

LEHRSTUHL FÜR BIOMEDIZINISCHE PHYSIK
TECHNISCHE UNIVERSITÄT MÜNCHEN

STATISTICAL ITERATIVE RECONSTRUCTION
FOR X-RAY PHASE-CONTRAST COMPUTED TOMOGRAPHY

DIETER HAHN
PHD THESIS

TECHNISCHE UNIVERSITÄT MÜNCHEN
Physik Department
Lehrstuhl für Biomedizinische Physik

Statistical Iterative Reconstruction for X-ray Phase-Contrast Computed Tomography

Dieter Hahn

Vollständiger Abdruck der von der Fakultät für Physik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr. M. Zacharias

Prüfer der Dissertation: 1. Univ.-Prof. Dr. F. Pfeiffer

2. Univ.-Prof. Dr. J. Wilkens

Die Dissertation wurde am 18.03.2014 bei der Technischen Universität München eingereicht und durch die Fakultät für Physik am 22.04.2014 angenommen.

Abstract

Computed tomography on the basis of grating-based phase-contrast imaging has the potential to significantly expand clinical diagnostic capabilities. This novel imaging method not only offers the conventional absorption signal, but also delivers information on the refraction and the scattering of X-rays. The additional contrast mechanisms provide increased soft-tissue contrast that is hardly accessible with conventional CT and new complementary information. Applications include, for example, early tumor detection, mammography or lung imaging. Correlated with this increased contrast is, however, an increase in radiation dose due to the grating structures, which are necessarily placed in the beam path for the measurement and which absorb about half of the incident X-rays.

Currently a lot of work in conventional CT research is put into reducing radiation exposure, so dose reduction has become an even more important topic in grating-based imaging. It can be achieved by either lowering the exposure time of individual acquisitions or by reducing the overall number of recorded angular views. However, both of these efforts decrease the applicability of the current state-of-the-art reconstruction technique, the filtered back projection, as they violate the requirements needed for an analytical solution to the reconstruction problem.

We address these problems with the development of a statistical reconstruction algorithm for differential, grating-based phase-contrast CT, which incorporates the statistical information contained in the scattering signal in order to achieve an increase in overall image quality of the phase reconstruction. Consequently, we demonstrate the algorithm's capability to reduce noise and artifacts caused by incomplete input data on exemplary data sets obtained by simulations and experiments. This generally applicable algorithm constitutes the first of the two main results of this thesis.

In a more sophisticated application we use the algorithm to reduce artifacts specific to differential phase contrast, which appear, when phase wrapping occurs at the boundary between structures with low and high density, e.g. between soft tissue and bones. These—appropriately termed—bone artifacts have a severe impact on the diagnostic usability of reconstructions. We show that, by also adding the information from the absorption signal to the statistical algorithm, the information obstructed by the artifacts can be recovered to a high degree.

As a result of the present thesis, statistical reconstruction for differential phase-contrast CT is considered a valuable technique for improving the quality of reconstructed images and thus expanding the current types of applications for grating-based imaging with the ultimate goal of translating the technique into a routinely used clinical tool.

Zusammenfassung

Computertomografie auf Basis der gitterbasierten Phasenkontrastbildung kann potentiell dazu verwendet werden, die Möglichkeiten der klinischen Diagnostik signifikant zu erweitern. Diese neuartige Bildgebungsmethode liefert nicht nur das konventionelle Absorptionssignal, sondern ist auch in der Lage, die Brechung und auch die Streuung von Röntgenstrahlen sichtbar zu machen. Dieser zusätzliche Kontrastmechanismus führt zu einem erheblich höheren Bildkontrast in Weichgewebe, der in konventioneller Computertomografie nur schwerlich zu erreichen ist und zu neuer komplementärer Information. Mögliche Anwendungen beinhalten zum Beispiel Früherkennung von Tumoren, Mammografie und Lungenbildung. Mit diesem höheren Kontrast ist allerdings auch eine erhöhte Strahlenbelastung für das untersuchte Objekt verbunden, die mit den Gitterstrukturen zusammenhängt. Diese müssen notwendigerweise für eine Messung im Strahlengang platziert werden und absorbieren ungefähr die Hälfte der einfallenden Strahlung.

Aktuell wird in der konventionellen Computertomografie viel Arbeit investiert, um die Strahlenbelastung von Patienten zu minimieren, und so kommt diesem Thema in der gitterbasierten Bildgebung eine noch größere Bedeutung zu. Diese Minimierung kann einerseits über eine Verringerung der Belichtungszeit der einzelnen Aufnahmen erreicht werden, andererseits indem weniger Projektionen über denselben Winkelbereich aufgenommen werden. Diese beiden Ansätze führen jedoch dazu, dass der momentane Standardalgorithmus für die Rekonstruktion, die gefilterte Rückprojektion, nurmehr schlecht angewendet werden kann, da sie die Voraussetzungen für eine analytische Lösung des Rekonstruktionsproblems verletzen.

Wir gehen diese Probleme durch die Entwicklung eines statistischen Rekonstruktionsalgorithmus für gitterbasierte Phasenkontrasttomografie an, der die im Streusignal enthaltenen Informationen mit verwendet, um die allgemeine Bildqualität des Phasensignales zu verbessern. Im Anschluss demonstrieren wir, wie der Algorithmus helfen kann, Rauschen und Artefakte durch unvollständige Daten zu verringern. Dazu bedienen wir uns sowohl simulierter als auch experimentell gemessener Rohdaten. Dieser allgemein anwendbare Algorithmus stellt das erste der zwei Hauptergebnisse dieser Arbeit dar.

In einer spezielleren Anwendung verwenden wir den Algorithmus zur Reduktion von Artefakten, die im differentiellen Phasenkontrast besonders an der Grenze zwischen Strukturen mit niedriger und hoher Dichte entstehen, zum Beispiel zwischen Weichgewebe und Knochen. Diese Knochenartefakte führen zu einer eingeschränkten diagnostischen Verwendung der finalen Bilder. Wir zeigen, dass durch die zusätzliche Einbindung des Absorptionssignals in den Algorithmus die durch die Artefakte überlagerte Bildinformation zu einem Großteil wieder sichtbar gemacht werden kann.

Im Ergebnis wurde in der vorliegenden Arbeit gezeigt, dass statistische Rekonstruktion für differentielle Phasenkontrasttomografie eine wertvolle Methode darstellt, um die Qualität rekonstruierter Bilder entscheidend zu verbessern und dadurch die Anwendungsbereiche gitterbasierter Bildgebung in Richtung des klinischen Einsatzes zu erweitern.

Table of Contents

Acronyms	1
Symbols	3
List of Figures	5
List of Tables	7
1 Introduction	9
1.1 Motivation	9
1.2 Outline	11
2 Theory	13
2.1 X-ray interaction with matter	13
2.2 Grating-based phase-contrast imaging	17
2.2.1 The Talbot effect	18
2.2.2 Grating interferometer	20
2.2.3 Signal extraction	22
2.3 Tomographic reconstruction	26
2.3.1 Fourier slice theorem	26
2.3.2 Filtered back projection	28
2.4 Iterative reconstruction methods	31
3 Data processing and analysis of projection data	35
3.1 Comparison of attenuation and phase projections	36
3.2 Statistical phase retrieval	44
3.3 Influence of the system PSF on reconstruction	48
3.4 Summary and outlook	55
4 Statistical reconstruction framework	57
4.1 Motivation	57
4.2 Statistical reconstruction algorithm (SIR)	58
4.2.1 Model	58
4.2.2 Regularization	66

4.2.3	Minimization	69
4.3	Evaluation of the algorithm	73
4.3.1	Synthetic data	73
4.3.2	Experimental data	80
4.3.3	Synchrotron data	88
4.3.4	Further analysis	89
4.4	Moving towards correct modelling of the cost function	93
4.5	Summary	96
5	Imaging application	97
5.1	Bone artifact reduction (BAR)	97
5.1.1	Motivation	97
5.1.2	Bone artifacts: causes and reduction	99
5.1.3	Reduction of bone artifacts on a whole mouse	102
5.1.4	Artifact reduction on a carotid artery	107
5.2	Reconstruction of large objects demonstrated on a sliced pig .	115
5.2.1	Experimental setup	115
5.2.2	Preprocessing and reconstruction results	117
5.3	Summary	121
6	Conclusions and outlook	123
A	Derivations and implementations	135
A.1	Implementation of the statistical phase retrieval	135
A.2	Derivation of Richardson-Lucy deconvolution	137
A.3	Gradient and step size calculation of regularization terms . . .	139
	Publications and scientific presentations	141
	Acknowledgements	145

Acronyms

ART	algebraic reconstruction technique
BAR	bone artifact reduction
CG	conjugate gradient
CNR	contrast-to-noise ratio
CT	computed tomography
DPC	differential phase contrast
ESRF	European Synchrotron Radiation Facility
FBP	filtered backprojection
GBI	grating-based imaging
GPGPU	general purpose graphics processing unit
IR	iterative reconstruction
LS	least-squares
MAP	maximum-a-posteriori
ML	maximum likelihood
NLCG	non-linear conjugate gradient
NPS	noise power spectrum
PCCT	phase-contrast computed tomography
PDF	probability density function
PSF	point spread function
RCG	relative contrast gain
RL	Richardson-Lucy
SART	simultaneous algebraic reconstruction technique
SIR	statistical iterative reconstruction
SIRT	simultaneous iterative reconstruction technique
SNR	signal-to-noise ratio
SPR	statistical phase retrieval
SPS	signal power spectrum
WLS	weighted least squares

Symbols

a	attenuation signal
β	imaginary part of the complex index of refraction
β_{PR}	Polak-Ribiere formula for calculation of new search direction in conjugate gradient (CG)
δ	real part of the complex index of refraction
G_0	source grating
G_1	phase grating
G_2	analyzer grating
$\mathcal{L}(A B)$	likelihood of event A given event B
$l(A B)$	logarithm of the likelihood of event A given event B
μ	linear attenuation coefficient
$p(A B)$	probability of event A given that event B has occurred
ϕ	differential phase signal
$\mathcal{P}_\theta(x')$	projection of $f(x', y')$ along y' in a coordinate system rotated by θ
T	transmission
v	visibility

List of Figures

2.1	Illustration of X-rays interacting with matter	14
2.2	Illustration of the Talbot effect	19
2.3	Schematic view of the functional principle of a grating interferometer	20
2.4	Phase-stepping curves for attenuation, refraction and scattering	21
2.5	Grating interferometer tomography setup	23
2.6	Illustration of the signal extraction from stepping curves with sample (object scan, blue) and without sample (blank scan, orange)	24
2.7	Illustration of the Fourier slice theorem.	26
2.8	Sampling of Fourier space in tomography.	28
3.1	Experimental data used in the relative contrast gain (RCG) analysis. (A) photograph of the sample, (B) attenuation signal, (C) differential phase-contrast signal, (D) differential attenuation signal.	39
3.2	Power spectra of differential phase-contrast and differential attenuation projections	40
3.3	Visual representation of the RCG analysis of regions A (left) and B (right).	42
3.4	Stepping curve with parameters of statistical phase retrieval .	45
3.5	Demonstration of deconvolution on raw stepping image.	50
3.6	Comparison of the three contrast channels retrieved from original and deconvolved data.	52
3.7	Comparison of phase contrast tomographic reconstructions from original and deconvolved raw data.	53
3.8	PSF from detector and source.	54
4.1	Implementation of forward and back projection.	60
4.2	Illustration of unconstrained reconstruction.	64
4.3	Illustration of the principal of local regularization.	67

4.4	Comparison of quadratic and Huber potentials.	68
4.5	Numerical phantom for algorithm evaluation.	74
4.6	Reconstruction results from simulated phantom (part 1).	76
4.7	Reconstruction results from simulated phantom (part 2).	79
4.8	Illustration of three-grating interferometer.	81
4.9	Overview of heart sample.	82
4.10	Reconstruction results from heart sample (part 1)	84
4.11	Reconstruction results from heart sample (part 2)	87
4.12	Tomographic reconstruction of a rat brain.	88
4.13	Convergence of phantom reconstructions.	89
4.14	Convergence of heart reconstructions.	91
4.15	Signal-to-noise spectrum of phantom.	92
4.16	Signal-to-noise spectrum of heart sample.	93
4.17	Wrapped normal distribution.	94
5.1	Comparison of metal artifacts and bone artifacts.	98
5.2	DPC sinogram of a mouse.	100
5.3	Calculation of modification to the statistical weights for phase- wrapping artifact reduction.	101
5.4	Bone artifact reconstruction results for mouse sample.	104
5.5	Analysis of streak artifact reduction.	105
5.6	Enlarged view of bone artifact reduction results.	106
5.7	DPC sinogram of mouse sample after reconstruction.	107
5.8	Sinogram and statistical weights of artery sample.	108
5.9	Comparison of FBP and SIR for artery sample.	110
5.10	Comparison of SIR and BAR for artery sample.	112
5.11	Comparison of SIR and BAR (coronal).	112
5.12	Comparison of SIR and BAR (sagittal).	113
5.13	Photograph of sliced pig sample.	115
5.14	Illustration of measurement setup for sliced pig sample.	116
5.15	DPC projections of sliced pig sample.	117
5.16	Comparison of FBP and SIR for sliced pig sample.	118
5.17	Comparison of SIR and FBP-SIR-mixture for sliced pig sample.	119
5.18	Comparison of FBP-SIR-mixture and clinical scan of sliced pig sample.	120

List of Tables

4.1	This table gives the numerical values of the phantom that was used in the simulations. The index corresponds to the symbols used in figure 4.5.	75
4.2	This table gives the reconstruction parameters for the algorithm verification with synthetic data.	78
4.3	This table gives the reconstruction parameters for the algorithm verification with experimental data.	83
5.1	Results of the standard deviation analysis of several regions of interest in figure 5.4	107

Chapter 1

Introduction

In this chapter we briefly introduce the context, in which this thesis is placed, and we describe its contribution to current research. In addition we give an outline of the structure of the thesis.

1.1 Motivation

Since their discovery by Röntgen in 1895 (Röntgen, 1898) the interest in X-rays has steadily increased because of their ability to penetrate matter. With the advent of computed tomography (CT) (Hounsfield, 1973) it became even possible to produce cross-sectional images of the inner structure of objects, taking X-ray imaging into the third dimension. Conventionally, X-ray CT makes use of the absorption of X-rays in different materials. X-ray images are taken from many angles around the investigated object and computer programs allow the reconstruction of the inner structure. This procedure is routinely used today for medical diagnosis, non-destructive testing and in scientific research. In general, good contrast is achieved between materials with a high density difference, for example between bones and soft tissue.

Contrast between different types of soft tissue, however, is limited due to similar attenuation properties. This soft-tissue contrast can be made visible by phase-sensitive imaging methods, that rely on measuring the refraction of X-rays instead of their absorption as they pass through matter. Several different techniques have been developed over the last decades, such as propagation-based methods (Snigirev et al., 1995), analyzer-based methods (Davis et al., 1995; Ingal and Beliaevskaya, 1995) and methods based on

interferometric effects (Momose, 2003; Weitkamp et al., 2005; Pfeiffer et al., 2006). The technique used in this thesis belongs to this last category and is called grating-based imaging (GBI). It uses gratings to translate phase changes into intensity variations that can be measured with conventional X-ray detectors. With this technique we can measure the refraction of an object, i.e. the phase gradient perpendicular to the gratings, giving the method its other common name: **differential phase-contrast imaging**. GBI has been shown to yield soft-tissue contrast that is significantly higher than in absorption imaging (Hahn et al., 2012; Sztrókay et al., 2013; Tapfer et al., 2013; Schleede et al., 2013) and to deliver additional and complementary information (Herzen et al., 2009; Tapfer et al., 2012; Willner et al., 2013).

Tomographic reconstruction from both attenuation and differential phase-contrast images is typically accomplished using the filtered backprojection (FBP) algorithm, adapted with an appropriate filter in the latter case (Pfeiffer et al., 2007b). The FBP is an analytical solution of the continuous reconstruction problem and several simplifying assumptions are made in its derivation, such as acquisitions being noiseless and continuously sampled. Thus, in reality, the FBP is sensitive to noise and especially ill-behaved when applied to irregularly sampled or incomplete data sets (Kak and Slaney, 1988; Hsieh, 2009). It is, however, well established that iterative algorithms—especially if statistical information is taken into account—can outperform the FBP to reconstruct both from projections with low statistics and few-view tomographic scans (Fessler, 2000; Kohler et al., 2011; Xu et al., 2012; Beister et al., 2012).

The aim of this thesis is the improvement of the image quality in phase-contrast computed tomography (PCCT) reconstructions for scenarios preventing the use of the FBP. We develop a **statistical iterative reconstruction (SIR)** algorithm for PCCT and for the inclusion of the statistical properties of the input data an alternative method for processing the measurements is presented. A demonstration of the performance of the algorithm is given with studies done on simulated and experimental PCCT data, incorporating typical cases of corrupted measurements.

The main motivation for the development of this algorithm—apart from general improvements of image quality—is the reduction of artifacts specific to grating-based PCCT. In appearance they are similar to metal artifacts known from conventional CT (see DeMan (2001) for an extensive analysis of these artifacts) and they share a similar cause, namely photon starvation due to a high degree of absorption and scattering. An additional cause in PCCT is related to the measurement of the gradient of the phase shifts, which is defined on the unit circle and thus intrinsically ambiguous. Projections of very large

gradients exhibit **phase wrapping**, leading to statistically unreliable values. As phase wrapping is prominently visible at boundaries between bones and soft tissue, the artifacts are given the name **bone artifacts**. We will demonstrate how SIR can be used to mitigate the effects of these artifacts.

1.2 Outline

This thesis is structured in the following way. **Chapter 2** will introduce the theoretical key concepts. It covers basic interactions of X-rays with matter and explains the principles of grating-based phase-contrast imaging. After that follows an introduction to tomographic reconstruction from an analytical standpoint as well as a brief overview of advanced numerical techniques. **Chapter 3** is the first of three results chapters and focuses on improvements in the domain of individual projections. These improvements include a technique to compare differential and non-differential signals, an improved signal extraction from the raw data and a study involving the effects of the measurement setup's point spread function on the differential phase-contrast signal. **Chapter 4** then represents the main part of this thesis and gives the derivation and evaluation of a statistical iterative reconstruction algorithm for grating-based phase-contrast CT. **Chapter 5** demonstrates the reconstruction algorithm on current typical imaging problems that arise in the process of developing grating-based phase-contrast CT in the direction of human clinical applications. First, the algorithm is used to reduce artifacts that arise specifically in the differential phase signal and are caused by bones and other dense materials. The second part deals with the imaging of large objects and demonstrates how advanced reconstruction algorithms can help to improve image quality. The thesis concludes with **Chapter 6**, which summarizes the scientific results and gives an outlook to further research.

Chapter 2

Theory

In this chapter we will give a summary of the theoretical key concepts needed for a basic understanding of the contents of the next chapters. First, we will briefly discuss the different ways, in which X-rays can interact with matter. After that we dedicate a section to describing grating-based phase-contrast imaging, the experimental technique used throughout the thesis. This section is followed by an introduction to the state-of-the-art reconstruction algorithm—the filtered back projection, which will serve as the reference reconstruction method for comparison with our new algorithm. The chapter ends with an overview of established methods for reducing metal artifacts in conventional CT.

2.1 X-ray interaction with matter

X-rays can interact with matter in different ways. These interactions include photoelectric absorption, elastic scattering, inelastic scattering and pair production. The last one does not play a role for the X-ray energies in the context of the thesis and a discussion is therefore omitted. Photoelectric absorption happens, when an X-ray photon is absorbed by an atom, transferring its energy to an electron and expelling it from the atom. The energy range of X-rays causes the interactions to occur mostly with the inner shell electrons of the atom. The photoelectric effect is the main mechanism behind the attenuation of X-rays for the energies that are in the scope of this thesis. Another source of attenuation is inelastic scattering or Compton scattering, where a photon loses part of its energy when scattered by an electron.




Figure 2.1: Phase shift, attenuation and refraction of a wave traveling through a medium with complex refractive index n . In a medium with a real part of n smaller than 1 the wave travels faster than in vacuum, leading to a phase shift of $\Delta\phi$ as indicated by the blue lines. The amplitude of the wave is decreased by ΔA due to the imaginary part of n . This is indicated by the orange lines. The phase difference of both waves causes a refraction by an angle α as illustrated by the red lines.

While photoelectric absorption and Compton scattering require a treatment in the particle picture, elastic scattering or Thomson scattering can be explained classically in the wave picture. When an X-ray wave interacts with one of the electrons of an atom, the electric field of the wave causes the electron to oscillate and in turn emit radiation of the same wavelength as the incoming wave. For typical X-ray energies the emitted wave is phase shifted by π in relation to the incident wave. In macroscopic terms the superposition of many of these scattering events causes the phase velocity of the wave traveling through the medium to be increased compared to a wave traveling through vacuum (Als-Nielsen and McMorrow, 2011). This difference in phase velocity in turn leads to the wave that passes the medium to be phase shifted by $\Delta\phi$ relative to the other wave and consequently to refraction. This is illustrated in figure 2.1 by the dashed blue lines. The refraction angle is indicated by α and the dashed red lines.

For a quantitative and combined description of refraction and attenuation

effects we can use the phenomenological notion of a complex index of refraction, which is well known from imaging with visible light. It is defined as

$$n = 1 - \delta + i\beta, \quad (2.1)$$

where the real part describes scattering and refraction effects and the imaginary part describes attenuation effects. The two parameters δ and β are called refractive index decrement and absorption index, respectively (Willmott, 2011). The main difference between visible light and X-rays is, however, that for X-rays the real part of n —i.e. $(1 - \delta)$ —is smaller than unity, corresponding to the increased phase velocity of a wave while penetrating matter.

That the real and imaginary part describe phase shift and absorption, respectively, can easily be verified by calculating the propagation of a plane wave with wave vector \mathbf{k} through a medium with refractive index n

$$\Psi(\mathbf{r}) = \Psi_0 \cdot e^{in\mathbf{k}\cdot\mathbf{r}} = \Psi_0 \cdot e^{i(1-\delta)\mathbf{k}\cdot\mathbf{r}} \cdot e^{-\beta\mathbf{k}\cdot\mathbf{r}}, \quad (2.2)$$

where the first exponential on the right hand side obviously represents a simple phase factor and the second exponential represents a decay of the wave's amplitude. The phase shift $\Delta\phi$ (and refraction α) and attenuation ΔA of a plane wave traveling through a medium are visualized in figure 2.1.

From equation (2.2) we can derive expressions for the absorption, the phase shift and the refraction angle as functions of β and δ , respectively.

Let $I_0 = |\Psi_0|^2$ be the intensity of the incident wave and $\Psi = \Psi_0 e^{-\beta kL}$ the amplitude of the attenuated wave after traveling a distance L inside the medium. Then

$$I = |\Psi|^2 = |\Psi_0 e^{-\beta kL}|^2 = |\Psi_0|^2 |e^{-\beta kL}|^2 = I_0 e^{-2\beta kL} \quad (2.3)$$

represents the intensity of the wave after attenuation. The relative transmission, i.e. the fraction of the wave's intensity that is left after attenuation, is then given accordingly as

$$T = \frac{I}{I_0} = e^{-2\beta kL} \quad (2.4)$$

and with the linear absorption coefficient $\mu = 2\beta k$ we get the well known Beer-Lambert law

$$I = I_0 e^{-\mu L}. \quad (2.5)$$

This relation holds for homogeneous media. For inhomogeneous media we have to take into account the spatial distribution of the linear absorption coefficient and can rewrite the previous equation as

$$I = I_0 e^{-\int \mu(x,y,z) dy}. \quad (2.6)$$

For the refraction we can find a similar expression. A slight reformulation of the phase-shift term of equation (2.2) gives

$$\Psi_P = \Psi_0 e^{i\mathbf{k}\cdot\mathbf{r}} e^{-i\delta\mathbf{k}\cdot\mathbf{r}}. \quad (2.7)$$

With the knowledge that a plane wave propagating through vacuum is given by

$$\Psi_v = \Psi_0 e^{i\mathbf{k}\cdot\mathbf{r}} \quad (2.8)$$

we can immediately see that the total phase shift is given by

$$\Delta\phi = \delta\mathbf{k}\cdot\mathbf{r} \quad (2.9)$$

Again, this relation only holds for a single material. For inhomogeneous media we get

$$\Delta\phi = \int \delta(x, y, z) k_y dy. \quad (2.10)$$

From figure 2.1 we can see that the phase difference between the two waves causes a change of the wave front's wave vector by the refraction angle α , which we can calculate by

$$\tan(\alpha(x, y, z)) \approx \alpha(x, y, z) = \frac{\lambda}{2\pi} \frac{\partial\Phi}{\partial x} \quad (2.11)$$

or by inserting equation (2.10)

$$\alpha(x, z) = \frac{\partial}{\partial x} \int \delta(x, y, z) dy. \quad (2.12)$$

A more in depth treatment of phase-shifting effects of X-rays can be found in Paganin (2006).

The last equation shows that refraction only occurs, when there is a gradient of the phase along the wave front. While it is impossible to directly measure the phase shifts with current detector technology, measuring the refraction angles is possible, albeit very difficult due to their small scale. To get a good estimate of the magnitude of the angles, we can use Snell's law

$$\frac{\sin \alpha_i}{\sin \alpha_r} \approx \frac{\alpha_i}{\alpha_r} = \frac{n_r}{n_i}, \quad (2.13)$$

where the subscripts i and r denote incoming and refracted wave, respectively. Correspondingly, n_i represents the refractive index of the medium before refraction occurs, n_r of the medium, into which the wave is refracted. Now let us assume that the wave is refracted, while propagating from a purely phase-shifting medium into vacuum, i.e. $n_r = 1$ and $n_i = 1 - \delta$. The refraction angle α —the change of propagation direction of the incoming wave—is then given by

$$\alpha = \alpha_r - \alpha_i \quad (2.14)$$

$$= n_i \alpha_i - \alpha_i \quad (2.15)$$

$$= (1 - \delta) \alpha_i - \alpha_i \quad (2.16)$$

$$= -\delta \alpha_i \quad (2.17)$$

For typical experimental X-ray energies on the order of 10 keV up to 100 keV, δ ranges from about 10^{-8} to 10^{-6} (Chantler, 1995). Subsequently, refraction angles for X-rays will be on the order of micro radians and even smaller and are not visible in conventional absorption-based X-ray imaging.

One particular method for accurately measuring these small angles is grating interferometry, which is the imaging technique solely used throughout this thesis. The basics of grating-based imaging (GBI) are described in the following section. Information on other phase-sensitive imaging methods can for example be found in Bech (2009), Willmott (2011), Als-Nielsen and McMorrow (2011) or Schleede (2013).

2.2 Grating-based phase-contrast imaging

In this section we will introduce the concepts and technology behind grating-based phase-contrast imaging. GBI uses interference effects to create a spatial reference pattern, relative to which the refraction angles can be determined. When refraction occurs in an object this reference pattern is distorted locally proportional to the magnitude of the refraction. Thus, phase changes are basically translated into changes of intensity, which can be measured with standard X-ray detectors and analyzed afterwards. In addition to this phase information, GBI also provides the conventional absorption signal and the so-called dark field signal, which is caused by ultra small-angle scattering.

2.2.1 The Talbot effect

Grating-based imaging uses the Talbot effect to create the spatial reference pattern needed to determine the refraction angles. This effect is named after Henry Fox Talbot, who discovered that self images of an absorbing periodic structure are formed at certain distances away from the structure, when illuminated with coherent light (Talbot, 1836). The first of these distances, denoted Talbot distance d_T , is given by

$$d_T = \frac{2p^2}{\lambda}, \quad (2.18)$$

where p is the periodicity of the structure and λ the wavelength of the incident radiation. However, the occurrence of the Talbot effect is not restricted to absorbing structures, but can also be observed for structures that periodically modulate the phase of the incoming wave. In this case the wave's intensity distribution directly behind the structure and at the Talbot distances is flat. Instead, periodic intensity modulations can be observed at fractional Talbot distances.

Figure 2.2 shows simulations of the intensity distribution created by an ideal absorption grating (A) and two ideal phase gratings that periodically shift the phase of the incoming wave by π (B) and $\pi/2$ (C), respectively. The intensity patterns, or Talbot carpets, are calculated by free-space propagation (see for example Goodman (2004); Paganin (2006); Als-Nielsen and McMorrow (2011)) of a monochromatic wave with full transverse coherence up to the full Talbot distance. As expected, in the case of an absorbing grating the self image is reproduced at the full Talbot distance, while for the phase-shifting gratings the periodic structure is reproduced at fractional Talbot distances. The blue and orange lines indicate exemplary fractional Talbot distances, where the intensity pattern is box-like and periodic. We can see that for the π -grating the period of the intensity modulation is half as large as the grating period and for the $\pi/2$ -grating both periods are exactly the same. The general formulation for these distances is given by (Weitkamp et al., 2006)

$$d_T = \frac{1}{\eta^2} \frac{np^2}{2\lambda} \quad \forall \text{ odd } n, \quad (2.19)$$

where n is called the order of the fractional Talbot distance and η distinguishes between the phase shift of the gratings and also determines the period of the intensity modulations

$$\eta = \begin{cases} 1 & \text{for } \Delta\phi = \pi/2 \\ 2 & \text{for } \Delta\phi = \pi \end{cases}. \quad (2.20)$$




Figure 2.2: Simulations of the intensity patterns that are formed when a plane X-ray wave propagates through a periodic structure. (A) interference pattern behind a pure absorption grating; (B) interference pattern behind a pure phase grating with a phase shift of π ; (C) interference pattern behind a pure phase grating with a periodic phase shift of $\pi/2$. Indicated with the blue and orange line are exemplary fractional Talbot distances, where a binary intensity modulation occurs.

With the above definition of n , figure 2.2 illustrates the 3rd fractional Talbot order for the π -grating with the blue line and the 1st order for the $\pi/2$ -grating with the orange line.

The intensity patterns at the fractional Talbot distances* provide the spatial reference for determining the refraction angles in GBI as we will see in the following section.

*The intensity modulations are binary only in the ideal case of full longitudinal and transverse coherence. In reality, a finite source size and a finite energy bandwidth and thus a limited coherence will lead to a smearing of the intensity pattern.

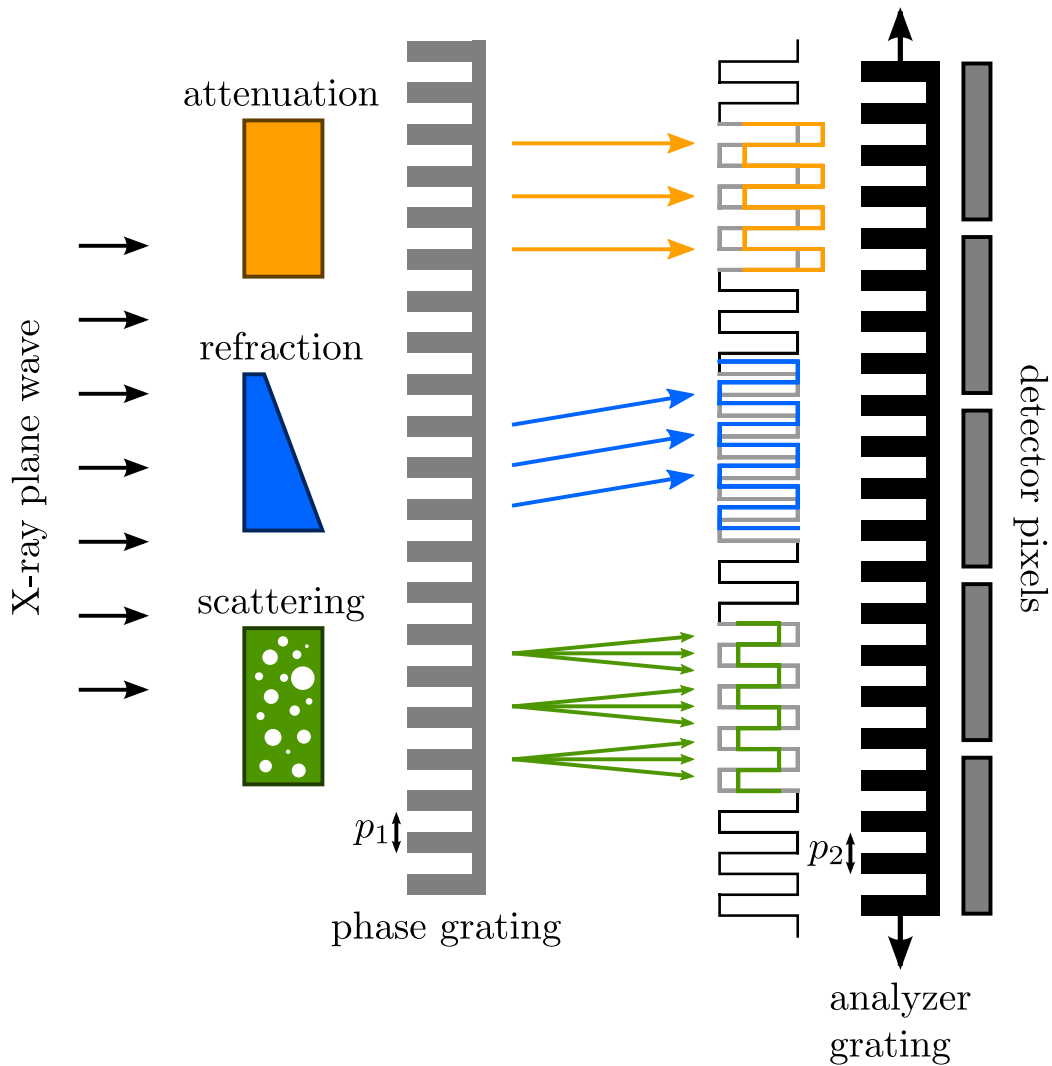


Figure 2.3: Functional principle of a grating interferometer. The figure illustrates the changes of the interference pattern created by the phase grating according to the Talbot effect, when the wave front is distorted by an ideal object. Attenuation leads to a drop of the average intensity of the pattern, phase shifts cause a lateral displacement of the pattern and scattering reduces the amplitude of the oscillation.

2.2.2 Grating interferometer

The experimental setup of choice for this thesis for measuring the refraction angles is the grating interferometer, also called Talbot interferometer..

The functional principle of the interferometer is visualized in figure 2.3. We first describe the way the interferometer works when there is no sample in

the beam. An X-ray wave propagating from the left passes the phase grating G_1 , represented in the figure by the gray grating structure, which creates a binary intensity modulation downstream according to the Talbot effect as explained in the previous section. This unmodified intensity pattern is sketched as the black box function to the right of the phase grating. The limited transverse coherence of typical X-ray sources require grating periods on the order of micrometers and, according to the Talbot effect, the period of the interference pattern is on the same scale. A detector with sufficiently small pixels can theoretically resolve the pattern directly, but usually X-ray detectors with a lot larger pixels—commonly on the order of $100\ \mu\text{m}$ —are used. To get information on the interference pattern, a second grating is inserted into the beam, which is accordingly termed analyzer grating and has a period matched to the period of the interference pattern. It is placed at a fractional Talbot distance, i.e. at the position where the intensity modulation is strongest, and consists of grating bars made from highly absorbing material, e.g. gold (Kenntner, 2012) to block parts of the intensity pattern. By moving this grating perpendicular to the beam direction and its grating lines over the distance of its periodicity and recording images at each position, the shape of the interference pattern can be sampled. This is commonly called phase-stepping scan and will result in a periodic intensity curve recorded for each detector pixel. In the ideal case of a box-shaped interference pattern, this intensity curve will represent the result of a convolution of the intensity pattern and the shape of the analyzer grating, i.e. the convolution of two box functions. The resulting intensity curve is then a triangular function. However, in a real application with limited coherence and finite source size the interference pattern will be blurred and the resulting curve of the phase-stepping scan then approximately represents a sine (Bech, 2009).

If a sample is placed in the beam the wave front is locally distorted depending on its index of refraction. This distortion then changes the shape of the




Figure 2.4: Visualization of phase-stepping curves for a purely attenuating (left), refracting (middle) and scattering (right) object. The black curve in each plot corresponds to a blank scan, the colored curves to the signal scans, respectively. The dashed lines indicate the quantity of interest for the three signals, i.e. mean value, phase and amplitude.

interference pattern. According to the fundamental interactions of X-rays with matter, three distinct effects can happen to the interference pattern. These effects are sketched in figure 2.3 for different idealized objects.

Attenuation: The orange sample locally attenuates the X-ray wave, causing the mean intensity of the interference pattern to decrease.

Refraction: The blue, purely phase-shifting wedge leads to a local refraction of the incoming wave, which manifests as a lateral shift of the interference pattern.

Scattering: The green sample scatters incoming X-rays on a very small scale, leading to a local loss of coherence and thus reduces the amplitude of the interference pattern.

Figure 2.4 shows exemplary stepping curves for these three effects. Plotted is the recorded intensity versus the grating position, which extends over one complete period of the analyzer grating. The black curve in each plot represents a scan without object in the beam—called flat-field or blank scan—and thus a sampling of the unmodified interference pattern. The colored curves show from left to right the results of scans with an attenuating, refracting and scattering object, respectively. The attenuation causes a change in mean value of the curve, refraction a shift of the phase and scattering a decrease of the amplitude as indicated by the dashed lines in the plots. These curves serve as a basis for extracting the three contrast signals available from grating interferometry. The signal extraction will be described in the next section.

The above description of the grating interferometer assumes a source, which itself provides enough coherence to allow for interference effects. This is mostly given when synchrotron sources are used. In 2006 it was shown that grating interferometry is also possible using conventional, incoherent X-ray tubes with large focal spots by adding a third, absorbing grating in front of the source. This source grating G_0 splits the large source into an array of small slit sources, each by itself providing the necessary coherence (Pfeiffer et al., 2006; Weitkamp et al., 2006). In this three-grating configuration the setup is then called Talbot-Lau interferometer.

For illustrative purposes, a 3D-rendering of a typical bench-top three-grating interferometer for computed tomography is shown in figure 2.5.

2.2.3 Signal extraction

To extract the three contrast images from the raw detector images captured during a phase-stepping scan, an additional processing step has to be per-




Figure 2.5: Setup used for doing tomographic imaging with a grating interferometer

formed. From the previous section we know that the intensity curves acquired during a phase-stepping scan are represented by periodic functions. Thus, we can write the recorded intensity as a Fourier series

$$I(x_g, x, z) = a_0(x, z) + \sum_{k=1}^{\infty} a_k(x, z) \cdot \cos(2\pi k x_g - \phi_k(x, z)), \quad (2.21)$$

where x_g is the grating position in units of the fraction of a full grating period and (x, z) are the coordinates on the detector. The propagation direction is again along the y -axis. It was also stated that usually the intensity curves are sine shaped, so the signal extraction can be done with only the 0th and 1st order terms of the series

$$I(x_g, x, z) = a_0(x, z) + a_1(x, z) \cos(2\pi x_g - \phi_1(x, z)) \quad (2.22)$$

The interesting parameters in this equation are the mean value a_0 , the phase ϕ_1 and the oscillation amplitude a_1 . The natural way to calculate the parameters of such a sinusoidal function is Fourier analysis. The average value and the amplitude are then given by the absolute value of the zeroth and first Fourier coefficients, respectively, and the phase by the argument of the first coefficient. These parameters can be related to the attenuation, refraction and scattering of a sample by comparing parameters retrieved from a

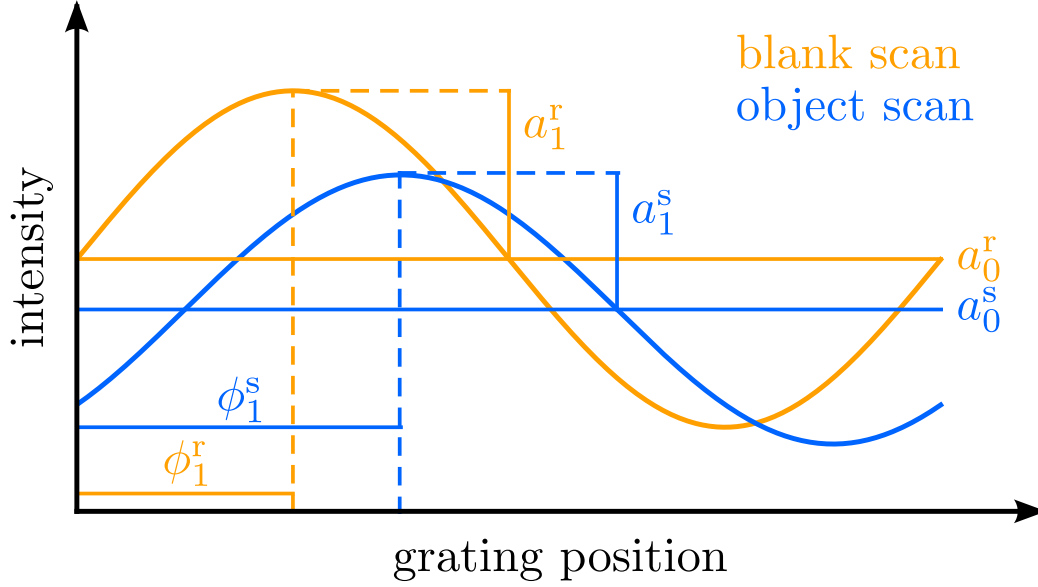


Figure 2.6: Illustration of the signal extraction from stepping curves with sample (object scan, blue) and without sample (blank scan, orange)

blank scan and a scan with sample. In the following the superscript 'r' will denote quantities from a blank—or reference—scan and 's' will be used for scans with sample. Stepping curves containing all three effects are shown in figure 2.6. The blank scan is represented by the orange curve, the scan with sample by the blue curve. Also marked are the three parameters from the Fourier analysis for both curves.

Attenuation: The transmission through a sample is given by the ratio of the mean intensities of signal and blank scans

$$T(x, z) = \frac{I(x, z)}{I_0(x, z)} = \frac{a_0^s(x, z)}{a_0^r(x, z)}. \quad (2.23)$$

Obviously the transmission is related to the attenuation of the sample, i.e. the linear attenuation coefficient μ , by the Beer-Lambert law (see equation (2.5)) and

$$a(x, z) = 1 - T(x, z) \quad (2.24)$$

Refraction: Refraction leads to a transverse shift of the interference pattern. This shift is proportional to the phase difference between the stepping curves with sample and the stepping curve of the blank scan. The exact relation is

$$S(x, z) = \frac{p}{2\pi} (\phi_1^s(x, z) - \phi_1^r(x, z)), \quad (2.25)$$

where p denotes the period of the grating that was moved during the phase-stepping scan. We call the result of the phase-stepping analysis $\phi = \phi_1^s - \phi_1^r$ the differential phase contrast (DPC) image. It is evident that the shift of the interference pattern cannot be determined without ambiguity. As the pattern is periodic, we can detect the phase of the stepping curve only in the interval $[0, 2\pi]$, values outside of this interval are simply wrapped back into it, thus this effect is called *phase wrapping*. We will see in chapter 5, that this ambiguity in determining the interference pattern shift—or the gradient of the phase—has dramatic effects in tomographic reconstruction, leading to the appearance of bone artifacts.

The transverse shift of the interference pattern is related to the refraction angle via

$$\alpha(x, z) = \frac{S(x, z)}{d} = \frac{p}{2\pi d} \phi(x, z), \quad (2.26)$$

with d denoting the distance between the phase grating and the analyzer grating. Together with the line integral equation for the refraction angle (see equation (2.12)) we can relate the phase difference from the stepping curves to the refractive index decrement δ

$$\phi(x, z) = \frac{2\pi d}{p} \frac{\partial}{\partial x} \int \delta(x, y, z) dy. \quad (2.27)$$

Scattering: The last signal we can extract is related to (ultra) small angle scattering with scattering angles much smaller than the resolution of the interferometer. This kind of scattering leads to a loss of the beam coherence, which then causes the amplitude of the interference pattern to decrease. We call the quantity that describes this coherence loss the visibility. It is defined as the ratio $(I_{\max} - I_{\min}) / (I_{\max} + I_{\min})$ or in terms of the Fourier coefficients

$$V(x, z) = \frac{a_1(x, z)}{a_0(x, z)}. \quad (2.28)$$

The final image signal incorporating both intensity curves is then given as the ratio of the visibility for the phase-stepping scan with sample and the visibility of the blank scan

$$v(x, z) = \frac{V^s(x, z)}{V^r(x, z)} = \frac{a_1^s(x, z) a_0^r(x, z)}{a_1^r(x, z) a_0^s(x, z)}. \quad (2.29)$$

This signal is called the dark-field image (Pfeiffer et al., 2008).




Figure 2.7: Illustration of the Fourier slice theorem.

2.3 Tomographic reconstruction

In this section the most widely used method for the reconstruction of tomographic measurements, the filtered back projection, is described. The reconstruction is the computational part of computed tomography and its purpose is to calculate the three-dimensional distribution of a quantity from two-dimensional measurements. For example, in conventional X-ray CT a large number of radiographic images (projections) of transmitted intensities are recorded from different angular views around an object and reconstruction algorithms are used to retrieve the three-dimensional map of the linear attenuation coefficient. In the following we first describe the Fourier slice theorem, which the FBP algorithm is based on and afterwards the FBP itself.

2.3.1 Fourier slice theorem

The Fourier slice theorem states that the Fourier transform of parallel projection of a two-dimensional function is equal to a slice through the two-

dimensional Fourier transform of the original function. If the projection is taken at an angle θ , its Fourier transform will correspond to the slice through the two-dimensional Fourier transform of the object along the line rotated by θ with respect to the q_x axis. The theorem is illustrated in figure 2.7.

We define the projection of a general two-dimensional distribution $f(x, y)$ in a coordinate system (x', y') rotated by an angle θ with respect to the original coordinates, i.e.

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}. \quad (2.30)$$

The projection along the propagation direction of the X-rays in the rotated coordinates is then

$$\mathcal{P}_\theta(x') = \int f(x', y') dy', \quad (2.31)$$

where the propagation direction is chosen as the y' axis. This projection is also called the Radon transform of $f(x, y)$ (Radon, 1917). We restrict this derivation to a parallel-beam geometry such that single slices at constant z positions can be treated independently. Stacking the projections for all measured angles creates a so-called sinogram, because a point in the object will follow a sine-shaped curve when tracked through all angular projections. Each point in the sinogram corresponds to a value of $\mathcal{P}_\theta(x')$.

The Fourier transform of the one-dimensional projection is

$$\tilde{\mathcal{P}}_\theta(q'_x) = \int_{-\infty}^{\infty} \mathcal{P}_\theta(x') e^{-2\pi i x' q'_x} dx' \quad (2.32)$$

The two-dimensional Fourier transform in the rotated coordinate system is

$$\tilde{f}(q'_x, q'_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x', y') e^{-i2\pi(x'q'_x + y'q'_y)} dx' dy'. \quad (2.33)$$

A slice through the origin of the frequency space defined by (q'_x, q'_y) tilted by the angle θ requires that $q'_y = 0$, which simplifies equation (2.33) to

$$\tilde{f}(q'_x, 0) = \int \int f(x', y') e^{-i2\pi x' q'_x} dx' dy' \quad (2.34)$$

$$= \int \left(\int f(x', y') dy' \right) e^{-i2\pi x' q'_x} dx' \quad (2.35)$$

$$= \int [\mathcal{P}_\theta(x')] e^{-i2\pi x' q'_x} dx' = \tilde{\mathcal{P}}_\theta(q'_x), \quad (2.36)$$




Figure 2.8: Illustration of the sampling of the Fourier space when recording projection images from different angles around an object.

thus proving the above stated Fourier slice theorem. If the values of all points in the (q_x, q_y) -plane are determined by measuring an infinite number of projections, the original object function can then be recovered by an inverse Fourier transform

$$f(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{f}(q_x, q_y) e^{i2\pi(xq_x + yq_y)} dq_x dq_y \quad (2.37)$$

2.3.2 Filtered back projection

The filtered backprojection (FBP) algorithm is currently the most widely used method for tomographic image reconstruction from projections because of its high speed and easy implementation. It can be directly derived from the Fourier slice theorem, which says that the frequency space can be sampled by measuring projection images from many directions around the object. The conclusion of the previous section was that a two-dimensional Fourier transform is sufficient to recover the original object function, if it were possible to measure an infinite number of projection images. However, in real experiments only a finite number of projections with a finite number of data points can be recorded. The frequency space is accordingly only sampled by a finite set of values along the radial lines corresponding to the angles θ used to capture the projections. The sampling of the Fourier space in this case is illustrated in figure 2.8. The orange points correspond to the measurements and it is obvious that for low frequencies the sampling points are much closer

than for high frequencies. The FBP compensates this by introducing a filtering step that gives more weight to the high frequency components and then smearing the filtered projections back over the field of view.

The derivation of the filtered back projection algorithm is based on expressing the inverse two-dimensional Fourier transform in polar coordinates with $(q_x, q_y) = (\omega \cos \theta, \omega \sin \theta)$ and $dq_x dq_y = \omega d\omega d\theta$

$$f(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{f}(q_x, q_y) e^{i2\pi(xq_x + yq_y)} dq_x dq_y \quad (2.38)$$

$$= \int_0^{\pi} \int_{-\infty}^{\infty} \tilde{f}(\omega, \theta) e^{i2\pi x' \omega} |\omega| d\omega d\theta, \quad (2.39)$$

with $x' = x \cos \theta + y \sin \theta$ as in the previous section. The last integral can be split into the integration over ω and over θ

$$\mathcal{H}(x' = x \cos \theta + y \sin \theta, \theta) = \int_{-\infty}^{\infty} \underbrace{\tilde{f}(\omega, \theta)}_{\tilde{P}_\theta(q'_x)} |\omega| e^{i2\pi x' \omega} d\omega, \quad (2.40)$$

where in the underbraced term ω and q'_x represent the same coordinates for a given angle θ , when we set $q'_y = 0$ as we did in equation (2.34). This first integration over ω constitutes a filtering operation with a ramp filter $\tilde{h}(\omega) = |\omega|$ defined in Fourier space. The second integration over θ

$$f(x, y) = \int_0^{2\pi} \mathcal{H}(x \cos \theta + y \sin \theta, \theta) d\theta \quad (2.41)$$

corresponds to a back projection of the filtered projection data, i.e. each filtered projection is smeared over the field of view. Due to the change of coordinates from a polar to a Cartesian system, i.e. the projections sample the Fourier space on a polar grid, whereas the reconstructed object is defined on a Cartesian grid, some kind of interpolation is required in the back projection step. Usually a simple linear interpolation is used. The polar sampling of frequency space also defines the lower bound on the number of projections that have to be measured for a reconstruction free of aliasing artifacts. This lower bound is given in Kak and Slaney (1988) and Mueller (1998) as

$$M_{\text{proj}} = \frac{\pi}{2} N_{\text{det}}, \quad (2.42)$$

where M_{proj} is the number of projections taken and N_{det} is the number of pixels in a single row of the detector.

Attenuation Conventional attenuation CT measures the relative transmission in each projection according to the Beer-Lambert law (see equation (2.5)). The object function to be reconstructed is the spatial distribution of the linear attenuation coefficient. We can define the projections $\mathcal{P}_\theta(x')$ as the logarithm of the relative transmission

$$\mathcal{P}_\theta^T(x') = -\ln\left(\frac{I_\theta(x')}{I_0}\right) = \int_{-\infty}^{\infty} \mu(x', y') dy'. \quad (2.43)$$

Applying the FBP algorithm on these projections yields $\mu(x, y)$ as desired.

$$\mathcal{H}(x', \theta) = \int_{-\infty}^{\infty} \tilde{\mathcal{P}}_\theta(x') |\omega| d\omega \quad (2.44)$$

$$\mu(x, y) = \int_0^\pi \mathcal{H}(x' = x \cos \theta + y \sin \theta, \theta) d\theta \quad (2.45)$$

Differential phase contrast In the case of the refraction data the projections are of a differential nature. The measurement—the phase difference of stepping curves from an object scan and a blank scan—is related to the spatial distribution of the decrement of the refractive index $\delta(x, y)$ via equation (2.27). The projections are then defined as

$$\mathcal{P}_\theta^\phi(x') = \phi_\theta(x') = -\frac{2\pi d}{p} \frac{\partial}{\partial x'} \int \delta(x', y') dy'. \quad (2.46)$$

Because of the partial derivative we cannot use these projections directly as input for the FBP but we can make use of the Fourier transform of the derivative of a function

$$\mathcal{FT} \left[\frac{\partial}{\partial x} f(x) \right] = 2\pi i \omega \tilde{f}(\omega), \quad (2.47)$$

so the ramp filter $|\omega|$ in the equation for the filtered projections (see equation (2.40)) has to be augmented to compensate for the additional factor from the derivative and yields

$$\mathcal{H}(x', \theta) = \int_{-\infty}^{\infty} -2\pi i \omega \tilde{\delta}(\omega, \theta) \frac{|\omega|}{2\pi i \omega} e^{i2\pi x' \omega} d\omega, \quad (2.48)$$

with the modified filter $\tilde{h}(\omega) = |\omega|/2\pi i\omega = -2\pi i \text{sgn}(\omega)$, which basically represents an integration of the differential projections in Fourier space. With this filter the differential projections can be directly used as input for the FBP reconstruction.

$$\delta(x, y) = \int_0^\pi -\mathcal{H}(x' = x \cos \theta + y \sin \theta, \theta) d\theta \quad (2.49)$$

Dark field The reconstruction procedure for the dark-field signal is analogous to the attenuation signal. Following Bech (2009) the dark-field signal can be explained by a linear diffusion coefficient $\epsilon(x, y)$, which is related to the measured quantity—the visibility—in the same way as the linear attenuation coefficient is related to the relative transmission. The projection function is then defined via the logarithm of the measured visibility $v(x, y)$

$$\mathcal{P}_\theta^v(x') = -\ln \frac{p^2}{2\pi^2 d^2} v_\theta(x') = \int_{-\infty}^{\infty} \epsilon(x', y') dy', \quad (2.50)$$

where the relation between the linear diffusion coefficient and the measured visibility is (Bech, 2009)

$$v(x) = \frac{V^s}{V^r} = e\left(-\frac{2\pi^2 d^2}{p^2} \int \epsilon(x, y) dy\right). \quad (2.51)$$

The measured projections can then be directly used in the FBP.

2.4 Iterative reconstruction methods

The FBP as a transform-based method relies on evenly distributed data sampling, i.e. projections measured uniformly over 180° or 360° . Also a sufficient number of projection views have to be recorded to achieve the desired spatial resolution for the specific application (see Kak and Slaney (1988) and Mueller (1998)). If these requirements are not met by, for example, non-uniform acquisition of projection data or not recording projections over the half or full circle, different reconstruction techniques have to be used. Basically, there are two main classes of algorithms, algebraic methods and model-based reconstruction.

The first group of methods relies on the algebraic solution of an inverse problem of the form

$$\mathbf{A}\mathbf{f} = \mathbf{p}, \quad (2.52)$$

where \mathbf{f} is the objective function in vector notation, \mathbf{p} the vector containing the measured projection data for all detector pixels and angles and the system matrix \mathbf{A} represents the so-called forward model. The forward model can, for example, be just the discretized Radon transform but it can also contain a complete physical model of the acquisition system, even for arbitrary scanning geometries. The system matrix \mathbf{A} is often defined to be an $N \times M$ matrix, where N is the number of discretized points—usually voxels—in the reconstructed object and M is the number of rays for all projections. For typical CT problems M and N are very large, so a direct inversion of the system matrix is usually not feasible. Numerical solutions for estimating the original object have to be used.

The first of these methods was proposed by Gordon et al. (1970) and is called **algebraic reconstruction technique (ART)**. In ART the elements of the system matrix are defined as

$$a_{ij} = \begin{cases} 1 & \text{if ray } i \text{ intersects voxel } j \\ 0 & \text{otherwise} \end{cases}, \quad (2.53)$$

i.e. a voxel contributes its value to a specific projection if it is intersected by a ray, independent of the intersection length. Equivalently the forward model reads

$$p_i = \sum_{j=0}^{N-1} a_{ij} f_j, \quad i = 1, 2, \dots, M, \quad (2.54)$$

i.e. it represents a linear system of M equations for N unknown values. The solution is based on the Kaczmarz method of projections (Kaczmarz, 1937), where the current estimate at step $(i - 1)$ is projected onto a hyperplane represented by the i^{th} equation. The update equation for the j^{th} voxel for such a subiteration reads

$$f_j^{(i)} = f_j^{(i-1)} + \frac{p_i - \sum_{k=0}^{N-1} f_k^{(i-1)} \cdot a_{ik}}{\sum_{k=0}^{N-1} a_{ik}^2} \cdot a_{ij}. \quad (2.55)$$

One iteration is complete as soon as every equation has been used for an update. One drawback of this method is that voxels in the reconstructed image are updated after each step, so it might occur that voxel values are

changed after considering one equation, i.e. have been set to fit to one particular measurement, and the same voxels are modified again to fulfill a subsequent equation. These competing voxel updates are avoided in a different implementation called **simultaneous iterative reconstruction technique (SIRT)** (Gilbert, 1972). In this approach the voxel updates are first calculated for all equations and then the average of the calculated updates for a particular voxel are written back, which leads to smoother reconstructions, albeit at the cost of slower convergence.

A third flavor of the algebraic methods seeks to combine the fast convergence of ART and the smooth reconstructions of SIRT. It is called **simultaneous algebraic reconstruction technique (SART)** and was first introduced by Andersen and Kak (1984). The voxel updates in this approach are also applied simultaneously as in SIRT, but not for all rays—one equation in the system represents the sum along a ray through the object—at once. The correction terms are calculated for all the rays in a particular angular view, then all voxels are updated using these terms. Again, a single full iteration is complete, when all ray-sums, i.e. equations, have been considered. In addition, angular projections are not traversed in sequential order, but rather angles far apart from each other are considered to avoid redundant updates caused by the high correlation between adjacent angular views.

A second improvement over ART and SIRT is the refinement of the forward projection model. Instead of using pixels as a basis for discretizing the object, bi-linear elements, pyramids-shaped objects extending over the area of four traditional pixels are introduced. In addition, the idea of binary elements for the system matrix, i.e. all voxels that are intersected by a ray are added to the particular equation of this ray, is dropped in favor of ray strips of finite width. The intersection area between the image sampling points and the rays is then calculated using bi-linear interpolation.

Lastly, the correction terms are no longer weighted uniformly (by the value of a_{ij} in equation (2.55)) for the back-distribution to the image, but a longitudinal hamming window is applied that gives more weight to points along a ray that are closer to the center of the object (Andersen and Kak, 1984).

Even more refinement of the forward projection model is found in statistical reconstruction methods. These most advanced iterative techniques use statistical models of the complete image acquisition process including noise properties to get an objective function measuring the deviation of the measured data from the model (e.g. minimum least-squares error, maximum likelihood, etc.). The objective function is then optimized to calculate an estimate of the original distribution of the desired quantity. A variety of

optimization algorithms have been developed and are used today. Examples include steepest descent, conjugate gradients—used within this thesis—, expectation maximization and many more. An exemplary derivation of such a SIR algorithm is given in detail in chapter 4 for a Gaussian distribution. In conventional CT most commonly the projection data is assumed to be following a Poisson distribution. The derivation of Poisson likelihood algorithms is not given here. The reader is referred to Fessler (2000) for an in-depth treatment of this topic.

Chapter 3

Data processing and analysis of projection data

This chapter is separated into three main sections, which all focus on methods or improvements that work on radiographic, i.e. projection, data. The first section focuses on a technique to compare the contrast-to-noise ratio of absorption and DPC radiographs without the need for an explicit integration of the differential data. This comparison is accomplished by instead calculating the gradient of the absorption projection and doing the analysis on the derivative projections. The results of this study have been published in D. Hahn et al., Numerical comparison of X-ray differential phase contrast and attenuation contrast, Biomedical Optics Express (2012). The next section, 3.2, introduces an alternative signal extraction, providing additional statistical information for the later development of statistical reconstruction techniques. It is based on weighted least-squares principles and supposed to replace the conventional Fourier analysis approach. In the final section, 3.3, we study how the system's point spread function affects raw projections, the three contrast channels and the DPC reconstruction and show a way to improve overall image sharpness with simple deconvolution tools.

3.1 Comparison of attenuation and phase signals in single projections

Introduction

In this section a method for comparing the performance of attenuation and differential phase contrast in single projections is presented. At a first glance this is not an easy task because of the differences in the image formation process. While an attenuation-contrast image can be interpreted as a simple projection of an object function (the distribution of linear absorption coefficient), in grating interferometry the differential phase-contrast image corresponds to the first derivative of the projected object function (the decrement of the complex refractive index).

One way to directly compare both contrast signals is, of course, quantitative PCCT, which is described in detail in chapter 2 and in (Herzen et al., 2009; Chen et al., 2011; Zambelli et al., 2010). In PCCT the differential phase-contrast images are inherently integrated by the reconstruction algorithm, resulting in an image proportional to the decrement of the refractive index. But for radiographic applications such as mammography, where the interest in increased soft tissue contrast and the requirement for low radiation dose is especially high, usually only a single projection is recorded. This gives rise to the question of how to assess and compare image quality and information content in the domain of single projections or in other words, how to compare a non-differential to a differential signal.

In this section we introduce a quantity, the relative contrast gain (RCG), which provides an estimate of the expected performance of a differential phase-contrast image compared to an attenuation projection. It is defined as the ratio of the contrast-to-noise ratio (CNR) in a DPC projection to the CNR in an attenuation projection and it is meant as a analysis tool that can be applied on real data. Engel et al. (2010) recently derived a theory of the CNR in differential phase-contrast imaging based on an integration of the phase (see also Kottler et al. (2007); Zanette et al. (2010); Thuring et al. (2011)). In contrast to the integration approach, we propose to solve the problem of comparing the two image modalities by differentiating the attenuation image along the same direction as the derivative of the phase.

To motivate the use of a differential attenuation signal for the analysis of the RCG, we have to go back to section 2.1, where the line integrals of the

attenuation and phase signals were introduced as

$$T(x, y) = \ln \left(\frac{I(x, y)}{I_0(x, y)} \right), \quad (3.1)$$

$$\partial_x \Phi(x, y) = \frac{2\pi}{\lambda} \alpha(x, y), \quad (3.2)$$

Note that from a single projection we can only obtain these line integrals of the two quantities through the sample.

In any experiment, noisy versions of T and $\partial_x \Phi$ are measured, and our aim is to compare the quality of these two signals. To obtain a physically meaningful comparison we can either differentiate the attenuation integral (eq. (3.1)) or integrate the differential phase (eq. (3.2)). The integration of the differential phase takes a lot of experimental effort to accomplish by using additional measurements (Kottler et al., 2007; Zanette et al., 2010). Numerical integration is equally difficult due to the missing constant of integration and an extremely low signal-to-noise ratio of the non-zero low frequency components, which can lead to strong artifacts (Thuring et al., 2011).

Equations (3.1) and (3.2) indicate that the only other possibility for comparison is to differentiate the attenuation signal along the same direction x as the derivative in eq. (3.2). As a result we get the following definition for the differential transmission:

$$\partial_x T = \frac{\partial}{\partial x} \ln \left(\frac{I}{I_0} \right). \quad (3.3)$$

The Relative Contrast Gain (RCG)

The main concept behind the definition of the Relative Contrast Gain is a quantification of the increase in feature contrast of differential phase-contrast projections compared to attenuation projections. For the analysis the following definition of a contrast-to-noise ratio (CNR) is used:

$$\text{CNR} = \frac{\max(A) - \min(A)}{\sigma_0}, \quad (3.4)$$

where A denotes a set of image values in a region of interest (ROI) in a projection and σ_0 is the standard deviation of image values in a flat region, i.e. the pure image noise. This definition can be seen as an effective dynamic range or in other words the number of grey levels needed to represent the

image without loss of information. For a comparison of differential phase-contrast projections to attenuation projections we then define the RCG as the ratio of the CNR in a ROI of the differential phase projection to the CNR in the same ROI of the differential attenuation projection,

$$\text{RCG} = \frac{\text{CNR}_\Phi}{\text{CNR}_T} = \frac{\Delta(\partial_x \Phi) / \sigma_{\partial_x \Phi}}{\Delta(\partial_x T) / \sigma_{\partial_x T}}, \quad (3.5)$$

where $\Delta(X)$ corresponds to $\max(X) - \min(X)$ and σ_X denotes the standard deviation.

The RCG is an estimate of the relative contrast of phase projections and attenuation projections. By normalizing with the standard deviation of the respective images, the pure image signals, that are not hidden by noise, are compared. In imaging terms we could also reformulate that by saying that the RCG is a quantity that determines which signal provides a larger dynamic range in the presence of noise. It is similar to the detective quantum efficiency (DQE), which quantifies the ability of an imaging system or detector to provide good signal-to-noise.

Differentiation of the attenuation image

The problem of differentiating a discrete image affected by noise presents a challenge that is difficult to solve. This is due to a low signal-to-noise ratio in high frequency components found in most images, where the derivative will lead to an amplification of the noise levels.

It is well known that the derivative corresponds to a linear function in Fourier space. In analogy to the Ram-Lak filter known from the reconstruction in attenuation computed tomography (Hsieh, 2009; Oppenheim and Schaffer, 1999) we have to bandlimit the filter response by applying a window function. Typical choices in CT reconstructions are the Dirichlet-window (rectangle-window), *sinc*-window, Hamming-window or the Hann-window. (Hsieh, 2009; Oppenheim and Schaffer, 1999) As images are functions that are sampled on a finite grid and spatially limited, their Fourier transform will be inherently bandlimited with a cutoff frequency of $1/2 \text{ px}^{-1}$. To prevent aliasing effects that may be introduced by applying a discrete Fourier transform, the filter also has to be bandlimited and, as in the case of CT reconstruction filters, there are certain degrees of freedom in choosing the shape of a bandlimiting window function. The two main concerns in choosing the appropriate window function are the accuracy of the resulting derivative and the best possible noise suppression without losing image information, i.e. spatial resolution.




Figure 3.1: Experimental data used in the RCG analysis. (A) photograph of the sample, (B) attenuation signal, (C) differential phase-contrast signal, (D) differential attenuation signal.

For the present application of differentiating the attenuation image we compared the Dirichlet-window to the Hamming- and a Gaussian-window. We chose the Gaussian-derivative filter (Gonzalez and Woods, 2008) because it is the most versatile one. In the following we will set the width $\sigma = 1/2\pi \text{ px}^{-1}$. This corresponds to a derivative filter combined with a smoothing operation

over one pixel in image space. Smoothing in y -direction by the same amount improves the noise suppression even further. Smoothing by this small amount is just enough to keep the noise at an acceptable level. Using a smaller σ , i.e. smoothing over a larger area of pixels in image space, leads to a visual degradation of the image and a loss of information.

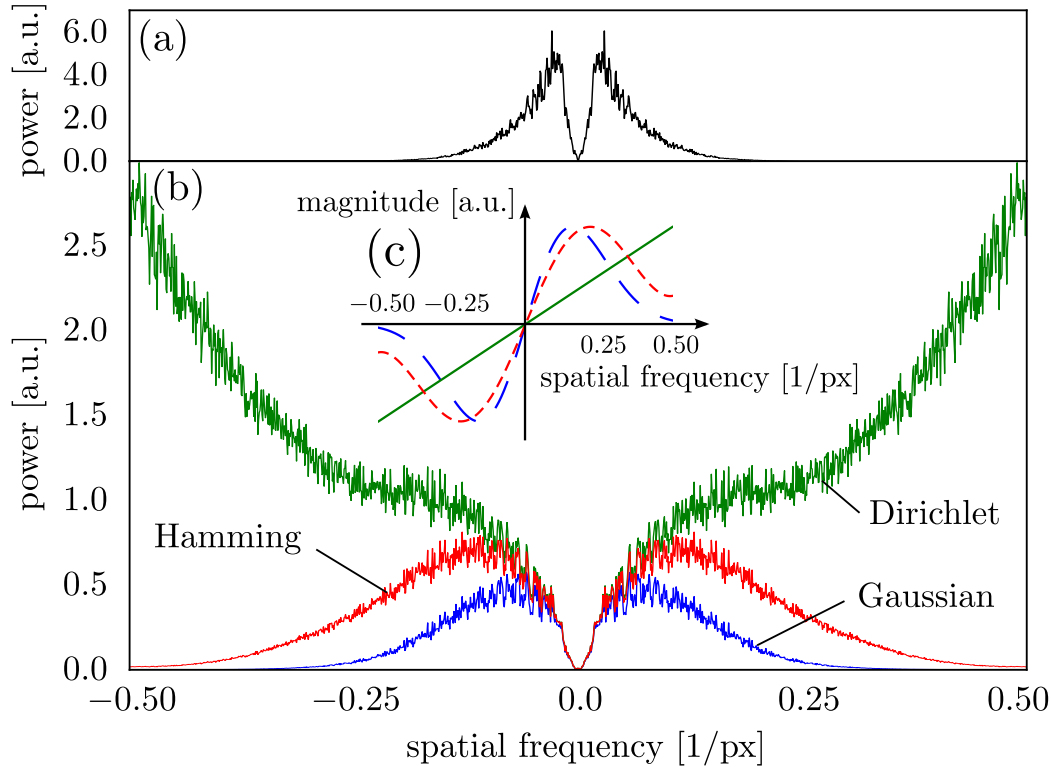


Figure 3.2: Power spectra of the experimental data shown in fig. 3.1, calculated by computing the absolute squared Fourier transform in x -direction for each image row and averaging in y -direction. (a) power spectrum of the DPC projection; (b) power spectra of the differential attenuation projection calculated with different filter functions as indicated in the figure; (c) corresponding filter functions; solid green: Dirichlet windowed; short-dashed red: Hamming windowed; long-dashed blue: Gaussian windowed with $\sigma = 1/2\pi \text{ px}^{-1}$.

Results

To validate the concept, we have applied the RCG analysis on experimental data and we present first results in the following. For that purpose, radiographs of a human breast sample were recorded using a two-grating Talbot interferometer (Weitkamp et al., 2005; Pfeiffer et al., 2007a; Schulz

et al., 2010; Weitkamp et al., 2010) at the ID19 beamline of the European Synchrotron Radiation Facility in Grenoble, France. The radiographs were taken with a monochromatic X-ray beam with an energy of 23 keV in the 9th fractional Talbot order. The distance between the phase grating G_1 and the analyzer grating G_2 was $d = 0.48$ m, G_1 had a period of $g_1 = 4.785$ μm and G_2 a period of $g_2 = 2.4$ μm . Phase stepping was performed in four steps over one period. The resulting raw images are processed with the commonly used Fourier phase-stepping analysis (Weitkamp et al., 2005) to extract the relative transmission signal T and the differential phase signal α . (Bech, 2009) The sample itself was a slice of human breast tissue about 1 cm thick, fixated in formalin. Figure 3.1 gives an overview of the experimental data used in the analysis and shows a photograph of the breast slice, an attenuation, differential phase-contrast and differential attenuation projection, respectively.

For the RCG analysis the attenuation signal was brought into the form introduced in the first section, more specifically eq. (3.3), to be able to compare the two signals. More precisely, we took the logarithm of the attenuation signal and differentiated it using the Gaussian-derivative filter defined in the previous section. We used eq. (3.2) as it is, as the differential phase signal is directly available from the phase-stepping analysis. To ensure a fair comparison, a Gaussian filter with the same width was also applied to the DPC signal. This introduces the same loss of high frequency information that occurs during the differentiation of the attenuation signal. The noise standard deviation of both signals was calculated from corresponding reference projections, i.e. from a blank scan without the sample in the beam.

Figure 3.2 shows power spectra for the experimental data shown in fig. 3.1 to assess the behaviour of the different filter functions when applied to the attenuation projection and to compare them to the DPC projection. The power spectrum of the differential phase contrast projection is plotted in panel a). Panel b) shows the power spectra of the attenuation projection differentiated using the three indicated window functions. Finally in panel c) the filter functions themselves are plotted. These plots show that the Gaussian-derivative filter with $\sigma = 1/2\pi$ px^{-1} performs reasonably well in making the power spectrum of the attenuation projection similar to that of the differential phase contrast projection.

A visual representation of the RCG is given in fig. 3.3 for the regions A (green box) and B (red box), which are defined in fig. 3.1 (A), respectively. In the main part of each subfigure the values of $\partial_x T$ and $\partial_x \Phi$ for each position (x, y) in the projection are presented in the form of a scatterplot for regions A (left) and B (right), respectively. The horizontal axis corresponds to the differential

phase-contrast projection, the vertical axis to the differential attenuation projection. Histograms showing the distribution of values for both signals for the sample and reference (i.e. blank scan) regions are plotted on top of the respective axes of the signals. Green color denotes the sample region and blue the corresponding reference regions. The definition of the RCG in eq. (3.5) can equivalently be stated in terms of the geometry of the ellipses formed by the pixel values of the sample region and the widths of the reference region histograms. The width of the ellipse corresponds to $\Delta\partial_x\Phi$ and the height corresponds to $\Delta\partial_x T$. This means that the RCG is inversely proportional to the slope of the ellipses major axis. Using eq. 3.5 we find for region A an $RCG_A = 9.5$ and for region B an $RCG_B = 5.8$.




Figure 3.3: Visual representation of the RCG analysis of regions A (left) and B (right).

Discussion and conclusion

The results of the RCG analysis from the last section indicate that phase-contrast projections, as expected, can provide substantially higher soft-tissue contrast than attenuation images. Two distinct regions of interest were analyzed, one with clearly visible features and one with features that are barely visible by the human observer. Looking at the scatterplots and histograms in fig. 3.3 and comparing the width and height of the ellipses formed by the sample (green) and reference (blue) pixel values and the widths of the corresponding projections in the 1-D histograms it is obvious that in both regions the fraction of the signals dynamic range covered by noise is larger for the differential attenuation. This observation is quantified by RCG values

of 9.5 and 5.8 for regions (A, green) and (B, red) respectively, confirming the improved CNR performance of differential phase projections. The RCG is different for these two regions, which indicates that the RCG analysis depends strongly on the properties of the analyzed region. We believe that this is mainly due to the different spatial frequencies contained in these two regions and the different ratios of δ/β for different materials. Therefore the RCG is a local measure that states for which regions the phase signal should be trusted more than attenuation. Because of the uncertainty in choosing the width of the Gaussian-derivative filter that is used to calculate $\partial_x T$, the RCG cannot be defined as an absolute scale, but rather a relative measure to compare different regions inside an image. Further work has to be done on calibrating the derivative filter to be able to absolutely quantify the contrast improvement of differential phase contrast compared to attenuation contrast, a task especially difficult because of the different noise power spectra in the attenuation and phase components.

The results we obtained for the RCG can, in principle, be further improved by optimizing the experimental setup. For the experimental data of the human breast tissue the visibility of the fringe oscillations during the phase-stepping was 54%. As the noise in a phase-contrast projection is inversely proportional to this visibility we can reach higher contrast-to-noise in the phase projection by increasing the visibility. Note that a change in visibility will not affect the CNR of the attenuation projection which will result in a higher relative contrast gain in favor of phase contrast. There are several ways to improve an experimental setup in terms of visibility. The manufacturing of phase and absorption gratings is constantly improving, leading to fewer defects in the gratings and to higher aspect ratios allowing X-rays with higher energies to be used. Also having the gratings matched to the desired X-ray energy and adjusting the Talbot distance between the phase grating and the absorption grating accordingly may improve the visibility.

Our results are consistent with those presented recently by (Engel et al., 2010) and we conclude that the RCG formalism can be used as a tool to assess the image quality in differential phase-contrast imaging. Also, the discussions in this paper and in (Engel et al., 2010) indicate that the details of an experimental setup, e.g. quality of the gratings, influence only the CNR of the phase projection but not the attenuation image, so from this fact we conclude that the RCG can be used as a tool to compare the performance of different experimental setups.

3.2 Statistical phase retrieval

This section describes an alternative procedure to retrieve the three contrast signals from the data measured with a grating interferometer. It uses the statistical properties of the raw projections to estimate not only the values of the three signals but also their uncertainties and is thus called statistical phase retrieval (SPR). As explained in section 2.2.2, the acquisition of projection data in grating interferometry is done by stepping one of the gratings and recording an image of the intensity distribution at each grating position and by that an intensity curve for each individual pixel. The common way to extract the three contrast signals—attenuation, differential phase and darkfield—from the stepping curves is the Fourier transform approach presented in detail in section 2.2.3. This method, however, has one drawback: the statistical uncertainties of the retrieved signals are not easily calculated.

Here we present an alternative implementation of the phase retrieval to accommodate the need for the projection's statistical information, that is used in the reconstruction algorithm throughout the next two chapters. A weighted least squares (WLS) estimation is used to directly fit a sinusoidal curve to the measured intensities for each pixel. If we recap the intensity curve model function used in the Fourier approach

$$I(x) = a_0 + a_1 \cdot \cos(x - \phi_1), \quad (3.6)$$

we can see, that estimating the three parameters a_0 , a_1 and ϕ_1 —corresponding to mean, amplitude and phase of the sinusoid—from this model function is not straightforward with a conventional least-squares (LS) approach, as this requires the function to be linear with respect to all fit parameters to work properly. Consequently the first step is to linearize the function, which is simple using basic trigonometric relations. The linearized model function of the stepping curve is then:

$$I(x) = A_0 + A_1 \cdot \cos(x) + B_1 \cdot \sin(x), \quad (3.7)$$

where the shifted cosine from the original function is replaced by a linear combination of a cosine and a sine. The fit parameters A_0 , A_1 and B_1 now denote the mean of the curve, the amplitude at $x = 0$ and at $x = 0.25 \cdot p$, respectively, where p is the period of the grating that was moved. The three original parameters can be recovered by

$$a_0 = A_0, \quad \phi_1 = \arctan\left(\frac{B_1}{A_1}\right), \quad a_1 = \frac{\sqrt{A_1^2 + B_1^2}}{A_0}. \quad (3.8)$$




Figure 3.4: Sample stepping curve. Indicated are the three fit parameters calculated by SPR as dashed lines and their uncertainties as the semi-transparent background. Additionally shown is their relation to the three contrast signals.

We can now introduce two vectors, $\mathbf{a} = (A_0, A_1, B_1)^\top$, containing the three estimated parameters, and $\mathbf{X}(x) = (1, \cos(x), \sin(x))^\top$, the vector of basis functions. Equation (3.7) can then be rewritten in general form as

$$I(x) = \sum_{j=0}^{M-1} a_j \cdot X_j(x), \quad (3.9)$$

where M is the total number of estimated parameters and basis functions. From this model a WLS cost function is derived that has to be solved for each pixel in each projection

$$S = \sum_{i=0}^{N-1} \frac{1}{\sigma_i^2} \left[I(x_i) - \sum_{j=0}^{M-1} a_j \cdot X_j(x_i) \right]^2, \quad (3.10)$$

where N denotes the total number of grating positions during a stepping scan and σ_i is the uncertainty of the intensity measurement at grating position x_i . Since the main source of noise is counting statistics, the natural model is the Poisson distribution, so the error on the measured intensity in a pixel

is estimated as $\sigma_I = \alpha\sqrt{I}$, with $\alpha = 1$ when the detector calibration is such that I represents photon counts.

To show that this assumption is not only true for photon counting detectors that measure the number of photons collected by a pixel directly, but that it can also be used in most cases for integrating detectors, e.g. CCDs, consider the following: The conversion from photon counts N to intensity I is linear in regimes far enough from saturation of the analog to digital converter of the detector. If readout noise can be neglected, this linearity also applies to the uncertainties σ_I and σ_N , as shown here:

$$\begin{aligned} I &= \alpha \cdot N_{ph} \\ \rightarrow \sigma_I^2 &= \sigma_N^2 \cdot \left(\frac{\partial I}{\partial N_{ph}} \right)^2 = \alpha^2 \cdot \sigma_N^2. \end{aligned} \quad (3.11)$$

To generally solve such a WLS problem, we have to compute the gradient of equation (3.10) with respect to each of the estimation parameters and set it to zero, leading to M equations that are indexed by k

$$0 = \sum_{i=0}^{N-1} \frac{1}{\sigma_i^2} \left[I(x_i) - \sum_{j=0}^{M-1} a_j \cdot X_j(x_i) \right] X_k(x_i) \quad k = 0, \dots, M-1 \quad (3.12)$$

In these equations the order of the summations can be reversed, which gives

$$\sum_{j=0}^{M-1} a_j \sum_{i=0}^{N-1} \frac{X_j(x_i) \cdot X_k(x_i)}{\sigma_i^2} = \sum_{i=0}^{N-1} \frac{I(x_i) \cdot X_k(x_i)}{\sigma_i^2}, \quad (3.13)$$

and with the definitions

$$B_{ij} = \frac{X_j(x_i)}{\sigma_i} \quad \text{and} \quad c_i = \frac{I(x_i)}{\sigma_i} \quad (3.14)$$

equation (3.13) can be rewritten in matrix notation as

$$(\mathbf{B}^T \cdot \mathbf{B}) \cdot \mathbf{a} = \mathbf{B}^T \cdot \mathbf{c} \quad (3.15)$$

or

$$\boldsymbol{\beta} \cdot \mathbf{a} = \boldsymbol{\gamma} \quad (3.16)$$

with $\boldsymbol{\beta} = (\mathbf{B}^T \cdot \mathbf{B})$ (system matrix) and $\boldsymbol{\gamma} = \mathbf{B}^T \cdot \mathbf{c}$ (solution vector).

The analytical exact solution is then given by

$$\mathbf{a} = \boldsymbol{\beta}^{-1} \cdot \boldsymbol{\gamma}, \quad (3.17)$$

where

$$\boldsymbol{\beta}^{-1} = \frac{1}{|\boldsymbol{\beta}|} \cdot \begin{pmatrix} (\beta_{11}\beta_{22} - \beta_{12}\beta_{21}) & (\beta_{12}\beta_{20} - \beta_{10}\beta_{22}) & (\beta_{10}\beta_{21} - \beta_{11}\beta_{20}) \\ (\beta_{02}\beta_{21} - \beta_{01}\beta_{22}) & (\beta_{00}\beta_{22} - \beta_{02}\beta_{20}) & (\beta_{01}\beta_{20} - \beta_{00}\beta_{21}) \\ (\beta_{02}\beta_{11} - \beta_{01}\beta_{12}) & (\beta_{02}\beta_{10} - \beta_{00}\beta_{12}) & (\beta_{00}\beta_{11} - \beta_{10}\beta_{01}) \end{pmatrix}, \quad (3.18)$$

with

$$|\boldsymbol{\beta}| = \det(\boldsymbol{\beta}) = \sum_{i,j,k=0}^2 \epsilon_{i,j,k} \cdot \beta_{0,i} \cdot \beta_{1,j} \cdot \beta_{2,k} \quad (3.19)$$

being the determinant of the system matrix. This generalized derivation is explicitly applied on the problem of statistical phase retrieval in appendix A.1. The actual computer implementation follows the equations in the appendix. After solving this system of equations for the three fit parameters in the solution vector \mathbf{a} , their variances and covariances can be extracted from the inverse of the system matrix $(\mathbf{B}^T \cdot \mathbf{B})^{-1}$, also called the covariance matrix of the system. However, as we do not estimate the three contrast signals directly—their relation to the estimated parameters is stated in equation (3.8)—further steps have to be applied to calculate the statistical uncertainties of the contrast signals from those of the estimated parameters. For this purpose we use standard Gaussian error propagation,* resulting in

$$\sigma_{a_0} = \sigma_{A_0}, \quad (3.21)$$

$$\sigma_{\phi_1} = \sqrt{\frac{A_1^2 \sigma_{B_1}^2 + B_1^2 \sigma_{A_1}^2}{A_1^2 + B_1^2}}, \quad (3.22)$$

$$\sigma_v = \sqrt{\frac{(A_1^2 + B_1^2) \sigma_{A_0}^2}{A_0^4} + \frac{A_1^2 \sigma_{A_1}^2 + B_1^2 \sigma_{B_1}^2}{A_0^2 (A_1^2 + B_1^2)}}, \quad (3.23)$$

where the covariance terms have been left out to make the display of the expressions more comprehensible. In practice the covariances are always taken into account.

Figure 3.4 shows the relation between the new parameters and the Fourier coefficients, as well as indicating their uncertainties. As usual, measurements

*Given $f(x_0, \dots, x_N)$ and $\sigma_{x_0}, \dots, \sigma_{x_N}$, σ_f follows

$$\sigma_f^2 = \sum_{i,j=0}^N \frac{\partial f}{\partial x_i} \frac{\partial f}{\partial x_j} \sigma_{x_i, x_j}, \quad (3.20)$$

where σ_{x_i, x_j} is the covariance between x_i and x_j and $\sigma_{x_i}^2 = \sigma_{x_i, x_i}$ denotes the variance of x_i .

with a sample in the beam and one without sample, i.e. a flatfield or reference image, are combined to get the resulting projections used as input for the tomographic reconstruction

$$a = -\ln\left(\frac{a_s}{a_r}\right), \quad \phi = \arg(e^{i(\phi_s - \phi_r)}), \quad v = -\ln\left(\frac{v_s}{v_r}\right), \quad (3.24)$$

where the subscript 'r' denotes the reference or flatfield scan and 's' the scan with sample in the beam. These are also functions of parameters, of which we know the uncertainty, so again applying standard error propagation yields the uncertainties of each of these signals

$$\sigma_a^2 = \frac{\sigma_{a_s}^2}{a_s^2} + \frac{\sigma_{a_r}^2}{a_r^2}, \quad \sigma_\phi^2 = \sigma_{\phi_s}^2 + \sigma_{\phi_r}^2, \quad \sigma_v^2 = \frac{\sigma_{v_s}^2}{v_s^2} + \frac{\sigma_{v_r}^2}{v_r^2}. \quad (3.25)$$

With equations (3.24) and (3.25) we now have analytical expressions for the three grating-interferometer signals and their statistical uncertainties, that can be explicitly and quickly calculated from the raw data. These expressions can later on directly be used in the statistical iterative reconstruction algorithm presented in chapter 4.

3.3 Influence of the system's point spread function on the tomographic reconstruction

Deconvolution is a common technique in many imaging fields, for example in microscopy or astronomical imaging. It is used whenever a measurement system distorts the gathered data along the optical path. This influence usually manifests as blurring of the measured images and is described by the so called point spread function (PSF). The PSF is a mathematical function that describes the response of the optical system to a point source, i.e. how a perfect point is spread or blurred because of the measurement device's imperfections. The PSF usually includes contributions from the source, any component in the beam path and the detection system. In the case of grating-based phase-contrast imaging these contributions are the size and shape of the incoming X-ray beam, imperfections in the gratings and the detector response. To my knowledge, in the context of GBI the effects of the PSF have so far not been taken into account. In this section we will present a qualitative analysis of how the PSF influences the sharpness of tomographic phase-contrast reconstructions and a simple but effective way to deconvolve the raw detector

images prior to further processing to increase the spatial resolution of the reconstructions. Especially if an experiment uses scintillator-coupled devices, such as CCD cameras or flat-panels, as a detection system, its PSF can have a significant influence on the spatial resolution of tomographic reconstructions.

In this section the effect of the PSF is investigated for an experimental measurement performed at the beamline ID19 of the European Synchrotron Radiation Facility (ESRF) in Grenoble, France. The grating interferometer at this beamline uses a FReLoN (Fast Readout Low Noise) camera coupled with a scintillator to convert incoming X-ray photons into visible light. The camera was first described by Labiche et al. (1996, 2007) and put into the context of grating interferometry at ID19 by Zanette (2011).

The deconvolution itself was performed using the Richardson-Lucy (RL) algorithm first introduced in Richardson (1972) and Lucy (1974). This algorithm is based on a maximum likelihood estimation of the latent or underlying image by an iterative update procedure. A derivation of the algorithm is given in appendix A.2. It leads to the following update equation

$$\mathbf{O}^{(i+1)} = \mathbf{O}^{(i)} \cdot \left(\bar{\mathbf{P}} * \frac{\mathbf{I}}{\mathbf{P} * \mathbf{O}^{(i)}} \right), \quad (3.26)$$

where \mathbf{I} is the measured image, \mathbf{O} is the underlying true object and \mathbf{P} is the known PSF. Furthermore, $*$ denotes the convolution operator, $\bar{\mathbf{P}}$ is the flipped PSF (as explained in appendix A.2) and the superscript (i) represents the iteration number.

A tomographic measurement of a mouse performed at the grating interferometer of beamline ID19 is used to demonstrate the effects of the PSF deconvolution on experimental data. As the true PSF of the system was not known, it was assumed as an isotropic two-dimensional Gaussian function. As we are not interested in a strictly quantitative analysis of the deconvolution, but only in the influence, that the PSF has on tomographic reconstructions, our criterion of image quality is solely based on visual inspection. In this regard, we achieved the visually best results with $\sigma = 1$ px, a size of 23×23 px² and 20 iterations of the RL algorithm. The deconvolution was performed on each raw projection of the stepping scan prior to the SPR. In the following, images derived from the raw images will be denoted by 'original' and images derived from the deconvolved raw data will simply be called 'deconvolved'.

Figure 3.5 shows two raw detector images. The upper panel (A) presents the image as it was originally captured by the camera. In the lower panel (B) the same image can be seen after deconvolution. For both images a zoom




Figure 3.5: This figure shows a comparison of a single raw detector image prior to deconvolution (A) and after deconvolving with the PSF (B). In addition to the full projection also a zoom of a part of the spine is presented as an inset. The deconvolved image is significantly clearer and sharper with respect to the fine details visible in the bone of the mouse.

of the region marked with the red rectangle containing parts of the spine is shown as an inset. When comparing these zoomed images it is obvious that the original image is blurred due to the effect of the PSF, whereas the deconvolved image is much clearer and sharper, especially when looking at edges or small details.

The same result can also be observed in the three contrast signals after the phase retrieval. Figure 3.6 shows enlargements of the same region used in figure 3.5 for all three contrast channels. The left column depicts projections retrieved from the original detector images, the right column the projections retrieved from the deconvolved raw data. The top row—panels (A) and (B)—displays the comparison for the attenuation signal, whereas the middle (C, D) and bottom (E, F) rows depict the differential phase- and dark-field signals, respectively. We can see that for all three modalities the image on the right, i.e. the one retrieved from the deconvolved data, is sharper and shows small details more clearly than the image on the left. So from this figure we can conclude that a deconvolution of the stepping images prior

to the phase retrieval not only significantly improves the visual quality and sharpness of the raw images, but these improvements are also propagated to the individual contrast signal's projections. The black and white pixels, which can be seen inside the bone structure in the differential phase-contrast projections (panels (C) and (D)), are caused by phase wrapping. This effect does not influence the conclusions drawn in this section and will be treated in chapters 4 and 5.

The last and most important part is now to analyze the effect of the deconvolution on the tomographic reconstruction. As an example, the differential phase projections were reconstructed with the conventional FBP algorithm, using a standard ramp filter in order to retain the maximum amount of sharpness in the resulting images. The reconstructions from both the raw and the deconvolved projections of the differential phase signal are shown in figure 3.7. The top row contains a slice of the full reconstructed volume from the original raw projections (A) and from the data that was deconvolved prior to the phase retrieval (B). The display window of both images is $\delta_{A,B} = [4.15 \cdot 10^{-7}, 5.00 \cdot 10^{-7}]$. As the difference in sharpness is barely visible in these full images, the figure also shows enlargements of two regions that contain fine detail. A zoom of the region marked with the red rectangle in panel (A), covering parts of the liver and the stomach, is shown in panel (C) for the reconstruction from the original projections and in panel (D) for the reconstruction from the deconvolved data. The display window for these two images is in the range $\delta_{C,D} = [4.40 \cdot 10^{-7}, 4.64 \cdot 10^{-7}]$. When comparing the two enlargements it is not only obvious that the deconvolved reconstruction is significantly sharper but also local contrast is improved in the deconvolution process. The same observation is true for the enlargement of the second region—marked with the green rectangle in panel (A)—, that shows parts of the intestines and also parts of the liver. This region is depicted in panels (E) and (F). Again, the reconstruction from the original dataset is by far blurrier and small details tend to be less visible than in the image reconstructed from the deconvolved data. The display window here is $\delta_{E,F} = [4.43 \cdot 10^{-7}, 4.64 \cdot 10^{-7}]$.

This analysis was done for a case, where the most probable cause of the blurring was the detector. This is shown in figure 3.8 (A). Depicted is the case of a parallel-beam geometry, i.e. the sample point is illuminated by a single ray. Due to the detector PSF, however, the projection of this single point is spread over multiple detector pixels.

Another possible origin of image distortion is discussed in the following. At the rotating anode setup of our laboratory, for example, a Pilatus 100K




Figure 3.6: Enlargements of spinal region for the three modalities. Top row: attenuation signal retrieved from original data (A) and from deconvolved data (B). Middle row: differential phase-contrast signal retrieved from original data (C) and from deconvolved data (D). Bottom row: darkfield signal retrieved from original data (E) and from deconvolved data (F).

detector is used, that has a boxlike PSF with a width smaller than the pixel size. Here the influence of the detector can be ruled out with a high certainty, but we still observe slight blurring of the raw images. In this case it is due to




Figure 3.7: Phase contrast reconstructions. Top row: full slice from original projections (A) and deconvolved projections (B). Middle row: zoom of region marked with red rectangle, (C) original, (D) deconvolved. Bottom row: zoom of region marked with green rectangle, (E) original, (F) deconvolved. Display windows: $\delta_{A,B} = [4.15 \cdot 10^{-7}, 5.00 \cdot 10^{-7}]$, $\delta_{C,D} = [4.40 \cdot 10^{-7}, 4.64 \cdot 10^{-7}]$, $\delta_{E,F} = [4.43 \cdot 10^{-7}, 4.64 \cdot 10^{-7}]$

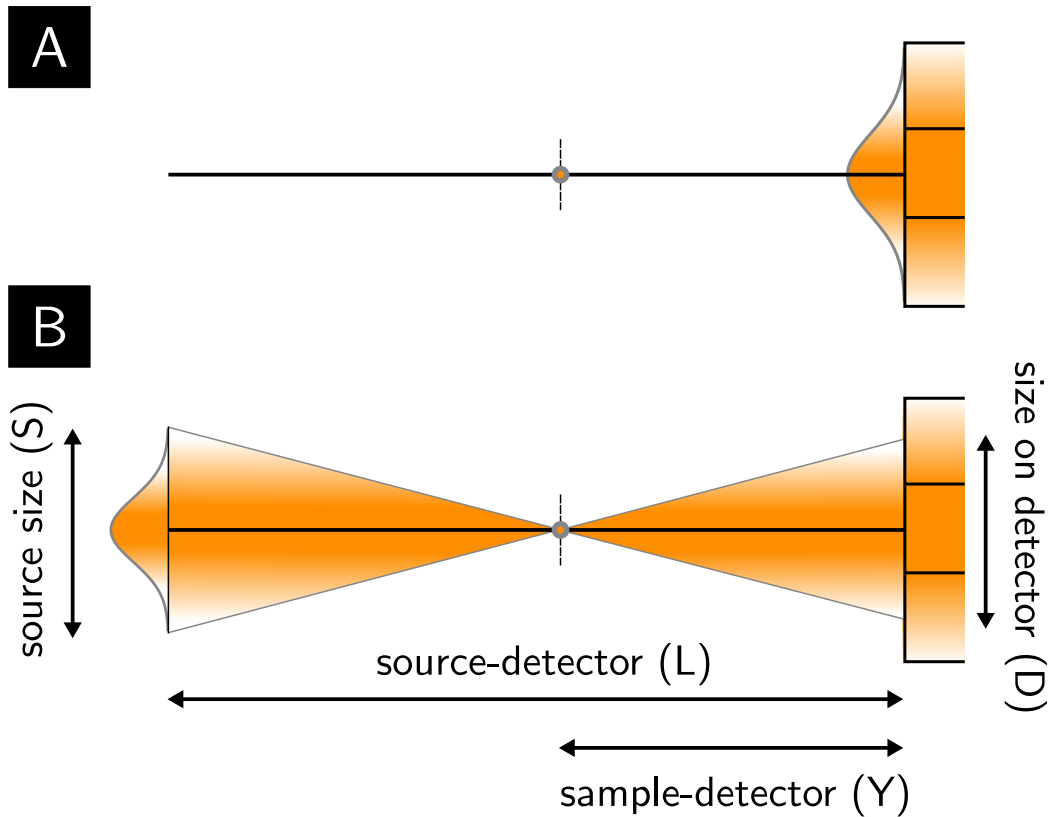


Figure 3.8: The influence of the PSF for the two extreme cases. (A) Only the detector PSF has an influence on how the projection of a point is spread over the detector pixels. (B) Only the finite size of the source contributes to a spread of the projection of a point on the detector.

the X-ray tube's focal spot, that is large enough to smear out the projection of a point-like structure over multiple detector pixels. This is illustrated in figure 3.8 (B). Neglecting effects from the detector, we concentrate only on the influence of the source size. As shown in the figure, the single sample point is illuminated by multiple rays from different directions and thus the projection of the point on the detector will also have finite size. Using the intercept theorem it is easy to estimate the blur due to the finite size of the focal spot. Consider

$$\frac{S}{L - Y} = \frac{D}{Y}, \quad (3.27)$$

where S is the source size, D the size of the projection of the sample point on the detector, L the distance between source and detector and Y the distance between sample point and detector. For illustration we will estimate this for the rotating anode lab setup mentioned above. Here the focal spot has

a size of $S = 0.3$ mm, a source-detector distance of $L = 2360$ mm and in typical experiments a sample-detector distance of $Y = 990$ mm. With the intercept theorem this amounts to a projection size on the detector of $D = 0.217$ mm = 217 μ m. Comparing this to the pixel size of the Pilatus 100K, which is $p_s = 172$ μ m, we see that even though the PSF of the detector itself does not have a significant effect on the spatial resolution, blurring is still observed due to the finite size of the source. Just like in the case of blurring due to the detector PSF, the image quality can significantly benefit from a deconvolution.

3.4 Summary and outlook

In this chapter we presented a collection of several methods and techniques based on radiographic images. The first section focused on the Relative Contrast Gain as a simple and practical tool to determine the relative contrast-to-noise-ratio of the attenuation- and differential phase-contrast signals that are available from grating-based X-ray imaging. The attenuation projection was differentiated to bring it into a comparable modality to the differential phase projection. The Relative Contrast Gain was then defined as the ratio of the contrast-to-noise ratio of a region of interest in the differential phase projection to the contrast-to-noise ratio in the same region in the differential attenuation projection. We applied the RCG analysis on experimental data of human breast tissue acquired on a grating-based imaging setup, yielding an RCG of 9.5 for a region containing strong features and 5.8 for a region with barely visible features. This indicates that the RCG is a local measure depending on the properties of the analyzed material. The last two sections demonstrate improvements of raw detector images from grating interferometry in preparation for statistical reconstruction discussed in the following chapter. The first of these two sections introduced an alternative method for processing phase-stepping sequences to extract the three contrast channels. The method uses a weighted least-squares approach to not only estimate the contrast signals, but simultaneously also their statistical uncertainties. Finally, in the last section we presented a study of how the influence of the system's point spread function propagates through the measurement, processing and reconstruction chain. We found that, by using a simple deconvolution approach and an estimate of the true PSF, we could significantly improve the image quality and sharpness in all three stages—raw acquisitions, processed contrast signals and reconstructed images.

Chapter 4

Statistical reconstruction framework

This chapter describes the main scientific contribution of this thesis. A statistical iterative reconstruction algorithm for grating-based phase-contrast CT is developed and tested with synthetic and experimental data. Section 4.1 motivates the use of iterative reconstruction—in particular for grating-based phase-contrast imaging. Section 4.2 introduces the statistical iterative reconstruction framework. Section 4.3 presents a visual validation of the algorithm for simulated and experimentally obtained data for a variety of cases of missing data, as well as qualitative convergence and noise analyses and section 4.4 gives an outlook to a more correct modelling of the statistical properties. The main results of this chapter are currently in preparation for publication.

4.1 Motivation

Iterative reconstruction techniques have seen a great deal of interest in the CT community for their ability to increase reconstruction quality—and thus diagnostic value—or to keep the quality at the same level while lowering patient doses. Grating-based phase contrast shares the same goals as conventional CT, but, due to the additional elements—the gratings—in the beam path and the required stepping procedure, the dose applied to a patient is intrinsically higher in grating interferometry. Even though it offers additional diagnostic information, we aim to produce diagnostically valuable reconstructions with as little dose as possible. The past decades have seen much effort in conventional CT reconstruction to develop and improve iterative reconstruction

with the goal of decreasing the patient radiation dose. We can learn from this experience and redevelop these already established techniques to make them applicable to our novel problems. Especially if the statistical properties of the raw data are taken into account, good reconstructions even from very noisy projections resulting from low-dose scans are still possible. In addition, iterative reconstruction can help to reduce artifacts that arise when the projection data is not complete or corrupt and thus contains unreliable data in some regions.

4.2 Statistical reconstruction algorithm (SIR)

In this section the different parts of the statistical iterative reconstruction (SIR) algorithm are derived and explained. The main concept of the algorithm is the application of maximum-a-posteriori (MAP) principles in a Bayesian sense. This principle is an extension of maximum likelihood (ML) estimation and combines the likelihood function with prior knowledge of the latent parameters.

4.2.1 Model

The most important part of any iterative reconstruction (IR) algorithm is the physical model that incorporates information on the image formation process and relates the measured projections to the reconstructed image values. In this part we will describe how different aspects have been modelled in our reconstruction framework and also the reasoning behind the choices we made. We loosely follow the structure used by Nuyts et al. (2013).

Discretization

The first aspect of the physical model is the discretization of the continuous distribution of reconstructed image values. IR algorithms are not restricted to a certain discretization. The grid size, i.e. the distance between two volume elements, can be chosen as finely as desired, only restricted by computational requirements. Also the shape of each volume element presents several degrees of freedom. In general, a reconstructed volume is represented as a weighted sum of spatial basis functions. Two of the most common choices in IR for such a basis set are voxels and blobs. Voxels are cubic and

non-overlapping basis functions that discretize a volume into a regular grid. Each voxel is isotropic and represents a single value. Blobs on the other hand model each discretized point in the volume with spherically symmetric Kaiser-Bessel functions (Lewitt, 1990; Matej and Lewitt, 1996; Ziegler et al., 2006) with several parameters to determine the shape and spatial extension. Those functions are bell-shaped, but, in contrast to a Gaussian function, they taper off to zero to prevent truncation of a volume element. The blobs are placed on a regular grid in such a way that they overlap, so that each element also contributes to its neighbors. Several authors have shown that using blobs in IR is advantageous in terms of image quality and noise suppression (Matej and Lewitt, 1996; Ziegler et al., 2006), but, in contrast to voxels, which are perfectly fit for a computer representation as arrays, the blob formulation is computationally more expensive and the parameters of the Kaiser-Bessel functions have to be chosen according to the reconstruction problem and geometry. In this thesis, due to the lower computational requirements, a voxel representation is used. In particular, the representation of a quantity $\delta(x, y, z)$ as a sum of basis functions is

$$\delta(\vec{x}) = \sum_{j=1}^{N_p} \delta_j \chi_j(\vec{x}), \quad (4.1)$$

where N_p is the number of elements δ_j in the volume and basis functions $\chi_j(\vec{x})$. For the voxel representation the basis functions can be written in three dimensions as

$$\chi_j(x, y, z) = \text{rect}\left(\frac{x - x_j}{\Delta}\right) \text{rect}\left(\frac{y - y_j}{\Delta}\right) \text{rect}\left(\frac{z - z_j}{\Delta}\right), \quad (4.2)$$

where the point (x_j, y_j, z_j) denotes the center of the j th voxel and Δ is the voxel size. The basis functions are defined in a way that they are 1 inside the j th voxel and 0 outside.

Forward and back projection

The decision to use the voxel model for the volume discretization stands in relation to the implementation of the forward- and back-projection operation. In this work we use a ray-driven model (Joseph, 1982; Siddon, 1984a,b) for the forward projection and a voxel-driven model for the back projection (see Mueller et al. (1998) for a comparison of ray-driven and voxel-driven). The reason for this choice is the implementation on general purpose graphics processing units (GPGPUs). For this massively parallel implementation it

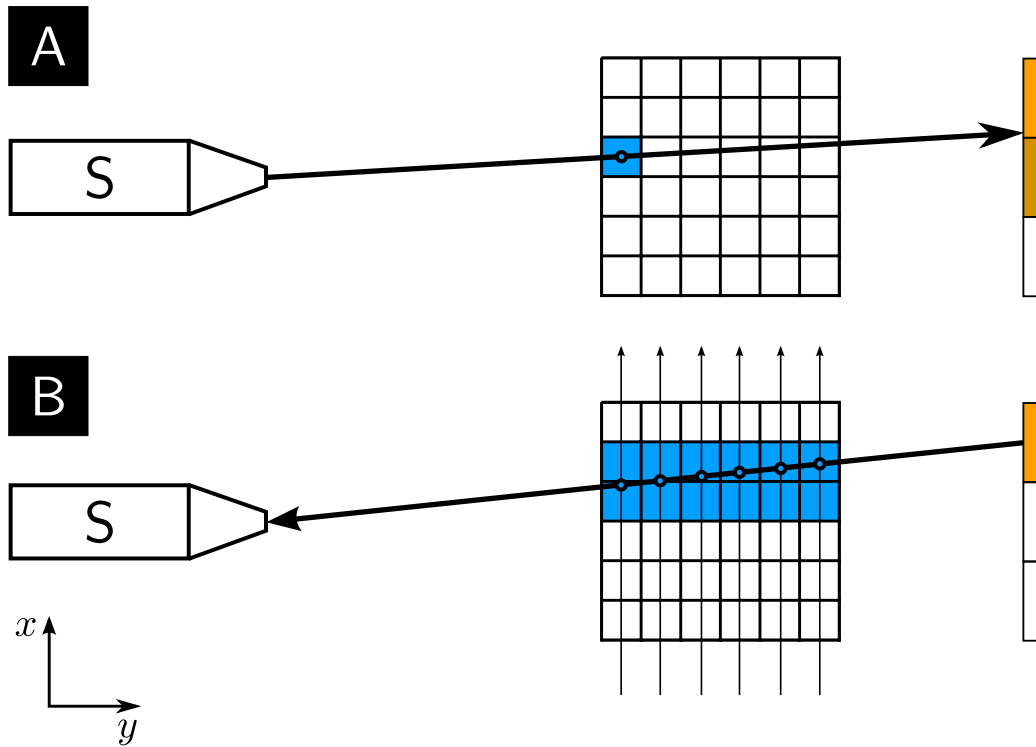


Figure 4.1: This figure visualizes the actual implementation of the way detector pixel values are projected back to an image voxel (A) and the way image voxels are projected forward to a detector pixel (B).

is advantageous to use a loop over the intended target volume, i.e. to treat each element of the target volume separately to prevent eventual race conditions (Fehringer et al., 2013). A schematic view of the implementation is shown in figure 4.1.

In the back projection operation the main loop runs over the image voxels and updates them separately using a voxel-driven projection model. Each of the image voxels is projected onto the detector and the linearly interpolated pixel value is added to the voxel value. The implementation of the back projection is illustrated in figure 4.1(A). The back projection is complete once each voxel and each projection angle have been processed.

For the forward projection the loop is running over the detector pixels in combination with ray tracing to determine the image voxels that intersect the particular ray and thus contribute to the update of the corresponding detector pixel. A ray starting from each detector pixel is traced back to the source position. The linearly interpolated values of all image voxels that

intersect this ray are added to the detector pixel value. The implementation of the forward projection is illustrated in figure 4.1(B). Again this operation is complete after processing each detector pixel and each projection angle.

Finite spatial resolution

The spatial resolution of reconstructed images is limited by several factors that have been described in detail in section 3.3. These factors include blurring by a detector PSF that is larger than the pixel size and the finite size of the X-ray source's focal spot. All of these effects originate from the physical properties of the measurement setup and manifest in the raw detector images. In conventional CT these effects can directly be modeled in the IR algorithm, as in this case the raw projections are used as input for the reconstruction. In grating-based phase-contrast CT an intermediate step—the signal extraction described in section 2.2.3 or the SPR described in section 3.2—has to be performed to acquire the input signal for the reconstruction. All the physical effects that lead to a limited spatial resolution have to be corrected for before the signal extraction step, and thus are not incorporated in the SIR algorithm presented in this chapter.

Statistical model

In grating-based phase-contrast imaging we measure the gradient of the line integrals of the decrement of the refractive index δ along the rays. The total phase shift along the i th ray with length L_i is given according to section 2.1

$$\Phi(x, z) = \int_{L_i} \delta(x, y, z) k \, dy, \quad (4.3)$$

where the ray travels along the y direction. This equation is related to the quantity that is directly measured in grating interferometry—the refraction angle—by

$$\alpha(x', z) = \frac{\lambda}{2\pi} \frac{\partial \Phi(x', y', z)}{\partial x'}, \quad (4.4)$$

where the prime now denotes coordinates in the coordinate system rotated around the z axis so that the beam direction points along y' . This equation together with the equation for the total phase shift yields

$$\alpha(x', z) = \partial_{x'} \int_{L_i} \delta(x', y', z) dy'. \quad (4.5)$$

We call this equation the forward model of our reconstruction problem. In the following, we leave out the primes for matters of clarity. Adding the discretization using voxel basis functions as discussed before this integral now becomes

$$\alpha(x, z) = \partial_x \int_{L_i} \sum_{j=0}^{N_p-1} \delta_j \chi_j(x, y, z) dy = \partial_x \sum_{j=0}^{N_p-1} \delta_j \int_{L_i} \chi_j(x, y, z) dy. \quad (4.6)$$

With the last result we can rewrite the forward model into a matrix equation for the i th measurement of the refraction angle vector

$$\alpha_i = [\partial_x \mathbf{A} \boldsymbol{\delta}]_i = \sum_{j=0}^{N_p-1} \partial_x a_{ij} \delta_j, \quad (4.7)$$

where the matrix \mathbf{A} denotes the forward projection operation defined earlier.

We can use this system of $N_{\mathcal{A}}$ equations to formulate a least-squares objective function to get an estimate of the parameter vector $\boldsymbol{\delta}$

$$S = \sum_{i=0}^{N_{\mathcal{A}}-1} \left(\alpha_i - \sum_{j=0}^{N_p-1} \partial_x a_{ij} \delta_j \right)^2. \quad (4.8)$$

However, from section 3.2 we know that each measurement α_i has its own statistical uncertainty σ_i . This can be included in equation (4.8) by attaching a weighting factor $w_i = 1/\sigma_i^2$ to each term in the sum, giving the following objective function, which, when minimized, results in the weighted least-squares estimator $\hat{\boldsymbol{\delta}}_{WLS}$

$$\hat{\boldsymbol{\delta}}_{WLS} = \arg \min_{\boldsymbol{\delta}} \sum_{i=0}^{N_{\mathcal{A}}-1} w_i \left(\alpha_i - \sum_{j=0}^{N_p-1} \partial_x a_{ij} \delta_j \right)^2. \quad (4.9)$$

This will make sure that measurements with a higher statistical uncertainty, i.e. that contain less reliable information, do not distort the resulting estimate as much as in the unweighted case. In addition, the weights w_i can be modified with a mask to exclude measurements from the reconstruction that are a-priori known to contain no reliable information. This can for example be used to mask out corrupt projections in a sinogram, which would lead to artifacts in the reconstruction if they were treated with the same weight as correct projections. This mechanism will be used extensively in section 4.3 when the algorithm is tested on different cases of corrupt data.

Log-likelihood formulation Equation (4.9) can also be derived in the sense of a log-likelihood formulation. If we assume that the α_i are realizations of a Gaussian distributed random variable \mathcal{A} , i.e. $\mathcal{A}_i \sim \mathcal{N}(\bar{\alpha}_i(\boldsymbol{\delta}), \sigma_i)$, the probability of observing a certain measurement for a particular set of parameters $\boldsymbol{\delta}$ is

$$\Pr(\mathcal{A} = \boldsymbol{\alpha} | \boldsymbol{\delta}) = \prod_{i=0}^{N_{\mathcal{A}}-1} \Pr(\mathcal{A}_i = \alpha_i | \boldsymbol{\delta}) = \prod_{i=0}^{N_{\mathcal{A}}-1} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\alpha_i - \bar{\alpha}_i(\boldsymbol{\delta}))^2}{2\sigma^2}\right), \quad (4.10)$$

where

$$\bar{\alpha}_i(\boldsymbol{\delta}) = [\partial_x \mathbf{A}\boldsymbol{\delta}]_i. \quad (4.11)$$

Equation (4.10) as the probability of observing a measurement $\boldsymbol{\alpha}$ given a fixed set of parameters $\boldsymbol{\delta}$ can also be viewed as a function of the parameters given a set of fixed measurements:

$$\mathcal{L}(\boldsymbol{\delta} | \boldsymbol{\alpha}) = \Pr(\mathcal{A} = \boldsymbol{\alpha} | \boldsymbol{\delta}). \quad (4.12)$$

In statistics this is then called the likelihood of $\boldsymbol{\delta}$ being the true parameters if a specific set of measurements has been observed. The method of maximum likelihood then seeks to maximize this function. Often, especially when working with exponential probability functions, it is easier to use the natural logarithm of the likelihood function, which is called the log-likelihood. For equation (4.10) the log-likelihood is then

$$l(\boldsymbol{\delta} | \boldsymbol{\alpha}) = - \sum_{i=0}^{N_{\mathcal{A}}-1} \log(\sqrt{2\pi\sigma^2}) - \sum_{i=0}^{N_{\mathcal{A}}-1} \frac{1}{2\sigma^2} (\alpha_i - [\partial_x \mathbf{A}\boldsymbol{\delta}]_i)^2 \quad (4.13)$$

As the first term of the log-likelihood does not depend on $\boldsymbol{\delta}$, it does not influence the parameter set that maximizes the function and thus can be dropped. In the ML formalism the reconstruction is then a matter of minimizing the negative log-likelihood (or maximizing the log-likelihood)

$$\hat{\boldsymbol{\delta}}_{ML} = \arg \min_{\boldsymbol{\delta}} \sum_{i=0}^{N_{\mathcal{A}}-1} \frac{1}{2\sigma^2} (\alpha_i - [\partial_x \mathbf{A}\boldsymbol{\delta}]_i)^2, \quad (4.14)$$

which is equivalent to minimizing the WLS objective of equation (4.9). In most cases, however, —especially when working with noisy measurements— estimating the true parameter set $\boldsymbol{\delta}$ by maximizing the log-likelihood alone is not sufficient. The problem of tomographic reconstruction is ill-posed, i.e. there are many parameter vectors $\hat{\boldsymbol{\delta}}$ that can lead to the specific measurements $\boldsymbol{\alpha}$. The result of an unconstrained reconstruction, i.e. a reconstruction




Figure 4.2: This figure illustrates an iterative reconstruction that only takes into account the likelihood term, i.e. represents the unconstrained reconstruction. Panel (A) shows the image, which was used as the initial guess for the reconstruction code and panel (B) the estimate after 10 iterations.

that only takes into account the likelihood term, is shown in figure 4.2. The left part of the figure represents the initial estimate, which the IR was started with and the right part demonstrates the estimate after 10 iterations of the unconstrained reconstruction. It is obvious that this nonsmooth solution is not the correct one, the algorithm simply reached the nearest solution that fulfills the maximum likelihood principle. To restrict the space of possible solutions to the one that best explains the observed data, the likelihood is augmented by an additional term that is commonly known as regularization or penalty function. More details on the regularization that is used in this work are presented in section 4.2.2. Usually this function is chosen in a way to enforce the reconstructed distribution of $\boldsymbol{\delta}$ to be smooth. The final objective function is

$$\hat{\boldsymbol{\delta}} = \arg \min_{\boldsymbol{\delta}} l(\boldsymbol{\delta}|\boldsymbol{\alpha}) + \lambda R(\boldsymbol{\delta}), \quad (4.15)$$

where λ is a tunable parameter that determines the strength of the regularization term with respect to the likelihood term. The right choice of λ will represent a compromise between spatial resolution and noise reduction of the reconstructed $\boldsymbol{\delta}$ vector. Equation (4.15) is also called penalized log-likelihood (Fessler, 2000).

Relation to Bayesian statistics This last equation can also be derived using a Bayesian formulation.

Bayesian statistics use probability distributions to test hypotheses based on experimental observation. In the case of tomographic reconstruction the hypothesis to test is: "given this specific measurement α , δ are the true reconstructed values". So, let

$$p(\delta|\alpha) \quad (\text{read: probability of } \delta \text{ given } \alpha) \quad (4.16)$$

be the probability that δ are the correct parameters for fixed experimental outcomes α . In Bayesian terminology this is called the posterior probability. According to Bayes' Theorem (Bayes and Price, 1763; Laplace, 1986) it is defined as

$$p(\delta|\alpha) = \frac{\overbrace{p(\alpha|\delta)}^{=\mathcal{L}(\delta|\alpha)} \cdot p(\delta)}{p(\alpha)}. \quad (4.17)$$

In this equation the probability of acquiring experimental results based on fixed model parameters is equal to the measure of how likely the parameters are the true ones, when given a specific experimental outcome. This is the same notation that was used in the previous section on the derivation of the ML formulation. In contrast to the ML formulation, the posterior probability contains two additional terms. The term $p(\alpha)$ denotes the probability distribution of the experimental data points, but as it does not depend on the variable parameters δ it represents a multiplicative constant, which can be disregarded in the following steps. The second additional term $p(\delta)$, however, describes the probability distribution of the estimated parameters. In Bayesian terminology this distribution is called the prior. It encapsulates every a-priori knowledge* that one has on the estimated parameters. In image reconstruction this prior knowledge can be for example the notion that the reconstructed image has to be locally smooth or individual voxels having specific values.

In this context the aim is to find the parameter set $\hat{\delta}$ that maximizes the posterior probability

$$\hat{\delta} = \arg \max_{\delta} p(\delta|\alpha) = \arg \max_{\delta} \frac{p(\delta)}{p(\alpha)} \cdot \mathcal{L}(\delta|\alpha) \quad (4.18)$$

Discarding the probability distribution of the experimental data and taking the natural logarithm of equation (4.18) leads to the following optimization

*In Bayesian inference this prior knowledge can also be purely subjective.

problem for estimating $\boldsymbol{\delta}$

$$\hat{\boldsymbol{\delta}} = \arg \max_{\boldsymbol{\delta}} \log(\mathcal{L}(\boldsymbol{\delta}|\boldsymbol{\alpha})) + \log(p(\boldsymbol{\delta})), \quad (4.19)$$

where $\hat{\boldsymbol{\delta}}_{\text{MAP}}$ is now accordingly called the MAP estimator, i.e. the argument that maximizes the posterior probability. Obviously this is equivalent to the formulation of the penalized likelihood objective function (4.15) stated earlier.

If there is no prior knowledge, i.e. $p(\boldsymbol{\delta})$ is constant, equation (4.17) will just state that the posterior probability is equal to the likelihood function and the MAP estimator becomes the ML estimator (4.14) again.

4.2.2 Regularization

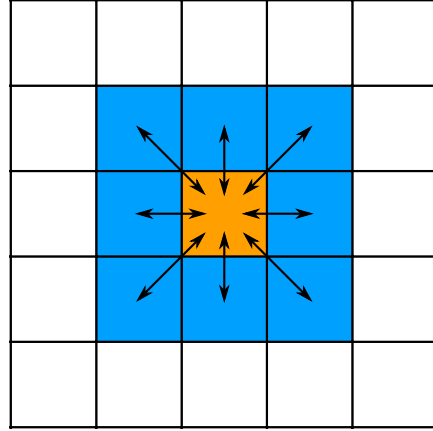
As already stated earlier, the maximum likelihood formulation by itself is numerically unstable, because it represents an ill-posed inverse problem, with the consequence that the space of possible solutions $\boldsymbol{\delta}_{\text{ML}}$ that maximize the likelihood is infinite. It was also noted already that a remedy for this problem is to add an additional term, which in mathematics is called regularization. In the Bayesian framework this regularization term represents prior knowledge about the solution and, in conjunction with the maximum likelihood estimator, the resulting statement constitutes the maximum-a-posteriori principle.

In this section several possibilities of implementing prior knowledge in the form of regularization terms are presented. Most regularization approaches in CT reconstruction are based on locally penalizing the differences between neighboring voxels as shown in figure 4.3. The general formulation is then (Fessler, 2000)

$$R(\mathbf{x}) = \sum_i \sum_{j \in \mathcal{N}_i} w_{ij} \Phi(x_i - x_j) \quad (4.20)$$

where the first sum index i runs over all image voxels and the second sum index j over all the direct neighbors of voxel i , with \mathcal{N}_i denoting the direct neighborhood of i . The coefficients w_{ij} are a weighting factor that is 1 for adjacent neighbors, $1/\sqrt{2}$ for diagonal neighbors in the same plane and $1/\sqrt{3}$ for diagonal neighbors in the adjacent planes to account for the varying inter-voxel distances in the individual differences. Finally, $\Phi(x)$ is an arbitrary scalar potential function, which—in the ideal case—is a convex function to keep the optimization stable. Two common choices for $\Phi(x)$ are a simple

Figure 4.3: Basic principle of locally restricted regularization. For each voxel (orange) in the current estimate of the reconstructed volume, the difference (arrows) of this voxel to its direct neighbors (blue) is computed. After applying a potential function to each difference, the sum of all function evaluations is added to the objective function as a penalty.



quadratic function and the Huber potential. For the quadratic potential function the regularization is defined as

$$R_Q(\mathbf{x}) = \sum_i \sum_{j \in \mathcal{N}_i} w_{ij} (x_i - x_j)^2, \quad (4.21)$$

so the summation is performed over the quadratic differences between the neighboring voxels. For large differences between neighboring voxels this results in a large value for the regularization term, preventing the reconstructed image to become non-smooth. This basically can be interpreted as a Gaussian prior distribution of the next-neighbor differences in the Bayesian formulation.

The second common regularization term is based on the Huber potential function (Huber, 1964). This function is piecewise defined with a quadratic and linear part as

$$\Phi(x, \gamma) = \begin{cases} \frac{(x)^2}{2\gamma^2} & \text{for } |x| \leq \gamma \\ \frac{|x| - \gamma/2}{\gamma} & \text{for } |x| > \gamma \end{cases}, \quad (4.22)$$

where γ is called the threshold and specifies the value that marks the transition from quadratic to linear behavior. This convex function and its first derivative are continuous at the threshold value.

The regularization based on this potential function is then defined as

$$R_H(\mathbf{x}, \mathbf{m}, \gamma) = \sum_i m_i \sum_{j \in \mathcal{N}_i} \begin{cases} \frac{(x_i - x_j)^2}{2\gamma^2} & \text{for } |x_i - x_j| \leq \gamma \\ \frac{|x_i - x_j| - \gamma/2}{\gamma} & \text{for } |x_i - x_j| > \gamma \end{cases}, \quad (4.23)$$




Figure 4.4: Comparison of the quadratic (orange curve) and Huber (blue curve) potential functions. The Huber threshold γ is shown in red and has a value of 3.0. The quadratic function is normalized to have the same scale as the quadratic part of the Huber function.

where again the first sum over i is performed over all image voxels and j is the index running over the direct neighbors of voxel i . The reason for using such a regularization term is to preserve edges that are present in the image that are not to be smoothed out. How this is accomplished is explained with the help of figure 4.4. For a more fine grained application of this regularization it can also be augmented by a mask \mathbf{m} that restricts the effect of the regularization to certain voxels. An example of where this is useful is given in section 5.1.3.

In figure 4.4 a comparison between the quadratic (orange curve) and the Huber (blue curve) potential functions is presented. Both functions are plotted for the range $x = (-10.0, 10.0)$. The threshold for the Huber function is set to $\gamma = 3.0$ and is marked with the red lines in the figure. The quadratic function is normalized to have the exact same values as the quadratic part of the Huber function. It can be clearly seen from the plot that for x larger than the threshold the values of the quadratic function are much larger than those of the Huber potential; thus, in terms of regularization, large differences between neighboring voxels are much more penalized in the quadratic case, forcing the voxel values to become much more similar during the minimization. This, in turn, leads to the statement that by using the Huber regularization term large voxel differences, which are likely to represent real sharp edges in the reconstructed volume, are penalized only linearly, while already more or less uniform areas are penalized quadratically.

A third regularization term used in this thesis is based on a slightly different kind of prior knowledge. It is now well known that, in contrast to conventional CT, where the noise power spectrum (NPS) is largest for high

frequencies, the noise in PCCT from grating interferometry is located more in the low frequency part of the spectrum (Köhler et al., 2011). This means that fluctuations in the reconstructed volume are of a longer range than usually observed in CT reconstructions and the concept of regularizing the differences between adjacent voxels as in the quadratic and Huber case is—while still effective—physically not perfectly suited for the noise properties of PCCT reconstructions. To remedy this problem another regularization term is proposed that can potentially be tuned to specific parts of the NPS. We call it the ‘mean regularization’. It is defined as

$$R_M(\mathbf{x}, \sigma) = \sum_i (x_i - \langle x \rangle_\sigma)^2, \quad (4.24)$$

where $\langle x \rangle_\sigma$ is the mean of the neighborhood of size σ around a voxel.

The last regularization term developed in the course of this thesis is called bone regularizer, as it was specifically designed for the purpose of bone artifact reduction (BAR). It is based on the assumption that the values in parts of the image are known a-priori. Consequently, the term is defined as the squared difference between the voxels in the reconstructed image and another set of corresponding voxels \mathbf{a} containing the a-priori known values. The definition is

$$R_B(\mathbf{x}, \mathbf{a}, \mathbf{b}, c) = \sum_i b_i (x_i - c \cdot a_i)^2, \quad (4.25)$$

where the mask \mathbf{b} specifies the parts of the image, which the regularization is supposed to have an effect on. The scalar factor c can be used to further scale the a-priori known values in the voxel set \mathbf{a} . This term will be used in sections 5.1.3 and 5.1.4 to reduce the artifacts originating from voxels containing no reliable DPC signal because of strong absorption, scattering and phase wrapping by replacing these voxels with the corresponding ones taken from an absorption reconstruction scaled by $c = \delta/\beta$.

4.2.3 Minimization

Just as a reminder, the complete cost function is restated here

$$\hat{\boldsymbol{\delta}} = \arg \min_{\boldsymbol{\delta}} \sum_{x,\theta} w_{x,\theta} \left(\alpha_{x,\theta} - \sum_j \partial_x A_{x,j}^\theta \delta_j \right)^2 + \sum_k \lambda_k R_k(\boldsymbol{\delta}), \quad (4.26)$$

with R_k being one or more of the regularization terms described above, each with its own strength parameter λ_k . The single index of the outer sum has

now been replaced with the sinogram coordinates: the projection angle θ and the detector channel x

For the minimization procedure we chose a standard conjugate gradient (CG) algorithm. This algorithm is similar to the method of steepest descent, where a function is minimized by calculating its gradient and then doing a line search to find the minimum along this direction and repeating this until the global minimum of the function is found. In CG after each line search one does not use the gradient at this point as a new search direction, but calculates a new direction that is conjugate to the previous one by using Gram-Schmidt orthogonalization. The CG algorithm, however, is in its basic form suitable only for the minimization of linear problems, i.e. purely quadratic functions that have a linear gradient. If the Huber term is used for regularization the cost function becomes partly non-quadratic, which means that its first derivative is non-linear. Because of this fact we use the non-linear variant of the conjugate gradient, which can be used for more general optimization problems. A pseudo code representation of the algorithm with comments is given in algorithm 1. The main difference between linear and non-linear CG is that the calculation of the step size to the minimum along the search direction is more complicated. In our non-linear conjugate gradient (NLCG) implementation we find the minimum along the search direction, i.e. the step size κ that minimizes the one-dimensional function along the search direction \mathbf{d} $f(\mathbf{x} + \kappa\mathbf{d})$, by setting the gradient of this function with respect to κ to 0 and following the Newton-Raphson method. This method is based on a truncated Taylor expansion of the function

$$f(\mathbf{x} + \kappa\mathbf{d}) \approx f(\mathbf{x}) + \kappa \left[\frac{d}{d\kappa} f(\mathbf{x} + \kappa\mathbf{d}) \right]_{\kappa=0} + \frac{\kappa^2}{2} \left[\frac{d^2}{d\kappa^2} f(\mathbf{x} + \kappa\mathbf{d}) \right]_{\kappa=0} \quad (4.27)$$

$$= f(\mathbf{x}) + \kappa [f'(\mathbf{x})]^T \mathbf{d} + \frac{\kappa^2}{2} \mathbf{d}^T f''(\mathbf{x}) \mathbf{d}, \quad (4.28)$$

and

$$\frac{d}{d\kappa} f(\mathbf{x} + \kappa\mathbf{d}) = 0 \approx [f'(\mathbf{x})]^T \mathbf{d} + \kappa \mathbf{d}^T f''(\mathbf{x}) \mathbf{d}. \quad (4.29)$$

To apply this we have to calculate the gradient and the Hessian of our cost function (4.26). The gradient calculation cannot be avoided because it is also needed for other parts of the algorithm. In order to not having to calculate the Hessian, which would be computationally very expensive considering the amount of data in CT, we use the relation that is directly given in equations (4.27) and (4.28)

$$\mathbf{d}^T f''(\mathbf{x}) \mathbf{d} \approx \left[\frac{d^2}{d\kappa^2} f(\mathbf{x} + \kappa\mathbf{d}) \right]_{\kappa=0}, \quad (4.30)$$

Algorithm 1 Pseudo code representation of the nonlinear conjugate gradient algorithm with Newton-Raphson line search and Polak-Ribiere formula for calculating conjugate directions

Precondition: $f(\mathbf{x}) = l(\mathbf{x}) + \lambda R(\mathbf{x})$ is a convex function not necessarily purely quadratic, \mathbf{x}_0 is a first guess of the solution

// preliminary initializations

$i \leftarrow 0$ *// initialize loop variable for main iterations*

$\mathbf{x} \leftarrow \mathbf{x}_0$

$\mathbf{r} \leftarrow -f'(\mathbf{x})$

$\mathbf{d} \leftarrow \mathbf{r}$ *// initial search direction equal to gradient*

$n_{\text{new}} \leftarrow \|\mathbf{r}\|^2$ *// gradient norm*

$\mathbf{g}_{\text{prev}} \leftarrow 0$ *// gradient in previous iteration*

while $i < i_{\text{max}}$ **do** *// begin iterative minimization*

$n_d \leftarrow \|\mathbf{d}\|^2$ *// norm of search direction vector*

$j \leftarrow 0$ *// initialize loop variable for Newton-Raphson iterations*

while $j < j_{\text{max}}$ **do**

$\kappa \leftarrow -[f'(\mathbf{x})]^\text{T} \mathbf{d} / \mathbf{d}^\text{T} f'' \mathbf{d}$

$\mathbf{x} \leftarrow \mathbf{x} + \kappa \mathbf{d}$ *// update of the solution vector by advancing along the current search direction by the step size*

$j \leftarrow j + 1$

$\mathbf{r} \leftarrow -f'(\mathbf{x})$

$n_{\text{old}} \leftarrow n_{\text{new}}$

$n_{\text{new}} \leftarrow \|\mathbf{r}\|^2$

$\beta_{\text{PR}} \leftarrow [n_{\text{new}} - [f'(\mathbf{x})]^\text{T} \mathbf{g}_{\text{prev}}] / n_{\text{old}}$ *// Polak-Ribiere formulation for calculating conjugate search directions*

$\mathbf{g}_{\text{prev}} \leftarrow f'(\mathbf{x})$ *// store gradient for next iteration*

$\mathbf{d} \leftarrow \mathbf{r} + \beta_{\text{PR}} \mathbf{d}$ *// update search direction*

$i \leftarrow i + 1$

which is much easier to calculate since we do not need the second derivative with respect to a vector but only with respect to a scalar. The resulting step size is then

$$\kappa = \frac{[f'(\mathbf{x})]^\text{T} \mathbf{d}}{\left[\frac{d^2}{d\kappa^2} f(\mathbf{x} + \kappa \mathbf{d}) \right]_{\kappa=0}}, \quad (4.31)$$

which is evaluated several times to make the line search more accurate.

What we do need to calculate is the first order gradient, i.e.

$$\frac{\partial}{\partial \delta_j} \left[l(\boldsymbol{\delta}) + \sum_k \lambda_k R_k(\boldsymbol{\delta}, \boldsymbol{\theta}) \right]. \quad (4.32)$$

Fortunately the cost function is a sum of multiple terms, so it is possible to calculate the gradient of the likelihood and the regularization terms separately. For the log-likelihood term the gradient is calculated as

$$\frac{\partial l(\boldsymbol{\delta})}{\partial \delta_j} = \frac{\partial}{\partial \delta_j} \sum_{x,\theta} w_{x,\theta} \left(\alpha_{x,\theta} - \sum_j \partial_x A_{x,j}^\theta \delta_j \right)^2 \quad (4.33)$$

$$= 2 \sum_{x,\theta} w_{x,\theta} \partial_x A_{x,j}^\theta \partial_x A_x^\theta \boldsymbol{\delta} - 2 \underbrace{\sum_{x,\theta} w_{x,\theta} \alpha_{x,\theta} \partial_x A_x^\theta}_{\mathbf{g}_\alpha} \quad (4.34)$$

or in vector notation

$$l'(\boldsymbol{\delta}) = 2 \sum_{x,\theta} w_{x,\theta} (\partial_x A_x^\theta)^\top (\partial_x A_x^\theta) \boldsymbol{\delta} - \mathbf{g}_\alpha, \quad (4.35)$$

where $(\partial_x \mathbf{A})$ is the differential forward projector. Its transpose represents the backward projection. Also note that the second term \mathbf{g}_α is constant and can thus be precomputed and stored. In addition we use the equality

$$(\partial_x \mathbf{A})^\top = \mathbf{A}^\top \partial_x^\top = \mathbf{A}^\top (-\partial_x), \quad (4.36)$$

with the interpretation of the partial derivative as a finite difference matrix.

The approximation of the Hessian, i.e. the denominator of κ , which is needed for the step size, is calculated as

$$\frac{\partial^2}{\partial \kappa^2} l(\boldsymbol{\delta} + \kappa \mathbf{d}) = \frac{\partial^2}{\partial \kappa^2} \sum_{x,\theta} w_{x,\theta} (\alpha_{x,\theta} - \partial_x A_x^\theta \boldsymbol{\delta} - \kappa \partial_x A_x^\theta \mathbf{d})^2 \Big|_{\kappa=0} \quad (4.37)$$

$$= 2 \sum_{x,\theta} w_{x,\theta} (\partial_x A_x^\theta \mathbf{d})^2, \quad (4.38)$$

so the complete step size definition is

$$\kappa = \frac{[l'(\boldsymbol{\delta})]^\top \mathbf{d}}{2 \sum_{x,\theta} w_{x,\theta} (\partial_x A_x^\theta \mathbf{d})^2}. \quad (4.39)$$

The same calculations have to be done for each regularization term. The explicit expressions for all regularization terms can be found in appendix A.3.

4.3 Evaluation of the algorithm

In this section the performance of the algorithm is tested against common cases of corrupt or incomplete data. The tests are performed on two different sets of data. The first part will deal with simulated projection data, whereas the second part will use an experimentally obtained data set to verify the algorithm. The reconstruction algorithm was tested on five different cases of corrupt or incomplete projection data and for reference, also a reconstruction from the complete data set was performed. In the order they are presented in the following sections, these six scenarios include:

(A) Full dataset: The reference reconstruction was performed on the complete data set.

(B) Noisy projections: For the simulations, Poisson noise was added to the raw detector images before doing the phase retrieval to simulate a scan with less exposure. The experimental data was measured with different exposure times, so one of the scans with low exposure was chosen for this test.

(C) Stationary gaps: In each projection, the information in several pixel columns was removed, i.e. those pixels were set to zero, to simulate stationary gaps in the projections that could be caused by tiling of individual detector modules or by seams created during the fabrication of larger gratings when stitching multiple smaller ones.

(D) Missing blocks of projections: The information in several contiguous projections in the angular range of 0° – 180° and in the corresponding projections in the range 180° – 360° was removed to simulate a CT scan that is angularly not fully sampled. This case corresponds to a slight tomosynthesis problem.

(E) Few projections: The number of projections is reduced to give a scan, which is undersampled by a factor of about 4, when taking into account that for an FBP reconstruction the ideal number of projections N_{proj} should be larger than $\pi/2 \cdot N_{\text{det}}$ (detector channels) to satisfy the Nyquist criterion (Kak and Slaney, 1988), i.e. to be free from aliasing artifacts.

(F) Very few projections: The number of projections is further reduced to an undersampling factor of about 8.5 with respect to the Nyquist criterion to produce even more aliasing artifacts for the algorithm to remove.

4.3.1 Synthetic data

For the first test of how well the algorithm can deal with corrupt or missing information, simulated data of a numerical phantom was used. In this section

the procedure of creating the test data is explained.

Simulation and preparation of projection data

For the simulation a framework based on wave propagation was used. A detailed description of this framework can be found in Malecki et al. (2012). This framework is able to model all components of an actual experimental setup. The simulation was started with a monochromatic plane wave with an energy of 23 keV. The plane wave was used to simulate a parallel beam geometry. For the grating interferometer two gratings were modelled. The phase grating G_1 consisted of silicon with a period of $4.80\ \mu\text{m}$ and a height of $29.5\ \mu\text{m}$. The analyzer grating G_2 consisted of gold with a period of $2.40\ \mu\text{m}$ and a height of $74.0\ \mu\text{m}$. The distance between the two gratings was 48.1 cm. Directly behind the analyzer grating a detector was placed. This detector was modelled with the attributes of a PILATUS photon counting detector with silicon as the sensor material and a pixel size of $15.36\ \mu\text{m} \times 15.36\ \mu\text{m}$. As a sample, a well defined numerical phantom with dimensions $1024 \times 1024 \times 128\ \text{px}^3$ was created, which is shown in

Figure 4.5: This figure shows a slice of the numerical phantom used as input for the simulation. The values for the different parts of the phantom, especially the small cylinders, are given in table 4.1. The large cylinder itself is simulated as plastic and filled with formalin. The region outside is made of water. The dimensions of the phantom are $1024 \times 1024 \times 128\ \text{px}^3$.




figure 4.5. It contains a plastic cylinder embedded in water and filled with formalin. Inside the large cylinder several small cylinders consisting of different materials found in breast tissue are placed. An overview of the materials and their δ -values is given in table 4.1 with symbols corresponding to those in the image of the phantom. The values were chosen to simulate a sample consisting purely of soft tissue with very weak contrast between the individual parts. For each of the simulated 360 projections a phase stepping with

Table 4.1: This table gives the numerical values of the phantom that was used in the simulations. The index corresponds to the symbols used in figure 4.5.

Index	δ	tissue name
I	4.059×10^{-7}	adipose tissue
II	4.643×10^{-7}	parenchyma
III	4.564×10^{-7}	tumor
IV	4.675×10^{-7}	ductal structure
V	4.360×10^{-7}	water
VI	4.086×10^{-7}	tube (plastic)
VII	4.466×10^{-7}	formalin

8 images over one period was performed. Every 60 projections a block of 20 simulated flatfield projections was inserted. A last, optional step was to add Poisson noise directly to the simulated detector images before continuing on with the statistical phase retrieval that was presented in section 3.2, to simulate noisy acquisitions. The combination of phantom dimensions, dimensions of the incoming plane wave and the detector pixel size was a compromise between getting reasonably large projection images, providing the simulation program enough resolution to actually acquire a phase signal and keeping memory consumption and simulation times feasible. The dimensions of the final projections due to these three parameters were $33 \times 257 \text{ px}^2$ and because of aliasing effects in the simulation code at the projection boundaries they were cropped in the phase retrieval to the final dimensions $28 \times 251 \text{ px}^2$.

Results

The results of the simulation study are summarized in figures 4.6 and 4.7. Shown from left to right are the sinogram, the FBP reconstruction and the SIR reconstruction, respectively. The parameters for the SIR reconstructions are summarized in table 4.2. Row A corresponds to the first of the six cases, the reference scan with full sampling and no noise. Before the phase retrieval only minimal Poisson noise was added to the simulated raw acquisitions, resulting in a mean photon count of about 1300 per pixel and stepping image, so it can be considered to be noise free for all practical purposes. Because there is neither corrupt nor missing data there is no visible difference between the FBP and SIR reconstructions.




Figure 4.6: First part of the results of the evaluation of the algorithm with simulated data. The left column shows the sinogram, the middle and right columns the reconstructions with FBP and SIR, respectively. A: full dataset; B: noisy data; C: stationary gaps. The second part is shown in figure 4.7.

For the second case—noisy projections—more Poisson noise was added to the raw projections, while reducing the number of counts, resulting in a mean photon count of around 13 photons per pixel and acquisition. The result is shown in row B. With this amount of noise the FBP begins to have trouble differentiating the tumorous tissue (the cylinders marked III in figure 4.5) from the surrounding formalin. However, in the SIR reconstruction the contrast between tumor and formalin is lower than in the full reconstruction, but it is still clearly visible. In addition, the noise has been completely removed from the reconstruction.

The third case—shown in row C—simulates stationary gaps, i.e. missing pixels at the same position for every projection. In this case we placed a 5 px wide gap every 11 px, which amounts to a fraction of almost 30% of pixels with no information. In the FBP reconstruction these gaps lead to severe ring artifacts and streaks radiating from the inserts with high contrast to the surrounding formalin. Although all of the inserts can still be clearly distinguished, the artifacts make it impossible to tell if there is another feature hidden in the surrounding formalin. The rings and streaks are successfully removed by the SIR reconstruction by modifying the statistical weights in the cost function to not consider the pixels that contain no information. The gaps are then filled according to the a-priori knowledge of the edges in the reconstructed volume, i.e. only governed by the regularization. However, due to gaps lying exactly on top of the boundary between plastic cylinder and formalin the algorithm cannot determine the values inside the container quantitatively anymore. While the values might not be quantitative any longer, the visible contrast between the inserts is faithfully reconstructed. Also, the formalin filled region around them is reconstructed as a flat area as it should be.

For the next case the angular views between 45° and 60° as well as between 225° and 240° of the 360 projections in the full scan were removed from the sinogram volume. The sinogram in row D gives an impression on how much of the data is missing. The FBP reconstruction exhibits the artifacts that are typical for the reconstruction of a tomosynthesis acquisition, where only a part of the angular range is recorded. These artifacts manifest as streaks in the angular direction of the missing projections because in the final step of the FBP, which is basically a sum of all the filtered and back projected projections, these missing projections do not contribute to the final image. The SIR algorithm on the other hand is able to slowly fill the gaps in the sinogram and create an almost flawless reconstruction. Here the missing projections are masked out in the cost function by modifying the statistical weights, i.e. we tell the algorithm that these projections contain

no information at all and it should fill the missing data with the help of the regularization alone.

The last two cases are related in that they both simulate different levels of undersampling artifacts. We know from Kak and Slaney (1988) that a certain minimum number of projections is needed for an artifact-free FBP reconstruction. This number is calculated by $N_{\text{proj}} \approx \pi/2 \cdot N_{\text{det}}$, so around 1.6 times the number of detector channels or number of beams per projection. For this phantom simulation with a sinogram width of 251 px this amounts to a minimum number of projections of 394, so theoretically the full scan is already undersampled. But we also have to consider what is contained in the field of view. The sample itself only occupies around 210 px of the full detector width. This means, that for an artifact-free reconstruction of the sample we need at least 330 projections, which is fulfilled in the full scan. The scan in row E was simulated with an undersampling of around 3.7, by reducing the number of projections from 360 to 90. The sinogram in figure 4.7 E is just stretched to have the same height as in the other rows. In the FBP reconstruction the undersampling—or aliasing—artifacts are clearly visible. The SIR reconstruction was done by first creating an empty sinogram with the dimensions of the original full scan and then distributing the remaining 90 projections evenly in there. The parts in between without information are masked out via the statistical weights. The result is a reconstruction that is closer to the original one from the full scan than the FBP reconstruction.

Table 4.2: This table gives the reconstruction parameters for the algorithm verification with synthetic data.

Case	# iterations	λ_{H}	γ_{H}
A	50	$1.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$
B	50	$2.0 \cdot 10^{-4}$	$3.0 \cdot 10^{-4}$
C	150	$1.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$
D	50	$1.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$
E	250	$1.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$
F	500	$1.0 \cdot 10^{-5}$	$1.0 \cdot 10^{-4}$

In the last case shown in row F the undersampling is increased even further. Now only 40 of the original 360 projections are used, giving an undersampling of around 8. In the FBP reconstruction we find that the severity of the aliasing artifacts is much more pronounced than in the previous case. However, the SIR algorithm is still able to reconstruct the flat, formalin-filled




Figure 4.7: Second part of the results of the evaluation of the algorithm with simulated data. The left column shows the sinogram, the middle and right columns the reconstructions with FBP and SIR, respectively. D: missing blocks of projections; E: short scan (90 out of 360 projections); F: ultrashort scan (40 out of 360 projections). The first part is shown in figure 4.6.

region and all of the inserts. The only drawback is that the relative contrast between the tumorous tissue (the cylinders marked III in figure 4.5) and the surrounding formalin, as well as between water (the cylinder marked V in figure 4.5) and formalin is drastically reduced. Still, all of the inserts can be clearly distinguished from the background.

4.3.2 Experimental data

However, software phantoms with distinct flat areas like the one used here favor every reconstruction algorithm and especially iterative algorithms that use any kind of edge preserving regularization or filtering. So using such a phantom can only be a part of testing the capabilities of an algorithm and one should not base the evaluation of algorithm performance solely on the study of analytical phantoms. These facts are the reason for performing the same study on medically relevant experimental data. So, in this section, the algorithm is tested with experimentally obtained measurements on the same cases of corrupt data as in the previous section. We use the PCCT scan of a human heart performed at a laboratory-based three-grating interferometer for this evaluation.

4.3.2.1 Experimental setup

The measurements were performed at a three-grating interferometer in a laboratory-based setup. It consists of a rotating anode X-ray tube as the source. The source grating G_0 is made of gold with a period of $5.4\ \mu\text{m}$ and a height of $65.0\ \mu\text{m}$. The phase grating G_1 consists of nickel and is $8.0\ \mu\text{m}$ high with a period of $5.4\ \mu\text{m}$. With these specifications the grating is optimized for an X-ray energy of $23.0\ \text{keV}$. Finally, the analyzer grating G_2 also consists of gold with a height of $70.0\ \mu\text{m}$ and a period of $5.4\ \mu\text{m}$. For this measurement the three gratings are arranged in a symmetrical setup with inter-grating distances of $80\ \text{cm}$. The sample is placed $130\ \text{cm}$ in front of the source, the phase grating $7\ \text{cm}$ further down the beam. The distance between the sample and the detector is $94\ \text{cm}$. Due to the placement of all these components the beam characteristic is that of a minimal cone with a sample magnification of 1.72 . The detector is a single-photon counting PILATUS 100K with a pixel size of $172.0\ \mu\text{m} \times 172.0\ \mu\text{m}$ and an effective area of $487\ \text{px} \times 195\ \text{px}$. The field of view, however, is restricted to a width of $420\ \text{px}$ due to the limited size of the gratings. A not-to-scale, schematic view of the three-grating setup is given in figure 4.8.




Figure 4.8: Schematic view of the three-grating interferometer used for the tomographic measurements of the heart sample for the verification of the statistical reconstruction algorithm. The different parts are: T - X-ray tube; G_0 , G_1 , G_2 - source, phase and analyzer gratings; S - rotation stage with attached sample; D - X-ray detector.

The sample that is presented contains parts of a human heart—mostly fat and muscle tissue—fixated in formalin and put in a plastic cylinder. In addition, the cylinder contains a PMMA rod to be able to quantitatively calibrate the resulting reconstructed values. To avoid phase wrapping at the edges of the container it was placed in a water bath. Figure 4.9 gives an overview of the different parts of the sample. The X-ray tube was operated at a voltage of 40 kV and a current of 70 mA. For the tomographic measurement 1200 angular views were recorded with 11 individual raw images during the phase-stepping scan. For every 20 projections with sample in the beam a block of 5 flatfield projections was recorded. Each raw image acquisition was taken with an exposure time of 3.6 s.

The complete tomographic scan was subsequently processed with the SPR procedure (cf. section 3.2) extracting the differential phase signal and the corresponding statistical uncertainties. For the verification of the new reconstruction algorithm the data set was prepared according to the six cases presented in the introduction to this section and reconstructed on one hand with a conventional FBP with a Hilbert filter and on the other hand with the SIR algorithm as presented in section 4.2. In the next section both reconstructions are compared for each of these cases.




Figure 4.9: Overview of the different parts contained in the human heart sample.

4.3.2.2 Results

The results of this study are presented in figures 4.10 and 4.11. As in the section about the synthetic data, the figures are presented in three columns, with the left one showing the sinogram, on which the reconstructions are based and the middle and right column displaying the FBP and SIR reconstructions, respectively. A summary of the parameters for the SIR reconstructions is given in table 4.3.

The reconstruction results of the full scan is shown in row A. The mean photon count per pixel and acquisition is around 2300, so it can be considered free of noise for all practical purposes. With a detector width of 420 px and the number of projections of 1200 it is also more than fully sampled and aliasing or undersampling artifacts are not expected to show up in the FBP reconstruction. A look at the figure confirms this. When comparing this reconstruction to the one performed with the SIR algorithm, however, it becomes clear that by using the new algorithm there is still some room for improvement. The iterative reconstruction appears much sharper than the FBP reconstruction. This demonstrates that even under ideal experimental conditions the prerequisites for the FBP can never be fully met.

The next case—shown in row B—incorporates a measurement with reduced statistics. To be precise, the exposure time of each acquisition was reduced from 3.6 s to 0.144 s, which amounts to 1/25 of the exposure time used in the fully sampled scan. Accordingly the mean number of photons detected per

pixel and acquisition is decreased to around 95. This decrease in statistics leads to a substantial amount of noise in the FBP reconstruction, obstructing most of the fine details. Especially in the fat tissue the distinction between noise and actual features is not clear. The SIR algorithm, however, is able to reduce the noise in the flat regions, while keeping the features intact. This works due to the combination of the statistical uncertainties in the likelihood term, which identifies pixels that most probably contain noise, and the edge preserving Huber regularization with carefully chosen parameters. Comparing the SIR reconstruction in row B with the one from the full scan in row A we see that apart from a slight reduction in overall contrast most of the fine details match between the two images.

For the case of stationary gaps the fraction of almost 30% of pixels with no information is kept the same as for the synthetic data in the previous section. Due to the larger detector width the gaps were set to 8px and placed every 19px. The distribution of the missing data in the sinogram, as well as the two reconstructions are shown in row C. The observation we can make from the FBP reconstruction is the same as for the synthetic data. The gaps lead to ring and streak artifacts, although not as pronounced as for the phantom. This can be explained by the distribution of image values. While for the phantom different δ values are concentrated in objects with very sharp edges, they are much more evenly distributed in the heart sample with much less sharp edges. Still, the artifacts make the identification of small features a lot harder than in the full scan. In addition, a look at the fat tissue suggests that the ring artifacts destroy the overall texture of the flat regions. All of these effects are mitigated by using the SIR algorithm. In a visual comparison between the SIR reconstructions from the sinogram with gaps and from the full scan no significant differences become apparent.

Table 4.3: This table gives the reconstruction parameters for the algorithm verification with experimental data.

Case #	# iterations	λ_H	γ_H
A	50	$7.0 \cdot 10^{-4}$	$1.5 \cdot 10^{-3}$
B	50	$2.0 \cdot 10^{-4}$	$3.0 \cdot 10^{-4}$
C	50	$1.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$
D	150	$1.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$
E	30	$1.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$
F	100	$1.0 \cdot 10^{-5}$	$1.0 \cdot 10^{-4}$




Figure 4.10: First part of the results from the experimental evaluation of the algorithm. The left column shows the sinogram, the middle and right columns the reconstructions with FBP and SIR, respectively. A: full dataset; B: noisy data; C: stationary gaps. The second part is shown in figure 4.11.

Row D presents the results of a scan with missing blocks of projections. In contrast to the case of the synthetic data, where blocks spanning 15° were missing, we removed two blocks from 45° to 75° and 225° to 255° from the original sinogram. This means that we doubled the fraction of pixels with no information. We did this in order to enhance the effect this has on the reconstruction and to make the resulting artifacts even more visible. While in the phantom case these tomosynthesis-like artifacts appeared as streaks in the direction of the missing projections originating from the different insets, the artifacts in the heart sample reconstruction with the FBP are even more similar to those encountered in tomosynthesis due to the increased angular range of the missing projections. Looking at the FBP reconstruction we can see that the artifacts, instead of just streaks, appear as horseshoe shapes originating from higher contrast, small features. In addition, there is strong shadowing around the PMMA rod, which represents the object with the highest contrast and the largest size in the reconstruction. In the SIR reconstruction there are still some distortions from the missing data but the overall quality of the image is significantly improved. Most of the horseshoe artifacts have been removed with only some of them still visible in the center of the image, where fat and muscle tissue meet. Also parts of the plastic container are still blurry and were not faithfully reconstructed. The shadows around the PMMA rod, however, have been perfectly removed. In summary, the SIR reconstruction represents a more useful image for detecting small features than the FBP reconstruction.

The last two cases are again scans with a reduced number of angular views. Row E gives the results of a scan with an undersampling factor of around 4.4, which is obtained by reducing the number of projections from the original 1200 to 150 at a detector width of 420 px. For a full sampling the minimum number of projections would be about 660. The sinogram in the figure is again just stretched to be aligned with the reconstructions. The FBP reconstruction exhibits the typical undersampling artifacts, which manifest as simple radial streaks. These, of course, overlay the original structure of the sample making the analysis of the image more difficult. In this case, however, the artifacts are not yet severe enough to significantly impair the reconstruction. Still, the result obtained with the SIR algorithm is much closer to the original reconstruction. As in the section on synthetic data a sinogram with the original dimensions was created and evenly filled with the remaining projections before the actual reconstruction. Again, this leaves enough room for the algorithm to fill the remaining pixels with meaningful data, resulting in a smooth, artifact-free image.

For the last case in row F the number of projections was further reduced to 75, giving an approximate undersampling factor of 8.8. With this amount of missing data the artifacts in the FBP reconstruction become more severe and start to look like noise. This leads to small features tending to get lost within this artificial texture. The SIR algorithm is able to remove the texture but in the process also starts to blur out some of the low-contrast features. Especially in this case the choice of the parameters for the Huber regularization plays an important role. It represents a trade-off between keeping the features, which will also retain artifacts, and removing the artifacts completely while sacrificing small features. From table 4.3 it is apparent that the parameters had to be set very differently from the ones in the other cases. In the light of this the summary is that this is probably the most challenging of all cases for the presented SIR algorithm.




Figure 4.11: Second part of the results from the experimental evaluation of the algorithm. The left column shows the sinogram, the middle and right columns the reconstructions with FBP and SIR, respectively. D: missing blocks of projections; E: short scan (150 out of 1200 projections); F: ultrashort scan (75 out of 1200 projections). The first part is shown in figure 4.10.

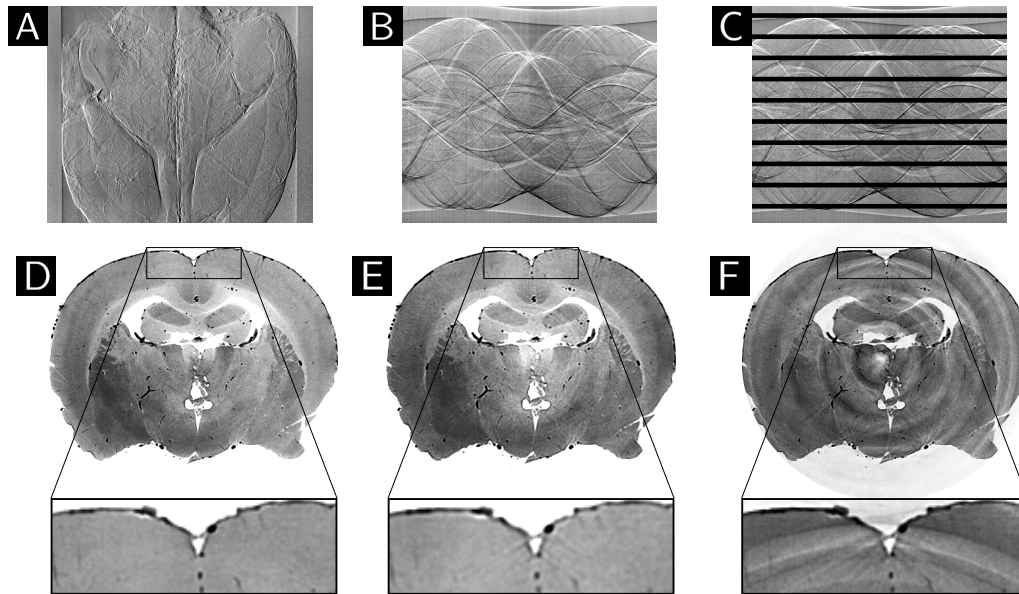


Figure 4.12: Tomography of a rat brain measured at a synchrotron. (A) DPC projection; (B) full DPC sinogram; (C) corrupt DPC sinogram with stationary gaps; (D) SIR reconstruction of the full sinogram; (E) SIR reconstruction of the sinogram with gaps; (F) FBP reconstruction of the same corrupt sinogram.

4.3.3 Synchrotron data

An additional validation of the algorithm is shown in figure 4.12 for a data set collected at a synchrotron facility. The reduction of artifacts arising from stationary gaps (labeled C in the previous sections) is demonstrated on the scan of a rat brain that was measured at the ESRF. For the tomographic scan 600 projections were acquired. In addition 25 blank scans each before and after the main projection block were recorded. The phase stepping was done in 8 steps over 2 grating periods. One exemplary DPC projection after the statistical phase retrieval is shown in panel (A) of the figure. Panel (B) presents one complete sinogram, while panel (C) demonstrates the placement of the stationary gaps. For reference, panel (D) illustrates the result of applying the SIR algorithm on the full data set. Finally, panel (E) and (F) show the reconstructions from the sinogram with gaps using SIR and FBP, respectively. The enlarged area of the FBP reconstruction clearly shows the severity of the artifacts originating from the gaps, manifesting as rings and small streaks. These artifacts are almost completely eliminated in the SIR reconstruction.




Figure 4.13: Plots of the likelihood, i.e. the value of the objective function, with respect to the iteration number for the six reconstructions from the synthetic data. Each graph is labeled according to the lettering in the previous section.

4.3.4 Further analysis

In this section we will analyze some more properties of the algorithm.

The first part deals with the rate of convergence. The convergence properties of the conjugate gradient algorithm have been thoroughly examined and discussed before (Shewchuk, 1994), so a mathematical analysis is omitted here. Instead we show how long it takes the SIR algorithm to reach a convincing reconstruction for each of the test cases from the previous section. The results are shown in figure 4.13. Visualized in the subfigures is the likelihood at each iteration. The likelihood is calculated by evaluating the complete objective function (4.26) including the regularization with the parameters given in table 4.2. The absolute value of the likelihood is not relevant—it depends strongly on the statistical weights and the regularization parameters—so the ordinate axis is labelled in arbitrary units. The curves for the first five cases (A–E) have the same overall shape, strongly decreasing for the initial 5–10 iterations with a decreasing slope afterwards. It is obvious that for different problems the number of iterations needed to obtain a good reconstruction varies between the cases.

Especially difficult to reconstruct is case (C), i.e. the detector gaps, because the gaps are positioned in a way to remove the information on the boundary between plastic container and formalin. The short scan (E) takes even longer to reach convergence due to the large amount of missing data (75% of data missing). For case (F), where even more data is missing (around

90% of the complete data), the likelihood curve takes on a completely different shape. The curve strongly decreases for approximately 60–70 iterations before it becomes more flat. It even increases slightly after iteration 150. According to the definition of the CG the value of the objective function should decrease at every iteration, but only if numerical errors are not taken into account. These errors can for example occur due to round off errors in the computer implementation or the inexact implementation of the forward and back projector. In addition, the presented algorithm is used on partly non-quadratic problems. Therefore a NLCG is used. It contains as one step a one-dimensional line search. This line search is another source of numerical inaccuracies.

The same plots of the likelihood at each iteration are shown for the experimental test cases in figure 4.14. The convergence behavior of (A) and (B) are almost identical to the corresponding cases with the synthetic data with convergence reached at about 10 and 30 iterations, respectively. We can conclude that for fully sampled data sets the specific structure of the reconstructed object does not play a significant role. The first difference in convergence speed between synthetic and experimental data occurs for the reconstruction from projection data with stationary gaps (case (D)). While it took the algorithm over 100 iterations to get to the final reconstruction for the synthetic data, the convergence behavior for the experimental data is not significantly different from the full or the noisy scan. We can explain this by the placement of the gaps on the sinogram. In the simulated sinogram gaps were located directly on the boundary between plastic container and formalin filling, making the reconstruction of the low frequency components inside the container very difficult. The algorithm has to be iterated many times before the estimate of the solution is visually satisfactory. The heart sample, however, was placed slightly off of the center of rotation, which leads to the container-formalin boundary to become visible under the gap for several projection angles. The algorithm has more intact data to work with and converges faster in this case.

While in the gap case the algorithm converged faster for the experimental data because there was more data available, the opposite is true for the reconstruction from the sinogram with missing angles (D). The size of the blocks of missing angles has been increased to 30° compared to the 15° blocks used for the simulated data. This increase in missing data is reflected in an increase in the number of iterations needed to reach convergence. The simulated data reconstruction is converged after about 20 iterations but the experimental reconstruction takes roughly 80 iterations.




Figure 4.14: Plots of the likelihood, i.e. the value of the objective function, with respect to the iteration number for the six reconstructions from the experimental data. Each graph is labeled according to the lettering in the previous section.

For the last two cases (E) and (F)—the undersampled scans—the opposite is true. For the simulated data set the algorithm is iterated 5–10 times longer than in the fully sampled case. The first of the two experimental reconstructions on the other hand converges only marginally slower than the fully sampled reconstruction and for the second reconstruction the number of iterations is still on the same order of magnitude. This can be explained with the content and frequency spectrum of the two samples. The phantom contains much less high frequency components than the heart sample and the fact that for the DPC signal the high frequency components are reconstructed much faster than the low frequencies (Köhler et al., 2011) leads to a faster convergence for the experimental data.

In the next part we highlight case (B) from the previous section by evaluating how the noise power spectrum behaves for both the FBP and the SIR algorithm. More specifically we look at the frequency dependent signal-to-noise ratio (SNR). For this evaluation we utilize the reconstructions from the fully sampled sinogram as noise-free ground truth and the noisy reconstructions for estimating the noise power spectrum. The noise power spectrum itself is estimated by first subtracting the noise-free reconstruction from the noisy one and calculating the NPS each for the FBP and the SIR reconstruction by

$$\text{NPS} = |\mathcal{F}(\boldsymbol{\delta}_{\text{noisy}} - \boldsymbol{\delta}_{\text{noise-free}})|^2, \quad (4.40)$$

i.e. the squared absolute value of the two-dimensional Fourier transform of the difference. The NPS calculated this way is two-dimensional, which does




Figure 4.15: Radial signal-to-noise spectrum of the FBP and SIR reconstructions for the simulated data set. The details of how the spectrum is calculated are given in the text.

not lend itself well for a qualitative and intuitive evaluation. Therefore, this two-dimensional quantity is transformed to a one-dimensional one by taking the average over the values in azimuthal direction. The resulting quantity then represents the average power spectrum in the radial direction and is correspondingly called the radial power spectrum.

The same calculations are performed on the noise-free reconstructions alone to get the radial signal power spectrum (SPS), i.e. the power spectrum of the noise-free signal. Dividing the SPS by the NPS we can get information on the frequency dependent SNR. For the simulated data this quantity is illustrated in figure 4.15 and for the experimental data in figure 4.16. A SNR below 1 means that the noise is so high that the underlying signal is not or only barely visible and a SNR above 1 indicates a signal, which is not completely obstructed by noise and thus features can be clearly distinguished. In both figures $\text{SNR} = 1$ is indicated by the horizontal black line. The main message of this analysis is that by reconstructing with the SIR algorithm the noise in the images is reduced and the detectability of features in the high frequency range is significantly improved. A look at figures 4.15 and 4.16 confirms that indeed the SNR of the high frequencies is way below 1 for the FBP and around or even above 1 for the line calculated from the SIR reconstructions.

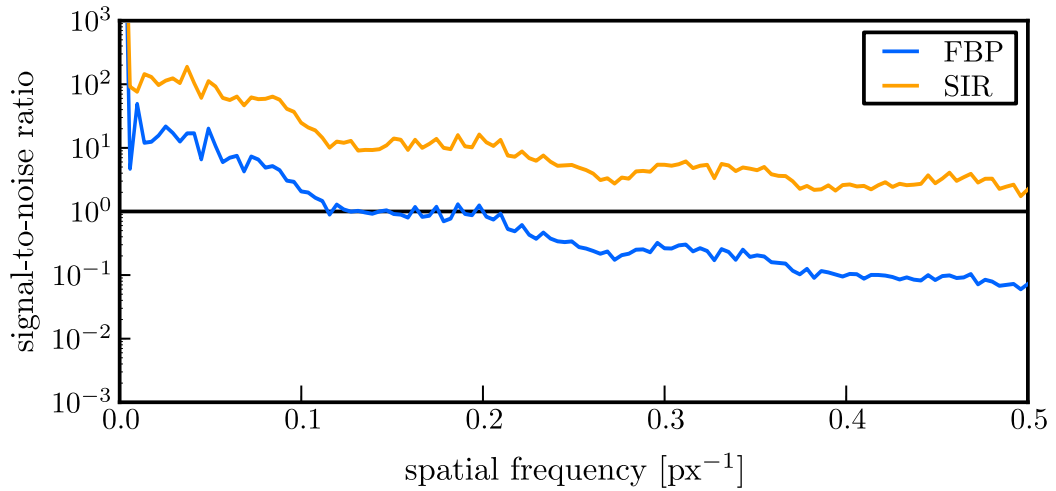


Figure 4.16: Radial signal-to-noise spectrum of the FBP and SIR reconstructions for the experimental data set. The details of how the spectrum is calculated are given in the text.

4.4 Moving towards correct modelling of the cost function

In section 4.2.1 it was stated that the basis for the development of the statistical model is the assumption that the measured and extracted DPC data is distributed following a Gaussian distribution. This, however, does not take into account the periodic nature of the data. We have seen that the phase gradient can only be measured in the interval $[-\pi, \pi)$. If a gradient exceeds this range it is wrapped back into it. Instead, Chabior (2011) proposes the so called wrapped normal distribution as correctly describing the DPC statistics, which is simply the normal distribution of a phase signal defined on the unit circle. In this section an objective function on the basis of this distribution in the log-likelihood formalism is derived. The probability density function (PDF) of the wrapped normal distribution can be written as

$$f(X|\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \sum_{n=-\infty}^{\infty} \exp\left(\frac{-(X - \mu + 2\pi n)^2}{2\sigma^2}\right), \quad (4.41)$$

where X is the random variable, μ the mean and σ the standard deviation. This definition means that the resulting distribution is composed of a sum of Gaussian distributions that are shifted by multiples of 2π around the mean. Figure 4.17 plots this distribution for different values of σ and $\mu = 0$. We can see that for small σ , i.e. high SNR, the function can be very well




Figure 4.17: Plot of the wrapped normal distribution (equation (4.41)) for different values of σ .

approximated by a simple Gaussian. The limit for large σ or low SNR, however, is the uniform distribution. This is also derived in Chabior (2011).

As this function is not easy to work with it is commonly approximated by the von-Mises distribution. It is defined for an angular random variable X as

$$f(X|\mu, \kappa) = \frac{\exp(\kappa \cos(X - \mu))}{2\pi J_0(\kappa)}, \quad (4.42)$$

where μ is again the mean, κ is equivalent to $1/\sigma^2$ in the normal distribution and J_0 is the zeroth order Bessel function of the first kind. This distribution defines the probability of measuring a realization of X on the unit circle when the mean and the variance-equivalent κ are given. In the likelihood formalism we seek the most likely values of these parameters given an already measured realization of X . The likelihood function for the von-Mises distribution is then written as

$$\mathcal{L}(\mu, \kappa|X) = \prod_{i=1}^n f(x_i|\mu, \kappa) = \prod_{i=1}^n \frac{\exp(\kappa \cos(x_i - \mu))}{2\pi J_0(\kappa)} \quad (4.43)$$

Now we can rewrite this equation in the notation already used in 4.2.1

$$\mathcal{L}(\boldsymbol{\delta}, \sigma | \boldsymbol{\alpha}) = \prod_{x, \theta} f(\alpha_x^\theta | \boldsymbol{\delta}, \sigma) \quad (4.44)$$

$$= \prod_{x, \theta} \frac{\exp\left(\frac{1}{(\sigma_x^\theta)^2} \cos\left(\alpha_x^\theta - \sum_j \partial_x A_{x,j}^\theta \delta_j\right)\right)}{2\pi J_0\left(\frac{1}{(\sigma_x^\theta)^2}\right)} \quad (4.45)$$

$$= \exp\left(\sum_{x, \theta} \frac{1}{(\sigma_x^\theta)^2} \cos(\alpha_x^\theta - \sum_j \partial_x A_{x,j}^\theta \delta_j)\right) \cdot \underbrace{\prod_{x, \theta} \frac{1}{2\pi J_0\left(\frac{1}{(\sigma_x^\theta)^2}\right)}}_{\text{constant}}, \quad (4.46)$$

where $\boldsymbol{\delta}$ is the vector of reconstructed values, $\boldsymbol{\alpha}$ the vector of measured projection values, x is the coordinate along the width of the detector and θ is the coordinate indicating the angular view in sinogram space. The second parameter κ of the von-Mises distribution has now been replaced by the inverse of the variance of the measurements $1/\sigma^2$. Of course we do not want to maximize the likelihood itself but rather its logarithm

$$l(\boldsymbol{\delta}, \sigma | \boldsymbol{\alpha}) = \sum_{x, \theta} \frac{1}{(\sigma_x^\theta)^2} \cos(\alpha_x^\theta - \sum_j \partial_x A_{x,j}^\theta \delta_j) \quad (4.47)$$

In its basic form this equation is very similar to equation (4.9) with the exception that the difference between measured data and the model is no longer squared but becomes the argument of a cosine function. This cosine function poses a problem regarding the optimization of the objective. While the cost function derived from the Gaussian distribution is squared and thus convex with a global optimum, equation (4.47) is periodic. What this means is that there is no longer a single optimum, but multiple ones. These multiple optima alone would not pose an unsolvable problem, one would just have to try to avoid local optima by, for example, a clever choice for the starting values. In the case of the cosine function, however, all optima have the same value, so it becomes a lot harder to find the correct one. Experiments with the non-linear conjugate gradient from section 4.2.3 show that this algorithm is not the ideal one to solve for this objective. Finding or developing an algorithm for this kind of problem, however, exceeds the scope of this thesis and is left for further research.

4.5 Summary

In this chapter the development and validation of a statistical iterative reconstruction algorithm for grating-based phase-contrast CT was presented. The goal of this algorithm was to be able to reconstruction PCCT data that is noisy or otherwise corrupt or incomplete. First, the physical model of the reconstruction problem was described in detail with an emphasis on the underlying statistical model. This statistical model was derived, starting with the line integrals for the total phase shift, resulting in an ill-posed minimization problem with an objective function that is called penalized weighted least-squares, penalized log-likelihood or posterior probability, depending on the formalism. We gave some examples of terms that can be used to constrain the solution space of this problem. In a next step the algorithm was validated first with the simulation of an analytical phantom and afterwards with real experimental data obtained at a laboratory-based grating interferometer. The validation was done on different scenarios of corrupt or missing data ranging from noisy projections to data sets with very few angular views. We found that in most simulated cases the algorithm performs as expected, giving well defined reconstructions. The more important finding, however, was that even for the experimental data set the SIR results were very close to the fully sampled, noise-free reconstruction for almost all of the different scenarios. Finally, the convergence and frequency dependent SNR analyses showed that the algorithm converges reasonably fast for all scenarios and both data sets and is able to improve the SNR especially for the high frequencies, where it is unsatisfying when analytical reconstruction algorithms are used. The chapter ended with an outlook to a more accurate design of the statistical model for the differential phase-contrast problem, which would require a totally different method for the minimization.

Chapter 5

Imaging application

In this chapter two individual applications of the algorithm developed in the previous chapter are presented. The first application is the reduction of artifacts caused by strongly absorbing and scattering objects (section 5.1). In section 5.1.3 BAR is demonstrated on a synchrotron scan of a mouse with bones. The main results of this section are currently under preparation for publication. In section 5.1.4 the algorithm is then applied to a laboratory-based scan of a carotid artery containing calcified plaque. The second application of the algorithm is presented in section 5.2 and deals with the reconstruction of the tomographic scan of a large object—the head of a pig.

5.1 Bone artifact reduction (BAR)

5.1.1 Motivation

Limited ability to record the full dynamic range of a signal is a common problem in imaging. The most prominent example is photography, where the camera sensor is often unable to capture the full contrast of a scene, especially when it contains very bright and very dark areas. This is usually overcome by taking multiple photos of the same scene with different exposure times and combining them with specialized software afterwards. This multiple-capture approach has an additional problem. The individual exposures are not perfectly registered with respect to each other.

The problem of limited dynamic range is, however, not restricted to visible light imaging. It also arises in X-ray imaging. Conventional X-ray absorp-




Figure 5.1: This figure visualizes the similarities between the artifacts created by metal implants in conventional CT (A) and those created by bones in grating-based phase-contrast CT (B). (Left image: courtesy of Dr. Peter B. Noël, Klinikum rechts der Isar)

tion computed tomography is routinely used in many fields to create three-dimensional maps of the linear attenuation coefficient μ —or the imaginary part of the complex index of refraction (Hounsfield, 1973). For example, in medical applications it can be used to retrieve three-dimensional information of the skeletal structure of a patient. However, problems arise when the patient’s body contains metal implants. The contrast between metal and bone is usually too high to capture information on both materials within the limited dynamic range of a conventional X-ray detector and clinically relevant measurement settings. In particular, with the settings of a CT scanner adjusted for good bone contrast, the absorption of metal implants is so high that virtually no photons will reach the detector behind the implants. This effect is called photon starvation and usually leads to streak artifacts in a tomographic reconstruction, commonly termed ‘metal artifacts’ in literature (see chapter 2 and De Man et al. (1999); Link et al. (2000); Lemmens et al. (2009); Meyer et al. (2009); Boas and Fleischmann (2011); Kachelriess et al. (2012); Meyer et al. (2012)). In addition, sufficient contrast in soft tissue is not easily achieved with conventional CT because of the small variations of the attenuation coefficients between the elements that make up soft tissue.

The part of the dynamic range, the soft-tissue contrast, is covered by phase-sensitive imaging methods, such as GBI. However, this high sensitivity to small density variations leads to a problem similar to the effects that metal introduces in conventional CT. This is mostly related to what is known as phase wrapping. Phase wrapping occurs in all phase-contrast techniques, as the quantity that is measured—the phase shift—is defined on the unit cir-

cle, i.e. in the interval $[0, 2\pi]$. If the phase of the X-rays is shifted by more than 2π , this value is wrapped into the original interval and thus leads to measurements that are no longer uniquely defined. This wrapping usually happens when the X-rays pass through dense objects, such as bones. This effect produces artifacts that look similar to those caused by metal in conventional CT and in analogy are termed 'bone artifacts'. Figure 5.1 shows a comparison of metal artifacts from a metallic hip implant in a conventional CT (A) and artifacts caused by bone material in the phase-contrast CT of the head of a mouse (B).

In this section we make use of the fact that grating interferometry provides three different, perfectly registered contrast signals. The SIR algorithm combines all these signals in order to improve the effective dynamic range of the final reconstructed volume with the aim to eliminate most of the artifacts originating from dense objects. In short, the soft-tissue parts provided by the phase-contrast signal are combined with information on the location of the dense parts, e.g. bones, taken from the absorption signal, complemented with statistical noise information, which is for the most part contained in the dark-field signal. In analogy to the reduction of metal artifacts, we call this method bone artifact reduction (BAR).

In the next two sections the capabilities of the BAR algorithm are demonstrated on two studies that present problems of preclinical relevance. The first study involves an ex-vivo scan of the abdominal region of a mouse measured at a synchrotron facility, where the spine causes severe artifacts. The subject of the second study is an excised carotid artery, which contains substantial amounts of calcified plaque. In contrast to the mouse, this sample was measured at a laboratory setup with a conventional rotating anode X-ray tube.

5.1.2 Bone artifacts: causes and reduction

The following three causes contribute to the appearance of bone artifacts:

1. The strong absorption in dense materials leads to photon starvation and loss of information.
2. Small-angle scattering inside dense, porous materials causes a loss of coherence and thus limits the ability to reliably determine the phase gradient.




Figure 5.2: Differential phase-contrast sinogram of a mouse (A). The strongest signal inside the container represents the spine, i.e. bone material. A part of the spine is shown enlarged in (B). The black and white pixels, which are visible inside the bone, are phase-wrapped.

3. The measurement of the phase gradient is intrinsically restricted to the interval $[0, 2\pi]$. If the gradient at a certain position exceeds this range, it will be wrapped into this interval, causing the value at this position to be undefined, thus lowering the statistical significance of this particular measurement. Phase wrapping usually occurs at strong edges, where refraction and thus the differential phase shift is especially high, e.g. at the boundary between soft tissue and bones. In combination with 1., also measurements close to the boundaries of the phase gradient interval can become wrapped due to higher noise fluctuations at positions where the count rate is lowered by strong absorption. This is called statistical phase wrapping (Chabior, 2011).

All of the above effects lead to a differential phase shift that is no longer uniquely defined at certain positions and thus does not represent reliable information for the tomographic reconstruction. An example of the visual appearance of phase wrapping is given in figure 5.2.

Up to this point the formulation of the SIR algorithm given in chapter 4 is very general. To apply it specifically on the reduction of bone artifacts, and especially on those caused by phase wrapping, additional steps have to be taken. The locations of pixels, which are wrapped with a high probability, have to be determined, with the goal of restricting their influence on the final reconstruction by modifying the statistical weights. This is accomplished by making use of the absorption signal, precisely delineating bones or other dense objects. The following steps are performed and illustrated in fig. 5.3:

1. reconstruction of absorption volume (fig. 5.3(A))




Figure 5.3: Steps for creating the mask m used to modify the statistical weights $1/\sigma_\phi^2$ in the cost function to alleviate the effects of phase-wrapping at the boundaries between bone and soft tissue. (A) absorption reconstruction serving as starting point; (B) zoom of red rectangle from (A) showing the soft-threshold segmentation of the bones; (C) the gradient magnitude of the segmented bones; (D) the resulting weights m/σ_ϕ^2 obtained after forward projection of the gradient volume and multiplication with the statistical weights.

2. thresholding of absorption volume to segment bones (fig. 5.3(B))
3. calculation of gradient magnitude of thresholded volume in three dimensions (fig. 5.3(C))
4. forward projection of gradient volume
5. multiplication of gradient projections onto statistical weights (fig. 5.3(D))

By thresholding the absorption reconstruction (step 2), only the information on the bones including their inner structure is preserved. The gradient calculation (step 3) is done to get the location of tissue-bone boundaries, where the phase gradient is known to be large. The forward projection operation of these boundaries is then used to pinpoint the location of most likely phase-wrapped pixels in projection space—the native space of the measurement data. Finally, the forward projection of the gradient volume is multiplied with the existing statistical weights.

Reconstructing with only the weighted least-squares term (see equation (4.9) in section 4.2.1) is, however, not sufficient. It represents an inverse problem with an infinite number of solutions. Choosing the desired solution is equivalent to regularizing or constraining the objective function. Most regularization terms are based on specific functions of neighboring voxels. A simple quadratic regularization, for example, is defined as the sum of quadratic differences between one voxel and all of its neighbors. The result is the penalty that is added to the objective function. In the optimization procedure this will lead to reduced value differences between adjacent voxels and thus helps to keep the resulting volume smooth. The same principle is used in edge-preserving regularization, where the quadratic function is replaced with a Huber potential function on the voxel differences. This function is piece-wise defined to be quadratic for differences below a choosable threshold, to further smooth out already flat regions, and linear above. In this work, the Huber term is supplied with a mask to be able to restrict its effect only to certain parts of the volume. For the treatment of the bones, a novel regularization term was designed. It is defined as the quadratic difference between a voxel of the phase reconstruction and the corresponding voxel of the absorption reconstruction, with the latter being scaled by a factor that can be chosen freely. This term forces the values of the bones in the phase reconstruction towards meaningful values, whereas before they were unreliable due to the ambiguity in the projection data, and effectively couples both signals. Just as the Huber term, the bone regularization contains a mask to restrict its effect to parts of the volume.

5.1.3 Reduction of bone artifacts on a whole mouse

Experimental setup

The abdominal region of a mouse cadaver—fixed in formalin and placed in a plastic container—was measured in a two-grating interferometer installed at beamline ID19 of the European Synchrotron Radiation Facility (ESRF) in Grenoble, France. The interferometer consists of a phase grating with a period of $4.78\ \mu\text{m}$ and an absorption grating with a period of $2.40\ \mu\text{m}$. The inter-grating distance of $48.1\ \text{cm}$ corresponds to the 9th Talbot order. The measurement was performed with monochromatic X-rays with an energy of $23\ \text{keV}$. The dataset was recorded in 902 projection views with four stepping images each. All images were recorded with the FReLoN camera, a scintillator-lens-coupled CCD, with an effective pixel size of $30.0\ \mu\text{m}$ and image dimensions of $1453 \times 433\ \text{pixel}^2$. Due to the nature of the detector, the raw

projection images were deconvolved before the signal extraction step to improve spatial resolution, as presented in section 3.3. Because the point spread function of the detector system was not exactly known, it was estimated as a two-dimensional Gaussian function with $\sigma_{x,y} = 1$ pixel. To prevent phase wrapping at the edges of the container during the measurement, it was placed in a water bath, that extended over the field of view. More information on this measurement can also be found in Tapfer et al. (2013).

The data set was subsequently reconstructed using both the standard filtered back projection algorithm with the Hilbert filter (Pfeiffer et al., 2007b), as well as the BAR algorithm. The statistical weights resulting from the least-squares processing step were modified using the procedure as explained in the last section and depicted in figure 5.3, using the absorption signal retrieved from the same measurement. The Huber and bone regularization each only worked on parts of the volume. The mask used to restrict the effects of these two regularizers was taken as the thresholded absorption reconstruction (figure 5.3 B). For the Huber regularizer the mask was inverted to only work on the soft-tissue regions outside of the bones and the original thresholded volume was used to restrict the bone regularizer to the voxels that contain bone material. The regularization parameters were chosen empirically, such that the result was visually and quantitatively most accurate. The strength of the Huber regularization was set to 0.5 with a threshold of 0.1, the strength for the bone regularizer was set to 10.0 with a factor of 1 and the values set to the absorption reconstruction. The results of both reconstructions were converted to units of the refractive index decrement δ by

$$\delta = \frac{p_2}{2\pi d} \rho + \delta_{H_2O,23keV}, \quad (5.1)$$

where the addition of the refractive index decrement of water takes into account the offset created by the water bath.

Results

Figure 5.4 depicts the results of reconstructions performed using the conventional filtered backprojection (FBP, left column) and the proposed bone artifact reduction method (BAR, right column) in the form of axial slices (A, B), as well as sagittal (C, D) and coronal (E, F) cuts through the center of the volume. All six images are windowed in the same range of $\delta = [4.067 \times 10^{-7}, 5.067 \times 10^{-7}]$. The FBP reconstruction exhibits strong streaking artifacts and shadowing around the bone in the axial view that manifest as horizontal lines in the sagittal view and a noise-like texture in




Figure 5.4: This figure shows the results of reconstructions of an ex-vivo mouse X-ray phase-contrast CT measurement using the conventional filtered backprojection (left) and the proposed BAR algorithm (right). Depicted are an exemplary axial slice (A, B) and sagittal (C, D) and coronal (E, F) cuts through the center of the volume. When comparing the filtered backprojection results with the BAR results, it is obvious that the strong streaking artifacts and the shadowing around the bone visible in the in the axial view of the FBP reconstruction (A) are clearly reduced in the BAR result (B). In the sagittal view (C) the streaking artifacts lead to horizontal lines which are strongest in the vicinity of the bones. Even relatively far away from the bone the artifacts affect the image quality, as can be seen in the coronal view (E). Both of these effects obstruct the underlying fine details, which become much clearer and easier to detect in the BAR reconstruction (D, F). Several regions of interest are marked with green rectangles in the left column images. There, a comparison of the standard deviation of both reconstructions is performed. The results are given in table 5.1.

the coronal view, obstructing most of the fine anatomical detail. Although there are still artifacts left in the reconstruction with the presented BAR

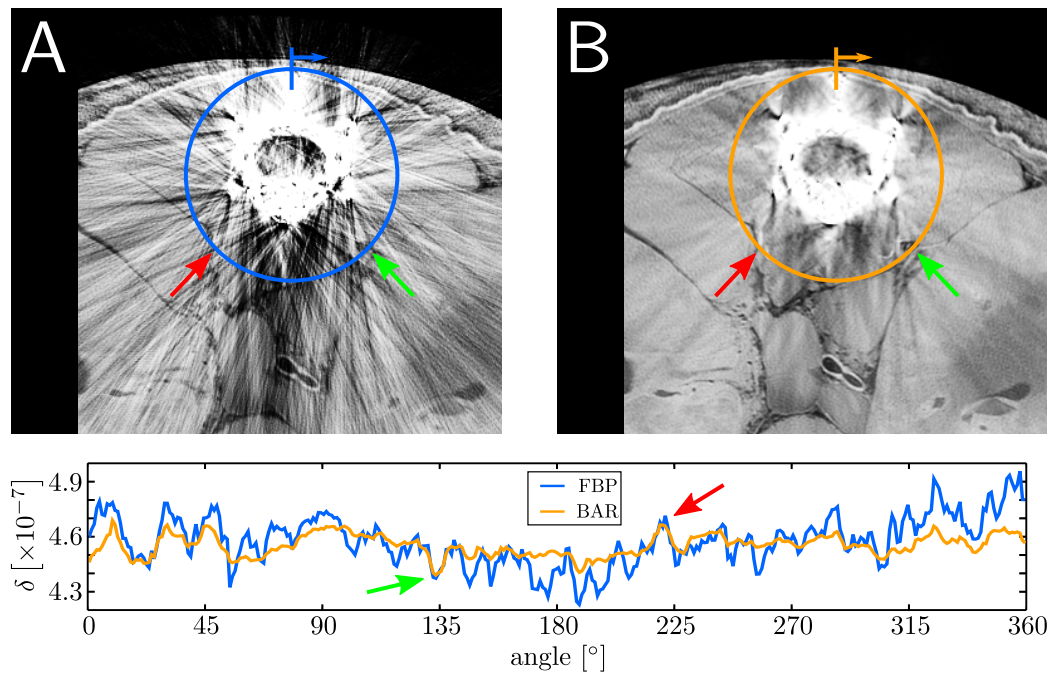


Figure 5.5: Analysis of streak artifact reduction. (A) Enlarged view of the spine in the axial slice of the FBP reconstruction; (B) the same image for the BAR result; the figure in the bottom shows ring plots along the circles drawn in the two axial views. The arrows indicate corresponding positions along the circle and in the line plot.

algorithm, their intensity and extension is drastically reduced, making the detection of small details a lot more feasible. This becomes most apparent in the sagittal and coronal views, where the images appear significantly clearer and almost free of artifacts. Figure 5.6 shows an enlarged view of the part of the sagittal cut, that is marked with a red, dashed rectangle in figure 5.4 (C). This makes the amount of artifact reduction and detail visibility of the BAR (B) compared to the FBP (A) even more obvious.

In order to quantify the reduction of artifacts, figure 5.5 presents an enlarged version of the axial slice from figure 5.4 showing the region around the spine that is effected the most by the streak artifacts. Quantitative values are plotted in the bottom of figure 5.5 along the circles drawn in blue (FBP, A) and orange (BAR, B), respectively. The 0° point is marked with a vertical bar and an arrow indicates the clockwise direction, in which the plot is drawn. The blue line corresponds to the FBP and the orange line to the BAR result. The standard deviation of these curves was determined as $1.35 \cdot 10^{-8}$ for FBP and $5.92 \cdot 10^{-9}$ for BAR, respectively, giving a good indication of the streak reduction. In addition, this plot shows the presented algorithm reduces the




Figure 5.6: Enlarged view of the part of the sagittal cut from figure 5.4(C) and (D), which is marked with the red dashed rectangle there. (A) FBP result; (B) BAR result.

streaks of the bone artifacts, while retaining the underlying details. Examples for these details are marked in the axial slice, as well as in the line plot. The green arrow, marking the point at around 135° on the circle, point to a feature that is darker than the surrounding tissue. In the line plot both the blue and the orange curve show this feature. At around 225° the red arrow points to a feature that is brighter than the surrounding tissue, which is also clearly visible as a spike of both curves in the line plot. In the other parts of the line plot it is obvious that the FBP curve contains many more and stronger variations than the BAR curve.

Furthermore, several homogeneous regions of interest (ROI) in different parts of the mouse and also in different views are studied with respect to their standard deviation to determine the degree of artifact reduction. They are marked with green rectangles in the left column images of figure 5.4. The results are summarized in table 5.1 and show, that a reduction of the artifacts by factors of 1.5–2 is achieved.

In a last figure, 5.7, we show the resulting sinogram after the reconstruction. This sinogram belongs to the same slice as the one shown in figure 5.2 and demonstrates that the phase wrapping has been completely removed by the BAR reconstruction.

Table 5.1: Results of the standard deviation analysis of several regions of interest in figure 5.4

region of interest	FBP	BAR	$\sigma_{\text{FBP}}/\sigma_{\text{BAR}}$
axial (ROI 1)	$5.01 \cdot 10^{-9}$	$3.78 \cdot 10^{-9}$	1.33
axial (ROI 2)	$2.27 \cdot 10^{-9}$	$1.56 \cdot 10^{-9}$	1.46
sagittal (ROI 3)	$2.42 \cdot 10^{-9}$	$1.16 \cdot 10^{-9}$	2.09
coronal (ROI 4)	$4.32 \cdot 10^{-9}$	$3.02 \cdot 10^{-9}$	1.43




Figure 5.7: The same DPC sinogram as shown in figure 5.2, but after finishing the BAR reconstruction, showing that the phase wrapping has been completely removed.

5.1.4 Reduction of artifacts caused by calcified plaque in a carotid artery

In this section the BAR algorithm is demonstrated on a different data set measured at a laboratory X-ray source. The sample under investigation is an excised carotid artery that contains large amounts of calcified plaque. This sample is different from the mouse of the previous section in that the soft-tissue parts contain much less contrast than the different organs in the mouse. In addition, the dense parts causing the bone artifacts—the calcified plaque—are much denser than the skeletal structure of the mouse and also more compactly distributed.




Figure 5.8: (A) Sinogram of the slice shown in the top row of figure 5.9. (B) Corresponding weights calculated from the statistical uncertainties.

Experimental setup

The artery was measured at a laboratory-based interferometer consisting of a conventional rotating anode X-ray source, a PILATUS 100K photon counting detector and three gratings as shown in figure 4.8 in the previous chapter. The source grating G_0 consists of gold bars with a period of $5.4\ \mu\text{m}$ and a height of $55\ \mu\text{m}$. The phase grating G_1 is made of nickel with a period of $5.4\ \mu\text{m}$ and a height of $8.5\ \mu\text{m}$, making it a π -shifting grating. Finally, the analyzer grating G_2 has the same specifications as the source grating. Taking into account the magnification of the setup of 1.72, the measurement was performed in the 6th Talbot order. The three gratings were arranged in a symmetric setup with inter-grating distances of 80 cm each. For the measurement the X-ray tube was operated at a voltage of 40 kV and a current of 70 mA with a molybdenum target, leading to an effective energy of about 27 keV. For the tomographic scan 1200 angular views over 360° were captured with 11 steps over 1 period per projection. The sensitive area of the detector covers $487 \times 195\ \text{px}^2$, but, due to the limited size of the gratings, the effective field of view is restricted to $380 \times 195\ \text{px}^2$. To prevent phase wrapping at the edge of the plastic container, the sample was measured in a 3.5 cm thick water bath.

For this study the data set is reconstructed three times. The first reconstruc-

tion uses the conventional FBP algorithm with the Hilber filter and serves as the reference result, which the other results are compared with. The second and third reconstructions are done with SIR and BAR, respectively. Before the statistical reconstructions—as for the mouse from the previous section—a mask \mathbf{m} for the regularization is created by thresholding a reconstructed absorption volume to separate the calcifications from the surrounding soft tissue and enabling the regularization to target both areas independently from each other. To soften the edges a little, the mask was filtered with a two-dimensional Gaussian with a σ of 2 px in both directions. Calculation of separate weights is not necessary for this data set, because the correct weighting is already achieved by the statistical uncertainties alone due to the strong concentration of calcified material in small volumes. The sinogram of one of the slices is shown in figure 5.8 (A), exhibiting strong phase wrapping in the parts, which contain the calcifications. Panel (B) of the same figure visualizes the corresponding statistical weights calculated from the phase variance. Dark pixels in the weight sinogram correspond to a lower weighting of the sinogram pixels in the data term of the reconstruction. It is obvious from a comparison of both images that all of the phase-wrapped pixels are weighted very low. The SIR reconstruction is performed without the use of the bone regularizer, solely using Huber regularization restricted with the mask \mathbf{m} to only affect the soft tissue parts surrounding the calcifications. The strength parameter is set to $\lambda_H = 0.001$ and the threshold to $\gamma_H = 0.005$. The algorithm is initialized with the FBP result and iterated 25 times. The second reconstruction uses the BAR algorithm, employing basically the same parameters with the differences that no mask is applied to the Huber regularization and the bone regularization is added with a strength of 100 and the absorption reconstruction as value. The additional factor of the term is empirically set to 8 for the best visual result.

Results

In this section the results of the three reconstructions are presented, compared and discussed.

A comparison of the FBP reconstruction and the SIR result is shown in figure 5.9. The top row contains the axial slice that corresponds to the sinogram in figure 5.8 with the FBP on the left (A) and SIR on the right (B). The two other rows show the transverse slices along the blue and orange line, respectively. The same colors are used to mark the corresponding slices in (C) and (E). The SIR slices along the same lines are given in (D) and (F).




Figure 5.9: Comparison of FBP and first SIR reconstructions of the artery with calcified plaque. (A) axial slice of the FBP volume; (B) corresponding axial slice of the SIR volume; (C) transverse slice through the FBP volume along the blue line as indicated in (A); (D) corresponding SIR slice; (E) another transverse slice through the FBP volume along the orange line as indicated in (A);(F) corresponding SIR slice.

Most apparent in the FBP reconstruction are the strong streak artifacts that obstruct most of the reconstruction. It is obvious that the strength of the streaks is larger in the direction along the elongated calcification. This effect happens, because in this direction the X-rays have to travel much longer through the dense material than in the other directions. Thus, they are

attenuated much more along their path and they are also scattered much stronger, leading to a lower signal on the detector for these rays. This in turn causes the variance of the phase signal to increase up to the point, where almost all pixels in the DPC signal are phase-wrapped. In addition to the streak artifacts we also observe dark shadows in front and behind the calcified parts, where the calcification is thickest. These shadows are also caused by the low signal that the detector receives. How strong and long-reaching the streak artifacts are is illustrated in the transverse slices. In (C) we can see a noise-like texture ranging over the complete width of the image. This texture is caused by the streak artifacts present in the corresponding slices. The image shown in (E) is even more illustrative. Here, no calcification is directly in the field of view, but still the extended streaks obstruct large parts of the soft tissue. A completely different picture is given in the right column. The SIR algorithm is able to significantly reduce the streak artifacts, revealing the soft-tissue structure that was formerly obstructed by the streaks. The conclusion from this comparison is that, for this kind of problem, the FBP is not suitable and produces unusable results and other algorithms have to be employed to get meaningful images.

The SIR results, however, are still not perfect. Despite nicely showing details in the soft-tissue in the regions far away from the calcifications, the dark shadows are still obstructing some of the features in the near vicinity of the plaque. In the following we concentrate on a direct comparison of the results obtained with the SIR and the BAR algorithms and highlight some of the areas, where one or the other algorithm produces more usable images.

Figure 5.10 shows the direct comparison of both statistical reconstructions in the axial view with (A) being the same image as in figure 5.9 with plain SIR and (B) representing the reconstruction with BAR. On the first glance both images appear very similar. The differences between them are visible only in small details. The orange and red arrows point to areas, where the SIR reconstruction exhibits dark shadows at the tips of the calcifications. The effect of the shadows is mitigated in the BAR reconstruction by applying a stricter constraint on the calcified regions. This constraint not only manifests in a reduction of the shadows, but overall the calcifications are much better delineated through the inclusion of the absorption signal in the phase reconstruction. The blue arrow points to a region, where this is most apparent. In the SIR reconstruction the bright area representing the large piece of calcification bleeds out into the surrounding soft tissue, decreasing the local contrast there. By using the BAR algorithm, the dark area between the fine bright line and the calcification itself is much better delineated. While the BAR algorithm is able to improve the image quality in small details right




Figure 5.10: Comparison of the axial slices reconstructed with SIR (A) and BAR (B) with the respective settings given in the text. The arrows point to the main differences between the images, which are discussed in the text.

next to the artifact-causing parts of the sample, the long streak artifacts in the outer parts of the reconstruction are a lot better controlled by just using the plain SIR algorithm.

This last statement is also clearly visualized in figure 5.11, which shows the transverse cut through the reconstructed volume corresponding to one in the second row of figure 5.9. The left image (A) is again the SIR reconstruction and the right image (B) depicts the BAR reconstruction. It contains a part of the artery that is heavily contaminated with calcifications. The dark shadows next to the calcified plaque is clearly not reduced by the SIR algorithm. The blue arrows point to two exemplary positions, where these shadows are very




Figure 5.11: Comparison of the first set of transverse slices reconstructed with SIR (A) and BAR (B) with the respective settings given in the text. The cut through the volume follows the blue line as indicated in figure 5.9 (A). The arrows point to the main differences between the images, which are discussed in the text.




Figure 5.12: Comparison of the second set of transverse slices reconstructed with SIR (A) and BAR (B) with the respective settings given in the text. The cut through the volume follows the orange line as indicated in figure 5.9 (A). The arrows point to the main differences between the images, which are discussed in the text.

strong. The left shadow is obviously caused by the calcification that lies in the same plane, whereas the right shadow stems from a calcification that itself cannot be seen in this cut. So, even in views that do not contain calcified plaque directly, the image quality can be impaired by the artifacts. The BAR algorithm, on the other hand, removes these shadows almost completely and is able to reveal much more detail, especially for the shadow on the right hand side. The orange arrow points to a location, where in fact the SIR reconstruction shows visibly better results than the one obtained with BAR. The long-range streak artifacts originating from the calcifications lead—in this view—to strong intensity variations over the whole width of the image in the BAR reconstruction. The SIR algorithm is able to suppress these artifacts almost completely.

As a last example we compare the SIR and BAR reconstructions on the same view as in the bottom row of figure 5.9 that does not contain any calcification. The comparison is shown in figure 5.12 with the SIR result on the left (A) and the BAR result on the right (B). As expected, in regions not directly in the vicinity of the calcifications, the SIR reconstruction provides better image quality than the BAR reconstruction. The orange arrow points at the most obvious location. In the BAR image the streak artifacts appear as a dark shadow obstructing the small corner of the vessel wall. Due to the reduction of the streaks by the SIR, this edge is much better delineated in the left image. The blue arrow points to a flat region outside the sample. In the left image this area is indeed flat, whereas in the BAR reconstruction we can see a texture that stems from streak artifacts that originate from out of plane. Lastly, the differences are visible in the general region, which the red arrow points to. Here we can also see the texture caused by the streak

artifacts overlaying large parts of the soft tissue in the BAR reconstruction. The SIR reconstruction on the left appears much clearer with small details being better visible.

From this study we can conclude that for these kinds of problems the FBP algorithm has limited use, due to the amount of missing and corrupt data. These limitations can be lifted by employing iterative reconstruction techniques. However, the final choice of which algorithm to use depends on the given task, i.e. which information we want to see in the reconstruction. Radiologists also face these kinds of task-dependencies in conventional CT, because state-of-the-art reconstruction programs today offer multiple protocols depending on, for example, whether the task is to image the skeletal structure of the patient or the diagnostic focus lies on the imaging of organs.

5.2 Reconstruction of large objects demonstrated on a sliced pig

This experiment is carried out to demonstrate that phase-contrast imaging with a grating interferometer in combination with SIR techniques is feasible even for large objects. For this purpose a slice of a frozen pig cadaver's head was measured at a synchrotron facility. In this section we apply the SIR algorithm to this challenging data set and compare the results to a state-of-the-art clinical scan.

5.2.1 Experimental setup

The frozen pig cadaver was cut into transverse slices with a height of about 5 cm. For the measurement a section of the head—containing a part of the brain—with a diameter of roughly 19 cm was used. This section was placed in a plastic container and fixated in formalin. This specific section was chosen to simulate a human head CT. A photograph of the prepared sample is shown in figure 5.13.




Figure 5.13: This figure shows a photograph of the 'Scheibenschwein' sample, a slice of a pig's head. The slice is inside a plexiglass container and fixated in formalin. The cylindrical part of the container has a diameter of roughly 19 cm and a height of about 5 cm.

The measurement was performed at beamline W2 of the storage ring DORIS III located at DESY, Hamburg and operated by Helmholtz-Zentrum Geesthacht (HZG). A schematic view—not to scale—of the measurement setup is shown in figure 5.14. The synchrotron X-ray source is represented by the tube marked 'T'. The beam characteristics of the source require a three-grating




Figure 5.14: Schematic view of the experimental setup used for measuring the sliced pig. The tube labelled 'T' represents the X-ray source (in this particular measurement a synchrotron beamline). The three gratings—source grating, phase grating and analyzer grating—are marked 'G0', 'G1' and 'G2', respectively. The label 'S' corresponds to the sample and sample holder, while 'D' marks the detector. The red arrow indicates the rotational movement for the tomographic acquisition and the green arrow the movement of the sample perpendicular to the beam direction to measure overlapping projection images of the sample.

interferometer to be used to get enough coherence for interference experiments. The source grating 'G0' has a period of $20.29\ \mu\text{m}$ and a height of $130\ \mu\text{m}$. The phase grating 'G1' is made of nickel with a period of $4.33\ \mu\text{m}$ and is $35\ \mu\text{m}$ high and the analyzer grating 'G2' is made of gold and has a period of $2.4\ \mu\text{m}$ and a height of $192\ \mu\text{m}$. The detection system—denoted 'D' in the figure—consists of a scintillating screen coupled with a CCD camera and a standard photographic lens to focus the visible light created in the scintillator on the sensor. The choice of focal length of the lens results in a field of view of about $14.42\ \mu\text{m} \times 14.42\ \mu\text{m}$. Due to this limited field of view, the whole sample had to be measured in parts while being moved perpendicular to the beam direction, creating 'panorama' projections. This is indicated by the green arrow in figure 5.14. The sample was mounted on a standard rotation stage, but, to prevent vibrations and bending of the setup during movement of the sample, a counterweight was installed to balance the perpendicular movement of the rotation stage.

The measurement was performed at an energy of $81\ \text{keV}$ for the beam to be




Figure 5.15: Two adjacent, overlapping DPC projections of the sliced pig sample. The vertical blue lines indicate the region of overlap in each of the projections. The orange lines mark one corresponding feature that can be used to align the two projections.

able to penetrate this large sample. Exposure times ranged from 0.9s for the blank scans and from 1.5s to 54s for the acquisitions with sample. The exposure times were varied between the perpendicular movements in order to keep mean count rates approximately equal as the thickness of the object varied. For each angular view a total of 16 overlapping individual images was recorded, to cover the complete width of the sample. Each individual acquisition consisted of a phase-stepping scan with 4 images each. The total number of angular views recorded for the tomographic scan was 240, covering the angular range from 0° to 180° .

5.2.2 Preprocessing and reconstruction results

All of the raw phase-stepping sequences are initially processed using the SPR to extract the differential-phase signal and the phase uncertainty to be later used in the reconstruction step.

The individual DPC projections belonging to a particular angular view are

then semi-manually stitched together. For the stitching the positions of the motor, which is responsible for the perpendicular movement of the sample, are used to calculate the corresponding pixel shift, by which the projections have to be moved to fit together. To reduce the effect of stitching seams the overlapping areas of neighboring projections are blended together by multiplying the overlapping areas with simple linear gradients in the interval $[0, 1]$. An example of two adjacent, overlapping DPC projections is given in figure 5.15. Projection (A) represents the left one, projection (B) is stitched to the right of (A). The vertical blue lines indicate the region, where both projections overlap and the orange line marks a distinct feature in both projections that can be used to align them. For these two projections the alignment is easy, even using automatic registration algorithms. For other pairs of projections, however, there are no such features present, making it very difficult to automatically align the projections using feature-based approaches. This is why we opted for using the motor positions for alignment. The angular views resulting from the stitching and blending are subsequently binned by a factor of 20, as the original angular views had a width of 21 591 px. There are two reasons for binning the projections: 1. reduce the amount of data that has to be handled by the reconstruction code; 2. reduce the undersampling (there are only 240 angular views) of the data set to a factor that is feasible for the SIR algorithm to manage.




Figure 5.16: Comparison of sliced pig reconstructions: (A) result of the FBP reconstruction from 240 projections; (B) result of the reconstruction with SIR of the same 240 projections.

The reconstruction itself is done on one hand with the FBP and the Hilber filter, on the other hand with SIR using the statistical weights from the signal extraction. The SIR algorithm is iterated 150 times with a Huber regularization. The Huber strength parameter is set to 2×10^{-5} and the




Figure 5.17: Comparison of sliced pig reconstructions: (A) previous SIR result; (B) resulting image when adding 30% of the FBP reconstruction to the SIR result.

threshold to 1×10^{-5} . In addition to the statistical weights another set of weights is created by masking out sinogram pixels that are close to π or $-\pi$ to filter unreliable pixels. Some of the flat-field images used to correct the signal projections were corrupted in some areas, leading to false values of pixels at those locations.

A comparison of both reconstructions is shown in figure 5.16, with the FBP result on the left (A) and the SIR result on the right (B). As expected, the FBP result exhibits very strong undersampling artifacts that make it hard to delineate the density differences in the soft tissue. These density differences are a lot better visible in the SIR reconstruction, the strong Huber regularization, however, smooths out most of the fine details. These details, especially the sharpness in the dense regions, e.g. the skeletal structure, are significantly better preserved in the FBP reconstruction. One way to combine the best of these advantages, i.e. fine details of the FBP and the strong soft-tissue delineation of the SIR, is to just mix the two reconstructions after the fact. Figure 5.17 compares the previous SIR result (A) to such a mixture, where roughly 30% of the FBP result have been added (B). The image of the mixture still clearly shows the variations in the soft tissue, while appearing sharper in the high-density parts.

For a final comparison the pig slice is also measured in a clinical CT scanner at Klinikum rechts der Isar. As this scan is performed with state-of-the-art soft- and hardware, we do not expect our results to come close to this reconstruction. The acquisition is done with a helical trajectory with the tube voltage set to 120 keV and a current of 404 mA in a manufacturer-specific

inner-ear protocol. The total exposure time of this scan is 1.437 s. The reconstruction of this data set is done with the manufacturer's implementation of the FBP algorithm.




Figure 5.18: Comparison of sliced pig reconstructions: (A) previous FBP-SIR mixture image; (B) result from a scan at a state-of-the-art clinical CT scanner.

Figure 5.18 shows the reconstruction of the clinical scan next to our FBP-SIR-mixture. As expected, when we compare our reconstruction to the clinical one, we can see that, even though the clinical scan only shows the absorption signal, image quality and soft-tissue delineation is a lot better than in our reconstruction. The main reason is, of course, the different sampling. While the clinical scan is fully sampled with about 2300 angular views for a detector width of 1024, the synchrotron scan, even with a 20-fold binning, is undersampled by a factor of 6.96. There are, however, also some grating-interferometry specific reasons, the first being that the scan was performed at a very low interferometer visibility of around 0.1, which is far from the maximum visibility* of 1.0. The visibility can be increased by improving grating fabrication, e.g. creating analyzer gratings with larger aspect ratios to increase absorption of the grating bars. The second reason is setup design. In a clinical CT scanner, all components are designed in a way to perfectly work together for a given scan geometry. Setups for grating interferometry, on the other hand, are for the most part optimized for performing experi-

*The maximum visibility is evident from

$$v = (I_{\max} - I_{\min}) / (I_{\max} + I_{\min}),$$

if $I_{\min} = 0$, i.e. the stepping curve goes to zero when the absorption grating bars cover the maxima of the interference pattern.

ments with different goals and geometries. This variability mostly comes at the cost of stability. In summary, there are currently many reasons, why clinical CT scanners outperform grating interferometry for the imaging of large objects. These reasons, however, are mainly of a technological nature, meaning that they can be solved by improving design and fabrication of grating interferometer components.

5.3 Summary

In this chapter two important applications were presented, where statistical iterative reconstruction techniques can significantly improve image quality of PCCT data sets. The first part demonstrated a new method for reducing the influence of strongly absorbing and scattering objects on the reconstruction of differential phase-contrast tomography data. By incorporating complementary information from the two other contrast channels available in grating interferometry, we were able to overcome the limited dynamic range inherently present in this measurement technique and considerably reduce the artifacts related to this effect. We presented the applicability of our bone artifact reduction algorithm in a difficult scenario with respect to the experimental parameters. This means, that both the energy of the measurement and the chosen Talbot distance contributed to stronger streak artifacts from the bones. The low energy leads to an increased attenuation, which in turn increased the information loss in the bones. The high Talbot order—that is, a large distance between the phase and absorption grating—produced larger transverse shifts of the interference pattern and thus increased the probability of phase wrapping. We have shown, that even under these challenging circumstances a significant reduction of bone artifacts and a convincing increase in image quality is possible. We limited this study to the case of a bone concentrated in a small part of the reconstructed volume. Further studies will have to show if our approach is also valid for more challenging anatomical examples, such as the ribcage with multiple sources of artifacts or the skull for artifact-free visualization of the brain. In addition, we demonstrated the BAR algorithm on a laboratory measurement of a carotid artery with calcified plaque with similar results of artifact reduction.

The second application is mainly geared towards human application of the method. We investigated the imaging and reconstruction of large objects on a slice of a pig's head, simulating a human head CT. We could show that by using statistical reconstruction a significant improvement in soft-

tissue delineation is possible compared to conventional reconstruction techniques, even though current experimental setups—without further advances in technology—are still lacking the image quality achieved with today’s generation of specialized and highly developed CT scanners.

Chapter 6

Conclusions and outlook

In this last chapter we summarize the main results of this PhD thesis and give a brief outlook. The structure follows the division of the thesis in three main results parts.

Projection images

In the domain of single projection images we worked on three distinct topics: the comparison between differential phase-contrast projections with non-differential absorption projections, the signal extraction of the three contrast images from raw acquisitions and the investigation of the influence of the point spread function on the imaging chain up to CT reconstructions. With the exception of the first topic, these results were further used in the course of the development of the statistical reconstruction algorithm.

- We presented a method to directly compare absorption projections with their differential phase-contrast counterparts without the need for a costly numerical integration of the phase gradient by carefully calculating the derivative of the absorption signal instead. The comparison involves the contrast-to-noise ratio in the same region of interest in the absorption and the differential phase-contrast projection resulting in the definition of the **relative contrast gain** as the ratio of CNRs. Application of this analysis tool on two distinct regions in a mammographic projection of a human breast resulted in a five- to ten-fold gain in relative contrast of the phase over the absorption projection.
- On the topic of signal extraction we introduced an alternative method to the well established Fourier analysis. Instead, using a weighted least-squares approach to directly fit a model to the measured stepping curves allows for a simple inclusion of the raw data's counting

statistics, resulting in the full covariance matrix in addition to the extracted contrast signals. Error propagation then gives access to the statistical uncertainties of absorption, differential phase-contrast and dark-field projections. Another advantage of this method is the possibility of extending the model for, e.g., future experimental setups with even better coherence and further developed gratings, to include higher order terms than the normally used first-order cosine.

- As a last part we briefly investigated the influence, which the setup's point spread function has on the projection data, as well as on CT reconstructions. We found that for a specific synchrotron setup the spatial resolution of acquired raw images is decreased, but can be recovered by deconvolution. Even though the point spread function was not known in advance, assuming a two-dimensional Gaussian led to a significant increase in image quality and sharpness in the raw stepping projections. We also compared untouched and deconvolved projections after the signal extraction and CT reconstructions from untouched and deconvolved data, with the result that the increase in spatial resolution and sharpness propagates through the processing and reconstruction. An estimate of the influence of the point spread function for setups at laboratory sources yielded the insight that, even if photon-counting detectors with a box-like point spread function are used, the finite size of the source's focal spot still leads to a degradation of the recorded images.

Statistical reconstruction algorithm

In this central part of the thesis we developed a statistical iterative reconstruction algorithm for differential phase-contrast computed tomography. The underlying statistical model was derived from the maximum-a-posteriori principle from Bayesian statistics based on the assumption of Gaussian distributed object values and takes into account the differential nature of the projections. This approach led to the formulation of a weighted least-squares objective function, which is iteratively minimized. For an appropriate weighting of the measurement values, the results from the alternative signal extraction procedure were used. This allowed for reducing the influence of measurements with a high statistical uncertainty on the reconstructed image. The objective function was complemented with prior knowledge about the reconstructed object in the form of regularization terms, two of which were developed during this thesis. An evaluation of the algorithm was not only performed with phantom simulations, but also with experimental data obtained at a laboratory setup and showed a significant reduction of noise as well as artifacts arising from different scenarios, where parts of the projection

information was missing. During this thesis we always assumed projections to be acquired with monochromatic X-rays, but further work will have to be put into taking into account polychromatic effects. Especially going into the direction of dual energy or even spectral phase-contrast CT is considered a viable option.

Bone artifact reduction

We showed for the first time how artifacts arising from phase wrapping at the interface between soft tissue and bone, so called bone artifacts, can be significantly reduced. The statistical reconstruction algorithm was augmented with information from the absorption signal. As the absorption projections are perfectly registered to the differential phase projections, we could use this information to extract the location of the edges of bones and decrease their influence on the final reconstruction. Together with a novel regularization term, which inserts absorption information about the bone structure back into the phase reconstruction, we were able to decrease the strength and extent of the bone artifacts. We showed that this method works on monochromatic synchrotron measurements, but also on measurements acquired at a conventional laboratory source.

Overall, we believe that the advances in data processing and in statistical reconstruction of differential phase-contrast projections worked out within this thesis will allow for a broad range of new applications of grating-based phase-contrast CT. This reconstruction method can easily be adapted for any kind of scan geometry, including cone beam acquisitions to allow for compact setups or helical geometries to speed up scan times. Together with the given results in noise reduction in combination with a reduction of undersampling artifacts the feasibility of dose compatible full body CT with a grating interferometer is greatly increased, by lifting the requirements on exposure of individual acquisitions and on angular sampling. Furthermore, by solving the problem of bone artifacts we pave the way for increased diagnostic value of full body phase-contrast CT reconstructions with the ultimate goal of bringing grating-based X-ray imaging to clinical application.

Bibliography

- J. Als-Nielsen and D. McMorrow. *Elements of Modern X-ray Physics*. John Wiley & Sons, Ltd, West Sussex, 2nd edition, 2011. [14, 17, 18]
- Andersen and Kak. A superior implementation of the ART algorithm. *Ultrasonic Imaging*, 94:81–94, 1984. [33]
- M. Bayes and M. Price. An Essay towards Solving a Problem in the Doctrine of Chances. By the Late Rev. Mr. Bayes, F. R. S. Communicated by Mr. Price, in a Letter to John Canton, A. M. F. R. S. *Philosophical Transactions of the Royal Society of London*, 53:370–418, Jan. 1763. [65]
- M. Bech. *X-ray imaging with a grating interferometer*. PhD thesis, University of Copenhagen, 2009. [17, 21, 31, 41]
- M. Beister, D. Kolditz, and W. A. Kalender. Iterative reconstruction methods in X-ray CT. *Physica medica : PM : an international journal devoted to the applications of physics to medicine and biology : official journal of the Italian Association of Biomedical Physics (AIFB)*, 28(2):94–108, Apr. 2012. [10]
- F. E. Boas and D. Fleischmann. Evaluation of two iterative techniques for reducing metal artifacts in computed tomography. *Radiology*, 259(3):894–902, June 2011. [98]
- M. Chabior. *Contributions to the characterization of grating-based x-ray phase-contrast imaging*. Phd thesis, TU Dresden, 2011. [93, 94, 100]
- C. T. Chantler. Theoretical form factor, attenuation, and scattering tabulation for $Z=1-92$ from $E=1-10$ eV to $E=0.4-1.0$ MeV. *Journal of Physical and Chemical Reference Data*, 24:71, 1995. [17]
- G.-H. Chen, J. Zambelli, K. Li, N. Bevins, and Z. Qi. Scaling law for noise variance and spatial resolution in differential phase contrast computed tomography. *Medical Physics*, 38(2):584, 2011. [36]

- T. J. Davis, D. Gao, T. E. Gureyev, A. W. Stevenson, and S. W. Wilkins. Phase-contrast imaging of weakly absorbing materials using hard X-rays. *Nature*, 373(6515):595–598, 1995. [9]
- B. De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. Metal streak artifacts in X-ray computed tomography: a simulation study. *IEEE Transactions on Nuclear Science*, 46(3):691–696, June 1999. [98]
- B. DeMan. *Iterative Reconstruction for Reduction of Metal Artifacts in Computed Tomography*. PhD thesis, 2001. [10]
- K. J. Engel, D. Geller, T. Köhler, G. Martens, S. Schusser, G. Vogtmeier, and E. Rössl. Contrast-to-noise in X-ray differential phase contrast imaging. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 648:202–207, 2010. [36, 43]
- A. Fehringer, T. Lasser, I. Zanette, P. B. Noël, and F. Pfeiffer. A versatile tomographic forward- and backprojection approach on Multi-GPUs. In *SPIE Medical Imaging*, page accepted for poster presentation, 2013. [60]
- J. A. Fessler. *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis*. SPIE, Bellingham, 2000. [10, 34, 64, 66]
- P. Gilbert. Iterative methods for the three-dimensional reconstruction of an object from projections. *Journal of Theoretical Biology*, 36(1):105–117, 1972. [33]
- R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Number 3 in Texts in Computer Science. Pearson Prentice Hall, Upper Saddle River, New Jersey, 3rd edition, 2008. [39]
- J. Goodman. *Introduction to Fourier Optics*. Roberts and Company Publishers, 3rd edition, 2004. [18]
- R. Gordon, R. Bender, and G. T. Herman. Algebraic Reconstruction Techniques (ART) for three-dimensional electron microscopy and X-ray photography. *Journal of Theoretical Biology*, 29(3):471–481, 1970. [32]
- D. Hahn, P. Thibault, M. Bech, M. Stockmar, S. Schleede, I. Zanette, A. Rack, T. Weitkamp, A. Sztrókey, T. Schlossbauer, F. Bamberg, M. Reiser, and F. Pfeiffer. Numerical comparison of X-ray differential phase contrast and attenuation contrast. *Biomedical optics express*, 3(6):105–108, 2012. [10]

- J. Herzen, T. Donath, F. Pfeiffer, O. Bunk, C. Padeste, F. Beckmann, A. Schreyer, and C. David. Quantitative phase-contrast tomography of a liquid phantom using a conventional x-ray tube source. *Optics Express*, 17(12):10010–10018, May 2009. [10, 36]
- G. N. Hounsfield. Computerized transverse axial scanning (tomography): Part 1. Description of system. *British Journal of Radiology*, 46(552):1016–1022, Dec. 1973. [9, 98]
- J. Hsieh. *Computed Tomography - Principles, Design, Artifacts, and Recent Advances*. John Wiley & Sons, Inc., Hoboken, New Jersey, 2nd edition, 2009. [10, 38]
- P. J. Huber. Robust Estimation of a Location Parameter. *Annals of Statistics*, 53:73–101, 1964. [67]
- V. N. Ingal and E. A. Beliaevskaya. X-ray plane-wave topography observation of the phase contrast from a non-crystalline object, 1999. *Journal of Physics D: Applied Physics*, 28(11):2314–2317, 1995. [9]
- P. M. Joseph. An Improved Algorithm for Reprojecting Rays through Pixel Images. *IEEE transactions on medical imaging*, 1(3):192–6, Jan. 1982. [59]
- M. Kachelriess, E. Meyer, and R. Raupach. METHOD, COMPUTING UNIT, CT SYSTEM AND C-ARM SYSTEM FOR REDUCING METAL ARTIFACTS IN CT IMAGE DATASETS, July 2012. URL <http://www.freepatentsonline.com/y2013/0039556.html>. [98]
- S. Kaczmarz. Angen{ä}herte aufl{ö}sung von systemen linearer gleichungen. *Bulletin International de l'Academie Polonaise des Sciences et des Lettres*, 35:355–357, 1937. [32]
- A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. IEEE Press, New York, 1988. [10, 29, 31, 73, 78]
- J. Kenntner. *Herstellung von Gitterstrukturen mit Aspektverhältnis 100 für die Phasenkontrastbildgebung in einem Talbot-Interferometer*. Phd thesis, Karlsruhe Institut für Technologie, 2012. [21]
- T. Köhler, K. Jürgen Engel, and E. Roessl. Noise properties of grating-based x-ray phase contrast computed tomography. *Medical physics*, 38 Suppl 1 (7):S106, July 2011. [69, 91]

- C. Kottler, C. David, F. Pfeiffer, and O. Bunk. A two-directional approach for grating based differential phase contrast imaging using hard x-rays. *Optics express*, 15(3):1175–81, Feb. 2007. [36, 37]
- T. Kohler, B. Brendel, and E. Roessl. Iterative reconstruction for differential phase contrast imaging using spherically symmetric basis functions. *Medical Physics*, 38(8):4542, July 2011. [10]
- J.-C. Labiche, J. S. Puchades, D. v. B. And, and J. Moy. No Title. *ESRF Newsletter*, 25:41–43, 1996. [49]
- J.-C. Labiche, O. Mathon, S. Pascarelli, M. A. Newton, G. G. Ferre, C. Curfs, G. Vaughan, A. Homs, and D. F. Carreiras. Invited article: the fast readout low noise camera as a versatile x-ray detector for time resolved dispersive extended x-ray absorption fine structure and diffraction studies of dynamic problems in materials science, chemistry, and catalysis. *The Review of scientific instruments*, 78(9):091301, Sept. 2007. [49]
- P. S. Laplace. Memoir on the Probability of the Causes of Events (1774). *Statistical Science*, 1(3):364–378, Aug. 1986. [65]
- C. Lemmens, D. Faul, and J. Nuyts. Suppression of metal artifacts in CT using a reconstruction procedure that combines MAP and projection completion. *IEEE transactions on medical imaging*, 28(2):250–60, Feb. 2009. [98]
- R. M. Lewitt. Multidimensional digital image representations using generalized Kaiser-Bessel window functions. *Journal of the Optical Society of America A*, 7(10):1834, Oct. 1990. [59]
- T. M. Link, W. Berning, S. Scherf, U. Joosten, A. Joist, K. Engelke, and H. E. Daldrup-Link. CT of metal implants: reduction of artifacts using an extended CT scale technique. *Journal of computer assisted tomography*, 24(1):165–72, 2000. [98]
- L. B. Lucy. An iterative technique for the rectification of observed distributions. *The Astronomical Journal*, 79:745, June 1974. [49]
- A. Malecki, G. Potdevin, and F. Pfeiffer. Quantitative wave-optical numerical analysis of the dark-field signal in grating-based X-ray interferometry. *EPL (Europhysics Letters)*, 99(4):48001, 2012. [74]
- S. Matej and R. M. Lewitt. Practical Considerations for 3-D Image Reconstruction Using Spherically Symmetric Volume Elements. *IEEE Transactions on Medical Imaging*, 15(1):68–78, 1996. [59]

- E. Meyer, F. Bergner, R. Raupach, T. Flohr, and M. Kachelriess. Normalized metal artifact reduction (NMAR) in computed tomography. In *Nuclear Science Symposium Conference Record (NSS/MIC), 2009 IEEE*, pages 3251–3255, 2009. [98]
- E. Meyer, R. Raupach, M. Lell, B. Schmidt, and M. Kachelrieß. Frequency split metal artifact reduction (FSMAR) in computed tomography. *Medical physics*, 39(4):1904–16, Apr. 2012. [98]
- A. Momose. Phase-sensitive imaging and phase tomography using X-ray interferometers. *Optics Express*, 11(19):2303–2314, Sept. 2003. [10]
- K. Mueller. *Fast and accurate three-dimensional reconstruction from cone-beam projection data using algebraic methods*. Phd, Ohio State University, 1998. [29, 31]
- K. Mueller, R. Yagel, and J. J. Wheller. A Fast and Accurate Projection Algorithm for 3D Cone-Beam Reconstruction with the Algebraic Reconstruction Technique (ART). *Proceedings of SPIE*, 3336(614):1–9, 1998. [59]
- J. Nuyts, B. D. Man, J. A. Fessler, W. Zbijewski, and F. J. Beekman. Modelling the physics in the iterative reconstruction for transmission computed tomography. *Physics in Medicine and Biology*, 58(12):R63, 2013. [58]
- A. V. Oppenheim and R. W. Schaffer. *Discrete-Time Signal Processing*. Prentice-Hall, Inc., Upper Saddle River, New Jersey, 2nd edition, 1999. [38]
- D. M. Paganin. *Coherent X-Ray Optics*. John Wiley & Sons, Ltd, 2006. [16, 18]
- F. Pfeiffer, T. Weitkamp, O. Bunk, and C. David. Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources. *Nature Physics*, 2(4):258–261, Mar. 2006. [10, 22]
- F. Pfeiffer, O. Bunk, C. David, M. Bech, G. Le Duc, A. Bravin, and P. Cloetens. High-resolution brain tumor visualization using three-dimensional x-ray phase contrast tomography. *Physics in Medicine and Biology*, 52(23):6923–30, Dec. 2007a. [40]
- F. Pfeiffer, C. Kottler, O. Bunk, and C. David. Hard X-Ray Phase Tomography with Low-Brilliance Sources. *Physical Review Letters*, 98(10):108105–108108, Mar. 2007b. [10, 103]

- F. Pfeiffer, M. Bech, O. Bunk, P. Kraft, E. F. Eikenberry, C. Brönnimann, C. Grünzweig, and C. David. Hard-X-ray dark-field imaging using a grating interferometer. *Nature Materials*, 7(2):134–137, 2008. [25]
- J. Radon. On the determination of functions from their integrals along certain manifolds. *Berichte über die Verhandlungen der Sächsische Akademie der Wissenschaften*, 69:262–277, 1917. [27]
- W. H. Richardson. Bayesian-Based Iterative Method of Image Restoration. *Journal of the Optical Society of America*, 62(1):55, Jan. 1972. [49]
- W. C. Röntgen. Ueber eine neue Art von Strahlen. *Annalen der Physik*, 300(1):1–11, 1898. [9]
- S. Schleede. *X-ray Phase-Contrast Imaging at a Compact Laser-Driven Synchrotron Source*. Phd thesis, Technische Universität München, 2013. [17]
- S. Schleede, M. Bech, S. Grandl, A. Sztrókay, J. Herzen, D. Mayr, M. Stockmar, G. Potdevin, I. Zanette, A. Rack, and Others. X-ray phase-contrast tomosynthesis for improved breast tissue discrimination. *European Journal of Radiology*, 83(3):531–536, Mar. 2014. [10]
- G. Schulz, T. Weitkamp, I. Zanette, F. Pfeiffer, F. Beckmann, C. David, S. Rutishauser, E. Reznikova, and B. Müller. High-resolution tomographic imaging of a human cerebellum: comparison of absorption and grating-based phase contrast. *Journal of the Royal Society, Interface / the Royal Society*, 7(53):1665–76, Dec. 2010. [40]
- J. Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. 1994. [89]
- R. L. Siddon. Calculation of the radiological depth. *Medical Physics*, 12(1):84–87, 1984a. [59]
- R. L. Siddon. Fast calculation of the exact radiological path for a three-dimensional CT array. *Medical Physics*, 12(2):252–255, 1984b. [59]
- A. Snigirev, I. Snigireva, V. Kohn, S. Kuznetsov, and I. Schelokov. On the possibilities of x-ray phase contrast microimaging by coherent high-energy synchrotron radiation. *Review of Scientific Instruments*, 66(12):5486–5492, 1995. [9]
- A. Sztrókay, J. Herzen, S. D. Auweter, S. Liebhardt, D. Mayr, M. Willner, D. Hahn, I. Zanette, T. Weitkamp, K. Hellerhoff, and Others. Assessment

- of grating-based X-ray phase-contrast CT for differentiation of invasive ductal carcinoma and ductal carcinoma in situ in an experimental ex vivo set-up. *European Radiology*, 23(2):381–387, 2013. [10]
- H. F. Talbot. LXXVI. Facts relating to optical science No. IV. *Philosophical Magazine Series 3*, 9:401, 1836. [18]
- A. Tapfer, M. Bech, A. Velroyen, J. Meiser, J. Mohr, M. Walter, J. Schulz, B. Pauwels, P. Bruyndonckx, X. Liu, A. Sasov, and F. Pfeiffer. Experimental results from a preclinical X-ray phase-contrast CT scanner. *Proceedings of the National Academy of Sciences of the United States of America*, 109(39):15691–6, Sept. 2012. [10]
- A. Tapfer, R. Braren, M. Bech, M. Willner, I. Zanette, T. Weitkamp, M. Trajkovic-Arsic, J. T. Siveke, M. Settles, M. Aichler, A. Walch, and F. Pfeiffer. X-Ray Phase-Contrast CT of a Pancreatic Ductal Adenocarcinoma Mouse Model. *PLoS ONE*, 8(3):e58439, Mar. 2013. [10, 103]
- T. Thuering, P. Modregger, B. R. Pinzer, Z. Wang, and M. Stampanoni. Non-linear regularized phase retrieval for unidirectional X-ray differential phase contrast radiography. *Optics Express*, 19(25):25545–25558, 2011. [36, 37]
- T. Weitkamp, A. Diaz, C. David, F. Pfeiffer, M. Stampanoni, P. Cloetens, and E. Ziegler. X-ray phase imaging with a grating interferometer. *Optics Express*, 13(16):6296–6304, 2005. [10, 40, 41]
- T. Weitkamp, C. David, C. Kottler, O. Bunk, and F. Pfeiffer. Tomography with grating interferometers at low-brilliance sources. *Proceedings of SPIE*, (6318):63180S–63180S–10, 2006. [18, 22]
- T. Weitkamp, I. Zanette, C. David, J. Baruchel, M. Bech, P. Bernard, H. Deyhle, T. Donath, J. Kenntner, S. Lang, J. Mohr, B. Müller, F. Pfeiffer, E. Reznikova, S. Rutishauser, G. Schulz, A. Tapfer, and J.-P. Valade. Recent developments in x-ray Talbot interferometry at ESRF-ID19. *Proceedings of SPIE*, 7804:780406–1–780406–10, 2010. [41]
- P. Willmott. *An introduction to synchrotron radiation: techniques and applications*. John Wiley & Sons, Ltd, West Sussex, 2011. [15, 17]
- M. Willner, M. Bech, J. Herzen, I. Zanette, D. Hahn, J. Kenntner, J. Mohr, A. Rack, T. Weitkamp, and F. Pfeiffer. Quantitative X-ray phase-contrast

- computed tomography at 82 keV. *Optics Express*, 21(4):4155, Feb. 2013. [10]
- Q. Xu, E. Y. Sidky, X. Pan, M. Stampanoni, P. Modregger, and M. A. Anastasio. Investigation of discrete imaging models and iterative image reconstruction in differential X-ray phase-contrast tomography. *Optics Express*, 20(10):10724, Apr. 2012. [10]
- J. Zambelli, N. Bevins, Z. Qi, and G.-H. Chen. Radiation dose efficiency comparison between differential phase contrast CT and conventional absorption CT. *Medical Physics*, 37(6):2473, 2010. [36]
- I. Zanette. *Interférométrie X à réseaux pour imagerie et lanalyse de front d'ondes au synchrotron*. Phd thesis, Université de Grenoble, 2011. [49]
- I. Zanette, T. Weitkamp, T. Donath, S. Rutishauser, and C. David. Two-Dimensional X-Ray Grating Interferometer. *Phys. Rev. Lett.*, 105:248102, 2010. [36, 37]
- A. Ziegler, T. Kohler, T. Nielsen, and R. Proksa. Efficient projection and backprojection scheme for spherically symmetric basis functions in divergent beam geometry. *Medical Physics*, 33(12):4653, 2006. [59]

Appendix A

Derivations and implementations

A.1 Implementation of the statistical phase retrieval

Following the derivation of the solution using normal equations for the WLS problem discussed in section 3.2 we now concretize the steps for the statistical phase retrieval. We start the derivation from the explicit WLS cost function for each pixel in each projection

$$S = \sum_{i=0}^{N-1} w_i \cdot (I_i - A_0 - A_1 \cdot \cos(x_i) - B_1 \cdot \sin(x_i))^2, \quad (\text{A.1})$$

where N is the total number of measurements on the stepping curve, I_i is the measured intensity at grating position x_i and $w_i = 1/\sigma_i^2$ represents the statistical weighting. $x_i = 2\pi ip/N$ is the position of the grating, where p is the number of periods, over which the stepping was performed.

The minimization of equation (A.1) involves setting the partial derivatives with respect to A_0 , A_1 and B_1 to zero. This yields a system of three equations for the three parameters. For matters of presentation the following

abbreviations are introduced:

$$\begin{aligned}
S &= \sum_i w_i & S_c &= \sum_i w_i \cos(x_i) & S_s &= \sum_i w_i \sin(x_i) \\
S_{cc} &= \sum_i w_i \cos^2(x_i) & S_{ss} &= \sum_i w_i \sin^2(x_i) & S_{cs} &= \sum_i w_i \cos(x_i) \sin(x_i) \\
S_I &= \sum_i w_i I_i & S_{cI} &= \sum_i w_i I_i \cos(x_i) & S_{sI} &= \sum_i w_i I_i \sin(x_i)
\end{aligned}$$

In matrix notation the system of equations looks like

$$\begin{pmatrix} S & S_c & S_s \\ S_c & S_{cc} & S_{cs} \\ S_s & S_{cs} & S_{ss} \end{pmatrix} \cdot \begin{pmatrix} A_0 \\ A_1 \\ B_1 \end{pmatrix} = \begin{pmatrix} S_I \\ S_{cI} \\ S_{sI} \end{pmatrix}$$

If the weights were equal for all pixels, i.e. not taking statistical properties into account, the coefficient matrix only depended on the grating positions and thus solving this system would only involve a single matrix inversion and multiplications of the inverse matrix with each pixel's solution vector. In the case of statistical phase retrieval, however, a matrix inversion has to be computed for each pixel, making an analytical expression more favorable over a numerical solution. Calculating the inverse of the coefficient matrix analytically yields

$$\begin{aligned}
\Delta &= \det(A) = (S_{cc}S_{ss} - S_{cs}^2) \cdot S + 2S_{cs}S_cS_s - S_c^2S_{ss} - S_{cc}S_s^2, \\
A^{-1} &= \frac{1}{\Delta} \cdot \begin{pmatrix} (S_{cc}S_{ss} - S_{cs}^2) & (S_{cs}S_s - S_cS_{ss}) & (S_{cs}S_c - S_{cc}S_s) \\ (S_{cs}S_s - S_cS_{ss}) & (SS_{ss} - S_s^2) & (S_cS_s - SS_{cs}) \\ (S_{cs}S_c - S_{cc}S_s) & (S_cS_s - S_{cs}S) & (S_{cc}S - S_c^2) \end{pmatrix},
\end{aligned}$$

leading to the following solution for the three fit parameters

$$\begin{aligned}
A_0 &= \frac{(S_{cc}S_{ss} - S_{cs}^2) \cdot S_I + (S_{cs}S_c - S_{cc}S_s) \cdot S_{sI} + (S_{cs}S_s - S_cS_{ss}) \cdot S_{cI}}{\Delta}, \\
A_1 &= \frac{(S_cS_{ss} - S_{cs}S_s) \cdot S_I + (S_{cs}S - S_cS_s) \cdot S_{sI} + (S_s^2 - S_{ss}S) \cdot S_{cI}}{\Delta}, \\
B_1 &= \frac{(S_{cs}S_c - S_{cc}S_s) \cdot S_I + (S_{cc}S - S_c^2) \cdot S_{sI} + (S_cS_s - S_{cs}S) \cdot S_{cI}}{\Delta},
\end{aligned}$$

their variances

$$\sigma_{A_0}^2 = \frac{S_{cc}S_{ss} - S_{cs}^2}{\Delta}, \quad \sigma_{A_1}^2 = \frac{S_{ss}S - S_s^2}{\Delta}, \quad \sigma_{B_1}^2 = \frac{S_{cc}S - S_c^2}{\Delta}$$

and covariances

$$\begin{aligned} COV(A_0, A_1) &= \frac{S_{cs}S_s - S_cS_{ss}}{\Delta} \\ COV(A_0, B_1) &= \frac{S_{cs}S_c - S_{cc}S_s}{\Delta} \\ COV(A_1, B_1) &= \frac{S_cS_s - S_{cs}S}{\Delta} \end{aligned}$$

These expressions now only consist of multiplications and summations, as the grating positions are the same for all pixels, allowing for a precalculation of the cosines and sines.

A.2 Derivation of Richardson-Lucy deconvolution

The Richardson-Lucy deconvolution algorithm is a maximum likelihood method based on Poisson statistics, which aims to recover the true object that has been blurred by a known PSF. In this appendix we derive the update equation used in section 3.3.

First, we define

- O** true object
- I** measured image
- P** a priori known PSF.

The maximum likelihood method is explained in section 4.2.1. The statistical model of choice in this case is the Poisson distribution, because the deconvolution is supposed to be applied to raw detector acquisitions that usually obey counting statistics. The maximization problem can then be formulated as

$$\hat{O} = \arg \max_O \mathcal{L}(O|I), \quad (\text{A.2})$$

i.e. the goal is to find the object O that maximizes the likelihood of O being the true object given the measured image I .

The Poisson likelihood function is given as

$$\mathcal{L}(O|I) = \prod_x \frac{e^{-M_x} M_x^{I_x}}{I_x!}, \quad (\text{A.3})$$

where M_x is the x^{th} value in the forward model vector, which we define as

$$M_x = (P * O)_x = \sum_{x'} P_{x-x'} \cdot O_{x'}, \quad (\text{A.4})$$

where x' represents the coordinate of the PSF vector. We can omit the factorial term in equation (A.3), because it evaluates to a constant that only changes the value of the maximum but not its location. Its logarithm is then

$$l(O|I) = \sum_x (M_x - I_x \ln M_x). \quad (\text{A.5})$$

Finding the maximum of this equation amounts to setting its gradient with respect to the object vector O to zero. The gradient of the likelihood with respect to each entry in the object vector is

$$\frac{\partial l(O|I)}{\partial O_x} = \sum_{x'} \frac{\partial l}{\partial M_{x'}} \frac{\partial M_{x'}}{\partial O_x} \quad (\text{A.6})$$

and

$$\frac{\partial M_{x'}}{\partial O_x} = \frac{\partial}{\partial O_x} \sum_{x''} P_{x'-x''} O_{x''} = \sum_{x''} P_{x'-x''} \overbrace{\frac{\partial O_{x''}}{\partial O_x}}^{\delta_{x'',x}} = P_{x'-x} = \bar{P}_{x-x'}. \quad (\text{A.7})$$

This result means that the PSF is traversed in reverse order, when used in a convolution. Inserting it in equation (A.6) gives

$$\frac{\partial l(O|I)}{\partial O_x} = \sum_{x'} \frac{\partial l}{\partial M_{x'}} \bar{P}_{x-x'} = \bar{P} * \frac{\partial l}{\partial M}. \quad (\text{A.8})$$

The last piece of equation (A.6) is the evaluation of the gradient of the likelihood with respect to the forward model M , which gives

$$\frac{\partial l}{\partial M_x} = 1 - \frac{I_x}{M_x}, \quad (\text{A.9})$$

so the final formulation of equation (A.6) is

$$\frac{\partial l(O|I)}{\partial O_x} = \bar{P} * \left(1 - \frac{I_x}{M_x} \right). \quad (\text{A.10})$$

The final formulation of the update equation requires the use of the prior knowledge that the true object is always positive, $O_x \geq 0$. To implement

this prior, we assume the object to be the result of an exponential function, so we substitute $O_x = \exp(u_x)$ and reevaluate the gradient with respect to u_x

$$\frac{\partial l}{\partial u_x} = \frac{\partial l}{\partial O_x} \frac{\partial O_x}{\partial u_x} = O_x \cdot \bar{P} * \left(1 - \frac{I_x}{M_x}\right). \quad (\text{A.11})$$

Setting this gradient to zero and reformulating slightly gives

$$O_x = O_x \cdot \left(\bar{P} * \frac{I_x}{M_x}\right) = O_x \cdot \left(\bar{P} * \frac{I_x}{(P * O)_x}\right), \quad (\text{A.12})$$

which exactly represents the update equation used in section 3.3.

A.3 Gradient and step size calculation of regularization terms

For the sake of completeness the expressions for the gradient and the denominator of the step-size calculation for the regularization terms are stated here (cf. section 4.2.3).

The gradient of the quadratic regularization term (eq. (4.21)) is

$$\frac{\partial}{\partial \delta_j} R_Q(\boldsymbol{\delta}) = 4 \sum_{i \in \mathcal{N}_j} w_{ij} (\delta_j - \delta_i) \quad (\text{A.13})$$

and the denominator term

$$\frac{\partial^2}{\partial \kappa^2} R_Q(\boldsymbol{\delta} + \kappa \mathbf{d}) \Big|_{\kappa=0} = \sum_j \sum_{i \in \mathcal{N}_j} w_{ij} (d_j - d_i)^2 \quad (\text{A.14})$$

The gradient of the Huber term (eq. (4.23)) is

$$\frac{\partial}{\partial \delta_j} R_H(\boldsymbol{\delta}, \mathbf{m}, \gamma) = m_j \sum_{i \in \mathcal{N}_j} \begin{cases} \frac{(\delta_i - \delta_j)}{\gamma^2} & \text{for } |\delta_i - \delta_j| \leq \gamma \\ \frac{-\text{sgn}(\delta_i - \delta_j)}{\gamma} & \text{for } |\delta_i - \delta_j| > \gamma \end{cases}, \quad (\text{A.15})$$

and the denominator term

$$\frac{\partial^2}{\partial \kappa^2} R_H(\boldsymbol{\delta} + \kappa \mathbf{d}, \mathbf{m}, \gamma) \Big|_{\kappa=0} = \sum_j m_j \sum_{i \in \mathcal{N}_j} w_{ij} \begin{cases} \frac{(d_i - d_j)^2}{\gamma^2} & \text{for } |\delta_i - \delta_j| \leq \gamma \\ 0 & \text{for } |\delta_i - \delta_j| > \gamma \end{cases} \quad (\text{A.16})$$

The same expressions for the mean regularizer (4.24) are

$$\frac{\partial}{\partial \delta_j} R_M(\boldsymbol{\delta}, \sigma) = 2(\delta_j - \langle \delta_j \rangle_\sigma) \quad (\text{A.17})$$

and

$$\frac{\partial^2}{\partial \kappa^2} R_M(\boldsymbol{\delta} + \kappa \mathbf{d}, \sigma) \Big|_{\kappa=0} = \sum_j (d_j - \langle d_j \rangle_\sigma). \quad (\text{A.18})$$

And finally for the bone regularization term (4.25)

$$\frac{\partial}{\partial \delta_j} R_B(\boldsymbol{\delta}, \mathbf{a}, \mathbf{b}, c) = 2b_j (\delta_j - c \cdot a_j) \quad (\text{A.19})$$

and

$$\frac{\partial^2}{\partial \kappa^2} R_B(\boldsymbol{\delta} + \kappa \mathbf{d}, \mathbf{a}, \mathbf{b}, c) \Big|_{\kappa=0} = 2 \sum_j b_j^2 \cdot d_j^2 \quad (\text{A.20})$$

Publications and scientific presentations

Here, all publications and scientific presentations (oral and poster), which are related to the work of this thesis are listed in chronological order. The list of publications is categorized into first-authored and co-authored publications and subdivided into peer-reviewed publications and conference proceedings.

First-authored publications (peer-reviewed)

D. Hahn, P. Thibault, M. Bech, M. Stockmar, S. Schleede, I. Zanette, A. Rack, T. Weitkamp, A. Sztrókay, T. Schlossbauer, F. Bamberg, M. Reiser and F. Pfeiffer, "Numerical comparison of X-ray differential phase contrast and attenuation contrast", *Biom. Opt. Express*, 3(6):1141-1148, 2012.

D. Hahn, P. Thibault, A. Fehringer, M. Bech, T. Köhler, F. Pfeiffer and P.B. Noël, "Bone artifact reduction for X-ray differential phase-contrast CT", *submitted to PNAS*, 2014.

D. Hahn, P. Thibault, A. Fehringer, M. Bech, P.B. Noël and F. Pfeiffer, "Statistical iterative reconstruction for X-ray differential phase-contrast CT", *in preparation*, 2014.

Co-authored publications (peer-reviewed)

J. Fu, P. Li, Q. L. Wang, S. Y. Wang, M. Bech, A. Tapfer, D. Hahn, and F. Pfeiffer, "A reconstruction method for equidistant fan beam differential phase contrast computed tomography", *Phys. Med. Biol.*, 56(14):452938, 2011.

A. Sztrókay, J. Herzen, S. Auweter, S. Liebhardt, D. Mayr, M. Willner, D. Hahn, I. Zanette, T. Weitkamp, K. Hellerhoff, F. Pfeiffer, M. Reiser and F. Bamberg, Assessment of grating-based X-ray phase-contrast CT for differentiation of invasive ductal carcinoma and ductal carcinoma in situ in an experimental ex vivo set-up, *European Radiology*, 23(2):381-7, 2013.

P.B. Noël, J. Herzen, A.A. Fingerle, M. Willner, M.K. Stockmar, D. Hahn, M. Settles, E. Drecoll, I. Zanette, T. Weitkamp, E.J. Rummeny and F. Pfeiffer, Evaluation of the potential of phase-contrast computed tomography for improved visualization of cancerous human liver tissue, *Z. Med. Phys.*, 23(3):204-11, 2013.

F. Schwab, S. Schleede, D. Hahn, M. Bech, J. Herzen, S. Auweter, F. Bamberg, K. Achterhold, A.Ö. Yildirim, A. Bohla, O. Eickelberg, R. Loewen, M. Gifford, R. Ruth, M.F. Reiser, K. Nikolaou, F. Pfeiffer and F.G. Meinel, "Comparison of contrast-to-noise ratios of transmission and dark-field signal in grating-based X-ray imaging for healthy murine lung tissue", *Z. Med. Phys.*, 23(3):236-42, 2013.

M. Willner, M. Bech, J. Herzen, I. Zanette, D. Hahn, J. Kenntner, J. Mohr, A. Rack, T. Weitkamp and F. Pfeiffer, Quantitative X-ray phase-contrast computed tomography at 82 keV, *Opt. Express*, 21(4):4155-4166, 2013.

A.A. Fingerle, M. Willner, P.B. Noël, J. Herzen, D. Munzel, D. Hahn, E.J. Rummeny and F. Pfeiffer, Quantitative x-ray phase-contrast computed tomography of cystic kidney lesions a phantom study, *Radiology*, (accepted)

Conference proceedings

P.B. Noël, M. Willner, A. Fingerle, J. Herzen, D. Munzel, D. Hahn, E. Rummeny and F. Pfeiffer, Improved diagnostic differentiation of renal cystic lesions with phase-contrast computed tomography, SPIE Proceedings Vol. 8313, 2012.

D. Hahn, P. Thibault, A. Fehringer, M. Bech, P. Noël and F. Pfeiffer, Bone artifact reduction in differential phase-contrast CT, Proceedings of Fully3D 2013.

Oral presentations

Title: Direct signal-to-noise comparison of radiographic attenuation- and differential phase-contrast X-ray images, *Frühjahrstagung der Deutschen Physikalischen Gesellschaft*, Dresden, Germany, March 2011.

Title: Bone artifact reduction for differential phase-contrast tomography based on a statistical iterative reconstruction algorithm, *Annual meeting of the Radiological Society of North America (RSNA)*, Chicago IL, USA, November/December 2012.

Title: Improvements in data processing and iterative tomographic reconstruction in grating-based phase-contrast imaging, *SPIE Medical Imaging*, Lake Buena Vista FL, USA, February 2013.

Title: Bone artifact reduction in differential phase-contrast CT, *Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine (Fully3D)*, Lake Tahoe CA, USA, June 2013.

Poster presentations

Title: Numerical comparison of X-ray differential phase contrast and attenuation contrast, *International Symposium on BioMedical Applications of X-Ray Phase-Contrast Imaging (IMXP)*, Garmisch, Germany, January 2012.

Acknowledgements

PAPA
MAMA
LINUS
EMIL

(Actually typed by Linus)

Doing the work on this thesis would not have been possible without the help of many colleagues and other supporters. In particular I would like to express my gratitude to

- my thesis adviser Prof. Franz Pfeiffer. Franz, thank you for giving me the opportunity to work on this interesting and challenging project.
- my scientific supervisors, Pierre Thibault, Peter Noël and Martin Bech. Thank you for your patience in answering all of my many questions and for sharing your vast knowledge and your time in numerous discussions. I learned a lot from you guys.
- the members of the examination committee for this thesis, Prof. Jan Wilkens as a second referee and Prof. Martin Zacharias as chairman.
- my co-workers and especially my direct and virtual office mates Arne Tapfer, Astrid Velroyen, Marco Stockmar, Marian Willner, Simone Schleede, Michael Epple, Andreas Fehringer, Björn Enders, Martin Dierolf, Andreas Malecki and all of you other guys for making my time in the group a truly awesome experience. Without you my PhD would have been less than half the fun.
- coffee for just existing.
- the open source community for giving up their spare time to create the awesome tools that made work so much more efficient.
- the most important people in my life, my family for always being there for me. Especially my two wonderful sons Linus (who even directly contributed to this thesis as seen above) and Emil, who not always made work easier but life so much richer, and my lovely wife Sabine for always supporting me and trying to keep distractions off my back. I love you all.