
Ph.D. Jury

Private defense: Friday May 13th 2005, 14h30 in TELIN meeting room, Technicum, Sint-Pietersnieuwstraat 41, 9000 Gent, Belgium.

Public defense: Friday June 10th 2005 at 16h00 in Jozef Plateau room, Jozef Plateau building, Jozef Plateaustraat 22, 9000 Gent, Belgium.

- Prof. dr. ir. W. Philips (promoter), Faculty of Engineering - dept. TELIN, Ghent University
- Prof. dr. ir. P. Kiekens, (chairman), Pro-dean Faculty of Engineering, Ghent University
- Dr. ir. A. Pižurica, (secretary), Faculty of Engineering - dept. TELIN, Ghent University
- Prof. dr. ir. J. Portilla, Dept. of Computer Science and Artificial Intelligence, Universidad de Granada, España
- Dr. ir. H. Van der Voort, Scientific Volume Imaging, Hilversum, Nederland
- Prof. dr. ir. H. Steendam, Faculty of Engineering - dept. TELIN, Ghent University
- Prof. dr. I. Lemahieu, Faculty of Engineering - dept. ELIS, Ghent University
- Prof. dr. ir. Y. D'Asseler, Faculty of Engineering - dept. ELIS, Ghent University
- Prof. dr. P. Van Oostveldt, Faculty of Bio-engineering - dept. Molecular Biotechnology, Ghent University

I would like to thank all members of the jury for the valuable feedback and for correction many mistakes and suggestions for clearer explanations of some parts of the text.

Acknowledgements

This work would never have been possible without help and support of friends, colleagues at the TELIN department of Ghent University and last but not least, my family (especially my father).

I would like to thank in particular:

- *my promoter, prof. dr. ir. Wilfried Philips for his everlasting support and patience with me during those five years he guided my work;*
- *prof. dr. ir. Javier Portilla for his support and patience with me during our collaboration in Granada and the time since then;*
- *dr. ir. Hans van de Voort, for raising my interest in the subject, and providing valuable advice and feedback;*
- *dr. ir. Aleksandra Pižurica, for guidance, useful discussions and leading to some of my ideas through discussions and leading to some of my ideas through her own work;*
- *dr. ir. Diane S. Lidke, for providing a set of biological test images, useful suggestions and feedback about the processed images.*
- *prof. dr. Patrick Van Oostveldt, for providing other biological test images, interesting discussions and feedback about the processed images.*
- *the Institute for the Promotion of Innovation through Science and technology in Flanders (IWT Vlaanderen) for financing this research with a scholarship for my PhD.*

Also I would like to thank dr. ir. Peter Verveer and dr. ir. Geert van Kempen, who kindly answered related questions by e-mail.

For proofreading certain chapters, I also would like to thank lic. Geert Vernaeve and ir. Bruno Huysmans.

For the L^AT_EX structure of this thesis, the framework of the GLMS course [GLMS, 2002] was used. This course was written to be an introduction of GNU/Linux in all its aspects. That course is distributed under the Free Documentation License, and is available at <http://savannah.nongnu.org/projects/glms/>.

Preface

According to the regulations of the Faculty of Engineering of Ghent University, this thesis contains a Dutch summary of about 50 pages, followed by the actual thesis in English (about 170 pages, not counting the bibliography). This thesis contains three introductory chapters. These chapters cover the following subjects:

- optical image formation;
- classical restoration techniques;
- multiresolution principles in image processing.

The reason that these topics are described so extensively, is because not everybody is familiar with all three topics together. For some people I hope to reach with this thesis, one or more of these chapters are probably redundant. For others, I hope it provides a first introduction in the covered topics.

We also we would like to explain how we used references in this thesis. Each reference in the text contains two parts: the author(s) and the year of publication. Sometimes, we refer to texts or figures from websites. In that case, the year was replaced by the key “WWW”, because most websites don’t mention the year when the content was published. So, when a reference in this thesis occurs which looks like:

[van der Voort and Strasters, 1995, UltraSharpen, WWW],

it means that we refer to two items: the first item is a paper from van der Voort and Strasters, which was published in the year 1995; the second item was a website, in this case of Ultrasharp, published on the WWW.

When a paper has more than two authors, only the name of the first author is mentioned, followed by “et al.”

Also, we only referenced to papers and material that was actually available to us. We didn’t use references that a lot of other authors refer to when we didn’t have access to those papers, books, ...

Contents

Samenvatting	1
S.1 Algemene situering	2
S.1.1 Toepassingen van beeldrestauratie	4
S.1.2 Doel en indeling van deze thesis	5
S.1.3 Overzicht van conceptuele bijdragen	6
S.1.4 Overzicht van publicaties	7
S.2 Beeldvorming	7
S.2.1 Geschiedenis: van lens tot microscoop	7
S.2.1.1 Van lens tot microscoop	7
S.2.1.2 Van klassieke microscoop tot fluorescentiemicroscoop	9
S.2.1.3 Van fluorescentiemicroscoop tot confocale microscoop	9
S.2.2 Een intuïtieve benadering tot de beperkingen van optische beeldvorming	10
S.2.3 Van Abbe naar Fourieroptica	11
S.2.4 Beeldvorming bij fluorescentiemicroscopie	12
S.2.4.1 Algemene fluorescentiemicroscoop	12
S.2.5 Confocale microscoop	12
S.2.6 Klassieke microscoop	12
S.3 Klassieke restauratietechnieken	13
S.3.1 Invers filter en Wienerfilter	13
S.3.1.1 Inverse filter	13
S.3.1.2 Wienerfilter	14
S.3.2 Bayesiaanse restauratie	14
S.3.2.1 Gaussiaans ruismodel	15
S.3.2.2 Poisson ruismodel	17
S.4 Multiresolutie en beeldverwerking	18
S.4.1 Algemene inleiding	18
S.4.2 Multiresolutie, wavelets en stuurbare piramides	19
S.4.2.1 Van Fouriertransformatie naar wavelets	19
S.4.3 De stuurbare piramide	22
S.4.3.1 Basisideeën achter de stuurbare piramide	22
S.4.3.2 Implementatiedetails	23

S.4.4	Toepassingen	23
S.4.4.1	Kwadratuurfilters en fase-onafhankelijkheid	23
S.4.4.2	Oriëntatie-analyse en randdetectie	24
S.4.4.3	Hoekdetectie	24
S.4.4.4	Beeldsegmentatie	25
S.4.4.5	Ruisreductie	25
S.4.4.6	Beeldrestauratie	27
S.4.5	3D multischalige transformaties	29
S.4.5.1	3D stuurbare piramide	29
S.4.5.2	Andere 3D multischalige transformaties	30
S.5	Twee methoden voor parametrische PSF schatting	30
S.5.1	Inleiding	30
S.5.2	Wave-Track: tracking van de waveletcoëfficiënten door- heen de schalen	31
S.5.2.1	Inleiding	31
S.5.2.2	Principe	32
S.5.2.3	Robuustheid tegen ruis	33
S.5.2.4	Evaluatie van Wave-Track	34
S.5.2.5	Toekomstig werk: uitbreiding van Wave-Track naar niet-isotrope PSF's	35
S.5.3	Sparse-Max: waveletgebaseerde ijlheidsmaximalisatie	35
S.5.3.1	Principe van Sparse-Max	35
S.5.3.2	Evaluatie en vergelijking	37
S.5.3.3	Conclusie	38
S.5.4	Toepassingen van vervagingsschatting	38
S.5.4.1	Beeldrestauratie	38
S.5.4.2	Autofocus	39
S.5.4.3	Diepte uit defocus	39
S.6	SPERRIL: schatting en restauratie van confocale beelden	40
S.6.1	Inleiding	40
S.6.2	Overzicht van het algoritme	41
S.6.2.1	Ruisreductiestap	41
S.6.2.2	Vervagingsschatting	42
S.6.2.3	Verscherpingsstap	43
S.6.2.4	Stopcriterium	44
S.6.2.5	Experimentele resultaten	44
S.6.2.6	Conclusies	47
S.7	Multiresolutie-GNC gebaseerde beeldrestauratie	47
S.7.1	Markov Random Velden	48
S.7.1.1	Energieminimalisatie	48
S.7.2	Multiresolutie MRV beeldrestauratie	49
S.7.2.1	Bestaande methoden	49
S.7.2.2	Wat achtergrond in verband met MMRV beeld- restauratie	50
S.7.3	Voorgestelde methode	51

S.7.3.1	Mogelijke randindicatoren	52
S.7.4	Conclusies en toekomstig onderzoek	53
S.8	Conclusies	53
1	General introduction	3
1.1	Image formation and optical limitations	3
1.2	Applications of image restoration	5
1.3	Scope of this thesis	7
1.4	Summary of contributions	8
1.5	Summary of publications	8
2	Image formation	11
2.1	A brief history: from single lens to confocal microscope	11
2.1.1	From lens to microscope	11
2.1.2	From classical microscope to fluorescence microscope . .	13
2.1.3	From fluorescence microscope to confocal microscope . .	15
2.2	An intuitive approach to the limitations of optical image formation	19
2.2.1	The basics	19
2.2.2	Ernst Abbe: the bridge between classical optics and Fourier optics	23
2.3	Fluorescence microscopy: A particular case of imaging	27
2.3.1	Three-dimensional PSF of an objective lens	28
2.3.2	General fluorescence microscope	28
2.3.3	Confocal microscope	29
2.3.4	Widefield microscope	30
3	Classical restoration techniques	37
3.1	Inverse filter and Wiener filter	38
3.1.1	Inverse filter	38
3.1.2	Wiener filter	39
3.2	Bayesian restoration	42
3.2.1	Gaussian noise model	44
3.2.1.1	Tikhonov-Miller restoration versus Wiener filter	44
3.2.1.2	Intermezzo: gradient-based minimization methods	46
3.2.1.3	Iterative Constrained Tikhonov-Miller	48
3.2.2	Poisson noise model	52
3.2.2.1	Richardson-Lucy restoration	52
3.2.2.2	Richardson-Lucy illustrated	58
3.2.2.3	Classical regularization of Richardson-Lucy . .	60
4	Multiresolution and image processing	65
4.1	General introduction	65
4.2	Multiresolution, wavelets and steerable pyramids	66
4.2.1	From the Fourier transform to wavelets	66
4.2.2	The steerable pyramid.	76

4.2.2.1	General concept of steerable pyramids.	76
4.2.2.2	An example of steerability.	77
4.2.2.3	Implementation issues of the steerable pyramid.	79
4.3	Applications	87
4.3.1	Image analysis	87
4.3.1.1	Quadrature filters and phase independence	87
4.3.1.2	Orientation analysis and edge detection	91
4.3.1.3	Corner detection	94
4.3.1.4	Image segmentation	95
4.3.2	Noise reduction	97
4.3.2.1	Gaussian noise	97
4.3.2.2	Poisson noise	102
4.3.3	Image restoration	104
4.3.3.1	Hybrid techniques: only regularization in the wavelet domain	105
4.3.3.2	Denoising and deblurring in the wavelet domain	105
4.4	3D multiscale transforms	106
4.4.1	3D steerable pyramid	107
4.4.2	Other 3D multiscale transforms	109
5	Two methods for parametric PSF estimation in images	113
5.1	Introduction	113
5.2	Wave-Track: tracking wavelet coefficients across scales	115
5.2.1	Introduction	115
5.2.2	Principle	116
5.2.2.1	Theoretical background	116
5.2.2.2	Robustness to noise	118
5.2.3	Evaluation of Wave-Track	121
5.2.4	Future work: extension of Wave-Track to anisotropic PSF's	122
5.3	Sparse-Max: wavelet-based sparseness maximization.	124
5.3.1	Principle of Sparse-Max	124
5.3.2	Evaluation and comparison	128
5.3.3	Conclusion	130
5.4	Applications of blur estimation	133
5.4.1	Image restoration	133
5.4.2	Autofocus	134
6	SPERRIL: Estimation and restoration of confocal images	139
6.1	Introduction	139
6.2	Outline of the algorithm	140
6.2.1	Noise reduction step	141
6.2.2	Blur estimation	143
6.2.3	Deblurring step	144
6.2.4	Stopping criterion	145
6.2.5	SPERRIL as RL with a prior?	145
6.3	Experimental results	146

6.3.1	Synthetic image	146
6.3.2	Biological experiments	150
6.3.2.1	Colocalization analysis: what and why?	150
6.3.2.2	Experimental setup	151
6.3.3	Typical running times	157
6.3.4	Some more visual results	159
6.4	Conclusions	166
7	Multiresolution Markov Random Field image restoration	169
7.1	Markov Random Fields	169
7.1.1	An example: the Ising model	172
7.1.2	Energy minimization	172
7.1.2.1	The Metropolis algorithm	173
7.1.2.2	Graduated Non-Convexity Minimization	174
7.2	Multiresolution MRF image restoration	175
7.2.1	Existing methods	175
7.2.2	Background on MMRF image restoration	176
7.3	Proposed method	177
7.3.1	Possible edge indicators	177
7.3.2	Results	178
7.4	Conclusions and future work	181
8	Conclusions	183
8.1	Summary of the conclusions	183
8.2	Future research	184

Symbols and acronyms used in this thesis

Symbols

Δ_i	difference in case i
$f(x, y)$	ideal image
f_i	Fourier frequency of event i
$g(x, y)$	degraded image
$h(x, y)$	Point Spread Function
κ	Kurtosis
λ_{cls}	Regularization parameter determined with CLS method (in context of classical image restoration)
λ_{ex}	Excitation wavelength (in context of image formation)
λ_{em}	Emission wavelength (in context of image formation)
λ_{gcv}	Regularization parameter determined with GCV method (in context of classical image restoration)
λ_{ml}	Regularization parameter determined with ML method (in context of classical image restoration)
μ	Expectancy Value
$N(\dots)$	Noise process
r_{focal}	radius of the focal spot
$R_{Pearson}$	Pearson correlation coefficient
$R_{Overlap}$	Overlap correlation coefficient
σ	standard deviation
σ_{blur}	standard deviation of the PSF
σ_{noise}	standard deviation of the noise

Acronyms

ADC	Analog to Digital Converter
AIF	Adaptive Interaction Function
AR	Auto Regressive

ARMA	Auto Regressive Moving Average
BLS	Bayes Least Square
CCD	Charge-Coupled Device
CG	(1) Center of Gravity (in context of blur estimation) (2) Conjugate Gradients (in context of ICTM)
CLS	Constrained Least Squares
COC	Circle of Confusion
CWT	Continuous Wavelet Transform
DWT	Discrete Wavelet Transform
EGF	Epidermal Growth Factor
ErbB1	a family member of Epidermal growth factor receptor proteins
EM	Expectation Maximization
FFT	Fast Fourier Transform
FITC	Fluorescein isothiocyanate (a fluorescent dye)
FWHM	Full Width at Half Maximum
GCV	Generalized Cross Validation
GFM	General Fluorescence Microscope
GFP	Green Fluorescent Protein (a fluorescent dye)
GNC	Graduated Non-Convexity
GRF	Gibbs Random Field
GSM	Gaussian Scale Mixture
HB	Highpass band
ICA	Independent Component Analysis
ICTM	Iterative Constrained Tikhonov-Miller
IID	Independent, Identically Distributed
LB	Lowpass band
MAP	Maximum A Posteriori
MEM	Maximum Entropy Methods
ML	Maximum Likelihood
MMRF	Multiscale Markov Random Field
MMSE	Minimum Mean Square Error
MPLE	Maximum Penalized Likelihood Estimation
MRF	Markov Random Field
OTF	Optical Transfer Function
PCA	Principal Component Analysis
PDF	probability density function
PMT	Photo Multiplier Tube
PSF	Point Spread Function
PSNR	Peak Signal to Noise Ratio
QD	Quantum Dots (=a fluorescent dye)
RA	Radially Averaged
RL	Richardson-Lucy
SNR	Signal to Noise Ratio
SPERRIL	Steerable Pyramid based Estimation and Restoration using Richardson-Lucy

CONTENTS

xiii

STFT	Short Term Fourier Transform
SURE	Stein Unbiased Risk Estimator
WF	WideField

Nederlandstalige samenvatting

The best way to predict the future is to invent it.
—Alan Kay

Vertaling van enkele wetenschappelijke termen

Hier geven we de gebruikte vertalingen van enkele typische sleutelwoorden die te maken hebben met ons werk. Woorden waarvoor we geen goede Nederlandse vertaling vonden, hebben we in hun Engelse vorm gebruikt in deze thesis. Voor de motivatie van de inleidende hoofdstukken en voor de conventies in verband met bibliografiereferenties verwijzen we naar “Preface” van deze thesis.

acceptance angle	ontvankelijkheidshoek
bias	
clique	kliek
conjugate gradient	toegevoegde gradiënt
constrained	beperkt, begrensd door randvoorwaarden
deblurring	verscherping
ground truth	grondwaarheid, referentiedata
likelihood	aannemelijkheid
probability	waarschijnlijkheid
rate	tempo
ringing	
smoothing	vergladding
sparseness	ijlheid
steepest descent	steilste helling
subsample	onderbemonsteren, decimeren
support	drager
thresholding	

S.1 Algemene situering

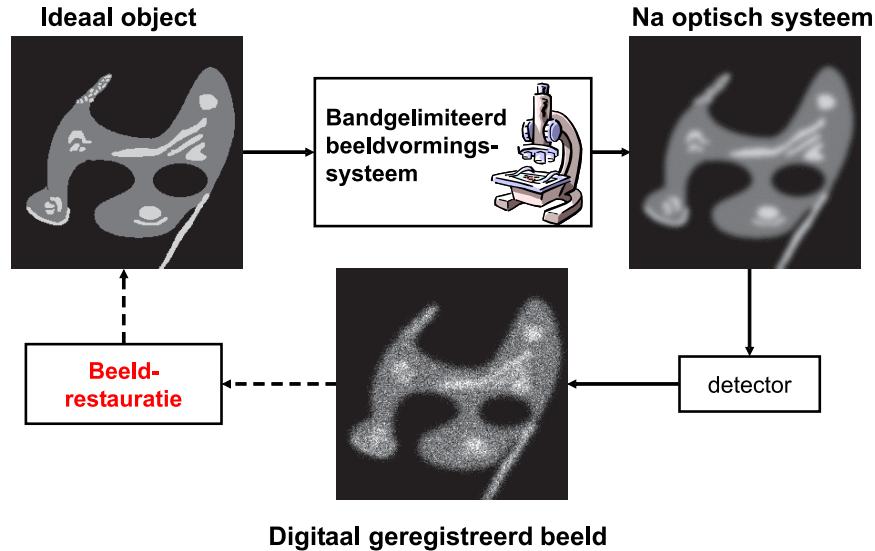
Mensen zijn sinds hun ontstaan uit op verkenning en ontdekkingen: ze verkennen planeten in het heelal, beklimmen de hoogste bergen, dalen af tot de bodem van de oceaan, ... Kortweg: mensen willen altijd weer verder, voorbij de grenzen van het huidige kunnen en kennen.

De apparaten die de mens bouwt, laten heel wat verborgen informatie zien. Daarenboven registreren ze ook nog informatie die niet direct voor de menselijke waarnemer zichtbaar is, maar toch in het beeld aanwezig is. Onderzoekers willen echter steeds tot op de limiet van hun instrumenten waarnemingen verrichten en wensen ook deze verborgen informatie zichtbaar te maken zodat de waarnemer ze kan interpreteren. Rekening houdend met de processen die plaatsvinden in de apparatuur bij het waarnemingsproces is het mogelijk om een model op te stellen welke informatie precies verloren gaat en welke informatie aan het oog onttrokken wordt, maar aanwezig blijft in het beeld. Daardoor wordt het ook mogelijk om methoden uit te denken om deze verborgen informatie te recupereren.

Hier in deze thesis zullen we ons toespitsen op optische systemen (zoals microscopen) verbonden met een detector. Elk optisch systeem heeft echter beperkingen: het beeldt een ideale puntbron af als een zekere ruimtelijke verdeling van licht doordat licht een golfkarakter heeft. Deze afbeelding van een ideale puntbron is de zogenaamde **PuntSpreidingsFunctie** (PSF). Omdat elk punt van het beeld op deze manier volgens de PSF uitgesmeerd afgebeeld wordt, ontstaat beeldvervaging, zelfs met optimaal afgestelde optische systemen. Daardoor heeft elk optisch systeem een resolutielimiet; natuurlijk wordt de situatie alleen maar erger wanneer ook de afstelling van de optische elementen niet optimaal is.

Om het beeld van het optische systeem op te slaan voor latere studie, analyse of bewerking, wordt het beeld vastgelegd door een detectorsysteem, zoals een CCD camera. Daarbij treedt ruis op. Deze ruis kan worden opgedeeld in volgende drie categorieën:

- **fotonruis:** door het deeltjeskarakter van licht zijn beelddetectors in feite apparaten die “fotonen tellen,” zeker in omstandigheden van lage lichtintensiteiten. De photonen die op de detector invallen, worden omgezet in ladingsdragers, zoals elektronen. Dit detectie- / omzettingsproces is echter niet deterministisch maar stochastisch, en kan dus best beschreven worden in termen van toevalsgrootheden.
- **sensormuis:** vat alle bronnen van ruis samen in de detector, waaronder de donkerstroom (ladingen die worden gegenereerd zonder dat er een signaal aanwezig is), de ladingen die geregistreerd worden op een verkeerde positie op de detector, variaties in de gevoeligheid van detectorelement tot detectorelement, ...



Figuur S.1: Blokdiagram van het beeldvormingsproces en de overeenkomstige beelddegradatie. Links laten we een afbeelding zien van een synthetische cel. Rechts hebben we het beeld van het object dat enkel vervaagd is door het optische systeem. Onderaan wordt het beeld getoond dat we uiteindelijk te zien krijgen.

- **uitlees- of versterkingsruis:** doordat het signaal dat van de ladingsdragers afkomstig is verder moet worden versterkt en gediscretiseerd naar een geheel getal dat overeenkomt met de intensiteit (door middel van een Analoog naar Digitaal Convertor (ADC)), treedt er nog extra ruis op in het beeld dat we uiteindelijk registreren.

We illustreren deze cyclus in Figuur S.1: een ideaal object (in feite het beeld van een object onder ideale opname-omstandigheden) is te zien bovenaan links. Dit ideaal object is vervaagd door gebruik te maken van een optisch (en dus bandbeperkt) beeldvormingssysteem. Door het digitaal registreren van het beeld komt er nog eens ruis bij, waardoor we uiteindelijk een vervaagd beeld met ruis verkrijgen.

Om samen te vatten: het digitaal beeld dat uiteindelijk opgenomen wordt, is in feite vervaagd en verstoord door ruis, waardoor er bepaalde informatie verloren is en waardoor andere informatie verborgen is voor de waarnemer, maar wel nog aanwezig is [Van der Voort, WWW]:

- bepaalde details zijn soms verborgen in de ruis, of ze worden gemaskeerd door andere beeldkenmerken;
- sommige artefacten kunnen de waarnemer in de war brengen;

- er kan informatie aanwezig zijn in impliciete vorm, wat wil zeggen dat die enkel kan worden gerecupereerd door middel van extra *a priori* kennis.

Het doel van beeldrestauratie is dus om zo goed mogelijk te compenseren voor het degradatieproces, om zo een beeld te bekomen waarin de “verborgen” informatie zichtbaar is, waardoor zowel interpretatie door een menselijke waarnemer als nabewerkingen op het beeld zoals segmentatie, objectherkenning, ... worden vergemakkelijkt.

S.1.1 Toepassingen van beeldrestauratie

Deze paragraaf is geïnspireerd op [Molina et al., 2001]. Beeldrestauratie werd oorspronkelijk ontwikkeld tijdens het Amerikaanse ruimtevaartprogramma in de jaren '50, waarbij onbemande ruimtesondes voor het eerst foto's maakten van de Aarde, de Maan en Mars. Deze beelden legden nooit eerder geziene details vast, maar de kwaliteit van deze beelden was zwaar aangetast door de moeilijke omstandigheden waarin ze waren opgenomen, zoals trillingen, slecht richten van het toestel, vervaging door beweging van de sonde tijdens de opname, ...

Gelet op de hoge kosten van deze beelden was de vraag om uit deze beelden toch zoveel mogelijk informatie te halen voor de hand liggend. Voor dat doel werden ééndimensionale technieken voor signaalrestauratie (zoals het Wiener-filter) uitgebreid naar twee dimensies. Het resultaat was dat de beelden leken opgenomen te zijn met apparatuur uit de toekomst; zoveel verschil was er tussen ruwe en gerestaureerde beelden. Sinds hun ontstaan worden technieken voor beeldrestauratie en -reconstructie algemeen gebruikt in disciplines als medische tomografie, seismologie, magnetische resonantie beeldvorming, enz. In optische sterrenkunde werden deze technieken tot voor kort beschouwd als overbodige luxe, maar daar kwam verandering in toen er een fout werd ontdekt in het optische systeem van de toen pas gelanceerde Hubble Space Telescope (HST) in 1990. Er werden toen twee conferenties georganiseerd om te zien wat de mogelijkheden waren om de HST beelden te restaureren [Hanisch and White, 1993].

Een andere historische toepassing van beeldrestauratie was de restauratie van de beelden die werden genomen tijdens de moord op John F. Kennedy. De beelden die toen werden genomen zijn zeer onscherp door de verwarring en de shock die er toen heerste onder de toeschouwers. Later werden er wel pogingen ondernomen om deze beelden te verscherpen. In het bijzonder was er veel belangstelling voor de restauratie van de zogenaamde *Zapruder* film.

Een tweede wetenschappelijke toepassing waar beeldrestauratie wordt gebruikt is fluorescentiemicroscopie [van der Voort and Strasters, 1995, Van Kempen et al., 1997, Verveer et al., 1999]. Een algoritme dat vaak wordt gebruikt om dit soort beelden te restaureren, is het *Expectation-Maximization* algoritme en werd voor het eerst op dit soort beelden toegepast in [Holmes, 1988]. Dit algoritme wordt in detail besproken in Sectie S.3.2.2. Doordat computers steeds maar krachtiger werden, werd het ook mogelijk om

steeds grotere datasets te restaureren, zoals driedimensionale beelden (zoals beelden opgenomen met een confocale microscoop; zie verder in deze thesis) en zelfs 4D (3D datasets in de tijd).

In feite kan beeldrestauratie worden toegepast op elke situatie waarin een digitaal beeld wordt verkregen met een optisch systeem. Tegenwoordig proberen ook steeds meer thuisgebruikers de kwaliteit van hun digitale foto's te verbeteren. Daarvoor zijn er ook reeds softwarepakketten beschikbaar: [Uman Software, WWW, Nik Multimedia, Inc., WWW, UltraSharpen, WWW], om er maar enkele te noemen.

S.1.2 Doel en indeling van deze thesis

In deze thesis bestuderen we een aantal algoritmen voor het restaureren van beelden opgenomen met een confocale microscoop. In het bijzonder onderzoeken we hoe de principes van multiresolutie kunnen bijdragen door verbeterde beeldkwaliteit na het restauratieproces. Beeldverwerking gebaseerd op multiresolutieconcepten heeft reeds aangetoond dat het beelden en signalen in het algemeen beter kan modelleren. Multiresolutieconcepten zijn dus een krachtig hulpmiddel om *a priori* kennis op te leggen. In wat nu volgt, worden in het kort de verschillende onderdelen van de thesis samengevat, met een structuur parallel aan de structuur van de Engelstalige tekst.

In deel S.2 wordt het beeldvormingsproces van een confocale microscoop besproken, waarbij de verschillende oorzaken van beelddegradatie in dit proces worden toegelicht.

Deel S.3 behandelt de meest gebruikte klassieke methoden voor het restaureren van confocale beelden, afgeleid van de basisprincipes. Het invers filter en het Wienerfilter worden besproken om de aandacht te vestigen op de belangrijkste problemen die optreden bij beeldrestauratie. Daarna worden meer geadvanceerde restauratietechnieken zoals Tikhonov-Miller en Richardson-Lucy behandeld, evenals hoe regularisatie *a priori* kennis in het restauratieproces invoert. Deze algoritmen worden behandeld in het Bayesiaanse kader.

Deel S.4 introduceert het multiresolutieprincipe, en geeft enkele toepassingen ervan in beeldverwerking, zoals ruisonderdrukking en beeldsegmentatie. In het bijzonder spitsen we ons toe op hoe de principes van multiresolutie kunnen worden toegepast om *a priori* kennis op te leggen aan het beeld tijdens het restauratieproces.

In deel S.5 worden twee originele methoden voor parametrische PSF schatting beschreven, die kunnen worden gebruikt wanneer de precieze vervaging niet bekend is. Daarbij gaan we uit van een parametrisch vervagingsmodel, waarbij de vereenvoudigde veronderstelling gemaakt wordt dat de PSF kan worden uitgedrukt als een functie van slechts enkele vrije parameters. In ons geval zullen we ons zelfs beperken tot één parameter, zoals een Gaussiaanse functie (die enkel afhangt van σ_{blur}) of de out-of-focus PSF (een uniforme schijf is, die enkel

afhangt van de straal van deze schijf). Onze technieken schatten de parameter van dit soort PSF's op basis van multiresolutieconcepten.

Deel S.6 gaat dieper in op het SPERRIL algoritme, wat een originele bijdrage van dit doctoraat is. Dit algoritme schat eerst de beelddegradatie en restaureert daarna het beeld. We geven ook een evaluatie van het algoritme op basis van experimenten met zowel synthetische beelden als met biologische beelden. Voor dat laatste soort beelden werd gebruik gemaakt van colocalisatie-analyse voor de evaluatie.

Deel S.7 bespreekt nog een andere bijdrage van dit doctoraat, namelijk een methode voor multiresolutiegebaseerde beeldrestauratie. Deze methode gebruikt multiresolutie Markov Random Velden, waarvan de energiefunctie geminimaliseerd wordt door de *Graduated Non-Convexity* (GNC) methode.

Deel S.7 bespreekt nog een andere bijdrage van dit doctoraat, namelijk een methode voor multiresolutiegebaseerde beeldrestauratie. Deze methode gebruikt multiresolutie Markov Random Velden, waarvan de energiefunctie geminimaliseerd wordt door de *Graduated Non-Convexity* (GNC) methode.

Deel S.8 geeft nog eens een kort overzicht van de thesis, en vat de conclusie samen.

S.1.3 Overzicht van conceptuele bijdragen

In dit doctoraat werd bestudeerd hoe multiresolutieconcepten kunnen bijdragen tot de restauratie van confocale beelden.

Ten eerste werden op basis van deze multiresolutieconcepten twee algoritmen ontwikkeld (Wave-track en Sparse-max) voor het schatten van de vervagingsparameter als we ervan uitgaan dat de PSF kan worden gemodelleerd door een parametrische uitdrukking.

Ten tweede werden er twee nieuwe restauratie-algoritmen voorgesteld, namelijk SPERRIL en de methode op basis van Markov Random Velden.

Tijdens dit onderzoek heb ik ook vier afstudeerwerken over beeldrestauratie begeleid:

- Tim Stevens (2000-01): “*Niet-blinde deconvolutie van 3D beelden opgenomen met een confocale microscoop*”
- Sofie De Wolf (2001-02): “*Kwantitatieve vergelijking van blinde deconvolutie-algoritmen*”
- Anthony Dewulf (2001-02): “*Waveletgebaseerde regularisatie van deconvolutie-algoritmen*”
- Hans Yperman (2004-05): “*Restauratie van beelden met een plaatsafhankelijke vervaging*”

S.1.4 Overzicht van publicaties

Tot nu toe heeft dit onderzoek geleid tot een populariserend tijdschriftartikel [Rooms and Philips, 2002a] en een tijdschriftartikel [Rooms et al., 2005].

Ook werden er acht gereviewde artikels in internationale conferenties met mezelf als eerste auteur gepubliceerd: [Rooms et al., 2004], [Rooms et al., 2003b], [Rooms et al., 2003c], [Rooms and Philips, 2002b], [Rooms et al., 2002c], [Rooms et al., 2002b], [Rooms et al., 2001a], [Rooms et al., 2001b], [Rooms et al., 2000b]

Daarnaast verschenen ook negen abstracten (≤ 2 pagina's) die niet inbegrepen zijn in bovenstaande lijst [Rooms and Philips, 2004], [Rooms et al., 2003e], [Rooms et al., 2003a], [Rooms et al., 2003d], [Rooms et al., 2002a], [Rooms et al., 2001c], [Rooms and Philips, 2001], [Rooms and Philips, 2000], [Rooms et al., 2000a].

Ik droeg ook bij tot werk van anderen, wat resulteerde in volgende artikels waar ik co-auteur ben [De Smet et al., 2004], [Stippel et al., 2001], [Stippel et al., 2000].

S.2 Beeldvorming

Deze sectie is bedoeld om duidelijk te maken wat de oorzaken zijn waardoor er bepaalde informatie verloren gaat en andere informatie gemaskeerd wordt. Daarom gaan we dieper in op hoe beeldkwaliteit wordt gewonnen door het optische beeldvormingsproces en hoe beeldkwaliteit verloren gaat. Verder wordt dit kwaliteitsverlies ook gemodelleerd, om het (zo goed mogelijk) te kunnen compenseren.

S.2.1 Geschiedenis: van lens tot microscoop

Om te beginnen geven we een korte geschiedenis van de microscoop. Het gedeelte over de geschiedenis van de klassieke microscoop is gebaseerd op [Bellis, WWW]; de onderstaande tekst over de geschiedenis en de werking van de confocale microscoop en de fysica van de beeldvorming daarentegen is gebaseerd op [van Kempen, 1999, Verveer et al., 1999, van der Voort, 1989, Plášek and Reischig, 1998, Lichtman, 1994].

S.2.1.1 Van lens tot microscoop

Lange tijd geleden raapte iemand toevallig een convex stuk kristal op en merkte op dat wanneer men objecten doorheen dit kristal keek, deze objecten groter leken; het gebruik van de eerste lens was een feit. Ook stelde men vast dat zulk kristal het licht en de warmte van de zon kon focussen in één enkel punt en zo

hout, papier, ... in brand kon steken. Zulke "brandglazen" of "vergrootglazen" werden reeds vernoemd door Seneca en Plinius de Oudere (Romeinse filosofen in de 1^{ste} eeuw na Christus). Maar blijkbaar werden ze niet echt courant gebruikt tot wanneer de bril werd uitgevonden aan het einde van de 13^{de} eeuw. De naam "lenzen" komt van de vorm, omdat die gelijkt op de vorm van zaden van de linze (een kleine erwrt, *Lens Culinaris* of *Lens Esculenta*).

Rond 1590 experimenteerden de twee Nederlandse opticiens Zacharias Janssen en zijn vader Hans met sets van lenzen in een cilindervormige buis. Ze ontdekten dat dit instrument dichtbij gelegen objecten sterker vergroot weergaf dan met enkelvoudige lenzen. Hun apparaat kan worden beschouwd als de voorloper van de samengestelde microscoop en de telescoop.

Antoni van Leeuwenhoek wordt als de vader van de microscoop beschouwd. Hij leefde in Nederland van 1632 tot 1723 en leerde met optiek omgaan in een textielwinkel waar de werknemers lenzen gebruikten om de draden in weefsels te tellen. Van Leeuwenhoek ontwikkelde nieuwe methoden voor het slijpen en polijsten van kleine lenzen met een grote kromming, die objecten tot 270x vergroot weergaven. Zijn lenzen waren de beste uit die tijd, en Antoni bouwde met deze lenzen microscopen, waarmee hij heel wat biologische ontdekkingen deed: hij zag en beschreef als eerste o.a. bacteriën, gistcellen en bloedlichaampjes. Hij wordt beschouwd als één van de grondleggers van wat later de celbiologie en de microbiologie zouden worden.

Robert Hooke leefde van 1635 tot 1703, en wordt beschouwd als de Engelse vader van de microscoop. Hij controleerde en bevestigde de ontdekkingen van Antoni van Leeuwenhoek met zelfgemaakte microscopen. Hij vertrok van Leeuwenhoeks ontwerp, om er dan eigen verbeteringen op aan te brengen.

Christiaan Huygens was in 1670 de eerste die de wetten van de geometrische optica afleidde uit de principes van de golftheorie. Huygens formuleerde een zeer belangrijk inzicht over de voortplanting van golven, wat nu bekend is als het principe van Huygens [Fitzpatrick, WWW]: *Elk punt van een golffront van licht kan worden beschouwd als een bron van secundaire golven, die voorwaarts uitdijen met de lichtsnelheid. Het nieuwe golffront is een oppervlak dat raakt aan al deze secundaire golven.* Met dit principe kon Huygens optische verschijnselen zoals breking en weerkaatsing van licht verklaren.

Ernst Abbe publiceerde zijn werk over de theorie van de microscoop in 1873. Hij was de eerste om een theoretische basis te leggen voor het ontwerp van microscopen, iets wat tot dan toe vooral gebeurde door *trial-and-error*. Hij formuleerde ook duidelijk het verschil tussen vergroting en resolutie, en schreef ook dat microscopen met een bepaalde objectief lens een maximale "nuttige" vergroting hebben; nog grotere vergroting noemde hij "lege vergroting." Gebaseerd op de golftheorie voor licht leidde hij zijn bekende formule af voor de resolutie van een microscoop.

S.2.1.2 Van klassieke microscoop tot fluorescentiemicroscoop

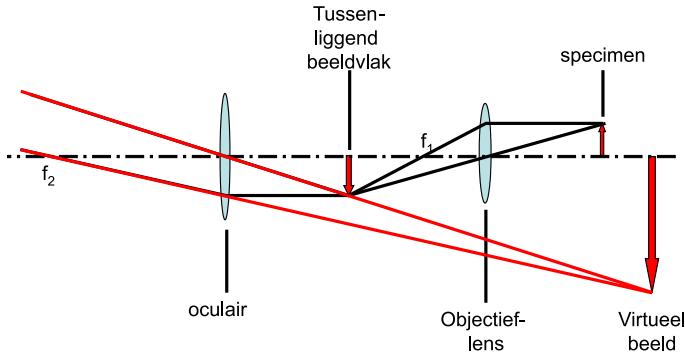
Niettegenstaande de enorme resolutiewinst die een microscoop biedt ten opzichte van het blote oog, zijn er toch enkele problemen:

- Het contrast in specimens is normaal gezien vrij laag. Dit kan worden verholpen door gebruik te maken van weefselafhankelijke contrastverhogende kleurstoffen om bepaalde onderdelen van het specimen in het specimen duidelijker te laten opvallen.
- Wanneer een klassieke lichtbron (met een breed spectrum) en klassieke kleurstoffen worden gebruikt, ontvangt de waarnemer licht dat wordt verstrooid vanuit het gebied dat hij wil bestuderen, maar ook licht dat in de nabijheid ervan wordt verstrooid; dat resulteert in daling van het contrast. Om daaraan te verhelpen worden fluorescerende kleurstoffen¹ of fluorochromen gebruikt om het contrast te verhogen. Deze fluorochromen worden geëxciteerd met licht van een bepaalde golflengte λ_{ex} en zenden licht uit op een andere, langere golflengte λ_{em} door de energiedissipatie die plaatsvond in de periode waarin het fluorochroom zich in de geëxciteerde toestand bevond (het verschil tussen deze twee golflengten wordt de *Stokes verschuiving* genoemd). Wanneer het specimen belicht wordt met één bepaalde golflengte (zoals met een laser), kan men het excitatielicht elimineren door middel van een filter, om zo enkel het emissielicht te laten bijdragen tot het beeld.
- Zelfs dan draagt emissielicht van gedeelten van het specimen boven en onder het focaalvlak bij tot het beeld van een enkele coupe, waardoor zogenaamde out-of-focus vervaging ontstaat.

S.2.1.3 Van fluorescentiemicroscoop tot confocale microscoop

Marvin Minsky is bekend als de vader van artificiële intelligentie, maar heeft nog een andere belangrijke verwezenlijking op zijn naam: hij vond namelijk de confocale microscoop uit in 1957 [Lichtman, 1994, van der Voort, 1989]. Dit revolutionaire type van microscoop liet namelijk toe om diepere lagen in een specimen te bestuderen zonder het specimen in dunne schijfjes te moeten snijden (wat een hele hoop werk was). In feite bracht Minsky slechts enkele eenvoudige maar belangrijke wijzigingen aan aan het ontwerp van een standaard lichtmicroscoop. Eerst reduceerde hij de verstrooiing van licht in het specimen door de belichting door de objectieflens te sturen, waardoor deze werd gefocuseerd in één punt in het specimen. Minsky's tweede verbetering was het plaatsen van een diafragma voor de detector, om het licht dat niet van het punt in focus afkomstig is, te elimineren bij het beeldvormingsproces. Deze twee aanpassingen zorgen ervoor dat een scherp beeld wordt bekomen van een

¹De namen fluorofoor (=fluorescerend bestanddeel) en fluorochroom (=fluorescerende kleurstof) worden door elkaar gebruikt voor deze stoffen.



Figuur S.2: Geometrische beeldvorming in een microscoop.

punt in het specimen. Door een horizontaal scanningsmechanisme toe te voegen kan een beeld van een dunne coupe van het specimen worden gemaakt; door daarenboven de focusdiepte op en neer te verplaatsen, kunnen coupes op verschillende dieptes in het specimen worden gemaakt, waardoor een 3D beeld bekomen wordt van het specimen.

S.2.2 Een intuïtieve benadering tot de beperkingen van optische beeldvorming

In [Plášek and Reischig, 1998] wordt er een intuïtieve uitleg gegeven over optische beeldvorming vanuit het oogpunt van de fysicus. Gebaseerd op dat artikel zullen we kort de beeldvorming in een microscoop bespreken, en de gevolgen ervan op de resolutie van het bekomen beeld.

De basiscomponenten van een microscoop zijn de *objectieflens* en het *oculair*. De rol van de objectieflens is om licht van het specimen te verzamelen en het te bundelen in een vergroot beeld (zie Figuur S.2). Dit beeld wordt dan verder vergroot door het oculair om het aan te passen aan het oplossend vermogen van het menselijk oog, of aan de detector van het beeldvormingssysteem.

Wanneer men een hypothetische puntvormige lichtbron in het specimen beschouwt, zal deze nooit als een ideaal punt worden waargenomen, maar als een zekere ruimtelijke lichtverdeling, namelijk de *Airy schijf* door de diffractie van de lichtgolven.

Voor een incoherente puntvormige lichtbron is de straal van de Airyschijf r_{Airy} afhankelijk van de golflengte van het gebruikte licht, de relatieve brekingsindex voor de overgang tussen het specimen en de objectieflens en de ontvankelijkheidshoek (*acceptance angle*) van de objectieflens; dit is voor een puntvormige lichtbron in het brandpunt van de lens gelijk aan de hoek waaronder licht van de puntbron nog door de objectieflens wordt onderschept. Het Rayleigh criterium zegt dat twee puntvormige lichtbronnen kunnen worden onderscheiden

als de afstand tussen de maxima van de twee Airy schijven ten minste r_{Airy} is. Deze limiet voor klassieke beeldvormingstechnieken is geldig wanneer de dip tussen de twee Airyschijven ongeveer 25% lager is dan de maxima van het diffractiepatroon. Voor technieken waar er bij de beeldvorming elektrische scanning komt kijken (zoals confocale microscopie) kan deze resolutie verder verbeterd worden met een factor 1.4 [Pawley, 1995, Plášek and Reischig, 1998] doordat de dip tussen de twee maxima dieper is².

S.2.3 Van Abbe naar Fourieroptica

Abbe bestudeerde de theorie van de vorming van een beeld uitgaande van een regelmatig rooster van aperturen als specimen. Dit rooster wordt dan belicht door een smalle lichtkegel, waardoor de invallende lichtstralen voor de eenvoud als parallel kunnen worden beschouwd. Het bekomen diffractiepatroon vertoont duidelijke maxima, die het best kunnen worden waargenomen in het focaalvlak. Het 0^{de} orde maximum heeft maximale intensiteit, en wordt gevormd in de richting van de parallelle straal, terwijl hogere orde maxima gevormd worden steeds verder weg van de optische as.

Abbe concludeerde dat om de periodiciteit van het rooster te kunnen onderscheiden, zowel 0^{de} als 1^{ste} orde diffractiemaxima moeten bijdragen tot de beeldvorming. De hoogste orde maxima in het bijzonder gaan verloren bij het beeldvormingsproces, omdat deze gevormd worden het verstuif van de optische as weg. Met andere woorden, bepaalde details (dus resolutie) gaan verloren tijdens het beeldvormingsproces.

Het aantal diffractie-orden dat kan bijdragen tot de vorming van het beeld, hangt af van de aperture van de objectieflens. Abbe leidde vanuit zijn theorie een formule af om het oplossend vermogen van een microscoop te berekenen; zijn uitdrukking is vrijwel identiek aan het Rayleigh criterium.

Abbe's theorie was een grote stap naar onze huidige inzichten over beeldvorming, en leidde uiteindelijk tot de Fourier optica. Deze laatste theorie beschrijft dat een convergerende lens een 2D Fouriertransformatie uitvoert [Clauws, 1998]. Als we een beeld van een object maken door middel van een convergerende lens, vinden we een amplitude- en fasedistributie die gegeven wordt door de Fouriertransformatie van het object (die de informatie bevat over de intensiteitsverdeling van het object), vermenigvuldigd met een fasefactor (die informatie bevat over de ruimtelijke structuur van het object). Als we enkel geïnteresseerd zijn in de intensiteit van een object, kan men deze fasefactor verwaarlozen.

²Het belangrijkste voordeel van de confocale microscoop is echter niet deze verbeterde laterale resolutie, maar dat het toelaat optische coupes te maken van een specimen.

S.2.4 Beeldvorming bij fluorescentiemicroscopie

De beeldvorming van een algemene fluorescentiemicroscoop werd beschreven in [van Kempen, 1999]. Wij vatten kort de voornaamste punten samen.

S.2.4.1 Algemene fluorescentiemicroscoop

Voor een objectieflens kan men, uitgaande van de scalaire diffractietheorie, een uitdrukking bekomen voor de driedimensionale PSF van een enkele lens. Er wordt uitgegaan van een uniforme lichtbron die licht uitzendt met golflengte λ_{ex} , waarvan licht passeert door een belichtingsapertuur $A_{ill}(x, y)$. Het licht passeert vervolgens door de objectieflens en wordt gefocusseerd in het specimen, waar het de fluoroforen exciteert. Deze zenden op hun beurt licht uit met golflengte λ_{em} , wat door de objectieflens wordt afgebeeld op de detector met aperture $A_{det}(x)$.

De PSF voor een algemene fluorescentiemicroscoop wordt bekomen als:

$$h_{GFM}(x, y) = A_{det}(x)h_{em}(x) \int d\xi h_{ex}(-\xi)A_{ill}(\xi)$$

waarbij h_{ex} en h_{em} de PSF van de objectieflens zijn bij λ_{ex} resp. λ_{em} .

S.2.5 Confocale microscoop

Voor een ideale confocale microscoop zijn zowel de belichtingsapertuur als de detectorapertuur ideale punten (dus deltafuncties). Daardoor kan de PSF van de ideale confocale microscoop beschreven worden als $h_{confocal,ideal}(x, y, z) = h_{em}(x, y, z)h_{ex}(-x, -y, -z)$.

In de praktijk is een confocaal diafragma echter geen ideaal punt. In dat geval wordt de uitdrukking voor de confocale PSF: $h_{confocal,real}(x, y, z) = (h_{em} * A_{det}(x, y, z))h_{ex}(-x, -y, -z)$. Voor meer details verwijzen we naar Sectie 2.3 van de Engelstalige tekst.

S.2.6 Klassieke microscoop

In een klassieke microscoop kan de belichtingsapertuur als “oneindig” worden beschouwd, waardoor $((A_{ill} * h_{ex})(x, y, z))$ een constante wordt. Daardoor wordt de PSF van een klassieke microscoop $h_{widefield}(x, y, z) = h_{em}(x, y, z)$

Doordat λ_{ex} en λ_{em} van dezelfde grootte-orde zijn, kan men bij benadering stellen dat de confocale PSF het kwadraat is van de klassieke PSF. Wanneer we de PSF’s voor de klassieke microscoop en de confocale microscoop vergelijken, zien dat we dat de PSF’s in het focaal vlak vergelijkbaar qua afmeting. De confocale PSF is echter veel minder uitgestrekt in de richting van de optische as, dus loodrecht op het focaal vlak (zie Figuur 2.22 van de Engelstalige tekst).

De Optische TransferFunctie (OTF, de Fouriergetransformeerde van de PSF) heeft in de richting loodrecht op het focaal vlak een gebied waar de OTF nul is. Deze zone is bekend als de “kegel van ontbrekende frequenties.” Objecten met ruimtelijke frequenties in deze kegel worden slecht weergegeven door een klassieke microscoop. Bij de OTF van de confocale microscoop is er daarentegen geen kegel van ontbrekende frequenties. Dat wil zeggen dat deze frequenties wel bijdragen tot het beeld, waardoor de resolutie langs de optische as beter is dan bij een klassieke microscoop.

S.3 Klassieke restauratietechnieken

S.3.1 Invers filter en Wienerfilter

S.3.1.1 Inverse filter

De eenvoudigste aanpak voor beeldrestauratie steunt op de eigenschap dat convolutie in het ruimtelijke domein wordt getransformeerd in vermenigvuldiging in het Fourierdomein. De vergelijking van de beeldvorming is³:

$$g(x, y) = (f * h)(x, y) + n(x, y),$$

waarbij $g(x, y)$ het gedegradeerde beeld is, $f(x, y)$ het ideale maar onbekende beeld, $h(x, y)$ de PSF en $n(x, y)$ de ruis in het beeld. Transformatie naar het Fourierdomein geeft:

$$G(u, v) = F(u, v)H(u, v) + N(u, v), \quad (\text{S.1})$$

waarbij de functies in hoofdletters de Fouriergetransformeerden zijn van de respectievelijke functies in kleine letters, en u en v zijn de coördinaten in het Fourierdomein.

$$\hat{F}(u, v) = \frac{G(u, v)}{H(u, v)} = F(u, v) + \frac{N(u, v)}{H(u, v)}$$

Deze uitdrukking illustreert een probleem van het invers filter; waar de Optische TransferFunctie (OTF) $H(u, v)$ zeer klein is (of zelfs nul), worden de spectrale componenten van het gedegradeerde signaal enorm versterkt omdat ze door zeer kleine waarden van de OTF worden gedeeld ($0 \leq H(u, v) \ll 1$). Een ruizig, vervaagd beeld verscherpen met een invers filter geeft een zeer onstabiel

³Voor de eenvoud van notaties beschouwen we hier 2D beelden; uitbreiding naar 3D is triviaal.

resultaat. Men zegt dat beeldrestauratie een *slecht geconditioneerd probleem* is, en dus zeer sterk afhankelijk is van de beginwaarden. Daarom is het noodzakelijk om methoden voor beeldrestauratie te stabiliseren door *a priori* kennis op te leggen. Daarvoor worden een aantal criteria geformuleerd waaraan een waarschijnlijke restauratie-oplossing moet voldoen.

Zo zijn beelden waarin de intensiteit voortdurend fluctueert tussen donker en licht (zoals "sneeuw" op televisie) minder waarschijnlijk dan beelden met grotere aaneengesloten gebieden van gelijkaardige intensiteit. Voorgaand criterium is een voorbeeld van een gladheidsbeperking, dat dient als een vorm van *a priori* kennis.

De eenvoudigste manier om het invers filter te stabiliseren, is om enkel die spectrale componenten van het signaal te versterken waarvoor de energie genoeg boven het ruisniveau uitsteekt, en andere spectrale componenten op nul te zetten.

S.3.1.2 Wienerfilter

De oorsprong van het Wienerfilter gaat terug naar de Wiener-Hopf vergelijkingen die werden afgeleid door Norbert Wiener en Eberhard Hopf in 1930. Het Wienerfilter is een oplossing van deze vergelijkingen.

In essentie gedraagt het Wienerfilter zich als het gestabiliseerde invers filter, waarbij echter een optimale afweging is gemaakt tussen enerzijds versterking van de spectrale componenten waar de signaalenergie domineert ten opzichte van de ruisenergie, en anderzijds attenuatie waar de ruisenergie van dezelfde grootte-orde wordt als de signaalenergie (of waar de ruisenergie zelfs volledig domineert). In feite kan het geclipte invers filter worden beschouwd als een globale *hard-thresholding* in het Fourier domein, terwijl het Wienerfilter een optimale kleinste-kwadraten thresholding van de spectrale componenten in het Fourier domein uitvoert.

S.3.2 Bayesiaanse restauratie

De meeste beeldrestauratie-algoritmen trachten de meest waarschijnlijke oplossing f te vinden gegeven de gedegradeerde waarneming g . Dit zogenaamde inverse probleem kan worden geformuleerd in termen van het voorwaartse probleem door middel van de regel van Bayes:

$$P(f|g) = \frac{P(g|f)P(f)}{P(g)} \quad (\text{S.2})$$

Het modelleren van beelddegradatie is een voorbeeld van een voorwaarts probleem: wanneer we een goed model voor de optiek en de detector hebben, kunnen we vrij gemakkelijk berekenen hoe het beeld van een bepaald object eruit zal zien met dit systeem. Het reconstrueren van het ideale object uit

de geobserveerde, gedegradeerde beeld is een zogenaamd invers probleem, wat heel wat moeilijker op te lossen is.

Om de meest waarschijnlijke oplossing f te vinden, moeten we die f zoeken die vergelijking (S.2) maximaliseert. In de praktijk is het vaak gemakkelijker om de negatieve logaritme van die uitdrukking te minimaliseren:

$$-\log P(f|g) = -\log P(g|f) - \log P(f) + \log P(g) \quad (\text{S.3})$$

waarbij de term $\log P(g)$ niet relevant is voor de minimalisatie. Meestal wordt de term $-\log P(g|f)$ de log-aannemelijkhedsfunctionaal genoemd. Deze term zorgt ervoor dat de oplossing aansluit bij de data; $-\log P(f)$ is de regularisatiefunctie, die de *a priori* informatie (zoals randvoorwaarden) bevat die we aan de oplossing opleggen. De meest gebruikte algoritmen zijn afgeleid uit deze uitdrukking uitgaande van Gaussiaanse of Poisson ruis, gecombineerd met een bepaalde regularisatiefunctie.

S.3.2.1 Gaussiaans ruismodel

(a) Wienerfilter versus Tikhonov-Miller restauratie

Wanneer het Gaussiaans ruismodel wordt gebruikt, dan is het minimaliseren van de negatieve log-aannemelijkhedsfunctionaal equivalent met het zoeken van de oplossing die het kwadratisch verschil tussen $g(x, y)$ en $(h * f)(x, y)$ minimaliseert. Als regularisatiefunctie wordt een uitdrukking gekozen die de ruwheid van de oplossing beperkt, om zo ruisversterking te reduceren.

λ is de regularisatieparameter, die de relatieve invloed afweegt van de log-aannemelijkhedsfunctionaal ten opzichte van de regularisatiefunctie. Hoe groter λ is, hoe meer belang er wordt gehecht aan de regularisatie; hoe kleiner λ , hoe meer belang er wordt gehecht aan een nauwe aansluiting aan de waargenomen data.

In feite is éénstaps Tikhonov-Miller restauratie equivalent met de oplossing van het Wienerfilter wanneer: $\lambda |C(u, v)|^2 = S_{nn}(u, v)/S_{ff}(u, v)$.

(b) Intermezzo: gradiëntgebaseerde minimalisatie

Een eenvoudige gradiëntgebaseerde minimalisatieprocedure telt een fractie van de gradiënt van de kostfunctie op bij de vorige schatting van de oplossing in elke iteratie. Deze methode heet de steilste-hellingsmethode, en heeft als nadeel dat ze traag convergeert, omdat de richting van de lokale gradiënt niet altijd naar het globale minimum leidt. Daardoor naderen de iteratiestappen het globale minimum langs een inefficiënt “zig-zag” traject.

Om dit probleem op te lossen werd de techniek van de Toegevoegde Gradiënten geïntroduceerd (een goede inleiding is te vinden in [Shewchuk, 1994]). Daarbij wordt een fractie van de vorige gradiëntrichting gecombineerd met de nieuwe

gradiëntrichting, om zo een betere nieuwe gradiëntrichting te kunnen bepalen dan bij de steilste-hellingsmethode.

(c) Iteratieve Tikhonov-Miller restauratie met randvoorwaarden (ICTM)

Het Wienerfilter en het Tikhonov-Miller restauratie-algoritme zijn voorbeelden van lineaire algoritmen, die twee belangrijke nadelen hebben:

1. Enkel die spectrale componenten van het signaal die niet volledig tot nul herleid zijn door het degradatieproces kunnen worden gerestaureerd.
2. Lineaire modellen kunnen de set van oplossingen niet beperken tot een zekere subruimte. Zo kan een beeld bijvoorbeeld enkel positieve intensiteiten hebben, terwijl de eerder besproken algoritmen in principe ook negatieve intensiteiten toelaten.

In deze paragrafen bespreken we een variant van het Tikhonov-Miller restauratie-algoritme dat niet-negativiteit oplegt aan de oplossing: het zogenaamde *Iteratieve Begrenste Tikhonov-Miller* algoritme. De praktische implementatie van dit algoritme berekent het minimum van de Tikhonov-Miller functionaal met behulp van de Toegevoegde Gradiëntenmethode. De niet-lineariteit wordt opgelegd als een harde begrenzing: na elke iteratie worden alle negatieve waarden van de huidige geschatte oplossing op nul gezet. Deze operatie is zeer sterk niet-lineair, en legt de extra *a priori* kennis op dat een beeld enkel positieve intensiteiten kan bevatten.

Wat we nog niet hebben vermeld tot nu toe, is de keuze van de regularisatieparameter λ . Voor de volledigheid vermelden we enkele van de meest gebruikte methoden om deze parameter te bepalen [Galatsanos and Katsaggelos, 1992, van Kempen, 1999]:

1. Methoden die uitgaan dat de ruisvariantie gekend is:

- **SNR methode:** deze methode stelt λ gelijk aan het inverse van de signaal/ruis verhouding.
- **Begrensde kleinste-kwadraten methode:** deze methode bepaalt de waarde voor λ die ervoor zorgt dat de kwadratische fout tussen het waargenomen gedegradeerde beeld en het hervervaagde gerestaureerde beeld minimaal wordt.

2. Methoden die geen kennis van de ruisvariantie nodig hebben:

- **GCV methode** (*Generalized Cross Validation*): daarbij wordt λ berekend met het *laat eenntje weg* principe: voor elke pixel wordt het restauratieresultaat berekend uitgaande van alle pixels, uitzonderd de beschouwde pixel. λ wordt dan zo gekozen dat het GCV voorspellingsverschil geminimaliseerd wordt.

- **Maximum Likelihood methode:** daarbij wordt λ bepaald onder de veronderstelling dat zowel de ruis als de regularisatiefunctionaal Gaussiaans verdeeld zijn met dezelfde variantie. De daaruit afgeleide aannemelijkhedsfunctionaal wordt dan geminimaliseerd.

S.3.2.2 Poisson ruismodel

(a) Richardson-Lucy restauratie

Het algoritme werd onafhankelijk afgeleid door [Lucy, 1974] en [Richardson, 1972]. Dit algoritme wordt zeer courant gebruikt voor de restauratie van astronomische en confocale beelden [van der Voort and Strasters, 1995, Van Kempen et al., 1997, Verveer et al., 1999]. Het Richardson-Lucy (RL) algoritme heeft volgend iteratief schema⁴:

$$f^{(t+1)}(x_1) = f^{(t)}(x_1) \left(\sum_{x_3} h(x_3 - x_1) \frac{g(x_3)}{\sum_{x_2} h(x_3 - x_2) f^{(t)}(x_2)} \right) \quad (\text{S.4})$$

waarbij $f^{(t+1)}(x)$ de nieuwe schatting is voor iteratie $t+1$, $f^{(t)}(x)$ is de schatting uit de vorige iteratie t . Opnieuw is $h(x)$ de PSF en $g(x)$ het gedegradeerde beeld. Shepp en Vardi [Shepp and Vardi, 1982] bewezen dat de aannemelijheid altijd groter wordt bij toenemend aantal iteraties, en maximaal wordt op voorwaarde dat alle pixels van het oorspronkelijke gedegradeerde beeld positief waren.

Het RL-algoritme heeft volgende eigenschappen [Bertero and Boccacci, 1998]:

- elke schatting $f^{(t)}(x)$ is niet-negatief;
- de totale beeldintensiteit van $(h * f^{(t+1)})(x)$ correspondeert met het totale aantal counts in het geobserveerde beeld $g(x)$;
- de log-aannemelijkhedsfunctie $L(f)$ is niet-dalend bij toenemende iteraties;
- $f^{(t)}(x)$ convergeert naar het maximum van $L(f)$.

Merk op dat het standaard RL-algoritme *geen* expliciete regularisatie oplegt (tenzij doordat de totale beeldintensiteit in elke interatie constant moet blijven, en dat waarden die op nul gezet in een iteratie nadien steeds nul moeten blijven). In S.3.2.2(b) bespreken we enkele regularisatiemethoden voor het standaard RL-algoritme.

(b) Klassieke regularisatie van Richardson-Lucy

Het referentiewerk bij uitstek is [Verveer, 1998]. Daar wordt de minimalisatie

⁴Voor de eenvoud van notatie in 1D; uitbreiding naar meer dimensies is triviaal.

behandeld van aannemelijkheidsfunctionalen zowel in het geval van data met Poisson ruis als data met Gaussiaanse ruis; deze functionalen worden telkens gecombineerd met verschillende regularisatiefunctionalen. Hier beperken we ons tot enkele veelgebruikte schema's voor de regularisatie van Richardson-Lucy:

- **RL met een beetje navervaging.** De eenvoudigste methode om het RL algoritme te regulariseren, is na elke iteratie een klein beetje te hervervagen, bijvoorbeeld door convolutie met een kleine Gaussiaanse kernel. Dit legt een zekere gladheid op aan de oplossing door singulariteiten te vervagen na elke iteratie. In [Silverman et al., 1990] werd een gelijkaardige aanpak voorgesteld voor de regularisatie van de reconstructie van emissie-tomografiebeelden. In [Slijpen and Beekman, 1999] werd aange- toond dat in bepaalde gevallen het filteren tussen de iteraties gelijkaardige resultaten oplevert als een bepaald filter toe te passen nadat de iteraties beëindigd zijn.
- **RL met Tikhonov-Miller regularisatie.** Door de log-aannemelijkheidsfunctionaal te combineren met de Tikhonov-Miller regularisatiefunctionaal, wordt er aan de oplossing een zekere gladheid opgelegd doordat de energie in de hoge frequenties van de oplossing aan banden wordt gelegd [Dey et al., 2004].
- **RL-Conchello.** Conchello stelde een regularisatiefunctionaal voor die pieken met te hoge intensiteit penaliseert, maar die niet echt gladheid aan het beeld oplegt [Conchello and McNally, 1996, van Kempen, 1999].
- **RL met Totale Variatie regularisatie.** De regularisatiefunctionaal voor Totale Variatie is zeer efficiënt voor objecten die bestaan uit gladde gebieden zonder textuur, omdat hij homogene gebieden gladder maakt, maar textuur en fijne details wegveegt. Anderzijds worden scherpe randen bewaard doordat enkel vervaagd wordt in de richting loodrecht op de gradiënt (dus lokaal evenwijdig met randen) en niet volgens de gradiënt [Dey et al., 2004]. Deze aanpak is verwant met anisotrope diffusie.

In de volgende sectie bespreken we de principes van multiresolutie, en een aantal toepassingen ervan. De sectie daarop wordt onze eigen methode besproken om het RL-algoritme te regulariseren, gebaseerd op de principes van multiresolutie.

S.4 Multiresolutie en beeldverwerking

S.4.1 Algemene inleiding

Tijdens de laatste 10 à 15 jaren zijn technieken gebaseerd op multiresolutie (zoals wavelets, stuurbare piramides, ...) steeds populairder geworden bij beeldverwerking. Deze sectie heeft niet als doel om al de wiskundige achtergrond

van multiresolutie te geven, enkel de basisideeën en de toepassingen ervan in beeldverwerking worden besproken.

In beeldverwerking steunen vele belangrijke technieken tegenwoordig op multiresolutieconcepten (waaronder scale-space, wavelets, stuurbare piramides, ...). Het basisidee is dat elk beeldkenmerk in een zeker bereik van schalen in het beeld terug te vinden is. Een voorbeeld uit het dagelijkse leven is de hiërarchie in een bloemenveld. Daarbij vinden we het bloemenveld terug op de ruwste schaal, terwijl het begrip “bloem” betekenis krijgt bij fijner schalen. Op nog fijnnere schalen wordt de bloem beschreven in termen van de cellen waaruit ze bestaat. Op ruwere schalen kunnen we spreken van het bos waarin het bloemenveld zich bevindt. Omdat een typisch beeld kenmerken op verschillende schalen bevat, is het nuttig om beelden te verwerken gebruik makend van verschillende schalen in plaats van slechts één schaal. Dat is precies wat multiresolutietechnieken doen [Lindeberg, 1994].

In de praktijk wordt een multiresolutievoorstelling van een beeld bekomen door steeds meer vervaging op een beeld toe te passen, waardoor steeds meer details uit het beeld verwijderd worden. Elke keer wordt het verschil genomen tussen het meest vervaagde en het minder vervaagde beeld. Deze verschilbeelden worden de detailbeelden van de overeenkomstige resolutieschaal genoemd.

S.4.2 Multiresolutie, wavelets en stuurbare piramides

S.4.2.1 Van Fouriertransformatie naar wavelets

De Fouriertransformatie is een belangrijk hulpmiddel bij signaalverwerking, omdat veel operaties (zoals berekening van het spectrum van een beeld, convolutie van twee beelden, ...) gebaseerd zijn op of gemakkelijker kunnen worden berekend met deze transformatie. De Fouriertransformatie geeft ons inzicht in de frequentie-inhoud van een beeld; met andere woorden, een beeld kan worden beschreven als de som van een reeks *sinusoïdale basisfuncties*, en de Fouriertransformatie laat ons zien hoe sterk de verschillende basisfuncties doorwegen. Gladde gebieden worden best beschreven in termen van lage frequenties, terwijl beeldkenmerken zoals punten, randen, lijnen, hoeken, ... best kunnen worden beschreven in termen van hoge frequenties.

Uit de Fouriergetransformeerde is het echter niet mogelijk om uit te maken van welke beeldkenmerken de verschillende frequenties afkomstig zijn, omdat de Fourierbasisfuncties zich oneindig uitstrekken. Vaak is het toch nuttig om te weten welke beeldkenmerken aanleiding gaven tot bepaalde spectrale componenten. Ook heeft het manipuleren van Fouriercoëfficiënten steeds invloed op het hele beeld.

Daarom werd de Korte-Termijn-Fourier-Transformatie geïntroduceerd. Hierbij wordt het signaal eerst vermenigvuldigd met een vensterfunctie met een zekere breedte vooraleer de Fouriertransformatie wordt toegepast. Zo wordt een lokale frequentie-analyse gemaakt. De basisfuncties voor deze representatie zijn

sinusoïdale golven, vermenigvuldigd met een vensterfunctie; een veelgebruikte keuze is een Gaussiaans venster. Deze procedure wordt toegepast op verschillende verschoven versies van het venster, waardoor de *lokale* frequentie-inhoud van het beeld voor elke beeldpositie wordt berekend.

Uit vorige paragraaf blijkt dat er een intrinsieke beperking bestaat op de gezamenlijke ruimtelijke resolutie (Δx) en de frequentieresolutie (Δf); deze beperking is een gevolg van dezelfde wiskunde die aan de basis ligt van het *Onzekerheidsprincipe van Heisenberg*. Dat vond zijn oorsprong in de quantummechanica, waar het stelde dat er voor een bewegend deeltje een intrinsieke beperking bestaat *op de gezamenlijke* nauwkeurigheid waarmee men de positie en de impuls van het deeltje op een zeker moment kan kennen: $\Delta x \Delta p \geq C$, waarbij C een constante is. Het Onzekerheidsprincipe zegt ook dat een functie en haar Fouriergetransformeerde niet beide een eindige drager⁵ kunnen hebben [Weisstein, WWWd]. In deze context betekent dit dat het onmogelijk is de plaats-frequentie voorstelling van een signaal met oneindige resolutie te kennen, oftewel, om de precieze spectrale componenten te kennen die op een exact ogenblik in een signaal aanwezig zijn: $\Delta x \Delta f \geq C$.

Vele nuttige signalen (zoals electrocardiogrammen, lijnen van beelden, ...) hebben hoogfrequente componenten die zich uitstrekken over korte intervallen, en laagfrequente componenten die zich uitstrekken over lange intervallen. Om discontinuïteiten in het signaal te lokaliseren, hebben we basisfuncties met zeer korte dragers (dus met een goede ruimtelijke resolutie) nodig. Om een gedetailleerde frequentie-analyse te bekomen hebben we dan weer zeer lange basisfuncties nodig. Daarom wordt bij waveletanalyse het volgende compromis gemaakt: er worden korte hoogfrequente basisfuncties gebruikt, samen met lange laagfrequente basisfuncties.

In beeldverwerking betekent dat het volgende: gladde beeldkenmerken (zoals vlakke gebieden) strekken zich uit over grotere gebieden in het beeld en geven vooral respons in de laagfrequente subbanden. In dit geval is het belangrijker om de frequentie-inhoud met een hoge frequentieresolutie te kennen dan de exacte locatie (omdat het kenmerk toch ruimtelijk uitgestrekt is).

In het geval van scherpe beeldovergangen (zoals punten, lijnen, stapronden, ...) is kennis over de exacte locatie van de overgang belangrijker en is een hoge ruimtelijke resolutie nodig. Daarom worden in dat geval hoogfrequente basisfuncties met een korte drager gebruikt, omdat die een goedgelokaliseerde respons produceren. De prijs die we daarvoor echter moeten betalen is een hogere onzekerheid over de frequentie-inhoud op deze positie.

Daarom wordt het ruimte-frequentie diagram als volgt onderverdeeld: voor de laagfrequente banden (ook wel de ruwere schalen genoemd) wordt er gekozen voor rechthoeken met een goede frequentieresolutie, terwijl voor de hoogfrequente banden rechthoeken met een goede ruimtelijke resolutie worden gekozen. Als we in het bijzonder kiezen voor de *dyadische* ruimte-frequentie voorstelling

⁵De drager van een functie (Eng: *support*) is het gebied waarin een functie verschillend is van nul.

halveert de frequentieresolutie wanneer we naar de volgende hogerliggende laag (oftewel: fijnere schaal) van het ruimte-frequentie vlak gaan, terwijl de ruimtelijke resolutie verdubbelt voor elke laag (zie Figuur 4.7 van de Engelstalige tekst).

De basisfuncties van de waveletdecompositie zijn de zogenaamde “wavelets,” oftewel *kleine golfjes*. Dit zijn oscillerende functies die corresponderen met een zekere plaats-frequentie resolutie. Als we de waveletdecompositie van een signaal willen berekenen, kunnen we kiezen tussen de Continue Wavelet Transformatie (CWT) of de Discrete Wavelet Transformatie (DWT).

De CWT berekent de inproducten van het signaal $f(x)$ met verschoven en geschaalde versies van de zogenaamde *moeder-waveletfunctie* $\psi(x)$:

$$Wf(u, s) = \langle f\psi_{u,s} \rangle = \int_{-\infty}^{+\infty} f(x) \cdot \frac{1}{\sqrt{s}} \psi^* \left(\frac{x-u}{s} \right) dx$$

$Wf(u, s)$ wordt berekend voor een continu bereik van de variabele u (die de verschuiving bepaalt) en s (die de schaal bepaalt, gerelateerd met de frequentie). Deze voorstelling is echter zeer redundant omdat ze een ééndimensionaal signaal omzet in een tweedimensionale functie $Wf(u, s)$. In het geval van digitale beeldverwerking maakt men vooral gebruik van de Discrete Wavelet Transformatie (DWT), waarbij zowel de verschuiving u als de schaal s over discrete bereiken worden gevarieerd. Deze transformatie kan efficiënt worden geïmplementeerd door middel van een filterbankschema.

Nu is het nuttig even in te gaan op het begrip “onderbemonstering” (*subsampling*): direct na één stap van de decompositie van een signaal in verschillende subbanden met een klassieke waveletdecompositie, bevat elke subband evenveel coëfficiënten als er coëfficiënten in het oorspronkelijke signaal aanwezig waren. Daarom zal een signaal met N pixels dat een decompositie met S schalen ondergaat, een waveletvoorstelling hebben met $N S$ waveletcoëfficiënten. Een deel van de informatie in deze voorstelling is echter redundant, en kan dus worden weggelaten. Zo is het in de klassieke waveletdecompositions voldoende om voor een zekere schaal slechts de helft van het aantal coëfficiënten van de vorige schaal te behouden⁶. Daardoor bevat een klassieke orthogonale waveletvoorstelling evenveel coëfficiënten als er coëfficiënten waren in het oorspronkelijke signaal; dit wordt een kritisch bemonsterde transformatie genoemd.

Om de wavelettransformatie ook op beelden te kunnen toepassen, wordt overgegaan van ééndimensionale naar tweedimensionale functies. Voor een beeld wordt de 1D wavelettransformatie afzonderlijk op de rijen en de kolommen van het beeld toegepast.

In het geval van 2D beelden worden vooral twee varianten gebruikt, namelijk die met twee (horizontale en verticale) details per resolutieschaal en die met drie

⁶Voor meerdimensionale signalen als beelden is per dimensie de helft van de coëfficiënten redundant. Voor een 2D beeld volstaat het zo om 1/4 van de coëfficiënten van de vorige schaal te behouden.

detailbeelden (horizontale, verticale en diagonale details) per resolutieschaal. Bij beide varianten kan men onderbemonsteren, waarmee men dan kritisch bemonsterde transformatie bekomt. Indien men niet onderbemonstert, heeft men een redundante transformatie; deze wordt gewoonlijk berekend met het zogenaamde *à trous* (=“met gaten”) algoritme [Mallat, 1999]).

S.4.3 De stuurbare piramide

S.4.3.1 Basisideeën achter de stuurbare piramide

Klassieke waveletdecomposities hebben enkele belangrijke nadelen: ten eerste is hun oriëntatieresolutie voor beelddetails beperkt. Zo kunnen bijvoorbeeld randen onder $\pi/4$ en onder $3\pi/4$ niet worden onderscheiden op basis van de verdeling van de energie over de verschillende subbanden. Een ander probleem met deze klassieke waveletvoorstellingen is hun verschuivingsvariante aard. Dit probleem doet zich voor in schema’s waarin er onderbemonstering wordt toegepast. Verschuivingsvariantie betekent dat de waveletcoëfficiënten voor twee lichtjes verschoven randen, volledig verschillend kunnen zijn (zie Figuur 4.12 in de Engelstalige tekst).

Om dit probleem te reduceren werd een variant van de klassieke wavelettransformatie voorgesteld, namelijk de stuurbare piramide [Simoncelli et al., 1992]. We bespreken nu kort enkele basisideeën omtrent deze transformatie.

Dikwijls is het nodig om georiënteerde structuren in een beeld te analyseren, zoals randen onder een bepaalde hoek. Daarvoor kan men het beeld convolveren met een bereik van georiënteerde filterkernels die elk reageren op een specifieke oriëntatie in het beeld. Zoiets heeft echter een enorme computationele kost.

Om dat te vermijden kan men een vaste set georiënteerde basiskernels definiëren. Deze kernels moeten voldoen aan de voorwaarde dat de responsen van het beeld gefilterd met deze kernels een set georiënteerde subbanden moet opleveren. Deze georiënteerde subbanden moet men dan kunnen interpoleren voor oriëntaties die liggen tussen de basisoriëntaties van de filterset [Freeman and Adelson, 1991, Simoncelli et al., 1992]. Zulk een set van basiskernels wordt een stuurbare filterset genoemd.

Bij een groot aantal vereiste oriëntaties is het computacioneel heel wat minder zwaar om te interpoleren tussen de basissubbanden dan telkens opnieuw te moeten convolveren met een aangepaste, geroteerde kernel. In [Freeman and Adelson, 1991] worden de algemene voorwaarden gegeven waaraan een set filterkernels moet voldoen om een stuurbare basisset te vormen. In [Simoncelli et al., 1992] wordt de stuurbare piramide beschreven. Deze decompositie is een multiresolutie-uitbreiding van stuurbare filters. Deze decompositie is gelijkaardig aan de klassieke waveletvoorstellingen, maar heeft een veel betere oriëntatieresolutie en heeft verschuivingsinvariantie.

S.4.3.2 Implementatiedetails

Er zijn twee methoden om deze dekompositie te implementeren: de ruimtelijke methode en de methode in het Fourierdomein. In essentie berekenen ze dezelfde opsplitsing in subbanden. De praktische implementatie is echter verschillend.

De ruimtelijke methode extraheert de verschillende subbanden door het beeld te convolveren in het ruimtelijke domein met een set filterkernels. Deze filterkernels werden vooraf speciaal ontworpen om een bepaalde compacte drager in het Fourierdomein te hebben. De ruimtelijke implementatie werd beschreven in [Simoncelli et al., 1992]; een vereenvoudigd schema is terug te vinden in [Castleman et al., 1998].

De methode in het Fourierdomein transformeert het beeld naar het Fourierdomein, en vermenigvuldigt de getransformeerde met een set transferfuncties die de subbanden van de dekompositie bepalen. Elke subband wordt dan weer naar het ruimtelijke domein teruggetransformeerd. Een voorbeeld van een set transferfuncties voor deze methode vindt men terug in [Portilla et al., 2003].

S.4.4 Toepassingen

S.4.4.1 Kwadratuurfilters en fase-onafhankelijkheid

Het gebruik van kwadratuurfilters gaat terug tot Denis Gabor in [Gabor, 1946] en zijn definitie van het “elementair signaal.” Later werden kwadratuurfilters gebruikt om de lokale energie en lokale fase van beeldkenmerken te berekenen. Deze twee kenmerken werden vaak gebruikt in beeldanalyse [Morrone and Owens, 1987, Venkatesh and Owens, 1990, Pudney et al., 1994, Granlund and Knutsson, 1995, Kovesi, 2002].

In een beeld zijn er heel wat structuren met verschillende fases aanwezig [Granlund and Knutsson, 1995]. De fase van een beeldkenmerk is een maat van hoe even of oneven een beeldkenmerk is. Even beeldkenmerken vertonen een spiegelsymmetrie ten opzichte van een centraal punt (bij 1D signalen), een rechte (bij 2D beelden) of een vlak (bij volumetrische beelden). Een witte lijn op een zwarte achtergrond in een beeld is een voorbeeld van een even structuur. Oneven beeldkenmerken vertonen een oneven symmetrie ten opzichte van een centraal punt (bij 1D signalen), een rechte (bij 2D beelden) of een vlak (bij volumetrische beelden). Een gebied begrensd door een staprond is een voorbeeld van een oneven beeldkenmerk (op een constante na) in een 2D beeld. Andere beeldcomponenten kunnen ontbonden worden in een even en een oneven component, en hebben dus een fase tussen deze van even en oneven kenmerken.

We bespreken nu kort de respons van verschillende filters op even en oneven kenmerken in een beeld. Een oneven filter is een filter met oneven symmetrie. Dit soort filter geeft voor even kenmerken (zoals een deltapiek) een nuldoorgang op de plaats van het kenmerk; voor oneven kenmerken (zoals een staprond) geeft

het een lokaal maximum. Een even filter daarentegen geeft voor even kenmerken (zoals een deltapiek) een lokaal maximum op de plaats van het kenmerk en voor oneven kenmerken (zoals een staprond) dan weer een nuldoorgang. Voor kenmerken die noch even, noch oneven zijn, is er in beide gevallen noch een maximum, noch een nuldoorgang terug te vinden op de plaats van de kenmerken. Elk beeldkenmerk kan echter ontbonden worden als de som van een even en een oneven component.

Wanneer we het kwadraat van de respons van het even filter nemen, en optellen bij het kwadraat van de respons van het oneven filter, dan vinden we lokale maxima bij alle interessante kenmerken. Deze som van de responsen van een even en een oneven filter voor een bepaalde plaats in een beeld wordt de *lokale energie* genoemd wanneer het even en het oneven filter in kwadratuur zijn.

Twee filters zijn in kwadratuur als ze elkaars Hilbertgetransformeerde zijn. De Hilberttransformatie is de transformatie die een functie omzet in een functie met hetzelfde Fourierspectrum, maar met alle fases over $\pi/2$ verschoven. Het resultaat van deze transformatie is dat een even functie in een oneven functie wordt omgezet, en omgekeerd. Wanneer het beeld met een set van filters in kwadratuur wordt gefilterd, is de respons niet langer afhankelijk van de lokale fase van de beeldkenmerken.

S.4.4.2 Oriëntatie-analyse en randdetectie

Met een set filters in kwadratuur is het mogelijk om de dominante oriëntatie en de lokale energie op een bepaalde plaats in een beeld te bepalen door de fase van de eerste harmonische te bepalen [Freeman and Adelson, 1991]. Deze maat is nauw verwant met de energieverdeling in het lokale Fourierspectrum. Als we de Fouriergetransformeerde nemen van een omgeving in het beeld waarin er slechts één dominante oriëntatie aanwezig is, dan zal de meeste energie van de Fourierspectrum van deze omgeving geconcentreerd zijn in een smalle sector die georiënteerd is volgens dezelfde hoek als de lokale beeldgradiënten. Hoe smaller deze sector is, hoe uitgesprokener een bepaalde oriëntatie in die buurt aanwezig is (zie Figuur 4.29(b) van de Engelstalige tekst).

S.4.4.3 Hoekdetectie

We schetsen hier kort een origineel idee voor een methode voor hoekdetectie. Voor het begrip bestaat er geen éénduidige wiskundige definitie. Men kan een hoek definiëren als een klein gebiedje in een beeld waarin een groep aangrenzende pixels (die bijvoorbeeld een kwart van de pixels in dit gebied omvat) een intensiteit heeft die significant hoger of lager is dan alle andere pixels in dit gebiedje. Een andere definitie noemt een hoek een punt met hoge kromming in een beeld. Wij gaan uit van deze laatste definitie.

Zoals reeds werd opgemerkt in [Simoncelli, 1999] zijn waveletcoëfficiënten verre van onafhankelijk van elkaar. Na de waveletdecompositie vertonen de wave-

letcoëfficiënten die corresponderen met een bepaalde positie in het beeld nog steeds correlatie doorheen de resolutieschalen en doorheen de verschillende oriëntatiesubbanden. Ook zijn waveletcoëfficiënten meestal ruimtelijk geclusterd. De correlatie doorheen resolutieschalen, en de correlatie tussen verschillende oriëntatiesubbanden gaat vooral op voor hoeken en vlekken (punten), die aanleiding geven tot een respons in verschillende georiënteerde subbanden.

Daarom stellen we voor om het product te nemen van de coëfficiënten van de verschillende georiënteerde subbanden per resolutieschaal, omdat hoeken wegens hun hoge kromming een respons hebben in de verschillende georiënteerde subbanden (in tegenstelling tot gewone randen). We stellen daarom voor om de georiënteerde subbanden per resolutieschaal coëfficiënt per coëfficiënt met elkaar te vermenigvuldigen. Het resultaat daarvan is dat er enkel respons is waar er structuren aanwezig zijn die respons hebben in de verschillende oriëntaties. Dit product wordt apart berekend voor alle resolutieschalen. Vervolgens wordt het resultaat daarvan coëfficiënt per coëfficiënt vermenigvuldigd over alle resolutieschalen.

Verder onderzoek zal nodig zijn om verfijndere en robuustere detectieregels op te stellen en te evalueren

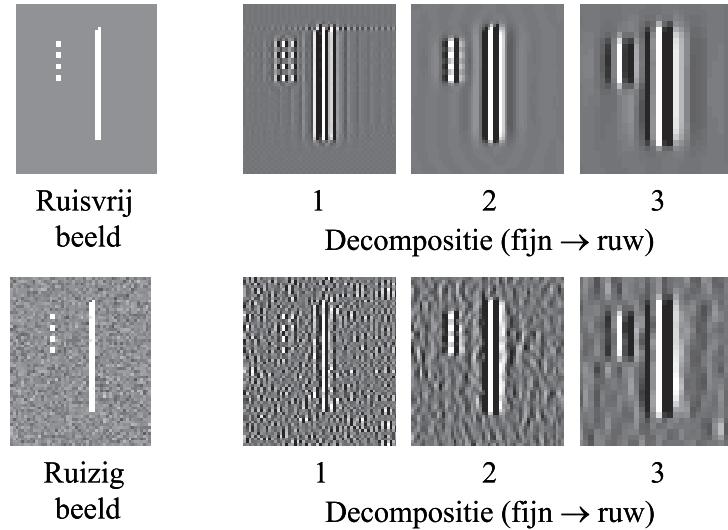
S.4.4.4 Beeldsegmentatie

Voor de volledigheid vermelden we ook dat multiresolutiegebaseerde technieken met succes worden toegepast in beeldsegmentatie, zoals onder andere beschreven in [Bouman and Shapiro, 1994]. Daarbij worden de eigenschappen van een multischalige decompositie direct in de segmentatiemethode gebruikt. In het bijzonder wordt in [Bouman and Shapiro, 1994] een MAP schatter gedefinieerd die de waarschijnlijkheid minimaliseert dat een pixel verkeerd wordt geïdentificeerd, uitgaande van een multischalig MRF model.

S.4.4.5 Ruisreductie

Gaussiaanse ruis

Multischalige decomposities (zoals de stuurbare piramide) hebben bewezen dat ze de scheiding van ruis en nuttige beeldkenmerken sterk kunnen vereenvoudigen, zoals we illustreren we in Figuur S.3. Daar is duidelijk te zien dat randen in het beeld aanleiding geven tot coëfficiënten met grote amplitude (hoge intensiteit betekent grote amplitude en positief; lage intensiteit betekent grote amplitude en negatief; grijs is nul) in alle schalen voor beide beelden. Die grote coëfficiënten zijn zowel ruimtelijk gecorreleerd in aaneengesloten clusters als doorheen de schalen. Ruis daarentegen resulteert in ongecorreleerde coëfficiënten met kleine amplitudes. Alle waveletgebaseerde algoritmen voor ruisreductie baseren zich in essentie op deze twee eigenschappen om ruis te onderscheiden van nuttige beeldkenmerken.



Figuur S.3: Vergelijking van de decompositie van een beeldstructuur zonder ruis, en met ruis. De eerste kolom is het beeld zelf, de andere kolommen tonen de opeenvolgende schalen van de decompositie.

Algoritmen voor ruisreductie in het waveletdomein hebben bewezen superieur te zijn aan klassieke ruisreductie-algoritmen. Deze algoritmen volgen in essentie dezelfde strategie:

- pas de voorwaartse wavelettransformatie toe;
- *threshold* of *shrink* sommige waveletcoëfficiënten volgens bepaalde regels of criteria;
- pas de inverse wavelettransformatie toe.

In [Donoho and Johnstone, 1994, Donoho and Johnstone, 1995] werd een statistische methode ontwikkeld (gebaseerd op SURE: Stein Unbiased Risk Estimator) om een optimale globale drempel (*threshold*) te berekenen voor alle waveletcoëfficiënten binnen een subband van een beeld dat gedegradeerd is door witte Gaussiaanse ruis. Alle coëfficiënten met een amplitude kleiner dan deze drempel worden op nul gezet, terwijl bij coëfficiënten met amplitude groter dan deze drempel de waarde van de drempel van hun amplitude wordt afgetrokken (*soft thresholding*).

Wij hebben echter gekozen om een recentere waveletgebaseerde ruisreductiemethode te gebruiken. Deze methode gaat ook uit van Gaussiaanse ruis, en werd ontwikkeld door [Sendur and Selesnick, 2002]. De MAP schatter die hij afleidde, gaat uit dat een waveletcoefficient en zijn “ouder” (= de waveletcoëfficiënt die overeenkomt met dezelfde positie in het beeld en met dezelfde oriëntatie, maar

in de volgende ruwere resolutieschaal) als gezamenlijke verdeling een bivariate Gaussiaanse distributie heeft. Dit algoritme is eenvoudig te implementeren, heeft een lage computationele kost en geeft niettemin *state-of-the-art* ruisreductieresultaten, omdat:

- het zich lokaal aanpast aan de aanwezigheid van randen;
- het (tenminste gedeeltelijk) rekening houdt met de correlatie tussen coëfficiënten doorheen de verschillende schalen.

Voor de volledigheid vermelden we nog enkele andere *state-of-the-art* methoden voor waveletgebaseerde ruisreductie. In [Pižurica et al., 2002] wordt een Markov Random Veld model gebruikt om filtermaskers te berekenen om nuttige beelddetails te behouden en ruis te onderdrukken in vlakke gebieden. In [Portilla et al., 2003] werd een Gaussian Scale Mixture model gebruikt om de relaties tussen de waveletcoëfficiënten te modelleren. Dit model laat toe om niet-Gaussiaanse coëfficiënten te behouden en Gaussiaans-achtige ruis te onderdrukken.

Poissonruis Ruisreductie die uitgaat van de SURE drempel in het waveletdomain, werd reeds toegepast in microscopie door [Boutet de Monvel et al., 2001, Stollberg et al., 2003]. Deze ruisreductie is suboptimaal omwille van twee redenen.

Ten eerste wordt er uitgegaan van Gaussiaanse ruis, terwijl in confocale fluorescentiebeeldvorming de voornaamste soort van ruis Poisson verdeeld is [Pawley, 1995, van Kempen et al., 1997]. In tegenstelling tot Gaussiaanse ruis is Poissonruis afhankelijk van de intensiteit van het onderliggende beeld, waardoor het maken van het onderscheid van de nuttige beeldkenmerken en ruis zeer moeilijk is [Nowak and Baraniuk, 1999].

Ten tweede wordt er enkel gebruik gemaakt van het feit dat nuttige coëfficiënten worden verondersteld grote amplitudes te hebben, en de correlaties van de coëfficiënten (zowel ruimtelijk als doorheen de resolutieschalen) niet worden gebruikt (waarbij we weer verwijzen naar Figuur S.3).

We pakken deze problemen als volgt aan: om met Poissonruis om te gaan, gebruiken we de Anscombe-transformatie [Starck et al., 2000], die de Poissondaten transformeert naar data met een Gaussiaanse distributie met standaardafwijking één. Deze transformatie laat toe om gekende technieken voor reductie van Gaussiaanse ruis toe te passen op de heel wat lastigere Poissonruis. We hebben de methode van [Şendur and Selesnick, 2002] aangepast voor implementatie met stuurbare piramiden, en geïntegreerd met de Anscombe-transformatie.

S.4.4.6 Beeldrestauratie

In deze subsectie geven we kort een overzicht van beeldrestauratiemethoden waar wavelets aan te pas komen. Die technieken kunnen in twee categorieën worden opgedeeld:

- regularisatie en verscherping worden beide in het waveletdomein uitgevoerd;
- enkel regularisatie wordt uitgevoerd in het waveletdomein en wordt gecombineerd met een verscherping in het ruimtelijke domein of het Fourierdomein.

(a) Hybride technieken: enkel regularisatie in het waveletdomein.

In feite is het het meest voor de hand liggend om de convolutie-operator voor te stellen in het Fourierdomein (aangezien daar de operator gediagonaleerd is) [Figueiredo and Nowak, 2003]. Het beeld zelf wordt daarentegen beter gemodelleerd in het waveletdomein. Algemeen is het echter zeer moeilijk om convolutie-operators in the waveletdomein voor te stellen (er bestaat een waveletgebaseerd algoritme [Beylkin et al., 1991, Beylkin, 1991], maar dat wordt niet vaak gebruikt in beeldverwerking omwille van de complexiteit). Dit suggerert om Fouriergebaseerde deconvolutie te combineren met waveletgebaseerde ruisonderdrukking als regularisatie. Het eerste artikel over dit onderwerp [Neelamani et al., 1999, Neelamani et al., 2004] combineerde een eenvoudig Wienerfilter voor verscherping met een waveletgebaseerd lokaal Wienerruisfilter als regularisatie.

Zowel [Kalifa et al., 2003] als [Jalobeanu et al., 2003] voeren een ongeregulariseerde inversie van de vervagingsoperator uit, gevolgd door een waveletgebaseerde onderdrukking van de ruis (die niet meer wit is door het feit dat de verscherping de ruis in de hoge frequenties meer versterkt dan de ruis in de lage frequenties). In [Kalifa et al., 2003] wordt er gebruik gemaakt van een zogenaamde “mirror wavelet basis” om het beeld zo voor te stellen dat de ruis kan worden gescheiden van de nuttige beeldkenmerken, zelfs bij de hoge frequenties, waar de ruis sterk versterkt werd. In [Jalobeanu et al., 2003] wordt een gelijkaardige aanpak gevuld, maar worden complexe wavelet packets gebruikt om ruis van nuttige kenmerken te kunnen onderscheiden.

In [Figueiredo and Nowak, 2003] wordt een FFT-gebaseerde Landweber-achtige deconvolutiestap gevolgd door een waveletgebaseerde MAP/MPLE soft-thresholding stap.

Onze eigen methode SPERRIL (beschreven in Hoofdstuk 6) valt ook onder deze categorie.

(b) Ruisonderdrukking en verscherping in het waveletdomein.

In [Banham et al., 1994, Banham and Katsaggelos, 1996] wordt een prefilter gecombineerd met een multischaal-Kalmanfilter. Het prefilter voert een kleinste-kwadratenfiltering met randvoorwaarden uit (in feite een waveletgebaseerd equivalent van de ICTM methode die we hebben beschreven in subsectie S.3.2.1), die een oplossing levert die scherp, maar ruzig is. Dit resultaat wordt dan gebruikt als invoer voor het multischaalig Kalman vergladdingsfilter. Het prefilter verwijdert ook afhankelijkheid tussen schalen binnen de toestandsvectoren. Daardoor wordt de lengte van deze toestandsvectoren sterk gereduceerd, wat ten goede komt aan de computationale efficiëntie van het algoritme.

In [Portilla and Simoncelli, 2003] wordt [Portilla et al., 2003] uitgebreid naar het volledige probleem van beeldrestauratie. Elke omgeving van coëfficiënten van de multischalige decompositie van een beeld wordt gemodelleerd als het product van een Gaussiaanse vector en een onafhankelijke verborgen toevalsvariabele. Het *Gaussian Scale Mixture* model werd uitgebreid om ook het vervagingsproces mee in rekening te brengen (hoe om te gaan met enkel ruis werd reeds beschreven in [Portilla et al., 2003]). Daaruit werd een lokale Wiernerschatting afgeleid.

In [Dragotti and Vetterli, 2003] wordt aangetoond dat elk stuksgewijs glad signaal kan worden uitgedrukt als een som van stuksgewijs polynomiale signalen en een uniform glad residu. Waveletvoetafdrukken (*footprints*) zijn schaal-ruimte vectoren die discontinuïteiten in stuksgewijs polynomiale signalen modelleren. Eerst wordt een woordenboek van waveletvoetafdrukken opgesteld uit de wavelettransformatie van het signaal. Gebaseerd op dit woordenboek is het dan mogelijk om operaties uit te voeren zoals ruisonderdrukking en deconvolutie. Ruisonderdrukking bijvoorbeeld wordt uitgevoerd door de thresholding uit te voeren op de wavelet voetafdrukken in plaats van op alle waveletcoëfficiënten. Deconvolutie wordt uitgevoerd door eerst de ruis te onderdrukken, waarna de wazige voetafdrukken worden vervangen door de overeenkomstige scherpe versies. Resultaten voor 1D signalen worden besproken, maar de uitbreiding naar meer dimensies is nog niet beschikbaar.

Onze eigen methode met Multischalige Markov Random Velden (beschreven in Hoofdstuk 7) valt in deze categorie.

S.4.5 3D multischalige transformaties

Om de ontwikkelde algoritmen te kunnen uitbreiden naar drie dimensies (zoals voor de restauratie van 3D confocale beeldenstacks), hebben we 3D uitbreidingen van de beschreven multiresolutie-decomposities nodig. Daarom geven we hier een kort overzicht van een aantal bestaande 3D multiresolutie-decomposities.

S.4.5.1 3D stuurbare piramide

Deze inleiding is losjes gebaseerd op [Delle Luche et al., 2003]. Het voornaamste probleem bij de uitbreiding van de stuurbare piramide van 2D naar 3D is niet het definiëren van de laagdoorlaatfilters en de hoogdoorlaatfilters, aangezien die filters omwentelingssymmetrie hebben (welke eenvoudig uit te breiden is naar sferische symmetrie in 3D). Het probleem is om de richtingen van de georiënteerde banddoorlaatfilters isotroop over de bol te verdelen, en daarbij bovendien perfecte reconstructie mogelijk maken (dus dat de som van de kwadranten van de transferfuncties over alle oriëntaties één is). In 2D kan men als richtingen van die georiënteerde transferfuncties de diagonalen of de middelloodlijnen van regelmatige veelhoeken kiezen.

In [Freeman and Adelson, 1991] werd een uitbreiding van het begrip stuurbare filters van 2D naar 3D voorgesteld op basis van een expansie in sferische harmonieken. In [Andersson, 1992] werd een methode voorgesteld die uitging van regelmatige veelvlakken (de logische uitbreiding naar 3D voor regelmatige veelhoeken) om de banddoorlaatfilters te oriënteren en zo een set filters te bekomen die uniform verdeeld is over de bol. In [Yu et al., 2000] werd een lokale methode voorgesteld die gebaseerd is op de stuurbaarheid bij benadering gebruik makend van kegelvormige Gaussiaanse filters.

In [Delle Luche et al., 2003] werd er een nieuwe methode voorgesteld die het idee van de kegelvormige filters combineert met het concept van regelmatige veelvlakken. Daarbij worden de diagonalen of de middelloodlijnen van regelmatige veelvlakken gekozen als centrale assen van de 3D georiënteerde banddoorlaatfilters. In [Delle Luche et al., 2003] werd er ook vermeld dat voor een kubus men drie richtingen kan definiëren (de middelloodlijnen) of vier (de diagonalen). Voor een icosaëder of een dodecaëder zijn er zo zes of tien richtingen. We hebben het algoritme geïmplementeerd voor het geval van drie centrale assen (waarbij we dus gewoon de drie coördinaatassen konden nemen).

S.4.5.2 Andere 3D multischalige transformaties

Natuurlijk bestaan er nog andere methoden voor een 3D decompositie/reconstructie. Een voorbeeld is de scheidbare transformatie van Mallat, die hoger werd beschreven. De uitbreiding naar drie dimensies is voor de hand liggend, maar dan treden natuurlijk wel weer de problemen op zoals verschuivingsvariantie en beperkte oriëntatieresolutie.

Er bestaat nog een andere methode die de eigenschappen van verschuivingsinvariantie en goede oriëntatiegevoeligheid heeft, namelijk de complexwaardige wavelettransformatie. Deze transformatie werd voorgesteld door [Kingsbury, 1998b, Kingsbury, 1998a] en uitgebreid door [Selesnick and Li, 1998, Selesnick, 2004]. Deze decompositie heeft enerzijds een lagere computationale kost dan de stuurbare piramide, maar is niet stuurbaar. De complexe waveletdecompositie is namelijk scheidbaar volgens de X en Y-assen, terwijl de stuurbare piramide scheidbaar is in poolcoördinaten.

S.5 Twee methoden voor parametrische PSF schatting

S.5.1 Inleiding

In voorgaande hoofdstukken hebben we het probleem besproken van beelddegradatie door vervaging en ruis. Om even kort te hernemen (we gaan hier uit van 2D beelden): beelddegradatie wordt gemodelleerd als $g(x, y) = N((h * f)(x, y))$ waarbij $g(x, y)$ het geobserveerde, gedegradeerde beeld is.

$f(x, y)$ is het onbekende, ideale beeld en $h(x, y)$ is de vervagingskernel (of PSF). $N(\cdot)$ vat de ruisprocessen samen die optreden bij de beeldvorming.

Om nu te kunnen compenseren voor de beelddegradatie, moeten we de ruisprocessen kunnen modelleren en moeten we $h(x, y)$ kennen. Deze vervagingskernel is echter vaak helemaal niet, of slechts gedeeltelijk gekend. In dat geval moet $h(x, y)$ worden geschat. Deze schatting kan gebeuren vóór het eigenlijke restauratieproces (wat *vervagingsschattung* wordt genoemd), of tijdens het restauratieproces (wat *blinde deconvolutie* wordt genoemd). Een goed overzicht van de tweede klasse van methoden is te vinden in [Kundur and Hatzinakos, 1996a, Kundur and Hatzinakos, 1996b]. Blinde deconvolutie-algoritmen zijn doorgaans complex en hebben een hoge computationale kost. Een goed overzicht van methoden voor *vervagingsschattung* is te vinden in de inleiding van [Flusser and Suk, 1998].

De methode uit [Savakis and Trussell, 1993a] is in feite niet echt een parametrische vervagingsschattingsmethode, maar een methode die beslist welke PSF uit een gegeven set van kandidaat PSF's de beste is voor restauratie (ook al zijn de kandidaat PSF's niet van een of andere parametrische vorm).

In de volgende twee secties bespreken we twee door ons ontwikkelde waveletgebaseerde methoden voor het schatten van parametrische vervaging.

S.5.2 Wave-Track: tracking van de waveletcoëfficiënten doorheen de schalen

S.5.2.1 Inleiding

In deze subsectie wordt een waveletgebaseerde methode voorgesteld om beeldvervaging te schatten, gebruik makend van informatie uit het beeld zelf. De methode is geïnspireerd op verwant onderzoek naar ruisreductie in beelden door dr. A. Pižurica van onze onderzoeksgroep. Onze methode werd gepubliceerd in [Rooms et al., 2001b, Rooms et al., 2002b, Rooms et al., 2002c], en werd geïncorporeerd in het automatische beeldrestauratie-algoritme SPERRIL, waarvan de details werden gepubliceerd in [Rooms et al., 2003c, Rooms et al., 2003b, Rooms et al., 2005] en in meer detail beschreven in Sectie S.6.

We kijken naar hoe scherp de scherpste randen in het vervaagde beeld zijn, die op hun beurt informatie over de PSF bevatten. In de praktijk wordt een gladheidsmaat berekend voor deze randen. Er wordt gebruik gemaakt van de relatie tussen deze parameter en de evolutie van de wavelet modulus maxima doorheen de verschillende resolutieschalen, die enkel van de beeldvervaging afhangt en niet van de beeldinhoud. In het bijzonder schat onze methode de variantieparameter σ_{blur} van een Gaussiaanse PSF.

Onze methode kan de Gaussiaanse PSF parameter σ_{blur} schatten met een nauwkeurigheid van enkele procenten. Deze methode kan ook worden uitgebreid naar

Gaussiaanse PSF's die niet axiaal symmetrisch zijn, en zelfs naar PSF's die zelfs niet Gaussiaans zijn. Voor out-of-focus vervaging wordt een uniforme schijf-vormige PSF gebruikt [Pavlović and Tekalp, 1992, Savakis and Trussell, 1993a]. Onze methode moet slechts minimaal aangepast worden voor dit soort PSF's, zoals verder in de tekst wordt aangetoond.

S.5.2.2 Principe

Onze methode voor parametrische vervagingsschatting is gebaseerd op het schatten van de scherpte van de scherpste randen in het beeld. Om die randen in het beeld te analyseren, wordt de evolutie van de maxima van de absolute waarden van de waveletcoëfficiënten doorheen de schalen gevolgd voor alle punten waar een intensiteitsgradiënt wordt gevonden in de horizontale of de verticale richting. De modulusmaxima zijn de lokale maxima in absolute waarde van de waveletcoëfficiënten

Deze evolutie doorheen de schalen hangt af van drie factoren, zoals Mallat heeft aangetoond [Mallat and Zhong, 1992, Mallat and Hwang, 1992, Mallat, 1999]:

- de Lipschitz-exponent van het onderliggende signaal;
- de basisfuncties gebruikt bij de waveletdecompositie;
- de vervaging van het signaal.

De eerste factor is de Lipschitz-exponent (ook de Hölder exponent genoemd), die een maat is voor de regulariteit in een bepaald punt. In feite is dit concept een uitbreiding van het aantal keer dat een functie afleidbaar is in een bepaald punt. Een functie die één keer afleidbaar is in een bepaald punt, heeft Lipschitz-exponent +1; een stapfunctie (niet afleidbaar) heeft Lipschitz 0 op de positie van de stap; een Diracdeltafunctie $\delta(x)$ heeft dan weer Lipschitz -1 op de positie van de deltapiek. Wanneer er echter Gaussiaanse vervaging wordt toegepast op een beeld, zijn alle beeldkenmerken in principe een oneindig aantal keer afleidbaar. Dat betekent dat de Lipschitz-exponent niets kan opleveren om de vervaging te schatten.

De derde factor omvat echter wel de beeldvervaging. Als we ons beperken tot beeldstructuren met dezelfde Lipschitz-exponent, dan worden voor een scherpe intensiteitsovergang de waveletcoëfficiënten kleiner in amplitude van fijne naar ruwere schalen en voor een vervaagde overgang worden de waveletcoëfficiënten groter in magnitude.

In feite zouden we enkel de scherpste randen met een vrij grote amplitude (dus met gelijke Lipschitz-exponenten) in het beeld in rekening mogen brengen om ruis en gladde beeldovergangen uit te sluiten bij onze vervagingsschatting. Daarom selecteren we in de fijnste resolutieschaal enkel de waveletcoëfficiënten met een magnitude groter dan een zekere drempel. De evolutie van deze waveletcoëfficiënten hangt dan enkel af van de vervaging van deze randkenmerken.

Voor scherpe stapronden hebben de corresponderende waveletcoëfficiënten een constante magnitude doorheen de resolutieschalen en voor een gladde overgang stijgen de magnituden van de waveletcoëfficiënten van fijnere naar ruwere schalen. Uit het tempo waarin de magnituden stijgen van fijnere naar ruwere schalen kunnen we de beeldvervaging bepalen.

Voor elke beschouwde pixel worden de maxima van de absolute waarden van de corresponderende waveletcoëfficiënten gevuld doorheen de schalen, en er wordt een exponentiële functie aangepast aan hun verloop in functie van de schaal:

$$|W(x, y; s)| = a \exp(bs),$$

waarbij $|W(x, y; s)|$ de maxima van de absolute waarden van de waveletcoëfficiënten in functie van de resolutieschaal s en de positie (x, y) zijn; a en b zijn de geschatte parameters, waarbij b gerelateerd is met de randvervaging. Daarom noemen we b de *vervagingsexponent*.

Met de vervagingsexponenten die corresponderen met alle pixels van significantie beeldranden wordt een histogram opgebouwd. Daarvoor wordt het bereik van de vervagingsexponenten in intervallen verdeeld met breedte 0.1. Omdat we de vervagingsexponenten enkel berekend hebben voor overgangen in het beeld met een grote amplitude, hebben we ons zo reeds tot de scherpste overgangen met een grote amplitude in het beeld beperkt.

Daarna wordt het zwaartepunt (*center of gravity*: CG) van het histogram berekend, wat een maat is voor de vervaging van het beeld. We berekenen de vervagingsexponent die overeenkomt met CG van het histogram, en bepalen de gemiddelde $\text{CG}_{\sigma_{blur}}$ over de hele set testbeelden die vervaagd werden met dezelfde σ_{blur} . Aan deze data $(\sigma_{blur}, \text{CG}_{\sigma_{blur}})$, werd een exponentiële curve aangepast: $\sigma_{blur} = a \exp(b \text{CG}_{\sigma_{blur}})$.

S.5.2.3 Robuustheid tegen ruis

De hierboven beschreven techniek is zeer ruisgevoelig, omdat deze uitgaat van het vinden van lokale maxima. In toepassingen zoals digitale camera's met voldoende belichting, tonen onze experimenten echter aan dat filteren met een mediaanfilter voldoende is (de ruis is er namelijk eerder impuls-achtig) voor betrouwbare vervagingsschatting. In het algemene geval zal een mediaanfilter nochtans niet voldoende zijn, en zal ruis de vervagingsschatting verstören. Er zijn daarvoor twee redenen. De eerste reden is dat randen niet betrouwbaar worden gedetecteerd in aanwezigheid van ruis. De tweede reden is dat de vervagingsexponenten voor gedetecteerde randen verstoord zijn, omdat beeldkenmerken door de aanwezigheid van ruis ruwer lijken dan ze in feite zijn.

In [Canny, 1986] werd het probleem van randdetectie in aanwezigheid van ruis aangepakt door Gaussiaanse vergladding. In [Elder and Zucker, 1998] werd deze aanpak ingebouwd in een probabilistisch schaal-ruimte kader en er werd

een uitdrukking afgeleid voor hoeveel vergladding nodig is om een rand betrouwbaar te kunnen detecteren in aanwezigheid van ruis. Deze hoeveelheid vergladding hangt af van het contrast (of de amplitude) van de rand, de oorspronkelijke randgladheid en het ruisniveau.

Deze minimale hoeveelheid vergladding die nodig is, wordt de *minimum betrouwbare schaal* voor dat beeldkenmerk genoemd. Om deze waarde te kunnen berekenen is het nodig de randkarakteristieken en het ruisniveau te kennen (wat vaak niet het geval is). In de praktijk worden de parameters geschat met een iteratieve procedure, wat vrij tijdrovend is.

Dit heeft geleid tot onze eerste poging om de vervagingsschatting in aanwezigheid van ruis te stabiliseren. We wilden weten hoeveel vergladding er nodig is om een betrouwbare vervagingsschatting uit te kunnen voeren in aanwezigheid van ruis. Deze extra vergladding werd toegepast als een voorbewerking bij onze eigen vervagingsschatting. De aldus geïntroduceerde extra vervaging werd op het einde terug afgetrokken van de geschatte vervaging in het beeld $\hat{\sigma}_{blur}$ om de echte vervaging te bekomen $\sigma_{blur} = \sqrt{\hat{\sigma}_{blur}^2 - \sigma_{postblur}^2}$.

In onze experimenten zagen we dat het zwaartepunt (*center of gravity*, CG) van het histogram van vervagingsexponenten voor de gevonden randen een betrouwbare parameter is om parametrisch gemodelleerde vervaging in een beeld te schatten. De standaardafwijking van de schatting groeit echter wel wanneer de echte vervaging toeneemt. Extra vergladding (in feite vervaging) toepassen reduceert dus de gevolgen van ruis in onze schatting in zekere mate, maar niet voldoende.

In zijn actuele vorm wordt de stabilisatie van de vervagingsschatting uitgevoerd met behulp van een klein heuristisch trucje: zoals beschreven in Sectie S.6 werd de vervagingsschatting geïntegreerd in de restauratiemethode. Door de ruisreductie, de vervagingsschatting en de verscherping te integreren werd de vervagingsschatting efficiënter gestabiliseerd dan door enkel extra vergladding toe te passen.

S.5.2.4 Evaluatie van Wave-Track

We hebben deze methode geëvalueerd op een set van natuurlijke testbeelden die kunstmatig werden gedegradeerd met verschillende hoeveelheden vervaging en ruis. Onze methode heeft de neiging om de vervaging systematisch te overschatten bij hogere ruisniveaus. Een mogelijke verklaring daarvoor is dat bij deze ruisniveaus de ruisreductie in SPERRIL niet alleen de ruis reduceert, maar ook het beeld wat vervaagt. Omdat we verondersteld hebben dat de ruis Gaussiaans is en het ruisniveau bekend was (of anders toch eenvoudig geschat kan worden [Olsen, 1993, Donoho and Johnstone, 1994]), kunnen we deze kennis gebruiken om die systematische afwijkingen in functie van de ruis te compenseren.

S.5.2.5 Toekomstig werk: uitbreiding van Wave-Track naar niet-isotrope PSF's.

Onze huidige implementatie veronderstelt isotrope vervaging, maar kan in principe worden uitgebreid naar anisotrope vervaging, zoals we nu zullen uitleggen. Voorlopige resultaten zijn veelbelovend, maar de volledige uitwerking is moeilijk omdat het volgen van de wavelet-modulusmaxima doorheen de schalen voor de verschillende oriëntaties een ingewikkelde boekhouding vereist.

Deze procedure voor de vervagingsschatting wordt dan onafhankelijk toegepast op de vier oriëntatiesubbanden van de piramide. Veronderstellend dat de PSF Gaussiaans (maar niet meer axiaal symmetrisch) en over het hele beeld invariant is, kunnen we de vervagingsexponenten uitmiddelen per oriëntatieband om zo een robuuste schatting van de breedte van de projectie van de PSF in de richting loodrecht op de oriëntatie van de subband te bekomen. Op die manier bekomen we een aantal waarden voor de breedte van de PSF geprojecteerd onder de verschillende oriëntaties. Uit deze projecties van de PSF is het dan mogelijk een terugprojectie van de echte PSF uit te voeren. In ons geval hebben we verondersteld dat de PSF een anisotrope Gaussiaanse functie was; in dat geval kan er zelfs een analytische vorm voor de PSF berekend worden.

S.5.3 Sparse-Max: waveletgebaseerde ijlheidsmaximalisatie.

Deze methode is gebaseerd op een idee van prof. dr. J. Portilla (DECSAI, Universidad de Granada, España), en werd uitgewerkt onder zijn supervisie tijdens mijn verblijf in Granada in 2004. De methode steunt op de maximalisatie van de ijheid van de waveletcoëfficiënten in de fijnste resolutieschalen van de waveletdecompositie. Om de ijheid van de wavelet subbanden te meten werd de kurtosis gebruikt. Deze methode heeft tot dusver geleid tot één conferentie-artikel [Rooms et al., 2004].

S.5.3.1 Principe van Sparse-Max

De aanpak die bij deze methode werd gevuld is het gebruik van een snelle, niet-blinde, geregelariseerde deconvolutiemethode (we hebben een eenvoudig Wienerfilter gebruikt) om schattingen te bekomen van het gerestaureerde beeld voor een groot bereik van de PSF parameter (de σ_{blur} van een isotrope 2D Gaussiaanse functie in dit geval). Bij het kijken naar de kenmerken van de schattingen van de (kandidaat) gerestaureerde beelden zoeken we naar een criterium dat bepaalt welk beeld "in focus" is. "In focus" heeft hier echter een andere betekenis dan in een fysisch focussysteem: in een fysisch focussysteem is het voldoende om te bepalen wanneer het beeld maximale scherpte bezit (met andere woorden, wanneer het een maximale energie in de hoge frequenties bezit). Wanneer we echter digitale deconvolutie toepassen, kan restauratie met

een “te grote” PSF aanleiding geven tot oververscherping en ringing rond de beeldranden.

In deze subsectie stellen we een waveletgebaseerd criterium voor om de optimale PSF parameter te kiezen. Voor signalen met sterke plaatselijke discontinuïteiten (zoals randen, hoeken, etc., in beelden) zijn waveletvoorstellingen krachtige, decorrelerende hulpmiddelen. Ook heeft men in het laatste decennium uitgebreid bestudeerd hoe een mengende lineaire transformatie (convolutie met de PSF in ons geval) de onderlinge informatie tussen de gemengde elementen vergroot, omdat een pixelintensiteit daarbij eigenlijk vervangen wordt door een soort van gewogen gemiddelde van de pixelintensiteiten uit de omgeving, en zo informatie deelt met zijn buren.

Een eerste benadering om uit gemengde waarnemingen de originele, ongemengde waarnemingen te recupereren, zou het toepassen van *Principale Componenten Analyse* (PCA) kunnen zijn. PCA is een wiskundige methode die een aantal gemengde en dus gecorreleerde variabelen in een (kleiner) aantal ongecorreleerde variabelen (of Principale componenten) transformeert. Een gelijkaardig probleem is *blinde bronseparatie* (*blind source separation*), waarbij een aantal lineaire combinaties (mengsels) van een aantal bronsignalen worden waargenomen, en men uit die lineaire combinaties de originele signalen probeert te bepalen. Een voorbeeld is het *cocktailparty probleem*, waar een aantal mensen door elkaar spreken, en men probeert te herkennen wat de afzonderlijke sprekers zeggen. Het is echter gebleken dat het toepassen van PCA op de gemengde waarnemingen niet voldoende is om de oorspronkelijke (niet-gemengde) elementen te recupereren.

Men kan dit probleem oplossen onder zekere voorwaarden door de inverse te zoeken van deze mengende lineaire transformatie, zodat de onderlinge informatie van de getransformeerde observaties wordt geminimaliseerd via *Onafhankelijke Componenten Analyse* (*Independent Component Analysis*, ICA). In de praktijk is een strikte minimalisatie vinden een hele uitdaging. Daarom werden er een aantal eenvoudiger benaderingen voorgesteld [Comon, 1994, Cardoso and Laheld, 1996, Bell and Sejnowsky, 1997, Hyvärinen, 1999]. In het geval van beelden werd reeds opgemerkt dat waveletrepresentaties zeer “ijl” zijn [Field, 1987, Mallat, 1989, Olshausen and Field, 1997]. Dat betekent dat de significante beeldkenmerken geconcentreerd in een relatief klein aantal coëfficiënten⁷ worden voorgesteld. Losweg betekent in deze context dat een hoge ijlheid van de waveletcoëfficiënten een teken is van kleine onderlinge informatie tussen de coëfficiënten. In andere woorden, door een maat voor de ijlheid te maximaliseren, verminderen we onderlinge informatie tussen de pixels.

Gebruik makend van deze achtergrond is dit werk verder gebaseerd op volgende twee vaststellingen:

1. vervaging, als een lokale lineaire menging vermindert meestal de ijlheid van de wavelet subbanden;

⁷In ons werk hebben we gebruik gemaakt van de à trous Haar wavelet, die de eenvoudigste decorrelerende multischalige transformatie is.

2. de vervaging overcompenseren (wat wil zeggen dat geprobeerd wordt een vervaging sterker dan de echte vervaging te inverteren) heeft typisch hetzelfde effect op de ijheid van de wavelet responsen.

Met andere woorden: de coëfficiënten van de waveletdecompositie zijn maximaal geconcentreerd als het beeld “in focus” is. Als maat voor de ijheid hebben we gekozen voor de kurtosis van de wavelet subbanden [Comon, 1994, Olshausen and Field, 1997, Bell and Sejnowsky, 1997], die een maat is die typisch aangeeft hoe gepiekt (of hoe sterk geconcentreerd rond de oorsprong) een distributie is. De kurtosis van een reeks datapunten y_i , $i = 1, \dots, N$ is gedefinieerd als:

$$\kappa = \frac{\sum_i (y_i - \mu)^4}{N\sigma^4},$$

met μ de verwachtingswaarde van de y_i en σ hun standaardafwijking. We hebben empirisch getest dat wanneer een beeld niet gedegradeerd is, de waveletsubbanden typisch maximale kurtosis hebben vergeleken met hetzelfde beeld waarop een licht laagdoorlaatfilter of het inverse ervan toegepast werd.

We zouden als referentie het liefst de fijnste detailsubbanden gebruiken die relatief ruisvrij zijn. De reden is dat we de wavelet subbanden willen gebruiken die zoveel mogelijk door de vervaging zijn beïnvloed. Empirisch kiezen we als referentiefrequentie de hoogst mogelijke ruimtelijke frequentie f_{ref} waarvoor de waargenomen energie in het spectrum 4.5 dB boven het ruisniveau ligt. We kiezen dan de wavelet schaal met ruimtelijke frequenties groter dan f_{ref} , en de schaal met ruimtelijke frequenties kleiner dan f_{ref} .

We noemen de fijnste van deze twee naburige schalen HB en de ruwste LB ; hun maximale banddoorlaatfrequenties zijn respectievelijk f_{HB} en f_{LB} . Telkens werd de kurtosis berekend voor deze twee schalen uit de decompositie van het beeld, gerestaureerd met het Wienerfilter met een bepaalde kandidaat PSF. Als $\sigma_{max,HB}$ de waarde van de PSF parameter is die de kurtosis in schaal HB maximaliseert, en $\sigma_{max,LB}$ de waarde die de kurtosis in schaal LB maximaliseert, dan nemen we een gewogen gemiddelde van deze twee waarden, waarbij we meer gewicht geven aan de schaal waar f_{ref} zich dichter tegen bevindt.

S.5.3.2 Evaluatie en vergelijking

Zoals bij Wave-Track hebben we deze methode geëvalueerd op dezelfde testbeelden, weer gedegradeerd met een bereik van verschillende hoeveelheden vervaging en ruis. Zoals reeds eerder werd vermeld, hebben we een Wienerfilter gebruikt om de schattingen van de beeldrestauraties te bekomen voor elke waarde van de PSF- en ruisparameters. Sparse-Max gaf bevredigende resultaten voor dit bereik van vervagings- en ruiswaarden.

Ter vergelijking werd de methode van Savakis [Savakis and Trussell, 1993a] geïmplementeerd. Die methode bepaalt welke PSF uit een zekere set kandi-

daat PSF's het verschil minimaliseert tussen het spectrum van het beeld geres-taureerd met de kandidaat PSF en het spectrum zoals het verwacht wordt na Wiener restauratie gegeven een *a priori* model voor het energiespectrum van het beeld.

Savakis' methode heeft de neiging om de vervaging voor alle vervagings- en ruisniveau te onderschatten, wat gedeeltelijk kan worden gecompenseerd door 0.5 op te tellen bij σ_{blur} . De afwijkingen in Savakis' methode hangen echter af van zowel de vervagings- en de ruisniveaus, en kunnen daarom niet zomaar worden gecompenseerd.

Sparse-Max langs de andere kant heeft de neiging om de vervaging systematisch te overschatten bij hogere ruisniveaus (de reden is niet helemaal duidelijk, en moet in de toekomst nog verder worden onderzocht). Omdat we het ruisniveau als gekend hebben beschouwd, kunnen we in functie daarvan echter wel compenseren. We hebben een empirische compensatie gevonden in functie van het ruisniveau (onafhankelijk van de vervaging) die de systematische afwijking bijna volledig compenseert (dit in tegenstelling tot de methode van Savakis). Dit is een belangrijk voordeel, maar doordat er toch nog wat variantie op de schattingssresultaten zit, is de MSE waarde over alle vervagings- en ruisniveaus toch niet spectaculair veel beter.

S.5.3.3 Conclusie

We hebben een eenvoudig en robuust criterium gebruikt om een PSF die beschreven wordt met één parameter te schatten uit een enkele waarneming van het gedegradeerde beeld. Door het maximaliseren van de kurtosis van bepaalde schalen om de meest waarschijnlijke PSF te kiezen, behaalden we bevredigende resultaten voor verschillende voorbeelden van natuurlijke beelden voor een groot bereik van vervagings- en ruisniveaus. Sparse-Max is robuust, heeft een betrekkelijk kleine variantie en een zeer kleine bias. Deze recentere methode is zelfs iets beter dan Wave-Track op gebied van de MSE. Als een nevenproduct geeft deze methode een semi-blind Wiener schatting van het originele beeld, die natuurlijk sterk kan verbeterd worden door de geschatte PSF te gebruiken in een krachtigere restauratiemethode.

S.5.4 Toepassingen van vervagingsschatting

S.5.4.1 Beeldrestauratie

De reden dat we op zoek zijn gegaan naar methoden om de PSF te schatten was om een gedegradeerd beeld te kunnen restaureren. In sectie S.6 bespreken we hoe we Wave-Track hebben ingebouwd in een beeldrestauratiemethode.

S.5.4.2 Autofocus

We hebben ook het probleem bestudeerd hoe de PSF te schatten voor out-of-focus vervaging. Dit soort vervaging doet zich voor bij focusseringstoestellen van beeldverbeeldingssystemen (zoals scherpstellen van een digitale camera, een microscoop, ...). Dit soort vervaging wordt gemodelleerd door een uniforme schijfachtige PSF [Pavlović and Tekalp, 1992, Savakis and Trussell, 1993a, Savakis and Trussell, 1993b]. Daarbij stellen we dat r_{focal} de straal van deze focale vlek voorstelt. Deze uniforme schijf wordt ook soms de Verwarringsschijf (*Circle of Confusion*) [Vivirito et al., 2002] genoemd.

Om r_{focal} te schatten uit een opgenomen beeld, moesten we Wave-Track aanpassen voor synthetische out-of-focus vervagingsschatting. Om Wave-Track aan te passen aan dit soort vervaging, moesten we de Gaussiaanse PSF's vervangen door uniforme schijf-PSF's, en dus het schatten van σ_{blur} vervangen door het schatten van r_{focal} (Sparse-Max kan op dezelfde manier worden aangepast naar dit soort vervaging toe).

Omdat we met een ander soort vervaging te maken hebben (namelijk out-of-focus vervaging), hebben we een andere relatie tussen het zwaartepunt van het histogram van vervagingsexponenten en de vervagingsparameter (hier nu r_{focal}), in dit geval gegeven een empirische polynomiale fit.

In de meeste autofocusstoepassingen is het niet nodig om de PSF te schatten, maar om te bepalen of de camera een beeld al dan niet scherp afbeeldt. Omdat we hier met een fysisch focussysteem te maken hebben, kunnen we dat nagegaan door de absolute waarde van de beeldgradiënt te sommeren over het hele beeld, zoals we reeds hebben vermeld in subsectie S.5.3.1). Wanneer die som maximaal is, dan is het resulterende beeld in focus. Het is dan wel nodig om het hele bereik van focusafstanden stapje per stapje te controleren om te zien of dat een scherp beeld oplevert of niet. Daarom zou het interessant zijn om te weten hoever de focusafstand in één keer moet verplaatst worden om direct de juiste correctie te kunnen uitvoeren. Dat zou kunnen door r_{focal} te schatten. Als we dan de openingshoek van het objectief kennen, kunnen we met driehoeksmeetkunde berekenen hoeveel de focusafstand moet worden bijgesteld om een scherp beeld te bekomen.

Er is dan wel nog de dubbelzinnigheid of het beeldvlak zich voor of achter het focaalvlak bevindt. Daarvoor zullen we toch een tweede testopname moeten maken of de toestand verbeterd of verslechtert.

S.5.4.3 Diepte uit defocus

Verder is het nog mogelijk om Wave-Track (en natuurlijk ook Sparse-Max) te gebruiken voor “diepte uit defocus” schattingen [RMA Belgium, WWW, Schechner and Kiryati, 1998]. Als een camera met een vaste focusafstand gebruikt wordt, produceert een object A op een afstand d_A van de camera een scherp beeld, terwijl een object B op afstand d_B onscherp wordt afgebeeld. Als

men dan scherpstelt op object B, wordt object A dan weer onscherp afgebeeld. Hoe groter de afstand van een object tot de positie waar het scherp zou worden weergegeven, hoe sterker het object vervaagd is in het beeld.

De Diepte uit Defocus methode maakt gebruik van de relatie tussen de afstand van een object tot de objectief lens, de optische parameters van de camera en de hoeveelheid vervaging van beeldkenmerken op verschillende plaatsen in het beeld om diepte in het beeld te schatten. Omdat vervaging in een beeld zowel kan veroorzaakt kan zijn door het beeldvormingsproces als door de omgeving zelf (zoals bijvoorbeeld door trillende hete lucht tussen waarnemer en camera), moeten er minstens twee beelden worden opgenomen met verschillende camera-instellingen (verschillende focus) om deze dubbelzinnigheid te elimineren.

S.6 SPERRIL: schatting en restauratie van confocale beelden

S.6.1 Inleiding

In deze Sectie beschrijven we onze eigen nieuwe methode voor gezamenlijke degradatieschatting en restauratie van confocale microscoopbeelden. De waargenomen beelden zijn gedegradeerd door twee oorzaken: ten eerste is er de vervaging door de bandbeperkte aard van het optische systeem (gemodelleerd door de PuntSpreidingsFunctie of PSF); ten tweede wordt de waarneming verstoord door Poissonruis door de discrete aard van het fotondetectieproces. De voorgestelde methode itereert tussen volgende stappen: ruisreductie, vervagingsschatting en verscherping. Deze stappen worden toegepast in twee fasen, namelijk een trainingsfase en een restauratiefase.

In de trainingsfase worden de drie stappen geïtereerd tot de vervagingsschatting convergeert. Ruisreductie en vervagingsschatting worden uitgevoerd met stuurbare piramiden, terwijl de verscherping wordt uitgevoerd door het Richardson-Lucy algoritme.

De tweede fase is de eigenlijke restauratiefase. In deze fase wordt de uiteindelijke vervagingsschatting uit de trainingsfase gebruikt voor de eigenlijke restauratie, maar nieuwe vervagingsschattingen worden ook berekend om de vooruitgang van de restauratie op te volgen. De iteraties worden gestopt wanneer de vervagingsschatting convergeert, wat in de praktijk altijd wel gebeurt.

Deze geïntegreerde methode is automatisch, en vereist geen *a priori* informatie over het beeld. We hebben onze methode SPERRIL genoemd: Stuurbare Piramide gebaseerde Estimatie en Regularizatie van Richardson-Lucy restauratie. vergeleken met bestaande technieken (zowel via objectieve maten als via visuele inspectie) presteert SPERRIL beter.

Deze Sectie is als volgt opgebouwd: eerst wordt een overzicht van SPERRIL gegeven. We zullen het in het bijzonder hebben over de ruisreductie, de schatting

van de beeldvervaging, de verscherpingsstap en het stopcriterium. Vervolgens worden enkele experimentele resultaten besproken. Tenslotte wordt er een conclusie gegeven.

S.6.2 Overzicht van het algoritme

Het overzicht van algoritme is als volgt opgebouwd: eerst wordt het beeld voor gefilterd om de Poissonruis te reduceren. Dit voorfilter is equivalent aan de ruisreductie verder in het algoritme, behalve dat hier het ruisreductiefilter wordt toegepast op de Anscombe-getransformeerde van het beeld om de Poissonruis in Gaussiaanse ruis te transformeren [Starck et al., 2000]. In verdere iteraties wordt deze transformatie weggelaten. De ruisreductiestap in het algoritme wordt uitgevoerd door de stuurbare piramide decompositie van het beeld te berekenen en dan ruisreductie toe te passen op de subbanden. De volgende stap is de schatting van de beeldvervaging uit de (gedeeltelijk) ruisgereduceerde subbanden. Na de ruisreductie worden de gefilterde subbanden van de stuurbare piramide weer samengevoegd, wat resulteert in een beeld met gereduceerde ruis, maar nog steeds vervaagd. Tenslotte wordt Richardson-Lucy verscherping toegepast. Deze stappen worden geïtereerd.

De vervagingsschatting is in het begin onstabiel (daarom is het geschatte gerestaureerde beeld dat wordt verkregen in deze fase onbetrouwbaar, en wordt het genegeerd bij de eigenlijke restauratiefase), maar convergeert typisch na een tweetal iteraties en wordt dan gebruikt om een PSF te genereren voor de verscherpingsstappen.

Na deze trainingsfase worden de iteraties opnieuw gestart met het oorspronkelijke gedegradeerde beeld en de stabiele schatting van de PSF, waarbij de vervagingsschatting enkel nog gebruikt wordt als een kwaliteitsmaat tijdens de iteraties: wanneer deze maat convergeert worden de iteraties beëindigd. De stuurbare piramide is reeds besproken in subsectie S.4.3. We gaan nu in op de verschillende stappen van SPERRIL.

S.6.2.1 Ruisreductiestap

De waveletgebaseerde ruisreductiemethode die werd beschreven in [Donoho and Johnstone, 1994, Donoho and Johnstone, 1995] werd reeds besproken in subsectie S.4.4.5. Deze methode werd reeds toepast in microscopie door [Boutet de Monvel et al., 2001] en [Stollberg et al., 2003]. Voor berekening van de globale SURE drempel moeten de coëfficiënten per subband echter gesorteerd worden, waar $O(N \log N)$ operaties voor nodig zijn (N is het aantal coëfficiënten per subband) [Mallat, 1999], wat vrij rekenintensief is.

Deze ruisreductie is daarenboven suboptimaal omwille van twee redenen:

1. Ze veronderstelt Gaussiaanse ruis. In confocale fluorescentiebeeldvorming is de belangrijkste bron van fouten Poissonruis [Pawley, 1995,

van Kempen et al., 1997]. In tegenstelling tot Gaussiaanse ruis is Poissonruis afhankelijk van de onderliggende beeldintensiteit, waardoor het scheiden van ruis en nuttige beeldkenmerken zeer moeilijk is.

2. Ze maakt enkel gebruik van het feit dat nuttige coëfficiënten een grote amplitude moeten hebben, en niet van de informatie die bevat is in de ruimtelijke correlatie tussen de waveletcoëfficiënten, noch van de informatie bevat in de correlatie tussen de verschillende schalen (zie weer Figuur S.3).

We pakken deze problemen aan op de volgende manier. Om met de Poissonruis om te gaan, gebruiken we de Anscombe-transformatie [Starck et al., 2000]), die de Poisson data transformeert in data met Gaussiaanse ruis met standaardafwijking één. Deze transformatie stelt ons in staat om om welbekende methoden voor Gaussiaanse ruisreductie toe te passen op data met de heel wat lastiger Poissonruis. We hebben gekozen om de bivariate waveletgebaseerde methode voor ruisreductie uit [Şendur and Selesnick, 2002] toe te passen; deze werd reeds besproken in subsubsection S.4.4.5.

We hebben dit algoritme aangepast voor gebruik met stuurbare piramiden. De eerste keer wordt de ruisreductie gecombineerd met de Anscombe-transformatie om de Poissonruis te reduceren. Tijdens het verdere verloop van het algoritme wordt de ruisreductie toegepast als regularisatie (onder de Anscombe-transformatie) na elke verscherpende stap.

S.6.2.2 Vervagingsschatting

Er bestaan twee veelgebruikte methoden om de PSF van een confocale microscoop te bepalen. De eerste methode steunt op het fysisch model van de beeldvorming. Wanneer alle parameters (zoals de numerieke aperture van het objectief, de brekingsindex, ...) bekend zijn dan kan de PSF berekend worden zoals beschreven werd in subsection S.2.4. Deze methode kan echter geen rekening houden met onbekende gebreken bij de optische afstellingen door bijvoorbeeld temperatuurfluctuaties, variaties in de brekingsindex in het specimen, ... (zie ook subsection S.2.6 voor meer uitleg over aberraties). Deze aberraties kunnen leiden tot afwijkingen tussen de berekende en de echte PSF. Sommige van deze afwijkingen kunnen worden gemodelleerd, maar ze zijn niet altijd bekend in detail.

De tweede methode is om een ijkbeeld op te nemen van een fluorescerend bolletje met een grootte rond de resolutielimiet van de microscoop. Methoden die zulk een ijkbeeld gebruiken om de PSF te bepalen zijn zeer betrouwbaar en laten toe om betere restauratieresultaten te bekomen dan mogelijk was met de theoretische PSF [van der Voort and Strasters, 1995, McNally et al., 1998], omdat ook de onbekende aberraties impliciet bijdragen tot het beeld van het bolletje, terwijl ze niet in rekening kunnen worden gebracht in het theoretisch model.

Het opnemen van beelden van zulke bolletjes vraagt echter heel wat extra tijd en deze ijkbeelden moeten regelmatig opnieuw opgenomen worden omdat de afstelling van de apparatuur degradeert in de tijd. Ook nemen onderzoekers niet altijd tijd om deze opnamen te maken om verschillende redenen. Confocale microscopen zijn namelijk zeer duur, en worden door verschillende onderzoeksteams afwisselend gebruikt; de tijd dat de microscoop kan gebruikt worden door een bepaalde onderzoeker is daarom beperkt.

Daarom hebben we Wave-Track gebruikt, onze eigen methode voor vervagings-schatting (zie subsectie S.5.2). Daardoor hebben we geen ijkbeeldje nodig, en de geschatte PSF past zich beter aan aan de beelddegradatie dan de methode met de theoretische PSF. De methode voert haar berekeningen uit in het stuurbare-piramide-domein (dus zijn er geen extra transformaties nodig na de ruisreductie/regularisatie).

Voor deze nieuw methode veronderstelden we dat de PSF kan worden gemodelleerd door een 2D Gaussiaanse functie. Men kan zich natuurlijk afvragen hoe nuttig een Gaussiaans PSF model is in het geval van microscopie, omdat dit type PSF in theorie niet bandbeperkt is, terwijl een optisch systeem wel bandbeperkt is. Dat is inderdaad zo wanneer het beeld enkel vervaagd is. In elk beeldvormingssysteem treedt echter ook ruis op, en deze ruis zorgt in feite voor een limiet voor welke frequenties in het gedegradeerde beeld kunnen worden gerecupereerd. Binnen deze frequentieband waar het signaal groter is dan het ruisniveau is het Gaussiaanse PSF model een redelijke benadering.

Deze aanpak heeft enkele onmiskenbare voordelen, ook al gaat ze uit van een sterk vereenvoudigd PSF model:

- er is geen extra opname van een ijkbeeldje nodig om de PSF te modelleren;
- de geschatte PSF past zich automatisch aan aan extra degradaties door aberraties, zonder dat deze expliciet gekend moeten zijn.

Natuurlijk kan men toch de echte opgenomen PSF gebruiken in ons algoritme als deze beschikbaar is; in dat geval moeten we zelf geen synthetische PSF genereren, en kunnen we onze vervagingsschatting gebruiken om de evolutie van de beeldscherpte tijdens de restauratie op te volgen.

S.6.2.3 Verscherplingsstap

Voor de eigenlijke verscherplingsstap hebben we het Richardson-Lucy (RL) algoritme gebruikt, dat reeds werd besproken in subsectie S.3.2.2. Dit algoritme is reeds goed ingeburgerd bij onderzoekers die werken met confocale beelden [Holmes et al., 1995, van Kempen et al., 1997, Verveer et al., 1999]. In diezelfde subsectie werden ook reeds enkele klassieke regularisatiemethoden voor RL vermeld. Deze regularisatiemethoden worden in SPERRIL vervangen door waveletgebaseerde ruisreductie, zoals beschreven in subsectie S.6.2.1.

S.6.2.4 Stopcriterium

In deze subsectie bespreken we kort wanneer de iteraties moeten worden beëindigd. We merkten al op dat de geschatte vervagingsmaat afneemt naarmate het aantal iteraties toeneemt (zoals verwacht), en naarmate de aannameelikheid van de oplossing toeneemt. Als een empirische vuistregel kiezen we daarom om de iteraties te stoppen wanneer de verhouding van de vervaging geschat tijdens de vorige iteratie en de vervaging geschat tijdens de huidige iteratie convergeert naar 1.

S.6.2.5 Experimentele resultaten

(a) Synthetisch beeld

Een eerste experiment was de restauratie van een synthetisch beeld met gesimuleerde degradatie. Voor dit experiment was het ideale beeld beschikbaar, dus kon de Piek Signaal Ruis Verhouding (*Peak Signal to Noise Ratio*, PSNR) berekend worden. Bij het gebruik van klassieke restauratie is er reeds een opmerkelijke verhoging in PSNR. Met RL-SURE (=RL gecombineerd met wavelet SURE thresholding, [Boutet de Monvel et al., 2001]) is er nog een lichte verbetering merkbaar in vergelijking met klassieke restauratie. Bij SPERRIL is er weer een lichte verbetering in PSNR in vergelijking met RL-SURE.

Als een test voor de resolutie van de methoden namen we een synthetisch rooster en pasten er synthetische degradatie op toe. Het rooster bevat lijnen van verschillende breedte, gescheiden door verschillende breedten achtergrond. Het rooster werd gedegradeerd (Gaussiaanse PSF, $\sigma_{blur} = 2$ pixels, Poissonruis 120 fotonen per pixel) en daarna gerestaureerd met de verschillende methoden: klassieke RL, RL-SURE en SPERRIL. Alle methoden geven gelijkaardige waarden voor de Volle Breedte op het Halve Maximum (*Full Width at Half Maximum*, FWHM) van het lijnprofiel. Daarom concluderen we dat de algoritmen in termen van resolutieverbetering vergelijkbaar zijn. Dat was natuurlijk te verwachten, aangezien de verscherpende stap in alle beschouwde algoritmen dezelfde is (namelijk RL).

Toch merkten we op dat SPERRIL beter artefacten zoals ringing in de achtergrond tussen de lijnstructuren onderdrukt, ten koste van lichtjes bredere staartjes van de lijnprofielen (vergelijkbaar met de profielen van RL-SURE). Ook heeft SPERRIL de neiging om beeldintensiteiten lichtjes te onderschatten, terwijl de andere methoden in het algemeen de neiging hebben om (vooral de hoogste) intensiteiten te overschatten.

(b) Biologische experimenten

De biologische geldigheid testen van de resultaten van eender welk restauratie-algoritme is altijd een hele uitdaging, omdat men geen referentiedata bezit over beeldkenmerken van de grootte-orde van de resolutielimiet van microscopen. Daarom hebben we gekozen om de verschillende algoritmen voor beeldrestauratie toe te passen als een voorbewerking, en dan de restauratieresultaten te

evalueren aan de hand van de resultaten van colocalisatie-analyse. In de volgende twee paragrafen wordt het principe van colocalisatie besproken. De uitleg is gebaseerd op [Manders et al., 1993, Casavan and Gaidoukevitch, 2003].

Specimens kunnen worden gekleurd met één of meerdere fluorochromen. Fluorochromen met één enkel paar van excitatie/emissie golflengten worden vaak gebruikt om celstructuren te visualiseren zoals microtubuli (buisvormige celstructuren die als transporttrails in de cel en als celskelet dienen), mitochondriën (voor de energievoorziening van de cel) of celkernen. Dat levert geen problemen als de waarnemer enkel geïnteresseerd is in een enkele specifieke structuur in het gezichtsveld. Vaak worden er echter twee of meer labelende kleurstoffen gebruikt om meerdere structuren tegelijk te bestuderen, maar ze toch gescheiden te houden (in verschillende golflengtebanden). Het specimen wordt dan onafhankelijk opgenomen in elke fluorescentiegolflengte en dan gecombineerd met de andere kleurkanalen.

Colocalisatie beschrijft de aanwezigheid van twee of meer soorten moleculen op dezelfde fysische lokatie. De reden zou kunnen zijn dat de fluorochromen bevestigd werden aan dezelfde receptor, of bevestigd zijn aan moleculen die met elkaar interageren. Bij beeldverwerking betekent dat dat intensiteiten uitgezonden door de verschillende fluorochromen bijdragen tot dezelfde pixel in het beeld.

In [Landmann, 2002] wordt beeldrestauratie als voorbewerking van colocalisatie-analyse gebruikt, waarbij geconcludeerd wordt dat er een belangrijke verbetering van de analyseresultaten verkregen wordt na beeldrestauratie. Ons werk werd uitgevoerd in samenwerking met dr. Diane S. Lidke van het Departement Moleculaire Biologie van het Max Planck Instituut voor Biofysische Chemie (Göttingen, Duitsland). Zij was zo vriendelijk om de beelden te leveren, en nadien de resultaten van de beeldverwerking te evalueren.

We evalueren de prestaties van SPERRIL door middel van colocalisatie-analyse van twee fluorescerend gelabelde moleculen in menselijke A431 huidkankercellen [Lidke et al., 2004].

De eerste fluorescerend gelabelde molecule is een transmembraan receptorproteïne, met andere woorden een molecule dat zich uitstrekkt doorheen het celmembraan. Deze receptor kan zich binden met moleculen buiten de cel die informatie brengen naar de cel. De erbB familie van receptoren bevat ook de Epidermale Groei Factor Receptors erbB1 en erbB3 die, als ze worden geactiveerd, de cel doen groeien. De erbB1 receptoren werden gelabeld met Groen Fluorescerend Proteïne (GFP), waardoor erbB1-eGFP gevormd werd; erbB3 werd gelabeld met mCitrine, wat een verbeterde mutant is van geel fluorescerend proteïne.

De tweede molecule die meespeelt is een zogenaamde ligand, een molecule zoals een hormoon die aan de cel bepaalde informatie overbrengt door met de receptorproteïnes te binden. De ligand die hier gebruikt werd is Epidermale Groei Factor (EGF). Wanneer Epidermale Groei Factor Receptor bindt met deze li-

gand krijgt de cel het signaal om te groeien. De ligand EGF werd gelabeld met fluorescerende Quantum Dots, wat resulteerde in EGF-QD.

Eén manier om te controleren of de receptor en de ligand met elkaar binden, is om beide moleculen fluorescerend te labelen zoals hierboven werd beschreven, en dan de colocalisatie (of overlap) van deze twee fluorescerende signalen van de receptor en de ligand in het opgenomen beeld opvolgen [Lidke et al., 2004].

De eerste reeks beelden bevat levende cellen die erbB1-eGFP tot expressie brengen, en waar EGF-QDs aan toegevoegd is. Na activatie door EGF-QDs zal het erbB1 internaliseren, met andere woorden: het wordt getransporteerd van het membraan naar de binnenkant van de cel. Omdat de QD-EGF vast blijft hangen aan het erbB1-eGFP na de internalisatie, zou de colocalisatie hoog moeten zijn.

De tweede soort cellen echter zijn A431 cellen die erbB3-mCitrine tot expressie brengen. In dat geval bindt EGF-QD niet direct met erbB3, maar bindt wel met het oorspronkelijke, ongelabelde erbB1 dat in de cel aanwezig is. Bij activatie van het erbB1 door EGF-QD worden er geen complexen met het gelabelde erbB3 gevormd; daarom zou in deze beelden de colocalisatie laag moeten zijn.

Ons doel is om de colocalisatie in deze twee gevallen experimenteel te bepalen, en te oordelen of de resultaten overeenstemmen met wat we verwachten. We tonen aan dat we een betere overeenkomst verkrijgen na restauratie.

De redenering van [Landmann, 2002] volgend, hebben we colocalisatie-analyse toegepast na beeldrestauratie. In ons geval hebben we de colocalisatie-resultaten vergeleken voor het onbewerkte beeld, met de resultaten na het toepassen van verschillende restauratiemethoden (klassieke RL, RL-SURE en SPERRIL) telkens gecombineerd met standaard achtergrondcorrectie (met ImageJ [ImageJ, WWW], een veelzijdig en gratis software pakket met talloze gratis plugins voor allerlei toepassingen). De verschillende kleurkanalen werden voor de eenvoud telkens onafhankelijk van elkaar bewerkt, en de PSF werd apart voor de verschillende kanalen geschat.

De colocalisatie-analyse levert een aantal parameters op die de graad van overlap tussen de verschillende kleurenbanden weergeven (1 betekent dat overal waar de ene kleur aanwezig is, de andere kleur ook aanwezig is; 0 betekent dat de twee kleuren nergens samen voorkomen). Voor de ruwe beelden is er nauwelijks een verschil te merken tussen de twee celtypes; het gemiddelde verschil in overlap tussen de twee klassen Δ is 0.18. Na restauratie met klassieke RL wordt het verschil tussen de twee klassen duidelijker: $\Delta = 0.22$. Na RL-SURE is $\Delta = 0.26$. Na toepassing van SPERRIL vooraleer de analyse uit te voeren wordt echter het grootste verschil in overlap tussen de twee klassen van cellen waargenomen: $\Delta = 0.30$. Dat is consistent met wat we verwachtten vanuit de onderliggende biochemische processen in de cellen [Lidke et al., 2004].

S.6.2.6 Conclusies

In deze subsectie bespraken we een geïntegreerd, stabiel en automatisch algoritme om gedegradeerde fotongelimiteerde beelden te restaureren. De degradatieparameters worden geschat uit het beeld zelf zonder een ijkbeeld nodig te hebben, wat een belangrijk voordeel is. Langs de andere kant gaat er geen restauratiekwaliteit verloren. De vervagingsschatting en de regularisatie van de restauratie worden berekend in het stuurbare piramide domein. De vervagingsschatting wordt daarnaast gebruikt om te bepalen wanneer de iteraties kunnen worden gestopt, waardoor de restauratie automatisch is.

Het algoritme werd geëvalueerd zowel op synthetische als op echte biologische confocale beelden. Ons algoritme presteert het best, zowel in termen van visuele kwaliteit als in termen van kwantitatieve maten (de PSNR voor de synthetische beelden, de colocalisatie-parameters voor de biologische beelden). Voor gedetailleerde resultaten verwijzen we naar de Engelstalige tekst.

Toekomstig werk zal erin bestaan om SPERRIL uit te breiden naar 3D. Het RL algoritme is reeds beschikbaar in 3D. Ook is de uitbreiding van de regularisatie naar 3D mogelijk op basis van 3D multischalige decomposities, omdat 3D uitbreidingen van stuurbare piramiden reeds bestaan [Delle Luche et al., 2003], zoals ook werd uitgetest in subsectie 4.4.1. De grootste uitdaging zou het schatten van de PSF in 3D worden. De PSF in 2D kan worden benaderd door een Gaussiaan, maar in de z-richting is dit duidelijk niet het geval. In dat geval moeten we een ander, realistischer confocaal PSF model gebruiken waarbij de parameters van het confocaal PSF model kunnen worden aangepast om te compenseren voor afwijkingen bij de geschatte PSF: stel dat bijvoorbeeld een zekere vervaging in de z-richting wordt geschat door onze methode, maar dat het theoretisch model resulteert in een kleinere uitgestrektheid van de PSF. Dan zouden de parameters van het confocale PSF model bijgestuurd kunnen worden om deze afwijking mee in rekening te brengen.

S.7 Multiresolutie-GNC gebaseerde beeldrestauratie

Markov Random Velden (MRV) zijn een krachtig hulpmiddel gebleken in beeldverwerking. In de literatuur zijn heel wat voorbeelden terug te vinden van gebruik van MRV's in beeldanalyse en -verbetering. Later werden de MRV's ook gecombineerd met de principes van multiresolutie, wat voor bepaalde problemen een nog krachtiger hulpmiddel was: er bestaan toepassingen in beeldsegmentatie [Bouman and Shapiro, 1994], ruisonderdrukking [Malfait and Roose, 1997, Jansen and Bultheel, 2001, Pižurica et al., 2002, Pižurica, 2002], ...

S.7.1 Markov Random Velden

MRV's kunnen worden gebruikt om a priori kennis over bepaalde beeld eigenschappen op te leggen zoals gladheid (namelijk dat vrij grote gebieden met gelijkaardige pixelintensiteiten voorkomen) en randcontinuïteit. De oorsprong van MRV's gaat terug tot de statistische mechanica voor ferro-magnetisme; in analogie met magnetische dipolen kunnen naburige pixels invloed op elkaar uitoefenen.

Statistische mechanica wordt formeel beschreven aan de hand van Gibbs random velden (GRV) [Li, 1995, Pižurica, 2002], die als waarschijnlijkheidsmodellen dienen om fluctuaties van grote fysische systemen rond hun evenwicht te beschrijven. Deze GRV's worden gebruikt om de systeemvariabelen globaal te modelleren: ze geven de gezamenlijke waarschijnlijkheid van een MRV. De Gibbs waarschijnlijkheidsfunctie wordt gegeven door: $P(\mathbf{x}) = (e^{-H(\mathbf{x})/T})/Z$, waarin $H(\mathbf{x})$ de *energiefunctie* is; Z wordt de partitiefunctie genoemd, welke een normaliserende constante is; T is de *temperatuur*, die in verband staat met de “gepiektheid” van de distributie. Wanneer T kleiner is, is de distributie gepiekter (wat betekent dat er een groter verschil is tussen twee verschillende beeldconfiguraties) waardoor het minimum gemakkelijker te vinden is.

Er bestaat een verband tussen het lokale MRV en het globale GRV. Dat verband laat ons toe om het globale systeem te modelleren door middel van lokale interacties. Dit verband tussen Gibbs en Markov random velden zijn de zogenaamde *klieken*. Een kliek is een groep elementen in het rooster, waarbij alle elementen directe *buren* zijn.

Een Markov Random Veld is een speciaal type Gibbs Random Veld, waarvan de energie kan worden uitgedrukt als een som van *kliekpotentiaLEN* $V_C(\mathbf{x})$: $H(\mathbf{x}) = \sum_{C \in \mathcal{C}} V_C(\mathbf{x})$, met \mathcal{C} de verzameling van alle mogelijke klieken in het beeld. De kliekpotentiaAL voor een bepaalde kliek in het rooster hangt af van de lokale configuratie van de kliek, en is enkel functie van de leden van de kliek. Onze a priori kennis over de lokale beeld eigenschappen vertaalt zich in het definiëren van de gepaste kliekpotentiaALfuncties.

Om de notaties compact te houden, duiden we elke pixel aan door middel van één enkele index (zoals in het scannen van een rooster). We veronderstellen ook dat er enkel klieken van pixelparen voorkomen. Met $\mathcal{L} = \{1, \dots, n\}$ duiden we de set van indices aan (de set van pixels op een rechthoekig rooster). Verder betekent de notatie $\langle k, l \rangle$ een paar buren (een kliek die bestaat uit pixels k en l) en de notatie $V(x_k, x_l)$ de potentiaAL van kliek $\langle k, l \rangle$.

S.7.1.1 Energieminimalisatie

Als we een gepaste energiefunctie hebben gedefinieerd, hebben we nog een efficiënt minimalisatie-algoritme nodig. Hier beschrijven we kort twee mogelijke algoritmen:

1. het Metropolis-algoritme, dat behoort tot de klasse van de randomzoekalgoritmen;
2. het Gradueel Niet-Convexiteitsalgoritme, dat behoort tot de klasse van de deterministische algoritmen.

(a) Het Metropolis-algoritme

Het Metropolis algoritme [Li, 1995] is een random-zoekalgoritme: sommige roosterelementen worden random verstoord, waarna de posteriori energie van het verstoerde rooster opnieuw wordt berekend. Deze nieuwe configuratie wordt dan aangenomen met een waarschijnlijkheid die afhangt van zowel de energieverandering van het rooster voor deze storing als van de temperatuur van het systeem. Hoe hoger de temperatuur van het systeem, hoe meer verstoringen aanvaard worden die eigenlijk energetisch ongunstig zijn. Het voordeel van random-zoekalgoritmen is dat ze beter kunnen vermijden vast te geraken in lokale minima van de energiefunctioaal (in tegenstelling tot bepaalde andere zoekmethoden). Het nadeel is dat random zoeken zeer rekenintensief is, en traag convergeert.

(b) Graduale Niet-Convexiteitsminimalisatie

De Graduale Niet-Convexiteitsminimalisatiemethode GNC zoals beschreven in [Blake and Zisserman, 1987, Li, 1995], is aantrekkelijk door zijn lage computatieve complexiteit. De naam is afkomstig van het feit dat de kliekpotentialen zo worden gekozen dat de resulterende energiefunctioaal in het begin convex is (maar niet realistisch). Wanneer het minimum van deze convex energiefunctioaal gevonden is, worden de potentialen aangepast zodat de energiefunctioaal minder convex wordt, maar wel al wat realistischer. Daarbij kunnen verschillende lokale minima ontstaan, maar we minimaliseren in de buurt van het vorige gevonden minimum (van de convexere energiefunctioaal). Deze procedure wordt zo voortgezet. Daarbij wordt de energiefunctioaal steeds minder convex, maar we vinden met grote waarschijnlijkheid een oplossing die in de buurt van het globale minimum ligt, doordat we steeds in de buurt blijven van de meer convex gemaakte versies van de energiefunctioaal.

S.7.2 Multiresolutie MRV beeldrestauratie

S.7.2.1 Bestaande methoden

Beeldrestauratie uitgaande van MRV's met een enkele schaal werden reeds uitgebreid beschreven in de literatuur. Daarbij werden verschillende methoden gebruikt om de energiefunctioaal te minimaliseren: de Gibbs sampler [Geman and Geman, 1984, Geman and Reynolds, 1992], GNC [Nikolova et al., 1998, Nikolova, 1999].

De literatuur over beeldrestauratie gebaseerd op Multiresolutie MRV's (MMRV) is echter beperkt. MMRV's werden reeds gebruikt in segmentatie [Bouman and Shapiro, 1994] en ruisonderdrukking [Malfait and Roose, 1997,

Jansen and Bultheel, 2001, Pižurica et al., 2002, Pižurica, 2002] maar heel wat minder voor restauratie van beelden die zowel *vervaagd* zijn als ruis bevatten. Twee voorbeelden zijn [Belge et al., 2000, Robini and Magnin, 2003]. In [Belge et al., 2000] werd een waveletgebaseerde aanpak voorgesteld, gecombineerd met een speciaal geval van de zogenaamde halfkwadratische regularisatie uit [Geman and Yang, 1995]. In [Robini and Magnin, 2003] werd een eenvoudige waveletgebaseerde penalisatieterm voorgesteld die de voorkeur geeft aan de vorming van horizontale en verticale randen (terwijl ze diagonale randen bewaart). De energieminimalisatie werd berekend met het Metropolis algoritme, gecombineerd met gesimuleerde annealing (met andere woorden, men laat de temperatuur dalen tijdens de iteraties van het Metropolis algoritme).

S.7.2.2 Wat achtergrond in verband met MMRV beeldrestauratie

In [Luettgen et al., 1993] werd er gesuggereerd dat een multiresolutieuitbreiding van MRV modellering een superieur middel voor beeldanalyse en -modellering zou zijn. Deze manier van voorstellen zou ook kunnen worden gebruikt om ook relaties tussen waveletcoëfficiënten binnen eenzelfde band, waveletcoëfficiënten uit verschillende resolutieschalen te modelleren.

Voor een directe toepassing van de klassieke MRV restauratie [Geman and Geman, 1984] in het ruimtelijke domein moeten we de volgende posteriori energie minimaliseren:

$$E(\mathbf{f}|\mathbf{g}) = \lambda \sum_{k \in \mathcal{L}} \left((h * f)_k - g_k \right)^2 + \mu \sum_{\langle k, l \rangle \in \mathcal{C}} V(f_k, f_l) \quad (\text{S.5})$$

met $\mathbf{f} = \{f_1, \dots, f_n\}$ het onbekende, niet-gedegradeerde beeld; $\mathbf{g} = \{g_1, \dots, g_n\}$ is het waargenomen, gedegradeerde beeld; h is de PSF en \mathcal{C} is de set van alle mogelijke (paarsgewijze) klieken. De eerste term in de posteriori energie geeft weer hoe nauw de oplossing bij de waargenomen data aansluit; de tweede term legt een randvoorwaarde op aan de gladheid van de oplossing (die een vertaling is van onze a priori kennis over de beeldstructuur). De parameters μ en λ drukken het relatieve gewicht van deze twee termen uit. De twee voornaamste problemen bij deze aanpak zijn:

1. moeilijkheden bij het specifiëren van de parameters μ en λ , oftewel een goede afweging vinden tussen de aanpassing aan de waargenomen data enerzijds, en gladheid van de oplossing anderzijds;
2. de keuze van goede kliekpotentiaLEN $V(x_k, x_l)$ die zich aanpassen aan *discontinuïteiten* door in de omgeving van randen de vervaging “uit te schakelen.”

Er werd een elegante oplossing van deze beide problemen voorgesteld in [Pižurica, 2002]. Daar werd voorgesteld om een waveletgebaseerde (dus ijle) voorstelling te gebruiken om de regularisatiegewichten μ en λ te bepalen. Daarnaast werd er voorgesteld om de energie te minimaliseren met de GNC methode. Als we de waveletgetransformeerde van de niet-gedegradeerde data noteren als $\mathbf{x} = \mathcal{W}(\mathbf{f})$ en de waveletgetransformeerde van de waargenomen, gedegradeerde data als $\mathbf{y} = \mathcal{W}(\mathbf{g})$ ⁸.

In [Pižurica, 2002] werd het volgende idee geïntroduceerd: uitgaande dat er enkel degradatie opgetreden is door de aanwezigheid van ruis, kan de posteriori energie in elke wavelet subband worden geschreven als:

$$E(x_k|\mathbf{y}) = f(\alpha_k) \sum_{k \in \mathcal{L}} (x_k - y_k)^2 + \sum_{\langle k, l \rangle \in \mathcal{C}} f(\alpha_l) V(x_k, x_l), \quad (\text{S.6})$$

waarbij α_k een waveletgebaseerde ‘‘randindicator’’ is voor positie k en $f(\alpha_k)$ een monotoon stijgende functie van α is. De tweede sommatie gaat over alle naburen x_l van een zekere coëfficiënt x_k , die elk gewogen worden met de maat $f(\alpha_l)$ volgens hoe betrouwbaar ze zijn.

Een randindicator is in het algemeen een lokale maat die berekend wordt met behulp van x_k en van zijn naburige coëfficiënten uit dezelfde en/of aangrenzende resolutieschalen. Bovenstaande formulering zorgt voor een plaatsafhankelijke afweging van aanpassen aan de data en de gladheid van de oplossing die afhankelijk is van het al dan niet aanwezig zijn van nuttige beeldinformatie op die plaats. De aanpassing aan discontinuïteiten of andere nuttige beeldkenmerken kan daarbij zelfs met een eenvoudige isotrope potentiaalfunctie gebeuren. In het bijzonder zal $f(\alpha_k)$ die in de eerste term staat meer gewicht geven aan de waargenomen data op positie k als het waarschijnlijk is dat y_k tot een beeldrand behoort. Ook kan er door de gewichten $f(\alpha_l)$ in de tweede term gladheid aan de oplossing worden opgelegd die zich lokaal aanpast aan de aanwezigheid van discontinuïteiten. In die tweede term zullen de naburen y_l die met grotere waarschijnlijkheid met randen overeenkomen een grotere invloed hebben op de schatting van de niet-gedegradeerde naburige coëfficiënt x_k .

S.7.3 Voorgestelde methode

We stellen een nieuwe, MMRV-gebaseerde restauratiemethode voor, die werd ontwikkeld in nauwe samenwerking met dr. ir. Aleksandra Pižurica. Deze nieuwe methode is gebaseerd op de formele beschrijving aan de hand van vergelijking (S.6) uit [Pižurica, 2002]. Die vergelijking breidt de klassieke aanpak voor beeldrestauratie uit naar een multiresolutiemethode die lokaal adaptief is.

⁸In tegenstelling tot elders in deze thesis slaat \mathbf{g} hier op data die *enkel* gedegradeerd is door ruis.

In [Pižurica, 2002] werd dit idee enkel theoretisch en in de context van ruisonderdrukking behandeld; hier stellen we praktisch bruikbaar algoritme voor dat daarenboven ook rekening houdt met beeldvervaging, en dat we kunnen toepassen voor echte beeldrestauratie. We breiden vergelijking (S.6) uit zodat ook het optreden van vervaging mee in rekening wordt gebracht. In dat geval moet x_k uit de eerste term⁹ van vergelijking (S.6) vervangen worden door $(h * x)_k$, waardoor we de volgende uitdrukking bekomen:

$$E(x_k|\mathbf{y}) = f(\alpha_k) \sum_{k \in \mathcal{L}} \left((h * x)_k - y_k \right)^2 + \sum_{\langle k, l \rangle \in \mathcal{C}} f(\alpha_l) V(x_k, x_l), \quad (\text{S.7})$$

We hebben deze methode geïmplementeerd met de stuurbare-piramidevoorstelling, waarbij $f(\alpha_k)$ de waarschijnlijkheid voorstelt dat een bepaalde waargenomen waveletcoëfficiënt te wijten is aan nuttig signaal.

S.7.3.1 Mogelijke randindicatoren

We hebben experimenten uitgevoerd met verschillende mogelijke randindicatoren, zoals de magnitude en de variantie van de waveletcoëfficiënten in een bepaalde omgeving van het beeld. Deze randindicatoren bleken in de praktijk echter niet bruikbaar.

De bivariate waveletinkrimpingsmethode uit [Sendur and Selesnick, 2002] is reeds enkele keren aan bod gekomen in deze thesis. Deze methode is van een multiplicatieve vorm, en is dus te schrijven als: $\hat{w}_1(k) = \mathcal{S}(k) \cdot y_1(k)$, waarbij $\hat{w}_1(k)$ de gefilterde coëfficiënt in een zekere schaal op positie k is, $y_1(k)$ is de waargenomen coëfficiënt gedegradeerd door ruis en $\mathcal{S}(k)$ is de corrigerende inkrimpingsfactor voor $y_1(k)$.

Deze inkrimpingsfactor $\mathcal{S}(k)$ kan worden gezien als een schatting van hoe betrouwbaar of hoe verstoord een zekere coëfficiënt $y_1(k)$ is: wanneer $\mathcal{S}(k)$ dichter bij 1 ligt, is de coëfficiënt waarschijnlijker afkomstig van nuttig signaal dan van ruis; wanneer $\mathcal{S}(k)$ echter dichter bij 0 ligt, is de coëfficiënt met grotere waarschijnlijkheid afkomstig van ruis dan van nuttig signaal.

In onze implementatie hebben we deze functies $\mathcal{S}(k)$ gebruikt als “randindicator” op positie k , die de rol overnemen van $f(\alpha_k)$ in vergelijking S.7. We hebben het algoritme reeds getest op enkele klassieke testbeelden (*Lena*, *Peppers*, ...) die synthetisch werden vervaagd en gecontamineerd met ruis. Restauratie geeft veelbelovende resultaten (zie Hoofdstuk 7 van de Engelstalige tekst), maar is nog vrij onstabiel voor hoge ruisniveaus.

⁹de term die de aanpassing van de oplossing aan de waargenomen data uitdrukt. De regularisatieterm, die onze voorkeur voor de oplossing voorstelt, moet niet worden aangepast: we verwachten nog steeds een beeld dat scherp is aan de randen.

S.7.4 Conclusies en toekomstig onderzoek

In deze Sectie hebben we enkele voorlopige resultaten van onze nieuwe methode gegeven, waarbij werd aangetoond dat het abstracte, theoretisch geformuleerde idee uit [Pižurica, 2002] in de praktijk ook werkt. Het is wel duidelijk dat deze methode nog in de kinderschoenen staat en er nog heel wat werk nodig zal zijn om de volle capaciteiten van deze methode te kunnen benutten. Voorlopige resultaten zijn echter veelbelovend.

Een idee voor toekomstig onderzoek is om de randindicator α_l in vergelijking (S.7) te baseren op de lokale energie en lokale oriëntatie¹⁰ die kunnen worden berekend uit de stuurbare piramide decompositie. Coëfficiënten waarvoor de overeenkomstige lokale energie groot is en die liggen in de richting aangeduid door de geschatte lokale oriëntatie zouden een sterkere invloed hebben op hun buren. Dit is een uitbreiding van de criteria voor randcontinuïteit: in plaats van coëfficiënten enkel op basis van hun magnitude te clusteren, kunnen we hier ook geometrische kenmerken zoals randsterkte en randoriëntatie mee in rekening brengen.

S.8 Conclusies

In het eerste deel werd een algemene situering van het probleem besproken en een indeling van de thesis gegeven. In delen twee tot vier werden in een tutorialvorm volgende onderwerpen achtereenvolgens besproken: ten eerste werd een schets gemaakt van de principes van optische beeldvorming en gevolgen op de beeldkwaliteit; ten tweede werden enkele klassieke restauratiemethoden besproken; het derde en laatste tutorial-hoofdstuk besprak de principes en toepassingen van multiresolutieconcepten in beeldverwerking.

In deel 5 werden twee originele methoden besproken om de vervaging van een beeld te schatten op basis van multiresolutieprincipes. Deze methoden werden getest aan de hand van een reeks beelden die kunstmatig werden gedegradeerd (vervaagd, en ruis toegevoegd). Motivatie om deze methoden te ontwikkelen was om informatie over de beelddegradatie te schatten als deze niet vorhanden is. We hebben deze informatie namelijk nodig om gedegradeerde beelden te kunnen restaureren.

Deel 6 incorporeert als *proof of concept* Wave-Track (de eerst besproken vervagingsschattingsmethode uit deel 5) in een volledige restauratietechniek. Deze techniek werd geëvalueerd aan de hand van synthetische en biologische beelden en blijkt veelbelovende resultaten te produceren in vergelijking met andere technieken.

Deel 7 gaf tenslotte een uitbreiding van een bestaande ruisonderdrukkingstechniek naar het volledige restauratieprobleem (compenseren voor zowel vervaging als ruis). Deze techniek staat nog in haar kinderschoenen, maar biedt wel de

¹⁰zie subsectie S.4.4.2

elegantie dat alle stappen (verscherping en regularisatie) in het waveletdomein worden berekend. Voorlopige resultaten zijn beloftevol.

Tenslotte werd nog eens een kort overzicht gegeven van de verschillende delen uit de thesis.

Nonlinear methods in image restoration applied to confocal microscopy

English chapters

Chapter 1

General introduction

For the motivation why we included certain chapters and for some explanation about the convention we used for references in this thesis, we refer the reader to the part “Preface” of this thesis.

1.1 Image formation and optical limitations

The job of image restoration is to figure out what the instrument is actually trying to tell you.
—Prof. E.R. Pike, King’s College, London

Humans have a sense of discovery: they explore planets in space, climb mountains, dive to the bottom of the ocean, . . . In one sentence: today, they always want to go one step further than yesterday.

This urge to always go further is the driving stimulus of innovation and technology. Man has built particle accelerators to study the properties of atoms and subatomic particles; he has built telescopes to study stars, planets, galaxies, . . . He has built microscopes that allow him to study the other extreme: microscopic cells and the mysterious processes that form the essence of life itself.

These devices reveal a lot of hidden information. However, they also capture information not readily visible for the human observer, but still present in the image in a hidden form. Researchers always want to explore the limits of their equipment, and also desire to recover this hidden information and visualize it to facilitate interpretation. Taking into account the processes that take place in the equipment that gather the information, we can model which information is “lost” and which information is still present in a “hidden” form. Then, we can design solutions to recover this hidden information.

Here, we will focus on optical systems (such as microscopes) connected to a detector. The optical system has its limitations: it images an ideal mathematical point not as a point but as a smeared-out spot because of the wave nature of light. This image of an ideal point is the so-called Point Spread Function (PSF). This causes image blurring, even with a well aligned system: all optical systems have a finite resolution, even when optimally tuned; of course, with misalignment of the optical components, the situation degrades even further.

So the optical system produces a blurred image of the object, which is then recorded by some detector system, like a CCD. The reasons that noise occurs in recorded images can be classified in three categories:

- **Photon or shot noise:** due to the particle nature of light, image detectors are in fact particle counting devices. The photons that hit the detector are converted in charge carriers, such as electrons. However, this detection/conversion process is not deterministic because not every photon is actually detected in the detector.
- **Sensor noise:** is caused by a variety of noise processes in the detector, like dark current (charges that are generated in the absence of a signal), charges that are captured in a wrong detection unit, spatial sensitivity variations between the detector pixels, ...
- **Read-out or amplifier noise:** since the signal from the charge carriers has to be amplified and converted into an intensity number (through an Analog to Digital Converter (ADC)), additional noise occurs in the final image.

We illustrate the flow of information during the imaging process in Figure 1.1: an ideal object (in fact, the image under ideal circumstances) is shown at the top left. The ideal object is blurred due to the band-limited properties of the optical system. After the recording process, we finally obtain a degraded image, blurred and noisy.

Summarizing, the final recorded image has in fact been blurred and corrupted by noise, which results in the loss of actual information. However, it is not because some feature is invisible that this feature is lost [Van der Voort, WWW]:

- often details are hidden in the noise or masked by other features;
- artifacts may confuse the viewer;
- information may be present in implicit form, i.e., it can only be retrieved when one imposes prior knowledge (i.e., knowledge that one already can assume about an image before observation, e.g., the fact that all pixels must be positive).

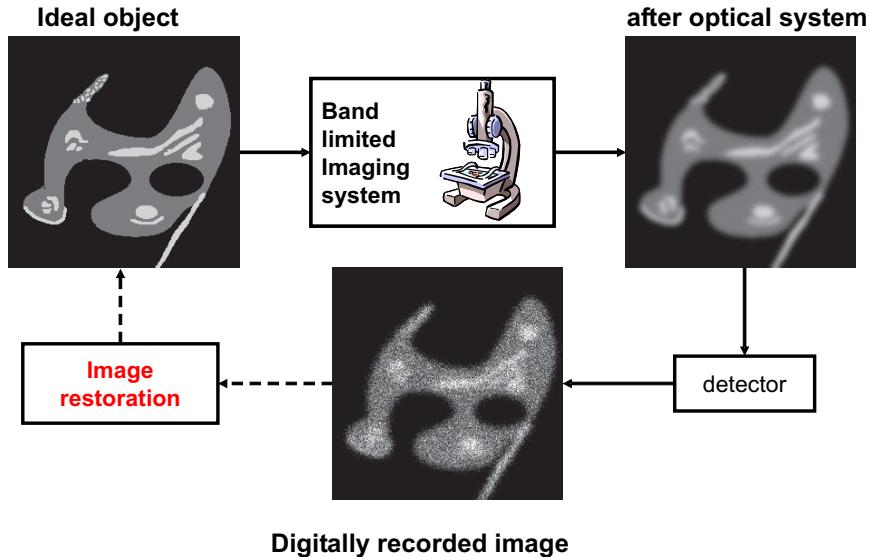


Figure 1.1: Block diagram of the flow of information during the imaging process and the corresponding degradation. Top left: a synthetic cell image; top right: the image, blurred due to band-limited optics. Bottom, the image as actually captured.

So, the goal of image restoration is to compensate as well as possible for the degradation process, and to recover an image that reveals some of the “hidden” information. This facilitates interpretation by a human observer, and post-processing, like segmentation, object recognition, ...

1.2 Applications of image restoration

This paragraph has been loosely based on [Molina et al., 2001]. Image restoration originated from the space program in the 1950’s, which provided us the first images by probes launched in space from the Earth, the Moon and Mars. These images had an unprecedented resolution, but were acquired under difficult circumstances, like vibrations, bad pointing, motion during the recording process, ... Because of the budgets invested in developing and launching these probes, the need to retrieve as much information as possible from these images was clear. Therefore, known one-dimensional signal restoration algorithms (like the Wiener filter) were extended to two dimensions. Applying these techniques to the space probe images resulted in a spectacular quality improvement.

Since their introduction, techniques of image restoration and reconstruction were commonly used in fields like medical tomography, seismology, magnetic resonance imaging, etc. In optical astronomy, it was considered to be an unnecessary luxury, until a defect was discovered in the optical system of

the Hubble Space Telescope (HST) shortly after it was launched in 1990. Two workshops were organized to consider the options for restoration of the HST images. The first workshop was quite general, the second workshop [Hanisch and White, 1993] was more focused on specific problems that arose for images captured with the HST.

Another historic use of image restoration techniques was the attempt to restore the pictures and movies taken during the assassination of John F. Kennedy. The images captured on the day of the event were blurred because the photographers were so shocked when they took them. Especially the restoration of the Zapruder movie could count on a lot of public interest.

A second scientific discipline where image restoration is common practice [van der Voort and Strasters, 1995, Van Kempen et al., 1997, Verveer et al., 1999], is fluorescence microscopy. An algorithm that is commonly used in this area is the Expectation-Maximization algorithm, first applied in microscopy by Holmes in 1988, and discussed in more detail in Section 3.2.2.1. As computers became more powerful, image restoration was also applied to three-dimensional images and even 4D (3D video sequences), like the ones acquired with confocal microscopes. This is a specific type of optical microscope that allows volumetric imaging (for more details, see Subsection 2.1.3).

This type of microscope is often used in research institutes for both fundamental as well as applied scientific research. In fundamental research, confocal microscopy is used for research in molecular biology; in chemistry, e.g., to study transport phenomena in polymer films, ... In applied research, it is used to study painting processes of textiles, wax layers of fruit skins to study storage life, ... In pharmaceutical companies, confocal microscopes are valuable tools to study the influence of medicines on cells. Other applications where confocal microscopes are used, are:

- strength studies of materials (e.g., to visualize cracks), studies of roughness of surfaces, ...;
- paper industry: studies of properties of wood fibers;
- textile: visualization of the structure of fibers for wear studies;
- electronics: inspection of micro-circuits.

In all of these applications, it is important to obtain the best image quality possible, and to extract as much information as possible from the observations.

In surveillance situations, image sequences are often of poor quality. Motion blur in particular introduces significant image degradation [Bascle et al., 1996]. Also there, restoration of images and image sequences is often desirable.

In fact, image restoration can be applied to every situation where a digital image is recorded by an optical system. Nowadays, improving the

quality of images from digital cameras is getting more and more popular, and plenty of software packages and Photoshop plugins are available for this purpose: [Uman Software, WWW, Nik Multimedia, Inc., WWW, UltraSharpen, WWW], just to name a few.

1.3 Scope of this thesis

In this thesis, we investigate a set of algorithms for restoration of confocal microscopy images. In particular, we investigate the application of multiresolution principles in the restoration process. Multiresolution methodologies have proven to be able to model images and other signals more accurately, making them powerful tools for modeling prior knowledge.

In Chapter 2, the image formation process of a confocal microscope is described, as well as the sources of degradation in the image formation and acquisition process.

Chapter 3 deals with the classical approaches for image restoration of confocal images from basic principles. The inverse filter and the Wiener filter are treated mainly to focus the attention on the basic problems involved in image restoration. Next, more advanced restoration techniques, like Tikhonov-Miller and Richardson-Lucy are explained, as well as how regularization introduces prior knowledge in the restoration process. These algorithms are discussed within the Bayesian framework. Since these topics were already discussed extensively, this Chapter is mainly based on the corresponding Chapters in [Verveer, 1998, van Kempen, 1999]

Chapter 4 introduces the principle of multiresolution, and presents some applications in image processing, like denoising and image segmentation. Specifically, we discuss how the principles of multiresolution can be applied in imposing prior knowledge to an image during restoration.

In Chapter 5, our two original methods for parametric PSF estimation are described, in case the characteristics of the blurring process are unknown, and have to be estimated. Parametric blur estimation starts from a simplified model for the PSF, assuming that the PSF can be expressed as a function depending on a single parameter, like a Gaussian function (with zero mean, and unknown variance) and an out-of-focus PSF (depending only on the radius of the focal spot). These estimation techniques are both based on properties of multiresolution.

Chapter 6 explains SPERRIL, which is our original algorithm for the estimation and restoration of confocal images. This Chapter includes results of experiments on synthetic images as well as results obtained with biological images.

Chapter 7 discusses our original multiresolution-GNC based algorithm for image restoration. This method is based on a multiresolution Markov Random Field (MRF), of which the energy function is minimized by the Graduated Non-Convexity (GNC) method.

Chapter 8 summarizes the thesis and its results. Also, a few suggestions for further research are given.

1.4 Summary of contributions

This work has lead to the following original contributions.

- A new parametric wavelet-based blur estimation method (called Wave-Track) has been developed;
- Wave-Track has been incorporated in our new restoration technique (SPERRIL) and evaluated with synthetical and biological images. For the colocalization analysis, we have worked together with dr. D.S. Lidke from the Max Plank Institute in Göttingen. For other confocal images, we have worked together with prof. dr. P. Van Oostveldt from Ghent University;
- Together with prof. dr. ir. J. Portilla, a second parametric wavelet-based blur estimation method (called Sparse-Max) has been developed;
- Together with dr. ir. A. Pižurica, a new restoration technique based on Markov Random Fields and wavelets was developed. This technique however is still under development, but the first results are promising.

During the work, four student graduation theses about image restoration were guided:

- Tim Stevens (2001): “*Non-blind deconvolution of three-dimensional images from a confocal microscope*”;
- Sofie De Wolf (2002): “*Quantitative comparison of blind deconvolution algorithms*”;
- Anthony Dewulf (2001-02): “*Wavelet-based regularization of deconvolution algorithms*”;
- Hans Yperman (2004-05): “*Restoration of images with spatially variant blur*”.

1.5 Summary of publications

So far, this research has resulted in a popularizing article in a magazine [Rooms and Philips, 2002a] and a journal paper [Rooms et al., 2005].

Also, eight peer-reviewed conference papers as a first author have been published: [Rooms et al., 2004, Rooms et al., 2003b, Rooms et al., 2003c,

Rooms and Philips, 2002b, Rooms et al., 2002c, Rooms et al., 2002b,
Rooms et al., 2001a, Rooms et al., 2001b, Rooms et al., 2000b]

Nine abstracts (papers \leq 2 pages) were published for presentations at international meetings:

[Rooms and Philips, 2004, Rooms et al., 2003e, Rooms et al., 2003a],

[Rooms et al., 2003d, Rooms et al., 2002a, Rooms et al., 2001c],

[Rooms and Philips, 2001, Rooms and Philips, 2000, Rooms et al., 2000a].

Related research has contributed to work of others, which resulted in the following papers and abstracts in national and international conferences [De Smet et al., 2004, Stippel et al., 2001, Stippel et al., 2000].

Chapter 2

Image formation

*All science begins with astonishment,
but the human behind the instrument is much
more important than the instrument itself.*

2.1 A brief history: from single lens to confocal microscope

To put this thesis about microscopical image processing in a historical context, we include an introduction on the history of the microscope. In this Chapter, the sections on the early days of the microscope have been loosely inspired on [Bellis, WWW], while the sections on confocal microscopy have been inspired on [van Kempen, 1999, Verveer et al., 1999, van der Voort, 1989, Plášek and Reischig, 1998, Lichtman, 1994].

2.1.1 From lens to microscope

Long ago, someone picked up a convex piece of transparent crystal, and noticed that objects appeared larger when observed through this crystal: the first lens was born. They were named lenses because they are shaped like the seeds of a lentil (*Lens culinaris* or *Lens esculenta*, see Figure 2.1). Even today, the words for the plant and the optical element in French, Dutch and English are related, as can be seen in the table below.



Figure 2.1: Seeds of a lentil [University Oklahoma, WWW]. Even today, the link between the word “lens” and the name of this plant still exists.



Figure 2.2: The Janssen microscope-tube (1595), which could magnify up to 20X [Moll, W.A.W., WWW].

	optics	plant
Dutch	lens	linze
English	lens	lentil
French	lentille	lentille
Spanish	lente	lenteja

Someone found that such a crystal could focus sunlight and heat into a single point, thus igniting wood, paper, ... by pure solar heat. These “burning glasses” or “magnifying glasses” are mentioned in the writings of Seneca and Pliny the Elder (Roman philosophers during the first century A.D.). However, they were not commonly used until the invention of glasses at the end of the 13th century.

Around 1590, two Dutch opticians, Zacharias Janssen and his father Hans were experimenting with sets of lenses in a tube, and discovered that nearby objects were observed greatly magnified through this device, even more than with single lenses. This device could be considered to be the prototype of the compound microscope and the telescope (see Figure 2.2).

Antoni van Leeuwenhoek (1632 - 1723) [Wikipedia, WWW, Moll, W.A.W., WWW] also lived in the Netherlands, and made some more significant improvements to the design of compound microscopes. He learned his skills in optics in a textile shop where people used magnifying glasses to count the threads in clothes. He developed new methods for grinding and polishing small lenses with great curvature. These new lenses magnified objects up to 270x, and were the finest lenses of their time. With these lenses, van Leeuwenhoek constructed his microscopes, with which he made many biological discoveries: he was the first to see and describe bacteria, yeast and the circulation of blood corpuscles in capillaries. He is considered as one of the pioneers of what would be known later as cellular biology and microbiology.

Robert Hooke lived between 1635 and 1703, and is considered to be the English father of microscopy. He verified and confirmed Antoni van Leeuwenhoek's discoveries, like the existence of extremely small organisms living in water. Hooke made his own microscopes, starting from the design of van Leeuwenhoek, and then made further improvements of his own.

The first person to explain the laws of geometric optics from the principles of wave theory, was Christiaan Huygens in 1670. Huygens formulated a very important insight about wave propagation, which is nowadays called Huygens' principle [Fitzpatrick, WWW]: *Every point on a wave-front may be considered a source of secondary spherical wavelets which spread out in the forward direction at the speed of light. The new wave-front is the tangential surface to all of these secondary wavelets.* Based on this principle, Huygens was able to explain phenomena like refraction and reflection of light.

Ernst Abbe published his work on the theory of the microscope in 1873. He was the first to lay a theoretical foundation for the design of microscopes, which had been mostly a trial and error process up to this point. He also formulated clearly the difference between magnification and resolution, and described that microscopes with a given objective lens had a maximal “useful” magnification; he called any larger magnification “empty magnification.” Based on the wave theory of light, he derived his widely used formula to calculate microscopic resolution.

2.1.2 From classical microscope to fluorescence microscope

The classical optical microscope has a few shortcomings. One specific shortcoming is that the contrast in specimens is usually quite low. This can be solved by coloring the specimen with dyes that attach to specific parts of the

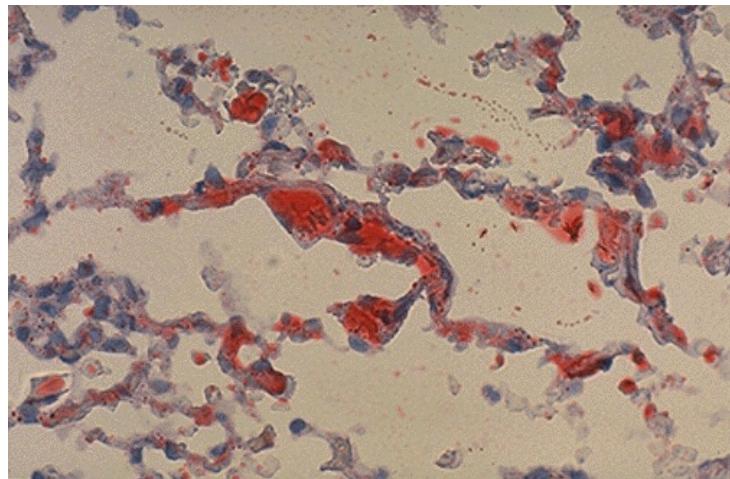


Figure 2.3: *Oil red O* stain (which highlights fat in tissues) of fat emboli (=free floating elements that are foreign to the blood stream) in lung [Florida State University College of Medicine, WWW].

specimen, which improves contrast with other parts of the specimen (see Figure 2.3 for an example where clumps of fat in lung tissue are highlighted by staining them red).

With a classical light source (broad spectrum) and classical dyes, the observer receives light scattered from the region of interest as well as light scattered from surrounding regions, which reduces the contrast. To enhance the contrast, fluorescent dyes¹ are used: fluorochromes are excited at one specific wavelength λ_{ex} and emit at another wavelength λ_{em} . This wavelength λ_{em} is always longer than λ_{ex} because of energy dissipation during the excited-state lifetime (the difference between these two wavelengths is known as the *Stokes shift*). Figure 2.4 shows excitation/emission spectra of two typical fluorochromes. Because of the Stokes shift, it is possible to separate scattered illumination light from the fluorescence light from the specimen.

In Figure 2.5, a schematic diagram of a widefield fluorescence microscope is shown. The illumination light from the light source to the specimen is shown in light blue: notice that the specimen is illuminated uniformly with excitation wavelength λ_{ex} , which then emits light of wavelength λ_{em} . The black lines show the path of light from a point in the focal plane. However, light emitted by dyed parts of the specimen above and below the focal plane also contributes to the image of a single slice (e.g., the red lines in Figure 2.5), causing the so-called out-of-focus blur.

In summary, the widefield microscope illuminates the whole specimen and im-

¹The terms fluorophore (=fluorescent component) and fluorochrome (=fluorescent dye) are both used.

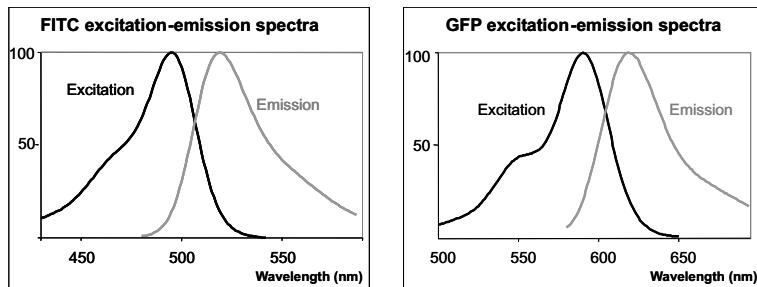


Figure 2.4: Excitation and emission spectra of two fluorochromes: FITC and GFP (Source: [Molecular Probes, WWW]). Note that in both cases the peaks of excitation and emission are close to each other, but that they can be almost completely separated (i.e., a filter can be chosen to eliminate excitation light, and to allow only emission light).

ages it as a whole, just as the common microscopes in biology classes in high school. The only difference with those microscopes is the fluorescence process involved. A filter placed before the camera that absorbs λ_{ex} , but that allows λ_{em} to pass, prevents scattered illumination light from contributing to the image of the specimen, and thus enhances the contrast.

2.1.3 From fluorescence microscope to confocal microscope

Marvin Minsky is famous as the father of artificial intelligence, but he is also known for another important achievement: the invention of the confocal microscope in 1957 [Lichtman, 1994, van der Voort, 1989]. This revolutionary type of microscope allows imaging of the deeper layers in a specimen without the hard labor of cutting the specimen in thin slices. In fact, he made just a few simple but very important modifications to the widefield microscope discussed above.

First, he reduced the scattering of the illumination light in the specimen by sending this light through the objective lens, thus reducing the illumination in principle to one point. In fact, the illuminated area is a double cone, and the total intensity of a section of the cone, orthogonal to the axis, is constant. This is illustrated at the top of Figure 2.6. At the bottom, the intensity profile along the optical axis for three different values of α is shown. Here, α is the acceptance angle of the objective, which is the maximum angle of the rays originating from the focal point to be intercepted by the objective lens; alternatively, it is the angle of the cone formed when a parallel light beam along the optical axis hits the objective (top of Figure 2.6). Typical values of α for confocal objectives are between $\pi/3$ and $\pi/4$. The larger α is, the more signal intensity is gathered from the focal point and the thinner the slices one can image.

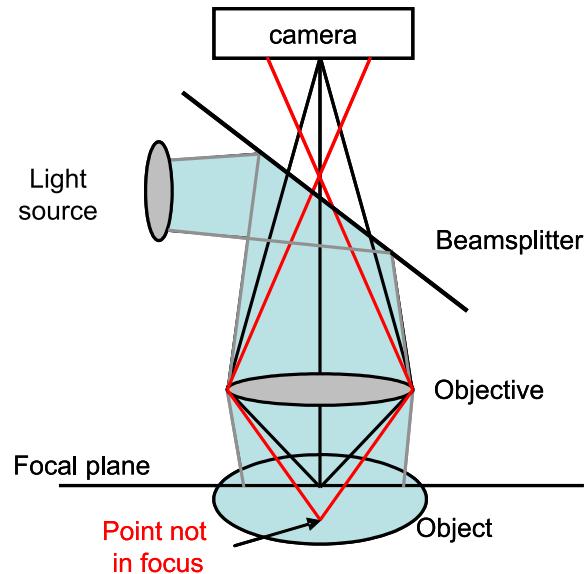


Figure 2.5: Block diagram of the image acquisition setup in a widefield fluorescence microscope. From this scheme, it is clear that this type of microscope cannot distinguish light from planes outside the focal plane from light coming from the focal plane because at a single instance, light from the whole specimen reaches the camera.

The second modification Minsky made, was adding a diaphragm in front of the detector (as shown in Figure 2.7). The purpose of this diaphragm is to prevent light from points in the specimen above and under the focal plane from contributing to the image. This is also illustrated in Figure 2.7, where the black lines represent light rays from the focal point, which can pass the diaphragm, while the red lines show light rays from a certain point below the focal plane, which are almost completely stopped by the diaphragm.

The result of these two modifications was that only a single point in the image is illuminated, and that only light from this point can reach the detector. This results in a sharp image of a single spot in the specimen. With a 3D scanning mechanism added to this optical system, it is possible to acquire 3D images of microscopical specimens.

An example is shown in Figure 2.8. There, a 3D image stack is shown of a so-called solid state nuclear track detector. This kind of detector is in fact a piece of plastic. The detector is then exposed to heavy particle radiation, like α particles (which are helium nuclei) or even heavier particles. Thus, the polymer structure is damaged where a particle crossed the plastic. After etching the plastic with a solvent, this damage is expanded and forms cone-shaped pits in the plastic. These pits are then coated with a fluorescent dye, and imaged with a confocal microscope. An image of such coated cones is shown in Figure 2.8. In Figure 2.8 (a), a 3D rendering of this object is shown. In Figure 2.8 (b), the

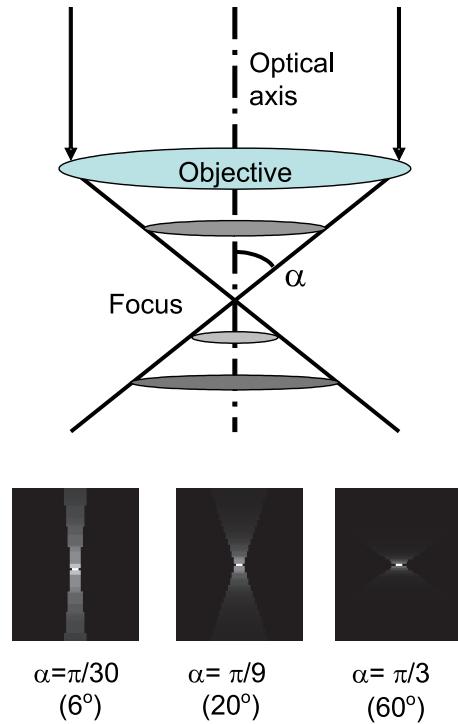


Figure 2.6: Illumination cone for a confocal microscope. The total intensity within the different ellipses is constant. Bottom: intensity profile for three different values of the acceptance angle α of the objective. The gray ellipses are cross-sections of the illumination cone.

confocal image stack is shown.

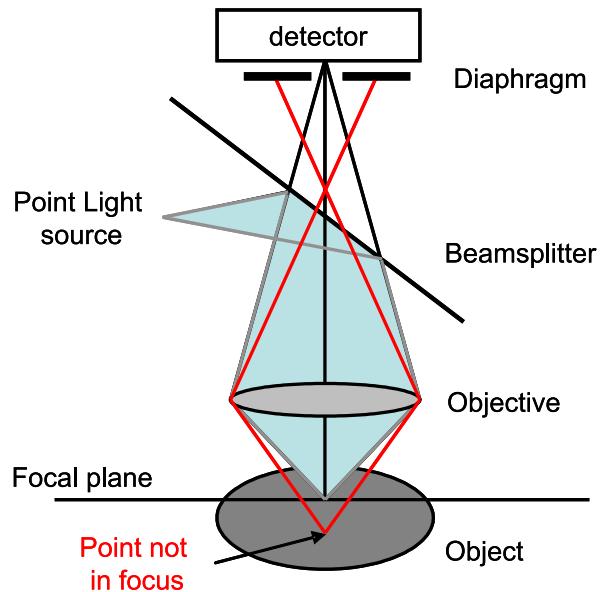


Figure 2.7: Block diagram of image acquisition setup in a confocal microscope. Only light from the plane of focus can pass to the detector without obstacles. Light from around the focal plane is almost completely eliminated by the illumination pinhole in front of the detector.

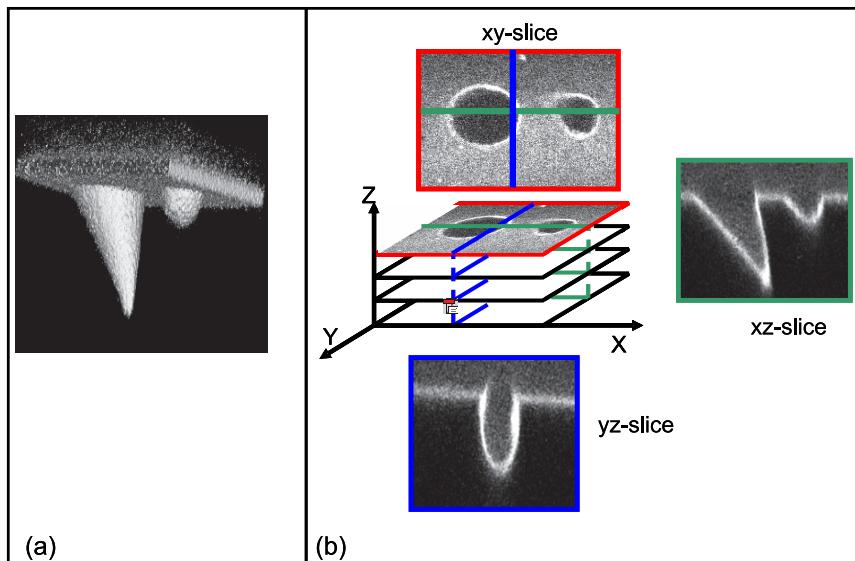


Figure 2.8: Example of a 3D confocal image: (a) 3D rendering from the confocal image; (b) confocal image stack, with three orthogonal sections through the stack.

2.2 An intuitive approach to the limitations of optical image formation 19

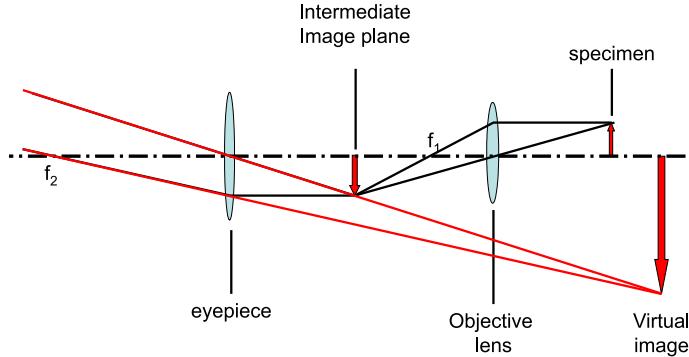


Figure 2.9: Geometric image formation in a microscope [Alonso and Finn, 1989], p190.

2.2 An intuitive approach to the limitations of optical image formation

2.2.1 The basics

In [Plášek and Reischig, 1998], a nice intuitive explanation is given on resolution of optical systems from the physicist's point of view. Based on this article, we will briefly summarize the basics of image formation in microscopy, and how these basics apply to the resolution of the image.

The basic components of the microscope are the *objective* and the *eyepiece*. The role of the objective is to collect the light from the specimen, and to focus it into a magnified image. This image is then further magnified by the eyepiece in order to adapt it to the resolving power of the human eye or the detector of the imaging system (i.e, the CCD, PMT, ...). See also Figure 2.9.

When one considers a hypothetical light source in the specimen shaped as a mathematical point, this will never be imaged as a mathematical point in the image, but as a spot of light, the so-called Point Spread Function (PSF) due to diffraction because of the wave nature of the light. The PSF of an optical lens is given by the *Airy disk*, of which the radial intensity profile is given by:

$$I_{\text{Airy}}(r) \sim \left[\frac{J_0(r)}{r} \right]^2,$$

where $J_0(x)$ is a Bessel function of the first kind of order 0, and r is the distance to the center of the pattern (Figure 2.10 (a)).

For a point source emitting incoherent light, the lateral resolution (in a plane orthogonal to the optical axis, see Figure 2.11) is given by the radius of the

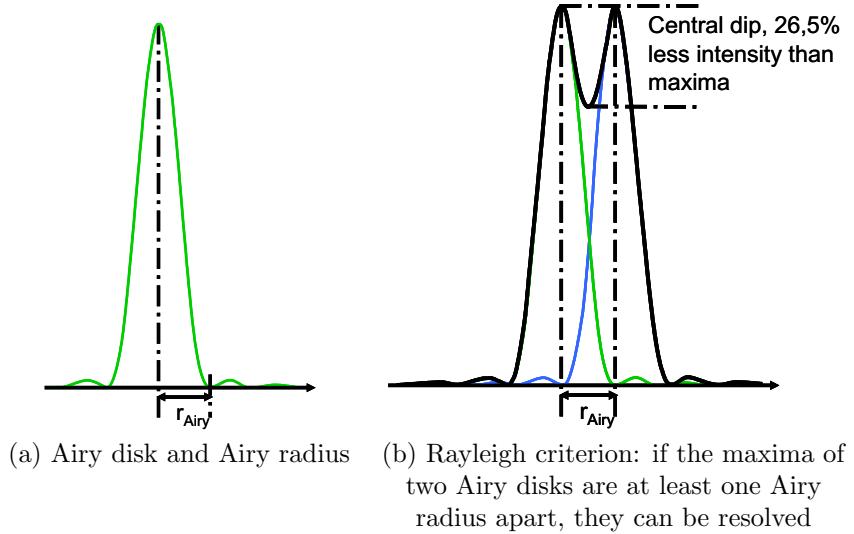


Figure 2.10: Airy disk, Airy radius and Rayleigh criterion.

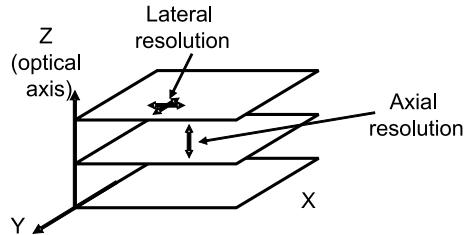


Figure 2.11: Axial versus lateral resolution.

Airy disk (r_{Airy}), which is given by:

$$r_{\text{lateral, widefield}} = r_{\text{Airy}} = 0.61 \frac{\lambda_0}{NA_{\text{objective}}}, \quad (2.1)$$

where λ_0 is the wavelength of the light that is used, and $NA_{\text{objective}}$ is the Numerical Aperture of the objective:

$$NA_{\text{objective}} = n \sin(\alpha),$$

where n is the relative refraction index for the transition between the specimen and the objective lens and α the acceptance angle of the objective lens (see

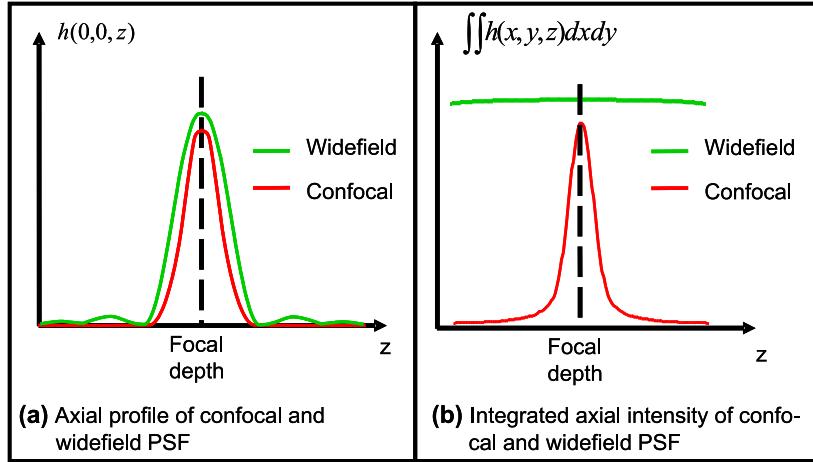


Figure 2.12: Axial intensity profile and integrated intensity per layer of depth for the PSF of a confocal microscope versus a widefield microscope.

again Figure 2.7). The Rayleigh criterion says that two point sources can be distinguished if the distance between the maxima of the two Airy disks is at least r_{Airy} (Figure 2.10 (b), see also [Pawley, 1995] p.1). At this distance, the dip between the two Airy disks is about 25% smaller than the maxima of the total diffraction pattern, but this is only true in classical imaging. In confocal imaging however, the illumination scans the specimen point-wise and the emission light is also point-wise detected, so that only that part of the specimen in the common volume of the illumination and detection point spread functions is detected. The result is that the lateral resolution is further improved:

$$r_{\text{lateral, confocal}} = 0.4 \frac{\lambda_0}{NA_{\text{objective}}}, \quad (2.2)$$

which yields an improvement by a factor 1.4 [Pawley, 1995, Plášek and Reischig, 1998]. As for the axial resolution (along the optical axis, see again Figure 2.11), the confocal microscope only obtains a small improvement in measured axial resolution over that of the widefield microscope (as illustrated in Figure 2.12 (a)):

$$r_{\text{axial}} = 1.4 \frac{\lambda_0 n}{(NA_{\text{objective}})^2} \quad (2.3)$$

The true advantage of the confocal microscope is that it allows optical sec-

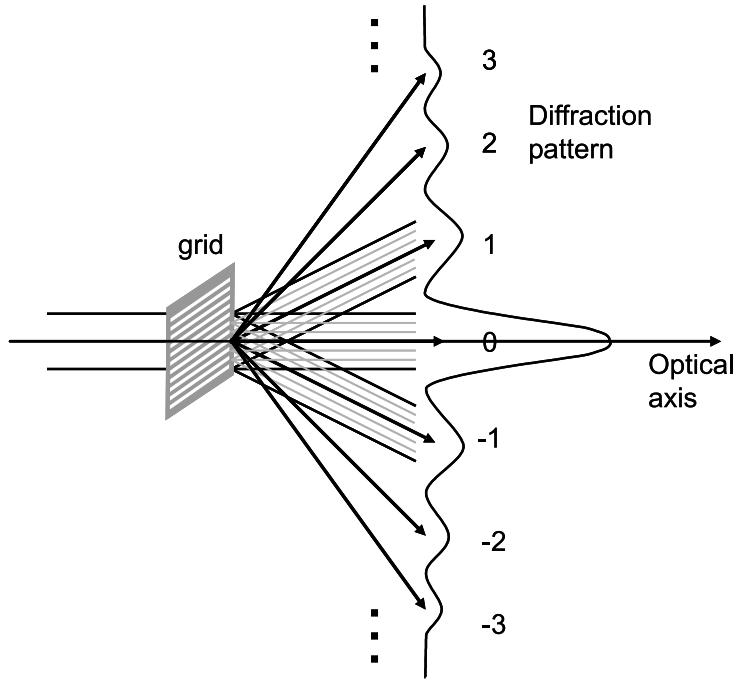


Figure 2.13: Diffraction pattern for a grid of apertures. The numbers indicate the orders of the diffraction maxima. Note that the higher the order, the higher the angle is between the direction under which the maximum is formed, and the optical axis. Only the rays corresponding with diffraction orders -1 , 0 and 1 are shown. For higher orders, only the directions in which the maxima are formed are shown.

tioning of thick specimens, which results in a huge improvement in effective axial resolution over that of the widefield microscope. The ability for optical sectioning of the confocal microscope results from the properties of the PSF, integrated over each xy-slice:

$$h_{\text{integrated}}(z) = \int \int h(x, y, z) dx dy,$$

so $h_{\text{integrated}}(z)$ is only dependent of the depth z . We can see in Figure 2.12 (b) that $h_{\text{integrated}}(z)$ is maximal in the focal plane, so only light from depths around the focal plane contributes to the image. The corresponding integrated intensity per slice for the PSF $h_{\text{integrated}}(z)$ of the widefield microscope is constant as a function of the depth z , so this type of microscope cannot perform optical sectioning since it collects light from all depths of the specimen [Olympus Fluo View, WWW].

A remark to be made here² is that Rayleigh formulated his criterion for line

²thanks to dr. H Van der Voort to make this comment.

patterns in a spectrometer: the resolution is defined as the two-line distance at which the maximum of the image of the first line coincides with the minimum of the second. Later, this criterion was extended to two-point resolution, were the diffraction patterns of both points are approximated by Airy functions (see again Figure 2.10). In the case of a widefield microscope, the Rayleigh resolution is in fact closer to the length of the highest spatial frequency in the pass band of the Optical Transfer Function than the factor of 1.4 which according to Pawley can be attained by contrast enhancement of a faint dip between the two objects.

In short, applied to widefield systems Pawley's statement is doubtful. Confocal systems have a twice as large bandwidth as a comparable widefield microscope, while their PSF is at most 1.4 times narrower, so in these cases the disparity between Rayleigh resolution and what can be obtained by contrast enhancement and for that matter image restoration is much larger. To a degree, this also explains the effectiveness of image restoration in confocal microscopy as shown in Chapters 3 and 6.

2.2.2 Ernst Abbe: the bridge between classical optics and Fourier optics

Abbe studied the theory of image formation starting from a regular grid (which is in fact a series of thin rectangular apertures) as specimen (see Figure 2.14 for the experimental setup and Figure 2.15 for the formed diffraction pattern). This grid is then illuminated with a narrow light cone, so that the incident light rays can be considered approximately to be parallel, for the sake of simplicity. Because of this grid of apertures, a diffraction pattern is formed behind the grid with distinct diffraction maxima (see Figure 2.13).

This diffraction pattern is best resolved in the focal plane (see Figure 2.15). The *zeroth* order diffraction maximum has maximum intensity, and is formed in the direction of the parallel beam, while diffraction maxima of increasing orders are formed at angular directions further and further away from the direction of the parallel beam. With a hypothetical perfect lens, the whole diffraction pattern would be imaged in the focal plane of the lens. However, due to the size limitations of the lens, only the central part of the diffraction pattern is imaged in the focal plane.

Abbe concluded that in order to image the true periodicity of the grid, at least the zeroth and the first order diffraction maxima $(-1, 1)$ must contribute to the image. The more diffraction maxima are intercepted by the lens and focused in the focal plane, the more accurate the formed image in the image plane is. However, when the angle between the parallel beam and a certain order diffraction maximum is bigger than the acceptance angle of the objective, this diffraction maximum is not intercepted by the objective and won't contribute to the formation of the image. The number of diffraction orders that contribute

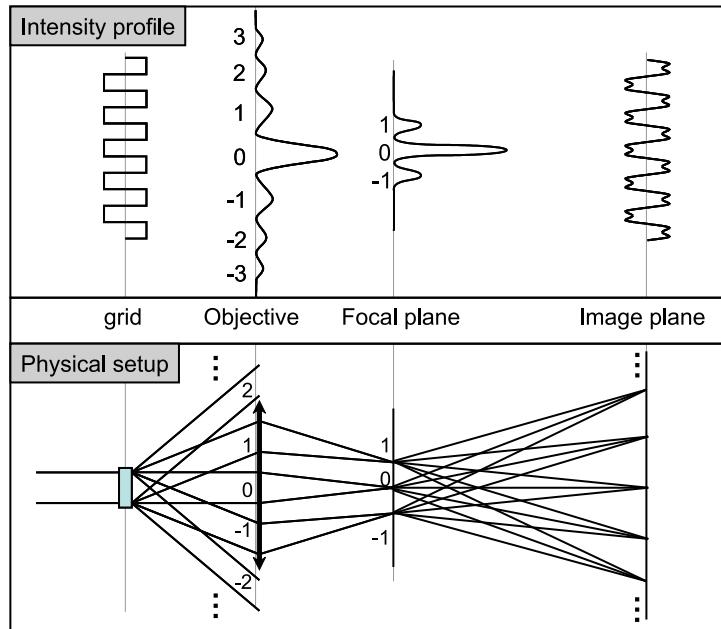


Figure 2.14: Setup of Abbe experiment: an object is illuminated by coherent light, which causes a diffraction pattern. This is imaged by an objective lens. Only those diffraction maxima that are intercepted by the objective can contribute to the formation of the image [Clauws, 1998].

to the formation of the final image depends on the aperture (i.e., the diameter) of the objective lens.

In Figure 2.15, the relation between the original grid, the diffraction maxima in the focal plane and the image of the grid with different numbers of contributing diffraction maxima is illustrated; also, the intensity profile of the final grid image with a certain number of diffraction maxima is shown. It is clear that when only the zeroth and the first order maxima contribute to the image, the periodicity of the grid is already visible, but the edges are still very unsharp. The more diffraction maxima contribute to the image, the more accurate the image formed in the image plane is. However, as explained before, the acceptance angle of the objective lens limits the number of diffraction maxima that can contribute to the image formation, and therefore also limits the resolution of the formed image.

From his theory, Abbe derived a formula to calculate the resolving power of a microscope, which is nearly identical to the Rayleigh limit; see equation (2.1).

Abbe's theory was an important step in our modern understanding of image formation: Fourier optics. The propagation of light is approximated by a scalar phenomenon (so its electromagnetic nature is neglected).

2.2 An intuitive approach to the limitations of optical image formation 25

If we consider monochromatic light propagating through vacuum, the variations in the electromagnetic field as a function of time and space can be separated, and the six field components (three for the electrical field and three for the magnetic field vectors) have identical variations in space. So the variation of the electrical and magnetical vector fields can therefore be written as the product of a scalar field, $U(\vec{r})$ (which captures the spatial variations), an exponential factor (which captures the variations in time) and the vectors $\vec{E}_{\omega 0}$ and $\vec{B}_{\omega 0}$ at time $t = 0$ [MZA Associates Corporation, 2004]:

$$\vec{E}_{\omega}(\vec{r}, t) = U(\vec{r}) \exp(j\omega t) \vec{E}_{\omega 0} \quad (2.4)$$

$$\vec{B}_{\omega}(\vec{r}, t) = U(\vec{r}) \exp(j\omega t) \vec{B}_{\omega 0} \quad (2.5)$$

with

$$U(\vec{r}) = A(\vec{r}) \exp(j\phi(\vec{r})) \quad (2.6)$$

$$\omega = c/\lambda \quad (2.7)$$

where U is the optical field, A is the amplitude, c is the light speed, λ is the wavelength of the light and ϕ is the phase. This scalar approach is only valid if [Noponen, 1994]:

- the diffracting aperture is large compared with the wavelength of light;
- the observation distance is large enough;
- the considered angles are small enough to be able to ignore the axial field components.

Scalar diffraction theory has the following limitations:

- it assumes monochromatic light;
- it does not include polarization effects;
- it is no longer valid when considering objects with sizes comparable with the light wavelength;
- it is no longer valid when considering very high and very low intensities.

A widely used simplified model in scalar optics is the so-called Fresnel diffraction formula, which makes the following additional assumptions (as illustrated in Figure 2.16):

- both the distance (light source - aperture) and the distance (observation point - aperture) are much larger than the size of the aperture;

- if the optical axis goes through the aperture and the aperture plane is orthogonal to the optical axis, then both the orthogonal distance (light source - optical axis) as well as the orthogonal distance (point of observation - optical axis) are of the same order as the size of the aperture.

The Fraunhofer diffraction formula is a further simplification of the Fresnel formula; it assumes on top of all the previous assumptions that the distance between observer and aperture is nearly infinite; this distance can be shortened by placing a converging lens between the observer and the aperture.

The rest of this Section is based on [Clauws, 1998]. In essence, the Fresnel-Kirchhoff integral is the quantitative form of the Huygens principle (which we mentioned before in Subsubsection 2.1.1). This integral can be simplified for most practical cases to what is known as the Fresnel diffraction formula. In case of optical imaging, often an even more simplified version of this formula can be used, i.e., the Fraunhofer diffraction theory. In this theory, light is described as a wave with a complex amplitude, as in equation (2.6).

The Fraunhofer diffraction formula in essence gives a general description for the diffraction pattern of wavefronts on a given aperture, and states that under the conditions for Fraunhofer diffraction, the diffraction pattern is given by the Fourier transform of the wavefront in the aperture, multiplied by a complex phase factor. In case of plane waves (plane waves have planes of points with identical phase), the diffraction pattern is just the Fourier transform of the aperture function (which describes where light can pass by 1 and where no light can pass by 0), multiplied by a complex phase factor. If we are only interested in the intensity of the diffraction pattern, this complex phase factor can be ignored.

We consider the diffraction of plane waves at an aperture $T(y, z)$, like in Figure 2.17. The diffraction pattern that is normally formed at infinity (lines that are (almost) parallel converge at infinity), is now imaged at the focal point F of lens L. Because of the presence of the aperture, the wave fronts behind the aperture are no longer planes and therefore, the resulting wavefront in F is no longer a point, but a Fraunhofer diffraction pattern $t(u, v) \sim FT(T(y, z))$, with $FT()$ the Fourier transform.

For instance, a rectangular aperture function can be considered as the convolution of a 1D box function in the y-direction with a 1D box function in the z-direction. The Fraunhofer diffraction pattern formed in a situation like in Figure 2.17 with incident plane waves will then be the product of a sinc-function in the y-direction and a sinc-function in the z-direction: the sinc-function is the Fourier transform of a box-function, and convolution in the spatial domain becomes multiplication in the Fourier domain. The intensity is given by the squared amplitude of this Fourier transform (see Figure 2.18).

The calculation of the time-dependent complex amplitude $U_{focus}(y, z)$ originating from an object illuminated with monochromatic light of wavelength λ

is performed in two steps (we assume that the object is illuminated by plane waves, which are diffracted by the object):

1. We need to calculate this diffracted complex scalar waveform $U_l(y, z)$ just before the lens. This is calculated with the Fresnel diffraction formula. In essence, this results in the Fourier transform of the intensity of the object, multiplied with a phase factor, which holds the 3D spatial information about the object (in holography, this phase is used).
2. We need to compute the complex scalar waveform $U_{focus}(y, z)$ from $U_l(y, z)$, which is again done using the Fresnel diffraction formula. In essence, the result is that the diffraction pattern of the image is clipped by the aperture of the lens, and is not formed at infinity, but in the focal plane of the lens.

When we come back to Figure 2.15 and consider the image formation of a regular transmission grid illuminated by plane waves, the resulting Fraunhofer diffraction pattern consists of a series of distinct maxima. The higher the order of the maximum, the higher the angle between the direction in which the order is formed, and the optical axis. When a converging lens is used, these maxima (up to a certain order, depending of the acceptance angle of the lens) are focused in the focal plane of the lens. From the previous discussion, we know that this diffraction pattern corresponds with the Fourier decomposition of the grid, and that the highest spatial frequencies are lost, because those are not intercepted by the lens, and thus do not contribute to the final image.

2.3 Fluorescence microscopy: A particular case of imaging

The image formation in a general fluorescence microscope is explained in [van Kempen, 1999]. We summarize the derivation of the image formation process based on Van Kempen's thesis for the sake of completeness, but refer to [van Kempen, 1999] for more details.

The whole derivation is divided in the following parts:

- discussion of the threedimensional PSF of an objective lens, the basic element of a microscope³;
- setup of a general fluorescence microscope, and the pathway of information;
- specification of the general model towards a confocal microscope;
- specification of the general model towards a widefield microscope.

³In this thesis, we assume that the PSF is spatially invariant, which is a common approximation in image restoration literature to reduce the complexity of the models.

2.3.1 Threedimensional PSF of an objective lens

This paragraph is based on [Booth and Wilson, 2001, van Kempen, 1999]. An expression for the 3D PSF of a lens can be obtained by calculating the diffraction pattern of an incident converging spherical wave on an aperture. In absence of this aperture, the light would be focused into one point. When the aperture function is rotationally invariant, the light intensity near the focal region of the objective lens (or, the 3D PSF of the lens) is given by:

$$I(v, u, d) = \left| \int_0^1 P(\rho, d) \exp(ju\rho^2/d) J_0(\rho v) \rho d\rho \right|^2 \quad (2.8)$$

with $J_0(\cdot)$ the zeroth order Bessel function of the first kind, $P(\cdot)$ the aperture function of the lens and ρ the radial coordinate in the aperture plane, normalized so that the objective aperture has unity radius. The distance d is the nominal focus depth. The normalized coordinates u and v are given by:

$$v = \frac{2\pi nr}{\lambda} \sin(\alpha) \quad (2.9)$$

$$u = \frac{8\pi nz}{\lambda} \sin^2(\alpha/2) \quad (2.10)$$

where r and z are the lateral and axial displacements from the nominal focus (see Figure 2.19), λ is the wavelength of the illumination light, n is the relative refractive index of the medium and α is the semi acceptance angle of the objective lens, which is related to the numerical aperture (NA) of the lens $NA = n \sin(\alpha)$.

2.3.2 General fluorescence microscope

This Subsection is based on [van Kempen, 1999]. An abstract diagram of a general fluorescence microscope is given in Figure 2.20. The light source has an illumination aperture $A_{ill}(\mathbf{x})$, and can be modeled as a set of point sources spread over the aperture, which are all projected by the objective on the sample. The illumination light passes the objective lens and is convolved with $h_{ex}(\mathbf{x})$; so $I_{ill}(\mathbf{x})$, the intensity illuminating the object of interest is given by:

$$I_{ill}(\mathbf{x}) = (A_{ill} * h_{ex})(\mathbf{x}) = \int h_{ex}(\mathbf{x} - \xi) A_{ill}(\xi) d\xi$$

with $h_{ex}(\mathbf{x})$ the PSF of the objective lens for excitation wavelength λ_{ex} , i.e., equation (2.8) for λ_{ex} . The illumination light excites the fluorochromes of the object (described by a fluorochrome density distribution $O(\mathbf{x})$, which represents the real physical object that is imaged), which emit fluorescence light $I_{em}(\mathbf{x})$ at wavelength λ_{em} :

$$I_{em}(\mathbf{x}) = O(\mathbf{x}) I_{ill}(\mathbf{x}) \quad (2.11)$$

$$= O(\mathbf{x}) (A_{ill} * h_{ex})(\mathbf{x}) \quad (2.12)$$

This light distribution from the specimen is imaged through the objective lens, and is thus spatially convolved with the PSF of the objective lens $h_{em}(\mathbf{x})$ for emission wavelength λ_{em} , and multiplied with the aperture of the detector $A_{det}(\mathbf{x})$:

$$I_{det}(\mathbf{x}) = A_{det}(\mathbf{x}) (I_{em} * h_{em})(\mathbf{x}) \quad (2.13)$$

The PSF for a General Fluorescence microscope (GFM) is obtained by replacing the object $O(\mathbf{x})$ in equation (2.12) by a point (i.e., a delta function), so it becomes:

$$\begin{aligned} h_{GFM}(\mathbf{x}) &= A_{det}(\mathbf{x}) \left(h_{em}(\mathbf{x}) * \left[\delta(\mathbf{x}) \int h_{ex}(\mathbf{x} - \xi) A_{ill}(\xi) d\xi \right] \right) \\ &= A_{det}(\mathbf{x}) \left(h_{em}(\mathbf{x}) * \left[\delta(\mathbf{x}) \int h_{ex}(-\xi) A_{ill}(\xi) d\xi \right] \right). \end{aligned}$$

This means that the point-object that receives the illumination light, acts as a point source of light itself, i.e., a delta function scaled by the intensity of the incident illumination light (expressed by the integral). Therefore, $h_{GFM}(\mathbf{x})$ becomes:

$$h_{GFM}(\mathbf{x}) = A_{det}(\mathbf{x}) h_{em}(\mathbf{x}) \int h_{ex}(-\xi) A_{ill}(\xi) d\xi. \quad (2.14)$$

2.3.3 Confocal microscope

In an ideal confocal microscope, both the illumination aperture and the detector aperture are points, i.e., delta functions⁴. So the ideal confocal PSF can be

⁴When a real point-like object (like a fluorescent bead) is imaged with a confocal microscope, then making $A_{det}(\mathbf{x})$ a delta function in Equation 2.14 would correspond with imaging

described as:

$$h_{\text{confocal,ideal}}(\mathbf{x}) = h_{em}(\mathbf{x}) h_{ex}(-\mathbf{x}). \quad (2.15)$$

However, a confocal diaphragm is not an ideal point in practice. In that case, it can be derived that the confocal PSF is:

$$h_{\text{confocal,real}}(\mathbf{x}) = ((h_{em} * A_{det})(\mathbf{x})) h_{ex}(-\mathbf{x}), \quad (2.16)$$

where the convolution $(h_{em} * A_{det})(\mathbf{x})$ is a 2D convolution in the plane of the aperture (i.e., the convolution is only carried out over the variables x and y).

2.3.4 Widefield microscope

In a widefield microscope, the illumination aperture is “infinite”, which means that $((A_{ill} * h_{ex})(\mathbf{x}))$ becomes a constant. So the widefield PSF becomes:

$$h_{\text{widefield}}(\mathbf{x}) = h_{em}(\mathbf{x}). \quad (2.17)$$

Since the excitation and emission wavelengths are of the same order of magnitude (e.g., see Figure 2.4 [Molecular Probes, WWW]), we can see that approximately:

$$h_{em}(\mathbf{x}) \cong h_{ex}(\mathbf{x}).$$

and since in an unaberrated system $h_{ex}(\mathbf{x})$ is point-symmetric, we can derive from equations (2.15) and (2.17) that the confocal PSF is approximately the square of the widefield PSF:

$$h_{\text{confocal}}(\mathbf{x}) \cong h_{\text{widefield}}^2(\mathbf{x}).$$

The following discussion of optical aberrations is based on [Weisstein, WWWc]: an optical aberration is a distortion in the image formed by an optical system compared to the original system. In our optical system, we have only taken into account the effects of diffraction, and have assumed that for the rest, the optical system is perfect. In reality, several other factors can distort the image:

(scanning) the actual PSF point by point. Or equivalently, a single intensity from a location within the PSF is obtained.

- Astigmatism occurs in lenses because a lens has different focal distances for rays hitting the lens under different angles; this results in a distortion of the image. In particular, rays of light from horizontal and vertical lines in a plane on the object are not focused to the same image plane.
- Chromatic aberration occurs in lenses because lenses focus different wavelengths of light at different distances.
- Spherical aberration happens with a spherical lens or mirror because those do not focus parallel rays to a point, but instead along a line.
- In confocal microscopy, a common aberration is the aberration due to refractive index mismatch [Hell et al., 1993, Booth et al., 1998, Booth and Wilson, 2001]. This means that the differences in refractive indices between the objective lens, the coverslip and the specimen (see Figure 2.21) cause the PSF to degrade even further, and to become less symmetrical (especially in the z-direction).
- ...

When we compare the PSF's for the confocal and the widefield microscope (as shown in Figure 2.22), we notice that the cross section of the PSF in the focal plane (xy-section, PSF_{xy} in Figure 2.22) is similar. However, the confocal PSF is much better confined in the z-direction (PSF_{xz} in Figure 2.22).

The Optical Transfer Function (OTF), which is the magnitude of the Fourier transform of the PSF, has a black area in the z-direction. This zone is known as the “cone of missing frequencies.” Objects with a lot of energy in this cone are poorly imaged by a widefield microscope.

However, when we compare with the OTF of the confocal microscope, we notice that there is no missing frequency cone. This means that these frequencies are imaged, thus allowing better resolution in the z-direction.

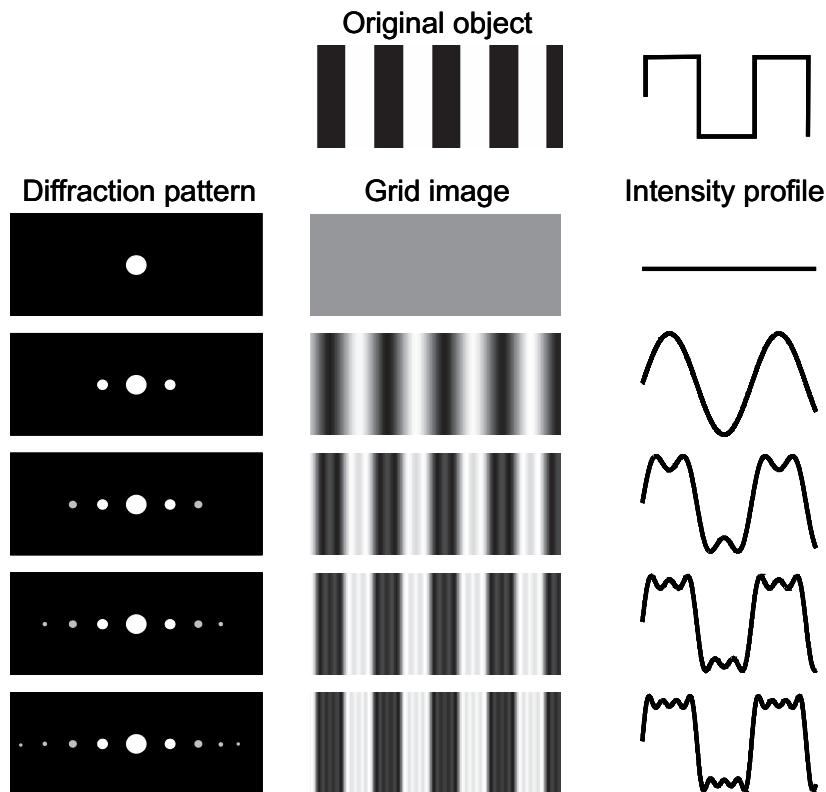


Figure 2.15: This diagram shows the principle of the image formation in the theory of Abbe [Plášek and Reischig, 1998]. On top, the original grid is shown. The first row shows what the image of the grid would be like when only the zeroth order diffraction maximum is intercepted by the lens, and thus contributes to the image. The lower row in the diagram, the more diffraction maxima are intercepted by the lens and contribute to the image and the more accurate the image resembles the original object.

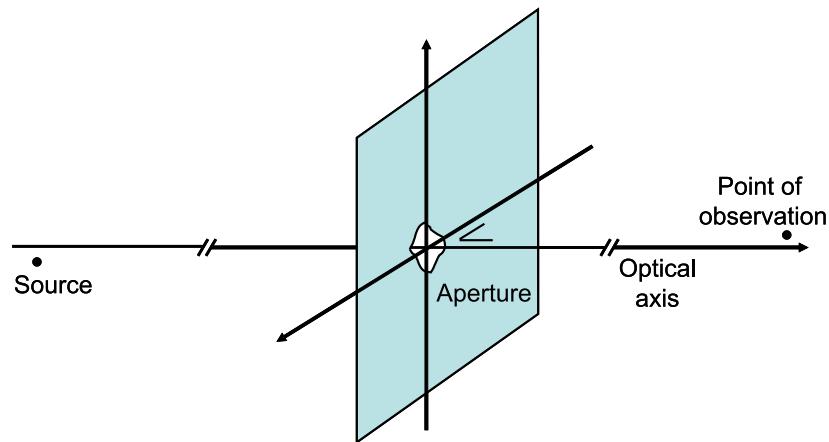
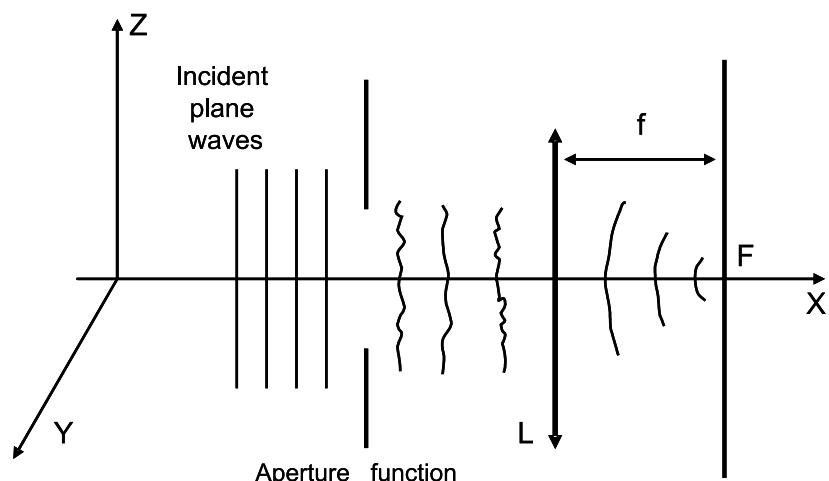


Figure 2.16: Conditions for the Fresnel approximation to be valid.



$$T(y, z) = \begin{cases} 1 & \text{inside the aperture} \\ 0 & \text{outside the aperture} \end{cases}$$

Figure 2.17: Fraunhofer diffraction.

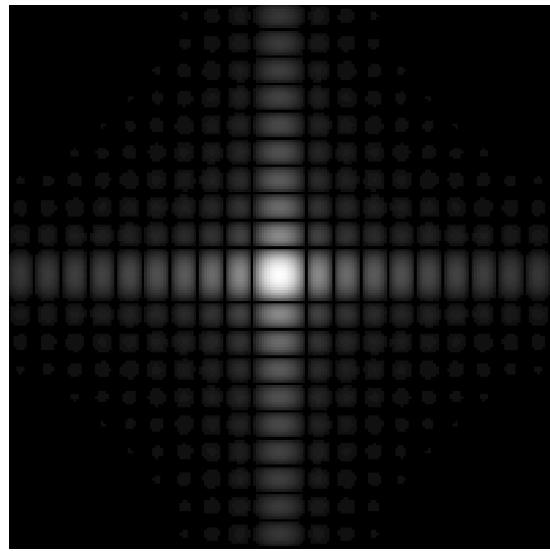


Figure 2.18: Fraunhofer diffraction of plane waves on a rectangular aperture.

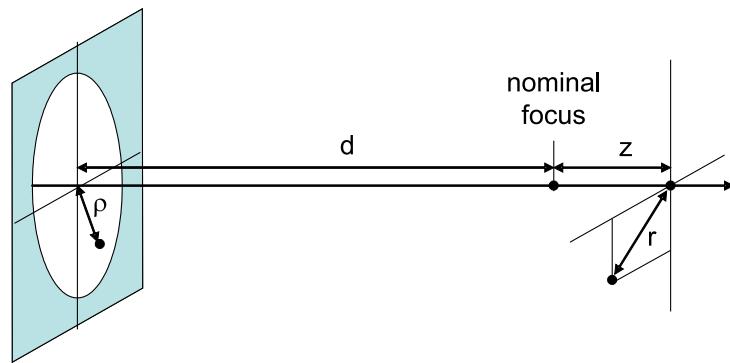


Figure 2.19: Illustration of the notation for the coordinates used in equations (2.8), (2.9) and (2.10).

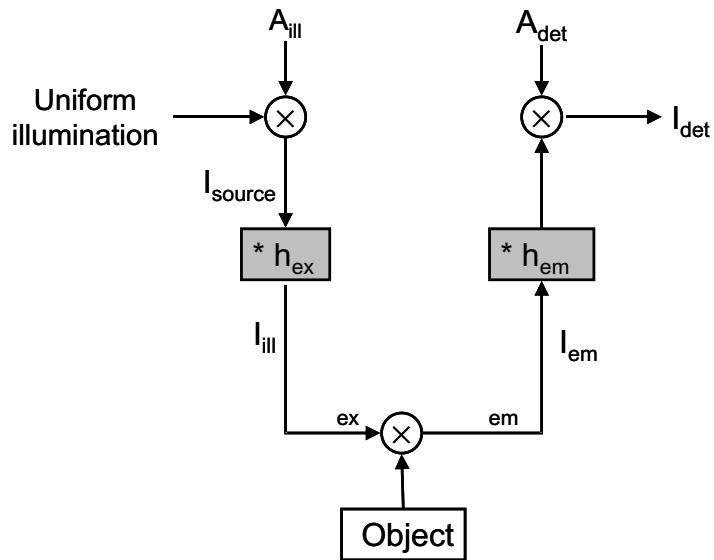


Figure 2.20: Abstract diagram of image formation of a general fluorescence microscope. The circles with the crosses indicate multiplication, the gray boxes indicate convolution with h_{ex} respectively h_{em} .

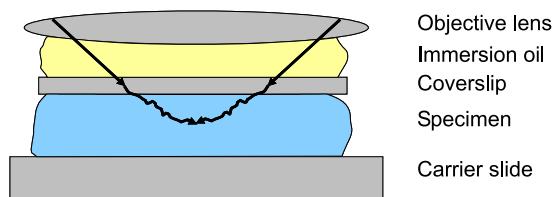


Figure 2.21: Typical specimen preparation in microscopy: between the objective lens and the cover slip (a small piece of glass that covers the specimen), a drop of immersion oil is added. Immersion oil has the same refractive index as glass, so from the objective until the cover slip, no mismatches occur. However, the specimen itself mainly consists of water and has a spatially variant refractive index.

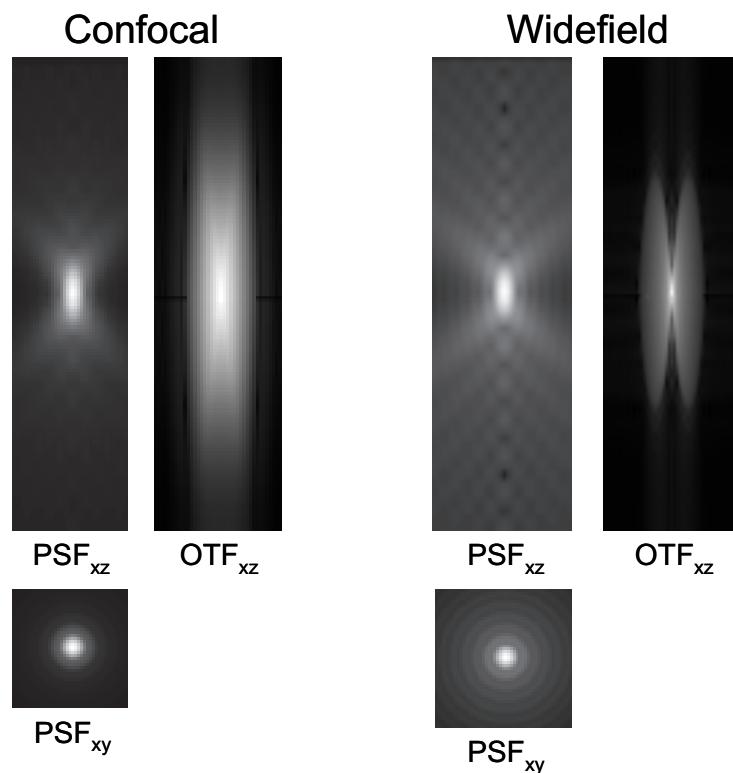


Figure 2.22: Comparison of the PSF and OTF of a confocal and a widefield fluorescence microscope. Images generated with Huygens software, SVI, the Netherlands, <http://www.svi.nl/>.

Chapter 3

Classical restoration techniques

enhance, v.: To tamper with an image, usually to its detriment.

In this Chapter, we discuss a number of classical image restoration algorithms. With “classical,” we mean “not based on multiresolution techniques.” We start from the inverse filter, and discuss the most common problem in image restoration, i.e., ill-posedness, or instability of the solution with respect to small changes in the input. The Chapter is divided in two sections. The first Section approaches the image restoration problem in an intuitive way, while Section 3.2 considers the problem of image restoration starting from Bayes’ rule.

A value that is frequently used to evaluate the performance of restoration algorithms, is the Peak Signal to Noise Ratio (PSNR). The PSNR is defined by:

$$PSNR = 10 \log_{10} \left(\frac{I_{\max}^2}{MSE} \right)$$

where MSE stands for the *mean square error* between the pixel intensities of the two images (ideal¹ and restored). I_{\max} is the maximal possible image intensity. For 8 bit images, e.g., $I_{\max} = 255$. For some results of restoration algorithms, we will use this measure.

¹So the ideal image must be available. This is the case, e.g., when synthetic degradation is applied to synthetic test images.

3.1 Inverse filter and Wiener filter

3.1.1 Inverse filter

The most basic approach towards image restoration is based on the mathematical property that convolution in the spatial domain is transformed into multiplication in the Fourier domain. The image formation equation in the spatial domain is given by²:

$$g(x, y) = (f * h)(x, y) + n(x, y),$$

where $g(x, y)$ is again the degraded image, $f(x, y)$ the ideal unknown image, $h(x, y)$ the PSF and $n(x, y)$ the noise in the image. Transformation to the Fourier domain yields:

$$G(u, v) = F(u, v)H(u, v) + N(u, v) \quad (3.1)$$

with $G(u, v)$, $F(u, v)$, $H(u, v)$ and $N(u, v)$ the Fourier transforms of $g(x, y)$, $f(x, y)$, $h(x, y)$ and $n(x, y)$, respectively, and u and v coordinates in the Fourier space. So, a naive estimate for the ideal (Fourier transform of the) image $\hat{F}(u, v)$ would be:

$$\hat{F}(u, v) = \frac{G(u, v)}{H(u, v)} \quad (3.2)$$

$$= F(u, v) + \frac{N(u, v)}{H(u, v)}. \quad (3.3)$$

From this expression, it is clear that this estimate does not contain useful information at frequencies where $H(u, v)$ is very small (or even zero), because $G(u, v)$ and $N(u, v)$ are then divided by these tiny values, which are in the order of magnitude of $0 \leq H(u, v) \ll 1$, i.e., the corresponding frequency components are amplified strongly. In Figure 3.1 (a) and (b), an image and its blurred version are shown. Figure 3.1 (c) is the result of the inverse filter in the case of deblurring the noise-free, blurred image. From Figure 3.1(e), it is clear that deblurring this same image in the presence of a little noise is highly unstable.

In scientific terms, it is said that this inversion is an ill-posed problem, or ill-conditioned. Therefore, it is highly desirable to stabilize the methods to solve the restoration problem by imposing prior knowledge. This is done by

²For the sake of compactness of notations, we consider 2D images here.

defining certain criteria that an image must satisfy to be a likely solution of the restoration problem. Images whose intensities fluctuate between black and white from pixel to pixel (e.g., Figure 3.1 (e)) are less likely than images where larger connected areas have similar intensities. The corresponding criterion is an example of a smoothness criterion, which is used under different forms as prior knowledge.

The simplest stabilization procedure for the inverse filter is to only allow the spectral components of the signal above the noise level to be amplified, and to set the other spectral components to zero. An example of an image restored with this technique is given in Figure 3.1(f).

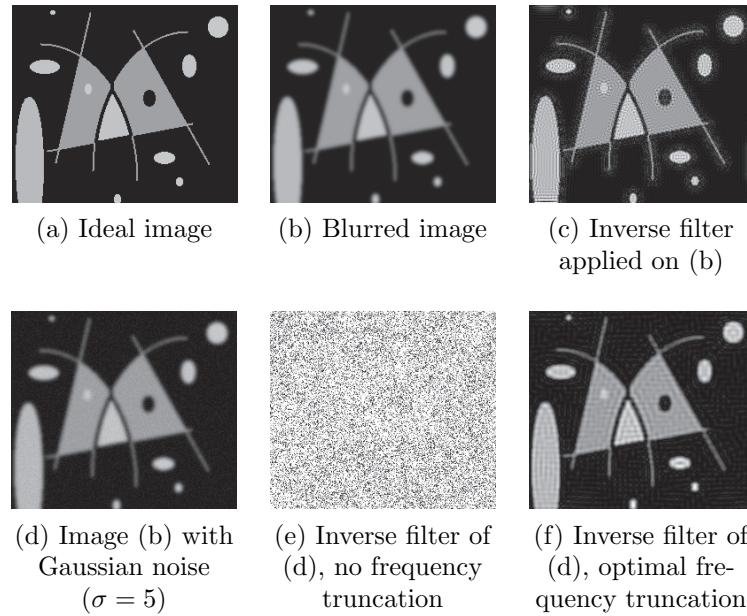


Figure 3.1: Some examples of inverse filtering.

3.1.2 Wiener filter

The Wiener filter has a long history that goes back to the Wiener-Hopf equations derived by Norbert Wiener and Eberhard Hopf in 1930. The Wiener filter is in fact a solution of these equations.

In essence, the Wiener filter behaves similar to the truncated inverse filter, except that an optimal trade-off is made in the area where the noise power becomes of the same order of magnitude as the signal power. In fact, the truncated inverse filter amplifies the amplitude of the spectral components of the signal for whose energy dominates that of the noise; spectral components

with an energy smaller than the noise energy are set to zero (which is fact a kind of global hard-thresholding in Fourier space).

On the other hand, the Wiener filter also amplifies the amplitude of the spectral components of the signal whose energy dominates the noise energy, but performs an optimal least-square shrinkage of the spectral components where the noise starts to dominate the signal (see Figure 3.2).

Let us recall the degradation equation:

$$g(x, y) = (f * h)(x, y) + n(x, y),$$

or in the Fourier domain:

$$G(u, v) = F(u, v)H(u, v) + N(u, v). \quad (3.4)$$

The Wiener filter searches for a least squares solution for the restoration, i.e., it computes an optimal filter $w(x, y)$ for which the solution (derivation based on [Hossack, WWW, van Kempen, 1999, Gonzalez and Woods, 1993]):

$$\hat{f}(x, y) = (g * w)(x, y)$$

minimizes the mean square error (MSE):

$$MSE = E \left(\left| \hat{f}(x, y) - f(x, y) \right|^2 \right)$$

where $\hat{f}(x, y)$ is the estimated solution, and $f(x, y)$ the unknown ideal image. Since the same information is present in the spatial domain as well as the Fourier domain, it is equivalent to minimizing:

$$\begin{aligned} MSE &= E \left(\left| \hat{F}(u, v) - F(u, v) \right|^2 \right) \\ &= E \left(\left| \hat{W}(u, v)G(u, v) - F(u, v) \right|^2 \right) \end{aligned}$$

So, we impose

$$\frac{\partial}{\partial W(u,v)} E \left(\left| \hat{F}(u,v) - F(u,v) \right|^2 \right) = 0.$$

When introduce the notation A^* for the complex conjugate of A , this means that

$$\frac{\partial}{\partial W(u,v)} E \left((W(u,v)G(u,v) - F(u,v))^* (W(u,v)G(u,v) - F(u,v)) \right) = 0$$

or

$$E \left(2G^*(u,v) (W(u,v)G(u,v) - F(u,v)) \right) = 0$$

Substituting equation (3.4) for $G(u,v)$, the last equation becomes:

$$E \left((H(u,v)F(u,v) + N(u,v))^* (W(u,v)(H(u,v)F(u,v) + N(u,v)) - F(u,v)) \right) = 0$$

Since the noise is assumed to be independent of the signal and to have expectancy value zero, we have:

$$E(N(u,v)F(u,v)) = E(N(u,v))E(F(u,v)) = 0$$

So we can rewrite the equation as

$$\begin{aligned} H^*(u,v)H(u,v) &= E(F^*(u,v)F(u,v))W(u,v) + E(N^*(u,v)N(u,v))W(u,v) \\ &= H^*(u,v)E(F^*(u,v)F(u,v)) \end{aligned} \tag{3.5}$$

Further more, the quantities

$$S_{ff}(u,v) \triangleq E(F^*(u,v)F(u,v)),$$

$$S_{nn}(u, v) \triangleq E(N^*(u, v)N(u, v))$$

are defined as the power spectrum of the ideal signal (which is unknown, and has to be approximated or estimated) and the power spectrum of the noise. So, equation (3.5) becomes:

$$W(u, v) \left(|H(u, v)|^2 S_{ff}(u, v) + S_{nn}(u, v) \right) = H^*(u, v)S_{ff}(u, v)$$

which is equivalent with:

$$W(u, v) = \frac{H^*(u, v)S_{ff}(u, v)}{|H(u, v)|^2 S_{ff}(u, v) + S_{nn}(u, v)}$$

or:

$$\hat{F}(u, v) = \frac{H^*(u, v)}{|H(u, v)|^2 + S_{nn}(u, v)/S_{ff}(u, v)} G(u, v)$$

which is the formula for the Wiener filter in the Fourier domain given in all text books. From this expression, it is clear that spectral components with $S_{nn}(u, v) \ll S_{ff}(u, v)$ are amplified because

$$\frac{S_{nn}(u, v)}{S_{ff}(u, v)} \ll |H(u, v)|^2$$

On the other hand, when $S_{nn}(u, v) \approx S_{ff}(u, v)$, all the spectral components (signal and noise) are attenuated to reduce the noise energy in the restored image (at the expense of the suppression of the signal components).

3.2 Bayesian restoration

Most image restoration algorithms aim to find the most probable solution f given a degraded observation g . In this case, f is the unknown, ideal image as if recorded in ideal circumstances (no blur, noise, ...), while g is the degraded image we actually can observe. To reconstruct f from g , we need a prior distribution $P(f)$, which is used to incorporate our knowledge about the expected structure of the image.

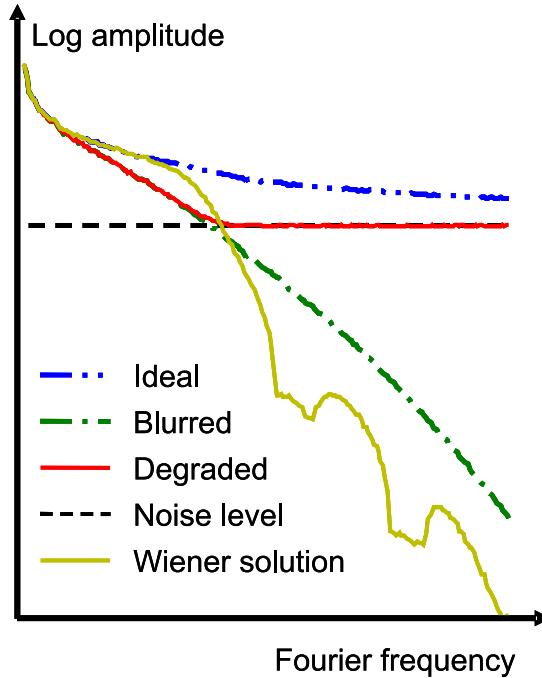


Figure 3.2: Spectra of ideal image, degraded image and Wiener restored image.

We also need a conditional distribution $P(g|f)$, i.e., a degradation model that relates the observed, degraded image to an ideal object as input (with a good model for the optics and the detector, it is fairly easy to calculate how an ideal object is imaged). On the other hand, $P(f)$ is more difficult to model or estimate, since it reflects what we expect for a plausible image to look like.

Knowing $P(g|f)$ and the prior probability $P(f)$, it is possible to infer the posterior probability $P(f|g)$ by using Bayes' rule:

$$P(f|g) = \frac{P(g|f)P(f)}{P(g)}$$

Bayesian methods for image restoration can use this information in several ways. Methods like *Minimum Mean Square Error* (MMSE) try to minimize the quadratic error between the estimate and the “true solution, while the *Maximum A Posteriori* (MAP) try to find a solution that maximizes the poster probability. In this thesis, we focus on the latter class of methods.

In the case of MAP methods, we wish to maximize $P(f|g)$, i.e., find the most probable ideal image f , given the observation g . In practice however, it is often easier to minimize minus the logarithm of the following expression:

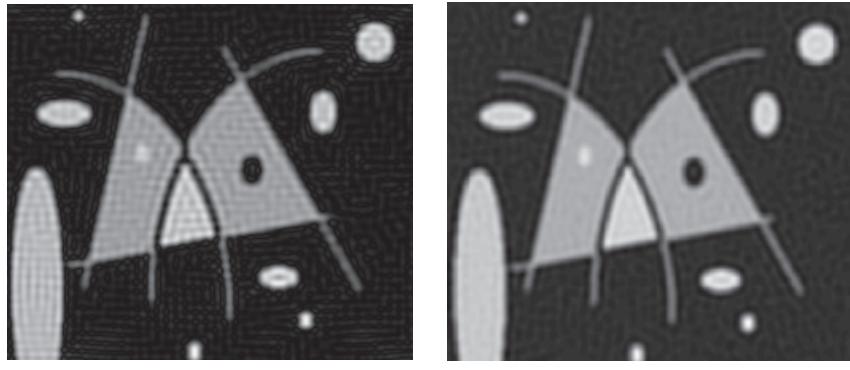


Figure 3.3: Comparison of truncated inverse filter and the Wiener filter.

$$-\log(P(f|g)) = -\log(P(g|f)) - \log(P(f)) + \log(P(g))$$

where the term $\log(P(g))$ is irrelevant in the minimization. It is common to call $L(f) = \log(P(g|f))$ the log-likelihood functional and $R(f) = \log(P(f))$ the regularization term, which carries the prior information we impose on the solution.

Maximizing only the log-likelihood (i.e., minimizing minus the log-likelihood) is called Maximum Likelihood estimation (ML); this is often used when no prior information is available. Maximizing the full expression (or minimizing minus the expression) is called Maximum A Posteriori estimation (MAP). Methods that use an entropy prior model for the image are referred to as Maximum Entropy methods (MEM). A good review paper on this class of methods is [Narayan and Nityananda, 1986], but the principles can be found in [Molina et al., 2001].

3.2.1 Gaussian noise model

The most common algorithms are derived from this expression for the cases of Gaussian and Poisson noise, combined with a variety of regularization functions. In this subsection, we start with Gaussian noise, which is the simplest case.

3.2.1.1 Tikhonov-Miller restoration versus Wiener filter

The likelihood functional in case of Gaussian noise is:

$$P(g|f) \sim \exp\left(-\frac{1}{2}\left\|g - (h * f)\right\|^2\right).$$

The corresponding log-likelihood function $-\log P(f|g)$ is then proportional to

$$-\log P(g|f) \sim \frac{1}{2}\left\|g - (h * f)\right\|^2 \quad (3.6)$$

In many image restoration methods (like here in Tikhonov-Miller restoration), a smoothness constraint is used as a regularization functional in order to reduce oscillations in the image due to noise amplification. In fact, regularization is imposing prior knowledge on the solution, i.e., we want to impose which kind of solutions, are more likely. E.g., we expect solutions with smooth areas to be more likely than solutions with random oscillations in the image. Therefore, we impose a limit on the high frequency content of the image by adding a fraction of high frequency energy as a penalty in the total functional to minimize, i.e. $\log P(f)$ is,

$$\log(P(f)) \sim \left\|(c * f)(x, y)\right\|^2$$

where f is the estimated solution, and c is a highpass filter (e.g., the Laplacian filter kernel is a popular choice for c) that is to be convolved with the solution f . The Tikhonov-Miller functional to be minimized is therefore:

$$\Phi(\hat{f}) = L(\hat{f}(x, y)) + \lambda R(\hat{f}(x, y)) \quad (3.7)$$

$$= \left\|g(x, y) - (h * \hat{f})(x, y)\right\|^2 + \lambda \left\|(c * \hat{f})(x, y)\right\|^2 \quad (3.8)$$

where $\|x\|^2$ is the Euclidean norm of x . The constant λ is the so-called regularization parameter, and determines the relative importance of the regularization to the importance of the data (the larger λ , the more importance is given to the regularization; the smaller λ , the more importance is given to the observed data).

Using Parseval's lemma, this is equivalent with the Fourier transform of equation (3.8):

$$\Phi(\hat{F}) = \left\|G(u, v) - H(u, v)\hat{F}(u, v)\right\|^2 + \lambda \left\|C(u, v)\hat{F}(u, v)\right\|^2$$

and its gradient with respect to $\hat{F}(u, v)$:

$$\begin{aligned}\nabla_{\hat{F}} \Phi(\hat{F}) &= -2H^*(u, v) \left(G(u, v) - H(u, v)\hat{F}(u, v) \right) \\ &\quad + 2\lambda |C(u, v)|^2 \hat{F}(u, v) \\ &= -2H^*(u, v)G(u, v) + 2A(u, v)\hat{F}(u, v)\end{aligned}\tag{3.9}$$

with

$$A(u, v) = |H(u, v)|^2 + \lambda |C(u, v)|^2.\tag{3.10}$$

So its maximum is reached for:

$$-H^*(u, v) \left(G(u, v) - H(u, v)\hat{F}(u, v) \right) + \lambda |C(u, v)|^2 \hat{F}(u, v) = 0\tag{3.11}$$

or, equivalently:

$$\hat{F}(u, v) = \frac{H^*(u, v)}{|H(u, v)|^2 + \lambda |C(u, v)|^2} G(u, v),\tag{3.12}$$

which is equivalent to the solution of the Wiener filter when λ and C are chosen such that

$$\lambda |C(u, v)|^2 = \frac{S_{nn}(u, v)}{S_{ff}(u, v)}.$$

3.2.1.2 Intermezzo: gradient-based minimization methods

Since the minimum of a function $f(x)$ coincides with a zero of the gradient:

$$\hat{x} = \arg \min_x f(x) \Rightarrow [\nabla_x f(x)]_{x=\hat{x}} = 0,$$

a simple gradient-based minimization procedure is to add a negative fraction of the gradient every iteration to the previous estimate in order to find a smaller $f(x)$. That results in the following scheme:

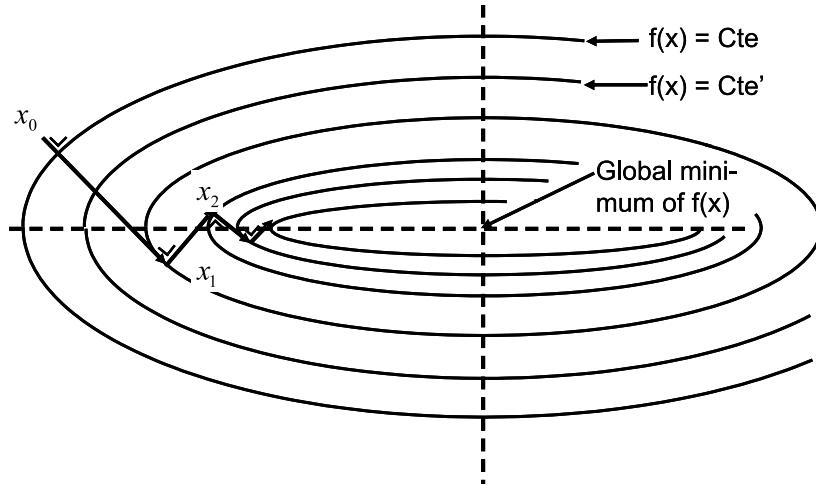


Figure 3.4: Classical steepest descent doesn't move directly towards the global minimum.

- Start with an arbitrary solution $f(x_0)$ ($k = 0$);
- Compute the gradient of the function at the k^{th} estimation of the minimum $\nabla_x f(x_k)$;
- Move in the direction of the steepest descent at position x_k to obtain a new estimate $x_{k+1} = x_k - \lambda \nabla_x f(x_k)$;
- Go to the second step until convergence.

The parameter λ is known as the “learning rate”; this parameter determines the convergence speed. When λ is taken too big, convergence might not be reached, since the final estimates can keep oscillating around the global minimum, or worse: the iterations could make the estimate leave the region of convergence.

This method is known as the *Steepest Descent* minimization method, and has as disadvantage that the convergence is often rather slow, because the local steepest descent doesn't always point in the direction of the global minimum, and the iteration steps approach the global minimum in an inefficient, “zig-zag” way. This is illustrated in Figure 3.4. The ellipses are contours which have the same value (iso-contours) of the function to be minimized. We start at position x_0 . The arrow points at the direction of steepest descent for x_0 (which is orthogonal to the tangent line of the iso-contour in x_0). It is clear that the directions of steepest descent for the different estimates do not lead to the global minimum of the function (which is located at the center of the ellipses).

In order to speed up this slow convergence, the technique of Conjugate Gradients (CG) was introduced (a nice introduction to the conjugate gradients meth-

ods is presented in [Shewchuk, 1994]). In this iterative optimization scheme, the new steepest descent direction is not followed, but rather a direction that is a linear combination of the steepest descent direction and the preceding direction. In the plane that is determined by these two directions, a new search direction is computed by optimization in this plane (so we have an extra degree of freedom, and thus have more possible solutions). The explicit scheme is summarized as follows: the direction of steepest descent during iteration k is given by:

$$r^{(k)}(x, y) = -\frac{1}{2} \nabla_f \Phi(f_k(x, y)) \quad (3.13)$$

with $\nabla_f \Phi(f)$ given in equation (3.9). Then the conjugate gradient direction for iteration k is given by:

$$d^{(k)}(x, y) = r^{(k-1)}(x, y) + \gamma^{(k)} d^{(k-1)}(x, y)$$

with

$$\gamma^{(k)} = \frac{\|r^{(k)}(x, y)\|^2}{\|r^{(k-1)}(x, y)\|^2}$$

A new estimate $\hat{f}^{(k+1)}$ is then obtained as follows:

$$\hat{f}^{(k+1)}(x, y) = \hat{f}^{(k)}(x, y) + \beta^{(k)} d^{(k)}(x, y)$$

3.2.1.3 Iterative Constrained Tikhonov-Miller

The Wiener filter and the Tikhonov-Miller restoration algorithm are examples of linear algorithms, which have two important disadvantages:

1. Only those spectral components of the signal that have not been removed completely by the degradation process, can be recovered. Recall Figure 2.22, where the 3D Optical Transfer Function (OTF) of a widefield microscope was given. This OTF has a cone of missing frequencies in the z-direction. The corresponding spectral components cannot be recovered with linear algorithms.

2. Linear models cannot restrict the set within the subspace of possible solutions. An image, e.g., can only have positive gray values. The algorithms we discussed previously do not guarantee positivity, even though negative gray values have no physical meaning. It would be advantageous to restrict the domain possible solutions even further to only those solutions that have a physical meaning (e.g., images which contain negative values, or which are too “rough”, like Figure 3.1(e) should not be allowed as real physical solutions).

In this Subsubsection, we discuss a variant of the Tikhonov-Miller restoration algorithm which imposes non-negativity: the so-called *Iterative Constraint Tikhonov-Miller* (ICTM) algorithm. The practical implementation of the ICTM computes the minimum of equation (3.6) in an iterative way using the CG technique discussed in the previous subsubsection. The non-linearity here is imposed as a hard constraint: after every iteration, the negative values of the current estimated solution are set to zero by the positivity constraint operator $P(x)$:

$$P(x) = \begin{cases} 0 & x \leq 0 \\ x & x \geq 0 \end{cases} \quad (3.14)$$

So, our new estimate $\hat{f}^{(k+1)}(x, y)$ becomes:

$$\hat{f}^{(k+1)}(x, y) = P \left(\hat{f}^{(k)}(x, y) + \beta^{(k)} d^{(k)}(x, y) \right) \quad (3.15)$$

The operator $P(x)$ is a highly nonlinear operator, and introduces the extra prior knowledge that an image contains only positive intensities.

In the absence of this positivity constraint operator, the value $\beta^{(k)}$ can be calculated analytically [Press et al., 1988, van Kempen, 1999]:

$$\beta^{(k)} = \frac{\sum_{x,y} d^{(k)}(x, y) r^{(k)}(x, y)}{\sum_{x,y} d^{(k)}(x, y) [(a * d^{(k)})(x, y)]}$$

where $a(x, y)$ is the *inverse Fourier transform* of $A(u, v)$.

However, in presence of the positivity constraint operator, this expression is no longer valid and cannot be calculated analytically. Nevertheless, it can be approximated by searching for the zero crossing of the gradient along the search direction by solving $q(\beta^{(k)}) = 0$ [Verveer and Jovin, 1997]:

$$q(\beta^{(k)}) = \sum_{x,y} d^{(k)}(x,y) \nabla_f \Phi(P(f(x,y)) + \beta^{(k)} d^{(k)}(x,y)) \quad (3.16)$$

with $\Phi(f)$ given in equation (3.8). This non-linear equation is linearized around $\beta = 0$, i.e., $\beta^{(k)}$ is the solution of

$$q(\beta^{(k)}) = q(0) + \beta^k q'(0) = 0$$

Solving equation (3.16) for β results in the following expression [Verveer and Jovin, 1997]:

$$\beta^{(k)} = \frac{\sum_{x,y} d^{(k)}(x,y) P(\hat{f}^{(k)}(x,y)) r^{(k)}(x,y)}{\sum_{x,y} \left(d^{(k)}(x,y) P(\hat{f}^{(k)}(x,y)) \right) \left(a(x,y) * \left(P(\hat{f}^{(k)}(x,y)) d^{(k)}(x,y) \right) \right)}$$

In Figure 3.5, we show our example image restored with the Wiener filter and with the ICTM algorithm.

What we didn't discuss yet, is the choice of the regularization parameter λ . For the sake of completeness, we will briefly mention some popular choices for λ [Galatsanos and Katsaggelos, 1992, van Kempen, 1999]:

1. Methods that require the knowledge of the noise variance:

- **SNR method:** this method sets λ equal to the inverse of the SNR (signal to noise ratio):

$$\lambda_{\text{SNR}} = \frac{\text{Energy}_{\text{noise}}}{\text{Energy}_{\text{signal}}}$$

with the energy of the noise given by $\sum_{x,y} n^2(x, y)$ and the energy of the signal by $\sum_{x,y} f^2(x, y)$. Both the nominator and the denominator are unknown, and are estimated. Using this regularization parameter has a clear intuitive interpretation: the higher the SNR, the more we need to suppress the noise. The inverse SNR is then bigger, so more weight is given to the regularization than to the data fit.

- **Method of constrained least squares (CLS):** this method determines the value λ_{CLS} for which the total squared error between the observed image and the blurred restored image is minimal:

$$\sum_{x,y} \left| g(x, y) - \left(h * \hat{f}_\lambda \right) (x, y) \right|^2$$

$\hat{f}(x, y)$ is estimated by the Tikhonov-Miller algorithm for various values of λ_{CLS} , and the minimum is found with a one-dimensional line-search method.

2. Methods not requiring knowledge of the noise variance:

- **Method of Generalized Cross-Validation:** this method determines the value λ_{GCV} from the so-called *leave one out* principle: for every pixel the Tikhonov-Miller restoration result is calculated using all pixels, except the one under consideration and the left-out pixel is predicted from this reduced, restored image. λ_{GCV} is chosen as the minimizer of the cross-validated prediction error.
- **Maximum likelihood method:** this method calculates the value λ_{ML} based on the stochastic assumption that both the noise and the regularization function $\sqrt{\lambda_{\text{ML}}} (c * f)(x, y)$ have a Gaussian distribution with variance σ_n^2 (where $\sqrt{\lambda_{\text{ML}}}$ scales the variance of $(c * f)(x, y)$ to σ_n^2). The derived likelihood function is then minimized.

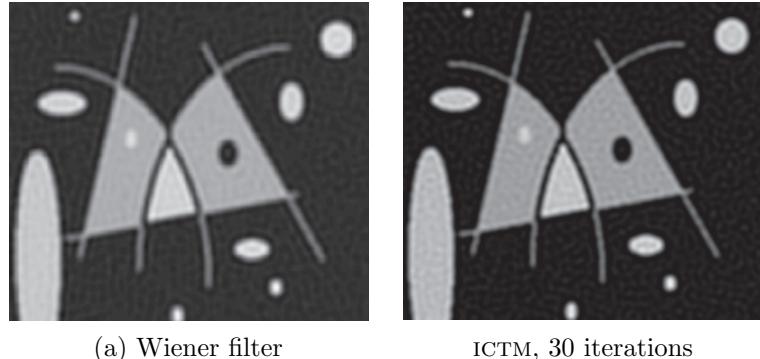


Figure 3.5: Comparison of the Wiener filter with the ICTM algorithm, 30 iterations, using the SNR method to compute λ .

3.2.2 Poisson noise model

3.2.2.1 Richardson-Lucy restoration

The algorithm was introduced independently by [Lucy, 1974] and [Richardson, 1972].

The related Expectation Maximization (EM) algorithm is found in [Dempster et al., 1977]. This method was first put in an iterative algorithm and applied in image reconstruction in [Shepp and Vardi, 1982], which lead to an algorithm equivalent to the Richardson-Lucy algorithm. Background information and references can be found in [Molina et al., 2001]. This algorithm is common practice in confocal image restoration [Holmes et al., 1995, van der Voort and Strasters, 1995, Van Kempen et al., 1997, Verveer et al., 1999, Schaeffer et al., 2001].

First, we derive the Richardson-Lucy algorithm according to [Richardson, 1972]. We use a one-dimensional notation for the sake of compactness; extension to more dimensions is trivial. In [Richardson, 1972], it is assumed that $g(x)$, $f(x)$ and $h(x)$ are discrete, probability frequency functions, not necessarily normalized. This means that the numerical value of a pixel of $g(x)$, $f(x)$ or $h(x)$ is a measure of the frequency that an event occurs in that pixel. Units of energy (usually photons) emitted from a certain location in $f(x)$ are distributed at positions in $g(x)$ according to frequencies indicated by $h(x)$.

As in [Richardson, 1972], our notation $f(x_i)$ here has *two* meanings:

1. $f(x_i)$ is an indicator function which is 1 at position (x_i) when an event (like photon detection) happens in the ideal image array $f(x)$;
2. $f(x_i)$ denotes the intensity value at position (x_i) in image array f .

If we denote according to the first notation convention the occurrence of an event like emission of a photon in the ideal image array by $f(x_1)$, and the actual observation of an event in the degraded image by $g(x_3)$, we can write that the probability that an event $f(x_1)$ took place, which is observed as $g(x_3)$, is given by:

$$P(f(x_1), g(x_3)) = P(f(x_1)|g(x_3)) P(g(x_3)) \quad (3.17)$$

$$= P(g(x_3)|f(x_1)) P(f(x_1)) \quad (3.18)$$

From these equations, it follows that the conditional probability of an event from $f(x_1)$ given the observation of an event at $g(x_3)$ is given by Bayes' rule:

$$P(f(x_1)|g(x_3)) = \frac{P(g(x_3)|f(x_1)) P(f(x_1))}{\sum_{x_2} P(g(x_3)|f(x_2)) P(f(x_2))}$$

We can write:

$$\begin{aligned} P(f(x_1)) &= \sum_{x_3} P(f(x_1), g(x_3)) \\ &= \sum_{x_3} P(f(x_1)|g(x_3)) P(g(x_3)) \\ &= \sum_{x_3} \left(\frac{P(g(x_3)|f(x_1)) P(f(x_1))}{\sum_{x_2} (P(g(x_3)|f(x_2)) P(f(x_2)))} P(g(x_3)) \right) \end{aligned}$$

In other words:

$$1 = \sum_{x_3} \left(\frac{P(g(x_3)|f(x_1))}{\sum_{x_2} [P(g(x_3)|f(x_2)) P(f(x_2))] } P(g(x_3)) \right) \quad (3.19)$$

We define:

$$I = \sum_{x_1} f(x_1) \quad (3.20)$$

$$= \sum_{x_3} g(x_3) \quad (3.21)$$

as the total number of events detected in the image, i.e., the total intensity of the image. We assume that the probability for a given photon to be detected at a given position in the image is uniform over the whole set of possible images: not just for a single observation at a single instant, but for all possible observations of different scenes. This means that a random photon does not prefer one position in the detector over another; in other words, counted photon intensities do not depend on the detector which images the object but on the imaged scene itself:

$$P(f(x_1)) = f(x_1)/I \quad (3.22)$$

$$P(g(x_3)) = g(x_3)/I \quad (3.23)$$

Let us take a closer look at $P(g(x_3)|f(x_1))$: we see that this is in fact the probability that an event $f(x_1)$ is detected as $g(x_3)$ in the observed image array [Hanisch, 1993]. This corresponds to what we previously defined as the PSF, i.e., the observed image of a mathematical point source. So we write $P(g(x_3)|f(x_1)) = h_{x_1, x_3}$. In most cases, it is assumed that h_{x_1, x_3} is spatially invariant, i.e. $h_{x_1, x_3} = h(x_1 - x_3)$, so the sum becomes a convolution:

$$\begin{aligned} f(x_1) &= f(x_1) \sum_{x_3} h(x_3 - x_1) \frac{g(x_3)}{\sum_{x_2} h(x_3 - x_2)f(x_2)} \\ &= M(f(x_1)) \end{aligned}$$

with $M(f(x_1))$ a short hand notation for the right hand side of the previous expression. The previous equation expresses $f(x_1)$ as a function of itself. This is called a fixed point equation [Bertero and Boccacci, 1998], and the Maximum Likelihood estimates are fixed points of $f(x_1) = M(f(x_1))$. If we apply the method of successive approximations to this fixed point equation, we obtain an iterative algorithm:

$$\hat{f}^{(k+1)}(x_1) = M(\hat{f}^{(k)}(x_1))$$

However, not all fixed points are maxima of the likelihood function: the equation has at least N fixed points (N is the number of pixels in the image) with only one component different from zero. It's easy to see that once a certain pixel in the estimate is set to zero, it cannot attain another value anymore. This can be seen in equation (3.24): a new estimate $f^{(k+1)}(x_1)$ is computed from the previous estimate $f^{(k)}(x_1)$ by multiplying this function pixel by pixel with a certain correction function $C^{(k)}(x_1)$. From this, we can also see that if a certain pixel $f^{(k+1)}(x_1)$ during iteration k has a value zero, it will stay zero from then on.

Nevertheless, Shepp and Vardi [Shepp and Vardi, 1982] proved that the likelihood of the iterations always increases, and converges to a maximum if all elements of the initial image are positive, so that this image cannot be a fixed point without being the Maximum Likelihood solution.

So the iterative formula becomes:

$$\begin{aligned} f^{(k+1)}(x_1) &= f^{(k)}(x_1) \sum_{x_3} \left(h(x_3 - x_1) \frac{g(x_3)}{\sum_{x_2} h(x_3 - x_2) f^{(k)}(x_2)} \right) \\ &= f^{(k)}(x_1) C^{(k)}(x_1) \end{aligned} \quad (3.24)$$

This formula can be seen as a forward-backward projection method, as used in the ML-EM³ reconstruction algorithm in tomography, as is illustrated in Figure 3.6.

Another way of obtaining the Richardson-Lucy algorithm is to minimize the log-likelihood function in the presence of Poisson noise [Hanisch, 1993]. The probability of getting an observed image intensity distribution $g(x)$ when the original fluorescence distribution in the object is given by $f(x)$ is given by [Verveer, 1998, van Kempen, 1999, Dey et al., 2004]:

$$P(g|f) = \prod_x \left(\frac{[(h * f)(x)]^{g(x)} e^{-(h * f)(x)}}{g(x)!} \right)$$

which is in fact the mathematical formulation for saying that the probability for observing an image g given an ideal “image” f is given by the probability that the first observed pixel intensity is $g(x_1)$ in case the “true” intensity was $f(x_1)$, multiplied by the probability that a second observed pixel intensity is $g(x_2)$ in case the “true” intensity there was $f(x_2)$, and so on for all pixel intensity probabilities.

³Maximum Likelihood-Expectation Maximization

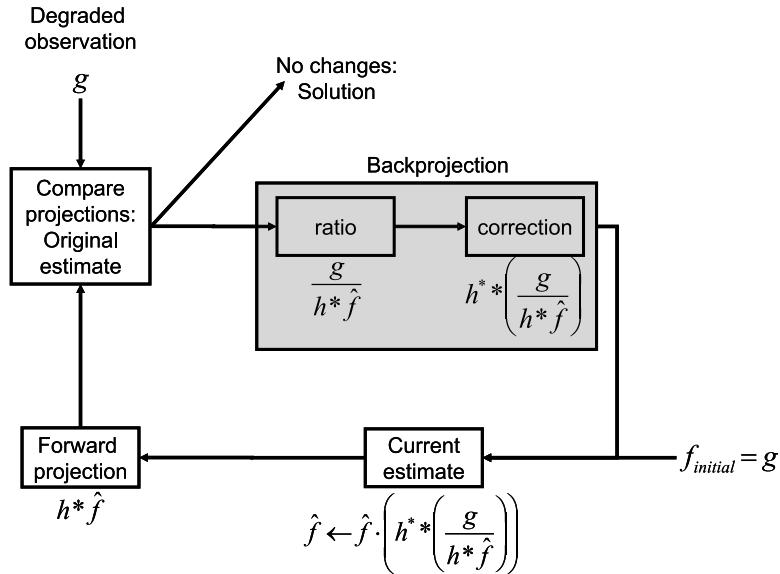


Figure 3.6: Block diagram of the Richardson-Lucy algorithm (based on [Hutton, 2002]).

The negative log-likelihood function $L(f)$ of getting the observed g when $(h * f)$ is expected, is then given by:

$$\log L(f) = \sum_x g(x) \log(h * f)(x) - (h * f)(x) - \log(g(x)!) \quad (3.25)$$

$$\approx \sum_x g(x) \log \left(\frac{(h * f)(x)}{g(x)} \right) + \left((h * f)(x) - g(x) \right) \quad (3.26)$$

where in the second formula the approximation formula of Stirling for factorials was used:

$$\begin{aligned} n! &= n^n e^{-n} \sqrt{2\pi n} \\ \Leftrightarrow \log(n!) &= -n \log(n) + n - \log(\sqrt{2\pi n}) \end{aligned}$$

and $\log(\sqrt{2\pi n})$ can be neglected, so

$$\Rightarrow \log(n!) \approx -n \log(n) + n.$$

Minimizing this log-likelihood with respect to f leads to the Richardson-Lucy algorithm. Note that the last term does not depend on f , and can therefore be dropped in the minimization.

In this respect, we mention the so-called *I-divergence*, used by Csiszár [Csiszár, 1991, Snyder et al., 1992, van Kempen, 1999]. Csiszár formulated a set of axioms (consistency, distinctness, continuity, locality and composition consistency). We briefly explain these axioms:

1. **consistency:** when a function f is selected because it satisfies a number of constraints of a certain selection rule, and f also satisfies some stronger constraints of another selection rule, then we should keep f when these constraints are combined;
2. **distinctness:** Different pieces of information must result in different conclusions about f , unless those pieces of information were irrelevant;
3. **continuity:** is an obvious regularity hypothesis;
4. **locality:** if the available information consists of two pieces of information that involve complementary sets of component of the vector we wish to infer, each component selected based on that information that applies to the component in question;
5. **composition consistency:** if information tells us that individual components in a system consist of two different components, but doesn't say anything about the interaction between these two components, then the selection rule should lead us to the correct result if the two components are not interacting.

Csiszár then examined a number of discrepancy measures between two positive functions $f_1(x)$ and $f_2(x)$, like the mean square error and the Kullback-Leibler distance (sometimes also referred to as cross entropy). He considered the possibility that $f_1(x)$ and $f_2(x)$ are not probability distributions, i.e., not normalized to one, but to arbitrary constants. Therefore, Csiszár extended the Kullback-Leibler distance measure and obtained thus the I-divergence [Csiszár, 1991, Snyder et al., 1992]:

$$I(f_1(x)|f_2(x)) = \sum_x f_1(x) \log \left(\frac{f_1(x)}{f_2(x)} \right) - \sum_x (f_1(x) - f_2(x))$$

Csiszár concluded that for functions that have both positive and negative values, the mean square error is the only discrepancy measure consistent with the axioms mentioned above, while for functions that have only nonnegative values, the I-divergence is the only consistent measure. In [Snyder et al., 1992], it is proven that minimizing $I(g(x)|(h * \hat{f})(x))$ is equivalent to minimizing the negative log-likelihood function for Poisson statistics, which leads to the Richardson-Lucy algorithm. So in a way, the Richardson-Lucy algorithm is also related to Maximum Entropy restoration[Csiszár, 1991].

The Richardson-Lucy algorithm has the following basic properties [Bertero and Boccacci, 1998]:

- every estimate $f^{(k)}(x)$ is non-negative;
- the total intensity of $h * f^{(k+1)}(x)$ corresponds to the total number of counts in the observed image $g(x)$;
- the log-likelihood function $L(f)$ is non-decreasing as the iterations increase;
- $f^{(k)}$ makes $L(f)$ converge to its maximum.

Note that the standard Richardson-Lucy algorithm does *not* include regularization (except for the fact that it keeps total mean intensity constant and it keeps zero values at zero, so noise amplification is somehow limited). In Subsubsection 3.2.2.3, we discuss some regularization schemes for the standard Richardson-Lucy algorithm.

3.2.2.2 Richardson-Lucy illustrated

In this paragraph, we illustrate the operation of the Richardson-Lucy algorithm.

Recall the RL iteration formula:

$$f^{(k+1)}(x_1) = f^{(k)}(x_1) \cdot \sum_{x_3} \left(h(x_3 - x_1) \frac{g(x_3)}{\sum_{x_2} h(x_3 - x_2) f^{(k)}(x_2)} \right) \quad (3.27)$$

$$= f^{(k)}(x_1) \cdot C^{(k)}(x_1) \quad (3.28)$$

where $C^{(k)}(x)$ is the correction function with which $f^{(k+1)}(x)$ has to be multiplied in order to obtain the next estimation.

In Figure 3.2.2.2, we show an example of a synthetic test function (blue) and a blurred observation (green). In Figure 3.8, we show in the left column the estimate $f^{(k+1)}(x)$ for iteration $k + 1$ and in the right column this correction function $C^{(k)}(x)$. It can be seen that where $f^{(k)}(x)$ is originally too small, the tendency is that $C^{(k)}(x) \geq 1$ and where $f^{(k)}(x)$ is too large, $C^{(k)}(x) \leq 1$.

Note also that the correction function becomes more oscillatory when the number of iterations increases, but that the amplitude of these oscillations decreases and that finally $C^{(k)}(x)$ converges to the constant value 1 everywhere.

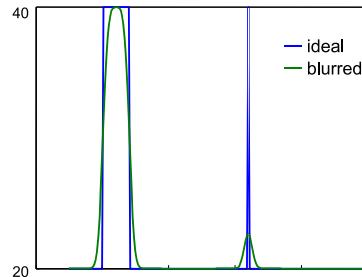


Figure 3.7: A synthetic ideal function and a blurred version of this ideal function.

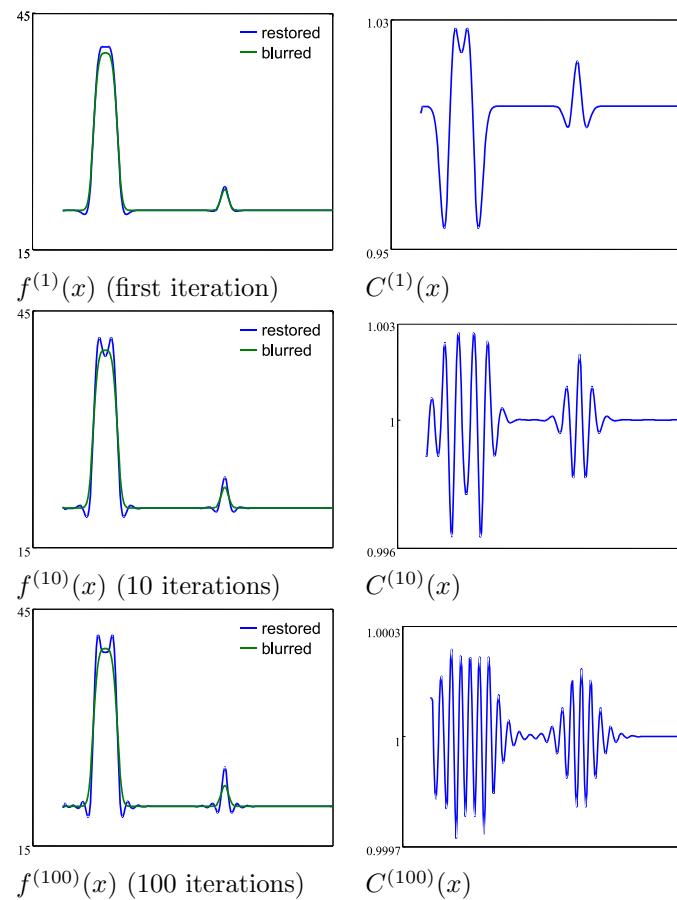


Figure 3.8: Illustration of the operation of the RL algorithm.

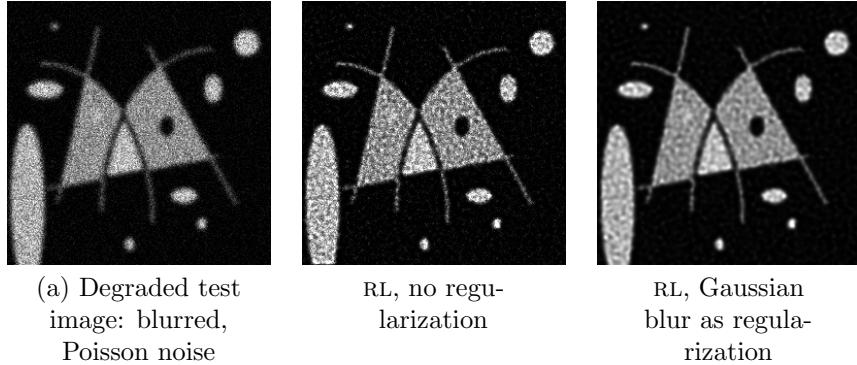


Figure 3.9: Comparison of unregularized versus regularized RL deconvolution.

3.2.2.3 Classical regularization of Richardson-Lucy

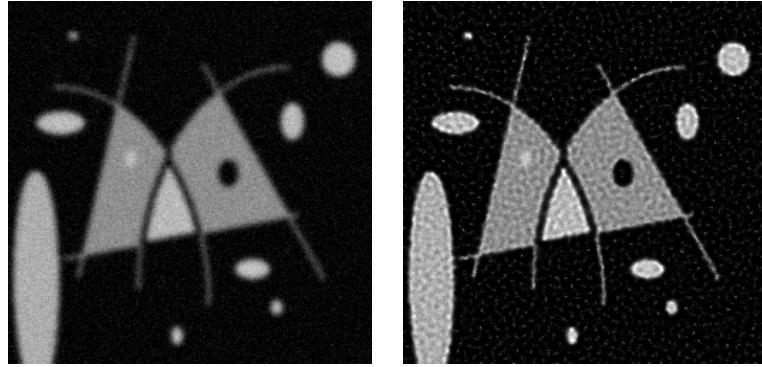
The reference work here is [Verveer, 1998]. There, an extended discussion is given for the minimization of likelihood functions both in case of Poisson data as well as Gaussian data, regularized with a set of different regularization priors.

Here, we will limit ourselves to a few common schemes of regularized Richardson-Lucy, based on [Conchello and McNally, 1996, Verveer, 1998, van Kempen, 1999, Dey et al., 2004]. All the schemes we discuss now were developed for confocal microscope image restoration.

- **RL with slight postblur** The simplest method to regularize the RL algorithm is to apply a blurring after every iteration with a small Gaussian kernel ($\sigma = 0.5$). This imposes some smoothness to the solution by blurring singularities after every iteration. In [Silverman et al., 1990], a similar approach was used for the regularization of the reconstruction of emission tomography images. In [Slijpen and Beekman, 1999], it was shown that in some situations, filtering between the iterations yields similar results to applying a filter after finishing the complete iterative reconstruction.
- **RL with Tikhonov-Miller regularization [Dey et al., 2004]:** The Tikhonov-Miller regularization functional is given by

$$R_{TM}(f) = \sum_{\mathbf{x}} |\nabla f(\mathbf{x})|^2 . \quad (3.29)$$

As mentioned before, this regularization functional imposes a degree of smoothness on the solution by limiting the high frequency content of the image. With this regularization functional, the iterative scheme becomes:



(a) Degraded synthetical test image, blurred and Gaussian noise

RL, Tikhonov-Miller regularization

Figure 3.10: Degraded test image and restored with RL deconvolution, Tikhonov-Miller regularized.

$$\hat{f}^{(k+1)}(x, y) = \left(\frac{g(x, y)}{\left(h * \hat{f}^{(k)} \right)(x, y)} * h^*(x, y) \right) \frac{\hat{f}^{(k)}(x, y)}{1 + 2\lambda\Delta\hat{f}^{(k)}(x, y)}$$

with Δ the Laplacian operator:

$$\Delta f(x, y) = \frac{\partial^2 f(x, y)}{\partial x^2} + \frac{\partial^2 f(x, y)}{\partial y^2}$$

- **RL Conchello** [Conchello and McNally, 1996, van Kempen, 1999] Conchello proposes a regularization functional that penalizes only the excessively bright spots, but doesn't impose smoothness on the restored image. Equation (3.25) gives the log-likelihood function for Poisson noise. Pure Richardson-Lucy deconvolution minimizes this function. In the Conchello algorithm, the following regularization functional is added to the functional to be minimized:

$$R(f) = \sum_{\mathbf{x}} |f(\mathbf{x})|^2 .$$

This results in a regularized iterative scheme of the form:

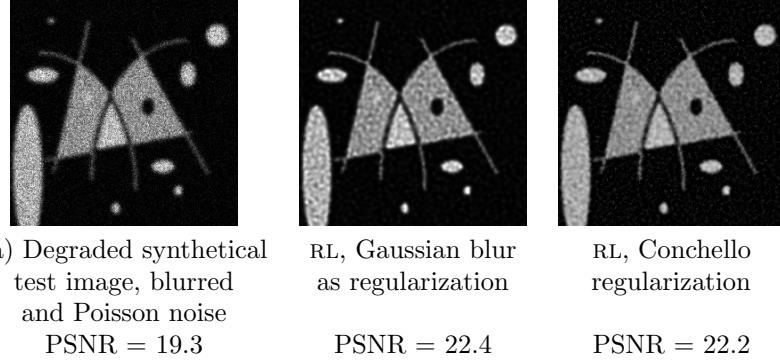


Figure 3.11: Comparison of unregularized versus regularized RL deconvolution.

$$\hat{f}_{\text{Conchello}}^{(k+1)}(x, y) = \frac{-1 + \sqrt{1 + 2\lambda \hat{f}_{\text{unregularized}}^{(k+1)}(x, y)}}{\lambda}$$

where $\hat{f}_{\text{unregularized}}^{(k+1)}(x, y)$ is the result applying a single unregularized RL iteration step to the previous regularized estimation $f_{\text{Conchello}}^{(k)}(x, y)$.

Conchello describes also an estimation for an appropriate value for λ in this scheme [Conchello and McNally, 1996]: $1/\lambda$ is chosen to be equal to the largest pixel value of the estimated specimen function. For images with bright pixels in isolated areas comparable in size to the diffraction limited spot, $1/\lambda$ can be approximated by the maximum value in the image divided by the maximum value of the Point Spread Function. For images with bright pixels in line-like areas, $1/\lambda$ can be approximated by the maximum value in the image divided by the maximum value of the Line Spread Function (i.e., the Point Spread Function integrated orthogonally to the direction of the line structure).

- **RL with total variation regularization [Dey et al., 2004]:** The total variation regularization functional is given by

$$R_{TV}(f) = \sum_{\mathbf{x}} |\nabla f(\mathbf{x})| .$$

Using the L_1 norm instead of the L_2 norm as in equation (3.29) allows edge preserving smoothing. It can be shown that the smoothing process introduced by this regularization functional acts only in the direction tangential to the edge gradients, and not in the direction orthogonal to edge

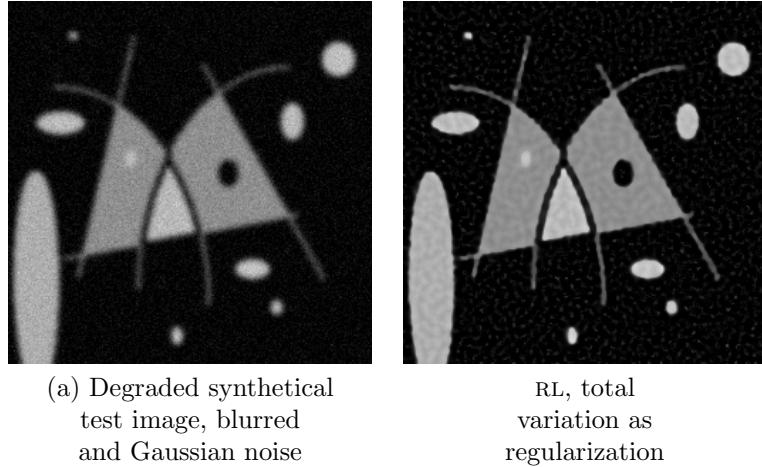


Figure 3.12: Degraded test image (Gaussian noise) and restored with RL deconvolution, regularized with total variation.

gradients [Dey et al., 2004]. In fact, this regularization method is closely related to anisotropic diffusion [Rudin et al., 1992, Chan et al., 2001], where smoothing is isn't allowed across large image gradients.

Total variation is very efficient for objects with smooth regions and no texture, since it smoothes homogeneous regions and restores sharp edges. However, fine texture is destroyed. Minimizing the functional composed of the likelihood functional of equation (3.25) with the total variation functional results in the next explicit regularized iterative scheme:

$$\hat{f}_{TV}^{(k+1)}(x, y) = \left[\frac{g(x, y)}{\left(h * \hat{f}^{(k)} \right)(x, y)} * h^*(x, y) \right] \frac{\hat{f}^{(k)}(x, y)}{1 - \lambda \nabla \cdot \left(\frac{\nabla \hat{f}^{(k)}(x, y)}{|\nabla \hat{f}^{(k)}(x, y)|} \right)} \quad (3.30)$$

The usage of $|\nabla f(\mathbf{x})|$ imposes edge-preserving smoothing of the solution.

In the next chapters, we will discuss a new method to regularize the RL algorithm, based on multiresolution. In general, multiresolution-based methods allow us to make a better distinction between signal and noise than classical restoration methods.

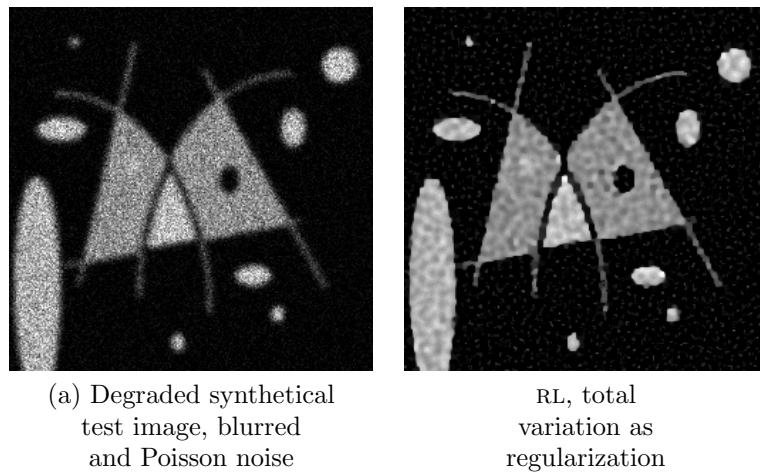


Figure 3.13: Degraded test image (Poisson noise) and restored with RL deconvolution, regularized with total variation.

Chapter 4

Multiresolution and image processing

When you study trees, do you want to study the forest, or the leaves of a single tree?

4.1 General introduction

During the last decade, multiresolution techniques (e.g., wavelets, steerable pyramids, ...) have acquired increasing popularity in the field of image processing. This Chapter doesn't intend to describe the full mathematics behind multiresolution, only the key ideas and their applications in image processing. More details about the theory and applications can be found in [Mallat, 1999, Strang and Nguyen, 1996, Daubechies, 1992].

In image processing, the concept of multiresolution is at the heart of many important techniques, e.g., scale-space, wavelets, steerable pyramids, ... The key idea is that each image feature exists on a certain range of scales in the image. A simple example is the hierarchy in a field of flowers. In this example, the field itself exists on the coarsest resolution scale, whereas the concept of "flower" only makes sense at a finer scale. At even finer scales, the flower is described in terms of its cells, and the molecules that make up the cells. Since a typical image contains objects of interest on various scales, it makes sense to process images at several resolution scales instead of just one. This is actually what multiresolution techniques do [Lindeberg, 1994].

In practice, a multiresolution representation of an image is obtained by applying increasing degrees of blur to the image, thus removing more and more details from the image. Each of the blurred images is subtracted from the previous, less blurred image. These difference images are called the details of

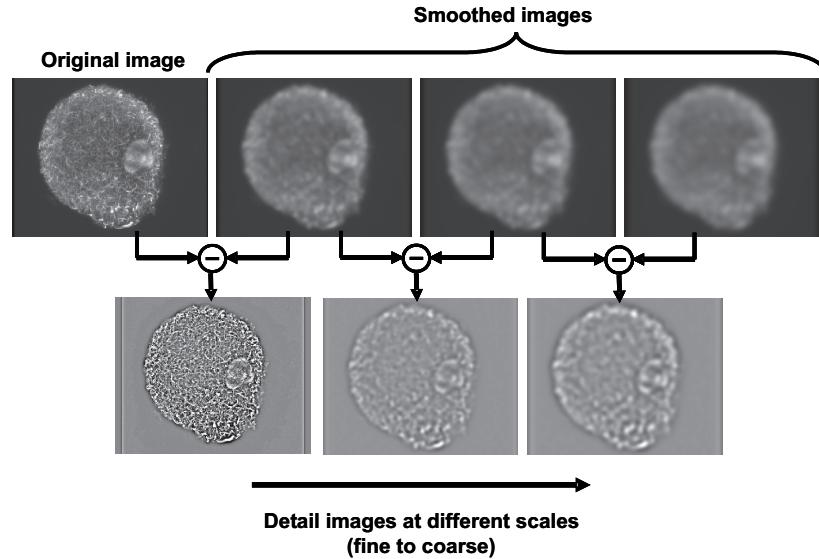


Figure 4.1: Multiscale decomposition example of a biological image (an ovary cell). The coarser the scale, the more details have been removed. The original image can be recovered by summing the coarsest lowpass image with all the detail images. This decomposition is very similar to the Laplacian pyramid [Burt and Adelson, 1983].

the corresponding resolution scale. This procedure is illustrated in Figure 4.1, where a confocal image of an ovary cell at different resolution scales is shown.

4.2 Multiresolution, wavelets and steerable pyramids

4.2.1 From the Fourier transform to wavelets

The Fourier transform is a valuable tool in signal processing, since many operations (like computation of the power spectrum, convolution, ...) rely on or are facilitated by this transform. The Fourier transform gives us insight in the frequency content of an image, i.e., an image can be described in terms of a sum of *sinusoidal basis functions*, and the Fourier transform shows us which basis functions are more prominent in this description. Smooth areas in images are best described in terms of low frequencies, while singularities like dots, edges, lines, corners, ... are described in terms of high frequencies.

However, from the Fourier transform, it is not obvious from which image features in the image the different frequencies originate, since the basis functions of the Fourier transform have infinite support. In Figure 4.2 e.g., adding a pair

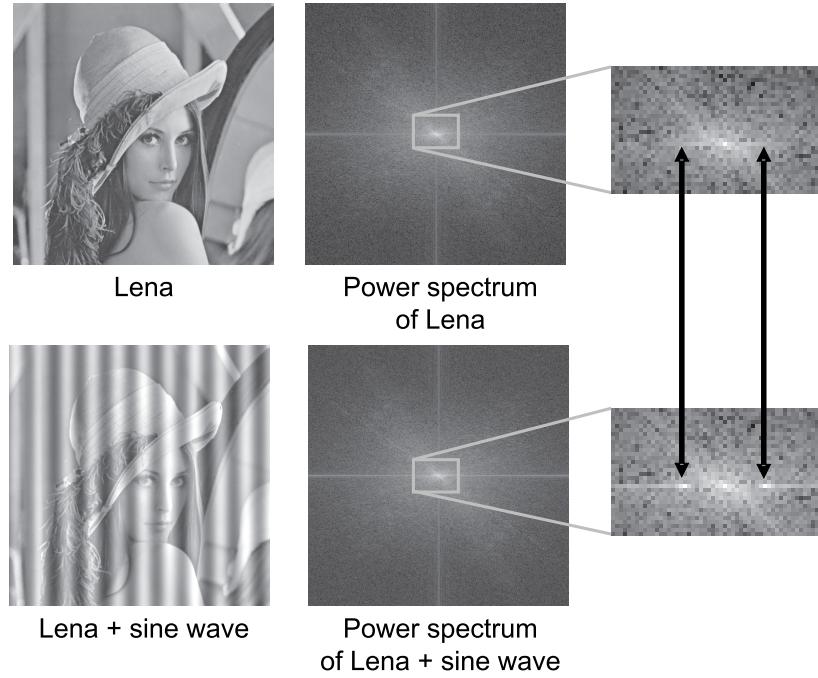


Figure 4.2: Illustration of the Fourier transform. The top row shows the Lena image (left) and its power spectrum (squared modulus of the Fourier transform coefficients). The center part is magnified on the right. The bottom row shows Lena, plus a sine wave. The power spectrum of this image shows two extra peaks near the center. The position of the two peaks (point-symmetric with respect to the origin) represents the frequency of the sine wave, while the intensity represents the amplitude of the sinusoidal wave. The phase information of the wave is included in the complex numbers (here not represented).

of complex conjugate Fourier coefficients (bright spots indicated with black arrows in the spectrum at the bottom left of Figure 4.2) to the Lena image: it results in adding a sinusoidal intensity pattern to the whole image.

Nevertheless, it is sometimes useful to know from which locations in the image certain Fourier frequencies originated. Also, manipulation of the Fourier coefficients has an impact on the whole image. In Figure 4.3 (a), we can see that the frequency-space resolution of the spatial representation is optimal in the spatial coordinates (i.e., without transformation), but the uncertainty in frequency is infinite. In the Fourier representation, it is exactly the other way around (Figure 4.3 (b)).

To obtain a representation in which both the spatial and the frequency content of a signal can be examined, the Short Term Fourier Transform (STFT) was proposed. Here, the signal is first multiplied with a window of a certain width, before the Fourier transform is applied (so we can consider the basis functions of

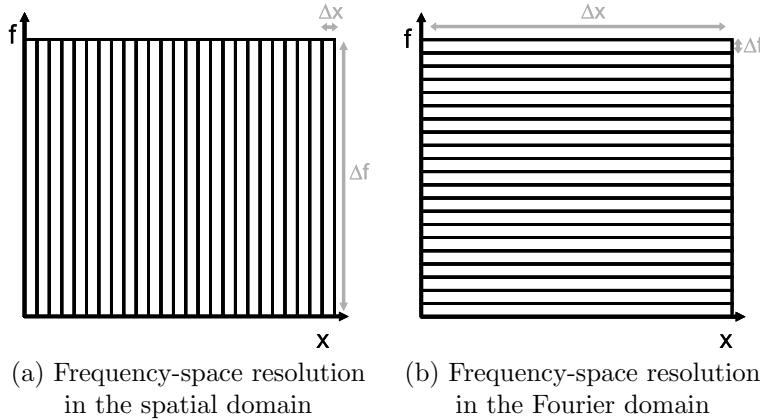


Figure 4.3: Comparison of frequency-space resolution in the spatial and the Fourier domain. The width of the boxes in the two diagrams (indicated with Δx) shows the uncertainty in space in both representations; the height of the boxes (indicated with Δf) shows the uncertainty in frequency in both representations. Note that in (a), one has no uncertainty about the spatial information, infinite uncertainty about the frequency content. In (b), it is exactly the opposite.

this representation to be sinusoidal waves multiplied with a window, usually a Gaussian window). This procedure is repeated for different translated versions of the window. This way, the *local* frequency content for every position of the window can be computed.

The following illustrative example in this context is loosely based on [Polikar, WWW]. In Figure 4.4, we show two functions which consist of three sinusoidal waves with different frequencies. One of them contains the three frequency components simultaneous in time (Figure 4.4 (a)), the other one separated in time (Figure 4.4 (c)). Their Fourier spectra show almost no differences (except for some smaller peaks due to the discontinuities between the segments containing the different frequencies in time).

In Figure 4.5, we show the STFT spectrum of Figure 4.4 (c). This illustration shows the local power spectrum of the signal on the vertical axis as a function of the position in the signal (horizontal axis, x). In Figure 4.5 (a), it clearly shows the three different frequency components separated in space. It shows that for a small window width, we obtain a good spatial resolution (the three different frequencies are clearly separated in space). The frequency resolution on the other hand is rather poor for this window width, indicated by the width of the white blocks in Figure 4.5 (a). For increasing window widths (shown in Figure 4.5 (b-d)), we see that the frequency resolution improves, at the expense of the spatial resolution, until in Figure 4.5 (d), the frequencies are clearly separated (narrow white lines) which almost lost all spatial resolution (they extend almost over the whole horizontal axis).

As we have seen in the previous paragraph, there is an intrinsic limitation to

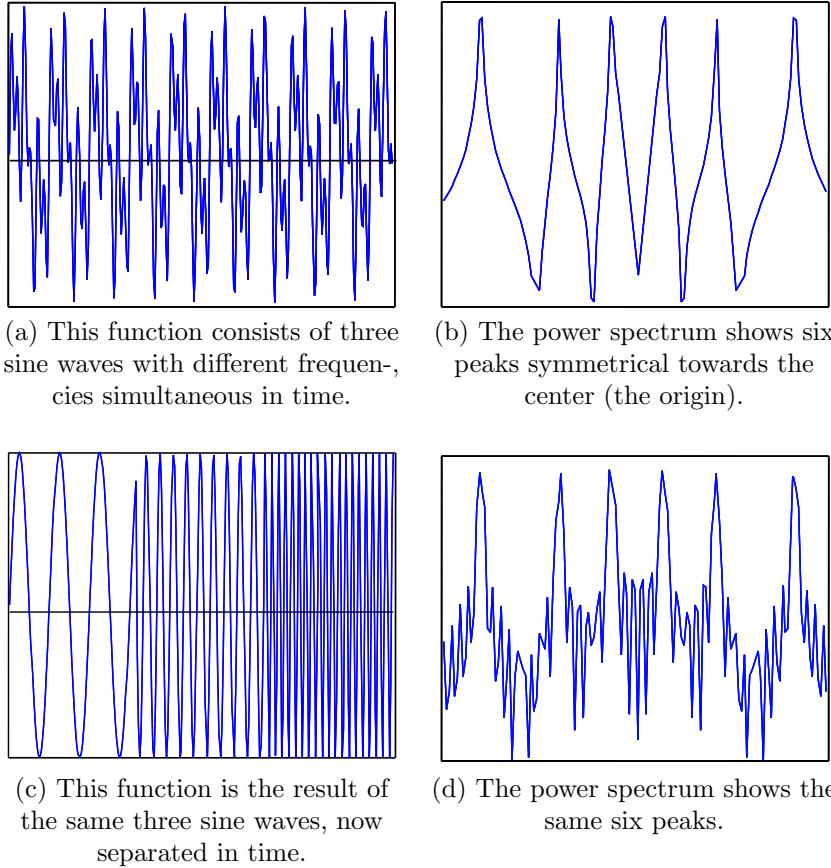


Figure 4.4: Fourier transform of two functions, each consisting of the same three sine waves with different frequencies. The first function is the sum of the three sine waves, the second function consists of these three sine waves, separated in time.

the joint resolution in space (Δx) and in frequency (Δf). This uncertainty is expressed in mathematical terms by the *Heisenberg Uncertainty Principle*. This principle originates from quantum mechanics, and originally stated that for moving particles, there is an implicit uncertainty on the joint information about their momentum (their mass, multiplied with their velocity vector) and their location:

$$\Delta x \Delta p \geq C$$

with C a fixed constant that depends on the detailed definition of Δx and Δp . Similar uncertainty applies to any pair of conjugate operators (generally even for any pair of Hermitian operators). The Uncertainty Principle in fact states

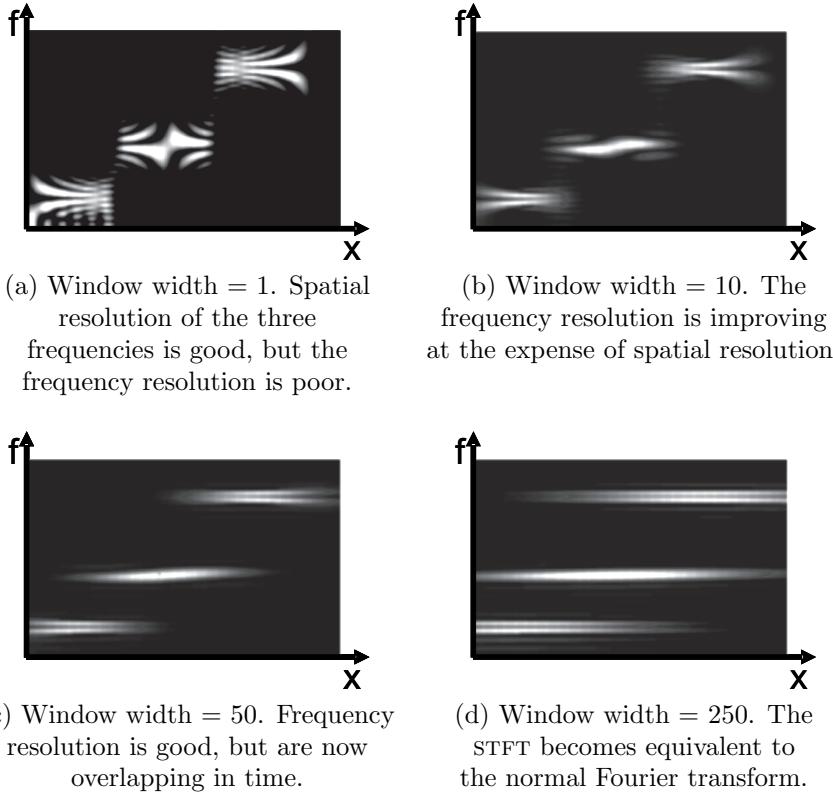


Figure 4.5: STFT the function in Figure 4.4 (c). From (a) to (d), the size of the window of the STFT is increased.

that a function and its Fourier transform cannot both have compact support [Weisstein, WWWd], so:

$$\Delta x \Delta f \geq C$$

In this context, it means that it is impossible to know the time-frequency representation of a signal with infinite detail, i.e., to know exactly which spectral components exist at a single moment in time (or space, in the case of an image). We can only know in which bands of frequencies a signal has a response in a given period of time (for images, the 1D time coordinate is replaced by 2D or more spatial coordinates). This is illustrated in Figure 4.6.

Many useful signals (e.g., electrocardiograms, lines of images, ...) have high frequency components that extend over small intervals and low frequency components that extend for long intervals. In order to localize signal discontinuities, very short basis functions (i.e., with a good spatial resolution) are needed. To

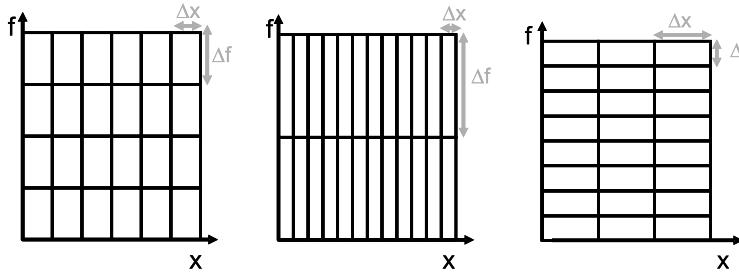


Figure 4.6: The space-frequency diagrams for the STFT with three different choices for the width of the window function. In each case, the joint resolution in frequency and space is represented by a rectangle. The surface of these rectangles are equal to a certain constant.

obtain a detailed frequency analysis, very long basis functions are needed. So in wavelet analysis, a trade-off is made: short, high-frequency basis functions are used together with long, low-frequency ones.

In image processing, this has following consequences: smooth features in an image (e.g., flat areas) are extended over larger areas in the image, and mainly produce responses in the low frequency subbands. In this case, to identify the frequency content with high frequency resolution is more important than to identify the precise location (after all, it is smooth and extended in space).

In case of sharp image transitions (like dots, lines, step edges, . . .), knowledge about the exact location of the edge is more important and a high spatial resolution is required. Therefore, high frequency basis functions with a short support are used, because they produce a well-localized response. The price we have to pay for this, is a higher uncertainty about the frequency content at this position.

Therefore, the space-frequency diagram is divided in the following way: for low frequency bands, rectangles with a good frequency resolution are chosen, while for high frequency bands rectangles with a good spatial resolution, are chosen. Specifically, if we choose the *dyadic* space-frequency representation (see Figure 4.7), the frequency resolution halves when we move up in the space-frequency plane, while the spatial resolution doubles for every layer (which are called “scales”).

We now translate these concepts in terms of basis functions. Recall that for the Fourier transform, these are infinite sinusoidal waves, while for the STFT, the basis functions are windowed sinusoidal waves. The basis functions of the wavelet decomposition are the so-called “wavelets”, i.e., *small waves*, or oscillating functions with a certain space-frequency resolution. An example is shown in Figure 4.8. To compute the wavelet decomposition of a signal, we can use the Continuous Wavelet Transform (CWT) or the Discrete Wavelet Transform (DWT).

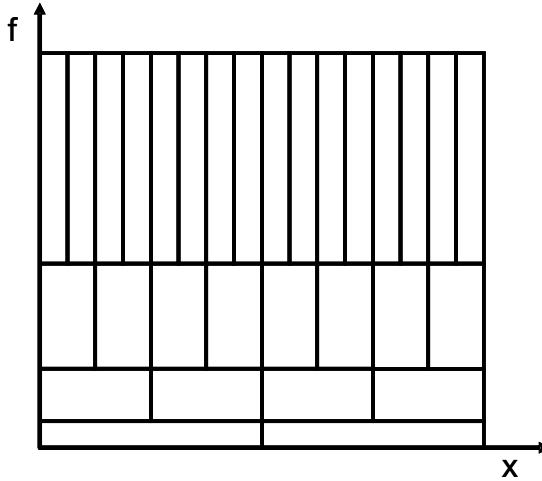


Figure 4.7: The space-frequency diagram in case of the wavelet decomposition, with the dyadic division of the space-frequency plane.

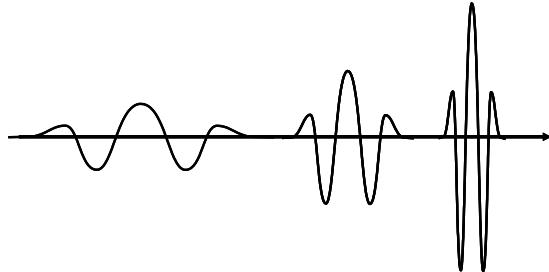


Figure 4.8: Example of a wavelet function, and scaled versions of this mother function.

The CWT computes the inner products of the signal $f(x)$ with shifted and scaled versions of the so-called *mother wavelet function* $\psi(x)$. Here, the term “mother” refers to the fact that the basis functions with different supports used in the decomposition are obtained from a single prototype function: the mother wavelet $\psi(x)$. Here, support means space-frequency support.

$$Wf(u, s) = \langle f\psi_{u,s} \rangle = \int_{-\infty}^{+\infty} f(x) \frac{1}{\sqrt{s}} \psi^* \left(\frac{x-u}{s} \right) dx$$

$Wf(u, s)$ is computed for a continuous range of the variable u (which determines the shift) and s (the variable which determines the scale, which is related to the frequency). This representation is highly redundant since it converts a one-dimensional signal into a two-dimensional function $Wf(u, s)$. In digital image processing, the Discrete Wavelet Transform (DWT) is more suited: it

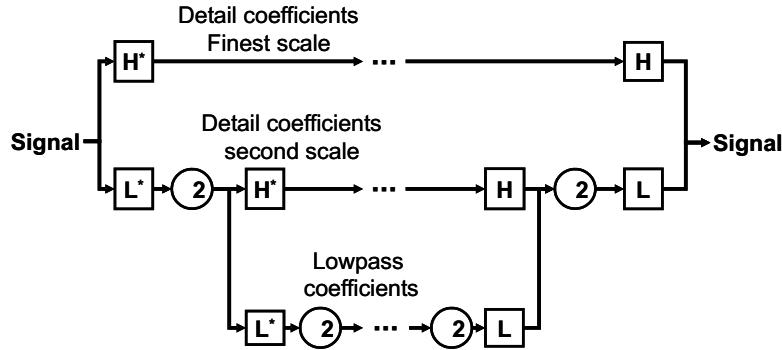


Figure 4.9: The filter bank scheme for a one-dimensional wavelet decomposition. H indicates a highpass filter, and L a lowpass filter (with X^* the filter kernel with the complex conjugate transfer function, i.e., the filter kernel which is a point reflected version of a filter X). The lowpass band is then recursively further subdivided by a highpass filter and a lowpass filter. $\downarrow 2$ means subsampling, which involves eliminating half of the samples per dimension. $\uparrow 2$ means upsampling.

only considers discrete values of s and u (with usually dyadic steps for s). It can be computed efficiently using a filter bank scheme (Figure 4.9 gives a filter bank scheme for a one-dimensional signal).

At this point, it is useful to mention the concept of subsampling¹: immediately after one step of the decomposition of a signal in different resolution scales, every scale contains the same number of coefficients as the number of pixels in the original signal. So, when a signal of N coefficients is decomposed in S resolution scales, the decomposition has NS coefficients. However, a great deal of this information is redundant, and can therefore be omitted. In classical wavelet decompositions, it is sufficient to keep only half the number of coefficients per scale per dimension². In classical orthogonal wavelet decompositions, this results in a decomposition with the same number of coefficients as there were coefficients in the original signal. This is called a critically sampled transform.

In the case of 2D images, there are two important variants (i.e., those with two or three detail images per scale). Both variants can involve subsampling (which results in a critically sampled transform in the case of classic wavelet decompositions with e.g. Daubechies wavelets) or not (which results in a redundant transform, usually computed with the à trous algorithm [Mallat, 1999]).

- A decomposition with two oriented detail bands per scale (horizontal and vertical details) is obtained by applying the one-dimensional scheme on the rows and columns separately (an example is given in Figure 4.10):

¹Sometimes, the terms downsampling and decimating are used.

²In the case of 2D image processing, it is sufficient to keep 1/4 of the coefficients of every subband per scale: keep only every second coefficient per row, and this for every second row.

- the horizontal detail band is obtained by filtering the image column by column with a 1D highpass filter;
 - the vertical detail band is obtained by filtering the image row by row with a 1D highpass filter;
 - the lowpass band is obtained by filtering row by row with a 1D lowpass filter, followed by filtering column by column with a 1D lowpass filter;
 - the lowpass band is subsampled with a factor two in every dimension, and the scheme is recursively applied on this lowpass band for the next scale.
- A decomposition with three oriented detail bands per scale (horizontal, vertical and diagonal details, see Figure 4.11 for an example):
 - the horizontal detail band is obtained by filtering the image column by column with a 1D highpass filter and then row by row with a 1D lowpass filter;
 - the vertical detail band is obtained by filtering the image row by row with a 1D highpass filter and filter the result column by column with a 1D lowpass filter;
 - the diagonal detail band is obtained by filtering the image row by row with a 1D highpass filter and filter the result column by column with a 1D highpass filter;
 - the lowpass image is again obtained by filtering the image row by row with a 1D lowpass filter, and to filter the result column by column with a 1D lowpass filter;
 - the lowpass band is subsampled with a factor two in every dimension, and the scheme is applied recursively on this lowpass image for the next scale.

	Vertical details	Horizontal details	Vertical details	Horizontal details	
Original image					Finest detail scale
					Detail scale 2
					Detail scale 3
					Lowpass band
	Wavelet subbands		Power spectra of the subbands		

Figure 4.10: An example of a wavelet decomposition of a synthetic test image: the left column shows the original image; the second column shows the different detail scales from top to bottom (both horizontal and vertical details). The right column shows the power spectra of the wavelet subbands.

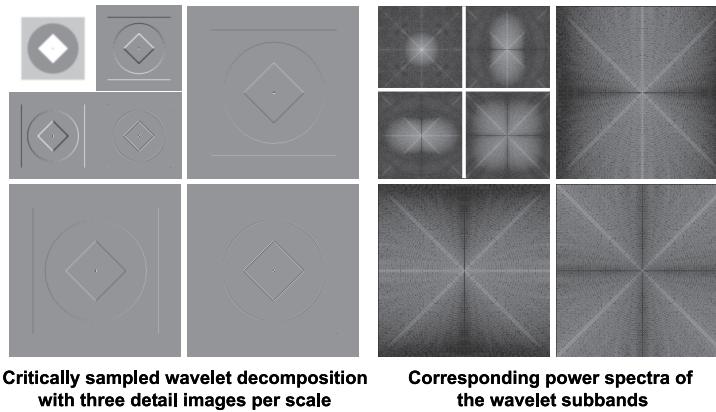


Figure 4.11: An example of a wavelet decomposition of the same synthetic test image, this time with three detail images per scale, and critically subsampled: the left part shows the wavelet decomposition; the right part shows the corresponding power spectra of the wavelet subbands.

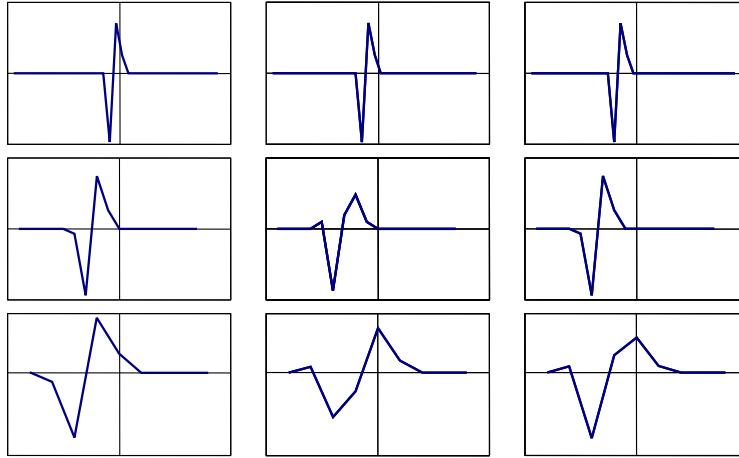


Figure 4.12: Wavelet decompositions of three step edges, each time shifted one pixel further.

4.2.2 The steerable pyramid.

4.2.2.1 General concept of steerable pyramids.

Classical wavelet decompositions have several disadvantages, such as limited orientation resolution. This means that edges under different orientations in an image aren't recognized as such, e.g., a classical separable wavelet decomposition doesn't distinguish edges at an angle of $\pi/4$ from edges at an angle of $3\pi/4$ according to their energy distribution over the differently oriented subbands. This is illustrated in Figures 4.10 and 4.11 (the power spectra of the horizontal and vertical orientation subbands in Figure 4.10 show a big overlap, which means that e.g., edges under $\pi/4$ and under $3\pi/4$ produce responses in the same subbands).

Another problem with some of these classical wavelet decomposition schemes is their shift-variant nature. This problem occurs in schemes where subsampling is involved ($\downarrow 2$ in Figure 4.9). Shift variance means that the wavelet coefficients corresponding with two edges at slightly different positions are completely different. In Figure 4.12, the wavelet decomposition of a step edge is shown. Each column shows the three scales of the wavelet decomposition of the same step edge, but each time shifted one pixel further. From this example, it is clear that the wavelet coefficients are not simply translated versions of one another in the three cases, especially at the coarser (more subsampled) scales.

To solve these problems, a variant on the classical wavelet transform was proposed, i.e., the steerable pyramid [Simoncelli et al., 1992]. We will now briefly explain some basic properties of this decomposition.

One often wants to analyze oriented image structures, like edges, at a certain

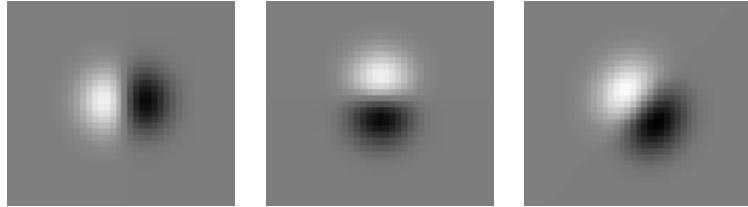


Figure 4.13: Examples of two steerable basis kernels, and a rotated version, interpolated from the two basis kernels.

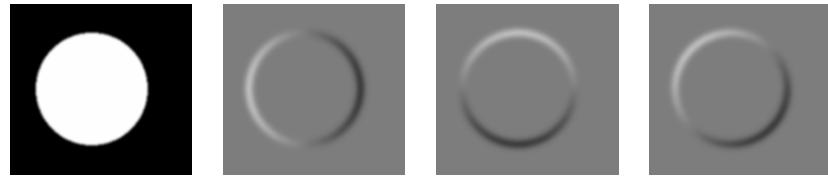


Figure 4.14: Example of an input image, filtered with the two basis kernels. The last image was interpolated from the two base images.

angle. In order to do so, one can convolve the image with a range of oriented filter kernels tuned to cover all orientations of interest in the image. However, this comes at a high computational cost. To avoid this, one can filter the image using a fixed set of basis kernels and interpolate for the other directions from the images filtered with the basis kernels [Freeman and Adelson, 1991, Simoncelli et al., 1992]. Such a set of basis kernels is called a steerable filter set.

4.2.2.2 An example of steerability.

An illustrative example of such a set of filter kernels are the first partial derivatives in x and y of the Gaussian function [Simoncelli et al., 1992], denoted respectively by $G'_x(x, y)$ and $G'_y(x, y)$ (see illustration in Figure 4.13):

$$G(x, y) = \frac{1}{2\pi\sigma^2} \exp\left(\frac{x^2 + y^2}{2\sigma^2}\right). \quad (4.1)$$

These two kernels $G'_x(x, y)$ and $G'_y(x, y)$ can be interpreted as filters for horizontal and vertical image features respectively. Let $I(x, y)$ be the original image and let the symbol $*$ denote the convolution operator. Then:

$$R_x(x, y) = (I * G'_x)(x, y) \quad (4.2)$$

is the image filtered for vertical features and

$$R_y(x, y) = (I * G'_y)(x, y) \quad (4.3)$$

for horizontal features (see Figure 4.14). If we want a rotated kernel $G'_\theta(x, y)$ tuned for features in the image under an angle θ , we just have to take a linear combination of the two basis kernels:

$$G'_\theta(x, y) = \cos(\theta)G'_x(x, y) + \sin(\theta)G'_y(x, y). \quad (4.4)$$

Because convolution is a linear operation, it is possible to calculate $R_\theta(x, y)$ by taking a linear combination of the result of the horizontal and vertical filters applied to the image:

$$R_\theta(x, y) = \cos(\theta)R_x(x, y) + \sin(\theta)R_y(x, y). \quad (4.5)$$

This is illustrated in the right image of Figure 4.14.

For a large number of orientations, this is computationally far less expensive (just two multiplications and one addition per pixel instead of a more complicated convolution) than a filtering with $G'_\theta(x, y)$ directly for an equal number of orientations. Steerable filters are not restricted to derivatives of Gaussians. In [Freeman and Adelson, 1991], the general conditions were given for a set of kernels to form a steerable basis set. In [Simoncelli et al., 1992], the steerable pyramid is described. This decomposition is a multiscale extension based on the concept steerable filters. It is similar to the classical wavelet decomposition, but has a much better orientation resolution and provides shift-invariance.

An example steerable pyramid decomposition of the *Zoneplate* test image (Figure 4.15) is shown in Figure 4.16. *Zoneplate* contains a rotationally symmetric sinusoidal wave, of which the frequency near the center is low, but increases towards the edges. The decomposition of this image results in responses of features with low spatial frequencies in the center, and responses of features with increasing frequencies towards the edges. The top row shows large coefficients at the borders, where the high spatial frequency response is strongest. The middle shows large coefficients away from the borders.

The bottom image is the low spatial frequency response of the image. In contrast with the transform we used, this decomposition is not subsampled (meaning that every subband has the same size as the original image). In practice however, subbands are subsampled, keeping for each dimension only half of the coefficients for the next lower resolution scale. Note also that image

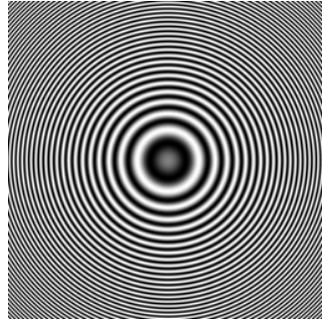


Figure 4.15: Zoneplate image.

features oriented under an angle $\pi/4$ and features oriented under an angle $3\pi/4$ produce responses in separate subbands.

4.2.2.3 Implementation issues of the steerable pyramid.

There are two methods to implement this decomposition: the spatial method, and the method in the Fourier domain. In essence, they achieve the same goal, but the practical implementation is different. The spatial method extracts the different subbands by convolving the image in the spatial domain with a set of filter kernels, designed for this purpose. The Fourier domain method implements these convolutions in the Fourier domain: the Fourier transform of the image is multiplied with the set of transfer functions, and each of these subbands is then after multiplication transformed back to the spatial domain.

Fourier domain implementation of the full steerable pyramid.

The difference between the fully steerable pyramid and the standard steerable pyramid is that in case of the standard steerable pyramid, the initial highpass transfer function $H_0(r)$ is not subdivided in oriented transfer functions. For the fully steerable pyramid, $H_0(r)$ is also subdivided into A_k 's ($k = 0, \dots, K - 1$, with K the number of orientations that is used).

Although this way of implementing the pyramid is more recent [Portilla et al., 2003], we describe this method first. To start the decomposition, the image is transformed to the Fourier domain and multiplied with a set of transfer functions (a transfer function is a mask to select certain spatial frequencies of the Fourier transform of an image). This defines different image subbands (a subband is an image which contains only spatial frequencies in a certain band of the Fourier domain, see also Figure 4.18). Each subband is then transformed back into the spatial domain.

We start by defining an initial highpass transfer function $H_0(r)$ (r and θ are polar coordinates in the Fourier domain), and a lowpass transfer function $L_0(r)$. This decomposition into $H_0(r)$ and $L_0(r)$ satisfies:

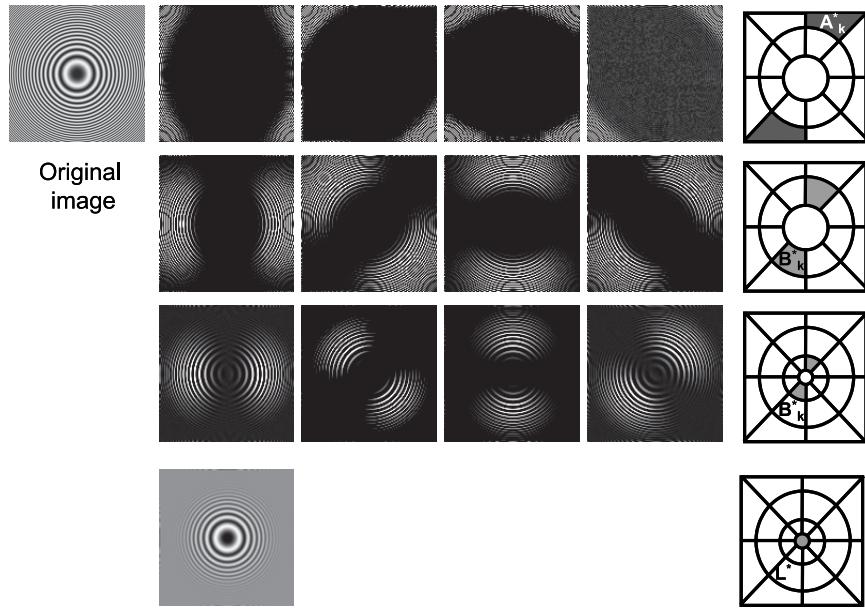


Figure 4.16: Steerable pyramid decomposition of Zoneplate image. For the notations, see Subsubsection 4.2.2.3. The last column shows a schematic representation of the corresponding transfer functions

$$|H_0(r)|^2 + |L_0(r)|^2 = 1$$

Next, $L_0(r)$ is subdivided in oriented transfer functions $B_k(r, \theta)$ (with $k = 0, \dots, K - 1$, with K the number of oriented subbands) and a residual lowpass transfer function $L(r)$ (these transfer functions are illustrated in Figure 4.21). This step of the decomposition satisfies:

$$|H_0(r)|^2 + |L_0(r)|^2 \left(|L(r)|^2 + \sum_{k=0}^{K-1} |B_k(r, \theta)|^2 \right) = 1$$

After every scale, subsampling is applied. The recursion relation that has to be satisfied (recursion in the dark rectangle in Figure 4.17) is:

$$\left| L\left(\frac{r}{2}\right) \right|^2 = \left| L\left(\frac{r}{2}\right) \right|^2 \left[|L(r)|^2 + \sum_{k=0}^{K-1} |B_k(r, \theta)|^2 \right] = 1$$

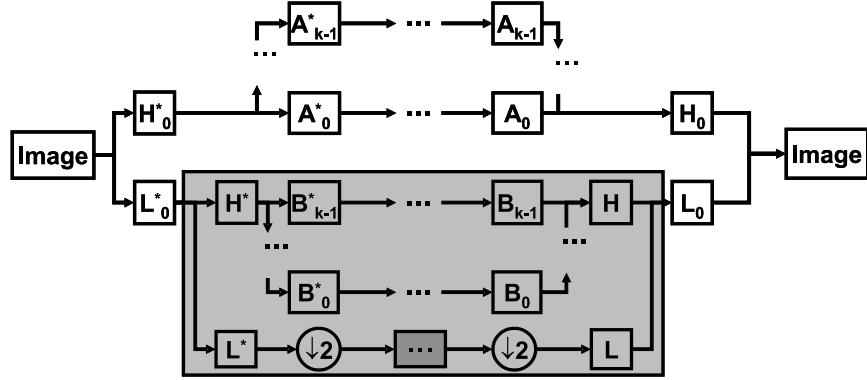


Figure 4.17: Here, the filter bank scheme of the decomposition is given. The light gray rectangle is the recursive part (applied in the darker gray rectangle, and recursively applied in the darkest smallest rectangle at the bottom). H is a highpass filter, L is a lowpass filter. The B_k 's are the oriented steerable filters for a resolution scale, the A_k 's are the oriented steerable filters for the finest resolution scale H_0 . $\downarrow 2$ means subsampling with a factor 2 in every dimension, and $\uparrow 2$ means upsampling.

Each of these transfer functions is multiplied with the Fourier transform of the image, and transformed back to the spatial domain, thus producing the different subbands of the image. In the example of the kernels based on derivatives of the Gaussian function, only two oriented subbands were used (which results in very limited orientation resolution). In our implementations, we normally use 4 oriented subbands per resolution scale (unless mentioned otherwise), which is sufficient to separate the useful features in the image from the noise³. For the next resolution scale, the subband resulting from transfer function $L(r)$ is subsampled and again subdivided in oriented subbands and a lowpass band. In Figure 4.17, the part in light gray is recursively applied in the dark rectangle of the bottom subband. A schematic illustration of the transfer functions is given in Figure 4.18 (the origin is at the center of each figure).

The explicit form of the transfer functions is given below:

$$L(r) = \begin{cases} \cos\left(\frac{\pi}{2} \log_2(4r/\pi)\right) & \frac{\pi}{4} < r < \frac{\pi}{2} \\ 1 & r \leq \frac{\pi}{4} \\ 0 & r \geq \frac{\pi}{2} \end{cases} \quad (4.6)$$

$$B_k(r, \theta) = H(r)G_k(\theta) \quad k = 0, \dots, K - 1 \quad (4.7)$$

³In [Portilla et al., 2003], it is shown that the increase in denoising quality for increasing the number of orientations from 4 to 8 is 0.3 dB) and still is computationally not too heavy.

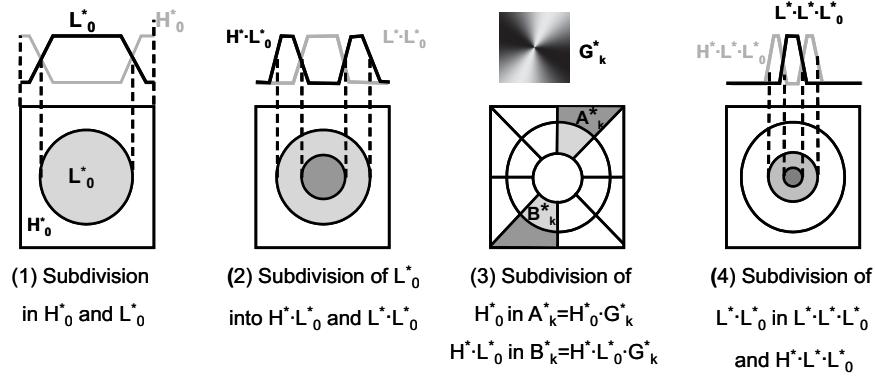


Figure 4.18: Here, a schematic representation of the transfer functions in the Fourier domain of the filters from Figure 4.17 decomposition of are shown.

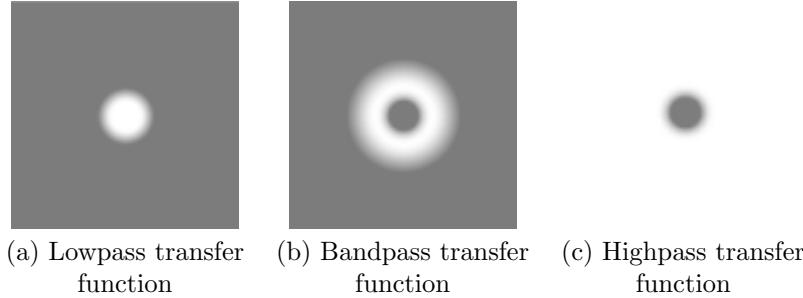


Figure 4.19: Design of a steerable pyramid: lowpass, highpass and bandpass transfer functions (black is negative, gray is zero and white is positive).

where

$$H(r) = \begin{cases} \cos\left(\frac{\pi}{2} \log_2(2r/\pi)\right) & \frac{\pi}{4} < r < \frac{\pi}{2} \\ 1 & r \geq \frac{\pi}{2} \\ 0 & r \leq \frac{\pi}{4} \end{cases} \quad (4.8)$$

and

$$G_k(\theta) = \frac{(K-1)!}{\sqrt{K[2(K-1)]!}} \left[2 \cos\left(\theta - \frac{\pi k}{K}\right) \right]^{K-1} \quad (4.9)$$

For the initial decomposition into lowpass subband L_0 and highpass subband H_0 , the formulae are given below:

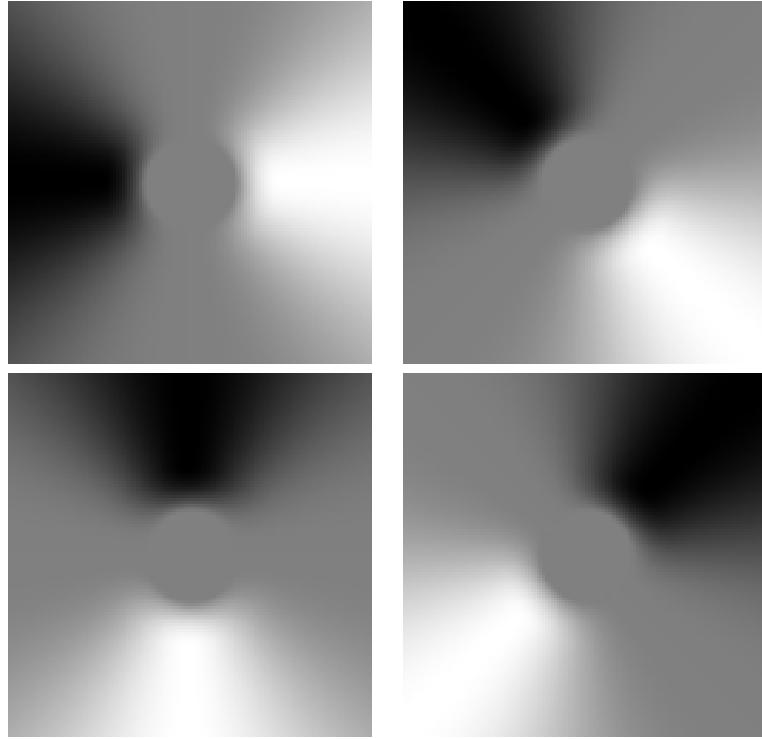


Figure 4.20: The four oriented transfer functions A_k . The squared sum of these transfer functions results back in Figure 4.19 (c) (black is negative, gray is zero and white is positive).

$$L_0(r) = L\left(\frac{r}{2}\right) \quad (4.10)$$

$$A_k(r, \theta) = H\left(\frac{r}{2}\right) G_k(\theta) \quad (4.11)$$

The transfer functions $H_0(r)$ and $L_0(r)$ in the initial step are given by equation (4.11), where the $A_k(r, \theta)$ form oriented subbands of $H_0(r)$ [Portilla et al., 2003]. The form of these subbands is illustrated in Figure 4.20. For the remainder of this thesis, when we use the term wavelet decomposition, we in fact use this special variant of the wavelet decomposition, i.e., the steerable pyramid. This exchange of terminology in case of the steerable pyramid has been used before in literature [Portilla et al., 2003].

Spatial implementation.

The spatial implementation of the steerable pyramid with quadrature fil-

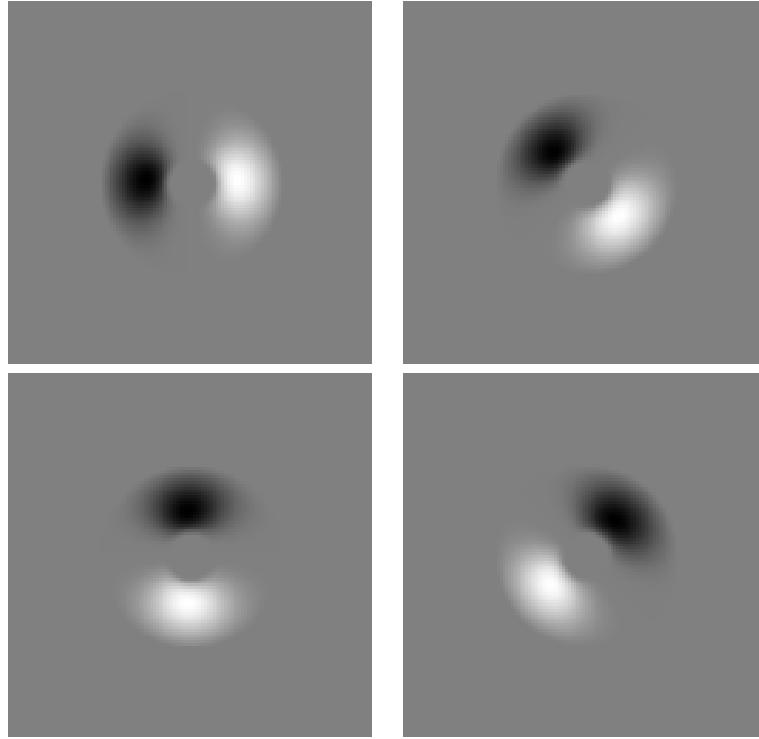


Figure 4.21: The four oriented transfer functions B_k . The squared sum of these transfer functions results back in Figure 4.19 (b) (black is negative, gray is zero and white is positive).

ters⁴ we use, is based on [Simoncelli et al., 1992] and the simplifications in [Castleman et al., 1998]. For this design, a set of filter kernels is computed as the inverse Fourier transform of a set of transfer functions in the Fourier domain. The definition of the transfer functions is equivalent to the transfer functions used in the Fourier-based implementation; that's why we discussed that more recent approach first.

The Fourier domain is again divided in an initial highpass band H_0 and a low-pass band L_0 . This lowpass band is then divided in a bandpass subband BP and a second lowpass subband L . For the next level, subband L is further subdivided in oriented bandpass subbands and a lowpass subband. The bandpass subband is then subdivided in oriented subbands B_k ($k = 0, \dots, K-1$), with K the number of oriented subbands per scale. In our case, $K = 4$. As explained before, this decomposition is completely analogue to what we described before with the Fourier-based implementation.

After the transfer functions are defined, the filter kernels are obtained by taking

⁴The concept of quadrature filters is explained and illustrated in Subsubsection 4.3.1.1.

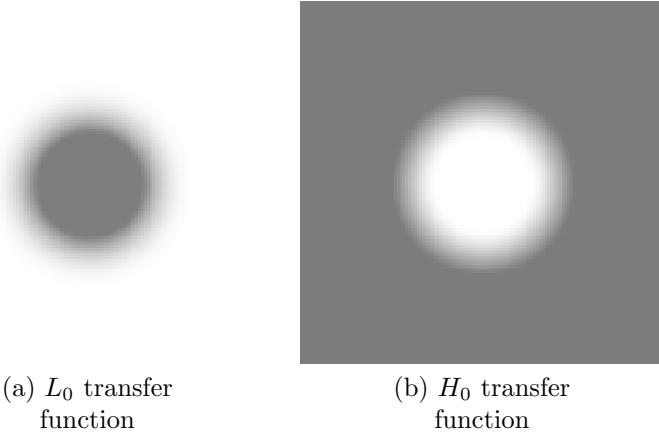


Figure 4.22: Design of a steerable pyramid: initial lowpass and highpass transfer functions (black is negative, gray is zero and white is positive).

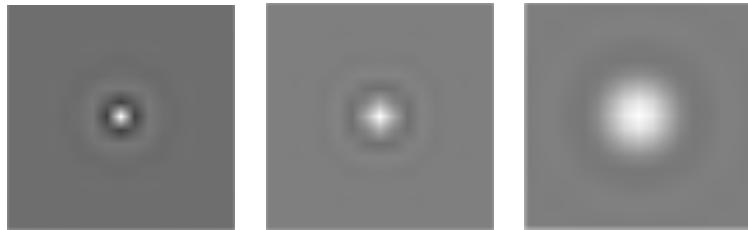


Figure 4.23: Filter kernels for the spatial implementation of the steerable pyramid: H_0 , L_0 and L_1 filter kernels. Again, we used black for negative, gray for zero and white for positive for these visualizations.

the inverse Fourier transform of these transfer functions. The image is then convolved with these transfer functions to obtain the different subbands. The resulting filter kernels are shown in Figure 4.23 and Figure 4.24. Again, we used black for negative, gray for zero and white for positive for these visualizations.

The pyramid with the quadrature filters was obtained in a similar way.

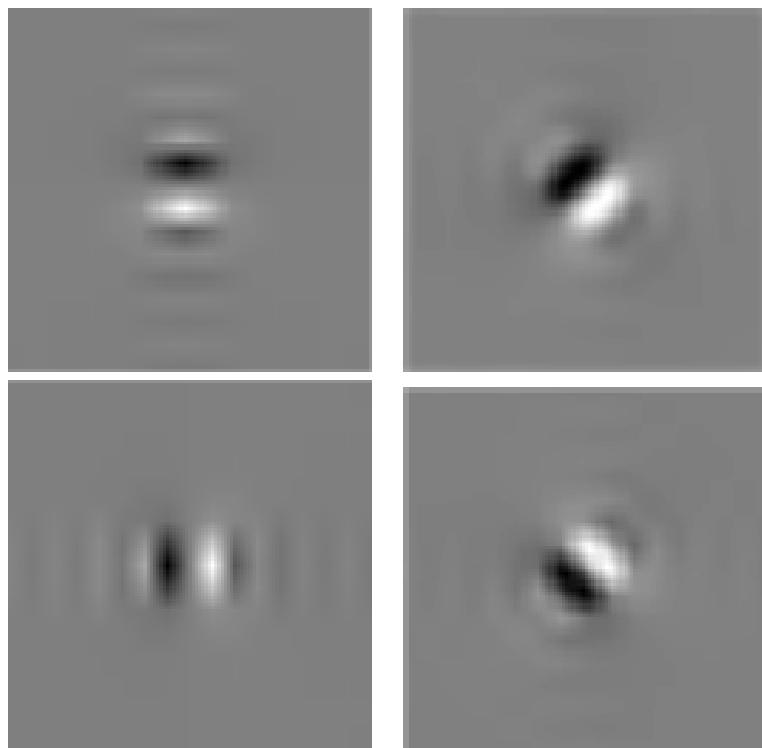


Figure 4.24: Filter kernels for the spatial implementation of the steerable pyramid: oriented bandpass kernels B_k , with $k = 0, \dots, 3$ ($K = 4$).

4.3 Applications

4.3.1 Image analysis

4.3.1.1 Quadrature filters and phase independence

The usage of quadrature filters goes back to Denis Gabor in [Gabor, 1946] with his definition of the “elementary signal”. Later, quadrature filters were used to calculate the so-called local energy and phase, two local image measures used in image analysis [Morrone and Owens, 1987, Venkatesh and Owens, 1990, Pudney et al., 1994, Granlund and Knutsson, 1995, Kovesi, 2002]. A set of filters is in quadrature if they are each others Hilbert transform, which is a transformation of a signal into one with the same Fourier spectrum, but with all phases shifted over $\pi/2$. The result is that all even components are transformed into odd components, and vice versa.

In an image, a lot of structures with different phases are present [Granlund and Knutsson, 1995]. The phase of an image feature represents the evenness or oddness of this feature (see Figure 4.25): even image features are symmetrical with respect to a central axis. An example in 1D is given in Figure 4.26 (a): the parts left and right of the vertical axis are symmetrical. In Figure 4.26 (b), an example of a 2D even feature is given.

In Figure 4.26 (c), an odd feature is shown: the part left of the vertical axis is point reflected with respect to the part right of the vertical axis. In Figure 4.26 (d), an example of a 2D odd feature is shown: the feature is even in one direction, but is odd in the orthogonal direction. In images, one can also have features that have a phase between a line and a step, as is illustrated in Figure 4.27. In fact, each image feature can be expressed as the sum of an even and an odd component.

We now discuss the output of different filters to analyze even and odd structures in a signal, and illustrate this with a one-dimensional synthetic example that contains a delta peak, a step edge and a step edge plus delta peak (Figure 4.28). If we use an odd filter, the response of the delta peak will be an odd structure, with a zero crossing at the delta peak. The response of the step edge will be a maximum. This filter is sometimes also referred to as an *edge detector*. If we filter the signal with an even filter, we have a maximum at the delta peak, while we will find a zero crossing at the step edge. Sometimes, an even filter is called a *line detector*. In both cases, the response at the combined edge is neither a zero crossing or a maximum.

However, if we take the sum of the squared responses of both filters, we find a maximum at all features of interest in the signal. The sum of the responses of an even and an odd filter is called the local energy of a signal at that position if both filters are in quadrature. So, filtering the image with a quadrature set of filters removes the phase dependency of the response of the analysis

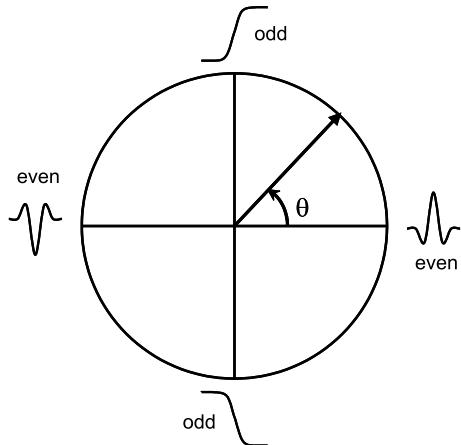


Figure 4.25: Illustration of local phase [Granlund and Knutsson, 1995].

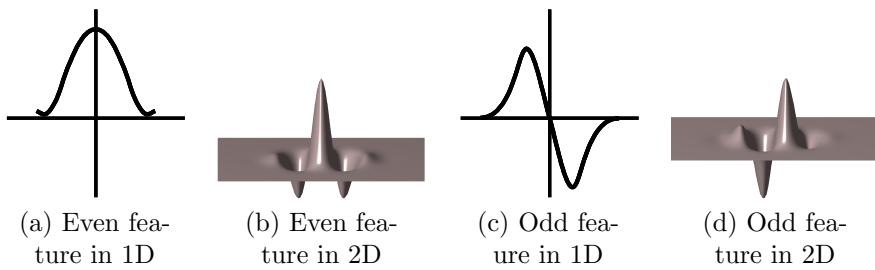


Figure 4.26: Even and odd features in 1D and 2D.

procedure. An extended discussion of using sets of quadrature filters is given in [Perona and Malik, 1991] and [Granlund and Knutsson, 1995].

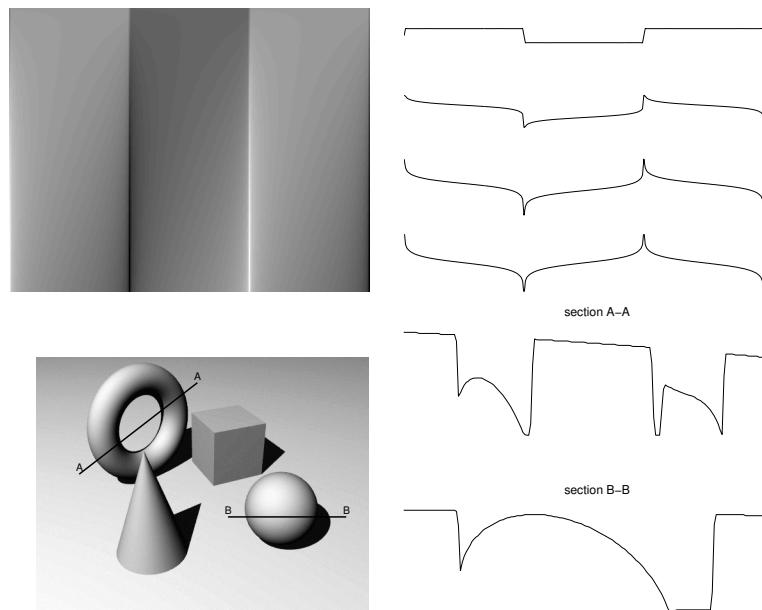


Figure 4.27: A line, a step edge and structures in between, with intensity profiles of some cross sections. Next is an example of a realistic scene, with some intensity profiles. Images were taken from [Kovesi, 2002] with permission of the author.

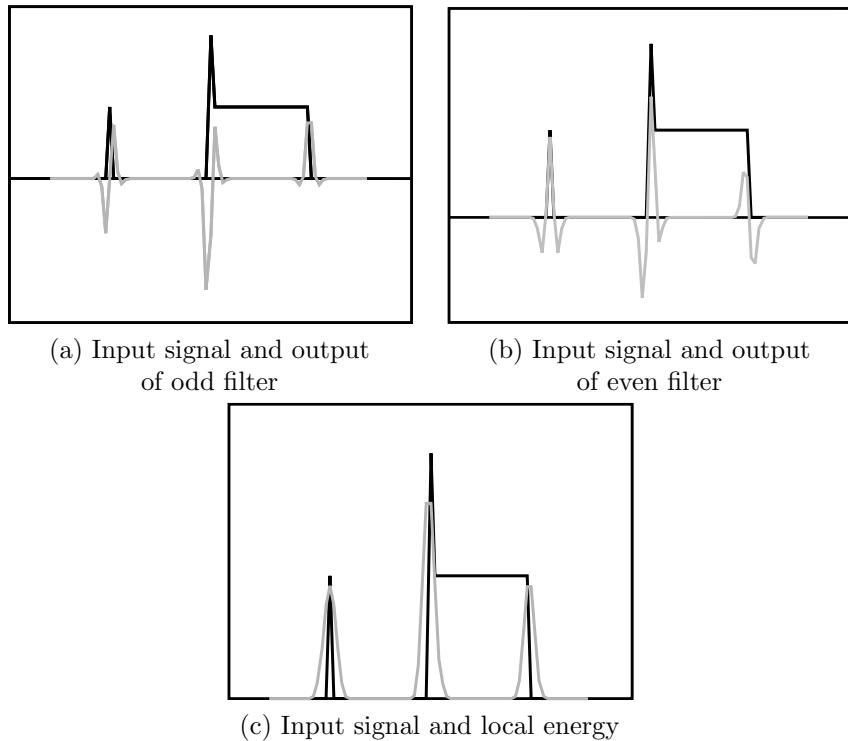


Figure 4.28: Even/odd filter responses (a,b) vs. local energy (c) of different 1D features.

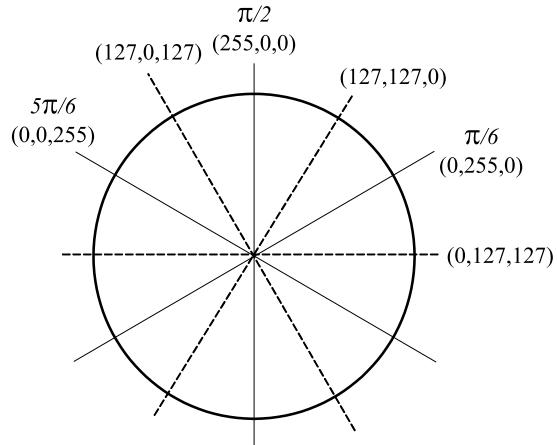


Figure 4.29: Color representation for edge orientations in 8 bit (R,G,B): red represents vertical edges ($\pi/2 \rightarrow (255, 0, 0)$), edges under $\pi/6$ are colored green ($\pi/6 \rightarrow (0, 255, 0)$) and blue represents edges under $5\pi/6$ ($5\pi/6 \rightarrow (0, 0, 255)$). Colors for edge orientations in between are interpolated in RGB-space between these main three directions.

4.3.1.2 Orientation analysis and edge detection

With a quadrature set of filters (the concept of quadrature filters is discussed in Section 4.3.1.1), it is possible to calculate the dominant orientation and local energy at a certain position in an image from the phase of the first harmonic [Freeman and Adelson, 1991], a measure closely related to the energy distribution in the local Fourier transform. If only a certain orientation is present in a certain neighborhood of the image, most of the energy of the Fourier transform of this neighborhood will be concentrated in a narrow sector oriented with the same direction as the local image gradients. The narrower this sector is, the

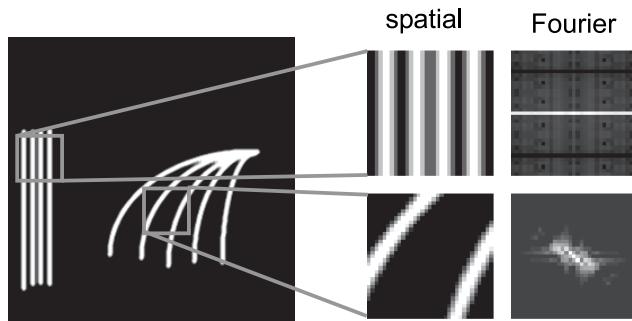


Figure 4.30: The local Fourier spectrum for two neighborhoods with oriented features in an image.

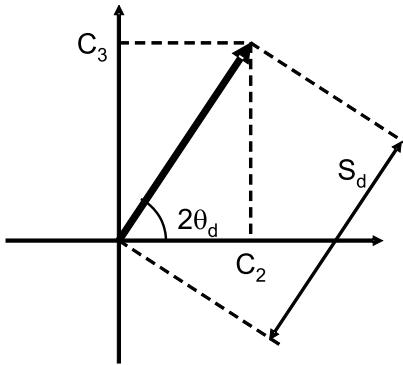


Figure 4.31: Coordinate vector (C_2, C_3) .

more prominent just a single orientation is present in this neighborhood. This is illustrated in Figure 4.30: the top cutout shows a neighborhood with just a single orientation, which leads to a concentration of almost all the energy of the power spectrum into a single line. The bottom cutout shows features with a varying orientations, which leads to energy of the power spectrum of this neighborhood spread into a broader sector.

In the method described by Freeman [Freeman and Adelson, 1991], the oriented energy is calculated from the output of two basic sets of steerable filters in quadrature. This oriented energy $E_\theta(x, y)$ under angle θ at position (x, y) is calculated from:

$$E_\theta(x, y) = [G_\theta(x, y)]^2 + [H_\theta(x, y)]^2$$

where $G_\theta(x, y)$ is the response of a bandpass filter steered under an angle θ at a certain scale, and $H_\theta(x, y)$ is its quadrature counterpart. Using the steering equations (to interpolate between the basis set of K oriented subbands $G_j(x, y)$ with $j = 1, \dots, K$ and the basis set of K' quadrature oriented subbands $H_k(x, y)$, with $k = 1, \dots, K'$), this expression can be re-written (using some trigonometric identities, or using Maple), and be reorganized in the form:

$$E_n(\theta) = C_1(x, y) + C_2(x, y) \cos(2\theta) + C_3(x, y) \sin(2\theta) + HOT$$

where the $C_i(x, y)$'s ($i = 1, 2, 3, \dots$) are functions of the $G_j(x, y)$'s and the $H_k(x, y)$'s. *HOT* stands for “higher order terms.”

Then, for every point in the image, the first harmonic of the local energy in the image is a vector with coordinates $(C_2(x, y), C_3(x, y))$ (Figure 4.31). Half of the angle of this vector (Figure 4.31) corresponds with the angle of the dominant orientation $\theta_d(x, y)$ at position (x, y) in the image:

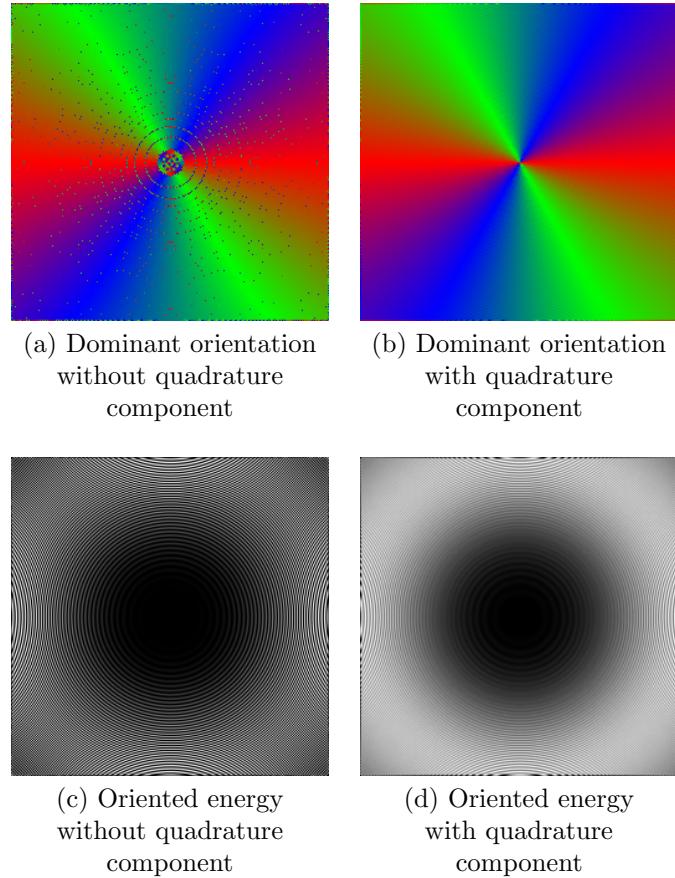


Figure 4.32: Application of quadrature filters in orientation analysis for the *Zone-plate* image. Colors are represented according to Figure 4.29.

$$\theta_d(x, y) = \frac{\arctan_2[C_3(x, y), C_2(x, y)]}{2} \quad (4.12)$$

with \arctan_2 the arctangent function that also determines the right quadrant, and the magnitude of this vector (Figure 4.31) with the orientation strength $S_d(x, y)$ at position (x, y) :

$$S_d(x, y) = \sqrt{C_2^2(x, y) + C_3^2(x, y)} \quad (4.13)$$

In Figure 4.32, the advantage of using quadrature sets of filters is illustrated:

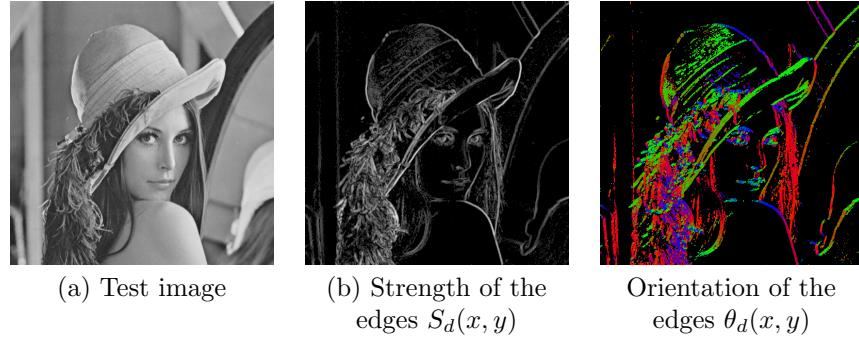


Figure 4.33: Application of using quadrature filters in orientation analysis for *Lena* image. The used color code in (c) for the orientations is the same as explained in the caption of Figure 4.32.

when only one set of filters is used, the results depend on the phase of the local feature in the image (in this case the phase of the radial sine wave present in the image). This is illustrated in the left column of Figure 4.32 (images (a) and (c)), where the orientation analysis and local energy without quadrature counterparts is performed. We clearly see some errors in the estimated local orientation and local energy which depend on the phase of the sine wave. In the right column (images (b) and (d)), the filters and their quadrature counterparts were used. Here, we see that both the orientation analysis and the local energy are now completely independent from the phase of the features in the image.

In Figure 4.33, we show an illustration of orientation analysis with a real image, i.e. *Lena*. In Figure 4.33 (b), $S_d(x, y)$ as in equation (4.13) is visualized: the higher the intensity, the higher the oriented energy. In Figure 4.33 (c), $\theta_d(x, y)$ as in equation (4.12) is visualized, using the same color codes for the gradient directions. A threshold based on the strength of $S_d(x, y)$ was applied in this Figure to highlight the gradient directions in the neighborhood of edges.

4.3.1.3 Corner detection

An original idea for corner detection is suggested in this part. As observed in [Simoncelli, 1999], wavelet coefficients are far from independent from each other. After the wavelet decomposition, wavelet coefficients still show correlations across scales, across orientations and over different positions. This is especially true for corners and for dots, which produce a response in subbands not adjacent in orientation.

Using wavelet products across scales was already suggested by [Olivo-Marin, 2002, Murtagh and Starck, 2003]. In [Olivo-Marin, 2002], multiscale products were used to detect intensity peaks to extract spots from biological fluorescence microscope images. In [Murtagh and Starck, 2003], a variant of this approach was used to detect edges in images. Both

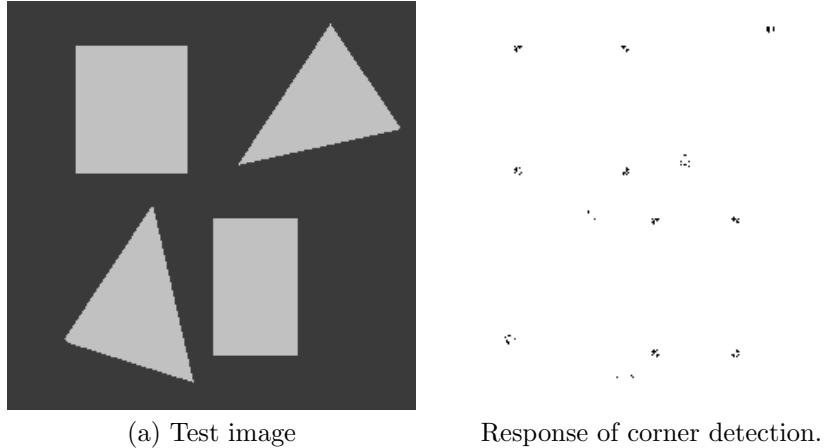


Figure 4.34: Example of corner detection in a synthetic test image using the proposed method. The resulted image is shown in negative (black is high response).

[Olivo-Marin, 2002, Murtagh and Starck, 2003] use the à trous wavelet transform.

Here, we suggest to use products of the coefficients in the differently oriented subbands per resolution scale, since corners have response in the differently oriented subbands, while normal edges don't. In Figure 4.34, we show a preliminary result of corner detection in a synthetic test image. Here, we just multiply all the oriented subbands, for three different resolution scales of the steerable pyramid. So we propose $P(\text{corner}(x, y))$ as probability of “corner-ness” at a position (x, y) :

$$P(\text{corner}(x, y)) = \prod_{j=1}^{\text{Maxscale}} \prod_{i=1}^K B_{i,j}(x, y)$$

Further research will be necessary to develop more complex and robust detection rules.

4.3.1.4 Image segmentation

For the sake of completeness, we mention that multiscale techniques were also successfully applied in image segmentation, like in [Bouman and Shapiro, 1994]. There, properties of a multiscale decomposition are incorporated directly in a segmentation method. In particular in Bouman's paper, a MAP estimator is defined in order to minimize the probability to misclassify a given pixel, based on properties derived from a multiscale MRF model. The source code for [Bouman and Shapiro, 1994] is available from <http://dynamo.ecn.purdue.edu/~bouman/>.

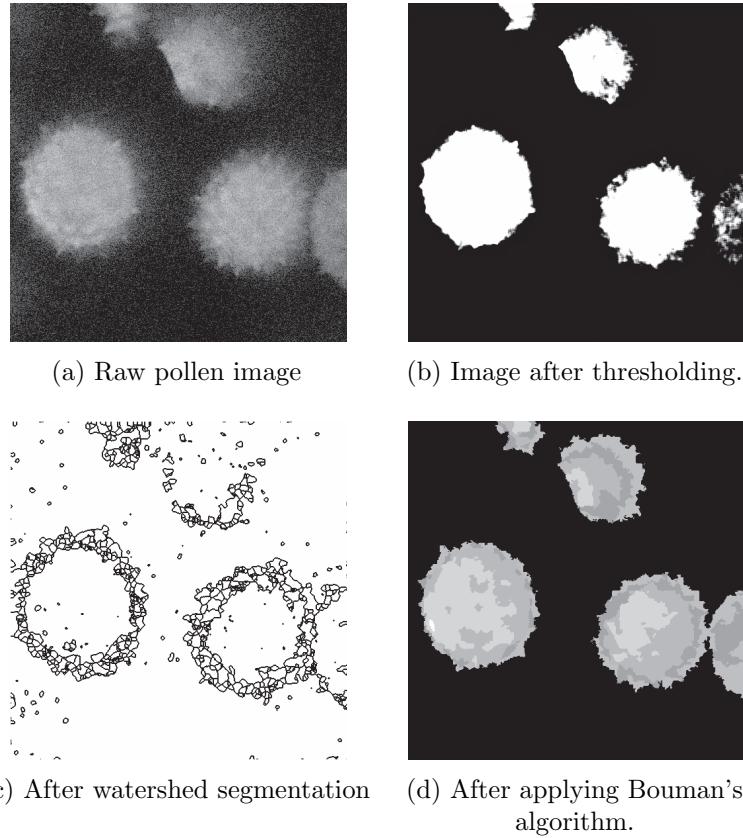


Figure 4.35: Illustration of some classical image analysis techniques compared to a multiscale segmentation technique [Bouman and Shapiro, 1994].

In Figure 4.35, we show the result of this segmentation technique applied to a confocal microscope image representing pollen of plants. Figure 4.35 (a) shows the raw image. This raw image clearly contains a lot of noise, which hinders post-processing like image analysis and image segmentation. A very basic technique often used in biological image analysis, is thresholding. In Figure 4.35 (b), it is clear that this basic technique is too simple, and doesn't really provide satisfactory results: the size of areas are under- or overestimated, connected regions are broken up, ...

In Figure 4.35 (c), the result of a classic segmentation technique (watershed, [Vincent and Soille, 1991]) is shown. This technique is based on taking the image gradient, and considering this gradient image as a topographic surface, with both valleys and mountains. Assumed that there is a hole in each of the local minima, we let water enter through those holes and flood the topographical image gradient surface, in which lakes form. When the water level rises,

water coming from different minima will meet. To avoid this, dams are defined whenever there would be a merge of the water. When the water level reaches a certain point, the only thing visible of the surface are the dams. These dam walls are called the watershed lines. However, due to noise and other local irregularities in the gradient image, the result of watershed segmentation is often oversegmented (i.e., produces too many small, insignificant segments), as is the case in Figure 4.35 (c).

Finally in (d), the result of the technique described in [Bouman and Shapiro, 1994] is given. Clearly, this is the best segmentation result, even with that much noise present. The pixels are classified in five different categories (indicated by the different gray levels) which form connected regions; this in contrast with the other techniques we discussed, which produced unconnected regions and small segments.

4.3.2 Noise reduction

4.3.2.1 Gaussian noise

To see how multiscale decompositions such as the steerable pyramid can facilitate the separation of the noise and image, we show an illustrative example in Figure 4.36. One can see that edges in the image result in large coefficients in all scales for both images. Bright means large amplitude and positive; dark means large amplitude and negative; gray means close to zero.

Those large coefficients (which are due to the useful image) are correlated both spatially in connected clusters, and across scales. Noise on the other hand results in small coefficients that are much less correlated. All wavelet-based noise reduction algorithms basically exploit these two properties to separate the noise from the useful features in an image.

Noise reduction algorithms implemented in the wavelet domain have proven to be superior to classical noise reduction algorithms. These algorithms basically follow the same strategy:

- apply the forward wavelet transform;
- threshold or shrink some of the wavelet coefficients according to some rules or criteria;
- apply the inverse wavelet transform.

In [Donoho and Johnstone, 1994, Donoho and Johnstone, 1995], a statistical method was developed (called SURE, Stein Unbiased Risk Estimator) to calculate an optimal global threshold for all the wavelet coefficients within a subband corrupted with white Gaussian noise, and to set all coefficients below this threshold to zero. For the larger ones, the value of the threshold is subtracted

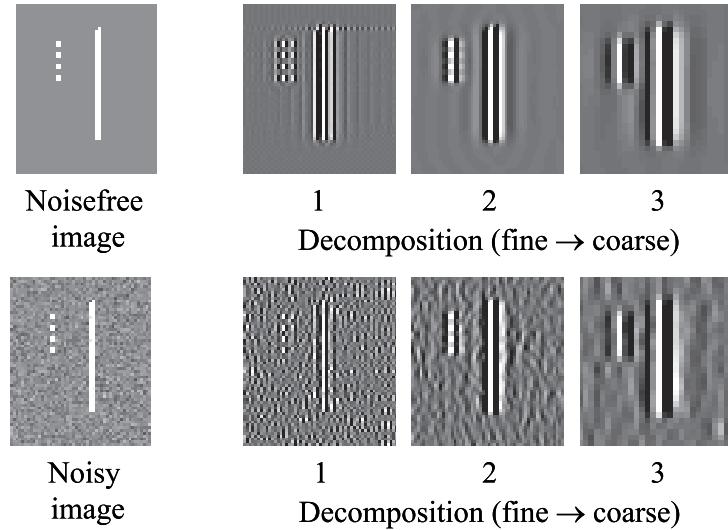


Figure 4.36: Comparison of decomposition of noise-free and noisy structure in the steerable pyramid domain. The first column is the image, the other columns show the vertical detail subbands of the decomposition. We can see in both cases that useful signal propagates well across scales (i.e., produces large coefficients at the same positions) in contrast with noise, which is much less correlated across scales.

(soft thresholding). The practical description of estimating this SURE threshold is loosely based on [Mallat, 1999].

The noisy data are represented in the spatial domain by:

$$g(x, y) = f(x, y) + n(x, y),$$

where $g(x, y)$ are the observed noisy data, $f(x, y)$ are the ideal, clean data and $n(x, y)$ is the noise. Let us denote the wavelet transform of these data in a certain scale s by:

$$v_s(x', y') = w_s(x', y') + n_s(x', y'),$$

with x' and y' the coordinates in the wavelet domain, and n_s the noise represented in the wavelet basis at scale s .

If $\tilde{f}(x, y) = D(g(x, y))$ is an estimation of the clean data, with $D(\cdot)$ the operator used to reduce the noise, the risk $r(D, f)$ of the estimation $\tilde{f}(x, y)$ is given by the expectancy value:

$$r(D, f) = E \left(\|D(g(x, y)) - f(x, y)\|^2 \right)$$

In soft thresholding, a coefficient with magnitude below a certain threshold T is set to zero. The corresponding risk in this case is $|f(x, y)|^2$, which can be approximated by $|g(x, y)|^2 - \sigma_n^2$. If the magnitude of the coefficient is larger than T , the value T is subtracted from the magnitude of the coefficient. In this case, the expected risk is the sum of the noise energy plus the bias introduced by reducing the magnitude of $g(x, y)$ by T . Then, the risk in this case is estimated by $\sigma_n^2 + T^2$.

To find the threshold T that minimizes this risk, the N wavelet coefficients $v(x, y)$ of a subband are sorted according to decreasing magnitude in $O(N \log N)$ operations. After sorting, let $v(k)$ be the wavelet coefficient⁵ of rank k . Let l be the index so that:

$$|v(l)| \geq T > |v(l+1)|$$

Then the risk to be minimized is:

$$r(f, T) = l(\sigma_n^2 + T^2) - (N - l)\sigma_n^2 + \sum_{k=l}^N |v(k)|^2 \quad (4.14)$$

where the first two terms of the right hand side correspond to the risk of the $N - l$ coefficients with magnitude below the threshold, and the last term corresponds to the risk of the clipped coefficients with magnitude above the threshold. To minimize $r(f, T)$, we set $T = |v(l)|$. To find T_{min} that minimizes equation (4.14), the results of evaluating equation (4.14) for the N values $v(k)_{1 \leq k < N}$ are compared.

As stated in [Mallat, 1999], this risk estimator is unbiased, but its variance could introduce errors that may lead to underestimations of T , especially when the signal energy is small compared to the noise energy:

$$\|f(x, y)\|^2 \ll \|n(x, y)\|^2 = N\sigma_n^2$$

When a minimum energy ϵ_N is defined as:

⁵Note that k doesn't refer to oriented subbands anymore.

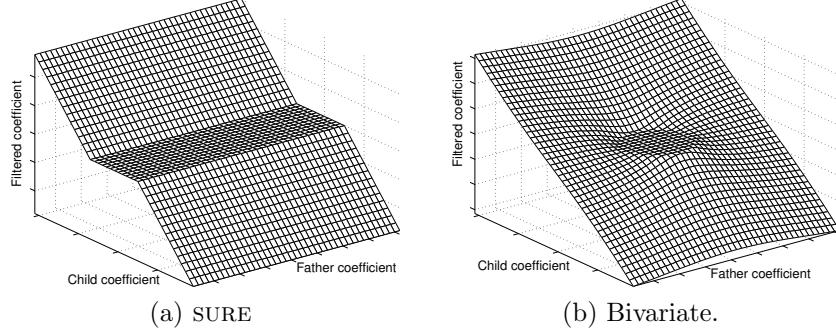


Figure 4.37: Illustration of the SURE versus the bivariate shrinkage function.

$$\epsilon_N = \sigma_n^2 N^{1/2} (\ln N)^{3/2}$$

the resulting threshold T_{SURE} becomes

$$T_{SURE} = \begin{cases} \sigma_n \sqrt{2 \ln N} & \text{if } \|f(x, y)\|^2 = \|g(x, y)\|^2 - N\sigma_n^2 \leq \epsilon_N \\ T_{min} & \text{if } \|f(x, y)\|^2 = \|g(x, y)\|^2 - N\sigma_n^2 > \epsilon_N \end{cases}$$

In our work, we have chosen to apply another wavelet shrinkage method developed for Gaussian data by [Sendur and Selesnick, 2002], where the following joint PDF for a wavelet coefficient $w_1(x)$ and its parent $w_2(x)$ (which is the coefficient at the same spatial position as $w_1(x)$, but in the next coarser resolution scale) is suggested:

$$p_{\mathbf{w}}(\mathbf{w}) = \frac{3}{2\pi\sigma^2} \exp\left(-\frac{\sqrt{3}}{\sigma} \sqrt{w_1^2(x) + w_2^2(x)}\right),$$

with $\mathbf{w} = (w_1(x), w_2(x))$. For the coefficients due to noise (assumed to be independent identically distributed Gaussian), the following joint PDF is assumed to be:

$$p_{\mathbf{n}}(\mathbf{n}) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{n_1^2(x) + n_2^2(x)}{2\sigma_n^2}\right).$$

with $\mathbf{n} = (n_1(x), n_2(x))$. The classical MAP estimator is given by:

$$\begin{aligned}\hat{\mathbf{w}}(\mathbf{v}) &= \arg \max_{\mathbf{w}} [p_{\mathbf{w}|\mathbf{v}}(\mathbf{w}|\mathbf{v})] \\ &= \arg \max_{\mathbf{w}} [p_{\mathbf{v}|\mathbf{w}}(\mathbf{v}|\mathbf{w})p_{\mathbf{w}}(\mathbf{w})] \\ &= \arg \max_{\mathbf{w}} [p_{\mathbf{n}}(\mathbf{v} - \mathbf{w})p_{\mathbf{w}}(\mathbf{w})]\end{aligned}$$

with $\mathbf{v} = (v_1(x), v_2(x))$. This MAP estimator combined with the expressions for $p_{\mathbf{w}}(\mathbf{w})$ and $p_{\mathbf{n}}(\mathbf{n})$ results then in the following shrinkage rule:

$$\hat{w}_1(x) = \frac{\left(\sqrt{v_1^2(x) + v_2^2(x)} - \frac{\sqrt{3}\sigma_n^2}{\sigma(x)} \right)_+ \cdot v_1(x)}{\sqrt{v_1^2(x) + v_2^2(x)}} \quad (4.15)$$

with $(a)_+ = \max(a, 0)$. A denoised coefficient $\hat{w}_1(x)$ is calculated from the corresponding noisy coefficient $v_1(x)$ and its parent $v_2(x)$. σ_n^2 denotes the noise variance, and $\sigma(x)$ denotes the marginal standard deviation for the coefficient at location x (in practice, $\sigma(x)$ is estimated as the standard deviation of the wavelet coefficients in a 7x7 window around the coefficient at position x). This algorithm is simple to implement, has low computational cost and yet provides state-of-the-art noise reduction, since

- it adapts locally to the presence of edges (due to the presence of $\sigma(x)$)
- the factor $(v_1^2(x) + v_2^2(x))$ captures the relation between coefficients across scales. It tends to keep wavelet coefficients unchanged when both the current coefficient $v_1(x)$ and its parent $v_2(x)$ are large, which means that they probably originate from a significant image feature, see again Figure 4.36.

In Figure 4.38, we show a comparison of SURE thresholding [Donoho and Johnstone, 1994, Donoho and Johnstone, 1995] versus bivariate shrinkage [Sendur and Selesnick, 2002]. In Figure 4.38 (a), the noisy image is shown. In Figure 4.38 (b), the result of SURE thresholding is shown. Note that the noise is reduced, but that due to the global nature of the threshold, there is no difference between edge areas and smooth areas. Also, the image gets a kind of “grainy” appearance due to residual noise. In Figure 4.38 (c), the result of bivariate shrinkage is shown. Note that here the noise is significantly better suppressed, while maintaining (most) of the sharpness in edge areas.

For the sake of completeness, we also mention some other state of the art methods in wavelet based noise reduction. In [Pižurica et al., 2002], a

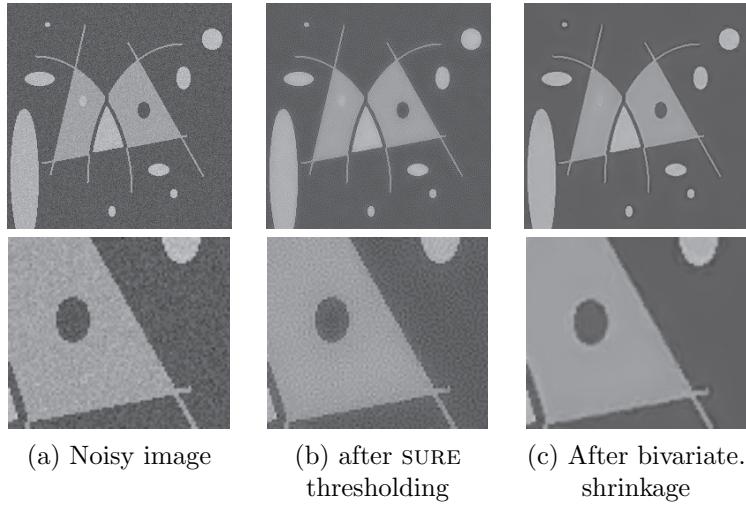


Figure 4.38: Illustration of the SURE versus the bivariate shrinkage results.

Markov Random Field (MRF) model was used to compute filter masks in order to preserve useful features, while suppressing noise in smooth areas. In [Portilla et al., 2003], the relations between wavelet coefficients were modeled by a Gaussian Scale Mixture model (GSM model). This model allows to maintain non-Gaussian coefficients, while suppressing Gaussian-like noise.

4.3.2.2 Poisson noise

The noise reduction based on wavelet SURE thresholding was already applied in microscopy by [Boutet de Monvel et al., 2001, Stollberg et al., 2003]. This noise reduction is suboptimal for two reasons. First, it assumes Gaussian noise. In confocal fluorescence imaging however, the major source of noise is Poisson distributed [Pawley, 1995, van Kempen et al., 1997]. Unlike Gaussian noise, Poisson noise is intensity dependent, which makes separating image from noise very difficult [Nowak and Baraniuk, 1999]. Secondly, it only exploits that useful coefficients are supposed to be large, and doesn't take into account spatial correlation and correlation across scales (see again Figure 4.36). We deal with these problems in the following way. To deal with Poisson data, we use the Anscombe transform [Starck et al., 2000], which approximately transforms the Poisson data into data with a Gaussian distribution with unit standard deviation. The Anscombe transform is given by

$$t(I(x, y)) = 2\sqrt{I(x, y) + \frac{3}{8}} \quad (4.16)$$

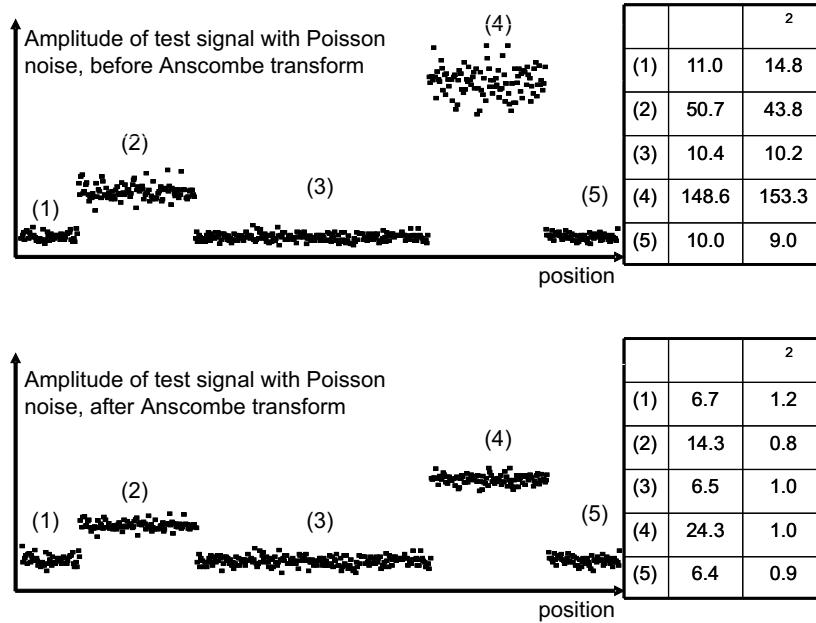


Figure 4.39: Illustration of the Anscombe transform on data with Poisson noise. The top half shows a synthetic test signal consisting of a piecewise constant function with Poisson noise. The table shows the expectancy value (μ) and variance (σ) over each originally constant interval. The bottom row shows the same data after the Anscombe transform. Note that now all variances are approximately one.

In Figure 4.39, we illustrate the effect of the Anscombe transform. We defined a piecewise constant function, and generated Poisson noise. It is clearly visible that the variance is higher in the intervals where the average intensity is higher. However, after applying the Anscombe transform, Poisson data with mean intensity $I_{poisson,1}$ and variance $\sigma_{poisson,1}^2$ is transformed into Gaussian data with mean $I_{gaussian,1}$ and variance 1.0, while Poisson data with intensity $I_{poisson,2}$ and variance $\sigma_{poisson,2}^2$ are transformed in Gaussian data with mean $I_{gaussian,2}$ but again with variance 1.0, so that the noise has become independent of the signal. This can be seen in Figure 4.39, where our piecewise constant dataset is shown before and after the Anscombe transform. For each constant segment, mean and average are given before and after the transform.

The Anscombe transform allows us to use well studied methods for Gaussian noise on data with the much trickier Poisson noise. We adapted this algorithm for implementation with steerable pyramids (by finetuning some parameters). After denoising, the inverse Anscombe transform was applied.

In Figure 4.40, the effect of the Anscombe transform on the noise reduction results of a synthetic test image contaminated with Poisson noise is demonstrated. In both cases, the noise in the background has been removed almost

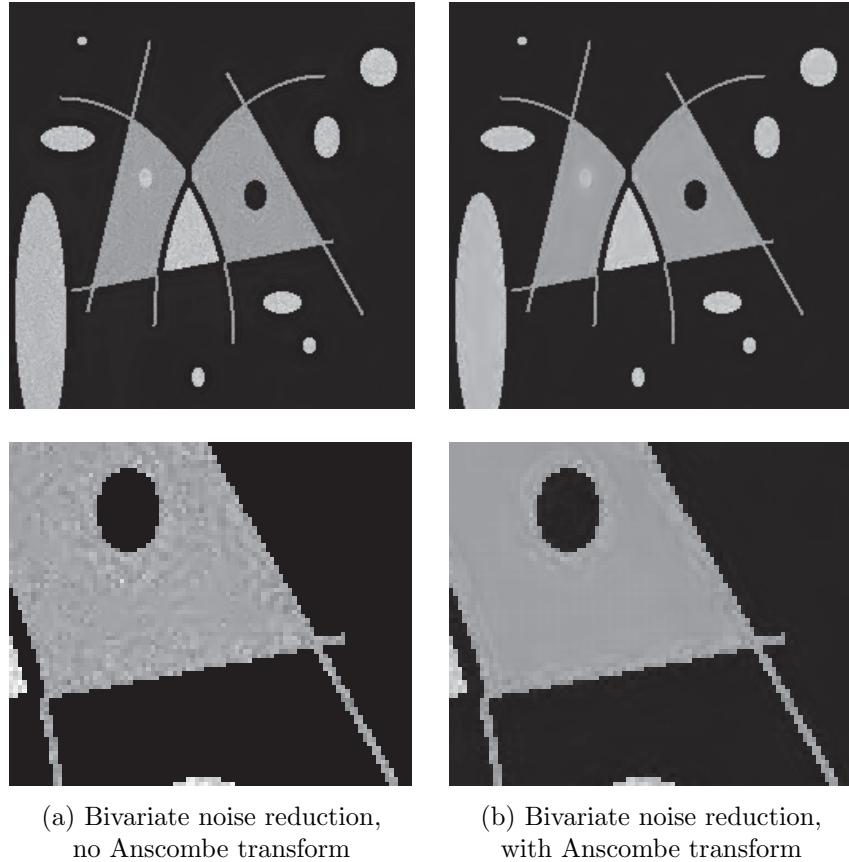


Figure 4.40: Illustration of the effect of the Anscombe transform on noise reduction on a synthetic test image with Poisson noise.

completely, but without the Anscombe transform (a) in the bright areas, noise is still present, because there the noise variance is a lot bigger than in the background. In (b), the Anscombe transform has normalized the noise variance prior to noise filtering, and noise has been removed much better in the bright areas. Note however that in the neighborhood of the borders, some noise is still present.

4.3.3 Image restoration

In this subsection, we give an overview of image restoration techniques involving wavelets. Those techniques can be divided into two categories:

- regularization and deblurring are both performed in the wavelet domain;

- only regularization is performed in the wavelet domain, but it is combined with a deblurring in the spatial or the Fourier domain.

4.3.3.1 Hybrid techniques: only regularization in the wavelet domain

It is most logical to represent the convolution operator in the Fourier domain (there the operator is diagonalized) [Figueiredo and Nowak, 2003]. However, an image is modeled better in the wavelet domain. On the other hand, it is generally quite difficult to express convolution operators in the wavelet domain (however, a wavelet-based algorithm has been suggested in [Beylkin et al., 1991, Beylkin, 1991], but this isn't often used in image processing due to its complexity). This suggests the combination of Fourier-based deconvolution with wavelet-based denoising as regularization. The first paper about this subject was [Neelamani et al., 1999] which was extended in [Neelamani et al., 2004]. In those papers, the authors combined a simple Wiener filter for deconvolution with a wavelet-based Wiener noise filter as regularization.

Both [Kalifa et al., 2003] and [Jalobeanu et al., 2003] perform an unregularized inversion of the blur operator, followed by a wavelet-based suppression of the noise (which is colored due to the fact that the noise in the high frequencies is amplified stronger than in the low frequencies). In [Kalifa et al., 2003], mirror wavelet bases were used to model the image in such a way that noise can be separated from the useful image, even in the high frequencies, where the noise has been amplified strongly. In [Jalobeanu et al., 2003], a similar approach is followed, but in this case, complex wavelet packets are used to distinguish noise from signal.

In [Figueiredo and Nowak, 2003], an FFT-based Landweber-like deconvolution step is followed by wavelet-based MAP/MPLE (*Maximum A Posteriori / Maximum Penalized Likelihood Estimator*) soft-thresholding step. Our own method SPERRIL (described in more details in Chapter 6) also falls in this category.

4.3.3.2 Denoising and deblurring in the wavelet domain

In [Banham et al., 1994, Banham and Katsaggelos, 1996], a prefilter combined with a multiscale Kalman filter has been applied to the deconvolution problem. The prefilter performs constrained least-square filtering (in fact a translation into the wavelet domain of the method we described in Subsubsection 3.2.1.1), which provides a candidate solution that is very sharp, but possibly noisy. This sharp, noisy intermediate image is then used as input for the multiscale Kalman smoothing filter (which is not just a regularization step, but also explicitly deals with the blur). The prefilter also has the effect to remove cross-scale dependencies within the state vectors (vectors consisting of a wavelet coefficient

and its children), thus substantially reducing the length of these state vectors, which is beneficial for the computational efficiency of the algorithm.

In [Portilla and Simoncelli, 2003], an extension of [Portilla et al., 2003] to the full image restoration problem is presented. There, the steerable pyramid decomposition is applied on an image. Every coefficient of the subbands of this decomposition is considered in a neighborhood of 3x3 pixels around the reference coefficient, plus the parent of the reference coefficient. This neighborhood is modeled using a Gaussian Scale Mixture model (GSM): it models the neighborhood as the product of a Gaussian vector and an independent hidden random variable. In [Portilla and Simoncelli, 2003], the GSM model is extended to take into account also image degradation due to blur (dealing with an image that is only corrupted with noise was already described in [Portilla et al., 2003]). Then, the reference coefficients are estimated from their neighborhoods using the Bayes Least Square (BLS) estimator.

In [Dragotti and Vetterli, 2003], it is shown that any piecewise smooth signal can be expressed as a sum of a piecewise polynomial signal and a uniformly smooth residual. Wavelet footprints are scale space vectors of coefficients across scales that model discontinuities in piecewise polynomial signals. E.g., for a piecewise constant signal that contains only one singularity at position x , the footprint would be a vector containing all the wavelet coefficients in the cone of influence of x (but not the lowpass coefficient associated with position x), followed by a normalization of this vector. First, a wavelet footprint dictionary is constructed from the wavelet transform of the signal. Based on this dictionary, operations like denoising and deconvolution can be performed. Denoising for example is performed by thresholding the wavelet footprints instead of all the wavelet coefficients. Deconvolution is performed by first denoising the signal, and then replacing the blurry footprints by the corresponding unblurred versions of the footprints. Results for 1D signals are presented, but the technique is not yet extended to higher dimensions.

Our second method, which is based on multiscale Markov Random Fields (described in Chapter 7), also falls in this category.

4.4 3D multiscale transforms

Since in this thesis, we aim to restore confocal microscope images, and since those images are three-dimensional, we need 3D extensions of the multiscale transforms we have discussed so far. These extensions are necessary if we want to use the the algorithms we have discussed before for, e.g., denoising of 3D images.

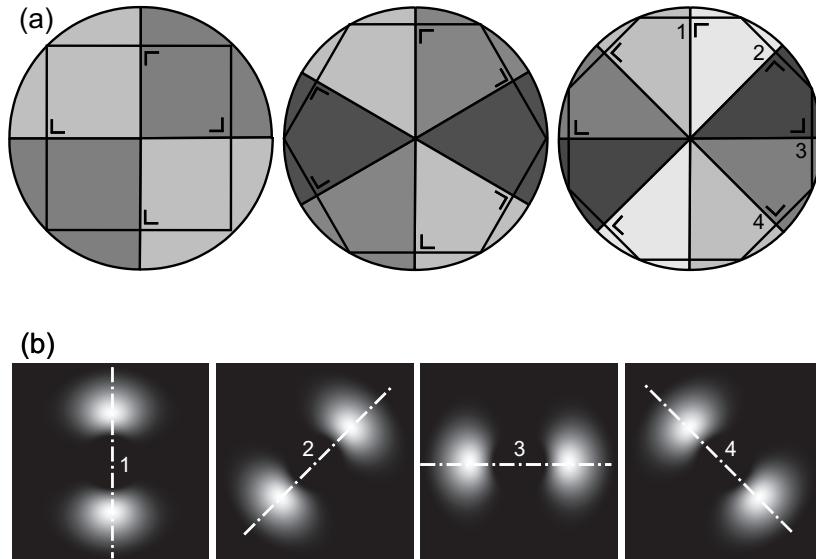


Figure 4.41: In (a), an isotropic division of a circle is shown. In order to divide the oriented subbands isotropically on a circle, it is sufficient to divide the subbands along the medians of a regular polygon. The example shows the case of two, three and four oriented subbands. In (b), a set of real oriented bandpass filters is shown (absolute value) for four orientations.

4.4.1 3D steerable pyramid

The 3D steerable pyramid was introduced in literature by Freeman [Freeman and Adelson, 1991]. Here, we give a brief introduction, loosely based on [Delle Luche et al., 2003]. The main problem to extend the steerable pyramid from 2D to 3D is not the definition of the lowpass and highpass filters, since those filters are circularly symmetric. The real problem is to distribute the oriented bandpass filters isotropically over the sphere. In 2D, the central axes of the oriented bandpass filters are oriented along the diagonals or the medians of regular polygons (for case of the medians, see Figure 4.41).

In [Freeman and Adelson, 1991], a 3D extension of steerable filters is proposed based on expansion in spherical harmonics. In [Andersson, 1992], a method is proposed to use regular polyhedra (the logical extension of regular polygons to 3D) to orient the bandpass filters and thus obtain a uniform set on the sphere. In [Yu et al., 2000], a local method based on approximate steerability using conic Gaussian filters is proposed.

In [Delle Luche et al., 2003], a new method was proposed which combines the conic filters idea with the regular polyhedron concept. There, the diagonals or medians (a median of a polyhedron is a segment that connects a vertex of the polyhedron to the midpoint of the opposite surface) of regular polyhedra are

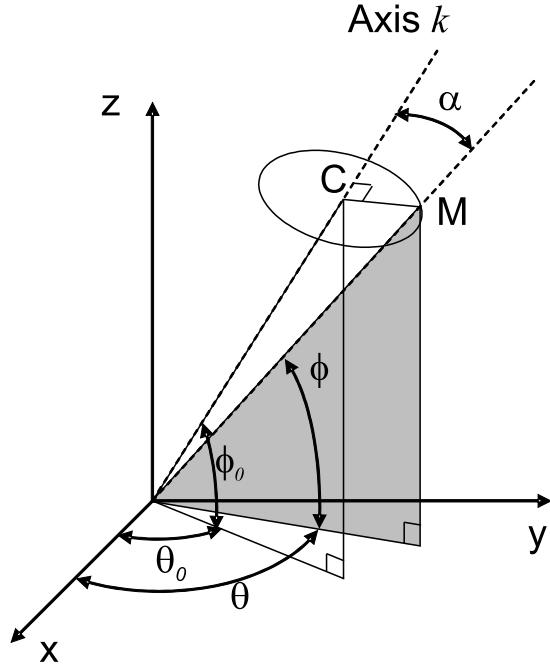


Figure 4.42: Some illustrations with the formulae regarding $\cos(\alpha)$.

chosen as central axes of the 3D oriented bandpass filters. These K bandpass filters can be expressed as:

$$B_k(\alpha) = \cos^n(\alpha) \quad k = 1, \dots, K \quad (4.17)$$

where n can take the values 1 or 2, and α is the angle with respect to the k^{th} central axis (see Figure 4.42):

$$\cos(\alpha) = \frac{\overrightarrow{OM} \cdot \overrightarrow{OC}}{\|\overrightarrow{OM}\| \|\overrightarrow{OC}\|}$$

In [Delle Luche et al., 2003], it is mentioned that for a cube, there can be three directions (if the medians are chosen as central axes) or four directions (if the diagonals are chosen). In that case, $n = 1$, and equation (4.17) becomes $B_k(\alpha) = \cos(\alpha)$. For an icosahedron or dodecahedron, [Delle Luche et al., 2003] mentions that there are six or ten directions. In that case, equation (4.17) becomes $B_k(\alpha) = \cos^2(\alpha)$ in order to cover the whole Fourier domain by this set of angular transfer functions. We implemented this algorithm for the case of three

central axes, which are in fact the three orthogonal coordinate axes. In Figure 4.43, we show the three angular transfer functions that correspond with $B_k(\alpha) = \cos(\alpha)$.

In Figure 4.44, we show an example of a 3D steerable pyramid decomposition of a synthetic image of a cubic shell (a cube with an empty cubic shaped center) in case of three oriented bandpass functions per scale. We show volume renderings⁶ of the three oriented subbands of the finest scale, and see indeed that the orthogonal surfaces of the cube have responses in different subbands.

4.4.2 Other 3D multiscale transforms

Of course, other approaches exist to formulate a 3D decomposition/reconstruction scheme. One obvious example is the separable transform of Mallat, described earlier. The extension to 3D is straightforward, but also has the same problems of shift-invariance and orientation resolution.

Another method exists that possesses the required properties of shift invariance and orientation selectivity, i.e., the complex wavelet transform. This transform was proposed by [Kingsbury, 1998b, Kingsbury, 1998a] and extended by [Selesnick and Li, 1998, Selesnick, 2004]. However, this decomposition is not steerable (see Figure 4.45 for a schematic representation of the oriented transfer functions in case of the steerable pyramid (a) and complex wavelets (b)). There, one can see that the transfer functions for the steerable pyramids are polarly separable, which means that orientation of these subbands (steering) is fairly easy. This is not the case for the complex wavelets, which have transfer functions shaped like rectangles.

⁶The volume rendering was calculated with VolumeJ, a free plugin for the free image processing program ImageJ [ImageJ, WWW]

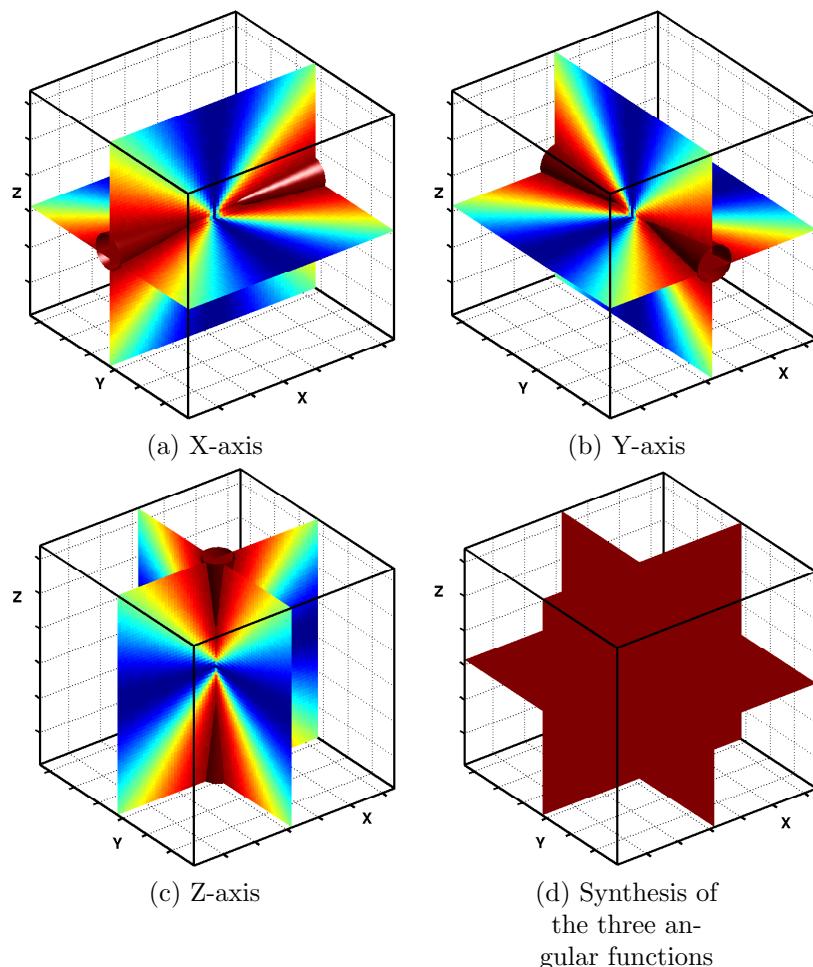


Figure 4.43: The 3 angular functions in case of three directions for the 3D steerable pyramid (a-c), and the recombined transfer function (d), which is unity everywhere. Three orthogonal slices through the transfer functions are shown, where red corresponds to the highest value and blue to the lowest value.

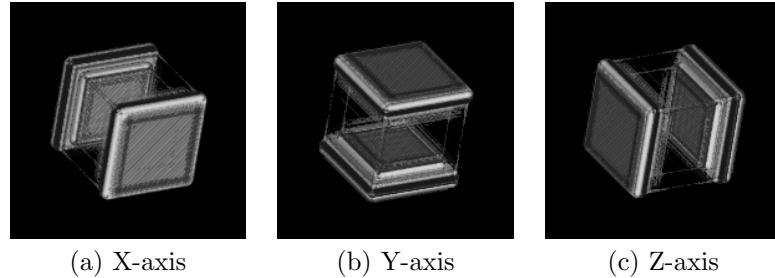


Figure 4.44: Example decomposition of a synthetic cubic shell image in case of three directions for the 3D steerable pyramid (rendered with VolumeJ, plugin of [ImageJ, WWW]).

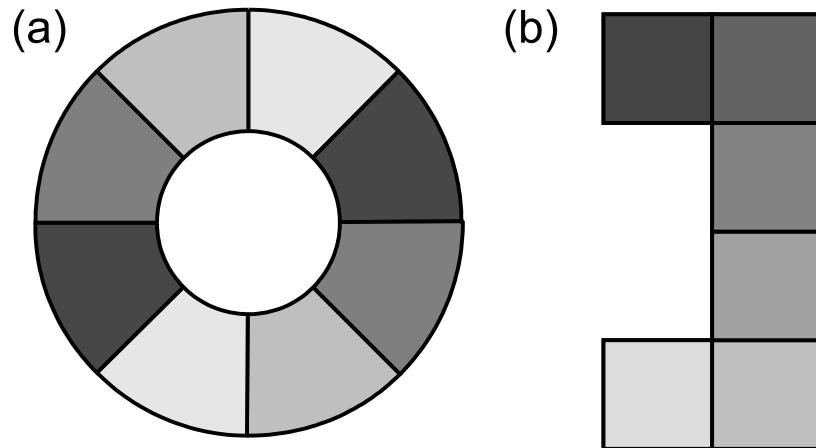


Figure 4.45: Oriented transfer functions of steerable pyramids and complex wavelets. The transfer functions for the steerable pyramids (left) are polarly separable, which is not true in case of the complex wavelets (right). Note that complex wavelets are analytical, and have only frequency response in half of the Fourier domain.

Chapter 5

Two methods for parametric PSF estimation in images

Q: Why don't statisticians like to model new clothes?

A: Lack of fit.

5.1 Introduction

In previous chapters, we discussed the problem of image degradation by blurring and noise. To repeat the main points, the degradation is modeled as follows:

$$g(x, y) = N((h * f)(x, y)),$$

where $g(x, y)$ is the observed, degraded image. $f(x, y)$ represents the unknown, ideal image we wish to recover and $h(x, y)$ is the blurring kernel or PSF. $N(\cdot)$ represents the noise processes that influence the imaging process. Often, noise is modeled to be additive, and independent of the underlying signal. In that case, we can rewrite the previous equation as:

$$g(x, y) = (h * f)(x, y) + n(x, y),$$

where $n(x, y)$ is the additive noise (in our case Gaussian noise).

Image blur is often unwanted and has to be compensated for (this is image restoration, and is applied in astronomy, medical imaging, microscopy,

...). A problem in image restoration is that the blurring kernel $h(x, y)$ is often unknown, or only partially known. In this case, $h(x, y)$ has to be estimated. This estimation can be performed prior to the restoration process (which we will refer to as *blur estimation*), or simultaneously during the restoration process (which we will refer to as *blind deconvolution*). A good review of the latter kind of methods can be found in [Kundur and Hatzinakos, 1996a, Kundur and Hatzinakos, 1996b]. In this class, there are also a few methods that alternatively estimate the image and the blurring kernel [Ayers and Dainty, 1988, Seldin and Fienup, 1990, Holmes, 1992, Fish et al., 1995]. Blind deconvolution algorithms usually are complex, and have a heavy computational load.

Sometimes, the blur contains extra information. For example, it can provide information about the settings of the camera. When dealing with autofocus cameras, one expects to find a sharp image, because ideally all natural images contain sharp objects in front of a background. When the camera is out of focus, blur estimation can help to estimate how much an out-of-focus camera needs to be adjusted.

Blurred edges can also provide information about the 3D nature of the scene itself. In those applications, depth is estimated from focus/defocus [Schechner and Kiryati, 1998, Xiong and Shafer, 1993]. Again, we assume that objects in front of a background have sharp edges. However, only objects in the focal plane are imaged with sharp edges. For objects not in the focal plane, these sharp edges are blurred in proportion to their distance from the focal plane, thus providing some depth information about the image. In this case, a local blur estimation provides information about the depth in the image scene.

We now give a summary of common methods for *blur estimation*. A part of this summary has been loosely based on the introduction of [Flusser and Suk, 1998], where blur estimation was considered for image analysis. A first approach is based on the analysis of the zero patterns in the Fourier spectrum [Gennery, 1973, Cannon, 2000, Chang et al., 1991]. The methods are particularly useful in estimating motion blur and out-of-focus blur, since those kind of blurs produce characteristic zero patterns in the Fourier spectrum. On the other hand, it is not suitable for blur with Gaussian PSF's. Another approach starts from the assumption of an autoregressive (AR) model for the image, while the blur is modeled by a moving average (MA) process. The blur estimation is therefore transformed into the estimation of the parameters of the AR-MA model. In [Lagendijk and Biemond, 1991, Pavlović and Tekalp, 1992] a maximum likelihood method is proposed to estimate the blur parameters. In [Reeves and Mersereau, 1992, Chardon et al., 1999], the estimation of the parameters relies on the method of Generalized Cross-Validation (GCV).

The method described in [Savakis and Trussell, 1993a] is in fact not really a parametric blur estimation method, but a method to decide which of a set of given blurring kernels is the best for restoration (even if those candidate PSF's are not of some parametric form). This method minimizes the spectral

difference between the image restored with the different candidate kernels, and the spectrum which should be found when the degraded image is restored with the correct kernel.

In the next two sections, we present two new wavelet-based techniques for parametric blur estimation. As far as we know, these are the first attempts to estimate image blur using the wavelet transform.

5.2 Wave-Track: tracking wavelet coefficients across scales

5.2.1 Introduction

In this section, a wavelet-based method is proposed to estimate the blur in an image using information contained in the image itself. It was inspired by related work on image denoising by dr. A. Pižurica from our research group. In her work [Pižurica, 2002], she used the regularity of image features (using histograms of the Lipschitz exponents) to distinguish useful image features from noise. We adapted this method for the estimation of image blur. This technique has been published in [Rooms et al., 2001b, Rooms et al., 2002b, Rooms et al., 2002c], and is incorporated in the automatic image restoration algorithm SPERRIL, which is extensively described in Chapter 6.

We look at the sharpness of the sharpest edges in the blurred image. This sharpness contains information about the blurring. Specifically, a smoothness measure is computed for these edges. A relation between this parameter and the evolution of the wavelet modulus maxima across the different resolution scales is shown, which is only dependent on the blur in the image and not on the image contents. In particular, our method estimates the variance parameter σ_{blur} of a Gaussian PSF:

$$\text{PSF}(x, y) = \frac{1}{2\pi\sigma_{blur}^2} e^{-(x^2+y^2)/(2\sigma_{blur}^2)} \quad (5.1)$$

Wave-Track can estimate the image blur parameter σ_{blur} with a reasonable accuracy. This method can also be extended to Gaussian PSF's that are not rotationally symmetrical and even to PSF's that aren't Gaussian. For out-of-focus blur, a uniform circular PSF is used [Pavlović and Tekalp, 1992, Savakis and Trussell, 1993a]. Wave-Track requires only minor modifications to adapt to this kind of PSF, as shown in this Section.

5.2.2 Principle

5.2.2.1 Theoretical background

Our method for parametric blur estimation is based on estimating a measure of sharpness for the sharpest edges in the image. To analyze edges in the image, the rate of evolution of the magnitude of the wavelet coefficients across scales is computed in all points where an intensity gradient is found either in the horizontal or vertical direction. This evolution across scales depends on three factors as Mallat has shown in [Mallat and Zhong, 1992, Mallat and Hwang, 1992, Mallat, 1999]:

- the Lipschitz exponent of the underlying signal;
- the basis functions used in the wavelet decomposition;
- the blurring of the signal.

The first factor is the Lipschitz exponent (also called Hölder exponent), which is a measure of the image regularity at a certain point. In fact, it is an extension of the degree of differentiability of the image at a certain point. For example, a signal that is differentiable once, has Lipschitz exponent +1, a step function (not differentiable) has Lipschitz 0 at the position of the step and a Dirac impulse $\delta(x)$ has Lipschitz -1 at the position of the impulse. The relation between regularity of a signal and the evolution of the wavelet coefficients across scales, is illustrated in Figure 5.1. For a sharp variation in amplitude, the wavelet coefficients decrease in magnitude, and for a smooth variation, the wavelet coefficients increase in magnitude with increasing scales. However, when Gaussian blur is applied to an image, all image features are infinitely differentiable, thus the Lipschitz exponent doesn't capture the image blur.

In fact, only the sharpest edges with a reasonably large amplitude in the image should be taken into account in order not to take into account wavelet coefficients originating from noise or smooth image features. Therefore, we select only the wavelet coefficients with magnitude larger than a certain threshold. The evolution of these wavelet coefficients then only depends on the blur of the corresponding image features. For unblurred step edges, e.g., the wavelet coefficients are constant in magnitude across scales, and for a smooth variation, the wavelet coefficients increase in magnitude from finer to coarser scales. From this rate of increase through the successive resolution scales, we can determine the blur of an edge in the image. This is illustrated in Figure 5.2.

For every pixel that is considered to be an edge pixel, the modulus maximum of the corresponding wavelet coefficients is tracked across scales, and a function is fitted to the wavelet coefficients $|W(x, y; s)|$ as a function of the scale s :

$$|W(x, y; s)| = a \exp(b s) \quad (5.2)$$

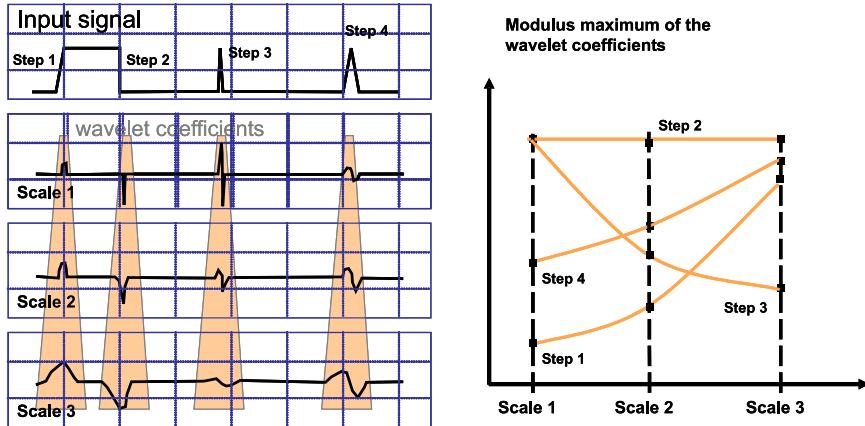


Figure 5.1: Evolution of wavelet coefficients across scales for intensity variations of a one-dimensional signal with different regularity.

with $|W(x, y; s)|$ the modulus of a wavelet coefficient at scale s and position (x, y) ; a and b are constants estimated from the data. The coefficient b is related to the blur of the edge, so we called it the *blur exponent*, in analogy with the concept of Lipschitz exponent. More recently, an analytic expression was derived that connects the evolution of the wavelet coefficients across scales with the amount of blur that was applied to the corresponding image features [Barat et al., 2003]. However, this method assumes knowledge of the type of the feature (line, edge, point, ...) and uses a continuous wavelet transform. In our case, we used a much simpler and faster approximation.

From the blur exponents corresponding with pixels from the significant edges in the image, a histogram is constructed. Therefore, we divided the range of blur exponents in intervals with a width of 0.1. Because we restricted the blur exponents to those corresponding with transitions with large amplitude, we already filtered out the sharpest transitions with a large amplitude in the image.

Then, the center of gravity (CG) of the histogram is a measure for the blur in the image. We calculate the blur exponent that corresponds to the CG of the histogram, and determine the average $\text{CG}_{\sigma_{\text{blur}}}$ over the whole set of test images blurred with the same σ_{blur} . To these data $(\sigma_{\text{blur}}, \text{CG}_{\sigma_{\text{blur}}})$, another exponential curve was fitted experimentally:

$$\sigma_{\text{blur}} = c \exp(d \text{CG}_{\sigma_{\text{blur}}}). \quad (5.3)$$

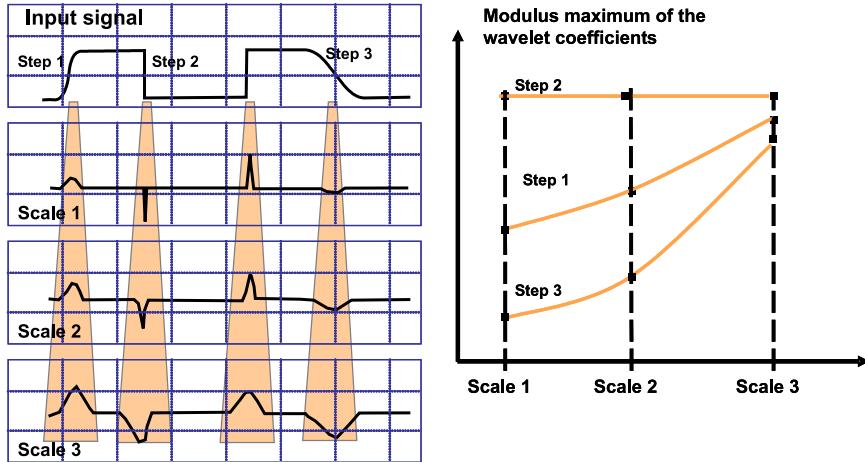


Figure 5.2: Evolution of wavelet coefficients across successive scales for a step edge, blurred with different amounts of blur.

The fitting of the exponential function to N points (x_i, y_i) , with $i = 1, \dots, N$ was computed by minimizing the function [Weisstein, WWWa]:

$$\sum_{i=1}^N y_i (\ln y_i - a - bx_i).$$

Setting the derivatives of this expression to a and b to zero, we can solve those expressions for a and b . Fitting this relation for a set of digital camera images synthetically degraded with a range of values for σ_{blur} is shown in Figure 5.3.

A schematic representation of the algorithm is shown in Figure 5.4.

Remark that in our implementation, we assumed isotropic blur. Therefore, we only used the vertical detail subbands and the corresponding cones of influence.

5.2.2.2 Robustness to noise

The technique described above is very noise sensitive, as could have been expected for any technique based on finding local maxima. Although the histogram of Lipschitz exponents of pixels located along edges should be (almost) identical in the case of a noise-free image and its noisy counterpart [Hsung et al., 1999, Pižurica, 2002]. However, this would require prior knowledge about the location of the edges. In our case, we just applied a threshold to the magnitude of the wavelet coefficients to select the wavelet coefficients that are more likely to originate from edges. In a noisy image, also noisy wavelet coefficients in smooth areas will have an amplitude higher than a certain threshold, thus also contributing to the histogram.

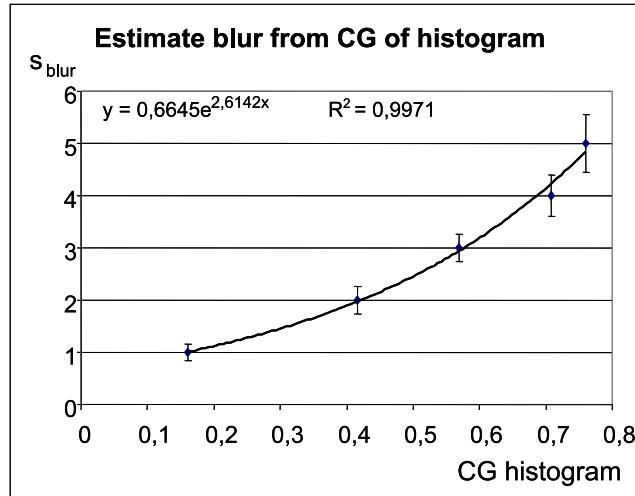


Figure 5.3: Fitting of the exponential function between the CG of the histogram, and the real σ_{blur} .

In applications like digital cameras under sufficient illumination conditions, our experiments show that applying a median filter is a sufficient preprocessing step for reliable blur estimation because the noise in cameras is impulse like (see Figure 5.5). In general cases however, when no precautions are taken, noise will disturb the blur estimation. There are two reasons for this. The first reason is because edges aren't detected accurately in the presence of noise. The second reason is that the blur exponents on detected edges are disturbed.

In [Canny, 1986], the problem of edge detection¹ in the presence of noise was handled by Gaussian smoothing. In [Elder and Zucker, 1998], this technique is incorporated in a probabilistic scale-space framework, and an expression is derived for how much smoothing is required for an edge to be detected reliably in the presence of noise. This amount of smoothing depends on:

- the contrast (or amplitude) of the edge;
- the original edge smoothness;
- the noise level.

The minimal required amount of smoothing is called the *minimum reliable scale* for that image feature. To actually compute this value, one needs the edge characteristics and the noise level, which are often unknown. In practice, these parameters are estimated iteratively until a certain reliability criterion is satisfied, which is a time-consuming procedure.

¹Here, the goal is just to find and localize edges and not yet to restore them.

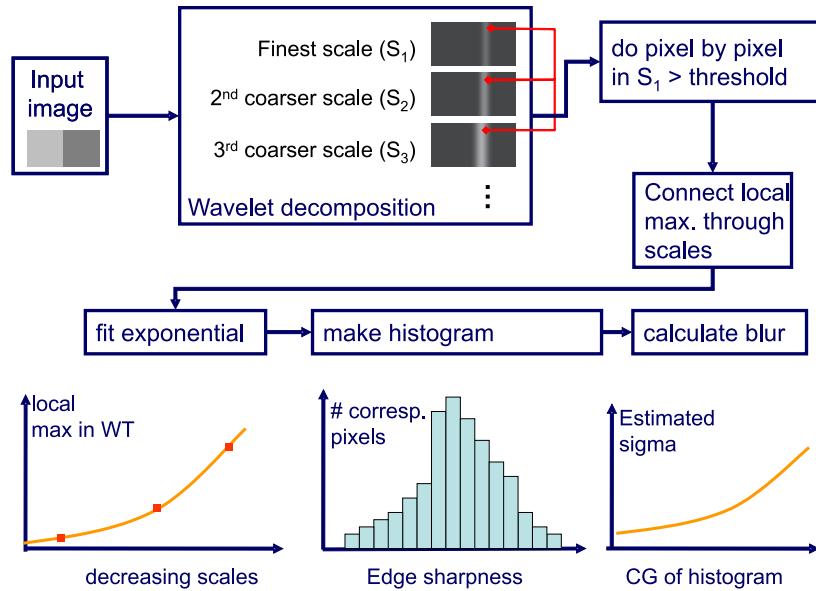


Figure 5.4: Schematic representation of the algorithm.

This leads to our initial attempt to stabilize our blur estimation method. We wanted to know how much smoothing was required in order to obtain a reliable blur estimation in the presence of noise. This additional smoothing was applied as a preprocessing step in our own blur estimation technique as described before. The additional blur value was later subtracted from the blur estimated in the smoothed image $\hat{\sigma}_{blur}$ to obtain the original blur in the image:

$$\sigma_{blur} = \sqrt{\hat{\sigma_{blur}}^2 - \sigma_{postblur}^2} \quad (5.4)$$

In Figure 5.6, we see the effect of the additional smoothing on a histogram of blur exponents from an image degraded by noise. Figure 5.6 (a) shows the histogram for an image that is only blurred; Figure 5.6 (b) shows the histogram for an image that is blurred and contaminated with noise; Figure 5.6 (c) is the result of applying smoothing as prefilter to the image used in Figure 5.6 (b).

In the experiments, we see that the Center of Gravity (CG) of the histogram of blur exponents calculated along the edges in the image is a reliable parameter to estimate parametric modeled blur in an image. However, the standard deviation on the estimate increases as the original image blur increases. Applying additional blur reduces the effect of noise in our blur estimation. However, this still didn't stabilize the blur estimation enough.

In its present form, the stabilization of the blur estimation was performed using

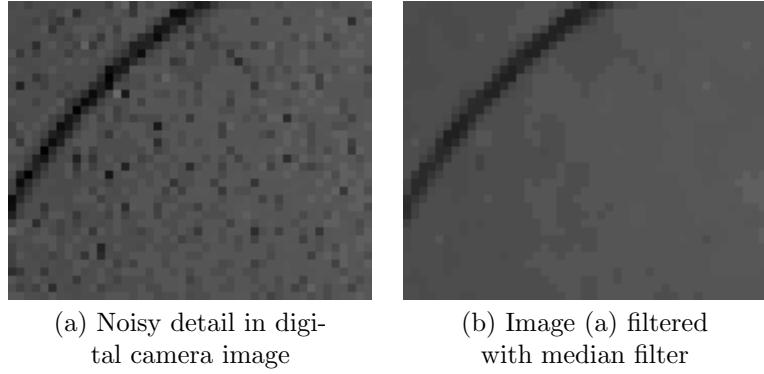


Figure 5.5: Prefilter of digital camera images.

a little heuristic trick: as is described in Chapter 6, the blur estimation was integrated in the restoration method. By integrating the noise reduction, the blur estimation and a deblurring step, the blur estimation is stabilized more efficiently than with the preblurring method.

5.2.3 Evaluation of Wave-Track

We have evaluated the method using a set of several natural images (*Airplane*, *Peppers*, *House*, *Flintstones* and *Lake*). These images were synthetically degraded with different amounts of blur and noise.

In the set of graphs in Figure 5.7 we show the average estimated blur values obtained for a set of test images. The left figure shows the results of our method without any compensation, while the right figure shows the results compensated to minimize the overall error in estimating σ_{blur} over the whole range of blur and noise levels. Our method tends to overestimate the blur systematically for higher noise levels. A possible reason is because for these high noise levels the prefilter of SPERRIL (see Figure 6.1, and will be described in more detail in Subsection 6.2.1) oversmoothes the image. Since we assumed the noise level to be known, we can compensate accordingly. By trial and error, we have hand-optimized a correction function to compensate for this overestimation in function of the noise only:

$$\sigma_{blur, \text{corrected}} = \sigma_{blur, \text{uncorrected}} - (0.055 \ln(\sigma_{noise}) + 0.1)$$

This way, it is possible to compensate for most of this overestimation.

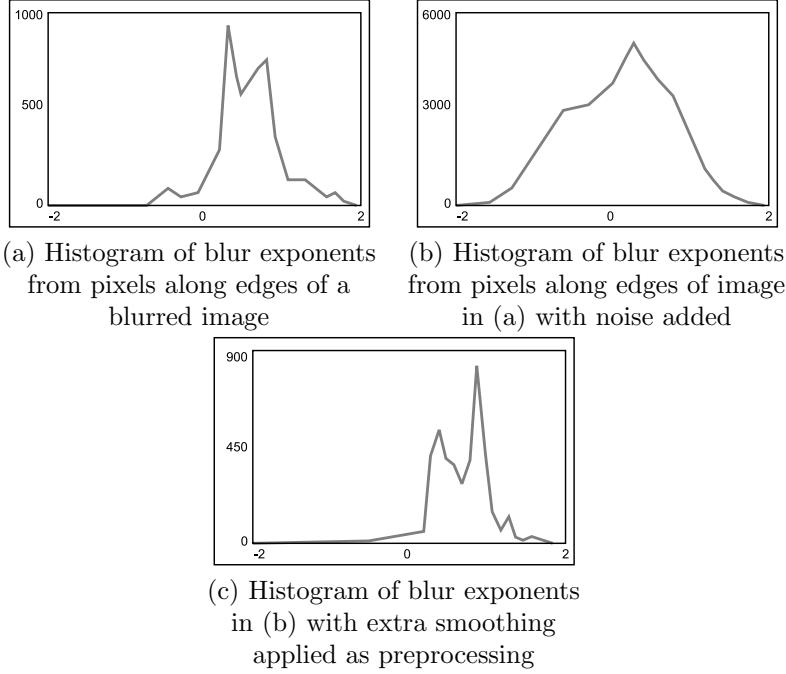


Figure 5.6: Initial stabilization attempt of blur estimation in the presence of noise.

5.2.4 Future work: extension of Wave-Track to anisotropic PSF's

Our implementation for now assumes isotropic blur, but can in principle be extended to anisotropic blur, as we explain now. Preliminary results are promising, however the practical implementation is quite hard since the tracking across scales of the wavelet modulus maxima for the different orientations requires a lot of book keeping.

This blur estimation procedure can be performed independently for the four orientations in the steerable pyramid. Assuming that the PSF is Gaussian and spatially invariant, we average out the blur measures for the different edge pixels for each of the four orientations to obtain a robust estimate for the width of the projection of the PSF in the direction perpendicular to the orientation of the subband. This way, we obtain $\sigma_{blur,k}$ for $k \in [0, K-1]$, where $\sigma_{blur,k}$ represents the width of the PSF projected under orientation $k\pi/4$. We know that for a projection of the PSF under an angle θ , we can write

$$\sigma_\theta^2 = E(\cos(\theta)\mathbf{x} + \sin(\theta)\mathbf{y})^2$$

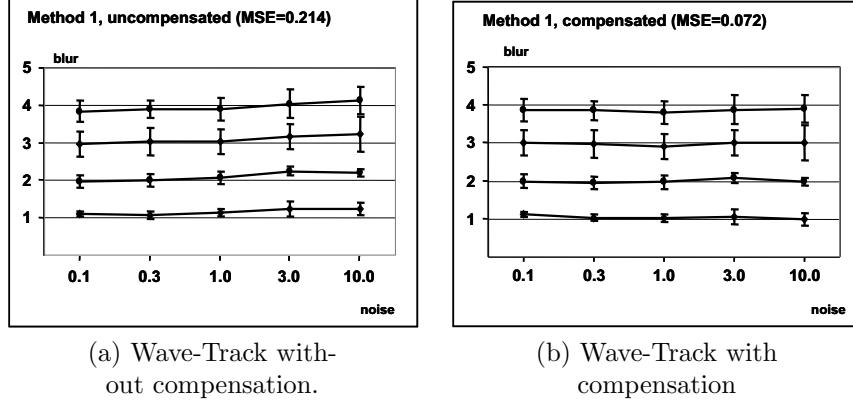


Figure 5.7: Summaries of blur estimation results for Wave-Track method, both uncompensated and compensated. In each graph, a horizontal line represents the estimations for 4 different blur levels for a range of noise levels. The error bars correspond to $(\mu + \sigma, \mu - \sigma)$, with μ the sample mean and σ the sample standard deviation over the test set. In both subfigures, the horizontal curves correspond with $\sigma_{blur,real} = 1$ (bottom curve) to $\sigma_{blur,real} = 4$ (top curve).

with \mathbf{x} and \mathbf{y} the data vectors obtained from the 2D distribution projected according to the x- and the y-axis. When $E(\mathbf{x})$ and $E(\mathbf{y})$ are zero (like for point-symmetric, zero-centered PSF's), this equation becomes

$$\sigma_{blur,\theta}^2 = \cos^2(\theta)\sigma_{blur,x}^2 + \sin^2(\theta)\sigma_{blur,y}^2 + 2\cos(\theta)\sin(\theta)\sigma_{blur,xy} \quad (5.5)$$

where $\sigma_{blur,x}$ is the width of the PSF projected on the x-axis; $\sigma_{blur,y}$ is the width projected on the y-axis and $\sigma_{blur,xy}$ is the covariance of the 2D distribution that corresponds to the PSF. We estimate these values from the oriented subbands for features in directions $\pi/4$ and $3\pi/4$, so we can calculate $\sigma_{blur,xy}$, which allows us to reconstruct the PSF with the following formula:

$$h(\mathbf{X}) = \frac{1}{2\pi|\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{X} - \mu)^T \Sigma^{-1} (\mathbf{X} - \mu)\right) \quad (5.6)$$

where \mathbf{X} is the vector containing the spatial coordinates, $\mu = 0$ and Σ the covariance matrix:

$$\Sigma = \begin{pmatrix} \sigma_x^2 & \sigma_{xy} \\ \sigma_{xy} & \sigma_y^2 \end{pmatrix} \quad (5.7)$$

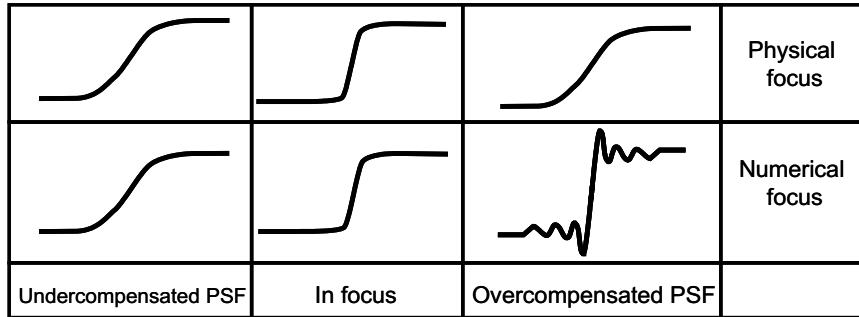


Figure 5.8: Physical versus numerical focus effect on the cross-section of an edge.

5.3 Sparse-Max: wavelet-based sparseness maximization.

This method is based on an idea of prof. dr. J. Portilla (DECSAI, Universidad de Granada, España), and has been worked out under his supervision during my stay in Granada. It relies on maximizing the sparseness of the wavelet coefficients in the finest resolution scales. To measure the sparseness of the wavelet subbands, the kurtosis is used. The kurtosis was used before to measure sparseness in Principal Component Analysis / Independent Component Analysis / Blind Source Separation [Comon, 1994, Cardoso and Laheld, 1996, Bell and Sejnowsky, 1997, Hyvärinen, 1999]. The sparseness of the wavelet representations is also exploited in almost all wavelet-based denoising methods (e.g., [Donoho and Johnstone, 1994, Mallat, 1999, Hyvärinen, 1999], to name just a few). As far as we know, this is the first attempt to use this measure for blur estimation. This technique resulted in a conference publication [Rooms et al., 2004].

5.3.1 Principle of Sparse-Max

The approach followed for this method has been to use a standard fast non-blind regularized deconvolution technique (we have used a simple Wiener implementation) to obtain image estimations for a wide range of values of the PSF parameter (the σ_{blur} of an isotropic 2D Gaussian function, in this case). Observing the features of those estimations, we search for a criterion that tells us when the image is “in focus.” “In focus” here has a different meaning than in a physical focusing system: in a physical focusing system, it is sufficient to judge when the image reaches maximal sharpness (e.g., has maximal high frequency energy content). However, when applying digital deconvolution, over-correction of the PSF effect in the presence of noise gives rise to overshoot and ringing of the image edges (see Figures 5.8 and 5.9).

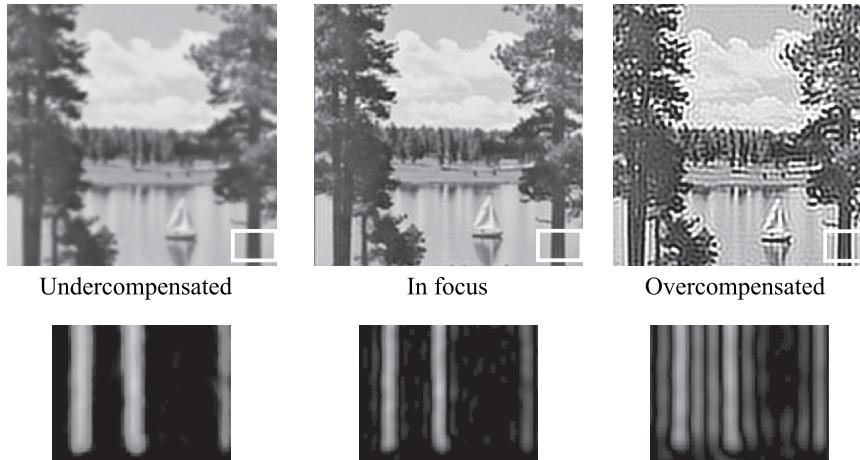


Figure 5.9: Physical versus numerical focus effect on a real image.

In this Section, we propose a wavelet-based criterion for choosing the optimal PSF parameter. For signals having strong localized discontinuities (such as edges, corners, . . . in images), the wavelet representations are powerful decorrelating tools. On the other hand, in the last decade, it has been widely studied how a mixing linear transformation (the PSF convolution, in our case) increases the mutual information between the mixed elements, and how applying Principal Component Analysis (PCA) to the mixed observations is not enough for recovering the original (unmixed) elements. In case of image blur, the convolution of the image with the PSF increases the mutual information between pixels because a pixel is replaced by a kind of weighted average of the pixel intensities in the neighborhood. A similar problem is blind source separation, where a number of linear combinations (mixes) of different signals are measured, and the goal is to retrieve the original signals.

That can be achieved, under certain conditions, by searching for the inverse of this mixing linear transform, so that the mutual information of the transformed observations is minimized (Independent Component Analysis, ICA). In practice, to achieve strictly that minimization is challenging, so several simpler approximated criteria have been proposed [Comon, 1994, Cardoso and Laheld, 1996, Bell and Sejnowsky, 1997, Hyvärinen, 1999]. In the case of images, it has been observed that wavelet coefficients are very “sparse” [Field, 1987, Mallat, 1989, Olshausen and Field, 1997], that is, the significant features are concentrated in relatively few coefficients (in this Section, we have used an à trous Haar wavelet, which is the simplest decorrelating multiscale transform). Loosely speaking, in this context a high sparseness of the wavelet coefficients is a sign of relatively low mutual information between them. Or, in other words, by maximizing a measurement of sparseness we decrease the mutual information. This is illustrated in Figure 5.10, where the edge in fo-

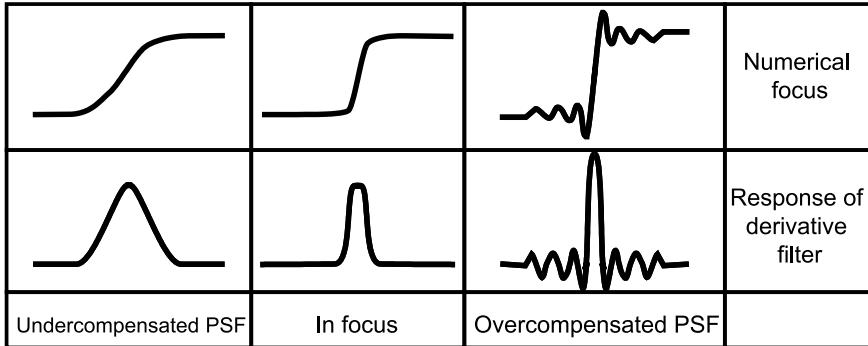


Figure 5.10: The three results for the edge profile in a numerical focus system (like in Figure 5.8), with their responses to a derivative filter. It is clear that the response is most concentrated when the edge is sharp (“in focus”).

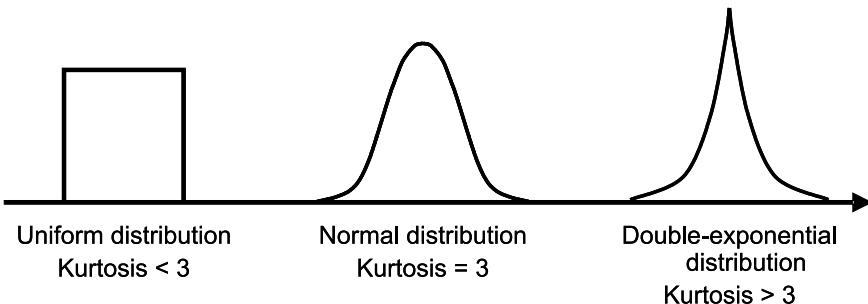


Figure 5.11: Examples of distributions and their kurtosis. It is clear in these cases that the higher the kurtosis, the more the distribution is concentrated around zero.

cus produces the sparsest response (i.e., with most coefficients equal or almost equal to zero) when a derivative filter is applied.

Using the previous background, this work is based on the two following observations:

1. blurring, as a local linear mixing, most often decreases the sparseness of the wavelet subbands;
2. overcompensating for the blur (that is, trying to invert a blur stronger than the actual blur) typically has a similar effect of reducing the sparseness of the wavelet responses.

This effect can be understood by applying a derivative filter (a similar effect to that of the wavelet) to the restored edges of Figure 5.8: the response is maximally concentrated when the image is “in focus.” As a means to characterize the sparseness [Comon, 1994, Olshausen and Field, 1997,

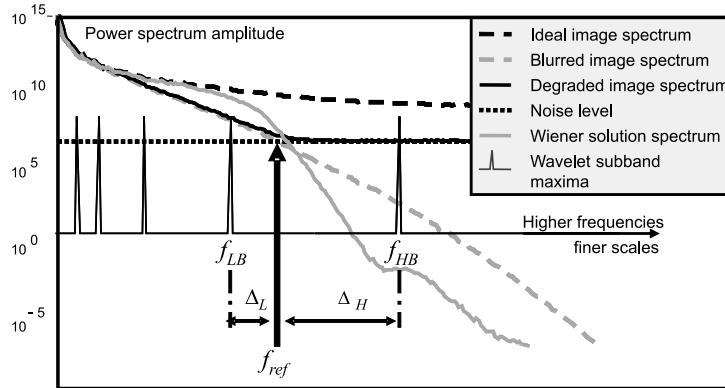


Figure 5.12: Reference frequency f_{ref} in the Fourier spectrum and choice of wavelet subbands to monitor for sparseness maximization.

Bell and Sejnowsky, 1997], we have used the kurtosis κ of each wavelet subband (see Figure 5.11, where some distributions and their kurtosis are shown). The kurtosis of a set data points y_i , $i = 1, \dots, N$ is defined as

$$\kappa = \frac{\sum_i (y_i - \mu)^4}{N\sigma^4}$$

with μ the expectancy value of the y_i and σ their standard deviation. We have tested empirically that when an image is not degraded, its wavelet subbands have typically maximal kurtosis, compared to the ones obtained with the same image convolved with a slight low-pass filter or its inverse (high-pass). Thus, the wavelet subbands of under- or oversharpened images have typically lower kurtosis values.

We would like to use as a reference the highest frequency subbands that are still relatively noise-free. The reason for preferring the high-frequency subbands instead of the low frequency subbands is that we want to use wavelet subbands as much affected by the image blur as possible. We have found empirically that a good compromise is to choose the two scales around the frequency f_{ref} for which the observed signal power spectral density is 4.5 dB above the noise level (Figure 5.12). We denote the finest of these two adjacent scales by HB and the coarsest by LB ; their maximal bandpass frequencies fall respectively at f_{HB} and f_{LB} . Then we checked which values of the PSF parameter maximized the kurtosis in these two neighboring scales after Wiener restoration with this candidate PSF. To calculate the kurtosis for a certain scale, the wavelet subbands with the horizontal and vertical detail coefficients were merged into one vector (the diagonal wavelet subband was discarded). If we introduce following notations for the wavelet coefficients of these two scales of the wavelet decomposition of the degraded image after Wiener restoration:

- $w_{HB,wiener,\sigma_{blur}}(x, y)$ are the wavelet coefficients of scale HB of the degraded image which is restored with a Wiener filter with a PSF with a certain σ_{blur} .
- $w_{LB,wiener,\sigma_{blur}}(x, y)$ are the wavelet coefficients of scale LB of the degraded image which is restored with a Wiener filter with a PSF with a certain σ_{blur} .

then we can write

$$\begin{aligned}\sigma_{max,HB} &= \arg \max_{\sigma_{blur}} \kappa(w_{HB,wiener,\sigma_{blur}}(x, y)) \\ \sigma_{max,LB} &= \arg \max_{\sigma_{blur}} \kappa(w_{LB,wiener,\sigma_{blur}}(x, y))\end{aligned}$$

Then we took a weighted average of $\sigma_{max,HB}$ and $\sigma_{max,LB}$, giving more weight when f_{ref} is closer to the maximum of the bandpass of a wavelet scale (see Figure 5.12). Δ_L is the difference in frequency between f_{LB} and f_{ref} , while Δ_H is the difference in frequency between f_{HB} and f_{ref} (also illustrated in Figure 5.12):

$$\begin{aligned}\Delta_L &= f_{ref} - f_{LB}; \\ \Delta_H &= f_{HB} - f_{ref}.\end{aligned}$$

So the weighted average to obtain $\sigma_{blurest}$ becomes:

$$\sigma_{blurest} = \frac{\sigma_{max,LB} \Delta_H + \sigma_{max,HB} \Delta_L}{f_{HB} - f_{LB}}.$$

5.3.2 Evaluation and comparison

Like with Wave-Track, we have evaluated the method using the same test images, again degraded with different amounts of blur and noise. As mentioned before, we used a simple Wiener implementation to obtain image estimates for each value of the PSF parameter, which we swept through the interval $\sigma_{blur} = 0.4$ to $\sigma_{blur} = 6.0$ with discrete steps of $\Delta\sigma_{blur} = 0.2$. The range of noise values considered in our experiments was $\sigma_{noise} = 0.1$ to $\sigma_{noise} = 10.0$. Sparse-Max provided satisfactory results for this range of blur and noise levels.

We implemented Savakis' method [Savakis and Trussell, 1993a] as a method for comparison. This method determines which $h_i(x, y)$ from a set of given candidates PSF's minimizes the difference between the spectrum $Q_{R_i}(u, v)$ of the image restored with the candidate PSF and the spectrum $Q_E(u, v)$ as it is supposed to be after Wiener restoration with the "right" PSF, under a prior

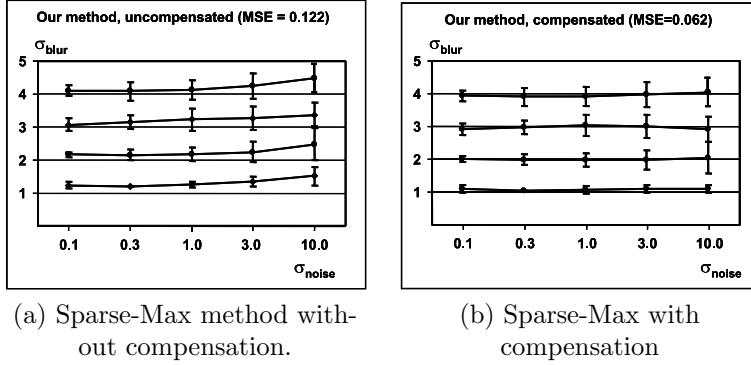


Figure 5.13: Summary of blur estimation results for Sparse-Max method uncompensated for the bias (left) and compensated (right). In each graph, a curve represents the estimations corresponding with the same $\sigma_{blur,real}$, but for different noise levels. The error bars correspond to $(\mu + \sigma, \mu - \sigma)$, with μ the sample mean and σ the sample standard deviation over the test set.

model for the image power spectral density. Savakis considered amongst others the *Chi square test*:

$$\chi^2 = \sum_{u,v} \frac{|S_E(u,v) - S_{R_i}(u,v)|^2}{S_E(u,v)}$$

where $S_E(u,v)$ and $S_{R_i}(u,v)$ are the respectively normalized spectra $Q_E(u,v)$ and $Q_{R_i}(u,v)$.

In the graphs in Figures 5.13 and 5.14, we show the average estimated blur values obtained for our set of test images (the same as the ones used to evaluate Wave-Track: *Airplane*, *Peppers*, *House*, *Flintstones* and *Lake*). Figure 5.14 shows the results obtained with Savakis' method, while Figure 5.13 shows the results obtained with Sparse-Max. The left graph in both figures shows the results of the method without any form of compensation for systematic biases in the estimation methods, while the right graph shows the results compensated to minimize the overall error. Savakis' method tends to underestimate the blur for all blur and noise levels, which can partially be corrected by adding 0.5 to σ_{blur} (when comparing the graphs in Figure 5.14, the compensated curves are closer to the real values). However, the errors in Savakis' method depend both on the blur and the noise level, and are therefore very difficult to compensate completely. On the other hand, our method tends to overestimate the blur systematically for higher noise levels (the reason why is still not clear, and has to be investigated further). Since we assumed the noise level to be known, we can compensate accordingly. We have hand-optimized the following heuristic expression to compensate for this overestimation as a function of the noise level

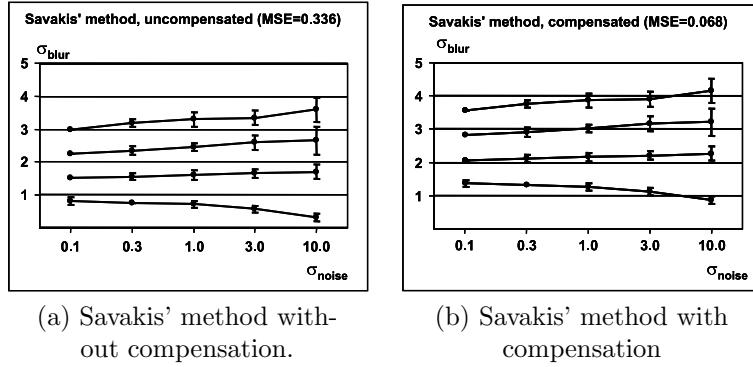


Figure 5.14: Summaries of blur estimation results for Savakis' method, uncompensated and compensated. The graphs were organized the same way as in Figure 5.13.

(see also Figure 5.13):

$$\sigma_{blur,compensated} = \sigma_{blur} + 0.0011\sigma_{noise}^2 - 0.0397\sigma_{noise} - 0.1594$$

Compared to Savakis' method, we have almost completely removed the estimation bias. This is an important advantage, which doesn't translate in a much lower MSE value (i.e., difference between the real blur parameter values and the estimated ones, over the range of blur and noise values), however.

To give an example of the robustness of the restoration result for small errors in σ_{blur} , we restored the *Lake* image. The real degradation parameters were $\sigma_{blur,real} = 2.0$, $\sigma_{noise} = 10.0$. We restored the degraded image with a Wiener filter using the real σ_{noise} , but swept σ_{blur} between 0.6 and 4.0. This way, we obtained a curve for the PSNR as a function of σ_{blur} (Figure 5.15). It shows a slow rise of the PSNR when σ_{blur} increases towards $\sigma_{blur,real}$. When σ_{blur} is close to $\sigma_{blur,real}$, the PSNR is more or less stable, which means that the restoration result is relatively insensitive for small errors of σ_{blur} in that region. When σ_{blur} is increased even further, the value of the PSNR drops steeply. In Figures 5.16 and 5.17, we show some restoration results on the *Lake* image using the estimated blur values, for two extreme situations (Figure 5.16: low blur, high noise level; Figure 5.17: high blur, low noise). Note how the PSNR values are close to the ones obtained using the real blur values.

5.3.3 Conclusion

We have used a simple and robust criterion for estimating a single-parameter PSF function from a single observation of a degraded image. Using the kurtosis of a set of selected subbands as the criterion to be maximized, we have obtained satisfactory results for several examples of natural images and for a

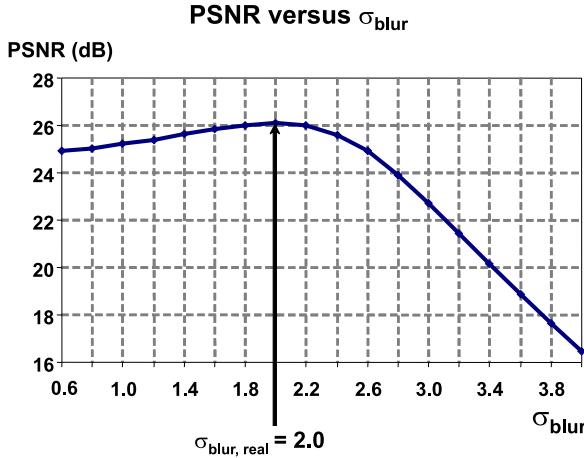


Figure 5.15: An example of the PSNR as a function of σ_{blur} . We artificially degraded the *Lake* image with $\sigma_{\text{blur}, \text{real}} = 2.0$ and $\sigma_{\text{noise}} = 10.0$. Then, the image was restored with the Wiener filter using a range of candidate PSF's for which σ_{blur} was swept between 0.6 and 4.0 with step size 0.2 (but using the real value for σ_{noise}). For each candidate σ_{blur} , the PSNR of the restoration result was computed.

wide range of blur and noise levels. Sparse-Max is robust, provides a reasonably small variance, and a very small bias (in contrast with previous methods (e.g., [Savakis and Trussell, 1993a]). This more recent method also outperforms Wave-Track in terms of MSE. As a by-product, this method provides a semi-blind Wiener estimation of the original image, which, obviously, could be greatly improved by using the estimated PSF with a more powerful restoration method.



(a) Degraded image
 $(\sigma_{blur,real} = 1.0, \sigma_{noise} = 10)$
 PSNR = 25.9dB



(b) Restored image
 with $\sigma_{blur,est} = 1.2$ PSNR = 28.7dB
 with $\sigma_{blur,real}$ PSNR= 28.9dB

Figure 5.16: Case example (low blur, high noise) of semi-blind Wiener restoration. It shows the degraded image (a) and as restored with our PSF estimation (b). The PSNR values (dB) are given for both our restoration and the restoration using the real blur.



(c) Degraded image
 $(\sigma_{blur,real} = 4.0, \sigma_{noise} = 0.1)$
 PSNR = 22.0dB



(d) Restored image
 with $\sigma_{blur,est} = 4.1$ PSNR = 25.2dB
 with $\sigma_{blur,real}$ PSNR= 25.3dB

Figure 5.17: Case example (high blur, low noise) of semi-blind Wiener restoration. The PSNR values (dB) are given for both our restoration and the restoration using the real blur.

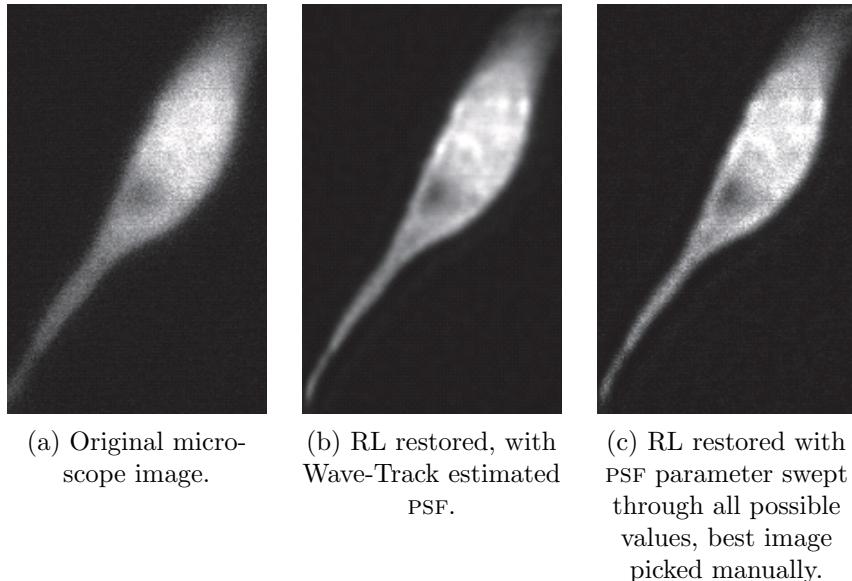


Figure 5.18: Blur estimation used in restoration of a real image.

5.4 Applications of blur estimation

5.4.1 Image restoration

In this Subsection, we give our first results of how the blur estimation was used as an input for classical (non-blind) image restoration techniques, in order to adapt them to unknown blur. In Figure 5.18, a confocal microscope image of a cell nucleus of *Arabidopsis Thaliana* stained with Green Fluorescent Protein (GFP) is shown. The left image shows the raw microscope image, the right image shows the image, restored with the Richardson-Lucy restoration algorithm (see Subsubsection 3.2.2.1), using the raw image and a synthetic PSF generated using the estimated parameter value as inputs. Chapter 6 deals in more detail with this integration into image deconvolution.

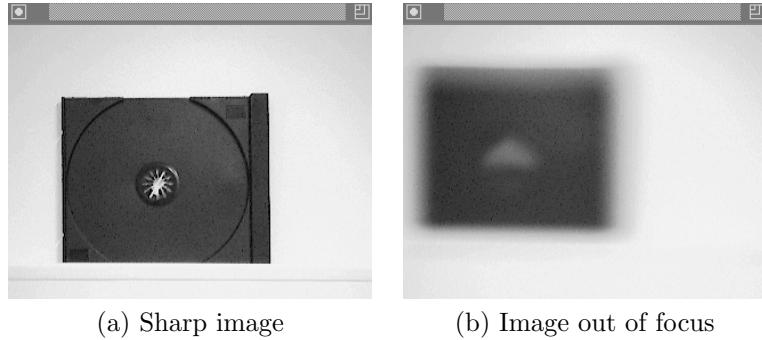


Figure 5.19: Autofocus illustration of a real image.

5.4.2 Autofocus

We also tried to estimate the PSF in case of out-of-focus blur. This kind of blur is encountered in autofocus applications, and is modeled by a uniform circular PSF [Pavlović and Tekalp, 1992, Savakis and Trussell, 1993a, Savakis and Trussell, 1993b].

$$\text{PSF}(x, y) = \begin{cases} K & \text{if } \sqrt{x^2 + y^2} < r_{focal} \\ 0 & \text{elsewhere} \end{cases} \quad (5.8)$$

with r_{focal} the radius of the focal spot and K a factor, chosen such that the norm of the PSF is 1.0. The r_{focal} is also sometimes referred to as the Circle of Confusion (COC) [Vivirito et al., 2002].

To estimate this r_{focal} from a captured image, we adapted Wave-Track for synthetic out-of-focus blur estimation. Adaptation of Sparse-Max can be done by generating uniform disc PSF's as candidate PSF's to maximize the sparseness criterion.

For out-of-focus blur, equation (5.3) of Wave-Track is not valid. However, we found by trial and error that the following polynomial provided a good fit:

$$r_{focal} = 19.7 \text{ CG}^3 - 19.1 \text{ CG}^2 + 17.3 \text{ CG} - 2.3$$

Using this relation, we can estimate r_{focal} . This is illustrated in Figure 5.19. These images were captured from a Sony DFW-VL500 camera. The diameter of the focal spot was estimated to be nearly zero in the first image (which is in focus), and 10 pixels in the second image, which proved to be correct by manual inspection of the edge width.

In most autofocus applications, one doesn't estimate the PSF of the blurring, but one only tries to determine whether an image is in focus or not. As this is

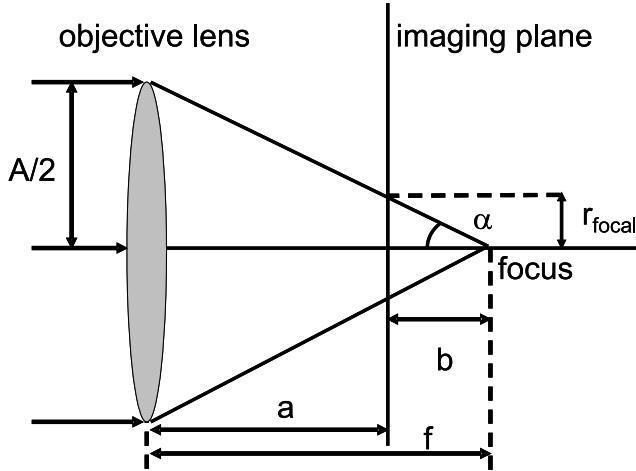


Figure 5.20: Autofocus application of blur estimation.

a real physical focus system, this can be checked by computing the total norm of the image gradient, as was explained in Subsection 5.3.1. When this norm is maximal, the system is in focus because its image is at its sharpest. However, this requires to check a whole range of focal distances, and the computation of the image gradient for every candidate focal distance. Therefore, it would be interesting to know how much the focal distance needs to be changed in a single step to obtain a sharp image.

In Figure 5.20, the possible usage of blur estimation in autofocus applications is illustrated: a is the distance from the lens to the imaging plane, which is at a distance b from the focus of the objective lens, which has an aperture A . The focal distance is then given by $f = a + b$, and one has to move the imaging plane by b in order to obtain a sharp image.

Nevertheless, it is possible to retrieve more information about the blurring, and to use it to adjust the focus more accurately. One can calculate the size of the focal spot r_{focal} and when one knows the focal distance f and the aperture A of the objective lens, one can calculate how much the focus needs to be adjusted from the following relation:

$$\tan \alpha = \frac{A/2}{f} = \frac{r_{focal}}{b} \Rightarrow b = \frac{r_{focal}}{\tan \alpha}$$

However, there's still an ambiguity whether the imaging plane is in front of or behind the focal plane. This can be solved by a second measurement of r_{focal} with a different focal distance, check the situation again, and correct the focal distance accordingly. A critical remark must be added nevertheless: many modern camera lenses suffer from considerable amounts of spherical aberration,

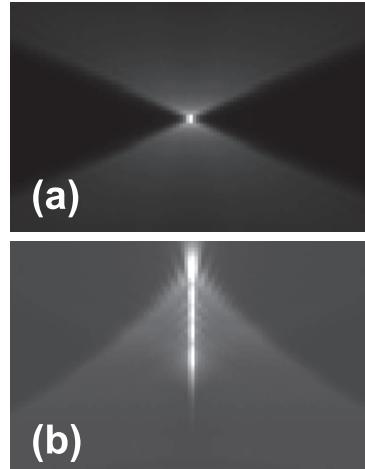


Figure 5.21: In (a), an XZ-section of an unaberrated PSF is shown; in (b), an XZ-section of a PSF is shown where severe negative spherical aberration occurs (in case of severe positive spherical aberration, image (b) is flipped upside down. Figures based on [Abu-Tarif, WWW].

causing an asymmetric blurring: on one side of the focus, the PSF is Gaussian shaped; on the other side of the focus, it is a ring with a central peak (see Figure 5.21). Since this can in principle be modeled with a single parameter, it should be possible to adapt the Wave-Track algorithm to this situation. In practice, I think that Sparse-Max would be more appropriate in this case of PSF's.

Further, Wave-Track (and also Sparse-Max) could also be used in depth from defocus estimation [RMA Belgium, WWW, Schechner and Kiryati, 1998]. If a camera with a fixed focal length is used, the image of the object placed at P produces a sharp picture of the object in the focal plane of P (see Figure 5.22). However, another object placed at Q is imaged unsharp in the focal plane of P , but is imaged sharply in the focal plane of Q . The larger the distance of an object from its position where it is imaged sharply, the more the object is blurred in the image. The Depth from Defocus method uses the relation between the distance of the object from the objective lens, the optical parameters of the camera and the amount of blurring of features at different positions in images to estimate the depth. Because the blurring in an image can be caused by either the imaging or the scene itself (like the turbulence-related blur in Figure 5.23), at least two images taken under different camera configurations (different focus) are generally required to eliminate this ambiguity.

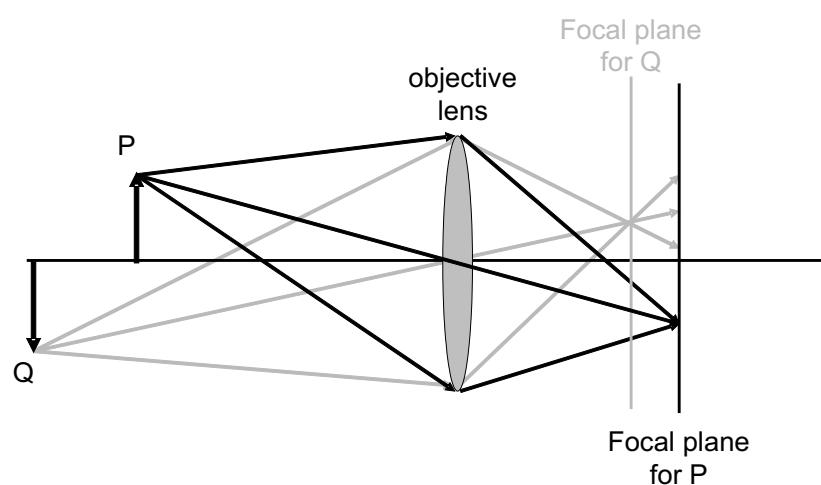


Figure 5.22: Illustration of principle of depth from defocus.



Figure 5.23: Illustration of blurring in a scene captured behind the wing of an airplane. A region of this scene is blurred due to turbulence of the hot gases from the jet engine of the airplane.

Chapter 6

SPERRIL: Estimation and restoration of confocal images

*“Begin at the beginning”, the King said, very gravely,
“and go on till you come to the end: then stop.”*
—Lewis Carroll, “Alice in Wonderland”

6.1 Introduction

In this Chapter, we describe our own original method for joint restoration and estimation of the degradation of confocal microscope images. The observed images are degraded due to two sources: first, we have blurring due to the band-limited nature of the optical system (modeled by the Point Spread Function or PSF); second: Poisson noise contaminates the observations due to the discrete nature of the photon detection process. The proposed method iterates between noise reduction, blur estimation and deblurring, and applies these steps in two phases, i.e., a training phase and a restoration phase.

In the first phase, the three steps are iterated until the blur estimate converges. Noise reduction and blur estimation are performed using steerable pyramids, while the deblurring is performed by the Richardson-Lucy algorithm.

The second phase is the actual restoration phase. In this phase, the blur estimate obtained from the training phase is used in the actual restoration, but new blur estimates are also made; the latter are only used to measure the progress of the restoration. The iterations are stopped when this measure converges, a result which is guaranteed.

This original integrated method is automatic, and no prior information on the image is required. We have nicknamed our method SPERRIL (Steerable Pyramid based Estimation and Regularized Richardson-Lucy restoration). Compared

with existing techniques by both objective measures and visual observation, SPERRIL performs comparable or even better. SPERRIL has been published in [Rooms et al., 2003c, Rooms et al., 2003b, Rooms et al., 2005].

Hybrid restoration methods (deblurring in the spatial or the Fourier domain, regularization in the wavelet domain) exist already quite some time [Neelamani et al., 1999, Boutet de Monvel et al., 2001, Figueiredo and Nowak, 2003]. Here, we tried to integrate the existing Richardson-Lucy algorithm for deblurring with the existing wavelet shrinkage algorithm of Šendur for regularization [Šendur and Selesnick, 2002] (he only used it for denoising). Our own blur estimation method Wave-Track was integrated in the method to estimate the PSF and to monitor the image sharpness during the iterations. Also, the Anscombe transform has already been used to deal with Poisson data [Starck et al., 2000], and has been applied in the field of astronomy, but rarely in microscopy (we only found one reference [Homem et al., 2002]).

This Chapter is organized as follows: first, the outline of our algorithm is presented. In particular, we discuss the noise reduction, the estimation of the image blur, the deblurring step and a stopping criterion. Next, experimental results are shown and discussed. Finally, a conclusion is given.

6.2 Outline of the algorithm

The outline of the algorithm is shown in a block diagram in Figure 6.1. First, the image is prefiltered to reduce Poisson noise. This prefilter is equivalent to the “denoise” steps further in the algorithm, except that the “denoise” is applied to the image transformed by the Anscombe transform to convert Poisson noise into Gaussian noise [Starck et al., 2000] instead of to the normal image. The “denoise” steps in the algorithm are performed by computing a steerable pyramid decomposition of the image and applying noise reduction to the subbands of the decomposed image. The next step is the estimation of the image blur from the (partially) denoised subbands. After denoising, the filtered subbands of the steerable pyramid are recombined, which produces a partially denoised, but not deblurred image. Finally, Richardson-Lucy deblurring is applied.

The estimate of the blur is unstable at first (therefore the estimated restored image obtained in this stage is unreliable and discarded in the actual restoration phase), but converges typically after two iterations and is then used to generate the PSF for the deblurring steps.

After this training process, the iterations are restarted with the original degraded image and the stable estimation of the PSF and the blur estimation is merely used as a quality measure to control the number of iterations: when it converges, the iterations are stopped. We now give a more detailed explanation

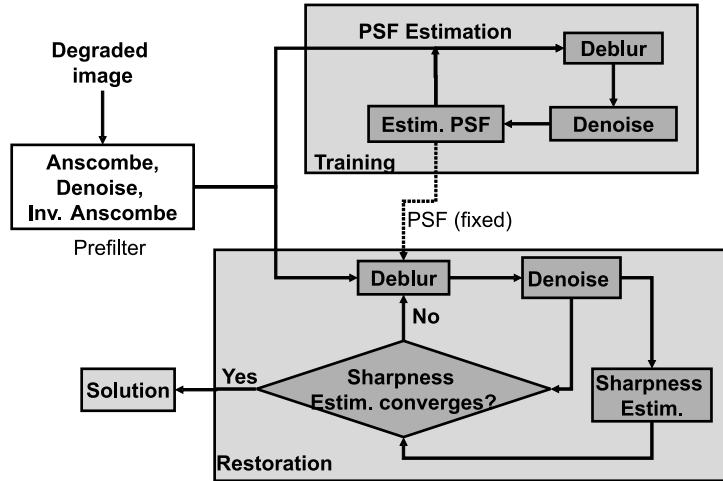


Figure 6.1: Block diagram of the algorithm. In the first block (the training phase), the PSF is estimated. This PSF is used as input for the restoration block. Both blocks consist of a denoise step (regularization) and a deconvolution step (deblurring). The PSF from the training phase is used unaltered, but a sharpness measure is calculated to determine when to stop the algorithm.

of the different steps of the algorithm. For more details about the steerable pyramid, we refer the reader to Subsection 4.2.2 of this thesis.

6.2.1 Noise reduction step

The wavelet-based denoising method of [Donoho and Johnstone, 1994, Donoho and Johnstone, 1995] is discussed in Section 4.3.2.1. This method has already been applied in microscopy by [Boutet de Monvel et al., 2001] and [Stollberg et al., 2003]. However, computation of the SURE threshold involves sorting the coefficients in the different subbands, which takes of the order of $N \log N$ operations, with N the number of coefficients in a subband [Mallat, 1999], which is pretty time-consuming.

This noise reduction is suboptimal for two reasons:

1. It assumes Gaussian noise. However, in confocal fluorescence imaging the major source of errors is Poisson noise [Pawley, 1995, van Kempen et al., 1997]. Unlike Gaussian noise, Poisson noise is intensity dependent, which makes separating image from noise very difficult.
2. It only exploits that useful coefficients should be large, and doesn't exploit information contained in spatial correlation between wavelet coefficients, and information contained in correlation across scales (see again Figure 4.36).

We approach these problems in the following way. To deal with Poisson noise, we used the Anscombe transform; see again equation (4.16) [Starck et al., 2000]. This transform converts the Poisson data into data with a Gaussian distribution with unit standard deviation. So, this transformation allows us to use well studied methods for Gaussian noise on data with the much trickier Poisson noise. We have chosen to apply the bivariate wavelet-based method (which has already been discussed in Subsubsection 4.3.2.1). We just resume the shrinkage formula and the declaration of the notations:

$$\hat{w}_1(x) = \frac{\left(\sqrt{v_1^2(x) + v_2^2(x)} - \frac{\sqrt{3}\sigma_n^2}{\sigma(x)} \right)_+ \cdot v_1(x)}{\sqrt{v_1^2(x) + v_2^2(x)}} \quad (6.1)$$

with $(a)_+ = \max(a, 0)$. In this expression $\hat{w}_1(x)$ represents a denoised wavelet coefficient in some resolution scale and orientation. This denoised coefficient is calculated from the corresponding noisy coefficient $v_1(x)$ in the same resolution scale and orientation, and its parent $v_2(x)$, which is the wavelet coefficient at the same spatial position as $v_1(x)$ and the same orientation band, but in next lower frequency resolution scale. σ_n^2 denotes the noise variance, and $\sigma(k)$ denotes the local standard deviation of the wavelet coefficients in the neighborhood of the coefficient at position x , (in practice, the standard deviation of the wavelet coefficients of the same subband in a 7x7 window around the coefficient at position x is used).

As we already mentioned before, this algorithm is simple to implement, has low computational cost (involves only order N operations per subband, with N the number of coefficients per subband) and yet provides noise reduction, comparable with other recent wavelet-based denoising techniques which also capture both inter- and intrascale dependencies of the wavelet coefficients with more complex models [Pižurica et al., 2003, Portilla et al., 2003].

The reasons why the noise reduction of [Sendur and Selesnick, 2002] outperforms that of [Donoho and Johnstone, 1994, Donoho and Johnstone, 1995] are

- it adapts locally to the presence of edges due to local estimation of $\sigma(x)$
- the expression $v_1^2(x) + v_2^2(x)$ captures some of the correlation between coefficients across scales.

We adapted this algorithm for implementation with steerable pyramids. The first time, the denoising step is combined with the Anscombe transform as a prefiltering step to reduce Poisson noise. Later, the denoising step is applied as regularization step (without the Anscombe transform) after each deblurring step.

6.2.2 Blur estimation

There are two common methods to determine the PSF of a confocal microscope. The first method relies on the physical model of the image formation. When all parameters like numerical aperture of the objective lens, the refraction index, ... are known, the PSF can be calculated, as described in Section 2.3. However, this method does not take into account optical misalignments due to temperature fluctuations, variations of the refractive index inside the specimen, ... (see Subsection 2.3.4, where the effect of aberrations on the image formation is discussed). These aberrations can cause deviations between the calculated PSF and the real PSF. Some of these deviations can be modeled, but they are not always known in detail.

Another method is imaging a small fluorescent bead with a size of the order of the resolution limit of the microscope. Methods using this bead image as a starting point to construct a PSF model have been proven to be very reliable and to allow better restoration results than with a theoretical PSF [van der Voort and Strasters, 1995, McNally et al., 1998], since not all possible causes of image aberration can be taken into account into a theoretical model. However, this bead imaging is very time consuming (one has to re-image these beads regularly as equipment calibration degrades in time) and researchers do not always take the time to do this for various reasons (at imaging time no intention to restore afterwards, no time, ...).

Therefore, we used Wave-Track, our own method for blur estimation as described in Section 5.2, because it does not require calibration images and better adapts to the image degradation than the method with the theoretical PSF. The method operates in the steerable pyramid domain (so no extra transforms are needed after the denoising/regularization).

In Wave-Track, we assumed the PSF to be modeled by a 2D Gaussian function. One can of course doubt how useful a Gaussian PSF model is in case of microscopy, since this type of PSF is not band limited (the Fourier transform of a Gaussian is again a Gaussian, which only becomes zero at infinity), and an optical system is definitely band-limited. This is indeed true when the image is degraded due to blur only. However, no imaging system is free of noise processes. The noise level in fact limits the band of useful frequencies in the degraded original. This is illustrated in Figure 6.2, where the radially averaged (RA) spectrum of an ideal test image, the RA spectrum of this image blurred with a Gaussian PSF and the RA spectrum of the blurred image, after which a Poisson noise process was simulated. The figure illustrates that for an image that is just blurred with a Gaussian PSF but without any noise, in principle all frequencies could be recovered. However, as soon as the image is contaminated by noise, an upper limit is imposed to which frequencies can be recovered (i.e., where the RA spectrum of the degraded image becomes flat). However, for the passband where the signal exceeds the noise level, the Gaussian PSF model is a reasonable approximation.

Nevertheless, there are some undeniable advantages to this approach:

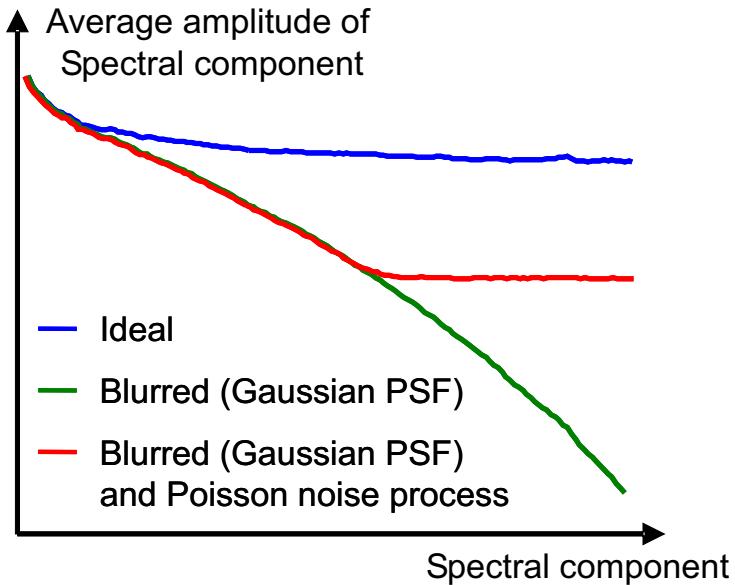


Figure 6.2: RA spectra of ideal test image (blue line), of the image blurred with a Gaussian PSF (green line) and of the blurred image, after which a Poisson noise process was simulated (red line). The vertical axis is scaled logarithmic.

- it does not require a calibration image of a bead to model the PSF;
- it adapts to PSF degradations due to aberrations, without assuming explicit knowledge of these aberrations.

Of course, when a real measured PSF is available, it can be used in our algorithm; in that case, we don't have to generate a synthetic PSF and our sharpness estimation can still be used to monitor the sharpening of the image.

6.2.3 Deblurring step

For the actual deblurring step, we used the Richardson-Lucy (RL) algorithm (which is discussed in Subsubsection 3.2.2.1). This algorithm is already common practice in confocal image restoration [Holmes et al., 1995, van Kempen et al., 1997, Verveer et al., 1999, Schaeffer et al., 2001].

Since the standard Richardson-Lucy algorithm is obtained by maximizing $L(f)$, no explicit regularization is applied. However, numerous authors have suggested different regularization schemes for the algorithm [Verveer, 1998, Verveer et al., 1999, Conchello and McNally, 1996]. One particularly simple regularization approach that provides remarkably nice results, is to apply a slight blurring after every RL iteration (e.g., with a Gaussian PSF with σ_{blur}

$= 0.5$). This regularization by postblurring is replaced in our algorithm by the wavelet-based denoising approach described in Subsection 6.2.1.

6.2.4 Stopping criterion

In this Subsection, we briefly discuss when to stop the iterations. We noticed that the estimation of the blur decreases with the number of iterations (as expected), as the likelihood of the solution is increasing. As an empirical rule, we stop the iterations when the ratio of the blur estimated during the previous iteration $\sigma_{blur,t-1}$ and the current iteration $\sigma_{blur,t}$ is converging to 1, i.e., when

$$\left| \frac{\sigma_{blur,t}}{\sigma_{blur,t-1}} - 1 \right| < \varepsilon \quad (6.2)$$

with ε a user-defined threshold. This criterion is the last ingredient of SPERRIL, resulting in a method that estimates the degradation from the original image (so no extra calibration is necessary) and automatically determines when the restoration is complete.

6.2.5 SPERRIL as RL with a prior?

This short Subsection tries to put the SPERRIL algorithm in terms of the Bayesian MAP framework. At first, we considered this algorithm to be some pure Richardson-Lucy restoration with a hard constraint applied after every iteration, like the positivity constraint in the ICTM algorithm. A similar approach was explained in [Starck et al., 2000] p80: the wavelet coefficients were used to extract the useful features after every iteration (without assuming a real prior model).

As far as we know, the first paper that applied this hybrid ad hoc combination of an existing deconvolution method with a wavelet-based regularization was [Neelamani et al., 1999]. Later, it was refined and founded more theoretically in [Neelamani et al., 2004]. Shortly after SPERRIL was born [Rooms et al., 2003d, Rooms et al., 2003c], papers appeared with attempts to find a more theoretical foundation for these hybrid algorithms. One of these papers was [Figueiredo and Nowak, 2003], where a method was derived from the Bayesian MAP framework, that resulted in a method that is very similar to ours (except that the deconvolution method is a Landweber-like algorithm, and the wavelet regularization is based on a very simple l_1 penalty term to induce sparseness of the wavelet coefficients).

The RL algorithm is a method that maximizes the likelihood of recovering data that are corrupted by a Poisson process, as expressed in Equation 3.26.

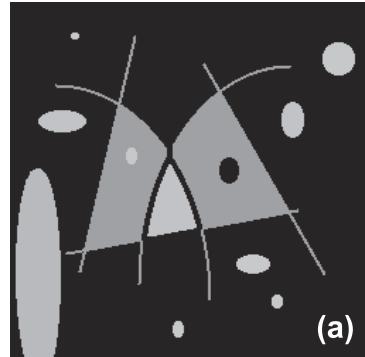


Figure 6.3: The ideal synthetic test image (a).

Sendur's method [Sendur and Selesnick, 2002] on the other hand maximizes the probability that clean wavelet coefficients follow a bivariate parent-child conditional probability function. Therefore, it should be possible to formulate a functional that maximizes the probability the Poisson probability and the prior probability that assumes a bivariate model for a “clean” solution. However, due to the fact that the likelihood functional is imposed to the spatial domain and the prior in the wavelet domain, this combination is not trivial and will be the subject of further research. Inspiration to combine both is in [van Kempen, 1999] p.47 and in [Figueiredo and Nowak, 2003].

6.3 Experimental results

6.3.1 Synthetic image

A first experiment was the restoration of a synthetic image with simulated degradation. Here, Figure 6.3 is the ideal image and Figure 6.4 (b) is a simulated degraded image, where first blurring was simulated by convolution with a synthetic PSF and then Poisson noise was applied (as is common use in synthetic experiments [van Kempen, 1999, Verveer, 1998]). In our case, a Gaussian PSF was used. In Figure 6.4 (c) the result with classical RL (with slight Gaussian blur applied after every iteration as regularization) is shown.

Figure 6.4 (d) is the result of restoration with RL-SURE [Boutet de Monvel et al., 2001], a deconvolution method with wavelet-based SURE regularization (see Subsubsection 4.3.2.1). This method doesn't perform optimally here, since SURE thresholding was designed for Gaussian noise and applies the same threshold for all coefficients in one wavelet subband, which means that it doesn't take into account the presence of local image features like edges. The SURE threshold is too large in background areas (since Poisson noise has a low variance in low intensity regions) and too small in bright areas

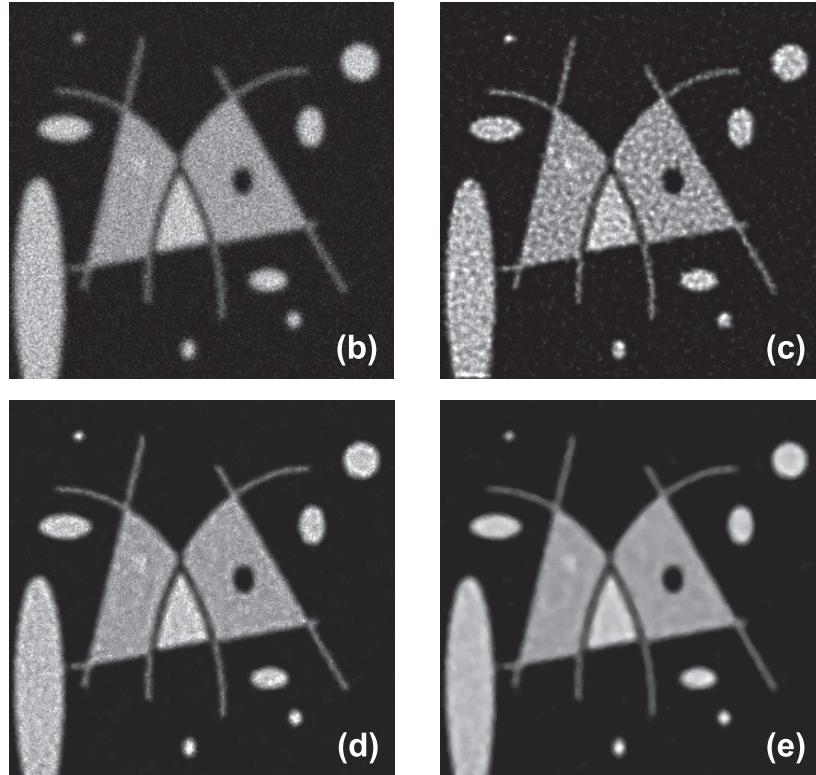


Figure 6.4: Comparison of different restoration algorithm on a synthetic test image. Image (b) shows simulated degradation (blurring with a Gaussian PSF of σ_{blur} of 2 pixels and simulated Poisson noise). In (c), Richardson-Lucy was applied, with slight Gaussian reblurring after each iteration as regularization; (d) is the result of the RL-SURE. Finally, (e) is the result of SPERRIL.

(since Poisson noise typically has a large variance in high intensity areas). We can see this in the restoration results because the noise in the background has been removed, but is still prominent in the bright image structures. Poisson noise could also be taken into account in this case by using the Anscombe transform here, but we only describe the technique as it was explained in [Boutet de Monvel et al., 2001].

Figure 6.4 (e) is the result with SPERRIL, which explicitly takes into account the Poisson nature of the noise (through the Anscombe transform). Notice also here that noise is best suppressed (less artifacts), resulting in homogeneous regions where the original image was homogeneous. For this image, the PSF was known and was used in all algorithms used in the comparison. In SPERRIL however, we used our own PSF estimation. So in fact, we compared our blind method with non-blind methods that used the correct PSF, which was a disadvantage

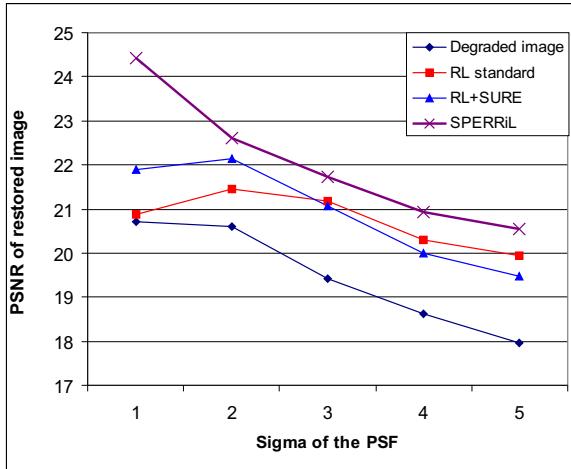


Figure 6.5: PSNR values for restoration of the synthetic test image for different amounts of blur. The image was blurred with a Gaussian PSF with different widths (σ_{blur} of the PSF), and each time, Poisson noise was simulated.

for our technique. Nevertheless, SPERRIL passed the test. For this experiment, also the ideal image was available, so we could calculate the Peak Signal to Noise Ratio (PSNR, see Chapter 3).

The values for the different restoration algorithms are shown in Figure 6.5. Using classical restoration shows already a significant increase in PSNR. With RL-SURE, there is a slight improvement (especially for low blur), and in case SPERRIL restoration is used, there's another slight increase in PSNR.

As a test for the resolution of the methods, we took a synthetic grid and applied synthetic degradation (Figure 6.6). The grid contains lines of different widths separated by different widths of background. The grid was degraded (Gaussian

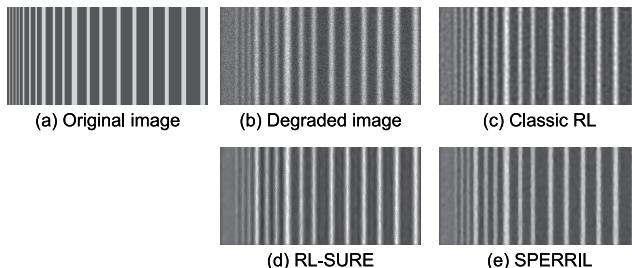


Figure 6.6: Restoration of a synthetic test-grid. (a) is the original grid, (b) is the degraded grid (blurred with a Gaussian with PSF, $\sigma_{blur} = 2$ pixels, and Poisson noise 120 photons per pixel). (c) is restoration with standard RL, (d) is the result of RL-SURE and (e) of SPERRIL.

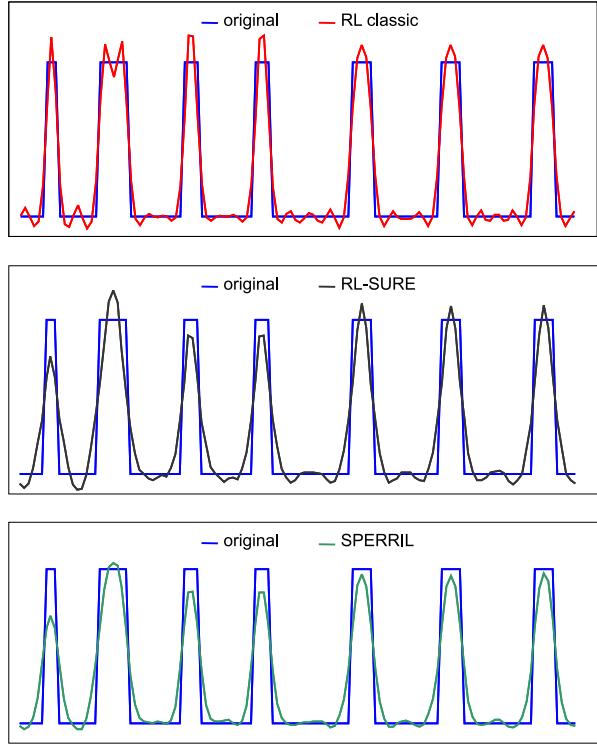


Figure 6.7: Line profile of Figure 6.6. The original grid is shown in blue. Note that classic RL (shown in red) tends to overcompensate the blur, giving higher intensities than original. RL-SURE (shown in black) overcorrects sometimes, and undercorrects in other cases. SPERRIL (shown in green) generally underestimates, but is closer on the peaks of the signal in most cases. Ringing is also less severe with SPERRIL.

PSF, $\sigma_{blur} = 2$ pixels, Poisson noise 120 photons per pixel) and restored with the different methods. In Figure 6.7, we plotted the line profile after restoration with the different methods. These line profiles were averaged over the different rows in the image. All methods indicate similar values for the Full Width at Half Maximum (FWHM) of the profiles, so in terms of resolution, the restoration methods are more or less comparable. This observation is to be expected, since the deblurring part is for all algorithms the same (i.e., Richardson-Lucy). However, we noticed that our method better suppresses the ringing in the background between the line structures, at the price of slightly wider tails of the profiles (which have a comparable width as with RL-SURE). Also, SPERRIL generally tends to slightly underestimate intensities, while the other methods generally tend to overestimate the highest intensities.

In Figure 6.8, we show a graph of the evolution of the PSNR as a function of the number of iterations during the restoration of our synthetic image (the one we always used in this thesis).

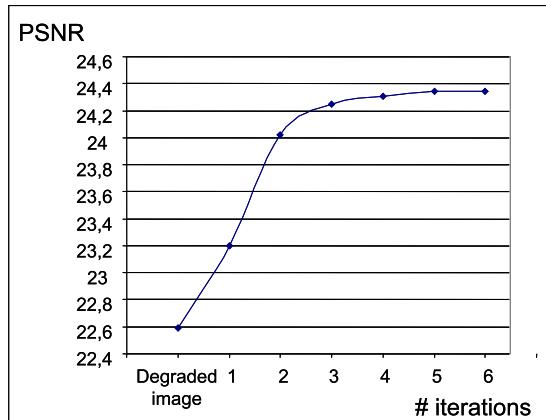


Figure 6.8: PSNR as a function of the number of iterations during SPERRIL restoration of a synthetic image ($\sigma_{blur} = 2$, 127 photons per pixel)

6.3.2 Biological experiments

6.3.2.1 Colocalization analysis: what and why?

To test the biological validity of any restoration algorithm is always challenging, because there is no so-called ground truth about the cell features of dimensions near the resolution limit of the microscopes. So, we have chosen to compare different image restoration algorithms as a preprocessing step before colocalization analysis, and evaluate the restoration results by the results of the colocalization analysis. In the following two paragraphs, the principle of colocalization is discussed. They are based on [Manders et al., 1993, Casavan and Gaidoukevitch, 2003].

Specimens can be labeled with one or more fluorochromes. Fluorochromes with a single duo of excitation/emission wavelengths are often used to highlight structures in the cell like microtubules (tube-like structures used for transport inside the cell, serves also like a kind of skeleton of the cell), mitochondria (for energy production in the cell) or nuclei (where the genetic information is stored). This is fine when a researcher is only interested in a specific structure within the field of view. However, often two or more labels are used to study many different structures at the same time, but keeping them separated (in different wavelength bands). The specimen is imaged at each fluorescent wavelength independently and then combined with the others.

Colocalization describes the presence of two or more types of molecules at the same physical location. The reason may be that the fluorochromes are attached to the same receptor, or are attached to molecules that are interacting with each other. In image processing, this translates to the fact that the intensities emitted by the fluorochromes contribute to the same pixel in the image.

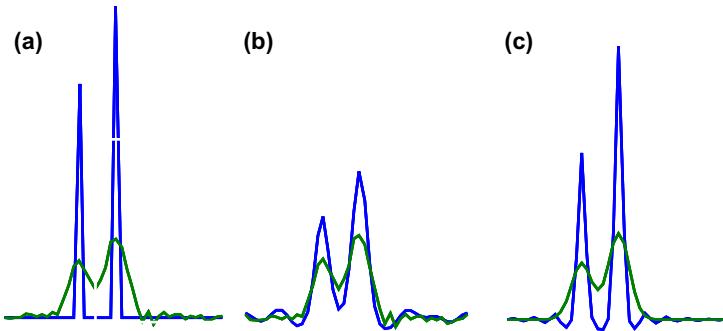


Figure 6.9: Degraded observation of separate objects. In (a), two separate objects are shown (blue) which are observed under non-ideal circumstances (green). In (b) we see that standard RL with slight reblurring (blue) already provides an improvement by eliminating part of the false overlap. In (c), we show the case where the two blurred peaks were restored in absence of noise (blue), thus completely separating the two peaks. The better the restoration technique, the closer we come to this case.

In [Landmann, 2002], the topic of deconvolution as preprocessing step to colocalization is discussed, and significant improvement in the analysis results were obtained when deconvolution was applied. In Figure 6.9, we illustrate the effect of image restoration on colocalization analysis. Two separate objects seem to have a certain overlap due to image blurring and the presence of noise in the image. When an appropriate restoration technique is applied, this false overlap is reduced, thus improving the accuracy of the colocalization analysis, which provides better results when the restoration technique is more reliable.

6.3.2.2 Experimental setup

This work has been carried out in collaboration with dr. Diane S. Lidke of the Department of Molecular Biology of the Max Planck Institute for Biophysical Chemistry (Göttingen, Germany), who kindly provided the images and feedback about the image processing results.

Here, we evaluate the performance of our algorithm on the colocalization analysis of two fluorescently tagged molecules in human A431 skin cancer cells [Lidke et al., 2004].

The first fluorescently labeled molecule is a transmembrane receptor protein, a molecule that spans the cell membrane (green in Figure 6.10), and can bind with molecules outside the cell to pass on information to the cell. The erbB family of receptors include the Epidermal Growth Factor Receptor erbB1 and erbB3 that, when activated, cause skin cells to grow. The erbB1 receptors were labeled with Green Fluorescent Protein (GFP), causing formation of erbB1-eGFP; erbB3 was labeled with mCitrine, which is an improved yellow fluorescent protein mutant.

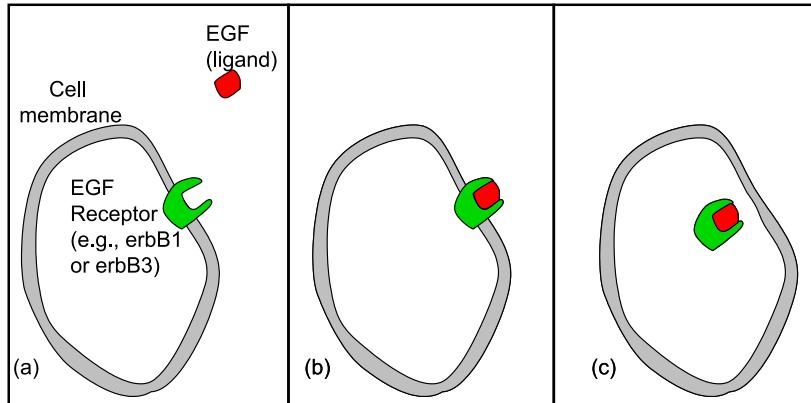


Figure 6.10: Schematic representation of the biological mechanism behind the colocalization: (a) The transmembrane receptor is drawn in green, the messenger ligand in red. In (b), binding of the ligand and the receptor is shown. In (c), internalization of the complex inside the cell is illustrated.

The second molecule is the so-called ligand (red in Figure 6.10), which is a molecule like a hormone that pass certain information to the cell by binding to the transmembrane proteins. The ligand used here is Epidermal Growth Factor (EGF). When the Epidermal Growth Factor Receptor binds with this ligand, the cell gets the signal to start growing. The ligand EGF was labeled with fluorescent Quantum Dots, thus resulting in EGF-QD.

One way to test if the receptor and the ligand interact, is to fluorescently label both molecules as described above, and monitor the colocalization (or overlap) of these two fluorescent signals in the image from the receptor and the ligand molecules [Lidke et al., 2004].

Figures 6.11 and 6.12 show the test set of images used in our analysis: Figure 6.11 (H1-H4) shows single confocal sections of living cells expressing erbB1-eGFP (green in these images) to which EGF-QDs (red in the images) have been added. After activation by EGF-QDs, the erbB1 will internalize, i.e. it is transported from the membrane to the inside of the cell (as illustrated in Figure 6.10). Since the QD-EGF remains attached to the erbB1-eGFP after internalization, the colocalization should be high.

In contrast, the cells in Figure 6.12 show A431 cells expressing erbB3-mCitrine. In this case, EGF-QD (red in images L1-L4) does not bind directly to erbB3 (green in L1-L4), but still binds to the native, unlabeled erbB1 (so invisible in L1-L4) present in the cell. Upon activation of erbB1 by EGF-QD, no complexes with the labeled erbB3 are formed; therefore in this image the colocalization should be low.

Our aim is to determine colocalization in these two cases experimentally, and to judge whether the results agree with what is expected. We will demonstrate

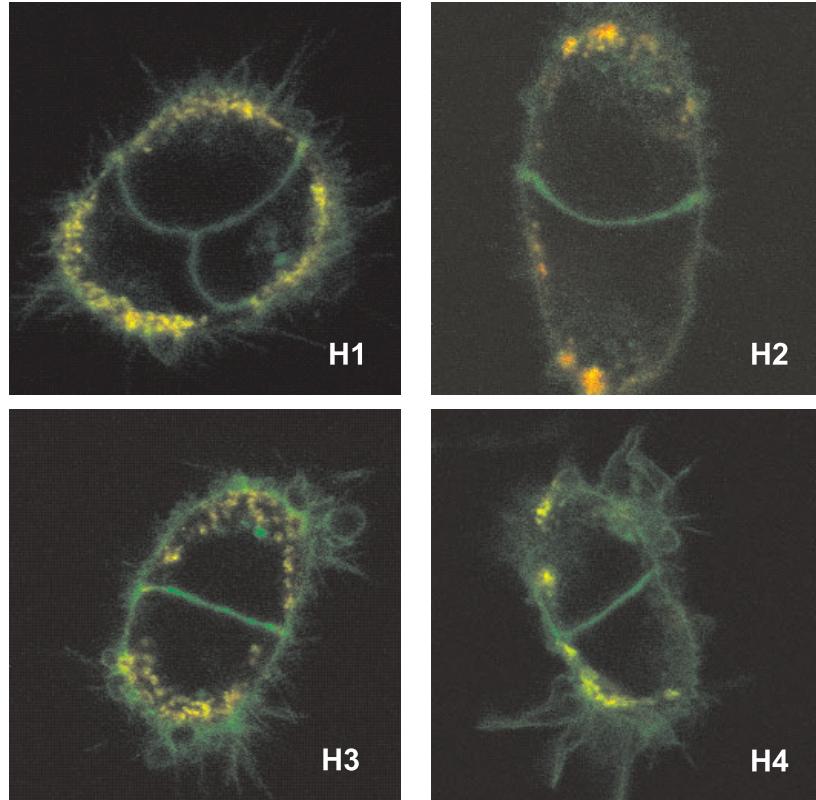


Figure 6.11: Set of test images (part 1) used in the colocalization analysis. In the cells in these images, erbB1-eGFP was visualized, thus resulting in high colocalization.

a better agreement after restoration.

Following [Landmann, 2002], we applied and compared the result of colocalization analysis on the raw image to the result of colocalization analysis applied after three different image restoration methods (classical RL, RL-SURE and SPERRIL) combined each time with standard background correction¹. The different color channels were each time processed independently for the sake of simplicity, and the PSF was estimated separately for the different color channels. In Table 6.14 and Table 6.15, we show the results of our colocalization analysis, which is based on the expressions in [Manders et al., 1993]. For the different methods, we calculate the overlap coefficient $R_{overlap}$, which is defined by:

¹with ImageJ [ImageJ, WWW], which is a versatile and multi-functional free software package with countless free plugins for different purposes.

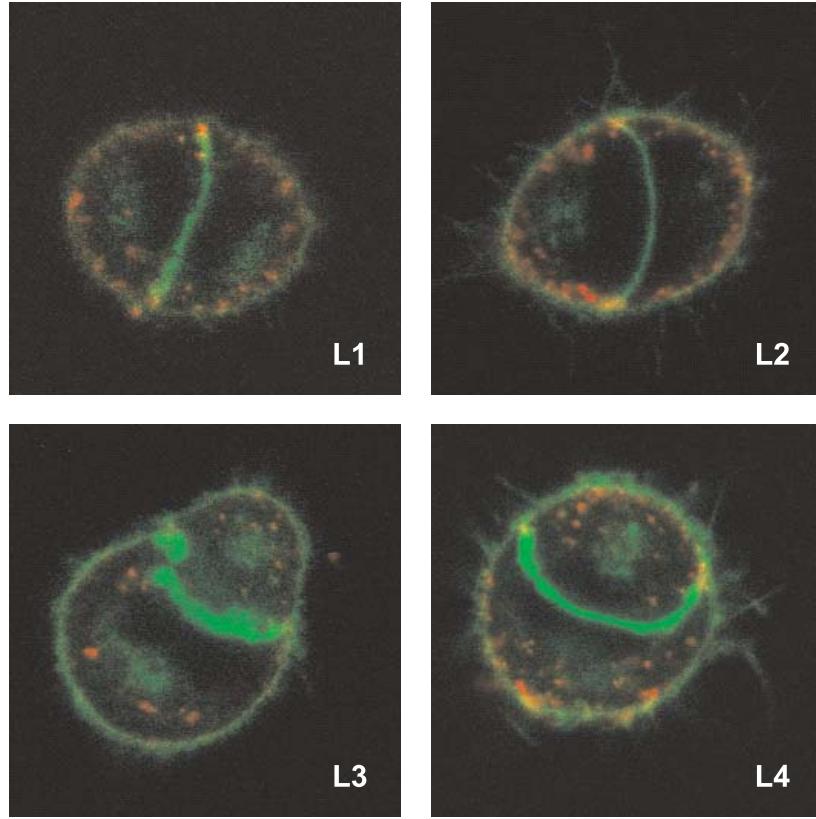


Figure 6.12: Set of test images (part 2) used in the colocalization analysis. In the cells in these images, erbB3-mCitrine was visualized, thus resulting in low colocalization.

$$R_{overlap} = \frac{\sum_{x,y} (I_1(x,y)I_2(x,y))}{\sqrt{\left(\sum_{x,y} (I_1(x,y))^2\right) \left(\sum_{x,y} (I_2(x,y))^2\right)}} \quad (6.3)$$

where $I_1(x,y)$ and $I_2(x,y)$ are the intensities in the red resp. the green channel. $R_{Pearson}$ is calculated in a similar way:

$$R_{Pearson} = \frac{\sum_{x,y} ((I_1(x,y) - I_{1,aver})(I_2(x,y) - I_{2,aver}))}{\sqrt{\left(\sum_{x,y} (I_1(x,y) - I_{1,aver})^2\right) \left(\sum_{x,y} (I_2(x,y) - I_{2,aver})^2\right)}} \quad (6.4)$$

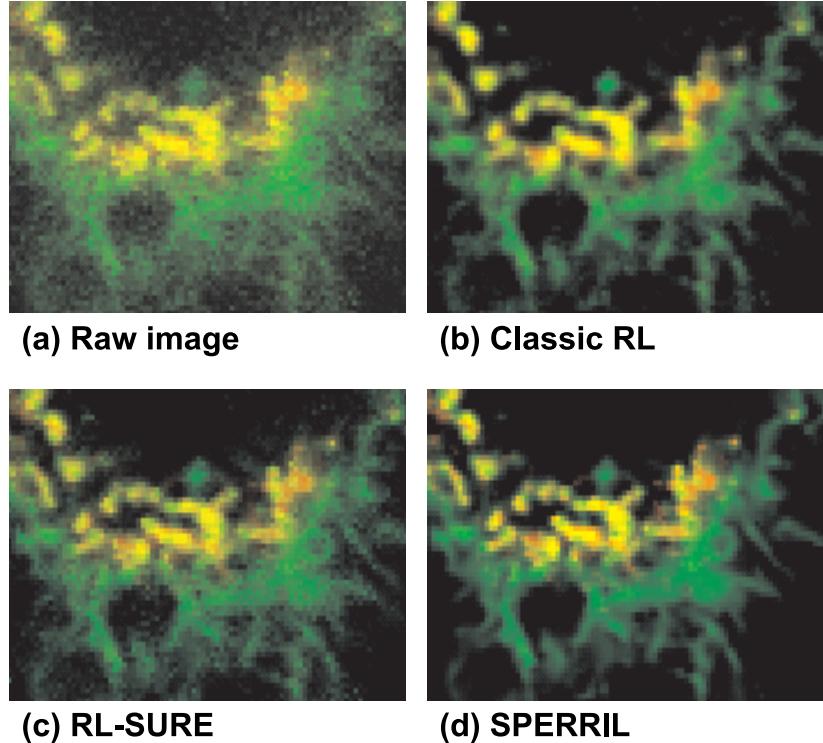


Figure 6.13: Detail views of the different restoration results for cell H3. In (a), the raw image is displayed; (b) is the results after classical RL; (c) is the result after RL-SURE and (d) is the result of SPERRIL. Note that (b) and (c) still contain noise, while (d) maintains similar sharpness and the noise is better suppressed.

Here, $I_{1,aver}$ and $I_{2,aver}$ are the average intensities of the first respectively the second color channel. This measure is related to the *correlation coefficient* [Weisstein, WWWb].

The $R_{overlap}$ and $R_{Pearson}$ values for the raw data (see Tables 6.14 and 6.15, plotted in a graph in Figure 6.16) show only a small difference between the two cell types. Restoration with the classical RL already improves the result of the colocalization analysis, in that the difference is clearer. However, when applying SPERRIL restoration prior to analysis, the largest difference between cell types is seen; this is consistent with what is expected from the underlying biochemical process in the cells [Lidke et al., 2004]. In Figure 6.13, a detail of cell H3 of Figure 6.11 is shown. Again, we can conclude that for classical RL, the regularization is rather poor. The results of RL-SURE also don't remove all the noise in the bright areas, while the result of SPERRIL provides a better suppression of the noise.

In Figure 6.17, we show a comparison of restoration of the same image as in

	Raw	Classic RL	RL-SURE	SPERRiL
H1	0.877	0.917	0.8674	0.8659
H2	0.685	0.771	0.6466	0.6622
H3	0.726	0.771	0.6828	0.6919
H4	0.758	0.872	0.7520	0.7382
average	0.761	0.833	0.7370	0.740
L1	0.532	0.604	0.5046	0.4525
L2	0.615	0.651	0.5057	0.4824
L3	0.434	0.544	0.3175	0.3344
L4	0.482	0.535	0.5519	0.3711
average	0.516	0.583	0.4700	0.4100
Δ	0.246	0.249	0.2670	0.3290

Figure 6.14: The $R_{Pearson}$ parameter obtained from the analysis the raw image, the image restored with classical Richardson-Lucy (classic RL), with RL-SURE and with SPERRIL. The values of $R_{Pearson}$ for the images with high colocalization (H1-4) are shown on top. Then, the values for the four images with low correlation are given (L1-4). In bold, the average of $R_{Pearson}$ over a certain class is given. Δ is the difference between the average for the group of images with high and the one with low colocalization. Δ is largest for SPERRIL, thus indicating that this method is best capable to separate low and high colocalization.

	Raw	Classic RL	RL-SURE	SPERRiL
H1	0.928	0.927	0.8785	0.8725
H2	0.868	0.807	0.805	0.6793
H3	0.831	0.798	0.7088	0.7183
H4	0.888	0.895	0.7543	0.7418
average	0.879	0.857	0.787	0.753
L1	0.832	0.670	0.5555	0.5076
L2	0.785	0.696	0.5519	0.5407
L3	0.685	0.626	0.3952	0.4285
L4	0.715	0.633	0.6102	0.4546
average	0.754	0.656	0.528	0.483
Δ	0.124	0.200	0.258	0.270

Figure 6.15: The $R_{overlap}$ parameter obtained from the analysis the raw image, the image restored with classical Richardson-Lucy (classic RL), with RL-SURE and with SPERRIL. The values of $R_{overlap}$ for the high colocalization images (H1-4) are shown on top. Then, the values for the low correlation images (L1-4) are given. In bold, the average of $R_{overlap}$ over a certain class is given. Δ is the difference between the average for the group of images with high colocalization and the group with low colocalization. Again, Δ is largest for SPERRIL, thus indicating that SPERRIL is most capable to separate low and high colocalization.

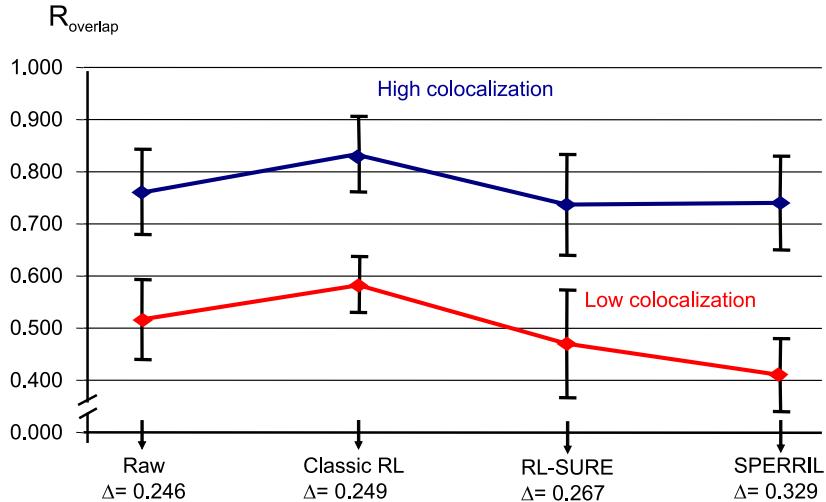
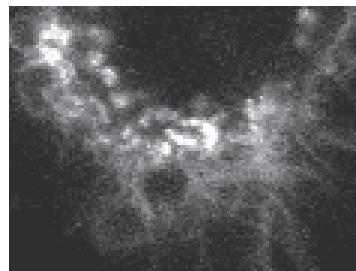


Figure 6.16: Results of Table 6.15 in a graph. For each class of cells (low/high colocalization) and for each restoration method, the sample mean is shown as a dot, and the sample variance is plotted as an error bar.

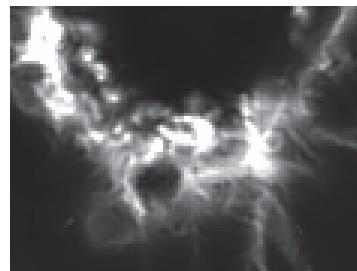
Figure 6.13, now only applying SPERRIL but using different PSF-models. In (a), we see the raw image; in (b), we show the result of restoration with a PSF as calculated theoretically according to what we explained in Section 2.3. Due to greater complexity when compared to WideField microscopes, confocal microscopes are susceptible to errors in alignment of the optical elements (thus introducing extra aberrations), which can occur due to, e.g., varying temperature in the room of the microscope in time. These may introduce severe shape distortions in the PSF [Van der Voort, WWW]. Since most confocal microscopes can suffer from these aberration effects, it is necessary to adapt the theoretical PSF for these distortions to obtain an accurate restoration result. However, here we used a simplified PSF model in (c) and (d), which yet adapts to the actual image blur (in contrast to a normal theoretically calculated PSF). This simplified adaptive model already provides satisfying results.

6.3.3 Typical running times

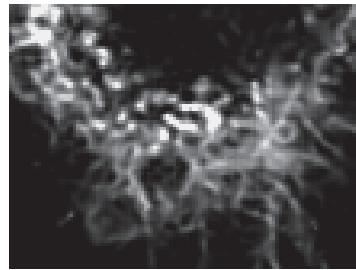
A typical run of SPERRIL takes 4 minutes for a 256 x 256 image (1 color channel) for 16 iterations. By comparison, the classical RL restoration schemes take a few seconds for the same number of iterations and the same image size. RL-SURE takes more than 40 minutes to run. The major part of this time is used to sort all the subbands of the steerable pyramid decomposition, although we are not sure if our implementation was programmed very efficiently (we used the standard C function `qsort`). All experiments were performed on an Athlon(TM) XP 2200+ with a CPU of 1.8 GHz.



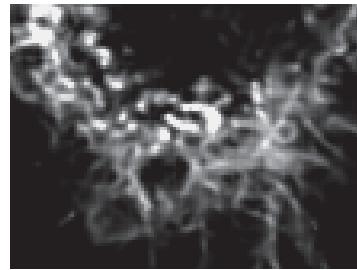
(a) Raw image



(b) Restored with theoretical PSF



(c) Restored with Gaussian PSF



(d) Restored with Airy PSF

Figure 6.17: Detail views of restoration results with different PSF's for cell H3. In (a), the raw image is displayed; (b) is the results after SPERRIL with a theoretically calculated PSF (as explained in Section 2.3); (c) is the result after SPERRIL with a Gaussian PSF and (d) is the result of SPERRIL with an Airy PSF. Note that (b) doesn't take into account optical misalignments etc, and should be rescaled (otherwise it performs suboptimally, like here). However, the difference between (c) and (d) is invisible.

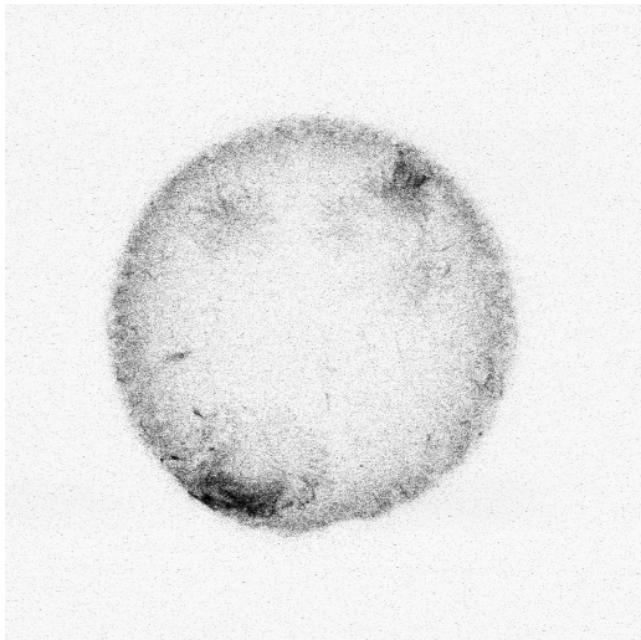
6.3.4 Some more visual results

In this Subsection, we show some more visual results. The first set of images were gray value images of ovary cells, provided and evaluated by prof. P. Van Oostveldt, dept. Molecular Biotechnology, faculty of Bioscience Engineering, Ghent University. These images represent ovary cells, and are shown in negative (black was brightest in the original images) to emphasize background structures.

Figures 6.18, 6.19 and 6.20 show labeled microtubuli in unfertilized ovary cells of mice. Microtubuli are microscopic filaments with a diameter of 24 nm (nanometer: $1\text{ nm} = 10^{-9}\text{m}$) which can bundle in more or lesser extent into larger structures under the influence of different external factors. They can be visualized by incubating antibodies with FITC attached. In young ovary cells, the microtubuli form a subcortical network, which has a strong dynamical structure. The images were recorded with a confocal microscope at a magnification of $0,278\mu\text{m} / \text{pixel}$ and image dimensions of 512×512 pixels (μm , or micrometer: $1\mu\text{m} = 10^{-6}\text{m}$). The objective lens that was used, had $NA = 1.4$, which results in a resolution of $0.61 \times 0.500\ \mu\text{m} / 1.4$ of $0.217\ \mu\text{m}$. A microtubulus of 24nm has a width below the resolution of a light microscope. The pixel dimensions are not sufficient to use the full resolution of the objective. An improvement of the images is visible, especially where bundles or fans of microtubuli are present. In general, the background has become less granular after restoration, which indicates that the SNR improved.

The second type of images (Figures 6.21, 6.22 and 6.23) is recorded under similar circumstances, but shows cortical granules (which are smaller vesicles that are located on the edge of the ovary cell). Also here a more homogeneous background is present after processing, but an improvement of the quality of strongly clustered structures is difficult to interpret. The question remains if weakly fluorescent objects don't disappear, and if they decrease in size. A qualitative impression is that most objects remain present, as interpreted from a false-color representation. There certainly is an effect on the area, which is clearly visible in clusters of particles.

Especially at the location of strong clusters of positive signals, it seems that both tubular as granular objects in the images improve strongly in quality, and have more contrast. There is no noticeable decrease of the number of objects present in the processed image compared to the raw image. If this can also be verified quantitatively, it should be concluded that this technique is promising, especially for live imaging. The more homogeneous background makes the image visually more pleasant, but doesn't produce a direct extra scientific value. It should be a fundamental improvement for additional segmentation or compression in later stages of automated processing and quantification.

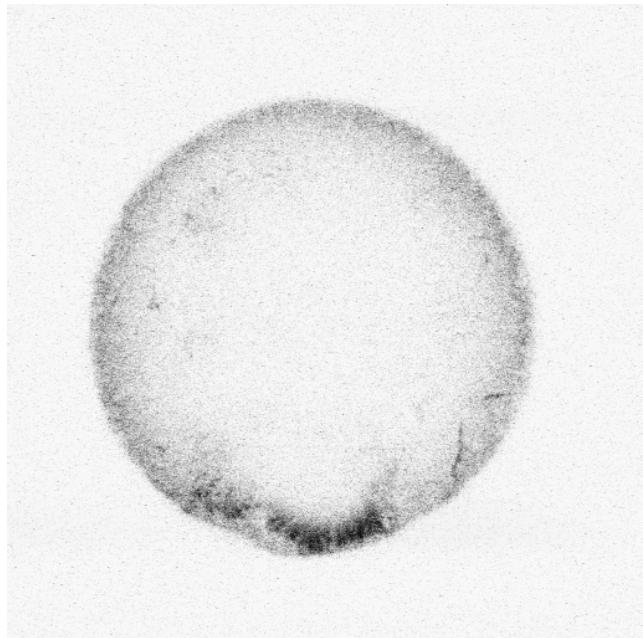


Raw image



Result of SPERRIL

Figure 6.18: Restoration of image of ovary cell: microtubuli (1).

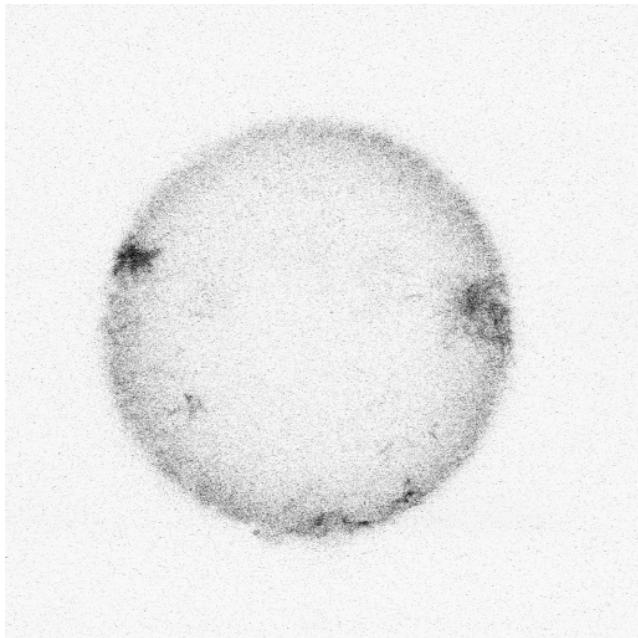


Raw image



Result of SPERRIL

Figure 6.19: Restoration of image of ovary cell: microtubuli (2).

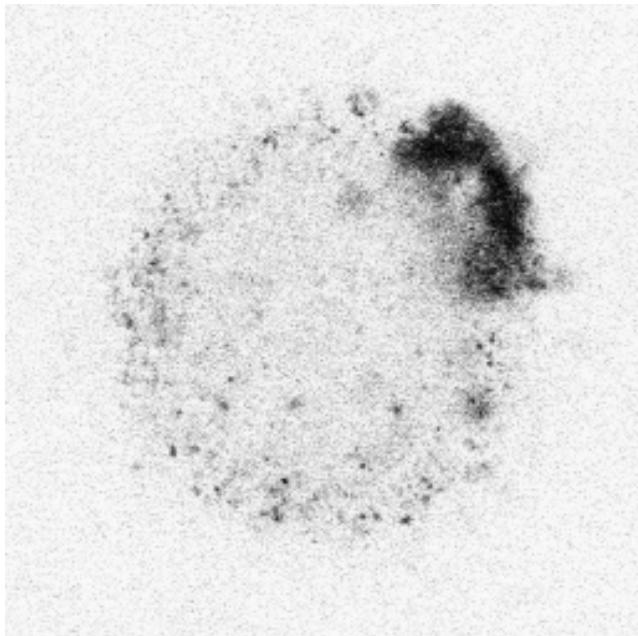


Raw image

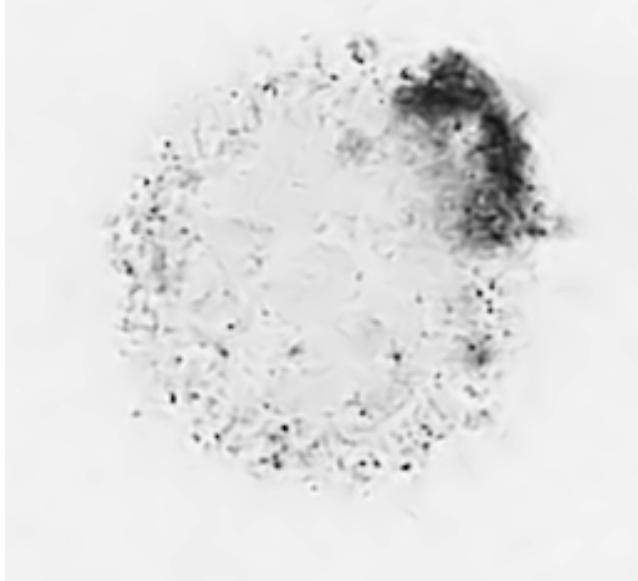


Result of SPERRIL

Figure 6.20: Restoration of image of ovary cell: microtubuli (3).

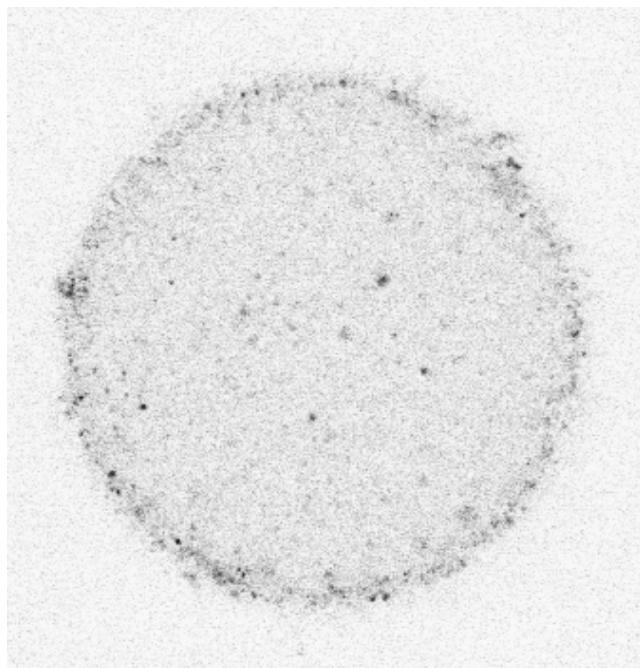


Raw image

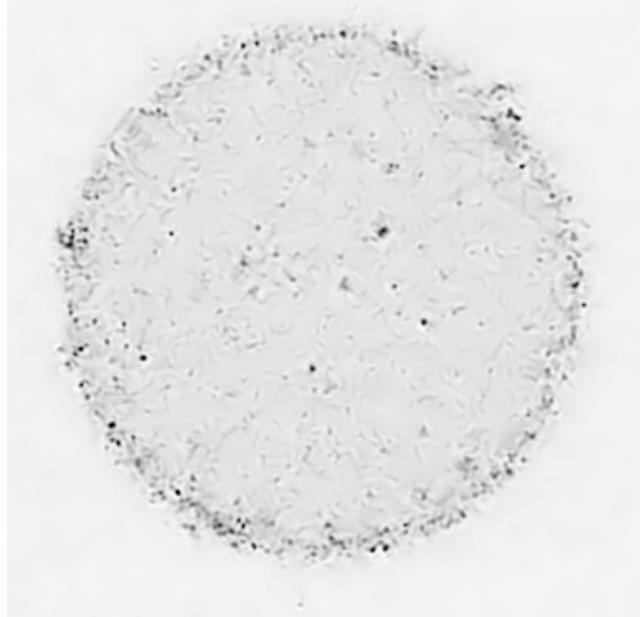


Result of SPERRIL

Figure 6.21: Restoration of image of ovary cell: vesicles (1).

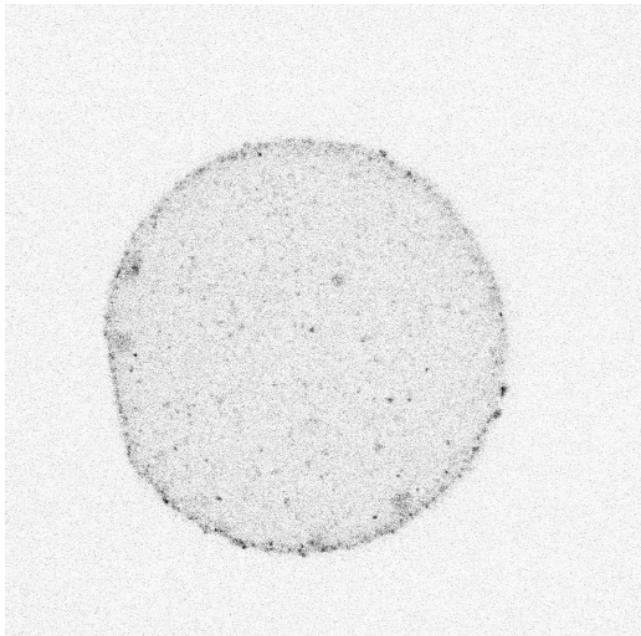


Raw image

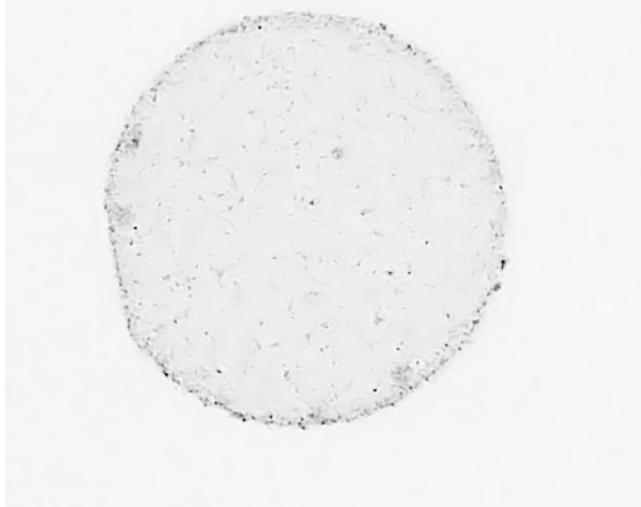


Result of SPERRIL

Figure 6.22: Restoration of image of ovary cell: vesicles (2).



Raw image



Result of SPERRIL

Figure 6.23: Restoration of image of ovary cell: vesicles (3).

Finally, we show a few more results (Figures 6.24 and 6.25). These images were provided by dr. Diane S. Lidke, Max Planck Institute for Biophysical Chemistry, Göttingen, Germany. These images were not included in the dataset used in Subsection 6.3.2.

6.4 Conclusions

In this Chapter, we presented an integrated, stable and automatic algorithm to restore degraded photon-limited images. The degradation parameters are estimated from the image itself without requiring a calibration image. This is an important advantage. The estimation as well as the regularization of the deconvolution are performed in the steerable pyramid domain. The sharpness estimation is used to formulate a stop criterion for the iterations, thus making the restoration automatic.

The algorithm was evaluated both on synthetic images and real confocal images. Our algorithm performs best, both in terms of visual quality as well as in terms of the PSNR. The algorithm was also evaluated on a set of biological images, where colocalization analysis was chosen as an objective way to evaluate the restoration result. In this case, we also compared our algorithm with other algorithms, and it performs best in terms of visual quality as well as in terms of the result of the colocalization analysis. In comparison with other restoration techniques, SPERRIL obtains promising results both on synthetic and real confocal images. Nevertheless, ringing still occurs quite strongly in the results of SPERRIL.

Future work involves extending SPERRIL to 3D. The Richardson-Lucy algorithm is already available to three dimensions. Also, the extension of the regularization in 3D is possible using 3D multiresolution decompositions, since 3D extensions of steerable pyramids are already available [Delle Luche et al., 2003], as demonstrated in Section 4.4.1. The biggest challenge would be the estimation of the PSF in 3D. The PSF in 2D can be approximated by a Gaussian. In the z-direction, this is obviously not the case. In that case, we would think of another, more realistic confocal PSF model where the parameters of the confocal PSF model can be adjusted to compensate any observed aberrations from the blur estimation (e.g., a certain blurring in the z-direction is observed by our method, but the theoretical model predicts a smaller deviation). Then the parameters of the confocal PSF model can be adjusted to take this deviation into account.

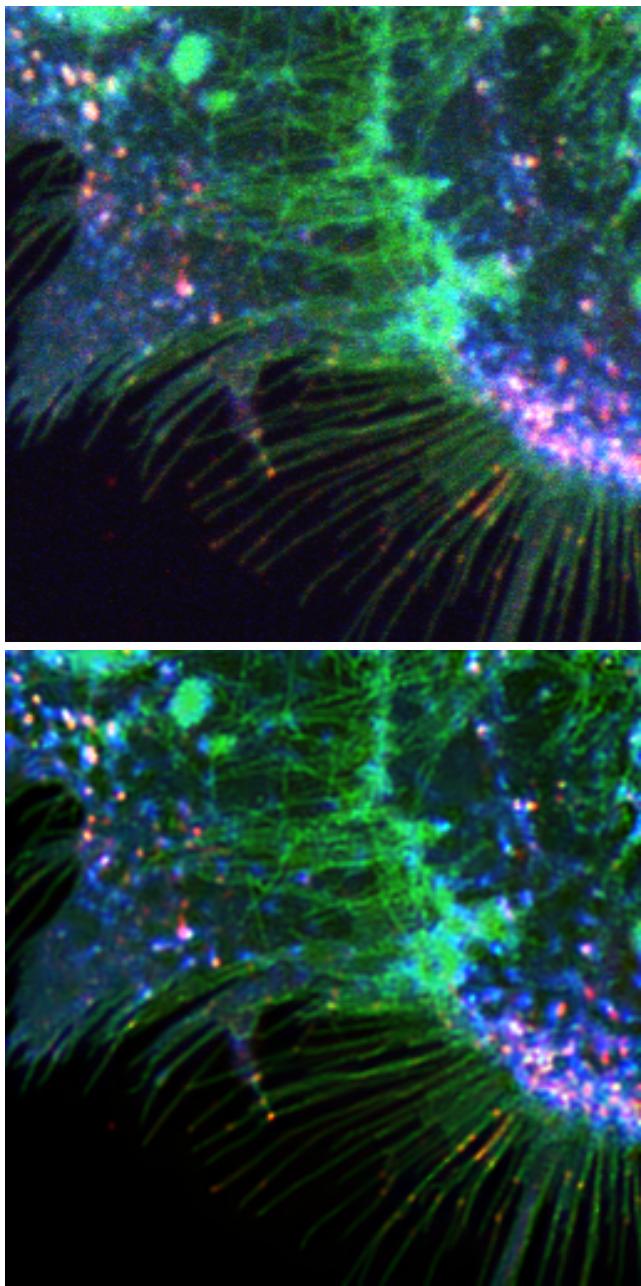


Figure 6.24: Restoration of image of epidermal cancer cell (1).

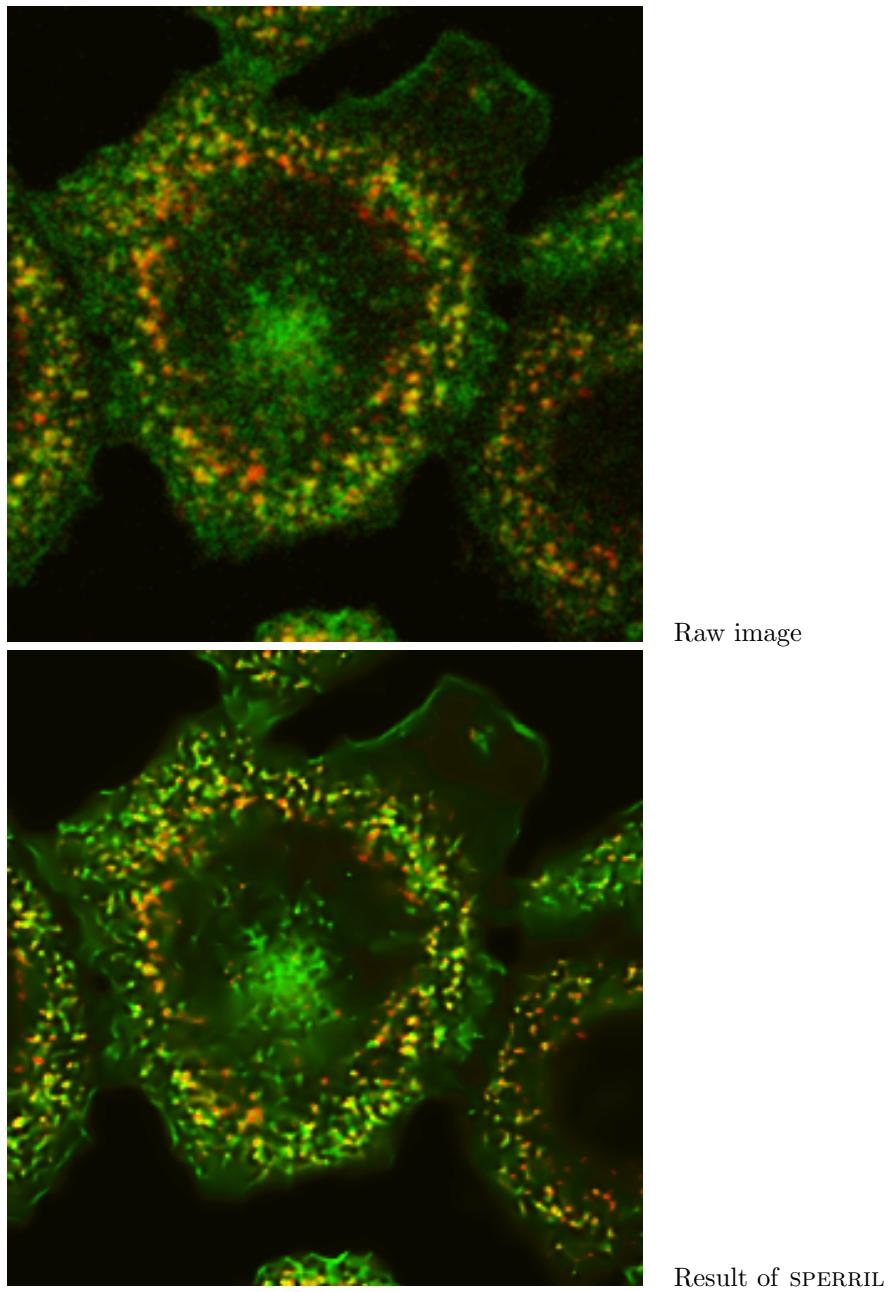


Figure 6.25: Restoration of image of epidermal cancer cell (2).

Chapter 7

Multiresolution Markov Random Field image restoration

The difference between ordinary and extraordinary is that little extra.

Markov Random Fields (MRF) have proven to be a very powerful tool in image processing. Numerous applications of single scale MRF's in image analysis and image enhancement are described in literature. Later, the combination of MRF's and multiresolution principles provided an even more powerful framework for image processing: applications can be found in image segmentation [Bouman and Shapiro, 1994] and denoising [Malfait and Roose, 1997, Jansen and Bultheel, 2001, Pižurica et al., 2002, Pižurica, 2002].

In this Chapter, we give some background on Markov Random Fields, and describe an original method for image restoration, based on a theoretical idea of dr. A. Pižurica in [Pižurica, 2002].

7.1 Markov Random Fields

MRF's provide a means to impose prior knowledge about image properties like smoothness (i.e., existence of relatively large areas with similar pixel intensities) and edge continuity. Figure 7.1 illustrates imposing a priori knowledge on edge continuity by giving a prior preference to spatially connected clusters of edge pixels. The origin of MRF's goes back to statistical mechanics describing ferromagnetism; in analogy, neighboring pixels can interact like groups of magnetic dipoles.

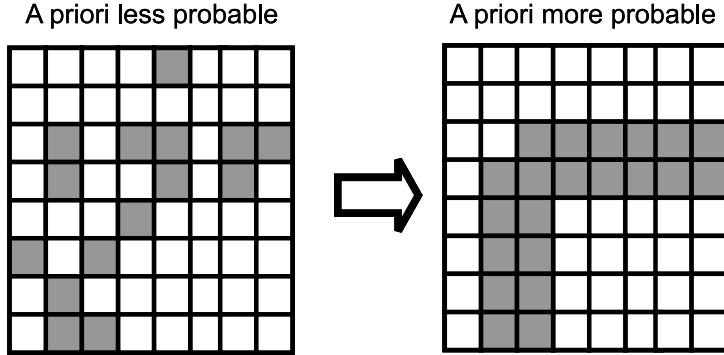


Figure 7.1: MRF prior and edge continuity.

More formally, statistical mechanics is described using Gibbs random fields (GRF) [Li, 1995, Pižurica, 2002] as probability models for the fluctuations of large physical systems around their equilibrium. These GRF's are used for global modeling of the system variables, and express the joint probability of a MRF. The Gibbs probability function is given by:

$$P(\mathbf{x}) = \frac{e^{-H(\mathbf{x})/T}}{Z}$$

where $H(\mathbf{x})$ is the *energy function*; Z is the so-called partition function, and is a normalization constant; T is the *temperature*, which controls the “peakedness” of the distribution. When T is smaller, the distribution is more peaked (i.e., the difference between two different image configurations is larger) so the minimum is easier to find.

The Hammersley-Clifford theorem links the local MRF with the global GRF [Li, 1995]. This link enables modeling the global system through local spatial interactions. This link between Gibbs and Markov random fields are the so-called *cliques*. A clique is a set of elements in the lattice, of which all members are *neighbors*. We give an illustration of first-order and second-order neighborhoods in Figure 7.2, and of their corresponding cliques in Figure 7.3. A Markov Random Field is a special type of Gibbs Random Field, the energy of which can be expressed as a sum of *clique potentials* $V_C(\mathbf{x})$:

$$H(\mathbf{x}) = \sum_{C \in \mathcal{C}} V_C(\mathbf{x}),$$

with \mathcal{C} the collection of all possible cliques in the image. The clique potential for a certain clique in the lattice depends on the local configuration on the

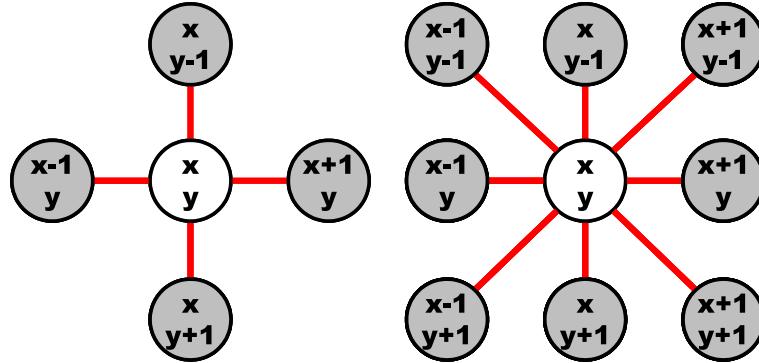


Figure 7.2: The concept of neighborhoods. Left, a first order neighborhood is shown; right, a second order neighborhood is shown. A first order neighborhood contains an element and all its neighbors with an Euclidian distance equal to one, so its horizontal and vertical neighbors. A second order neighborhood contains an element and all its neighbors with a distance smaller than or equal to $\sqrt{2}$, so its horizontal, vertical as well as its diagonal neighbors.

clique, and is only function of the members of the clique. Our prior knowledge about the local image properties is encoded by defining the appropriate clique potential functions.

For the sake of compactness, each pixel is denoted with a single index (like in a raster scanning) and a MRF prior with pair-wise cliques only is assumed. Let $\mathcal{L} = \{1, \dots, n\}$ denote the index set (a set of pixels on a regular rectangular lattice). Further on, let $\langle k, l \rangle$ denote a pair of neighbors (i.e., a clique composed of pixels k and l) and let $V(x_k, x_l)$ denote the potential of the clique $\langle k, l \rangle$.

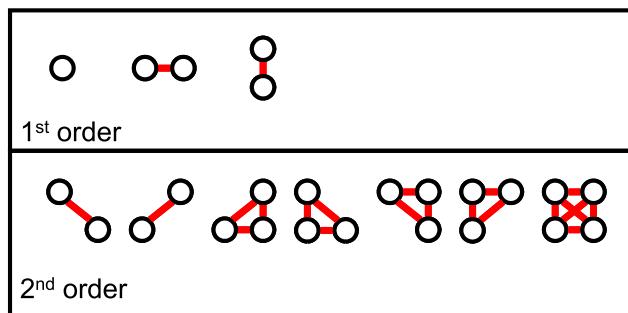


Figure 7.3: Cliques in first order and second order neighborhoods. The second order neighborhood also includes all the clique configurations from the first neighborhood.

7.1.1 An example: the Ising model

To illustrate this with a practical example, we will now discuss the simplest MRF model, which is the so-called Ising model (introduced in [Ising, 1925], described practically in [Pižurica, 2002]). It assumes a regular lattice x_k with spins, where every grid point can have two values, i.e., $S(x, y) = -1$ or 1 , and the first-order neighborhood from Figure 7.2. The cliques consist of up to two pixels. For the Ising model, the potential of the single-site clique is:

$$V_1(x_k) = \alpha x_k$$

and the potential of a pair-site clique is:

$$V_2(x_k, x_l) = \beta x_k x_l$$

The resulting energy $H(\mathbf{x})$ associated with a certain lattice configuration in the Ising model is given by:

$$H(\mathbf{x}) = \sum_k V_1(x_k) + \sum_k \sum_{(x_l) \in \partial(x_k)} V_2(x_k, x_l),$$

where $\partial(x_k)$ denotes the first-order neighborhood of x_k . The first term expresses the cost of an occurrence of a value in the lattice, no matter the values of its neighbors. If no value has *a priori* preference, $\alpha = 0$. The second term expresses the cost of the occurrence of couples of values for two neighboring sites in the lattice. It describes the interactions of pairs of neighbors. If $\beta < 0$, equal neighbors are more likely (Figure 7.4); if $\beta > 0$, different neighbors are more likely. The larger β is, the larger the areas become where the pixel values alternate constantly (see Figure 7.5).

7.1.2 Energy minimization

After having defined a proper energy function, we need to choose an efficient minimization algorithm. Here we describe briefly two examples of such algorithms:

1. the Metropolis algorithm, which belongs to the class of the random search algorithms;
2. the Graduated Non Convexity algorithm (GNC), which belongs to the class of the deterministic algorithms.

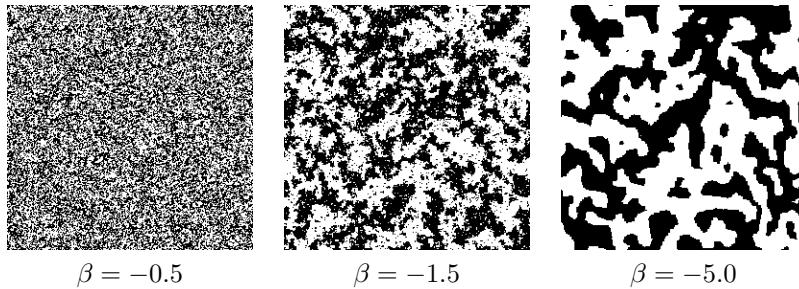


Figure 7.4: Influence of the β parameter in the Ising model. Each sample of the random field was obtained for $\alpha = 0$ and $T = 2$ (20 iterations).

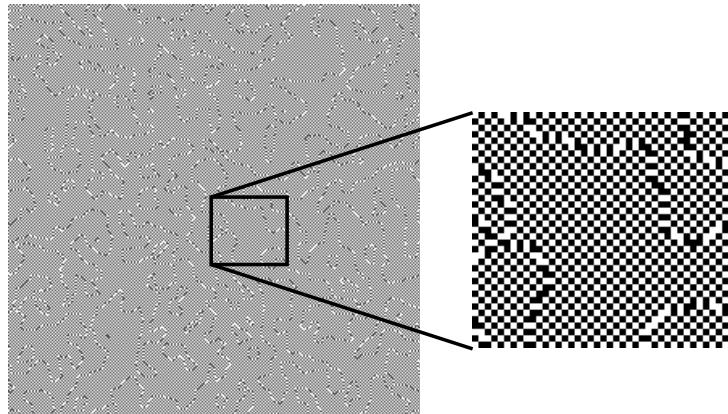


Figure 7.5: Influence of a positive β parameter in the Ising model: the larger β becomes, the larger the areas of chess board patterns.

7.1.2.1 The Metropolis algorithm

The Metropolis algorithm [Li, 1995] is a random search algorithm: some elements of the lattice are disturbed randomly, after which the posterior energy of this disturbed lattice is calculated. This new configuration is then accepted with a probability which depends both on the change in the posterior energy, and on the temperature of the system. The higher the temperature of the system, the more disturbances are accepted that are energetically “disadvantageous.”

We can use the Metropolis in simulated annealing by lowering the temperature during the iterations. In Figure 7.6, we show how a fixed α and β produce a random scattering of the labels when the temperature is high. When the temperature is lowered, labels start to cluster in smooth areas.

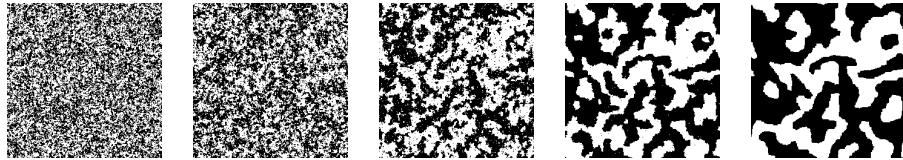


Figure 7.6: Ising model combined with the Metropolis algorithm for random sampling in a simulated annealing scheme ($\alpha = 0.0$; $\beta = -1$; $T_0 = 4$; $\Delta T = 0.04$, snapshot every 10 iterations)



Figure 7.7: Principle of Graduated Non Convexity: at first, the energy functional is “convexified”; when the minimum for this energy functional is reached, its convexity is decreased, and the minimum of this functional is computed. This procedure is repeated until a reliability criterion is satisfied.

The advantage of random search algorithms is that they are less likely to get trapped in local minima compared to other search procedures. The disadvantage is that random searches are computationally very intensive, and converge slowly.

7.1.2.2 Graduated Non-Convexity Minimization

The Graduated Non-Convexity (GNC) minimization algorithm, as described in [Blake and Zisserman, 1987, Li, 1995], is attractive due to the low computation complexity. It derives its name from the fact that the clique potentials are chosen in such a way that the resulting energy functional is convex in the beginning (but not very realistic, see Figure 7.7). When the minimum for this energy functional is found, the potentials are modified so that the energy functional becomes less convex, but already more realistic. Now, several local minima might appear, but we minimize in the neighborhood of the previously found minimum. Thus the procedure is continued: the energy functional becomes less and less convex, but we are likely to find a solution close to the global minimum because we stay close to the minimum of the more “convexified” energy

functional (i.e., made more convex).

In this implementation, the pairwise potential $V_2(x_k, x_l)$ is represented as $V(x_k, x_l) = g_\gamma(x_k - x_l)$ where the potential function $g_\gamma(\eta)$ is designed via the so called *adaptive interaction function* (AIF) $h_\gamma(\eta)$ as follows: $g'_\gamma(\eta) = 2\eta h_\gamma(\eta)$. In our implementation, we have chosen to use ([Li, 1995], p.129):

$$h_\gamma(\eta) = \frac{1}{1 + \frac{|\eta|}{\gamma}}$$

The larger the parameter γ , the more convex the resulting energy functional is. The minimization is started with a choice of γ that is large enough. The GNC algorithm applied to our problem seeks the solution by the following iterations:

$$x_k^{(t+1)} \leftarrow x_k^{(t)} - f(\alpha_k) \left((h * x^{(t)})_k - y_k \right) + \sum_{\langle k, l \rangle \in C_l} f(\alpha_l) \left(x_k^{(t)} - x_l^{(t)} \right) h_\gamma \left(x_k^{(t)} - x_l^{(t)} \right) \quad (7.1)$$

which is gradient descent minimization method. While this gradient descent can still decrease the energy, γ is kept constant. When a minimum of the energy for a certain value of γ is found, γ is lowered, and the iterations are restarted. This is repeated until the minimum is found.

7.2 Multiresolution MRF image restoration

7.2.1 Existing methods

Image restoration based on single scale MRF's are studied extensively in literature, using different approaches for minimizing the energy function: the Gibbs sampler [Geman and Geman, 1984, Geman and Reynolds, 1992], GNC [Nikolova et al., 1998, Nikolova, 1999].

The literature about image restoration based on Multiresolution MRF's (MMRF) however is very limited. The MMRF's have been used in segmentation [Bouman and Shapiro, 1994] and denoising [Malfait and Roose, 1997, Jansen and Bultheel, 2001, Pižurica et al., 2002, Pižurica, 2002] but much less for the restoration of noisy and *blurred* images. Two of the few examples are [Belge et al., 2000, Robini and Magnin, 2003]. In [Belge et al., 2000], a multi-scale approach was proposed in combination with a half-quadratic regularization scheme [Geman and Yang, 1995]. In [Robini and Magnin, 2003], a simple wavelet-based penalty term was proposed that favors the formation of horizontal and vertical boundaries, while preserving diagonal boundaries. The minimization of the corresponding energy functional was computed using the Metropolis sampler combined with simulated annealing.

7.2.2 Background on MMRF image restoration

In [Luettgen et al., 1993], it is suggested that a multiscale extension of the MRF modeling framework provides a superior basis for image analysis and modeling. This kind of representation could be used to take also into account inter- and intraband dependencies of the wavelet coefficients.

A direct application of the classical MRF restoration approach [Geman and Geman, 1984] in the image domain requires the minimization of the following posterior energy:

$$E(\mathbf{f}|\mathbf{g}) = \sum_{k \in \mathcal{L}} \left((h * f)_k - g_k \right)^2 + \lambda \sum_{k \in \mathcal{L}} \sum_{\langle k, l \rangle \in \mathcal{C}} V(f_k, f_l) \quad (7.2)$$

where $\mathbf{f} = \{f_1, \dots, f_n\}$ is the unknown degradation-free image, $\mathbf{g} = \{g_1, \dots, g_n\}$ is the observed image, h is the point spread function and \mathcal{C} denotes the set of all possible (pair-wise) cliques. The first term in the posterior energy defined above represents “closeness to the data” constraint and the second term represents the “smoothness” constraint (i.e., the prior knowledge about the image structure). The parameter λ expresses the relative importance given to each of these two terms, respectively. The main two problems with this approach are:

1. difficulties in specifying the parameters (i.e., in balancing between the data fit and the smoothness)
2. the construction of *discontinuity adaptive* potentials $V(x_k, x_l)$ which would “turn-off” smoothing at edges.

An elegant solution to both of these problems is offered by implementing the energy minimization in a sparse multiresolution representation and using a Graduated Non-Convexity method (GNC) to solve it, as was proposed in [Pižurica, 2002]. Let us denote the wavelet transform of the non-degraded data by $\mathbf{x} = \mathcal{W}(\mathbf{f})$ and the wavelet transform of the observed data by $\mathbf{y} = \mathcal{W}(\mathbf{g})$ ¹. Treating only the presence of noise, the idea introduced in [Pižurica, 2002] was to express the posterior energy in each subband as:

$$E(\mathbf{x}|\mathbf{y}) = \sum_{k \in \mathcal{L}} f(\alpha_k) \left(x_k - y_k \right)^2 + \sum_{k \in \mathcal{L}} \sum_{\langle k, l \rangle \in \mathcal{C}} f(\alpha_l) V(x_k, x_l), \quad (7.3)$$

where α_k is a wavelet domain “edge indicator” at position k and $f(\alpha_k)$ is a monotonic increasing function of α . The second sum is over all neighbors x_l of a certain coefficient x_k , and is weighted by its importance factor $f(\alpha_l)$.

¹In contrast to previous notations, \mathbf{g} means data *only* contaminated by noise.

An edge indicator is in general a local measurement computed from x_k and from its neighboring coefficients within the same and/or adjacent scales. The above defined formulation takes care of both balancing the data fit and the smoothness on a local basis (from point to point) and achieving a discontinuity adaptation using even a simple isotropic potential function. In particular, $f(\alpha_k)$ that appears in the first term will give more weight to the measured data at position k if it is likely that y_k is an “edge coefficient”. Also, through weights $f(\alpha_l)$ in the second term, a discontinuity-adaptive smoothness constraint is imposed (where those neighbors y_l that are more likely to represent edges will have a stronger influence on the estimate of the non-degraded neighboring value x_k).

7.3 Proposed method

We propose a new MMRF based restoration method, which was developed in close collaboration with dr. Aleksandra Pižurica. This new method is based on the formal description by equation (7.3) of [Pižurica, 2002], which extends the classical image restoration approach to a locally adaptive multiresolution method. While in [Pižurica, 2002] this idea was discussed only theoretically and in the context of denoising, we propose a practical realization for degradation due to noise *and blur* in this work, and apply it to real image restoration problems. We extend equation (7.3) to take into account blur as well. In that case, x_k in first term² of equation (7.3) has to be replaced by $(h * x)_k$, and the expression becomes:

$$E(\mathbf{x}|\mathbf{y}) = \sum_{k \in \mathcal{L}} f(\alpha_k) \left((h * x)_k - y_k \right)^2 + \sum_{k \in \mathcal{L}} \sum_{\langle k, l \rangle \in \mathcal{C}} f(\alpha_l) V(x_k, x_l), \quad (7.4)$$

We have implemented this method with the steerable pyramid representation where $f(\alpha_k)$ is representing the probability that a certain observed wavelet coefficient is due to signal.

7.3.1 Possible edge indicators

We have experimented with several possible edge indicators, such as the magnitude of the wavelet coefficients in a certain neighborhood, as well as the local variance of the wavelet coefficients in a certain neighborhood. However, these edge indicators haven't shown to be useful in our image restoration attempts.

²i.e., the data fidelity term, which imposes that solution should fit with the measured data. The second term or regularization term imposes our prior knowledge of the *solution*, and doesn't need to be adapted.

Let us resume the bivariate wavelet shrinkage method [Sendur and Selesnick, 2002]:

$$\hat{w}_1(k) = \frac{\left(\sqrt{y_1^2(k) + y_2^2(k)} - \frac{\sqrt{3}\sigma_n^2}{\sigma(k)} \right)_+}{\sqrt{y_1^2(k) + y_2^2(k)}} y_1(k) \quad (7.5)$$

$$= \mathcal{S}(k) y_1(k) \quad (7.6)$$

with $(x)_+ = \max(x, 0)$, $\hat{w}_1(k)$ the denoised coefficient corresponding with the noisy coefficient $y_1(k)$ and its parent $y_2(k)$, σ_n^2 the noise variance and $\sigma(k)$ the standard deviation of the wavelet around the coefficient at position k). With $\mathcal{S}(k)$, we denote the shrinkage factor for $y_1(k)$.

The shrinkage factor $\mathcal{S}(k)$ can be seen as an estimation how clean or how noisy coefficient $y_1(k)$ is: when $\mathcal{S}(k)$ is closer to 1, the coefficient is more likely to contain more signal than noise; when $\mathcal{S}(k)$ is closer to 0 however, the coefficient is more likely to contain more noise than signal, and is suppressed.

In our implementation, we used these fields $\mathcal{S}(k)$ as functions of the “edge indicator” at position k , and take over the role of $f(\alpha_k)$ in equation 7.3.

7.3.2 Results

In Figure 7.9, the image top left is blurred, and white Gaussian noise was added. The image top right was obtained with a GNC algorithm implemented in the spatial domain, with λ and μ fixed. The bottom image is obtained with our multiscale GNC method. We can conclude that this result is rather promising in terms of sharpening and noise reduction, at the cost of a significant amount of ringing. Also, the algorithm in its current form isn’t stable for high noise levels (like Gaussian noise with $\sigma_{noise} \gg 5$ for 8 bits images).

In Figure 7.10, we show another result of this preliminary algorithm. Again, we can draw the same conclusions: significant sharpening, no significant noise amplification in the smooth areas, but again a significant amount of ringing artifacts around sharpened edges.

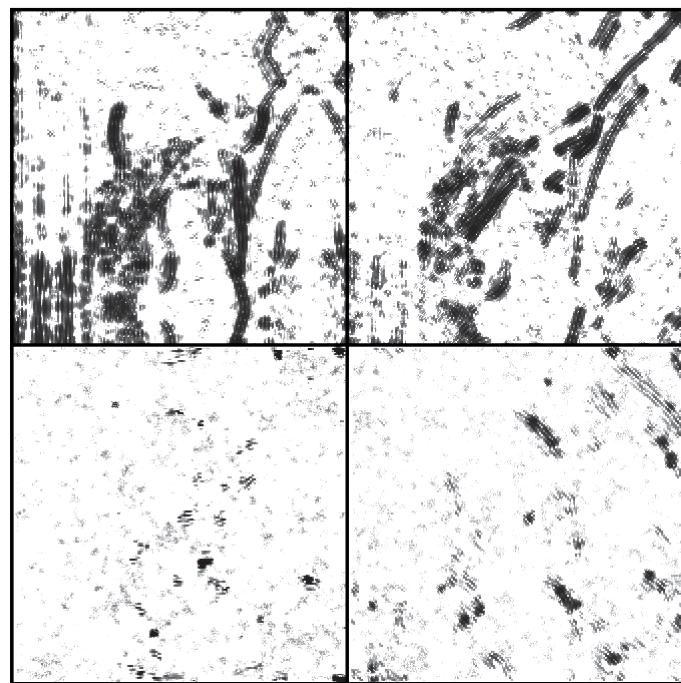


Figure 7.8: The $f(\alpha_k)$ as used in our GNC algorithm as regularization weight map. The four squares show the edge indicators for our algorithm applied to the restoration of the *Lena* image for the four oriented subbands.



Figure 7.9: A preliminary result of our GNC-based method. Top: the degraded image. Bottom left: result of the spatial version of the GNC-based method (with constant λ and μ). Bottom right: result of our multiscale GNC-based method.

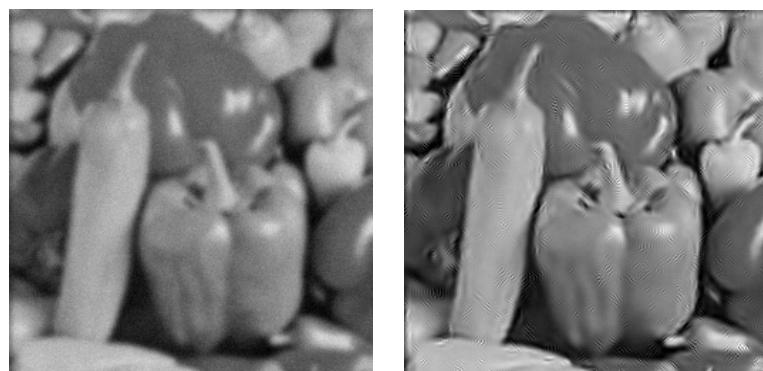


Figure 7.10: Another preliminary result of our GNC-based method. Left: the degraded image. Right: result of our multiscale GNC-based method.

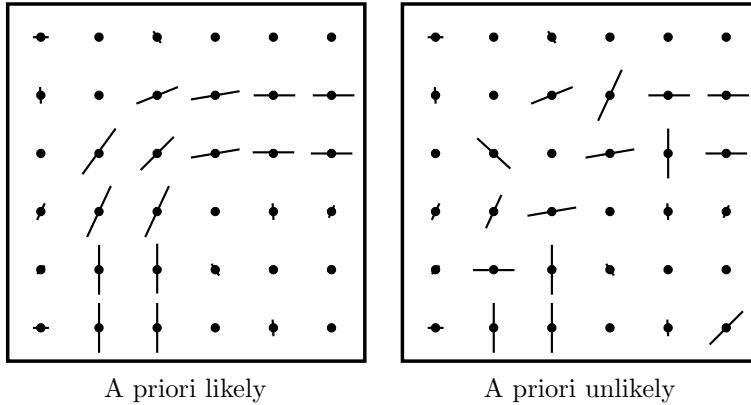
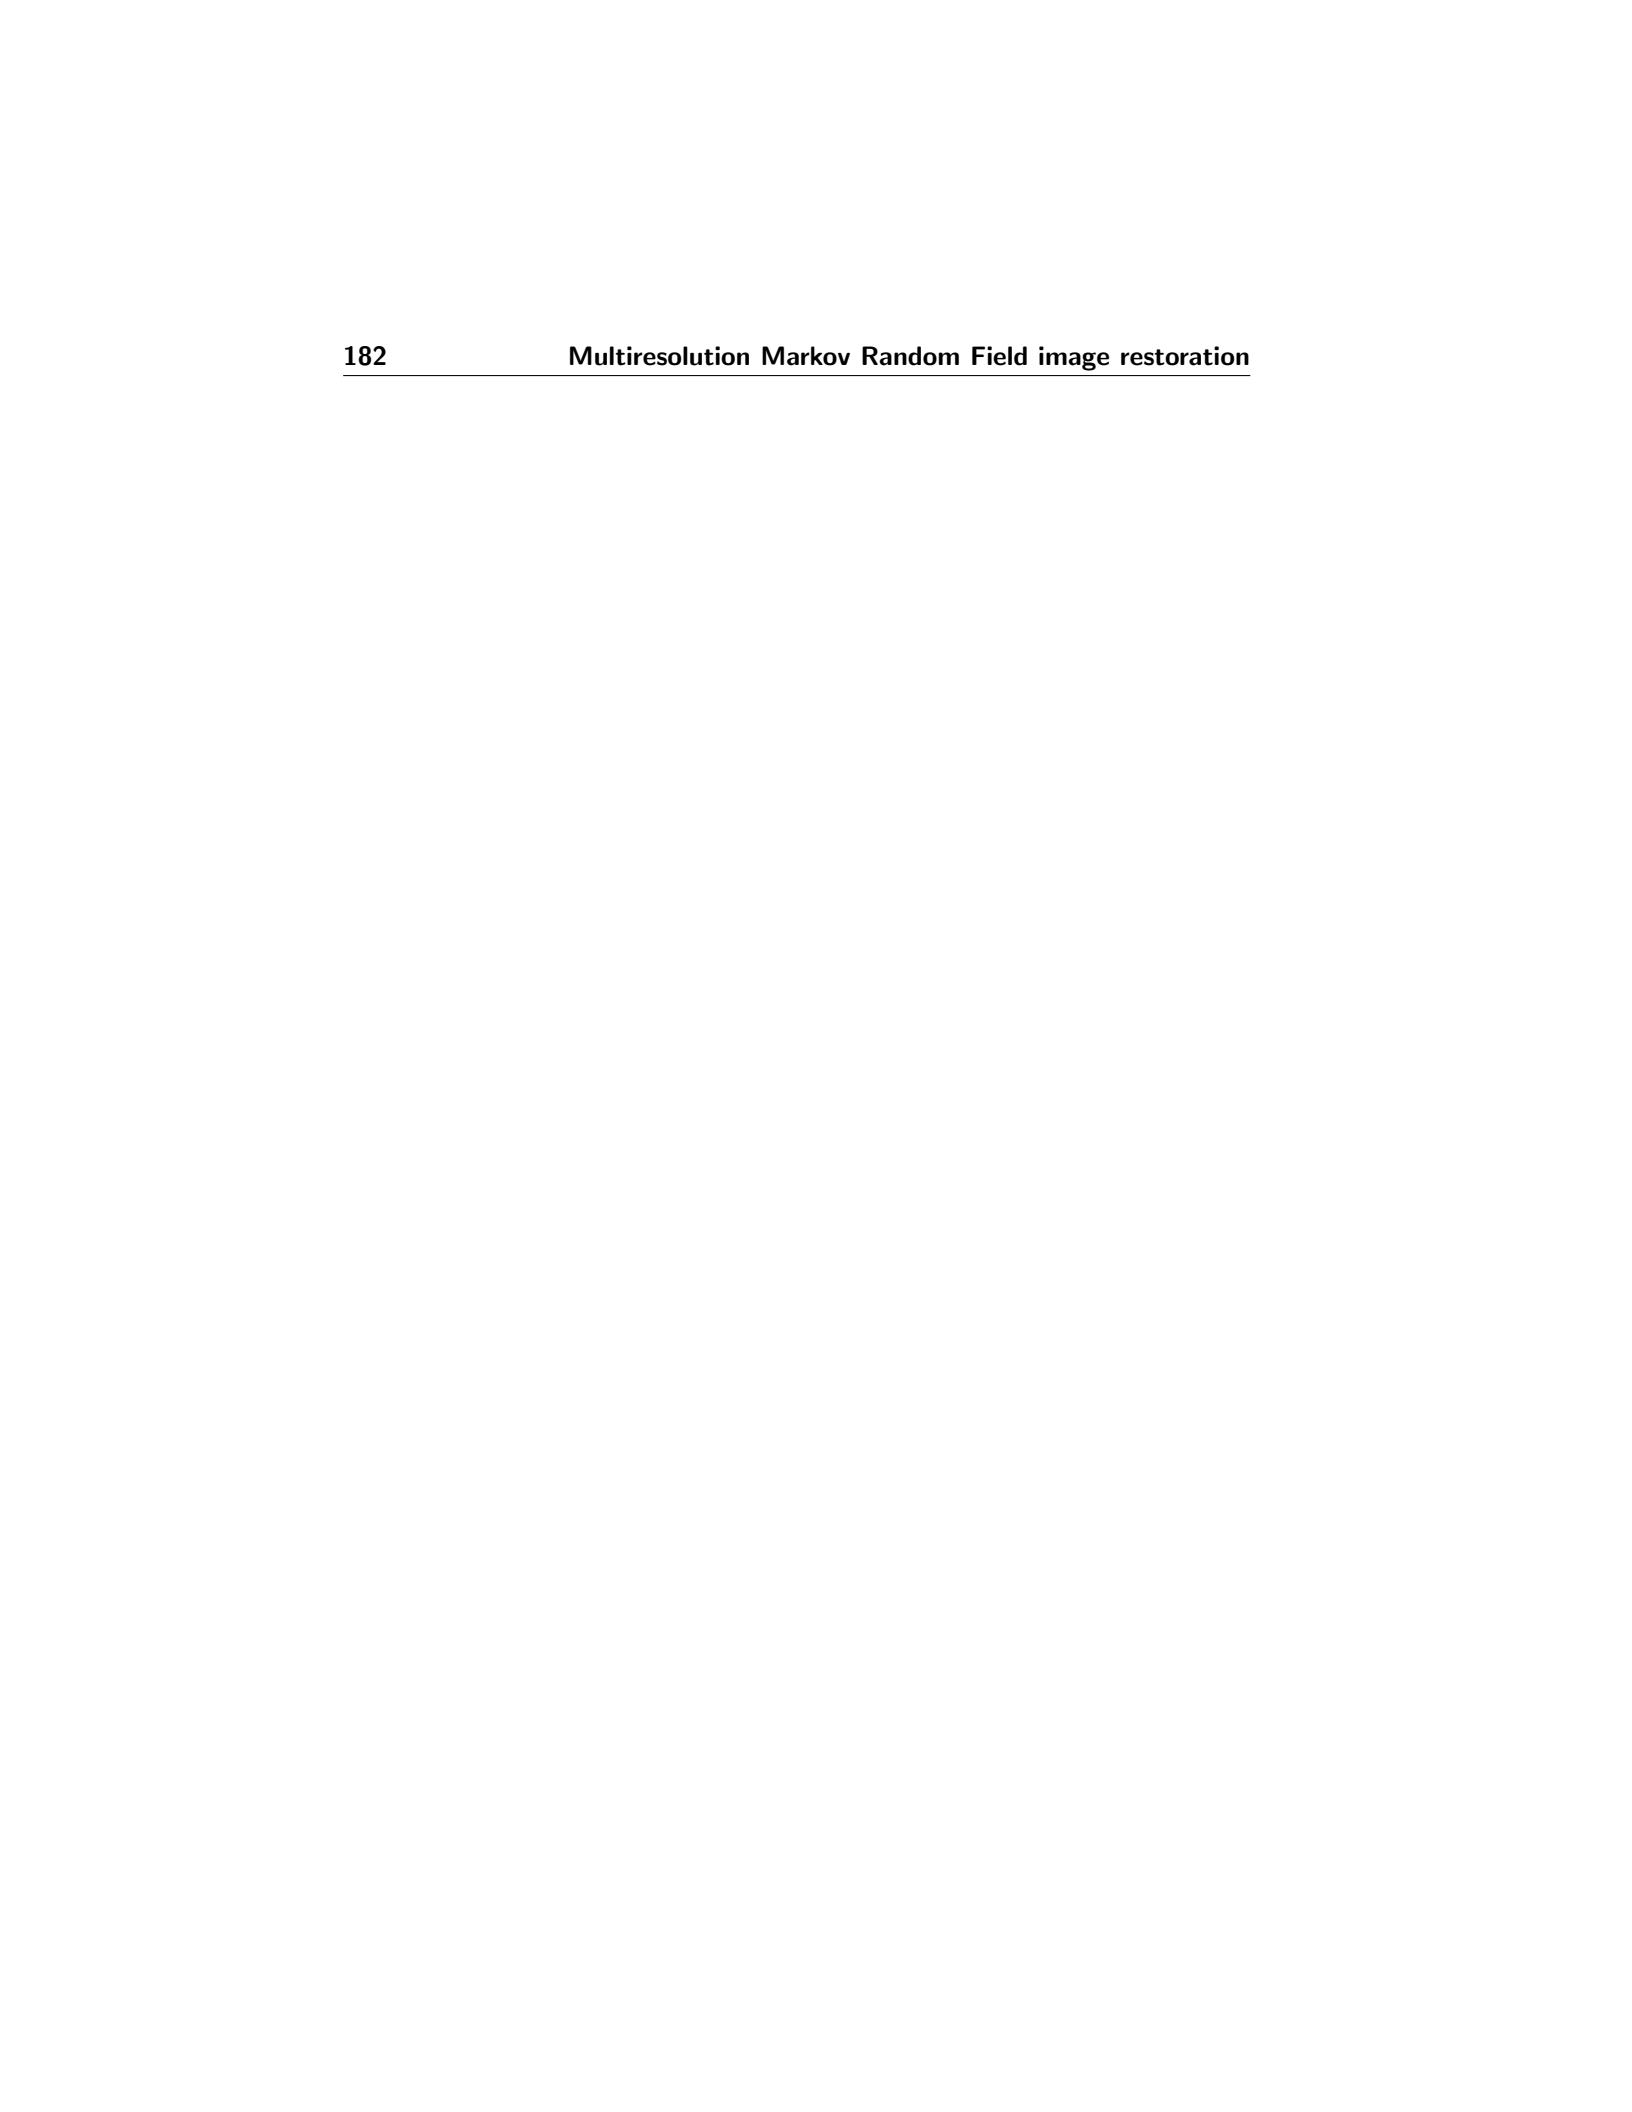


Figure 7.11: Local edge strength and orientation of an image. We would like to investigate how these local measures can be exploited in MMRF image restoration. Left: a configuration that is a priori likely. Right: a configuration that is a priori unlikely.

7.4 Conclusions and future work

In this Chapter, we described some preliminary results of a new method, providing a proof of concept of the idea described in [Pižurica, 2002]. Clearly, it still needs a lot of work to reach a more mature state, but the results are already quite promising.

An idea for future research would be to base the α_l in equation (7.3) on the local energy and the local orientation (see Subsubsection 4.3.1.2), as obtained from the steerable pyramid decomposition. Elements with a large local energy and oriented in the direction indicated by the calculated local orientation at a certain position in the image, would impose a stronger influence on their neighbors. That is illustrated in Figure 7.11, where a lattice is shown with the local orientation and strength in each lattice point. It extends the concept of Figure 7.1: instead of stimulating connected edges based on the criterium of large magnitude coefficients alone, we propose to incorporate also knowledge about the edge strength and orientation.



Chapter 8

Conclusions

Of course there's no reason for it; it's just our policy.

8.1 Summary of the conclusions

In this thesis, we explore the use of multiresolution concepts in image restoration.

In Chapter 2, we described the process of image formation in optical systems and the problem of image degradation in such systems. As a case study, we focused on the image formation in a confocal microscope. In particular, we described the problem of image degradation in such systems.

In Chapter 3, we described some commonly used techniques for image restoration and their limitations.

Chapter 4 is the last tutorial chapter aimed to be a general introduction to the concepts of multiresolution, and examples how these concepts are used in image processing.

In Chapter 5, we described two original multiresolution-based methods (*Wave-Track* and *Sparse-Max*) for parametric blur estimation. We tested these methods on several natural images that were synthetically degraded. These methods have proven to be able to estimate parametric image blur in a wide range of blur and noise levels.

In Chapter 6, we described how *Wave-Track* (the oldest of our two blur estimation methods) was incorporated in an image restoration method, called SPERRIL. We tested SPERRIL both on synthetical and biological images, and concluded that it provided state-of-the-art restoration results with reduced artifacts. The biological validation was performed by colocalization analysis, where the results showed that SPERRIL was best in separating images of cells with high colocalization from images of cells with low colocalization.

Another original restoration algorithm was proposed in Chapter 7. It is a practical algorithm based on an idea from dr. Pižurica, which has proven to work. Results of this algorithm however are still preliminary.

8.2 Future research

The blur estimation methods all assume a Gaussian PSF, and estimate the standard deviation of the Gaussian function. We already proposed extending these methods to out-of-focus blur. Future research could also include estimating the blur parameters of PSF's other than the Gaussian function or the uniform disc.

All methods we discussed are tested on 2D images. However, confocal images are 3D in nature. So, future research has to extend these methods to practical algorithms in 3D. We already explained in Chapter 6 some key ideas to extend SPERRIL to 3D.

Finally, the restoration method proposed in Chapter 7 is based on Šendur's method for denoising. We used his method to determine the spatially-variant regularization parameter. However, we are convinced that there are many ways that can be used to determine these spatially-variant regularization parameter. In Chapter 7, we also propose a new idea to determine these parameters using edge orientation and edge strength as estimated with steerable pyramids (estimation of edge orientation and edge strength is explained in Chapter 4).

Bibliography

- [Abu-Tarif, WWW] Abu-Tarif, A. (WWW). Note 4: Description of spherical aberrations (from autoquant imaging). <http://www.aqi.com/pdf/technote4.pdf>.
- [Alonso and Finn, 1989] Alonso, M. and Finn, E. J. (1989). *Fundamentele Natuurkunde 3: Golven*. Delta Press.
- [Andersson, 1992] Andersson, M. (1992). *Controllable Multidimensional Filters in Low Level Computer Vision*. PhD thesis, Linköping University, Sweden, SE-581 83 Linköping, Sweden. Dissertation No 282, ISBN 91-7870-981-4.
- [Ayers and Dainty, 1988] Ayers, G. and Dainty, J. (1988). Iterative blind deconvolution method and its applications. *Optics Letters*, 13(7):547–549.
- [Banham et al., 1994] Banham, M., Galatsanos, N., Gonzalez, H., and Katsaggelos, A. (1994). Multichannel restoration of single channel images using a wavelet-based subband decomposition. *IEEE Trans. Image processing*, 3(6):821–833.
- [Banham and Katsaggelos, 1996] Banham, M. and Katsaggelos, A. (1996). Spatially adaptive wavelet-based multiscale image restoration. *IEEE Trans. Image processing*, 5(4):619–634.
- [Barat et al., 2003] Barat, C., Ducottet, C., Bilgot, A., and Desbat, L. (2003). Segmentation of blurred objects using wavelet transform: application to x-ray images. In *Proc. SPIE Wavelet Applications in Industrial Processing*, pages 191–202.
- [Bascle et al., 1996] Bascle, B., Blake, A., and Zisserman, A. (1996). Motion deblurring and super-resolution from an image sequence. In *Proc. European Conference on Computer Vision, vol. 2*, pages 573–582.
- [Belge et al., 2000] Belge, M., Kilmer, M., and Miller, E. (2000). Wavelet domain image restoration with adaptive edge-preserving regularization. *IEEE Trans. Image processing*, 9(4):597–608.

- [Bell and Sejnowsky, 1997] Bell, A. and Sejnowsky, T. (1997). The independent components of natural scenes are edge filters. *Vision Research*, 37(23):3201–3439.
- [Bellis, WWW] Bellis, M. (WWW). History of the microscope. <http://inventors.about.com/library/inventors/blmicroscope.htm>.
- [Bertero and Boccacci, 1998] Bertero, M. and Boccacci, P. (1998). *Introduction to inverse problems in imaging*. Institute of Physics Publishing, Bristol and Philadelphia.
- [Beylkin, 1991] Beylkin, G. (1991). Wavelets, multiresolution analysis and fast numerical algorithms. *A draft of INRIA lectures*. available from <http://amath.colorado.edu/faculty/beylkin/>.
- [Beylkin et al., 1991] Beylkin, G., Coifman, R., and Rokhlin, V. (1991). Fast wavelet transforms and numerical algorithms I. *Comm. Pure and Appl. Math.*, 44:141–183. Yale University Technical Report YALEU/DCS/RR-696, August 1989.
- [Blake and Zisserman, 1987] Blake, A. and Zisserman, A. (1987). *Visual reconstruction*. MIT Press, Cambridge, Massachusetts.
- [Booth et al., 1998] Booth, M., Neil, M., and Wilson, T. (1998). Aberration correction for confocal imaging in refractive-index-mismatched media. *Journal of Microscopy*, 192(2):90–98.
- [Booth and Wilson, 2001] Booth, M. and Wilson, T. (2001). Refractive-index-mismatch induced aberrations in single-photon and two-photon microscopy and the use of aberration correction. *Journal of Biomedical Optics*, 6(3):266–272.
- [Bouman and Shapiro, 1994] Bouman, C. and Shapiro, M. (1994). A multiscale random field model for bayesian image segmentation. *IEEE Trans. Image Processing*, 3(2):162–177.
- [Boutet de Monvel et al., 2001] Boutet de Monvel, J., Le Calvez, S., and Ulfendahl, M. (2001). Image restoration for confocal microscopy : Improving the limits of deconvolution, with application to the visualization of the mammalian hearing organ. *Biophysical Journal*, 80(5):2455–2470.
- [Burt and Adelson, 1983] Burt, P. J. and Adelson, E. H. (1983). The laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31(4):532–540.
- [Cannon, 2000] Cannon, M. (2000). Blind deconvolution of spatially invariant image blurs with phase. *IEEE Trans. on Acoustics, Speech and Signal Processing*, ASSP-24(1):58–63.

- [Canny, 1986] Canny, J. (1986). A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8(6):679–698.
- [Cardoso and Laheld, 1996] Cardoso, J.-F. and Laheld, B. (1996). Equivariant adaptive source separation. *IEEE Trans. on Signal Processing*, 44(12):3017–3030.
- [Casavan and Gaidoukevitch, 2003] Casavan, W. and Gaidoukevitch, Y. (2003). Colocalization of fluorescent probes using image-pro plus v. 5.0. *Microscopy Today*, 11(6):48–50.
- [Castleman et al., 1998] Castleman, K. R., Schulze, M. A., and Wu, Q. (1998). Simplified design of steerable pyramid filters. In *Proc. IEEE ISCAS*.
- [Chan et al., 2001] Chan, T. F., Osher, S., and Shen, J. (2001). The digital tv filter and nonlinear denoising. *IEEE Trans. Image Processing*, 10(2):231–241.
- [Chang et al., 1991] Chang, M., Murat Tekalp, A., and Tanju Erdem, A. (1991). Blur identification using the bispectrum. *IEEE Trans. Signal Processing*, 39(10):2323–2325.
- [Chardon et al., 1999] Chardon, S., Vozel, B., and Chehdi, K. (1999). Parametric blur estimation using the generalized cross-validation criterion and a smoothness constraint on the image. *Multidimensional Systems and Signal Processing*, 10(4):395–414.
- [Clauws, 1998] Clauws, P. (1998). *Syllabus Optica en Kwantumelectronica*. Course at Ghent University.
- [Comon, 1994] Comon, P. (1994). Independent component analysis, a new concept? *Signal Processing*, 36(3):287–314.
- [Conchello and McNally, 1996] Conchello, J.-A. and McNally, J. G. (1996). Fast regularization technique for expectation maximization algorithm for optical sectioning microscopy. In *Proc. SPIE Vol. 2655, Three-Dimensional Microscopy: Image Acquisition and Processing III*, pages 199–208.
- [Şendur and Selesnick, 2002] Şendur, L. and Selesnick, I. W. (2002). Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency. *IEEE Trans. Signal Processing*, 50(11):2744–2756.
- [Csiszár, 1991] Csiszár, I. (1991). Why least squares and maximum entropy? an axiomatic approach to inference for linear inverse problems. *Annals of Statistics*, 19(4):2032–2066.
- [Daubechies, 1992] Daubechies, I. (1992). *Ten Lectures in Wavelets*. SIAM.
- [De Smet et al., 2004] De Smet, P., Rooms, F., Luong, H. Q., and Philips, W. (2004). Do not zero-pute: an efficient homespun mpeg audio layer ii decoding and optimization strategy. In *Proc. of ACM Multimedia 2004 conference, New York, USA*, pages 376–379.

- [Delle Luche et al., 2003] Delle Luche, C., Denis, F., and Baskurt, A. (2003). Three-dimensional steerable pyramid based on conic filters. In *Proc. SPIE Vol. 5266, Wavelet Applications in Industrial Processing*, pages 199–208.
- [Dempster et al., 1977] Dempster, A. O., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Soc. B*, 39(1):1–38.
- [Dey et al., 2004] Dey, N., Blanc-Féraud, L., Zimmer, C., Roux, P., Kam, Z., Olivo-Marin, J., and Zerubia, J. (2004). 3d microscopy deconvolution using richardson-lucy algorithm with total variation regularization. <http://www-sop.inria.fr/ariana/BIBLIO-ENG/Author/DEY-N.html>.
- [Donoho and Johnstone, 1994] Donoho, D. and Johnstone, I. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455.
- [Donoho and Johnstone, 1995] Donoho, D. and Johnstone, I. (1995). Denoising by soft-thresholding. *IEEE Trans. Information Theory*, 41(3):613–627.
- [Dragotti and Vetterli, 2003] Dragotti, P. and Vetterli, M. (2003). Wavelet footprints: theory, algorithms, and applications. *IEEE Trans. Signal Processing*, 51(5):1306–1323.
- [Elder and Zucker, 1998] Elder, J. and Zucker, S. (1998). Local scale control for edge detection and blur estimation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 20(7):699–716.
- [Field, 1987] Field, D. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4(12):2379–2394.
- [Figueiredo and Nowak, 2003] Figueiredo, M. and Nowak, R. (2003). An em algorithm for wavelet-based image restoration. *IEEE Trans. Image Processing*, 12(8):906–916.
- [Fish et al., 1995] Fish, D., Brinicombe, A., and Pike, E. (1995). Blind deconvolution by means of the richardson-lucy algorithm. *Journal of the Optical Society of America A*, 12(1):58–65.
- [Fitzpatrick, WWW] Fitzpatrick, R. (WWW). Huygens' principle. <http://farside.ph.utexas.edu/teaching/3021/lectures/node135.html>.
- [Florida State University College of Medicine, WWW] Florida State University College of Medicine (WWW). Tutorial: Special stains in histology. <http://www-medlib.med.utah.edu/WebPath/HISTHTML/STAINS/STAINS.html>.
- [Flusser and Suk, 1998] Flusser, J. and Suk, T. (1998). Degraded image analysis, an invariant approach. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 20(6):1–14.

- [Freeman and Adelson, 1991] Freeman, W. and Adelson, E. (1991). The design and use of steerable filters. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(9):891–906.
- [Gabor, 1946] Gabor, D. (1946). Theory of communication. *J. IEE (London)*, 93(26):429–457.
- [Galatsanos and Katsaggelos, 1992] Galatsanos, N. and Katsaggelos, A. (1992). Methods for choosing the regularization parameter and estimating the noise variance in image restoration and their relation. *IEEE Trans. Image Processing*, 1(3):322–336.
- [Geman and Reynolds, 1992] Geman, D. and Reynolds, G. (1992). Constrained restoration and the recovery of discontinuities. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(3):367–383.
- [Geman and Yang, 1995] Geman, D. and Yang, C. (1995). Nonlinear image recovery with half-quadratic regularization. *IEEE Trans. Image Processing*, 4(7):932–946.
- [Geman and Geman, 1984] Geman, S. and Geman, D. (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6(6):721–741.
- [Gennery, 1973] Gennery, D. (1973). Determination of optical transfer function by inspection of frequency domain plot. *Journal of the Optical Society of America*, 63(12):1571–1577.
- [GLMS, 2002] GLMS (2002). *GNU/Linux voor middelbare scholen*. <http://savannah.nongnu.org/projects/glms/>.
- [Gonzalez and Woods, 1993] Gonzalez, R. and Woods, R. (1993). *Digital Image Processing*. Addison-Wesley.
- [Granlund and Knutsson, 1995] Granlund, G. and Knutsson, H. (1995). *Signal Processing for Computer Vision*. Kluwer Academic Publishers.
- [Hanisch and White, 1993] Hanisch, R. and White, R. e. (1993). The restoration of hst images and spectra - ii (workshop proceedings). <http://www.stsci.edu/stsci/meetings/irw/tablecon.html>.
- [Hanisch, 1993] Hanisch, R. e. (1993). *Restoration: Newsletter of STScI's Image Restoration Project*. Space Telescope Science Institute.
- [Hell et al., 1993] Hell, S., Reiner, G., Cremer, C., and Stelzer, E. (1993). Aberrations in confocal fluorescence microscopy induced by mismatches in refractive index. *Journal of Microscopy*, 169(3):391–405.
- [Holmes, 1988] Holmes, T. (1988). Maximum-likelihood image restoration adapted for noncoherent optical imaging. *Journal of the Optical Society of America A*, 5(5):666–673.

- [Holmes, 1992] Holmes, T. (1992). Blind deconvolution of quantum limited incoherent imagery: maximum likelihood approach. *Journal of the Optical Society of America A*, 9(7):1052–1061.
- [Holmes et al., 1995] Holmes, T., Bhattacharyya, S., Cooper, J., Hanzel, D., and Krishnamurti, V. (1995). *Handbook of biological Confocal Microscopy (2nd ed.)*, chapter Light Microscopic Images reconstructed by Maximum Likelihood deconvolution, pages 389–402. Plenum Press, New York.
- [Homem et al., 2002] Homem, M., N.D.A., M., Costa, L., and Preza, C. (2002). Biological image restoration in optical-sectioning microscopy using prototype image constraints. *Real-Time Imaging*, 8(6).
- [Hossack, WWW] Hossack, W. (WWW). Digital image analysis and image processing 1. <http://www.ph.ed.ac.uk/~wjh/teaching/dia/index.html>.
- [Hsung et al., 1999] Hsung, T., Pak-Kong, D., and Siu, W. (1999). Denoising by singularity detection. *IEEE Trans. Signal Processing*, 47(11).
- [Hutton, 2002] Hutton, B. F. (2002). An introduction to iterative reconstruction. http://www2.alasbimnjournal.cl/alasbimn/CDA/sec_b/0,1206,SCID%253D455,00.html.
- [Hyvärinen, 1999] Hyvärinen, A. (1999). Sparse code shrinkage: Denoising of nongaussian data by maximum likelihood estimation. *Neural Comput.*, 11(7):1739–1768.
- [ImageJ, WWW] ImageJ (WWW). <http://rsb.info.nih.gov/ij/>.
- [Ising, 1925] Ising, E. (1925). Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik*, 31:253–258.
- [Jalobeanu et al., 2003] Jalobeanu, A., Blanc-Féraud, L., and Zerubia, J. (2003). Satellite image deblurring using complex wavelet packets. *International Journal of Computer Vision*, 51(3):205–217.
- [Jansen and Bultheel, 2001] Jansen, M. and Bultheel, A. (2001). Empirical bayes approach to improve wavelet thresholding for image noise reduction. *Journal of the American Statistical Association*, 96(454):629–639.
- [Kalifa et al., 2003] Kalifa, J., Mallat, S., and Rouge, B. (2003). Deconvolution by thresholding in mirror wavelet bases. *IEEE Trans. Image Processing*, 12(4):446–457.
- [Kingsbury, 1998a] Kingsbury, N. (1998a). The dual-tree complex wavelet transform: a new efficient tool for image restoration and enhancement. In *Proc. European Signal Processing Conference, EUSIPCO 98*, pages 319–322.
- [Kingsbury, 1998b] Kingsbury, N. (1998b). The dual-tree complex wavelet transform: a new technique for shift invariance and directional filters. In *Proc. IEEE Digital Signal Processing Workshop (Bryce Canyon), DSP 98*.

- [Kovesi, 2002] Kovesi, P. (2002). Edges are not just steps. In *Proc. ACCV 2002*, pages 822–827.
- [Kundur and Hatzinakos, 1996a] Kundur, D. and Hatzinakos, D. (1996a). Blind image deconvolution. *IEEE Signal Processing Magazine*, 13(3):43–64.
- [Kundur and Hatzinakos, 1996b] Kundur, D. and Hatzinakos, D. (1996b). Blind image deconvolution revisited. *IEEE Signal Processing Magazine*, 13(6):61–63.
- [Lagendijk and Biemond, 1991] Lagendijk, R. L. and Biemond, J. (1991). *Iterative Identification and Restoration of Images*. Kluwer Academic Publishers, Boston, Dordrecht, London.
- [Landmann, 2002] Landmann, L. (2002). Deconvolution improves colocalization analysis of multiple fluorochromes in 3d confocal data sets more than filtering techniques. *Journal of Microscopy*, 208(2):134–147.
- [Li, 1995] Li, S. (1995). *Markov Random Field Modeling in Computer Vision*. Springer-Verlag.
- [Lichtman, 1994] Lichtman, J. W. (1994). Confocal microscopy. *Scientific American*, 271(2):30–35.
- [Lidke et al., 2004] Lidke, D. S., Nagy, P., Heintzmann, R., Arndt-Jovin, D., Post, J., Grecco, H., Jares-Erijman, E., and Jovin, T. (2004). Quantum dot ligands provides new insights into erbB/her receptor-mediated signal transduction. *Nature Biotechnology*, 22(2):198–203.
- [Lindeberg, 1994] Lindeberg, T. (1994). Scale-space theory: A basic tool for analysing structures at different scales. *J. Of Applied Statistics*, 21(2):224–270.
- [Lucy, 1974] Lucy, L. B. (1974). An iterative technique for the rectification of observed distributions. *The Astronomical Journal*, 79(6):745–754.
- [Luettgen et al., 1993] Luettgen, M., Karl, W., Willskey, A., and Tenney, R. (1993). Multiscale representations of markov random fields. *IEEE Trans. on Signal Processing*, 41(12):3377–3395.
- [Malfait and Roose, 1997] Malfait, M. and Roose, D. (1997). Wavelet-based image denoising using a markov random field a priori model. *IEEE Trans. on Image Processing*, 6(4):549–565.
- [Mallat, 1989] Mallat, S. (1989). A theory for multiresolution signal decomposition. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(7):674–693.
- [Mallat, 1999] Mallat, S. (1999). *A wavelet tour of signal processing (2nd ed.)*. Academic Press, Oval Road, London.

- [Mallat and Hwang, 1992] Mallat, S. and Hwang, W. L. (1992). Singularity detection and processing with wavelets. *IEEE Trans. Information Theory*, 38(8):617–643.
- [Mallat and Zhong, 1992] Mallat, S. and Zhong, S. (1992). Characterization of signals from multiscale edges. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(7):710–732.
- [Manders et al., 1993] Manders, E. M., Verbeek, F. J., and Aten, J. A. (1993). Meaurement of co-localization of objects in dual colour confocal images. *Journal of Microscopy*, 169:375–382.
- [McNally et al., 1998] McNally, J., Markham, J., and Conchello, J. (1998). Comparing methods for 3d microscopy. In Cogswell, C., Conchello, J.-A., Lerner, J., Lu, T., and Wilson, T., editors, *Proc. SPIE vol 3261: Three-dimensional and multidimensional microscopy: Image Aquisition and Processing V*, pages 108–116.
- [Molecular Probes, WWW] Molecular Probes (WWW). Invitrogen detection technologies. <http://www.probes.com/>.
- [Molina et al., 2001] Molina, R., Núñez, J., Cortijo, F. J., and Mateos, J. (2001). Image restoration in astronomy, a bayesian perspective. *IEEE Signal Processing Magazine*, 18(2):11–29.
- [Moll, W.A.W., WWW] Moll, W.A.W. (WWW). Antoni van leeuwenhoek. <http://www.euronet.nl/users/warnar/leeuwenhoek.html>.
- [Morrone and Owens, 1987] Morrone, M. and Owens, R. (1987). Feature detection from local energy. *Pattern Recognition Letters*, 6(5):303–313.
- [Murtagh and Starck, 2003] Murtagh, F. and Starck, J. (2003). Bayes factors for edge detection from wavelet product spaces. *Optical Engineering*, 42(5):1375–1382.
- [MZA Associates Corporation, 2004] MZA Associates Corporation (2004). Modeling and simulation of beam control systems seventh annual directed energy symposium short course. <http://www.mza.com/doc/MZADEPSBCSMSC2004/>.
- [Narayan and Nityananda, 1986] Narayan, R. and Nityananda, R. (1986). Maximum entropy image restoration in astronomy. *Annual review of astronomy and astrophysics*, 24:127–170.
- [Neelamani et al., 2004] Neelamani, R., Choi, H., and Baraniuk, R. (2004). Forward: Fourier-wavelet regularized deconvolution for ill-conditioned systems. *IEEE Trans. Signal Processing*, 52(2):418–433.

- [Neelamani et al., 1999] Neelamani, R., Choi, H., and Baraniuk, R. G. (1999). Wavelet-based deconvolution for ill-conditioned systems. In *Proc. IEEE Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 6, pages 3241–3244.
- [Nik Multimedia, Inc., WWW] Nik Multimedia, Inc. (WWW). Nik sharpener pro. <http://www.nikmultimedia.com/sharpenerpro/usa/entry.php>.
- [Nikolova, 1999] Nikolova, M. (1999). Markovian reconstruction using a gnc approach. *IEEE Trans. Image Processing*, 8(9):1204–1220.
- [Nikolova et al., 1998] Nikolova, M., Idier, J., and Mohammad-Djafar, A. (1998). Inversion of large-support ill-posed linear operators using a piecewise gaussian mrf. *IEEE Trans. Image Processing*, 7(4):571–585.
- [Noponen, 1994] Noponen, E. (1994). *Electromagnetic Theory of Diffractive Optics*. PhD thesis, Helsinki University of Technology, <http://focus.hut.fi/staff/en.shtml>.
- [Nowak and Baraniuk, 1999] Nowak, R. D. and Baraniuk, R. G. (1999). Wavelet-domain filtering for photon imaging systems. *IEEE Trans. Image Processing*, 8(5):666–678.
- [Olivo-Marin, 2002] Olivo-Marin, J. (2002). Extraction of spots in biological images using multiscale products. *Pattern Recognition*, 35(9):1989–1996.
- [Olsen, 1993] Olsen, S. I. (1993). Estimation of noise in images: An evaluation. *CVGIP: Graphical Models and Image Processing*, 55(4):319–323.
- [Olshausen and Field, 1997] Olshausen, B. and Field, D. (1997). Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23):3311–3325.
- [Olympus Fluo View, WWW] Olympus Fluo View (WWW). Theory of confocal microscopy: Resolution and contrast in confocal microscopy. <http://www.olympusfluoview.com/theory/resolutionintro.html>.
- [Pavlović and Tekalp, 1992] Pavlović, G. and Tekalp, A. (1992). Maximum likelihood parametric blur identification based on a continuos spatial domain model. *IEEE Trans. Image Processing*, 1(4):496–504.
- [Pawley, 1995] Pawley, J. e. (1995). *Handbook of biological confocal microscopy (2nd ed.)*. Kluwer Academic/Plenum Press, New York.
- [Perona and Malik, 1991] Perona, P. and Malik, J. (1991). Detecting and localizing edges composed of steps, peaks and roofs. *MIT Technical Report LIDS-P-2075*.
- [Pižurica, 2002] Pižurica, A. (2002). *Image Denoising Using Wavelets and Spatial Context Modeling*. PhD thesis, Ghent University.

- [Pižurica et al., 2002] Pižurica, A., Philips, W., Lemahieu, I., and Acheroj, M. (2002). A joint inter- and intrascale statistical model for bayesian wavelet based image denoising. *IEEE Trans. on Image Processing*, 11(5):545–557.
- [Pižurica et al., 2003] Pižurica, A., Philips, W., Lemahieu, I., and Acheroj, M. (2003). A versatile wavelet domain noise filtration technique for medical imaging. *IEEE Trans. on Medical Imaging*, 22(3):323–331.
- [Plášek and Reischig, 1998] Plášek, J. and Reischig, J. (1998). Transmitted-light microscopy for biology: A physicist’s point of view part i. *Proceedings of the Royal Microscopical Society*, 33(2):121–127.
- [Polikar, WWW] Polikar, R. (WWW). The wavelet tutorial. <http://users.rowan.edu/~polikar/WAVELETS/WTTutorial.html>.
- [Portilla and Simoncelli, 2003] Portilla, J. and Simoncelli, E. (2003). Image restoration using gaussian scale mixtures in the wavelet domain. In *Proc. 9th IEEE Int’l Conf on Image Processing (ICIP)*, volume II, pages 965–968.
- [Portilla et al., 2003] Portilla, J., Strela, V., Wainwright, M., and Simoncelli, E. P. (2003). Image denoising using gaussian scale mixtures in the wavelet domain. *IEEE Trans. on Image Processing*, 12(11):1338–1351.
- [Press et al., 1988] Press, W., Teukolsky, S., Vetterling, T., and Flannery, B. (1988). Numerical recipes in c. <http://www.nr.com/>.
- [Pudney et al., 1994] Pudney, C., Kovesi, P., and Stretch, D. (1994). Surface detection using local energy for 3d confocal microscope images. In *Proc. Image Analysis 1994, Perth, Western Australia*, pages 68–73.
- [Reeves and Mersereau, 1992] Reeves, S. and Mersereau, R. (1992). Blur identification by the method of generalized cross-validation. *IEEE Trans. on Image Processing*, 1(3):301–311.
- [Richardson, 1972] Richardson, W. H. (1972). Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America*, 62(1):55–59.
- [RMA Belgium, WWW] RMA Belgium (WWW). Depth from defocus. <http://mecatron.rma.ac.be/dfd.html>.
- [Robini and Magnin, 2003] Robini, M. and Magnin, I. (2003). Stochastic non-linear image restoration using the wavelet transform. *IEEE Trans. Image Processing*, 12(8):141–151.
- [Rooms and Philips, 2000] Rooms, F. and Philips, W. (2000). Non-linear techniques in image restoration. In *Proceedings 1st Phd symposium of the Faculty of Applied Sciences at Ghent University*, page 26.
- [Rooms and Philips, 2001] Rooms, F. and Philips, W. (2001). Estimating the blur in confocal images. In *Abstract book Focus on Microscopy 2001 (Amsterdam, the Netherlands)*, page 81.

- [Rooms and Philips, 2002a] Rooms, F. and Philips, W. (2002a). Multiresolution techniques in microscopical image processing. *G.I.T. Imaging & Microscopy*, pages 60–61.
- [Rooms and Philips, 2002b] Rooms, F. and Philips, W. (2002b). Regularization of deconvolution using steerable pyramids. In *14th International Conference on Digital Signal Processing Proceedings (IEEE/DSP 2002), Santorini, Greece*, pages 943–946.
- [Rooms and Philips, 2004] Rooms, F. and Philips, W. (2004). Estimation and restoration of confocal images. In *Proc. of 5th Phd symposium of the Faculty of Applied Sciences at Ghent University*, page 122.
- [Rooms et al., 2005] Rooms, F., Philips, W., and Lidke, D. (2005). Simultaneous estimation and restoration of confocal images and performance evaluation by colocalization analysis. *Journal of Microscopy*, 218(1):22–36.
- [Rooms et al., 2000a] Rooms, F., Philips, W., Meesen, G., Poffijn, A., and Van Oostveldt, P. (2000a). 3d analysis of nuclear tracks with confocal microscopy. In *Newsletter for the Belgian Association for Microscopy nr4*, page 18.
- [Rooms et al., 2000b] Rooms, F., Philips, W., Meesen, G., Poffijn, A., and Van Oostveldt, P. (2000b). Using confocal microscopy and image restoration for the analysis of nuclear tracks. In *Proc. of 11th annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, pages 479–482.
- [Rooms et al., 2002a] Rooms, F., Philips, W., Meesen, G., Poffijn, A., and Van Oostveldt, P. (2002a). Estimating image blur in the wavelet domain. In *Student Forum abstract at ICASSP 2002, Orlando*, pages VI–4190.
- [Rooms et al., 2003a] Rooms, F., Philips, W., Meesen, G., Poffijn, A., and Van Oostveldt, P. (2003a). Simultaneous estimation and restoration of confocal images. In *Abstract book Bioforum 2003 (Liège, Belgium)*, page 20.
- [Rooms et al., 2001a] Rooms, F., Philips, W., Meesen, G., and Van Oostveldt, P. (2001a). The analysis of nuclear tracks with confocal microscopy and image restoration. In *Proc. of 5th WSES/IEEE World Multiconference on Circuits, Systems, Communications and Computers (CSCC 2001)*, pages 4871–4876.
- [Rooms et al., 2004] Rooms, F., Philips, W., and Portilla, J. (2004). Parametric psf estimation via sparseness maximization in the wavelet domain. In Truchetet, F. and Laligand, O., editors, *Proc. of SPIE Wavelet Applications in Industrial Processing II*, pages 26–33.
- [Rooms et al., 2003b] Rooms, F., Philips, W., and Van Oostveldt, P. (2003b). Estimation of anisotropic blur for the restoration of confocal images. In Truchetet, F., editor, *Proc. of SPIE Wavelet Applications in Industrial Processing I*, pages 131–136.

- [Rooms et al., 2003c] Rooms, F., Philips, W., and Van Oostveldt, P. (2003c). Integrated approach for estimation and restoration of photon-limited images based on steerable pyramids. In Grcić, M. and Grcić, S., editors, *Proc. of the 4th EURASIP Conference focused on Video / Image Processing and Multimedia Communications, Zagreb, Croatia*, pages 131–136.
- [Rooms et al., 2003d] Rooms, F., Philips, W., and Van Oostveldt, P. (2003d). Joint estimation and restoration of confocal images. In *Abstract book Focus on Microscopy 2003 (Genova, Italia)*, page 117.
- [Rooms et al., 2003e] Rooms, F., Philips, W., and Van Oostveldt, P. (2003e). Simultaneous estimation and restoration of confocal images. In *Abstract book of the 1st International meeting on Applied Physics (APHYS 2003) (Badajoz, Spain)*, page 499.
- [Rooms et al., 2001b] Rooms, F., Pizurica, A., and Philips, W. (2001b). Estimating image blur in the wavelet domain. In *Proc. of 12th annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, pages 568–572.
- [Rooms et al., 2001c] Rooms, F., Pizurica, A., and Philips, W. (2001c). Wavelet-based estimation of image blur. In *Proceedings 2nd Phd symposium of the Faculty of Applied Sciences at Ghent University*, page 62.
- [Rooms et al., 2002b] Rooms, F., Pizurica, A., and Philips, W. (2002b). Estimating image blur in the wavelet domain. In Suter, D. and Bab-Hadiashar, A., editors, *Proceedings of the Fifth Asian Conference on Computer Vision (ACCV)*, pages 210–215.
- [Rooms et al., 2002c] Rooms, F., Ronsse, M., Pizurica, A., and Philips, W. (2002c). Psf estimation with applications in autofocus and image restoration. In *Proceedings of the 3rd IEEE Benelux Signal Processing Symposium (SPS)*, pages 13–16.
- [Rudin et al., 1992] Rudin, L. I., Osher, S., and Fatami, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268.
- [Savakis and Trussell, 1993a] Savakis, A. and Trussell, H. J. (1993a). Blur identification by residual spectral matching. *IEEE Trans. Image Processing*, 2(2):141–151.
- [Savakis and Trussell, 1993b] Savakis, A. and Trussell, H. J. (1993b). On the accuracy of psf representation in image restoration. *IEEE Trans. Image Processing*, 2(2):252–259.
- [Schaeffer et al., 2001] Schaeffer, L., Schuster, D., and Hertz, H. (2001). Generalized approach for accelerated maximum likelihood based image restoration applied to three-dimensional fluorescence microscopy. *Journal of Microscopy*, 204(2):99–107.

- [Schechner and Kiryati, 1998] Schechner, Y. and Kiryati, N. (1998). Depth from defocus vs. stereo: How different really are they? In *Proc. International Conf. on Pattern Recognition*.
- [Seldin and Fienup, 1990] Seldin, J. and Fienup, J. (1990). Iterative blind deconvolution algorithm applied to phase retrieval. *Journal of the Optical Society of America A*, 7(3):428–433.
- [Selesnick, 2004] Selesnick, I. (2004). The double-density dual-tree dwt. *IEEE Trans. on Signal Processing*, 52(5):1304–1314.
- [Selesnick and Li, 1998] Selesnick, I. and Li, K. Y. (1998). Video denoising using 2d and 3d dual-tree complex wavelet transforms. In *Proc. SPIE vol. 5207: Wavelet Applications in Signal and Image Processing X*, pages 607–618.
- [Shepp and Vardi, 1982] Shepp, L. A. and Vardi, Y. (1982). Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Medical Imaging*, 1(2):113–122.
- [Shewchuk, 1994] Shewchuk, J. (1994). An introduction to the conjugate gradient method without the agonizing pain. <http://www-2.cs.cmu.edu/~jrs/jrspapers.html>.
- [Silverman et al., 1990] Silverman, B. W., Jones, M. C., Wilson, J. D., and Nychka, D. W. (1990). A smoothed em approach to indirect estimation problems, with particular reference to stereology and emission tomography. *Journal of the Royal Statistical Society, Series B*, 52(2):271–324.
- [Simoncelli, 1999] Simoncelli, E. (1999). Modeling the joint statistics of images in the wavelet domain. In *Proc. SPIE, vol. 3813: 44th Annual Meeting*, pages 188–195.
- [Simoncelli et al., 1992] Simoncelli, E. P., Freeman, W. T., Adelson, E. H., and Heeger, D. J. (1992). Shiftable multi-scale transforms. *IEEE Trans. Information Theory*, 38(2):587–607.
- [Slijpen and Beekman, 1999] Slijpen, E. T. P. and Beekman, F. J. (1999). Comparison of post-filtering and filtering between iterations for spect reconstruction. *IEEE Trans. Nuclear Science*, 46(6):2233–2238.
- [Snyder et al., 1992] Snyder, D., Schultz, T., and O’Sullivan, J. (1992). Deblurring subject to nonnegativity constraints. *IEEE Trans. Signal Processing*, 40(5):1143–1150.
- [Starck et al., 2000] Starck, J.-L., Murtagh, F., and Bijaoui, A. (2000). *Image Processing and Data Analysis, the multiscale approach*. Cambridge University Press.

- [Stippel et al., 2001] Stippel, G., Duskunovic, I., Pizurica, A., Rooms, F., Philips, W., and Lemahieu, I. (2001). A speckle suppression method for medical ultrasound images based on local statistics. In *Proc. of 11th annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, pages 525–530.
- [Stippel et al., 2000] Stippel, G., Duskunovic, I., Rooms, F., Philips, W., and Lemahieu, I. (2000). Segmenting 'flares' in ultrasound images. In *Proceedings 1st Phd symposium of the Faculty of Applied Sciences at Ghent University*, page 30.
- [Stollberg et al., 2003] Stollberg, H., Boutet de Monvel, J., Holmberg, A., and Hertz, H. (2003). Wavelet-based image restoration for compact x-ray microscopy. *Journal of Microscopy*, 211(2):154–160.
- [Strang and Nguyen, 1996] Strang, G. and Nguyen, T. (1996). *Wavelets and filter banks*. Wellesley - Cambridge Press.
- [UltraSharpen, WWW] UltraSharpen (WWW). Ultrasharpen. <http://www.ultrasharpen.com/>.
- [Uman Software, WWW] Uman Software (WWW). Autofocus. <http://www.humansoftware.com/pages1200/AutoFocus/HSAutofocus0.html>.
- [University Oklahoma, WWW] University Oklahoma (WWW). <http://teachplants.okstate.edu/lgsdlgm/lentil/lentil.html>.
- [van der Voort, 1989] van der Voort, H. (1989). *Three dimensional Image Formation and Processing in Confocal Microscopy*. PhD thesis, University of Amsterdam.
- [Van der Voort, WWW] Van der Voort, H. (WWW). Image restoration: getting it right. Online presentation at <http://www.svi.nl/education/>.
- [van der Voort and Strasters, 1995] van der Voort, H. and Strasters, K. (1995). Restoration of confocal images for quantitative image analysis. *Journal of Microscopy*, 178:165–181.
- [van Kempen, 1999] van Kempen, G. (1999). *Image restoration in Fluorescence Microscopy*. PhD thesis, Delft University.
- [van Kempen et al., 1997] van Kempen, G., van Vliet, L., and Verveer, P. (1997). Application of image restoration methods for confocal fluorescence microscopy. In Cogswell, C., Conchello, J.-A., and Wilson, T., editors, *Proc. SPIE, vol. 2984: 3-D Microscopy: Image Acquisition and Processing IV*, pages 114–124.
- [Van Kempen et al., 1997] Van Kempen, G., van Vliet, L., Verveer, P., and Van der Voort, H. (1997). A quantitative comparison of image restoration methods for confocal microscopy. *Journal of Microscopy*, 185(3):354–365.

- [Venkatesh and Owens, 1990] Venkatesh, S. and Owens, R. (1990). On the classification of image features. *Pattern Recognition Letters*, 11:339–349.
- [Verveer, 1998] Verveer, P. (1998). *Computational and optical methods for improving resolution and signal quality in fluorescence microscopy*. PhD thesis, Delft University.
- [Verveer et al., 1999] Verveer, P., Gemkow, M., and Jovin, T. (1999). A comparison of image restoration approaches applied to three-dimensional confocal and wide-field fluorescence microscopy. *Journal of Microscopy*, 193:50–61.
- [Verveer and Jovin, 1997] Verveer, P. and Jovin, T. (1997). Acceleration of the ictm algorithm. *Journal of Microscopy*, 188:191–195.
- [Vincent and Soille, 1991] Vincent, L. and Soille, P. (1991). Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(6):583–597.
- [Vivirito et al., 2002] Vivirito, P., Battiatto, S., Curti, S., Cascia, M., and Pirrone, R. (2002). Restoration of out of focus images based on circle of confusion estimate. In *Proceedings of SPIE vol 4790, 47th Annual Meeting 2002.*, pages 408–416.
- [Weisstein, WWWa] Weisstein, E. W. (WWWa). <http://mathworld.wolfram.com/LeastSquaresFittingExponential.html>.
- [Weisstein, WWWb] Weisstein, E. W. (WWWb). Correlation coefficient (from mathworld—a wolfram web resource). <http://mathworld.wolfram.com/CorrelationCoefficient.html>.
- [Weisstein, WWWc] Weisstein, E. W. (WWWc). Optical aberration (from mathworld—a wolfram web resource). <http://scienceworld.wolfram.com/physics/OpticalAberrations.html>.
- [Weisstein, WWWd] Weisstein, E. W. (WWWd). Uncertainty principle (from mathworld—a wolfram web resource). <http://scienceworld.wolfram.com/physics/UncertaintyPrinciple.html>.
- [Wikipedia, WWW] Wikipedia (WWW). Antoni van Leeuwenhoek. http://nl.wikipedia.org/wiki/Antoni_van_Leeuwenhoek.
- [Xiong and Shafer, 1993] Xiong, Y. and Shafer, S. (1993). Depth from focusing and defocusing. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*.
- [Yu et al., 2000] Yu, W., Daniilidis, K., and Sommer, G. (2000). A new 3d orientation steerable filter. In *DAGM-Symposium Mustererkennung Kiel*, pages 203–212.

