}$$

which leads us to the *natural boundary conditions*

$$\boldsymbol{\eta}^\top \boldsymbol{F}_{\boldsymbol{\nabla} u} \stackrel{!}{=} 0 \, . \tag{2.33}$$

Let us now come back to our variational image restoration approach (cf. Equation (2.24)). Following Equation (2.31), eventually, the necessary Euler-Lagrange condition reads

$$\underbrace{(u - f)}_{\text{similarity}} - \alpha \cdot \underbrace{\operatorname{div}\left(\Psi'(|\boldsymbol{\nabla} u|^2) \cdot \boldsymbol{\nabla} u\right)}_{\text{smoothness}} \stackrel{!}{=} 0 \tag{2.34}$$

with boundary conditions (cf. Equation (2.33))

$$\boldsymbol{\eta}^\top \boldsymbol{\nabla} u \stackrel{!}{=} 0 \, , \tag{2.35}$$

which can be realised, e.g. by mirroring the image at its boundaries. As we can see, while the data term results in a *similarity* expression in the

Figure 2.5: Inpainting. *From left to right:* **(a)** Input image. **(b)** inpainting mask. **(c)** PDE-based inpainting (2.36) with Whittaker-Tikhonov penaliser [Tikhonov, 1963; Whittaker, 1923]. **(d)** Ground truth.

associated Euler-Lagrange equation, regularisation finds its analogue in a term that is related to isotropic diffusion as described in Equation (2.6). The relation between regularisation and diffusion is investigated in detail by Scherzer and Weickert [2000]. This connection not only allows us to modify the smoothing behaviour of variational methods in the same flexible way as in diffusion filtering, it also enables us to follow similar numerical strategies.

In Figure 2.4, we compare the results of variational image restoration with different regularisation strategies. As we can see, Whittaker-Tikhonov [Whittaker, 1923; Tikhonov, 1963] regularisation resembles homogeneous diffusion i.e. noise is removed but edges become blurred. With Charbonnier penalisation [Charbonnier et al., 1997] edge information can be preserved.

Let us now explain the way from variational image restoration to a PDE or diffusion-based inpainting approach. After that we will also show how to derive a joint inpainting and restoration approach [Chan and Shen, 2002; Weickert and Welk, 2006].

2.3 PDE-Based Inpainting

In the previous sections, we have explained the principles of diffusion filtering and variational image restoration. Both techniques utilise a smoothing process in order to denoise an image. The strength of smoothness can be steered by the diffusion time t and parameter α respectively. This is motivated by the assumption that neighbouring values of the undisturbed image belong together, i.e. their intensities should not vary strongly. This aspect should not only be considered for denoising given data, it may also help to estimate missing information. However, instead of weighting globally be-

tween similarity and smoothness one distinguishes two cases: (i) in regions where information is known, one solely demands similarity while refraining from any smoothness, (ii) in regions where no data is available, one proceeds vice versa and requires smoothness without any similarity constraints [Galić et al., 2005; Weickert and Welk, 2006]. This can be realised by modifying Equation (2.34) to

$$\underbrace{\chi_D \cdot (u - f)}_{\text{similarity}} - (1 - \chi_D) \cdot \underbrace{\operatorname{div}\left(\Psi'(|\boldsymbol{\nabla} u|^2) \cdot \boldsymbol{\nabla} u\right)}_{\text{filling-in}} \stackrel{!}{=} 0 \tag{2.36}$$

whereas the boundary conditions remain unchanged. By $\chi_D := \chi_D(\boldsymbol{x})$ we denote the *characteristic* or *confidence* function

$$\chi_D(\boldsymbol{x}) := \begin{cases} 1 & \text{for } \boldsymbol{x} \in D \,, \\ 0 & \text{for } \boldsymbol{x} \in \Omega_n \backslash D \,. \end{cases} \tag{2.37}$$

Here $D \subset \Omega_n$ denotes the region where the information is known and $\Omega_n \backslash D$ the region where no information is available. Hence, χ_D switches between confidence and smoothness depending on whether information is available or not. Figure 2.5 provides a small impression of the quality of PDE-based inpainting. The characteristic function χ_D is given by Figure 2.5(b). There, black indicates regions where information is missing, i.e. $\chi_D = 0$ and white where it is known, i.e. $\chi_D = 1$.

However, please note that plugging the characteristic function directly into the *partial differential equation* (PDE) is different to that of embedding it into the functional (2.24). This is due to the fact that χ_D can not be treated as a constant within the divergence operator while deriving the minimality condition. Moreover, please keep in mind that such inpainting approaches imply the correctness of the known data, i.e. no degradations such as noise should be present. Otherwise, in case that we also have to denoise the given information, we refer to the next section.

Due to its considerable capabilities, in the meantime, PDE-based inpainting has found its way into image compression [Galić et al., 2005, 2008; Mainberger and Weickert, 2009; Hoffmann et al., 2013]. There, only a small number of cleverly selected pixels are necessary to reconstruct the complete original image with high accuracy.

2.4 Joint Inpainting and Restoration

In the last section, we have illustrated the idea behind PDE-based inpainting. This technique is a valuable tool to fill-in missing information while

Figure 2.6: Joint inpainting and restoration (2.39). *From left to right:* **(a)** Input image with mask from Figure 2.5(b). **(b)** With Whittaker-Tikhonov penaliser ($\alpha = 20$) [Whittaker, 1923; Tikhonov, 1963]. **(c)** With Charbonnier penaliser ($\lambda = 2$, $\alpha = 20$) [Charbonnier et al., 1997]. **(d)** Ditto ($\lambda = 3$, $\alpha = 25$).

given image information is preserved. However, this preservation only makes sense, of course, if the given data is correct. Otherwise, we not only have a preservation of corrupted data but also a filling-in of it. On the other hand, with variational image restoration we have shown a method for denoising and restoring perturbed data.

In this section, we now want to illustrate the combination of both ideas. One strategy for this purpose is presented by Weickert and Welk [2006]. They propose the use of a confidence function in Equation (2.36) being between 0 and 1 in order to perform also regularisation of the known data. A variational strategy is suggested by Chan and Shen [2002]. There, similarity enforced by the data term is only required on the subset D where information is given while regularisation is performed over the whole n-dimensional image domain Ω_n including the known and the missing regions. This can be formulated as

$$E(u) := \int_D (u - f)^2 \, \mathrm{d}\boldsymbol{x} + \alpha \cdot \int_{\Omega_n} \Psi(|\boldsymbol{\nabla} u|^2) \, \mathrm{d}\boldsymbol{x} \ . \tag{2.38}$$

With the help of the characteristic function χ_D (cf. Equation (2.37)), both integrals can be summarised to one:

$$E(u) := \int_{\Omega_n} \left(\chi_D \cdot (u - f)^2 + \alpha \cdot \Psi(|\boldsymbol{\nabla} u|^2) \right) \mathrm{d}\boldsymbol{x} \ . \tag{2.39}$$

Following classical Euler-Lagrange formalism, the scheme of (2.31) eventually leads us to the minimality condition:

$$\chi_D \cdot \underbrace{(u - f)}_{\text{similarity}} - \alpha \cdot \underbrace{\operatorname{div}\left(\Psi'(|\boldsymbol{\nabla} u|^2) \cdot \boldsymbol{\nabla} u \right)}_{\text{smoothness}} \stackrel{!}{=} 0 \ . \tag{2.40}$$

Note that the characteristic function is not applied to the smoothness term, such that no derivatives of it are needed in the minimisation step.

Figure 2.6 compares joint inpainting and restoration approaches with different regularisation strategies. As we can see, using Whittaker-Tikhonov [Whittaker, 1923; Tikhonov, 1963] regularisation already provides sufficient inpainting and denoising performance. Also in this setting, edges can be preserved by applying Charbonnier diffusivity [Charbonnier et al., 1997].

2.5 Deconvolution

This section is devoted to reconstructing information that suffers from a blurred acquisition. This means, we want to invert the imaging model (1.4) in the spatially variant case or (1.5) in the spatially invariant one. Such methods – referred to as *deblurring* or *deconvolution* – belong to the fundamental tasks in image processing.

The task of deblurring or deconvolution is severely ill-posed not offering a unique solution. If not only the sharp image but, at the same time, the point-spread function (PSF), which describes the blur, has to be estimated, the problem will become considerably harder. Such methods are called *blind* deconvolution techniques and among others considered by Ayers and Dainty [1988]; Fish et al. [1995]; Chan and Wong [1998]; You and Kaveh [1996b,a]. In contrast to that, *non-blind* deconvolution methods assume the blurring kernel either to be known or to be estimated in a preprocessing step. In this dissertation, only non-blind deconvolution methods are relevant.

Postulating a spatially invariant blurring process, this section briefly explains the convolution theorem and, based on this, revisits *inverse filtering* as well as *Wiener* deconvolution [Wiener, 1949] as two linear deconvolution models. After that, we discuss the iterative *Richardson-Lucy* (RL) deconvolution model [Richardson, 1972; Lucy, 1974] that also allows a spatially variant blurring kernel. This nonlinear technique is based on Bayes' theorem of conditional probabilities and presumes only positive intensities. Besides that, we illustrate the idea behind variational deconvolution as proposed by e.g. Osher and Rudin [1994] and Marquina and Osher [1999].

2.5.1 Convolution Theorem and Wiener Deconvolution

If the PSF can be assumed to be spatially invariant, the process of a noise-free blurring can be described by $f = h * g$ (cf. Equation (1.5)). However, especially in case of a large support of the PSF, computing the convolution

Figure 2.7: Wiener deconvolution (2.45). *From left to right:* **(a)** Input image and PSF ($\sigma = 3.0$). **(b)** Wiener deconvolution with $K = 0.1$. **(c)** Ditto with $K = 0.01$. **(d)** Ditto with $K = 0.001$.

operation in a straightforward way may be very expensive. As a remedy, one considers the components not in the spatial domain but in the *Fourier* or frequency domain. The Fourier transformation of a 1-D signal ν reads

$$\hat{\nu}(\xi) := \int_{-\infty}^{\infty} \nu(x)\ e^{-i2\pi\xi x} \mathrm{d}x\ , \tag{2.41}$$

where ξ describes the frequency and i is the *imaginary unit* (see e.g. [Gasquet et al., 1998; Bracewell, 1999]). Its inverse is defined by

$$\nu(x) := \int_{-\infty}^{\infty} \hat{\nu}(\xi)\ e^{i2\pi\xi x} \mathrm{d}\xi\ . \tag{2.42}$$

To transfer a higher dimensional signal, one exploits the separability of this transformation. Further, in the discrete setting, the *fast Fourier transform* (FFT), that is based on a divide-and-conquer strategy, allows to reduce the complexity from $\mathcal{O}(N^2)$ of the *discrete Fourier transform* (DFT) to $\mathcal{O}(N\ log(N))$, where N denotes the signal length. However, the FFT requires N to be a power-of-two.

Following the *convolution theorem*, the expensive convolution operation in the spatial domain becomes a cheap multiplication

$$\hat{f} = \hat{h} \cdot \hat{g} \tag{2.43}$$

in the Fourier domain, where $\hat{f}, \hat{h}, \hat{g}$ denote the signals f, h, g respectively given in the Fourier space (see e.g. [Gasquet et al., 1998; Bracewell, 1999]). Besides a more efficient computation, this formulation enables us to solve directly for the unknown $\hat{g}$ by $\hat{g} = \hat{f}/\hat{h}$. However, this deconvolution method, known as *inverse filtering*, diverges for $\hat{h} \to 0$. Hence, as a remedy one can consider the stabilised variant

$$\hat{u} = \begin{cases} \frac{\hat{f}}{\hat{h}} & \text{if} \quad |\hat{h}| > \epsilon \ , \\ 0 & \text{else} \ , \end{cases} \tag{2.44}$$

where u constitutes an approximation of the unknown sharp image g. Since this scheme is still sensitive w.r.t. noise, Wiener [1949] proposed the so-called *Wiener* deconvolution method

$$\hat{u} = \frac{1}{\hat{h}} \ \frac{|\hat{h}|^2}{|\hat{h}|^2 + K} \cdot \hat{f} \ , \tag{2.45}$$

where K acts as a stability parameter. This deconvolution technique aims at minimising the *mean squared error* (MSE) (cf. Section 4.5.1) between u and f. Thus, it is especially suited in the presence of additive Gaussian noise. The influence of the parameter K is demonstrated in Figure 2.7.

2.5.2 Richardson-Lucy Deconvolution

A quite different deconvolution strategy has been proposed independently by Richardson [1972] and Lucy [1974]. To illustrate the basic idea of Richardson [1972], we assume an imaging model with space-variant blur according to Equation (1.4). We denote by f the observed discrete 1-D signal, by g the original undisturbed one, and by H the spatially variant PSF. Instead of being interpreted as grey values, intensities of a signal may be interpreted as frequency values of an event. This means that one can estimate the probability $P(f_i)$ that an event occurs in f at position i by

$$P(f_i) = \frac{f_i}{F} \ , \tag{2.46}$$

where $F := \sum_i f_i$ denotes the total number of events of f. Further, following the law of total probability, it holds that:

$$P(f_i) = \sum_\ell P(f_i \cap g_\ell) = \sum_\ell P(f_i|g_\ell) \cdot P(g_\ell) \ , \tag{2.47}$$

where $P(g_\ell)$ is the probability that an event occurs in g at position ℓ and the conditional probability $P(f_i|g_\ell)$ defines the probability of an event in signal f at position i given an event in g at position ℓ. Analogously, we have

$$P(g_i) = \sum_\ell P(g_i \cap f_\ell) = \sum_\ell P(g_i|f_\ell) \cdot P(f_\ell) \ . \tag{2.48}$$

Figure 2.8: Richardson-Lucy (RL) deconvolution (2.52). *From left to right:* **(a)** Input image and PSF ($\sigma = 3.0$). **(b)** Result after 1 RL deconvolution iteration. **(c)** After 30 iterations. **(d)** After 100 iterations.

With the help of Bayes' theorem

$$P(g_i|f_\ell) = \frac{P(f_\ell|g_i) \cdot P(g_i)}{P(f_\ell)} = \frac{P(f_\ell|g_i) \cdot P(g_i)}{\sum_j P(f_\ell|g_j) \cdot P(g_j)} \tag{2.49}$$

one can reformulate Equation (2.48) to

$$P(g_i) = \sum_\ell \frac{P(f_\ell|g_i) \cdot P(g_i)}{\sum_j P(f_\ell|g_j) \cdot P(g_j)} \cdot P(f_\ell) \ . \tag{2.50}$$

The conditional probability $P(f_\ell|g_i)$ can be described by the PSF itself. Hence, we can replace $P(f_\ell|g_i)$ by $H_{\ell,i}$. Additionally assuming energy preservation, i.e. $P(g_i) = g_i/G$ with $G = F$, brings us to

$$g_i = g_i \cdot \sum_\ell \frac{H_{i,\ell} \cdot f_\ell}{\sum_j H_{i,j} \cdot g_j} \ . \tag{2.51}$$

Introducing a fixed point iteration over k and making use of the $\circledast$ operator yields

$$u^{k+1} = u^k \cdot \left(H^* \circledast \frac{f}{H \circledast u^k} \right) \ , \tag{2.52}$$

where H^* is the adjoint of H, i.e. $H^*_{i,\ell} := H_{\ell,i}$. Initialised with the observed image, Richardson-Lucy deconvolution produces successively sharpened images u^k as approximations of g. If no noise is involved in the imaging model, g constitutes a fixed point of this iteration scheme. However, if the observed image suffers from strong noise, this scheme tends to diverge. This is because the degree of sharpening only depends on the number of iterations in the absence of any additional regularisation. How to extend RL deconvolution to a stabilised and regularised version will be the subject of later sections. Further, we are going to illustrate its justification in terms of Poisson statistics.

Figure 2.9: Variational deconvolution (2.53). *From left to right:* **(a)** Input image and PSF ($\sigma = 3.0$). **(b)** With Whittaker-Tikhonov regularisation ($\alpha = 0.2$) [Whittaker, 1923; Tikhonov, 1963]. **(c)** ditto ($\alpha = 0.5$). **(d)** With Charbonnier Regularisation ($\lambda = 1.0$, $\alpha = 0.2$) [Charbonnier et al., 1997].

Due to this, RL deconvolution is a very popular deblurring technique for low intensity imagery such as astronomical image acquisition imaging or confocal microscopy [Bratsolis and Sigelle, 2001; Dey et al., 2006; Prato et al., 2012; Bertero et al., 2009]. Figure 2.8 illustrates the sharpening process of RL deconvolution.

2.5.3 Variational Deconvolution

In Section 2.2, the reconstruction of information by means of variational methods has been discussed. There, enforced by the data term, the searched u should resemble the acquired signal. For the task of deconvolution, one goes a slightly different way. This time, u should not resemble the observed signal, what of course would end in a blurred estimate. Instead, one is interested in a reconstruction that approximates the observed image only after convolution with the PSF H. The PSF H is assumed to be known. That is, we want to minimise the distance between f and $H \circledast u$. Deviations between both are commonly penalised in a quadratic way. Since deconvolution constitutes an ill-posed problem and noise is still a severe issue, one follows established regularisation strategies. Osher and Rudin [1994] and Marquina and Osher [1999] propose a simultaneous deconvolution and denoising approach. A corresponding energy can be formulated as

$$E(u) := \underbrace{\int_{\Omega_n} (H \circledast u - f)^2 \, \mathrm{d}\boldsymbol{x}}_{\text{data term}} + \alpha \cdot \underbrace{\int_{\Omega_n} \Psi(|\boldsymbol{\nabla} u|^2) \, \mathrm{d}\boldsymbol{x}}_{\text{smoothness term}} . \tag{2.53}$$

Applying additive Euler-Lagrange formalism (cf. Equation (2.31)), we obtain the minimality condition:

$$H^* \circledast (H \circledast u - f) - \alpha \cdot \operatorname{div}\left(\Psi'(|\boldsymbol{\nabla} u|^2) \cdot \boldsymbol{\nabla} u\right) \stackrel{!}{=} 0 , \tag{2.54}$$

where H^* denotes the adjoint of H, i.e. $H^*(\boldsymbol{x}, \boldsymbol{y}) := H(\boldsymbol{y}, \boldsymbol{x})$. The boundary conditions remain unchained as given in Equation (2.35).

Figure 2.9 shows the results of variational deconvolution with different regularisation strategies. As we can see, due to its homogeneous smoothing behaviour Whittaker-Tikhonov regularisation [Whittaker, 1923; Tikhonov, 1963] works against the deblurring process. Instead performing Charbonnier regularisation [Charbonnier et al., 1997] results in a very sharp reconstruction. Besides denoising, regularisation is beneficial to tackle oscillations or ringing artefacts such as shadow-like over- and undershoots near image edges. These artefacts are typical for deconvolution methods. This is demonstrated in Figure 2.10: Standard deconvolution techniques such as Wiener deconvolution [Wiener, 1949] and Richardson-Lucy deconvolution [Lucy, 1974; Richardson, 1972] suffer from oscillation artefacts (cf. Figure 2.10(a),(b) and (c)). These artefacts are successfully damped in Figure 2.10(d), where the result of variational deconvolution with Charbonnier regularisation [Charbonnier et al., 1997] is shown.

Figure 2.10: Direct comparison of the proposed deconvolution methods (2.52). The white rectangle in the top image indicates the origin of the zoom-in shown in the bottom image. *From left to right:* **(a) Column 1:** Wiener deconvolution ($K = 0.01$) [Wiener, 1949]. **(b) Column 2:** Ditto ($K = 0.001$). **(c) Column 3:** RL deconvolution after 100 iterations. **(d) Column 4:** Variational deconvolution with Charbonnier regularisation ($\lambda = 3$, $\alpha = 25$).

Chapter 3

Cell Reconstruction

This chapter is devoted to the enhancement of 3-D cell images recorded with *confocal laser scanning* [Minsky, 1988] or *stimulated emission depletion (STED)* [Hell and Wichmann, 1994] microscopy. Such challenging images play an important role in our cooperation within the *Nano-Cell Interaction Group*. There physicists, biologists, and computer scientists work together and analyse the harmfulness of nano-particles and nano-technology. Nanotechnology means the manipulation of matter on a scale between 1 and 100 nanometres. Since this technology allows the design of novel materials with completely new properties (e.g. dirt- and smell-resistant), it has found its way into modern life in the past years and is well on track to become omnipresent [The Royal Society & The Royal Academy of Engineering, 2004; BMBF, 2010]. In the meantime, one can find nanotechnology in some modern sun-blockers, toothpaste, washing powder, cosmetics, and even food. However, due to their direct contact with the human body and very small size, nano-scale particles might penetrate cell membranes and embed themselves into the cells of our bodies and organs. Consequently, the obvious question about the harmfulness of such a technology arises and researchers have started observing the cell filament network [Maynard, 2006; Lewinski et al., 2008; Krug and Wick, 2011]. If abnormal changes can be detected, this might serve as an indicator of such inflammatory effects (see e.g. [Weber, 2010]). Adequate imaging tools for such investigations are CLSM and STED microscopes. Both allow the imaging of living material and STED additionally offers a very high lateral resolution. However, one has to be aware that these techniques have several systematic drawbacks. For this reason, we will now first perform a detailed analysis of these image acquisition techniques and their individual drawbacks. After that, the aim of this chapter is the development of an enhancement method that is tailored to exactly these inherent weaknesses.

This chapter is based on our journal publication [Persch et al., 2013], which extends the conference publication of Elhayek et al. [2011] in several aspects.

Related work. Before presenting our strategies, first, let us discuss some related approaches. In the context of deconvolution, van Cittert is regarded as one of the first researchers in this field with his more than 80 years old work [Cittert, 1931]. Later, Wiener [1949] popularises the discipline with his seminal deconvolution method (cf. Section 2.5.1). Since then, deconvolution has been the topic of numerous publications. Among those, the Bayesian-based iterative Richardson-Lucy (RL) deconvolution scheme [Richardson, 1972; Lucy, 1974] (cf. Section 2.5.2) and the maximum likelihood-based equivalent proposed by Shepp and Vardi [1982], respectively, are closely related to our approach. To avoid typical over- and undershoots, Bratsolis and Sigelle [2001] propose a regularised RL deconvolution approach where regularisation takes place as a convex combination directly within the iterative scheme. Green [1990] and Panin et al. [1998] propose reconstruction approaches for *single-photon emission computerised tomography* imagery suffering under Poisson noise by means of *expectation-maximisation* (EM) algorithms. Dey et al. [2004, 2006] derive a variational model for the RL scheme and supplement it with an additional *total variation* (TV) [Rudin et al., 1992] regularisation term. They further compare it with a Whittaker-Tikhonov [Whittaker, 1923; Tikhonov, 1963] regularisation term in [Dey et al., 2004]. Using alternating split Bregman techniques, Figueiredo and Bioucas-Dias [2009] consider a discrete TV regularised Csiszár's [Csiszár, 1991] (or Kullback-Leibler) divergence term. In [Setzer et al., 2010], this approach is supplemented by embedding an explicit non-negativity constraint into the functional to be minimised. Sawatzky et al. [2008, 2009], Sawatzky and Burger [2010], and Brune et al. [2011] consider primal-dual algorithms to solve regularised EM or RL deconvolution models. The work of Zanella et al. [2013] focuses on the computational efficiency of RL-based deconvolution techniques. They suggest a *scaled-gradient-projection* method as well as a *graphics processing unit* (GPU) implementation.

A wavelet-regularised RL approach is proposed by Starck and Murtagh [1994] in the context of astronomical imaging. In the context of 3-D confocal microscopy imaging, De Monvel et al. [2001] present a combination of RL deconvolution and wavelet denoising. Also regarding this type of imaging, Vonesch and Unser [2007] propose a wavelet-regularised approach, however, in conjunction with a quadratic similarity term.

Earlier variational non-blind deconvolution models with TV regularisa-

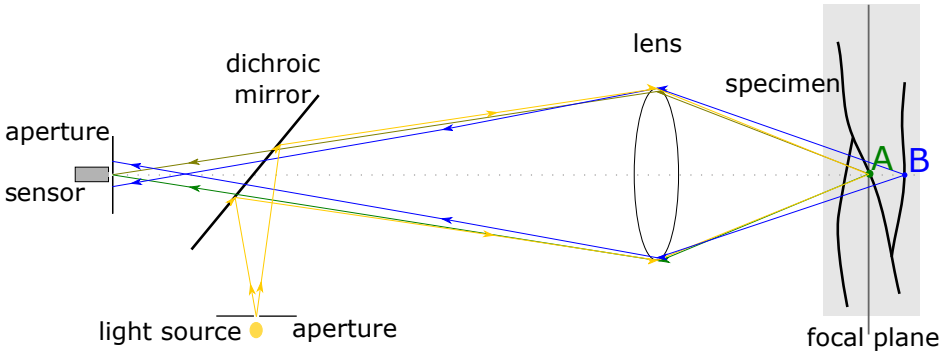

Figure 3.1: Schematic view of a confocal microscope (CM): The aperature in front of the light source allows a point-wise illumination of the focal region. With a second aperture in front of the sensor, only in-focus information reaches the sensor. A dichroic mirror allows to separate excitation and emitted light using the fluorescence technique.

tion date back to the work by Marquina and Osher [1999], which also fits into the more general model of Osher and Rudin [1994]. Furthermore, blind deconvolution models with simultaneous regularisation are proposed by Chan and Wong [1998] and You and Kaveh [1999]. The anisotropic diffusion concepts of Weickert [1998] are introduced to deconvolution by Welk et al. [2005]. In the recent works of Ben Hadj et al. [2013, 2014], blind deconvolution methods accounting for Poisson noise and spatial varying PSF are proposed.

In the context of image interpolation and inpainting [Bertalmío et al., 2000; Galić et al., 2005], a joint model of 2-D image restoration and inpainting is proposed by Chan and Shen [2002] and Weickert and Welk [2006]. In [Chan et al., 2005] inpainting and blind deconvolution are combined in one joint model.

3.1 Image Acquisition

3.1.1 Confocal Laser Scanning Microscope (CLSM)

Figure 3.1 provides a schematic view of the *confocal microscopy* technique. Its functional principle was developed in 1955 and patented in 1957 by Minsky [Minsky, 1988]. Although working with visible light, so that confocal microscopes can be classified into the group of optical or so-called light microscopes, their acquisition strongly differs from that of standard microscopes known from biology lectures in school. Based on the so-called *wide-field* tech-

nique, conventional microscopy illuminates and captures a translucent specimen at once. However when using a lens system, only points lying within the focal plane can be mapped sharply to the image plane (sensor). Regions in front or behind this plane are mapped in a blurred way (cf. the thin lens model of Section 4.1.2). Consequently, a captured 2-D image constitutes a superposition of the sharp information of the focal plane with intensities of the blurred background. In contrast to that, confocal microscopy performs a point-wise illumination as well as a point-wise acquisition by blocking out-of-focus information [Pawley, 2006; Semwogerere and Weeks, 2008]. This is achieved with the help of two apertures (pinholes). Together with the optical system, an aperture in front of the light source bundles a light beam to the focus point of the lens system. By this quite locally restricted illumination, only that part of the translucent specimen which is located near the focus is illuminated brightly. Regarding Figure 3.1, such a situation is given in Point A. However due to the fact that some out-of-focus light may survive, the remaining parts (e.g. Point B) may also receive some weak light and, thus, are not completely dark. As a remedy, a second aperture is placed in front of the imaging sensor. Only light coming from the in-focus region of the lens system can pass this second aperture and arrive at the sensor. This is illustrated by the green beam. Scattered light arriving from out-of-focus areas (blue beam) is blocked. In this way, a very thin *optical sectioning* can be reached such that only a small region around the focal plane is imaged [Semwogerere and Weeks, 2008; Smith, 2008]. Sampling of the probe can be accomplished, e.g. by moving the specimen (*stage-scanning*) or by guiding the beams via small moving mirrors (*beam-scanning*). Eventually, the final image is composed. This way, a confocal microscope is capable of composing a 3-D stack of slices, each showing its actual focal plane. The fact that the aperture in front of the sensor and the illuminated point are confocal to each other explains the origin of the microscope's name. Nowadays, *confocal laser scanning microscopy (CLSM or LSCM)* (see e.g. [Pawley, 2006; Smith, 2008]) constitutes a further stage of this technique. There, the point illumination is accomplished by a laser. It provides a higher intensity than other light sources which becomes especially important when CLSM is combined with the fluorescence technique [Lichtman and Conchello, 2005; Pawley, 2006; Semwogerere and Weeks, 2008]. This method uses fluorophores to label the sub-cellular structure of the cell one is interested in. By adjusting the laser beam to the specific wavelength of the fluorophores, their electrons become excited and reach a higher energetic quantum state. After a short time, the electrons return back to the ground state while emitting photons of lower energy (longer wavelength) than the excitation light (see e.g. [Lichtman and Conchello, 2005]). This process is known as *spontaneous relaxation*. Using

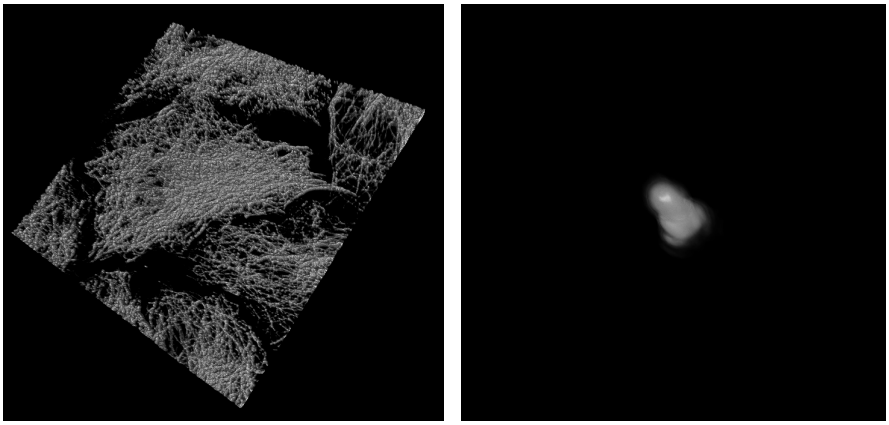

Figure 3.2: **(a) Left:** Volumetric visualisation of the filament network of a cell, recorded with a 3-D confocal laser scanning microscope (CLSM). **(b) Right:** Scaled 3-D point-spread function (PSF). For the visualisation we used *imagevis3d* software package (*http://www.imagevis3d.org*).

a dichroic mirror, excitation and fluorescence light can be separated. This way, only light coming from the structure of interest finally arrives in the detector.
CLSM provides a higher (especially axial) resolution than conventional optical microscopes. By incorporating the fluorescence technique, high contrast images offering better separability of adjacent structures can be acquired. Moreover, while other microscopy techniques such as electron microscopy [Knoll and Ruska, 1932] or scanning tunnelling microscopy [Binnig and Rohrer, 1983] already reached much higher resolutions, CLSM offers the essential advantage of being applicable on living materials.

However, due to the two apertures, only a few photons finally arrive at the detector. This low light intensity naturally leads to Poisson distributed noise [Pawley, 2006; Semwogerere and Weeks, 2008]. By increasing the size of the pinholes, it is possible to reduce the noise level at the expense of more blur. As a consequence, there exists a natural trade-off between the noise and blur level. Moreover, 3-D records of CLSM suffer from a relatively low axial resolution which constitutes a typical problem for light microscopes [Abbe, 1873; Murphy and Davidson, 2012]. A volumetric visualisation of an exemplary 3-D CLSM image together with the estimated PSF is depicted in Figure 3.2. How an associated PSF can be estimated is the subject of Section 3.1.3

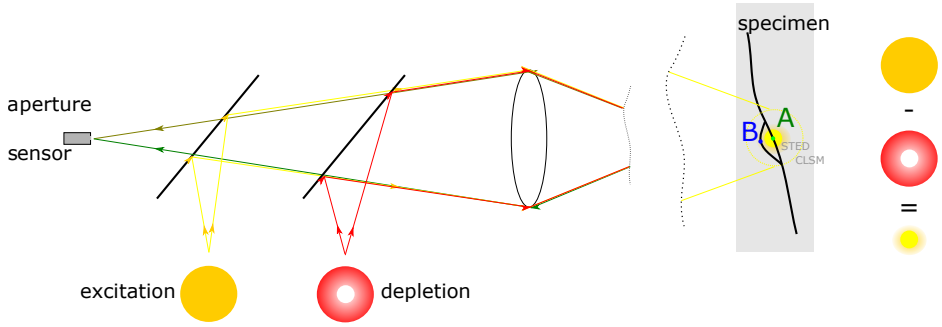

Figure 3.3: Schematic view of a STED microscope: With the help of a depletion beam, already excited fluorophores can be switched off. The effective area of excitation is given as the difference between the excitation and depletion spot. While in CLSM point A and B are excited at the same time, in STED only Point A can remain excited while B can be turned off (and vice versa).

3.1.2 Stimulated Emission and Depletion (STED) Microscopy

Although the CLSM principle already reaches a higher resolution than conventional light microscopes, the beam to excite the fluorophores cannot be focussed arbitrarily sharply. The physical restriction for this lies in the wave nature of light [Hooke, 1665; Ango, 1682; Huygens, 1690]. By that, the size of both the illumination spot and the acquisition point is bounded from below by the diffraction law, which is described by the *Abbe-limit* [Abbe, 1873]. It limits the physical resolution of CLSM to approximately half the used wavelength, i.e. approx 200 nm. Features of the specimen which are located closer to each other than 200 nm appear as a single light spot and are no longer distinguishable.

Eventually, by presenting the STED microscopy, Hell and Wichmann [1994] found a way to bypass the diffraction limit: Besides spontaneous relaxation, fluorophores can also be pushed back to a lower energetic quantum state by *stimulated depletion*. This means, exposing already excited fluorophores to light having a sufficient intensity and a specific wavelength, immediately leads to their de-excitation (depletion). The wavelength of the depletion beam has to be close to that of the fluorescence light. At the same time, photons having the same direction, polarisation, and wavelength as the incident de-excitation light are emitted (see e.g. [Farahani et al., 2010]). Hence, by using a second (red-shifted) depletion beam, which is doughnut-

shaped, already excited molecules can be switched off again. This way, the area emitting photons can be made much smaller. Since the wavelength of the de-excitation light can be tuned to differ slightly to that of the fluorescence light, one can separate both. Consequently, a specific isolation of the region of interest is possible. Regarding Figure 3.3, in CLSM, the adjacent Points A and B lying within the excitation spot (dotted circle) are excited to fluoresce. However, due to diffraction phenomena, they arrive as one blurred spot at the sensor and are thus not separable. Instead, in the STED case, Point B is switched off by stimulated depletion. Therefore, only light emitted by Point A arrives at the sensor. Knowing the position of both beams, the origin of the received light is very well localised. This way, STED is able to reduce the diffraction barrier of Abbe [1873] to a lower one that is formulated by Westphal and Hell [2005]. As a result, nowadays, a lateral resolution of a few nanometres can be reached. However, it is clear that incorporating a depletion beam that goes parallel to the optical axis cannot help to increase the resolution in that direction. Moreover, by de-exciting fluorophores, the overall intensity further shrinks. This in turn brings us back to the compromise between blur and noise. Additionally, due to the high intensity of the depletion beam, fluorophores may permanently lose their ability to fluoresce (so-called *photobleaching*, see e.g. [Bradshaw and Stahl, 2015]).

3.1.3 Physical Estimation of the Point-Spread Function (PSF)

The point-spread function (PSF) delivers a measure for the blurriness of the acquired signal. This constitutes essential information required within our reconstruction task. Since only non-blind deconvolution techniques are considered in this dissertation, we need a way to estimate the PSF of the microscope within a separate preprocessing step. This estimation can be performed either completely theoretically or practically with the help of physical measurements. While in the theoretical case, the estimation is done solely on the basis of the configurations of the microscope, the physical estimation experimentally measures the redistribution of bead shaped light sources within the microscope [Müller, 2006]. Since, here, the light passes the same optical system as during the specimen acquisition, one usually obtains a more realistic PSF. This motivates us to follow the physical way and to analyse 3-D recordings of some small fluorescence beads. Here, one has to pay attention that the recordings of the beads should be chronologically very close to those of the specimen and to avoid any changes of the camera configurations. Instead, the estimation would not reflect the PSF of the original

acquisition process. Moreover, captured bead aggregates should be rejected. Eventually, to extract a PSF, we apply Huygens[2] software package. Besides a few parameters of the device, Huygens requires the diameter of fluorescence beads. Then, Huygens searches the image for suitable beads. Although one acquisition can capture several beads, they should be not too close to each other, and not be located at image boundaries. Images of single beads are extracted, aligned, and averaged to obtain a better noise ratio and attenuate outliers. The connection between a fluorescence bead θ and its appearance $\tilde{\theta}$ can be described by

$$\tilde{\theta} = h * \theta \,, \tag{3.1}$$

where h denotes the searched PSF. Hence, the PSF can be estimated via deconvolution of $\tilde{\theta}$ and θ. To this end, Huygens models theoretically the shape of θ based on the device's settings and the given bead diameter.

Besides showing how to determine the PSF of a device, this section illustrates that our reconstruction framework has to cope with a more or less accurate estimation of the PSF. Further, the PSF is regarded to be spatially invariant which, for instance, ignores diffraction phenomena within the translucent specimen. This motivates the incorporation of robustness ideas into our variational model (see Section 3.2.4).

3.2 Simultaneous Interpolation and Deconvolution

Now that we have understood the physical imaging process of confocal and STED microscopes, let us consider the three main systematic weaknesses of these microscopy techniques: First, we are confronted with image material whose axial resolution (z-direction) is significantly lower than its lateral one (x- and y-direction) [Abbe, 1873; Murphy and Davidson, 2012]. Second, the acquired data is perturbed by strong Poisson distributed noise [Pawley, 2006]. As already mentioned, reasons for that are the low intensity light emitted by fluorescence dyes, the two apertures, and the de-excitation technique of STED. Third, the raw acquired data suffers from out-of-focus blur which can be described by the PSF (as discussed in the previous section) [Müller, 2006].

To counteract degradations such as noise, blur, and partial loss of information, some established methods have already been revisited in the previous chapter. So, let $\Omega_3 \subset \mathbb{R}^3$ denotes the whole three-dimensional image domain, where the image shall be reconstructed and $D \subset \Omega_3$ the region where

[2]Scientific Volume Imaging b.v., Huygens Software, *http://www.svi.nl*

measured image data is available. In practice, usually Ω_3 is a cuboid, and D consists of x-y-slices which are equidistantly spaced in the z-direction. Given the recorded microscopy data $f : D \to \mathbb{R}_+$, to find a denoised version $u : \Omega_3 \to \mathbb{R}_+$ having the desired resolution, one could, for instance, follow Section 2.4 and minimise a functional similar to the one in (2.39).

Although the functional in (2.39) already realises denoising and interpolation, it does not incorporate the handling of out-of-focus blur. Moreover, quadratic penalisation within the data term is justified to Gaussian distributed noise [Steidl and Teuber, 2010; Welk, 2015]. It is inappropriate for Poisson distributed noise which is the primary noise type in such low photon techniques. For this reason, the following sections are devoted to the development of a variational approach that is tailored to exactly these issues: After defining a forward operator which describes the imaging process of the acquired data in a mathematical way, we illustrate how to derive a data term that is justified to Poisson statistics. To find a suitable minimiser, we follow Welk [2010] and revisit the *multiplicative* Euler-Lagrange formalism. Besides restricting the solution to the physical plausible positive range, multiplicative Euler-Lagrange relates the considered Poisson justified data term to iterative Richardson-Lucy (RL) deconvolution. Such a connection allows the extension to a robust and regularised version of RL deconvolution as proposed by Welk [2010]. Moreover, it can be extended to a variational simultaneous interpolation and deconvolution approach [Elhayek et al., 2011].

3.2.1 Forward Operator

The aim of a forward operator is to simulate the mapping of the original undisturbed information to the acquired data. Here, the investigation is restricted to the subset D where the data is actually known. As already mentioned, in confocal microscopy, this data suffers from blur and Poisson distributed noise. Since blur is a redistribution of light energy which we assume to be independent of the location, it can be approximated in terms of a convolution operation defined by $(h*g)(\boldsymbol{x}) := \int_{\Omega_n} h(\boldsymbol{x}-\boldsymbol{s})\cdot g(\boldsymbol{s})\,\mathrm{d}\boldsymbol{s}$ for the n-dimensional case, where $g : \mathbb{R}^n \to \mathbb{R}_+$ represents the original undisturbed information and $h : \mathbb{R}^n \to \mathbb{R}_+$ the PSF. The PSF is estimated physically in a separate processing step according to Section 3.1.3.

Further, to describe the noise, we make use of a spatially independent function $\eta : \mathbb{R}_+ \to \mathbb{R}_+$. Then, the image acquisition can be formulated by the following forward operator:

$$\tilde{\mathcal{F}}[g](\boldsymbol{x}) = \eta\left((h * g)(\boldsymbol{x})\right), \quad \boldsymbol{x} \in D\,. \tag{3.2}$$

As we can see, the forward operator not only consists of a deterministic convolution operation but also of a noise function which acts stochastically. Hence, particularly with respect to our reconstruction task, the question arises how to incorporate such a stochastic part. One strategy is to completely refrain from its modelling within the forward process. Instead, one chooses an adequate penalisation strategy within the data term. This means one exploits a suitable discrepancy measurement between the given and sought data. Thus, while the forward operator reduces to just a convolution operation

$$\mathcal{F}[g](\boldsymbol{x}) = (h * g)(\boldsymbol{x}), \quad \boldsymbol{x} \in D \,, \tag{3.3}$$

the task turns into the modelling of a data term tailored especially towards Poisson distributed noise.

3.2.2 Towards a Physically Justified Data Term

In confocal microscopy mainly two types of noise appear: On the one hand, caused by the imaging sensor and by the signal amplification, additive white Gaussian noise appears which is independent of the light intensity. On the other hand, the dominant type of noise in this setting is Poisson noise. This is due to the fact that the probability of $k \in \mathbb{N}$ photon impacts at the sensor complies with the Poisson distribution [Bovik, 2009]

$$P_\lambda(X = k) = \frac{\lambda^k}{k!} e^{-\lambda} \,, \tag{3.4}$$

where e is the exponential function and $\lambda \in \mathbb{R}_+$ is both the mean and the variance of the distribution. Since the mean can be seen as a long-time average of the acquisition, it corresponds to a capture where no noise is involved. Therefore, expecting a blurred acquisition according to Equation (3.3), one can set $\lambda = \mathcal{F}[g] = h * g$. Moreover, the probability of k photon impacts at a specific location $\boldsymbol{x}_i$ of the sensor is regarded to be independent of the q impacts at a different location $\boldsymbol{x}_j$. Therefore, the joint probability of both events can be calculated by

$$P_\lambda(X = k \text{ and } Y = q) \;=\; P_\lambda(X = k) \cdot P_\lambda(Y = q) \,, \tag{3.5}$$

where the random variables X and Y count the photon impacts at location $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, respectively. Hence, according to Shepp and Vardi [1982], Bratsolis and Sigelle [2001], Dey et al. [2006] and Le et al. [2007], under the assumption that the photon impacts follow a Poisson distribution and all points $\boldsymbol{x} \in D$

are independent, the probability $p(f|g)$ of acquiring image f, given that the true object is g, reads

$$p(f|g) = \prod_{\boldsymbol{x}\in D} \left(\frac{(\mathcal{F}[g](\boldsymbol{x}))^{f(\boldsymbol{x})}}{f(\boldsymbol{x})!}\ e^{-\mathcal{F}[g](\boldsymbol{x})} \right) , \tag{3.6}$$

where $f(\boldsymbol{x})!$ denotes the gamma function applied to $f(\boldsymbol{x})$ (since this term vanishes below, we do not go into detail here). Then, an estimate u of the original undisturbed signal g that leads to the given observation f can be found by following a maximum-likelihood approach. This means that one searches an u that maximises the probability p w.r.t. a fixed f:

$$u = \operatorname*{argmax}_{\tilde{u}} \left(\prod_{\boldsymbol{x}\in D} \left(\frac{(\mathcal{F}[\tilde{u}](\boldsymbol{x}))^{f(\boldsymbol{x})}}{f(\boldsymbol{x})!}\ e^{-\mathcal{F}[\tilde{u}](\boldsymbol{x})} \right) \right) . \tag{3.7}$$

However, searching for a maximiser of such a probability equals the search for a minimiser of its negative. Thus, the problem can be reformulated to

$$u = \operatorname*{argmin}_{\tilde{u}} \left(-\prod_{\boldsymbol{x}\in D} \left(\frac{(\mathcal{F}[\tilde{u}](\boldsymbol{x}))^{f(\boldsymbol{x})}}{f(\boldsymbol{x})!}\ e^{-\mathcal{F}[\tilde{u}](\boldsymbol{x})} \right) \right) . \tag{3.8}$$

Further, by taking the logarithm, multiplications become summations and integration, respectively. This way, the problem can be interpreted as the minimisation of the variational functional

$$\begin{aligned} E_{\mathsf{D}}(f, u, \mathcal{F}) &:= \int_D \left(-\ln \frac{\mathcal{F}[u]^f}{f!} - \ln\left(e^{-\mathcal{F}[u]}\right) \right) \mathrm{d}\boldsymbol{x} \\ &= \int_D \Big(\mathcal{F}[u] + \ln(f!) - f \cdot \ln(\mathcal{F}[u]) \Big) \mathrm{d}\boldsymbol{x} . \end{aligned} \tag{3.9}$$

Based on the fact that any constant additive change of a functional does not affect its minimiser, the term $\ln(f!)$ can be left out as done by Dey et al. [2004, 2006] or be replaced by $\ln(f^f) - f$. The last in turn reveals the equivalence of (3.9) to

$$E^{\mathrm{RL}}(f, u, \mathcal{F}) := \int_D \underbrace{\left(\mathcal{F}[u] - f - f \cdot \ln\left(\frac{\mathcal{F}[u]}{f} \right) \right)}_{=r_{2,f}(\mathcal{F}[u])} \mathrm{d}\boldsymbol{x} , \tag{3.10}$$

which is known as *Csiszár's information divergence* (I-divergence) [Csiszár, 1991]. Due to its justification to Poisson statistics, Csiszár's I-divergence

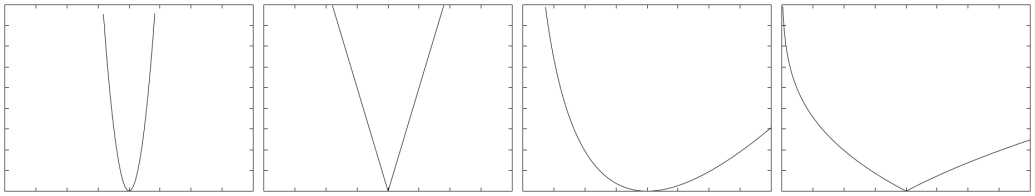

Figure 3.4: Comparison of different discrepancy measures (x-axis: s and y-axis: value of penalisation function). *From left to right:* **(a)** Quadratic penaliser $r_{1,f}(s)$. **(b)** Robust penaliser $\Phi(r_{1,f}(s))$ with $\Phi(r) := 2\sqrt{r}$. **(c)** Asymmetric penaliser $r_{2,f}(s)$. **(d)** Robust asymmetric penaliser $\Phi(r_{2,f}(s))$.

is especially suitable for low-intensity imagery, and, among others, used in [Panin et al., 1998; Anconelli et al., 2005; Sawatzky et al., 2009; Dey et al., 2006; Welk, 2010]. As we can see, instead of penalising the discrepancy between two measures $\mathcal{F}$ and f quadratically, i.e. via a penalty function $r_{1,f}(\mathcal{F}) := (\mathcal{F}-f)^2$ as it is done in Section 2.2.1 or 2.5.3, Csiszár considers an asymmetric penaliser $r_{2,f}(\mathcal{F}) := (\mathcal{F}-f-f\,\ln(\mathcal{F}/f))$. A direct comparison of different penalisers is given in Figure 3.4. Although, $r_{2,f}(\mathcal{F})$ is strictly convex in $\mathcal{F}$ for $\mathcal{F}, f > 0$ and attains its minimum for $\mathcal{F} = f$, it is obvious that a minimiser u of the functional (3.10) is not unique: Assume that u minimises the functional in (3.10), then any ξ that vanishes under convolution with h, i.e. $h * \xi = 0$ induces an additional minimiser $\tilde{u} := u + \xi$ with the same minimal energy. This is because the forward operator would produce the same result in both cases:

$$\mathcal{F}[\tilde{u}] = h * \tilde{u} = h * (u + \xi) = h * u + h * \xi = h * u + 0 = h * u = \mathcal{F}[u]\ .$$

Now that we have found a Poisson justified deconvolution functional or data term (since there is no regularisation involved), we need a strategy to estimate a suitable minimiser. Without the task of inpainting, i.e. $D := \Omega_n$ and assuming an imaging model of $\mathcal{F}[u] := h * u$, different approaches already exist: Snyder et al. [1992] consider Equation (3.10) within a discrete setting. They show that the iterative Richardson-Lucy scheme constitutes a fixed point iteration converging to a minimiser of Csiszár's information divergence. Dey et al. [2004] derive the variational gradient of an equivalent of (3.10) for the 2-D ($\Omega_2 \subset \mathbb{R}^2$) and 3-D case ($\Omega_3 \subset \mathbb{R}^3$) according to an additive perturbation. Subsequently, the resulting Euler-Lagrange equation is solved iteratively in a multiplicative way which in turn corresponds to RL deconvolution. In other words, if $\mathcal{F}[u] := h * u$ and $D := \Omega_n$, Csiszár's I-divergence (3.10) provides a variational interpretation of RL deconvolution.

This fact has been exploited among others by Welk [2010] and allows the enhancement of RL deconvolution by established variational robustification and regularisation ideas. While such ideas are the subject of later sections, we first follow [Welk, 2010, 2015; Elhayek et al., 2011] and illustrate the relation between Csiszár's I-divergence and RL deconvolution. To this end, a minimiser of Equation (3.10) is estimated according to the Euler-Lagrange (EL) formalism. However, instead of classical additive EL, this time, we consider its multiplicative counterpart.

3.2.3 Multiplicative Euler-Lagrange Formalism

To derive the variational gradient of a considered energy functional, the classical Euler-Lagrange formalism [Gelfand and Fomin, 2000], which is based on an additive perturbation, provides a very common procedure (cf. Section 2.2.2). However, deconvolution states a severely ill-posed problem with a non-unique minimiser. To reduce the problem of ill-posedness, a common remedy is given by imposing additional inequality constraints. This way, the solution can be restricted to only physically plausible values. Regarding the imaging process, we know that the number of photons arriving at the image sensor is larger than zero. Hence, the intensity values have to be positive. While the classical additive Euler-Lagrange formalism *does not* impose any constraints on the estimation, in this section, we now follow *multiplicative Euler-Lagrange* formalism as one way to retain the positivity of the solution [Welk, 2010, 2015]. Compared to the classical Euler-Lagrange formalism, here, the additive perturbation is replaced by a multiplicative one. More precisely, within the multiplicative formalism, the variational gradient $\frac{\delta^* E}{\delta u}$ of a functional E is derived by requiring

$$\left\langle \frac{\delta^* E}{\delta u}, v \right\rangle \stackrel{!}{=} \frac{\partial}{\partial \epsilon} E(u \cdot (1 + \epsilon\, v))\Big|_{\epsilon=0}, \quad \forall v\ , \tag{3.11}$$

where v, again, denotes any differentiable perturbation function and $\langle \cdot, \cdot \rangle$ the standard inner product. Assuming a general functional (2.25) and proceeding analogously to the additive case, the right-hand side expands to

$$\begin{aligned} &\frac{\partial}{\partial \epsilon} E(u \cdot (1 + \epsilon\, v)\Big|_{\epsilon=0} \\ &= \frac{\partial}{\partial \epsilon} \int_{\Omega_n} F(x_1, \ldots, x_n, u(1+\epsilon v), (u(1+\epsilon v))_{x_1}, \ldots, (u(1+\epsilon v))_{x_n})\ \mathrm{d}\boldsymbol{x} \Big|_{\epsilon=0} \end{aligned}$$

$$\stackrel{*}{=} \frac{\partial}{\partial\epsilon}\Big($$

$$\int_{\Omega_n} F(x_1,\dots,x_n,u(1+\epsilon v),u_{x_1}+\epsilon u_{x_1}v+\epsilon u v_{x_1},\dots,u_{x_n}+\epsilon u_{x_n}v+\epsilon u v_{x_n})\ \mathrm{d}\boldsymbol{x}\Big)\Big|_{\epsilon=0}$$

$$\stackrel{**}{=} \int_{\Omega_n}\Big(F_u(x_1,\dots,x_n,u(1+\epsilon v),u_{x_1}+\epsilon u_{x_1}v+\epsilon u v_{x_1},\dots,u_{x_n}+\epsilon u_{x_n}v+\epsilon u v_{x_n})\cdot vu$$

$$+\sum_{i=1}^{n}\Big(F_{u_{x_i}}(x_1,\dots,x_n,u(1+\epsilon v),u_{x_1}+\epsilon u_{x_1}v+\epsilon u v_{x_1},\dots,u_{x_n}+\epsilon u_{x_n}v+\epsilon u v_{x_n})$$

$$\cdot\,(u_{x_i}v+v_{x_i}u)\Big)\Big)\ \mathrm{d}\boldsymbol{x}\Big|_{\epsilon=0} \tag{3.12}$$

Since both, the searched function u, as well as the perturbation function v, depend on $\boldsymbol{x}$, we follow the product rule in (*). The chain rule is applied at (**). Now, by setting $\epsilon = 0$, we proceed with

$$\frac{\partial}{\partial\epsilon}E(u\cdot(1+\epsilon\,v)\Big|_{\epsilon=0} = \int_{\Omega_n}\Big(F_u(x_1,\dots,x_n,u,u_{x_1},\dots,u_{x_n})\cdot vu$$

$$+\sum_{i=1}^{n}F_{u_{x_i}}(x_1,\dots,x_n,u,u_{x_1},\dots,u_{x_n})\cdot(u_{x_i}v+v_{x_i}u)\Big)\mathrm{d}\boldsymbol{x}$$

$$\stackrel{***}{=} \int_{\Omega_n}\left(F_u\cdot v+\sum_{i=1}^{n}F_{u_{x_i}}\cdot(uv)_{x_1}\right)\mathrm{d}\boldsymbol{x}\ , \tag{3.13}$$

where we follow the product rule and omit the arguments in (***). By using our notation $\boldsymbol{F}_{\boldsymbol{\nabla}u} := (F_{u_{x_1}},\dots,F_{u_{x_n}})^\top$, we obtain

$$\frac{\partial}{\partial\epsilon}E(u\cdot(1+\epsilon\,v)\Big|_{\epsilon=0} = \int_{\Omega_n}\left(F_u\cdot v+\boldsymbol{F}_{\boldsymbol{\nabla}u}^\top\boldsymbol{\nabla}(uv)\right)\ \mathrm{d}\boldsymbol{x}\ , \tag{3.14}$$

such that Equation (3.11) becomes (by following integration by parts analogously to (2.30))

$$\left\langle\frac{\delta^* E}{\delta u},v\right\rangle \stackrel{!}{=} \int_{\Omega_n}(F_u-\operatorname{div}(\boldsymbol{F}_{\boldsymbol{\nabla}u}))\cdot uv\ \mathrm{d}\boldsymbol{x}$$

$$+\int_{\partial\Omega_n}\left(\boldsymbol{\eta}^\top\boldsymbol{F}_{\boldsymbol{\nabla}u}\right)\cdot uv\ \mathrm{d}\boldsymbol{x}\ ,\quad\forall v\ . \tag{3.15}$$

Here, the image boundary is again described by $\partial\Omega_n$ where $\boldsymbol{\eta}$ defines its unit normal vector. Continuing analogously to Equation (2.31) and (2.32), we first consider the subset of perturbation functions where $v(\boldsymbol{x}) = 0$ for $\boldsymbol{x} \in \partial\Omega_n$ and obtain

$$\begin{aligned} \left\langle \frac{\delta^* E}{\delta u}, v \right\rangle &\stackrel{!}{=} \int_{\Omega_n} \left(F_u - \operatorname{div}\left(\boldsymbol{F}_{\boldsymbol{\nabla} u}\right)\right) \cdot uv \, \mathrm{d}\boldsymbol{x} \\ \iff \qquad \frac{\delta^* E}{\delta u} &= \left(F_u - \operatorname{div}\left(\boldsymbol{F}_{\boldsymbol{\nabla} u}\right)\right) \cdot u \, . \end{aligned} \tag{3.16}$$

If we plug this variational gradient into Equation (3.15), we arrive at the boundary conditions:

$$\begin{aligned} \left\langle \left(F_u - \operatorname{div}\left(\boldsymbol{F}_{\boldsymbol{\nabla} u}\right)\right) \cdot u, v \right\rangle &\stackrel{!}{=} \int_{\Omega_n} \left(F_u - \operatorname{div}\left(\boldsymbol{F}_{\boldsymbol{\nabla} u}\right)\right) \cdot uv \, \mathrm{d}\boldsymbol{x} \\ &\quad + \int_{\partial\Omega_n} \left(\boldsymbol{\eta}^\top \boldsymbol{F}_{\boldsymbol{\nabla} u}\right) \cdot uv \, \mathrm{d}\boldsymbol{x} \, , \quad \forall v \\ \iff \qquad 0 &\stackrel{!}{=} \int_{\partial\Omega_n} \left(\boldsymbol{\eta}^\top \boldsymbol{F}_{\boldsymbol{\nabla} u}\right) \cdot uv \, \mathrm{d}\boldsymbol{x} \, , \quad \forall v \\ \iff \qquad 0 &\stackrel{!}{=} \boldsymbol{\eta}^\top \boldsymbol{F}_{\boldsymbol{\nabla} u} \cdot u \, . \end{aligned} \tag{3.17}$$

Compared to the classical additive Euler-Lagrange formalism, its multiplicative variant now differs through a multiplication with the unknown function for both the variational gradient, i.e. $\frac{\delta^* E}{\delta u} = \frac{\delta E}{\delta u} \cdot u$ and the boundary conditions.

Welk [2010, 2015] illustrates in two different ways why the multiplicative Euler-Lagrange formalism restricts the solution to positive values: On the one hand, one can observe that the multiplicative functional gradient $\frac{\delta^* E}{\delta u}$ occurs within the additive formalism if one replaces the Euclidean metric $\mathrm{d}u$ by a hyperbolic one, i.e. $\mathrm{d}u/u$. Thus, one effectively moves unwanted values to infinite distance. On the other hand, it can also be shown that the multiplicative Euler-Lagrange formalism corresponds to the reparametrisation $u = \exp(w)$. Such a reparametrisation concerning deconvolution in a discrete setting is proposed by Nagy and Strakoš [2000]. Welk and Nagy [2007] demonstrate the advantages of such a reparametrisation in the context of variational deconvolution: While commonly *Gibbs phenomena* – which are characteristic artefacts for deconvolution – are counteracted by regularisation, Welk and Nagy [2007] illustrate that these artefacts are mitigated in their reparametrised approach. As a consequence, the need for regularisation

may be reduced. This implies a possible higher accuracy within such a constrained model. Besides the positivity of scalar values, such a technique is also suggested if the preservation of positive definiteness w.r.t. tensor-valued images is relevant.

From Csiszár's *I*-divergence to RL Deconvolution

With Csiszár's information divergence [Csiszár, 1991], we have found a variational functional that is especially tailored towards Poisson statistics and, therefore, to the considered imagery. With multiplicative Euler-Lagrange formalism, we have a strategy to restrict the solution to the plausible positive range. In this section, we now combine both ideas: Assuming a forward process of $\mathcal{F}[u] := h * u$, we derive the Euler-Lagrange equation of the functional in (3.10) by following the multiplicative formalism. Here, we ignore for the moment the problem of a low axial resolution and assume that acquired information is complete, i.e. we set $D := \Omega_n$. According to Welk [2010, 2015] and Elhayek et al. [2011] besides deriving the associated Euler-Lagrange equation, we illustrate the way from the multiplicative-based minimality condition to the iterative Richardson-Lucy deconvolution scheme. So, let us derive the associated variational gradient $\frac{\delta^* E}{\delta u}$ by requiring the equality in Equation (3.11). For the functional (3.10), the right-hand side of (3.11) expands to

$$
\begin{aligned}
&\frac{\partial}{\partial \epsilon} E^{\mathrm{RL}}(f, u \cdot (1+\epsilon v), \mathcal{F}) \,\Big|_{\epsilon=0} \\
&= \frac{\partial}{\partial \epsilon} \int_{\Omega_n} \left(\mathcal{F}[u \cdot (1+\epsilon v)] - f - f \,\ln\left(\frac{\mathcal{F}[u \cdot (1+\epsilon v)]}{f} \right) \right) \,\mathrm{d}\boldsymbol{x} \,\Big|_{\epsilon=0} \\
&= \frac{\partial}{\partial \epsilon} \int_{\Omega_n} \left(h * (u \cdot (1+\epsilon v)) - f - f \,\ln\left(\frac{h * (u(1+\epsilon v))}{f} \right) \right) \,\mathrm{d}\boldsymbol{x} \,\Big|_{\epsilon=0} \\
&= \frac{\partial}{\partial \epsilon} \int_{\Omega_n} \left(h * u + \epsilon\, h * (uv) - f - f \,\ln\left(\frac{h * u + \epsilon\, h * (uv)}{f} \right) \right) \,\mathrm{d}\boldsymbol{x} \,\Big|_{\epsilon=0} \\
&= \int_{\Omega_n} \left(h * (uv) - \left(\frac{f}{h * u + \epsilon\, h * (uv)} \right) \cdot (h * (uv)) \right) \,\mathrm{d}\boldsymbol{x} \,\Big|_{\epsilon=0}
\end{aligned}
$$

$$= \int_{\Omega_n} \left(\left(1 - \frac{f}{h * u + \epsilon\ h * (uv)} \right) \cdot (h * (uv)) \right)\ \mathrm{d}\boldsymbol{x} \Bigg|_{\epsilon=0}$$

$$= \int_{\Omega_n} \left(\left(1 - \frac{f}{h * u} \right) \cdot (h * (uv)) \right)\ \mathrm{d}\boldsymbol{x}\ . \tag{3.18}$$

By applying the definition of convolution, we can proceed with

$$\frac{d}{d\epsilon} E^{\mathrm{RL}}(f, u \cdot (1 + \epsilon v), \mathcal{F}) \Big|_{\epsilon=0}$$

$$= \int_{\Omega_n} \left(\left(1 - \frac{f}{h * u} \right)(\boldsymbol{x}) \cdot \int_{\Omega_n} \left(h(\boldsymbol{x} - \boldsymbol{s})\ u(\boldsymbol{s})\ v(\boldsymbol{s}) \right) \mathrm{d}\boldsymbol{s} \right)\ \mathrm{d}\boldsymbol{x}$$

$$= \int_{\Omega_n} \int_{\Omega_n} \left(\left(1 - \frac{f}{h * u} \right)(\boldsymbol{x}) \cdot h(\boldsymbol{x} - \boldsymbol{s})\ u(\boldsymbol{s})\ v(\boldsymbol{s}) \right)\ \mathrm{d}\boldsymbol{s}\ \mathrm{d}\boldsymbol{x}$$

$$= \int_{\Omega_n} \int_{\Omega_n} \left(h^*(\boldsymbol{s} - \boldsymbol{x}) \cdot \left(1 - \frac{f}{h * u} \right)(\boldsymbol{x}) \cdot u(\boldsymbol{s})\ v(\boldsymbol{s}) \right)\ \mathrm{d}\boldsymbol{s}\ \mathrm{d}\boldsymbol{x}\ , \tag{3.19}$$

where $h^*(\boldsymbol{x}) := h(-\boldsymbol{x})$. Now, if we change the order of integration as follows

$$\frac{d}{d\epsilon} E^{\mathrm{RL}}(f, u \cdot (1 + \epsilon v), \mathcal{F}) \Big|_{\epsilon=0}$$

$$= \int_{\Omega_n} \left(\int_{\Omega_n} \left(h^*(\boldsymbol{s} - \boldsymbol{x}) \cdot \left(1 - \frac{f}{h * u} \right)(\boldsymbol{x})\ \mathrm{d}\boldsymbol{x} \right) \cdot u(\boldsymbol{s}) v(\boldsymbol{s}) \right)\ \mathrm{d}\boldsymbol{s}\ , \tag{3.20}$$

one can in turn make use of the convolution notation. This leads us to

$$\frac{d}{d\epsilon} E^{\mathrm{RL}}(f, u \cdot (1 + \epsilon v), \mathcal{F}) \Big|_{\epsilon=0} = \int_{\Omega_n} \left(\left(h^* * \left(1 - \frac{f}{h * u} \right) \right)(\boldsymbol{s}) \cdot u(\boldsymbol{s}) v(\boldsymbol{s}) \right)\ \mathrm{d}\boldsymbol{s}\ . \tag{3.21}$$

Embedding this result into the requirement in (3.11), and substituting the integration variable $\boldsymbol{s}$ by $\boldsymbol{x}$ brings us to

$$\left\langle \frac{\delta^* E}{\delta u}, v \right\rangle \stackrel{!}{=} \int_{\Omega_n} \left(\left(h^* * \left(1 - \frac{f}{h * u} \right) \right)(\boldsymbol{x}) \cdot u(\boldsymbol{x}) v(\boldsymbol{x}) \right)\ \mathrm{d}\boldsymbol{x}\ , \quad \forall v\ . \tag{3.22}$$

Following the definition of the scalar product immediately reveals the functional derivative

$$\frac{\delta^* E}{\delta u} = \left(h^* * \left(1 - \frac{f}{h * u}\right)\right) \cdot u \; . \tag{3.23}$$

If the support of the PSF h exceeds the image boundaries, we impose mirrored boundary conditions on u.

According to the minimality condition (2.26), a minimiser of (3.10) has to fulfil a vanishing variational gradient, i.e.

$$\begin{aligned} & \frac{\delta^* E}{\delta u} = 0 \\ \iff \quad & \left(h^* * \left(1 - \frac{f}{h * u}\right)\right) \cdot u = 0 \\ \iff \quad & \left(h^* * 1 - h^* * \frac{f}{h * u}\right) \cdot u = 0 \; . \end{aligned} \tag{3.24}$$

Since one additionally assumes energy preservation, i.e. $h^* * 1 = 1$ where 1 is the constant function 1, we obtain

$$\begin{aligned} & \left(1 - h^* * \frac{f}{h * u}\right) \cdot u = 0 \\ \iff \quad & u - \left(h^* * \frac{f}{h * u}\right) \cdot u = 0 \; . \end{aligned} \tag{3.25}$$

Eventually, introducing a fixed point iteration in k yields the well known Richardson-Lucy (RL) deconvolution algorithm [Lucy, 1974; Richardson, 1972] of Section 2.5.2:

$$u^{k+1} = \left(h^* * \frac{f}{h * u^k}\right) \cdot u^k \; . \tag{3.26}$$

As already mentioned, initialised with the observed image, i.e. $u^0 := f$, this scheme produces successively sharpened images u^k. Further, in the absence of noise, i.e. $f = h * g$, one fixed point of this scheme is given by g. This can be shown as follows: Assume g be given as an intermediate solution at iteration step k, i.e. $u^k = g$, then any further step

$$u^{k+1} = \left(h^* * \frac{f}{h * g}\right) \cdot g = \left(h^* * \frac{f}{f}\right) \cdot g = (h^* * 1) \cdot g = g \; .$$

would not change anything so that $u^{k+1} = u^k = g$ which is the definition of a fixed point. However, in practice if noise is present, the scheme has to be stopped after a certain number of iterations since it is known to diverge for $k \to \infty$. The fact that noise can be tackled by limiting the total number of iterations can be interpreted as some kind of regularisation [Welk, 2010]. However, this, of course, would not only suppress noise, but it would also inhibit the sharpening process. For this reason, one is interested in an explicit (edge preserving) regularisation.

3.2.4 Robust Regularised Richardson-Lucy Deconvolution

In the last section, we inspected the relation between the iterative Richardson-Lucy deconvolution technique and the variational Csiszár's information divergence (for $D = \Omega_n$, $\mathcal{F}[u] := h * u$) by means of the multiplicative Euler-Lagrange formalism. Besides illustrating the tailoring of RL deconvolution towards Poisson statistics, this variational interpretation gives direct access to established regularisation and robustification ideas. Hence, the goal of this section is to combine the ideas from Section 2.2 and 2.5.3 with the variational RL interpretation from the last section. Analogously to Dey et al. [2004, 2006], Sawatzky et al. [2009]; Sawatzky and Burger [2010], Welk [2010, 2015], and Elhayek et al. [2011], we thus extend the energy functional E^{RL} (3.10) with an explicit regularisation term E_{S} to:

$$E^{\mathrm{RRL}}(u) := \underbrace{E^{\mathrm{RL}}(f, u, \mathcal{F})}_{\text{RL-deconvolution}} + \alpha \cdot \underbrace{E_{\mathsf{S}}(u)}_{\text{smoothness}} \quad , \tag{3.27}$$

where we chose

$$E_{\mathsf{S}}(u) := \int_D \Psi(|\boldsymbol{\nabla} u|^2) \, \mathrm{d}\boldsymbol{x} \; , \tag{3.28}$$

where $\boldsymbol{\nabla}$ is the n-dimensional gradient operator.

As we can see, the variational interpretation E^{RL} of RL deconvolution represents the data term E_{D}. As a result, instead of limiting the total number of iterations, now the amount of smoothness can be steered by the parameter α. This offers a better balance between deblurring and regularisation. This way, we can counteract the problem of non-uniqueness, noise, as well as the occurrence of ringing artefacts, without stopping the sharpening process.

According to Equation (3.16) and (3.23), under the assumption that the forward operator is defined by $\mathcal{F}[u] := h * u$, the associated minimality condition [Welk, 2010; Elhayek et al., 2011; Welk, 2015] of this supplemented

approach reads

$$\left(h^* * \left(1 - \frac{f}{h * u} \right) - \alpha \cdot \operatorname{div} \left(\Psi'(|\boldsymbol{\nabla} u|^2) \; \boldsymbol{\nabla} u \right) \right) \cdot u \; = \; 0 \; , \tag{3.29}$$

where the boundary conditions (3.17) have to be fulfilled (e.g. by mirroring the image at its boundaries).

Besides regularisation, this variational interpretation of Richardson-Lucy deconvolution allows to introduce the concept of robust statistics [Huber, 2004] to the data term as suggested by Welk [2010, 2015] and Elhayek et al. [2011] or by Zervakis et al. [1995], Bar et al. [2005], and Welk and Nagy [2007] in the context of variational deconvolution. To this end, we apply a non-negative sub-linear increasing penalisation function $\Phi : \mathbb{R}_{0,+} \to \mathbb{R}_{0,+}$ within the RL deconvolution expression E^{RL}. Together with regularisation, we arrive at the *robust regularised Richardson-Lucy (RRRL)* scheme of Welk [2010, 2015]:

$$E^{\mathrm{RRRL}}(u) \; := \; \int_D \Bigg(\Phi \Bigg(\underbrace{\mathcal{F}[u] - f - f \, \ln \left(\frac{\mathcal{F}[u]}{f} \right)}_{=: r_{2,f}(\mathcal{F}[u])} \Bigg) + \; \alpha \cdot \Psi(|\boldsymbol{\nabla} u|^2) \Bigg) \, \mathrm{d}\boldsymbol{x} \; . \tag{3.30}$$

By that, we gain robustness against outliers and imprecisions in the deconvolution model. Such imprecisions can be caused, e.g. due to the fact that we do not have the correct PSF, but only an estimation of it which is additionally simplified to be a spatially-invariant one (cf. Section 3.1.3). Moreover, not all the details of the imaging process can be incorporated or modelled correctly, such as variations in sensor gain or photobleaching.

To grow less than linearly and thereby penalise outliers less severely, we chose the penalising function $\Phi(s) := 2\sqrt{s + \beta}$ (cf. Figure 3.4(d)) with the small regularisation constant $\beta > 0$. The associated minimality condition reads

$$\left(h^* * \left(\Phi'(r_{2,f}(\mathcal{F}[u])) \left(1 - \frac{f}{h * u} \right) \right) - \alpha \cdot \operatorname{div}(\Psi'(|\boldsymbol{\nabla} u|^2) \boldsymbol{\nabla} u) \right) \cdot u \; = \; 0 \; . \tag{3.31}$$

In the following Section, we will discuss the derivation of the minimality condition in detail.

3.2.5 Joint Variational Approach

As a simplification in the previous sections, we have assumed that the acquired data is complete, i.e. we have set $D := \Omega_n$. Now, let us come back to

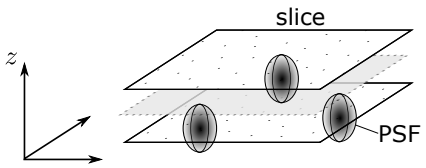

Figure 3.5: Support of the PSF overlaps missing slices to be inpainted (light grey). Consequently, the blurring process carries partially missing information to acquired slices (solid black).

the additional problem of a low axial resolution. Hence, assuming that the data is only given at a subset $D \subset \Omega_n$, we want to reconstruct the information on the whole image domain Ω_n. To solve both tasks of deconvolution and interpolation, clearly we could apply the methods for inpainting from Section 2.3 and deconvolution from Section 3.2.4 sequentially. However, as Chan et al. [2005] point out, such a naive solution is disadvantageous: Interpolation followed by deconvolution propagates blurred information into missing regions. By reversing the order of application, the support of the PSF overlaps with the unknown areas. Moreover, as illustrated in Figure 3.5, the acquired data contains information about the signal in the domain to be inpainted due to the convolution that is part of the image acquisition process.

This motivates a joint variational solution as advocated by Chan et al. [2005] in the context of blind deconvolution, and by Elhayek et al. [2011] in the context of RRRL deconvolution. The latter proposes the following energy functional for *simultaneous interpolation and RRRL deconvolution (IRRRL)*:

$$E^{\mathrm{IRRRL}}(u) := \int_D \Phi\Bigg(\underbrace{\mathcal{F}[u] - f - f \, \ln\left(\frac{\mathcal{F}[u]}{f}\right)}_{=: r_{2,f}(\mathcal{F}[u])} \Bigg) \mathrm{d}\boldsymbol{x} + \alpha \int_{\Omega_n} \left(\Psi(|\boldsymbol{\nabla} u|^2)\right) \, \mathrm{d}\boldsymbol{x} \, . \tag{3.32}$$

By restricting the data term (or deconvolution term) to only the acquired subset D but at the same time demanding smoothness over the whole image domain Ω_n, the functional (3.32) performs simultaneous interpolation and deconvolution.

Since, compared to the functional in (3.27), nothing is changed w.r.t. the regularisation term E_{S}, let us for the moment focus on the data term $E_{\mathsf{D}}^{\mathrm{IRRRL}}$ of E^{IRRRL}. To derive the variational gradient of $E_{\mathsf{D}}^{\mathrm{IRRRL}}$, we proceed analo-

gously to (3.18) et seqq. and obtain:

$$\frac{\partial}{\partial\epsilon} E_{\mathsf{D}}^{\mathrm{IRRRL}}(f, u\cdot(1+\epsilon v), \mathcal{F}) \Big|_{\epsilon=0}$$

$$= \frac{\partial}{\partial\epsilon} \int_D \Phi\Big(r_{2,f}(\mathcal{F}[u(1+\epsilon v)])\Big)\, \mathrm{d}\boldsymbol{x} \Big|_{\epsilon=0}$$

$$= \frac{\partial}{\partial\epsilon} \int_D \Phi\left(h*(u\cdot(1+\epsilon v)) - f - f\,\ln\left(\frac{h*(u(1+\epsilon v))}{f}\right)\right)\, \mathrm{d}\boldsymbol{x} \Big|_{\epsilon=0}$$

$$= \int_D \Phi'\Big(r_{2,f}(\mathcal{F}[u(1+\epsilon v)])\Big)\cdot\left(\left(1-\frac{f}{h*u+\epsilon\, h*(uv)}\right)\cdot(h*(uv))\right)\, \mathrm{d}\boldsymbol{x} \Big|_{\epsilon=0}$$

$$= \int_D \Phi'\Big(r_{2,f}(\mathcal{F}[u])\Big)\cdot\left(\left(1-\frac{f}{h*u}\right)\cdot(h*(uv))\right)\, \mathrm{d}\boldsymbol{x}\,. \tag{3.33}$$

At this point, one has to take care that now the convolution domain (which is given by Ω_n) in the forward operator $\mathcal{F}$ differs from the integration domain D. Using the definition of convolution and the abbreviation $\Phi' := \Phi'\left(r_{2,f}(\mathcal{F}[u])\right)$, we thus have

$$\frac{d}{d\epsilon} E_{\mathsf{D}}^{\mathrm{IRRRL}}(f, u\cdot(1+\epsilon v), \mathcal{F})\Big|_{\epsilon=0}$$

$$= \int_D \left(\Phi'(\boldsymbol{x})\cdot\left(1-\frac{f}{h*u}\right)(\boldsymbol{x})\cdot\int_{\Omega_n}\Big(h(\boldsymbol{x}-\boldsymbol{s})\; u(\boldsymbol{s})\; v(\boldsymbol{s})\Big)\, \mathrm{d}\boldsymbol{s}\right)\, \mathrm{d}\boldsymbol{x} \tag{3.34}$$

However, with the help of the characteristic function χ_D, defined in Equation (2.37), both integration domains can be harmonised as follows:

$$\frac{d}{d\epsilon} E_{\mathsf{D}}^{\mathrm{IRRRL}}(f, u\cdot(1+\epsilon v), \mathcal{F})\Big|_{\epsilon=0}$$

$$= \int_{\Omega_n}\left(\chi_D(\boldsymbol{x})\cdot\Phi'(\boldsymbol{x})\cdot\left(1-\frac{f}{h*u}\right)(\boldsymbol{x})\cdot\int_{\Omega_n}\Big(h(\boldsymbol{x}-\boldsymbol{s})\; u(\boldsymbol{s})\; v(\boldsymbol{s})\Big)\, \mathrm{d}\boldsymbol{s}\right)\, \mathrm{d}\boldsymbol{x}$$

$$= \int_{\Omega_n}\left(\int_{\Omega_n}\left(h^*(\boldsymbol{s}-\boldsymbol{x})\cdot\left(\chi_D(\boldsymbol{x})\cdot\Phi'(\boldsymbol{x})\left(1-\frac{f}{h*u}\right)(\boldsymbol{x})\, \mathrm{d}\boldsymbol{x}\right)\right) u(\boldsymbol{s})v(\boldsymbol{s})\right)\, \mathrm{d}\boldsymbol{s}$$

$$= \int_{\Omega_n} \left(h^* * \left(\chi_D \cdot \Phi' \cdot \left(1 - \frac{f}{h * u} \right) \right)(\boldsymbol{s}) \cdot u(\boldsymbol{s}) v(\boldsymbol{s}) \right) \, \mathrm{d}\boldsymbol{s} \; . \tag{3.35}$$

Hence, the corresponding minimality condition for the functional (3.32) reads [Elhayek et al., 2011]

$$\left(h^* * \left(\chi_D \cdot \Phi'(r_{2,f}(\mathcal{F}[u])) \cdot \left(1 - \frac{f}{u * h} \right) \right) - \alpha \cdot \operatorname{div}(\Psi'(|\boldsymbol{\nabla} u|^2) \boldsymbol{\nabla} u) \right) \cdot u = 0 \; . \tag{3.36}$$

Similarly to Equation (2.39) the characteristic function χ_D is used to impose the fidelity condition $\mathcal{F}[u] := h * u \approx f$ only at those locations where data is available. In the minimisation process, information is transferred from D to $\Omega \setminus D$ by both the enforced smoothness in the regularisation term and the convolution operations in the data term. A coupling of neighbours thus not only exists by regularisation but also within the data term.

Even though this model already achieves an interpolation compatible with the blur model, it does not respect the directional information of the underlying cell structure. Hence, let us discuss in the next section, how to turn the smoothing behaviour induced by regularisation into a fibre enhancement technique with the help of anisotropic diffusion (cf. Section 2.1.2).

3.3 Fibre Enhancement with Anisotropic Regularisation

In this section, we want to illustrate how we can exploit the physical process of anisotropic diffusion in order to enhance 3-D images capturing the filament network of cells. In previous sections, we have already discussed variational restoration and deconvolution models, which utilise the concept of regularisation in some sense. The need of regularisation has been argued by the problem of non-uniqueness, noise, and ringing artefacts. Moreover, regularisation has been extended to achieve an inpainting approach. As already mentioned in Section 2.2.2, mathematically, such variational regularisation terms, which penalise the derivative of the unknown, cause a divergence expression of type

$$\operatorname{div} \left(\, \Psi'(|\boldsymbol{\nabla} u|^2) \, \boldsymbol{\nabla} u \, \right) \tag{3.37}$$

in the associated minimality conditions. Such expressions are known to lead to scalar-valued (isotropic) diffusion [Perona and Malik, 1987]. However, the aim of this study is the enhancement and reconstruction of tube-like patterns such as the microtubules, which are part of the intracellular structure of cells.

For this particular application, the isotropic behaviour of the latter class of processes is not well suited.

In the context of image denoising and smoothing, we have revisited the physical phenomena of diffusion in Section 2.1. Besides isotropic diffusion, we have also recapitulated the idea behind anisotropic diffusion [Weickert, 1998] in Section 2.1.2. By incorporating directional information of the underlying image structure, anisotropic diffusion allows one to steer the flux and thereby the smoothing process along such elongated structures. This aspect makes anisotropic diffusion in particular attractive for our cell enhancement task.

Hence, let us now exchange the isotropic smoothing behaviour by an anisotropic one. To this end, we proceed as in Section 2.1.2, and replace the scalar-valued diffusivity $\Psi' \in \mathbb{R}_+$ in Equation (3.37) by a diffusion tensor $\boldsymbol{D}$. This leads to the right-hand side of Equation (2.21):

$$\operatorname{div}\Big(\boldsymbol{D}(\boldsymbol{J}_\rho(u_\sigma)) \cdot \boldsymbol{\nabla} u\Big) \,. \tag{3.38}$$

In this way, our strategy is related to the work of Welk et al. [2005] in the context of deblurring and to Galić et al. [2008] in the context of inpainting and compression, respectively. To gather local directional information, we again choose the structure tensor $\boldsymbol{J}$ of Förstner and Gülch [1987]. Since we are dealing with 3-D microscopy data sets, $\boldsymbol{D}$ as well as $\boldsymbol{J}$ are 3×3 tensor fields. Moreover, we assume that the recorded data only consist of one channel. In this setting, the structure tensor is given by

$$\begin{aligned}
\boldsymbol{J}_\rho(\boldsymbol{\nabla} u_\sigma) :=& \mathcal{K}_\rho * \left(\begin{pmatrix} u_{\sigma x_1} \\ u_{\sigma x_2} \\ u_{\sigma x_3} \end{pmatrix} \begin{pmatrix} u_{\sigma x_1} \\ u_{\sigma x_2} \\ u_{\sigma x_3} \end{pmatrix}^\top \right) \\
=& \mathcal{K}_\rho * \begin{pmatrix} u_{\sigma x_1}^2 & u_{\sigma x_1} u_{\sigma x_2} & u_{\sigma x_1} u_{\sigma x_3} \\ u_{\sigma x_1} u_{\sigma x_2} & u_{\sigma x_2}^2 & u_{\sigma x_2} u_{\sigma x_3} \\ u_{\sigma x_1} u_{\sigma x_3} & u_{\sigma x_2} u_{\sigma x_3} & u_{\sigma x_3}^2 \end{pmatrix} \,.
\end{aligned} \tag{3.39}$$

Since the directional information of the underlying image structure is contained in the eigenvectors and eigenvalues of this symmetric, positive semi-definite tensor, we consider the eigendecompositon

$$\boldsymbol{J}_\rho(\boldsymbol{\nabla} u_\sigma) = (\boldsymbol{v}_1 | \boldsymbol{v}_2 | \boldsymbol{v}_3) \begin{pmatrix} \mu_1 & & \\ & \mu_2 & \\ & & \mu_3 \end{pmatrix} \begin{pmatrix} \boldsymbol{v}_1^\top \\ \boldsymbol{v}_2^\top \\ \boldsymbol{v}_3^\top \end{pmatrix} \,, \tag{3.40}$$

where $\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3$ are orthonormal eigenvectors.

Now, the essential idea is that the image intensity does not vary much along a cell fibre. Hence, we are searching for the direction of lowest contrast, which is given by the eigenvector $\boldsymbol{v_3}$ corresponding to the smallest eigenvalue μ_3. In order to enhance such elongated approximately one-dimensional cell structures in 3-D space, we demand smoothness only in the direction of the lowest contrast, i.e. along $\boldsymbol{v_3}$. This smoothing process should also be able to close small gaps caused by, e.g. an imperfect labelling with the fluorescence dyes. On the other hand, the reconstruction should be sharp so that the single fibres are clearly differentiable against their background. Thus, smoothing perpendicular to $\boldsymbol{v_3}$, i.e. along $\boldsymbol{v_1}$ and $\boldsymbol{v_2}$ respectively should be penalised in dependence to the respective contrast described by μ_1 and μ_2. To achieve such a directionally dependent smoothing behaviour, we assemble the diffusion tensor as follows

$$\boldsymbol{D}(\boldsymbol{J}_\rho(\boldsymbol{\nabla} u_\sigma)) := (\boldsymbol{v}_1|\boldsymbol{v}_2|\boldsymbol{v}_3) \begin{pmatrix} \Psi'(\mu_1) & & \\ & \Psi'(\mu_2) & \\ & & 1 \end{pmatrix} \begin{pmatrix} \boldsymbol{v}_1^\top \\ \boldsymbol{v}_2^\top \\ \boldsymbol{v}_3^\top \end{pmatrix} , \tag{3.41}$$

where we apply the Charbonnier diffusivity [Charbonnier et al., 1997], i.e. $\Psi'(s^2) = 1/\sqrt{1 + s^2/\lambda^2}$ to the two largest eigenvalues μ_1 and μ_2. The third eigenvalue is strictly set to 1 in order to perform homogeneous diffusion along the eigenvector $\boldsymbol{v}_3$, i.e. along the fibres. As reasoned above, this choice is suitable to enhance the tube-like structures of microtubules. Eventually, replacing the diffusivity in Equation (3.36) by this anisotropic term, we obtain

$$\begin{aligned} u \cdot \Bigg(h^* * \Bigg(\chi_D \cdot \Phi'(r_f(\mathcal{F}[u])) \left(1 - \frac{f}{u * h}\right) \Bigg) \\ - \alpha \cdot \operatorname{div}\left(\boldsymbol{D}(\boldsymbol{J}_\rho(\boldsymbol{\nabla} u_\sigma))\ \boldsymbol{\nabla} u\right) \Bigg) \ = \ 0 \, . \end{aligned} \tag{3.42}$$

We call this *interpolating robust and (an)isotropic regularised Richardson-Lucy* (IRARRL) method. How to solve such a PDE efficiently, as well as its application to discrete digital images, is the central topic of the following sections.

3.4 Efficient and Stabilised Iteration Scheme

This section is devoted to the task of solving Equation (3.42) by means of an efficient and stabilised gradient descent scheme. We want to illustrate

that a clever evaluation of the unknown w.r.t. the old and the new time step may drastically improve the performance of the overall method. Generally, computing a solution of isotropic regularised RL (or equivalent models) is a non-trivial task and different strategies for this exist in the literature [Green, 1990; Panin et al., 1998; Sawatzky et al., 2008; Dey et al., 2004; Sawatzky et al., 2009; Bertero et al., 2009; Sawatzky and Burger, 2010; Setzer et al., 2010; Welk, 2010].

Dey et al. [2004, 2006] propose for their related isotropic approach without any inpainting or robustification strategy the fixed point scheme

$$u^{k+1} - \left(h^* * \frac{f}{h * u^k} \right) \cdot u^k - \alpha \ \mathrm{div} \left(\Psi'(|\boldsymbol{\nabla} u^k|^2) \boldsymbol{\nabla} u^k \right) \cdot u^{k+1} = 0 \ , \tag{3.43}$$

where the second term, as well as the divergence term, are evaluated completely at the old time step k. Only the first term and the factor behind the divergence term are evaluated at the new time step $k+1$. Eventually, solving for u^{k+1} yields

$$u^{k+1} = \frac{u^k}{1 - \alpha \ \mathrm{div} \left(\Psi'(|\boldsymbol{\nabla} u^k|^2) \boldsymbol{\nabla} u^k \right)} \cdot \left(h^* * \frac{f}{h * u^k} \right) \ . \tag{3.44}$$

However, as we can see, the divergence term has moved into the denominator. This, in turn, leads to the risk of evaluating to negative values or even divisions by zero. To prevent such a scenario, Dey et al. [2004, 2006] suggest the restriction of regularisation by allowing only small values for α. Such a restriction has already been proposed by Green [1990] and Panin et al. [1998], where an expectation-maximisation (EM) algorithm is performed. In order to avoid such a restriction of the regularisation parameter, Welk [2010] proposes a different strategy. He refrains from strictly multiplying the divergence term in Equation (3.31) with u^{k+1}. Instead, the multiplication is made conditionally, depending on the sign of the divergence term: If the divergence term has a negative sign, it is multiplied with u^{k+1} from the new time step. Otherwise, the old one u^k is chosen. This can be formulated as

$$\left(h^* * \Phi_D'^k \right) u^{k+1} - \left(h^* * \left(\Phi_D'^k \cdot \frac{f}{h * u^k} \right) \right) u^k - \alpha \ \mathrm{div} \left(\Psi'(|\boldsymbol{\nabla} u^k|^2) \boldsymbol{\nabla} u^k \right) u^{k+\nu} = 0 \ . \tag{3.45}$$

Here, we use the abbreviation $\Phi_D'^k := \chi_D \cdot \Phi'(r_f(\mathcal{F}[u^k]))$ and set $\nu := 1$, if the divergence term is negative, and $\nu := 0$ else. Solving for u^{k+1} leads to the iteration scheme

$$u^{k+1} = \frac{h^* * \left(\Phi_D'^k \cdot \frac{f}{h * u^k} \right) + \alpha \ [\mathrm{div} \left(\Psi'(|\boldsymbol{\nabla} u^k|^2) \ \boldsymbol{\nabla} u^k \right)]_+}{h^* * \Phi_D'^k - \alpha \ [\mathrm{div} \left(\Psi'(|\boldsymbol{\nabla} u^k|^2) \ \boldsymbol{\nabla} u^k \right)]_-} \cdot u^k \ , \tag{3.46}$$

where we adopt the notation $[z]_\pm := \frac{1}{2}(z \pm |z|)$ from Welk [2010].

Although this case distinction ensures the non-negativity of the overall scheme, our experiments in Section 3.6 (cf. Figure 3.8) show that it still behaves unsatisfactory when high levels of regularisation are used.

This motivates us to a novel fixed point iteration with better stability properties: Instead of evaluating the divergence expression completely at the old time step, we propose a semi-implicit realisation. This means, within the divergence term, we propose to apply the diffusivity function (and diffusion tensor respectively) to u^k, i.e. at the old time step. However, the gradient $\boldsymbol{\nabla} u$ is taken from the new time step $k+1$. At the end, the diffusion term is multiplied with u^k. The other terms are untouched and the remaining nonlinear parts are still solved by the lagged diffusivity or Kačanov-method [Fučik et al., 1973; Chan and Mulet, 1999; Vogel, 2002], i.e. at the old iteration step k. Applying this novel fixed point strategy to Equation (3.42), the more general anisotropic model can be written as

$$(h^* * \Phi_D'^k)\; u^{k+1} - \left(h^* * \left(\Phi_D'^k\; \frac{f}{h * u^k}\right)\right) u^k - \alpha\; \mathrm{div}\left(\boldsymbol{D}(\boldsymbol{J}_\rho(\boldsymbol{\nabla} u_\sigma^k))\; \boldsymbol{\nabla} u^{k+1}\right) u^k = 0\;. \tag{3.47}$$

This scheme already reaches a higher stability w.r.t. large amount of regularisation. However, in each iteration, three convolution operations have to be performed. Since each convolution is computationally expensive, one should try to minimise their occurrence. To this end, we propose to replace the factor u^k in the first subtrahend by u^{k+1}, leading to

$$(h^* * \Phi_D'^k)\; u^{k+1} - \left(h^* * \left(\Phi_D'^k\; \frac{f}{h * u^k}\right)\right) u^{k+1} - \alpha\; \mathrm{div}\left(\boldsymbol{D}(\boldsymbol{J}_\rho(\boldsymbol{\nabla} u_\sigma^k))\; \boldsymbol{\nabla} u^{k+1}\right) u^k = 0\;, \tag{3.48}$$

where the distributivity property of convolution, i.e. $h*f + h*g = h*(f+g)$, allows us to reformulate it as

$$\left(h^* * \left(\Phi_D'^k \cdot \left(1 - \frac{f}{h * u^k}\right)\right)\right) u^{k+1} - \alpha\; \mathrm{div}\left(\boldsymbol{D}(\boldsymbol{J}_\rho(\boldsymbol{\nabla} u_\sigma^k))\; \boldsymbol{\nabla} u^{k+1}\right) u^k = 0\;. \tag{3.49}$$

On the one hand, as we can see, this strategy reduces the number of necessary convolutions per iteration by one. On the other hand, our scheme becomes even more semi-implicit.

The actual solution of the latter equation is carried out using the steepest descent method (see e.g. [Luenberger and Yinyu, 2015]), to which end we

introduce the relaxation parameter τ:

$$-\left(h^* * \left(\Phi_D'^k \cdot \left(1 - \frac{f}{h * u^k}\right)\right)\right) u^{k+1} + \alpha \operatorname{div}(\boldsymbol{D}(\boldsymbol{J}_\rho(\boldsymbol{\nabla} u_\sigma^k))\ \boldsymbol{\nabla} u^{k+1})\ u^k = \frac{u^{k+1} - u^k}{\tau}\ . \tag{3.50}$$

To apply this fixed point iteration to 3-D digital images, we discuss in Section 3.5 how this scheme above can be discretised.

A Semi-implicit Relaxation Scheme for RL Deconvolution

Of course, the latter semi-implicit iteration scheme can also be transferred to the original Richardson-Lucy deconvolution approach just by omitting interpolation and setting $\Phi(s) = s$, $\alpha = 0$. Hence, we proceed analogously to Section 3.4, and consider the fixed point iteration (3.49) for $\Phi_D'^k = 1$ and $\alpha = 0$. This way, Equation (3.49) comes down to

$$\left(h^* * \left(1 - \frac{f}{h * u^k}\right)\right) \cdot u^{k+1} = 0\ . \tag{3.51}$$

After that, we apply again the steepest descent method with relaxation parameter τ. This leads us to

$$-\left(h^* * \left(1 - \frac{f}{h * u^k}\right)\right) \cdot u^{k+1} = \frac{u^{k+1} - u^k}{\tau}\ , \tag{3.52}$$

which is a semi-implicit version of the scheme of Holmes and Liu [1991]. Finally, we solve for u^{k+1} and obtain

$$u^{k+1} = \left(1 + \tau\left(1 - h^* * \frac{f}{h * u^k}\right)\right)^{-1} \cdot u^k\ . \tag{3.53}$$

As for the standard RL scheme (3.26), if noise is negligible, i.e. $f = h * g$, the undisturbed signal g is a fixed point of the latter scheme.

3.5 Space Discretisation

In this section, we want to discretise the IRARRL deconvolution scheme from Equation (3.50). To this end, according to the discretisation scheme in Section 2.1.1, we assume a 3-D cell image to be sampled on a regular grid of

size $N_{x_1} \times N_{x_2} \times N_{x_3} =: N$ with sampling distances $(h_{x_1}, h_{x_2}, h_{x_3})^\top =: \boldsymbol{h}_3$ in horizontal-, vertical- and depth-direction, respectively. Further, we denote by $u^k_{i,j,\ell}$ the approximation of u at voxel (i, j, ℓ) at evolving time $t = k \cdot \tau$. Hence, a discrete version of Equation (3.50) has to fulfil the following equation in each voxel $(i, j, \ell) \in \{1, \ldots, N_{x_1}\} \times \{1, \ldots, N_{x_2}\} \times \{1, \ldots, N_{x_3}\}$:

$$\underbrace{-\left[h^* * \left(\Phi'^k_D \cdot \left(1 - \frac{f}{h * u^k}\right)\right)\right]_{i,j,\ell}}_{\boldsymbol{D}_1} \cdot u^{k+1}_{i,j,\ell} + \alpha \cdot \underbrace{[\operatorname{div}(\boldsymbol{D}(\boldsymbol{J}_\rho(\boldsymbol{\nabla} u^k_\sigma))\boldsymbol{\nabla} u^{k+1})]_{i,j,\ell}}_{\boldsymbol{A}(\boldsymbol{u}^k)\boldsymbol{u}^{k+1}} \cdot \underbrace{u^k_{i,j,\ell}}_{\boldsymbol{D}_2} = \frac{u^{k+1}_{i,j,\ell} - u^k_{i,j,\ell}}{\tau} \quad . \tag{3.54}$$

At this point, let us again change to a single-index notation (e.g. row-major ordering) and arrange 3-D signals $u : \mathbb{R}^3 \to \mathbb{R}$ in vectors $\boldsymbol{u} \in \mathbb{R}^N$. Then, the diffusion term can be expressed in terms of the matrix-vector multiplication $\boldsymbol{A}(\boldsymbol{u}^k)\boldsymbol{u}^{k+1}$. This reminds us of the semi-implicit diffusion scheme of Equation (2.14). For the discretisation of the anisotropic diffusion process, we use a 3-D extension[3] of the scheme presented by Weickert et al. [2013]. While, in the isotropic case, $\boldsymbol{A}$ has a straightforward structure and can be described, e.g. by (2.10), the anisotropic case is much more complex. Therefore it is not shown explicitly here. Further, in order to formulate the point-wise multiplications in Equation (3.54), we introduce the $N \times N$ diagonal matrices

$$\boldsymbol{D}_1 := \operatorname{diag}\Bigg(-\left[h^* * \left(\Phi'^k_D \cdot \left(1 - \frac{f}{h * u^k}\right)\right)\right]_1, \ldots, \ldots, -\left[h^* * \left(\Phi'^k_D \cdot \left(1 - \frac{f}{h * u^k}\right)\right)\right]_N \Bigg) \tag{3.55}$$

and

$$\boldsymbol{D}_2 := \operatorname{diag}\left(u^k_1, \ldots, u^k_N\right) \ , \tag{3.56}$$

where the two convolution operations are realised by exploiting the convolution theorem (cf. Section 2.5.1) [Gasquet et al., 1998; Bracewell, 1999]. However, to reduce wraparound errors and to obtain a power-of-two image size that is well-suited for a *fast Fourier transformation (FFT)* application, we first mirror the image at the boundary according to Figure 3.6.

[3]Special thanks go to Prof. Dr. Joachim Weickert and Prof. Dr. Martin Welk for providing their implementation.

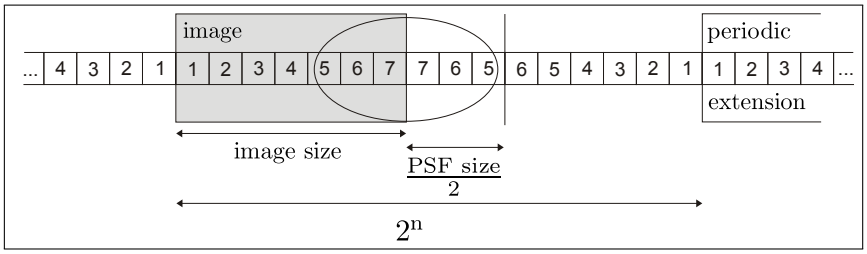

Figure 3.6: Boundary treatment to reduce wraparound errors, By mirroring the image depending on the size of the PSF support, we obtain a power-of-two image size. The remaining voxels are set as the mirror of the periodic image extension.

Eventually, Equation (3.54) can be written in the following vector-valued scheme:

$$\Big(\boldsymbol{D}_1(\boldsymbol{u}^k) + \alpha \cdot \boldsymbol{D}_2(\boldsymbol{u}^k) \cdot \boldsymbol{A}(\boldsymbol{u}^k)\Big) \cdot \boldsymbol{u}^{k+1} = \frac{\boldsymbol{u}^{k+1} - \boldsymbol{u}^k}{\tau} \ . \tag{3.57}$$

This leads us to a system of equations $\boldsymbol{\mathcal{B}}\boldsymbol{x} = \boldsymbol{b}$ to be solved in every iteration step:

$$\underbrace{\Big(\boldsymbol{I} - \tau(\boldsymbol{D}_1(\boldsymbol{u}^k) + \alpha \cdot \boldsymbol{D}_2(\boldsymbol{u}^k) \cdot \boldsymbol{A}(\boldsymbol{u}^k))\Big)}_{\boldsymbol{\mathcal{B}}} \cdot \underbrace{\boldsymbol{u}^{k+1}}_{\boldsymbol{x}} = \underbrace{\boldsymbol{u}^k}_{\boldsymbol{b}} \ . \tag{3.58}$$

Please note that due to the multiplication with diagonal matrices, the resulting system matrix $\boldsymbol{\mathcal{B}}$ is not symmetric.

Let us now illustrate how to solve this system of equations with the help of the *Jacobi relaxation* or *weighted Jacobi* method [Morton and Mayers, 2005].

The Jacobi Relaxation Method

In every iteration k of the presented approach, the system of equations in (3.58) has to be solved. To this end, let us split the system matrix $\boldsymbol{\mathcal{B}} = \boldsymbol{\mathcal{D}} - \boldsymbol{\mathcal{T}}$ into a diagonal component $\boldsymbol{\mathcal{D}}$ that is easily invertible and a remaining off-diagonal part $\boldsymbol{\mathcal{T}}$ by

$$\begin{aligned} \boldsymbol{\mathcal{D}} &= \boldsymbol{I} - \tau \cdot \Big(\boldsymbol{D}_1(\boldsymbol{u}^k) + \alpha \cdot \boldsymbol{D}_2(\boldsymbol{u}^k) \cdot \boldsymbol{A}_{\text{diag}}(\boldsymbol{u}^k)\Big) \ , \\ \boldsymbol{\mathcal{T}} &= \tau \cdot \alpha \cdot \boldsymbol{D}_2(\boldsymbol{u}^k) \cdot \boldsymbol{A}_{\text{rest}}(\boldsymbol{u}^k) \ , \end{aligned} \tag{3.59}$$

where $\boldsymbol{A}_{\text{diag}}$ and $\boldsymbol{A}_{\text{rest}}$ denotes the diagonal and off-diagonal part of $\boldsymbol{A}$ respectively. Then, following the Jacobi relaxation method [Morton and Mayers,

2005] with parameter $\omega > 0$ and iteration index m, the solution vector $\boldsymbol{x}$ can be iteratively determined by

$$\boldsymbol{x}^{m+1} = (1-\omega)\boldsymbol{x}^m + \omega\mathcal{D}^{-1}(\mathcal{T}\boldsymbol{x}^m + \boldsymbol{b}) \; . \tag{3.60}$$

In detail, we compute the solution $\boldsymbol{u}^{k+1} = (u_1^{k+1}, \ldots, u_N^{k+1})^\top$ for the discrete version of our IRARRL deconvolution scheme in every gradient descent step k via iterating the following scheme over m:

$$\begin{aligned} u_p^{k+1,m+1} &= (1-\omega) \cdot u_p^{k+1,m} \\ &+ \omega \cdot \frac{\left(1 + \tau \cdot \alpha \cdot \sum\limits_{\substack{q=1 \\ q \neq p}}^{N} a_{pq}^k \cdot u_q^{k+1,m}\right) \cdot u_p^k}{1 + \tau \left(\left[h^* * \left(\Phi_D'^k \cdot \left(1 - \frac{f}{h * u^k}\right)\right)\right]_p - \alpha \cdot \; u_p^k \cdot a_{pp}^k \right)} \; , \quad p = 1 \ldots N \; , \end{aligned} \tag{3.61}$$

where a_{pq}^k denotes an entry of the discretised diffusion matrix $\boldsymbol{A}(\boldsymbol{u}^k)$. Initialised with the observed image f, i.e. $\boldsymbol{u}^{0,0} := f$, after a fixed number of inner iterations, the outer iteration process is updated and the inner one starts again. This is continued until a certain stopping criterion is satisfied. In our work, we have considered the relative L_2-norm of the residuum of the system of equations.

Regarding the latter equation, we can recognise that during the inner iteration, the coupling of neighbours is accomplished only by regularisation. This is because the diffusion term is considered semi-implicitly. In contrast, $\boldsymbol{D}_1$ (cf. Equation (3.54)) is evaluated in a lagged diffusivity manner (Kačanov-method) [Fučik et al., 1973; Chan and Mulet, 1999; Vogel, 2002], i.e. considered completely at the old outer iteration step. Therefore, the coupling performed by the data term comes only into play when the outer iteration is updated, i.e. during the gradient descent process.

Now, it remains to prove that if $\boldsymbol{x}^m$ contains only physical plausible positive entries, the same holds for $\boldsymbol{x}^{m+1}$. To this end, let us first consider the factor $\boldsymbol{\mathcal{D}}^{-1}$ in Equation (3.60). Negative entries in this diagonal matrix can appear for large τ, since $\boldsymbol{D}_1$ can have entries of arbitrary sign. Hence, we derive the following bound on τ:

$$\begin{aligned} \forall\, p \in \{1, \ldots, N\} : \quad & \boldsymbol{\mathcal{D}}_{p,p} > 0 \\ & \Leftrightarrow \left[\boldsymbol{I} - \tau \left(\boldsymbol{D}_1 + \alpha \cdot \boldsymbol{D}_2 \cdot \boldsymbol{A}_{\text{diag}}\right)\right]_{p,p} > 0 \\ & \Leftrightarrow \left[-\boldsymbol{D}_1 - \alpha \cdot \boldsymbol{D}_2 \cdot \boldsymbol{A}_{\text{diag}}\right]_{p,p} > -\frac{1}{\tau} \end{aligned} \tag{3.62}$$

The latter inequality must hold for all $\alpha > 0$, and since $-\alpha \cdot \boldsymbol{D}_2 \cdot \boldsymbol{A}_{\text{diag}} > 0$, we obtain:

$$\min_p \left[-\boldsymbol{D}_1 - \alpha \cdot \boldsymbol{D}_2 \cdot \boldsymbol{A}_{\text{diag}}\right]_{p,p} > \min_p \left[-\boldsymbol{D}_1\right]_{p,p} > -\frac{1}{\tau} \,. \tag{3.63}$$

At this point we distinguish two cases: If $\min_p[-\boldsymbol{D}_1]_{p,p} \geq 0$, the condition is fulfilled for all $\tau > 0$ anyway. Otherwise, the theoretical step-size of the gradient descent is restricted by the inequality

$$\tau < \frac{-1}{\min_p \Big[-\boldsymbol{D}_1 \Big]_{p,p}} \,, \tag{3.64}$$

or expressed with the values of $\boldsymbol{D}_1$:

$$\tau < \frac{-1}{\min_p \left[h^* * \left(\Phi_D'^k \cdot \left(1 - \frac{f}{h * u^k}\right)\right)\right]_p} \,. \tag{3.65}$$

Please note that this bound is necessary to ensure the positivity of the solution.

Next, let us discuss the remaining component of Equation (3.60) that might become negative: Since the matrix $\boldsymbol{\mathcal{T}}$ contains the off-diagonal entries of the discrete anisotropic diffusion operator $\boldsymbol{A}$, the positivity of the term $\boldsymbol{\mathcal{T}}\boldsymbol{x}^m + \boldsymbol{b}$ cannot be guaranteed. Thus, we have to ensure for all $p \in \{1, \ldots, N\}$

$$(1-\omega)[\boldsymbol{x}^m]_p > -\omega \left[\boldsymbol{\mathcal{D}}^{-1}(\boldsymbol{\mathcal{T}}\boldsymbol{x}^m + \boldsymbol{b})\right]_p \,. \tag{3.66}$$

The latter inequality is only critical for voxels, where the right-hand side is positive. Hence, the following bound on the relaxation parameter can be established:

$$\omega \;<\; \min_p \frac{[\boldsymbol{x}^m]_p}{[\boldsymbol{x}^m]_p - \min\{\left[\boldsymbol{\mathcal{D}}^{-1}(\boldsymbol{\mathcal{T}}\boldsymbol{x}^m + \boldsymbol{b})\right]_p, 0\}} \,, \tag{3.67}$$

or with the components of the matrices

$$\omega \;<\; \min_p \frac{u_p^{k+1,m}}{u_p^{k+1,m} - \min\left\{ \dfrac{\left(1+\tau\cdot\alpha\cdot \sum\limits_{\substack{q=1 \\ q\neq p}}^{N} a_{pq}^k \cdot u_q^{k+1,m}\right)\cdot u_p^k}{1+\tau\left(\left[h^* * \left(\Phi_D'^k \cdot \left(1-\frac{f}{h*u^k}\right)\right)\right]_p - \alpha \cdot u_p^k \cdot a_{pp}^k\right)}, 0 \right\}} \,. \tag{3.68}$$

Note that in the isotropic case, the positivity of the off-diagonals of $\boldsymbol{A}$ is guaranteed anyway, so no relaxation is necessary ($\omega = 1$). Furthermore, the latter bound is derived for arbitrarily negative entries in $\boldsymbol{A}$. In practice, we can safely set ω in $[0.9, 1.0]$ without observing any experimental problems in practice.

Now, let us briefly come back to the semi-implicit relaxation scheme for RL deconvolution. To show that the equivalence for $\Phi_D'^k = 1$, and $\alpha = 0$ also holds in the discrete case, we plug these values into the iteration scheme (3.61) and obtain:

$$u_p^{k+1} = \frac{u_p^k}{1 + \tau \left(\left[h^* * \left(1 - \frac{f}{h * u^k} \right) \right]_p \right)} , \quad p = 1 \dots N . \tag{3.69}$$

Here, we can refrain from the Jacobi method since the off-diagonal part of the system matrix $\boldsymbol{\mathcal{B}}$ vanishes by setting $\alpha := 0$. As we can see, the iteration scheme above constitutes a discretised version of Equation (3.53). The derived condition for τ (Equation (3.64)) for preserving the positivity carries over as well. This accelerated Richardson-Lucy scheme has not been tested extensively, but first experiments demonstrate that half as many iterations are needed compared to standard Richardson-Lucy deconvolution.

3.6 Experiments

In this section, we demonstrate the performance of our reconstruction approach on a 3-D confocal laser scanning microscopy (CLSM) image as well as on a stimulated emission depletion (STED) microscopy image of cells. These data sets are provided by the Nano-Cell Interaction group of the Leibniz Institute for New Materials (INM) in Saarbrücken. As illustrated in Section 3.1.1 and 3.1.2, respectively, these data sets comprise degradations such as blur, Poisson noise and a relatively low axial resolution. To estimate the individual point-spread function (PSF) of each record, we proceed as explained in Section 3.1.3. For this purpose, to each cell image, a second 3-D record showing some small fluorescence beads has been provided. The image of the 3-D confocal microscope has a resolution of $1024 \times 1024 \times 50$ voxels with a grid size of (62/62/126) nm. A corresponding PSF is estimated at a grid size of (25/25/126) nm. Since the PSF must have the same grid size as the reconstruction, we have to resample it accordingly: For doubling the depth (axial) resolution and leaving the grid in x and y direction unaltered,

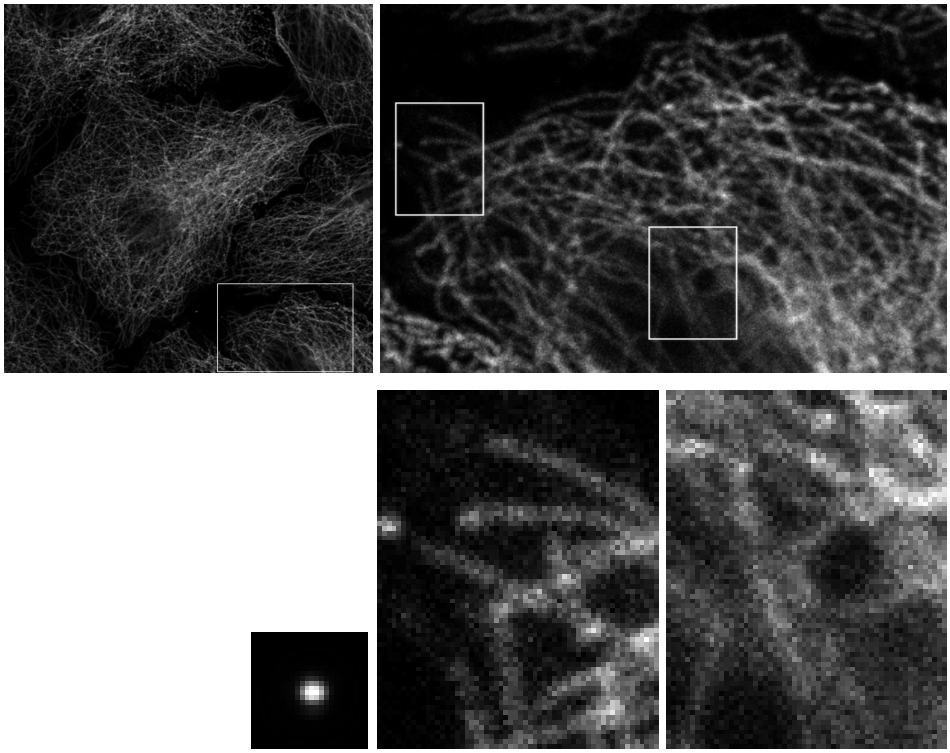

Figure 3.7: Slice 13 of the 3-D CLSM image (grey values rescaled to $[0, 255]$). **(a) Top left:** Slice of the complete volume of size $1024 \times 1024 \times 50$. The white rectangle indicates the origin of the first zoom-in. **(b) Top right:** Slice of the volume segment ($376 \times 244 \times 24$) of (a). Again, the white rectangles show the position of the second level magnifications. **(c) Bottom left:** Central slice of the estimated PSF ($24 \times 24 \times 33$). Its scale fits to the second magnification level. **(d) Bottom centre and (e) right:** Two second-level magnifications.

we need (62/62/63) nm. For this reason, the given PSF must be subsampled in the lateral direction and supersampled in axial direction. Since the PSF is a smooth function, we found simple linear interpolation to be sufficient for this task.

Although all our experiments are performed in 3-D, we depict distinctive 2-D slices of the processed volumes for the sake of a better recognisability. Figure 3.7 shows one such slice of the original CLSM image, along with two levels of magnification. For the rest of this section, we will discuss our results on the basis of these magnifications, since the fine-scale details are better comparable on high magnification levels.

Our first experiment addresses the stability of the semi-implicit scheme

(3.50) for the interesting case of relatively large α. For a direct comparison, we restrict our method to isotropic regularisation in this experiment. In the isotropic setting, typical choices for α are in the interval $[0.001, 0.05]$. Figure 3.8 compares our scheme (3.50) against the one of Elhayek et al. [2011] and Welk [2010] (see Equation (3.46)) respectively. While the numerical scheme of the latter method becomes unstable and exhibits chequerboard-like artefacts with increasing α, our semi-implicit scheme remains stable even for very high regularisation weights.

Before we demonstrate the performance of our anisotropic fibre enhancing approach, let us first illustrate the behaviour of the anisotropic smoothing strategy with regard to an underlying cell filament network. To this end, in our next experiment we superimpose a 2-D image slice with a visualisation of the main estimated smoothing orientations in Figure 3.9. This visualisation is computed by projecting the 3-D eigenvector of the structure tensor corresponding to the smallest eigenvalue into the 2-D slice. One can see that the anisotropic strategy yields good estimates for the sought smoothing directions and seems to be particularly applicable for such a task.

The third experiment compares our anisotropic fibre enhancement approach (3.50) with the original Richardson-Lucy method (2.52) [Richardson, 1972; Lucy, 1974]), as well as the isotropic (TV) regularisation strategy (3.46) of Elhayek et al. [2011]. To assess the interpolation quality of these techniques, we show in Figure 3.10 an in-between slice of the data set. There are a number of equivalences between wavelet-based approaches and techniques using variational principles or partial differential equations (such as TV approaches) [Steidl et al., 2004; Welk et al., 2008]. Since one can expect a similar quality as with the TV-like approach of Elhayek et al. [2011], we do not explicitly compare with wavelet-based methods here.

For the special case of Richardson-Lucy (RL) deconvolution [Richardson, 1972; Lucy, 1974], we have to fill in the missing slice by linear interpolation in advance. RL deconvolution results in very sharp contrasts after 200 iterations. However, due to absence of any regularisation, the main drawback of the RL algorithm is its sensitivity to noise and the deconvolution typical creation of over- and undershoots, i.e. oscillation artefacts [Bratsolis and Sigelle, 2001]. As a consequence, the deconvolved images have sharp contrasts, but the fibre structures are inhomogeneous and rough, particularly in fibre crossings. The regularisation component in the methods of Dey et al. [2004, 2006] and Elhayek et al. [2011] suppresses these over- and undershoots. With increasing degree of regularisation, the fibres become more and more smooth and homogeneous. However, at the same time, the radii of single fibres grow and fine structured details melt together, which is shown in Figure

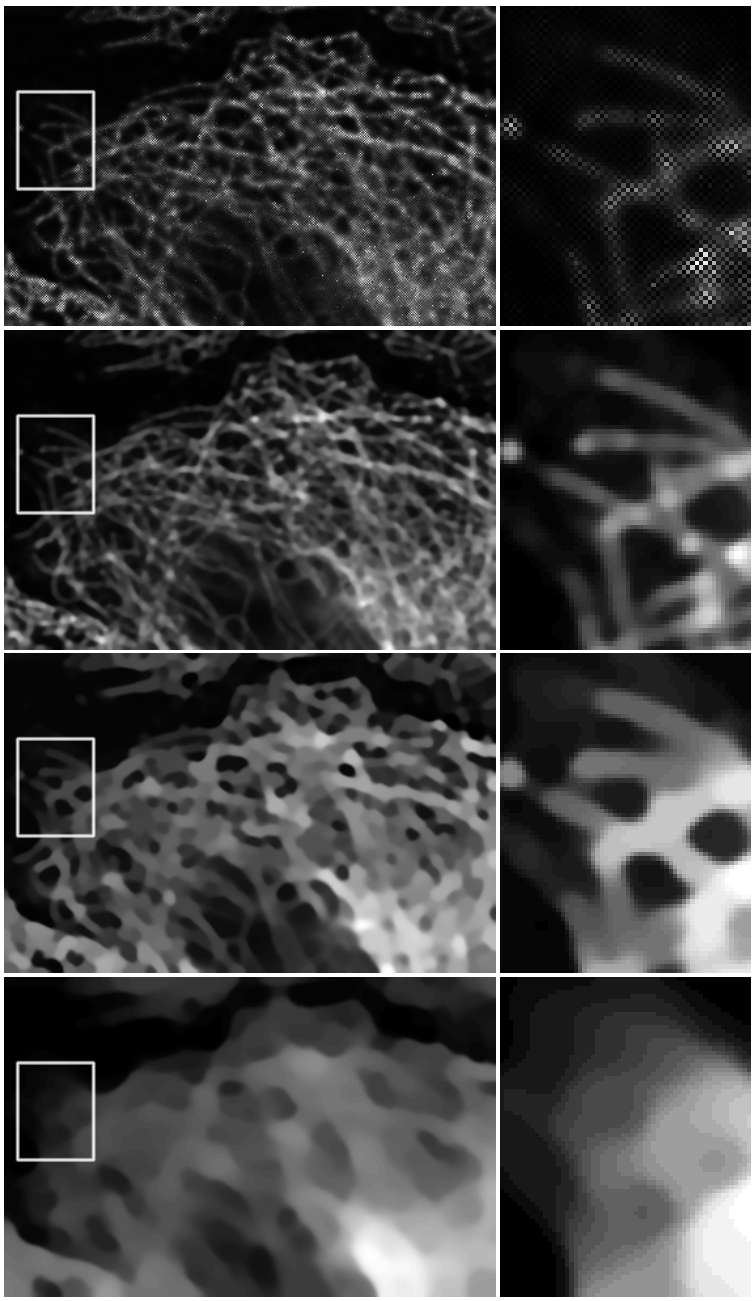

Figure 3.8: Numerical stability with respect to large amount of isotropic (TV) regularisation (rescaled to $[0, 255]$). *From Top to Bottom* **(a) Row 1:** Result of the method of Welk [2010] and Elhayek et al. [2011] respectively, $\alpha = 0.1$, 11 iterations. **(b) Row 2:** Semi-implicit approach with isotropic regularisation, $\alpha = 0.1$, 11 iterations. **(c) Row 3:** Same with 900 iterations and $\alpha = 0.1$. **(d) Row 4:** With 900 iterations and $\alpha = 0.5$.

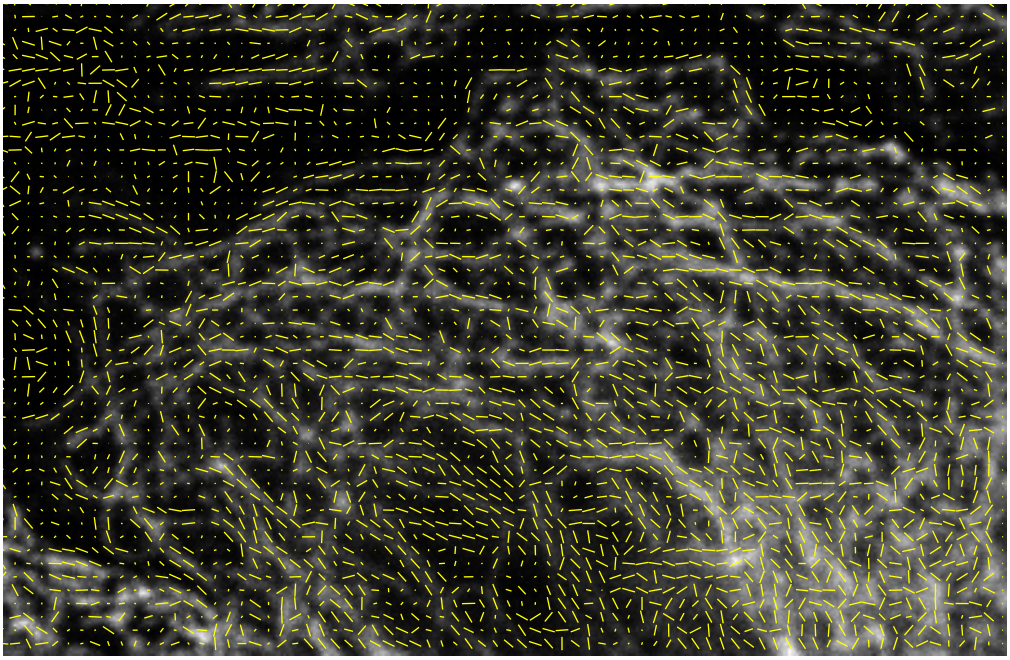

Figure 3.9: Orientations of the eigenvectors corresponding to the smallest eigenvalues of the structure tensor (3.40) projected into the actual slice.

3.10(b),(c) and Figure 3.11(b),(c), respectively. It partially results in small over-segmentation of the fibres. The benefit of our anisotropic regularisation becomes obvious by considering Figure 3.10(d) and Figure 3.11(d), respectively. As we can see, we obtain smoothness only along the fibres but sharp edges against the background.

Let us now come to the STED data set. The original record has a resolution of $1057 \times 1059 \times 65$, acquired with a sampling distance of (40/40/168) nm. The corresponding beads for the PSF estimation are recorded with a sampling distance of (40/40/84) nm, which is already the required grid size for doubling the z-resolution. In Figure 3.12, we illustrate the suitability of our IRARRL scheme for STED images. Moreover, we compare the robust data term against its non-robust variant. One can recognise that the robust data term results in a slightly more sharpened structure.

In the final experiment, we analyse the runtime of our scheme (3.50) and compare it against the numerical interpolation scheme of Elhayek et al. [2011]. To have a fair comparison, we compute an approximate steady state using 10,000 iterations of the scheme of Elhayek et al. [2011]. To reduce the experimentation time, we restrict our computations to a region of size $105 \times 89 \times 47$. In Figure 3.13 we show the input region along with the

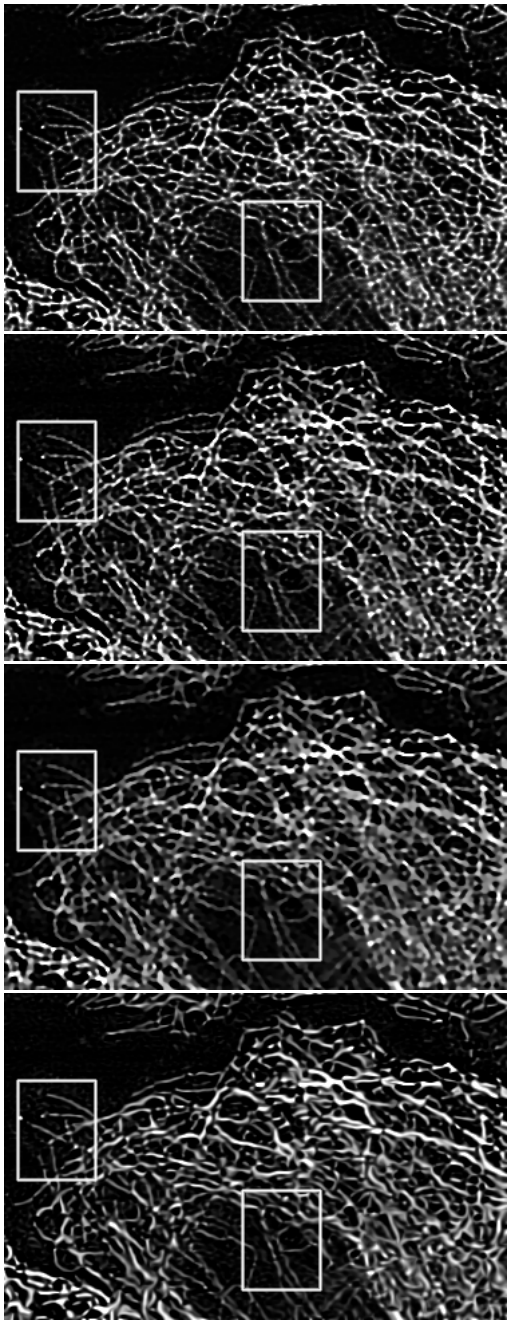

Figure 3.10: Reconstructed slice between slices 13 and 14 of the CLSM cell image (rescaled to $[0, 255]$). Zoom-ins of the white rectangle are shown in Fig. 3.11. *From top to bottom* **(a) Row 1:** RL deconvolution (200 iterations) with preceding linear interpolation. **(b) Row 2:** Isotropic (TV) method of Elhayek et al. [2011], $\alpha = 0.0006$. **(c) Row 3:** Ditto, $\alpha = 0.002$. **(d) Row 4:** Our anisotropic method (3.42), $\alpha = 0.0004$, $\rho = 1.5$, $\sigma = 0.6$, $\lambda = 0.1$.

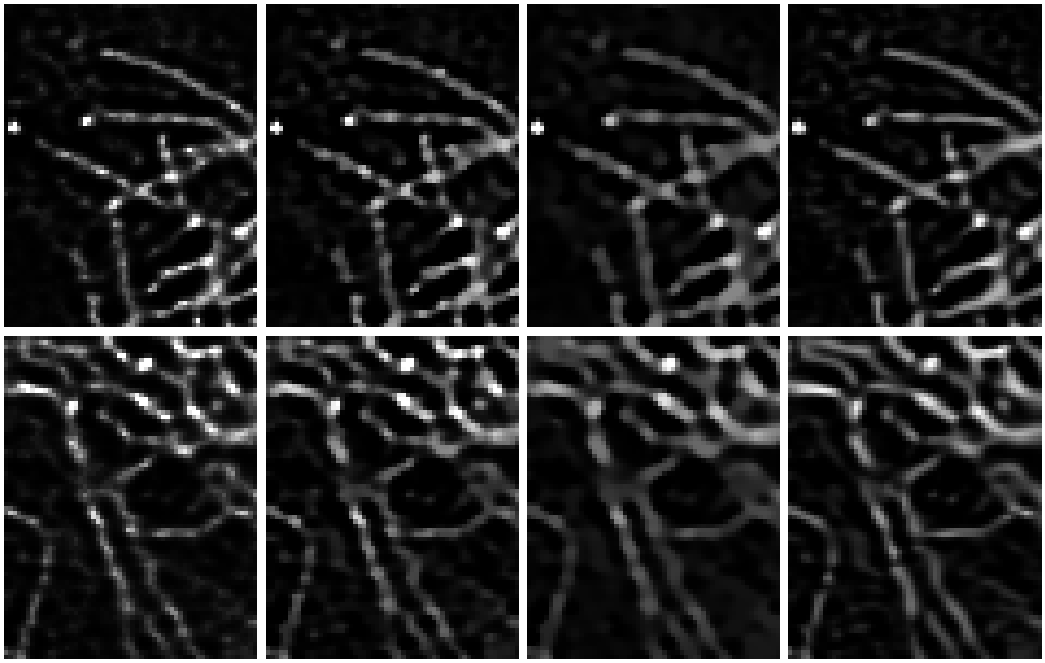

Figure 3.11: *From left to right* **(a) Column 1:** Zoom-in of Figure 3.10(a): RL deconvolution. **(b) Column 2:** Zoom-in of Fig. 3.10(b): Isotropic (TV) method, $\alpha = 0.0006$. **(c) Column 3:** Zoom-in of Fig. 3.10(c): Ditto with $\alpha = 0.002$. **(d) Column 4:** Zoom-in of Fig. 3.10(d): Our anisotropic method.

approximate solution. Next, we once more run the scheme of Elhayek et al. [2011] as well as our scheme (with a diffusion tensor that realises the same isotropic behaviour) and plot the average difference per voxel between the actual iterate and the approximate solution in Figure 3.14. One can clearly see the superior convergence rate of our method. Furthermore, we measure the time until the average difference per voxel drops below 0.1 and summarise these computation times in Table 3.1. All runtime experiments are performed with a C implementation on an Intel Xeon Processor W3565 (8M Cache, 3.20 GHz, 4.80 GT/s Intel) CPU with 24 GB RAM using a single-threaded implementation. The upper bound for the relaxation parameter $\tau = 1.5$ is determined experimentally. This speed-up can be explained by two facts: On the one hand, the novel semi-implicit scheme only performs two instead of three convolutions per iteration, which leads to a speed-up factor of 1.5. On the other hand, the number of required iterations is decreased by a factor of 2.2.

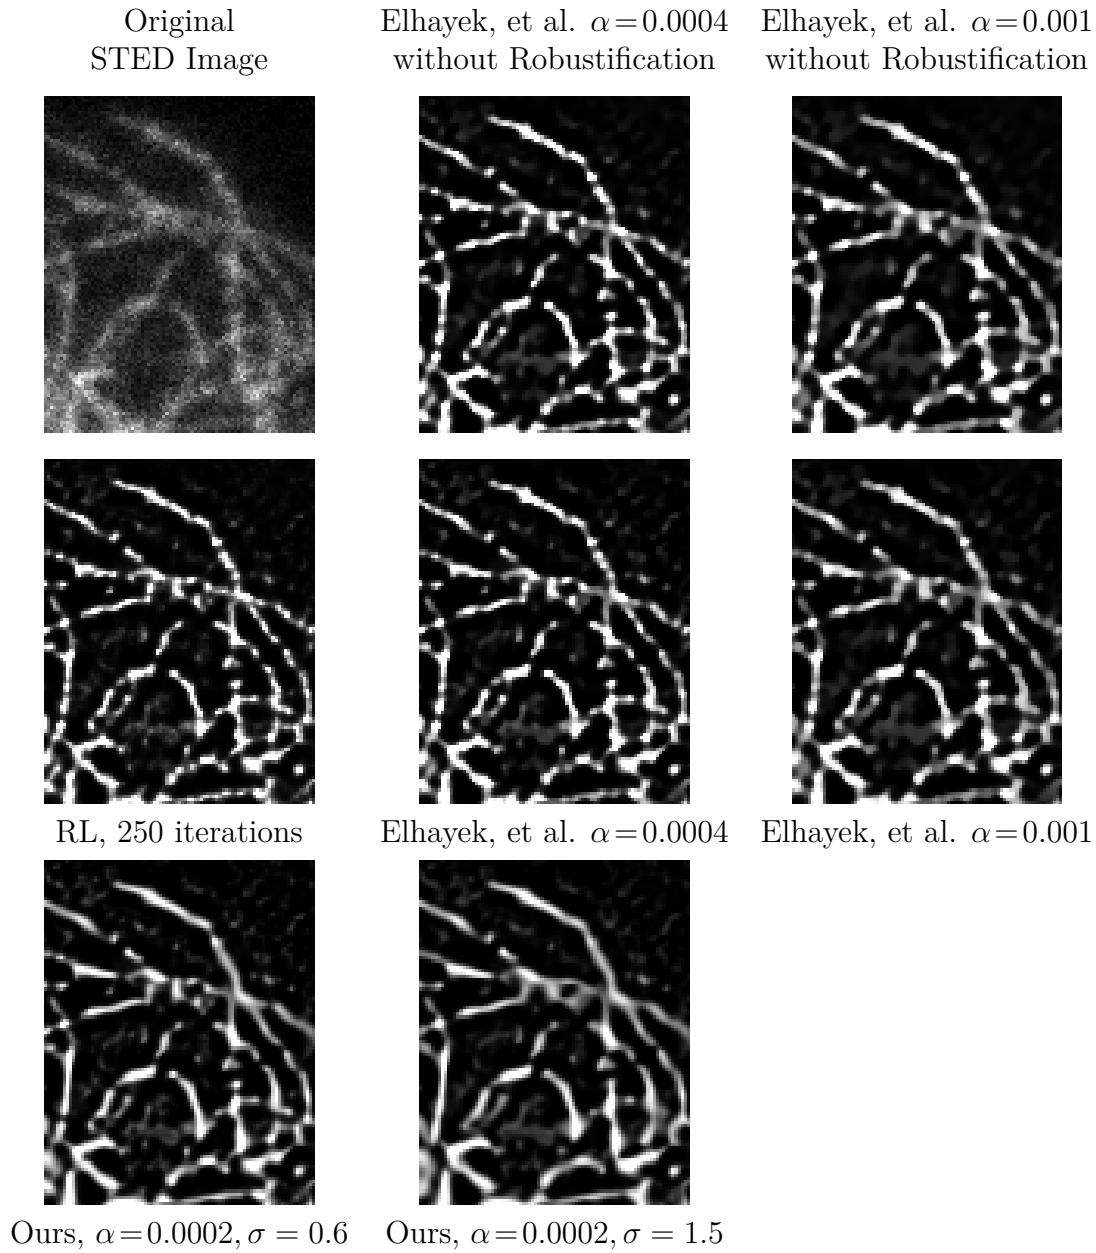

Figure 3.12: Reconstruction of cell fibres recorded with a STED microscope. Cut-outs (87×111 px) of an interpolated slice (Images are rescaled to $[0, 255]$). *In reading order:* **(a)** Original input slice. **(b)** Method of Elhayek et al. [2011], (TV) regularisaton $\alpha = 0.0004$, without robustification. **(c)** Ditto, $\alpha = 0.001$. **(d)** RL deconvolution (250 iterations). **(e)** Result of Elhayek et al. [2011], $\alpha = 0.0004$. **(f)** Ditto, $\alpha = 0.001$. **(g)** Our anisotropic scheme (3.42), $\alpha = 0.0002$, $\rho = 2.0$, $\sigma = 0.6$, $\lambda = 0.1$. **(h)** Ditto, $\sigma = 1.5$.

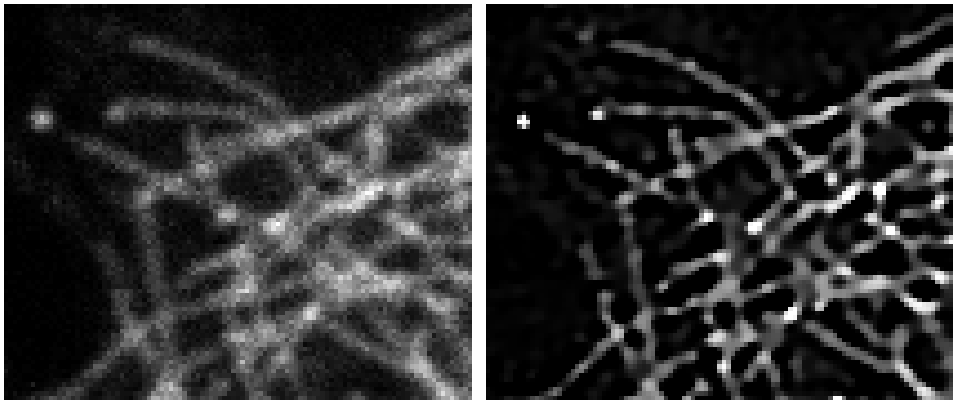

Figure 3.13: **(a) Left:** Original input CLSM region (rescaled to $[0, 255]$). **(b) Right:** Result after 10,000 iterations using the scheme of Elhayek et al. [2011] with $\alpha = 0.001$.

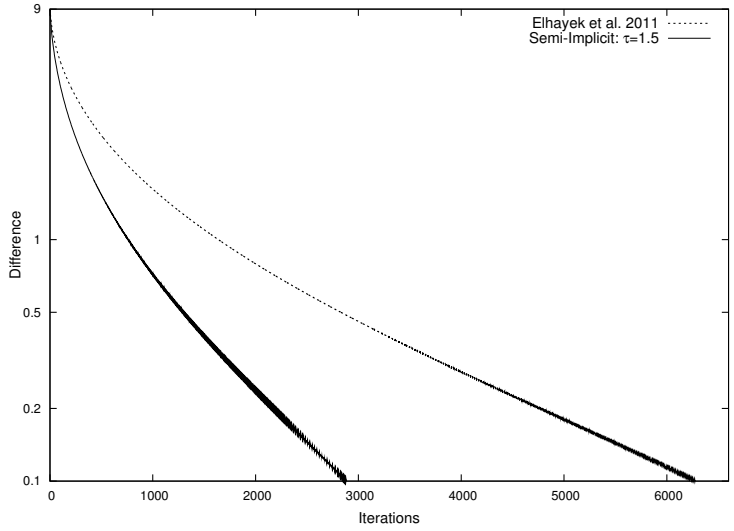

Figure 3.14: Average L_1 difference between the approximate solution and the iterated signal. Note the logarithmic scaling of the y-axis. **Dotted line:** Elhayek et al. [2011]. **Solid line:** Our semi-implicit iteration scheme, $\tau = 1.5$.

Table 3.1: Computation times of the approach of Elhayek et al. [2011] and our novel semi-implicit scheme (3.50).

Method	Iterations	Time	speed up factor
Elhayek et al. [2011]	6276	9827 s	1
Semi-implicit $\tau = 1.5$	2879	2982 s	3.3

3.7 Summary

In the beginning of this thesis, we revisited some established methods for image restoration. In this chapter, we have shown how to combine and extend these concepts in order to reach a tailoring towards the specific physical characteristics of modern low photon light microscope imagery. We have illustrated the development of a PDE-based method aiming especially at the restoration and enhancement of 3-D cell images recorded by CLSM and STED microscopes. To this end, we started with a detailed discussion of these microscopy techniques and their functional principles. Moreover, we have revealed their physical limitations and typical weaknesses: blur, Poisson noise and low axial resolution. To provide access to non-blind deconvolution techniques, we employed an estimation of the PSF of the participating optical system based on physical measurements. Concerning deblurring under Poisson distributed noise, we have shown that Richardson-Lucy deconvolution (RL) [Richardson, 1972; Lucy, 1974] is well suited for this task. By following the multiplicative Euler-Lagrange formalism, we have illustrated the connection between RL deconvolution and Csiszár's information divergence [Csiszár, 1991]. With this variational representation, we have shown the way to the robust and regularised RL approach of Welk [2010] and its extension to inpainting [Elhayek et al., 2011]. Eventually, we exploited the physical process of anisotropic diffusion [Weickert, 1998]. To this end, we replaced the scalar-valued diffusivity by a diffusion tensor. In this way, the smoothing behaviour is aligned to the tube-like structure of the cell fibres and the reconstruction quality is increased.

Compared to the work by Elhayek et al. [2011] and Dey et al. [2004, 2006], our novel proposed semi-implicit iteration scheme provides higher stability w.r.t. large regularisation weights. Besides that, we achieved a higher efficiency than Elhayek et al. [2011], because first, our novel scheme needs less iterations and second, it saves one out of three convolutions in each iteration step.

This chapter illustrates the way of tailoring a reconstruction framework

towards the specific physical characteristics and weaknesses of a considered image acquisition technique. Additionally, this chapter shows once more how one can get inspired by physical phenomena appearing in nature (here anisotropic diffusion) and how to adjust it according to specific requirements. Moreover, it demonstrates the benefits of an integrated view and the handling of all degradations simultaneously. By an adequate numerical algorithm, we combine stability and efficiency.

Chapter 4

Depth-from-Defocus

In this chapter, we address the so-called *depth-from-defocus* problem. Our goal is to exploit the physical properties of an optical system, which appear to be a weakness or limitation at first glance. More precisely, *depth-from-defocus* methods are based on the fact that a lens can only focus points at a certain distance. This distance is given by the *focal plane*. Points displaced from it are imaged in a blurred way whereas the amount of blur increases as the offset becomes larger. Although, strictly speaking, all points not lying within the focal plane are projected blurred, they still appear sharp, if they are within a certain distance range around the focal plane. This distance interval is called *depth-of-field* (DOF) [Träger, 2007; Wayne, 2013]. The position of the focal plane as well as the width of the depth-of-field depend on the individual optical properties of the camera. For instance, optical systems having a small *focal ratio* (also called *f-number* or *relative aperture*) – meaning the ratio of the *focal length* of the lens to the aperture diameter, see e.g. [Barsky et al., 2003] – suffer from a very limited depth-of-field [Langford, 2000; Stroebel, 1999; Wayne, 2013]. Such a situation is present in close-up and macro photography [Davies, 2012]. To imitate an acquisition as it would be done by such a limited DOF imagery, computer graphics methods simulate the local blur on the basis of the local depth information. This is called *depth-of-field simulation* and is usually applied in order to increase the realism of artificially created images [Pharr and Humphreys, 2004; Barsky and Kosloff, 2008].

In principle, depth-from-defocus is the inverse operation to the depth-of-field simulation. This means its goal is to reconstruct the depth map by means of estimating the local amount of out-of-focus blur. As a by-product, one is also able to compute the completely sharp image as it would be recorded by a camera having an infinite depth-of-field. However, regarding a single image, one can neither estimate the amount of blur nor distinguish

between out-of-focus blur and a blurred texture. Therefore, several images of the same static scene but varying focal settings are required. Then each of these images, usually arranged in a so-called *focal stack*, varies locally in their amount of out-of-focus blur (cf. Figure 4.1).

While the forward problem of simulating the depth-of-field effect always has a unique solution, depth-from-defocus is an ill-posed inverse problem: For instance, blurring a homogeneous region creates a focal stack whose images do not differ in the amount of blur. Thus, it is not possible to reconstruct a unique depth map. As in previous chapters, in order to handle such an ill-posed problem, we make use of variational methods and regularisation. This allows to extract a solution as the minimiser of some energy functional that involves an additional smoothness assumption [Bertero et al., 1988; Gelfand and Fomin, 2000; Aubert and Kornprobst, 2006]. Hence, the task turns into the design of a suitable variational functional to which end we first have to find an adequate forward operator. To keep the numerical complexity reasonable, many variational models for the depth-from-defocus problem involve only a relatively small set of simplified assumptions in their forward model, e.g. requiring local equifocality, i.e. a locally constant depth [Rajagopalan and Chaudhuri, 1997; Favaro and Soatto, 2000; Favaro et al., 2003a; Lin et al., 2015]. As a consequence, their solutions reflect the physical reality only to a very limited extent in the general case. Approaches that do not rely on any equifocal assumption at all can also violate important physical principles such as a maximum-minimum principle for the intensity values. For that purpose, the goal of Section 4.1 is to find a forward operator that approximates the depth-of-field effect. By refraining from any equifocal assumption, our operator comes closer to the physical reality. An additional requirement is that the forward operator should fit well into a variational framework. In Section 4.2 we then formulate our variational framework which enables us to invert this imaging model. In Section 4.2.2 we illustrate how to find a suitable minimiser of it. We further show how to handle multi-channel images in Section 4.2.3 and how to gain benefit from a robustification in Section 4.2.4. In Section 4.3 we present our novel *denoising and depth-from-defocus* approach in order to handle focal stacks suffering from severe noise. How to discretise the presented models is the topic of Section 4.4 and the following Section 4.5 demonstrates the performance of the different approaches on synthetic and real-world experiments.

This chapter is based on our conference publication [Persch et al., 2014] and its follow-up technical report [Persch et al., 2015] which has been submitted to the journal '*Image and Vision Computing*' and is still under review.

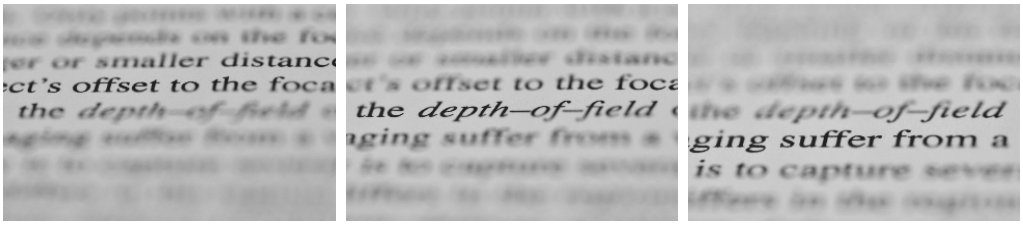

Figure 4.1: Focal stack. Each slice captures the same static scene, but differs in its focal settings. The gradual transition from sharp to blurred regions within each image corresponds to the depth profile.

Related Work. Generally, depth estimation based on differently focused images can be separated into two approaches: While *depth-from-defocus* methods also incorporate out-of-focus information, *depth-from-focus* methods estimate the depth by means of in-focus information only. To this end, a local sharpness criterion [Pertuz et al., 2013] such as the local variance of intensities [Sugimoto and Ichioka, 1985] is applied. Locally, the slice with maximal sharpness is assumed to be in focus and to match the depth. Boshtayeva et al. [2015] analyse different sharpness measures and combine them with anisotropic smoothing strategies.

Since the estimation of depth is usually combined with the recovery of the sharp pinhole image, *image fusion* methods can also be seen as related. Such methods fuse several defocused images to one sharp image. A simultaneous multi-focus image fusion and denoising approach is proposed by Ludusan and Lavialle [2012].

Inferring the depth based on the amount of out-of-focus blur goes back to Pentland [1987] as one of the pioneers in the field of depth-from-defocus. In his work, the blurriness of sharp edges or patches in different recordings serves for the depth estimation. To also incorporate diffraction effects, he discusses the use of a Gaussian instead of a pinhole as a suitable point-spread function (PSF). However, Pentland [1987] assumes one image to be in focus acting as a reference. This restriction is resolved in the work of Subbarao [1988]. He continues his work by suggesting a depth-from-defocus approach [Subbarao and Surya, 1994] using a convolution/deconvolution transform in the spatial domain and investigating its noise sensitivity [Subbarao and Tyan, 1997].

Rajagopalan and Chaudhuri [1997] approach the depth-from-defocus problem by means of a space-frequency representation technique based on the *Wigner distribution* and the *complex spectrogram*. Furthermore spatial regularisation of the local blur parameter is imposed.

The approach of Bailey et al. [2014] consists of two main steps. First, the

level of blurriness is determined using the method of Hu and de Haan [2007]. Next, to estimate the depth, the relation between blurriness, position of the focal plane and depth is exploited.

Making use of the equivalence of Gaussian blurring and linear homogeneous diffusion in the context of depth-from-defocus is proposed by Namboodiri and Chaudhuri [2004]. Since this implies a constant depth, images are partitioned into equifocal patches. In [Favaro et al., 2008; Namboodiri and Chaudhuri, 2007; Namboodiri et al., 2008; Wei et al., 2009] the idea is extended to nonlinear isotropic diffusion that allows spatial changes of the diffusivity corresponding to the depth profile. Favaro et al. [2003b] and Hong et al. [2009] extend this diffusion strategy by also involving directional information via anisotropic diffusion.

Markov random fields can also be employed to handle the depth-from-defocus problem [Chaudhuri and Rajagopalan, 1999; Bhasin and Chaudhuri, 2001; Namboodiri et al., 2008]. Among those approaches, Bhasin and Chaudhuri [2001] investigate how the PSF has to be iteratively corrected at strong depth changes, where partial occlusions occur. However, their work is limited to only two focal planes (two defocussed images).

Favaro [2010] and Wu et al. [2014] suggest to estimate the relative spread, i.e. the discrepancy in the standard deviations or width between the PSFs of two images in order to infer the depth values. This way, the computation of the sharp pinhole image can be avoided. While in the work of Favaro [2010] a nonlocal-means regularisation w.r.t. the depth map is proposed, Wu et al. [2014] considers geometric constraints to the relative spread.

Also restricting themselves to only two images, Ben-Ari and Raveh [2011] address the performance of depth-from-defocus. They suggest the use of *fast explicit diffusion* (FED) [Grewenig et al., 2010], as well as a *graphics processing unit* GPU implementation. Lin et al. [2015] propose a depth-from-defocus approach working with video pairs.

Most depth-from-defocus methods are restricted to imaging models based on geometric optics describing light propagation via straight lines. In the recent work of Wei et al. [2014] and Wei and Wu [2015], respectively, this restriction is removed by incorporating also *Fresnel* diffraction.

Modelling the image formation (forward operator), embedding it into a suitable energy functional, and posing depth-from-defocus as a variational minimisation problem, is most closely related to our work [Favaro and Soatto, 2000; Jin and Favaro, 2002; Aguet et al., 2008]. As a suitable discrepancy measure between the given data and the forward operator, the first two approaches suggest Csiszár's information divergence [Csiszár, 1991]. Aguet et al. [2008] penalise deviations from the model assumption in a quadratic way. Favaro and Soatto [2000] assumes a locally equifocal surface, which

implies a shift-invariant PSF. Instead, Jin and Favaro [2002] refrain from such a restriction and embed a shift-variant PSF in their model.

Interpreting a shift-invariant PSF as a 2-D function and a shift-variant PSF as a 4-D function, the approach of Aguet et al. [2008] can be seen as a compromise between both. They propose the use of a 3-D PSF whose slices are 2-D normalised Gaussians with varying standard deviation. Each slice represents the blur level according to a specific depth value. On the one hand, this strategy reduces the complexity by incorporating knowledge about the depth dependence on the PSF. On the other hand, especially at strong depth changes, it may result in a convolution operation with a non-normalised kernel. As a consequence, an essential physical property, namely the maximum-minimum principle w.r.t. the image intensities may be violated which leads to a wrong model assumption.

4.1 Image Formation Models

In this section, we want to approach the depth-from-defocus problem by means of inverting an approximation of the physical imaging process. Hence, as a first step, we have to find a forward operator that describes the imaging process reasonably well. Given the depth information of the scene and a completely sharp image, the sought forward operator should generate a result as it would be produced by a camera with a limited depth-of-field. In literature, in particular in the field of computer graphics, several suggestions already exist [Barsky and Kosloff, 2008]. The most famous model is the thin lens camera model which is often used in the context of raytracing to simulate the depth-of-field effect [Pharr and Humphreys, 2004; Cook et al., 1984]. There, the sharp information is given by the object texture. However, even though these approaches already produce very realistic results, their inversion in the sense of depth-from-defocus, is not possible or at least very difficult. Consequently, we have to develop a novel forward operator with a better trade-off between being invertible and simulating a real camera well.

To understand the relation between depth and out-of-focus blur, let us briefly discuss the pinhole camera model and the thin lens camera model in this section. While the first one produces completely sharp images without suffering from any depth-of-field effect, the second is a physical model to simulate such depth dependent out-of-focus blur [Pharr and Humphreys, 2004; Shirley and Morley, 2008]. Differences in their results allow to infer the amount of out-of-focus blur and thus the value of depth. Hence, the main idea is to relate both models. To this end, we express the thin lens camera model first by means of a spatially variant PSF applied to the pinhole

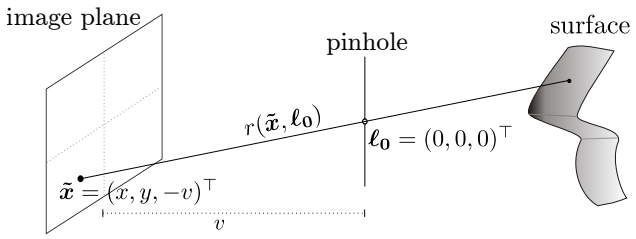

Figure 4.2: Pinhole camera model.

image. Next, we follow Aguet et al. [2008], and approximate the spatially variant PSF by a three-dimensional but spatially invariant one. Assuming local equifocal patches, i.e. patches that are parallel to the lens, such an approach may already yield good approximations. However, in case of strong depth changes, it leads to a convolution with a non-normalised kernel. Consequently, it violates the maximum-minimum principle of the intensity values at those places. Therefore, its use within a depth reconstruction framework is not advisable and motivates us to develop a novel forward operator which incorporates a normalisation function.

4.1.1 Pinhole Camera Model

The *pinhole camera model* (cf. Figure 4.2) is the standard imaging model. It refrains from any lenses and consists only of a very small pinhole and an image plane (or screen) [Pharr and Humphreys, 2004; Shirley and Morley, 2008]. The pinhole is placed in the *optical centre* $\boldsymbol{\ell}_0 := (0,0,0)^\top$ and has a distance $v \in \mathbb{R}_+$ to the image plane. The functional principles of the pinhole imaging model are based only on geometrical optics. This means, it describes the propagation of light only via rays, i.e. as linear subsets of $\mathbb{R}^3$. In this way, it completely ignores the wave character of light so that no diffraction phenomena come into play. For the sake of notational convenience, let us parametrise a ray by two points $\boldsymbol{a}, \boldsymbol{b} \in \mathbb{R}^3$ it passes through. Hence, we can fully describe a ray by

$$r(\boldsymbol{a}, \boldsymbol{b}) := \{\boldsymbol{y} \in \mathbb{R}^3 \mid \boldsymbol{y} = (1-\lambda) \cdot \boldsymbol{a} + \lambda \cdot \boldsymbol{b},\ \lambda \in \mathbb{R}\}. \tag{4.1}$$

The set of all rays is denoted by $\mathcal{R}$. If $\boldsymbol{x} := (x_1, x_2)^\top$ denotes the location within the image domain $\Omega_2 \subset \mathbb{R}^2$, then for each point $\tilde{\boldsymbol{x}} := (x_1, x_2, -v)^\top$ on the image plane, there exists exactly one optical ray $r(\tilde{\boldsymbol{x}}, \boldsymbol{\ell}_0) \in \mathcal{R}$ going through the pinhole. Let $d : \Omega_2 \to \mathbb{R}_+$ denote the depth map of a surface $\mathcal{S} \subset \mathbb{R}^3$, which we assume to be opaque. Then the resulting *pinhole operator*

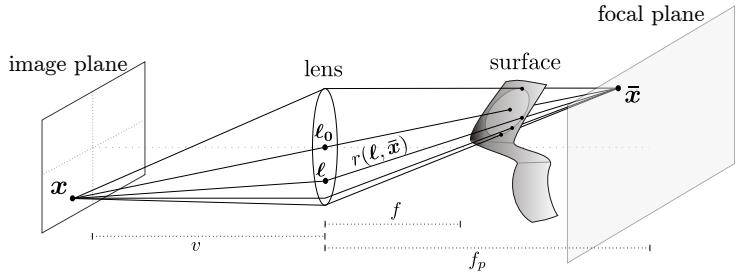

Figure 4.3: Thin lens camera model.

$\mathcal{F}_{\mathrm{P}}$ can be expressed as

$$\mathcal{F}_{\mathrm{P}}[\phi, d](\boldsymbol{x}) := \phi\Big(Z_d\Big(\tilde{\boldsymbol{x}}, \boldsymbol{\ell_0}\Big)\Big) \;, \tag{4.2}$$

where, $\phi : \mathcal{S} \to \mathbb{R}_+$ denotes the intensity value of a surface point and $Z_d(\tilde{\boldsymbol{x}}, \boldsymbol{\ell_0})$ yields the first (i.e. closest) intersection point of the ray $r(\tilde{\boldsymbol{x}}, \boldsymbol{\ell_0})$ with the surface $\mathcal{S}$. Since there exists at most one optical ray per image point hitting the surface, the object will be imaged completely sharp. Thus, no out-of-focus blur arises in this image formation model, and the depth-of-field is infinite.

4.1.2 Thin Lens Model

Figure 4.3 shows the thin lens camera model, which is a classical model to simulate the depth-of-field effect physically. In this model, instead of a pinhole, an infinitely thin (circular) lens with focal length f is placed in the optical centre $\boldsymbol{\ell}_0$. Lens and image plane are parallel and have a distance v to each other. As already mentioned in the beginning of this chapter, a lens can only focus points to the image plane that lie within the *focal plane*. To obtain the distance f_p of focal plane to the lens one can follow the thin lens equation (see e.g. [Born and Wolf, 1970]):

$$\frac{1}{f_p} = \frac{1}{f} - \frac{1}{v} \;. \tag{4.3}$$

Intersecting the *pinhole ray* $r(\tilde{\boldsymbol{x}}, \boldsymbol{\ell}_0)$ with the focal plane yields for each image point $\boldsymbol{x}$ its corresponding point $\bar{\boldsymbol{x}}$ within the focal plane. As the name of the ray already implies, this mapping exactly corresponds to the one of the pinhole camera model.

The more interesting case, of course, involves points that lie outside of the focal plane. They spread their intensity to a *circle of confusion* onto

the image plane [Horn, 1968; Subbarao, 1988; Barsky et al., 2003; Pharr and Humphreys, 2004]. In other words, if the object is not lying in the focal plane, the intensity of several surface points may blend into one single image point causing blurred information. These surface points lie within the intersection area of the bundle of lens rays with the surface. For an image point $\boldsymbol{x}$, this bundle can be described using $\bar{\boldsymbol{x}}$ and all points on the lens. Following Pharr and Humphreys [2004], one can formulate this imaging process via the *thin lens operator*

$$\mathcal{F}_{\mathrm{L}}[\phi, d](\boldsymbol{x}) := \frac{1}{|\mathcal{A}|} \int_{\mathcal{A}} \phi\Big(Z_d\Big(\boldsymbol{\ell}, \bar{\boldsymbol{x}}\Big)\Big) \, \mathrm{d}\boldsymbol{\ell} \,, \tag{4.4}$$

where $|\mathcal{A}|$ denotes the area of the lens. For the limiting case that $|\mathcal{A}| \to 0$, only the pinhole ray $r(\tilde{\boldsymbol{x}}, \boldsymbol{\ell}_0)$ survives which again leads to the pinhole camera model with no out-of-focus blur.

The thin lens camera model also complies with geometric optics and can thus be simulated in a straightforward way using raytracing techniques [Cook et al., 1984]. However, since a large amount of blur requires processing a huge number of rays per pixel, raytracing is computationally very expensive. Furthermore, it is not well suited for our variational inversion strategy. Hence, the goal of the following sections is the development of a forward operator that approximates the thin lens camera model, but can additionally be embedded into a variational framework. As a first step in that direction, let us express the thin lens camera model with the help of a spatially variant point-spread function (PSF) and the completely sharp pinhole image.

4.1.3 Spatially Variant Point-Spread Function

In the thin lens camera model, the sharp information is given by the intensity values of the surface. However, we already know that in the pinhole camera model each surface point is represented by exactly one image point because there exists a sharp one-to-one mapping. Hence, instead of integrating over the lens, we can integrate over the sharp pinhole image u in order to express the thin lens camera model. To this end, though, one has to weight each pinhole image point correctly. This can be realised with the help of a spatially variant PSF $H_d : \Omega_2 \times \Omega_2 \to \mathbb{R}_{0+}$ that depends on the depth profile d. In the context of computer graphics, such a postprocessing depth-of-field approach has been proposed by Potmesil and Chakravarty [1981]. Accordingly, instead of pursuing the intersection of light rays, we can express the imaging process as

$$\mathcal{F}_{\mathrm{H}}[u, d](\boldsymbol{x}) := \int_{\Omega_2} H_d(\boldsymbol{x}, \boldsymbol{y}) \; u(\boldsymbol{y}) \; \mathrm{d}\boldsymbol{y} \,, \tag{4.5}$$

where $\boldsymbol{x}$ describes the location within the 2-D image plane.

One important property of the thin lens camera model is its preservation of a maximum-minimum principle w.r.t. ϕ. This means that the intensity value of an image point lies between the minimum and maximum intensity value of any underlying surface point:

$$\phi_{\min} \leq \mathcal{F}_{\mathrm{L}}[\phi, d](\boldsymbol{x}) \leq \phi_{\max} \ , \quad \forall \boldsymbol{x} \in \Omega_2 \ . \tag{4.6}$$

Consequently, in order to imitate the thin lens camera model, the used spatially variant PSF H_d has to preserve this property w.r.t. the intensity values of the sharp pinhole image u. Therefore, we have to guarantee that the PSF is normalised for each image point $\boldsymbol{x}$:

$$\int_{\Omega_2} H_d(\boldsymbol{x}, \boldsymbol{y}) \ \mathrm{d}\boldsymbol{y} = 1 \ , \quad \forall \boldsymbol{x} \in \Omega_2 \ . \tag{4.7}$$

As this imaging model constitutes a weighted average of the sharp image intensities, the main issue turns into the computation of the weights of the PSF H_d. A straightforward solution for that would be the use of raytracing techniques and to apply the thin lens camera model. However, this is exactly what we want to avoid. Thus, we have to find a more efficient way to approximate the weights of the PSF H_d.

4.1.4 Approximation of the PSF

In this section, we want to approximate the weights of the PSF in order to avoid the need of expensive raytracing techniques. To this end, let us take a closer look at the intensity distribution within the thin lens camera model. From Horn [1968] and Subbarao [1988] we know that the intensity emitted by a non-occluded surface point lying outside of the focal plane is (uniformly) spread to a circle – assuming a circular lens – of confusion on the image plane. The size of this circle of confusion depends on the distance d of the surface point to the lens. Since we know from previous sections that surface intensities can be represented by grey values of the pinhole image, the weights of the PSF in Equation (4.5) have to represent the intensity distribution caused by such circle of confusion.

For simplification, let us for the moment assume that the surface is equifocal. This means that it is aligned parallel to the lens so that the depth d is constant. Then the circle of confusion do not change for any surface point and one has to distinguish only two cases: (i) surface points whose circle of confusion overlaps the actual image point are weighted by the reciprocal of the circle area (assuming a uniform intensity distribution). (ii) surface

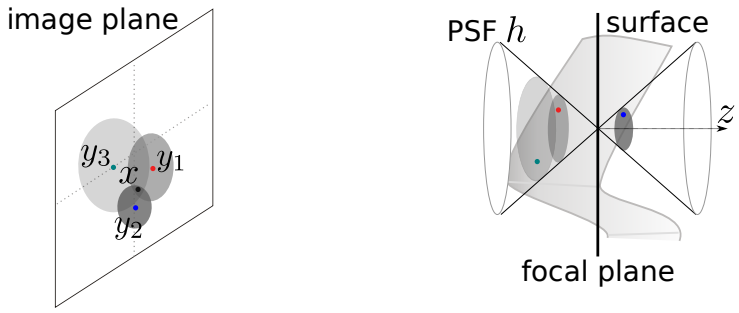

Figure 4.4: **Left:** Circles of confusion for different surface points appearing on the image plane. **Right:** 3-D PSF composed of 2-D normalised Gaussians.

points whose circle of confusion does not reach the actual image point are weighted by zero. This equifocal assumption, moreover, allows the simplification of Equation (4.5) to a convolution operation [Horn, 1968]. The spatially variant PSF H_d thus can be replaced by a spatially invariant kernel $h_d : \Omega_2 \subset \mathbb{R}^2 \to \mathbb{R}_{0+}$. The case distinction above can be realised with the help of a 2-D pillbox PSF whose radius is related to the constant depth.

However, in the general case of a non-constant depth map, the radius changes with the depth of each surface point. Thus, to estimate the intensity of an image point $\boldsymbol{x}$, Aguet et al. [2008] weight each neighbouring point $u(\boldsymbol{y})$ corresponding to its circle of confusion, where a point having a large circle of confusion will get a small weight and vice versa. To achieve this, they introduce a 3-D spatially invariant PSF $h : \Omega_3 \subset \mathbb{R}^3 \to \mathbb{R}_{0+}$ as an approximation of spatially variant H_d:

$$\mathcal{F}_{\mathrm{U}}[u, d](\boldsymbol{x}, z) := \int_{\Omega_2} \underbrace{h(\boldsymbol{x} - \boldsymbol{y}, z - d(\boldsymbol{y}))}_{\approx H_d(\boldsymbol{x}, \boldsymbol{y})} \, u(\boldsymbol{y}) \, \mathrm{d}\boldsymbol{y} \, , \tag{4.8}$$

where z represents a given focal plane. In this model, the weight not only depends on the distance of two points, but also on the actual depth value of the neighbouring points.

Until now, we have assumed that an out-of-focus point spreads its intensity uniformly to a circle of confusion. This results in PSF that is pillbox-shaped. However, to additionally take into account the wave character of light, Aguet et al. [2008] choose a 2-D Gaussian PSF instead of a pillbox one. This has been already proposed by Pentland [1987]. The standard deviation of the Gaussian replaces the radius of the pillbox. The approach of Aguet et al. [2008] can be illustrated as in Figure 4.4: First, the 2-D function h_d is lifted to a 3-D one h, composed of 2-D normalised Gaussians. The standard

deviation of each Gaussian increases with increasing distance to the focal plane. Next, one cuts a slice out of this 3-D PSF h corresponding to the depth. This slice serves as a local approximation of H_d for a specific image point $\boldsymbol{x}$. With the help of this 3-D PSF, a second interpretation of the equation above is also possible: For this, we assume that the sharp pinhole image u lies in a dark volume $g : \Omega_3 \to \mathbb{R}$ corresponding to the depth profile. It can be defined as

$$g(\boldsymbol{x}, z) := u(\boldsymbol{x}) \cdot \delta(z - d(\boldsymbol{x})) \quad \text{with} \quad \delta(x) := \begin{cases} 1 & \text{if } x = 0 \,, \\ 0 & \text{else} \,. \end{cases} \tag{4.9}$$

Then Equation (4.8) is just a standard 3-D convolution of g with the PSF h:

$$\begin{aligned} (g * h)(\boldsymbol{x}, z) &:= \int_{\Omega_3} g(\boldsymbol{y}, z') \cdot h(\boldsymbol{x} - \boldsymbol{y}, z - z') \, \mathrm{d}\boldsymbol{y} \, \mathrm{d}z' \\ &= \int_{\Omega_3} u(\boldsymbol{y}) \cdot \delta(z' - d(\boldsymbol{y})) \cdot h(\boldsymbol{x} - \boldsymbol{y}, z - z') \, \mathrm{d}\boldsymbol{y} \, \mathrm{d}z' \,. \end{aligned} \tag{4.10}$$

Here, the integrand vanishes everywhere except for $z' = d(\boldsymbol{y})$. Thus, the equation above can be simplified to

$$(g * h)(\boldsymbol{x}, z) = \int_{\Omega_2} u(\boldsymbol{y}) \cdot h(\boldsymbol{x} - \boldsymbol{y}, z - d(\boldsymbol{y})) \, \mathrm{d}\boldsymbol{y} \,, \tag{4.11}$$

where the integration domain has changed from Ω_3 to Ω_2. This illustrates that the forward operator of Aguet et al. [2008] performs a spatially invariant 3-D convolution.

Although the forward operator in (4.8) uses a 3-D PSF which is composed of 2-D normalised Gaussians, one has to keep in mind that the slice to be cut out – as local approximation of H_d – is not necessarily normalised. This would violate the requirement in Equation (4.7). Therefore, in the next section, we suggest a novel forward operator that counters exactly this problem.

4.1.5 Our Modification

In the last section, we have illustrated how to approximate locally the weights of H_d by following Aguet et al. [2008]. The idea is to compose a 3-D PSF of 2-D Gaussians varying in their standard deviation and taking a slice out of it corresponding to the depth profile. Assuming the equifocal case, this slice is just a 2-D normalised Gaussian of a certain standard deviation again. Accordingly, the requirement in Equation (4.7) is fulfilled, the maximum-minimum principle is automatically guaranteed, and the imaging model is

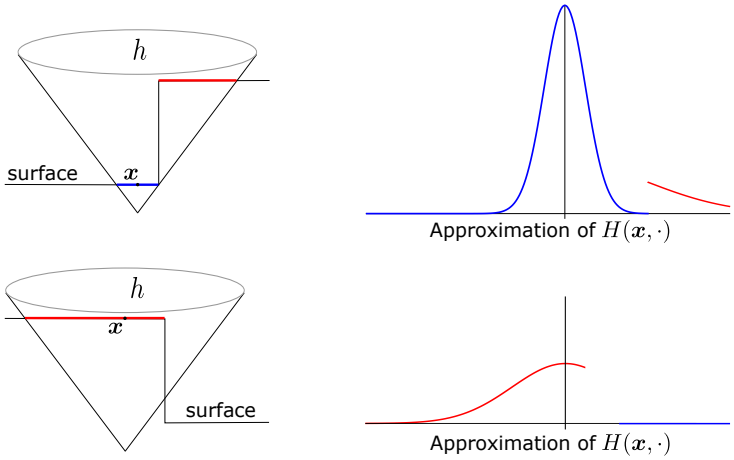

Figure 4.5: Unnormalised kernel. In the presence of strong depth changes, the local approximation of H_d by Aguet et al. [2008] is composed of different Gaussians (blue and red). **(a) Top row:** Overshoot: While the blue part is nearly a complete Gaussian, with the second (red) part an integration weight of 1 is exceeded. **(b) Bottom row:** Undershoot: The local composed PSF consists only of a part of a normalised Gaussian (red) and a second (blue) Gaussian already reaching negligible values. The resulting integration weight becomes smaller than 1.

approximated. However, the formulation above becomes problematic if partial occlusions occur, which is expected to happen due to depth changes. In this case, the slice is a composition of weights of several Gaussians with different standard deviations. As a direct consequence the normalisation cannot be guaranteed anymore. The forward operator then effectively performs spatially invariant 2-D integration with an unnormalised kernel as a local approximation of H_d (see Figure 4.5). Hence, the requirement in Equation (4.7) is not fulfilled. This leads to a violation of the maximum-minimum principle w.r.t. the image intensities (4.6) and thus to a violation of the imaging model.

To avoid this, the idea is to consider the local approximation of H_d and to use it as a normalisation function. That way, our novel forward operator reads

$$\mathcal{F}_{\mathrm{N}}[u, d](\boldsymbol{x}, z) := \frac{\mathcal{F}_{\mathrm{U}}[u, d](\boldsymbol{x}, z)}{\int_{\Omega_2} h(\boldsymbol{x} - \boldsymbol{x'}, z - d(\boldsymbol{x'}))\, \mathrm{d}\boldsymbol{x'}} \,. \tag{4.12}$$

This forward operator guarantees the requirement in Equation (4.7) per definition and consequently preserves the maximum-minimum principle of Equation (4.6).

Although this normalisation function may look like a small modification

at first glance, it can have a large impact on the simulation quality of the depth-of-field effect. To demonstrate this, we compare different forward operators in Figure 4.6. As one can see, the result of the forward operator of Aguet et al. [2008] shown in Figure 4.6(c) is very close to simple 3-D convolution (cf. Equation (4.10)) given in Figure 4.6(b). Differences are caused due to the discretisation required when embedding the pinhole image into the discrete 3-D volume according to Equation (4.9). Since both forward models perform convolutions with an unnormalised kernel h on strong depth changes, they produce bright overshoots followed by dark shadows. This behaviour is illustrated in Figure 4.5. In Figure 4.5(a) the slice taken out of the PSF consists of nearly a complete Gaussian caused by the blue surface and large part of a second Gaussian caused by the red surface. Hence, the integration weight of the slice exceeds 1 which leads to overshoots. In Figure 4.5(b) the slice only consists of approximately three-quarter of a Gaussian caused by the red surface. The rest is set to 0. Thus, the integration weight is lower than 1 which results in undershoots. Indeed, regarding Figure 4.6(c) and 4.6(b) respectively these local violations of the maximum-minimum principle on strong depth changes are violations of the model assumption and produce results that are not photorealistic. In contrast, comparing Figure 4.6(d) and (e) demonstrates that our normalised approach comes very close to the physically well-founded thin lens camera model, and allows to create realistic depth-of-field effects.

However, we have to keep in mind that the proposed normalisation is not a constant but a function depending on the position and on the depth profile of the scene. This make the model and the mathematical equations more complicated. Hence, the topic of the next sections is the embedding of our novel forward operator into a variational framework and how to find a suitable minimiser of it.

4.2 Variational Formulation

In the last section, we have proposed a novel forward operator that approximates the thin lens camera model simulating the depth-of-field effect. Now, as the next step of our reconstruction framework, we illustrate how to invert this novel forward operator within a variational framework. Given a stack of blurred images, we want to jointly estimate the depth map and the sharp pinhole image.

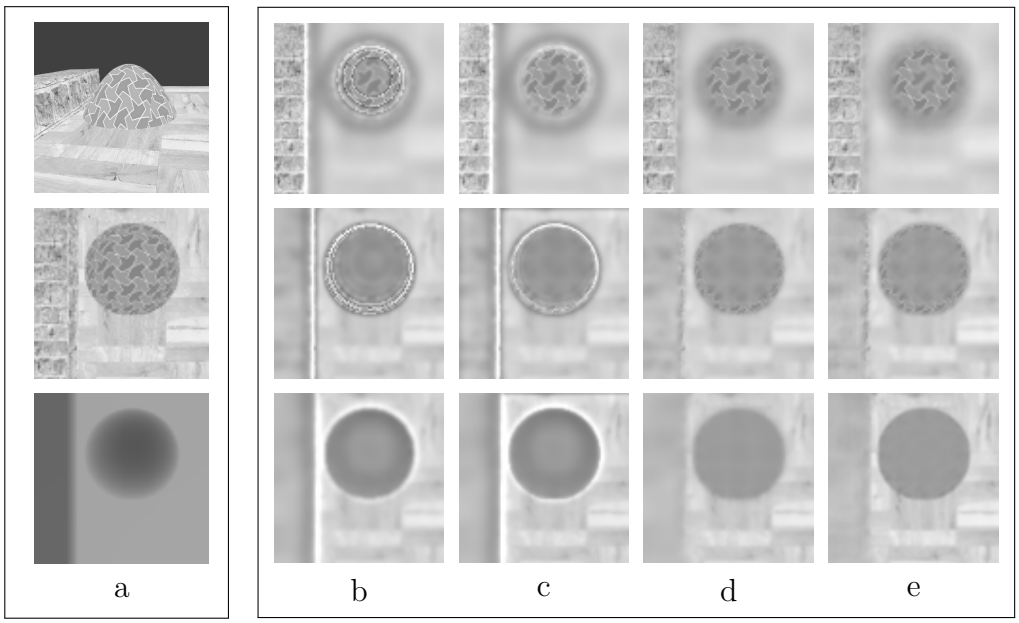

Figure 4.6: **Left box: (a) Top:** Synthetic 3-D test model. **(a) Middle:** Sharp image obtained with a pinhole camera renderer. **(a) Bottom:** Corresponding grey-valued coded depth map (the brighter the larger the depth value). **Right box:** Comparison of different forward operators. *From top to bottom:* Increasing focal plane distance. **(b)** Standard spatially invariant 3-D convolution (4.10). **(c)** Forward operator of Aguet et al. [2008] without normalisation preservation. **(d)** Our normalised forward operator (4.12). **(e)** Thin lens camera model (4.4) realised by raytracing technique.

4.2.1 Variational Model

As already mentioned, our depth-from-defocus framework needs a set of 2-D images capturing the same static scene but varying in their focal settings. Thus, let us assume such a set being arranged as a 3-D focal stack $s : \Omega_3 \to \mathbb{R}_+$ where $\Omega_3 \subset \mathbb{R}^3$ denotes the stack volume. In contrast to the last chapter, our forward operator is now applied to two arguments instead of one, namely the depth map $d : \Omega_2 \to \mathbb{R}_+$ of the scene and the sharp image $u : \Omega_2 \to \mathbb{R}_+$ as it would be recorded by a pinhole camera. Moreover, in this chapter we do not assume the input data to be acquired by a low intensity image acquisition technique such as CLSM or STED microscopy. Hence, it is no longer necessary to assume Poisson distributed noise. Thus, we can refrain from Csiszár's information divergence [Csiszár, 1991]. Instead, we demand similarity between the recorded focal stack and our forward operator $\mathcal{F}_{\mathrm{N}}$

applied to u and d by considering the *residual error*

$$r[u,d](\boldsymbol{x},z) = s(\boldsymbol{x},z) - \mathcal{F}_{\mathrm{N}}[u,d](\boldsymbol{x},z)\ . \tag{4.13}$$

Then, we follow the most common choice in variational calculus and penalise the residual error in a quadratic way (least-squares sense). This is especially suitable if Gaussian distributed noise is involved and leads us to the data term

$$E_{\mathsf{D}}(s,u,d,\mathcal{F}_{\mathrm{N}}) := \int_{\Omega_3} \Big(r[u,d]\Big)^2 \mathrm{d}\boldsymbol{x}\ \mathrm{d}z\ . \tag{4.14}$$

Please note here that given a depth map and a completely sharp image as arguments, $\mathcal{F}_{\mathrm{N}}$ produces only one 2-D image of a certain focal plane z. Only by integrating over z, we enforce similarity to each of the given slices of the focal stack. Exactly like $\mathcal{F}_{\mathrm{U}}$, our forward operator $\mathcal{F}_{\mathrm{N}}$ from Equation (4.12) is linear in u but nonlinear in d, and the data term is convex in u but nonconvex in d.

A minimiser of the data term alone is not unique. On the one hand, this is due to the involved convolution operation as explained in the previous chapter (cf. Equation (3.11)). On the other hand, regarding a homogeneous region where u is constant, it is not possible to infer any amount of blur. Accordingly, in such regions, the depth d cannot be determined. To avoid such ambiguities and to cope with the problem of ill-posedness, we supplement our variational approach by a regularisation term E_{S} [Bertero et al., 1988; Gelfand and Fomin, 2000; Aubert and Kornprobst, 2006]. Since we want to reconstruct a completely sharp image u, here, we restrict the (piecewise) smoothness assumption to the sought depth map and refrain from any smoothness on u. This leads to the functional

$$E(u,d) = E_{\mathsf{D}}(s,u,d,\mathcal{F}_{\mathrm{N}}) + \alpha \cdot E_{\mathsf{S}}(d)\ . \tag{4.15}$$

Although the integration domain of the data term is Ω_3, the depth map constitutes a 2-D signal. This implies an integration over Ω_2. Hence, we can achieve a penalisation of large gradient magnitudes of the depth profile by employing the smoothness term

$$E_{\mathsf{S}}(d) := \int_{\Omega_2} \Psi(|\boldsymbol{\nabla} d|^2)\ \mathrm{d}\boldsymbol{x}\ . \tag{4.16}$$

As usual, $\Psi : \mathbb{R} \to \mathbb{R}_+$ denotes a positive increasing function. Our experiments in Section 4.5 are obtained with the Whittaker-Tikhonov [Tikhonov, 1963; Whittaker, 1923] penaliser $\Psi(x^2) = x^2$.

4.2.2 Minimisation

Euler-Lagrange Equations

In previous chapters, we have already illustrated the way of finding a minimiser of a variational energy by means of the Euler-Lagrange formalism. No matter whether we follow the additive or multiplicative Euler-Lagrange formalism, a minimiser has to fulfil the Euler-Lagrange equation. So far, however, the considered approaches only depend on one unknown leading to one Euler-Lagrange equation. In this chapter, we now want to estimate jointly the depth map d along with the sharp pinhole image u. Hence, our variational framework depends on two unknown functions and we have to consider the two Euler-Lagrange equations

$$\frac{\delta E}{\delta u} = 0 \quad \text{and} \quad \frac{\delta E}{\delta d} = 0 \,. \tag{4.17}$$

To derive the variational gradient $\frac{\delta E}{\delta u}$ w.r.t. the sharp pinhole image u as well as the variational gradient $\frac{\delta E}{\delta d}$ w.r.t. the depth map d, we can again follow classical additive Euler-Lagrange formalism [Gelfand and Fomin, 2000] (cf. Section 2.2.2). Since the novelties take place only within the data term E_{D}, we can leave out the smoothness term E_{S} for a moment. Regarding the variational gradient $\frac{\delta E}{\delta u}$ no smoothness assumption is affected anyway. Hence, let us consider an additive perturbation with a test function v in E_{D} according to Equation (2.27), first w.r.t. u:

$$\begin{aligned}
&\frac{\partial}{\partial \epsilon} E_{\mathsf{D}}(s, u + \epsilon v, d, \mathcal{F}_{\mathrm{N}}) \\
&= \frac{\partial}{\partial \epsilon} \int_{\Omega_3} \Big(s(\boldsymbol{x}, z) - \mathcal{F}_{\mathrm{N}}[u + \epsilon v, d](\boldsymbol{x}, z) \Big)^2 \, \mathrm{d}\boldsymbol{x} \, \mathrm{d}z \\
&= \frac{\partial}{\partial \epsilon} \int_{\Omega_3} \left(s(\boldsymbol{x}, z) - \int_{\Omega_2} \frac{(u + \epsilon v)(\boldsymbol{y}) \cdot h(\boldsymbol{x} - \boldsymbol{y}, z - d(\boldsymbol{y}))}{\int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}})) \, \mathrm{d}\hat{\boldsymbol{x}}} \, \mathrm{d}\boldsymbol{y} \right)^2 \mathrm{d}\boldsymbol{x} \, \mathrm{d}z \\
&= -2 \int_{\Omega_3} \left(\left(s(\boldsymbol{x}, z) - \int_{\Omega_2} \frac{(u + \epsilon v)(\boldsymbol{y}) \cdot h(\boldsymbol{x} - \boldsymbol{y}, z - d(\boldsymbol{y}))}{\int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}})) \, \mathrm{d}\hat{\boldsymbol{x}}} \, \mathrm{d}\boldsymbol{y} \right) \right. \\
&\qquad \left. \cdot \int_{\Omega_2} \frac{v(\tilde{\boldsymbol{x}}) \cdot h(\boldsymbol{x} - \tilde{\boldsymbol{x}}, z - d(\tilde{\boldsymbol{x}}))}{\int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}})) \, \mathrm{d}\hat{\boldsymbol{x}}} \, \mathrm{d}\tilde{\boldsymbol{x}} \right) \mathrm{d}\boldsymbol{x} \, \mathrm{d}z
\end{aligned} \tag{4.18}$$

If we now set $\epsilon = 0$ and change the order of integration, we proceed with

$$\begin{aligned}
&\frac{\partial}{\partial \epsilon} E_{\mathsf{D}}(s, u + \epsilon v, d, \mathcal{F}_{\mathrm{N}})\Big|_{\epsilon=0} \\
&= -2 \int_{\Omega_2} \left(\int_{\Omega_3} \frac{s(\boldsymbol{x}, z) - \int_{\Omega_2} u(\boldsymbol{y}) \cdot \frac{h(\boldsymbol{x}-\boldsymbol{y}, z-d(\boldsymbol{y}))}{\int_{\Omega_2} h(\boldsymbol{x}-\hat{\boldsymbol{x}}, z-d(\hat{\boldsymbol{x}}))\ \mathrm{d}\hat{\boldsymbol{x}}}\ \mathrm{d}\boldsymbol{y}}{\int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}}))\ \mathrm{d}\hat{\boldsymbol{x}}} \right. \\
&\qquad \left. \cdot\, h(\boldsymbol{x} - \tilde{\boldsymbol{x}}, z - d(\tilde{\boldsymbol{x}}))\ \ \mathrm{d}\boldsymbol{x}\ \mathrm{d}z \right) \cdot v(\tilde{\boldsymbol{x}})\ \mathrm{d}\tilde{\boldsymbol{x}}\ . \qquad (4.19)
\end{aligned}$$

The variational gradient w.r.t. u can then be obtained via the requirement of Equation (2.27). Moreover, we use the notation of the adjoint of h, i.e. $h^*(\boldsymbol{x}, z) = h(-\boldsymbol{x}, -z)$ and obtain:

$$\begin{aligned}
\frac{\delta E_{\mathsf{D}}}{\delta u} = -2 \int_{\Omega_3} \left(\frac{s(\boldsymbol{x}, z) - \int_{\Omega_2} u(\boldsymbol{y}) \cdot \frac{h(\boldsymbol{x}-\boldsymbol{y}, z-d(\boldsymbol{y}))}{\int_{\Omega_2} h(\boldsymbol{x}-\hat{\boldsymbol{x}}, z-d(\hat{\boldsymbol{x}}))\ \mathrm{d}\hat{\boldsymbol{x}}}\ \mathrm{d}\boldsymbol{y}}{\int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}}))\ \mathrm{d}\hat{\boldsymbol{x}}} \right. \\
\left. \cdot\, h^*(\tilde{\boldsymbol{x}} - \boldsymbol{x}, d(\tilde{\boldsymbol{x}}) - z)\ \right) \mathrm{d}\boldsymbol{x}\ \mathrm{d}z\ . \qquad (4.20)
\end{aligned}$$

Let us now introduce the abbreviation $N(\boldsymbol{x}, z) := \int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}}))\ \mathrm{d}\hat{\boldsymbol{x}}$ that corresponds to the normalisation function, i.e. the denominator in Equation (4.12). Then, we can shorten the equation above to

$$\frac{\delta E_{\mathsf{D}}}{\delta u} = -2 \int_{\Omega_3} \Big(N^{-1}(\boldsymbol{x}, z) \cdot r[u, d](\boldsymbol{x}, z) \cdot h^*(\tilde{\boldsymbol{x}} - \boldsymbol{x}, d(\tilde{\boldsymbol{x}}) - z)\ \Big) \mathrm{d}\boldsymbol{x}\ \mathrm{d}z\ , \quad (4.21)$$

where $r[u, d]$ denotes the residual error defined in Equation (4.13). The remaining integration constitutes a three-dimensional convolution. Accordingly, we can make use of the convolution operator $*$ and thereby express the variational gradient w.r.t. to the sharp pinhole image u as

$$\frac{\delta E}{\delta u}(\boldsymbol{x}) = -2\ \Big(\big(N^{-1} \cdot r[u, d] \big) * h^* \Big)(\boldsymbol{x}, d(\boldsymbol{x}))\ , \qquad (4.22)$$

where $\tilde{\boldsymbol{x}}$ has been substituted by $\boldsymbol{x}$, in the last step. Further, we have replaced E_{D} by E since no regularisation on u is involved in this model.

Let us now derive the functional derivative w.r.t. the depth map d. To this end, we again additively perturb within the data term, but this time the

perturbation is applied to the third argument of E_{D}:

$$\begin{aligned}
&\frac{\partial}{\partial\epsilon} E_{\mathsf{D}}(s, u, d + \epsilon v, \mathcal{F}_{\mathrm{N}}) \\
&= \frac{\partial}{\partial\epsilon} \int_{\Omega_3} \Big(s(\boldsymbol{x}, z) - \mathcal{F}_{\mathrm{N}}[u, d + \epsilon v](\boldsymbol{x}, z) \Big)^2 \mathrm{d}\boldsymbol{x}\ \mathrm{d}z \qquad (4.23) \\
&= \frac{\partial}{\partial\epsilon} \int_{\Omega_3} \left(s(\boldsymbol{x}, z) - \int_{\Omega_2} \frac{u(\boldsymbol{y}) \cdot h(\boldsymbol{x} - \boldsymbol{y}, z - (d + \epsilon v)(\boldsymbol{y}))}{\int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - (d + \epsilon v)(\hat{\boldsymbol{x}}))\ \mathrm{d}\hat{\boldsymbol{x}}}\ \mathrm{d}\boldsymbol{y} \right)^2 \mathrm{d}\boldsymbol{x}\ \mathrm{d}z\,.
\end{aligned}$$

Since ϵ is placed in the nominator as well as in the denominator, we have to follow the quotient rule. This is aggravated by the fact that it is also placed within the argument of the PSF h. Consequently, we additionally have to consider the chain rule. In doing so, we obtain:

$$\begin{aligned}
&\frac{\partial}{\partial\epsilon} E_{\mathsf{D}}(s, u, d + \epsilon v, \mathcal{F}_{\mathrm{N}}) \\
&= -2 \int_{\Omega_3} \left(s(\boldsymbol{x}, z) - \int_{\Omega_2} u(\boldsymbol{y}) \frac{h(\boldsymbol{x} - \boldsymbol{y}, z - (d + \epsilon v)(\boldsymbol{y}))}{\int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - (d + \epsilon v)(\hat{\boldsymbol{x}}))\ \mathrm{d}\hat{\boldsymbol{x}}}\ \mathrm{d}\boldsymbol{y} \right) \cdot \\
&\quad \left(\int_{\Omega_2} \frac{1}{\left(\int_{\Omega_2} h(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - (d + \epsilon v)(\hat{\boldsymbol{x}}))\ \mathrm{d}\hat{\boldsymbol{x}} \right)^2} \cdot \right. \\
&\quad \Big(-h_z(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - (d + \epsilon v)(\bar{\boldsymbol{x}})) \cdot v(\bar{\boldsymbol{x}}) \cdot \int_{\Omega_2} h(\boldsymbol{x} - \tilde{\boldsymbol{x}}, z - (d + \epsilon v)(\tilde{\boldsymbol{x}}))\ \mathrm{d}\tilde{\boldsymbol{x}} \\
&\quad + h(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - (d + \epsilon v)(\bar{\boldsymbol{x}})) \cdot \int_{\Omega_2} h_z(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - (d + \epsilon v)(\hat{\boldsymbol{x}}))\ v(\hat{\boldsymbol{x}})\ \mathrm{d}\hat{\boldsymbol{x}} \Big) \\
&\quad \left. \cdot\ u(\bar{\boldsymbol{x}})\ \mathrm{d}\bar{\boldsymbol{x}} \right) \mathrm{d}\boldsymbol{x}\ \mathrm{d}z\ . \qquad (4.24)
\end{aligned}$$

Here h_z denotes the partial derivative of h in z-direction. If we set $\epsilon = 0$ and

use the introduced abbreviations, we proceed with:

$$\begin{aligned}
&\frac{\partial}{\partial\epsilon} E_{\mathsf{D}}(s, u, d + \epsilon v, \mathcal{F}_{\mathrm{N}})\Big|_{\epsilon=0} \\
&= -2 \int_{\Omega_3} \frac{r[u,d](\boldsymbol{x}, z)}{N(\boldsymbol{x}, z)^2} \cdot \Bigg(\int_{\Omega_2} u(\bar{\boldsymbol{x}}) \cdot \Big(-h_z(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}})) \cdot v(\bar{\boldsymbol{x}}) \cdot N(\boldsymbol{x}, z) \\
&\quad + h(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}})) \cdot \int_{\Omega_2} h_z(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}}))\ v(\hat{\boldsymbol{x}})\ \mathrm{d}\hat{\boldsymbol{x}} \Big)\ \mathrm{d}\bar{\boldsymbol{x}} \Bigg) \mathrm{d}\boldsymbol{x}\ \mathrm{d}z \\
&= -2 \int_{\Omega_3} \frac{r[u,d](\boldsymbol{x}, z)}{N(\boldsymbol{x}, z)^2} \cdot \Bigg(\int_{\Omega_2} -u(\bar{\boldsymbol{x}}) \cdot h_z(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}})) \cdot v(\bar{\boldsymbol{x}}) \cdot N(\boldsymbol{x}, z)\ \mathrm{d}\bar{\boldsymbol{x}} \\
&\quad + \iint_{\Omega_2} u(\bar{\boldsymbol{x}}) \cdot h(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}})) \cdot h_z(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}}))\ v(\hat{\boldsymbol{x}})\ \mathrm{d}\hat{\boldsymbol{x}}\ \mathrm{d}\bar{\boldsymbol{x}} \Bigg) \mathrm{d}\boldsymbol{x}\ \mathrm{d}z \\
&= -2 \int_{\Omega_3} \frac{r[u,d](\boldsymbol{x}, z)}{N(\boldsymbol{x}, z)} \cdot \Bigg(\int_{\Omega_2} -u(\bar{\boldsymbol{x}}) \cdot h_z(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}})) \cdot v(\bar{\boldsymbol{x}})\ \mathrm{d}\bar{\boldsymbol{x}} \\
&\quad + \int_{\Omega_2} \frac{\int_{\Omega_2} u(\bar{\boldsymbol{x}})\ h(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}}))\ \mathrm{d}\bar{\boldsymbol{x}}}{N(\boldsymbol{x}, z)} \cdot h_z(\boldsymbol{x} - \hat{\boldsymbol{x}}, z - d(\hat{\boldsymbol{x}})) \cdot v(\hat{\boldsymbol{x}})\ \mathrm{d}\hat{\boldsymbol{x}} \Bigg) \mathrm{d}\boldsymbol{x}\ \mathrm{d}z\ .
\end{aligned}$$

The fraction in the last row just describes our forward operator $\mathcal{F}_{\mathrm{N}}$. Moreover, we can harmonise the variables of integration. This yields

$$\begin{aligned}
&\frac{\partial}{\partial\epsilon} E_{\mathsf{D}}(s, u, d + \epsilon v, \mathcal{F}_{\mathrm{N}})\Big|_{\epsilon=0} \\
&= -2 \int_{\Omega_3} \frac{r[u,d](\boldsymbol{x}, z)}{N(\boldsymbol{x}, z)} \cdot \Bigg(\int_{\Omega_2} -u(\bar{\boldsymbol{x}})\ h_z(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}})) \cdot v(\bar{\boldsymbol{x}})\ \mathrm{d}\bar{\boldsymbol{x}} \\
&\qquad\qquad + \int_{\Omega_2} \mathcal{F}_{\mathrm{N}}[u,d](\boldsymbol{x}, z) \cdot h_z(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}})) \cdot v(\bar{\boldsymbol{x}})\ \mathrm{d}\bar{\boldsymbol{x}} \Bigg) \mathrm{d}\boldsymbol{x}\ \mathrm{d}z \\
&\overset{*}{=} 2 \int_{\Omega_2} \int_{\Omega_3} \frac{r[u,d](\boldsymbol{x}, z)}{N(\boldsymbol{x}, z)}\ \cdot \\
&\qquad \Bigg(\Big(u(\bar{\boldsymbol{x}}) - \mathcal{F}_{\mathrm{N}}[u,d](\boldsymbol{x}, z) \Big) \cdot h_z(\boldsymbol{x} - \bar{\boldsymbol{x}}, z - d(\bar{\boldsymbol{x}})) \Bigg)\ \mathrm{d}\boldsymbol{x}\ \mathrm{d}z \cdot v(\bar{\boldsymbol{x}})\ \mathrm{d}\bar{\boldsymbol{x}}\ ,
\end{aligned}$$

where we have changed the order of integration in (*). At this point, we follow once more the requirement in Equation (2.27) giving the variational gradient

$$\frac{\delta E_{\mathrm{D}}}{\delta d}(\bar{\boldsymbol{x}}) = 2\cdot$$

$$\int_{\Omega_3} \frac{r[u,d](\boldsymbol{x},z)}{N(\boldsymbol{x},z)} \cdot \left(\Big(u(\bar{\boldsymbol{x}}) - \mathcal{F}_{\mathrm{N}}[u,d](\boldsymbol{x},z)\Big) \cdot h_z(\boldsymbol{x}-\bar{\boldsymbol{x}}, z-d(\bar{\boldsymbol{x}}))\right) \mathrm{d}\boldsymbol{x}\ \mathrm{d}z$$

$$= 2\cdot \int_{\Omega_3} \frac{r[u,d](\boldsymbol{x},z)}{N(\boldsymbol{x},z)} \cdot h_z^*(\bar{\boldsymbol{x}}-\boldsymbol{x}, d(\bar{\boldsymbol{x}})-z)\ \mathrm{d}\boldsymbol{x}\ \mathrm{d}z \cdot u(\bar{\boldsymbol{x}})$$

$$- 2\cdot \int_{\Omega_3} \frac{r[u,d](\boldsymbol{x},z)}{N(\boldsymbol{x},z)}\ \mathcal{F}_{\mathrm{N}}[u,d](\boldsymbol{x},z) \cdot h_z^*(\bar{\boldsymbol{x}}-\boldsymbol{x}, d(\bar{\boldsymbol{x}})-z)\ \mathrm{d}\boldsymbol{x}\ \mathrm{d}z\ . \tag{4.25}$$

As one can see, we can now simplify things by using the convolution operator. If we further incorporate regularisation to the depth map, the functional derivative w.r.t. d can be written as

$$\frac{\delta E}{\delta d}(\boldsymbol{x}) = 2\left(\left(\frac{r[u,d]}{N} * h_z^*\right)(\boldsymbol{x}, d(\boldsymbol{x}))\cdot u - \left(\left(\frac{r[u,d]}{N}\cdot \mathcal{F}_{\mathrm{N}}[u,d]\right) * h_z^*\right)(\boldsymbol{x}, d(\boldsymbol{x}))\right.$$

$$\left. - \alpha \cdot \mathrm{div}\Big(\Psi'(|\boldsymbol{\nabla} d|^2)\ \boldsymbol{\nabla} d\Big)\right)\ . \tag{4.26}$$

Since $\Psi'_d > 0$, the natural boundary condition reads $\boldsymbol{n}^\top \boldsymbol{\nabla} d = 0$, where $\boldsymbol{n}^\top$ is the outer normal vector at the image boundary.

Enforcing Positivity

In Section 3.2.3 we have discussed the multiplicative Euler-Lagrange (EL) formalism. Its usage has been motivated by the obvious assumption that the number of arriving photons at the imaging sensor is larger than zero. Due to the positivity preservation property of the multiplicative EL formalism, the solution is restricted to the plausible positive range, and the ill-posedness of the incorporated deconvolution part can be mitigated [Welk, 2010; Welk and Nagy, 2007]. Regarding our depth-from-defocus approach, we are in a very similar situation. We also have to handle a deconvolution part here and the accompanying ill-posedness. However, we can still assume that the acquired

intensities are in the positive range. Besides that, the considered surface is lying in front of the lens which also implies a positive range concerning the depth values. Hence, it should make sense to derive the variational gradients via this multiplicative formalism. For this purpose, we can exploit the fact that a variational gradient of the multiplicative formalism can be obtained by the classical functional derivative multiplied with the unknown function (cf. Section 3.2.3). Accordingly, in this multiplicative setting, the functional derivative w.r.t. the pinhole image reads

$$\frac{\delta^* E}{\delta u}(\boldsymbol{x}) = -2u\Big(\Big(N^{-1} \cdot r[u,d]\Big) * h^*\Big)(\boldsymbol{x}, d(\boldsymbol{x}))\ , \tag{4.27}$$

and the one w.r.t. the depth d is given by

$$\frac{\delta^* E}{\delta d}(\boldsymbol{x}) = 2d\Bigg(\Big(\frac{r[u,d]}{N} * h_z^*\Big)(\boldsymbol{x}, d(\boldsymbol{x})) \cdot u - \Big(\Big(\frac{r[u,d]}{N} \cdot \mathcal{F}_{\mathrm{N}}[u,d]\Big) * h_z^*\Big)(\boldsymbol{x}, d(\boldsymbol{x})) - \alpha \cdot \operatorname{div}\Big(\Psi'(|\boldsymbol{\nabla} d|^2)\ \boldsymbol{\nabla} d\Big)\Bigg)\ . \tag{4.28}$$

The boundary condition remains the same as in the additive formalism.

4.2.3 Multi-Channel Images

Until now, our depth-from-defocus approach has been restricted to single-channel images. This has been done to simplify things in illustrating the main ideas and strategies such as the derivation of the functional derivatives. In this section, we now explain how to extend our method to the more general multi-channel case. To this end, let us assume the refraction index of the lens to be independent of the wavelength of the light. Consequently, the lens treats all channels equally. This results in a uniform, channel-invariant PSF. Additionally, the focal length and therewith the distance of the focal plane does not change between different channels. To approximate the depth-of-field effect given a depth map and a sharp pinhole image with channel index set $\mathcal{C}$, we thus can apply our forward operator $\mathcal{F}_{\mathrm{N}}$ channel-wise in a straightforward way (cf. Figure 4.7).

To solve the inverse problem, Aguet et al. [2008] propose to convert a multi-channel image into a single-channel one before performing their framework. To obtain a multi-channel texture, only within the last processing step, after the depth estimation is completed, the texture is reconstructed

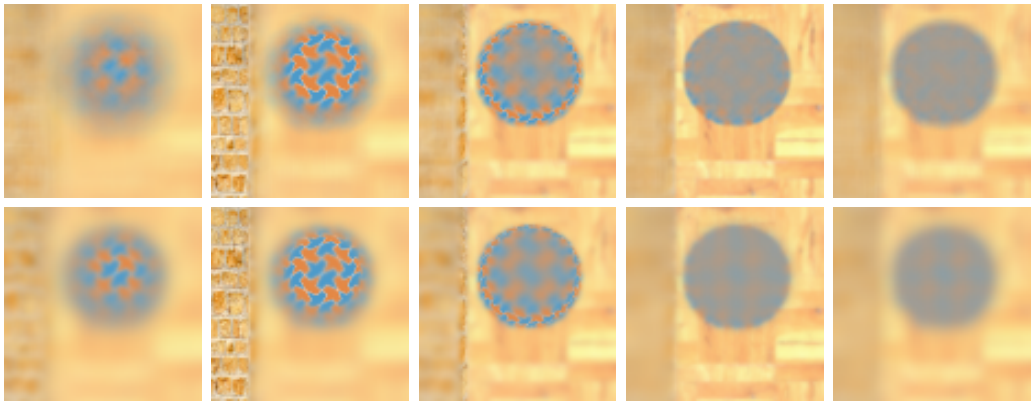

Figure 4.7: Multi-channel focal stack. Simulating the depth-of-field effect by applying our novel forward operator channel-wise to an RGB pinhole image. 5 out of 20 slices. **(a) Top:** Result of a thin lens camera renderer. **(b) Bottom:** Our result.

by considering all channels. Although this also constitutes a way to reconstruct a sharp multi-channel image, the applied grey scale conversion entails a loss of information. Here, we want to prevent this loss and improve the reconstruction quality by incorporating the information of all channels of the recorded focal stack $\boldsymbol{s} = (s_c)_{c\in\mathcal{C}}$. To this end, we consider the residual error

$$r_c(\boldsymbol{u}, d) = s_c - \mathcal{F}_{\mathrm{N}}[u_c, d] \;, \quad \forall c \in \mathcal{C} \;. \tag{4.29}$$

To determine the sharp multi-channel image $\boldsymbol{u} = (u_c)_{c\in\mathcal{C}}$ along with the depth map d, we follow Welk and Nagy [2007] in the context of variational deconvolution of multi-channel images. Accordingly, we sum up the squared residual errors over all channels $c \in \mathcal{C}$ within the data term

$$E_{\mathsf{D2}}(s, \boldsymbol{u}^\top, d, \mathcal{F}_{\mathrm{N}}) := \int_{\Omega_3} R(\boldsymbol{u}, d) \; \mathrm{d}\boldsymbol{x} \; \mathrm{d}z \;, \tag{4.30}$$

where we use the abbreviation

$$R(\boldsymbol{u}, d) := \sum_{c\in\mathcal{C}} \Big(r_c(\boldsymbol{u}, d) \Big)^2 \;. \tag{4.31}$$

Since we estimate a joint depth map for all channels, it remains a one channel signal, and there is no change required in the smoothness term.

4.2.4 Robustification

In the last chapter, we have already benefited from robustification strategies. Hence, it is a straightforward idea to also improve the results of our depth-from-defocus approach in this way. For this purpose, let us now have a closer look at the data term. It measures the distance between the forward operation and the given data. Although a quadratic penalisation of deviations is especially suited in the presence of Gaussian distributed noise, it may penalise outliers or deviations of the model assumptions from the recorded stack too severely. As already discussed in the last chapter, the response of an optical system to a point light source depends on a lot of different factors such as optical imperfections of the lenses or diffraction phenomena. In our cell reconstruction framework, the PSF has been estimated by means of physical measurements. Hence, at least partially, the response of the used optical system has been incorporated into our reconstruction framework. Instead, in this chapter, the forward process and thereby the employed PSF is derived completely theoretically. Choosing a pillbox or Gaussian kernel can thus only be a rough approximation of the true PSF. For this reason, the robustification strategy of Zervakis et al. [1995], Bar et al. [2005], and Welk [2010] that we followed already in the last chapter, should even more show its potential for improvements in our depth-from-defocus approach. Hence, let us replace the quadratic data term above by a robust one:

$$E_{\mathrm{D3}}(s, \boldsymbol{u}^\top, d, \mathcal{F}_{\mathrm{N}}) := \int_{\Omega_3} \Phi(R(\boldsymbol{u}, d)) \, \mathrm{d}\boldsymbol{x} \, \mathrm{d}z \;, \tag{4.32}$$

with the non-negative, subquadratic penaliser function $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$ in order to give large outliers less influence. More precisely, we apply the regularised L^1-norm $\Phi(x^2) = \sqrt{x^2 + \epsilon}$ (cf. Figure 3.4(b)) with some small stabilisation $\epsilon > 0$ to avoid singularities at 0.

Regarding the more general multi-channel case, an associated minimiser $(\boldsymbol{u}, d)$ has now to fulfil the set of Euler-Lagrange equations $\frac{\delta^* E}{\delta u_c} = 0, \forall c \in \mathcal{C}$ as well as $\frac{\delta^* E}{\delta d} = 0$. If we further incorporate the proposed robustification and follow multiplicative EL formalism, we obtain for each channel

$$\frac{\delta^* E}{\delta u_c}(\boldsymbol{x}) = -2 \; u_c \cdot \Big(\Big(\Phi' \cdot \overline{r}_c \Big) * h^* \Big) (\boldsymbol{x}, d(\boldsymbol{x})) \tag{4.33}$$

w.r.t. $\boldsymbol{u}$ and

$$\frac{\delta^* E}{\delta d}(\boldsymbol{x}) = 2d \cdot \Bigg(\sum_{c \in \mathcal{C}} \Bigg(\Big(\big(\Phi' \cdot \overline{r}_c \big) * h_z^* \Big)(\boldsymbol{x}, d(\boldsymbol{x})) \cdot u_c - \Big(\Big(\Phi' \cdot \overline{r}_c \cdot \mathcal{F}_{\mathrm{N}}[u_c, d] \Big) * h_z^* \Big)(\boldsymbol{x}, d(\boldsymbol{x})) \Bigg) - \alpha \cdot \operatorname{div} \Big(\Psi'(|\boldsymbol{\nabla} d|^2) \; \boldsymbol{\nabla} d \Big) \Bigg) \tag{4.34}$$

w.r.t. d where we have introduced the abbreviations $\Phi' := \Phi'(R(\boldsymbol{u}, d))$ and $\overline{r}_c := N^{-1} \cdot r_c$. For the channel-wise residuum r_c (4.29) we omit the arguments for a better readability. Note that setting $\Phi' := 1$ comes down to the minimality conditions of the multi-channel approach without robustification.

4.3 Joint Denoising and Depth-from-Defocus

Until now, we have assumed that our recorded focal stack is free from any kind of noise. The only perturbation involved was blur, which can be exploited to estimate the depth. Therefore, and since the ill-posedness of the problem has been counteracted by imposing smoothness on d, there has been no need to postulate any smoothness assumption on the sought pinhole image u. Moreover, smoothness of u, of course, would counteract the deblurring process and we want to recover a sharp pinhole image. However, especially in the context of microscopy at low light intensity or due to signal processing in general, the recorded stacks can contain noise. Then, it can be beneficial to demand (piecewise) smoothness of the reconstructed sharp image. Of course considering denoising and depth-from-defocus as two separate tasks offers a straightforward strategy. Techniques for image denoising have already been discussed in the Chapter 2 of this dissertation: In Section 2.1, we have revisited diffusion-based image filtering. Besides that, variational image restoration has been recapitulated in Section 2.2. Both methods allow to denoise each 2-D image $\tilde{w} : \Omega_2 \to \mathbb{R}_+$ of a focal stack before performing a standard depth-from-defocus method. For the multi-channel case, let $\tilde{\boldsymbol{w}} = (\tilde{w}_c)_{c \in \mathcal{C}}$ denotes the input slice of a noisy focal stack. Then the filtered slice $\boldsymbol{w}$ can be determined, for instance, as minimiser of

$$E(\boldsymbol{w}^\top) = \int_{\Omega_2} \sum_{c \in \mathcal{C}} \big(\tilde{w}_c - w_c \big)^2 \, \mathrm{d}\boldsymbol{x} + \gamma \, E_{\mathsf{S}_2}(\boldsymbol{w}^\top) \; . \tag{4.35}$$

Here the smoothness term E_{S_2} is applied to the evolving image $\boldsymbol{w}$ and γ balances between smoothness and accuracy.

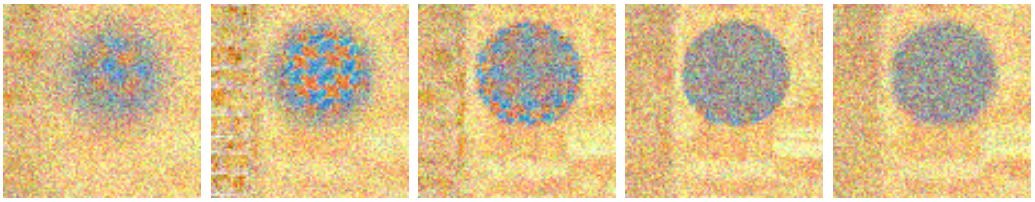

Figure 4.8: Focal stack disturbed by artificial Gaussian noise with standard deviation $\sigma_{\text{noise}} = 30$ and mean 0.

However, a more promising idea is to couple these tasks into one joint model instead of treating them separately. As already mentioned in Section 2.5.3, in the context of deblurring noisy data, in [Osher and Rudin, 1994; Marquina and Osher, 1999] a simultaneous approach is recommended. Additionally, Figure 2.10 demonstrates that regularisation allows to suppress the occurrence of oscillation artefacts in deconvolution approaches. Since our depth-from-defocus approach also contains a deconvolution part, it is promising to follow these ideas. Moreover, in our approach, we treat the unknown depth as a parameter of the blur kernel. Therefore blind deconvolution approaches such as the one by Chan and Wong [1998] can also be seen as related. An extension to a joint restoration and blind deconvolution model was presented by You and Kaveh [1996b,a]. Following these approaches, we extend our method to

$$E(\boldsymbol{u}^\top, d) = E_{\mathsf{D}}(s, \boldsymbol{u}^\top, d, \mathcal{F}_{\mathrm{N}}) + \alpha\, E_{\mathsf{S}}(d) + \beta\, E_{\mathsf{S}_2}(\boldsymbol{u}^\top)\,, \tag{4.36}$$

where a second regularisation term

$$E_{\mathsf{S}_2}(\boldsymbol{u}) = \int_{\Omega_2} \Psi_u(|\boldsymbol{\nabla} \boldsymbol{u}|^2)\, \mathrm{d}\boldsymbol{x}\,, \tag{4.37}$$

is applied to the evolving sharp pinhole image $\boldsymbol{u}$. Here $\alpha, \beta \geq 0$ balance the two regularisation terms against the data term where $|\boldsymbol{\nabla} \boldsymbol{u}|^2 := \sum_{c \in \mathcal{C}} |\boldsymbol{\nabla} u_c|^2$. The functions Ψ (see Equation (4.16)) and Ψ_u allow different smoothing behaviours for each regularisation.

Since the second regularisation term only depends on the sharp pinhole image u, the functional derivative w.r.t. d remains the same as in Equation (4.34). The variational gradient w.r.t. u_c changes to

$$\frac{\delta^* E}{\delta u_c}(\boldsymbol{x}) = -2\, u_c \cdot \Big(\Big(\big(\Phi' \cdot \overline{r}_c \big) * h^* \Big)(\boldsymbol{x}, d(\boldsymbol{x})) + \beta \cdot \operatorname{div}\Big(\Psi'_u\, \boldsymbol{\nabla} u_c \Big) \Big)\,, \tag{4.38}$$

where the abbreviation $\Psi'_u := \Psi'_u(|\boldsymbol{\nabla} \boldsymbol{u}|^2)$ is used. Since also $\Psi'_u > 0$ the boundary conditions read

$$\boldsymbol{n}^\top \boldsymbol{\nabla} u_c = 0 \,, \quad \forall c \in \mathcal{C} \qquad \text{and} \qquad \boldsymbol{n}^\top \boldsymbol{\nabla} d = 0 \,. \tag{4.39}$$

Please note that in case of very low intensity levels where Poisson distributed noise becomes relevant, one should think about coming back to Csiszár's information divergence [Csiszár, 1991]. Then our forward operator $\mathcal{F}_{\mathrm{N}}$ should be considered within the data term (3.10) or its robust variant. While Csiszár's I-divergence is regarded, e.g. by Favaro and Soatto [2000] and Jin and Favaro [2002], we are not employing it in our work.

4.4 Discretisation and Implementation

In the previous sections, we have proposed a novel variational depth-from-defocus approach and suggested further ideas for its improvement. Besides that, we have presented a joint denoising and depth-from-defocus approach. So far, however, only continuous ideas have been formulated. In this section, we now show how to transfer them into the discrete setting. On the one hand, this is required for the application of these ideas to digital images. On the other hand, we can thereby again follow iterative methods to solve the presented PDEs.

Regarding the last point, we once more profit from applying the multiplicative Euler-Lagrange formalism. In addition to constraining the solution to the plausible positive range, it also allows an iteration strategy that is in accordance with our fast and stabilised iteration scheme of our cell reconstruction framework (cf. Section 3.4). Accordingly, concerning the multiplicative functional derivative from Equation (4.38) w.r.t. $\boldsymbol{u} = (u_c)_{c \in \mathcal{C}}$, we suggest the following gradient descent scheme:

$$\frac{u_c^{k+1} - u_c^k}{\tau} = 2\, u_c^{k+1} \Big(\Big(\Phi'^k \cdot \overline{r}_c^{\,k} \Big) * h^* \Big)(\boldsymbol{x}, d) + 2\beta \cdot \operatorname{div} \Big(\Psi'^k_u \cdot \boldsymbol{\nabla} u_c^{k+1} \Big) \cdot u_c^k \,, \tag{4.40}$$

for all $c \in \mathcal{C}$, where τ is the relaxation parameter, and k denotes the iteration level. Furthermore, we have used the abbreviations $\Phi'^k := \Phi'(R(\boldsymbol{u}^k, d))$ and $\Psi'^k_u := \Psi'_u(|\boldsymbol{\nabla} \boldsymbol{u}^k|^2)$ evaluated in a lagged diffusivity manner (Kačanov-method) [Fučik et al., 1973; Chan and Mulet, 1999; Vogel, 2002]. Exactly like the gradient descent scheme of Equation (3.50), the unknown that appears as factor in the similarity expression (the first summand), is evaluated at the new time step $k+1$, while the smoothness expression, i.e. last summand

of (4.40) is evaluated in a semi-implicit fashion. Proceeding analogously in estimating the depth map d, we obtain

$$\begin{aligned}\frac{d^{k+1}-d^k}{\tau} = -2 \sum_{c\in\mathcal{C}} \Bigg(& \Big(\Big(\Phi'^k \cdot \overline{r}_c^k \Big) * h_z^* \Big)(\boldsymbol{x}, d^k) \cdot u_c \\ & - \Big(\Big(\Phi'^k \cdot \overline{r}_c^k \cdot \mathcal{F}_{\mathrm{N}}[u_c, d^k] \Big) * h_z^* \Big)(\boldsymbol{x}, d^k) \Big) \cdot d^{k+1} \\ & + 2\alpha \cdot \operatorname{div}\Big(\Psi'^k \cdot \boldsymbol{\nabla} d^{k+1} \Big) \cdot d^k \,. \end{aligned} \tag{4.41}$$

Here, we abbreviate $\Psi'^k := \Psi'(|\boldsymbol{\nabla} d^k|^2)$ and $\Phi'^k := \Phi'(R(\boldsymbol{u}, d^k))$. Using the standard additive Euler-Lagrange formalism requires to adapt the relaxation parameter in each iteration step. For this purpose, the computation of a suitable step-size has to be done by time-expensive algorithms such as the backtracking line-search method (see e.g. [Nocedal and Wright, 2006]). Since the suggested semi-implicit scheme above gives a higher stability range w.r.t. the relaxation parameter τ, the step-size can be chosen fixed in advance. Therefore, we can refrain from such computationally expensive algorithms.

To apply the presented method on a focal stack consisting of N_z digital images, each sampled on a rectangular regular grid of size $N_{x_1} \times N_{x_2} =: N$, we replace continuous functions by their discrete approximations. Hence, in each pixel $(i,j) \in \{1,...,N_{x_1}\} \times \{1,...,N_{x_2}\}$, we have to fulfil

$$\frac{u_{c,i,j}^{k+1} - u_{c,i,j}^k}{\tau} = \underbrace{2 \left[\Big(\Big(\Phi'^k \cdot \overline{r}_c^k \Big) * h^* \Big) \right]_{i,j,d_{i,j}}}_{\boldsymbol{D}_1} \cdot u_{c,i,j}^{k+1} + 2\beta \cdot \underbrace{\Big[\operatorname{div}\Big(\Psi_u'^k \, \boldsymbol{\nabla} u_c^{k+1} \Big) \Big]_{i,j}}_{\boldsymbol{A}(\boldsymbol{u}^k)\boldsymbol{u}^{k+1}} \cdot \underbrace{u_{c,i,j}^k}_{\boldsymbol{D}_2} \tag{4.42}$$

w.r.t. the sharp pinhole image $\boldsymbol{u} = (u_c)_{c\in\mathcal{C}}$. For the 3-D convolution denoted by the operator $*$, we again exploit the convolution theorem (cf. Section 2.5.1) [Gasquet et al., 1998; Bracewell, 1999]. If we refrain from regularisation of the sharp pinhole image, i.e. $\beta = 0$, Equation (4.40) can be solved directly by

$$u_{c,i,j}^{k+1} = \frac{u_{c,i,j}^k}{1 - \tau \left[\Big(\Big(\Phi'^k \cdot \overline{r}_c^k \Big) * h^* \Big) \right]_{i,j,d_{i,j}}} \,. \tag{4.43}$$

Otherwise, we have to solve a linear system of equations to which end we again switch to a single-index notation. Accordingly, let us describe 2-D multi-channel images $\boldsymbol{u} : \mathbb{R}^2 \to \mathbb{R}_+^{|\mathcal{C}|}$ by vectors $\boldsymbol{u} \in \mathbb{R}^{N_C}$ and the point-wise

multiplications with the help of the diagonal matrices

$$\boldsymbol{D}_1 := \operatorname{diag}\left(2\left[\left(\left(\Phi'^k \cdot \overline{r}_c^k\right) * h^*\right)\right]_1, \ldots, 2\left[\left(\left(\Phi'^k \cdot \overline{r}_c^k\right) * h^*\right)\right]_{N_C}\right), \tag{4.44}$$

and $\boldsymbol{D}_2 := \operatorname{diag}\left(u_1^k, \ldots, u_{N_c}^k\right)$, where $N_C := N \cdot |\mathcal{C}|$. The discrete regularisation term is again accomplished by $\boldsymbol{A}(\boldsymbol{u}^k)\boldsymbol{u}^{k+1}$. The diffusion matrix $\boldsymbol{A}$ can be defined according to Equation (2.10) where the diffusivity function g is represented by Ψ'. With that, the equation above can be written in the same form as in (3.58):

$$\underbrace{\left(\boldsymbol{I} - \tau(\boldsymbol{D}_1(\boldsymbol{u})^k + 2\beta \cdot \boldsymbol{D}_2(\boldsymbol{u}^k) \cdot \boldsymbol{A}(\boldsymbol{u}^k))\right)}_{\boldsymbol{B}} \cdot \underbrace{\boldsymbol{u}^{k+1}}_{\boldsymbol{x}} = \underbrace{\boldsymbol{u}^k}_{\boldsymbol{b}} . \tag{4.45}$$

Since we only employ isotropic regularisation in this chapter, we choose the Jacobi algorithm [Morton and Mayers, 2005] with iteration index m to solve this system of equations. In this way, we obtain for $p = 1 \ldots N_C$:

$$u_p^{k+1,m+1} = \frac{\left(1 + 2\tau \cdot \beta \cdot \sum\limits_{\substack{\ell \in \\ \{x_1, x_2\}}} \sum\limits_{\substack{q \in \\ \mathcal{N}_\ell(p)}} \frac{\Psi'^k_{u_p} + \Psi'^k_{u_q}}{2h_\ell^2} \cdot u_q^{k+1,m}\right) \cdot u_p^k}{1 - 2\tau\left(\left[\left(\left(\Phi'^k \cdot \overline{r}_c^k\right) * h^*\right)\right]_p + \beta \cdot \; u_p^k \cdot \sum\limits_{\substack{\ell \in \\ \{x_1, x_2\}}} \sum\limits_{\substack{q \in \\ \mathcal{N}_\ell(p)}} \frac{\Psi'^k_{u_p} + \Psi'^k_{u_q}}{2h_\ell^2}\right)} . \tag{4.46}$$

Concerning the estimation of the depth map d, we proceed analogously. To formulate a system of equations like (4.45), we set $\boldsymbol{D}_2 := \operatorname{diag}\left(d_1^k, \ldots, d_N^k\right)$, exchange β by α, and set

$$\begin{aligned}\boldsymbol{D}_1 := \operatorname{diag}\Bigg(&-2\sum_{c\in\mathcal{C}}\left(\left[\left(\left(\Phi'^k \cdot \overline{r}_c^k\right) * h_z^*\right] \cdot u_1 - \left[\left(\Phi'^k \cdot \overline{r}_c^k \; \cdot \mathcal{F}_{\mathrm{N}}[u_c, d^k]\right) * h_z^*\right]_1\right), \ldots \\ &\ldots, -2\sum_{c\in\mathcal{C}}\left(\left[\left(\left(\Phi'^k \cdot \overline{r}_c^k\right) * h_z^*\right] \cdot u_N - \left[\left(\Phi'^k \cdot \overline{r}_c^k \; \cdot \mathcal{F}_{\mathrm{N}}[u_c, d^k]\right) * h_z^*\right]_N\right)\Bigg) .\end{aligned} \tag{4.47}$$

This leads us to

$$d_p^{k+1,m+1} = \left(1 + 2\tau \cdot \alpha \cdot \sum_{\substack{\ell \in \\ \{x_1,x_2\}}} \sum_{\substack{q \in \\ \mathcal{N}_\ell(p)}} \frac{\Psi_p'^k + \Psi_q'^k}{2h_\ell^2} \cdot d_q^{k+1,m}\right) \cdot d_p^k$$
$$\cdot \left(1 + 2\tau \sum_{c \in \mathcal{C}} \left(\left[\left(\Phi'^k \cdot \overline{r}_c^k \right) * h_z^* \right]_1 \cdot u_1 - \left[\left(\Phi'^k \cdot \overline{r}_c^k \ \cdot \mathcal{F}_{\mathrm{N}}[u_c, d^k] \right) * h_z^* \right]_1 \right) \right.$$
$$\left. - \alpha \cdot d_p^k \sum_{\substack{\ell \in \\ \{x_1,x_2\}}} \sum_{\substack{q \in \\ \mathcal{N}_\ell(p)}} \frac{\Psi_p'^k + \Psi_q'^k}{2h_\ell^2} \right)^{-1}. \quad (4.48)$$

Since both systems of equations depend on each other, we perform an alternating minimisation strategy [Luenberger and Yinyu, 2015]. While solving the first problem (e.g. the estimation of the sharp image $\boldsymbol{u}$), the second one (e.g. the estimation of the depth d) remains fixed. After a fixed number of gradient descent steps, the roles are exchanged. Such a strategy is very common, for example in blind deconvolution problems [Chan and Wong, 1998], where both the sharp image and the blur kernel have to be estimated. To handle the problem of nonconvexity, we apply a coarse-to-fine strategy where the solution of the coarser level provides the initialisation for the next finer one [Bornemann and Deuflhard, 1996]. This strategy cannot remove the nonconvexity but at least allows a stable and repeatable solution of the system.

To guarantee the positivity of our solution under the condition that the result of the previous iteration step is positive, we have to restrict the relaxation parameter τ. An upper bound can be found by plugging-in the corresponding $\boldsymbol{D}_1$ (in dependency whether we estimate d or $\boldsymbol{u}$) into the formula (3.64).

4.5 Experiments

Now that we have discussed different strategies for depth-from-defocus methods and have illustrated how to discretise them as well as how to solve the associated PDE, in this section, let us compare the performance of these techniques. To this end, we first consider experiments on synthetic data which also allows a quantitative comparison. To substantiate the practical applicability of our methods, we perform further experiments on real-world data.

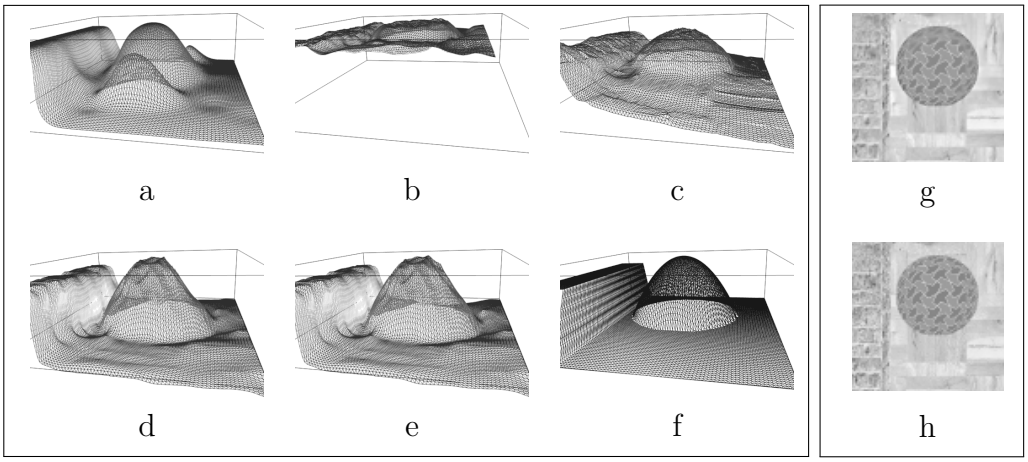

Figure 4.9: Comparison of different reconstruction methods. **Left box:** *In reading order* **(a)** Variance Method (VM) with a subsequent Gaussian smoothing step (patch-size = 6, $\sigma = 4.0$). **(b)** Without normalisation $\mathcal{F}_{\mathrm{U}}$, initialised with constant depth ($\alpha = 45$). **(c)** ditto, initialisation provided by the variance method. **(d)** Our normalised approach $\mathcal{F}_{\mathrm{N}}$, initialised with constant depth ($\alpha = 150$). **(e)** ditto initialisation provided by the variance method. **(f)** Ground truth of the depth profile. **Right box: (g) Top:** Estimated pinhole image. **(h) Bottom:** Ground truth of the pinhole image.

4.5.1 Error Measures

Before starting with synthetic experiments, let us first explain the applied error measures to assess the achieved reconstruction quality. In this work, we consider the *mean squared error* (MSE) as well as the *structure similarity error* (SSIM) of Wang et al. [2004].

Mean Squared Error. Regarding two discrete signals $\boldsymbol{u} \in \mathbb{R}^N$ and $\boldsymbol{f} \in \mathbb{R}^N$, each having a signal length N, the *mean squared error* (MSE) is given by the averaged squared differences in sample values between $\boldsymbol{u}$ and $\boldsymbol{f}$:

$$\mathrm{MSE}(\boldsymbol{u}, \boldsymbol{f}) := \frac{1}{N} \sum_{i=1}^{N} (u_i - f_i)^2 \,. \tag{4.49}$$

The MSE can be straightforwardly implemented and it is one of the most common error measure to judge the signal quality w.r.t. a reference signal. However, with regard to image processing, it does not incorporate any characteristics of the human visual system. Therefore, reconstructed signals having the same MSE may provide completely different visual quality.

Table 4.1: Error measurements. To compare the estimated pinhole image as well as the depth map against their ground truth, we consider the mean squared error (MSE) and the mean structural similarity (MSSIM) [Wang et al., 2004] as similarity metrics. We show the error of the pure variance method (VM) as well as the one with an additional post-smoothing step with variance σ. Further, the operator $\mathcal{F}_{\mathrm{U}}$ [Aguet et al., 2008] and our normalised imaging model $\mathcal{F}_{\mathrm{N}}$ is considered. The latter two are initialised once with a constant depth and once with an estimation of the VM.

Method		VM		$\mathcal{F}_{\mathrm{U}}$		Ours	
		$\sigma = 0$	$\sigma = 4$	const.	VM	const.	VM
Depth	**MSE**	2.83	1.10	21.66	2.26	0.77	**0.58**
	MSSIM	0.98	0.93	0.94	0.99	1.00	**1.00**
Image	**MSE**	48.33	46.38	124.52	52.17	49.09	**45.17**
	MSSIM	0.87	0.87	0.67	0.87	0.90	**0.92**

Structure Similarity Error. As its name already implies, with the *structure similarity error* (SSIM), Wang et al. [2004] judge the local structural similarity between two discrete signals $\boldsymbol{u} \in \mathbb{R}^N$ and $\boldsymbol{f} \in \mathbb{R}^N$ with respect to the human visual system. To this end, they compose the SSIM of three components:

$$\mathrm{SSIM}(\boldsymbol{u}, \boldsymbol{f}) := \underbrace{\left(\frac{2\mu_u\mu_f + C_1}{\mu_u^2 + \mu_f^2 + C_1}\right)^{\alpha}}_{\ell(\boldsymbol{u},\boldsymbol{f})} \cdot \underbrace{\left(\frac{2\sigma_u\sigma_f + C_2}{\sigma_u^2 + \sigma_f^2 + C_2}\right)^{\beta}}_{c(\boldsymbol{u},\boldsymbol{f})} \cdot \underbrace{\left(\frac{\sigma_{uf} + C_3}{\sigma_u\sigma_f + C_3}\right)^{\gamma}}_{s(\boldsymbol{u},\boldsymbol{f})}, \tag{4.50}$$

where μ_u, μ_f denote the mean intensities, σ_u,σ_f the standard deviations, and σ_{uf} the cross-covariance for signals $\boldsymbol{u}$ and $\boldsymbol{f}$ respectively within a certain patch. The constants are defined by $C_1 := (K_1 \cdot L)^2$, $C_2 := (K_2 \cdot L)^2$, and $C_3 := C2/2$, where L denotes the dynamic range of the images, i.e. 255 in our case and K_1, K_2 is suggested to be 0.01 and 0.03 respectively.

In this way, $\ell(\boldsymbol{u}, \boldsymbol{f})$ acts as a comparison of $\boldsymbol{u}$ and $\boldsymbol{f}$ w.r.t. the luminance, $c(\boldsymbol{u}, \boldsymbol{f})$ w.r.t. the contrast, and $s(\boldsymbol{u}, \boldsymbol{f})$ as a structure comparison. The weighting parameters α, β, γ are usually set to 1.

The range of SSIM lies between -1 and 1 where the latter one is only obtained if and only if both signal are identical. To obtain a single error value

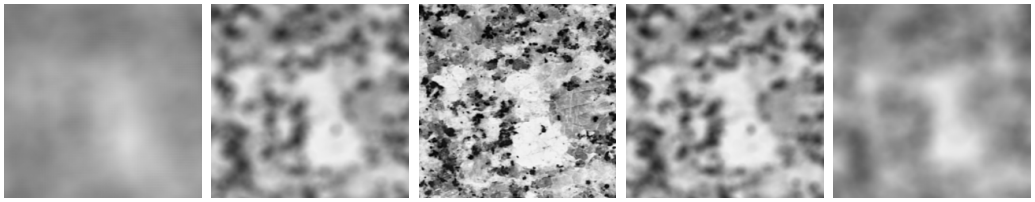

Figure 4.10: Equifocal case: Highly textured plane parallel to the lens. 5 out of 9 images of a focal stack rendered by a thin lens camera renderer. The third image is in focus.

for two images, one considers the *mean structure similarity error* (MSSIM):

$$\mathrm{MSSIM}(u, f) := \frac{1}{N} \sum_{i=1}^{N} \mathrm{SSIM}(\boldsymbol{u}_i, \boldsymbol{f}_i) \; . \tag{4.51}$$

where $\boldsymbol{u}_i, \boldsymbol{f}_i$ denote the i-th patches within $\boldsymbol{u}$ and $\boldsymbol{f}$ respectively.

4.5.2 Synthetic Data

In Section 4.1.4, Equation (4.10) illustrates the equivalence of the forward operator of Aguet et al. [2008] to standard 3-D convolution. The only requirement is that the sharp pinhole image is placed in a dark volume corresponding to the depth profile as defined in Equation (4.9). This aspect has also been confirmed by our comparison of different forward operators in Figure 4.6. As we have seen, both forward models do not preserve the maximum-minimum principle w.r.t. the intensities of the pinhole image at strong depth changes. Nevertheless, the natural question arises whether a 3-D deconvolution can be used to recover the sharp pinhole image from a recorded focal stack. Hence, our first experiment is devoted to exactly this question. To this end, we generate a focal stack with the help of a thin lens camera renderer with the following parameters: lens diameter $D = 2.69$ cm, lens distance to image plane $v = 35$ mm. The distance of the focal plane to the lens varies equidistantly from $f_p = 3$ cm to $f_p = 11$ cm. In this way, we produce 9 images of size 250×250 in total, 5 of them are shown in Figure 4.10. For this experiment, a very simple 3-D model which consists only of a highly textured equifocal plane at distance $d(\boldsymbol{x}) = 7$ cm is entirely sufficient. We deblur the generated focal stack with variational deconvolution from Section 2.5.3. To this end, we consider Equation (2.53) for the three-dimensional case where the 3-D PSF is given by h. Figure 4.11(a) demonstrates that variational 3-D deconvolution is not able to reconstruct the sharp slice in a reasonable way.

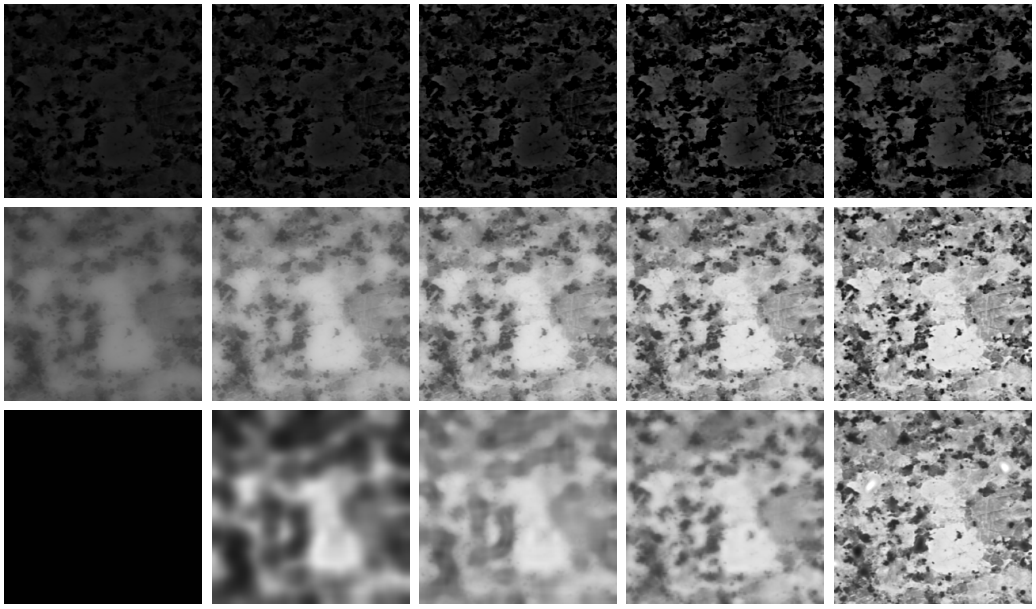

Figure 4.11: Variational 3-D deconvolution applied to the focal stack of Figure 4.10. The slice corresponding to the true depth $d(\boldsymbol{x}) = 7$ cm with increasing gradient descent steps (*from left to right*) is shown. **(a) Top:** Without further assumptions. **(b) Middle:** Additionally all slices not corresponding to the depth profile are set to zero. **(c) Bottom:** Our method.

This is because the standard 3-D deconvolution does not incorporate the fact that the focal stack has to originate from a dark volume with only a single sharp slice according to a depth profile. However, if we incorporate depth information, e.g. by setting all values to zero that do not correspond to the actual depth in each iteration, we obtain the result in Figure 4.11(b). Of course this is not a practical solution since it requires knowledge of the correct depth. Since such an approach is not useful even in this simple scenario, we do not consider the standard 3-D deconvolution any further. Figure 4.11(c) shows the result of our method from Section 4.2.1. Here, we have used a coarse-to-fine approach (unidirectional multigrid [Bornemann and Deuflhard, 1996]) where the coarsest grid is initialised with a constant depth $d(\boldsymbol{x}) = 3$ cm and constant texture intensity $\boldsymbol{u}^0 = 1$.

Our second experiment addresses the direct comparison of our novel depth-from-defocus approach against the variance method (VM) as well as the approach of Aguet et al. [2008]. The variance method belongs to *depth-from-focus* approaches. First, for each slice of the focal stack, the VM computes the local variances of the intensities in a patch-based manner. Next,

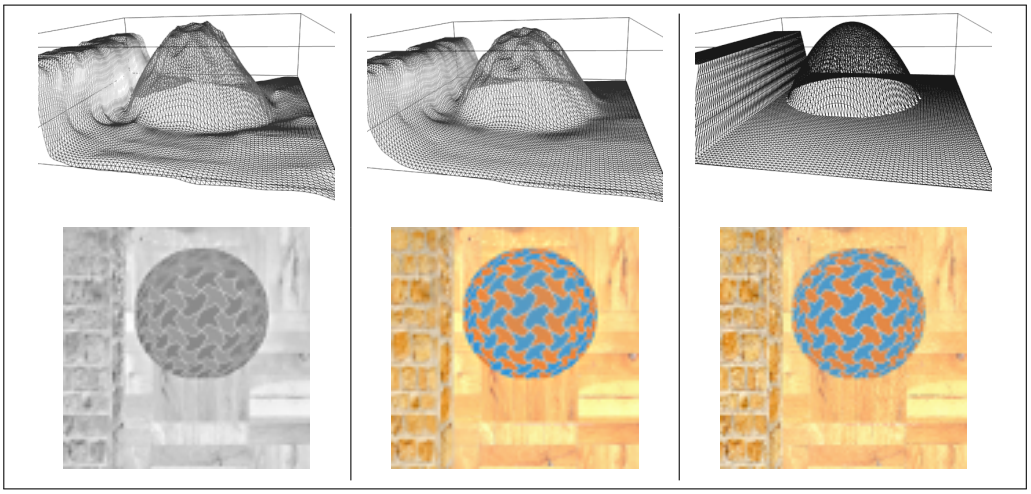

Figure 4.12: Visual comparison. **(a) Left:** Single channel approach. **(b) Middle:** Incorporating all channels of an RGB focal stack. **(c) Right:** Ground truth.

locally, the slice with the highest variance is assumed to represent the in-focus slice and thus to describe the relative depth value. In this way, we investigate the impact of our normalisation strategy and compare both approaches against a depth-from-focus method. For this experiment, we use a 3-D model that is a bit more complex than in our first experiment. It is shown in Figure 4.6(a). We render this scene with the thin lens camera renderer where the camera is placed perpendicular above the model. The optical settings remain the same as in the first experiment and the focal plane moves equidistantly from $f_p = 3$ cm to $f_p = 7$ cm to render a focal stack of 20 images. In this experiment, we restrict ourselves to the one-channel case since we only want to demonstrate the impact of normalisation here. Figure 4.6(e) shows 3 different slices of the created focal stack. The local variation of blur between each slice is clearly recognisable.

Besides that, Figure 4.9 also provides the reconstruction results of the different approaches. The result of the variance method is shown in Figure 4.9(a). As we can see, the reconstruction of the depth suffers from two undesired hills in front of and behind the hemisphere. That is because the VM misinterprets the blur circle of the hemisphere as a higher local sharpness than the flat contrast of the textured floor which corresponds to the actual slice being in focus.

The consequence of ignoring normalisation can be clearly seen in Figure 4.9(b) and (c). Applying the forward operator $\mathcal{F}_U$ to a depth map produces severe over- and undershoots at strong depth changes. This is a direct con-

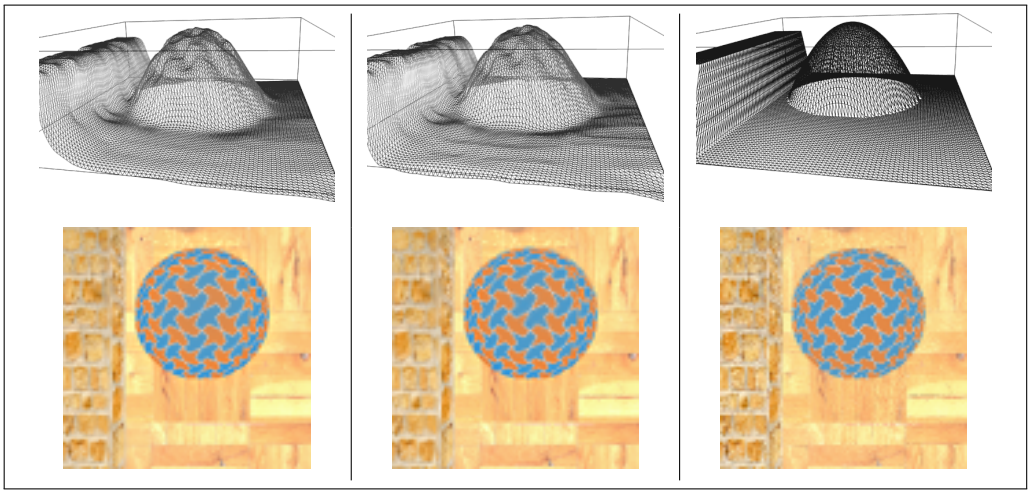

Figure 4.13: **(a) Left:** Without robustification. **(b) Middle:** With robustification. **(c) Right:** Ground truth.

sequence of ignoring the requirement in (4.7) that leads to the discussed violation of the imaging model. This in turn implies that keeping strong depth changes in the inverse operation would increase the residual error drastically at those locations since such over- and undershoots do not occur in natural recordings. Thus, when minimising the residual error with $\mathcal{F}_U$ as forward operator, strong changes of the depth are avoided and smooth ones are preferred. This can be seen as an unwanted regularisation of the depth reconstruction implied by the forward operator. Furthermore, comparing Figure 4.9(b) and (c), one observes that the result is strongly affected by the initialisation. While in the first one a constant depth map is used, the second one was initialised with the result of the variance method. Due to the strong regularisation implied by $\mathcal{F}_U$, initialising with a constant depth, the method is not able to converge to a reasonable solution. In contrast to that,

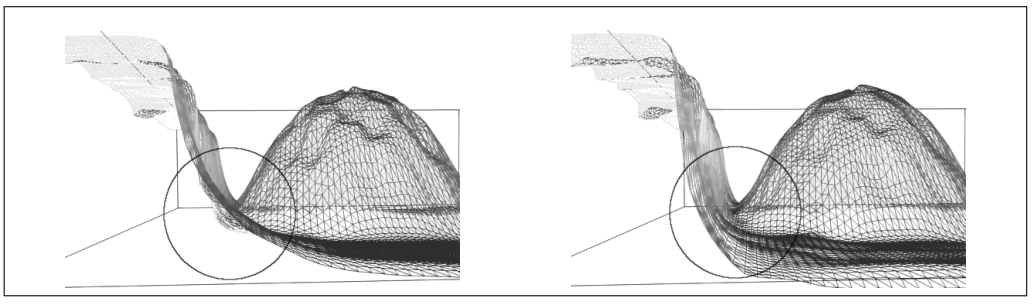

Figure 4.14: **(a) Left:** Without robustification. **(a) Right:** With robustification.

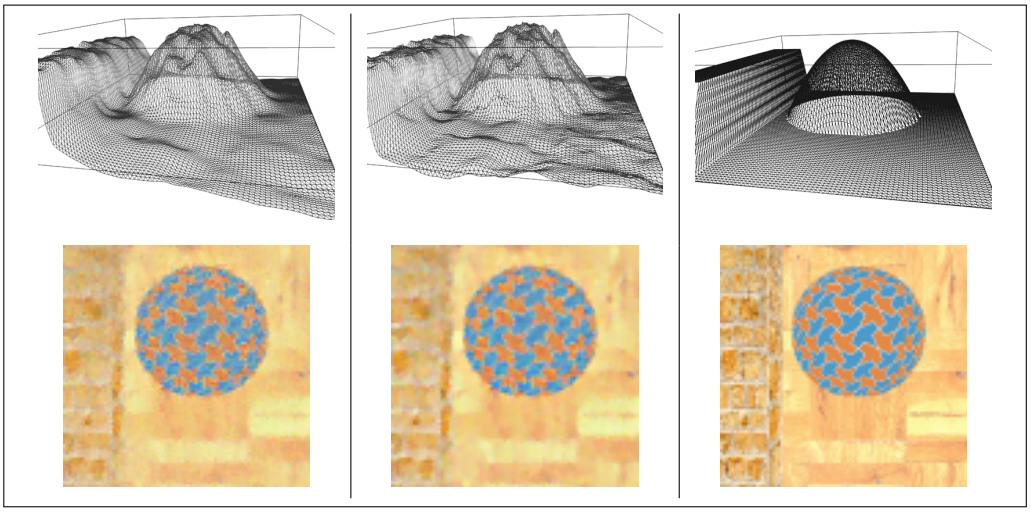

Figure 4.15: **(a) Left:** Sequential approach. **(b) Middle:** Joint approach. **(c) Right:** Ground truth.

our forward operator $\mathcal{F}_{\mathrm{N}}$ approximates the thin lens camera model very well. With its incorporated normalisation function, the requirement in (4.7) is locally preserved leading to a guarantee of the maximum-minimum principle. Thus, our forward operator comes much closer to the physical imaging process, especially at strong depth changes. Consequently, it is better suitable for variational depth-from-defocus approaches. Embedded into a variational framework, our forward operation improves the estimation of the unknown depth as well as the sharp pinhole image substantially. Indeed, as we can see in the Figures 4.9(d) and (e), the hemisphere as well as the strong depth change at the wall are well reconstructed and no smoothing effect implied by the forward operator exists. Besides that, our reconstruction does not suffer from misinterpretations like the one of the variances method. Also regarding the reconstruction of the sharp pinhole image shown in Figure 4.9(g), our results match the ground truth more closely. This can also be seen quantitatively in Table 4.1. Moreover, the initialisation of our approach does not affect the solution severely.

With the third synthetic experiment we investigate the benefit of incorporating all channels given a multi-channel focal stack. To this end, we apply our algorithm with data term E_{D_2} (cf. Equation 4.30) to an RGB version of the focal stack from the last experiment (see Figure 4.7(a)). In Figure 4.12 we compare the multi-channel to the single channel approach. As one can recognise, incorporating the information of all channels not only leads to a more appealing and more accurate estimation of the sharp pinhole image,

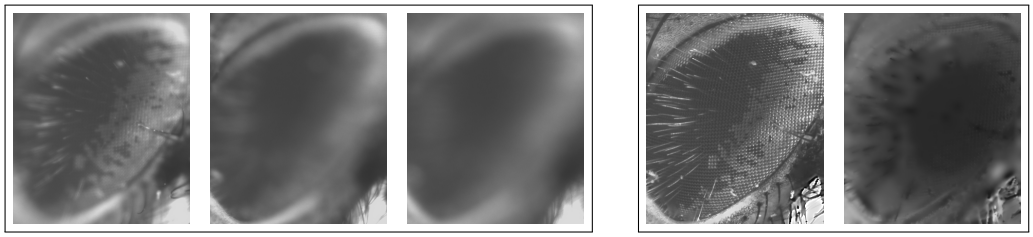

Figure 4.16: Focal stack of a house fly eye (grey scaled). This focal stack was provided by the Biomedical Imaging Group EPFL, Lausanne, Switzerland. **(a) Left box:** 3 of 21 images of the focal stack. **(b) Right box:** Recovered pinhole image and depth profile ($\alpha = 25$).

but also the reconstruction of the depth map improved. The quantitative comparison given in Table 4.2 confirms this visual impression.

In our cell enhancement framework from Chapter 3, the reconstruction quality could be clearly improved by robustification strategies that penalise strong deviations from the model assumption less severely. Hence, the natural question comes up if such a strategy also improves the performance of our depth-from-defocus approach. Therefore, the fourth experiment compares the results using a quadratic data term E_{D_2} against the robust data term E_{D_3} (cf. Equation (4.32)). In Figure 4.13 both results are shown. The data term E_{D_3} gives a higher reconstruction quality, especially at the strong depth change at the wall (cf. Figure 4.14). In Table 4.2 the results with different data terms are summarised. We see that both modifications (colour and robustification) lead to better depth estimates. In this experiment, the parameters are tuned w.r.t the depth reconstruction. For all approaches, the quality of the pinhole image is visually similar. If the main focus lies on the quality of the pinhole image, the parameters can be adapted to this end.

With the last experiment on synthetic data, we want to demonstrate the potential of our novel joint denoising and depth-from-defocus approach. Therefore, we first add artificially created Gaussian noise with zero mean and standard deviation $\sigma_{\text{noise}} = 30$ (cf. Figure 4.8) according to Equation (1.1) to the focal stack of the second experiment. As a baseline for comparison, we consider a sequential framework where each slice of the focal stack is denoised by image restoration according to Equation (4.35) in advance before performing depth-from-defocus. The comparisons in Figure 4.15 and Table 4.3 show that our novel joint approach outperforms the sequential one qualitatively and quantitatively.

Table 4.2: Quantitative comparison: Colour and robustification. Again mean squared error (MSE) and the mean structural similarity (MSSIM) [Wang et al., 2004] are used. We consider the benefit of incorporating all image channels as well as the influence of a robustification. We use the sharp grey scale pinhole image as ground truth. Therefore, we convert the reconstructed sharp colour image to grey scale before measuring the MSE and MSSIM.

Method		Grey value	Colour	Colour + Robust
		E_{D}	E_{D_2} $(\alpha = 780)$	E_{D_3} $(\alpha = 60)$
Depth	**MSE**	0.58	0.23	**0.18**
	MSSIM	0.9973	0.9985	**0.9986**
Image	**MSE**	45.17	**33.73**	39.79
	MSSIM	0.9175	**0.9301**	0.9167

4.5.3 Real-World Data

In the last section, we have demonstrated that all of the devised concepts are well suited for the depth-from-defocus task and each one clearly improves the reconstruction results. However, we have restricted our considerations only on synthetic data, so far. While such experiments offer the advantage of having a ground truth and therewith offer the possibility to judge the results quantitatively, it is also important to consider the results on focal stacks captured by a real optical system. Besides demonstrating the practical applicability of the method, this shows the ability to handle characteristics of the imaging process that are not incorporated in the model assumption. Here, especially the impact of a real PSF against an estimated one comes into play.

To this end, we consider in our next experiment a real-world focal stack showing a house fly eye. This stack consists of 21 slices where 3 of them are shown in Figure 4.16(a). For this experiment, our approach uses a coarse-to-fine strategy [Bornemann and Deuflhard, 1996]. On the coarsest grid the method is initialised with the variance method. Figure 4.16(b) shows the estimated sharp pinhole image along with the recovered depth map of the scene. The determined pinhole image is reconstructed well and the depth-of-field appears infinite: the small hairs in the front as well as the compound eye are entirely sharp. The level of detail can also be seen in the depth map. Also here small structures such as hairs are clearly recognisable. For the depth map a grey value coding is used: the brighter the grey value, the

Table 4.3: Quantitative comparison of sequential and joint approach. To have a better comparison against previous experiments, we again convert the reconstructed colour results to grey scale before measuring the MSE and MSSIM.

Method		Seq. ($\alpha = 500$, $\gamma = 20$)	Joint ($\alpha = 350$, $\beta = 2.0$)
Depth	**MSE**	0.52	**0.32**
	MSSIM	0.9971	**0.9976**
Image	**MSE**	110.54	**103.18**
	MSSIM	0.7421	**0.7649**

Figure 4.17: Focal stack of a coffee bean. This stack was provided by the Computer Graphics Group, MPI for Informatics, Saarbrücken, Germany. **(a) Left box:** 3 of 22 images of the focal stack. **(b) Right box:** Estimated pinhole image and reconstructed depth map ($\alpha = 20$).

larger the depth.

For the second real-world experiment, we apply our robust approach to an RGB focal stack capturing a coffee bean. The focal stack consists of 22 frames where 3 of them are depicted in Figure 4.17(a). Also for this experiment, a coarse-to-fine strategy is used. This time, the challenge is made slightly more difficult: We refrain from the variance method as initialisation on the coarsest grid, but use a constant depth value instead. The results are shown in Figure 4.17(b).

4.6 Summary

In this chapter, we have illustrated how to exploit the physical effect of a limited depth-of-field to infer jointly the depth map of the scene along with the sharp image. To this end, we first made a detailed investigation of different imaging processes. Next, we showed how to model a suitable forward operator and demonstrated the improvement of existing approaches by in-

corporating essential physical characteristics such as the maximum-minimum principle. To accomplish the ill-posed inversion, we advocate variational methods. Our experiments show that our novel forward operator achieves a depth-of-field simulation very close to that of the thin lens camera model constituting a physical model to describe such an effect.

Besides an improved forward operator, we have demonstrated that the benefits of respecting physical constraints generally carries over to a significantly better reconstruction quality in the depth-from-defocus problem, especially at strong depth changes. In the context of modelling, we showed once more that robustification can be a fruitful strategy to limit the influence of remaining imperfections in the model assumptions. Moreover, we made full usage of the colour information and proposed a joint denoising and depth-from-defocus approach.

In finding a suitable minimiser, as in the previous chapter, we have advocated to replace the traditional Euler-Lagrange formalism by a multiplicative variant. Besides the plausible positivity preservation, this allowed us to follow a more efficient semi-implicit iteration strategy.

This chapter illustrate the advantages of a physically refined modelling in combination with appropriate mathematics. Again, the superiority of a joint handling of different reconstruction task in contrast to a separate one has been demonstrated successfully.

Chapter 5

Summary and Outlook

5.1 Summary

In this dissertation, we have pursued the strategy of interpreting reconstruction as the inversion of the physical imaging processes. To this end, we showed the way from a detailed analysis of the image acquisition technique to the modelling of a suitable forward operator. Such a forward operator serves as a mathematically sound formulation of the physical imaging process and should allow a reasonably accurate simulation. While in general, the forward process is well-posed, its inversion confronts us with an ill-posed problem. For this purpose, we have made use of variational methods. The aim was the combination of an adequate forward operator with a suitable functional where we focused the tailoring to the peculiarities and deficiencies of the considered image acquisition technique and the exploitation of physical principles. To clarify our ideas, we considered two concrete examples. While the first one lies in the area of image processing, the second one falls into the scope of computer vision.

In the first part of this dissertation, we developed a reconstruction framework aiming at the enhancement of 3-D cell images recorded with confocal laser scanning (CLSM) and stimulated emission depletion (STED) microscopy techniques. While such microscopes provide a very high spatial resolution – and in the case of STED a resolution beyond the diffraction limit – they have to cope with very low light intensities. Therefore, the recordings suffer from blur and Poisson distributed noise. Besides that, as they belong to the class of optical 3-D microscopes, the captured images provide only a relatively low axial resolution. Moreover, working with fluorescence dyes, the region of interest may suffer from a defect labelling. To handle all these issues, we presented a joint model that counteracts these prob-

lems simultaneously: By choosing an adequate forward operator in combination with Csiszár's information divergence [Csiszár, 1991], we incorporated Richardson-Lucy [Richardson, 1972; Lucy, 1974] deconvolution as an appropriate deblurring technique under Poisson distributed noise. Ill-posedness of deconvolution and remaining imperfections in the model assumptions, as well as in the estimation of the point-spread functions, have led us to the robust and regularised variant of RL deconvolution [Welk, 2010]. The relatively low z-resolution has been counteracted by extending the regularisation domain according to the approach of Elhayek et al. [2011]. Also with respect to a defect labelling, we incorporated anisotropic diffusion. To this end, the isotropic divergence term has been replaced by an anisotropic one which is also well suited for inpainting. Moreover, its anisotropy has been designed in such a way that it enhances the filament structures of cells. As result, our anisotropic model yields a higher reconstruction quality than its isotropic predecessor from Elhayek et al. [2011].

In the second part of this dissertation, we presented a novel variational depth-from-defocus approach. There, a detailed discussion of different forward operators describing depth-of-field simulations brought us to the necessity of incorporating specific physical principles. We demonstrated that ignoring principles such as the maximum-minimum principle w.r.t. the intensity values not only results in a wrong model assumption in the forward direction but may also erroneously implicitly regularise the estimation process of the depth values. The result was an insufficient reconstruction especially at strong depth changes. To handle this problem, we advocated a novel forward operator. It incorporates a normalisation function and thereby inherently guarantees the maximum-minimum principle. Our novel forward operator provides a simulation very close to that of the thin lens camera model and is designed in such a way that it fits well into a variational framework. In that context, we illustrated the benefit of robustification ideas and the full integration of colour information. All these ideas successively ameliorated the reconstruction quality. As a further contribution, regarding the handling of noisy focal stacks, we supplemented our method to a novel joint variational depth-from-defocus and denoising approach.

Besides addressing modelling aspects in this dissertation, we also consider adequate variational minimisation strategies. With the help of the multiplicative Euler-Lagrange formalism we explained the tailoring of Richardson-Lucy deconvolution [Richardson, 1972; Lucy, 1974] to Poisson statistics and its relation to Csiszár's information divergence [Csiszár, 1991]. Besides that, in contrast to the well-known classical additive formalism, the multiplicative variant constrains the solution to the positive range. This is reasonable,

since both the number of arriving photons at an imaging sensor and the values of the depth map are positive. Moreover, it allowed us to develop more efficient gradient descent schemes also offering better stability. This was realised by using explicit and implicit terms in a more powerful way.

5.2 Conclusions and Future Work

This thesis advocates the incorporation of physical principles into the task of reconstruction. The more such laws are respected, the better the resulting quality could potentially be. However, accurate physics most often comes at the cost of computational expense. To date, it is just not feasible to invert any physically exact forward operator. Hence, this thesis balances these opposite requirements and presents two approaches that at the same time incorporate more physics and are efficiently solvable. Eventually, the attained reconstruction performance could only be realised through the harmony of appropriate physics and established mathematical concepts.

Obviously, one goal of future work should be the development of even more realistic forward operators. On the one hand, this means the incorporation of additional physical principles and physical refinements of the considered model assumptions, respectively. For instance, an important problem that has not been solved to date are partial occlusions. While our normalisation and robustification efforts limit their influence, an explicit incorporation of occlusion handling into the operator might perform much better. On the other hand, generally, the big challenge will be to manage the proper minimisation of such models. In order to make significant progress, however, a novel strategy in that respect will be indispensable.

A further example concerns the estimation of the point-spread function (PSF). Until now, as discussed, e.g. in our cell reconstruction approach, our framework depends on an external estimation of the PSF. Here, it can be advantageous to involve this estimation process directly into our framework. In that respect, we also believe that a joint semi-blind approach that optimises the PSF based on bead measurements, should be advantageous and further improve the results.

Own Publications

Persch, N., Elhayek, A., Welk, M., Bruhn, A., Grewenig, S., Böse, K., Kraegeloh, A., and Weickert, J. (2013). Enhancing 3-D cell structures in confocal and STED microscopy: a joint model for interpolation, deblurring and anisotropic smoothing. *Measurement Science and Technology*, 24(12):125703.

Persch, N., Schroers, C., Setzer, S., and Weickert, J. (2014). Introducing more physics into variational depth-from-defocus. In Jiang, X., Hornegger, J., and Koch, R., editors, *Pattern Recognition*, volume 8753 of *Lecture Notes in Computer Science*, pages 15–27. Springer, Berlin.

Persch, N., Schroers, C., Setzer, S., and Weickert, J. (2015). Physically inspired depth-from-defocus. Technical Report 355, Department of Mathematics, Saarland University, Saarbrücken, Germany. Submitted to *Image and Vision Computing* Journal on Feb. 3, 2015.

Bibliography

Abbe, E. (1873). Beiträge zur Theorie des Mikroskops und der Mikroskopischen Wahrnehmung. *Archiv für mikroskopische Anatomie*, Bd. 9:413–468. (Cited on pages 1, 45, 46, 47, and 48)

Aguet, F., Van De Ville, D., and Unser, M. (2008). Model–based 2.5-D deconvolution for extended depth of field in brightfield microscopy. *IEEE Transactions on Image Processing*, 17(7):1144–1153. (Cited on pages 88, 89, 90, 94, 95, 96, 97, 98, 105, 115, 116, and 117)

Alvarez, L., Deriche, R., Sánchez, J., and Weickert, J. (2002). Dense disparity map estimation respecting image derivatives: A PDE and scale-space based approach. *Journal of Visual Communication and Image Representation*, 13(1/2):3–21. (Cited on page 6)

Ambrosio, L. and Tortorelli, V. (1992). Approximation of functionals depending on jumps by elliptic functionals via Γ-convergence. *Bolletino della Unione Matematica Italiana*, 7:105–123. (Cited on page 10)

Anconelli, B., Bertero, M., Boccacci, P., Carbillet, M., and Lanteri, H. (2005). Restoration of interferometric images - III. Efficient Richardson-Lucy methods for LINC-NIRVANA data reduction. *Astronomy & Astrophysics*, 430(2):731–738. (Cited on page 52)

Ango, P. (1682). *L' optique*. Chez Estienne Michallet. (Cited on pages 1 and 46)

Aubert, G. and Kornprobst, P. (2006). *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*, volume 147 of *Applied Mathematical Sciences*. Springer, New York, second edition. (Cited on pages 7, 10, 86, and 99)

Ayers, G. and Dainty, J. C. (1988). Iterative blind deconvolution method and its applications. *Optics letters*, 13(7):547–549. (Cited on page 34)

Bailey, S. W., Echevarria, J. I., Bodenheimer, B., and Gutierrez, D. (2014). Fast depth from defocus from focal stacks. *The Visual Computer*, 31:1–12. (Cited on pages 87 and 88)

Bar, L., Sochen, N., and Kiryati, N. (2005). Image deblurring in the presence of salt-and-pepper noise. In Kimmel, R., Sochen, N., and Weickert, J., editors, *Scale Space and PDE Methods in Computer Vision*, volume 3459 of *Lecture Notes in Computer Science*, pages 107–118. Springer, Berlin. (Cited on pages 11, 60, and 107)

Barsky, B. A., Horn, D. R., Klein, S. A., Pang, J. A., and Yu, M. (2003). Camera models and optical systems used in computer graphics: Part I, object-based techniques. In Kumar, V., Gavrilova, M., Tan, C. J. K., and L' Ecuyer, P., editors, *Computational Science and Its Applications - ICCSA 2003*, volume 2669 of *Lecture Notes in Computer Science*, pages 246–255. Springer, Berlin. (Cited on pages 85 and 92)

Barsky, B. A. and Kosloff, T. J. (2008). Algorithms for rendering depth of field effects in computer graphics. In *Proc. WSEAS International Conference on Computers*, pages 999–1010, Heraklion, Greece. World Scientific and Engineering Academy and Society. (Cited on pages 85 and 89)

Ben-Ari, R. and Raveh, G. (2011). Variational depth from defocus in real-time. In *IEEE International Conference on Computer Vision (ICCV) Workshops*, pages 522–529, Barcelona, Spain. IEEE. (Cited on page 88)

Ben Hadj, S., Blanc-Féraud, L., and Aubert, G. (2014). Space variant blind image restoration. *SIAM Journal on Imaging Sciences*, 7(4):2196–2225. (Cited on page 43)

Ben Hadj, S., Blanc-Féraud, L., Aubert, G., and Engler, G. (2013). Blind restoration of confocal microscopy images in presence of a depth-variant blur and Poisson noise. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 915–919, Vancouver, Canada. IEEE. (Cited on page 43)

Bertalmío, M., Sapiro, G., Caselles, V., and Ballester, C. (2000). Image inpainting. In *Proc. SIGGRAPH*, pages 417–424, New Orleans, USA. ACM. (Cited on pages 16 and 43)

Bertero, M., Boccacci, P., Desiderà, G., and Vicidomini, G. (2009). Image deblurring with Poisson data: from cells to galaxies. *Inverse Problems*, 25(12):123006. (Cited on pages 38 and 66)

Bertero, M., Poggio, T. A., and Torre, V. (1988). Ill-posed problems in early vision. *Proceedings of the IEEE*, 76(8):869–889. (Cited on pages 10, 15, 86, 99, and 7)

Bhasin, S. and Chaudhuri, S. (2001). Depth from defocus in presence of partial self occlusion. In *Proc. IEEE International Conference on Computer Vision*, volume 1, pages 488–493, Vancouver. (Cited on page 88)

Binnig, G. and Rohrer, H. (1983). Scanning tunneling microscopy. *Surface science*, 126(1):236–244. (Cited on pages 1 and 45)

BMBF (2010). Nano-initiative action plan. *German Federal Ministry of Education and Research*. (Cited on page 41)

Born, M. and Wolf, E. (1970). *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Pergamon Press, Oxford, 4th edition. (Cited on page 91)

Bornemann, F. and Deuflhard, P. (1996). The cascadic multigrid method for elliptic problems. *Numerische Mathematik*, 75:135–152. (Cited on pages 113, 117, and 122)

Boshtayeva, M., Hafner, D., and Weickert, J. (2015). A focus fusion framework with anisotropic depth map smoothing. *Pattern Recognition*, 48(11):3310 – 3323. (Cited on page 87)

Bovik, A. (2009). *The Essential Guide to Image Processing*. Elsevier Science, Philadelphia. (Cited on pages 3, 4, and 50)

Bracewell, R. N. (1999). *The Fourier Transform and its Applications*. McGraw-Hill, New York. (Cited on pages 12, 16, 35, 69, and 111)

Bradshaw, R. and Stahl, P. (2015). *Encyclopedia of Cell Biology*. Elsevier Science, Philadelphia. (Cited on page 47)

Bratsolis, E. and Sigelle, M. (2001). A spatial regularization method preserving local photometry for Richardson-Lucy restoration. *Astronomy & Astrophysics*, 375(3):1120–1128. (Cited on pages 38, 42, 50, and 75)

Brox, T., Bruhn, A., Papenberg, N., and Weickert, J. (2004). High accuracy optical flow estimation based on a theory for warping. In Pajdla, T. and Matas, J., editors, *Computer Vision – ECCV 2004*, volume 3024 of *Lecture Notes in Computer Science*, pages 25–36. Springer, Berlin. (Cited on page 10)

Bruhn, A., Weickert, J., and Schnörr, C. (2005). Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. *International Journal of Computer Vision*, 61(3):211–231. (Cited on page 11)

Brune, C., Sawatzky, A., and Burger, M. (2011). Primal and dual Bregman methods with application to optical nanoscopy. *International Journal of Computer Vision*, 92(2):211–229. (Cited on page 42)

Cannon, J. (1984). *The One-Dimensional Heat Equation*. Encyclopedia of Mathematics and its Applications. Cambridge University Press. (Cited on page 18)

Carslaw, H. S. and Jaeger, J. C. (1959). *Conduction of Heat in Solids*. Oxford University Press, Oxford, second edition. (Cited on page 15)

Chan, T. F. and Mulet, P. (1999). On the convergence of the lagged diffusivity fixed point method in total variation image restoration. *SIAM Journal on Numerical Analysis*, 36(2):354–367. (Cited on pages 20, 67, 71, and 110)

Chan, T. F. and Shen, J. (2001). Non-texture inpainting by curvature-driven diffusions (CDD). *Journal of Visual Communication and Image Representation*, 12(4):436–449. (Cited on page 16)

Chan, T. F. and Shen, J. (2002). Mathematical models for local nontexture inpaintings. *SIAM Journal on Applied Mathematics*, 62:1019–1043. (Cited on pages 16, 31, 33, and 43)

Chan, T. F. and Vese, L. A. (2001). Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277. (Cited on page 10)

Chan, T. F. and Wong, C. K. (1998). Total variation blind deconvolution. *IEEE Transactions on Image Processing*, 7:370–375. (Cited on pages 6, 10, 34, 43, 109, and 113)

Chan, T. F., Yip, A. M., and Park, F. E. (2005). Simultaneous total variation image inpainting and blind deconvolution. *International Journal of Imaging Systems and Technology*, 15:92–102. (Cited on pages 43 and 61)

Charbonnier, P., Blanc-Féraud, L., Aubert, G., and Barlaud, M. (1994). Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proc. IEEE International Conference on Image Processing*, volume 2, pages 168–172. (Cited on page 15)

Charbonnier, P., Blanc-Féraud, L., Aubert, G., and Barlaud, M. (1997). Deterministic edge–preserving regularization in computed imaging. *IEEE Transactions on Image Processing*, 6(2):298–311. (Cited on pages 19, 22, 23, 25, 27, 28, 31, 33, 34, 38, 39, 65, and 40)

Chaudhuri, S. and Rajagopalan, A. (1999). *Depth From Defocus: A Real Aperture Imaging Approach*. Springer, Berlin. (Cited on page 88)

Cittert, P. H. (1931). Zum Einfluß der Spaltbreite auf die Intensitätsverteilung in Spektrallinien. II. *Zeitschrift für Physik*, 69:298–308. (Cited on page 42)

Conn, P. (2012). *Imaging and Spectroscopic Analysis of Living Cells: Optical and Spectroscopic Techniques*. Academic Press. Elsevier/Academic Press. (Cited on page 2)

Cook, R. L., Porter, T., and Carpenter, L. (1984). Distributed ray tracing. In *Computer Graphics*, SIGGRAPH, pages 137–145, Minneapolis. ACM. (Cited on pages 89 and 92)

Cottet, G.-H. and El Ayyadi, M. (1998). A Volterra type model for image processing. *IEEE Transactions on Image Processing*, 7(3):292–303. (Cited on page 25)

Csiszár, I. (1991). Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems. *Annals of Statistics*, 19(4):2032–2066. (Cited on pages 11, 42, 51, 56, 82, 88, 98, 110, 126, 10, and 99)

Cussler, E. (1997). *Diffusion: Mass Transfer in Fluid Systems*. Cambridge Series in Chemical Engineering. Cambridge University Press. (Cited on pages 15 and 17)

Davies, A. (2012). *Close-Up and Macro Photography*. The Focus On Series. Taylor & Francis, Oxford. (Cited on page 85)

De Monvel, J. B., Le Calvez, S., and Ulfendahl, M. (2001). Image restoration for confocal microscopy: improving the limits of deconvolution, with application to the visualization of the mammalian hearing organ. *Biophysical Journal*, 80(5):2455–2470. (Cited on page 42)

Demoment, G. (1989). Image reconstruction and restoration: Overview of common estimation structures and problems. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(12):2024–2036. (Cited on page 15)

Dey, N., Blanc-Féraud, L., Zimmer, C., Kam, Z., Roux, P., Olivo-Marin, J., and Zerubia, J. (2006). Richardson-Lucy algorithm with total variation regularization for 3D confocal microscope deconvolution. *Microscopy Research Technique*, 69:260–266. (Cited on pages 38, 42, 50, 51, 52, 59, 66, 75, and 82)

Dey, N., Blanc-Féraud, L., Zimmer, C., Roux, P., Kam, Z., Olivo-Marin, J., and Zerubia, J. (2004). 3D microscopy deconvolution using Richardson-Lucy algorithm with total variation regularization. Research Report RR-5272, INRIA, Paris. (Cited on pages 42, 51, 52, 59, 66, 75, and 82)

Di Zenzo, S. (1986). A note on the gradient of a multi-image. *Computer Vision, Graphics and Image Processing*, 33:116–125. (Cited on page 23)

Dijksterhuis, F. J. (2006). *Lenses and waves: Christiaan Huygens and the mathematical science of optics in the seventeenth century*, volume 9. Springer, Berlin. (Cited on pages 1 and 46)

Dongmo, S., Troyon, M., Vautrot, P., Delain, E., and Bonnet, N. (1996). Blind restoration method of scanning tunneling and atomic force microscopy images. *Journal of Vacuum Science & Technology B*, 14(2):1552–1556. (Cited on page 2)

Einstein, A. (1905). Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik*, 322(6):132–148. (Cited on page 1)

Elhayek, A., Welk, M., and Weickert, J. (2011). Simultaneous interpolation and deconvolution model for the 3-D reconstruction of cell images. In Mester, R. and Felsberg, M., editors, *Pattern Recognition*, volume 6835 of *Lecture Notes in Computer Science*, pages 316–325. Springer, Berlin. (Cited on pages 8, 13, 42, 49, 53, 56, 59, 60, 61, 63, 75, 76, 77, 78, 79, 80, 81, 82, and 126)

Evans, G., Blackledge, J., and Yardley, P. (1999). *Numerical Methods for Partial Differential Equations*. Springer, Berlin. (Cited on page 15)

Farahani, J. N., Schibler, M. J., and Bentolila, L. A. (2010). Stimulated emission depletion (STED) microscopy: from theory to practice. In Méndez-Vilas, A. and Díaz, J., editors, *Microscopy: Science, Technology, Applications and Education*, volume 2, pages 1539–1547. Formatex Research Center, Badajoz, Spain. (Cited on page 46)

Favaro, P. (2010). Recovering thin structures via nonlocal-means regularization with application to depth from defocus. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 1133–1140, San Francisco. (Cited on page 88)

Favaro, P., Mennucci, A., and Soatto, S. (2003a). Observing shape from defocused images. *International Journal of Computer Vision*, 52(1):25–43. (Cited on page 86)

Favaro, P., Osher, S., Soatto, S., and Vese, L. (2003b). 3D shape from anisotropic diffusion. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 179–186, Madison, USA. (Cited on page 88)

Favaro, P. and Soatto, S. (2000). Shape and radiance estimation from the information divergence of blurred images. In Vernon, D., editor, *Computer Vision – ECCV 2000*, volume 1842 of *Lecture Notes in Computer Science*, pages 755–768. Springer, Berlin. (Cited on pages 86, 88, 110, and 89)

Favaro, P., Soatto, S., Burger, M., and Osher, S. (2008). Shape from defocus via diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):518–531. (Cited on page 88)

Feng, X. and Prohl, A. (2002). Analysis of total variation flow and its finite element approximations. *ESAIM: Mathematical Models and Methods in the Applied Sciences*, 37(3):533–556. (Cited on page 20)

Fick, A. (1855). Über Diffusion. *Annalen der Physik*, 170(1):59–86. (Cited on pages 15, 16, and 17)

Figueiredo, M. A. and Bioucas-Dias, J. M. (2009). Deconvolution of Poissonian images using variable splitting and augmented Lagrangian optimization. In *IEEE/SP 15th Workshop on Statistical Signal Processing*, pages 733–736, Cardiff. (Cited on page 42)

Fish, D. A., Brinicombe, A. M., Pike, E. R., and Walker, J. G. (1995). Blind deconvolution by means of the Richardson-Lucy algorithm. *Journal of the Optical Society of America A*, 12(1):58–65. (Cited on page 34)

Förstner, W. and Gülch, E. (1987). A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *Proc. ISPRS Intercommission Conference on Fast Processing of Photogrammetric Data*, pages 281–305, Interlaken, Switzerland. (Cited on pages 23 and 64)

Fourier, J. (1822). *Théorie analytique de la chaleur, par M. Fourier.* Chez Firmin Didot, père et fils. (Cited on page 15)

Fritsch, D. S. (1992). A medial description of greyscale image structure by gradient-limited diffusion. In Robb, R. A., editor, *Visualization in Biomedical Computing '92*, volume 1808 of *Proceedings of SPIE*, pages 105–117. SPIE Press, Bellingham. (Cited on pages 15 and 19)

Fučik, S., Kratochvil, A., and Nečas, J. (1973). Kačanov–Galerkin method. *Commentationes Mathematicae Universitatis Carolinae*, 14(4):651–659. (Cited on pages 20, 67, 71, and 110)

Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., and Seidel, H.-P. (2005). Towards PDE-based image compression. In Paragios, N., Faugeras, O., Chan, T., and Schnörr, C., editors, *Variational, Geometric and Level-Set Methods in Computer Vision*, volume 3752 of *Lecture Notes in Computer Science*, pages 37–48. Springer, Berlin. (Cited on pages 16, 18, 32, and 43)

Galić, I., Weickert, J., Welk, M., Bruhn, A., Belyaev, A., and Seidel, H.-P. (2008). Image compression with anisotropic diffusion. *Journal of Mathematical Imaging and Vision*, 31(2–3):255–269. (Cited on pages 18, 32, and 64)

Gasquet, C., Ryan, R., and Witomski, P. (1998). *Fourier Analysis and Applications: Filtering, Numerical Computation, Wavelets.* Texts in Applied Mathematics. Springer, New York. (Cited on pages 12, 16, 35, 69, and 111)

Gelfand, I. M. and Fomin, S. V. (2000). *Calculus of Variations.* Dover, New York. (Cited on pages 7, 10, 11, 15, 28, 53, 86, 99, and 100)

Gerig, G., Kübler, O., Kikinis, R., and Jolesz, F. A. (1992). Nonlinear anisotropic filtering of MRI data. *IEEE Transactions on Medical Imaging*, 11:221–232. (Cited on page 23)

Goldstein, J., Newbury, D., Echlin, P., Joy, D., Romig, A., Lyman, C., Fiori, C., and Lifshin, E. (2012). *Scanning Electron Microscopy and X-Ray Microanalysis: A Text for Biologists, Materials Scientists, and Geologists.* Springer, New York. (Cited on page 2)

Gonzalez-Velasco, E. (1996). *Fourier Analysis and Boundary Value Problems.* Elsevier Science. (Cited on page 18)

Graham, T. (1829). A short account of experimental researches on the diffusion of gases through each other, and their separation by mechanical means. *Quarterly Journal of Science, Literature and Art*, 27:74–83. (Cited on page 15)

Green, P. J. (1990). Bayesian reconstructions from emission tomography data using a modified em algorithm. *IEEE Transactions on Medical Imaging*, 9(1):84–93. (Cited on pages 42 and 66)

Grewenig, S., Weickert, J., and Bruhn, A. (2010). From box filtering to fast explicit diffusion. In Goesele, M., Roth, S., Kuijper, A., Schiele, B., and Schindler, K., editors, *Pattern Recognition*, volume 6376 of *Lecture Notes in Computer Science*, pages 533–542. Springer, Berlin. (Cited on page 88)

Hall, A. R. (1990). Beyond the fringe: Diffraction as seen by Grimaldi, Fabri, Hooke and Newton. *Notes and Records of the Royal Society of London*, 44(1):13–23. (Cited on pages 1 and 46)

Hell, S. W. and Wichmann, J. (1994). Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy. *Optics Letters*, 19(11):780–782. (Cited on pages 1, 8, 41, and 46)

Hoffmann, S., Mainberger, M., Weickert, J., and Puhl, M. (2013). Compression of depth maps with segment-based homogeneous diffusion. In Kuijper, A., Bredies, K., Pock, T., and Bischof, H., editors, *Scale-Space and Variational Methods in Computer Vision*, volume 7893 of *Lecture Notes in Computer Science*, pages 319–330. Springer, Berlin. (Cited on pages 18 and 32)

Holmes, T. J. and Liu, Y.-H. (1991). Acceleration of maximum–likelihood image restoration for fluorescence microscopy and other noncoherent imagery. *Journal of the Optical Society of America A*, 8(6):893–907. (Cited on page 68)

Hong, L., Yu, J., Hong, C., and Sui, W. (2009). Depth estimation from defocus images based on oriented heat–flows. In *Proc. IEEE International Conference on Machine Vision*, pages 212–215, Dubai, UAE. (Cited on page 88)

Hooke, R. (1665). *Micrographia: or Some physiological descriptions of minute bodies made by magnifying glasses with observations and inquiries thereupon*. Royal Society, J. Martyn and J. Allestry, London. (Cited on pages 1 and 46)

Horn, B. (1968). Focusing. *MIT Artificial Intelligence Laboratory*, Memo no. 160. (Cited on pages 92, 93, and 94)

Horn, B. and Brooks, M. J., editors (1989). *Shape from Shading*. MIT Press, Cambridge, MA, USA. (Cited on page 6)

Horn, B. and Schunck, B. (1981). Determining optical flow. *Artificial Intelligence*, 17:185–203. (Cited on page 10)

Hu, H. and de Haan, G. (2007). Adaptive image restoration based on local robust blur estimation. In Blanc-Talon, J., Philips, W., Popescu, D., and Scheunders, P., editors, *Advanced Concepts for Intelligent Vision Systems*, volume 4678 of *Lecture Notes in Computer Science*, pages 461–472. Springer, Berlin. (Cited on page 88)

Huber, P. J. (2004). *Robust Statistics*, volume 1. Wiley, Chichester. (Cited on pages 11 and 60)

Huygens, C. (1690). *Traité de la lumière: où sont expliquées les causes de ce qui luy arrive dans la reflexion, & dans la refraction, et particulièrement dans l'étrange refraction du Cristal d'Islande*. Pierre Vander Aa. (Cited on pages 1 and 46)

Iijima, T. (1959). Basic theory of pattern observation. In *Papers of Technical Group on Automata and Automatic Control*. IECE, Japan. In Japanese. (Cited on page 18)

Jin, H. and Favaro, P. (2002). A variational approach to shape from defocus. In Heyden, A., Sparr, G., Nielsen, M., and Johansen, P., editors, *Computer Vision – ECCV 2002*, volume 2351 of *Lecture Notes in Computer Science*, pages 18–30. Springer, Berlin. (Cited on pages 88, 89, and 110)

King, H. C. (2011). *The history of the telescope*. Dover Publications, New York. (Cited on page 1)

Knoll, M. and Ruska, E. (1932). Das Elektronenmikroskop. *Zeitschrift für Physik*, 78(5-6):318–339. (Cited on pages 1 and 45)

Köser, K., Zach, C., and Pollefeys, M. (2011). Dense 3D reconstruction of symmetric scenes from a single image. In *Pattern Recognition*, volume 6835 of *Lecture Notes in Computer Science*, pages 266–275. Springer, Berlin. (Cited on page 6)

Krug, H. F. and Wick, P. (2011). Nanotoxicology: An interdisciplinary challenge. *Angewandte Chemie International Edition*, 50(6):1260–1278. (Cited on page 41)

Langford, M. (2000). *Basic Photography*. Taylor & Francis, Oxford. (Cited on page 85)

Le, T., Chartrand, R., and Asaki, T. J. (2007). A variational approach to reconstructing images corrupted by Poisson noise. *Journal of Mathematical Imaging and Vision*, 27(3):257–263. (Cited on page 50)

Lehmann, T., Gönner, C., and Spitzer, K. (1999). Survey: Interpolation methods in medical image processing. *IEEE Transactions on Medical Imaging*, 18(11):1049–1075. (Cited on page 16)

Lewinski, N., Colvin, V., and Drezek, R. (2008). Cytotoxicity of nanoparticles. *Small*, 4(1):26–49. (Cited on page 41)

Lichtman, J. W. and Conchello, J. A. (2005). Fluorescence microscopy. *Nature Methods*, 2:910–919. (Cited on page 44)

Lin, X., Suo, J., and Dai, Q. (2015). Extracting depth and radiance from a defocused video pair. *IEEE Transactions on Circuits and Systems for Video Technology*, 25(4):557–569. (Cited on pages 86 and 88)

Lucy, L. B. (1974). An iterative technique for rectification of observed distributions. *The Astronomical Journal*, 79(6):745–765. (Cited on pages 6, 11, 12, 16, 34, 36, 39, 42, 58, 75, 82, and 126)

Ludusan, C. and Lavialle, O. (2012). Multifocus image fusion and denoising: A variational approach. *Pattern Recognition Letters*, 33(10):1388 – 1396. (Cited on page 87)

Luenberger, D. and Yinyu, Y. (2015). *Linear and Nonlinear Programming*. Springer, New York, 4rd edition. (Cited on pages 67 and 113)

Mainberger, M. and Weickert, J. (2009). Edge-based image compression with homogeneous diffusion. In Jiang, X. and Petkov, N., editors, *Computer Analysis of Images and Patterns*, volume 5702 of *Lecture Notes in Computer Science*, pages 476–483. Springer, Berlin. (Cited on pages 18 and 32)

Marquina, A. and Osher, S. (1999). A new time dependent model based on level set motion for nonlinear deblurring and noise removal. In Nielsen, M.,

Johansen, P., Olsen, O. F., and Weickert, J., editors, *Scale-Space Theories in Computer Vision*, volume 1682 of *Lecture Notes in Computer Science*, pages 429–434. Springer, Berlin. (Cited on pages 6, 10, 16, 34, 38, 43, 109, and 42)

Marr, D. and Poggio, T. (1976). Cooperative computation of stereo disparity. *Science*, 194(4262):283–287. (Cited on page 6)

Masnou, S. and Morel, J.-M. (1998). Level lines based disocclusion. In *Proc. IEEE International Conference on Image Processing*, volume 3, pages 259–263, Chicago. (Cited on page 16)

Maynard, A. (2006). Nanotechnology: A research strategy for addressing risk. *Woodrow Wilson International Center for Scholars, PEN*, 3. (Cited on page 41)

Minsky, M. (1988). Memoir on inventing the confocal scanning microscope. *Scanning*, 10:128–138. (Cited on pages 8, 41, and 43)

Mitchell, A. R. and Griffiths, D. F. (1980). *The Finite Difference Method in Partial Differential Equations*. Wiley, Chichester. (Cited on pages 15 and 20)

Morel, J.-M. and Solimini, S. (1994). *Variational Methods in Image Segmentation*. Birkhäuser, Basel. (Cited on page 10)

Morton, K. W. and Mayers, L. M. (2005). *Numerical Solution of Partial Differential Equations*. Cambridge University Press, Cambridge, UK, second edition. (Cited on pages 15, 20, 70, and 112)

Müller, M. (2006). *Introduction to Confocal Fluorescence Microscopy*. Tutorial Text Series. SPIE Press, Bellingham. (Cited on pages 47 and 48)

Mumford, D. and Shah, J. (1989). Optimal approximation of piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42:577–685. (Cited on page 10)

Murphy, D. and Davidson, M. (2012). *Fundamentals of Light Microscopy and Electronic Imaging*. Wiley, New York. (Cited on pages 45 and 48)

Nagy, J. G. and Strakoš, Z. (2000). Enforcing nonnegativity in image reconstruction algorithms. In Wilson, D. C., Tagare, H. D., Bookstein, F. L., Preteux, F. J., and Dougherty, E. R., editors, *Mathematical Modeling, Estimation, and Imaging*, volume 4121, pages 182–190, San Diego. (Cited on page 55)

Namboodiri, V. P. and Chaudhuri, S. (2004). Use of linear diffusion in depth estimation based on defocus cue. In *Proc. Indian Conference on Computer Vision, Graphics and Image Processing*, pages 133–138, Kolkata. Allied Publishers. (Cited on page 88)

Namboodiri, V. P. and Chaudhuri, S. (2007). On defocus, diffusion and depth estimation. *Pattern Recognition Letters*, 28(3):311–319. (Cited on page 88)

Namboodiri, V. P., Chaudhuri, S., and Hadap, S. (2008). Regularized depth from defocus. In *Proc. IEEE International Conference on Image Processing*, pages 1520–1523, San Diego, USA. (Cited on page 88)

Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, Berlin. (Cited on page 111)

Osher, S. and Rudin, L. (1994). Total variation based image restoration with free local constraints. In *Proc. IEEE International Conference on Image Processing*, volume 3, pages 31–35, Austin, Texas. (Cited on pages 6, 16, 34, 38, 43, 109, and 42)

Panin, V., Zeng, G., and Gullberg, G. (1998). Total variation regulated EM algorithm. In *Nuclear Science Symposium, 1998. Conference Record*, volume 3, pages 1562–1566. IEEE. (Cited on pages 42, 52, and 66)

Pawley, J. B., editor (2006). *Handbook of Biological Confocal Microscopy*. Springer, Berlin, 3rd edition. (Cited on pages 2, 44, 45, and 48)

Pentland, A. P. (1987). A new sense for depth of field. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(4):523–531. (Cited on pages 87 and 94)

Perona, P. and Malik, J. (1987). Scale space and edge detection using anisotropic diffusion. In *Proc. Computer Society Workshop on Computer Vision*, pages 16–22. IEEE Computer Society Press. (Cited on pages 15, 18, 19, and 63)

Pertuz, S., Puig, D., and García, M. A. (2013). Analysis of focus measure operators for shape-from-focus. *Pattern Recognition*, 46:1415–1432. (Cited on page 87)

Pharr, M. and Humphreys, G. (2004). *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann, San Francisco. (Cited on pages 85, 89, 90, and 92)

Potmesil, M. and Chakravarty, I. (1981). A lens and aperture camera model for synthetic image generation. In *Computer Graphics*, SIGGRAPH, pages 297–305, Dallas. ACM. (Cited on page 92)

Prato, M., R., R. C., Zanni, L., Boccacci, P., and Bertero, M. (2012). Efficient deconvolution methods for astronomical imaging: algorithms and IDL–GPU codes. *Astronomy & Astrophysics*, 539:A133. (Cited on page 38)

Rajagopalan, A. and Chaudhuri, S. (1997). A variational approach to recovering depth from defocused images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(10):1158–1164. (Cited on pages 86 and 87)

Rangayyan, R. (2004). *Biomedical Image Analysis*. Biomedical Engineering. CRC Press, Boca Raton, USA. (Cited on page 4)

Richardson, W. H. (1972). Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America*, 62:55–59. (Cited on pages 6, 11, 12, 16, 34, 36, 39, 42, 58, 75, 82, and 126)

Roussos, A. and Maragos, P. (2007). Vector-valued image interpolation by an anisotropic diffusion-projection PDE. In Sgallari, F., Murli, F., and Paragios, N., editors, *Scale Space and Variational Methods in Computer Vision*, volume 4485 of *Lecture Notes in Computer Science*, pages 104–115. Springer, Berlin. (Cited on page 16)

Rudin, L. I., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268. (Cited on pages 10, 15, 20, 28, and 42)

Sawatzky, A., Brune, C., Mueller, J., and Burger, M. (2009). Total variation processing of images with Poisson statistics. In Jiang, X. and Petkov, N., editors, *Proceedings of the 13th International Conference on Computer Analysis of Images and Patterns*, volume 5702 of *Lecture Notes in Computer Science*, pages 533–540. Springer, Berlin. (Cited on pages 42, 52, 59, and 66)

Sawatzky, A., Brune, C., Wubbeling, F., Kosters, T., Schafers, K., and Burger, M. (2008). Accurate EM-TV algorithm in PET with low SNR. In *Nuclear Science Symposium Conference Record, NSS*, pages 5133–5137. (Cited on pages 42 and 66)

Sawatzky, A. and Burger, M. (2010). Edge–preserving regularization for the deconvolution of biological images in nanoscopy. In T. E. Simos, G. P.

and Tsitouras, C., editors, *Proc. International Conference of Numerical Analysis and Applied Mathematics*, volume 1281 of *Conference Proceedings*, pages 1983–1986. (Cited on pages 42, 59, and 66)

Scherzer, O. and Weickert, J. (2000). Relations between regularization and diffusion filtering. *Journal of Mathematical Imaging and Vision*, 12(1):43–63. (Cited on pages 18 and 31)

Schnörr, C. (1994). Unique reconstruction of piecewise smooth images by minimizing strictly convex non-quadratic functionals. *Journal of Mathematical Imaging and Vision*, 4:189–198. (Cited on page 15)

Semwogerere, D. and Weeks, E. R. (2008). Confocal microscopy. In Wnek, G. and Bowlin, G., editors, *Encyclopedia of Biomaterials and Biomedical Engineering*, volume 2, pages 705–714. Informa Healthcare, London. (Cited on pages 44 and 45)

Setzer, S., Steidl, G., and Teuber, T. (2010). Deblurring Poissonian images by split Bregman techniques. *Journal of Visual Communication and Image Representation*, 21(3):193–199. (Cited on pages 42 and 66)

Shepp, L. A. and Vardi, Y. (1982). Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging*, 1(2):113–122. (Cited on pages 42 and 50)

Shirley, P. and Morley, R. (2008). *Realistic Ray Tracing, Second Edition.* Taylor & Francis, Oxford. (Cited on pages 89 and 90)

Slesareva, N., Bruhn, A., and Weickert, J. (2005). Optic flow goes stereo: A variational method for estimating discontinuity-preserving dense disparity maps. In Kropatsch, W., Sablatnig, R., and Hanbury, A., editors, *Pattern Recognition*, volume 3663 of *Lecture Notes in Computer Science*, pages 33–40. Springer, Berlin. (Cited on page 6)

Smith, C. L. (2008). *Basic Confocal Microscopy*, pages 14.11.1–14.11.18. Wiley, Chichester. (Cited on page 44)

Snyder, D., Schulz, T. J., and O'Sullivan, J. A. (1992). Deblurring subject to nonnegativity constraints. *IEEE Transactions on Image Processing*, 40(5):1143–1150. (Cited on page 52)

Starck, J.-L. and Murtagh, F. (1994). Image restoration with noise suppression using the wavelet transform. *Astronomy and Astrophysics*, 288:342–348. (Cited on page 42)

Steidl, G. and Teuber, T. (2010). Removing multiplicative noise by Douglas-Rachford splitting methods. *Journal of Mathematical Imaging and Vision*, 36(2):168–184. (Cited on pages 11, 28, 49, and 10)

Steidl, G., Weickert, J., Brox, T., Mrázek, P., and Welk, M. (2004). On the equivalence of soft wavelet shrinkage, total variation diffusion, total variation regularization, and SIDEs. *SIAM Journal on Numerical Analysis*, 42(2):686–713. (Cited on page 75)

Stroebel, L. (1999). *View Camera Technique*. Taylor & Francis, Oxford. (Cited on page 85)

Subbarao, M. (1988). Parallel depth recovery by changing camera parameters. In *Proc. IEEE International Conference on Computer Vision*, pages 149–155, Washington, USA. (Cited on pages 87, 92, and 93)

Subbarao, M. and Surya, G. (1994). Depth from defocus: A spatial domain approach. *International Journal of Computer Vision*, 13:271–294. (Cited on page 87)

Subbarao, M. and Tyan, J.-K. (1997). Noise sensitivity analysis of depth-from-defocus by a spatial-domain approach. In *Proceedings of SPIE 3174*, pages 174–187, San Diego, USA. (Cited on page 87)

Sugimoto, S. A. and Ichioka, Y. (1985). Digital composition of images with increased depth of focus considering depth information. *Applied Optics*, 24(14):2076–2080. (Cited on page 87)

The Royal Society & The Royal Academy of Engineering (2004). Nanoscience and nanotechnologies: opportunities and uncertainties. (Cited on page 41)

Tikhonov, A. N. (1963). Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics Doklady*, 4:1035–1038. (Cited on pages 10, 27, 28, 31, 33, 34, 38, 39, 42, and 99)

Träger, F. (2007). *Springer Handbook of Lasers and Optics*. Springer, New York. (Cited on page 85)

Twyman, F. (1988). *Prism and Lens Making: A Textbook for Optical Glassworkers*. CRC Press, Boca Raton, USA. (Cited on page 1)

Van Etten, W. (2006). *Introduction to Random Signals and Noise*. Wiley, Chichester. (Cited on page 3)

Vogel, C. R. (2002). *Computational Methods for Inverse Problems*. SIAM, Philadelphia. (Cited on pages 20, 67, 71, and 110)

Vogel, O., Leichtweis, T., Breuß, M., and Weickert, J. (2009). Fast shape from shading for Phong-type surfaces. In Tai, X.-C., Mørken, K., Lysaker, M., and Lie, K.-A., editors, *Scale Space and Variational Methods in Computer Vision*, volume 5567 of *Lecture Notes in Computer Science*, pages 733–744. Springer, Berlin. (Cited on page 6)

Vonesch, C. and Unser, M. (2007). A fast iterative thresholding algorithm for wavelet-regularized deconvolution. In *Optical Engineering + Applications*, volume 6701, pages 1–5, San Diego, CA. (Cited on page 42)

Wang, Z., Bovik, A., Sheikh, H., and Simoncelli, E. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612. (Cited on pages 114, 115, and 122)

Wayne, R. (2013). *Light and Video Microscopy*. Elsevier Science, Philadelphia. (Cited on page 85)

Weber, I. (2010). Analyse und Modellierung von Fluktuationen intrazellulärer Filamente. Diploma thesis, Department of Theoretical Physics, Saarland University, Saarbrücken. (Cited on page 41)

Wei, Y., Dong, Z., and Wu, C. (2009). Global depth from defocus with fixed camera parameters. In *Proc. IEEE International Conference on Mechatronics and Automation*, pages 1887–1892, Changchun, China. (Cited on page 88)

Wei, Y. and Wu, C. (2015). Fast depth reconstruction with a defocus model on micro scale. In *IEEE International Conference on Cyber Technology in Automation, Control, and Intelligent Systems (CYBER)*, pages 947–952. IEEE. (Cited on page 88)

Wei, Y., Wu, C., and Dong, Z. (2014). Nanoscale depth reconstruction from defocus: within an optical diffraction model. *Optics express*, 22(21):25481–25493. (Cited on page 88)

Weickert, J. (1996). Theoretical foundations of anisotropic diffusion in image processing. *Computing Supplement*, 11:221–236. (Cited on page 15)

Weickert, J. (1997). A review of nonlinear diffusion filtering. In ter Haar Romeny, B., Florack, L., Koenderink, J., and Viergever, M., editors, *Scale-Space Theory in Computer Vision*, volume 1252 of *Lecture Notes in Computer Science*, pages 3–28. Springer, Berlin. (Cited on page 15)

Weickert, J. (1998). *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart. (Cited on pages 8, 15, 17, 20, 21, 22, 23, 25, 43, 64, and 82)

Weickert, J. (1999a). Coherence-enhancing diffusion of colour images. *Image and Vision Computing*, 17(3–4):199–210. (Cited on page 23)

Weickert, J. (1999b). Nonlinear diffusion filtering. In Jähne, B., Haußecker, H., and Geißler, P., editors, *Handbook on Computer Vision and Applications, Vol. 2: Signal Processing and Pattern Recognition*, pages 423–450. Academic Press, San Diego. (Cited on page 25)

Weickert, J. and Welk, M. (2006). Tensor field interpolation with PDEs. In Weickert, J. and Hagen, H., editors, *Visualization and Processing of Tensor Fields*, pages 315–325. Springer, Berlin. (Cited on pages 16, 31, 32, 33, and 43)

Weickert, J., Welk, M., and Wickert, M. (2013). L2-stable nonstandard finite differences for anisotropic diffusion. In Kuijper, A., Bredies, K., Pock, T., and Bischof, H., editors, *Scale-Space and Variational Methods in Computer Vision*, volume 7893 of *Lecture Notes in Computer Science*, pages 380–391. Springer, Berlin. (Cited on pages 25 and 69)

Welk, M. (2010). Robust variational approaches to positivity-constrained image deconvolution. Technical Report 261, Department of Mathematics, Saarland University, Saarbrücken, Germany. (Cited on pages 11, 13, 49, 52, 53, 55, 56, 59, 60, 66, 67, 75, 76, 82, 104, 107, 126, and 12)

Welk, M. (2015). A robust variational model for positive image deconvolution (online first: 07 February). *Signal, Image and Video Processing*. (Cited on pages 11, 28, 49, 53, 55, 56, 59, 60, and 10)

Welk, M. and Nagy, J. (2007). Variational deconvolution of multi–channel images with inequality constraints. In Martí, J., Benedí, J., Mendonça, A., and Serrat, J., editors, *Pattern Recognition and Image Analysis*, volume 4477 of *Lecture Notes in Computer Science*, pages 386–393. Springer, Berlin. (Cited on pages 55, 60, 104, and 106)

Welk, M., Steidl, G., and Weickert, J. (2008). Locally analytic schemes: A link between diffusion filtering and wavelet shrinkage. *Applied and Computational Harmonic Analysis*, 24:195–224. (Cited on page 75)

Welk, M., Theis, D., Brox, T., and Weickert, J. (2005). PDE–based deconvolution with forward–backward diffusivities and diffusion tensors. In Kimmel, R., Sochen, N., and Weickert, J., editors, *Scale Space and PDE*

Methods in Computer Vision, volume 3459 of *Lecture Notes in Computer Science*, pages 585–597. Springer, Berlin. (Cited on pages 43 and 64)

Westphal, V. and Hell, S. (2005). Nanoscale resolution in the focal plane of an optical microscope. *Physical Review Letters*, 94(14):571–579. (Cited on page 47)

Whittaker, E. T. (1923). A new method of graduation. *Proceedings of the Edinburgh Mathematical Society*, 41:65–75. (Cited on pages 10, 27, 28, 31, 33, 34, 38, 39, 42, and 99)

Wiener, N. (1949). *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. MIT Press, Cambridge. (Cited on pages 6, 12, 16, 34, 36, 39, 40, and 42)

Wu, Q., Wang, K., and Zuo, W. (2014). Depth from defocus based on geometric constraints. *Journal of Computers*, 9(1):44–51. (Cited on page 88)

You, Y.-L. and Kaveh, M. (1996a). Anisotropic blind image restoration. In *Proc. IEEE International Conference on Image Processing*, volume 2, pages 461–464, Lausanne, Switzerland. (Cited on pages 6, 10, 34, and 109)

You, Y.-L. and Kaveh, M. (1996b). A regularization approach to joint blur identification and image restoration. *IEEE Transactions on Image Processing*, 5(3):416–428. (Cited on pages 34 and 109)

You, Y.-L. and Kaveh, M. (1999). Blind image restoration by anisotropic regularization. *IEEE Transactions on Image Processing*, 8(3):396–407. (Cited on page 43)

Zanella, R., Zanghirati, G., Cavicchioli, R., Zanni, L., Boccacci, P., Bertero, M., and Vicidomini, G. (2013). Towards real-time image deconvolution: application to confocal and STED microscopy. *Scientific reports*, 3. (Cited on page 42)

Zervakis, M., Katsaggelos, A., and Kwon, T. (1995). A class of robust entropic functionals for image restoration. *IEEE Transactions on Image Processing*, 4(6):752–773. (Cited on pages 11, 60, and 107)

Zhang, R., Tsai, P.-S., Cryer, J. E., and Shah, M. (1999). Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706. (Cited on page 6)

Ziggelaar, A. (1980). How did the wave theory of light take shape in the mind of Christiaan Huygens? *Annals of Science*, 37(2):179–187. (Cited on page 1)