

Nonlocal Smoothing and Adaptive Morphology for Scalar- and Matrix-Valued Images

Thesis for obtaining the degree of
Doctor of Natural Sciences (Doctor rerum naturalium, Dr. rer. nat.)
of the Faculties of Natural Sciences and Technology of Saarland University

Luis Pizarro

Saarbrücken, Germany
2010

Thesis Advisor: Prof. Dr. Joachim Weickert

Referees: Prof. Dr. Joachim Weickert, Prof. Dr. Charles Kervrann



Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science
Saarland University, 66041 Saarbrücken, Germany

This research was supported by the
Deutscher Akademischer Austausch Dienst (DAAD), grant no. A/05/21715

Tag des Kolloquiums

19.05.2011

Dekan

Prof. Dr. Holger Hermanns

Prüfungsausschuss

Prof. Dr. Matthias Hein (Ausschussvorsitzender)
Universität des Saarlandes

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

Prof. Dr. Charles Kervrann (2. Gutachter)
INRIA Rennes, Frankreich

Prof. Dr. Bernhard Burgeth (Protokollführer)
Universität des Saarlandes

Affidavit

I hereby swear in lieu of an oath that I have independently prepared this thesis and without using other aids than those stated. The data and concepts taken over from other sources or taken over indirectly are indicated citing the source.

The thesis was not submitted so far either in Germany or in another country in the same or a similar form in a procedure for obtaining an academic title.

Saarbrücken, May 24th, 2011

Luis Pizarro

Copyright © by Luis Pizarro 2010. All rights reserved. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photography, recording, or any information storage or retrieval system, without permission in writing from the author. An explicit permission is given to Saarland University to reproduce up to 100 copies of this work and to publish it online. The author confirms that the electronic version is equal to the printed version. It is currently available at <http://www.mia.uni-saarland.de/pizarro/pizarro10-diss.pdf>.

Kurze Zusammenfassung

In dieser Arbeit beschäftigen wir uns mit zwei klassischen Störungsquellen in der Bildanalyse, nämlich mit Rauschen und unvollständigen Daten. Klassische Grauwert- und Farb-Fotografien wie auch matrixwertige Bilder, zum Beispiel Diffusionstensor-Magnetresonanz-Aufnahmen, können durch Gauß- oder Impulsrauschen gestört werden, oder können durch fehlende Daten gestört sein. In dieser Arbeit entwickeln wir neue Rekonstruktionsverfahren zum zur Bildglättung und zur Bildvervollständigung, die sowohl auf skalar- als auch auf matrixwertige Bilddaten anwendbar sind. Zur Lösung des Bildglättungsproblems schlagen wir diskrete Variationsverfahren vor, die aus nichtlokalen Daten- und Glattheitstermen bestehen und allgemeine auf Bildausschnitten definierte Unähnlichkeitsmaße bestrafen. Kantenerhaltende Filter werden durch die gemeinsame Verwendung solcher Maße in stark texturierten Regionen zusammen mit robusten nichtkonvexen Straffunktionen möglich. Für das Problem der Datenvervollständigung führen wir adaptive anisotrope morphologische partielle Differentialgleichungen ein, die Dilatations- und Erosionsprozesse modellieren. Diese passen sich der lokalen Geometrie an, um adaptiv fehlende Daten aufzufüllen, unterbrochene gerichtet Strukturen zu schließen und sogar flussartige Strukturen anisotrop zu verstärken. Die ausgezeichneten Rekonstruktionseigenschaften der vorgestellten Techniken werden anhand verschiedener synthetischer und realer Datensätze demonstriert.

Short Abstract

In this work we deal with two classic degradation processes in image analysis, namely noise contamination and incomplete data. Standard greyscale and colour photographs as well as matrix-valued images, e.g. diffusion-tensor magnetic resonance imaging, may be corrupted by Gaussian or impulse noise, and may suffer from missing data. In this thesis we develop novel reconstruction approaches to image smoothing and image completion that are applicable to both scalar- and matrix-valued images. For the image smoothing problem, we propose discrete variational methods consisting of nonlocal data and smoothness constraints that penalise general dissimilarity measures. We obtain edge-preserving filters by the joint use of such measures rich in texture content together with robust non-convex penalisers. For the image completion problem, we introduce adaptive, anisotropic morphological partial differential equations modelling the dilation and erosion processes. They adjust themselves to the local geometry to adaptively fill in missing data, complete broken directional structures and even enhance flow-like patterns in an anisotropic manner. The excellent reconstruction capabilities of the proposed techniques are tested on various synthetic and real-world data sets.

Zusammenfassung

Das Problem der Rekonstruktion eines beschädigten Bildes stellt eine sehr große Herausforderung dar, in Abhängigkeit von der Art des Bildes selbst, der Art der Beschädigung, die das Bild erlitten hat, und der speziellen Rekonstruktionsanforderungen. In dieser Dissertation arbeiten wir mit klassischen Grauwert- und Farb-Fotografien – als *skalarwertige* Bilder bezeichnet – und mit matrixwertigen Bildern, wie sie etwa durch Diffusionstensor-Magnetresonanz-Bildgebung (DT-MRI) erzeugt werden. Wir betrachten die Situation gestörter Bilddaten durch Gauß- oder Impulsrauschen und das Problem fehlender oder unvollständiger Daten. Ziel ist die Entwicklung neuer Verfahren zur Bildglättung, zum Entrauschen und zur Vervollständigung bei folgenden Anforderungen: Die Methoden sollen kanten erhalten und strukturadaptive sein, und sie sollen mit den nötigen Anpassungen sowohl auf skalarwertige als auch auf matrixwertige Daten anwendbar sein. Darüber hinaus erzeugt DT-MRI physikalisch interpretierbare positiv semidefinite Matrixfelder, eine Eigenschaft, die unter Anwendung der entwickelten Verfahren erhalten werden soll.

Diese Arbeit besteht aus zwei Teilen: Im ersten Teil werden diskrete Variationsansätze zum Vereinfachen von Bilddaten untersucht. Im skalarwertigen Fall schlagen wir ein Energiefunktional vor, das aus nichtlokalen Daten- und Glattheitstermen besteht, die allgemeine auf Bildausschnitten definierte Unähnlichkeitsmasse bestrafen. Das Erhalten von Kanten wird durch die gemeinsame Verwendung solcher Masse bei stark texturierten Inhalten zusammen mit robusten nichtkonvexen Straffunktionen erreicht. Unser Modell verallgemeinert frühere Arbeiten in der Literatur zu nichtlokalen Methoden, liefert bessere Ergebnisse als andere Ausschnittsbasierte Verfahren und ist qualitativ vergleichbar mit dem aktuellen Stand der Forschung. Wir erweitern diese nichtlokale Herangehensweise auf zwei verschiedenen Arten auf matrixwertige Daten: zunächst isotrop, indem die Energieterme auf skalaren Unähnlichkeitsmassen für Matrizen wirken, danach anisotrop, durch Vermeiden solcher skalarwertiger Masse, durch direktes Einsetzen von Matrizendifferenzen in die Straffunktionen. In beiden Fällen werden Bedingungen für den Erhalt positiver Definitheit diskutiert. Schließlich gehen wir auf die Verbindungen zwischen beiden hier vorgestellten Verfahren und anderen aktuellen skalar- und matrixwertigen Ansätzen ein.

Im zweiten Teil der Arbeit lösen wir das Vervollständigungsproblem mit Hilfe des Formalismus der mathematischen Morphologie. Modelliert durch partielle Differentialgleichungen (PDEs) schlagen wir skalarwertige Dilatations- und Erosionsprozesse vor, die an die lokale Geometrie adaptiert werden können. Auf diese Weise können fehlende Daten erfolgreich rekonstruiert, unterbrochene gerichtete Strukturen vervollständigt und sogar flussartige Strukturen anisotrop verstärkt werden. Der Schlüssel zu solch einer Struktur-Adaptivität ist die Einführung einer Steuerungsma-

trix, die die Richtung ändert, in die sich die Wellenfront ausbreitet. Die Verwendung des Strukturensors macht eine zuverlässige Schätzung der Vorzugsrichtung lokaler Strukturen und damit eine angepasste Steuerung der morphologischen Operationen möglich. Die adaptiven PDEs werden durch neue numerische Methoden gelöst, die als gerichtete Realisierungen des Rouy-Tourin- und FCT-Schemas konstruiert werden. Um dieses Modell auf Matrixfelder zu erweitern, verwenden wir ein vor kurzem vorgeschlagenes operator-algebraisches Konzept, das es erlaubt, die reellwertige (skalare) Algebra durch die Algebra der symmetrischen Matrizen in der Formulierung skalarwertiger PDEs auszutauschen. Zusammen mit einer Verallgemeinerung des Strukturensors auf Matrixfelder erlaubt uns dieses Konzept, matrixwertige PDEs für die adaptive Dilatation/Erosion sowie matrixwertige gerichtete numerische Schemata zu formulieren. Basierend auf diesen zwei Basisoperationen konstruieren wir schließlich adaptive Versionen komplexerer morphologischer Operationen wie Öffnung, Schließung, Gradienten und Schockfiltern für skalar- und matrixwertige Daten.

Die ausgezeichneten Rekonstruktionseigenschaften der vorgestellten Techniken werden anhand verschiedener synthetischer und real gemessener Datensätze demonstriert. Da es in den Ingenieurwissenschaften und der Physik auch andere Quellen von Matrix- und Tensorfeldern neben DT-MRI gibt, ist unsere Hoffnung, dass diese Methoden weitere Anwendung über Bildverarbeitung hinaus finden werden.

Abstract

The problem of reconstructing a degraded image imposes significant challenges depending on the image type itself, the kind of degradation the image has undergone, and the specific reconstruction requirements. In this dissertation we work with standard greyscale and colour photographs – referred to as *scalar-valued* images – and with *matrix-valued* images such as those obtained by diffusion-tensor magnetic resonance imaging (DT-MRI). We confront the situation where these images are contaminated by Gaussian or impulse noise and suffer from missing or incomplete data. We therefore pursue to develop novel methods for image smoothing and image completion under the following requirements: they should exhibit edge-preserving and structure-adaptive properties, and they should be applicable to both scalar- and matrix-valued images with the necessary adjustments. Moreover, DT-MRI produces physically meaningful positive semidefinite matrix fields, property that should be maintained under the action of the developed techniques.

This thesis consists of two parts. The first part is devoted to discrete variational methods to solve the image smoothing problem. In the scalar-valued setting we propose an energy functional consisting of nonlocal data and smoothness constraints that penalise general dissimilarity measures defined on image patches. Edge-preservation is achieved by the joint use of such measures rich in texture content together with robust non-convex penalisers. Our model generalises previous work in the literature on nonlocal methods, outperforms other patch-based approaches and compares very well to the state-of-the-art. We further extend the nonlocal framework to the matrix-valued setting in two different ways: isotropically, by allowing the energy constraints act on scalar dissimilarity measures for matrices; and anisotropically, refraining from using such scalar measures and letting the penalisers act directly on the matrix differences. In both cases we discuss the conditions for the preservation of positive semidefiniteness. Last but not least, we report on the connections between the proposed approaches and other recent scalar- and matrix-valued methods.

In the second part of this thesis we solve the image completion problem via the formalism of mathematical morphology. Focusing on the framework of partial differential equations (PDEs), we propose scalar-valued dilation/erosion processes that can be adapted to the local geometry. This way, we can successfully fill in missing data, complete broken directional structures and even enhance flow-like patterns in an anisotropic manner. The key element to achieve such a structure-adaptiveness is to introduce a steering matrix that changes the direction along which the wavefronts propagate. By making use of the structure tensor concept we can reliably estimate the orientation of the local structures and hence steer the morphological operations conveniently. The adaptive PDEs are solved by novel numerical methods constructed

as directional realisations of the Rouy-Tourin and the FCT schemes. In order to extend our model to matrix fields, we use a recently proposed operator-algebraic framework that allows one to replace the real (scalar) algebra by the algebra of symmetric matrices in the formulation of scalar-valued PDEs. This framework, together with a previous generalisation of the structure tensor to matrix fields, enables us to formulate matrix-valued PDEs for adaptive dilation/erosion as well as matrix-valued directional numerical schemes. Finally, based on these two basic processes we build up adaptive versions of more complex operations such as opening, closing, gradients and shock filtering for scalar and matrix fields.

The excellent reconstruction capabilities of the proposed techniques are tested on various synthetic and real-world data sets. It is our hope that these methods find further applications beyond image processing, considering that there exist other sources of matrix and tensor fields aside from DT-MRI in engineering and the physical sciences.

To Ivanka and Luis Angelo

Acknowledgements

First of all, I want to thank my supervisor *Joachim Weickert* for giving me the opportunity to work in his group and for having supported my scholarship, kindly funded by the *German Academic Exchange Service* (DAAD). I am most grateful to him for many inspiring conversations and for his excellent teaching abilities that strengthened my commitment to the fascinating world of image analysis. Secondly, I would like to thank *Charles Kervrann* for serving as external reviewer of this thesis.

I owe my heartfelt gratitude to *Bernhard Burgeth*, *Stephan Didas* and *Michael Breuß* for the fruitful collaborations we had during my studies, but specially for their sincere friendship, advice and encouragement that I appreciate very much. I wish to thank *Pavel Mrázek* from UPEK for our productive collaboration and for the nice time we spent at the ECCV 2008 conference in Marseille. I am grateful to all co-authors of joint publications written during these years, and to *Thomas Brox* from University of Freiburg and *Thomas Schultz* from University of Chicago for several scientific discussions. I reiterate my deep gratitude to the DAAD for the financial support of my studies and to its *Referat*-team in Bonn *Maria Hartmann*, *Elfie Firzlauff* and *Britta Bauch*.

For the open and inspiring working atmosphere, for the barbecues, the beer tastings and the sport evenings, I wish to thank all former and current members of our group: *Andrés*, *Bernhard*, *Christian*, *Henning*, *Irena*, *Joachim*, *Kai*, *Levi*, *Markus*, *Martin*, *Michael*, *Natalia*, *Oliver*, *Pascal*, *Rodrigo*, *Stephan* and *Sven*. My special thanks to *Ellen Wintringer* and *Marcus Hargarter*, two exceptional people whose administrative and computational support greatly alleviated my life.

Quiero agradecer a mi numerosa familia en Chile, en especial a mis padres *Violeta* y *Manuel*, y a mis hermanos *Franco*, *Moisés* y *Daniel*, porque siempre han confiado en mí y me han brindado su apoyo y amor durante estos duros años en Europa. Los adoro, junto con todos mis familiares que no alcanzo a nombrar en estas líneas.

Final y muy especialmente, quiero dedicar este esfuerzo a mi esposa *Ivanka* y a mi hijo *Luis Angelo*, por su paciencia infinita y porque aun en los días más oscuros y lluviosos me regalan su amor y compañía, mi sol de cada día.

Saarbrücken, October 6th, 2010

Luis Pizarro

Contents

1	Introduction	1
1.1	Nonlocal Smoothing	1
1.2	Adaptive Morphology	4
1.3	Definition and Visualisation of Matrix Fields	6
1.4	Operator-Algebraic Framework for Matrix Fields	8
1.5	Outline	10
I	Nonlocal Smoothing for Scalar- and Matrix-Valued Images	13
2	Nonlocal Data and Smoothness (NDS) Framework	15
2.1	NDS Functional and its Minimisation	15
2.2	Important Special Cases	20
2.2.1	M-Estimators and Local M-Smothers	20
2.2.2	Bilateral Filtering	22
2.2.3	Regularisation Methods	22
2.2.4	Histogram Quantisation	24
2.3	NDS vs Graph Regularisation	25
2.4	Tuning the NDS Parameters	27
2.4.1	Statistically-Based Selection of Tonal Kernels	27
2.4.2	Linear Combination of Kernels	28
2.4.3	Smoothing Effects	30
2.5	Summary	33
3	Generalised NDS Framework	35
3.1	GNDS Functional and its Minimisation	35
3.2	Double Weighting	37
3.3	GNDS Filter Family	39
3.4	Alternative Formulations of the NL-Means Filter	42
3.5	(Non-) Iterative and Steady-State Solutions	43
3.6	Extension to Multichannel Images	43
3.7	Extension to Other Distance Measures	44
3.8	Numerical Experiments	44
3.8.1	Comparison of similarity measures	44
3.8.2	Comparison of several patch-based methods	46
3.9	Summary	55

4	Matrix-Valued NDS Framework	57
4.1	Isotropic Matrix-Valued NDS Model	58
4.2	Anisotropic Matrix-Valued NDS Model	61
4.3	Related Filters within This Framework	62
4.4	Numerical Experiments	63
4.4.1	Two Isotropic Models: iMNDS-E and iMNDS-LE	63
4.4.2	Influence of Parameters	65
4.4.3	Performance	65
4.4.4	3D DT-MRI data	69
4.5	Summary	70
II	Adaptive Morphology for Scalar- and Matrix-Valued Images	75
5	Introduction to Mathematical Morphology	77
5.1	Scalar-Valued Morphology	77
5.1.1	Set-Theoretical Approach	77
5.1.2	Continuous-Scale Approach	78
5.1.3	Set-Theoretical vs. Continuous-Scale Approach	79
5.1.4	Numerical Schemes for Continuous-Scale Morphology	79
5.2	Matrix-Valued Morphology	83
5.2.1	PDE-Based Morphology for Matrix Fields	83
5.2.2	Matrix-Valued Numerical Schemes	83
5.3	Summary	85
6	Adaptive Continuous-Scale Morphology	91
6.1	PDE-Driven Adaptive Morphology	92
6.2	Full Structure Tensor for Matrix Fields	94
6.3	Steering Matrix for Matrix Fields	95
6.4	Directional Numerical Schemes	96
6.4.1	Directional Rouy-Tourin scheme	96
6.4.2	Directional FCT scheme	98
6.5	Structure-Adaptive Morphological Operations	100
6.6	Experiments	101
6.6.1	Scalar-Valued Images	101
6.6.2	Synthetic 2D and 3D Matrix Fields	102
6.6.3	Real World Data: 3D DT-MRI	103
6.7	Summary	104

7	Conclusions and Perspectives	117
7.1	Conclusions	117
7.2	Perspectives	119
8	Further Contributions and Publications	121
8.1	Further Contributions	121
8.2	Publications	121
	Bibliography	125

In this dissertation we deal with two fundamental problems in image processing and analysis: *image smoothing* and *image completion*. Despite the enormous volume of research in these areas, novel imaging modalities rendering compelling applications impose new theoretical and practical challenges. Such is the case of diffusion-tensor magnetic resonance imaging (DT-MRI) that provides *tensor-* or *matrix-valued* images carrying information about the microstructure and connectivity of the brain, which drives numerous clinical applications. As in standard greyscale/colour photographs – referred to as *scalar-valued* images – DT images are as well affected by noise and suffer from missing or incomplete data. In this work we develop novel methods for image smoothing and image completion that are applicable to both scalar- and matrix-valued images. We address these challenges from the perspectives of *nonlocal smoothing* and *adaptive morphology*.

In this chapter we present a brief literature background to the techniques that will be at the core of this thesis. Further bibliographic aspects regarding connections between the proposed approaches and other methods from the literature will be detailed in subsequent chapters.

1.1 Nonlocal Smoothing

Image smoothing is a fundamental task in image processing. It serves as a noise removal tool for improving the visual quality of noisy images taken from digital cameras or scanners, as well as for providing simplified input images that are further processed in tasks such as segmentation, feature extraction and texture analysis. There exist numerous approaches to image smoothing emerging from statistical methods, information theory, transforms in the frequency domain, partial differential equations (PDEs) and variational methods [220, 7, 234, 58]. Establishing equivalences and relations between the different approaches has been focus of intense research in recent years [14, 76, 77, 148, 180, 184, 213, 198, 236].

In this thesis we concentrate on the framework of discrete variational methods. We follow a recent model for nonlocal smoothing due to Mrázek *et al.* [148] that unifies different and widely used filters in the literature. More precisely, Mrázek *et al.* pointed out the relations between several nonlinear smoothing methods such as M-estimators [66, 236], bilateral filtering [206], diffusion filters [167, 219], and regularisation/Bayesian techniques [19, 84, 149, 236]. Although these methods seem

very different at the first glance and originate in different mathematical theories, Mrázek *et al.* showed that they lead to highly similar discrete algorithms, and that all these methods can be cast in a single unified framework of discrete regularisation theory. The unifying model is formulated as an energy functional with *nonlocal data* and *smoothness* (NDS) terms – hence called NDS model. The data term rewards similarity of the filtered image to the input (noisy) image, while the smoothness term penalises high deviations from regularity on the solution. These terms can consider not only information from a small region around a pixel but also make it possible to involve large neighbourhoods. Nevertheless, the integration of nonlocal information is limited in practice since this model relies on single pixel differences, which do not carry reliable structural information too far away from a chosen pixel. This is actually the main drawback of all *single differences*-based approaches.

Two equivalent and simultaneously proposed methods, namely the *non-local means* (NL-means) filter [40, 41] and the *unsupervised, information-theoretic, adaptive* (UINTA) filter [9, 10] are able to overcome this problem. Both methods consider a whole neighbourhood (or patch) around a pixel to measure similarity. In this way, if the corresponding neighbourhoods of two pixels are similar, the pixels themselves will be considered alike even if they are spatially distant from each other. This simple idea allows a real incorporation of nonlocal pixel interactions in the smoothing process, providing impressive denoising results. The NL-means filter belongs to the class of neighbourhood filters [125, 237, 195, 206, 42] that average similar pixels based on their photometric and spatial proximities – where the spatial distance does not play a role in NL-means. In particular, it can be seen as a bilateral filter [206] with a patch-based photometric similarity measure. Several variational formulations of the NL-means filter have been proposed [117, 88, 11, 38, 129, 207] together with acceleration techniques [134, 23, 69, 71, 38, 158] and invariant patch similarity measures [217, 119, 242, 130]. This method has inspired numerous *patch-based* approaches for image smoothing, deblurring, segmentation, inpainting, super-resolution, texture analysis, among others.

Our contribution is threefold: First, we propose a generalisation of Mrázek *et al.*'s approach to the space of patches by introducing a discrete variational approach with nonlocal constraints that penalise general dissimilarity measures defined on image patches. This leads to a more robust smoothing process since the pixels selected for averaging are more coherent with the local image structure. Our framework leads to a novel family of filters suitable for removing Gaussian and impulse noise in scalar-valued images. Second, we further extend the NDS approach to the setting of matrix fields in two different ways, namely *isotropically* and *anisotropically* by allowing the energy penalisers to act on scalar- and matrix-valued similarity measures, respectively. We test our novel approaches on synthetic and DT-MRI data sets. We also discuss the conditions for preservation of positive semidefiniteness of DT-MRI fields. Third, we establish connections between the proposed approaches and other recent smoothing methods for scalar and matrix fields.

Figure 1.1 exemplifies the application of the proposed nonlocal smoothing methods to scalar- and matrix-valued images.




Figure 1.1: Nonlocal smoothing. *Left column:* Original image, noisy version with white Gaussian noise, and smoothed version with our scalar-valued nonlocal approach. *Right column:* Synthetic 2D tensor field of 3×3 symmetric matrices, noisy version with random positive definite matrices, and smoothed version with the our matrix-valued nonlocal approach. Our smoothing methods successfully reconstruct subtle image details. Tensors are visualised as ellipsoids (cf. Section 1.3).

1.2 Adaptive Morphology

Mathematical morphology aims at extracting and enhancing and shape information from image objects. It originated in 1964 at the École des Mines de Paris in the works of Matheron [140] and Serra [187]. A detailed chronology of the birth of mathematical morphology has recently been published by Matheron and Serra in [142]. Initially developed to study the structure of materials in geological samples, current application domains of this technique range from biology to astronomy.

Classical morphology refers to the so-called *set-theoretical* approach developed by its founders and is based on set theory, integral geometry and lattice algebra. It consists of matching an image with a bounded set of known shape called *structuring element* (SE) such as squares, diamonds, hexagons, discs or ellipses defined on the discrete pixel grid. The two fundamental morphological operators are *dilation* and *erosion*, which respectively expand and shrink the boundaries of an object to an extent specified by the shape of the chosen SE. These operators are utilised as the building blocks of more complex morphological processes upon which advanced algorithms can be constructed for image filtering, segmentation, classification and analysis of textures, among others.

Until the end of the '80s, all morphological processes and their applications were modelled and implemented under the umbrella of the discrete lattice theory. At the beginning of the '90s, inspired by the Gaussian scale-space generated by the linear diffusion PDE, several groups of researchers [35, 211, 4, 2] independently published nonlinear PDEs modelling the continuous morphological scale-space, referred to as *continuous-scale* approach. These morphological PDEs describe a new representation for the evolution of the object's boundaries during dilation and erosion with digitally scalable SEs, which provides sub-pixel accuracy in the resulting morphological operations/applications. For a comprehensive account of the classical aspects of mathematical morphology we refer the reader to the monographs [141, 188, 189, 100, 196], for the continuous-scale framework to the texts [181, 54, 95], and for the successive developments and applications in the area we recommend the proceedings of the nine versions of the International Symposium on Mathematical Morphology held from 1993 to 2009 [190, 191, 136, 102, 91, 203, 177, 12, 233].

In this thesis we focus on the PDE framework which is conceptually attractive since it allows for digital scalability and even adaptivity of the represented SE. Our contribution is threefold: First, we propose scalar-valued morphological PDEs with spatially-variant elliptical SEs. They can adaptively change the direction along which the wavefronts propagate depending on the orientation of the local structures. Second, we introduce directional solution schemes for these PDEs, avoiding the numerical bias to the coordinate axes of classical schemes. Third, we transfer the proposed PDEs and numerical schemes to the matrix-valued setting using a recently proposed operator-algebraic framework (Section 1.4) that allows one to replace the real (scalar) algebra by the algebra of symmetric matrices in the formulation of scalar-valued PDEs.

Figure 1.2 exemplifies the application of the proposed adaptive morphological PDEs to scalar- and matrix-valued images.




Figure 1.2: Adaptive morphology. *Left column:* Original image, isotropically eroded version, and anisotropically eroded version with our scalar-valued adaptive approach. *Right column:* Synthetic 2D spiral of 3×3 symmetric matrices with missing tensors, isotropically dilated version, and anisotropically dilated version with our matrix-valued adaptive approach. Our steerable methods are able to complete directional structures accurately. Tensors are visualised as ellipsoids (cf. Section 1.3).

1.3 Definition and Visualisation of Matrix Fields

In this thesis we regard a matrix or tensor field as a mapping

$$U : \Omega \subset \mathbb{R}^d \longrightarrow \text{Sym}_m(\mathbb{R}) \quad (1.1)$$

from a d -dimensional image domain into the set of symmetric $m \times m$ -matrices with real entries, $U(x) = (U_{p,q}(x))_{p,q=1,\dots,m}$. The set of positive (semi-)definite matrices, denoted by $\text{Sym}_m^{++}(\mathbb{R})$ (resp. $\text{Sym}_m^+(\mathbb{R})$), consists of all symmetric matrices A with $\langle v, Av \rangle := v^\top Av > 0$ (resp., ≥ 0) for $v \in \mathbb{R}^m \setminus \{0\}$. Note that at each point x the matrix $U(x)$ of a field of symmetric matrices can be diagonalised yielding

$$U(x) = V(x)^\top D(x) V(x), \quad (1.2)$$

where $V(x)$ is a real orthogonal matrix whose columns correspond to the m eigenvectors of $U(x)$ denoted by $v_1, \dots, v_m \in \mathbb{R}^m$, and $D(x)$ is a diagonal matrix whose entries are the corresponding m eigenvalues denoted by $\lambda_1, \dots, \lambda_m \in \mathbb{R}$. Without loss of generality, let us assume that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_m$.

The matrix data are visualised as an ellipsoid in each voxel via the level sets of quadratic form

$$u^\top U^{-2}(x) u = c, \quad (1.3)$$

with $u \in \mathbb{R}^3, c \in \mathbb{R}^+$. By using U^{-2} the length of the semi-axes of the ellipsoid correspond directly with the three eigenvalues of the matrix. Changing the constant c amounts to a mere scaling of the ellipsoids. Note that only positive definite matrices produce ellipsoids as level sets of its quadratic form.

Matrix fields offer the opportunity of describing anisotropy in physical measurements and in image processing models, see [225],[123] for an overview. In diffusion-tensor magnetic resonance imaging (DT-MRI), for example, information about the diffusive properties of water molecules is captured in symmetric positive definite matrices. The corresponding matrix field reflects the structure of the tissue under examination [16]. Figure 1.3 shows the axial view of a full DT-MRI data set of a human head ($128 \times 128 \times 30$ voxels) and a 2D slice (32×32 voxels). We use the most common colour scheme in the literature [164] which is driven by the orientation of the principal eigenvector. Ellipsoids with normalised $v_1 := (v_1^x, v_1^y, v_1^z)^\top$ oriented along the x -, y -, z -axis are shown in red, green and blue tones, respectively. Therefore, tensors are displayed in the (normalised) RGB colour space according to

$$\begin{aligned} R &= 1 - W(1 - |v_1^x|), \\ G &= 1 - W(1 - |v_1^y|), \\ B &= 1 - W(1 - |v_1^z|), \end{aligned} \quad (1.4)$$

where W is a weighting factor influencing the colour purity by the degree of




Figure 1.3: Diffusion-tensor magnetic resonance imaging (DT-MRI) data set. Tensors are visualised as ellipsoids with *principal eigenvector* colour coding.

anisotropy. We use the so-called *fractional anisotropy* (FA) measure [17]

$$\text{FA} = \sqrt{\frac{3}{2}} \sqrt{\frac{(\lambda_1 - \bar{\lambda})^2 + (\lambda_2 - \bar{\lambda})^2 + (\lambda_3 - \bar{\lambda})^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}, \quad (1.5)$$

with $\bar{\lambda} := (\lambda_1 + \lambda_2 + \lambda_3)/3$. Note that $\text{FA} \in [0, 1]$: $\text{FA} = 0$ means that the diffusion is isotropic at that voxel, while $\text{FA} = 1$ indicates that diffusion occurs only along v_1 .

To conclude this section, let us make the observation that aside from DT-MRI there exist other sources of matrix/tensor fields in engineering and the physical sciences: stress-strain tensors in material science, geomechanics, elastography and fluid dynamics; permittivity tensors in electromagnetism and optics; permeability tensors in fluid mechanics and earth sciences; among others.

1.4 Operator-Algebraic Framework for Matrix Fields

The operator-algebraic framework proposed by Burgeth *et al.* in [48, 47] will help us extend the proposed scalar-valued methods for nonlocal image smoothing and adaptive morphology to the setting of matrix fields. The basic idea behind this approach is to consider symmetric matrices as a natural generalisation of real numbers with a rich algebraic structure. Provided that the matrices are self-adjoint operators that can be added, multiplied with a scalar, and concatenated, they thus form an algebra. Consequently, this approach aims at replacing the field \mathbb{R} by the algebra $\text{Sym}_m(\mathbb{R})$ in the scalar, that is, \mathbb{R} -based formulation of PDEs used in image processing. This framework provides the grounds to translate scalar-valued techniques to the matrix-valued setting, as well as to directly define novel processing techniques for matrix fields, positive definite or not.

Existing scalar PDE-based processes that have been successfully transferred to the matrix-valued case using this approach are: TV diffusion, mean curvature motion, self-snakes [48], Perona-Malik diffusion [47], coherence-enhancing diffusion [49, 52], relaxed inverse scale-space flow [127] and isotropic morphology [43]. The obtained PDEs are truly matrix-valued and they do not require the additional coupling of the matrix channels as it is necessary in other PDE approaches for tensor fields [224]. This framework, however, is not restricted to PDEs.

Here we present a brief description of the key elements we utilise in Chapters 4 and 6. For an extended exposition the reader is referred to [48, 47]. With the definition of matrix fields given in the previous section, let us consider the following aspects of the algebra of symmetric matrices:

Functions of matrices. The extension of a function $h : \mathbb{R} \rightarrow \mathbb{R}$ to $\text{Sym}_m(\mathbb{R})$ is standard [107]: With a slight abuse of notation we set

$$h(U) := V^\top \text{diag}(h(\lambda_1), \dots, h(\lambda_m))V \in \text{Sym}_m^+(\mathbb{R}), \quad (1.6)$$

h denoting now a function acting on matrices as well. Observe that the systems of eigenvectors remains untouched. Specifying $h(s) = |s|$, $s \in \mathbb{R}$ as the absolute value function leads to the absolute value $|A| \in \text{Sym}_m^+(\mathbb{R})$ of a matrix A . Similarly, multivariate functions $H : \mathbb{R}^m \rightarrow \mathbb{R}^m$ can be extended to symmetric matrices via

$$H(U) := V^\top \text{diag}(H_1(\lambda_1 \dots \lambda_m), \dots, (H_m(\lambda_1 \dots \lambda_m)))V \in \text{Sym}_m^+(\mathbb{R}). \quad (1.7)$$

Partial derivatives and generalised gradient. It is natural to define the partial derivative for matrix fields *componentwise*:

$$\bar{\partial}_\omega U = (\partial_\omega U_{p,q})_{p,q=1,\dots,m} \quad (1.8)$$

where $\omega \in \{t, x_1, \dots, x_d\}$, that is, $\bar{\partial}_\omega$ stands for a spatial or temporal derivative. Viewing a matrix as a tensor (of second order), its gradient would be a third order

tensor according to the rules of differential geometry. However, we adopt a more operator-algebraic point of view by defining the *generalised gradient* $\overline{\nabla}U(x)$ at a voxel $x = (x_1, \dots, x_d)$ by

$$\overline{\nabla}U(x) := (\overline{\partial}_{x_1}U(x), \dots, \overline{\partial}_{x_d}U(x))^\top \quad (1.9)$$

which is an element of $(\text{Sym}_m(\mathbb{R}))^d$, in close analogy to the scalar setting where $\nabla u(x) \in \mathbb{R}^d$.

Symmetrised product of matrices. There will be the need for a symmetric multiplication of symmetric matrices. The product of two symmetric matrices $A, B \in \text{Sym}_m(\mathbb{R})$ is not symmetric unless the matrices commute. The following alternatives are possible [107, 47, 6]:

$$\text{Jordan product (JP): } A \bullet_J B := (AB + BA)/2, \quad (1.10)$$

$$\text{Pre-conditioning product (PP): } A \bullet_P B := A^{\frac{1}{2}}BA^{\frac{1}{2}}, \quad (1.11)$$

$$\text{Logarithmic product (LP): } A \bullet_L B := \exp(\log(A) + \log(B)). \quad (1.12)$$

In general, the JP does not preserve the positive (semi-)definiteness of its arguments. The PP preserves the property in B if A is non-singular. The LP preserves positive definiteness, but it is only defined for matrices holding that property – excluding its application to indefinite matrices.

Maximum and minimum of symmetric matrices. For later use in numerical schemes we have to clarify the notion of maximum and minimum of two symmetric matrices A, B . In direct analogy with relations known to be valid for real numbers one defines [45]:

$$\max(A, B) = \frac{1}{2}(A + B + |A - B|) \quad (1.13)$$

$$\min(A, B) = \frac{1}{2}(A + B - |A - B|) \quad (1.14)$$

where $|F|$ stands for the absolute value of the matrix F . For $F \in (\text{Sym}_m(\mathbb{R}))^d$ we set $|F|_p := \sqrt[p]{|F_1|^p + \dots + |F_d|^p}$ for $0 < p < +\infty$. It results in a positive semidefinite matrix from $\text{Sym}_m^+(\mathbb{R})$, the direct counterpart of a nonnegative real number as the length of a vector in \mathbb{R}^d .

1.5 Outline

The novel approaches developed in this dissertation have been published at international conferences [173, 171, 44, 170] and journals [174, 51]. For clarity of presentation, we have divided the exposition of our work in two parts.

Part I: Nonlocal Smoothing for Scalar- and Matrix-Valued Images

Chapter 2

We provide an overview of well-known smoothing approaches from the literature, which have recently been cast in a single unified framework of discrete regularisation theory. We also report on new connections we have discovered between this framework and classical diffusion/regularisation methods and graph regularisation methods. Moreover, we study its parameter configuration and show that it indeed outperforms its special cases filters for the task of denoising scalar-valued images degraded with Gaussian and impulse noise.

Chapter 3:

Based on the previous approach, we formulate a discrete variational framework for nonlocal smoothing whose data and regularisation constraints penalise general dissimilarity measures defined on image patches. We describe how our model generalises previous work on nonlocal methods and show that it outperforms other patch-based approaches and compares very well to the current state-of-the-art, being our approach much simpler.

Chapter 4:

Also in the framework of discrete variational methods, we further propose nonlocal *isotropic* and *anisotropic* approaches for smoothing matrix fields. We discuss the conditions for the preservation of positive semidefiniteness of the original data sets as it occurs in DT-MRI. We additionally show that several filters for symmetric matrices recently appeared in the literature can be seen as particular cases of the proposed approaches.

Part II: Adaptive Morphology for Scalar- and Matrix-Valued Images

Chapter 5:

We first review the elementary concepts of scalar-valued *isotropic* morphology from the two points of view it has been treated in the literature: the classical discrete *set-theoretical* framework and the modern *continuous-scale* framework, providing a summary of the advantages and shortcomings of both approaches. Focusing on the latter PDE-based method, we report on its recent extension to the setting of matrix fields and we further describe two numerical schemes to solve these PDEs.

Chapter 6:

In the context of continuous-scale morphology, we concentrate on adaptive, *anisotropic* morphological operators for scalar- and matrix-valued images. We propose a steerable PDE-driven approach that can adapt itself according to the orientation of the local geometry. This way, it is possible to complete directional structures, filling in missing data and even enhance flow-like patterns in an adaptive, directional manner. We as well propose directional versions of the numerical solution schemes that allow a more accurate implementation of the steerable morphological PDEs.

Finally, in Chapter 7 we briefly summarise our contributions and sketch future research directions.

Part I

**Nonlocal Smoothing
for Scalar- and Matrix-Valued Images**

2

Nonlocal Data and Smoothness (NDS) Framework

In this chapter, we describe in detail the *Nonlocal Data and Smoothness* (NDS) framework proposed by Mrázek *et al.* [148], whose numerical aspects were further treated in Didas *et al.* [74]. We provide an overview of the various smoothing methods that can be considered special cases of this unifying approach for discrete regularisation theory. We as well report novel insights into relations of this model with diffusion/regularisation methods as well as with some recently proposed graph regularisation techniques. We explore the generality of the NDS framework in combining nonlocal tonal and spatial kernels to tune hybrid filters with better performance than the previously mentioned special cases. We show that Bayesian analysis provides suitable foundations for restricting the parametric space in a noise-dependent way. Moreover, we show that the existing trade-off between the parameters controlling similarity and smoothness leads to similar results under different settings. We particularly focus on degradation processes governed by Gaussian and impulse noise. This chapter is based on [173, 174].

2.1 NDS Functional and its Minimisation






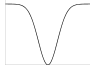
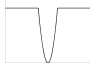
Let $f, u : \Omega \rightarrow \mathbb{R}$ be scalar images defined on the discrete image domain Ω . f stands for the (noisy) original image while u represents a processed version of it. Let $J = \{1, \dots, N\}$ be the index set of all pixels in the images. The pixel position in the bi-dimensional grid is indicated by x_i ($i \in J$). The discrete energy function E of the NDS filter presented in [148] is a convex combination of a nonlocal data (or similarity) term E_D and a nonlocal smoothness term E_S :

$$E_D(u) = \sum_{i,j \in J} \Psi_D(|u_i - f_j|^2) w_D(|x_i - x_j|^2), \quad (2.1)$$

$$E_S(u) = \sum_{i,j \in J} \Psi_S(|u_i - u_j|^2) w_S(|x_i - x_j|^2). \quad (2.2)$$

Here $\Psi_\circ : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ are increasing functions that penalise large (greyvalue) *tonal distances*. The weights $w_\circ : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ are nonnegative functions downweighting large *spatial distances*. See Table 2.1 and Table 2.2 for a non-exhaustive list of kernels Ψ and w proposed in literature. For a more comprehensive list of penalisers, see [153].

Table 2.1: Popular choices for tonal weights Ψ .

$\Psi(s^2)$		$\Psi'(s^2)$	source
(a) s^2		1	Whittaker [231] Tikhonov [204]
(b) $2 s $		$ s ^{-1}$	Rudin <i>et al.</i> [179]
(c) $2(\sqrt{s^2 + \varepsilon^2} - \varepsilon)$		$(s^2 + \varepsilon^2)^{-\frac{1}{2}}$	Acar and Vogel [1]
(d) $2\lambda^2 \left(\sqrt{1 + \frac{s^2}{\lambda^2}} - 1 \right)$		$\left(1 + \frac{s^2}{\lambda^2}\right)^{-\frac{1}{2}}$	Charbonnier <i>et al.</i> [59]
(e) $\lambda^2 \log \left(1 + \frac{s^2}{\lambda^2}\right)$		$\left(1 + \frac{s^2}{\lambda^2}\right)^{-1}$	Holland and Welsch [106] Perona and Malik [167]
(f) $2\lambda^2 \left(1 - \exp\left(-\frac{s^2}{2\lambda^2}\right)\right)$		$\exp\left(-\frac{s^2}{2\lambda^2}\right)$	Perona and Malik [167]
(g) $\min(s^2, \lambda^2)$		$\begin{cases} 1 & s < \lambda \\ 0 & \text{else} \end{cases}$	Hampel <i>et al.</i> [99] Mumford and Shah [150]

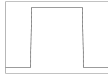

The data term E_D rewards similarity of the filtered image to the input (noisy) image, while the smoothness term E_S penalises high deviations from regularity on the solution. These terms can consider not only information from a small region around a pixel but also make it possible to involve large neighbourhoods by adjusting the extent of the spatial weights.

The complete NDS model can be regarded as a discrete nonlocal variational method combining both the data (2.1) and the smoothness (2.2) terms:

$$\begin{aligned}
E(u) &= (1 - \alpha) E_D(u) + \alpha E_S(u) & (2.3) \\
&= (1 - \alpha) \sum_{i,j \in J} \Psi_D(|u_i - f_j|^2) w_D(|x_i - x_j|^2) \\
&\quad + \alpha \sum_{i,j \in J} \Psi_S(|u_i - u_j|^2) w_S(|x_i - x_j|^2)
\end{aligned}$$

with regularisation parameter $\alpha \in [0, 1]$.

Table 2.2: Possible choices for spatial weights w .

$w(s^2)$		source
(a) $\begin{cases} 1 & s^2 \leq \varpi^2 \\ 0 & \text{else} \end{cases}$		hard window Koenderink and van Doorn [120]
(b) $\exp\left(-\frac{s^2}{\varpi^2}\right)$		soft window Chu <i>et al.</i> [66]

We now consider a robust and stable iterative procedure for minimising the energy functional. Even if the presented iterative fixed point approach is very simple, we will see that it satisfies a maximum-minimum principle for a general set of penaliser functions, and we will prove the existence of a fixed point.

Taking the partial derivatives of the data term (2.1) yields

$$\frac{\partial E_D}{\partial u_k} = 2 \sum_{j \in J} \Psi'_D(|u_k - f_j|^2) (u_k - f_j) w_D(|x_k - x_j|^2), \quad (2.4)$$

where Ψ' denotes the derivative of Ψ w.r.t. its argument. In a similar way we calculate the derivatives of the smoothness term (2.2) which leads to

$$\frac{\partial E_S}{\partial u_k} = 4 \sum_{j \in J} \Psi'_S(|u_k - u_j|^2) (u_k - u_j) w_S(|x_k - x_j|^2). \quad (2.5)$$

It is clear that the complete derivatives then have the form

$$\frac{\partial E}{\partial u_i} = (1 - \alpha) \frac{\partial E_D}{\partial u_i} + \alpha \frac{\partial E_S}{\partial u_i}. \quad (2.6)$$

For a critical point u of the energy functional E we have

$$\nabla E(u) = 0 \iff \frac{\partial E}{\partial u_i} = 0 \quad \text{for all } i \in J. \quad (2.7)$$

We define the abbreviations

$$d_{i,j} := \Psi'_D(|u_i - f_j|^2) w_D(|x_i - x_j|^2), \quad (2.8)$$

$$s_{i,j} := 2\Psi'_S(|u_i - u_j|^2) w_S(|x_i - x_j|^2) \quad (2.9)$$

which help us to rewrite (2.7) as

$$0 = (1 - \alpha) \sum_{j \in J} d_{i,j} (u_i - f_j) + \alpha \sum_{j \in J} s_{i,j} (u_i - u_j) \quad (2.10)$$

where we use the partial derivatives shown in (2.4) and (2.5). This can be transformed into fixed point form

$$u_i = \frac{(1 - \alpha) \sum_{j \in J} d_{i,j} f_j + \alpha \sum_{j \in J} s_{i,j} u_j}{(1 - \alpha) \sum_{j \in J} d_{i,j} + \alpha \sum_{j \in J} s_{i,j}}. \quad (2.11)$$

To have a positive denominator we assume that $\Psi'_D(s^2) > 0$ and $\Psi'_S(s^2) > 0$, i. e., the penalisers are monotonically increasing. Furthermore we assume that $w_D(s^2) \geq 0$, $w_S(s^2) \geq 0$ as well as $w_D(0) > 0$ and $w_S(0) > 0$ for the spatial weights. We use this equation to build up a first iterative method to minimise the value of E where the upper index k denotes the iteration number. Note that $d_{i,j}$ and $s_{i,j}$ also depend on the evolving image u^k and thus also get a superscript to denote the iteration level involved. The corresponding fixed point iteration then reads as

$$u_i^0 := f_i, \quad (2.12)$$

$$u_i^{k+1} := \frac{(1 - \alpha) \sum_{j \in J} d_{i,j}^k f_j + \alpha \sum_{j \in J} s_{i,j}^k u_j^k}{(1 - \alpha) \sum_{j \in J} d_{i,j}^k + \alpha \sum_{j \in J} s_{i,j}^k}. \quad (2.13)$$

In the following we will write this scheme (2.13) in the form $u^{k+1} = F(u^k)$ with $F : \mathbb{R}^N \rightarrow \mathbb{R}^N$. We note that we calculate u^{k+1} using only components of the vector u^k of the old iteration level:

$$u_i^{k+1} := F^i(u^k) \quad \text{for all } i \in J, k \in \mathbb{N}_0. \quad (2.14)$$

Such a method can also be called a *nonlinear Jacobi method*.

Let us now state two important results.

Proposition 2.1.1 (Maximum-Minimum Principle) [73]

With the assumptions on Ψ_D , Ψ_S , w_D , and w_S as above, the scheme (2.13) satisfies a maximum-minimum principle:

$$\min_{j \in J} f_j \leq u_i^k \leq \max_{j \in J} f_j \quad \text{for all } i \in J, k \in \mathbb{N}. \quad (2.15)$$

Proof: With our assumptions on the tonal and spatial weights from above we know that $d_{i,j}^k \geq 0$ and $s_{i,j}^k \geq 0$ for all i, j, k . That means in (2.13), u_i^{k+1} is calculated as a convex combination of grey values of the initial image f and of the last iteration step u^k . Thus we have

$$\min_{j \in J} \{u_j^k, f_j\} \leq u_i^{k+1} \leq \max_{j \in J} \{u_j^k, f_j\} \quad \text{for all } i \in J, k \in \mathbb{N}. \quad (2.16)$$

Induction shows that the fixed point scheme (2.13) satisfies a maximum-minimum principle, i. e.

$$\min_{j \in J} f_j \leq u_i^k \leq \max_{j \in J} f_j \quad \text{for all } i \in J, k \in \mathbb{N}. \quad (2.17)$$

□

In the next proposition, we see that this property is not only useful from a practical point of view: Together with continuity, it gives us the existence of a fixed point.

Proposition 2.1.2 (Existence of a Fixed Point) [73]

The fixed point equation (2.11) has a solution.

Proof: Let us consider the set $M := \{u \in \mathbb{R}^N \mid \|u\|_\infty \leq \|f\|_\infty\}$ with the norm $\|u\|_\infty := \max_{j \in J} |u_j|$. M is nonempty, compact and convex. Then the maximum-minimum stability implies that $F(M) \subseteq M$. With our requirements on the tonal and spatial weights, the denominator in (2.13) is always larger than zero. This means that each component $F_i : \mathbb{R}^N \rightarrow \mathbb{R}$ is continuous with respect to the norm $\|\cdot\|_\infty$. Since this holds for all i , we know that $F : (\mathbb{R}^N, \|\cdot\|_\infty) \rightarrow (\mathbb{R}^N, \|\cdot\|_\infty)$ is continuous. Then Brouwer's fixed point theorem (see [36] or [238, page 51], for example) shows that F has a fixed point in M . □

From the derivation it is clear that a fixed point corresponds to a critical point of E . If we have chosen our penaliser functions such that the energy functional is strictly convex, this is equivalent to the unique minimum of E .

Alternatively, the solution of the NDS energy (2.3) can be obtained by gradient descent optimisation:

$$\frac{u_i^{k+1} - u_i^k}{\tau} = -\frac{\partial E}{\partial u_i^k} \quad \text{for all } i \in J, \quad (2.18)$$

with step size $\tau > 0$. Considering (2.7)–(2.11), the energy minimiser is computed as

$$u_i^0 := f_i, \quad (2.19)$$

$$u_i^{k+1} := (1 - \tau)u_i^k + \tau \frac{(1 - \alpha) \sum_{j \in J} d_{i,j}^k f_j + \alpha \sum_{j \in J} s_{i,j}^k u_j^k}{(1 - \alpha) \sum_{j \in J} d_{i,j}^k + \alpha \sum_{j \in J} s_{i,j}^k}. \quad (2.20)$$

Note that by setting $\tau = 1$ one obtains the fixed point iteration (2.13).

2.2 Important Special Cases

Recall that the NDS functional (2.3) can be optimised using the fixed point iterations (2.11). Let us introduce the following notation for the tonal weights,

$$g_{i,j}^D := \Psi'_D(|u_i - f_j|^2), \quad (2.21)$$

$$g_{i,j}^S := 2\Psi'_S(|u_i - u_j|^2), \quad (2.22)$$

and for the spatial weights,

$$w_{i,j}^{D,\varpi} := w_D(|x_i - x_j|^2), \quad (2.23)$$

$$w_{i,j}^{S,\varpi} := w_S(|x_i - x_j|^2), \quad (2.24)$$

where the spatial weights w implicitly contain a scale parameter ϖ specifying how quickly the weight decreases: Small ϖ means a local operation (or a smaller window), larger ϖ leads to operations with large-scale effects. The window sizes for the data and smoothness terms may differ.

Using this notation, equation (2.11) can be written as

$$u_i = \frac{(1 - \alpha) \sum_{j \in J} g_{i,j}^D w_{i,j}^{D,\varpi} f_j + \alpha \sum_{j \in J} g_{i,j}^S w_{i,j}^{S,\varpi} u_j}{(1 - \alpha) \sum_{j \in J} g_{i,j}^D w_{i,j}^{D,\varpi} + \alpha \sum_{j \in J} g_{i,j}^S w_{i,j}^{S,\varpi}}. \quad (2.25)$$

In the following sections we show that many well known filtering and estimation methods can be derived from equation (2.25) (and thus from the NDS functional) by a simple choice of the parameter α which balances the smoothness and data terms, the window size ϖ , and by an appropriate selection of the weighting functions g^D , g^S , w^D , and w^S . Fig. 2.1 gives an overview of the NDS landscape and the methods covered below.

2.2.1 M-Estimators and Local M-Smothers

When estimating the underlying constant signal from noisy samples, the selected method should depend on the type of noise present in the data. For Gaussian noise, taking the sample mean is a good choice, providing the maximum a posteriori (MAP) estimate. For noise with heavier tails (caused either by the noise properties themselves, or because e.g. the samples were mixed from two distributions due to signal discontinuity), one has to use methods from robust statistics that are less effected by outliers, such as an M-estimator [108, 99]. An M-estimate of a constant value u from noisy data f_j is found by minimising

$$E(u) = \sum_{j=1}^K \Psi(|u - f_j|^2). \quad (2.26)$$

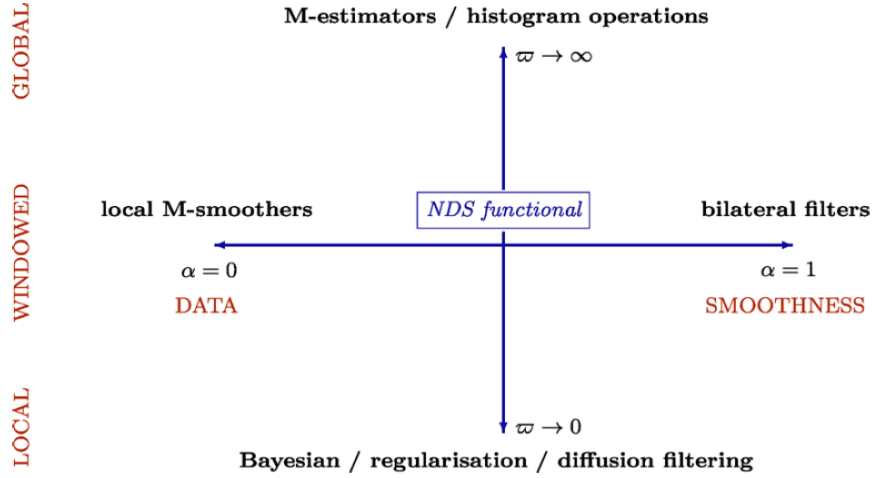


Figure 2.1: Overview of methods covered by the NDS functional. *Left to right*, the parameter α balances the data and smoothness terms. *Bottom to top*, the window size ϖ determines the spatial extent of the methods, from local to global.

Some of the error penalisers Ψ previously shown in Table 2.1 lead to a direct statistical interpretation of the element that minimises the functional (2.26). For the L_2 norm (a), the solution is the mean of the noisy samples. The L_1 norm (b) is minimised by the median. For the robust error norms (f) and (g), the influence of outliers is very much reduced, and the solution u minimising (2.26) approximates a mode (maximum) of the probability density underlying the noisy samples. The mode ideally corresponds to the most frequent value present in the data. For the discrete noisy samples, the maximum of the density can be only estimated e.g. using suitable smoothing kernels; see [65] for some examples and a connection to iterative solvers. Note that while the L_2 and L_1 norms lead to a convex functional minimisation, the robust error norms (f) and (g) in Table 2.1 are nonconvex, and their corresponding functionals $E(u)$ may exhibit multiple local minima.

The M-estimators were introduced to robustly estimate a single value from noisy samples. For images, we have to consider also the spatial distribution of the data. Such a generalisation is known as *local M-smoothers*, and the functional to minimise has the following structure [66, 236]:

$$E(u) = \sum_{i=1}^N \sum_{j \in \mathcal{B}(i)} \Psi(|u_i - f_j|^2) w(|x_i - x_j|^2) \quad (2.27)$$

where w represents the spatial weight depending on sample distance. The local window $\mathcal{B}(i)$ is introduced in (2.27) for computational convenience only, to make the index j run through the neighbourhood of x_i where $w(|x_i - x_j|^2)$ exceeds some threshold of contribution importance.

The energy functional (2.27) can be minimised using an iterative scheme called *W-estimator* [236],

$$u_i^{k+1} = \frac{\sum_{j \in \mathcal{B}(i)} g(|u_i^k - f_j|^2) w(|x_i - x_j|^2) f_j}{\sum_{j \in \mathcal{B}(i)} g(|u_i^k - f_j|^2) w(|x_i - x_j|^2)} \quad (2.28)$$

where the process is initialised with $u_i^0 := f_i$. This iterative scheme converges to a local minimum of (2.27) close to the input data. Depending on the penaliser Ψ , the iterations may lead e.g. to a local mode approximation [92, 214, 210, 213, 67], or to an approximation of a windowed median filter or Gaussian smoothing.

Comparing equation (2.28) with our scheme (2.25), we observe that the local M-smoothers and the W-estimator correspond to the data term of the NDS model. To obtain the W-estimator from (2.25), simply set the smoothness parameter $\alpha = 0$. The spatial weight w will be chosen so that it covers some area around the current pixel, typically larger than the immediate neighbourhood.

2.2.2 Bilateral Filtering

Contrary to the previous section, let us analyse the situation for the maximum smoothness parameter, $\alpha = 1$. Then, the data term from equations (2.3) and (2.25) vanishes, and the full scheme consists of the smoothness term only. The resulting energy functional

$$E_S(u) = \sum_{i,j \in J} \Psi_S(|u_i - u_j|^2) w_S(|x_i - x_j|^2) \quad (2.29)$$

can be minimised by the fixed point iterations

$$u_i^{k+1} = \frac{\sum_{j \in J} g_S(|u_i^k - u_j^k|^2) w_S(|x_i - x_j|^2) u_j^k}{\sum_{j \in J} g_S(|u_i^k - u_j^k|^2) w_S(|x_i - x_j|^2)}. \quad (2.30)$$

Equation (2.30) is known as *bilateral filter* [8, 195, 206]. While bilateral filtering was originally proposed as a heuristic algorithm, we have shown that it can be derived as a special case from the NDS energy functional (2.3) where only the smoothness term is considered, and the local smoothness of the signal u is evaluated in a nonlocal window w_S .

2.2.3 Regularisation Methods

In this section we reveal novel insights into relations between the NDS filter and diffusion/regularisation methods.

Consider the optimality condition $\frac{\partial E_S}{\partial u_k} = 0$ for (2.2) written as

$$0 = \sum_{j \in \mathcal{B}(k)} \Psi_S(|u_k - u_{k+j}|^2) (u_k - u_{k+j}) w_S(|x_k - x_{k+j}|^2), \quad (2.31)$$

with $\mathcal{B}(k) = \{j \in \mathbb{Z} : |x_k - x_{k+j}| \leq \varpi_S\}$, and the hard window

$$w_S(|x_k - x_{k+j}|^2) = \begin{cases} 1 & \text{if } |x_k - x_{k+j}| \leq \varpi_S \\ 0 & \text{otherwise} \end{cases}.$$

Then, following [218, 75], equation (2.31) can be regarded as a crude approximation of the steady state of the rotationally invariant PDE

$$\partial_t u = \frac{2}{\pi} \int_0^\pi \partial_{e_\varphi} \left(g_S(|\partial_{e_\varphi} u_\delta|^2) \partial_{e_\varphi} u \right) d\varphi \quad (2.32)$$

when the kernel size δ in $u_\delta := G_\delta * u$ vanishes, $g_S := \Psi'_S$, and $e_\varphi = (\cos \varphi, \sin \varphi)^\top$. Furthermore, equation (2.32) is equivalent to the anisotropic model

$$\partial_t u = \operatorname{div}(D \cdot \nabla u) \quad (2.33)$$

with the diffusion tensor

$$D := \frac{2}{\pi} \int_0^\pi e_\varphi e_\varphi^\top g_S(|\partial_{e_\varphi} u_\delta|^2) d\varphi.$$

In [218] it is shown that the eigenvectors and eigenvalues of D are given by

$$v_{\parallel}(\varphi) = \begin{pmatrix} -\sin \varphi \\ \cos \varphi \end{pmatrix}, \quad v_{\perp}(\varphi) = \begin{pmatrix} \cos \varphi \\ \sin \varphi \end{pmatrix};$$

$$\lambda_{\parallel}(r^2) = \frac{4}{\pi} \int_0^{\pi/2} \sin^2 \varphi g_S(|r \cos \varphi|^2) d\varphi,$$

$$\lambda_{\perp}(r^2) = \frac{4}{\pi} \int_0^{\pi/2} \cos^2 \varphi g_S(|r \cos \varphi|^2) d\varphi,$$

where (r, φ) are the polar coordinates of ∇u . In our case, i.e. $\delta \rightarrow 0$, the process (2.33) becomes isotropic with scalar diffusivity $\tilde{g} := \lambda_{\perp}$:

$$\partial_t u = \operatorname{div}(\tilde{g} \cdot \nabla u). \quad (2.34)$$

This means that the solution of the smoothness term (2.2) approximates a Perona-Malik filter that diffuses in direction ∇u^\perp perpendicular to the gradient, i.e. along edges.

If we now include the data term (2.1) with a local spatial window

$$w_D(|x_k - x_{k+j}|^2) = \begin{cases} 1 & \text{if } x_k = x_{k+j} \\ 0 & \text{otherwise} \end{cases},$$

the resulting process approximates

$$\frac{u - f}{\frac{2\alpha}{1-\alpha}} = \operatorname{div}(\tilde{g} \cdot \nabla u), \quad (2.35)$$

which can be regarded as a fully implicit time discretisation of the diffusion process (2.34) with a single time step of size $\frac{2\alpha}{1-\alpha} > 0$. Following Scherzer and Weickert [184], it can be shown that (2.35) corresponds to the Euler-Lagrange equation of the continuous functional

$$E(u) = \int_{\Omega} \left((1 - \alpha) \Psi_D(|u - f|^2) + \alpha \tilde{\Psi}_S(|\nabla u|^2) \right) dx, \quad (2.36)$$

where $\tilde{\Psi}'_S := \tilde{g}$. The continuous functional is the classical energy functional from regularisation or Bayesian frameworks; see e.g. [19, 84, 149, 236]. As an example, the continuous Mumford–Shah functional fits into this framework if we choose $\Psi_D(s^2) := s^2$ and $\Psi_S(s^2) := \min(s^2, \lambda^2)$. Also, the diffusion filters [167, 219] and diffusion-reaction processes [156, 185, 199, 60] can be derived from equation (2.36).

2.2.4 Histogram Quantisation

For the sake of completeness, let us consider the case when the spatial support window grows to ‘infinite’ size, and all the pixels are connected with the same weight regardless of their position in the image, $w_{D,S} \equiv 1$. Then, the NDS functional simplifies to

$$E(u) = (1 - \alpha) \sum_{i,j \in J} \Psi_D(|u_i - f_j|^2) + \alpha \sum_{i,j \in J} \Psi_S(|u_i - u_j|^2).$$

Because the spatial information does not appear in the formula, the solution can be equivalently found in a space where the spatial information was omitted and only the tonal information remains: the image histogram. For example, minimizing the functional for the robust penaliser Ψ from Table 2.1 (f) or (g) corresponds to replacing each pixel value with the local mode of the corresponding image histogram. The resulting image will have a smaller number of gray values, adaptively quantised. The data and smoothness terms in this context correspond to the non-blurring or blurring mean shift process, respectively [65].

2.3 NDS vs Graph Regularisation

In this section we show that the discrete NDS framework is closely related to graph regularisation techniques and that it extends recent developments in the context of image and manifold regularisation on weighted graphs.

A discrete image is usually defined on a regular domain, e.g. on a rectangular grid. However, for more general image domains it is more appropriate to represent an image as a graph with arbitrary topology. Every vertex (pixel) i of the graph encodes both the pixel location x_i and the pixel intensity f_i . The edge connecting two vertices i and j represents the similarity between both pixels, expressed as a weight function $w(i, j) > 0$. Employing such graph representation and special calculus on graphs [240, 241], several regularisation models for general data living on discrete spaces have been recently proposed. In the context of image denoising Weickert [220] developed a space-discrete theory for diffusion filtering that is directly applicable to functions defined on graphs, and Chan *et al.* [57] introduced the digital TV filter as a discrete version of the continuous ROF model [179]. In the context of semi-supervised learning Zhou and Schölkopf [240, 241] proposed a discrete analogue of classical regularisation [205] with a p -Dirichlet regulariser; and Zhou and Burges [239] introduced a discrete analogue of the Laplace-de Rham operator as a regulariser.

Following the ideas from graph theory presented in [240, 241], Gilboa and Osher [88] proposed the use of nonlocal operators to extend some known PDEs and variational techniques in image processing to a nonlocal framework. In particular, they use discretised differential operators such as gradient and divergence. The discretisations involve pixel differences that are weighted by a patch-based similarity between pixels as in [41]. Bougleux *et al.* [29, 78, 30] designed a discrete graph regularisation framework that can be seen as a digital extension of the continuous framework [88] employing a p -Dirichlet regulariser. The same discrete framework has been applied in image segmentation tasks [202]. Furthermore, nonlocal differential operators have been used to derive nonlocal morphological PDEs [79].

We now show that the discrete variational NDS model (2.3) can be regarded as a common regularisation method for general data defined on discrete spaces. Let us consider the smoothness term (2.2) of the NDS model using $\Psi(s^2) = \frac{1}{p}|s|^p$, $p > 0$, as penaliser:

$$\begin{aligned}
 E_S(u) &= \sum_{i \in J} \sum_{j \in J} \Psi(|u_i - u_j|^2) w(i, j) \\
 &= \frac{1}{p} \sum_{i \in J} \sum_{j \in J} |u_i - u_j|^p w(i, j) \\
 &= \frac{1}{p} \sum_{i \in J} \|\nabla_w u_i\|_p^p, \tag{2.37}
 \end{aligned}$$

where $\|\nabla_w u_i\|_p = \left(\sum_{j \in J} |u_i - u_j|^p w(i, j) \right)^{\frac{1}{p}}$ is the weighted L_p norm. Other definitions of the weighted gradient norm are possible using alternative weighted difference operators (see [104] and references therein). The regulariser (2.37) has been used in [241, 29, 78, 30] for regularisation on arbitrary graphs. In particular, the following energy functionals have been proposed in [30]:

$$E_{iso}(u) = \sum_{i \in J} \left(\frac{\lambda}{2} (u_i - f_i)^2 + \frac{1}{p} \|\nabla_w u_i\|_2^p \right), \quad (2.38)$$

$$E_{ani}(u) = \sum_{i \in J} \left(\frac{\lambda}{2} (u_i - f_i)^2 + \frac{1}{2p} \|\nabla_w u_i\|_p^p \right). \quad (2.39)$$

The functional (2.38) corresponds to an isotropic model whose minimiser is obtained by solving a linear system, whereas (2.39) is an anisotropic model leading to a nonlinear system. The nonlocal interactions between graph nodes are introduced via the weight function w . In the general case the weight $w(i, j) := w(\mathcal{F}_i, \mathcal{F}_j)$ measures the similarity between the nodes i and j with respect to a certain feature vector \mathcal{F} . For instance, a weighted L_2 norm between image patches [41] can be used for the task of image smoothing.

There exist three main differences between the NDS framework and the graph regularisation (GR) approach:

- (i) in the NDS we allow the use of any penaliser for both the data similarity and the smoothness term, whereas GR only considers penalisers of the form $\Psi(s^2) = \frac{1}{p} |s|^p$ for $p \in]0, 2]$;
- (ii) in the NDS model nonlocal interactions are present in both the data and the smoothness term, while in the GR techniques the non-localities are only considered in the regularisation term; and
- (iii) in the NDS framework the functions w only depend on the spatial node/pixel locations, whereas in the GR approaches w can be defined in terms of several node characteristics.

The point (iii) suggests that the NDS model (2.3) can be generalised by extending the definition of the weighting functions w . However, we do not further develop this idea here. That will be part of future work. In Chapter 3 we will consider another generalisation of the NDS framework, where we rather concentrate on the penalisers Ψ , which we allow to act on more general constraints.

2.4 Tuning the NDS Parameters

In [148], Mrázek *et al.* performed preliminary experiments to show the smoothing properties of the NDS model with data and smoothness constraints acting over neighbourhoods of varying size. In [74], Didas *et al.* compared several minimisation strategies for the energy functional. Here we focus on two main issues: (i) show that the NDS model can outperform a wide range of well-known filters, and (ii) study the relations among its smoothing parameters. Before developing these points, we briefly describe the classical statistical interpretation of energy-based image restoration models, which sheds some light on the selection of the energy penalisers [173].

2.4.1 Statistically-Based Selection of Tonal Kernels

The most common degradation model is given by $f = u + \eta$, where u is the true image, η represents a zero-mean additive noise with standard deviation σ , and f is the recorded image. In the following we consider these quantities as realisations of the random variables U , η , and F , denoting as p_U , p_η , and p_F their probability density function (pdf), respectively. In Bayesian analysis [109, 84, 235], the *maximum a posteriori* (MAP) estimator

$$\begin{aligned}\hat{u}_{\text{MAP}} &= \arg \max_u \log p_{U|F}(u|f) \\ &= \arg \min_u (-\log p_{F|U}(f|u) - \log p_U(u))\end{aligned}\quad (2.40)$$

yields the most likely image u given f . The conditional distribution $p_{F|U}$, also called *likelihood*, models the degradation process of U to F and is therefore considered as the noise distribution, i. e. $p_{F|U}(f|u) = p_\eta(\eta) = \prod p_\eta(f_i - u_i)$, assuming that the noise is independent and identically distributed (i.i.d.). We will focus on Gaussian and impulse noise. These types of noise are well modelled by the *Gaussian* and *Laplacian* distributions, being respectively their pdf's

$$p_{\eta_G} = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^n \exp\left(-\frac{1}{2\sigma^2} \sum_{i \in \Omega} |\eta_i|^2\right), \quad (2.41)$$

$$p_{\eta_I} = \left(\frac{1}{2\sigma}\right)^n \exp\left(-\frac{1}{\sigma} \sum_{i \in \Omega} |\eta_i|\right). \quad (2.42)$$

Plugging these noise models into the MAP estimator (2.40), and observing the structural resemblance to our NDS function (2.3), suggests that the penalisers for the data term can be instantiated as $\Psi_D(s^2) = s^2$ for Gaussian noise, and as $\Psi_D(s^2) = |s|$ for impulse noise.

The previous noise distributions are special cases of a more general probabilistic law: the *generalised Gaussian* distribution [105], with parameters mean μ , variance

σ^2 , and $\nu > 0$ (Gaussian case $\nu = 2$, Laplacian case $\nu = 1$); and pdf

$$p_Z(z) = \frac{\nu\Gamma(3/\nu)^{1/2}}{2\sigma\Gamma(1/\nu)^{3/2}} \exp\left(-\frac{|z-\mu|^\nu}{\sigma^\nu} \left(\frac{\Gamma(3/\nu)}{\Gamma(1/\nu)}\right)^{\nu/2}\right), \quad (2.43)$$

where $\Gamma(\cdot)$ is the Euler Gamma function. This distribution has been also utilised for modelling probabilistic *prior* knowledge about the signal u to recover. In (2.40), this information is represented in terms of the *prior* distribution p_U of the grey values of U . Besag [20] proposed the Laplacian law as model for p_U , which was later extended by Bouman & Sauer with their *generalised Gaussian Markov random field* [31] based on the distribution (2.43) for $\nu \in [1, 2]$. Since choosing a particular model for the prior distribution is essentially equivalent to specify the penaliser Ψ_S for the smoothness term in our NDS framework, we can instantiate such tonal kernel as $\Psi_S(s^2) = |s|^\nu$. This function is nonconvex for $0 < \nu < 1$, what may give rise to local minima in (2.3). However, nonconvex penalisers can allow almost exact restoration quality [122, 152, 153].

The Bayesian framework thus provides a founded basis for choosing appropriate tonal kernels $\Psi(\cdot)$ for the data and smoothness terms in (2.3). Studying other types of noise and the properties of the signal to recover will lead to different criteria for selecting the penalisers. It is interesting to note that Nikolova [154, 155] has recently shown that classical MAP estimators introduce distortions to the assumed models for the noise and the signal to recover.

2.4.2 Linear Combination of Kernels

The problem of determining the regularisation parameter α in (2.3) is crucial to obtain an optimal balance between data similarity and smoothness. We intend to justify the use of such framework for $\alpha \notin \{0, 1\}$, i. e. for a wider spectrum of filters than those special cases outlined in Table 2.1.

A function $\varphi : [0, 1] \rightarrow \mathbb{R}$ is called *unimodal* on $[0, 1]$ if it contains a single minimum in that interval. Then, we obtain an estimate of the true image u as

$$\hat{u} = \arg \min_{\alpha} \varphi(\alpha), \quad (2.44)$$

assuming that we deal with the unimodal function $\varphi(\alpha) := \|u - u_\alpha\|_1$, where u_α is the solution image for a specific value of α from (2.3). Exploiting the empirical unimodality¹ of φ on $[0, 1]$, we employ the *Fibonacci Search* method to find an optimal value for α that solves (2.44). This line-search strategy ensures fast convergence. For *multimodal* functions it is better to utilise the *Simulated Annealing* [118, 55] technique as minimization strategy, which guarantees finding a global minimum in finite time. See [15] for the implementation details.

1. Even though we can guarantee neither the continuous dependence of φ with respect to α nor a unique solution, we have observed this behavior in most of our experiments.




Figure 2.2: Signal denoising with the NDS functional. Original signals in solid lines, noisy and denoised signals in dashed lines. (a) Noisy signal perturbed by zero-mean Gaussian noise with $\sigma = 40$, $L_1 = 27.30$ and (b) its denoised version $L_1 = 13.83$. (c) Noisy signal perturbed by 40% of salt-and-pepper noise, $L_1 = 48.04$ and (d) its denoised version $L_1 = 4.61$.

As suggested by the statistical framework, our designed nonlocal filters for Gaussian noise and impulse noise read

$$E(u) = (1 - \alpha) \sum_{i \in J, j \in \mathcal{B}_D(i)} |u_i - f_j|^2 + \alpha \sum_{i \in J, j \in \mathcal{B}_S(i)} |u_i - u_j|^\nu \quad (2.45)$$

and

$$E(u) = (1 - \alpha) \sum_{i \in J, j \in \mathcal{B}_D(i)} |u_i - f_j| + \alpha \sum_{i \in J, j \in \mathcal{B}_S(i)} |u_i - u_j|^\nu, \quad (2.46)$$

respectively, where \mathcal{B}_\varnothing is the disk-shaped hard window function used as spatial kernel with radius ϖ . Here we focus on $\nu \in \{1, 2\}$. We apply these models to reconstruct the noisy signals depicted in Fig. 2.2. All parameters were optimised and the best five parameterisation for each model are shown in Table 2.3, with L_1 denoting the absolute difference between the original (uncorrupted) signal and the denoised version. We

Table 2.3: Quantitative comparison of different filters. *Left*: Denoising results of the signal perturbed with Gaussian noise shown in Fig. 2.2(a). *Right*: Denoising results of the signal perturbed with salt-and-pepper noise shown in Fig. 2.2(c). The best results are written in bold letters and plotted in Fig. 2.2(b) and Fig. 2.2(d), respectively.

Filter	ϖ_D	ϖ_S	α	L_1	Filter	ϖ_D	ϖ_S	α	L_1
model (2.45), $\nu = 2$	3	1	0.78	14.18	model (2.46), $\nu = 2$	0	1	0.10	4.61
	3	2	0.43	14.24		3	1	0.19	4.67
	2	1	0.79	14.26		3	2	0.06	4.80
	2	2	0.46	14.30		4	1	0.21	4.90
	3	3	0.21	14.32		2	1	0.24	4.91
model (2.45), $\nu = 1$	2	2	0.93	13.83	model (2.46), $\nu = 1$	3	8	0.11	5.20
	2	3	0.87	13.83		3	9	0.09	5.27
	2	4	0.82	13.85		3	10	0.08	5.38
	2	5	0.78	13.93		4	9	0.08	5.59
	3	2	0.90	14.00		3	7	0.11	5.51
mean	4	-	0.00	14.93	mean	6	-	0.00	23.95
median	4	-	0.00	14.90	median	6	-	0.00	6.98
Tikhonov	0	1	0.67	14.57	Tikhonov	0	1	0.90	23.22
TV	0	1	0.99	15.62	TV	0	1	0.99	35.04
Perona-Malik	0	1	0.70	14.47	Perona-Malik	0	1	0.90	23.21
Charbonnier	0	1	0.69	14.53	Charbonnier	0	1	0.90	23.21

also report on the performance of the mean and median filters as representatives of M-smoothers (Section 2.2.1), and classical regularisation filtering (Section 2.2.3) with four different penalisers. Without exceptions, our designed models outperform all the well known filters obtained as particular cases of the NDS framework. Note that we have deliberately chosen penalisers that do not require any contrast parameter. This keeps our models simple and efficient.

2.4.3 Smoothing Effects

Trade-off between α and ϖ_S

As it is noticeable in Table 2.3 there exists a trade-off between the parameter α and the radii ϖ_D, ϖ_S of the spatial kernels. For example, it is possible to achieve similar filtering results either by decreasing α or by increasing ϖ_S . On the one hand decreasing α reduces the influence of the smoothness term, but on the other, increasing ϖ_S considers contributions to the smoothness term from a larger neighbourhood. To illustrate this effect let us consider the original image shown in Fig. 2.3(a)-top left

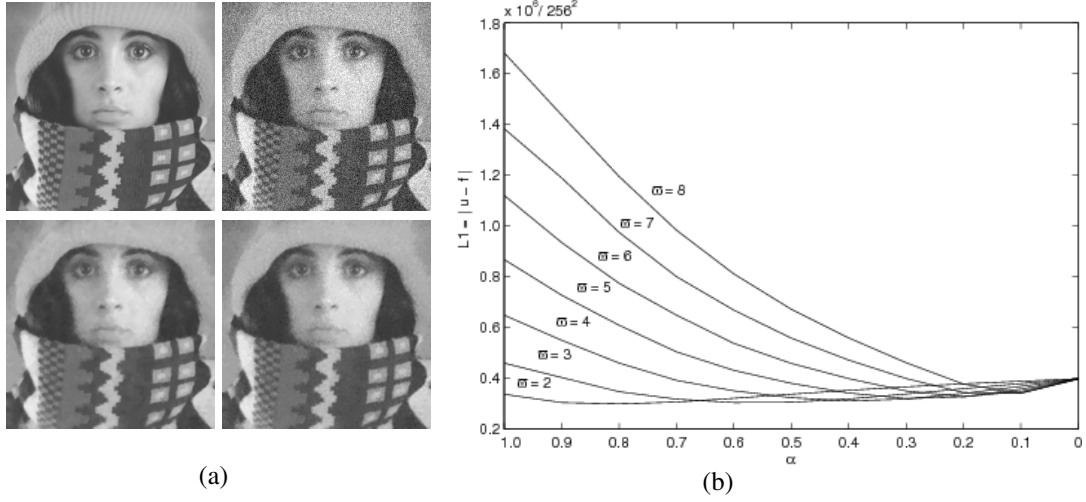


Figure 2.3: Trade-off between α and ϖ_S . (a) Original image (*top left*); noisy image perturbed with zero-mean Gaussian with $\sigma = 20$, $L_1 = 16.02$ (*top right*); restored image with $\alpha = 0.8$, $\varpi_S = 2$, $L_1 = 4.88$ (*bottom left*); restored image with $\alpha = 0.2$, $\varpi_S = 6$, $L_1 = 5.18$ (*bottom right*). (b) L_1 distance between the original and the denoised image for different values of α and ϖ_S .

and its degraded version with Gaussian noise of Fig. 2.3(a)-*top right* that we restore employing the model (2.45) with $\nu = 1$. The radius ϖ_D of the spatial kernel in the data term was fixed to 1. Fig. 2.3(b) displays the filtering results for a range of values α and ϖ_S . Fig. 2.3(a)-*bottom* shows examples where similar restoration quality is achieved under different parameterisation. Moreover, slightly better results are attained for α large and ϖ_S small, which implies less operations and more efficiency.

Trade-off between α and ϖ_D

If we consider a functional which only consists of a data term, we notice that increasing the support of the spatial window leads to smoothing. On the other hand, if we leave the spatial window of the data term small and add a smoothness term, this has visually almost the same effect. In this experiment we want to quantify the difference more accurately and search for α corresponding to a certain support for w_D . To this end we consider the two following functions. The first function

$$E_D(u) = \sum_{i,j \in \Omega} (u_i - f_j)^2 w_D(|x_i - x_j|^2) \quad (2.47)$$

consists only of a data term, but allows for a larger window given by the disc-shaped hard window function w_D with radius ϖ_D . The second function has a local data term and a smoothness term which only takes the direct neighbours $\mathcal{N}(i)$ of pixel i into




Figure 2.4: *Top left:* Original image, 256×256 pixels. *Top middle:* Smoothed version with E_D , radius $\varpi_D = 5$. *Top right:* Smoothing with E_C , $\alpha = 0.50$. *Bottom left:* Image with additive Gaussian noise, standard deviation $\sigma = 50$. *Bottom middle:* Denoising with E_D , radius $\varpi_D = 5$. *Bottom right:* Denoising with E_C , $\alpha = 0.58$.

consideration

$$E_C(u) = (1 - \alpha) \sum_{i \in \Omega} (u_i - f_i)^2 + \alpha \sum_{i \in \Omega, j \in \mathcal{N}(i)} (u_i - u_j)^2 \quad (2.48)$$

Here, only changing the value α is used to steer the amount of smoothness. Fig. 2.4 shows two examples of the trade-off between these parameters. We see that in both cases with and without noise it is possible to obtain very similar results with both functions. We are interested in knowing how far away the results obtained with the functions E_D and E_C are from each other, and how they approach each other by means of tuning ϖ_D and α , respectively. In Fig. 2.5, we display some measurements to quantify the trade-off between these parameters. In the left column, each curve stands for a certain radius size, and there is one value of α that minimises the L_1 distance between their estimates. The minimum distance achieved for every pair (ϖ_D, α) is displayed in the right column. These curves confirm our intuition that although the functions E_D and E_C are completely different, it is possible that both models produce similar results when their tuning parameters are accordingly adjusted.




Figure 2.5: Trade-off between α and ϖ_D . *Top row:* Noiseless case. L_1 -error curves between the models E_D and E_C for varying α and ϖ_D (*left*), and optimal values that minimise the L_1 -error (*right*). *Bottom row:* Same plots for the noisy case.

2.5 Summary

In this chapter we have shown the NDS model as a framework that unifies a number of filtering approaches from the literature and have provided novel connections to recent graph regularisation methods. The flexibility of the NDS model allowed us: First, to design hybrid filters that outperform the others obtained as special cases. Second, to establish parametric correspondences between different filters.

Despite the fact that the NDS framework allows for nonlocal pixel interactions by extending the support of the spatial windows $w_{\{D,S\}}$, note in Table 2.3 that the radii ϖ_D and ϖ_S do not take very large values for the best denoising results. This suggests that the effective utilisation of larger neighbourhoods is hindered in practice. In Chapter 3 this issue will be further discussed and a generalisation of the NDS framework will be proposed that overcomes this limitation.

3

Generalised NDS Framework

The NDS model of the previous chapter was termed *nonlocal* data and smoothness (NDS) because of the interactions between more distant pixels than the immediate neighbourhood. However, the tonal weights in depend on the single differences between pairs of connected pixels. These single differences have a limited ability to express local image structure and geometry, and for practical purposes, the pixel interactions have to be kept to a relatively small neighbourhood.

Many recent approaches for image denoising make use of self-similarity of the whole image, or similarity between several images. For filtering, pixels from very distant locations could also contribute to the result. To distinguish which pixels are compatible, a more powerful measure is needed to evaluate the similarity: Not just pixel difference, but the similarity of a whole region of interest, or *image patch* around the central pixel, is considered. Two equivalent filters such as NL-means [40, 41] and UINTA [9, 10] are typical examples of this class of filters.

In this chapter, we combine the idea of patch similarity with the NDS functional, which leads to a *Generalised Nonlocal Data and Smoothness*, or GNDS model. We keep the discrete variational framework involving both data and smoothness terms, and allow for different ways to calculate the distance of the image patches. We will show which iterative filter can be derived as a minimiser of the GNDS energy functional. Inspired by its form, we will relax a constraint and present a new family of patch-based GNDS filters. This chapter is based on [174].

3.1 GNDS Functional and its Minimisation

First, let us introduce the tonal distance functions $d_D, d_S : \mathbb{R}^{2p} \rightarrow \mathbb{R}_0^+$ in the data and the smoothness term. For example, in the data term, such a function calculates the distance between two image patches $u(\mathcal{P}_i)$ of the evolving image and $f(\mathcal{P}_j)$ of the initial image. The index sets \mathcal{P}_i and \mathcal{P}_j define image patches as neighbourhoods of the pixels i and j , respectively. Both patches are assumed to have the same size $p \in \mathbb{N}$ and the same shape.

As distance function, for example the weighted L_2 norm can be used, i. e.

$$|d(u(\mathcal{P}_i), f(\mathcal{P}_j))|^2 = \sum_p G_\sigma(p) (u_{i+p} - f_{j+p})^2, \quad (3.1)$$

where $G_\sigma(p) := \exp(-p^2/(2\sigma^2))$. This has also been used as a patch distance in the NL-means algorithm.

With these definitions, the Generalised Nonlocal Data and Smoothness (GNDS) model reads

$$\begin{aligned} E_G(u) &= (1 - \alpha) E_{GD}(u) + \alpha E_{GS}(u) \\ &= (1 - \alpha) \sum_{i,j \in J} \Psi_D \left(|d_D(u(\mathcal{P}_i), f(\mathcal{P}_j))|^2 \right) w_D(|x_i - x_j|^2) \\ &\quad + \alpha \sum_{i,j \in J} \Psi_S \left(|d_S(u(\mathcal{P}_i), u(\mathcal{P}_j))|^2 \right) w_S(|x_i - x_j|^2). \end{aligned} \quad (3.2)$$

As we did for the NDS model, we now obtain the corresponding fixed point form for (3.2). The minimiser u of (3.2) necessarily satisfies

$$\frac{\partial E_G}{\partial u_i} = (1 - \alpha) \frac{\partial E_{GD}}{\partial u_i} + \alpha \frac{\partial E_{GS}}{\partial u_i} = 0 \quad \text{for all } i \in J. \quad (3.3)$$

Using the distance function (3.1) in both the data and the smoothness terms, we have:

$$\frac{\partial E_{GD}}{\partial u_k} = 2 \sum_{j \in J} G_\sigma * \Psi'_D \left(d_{D;k-\cdot, j-\cdot}^2 \right) (0) (u_k - f_j) w_D(|x_k - x_j|^2) \quad (3.4)$$

$$\frac{\partial E_{GS}}{\partial u_k} = 4 \sum_{j \in J} G_\sigma * \Psi'_S \left(d_{S;k-\cdot, j-\cdot}^2 \right) (0) (u_k - u_j) w_S(|x_k - x_j|^2) \quad (3.5)$$

where the operator ‘*’ stands for convolution. Then, with help of the abbreviations

$$g_{i,j}^{GD} := G_\sigma * \Psi'_D \left(|d_D(u(\mathcal{P}_{i-\cdot}), f(\mathcal{P}_{j-\cdot}))|^2 \right) (0) \quad (3.6)$$

$$g_{i,j}^{GS} := 2 G_\sigma * \Psi'_S \left(|d_S(u(\mathcal{P}_{i-\cdot}), u(\mathcal{P}_{j-\cdot}))|^2 \right) (0), \quad (3.7)$$

and with the spatial weights defined as in (2.23)-(2.24), the fixed point for the GNDS model reads

$$u_i = \frac{(1 - \alpha) \sum_{j \in J} g_{i,j}^{GD} w_{i,j}^{D,\varpi} f_j + \alpha \sum_{j \in J} g_{i,j}^{GS} w_{i,j}^{S,\varpi} u_j}{(1 - \alpha) \sum_{j \in J} g_{i,j}^{GD} w_{i,j}^{D,\varpi} + \alpha \sum_{j \in J} g_{i,j}^{GS} w_{i,j}^{S,\varpi}}, \quad (3.8)$$

for all $i \in J$. This equation can be embedded in a fixed point iteration scheme similar to (2.12)-(2.14). A maximum-minimum principle and the existence of a fixed point can be proven following Propositions 2.1.1 and 2.1.2. Analogously to (2.18), the energy minimiser can be obtained via gradient descent.

The data similarity and smoothness constraints in our generalised model (3.2) penalise tonal distances between patches rather than between single pixels as in the original NDS approach (2.3). Comparing (3.8) with the fixed point form of the NDS model (2.25) we note that the *patch distances induce convolutions with the neighbouring tonal weights*. In Sections 3.2 and 3.3 we discuss the implications of this fact and how it inspires the modelling of new filters.

For the sake of completeness, we provide a detailed derivation of the term (3.4), which holds analogously for (3.5):

$$\begin{aligned}
& \frac{\partial E_{GD}}{\partial u_k} \\
&= \frac{\partial}{\partial u_k} \sum_{i,j \in J} \Psi_D (d_{D;i,j}^2) w_D (|x_i - x_j|^2) \\
&= \sum_{i,j \in J} \Psi'_D (d_{D;i,j}^2) \frac{\partial}{\partial u_k} (d_{D;i,j}^2) w_D (|x_i - x_j|^2) \\
&= \sum_{i,j \in J} \Psi'_D (d_{D;i,j}^2) \frac{\partial}{\partial u_k} \left(\sum_p G_\sigma(p) (u_{i+p} - f_{j+p})^2 \right) w_D (|x_i - x_j|^2) \\
&\stackrel{(i=k-p)}{=} 2 \sum_{j \in J} \sum_p \Psi'_D (d_{D;k-p,j}^2) G_\sigma(p) (u_k - f_{j+p}) w_D (|x_{k-p} - x_j|^2) \\
&\stackrel{(j=l-p)}{=} 2 \sum_{l \in J} \sum_p G_\sigma(p) \Psi'_D (d_{D;k-p,l-p}^2) (u_k - f_l) w_D (|x_{k-p} - x_{l-p}|^2) \\
&= 2 \sum_{j \in J} G_\sigma * \Psi'_D (d_{D;k-\cdot,j-\cdot}^2) (0) (u_k - f_j) w_D (|x_k - x_j|^2) .
\end{aligned}$$

3.2 Double Weighting

Considering the data term of eq. (3.8) only (the situation for the smoothness term is analogous), and expanding the convolution (3.6), the fixed point equation for the filtered pixel u_i becomes

$$u_i = \frac{1}{M_{i,j}} \sum_{j,p} G_\sigma(p) \Psi' \left(\sum_q G_\sigma(q) |u_{i+p+q} - f_{j+p+q}|^2 \right) w_{i,j} f_j \quad (3.9)$$

where $M_{i,j}$ is the usual normalisation by the sum of all applied weights:

$$M_{i,j} = \sum_{j,p} G_\sigma(p) \Psi' \left(\sum_q G_\sigma(q) |u_{i+p+q} - f_{j+p+q}|^2 \right) w_{i,j} .$$

In (3.9), G_σ is the Gaussian of radius r_σ which represents the patch size in the patch similarity computation (3.1). Note that this weighting appears twice in formula (3.9): Once during the patch similarity calculation (summed over q) before the nonlinearity Ψ' is applied. We call this G_σ the *inner* weighting of patch pixels. Moreover, G_σ appears also for a second time in (3.9), in the sum over p . We call this the *outer* weighting which is applied when summing the results of the function Ψ' after it is applied to individual patch distances. Figure 3.1 demonstrates this: the tonal weight (3.6) entering in (3.9) not only involves the comparison of the patches about the pixels i and j , but also the patch similarity between their corresponding neighbours is considered.

Equation (3.9), and particularly this double weighting, deserve a detailed discussion. The estimated pixel value u_i in (3.9) is obtained as a weighted average of some data samples f_j . Let us consider a single data pixel f_j , and analyse what is the weight by which this pixel contributes to the weighted result. For a single value of the dummy variable p , the sum

$$\sum_q G_\sigma(q) |u_{i+p+q} - f_{j+p+q}|^2$$

evaluates the weighted L_2 distance between an image patch around pixel u_{i+p} on one hand, and an image patch around pixel f_{j+p} on the other hand (where the size of the patches is given by the weighting function G_σ). In the notation used earlier in this chapter, this patch distance is denoted $d(u(\mathcal{P}_{i+p}), f(\mathcal{P}_{j+p}))$. Note that the compared patches are offset with respect to the estimation and data positions i and j , respectively, by a common shift p .

Coming back to equation (3.9), after evaluating the patch distance, the nonlinearity Ψ' is applied next. We remark that this nonlinearity can be related to robust statistical estimation; its role is to downweight outliers, and convert patch distance to (robust) patch similarity. Then, the resulting patch similarities are summed over variable p in




Figure 3.1: The weight computation in (3.9) between two pixels i and j involves the patch comparison – using patches G_σ (dashed lines) – between every pair of corresponding neighbours $i + p$ and $j + p$ within a neighbourhood G_σ (solid lines).

a second patch neighbourhood, again defined by the weighing function G_σ . Note that the inner and outer weighing functions are identical, which originates in the functional E_{GD} of (3.2) and the derivatives with respect to u_i which duplicated the inner weight also out of the nonlinearity.

Summarising it in words, the equation (3.9) has the following meaning: For pixels u_i and f_j , calculate the patch distances of all patches at positions $i+p$ and $j+p$ taken with the offset p around u_i and f_j , respectively. Then, average these patch distances (transformed first by the nonlinearity Ψ') using the outer weighting G_σ . Thus, the pixel f_j will contribute to the result u_i with a high weight not only if the patches around u_i and f_j are similar, but also if the neighbouring patches u_{i+p} and f_{j+p} resemble each other.

3.3 GNDS Filter Family

In the previous section we discussed the roles of the inner (patch) weighting G_σ and the outer (similarity integration) weighting. Derived from the energy functional, these two weightings are identical. In the fixed point iteration though, these two weighting functions have a different role, and it is instructive to analyse what changes if they are decoupled.

In the following, we keep the parameter r_σ for the radius of the Gaussian G_σ of the inner pixel weighing for patch similarity calculation. The outer integration scale will use a different weighting function G_ρ of radius r_ρ , and the pixel averaging equation becomes

$$u_i = \frac{1}{M_{i,j}} \sum_{j,p} G_\rho(p) \Psi' \left(\sum_q G_\sigma(q) |u_{i+p+q} - f_{j+p+q}|^2 \right) w_{i,j} f_j \quad (3.10)$$

where $M_{i,j}$ is the corresponding normalisation factor.

Let us now study what is the effect of varying the parameters σ and ρ which determine the size of the inner and outer weighting windows, respectively.

First, let $\rho \rightarrow 0$, leading to the following outer weighting:

$$G_0(p) = \begin{cases} 1 & \text{if } p = 0 \\ 0 & \text{if } p \neq 0 \end{cases}. \quad (3.11)$$

The equation (3.10) then simplifies to

$$u_i = \frac{1}{M_{i,j}} \sum_j \Psi' \left(\sum_q G_\sigma(q) |u_{i+q} - f_{j+q}|^2 \right) w_{i,j} f_j, \quad (3.12)$$

which, using $\Psi(s^2) = 2\lambda^2 (1 - \exp(-s^2/(2\lambda^2)))$, corresponds to the non-iterative NL-means filter introduced by Buades *et al.* in [40, 41]. NL-means weights the

contribution of the pixel f_j using a single patch distance comparing patches around u_i and f_j , and omits any additional integration of these patch similarities using the outer summation.

Second, let $\sigma \rightarrow 0$. This leads to

$$u_i = \frac{1}{M_{i,j}} \sum_{j,p} G_\rho(p) \Psi'(|u_{i+p} - f_{j+p}|^2) w_{i,j} f_j. \quad (3.13)$$

Comparing (3.13) with (3.12), we observe that these two equations have a highly similar structure, with a single difference: The position where the nonlinearity Ψ' is applied. For NL-means (3.12), we first sum the differences of individual pixels, thus evaluating the weighted L_2 similarity, and then apply the robust weighting Ψ' . In the other case of equation (3.13), we apply the nonlinearity Ψ' to individual pixel differences, and then integrate the result over the window G_ρ . Even in this case, the weight of pixel f_j is influenced by the whole patches around u_i and f_j . The difference lies in the way the patch similarity is evaluated. Due to the structural resemblance of the filters (3.12) and (3.13) to isotropic and anisotropic penalisation [226] we call (3.13) *anisotropic NL-means*.

As a third example, let both $\sigma \rightarrow 0$ and $\rho \rightarrow 0$. Then, the generalised NDS scheme (3.10) simplifies to the classical NDS scheme (2.11) which is based on simple pixel differences instead of patch distances.

Following the previous analysis, if we allow the inner and outer Gaussian kernels in (3.6)-(3.7) to operate on different integration scales, i.e. σ (inner), ρ (outer),

$$\tilde{g}_{i,j}^{GD} := G_\rho * \Psi'_D \left(|d_D(u(\mathcal{P}_{i-}), f(\mathcal{P}_{j-}))|^2 \right) (0) \quad (3.14)$$

$$\tilde{g}_{i,j}^{GS} := 2 G_\rho * \Psi'_S \left(|d_S(u(\mathcal{P}_{i-}), u(\mathcal{P}_{j-}))|^2 \right) (0), \quad (3.15)$$

the modified fixed point equation (3.8)

$$\tilde{u}_i = \frac{(1 - \alpha) \sum_{j \in J} \tilde{g}_{i,j}^{GD} w_{i,j}^{D,r} f_j + \alpha \sum_{j \in J} \tilde{g}_{i,j}^{GS} w_{i,j}^{S,r} u_j}{(1 - \alpha) \sum_{j \in J} \tilde{g}_{i,j}^{GD} w_{i,j}^{D,r} + \alpha \sum_{j \in J} \tilde{g}_{i,j}^{GS} w_{i,j}^{S,r}} \quad (3.16)$$

can be regarded as a full family of highly nonlinear and robust filters. A single member of this family with $\rho = \sigma$ can be derived from the energy functional (3.2). The well known NL-means method belongs to this family. It represents the case when $\alpha \in \{0, 1\}$ and the outer scale vanishes. Some of these special cases are summarised in Table 3.1.

Table 3.1: Examples of filtering methods belonging to the GNDS family (3.16) with varying regularisation parameter α , patch size σ and integration scale ρ .

Regularisation parameter	Patch size	Integration scale	Method
$0 \leq \alpha \leq 1$	$\sigma > 0$	$\rho = \sigma$	Generalised NDS (3.10)
$\alpha \in \{0, 1\}$	$\sigma > 0$	$\rho \rightarrow 0$	NL-means (3.12), [41]
$\alpha \in \{0, 1\}$	$\sigma \rightarrow 0$	$\rho > 0$	Anisotropic NL-means (3.13)
$0 \leq \alpha \leq 1$	$\sigma \rightarrow 0$	$\rho \rightarrow 0$	Classical NDS (2.3), [148]

Practically, the inner and outer scales both act in the direction that by increasing them, we increase the area used to evaluate image similarity: Higher values lead to a more thorough (and costly) patch comparison. Consider the special situation when this combined scale is kept constant, the amount of integration just shifts between the inner and outer scales. Such a setting forms a family of filters with approximately the same spatial extent of operations. What changes is the position at which the nonlinearity Ψ' enters the chain. The NL-means (3.12) and the summation of robust pixel similarities (3.13) represent the two extremes of this family.

Considering an image of N pixels, a squared search window w of s^2 pixels, and circular patch of radius r , the computational complexity of the filter family (3.16) is $O(N \times s^2 \times r_\rho^2 \times r_\sigma^2)$. Fig. 3.2 illustrates the effect of varying r_ρ and r_σ while keeping $r_\rho + r_\sigma$ constant. The different configurations correspond to different ways of computing pixel similarity.




Figure 3.2: The tonal weight in (3.10) is computed over an area determined by the integration neighbourhood (solid lines) – Gaussian G_ρ of radius r_ρ – and the patch size (dashed lines) – Gaussian G_σ of radius r_σ . From left to right, different configurations where the total area described by a Gaussian of radius $r_\rho + r_\sigma$ (dotted line) is kept constant. The first and the last configurations correspond to the weighting scheme of the filters (3.12) and (3.13), respectively.

3.4 Alternative Formulations of the NL-Means Filter

Using the Whittaker-Tikhonov penaliser $\Psi(s^2) = s^2$ we obtain $\Psi'(s^2) := \partial_{s^2}\Psi(s^2) = 1$ and both filters (3.12) and (3.13) become equivalent to

$$u_i = \frac{1}{M_{i,j}} \sum_j w_{i,j} f_j. \quad (3.17)$$

In our setting the spatial function w acts uniquely as a search window, i.e. it delimits the spatial extent where the pixels j , neighbours of i , are taken from. However, in various works [87, 29, 61] it is argued that (3.17) can be regarded as the NL-means filter by redefining the weights via

$$\tilde{w}_{i,j} := w_{i,j} \cdot \exp\left(-\frac{|d(f(\mathcal{P}_i), f(\mathcal{P}_j))|^2}{h}\right), \quad (3.18)$$

with $h > 0$ as a filter parameter. Note that the additional weighting term is constant as it depends on the input image f . This indicates that the filter (3.17) could be directly derived from the data term (2.1) of the NDS functional employing $\Psi(s^2) = s^2$ and \tilde{w} . Analogously, a filter that averages over the evolving image u can be obtained from the smoothness term (2.2). Similar ideas have been considered in [86, 11, 38].

In [114, 64, 169] energy functionals with weights depending on the unknown solution u via $d(u(\mathcal{P}_i), u(\mathcal{P}_j))$ have been considered. However, all these methods assume constant weights in the computation of the optimality conditions $\nabla E(u) = 0$. The variational filter proposed by Brox *et al.* in [38] also considers nonlocal weights depending on u . Although they do not assume constant weights in the derivation of the Euler-Lagrange equations, these become very complex and computationally expensive. To a certain extent, the mentioned filters could be obtained from the original NDS framework (2.3) by extending the definition of the weights w as in (3.18) (see also a related discussion in Section 2.3). In the proposed GNDS framework we keep regarding the weights w only as (nonlocal) spatial functions. Instead, we generalise the constraints being penalised in the energy functional. That is, in (3.2) we have replaced the *single pixel* similarity constraints of (2.3) by *patch similarity* constraints using the weighted L_2 distance between patches, obtaining a new family of neighbourhood filters. The use of other similarity measures is discussed in Section 3.7. It is important to mention that we do take into account the dependency of the distance measures on the solution u when deriving the optimality conditions. As a result, the classical and also some iterative versions of the NL-means filter can be obtained as special cases of the proposed filter family (3.16) without need of redefining the spatial weights w .

3.5 (Non-) Iterative and Steady-State Solutions

In Section 3.3 we explored the full family of filters that can be obtained from the proposed GNDS model by varying the inner and outer scales in the patch similarity computation. This entails the immediate extension of the filters presented in Section 2.2 to work with image patches rather than with single pixel differences.

Let us consider the fixed point (3.8) that iteratively minimises the energy functional (3.2). For $0 \leq \alpha < 1$ this process will converge to a stationary state due to the data term dependency on the input image. Note that for $\alpha = 0$ we obtain a *generalised nonlocal M-smoothing* process. In this case we can think, for instance, of a novel NL-means filter with a steady-state solution. For $\alpha = 1$ we obtain a *generalised nonlocal Bilateral filter*, which needs to be stopped after certain number of iterations before the image gets completely smoothed away. This can be done by using the decorrelation criterion devised by Mrázek and Navara [147].

3.6 Extension to Multichannel Images

The extension of the GNDS model to multichannel images is straightforward. Let $\mathbf{f}, \mathbf{u} : \Omega \rightarrow \mathbb{R}^d$ be the noisy image and the unknown noise-free image, respectively, both with d channels. To obtain the multichannel counterpart of the scalar GNDS model (3.2) we just need to redefine the patch distance (3.1) as

$$|d(\mathbf{u}(\mathcal{P}_i), \mathbf{f}(\mathcal{P}_j))|^2 = \sum_p G_\sigma(p) \|\mathbf{u}_{i+p} - \mathbf{f}_{j+p}\|_2^2, \quad (3.19)$$

where $\|\cdot\|_2$ is the Euclidean norm. Computing the optimality conditions $\nabla E(\mathbf{u}) = 0$ we obtain a fixed point for every channel u^m ($m = 1, \dots, d$), cf. (3.8):

$$u_i^m = \frac{(1 - \alpha) \sum_{j \in J} g_{i,j}^{GD} w_{i,j}^{D,\varpi} f_j^m + \alpha \sum_{j \in J} g_{i,j}^{GS} w_{i,j}^{S,\varpi} u_j^m}{(1 - \alpha) \sum_{j \in J} g_{i,j}^{GD} w_{i,j}^{D,\varpi} + \alpha \sum_{j \in J} g_{i,j}^{GS} w_{i,j}^{S,\varpi}}. \quad (3.20)$$

All channels are coupled via the tonal weights

$$g_{i,j}^{GD} := G_\sigma * \Psi'_D \left(|d_D(\mathbf{u}(\mathcal{P}_{i-}), \mathbf{f}(\mathcal{P}_{j-}))|^2 \right) (0) \quad (3.21)$$

$$g_{i,j}^{GS} := 2 G_\sigma * \Psi'_S \left(|d_S(\mathbf{u}(\mathcal{P}_{i-}), \mathbf{u}(\mathcal{P}_{j-}))|^2 \right) (0), \quad (3.22)$$

which avoid the formation of discontinuities at different locations for the different image channels. Note that the fixed point (3.20) can be modified as in (3.16) to obtain a more flexible and robust filter.

3.7 Extension to Other Distance Measures

The proposed energy functional (3.2) is very general in the sense that one could choose any suitable distance measures d_D , d_S to impose similarity of particular image characteristics. Once the distances have been chosen, the optimality conditions $\nabla E(u) = 0$ need to be derived in order to prescribe the corresponding energy minimiser, for instance, via a fixed point or a gradient descent scheme.

In the proposed GNDS model we have used the weighted L_2 norm (3.1) to measure similarity between image patches. However, one can employ different distance measures as well. For example, Kervrann and Boulanger [114] use

$$|d(\mathbf{u}(\mathcal{P}_i), \mathbf{u}(\mathcal{P}_j))|^2 = \text{vec}(\mathbf{u}(\mathcal{P}_i) - \mathbf{u}(\mathcal{P}_j))^{\top} V_{ij}^{-1} \text{vec}(\mathbf{u}(\mathcal{P}_i) - \mathbf{u}(\mathcal{P}_j)), \quad (3.23)$$

where V_{ij} is a diagonal matrix whose entries are averaged local variances of the image patches. Similarly, Goossens *et al.* [90] replace V_{ij} by a local estimation of the noise covariance matrix to filter images corrupted by correlated noise.

Another example where the selection of the patch distance is driven by the noise type corrupting the image data is due to Coupé *et al.* [68]. Based on the Bayesian non-local means filter [115] and on the Speckle noise model introduced in [132], the authors propose a non-local filter for ultrasound images that uses the so-called Pearson distance for computing patch similarity:

$$|d(u(\mathcal{P}_i), u(\mathcal{P}_j))|^2 = \sum_p \frac{(u_{i+p} - u_{j+p})^2}{u_{j+p}}. \quad (3.24)$$

These and other measures of similarity can be utilised in the proposed functional (3.2) with accordingly derived minimisation algorithms. In addition, as classically done in variational methods, the choice of the data similarity constraint can be driven by the statistical properties of the type of noise present, whereas the smoothness term must reflect desirable properties of the solution. Therefore, the issue of selecting appropriate patch distances for the data and smoothness terms of the proposed GNDS model is still open. We will explore these issues in a future work.

3.8 Numerical Experiments

3.8.1 Comparison of similarity measures

The filter (3.10) induces a novel similarity measure between two pixels u_i, u_j that can be considered as an *extended* patch similarity measure

$$S_{ext}(u_i, u_j) := \sum_p G_\rho(p) \cdot \Psi' \left(\sum_q G_\sigma(q) \cdot |u_{i+p+q} - u_{j+p+q}|^2 \right). \quad (3.25)$$

Choosing $\rho \rightarrow 0$ one obtains an *isotropic* similarity measure

$$S_{iso}(u_i, u_j) := \Psi' \left(\sum_q G_\sigma(q) \cdot |u_{i+q} - u_{j+q}|^2 \right), \quad (3.26)$$

while with $\sigma \rightarrow 0$, equation (3.25) becomes an *anisotropic* similarity measure

$$S_{ani}(u_i, u_j) := \sum_p G_\rho(p) \cdot \Psi' (|u_{i+p} - u_{j+p}|^2). \quad (3.27)$$

Considering the penaliser of Leclerc [124] and Perona/Malik [167]

$$\Psi(s^2) = 2\lambda^2 \left(1 - \exp \left(-\frac{s^2}{2\lambda^2} \right) \right) \quad (3.28)$$

with filter parameter λ , (3.26) corresponds exactly to the similarity measure used by Buades *et al.* [41] in their NL-means filter. A couple of recent works have proposed the use of other robust penalisers as well [90, 168]. We test these three measures on the noisy images displayed in Fig. 3.3. For each one of the 16 textures we select 30 random pixels and compute their similarity to all other pixels in the image. For every chosen pixel we take its best 20 matches (pixels with the largest similarity) and check whether they belong to the same texture or not. Table 3.2 shows the average number of matches within the same texture and the overall performance of each similarity measure. In the case of Gaussian noise we used the Leclerc penaliser and for salt-and-pepper noise the regularised L_1 norm $\Psi(s^2) = \sqrt{s^2 + \epsilon^2}$. The radii of the Gaussians were set to $r_\rho = r_\sigma = 4$.

The results show that the extended similarity measure is more robust and perform best under Gaussian degradation. This is due to the fact that, via the outer Gaussian weighting, the selection of similar pixels relies more strongly on the underlying image structures. On the other hand, it performs poorly under impulse noise. In this case




Figure 3.3: *Left:* Original image with 16 textures, each one identified by its coordinates in matrix notation $T(x, y)$, $x, y = \{1, 2, 3, 4\}$. *Middle left:* Zoom of the original image. *Middle right:* Added zero-mean Gaussian noise with $\sigma = 20$. *Right:* Added 20% of salt-and-pepper noise.

the best choice is the anisotropic similarity measure, which acts as a noise detector at every pixel location. The same holds for higher levels of noise. Smoothing experiments will be presented in the following sections.

Table 3.2: Quantitative comparison of the three similarity measures S_{iso} , S_{ani} and S_{ext} induced from the GNDS filter (3.10). S_{ext} is more suitable for images degraded with Gaussian noise, while S_{ani} is more robust under salt-and-pepper noise.

	Gaussian noise ($\sigma = 20$)			salt-and-pepper noise (20%)		
	S_{iso}	S_{ani}	S_{ext}	S_{iso}	S_{ani}	S_{ext}
T(1,1)	3	4	13	1	3	2
T(1,2)	19	19	19	19	19	19
T(1,3)	20	16	20	10	19	13
T(1,4)	10	8	12	11	10	8
T(2,1)	20	18	20	20	20	18
T(2,2)	14	5	19	9	18	12
T(2,3)	20	13	20	1	11	2
T(2,4)	16	6	15	4	11	6
T(3,1)	12	5	20	4	16	7
T(3,2)	18	17	19	11	20	14
T(3,3)	14	7	20	3	12	3
T(3,4)	12	13	18	1	13	3
T(4,1)	10	8	14	8	10	8
T(4,2)	20	14	20	4	20	5
T(4,3)	9	7	11	1	5	1
T(4,4)	17	10	16	1	18	6
matches	234	170	276	108	225	127
%	73.1	53.1	86.2	33.7	70.3	39.7

3.8.2 Comparison of several patch-based methods

We now evaluate the smoothing capabilities of the proposed GNDS model on the set of test images *Barbara*, *House*, *Lena*, *Peppers*, *Boats* from Portilla *et al.* [175] which already contain Gaussian noise. The proposed GNDS filter is run iteratively via a gradient descent scheme. In all experiments we use the penaliser (3.28) with fixed contrast parameter λ for successive iterations of the filter, a search window of size

21×21 and patches of radius $r_\sigma = 5$ (implemented as squares of $(2r_\sigma - 1)^2$ pixels). The radius r_ρ of the outer patch weighting was chosen between 0 and 2 pixels. With this configuration, one iteration of (3.16) on a 256×256 image took between 17 and 95 seconds on a Pentium IV 2.8GHz implemented in C. Table 3.3 juxtaposes several patch-based filters proposed in the literature. We employ the peak signal-to-noise ratio (PSNR) as criterion for quality measure:

$$\text{PSNR (dB)} = 10 \log_{10} \left(\frac{255^2}{\frac{1}{|J|} \sum_{i \in J} (o_i - u_i)^2} \right), \quad (3.29)$$

where o denotes the original noise free image and u the estimated denoised version. The shown results for [41, 10, 86, 87] were taken from [38]. From those most competitive methods related to the proposed GNDS filter: Brox *et al.* [37, 38] run an iterative NL-means algorithm that uses the noisy image for averaging and updates the weights from the estimated solution u of the previous iteration. A similar strategy is due to Kervrann *et al.* [113, 114] who additionally adapt the size of the averaging neighbourhood at each pixel location to better capture local geometries. Azzabou *et al.* [11] developed a variational filter structurally similar to [88] and [38] that adapts the spatial extent of the local neighbourhoods. We also compare with the nonlocal TV filter as in [30] (see [88] as well). Although the proposed GNDS filter does not utilise sophisticated adaptive strategies, it also allows for a robust selection of similar pixels by making use of the extended patch similarity measure defined in (3.25). Note that in some cases the GNDS filter outperforms the more elaborated methods, though it is still below the state-of-the-art results provided by Dabov *et al.* [70].

It is worth mentioning that the nonlocal smoothness term of GNDS model (3.2) reaches higher PSNRs than the nonlocal data term, which is more pronounced for higher levels of noise. Interestingly, the combined use of both terms leads to slightly better results than the smoothness term alone. We also run experiments considering models such as (i) local data terms $\sum_{i \in J} \Psi(|u_i - f_i|^2)$ with a nonlocal smoothness term, and (ii) a nonlocal data term combined with semilocal smoothness terms $\sum_{i \in J, j \in \mathcal{N}(i)} \Psi(|u_i - u_j|^2)$, where the set $\mathcal{N}(i)$ contains the 4 direct neighbours of pixel i . However, both models led to poorer results. This is in concordance with the findings in [86, 87], where the proposed variational filters perform better when a nonlocal regulariser is used and the data fidelity term is disregarded.

As was mentioned above, the GNDS filter was implemented using a steepest descent algorithm. We run the iterative scheme for different time-step size $\tau = 0.1, 0.2, \dots, 1.0$. Fig. 3.4 shows the performance of the GNDS-D filter applied to the noisy test image *House* as a function of the time step τ and the filter parameter λ . Similar curves are obtained with the GNDS-S filter. As noted from Table 3.3 the best denoising results are attained with τ in the range $[0.8, 1.0]$, in which case the number of iterations needed to reach the highest PSNR ranges between 1 and 3.

Fig. 3.5 shows a visual comparison of the proposed GNDS filter with the two most competitive methods [70, 114]. The *absolute method noise* (AMN) $|o - u| (\times 5)$

Table 3.3: Denoising results of several patch-based filters on standard test images degraded with additive zero-mean Gaussian noise with standard deviation $\{20, 50\}$. The method of Dabov *et al.* [70] reaches the highest PSNRs. The second best results are highlighted. The parameters (λ, τ , iterations) are displayed for the GNDS-D and GNDS-S filters (data and smoothness terms of (3.16), respectively). GNDS, i.e. the combination of both terms, yields better results than applying each of them independently, outperforming most of the listed methods.

	Gaussian noise ($\sigma = 20$)				
Filter \ PSNR (dB)	<i>Barbara</i>	<i>House</i>	<i>Lena</i>	<i>Peppers</i>	<i>Boats</i>
	22.18	22.11	22.13	22.19	22.17
Buades <i>et al.</i> [41]	30.31	32.49	31.78	29.62	29.34
Awate <i>et al.</i> [10]	30.14	32.59	31.79	29.75	29.54
Gilboa <i>et al.</i> [86]	29.43	32.17	31.39	30.04	29.53
Gilboa <i>et al.</i> [87]	30.20	32.55	31.95	30.28	29.89
Dabov <i>et al.</i> [70]	31.78	33.77	33.05	31.29	30.88
Azzabou <i>et al.</i> [11]	30.46	32.34	32.12	30.67	29.94
Brox <i>et al.</i> [38]	30.33	32.74	32.08	30.04	29.69
Kervrann <i>et al.</i> [114]	30.37	32.90	32.64	30.59	30.12
Bougleux <i>et al.</i> [30]	30.41	32.72	31.95	30.17	29.57
GNDS-D	30.62 (14,1.0,1)	32.66 (15,1.0,1)	31.98 (14,1.0,1)	30.21 (13,1.0,1)	29.78 (12,1.0,1)
GNDS-S	30.62 (14,1.0,1)	32.75 (11,0.8,2)	32.03 (10,0.8,2)	30.21 (13,1.0,1)	29.78 (12,1.0,1)
GNDS	30.64	32.78	32.05	30.22	29.80
	Gaussian noise ($\sigma = 50$)				
Filter \ PSNR (dB)	<i>Barbara</i>	<i>House</i>	<i>Lena</i>	<i>Peppers</i>	<i>Boats</i>
	14.76	14.56	14.62	14.68	14.59
Dabov <i>et al.</i> [70]	27.17	29.37	28.86	26.41	26.64
Kervrann <i>et al.</i> [114]	24.09	28.67	28.38	25.29	25.93
GNDS-D	25.40 (23,1.0,2)	27.66 (24,1.0,2)	27.30 (24,1.0,2)	25.25 (24,1.0,2)	25.16 (23,1.0,2)
GNDS-S	25.75 (20,0.9,2)	28.38 (19,1.0,2)	27.77 (19,1.0,2)	25.64 (19,1.0,2)	25.58 (20,0.9,2)
GNDS	25.78	28.40	27.81	25.67	25.60




Figure 3.4: Performance of the GNDS-D filter (see Table 3.3) on the test image *House* degraded by Gaussian noise with standard deviation 20 (*Left*) and 50 (*Right*). The PSNR curves are displayed as functions of the time step size τ and the filter parameter λ . The plots show that there exists an optimally global λ^* for a chosen time step size. Considering all the experiments carried out, the best results were obtained with τ in the range $[0.8, 1.0]$.

between the noise free images o and the restored versions u are shown in Fig. 3.6. All three approaches provide very good results, while the method of Dabov *et al.* [70] gives the highest PSNRs. Although our GNDS approach does not outperform these two methods in terms of PSNR, our results look much more pleasant and natural than those from Kervrann *et al.* [114]. That filter tends to over-enhance edges, creating staircasing artifacts that make the images look less natural, which can be observed in Fig. 3.7. Another visible effect of the method of Kervrann *et al.* is noticeable in the AMN images of Fig. 3.6. The black areas reveal that many edges remain untouched in the filtering process, i.e. no noise is removed at those locations. Our GNDS results do not show any visible artifacts and almost no loss of structures is perceived in the method noise images. These findings suggest that PSNR is not a fully reliable measure for denoising capability and perceptual quality altogether. Alternative ways of assessing these criteria are necessary, but this goes beyond the scope of our work.

Fig. 3.8 demonstrates the application of our GNDS filter to denoising colour images. The noisy *Boy* images were created adding zero-mean Gaussian noise in every $\{R,G,B\}$ channel independently. As it was indicated in Section 3.6 we apply the filter (3.20) on every image channel using the so-called *channel coupling* technique in order to avoid the formation of false colours and the dislocation of edges. That is, the same tonal weights (3.21)–(3.22) are used in all channels. The accurate localisation and restoration of edges can be observed in the zoomed images of Fig. 3.9. This is especially visible in the transition between the *Boy*'s cheek and the red collar. Our filter is able to restore gentle facial features and to preserve small details such as the pullover's zip.



Figure 3.5: Comparison to state-of-the-art methods. *Top Row*: Test images degraded by Gaussian noise with standard deviation 20. *2nd Row*: Restored images by Dabov *et al.* [70]. *3rd Row*: Restored images by Kervrann *et al.* [114]. *Bottom Row*: Restored images by the proposed GNDS model (3.2).




(a) AMN for the restored *Barbara*, *Boats* and *Lena* by Dabov *et al.* [70].(b) AMN for the restored *Barbara*, *Boats* and *Lena* by Kervrann *et al.* [114].(c) AMN for the restored *Barbara*, *Boats* and *Lena* by the proposed GNDS model (3.2).

Figure 3.6: *Absolute method noise* (AMN) $|o - u|$ ($\times 5$) for the smoothing results shown in Fig. 3.5 obtained by (a) Dabov *et al.* [70], (b) Kervrann *et al.* [114], and (c) the proposed GNDS filter.

We finally test our approach for restoring impulse noise. Fig. 3.10 compares several filters: (c) iterative median filtering within a 3×3 window; (d) NDS model (2.46) with parameters $(\alpha; \varpi_D; \varpi_S) = (0.05; 2; 1)$; (e) GNDS model (3.16) with penalisers $\Psi_D(s^2) = |s|$, $\Psi_S(s^2) = \lambda^2 \ln(1 + s^2/\lambda^2)$ and parameters $(\alpha; \varpi_D, r_{D\rho}, r_{D\sigma}; \varpi_S, r_{S\rho}, r_{S\sigma}) = (0.05; 2, 0, 0; 2, 0, 3)$; and (f) the method of Chan



(a) Noisy test images degraded by Gaussian noise with standard deviation 20.



(b) Restored images by Dabov *et al.* [70].



(c) Restored images by Kervrann *et al.* [114].



(d) Restored images by the proposed GNDS model (3.2).

Figure 3.7: Zoomed restored images from Fig. 3.5. All methods provide pleasant visual results, although the method by Kervrann *et al.* [114] also shows some staircasing artifacts that make the images look less natural.




Figure 3.8: GNDS filtering on colour images. *Top Row:* From left to right: Original noise free image and noisy versions degraded by Gaussian noise with standard deviation $\{20, 40\}$. *Bottom Row:* From left to right: Slightly smoothed original image and the corresponding restored images by the proposed GNDS filter.

et al. [56] with parameters as described in their paper. Note that we have used the anisotropic similarity measure (3.27) in the smoothness term of the GNDS model. The other similarity measures (3.25)–(3.26) do not provide better results. Our approach slightly improves the results obtained by the NDS model. It is worth mentioning that the method of Chan *et al.* initially detects the noisy pixels (salt or pepper) which are subsequently restored, while the other pixels remain unchanged. Although our approach does not recourse to a noise detector as a pre-processing step, it provides reasonable results for high levels of noise. Other interesting approaches dealing with simultaneous restoration and deblurring can be found in [13, 53, 128, 111].




Figure 3.9: Zoomed images from Fig. 3.8. Gentle facial features are well restored, and the edges are well localised thanks to the channel coupling.




Figure 3.10: Filtering impulse noise. (a) Original image. (b) Degraded with 50% salt-and-pepper noise. (c) Restored by iterative median filtering. (d) Restored by the NDS model. (e) Restored by our GNDS model. (f) Restored by Chan *et al.*'s method [56].

3.9 Summary

We have introduced a general nonlocal discrete variational framework for image smoothing. It arises as a generalisation of the *Nonlocal Data and Smoothness* (NDS) filtering approach of Mrázek *et al.* [148]. Although the NDS model allows nonlocal interactions between pixels, these are effective only semi-locally. This is caused by that fact that its model constraints just penalise single pixel differences that cannot propagate reliable information about the local geometry too far away from a chosen pixel. Therefore, we propose the *Generalised NDS* (GNDS) model with data and smoothness terms penalising general dissimilarity measures defined on image patches. They allow us to incorporate structured pixel information from truly nonlocal neighbourhoods in the smoothing process. We showed that by using the weighted L_2 norm as distance measure the energy minimiser results in a robust and versatile neighbourhood filter that can be adjusted to restore vector-valued images corrupted by Gaussian and salt-and-pepper noise. With respect to restoration quality our GNDS approach can outperform other related patch-based methods and compares fairly well to more advanced approaches [70, 114].

Our discrete variational framework includes as special cases patch-based generalisations of M-smoothers and bilateral filtering. We showed that a slight modification of the fixed-point solution leads to a more general family of nonlocal nonlinear filters, from which the NL-means filter of Buades *et al.* [41] and some of its iterative variants can be obtained. The proposed smoothing framework is closely related to the methods of Kervrann *et al.* [114], Azzabou *et al.* [11], Brox *et al.* [38] as well as to the approaches of Gilboa *et al.* [86, 87] and Bogleux *et al.* [30] inspired from graph regularisation techniques. Some of these filters can be derived from our energy model by employing a different similarity measure and/or by redefining the spatial weight functions that we use as search windows.

In this work we have mainly exploited the use of the weighted L_2 norm to compute patch distances. However, there is a rich opportunity for future work concerning alternative similarity measures better suited for different types of noise contamination as well as for other applications such as deblurring, inpainting, super-resolution and segmentation.

4

Matrix-Valued NDS Framework

In Chapter 2 we showed that the NDS model, thanks to its structural flexibility, outperforms all other methods that can be obtained as special cases from it. Later, in Chapter 3 we obtained a more nonlocal general framework, the GNDS filter, by combining the NDS model with the idea of penalising different dissimilarity measures. We are now interested in extending these ideas to the formulation of a nonlocal variational approach for matrix fields. In this chapter, we propose two different matrix-valued counterparts of the NDS model: The *isotropic* matrix-valued NDS (iMNDS) approach and the *anisotropic* matrix-valued NDS (aMNDS) approach. The distinction between the isotropic and the anisotropic case is in the sense of [227, 223]. As an inherited property from the original NDS model, our approaches generalise several methods for filtering, regularisation and interpolation of matrix fields proposed in the literature, which can be seen as special cases. We demonstrate the reconstructing capabilities and the differences between our two approaches on several experiments with synthetic and real-world data sets such as those displayed in Figure 4.1. This chapter is based on [171, 172].

There exist multiple sources of matrix and tensors fields in engineering and the physical sciences [139]. In this chapter we work with matrix-valued images obtained by diffusion-tensor magnetic resonance imaging (DT-MRI). This modern medical imaging modality associates a real symmetric positive-definite 3×3 -matrix to each voxel of the volume under consideration. These matrices, visualised by ellipsoids, indicate the diffusive behaviour of water molecules under thermal Brownian motion, and as such reflect the structure of the surrounding tissue. However, symmetric but possibly indefinite matrix-fields also appear, for example in physics and engineering as general descriptors of anisotropic behaviour. In any case, the data are often corrupted by noise and therefore developing filtering and simplification techniques for matrix fields is compelling. In the literature, smoothing processes for positive definite matrix-fields, namely DT-MRI data, based on diffusion and regularisation concepts have been proposed in [208, 223], while based on differential geometric considerations in [63, 18, 144, 145, 166, 6, 82, 96]. Other approaches relying on the operator-algebraic view of symmetric matrices (cf. Section 1.4) have been proposed for filtering and regularisation of matrix fields, positive definite or not, in [48, 47, 197, 193].




Figure 4.1: Examples of matrix fields of 3×3 symmetric positive semidefinite matrices. *Left*: Synthetic data with 32×32 voxels showing a blend between the original field and its noisy version with random positive definite matrices. *Right*: 2D slice extracted from a 3D DT-MRI data set, 128×128 voxels.

4.1 Isotropic Matrix-Valued NDS Model

Matrix-valued images are mappings $U : \Omega \subset \mathbb{R}^d \rightarrow \text{Sym}_m(\mathbb{R})$, where $\text{Sym}_m(\mathbb{R})$ is the set of all symmetric real-valued $m \times m$ matrices. Here, we typically set $d = 2, 3$ and $m = 3$. We denote by $U = (U_i)_{i \in J}$ a processed version an original (noisy) field $F = (F_i)_{i \in J}$, where $J = \{1, \dots, N\}$ and N is the total number of voxels.

Recall from the NDS model 2.3 in Chapter 2 that the penalisers Ψ in the similarity and in the smoothness term act on scalar differences of the form $|u_i - f_j|$ and $|u_i - u_j|$, respectively. These quantities were termed *tonal distances* as they measure the distance between the intensity values at the pixel locations x_i and x_j . Transferring this concept to the matrix-valued setting amounts to computing distances between real symmetric matrices A, B :

$$\begin{aligned} d : \text{Sym}_m(\mathbb{R}) \times \text{Sym}_m(\mathbb{R}) &\longrightarrow \mathbb{R}_0^+ \\ (A, B) &\mapsto d(A, B). \end{aligned} \quad (4.1)$$

This allows us to redefine the penalisers Ψ as a composition of two functions:

$$\begin{aligned} \Psi \circ d : \text{Sym}_m(\mathbb{R}) \times \text{Sym}_m(\mathbb{R}) &\longrightarrow \mathbb{R}_0^+ \\ (A, B) &\mapsto \Psi(d(A, B)^2), \end{aligned} \quad (4.2)$$

where Ψ can be any of the tonal functions listed in Table 2.1, Chapter 2.

With these definitions, we propose the following *isotropic* matrix-valued NDS (iMNDS) model for nonlocal regularisation of matrix-valued images

$$E_I(U) = \beta \sum_{i,j \in J} \Psi_D(d(U_i, F_j)^2) w_D(|x_i - x_j|^2) + (1 - \beta) \sum_{i,j \in J} \Psi_S(d(U_i, U_j)^2) w_S(|x_i - x_j|^2), \quad (4.3)$$

where we have replaced $\beta := (1 - \alpha)$ to give a more general interpretation of this weight. In fact, β could be taken inside the sum to make it a spatially-variant function. For classical regularisation with a local similarity term [223], β is constant everywhere. For interpolation, a function $\beta(x) : \Omega \rightarrow [0, 1]$ allows to fill in missing data at voxels x_j where $\beta(x_j) = 0$, while reproducing F_j where $\beta(x_j) = 1$. At locations where $0 < \beta(x_j) < 1$ one obtains regularised estimates. The function β can be seen as an indicator of the *local confidence* in the data [228]. For example, if there exists a noise model from which local noise statistics can be estimated [163, 97], β can be chosen inversely proportional to the local noise variance of the input tensor field.

The minimiser of (4.3) must satisfy the conditions $\frac{\partial}{\partial U_i} E_I(U) = 0$ for all $i \in J$. For a tensor U_i at voxel x_i this condition can be written in a fixed-point form as

$$U_i = \frac{1}{c_i} H^{-1} \left(\beta \sum_{j \in J} \Psi'_D(d(U_i, F_j)^2) w_D(x_{i,j}^2) H(F_j) + 2(1 - \beta) \sum_{j \in J} \Psi'_S(d(U_i, U_j)^2) w_S(x_{i,j}^2) H(U_j) \right), \quad (4.4)$$

where c_i is the normalisation constant given by

$$c_i = \beta \sum_{j \in J} \Psi'_D(d(U_i, F_j)^2) w_D(x_{i,j}^2) + 2(1 - \beta) \sum_{j \in J} \Psi'_S(d(U_i, U_j)^2) w_S(x_{i,j}^2). \quad (4.5)$$

Again, $\Psi'(s^2)$ denotes the derivative of $\Psi(s^2)$ w.r.t. its argument. The function $H : \text{Sym}_m(\mathbb{R}) \rightarrow \text{Sym}_m(\mathbb{R})$ is associated with a specific distance measure d . Please refer to Section 1.4 for the definition of *functions of matrices*. For the *Euclidean* and *log-Euclidean* distance measures shown in Table 4.1, the function H and its inverse can be derived straightforwardly. In the case of the *affine-invariant* distance, these functions need to be coupled with the pivot point U_i in a *geodesic marching* style as motivated in [166].

Table 4.1: Distance measures: Euclidean (E), log-Euclidean (LE), and affine-invariant (AI), with their corresponding mapping H and its inverse H^{-1} . $\|C\| := \sqrt{\text{trace}(C^\top C)}$ denotes the Frobenius norm of the matrix $C \in \text{Sym}_m(\mathbb{R})$.

	$d(U_i, F_j)$	$H(F_j)$	$H^{-1}(F_j)$
E	$\ U_i - F_j\ $	F_j	F_j
LE	$\ \ln U_i - \ln F_j\ $	$\ln F_j$	$\exp F_j$
AI	$\left\ \ln \left(U_i^{-\frac{1}{2}} F_j U_i^{-\frac{1}{2}} \right) \right\ $	$U_i^{\frac{1}{2}} \ln \left(U_i^{-\frac{1}{2}} F_j U_i^{-\frac{1}{2}} \right) U_i^{\frac{1}{2}}$	$U_i^{\frac{1}{2}} \exp \left(U_i^{-\frac{1}{2}} F_j U_i^{-\frac{1}{2}} \right) U_i^{\frac{1}{2}}$

By writing equation (4.4) in the form $U^{k+1} := \Upsilon(U^k)$ with $\Upsilon : \mathbb{R}^N \rightarrow \mathbb{R}^N$, we can build a matrix fixed-point iteration scheme. A restored version $U = (U_i)$ of a noisy tensor field $F = (F_i)$ can be thus obtained via

$$\begin{aligned} U_i^0 &:= F_i, \\ U_i^{k+1} &:= \Upsilon^i(U^k), \quad \text{for all } i \in J, k \in \mathbb{N}_0. \end{aligned} \tag{4.6}$$

An important property of this fixed-point scheme (4.6) is that it preserves positive semidefiniteness of the input field F . This can be easily seen in equation (4.4) where the filtered tensor U_i is obtained as a convex combination of positive semidefinite tensors with $\beta, \Psi', w > 0$.

In the context of classical variational regularisation, Weickert and Welk [228] proposed a continuous scalar-valued approach that is structurally similar to our model (4.3) without the nonlocal spatial weights. They obtain a PDE for interpolation and regularisation of scalar-valued images. In its extension to matrix fields, the PDE is applied to each matrix channel separately, but all channels are coupled via the *joint diffusion tensor* to ensure preservation of discontinuities. In our framework we directly work on the space of symmetric matrices. Our scheme (4.4) is truly matrix-valued and automatically ensures proper interaction of the matrix channels without additional channel coupling techniques. Other approaches related to ours will be discussed in Section 4.3.

To explain why we regard our method as *isotropic*, let us consider the matrices $A, B \in \text{Sym}_m(\mathbb{R})$ and the squared difference matrix $C^2 = (A - B)^2$, which results in a positive (semi-)definite matrix. The spectral decomposition of C^2 reveals directional and shape information of the local structure between the matrices A and B . Nevertheless, such information is reduced to a scalar quantity when a distance measure, such as those shown in Table 4.1, is applied to the difference matrix. In that sense, our model (4.3) can be regarded as isotropic.

4.2 Anisotropic Matrix-Valued NDS Model

We now present an alternative approach where we refrain from using distance measures and directly penalise squared difference matrices in the energy functional, keeping thus the structural local information. We need to redefine the penalisers as

$$\begin{aligned} \tilde{\Psi} &: \text{Sym}_m(\mathbb{R}) \times \text{Sym}_m(\mathbb{R}) \longrightarrow \mathbb{R} \\ (A, B) &\longmapsto \text{trace}\left(\Psi((A - B)^2)\right), \end{aligned} \quad (4.7)$$

where Ψ can be any of the tonal functions listed in Table 2.1, Chapter 2. Note that we have used the linear trace operator as the simplest way of getting a scalar energy that can be properly minimised, which is standard in the literature [227]. Then, the proposed *anisotropic* matrix-valued NDS (aMNDS) model for matrix fields reads

$$\begin{aligned} E_A(U) &= \beta \sum_{i,j \in J} \tilde{\Psi}_D(U_i, F_j) w_D(|x_i - x_j|^2) \\ &\quad + (1 - \beta) \sum_{i,j \in J} \tilde{\Psi}_S(U_i, U_j) w_S(|x_i - x_j|^2). \end{aligned} \quad (4.8)$$

The minimiser of (4.8) must satisfy the conditions $\frac{\partial}{\partial U_i} E_A(U) = 0$ for all $i \in J$. Following the rules of matrix differential calculus [133, 197] we have that

$$\frac{\partial}{\partial U_i} \text{trace}\left(\Psi_D((U_i - F_j)^2)\right) = 2 \Psi'_D((U_i - F_j)^2) \bullet_J (U_i - F_j), \quad (4.9)$$

where the symbol “ \bullet_J ” denotes the Jordan product of matrices defined by $A \bullet_J B := (AB + BA)/2$. Using the distributive property of this product, the result in (4.9), and the abbreviations $w_{i,j}^{(\cdot)} := w_{(\cdot)}(|x_i - x_j|^2)$, the optimality conditions read

$$\frac{\partial}{\partial U_i} E_A(U) = X_i \bullet_J U_i - Y_i = 0, \quad \text{for all } i \in J, \quad (4.10)$$

with the following matrix expressions

$$\begin{aligned} X_i &:= \beta \sum_{j \in J} \Psi'_D((U_i - F_j)^2) w_{i,j}^D + 2(1 - \beta) \sum_{j \in J} \Psi'_S((U_i - U_j)^2) w_{i,j}^S, \\ Y_i &:= \beta \sum_{j \in J} \Psi'_D((U_i - F_j)^2) \bullet_J F_j w_{i,j}^D \\ &\quad + 2(1 - \beta) \sum_{j \in J} \Psi'_S((U_i - U_j)^2) \bullet_J U_j w_{i,j}^S. \end{aligned}$$

Finally, we solve (4.8) by the following iterative solution scheme

$$\begin{aligned} U_i^0 &:= F_i, \\ U_i^{k+1} &= (X_i^k)^{-1} \bullet_J Y_i^k, \end{aligned} \quad (4.11)$$

for all $i \in J, k \in \mathbb{N}_0$. Despite the fact that the Jordan product does not ensure the preservation of positive semidefiniteness (cf, Section 1.4) in the solution (4.11), we did not experience violations of this property in our experiments. Moreover, using this product enables us to apply our approach to indefinite matrices also, which may occur at locations with missing and noisy data in DT-MRI. The alternative logarithmic product is however not defined for indefinite matrices.

4.3 Related Filters within This Framework

An important consequence of our matrix-valued models (4.3) and (4.8) is that they represent a generalisation of several filters for matrix fields proposed in the literature. Before referring to those approaches, note that our models provide matrix-valued representatives of M-smoothing when $\beta = 1$ and bilateral filtering when $\beta = 0$, cf. Sections 2.2.1 and 2.2.2. Below we list various methods that can either be obtained as particular cases or be seen as related approaches to our models.

M1. *Affine-invariant weighted average* [18, 144, 166, 82], obtained from the iMNDS

- $\beta = 0$, Ψ_D and w_D do not play any role,
- $\Psi_S(d_{AI}^2) = d_{AI}^2$, $w_S = \text{Gaussian}$.

M2. *Log-Euclidean weighted average* [6], obtained from our iMNDS with

- as in M1, but with d_{LE} .

M3. *Affine-invariant regularisation and interpolation via a discrete geodesic marching scheme* [166], obtained from our iMNDS with

- $\beta = \frac{2}{2+\lambda}$, $\Psi_D(d^2) = d^2$, $w_D = \text{Gaussian}$,
- $\Psi_S(d_{AI}^2) = \text{any}$, $w_S = \text{unit disk}$.

M4. *Log-Euclidean regularisation and interpolation via a discrete geodesic marching scheme* [81], obtained from our iMNDS with

- as in M3, but with d_{LE} .

M5. A modified version of *bilateral filtering* [98], obtained from our iMNDS with

- $\beta = 0$, Ψ_D and w_D do not play any role,
- $\Psi_S(d_{LE}^2) = d_{LE}^2$, $w_S = \mu_1 \cdot d_T + \mu_2 \cdot |x_i - x_j|$,
with several possibilities of tensor distances d_T and weights $\mu_1, \mu_2 > 0$.

M6. *NL-means for DT-MRI* [232], related to our iMNDS with

- $\beta = 0$, Ψ_D and w_D do not play any role,

- $\Psi_S(d_{LE}^2) = d_{LE}^2$, $w_S = \frac{1}{z} \exp(-d_P^2/(\sigma^2))$,
where d_P is a patch distance as in (3.1), using d_{AI} between two tensors.

M7. *Component-based regularisation* [197, 193], continuous approach related to our iMNDS with

- $\beta = \frac{2}{2+\alpha}$, $\Psi_D(d_E^2) = d_E^2$, $w_D = \text{local}$,
- $\Psi_S(d_E^2) = \sqrt{d_E^2 + \epsilon^2}$, $w_S = \text{unit disk}$.

M8. *Operator-based regularisation* [193], continuous approach related to our aMNDS with

- $\beta = \frac{2}{2+\alpha}$, $\tilde{\Psi}_D(C^2) = \text{trace}(C^2)$, $w_D = \text{local}$,
- $\tilde{\Psi}_S(C^2) = \text{trace}(\sqrt{C^2 + \epsilon^2})$, $w_S = \text{unit disk}$.

Note that most of these methods do not exploit the utilisation of nonlocal information in the data similarity term ($\beta = 0$) or the radius of action of their smoothness term is restricted to the unit circle. Our MNDS approaches are more general and provide additional degrees of freedom to set up specialised filters for particular applications. In the following section we demonstrate the capabilities of our approaches for regularising synthetic and real-world matrix fields.

4.4 Numerical Experiments

Let us consider a 2D synthetic field $U = (U_i)_{i \in J} \in \text{Sym}_3^+(\mathbb{R})$ and its degraded version $F = (F_i)_{i \in J} \in \text{Sym}_3^+(\mathbb{R})$ obtained according to the noise model $F_i = |U_i + N_i|$. The eigenvalues of the noise matrix N_i stem from a Gaussian distribution with vanishing mean and standard deviation σ . Its eigenvectors are obtained by choosing three uniformly distributed angles and rotating N_i by these angles about the coordinate axes. Finally, we take its absolute value for positive definiteness. Considering that the eigenvalues of the original field U are in the range [1000, 4000], we create three noisy fields with $\sigma = 500, 1000, 2000$. Figure 4.2 shows our test data sets.

4.4.1 Two Isotropic Models: iMNDS-E and iMNDS-LE

We first focus on our isotropic MNDS (iMNDS) approach, which offers more degrees of freedom than the anisotropic model. Let us consider two realisations: The models iMNDS-E and iMNDS-LE when choosing the distances d_E and d_{LE} from Table 4.1, respectively. We do not test our approach with the distance d_{AI} as it has been shown in the literature [6] that it provides very similar results compared to using d_{LE} , although it is much more computationally demanding.




Figure 4.2: Detail of the synthetic data shown in Figure 4.1. *Left to right*: Original matrix field displaying four textures built up from tensors with different degrees of anisotropy, and its degraded versions by our noise model with $\sigma = 500, 1000, 2000$, which correspond to the 12.5%, 25% and 50% of the largest eigenvalue in the data.




Figure 4.3: Influence of the parameter β on the reconstruction error $\|F - U\|$ for varying levels of noise $\sigma = 500, 1000, 2000$. The penalisers P2 are used in both models. The spatial support of the weight functions were set to $\varpi_D = \varpi_S = 1$.

Independently of the model, the choice for the tonal penalisers Ψ_D and Ψ_S is made following two strategies:

- P1. Penalisers requiring no parameters at all. We use the Whittaker-Tikhonov penaliser for the data term $\Psi_D(d^2) = d^2$, and the Nashed-Scherzer penaliser [151] for the smoothness term $\Psi_S(d^2) = \gamma d^2 + \sqrt{d^2 + \epsilon^2}$, with $\epsilon = 1$ and $\gamma = \frac{1}{10\epsilon}$.
- P2. Penalisers with better edge-preservation properties, paying the price of introducing an additional *contrast parameter* λ . We use the classic Perona-Malik penaliser $\Psi(d^2) = \lambda^2 \ln \left(1 + \frac{d^2}{\lambda^2} \right)$ in both energy terms, choosing λ as the 1%-quantile of the distribution of distances in the noisy matrix field.

We use the soft window $w(h^2) = \exp \left(-\frac{h^2}{2\varpi^2} \right)$ for the spatial weights w_D and w_S , with support ϖ_D and ϖ_S respectively.

4.4.2 Influence of Parameters

Although we have specified the iMNDS-E and the iMNDS-LE models in the previous section, note that there are still some free parameters that will directly influence the denoising capabilities of our filters. Namely, the parameter β that counterbalances the contributions of the data and the smoothness term in (4.3)–(4.4), and the supports ϖ_D and ϖ_S of the spatial weighting functions that allow smoothing within larger neighbourhoods. Figure 4.3 shows the influence of the parameter β on the models iMNDS-E and iMNDS-LE with respect to the *reconstruction error* measured by

$$\|F - U\| := \left(\sum_{i=1}^N \|F_i - U_i\|_F \right)^{1/2}, \quad (4.12)$$

where F denotes the original, uncorrupted matrix field and U stands for its denoised version. The non-trivial steady-state is shown for $\beta \in (0, 1]$. We see that there is a value $\hat{\beta}$ for which the restoration is optimal.

We now want to quantify the influence of the spatial supports ϖ_D, ϖ_S . Naturally, making them larger will increase the number of arithmetic operations in (4.4). However, the restoration quality could be improved and the steady-state solution could be reached in a shorter time. If we vary these parameters in the range $[0, 4]$, there are 25 possible combinations for the pair (ϖ_D, ϖ_S) , which we arrange as O_0, \dots, O_{24} following the ordering shown in Figure 4.4 (*top*). The diagonal lines in the figure group the combinations according to *complexity order* (CO), that is, configurations with equal or increasing number of operations. Figure 4.4 (*bottom*) shows the restoration error (*left*), the logarithmic computational time (*middle*), and the overall performance (*right*) of the iMNDS-E model. The last measurement is simply the mean between the first two normalised measurements. We see that the configuration with the best performance in terms of quality and convergence is $O_8 = (\varpi_D, \varpi_S) = (1, 2)$ for $\beta = 0.9$. It is worth mentioning that the configurations $O_5 = (2, 0)$, $O_6 = (3, 0)$ and $O_{14} = (4, 0)$, lead to good results despite the fact that they allow only for the incorporation of neighbourhood information in the data term. This is actually in agreement with our findings in Subsection 2.4.3, where we showed that nonlocal M-smoothers produce similar results compared to classical regularisation methods. These observations also hold for the iMNDS-LE model.

4.4.3 Performance

Now that we have seen how the different parameters influence our filtering approach, let us now report on the best results achieved by both models iMNDS-E and iMNDS-LE, as well as our anisotropic MNDS (aMNDS) model. All results are computed with optimised free parameters $\beta, \varpi_D, \varpi_S$.

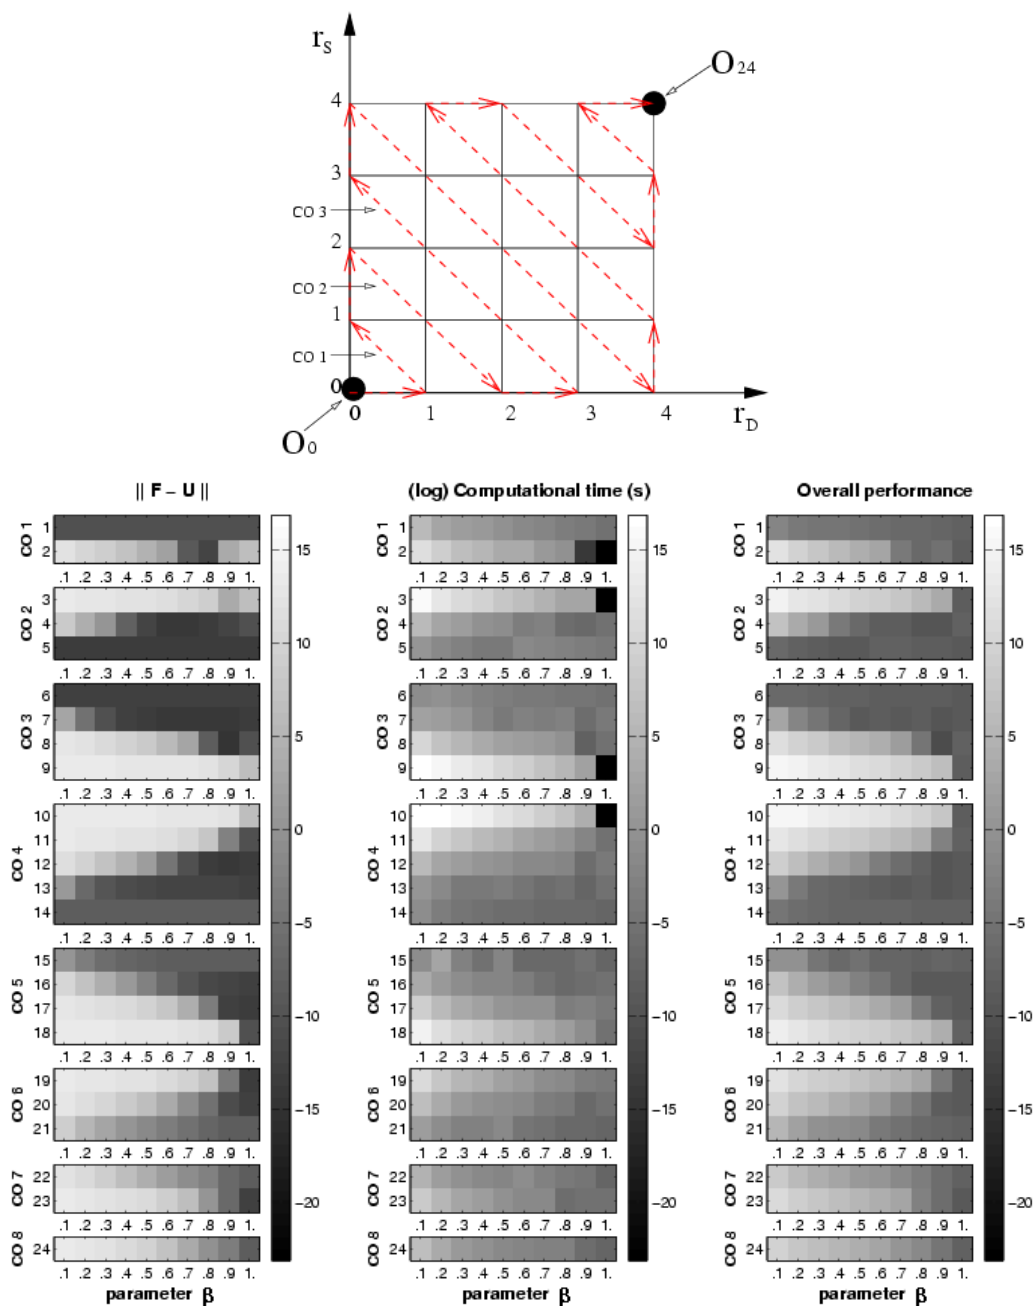


Figure 4.4: *Top*: Ordering O_0, \dots, O_{24} for the different combinations for the pair (ϖ_D, ϖ_S) , grouped according to complexity order CO_i ($i = 1, \dots, 8$). *Bottom, left to right*: Normalised restoration quality, computational time, and overall performance of the iMNSD-E model with P2 penalisers. Test run on the noisy tensor field with $\sigma = 1000$ from Figure 4.2.




Figure 4.5: Restoration error for both models iMNDS-E and iMNDS-LE. They are tested with both set of penalisers P1 and P2. All other parameters have been optimised. As expected, the edge-preserving penalisers P2 provide better results.

Table 4.2: Best filtering results for our models iMNDS-E, iMNDS-LE and aMNDS under the noise level $\sigma = 500, 1000, 2000$. All parameters have been optimised. Figure 4.6 displays the corresponding filtered images.

σ	model	ϖ_D	ϖ_S	β	$\ F - U\ $	#iter	time(sec.)
500	iMNDS-E	2	3	0.9	167	33	0.69
	iMNDS-LE	1	2	0.9	264	71	19.31
	aMNDS	2	3	0.9	219	29	18.00
1000	iMNDS-E	1	2	0.9	409	72	0.71
	iMNDS-LE	2	0	0.2	568	236	57.97
	aMNDS	1	2	0.9	494	120	35.49
2000	iMNDS-E	2	0	0.1	1238	256	1.91
	iMNDS-LE	2	1	0.9	1214	32	9.50
	aMNDS	2	0	0.1	1230	40	9.10

Figure 4.5 depicts the restoration error achieved by our models iMNDS-E and iMNDS-LE with optimised parameters. We notice that in both cases the best performance is obtained with the Perona-Malik penalisers P2. Numerically, these results are outlined in Table 4.2. The table also shows the best results of the aMNDS




Figure 4.6: *First row, from left to right:* Degraded matrix fields with noise level $\sigma = 500, 1000, 2000$. *Second row:* Corresponding best filtering results for the iMNDS-E model, cf. Table 4.2. *Third row:* Corresponding results for the iMNDS-LE model. *Fourth row:* Corresponding results for the aMNDS model.

model obtained with Perona-Malik penalisers. It is clear that all models benefit from nonlocal smoothing by considering $\varpi_D, \varpi_S > 0$. Note that the iMNDS-E model is considerably faster than the other two models: iMNDS-LE is burdened with computation of logarithms and exponentials of matrices, while aMNDS needs to perform spectral decompositions and multiplications of matrices. Computations have been performed on a 1.86 GHz Intel Core 2 Duo processor (without exploiting multitasking) executing C code.

Figure 4.6 shows the filtered matrix fields for the results presented in Table 4.2. At any noise level, the iMNDS-E model performs slightly better than the iMNDS-LE and the aMNDS models. This is clearly observable for the case $\sigma = 1000$, but it can also be seen when $\sigma = 500$ at the lower part of the inner ring. Particularly noticeable in the example $\sigma = 1000$ is that both the edges of the image structures and the anisotropy of the matrices are better preserved with the iMNDS-E model. Moreover, the eigenvalue-swelling-effect at the edges is less perceptible in the iMNDS-E model than in the other two models.

4.4.4 3D DT-MRI data

In DT-MRI, noisy *diffusion weighted images* (DWIs) are used to estimate the diffusion tensors via regression analysis. It is known that DWIs are perturbed by Rician noise [93]. However, the noise distribution of the diffusion tensors obeys a multivariate Gaussian distribution, as it has been statistically proven by Pajevic and Basser [163]. Here, as it was done in the previous section, we directly apply our filtering framework to the tensor field, and not to the scalar DWIs. We consider the 3D DT-MRI dataset of a human head consisting of a $128 \times 128 \times 30$ -field of positive definite matrices displayed earlier in Section 1.3.

The left column in Figure 4.7 shows two magnified areas about the *genu* and the *corpus callosum* where it can be observed that the original data is significantly affected by noise. Moreover, there exist several locations where the data is completely missing, specially in the anisotropic regions. The FA map, computed according to (1.5), shows isotropic areas in dark, while the anisotropic areas appear bright. It can be noticed that the FA map is also affected by the presence of noise.

The middle and right columns of Figure 4.7 exhibit the smoothing results provided by our isotropic MNDS model along 2D slices and in full 3D, respectively. In both settings, we used Perona-Malik type penalisers, the Euclidean distance d_E and the parameters $(\varpi_D, \varpi_S, \beta) = (1, 2, 0.9)$. Both results look visually pleasant, with edges well preserved and localised. The 3D iMNDS however can better reconstruct the missing data thanks to the utilisation of the full 3D structural information. Despite the poor resolution (40×55 pixels) of the FA maps, the distinction between isotropic and anisotropic areas seems much more notorious now. The FA map corresponding to the 3D filtering looks better smoothed and it still preserves the edge-structures along the anisotropic (bright) areas.

We now compare our results to those obtained with another matrix-valued approach, namely the 2D and 3D *coherence-enhancing diffusion* filter for matrix fields (MCED) proposed in [49, 52]. Figure 4.8 shows the corresponding results under of the MCED filter with evolution time $t = 2$. Note that the 3D filtering respects better the underlying structure of the ventricles, which can also be observed in its FA map. Compared to our approach, the MCED filter produces slightly more diffusive structures, but thanks to its coherence-enhancing properties the anisotropic areas look stronger (brighter) than in our approach.

Finally, we show in Figure 4.9 that our MNDS approach is a powerful preprocessing technique that provides algorithms such as tractography [146] with filtered and gap-completed tensors fields, which results in much better fibre tracts. These characteristics are important for the study of diseases associated with certain abnormalities in brain anatomy [157].

4.5 Summary

In this chapter, the original scalar-valued NDS filtering framework has been extended in full generality to the matrix-valued setting. We presented two alternative matrix-valued filters, namely an isotropic and an anisotropic approach. The isotropic model provides more degrees of freedom to tune a desired filter and it is computationally inexpensive compared to its anisotropic model. In practice, both models provide similar filtering results on synthetic and DT-MRI data sets.

It is important to mention that our matrix-valued NDS (MNDS) framework generalise several known approaches suggested in the literature for the filtering and interpolation of DT-MRI data, including those employing the log-Euclidean framework to preserve positive definiteness of the data. We emphasise that our methodology is generic and thus not restricted to DT-MRI denoising. It can be applied to any multi-valued image with values in the space of symmetric matrices, positive (semi-)definite or not.




Figure 4.7: 2D vs 3D filtering with our iMNDS approach. *Left column:* Scaled-up images showing details of the anisotropic zones around the genu (top) and the corpus callosum (middle), which appear with high intensity in the FA map (bottom). *Middle column:* Smoothing results with the iMNDS approach in 2D, with its corresponding FA map. *Right column:* Results with the 3D iMNDS.




Figure 4.8: 2D vs 3D filtering with the matrix-valued CED (MCED) approach [49, 52]. *Left column:* Scaled-up images showing details of the anisotropic zones around the genu (top) and the corpus callosum (middle), which appear with high intensity in the FA map (bottom). *Middle column:* Smoothing results with MCED in 2D, and its corresponding FA map. *Right column:* Results with 3D MCED.




Figure 4.9: Tractography performed on the original 3D DT-MRI volume (*top*) and after smoothing with our iMNDS model (*bottom*). Fibre tracking algorithm courtesy of Thomas Schultz, University of Chicago.

Part II

**Adaptive Morphology
for Scalar- and Matrix-Valued Images**

5

Introduction to Mathematical Morphology

Mathematical morphology aims at extracting and enhancing and shape information from image objects. It originated in 1964 at the École des Mines de Paris in the works of Matheron [140] and Serra [187]. The so-called *set-theoretical* approach developed by its founders is based on set theory, integral geometry and lattice algebra. All morphological operations are based on two basic processes: *dilation* and *erosion*, which respectively expand and shrink the boundaries of an object to an extent specified by the shape of a pre-defined convex set called *structuring element* (SE).

A different approach to morphology is the so-called *continuous-scale* framework. In this context, the morphological operators are described by partial differential equations (PDEs). This approach allows for digital scalability of the SE, which provides all morphological operators and their applications with *sub-pixel accurate* results. In this chapter we briefly review both frameworks, commenting on their advantages and shortcomings. After describing these methods for scalar-valued images we focus on the extension of the PDE-based approach to matrix fields, which will be the basis for our work in the subsequent Chapter 6.

5.1 Scalar-Valued Morphology

5.1.1 Set-Theoretical Approach

Classical morphology consists of a number of operations that match a d -dimensional scalar-valued signal $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with a bounded set of known shape $B \subset \mathbb{R}^d$ called *structuring element* (SE). For ease of presentation, let us focus on bi-dimensional images ($d = 2$). In this case, the pre-defined shapes for B could be squares, diamonds, discs or ellipses defined on the discrete pixel grid. The fundamental dilation and erosion processes are defined by [188, 194, 135, 35, 196]

$$\text{dilation:} \quad (f \oplus B)(x) := \sup_{z \in B} \{f(x - z)\}, \quad (5.1)$$

$$\text{erosion:} \quad (f \ominus B)(x) := \inf_{z \in B} \{f(x + z)\}. \quad (5.2)$$

On discrete lattices, the dilated (resp. eroded) value at a given pixel x is the maximum (resp. minimum) value of the image in the window defined by the SE centred at x . Erosion is the dual operation of dilation. That is, eroding an image corresponds to

dilating its background (or the inverse image). Applying dilation/erosion to an object leads to the expansion/shrinking of its boundaries.

These two operations are the basis to define more complex operators. For instance, another pair of dual operations, *opening* and *closing*, can be obtained by concatenation of dilation and erosion:

$$\text{opening: } (f \circ B)(x) := ((f \ominus B) \oplus B)(x), \quad (5.3)$$

$$\text{closing: } (f \bullet B)(x) := ((f \oplus B) \ominus B)(x). \quad (5.4)$$

These operators act as smoothers, allowing the removal of small objects and small holes, respectively. They are also the basis of other morphological filters [196]. The first row of Figures 5.2 and 5.3 show the dilation, erosion, opening and closing of a greyscale image with a disc of radius 3 pixels as SE.

5.1.2 Continuous-Scale Approach

For convex structuring elements there exists an alternative formalisation of the dilation and erosion processes in terms of partial differential equations (PDEs) due to [35, 211, 4, 2]. In this context, the SE is represented as tB with a parameter $t \geq 0$ that allows for continuous-scaling. Thus, the solutions to $u(x, t) = (f \oplus tB)(x)$ and $u(x, t) = (f \ominus tB)(x)$ can be equivalently obtained by solving respectively

$$\text{dilation: } \partial_t u(x, t) = \sup_{z \in B} \langle z, \nabla u(x, t) \rangle, \quad (5.5)$$

$$\text{erosion: } \partial_t u(x, t) = \inf_{z \in B} \langle z, \nabla u(x, t) \rangle, \quad (5.6)$$

with initial condition $u(x, 0) = f(x)$ and homogeneous Neumann boundary conditions at the image boundary, i.e. $\langle \nabla u, n \rangle = 0$. Here n is the unit normal vector pointing outside the image domain, $\nabla = (\partial_x, \partial_y)^\top$ is the spatial nabla operator, and $\langle \cdot, \cdot \rangle$ denotes the Euclidean vector product. In particular, choosing the unitary disc $B := \{z \in \mathbb{R}^2, \|z\|_2 \leq 1\}$, the processes (5.5)–(5.6) become

$$\text{dilation: } \partial_t u(x, t) = \|\nabla u(x, t)\|_2, \quad (5.7)$$

$$\text{erosion: } \partial_t u(x, t) = -\|\nabla u(x, t)\|_2. \quad (5.8)$$

That is, the solution $u(\cdot, t)$ at “time” t is the dilation/erosion of f with a disc of radius t and centre 0 as SE. Similarly, other PDE evolutions can be obtained for squares and diamonds as SE, see [35, 182]. The operations of opening and closing in the PDE framework are obtained by concatenating the solutions of (5.5)–(5.6) in (5.3)–(5.4). The second and third rows of Figures 5.2 and 5.3 show the PDE-based dilation, erosion, opening and closing of a greyscale image with a disc of radius $t = n\tau = 3$ ($n = 15, \tau = 0.2$) using two different numerical schemes to discretise the PDEs, which are described later in Subsection 5.1.4.

5.1.3 Set-Theoretical vs. Continuous-Scale Approach

In the previous two sections we showed that the set-theoretical and the continuous-scale approaches provide equivalent formulations of the dilation and erosion processes. However, they have important differences that make one approach or the other more suitable for certain applications.

The set-theoretical framework is computationally inexpensive and very easy to implement. It just amounts to computing the maximum/minimum value within the structuring element at each pixel location. However, as this approach is based on set and lattice theory, the structuring elements are defined in the discrete setting. This leads to the problem that shapes such as circles and ellipses can only be represented by discrete approximations, e.g. discs represented by squares or hexagons, which is eventually reflected on the dilated/eroded edges. See Figure 5.2(a)–(b) and Figure 5.3(a)–(b).

The PDE-based framework allows for a proper digital representation of structuring elements via the continuous scaling parameter t , which provides sub-pixel accuracy in the resulting morphological operations. This approach originates the concept of morphological scale-space, opening up the analysis of shapes under continuous morphological evolutions. These benefits come at the price of more complex and demanding implementations of the PDEs, which usually introduce an undesirable blurring of edges. See Figure 5.2(c)–(d) and Figure 5.3(c)–(d).

In the sequel, we adopt the PDE-based framework due to its several advantages, versatility, and because it will be the main topic of Chapter 6. In the following section we review two algorithms to solve the morphological PDEs. The first one is a classical scheme that introduces dissipative numerical artifacts, and the second one is a predictor-corrector scheme that removes these blurring effects. See Figure 5.2(e)–(f) and Figure 5.3(e)–(f).

5.1.4 Numerical Schemes for Continuous-Scale Morphology

The evolutionary equations (5.5)–(5.6) belong to the class of first-order hyperbolic PDEs that describe the propagation of wavefronts [126, 116]. Numerical algorithms to solve these PDEs with first-order accuracy are the *Osher-Sethian (OS) scheme* [162, 192, 159] and the *Rouy-Tourin (RT) scheme* [178, 212]. Both alternatives suffer from significant blurring at discontinuities. Higher-order accuracy algorithms that correct these dissipative artifacts are the *high-resolution OS scheme* [162] and the *flux-corrected transport (FCT) scheme* [33].

5.1.4.1 The Rouy-Tourin Scheme

To describe this scheme we first need to introduce some basic notation. Let h_x, h_y by the spatial grid size and $u_{i,j}^n$ the grey value of a 2D image u at the pixel centred in

$(ih_x, jh_y) \in \mathbb{R}^2$ at the time-level $t = n\tau$ of the evolution, with n the iteration number and τ the time step. Furthermore, we employ standard forward, backward, and central difference operators, i.e.,

$$D_+^x u_{i,j}^n := u_{i+1,j}^n - u_{i,j}^n \quad (5.9)$$

and

$$D_-^x u_{i,j}^n := u_{i,j}^n - u_{i-1,j}^n \quad (5.10)$$

and finally

$$D_c^x u_{i,j}^n := (u_{i+1,j}^n - u_{i-1,j}^n) / 2 \quad (5.11)$$

here in x – but analogously in y –direction.

Rouy and Tourin [178] proposed an *upwinding* first-order finite difference method to discretise the dilation process (5.7) as

$$\partial_t u = ((\partial_x u)^2 + \partial_y u)^2)^{1/2} \quad (5.12)$$

$$\begin{aligned} \frac{u_{i,j}^{n+1} - u_{i,j}^n}{\tau} = & \left(\max \left(\frac{1}{h_x} \max(-D_-^x u_{i,j}^n, 0), \frac{1}{h_x} \max(D_+^x u_{i,j}^n, 0) \right)^2 \right. \\ & \left. + \max \left(\frac{1}{h_y} \max(-D_-^y u_{i,j}^n, 0), \frac{1}{h_y} \max(D_+^y u_{i,j}^n, 0) \right)^2 \right)^{1/2}, \end{aligned}$$

and the erosion process (5.8) as

$$\partial_t u = -((\partial_x u)^2 + \partial_y u)^2)^{1/2} \quad (5.13)$$

$$\begin{aligned} \frac{u_{i,j}^{n+1} - u_{i,j}^n}{\tau} = & - \left(\min \left(\frac{1}{h_x} \min(-D_-^x u_{i,j}^n, 0), \frac{1}{h_x} \min(D_+^x u_{i,j}^n, 0) \right)^2 \right. \\ & \left. + \min \left(\frac{1}{h_y} \min(-D_-^y u_{i,j}^n, 0), \frac{1}{h_y} \min(D_+^y u_{i,j}^n, 0) \right)^2 \right)^{1/2}, \end{aligned}$$

where the time derivative $\partial_t u$ is discretised by Euler forward differences. The spatial derivatives $\partial_x u, \partial_y u$ are approximated by *upwind* schemes: one-sided differences in direction from which information comes, see details in [178, 192]. This algorithm is numerically stable for $\tau \leq 1/\sqrt{2}$. Its performance is very similar to that of the first-order OS scheme, hence we refrain from describing the latter.

In the second row of Figures 5.2 and 5.3 we can observe the blurring artifacts produced by the RT-scheme. Despite this shortcoming, these blurred solutions can be corrected by the FCT-scheme that we describe in the following subsection.

5.1.4.2 The Flux-Corrected Transport Scheme

As mentioned earlier, first-order difference schemes introduce dissipative artifacts at discontinuities. Osher and Sethian [162] proposed a high-resolution scheme that compensates for the blurring with a second-order correction step. This method however does not completely remove the dissipative artifacts. Breuß and Weickert [33] introduced the *flux-corrected transport (FCT) scheme* as an alternative predictor-corrector method that achieves much sharper edges. The FCT scheme results in accurate and (largely) rotationally invariant discrete representations of continuous-scale morphological dilation/erosion. It can be realised as a new variant of a technique originally proposed by Boris and Book [24, 25, 26] in the context of fluid flow simulation.

The FCT scheme also relies on one-sided upwind differences, but it additionally quantifies the undesirable blurring effects introduced by the upwinding which are subsequently corrected by *stabilised inverse diffusion*. To see this, let us consider the quantities (cf. [33])

$$\lceil^x u_{i,j}^n := \frac{\tau}{2h_x} |D_c^x u_{i,j}^n| + \frac{\tau}{2h_x} D_+^x u_{i,j}^n - \frac{\tau}{2h_x} D_-^x u_{i,j}^n, \quad (5.14)$$

$$\lceil^y u_{i,j}^n := \frac{\tau}{2h_y} |D_c^y u_{i,j}^n| + \frac{\tau}{2h_y} D_+^y u_{i,j}^n - \frac{\tau}{2h_y} D_-^y u_{i,j}^n, \quad (5.15)$$

which describe the upwinding incorporated in the spatial derivatives of the dilation process (5.12). Note that the central differences above incorporate a second-order error which is nondiffusive, while the one-sided differences are discrete diffusive fluxes. Thus, a spatial discretisation relying on (5.14)–(5.15) such as the Rouy-Tourin scheme introduces exactly these diffusive fluxes. The FCT procedure then inverts the corresponding numerical diffusion using the predicted data in the corrector step.

Focusing on the dilation process, let us denote by $u_{i,j}^{n+1,\text{pred}}$ the *predicted* solution computed with the Rouy-Tourin method (5.12). Then, the *corrected* solution obtained by the FCT method reads

$$u_{i,j}^{n+1} = u_{i,j}^{n+1,\text{pred}} + q_h^{n+1,\text{pred}} - q_d^{n+1,\text{pred}}. \quad (5.16)$$

One can identify the higher-order term $q_h^{n+1,\text{pred}}$ in (5.16) as

$$q_h^{n+1,\text{pred}} := \left(\left(\frac{\tau}{h_x} |D_c^x u_{i,j}^{n+1,\text{pred}}| \right)^2 + \left(\frac{\tau}{h_y} |D_c^y u_{i,j}^{n+1,\text{pred}}| \right)^2 \right)^{1/2}. \quad (5.17)$$

For the lower-order term $q_d^{n+1,\text{pred}}$ in (5.16) we make use of the quantities

$$g_{i+1/2,j} := \text{mm} \left(D_-^x u_{i,j}^{n+1,\text{pred}}, \frac{\tau}{2h_x} D_+^x u_{i,j}^{n+1,\text{pred}}, D_+^x u_{i+1,j}^{n+1,\text{pred}} \right) \quad (5.18)$$

$$g_{i,j+1/2} := \text{mm} \left(D_-^y u_{i,j}^{n+1,\text{pred}}, \frac{\tau}{2h_y} D_+^y u_{i,j}^{n+1,\text{pred}}, D_+^y u_{i,j+1}^{n+1,\text{pred}} \right) \quad (5.19)$$

where $\text{mm}(\cdot, \cdot, \cdot)$ is the *scalar* minmod-function defined for three arguments as

$$\text{mm}(a_1, a_2, a_3) := \begin{cases} \inf(a_1, a_2, a_3) & \text{for } a_1, a_2, a_3 > 0, \\ \sup(a_1, a_2, a_3) & \text{for } a_1, a_2, a_3 < 0, \\ 0 & \text{else.} \end{cases} \quad (5.20)$$

With these abbreviations we set

$$\delta^x u_{i,j}^{n+1,\text{pred}} := \frac{\tau}{h_x} \left| D_c^x u_{i,j}^{n+1,\text{pred}} \right| + g_{i+1/2,j} - g_{i-1/2,j} \quad (5.21)$$

$$\delta^y u_{i,j}^{n+1,\text{pred}} := \frac{\tau}{h_y} \left| D_c^y u_{i,j}^{n+1,\text{pred}} \right| + g_{i,j+1/2} - g_{i,j-1/2} \quad (5.22)$$

which finally yields the second new term in (5.16) as

$$q_d^{n+1,\text{pred}} := \left(\left(\delta^x u_{i,j}^{n+1,\text{pred}} \right)^2 + \left(\delta^y u_{i,j}^{n+1,\text{pred}} \right)^2 \right)^{1/2}. \quad (5.23)$$

For the sake of completeness, let us describe the FCT scheme for the erosion operation. Using the *predicted* solution $u_{i,j}^{n+1,\text{pred}}$ of the Rouy-Tourin scheme (5.13), the FCT method provides a *corrected* solution $u_{i,j}^{n+1}$ as defined in (5.16) utilising the definitions

$$q_h^{n+1,\text{pred}} := - \left(\left(\frac{\tau}{h_x} \left| D_c^x u_{i,j}^{n+1,\text{pred}} \right| \right)^2 + \left(\frac{\tau}{h_y} \left| D_c^y u_{i,j}^{n+1,\text{pred}} \right| \right)^2 \right)^{1/2}, \quad (5.24)$$

$$q_d^{n+1,\text{pred}} := - \left(\left(\delta^x u_{i,j}^{n+1,\text{pred}} \right)^2 + \left(\delta^y u_{i,j}^{n+1,\text{pred}} \right)^2 \right)^{1/2}, \quad (5.25)$$

with

$$\delta^x u_{i,j}^{n+1,\text{pred}} := \frac{\tau}{h_x} \left| D_c^x u_{i,j}^{n+1,\text{pred}} \right| - g_{i+1/2,j} + g_{i-1/2,j}, \quad (5.26)$$

$$\delta^y u_{i,j}^{n+1,\text{pred}} := \frac{\tau}{h_y} \left| D_c^y u_{i,j}^{n+1,\text{pred}} \right| - g_{i,j+1/2} + g_{i,j-1/2}. \quad (5.27)$$

As the method of Rouy and Tourin, the FCT approach is numerically stable for $\tau \leq 1/\sqrt{2}$. We refer to [33] for the complete details of this scheme and to [34] for its extension to dilation/erosion processes with diamonds and ellipses as structuring elements.

In the third row of Figures 5.2 and 5.3 we show the ability of the FTC-scheme to correct the blurring artifacts produced by the RT-scheme. This results in very sharp edges. Furthermore, note the ball-shaped structuring element has a much better digital representation in the PDE-based approach compared to the set-theoretical framework.

5.2 Matrix-Valued Morphology

Defining morphological operations for matrix fields is not straightforward. In the set-theoretical approach (Section 5.1.1), the notions of maximum and minimum in the scalar-valued case do not carry over directly to matrix-valued setting. Based upon the *Loewner ordering* for symmetric matrices, Burgeth *et al.* [50, 46] extended the maximum/minimum operations to matrix fields, allowing a proper definitions of the dilation/erosion processes in a set-theoretic style. In the continuous-scale approach (Section 5.1.2), Burgeth *et al.* [45, 43] proposed matrix-valued versions of the PDEs governing dilation/erosion. In the following two subsections we describe these PDEs and the numerical schemes to solve them.

5.2.1 PDE-Based Morphology for Matrix Fields

Recalling our notation, a matrix field is considered as a mapping $U : \Omega \subset \mathbb{R}^d \rightarrow \text{Sym}_m(\mathbb{R})$ from a d -dimensional image domain into the set of symmetric $m \times m$ -matrices with real entries. Our aim is to transfer the scalar-valued PDEs

$$\partial_t u = \pm \|\nabla u\|_2 \quad (5.28)$$

governing dilation (+) and erosion (−) with a ball-shaped structuring element to the matrix-valued setting. Equations (5.28) contain three operators that need to be analogously formalised for matrix fields: the partial derivatives $\partial_{(\cdot)}$, the gradient operator ∇ and the Euclidean vector norm $\|\cdot\|_2$. An elegant tool for doing so is *operator-algebraic framework* for matrix fields proposed by Burgeth *et al.* [48, 47], which we summarised in Section 1.4. The basic idea behind this approach is to consider symmetric matrices as a natural generalisation of real numbers with a rich algebraic structure. Following such principles, the 3D matrix-valued counterpart of (5.28) reads

$$\begin{aligned} \bar{\partial}_t U &= \pm |\bar{\nabla} U|_2 \\ &= \pm \sqrt{|\bar{\partial}_x U|^2 + |\bar{\partial}_y U|^2 + |\bar{\partial}_z U|^2}, \end{aligned} \quad (5.29)$$

where $\bar{\partial}_{(\cdot)}$, $\bar{\nabla}$, $|\cdot|_p$ and $|\cdot|^q$ are characterised in Section 1.4. The initial condition of these PDEs is $U(x, 0) = F(x)$, with $F : \Omega \rightarrow \text{Sym}_m(\mathbb{R})$ the input tensor field.

5.2.2 Matrix-Valued Numerical Schemes

From Section 5.1.4 we know that the nonlinear PDEs (5.28) can be numerically solved by first-order finite difference methods such as the *Osher-Sethian (OS) scheme* [162, 192, 159] and the *Rouy-Tourin (RT) scheme* [178, 212]. They are capable of correctly capturing propagating shocks, however, at the price of introducing some dissipation

and blurring of edges. A remedy is provided by the *high-resolution OS scheme* [162] and the *flux-corrected transport (FCT) scheme* [33]. All these methods have been transferred to the matrix-valued setting to solve (5.29), see [45, 43].

Here we provide the necessary elements to generalise the scalar-valued RT-scheme (Section 5.1.4.1) and FCT-scheme (Section 5.1.4.2) to matrix fields. These numerical methods include several scalar operations that need to be analogously formulated for symmetric matrices. Once again we make use of the operator-algebraic framework described in Section 1.4 to define functions operating on matrices such as square, square-root and absolute value. Moreover, in correspondence with (5.9)–(5.11) it is straightforward to define one-sided and central differences of 3D matrix fields in $\text{Sym}_m(\mathbb{R})$ as

$$D_+^x U^n(ih_x, jh_y, kh_z) \quad (5.30)$$

$$:= U^n((i+1)h_x, jh_y, kh_z) - U^n(ih_x, jh_y, kh_z),$$

$$D_-^x U^n(ih_x, jh_y, kh_z) \quad (5.31)$$

$$:= U^n(ih_x, jh_y, kh_z) - U^n((i-1)h_x, jh_y, kh_z),$$

$$D_c^x U^n(ih_x, jh_y, kh_z) \quad (5.32)$$

$$:= (U^n((i+1)h_x, jh_y, kh_z) - U^n((i-1)h_x, jh_y, kh_z)) / 2,$$

here in x - but equivalently in y - and z -direction. To avoid confusion with the subscript notation for matrix components we wrote $U(ih_x, jh_y, kh_z)$ to indicate the (matrix-) value of the matrix field evaluated at the voxel centred at $(ih_x, jh_y, kh_z) \in \mathbb{R}^3$. We also need the notion of supremum/infimum (maximum/minimum) of two matrices provided in Section 1.4 to generalise the minmod function as

$$\text{mm}(A_1, A_2) := \begin{cases} \inf(A_1, A_2) & \text{for } A_1 > 0 \text{ and } A_2 > 0, \\ \sup(A_1, A_2) & \text{for } A_1 < 0 \text{ and } A_2 < 0, \\ 0 & \text{else,} \end{cases} \quad (5.33)$$

in the case of two matrices, while for three matrices we define

$$\text{mm}(A_1, A_2, A_3) := \begin{cases} \inf_{opt}(A_1, A_2, A_3) & \text{for } A_i > 0, i = 1, 2, 3, \\ \sup_{opt}(A_1, A_2, A_3) & \text{for } A_i < 0, i = 1, 2, 3, \\ 0 & \text{else.} \end{cases} \quad (5.34)$$

The functions \inf_{opt} and \sup_{opt} are given by

$$\inf_{opt}(A_1, A_2, A_3) = -\sup_{opt}(-A_1, -A_2, -A_3), \quad (5.35)$$

$$\sup_{opt}(A_1, A_2, A_3) := S_m - \lambda_{opt} I_m, \quad (5.36)$$

with I_m the $m \times m$ -identity matrix and

$$S_m := (S_1 + S_2 + S_3)/3, \quad (5.37)$$

$$S_1 := \sup(A_1, \sup(A_2, A_3)), \quad (5.38)$$

$$S_2 := \sup(A_2, \sup(A_1, A_3)), \quad (5.39)$$

$$S_3 := \sup(A_3, \sup(A_1, A_2)), \quad (5.40)$$

$$\lambda_{opt} := \min_{\substack{i=1,2,3 \\ j=1,\dots,m}} (\mu_{ij}), \quad (5.41)$$

where μ_{ij} is the j -th eigenvalue of the matrix $S_m - A_i$, $i = 1, 2, 3$ (cf. [43]).

With these definitions at our disposal, the original scalar-valued RT-scheme and FCT-scheme are directly translated to the 2D and 3D matrix-valued setting. Let us now see how the new schemes work in practice. Figures 5.4 and 5.5 show the operations of dilation, erosion, opening and closing applied to a matrix field of 3×3 -symmetric matrices using the PDE approach and solved by both the RT-scheme and FCT-scheme. One can observe that the behaviour of both methods is as in the scalar-valued setting. That is, the FCT-scheme produces sharp results, being able to correct the blurring artifacts of the RT-scheme.

5.3 Summary

In this chapter we have introduced the elementary notions of mathematical morphology in the classical set-theoretical context as well as in the scope of partial differential equations (PDEs), commenting on the advantages and shortcomings of both frameworks. We covered both the scalar- and matrix-valued realisations of the two fundamental processes, dilation and erosion, that serve as building blocks for more complex morphological operations. We focused on the PDE-based approach, particularly on the matrix-valued dilation/erosion with a ball-shaped structuring element. We showed how to solve these PDEs with matrix-valued versions of the popular *Rouy-Tourin* scheme and the high-resolution *flux-corrected transport* scheme.

The approaches to mathematical morphology presented in this chapter can be considered as *isotropic* processes where all pixel/voxel locations are treated equally. For example, the operation of dilation (resp. erosion) grows (resp. shrinks) the object's boundaries by the same factor in all directions. Despite the numerous applications of these operations, they are not suitable for completing broken lines or curves that underlie anisotropic structures. In the next chapter we propose a PDE-based formulation for *adaptive, anisotropic* morphology for scalar and matrix fields that allows for reconstructing/enhancing broken lines and other coherent, flow-like structures.




Figure 5.1: MRI slice of a human head

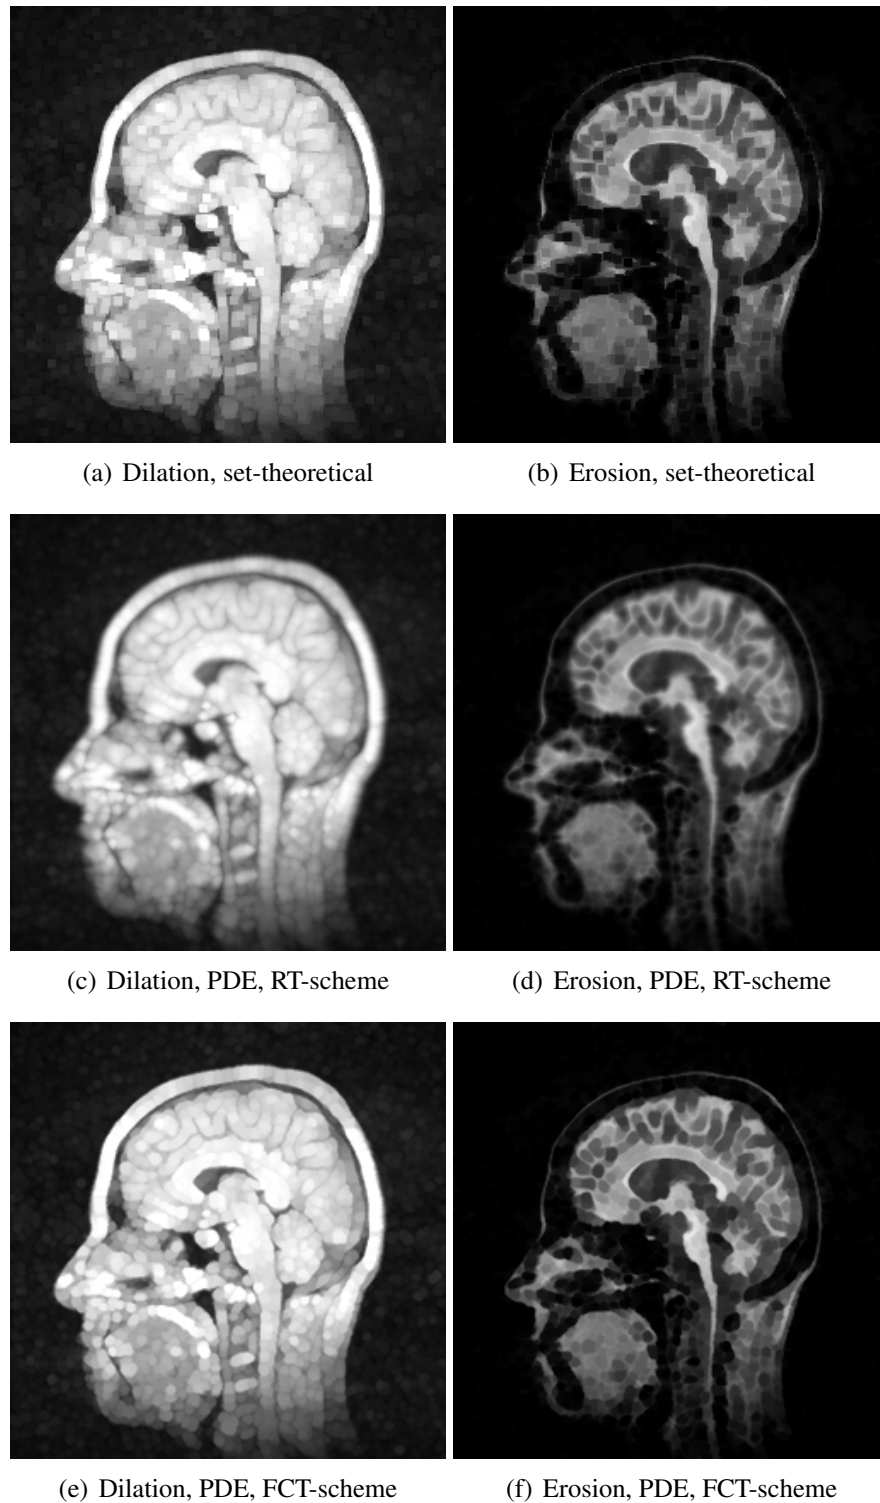


Figure 5.2: Scalar-valued morphology: Set-theoretical vs PDE-based dilation/erosion with a dis of radius 3 as structuring element. Original image shown in Fig. 5.1.

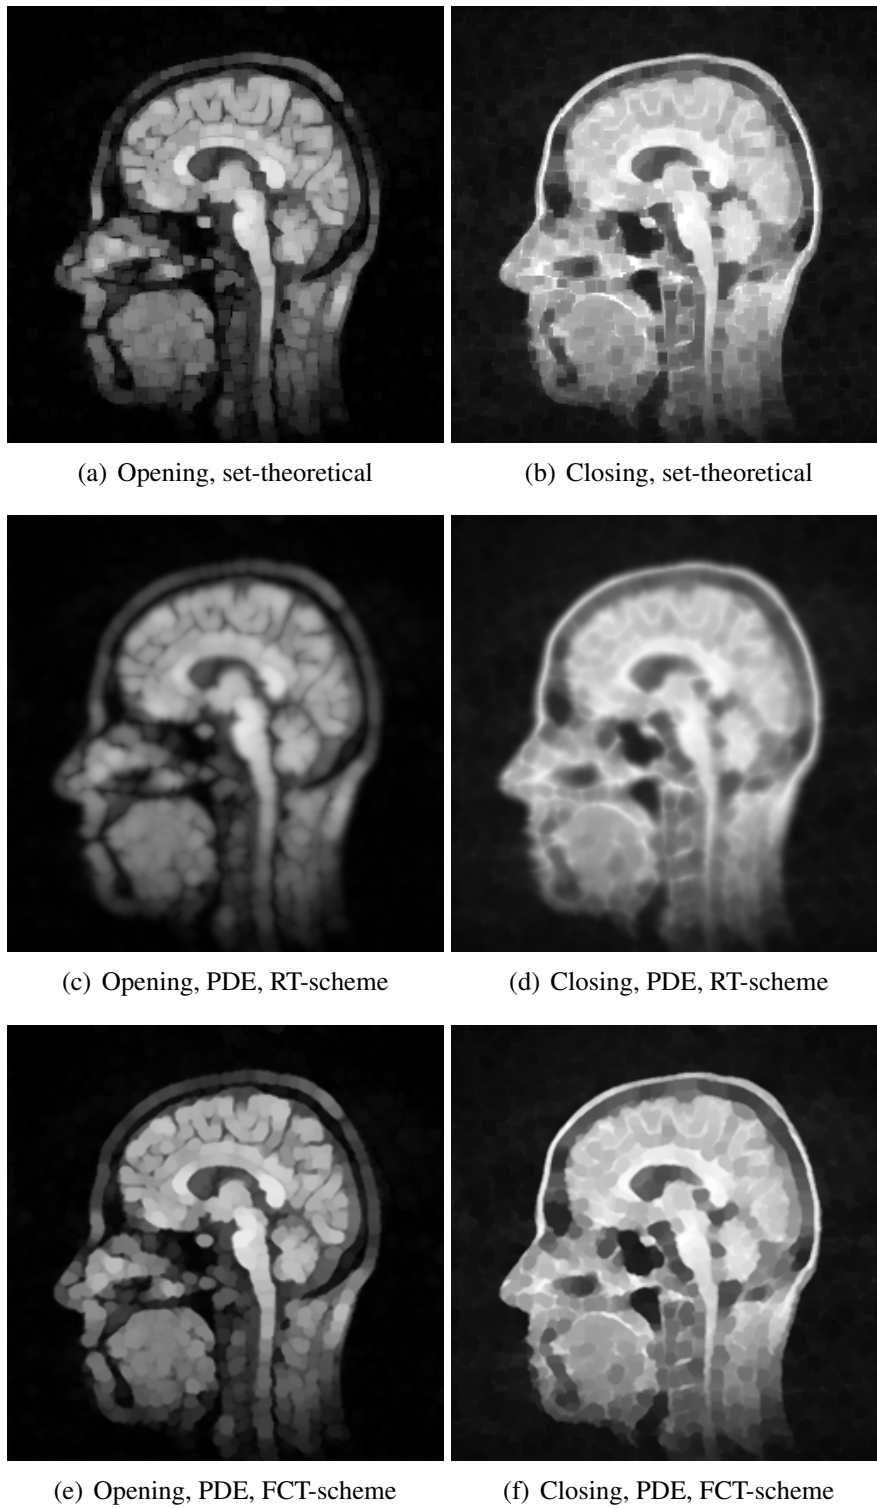





Figure 5.3: Scalar-valued morphology: Set-theoretical vs PDE-based opening/closing with a disc of radius 3 as structuring element. Original image shown in Fig. 5.1.




(a) Original matrix field




(b) Dilation, RT-scheme



(c) Erosion, RT-scheme




(d) Dilation, FCT-scheme




(e) Erosion, FCT-scheme


Figure 5.4: Matrix-valued morphology. PDE-based dilation and erosion with a dis of radius 1.5 as structuring element. FCT-scheme successfully corrects the dissipative artifacts produced by the RT-scheme.




(a) Original matrix field




(b) Opening, RT-scheme



(c) Closing, RT-scheme



(d) Opening, FCT-scheme



(e) Closing, FCT-scheme

Figure 5.5: Matrix-valued morphology. PDE-based opening and closing with a dis of radius 1.5 as structuring element. FCT-scheme successfully corrects the dissipative artifacts produced by the RT-scheme.

6

Adaptive Continuous-Scale Morphology

The PDE-based approach is conceptually attractive since it allows for digital scalability and even adaptivity of the represented structuring element (SE). The current chapter, based on [51, 170, 44], exploits this versatility to propose an anisotropic approach to continuous-scale morphology that adapts itself to the local spatial structure. The concept of *adaptivity* in mathematical morphology is not new. Spatially-variant SEs were first proposed in [189] in the set-theoretical framework. For a recent account of the advances in this context we refer to [27, 28] and the references therein. In general, research on adaptive morphology follows three trends [138]: (i) adaptivity w.r.t. the spatial neighbourhood of morphological operators, (ii) adaptivity w.r.t. how the operators process the different image level sets, and (iii) algebraic principles such as group and representation theory to unify concepts for adaptive operators.

In the continuous-scale framework, it can be shown [181, 54, 95] that the dilation/erosion PDEs are related to the PDE describing the geometric deformation of a curve $\mathcal{C}(p, t) : [0, 2\pi] \times [0, \infty+) \rightarrow \mathbb{R}^2$, p the curve parameterisation and t the evolution parameter:

$$\partial_t \mathcal{C} = \beta \cdot \mathcal{N}, \quad (6.1)$$

with $\mathcal{C}_0(p)$ as initial condition. The curve is propagated in outer/inner normal direction \mathcal{N} with speed β . Representing the curve $\mathcal{C}(p, t)$ as the level set of an image $u(x, t)$ transforms (6.1) into the image evolution PDE

$$\partial_t u = \beta \cdot \|\nabla u\|. \quad (6.2)$$

The simplest case $\beta = \pm 1$ corresponds to the dilation/erosion PDEs (5.7)–(5.8). The evolution (6.2) can as well be adapted to the underlying image structure. For instance, letting the speed be a function of the image curvature $\beta := (\text{curv}(u))^\alpha$ leads to curvature-based morphological processes. When $\alpha = 1$ and $\alpha = 1/3$ these processes are called *mean curvature motion* (MCM) and *affine morphological scale-space* (AMSS), respectively. Other adaptations of the speed function in (6.2) are:

- $\beta := -\text{sign}(\Delta u)$ leads to *shock filtering*, i.e. shocks are created at edge locations where the Laplacian (Δu) vanishes [161];
- $\beta := \pm(f_* - u)$ leads to *viscous*, intensity-level adaptive dilation/erosion, with f_* a chosen (e.g. maximum/minimum) grey value from the input image f [137];
- $\beta := \pm|\cos \theta|^\gamma$ leads to *anisotropic* dilation/erosion, with $\gamma \in \mathbb{Z}^+$ and θ the angle between the image gradient and the direction of minimum contrast [32].

6.1 PDE-Driven Adaptive Morphology

We describe our approach to adaptive morphology starting from a scalar formulation for d -dimensional data $u : \Omega \subset \mathbb{R}^d \rightarrow \mathbb{R}$ in form of the dilation PDE

$$\partial_t u = \|M(u) \cdot \nabla u\| \quad (6.3)$$

with a data dependent, symmetric, positive semidefinite $d \times d$ -matrix $M = M(u)$. Equation (6.3) describes a dilation with an ellipsoidal structuring element since an application of the mapping $(x, y, z)^\top \mapsto M \cdot (x, y, z)^\top$ transforms a sphere centred around the origin into an ellipsoid. In contrast to [32] our steering matrix M allows us to control the amount of dilation/erosion along different directions. When considering three-dimensional matrix-fields, for example, in DT-MRI data sets ($d = 3$) one has

$$M = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} = \begin{pmatrix} \|(a_{11}, a_{12}, a_{13})\| \nu^\top \\ \|(a_{21}, a_{22}, a_{23})\| \mu^\top \\ \|(a_{31}, a_{32}, a_{33})\| \eta^\top \end{pmatrix} \quad (6.4)$$

with unit vectors ν, μ , and η where, e.g.

$$\nu = \frac{1}{\|(a_{11}, a_{12}, a_{13})\|} \begin{pmatrix} a_{11} \\ a_{12} \\ a_{13} \end{pmatrix}. \quad (6.5)$$

This turns (6.3) into

$$\begin{aligned} \partial_t u &= \left((a_{11}\partial_x u + a_{12}\partial_y u + a_{13}\partial_z u)^2 \right. \\ &\quad + (a_{21}\partial_x u + a_{22}\partial_y u + a_{23}\partial_z u)^2 \\ &\quad \left. + (a_{31}\partial_x u + a_{32}\partial_y u + a_{33}\partial_z u)^2 \right)^{\frac{1}{2}} \end{aligned} \quad (6.6)$$

$$\begin{aligned} &= \left(\|(a_{11}, a_{12}, a_{13})\|^2 (\partial_\nu u)^2 \right. \\ &\quad + \|(a_{21}, a_{22}, a_{23})\|^2 (\partial_\mu u)^2 \\ &\quad \left. + \|(a_{31}, a_{32}, a_{33})\|^2 (\partial_\eta u)^2 \right)^{\frac{1}{2}} \end{aligned} \quad (6.7)$$

In [44] the partial derivatives $\partial_x u$, $\partial_y u$, and $\partial_z u$ in (6.6) were approximated with the standard Rouy-Tourin scheme [178] in its two-dimensional version to obtain a directional derivative. However, in [170] the directional derivatives necessary for the steering process were realised directly by means of equation (6.7) with better results than in [44]. Hence it is decisive for our approach to implement the directional derivatives $\partial_\nu u$, $\partial_\mu u$, and $\partial_\eta u$ in (6.7) via a directional version of the Rouy-Tourin

scheme as a suitable upwind numerical scheme to solve the transport equation (6.7). As it will be explained in Section 6.4 an important feature of the proposed approach is the fact that the upwind schemes are employed only in their simplest one-dimensional variant regardless of the dimensionality of the data set. No specially designed higher-dimensional versions or operator splitting methods have to be engaged.

This opens the path for using a high resolution method such as the flux-corrected transport (FCT) scheme of [33] for which its 3D-version is not easily obtained in adaptive form in the setting of matrix fields. In total the novel features over [44] and [170] are the realisation of higher morphological operators based on an adaptive directional version of the FCT scheme in three spatial dimensions.

The necessary directional information of the evolving u contained in the matrix $M(u)$ may be derived from the so-called *structure tensor*, dating back to [83, 22]. It allows to extract directional information from an image and is given by

$$S_\rho(u(x)) := G_\rho * (\nabla u(x) \cdot (\nabla u(x))^\top) \quad (6.8)$$

$$= (G_\rho * (\partial_{x_i} u(x) \cdot \partial_{x_j} u(x)))_{i,j=1,\dots,d} \quad (6.9)$$

Here $G_\rho*$ indicates a convolution with a Gaussian of standard deviation ρ . For more details the reader is referred to [21] and the literature cited therein. In [39, 80] Di Zenzo's approach [72] to construct a structure tensor for multi-channel images has been extended to matrix fields yielding a *standard structure tensor*

$$J_\rho(U(x)) := \sum_{p,q=1}^m S_\rho(U_{p,q}(x)) \quad (6.10)$$

with $U = (U_{p,q})_{p,q=1,\dots,m} \in \text{Sym}_m(\mathbb{R})$, and $\text{Sym}_m(\mathbb{R})$ denotes the set of symmetric $m \times m$ -matrices with real entries. This tensor is a special case of the *full structure tensor concept* for matrix fields [49] which will be reviewed in Section 6.2.

Matrix-valued PDE. After presenting our PDE-based approach to anisotropic morphology in the scalar-valued setting, we now formulate the matrix-valued counterpart of (6.3). Building on the same principles described in Section 5.2.1, we employ the operator-algebraic framework developed by Burgeth *et al.* [48, 47] (see Section 1.4) to proposed the following matrix-valued PDE for adaptive, anisotropic dilation

$$\bar{\partial}_t U = |\bar{M}(U) \bullet \bar{\nabla} U|_2, \quad (6.11)$$

with an initial matrix field $F(x) = U(x, 0)$. Here $\bar{M}(U)$ denotes a symmetric $md \times md$ -block matrix with d^2 blocks of size $m \times m$ that is multiplied block-wise with $\bar{\nabla} U$ employing the Jordan product “ \bullet ”. Note that $|\cdot|_2$ stands for the length of $\bar{M}(U) \bullet \bar{\nabla} U$ in the matrix valued sense, which makes the right-hand side and thus the solution of (6.11) positive semidefinite. An alternative to the Jordan product is the *logarithmic product* [6], but it is computationally more expensive and only defined for positive definite matrices, cf. Section 1.4. The construction of $\bar{M}(U)$ is detailed in Section 6.3 and it relies on the full structure tensor described in the next section.

6.2 Full Structure Tensor for Matrix Fields

The full structure tensor $\bar{\mathcal{S}}_L$ for matrix fields as defined in [49] is used to extract directional information and it plays a vital role in the steering of evolution processes (6.11). The full structure tensor, as a generalisation of (6.8)–(6.9), reads

$$\bar{\mathcal{S}}_L(U) := G_\rho * (\bar{\nabla}U \cdot (\bar{\nabla}U)^\top) \quad (6.12)$$

$$= (G_\rho * (\bar{\partial}_{x_i}U \cdot \bar{\partial}_{x_j}U))_{i,j=1,\dots,d}, \quad (6.13)$$

with $G_\rho*$ indicating a convolution with a Gaussian of standard deviation ρ . $\bar{\mathcal{S}}_L(U(x))$ is a symmetric $md \times md$ -block matrix with d^2 blocks of size $m \times m$, $\bar{\mathcal{S}}_L(U(x)) \in \text{Sym}_d(\text{Sym}_m(\mathbb{R})) = \text{Sym}_{md}(\mathbb{R})$. Typically for the 3D medical DT-MRI data one has $d = 3$ and $m = 3$, yielding a 9×9 -matrix $\bar{\mathcal{S}}_L$. It can be diagonalised as $\bar{\mathcal{S}}_L(U) = \sum_{k=1}^{md} \lambda_k w_k w_k^\top$ with real eigenvalues λ_k (w.l.o.g. arranged in decreasing order) and an orthonormal basis $\{w_k\}_{k=1,\dots,md}$ of \mathbb{R}^{md} .

In order to extract useful d -dimensional directional information, $\bar{\mathcal{S}}_L(U) \in \text{Sym}_{md}(\mathbb{R})$ is reduced to a structure tensor $S(U) \in \text{Sym}_d(\mathbb{R})$ in a generalised projection step [49] using the block operator matrix $\text{Tr}_A := \text{diag}(\text{tr}_A, \dots, \text{tr}_A)$ containing the trace operation. We set $\text{Tr} := \text{Tr}_{I_m}$ where I_m denotes the $m \times m$ unit matrix. This operator matrix acts on elements of the space $(\text{Sym}_m(\mathbb{R}))^d$ as well as on block matrices via formal block-wise matrix multiplication

$$\begin{pmatrix} \text{tr}_A & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \text{tr}_A \end{pmatrix} \begin{pmatrix} M_{11} & \cdots & M_{1d} \\ \vdots & \ddots & \vdots \\ M_{d1} & \cdots & M_{dd} \end{pmatrix} = \begin{pmatrix} \text{tr}_A(M_{11}) & \cdots & \text{tr}_A(M_{1d}) \\ \vdots & \ddots & \vdots \\ \text{tr}_A(M_{d1}) & \cdots & \text{tr}_A(M_{dd}) \end{pmatrix}, \quad (6.14)$$

provided that the square blocks M_{ij} have the same size as A . The projection that is conveyed by the reduction process condenses the directional information contained in $\bar{\mathcal{S}}_L(U)$, for a more detailed reasoning we refer the reader to [49]. The reduction operation is accompanied by an extension operation: The I_m -extension is the mapping from $\text{Sym}_d(\mathbb{R})$ to $\text{Sym}_{md}(\mathbb{R})$ conveyed by the *Kronecker product* \otimes :

$$\begin{pmatrix} v_{11} & \cdots & v_{1d} \\ \vdots & \ddots & \vdots \\ v_{d1} & \cdots & v_{dd} \end{pmatrix} \mapsto \begin{pmatrix} v_{11} & \cdots & v_{1d} \\ \vdots & \ddots & \vdots \\ v_{d1} & \cdots & v_{dd} \end{pmatrix} \otimes I_m := \begin{pmatrix} v_{11}I_m & \cdots & v_{1d}I_m \\ \vdots & \ddots & \vdots \\ v_{d1}I_m & \cdots & v_{dd}I_m \end{pmatrix}.$$

This resizing step renders a proper matrix-vector multiplication with the large generalised gradient $(\bar{\nabla}U(x))$ possible. By specifying the matrix A in (6.14) one may invoke a priori knowledge into the direction estimation [49]. The research

on these structure-tensor concepts has been initiated by [223, 39]. The approaches to matrix field regularisation suggested in [62] are based on differential geometric considerations. Comprehensive survey articles on the analysis of matrix fields using various techniques can be found in [225].

6.3 Steering Matrix for Matrix Fields

We are now in the position to propose the steering matrix \overline{M} for the adaptive dilation process for matrix fields (6.11). We proceed in four steps:

- (i) The matrix field $\mathbb{R}^d \ni x \mapsto U(x)$ provides us with a module field of generalised gradients $\overline{\nabla}U(x)$ from which we construct the generalised structure tensor $\overline{\mathcal{S}}_L(U(x))$ possibly with a certain integration scale ρ . This step corresponds exactly to the scalar case.
- (ii) We infer d -dimensional directional information by reducing $\overline{\mathcal{S}}_L(U(x))$ with tr_A with the help of the block operator matrix given in (6.14). This leads to a symmetric $d \times d$ -matrix S , for example $S = J_\rho$ if $A = I_m$:

$$S(x) := \text{Tr}_A(\overline{\mathcal{S}}_L(U(x))). \quad (6.15)$$

- (iii) The symmetric $d \times d$ -matrix S is spectrally decomposed, and the following mapping is applied:

$$H : \begin{cases} \mathbb{R}_+^d & \longrightarrow \mathbb{R}^d \\ (\lambda_1, \dots, \lambda_d) & \longmapsto \frac{1}{\sum_{i=1}^d \lambda_i} (c_1 \lambda_1, \dots, c_d \lambda_d) \end{cases} \quad (6.16)$$

where $c = (c_1, \dots, c_d)$ is a vector with nonnegative entries. With the choice of the vector c we select the eigendirection in which the process is steered. For instance, specifying $c_1 = \dots = c_{d-1} = k$ and $c_d = K \gg k$ one obtains an ellipsoid associated with the matrix M which is flipped if compared with S . Depending on the choice of K it can be more excentric than the one accompanying S . H applied to S yields the steering matrix M ,

$$M := H(S). \quad (6.17)$$

- (iv) Finally we enlarge the $d \times d$ -matrix M to a $md \times md$ -matrix \overline{M} by the extension operation

$$\overline{M} = M \otimes I_m. \quad (6.18)$$

The steering matrix \overline{M} as defined above completes our description of the proposed adaptive, anisotropic dilation process (6.11). What remains to be characterised is the numerical schemes we employ to discretise and solve this PDE. This is the subject of the next section.

6.4 Directional Numerical Schemes

We now outline the numerical schemes we use to solve the proposed nonlinear PDEs that describe the proposed scalar- and matrix-valued anisotropic dilation, respectively. The numerics for the corresponding erosion processes follow analogously.

As discussed earlier, the PDE (6.3) can be rewritten as in (6.7) where it becomes evident that the steering matrix M introduces directional components into the dilation process. Therefore, in the following two subsections we sketch the *directional* versions of the Rouy-Tourin scheme and the FCT scheme and their extensions to the matrix-valued setting to solve (6.3) and (6.11).

6.4.1 Directional Rouy-Tourin scheme

The first-order finite difference method of *Rouy and Tourin* [178] (cf. Section 5.1.4.1) may be used to solve the scalar PDE (6.7) in the *isotropic* case with $M = I_d$. However, in the anisotropic case we need to introduce a few modification to this scheme. Let us denote by h_x, h_y, h_z the spatial grid size and by $u_{i,j,k}^n$ the grey value of a scalar 3D image data set u at the pixel centred in $(ih_x, jh_y, kh_z) \in \mathbb{R}^3$ at the time-level $t = n\tau$. In particular, we need to formulate directional adaptations of the standard forward and backward difference operators

$$D_+^x u_{i,j,k}^n := u_{i+1,j,k}^n - u_{i,j,k}^n, \quad (6.19)$$

$$D_-^x u_{i,j,k}^n := u_{i,j,k}^n - u_{i-1,j,k}^n, \quad (6.20)$$

as well as of the *upwind approximation* of the partial derivative u_x

$$u_x \approx \frac{1}{h_x} \max \left(\max \left(-D_-^x u_{i,j,k}^n, 0 \right), \max \left(D_+^x u_{i,j,k}^n, 0 \right) \right), \quad (6.21)$$

here in x -, but analogously in y - and z -direction.

For a unit vector $\nu = (\nu_1, \nu_2, \nu_3)^\top$ the directional derivative $\partial_\nu u$ of u may be approximated by $\partial_\nu u = \langle \nu, \nabla u \rangle = \nu_1 \partial_x u + \nu_2 \partial_y u + \nu_3 \partial_z u$. Hence it is close at hand to approximate numerically equation (6.6) directly. However, this favours mass transport along the directions of the x -, y -, and z -axis leading to a poor representation of the directional derivative. Instead we take advantage of equation (6.7) and propose an alternative involving an interpolated function value $u_{i+\nu_1, j+\nu_2, k+\nu_3}$ defined by the subsequent well-known tri-linear¹ approximation (6.23) employing as weights the volumes of cuboids associated with the grid points (Figure 6.1 shows the bi-linear case in 2D):

1. For the sake of efficiency we use tri-linear interpolation, although higher order alternatives such as tri-cubic or spline interpolation can be employed as well.

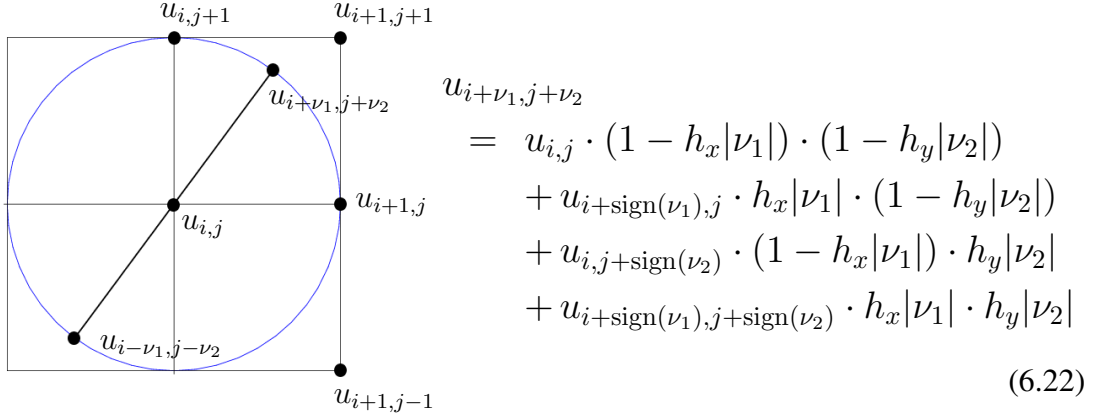


Figure 6.1: 2D interpolation of the image value $u_{i+\nu_1, j+\nu_2}$ with $\sqrt{\nu_1^2 + \nu_2^2} = 1$. It allows for backward and forward finite differences in the direction of $(\nu_1, \nu_2)^\top$.

$$\begin{aligned}
u_{i+\nu_1, j+\nu_2, k+\nu_3} &= u_{i,j,k} \cdot (1 - h_x |\nu_1|) \cdot (1 - h_y |\nu_2|) \cdot (1 - h_z |\nu_3|) \\
&\quad + u_{i+\text{sign}(\nu_1), j, k} \cdot h_x |\nu_1| \cdot (1 - h_y |\nu_2|) \cdot (1 - h_z |\nu_3|) \\
&\quad + u_{i, j+\text{sign}(\nu_2), k} \cdot (1 - h_x |\nu_1|) \cdot h_y |\nu_2| \cdot (1 - h_z |\nu_3|) \\
&\quad + u_{i+\text{sign}(\nu_1), j+\text{sign}(\nu_2), k} \cdot h_x |\nu_1| \cdot h_y |\nu_2| \cdot (1 - h_z |\nu_3|) \\
&\quad + u_{i, j, k+\text{sign}(\nu_3)} \cdot (1 - h_x |\nu_1|) \cdot (1 - h_y |\nu_2|) \cdot h_z |\nu_3| \\
&\quad + u_{i+\text{sign}(\nu_1), j, k+\text{sign}(\nu_3)} \cdot h_x |\nu_1| \cdot (1 - h_y |\nu_2|) \cdot h_z |\nu_3| \\
&\quad + u_{i, j+\text{sign}(\nu_2), k+\text{sign}(\nu_3)} \cdot (1 - h_x |\nu_1|) \cdot h_y |\nu_2| \cdot h_z |\nu_3| \\
&\quad + u_{i+\text{sign}(\nu_1), j+\text{sign}(\nu_2), k+\text{sign}(\nu_3)} \cdot h_x |\nu_1| \cdot h_y |\nu_2| \cdot h_z |\nu_3|.
\end{aligned} \tag{6.23}$$

This leads to forward and backward difference operators in the direction of ν with $\|\nu\| = \sqrt{\nu_1^2 + \nu_2^2 + \nu_3^2} = 1$:

$$D_+^\nu u_{i,j,k}^n := u_{i+\nu_1, j+\nu_2, k+\nu_3}^n - u_{i,j,k}^n \tag{6.24}$$

$$D_-^\nu u_{i,j,k}^n := u_{i,j,k}^n - u_{i-\nu_1, j-\nu_2, k-\nu_3}^n \tag{6.25}$$

and to a direct approximation of the directional derivative

$$\begin{aligned}
\partial_\nu u &= u_\nu \\
&\approx \frac{1}{h} \max \left(\max \left(-D_-^\nu u_{i,j,k}^n, 0 \right), \max \left(D_+^\nu u_{i,j,k}^n, 0 \right) \right)
\end{aligned} \tag{6.26}$$

where $h := \min(h_x, h_y, h_z)$. Furthermore, the resulting approximation of the directional derivatives is also *consistent*: tri-linear approximation implies

$$\begin{aligned}
u_{i+\nu_1, j+\nu_2, k+\nu_3} &= u((i + \nu_1)h_x, (j + \nu_2)h_y, (k + \nu_3)h_z) \\
&\quad + \mathcal{O}(\max(h_x, h_y, h_z)),
\end{aligned} \tag{6.27}$$

and hence we have

$$\begin{aligned} & \frac{1}{h} D_+^\nu u_{i,j,k} \\ &= \frac{1}{h} (u((i + \nu_1)h_x, (j + \nu_2)h_y, (k + \nu_3)h_z) - u(ih_x, jh_y, kh_z)) \\ & \quad + \mathcal{O}(\max(h_x, h_y, h_z)) \end{aligned} \quad (6.28)$$

$$= u_\nu + \mathcal{O}(\max(h_x, h_y, h_z)). \quad (6.29)$$

Analogous reasoning applies to $D_-^\nu u_{i,j,k}$.

With the notation introduced in Section 5.2.2 and the calculus concept presented in Section 1.4 it is now straightforward to define directional one-sided differences in ν -direction for fields of $m \times m$ -matrices:

$$D_+^\nu U^n(ih_x, jh_y, kh_z) \quad (6.30)$$

$$:= U^n((i + \nu_1)h_x, (j + \nu_2)h_y, (k + \nu_3)h_z) - U^n(ih_x, jh_y, kh_z),$$

$$D_-^\nu U^n(ih_x, jh_y, kh_z) \quad (6.31)$$

$$:= U^n(ih_x, jh_y, kh_z) - U^n((i - \nu_1)h_x, (j - \nu_2)h_y, (k - \nu_3)h_z),$$

where $D_+^\nu U^n, D_-^\nu U^n \in \text{Sym}_m(\mathbb{R})$. The directions μ and η are treated accordingly. The notion of supremum and infimum of two matrices – as needed in a matrix variant of Rouy-Tourin – has been provided in Section 1.4 as well. Hence, having these generalisations at our disposal a directionally adaptive version of the Rouy-Tourin scheme is available now in the setting of matrix fields simply by replacing grey values $u_{i,j,k}^n$ by matrices $U^n(ih_x, jh_y, kh_z)$ and utilising the directional derivative approximations.

The numerical scheme for the PDE describing the erosion process can be obtained analogously by considering the directional version of the Rouy-Tourin approximation described in Subsection 5.1.4.1 in the previous chapter.

6.4.2 Directional FCT scheme

Here we provide the directional modifications of the FCT method (cf. Section 5.1.4.2) to solve (6.3) and (6.11) in the three-dimensional case. Let us start with the scalar setting. Using the solution $u_{i,j}^{n+1,\text{pred}}$ obtained with the directional Rouy-Tourin scheme as a *predictor*, the FCT method relies on a *corrector step*, which will finally read as

$$u_{i,j,k}^{n+1} = u_{i,j,k}^{n+1,\text{pred}} + q_h^{n+1,\text{pred}} - q_d^{n+1,\text{pred}}. \quad (6.32)$$

The directional version of the correction step (6.32) is now obtained by replacing the finite differences $D_{(\cdot)}^x$ in x -direction in the terms defining $q_h^{n+1,\text{pred}}$ and $q_d^{n+1,\text{pred}}$ by the weighted finite differences $\|\hat{\nu}\| D_{(\cdot)}^\nu$ in ν -direction with $\hat{\nu} = (a_{11}, a_{12}, a_{13})$, see (6.4).

We proceed in the same way with the other directions, substituting $D_{(\cdot)}^y$ by $\|\hat{\mu}\| D_{(\cdot)}^\mu$, and in the three-dimensional case, exchanging $D_{(\cdot)}^z$ by $\|\hat{\eta}\| D_{(\cdot)}^\eta$ with corresponding interpretations of $\hat{\mu}$ and $\hat{\eta}$. Doing so, the higher-order term $q_h^{n+1,\text{pred}}$ in (6.32) reads

$$q_h^{n+1,\text{pred}} := \frac{\tau}{h} \left(\left(\|\hat{\nu}\| \left| D_c^\nu u_{i,j,k}^{n+1,\text{pred}} \right| \right)^2 + \left(\|\hat{\mu}\| \left| D_c^\mu u_{i,j,k}^{n+1,\text{pred}} \right| \right)^2 + \left(\|\hat{\eta}\| \left| D_c^\eta u_{i,j,k}^{n+1,\text{pred}} \right| \right)^2 \right)^{1/2}, \quad (6.33)$$

with $h := \min(h_x, h_y, h_z)$ and the directional central differences

$$D_c^\nu u_{i,j,k}^n := \left(u_{(i,j,k)+\nu/2}^n - u_{(i,j,k)-\nu/2}^n \right) / 2. \quad (6.34)$$

Consequently, the lower-order term in (6.32) is given by

$$q_d^{n+1,\text{pred}} := \left(\left(\delta^\nu u_{i,j,k}^{n+1,\text{pred}} \right)^2 + \left(\delta^\mu u_{i,j,k}^{n+1,\text{pred}} \right)^2 + \left(\delta^\eta u_{i,j,k}^{n+1,\text{pred}} \right)^2 \right)^{1/2}, \quad (6.35)$$

using the abbreviations

$$\delta^\nu u_{i,j,k}^{n+1,\text{pred}} := \frac{\tau}{h} \|\hat{\nu}\| \left| D_c^\nu u_{i,j,k}^{n+1,\text{pred}} \right| + g_{(i,j,k)+\nu/2} - g_{(i,j,k)-\nu/2}, \quad (6.36)$$

$$g_{(i,j,k)+\nu/2} := \|\hat{\nu}\| \text{mm} \left(D_-^\nu u_{(i,j,k)}^{n+1,\text{pred}}, \frac{\tau}{2h} D_+^\nu u_{(i,j,k)}^{n+1,\text{pred}}, D_+^\nu u_{(i,j,k)+\nu}^{n+1,\text{pred}} \right), \quad (6.37)$$

$$g_{(i,j,k)-\nu/2} := \|\hat{\nu}\| \text{mm} \left(D_-^\nu u_{(i,j,k)-\nu}^{n+1,\text{pred}}, \frac{\tau}{2h} D_+^\nu u_{(i,j,k)-\nu}^{n+1,\text{pred}}, D_+^\nu u_{(i,j,k)}^{n+1,\text{pred}} \right), \quad (6.38)$$

and the minmod-function defined as in (5.20). The directions μ and η are treated analogously.

The directional version of the FCT-scheme for matrix fields is obtained by replacing the scalar values $u_{i,j,k}^n$ by matrices $U^n(ih_x, jh_y, kh_z)$, utilising the matrix-valued directional derivative approximations, and the definition of the minmod-function as in (5.34). We emphasise that the extension to tensor fields follows from the operator-algebraic framework of Burgeth *et al.* [48, 47] outlined in Section 1.4, see also Subsection 5.2.2.

Finally, the numerical scheme for the erosion PDE is obtained analogously by considering the directional version of the FCT approximation described in Subsection 5.1.4.2 in the preceding chapter.

6.5 Structure-Adaptive Morphological Operations

As indicated in Section 6.1, the solution u of equation (6.3) mimics the dilation process with an adaptive ellipsoidal structuring element B which changes according to the spatial local structure as well as in time since it depends on u : $B = B_u$. A minus sign on the right-hand-side of (6.3) and (6.11) gives the PDE-formulation of the corresponding adaptive erosion process. Using a common notation we express the dilation and the erosion of an original image f with such a structuring element B_u by

$$f \oplus B_u \quad \text{and} \quad f \ominus B_u. \quad (6.39)$$

The combinations of dilation and erosion lead to various morphological operators such as *opening* and *closing*,

$$f \circ B_u := (f \ominus B_u) \oplus B_u, \quad (6.40)$$

$$f \bullet B_u := (f \oplus B_u) \ominus B_u. \quad (6.41)$$

In an image, boundaries of objects are loci of high grey value variations, and as such they can be detected by derivative operators. The so-called *Beucher gradient*

$$\varrho_{B_u}(f) := (f \oplus B_u) - (f \ominus B_u), \quad (6.42)$$

as well as the *internal* and *external gradient*,

$$\varrho_{B_u}^-(f) := f - (f \ominus B_u), \quad \varrho_{B_u}^+(f) := (f \oplus B_u) - f, \quad (6.43)$$

are morphological counterparts of the norm of the gradient f , $\|\nabla f\|$, if f is considered as a differentiable image.

In [215] a *morphological Laplacian* has been introduced. We define a variant by

$$\Delta_{B_u} f := \varrho_{B_u}^+(f) - \varrho_{B_u}^-(f) \quad (6.44)$$

$$= (f \oplus B_u) - 2 \cdot f + (f \ominus B_u). \quad (6.45)$$

This Laplacian is a morphological equivalent of the second derivative $\partial_{\eta\eta} f$ where η stands for the unit vector in the direction of the steepest slope. It allows us to distinguish between influence zones of minima and maxima of the image f . This is a vital property for the construction of so-called shock filters [94, 121, 161].

Shock filtering amounts to applying either a dilation or an erosion to an image, depending on whether the pixel is located within the influence zone of a minimum or a maximum:

$$S_{B_u} f := \begin{cases} f \oplus B_u, & \Delta_{B_u} f < 0, \\ f, & \Delta_{B_u} f = 0, \\ f \ominus B_u, & \Delta_{B_u} f > 0. \end{cases} \quad (6.46)$$

A considerable number of variants of shock filters have been considered in the literature [3, 89, 160, 176, 183, 222]. When they are applied iteratively, experiments show that their steady state is given by a piecewise constant image with discontinuities (“shocks”) between adjacent segments of constant grey value. For more details about the morphological shock filter as introduced above, see [46].

In the experimental Section 6.6 we will see the results obtained by the various adaptive, PDE-driven morphological operators when applied to synthetic and real-world scalar- and matrix-valued images.

6.6 Experiments

6.6.1 Scalar-Valued Images

We first consider scalar-valued images to test our approach to adaptive, anisotropic continuous-scale morphology (6.3) implemented with both the classical (non-directional) and the proposed directional Rouy-Tourin scheme. We compare our method to the standard case of isotropic morphology ($M = I_d$) and to the coherence-enhancing diffusion (CED) approach of Weickert [221] as another PDE-based method useful for completing and enhancing flow-like structures. We juxtapose all these methods in the following figures: Figure 6.2 is a binary image of a ring with missing areas. Figure 6.3 shows a fingerprint with numerous interrupted lines. In both cases we are interested in reconstructing the missing parts. Figure 6.4 exhibits the ‘Selfportrait’ of van Gogh where coherent flow-like structures can be distinguished. Interestingly, we shall see that the proposed directional approach leads to an anisotropic ‘enhancement’ of these structures.

It is noticeable in all three experiments that standard isotropic morphology does not provide the expected results, broadening/shrinking the structures in all directions. On the other side, our adaptive directional approach successfully completes and enhances the structures. It is worth to emphasise the importance of working with a directional numerical schemes as standard, non-directional schemes suffer from numerical bias to the coordinate axes, whose consequences result evidently disastrous in all our experiments. Regarding the CED approach, it is well-known the suitability of this method for denoising and enhancement of flow-like structures. As observed in the fingerprint and in the ‘Selfportrait’ our adaptive approach delivers similar results, although noise is as well enhanced by any morphological method. Note also that CED takes considerably more time than our approach to complete/enhance the underlying structures.

6.6.2 Synthetic 2D and 3D Matrix Fields

We now test our matrix-valued approach to adaptive, anisotropic continuous-scale morphology (6.11) and consider the matrix-valued extensions of both the Rouy-Tourin and the FCT numerical schemes.

Figure 6.5(a) exhibits a 32×32 matrix field composed of two interrupted diagonal stripes with different thickness, both built with cigar-shaped ellipsoids of equal size but different orientation. The line-like structures are tilted with respect to the x -axis by an angle of about 117 degrees. Figure 6.5(b) shows the result of applying coherence-enhancing diffusion (CED) [49]. Figure 6.5(c) contains the result of isotropic (classical) dilation [45] using the standard (non-directional) Rouy-Tourin scheme. Figure 6.5(d) and Figure 6.5(e) show the result of the proposed adaptive anisotropic dilation using the Rouy-Tourin scheme in its standard and directional versions, respectively. The parameters used were chosen in a way that every method fills in the missing tensors of both stripes. Our approach is able to complete the line-like structures much faster and more accurate than the other methods. It is noticeable the improvement introduced by the directional numerical scheme, which overcomes the numerical bias to the coordinate axes of classical upwind schemes (e.g. Rouy-Tourin). Moreover, note that the direction and amount of adaptive anisotropic dilation does not depend on the orientation of the ellipsoids, but on the orientation and width of the structures. It is worth mentioning that the CED approach decreases the overall size of matrices since the total mass, that is, the volume of the ellipsoids is only redistributed due to the property of mass conservation. The same experiment is performed on a 2D spiral data set with missing information, whose outcome is depicted in Figure 6.6. Again, only the anisotropic dilation succeeds to close the gaps satisfactorily preserving the spiral structure of the object. As expected, the adaptive dilation process is faster than the diffusion based method and the classical isotropic dilation. Furthermore, the effect of employing a directional numerical scheme becomes more evident when the anisotropic dilation is steered along circular structures.

We now use both the directional Rouy-Tourin scheme and the directional FCT scheme for dilating the test image with an interrupted circular structure shown at the top of Figure 6.7. In the first test the dilation process is steered in tangential direction while in a second test the radial direction is selected via the choice of the parameter $c = (c_1, c_2)$, namely $c = (0.1, 10)$ in the first case and $c = (10, 0.1)$ in the second one. Both directional schemes were applied, the results together with a scaled ($\times 5$) absolute difference image are depicted in Figure 6.7. As expected the directional FCT method performs favourable in terms of edge preservation over the directional Rouy-Tourin scheme. This can be seen in the difference images at the bottom of Figure 6.7, which display the (scaled) dissipation introduced by the Rouy-Tourin scheme that has been corrected by the FCT scheme.

A much more elaborate matrix field can be seen in Figure 6.8(a). This 3D data set² is sparsified by removing 80% of the matrices (Figure 6.8(b)). Both adaptive anisotropic dilation (Figure 6.8(c)) and closing (Figure 6.8(d)) performed with the superior directional FCT scheme, provide a reasonable reconstruction of the original data.

6.6.3 Real World Data: 3D DT-MRI

We also tested the proposed method on a real DT-MRI data set of a human head consisting of a $128 \times 128 \times 38$ -field of positive definite matrices. In the experiments on real-world data we will always juxtapose the results of various adaptive morphological operations when applied to 2D-slices and truly 3D data. However, in order to avoid visual cluttering, we will in general extract and depict from the processed 3D data an appropriate 2D-slice only.

Figure 6.9(a) displays part of the lateral ventricles as an actual three-dimensional $40 \times 55 \times 3$ -data set while Figure 6.9(b) shows only a 2D-slice. Adaptive dilation with the classical Rouy-Tourin scheme [44] and the proposed directional Rouy-Tourin scheme are shown in Figures 6.9(c) and 6.9(d), respectively. It is evident that the classical scheme is affected by numerical artifacts, problem that is overcome by using its directional version.

In Figure 6.10 we zoom into the lateral ventricles to show the effect of applying adaptive dilation and erosion in both the 3D and 2D setting. We see that the adaptive dilation and erosion processes on matrix fields respect the underlying shape of the ventricles if compared to the isotropic case [43]. We notice that the results are quite similar in the 3D and in the 2D setting. However, the 3D process seems to be more accurate at the price of a longer evolution time ($t = 3$), than in the 2D case ($t = 1$).

The lateral ventricles serve also as a test case for the reconstructing operations of adaptive opening and closing, Figure 6.11. In 3D the lateral ventricles are nicely recovered in a slightly simplified form, as expected, since it incorporates also information from neighbouring slices.

The difference in processing of 2D and 3D data sets becomes prominent in the case of the morphological derivatives, e.g. the Beucher gradient. The gradient operations detect the boundary of the ventricles, which are three-dimensional anatomical structures. This boundary in a cross-section can be seen clearly in Figure 6.12(a). In contrast to this, the boundaries are less localised in the 2D case, Figure 6.12(b).

For the adaptive version of the morphological shock filter we obtain the matrix valued equivalent of a piece-wise constant image both in the 3D and the 2D case. In the three-dimensional setting, Figure 6.12(c), we observe a slightly better localisation of the shock segments than in 2D, Figure 6.12(d).

2. The 3D spiral data set is freely available at <http://teem.sourceforge.net>.

6.7 Summary


We have presented a method for an adaptive, PDE-based dilation and erosion processes in the setting of scalar- and matrix-valued images. The evolution governed by matrix-valued PDEs is guided by a steering tensor whose construction relies on the full structure tensor for matrix data.

In order to enable proper directional steering we extended the schemes of Rouy-Tourin and the FCT method in two ways: First, turning them into directional schemes based on directional finite differences via interpolation. Second, by means of advanced matrix calculus, extending these directional variants to matrix fields solving the matrix-valued adaptive PDEs of dilation and erosion. Having these two key operations at our disposal we were able to propose higher order morphological operators such as morphological derivatives that are adaptive and act on matrix fields.

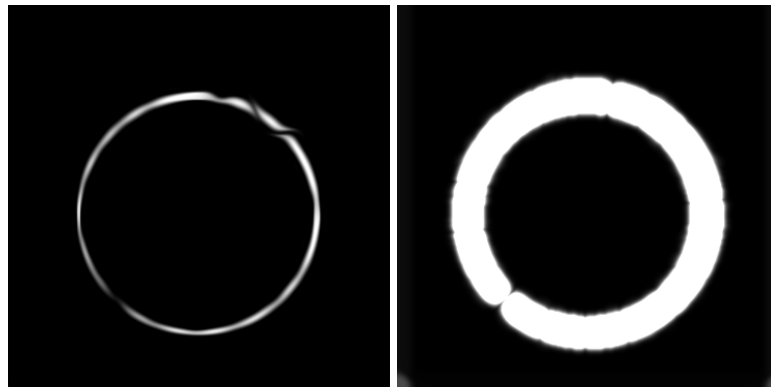
A special advantage of our approach is that the directional numerical schemes are utilised only in their basic one-dimensional version, hence avoiding grid effects and leading to an accurate algorithm to complete or enhance anisotropic structures effectively. No higher dimensional variants of the schemes themselves are required.

As a proof-of-concept we applied these adaptive morphological operations to synthetic and real DT-MRI data. The tests reveal that the various adaptive morphological operators behave as one might expected from their scalar counterparts. For instance, the adaptive dilation and closing are indeed capable of filling in missing data and to complete directional structures. We also confirmed that the FCT performs preferable over the scheme of Rouy and Tourin.

The direct application of an elementary morphological operation is usually not suitable for improving an image or extracting useful information from it. Morphology (adaptive or not) gains its power from the capability of concatenating and combining elementary operations according to the task at hand. We hope to have just opened the adaptive morphological toolbox for matrix fields. In this respect the proposed approach to adaptive morphology for matrix fields may have its merits, for example, in the segmentation of directional structures or as a preprocessing step for fibre tracking algorithms in medical imaging.




(a) Original image



(b) Coherence-enhancing diffusion

(c) Isotropic dilation



(d) Anisotropic dilation

(e) Directional anisotropic dilation

Figure 6.2: Comparison of different methods for binary images. (a) Original image showing a ring with missing areas [216]. (b) Coherence-enhancing diffusion (CED) as in Weickert [221] with $\sigma = 0, \rho = 10, t = 48$. (c) Isotropic (classical) dilation at $t = 12$ using the Rouy-Tourin scheme. (d) and (e) show the proposed adaptive, anisotropic dilation using the classical Rouy-Tourin scheme and the directional Rouy-Tourin scheme, respectively, with $\sigma = 0, \rho = 10, c = (0.05, 5), t = 12$.

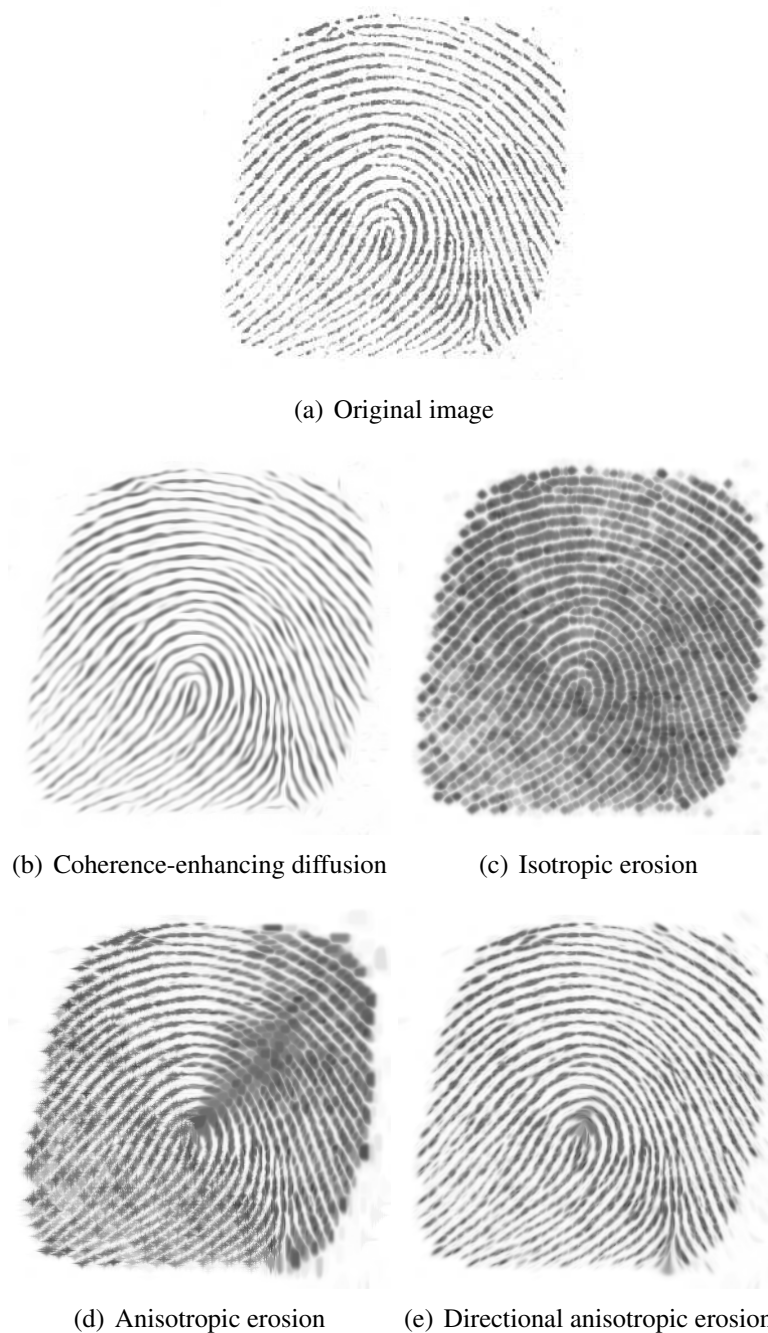



Figure 6.3: Comparison of different methods for greyscale images. (a) Original image showing a fingerprint. (b) Coherence-enhancing diffusion (CED) as in Weickert [221] with $\sigma = 0.5$, $\rho = 10$, $t = 6$. (c) Isotropic (classical) erosion at $t = 3$ using the Rouy-Tourin scheme. (d) and (e) show the proposed adaptive, anisotropic erosion using the classical Rouy-Tourin scheme and the directional Rouy-Tourin scheme, respectively, with $\sigma = 0.5$, $\rho = 10$, $c = (0.1, 10)$, $t = 3$.



(a) Original image




(b) CED



(c) Isotropic erosion




(d) Anisotropic erosion




(e) Dir. anisotropic erosion


Figure 6.4: Comparison of different methods for greyscale images. (a) ‘Selfportrait’ by van Gogh. (b) Coherence-enhancing diffusion (CED) as in Weickert [221] with $\sigma = 0.5, \rho = 5, t = 2$. (c) Isotropic (classical) erosion at $t = 2$ using the Rouy-Tourin scheme. (d) and (e) show the proposed adaptive, anisotropic erosion using the classical Rouy-Tourin scheme and the directional Rouy-Tourin scheme, respectively, with $\sigma = 0.5, \rho = 5, c = (0.1, 10), t = 2$.




(a) Original 2D line-like structures




(b) Coherence-enhancing diffusion



(c) Isotropic dilation




(d) Anisotropic dilation




(e) Directional anisotropic dilation


Figure 6.5: Comparison of different methods in 2D. (a) Original matrix field with ellipsoids in a line-like arrangement. (b) Coherence-enhancing diffusion (CED) as in Burgeth *et al.* [49] with $\sigma = 0, \rho = 4, t = 3$. (c) Isotropic (classical) dilation at $t = 3$ using the Rouy-Tourin scheme. (d) and (e) show the proposed adaptive, anisotropic dilation using the classical Rouy-Tourin scheme and the directional Rouy-Tourin scheme, respectively, with $\sigma = 0, \rho = 4, c = (0.2, 20), t = 1$.




(a) Original 2D spiral




(b) Coherence-enhancing diffusion



(c) Isotropic dilation



(d) Anisotropic dilation



(e) Directional anisotropic dilation

Figure 6.6: Comparison of different methods in 2D. (a) Original spiral with missing tensors. (b) Coherence-enhancing diffusion (CED) as in Burgeth *et al.* [49] with $\sigma = 0, \rho = 3, t = 6$. (c) Isotropic dilation at $t = 3$ using the Rouy-Tourin scheme. (d) and (e) show the proposed adaptive, anisotropic dilation using the classical Rouy-Tourin scheme and the directional Rouy-Tourin scheme, respectively, with $\sigma = 0, \rho = 3, c = (0.2, 20), t = 1$.




Figure 6.7: *Top row*: Synthetic 2D circle with missing information. *Second row*: From left to right: Dilation with directional Rouy-Tourin scheme in the tangential direction and in the radial direction. *Third row*: The same using the directional FCT scheme. *Bottom row*: Scaled absolute differences between both schemes.




Figure 6.8: Adaptive, anisotropic dilation and closing in 3D using the directional FCT scheme with parameters $\sigma = 2$, $\rho = 2$, $c = (0.2, 0.2, 20)$, $t = 5$.




Figure 6.9: Classical versus directional Rouy-Tourin scheme. (a) Original data set. (b) 2D section. (c) and (d) show the proposed adaptive, anisotropic dilation using the classical Rouy-Tourin scheme and the directional Rouy-Tourin scheme, respectively, with $\sigma = 0$, $\rho = 1$, $c = (0.1, 10)$, $t = 1.5$.

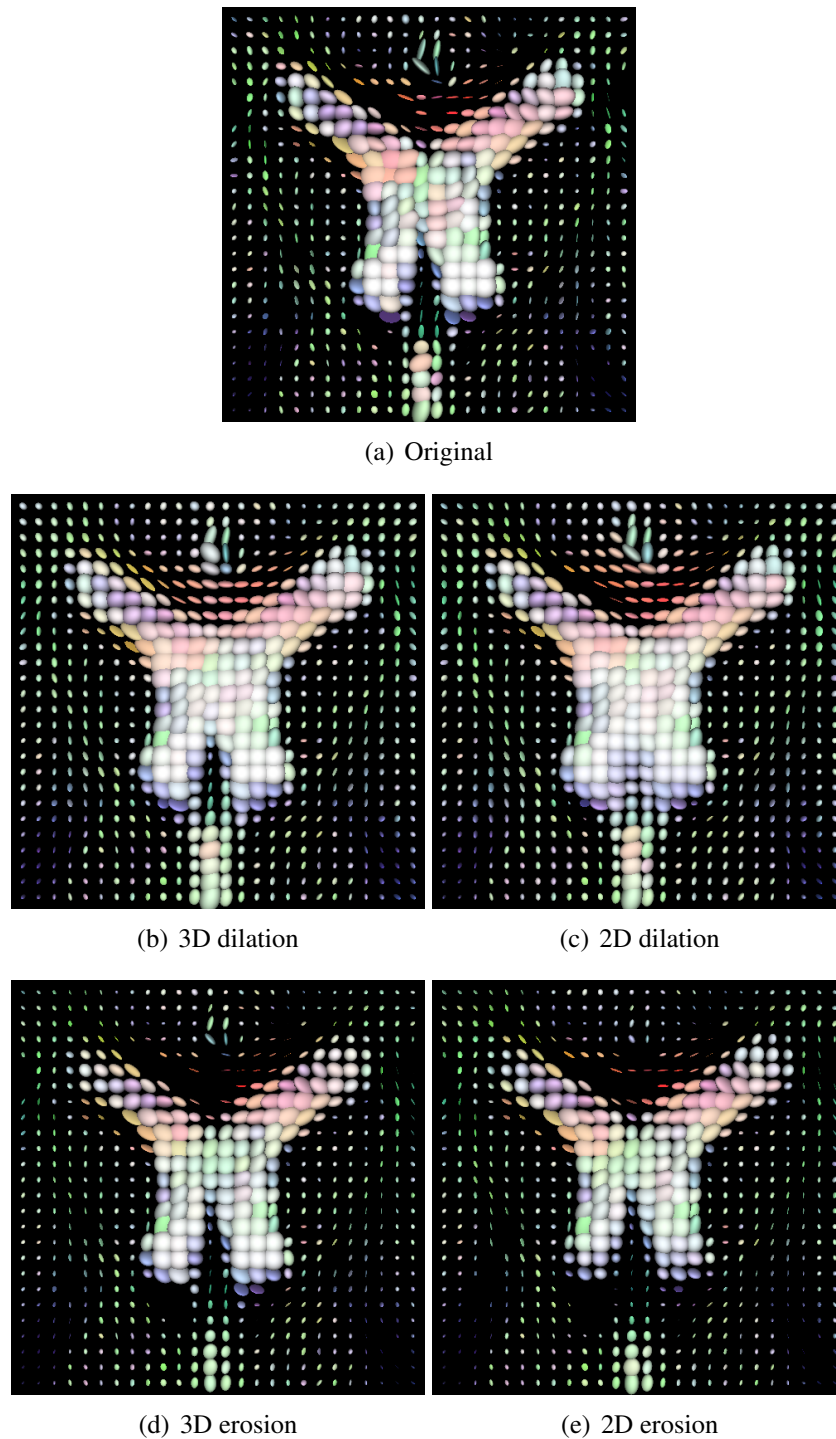





Figure 6.10: Adaptive, anisotropic dilation and erosion in 3D and 2D using the directional FCT scheme with parameters $\sigma = 1, \rho = 1, c = (0.05, 0.05, 5), t = 3$ in the 3D case, and $\sigma = 1, \rho = 1, c = (0.05, 5), t = 1$ in the 2D case.




(a) Original




(b) 3D opening



(c) 2D opening



(d) 3D closing



(e) 2D closing

Figure 6.11: Adaptive, anisotropic opening and closing in 3D and 2D using the directional FCT scheme with parameters $\sigma = 1, \rho = 1, c = (0.05, 0.05, 5), t = 3$ in the 3D case, and $\sigma = 1, \rho = 1, c = (0.05, 5), t = 1$ in the 2D case.




Figure 6.12: Adaptive, anisotropic Beucher gradient and shock filtering in 3D and 2D using the directional FCT scheme with parameters $\sigma = 1, \rho = 1, c = (0.1, 0.1, 10), t = 5$ in the 3D case, and $\sigma = 1, \rho = 1, c = (0.1, 10), t = 2$ in the 2D case.

7

Conclusions and Perspectives

7.1 Conclusions

The main motivation of this work was to develop novel reconstruction methods for scalar- and matrix-valued images in order to ameliorate the effects of noise degradation and missing data in practical applications. We addressed these challenges from the perspectives of nonlocal smoothing and adaptive morphology. In this section we briefly summarise our contributions in these contexts.

Nonlocal Smoothing

In Chapter 2, after describing the constituent components of the *Nonlocal Data and Smoothness* (NDS) filtering approach of Mrázek *et al.*, we established new connections between this framework and other smoothing approaches such as diffusion/regularisation methods (Subsection 2.2.3) and graph regularisation methods (Section 2.3). Moreover, our experiments in Section 2.4 showed that NDS models with L_1 and L_2 penalisers outperform the classical filters – special cases of the NDS approach – for the task of denoising greyscale images degraded with Gaussian and impulse noise. To guide the process of setting up a desired filter, we also studied parameter configuration from an experimental point of view, showing that there exists a trade-off between several model parameters that allows one to achieve similar smoothing results under different configurations.

Despite the fact that the NDS approach allows for nonlocal pixel interactions, in practice they do not occur within large neighbourhoods. This is due to the fact that single differences in intensity do not carry reliable information about the local image structure too far away from a chosen pixel. This problem is overcome by considering patch-based differences. We thus proposed the *Generalised NDS* (GNDS) approach in Chapter 3 as an extension of the NDS framework that allows the data and smoothness penalisers to act on general dissimilarity measures defined on image patches. In particular, using the weighted L_2 patch distance leads to a filter that induces a new similarity measure that evaluates the patch similarity of two chosen pixels and their corresponding neighbours. This similarity measure has three special cases that lead to the NDS filter, the NL-means filter and a novel filter for removing impulse noise. We also discussed extensions to multichannel images, the use of other similarity measures and showed that the GNDS model outperforms other nonlocal patch-based methods and compares very well to more sophisticated approaches.

We have further studied smoothing approaches for fields of positive semidefinite matrices (e.g. DT-MRI data). In Chapter 4 we have investigated two different approaches that extend the NDS framework to matrix fields. Firstly, we introduced the *isotropic matrix-valued NDS* (iMNDS) model in Section 4.1 where pixel similarity in the scalar setting is replaced by matrix similarity, measured by suitable distances between matrices. Secondly, we proposed the *anisotropic matrix-valued NDS* (aMNDS) model in Section 4.2 where we refrain from using such matrix distances that reduce the directional information contained in the pair of matrices compared to a single scalar value. Instead, the data and smoothness penalisers act directly on the matrix differences. We have additionally shown in Section 4.3 that several filters for symmetric matrices proposed in the literature can be seen as special or related cases of our matrix-valued approaches, which is an inherited property from the scalar NDS model. Finally, our experiments in Section 4.4 demonstrated the suitability of our framework for the regularisation of synthetic and real-world DT-MRI fields. In particular, we showed that tractography algorithms can provide much better fibre tracks when the noisy DT-MRI data is preprocessed with our smoothing approach.

Adaptive Morphology

In Chapter 5 we have reviewed the elementary concepts of mathematical morphology from the discrete, set-theoretical point of view and from the modern continuous-scale approach. Focusing on the latter, we showed how the scalar-valued dilation and erosion processes have been recently transferred to the matrix-valued setting by means of the elegant *operator-algebraic framework* for matrix fields developed by Burgeth *et al.* These isotropic processes treat all pixel/voxel locations equally by growing/shrinking the object's boundaries by the same factor in all directions.

Seeking for morphological operations that could affect the object's boundaries anisotropically, in Chapter 6 we have proposed a steerable PDE-driven approach for scalar- and matrix-valued images that adapts itself to the object's geometry. The steering process relies on the structure tensor concept for scalar and matrix data that allows us to accurately estimate the orientation of local structures. In order to enable proper anisotropic steering we have proposed directional versions of the Rouy-Tourin and the FCT numerical schemes (Section 6.4). A special advantage is that these novel schemes are utilised only in their basic one-dimensional version – no higher dimensional variants required – hence avoiding grid artifacts. Our approach is capable of completing directional structures, filling in missing data and enhancing flow-like patterns in an adaptive, anisotropic manner.

We hope that the methods developed in this thesis can serve as building or preprocessing blocks for other higher-level tasks in image processing and computer vision. For instance, in the field of biomedical image computing these methods can help improve the interpretability of 2D/3D datasets affected by missing and corrupted observations.

7.2 Perspectives

The fundamental ideas for image smoothing and mathematical morphology studied in this thesis can be combined and extended to tackle other problems in image and video analysis. In this section we sketch some future research directions.

- *Nonlocal M-smoothers vs regularisation methods.* We have experimentally shown that the NDS framework can be set up with different configurations yielding similar smoothing results. We observed parametric trade-offs between nonlocal M-smoothers and regularisation/diffusion methods. It would be interesting to investigate theoretical connections between both approaches.
- *Joint nonlocal smoothing and deconvolution.* The discrete variational description of the GNDS approach could be adapted to jointly solve for image smoothing and deconvolution. Different models can be investigated for non-blind and blind deconvolution, Gaussian and impulse noise. Recent variational approaches with nonlocal regularisers have been studied in [112, 110, 131]. Similarly, our approaches for smoothing of matrix-valued images (MNDS) can also be extended to nonlocal deconvolution.
- *Joint nonlocal smoothing and inpainting.* The GNDS framework could also be suitable for the joint smoothing and inpainting of greyscale and colour images. This can be achieved by utilising a nonlocal data term or a (semi-)local data term with a nonlocal regulariser. Related studies can be found in [85, 5].
- *Nonlocal optic flow methods.* Recent papers have shown that modelling optic flow with a nonlocal brightness constancy assumption (data term) [229] and with a nonlocal regulariser [230] results in better motion estimations. As the structure of the GNDS approach fits both models, it would be interesting to employ our framework to investigate more general optic flow methods with different nonlocal constancy assumptions and regularisation terms.
- *Joint nonlocal estimation and regularisation of tensor fields.* In the context of diffusion tensor imaging, our variational MNDS framework could be used to jointly estimate and regularise a single or multiple diffusion tensors per voxel from diffusion weighted images, extending the original methods presented in [209, 165].
- *Adaptive morphology on arbitrary graphs.* Isotropic morphological operations have been extended to graphs of arbitrary topology in the set-theoretical approach [101, 103] and in the PDE-based approach [200, 201]. Following this trend, our framework for adaptive morphology could enable anisotropic operations on 3D surfaces/objects defined on non-uniform meshes. This includes the investigation of corresponding numerical schemes on graphs to ensure proper directional steering.

- *Non-orthogonal adaptive morphology.* The use of the structure tensor in our morphological framework helps steer the dilation/erosion processes along orthogonal orientations. To make our approach suitable for more complex textures outlining multiple non-orthogonal orientations, we can replace the standard second-order structure tensor by its higher-order generalisation [186] or by the tensor voting toolbox [143], which allow to estimate multiple local orientations.
- *Reduction-extension adaptivity.* In order to extract directional information from the full structure tensor for matrix fields we applied a projection (reduction) step followed by an extension step, both parameterised by matrices that we simply set to the identity. However, one could adjust these matrices in a spatially-variant manner, making them dependent on the local structure as well or on other user-specific criteria. This would provide our adaptive morphology approach with an additional degree of freedom that needs to be further explored.

We further hope that our contributions can be exploited beyond the image processing arena, considering the generality and adaptiveness of the developed techniques. For instance, in other areas where tensor fields are commonly used: stress-strain tensors in material science, geomechanics, elastography and fluid dynamics; permittivity tensors in electromagnetism and optics; permeability tensors in fluid mechanics and earth sciences; among others.

8

Further Contributions and Publications

8.1 Further Contributions

Programme Committee Member:

- 3rd IEEE Pacific-Rim Symposium on Image and Video Technology. Tokyo, Japan, January 13–16, 2009.
- 2nd IEEE Pacific-Rim Symposium on Image and Video Technology. Santiago, Chile, December 17–19, 2007. *Co-Chair of the Image Analysis track.*
- 1st IEEE Pacific-Rim Symposium on Image and Video Technology. Hsinchu, Taiwan, December 11–13, 2006.

8.2 Publications

Book Chapters:

1. B. Burgeth, L. Pizarro, S. Didas. Edge-enhancing diffusion filtering for matrix fields. In D. H. Laidlaw and A. Vilanova, editors, *New Developments in the Visualization and Processing of Tensor Fields*. Springer, 2011. To appear.
2. R. Moreno, L. Pizarro, B. Burgeth, J. Weickert, M. A. Garcia, D. Puig. Adaptation of tensor voting to image structure estimation. In D. H. Laidlaw and A. Vilanova, editors, *New Developments in the Visualization and Processing of Tensor Fields*. Springer, 2011. To appear.
3. B. Burgeth, L. Pizarro, S. Didas, J. Weickert. 3D-coherence-enhancing diffusion filtering for matrix fields. In L. Florack, R. Duits, G. Jongbloed, M-C. van Lieshout and L. Davies, editors, *Locally Adaptive Filters in Signal and Image Processing*. Springer, 2010. To appear.

Journal Papers:

4. R. Moreno, M. A. Garcia, D. Puig, L. Pizarro, B. Burgeth, J. Weickert. On improving the efficiency of tensor voting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2011. To appear.

5. B. Burgeth, L. Pizarro, M. Breuß, J. Weickert. Adaptive continuous-scale morphology for matrix fields. *International Journal of Computer Vision*, Vol. 92, No. 2, 146–161, 2011.
6. L. Pizarro, P. Mrázek, S. Didas, S. Grewenig, J. Weickert. Generalised nonlocal image smoothing. *International Journal of Computer Vision*, Vol. 90, No. 1, 62–87, 2010.
7. M. Carrasco, L. Pizarro, D. Mery. Visual inspection of glass bottlenecks by multiple-view analysis. *International Journal of Computer Integrated Manufacturing*, Vol. 23, No. 10, 925–941, 2010.
8. L. Pizarro, D. Mery, R. Delpiano, M. Carrasco. Robust automated multiple view inspection. *Pattern Analysis and Applications*, Vol. 11, No. 1, 21–32, 2008.
9. H. Allende, A. Frery, J. Galbiati, L. Pizarro. M-estimators with asymmetric influence functions: the \mathcal{G}_A^0 distribution case. *Journal of Statistical Computation and Simulation*, Vol. 76, No. 11, 941–956, 2006.

Conference Papers:

10. L. Pizarro, J. Delpiano, P. Aljabar, J. Ruiz-del-Solar, D. Rueckert. Towards dense motion estimation in light and electron microscopy. *Proc. of 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2011.
11. R. Garg, L. Pizarro, D. Rueckert, L. Agapito. Dense multi-frame optic flow for non-rigid objects using subspace constraints. In R. Kimmel, R. Klette and A. Sugimoto (Eds.): *Computer Vision – ACCV 2010*. Lecture Notes in Computer Science, Vol. 6495, 460–473, Springer, Berlin, 2011.
12. L. Pizarro, B. Burgeth, M. Breuß, J. Weickert. A directional Rouy-Tourin scheme for adaptive matrix-valued morphology. In M.H.F. Wilkinson and J.B.T.M. Roerdink (Eds.): *Mathematical Morphology and Its Application to Signal and Image Processing*. Lecture Notes in Computer Science, Vol. 5720, 250–260, Springer, Berlin, 2009.
13. B. Burgeth, M. Breuß, L. Pizarro, J. Weickert. PDE-driven adaptive morphology for matrix fields. In X.-C. Tai, K. Mørken, M. Lysaker and K.-A. Lie (Eds.): *Scale Space and Variational Methods in Computer Vision*. Lecture Notes in Computer Science, Vol. 5567, 247–258, Springer, Berlin, 2009.
14. L. Pizarro, B. Burgeth, S. Didas, J. Weickert. A generic neighbourhood filtering framework for matrix fields. In D. Forsyth and P. Torr and A. Zisserman (Eds.): *Computer Vision – ECCV 2008*. Lecture Notes in Computer Science, Vol. 5304, 521–532. Springer, Berlin, 2008.

15. M. Carrasco, L. Pizarro, D. Mery. Image acquisition and automated inspection of wine bottlenecks by tracking in multiple views. *Proc. of the 8th International Conference on Signal Processing, Computational Geometry and Artificial Vision*, 82–89, 2008.
16. M. Carrasco, L. Pizarro, D. Mery. Bimodal biometric person identification system under perturbations. In D. Mery and L. Rueda (Eds.): *Advances in Image and Video Technology*. Lecture Notes in Computer Science, Vol. 4872, 114–127, Springer, Berlin, 2007.
17. L. Pizarro, S. Didas, F. Bauer, J. Weickert. Evaluating a general class of filters for image denoising. In B.K. Ersbøll and K.S. Pedersen (Eds.): *Image Analysis*. Lecture Notes in Computer Science, Vol. 4522, 601–610, Springer, Berlin, 2007.
18. N. Pérez, L. Pizarro, J. Zaegel. Neural networks in blasting process managing. *6th Chilean Blasting Days*, Viña del Mar, Chile, May 24–27, 2005.
19. H. Allende, L. Pizarro. Robust estimation of roughness parameter in SAR amplitude images. In A. Sanfeliu and J. Ruiz-Shulcloper (Eds.): *Progress in Pattern Recognition, Speech and Image Analysis*. Lecture Notes in Computer Science, Vol. 2905, 129–136, Springer, Berlin, 2003.
20. H. Allende, J. Galbiati, L. Pizarro. GM-estimates for the \mathcal{G}_A^0 distribution in SAR imagery. *10th Latin American Symposium on Remote Sensing*, Cochabamba, Bolivia, Nov. 11–15, 2002.
21. H. Allende, J. Galbiati, L. Pizarro. M-estimadores simétricos de la distribución \mathcal{G}_A^0 . *5th Congreso Latinoamericano de Sociedades de Estadística*, Buenos Aires, Argentina, Oct. 28 – Nov. 1, 2002.
22. H. Allende, J. Galbiati, L. Pizarro. Simulación de la distribución \mathcal{G}_A^0 en imágenes SAR. *4th Congreso Nacional de Investigación de Operaciones*, Curicó, Chile, Oct. 10–12, 2001.

Theses:

23. L. Pizarro. Robust parameter estimation on SAR data distributions (in Spanish). M.Sc. Thesis, Departamento de Informática, Universidad Técnica Federico Santa María, Valparaíso, Chile, December 2003.

Bibliography

- [1] R. Acar and C. R. Vogel. Analysis of bounded variation penalty methods for ill-posed problems. *Inverse Problems*, 10:1217–1229, 1994.
- [2] L. Alvarez, F. Guichard, P.-L. Lions, and J.-M. Morel. Axioms and fundamental equations in image processing. *Archive for Rational Mechanics and Analysis*, 123:199–257, 1993.
- [3] L. Alvarez and L. Mazorra. Signal and image restoration using shock filters and anisotropic diffusion. *SIAM Journal on Numerical Analysis*, 31(2):590–605, April 1994.
- [4] A. B. Arehart, L. Vincent, and B. B. Kimia. Mathematical morphology: The Hamilton–Jacobi connection. In *Proc. Fourth International Conference on Computer Vision*, pages 215–219, Berlin, May 1993. IEEE Computer Society Press.
- [5] P. Arias, G. Facciolo, V. Caselles, and G. Sapiro. A variational framework for exemplar-based image inpainting. *International Journal of Computer Vision*, 93(3):319–347, 2011.
- [6] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56(2):411–421, 2006.
- [7] G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*, volume 147 of *Applied Mathematical Sciences*. Springer, New York, second edition, 2006.
- [8] V. Aurich and J. Weule. Non-linear Gaussian filters performing edge preserving diffusion. In G. Sagerer, S. Posch, and F. Kummert, editors, *Mustererkennung 1995*, Informatik Aktuell, pages 538–545. Springer, Berlin, 1995.
- [9] S. P. Awate and R. T Whitaker. Higher-order image statistics for unsupervised, information-theoretic, adaptive, image filtering. In *Proc. of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 44–51, 2005.
- [10] S. P. Awate and R. T Whitaker. Unsupervised, information-theoretic, adaptive image filtering for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(3):364–376, March 2006.

- [11] N. Azzabou, N. Paragios, F. Guichard, and F. Cao. Variable bandwidth image denoising using image-based noise models. In *Proc. of the 2007 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1–7, 2007.
- [12] G.J.F. Banon, J. Barrera, U.D.M. Braga-Neto, and N.S.T. Hirata, editors. *Proceedings of the 8th International Symposium on Mathematical Morphology*. Instituto Nacional de Pesquisas Espaciais (INPE), São José dos Campos, 2007. <http://urlib.net/dpi.inpe.br/ismm@80/2007/07.16.14.39>.
- [13] L. Bar, N. Kiryati, and N. Sochen. Image deblurring in the presence of impulsive noise. *International Journal of Computer Vision*, 70(3):279–298, 2006.
- [14] D. Barash and D. Comaniciu. A common framework for nonlinear diffusion, adaptive smoothing, bilateral filtering and mean shift. *Image and Video Computing*, 22(1):73–81, 2004.
- [15] J. Bard. *Practical Bilevel Optimization. Algorithms and Applications*. Springer, 1999.
- [16] P. J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.
- [17] P. J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance*, 11:209–219, 1996.
- [18] P. G. Batchelor, M. Moakher, D. Atkinson, F. Calamante, and A. Connelly. A rigorous framework for diffusion tensor calculus. *Magnetic Resonance in Medicine*, 53(1):221–225, 2005.
- [19] M. Bertero, T. A. Poggio, and V. Torre. Ill-posed problems in early vision. *Proceedings of the IEEE*, 76(8):869–889, August 1988.
- [20] J. Besag. Toward bayesian image analysis. *Journal of Applied Statistics*, 16(3):395–407, 1989.
- [21] J. Bigün. *Vision with Direction*. Springer, Berlin, 2006.
- [22] J. Bigün, G. H. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790, August 1991.

- [23] R. C. Bilcu and M. Vehvilainen. Fast nonlocal means for image denoising. In R. A. Martin, J. M. DiCarlo, and N. Samped, editors, *Digital Photography III*, volume 6502 of *Proc. of SPIE*, page 65020R, 2007.
- [24] J. P. Boris and D. L. Book. Flux corrected transport. I. SHASTA, a fluid transport algorithm that works. *Journal of Computational Physics*, 11(1):38–69, 1973.
- [25] J. P. Boris and D. L. Book. Flux corrected transport. III. Minimal error FCT algorithms. *Journal of Computational Physics*, 20:397–431, 1976.
- [26] J. P. Boris, D. L. Book, and K. Hain. Flux corrected transport. II. Generalizations of the method. *Journal of Computational Physics*, 18:248–283, 1975.
- [27] N. Bouaynaya, M. Charif-Chefchaoui, and D. Schonfeld. Theoretical foundations of spatially-variant mathematical morphology part I: Binary images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(5):823–836, May 2008.
- [28] N. Bouaynaya and D. Schonfeld. Theoretical foundations of spatially-variant mathematical morphology part II: Gray-level images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(5):837–850, May 2008.
- [29] S. Bougleux, A. Elmoataz, and M. Melkemi. Discrete regularization on weighted graphs for image and mesh filtering. In F. Sgallari, A. Murli, and N. Paragios, editors, *Scale Space and Variational Methods in Computer Vision*, volume 4485 of *Lecture Notes in Computer Science*, pages 128–139. Springer, Berlin, 2007.
- [30] S. Bougleux, A. Elmoataz, and M. Melkemi. Local and nonlocal discrete regularization on weighted graphs for image and mesh processing. *International Journal of Computer Vision*, 84(2):220–236, 2009.
- [31] C. Bouman and K. Sauer. A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Transactions on Image Processing*, 2(3):296–310, July 1993.
- [32] M. Breuß, B. Burgeth, and J. Weickert. Anisotropic continuous-scale morphology. In J. Martí, J. M. Benedí, A. M. Mendonça, and J. Serrat, editors, *Pattern Recognition and Image Analysis*, volume 4478 of *Lecture Notes in Computer Science*, pages 515–522. Springer, Berlin, 2007.
- [33] M. Breuß and J. Weickert. A shock-capturing algorithm for the differential equations of dilation and erosion. *Journal of Mathematical Imaging and Vision*, 25(2):187–201, September 2006.

- [34] M. Breuß and J. Weickert. Highly accurate schemes for PDE-based morphology with general convex structuring elements. *International Journal of Computer Vision*, 92(2):132–145, 2011.
- [35] R. W. Brockett and P. Maragos. Evolution equations for continuous-scale morphological filtering. *IEEE Transactions on Signal Processing*, 42:3377–3386, 1994.
- [36] L. E. J. Brouwer. Über Abbildungen von Mannigfaltigkeiten. *Mathematische Annalen*, 71(1):97–115, March 1911.
- [37] T. Brox and D. Cremers. Iterated nonlocal means for texture restoration. In F. Sgallari, A. Murli, and N. Paragios, editors, *Scale Space and Variational Methods in Computer Vision*, volume 4485 of *Lecture Notes in Computer Science*, pages 13–24. Springer, Berlin, 2007.
- [38] T. Brox, O. Kleinschmidt, and D. Cremers. Efficient nonlocal means for denoising of textural patterns. *IEEE Transactions on Image Processing*, 17(7):1083–1092, July 2008.
- [39] T. Brox, J. Weickert, B. Burgeth, and P. Mrázek. Nonlinear structure tensors. *Image and Vision Computing*, 24(1):41–55, January 2006.
- [40] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *Proc. of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 60–65, 2005.
- [41] A. Buades, B. Coll, and J.-M. Morel. A review of image denoising algorithms, with a new one. *SIAM Multiscale Modeling and Simulation*, 4(2):490–530, 2005.
- [42] A. Buades, B. Coll, and J.-M. Morel. Neighborhood filters and PDEs. *Numerische Mathematik*, 105(1):1–34, 2006.
- [43] B. Burgeth, M. Breuß, S. Didas, and J. Weickert. PDE-based morphology for matrix fields: Numerical solution schemes. In S. Aja-Fernández, R. de Luis García, D. Tao, and X. Li, editors, *Tensors in Image Processing and Computer Vision*, Advances in Pattern Recognition, pages 125–150. Springer, London, 2009.
- [44] B. Burgeth, M. Breuß, L. Pizarro, and J. Weickert. PDE-driven adaptive morphology for matrix fields. In X.-C. Tai, K. Mørken, M. Lysaker, and K.-A. Lie, editors, *Scale Space and Variational Methods in Computer Vision*, volume 5567 of *Lecture Notes in Computer Science*, pages 247–258. Springer, Berlin, 2009.

- [45] B. Burgeth, A. Bruhn, S. Didas, J. Weickert, and M. Welk. Morphology for matrix-data: Ordering versus PDE-based approach. *Image and Vision Computing*, 25(4):496–511, April 2007.
- [46] B. Burgeth, A. Bruhn, N. Papenberg, M. Welk, and J. Weickert. Mathematical morphology for tensor data induced by the Loewner ordering in higher dimensions. *Signal Processing*, 87(2):277–290, February 2007.
- [47] B. Burgeth, S. Didas, L. Florack, and J. Weickert. A generic approach to diffusion filtering of matrix-fields. *Computing*, 81(2/3):179–197, 2007.
- [48] B. Burgeth, S. Didas, L. Florack, and J. Weickert. A generic approach to the filtering of matrix fields with singular PDEs. In F. Sgallari, F. Murli, and N. Paragios, editors, *Scale Space and Variational Methods in Computer Vision*, volume 4485 of *Lecture Notes in Computer Science*, pages 556–567. Springer, Berlin, 2007.
- [49] B. Burgeth, S. Didas, and J. Weickert. A general structure tensor concept and coherence-enhancing diffusion filtering for matrix fields. In D. H. Laidlaw and J. Weickert, editors, *Visualization and Processing of Tensor Fields: Advances and Perspectives*, Mathematics and Visualization, pages 305–323. Springer, Berlin, 2009.
- [50] B. Burgeth, N. Papenberg, A. Bruhn, M. Welk, C. Feddern, and J. Weickert. Morphology for higher-dimensional tensor data via Loewner ordering. In C. Ronse, L. Najman, and E. Decenci ere, editors, *Mathematical Morphology: 40 Years On*, volume 30 of *Computational Imaging and Vision*, pages 407–418. Springer, Dordrecht, 2005.
- [51] B. Burgeth, L. Pizarro, M. Breu , and J. Weickert. Adaptive continuous-scale morphology for matrix fields. *International Journal of Computer Vision*, 92(2):146–161, 2011.
- [52] B. Burgeth, L. Pizarro, S. Didas, and J. Weickert. 3D-coherence-enhancing diffusion filtering for matrix fields. In L. Florack, R. Duits, G. Jongbloed, M.-C. van Lieshout, and L. Davies, editors, *Locally Adaptive Filters in Signal and Image Processing*. Springer, 2010. To appear.
- [53] J.-F. Cai, R. H. Chan, and M. Nikolova. Two-phase approach for deblurring images corrupted by impulse plus gaussian noise. *Inverse Problems and Imaging*, 2(2):187–204, May 2008.
- [54] F. Cao. *Geometric Curve Evolution and Image Processing*, volume 1805 of *Lecture Notes in Mathematics*. Springer, Berlin, 2003.

- [55] V. Cerney. A thermodynamic approach to the travelling salesman problem: An efficient simulation algorithm. *Journal of Optimization Theory and Applications*, 45:41–51, 1985.
- [56] R. H. Chan, C. W. Ho, and M. Nikolova. Salt-and-pepper noise removal by median-type noise detectors and detail-preserving regularization. *IEEE Transactions on Image Processing*, 14(10):1479–1485, October 2005.
- [57] T. F. Chan, S. Osher, and J. Shen. The digital TV filter and nonlinear denoising. *IEEE Transactions on Image Processing*, 10(2):231–241, February 2001.
- [58] T. F. Chan and J. Shen. *Image Processing and Analysis: Variational, PDE, Wavelet, and Stochastic Methods*. SIAM, Philadelphia, 2005.
- [59] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proc. 1994 IEEE International Conference on Image Processing*, volume 2, pages 168–172, Austin, TX, November 1994. IEEE Computer Society Press.
- [60] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud. Deterministic edge-preserving regularization in computed imaging. *IEEE Transactions on Image Processing*, 6(2):298–311, February 1997.
- [61] P. Chatterjee and P. Milanfar. A generalization of non-local means via kernel regression. In C. A. Bouman, E. L. Miller, and I. Pollak, editors, *Computational Imaging VI*, volume 6814 of *Proc. of SPIE*, page 68140P, 2008.
- [62] C. Ched'Hotel, D. Tschumperlé, R. Deriche, and O. Faugeras. Constrained flows of matrix-valued functions: Application to diffusion tensor regularization. In A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, editors, *Computer Vision – ECCV 2002*, volume 2350 of *Lecture Notes in Computer Science*, pages 251–265. Springer, Berlin, 2002.
- [63] C. Ched'hotel, D. Tschumperlé, R. Deriche, and O. Faugeras. Regularizing flows for constrained matrix-valued images. *Journal of Mathematical Imaging and Vision*, 20(1-2):147–162, 2004.
- [64] Y. Chen, J. Ma, Q. Feng, L. Luo, P. Shi, and W. Chen. Nonlocal prior Bayesian tomographic reconstruction. *Journal of Mathematical Imaging and Vision*, 30(2):133–146, 2008.
- [65] Y. Cheng. Mean shift, mode seeking, and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):790–799, August 1995.

- [66] C. K. Chu, I. K. Glad, F. Godtlielsen, and J. S. Marron. Edge-preserving smoothers for image processing. *Journal of the American Statistical Association*, 93(442):526–541, June 1998.
- [67] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, May 2002.
- [68] P. Coupé, P. Hellier, C. Kervrann, and C. Barillot. Bayesian non local means-based speckle filtering. In *Proc. of the 2008 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1291–1294, 2008.
- [69] P. Coupé, P. Yger, S. Prima, P. Hellier, C. Kervrann, and C. Barillot. An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images. *IEEE Transactions on Medical Imaging*, 27(4):425–441, April 2008.
- [70] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3D transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, August 2007.
- [71] J. Darbon, A. Cunha, T. F. Chan, S. Osher, and G. J. Jensen. Fast nonlocal filtering applied to electron cryomicroscopy. In *Proc. of the 2008 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1331–1334, 2008.
- [72] S. Di Zenzo. A note on the gradient of a multi-image. *Computer Vision, Graphics and Image Processing*, 33:116–125, 1986.
- [73] S. Didas. *Denoising and Enhancement of Digital Images – Variational Methods, Integrodifferential Equations, and Wavelets*. PhD thesis, Saarland University, Germany, 2008.
- [74] S. Didas, P. Mrázek, and J. Weickert. Energy-based image simplification with nonlocal data and smoothness terms. In A. Iske and J. Levesley, editors, *Algorithms for Approximation*, pages 50–60. Springer, Berlin, 2007.
- [75] S. Didas and J. Weickert. From adaptive averaging to accelerated nonlinear diffusion filtering. In K. Franke, K.-R. Müller, B. Nickolay, and R. Schäfer, editors, *Pattern Recognition*, volume 4174 of *Lecture Notes in Computer Science*, pages 101–110. Springer, Berlin, September 2006.
- [76] S. Didas and J. Weickert. Integrodifferential equations for continuous multi-scale wavelet shrinkage. *Inverse Problems and Imaging*, 1(1):47–62, February 2007.

- [77] M. Elad. On the origin of the bilateral filter and ways to improve it. *IEEE Transactions on Image Processing*, 11(10):1141–1151, October 2002.
- [78] A. Elmoataz, O. Lézoray, and S. Bougleux. Nonlocal discrete regularization on weighted graphs: A framework for image and manifold processing. *IEEE Transactions on Image Processing*, 17(7):1047–1060, July 2008.
- [79] A. Elmoataz, O. Lézoray, S. Bougleux, and V.-T. Ta. Unifying local and nonlocal processing with partial difference operators on weighted graphs. In A. Foi and A. Gotchev, editors, *Proc. International Workshop on Local and Non-Local Approximation in Image Processing*, pages 11–26, 2008.
- [80] C. Feddern, J. Weickert, B. Burgeth, and M. Welk. Curvature-driven PDE methods for matrix-valued images. *International Journal of Computer Vision*, 69(1):91–103, August 2006.
- [81] P. Fillard, X. Pennec, V. Arsigny, and N. Ayache. Clinical DT-MRI estimation, smoothing, and fiber tracking with log-euclidean metrics. *IEEE Transactions on Medical Imaging*, 26(11):1472–1482, November 2007.
- [82] P. T. Fletcher and S. Joshi. Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing*, 87(2):250–262, 2007.
- [83] W. Förstner and E. Gülch. 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, June 1987.
- [84] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- [85] M. Ghoniem, Y. Chahir, and A. Elmoataz. Nonlocal video denoising, simplification and inpainting using discrete regularization on graphs. *Signal Processing*, 90(8):2445–2455, 2010.
- [86] G. Gilboa, J. Darbon, S. Osher, and T. F. Chan. Nonlocal convex functionals for image regularization. Technical Report CAM-06-57, Department of Mathematics, University of California at Los Angeles, CA, U.S.A., 2006.
- [87] G. Gilboa and S. Osher. Nonlocal linear image regularization and supervised segmentation. *Multiscale Modeling and Simulation*, 6(2):595–630, 2007.
- [88] G. Gilboa and S. Osher. Nonlocal operators with applications to image processing. *Multiscale Modeling and Simulation*, 7(3):1005–1028, 2008.

- [89] G. Gilboa, N. A. Sochen, and Y. Y. Zeevi. Regularized shock filters and complex diffusion. In A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, editors, *Computer Vision – ECCV 2002*, volume 2350 of *Lecture Notes in Computer Science*, pages 399–413. Springer, Berlin, 2002.
- [90] B. Goossens, H. Luong, A. Pizurica, and W. Philips. An improved non-local denoising algorithm. In A. Foi and A. Gotchev, editors, *Proc. International Workshop on Local and Non-Local Approximation in Image Processing*, pages 143–156, 2008.
- [91] J. Goutsias, L. Vincent, and D. S. Bloomberg, editors. *Mathematical Morphology and its Applications to Image and Signal Processing*, volume 18 of *Computational Imaging and Vision*. Kluwer, Dordrecht, 2000.
- [92] L. D. Griffin. Mean, median and mode filtering of images. *Proceedings of the Royal Society of London, Series A*, 456(2004):2995–3004, 2000.
- [93] H. Gudbjartsson and S. Patz. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34(6):910–914, December 1995. [published erratum appears in *Magnetic Resonance in Medicine* 36 (1996) 332].
- [94] F. Guichard and J.-M. Morel. A note on two classical enhancement filters and their associated PDE’s. *International Journal of Computer Vision*, 52(2/3):153–160, May 2003.
- [95] F. Guichard, J.-M. Morel, and R. Ryan. *Contrast invariant image analysis and PDE’s*. http://mw.cmla.ens-cachan.fr/~morel/JMMBookOct04_complete.pdf.
- [96] Y. Gur and N. Sochen. Coordinate-free diffusion over compact Lie-groups. In F. Sgallari, A. Murli, and N. Paragios, editors, *Scale Space and Variational Methods in Computer Vision*, volume 4485 of *Lecture Notes in Computer Science*, pages 580–591. Springer, Berlin, 2007.
- [97] K. R. Hahn, S. Prigarin, S. Heim, and H. Hasan. Random noise in diffusion tensor imaging, its destructive impact and some corrections. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, Mathematics and Visualization, pages 107–119. Springer, Berlin, 2006.
- [98] G. Hamarneh and J. Hradsky. Bilateral filtering of diffusion tensor magnetic resonance images. *IEEE Transactions on Image Processing*, 16(10):2463–2475, October 2007.
- [99] F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. Probability and Mathematical Statistics. Wiley & Sons, New York, 1986.

- [100] H. J. A. M. Heijmans. *Morphological Image Operators*. Academic Press, Boston, 1994.
- [101] H. J. A. M. Heijmans, P. Nacken, A. Toet, and L. Vincent. Graph morphology. *Journal of Visual Communication and Image Representation*, 3(1):24–38, 1992.
- [102] H. J. A. M. Heijmans and J. B. T. M. Roerdink, editors. *Mathematical Morphology and its Applications to Image and Signal Processing*, volume 12 of *Computational Imaging and Vision*. Kluwer, Dordrecht, 1998.
- [103] H. J. A. M. Heijmans and L. Vincent. Graph morphology in image analysis. In E. Dougherty, editor, *Mathematical Morphology in Image Processing*, volume 34 of *Optical Engineering*, pages 171–203. Marcel Dekker, New York, 1993.
- [104] M. Hein, J.-Y. Audibert, and U. von Luxburg. Graph Laplacians and their convergence on random neighborhood graphs. *Journal of Machine Learning Research*, 8:1325–1370, 2007.
- [105] G.T. Herman, H. Hurwitz, A. Lent, and H-P. Lung. On the bayesian approach to image reconstruction. *Information and Control*, 42:60–71, 1979.
- [106] P. W. Holland and R. E. Welsch. Robust regression using iteratively reweighted least-squares. *Communications in Statistics*, 6:813–827, 1977.
- [107] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1990.
- [108] P. J. Huber. *Robust Statistics*. Wiley, New York, 1981.
- [109] B. R. Hunt. Bayesian methods in nonlinear digital image restoration. *IEEE Transactions on Computers*, 26:219–229, 1977.
- [110] M. Jung, X. Bresson, T. Chan, and L. Vese. Color image restoration using nonlocal mumford-shah regularizers. In D. Cremers, Y. Boykov, A. Blake, and F. Schmidt, editors, *Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 5681 of *Lecture Notes in Computer Science*, pages 373–387. Springer Berlin, 2009.
- [111] M. Jung and L. Vese. Image restoration via nonlocal Mumford-Shah regularizers. Technical Report CAM-09-09, Department of Mathematics, University of California at Los Angeles, CA, U.S.A., 2009.
- [112] M. Jung and L. Vese. Nonlocal variational image deblurring models in the presence of Gaussian or impulse noise. In X.-C. Tai, K. Mørken, M. Lysaker, and K.-A. Lie, editors, *Scale Space and Variational Methods in Computer*

- Vision*, volume 5567 of *Lecture Notes in Computer Science*, pages 401–412. Springer, Berlin, 2009.
- [113] C. Kervrann and J. Boulanger. Optimal spatial adaptation for patch-based image denoising. *IEEE Transactions on Image Processing*, 15(10):2866–2878, October 2006.
- [114] C. Kervrann and J. Boulanger. Local adaptivity to variable smoothness for exemplar-based image regularization and representation. *International Journal of Computer Vision*, 79(1):45–69, August 2008.
- [115] C. Kervrann, J. Boulanger, and P. Coupé. Bayesian non-local means filter, image redundancy and adaptive dictionaries for noise removal. In F. Sgallari, A. Murli, and N. Paragios, editors, *Scale Space and Variational Methods in Computer Vision*, volume 4485 of *Lecture Notes in Computer Science*, pages 520–532. Springer, Berlin, 2007.
- [116] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Shapes, shocks, and deformations I: The components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision*, 15:189–224, 1995.
- [117] S. Kindermann, S. Osher, and P. W. Jones. Deblurring and denoising of images by nonlocal functionals. *SIAM Multiscale Modeling and Simulation*, 4(4):1091–1115, 2005.
- [118] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.
- [119] O. Kleinschmidt, T. Brox, and D. Cremers. Nonlocal texture filtering with efficient tree structures and invariant patch similarity measures. In A. Foi and A. Gotchev, editors, *Proc. International Workshop on Local and Non-Local Approximation in Image Processing*, pages 103–113, 2008.
- [120] J. J. Koenderink and A. L. Van Doorn. The structure of locally orderless images. *International Journal of Computer Vision*, 31(2/3):159–168, 1999.
- [121] H. P. Kramer and J. B. Bruckner. Iterations of a non-linear transformation for enhancement of digital images. *Pattern Recognition*, 7:53–58, 1975.
- [122] H. R. Künsch. Robust priors for smoothing and image restoration. *Annals of the Institute of Statistical Mathematics*, 46:1–19, 1994.
- [123] D. Laidlaw and J. Weickert, editors. *Visualization and Processing of Tensor Fields*. Springer, Berlin, 2009.

- [124] Y. G. Leclerc. Constructing simple stable descriptions for image partitioning. *International Journal of Computer Vision*, 3(1):73–102, 1989.
- [125] J.-S. Lee. Digital image smoothing and the sigma filter. *Computer Vision, Graphics, and Image Processing*, 24:255–269, 1983.
- [126] R. J. LeVeque. *Numerical Methods for Conservation Laws*. Birkhäuser, Basel, 1992.
- [127] J. Lie, B. Burgeth, and O. Christiansen. An operator algebraic inverse scale space method for symmetric matrix valued images. In D. H. Laidlaw and J. Weickert, editors, *Visualization and Processing of Tensor Fields: Advances and Perspectives*, Mathematics and Visualization, pages 361–376. Springer, Berlin, 2009.
- [128] J. Liu, Z. Huan, H. Huang, and H. Zhang. An adaptive method for recovering image from mixed noisy data. *International Journal of Computer Vision*, 85(2):182–191, 2009.
- [129] Y.-L. Liu, J. Wang, X. Chen, and Y.-W. Guo. A robust and fast non-local means algorithm for image denoising. *Journal of Computer Science and Technology*, 23(2):270–279, March 2008.
- [130] Y. Lou, P. Favaro, S. Soatto, and A. Bertozzi. Nonlocal similarity image filtering. In P. Foggia, C. Sansone, and M. Vento, editors, *Image Analysis and Processing - ICIAP 2009*, volume 5716 of *Lecture Notes in Computer Science*, pages 62–71. Springer, Berlin, 2009.
- [131] Y. Lou, X. Zhang, S. Osher, and A. Bertozzi. Image recovery via nonlocal operators. *Journal of Scientific Computing*, 42:185–197, 2010.
- [132] T. Loupas, W. N. McDicken, and P. L. Allan. An adaptive weighted median filter for speckle suppression in medical ultrasonic images. *IEEE Transactions on Circuits and Systems*, 36(1):129–135, January 1989.
- [133] J. R. Magnus and H. Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. J. Wiley and Sons, Chichester, 1988.
- [134] M. Mahmoudi and G. Sapiro. Fast image and video denoising via nonlocal means of similar neighborhoods. *IEEE Signal Processing Letters*, 12(12):839–842, December 2005.
- [135] P. Maragos and R. W. Schafer. Morphological systems for multidimensional signal processing. *Proceedings of the IEEE*, 78(4):690–710, April 1990.

- [136] P. Maragos, R. W. Schafer, and M. A. Butt, editors. *Mathematical Morphology and its Applications to Image and Signal Processing*, volume 5 of *Computational Imaging and Vision*. Kluwer, Dordrecht, 1996.
- [137] P. Maragos and C. Vachier. A pde formulation for viscous morphological operators with extensions to intensity-adaptive operators. In *Proc. of the 15th IEEE International Conference on Image Processing*, pages 2200–2203, 2008.
- [138] P. Maragos and C. Vachier. Overview of adaptive morphology: Trends and perspectives. In *Proc. of the 16th IEEE International Conference on Image Processing*, pages 2241–2244, 2009.
- [139] J. E. Marsden, T. Ratiu, and R. Abraham. *Manifolds, Tensor Analysis, and Applications*. Applied Mathematical Sciences. Springer, New York, NY, USA, third edition, 2001.
- [140] G. Matheron. *Éléments pour une théorie des milieux poreux*. Masson, Paris, 1967.
- [141] G. Matheron. *Random Sets and Integral Geometry*. Wiley, New York, 1975.
- [142] G. Matheron and J. Serra. The birth of mathematical morphology. In H. Talbot and R. Beare, editors, *Proc. of the Sixth International Symposium on Mathematical Morphology*, pages 1–16. CSIRO Publishing, Sydney, Australia, April 2002.
- [143] G. Medioni, M.-S. Lee, and C.-K. Tang. *A Computational Framework for Segmentation and Grouping*. Elsevier, Amsterdam, 2000.
- [144] M. Moakher. A differential geometry approach to the geometric mean of symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications*, 26(3):735–747, June 2005.
- [145] M. Moakher and P. G. Batchelor. Symmetric positive definite matrices: From geometry to applications and visualization. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, Mathematics and Visualization, pages 285–298. Springer, Berlin, 2006.
- [146] S. Mori and P. C. van Zijl. Fiber tracking: principles and strategies - a technical review. *NRM in Biomedicine*, 15:468–480, 2002.
- [147] P. Mrázek and M. Navara. Selection of optimal stopping time for nonlinear diffusion filtering. *International Journal of Computer Vision*, 52(2/3):189–203, 2003. stopping time.

- [148] P. Mrázek, J. Weickert, and A. Bruhn. On robust estimation and smoothing with spatial and tonal kernels. In R. Klette, R. Kozera, L. Noakes, and J. Weickert, editors, *Geometric Properties from Incomplete Data*, volume 31 of *Computational Imaging and Vision*, pages 335–352. Springer, Dordrecht, 2006.
- [149] D. Mumford. The Bayesian rationale for energy functionals. In B. M. ter Haar Romeny, editor, *Geometry-Driven Diffusion in Computer Vision*, volume 1 of *Computational Imaging and Vision*, pages 141–153. Kluwer, Dordrecht, 1994.
- [150] D. Mumford and J. Shah. Optimal approximation of piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42:577–685, 1989.
- [151] M. Z. Nashed and O. Scherzer. Least squares and bounded variation regularization with nondifferentiable functionals. *Numerical Functional Analysis and Optimization*, 19:873–901, 1998.
- [152] M. Nikolova. Minimizers of cost-functions involving nonsmooth data-fidelity terms. Application to the processing of outliers. *SIAM Journal on Numerical Analysis*, 40(3):965–994, 2002.
- [153] M. Nikolova. Analysis of the recovery of edges in images and signals by minimizing nonconvex regularized least-squares. *SIAM Multiscale Modeling and Simulation*, 4(3):960–991, 2005.
- [154] M. Nikolova. Counter-examples for Bayesian MAP restoration. In F. Sgallari, A. Murli, and N. Paragios, editors, *Scale Space and Variational Methods in Computer Vision*, volume 4485 of *Lecture Notes in Computer Science*, pages 140–152. Springer, Berlin, 2007.
- [155] M. Nikolova. Model distortions in Bayesian MAP reconstruction. *Inverse Problems and Imaging*, 1(2):399–422, May 2007.
- [156] N. Nordström. Biased anisotropic diffusion – a unified regularization and diffusion approach to edge detection. *Image and Vision Computing*, 8:318–327, 1990.
- [157] P. G. Nucifora, R. Verma, S.-K. Lee, and E. R. Melhem. Diffusion-tensor MR imaging and tractography: Exploring brain microstructure and connectivity. *Radiology*, 245:367–384, 2007.
- [158] J. Orchard, M. Ebrahimi, and A. Wong. Efficient non-local-means denoising using the SVD. In *Proc. of the 15th IEEE International Conference on Image Processing*, pages 1732–1735. IEEE Computer Society Press, 2008.

- [159] S. Osher and R. P. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*, volume 153 of *Applied Mathematical Sciences*. Springer, New York, 2002.
- [160] S. Osher and L. Rudin. Shocks and other nonlinear filtering applied to image processing. In A. G. Tescher, editor, *Applications of Digital Image Processing XIV*, volume 1567 of *Proceedings of SPIE*, pages 414–431. SPIE Press, Bellingham, 1991.
- [161] S. Osher and L. I. Rudin. Feature-oriented image enhancement using shock filters. *SIAM Journal on Numerical Analysis*, 27(4):919–940, August 1990.
- [162] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [163] S. Pajevic and P. J. Basser. Parametric and non-parametric statistical analysis of DT-MRI data. *Journal of Magnetic Resonance*, 161(1):1–14, March 2003.
- [164] S. Pajevic and C. Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, 42:526–540, 1999.
- [165] O. Pasternak, Y. Assaf, N. Intrator, and N. Sochen. Variational multiple-tensor fitting of fiber-ambiguous diffusion-weighted magnetic resonance imaging voxels. *Magnetic Resonance Imaging*, 26(8):1133–1144, 2008.
- [166] X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66, 2006.
- [167] P. Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:629–639, 1990.
- [168] J. D. Peter, V. K. Govindan, and A. T. Mathew. Robust estimation approach for NL-means filter. In G. Bebis, R. Boyle, B. Parvin, D. Koracin, P. Remagnino, F. Porikli, J. Peters, J. Klosowski, L. Arns, K.C. Yu, T-M. Rhyne, and L. Monroe, editors, *Advances in Visual Computing*, volume 5359 of *Lecture Notes in Computer Science*, pages 571–580. Springer, Berlin, 2008.
- [169] G. Peyré, S. Bougleux, and L. Cohen. Non-local regularization of inverse problems. In D. Forsyth, P. Torr, and A. Zisserman, editors, *Computer Vision – ECCV 2008, Part III*, volume 5304 of *Lecture Notes in Computer Science*, pages 57–68. Springer, Berlin, October 2008.

- [170] L. Pizarro, B. Burgeth, M. Breuß, and J. Weickert. A directional Rouy-Tourin scheme for adaptive matrix-valued morphology. In M.H.F. Wilkinson and J.B.T.M. Roerdink, editors, *Mathematical Morphology and Its Application to Signal and Image Processing*, volume 5720 of *Lecture Notes in Computer Science*, pages 250–260, Berlin, 2009. Springer.
- [171] L. Pizarro, B. Burgeth, S. Didas, and J. Weickert. A generic neighbourhood filtering framework for matrix fields. In D. Forsyth, P. Torr, and A. Zisserman, editors, *Computer Vision – ECCV 2008, Part III*, volume 5304 of *Lecture Notes in Computer Science*, pages 521–532. Springer, Berlin, October 2008.
- [172] L. Pizarro, B. Burgeth, S. Didas, and J. Weickert. A general variational framework for smoothing matrix fields. In preparation, 2011.
- [173] L. Pizarro, S. Didas, F. Bauer, and J. Weickert. Evaluating a general class of filters for image denoising. In B. K. Ersbøll and K. S. Pedersen, editors, *Image Analysis*, volume 4522 of *Lecture Notes in Computer Science*, pages 601–610. Springer, Berlin, 2007.
- [174] L. Pizarro, P. Mrázek, S. Didas, S. Grewenig, and J. Weickert. Generalised non-local image smoothing. *International Journal of Computer Vision*, 90(1):62–87, 2010.
- [175] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli. Image denoising using Gaussian scale mixtures in the wavelet domain. *IEEE Transactions on Image Processing*, 12(11):1338–1351, November 2003.
- [176] L. Remaki and M. Cheriet. Numerical schemes of shock filter models for image enhancement and restoration. *Journal of Mathematical Imaging and Vision*, 18(2):153–160, March 2003.
- [177] C. Ronse, L. Najman, and E. Decencière, editors. *Mathematical Morphology: 40 Years On*, volume 30 of *Computational Imaging and Vision*. Springer, Dordrecht, 2005.
- [178] E. Rouy and A. Tourin. A viscosity solutions approach to shape-from-shading. *SIAM Journal on Numerical Analysis*, 29:867–884, 1992.
- [179] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
- [180] P. Saint-Marc, J. S. Chen, and G. Medioni. Adaptive smoothing: A general tool for early vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):514–529, June 1991.

- [181] G. Sapiro. *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press, Cambridge, UK, 2001.
- [182] G. Sapiro, R. Kimmel, D. Shaked, B. B. Kimia, and A. M. Bruckstein. Implementing continuous-scale morphology via curve evolution. *Pattern Recognition*, 26:1363–1372, 1993.
- [183] J. G. M. Schavemaker, M. J. T. Reinders, and R. van den Boomgaard. Image sharpening by morphological filtering. In *Proc. 1997 IEEE Workshop on Nonlinear Signal and Image Processing*, Mackinac Island, MI, September 1997. www.ecn.purdue.edu/NSIP/.
- [184] O. Scherzer and J. Weickert. Relations between regularization and diffusion filtering. *Journal of Mathematical Imaging and Vision*, 12:43–63, 2000.
- [185] C. Schnörr. Unique reconstruction of piecewise smooth images by minimizing strictly convex non-quadratic functionals. *Journal of Mathematical Imaging and Vision*, 4:189–198, 1994.
- [186] T. Schultz, J. Weickert, and H.-P. Seidel. A higher-order structure tensor. In D. H. Laidlaw and J. Weickert, editors, *Visualization and Processing of Tensor Fields: Advances and Perspectives*, Mathematics and Visualization, pages 263–279. Springer, Berlin, 2009.
- [187] J. Serra. *Echantillonnage et estimation des phénomènes de transition minier*. PhD thesis, University of Nancy, France, 1967.
- [188] J. Serra. *Image Analysis and Mathematical Morphology*, volume 1. Academic Press, London, 1982.
- [189] J. Serra. *Image Analysis and Mathematical Morphology*, volume 2. Academic Press, London, 1988.
- [190] J. Serra and P. Salembier, editors. *Proc. International Workshop on Mathematical Morphology and its Applications to Signal Processing*. UPC Publications Office, Barcelona, Spain, May 1993.
- [191] J. Serra and P. Soille, editors. *Mathematical Morphology and its Applications to Image Processing*, volume 2 of *Computational Imaging and Vision*. Kluwer, Dordrecht, 1994.
- [192] J. A. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge University Press, Cambridge, UK, second edition, 1999. Paperback edition.

- [193] S. Setzer, G. Steidl, B. Popilka, and B. Burgeth. Variational methods for denoising matrix fields. In D. H. Laidlaw and J. Weickert, editors, *Visualization and Processing of Tensor Fields: Advances and Perspectives*, Mathematics and Visualization, pages 341–360. Springer, Berlin, 2009.
- [194] I. H. Sloan and V. Thomée. Time discretisation of an integro-differential equation of parabolic type. *SIAM Journal on Numerical Analysis*, 23(5):1052–1061, October 1986.
- [195] S. M. Smith and J. M. Brady. SUSAN – A new approach to low-level image processing. *International Journal of Computer Vision*, 23(1):45–78, May 1997.
- [196] P. Soille. *Morphological Image Analysis*. Springer, Berlin, second edition, 2003.
- [197] G. Steidl, S. Setzer, B. Popilka, and B. Burgeth. Restoration of matrix fields by second order cone programming. *Computing*, 81(2/3):161–178, 2007.
- [198] G. Steidl, J. Weickert, T. Brox, P. Mrázek, and M. Welk. On the equivalence of soft wavelet shrinkage, total variation diffusion, total variation regularization, and SIDEs. *SIAM Journal on Numerical Analysis*, 42(2):686–713, 2004.
- [199] R. L. Stevenson, B. E. Schmitz, and E. J. Delp. Discontinuity preserving regularization of inverse visual problems. *IEEE Transactions on Systems, Man and Cybernetics*, 24:455–469, 1994.
- [200] V.-T. Ta, A. Elmoataz, and O. Lézoray. Partial difference equations over graphs: Morphological processing of arbitrary discrete data. In D. Forsyth, P. Torr, and A. Zisserman, editors, *Computer Vision – ECCV 2008, Part III*, volume 5304 of *Lecture Notes in Computer Science*, pages 668–680. Springer, Berlin, October 2008.
- [201] V.-T. Ta, A. Elmoataz, and O. Lézoray. Mathematical morphology and eikonal equations on graphs for nonlocal image and data processing. Technical Report GREYC CNRS UMR 6072, Université de Caen Basse-Normandie - ENSICAEN, Caen, France, March 2009.
- [202] V.-T. Ta, O. Lézoray, A. Elmoataz, and S. Schüpp. Graph-based tools for microscopic cellular image segmentation. *Pattern Recognition*, 42(6):1113–1125, 2009.
- [203] H. Talbot and R. Beare, editors. *Proceeding of the Sixth International Symposium on Mathematical Morphology*. CSIRO Publishing, Sydney, Australia, April 2002. <http://www.cmis.csiro.au/ismm2002/proceedings/>.

- [204] A. N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics Doklady*, 4(2):1035–1038, 1963.
- [205] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-Posed Problems*. Wiley, Washington, DC, 1977.
- [206] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. Sixth International Conference on Computer Vision*, pages 839–846, Bombay, India, January 1998. Narosa Publishing House.
- [207] D. Tschumperlé and L. Brun. Image denoising and registration by PDE's on the space of patches. In A. Foi and A. Gotchev, editors, *Proc. International Workshop on Local and Non-Local Approximation in Image Processing*, pages 3–10, 2008.
- [208] D. Tschumperlé and R. Deriche. Diffusion tensor regularization with constraints preservation. In *Proc. 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 948–953, Kauai, HI, December 2001. IEEE Computer Society Press.
- [209] D. Tschumperlé and R. Deriche. Variational frameworks for DT-MRI estimation, regularization and visualization. In *Proc. Ninth International Conference on Computer Vision*, volume 1, pages 116–121, Nice, France, October 2003. IEEE Computer Society Press.
- [210] J. van de Weijer and R. van den Boomgaard. Local mode filtering. In *Proc. 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 428–433, Kauai, HI, December 2001. IEEE Computer Society Press.
- [211] R. van den Boomgaard. *Mathematical Morphology: Extensions Towards Computer Vision*. PhD thesis, University of Amsterdam, The Netherlands, 1992.
- [212] R. van den Boomgaard. Numerical solution schemes for continuous-scale morphology. In M. Nielsen, P. Johansen, O. F. Olsen, and J. Weickert, editors, *Scale-Space Theories in Computer Vision*, volume 1682 of *Lecture Notes in Computer Science*, pages 199–210. Springer, Berlin, 1999.
- [213] R. van den Boomgaard and J. van de Weijer. On the equivalence of local-mode finding, robust estimation and mean-shift analysis as used in early vision tasks. In *Proc. 16th International Conference on Pattern Recognition*, volume 3, pages 927–930, Quebec City, Canada, August 2002.

- [214] B. van Ginneken and B. M. ter Haar Romeny. Applications of locally orderless images. *Journal of Visual Communication and Image Representation*, 11(2):196–208, June 2000.
- [215] L. J. van Vliet, I. T. Young, and A. L. D. Beckers. A nonlinear Laplace operator as edge detector in noisy images. *Computer Vision, Graphics and Image Processing*, 45(2):167–195, 1989.
- [216] R. Verdú-Monedero and J. Angulo. Spatially-variant directional mathematical morphology operators based on a diffused average squared gradient field. In J. Blanc-Talon et al., editor, *Advanced Concepts for Intelligent Vision Systems*, volume 5259 of *Lecture Notes in Computer Science*, pages 542–553, Berlin, 2008. Springer.
- [217] E. R. Vrscay. A simple model for affine self-similarity of images and its applications. In A. Foi and A. Gotchev, editors, *Proc. International Workshop on Local and Non-Local Approximation in Image Processing*, pages 61–76, 2008.
- [218] J. Weickert. Anisotropic diffusion filters for image processing based quality control. In A. Fasano and M. Primicerio, editors, *Proc. Seventh European Conference on Mathematics in Industry*, pages 355–362. Teubner, Stuttgart, 1994.
- [219] J. Weickert. A review of nonlinear diffusion filtering. In B. ter Haar Romeny, L. Florack, J. Koenderink, and M. Viergever, editors, *Scale-Space Theory in Computer Vision*, volume 1252 of *Lecture Notes in Computer Science*, pages 3–28. Springer, Berlin, 1997.
- [220] J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart, 1998. <http://www.mia.uni-saarland.de/weickert/Papers/book.pdf>.
- [221] J. Weickert. Coherence-enhancing diffusion filtering. *International Journal of Computer Vision*, 31(2/3):111–127, April 1999.
- [222] J. Weickert. Coherence-enhancing shock filters. In B. Michaelis and G. Krell, editors, *Pattern Recognition*, volume 2781 of *Lecture Notes in Computer Science*, pages 1–8. Springer, Berlin, 2003.
- [223] J. Weickert and T. Brox. Diffusion and regularization of vector- and matrix-valued images. In M. Z. Nashed and O. Scherzer, editors, *Inverse Problems, Image Analysis, and Medical Imaging*, volume 313 of *Contemporary Mathematics*, pages 251–268. AMS, Providence, 2002.

- [224] J. Weickert, C. Feddern, M. Welk, B. Burgeth, and T. Brox. PDEs for tensor image processing. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, Mathematics and Visualization, pages 399–414. Springer, Berlin, 2006.
- [225] J. Weickert and H. Hagen, editors. *Visualization and Processing of Tensor Fields*. Mathematics and Visualization. Springer, Berlin, 2006.
- [226] J. Weickert and C. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264, 2001.
- [227] J. Weickert and C. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264, December 2001.
- [228] J. Weickert and M. Welk. Tensor field interpolation with PDEs. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, Mathematics and Visualization, pages 315–325. Springer, Berlin, 2006.
- [229] Y. Weiss, E.P. Simoncelli, and E.H. Adelson. Motion illusions as optimal percepts. *Nature Neuroscience*, 5(6):598–604, May 2002.
- [230] M. Werlberger, T. Pock, and H. Bischof. Motion estimation with non-local total variation regularization. In *Proc. of the 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2464–2471, 2010.
- [231] E. T. Whittaker. On a new method of graduation. *Proceedings of the Edinburgh Mathematical Society*, 41:63–75, 1923.
- [232] N. Wiest-Daesslé, S. Prima, P. Coupé, S. P. Morissey, and C. Barillot. Non-local means variants for denoising of diffusion-weighted and diffusion tensor MRI. In N. Ayache, S. Ourselin, and A. Maeder, editors, *MICCAI 2007, Part II*, volume 4792 of *Lecture Notes in Computer Science*, pages 344–351. Springer, Berlin Heidelberg, 2007.
- [233] M. H. F. Wilkinson and J. B. T. M. Roerdink, editors. *Mathematical Morphology and Its Application to Signal and Image Processing*, volume 5720 of *Lecture Notes in Computer Science*. Springer, Berlin, 2009.
- [234] G. Winkler. *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods*. Springer, second edition, 2003.
- [235] G. Winkler. *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods: A Mathematical Introduction*. Springer, 2006.

- [236] G. Winkler, V. Aurich, K. Hahn, and A. Martin. Noise reduction in images: Some recent edge-preserving methods. *Pattern Recognition and Image Analysis*, 9(4):749–766, 1999.
- [237] L. P. Yaroslavsky. *Digital Picture Processing*. Springer, New York, 1985.
- [238] E. Zeidler. *Nonlinear Functional Analysis and Applications I: Fixed-Point Theorems*. Springer, New York, 1986.
- [239] D. Zhou and C. Burges. High-order regularization on graphs. In *Proc. of the 6th International Workshop on Mining and Learning with Graphs*, Helsinki, Finland, 2008.
- [240] D. Zhou and B. Schölkopf. A regularization framework for learning from graph data. In *Proc. ICML Workshop on Statistical Relational Learning and Its Connections to Other Fields*, pages 132–137, 2004.
- [241] D. Zhou and B. Schölkopf. Regularization on discrete spaces. In W. G. Kropatsch, R. Sablatnig, and A. Hanbury, editors, *Pattern Recognition*, volume 3663 of *Lecture Notes in Computer Science*, pages 361–368. Springer, Berlin, September 2005.
- [242] S. Zimmer, S. Didas, and J. Weickert. A rotationally invariant block matching strategy improving image denoising with non-local means. In A. Foi and A. Gotchev, editors, *Proc. International Workshop on Local and Non-Local Approximation in Image Processing*, pages 135–142, 2008.