Chair of Biomedical Physics
Department of Physics
Technical University of Munich

Model-Based Iterative Reconstruction for X-ray Computed Tomography Using Attenuation and Propagation-Based Phase Contrast

Lorenz Hehn

Dissertation

December 2019

Supervisors:
Prof. Dr. Franz Pfeiffer
Dr. Martin Dierolf

Technische Universität München
Physik Department
Lehrstuhl für Biomedizinische Physik

Model-Based Iterative Reconstruction for X-ray Computed Tomography Using Attenuation and Propagation-Based Phase Contrast

Lorenz Markus Tobias Hehn

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:	Prof. Dr. Martin Zacharias
Prüfer der Dissertation:	1. Prof. Dr. Franz Pfeiffer
	2. Priv.-Doz. Dr. Tobias Lasser

Die Dissertation wurde am 09.12.2019 bei der Technischen Universität München eingereicht und durch die Fakultät für Physik am 23.12.2019 angenommen.

Abstract

This work is focused on conventional and emerging methods and modalities of X-ray imaging, which can be implemented at laboratory environments, describing a number of new algorithms designed to improve image quality and reduce acquisition time and dose. These imaging modalities include conventional computed tomography (CT), which is solely sensitive to the attenuating properties of the sample. CT is widely used in the field of medical imaging and non-destructive testing. One focus of this thesis lies on high-resolution CT systems, which achieve spatial resolutions in the nanometer regime enabling three-dimensional visualization of small samples. Another imaging modality explored in this work is propagation-based phase-contrast imaging (PBI) in particular in combination with CT. PBI is sensitive to not only absorption, but also the phase-shifting properties of the sample. Employing this phase information yields significant advantages when imaging small or weakly attenuating samples or ones that consist of materials that have similar attenuation properties, such as biological structures made from soft-tissue.

First, reconstruction approaches for PBI and propagation-based phase-contrast computed tomography (PB-CT) are investigated, which combine a model of the image formation, a statistical description of the measurements and regularization techniques. Using a laboratory CT system as well as synchrotron radiation, it is demonstrated how these algorithms benefit image quality, in particular regarding artifacts arising from materials that do not fit within the assumptions imposed by these approaches.

Second, algorithms for PBI and PB-CT are developed that further account for the reduced spatial coherence and low flux found at laboratory sources, and the response of efficient scintillator-based detectors. Thereby, an approximate analytical solution of the proposed algorithm is derived for PBI, which generalizes the conventional single-material phase-retrieval algorithm and its extension. For PB-CT, a simulation study and an experimental study are conducted, which clearly show the improvements with respect to spatial resolution and the benefits related to an accurate modeling of the noise covariance statistics.

Third, in the context of conventional CT, an optimization approach is proposed, which enables the estimation of parameterized models of the system blur during tomographic reconstruction. This has the potential to replace laborious characterizations of all components that add to the system blur and to account for additional contributions such as interpolation artifacts from the projection operations. An extensive simulation study is performed and the feasibility of this approach is demonstrated experimentally at a test-bench setup.

Furthermore, means to estimate the positions of the X-ray source, rotation axis and detector in CT applications at the nanometer scale are investigated. Mechanical instabilities and drifts of different components limit the spatial resolution at this scale. As the geometry is non-deterministic, established calibration procedures cannot be employed. Thus, two optimization-based approaches are examined, which estimate the geometry parameters from the measurements alone. The respective methods are validated in a simulation study as well as on measured data revealing distinct improvements in resolution.

Finally, an outlook is presented, which discusses the potential of combining the different methods investigated in this work for various imaging scenarios.

Zusammenfassung

Die vorliegende Arbeit beschäftigt sich mit konventionellen und neuartigen Röntgenbildgebungsverfahren, die im Labor realisiert werden können. Im Speziellen wird eine Reihe neuer Algorithmen entwickelt, um die Bildqualität zu verbessern beziehungsweise die Aufnahmezeit zu reduzieren und Dosis einzusparen. Eines dieser Verfahren ist die konventionelle Röntgencomputertomographie (CT), die einzig auf die absorbierenden Eigenschaften der Probe sensitiv ist. CT ist in der medizinischen Bildgebung sowie im Bereich der zerstörungsfreien Materialprüfung weit verbreitet. Ein besonderes Augenmerk dieser Arbeit liegt auf hochauflösenden CT Systemen, die räumliche Auflösungen im Nanometerbereich erreichen und somit in der Lage sind, kleinste Proben dreidimensional darzustellen. Ein weiteres Verfahren ist die propagationsbasierte Phasenkontrastbildgebung (PBI) und ihre Kombination mit CT. PBI ist dabei zusätzlich sensitiv auf die phasenschiebenden Eigenschaften der Probe. Diese Phaseninformationen zu verwenden birgt erhebliche Vorteile, um kleine oder schwach absorbierende Proben darzustellen oder Proben, die aus Materialien zusammengesetzt sind, die ähnlich absorbierende Eigenschaften haben, wie beispielsweise biologisches Weichteilgewebe.

Als Erstes werden Rekonstruktionsmethoden für PBI und propagationsbasierte Phasenkontrastcomputertomographie (PB-CT) untersucht, welche ein Modell der Bildentstehung, eine statistische Beschreibung der Messungen und Regularisierungstechniken kombinieren. Anhand eines kompakten CT Systems sowie mit Synchrotronstrahlung wird gezeigt, wie diese Algorithmen die Bildqualität verbessern, besonders bezüglich Artefakten, die von Materialien ausgehen, welche die bei diesen Methoden gemachten Annahmen nicht erfüllen.

Als Zweites werden Algorithmen für PBI and PB-CT entwickelt, welche zusätzlich die reduzierte partielle Kohärenz und den niedrigen Fluss von Laborquellen sowie den Einfluss effizienter szintillationsbasierter Detektoren erfassen. Für PBI wird dabei eine genäherte analytische Lösung des entwickelten Algorithmus hergeleitet, die den konventionellen Phasenrekonstruktionsalgorithmus von Paganin und seine Erweiterung verallgemeinert. Zusätzlich werden eine Simulationsstudie sowie eine experimentelle Studie durchgeführt, welche die Verbesserungen der räumlicher Auflösung durch die vorgeschlagenen Ansätze aufzeigen sowie die Vorteile, die sich aus der präzisen Modellierung der Rauschkovarianzstatistik ergeben, nachweisen.

Als Drittes wird in Bezug auf die konventionelle CT ein Optimierungsansatz vorgeschlagen, der es ermöglicht, parameterisierte Modelle, welche die Gesamtverschmierung durch das System beschreiben, während der tomografischen Rekonstruktion zu schätzen. Dies hat das Potential, aufwändige Charakterisierungen aller Komponenten, die zur Systemverschmierung beitragen, zu ersetzen und dabei gleichzeitig zusätzliche Komponenten, wie beispielsweise Interpolationsartefakte der Projektionsoperationen, zu erfassen. Eine umfangreiche Simulationsstudie wird durchgeführt und die Realisierbarkeit des Ansatzes wird experimentell anhand eines Testaufbaus demonstriert.

Des Weiteren werden Möglichkeiten untersucht, für CT Anwendungen im Nanometerbereich die Positionen der Röntgenquelle, der Rotationsachse und des Detektors zu schätzen. Mechanische Instabilitäten und Drifts der verschiedenen Komponenten limitieren die

räumliche Auflösung auf dieser Größenskala. Da die Geometrie nicht reproduzierbar ist, können außerdem keine etablierten Kalibrierungsverfahren verwendet werden. Aus diesem Grund werden zwei Optimierungsansätze untersucht, welche die Geometrie aus den Messungen selbst schätzen. Die jeweiligen Methoden werden anhand einer Simulationsstudie sowie anhand von Messdaten validiert, welche die klaren Verbesserungen in der Auflösung demonstrieren.

Zum Schluss dieser Arbeit wird ein Ausblick gegeben, der das Potential der Kombination der verschiedenen Methoden, die in dieser Arbeit untersucht werden, für verschiedene Bildgebungsszenarien diskutiert.

Contents

1. Introduction

X-rays are electromagnetic waves with shorter wavelengths than visible light. Typical wavelengths of X-rays are in the order of an ångström ($1\,\text{Å} = 10^{-10}\,\text{m}$). From a quantum mechanical perspective, X-rays are quantized into photons. In the latter case, it is more natural to refer to the (quantized) energy rather than the wavelength. Within this work, the energies of interest are in the region just below 10 keV up to over 100 keV. In this picture, the intensity of the X-rays is simply given by the number of photons passing through a given area per unit time, whereas in the wave picture the amplitude relates to the intensity. [Als-Nielsen and McMorrow, 2011, Attwood and Sakdinawat, 2017]

1.1. X-ray radiography and computed tomography (CT)

X-rays and their ability to penetrate opaque materials were first described by Wilhelm Conrad Röngten in 1895. As the attenuation of the X-ray intensity varies between different materials, the interior structure of objects can be depicted. For example, the intensity of X-rays is more strongly attenuated by bone than by soft-tissue. This resulted in the first radiograph of a human hand in the same year, already indicating the significance of X-rays in the field of medicine. Thereby, the hand was placed in between an X-ray source and an X-ray detector (fluorescent screen), which is sensitive to the intensity of the X-rays behind the object. [Röntgen, 1895, Als-Nielsen and McMorrow, 2011]

Considering radiography the first revolution in medical imaging, the second revolution was the invention of computed tomography (CT) in the 1970's by Godfrey Hounsfield. The theory of CT had been independently worked out by Allan McLeod Cormack in the previous decade. In contrast to conventional radiography, which is only sensitive to the projected attenuation of the object along the X-ray paths, CT enables the reconstruction of the full three-dimensional interior of the object by acquiring radiographs from various positions all around the object. Since then, medical CT scanners have improved significantly to allow for higher spatial resolution down to 0.5 mm, faster acquisition times and improved image quality. The most important aspect for improving CT is to reduce radiation dose given to the patient, which is in the range of $1 - 10\,\text{mSv}$ depending on the scanning protocol. In comparison, in radiography, the patients are exposed to approximately 0.1 mSv. [Als-Nielsen and McMorrow, 2011]

Beyond medical imaging, CT is also widely applied to many problems in materials and biological sciences, where dose requirements are often less strict. State-of-the-art laboratory high-resolution CT devices that utilize the same projection-based imaging principle (without using X-ray optics) are able to resolve structures down to 100 nm. However, the resolution is ultimately limited by the extent of the X-ray source spot, due to the high magnifications as well as by mechanical instabilities and thermal drifts of the individual components at this scale. Moreover, with higher resolution, the field-of-view (FOV) is also limited and acquisition times increase significantly. [Als-Nielsen and McMorrow, 2011, Müller et al., 2017]

X-ray radiography and CT are only sensitive to the attenuating properties of the object. Thus, these techniques are not particularly well suited for visualizing weakly absorbing

materials without additional contrast agents or objects that consist of materials with similar attenuating properties, such as soft-tissue, most commonly encountered in biology and medicine. [Bravin et al., 2013]

1.2. Phase-contrast imaging (PCI)

Electromagnetic waves are not only attenuated when interacting with an object, but their phase is shifted as well. However, this shift in phase cannot be resolved directly, as only the intensity can be measured. Early methods to indirectly visualize this phase shift include Schlieren phase contrast and out-of-focus contrast. Most notably, in the early 1930's Frits Zernike developed a phase-contrast method for visible light microscopy enabling high contrast visualization of living tissue cultures. [Zernike, 1942, Zernike, 1955].

The first X-ray phase-contrast images were acquired in 1965 by Ulrich Bonse and Michael Hart using a crystal interferometer [Bonse and Hart, 1965, Momose, 1995]. Today, various different phase-contrast imaging (PCI) methods have been developed that are able to utilize the X-ray phase shift for imaging [Wilkins et al., 2014, Endrizzi, 2018]. These include analyzer-based methods, which use single-crystals [Förster et al., 1980, Davis et al., 1995, Ingal and Beliaevskaya, 1995], and propagation-based phase-contrast imaging (PBI) methods, which do not rely on any additional X-ray optics in the beam path [Snigirev et al., 1995, Cloetens et al., 1996]. Other methods introduce different known structures into the beam path to track the changes of the radiation field induced by the sample. Among these structures are lenslet arrays [Mayo and Sexton, 2004], two-dimensional grids [Morgan et al., 2011], speckle patterns [Morgan et al., 2012] or coded apertures in the edge-illumination method [Olivo et al., 2001]. In addition, grating-based phase-contrast imaging (GBI) methods are popular, which use two gratings and employ the Talbot effect to resolve the induced phase shifts [David et al., 2002, Momose et al., 2003, Weitkamp et al., 2005].

These methods work best at synchrotron facilities, which create highly coherent monochromatic X-rays with high intensity. Using synchrotron radiation, additional phase-contrast imaging methods beyond the aforementioned projection imaging are feasible, such as holography [Eisebitt et al., 2004], coherent diffraction imaging [Miao et al., 1999] and ptychography [Rodenburg et al., 2007]. However, synchrotrons are large scale facilities, due to the requirements on the size of the electron storage rings, and are expensive to build, operate and maintain. This limits the accessibility of X-ray phase-contrast imaging methods significantly.

Some of the methods mentioned above have been successfully transferred to laboratory environments. For instance, with the introduction of a source grating, the GBI method could be transferred to conventional laboratory X-ray sources in radiography [Pfeiffer et al., 2006] and CT [Pfeiffer et al., 2007]. Since then, this method has become a promising candidate for translating phase contrast to clinical environments [Scherer et al., 2015, Koehler et al., 2015]. Also methods using coded apertures such as edge-illumination have been implemented using laboratory sources [Olivo and Speller, 2007]. However, both methods require additional optics, which reduce the incident flux due to the source grating in the GBI method or the sample mask in the edge-illumination method. Moreover, higher (sample) dose is required for both techniques, due to the analyzer grating in the GBI method or the detector mask in the edge illumination method.

Regarding phase-contrast techniques, this work focuses on the PBI method, which does not directly suffer from flux reductions and increased dose requirements as there are no additional apertures in the beam path. However, PBI requires X-rays which are suffi-

ciently spatially coherent, such that the phase shifts induced by the sample lead to distinct variations in the measured intensity [Wilkins et al., 1996]. These variations originate from self-interference effects of the wave during free-space propagation between the sample and the detector. Different PBI methods exist that utilize one or more images at different distances behind the sample or at different energies. With advances of laboratory X-ray sources including liquid-metal jet sources [Hemberg et al., 2003] and compact synchrotrons [Eggl et al., 2016], the available spatial coherence is sufficient to transfer PBI from synchrotron facilities to laboratories. Mostly due to the reduced spatial coherence of laboratory sources [Wilkins et al., 1996, Wilkins et al., 2014], but also due to difficulties in image alignment because of magnification effects, PBI methods at laboratory sources typically acquire images only at a single distance. This makes this method also particularly well suited for the extension to CT referred to as propagation-based phase-contrast computed tomography (PB-CT). [Endrizzi, 2018]

In general, it is not possible to recover the attenuating and phase-shifting properties independently by using only a single image [Endrizzi, 2018]. Thus, additional constraints about the sample have to be imposed. One of the most prominent assumptions is the homogeneity assumption which couples the attenuating and phase-shifting properties of the sample [Paganin et al., 2002]. Recently, it has been demonstrated that under this assumption, noise levels in PBI and PB-CT can be reduced significantly compared to conventional radiography or CT without loss of resolution [Kitchen et al., 2017, Gureyev et al., 2017]. Therefore, this method has the potential to reduce dose requirements, making it a candidate for medical imaging in the future. Beyond the dose argument, PBI and PB-CT have become extremely popular for various applications, due to the ease of implementation and the benefits of utilizing phase contrast for visualization including high-resolution CT applications, (dynamic) lung imaging and breast tissue imaging [Bravin et al., 2013, Endrizzi, 2018].

However, there are still several challenges regarding PBI, in particular when using laboratory X-ray sources. Most importantly, the image quality deteriorates when the homogeneity assumption is violated. In addition, compared to synchrotron facilities, the reduced spatial coherence of the X-rays influences the image quality. Due to the relatively low flux of laboratory sources, high noise levels are present. Therefore, efficient detection systems are required, which include thicker scintillators and thus result in a larger spread of the signal reducing image quality further.

1.3. Tomographic reconstruction techniques

Ever since the invention of CT, algorithms have become a key tool for advances in X-ray imaging. These algorithms are required to reconstruct the three-dimensional structure of the sample from the radiographs acquired around the object. Although early CT scanners acquired only a limited amount of data, the early reconstruction algorithms were inaccurate as well as time-consuming due to their iterative nature and the limited performance of computer systems. [Beckmann, 2006].

The first breakthrough was the invention of analytical reconstruction algorithms that could directly reconstruct the three-dimensional structure of the sample from the acquired radiographs. These algorithms are still utilized today, as they achieve good image quality, are fast and can thus easily handle large amounts of data. [Nuyts et al., 2013]

With the increase of computational power and advances in computer architectures, more advanced iterative CT algorithms are heavily researched. One of the main reasons is

the growing concern about radiation dose and the success of these algorithms to reduce dose requirements, while maintaining image quality. In addition, these algorithms are robust and more flexible to different geometries, enabling new CT designs. Furthermore, these algorithms allow for precise models of the whole image formation process, including photon transport and detection physics. Moreover, one can account for various effects that deteriorate image quality, such as the finite spatial resolution, Compton scatter and noise. [Nuyts et al., 2013]

1.4. Aim and outline

The aim of the work presented in this thesis is to investigate and advance modern reconstruction approaches for medical CT, high-resolution CT and PB-CT at laboratory sources. The focus lies on incorporating physical models of the image formation process to improve image quality. This includes reconstruction approaches for PBI and PB-CT that model the interference effects and that are capable of evaluating the validity of the underlying assumptions for individual data points. In addition, models for the reduced spatial coherence are developed and recent advances in medical CT reconstruction for incorporating detector models including covariance noise statistics are transferred to high-resolution CT and PB-CT. Moreover, general algorithmic developments are proposed to estimate system blur during reconstruction in the context of medical CT. Finally, algorithms for estimating geometric parameters for high-resolution CT are evaluated in the context of biological sciences.

The structure of this thesis is divided into two main parts. The first part summarizes the theoretical background required for the understanding and the derivations of the methods and results presented in the second part.

The theoretical background itself is divided into three chapters. In Chapter 2, starting with the wave equation, a consistent description of X-rays is outlined using scalar wave theory. Thereby, the underlying approximations, the interactions of the X-rays with matter, free-space propagation and a description of coherence, relevant to the second part, are discussed. Subsequently, Chapter 3 evolves around PBI focusing on various phase-retrieval algorithms. In particular, the single-material phase-retrieval algorithm and its underlying assumptions and implications are discussed in detail. In Chapter 4, tomographic reconstruction techniques are outlined. After discussing the discretization of the data and operations, the fundamentals of iterative reconstruction are presented in the context of maximum a posteriori (MAP) estimation.

The second part describes the developed methods and obtained results. In Chapter 5, a general MAP approach for PBI and PB-CT is investigated. Thereby, various special cases are discussed, in particular, a formulation which obeys the homogeneity assumption is addressed in greater detail. Subsequently, the proposed approach for PBI is validated using a laboratory micro CT system and compared to the single-material phase-retrieval algorithm. Afterwards, the benefits of the proposed approach for PB-CT are investigated for two applications. First, various reconstruction approaches are compared by means of a test sample, which violates the homogeneity assumption, and was measured at a laboratory X-ray micro CT system. The second application details benefits of the proposed approach for visualizing an animal cochlea with an implant using synchrotron radiation.

In Chapter 6, incorporating physical models of the spatial coherence of the X-ray source, the detector response and noise correlations into a MAP reconstruction approach is analyzed in the context of homogeneous PBI and PB-CT. An approximate analytical solution of the

proposed algorithm is derived which further generalizes the single-material phase-retrieval algorithm and its extensions. Furthermore, a relation between the single-material phase-retrieval algorithm and regularization techniques is discussed. A simulation study as well as an experimental study at a compact synchrotron are performed detailing significant advantages in terms of resolution and noise in the context of the proposed approach, which accounts for the covariance statistics and system blur.

In order to leverage the full potential of the previous approach, exact models describing the system blur have to be available. This usually requires the exact characterization of all components that add to the system blur, which can be laborious. In Chapter 7, a reconstruction approach is proposed, which jointly estimates the system blur during tomographic reconstruction to eliminate the need for exact characterizations of the system blur. This work is done in the context of conventional CT.

Finally, Chapter 8 is concerned with high-resolution CT applications. At this scale, phase effects are often unavoidable, but can also be used to increase contrast. One of the most crucial challenges is to properly characterize the geometry of the system, which is corrupted by random drifts and thus can not be fully characterized beforehand. Here, two optimization-based approaches to estimate geometry parameters are investigated. The first method uses a metric proposed in the previous chapter and is applied to the reconstruction of a sea cucumber. The second approach jointly estimates the geometry parameters during tomographic reconstruction. This method is validated using a simulation study and an experimental test sample.

Chapter 9 gives a general conclusion and outlook for future developments and shows how the investigated approaches can play together for various applications, detailing their limitations and potentials.

Part I.

Theoretical background

2. Image formation description using scalar wave theory

This chapter introduces the underlying formalism to describe the properties of X-ray wavefields. In particular, the notion of paraxial wavefields is introduced, their interaction with matter is outlined and their evolution in free-space is discussed. The chapter concludes with a short introduction to the concept of partial coherence. The aim of the chapter is a self-consistent derivation of the formalism relevant for this thesis.

Various theories describing electromagnetic radiation exist, each with its own benefits and drawbacks. For instance, in geometrical optics, X-rays are regarded as rays propagating along straight lines in a homogeneous medium. At interfaces between two media, the reflecting and refracting properties can be described in terms of the refractive index of each medium. However, effects like diffraction, interference, absorption and scattering are usually not described using geometrical optics, although there exists a geometrical theory of diffraction, which extends geometrical optics [Keller, 1962].

To model these effects, X-rays can be regarded as electromagnetic waves. In the following, X-rays are described as a complex scalar wavefield obeying the wave equation. The properties of a medium is given in terms of its complex refractive index. Although more accurately, one can characterize electromagnetic waves in terms of their coupled electric and magnetic fields obeying the Maxwell equations, the treatment of X-rays as a scalar wavefield is sufficient to describe diffraction, interference, absorption and scattering relevant for this thesis.

However, some important aspects, such as the noise properties inherent to measuring certain properties of the X-ray wavefield are not described using electromagnetic waves. This requires the notion of wave-particle dualism of quantum mechanics. These aspects will be discussed in Section 4.3, Chapter 5 and 6.

Most of the derivations of this chapter follow [Paganin, 2006]. Many concepts can also be found elsewhere, including [Cowley, 1995, Goodman, 2005, Als-Nielsen and McMorrow, 2011, Peatross and Ware, 2015, Attwood and Sakdinawat, 2017].

2.1. Paraxial wavefields

In scalar wave theory, the electromagnetic disturbance is described by a single complex scalar field $\Psi(\mathbf{r}, t)$ at each point in space $\mathbf{r} = (x, y, z)$ and time t. In free-space this field is governed by the so-called *d'Alembert equation*

$$\boxed{\left(\frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \nabla^2\right)\Psi(\mathbf{r}, t) = 0,} \tag{2.1}$$

which is used as a starting point for the subsequent derivations. This equation can be derived from the Maxwell equations in vacuum. [Paganin, 2006, p.2ff]

2.1.1. Helmholtz equations

To obtain a time-independent wave-equation, the wavefield can be decomposed into its spectral components characterized by its angular frequencies ω according to

$$\Psi(\mathbf{r},t) = \frac{1}{\sqrt{2\pi}} \int_0^\infty \psi_\omega(\mathbf{r}) e^{-\mathrm{i}\omega t} \mathrm{d}\omega. \tag{2.2}$$

It is convenient to omit the negative angular frequencies in the definition, as all information is already encoded in the positive angular frequencies. Here $\psi_\omega(\mathbf{r})$ denote the monochromatic components of the decomposition. Substituting Eq. (2.2) into Eq. (2.1), interchanging the order of differentiation and integration and multiplying by $-\sqrt{2\pi}$ on both sides, results in

$$\int_0^\infty \left[\left(\nabla^2 + \frac{\omega^2}{c^2} \right) \psi_\omega(\mathbf{r}) \right] e^{-\mathrm{i}\omega t} \mathrm{d}\omega = 0. \tag{2.3}$$

Thus, the quantity in square brackets must be zero everywhere. Consequently, each monochromatic component must obey

$$\boxed{\left(\nabla^2 + k^2 \right) \psi_\omega(\mathbf{r}) = 0,} \tag{2.4}$$

where the wave number $k = \omega/c$ is introduced. This equation is called *Helmholtz equation*. By solving the Helmholtz equation for each monochromatic component and using Eq. (2.2), the time-dependency is recovered. [Paganin, 2006, p.4ff]

Within this framework, the influence of a time-independent medium on the wavefield is described by the energy-dependent *complex refractive index* $n_\omega(\mathbf{r})$ of the medium. The equivalent of the Helmholtz equation within a medium is called the *inhomogeneous Helmholtz equation*. It follows by adapting the wave number $k^2 \to k^2 n_\omega^2(\mathbf{r})$ resulting in

$$\boxed{\left[\nabla^2 + k^2 n_\omega^2(\mathbf{r}) \right] \psi_\omega(\mathbf{r}) = 0.} \tag{2.5}$$

A more rigorous derivation of this equation would start from the Maxwell equations in matter instead. [Paganin, 2006, p.69ff]

2.1.2. Paraxial equations

It is convenient to decompose the wavefield into an unscattered plane wave propagating without loss of generality along the optical axis z and an envelope $\tilde{\psi}_\omega(\mathbf{r}_\perp, z)$ according to

$$\psi_\omega(\mathbf{r}) = \tilde{\psi}_\omega(\mathbf{r}_\perp, z) e^{\mathrm{i}kz}, \tag{2.6}$$

with $\mathbf{r} = (\mathbf{r}_\perp, z)$, where $\mathbf{r}_\perp = (x, y)$ denotes the spatial dimensions orthogonal to the optical axis. Inserting Eq. (2.6) into the inhomogeneous Helmholz equation given by Eq. (2.5) results in

$$\left\{ 2\mathrm{i}k \frac{\partial}{\partial z} + \nabla_\perp^2 + \frac{\partial^2}{\partial z^2} + k^2 \left[n_\omega^2(\mathbf{r}_\perp, z) - 1 \right] \right\} \tilde{\psi}_\omega(\mathbf{r}_\perp, z) = 0, \tag{2.7}$$

defining $\nabla^2 = (\nabla_\perp^2, \partial^2/\partial z^2)$ and using

$$\nabla^2 \psi = \nabla \cdot \nabla \tilde{\psi} e^{\mathrm{i}kz} = e^{\mathrm{i}kz} \nabla_\perp^2 \tilde{\psi} + e^{\mathrm{i}kz} \frac{\partial^2 \tilde{\psi}}{\partial z^2} + 2\mathrm{i}k e^{\mathrm{i}kz} \frac{\partial \tilde{\psi}}{\partial z} - k^2 e^{\mathrm{i}kz} \tilde{\psi}, \tag{2.8}$$

where for convenience the subscript and dependencies of the wavefield are not explicitly denoted. In free space $n_\omega^2(\mathbf{r}_\perp, z) = 1$ holds. [Paganin, 2006, p.71ff]

The paraxial approximation refers to neglecting the second derivative with respect to z in Eq. (2.7), implying that the envelope $\tilde{\psi}_\omega(\mathbf{r}_\perp, z)$ is less strongly varying in the direction of the optical axis compared to its orthogonal components. Thus, the *inhomogeneous paraxial equation* is given by

$$\boxed{\left\{2\mathrm{i}k\frac{\partial}{\partial z} + \nabla_\perp^2 + k^2\left[n_\omega^2(\mathbf{r}_\perp, z) - 1\right]\right\}\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = 0} \tag{2.9}$$

and consequently, the *paraxial equation*, describing the evolution of the envelope wavefield in free space, is given by

$$\boxed{\left(2\mathrm{i}k\frac{\partial}{\partial z} + \nabla_\perp^2\right)\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = 0} \tag{2.10}$$

using $n_\omega^2(\mathbf{r}_\perp, z) = 1$. [Paganin, 2006, p.69ff]

2.2. Interaction of X-rays with matter

Modeling the interactions of the X-rays with matter usually employs the projection approximation. This approximation requires that the medium is sufficiently thin or that the scatterers within the medium scatter sufficiently weakly, such that the Laplacian in the inhomogeneous paraxial equation given by Eq. (2.9) can be neglected. This results in

$$\left\{2\mathrm{i}k\frac{\partial}{\partial z} + k^2\left[n_\omega^2(\mathbf{r}_\perp, z) - 1\right]\right\}\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = 0. \tag{2.11}$$

Assuming that the medium is only located between $z = 0$ and $z = z_0$ along the optical axis, Eq. (2.11) can be solved for the exit wavefield envelope $\tilde{\psi}_\omega(\mathbf{r}_\perp, z_0)$ at $z = z_0$ according to

$$\tilde{\psi}_\omega(\mathbf{r}_\perp, z_0) = \exp\left\{\frac{k}{2\mathrm{i}}\int_0^{z_0}\left[1 - n_\omega^2(\mathbf{r}_\perp, z)\right]\mathrm{d}z\right\}\tilde{\psi}_\omega(\mathbf{r}_\perp, 0), \tag{2.12}$$

given the incoming wavefield envelope $\tilde{\psi}(\mathbf{r}_\perp, 0)$ at $z = 0$. In the following, the plane behind the sample, in this case at $z = z_0$, is referred to as the *object plane*. For X-rays, the complex refractive index is conventionally written as

$$n_\omega(\mathbf{r}_\perp, z) = 1 - \delta_\omega(\mathbf{r}_\perp, z) + \mathrm{i}\beta_\omega(\mathbf{r}_\perp, z), \tag{2.13}$$

where δ_ω refers to the *refractive index decrement* and β_ω to the *imaginary part*. As for X-rays, the refractive index decrement as well as the imaginary part are usually distinctively less than unity, the expansion

$$1 - n_\omega(\mathbf{r}_\perp, z)^2 \approx 2\left[\delta_\omega(\mathbf{r}_\perp, z) - \mathrm{i}\beta_\omega(\mathbf{r}_\perp, z)\right], \tag{2.14}$$

where only the first order terms in δ_ω and β_ω are kept, holds reasonably well. Consequently, the wavefield in the object plane is given by

$$\tilde{\psi}_\omega(\mathbf{r}_\perp, z_0) = \exp\left\{-\mathrm{i}k\int_0^{z_0}\left[\delta_\omega(\mathbf{r}_\perp, z) - \mathrm{i}\beta_\omega(\mathbf{r}_\perp, z)\right]\mathrm{d}z\right\}\tilde{\psi}_\omega(\mathbf{r}_\perp, 0), \tag{2.15}$$

under the assumptions outlined above. [Paganin, 2006, p.71ff]

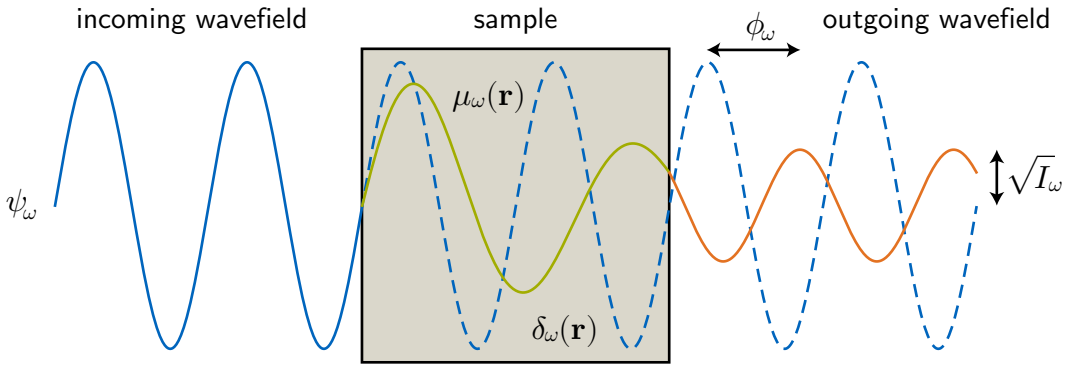

Figure 2.1.: Interactions of an X-ray wavefield with matter. The incoming wavefield ψ_ω is depicted on the left in blue. In green, the wavefield as it traverse the sample is depicted and the outgoing wavefield is shown in orange. The attenuating and phase-shifting properties of the sample are described by the linear attenuation coefficient $\mu_\omega(\mathbf{r})$ and refractive index decrement $\delta_\omega(\mathbf{r})$ respectively. The envelope wavefield behind the sample is defined by its intensity $\sqrt{I_\omega}$ and phase shift ϕ_ω compared to the undisturbed wavefield illustrated by the dashed blue line.

In the following, it will be useful to decompose different parts of the envelope according to

$$\boxed{\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = \sqrt{I_\omega(\mathbf{r}_\perp, z)}\exp\left[\mathrm{i}\phi_\omega(\mathbf{r}_\perp, z)\right],} \tag{2.16}$$

where the intensity $I_\omega(\mathbf{r}_\perp, z) = |\tilde{\psi}_\omega(\mathbf{r}_\perp, z)|^2$ is the squared modulus of the wavefield and the phase $\phi_\omega(\mathbf{r}_\perp, z) = \arg\tilde{\psi}_\omega(\mathbf{r}_\perp, z)$ is the complex argument of the wavefield. [Paganin, 2006, p.279]

Defining the *linear attenuation coefficient* as $\mu_\omega(\mathbf{r}_\perp, z) = 2k\beta_\omega(\mathbf{r}_\perp, z)$, the intensity in the object plane is obtained by squaring Eq. (2.15) resulting in

$$\boxed{I_\omega(\mathbf{r}_\perp, z_0) = I_\omega(\mathbf{r}_\perp, 0)\exp\left[-\int_0^{z_0}\mu_\omega(\mathbf{r}_\perp, z)\mathrm{d}z\right].} \tag{2.17}$$

This equation is referred to as the *Lambert-Beer law* and describes how the intensity is attenuated by the sample according to the projection of the three-dimensional distribution of the linear attenuation coefficient along the X-ray paths. Figure 2.1 illustrates how these quantities are related. In conventional CT, this is the property which is to be reconstructed. In the particle picture, the attenuation in intensity in the energy range of interest is mostly due to two processes. First of all, photoelectric absorption, where an entire X-ray photon is absorbed by an atom. This is possible, if the binding energy of the atomic electron is smaller than the energy of the photon. As a consequence, the electron, with which the X-ray photon interacts, is released from the atom. Subsequent processes are X-ray fluorescence or the Auger process. The probability of photoelectric absorption increases with the atomic number and decreases with increasing energy. The second process is Compton scattering. Here, the X-ray photon interacts with an electron mostly in the outer shell of the atom. After the interaction, the energy of the X-ray photon is only reduced and the photon continues to traverse the object. The complementary part of the energy is transferred to the electron, which is released from the atom. The probability of Compton scattering increases with electron density. [Buzug, 2008]

The other quantity that is altered is the phase. The total phase shift imposed by the medium can be obtained by comparing Eq. (2.15) and Eq. (2.16) according to

$$\boxed{\phi_\omega(\mathbf{r}_\perp, z_0) = -k \int_0^{z_0} \delta_\omega(\mathbf{r}_\perp, z)\mathrm{d}z,} \tag{2.18}$$

which is proportional to the projection of the refractive index decrement along the X-ray paths. The phase shift is also illustrated in Figure 2.1. The underlying processes are not as intuitively described in the particle picture as the attenuation processes. [Paganin, 2006, p.71ff]

2.3. Free-space propagation

In the following, the evolution of the X-ray wavefield in free space, referred to as free-space propagation, is outlined, focusing on the limiting case of small propagation distances, relevant for this work.

2.3.1. Fresnel propagator

Given a monochromatic paraxial wavefield as specified above, the evolution of the envelope $\tilde{\psi}_\omega$ in free space is governed by the paraxial equation given by Eq. (2.10). Assuming elementary plane waves

$$\tilde{\psi}_\omega^{(\mathrm{PW})}(\mathbf{r}_\perp, z) = \exp\left[\mathrm{i}\left(\mathbf{k}_\perp \cdot \mathbf{r}_\perp + k_z z\right)\right], \tag{2.19}$$

the z-component of the wavevector k_z can be calculated by inserting Eq. (2.19) into Eq. (2.10) resulting in

$$k_z = -\frac{\mathbf{k}_\perp^2}{2k}. \tag{2.20}$$

Thus, the plane wave at any point z along the optical axis is given by

$$\tilde{\psi}_\omega^{(\mathrm{PW})}(\mathbf{r}_\perp, z) = \tilde{\psi}_\omega^{(\mathrm{PW})}(\mathbf{r}_\perp, 0) \exp\left[-\frac{\mathrm{i}z}{2k}\mathbf{k}_\perp^2\right], \tag{2.21}$$

where $\tilde{\psi}_\omega^{(\mathrm{PW})}(\mathbf{r}_\perp, 0)$ is obtained from Eq. (2.19) for $z = 0$. Although this is only valid for plane waves, similarly to the decomposition of spectral components given by Eq. (2.2), one can decompose the unpropagated wavefield $\tilde{\psi}_\omega(\mathbf{r}_\perp, 0)$ into a linear combination of plane waves $\tilde{\psi}_\omega(\mathbf{k}_\perp, 0)$ according to

$$\tilde{\psi}_\omega(\mathbf{r}_\perp, 0) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{\psi}_\omega(\mathbf{k}_\perp, 0) \exp\left[\mathrm{i}\mathbf{k}_\perp \cdot \mathbf{r}_\perp\right] \mathrm{d}\mathbf{k}_\perp. \tag{2.22}$$

Mathematically $\tilde{\psi}_\omega(\mathbf{k}_\perp, 0)$ is the two-dimensional Fourier transform (FT) of $\tilde{\psi}_\omega(\mathbf{r}_\perp, 0)$ with respect to $\mathbf{r}_\perp$. Each plane wave propagates according to Eq. (2.21), thus the propagated wavefield is given by

$$\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{\psi}_\omega(\mathbf{k}_\perp, 0) \exp\left[-\frac{\mathrm{i}z}{2k}\mathbf{k}_\perp^2\right] \exp\left[\mathrm{i}\mathbf{k}_\perp \cdot \mathbf{r}_\perp\right] \mathrm{d}\mathbf{k}_\perp. \tag{2.23}$$

Finally, this can be rewritten using an operator formulation according to

$$\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = \mathcal{D}_z \tilde{\psi}_\omega(\mathbf{r}_\perp, 0), \tag{2.24}$$

where the so-called *Fresnel operator*

$$\boxed{\mathcal{D}_z = \mathcal{F}_\perp^{-1} \exp\left[-\frac{\mathrm{i}z}{2k}\mathbf{k}_\perp^2\right] \mathcal{F}_\perp} \tag{2.25}$$

is given by Eq. (2.23). In this notation, $\mathcal{F}_\perp$ maps the $\tilde{\psi}_\omega(\mathbf{r}_\perp, 0)$ to its FT $\tilde{\psi}_\omega(\mathbf{k}_\perp, 0)$ and $\mathcal{F}_\perp^{-1}$ denotes the inverse Fourier transform (IFT) and these operators act from right to left. Formally, these operators are defined according to

$$f(\mathbf{r}_\perp) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\mathbf{k}_\perp) \exp\left[\mathrm{i}\mathbf{k}_\perp \cdot \mathbf{r}_\perp\right] \mathrm{d}\mathbf{k}_\perp, \tag{2.26}$$

such that $f(\mathbf{k}_\perp) \equiv \mathcal{F}_\perp[f(\mathbf{r}_\perp)]$ and

$$f(\mathbf{k}_\perp) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\mathbf{r}_\perp) \exp\left[-\mathrm{i}\mathbf{k}_\perp \cdot \mathbf{r}_\perp\right] \mathrm{d}\mathbf{r}_\perp, \tag{2.27}$$

such that $f(\mathbf{r}_\perp) \equiv \mathcal{F}_\perp^{-1}[f(\mathbf{k}_\perp)]$. The Fresnel operator is not explicitly derived this way in [Paganin, 2006, p.6ff], but follows analogous to the derivation of the angular spectrum method. [Paganin, 2006, p.6ff]

2.3.2. Limit of small propagation distances

The main application of the work evolves around small propagation distances. Small in the sense that one can make the following approximation in the exponent of the Fresnel operator

$$\exp\left[-\frac{\mathrm{i}z}{2k}\mathbf{k}_\perp^2\right] \approx 1 - \frac{\mathrm{i}z}{2k}\mathbf{k}_\perp^2, \tag{2.28}$$

ignoring terms of the order $\mathcal{O}(z^2)$. Substituting this in Eq. (2.23) results in

$$\tilde{\psi}_\omega(\mathbf{r}_\perp, z) \approx \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{\psi}_\omega(\mathbf{k}_\perp, 0) \left[1 - \frac{\mathrm{i}z\mathbf{k}_\perp^2}{2k}\right] \exp\left[\mathrm{i}\mathbf{k}_\perp \cdot \mathbf{r}_\perp\right] \mathrm{d}\mathbf{k}_\perp. \tag{2.29}$$

Using the Fourier derivative theorem according to

$$\nabla_\perp^2 f(\mathbf{r}_\perp) = \nabla_\perp^2 \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\mathbf{k}_\perp) \exp\left[\mathrm{i}\mathbf{k}_\perp \cdot \mathbf{r}_\perp\right] \mathrm{d}\mathbf{k}_\perp \tag{2.30}$$

$$= -\frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{k}_\perp^2 f(\mathbf{k}_\perp) \exp\left[\mathrm{i}\mathbf{k}_\perp \cdot \mathbf{r}_\perp\right] \mathrm{d}\mathbf{k}_\perp \tag{2.31}$$

or equivalently $\mathcal{F}_\perp\left[\nabla_\perp^2 f(\mathbf{r}_\perp)\right] = -\mathbf{k}_\perp^2 \mathcal{F}_\perp\left[f(\mathbf{r}_\perp)\right]$, the propagated wavefield for small distances can then be calculated according to

$$\boxed{\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = \left(1 + \frac{\mathrm{i}z}{2k}\nabla_\perp^2\right) \tilde{\psi}_\omega(\mathbf{r}_\perp, 0),} \tag{2.32}$$

where $\nabla_\perp^2$ denotes the Laplacian operator orthogonally to the optical axis. [Paganin, 2006, p.278ff]

The intensity of the propagated wavefield can be calculated taking the squared modulus of Eq. (2.32). Although the wavefield evolves linearly with the propagation distance, this

does not hold for the intensity of the wavefield. The following expression of the propagated intensity is central to this thesis and is thus derived explicitly. Starting with

$$I_\omega(\mathbf{r}_\perp, z) = |\tilde{\psi}_\omega(\mathbf{r}_\perp, z)|^2 \tag{2.33}$$

$$\approx |\tilde{\psi}_\omega(\mathbf{r}_\perp, 0)|^2 + \tilde{\psi}^*_\omega(\mathbf{r}_\perp, 0)\frac{\mathrm{i}z}{2k}\nabla^2_\perp \tilde{\psi}_\omega(\mathbf{r}_\perp, 0) \tag{2.34}$$

$$- \tilde{\psi}_\omega(\mathbf{r}_\perp, 0)\frac{\mathrm{i}z}{2k}\nabla^2_\perp \tilde{\psi}^*_\omega(\mathbf{r}_\perp, 0) \tag{2.35}$$

$$= |\tilde{\psi}_\omega(\mathbf{r}_\perp, 0)|^2 + 2\Re\left[\tilde{\psi}^*_\omega(\mathbf{r}_\perp, 0)\frac{\mathrm{i}z}{2k}\nabla^2_\perp \tilde{\psi}_\omega(\mathbf{r}_\perp, 0)\right] \tag{2.36}$$

$$= |\tilde{\psi}_\omega(\mathbf{r}_\perp, 0)|^2 - \frac{z}{k}\Im\left[\tilde{\psi}^*_\omega(\mathbf{r}_\perp, 0)\nabla^2_\perp \tilde{\psi}_\omega(\mathbf{r}_\perp, 0)\right], \tag{2.37}$$

where terms of the order $\mathcal{O}(z^2)$ are discarded and the identity $2\Re(w) = w + w^*$ for w being a complex number and w^* being its complex conjugate is employed. Here $\Re$ denotes the real part of the complex number and $\Im$ the imaginary part respectively. The relation above states the connection between the initial wavefield and the propagated intensity. To establish the relation between the propagated intensity and the initial intensity and phase, additional manipulations have to be performed. The first term already corresponds to the intensity of the initial wavefield. The decomposition of the second term requires the intermediate computation of

$$\tilde{\psi}^*\nabla^2\tilde{\psi} = \sqrt{I}e^{-\mathrm{i}\phi}\nabla\cdot\left(\frac{e^{\mathrm{i}\phi}}{2\sqrt{I}}\nabla I + \mathrm{i}\sqrt{I}e^{\mathrm{i}\phi}\nabla\phi\right) \tag{2.38}$$

$$= \sqrt{I}e^{-\mathrm{i}\phi}\left[\frac{e^{\mathrm{i}\phi}}{2\sqrt{I}}\nabla^2 I + \left(-\frac{e^{\mathrm{i}\phi}}{4I^{3/2}}\nabla I + \mathrm{i}\frac{e^{\mathrm{i}\phi}}{2\sqrt{I}}\nabla\phi\right)\cdot\nabla I\right. \tag{2.39}$$

$$\left. + \mathrm{i}\sqrt{I}e^{\mathrm{i}\phi}\nabla^2\phi + \left(\mathrm{i}\frac{e^{\mathrm{i}\phi}}{2\sqrt{I}}\nabla I - \sqrt{I}e^{\mathrm{i}\phi}\nabla\phi\right)\cdot\nabla\phi\right] \tag{2.40}$$

$$= \frac{1}{2}\nabla^2 I - \frac{1}{4I}\nabla I\cdot\nabla I - I\nabla\phi\cdot\nabla\phi + \mathrm{i}\left(\nabla I\cdot\nabla\phi + I\nabla^2\phi\right) \tag{2.41}$$

$$= \sqrt{I}\nabla^2\sqrt{I} - I\nabla\phi\cdot\nabla\phi + \mathrm{i}\nabla\cdot(I\nabla\phi), \tag{2.42}$$

where according to Eq. (2.16) the initial wavefield was decomposed into its intensity and phase. For simplicity, the subscripts and arguments are omitted. In the end, the product rule is applied in reverse. Keeping the imaginary part and recovering the subscript and arguments, the propagated intensity is given in terms of the intensity and phase of the initial wavefield according to

$$\boxed{I_\omega(\mathbf{r}_\perp, z) = I_\omega(\mathbf{r}_\perp, 0) - \frac{z}{k}\nabla_\perp\cdot\left[I_\omega(\mathbf{r}_\perp, 0)\nabla_\perp\phi_\omega(\mathbf{r}_\perp, 0)\right].} \tag{2.43}$$

This equation can be derived in different ways. The derivation extends [Paganin, 2006, p.278ff] with the explicit computation of the stated formulas. This equation can be interpreted geometrically. The stronger the wavefield is curved at a given point in the sample plane, the stronger the change in intensity at that corresponding point in the detector plane. [Paganin, 2006, p.281ff]

2.3.3. Transport-of-intensity and transport-of-phase equations

The transport-of-intensity equation (TIE) and the transport-of-phase equation (TPE) state how the intensity and phase change along the optical axis, given the intensity and phase

at any point along the optical axis. The TPE is also referred to as the eikonal equation. Especially the TIE has often been used as a starting point for deriving various phase retrieval algorithms [Gureyev and Nugent, 1996, Paganin et al., 2002].

The derivation of the TPE starts from the paraxial equation given by Eq. (2.10). First, the complex conjugate of the wavefield is multiplied to the left of the paraxial equation resulting in

$$\tilde{\psi}_\omega^*(\mathbf{r}_\perp, z)\left(2\mathrm{i}k\frac{\partial}{\partial z} + \nabla_\perp^2\right)\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = 0 \tag{2.44}$$

$$2\mathrm{i}k\tilde{\psi}_\omega^*(\mathbf{r}_\perp, z)\frac{\partial\tilde{\psi}_\omega(\mathbf{r}_\perp, z)}{\partial z} + \tilde{\psi}_\omega^*(\mathbf{r}_\perp, z)\nabla_\perp^2\tilde{\psi}_\omega(\mathbf{r}_\perp, z) = 0. \tag{2.45}$$

The second term has already been evaluated in Eq. (2.42). The first term follows according to

$$\psi^*\frac{\partial\psi}{\partial z} = \sqrt{I}e^{-\mathrm{i}\phi}\frac{\partial\left(\sqrt{I}e^{\mathrm{i}\phi}\right)}{\partial z} \tag{2.46}$$

$$= \sqrt{I}e^{-\mathrm{i}\phi}\left(\frac{e^{\mathrm{i}\phi}}{2\sqrt{I}}\frac{\partial I}{\partial z} + \mathrm{i}\sqrt{I}e^{\mathrm{i}\phi}\frac{\partial\phi}{\partial z}\right) \tag{2.47}$$

$$= \frac{1}{2}\frac{\partial I}{\partial z} + \mathrm{i}I\frac{\partial\phi}{\partial z}. \tag{2.48}$$

Combining Eq. (2.42) and Eq. (2.48) in the context of Eq. (2.45) and isolating the real part results in the TPE, given by

$$-2kI(\mathbf{r}_\perp, z)\frac{\partial\phi(\mathbf{r}_\perp, z)}{\partial z} + \sqrt{I(\mathbf{r}_\perp, z)}\nabla_\perp^2\sqrt{I(\mathbf{r}_\perp, z)} \tag{2.49}$$

$$- I(\mathbf{r}_\perp, z)\nabla_\perp\phi(\mathbf{r}_\perp, z)\cdot\nabla_\perp\phi(\mathbf{r}_\perp, z) = 0 \tag{2.50}$$

$$\boxed{2k\frac{\partial\phi(\mathbf{r}_\perp, z)}{\partial z} = -\nabla_\perp\phi(\mathbf{r}_\perp, z)\cdot\nabla_\perp\phi(\mathbf{r}_\perp, z) + \frac{1}{\sqrt{I(\mathbf{r}_\perp, z)}}\nabla_\perp^2\sqrt{I(\mathbf{r}_\perp, z)}.} \tag{2.51}$$

This equation states how the phase changes in the direction of the optical axis given the intensity and phase at some point z along the optical axis. The final form of the equation is for instance stated in [Gureyev and Nugent, 1996].

The more prominent equation is the TIE. It can be derived by again inserting the results of Eq. (2.42) and Eq. (2.48) into Eq. (2.45), but solving for the imaginary part instead. This results in

$$k\frac{\partial I(\mathbf{r}_\perp, z)}{\partial z} + \nabla_\perp\cdot\left[I(\mathbf{r}_\perp, z)\nabla_\perp\phi(\mathbf{r}_\perp, z)\right] = 0 \tag{2.52}$$

$$\boxed{k\frac{\partial I(\mathbf{r}_\perp, z)}{\partial z} = -\nabla_\perp\cdot\left[I(\mathbf{r}_\perp, z)\nabla_\perp\phi(\mathbf{r}_\perp, z)\right].} \tag{2.53}$$

This equation was suggested for phase retrieval in [Teague, 1983, Gureyev and Nugent, 1996]. The TIE implies energy conservation along the optical axis. [Sziklas and Siegman, 1974, Paganin, 2006, p.297]

Finally, using the finite difference approximation of the gradient of the intensity with the propagation distance

$$\frac{\partial I(\mathbf{r}_\perp, z)}{\partial z} \approx \frac{I(\mathbf{r}_\perp, z) - I(\mathbf{r}_\perp, 0)}{z} \tag{2.54}$$

one can recover Eq. (2.43), previously derived by expanding the Fresnel operator with respect to the propagation distance. [Paganin, 2006, p.296f]

2.3.4. Fresnel scaling theorem

The previous derivations assume that plane wave illumination, e.g. the source located at $z = -\infty$. While this can be a reasonable assumption if the distance between the source and the object is large as it is usually the case for experiments conducted at synchrotron facilities, for laboratory environments, this assumption is usually violated. Here, one has to explicitly take into account that the X-rays originate from a point source located at some position $z = -R$.

The *Fresnel scaling theorem* connects the intensity pattern $I_\omega^{(\mathrm{R})}$ that one acquires from a point-source illumination with a point located at $z = -R$ and relates it to a virtual intensity pattern $I_\omega^{(\infty)}$ that one would have acquired using plane wave illumination. It is given by

$$\boxed{I_\omega^{(\mathrm{R})}(\mathbf{r}_\perp, z) = \frac{1}{M^2} I_\omega^{(\infty)}\left(\frac{\mathbf{r}_\perp}{M}, \frac{z}{M}\right),} \tag{2.55}$$

where the geometric magnification

$$M = \frac{R+z}{R} \tag{2.56}$$

is introduced. Essentially, an intensity pattern obtained by point-illumination can be thought of as an intensity pattern obtained by plane-wave illumination with an effective propagation distance of z/M. The preceded factor $1/M^2$ is due to energy conservation. In addition, $\mathbf{r}_\perp/M$ takes the geometric distortions of the lengths into account. [Paganin, 2006, p.397ff]

2.3.5. Towards larger propagation distances

For larger propagation distances, different expansions have been studied in literature. For instance, instead of decomposing the wavefield into its intensity and phase according to Eq. (2.6), one can use the following decomposition

$$\tilde{\psi}_\omega(\mathbf{r}_\perp) = \exp\left[-B(\mathbf{r}_\perp) + \mathrm{i}\phi(\mathbf{r}_\perp)\right]. \tag{2.57}$$

Assuming a weakly absorbing object $B \ll 1$ and only slowly varying phase, one can devise the following linearization

$$\tilde{\psi}_\omega(\mathbf{r}_\perp) \approx 1 - B(\mathbf{r}_\perp) + \mathrm{i}\phi(\mathbf{r}_\perp), \tag{2.58}$$

which is the starting point for the so-called contrast transfer function (CTF) formulation [Cowley, 1995]. This has the benefit that this linear equation can be computed more efficiently. Other methods like a TIE variant, which assumes weakly absorbing objects or mixed approaches between the CTF and TIE regime, have been proposed. For this, the reader is referred to [Langer et al., 2008] as a starting point. [Langer et al., 2008]

There are also various formulations beyond the Fresnel regime. In particular, the angular spectrum formulation is derived similarly to the derivation of the Fresnel operator, but without imposing the paraxial approximation. In addition, Fraunhofer diffraction refers to particularly large propagation distances in the so-called far-field. Here the propagated wavefield is elegantly connected to the initial wavefield by a Fourier transform. [Paganin, 2006, p.6ff]

However, this work is concerned with imaging regimes accessible by laboratory X-ray sources and thus, these formulations will be discarded in the following.

2.4. Space-frequency description of partial coherence

The *mutual coherence function* is a useful tool for describing partial coherence of wavefields. It can be defined as

$$\Gamma(\mathbf{r}_1, \mathbf{r}_2, \tau) = \langle \Psi(\mathbf{r}_1, t+\tau)\Psi^*(\mathbf{r}_2, t)\rangle, \tag{2.59}$$

where τ is the time difference between the wavefield originating at $\mathbf{r}_1$ and the wavefield originating at $\mathbf{r}_2$ to a given point $\mathbf{r}$. In this definition, Ψ is assumed to be an ergodic, stationary random process. Stationarity implies that the mutual coherence function is independent of the time t. Ergodicity on the other hand allows the time average denoted by $\langle\rangle$ to be interpreted as an ensemble average over all possible fields. The mutual coherence function is a measure for the correlation between two wavefields originating at $\mathbf{r}_1$ and $\mathbf{r}_2$. The respective interference term at $\mathbf{r}$ (defined by τ) is proportional to the real part of the mutual coherence function. Consequently, if the two wavefields are uncorrelated, the real part of the mutual coherence function is zero. Under the assumptions outlined above, the following relation

$$I(\mathbf{r}) = \Gamma(\mathbf{r}, \mathbf{r}, 0) = \langle |\Psi(\mathbf{r}, t)|^2\rangle \tag{2.60}$$

connects the mutual coherence function to the observable intensity I at $\mathbf{r}$. [Carter, 1993, Paganin, 2006, p.40ff]

Similarly to the spectral decomposition of Ψ, it is convenient to also decompose the mutual coherence function into its monochromatic components. Inserting Eq. (2.2) into Eq. (2.59) yields

$$\begin{aligned}
\Gamma(\mathbf{r}_1, \mathbf{r}_2, \tau) &= \langle \Psi(\mathbf{r}_1, t+\tau)\Psi^*(\mathbf{r}_2, t)\rangle && (2.61)\\
&= \left\langle \frac{1}{\sqrt{2\pi}}\int_0^\infty \psi_\omega(\mathbf{r}_1)e^{-\mathrm{i}\omega(t+\tau)}\mathrm{d}\omega \frac{1}{\sqrt{2\pi}}\int_0^\infty \psi^*_{\omega'}(\mathbf{r}_2)e^{\mathrm{i}\omega' t}\mathrm{d}\omega' \right\rangle && (2.62)\\
&= \int_0^\infty\int_0^\infty \left\langle \frac{1}{2\pi}e^{-\mathrm{i}(\omega-\omega')t}\right\rangle \langle \psi_\omega(\mathbf{r}_1)\psi^*_{\omega'}(\mathbf{r}_2)\rangle\, e^{-\mathrm{i}\omega\tau}\mathrm{d}\omega'\mathrm{d}\omega && (2.63)\\
&= \int_0^\infty\int_0^\infty \delta(\omega-\omega')\, \langle \psi_\omega(\mathbf{r}_1)\psi^*_{\omega'}(\mathbf{r}_2)\rangle\, e^{-\mathrm{i}\omega\tau}\mathrm{d}\omega'\mathrm{d}\omega && (2.64)\\
&= \int_0^\infty \langle \psi_\omega(\mathbf{r}_1)\psi^*_{\omega}(\mathbf{r}_2)\rangle\, e^{-\mathrm{i}\omega\tau}\mathrm{d}\omega && (2.65)\\
&= \int_0^\infty W_\omega(\mathbf{r}_1, \mathbf{r}_2)e^{-\mathrm{i}\omega\tau}\mathrm{d}\omega, && (2.66)
\end{aligned}$$

where the *cross-spectral density* W_ω [Wolf, 1982] is introduced and defined according to

$$W_\omega(\mathbf{r}_1, \mathbf{r}_2) = \langle \psi_\omega(\mathbf{r}_1)\psi^*_\omega(\mathbf{r}_2)\rangle. \tag{2.67}$$

Thus, the cross-spectral density is the FT of the mutual coherence function[1] with respect to the time difference τ. The intensity of a single monochromatic component is given by

$$I_\omega(\mathbf{r}) = W_\omega(\mathbf{r}, \mathbf{r}) = \langle |\psi_\omega(\mathbf{r})|^2\rangle \tag{2.68}$$

and the total intensity is given by

$$I(\mathbf{r}) = \int_0^\infty I_\omega(\mathbf{r})\mathrm{d}\omega \tag{2.69}$$

[1] In the above definition the Fourier integrals extends only over the positive frequencies similar to the spectral decomposition in Eq. (2.2) and unlike the previously stated definition of the FT given by Eq. (2.26) and Eq. (2.27).

which follows from the fact that different monochromatic components are mutually incoherent. [Carter, 1993, Paganin, 2006, Wolf, 2007]

As a special case, for strictly monochromatic wavefields, and further assuming that the wavefield ensemble consists of only plane wavefields, the intensity at the plane $z = 0$ can be written as

$$I_\omega(\mathbf{r}_\perp, 0) = \langle |\psi^\theta_\omega(\mathbf{r}_\perp, 0)|^2 \rangle_\theta. \tag{2.70}$$

Here, the ensemble average denoted by $\langle\rangle_\theta$ averages over all possible fields defined by their wavevector $\mathbf{k}_\theta$, where θ indicates the angle of the wavevector with respect to the optical axis. Assuming purely parallel rays, such an averaging can be expressed as a convolution according to

$$I_\omega(\mathbf{r}_\perp, 0) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} w_\omega(\mathbf{r}'_\perp)|\psi_\omega(\mathbf{r}_\perp - \mathbf{r}'_\perp, 0)|^2 \mathrm{d}\mathbf{r}'_\perp, \tag{2.71}$$

where w_ω defines the abundance of the respective wavefields. This resembles blurring due to the extent of the source. [Pfeiffer et al., 2005, Gureyev et al., 2009, Beltran et al., 2018]

3. Propagation-based phase-contrast imaging (PBI)

In the following, the underlying principles of PBI are outlined and the most commonly used assumptions and methods to retrieve information about the phase-shifting properties of the sample are detailed.

Figure 3.1 illustrates the image formation in PBI by showing the intensity profiles of the X-ray wavefield, propagating from left to right, at different stages along the optical axis. In addition, the notation used within this chapter is introduced. The intensity of the incoming wavefield, created by an idealized X-ray source at infinity, is assumed to have constant illumination, described by the scalar I_0. It is further assumed that the wavefield is monochromatic. To simplify notation, the subscript ω used in the previous chapter to indicate a single spectral component is omitted in the following. The sample is defined by the three-dimensional distribution of its linear attenuation coefficient $\mu(\mathbf{r})$ and refractive index decrement $\delta(\mathbf{r})$. The interaction of the sample results in a lower amplitude and shifted phase of the wavefield. Directly behind the sample in the object plane the intensity of the wavefield $I(\mathbf{r}_\perp, 0)$ is solely influenced by the attenuating properties of the sample, resulting in strongest attenuation where the projection of the linear attenuation coefficient of the sample along the propagation direction is largest. However, with increasing propagation distance, the phase shifts induced by the sample lead to increasingly prominent self-interference effects that manifest in the intensity profile. These effects are predominantly present at the edges of the sample, where there is a significant change in phase between neighboring regions of the wavefield. The intensity profile $I(\mathbf{r}_\perp, z)$ at some distance behind the sample is then measured using an X-ray detector.

The core idea behind PBI is to deduce information about the phase-shifting properties of the sample from the measured interference effects. This procedure is referred to as phase retrieval. As discussed in the introduction, in general, it is not possible to recover the attenuating properties as well as the phase-shifting properties independently from a single measurement and acquiring multiple measurements at different positions along the optical axis or at different energies is difficult using laboratory sources. Thus, algorithms which deduce information about the phase-shifting properties of the sample from a single measurement require additional assumptions about the sample, but are also particularly suitable for tomography. The following sections outline the two most widely used assumptions and the various phase-retrieval algorithms that build upon these assumptions. First, the special case of a pure phase object will be discussed and second, the homogeneity assumption and its implications are outlined in detail. [Bremmer, 1952, Bronnikov, 1999, Paganin et al., 2002]

A large variety of phase-retrieval algorithms have been proposed, suitable for different samples and different imaging regimes. These include different Fourier methods derived from the Fresnel diffraction integrals using various approximations, such as the Born approximation, the Rytov approximation [Burvall et al., 2011] or phase-retrieval algorithms that use the CTF formulation [Langer et al., 2008], which will not be covered.

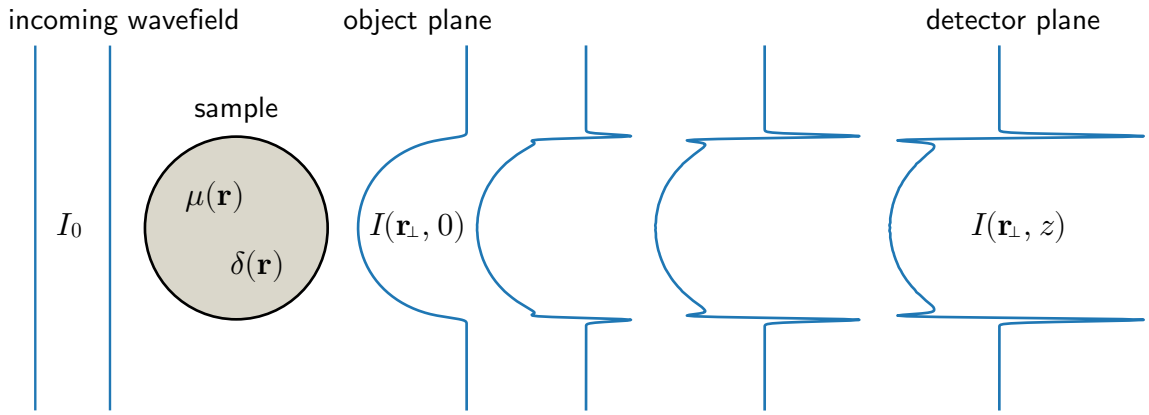

Figure 3.1.: Illustration of the evolution of the wavefield intensity in PBI. The incoming wavefield from the left is defined by its intensity I_0. The cylindrical sample is described by its linear attenuation coefficient $\mu(\mathbf{r})$ and refractive index decrement $\delta(\mathbf{r})$. The intensity behind the sample in the object plane is denoted by $I(\mathbf{r}_\perp, 0)$ and in the detector plane by $I(\mathbf{r}_\perp, z)$.

3.1. Pure phase objects

The first method derived in the context of PB-CT assumes pure phase objects. Pure phase objects do not attenuate the intensity of the X-rays, thus requiring $\mu(\mathbf{r}) = 0$. Consequently, the object is entirely defined by the three-dimensional distribution of the refractive index decrement $\delta(\mathbf{r})$. This can hold reasonably well for some biological samples or low-density materials measured at high X-ray energies, such that almost no attenuation contrast is present. These experiments are predominantly performed at synchrotron facilities. [Bronnikov, 1999, Bronnikov, 2002]

The algorithm was derived for CT applications, combining tomography reconstruction and phase retrieval [Bronnikov, 1999]. Here, only the part that accounts for the phase retrieval is outlined. Although the original derivation starts from the so-called weak focusing condition [Cowley, 1995], which is similar to the TIE [Burvall et al., 2011], one can also employ Eq. (2.43) directly, according to

$$I(\mathbf{r}_\perp, z) = I(\mathbf{r}_\perp, 0) - \frac{z}{k}\nabla_\perp \cdot [I(\mathbf{r}_\perp, 0)\nabla_\perp \phi(\mathbf{r}_\perp, 0)] \tag{3.1}$$

$$= I_0 - \frac{z}{k}\left[\nabla_\perp I_0 \cdot \nabla_\perp \phi(\mathbf{r}_\perp, 0) + I_0 \nabla_\perp^2 \phi(\mathbf{r}_\perp, 0)\right] \tag{3.2}$$

$$= I_0 \left(1 - \frac{z}{k}\nabla_\perp^2 \phi(\mathbf{r}_\perp, 0)\right), \tag{3.3}$$

where $I(\mathbf{r}_\perp, 0) = I_0$ was used, which follows from $\mu(\mathbf{r}) = 0$ and Eq. (2.17) assuming uniform illumination for the incoming wavefield. As I_0 is a constant scalar, it follows $\nabla_\perp I_0 = 0$. This connection between intensity and phase was already described in [Bremmer, 1952].

This equation can be solved using Fourier decomposition according to Eq. (2.26) and Eq. (2.27) and employing the Fourier derivative theorem given by Eq. (2.31). Thereby, the term, which includes the phase, is isolated and subsequently, the phase is decomposed into its Fourier components $\mathcal{F}_\perp[\phi(\mathbf{r}_\perp, 0)]$ resulting in

$$\mathcal{F}_\perp\left[\frac{I(\mathbf{r}_\perp, z)}{I_0} - 1\right] = \frac{z}{k}\mathbf{k}_\perp^2 \mathcal{F}_\perp[\phi(\mathbf{r}_\perp, 0)]. \tag{3.4}$$

This equation can be solved for $\phi(\mathbf{r}_\perp, 0)$ by applying the IFT using

$$\boxed{\phi(\mathbf{r}_\perp, 0) = \mathcal{F}_\perp^{-1}\left[\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)/I_0 - 1\right]}{\frac{z}{k}\mathbf{k}_\perp^2}\right]} \tag{3.5}$$

and thus the phase shift in the object plane is recovered. [Bronnikov, 1999]

Due to the assumption of the sample being not attenuating, the applications of this algorithm are very limited. In practice, residual absorption corrupts the phase signal severely. In [Groso et al., 2006] it was shown that even if the absorption level is only 2%, the pure phase object assumption does not hold. They proposed to add an absorption dependent correction factor α in the denominator of the Eq. (3.5) resulting in

$$\boxed{\tilde{\phi}(\mathbf{r}_\perp, 0) = \mathcal{F}_\perp^{-1}\left[\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)/I_0 - 1\right]}{\frac{z}{k}\mathbf{k}_\perp^2 + \alpha}\right],} \tag{3.6}$$

which is referred to as the modified Bronnikov algorithm (MBA). [Groso et al., 2006]

Another approach tries to remove the phase effects on the projections and then reconstructs the linear attenuation coefficients instead. The phase effects $\tilde{\phi}$ are estimated using the MBA algorithm derived above and then a correction function is calculated according to

$$C(\mathbf{r}_\perp) = 1 - \gamma\tilde{\phi}(\mathbf{r}_\perp, 0), \tag{3.7}$$

where γ is a parameter to tune the strength of the correction. Finally, the intensity without the phase effects, which can be thought of as the intensity in the object plane as illustrated in Figure 3.1, is given by

$$I(\mathbf{r}_\perp, 0) = I(\mathbf{r}_\perp, z)/C(\mathbf{r}_\perp), \tag{3.8}$$

which ideally should only contain information about the attenuating properties of the sample. This algorithm is referred to as Bronnikov-aided correction (BAC) and is preferably applied in a laboratory environment using samples with prominent attenuating properties. [De Witte et al., 2009]

3.2. Homogeneity assumption and the single-material phase-retrieval algorithm

The *single-material phase-retrieval algorithm* is one of the most widely used phase-retrieval algorithms in PBI to recover phase information qualitatively as well as quantitatively. In the following, the underlying assumptions and extensions are discussed and the benefits of this approach with respect to noise are outlined. [Paganin et al., 2002]

3.2.1. Homogeneity assumption

The single-material phase-retrieval algorithm evolves around the so-called *homogeneity assumption*, where it is assumed that the imaged object consists only of one material. Such objects are sometimes also called *monomorphous*. Although in practice most samples consist of multiple materials, this assumption holds reasonably well for a wide variety of applications. This can be explained to some extent by the fact that, for light atoms or high energies, the interactions of X-rays with different materials can be accurately described in terms of

their electron densities, as the influence of the nuclei becomes less stringent. Thus, one reconstructs a single material of electron density. On the one hand, for most biological samples that consist mostly of light atoms (e.g. soft-tissue), the homogeneity assumption holds well in practice. On the other hand, for medical applications comparably high X-ray energies are required, which are able to penetrate the sample, making the homogeneity assumption again reasonable. [Thompson et al., 2019]

In the previous chapter, the interactions of the wavefield were described in terms of the refractive index decrement $\delta(\mathbf{r}_\perp, z)$ and linear attenuation coefficient $\mu(\mathbf{r}_\perp, z)$. The homogeneity assumption now couples these two quantities to the density $t(\mathbf{r}_\perp, z)$ according to

$$\mu(\mathbf{r}_\perp, z) = \mu t(\mathbf{r}_\perp, z) \qquad \text{and} \qquad \delta(\mathbf{r}_\perp, z) = \delta t(\mathbf{r}_\perp, z), \tag{3.9}$$

where δ and μ are now constant scalar quantities, which describe the phase-shifting and attenuating properties of the single material respectively. Thus, the number of unknowns is reduced by a factor of two. By knowing $t(\mathbf{r}_\perp, z)$, the above equations are sufficient to recover both the attenuating and phase-shifting properties of the sample. [Paganin et al., 2002]

The intensity and the phase shift of the wavefield behind the sample in the object plane are given by Eq. (2.17) and Eq. (2.18). Applying the homogeneity assumption using Eq. (3.9) results in

$$I(\mathbf{r}_\perp, 0) = I_0 \exp\left[-\mu \int t(\mathbf{r}_\perp, z)\mathrm{d}z\right] = I_0 e^{-\mu T(\mathbf{r}_\perp)} \tag{3.10}$$

$$\phi(\mathbf{r}_\perp, 0) = -k\delta \int t(\mathbf{r}_\perp, z)\mathrm{d}z = -k\delta T(\mathbf{r}_\perp) \tag{3.11}$$

assuming constant illumination of the incoming wavefield given by I_0. Here $z = 0$ denotes the object plane and the sample is located at $z < 0$. The *projected thickness* or *trace* of the sample is defined according to

$$\boxed{T(\mathbf{r}_\perp) = \int t(\mathbf{r}_\perp, z)\mathrm{d}z} \tag{3.12}$$

as the projection of the density along the X-ray paths. Consequently, if one is able to recover the projected thickness from various angles, the three-dimensional distribution can be recovered using tomographic methods, outlined in Chapter 4. [Paganin et al., 2002]

3.2.2. Single-material phase-retrieval algorithm

Although the original derivation [Paganin et al., 2002] starts from the TIE, as given by Eq. (2.53), one can also use Eq. (2.43) as a starting point written as

$$I(\mathbf{r}_\perp, z) = I(\mathbf{r}_\perp, 0) - \frac{z}{k}\nabla_\perp \cdot \left[I(\mathbf{r}_\perp, 0)\nabla\phi(\mathbf{r}_\perp, 0)\right], \tag{3.13}$$

which directly relates the intensity in the detector plane to the wavefield defined by its intensity and phase in the object plane. By enforcing the homogeneity assumption, one can substitute Eq. (3.10) and Eq. (3.11) in Eq. (3.13), which is sometimes referred to as the TIE-Hom in literature [Gureyev et al., 2017, Kitchen et al., 2017]. The derivation of the

single-material phase-retrieval algorithm follows as

$$I(\mathbf{r}_\perp, z) = I_0 e^{-\mu T(\mathbf{r}_\perp)} + z\delta\nabla_\perp \cdot \left(I_0 e^{-\mu T(\mathbf{r}_\perp)} \nabla_\perp T(\mathbf{r}_\perp)\right) \tag{3.14}$$

$$= I_0 e^{-\mu T(\mathbf{r}_\perp)} - z\frac{\delta}{\mu}\nabla_\perp^2 I_0 e^{-\mu T(\mathbf{r}_\perp)} \tag{3.15}$$

$$= \left(1 - z\frac{\delta}{\mu}\nabla_\perp^2\right) I_0 e^{-\mu T(\mathbf{r}_\perp)}, \tag{3.16}$$

where in the second line the chain-rule is applied in reverse. Notably, given $I(\mathbf{r}_\perp, 0) = I_0 e^{-\mu T(\mathbf{r}_\perp)}$, the intensity evolves linearly with the intensity at $z = 0$ and does not directly depend on the phase at $z = 0$. Due to the coupling of attenuating and phase-shifting properties, this term is already encoded in the intensity at the object plane. [Paganin et al., 2002, Beltran et al., 2018]

Fourier decomposition remains a powerful tool and is also used to solve this equation for the projected thickness explicitly. Therefore, the intensity at the detector plane is decomposed into its Fourier components according to $\mathcal{F}_\perp[I(\mathbf{r}_\perp, z)]$ using the FT defined by Eq. (2.26). In addition, the Fourier components of the intensity in the object plane are computed according to $\mathcal{F}_\perp[I_0 e^{-\mu T(\mathbf{r}_\perp)}]$. Finally, using the Fourier derivative theorem given by Eq. (2.31), Eq. (3.16) can be written according to

$$\mathcal{F}_\perp[I(\mathbf{r}_\perp, z)] = \left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) \mathcal{F}_\perp[I_0 e^{-\mu T(\mathbf{r}_\perp)})]. \tag{3.17}$$

Using the IFT and elementary manipulation, one can solve this for the projected thickness $T(\mathbf{r}_\perp)$ according to

$$\boxed{T(\mathbf{r}_\perp) = -\frac{1}{\mu}\log\left(\mathcal{F}_\perp^{-1}\left\{\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)\right]/I_0}{z\frac{\delta}{\mu}\mathbf{k}_\perp^2 + 1}\right\}\right),} \tag{3.18}$$

where the terms in the denominator were flipped to more closely resemble literature. Apart from the scaling factor μ^{-1}, the phase-retrieval algorithm is entirely defined by the factor $z\delta/\mu$. If this factor is too large for the given material, the resulting trace will be smeared out. By contrast, if this factor is too small, residual fringes from the interference effects remain. [Paganin et al., 2002]

For qualitative phase retrieval, the linear attenuation coefficient can be set to $\mu = 1$ and the remaining parameter is defined according to $\xi = z\delta/\mu = z\delta$.

3.2.3. Noise considerations

Besides the fact that this phase-retrieval algorithm is computationally very efficient and stable with respect to noise, as it acts as a low-pass filter, the outstanding advantages of this method are related to its noise and resolution properties, which are outlined in the following using Figure 3.2. The top row illustrates the evolution of the intensity and signal for X-ray radiography and the lower row illustrates the same information in the context of PBI. The X-ray generation and the interaction of the wavefield with the sample are the same for both methods. For simplicity, it is assumed that the generated wavefields have constant illumination, as shown in (a) and (e) respectively. In both methods, the interactions of the wavefields with the sample lead to an attenuation in intensity and a shift in phase. In the

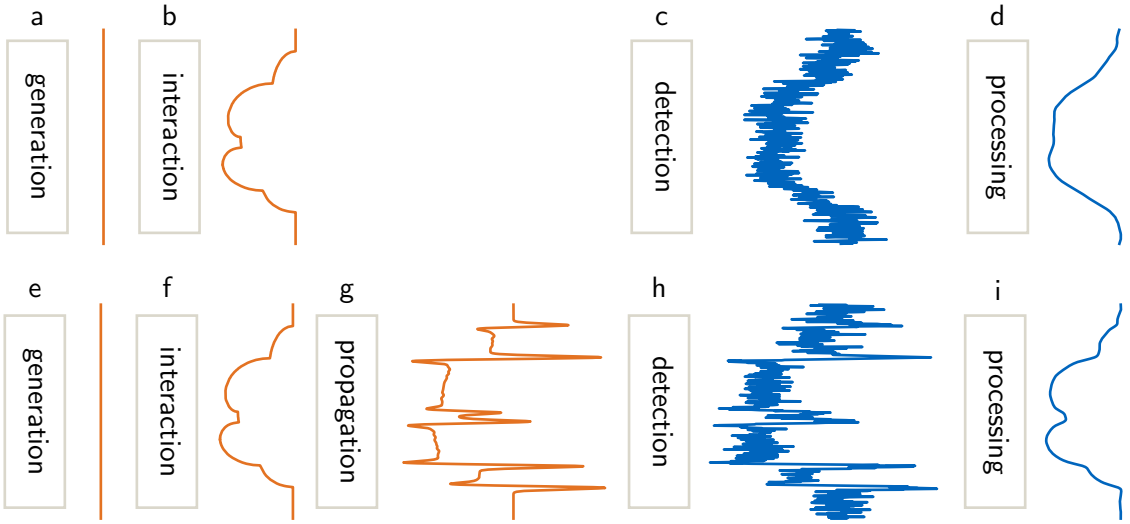

Figure 3.2.: Noise and resolution properties in PBI. On the top row, the evolution of the X-ray intensity and the detected signal is depicted for radiography. On the lower row, the same information for PBI is shown. In orange, the (undetected) X-ray intensity at different stages is shown and in blue, the detected intensity in form of the electronic signal is depicted. In radiography, the propagation part is omitted and the processing of the detected signal is optional.

object plane, the intensity of the wavefields is solely dependent on the attenuating properties of the sample as depicted in (b) and (f) respectively.

In conventional radiography, this intensity distribution is measured. However, in PBI the wavefield is subsequently propagating in free-space over a certain distance resulting in distinct variations of the wavefield's intensity, due to self-interference effects caused by the phase shifts induced by the sample. The resulting intensity distribution of the propagated wavefield is shown in (g). Compared to the intensity in conventional X-ray radiography, the high-frequency components are enhanced. Applying the single-material phase-retrieval algorithm to this intensity would decrease the high-frequency components to recover the traces that perfectly represent the intensity at the object plane. Consequently, the intensity measured in conventional radiography would be recovered without any advantages.

The key advantage of PBI over conventional radiography is with respect to noise inherent to measuring the intensity. However, the detection process and the noise properties cannot be understood using scalar wave theory. This requires the wavefield to be described quantum mechanically. Thereby, during the detection process (assuming an ideal detector), the quantum mechanical wavefield collapses and the measured photons are distributed according to the Poisson distribution. In essence, the number of detected photons is distributed around the corresponding value of the intensity of the classical wavefield with a variance equal to this intensity. The detected signals are depicted in (c) and (h) respectively. As the number of photons is conserved during free-space propagation, the noise level can be considered very similar. However, applying the single-material phase-retrieval algorithm for processing the signal in (h) recovers the trace, which describes the underlying intensity profile at the object plane. In that sense, the resolution properties of the measured signal in conventional radiography shown in (c) and the processed signal in (i) are identical. However, the noise level

of (i) is highly reduced compared to (c). This reduction in noise level is dependent on the frequency components, which are more prominently reduced for high-frequency components, which is evident from the denominator of the single-material phase-retrieval filter shown in Eq. (3.18). Similarly, one could apply the single-material phase-retrieval algorithm to reduce the noise level in (c) as shown in (d). Although the noise level of each component is the same as for (i), the resolution properties of the underlying sample are reduced. This, however, is not the case in PBI as these frequencies are enhanced due to free-space propagation beforehand.

In summary, the fact that during free-space propagation only the high frequencies of the underlying true signal are enhanced independently of noise and the single-material phase-retrieval algorithm allows for reverting this processes, the noise inherent to the detection process can be reduced without compromising the resolution of the underlying signal. This is somewhat contradictory from a signal processing point of view, but this purely quantum mechanical effect allows for a reduction in noise without compromising resolution.

This effect has recently been demonstrated experimentally in [Kitchen et al., 2017] and thoroughly analyzed mathematically in [Gureyev et al., 2017], based on related work analyzing the image quality in PBI [Gureyev et al., 2014a, Nesterets and Gureyev, 2014] and investigating the trade-off between noise and spatial resolution in general linear shift-invariant systems [Gureyev et al., 2014b].

3.2.4. Extensions of the single-material phase-retrieval algorithm

Multiple applications exist, where the single-material phase-retrieval algorithm performs poorly because the object is not well approximated by a single material. These applications include biological samples that consist of bone and soft-tissue, like for instance, small animals. If one assumes the sample to consist only of soft-tissue, the bones will get smeared out as $\delta_{\text{bone}}/\mu_{\text{bone}} \ll \delta_{\text{soft-tissue}}/\mu_{\text{soft-tissue}}$. Consequently, the filtering strength is too strong for bone. On the other hand, if one assumes the sample to only consist of bone, residual edge-enhancement will be present around soft-tissue. In the following, a few extensions to alleviate this problem are highlighted.

The single-material phase-retrieval algorithm was extended to multi-material objects which are embedded into another object. Considering the object is made of j materials with projected thickness $T_j(\mathbf{r}_\perp)$ inside the material 1 with projected thickness $T_1(\mathbf{r}_\perp)$ and introducing the total projected thickness $A(\mathbf{r}_\perp) = T_j(\mathbf{r}_\perp) + T_1(\mathbf{r}_\perp)$, the extended single-material phase-retrieval algorithm is given by

$$T_j(\mathbf{r}_\perp) = -\frac{1}{\mu_j - \mu_1} \log\left(\mathcal{F}_\perp^{-1}\left\{\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)\right] / \left(I_0 e^{-\mu_1 A(\mathbf{r}_\perp)}\right)}{z\frac{\delta_j - \delta_1}{\mu_j - \mu_1}\mathbf{k}_\perp^2 + 1}\right\}\right). \tag{3.19}$$

The biggest drawback of this method is that the reconstructions obtained from reconstructing each pair of material have to be sliced together manually. This approach has given good results, for instance, for imaging mice lungs. [Beltran et al., 2010, Beltran et al., 2011]

Other approaches to improve image quality from multi-material objects rely on merging reconstructions obtained without phase retrieval and those with phase retrieval in order to leverage the contrast added by phase retrieval and the resolution from the reconstruction without phase retrieval, which however is not suitable for quantitative imaging. [Irvine et al., 2014]

For CT applications, methods exist that first perform tomographic reconstruction and then apply a phase-retrieval filter on the volume. The advantage of this approach is that within the volume different materials are better separated than on the projections. As in the previous approach, multiple filters with different strengths are applied and then manually sliced together on the volume. However, interchanging the order of phase-retrieval and tomographic reconstruction is not valid, due to the non-linearity of the phase-retrieval step. In essence, these methods rely on the fact that for the given scenario, the logarithm in the phase-retrieval step can be linearized. The authors acknowledge that this is relatively crude, but have shown that it can work well in practice. [Ullherr and Zabler, 2015, Thompson et al., 2019]

4. Fundamentals of tomographic reconstruction

The underlying idea of conventional X-ray CT is to reconstruct the three-dimensional distribution of the linear attenuation coefficient from a set of radiographs obtained at various angles around the sample. However, the measured intensities are only sensitive to the projection of the linear attenuation coefficients along the X-ray paths according to Eq. (2.17). Consequently, the fundamental problem of CT is to reconstruct the three-dimensional distribution of a quantity from its projections. This mathematical problem was first described in detail by Johann Radon in 1917 [Radon, 1917], despite earlier work dating back to Hendrik Antoon Lorentz. [Buzug, 2008]

In the first section, the most-widely used analytical tomographic reconstruction formula that connects the projections of a quantity to its interior distribution is derived using continuous functions. The second section describes the transition from continuous functions to discrete quantities. Lastly, iterative reconstruction approaches are discussed, which are derived from a MAP principle and have become an active field of research promising improved image quality or lower acquisition time and dose requirements [Nuyts et al., 2013].

4.1. Analytical reconstruction

Various textbooks exist that cover analytical tomographic reconstruction techniques, such as [Natterer, 1986, Kak and Slaney, 1988, Buzug, 2008, Fessler, 2009, Zeng, 2010]. However, there are various conventions to define the coordinate systems and parameters. In the following, the conventions of the open-source Reconstruction Toolkit (RTK) are adapted, which are detailed in [Rit et al., 2014] and are based on the international standard IEC 61217 [Commission, 2008].

For simplicity, the derivation of the analytical algorithm is demonstrated for a two-dimensional parallel-beam geometry. Thereby, the X-ray source moves in a plane on a circular trajectory around the object in a sufficiently large distance such that the X-ray paths are in parallel with respect to each other. This is referred to as the parallel-beam geometry. Extensions to more realistic geometries are discussed afterwards.

To illustrate the concepts of tomographic reconstruction, the Shepp-Logan phantom is used, which was first introduced in 1974 [Shepp and Logan, 1974a, Shepp and Logan, 1974b]. The phantom consists of 10 ellipses, specified by their size and magnitude (value) and has been widely used for bench-marking reconstruction algorithms [Andersen and Kak, 1984, Kak and Slaney, 1988, Katsevich, 2002]. Nowadays, this phantom is usually discouraged for most simulation studies due to its limited complexity.

4.1.1. Forward- and backprojection

Let $f(x, y)$ be a two-dimensional (scalar) function, which represents the quantity to be reconstructed. The projections of this quantity $p(s, \theta)$ are the integrals along straight lines L_i

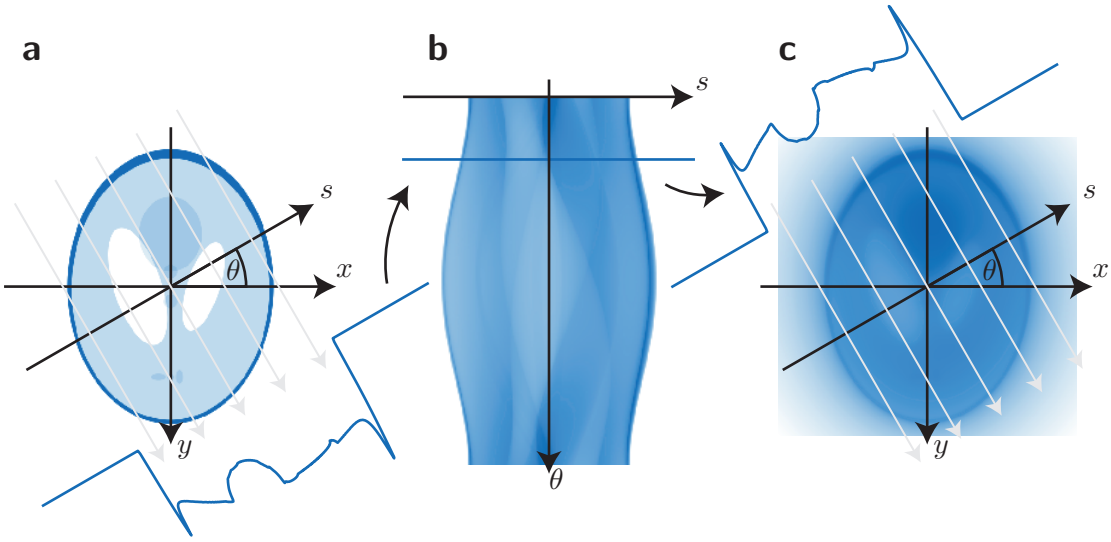

Figure 4.1.: Radon transform, sinogram representation and backprojection. In (a) the Radon transform is illustrated. The values of the Shepp-Logan phantom are integrated along the gray lines under a certain angle resulting in the blue projection. In (b) the projections for all angles are depicted. This representation is referred to as a sinogram. The backprojection is illustrated in (c). Thereby, the line-integrals are smeared out along the gray lines under the respective angles.

under different angles around the object. As illustrated in Figure 4.1 (a), a one-dimensional projection parameterized by s is obtained under each angle θ. Mathematically, the projections are given by

$$p(s,\theta) = \int_{L_i} f(x,y)\mathrm{d}l \tag{4.1}$$

$$= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\delta(x\cos\theta - y\sin\theta - s)\mathrm{d}x\mathrm{d}y. \tag{4.2}$$

This mapping of a function $f(x,y)$ to its projections $p(\theta,s)$ is called *(forward)projection* or *Radon transform* in honor of [Radon, 1917]. The angle θ is defined such that for $\theta = 0$, the s-coordinate of the projection coincides with the x coordinate of the sample according to $p(s,\theta=0) = \int f(s,y)\mathrm{d}y$. The representation of the projections according to Figure 4.1 (b) is called *sinogram*. [Zeng, 2010]

The transpose operation of the Radon transform is the so-called *backprojection* defined by

$$b(x,y) = \int_0^{\pi} p(x\cos\theta - y\sin\theta,\theta)\mathrm{d}\theta \tag{4.3}$$

and illustrated in Figure 4.1 (c). Thereby, the individual projections are smeared out along the corresponding straight lines and accumulated for each angle. The backprojection is not the inverse of the forwardprojection and thus, the backprojection alone cannot be used to recover $f(x,y)$ from its projections $p(s,\theta)$. However, the backprojection plays a crucial role for almost all reconstruction algorithms. [Zeng, 2010]

4.1.2. Filtered backprojection

The *Fourier slice theorem* or *central slice theorem* is the core concept of analytical tomographic reconstruction. It states that the one-dimensional FT[1] $P(\nu,\theta)$ of a projection $p(s,\theta)$ along s is equal to a one-dimensional profile through the origin of the two-dimensional FT $F(\nu,\theta)$ of $f(x,y)$. The proof follows by Fourier transforming the projections given by Eq. (4.2) and rearranging the order of integration according to

$$\begin{aligned} P(\nu,\theta) &= \int_{-\infty}^{\infty} p(s,\theta) e^{-2\pi \mathrm{i}\nu s} \mathrm{d}s && (4.4)\\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\delta(x\cos\theta - y\sin\theta - s) e^{-2\pi\mathrm{i}\nu s}\mathrm{d}x\mathrm{d}y\mathrm{d}s && (4.5)\\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y) e^{-2\pi\mathrm{i}\nu(x\cos\theta - y\sin\theta)}\mathrm{d}x\mathrm{d}y && (4.6)\\ &= F(u,v) && (4.7)\\ &= F_\circ(\nu,\theta), && (4.8) \end{aligned}$$

where $u = \nu\cos\theta$ and $v = -\nu\sin\theta$ are defined as the spatial frequencies in Cartesian coordinates corresponding to x and y. To emphasize that the two-dimensional FT is given in polar coordinates, a corresponding subscript is added. [Fessler, 2009, Zeng, 2010].

In principle, $f(x,y)$ can be recovered by a two-dimensional IFT of $P(\nu,\theta)$ directly (without the need for a backprojection). This is referred to as the *direct Fourier reconstruction* [Fessler, 2009]. However, prior to the two-dimensional IFT the polar representation of $F_\circ(\nu,\theta)$ has to be converted to a Cartesian grid to give $F(u,v)$ as the IFT requires the data to be lying on a regular rectangular grid. This approach was first proposed in the context of electron microscopy [De Rosier and Klug, 1968]. This gridding (or regridding [Buzug, 2008]) step from polar to Cartesian coordinates is very delicate and has been addressed in various papers [O'Sullivan, 1985, Schomberg and Timmer, 1995]. Most notably, employing non-uniform fast Fourier transform (FFT)s (NUFFT) methods with good interpolation kernels that directly yield $F(u,v)$ from the projections $p(s,\theta)$ has shown some success [Fourmont, 2003]. After the gridding step a two-dimensional IFT is sufficient to reconstruct $f(x,y)$. In practice, these methods are not widely used, mostly because the gridding causes interpolation artifacts [Fessler, 2009].

The de-facto standard analytical method for reconstructing tomographic data is referred to as the FBP algorithm. The derivation evolves around the Fourier slice theorem, where the coordinate transformation from polar to Carthesian coordinates is performed analytically. This leads to

$$\begin{aligned} f(x,y) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F(u,v) e^{2\pi\mathrm{i}(ux+vy)}\mathrm{d}u\mathrm{d}v && (4.9)\\ &= \int_{0}^{\pi}\int_{-\infty}^{\infty} F_\circ(\nu,\theta) e^{2\pi\mathrm{i}\nu(x\cos\theta - y\sin\theta)}|\nu|\mathrm{d}\nu\mathrm{d}\theta && (4.10)\\ &= \int_{0}^{\pi}\int_{-\infty}^{\infty} P(\nu,\theta) e^{2\pi\mathrm{i}\nu(x\cos\theta - y\sin\theta)}|\nu|\mathrm{d}\nu\mathrm{d}\theta && (4.11)\\ &= \int_{0}^{\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} p(s,\theta) e^{2\pi\mathrm{i}\nu(x\cos\theta - y\sin\theta - s)}\mathrm{d}s|\nu|\mathrm{d}\nu\mathrm{d}\theta, && (4.12) \end{aligned}$$

[1] The definition of the FT differs from the FT used in the previous chapters to be consistent with literature. Therefore, the precise form of the FT within this chapter is always explicitly denoted.

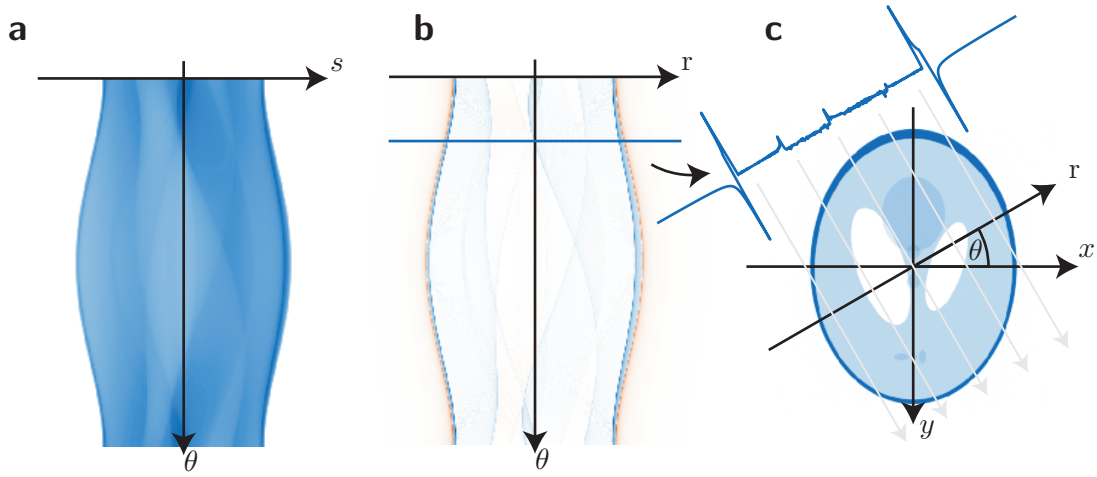

Figure 4.2.: Illustration of the FBP. In (a) the sinogram is depicted. After filtering, the high-frequencies of the sinogram are emphasized (red corresponds to negative values) as shown in (b). Finally, the filtered projections are backprojected under their respective angles as illustrated in (c) to yield the reconstructed quantity.

where the additional factor $|\nu|$ comes from the absolute value of the determinant of the Jacobian of the coordinate transformation from Cartesian coordinates (u, v) to polar coordinates (ν, θ). [Zeng, 2010]

The integration over s can be identified as a convolution with a filter h according to

$$q(r,\theta) = \int_{-\infty}^{\infty} p(s,\theta) \left[\int_{-\infty}^{\infty} e^{2\pi \mathrm{i}\nu(r-s)} |\nu| \mathrm{d}\nu \right] \mathrm{d}s \tag{4.13}$$

$$= \int_{-\infty}^{\infty} p(s,\theta) h\left(r-s\right) \mathrm{d}s, \tag{4.14}$$

where the convolution kernel is defined according to

$$h(s) = \int_{-\infty}^{\infty} e^{2\pi \mathrm{i}\nu s} |\nu| \mathrm{d}\nu. \tag{4.15}$$

This is referred to as the filtering step, within the FBP. The remaining integration over θ then denotes the backprojection step according to

$$f(x,y) = \int_0^{\pi} q(x\cos\theta - y\sin\theta, \theta) \mathrm{d}\theta, \tag{4.16}$$

which coincides with Eq. (4.3). The different steps of the FBP algorithm are illustrated in Figure 4.2. [Zeng, 2010]

4.1.3. Extensions to other geometries

In the fan-beam geometry, a point-like X-ray source moves at some finite distance around the sample on a circular trajectory inside the plane. As a consequence, the respective X-ray paths are not parallel with respect to each other. To obtain a reconstruction algorithm for this geometry, one converts the fan-beam geometry into a parallel-beam geometry by accordingly

substituting the parameterization of the projections. This results in slight modifications of the filtering as well as the backprojection step. [Turbell, 2001, Zeng, 2010, Hehn, 2015]

In practice, one is interested in the three-dimensional distribution of the linear attenuation coefficient. Thus, under each view, a two-dimensional projection is obtained. The extension of the parallel-beam geometry to three dimensions is trivial as one can think of this geometry as a stack of independent two-dimensional parallel-beam geometries. The extension of the fan-beam geometry to three dimensions is referred to as the cone-beam geometry. As for the fan-beam case, a point-like source moves at some finite distance around the object on a circular trajectory in a plane. In general, one cannot reconstruct the three-dimensional distribution of the linear attenuation coefficient using arbitrary geometries. However, there exists a set of conditions, which the source trajectory has to fulfill in order to reconstruct the quantity correctly, which is referred to as *Tuy's criterion* [Tuy, 1983]. An equivalent formulation states that every plane that intersects the object must contain a cone-beam focal point [Zeng, 2010]. For the cone-beam geometry defined above, this criterion is only fulfilled for the plane intersecting the object, in which the source is moving around the object in a circular trajectory. Thus, only approximate reconstruction algorithms exist in this case. The de-facto standard reconstruction algorithm for this geometry is the so-called Feldkamp, Davis, Kress (FDK) algorithm [Feldkamp et al., 1984]. Thereby, the formula for the fan-beam geometry is heuristically extended to three dimensions. For the central plane, the algorithm reduces to the fan-beam algorithm. For increasing distances between the source and the sample, the algorithm reduces to the reconstruction algorithm of the three-dimensional parallel-beam geometry. Finally, if the object is constant along the additional dimension, the Feldkamp algorithm obtains the correct result [Feldkamp et al., 1984]. Implementation details of the algorithm used in this thesis can be found in [Fehringer et al., 2014, Hehn, 2015, Fehringer, 2019]. For clinical applications spiral CT is most commonly used, for which the X-ray source moves along a helical trajectory. If the object lies entirely inside the source trajectory, Tuy's criterion is fulfilled and the object can be reconstructed exactly. An exact inversion formula exists [Katsevich, 2002]. This approach has also been extended to more general trajectories [Pack et al., 2004, Ye and Wang, 2005]. In practice, approximate algorithms are used, which are based on pre-processing steps such as rebinning or weighting of the projection [Schaller et al., 2000, Fuchs et al., 2000].

4.2. Discretization

For an actual X-ray experiment, the continuous function representing the projections of the three-dimensional distribution of the linear attenuation coefficient is not accessible. The amount of measured intensity values is restricted by the number of pixels of the X-ray detector and the finite number of projections around the sample. In addition, numerical methods only allow for the manipulation of a finite amount of data. Thus, the continuous functions and variables encountered in the previous section need to be discretized in order to be implemented numerically. This processes has to be carried out carefully as it comes with many subtleties that are addressed in this chapter. [Gonzalez and Woods, 2001, Seibt, 2006]

4.2.1. Discretization of measurement and volume

As the intensity measurements are discrete by default, the projections of the linear attenuation coefficients related to the intensities according to Eq. (2.17) are discrete as well. The projections are represented by a vector $\mathbf{p}$, merging the spatial dimensions and projections

from different angles into a single dimension. Again, restricted to the two-dimensional CT problem discussed in the previous section, the discretization is defined by the number of pixels N_s with a sampling distance (pixelsize) of Δ_s and the number of projections by N_θ with an angular interval of Δ_θ.

In the following, criteria about the choice of these parameters are outlined. The optimal pixelsize Δ_s depends on the spatial resolution of the system, which includes the effects of the focal spot blur and the detection system. A practical rule-of-thumb connects the pixelsize to the full width at half maximum (FWHM) according to $\Delta_s = \mathrm{FWHM}/2$, which is motivated by the Nyquist–Shannon sampling theorem [Seibt, 2006]. The number of pixels N_s is then determined by the size of the sample to be measured. [Fessler, 2009]

A criterion for the number of projections can be derived from the Fourier slice theorem. Given the pixelsize of Δ_s for the projections $p(s,\theta)$, the highest frequency in $P(\nu,\theta)$ is given by $\nu_{\max} = 1/(2\Delta_s)$. Accounting for negative frequencies, the N_s sampling points are thus spaced by $\Delta_\nu = 1/(N_s\Delta_s)$. Now, the angular sampling Δ_θ can be chosen such that for $F_\circ(\nu,\theta)$ all sampling points are spaced by not more than this amount. The angular spacing is worst for the highest frequency, where the angular spacing is given by

$$\Delta_\theta = \frac{\Delta_\nu}{\frac{N_s}{2}\Delta_\nu} = \frac{2}{N_s}, \tag{4.17}$$

using the above assumption. Thus, the total number of angles over the interval of 180 ° is

$$N_\theta = \frac{\pi}{\Delta_\theta} = \frac{\pi}{2}N_s. \tag{4.18}$$

This derivation can for instance be found in slightly different variations in the cited literature. [Kak and Slaney, 1988, Fessler, 2009]

Although $f(x,y)$ is continuous in reality, only a discrete representation can be reconstructed numerically. Therefore, $f(x,y)$ is expanded in terms of some finite basis function $\chi_i(x,y)$ according to

$$f(x,y) = \sum_{i=1}^{N_x \times N_y} f_i \chi_i(x,y), \tag{4.19}$$

where $N_x \times N_y$ is the total number of coefficients f_i. The most common choice of $\chi_i(x,y)$ is the so-called voxel basis. This basis is 1 inside the i-th voxel and 0 everywhere else. In the two-dimensional case, the basis function can be written as

$$\chi_i(x,y) = \mathrm{rect}\left(\frac{x-x_i}{\Delta_x}\right)\mathrm{rect}\left(\frac{y-y_i}{\Delta_y}\right). \tag{4.20}$$

Thereby, (x_i, y_i) is the center of the i-th voxel and Δ_x and Δ_y are the width and height of the voxels respectively. [Fessler, 2000]

4.2.2. Projection operations

The two most crucial and computational most demanding operations in CT reconstruction are the projection operations, i.e. the forwardprojection given by Eq. (4.2) and the backprojection given by Eq. (4.3). To discretize the forwardprojection, the analytic representation

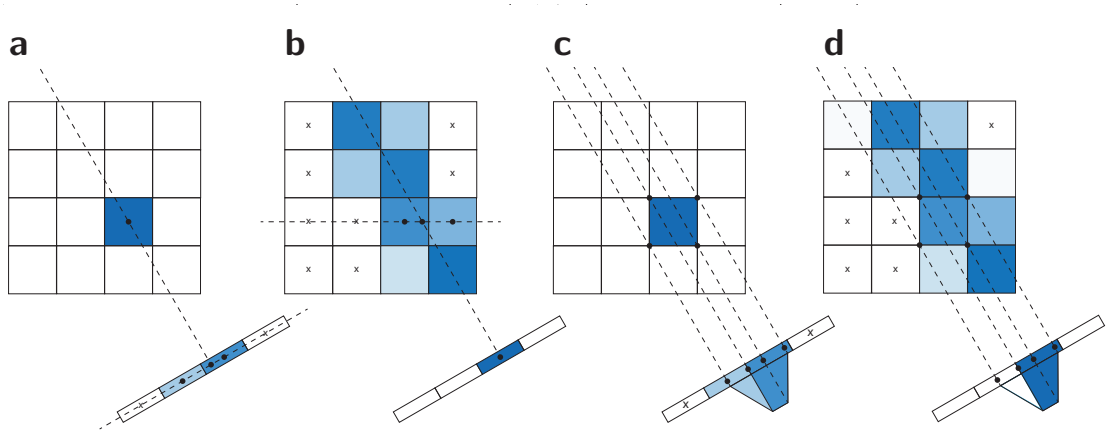

Figure 4.3.: Illustration of different projector models. In (a) the working principle of a pixel-driven projector with linear interpolation is depicted. In (b) the ray-driven projector with linear interpolation is illustrated. In (c) and (d) the working principle of the footprint projector is shown.

of f in terms of its finite basis function given by Eq. (4.19) is inserted into the analytical representation of the forwardprojection given by Eq. (4.3). This results in

$$p_i = \int_{L_i} f(x,y)\mathrm{d}l = \int_{L_i} \sum_{j=1}^{N_x \times N_y} f_j \chi_j(x,y)\mathrm{d}l = \sum_{j=1}^{N_x \times N_y} \int_{L_i} f_j \chi_j(x,y)\mathrm{d}l = \sum_j a_{ij} f_j, \tag{4.21}$$

where the order of summation and integration is interchanged and

$$a_{ij} = \int_{L_i} \chi_j(x,y)\mathrm{d}l \tag{4.22}$$

defines the matrix elements of the projection matrix. Thus, the forwardprojection given analytically by Eq. (4.2) can be written as a matrix-vector multiplication according to

$$\mathbf{p} = \mathbf{A}\mathbf{f}, \tag{4.23}$$

where the elements of $\mathbf{A}$ are defined by Eq. (4.22). [Fessler, 2009]

The discretization of the backprojection given by Eq. (4.3) follows from the fact that the backprojection is the transpose of the forwardprojection. Thus,

$$\mathbf{g} = \mathbf{A}^T\mathbf{p} \tag{4.24}$$

describes the backprojection. [Zeng, 2010]

In practice, the forwardprojection as well as the backprojection are not implemented as a matrix-vector multiplications. If 1000 projections around the object are acquired and the detector holds 1000×1000 pixels, the measurement vector $\mathbf{p}$ holds $1000^3 = 10^9$ elements. A natural choice for the number of voxels that represent the three-dimensional distribution of the linear attenuation coefficient would then be $1000^3 = 10^9$. Thus, the total number of elements a_{ij} would be 10^{18}. Assuming single-precision floating-point numbers, storing the matrix $\mathbf{A}$ explicitly would consume $10^{18} \cdot 4\,\mathrm{B} = 4\,\mathrm{EB}$, which cannot be stored on a computer. However, the fact that $\mathbf{A}$ is sparse (has many elements that are zero) makes it feasible to

calculate the matrix multiplication on the fly. This means that only those elements of a_{ij} that are non-zero are calculated and directly summed accordingly, without storing these entries explicitly. Consequently, only $\mathbf{p}$ and $\mathbf{f}$ have to be stored explicitly, which accounts in the previous case for only $2 \cdot 10^9 \cdot 4\,\mathrm{B} = 8\,\mathrm{GB}$. However, for each forwardprojection (or backprojection), the matrix elements have to be recomputed. [Fehringer, 2019]

The last step is to find a way to explicitly calculate the elements a_{ij} given analytically by Eq. (4.22). This integral is, however, impractical to compute. Several different projector models exist, which compute these elements differently, depending on the geometry and the requirements for accuracy and speed. [De Man and Basu, 2004, Long et al., 2010]

The projector models that are used within this thesis are illustrated in Figure 4.3 for a two-dimensional parallel-beam geometry. Thereby, the elements of the object are referred to as pixels and the elements of the detector as detector elements in order to be consistent with the nomenclature in literature. [De Man and Basu, 2004, Fessler, 2009]

In (a) the working principle of a simple pixel-driven projector is shown. Thereby, one connects a line from the source through the center of the pixel to the detector [Peters, 1981]. This method is most commonly used for implementing the backprojection. The value of the pixel is then obtained by linear interpolation according to the position of the projected center of the pixel on the detector. The backprojection can be implemented very efficiently as it can be parallelized for all pixels in the object. Using the same model for the forwardprojection is discouraged as it introduces high-frequency artifacts. In addition, the parallelization becomes more difficult as the values of different pixels need to be accumulated to the same detector element at the same time (output dependency). [De Man and Basu, 2004]

In (b) another projector model is illustrated, which is referred to as the ray-driven method. Thereby, the ray from the source is connected to the center of the detector element. This projector model, usually referred to as Joseph's method [Joseph, 1982], is suited for implementing the forwardprojection. Thereby, for every row within the object (or column depending on the projection angle) linear interpolation is performed and the values are added to the corresponding detector element. This forwardprojection can also be parallelized for every detector element and every projection angle and thus implemented very efficiently. Again, this model is discouraged for the transpose operation (in this case the backprojection) as it introduces Moiré patterns and suffers again from the output dependency. [De Man and Basu, 2004]

Most parts of this thesis use the extension to the cone-beam geometry of the aforementioned ray-driven projector model for forwardprojection and the above pixel-driven model for backprojection as they are computationally most efficient. More details about the implementation can be found in [Fehringer et al., 2014, Fehringer, 2019]. However, there are two drawbacks. One is that the projector models are not very accurate leading to artifacts for various scenarios. Secondly, as different models are used for forward and backprojection, these two operations are not consistent. This means that the forwardprojection is no longer the transpose of the backprojection. The implications of this mismatch are quite controversial [Zeng, 2010]. However, there is a criterion that can guarantee the convergence of iterative algorithms if the composite backprojection-projector matrix $\mathbf{A}^T\mathbf{A}$ does not have negative eigenvalues [Zeng and Gullberg, 2000]. A more recent study has shown that the reconstruction quality as well as the convergence of any iterative algorithm greatly relies on a good matching between the forward and backprojection [Arcadu et al., 2016].

To obtain matched projectors with an accurate projector model, a so-called footprint projector has been implemented for this thesis. The working principle is illustrated in Figure 4.3 (c) and (d). Thereby, four lines from the source through the edges of the object pixel are drawn to get the corresponding intersection points on the detector. From these points, the

actual shape of the projected pixel (its footprint) can be calculated. The contribution of the object pixel with each detector element is then deduced by the amount of overlap between the pixel and the true footprint of the projected pixel. As for the above pixel-driven projector, the backprojection can be computed for all pixels of the object in parallel, which again makes it very efficient to compute. The forwardprojection cannot be parallelized over all pixels due to the same output dependency as before. However, it is possible to parallelize the calculation for each projection angle, which makes this projector model sufficiently fast.

The extension to three dimensions is called the separable footprint projector model, which is both computationally efficient and accurate [Long et al., 2010]. In this case, the footprints in horizontal and vertical direction are calculated independently. Several publications have addressed efficient and highly parallelized implementations for this projector model [Wu and Fessler, 2011, Xie et al., 2017].

4.2.3. Filtered backprojection algorithm

In order to numerically implement the FBP algorithm, the filtering step given by Eq. (4.14) remains to be discretized. Therefore, the tomographic filter given by Eq. (4.15) is band-limited by the Nquist frequency $\nu_{\max} = 1/(2\Delta_s)$ discussed in Subsection 4.2.1. Consequently, the tomographic filter is given by

$$h_{\text{band-limited}}(s) = \int_{-\frac{1}{2\Delta_s}}^{\frac{1}{2\Delta_s}} e^{2\pi \mathrm{i}\nu s}|\nu|\mathrm{d}\nu = \frac{1}{2(\Delta_s)^2}\operatorname{sinc}\frac{2s}{2\Delta_s} - \frac{1}{4(\Delta_s)^2}\operatorname{sinc}^2\frac{s}{2\Delta_s}, \tag{4.25}$$

which reduces to

$$h_i = \Delta_s h_{\text{band-limited}}(i\Delta_s) = \begin{cases} \frac{1}{4\Delta_s}, & i = 0 \\ 0, & i \text{ even} \\ -\frac{1}{i^2\pi^2\Delta_s}, & i \text{ odd} \end{cases}. \tag{4.26}$$

Moreover, the convolution in Eq. (4.14) has to be discretized. A convolution can be expressed as a matrix multiplication with a Toeplitz matrix, thus the filtering step can be written as

$$\mathbf{q} = \mathbf{H}\mathbf{p}. \tag{4.27}$$

Finally, the backprojection step given by

$$\hat{\mathbf{f}} = \mathbf{A}^T\mathbf{q} \tag{4.28}$$

enables the reconstruction of the discrete object from its discrete projections. [Turbell, 2001, Fessler, 2009]

4.3. Statistical iterative reconstruction (SIR)

Various iterative tomographic reconstruction approaches have been developed ever since the success of CT [Qi and Leahy, 2006, Nuyts et al., 2013]. Early iterative methods are algebraic reconstruction techniques (ART) [Gordon et al., 1970]. Later, variations like the simultaneous iterative reconstruction technique (SIRT) [Gilbert, 1972] or the simultaneous ART (SART) [Andersen and Kak, 1984] were developed. These algorithms provide means to solve Eq. (4.23) iteratively. However, these techniques only converge if the data is consistent,

such that an exact solution exists, which is however commonly not the case due to noise (and other inconsistencies) in the measurements [Qi and Leahy, 2006]. With the introduction of the FBP algorithm, iterative approaches became less popular. [Feldkamp et al., 1984, Zhang et al., 2018].

Nowadays with the increasing performance of modern computer systems, iterative reconstruction techniques are again heavily researched. While previously only the projection operation was modeled by the matrix $\mathbf{A}$ and all rays were treated equally, including a statistical description of the noise and modeling additional aspects of the imaging system have gained a lot of attention. These algorithms are referred to as statistical iterative reconstruction (SIR) techniques. Such algorithms have shown significant advantages in terms of lowering dose requirements and improving image quality. All major CT vendors have implemented some kind of SIR method for image reconstruction. Various names for SIR methods exist in literature and throughout different vendors, such as model-based iterative reconstruction (MBIR), iterative model reconstruction (IMR), advanced modeled iterative reconstruction (ADMIRE) or forward-projected model-based iterative reconstruction solution (FIRST). [Nuyts et al., 2013, Zhang et al., 2018]

Within this work, the proposed algorithms are referred to as SIR, if the novelty is from the reconstruction approach, and as MBIR, if the novelty stems from the physical models that describe the image formation process. In the following, the basic statistical description of an X-ray experiment is outlined followed by the derivation of the different parts of the SIR algorithm starting from the MAP principle.

4.3.1. Maximum a posteriori (MAP) estimation

One can think of a CT scan (or any X-ray imaging experiment) as a probabilistic experiment due to the statistical nature of measuring X-rays. Thereby, the measurements (the intensities measured at the detector) can be modeled by a *random vector* $\mathbf{Y}$. [Fessler, 2000]

Mathematically, an estimate of the three-dimensional distribution of the linear attenuation coefficient $\hat{\boldsymbol{\mu}}$ can be obtained by maximizing

$$\hat{\boldsymbol{\mu}} = \arg\max_{\boldsymbol{\mu}} f(\boldsymbol{\mu}|\mathbf{Y}), \tag{4.29}$$

where $f(\boldsymbol{\mu}|\mathbf{Y})$ denotes the *posterior distribution*. This distribution describes the probability of $\boldsymbol{\mu}$ being the actual linear attenuation coefficients given the measurements $\mathbf{Y}$. Thus, $\hat{\boldsymbol{\mu}}$ is called the MAP estimator. However, the actual form of the posterior distribution $f(\boldsymbol{\mu}|\mathbf{Y})$ is not directly accessible. Therefore, *Bayes' theorem* is employed given by

$$f(\boldsymbol{\mu}|\mathbf{Y}) = \frac{f(\mathbf{Y}|\boldsymbol{\mu})f(\boldsymbol{\mu})}{f(\mathbf{Y})}. \tag{4.30}$$

Here, $f(\mathbf{Y}|\boldsymbol{\mu})$ denotes the likelihood of the measurements that describes the probability of the measurements $\mathbf{Y}$ under the assumption of $\boldsymbol{\mu}$ being the linear attenuation coefficients. This likelihood can, in principle, be computed if one can model the image formation as well as the statistical properties of the measurements given the linear attenuation coefficients. Furthermore, $f(\boldsymbol{\mu})$ denotes the prior distribution that described how probable a certain realization of $\boldsymbol{\mu}$ independent of the measurements is. Finally, $f(\mathbf{Y})$ denotes the probability of the measurements independent of the linear attenuation coefficients. [Fessler, 2000]

Instead of maximizing the posterior distribution, one can equivalently minimize the negative log-likelihood of the posterior distribution, resulting in

$$\hat{\boldsymbol{\mu}} = \arg\min_{\boldsymbol{\mu}} \mathcal{C}(\boldsymbol{\mu}|\mathbf{Y}), \tag{4.31}$$

where the objective function $\mathcal{C}(\boldsymbol{\mu}|\mathbf{Y}) = -\log f(\boldsymbol{\mu}|\mathbf{Y})$ is defined. Using Eq. (4.30), the explicit form of the objective function is given by

$$\mathcal{C}(\boldsymbol{\mu}|\mathbf{Y}) = \mathcal{L}(\mathbf{Y}|\boldsymbol{\mu}) + \beta\mathcal{R}(\boldsymbol{\mu}), \tag{4.32}$$

where the negative log-likelihood $\mathcal{L}(\mathbf{Y}|\boldsymbol{\mu}) = -\log f(\mathbf{Y}|\boldsymbol{\mu})$, the regularization strength β and the regularizer $\mathcal{R}$ are defined according to $\beta\mathcal{R}(\boldsymbol{\mu}) = -\log f(\boldsymbol{\mu})$. As $f(\mathbf{Y})$ does not depend on $\boldsymbol{\mu}$, this term is omitted from the definition of the objective function, as it does not alter the estimate according to Eq. (4.31). [Fessler, 2000]

4.3.2. Likelihood terms

Firstly, the probability distribution $f(\mathbf{Y}|\boldsymbol{\mu})$ is defined. As discussed in Chapter 3, X-rays follow a Poisson distribution according to the expected intensity $\bar{\mathbf{y}}(\boldsymbol{\mu})$, which is discussed in the next subsection. The likelihood of a given measurement $\mathbf{y}$, which is a realization of the random vector $\mathbf{Y}$, is given by

$$f(\mathbf{Y} = \mathbf{y}|\boldsymbol{\mu}) = \prod_{i=1}^{p} \mathcal{P}(y_i|\bar{y}_i(\boldsymbol{\mu})) = \prod_{i=1}^{p} \frac{\bar{y}_i(\boldsymbol{\mu})^{y_i}}{y_i!} e^{-\bar{y}_i(\boldsymbol{\mu})} = \mathcal{P}(\mathbf{y}|\bar{\mathbf{y}}(\boldsymbol{\mu})). \tag{4.33}$$

Here, y_i denote the individual entries of $\mathbf{y}$. Following Eq. (4.31), one would compute the negative logarithm of this probability distribution, which omits the numerically unstable calculation of the products. [Fessler, 2000]

However, within this thesis, another probability distribution is mostly used, namely the normal distribution. For sufficiently high counts $\bar{y}_i(\boldsymbol{\mu})$, the Poisson distribution $\mathcal{P}$ can be approximated by a normal distribution $\mathcal{N}$ according to

$$\lim_{\bar{y}_i(\boldsymbol{\mu})\to\infty} \mathcal{P}(y_i|\bar{y}_i(\boldsymbol{\mu})) = \frac{1}{\sqrt{2\pi\bar{y}_i(\boldsymbol{\mu})}} \exp\left\{-\frac{(y_i - \bar{y}_i(\boldsymbol{\mu}))^2}{2\bar{y}_i(\boldsymbol{\mu})}\right\} = \mathcal{N}(y_i|\bar{y}_i(\boldsymbol{\mu}), \bar{y}_i(\boldsymbol{\mu})). \tag{4.34}$$

Consequently, the likelihood of Eq. (4.33) is given by

$$\begin{aligned}
f(\mathbf{Y} = \mathbf{y}|\boldsymbol{\mu}) &= \prod_{i=0}^{p} \mathcal{N}(y_i|\bar{y}_i(\boldsymbol{\mu}), \bar{y}_i(\boldsymbol{\mu})) && (4.35)\\
&= \frac{1}{\sqrt{(2\pi)^p \prod_{i=0}^{p} \bar{y}_i(\boldsymbol{\mu})}} \exp\left[-\sum_{i=0}^{p} \frac{(y_i - \bar{y}_i(\boldsymbol{\mu}))^2}{2\bar{y}_i(\boldsymbol{\mu})}\right] && (4.36)\\
&= \frac{1}{\sqrt{(2\pi)^p \det\left[\mathcal{D}(\bar{\mathbf{y}}(\boldsymbol{\mu}))\right]}} \exp\left[-\frac{1}{2}(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu}))^T \mathcal{D}^{-1}\left[\bar{\mathbf{y}}(\boldsymbol{\mu})\right](\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu}))\right] && (4.37)\\
&= \mathcal{N}(\bar{\mathbf{y}}(\boldsymbol{\mu})|\mathcal{D}(\bar{\mathbf{y}}(\boldsymbol{\mu}))), && (4.38)
\end{aligned}$$

where $\mathcal{D}$ denotes a diagonal matrix with the vector in brackets on its diagonal. The negative log-likelihood is then given by

$$\begin{aligned}
\mathcal{L}(\mathbf{y}|\boldsymbol{\mu}) &= -\log \mathcal{N}(\bar{\mathbf{y}}(\boldsymbol{\mu})|\mathcal{D}(\bar{\mathbf{y}}(\boldsymbol{\mu}))) && (4.39)\\
&= \frac{p}{2}\log[2\pi] + \frac{1}{2}\log\det\mathcal{D}\left[\bar{\mathbf{y}}(\boldsymbol{\mu})\right] + \frac{1}{2}(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu}))^T \mathcal{D}^{-1}\left[\bar{\mathbf{y}}(\boldsymbol{\mu})\right](\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu})) && (4.40)\\
&\equiv \frac{1}{2}\log\det\mathcal{D}\left[\bar{\mathbf{y}}(\boldsymbol{\mu})\right] + \frac{1}{2}(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu}))^T \mathcal{D}^{-1}\left[\bar{\mathbf{y}}(\boldsymbol{\mu})\right](\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu})), && (4.41)
\end{aligned}$$

with the term independent of $\boldsymbol{\mu}$ being omitted as indicated by the $\equiv$ symbol. Finally, by estimating the statistical weights from the measurements directly to simplify the optimization problem according to $\mathcal{D}\left[\bar{\mathbf{y}}(\boldsymbol{\mu})\right] \approx \mathcal{D}\left[\mathbf{y}\right]$, the negative log-likelihood reduces to

$$\mathcal{L}(\mathbf{y}|\boldsymbol{\mu}) = \frac{1}{2}\log\det\mathcal{D}\left[\mathbf{y}\right] + \frac{1}{2}(\mathbf{y}-\bar{\mathbf{y}}(\boldsymbol{\mu}))^T\mathcal{D}^{-1}\left[\mathbf{y}\right](\mathbf{y}-\bar{\mathbf{y}}(\boldsymbol{\mu})) \tag{4.42}$$

$$\boxed{\equiv \frac{1}{2}(\mathbf{y}-\bar{\mathbf{y}}(\boldsymbol{\mu}))^T\mathbf{W}(\mathbf{y}-\bar{\mathbf{y}}(\boldsymbol{\mu})),} \tag{4.43}$$

where $\mathbf{W} = \mathcal{D}^{-1}\left[\mathbf{y}\right]$ is defined and again the term independent of $\boldsymbol{\mu}$ is omitted. This formulation coincides with a penalized weighted least squares approach. [Zhang et al., 2018]

4.3.3. Physical mean models

The likelihood requires the computation of the expected intensity values defined as the expectation value of the measurements according to $\bar{\mathbf{y}}(\boldsymbol{\mu}) = \mathrm{E}[\mathbf{Y}]$. In the particle picture this coincides with the expected number of detected X-ray photons. In the following, $\bar{\mathbf{y}}(\boldsymbol{\mu})$ is referred to as the physical mean model, which models the image formation given the linear attenuation coefficients. Improving the physical mean model is still an active field of research [Nuyts et al., 2013]. Finding a suitable model strongly depends on the applications and setups. Various different forms of the basic mean model exist in literature [Fessler, 2009, Nuyts et al., 2013]. Here, the physical mean model matches [Gang et al., 2014] written as

$$\boxed{\bar{\mathbf{y}}(\boldsymbol{\mu}) = I_0 e^{-\mathbf{A}\boldsymbol{\mu}},} \tag{4.44}$$

where I_0 denotes the mean intensity without any sample in the beam. The exponential function operates element-wise. Thus, it can be thought of as a discrete version of the Lambert-Beer law given by Eq. (2.17). This model is suited for the likelihood terms in Eq. (4.33) as well as Eq. (4.43). [Fessler, 2000, Gang et al., 2014]

However, for many applications a different mean model is used, which acts on the line-integrals instead of the intensities [Nuyts et al., 2013]. The random vector transforms according to

$$\mathbf{L} = -\log\left(\frac{\mathbf{Y}}{I_0}\right). \tag{4.45}$$

Using first order Taylor approximation for the expectation value according to

$$\boxed{\bar{\boldsymbol{\ell}}(\boldsymbol{\mu}) = \mathrm{E}[\mathbf{L}] = \mathrm{E}[f(\mathbf{Y})] \approx f\left(\mathrm{E}[\mathbf{Y}]\right) = f\left(I_0 e^{-\mathbf{A}\boldsymbol{\mu}}\right) = \mathbf{A}\boldsymbol{\mu}} \tag{4.46}$$

defines the line-integrals as the projections of the linear attenuation coefficient. However, this expectation value is biased due to Jensen's inequality [Lehmann and Casella, 2006]. From the fact that $-\log x$ is a concave function it follows

$$\mathrm{E}[-\log X] \geq -\log\mathrm{E}[X], \tag{4.47}$$

which systematically overestimates the line-integrals [Fessler, 2000]. This has been verified empirically and furthermore it has been shown that this bias increases with decreasing counts [Fessler, 2000]. On the other hand, the variance is given by

$$\mathrm{Var}[\mathbf{L}] = \mathrm{Var}[f(\mathbf{Y})] \approx \frac{1}{\mathrm{E}[\mathbf{Y}]^2}\mathrm{Var}[\mathbf{Y}] = \frac{I_0 e^{-\mathbf{A}\boldsymbol{\mu}}}{\left(I_0 e^{-\mathbf{A}\boldsymbol{\mu}}\right)^2} = \frac{1}{I_0 e^{-\mathbf{A}\boldsymbol{\mu}}} = \mathcal{D}[\bar{\mathbf{y}}^{-1}] \approx \mathcal{D}[\mathbf{y}^{-1}], \tag{4.48}$$

where the division is to be understood element-wise. In the last step, the mean model was approximated by the measurements to make the variance independent of the linear attenuation coefficients. The likelihood is then modeled by $\mathcal{N}(\bar{\boldsymbol{\ell}}, \mathcal{D}[\mathbf{y}^{-1}])$ resulting in

$$\boxed{\mathcal{L}(\boldsymbol{\ell}|\boldsymbol{\mu}) = \frac{1}{2}(\boldsymbol{\ell} - \bar{\boldsymbol{\ell}}(\boldsymbol{\mu}))^T \mathcal{D}\left[\mathbf{y}\right](\boldsymbol{\ell} - \bar{\boldsymbol{\ell}}(\boldsymbol{\mu})),} \tag{4.49}$$

with $\boldsymbol{\ell} = -\log(\mathbf{y}/I_0)$. As the mean model as stated in Eq. (4.46) is linear, the optimization procedure of the likelihood is simpler compared to the previous case. It also coincides with a penalized weighted least squares formulation. [Fessler, 2000, Nuyts et al., 2013, Zhang et al., 2018]

4.3.4. Prior distributions

Finally, the prior distribution $f(\boldsymbol{\mu})$ in Eq. (4.30) has to be defined. Therefore, one assumes $\boldsymbol{\mu}$ to be drawn from a distribution of the form

$$f(\boldsymbol{\mu}) \propto e^{-\beta \mathcal{R}(\boldsymbol{\mu})}, \tag{4.50}$$

where β denotes the so-called regularization strength, which determines how much influence the prior information has on the posterior $f(\boldsymbol{\mu}|\mathbf{Y})$. Usually, this parameter has to be chosen by hand. The term $\mathcal{R}$ is referred to as the *regularizer*. A good choice of the regularizer is still an active field of research and depends on the particular application. Early regularization techniques assumed spatial independence among different voxels. For instance, as regularizer, one can use the ℓ_2 norm of the image directly or devise priors that are based on the maximum entropy criterion [Nunez and Llacer, 1990]. However, this biases the image tremendously and alleviating this by the introduction of a mean image makes these priors not longer truly spatially independent [Qi and Leahy, 2006]. [Zhang et al., 2018]

Most commonly used are Markov random field (MRF) priors, which assume that neighboring voxels have similar values. Mathematically, this can be described by a Gibbs distribution [Geman and Geman, 1984]. The regularization techniques used within this thesis belong to this class of priors. In particular, pair-wise Gibbs priors are a special case of Gibbs priors which only penalize the difference in neighboring voxel values. The corresponding regularization term can be written as

$$\boxed{\mathcal{R}(\boldsymbol{\mu}) = \frac{1}{2}\sum_i \sum_{n \in \mathcal{N}_i} \frac{1}{\Delta_{in}} \psi\left(\frac{\mu_i - \mu_n}{\Delta_{in}}\right),} \tag{4.51}$$

where $\mathcal{N}_i$ refers to the neighboring voxels of the i-th voxel (usually 26 next neighbors in three dimensions) and ψ is the penalty function. The additional factor Δ_{in} denotes the Euclidean distance between two voxels. When denoting the 3D coordinates of voxel i by (i_x, i_y, i_z) and the corresponding coordinates of the voxel n by (n_x, n_y, n_z), this distance can be written as

$$\Delta_{in} = \sqrt{(i_x - n_x)^2 + (i_y - n_y)^2 + (i_z - n_z)^2}. \tag{4.52}$$

In literature, the regularization term usually omits the additional weighting within the argument of the penalty function [Fessler, 2000, Zhang et al., 2018]. However, the optimal choice of Δ_{in} is still unclear. For instance, it was noted that the Euclidean distance is suboptimal for producing isotropic spatial resolution [Fessler, 2003]. [Zhang et al., 2018]

Various penalty functions exist in literature [Zhang et al., 2018]. Within this thesis, there are three different penalty functions, which are used depending on the application. First of all, the quadratic penalty, which uses the ℓ_2 norm, has the form

$$\psi(t) = t^2. \tag{4.53}$$

A disadvantage of this prior is that it leads to oversmoothing at boundaries between different materials and fine structures. This has a negative effect on the resolution of the reconstructed object. [Zhang et al., 2018]

One approach to alleviate this issue is the so-called Huber penalty [Huber, 2011] given by

$$\psi_\gamma(t) = \begin{cases} t^2/(2\gamma) & \text{for } t < \gamma \\ |t| - \gamma/2 & \text{else} \end{cases}. \tag{4.54}$$

Thereby, the penalty coincides (up to the normalization) with the quadratic penalty given by Eq. (4.53) if the weighted difference between neighboring voxels is smaller than γ. If this difference exceeds γ, the penalty function corresponds to a ℓ_1 norm (up to an offset). The reason is that one would like to mitigate the penalty for high differences in neighboring voxel values, implying that these differences are due to an edge in the sample and not due to noise. By contrast, small differences in neighboring voxel values are most probably due to noise and therefore penalized quadratically. [Zhang et al., 2018]

Finally, total-variation (TV) regularization [Rudin et al., 1992] uses the ℓ_1 norm as penalty. Thereby, the absolute difference of the weighted difference of neighboring voxel values is penalized according to

$$\psi(t) = |t|. \tag{4.55}$$

This penalty coincides (up to an offset) with the Huber penalty given by Eq. (4.54) for weighted differences that exceed γ. This function is still convex, however not differentiable at zero. TV regularization also plays a crucial role in the field of compressed sensing as the penalty enforces sparsity [Sidky et al., 2006, Sidky and Pan, 2008]. In compressed sensing, one tries to enforce sparse solutions. This can for instance be achieved by employing the ℓ_0 norm [Hu et al., 2011] or other ℓ_p norms with $0 < p < 1$ (which is actually not a norm) in the regularization term. However, the optimization problem becomes increasingly difficult as these norms are non-convex. [Zhang et al., 2018]

Finally, one may generalize the regularizer given by Eq. (4.51) and make the weights spatially varying. This can for instance be used to construct a regularization term that yields uniform resolution. [Stayman and Fessler, 2000]

Various other regularization techniques exist that cannot be formulated according to Eq. (4.51). For instance, nonlinear neighborhood filters such as the bilateral filter have been reformulated as priors [Tomasi and Manduchi, 1998, Elad, 2002]. Patched-based regularization techniques do not compare individual neighboring voxels, but patches of voxels in order to better distinguish intensity differences originating from noise or from actual structures [Wang and Qi, 2012]. Another technique is based on dictionary learning, where the images are sparsely represented from a linear combination of atoms. One expects that if the atoms are chosen (or learned) accordingly, this method captures local image features more effectively [Xu et al., 2012]. If prior images (normal dose) of the patient are available, such images can be used as regularization as well, for instance to reduce the patient dose of follow-up scans [Zhang et al., 2014a]. [Zhang et al., 2018]

4.3.5. Optimization algorithms

Having defined all parts of the objective function $\mathcal{C}(\boldsymbol{\mu}|\mathbf{Y})$ given by Eq. (4.32), the remaining task is to solve Eq. (4.31) by minimizing the objective function with respect to $\boldsymbol{\mu}$. To simplify notation $\mathcal{C}(\boldsymbol{\mu}) \equiv \mathcal{C}(\boldsymbol{\mu}|\mathbf{Y})$ is defined and thus, Eq. (4.31) can be written as

$$\hat{\boldsymbol{\mu}} = \arg\min_{\boldsymbol{\mu}} \mathcal{C}(\boldsymbol{\mu}). \tag{4.56}$$

The objective function is a scalar function according to $\mathcal{C} : \mathbb{R}^p \to \mathbb{R}$. In general, p is very large as it holds the number of unknowns (in CT p corresponds to the number of elements for the linear attenuation coefficient). A lot of research evolved around the field of numerical optimization, where one tries to optimize such objectives. These problems play a crucial role in various fields in science, engineering, economics and industry. [Nocedal and Wright, 2006]

In many cases, the objective function cannot be minimized analytically. Either the resulting problem would be too expensive to compute (as it would for instance require the inversion of the projection matrix given by Eq. (4.23)), or the objective function can simply not be inverted analytically. Thus, the objective function is minimized iteratively by defining an initial guess $\boldsymbol{\mu}^{(0)}$ and then computing a sequence of iterates $\{\boldsymbol{\mu}^{(n)}\}_{i=0}^{N}$ all of which have a lower function value than the value of the preceding iterate, denoted by $\mathcal{C}(x^{(n+1)}) < \mathcal{C}(x^{(n)})$, until the number of iterations N is exceeded. [Nocedal and Wright, 2006]

Line search

One strategy to compute this sequence is referred to as line search. Thereby, in each iteration, a search direction $\mathbf{p}$ is computed from the current iterate $\boldsymbol{\mu}^{(n)}$ along which a new iterate with a lower objective function is searched according to

$$\min_{\alpha>0} \mathcal{C}(\boldsymbol{\mu}^{(n)} + \alpha\mathbf{p}^{(n)}), \tag{4.57}$$

where α is referred to as the step length. Various different algorithms exist that define the search direction as well as the procedure to find a suitable step length. The methods differ in terms of robustness (they should perform well on a variety of problems), efficiency (they should only require a small amount of computational time or storage) and accuracy (they should identify a solution accurately). Thus, different methods are suited for different problems. In the following, most of the algorithms used within this thesis are outlined. [Nocedal and Wright, 2006]

Search direction

The fundamental tool for the subsequent algorithms is Taylor's theorem. In this case the objective $\mathcal{C}(\boldsymbol{\mu}^{(n)})$ at some point $\mathbf{x}^{(\mathbf{n})}$ for any search direction $\mathbf{p}^{(n)}$ and step length α can be approximated by

$$\mathcal{C}(\boldsymbol{\mu}^{(n)} + \mathbf{p}^{(n)}) \approx \mathcal{C}(\boldsymbol{\mu}^{(n)}) + \left(\mathbf{p}^{(n)}\right)^T \nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n)}} + \frac{1}{2}\left(\mathbf{p}^{(n)}\right)^T \nabla^2\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}\mathbf{p}^{(n)}. \tag{4.58}$$

Here, $\nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}$ denotes the gradient of $\mathcal{C}$ with respect to $\boldsymbol{\mu}$ evaluated at $\boldsymbol{\mu}^{(n)}$ and $\nabla^2\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}$ is the Hessian of $\mathcal{C}$ evaluated at the same point. [Nocedal and Wright, 2006]

Ignoring the second order term in Eq. (4.58), one can choose the search direction by minimizing the objective function value according to

$$\min_{\mathbf{p}^{(n)}} \left(\mathbf{p}^{(n)}\right)^T \nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}, \tag{4.59}$$

enforcing unit length according to $\|\mathbf{p}\| = 1$. By explicitly writing the scalar product as $a^T b = |a||b| \cos\theta$, realizing that this expression is minimal for $\cos\theta = -1$, the search direction is given by

$$\mathbf{p}_{\mathrm{SD}}^{(n)} = -\frac{\nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}}{\|\nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}\|}. \tag{4.60}$$

This is equivalent to the direction of the negative gradient normalized accordingly. By definition, this is the direction along which the objective decreases most rapidly. This search direction is therefore referred to as steepest descent direction. Most importantly, this choice of the search direction does not require any second order information and is thus computationally very efficient.

Another choice for the search direction is referred to as the Newton direction. For this approach, Eq. (4.58) is minimized directly by setting the gradient of the objective with respect to the search direction to zero resulting in

$$\mathbf{p}_{\mathrm{ND}}^{(n)} = -\left(\nabla^2\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}\right)^{-1} \nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}. \tag{4.61}$$

This method is very powerful, if the Hessian of the objective can be easily inverted and the objective is well approximated by a quadratic function. In addition, the search direction is already properly scaled, thus the computation of the step length might be omitted. [Nocedal and Wright, 2006]

The Hessian holds $p \times p$ values and consequently can be very large. This makes the explicit inversion of such a matrix computationally very expensive. Therefore, various methods exist that try to approximate the inverse Hessian. The corresponding search directions are referred to as Quasi-Newton directions. In essence, one tries to approximate the Hessian by differences of previously calculated gradients according to

$$\nabla^2\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}(\boldsymbol{\mu}^{(n+1)} - \boldsymbol{\mu}^{(n)}) \approx \nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n+1)}} - \nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}. \tag{4.62}$$

Consequently, one tries to find an approximate Hessian matrix denoted by $\mathbf{B}$ which fulfills

$$\mathbf{B}^{(n+1)}\mathbf{s}^{(n)} = \mathbf{t}^{(n)} \tag{4.63}$$

using $\mathbf{s}^{(n)} = \boldsymbol{\mu}^{(n+1)} - \boldsymbol{\mu}^{(n)}$ and $\mathbf{t}^{(n)} = \nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n+1)}} - \nabla\mathcal{C}|_{\boldsymbol{\mu}^{(n)}}$. Various different approximations of $\mathbf{B}^{(n)}$ exist, among others the Symmetric rank-one (SR1) formula and the Broyden–Fletcher–Goldfarb–Shanno (BFGS) formula. In practice, these formulas do not update $\mathbf{B}^{(n)}$ directly but its inverse denoted by $\mathbf{H}^{(n)} \equiv (\mathbf{B}^{(n)})^{-1}$. For the BFGS formula this update step can be written as

$$\mathbf{H}^{(n+1)} = \left(\mathbf{V}^{(n)}\right)^T \mathbf{H}^{(n)}\mathbf{V}^{(n)} + \rho^{(n)}\mathbf{s}^{(n)}\left(\mathbf{s}^{(n)}\right)^T \tag{4.64}$$

with $\mathbf{V}^{(n)} = \mathbf{I} - \rho^{(n)}\mathbf{s}^{(n)}(\mathbf{t}^{(n)})^T$ and $1/\rho^{(n)} = (\mathbf{t}^{(n)})^T\mathbf{s}^{(n)}$, where $\mathbf{I}$ denotes the identity matrix. As the above procedure would require to store the whole approximation of the Hessian matrix of the size $p \times p$, one may want to save only a few vectors of length m that represent this matrix explicitly. These methods are referred to as limited-memory methods. The following method is referred to as the limited-memory BFGS (L-BFGS) method [Nocedal, 1980, Liu and Nocedal, 1989] for which only m vector pairs $\{\mathbf{s}^{(i)}, \mathbf{t}^{(i)}\}$ for

$i = n-m, \ldots, n-1$ are stored. For completeness, this approximation is given by

$$\mathbf{H}^{(n)} = \left(\left(\mathbf{V}^{(n-1)}\right)^T \ldots \left(\mathbf{V}^{(n-m)}\right)^T\right) \mathbf{H}_0^{(n)} \left(\mathbf{V}^{(n-m)} \ldots \mathbf{V}^{(n-1)}\right) \tag{4.65}$$

$$+ \rho^{(n-m)} \left(\left(\mathbf{V}^{(n-1)}\right)^T \ldots \left(\mathbf{V}^{(n-m+1)}\right)^T\right) s^{(n-m)} \left(s^{(n-m)}\right)^T \left(\mathbf{V}^{(n-m+1)} \ldots \mathbf{V}^{(n-1)}\right) \tag{4.66}$$

$$+ \rho^{(n-m+1)} \left(\left(\mathbf{V}^{(n-1)}\right)^T \ldots \left(\mathbf{V}^{(n-m+2)}\right)^T\right) s^{(n-m+1)} \left(s^{(n-m+1)}\right)^T \left(\mathbf{V}^{(n-m+2)} \ldots \mathbf{V}^{(n-1)}\right) \tag{4.67}$$

$$+ \ldots \tag{4.68}$$

$$+ \rho^{(n-1)} s^{(n-1)} \left(s^{(n-1)}\right)^T \tag{4.69}$$

and for the initial approximation of the Hessian matrix, one may choose

$$\mathbf{H}_0^{(n)} = \frac{s_{k-1}^T y_{k-1}}{y_{k-1}^T y_{k-1}} \mathbf{I}, \tag{4.70}$$

which has proven to be effective in practice. Within this thesis, the memory was set to $m = 6$. Finally, the quasi-Newton search direction is given by

$$\mathbf{p}_{\mathrm{QN}}^{(n)} = -\mathbf{H}^{(n)} \nabla \mathcal{C}|_{\boldsymbol{\mu}^{(n)}}, \tag{4.71}$$

which can be computed efficiently and implicitly includes approximate second order information. [Nocedal and Wright, 2006]

Furthermore, there are other methods to define the search direction. Most prominent are conjugate gradient (CG) methods, which ensure that subsequent search directions are conjugate to each other. Although originally designed to solve linear systems, this method has been generalized to nonlinear problems. [Hager and Zhang, 2006]

Step length

In order to minimize the objective function, the search direction $\mathbf{p}$ is not sufficient. The step length α is used to decent on the objective function in search direction. The new iterate is then given by

$$\boldsymbol{\mu}^{(n+1)} = \boldsymbol{\mu}^{(n)} + \alpha^{(n)} \mathbf{p}^{(n)}. \tag{4.72}$$

While computing the optimal step length, a substantial reduction in the objective value should be achieved on the one hand (accuracy), but on the other hand, the method to obtain this step length should be fast (efficiency). Thus, usually a less accurate minimization of Eq. (4.57) is desirable, if it can be done computationally efficient. This is commonly performed in two steps. First, an interval containing desirable step lengths needs to be found. Then, by bisection or interpolation, a good step length is computed within this interval. In the following, two conditions are introduced, which should be fulfilled in order to accept the current step length. [Nocedal and Wright, 2006]

The first condition is the so-called Armijo condition, which requires that the chosen step length should give a sufficient decrease in the objective value according to

$$\mathcal{C}(\boldsymbol{\mu}^{(n)} + \alpha \mathbf{p}) \leq \mathcal{C}(\boldsymbol{\mu}^{(n)}) + c_1 \alpha \mathbf{p}^T \nabla \mathcal{C}|_{\boldsymbol{\mu}^{(n)}}, \tag{4.73}$$

where the scalar $c_1 \in (0, 1)$ has to be chosen manually. Within this thesis $c_1 = 10^{-4}$ is chosen if not stated otherwise. To ensure that the algorithm makes sufficient progress, the so-called curvature condition has to be satisfied given by

$$\mathbf{p}^T \nabla \mathcal{C}(\boldsymbol{\mu}^{(n)} + \alpha \mathbf{p}) \geq c_2 \mathbf{p}^T \nabla \mathcal{C}|_{\boldsymbol{\mu}^{(n)}}, \tag{4.74}$$

with another scalar $c_2 \in (c_1, 1)$ chosen manually. Here, this value is set according to $c_2 = 0.9$. These two conditions are referred to as the Wolfe conditions. A variation of this criterion is

$$|\mathbf{p}^T \nabla \mathcal{C}(\boldsymbol{\mu}^{(n)} + \alpha \mathbf{p})| \leq |c_2 \mathbf{p}^T \nabla \mathcal{C}|_{\boldsymbol{\mu}^{(n)}}|. \tag{4.75}$$

This equation and Eq. (4.73) form the *strong* Wolfe conditions. Various different algorithms have been developed to compute the step length efficiently and accurately to achieve the above conditions. Within this thesis, the algorithm of Moré and Thuente requiring the strong Wolfe conditions is used in most cases [Moré and Thuente, 1994]. [Nocedal and Wright, 2006]

Several algorithms have been developed that are tailored to the CT reconstruction using the mean model of Eq. (4.44) or Eq. (4.46). For the Poisson model given by Eq. (4.33), monotonic algorithms have been developed based on paraboloidal surrogates of the log-likelihood function, which can be optimized efficiently [Erdogan and Fessler, 1998]. Also accelerations using the ordered subsets principle have been introduced [Erdogan and Fessler, 1999a, Erdogan and Fessler, 1999b]. In addition, Nesterov-acceleration methods have been proposed [Kim et al., 2015]. In [Tilley et al., 2018a], a combination of these methods for the mean model similar to Eq. (4.44) and the Gaussian likelihood given by Eq. (4.43), has been proposed, which is discussed in more detail in Chapter 6.

Part II.

Methods and Results

5. Non-linear SIR framework for PBI and PB-CT

In the previous chapters, PBI and CT have been treated separately. In PB-CT, one acquires projection measurements like in PBI, but from various angles around the object. This enables three-dimensional visualization of thick and complex samples like in conventional CT, but additionally leveraging the phase-shifting properties of the object. Thus, PB-CT has become a valuable tool for examining biological samples, in particular due to its potential to increase contrast between materials with similar absorption properties, like soft-tissue. [Mayo et al., 2003, Bravin et al., 2013, Longo et al., 2016]

To obtain three-dimensional information about the measured sample, two steps are conventionally performed. First, the attenuation or phase properties are retrieved from each projection measurement separately as discussed in Chapter 3 for pure phase and monomorphous objects. Subsequently, this information is used to obtain the corresponding three-dimensional distributions of the recovered properties by means of tomographic reconstruction techniques discussed in Chapter 4. Several iterative approaches that integrate phase retrieval and tomographic reconstruction have been proposed and successfully applied for various PB-CT techniques. For instance, in [Ruhlandt et al., 2014], for larger propagation distances in the Fresnel regime, it has been shown that the phase information can be retrieved using iterative algorithms that cycle between the object and detector plane applying several constraints, similar to phase-retrieval approaches in the Fraunhofer regime [Paganin, 2006, p.289ff]. Other iterative approaches rely on the linearization of the image formation process, as discussed in Subsection 2.3.5. Thereby, multiple images are acquired at each view at different distances. Using this technique, it has been shown that one can accurately recover the phase information from multi-material and heterogeneous objects [Langer et al., 2012, Kostenko et al., 2013, Langer et al., 2014]. Various reasons exist why these approaches are limited to synchrotron radiation facilities. Among others, multi-distance methods seem to be more sensitive to noise [Arhatari et al., 2010]. In addition, the spatial coherence of laboratory sources is too low to access wider ranges of the CTF [Wilkins et al., 2014]. Moreover, due to the cone-beam geometry, alignment is more difficult in a laboratory setup [Weber et al., 2017].

In the following, SIR approaches are investigated for PBI and PB-CT, which use measurements acquired only at a single distance making these techniques suitable for laboratory environments. A significant amount of research evolves around SIR techniques for conventional attenuation-based CT to improve image quality or to reduce acquisition time and dose. As outlined in Section 4.3, these approaches model the attenuation in intensity, using for instance the mean model given by Eq. (4.44) or Eq. (4.46), with an appropriate noise model and regularization. Compared to the conventional analytical reconstruction approaches, SIR approaches have already demonstrated significant advantages for many applications [Thibault et al., 2007, Noël et al., 2013]. In this chapter, SIR concepts that have proven successful in conventional attenuation-based CT are transferred to PB-CT. In particular, the statistical description of the measurements inherent to SIR approaches is a promising concept for the ongoing transition of PB-CT to laboratory sources, which naturally suffers

from low flux levels.

This chapter is structured as follows. First, physical mean models for PBI and PB-CT are proposed using the results derived in Chapter 2. Based on these physical mean models, SIR reconstruction approaches are derived building upon Section 4.3, which include a description of the likelihood of the measurements as well as prior knowledge in the form of regularization techniques. In addition, concepts introduced in Chapter 3 are utilized and their relation to the previously discussed phase-retrieval and reconstruction algorithms is outlined. The methods are validated using experimental datasets acquired at laboratory sources and synchrotron facilities.

Parts of the results discussed in this chapter have been published in [Hehn et al., 2018a, Hehn et al., 2018b].

5.1. Image formation

First, the image formation in PBI suitable for a laboratory environment needs to be described to obtain a physical mean model according to Subsection 4.3.3, which connects the measured intensity at the detector plane to the different components of the refractive index describing the object. The theoretical background has been outlined in Chapter 2. The notation is adapted from Chapter 3, in particular using Figure 3.1. Parts of this section have been published in [Hehn et al., 2018b].

The X-ray wavefield generated by the source is described as a monochromatic forward propagating paraxial envelope $\tilde{\psi}(\mathbf{r}_\perp, z)$ defined by Eq. (2.6). To shorten notation, the subscript ω to specify a certain spectral component is again omitted. The wavefield is described in terms of its intensity and phase according to Eq. (2.16) at different stages along the optical axis.

For simplicity, it is assumed that the incoming wavefield has constant illumination given by I_0. The sample located at $z < 0$ is defined by its attenuating and phase-shifting properties. The attenuating properties are described by the three-dimensional distribution of the linear attenuation coefficient $\mu(\mathbf{r})$. The phase-shifting properties are given by the corresponding distribution of the refractive index decrement $\delta(\mathbf{r})$. Under the projection approximation discussed in Section 2.2, the intensity of the wavefield behind the object $\tilde{\psi}(\mathbf{r}_\perp, 0)$ is given by the Lambert-Beer law according to Eq. (2.17) written as

$$I(\mathbf{r}_\perp, 0) = I_0 \exp\left[-\int \mu(\mathbf{r}_\perp, z)\mathrm{d}z\right] = I_0 e^{-\ell(\mathbf{r}_\perp, 0)}, \tag{5.1}$$

where the line-integrals $\ell(\mathbf{r}_\perp, 0)$ are defined as the projection of the linear attenuation coefficient of the object along the X-ray paths. The corresponding phase shift is given according to Eq. (2.18) by

$$\phi(\mathbf{r}_\perp, 0) = -k \int \delta(\mathbf{r}_\perp, 0)\mathrm{d}z \tag{5.2}$$

and is proportional to the projection of the refractive index decrement of the object along the X-ray path[1]. In the context of PBI, the validity of the projection approximation has been discussed in [Morgan et al., 2010], demonstrating that this approximation is accurate for describing the corresponding interference effects.

[1] Technically, the phase shift can also be thought of as line-integrals. However, within this work, line-integrals refer to the projection of the linear attenuation coefficient as defined by Eq. (4.45) and phase shift refers to the projection of the refractive index decrement.

For laboratory environments, the limit of small propagation distances derived in Section 2.3 holds, which relates the intensity at the detector plane to the intensity and phase behind the sample according to Eq. (2.43), which can be written as

$$I(\mathbf{r}_\perp, z) = I(\mathbf{r}_\perp, 0) - \frac{z}{k}\nabla_\perp \cdot [I(\mathbf{r}_\perp, 0)\nabla_\perp \phi(\mathbf{r}_\perp, 0)] \tag{5.3}$$

$$= I(\mathbf{r}_\perp, 0) - \frac{z}{k}\left\{[\nabla_\perp I(\mathbf{r}_\perp, 0)] \cdot \nabla_\perp \phi(\mathbf{r}_\perp, 0) + I(\mathbf{r}_\perp, 0)\nabla_\perp^2 \phi(\mathbf{r}_\perp, 0)\right\}, \tag{5.4}$$

where the terms are expanded explicitly. It is further assumed that at a particular position in the wavefield, the product of the intensity gradient and the phase gradient is comparably small, written as

$$[\nabla_\perp I(\mathbf{r}_\perp, 0)] \cdot \nabla_\perp \phi(\mathbf{r}_\perp, 0) \ll I(\mathbf{r}_\perp, 0)\nabla_\perp^2 \phi(\mathbf{r}_\perp, 0). \tag{5.5}$$

This approximation can also be found in [Paganin, 2006, p.281]. Consequently, the intensity at the detector plane is given by

$$\boxed{I(\mathbf{r}_\perp, z) \approx I(\mathbf{r}_\perp, 0)\left(1 - \frac{z}{k}\nabla_\perp^2 \phi(\mathbf{r}_\perp, 0)\right),} \tag{5.6}$$

where the attenuation is modeled by the Lambert-Beer law and the interference effects are proportional to the Laplacian of the phase shift. This allows to interpret the intensity at the detector plane as a product of the influence of the attenuation and diffraction. In [Bremmer, 1952], thorough derivations on expressing the wave function in the detector plane given the wave function in the object plane have been discussed, among others the expansion in powers of the Laplacian operator.

5.2. Radiography

Having defined a physical model describing the image formation in PBI, several MAP approaches are formulated to recover information about the object. Parts of this section have been published in [Hehn et al., 2018b].

5.2.1. Derivation of the physical mean model and likelihood

Following Section 4.2.1, the intensity at the detector plane given by Eq. (5.6) is recorded by a detector with finite pixelsize. Thus, the measured intensity is sampled by default. To reduce notation overhead, the measured intensity is denoted by a (one-dimensional) vector $\mathbf{y}$. The intensity behind the object is expressed in terms of the line-integrals by a vector $\boldsymbol{\ell}$ according to Eq. (5.1). The discrete version of the phase shift is denoted by $\boldsymbol{\phi}$. Moreover, the Laplacian operator can be modeled as a Toeplitz matrix $\mathbf{L}$. Finally, the physical mean model connecting the intensity at the detector plane to the line-integrals and the phase shift is given in matrix notation by

$$\bar{\mathbf{y}}(\boldsymbol{\ell}, \boldsymbol{\phi}) = I_0 \mathcal{D}\left[e^{-\boldsymbol{\ell}}\right]\left(\mathbf{1} - \frac{z}{k}\mathbf{L}\boldsymbol{\phi}\right), \tag{5.7}$$

where $\mathcal{D}$ denotes again the diagonal matrix with the vector in brackets on its diagonal. Several discretizations of the Laplacian operator are possible. In practice, the five-point stencil finite-difference method is used, which only includes the four nearest neighbors. This provides good performance and reasonable accuracy. Another possibility would be to express

this model in terms of the intensity behind the object instead of the line-integrals. However, the expression in terms of the line-integrals will become more intuitive when extending this model to tomography.

Next, the negative log-likelihood term is derived. As the propagation of the X-rays does not correlate noise, one can utilize the same formulation as for attenuation-based imaging. Thus, the negative log-likelihood of having obtained the actual measurements $\mathbf{y}$ given the line-integrals $\boldsymbol{\ell}$ and the phase shift $\boldsymbol{\phi}$ can be modeled by a Gaussian noise model according to Eq. (4.43) given by

$$\mathcal{L}(\mathbf{y}|\boldsymbol{\ell},\boldsymbol{\phi}) = \frac{1}{2}\left(\mathbf{y}-\bar{\mathbf{y}}(\boldsymbol{\ell},\boldsymbol{\phi})\right)^T \mathbf{W}\left(\mathbf{y}-\bar{\mathbf{y}}(\boldsymbol{\ell},\boldsymbol{\phi})\right). \tag{5.8}$$

In principle, one could also choose the Poisson noise model given by Eq. (4.33), but as discussed in detail in Chapter 6, the Gaussian noise model is more suited for describing detector systems used in PBI.

Ideally, one would estimate both the line-integrals as well as the phase shift by optimizing the negative log-likelihood according to

$$\left\{\hat{\boldsymbol{\ell}},\hat{\boldsymbol{\phi}}\right\} = \arg\min_{\boldsymbol{\ell},\boldsymbol{\phi}} \mathcal{L}(\mathbf{y}|\boldsymbol{\ell},\boldsymbol{\phi}). \tag{5.9}$$

For every measurement, this leaves twice as many unknowns $(\boldsymbol{\ell},\boldsymbol{\phi})$ as knowns $(\mathbf{y})$. In this case, there are many solutions to solve Eq. (5.9) including

$$\boldsymbol{\ell} = -\log\left(\frac{\mathbf{y}}{I_0}\right) \qquad \text{and} \qquad \boldsymbol{\phi} = \mathbf{0}, \tag{5.10}$$

which can be regarded as the attenuation-based solution as it ignores phase effects. However, the above solution implies negative line-integrals where the measured intensity is higher than I_0 due to interference effects. Although not the focus of this thesis, one may impose prior knowledge in form of (physical) constraints, such as requiring the line-integrals to be positive, to penalize non-physical solutions. This can for instance be achieved by regularization techniques.

Although the above formulations are not used to estimate the line-integrals and phase shift independently, they build the foundation of the subsequent algorithms. To minimize the likelihood given by Eq. (5.8) according to Eq (5.9), gradient-based optimization algorithms are used. These algorithms require the gradients[2] of the negative log-likelihood with respect to its arguments, which are derived in the following. First of all, deriving the negative log-likelihood of Eq. (5.8) with respect to the physical mean model of Eq. (5.7) results in

$$\frac{\partial\mathcal{L}(\mathbf{y}|\boldsymbol{\ell},\boldsymbol{\phi})}{\partial\bar{\mathbf{y}}(\boldsymbol{\ell},\boldsymbol{\phi})} = -\frac{1}{2}\left(\mathbf{W}+\mathbf{W}^T\right)\left(\mathbf{y}-\bar{\mathbf{y}}(\boldsymbol{\ell},\boldsymbol{\phi})\right) \tag{5.11}$$

$$= -\mathbf{W}\left(\mathbf{y}-\bar{\mathbf{y}}(\boldsymbol{\ell},\boldsymbol{\phi})\right), \tag{5.12}$$

assuming that $\mathbf{W}$ is symmetric, which holds for diagonal matrices or covariance matrices by definition. Next, the derivatives of the mean model with respect to the line-integrals as well as with respect to the phase shift are calculated according to

$$\frac{\partial\bar{\mathbf{y}}(\boldsymbol{\ell},\boldsymbol{\phi})}{\partial\boldsymbol{\ell}} = -\mathcal{D}\left[\mathbf{1}-\frac{z}{k}\mathbf{L}\boldsymbol{\phi}\right]\mathcal{D}\left[e^{-\boldsymbol{\ell}}\right]I_0 \tag{5.13}$$

[2] For computation of the gradients, the denominator-layout notation is used, which defines the gradient of the (scalar) objective function with respect to a column vector to be a column vector as well. This notation seems most intuitive for gradient-based approaches.

and

$$\frac{\partial \bar{\mathbf{y}}(\boldsymbol{\ell}, \boldsymbol{\phi})}{\partial \boldsymbol{\phi}} = -\frac{z}{k}\mathbf{L}^T \mathcal{D}\left[e^{-\boldsymbol{\ell}}\right] I_0. \tag{5.14}$$

Thus, using Eq. (5.12), Eq. (5.13) and Eq. (5.14), and applying the chain rule, the derivatives of the negative log-likelihood with respect to its arguments are given by

$$\frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}, \boldsymbol{\phi})}{\partial \boldsymbol{\ell}} = \frac{\partial \bar{\mathbf{y}}(\boldsymbol{\ell}, \boldsymbol{\phi})}{\partial \boldsymbol{\ell}} \frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}, \boldsymbol{\phi})}{\partial \bar{\mathbf{y}}(\boldsymbol{\ell}, \boldsymbol{\phi})} = \mathcal{D}\left[\mathbf{1} - \frac{z}{k}\mathbf{L}\boldsymbol{\phi}\right] \mathcal{D}\left[e^{-\boldsymbol{\ell}}\right] I_0 \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\ell}, \boldsymbol{\phi})\right) \tag{5.15}$$

and

$$\frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}, \boldsymbol{\phi})}{\partial \boldsymbol{\phi}} = \frac{\partial \bar{\mathbf{y}}(\boldsymbol{\ell}, \boldsymbol{\phi})}{\partial \boldsymbol{\phi}} \frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}, \boldsymbol{\phi})}{\partial \bar{\mathbf{y}}(\boldsymbol{\ell}, \boldsymbol{\phi})} = \frac{z}{k}\mathbf{L}^T \mathcal{D}\left[e^{-\boldsymbol{\ell}}\right] I_0 \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\ell}, \boldsymbol{\phi})\right). \tag{5.16}$$

5.2.2. Homogeneity assumption

In Section 3.2, the homogeneity assumption was introduced, which couples the phase and amplitude of the wavefield behind the sample. For monomorphous objects, this enables the recovery of the phase shift from a single measurement, resulting in the single-material phase-retrieval algorithm given by Eq. (3.18). Using Eq. (3.10) and Eq. (3.11) according to

$$\boldsymbol{\ell}(\mathbf{t}) = \mu\mathbf{t} \qquad \text{and} \qquad \boldsymbol{\phi}(\mathbf{t}) = -k\delta\mathbf{t} \tag{5.17}$$

with the material specific constants μ and δ and the wave number k, the line-integrals and the phase shift are expressed in terms of the discrete representation of the trace $\mathbf{t}$. Inserting this into the physical mean model given by Eq. (5.7) results in

$$\boxed{\bar{\mathbf{y}}(\mathbf{t}) = I_0 \mathcal{D}\left[e^{-\mu\mathbf{t}}\right](\mathbf{1} + z\delta\mathbf{L}\mathbf{t}),} \tag{5.18}$$

where the mean intensity depends now solely on the trace of the object. Thus, there are as many unknowns ($\mathbf{t}$) as knowns ($\mathbf{y}$). The likelihood term of Eq. (5.8) remains unchanged and is minimized according to

$$\hat{\mathbf{t}} = \arg\min_{\mathbf{t}} \mathcal{L}\left(\mathbf{y}|\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t})\right). \tag{5.19}$$

The gradients relating the line-integrals and the phase shift to the trace are constant and given by

$$\frac{\partial \boldsymbol{\ell}(\mathbf{t})}{\partial \mathbf{t}} = \mu\mathbf{I} \qquad \text{and} \qquad \frac{\partial \boldsymbol{\phi}(\mathbf{t})}{\partial \mathbf{t}} = -k\delta\mathbf{I}, \tag{5.20}$$

where $\mathbf{I}$ denotes the identity matrix. Accounting for the fact that the line-integrals and the phase shift are now functions of the trace denoted by $\boldsymbol{\ell} = \boldsymbol{\ell}(\mathbf{t})$ and $\boldsymbol{\phi} = \boldsymbol{\phi}(\mathbf{t})$, the gradient of the likelihood with respect to the trace is given by the chain rule using Eq. (5.20), Eq. (5.15), Eq. (5.16) and Eq. (5.12) according to

$$\frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))}{\partial \boldsymbol{t}} = \frac{\partial \boldsymbol{\ell}(\mathbf{t})}{\partial \mathbf{t}} \frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))}{\partial \boldsymbol{\ell}(\boldsymbol{t})} + \frac{\partial \boldsymbol{\phi}(\mathbf{t})}{\partial \mathbf{t}} \frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))}{\partial \boldsymbol{\phi}(\boldsymbol{t})} \tag{5.21}$$

$$= \left(\frac{\partial \boldsymbol{\ell}(\mathbf{t})}{\partial \mathbf{t}} \frac{\partial \bar{\mathbf{y}}(\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))}{\partial \boldsymbol{\ell}(\mathbf{t})} + \frac{\partial \boldsymbol{\phi}(\mathbf{t})}{\partial \mathbf{t}} \frac{\partial \bar{\mathbf{y}}(\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))}{\partial \boldsymbol{\phi}(\mathbf{t})}\right) \frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))}{\partial \bar{\mathbf{y}}(\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))} \tag{5.22}$$

$$= \left(\mu\mathcal{D}\left[\mathbf{1} - \frac{z}{k}\mathbf{L}\boldsymbol{\phi}(\mathbf{t})\right] - z\delta\mathbf{L}^T\right) \mathcal{D}\left[e^{-\boldsymbol{\ell}(\mathbf{t})}\right] I_0 \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))\right) \tag{5.23}$$

$$= \left(\mu\mathcal{D}\left[\mathbf{1} + z\delta\mathbf{L}\mathbf{t}\right] - z\delta\mathbf{L}^T\right) \mathcal{D}\left[e^{-\mu\mathbf{t}}\right] I_0 \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{t})\right), \tag{5.24}$$

where the dependency of the line-integrals and the phase shift on the trace given by Eq. (5.17) and $\bar{\mathbf{y}}(\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t})) = \bar{\mathbf{y}}(\mathbf{t})$ are used.

5.2.3. Limit of pure phase objects

For the limiting case of a pure phase object (no absorption), $\boldsymbol{\ell} = \mathbf{0}$ holds. Thus, the mean model of Eq. (5.7) simplifies to

$$\bar{\mathbf{y}}(\boldsymbol{\phi}) = I_0 \left(\mathbf{1} - \frac{z}{k} \mathbf{L} \boldsymbol{\phi} \right). \tag{5.25}$$

In this case, the log-likelihood can be solved directly by solving $\mathbf{y} = \bar{\mathbf{y}}(\boldsymbol{\phi})$, resulting in

$$\boldsymbol{\phi} = \boldsymbol{\mathcal{F}}^{-1} \left[\frac{\boldsymbol{\mathcal{F}} \left[\mathbf{y}/I_0 - 1 \right]}{\frac{z}{k} \mathbf{k}_\perp^T \mathbf{k}_\perp} \right], \tag{5.26}$$

where $\boldsymbol{\mathcal{F}}_\perp$ denotes the two-dimensional discrete Fourier transform (DFT) with the spatial frequencies $\mathbf{k}_\perp$, given that $\mathbf{L}$ approximates the continuous Laplacian such that the Fourier derivative theorem holds. In practice, the FFT algorithm is used to evaluate the DFT numerically. Here, the convention of the continuous FT given by Eq. (2.26) and Eq. (2.27) is adapted for the DFT to be consistent with the analytical formulas derived in Chapter 3. The conventional definition of the DFT can for instance be found in [Oliphant, 2006]. The above equation coincides with the phase retrieval algorithm derived analytically in Eq. (3.5). It also coincides with [Burvall et al., 2011] using $k = 2\pi\lambda^{-1}$ and $\mathbf{k}_\perp \to 2\pi\mathbf{k}_\perp$, due to different definitions of the DFT. [Gureyev and Nugent, 1996, Bronnikov, 1999]

5.3. Tomography

As discussed above, in general one cannot solve for the line-integrals and the phase shift independently as this problem is ill-posed. In the following, the physical mean model and the likelihood are extended to tomography. Ideally, due to the additional constraints imposed by tomography, a more independent reconstruction of the complex refractive index can be achieved. Parts of this section have been published in [Hehn et al., 2018b].

5.3.1. Extension of the physical mean model and likelihood

In tomography, the quantities of interest are not the line-integrals and the phase shift, but the three-dimensional distribution of the linear attenuation coefficient $\boldsymbol{\mu}$, which describes the attenuating properties of the object, and the three-dimensional distribution of the refractive index decrement $\boldsymbol{\delta}$ describing the phase-shifting properties of the sample. These quantities relate to the line-integrals and the phase shift according to

$$\boldsymbol{\ell}(\boldsymbol{\mu}) = \mathbf{A}\boldsymbol{\mu} \qquad \text{and} \qquad \boldsymbol{\phi}(\boldsymbol{\delta}) = -k\mathbf{A}\boldsymbol{\delta}, \tag{5.27}$$

as discussed in Section 2.2. The matrix $\mathbf{A}$ denotes the operator describing the forwardprojection. The entries of this matrix in the voxel basis are given by Eq. (4.22). Now, $\boldsymbol{\ell}$ and $\boldsymbol{\phi}$ hold the line-integrals and phase shift under every view. The physical mean model for tomography is then given by inserting Eq. (5.27) into Eq. (5.7) resulting in

$$\bar{\mathbf{y}}(\boldsymbol{\mu}, \boldsymbol{\delta}) = I_0 \mathcal{D} \left[e^{-\mathbf{A}\boldsymbol{\mu}} \right] (\mathbf{1} + z\mathbf{L}\mathbf{A}\boldsymbol{\delta}). \tag{5.28}$$

Again, the measurements $\mathbf{y}$, which now hold the measured intensities under every view and pixel, can again be modeled by a normal distribution according to Eq. (4.43) resulting in

$$\mathcal{L}(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\delta}) = \frac{1}{2} \left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu}, \boldsymbol{\delta}) \right)^T \mathbf{W} \left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu}, \boldsymbol{\delta}) \right). \tag{5.29}$$

Here, the goal is to estimate the three-dimensional distribution of the linear attenuation coefficient as well as the refractive index decrement according to

$$\left\{\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\delta}}\right\} = \arg\min_{\boldsymbol{\mu},\boldsymbol{\delta}} \mathcal{L}(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\delta}). \tag{5.30}$$

In this case, the number of unknowns $(\boldsymbol{\mu}, \boldsymbol{\delta})$ and knowns $(\mathbf{y})$ depends on the number of acquired views. However, in general an independent reconstruction of these two quantities is not feasible.

The gradients of the line-integrals and the phase shift with respect to the linear attenuation coefficient and the refractive index decrement are given by

$$\frac{\partial \boldsymbol{\ell}(\boldsymbol{\mu})}{\partial \boldsymbol{\mu}} = \mathbf{A}^T \qquad \text{and} \qquad \frac{\partial \boldsymbol{\phi}(\boldsymbol{\delta})}{\partial \boldsymbol{\delta}} = -k\mathbf{A}^T, \tag{5.31}$$

where $\mathbf{A}^T$ denotes the operator describing the backprojection. The corresponding gradients of the likelihood using $\mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\delta}) = \mathcal{L}(\boldsymbol{\ell}(\boldsymbol{\mu}), \boldsymbol{\phi}(\boldsymbol{\delta}))$ are then given by

$$\begin{aligned}
\frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\boldsymbol{\mu}), \boldsymbol{\phi}(\boldsymbol{\delta}))}{\partial \boldsymbol{\mu}} &= \frac{\partial \boldsymbol{\ell}(\boldsymbol{\mu})}{\partial \boldsymbol{\mu}} \frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\boldsymbol{\mu}), \boldsymbol{\phi}(\boldsymbol{\delta}))}{\partial \boldsymbol{\ell}(\boldsymbol{\mu})} && (5.32)\\
&= \mathbf{A}^T \mathcal{D}\left[\mathbf{1} - \frac{z}{k}\mathbf{L}\boldsymbol{\phi}(\boldsymbol{\delta})\right] \mathcal{D}\left[e^{-\boldsymbol{\ell}(\boldsymbol{\mu})}\right] I_0 \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\ell}(\boldsymbol{\mu}), \boldsymbol{\phi}(\boldsymbol{\delta}))\right) && (5.33)\\
&= \mathbf{A}^T \mathcal{D}\left[\mathbf{1} + z\mathbf{L}\mathbf{A}\boldsymbol{\delta}\right] \mathcal{D}\left[e^{-\mathbf{A}\boldsymbol{\mu}}\right] I_0 \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu}, \boldsymbol{\delta})\right) && (5.34)
\end{aligned}$$

and

$$\begin{aligned}
\frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\boldsymbol{\mu}), \boldsymbol{\phi}(\boldsymbol{\delta}))}{\partial \boldsymbol{\delta}} &= \frac{\partial \boldsymbol{\phi}(\boldsymbol{\delta})}{\partial \boldsymbol{\delta}} \frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\boldsymbol{\mu}), \boldsymbol{\phi}(\boldsymbol{\delta}))}{\partial \boldsymbol{\phi}(\boldsymbol{\delta})} && (5.35)\\
&= -\mathbf{A}^T z \mathbf{L}^T \mathcal{D}\left[e^{-\boldsymbol{\ell}(\boldsymbol{\mu})}\right] I_0 \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\ell}(\boldsymbol{\mu}), \boldsymbol{\phi}(\boldsymbol{\delta}))\right) && (5.36)\\
&= -\mathbf{A}^T z \mathbf{L}^T \mathcal{D}\left[e^{-\mathbf{A}\boldsymbol{\mu}}\right] I_0 \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\boldsymbol{\mu}, \boldsymbol{\delta})\right). && (5.37)
\end{aligned}$$

5.3.2. Homogeneity assumption

As for radiography, one can impose the homogeneity assumption according to

$$\mathbf{t}(\mathbf{x}) = \mathbf{A}\mathbf{x}, \tag{5.38}$$

where $\mathbf{x}$ denotes the three-dimensional distribution of the (line) density, referred to, for simplicity, as the three-dimensional distribution of the sample. Inserting Eq. (5.38) into Eq. (5.18) results in the following physical mean model

$$\boxed{\bar{\mathbf{y}}(\mathbf{x}) = I_0 \mathcal{D}\left[e^{-\mu\mathbf{A}\mathbf{x}}\right](\mathbf{1} + z\delta\mathbf{L}\mathbf{A}\mathbf{x}),} \tag{5.39}$$

which models the measured intensity depending only on the three-dimensional distribution of the sample. This model is valid if the sample under consideration obeys the homogeneity assumption. The number of unknowns $(\mathbf{x})$ and the number of knowns $(\mathbf{y})$ coincide with conventional tomographic reconstruction.

The gradient of the trace with respect to the density gives

$$\frac{\partial \mathbf{t}(\mathbf{x})}{\partial \mathbf{x}} = \mathbf{A}^T \tag{5.40}$$

and consequently, the respective gradient of the likelihood term can be computed using Eq. (5.40) and Eq. (5.24) and applying the chain rule resulting in

$$\frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\mathbf{t}(\mathbf{x})), \boldsymbol{\phi}(\mathbf{t}(\mathbf{x})))}{\partial \boldsymbol{t}(\mathbf{x})} \tag{5.41}$$

$$= \frac{\partial \mathbf{t}(\mathbf{x})}{\partial \mathbf{x}} \frac{\partial \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\mathbf{t}(\mathbf{x})), \boldsymbol{\phi}(\mathbf{t}(\mathbf{x})))}{\partial \boldsymbol{t}(\mathbf{x})} \tag{5.42}$$

$$= \mathbf{A}^T \left(\mu \mathcal{D} \left[\mathbf{1} + z\delta \mathbf{L}\mathbf{t}(\mathbf{x}) \right] - z\delta \mathbf{L}^T \right) \mathcal{D} \left[e^{-\mu \mathbf{t}(\mathbf{x})} \right] I_0 \mathbf{W} \left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{t}(\mathbf{x})) \right) \tag{5.43}$$

$$= \mathbf{A}^T \left(\mu \mathcal{D} \left[\mathbf{1} + z\delta \mathbf{L}\mathbf{A}\mathbf{x} \right] - z\delta \mathbf{L}^T \right) \mathcal{D} \left[e^{-\mu \mathbf{A}\mathbf{x}} \right] I_0 \mathbf{W} \left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x}) \right) \tag{5.44}$$

and making use of Eq. (5.38) and $\bar{\mathbf{y}}(\mathbf{t}(\mathbf{x})) = \bar{\mathbf{y}}(\mathbf{x})$.

5.3.3. Limit of conventional tomography

In the limit of no phase effects such that $z = 0$, the physical mean model of Eq. (5.28) reduces to

$$\bar{\mathbf{y}}(\boldsymbol{\mu}) = I_0 e^{-\mathbf{A}\boldsymbol{\mu}}, \tag{5.45}$$

which coincides with the physical mean model of conventional attenuation-based CT given by Eq. (4.44), modeling solely the attenuation of intensity.

In summary, it was shown how the framework introduced above can be used for attenuation-based CT as well as PB-CT employing the homogeneity assumption. In addition, this framework has the potential to investigate means to more independently reconstruct the attenuating and phase-shifting properties of the sample from a single distance in PB-CT using additional constraints in the form of regularization techniques.

5.4. Validation of the homogeneous phase-retrieval algorithm at a laboratory source

First of all, the physical mean model of Eq. (5.18) describing the interference effects under the homogeneity assumption is validated using an X-ray micro-tomography system. In addition, a comparison to the single-material phase-retrieval given in its analytical form by Eq. (3.18) is performed. This comparison is crucial as the underlying physical mean model of Eq. (5.7) describing the interference effects in terms of the line-integrals as well as the phase shift does not coincide with the model used in the single-material phase-retrieval algorithm, due to the approximation given by Eq. (5.5). Thereby, it was assumed that the product of the intensity gradient and the phase gradient can be neglected. Parts of this validation have been published in [Hehn et al., 2018b].

5.4.1. Phase retrieval algorithms

Before stating the optimization-based approach for phase retrieval building upon the results of Section 5.2, the discrete version of the single-material phase-retrieval algorithm is outlined.

Single-material phase-retrieval algorithm

The discretization of the single-material phase-retrieval algorithm follows analogous to the Bronnikov algorithm as stated in Eq. (5.26). Thus, from the analytical form of Eq. (3.18), the FT is replaced by the DFT and the continuous measurements $I(\mathbf{r}_\perp, z)$ are represented by the vector $\mathbf{y}$, resulting in

$$\boxed{\mathbf{t} = -\frac{1}{\mu}\log\left(\boldsymbol{\mathcal{F}}_\perp^{-1}\left[\frac{\boldsymbol{\mathcal{F}}_\perp\left[\mathbf{y}\right]/I_0}{z\frac{\delta}{\mu}\mathbf{k}_\perp^T\mathbf{k}_\perp + 1}\right]\right),} \tag{5.46}$$

which again coincides with the notation found in [Burvall et al., 2011] substituting $\mathbf{k}_\perp \to 2\pi\mathbf{k}_\perp$ accordingly.

Optimization-based phase retrieval approach

The proposed phase retrieval method minimizes the negative log-likelihood of the measurements with respect to the trace of the sample using the physical mean model describing the image formation under the homogeneity assumption, which is given by Eq. (5.18). Using $\mathcal{L}(\mathbf{y}|\mathbf{t}) = \mathcal{L}(\mathbf{y}|\boldsymbol{\ell}(\mathbf{t}), \boldsymbol{\phi}(\mathbf{t}))$, which can be derived from Eq. (5.8), the corresponding negative log-likelihood is given by

$$\mathcal{L}(\mathbf{y}|\mathbf{t}) = \frac{1}{2}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{t})\right)^T\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{t})\right), \tag{5.47}$$

assuming $\mathbf{W} = \mathbf{I}$ to be the identity. No statistical weighting and no additional prior knowledge in the form of regularization terms is included into the phase-retrieval algorithm in order not to bias the comparison to the single-material phase-retrieval algorithm. Formally, the optimal trace is then given by

$$\hat{\mathbf{t}} = \arg\min_{\mathbf{t}} \mathcal{L}(\mathbf{t}|\mathbf{y}). \tag{5.48}$$

5.4.2. Xradia 500 Versa

The X-ray microtomography system (Xradia 500 Versa) is a commercial device sold by Carl Zeiss AG. The X-ray source is a transmission tube with a 5 µm thick tungsten target. The tube current ranges between $30 - 160$ keV with a power range of $2 - 10$ W. The focal spot size depends on the tube current but is overall stated to be less than 5 µm. The focal spot size is connected to the spatial coherence of the X-rays and thus one of the most important characteristics for PBI. The detector system employs indirect detection of the X-rays using a scintillator which converts the X-rays to visible light. This light travels through lenses for magnification or demagnification. Afterwards, the light is converted to a digital signal inside a charged-coupled device (CCD) camera (Andor iKon-L) with 2048×2048 pixels. The pixelsize is 13.5 µm. [Bidola, 2017]

5.4.3. Sample and experimental parameters

The data used for validation has been measured as part of a study to optimize PBI at laboratory setups, published in [Bidola et al., 2015], where additional details can be found.

The source-to-axis distance (SAD) was set to $d_{\mathrm{SAD}} = 10$ mm and the source-to-detector distance (SDD) was set to $d_{\mathrm{SDD}} = 70$ mm, resulting in a propagation distance of $z = d_{\mathrm{SDD}} -$

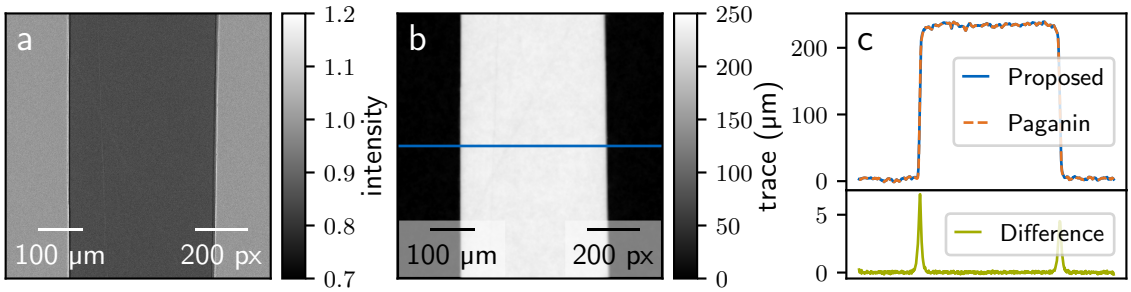

Figure 5.1.: Validation of the proposed phase-retrieval algorithm using a Teflon plate. In (a) the measured intensity behind the Teflon plate relative to the reference intensity is depicted. In (b) the trace recovered by minimizing the likelihood according to Eq. (5.48) is shown. A comparison to the conventional single-material phase-retrieval algorithm is shown in (c). In the upper part, the line profile of the middle row indicated by a blue line in (b) is shown in blue. In orange, the same row obtained with single-material phase-retrieval algorithm is shown, referred to as 'Paganin'. Finally, in the lower part, the difference to the single-material phase-retrieval algorithm is shown. This figure was adapted from [Hehn et al., 2018b].

$d_{\mathrm{SAD}} = 60\,\mathrm{mm}$. The X-ray source was operated at 40 keV and a power of 3 W, resulting in a spot size of around 1.8 µm. The average energy of the X-ray beam is 13 keV, as calculated in [Bidola et al., 2015]. Lenses with four times magnification were used in the detection system. To achieve lower exposure times, the detector pixels were binned by a factor of two. This results in an effective pixelsize of $p = 1/4 \cdot 2 \cdot 13.5\,\mu\mathrm{m} = 6.75\,\mu\mathrm{m}$. [Bidola et al., 2015]

A head-on projection of a 250 µm thick Teflon plate was acquired, which can be seen in Figure 5.1 (a). The attenuating properties of the sample as well as edge-enhancement effects at the interfaces from Teflon to air, where information about the phase-shifting properties of the sample is encoded, can be seen.

5.4.4. Results and discussion

The phase-retrieval algorithms require several parameters. First of all, the material specific scalar constants μ and δ are extracted from the xraylib library [Schoonjans et al., 2011] according to the material, density and energy. As the physical mean model is derived for a parallel-beam geometry, the Fresnel scaling theorem given by Eq. (2.55) is employed to translate the parameters of the cone-beam geometry to the corresponding parallel-beam geometry using the magnification given by $M = d_{\mathrm{SDD}}/d_{\mathrm{SAD}} = 7.0$. This results in an effective propagation distance of $z \to z/M \approx 8.6\,\mathrm{mm}$ and an effective pixelsize of $p \to p/M \approx 0.96\,\mu\mathrm{m}$. These parameters define the single-material phase-retrieval algorithm of Eq. (5.46) and the physical mean model of Eq. (5.18).

The proposed approach for phase retrieval obtained by solving Eq. (5.48) is shown in Figure 5.1 (b). The edges of the sample are sharply depicted without any overshoots, as the interference effects are accounted for in the physical mean model. In addition, the noise level is reduced compared to the measured intensity shown in Figure 5.1 (a). In Figure 5.1 (c) a

comparison to the trace obtained by the single-material phase-retrieval algorithm according to Eq. (5.46) is shown. In the upper part, the line profile of the central row of the trace as depicted in Figure 5.1 (b) is plotted in blue. In comparison, the same profile of the trace obtained with the single-material phase-retrieval algorithm is plotted in red. At the lower part, the difference between both profiles is depicted. Small differences are visible at the edges of the sample, which can be explained by neglecting the cross terms given by Eq. (5.5) in the optimization-based approach.

In conclusion, the proposed optimization-based approach can accurately recover the trace of the sample. The differences to the single-material phase-retrieval algorithm are small. In Chapter 6, a physical mean model is discussed which coincides with the single-material phase-retrieval algorithm and thus does not rely on the simplification introduced by Eq. (5.5). Compared to the analytical single-material phase-retrieval algorithm, the above iterative framework allows for more versatile physical mean models, which model additional effects of the image formation process, like for instance source or detector blur. Again, extensions of the physical mean model are discussed in Chapter 6. In addition, prior knowledge in the form of regularization techniques are possible as well as a more accurate modeling of the noise properties in the likelihood by including statistical weights.

5.5. PB-CT using the homogeneity assumption at a laboratory source

After validating the physical mean model in the previous section, a reconstruction algorithm for PB-CT is constructed, based on the results derived in Section 5.3. A comparison to two other reconstruction approaches for PB-CT is performed, by measuring a self-built sample at the microtomography system introduced in the previous section. Parts of the following results have been published in [Hehn et al., 2018b].

5.5.1. Phase retrieval and reconstruction algorithms

Again, before deriving the MAP reconstruction approach for PB-CT, the conventional two-step reconstruction approaches are introduced, which are used for comparison.

Two-step reconstruction algorithms

As already outlined in the introduction of this chapter, the conventional reconstruction approach consists of two parts. First, the traces are recovered for every measurement independently using Paganin's single-material phase-retrieval algorithm (PAG) given by Eq. (5.46). Subsequently, tomographic reconstruction of the traces is applied. This is analogous to the reconstruction of the line-integrals in attenuation-based CT. Thereby, one uses the FBP algorithm given by Eq. (4.27) and Eq. (4.28). This approach is illustrated in Figure 5.2 by the black arrows.

However, for tomographic reconstruction of the traces, SIR approaches can also be employed as illustrated in Figure 5.2. Thereby, one models only the projection operation in the physical mean model given by Eq. (4.46) and the objective function is given according to Eq. (4.49) by

$$\mathcal{C}(\mathbf{x}|\mathbf{t}) = \frac{1}{2}\left(\mathbf{t} - \mathbf{A}\mathbf{x}\right)^T \mathbf{W}\left(\mathbf{t} - \mathbf{A}\mathbf{x}\right) + \beta\mathcal{R}_{\ell_1}(\mathbf{x}), \tag{5.49}$$

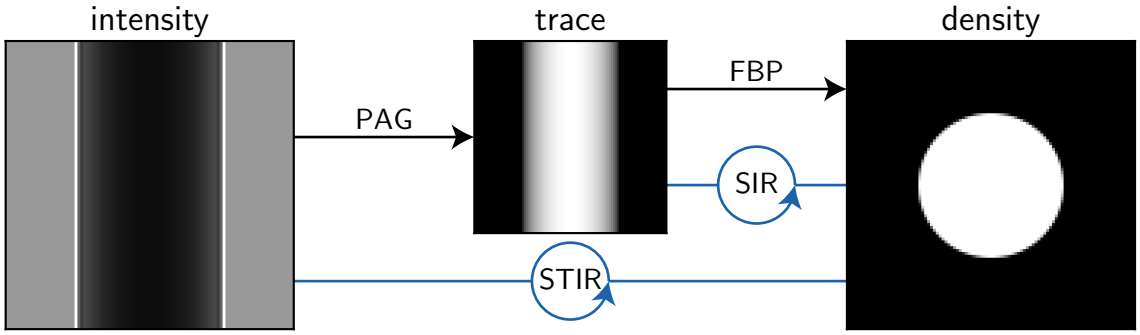

Figure 5.2.: Overview of reconstruction algorithms for tomography. Depicted from left to right is the measured intensity, the retrieved trace and the reconstructed density of a homogeneous cylinder. The conventional reconstruction approach shown in black consists of two parts. First, the trace is retrieved from the measured intensity using the single-material phase-retrieval algorithm (PAG) given by Eq. (5.46). Second, tomographic reconstruction is performed on the traces for every view using the FBP given by Eq. (4.27) and Eq. (4.28). However, tomographic reconstruction can also be performed using iterative methods by minimizing Eq. (5.49) (SIR). Finally, modeling the interactions within the sample as well as the propagation effects, the density of the sample can directly be reconstructed from the measured intensities by solving Eq. (5.51) (STIR). This figure was adapted from [Hehn et al., 2018b].

where $\mathbf{x}$ in this case is the three-dimensional distribution of the sample. As outlined subsequently, TV regularization is employed. One distinct problem with this approach is that the weights $\mathbf{W}$ cannot be approximated by the intensity measurements, as the previous phase retrieval step alters the statistical properties of the measurements significantly. Most importantly, the statistical weights cannot be approximated by a diagonal matrix as the phase retrieval correlates noise significantly. Due to the non-linearity of the logarithm, the noise properties of the traces can only be approximated, as shown for attenuation-based CT in the appendix of [Tilley et al., 2016a]. The topic of modeling noise correlations will be further discussed in Chapter 6.

Optimization-based reconstruction approach

In the following, a MAP approach suitable for recovering the three-dimensional distribution of the sample in PB-CT is proposed. As outlined in Section 4.3, the core idea of the MAP approach is to maximize the posterior distribution to give the most probable sample given the acquired measurements. Therefore, the negative log-likelihood is modeled according to Eq. (5.29) using $\mathcal{L}(\mathbf{y}|\mathbf{x}) = \mathcal{L}(\mathbf{y}|\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\delta}(\mathbf{x}))$ and the measurements are modeled using the physical mean model of Eq. (5.39), which employs the homogeneity assumption. In addition, a regularization term, given by Eq. (4.51), is required, as the tomographic reconstruction problem is ill-posed. As the regularization term is applied on the volume, the TV penalty given by Eq. (4.55) is well suited as it enforces sparse solutions. By employing the homogeneity assumption, the assumption of having piece-wise constant structures, which TV regularization evolves around, seems to be particularly well suited. By contrast, as discussed in Chapter 6, quadratic regularization might bias the interference term. The

objective function can be written as

$$\boxed{\mathcal{C}(\mathbf{x}|\mathbf{y}) = \frac{1}{2}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x})\right)^T \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x})\right) + \beta \mathcal{R}_{\ell_1}(\mathbf{x})} \tag{5.50}$$

and the MAP estimate of the volume is given by

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \mathcal{C}(\mathbf{x}|\mathbf{y}). \tag{5.51}$$

In contrast to the above two step approach, the statistical weights can be directly estimated from the measured intensities according to $\mathbf{W} = \mathcal{D}[\mathbf{y}^{-1}]$. In the following, this approach is referred to as statistical TIE-based iterative reconstruction (STIR) as it integrates a model for the interference effects based on the TIE within a SIR framework for tomographic reconstruction.

This reconstruction approach has several advantages compared to the above mentioned two-step approaches. First, the phase retrieval of the individual views is coupled over the volume. Second, the noise properties of the measured intensities can be modeled directly. As noise is realized not until the actual detection processes in the detector and thus the propagation of the X-rays does not influence the photon statistics, one does not rely on additional approximation modeling the noise properties after phase retrieval. Consequently, the assumption of having uncorrelated noise is a good approximation. In Chapter 6, noise correlations in PB-CT will be addressed thoroughly. Third, compared to the PAG + FBP approach, prior knowledge about the sample can be imposed further, such as the assumption of piece-wise constant structures.

5.5.2. Sample and experimental parameters

For comparing the proposed STIR algorithm to the two-step approaches (PAG + FBP and PAG + SIR), the following sample was created. A perfusion tube made of polyethylene (PE) defines the boundaries of the sample. Within the tube, poly(methyl methacrylate) (PMMA) spheres have been crushed with pliers to yield more complex structures. In addition, a tungsten thread was added, which has different attenuating and phase-shifting properties compared to PE and PMMA, and thus violates the homogeneity assumption.

The sample was measured at the Xradia 500 Versa microtomography system, detailed in Subsection 5.4.2. The SAD was set to $d_{\mathrm{SAD}} = 45\,\mathrm{mm}$ and the SDD was set to $d_{\mathrm{SDD}} = 119\,\mathrm{mm}$, resulting in a propagation distance of $z = 74\,\mathrm{mm}$ and a magnification of $M \approx 2.64$. The detector was again used with the fourfold magnification lenses and with 2×2 spatial binning of the detector pixels, resulting in the effective pixelsize of $p = 6.75\,\mu\mathrm{m}$.

The reference corrected intensity of the first view is depicted in Figure 5.3 (a). PE and PMMA have similar attenuating and phase-shifting properties, while the tungsten thread is highly attenuating.

5.5.3. Results and discussion

To compare the different reconstruction approaches detailed in Figure 5.2, the parameters of the cone-beam geometry have to be transferred to the parallel-beam geometry according to the Fresnel scaling theorem using the magnification. The effective propagation distance is given by $z \to z/M \approx 28\,\mathrm{mm}$ and the effective pixelsize is $p \to p/M \approx 2.56\,\mu\mathrm{m}$. The material specific constants were chosen for PMMA, which were again extracted from the xraylib library [Schoonjans et al., 2011].

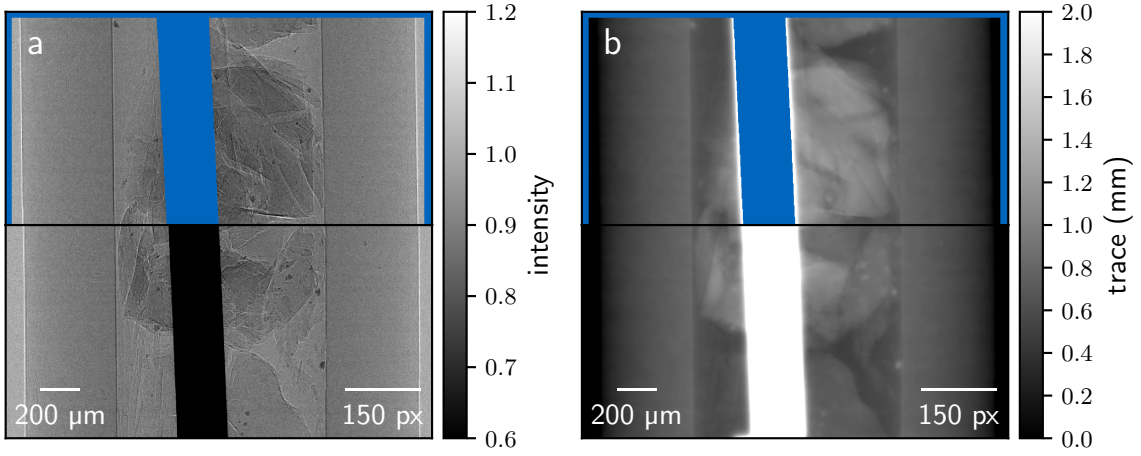

Figure 5.3.: Intensity and trace of the tomographic dataset for the first projection. In (a) the measured intensity of the first projection relative to the reference intensity is shown. In (b) the trace recovered by the single-material phase-retrieval algorithm is depicted. In addition, in the upper parts, the regions, where the statistical weights are set to zero, are shown. This figure was adapted from [Hehn et al., 2018b].

With these parameters, the traces can be retrieved using PAG given by Eq. (5.46). The respective trace of the first view depicted in Figure 5.3 (a) is shown in Figure 5.3 (b). The PE and PMMA parts look sharp, however the tungsten thread in smeared out, covering features in its vicinity, which cannot be recovered in a subsequent tomographic reconstruction. By assuming the attenuating and phase-shifting properties of PMMA, the ratio of δ/μ in the denominator of the phase-retrieval algorithm in Eq. (5.46) is too high for tungsten, which results in the smearing of the thread.

In Figure 5.4, a transverse and longitudinal cut through the reconstructed volume obtained by reconstructing the previously recovered traces using the FBP algorithm is shown. From the tungsten thread, severe streak artifacts arise diminishing the image quality of the whole reconstruction. In addition, due to the smearing of the tungsten thread, features in its vicinity cannot be visualized.

Within the likelihood terms, the statistical weights have been introduced as means to account for the reliability of the acquired measurements with respect to noise arising mainly from the photon noise. However, this is only valid if the physical mean model correctly models the image formation. In the presented case, the physical mean model does not correctly model the measurements[3] at regions where the homogeneity assumption is violated, e.g. where the measurements are corrupted by the influence of the tungsten thread. Thus, no statements about the likelihood of the respective measurements can be made for this case. However, the statistical weights can also be interpreted as a measure for the reliability of the physical mean model. Mathematically, this can be achieved by setting the weights to zero for all measurements, which have been corrupted by the tungsten thread. As a consequence, the resulting likelihood again correctly models the likelihood of the measurements and measurements, which cannot be described by the physical mean model, are excluded from the

[3] The term “measurements” refers to the measured intensities at individual pixels.

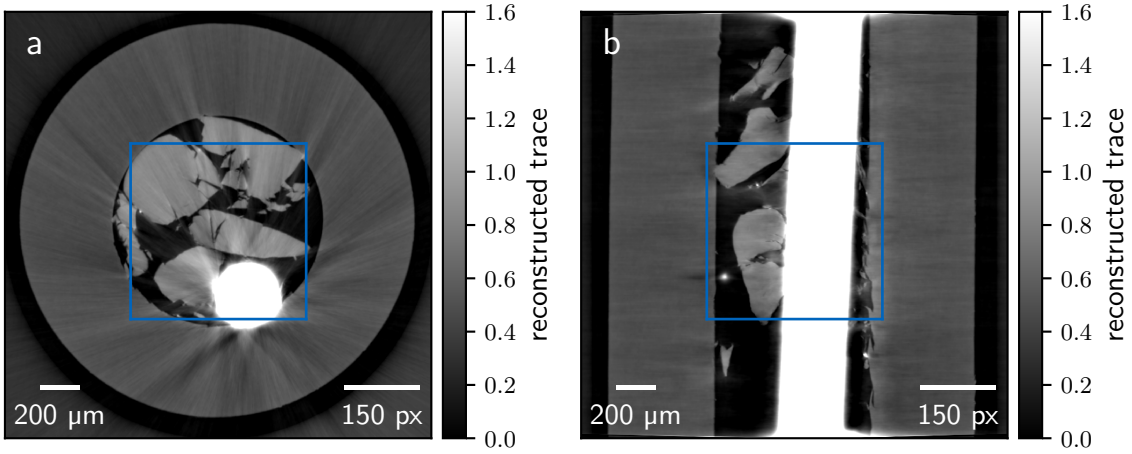

Figure 5.4.: Conventional two-step reconstruction approach. In (a) a transverse slice of the FBP reconstruction of the traces is shown. In (b) a corresponding longitudinal slice is depicted. The blue squares denote the zoomed regions of Figure 5.5. This figure was adapted from [Hehn et al., 2018b].

reconstruction.

Regions where tungsten is present can be obtained by simple thresholding. Thus, the weights are set to zero at regions, where the reference corrected intensities have transmission values lower than 0.5. In order to ensure that the whole tungsten thread has been selected, the weights in the immediate neighborhood are also set to zero. In addition, to avoid effects at the borders, the weights for the 15 outermost pixels at the detector borders are also set to zero. The masked regions are depicted in Figure 5.3 in blue in the upper halves for both the measured intensities as well as the recovered traces.

To avoid biasing the comparison between the PAG + SIR and STIR approach, the statistical weights outside the masked regions were set to one. The regularization strength for both SIR and STIR was chosen such that the noise level in the reconstruction is comparable. This was done empirically as the reconstructions are diminished by artifacts and thus quantitative measures of the noise levels are corrupted. An empty volume was used as initialization of the volume as the PAG + FBP reconstruction shown in Figure 5.4 is corrupted by artifacts. For minimizing the respective objective functions, the L-BFGS algorithm outlined in Subsection 4.3.5 was used. However, as an initial guess of zeros was chosen for the volume, a comparably large amount of iterations is required, namely 100 iterations for the SIR approach and 800 for the STIR approach.

The regions marked in Figure 5.4 by blue rectangles are shown in Figure 5.5 for the three reconstruction approaches. The reconstructions shown in (a) and (d) were obtained by the PAG + FBP as already depicted in Figure 5.4. In (b) and (e), the reconstructions obtained by PAG + SIR are shown. For tomographic reconstruction, Eq. (5.49) was minimized with the parameters specified above. Finally, the integrated STIR approach results in the reconstructions shown in (c) and (f). The corresponding objective function is given by Eq. (5.50).

As discussed above, in the reconstruction using the FBP, the tungsten thread is smeared out covering features in its vicinity, which cannot be depicted anymore. In addition, streak like artifacts arise from the tungsten thread affecting the whole reconstruction. By replacing

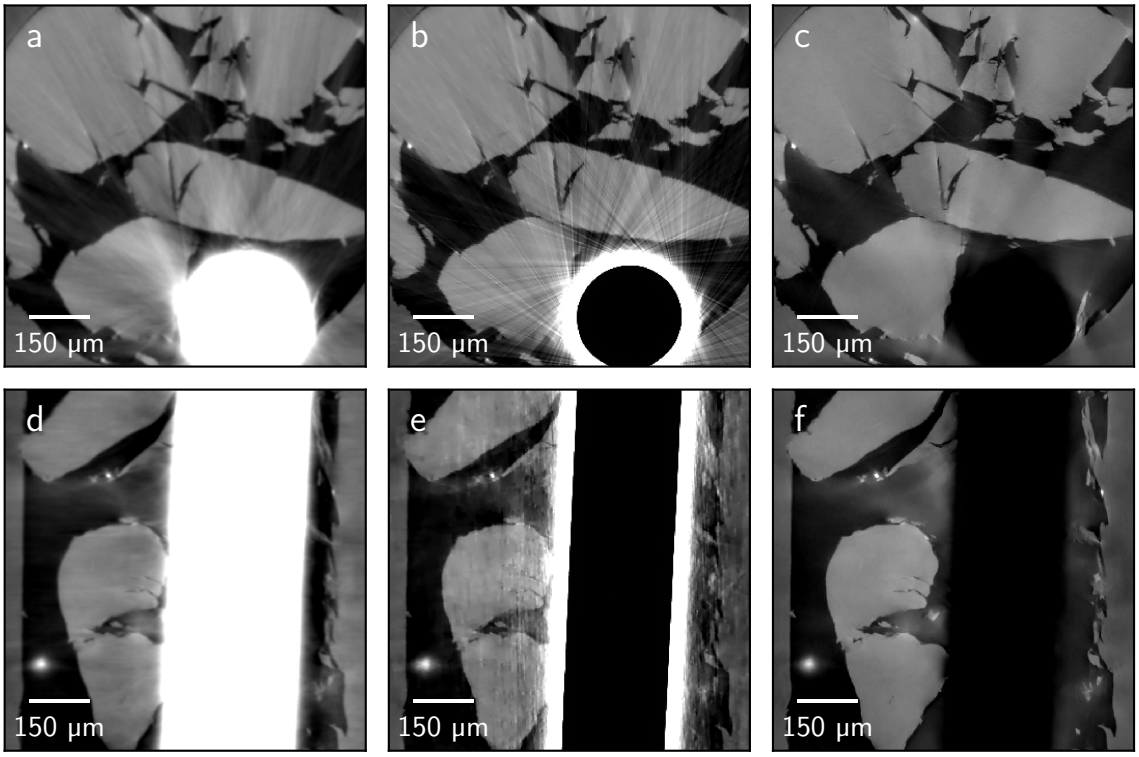

Figure 5.5.: Comparison of different reconstruction techniques. In (a)-(c) the transverse slice marked in Figure 5.4 (a) is depicted for different reconstruction techniques. Likewise, in (d)-(f) the corresponding longitudinal slice marked in Figure 5.4 (b) is shown. The reconstructions of the first column (a), (d) were obtained by an FBP subsequent to single-material phase retrieval. This corresponds to the results depicted in Figure 5.4. In the middle column (b), (e) statistical iterative reconstruction was used to recover the distribution of the sample from the traces obtained by the single-material phase-retrieval algorithm. The statistical weights as depicted in Figure 5.3 were used. Finally, in (c), (f) the measured intensities were directly reconstructed by minimizing the objective function given by Eq. (5.50) with the same statistical weights. This figure was adapted from [Hehn et al., 2018b].

FBP by SIR, the weights depicted in Figure 5.3 (b) are leveraged to mask those parts, where the tungsten is present on the measurements. However, the smearing of this thread is a consequence of the earlier phase retrieval. Thus, information in the vicinity of the tungsten thread is already lost and cannot be recovered in the subsequent tomographic reconstruction. Due to TV regularization, the edges are sharper. However, the streak artifacts are more prominent as well, due to the edge-preserving properties of the regularization. Finally, STIR, which utilizes the same mask as shown in Figure 5.3 (a) directly on the acquired measurements, removes the tungsten thread already in the integrated phase-retrieval step. Thus, features in its vicinity can be recovered. As a consequence, no streak artifacts are present. In addition, the assumption of having piece-wise constant structures can be fully leveraged, resulting in a sharp depiction of the object.

In conclusion, benefits of the proposed STIR approach compared to two-step reconstruc-

tion approaches have been demonstrated using a microtomography system. In particular, the one-step approach can better account for artifacts introduced by materials, which violate the homogeneity assumption. In addition, compared to the conventional PAG + FBP approach, the image quality benefits significantly from TV regularization. Furthermore, it could be shown that the proposed approach even works for a polychromatic source.

5.6. Application of PB-CT for imaging biological samples using synchrotron radiation

One prominent application of PB-CT is the visualization of biological samples, as this imaging technique provides high resolution and high contrast even for small samples. However, if highly attenuating components such as implants are present in the sample, the image quality suffers significantly, as demonstrated in the previous section.

The visualization of animal cochlear implants is one such imaging scenario, as the electrode array within the implant has distinctly different attenuating and phase-shifting properties than the cochlear itself. The visualization of the exact position of the implant within the cochlea is crucial as the insertion of the implant into the narrow turns of the cochlea often causes subtle damage, which influences the experimental results. For instance, the electrode of the implant penetrates the basilar membrane and disrupts the organ of Corti, which lies on top. Or trauma caused by inserting the electrode in chronic experiments triggers soft-tissue growth. Unfortunately, the regions in the vicinity of the implant, where damage is most likely to occur and thus highest resolution is desired, are most corrupted by imaging artifacts. Retracting the implant from the cochlea is likely to cause additional damages, therefore, in-situ imaging of inserted implants is highly desirable. [Wanna et al., 2014, O'Connell et al., 2017, McJunkin et al., 2018]

In this section the results of this chapter are applied to the reconstruction of a guinea pig with cochlea implant using synchrotron radiation (SR). SR employs different characteristics compared to X-rays, generated at laboratory sources, which are particularly well suited to the developed physical mean model. The following results have been published in [Hehn et al., 2018a].

5.6.1. Large-scale synchrotron facilities

Synchrotron facilities have emerged to provide monochromatic X-rays with high flux and spatial coherence enabling a variety of X-ray methods, which can not be implemented using laboratory X-ray sources, including PBI using multiple distances or PBI methods utilizing imaging regimes beyond the edge-enhancement regime.

SR is created by accelerating electrons to energies of a few GeV. The electrons are then kept in a storage ring, which can have circumferences of several hundreds of meters, using bending magnets with field strengths of just below one Tesla. The radiation employed for X-ray experiments is created using so-called insertion devices, such as undulators, which are located at straight sections between the bending magnets. These undulators consist of alternating magnetic fields resulting in rapid oscillations of the electrons. These oscillations cause the emission of X-rays. The magnetic fields are designed in a way that peaks with a bandwidth of a fraction of a keV at specific energies are obtained. To further refine the energy of the X-rays, monochromator crystals are used that reflect only the desired energy using Bragg crystals. [Als-Nielsen and McMorrow, 2011, Schaff, 2017, Willmott, 2011]

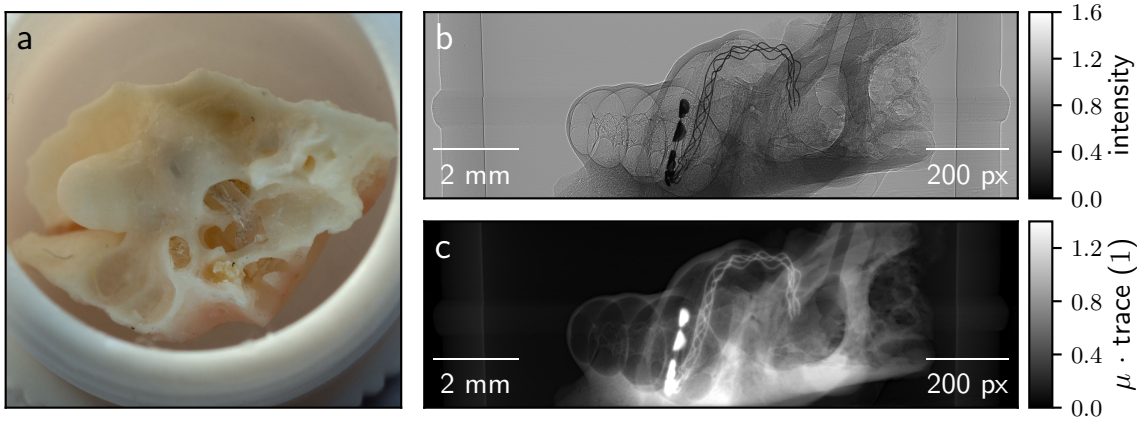

Figure 5.6.: Cochlea sample and virtual projections. In (a) a photograph of the guinea pig cochlea with implant is depicted. In (b) a virtual intensity measurement is depicted after preprocessing. In (c) the trace of the measurement in (b) recovered with PAG is shown. This figure was adapted from [Hehn et al., 2018a].

5.6.2. Sample, experimental setup and preprocessing

For sample preparation, the cochlea was boiled out and bleached with hydrogen peroxide. Afterwards, the implant was inserted into the cochlea. The electrodes of the implant are made from Platinum and the wire material consists of Platinum-Iridium embedded in Silicon. The implant was fabricated by MED-EL GmbH and is identical to cochlear implants used in patients. A photograph of the sample with the cochlear implant can be seen in Figure 5.6 (a).

The sample was measured at the micro-tomography end-station of the imaging beamline P05 at PETRA III at DESY (Hamburg, Germany) operated by the Helmholtz-Zentrum Geestacht. Details about the experimental setup can be found in [Wilde et al., 2016]. The monochromatic X-ray energy was set to 35 keV. The detection system consists of a CCD camera with 12.0 µm. Optics with a magnification factor of 4.96 were used, resulting in an effective pixelsize of around 2.42 µm. The propagation distance from the sample to the detector was set to 1 m. The exposure time was set to 0.1 s. An off-axis tomography was performed to extend the limited FOV of 7.2 mm in the horizontal direction. Therefore, the detector was shifted such that the projected rotation-axis was located near the edge of the detector and a full-scan was performed. Due to the parallel-beam geometry, one can obtain a virtual half-scan with twice the FOV in horizontal direction by merging opposed projections. In addition, to double the FOV in vertical direction, which was limited to 2 mm, two off-axis tomographies were performed at different heights.

The intensity of the incident beam drifted over time due to drifts in the monochromator. To correct for these fluctuations, a series of reference intensities without the sample were acquired. For each projection, the variance of all available reference intensities was minimized in a sample-free region to select the best reference intensity to correct for the intensity inhomogeneities. The rotation axis was slightly rotated around the optical axis with respect to the detector. To correct for this rotation, cross-correlations between overlapping regions of opposing views were performed. This information was also used for creating the virtual half-scan projection from each pair of opposed views. In addition, the projections obtained at different heights were merged by maximizing the sum of overlapping regions. Finally, the merged virtual half-scan projections were cropped and binned by a factor of four, resulting

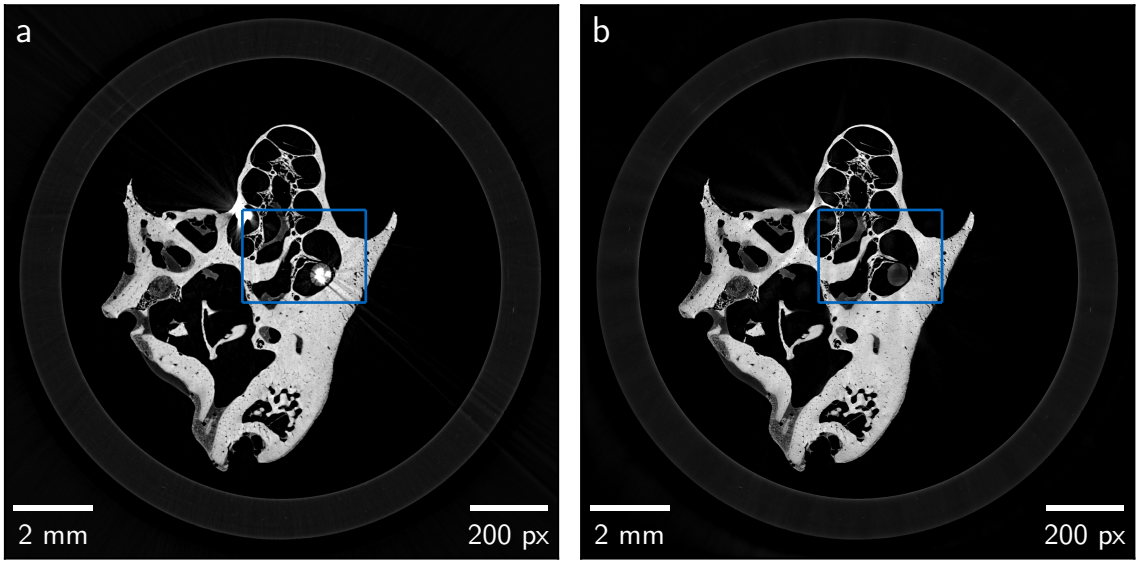

Figure 5.7.: Reconstructions of the cochlea sample. In (a) a slice of the reconstruction is depicted for the conventional PAG + FBP reconstruction approach. In (b) the corresponding slice reconstructed using STIR is shown. The blue rectangles denote the zoomed regions of Figure 5.8. The results are depicted using the same linear gray scale. This figure was adapted from [Hehn et al., 2018a].

in 1434×450 pixels and 1500 virtual views with a pixelsize of 9.68 µm. In Figure 5.6 (b), an example for the virtual projection is depicted. This virtual projection consists of four measurements, individually corrected for the incident flux, rotated, stitched, cropped and binned.

5.6.3. Results and discussion

The qualitative parameters for the homogeneous material of the cochlea were set to $\mu = 1\,\mathrm{m}^{-1}$ and $\delta = 10^{-9}$. The choice for setting μ to unity results in no further scaling of the recovered traces.

The traces recovered by PAG, using Eq. (5.46) with the parameters specified above are shown in Figure 5.6 (c) for the intensity measurement depicted in Figure 5.6 (b). The traces at the implant are slightly enlarged and have very high values. The subsequent tomographic reconstruction of the traces is performed by FBP using Eq. (4.27) and Eq. (4.24). In Figure 5.7 (a) a reconstructed slice is depicted and the zoom indicated by the blue rectangle is shown in Figure 5.8 (a). One finds streak artifacts arising from the implant.

For the STIR approach, the objective is similar to Eq. (5.50) used for reconstruction in the previous section. It reads

$$\mathcal{C}(\mathbf{x}|\mathbf{y}) = \frac{1}{2}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x})\right)^T \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x})\right) + \beta\mathcal{R}_\gamma(\mathbf{x}), \tag{5.52}$$

where the TV penalty function was replaced by the Huber penalty, defined by Eq. (4.54), where the transition parameter γ was estimated from the background noise using the conventional reconstruction approach. The strength of the regularization was set to $5 \cdot 10^{-5}$ by visual inspection. Again, as optimization routine the L-BFGS algorithm with 2000 iterations

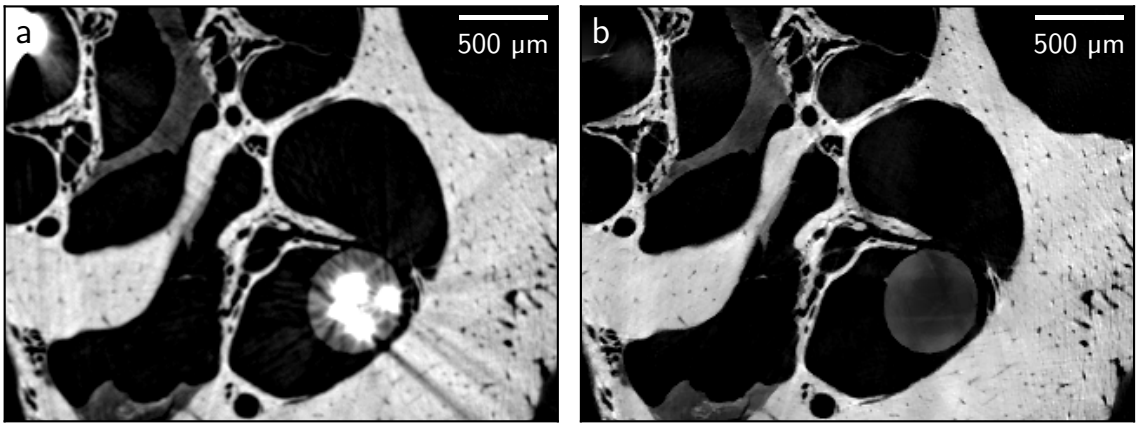

Figure 5.8.: Zooms at regions around the implant. In (a) and (b) the corresponding regions indicated in Figure 5.7 by the blue rectanges are depicted. In (a) the conventional PAG + FBP approach was used, while in (b) STIR was employed. This figure was adapted from [Hehn et al., 2018a].

was used to ensure sufficiently converged estimates. As the conventional reconstruction is corrupted by streak artifacts, additional 250 iterations of the L-BFGS were performed initialized with zeros and a strong regularization strength of $\beta = 10^{-3}$ to obtain a low-frequency representation of the sample. The weights were specified by $\mathbf{W} = \mathcal{D}^{-1}[\mathbf{y}]$ according to the assumption of the measured intensities being Poisson distributed. As discussed in the previous section in detail, the negative log-likelihood term is not valid for measurements that are not described by the established physical mean model. Thus, the likelihood term was set to zero for measurements that are corrupted by the implant. The corresponding measurements were selected by thresholding the direct FBP reconstruction of the measured intensities (not shown here). The selected region was then slightly dilated and forward projected using Eq. (4.23). One could also have chosen to threshold the measured intensities directly. However, the thresholding of the volume is less prone to noise and consistent with every view. The respective reconstructed slice and the corresponding zoom are depicted in Figure 5.7 (b) and Figure 5.8 (b) respectively.

Compared to the conventional reconstruction approach, the electrodes of the implant are removed from the reconstruction and thus there are no streak artifacts present. In addition, the noise and resolution properties of the reconstruction using the STIR algorithm are superior, mainly due to the regularization, which suppresses noise while maintaining the resolution at the edges, which can best be seen from Figure 5.8.

Lastly, renderings from both reconstruction approaches were created[4] using Avizo Fire 9.1 (Thermo Fisher Scientific, USA). The renderings are depicted in Figure 5.9. The position of the electrode can be segmented from the direct FBP reconstruction of the intensities and is depicted in red. The rendering based on the reconstruction using the conventional reconstruction approach shows streak artifacts and the position of the implant cannot be visualized as emphasized by the zoom. By contrast, using the reconstruction from the STIR approach yields an artifact-free rendering and the position of the implant can be accurately depicted. Most importantly, regions around the implant are accurately visualized. Moreover, the resolution properties are enhanced.

[4] The renderings were created by Regine Gradl.

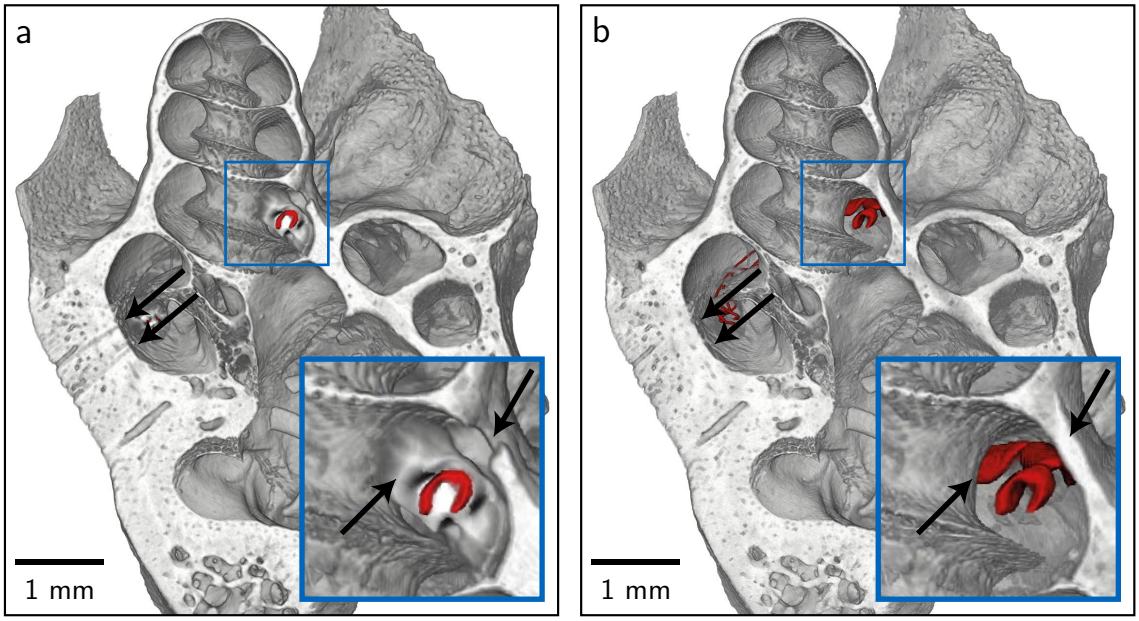

Figure 5.9.: Three-dimensional renderings of the cochlea sample with implant. In red, the position of the implant segmented from a direct FBP reconstruction of the measured intensities is shown for both reconstruction approaches. In (a), the reconstruction obtained by the conventional PAG + FBP approach was used for the rendering, while for rendering shown in (b), the iterative STIR approach was used. The inserts zoom at regions around the implant detailing degradation caused by the artifacts arising from the PAG step in the conventional reconstruction approach. This figure was adapted from [Hehn et al., 2018a].

The streak artifacts are predominantly related to the violation of the homogeneity assumption, as they are not present in the direct FBP reconstruction of the measured intensities. However, many effects can cause similar streak artifacts. These effects include beam hardening, photon starvation or noise from low photon counts, scattering, motion artifacts and nonlinear partial volume effects (NLPV), also known as exponential edge-gradient effects. In literature, these effects have been studied thouroughly in particular in the context of medical imaging [Boas and Fleischmann, 2011, Man et al., 1999, Rinkel et al., 2008, Stayman et al., 2013]. Beam hardening should be negligible, due to the monochromaticity of the source. With sufficient transmission behind the implants, the artifacts should also not be attributed to beam starvation. However, as seen in the previous section, the proposed approach can naturally account for beam starvation as well. Scatter should also be reduced in PBI, due to the increased distance between the object and the detector. The inanimate sample ideally should also not introduce motion artifacts. For the presented sample, NLPV effects have an influence on the image quality. However, it was shown in [Stayman et al., 2013] for attenuation-based imaging that MAP approaches are capable of alleviating these effects by a finer sampling of the image volume at regions around the implant. Although resulting in an increase of computational cost, this could be integrated in the proposed approach. Closely related to NLPV effects are under-sampling artifacts, if the number of views is small. In Section 4.2.1 the Nyquist condition was introduced as a guideline for the required number of views. However, several publications have demon-

strated that iterative reconstruction techniques, which employ regularization techniques, can accurately reconstruct images from a limited number of views [Sidky and Pan, 2008], also in the content of PBI [Sidky et al., 2010]. In [Vidal et al., 2005] an overview of additional sources for artifacts in synchrotron micro-tomography, not restricted to streak artifacts, can be found.

In summary, a guinea cochlea with implant was examined using PB-CT and SR detailing the challenges arising from the implant, which violated the homogeneity assumption. It was demonstrated that by integrating the phase retrieval into tomographic reconstruction inside a MAP approach, the artifacts arising from the implant can be removed. Overall, this approach resulted in superior image quality in terms of resolution compared to the conventional two-step reconstruction approach. In the particular case, this might prove beneficial to evaluate whether the insertion of the implant resulted in subtle damages of the cochlea. This approach can be transferred to similar imaging tasks that involve biological samples, which have additional features violating the homogeneity assumption such as metal implants or bones. Thereby, one can leverage the high-contrast between materials with similar attenuating properties such as soft-tissue using PB-CT without being diminished by the presence of highly absorbing features.

5.7. Conclusion

In this chapter, various SIR algorithms have been investigated for PBI and PB-CT. Therefore, a physical mean model was developed, consisting of a term accounting for the attenuating properties of the sample and a term describing the interference effects due to the phase shifts induced by the sample. First, a SIR framework for PBI was developed, which can potentially be used to independently recover the amplitude and phase in the object plane. As this problem, however, is ill-posed, two special cases were investigated. First, the homogeneity assumption was employed. The resulting algorithm is in good agreement with the results obtained by the single-material phase-retrieval algorithm [Paganin et al., 2002] as validated experimentally. In addition, in the limit of pure phase objects, the analytic solution of the proposed algorithm coincides with the algorithm proposed by Bronnikov [Bronnikov, 1999]. The above approach was then extended to PB-CT, thereby incorporating phase retrieval within tomographic reconstruction. Again, two special cases were investigated. For an incoherent source, the proposed physical mean model reduces to the mean model given by Eq. (4.44). Furthermore, two applications of the proposed algorithm for tomography, which incorporates the homogeneity assumption, were shown at a laboratory source and at a synchrotron. Thereby, the versatility of the proposed SIR approach allowed to individually exclude materials violating the homogeneity assumption from the reconstruction leading to distinct improvements in image quality.

The proposed algorithms have various advantages compared to two step methods, which could be investigated in more detail. For instance, the bias introduced due to Jensen's inequality given by Eq. (4.47) should be reduced in the proposed framework. Moreover, two step approaches do not allow to easily use the noise properties after phase retrieval as they are highly correlated. Means to leverage this information in a subsequent tomographic reconstruction would be interesting. In [Tilley et al., 2016a], the influence of the logarithm and an additional deblurring operation has been investigated for conventional CT, which could potentially be transferred to PB-CT.

In addition, a challenging task would be to investigate means to more independently re-

construct the linear attenuation coefficient and refractive index decrement. One approach would be to investigate different regularization approaches to mitigate the homogeneity assumption. Thereby, regularization techniques in the field of dual-energy CT are promising, which require that the two reconstructed images have many regions, where edges are located at the same position [Huh and Fessler, 2011]. On the other hand, regularization techniques of multi-distance methods could be investigated further [Langer et al., 2014].

6. Modeling the source and the detector in homogeneous PBI and PB-CT

In the previous chapter, SIR approaches for PBI and PB-CT have been investigated. In particular, incorporating the interference effects into a physical mean model and including a statistical description of the measurements enabled removing materials not obeying the homogeneity assumption from the reconstruction. The artifacts induced by such materials hinder the effective use of PB-CT for various applications at laboratory environments as well as synchrotron facilities.

This chapter addresses the challenges that arise when translating PBI and PB-CT to laboratory environments. These include the reduced spatial coherence of available X-ray sources and the reduced flux resulting in reduced spatial resolution and high noise levels. The reduced flux makes the use of efficient detectors desirable. These detectors usually feature thicker scintillators, which spread out the signal and further reduce spatial resolution. The blurring of the source and the detector diminishes the interference effects crucially. However, the interference effects hold the information about the phase-shifting properties of the sample, which PBI and PB-CT try to leverage. Thus, methods to deconvolve the blurring of the source and the detector are desired to enhance the quality of the interference effects. However, deconvolution approaches struggle in particular in the presence of high noise levels inherent to laboratory environments.

In conventional attenuation-based CT, most prior efforts to model and correct for system blur have concentrated on preprocessing techniques, applied on the measured intensities prior to tomographic reconstruction [La Rivière et al., 2006, Riviere and Vargas, 2008, Zhang et al., 2014b]. However, during preprocessing not all information present in the measured intensities is preserved and thus the lost information cannot be leveraged in the subsequent tomographic reconstruction. In addition, the noise properties of the preprocessed intensities are usually altered significantly. Thus, in order to include the noise properties in the tomographic reconstruction, these properties have to be carefully propagated through all preprocessing steps, which may require additional simplification [Tilley et al., 2016a].

MBIR methods that account for some kind of system blur by directly including a model of the blur into the physical mean model have already been developed for single-photon emission computed tomography (SPECT) [Yu et al., 2000, Feng et al., 2006] and recently also for digital breast tomosynthesis [Zheng et al., 2018] as well as flat-panel cone-beam CT [Tilley et al., 2016a, Hashemi et al., 2017, Tilley et al., 2018a]. In the following, building upon the results of the last chapter, MBIR approaches for PBI and PB-CT are developed, which account for the influence of the reduced spatial coherence, the spread of scintillator-based detectors as well as the high noise levels. On the one hand, the developed algorithms build upon and extend recent advances in analytical phase-retrieval approaches for PBI, which take the effects of spatial coherence into account [Gureyev et al., 2009, Beltran et al., 2018]. Furthermore, a connection between analytical phase retrieval and regularized image denoising algorithms is outlined [Rudin et al., 1992]. On the other hand, recent advances in MBIR methods for conventional CT are transferred to PB-CT, which

among others model the response of the detector and its influence on noise. For the proposed algorithms, a new physical mean model is devised, which includes models of the source, the detector and interference effects. In addition, an expression for the full noise covariance statistics is derived and utilized in the reconstruction algorithms.

First, an analytical model of the image formation process is developed, which directly integrates the homogeneity assumption. Next, a physical mean model as well as the full covariance statistics are derived. Subsequently, a model-based iterative phase-retrieval algorithm for PBI is devised, which analytically extends the single-material phase-retrieval algorithm to include the effects of the source, detector and prior knowledge about the sample in the form of regularization techniques. Moreover, a connection between regularized image denoising and phase retrieval is shown. Finally, an MBIR algorithm for PB-CT is derived to tackle the above-mentioned challenges of translating this technique into laboratory environments. The algorithms are validated using a simulation study as well as an experimental study at a compact inverse Compton source.

Parts of this chapter were submitted for publication [Hehn et al., 2019a]. Other parts are currently being prepared for publication.

6.1. Image formation description at laboratory environments

In the following, the different steps $(0)-(4)$ of the image formation as illustrated in Figure 6.1 are described. As in the previous chapter, the generated wavefield is initially assumed to be a spatially coherent monochromatic forward propagating wavefield. This wavefield is again decomposed into the envelope and the unscattered wave according to Eq. (2.6). In the following, the intensity and phase, which define the envelope according to Eq. (2.16), are discussed at different stages along the optical axis.

Between the source and the sample at position (0), the envelope of the generated wavefield is defined by its intensity $I_0(\mathbf{r}_\perp)$ as well as by $\phi_0(\mathbf{r}_\perp) = 0$ as there is no phase shift compared to the unscattered wave. The corresponding subscript indicates the position along the optical axis as shown in Figure 6.1.

Behind the sample at (1), the intensity of the envelope is attenuated due to photoelectric absorption and Compton scattering. Again, as discussed in Section 2.2, using the projection approximation, the intensity behind the sample is given by the Lambert-Beer law according to

$$I_1(\mathbf{r}_\perp) = I_0 e^{-\ell(\mathbf{r}_\perp)}. \tag{6.1}$$

In order to align notation with reconstruction algorithms for conventional attenuation-based CT, the influence of the sample is described in terms of the line-integrals, which relate to the linear attenuation coefficient $\mu(\mathbf{r}_\perp, z)$ according to

$$\ell(\mathbf{r}_\perp) = \int \mu(\mathbf{r}_\perp, z)\mathrm{d}z. \tag{6.2}$$

By employing the homogeneity assumption, discussed in Section 3.2, given by

$$\delta(\mathbf{r}) = \frac{\delta}{\mu}\mu(\mathbf{r}), \tag{6.3}$$

the phase shift induced by the sample can as well be directly expressed in terms of the line-integrals according to

$$\phi_1(\mathbf{r}_\perp) = -k \int \delta(\mathbf{r}_\perp, z)\mathrm{d}z = -k\frac{\delta}{\mu}\ell(\mathbf{r}_\perp). \tag{6.4}$$

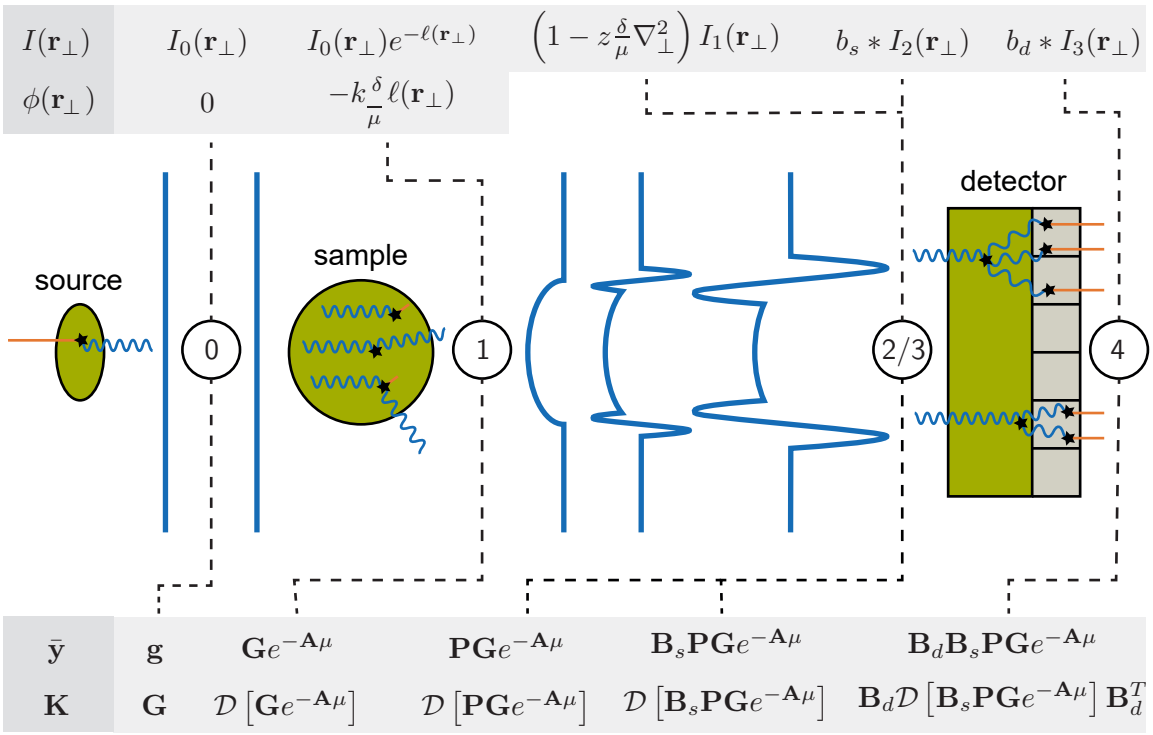

Figure 6.1.: Illustration of the image formation in PBI. From left to right, the source, sample and detector are depicted. In blue the intensity profile at different stages is sketched. On the top of the image, the analytic expressions for the intensity $I(\mathbf{r}_\perp)$ and phase $\phi(\mathbf{r}_\perp)$ of the wavefield are given at different stages. On the bottom, the mean vector $\bar{\mathbf{y}}$ and covariance matrix $\mathbf{K}$ used to describe the random vector of the measurements are denoted at the different positions. A similar depiction of the image formation in terms of the mean vector and covariance matrix for conventional attenuation-based CT has been given in [Tilley et al., 2016a]. An adapted version of this figure was submitted for publication [Hehn et al., 2019a].

Thus, similarly to conventional attenuation-based radiography, the influence of the sample on the wavefield behind the sample is entirely defined by its line-integrals. [Paganin et al., 2002]

Next, the evolution of the intensity from the sample to the detector at (2) is described by Eq. (2.43) according to

$$I_2(\mathbf{r}_\perp) = I_1(\mathbf{r}_\perp) - \frac{z}{k}\nabla_\perp \cdot \left(I_1(\mathbf{r}_\perp)\nabla_\perp\phi_1(\mathbf{r}_\perp)\right). \tag{6.5}$$

In Eq. (5.4), the second term was explicitly expanded to get an expression that entangles the influence of the attenuating and phase-shifting properties of the intensity at the detector

plane. Instead, one can directly employ Eq. (6.1) and Eq. (6.4) according to

$$I_2(\mathbf{r}_\perp) = I_0 e^{-\ell(\mathbf{r}_\perp)} + z\frac{\delta}{\mu}\nabla_\perp \cdot \left(I_0 e^{-\ell(\mathbf{r}_\perp)}\nabla_\perp \ell(\mathbf{r}_\perp)\right) \tag{6.6}$$

$$= I_0 e^{-\ell(\mathbf{r}_\perp)} - z\frac{\delta}{\mu}\nabla_\perp^2 I_0 e^{-\ell(\mathbf{r}_\perp)} \tag{6.7}$$

$$= \left(1 - z\frac{\delta}{\mu}\nabla_\perp^2\right) I_1(\mathbf{r}_\perp). \tag{6.8}$$

Due to the homogeneity assumption, the phase information of the wavefield in the sample plane is already encoded in $I_1(\mathbf{r}_\perp)$. This formulation coincides with the model used to derive the single-material phase-retrieval algorithm. [Paganin et al., 2002]

At position (3), the effect of the reduced spatial coherence of the source is now included. In Section 2.4, starting from the mutual coherence function, a relation for strictly monochromatic wavefields assuming an ensemble of plane wavefields was derived that captures the influence of partial coherence using the ensemble average according to

$$I_3(\mathbf{r}_\perp) = \langle I_2(\mathbf{r}_\perp)\rangle_\theta \tag{6.9}$$

$$= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} b_s(\mathbf{r}'_\perp) I_2(\mathbf{r}_\perp - \mathbf{r}'_\perp)\mathrm{d}\mathbf{r}'_\perp, \tag{6.10}$$

which can be expressed as a convolution, where b_s defines the respective abundance of the corresponding wavefields. As discussed later, this formulation has recently been used to analytically extend the single-material phase-retrieval algorithm to include effects of partial coherence. [Pfeiffer et al., 2005, Gureyev et al., 2009, Beltran et al., 2018]

Finally at (4), the influence of scintillator-based detectors is included in the image formation. These detectors consist of a scintillator material and a photodetector [Nikl, 2006]. Thereby, within the scinillator material, the individual X-rays are converted into visible light. However, the visible light is emitted with a broad angular distribution. This results in a signal spread over multiple pixels, which in return leads to additional blurring of the system. This blurring can again be modeled by a convolution according to

$$I_4(\mathbf{r}_\perp) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} b_d(\mathbf{r}'_\perp) I_3(\mathbf{r}_\perp - \mathbf{r}'_\perp)\mathrm{d}\mathbf{r}'_\perp, \tag{6.11}$$

where the response of the detector is described by b_d, which includes the blur and additional gain terms. More advanced approaches have been proposed to accurately model the response of the detector system. [Tilley et al., 2016a]

6.2. Discretization and full covariance statistics

The above analytical description of the image formation has not considered noise. In the following, the measurements are modeled as a random vector. At each position in the image formation process, the mean as well as the covariance matrix of the random vector are derived. This approach has been used for conventional attenuation-based CT in [Tilley et al., 2016a].

At position (0), the X-rays are assumed to be generated with a spatial distribution denoted by a vector $\mathbf{g}$. The generated X-rays are uncorrelated and follow a Poisson distribution resulting in a diagonal covariance matrix. Thus, the mean vector and the covariance matrix are given by

$$\bar{\mathbf{y}}_0 = \mathbf{g} \qquad \text{and} \qquad \mathbf{K}_0 = \mathbf{G}, \tag{6.12}$$

where $\mathbf{G} = \mathcal{D}[\mathbf{g}]$ is defined. After the interaction of the wavefield with the sample, the intensity is attenuated and the phase is shifted according to Eq. (6.1) and Eq. (6.4). As discussed in the previous section, it is sufficient to describe the intensity behind the sample, as it already encodes the information about the phase-shifting properties of the sample. The interaction with the sample does not correlate noise. Thus, the mean vector and covariance matrix at position (1) are given by

$$\bar{\mathbf{y}}_1 = \mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}} \qquad \text{and} \qquad \mathbf{K}_1 = \mathcal{D}\left[\mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}}\right], \tag{6.13}$$

where the discrete line-integrals are given by $\boldsymbol{\ell} = \mathbf{A}\boldsymbol{\mu}$, where $\mathbf{A}$ again models the projection operation according to Eq. (4.22). At position (2), the free-space propagation of the wavefield is included, which results in interference effects due to the phase-shifts induced by the sample. According to Eq. (6.8), this can be modeled by the linear operator

$$\mathbf{P} = \mathbf{I} - z\frac{\delta}{\mu}\mathbf{L}, \tag{6.14}$$

where $\mathbf{I}$ denotes the identity matrix and $\mathbf{L}$ again approximates the continuous Laplacian operator. The propagation of the wavefield does also not correlate noise, thus the resulting random vector is still Poisson distributed and its mean vector and covariance matrix are defined according to

$$\bar{\mathbf{y}}_2 = \mathbf{P}\mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}} \qquad \text{and} \qquad \mathbf{K}_2 = \mathcal{D}\left[\mathbf{P}\mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}}\right]. \tag{6.15}$$

Again, at position (3), the influence of the source is included. As shown in Eq. (6.10), the influence of the reduced spatial coherence can be described by a convolution. As the convolution is a linear operation, a Toeplitz matrix denoted by $\mathbf{B}_s$ can be used to describe this effect. Due to the fact that the source blur does not correlate noise, the respective mean vector and covariance matrix are given by

$$\bar{\mathbf{y}}_3 = \mathbf{B}_s\mathbf{P}\mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}} \qquad \text{and} \qquad \mathbf{K}_3 = \mathcal{D}\left[\mathbf{B}_s\mathbf{P}\mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}}\right]. \tag{6.16}$$

Finally, the X-rays are converted to visible light in the scintillator, which results in an additional blurring of the signal. However, as visible light photons created from a single X-ray photon are fully correlated, the subsequent signal is not Poisson distributed anymore. The respective mean vector and covariance matrix are consequently given by

$$\boxed{\bar{\mathbf{y}}_4 = \mathbf{B}_d\mathbf{B}_s\mathbf{P}\mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}}} \qquad \text{and} \qquad \boxed{\mathbf{K}_4 = \mathbf{B}_d\mathcal{D}\left[\mathbf{B}_s\mathbf{P}\mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}}\right]\mathbf{B}_d^T.} \tag{6.17}$$

Although not the focus of this work, one can include additional readout noise $\mathbf{K}_{\text{ro}}$ in the above framework as additive term to $\mathbf{K}_4$ [Tilley et al., 2016a]. It is further assumed that influence of the stochastic nature of the conversion from an X-ray photon to light is negligible. This effect has been analyzed in [Elbakri and Fessler, 2003]. Moreover, the effect of the rectangular shape of the pixels is assumed to be small compared to the scintillator blur and aliasing is not a dominant effect. [Tilley et al., 2016a]

6.3. Model-based iterative phase retrieval

In this section[1], the framework derived in the previous section is used for phase retrieval. The results are compared to the single-material phase-retrieval algorithm [Paganin et al., 2002]

[1] Significant parts of this section have been developed together with Wolfgang Noichl.

and its extension, which includes the effect of partial coherence [Beltran et al., 2018]. Subsequently, the connection to MAP approaches in the context of regularized denoising approaches [Rudin et al., 1992] is discussed.

Here, phase retrieval refers to the process of recovering the intensity in the object plane, denoted by $\mathbf{z}$ in the following, from the intensity in the detector plane, denoted by $\mathbf{y}$. Knowing the intensity in the object plane is sufficient to recover the phase information due to the homogeneity assumption. The physical mean model and the covariance matrix, which are examined in the following, are given by Eq. (6.17) according to

$$\bar{\mathbf{y}}(\mathbf{z}) = \mathbf{B}_d\mathbf{B}_s\mathbf{P}\mathbf{z} \qquad \text{and} \qquad \mathbf{K}(\mathbf{z}) = \mathbf{B}_d\mathcal{D}[\mathbf{B}_s\mathbf{P}\mathbf{z}]\mathbf{B}_d^T. \tag{6.18}$$

Using this formulation, the connection between the intensity in the object plane and the intensity in the detector plane is linear. The line-integrals and the phase shift are related to the intensity in the object plane by $\boldsymbol{\ell} = \mathbf{A}\boldsymbol{\mu} = -\log(\mathbf{z}/I_0)$ and $\boldsymbol{\phi} = k\delta/\mu\log(\mathbf{z}/I_0)$ respectively, assuming $\mathbf{G} = I_0$ to be a scalar to align notation with the analytical algorithms.

Finally, the objective function consisting of a log-likelihood term based on the multivariate normal distribution $\mathcal{N}(\bar{\mathbf{y}}, \mathbf{K})$ and the regularization term $\mathcal{R}$ is given by

$$\mathcal{C}(\mathbf{z}) = \frac{1}{2}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{z})\right)^T \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{z})\right) + \beta\mathcal{R}(\mathbf{z}), \tag{6.19}$$

where $\mathbf{W} \approx \mathbf{K}^{-1}(\mathbf{z})$ approximates the (non-diagonal) inverse covariance matrix independently of $\mathbf{z}$ and β denotes the regularization strength.

6.3.1. Relation to the single-material phase-retrieval algorithm

To examine how the above framework relates to the single-material phase-retrieval algorithm, an ideal source and detector are assumed according to $\mathbf{B}_s = \mathbf{B}_d = \mathbf{I}$. Thus, the physical mean model reduces to

$$\bar{\mathbf{y}}(\mathbf{z}) = \mathbf{P}\mathbf{z}. \tag{6.20}$$

In addition, no prior knowledge in the form of regularization techniques is incorporated in the objective function. Thus, the objective function given by Eq. (6.19) reduces to the following log-likelihood function

$$\mathcal{L}(\mathbf{z}) = \frac{1}{2}\left(\mathbf{y} - \mathbf{P}\mathbf{z}\right)^T \mathbf{W}\left(\mathbf{y} - \mathbf{P}\mathbf{z}\right). \tag{6.21}$$

This log-likelihood function coincides with a weighted least squares objective. To obtain an estimate for the intensity in the object plane, this function is minimized. The gradient of $\mathcal{L}$ with respect to $\mathbf{z}$ reads

$$\nabla_{\mathbf{z}}\mathcal{L} = -\mathbf{P}^T\mathbf{W}\left(\mathbf{y} - \mathbf{P}\mathbf{z}\right). \tag{6.22}$$

By setting the gradient to zero and solving for $\mathbf{z}$ explicitly, the estimate for the intensity in the object plane is obtained according to

$$\mathbf{z} = \left(\mathbf{P}^T\mathbf{W}\mathbf{P}\right)^{-1}\mathbf{P}^T\mathbf{W}\mathbf{y} = \mathbf{P}^{-1}\mathbf{y}. \tag{6.23}$$

As the term in the parenthesis can be inverted, the solution is independent of the statistical weights. This leads to the trivial solution of inverting the propagator to obtain the intensity in the object plane given the intensity in the detector plane and accounting for the incident flux accordingly.

This inversion can be computed explicitly in Fourier space. To match the notation of the single-material phase-retrieval algorithm, the measurements $\mathbf{y}$ in their continuous representation are denoted by $I(\mathbf{r}_\perp, z)$. In addition, the intensity in the object plane denoted by $\mathbf{z}$ above is given in its continuous representation by $I(\mathbf{r}_\perp, 0)$. The propagator $\mathbf{P}$ is defined by Eq. (6.14) and its continuous representation is given by Eq. (6.8), which includes the continuous Laplacian. Thus, the Fourier derivative theorem given by Eq. (2.31) can be used again. The propagator $\mathbf{P}$ in its continuous representation acting on some function f can be expressed in Fourier space according to

$$\mathcal{F}_\perp \left[\left(1 - z\frac{\delta}{\mu}\nabla_\perp^2 \right) f(\mathbf{r}_\perp) \right] = \left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2 \right) \mathcal{F}_\perp \left[f(\mathbf{r}_\perp) \right]. \tag{6.24}$$

This is valid if the operator $\mathbf{L}$ within the propagator $\mathbf{P}$ models the continuous Laplacian. If $\mathbf{L}$ is implemented using the finite difference approximation, for instance using a five-point stencil, according to

$$[\mathbf{L}\mathbf{x}]_{i,j} = x_{i+1,j} + x_{i,j+1} - 4x_i + x_{i-1,j} + x_{i,j-1}, \tag{6.25}$$

where the different spatial dimensions are denoted explicitly by two indices, the Fourier derivative theorem has to be modified. Denoting the two-dimensional DFT by $\boldsymbol{\mathcal{F}}_\perp$ and $\mathbf{x}$ being some vector, the analogous Fourier derivative theorem reads

$$\boldsymbol{\mathcal{F}}_\perp \left[\mathbf{L}\mathbf{x} \right] = -\boldsymbol{\kappa}_\perp^2 \boldsymbol{\mathcal{F}}_\perp \left[\mathbf{x} \right], \tag{6.26}$$

where the analogous discrete frequencies $\boldsymbol{\kappa}_\perp^2$ are given by

$$\boldsymbol{\kappa}_\perp^2 = 4 \left(\sin^2 \left(\frac{\mathbf{k}_x}{2} \right) + \sin^2 \left(\frac{\mathbf{k}_y}{2} \right) \right). \tag{6.27}$$

The trigonometric functions operate element-wise. By Taylor expanding $\sin^2(x) = x^2 + \mathcal{O}(x^4)$, the discrete frequencies $\boldsymbol{\kappa}_\perp^2$ coincide with the analytic frequencies $\mathbf{k}_\perp^2$ according to

$$\boldsymbol{\kappa}_\perp^2 \approx 4 \left(\left(\frac{\mathbf{k}_x}{2} \right)^2 + \left(\frac{\mathbf{k}_y}{2} \right)^2 \right) = \mathbf{k}_x^2 + \mathbf{k}_y^2 = \mathbf{k}_\perp^2. \tag{6.28}$$

This subtle difference will occur for the derivation of the generalized single-material phase-retrieval algorithm as well as the analytic solution for the denoising algorithm, which will be discussed in this section.

Having defined the analytic expression of the propagator $\mathbf{P}$, Eq. (6.23) has the following representation in Fourier space

$$\mathcal{F}_\perp \left[I(\mathbf{r}_\perp, 0) \right] = \left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2 \right)^{-1} \mathcal{F}_\perp \left[I(\mathbf{r}_\perp, z) \right] \tag{6.29}$$

and can be solved for the intensity in the object plane according to

$$I(\mathbf{r}_\perp, 0) = \mathcal{F}_\perp^{-1} \left[\frac{\mathcal{F}_\perp \left[I(\mathbf{r}_\perp, z) \right]}{1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2} \right]. \tag{6.30}$$

This result coincides with the single-material phase-retrieval algorithm as stated in Eq. (3.18) using $T(\mathbf{r}_\perp) = -1/\mu \log(I(\mathbf{r}_\perp, 0)/I_0)$ to recover the traces explicitly. From Eq. (6.23) it is obvious that the above formulation recovers the single-material phase-retrieval algorithm as the models for the interference effects coincide, at least in its analytical representation.

6.3.2. Generalization of the single-material phase-retrieval algorithm

By extending the physical mean model to additionally account for the spatial coherence of the source as well as the detector response as given by Eq. (6.18) and including a regularization term, the objective function based on Eq. (6.19) reads

$$\mathcal{C}(\mathbf{z}) = \frac{1}{2}\left(\mathbf{y} - \mathbf{B}_d\mathbf{B}_s\mathbf{P}\mathbf{z}\right)^T \mathbf{W}\left(\mathbf{y} - \mathbf{B}_d\mathbf{B}_s\mathbf{P}\mathbf{z}\right) + \beta\mathcal{R}(\mathbf{z}). \tag{6.31}$$

First, the regularization term $\mathcal{R}$ is specified. One of the most successful regularization techniques in image denoising of conventional images as well as in CT is TV regularization, which uses the ℓ_1 norm as penalty function as defined by Eq. (4.55). TV regularization evolves around the assumption of having piece-wise constant structures and thus favors sparse solutions (in terms of neighboring pixel/voxel values), which, however, is not a good assumption in the projection domain. For instance, the projection of a homogeneous cylinder is by no means piece-wise constant. However, in the volume domain, the three-dimensional distribution consists only of cylinder or no cylinder, which is a piece-wise constant structure. In the projection domain, quadratic regularization can be employed instead, which uses the ℓ_2 norm as penalty function given by Eq. (4.53). Quadratic regularization only penalizes the magnitude without enforcing sparsity. TV regularization in the context of sparsity and the comparison to quadratic regularization will be discussed in greater detail in Chapter 7. The regularizer $\mathcal{R}$ given by Eq. (4.51) with the quadratic penalty given by Eq. (4.53) can be written as

$$\mathcal{R}(\mathbf{z}) = \frac{1}{2}\sum_{i,j}\sum_{m,n\in\mathcal{N}_{i,j}} \frac{1}{\Delta_{im,jn}}\left(\frac{z_{m,n} - z_{i,j}}{\Delta_{im,jn}}\right)^2, \tag{6.32}$$

where the indices explicitly refer to the different spatial components. If the neighborhood only accounts for the next four neighbors, the additional weights $\Delta_{im,jn}$ can be omitted as they are identical[2]. Given appropriate border handling, it is furthermore sufficient to only consider the next two neighbors due to the symmetry of the quadratic penalty. Thus, with $\mathcal{N}_{i,j} = \{(i+1, j), (i, j+1)\}$, Eq. (6.32) reduces to

$$\mathcal{R}(\mathbf{z}) = \frac{1}{2}\sum_{i,j}\left((z_{i+1,j} - z_{i,j})^2 + (z_{i,j+1} - z_{i,j})^2\right). \tag{6.33}$$

This equation can be reformulated as a quadratic equation using matrix-vector notation according to

$$\mathcal{R}(\mathbf{z}) = \frac{1}{2}\left(\mathbf{z}^T\mathbf{D}_x^T\mathbf{D}_x\mathbf{z} + \mathbf{z}^T\mathbf{D}_y^T\mathbf{D}_y\mathbf{z}\right) \tag{6.34}$$

$$= \frac{1}{2}\mathbf{z}^T\left(\mathbf{D}_x^T\mathbf{D}_x + \mathbf{D}_y^T\mathbf{D}_y\right)\mathbf{z} \tag{6.35}$$

$$= \frac{1}{2}\mathbf{z}^T\mathbf{R}\mathbf{z}, \tag{6.36}$$

defining $\mathbf{R} = \mathbf{D}_x^T\mathbf{D}_x + \mathbf{D}_y^T\mathbf{D}_y$, where $\mathbf{D}_x$ and $\mathbf{D}_y$ denote the forward differences in the two spatial dimensions respectively. Formulating the quadratic penalty according to Eq. (6.36) has for instance been done in [Gang et al., 2014].

[2] In this case, square pixels are assumed.

The resulting objective function is then given by

$$\mathcal{C}(\mathbf{z}) = \frac{1}{2}\left(\mathbf{y} - \mathbf{B}_d\mathbf{B}_s\mathbf{P}\mathbf{z}\right)^T \mathbf{W}\left(\mathbf{y} - \mathbf{B}_d\mathbf{B}_s\mathbf{P}\mathbf{z}\right) + \frac{\beta}{2}\mathbf{z}^T\mathbf{R}\mathbf{z} \tag{6.37}$$

and solving for $\mathbf{z}$ by setting the gradient to zero results in

$$\mathbf{z} = \left(\mathbf{P}^T\mathbf{B}_s^T\mathbf{B}_d^T\mathbf{W}\mathbf{B}_d\mathbf{B}_s\mathbf{P} + \beta\mathbf{R}\right)^{-1}\mathbf{P}^T\mathbf{B}_s^T\mathbf{B}_d^T\mathbf{W}\mathbf{y}. \tag{6.38}$$

Due to the size of the respective operators, this direct inversion cannot be computed for realistically sized images. Nevertheless, gradient-based approaches for minimizing the above objective function to obtain this solution would be feasible, but are not addressed here.

However, given certain approximations, this equation can be solved efficiently in Fourier space. Therefore, the statistical weights will be omitted according to $\mathbf{W} = \mathbf{I}$. Thus, Eq. (6.38) reduces to

$$\mathbf{z} = \left(\mathbf{P}^T\mathbf{B}_s^T\mathbf{B}_d^T\mathbf{B}_d\mathbf{B}_s\mathbf{P} + \beta\mathbf{R}\right)^{-1}\mathbf{P}^T\mathbf{B}_s^T\mathbf{B}_d^T\mathbf{y}. \tag{6.39}$$

In the following, the respective operations will be expressed analytically and their representations in Fourier space will be discussed. Again, the intensity in the object plane and the detector plane will be denoted by $I(\mathbf{r}_\perp, 0)$ and $I(\mathbf{r}_\perp, z)$ respectively. The Fourier space representation of the propagator $\mathbf{P}$ has already been introduced in Eq. (6.29). For the blur operators $\mathbf{B}_s$ and $\mathbf{B}_d$ it is assumed that they can be approximated by a Gaussian blur. In accordance with Eq. (6.10) and Eq. (6.11), the blur operators can be expressed as a convolution. Employing the convolution theorem, a convolution in real space can be equivalently expressed as an element-wise multiplication in Fourier space. Thus, the real-space and Fourier-space representations of these operations acting on some function f can be written as

$$\mathcal{F}_\perp\left[\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\mathbf{r}'_\perp}{2\sigma_s^2}} f(\mathbf{r}_\perp - \mathbf{r}'_\perp)\mathrm{d}\mathbf{r}'_\perp\right] = e^{-2\sigma_s^2\mathbf{k}_\perp^2}\mathcal{F}_\perp\left[f(\mathbf{r}_\perp)\right], \tag{6.40}$$

$$\mathcal{F}_\perp\left[\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\mathbf{r}'_\perp}{2\sigma_d^2}} f(\mathbf{r}_\perp - \mathbf{r}'_\perp)\mathrm{d}\mathbf{r}'_\perp\right] = e^{-2\sigma_d^2\mathbf{k}_\perp^2}\mathcal{F}_\perp\left[f(\mathbf{r}_\perp)\right]. \tag{6.41}$$

Finally, the discrete regularization term has to be converted into an analytic expression. The explicit form of the matrix $\mathbf{R}$ is illustrated on a small one-dimensional problem with only four entries and information at the borders is wrapped around. This results in

$$\mathbf{D}_x^T\mathbf{D}_x = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \\ 1 & 0 & 0 & -1 \end{pmatrix}^T \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \\ 1 & 0 & 0 & -1 \end{pmatrix} = \begin{pmatrix} 2 & 1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{pmatrix} = -\mathbf{L}_x, \tag{6.42}$$

which is the negative one-dimensional Laplacian operator. Consequently, the regularization matrix can be written as

$$\mathbf{R} = \mathbf{D}_x^T\mathbf{D}_x + \mathbf{D}_y^T\mathbf{D}_y = -\mathbf{L}_x - \mathbf{L}_y = -\mathbf{L}, \tag{6.43}$$

and is equal to the negative Laplacian operator $\mathbf{L}$, which is also present in the propagator $\mathbf{P}$. The corresponding analytic Fourier space representation is then given by

$$\mathcal{F}_\perp\left[-\nabla_\perp^2 f(\mathbf{r}_\perp)\right] = \mathbf{k}_\perp^2 \mathcal{F}_\perp\left[f(\mathbf{r}_\perp)\right], \tag{6.44}$$

using again the Fourier derivative theorem. Here, the approximation $\boldsymbol{\kappa}_\perp^2 \approx \mathbf{k}_\perp^2$ is used, which has been discussed previously.

Having defined all operators in Fourier space, Eq. (6.39) can be expressed analytically in Fourier space according to

$$\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, 0)\right] = \Bigg(\left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) e^{-2\sigma_s^2\mathbf{k}_\perp^2} e^{-2\sigma_d^2\mathbf{k}_\perp^2} e^{-2\sigma_d^2\mathbf{k}_\perp^2} e^{-2\sigma_s^2\mathbf{k}_\perp^2} \left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) + \tag{6.45}$$

$$+ \beta\mathbf{k}_\perp^2 \Bigg)^{-1} \left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) e^{-2\sigma_s^2\mathbf{k}_\perp^2} e^{-2\sigma_d^2\mathbf{k}_\perp^2} \mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)\right]. \tag{6.46}$$

Rearranging the terms and solving for the intensity in the object plane results in

$$I(\mathbf{r}_\perp, 0) = \mathcal{F}_\perp^{-1}\left[\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)\right]}{\left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) e^{-2\left(\sigma_s^2+\sigma_d^2\right)\mathbf{k}_\perp^2} + \frac{\beta\mathbf{k}_\perp^2}{\left(1+z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) e^{-2\left(\sigma_s^2+\sigma_d^2\right)\mathbf{k}_\perp^2}}}\right]. \tag{6.47}$$

Expressing the intensity in the object plane in terms of the trace according to $I(\mathbf{r}_\perp, 0) = I_0 \exp[-\mu T(\mathbf{r}_\perp)]$ and solving for the trace explicitly, reads

$$\boxed{T(\mathbf{r}_\perp) = -\frac{1}{\mu}\log\left[\mathcal{F}_\perp^{-1}\left[\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)\right]/I_0}{\left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) e^{-2\left(\sigma_s^2+\sigma_d^2\right)\mathbf{k}_\perp^2} + \frac{\beta\mathbf{k}_\perp^2}{\left(1+z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) e^{-2\left(\sigma_s^2+\sigma_d^2\right)\mathbf{k}_\perp^2}}}\right]\right].} \tag{6.48}$$

Similarly to the single-material phase-retrieval algorithm, the above expressions can be computed efficiently and have little memory requirements. For $\sigma_s = \sigma_d = 0$ and $\beta = 0$, Eq. (6.48) reduces to the single-material phase-retrieval algorithm as stated in Eq. (3.18).

In order to interpret the different components of Eq. (6.48), it is useful to expand the denominator in terms of $\mathbf{k}_\perp^2$. Employing the Taylor expansion of the Gaussian blur in Fourier space has been proposed in a similar context as discussed below in [Beltran et al., 2018]. The Taylor expansion reads

$$\left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) e^{-2\left(\sigma_s^2+\sigma_d^2\right)\mathbf{k}_\perp^2} + \frac{\beta\mathbf{k}_\perp^2}{\left(1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right) e^{-2\left(\sigma_s^2+\sigma_d^2\right)\mathbf{k}_\perp^2}} \tag{6.49}$$

$$= 1 + \left(z\frac{\delta}{\mu} - 2\left(\sigma_s^2 + \sigma_d^2\right) + \beta\right)\mathbf{k}_\perp^2 + \mathcal{O}(\mathbf{k}_\perp^4). \tag{6.50}$$

Neglecting any terms $\mathcal{O}(\mathbf{k}_\perp^4)$ also eliminates the discrepancies between $\mathbf{k}_\perp^2$ and $\boldsymbol{\kappa}_\perp^2$. Eq. (6.48) then becomes

$$\boxed{T(\mathbf{r}_\perp) = -\frac{1}{\mu}\log\left[\mathcal{F}_\perp^{-1}\left[\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)\right]/I_0}{\left(z\frac{\delta}{\mu} - 2\left(\sigma_s^2 + \sigma_d^2\right) + \beta\right)\mathbf{k}_\perp^2 + 1}\right]\right],} \tag{6.51}$$

which is referred to as the *generalized single-material phase-retrieval algorithm* in the following. This filter coincides with the single-material phase-retrieval algorithm given by Eq. (3.18) substituting

$$z\frac{\delta}{\mu} \to z\frac{\delta}{\mu} - 2\left(\sigma_s^2 + \sigma_d^2\right) + \beta, \tag{6.52}$$

which is discussed below.

The first term $z\delta/\mu$ accounts for the interference effects due to the phase-shifting and attenuating properties of the material. In an ideal system this term is sufficient to retrieve the intensity in the object plane as derived in [Paganin et al., 2002]. The second term $-2\left(\sigma_s^2 + \sigma_d^2\right)$ accounts for the reduced spatial coherence and blur in the detection system. Both the effects of the source and the detector have the same influence on the phase retrieval, although their properties with respect to noise vary, as addressed in more detail in the next section. The negative sign indicates that the blur of the system already accounts for a certain part of the phase retrieval. This counteracting effect has already been utilized experimentally in [Gureyev et al., 2004]. However, the equivalence of phase retrieval and system blur is only valid due to the approximations of Eq. (6.50). Recently in [Beltran et al., 2018], the single-material phase-retrieval algorithm has been extended using the formalism of scalar wave theory to similarly include the effect of partial coherence. Their result can be obtained by setting $\sigma_d = \beta = 0$, although they argue that the effect of the detector blur can also be incorporated. Finally, the third term accounts for the denoising properties due to quadratic regularization. In contrast to the system blur, incorporating stronger regularization is equivalent to stronger phase retrieval. Consequently, this results in lower noise, but also reduces spatial resolution. In practice, the value of $z\delta/\mu$ as used in the conventional single-material algorithm is chosen by hand. Thereby, one tries to balance the suppression of the interference effects to give best resolution as well as low noise level. As shown above, the manual search for this factor includes various physical effects beyond the classical phase retrieval due to the interference effects, including the effects of partial coherence and the blurring of the detection system as well as noise suppression. Equivalently, one can also think about these effects as a reduction of the effective propagation distance according to

$$z \to z - 2\frac{\mu}{\delta}\left(\sigma_s^2 + \sigma_d^2\right) + \frac{\mu}{\delta}\beta. \tag{6.53}$$

The effective propagation distance has been discussed in the context of the Fresnel scaling theorem given by Eq. (2.55) to account for cone-beam effects. In summary, the effective propagation distance is reduced by the blur of the system due to the source [Beltran et al., 2018] or the detector and increased if more noise suppression is desired.

Recently in [Paganin and Morgan, 2019], the reduction of the term $z\delta/\mu$ has been analyzed in the context of the Fokker-Planck equation. Thereby, this factor is reduced by a so-called effective diffusion coefficient $D_{\text{eff}}(\mathbf{r}_\perp)$, which can account for small-angle X-ray scattering (referred to as dark-field signal), edge-scattering signal or incoherent aberrations such as the influence of the source or the detector.

The proposed framework using the objective function given by Eq. (6.37) allows to include various effects that influence the phase retrieval. Among others, more complicated blur models can be used, which do not necessarily have to be expressible as a convolution. In addition, the statistical properties in the image formation can be utilized including noise correlations, which will be discussed in section 6.4 in the context of PB-CT. Furthermore, also non-quadratic regularization techniques such as TV regularization can be employed. Finally, the effects at the borders of the image can be accounted for differently.

6.3.3. Relation to image denoising algorithms

The generalized single-material phase-retrieval algorithm already hinted at a similarity between phase retrieval described by $z\delta/\mu$ and denoising using quadratic regularization parameterized by β. The connection of these seemingly unrelated operations is outlined in the following. Again an ideal source and detector are assumed according to $\sigma_s = \sigma_d = 0$. Consequently, the analytic solution given by Eq. (6.47) of the objective function given by Eq. (6.37) reduces to

$$I(\mathbf{r}_\perp, 0) = \mathcal{F}_\perp^{-1}\left[\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)\right]}{1 + z\frac{\delta}{\mu}\mathbf{k}_\perp^2 + \frac{\beta\mathbf{k}_\perp^2}{\left(1+z\frac{\delta}{\mu}\mathbf{k}_\perp^2\right)}}\right]. \tag{6.54}$$

Assuming no interference effects due to free-space propagation according to $z\delta/\mu = 0$, the above equation reads

$$I(\mathbf{r}_\perp, 0) = \mathcal{F}_\perp^{-1}\left[\frac{\mathcal{F}_\perp\left[I(\mathbf{r}_\perp, z)\right]}{\beta\mathbf{k}_\perp^2 + 1}\right], \tag{6.55}$$

which has the same form as the single-material phase-retrieval algorithm for $\beta = z\delta/\mu$ using $T(\mathbf{r}_\perp) = -1/\mu \log(I(\mathbf{r}_\perp, 0)/I_0)$. One the other hand, the corresponding objective function, which has the above solution, is given by

$$\mathcal{C}(\mathbf{z}) = \frac{1}{2}(\mathbf{y} - \mathbf{z})^T(\mathbf{y} - \mathbf{z}) + \frac{\beta}{2}\mathbf{z}^T\mathbf{R}\mathbf{z}, \tag{6.56}$$

where the mean model $\bar{\mathbf{y}}(\mathbf{z}) = \mathbf{z}$ does not account for free-space propagation. Recovering the regularizer $\mathcal{R}(\mathbf{z}) = 1/2\,\mathbf{z}^T\mathbf{R}\mathbf{z}$, the objective function is given by

$$\mathcal{C}(\mathbf{z}) = \frac{1}{2}(\mathbf{y} - \mathbf{z})^T(\mathbf{y} - \mathbf{z}) + \beta\mathcal{R}(\mathbf{z}). \tag{6.57}$$

This coincides with the widely used TV denoising approach if the penalty function of the regularization term is given by the ℓ_1 norm according to Eq. (4.55). Thereby, $\mathbf{y}$ denotes the acquired noisy image and $\mathbf{z}$ is the denoised image. [Rudin et al., 1992]

In conclusion, the single-material phase-retrieval algorithm can be thought of as a quadratic denoising algorithm, which omits an explicit physical description of the interference effects. The regularization strength is chosen by hand to give the best depiction of the image. This is similar to most applications using the single-material phase-retrieval algorithm, where the factor $z\delta/\mu$ is commonly also chosen empirically.

6.4. Objective function for homogeneous PB-CT

After discussing how the framework derived in Section 6.2 can be utilized for phase retrieval in PBI, this section introduces the main result of this chapter, namely a MBIR algorithm for PB-CT suitable for laboratory sources.

The physical mean model and covariance matrix, according to Eq. (6.17), are given by

$$\bar{\mathbf{y}}(\boldsymbol{\mu}) = \mathbf{B}e^{-\mathbf{A}\boldsymbol{\mu}} \qquad \text{and} \qquad \mathbf{K}(\boldsymbol{\mu}) = \mathbf{B}_d\mathcal{D}[\mathbf{B}_s\mathbf{P}\mathbf{G}e^{-\mathbf{A}\boldsymbol{\mu}}]\mathbf{B}_d^T, \tag{6.58}$$

where $\mathbf{B} = \mathbf{B}_d\mathbf{B}_s\mathbf{P}\mathbf{G}$ is defined to reduce notation overhead and to make the mean model consistent with the mean model recently proposed for conventional attenuation-based CT in [Tilley et al., 2018a].

The objective function consists again of a log-likelihood term and a regularization term. In the previous chapter, the log-likelihood terms assumed uncorrelated measurements, as the underlying noise models were assumed to follow a Poisson distribution. Now, the measurements are modeled by a multivariate normal distribution according to $\mathcal{N}(\bar{\mathbf{y}}(\boldsymbol{\mu}), \mathbf{K}(\boldsymbol{\mu}))$. Thus, with some regularization term $\mathcal{R}(\boldsymbol{\mu})$, the objective reads

$$\mathcal{C}(\boldsymbol{\mu}) = \frac{1}{2}(\mathbf{y} - \mathbf{B}\mathrm{e}^{-\mathbf{A}\boldsymbol{\mu}})^T\mathbf{W}(\mathbf{y} - \mathbf{B}\mathrm{e}^{-\mathbf{A}\boldsymbol{\mu}}) + \beta\mathcal{R}(\boldsymbol{\mu}), \tag{6.59}$$

where here $\mathbf{W} \approx \mathbf{K}^{-1}(\boldsymbol{\mu})$ approximates the inverse non-diagonal covariance matrix. The gradient of the objective with respect to $\boldsymbol{\mu}$ is given by

$$\begin{aligned}\nabla\mathcal{C} &= \mathbf{A}^T\mathbf{D}\left[e^{-\mathbf{A}\boldsymbol{\mu}}\right]\mathbf{B}^T\mathbf{W}(\mathbf{y} - \mathbf{B}\mathrm{e}^{-\mathbf{A}\boldsymbol{\mu}}) + \beta\nabla\mathcal{R} &(6.60)\\ &= \mathbf{A}^T\mathbf{D}\left[e^{-\mathbf{A}\boldsymbol{\mu}}\right]\left(\mathbf{B}^T\mathbf{W}\mathbf{y} - \mathbf{B}^T\mathbf{W}\mathbf{B}\mathrm{e}^{-\mathbf{A}\boldsymbol{\mu}}\right) + \beta\nabla\mathcal{R}. &(6.61)\end{aligned}$$

The main problem of the non-diagonal covariance matrix is that it cannot be inverted efficiently. Thus, in the following, various approximations of $\mathbf{W}$ are discussed.

Using the exact inversion of the covariance matrix according to $\mathbf{W} = \mathbf{K}^{-1}(\boldsymbol{\mu})$ is not feasible. On the one hand, assuming 1000^3 pixels and 1000^3 voxels, this matrix holds 10^{18} elements, similarly to the projection matrix $\mathbf{A}$, as discussed in Subsection 4.2.2. On the other hand, computing the respective matrix-vector product using iterative methods would be computationally very expensive as this has to be performed in every iteration. However, similar nested iterations have been demonstrated in [Tilley et al., 2016a] in the context of attenuation-based CT using a linearized mean model.

The simplest choice would be to approximate the covariance matrix by

$$\mathbf{K}(\boldsymbol{\mu}) \approx \mathcal{D}\left[\mathbf{y}\right] \tag{6.62}$$

by assuming uncorrelated noise. Thus, the approximation $\mathbf{W} = \mathcal{D}^{-1}[\mathbf{y}]$ would make the weights independent of $\boldsymbol{\mu}$. However, in this case the noise correlations imposed by the blur of the detection system are not accounted for.

In [Tilley et al., 2016b], a different approach has been proposed, which includes correlated noise into the reconstruction. Thereby, the covariance matrix is approximated according to

$$\mathbf{K}(\boldsymbol{\mu}) \approx \mathbf{B}_d\mathcal{D}\left[\mathbf{y}\right]\mathbf{B}_d^T. \tag{6.63}$$

The reason for this choice can be seen from the computation of the gradient given by Eq. (6.61). The first term $\mathbf{B}^T\mathbf{W}\mathbf{y}$ has to be computed only once. This can for instance be achieved by an iterative algorithm such as the CG algorithm [Shewchuk et al., 1994]. The second term can be computed according to

$$\begin{aligned}\mathbf{B}^T\mathbf{W}\mathbf{B} &= \mathbf{B}^T\left(\mathbf{B}_d\mathcal{D}\left[\mathbf{y}\right]\mathbf{B}_d^T\right)^{-1}\mathbf{B} &(6.64)\\ &= \mathbf{G}^T\mathbf{P}^T\mathbf{B}_s^T\mathbf{B}_d^T\left(\mathbf{B}_d^{-T}\mathcal{D}^{-1}\left[\mathbf{y}\right]\mathbf{B}_d^{-1}\right)\mathbf{B}_d\mathbf{B}_s\mathbf{P}\mathbf{G} &(6.65)\\ &= \mathbf{G}^T\mathbf{P}^T\mathbf{B}_s^T\mathcal{D}^{-1}\left[\mathbf{y}\right]\mathbf{B}_s\mathbf{P}\mathbf{G}, &(6.66)\end{aligned}$$

thereby omitting any explicit inversions of non-diagonal matrices. Thus, the gradient can be computed efficiently in every iteration and the first term can be precomputed at the beginning of the reconstruction.

The proposed physical mean model has been rewritten in Eq. (6.58) such that it matches the mean model examined in [Tilley et al., 2016b] for conventional attenuation-based CT. For their model a separable surrogate approach to minimize the according objective function efficiently has been developed recently [Tilley et al., 2018a]. However, this requires the matrix $\mathbf{B}$ to be non-negative in order to fulfill the optimal curvature conditions derived in [Erdogan and Fessler, 1999a]. In the proposed model, the operator modeling the interference effects does not fulfill this condition and thus, this optimization approach cannot be leveraged.

6.5. Simulation study tailored to laboratory environments

The MBIR algorithms developed in the previous section were evaluated using a simulation study and are compared to several analytical approaches. The parameters of the Munich Compact Light Source (MuCLS) described below were used as reference for the simulation study.

6.5.1. The Munich Compact Light Source (MuCLS)

The MuCLS consists of an inverse Compton source (manufactured by Lyncean Technologies Inc., USA) and an imaging beamline with two endstations.

The X-rays are generated by inverse Compton scattering of laser light with electrons. The generated X-rays are quasimonochromatic with a tune-able energy between 15 keV and 35 keV. The flux depending on the energy is in the order of 10^{10} ph/s with a divergence of around 4 mrad. The X-ray source spot has a Gaussian profile with a standard deviation of around 50 µm. [Achterhold et al., 2013, Eggl et al., 2016, Gradl et al., 2017]

A setup for PBI and PB-CT is located in the first endstation at approximately 4 m distance from the X-ray source spot. The low divergence of the generated X-rays leads to a FOV of approximately 16 mm in diameter. The combination of a small source size and large source to sample distances provides sufficient spatial coherence for these imaging techniques. [Gradl et al., 2017]

6.5.2. Parameter selection and reconstruction algorithms

The parameters for the simulation study are based on the characteristics of the MuCLS. The energy of the X-rays is set to 25 keV. The extent of the projected source is assumed to be Gaussian with $\sigma_s = 2$ px and the X-ray flux is given according to $I_0 = 10^3$ assuming homogeneous illumination. The scintillator-based detector is simulated with 512×64 pixels and a pixelsize of $p = 20$ µm. For simplicity, additional gain factors are omitted such that every incoming X-ray photon is assumed to be detected. The detector is positioned 1 m behind the sample.

The sample used for simulation is based on the FORBILD head phantom as defined in [FORBILD phantoms, 1999]. The voxel values of the phantom are given by water equivalent density, which were used directly to obtain a phantom obeying the homogeneity assumption. The corresponding scalar constants μ and δ of water were extracted from the x-raylib library [Schoonjans et al., 2011]. Finally, the phantom was shrunk by a factor of 25 to fit the FOV of the detector, resulting in a diameter of around 10 mm. The volume consists of the central 64 slices with 512×512 voxels each. In Figure 6.2 (a), the central slice of the phantom is depicted.

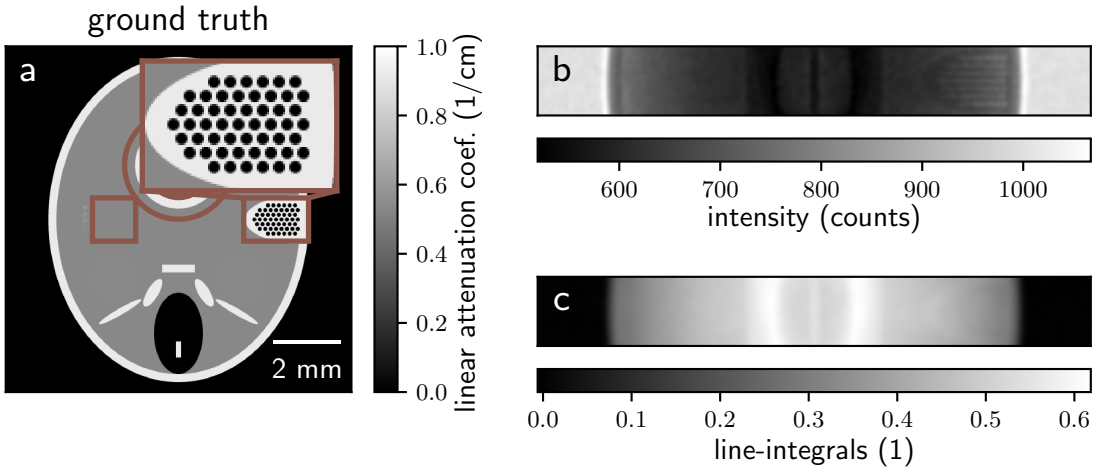

Figure 6.2.: Simulation study of PB-CT. In (a) the central slice of the water phantom is depicted. The rectangular region on the left is used for the computation of the noise level. In addition, the features mimicking the inner-ear are used to evaluate the resolution properties of the reconstruction algorithms. The region between the two partially visible brown circles is used later on for the validation of the edge sharpness. In (b) the simulated measurement for the first view is depicted. The line-integrals recovered by the single-material phase-retrieval algorithm are shown in (c). An adapted version of this figure was submitted for publication [Hehn et al., 2019a].

A half-scan with 512 projections equidistantly distributed between 0° and 180° was simulated in a parallel-beam geometry. The measurement was simulated using

$$\mathbf{y} \sim \mathbf{B}_d \mathcal{P}\left[\mathbf{B}_s \mathbf{P} I_0 e^{-\mathbf{A}\boldsymbol{\mu}}\right], \tag{6.67}$$

where $\mathcal{P}$ draws from a Poisson distribution. The resulting measurement for the first view is depicted in Figure 6.2 (b). In addition to the attenuating properties, edge-enhancement effects in particular at the transition from sample to air can be seen, which are smeared out due to the effects of the source and detector.

In the following, five different reconstruction approaches are investigated, of which the first three methods are analytic and the remaining methods are iterative. If included in the reconstruction, the model for the interference effects of all reconstruction approaches coincide based on Eq. (3.16).

The first reconstruction approach does not take the interference effects into account. According to Eq. (4.45), the line-integrals are given by

$$\boldsymbol{\ell} = -\log\left(\frac{\mathbf{y}}{I_0}\right), \tag{6.68}$$

which are then reconstructed tomographically using the FBP technique defined by Eq. (4.27) and Eq. (4.28). This reconstruction approach is simply referred to as the 'FBP' method and coincides with the conventional reconstruction approach for attenuation-based CT. The second reconstruction approach is the conventional reconstruction approach for PB-CT, as outlined in Figure 5.2. Here, the line-integrals are recovered from each view individually

using the single-material phase-retrieval algorithm as proposed by [Paganin et al., 2002]. The algorithm given by Eq. (5.46) can be expressed in terms of line-integrals according to

$$\boldsymbol{\ell} = -\log\left(\boldsymbol{\mathcal{F}}^{-1}\left[\frac{\boldsymbol{\mathcal{F}}\left[\mathbf{y}\right]/I_0}{z\frac{\delta}{\mu}\mathbf{k}_\perp^2+1}\right]\right), \tag{6.69}$$

using $\boldsymbol{\ell} = \mu\mathbf{t}$. The values for z, δ, μ as well as the frequencies $\mathbf{k}_\perp$ are matched with the simulations. The line-integrals are subsequently reconstructed using the FBP algorithm. This method is therefore referred to as the 'Paganin + FBP' method. The third method uses the extension of the single-material phase-retrieval algorithm that takes the blur of the system into account as proposed by [Beltran et al., 2018]. The algorithm to obtain the line-integrals can be obtained by discretizing Eq. (6.51) and setting $\beta = 0$ resulting in

$$\boldsymbol{\ell} = -\log\left(\boldsymbol{\mathcal{F}}^{-1}\left[\frac{\boldsymbol{\mathcal{F}}\left[\mathbf{y}\right]/I_0}{\left(z\frac{\delta}{\mu}-2\left(\sigma_s^2+\sigma_d^2\right)\right)\mathbf{k}_\perp^2+1}\right]\right) \tag{6.70}$$

using the same parameters for z, δ, μ and $\mathbf{k}_\perp$. In addition, the effects of the source and the detector are accounted for by the parameters σ_s and σ_d, which are assumed to be known. Again, FBP is used for tomographic reconstruction of the line-integrals. This method is referred to as the 'Beltran + FBP' method.

The two remaining reconstruction approaches use the MBIR approaches derived in Section 6.4. The algorithms minimize the objective function given by Eq. (6.59). The regularization term given by Eq. (4.51) uses the quadratic penalty given by Eq. (4.53) according to

$$\mathcal{R}(\boldsymbol{\mu}) = \frac{1}{4}\sum_i \sum_{n\in\mathcal{N}} \frac{1}{\Delta_{in}}\left(\frac{\mu_i-\mu_n}{\Delta_{in}}\right)^2, \tag{6.71}$$

where the sum over n includes the next 26 neighboring voxels and $\Delta_{in} \in \{1, 2^{1/2}, 3^{1/2}\}$ denotes the distances to the respective neighbors. The first algorithm referred to as the 'uncorrelated MBIR' method approximates the covariance matrix directly with the measurements given by Eq. (6.62). The second algorithm approximates the covariance matrix according to Eq. (6.63), which includes noise correlations and is therefore referred to as the 'correlated MBIR' method. Reconstructions for both methods are performed for different values of the regularization strength β. For the 'uncorrelated MBIR' method the regularization strength is evaluated from 10^{-2}cm^2 to 10^{2}cm^2 using 17 steps, which are equidistantly distributed on a logarithmic scale. For the 'correlated MBIR' method, the 17 different values of the regularization strength were logarithmically distributed between $10^{-1.25}\text{cm}^2$ and $10^{2.75}\text{cm}^2$. Optimization is performed using the L-BFGS algorithm and only evaluating the curvature conditions in the line-search routine. In addition, the term $\mathbf{B}^T\mathbf{W}\mathbf{y}$ was precomputed once using the CG algorithm. This term could be used independently of the regularization strength. As convergence criterion, the gradient norm is computed according to

$$\|\nabla\mathcal{C}(\boldsymbol{\mu})\|_2 < \epsilon \max\left(1, \|\boldsymbol{\mu}\|_2\right) \tag{6.72}$$

with $\epsilon = 10^{-3}$ chosen empirically such that by visual inspection the solution does not change noticeably anymore.

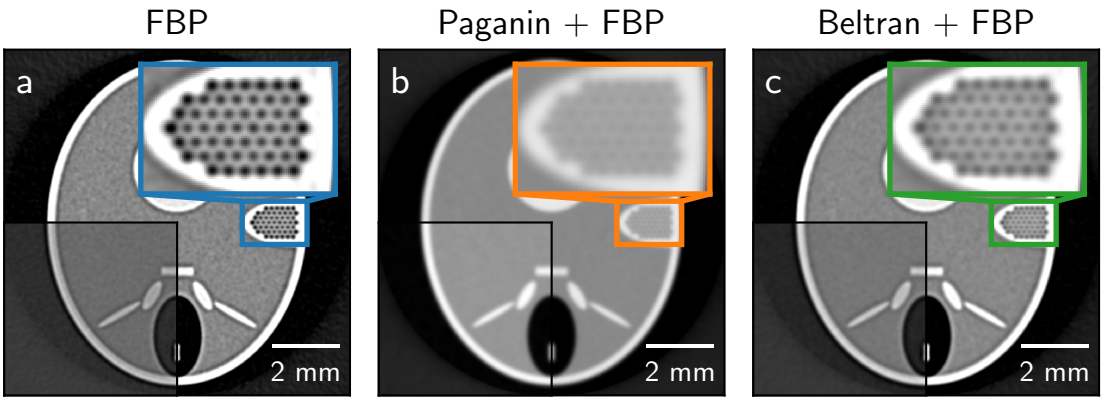

Figure 6.3.: Comparison of the analytic reconstruction methods. The slice shown in Figure 6.2 (a) for the ground truth is depicted with the same grayscale for the different analytic reconstruction methods. Moreover, in the lower left quarter the linear grayscale ranges from the minimal to maximal value in this region to better assess the quality of the sample edges. In (a) the central slice obtained by the 'FBP' method is shown. In (b), the corresponding slice obtained by the 'Paganin + FBP' method is depicted. Finally, in (c) the result of the 'Beltran + FBP' method is shown. An adapted version of this figure was submitted for publication [Hehn et al., 2019a].

6.5.3. Results and discussion

In Figure 6.3 the central slices obtained by the three analytical reconstruction methods are depicted. The linear attenuation coefficients are shown with the same linear grayscale used for the depiction of the ground truth in Figure 6.2 (a). In addition, the lower left quarter shows the reconstructions in a different linear grayscale ranging from the minimal to the maximal value of the attenuation coefficients in this region, to better assess the overshoots at the transition from the sample to the background. The high-resolution features representing the inner ear structure are magnified to evaluate visually the resolution properties of the different methods.

The central slice obtained with the 'FBP' method is shown in Figure 6.3 (a). As this method does not explicitly account for the interference effects, overshoots at the edges are visible, which can be seen in particular in the lower left window. However, due to the smearing of the source and the detector, partial phase retrieval is implicitly performed, resulting in overall good contrast.

In Figure 6.3 (b) the reconstructed slice obtained by the 'Paganin + FBP' method is shown. This method applies the single-material phase-retrieval algorithm for each view independently prior to tomographic reconstruction. Due to the properties of the phase-retrieval algorithm, excellent contrast and low noise levels are obtained. As the interference effects are fully accounted for, no residual edge-enhancement is visible. However, the effects of the source and the detector are not accounted for, resulting in bad resolution properties as detailed in the magnified region.

Finally, Figure 6.3 depicts the respective slice obtained by the 'Beltran + FBP' method. The blur of the source and the detector is taken into account in the phase-retrieval step. Thereby, the strength of the phase retrieval is reduced by $2\sigma^2 = 16p^2$ from $z\delta/\mu \approx 18.1p^2$

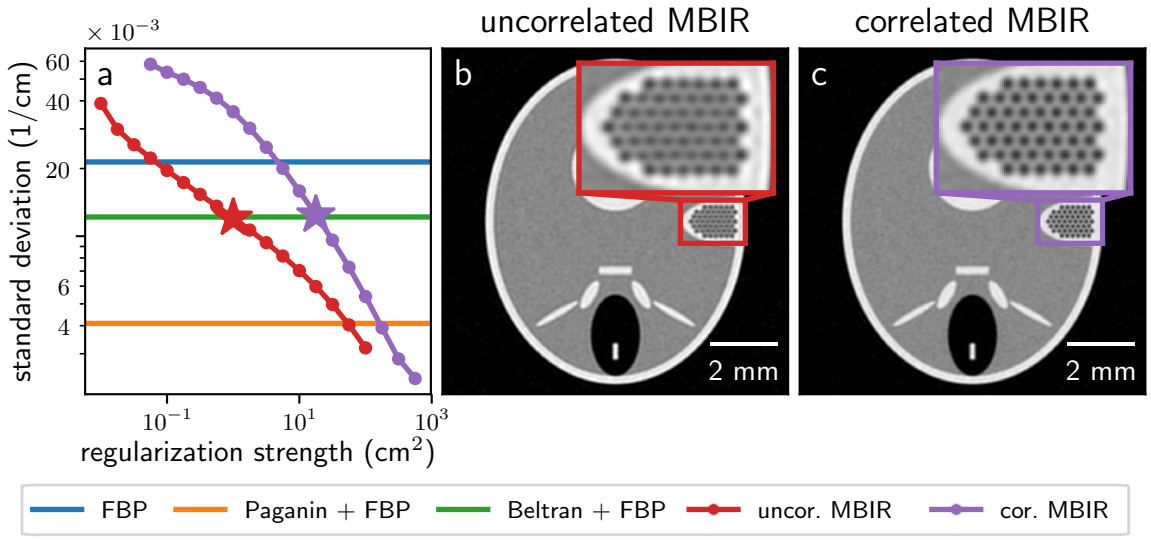

Figure 6.4.: MBIR reconstructions with matched noise levels. In (a) the standard deviation of the reconstructions obtained by the two MBIR methods are plotted over the regularization parameters. In addition, the standard deviation of the results obtained by the analytical reconstructions are included. In (b), a reconstructed slice obtained with the 'uncorrelated MBIR' method is shown which closely matches the noise level of the 'Beltran + FBP' method. In (c), a reconstructed slice obtained with the 'correlated MBIR' method is depicted. Again, the regularization strength of this reconstruction is chosen such that the noise level compares to the noise level of the 'Beltran + FBP' method. An adapted version of this figure was submitted for publication [Hehn et al., 2019a].

resulting in only a weak phase retrieval. However, similar to the 'FBP' method, residual edge-enhancement effects are visible. Thus, a reduction of the phase retrieval strength is not able to accurately model the effects of the source and detector.

Depending on the choice of the regularization strength β, different reconstructions are obtained for the two MBIR methods. For lower values of the regularization strength the resolution is improved, but the noise level increases, while for higher values the noise level is reduced, but the resolution is reduced as well. Therefore, the regularization strengths of the MBIR methods are selected using different criteria.

The first criterion selects the regularization strength of the two MBIR methods such that the variances of the reconstructions match the variance obtained by the 'Beltran + FBP' method best. The variances are calculated from the homogeneous region indicated by the left square shown in Figure 6.2 (a). The region was chosen reasonably big, such that the noise level of all relevant frequencies is incorporated in the evaluation of the variance. The results of the two MBIR methods are depicted in Figure 6.4. In Figure 6.4 (a), the variances of the two methods are plotted over the regularization strength, together with the variances of the analytical methods depicted by the horizontal lines. As already discussed qualitatively, this figure shows quantitatively that the 'FBP' method results in the highest noise level followed by the 'Beltran + FBP' and 'Paganin + FBP' methods. In Figure 6.4 (b), the slice obtained with the 'uncorrelated MBIR' method and in Figure 6.4 (c) the corresponding slice obtained with the 'correlated MBIR' method are depicted. Compared to the analytic methods shown

in Figure 6.3, the MBIR method results in a more accurate reconstruction of the phantom. In particular, compared to the reconstruction obtained with the 'Beltran + FBP' method, which has the same variance, the resolution is improved significantly as seen by the inner ear structures. Moreover, the approximation of the inverse covariance matrix within the two MBIR methods has a crucial impact on the reconstruction quality, which can be seen best from the reconstruction of the inner ear structure. While the reconstruction obtained with the 'uncorrelated MBIR' method has strong overshoots at the edges, the 'correlated MBIR' method reduces these artifacts significantly. This also results in a better overall visualization of the small features using the 'correlated MBIR' method.

The second criterion evaluates the mean squared error (MSE) of the different reconstructions to the ground truth phantom over the whole volume. In Figure 6.5 (a) the MSE is plotted over the regularization strength for the two MBIR methods. In general, if the regularization strength is too low, the MSE increases due to the increased noise level. By contrast, if the regularization strength is too high, the resolution degrades significantly, resulting again in an increase of the MSE as the ground truth cannot be represented accurately. In addition, the figure shows the MSE of the analytical reconstruction approaches by horizontal lines. The 'FBP' method performs worst as it cannot reconstruct the phantom accurately and has high noise levels. The 'Paganin + FBP' and 'Beltran + FBP' methods perform similarly. While the 'Paganin + FBP' method results in a lower noise level, the value of the MSE is partially compensated by the worse resolution compared to the 'Beltran + FBP' method. In Figure 6.5 (b) and Figure 6.5 (c), the reconstructions of the 'uncorrelated MBIR' method and the 'correlated MBIR' method are shown using the regularization strengths that result in the lowest MSE. This results in a good trade-off between noise level and resolution. Compared to the reconstructions at matched variance depicted in Figure 6.4, the resolution is increased and thus the noise level is overall higher. As already seen in Figure 6.4, the reconstruction obtained with the 'uncorrelated MBIR' method has overshoots at the edges which diminish the reconstruction quality considerably. The 'correlated MBIR' method gives the most accurate reconstruction of the phantom without severe overshoots at the edges of the structures.

To better assess the resolution properties of the different reconstruction approaches, line profiles are employed. The line profiles are obtained from the edge of the cylindrical structure located in the upper center of the phantom. This structure is partially visible in Figure 6.2 (a). To achieve subpixel resolution and a low noise level, the symmetry of the structure is employed. Therefore, the values of all voxels are taken into account that lie inside the two brown circles shown in Figure 6.2 (a). In Figure 6.6, the values of the respective voxels are plotted over the distance from the center of the cylindrical structure. As multiple voxel values have the same distance from the center, the corresponding mean values and standard deviations are computed. To improve the visualization, the mean values are interpolated using smoothed spline interpolation by utilizing the standard deviations for appropriate weighting.

In Figure 6.6 (a) the corresponding line profiles of the three analytic reconstruction approaches as depicted in Figure 6.3 are shown together with the MBIR methods at matched variances as depicted in Figure 6.4 and the ground truth line profile. The line profile of the 'FBP' results in the steepest slope at the edge, but also the highest overshoots, followed by the 'Beltran + FBP' method. By contrast, the conventional 'Paganin + FBP' method gives the most shallow slope, but no overshoots. The line profiles of the MBIR methods lie between the 'Paganin + FBP' and the 'Beltran + FPB' methods, both in terms of slope steepness as well as overshoots. The 'uncorrelated MBIR' method results in a steeper slope,

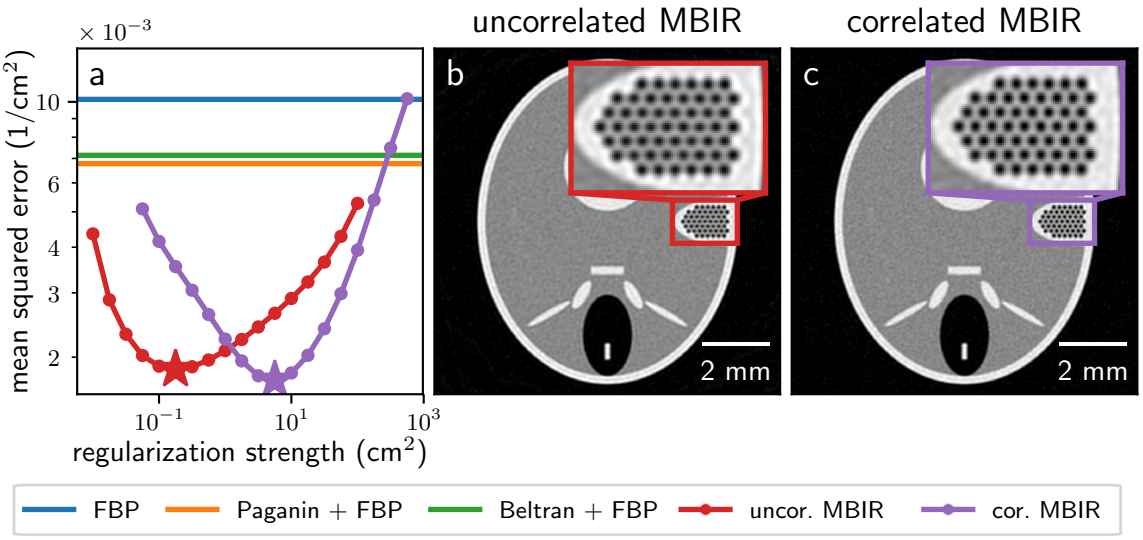

Figure 6.5.: MBIR reconstructions with lowest MSE. In (a) the MSE of the reconstructions obtained by the two MBIR methods are plotted over the regularization parameters. In addition, the MSE of the results obtained by the analytical reconstructions are included. In (b) and (c), reconstructed slices of the two MBIR methods are shown using a regularization strength resulting in the lowest MSE. In (b) the reconstruction is obtained with the 'uncorrelated MBIR' method and in (c) the reconstruction is obtained with the 'correlated MBIR' method. An adapted version of this figure was submitted for publication [Hehn et al., 2019a].

but also in higher overshoots compared to the 'correlated MBIR' method. In general, the steepness of the slope is not a good criterion to assess the resolution of the reconstruction as it is biased by the overshoots. However, the MBIR methods approximate the ground truth phantom best, as also shown by visual inspection of Figure 6.3 and Figure 6.4.

To further compare the differences between the 'uncorrelated MBIR' method and the 'correlated MBIR' method, in Figure 6.6 (b), the line profiles obtained by the two MBIR methods with lowest MSE are shown together with the ground truth line profile. The corresponding slices are shown in Figure 6.5. Again, the 'correlated MBIR' method has less overshoots compared to the 'uncorrelated MBIR' method. However, the steepness of the two line profiles is almost identical, confirming that the 'correlated MBIR' method recovers the ground truth best.

In summary, the proposed MBIR approaches can improve image quality significantly compared to the analytical reconstruction approaches by directly modeling the effects of the reduced spatial coherence and detector response. In particular, accounting for noise correlations reduces overshoots significantly.

For this simulation study, the improvements due to the new physical mean model were of interest. As the phantoms are by definition piece-wise constant, using TV regularization, which additionally favors these structures, would improve reconstruction quality further. However, to obtain a more unbiased comparison to the analytical approaches, quadratic regularization was employed, which only penalizes the magnitude of neighboring voxel values. As discussed in Section 6.3, given a certain choice for the neighborhood in the regularization term, quadratic regularization in the projection domain is connected to phase retrieval. How-

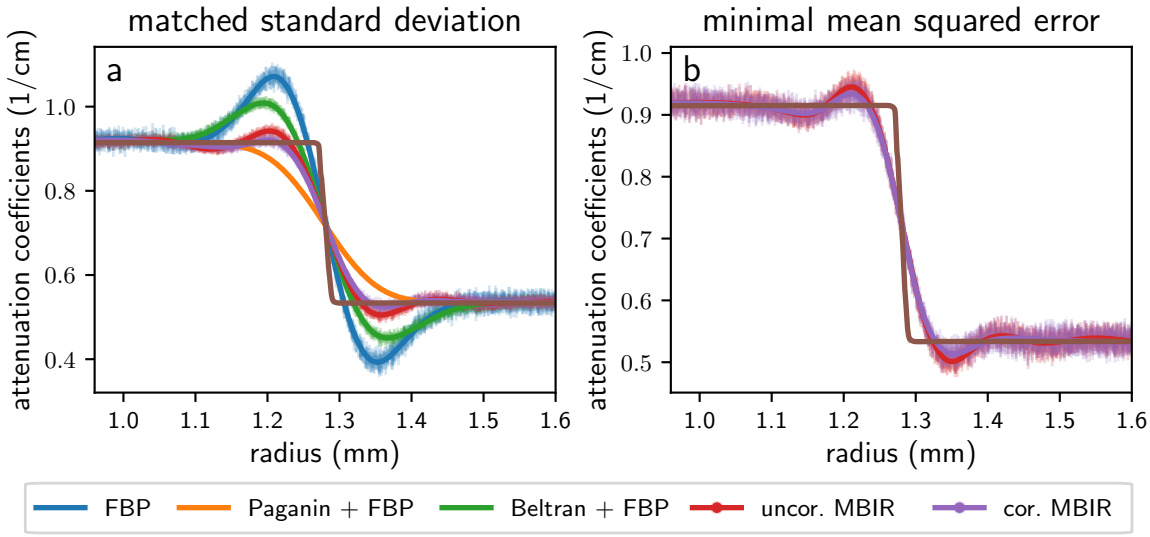

Figure 6.6.: Line profiles to evaluate edge sharpness. The line profiles are obtained around the cylindrical structure in the upper center of the phantom. Thereby, the voxel values are plotted over the distance from the center of the cylinder. The solid lines are obtained with smooth spline interpolation. In (a) the line profiles obtained from the three analytical approaches, the two MBIR methods at matched standard deviation and the ground truth are shown. The respective reconstruction can be seen in Figure 6.3 and Figure 6.4. In (b) the line profiles obtained by the two MBIR methods with lowest MSE are compared to the ground truth. The corresponding reconstruction can be seen in Figure 6.5. An adapted version of this figure was submitted for publication [Hehn et al., 2019a].

ever, regularization in CT is applied in the volume domain. Furthermore, this equivalence is only valid if no statistical noise modeling is performed.

6.6. Experimental verification at the MuCLS

In this section, the five algorithms investigated in the previous section are applied to an experimental measurement acquired at the MuCLS. Thereby, the accuracy of the physical mean model can be validated. Moreover, the effects of not having exact representations of the different operators (source blur, propagation effects and detector blur) can be analyzed.

The sample was extracted from the large-pore part of a kitchen sponge. It was chosen such that it obeys the homogeneity assumption reasonably well, has relatively complex structures of different sizes and attenuates sufficiently such that the contrast is given by a balanced combination of its attenuating and phase-shifting properties. The transmission values behind the thickest parts of the sample were around 50%.

The MuCLS was described in Subsection 6.5.1. The source was operated to generate quasi-monochromatic X-rays at an energy of 25 keV. The sample was placed approximately 4 m behind the X-ray source. The distance between the sample and the detector was set to 1 m. As detector system, a Andor Zyla sCMOS 5.5 camera (Oxford Instruments, United Kingdom) with a fiber-optically coupled gadolinium-oxysulfid scintillator and a pixelsize of 6.5 µm was used. The scan was acquired with 800 views distributed equidistantly between

0 ° and 380 °.

For preprocessing, the data was binned by a factor of two resulting in an effective pixelsize of 13.0 µm. Furthermore, individual pixel responses were corrected by reference projections without the sample in the beam. Afterwards, the incident flux was corrected by computing the respective mean values of a sample free region (resulting formally in $I_0 = 1$). Subsequently, the projections were cropped to fit the extent of the sample and the horizontal detector offset was computed from opposing projections.

The voxelsize of 10.4 µm was given by the magnification of $M = (4\,\mathrm{m} + 1\,\mathrm{m})/4\,\mathrm{m} = 1.25$. Due to the small extent of the sample of around 5 mm and the small divergence of the X-ray beam of 4 mrad, a parallel-beam geometry with unit pixelsize was assumed for simplicity. This already defines the analytical 'FBP' method given by Eq. (6.68).

The source was assumed to be described by a symmetric two-dimensional Gaussian with $\sigma_s = 50\,\mu\mathrm{m}$. In reality, the source was not perfectly symmetric [Eggl et al., 2016] and the size of the source was not known precisely for this experiment. The projected size of the source at the detector in unit pixelsize can be calculated using $\sigma_s = 50\,\mu\mathrm{m}/4\,\mathrm{m} \cdot 1\,\mathrm{m}/13.0\,\mu\mathrm{m} = 0.96$.

Similarly, an exact model of the detector response was not available. As the derivation of the 'Beltran + FBP' method assumes Gaussian blur, the detector response was also approximated by a two-dimensional symmetric Gaussian. By analyzing the noise power spectrum (NPS) of a reference corrected projection without the sample, a value of $\sigma_d = 1.14$ was found to be in best agreement with the NPS. However, in reality the response of this camera is better described by a radial power-law [Cont, 2016].

Finally, a value for $z\delta/\mu$, which accounts for the attenuating and phase-shifting properties of the sample, has to be estimated. Therefore, the 'Beltran + FBP' method was applied for different values of $z\delta/\mu - 2(\sigma_s^2 + \sigma_d^2)$. By visual inspection, a value of 15.8 was found to give the best depiction of the object. Thus, the value of $z\delta/\mu = 20.2$ by not accounting for source and detector blur was chosen for the conventional 'Paganin + FBP' method. With the values stated above, the remaining analytical algorithms 'Paganin + FBP' and 'Beltran + FBP' given by Eq. (6.69) and Eq. (6.70) are defined.

The two iterative algorithms 'uncorrelated MBIR' and 'correlated MBIR' are defined by the objective given by Eq. (6.59). The weights for the uncorrelated version are given by Eq. (6.62) and for the correlated version by Eq. (6.63). The regularization term is the same as in the simulation study given by Eq. (6.71). The L-BFGS algorithm was again used for optimization and only the curvature conditions were evaluated. By visual inspection, the number of iterations was set to 200, which provided sufficiently converged estimates. For the uncorrelated model, seven different regularization strengths were used ranging from 10^3 and 10^6 equidistantly distributed on a logarithmic scale. Similarly, for the correlated model, the seven different regularization strengths were respectively distributed between 10^1 and 10^4. As in the simulation study, a CG algorithm for precomputing the term $\mathbf{B}^T\mathbf{W}\mathbf{y}$ was used for the correlated version.

The results obtained by the five different reconstruction approaches are summarized in Figure 6.7. In the top row, the analytical reconstruction approaches are shown. In the bottom row, the standard deviations of the five reconstruction approaches computed in a sample-free region are shown with the results obtained from the MBIR methods. The reconstructions shown for the two MBIR methods are chosen such that they have mutually comparable noise levels (although having different noise characteristics). Their noise level is smaller than for the 'FBP' method, but higher than for the 'Paganin + FBP' or 'Beltran + FBP' methods. The reconstruction obtained by the 'FBP' method gives strong overshoots and highest noise levels. As already discussed in the simulation study, the sample is not

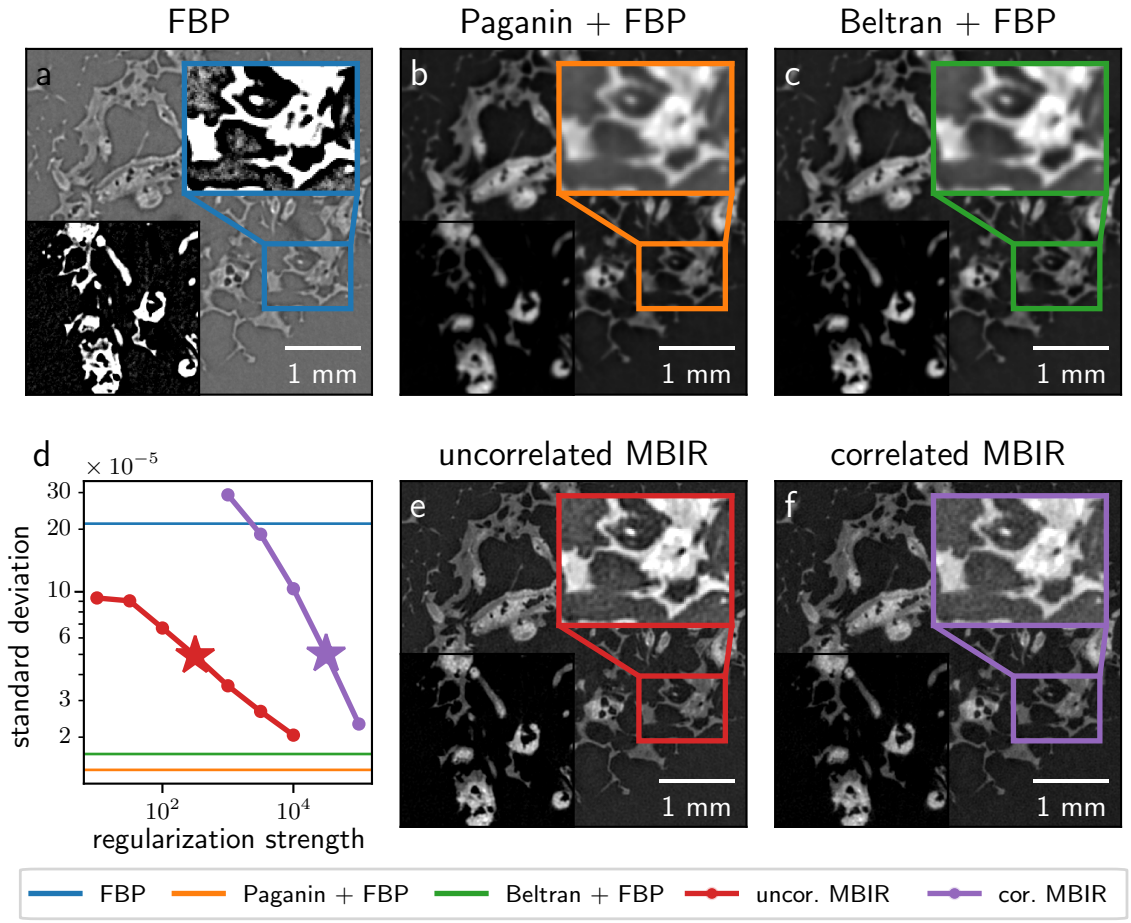

Figure 6.7.: Experimental results of the MBIR methods for PB-CT. In (a)-(c), the results obtained by the analytical reconstruction approaches are depicted. In (d), the standard deviations measured in a sample-free region in the background for all different methods and regularization strengths are depicted. In (e), (f), two MBIR reconstructions with mutually comparable noise levels are shown. The quantitative values of the reconstructions cropped to the sample are shown using a linear grayscale ranging from their minimal value to their maximum value in the region depicted. In addition, on the lower left quarters, identical linear grayscale windows are shown ranging from 0 (air) to a value chosen by visual inspection. Finally, zooms are provided with a more narrow but mutually identical grayscale window to emphasize residual overshoots at the edges of the structures. An adapted version of this figure was submitted for publication [Hehn et al., 2019a].

well represented by this method as it does not take interference effects into account. The reconstructions obtained by the 'Paganin + FBP' and 'Beltran + FBP' look very similar as the phase-retrieval algorithm in the 'Beltran + FBP' method is dominated by the factor $z\delta/\mu$. However, as already shown in the simulation study, the resolution of the 'Paganin + FBP' is worse, but also gives less overshoots at the sample edges compared to the 'Beltran + FBP' method. The MBIR methods overall give the most accurate representations of the

sample having relatively sharp edges (compared to the 'Paganin + FBP' and 'Beltran + FBP' methods) without the strong overshoots of the 'FBP' method. Again, the correlated version exhibits less overshoots at sample edges with similar sharpness compared to the uncorrelated version, resulting in an even more accurate representation of the sample.

In conclusion, the results of the experimental study are in excellent agreement with the results of the simulation study discussed in the previous section. Thus, it could be verified that the physical mean model derived in this chapter accurately describes the image formation in PB-CT using laboratory sources with reduced spatial coherence, efficient detectors and high noise levels. The crucial point is demonstrated that exact models of all components of the experimental setup are not required to improve image quality significantly. Finally, the strong impact of modeling the full covariance statistics of the noise could be demonstrated.

6.7. Conclusion

In this chapter model-based iterative phase retrieval and reconstruction algorithms for PBI and PB-CT suitable for laboratory sources were investigated. The physical mean model was extended to include the effects of the source and the detector. The formulation for the interference effects now coincides with the single-material phase-retrieval algorithm by directly imposing the homogeneity assumption from the start. In addition, the covariance description of the noise was derived and employed for the reconstruction algorithms.

This framework was used to develop an analytic phase-retrieval algorithm, which accounts for the influence of the source, detector and prior knowledge. The single-material phase-retrieval algorithm [Paganin et al., 2002] as well as a recent extension [Beltran et al., 2018], which includes the effects of partial coherence, can be identified as a special case of the derived more general algorithm. In addition, a relation between the phase-retrieval using the homogeneity assumption and regularized image denoising techniques [Rudin et al., 1992] was derived.

Moreover, MBIR approaches for PB-CT were proposed, which incorporate deconvolution of the source and the detector as well as phase retrieval in tomographic reconstruction. These approaches were validated and compared to analytical reconstruction approaches in a simulation study showing distinct improvements in image quality, in particular regarding residual overshoots at the interfaces between different materials. These results could furthermore be transferred to an experimental study performed at the MuCLS. Thus, the validity of the physical mean model and the modeling of noise correlations could be demonstrated experimentally.

Including models of the system blur in the physical mean model has been recently applied in the context of conventional attenuation-based CT for medical image applications showing distinct improvements [Tilley et al., 2018a]. However, for most medical applications the influence of the source is typically small [Hofmann et al., 2014] and also detector systems such as photon-counting detectors [Taguchi and Iwanczyk, 2013] exist that hardly spread out the signal. By contrast, for micro CT applications, where a resolution in the micrometer and sub-micrometer regime is desired, the influence of the source and the detector becomes increasingly prominent. High resolution can be achieved by using detectors with small pixelsizes, which however spread out the signal more, or by using high magnifications of the sample. However, employing higher magnification is ultimately limited by the extent of the source.

These points are even more prominent for PBI and PB-CT, which usually require high

spatial resolution. In particular, the interference effects, which hold the information about the phase-shifting properties of the sample, are located at the high-frequency components of the measurements. These high-frequency components, however, are predominantly diminished by any blurring of the system. Consequently, recovering these high-frequency components by accounting for the source and the detector directly benefits the quality of the phase information. Thus, such approaches should be particularly beneficial for these imaging techniques.

Compared to analytic reconstruction approaches, the proposed algorithms do not impose restrictions on the shape of the source or the detector. In principle, the depth dependency of the source could be included as well, which can prove useful for cone-beam systems [Tilley et al., 2016c].

Although the simulation study as well as the experimental study focused on the MuCLS, the proposed algorithms can be applied to other X-ray imaging environments with micrometer and sub-micrometer resolution, which are enhanced by phase effects. This includes measurements at laboratory microtomography systems such as the Xradia 500 Versa, discussed in Subsection 5.4.2 or the NanoCT setup detailed in Subsection 8.3.1, which uses high magnification to achieve sub-micron resolution.

Further improvements on the physical mean model could include polychromatic effects. As indicated in Section 2.4, the total intensity is given by summing up the individual monochromatic intensity components according to Eq. (2.69). This could potentially be approximated by an additional convolution. However, the effects on the monochromaticity of the source are not very strict for PBI. In [Tilley et al., 2018b], a similar framework has been used for dual energy applications, which could potentially be extended to PB-CT.

Finally, the above framework can be transferred to other phase-contrast imaging techniques that are already applied in laboratory environments. For instance, in GBI and grating-based phase-contrast computed tomography (GB-CT), the first derivative of the phase-information is encoded in the measurements. A physical mean model can be found in [Brendel et al., 2016]. In matrix-vector notation, this mean model could be adapted according to

$$\bar{\mathbf{y}}(\boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\epsilon}) = \mathbf{B}_d \mathbf{B}_s \mathbf{G} \mathcal{D}\left[e^{-\mathbf{A}\boldsymbol{\mu}}\right]\left(\mathbf{1} + \mathbf{V}\mathcal{D}\left[e^{-\mathbf{A}\boldsymbol{\epsilon}}\right]\cos\left[\mathbf{A}_\delta \boldsymbol{\delta} - \boldsymbol{\phi}\right]\right) \tag{6.73}$$

to include the effects of the source modeled by $\mathbf{B}_s$ and the detector given by $\mathbf{B}_d$. Here $\boldsymbol{\mu}, \boldsymbol{\delta}, \boldsymbol{\epsilon}$ denote the linear attenuation coefficient, refractive index decrement as well as the linear diffusion coefficient, which is related to the so-called dark-field signal, respectively. In addition, the reference intensity $\mathbf{G}$, the visibility $\mathbf{V}$ as well as the reference phase $\boldsymbol{\phi}$ are required. Moreover, $\mathbf{A}_\delta$ models a differential projection operation. The exponential and cosine functions are to be understood element-wise. The resulting covariance matrix can be written as

$$\mathbf{K}(\boldsymbol{\mu}, \boldsymbol{\phi}, \boldsymbol{\epsilon}) = \mathbf{B}_d \mathcal{D}\left[\mathbf{B}_s \mathbf{G} \mathcal{D}\left[e^{-\mathbf{A}\boldsymbol{\mu}}\right]\left(\mathbf{1} + \mathbf{V}\mathcal{D}\left[e^{-\mathbf{A}\boldsymbol{\epsilon}}\right]\mathcal{D}\left[\cos\left[\mathbf{A}_\delta \boldsymbol{\delta} - \boldsymbol{\phi}\right]\right]\right)\right]\mathbf{B}_d^T, \tag{6.74}$$

which includes the correlating noise properties of the detector blur. This approach has the potential to further increase spatial resolution in a laboratory environment, which is ultimately limited by the blur of the source and the response of the detector [Birnbacher, 2018].

7. Blind deconvolution CT reconstruction

In the previous chapter, MBIR algorithms for PBI and PB-CT have been developed, which include models of components that add to the system blur, such as the source and the detector. However, the respective components have to be characterized precisely as a slight mismatch can significantly deteriorate image quality. An underestimation of the blur leads to a lower resolution of the reconstructed volume and an overestimation of the blur results in noise amplifications and overshoots at sample edges. This also holds true for previously mentioned related methods that integrate the system blur into the reconstruction in SPECT [Yu et al., 2000, Feng et al., 2006], digital breast tomosynthesis [Zheng et al., 2018] and flat-panel cone-beam CT [Tilley et al., 2016a, Hashemi et al., 2017, Tilley et al., 2018a].

However, characterizing all components that add to the system blur can be difficult and time-consuming. This hold true, in particular, if the system blur changes over time as for instance due to a change in the system geometry, tube current, or for different samples, due to the depth dependency of the blur [Riviere and Vargas, 2008, Behling, 2015, Tilley et al., 2016c]. In addition, accurate models for the detector response are hard to obtain without exact knowledge of all components or for high-resolution detectors, which include additional optical elements that have to be aligned. Finally, blur, which is attributed to the reconstruction algorithm itself, can be difficult to characterize, such as interpolation in the projection operations or the regularization [Tward and Siewerdsen, 2008, Gang et al., 2014, Hashemi et al., 2017]. These points can hinder an effective use of the above reconstruction approaches.

Blind deconvolution methods, which are able to simultaneously estimate the system blur and recover the underlying image without prior characterization of the instrumentation, are already known from optical imaging [Chan and Wong, 1998]. These methods use regularization techniques that are similar to those encountered in CT [Rudin et al., 1992] making them promising candidates for the transfer to tomographic reconstruction. This chapter is concerned with jointly estimating parameterized models of the system blur during tomographic reconstruction in the context of flat-panel cone-beam CT. This would replace an accurate characterization of every component that adds to the system blur and could potentially be transferred to SPECT, digital breast tomosynthesis or even PB-CT.

The following chapter is structured as follows. First, the reconstruction framework, which includes a parametric blur model, is introduced. Then, blind deconvolution algorithms for optical imaging are investigated and based on these algorithms, a new regularization term for CT is proposed. Afterwards, the framework for parametric deconvolution CT reconstruction is derived and means to jointly estimate both the parameters of the blur model as well as the reconstructed object are developed. The new regularization term is validated and an extensive simulation study is implemented detailing the properties of the proposed approach. Finally, the feasibility is experimentally demonstrated at a test-bench setup.

The following studies were conducted as part of a research exchange to Johns Hopkins University, Baltimore MD, USA, at the Advanced Imaging Algorithms and Instrumentation Laboratory (Department of Biomedical Engineering). A manuscript covering most of the results presented in this chapter was published in [Hehn et al., 2019b].

7.1. CT reconstruction using a parametric blur model

This section introduces the key components towards a formulation to jointly estimate the system response during tomographic reconstruction. In the following, a parametric version of the system blur is used. This can reduce the number of free parameters significantly. However, one can imagine every pixel of the system response to be a free parameter, which would be equivalent to estimating the whole system response without any prior knowledge.

7.1.1. Objective function and physical mean model

As a first step, the physical mean model given by Eq. (4.44), which models the exponential decay in intensity is extended by a parametric blur model denoted by $\mathbf{B}(\boldsymbol{\sigma})$, which is parameterized by some blur parameters $\boldsymbol{\sigma}$. The physical mean model is thus given by

$$\bar{\mathbf{y}}(\mathbf{x}, \boldsymbol{\sigma}) = \mathbf{B}(\boldsymbol{\sigma}) I_0 e^{-\mathbf{A}\mathbf{x}}. \tag{7.1}$$

In contrast to the blur model examined in the MBIR approach of [Tilley et al., 2018a] or in the previous chapter, the blur model does not have to be known exactly, but only a parametric version of it. The actual blur parameters $\boldsymbol{\sigma}$ are not known a priori. To distinguish this approach from MBIR approaches with known models of the system response, this approach is referred to as *parametric deconvolution CT reconstruction*.

The likelihood term given by Eq. (4.43) is derived from a normal distribution. With the pair-wise Gibbs prior of Eq. (4.51), the objective reads

$$\mathcal{C}(\mathbf{x}, \boldsymbol{\sigma}) = \frac{1}{2}(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x}, \boldsymbol{\sigma}))^T \mathbf{W}(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x}, \boldsymbol{\sigma})) + \beta \mathcal{R}(\mathbf{x}), \tag{7.2}$$

with β denoting the regularization strength. One could also use the Poisson likelihood given by Eq. (4.33) instead. With the possibility of including noise correlation in the above formulation, this likelihood can be considered more versatile [Tilley et al., 2018a].

However, estimating the blur parameters in this context imposes some challenges. First, the regularization term is biased in terms of blur as the difference in neighboring voxel values is penalized [Krishnan et al., 2011]. Consequently, the response of the regularizer depends on the high-frequency components of the object. In general, if the blur parameters are chosen such that least deblurring is performed during reconstruction, a blurred estimate of the volume will be obtained. This also results in low noise levels, which in turn leads to a small value of the regularization term. By contrast, if the blur parameters correctly account for the blur of the system, the reconstructed volume has a higher resolution, but also has higher noise levels compared to the previous case. Due to higher noise levels, the value of the regularization term is higher as well. In summary, the regularizer would prefer reconstructions with low resolution and low noise levels over the desired high-resolution reconstructions with higher noise levels. In addition, the likelihood term might also be biased in terms of blur, at least to some degree, as most of the energy is contained in the low and mid frequencies [Krishnan et al., 2011]. Besides, the likelihood term can be thought of as a weighted ℓ_2 norm and in general a noiseless residuum has a lower value compared to a noisy residuum.

7.1.2. Blind deconvolution algorithms for photographs

Ideally, the regularization term should give minimal response (value) to the true sharp reconstruction. Thus, the regularization term should be independent of noise and only evaluate

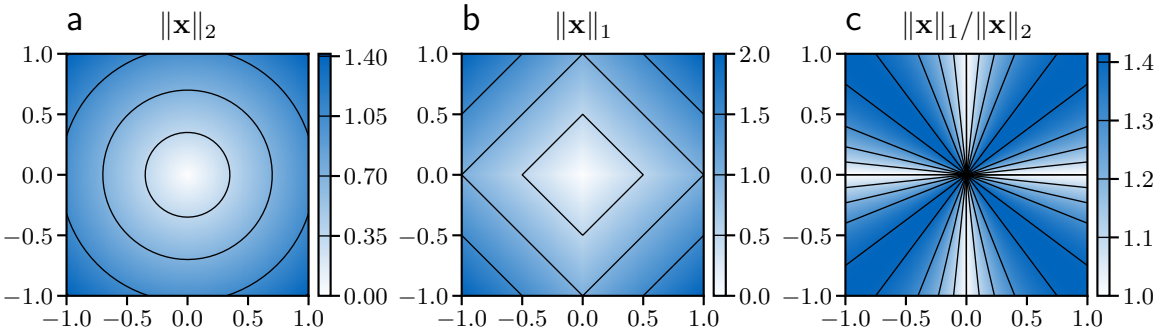

Figure 7.1.: Comparison of the ℓ_2, ℓ_1 and ℓ_1/ℓ_2 norms. The different norms are illustrated for a two-dimensional problem $\mathbf{x} = (x_1, x_2)^T$. The potentials are color-coded and the contour lines are drawn in black. In (a) the ℓ_2 norm is visualized. The potential only depends on the magnitude (distance to the origin), which makes this norm scale-dependent. In (b) the ℓ_1 norm is depicted. Again, solutions with a smaller magnitude are preferred making this norm also scale dependent. However, at the same magnitude, sparse solutions (located at the axis) are preferred. In (c), the ℓ_1 norm normalized by the ℓ_2 norm as proposed in [Krishnan et al., 2011] is shown. The potential is independent of the magnitude (radial contour lines) enforcing sparse solutions located at the axis. A similar version of the figure was published in [Hehn et al., 2019b].

the properties of the underlying noise-free sample. To find means to address the potential bias of the regularization term, regularization approaches for blind deconvolution in optical imaging are analyzed.

For instance, TV regularization[1] is also widely used in the context of digital image processing of conventional two-dimensional images [Rudin et al., 1992], where it was originally introduced. In particular, TV-based methods have become popular in the field of noise removal [Rudin et al., 1992]. In addition, these methods have been extended to blind deconvolution [Chan and Wong, 1998]. In blind deconvolution, the system blur is also unknown. Thus, the system blur and the noise-free and undistorted image have to be estimated.

Since then, many new algorithms have adapted the work of [Chan and Wong, 1998] for blindly estimating system blur. However, it was shown in [Levin et al., 2011] that any algorithm that would actually minimize the respective objective functions would yield the no-blur solution, which refers to the solution, where no (least) deblurring is performed. The reason is the aforementioned bias of the response of the regularization term that favors noiseless images with low resolution over sharp results with higher noise levels. In [Perrone and Favaro, 2014] it was shown how subtle choices in the optimization algorithm proposed in the original work of [Chan and Wong, 1998] eluded the no-blur solution, despite having lowest cost. Various approaches have been proposed to yield the desired solution, without having lowest cost, such as marginalization of the cost function, adaptive cost functions, alpha-matte extraction or edge location [Krishnan et al., 2011].

In [Krishnan et al., 2011] a metric was proposed, which evaluates the sparsity (number of

[1] There are subtle differences to the definition of the TV normalization given by Eq. (4.51) using the ℓ_1 penalty of Eq. (4.55). For instance, TV regularization in the context of [Rudin et al., 1992] is isotropic, while here the anisotropic version is used as it is easier to optimize within this framework. However, isotropic TV regularization has been used for CT reconstruction [Sidky et al., 2006].

values equal to zero) and is independent of the magnitude. Figure 7.1 shows how such a metric can be constructed for a two-dimensional problem $\mathbf{x} = (x_1, x_2)^T$. In (a) the response of the ℓ_2 norm is depicted. Therefore, $\|\mathbf{x}\|_2$ is evaluated for $-1 < x_1 < 1$ and $-1 < x_2 < 1$. The response of the ℓ_2 norm is minimal for $\mathbf{x} = \mathbf{0}$. This solution is obviously sparse. The value of this norm increases with the magnitude (distance from the center) without any preferences for sparse solutions (located at the axis). In comparison, the corresponding values for the ℓ_1 norm are depicted in (b). Again, the minimum is found for $\mathbf{x} = \mathbf{0}$. In contrast to the ℓ_2 norm some sparsity is imposed on the other solutions. For solutions with the same magnitude, sparse solutions are preferred, where one of the components is near the axis. However, solutions with less magnitude are preferred in general. The metric proposed in [Krishnan et al., 2011] scales the ℓ_1 norm by the ℓ_2 norm, which is referred to as the *normalized sparsity measure*. The values for the corresponding two-dimensional problem are depicted in (c). Here, sparse solutions near the axis are preferred. The value of this norm is independent of the magnitude resulting in radial contour lines.

If the normalized sparsity measure is employed as penalty function in the regularization term, $\mathbf{x}$ holds the difference in neighboring voxel values. In this case, sparsity implies that many neighboring pixels have the same value. On the other hand, the magnitude refers to the average difference in neighboring pixel values. A high noise level implies a large magnitude and a low noise level implies a small magnitude respectively. In that sense, the normalized sparsity measure enforces solutions that are sparse and independent of the noise level.

Employing the normalized sparsity measure has been successfully applied to blind image deconvolution of photographs [Krishnan et al., 2011] using a two-step process. First, the likelihood acts only on the high-frequency components (the concatenation of the forward difference in both dimensions of the actual images) as most energy is contained in the low and mid frequencies. Two regularization terms and therefore two parameters for the respective regularization strengths are required. One term is the normalized sparsity measure and one is the ℓ_1 norm of the kernel modeling the system blur. With this approach the kernel for the system blur is obtained in a first step. As a second step, a non-blind deconvolution is applied to recover the desired image.

7.1.3. Normalized sparsity measure for CT

Although there are substantial differences between photographs and tomographic datasets, the concept of [Krishnan et al., 2011] is used to devise a regularizer, which is less biased in terms of blur.

For TV regularization, Eq. (4.51) and the ℓ_1 penalty of Eq. (4.55) are used. To compensate for the reduction in magnitude imposed by conventional TV regularization, a normalization factor is introduced given by the square-root of the value of the quadratic regularization, which uses the ℓ_2 penalty of Eq. (4.53). This results in

$$\mathcal{S}(\mathbf{x}) = \frac{\mathcal{R}_{\ell_1}(\mathbf{x})}{\sqrt{\mathcal{R}_{\ell_2^2}(\mathbf{x})}} = \frac{\frac{1}{2}\sum_i \sum_{n\in\mathcal{N}_i} \frac{1}{\Delta_{in}} \left|\frac{x_i - x_n}{\Delta_{in}}\right|}{\sqrt{\frac{1}{2}\sum_i \sum_{n\in\mathcal{N}_i} \frac{1}{\Delta_{in}} \left(\frac{x_i - x_n}{\Delta_{in}}\right)^2}}, \tag{7.3}$$

where in the following $\mathcal{S}$ is referred to as the *normalized sparsity measure* for CT. This measure is designed to capture the sparsity of the volume and to account to some extent for the reduction in magnitude due to different noise levels.

Compared to quadratic or TV regularization, the normalized sparsity measure is non-convex. In addition, the normalized sparsity measure cannot be interpreted as a probabilistic prior. The regularizer is connected to the prior distribution over Eq. (4.50), which due to the denominator of Eq. (7.3) is not bounded. A similar argument has been given in [Krishnan et al., 2011] for the ℓ_1/ℓ_2 norm.

7.1.4. Parametric deconvolution CT reconstruction

Substituting the normalized sparsity measure for the regularizer in the objective function given by Eq. (7.2) results in

$$\mathcal{C}(\mathbf{x}, \boldsymbol{\sigma}) = \frac{1}{2}(\mathbf{y} - \bar{\mathbf{y}})^T \mathbf{W}(\mathbf{y} - \bar{\mathbf{y}}) + \beta \mathcal{S}(\mathbf{x}). \tag{7.4}$$

Due to the fact that the normalized sparsity measure is non-convex, optimizing the respective objective function is challenging. However, as proposed by [Krishnan et al., 2011] for the ℓ_1/ℓ_2 norm, one can compute the denominator of the normalized sparsity measure from the previous iterate. Thus, the objective function in each iteration can be treated as conventional TV regularization. Formally, one can define

$$\mathcal{C}_\gamma(\mathbf{x}, \boldsymbol{\sigma}) = \frac{1}{2}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x}, \boldsymbol{\sigma})\right)^T \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x}, \boldsymbol{\sigma})\right) + \beta \mathcal{S}(\mathbf{x}), \tag{7.5}$$

$$= \frac{1}{2}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x}, \boldsymbol{\sigma})\right)^T \mathbf{W}\left(\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x}, \boldsymbol{\sigma})\right) + \frac{\beta}{\gamma}\mathcal{R}_{\ell_1}(\mathbf{x}), \tag{7.6}$$

where the subscript γ indicates that the objective also depends on a normalization factor γ defined by

$$\gamma(\mathbf{x}) = \sqrt{\mathcal{R}_{\ell_2^2}(\mathbf{x})}. \tag{7.7}$$

The objective of Eq. (7.6) coincides with Eq. (7.2) with a TV regularizer up to the normalization factor γ.

Algorithm 7.1 Algorithm for parametric deconvolution reconstruction. The algorithm requires the measured intensities $\mathbf{y}$, an initial guess of the volume $\mathbf{x}^{(0)}$, the parameters $\boldsymbol{\sigma}$, which describe the blur model, the regularization strength β as well as the number of iterations N. The algorithm returns the estimate of the volume $\mathbf{x}^{(N)}$ after N iterations.

Require: $\boldsymbol{y}, \mathbf{x}^{(0)}, \boldsymbol{\sigma}, \beta, N$
  for $n \leftarrow 0$ **to** $N-1$ **do**
    $\gamma \leftarrow \sqrt{\mathcal{R}_{\ell_2^2}(\mathbf{x}^{(n)})}$
    $\Delta\mathbf{x} \leftarrow$ Using $\mathcal{L}_\gamma(\mathbf{x}, \boldsymbol{\sigma})$ and $\nabla_{\mathbf{x}}\mathcal{L}_\gamma(\mathbf{x}, \boldsymbol{\sigma})$ around $\mathbf{x}^{(n)}$
    $\mathbf{x}^{(n+1)} \leftarrow \mathbf{x}^{(n)} + \Delta\mathbf{x}$
  end for
  return $\mathbf{x}^{(N)}$

Algorithm 7.1 summarizes how this objective function can be optimized for $\mathbf{x}$ according to

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \mathcal{C}_\gamma(\mathbf{x}, \boldsymbol{\sigma}) \tag{7.8}$$

assuming that the blur parameters $\boldsymbol{\sigma}$ are known. The algorithm does not require any additional information compared to the conventional approach. Thus, the projection measurements $\mathbf{y}$, an initial guess of the volume $\mathbf{x}^{(0)}$, the blur parameters $\boldsymbol{\sigma}$, the regularization

strength β as well as the number of iterations N need to be provided. However, the initial guess of the volume must ensure $\gamma(\mathbf{x}^{(0)}) \neq 0$. In practice, an analytical reconstruction should fulfill this criterion. In every iteration $n \in \{0, \dots, N-1\}$, the normalization is computed from the current estimate $\mathbf{x}^{(n)}$. Afterwards, a gradient-based optimization algorithm is used to compute the update $\Delta\mathbf{x}$ from the objective function, which is subsequently added to obtain the new estimate of the volume $\mathbf{x}^{(n+1)}$. After N iterations the current estimate $\mathbf{x}^{(N)}$ is returned. However, the above algorithm still requires the exact blur parameters $\boldsymbol{\sigma}$.

7.2. Blind deconvolution CT reconstruction algorithm

Blind deconvolution reconstruction refers to estimating the blur parameters $\boldsymbol{\sigma}$ jointly with the volume $\mathbf{x}$. Formally, this can be written as

$$\{\hat{\mathbf{x}}, \hat{\boldsymbol{\sigma}}\} = \arg\min_{\mathbf{x},\boldsymbol{\sigma}} \mathcal{C}_\gamma(\mathbf{x}, \boldsymbol{\sigma}). \tag{7.9}$$

If the volume is defined by N^3 voxels and the blur matrix is parameterized by M parameters, this problem has $N^3 + M$ dimensions. In principle, by calculating the gradients

$$\nabla_{\mathbf{x}} \mathcal{C}_\gamma(\mathbf{x}, \boldsymbol{\sigma}) = \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\sigma}) + \nabla_{\mathbf{x}} \mathcal{S}(\mathbf{x}) \tag{7.10}$$

$$\nabla_{\boldsymbol{\sigma}} \mathcal{C}_\gamma(\mathbf{x}, \boldsymbol{\sigma}) = \nabla_{\boldsymbol{\sigma}} \mathcal{L}(\mathbf{x}, \boldsymbol{\sigma}) + \nabla_{\boldsymbol{\sigma}} \mathcal{S}(\mathbf{x}) \tag{7.11}$$

the above objective could be optimized for the blur parameters as well as the volume using gradient-based optimization approaches. However, the gradient of the normalized sparsity measure with respect to the blur parameters gives

$$\nabla_{\boldsymbol{\sigma}} \mathcal{S}(\mathbf{x}) = \mathbf{0} \tag{7.12}$$

as this measure does not directly depend on the blur parameters. The volume only indirectly depends on the blur parameter via the objective $\mathcal{C}_\gamma(\mathbf{x}, \boldsymbol{\sigma})$. Thus, the normalized sparsity measure cannot be utilized for estimating the blur parameters.

As a first step towards optimizing the objective in blind deconvolution reconstruction the following new objective can be defined

$$\mathcal{C}_\gamma(\boldsymbol{\sigma}) = \min_{\mathbf{x}} \mathcal{C}_\gamma(\mathbf{x}, \boldsymbol{\sigma}), \tag{7.13}$$

which holds the values of the non-blind reconstruction problem in convergence for different blur parameters. Thus, the dimensionality of $\mathcal{C}_\gamma(\mathbf{x}, \boldsymbol{\sigma})$ is reduced from $N^3 + M$ dimensions to M dimensions. It is reasonable to assume that for the true blur parameters this objective is optimal, as the respective volume would be most sparse and thus the response of the normalized sparsity measure would be minimal and the volume would be consistent with the measurements resulting in a small response of the likelihood as well. Thereby it is assumed that the influence of the likelihood term does not affect the optimum of the objective function. Thus, to solve

$$\hat{\boldsymbol{\sigma}} = \arg\min_{\boldsymbol{\sigma}} \mathcal{C}_\gamma(\boldsymbol{\sigma}) \tag{7.14}$$

one can use an extensive search over the blur parameters to find the optimum or other direct search algorithms such as the Nelder-Mead algorithm [Nelder and Mead, 1965].

However, to obtain $\mathcal{C}_\gamma(\boldsymbol{\sigma})$ at any point, a whole iterative reconstruction has to be performed, which is time consuming. Moreover, the computational cost increases exponentially

with the number of unknown blur parameters. Consequently, this approach is not feasible for any realistically sized datasets.

Assuming a low-dimensional blur model defined by $\boldsymbol{\sigma} = \sigma$, which is a scalar quantity, the objective function given by Eq. (7.13) becomes one-dimensional. Thus, one may solve this objective according to

$$\hat{\sigma} = \arg\min_{\sigma} \mathcal{C}_\gamma(\sigma). \tag{7.15}$$

However, this approach is not particularly efficient, as a whole non-blind reconstruction using Algorithm 7.1 would have to be performed to yield $\mathcal{C}_\gamma(\sigma)$ at any point making this approach still unfeasible in practice.

Algorithm 7.2 Algorithm for blind deconvolution reconstruction. As in Algorithm 7.1, the intensity measurement $\mathbf{y}$, an initial guess of the volume $\mathbf{x}^{(0)}$, the regularization strength β and the number of iterations N are required. However, only an initial guess of the blur parameter $\sigma^{(0)}$ has to be provided, a parameter $\Delta\sigma$ defining the accuracy of the blur parameter and the number of global iterations M. The algorithm returns both the estimate of the volume $\mathbf{x}^{(M)}$ as well as an estimate of the blur parameter $\sigma^{(M)}$ after M iterations.

Require: $\boldsymbol{y}, \mathbf{x}^{(0)}, \sigma^{(0)}, \beta, N, \Delta\sigma, M$
  for $m \leftarrow 0$ **to** $M-1$ **do**
    forall $\tilde{\sigma} \in \{\sigma^{(m)} - \Delta\sigma, \sigma^{(m)}, \sigma^{(m)} + \Delta\sigma\}$ **do in parallel**
      $\mathcal{L}_{\tilde{\sigma}}, \mathbf{x}_{\tilde{\sigma}} \leftarrow$ Compute $\min_{\mathbf{x}} \mathcal{L}_\gamma(\mathbf{x}, \tilde{\sigma})$ using $\boldsymbol{y}, \mathbf{x}^{(m)}, \tilde{\sigma}, \beta, N$ {See algorithm (7.1)}
    end for
    $\mathbf{x}^{(m+1)}, \sigma^{(m+1)} \leftarrow$ Choose according to $\min\{\mathcal{L}_{\sigma^{(m)}-\Delta\sigma}, \mathcal{L}_{\sigma^{(m)}}, \mathcal{L}_{\sigma^{(m)}+\Delta\sigma}\}$
  end for
  return $\mathbf{x}^{(M)}, \sigma^{(M)}$

Instead, a computationally more efficient approach to solve the objective function given by Eq. (7.13) is proposed for a one-dimensional blur parameter $\boldsymbol{\sigma} = \sigma$. The aim is to compute Eq. (7.15) only approximately and locally and update the volume and blur parameter dynamically. The corresponding algorithm is outlined in Algorithm 7.2. Compared to Algorithm 7.1 only a rough estimate of the blur parameter $\sigma^{(0)}$ is required and additionally a parameter $\Delta\sigma$ defining the resolution of the blur parameters as well as the number of global iterations M, discussed below. In every global iteration $m \in \{0, \ldots, M-1\}$, three non-blind reconstructions are performed in parallel for the current estimate of the blur parameter, a blur parameter reduced by $\Delta\sigma$ and a blur parameter increased by the same amount respectively. The respective non-blind reconstructions are not iterated until convergence but only for N intermediate iterations. From these three reconstructions, new best estimates for the volume and the blur parameter are chosen according to the lowest objective value of these reconstructions. Consequently, in the next global iteration, three new non-blind reconstructions are performed distributed around the (new) best estimate for the blur parameters. In addition, the new best estimate for the volume is used in the respective reconstruction as initialization. This procedure is repeated until the number of global iterations M is exceeded. The algorithm then returns the best estimate of the blur parameter as well as the volume.

7.3. Validating of the normalized sparsity measure for CT

To validate the properties of the normalized sparsity measure and to work out the differences to TV regularization, the following restoration problem is devised.

Specifically, the physical mean model follows Eq. (7.1). Thereby, $\mathbf{A} = \mathbf{I}$ is set to the identity and $\mathbf{B}(\boldsymbol{\sigma}) = \mathbf{B}_{3\mathrm{d}}(\sigma)$ denotes a three-dimensional symmetric Gaussian blur kernel. In this case $\mathbf{x}$ is the three-dimensional volume without the blur, which has to be estimated. The measurement $\mathbf{y}$, which is a three-dimensional volume which is corrupted by the Gaussian blur and noise, is generated according to

$$\mathbf{y} \sim \mathbf{B}_{3\mathrm{d}}(\sigma_{\mathrm{true}})\,\mathcal{N}\left(I_0 e^{-\mathbf{x}}, \mathbf{D}\left[I_0 e^{-\mathbf{x}}\right]\right), \tag{7.16}$$

where $\mathcal{N}$ denotes the multivariate normal distribution with equal mean and variance to approximate a Poisson distribution. The true blur parameter is denoted by σ_{true}. In this model, the volume $\mathbf{x}$ can be directly restored from the measurement according to

$$\mathbf{x} = -\log\left(\frac{\mathbf{B}_{3\mathrm{d}}^{-1}(\sigma)\mathbf{y}}{I_0}\right), \tag{7.17}$$

where the inversion of the blur $\mathbf{B}_{3\mathrm{d}}^{-1}$ can be performed by Fourier inversion, as the blur operator acts after the values have been drawn from the normal distribution. This results in fully correlated noise.

For the following simulation study, the volume $\mathbf{x}$ was chosen according to the FORBILD head phantom [FORBILD phantoms, 1999] with 256×256 voxels and 32 slices at an energy of 100 keV. To avoid partial volume effects the phantom was generated with 2048×2048 voxels and 256 slices and then downsampled by a factor of 8 to the aforementioned size. The true blur parameter was chosen to be $\sigma = 1$ px. Three different measurements were simulated with different noise levels defined by $I_0 \in \{10^5, 10^6, \infty\}$ according to Eq. (7.16)[2]. For each of the three measurements, restorations in the range $\sigma \in [0.8\,\mathrm{px}, 1.1\,\mathrm{px}]$ were calculated according to Eq. (7.17).

For all three measurements and all restored volumes, the response of the TV regularization and the normalized sparsity measure given by Eq. (7.3) were calculated. The results are depicted in Figure 7.2. In Figure 7.2 (a)-(c), the response of the TV regularizer is depicted in blue as a function of the blur parameters used for restoration. The values of the TV regularizer are attributed to the left axis. In addition, the corresponding response of the normalized sparsity measure is visualized in orange and attributed to the right axis. In Figure 7.2 (a) the high noise scenario with $I_0 = 10^5$ is shown, in Figure 7.2 (b) the medium noise scenario with $I_0 = 10^6$ is depicted and Figure 7.2 (c) shows the case without noise for $I_0 = \infty$. In Figure 7.2 (a), the response of both the TV regularizer and the normalized sparsity measure increase with increasing blur parameter. There is no distinct optimum visible at the true blur parameters. In Figure 7.2 (b), the response of the TV regularizer still increases with increasing blur parameter without any distinct feature at the true blur parameter. By contrast, the normalized sparsity measure has a distinct optimum for the true blur parameter. Thus, one could recover the true blur parameter used for the simulation by minimizing the response of the normalized sparsity measure. Finally, for the noiseless scenario depicted in Figure 7.2 (c), the TV regularizer remains overall increasing with increasing blur parameter. The optimum of the normalized sparsity measure becomes very distinct. For an underestimation of the blur parameter the curvature of the response is negative and for an overestimation the curvature is positive. A gradient-based approach would allow to find the optimum easily. In Figure 7.2 (d), an extract of the ground truth

[2] As the measurements are normalized afterwards, $I_0 = \infty$ is implemented according to $\mathbf{y} = \mathbf{B}_{3\mathrm{d}}(\sigma_{\mathrm{true}}) I_0 e^{-\mathbf{x}}$ with an abitrary value for I_0.

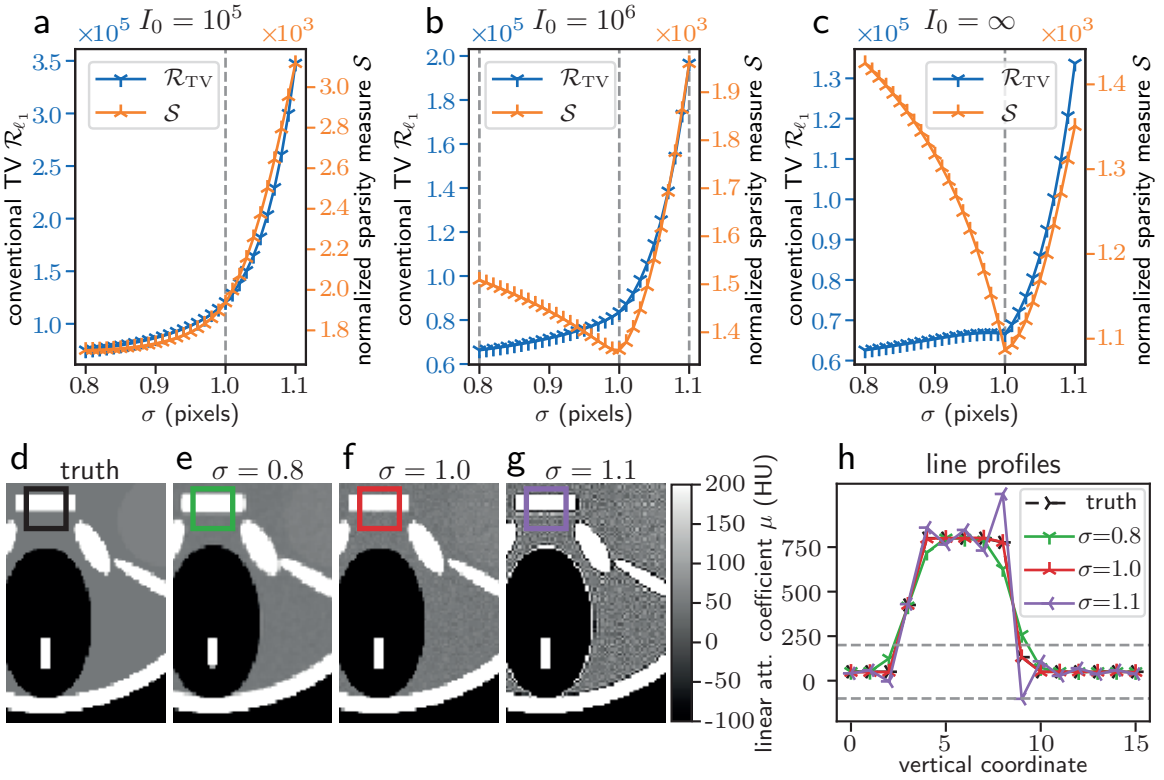

Figure 7.2.: Three-dimensional parametric deconvolution restoration. In (a)-(c), the response of the conventional TV regularizer and the normalized sparsity measure is plotted over the blur parameter used in the restoration. The response of the conventional TV regularizer is attributed to the left axis, while the response of the normalized sparsity measure refers to the right axis respectively. In (a), the flux was set according to $I_0 = 10^5$ and in (b) the flux was given by $I_0 = 10^6$. In (c) no noise is present. An extent of the ground truth phantom is shown in (d). In (e)-(f) different restorations for the data generated with $I_0 = 10^6$ using different blur parameters are depicted. To emphasize the sharpness of the edges of the respective restorations in (e)-(g), the horizontally averaged line profiles in the regions marked by the respective rectangles are shown. A similar version of the figure was published in [Hehn et al., 2019b].

phantom is depicted. In Figure 7.2 (e)-(f), three restorations for $I_0 = 10^6$ using $\sigma = 0.8\,\mathrm{px}$, $\sigma = 1.0\,\mathrm{px}$ and $\sigma = 1.1\,\mathrm{px}$ are depicted respectively. One can see that for underestimated blur parameters the edges are blurred out, for the correct blur parameters sharp edges are recovered and for an overestimation high noise levels and overshoots at edges are visible. The line profiles in Figure 7.2 (h), obtained by vertically averaging the corresponding regions marked in Figure 7.2 (e)-(f), further illustrate the noise level and edge sharpness.

In conclusion, if the noise level is not too high, the true blur parameter can be recovered by minimizing the response of the normalized sparsity measure. By contrast, the response of the TV regularizer increases with increasing blur parameter as the noise increases in the deconvolution step.

7.4. Evaluation of blind deconvolution reconstruction using a simulation study

After validating the normalized sparsity measure for three-dimensional volumes in a highly simplified restoration problem, this section addresses the main objective of the chapter, namely to demonstrate the feasibility of blind deconvolution reconstruction for CT applications. Therefore, three simulation studies were designed to validate if the objective with the normalized sparsity measure given by Eq. (7.6) has an optimum at the true blur parameter as this would imply that the volume and the blur parameters could, in principle, be estimated simultaneously during reconstruction. In addition, the feasibility of Algorithm 7.2 to blindly estimate low-dimensional blur parameters during reconstruction was examined.

As in the previous section, the FORBILD head phantom [FORBILD phantoms, 1999] with 256×256 voxels and 32 slices at an energy of 100 keV was used as the volume. Again, the volume was initially oversampled by a factor of 8 to avoid partial volume effects. The number of voxels defined the voxel size of 0.1 cm. Here, the mean model given by Eq. (7.1) was used according to

$$\bar{\mathbf{y}}(\mathbf{x}, \sigma) = \mathbf{B}(\sigma) I_0 e^{-\mathbf{A}\mathbf{x}}. \tag{7.18}$$

The matrix $\mathbf{A}$ modeled a parallel-beam geometry with 403 projections, which were equidistantly distributed between 0 ° and 180 °. This is in accordance with the Nyquist criterion given by Eq. (4.18). The flux was given by $I_0 = 10^4$ and the matrix $\mathbf{B}$ accounted for a two-dimensional symmetric Gaussian blur parameterized by the scalar σ_{true}, which models the detector response. The measurements were generated according to

$$\mathbf{y} \sim \mathbf{B}(\sigma_{\text{true}}) \mathcal{N}\left(I_0 e^{-\mathbf{A}\mathbf{x}}, \mathbf{D}\left[I_0 e^{-\mathbf{A}\mathbf{x}}\right]\right). \tag{7.19}$$

Again the values were drawn from a normal distribution with equal mean and variance defined by the expected number of photons behind the sample.

7.4.1. Comparison to conventional MBIR

The first study was concerned with the differences between CT reconstruction using TV regularization and the normalized sparsity measure. The former uses Eq. (7.2) with $\mathcal{R} = \mathcal{R}_{\ell_1}$ and the mean model defined by Eq. (7.18), the latter the objective given by Eq. (7.6) with the normalized sparsity measure and the same mean model as above.

The true blur parameter used for simulation of the measurements according to Eq. (7.19) was set to $\sigma_{\text{true}} = 1$ px. For both approaches, the respective objective functions were solved for all $\sigma \in \{0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4\}$ px using Algorithm 7.1. In addition, the reconstructions were performed for different regularization strengths $\beta \in \{0.5, 1, 2, 4\}\beta_0$. Here, β_0 denotes the regularization strength for $\sigma = \sigma_{\text{true}}$ that yields the same variance in a uniform region of the sample as the analytical reconstruction. Consequently, the value of β_0 for the reconstruction approach using TV regularization is smaller than the value of β_0 for the normalized sparsity regularization. For optimizing the respective objective functions, the L-BFGS algorithm with 2000 iterations was used to ensure sufficiently converged estimates of the volume. Instead of the L-BFGS algorithm, other algorithms to solve non-linear problems with many unknowns could be used as well. For the approach using the normalized sparsity measure the L-BFGS algorithm was utilized with Algorithm 7.1. In both cases, the analytical FBP reconstruction was used as initialization.

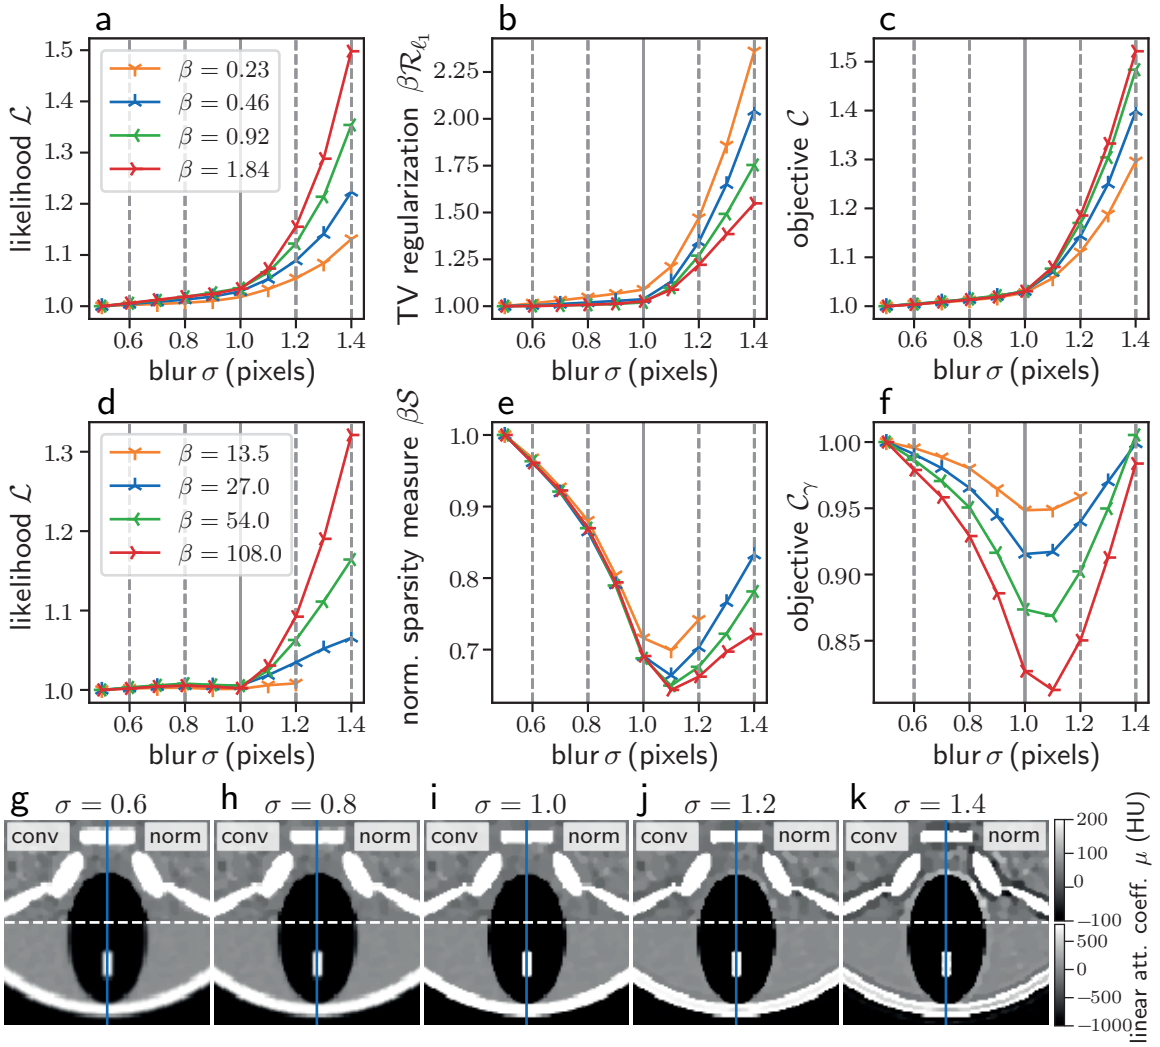

Figure 7.3.: Comparison of the TV regularizer with the normalized sparsity measure in CT. In (a)-(c) the different parts of the objective with TV regularization are plotted over the blur parameters for different regularization strengths. In (a) the likelihood term is depicted, in (b) the response of the TV regularizer and in (c) the whole objective function. The values are normalized to the respective first data-points. In (d)-(f) the same information about the objective function using the normalized sparsity measure is shown. In (g)-(k) extracts of reconstructed slices are shown for different blur parameters for $\beta = \beta_0$. The upper halves are depicted in a more narrow window, therefore focusing on the noise properties. The lower halves show the full range of values to examine the resolution properties at the transition from bone to air or soft-tissue. The left halves show the reconstructed slices using the objective function with TV regularization, while the right halves were obtained by the objective using the normalized sparsity measure. A similar version of the figure was published in [Hehn et al., 2019b].

In Figure 7.3 the two objective functions are compared. In Figure 7.3 (a)-(c) the different parts of the objective function with the TV regularization are shown. The likelihood term shown in (a) as well as the response of the TV regularizer depicted in (b) increase with increasing blur parameter for all regularization strengths. Thus, the total objective function is also increasing with increasing parameters for all regularization strength. Consequently, jointly estimating the blur parameter during reconstruction would potentially recover the no-blur solution. Thus, with this objective function a joint reconstruction is not feasible.

In Figure 7.3 (d)-(f) the corresponding information of the objective function with the normalized sparsity is shown. In contrast to the previous case, the likelihood term of this objective function as shown in (a) is not strictly increasing for increasing blur parameters. Most notably, for an underestimation of the blur parameter the likelihood term remains rather constant. Thus, the likelihood term does not have a strong influence on the total objective function in this regime. In (b) the corresponding responses of the normalized sparsity measure are depicted. Similar to Figure 7.2 (b)-(c), these responses have a distinct optimum around the true blur parameter for all regularization strengths. However, the optima of the responses of the normalized sparsity measure are slightly overestimated, which hints towards additional blur components of the system such as interpolation blur in the projection operations or the fact that the non-binary edges of the (ground truth) phantom are compensated. Finally, as shown in Figure 7.3 (f), the values of the total objective functions also have distinct optima around the true blur parameter for all regularization strengths. As for the response of the normalized sparsity measure, the blur parameters are slightly overestimated for the respective regularization strengths. However, this effect is alleviated by the influence of the likelihood term, which increases for an overestimation of the blur parameters. Furthermore, it can be observed that for decreasing regularization strength, the overestimation of the blur parameter becomes smaller. The fact that there are distinct optima in the total objective functions for all regularization strengths makes a joint estimation of the blur parameter, in principle, feasible.

Extracts of reconstructed slices for both approaches and different blur parameters are shown for $\beta = \beta_0$ in Figure 7.3 (g)-(k). The left halves show the reconstructed slices obtained by minimizing the objective with TV regularization, while the right halves were obtained by optimizing the objective function with the normalized sparsity measure. As the regularization strength $\beta = \beta_0$ for $\sigma = 1.0$ is matched for both approaches, the reconstructions of both approaches look similar. For both approaches, the noise level increases with increasing blur parameter as seen from the upper halves with the more narrow window. However, if the blur parameter is underestimated, the noise level of the approach using the normalized sparsity measure is lower compared to the approach that uses TV regularization. For an overestimation of the blur parameters, higher noise levels are present if the normalized sparsity measure is used. This results from the normalization in the normalized sparsity measure, as for increasing blur parameter, the noise level of the reconstructed volume increases and thus the value for the normalization increases, which in return results in a lower effective regularization strength. From the lower halves, which show the full range of the linear attenuation coefficients, the resolution properties of different blur parameters can be analyzed. As expected, with increasing underestimation of the blur parameter, the interface between bone and air smears out as the resolution decreases. For the reconstruction with the true blur parameter a sharp transition can be observed and with increasing overestimation, increasing overshoots at the edges become visible.

In conclusion, this study shows that using the objective function with TV regularization cannot be used to estimate the blur parameter. In contrast, the objective function with the

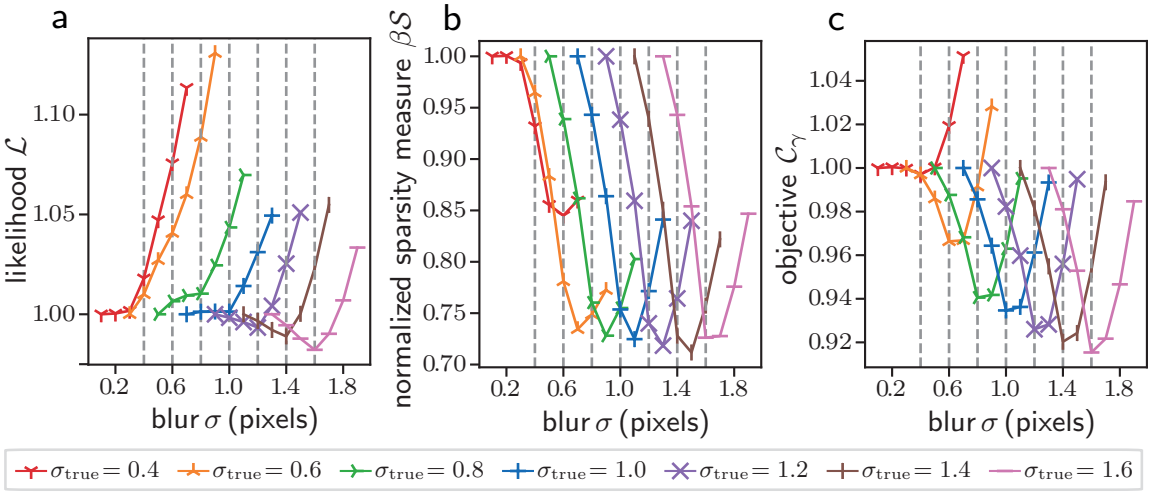

Figure 7.4.: Validation of blind deconvolution reconstruction for different system blurs. In (a)-(c) the different parts of the objective function using the normalized sparsity measure are depicted over the blur parameter for different measurements simulated with varying blur parameters. For better visualization, the data is normalized to the value with the smallest blur parameter for each of the seven measurements. Data points belonging to the same measurement are depicted in the same color. In (a) the likelihood terms are shown. In (b) the response of the normalized sparsity measure is depicted and the total objective function is shown in (c). A similar version of the figure was published in [Hehn et al., 2019b].

normalized sparsity measure yields a distinct optimum around the true blur parameter for all regularization strengths and can thus be used for jointly estimating the blur parameter during reconstruction.

7.4.2. Dependency on the blur size

To further validate the feasibility of jointly estimating the blur parameter during reconstruction using the objective function given by Eq. (7.6) with the mean model given by Eq. (7.18), seven independent measurements were simulated. The measurements were generated according to Eq. (7.19) with $\sigma_{\text{true}} \in \{0.4, 0.6, 0.8, 1.0, 1.2, 1.4, 1.6\}$ px. For each of the seven measurements, seven reconstructions were preformed with $\sigma \in \sigma_{\text{true}} \pm \{0.0, 0.1, 0.2, 0.3\}$ px. To account for the fact that the noise characteristics change for the different simulated measurements, the regularization strength $\beta = \beta_0$ used for $\sigma_{\text{true}} = 1.0$ px was scaled with the relative difference in variance of the other measurements to make the noise levels in the reconstruction of different measurements comparable. The L-BFGS algorithm was used in combination with Algorithm 7.1 with $N = 1000$ initialized by an FBP reconstruction.

In Figure 7.4, the different parts of the objective function are shown as a function of the blur parameter for the seven different measurements. The likelihood terms are shown in (a). For measurements simulated with small blur parameters the likelihood terms are increasing monotonously with the blur parameter. However, for measurements simulated with larger blur parameters, a distinct optimum can be observed at the true blur parameter. In (b) the responses of the normalized sparsity measure are depicted. For all measurements, a distinct optimum is observed around the true blur parameter. As in the previous case, the

true blur parameter is slightly overestimated. This overestimation becomes less significant for measurements simulated with increasing blur parameter. The total objective function is shown in (c). For all measurements the respective true blur parameters can be recovered as a distinct optimum can be observed. For measurements simulated with increasing blur parameter, the optimum becomes more prominent and the overestimation of the true blur parameter less significant.

This simulation study further shows that the true blur parameter used for simulating the measurements can be recovered over a wide range, indicating the feasibility of jointly estimating the parameter during reconstruction.

7.4.3. Optimization algorithm for blind deconvolution reconstruction

Finally, Algorithm 7.2 is evaluated, which omits an extensive search over the whole parameter space and blindly estimates the true blur parameter during reconstruction. Therefore, measurements were simulated with $\sigma_{\text{true}} = 1.0\,\text{px}$. Algorithm 7.2 was provided with the simulated measurements, an initial guess for the volume obtained by an analytical FBP reconstruction, the regularization strength $\beta = \beta_0$, $N = 100$ iterations, $\Delta\sigma = 0.1\,\text{px}$ defining the resolution of the blur parameter and $M = 5$, the number of global iterations. The algorithm was executed twice, using two different values for the initial guess of the blur parameter, namely $\sigma^{(0)} \in \{0.7, 1.3\}\,\text{px}$ respectively. Independently of Algorithm 7.2, the values of the objective function were tracked over $100, 200, 300, 400, 500$ and 2000 iterations for $\sigma \in \{0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4\}\,\text{px}$ for reference using Algorithm 7.1.

In Figure 7.5, the sequence of events of Algorithm 7.2 is illustrated. In (a) the initial guess of the blur parameter was set to $\sigma^{(0)} = 0.7\,\text{px}$. Thus, in the first global iteration $m = 0$, three reconstructions were performed for $N = 100$ iterations in parallel for $\tilde{\sigma} \in \{0.6, 0.7, 0.8\}\,\text{px}$ as indicated by the colored dots. The values of the objective function for these three reconstructions coincide perfectly with the values of the objective function obtained by the reference reconstructions shown in the background. As the value of the objective function for $\tilde{\sigma} = 0.8\,\text{px}$ is the lowest of the three reconstructions, further indicated by a large colored dot, the respective volume and blur parameter are kept for the subsequent global iteration. In the next global iteration $m = 1$, reconstructions for $\tilde{\sigma} \in \{0.7, 0.8, 0.9\}\,\text{px}$ were performed using the volume kept from the previous global iteration. As a consequence of the initialization, the values of the objective function for $\tilde{\sigma} \in \{0.7, 0.9\}\,\text{px}$ do not match the values of the objective function obtained by the reference reconstructions anymore. However, the value of the objective function for $\tilde{\sigma} = 0.8\,\text{px}$ matches the value of the objective function of the reference reconstruction rather nicely. The only difference between this reconstruction and the reference reconstruction is that the optimization algorithm was re-initialized after $N = 100$ iterations. Subsequently, the volume corresponding to $\tilde{\sigma} = 0.9\,\text{px}$ is kept and a new global iteration $m = 2$ is started. In this global iteration, the correct blur parameter $\tilde{\sigma} = 1.0\,\text{px}$ was already obtained and did not change in the subsequent global iterations. In the following iterations only the estimate of the volume was improved. Similarly, as shown in (b) for a different initialization of the blur parameter according to $\sigma^{(0)} = 1.3\,\text{px}$, Algorithm 7.2 recovers the true blur parameter after four global iterations, as well.

Thus, Algorithm 7.2 was able to jointly estimate the blur parameter during reconstruction, without an extensive search over the whole parameter space independently of the exact value of the initial guess for the blur parameter.

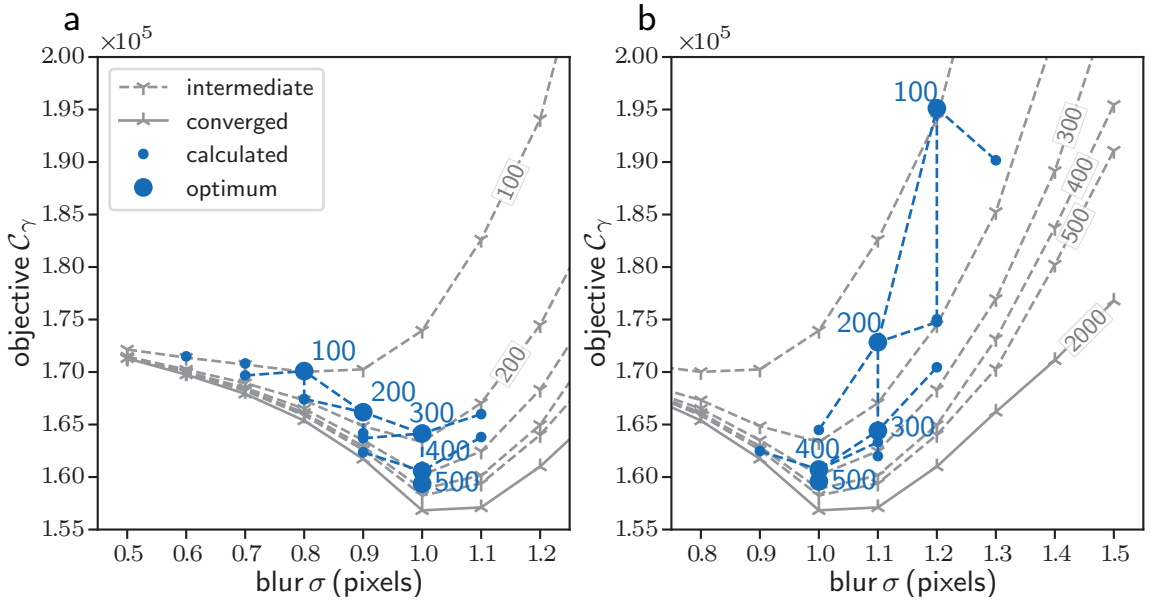

Figure 7.5.: Optimization algorithm for blind deconvolution reconstruction. In (a) and (b) the values of the objective function are plotted for reference over the blur parameter after 100, 200, 300, 400 and 500 iterations using dashed gray lines. After 2000 iterations, ensuring sufficiently converged estimates, the values of the objective function are plotted using a solid gray line. In (a) the steps of Algorithm 7.2 are visualized for $\sigma^{(0)} = 0.7\,\text{px}$ and in (b) the corresponding information is shown for $\sigma^{(0)} = 1.3\,\text{px}$. Colored dots indicate which reconstructions have been explicitly calculated in which global iteration. Large colored dots indicate the best reconstruction in the respective global iteration. Dashed colored lines indicate which information of the previous global iterations was used for the reconstructions in the subsequent global iteration. A similar version of the figure was published in [Hehn et al., 2019b].

7.5. Experimental verification at a test-bench setup

Finally, blind deconvolution reconstruction is applied to a human wrist phantom measured at a cone-beam CT test bench. This study is designed to validate the applicability of blind deconvolution reconstruction to real datasets.

7.5.1. Sample, test-bench setup and preprocessing

The sample used for the validation was a human wrist phantom. The test-bench setup consists of a flat-panel detector (4030CB, Varian, Palo Alto CA, USA) with a pixelsize of 0.388 mm. The X-ray tube (Rad-94, Varian, Salt Lake City UT, USA) was operated at 100 keV with 100 mA. The exposure time was set to 6.3 ms per projection. The source-axis distance was 595 mm and the source-detector distance was 1184 mm.

For preprocessing, the acquired measurements were gain and offset corrected and afterwards binned by a factor of two to 340×100 pixels. In total, 360 measurements were acquired equidistantly distributed around the sample. The dimensions of the volume were 340×340 voxels in 120 slices. The 20 additional slices account for limited data in the cone-beam

direction.

7.5.2. Reconstructions of the human wrist phantom

The physical model is given by Eq. (7.18), which only models a symmetric Gaussian (detector) blur using a single scalar blur parameter. Algorithm 7.1 was used for reconstruction. Thereby, the 360 preprocessed measurements were provided as well as an FBP reconstruction as initial guess for the volume. The regularization strength was set to $\beta = 0.01$ according to visual inspection and the number of iterations was set to 400. To further validate the dependency of the value of the different parts of the objective function on the blur parameter, a series of reconstructions with $\sigma \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$ px was performed.

The results are summarized in Figure 7.6. In (a)-(c) the values of the different parts of the objective function are shown as a function of the blur parameter for different numbers of iterations. In (a) the values of the likelihood term are shown, which increase with increasing blur parameter as already shown in Figure 7.4 (a) in the simulation study for small blur parameters. For these small blur parameters, the values after the first 100 iterations already closely resemble the values after 400 iterations. For larger values of the blur parameter, the values of the likelihood term still decrease with increasing number of iterations. In (b) the responses of the normalized sparsity measure are shown. For increasing numbers of iterations an optimum at $\sigma = 0.8\,\text{px}$ can be observed. As noted in Figure 7.4 (b) for the simulation study, this value usually overestimated the true blur parameter. In (c) the total objective function is shown. The overall trend of the values of the objective function closely resembles the results from the simulation study shown in Figure 7.5. As in Figure 7.5, the optimum slightly increases within the first 200 iterations, before it remains at $\sigma = 0.6\,\text{px}$. As validated on the simulation study, Algorithm 7.2 would, in principle, return this value for the blur parameter.

Extracts of the reconstructed slice obtained with different blur parameters are depicted in Figure 7.6 (d)-(i). From the additional zooms at the edges of the sample, the resolution properties can be evaluated. In (d) and (e), the slices obtained with $\sigma \in \{0.4, 0.5\}$ px show increasingly blurred edges for decreasing blur parameter. This indicates too small values for the blur parameter. The reconstruction with the lowest overall objective function value at $\sigma = 0.6\,\text{px}$ is depicted in (f). A sharp transition from tissue to air is reconstructed. For larger blur parameters, as depicted in (g)-(i), increasing overshoots at the edges are visible.

In conclusion, from visual inspection the best reconstruction is obtained for $\sigma = 0.6\,\text{px}$, as the reconstructed slices have sharp edges and no overshoots and thus the most realistic representation of the human wrist phantom. This blur parameter coincides with the optimum of the total objective function. Thus, it is possible to jointly estimate the blur parameter during reconstruction using Algorithm 7.2. In addition, it could be shown that this approach also works if the blur model is mismatched as the assumption of symmetric Gaussian blur does not match the physical system perfectly.

7.6. Discussion and conclusion

In this chapter an objective function was derived, which enables the joint estimation of parameters describing the system blur during tomographic reconstruction. A normalized sparsity measure for CT was devised, which captures the prior information about the sample. This measure favors sparse solutions and is less influenced by the noise level compared to conventional regularization approaches. An extensive simulation study was performed

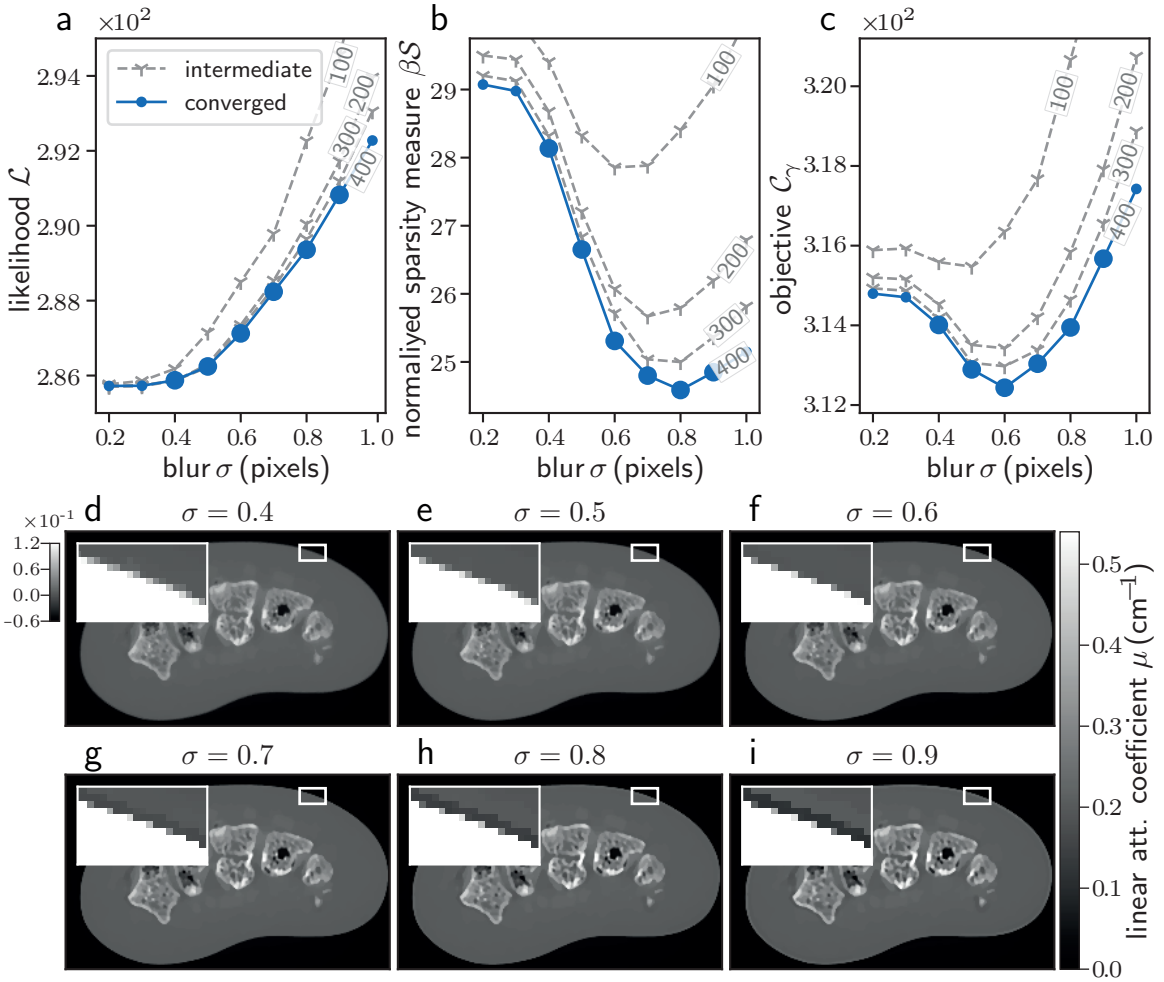

Figure 7.6.: Blind deconvolution reconstruction at an experimental test-bench setup. In (a)-(c), the different parts of the objective function are plotted over the blur parameter at different iterations. In gray, the values of the different parts of the objective function are plotted after 100, 200 and 300 iterations. The corresponding values after 400 iterations are depicted in blue. In (a) the values of the likelihood term are shown. In (b) the responses of the normalized sparsity measure are depicted. Finally, in (c) the values of the total objective function are plotted. In (d)-(i), extracts of the reconstructed slices obtained with different blur parameters are shown. In addition, zooms at the tissue to air interfaces are provided. To better asses the resolution properties, a more narrow window is used for these zooms. A similar version of the figure was published in [Hehn et al., 2019b].

showing that the objective function using conventional TV regularization would estimate the no-blur solution. By contrast, the proposed objective function has a global optimum at the true blur parameter. This holds true independently of the regularization strength and system blur. In addition, an optimization algorithm was proposed to jointly estimate low-dimensional blur parameters efficiently. Finally, this approach was successfully validated using an experimental test-bench setup.

The simulation study as well as the experimental test-bench study were carried out using a symmetric Gaussian system blur model, parameterized by a single blur parameter. To better approximate the system blur, more comprehensive models could be beneficial. For instance, a Lorentzian of Gaussian model could potentially better approximate detector blur, which however would require additional blur parameters. For parametric deconvolution reconstruction, no requirements on the dimensionality of the blur model were imposed and it is reasonable to assume that a global minimum at the desired blur parameters would be present. However, for additional blur parameters, the computational cost of blind deconvolution reconstruction increases significantly. For low-dimensional blur models, one could extend Algorithm 7.2 to update the blur parameters in an alternating fashion, while still only requiring three reconstructions in parallel.

Deconvolution algorithms are in general very sensitive to the noise level. As seen in Figure 7.2, for increasing noise levels, the performance of the normalized sparsity measure gets worse as the sample edges become increasingly more corrupted. One could devise a regularization term, which accounts differently for small differences in neighboring voxel values in order to be more sensitive to the actual sample edges. For conventional regularization techniques, this would relate to a Huber penalty given by Eq. (4.54) [Huber, 2011].

As seen for small blur parameters in Figure 7.4 (a) and Figure 7.6 (a), the likelihood term is biased in terms of blur. As proposed in [Krishnan et al., 2011], where the estimation of the blur kernel is performed only on the high frequencies of the photographs, one could modify the likelihood term to emphasize the high-frequency components. Such a behaviour would arise naturally from a likelihood term that includes noise covariances, as introduced in [Tilley et al., 2016a]. Thereby, the covariance matrix acts as a high-pass filter suppressing low and mid frequencies, where most energy is contained.

In summary, blind deconvolution has distinct benefits over conventional MBIR approaches. Firstly, time-consuming prior characterizations of the system blur can be omitted if a parameterized blur model is available. This may prove particularly beneficial if the system geometry or X-ray source parameters change frequently. Secondly, compared to prior characterizations, the blur is directly estimated from the reconstructed volume. Thus, one can also account for system blur due to algorithmic components, such as interpolations in the projection operations. Thirdly, although the above examinations were applied to conventional attenuation-based CT, the proposed approach is not restricted to this particular form of the physical mean model or noise model. Thus, additional parameters of the physical mean model that change the noise characteristics could be estimated in principle. For PB-CT, the model for the interference effects is parameterized by a single scalar, describing the monomorphous material. As a consequence, PB-CT would not require any prior knowledge about the materials and, as outlined in the Chapter 6, estimating a single parameter already implicitly accounts for the effects of the source, interference effects, detector and noise suppression. In general, high-resolution applications like CT in the nanometer regime might benefit particularly from the proposed approach, as higher resolution is fundamentally limited by the system blur.

8. Optimization-based geometry estimation for CT at the nanometer scale

Precise knowledge about the geometry, which is defined by the relative position of the source, rotation axis and detector, is crucial for image quality in CT as it defines the projection and backprojection operations. With increasing resolution, uncertainties in the geometry have an increasing influence on image quality and thus the geometry has to be known more accurately. This chapter addresses means to estimate the parameters which define the geometry, focusing on high-resolution cone-beam CT systems with sub-micron resolution.

Most geometry calibration methods require phantoms or other markers. In the following, two different approaches belonging to this class of methods are highlighted. An overview of the large amount of available literature is for instance presented in [Ferrucci et al., 2015].

The first approach uses a precisely defined calibration phantom enabling the estimation of geometry parameters for every view individually [Cho et al., 2005]. This phantom consists of 24 point markers (steel ball bearings). These markers are distributed on two parallel planes separated by a known distance along the rotation axis. The twelve markers at the respective planes are aligned evenly-spaced in a circle. For each view, the center of each individual marker is determined, which is sufficient to extract the corresponding geometry parameters. This method is highly sensitive to how accurately the calibration phantom is fabricated. Consequently, although this method is powerful and widely used in conventional non-microscopic systems, especially for medical imaging, for high-resolution microscopic systems accurate fabrication of such phantoms is difficult. In addition, the calibration scan has to be performed every time the geometry changes. Moreover, it is required that the geometry remains constant between measurements. Due to external disturbances such as thermal drifts, the geometry of high-resolution systems is not guaranteed to be constant over time. Therefore, discrepancies certainly occur between the geometry parameters determined in the phantom calibration and the geometry parameters of a subsequent sample scan.

Other approaches do not require known phantoms, but use markers with unknown positions [Gullberg et al., 1990, Noo et al., 2000, Von Smekal et al., 2004, Dittmann, 2018]. Instead, precise circular motion of the gantry (or rotation stage) is assumed. Thereby, one calculates the geometry parameters of the detector as well as the positions of each marker within the volume. Although most of the time these geometry parameters can only be computed globally (the respective geometric parameters are the same for every view) and, due to the assumption of precise circular motion, only a subset of all geometric parameters can be estimated, these approaches provide several advantages for microscopic systems. First of all, they do not require the fabrication of precise calibration phantoms. In addition, they do not rely on a separate calibration scan as the markers can be placed within the sample itself, not requiring a constant geometry between different scans. However, placing markers within the sample may not always be desirable.

This chapter focuses on calibration procedures that do not require a dedicated phantom or additional markers. These methods are also often referred to as on-line methods in literature [Muders and Hesser, 2014, Ferrucci et al., 2015]. Some commonly used methods

are mentioned below.

For instance, one can use the consistency conditions of opposing projections to determine the axis of rotation [Patel et al., 2009]. In addition, iterative alignment procedures exist that optimize the detector offset under every view [Mayo et al., 2007]. Here, one reconstructs the projections with some approximate geometry and simulates the corresponding projections afterwards. Then, the actual projections and simulated projections are registered to obtain the mismatch of the detector offset. An optimization-based approach for geometrical calibration has been proposed in [Panetta et al., 2008], where the objective function is based on the redundancy conditions of cone-beam data.

Other algorithms use some criterion to asses the image quality of the reconstruction itself. For instance, in [Donath et al., 2006] the center of rotation is determined using various different metrics. In [Kyriakou et al., 2008], the entropy of the reconstructed volume is minimized with respect to some geometry parameters using a simplex algorithm for multi-parameter optimization. A similar approach is proposed in [Kingston et al., 2011], where the ℓ_2 norm of the (spatial) gradient images is maximized with respect to the geometry parameters.

Little research revolves around jointly estimating the geometry parameters directly from the objective function used for iterative tomographic reconstruction techniques. The most likely explanation for this is that such an approach represents a heavily unconstrained optimization problem. This argument is also stated in [Bui et al., 2017], where deep-learning methods and CAT models are used to estimate the geometry parameters. There are however two conference proceedings [Rottman et al., 2015a, Rottman et al., 2015b], which describe, how geometry parameters can be estimated from the objective function in the context of mobile C-arms.

Within this context of geometry calibration methods, this chapter presents two on-line approaches for geometric calibration. The first approach falls into the category of optimizing a criterion that evaluates the image quality of the reconstruction directly. For this purpose, the normalized sparsity measure introduced in Chapter 7 is used. In theory, this metric gives advantages over commonly used metrics (such as the ℓ_2 norm of the gradient images) as it enforces sparsity. In contrast to most of the above-mentioned methods, here, the geometry parameters are estimated for every view individually. The second approach tries to optimize the geometry parameters jointly during tomographic reconstruction inside a statistical iterative reconstruction framework. The key contribution is a thorough derivation of the optimization algorithm used for jointly estimating the geometry parameters independently of the projector model. This approach is computationally efficient and no explicit calculations of the gradients of the projection matrix with respect to the geometry parameters are needed as in [Rottman et al., 2015a].

8.1. Projection matrices for two- and three-dimensional geometries

The projective geometry, which defines how voxel coordinates (x, y, z) of the object are mapped to the corresponding pixel coordinates (u, v) of the detector can be described by a so-called *projection matrix* $\mathbf{P}$ according to

$$\begin{pmatrix} uw \\ vw \\ w \end{pmatrix} = \mathbf{P} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}. \tag{8.1}$$

Consequently $\mathbf{P}$ is a 3×4 matrix for a three-dimensional geometry. To obtain the actual pixel coordinates, the corresponding entries of the resulting vector have to be normalized by w^{-1}. Such projection matrices are for instance discussed in detail in [Hartley and Zisserman, 2003] in the context of computer vision.

Countless different conventions to define the geometry in cone-beam CT exist. However, in [Commission, 2008] standardized conventions for circular cone-beam geometries are defined. These conventions are for instance adopted in a popular open-source CT reconstruction toolkit [Rit et al., 2014] and will also be used in the following. In this convention, the optical axis is the z-axis and the tomographic axis aligns with the y-axis. There are no particular units enforced on the distances. However, it must be ensured that the distances are consistent (have the same units).

8.1.1. Parallel- and fan-beam geometry

For two-dimensional geometries, the y coordinate is ignored. Consequently, the projection matrix is a 2×3 matrix. The conventions found in [Commission, 2008, Rit et al., 2014] for the cone-beam geometry are adapted in the following.

Parallel-beam geometry

First of all, the two-dimensional parallel-beam geometry is defined in (continuous) world coordinates, e.g. from continuous world coordinates of the volume to continuous projected world coordinates of the detector. Therefore, the geometry for each view is described by two parameters: the offset of the detector d_x (relative to different view) and the tomographic angle ϕ. The projection matrix can then be defined as

$$\mathbf{P}_{\text{parallel}}^{\text{world}} = \mathbf{M}\mathbf{R}, \tag{8.2}$$

with

$$\mathbf{M} = \begin{bmatrix} 1 & 0 & -d_x \\ 0 & 0 & 1 \end{bmatrix} \qquad \mathbf{R} = \begin{bmatrix} \cos\phi & -\sin\phi & 0 \\ \sin\phi & \cos\phi & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{8.3}$$

Fan-beam geometry

For the (two-dimensional) fan-beam geometry, there are three additional parameters required to describe the geometry under each view, namely the SAD d_{SAD}, the SDD d_{SDD} and the offset of the source s_x with respect to the optical axis. The corresponding projection matrix is then given by

$$\mathbf{P}_{\text{fan}}^{\text{world}} = \mathbf{M}_{\text{det}}\mathbf{M}_{\text{dst}}\mathbf{M}_{\text{src}}\mathbf{R}, \tag{8.4}$$

where the three new matrices are defined according to

$$\mathbf{M}_{\text{det}} = \begin{bmatrix} 1 & s_x - d_x \\ 0 & 1 \end{bmatrix} \qquad \mathbf{M}_{\text{dst}} = \begin{bmatrix} -d_{\text{SDD}} & 0 & 0 \\ 0 & 1 & -d_{\text{SAD}} \end{bmatrix} \qquad \mathbf{M}_{\text{src}} = \begin{bmatrix} 1 & 0 & -s_x \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{8.5}$$

Transformation to voxel and pixel coordinates

The final projection matrices map the discrete voxel coordinates represented by $(x, z, 1)^T$ to the (discrete) pixel coordinates represented by $(uw, w)^T$. Therefore, two additional affine transformations have to be defined. An affine projection $\mathbf{P}_{a,2\mathrm{d}}$ that converts projected world coordinates to pixel coordinates and an affine projection $\mathbf{V}_{a,2\mathrm{d}}^T$ that maps voxel coordinates to world coordinates are defined. With N_s being the number of pixels of size p_s on the detector and N_x the number of voxels of size v_x in one dimension, these matrices[1] can be defined as

$$\mathbf{P}_{a,2\mathrm{d}} = \begin{bmatrix} 1/p_s & (N_s-1)/2 \\ 0 & 1 \end{bmatrix} \qquad \mathbf{V}_{a,2\mathrm{d}}^T = \begin{bmatrix} v_x & 0 & -v_x(N_x-1)/2 \\ 0 & v_x & -v_x(N_x-1)/2 \\ 0 & 0 & 1 \end{bmatrix}. \tag{8.6}$$

Consequently, the final projection matrices for the parallel-beam and the fan-beam geometry are

$$\mathbf{P}_{\text{parallel}} = \mathbf{P}_{a,2\mathrm{d}}\mathbf{P}_{\text{parallel}}^{\text{world}}\mathbf{V}_{a,2\mathrm{d}}^T \qquad \text{and} \qquad \mathbf{P}_{\text{fan}} = \mathbf{P}_{a,2\mathrm{d}}\mathbf{P}_{\text{fan}}^{\text{world}}\mathbf{V}_{a,2\mathrm{d}}^T \tag{8.7}$$

respectively.

8.1.2. Cone-beam geometry

To define the (three-dimensional) cone-beam geometry, four additional parameters compared to the fan-beam geometry are required for each view. These include the rotation angle θ of the object around the optical axis (z-axis) and the rotation angle ψ that defines the tip in direction to the optical axis (around the x-axis). The last two parameters are s_y, the offset of the source in the orthogonal direction to s_x and d_y, the detector offset orthogonal to d_x. Following the conventions of [Commission, 2008, Rit et al., 2014], the projection matrix in world coordinates is given by

$$\mathbf{P}_{\text{cone}}^{\text{world}} = \mathbf{M}_{\text{det}}\mathbf{M}_{\text{dst}}\mathbf{M}_{\text{src}}\mathbf{R}_\theta\mathbf{R}_\psi\mathbf{R}_\phi \tag{8.8}$$

with

$$\mathbf{M}_{\text{det}} = \begin{bmatrix} 1 & 0 & s_x - d_x \\ 0 & 1 & s_y - d_y \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{M}_{\text{src}} = \begin{bmatrix} 1 & 0 & 0 & -s_x \\ 0 & 1 & 0 & -s_y \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{8.9}$$

$$\mathbf{M}_{\text{dst}} = \begin{bmatrix} -d_{\text{SDD}} & 0 & 0 & 0 \\ 0 & -d_{\text{SDD}} & 0 & 0 \\ 0 & 0 & 1 & -d_{\text{SAD}} \end{bmatrix} \tag{8.10}$$

and the modified rotation matrices[2]

$$\mathbf{R}_\theta = \begin{bmatrix} \cos\theta & \sin\theta & 0 & 0 \\ -\sin\theta & \cos\theta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{R}_\psi = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\psi & \sin\psi & 0 \\ 0 & -\sin\psi & \cos\psi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{8.11}$$

[1] More general definitions which do not require squared voxels are possible.

[2] For the existing conventions used in our software, the sign on the distances and angles have to be inverted. In addition, the rotation matrices are defined differently. The above convention for the rotations is referred to as the ZXY convention which is the natural choice for circular cone-beam geometries, where previous work at our chair used the YZY convention suitable for Euler cradles [Als-Nielsen and McMorrow, 2011, Schaff, 2017].

$$\mathbf{R}_\phi = \begin{bmatrix} \cos\phi & 0 & -\sin\phi & 0 \\ 0 & 1 & 0 & 0 \\ \sin\phi & 0 & \cos\phi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{8.12}$$

To obtain the final projection matrix, the affine transformation matrices are extended to three dimensions with N_t being the number of detector rows, p_t the pixel height, N_y the number of voxels parallel to the rotation axis with voxel size v_y, according to

$$\mathbf{P}_a = \begin{bmatrix} 1/p_s & 0 & (N_s-1)/2 \\ 0 & 1/p_t & (N_t-1)/2 \\ 0 & 0 & 1 \end{bmatrix} \qquad \mathbf{V}_a^T = \begin{bmatrix} v_x & 0 & 0 & -v_x(N_x-1)/2 \\ 0 & v_y & 0 & -v_y(N_y-1)/2 \\ 0 & 0 & v_x & -v_x(N_x-1)/2 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{8.13}$$

resulting in

$$\mathbf{P}_{\text{cone}} = \mathbf{P}_a \mathbf{P}_{\text{cone}}^{\text{world}} \mathbf{V}_a^T. \tag{8.14}$$

8.2. Metric-guided geometry optimization

In metric-guided geometry estimation, the geometry parameters are altered and the quality of the reconstruction is evaluated by a certain criterion. If the quality of the reconstruction improves by a given change of geometry parameters, the old parameters are updated by the new ones. This approach consists of two crucial components: the criterion that evaluates the quality of the reconstruction and the procedure by which the geometry parameters are altered.

The following algorithms only require the filtered projections $\tilde{\mathbf{y}}$, obtained by Eq. (4.27). Given the set of geometry parameters $\boldsymbol{\theta}$ that define the geometry, the analytic reconstruction can be denoted by

$$\hat{\mathbf{x}} = \mathbf{A}^T(\boldsymbol{\theta})\tilde{\mathbf{y}}, \tag{8.15}$$

using backprojection according to Eq. (4.28)[3]. Here $\boldsymbol{\theta} = [\boldsymbol{\theta}^0, \boldsymbol{\theta}^1, \dots]$ is a matrix that holds the geometry parameters for every view, e.g. $\boldsymbol{\theta} = [\mathbf{d}_{\text{SAD}}, \mathbf{d}_{\text{SDD}}, \dots]$ for a cone-beam geometry and $\mathbf{d}_{\text{SAD}} = (d_{\text{SAD},0}, d_{\text{SAD},1}, \dots)^T$ contains the respective geometry parameter for every view. As the criterion for the quality of the reconstruction, the normalized sparsity metric for CT developed in Chapter 7 is used, which is given by Eq. (7.3) according to

$$\mathcal{S}(\mathbf{x}) \equiv \frac{\mathcal{R}_{\ell_1}(\mathbf{x})}{\sqrt{\mathcal{R}_{\ell_2^2}(\mathbf{x})}}. \tag{8.16}$$

In this case, a lower value of this metric corresponds to higher reconstruction quality. Formally, an objective function

$$\boxed{\mathcal{C}(\boldsymbol{\theta}) = \mathcal{S}(\mathbf{A}^T(\boldsymbol{\theta})\tilde{\mathbf{y}})} \tag{8.17}$$

can be defined and optimized according to

$$\boxed{\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} \mathcal{C}(\boldsymbol{\theta}),} \tag{8.18}$$

[3] This notation has subtle differences to the actual cone-beam reconstruction techniques. For instance, for the FDK algorithm, this equation refers to the backprojection step, which is not the same as the transpose of the projection operation as it includes additional weights.

resulting in an estimate of geometry parameters $\hat{\boldsymbol{\theta}}$ with the lowest value of the normalized sparsity measure. In practice, however, solving this optimization problem directly is not feasible. No obvious gradient information is available and manually searching this high-dimensional objective is computationally very demanding, as nine geometry parameters (for the cone-beam geometry) are needed for every view.

However, for estimating a single geometry parameter (for instance, the global offset of the detector in x-direction) this approach is reduced to a one-dimensional minimization of

$$f(\theta) = \mathcal{S}(\mathbf{A}^T(\bar{\boldsymbol{\theta}} + \theta\bar{\mathbf{e}}_m)\tilde{\mathbf{y}}). \tag{8.19}$$

Here, $\bar{\boldsymbol{\theta}} = (\bar{\theta}_0, \bar{\theta}_1, \dots)$ holds the different global geometry parameters (which are the same for every view) and $\bar{\mathbf{e}}_m$ is a vector with zeros for all entries of $\bar{\boldsymbol{\theta}}$ except the m-th global geometry parameter, where it is one. A horizontal line is put over parameters to indicate that they are the same for every view. This objective is then minimized according to

$$\hat{\theta} = \arg\min_{\theta} f(\theta), \tag{8.20}$$

which can be solved by various line-search algorithms that do not require the first order derivatives, such as Brent's method [Brent, 1971].

This approach can be extended to estimate multiple global geometry parameters by applying Eq. (8.20) sequentially for different parameters as illustrated in Algorithm 8.1. The algorithm requires the filtered projections $\tilde{\mathbf{y}}$, an initialization of the global geometry parameters $\bar{\boldsymbol{\theta}}$ and a set M specifying which parameters to optimize for. The optimization approach given by Eq. (8.20) is then solved in line 2 for each geometry parameter respectively. Finally, the reconstruction $\mathbf{x}$ and the updated geometry parameters $\bar{\boldsymbol{\theta}}$ are returned.

Algorithm 8.1 Algorithm to optimize global geometry parameters. It requires the filtered projection measurements $\tilde{\mathbf{y}}$, an initial guess of the global geometry parameters $\bar{\boldsymbol{\theta}} = (\bar{\theta}^0, \bar{\theta}^1, ...)$ and a set M specifying which parameters to optimize for. The algorithm returns the reconstructed volume $\mathbf{x}$ and the optimized global geometry parameters $\bar{\boldsymbol{\theta}}$.

Require: $\tilde{\mathbf{y}}, \bar{\boldsymbol{\theta}}, M$
1: **for** $m \in M$ **do**
2: $\quad\hat{\theta} \leftarrow \arg\min_\theta \mathcal{S}(\mathbf{A}^T(\bar{\boldsymbol{\theta}} + \theta\bar{\mathbf{e}}_m)\tilde{\mathbf{y}})$
3: $\quad\bar{\theta}_m \leftarrow \bar{\theta}_m + \hat{\theta}$
4: **end for**
5: $\mathbf{x} = \mathbf{A}^T(\bar{\boldsymbol{\theta}})\tilde{\mathbf{y}}$
6: **return** $\mathbf{x}, \bar{\boldsymbol{\theta}}$

Moreover, this approach can even be extended to optimizing the geometry parameters for every view individually, if the increase in computational time by the number of views is acceptable. This, however, requires the metric to be sensitive enough to capture slight changes in reconstruction quality if only a single geometry parameter of an individual view is slightly altered. Algorithm 8.2 illustrates this approach. The algorithm again requires the filtered projections $\tilde{\mathbf{y}}$, an initial (reasonable) guess of geometry parameters $\boldsymbol{\theta}$ and the set M specifying which parameters are to be updated. Then, for every geometry parameter specified by M and for every of the P projection views, a one dimensional objective function is solved in line 3. In contrast to the previous algorithm and Eq. (8.19), only one parameter for an individual view is altered. Again, this can be solved by using a gradient-free line-search algorithm. Afterwards, the correction of this particular geometry parameter is used

to refine the geometry. After this procedure is repeated for all projection views and all desired geometry parameters, the reconstruction $\mathbf{x}$ is computed and returned with the refined geometry parameters $\boldsymbol{\theta}$.

Algorithm 8.2 Algorithm to optimize geometry parameters for every view. It requires the filtered projection measurements $\tilde{\mathbf{y}}$, the initial guess of all geometry parameters at every view $\boldsymbol{\theta} = (\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, ...)$ and a set M specifying which parameters to include in the optimization. As a result, the reconstructed volume $\mathbf{x}$ and the optimized geometry parameters for every view $\boldsymbol{\theta}$ are returned.

Require: $\tilde{\mathbf{y}}, \boldsymbol{\theta}$
for $m \in M$ **do**
  for $p \leftarrow 0..P-1$ **do**
    $\hat{\theta} \leftarrow \arg\min_\theta \mathcal{S}(\mathbf{A}^T(\boldsymbol{\theta} + \theta \mathbf{e}_m^p)\tilde{\mathbf{y}})$
    $\theta_m^p \leftarrow \theta_m^p + \hat{\theta}$
  end for
end for
$\mathbf{x} = \mathbf{A}^T(\boldsymbol{\theta})\tilde{\mathbf{y}}$
return $\mathbf{x}, \boldsymbol{\theta}$

The advantages of the approaches introduced above are that they are not restricted to a particular set of geometry parameters. In addition, one can utilize different metrics to capture different aspects of the reconstruction quality. However, these approaches (especially for per-view optimization) are computationally very expensive and not guaranteed to converge globally. In practice, there are two strategies one could apply to reduce computational cost. Firstly, one could include only a subset of critical parameters that has the greatest impact on image quality. Secondly, the implementation of the geometry parameters in world coordinates allows to apply the optimization to a subsampled representation of the data. The resulting parameters can then be used for a reconstruction of the full resolution without any modifications. In addition, these methods are restricted to geometries where fast (analytical) reconstruction algorithms are available, such as the FDK algorithm for circular source trajectories. Furthermore, intermediate calculations within the backprojection, which are independent of the geometry parameter that is currently optimized for, can be reused.

8.3. Application of metric-guided geometry estimation at the NanoCT

In this section, metric-guided geometry estimation employing the normalized sparsity measure was applied at an experimental setup, which achieves submicron resolution using geometric magnification. At this resolution, the geometry is prone to misalignment, for instance, due to thermal drifts or vibrations. Since these misalignments cannot be characterized beforehand, on-line methods to correctly estimate the geometry are desired.

Parts of the following results are prepared for submission to an international peer-reviewed journal [Ferstl et al., 2019].

8.3.1. The NanoCT

The proposed method was validated at the NanoCT. This experimental setup consists of a prototype nanofocus source, a photon-counting detector and a high-precision rotary stage.

The prototype source (Excillum AB, Sweden) consists of a thin tungsten transmission target on a diamond layer with integrated water cooling enabling X-ray spot sizes down to 200 – 300 nm FWHM. The X-ray detector is a PILATUS 300K-W (DECTRIS Ltd., Switzerland) with a 1 mm thick silicon sensor with a pixelsize of 172 µm. Due to the direct conversion of X-ray photons, the point spread function (PSF) of the system can be neglected. In addition, by only counting single X-ray photons above a certain energy threshold, no readout noise is present. The sample stage is mounted as an overhead construction with an air-bearing rotary stage for sample rotation. [Müller et al., 2017]

For the following measurement, the SAD was set to $d_{\mathrm{SAD}} = 1.03\,\mathrm{mm}$ and the SDD to 600 mm. This results in a magnification of around $M = 583$, and thus in an effective voxel size of $v_x = v_y = 295\,\mathrm{nm}$. The detector was shifted after every projection to better account for gaps between different detector modules. [Ferstl et al., 2019]

8.3.2. Sample and preprocessing

The sample is a sea cucumber (Leptosynapta Minuta), which was stained with osmium tetroxide and embedded in epoxy resin. [Ferstl et al., 2019]

Before tomographic reconstruction, a sequence of preprocessing steps was performed. First of all, dead and corrupted detector pixels were replaced by the median of the neighboring pixel values. Afterwards, from a series of reference projections without the sample, the inhomogeneous illumination as well as the individual pixel response were corrected (details can be found in [Allner, 2019]). After correcting for the shift of the detector under each view, the missing data from the module gaps of the detector were interpolated using neighboring projection images. Subsequently, an intensity ramp was fitted onto the projections to alleviate artifacts originating from the fact that the sample is significantly larger than the FOV. Finally, qualitative single-material phase-retrieval is performed according to Eq. (3.18) with $\xi = 5$.

8.3.3. Alignment and reconstruction

For geometry optimization, the horizontal detector offset (simply referred to as detector offset in the following) was estimated. This parameter has a crucial impact on the image quality of the reconstruction.

Firstly, to reduce computational cost, the projections were subsampled by a factor of three in the angular dimension and by a factor of two in each of the spatial dimensions by calculating the respective mean values. The number of data points was thus reduced by a factor of twelve.

Next, the global detector offset was computed using Algorithm 8.1 with the (filtered) downsampled projections. According to visual inspection, the initial guess of the detector offset was set to 1 mm. The other geometry parameters were assumed to be correct. For the optimization in line 2 of the algorithm, Brent linesearch was used with the initial bracket of $(1\,\mathrm{mm}, 1.5\,\mathrm{mm})$ until the respective updates were smaller than $10^{-3}\,\mathrm{mm}$. These values were chosen by hand. The value for the estimated global detector offset was $\bar{d}_x^{\mathrm{ds}} \approx 1.19\,\mathrm{mm}$. The superscript indicates that this estimate was obtained from the downsampled projections. A reconstructed slice using this estimated detector offset can be seen in Figure 8.1 (a).

To estimate the detector offset for each view individually, Algorithm 8.2 was used. Again, the downsampled (filtered) projections were used. The detector offset for all projections was initialized with the global value obtained by Algorithm 8.1 as described above. For

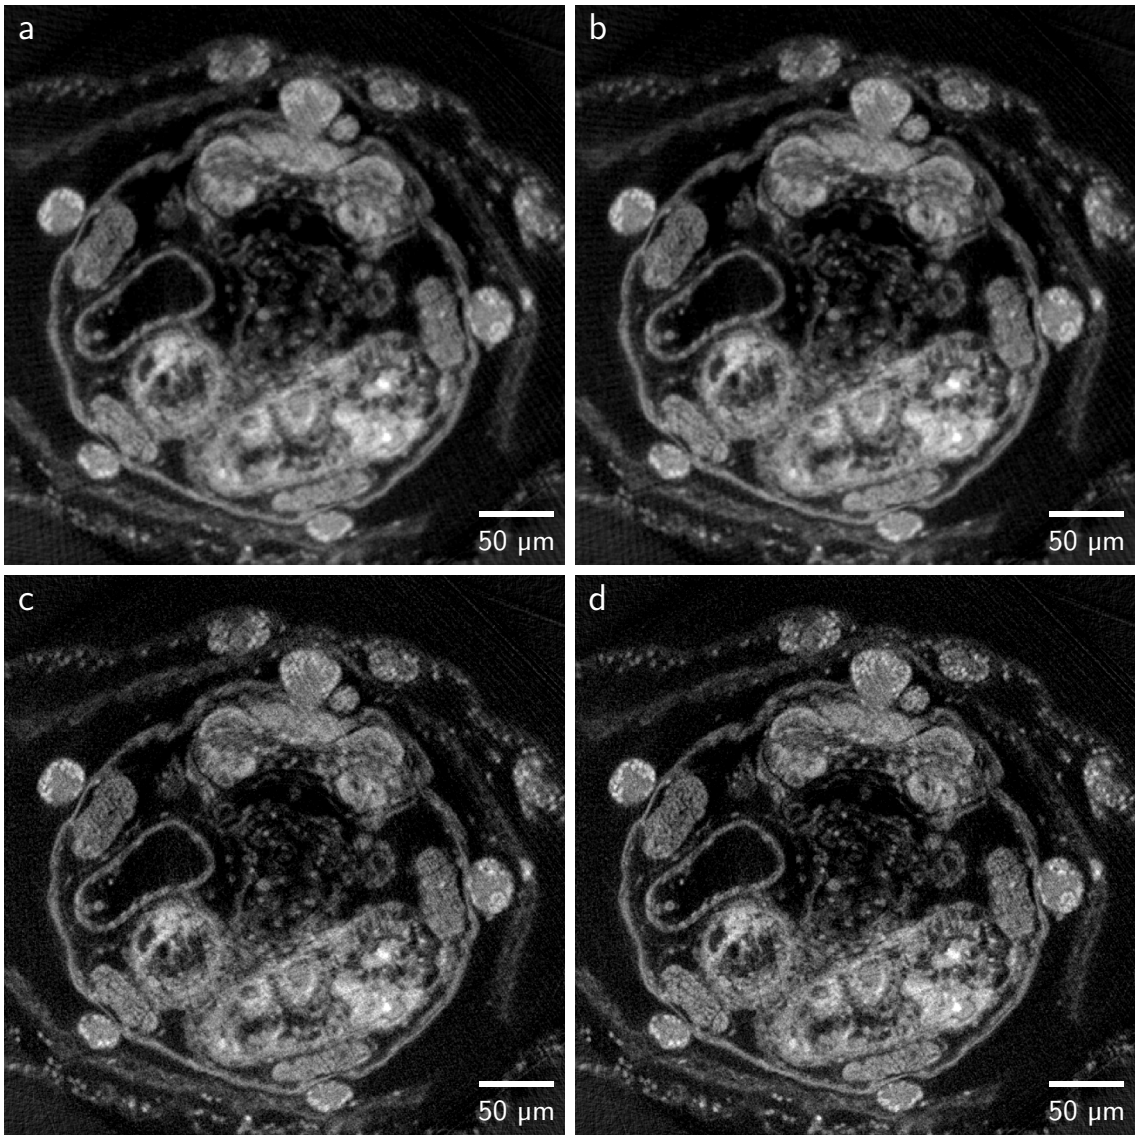

Figure 8.1.: Reconstructions of a sea cucumber specimen measured at the NanoCT. An extract of the central slice of the specimen is depicted using different sampling of the projections and different geometry parameters. In (a) and (b) the reconstructions are depicted in the same linear gray scale using the undersampled phase-retrieved projections. In (a) only the global horizontal detector offset has been optimized. In (b) the horizontal detector offset has been optimized for every view. The reconstruction using the full resolution of the projections preprocessed with weaker phase-retrieval can be seen in (c) and (d) using the same linear gray scale (but a different gray scale than above). In (c) the same global detector offset is used that has been estimated for (a). By contrast, in (d) the horizontal detector shifts for every view are linearly interpolated (and extrapolated) from the horizontal detector shifts estimated in (b) as shown in Figure 8.2. A manuscript is under preparation, where a similar axial slice of the sample is presented. [Ferstl et al., 2019]

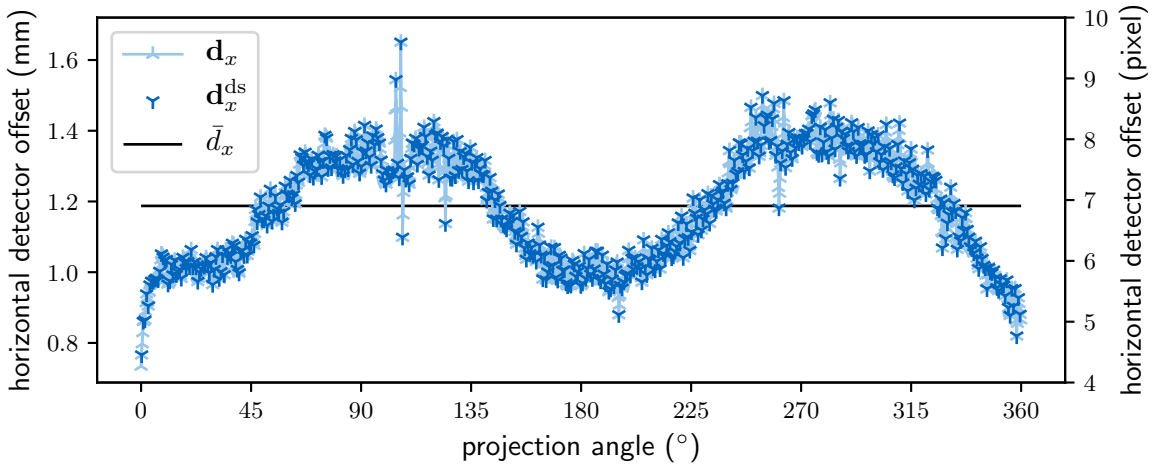

Figure 8.2.: Estimated horizontal detector offsets. The global detector offset, which has the same value for the subsampled as well as the original geometry, is shown by the black horizontal line. In (dark) blue the detector offsets for every view as estimated from the subsampled projections are shown. The respective detector offsets used for the original geometry, which are obtained by linear interpolation (and extrapolation) of the subsampled offsets, are depicted in a lighter color.

optimization in line 4, Brent linesearch was used again with a bracket of $(-0.5\,\mathrm{mm}, 0.5\,\mathrm{mm})$. The stopping criterion is fulfilled if the updates were smaller than $10^{-2}\,\mathrm{mm}$ or the number of iterations exceeds 10 (which however was very uncommon). In Figure 8.1 (b) the resulting reconstructed slice is depicted. Compared to the previous reconstruction obtained with the global detector offset, the resolution is improved and features are more clearly visible. The estimates for every view referred to as $\mathbf{d}_x^{\mathrm{ds}}$ are shown in Figure 8.2 together with the global estimate obtained before. As expected, the global detector offset is approximately the mean of the detector offsets obtained for every view.

As all geometry parameters are given in (continuous) world units and are thus independent of the particular pixel or voxel sizes, the values of the geometry parameters can be directly used for the original projections with the full resolution. Thereby, only the affine matrices that describe the conversion from voxel to world coordinates and from world coordinates to pixel coordinates have to be updated. To further improve the resolution, a weaker single-material phase-retrieval with $\xi = 3$ was applied. In Figure 8.1 (c) the reconstruction with the global estimate of the detector offset $\bar{d}_x$ is shown.

Finally, to transfer the estimates of the individual detector offsets $\mathbf{d}_x^{\mathrm{ds}}$ to the original projections, referred to as $\mathbf{d}_x$, an interpolation has to be performed as three times as many projection views are available. As the individual detector offsets vary slowly between neighboring projection views, as indicated in Figure 8.2, linear interpolation can be used to estimate the individual detector offsets for the original projection. In Figure 8.2 the interpolated detector offsets and the undersampled offsets are shown. The final reconstruction is shown in Figure 8.1 (d) improving the results of Figure 8.1 (c) further.

In conclusion, the normalized sparsity measure for CT can be used to refine the horizontal detector offset for each view. Although a comparison with other methods is beyond the scope of this thesis, this method seems to be be more robust than iterative alignment procedures that correlate the measurements to virtual projections [Mayo et al., 2007], in particular,

if the sample exceeds the FOV and features only weak contrast. Most importantly, the proposed approach can, in principle, be used to estimate other parameters, such as the alignment of the rotation axis or the distances between the source, axis and detector.

8.4. Joint geometry estimation and tomographic reconstruction

In joint geometry estimation and tomographic reconstruction, the geometry parameters are not assessed according to some criterion applied to the reconstruction, but directly to the objective function (posterior distribution) used for tomographic reconstruction. Here, the objective function $\mathcal{C}$ is not only understood as function of the sample $\mathbf{x}$, but also of the geometry parameters $\boldsymbol{\theta}$, which parameterize the forward and backprojection operations. Formally, this can be written as

$$\left\{\hat{\mathbf{x}}, \hat{\boldsymbol{\theta}}\right\} = \arg\min_{\mathbf{x},\boldsymbol{\theta}} \mathcal{C}(\mathbf{x}, \boldsymbol{\theta}). \tag{8.21}$$

In the following, a strategy is derived to optimize Eq. (8.21), which is independent of the projector model. This implies that the following approach can be used with any existing implementation of the projection operations.

8.4.1. Objective function and derivation of the gradients

In the following, only the projection operations are modeled in the mean model according to Eq. (4.46). Thus, the matrix $\mathbf{A}$, defining the forward and backprojection operations can be referred to as the system matrix. The quantity from which to reconstruct the volume $\mathbf{x}$ are the line-integrals denoted by $\boldsymbol{\ell}$. The following notation is used to explicitly differentiate between different views and different pixels. Thereby, ℓ_i^p refers to the line-integrals of the i-th pixel and the p-th view. The entries of the system matrix are referred to as a_{ij}^p, where j refers to the voxels. In addition, $\boldsymbol{\theta}$ holds the geometry parameters for every view. The entries are indexed according to θ_m^q, where q refers to the views and m runs over different geometry parameters. For simplicity, the likelihood is written as a least squares objective and additional regularization terms are omitted. Thus, the objective function equals the likelihood and can be written as

$$\boxed{\mathcal{L}(\mathbf{x}, \boldsymbol{\theta}) = \frac{1}{2} \sum_p \sum_i \left(\ell_i^p - \sum_j a_{ij}^p(\boldsymbol{\theta}) x_j \right)^2,} \tag{8.22}$$

where the dependencies are explicitly denoted. This formulation coincides with Eq. (4.49) omitting the statistical weights for convenience to keep the notation short. The gradient with respect to $\mathbf{x}$ indexed by k is given by

$$\frac{\partial \mathcal{L}(\mathbf{x}, \boldsymbol{\theta})}{\partial x_k} = -\sum_p \sum_i a_{ik}^p(\boldsymbol{\theta}) \left(\ell_i^p - \sum_j a_{ij}^p(\boldsymbol{\theta}) x_j \right). \tag{8.23}$$

To jointly estimate $\mathbf{x}$ and $\boldsymbol{\theta}$ the objective function also has to be derived with respect to $\boldsymbol{\theta}$. This results in

$$\frac{\partial \mathcal{L}(\mathbf{x}, \boldsymbol{\theta})}{\partial \theta_m^q} = -\sum_p \sum_i \left(\sum_k \frac{\partial a_{ik}^p(\boldsymbol{\theta})}{\partial \theta_m^q} x_k \right) \left(\ell_i^p - \sum_j a_{ij}^p(\boldsymbol{\theta}) x_j \right), \tag{8.24}$$

where the gradient of the system matrix $\mathbf{A}$ with respect to the geometry parameters $\boldsymbol{\theta}$ has to be computed. This gradient depends on the individual projector model. An explicit derivation is not straightforward as the projector models can become complex. Therefore, to alleviate this problem, this gradient is evaluated as a finite difference according to

$$\sum_k \frac{\partial a_{ik}^p(\boldsymbol{\theta})}{\partial \theta_m^q} x_k \approx \frac{1}{\epsilon}\left(\sum_k a_{ik}^p(\boldsymbol{\theta} + \epsilon \mathbf{e}_m^q) x_k - \sum_k a_{ik}^p(\boldsymbol{\theta}) x_k\right), \tag{8.25}$$

where $\mathbf{e}_q^m$ denotes a matrix, with all entries being zero except for the element of the q-th view of the m-th geometry parameter and $\epsilon > 0$ being sufficiently small. However, the first sum requires an individual forwardprojection for every view and geometry parameter, making this approach unfeasible due to its computational cost.

To reduce computational cost, the argument of the system matrix needs to become independent of the respective projection view. This can be achieved by employing the fact that a change of the m-th geometry parameter of the q-th view only influences the q-th view. Mathematically, this can be written as

$$\sum_k a_{ik}^p(\boldsymbol{\theta} + \epsilon \mathbf{e}_m^q) x_k = \delta^{pq} \sum_k a_{ik}^p(\boldsymbol{\theta} + \epsilon \bar{\mathbf{e}}_m) x_k + (1 - \delta^{pq}) \sum_k a_{ik}^p(\boldsymbol{\theta}) x_k. \tag{8.26}$$

Here $\bar{\mathbf{e}}_m$ denotes a matrix, where all entries are 0 except for those of the m-th projection, where the entries are one. In addition, δ^{pq} denotes the Kronecker delta which is zero everywhere except for $p = q$, where it is one. Inserting Eq. (8.26) into Eq. (8.25) results in

$$\sum_k \frac{\partial a_{ik}^p(\boldsymbol{\theta})}{\partial \theta_m^q} x_k \approx \frac{\delta^{pq}}{\epsilon}\left(\sum_k a_{ik}^p(\boldsymbol{\theta} + \epsilon \bar{\mathbf{e}}_m) x_k - \sum_k a_{ik}^p(\boldsymbol{\theta}) x_k\right). \tag{8.27}$$

Thereby, the number of required forwardprojections is reduced by the number of views. Finally, the gradient given by Eq. (8.24) reduces with Eq. (8.27) to

$$\boxed{\frac{\partial \mathcal{L}(\mathbf{x}, \boldsymbol{\theta})}{\partial \theta_m^q} = -\frac{1}{\epsilon} \sum_i \left(\sum_k a_{ik}^q(\boldsymbol{\theta} + \epsilon \bar{\mathbf{e}}_m) x_k - \sum_k a_{ik}^q(\boldsymbol{\theta}) x_k\right)\left(\ell_i^q - \sum_j a_{ij}^q(\boldsymbol{\theta}) x_j\right),} \tag{8.28}$$

where δ^{pq} cancels the summation over p and one can simultaneously calculate the updates of a particular geometry parameter for all projection angles q using only one additional forwardprojection, as the second forwardprojection is already computed in the evaluation of the gradient with respect to the volume given by Eq. (8.23).

8.4.2. First-order optimization using ADAM

For minimizing Eq. (8.22), using the steepest descent search direction given by Eq. (4.60) with constant step length gave poor performance. Thus, Adaptive Momentum (ADAM) estimation was chosen, which includes adaptive step lengths (or learning rates) based on the estimates of the first and second moments of the gradients, as defined in Algorithm 8.3. The ADAM solver was introduced for stochastic optimization problems and is nowadays widely used in the field of machine learning. This solver might not be particularly well suited for the proposed problem, but it is computationally efficient, has little memory requirements and only requires first-order derivatives [Kingma and Ba, 2014]. For instance, one could make use of the fact that minimizing the likelihood with respect to x is a linear problem or one

Algorithm 8.3 Algorithm for joint geometry estimation and reconstruction. The number of iterations is given by N and the set M indicates the geometry parameters to optimize for. The Hadamard product is denoted by $\odot$. The algorithm requires the line-integrals $\boldsymbol{\ell}$, initial guesses for $\mathbf{x}^{(0)}$ and $\boldsymbol{\theta}^{(0)} = (\boldsymbol{\theta}_0^{(0)}..\boldsymbol{\theta}_{M-1}^{(0)})$ as well as the step lengths α_x and $\boldsymbol{\alpha} = (\alpha_0..\alpha_{M-1})$. The algorithm returns the reconstructed object $\mathbf{x}^{(N)}$ and the optimized geometry parameters $\boldsymbol{\theta}^{(N)}$.

Require: $\boldsymbol{\ell}, \mathbf{x}^{(0)}, \boldsymbol{\theta}^{(0)}, \alpha_x, \boldsymbol{\alpha}, M, N$
$\mathbf{m}_x^{(0)} = \mathbf{v}_x^{(0)} = \mathbf{0}$
for $m \in M$ **do**
  $\mathbf{m}_m^{(0)} = \mathbf{v}_m^{(0)} = \mathbf{0}$
end for
$\beta_1 = 0.9$; $\beta_2 = 0.999$; $\epsilon = 10^{-8}$; $\epsilon' = 10^{-10}$
for $n \leftarrow 0..N-1$ **do**
  $\mathbf{p} = \mathbf{A}(\boldsymbol{\theta}^{(n)})\mathbf{x}^{(n)}$
  $\mathbf{r} \leftarrow \boldsymbol{\ell} - \mathbf{p}$
  $\mathbf{D}_x \leftarrow -\mathbf{A}^T(\boldsymbol{\theta}^{(n)})\mathbf{r}^{(n)}$
  $\mathbf{m}_x^{(n+1)} \leftarrow \beta_1 \mathbf{m}_x^{(n)} + (1-\beta_1)\mathbf{D}_x$
  $\mathbf{v}_x^{(n+1)} \leftarrow \beta_2 \mathbf{v}_x^{(n)} + (1-\beta_2)\mathbf{D}_x \odot \mathbf{D}_x$
  $\hat{\mathbf{m}}_x \leftarrow \mathbf{m}_x^{(n+1)}/(1-\beta_1^{n+1})$
  $\hat{\mathbf{v}}_x \leftarrow \mathbf{v}_x^{(n+1)}/(1-\beta_2^{n+1})$
  for $m \in M$ **do**
    $\mathbf{D}_m \leftarrow -1/\epsilon' \left(\mathbf{A}(\boldsymbol{\theta}^{(n)} + \epsilon'\bar{\mathbf{e}}_m)\mathbf{x}^{(n)} - \mathbf{p}\right)^T \mathbf{r}$
    $\mathbf{m}_m^{(n+1)} \leftarrow \beta_1 \mathbf{m}_m^{(n)} + (1-\beta_1)\mathbf{D}_m$
    $\mathbf{v}_m^{(n+1)} \leftarrow \beta_2 \mathbf{v}_m^{(n)} + (1-\beta_2)\mathbf{D}_m \odot \mathbf{D}_m$
    $\hat{\mathbf{m}}_m \leftarrow \mathbf{m}_m^{(n+1)}/(1-\beta_1^{n+1})$
    $\hat{\mathbf{v}}_m \leftarrow \mathbf{v}_m^{(n+1)}/(1-\beta_2^{n+1})$
  end for
  $\mathbf{x}^{(n+1)} \leftarrow \mathbf{x}^{(n)} - \alpha_x \hat{\mathbf{m}}_x / \left(\sqrt{\hat{\mathbf{v}}_x} + \epsilon\right)$
  for $m \in M$ **do**
    $\boldsymbol{\theta}_m^{(n+1)} \leftarrow \boldsymbol{\theta}_m^{(n)} - \alpha_m \hat{\mathbf{m}}_m / \left(\sqrt{\hat{\mathbf{v}}_m} + \epsilon\right)$
  end for
end for
return $\mathbf{x}^{(N)}, \boldsymbol{\theta}^{(N)}$

could use ordered subset algorithms developed for transmission tomography applications as discussed in Section 4.3.

The optimization scheme is sketched in Algorithm 8.3. The line-integrals $\boldsymbol{\ell}$ as well as the initial guesses $\mathbf{x}^{(0)}$ and $\boldsymbol{\theta}^{(0)}$ are required. In addition, one has to choose the step length manually for the update in $\mathbf{x}$, referred to as α_x and the step lengths $\boldsymbol{\alpha}$ for each geometry parameter one chooses to optimize for. Which parameters are estimated is again specified by the set M. In addition, a termination criterion such as the number of iterations N has to be stated. The first momentum vectors $\mathbf{m}$ and second momentum vectors $\mathbf{v}$ are all initialized with zeros. The parameters β_1, β_2 and α are chosen to be the default parameters according to [Kingma and Ba, 2014]. In addition, ϵ', which is used for the finite difference approximation is set depending on the machine precision. In every iteration indexed by n the virtual projections $\mathbf{p}$ are calculated from the current estimate of the object $\mathbf{x}^{(n)}$ and geometry parameters $\boldsymbol{\theta}^{(n)}$. This variable is then stored. Subsequently, the residuum $\mathbf{r} = \boldsymbol{\ell} - \mathbf{p}$ is calculated and also stored. Afterwards, the gradient of the objective with respect to $\mathbf{x}$ is calculated by backprojecting the residuum. The first and second moments $\mathbf{m}_x$ and $\mathbf{v}_x$ are then calculated. To compensate the bias from the initialization of the moments, the moments are corrected depending on the iteration resulting in $\hat{\mathbf{m}}_x$ and $\hat{\mathbf{v}}_x$. Afterwards, for every geometry parameter $\boldsymbol{\theta}_m$, the respective gradient is calculated for all projection angles simultaneously by computing one additional forwardprojection and using the previously stored projections $\mathbf{p}$ and residuum $\mathbf{r}$. Then, the respective moments $\mathbf{m}_m, \mathbf{v}_m, \hat{\mathbf{m}}_m, \hat{\mathbf{v}}_m$ are computed. Finally, the new estimates $\mathbf{x}^{(n+1)}$ and $\boldsymbol{\theta}^{(n+1)}$ are calculated using the corrected first and second raw moments. After N iterations, the estimates for $\mathbf{x}^{(N)}$ and $\boldsymbol{\theta}^{(N)}$ are returned.

In addition, one could include further improvements such as leveraging the fact that ADAM is built upon stochastical optimization approaches. Thus, including only an (ordered) subset (sometimes also called minibatch) of projections in each iteration might lead to faster convergence. However, for the update of the geometry parameters this would also result in only an update of a subset of these parameters.

8.5. Evaluation of joint geometry estimation using a simulation study

In the following, a simulation study is performed to evaluate the performance and limitations of jointly estimating the geometry parameters during tomographic reconstruction using Algorithm 8.3. This approach is evaluated for a two-dimensional parallel-beam geometry and a fan-beam geometry using the Shepp-Logan phantom [Shepp and Logan, 1974b], which was downsampled to 64×64 px to reduce partial volume effects and to reduce computational cost. Forward and backprojection are matched using the footprint projector model, described in Figure 4.3 (c)-(d). The simulations were performed without noise.

8.5.1. Two-dimensional parallel-beam geometry

In a two-dimensional parallel-beam arrangement the geometry of the tomographic system is described by the offset of the detector $\mathbf{d}_x$ orthogonal to the beam axis and the tomographic angles $\boldsymbol{\phi}$. For the uncorrupted (reference) geometry, no detector offset $\mathbf{d}_x = \mathbf{0}$ was assumed. The number of tomographic angles was set to 101 according to Nyquist's criterion given by Eq. (4.18) and distributed equidistantly between $0\,^\circ$ and $360\,^\circ$. The voxel as well as the

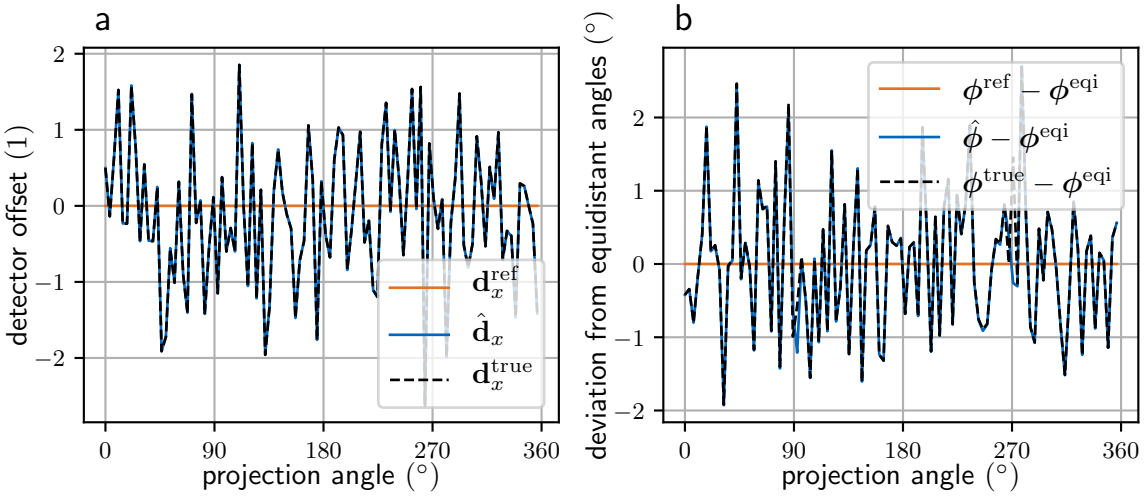

Figure 8.3.: Geometry parameters for the parallel-beam geometry. In (a) the detector shift is depicted for every angle. In orange the initial guess $\mathbf{d}_x^{\text{ref}}$ of having no detector shift is shown. In blue the estimated detector shift $\hat{\mathbf{d}}_x$ is shown for every angle. In comparison, the dashed black line shows the ground truth values $\mathbf{d}_x^{\text{true}}$ used for simulation, which match the estimated values extremely well. In (b) the deviation of the projection angle from the evenly-spaced initial guess is depicted. Again, in orange the initial guess $\boldsymbol{\phi}_x^{\text{ref}}$, in blue the estimate $\hat{\boldsymbol{\phi}}_x$ and in black the ground truth $\boldsymbol{\phi}_x^{\text{true}}$ are shown.

pixelsize were set to 1.

For the actual geometry, the detector offset of each projection was corrupted by a random shift drawn from a normal distribution with zero-mean and standard deviation of 1. Similarly, the tomographic angles were corrupted by random rotations drawn from a zero-mean normal distribution with standard deviation of 1 °. The actual values can be read from the black dashed lines in Figure 8.3. The line-integrals depicted in Figure 8.4 (a) were simulated by forward projecting the phantom. In Figure 8.4 (b) the FBP reconstruction $\mathbf{x}_{\text{fbp}}^{\text{ref}}$ is visualized assuming the uncorrupted geometry. Severe geometry artifacts arising predominantly from the jitter of the detector are visible.

For joint reconstruction $\mathbf{x}$, $\mathbf{d}_x$ and $\boldsymbol{\phi}$ were optimized for every view. Thus, in addition to the $64^2 = 4096$ unknown parameters of the object, $101 \times 2 = 202$ additional geometry parameters were estimated. Algorithm 8.3 was initialized with the simulated projections $\boldsymbol{\ell}$ depicted in Figure 8.4 (a), the FBP $\mathbf{x}_{\text{fbp}}^{\text{ref}}$ depicted in Figure 8.4 (b) blurred by a Gaussian with a standard deviation of 1, as well as the uncorrupted geometry parameters $\mathbf{d}_x^{\text{ref}}$ and $\boldsymbol{\phi}_x^{\text{equi}}$ depicted in Figure 8.3. The step lengths were manually chosen to be $\epsilon_{\mathbf{x}} = 10^{-3}$, $\epsilon_{\mathbf{d}_x} = 10^{-2}$ and $\epsilon_{\boldsymbol{\phi}} = 10^{-3}$ to balance the influence of these parameters on the objective function. In total 10000 iterations were performed.

In Figure 8.3, the estimated as well as the initial (reference) and corrupted (ground truth) geometry parameters are plotted. In (a), the estimates of the detector offsets are shown. One finds that the corrupted offsets used for simulation can be perfectly recovered. In addition, as shown in (b), the deviations from the equidistantly sampled projection angles are also recovered. However, there are still slight variations from the ground truth, especially around a projection angle of 90 °. The reconstructed volume $\hat{\mathbf{x}}$ after 10000 iterations can be seen in

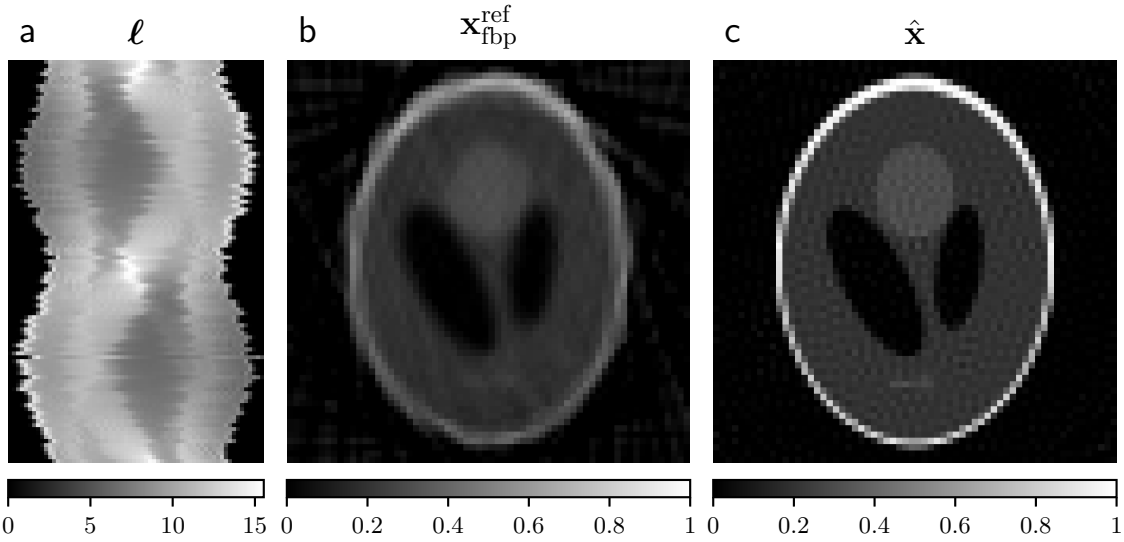

Figure 8.4.: Sinogram and reconstructions for the parallel-beam geometry. In (a) the sinogram $\boldsymbol{\ell}$ with the corrupted geometry is shown. The FBP reconstruction $\mathbf{x}_{\text{fbp}}^{\text{ref}}$ obtained with the uncorrupted geometry parameter is shown in (b). In (c) the estimate for the volume $\hat{\mathbf{x}}$ using joint geometry estimation and tomographic reconstruction is shown.

Figure 8.4 (c). The geometry artifacts of the FBP reconstruction are not visible anymore, resulting in a sharp depiction of the object. As no regularization is used, slight noise levels are visible (overfitting).

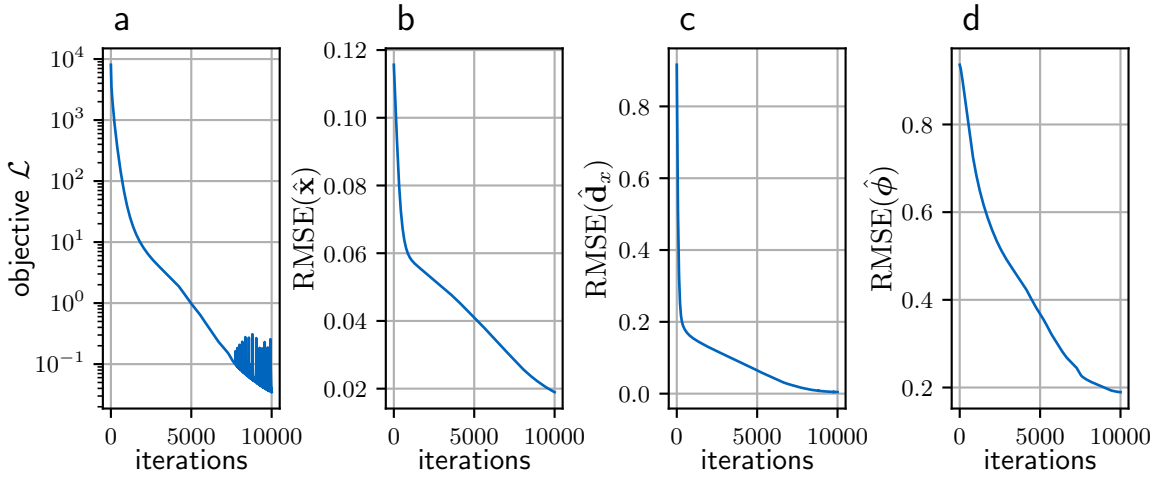

Figure 8.5.: Objective and deviations to ground truth for the parallel-beam geometry. In (a) the value of the likelihood is plotted over the iterations. In (b), (c) and (d), the differences to the ground truth of the reconstructed volume, the detector offsets and the projection angles are depicted.

Finally, Figure 8.5 shows additional information about the optimization process. In (a) the value of the objective function is shown as a function of the number the iterations. The value monotonously decreases over several orders of magnitude apart from small instabilities

at the end. In (b), the root MSE (RMSE) between the estimated object and the ground truth is shown over the number of iterations. For the first 1000 iterations, a strong decrease is visible. In (c), the RMSE of the estimated detectors offsets to the ground truth is depicted. As the geometry and thus the objective function is very sensitive to the detector offsets, this geometry parameter is optimized first. The detector offsets also have a strong influence on the representation of the object and thus the difference to the ground truth of the object decreases most significantly in the beginning. Finally, the comparison of the projection angles to the ground truth is shown. In this case, the influence on the objective function is less prominent than for the detector offsets. In comparison to the detectors offsets, the error decreases more slowly.

In conclusion, jointly estimating all geometry parameters during tomographic reconstruction for this idealized setup is feasible.

8.5.2. Fan-beam geometry

The number of geometry parameters for each view increases from two parameters in the (two-dimensional) parallel-beam geometry to five parameters in the fan-beam geometry. In addition to the detector offset and the projection angle, the SAD and SDD as well as the offset of the source orthogonal to the optical axis have to be accounted for. For the uncorrupted (reference) geometry no source and detector offsets were assumed. The SAD was set to $\mathbf{d}_{\text{SAD}} = 60$ and the SDD was set to $\mathbf{d}_{\text{SDD}} = 100$ for all views. As previously, 101 projections were acquired equidistantly distributed over 360 °. Again, unit voxel size was used. To account for the magnification, the pixelsize was set to 2.

For the corrupted (true) geometry, the source offsets were drawn from a zero-mean normal distribution with standard deviation 1, the detector offsets were drawn from a zero-mean normal distribution of standard deviation 2 and the tomographic angles were corrupted by adding random shifts drawn from a zero-mean normal distribution with standard deviation of 1 °. Finally, the SAD and SDD were corrupted for each projection by a normal distribution with a standard deviation of 1 % of the distances of the uncorrupted geometry. The geometry parameters are depicted in Figure 8.6 (b)-(f). The simulated sinogram is shown in Figure 8.7 (a) and the FBP reconstruction using the uncorrupted geometry parameters is plotted in Figure 8.7 (b). The quality of the reconstruction is significantly reduced by not correctly accounting for the geometry.

For iterative reconstruction, the volume as well as all five geometry parameters for each projection were estimated simultaneously. Therefore, Algorithm 8.3 was initialized with the simulated projections $\boldsymbol{\ell}$ shown in Figure 8.7 (a), the FBP $\mathbf{x}_{\text{fbp}}^{\text{ref}}$ obtained with the uncorrupted geometry parameters depicted in Figure 8.7 (b), which was afterwards blurred with a Gaussian filter with a standard deviation of 1, as well as the uncorrupted geometry parameters as defined above and shown in Figure 8.6 (b)-(f). The step length for the volume update was set to $\alpha_{\mathbf{x}} = 10^{-3}$. The step length for each geometry parameter was tuned manually to $\epsilon_{\mathbf{d}_{\text{SAD}}} = 10^{-3}, \epsilon_{\mathbf{d}_{\text{SDD}}} = 10^{-4}, \epsilon_{\mathbf{s}_x} = 5 \cdot 10^{-3}, \epsilon_{\mathbf{d}_x} = 10^{-2}$ and $\epsilon_{\boldsymbol{\phi}} = 2 \cdot 10^{-3}$.

Figure 8.6 (a) shows the value of the objective function as a function of the number of iterations. The value decreases again over several orders of magnitude, with some outliers at the end, which are attributed to the optimization algorithm. In (b)-(f), the estimated, true and initial geometry parameters are shown. The tomographic angle, the source offset as well as the detector offset can be estimated very accurately. However, there are inaccuracies for the estimation of the SAD as well as the SDD. Nevertheless, a slight change of these values hardly influences the likelihood term and their influence on the image quality is

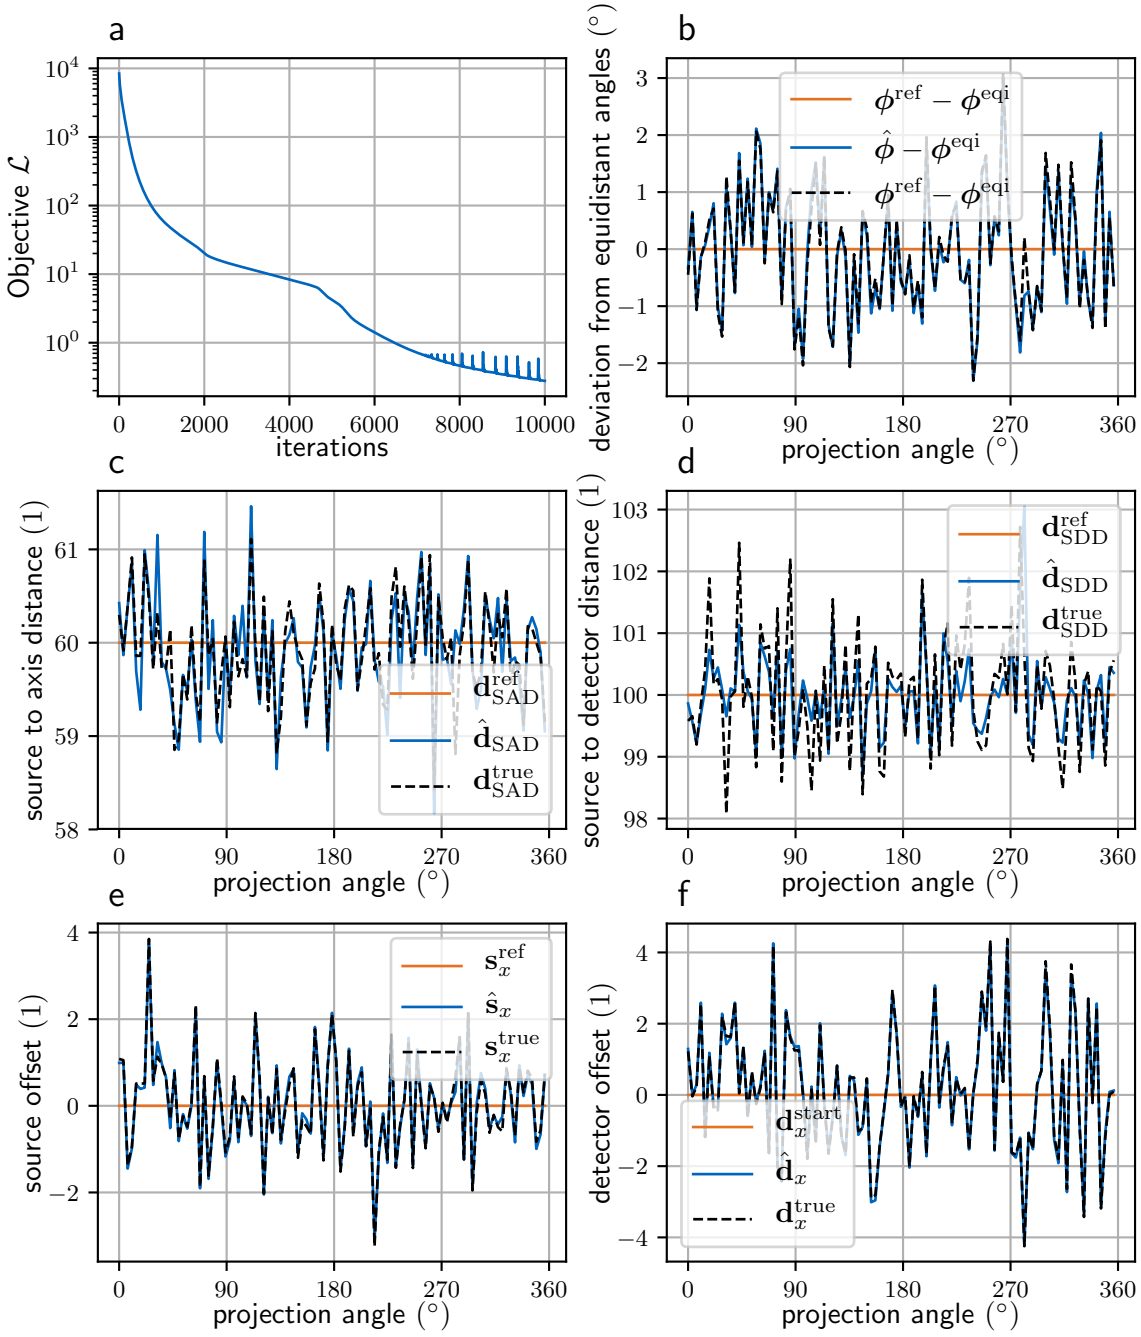

Figure 8.6.: Likelihood and geometry parameters for the fan-beam geometry. In (a) the value of the likelihood term is plotted over the iterations. In (b), the deviation of the projection angle from the evenly-spaced initial guess $\boldsymbol{\phi}^{\text{eqi}}$ is depicted. In (c)-(f), the values of the SAD, SDD, source offset and detector offset are shown respectively. In orange, the values of the uncorrupted geometry, in blue, the estimated values, and in back the true values for each of the geometry parameters are illustrated.

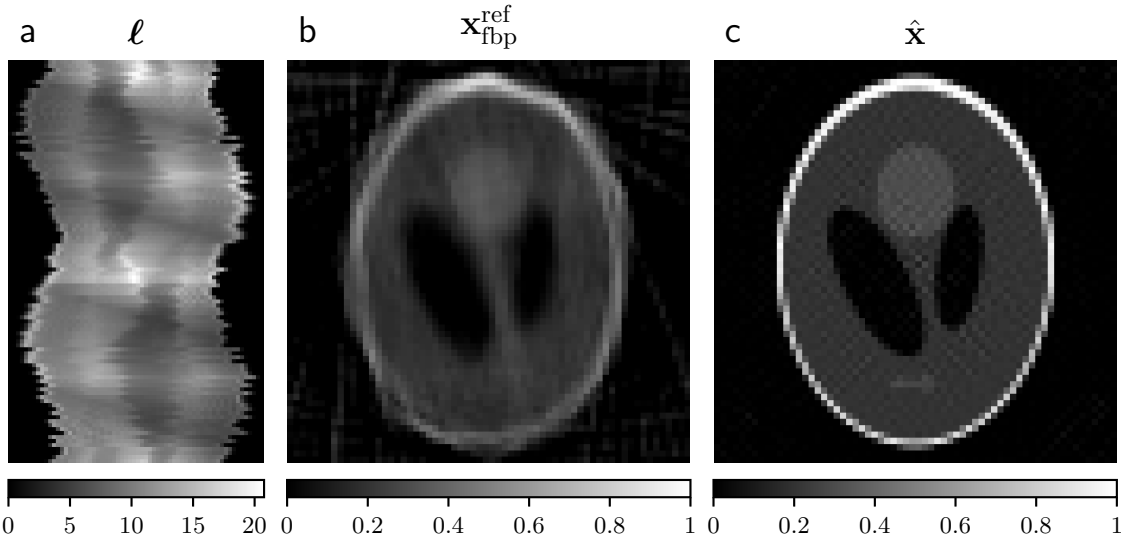

Figure 8.7.: Sinogram and reconstruction for the fan-beam geometry. In (a), the simulated measurement generated with the unknown geometry is shown. In (b) the conventional FBP reconstruction $\hat{\mathbf{x}}_{\text{fbp}}^{\text{ref}}$ using the uncorrupted geometry is shown. In (c) the estimate for $\hat{\mathbf{x}}$ using the iterative reconstruction approach is depicted.

comparably small. The reconstructed volume is depicted in Figure 8.7 (c). Despite the strong manipulations performed on the geometry, a sharp object is obtained. The resulting reconstruction depicts slight noise levels, as no regularization was used.

In Figure 8.8, the RMSE of the volume and the geometry parameters are depicted as a function of the number of iterations. The RMSEs are monotonically decreasing for the volume as well as all geometry parameters. However, increasing the number of iterations beyond 10000 has shown to result in an increase of the RMSE for the SAD and SDD, as for certain views the estimated values diverge from the ground truth. As for the parallel-beam case, the detector offset has the greatest influence on the likelihood and is thus optimized first.

This study shows that, in principle, all geometry parameters can be estimated even for a fan-beam geometry. However, the exact values of the step lengths for each geometry parameter have to be determined manually, which can be difficult. On the one hand, it was also observed that accurately estimating the SAD and SDD was difficult, as their impact on the value of the likelihood is small. On the other hand, this also implies that an inaccurate estimation of these parameters does not corrupt the reconstruction quality significantly.

8.6. Application of joint geometry estimation at the NanoCT

In this section, joint geometry optimization and tomographic reconstruction is applied to a dataset measured at the NanoCT. The experimental setup is described in Subsection 8.3.1. Again, the geometry estimation is restricted to refining the horizontal detector offsets, simply referred to as detector offsets in the following. The cone-beam projector, which is elaborated in [Fehringer et al., 2014, Fehringer, 2019], is used consisting of a pixel-driven backprojector described in Figure 4.3 (a) and a ray-driven forward projector illustrated in Figure 4.3 (b).

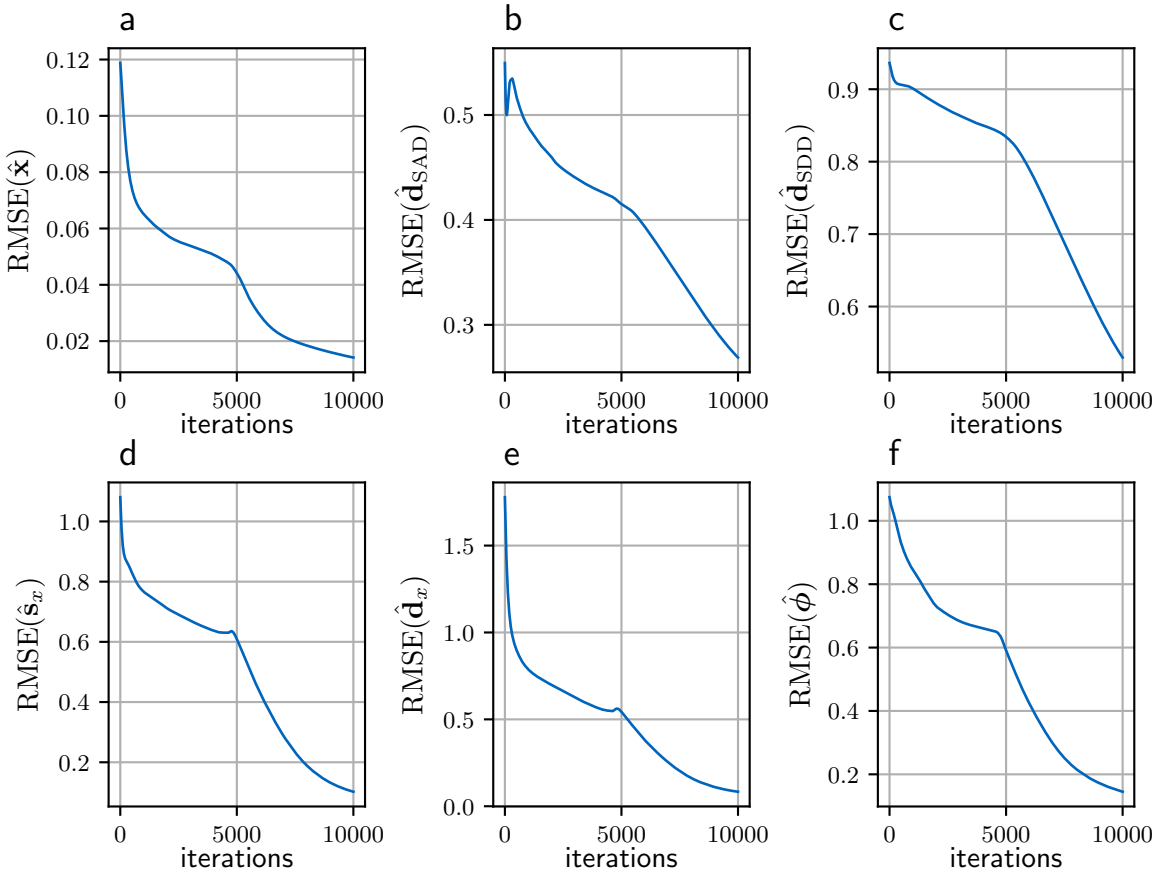

Figure 8.8.: Deviations to ground truth for the fan-beam geometry. In (a) the RMSE of the volume to the ground truth is depicted over the number of iterations. In (b)-(f) the corresponding information for the geometry parameters is depicted.

8.6.1. Sample, acquisition and preprocessing

The sample[4] consists of sand grains attached to a PMMA rod using superglue. The SAD was set to 3.36 mm and the SDD was set to 200 mm, resulting in a magnification of around 60. Given the detector pixelsize of 172 µm, the effective voxel size is given as 2.9 µm. 1599 projections were acquired[5] equidistantly around the sample. In addition, reference projections without the sample were acquired. The exposure time was set to 2 s.

For preprocessing, dead and corrupted pixels were removed by replacing their value with the median of neighboring pixel values. Using the reference projections, the inhomogeneous illumination and the individual pixel responses were corrected (details can be found in [Allner, 2019]). As the detector was shifted during acquisition, the missing data from the module gaps of the detector were interpolated from neighboring projections. Afterwards, the detector shifts were corrected. In addition, the intensity was adjusted by an intensity ramp and finally the projections were cropped to reduce computational time. Although slight edge-enhancement effects are visible, no phase retrieval was performed and only the line-integrals were computed.

[4] Sample preparation was performed by Simone Ferstl and Katharina Scheidt.
[5] Data acquisition was done by Katharina Scheidt.

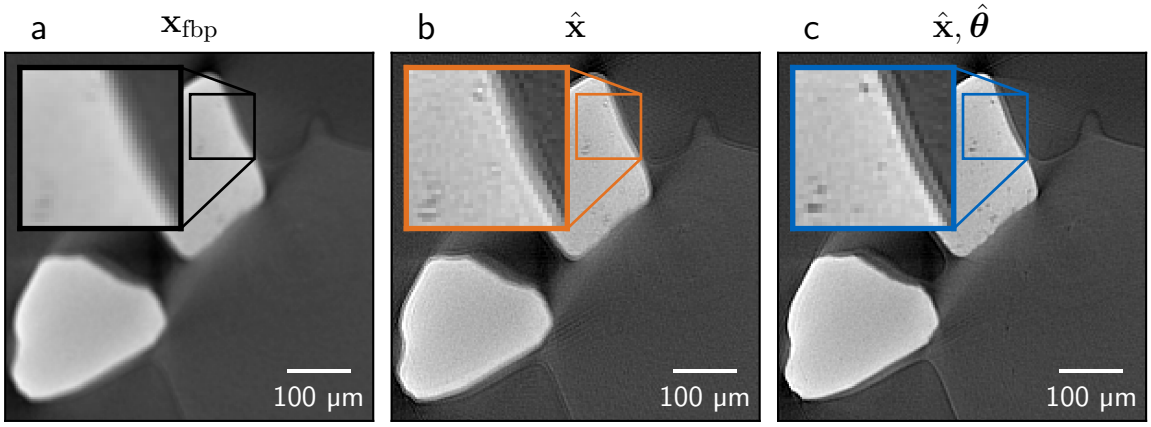

Figure 8.9.: Reconstructions using joint geometry estimation at the NanoCT. In (a), an extract of the reconstructed sea sand sample using FBP is depicted with an additional zoom to the interface between the sand, glue and air. In (b), the corresponding information is shown for the reference reconstruction, which uses iterative reconstruction. For the analytical reconstruction and the reference reconstruction no additional geometry refinement was performed during reconstruction. In (c) the horizontal detector offset was estimated during tomographic reconstruction.

8.6.2. Reference reconstructions and joint geometry estimation

The geometry was defined according to Eq. (8.8). It was assumed that there is no rotation of the object around the optical axis, no tip in direction of the optical axis, and the tomographic angles are equidistantly distributed between $0\,^\circ$ and $360\,^\circ$. In addition, it was assumed that there is (on average) no offset of the source and the detector. This is reasonable as the cropping was performed such that the sample is aligned to the center.

First, an FBP was computed with the geometry defined above. A region of a reconstructed slice can be seen in Figure 8.9 (a). From the provided zoom one can observe that the interface between the sand and air is blurred significantly. However, the artifacts at the reconstructed slice do not obviously hint at a wrong value for the detector offset, as other sources of system blur could potentially explain this effect as well.

Next, an iterative reconstruction with the same geometry is performed using the objective given by Eq. (8.22). However, the geometry parameters are fixed and thus, this objective is only optimized with respect to the volume. Algorithm 8.3 was used with the preprocessed projections, the FBP as initial guess and the same geometry parameters as defined above. The step length is set to $\alpha_x = 10^{-4}$ and the number of iterations to $N = 250$. As no geometry parameters are updated, the other parameters are not provided. The respective extract of the reconstruction is shown in Figure 8.9 (b). The edges are less blurred as expected from the iterative algorithm. As the noise levels are moderate, the noise does not increase significantly with increasing number of iterations. From this extract, double edges are visible which hint at a wrong estimate of the detector offset.

Finally, the detector offset is estimated during tomographic reconstruction using again Algorithm 8.3 with the same arguments as provided before. In addition, a set $M = \{d_x\}$ was

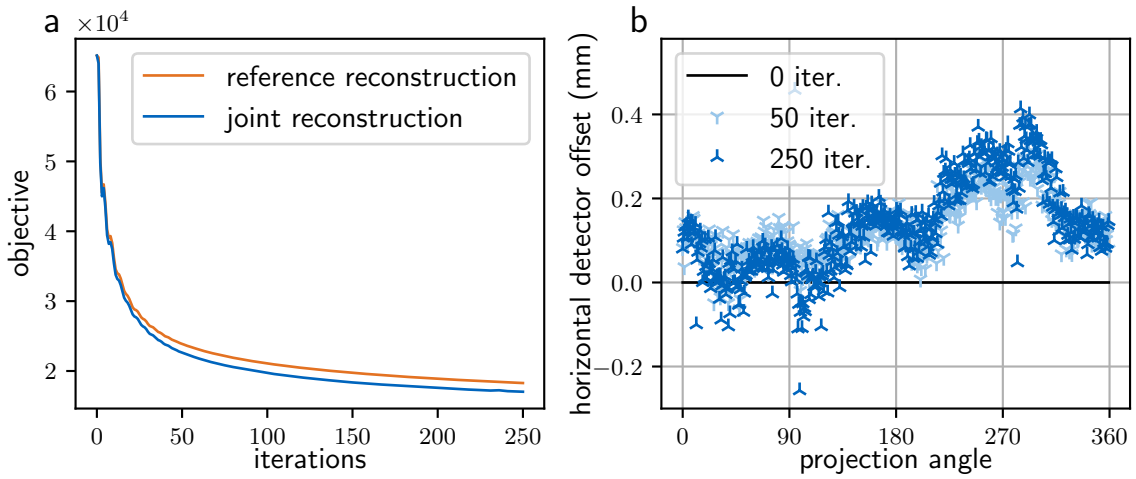

Figure 8.10.: Likelihood and detector offsets at the NanoCT. In (a), the value of the likelihood term of the joint reconstruction approach is depicted in blue as a function of the number of iterations. For reference, the corresponding value of the likelihood term, where no additional geometry optimization was performed, is shown in orange. In (b), the horizontal detector offsets are plotted over the projection angles. In black, the initial guess is shown. The detector offsets estimated after 50 iterations are shown in light blue and the estimated detector offsets after 250 iterations are depicted in (dark) blue.

provided, referring to the detector offset, and the step length was chosen according to $\alpha_{d_x} = 0.1$. The corresponding extract of the reconstruction can be seen in Figure 8.9 (c). The edges are clearly sharper and the edge-enhancement effects become visible in the reconstruction as dark contour lines at the interfaces to air. In the provided zoom one can precisely distinguish the additional layer of glue between the sand and the air. The noise level is comparable with the previous iterative reconstruction.

In Figure 8.10 (a), the value of the likelihood attributed to the reference reconstruction and the value of the joint geometry estimation and tomographic reconstruction are depicted as a function of the number of iterations. The ADAM optimization does not guarantee that the values decrease monotonically and small bumps can be observed, which may be resolved by specifying a smaller step length, which however may slow down convergence. In general, the value of the likelihood of the joint reconstruction is lower than the corresponding value of the reference reconstruction. This suggests that the model for the joint reconstruction can better approximate the actual measurements. In (b) the detector offset is plotted as a function of the projection angle. The initial guess corresponds to no detector offset for all projections. After 50 iterations a nonuniform offset can be observed, which already approximates the estimated detector offset after 250 iterations reasonably well.

In summary, joint geometry optimization and reconstruction can, in principle, be performed at the NanoCT to estimate geometry parameters. Thus, even with additional noise, an imperfect knowledge of the remaining geometry parameters, an incomplete physical mean model (as it does not account for the interference effects) and an unmatched projector model, joint geometry optimization and reconstruction is feasible.

8.7. Discussion and conclusion

In this chapter, two methods for estimating geometry parameters in high-resolution CT were investigated. In the following, the benefits and drawbacks of the respective methods and their potential to improve image quality for reconstructions at the NanoCT are discussed. This is particularly helpful, as the error in the geometry parameters for this application is not predictable and conventional geometry characterization methods performed prior to data acquisition do not suffice.

The first method used the normalized sparsity metric for CT, which was developed in Chapter 7, to estimate geometry parameters. This method was found to be robust to noise and compatible with region of interest (ROI) reconstructions. Most importantly, in principle, every geometry parameter can be estimated using this method. In addition, it was shown that the geometry parameters can be estimated for every view individually. This method was applied to a complex tomographic dataset at the NanoCT, where the sample exceeded the FOV significantly. Computational cost is the main disadvantage of this approach. However, a computationally more efficient approach was introduced, which performs the estimation of the geometry parameters on a coarser grid and uses interpolation to estimate the geometry parameters for the original geometry. Assuming the increase in computational time is acceptable, this approach holds the potential to estimate multiple parameters at once. Furthermore, by using multiple iterations of this approach, the geometry parameters can be further refined.

The second method jointly estimates the geometry parameters during tomographic reconstruction. In contrast to the first method, where geometric alignment has to be performed as a separate step prior to a MBIR, this method integrates the refinement of the geometry parameters in MBIR. This is conceptually elegant as the geometry parameters can be interpreted as additional parameters of the physical mean model, similarly to the blur parameters in blind deconvolution CT reconstruction introduced in Chapter 7. A key contribution is the efficient evaluation of the gradient of the physical mean model with respect to the geometry parameters. It was shown in a simulation study that under ideal conditions, the geometry parameters for two-dimensional geometries can be estimated for every view during tomographic reconstruction. In addition, this approach was evaluated on a simple phantom measurement acquired at the NanoCT. It could be demonstrated that the horizontal detector offset can be recovered during reconstruction. However, it remains unclear if this approach is feasible for datasets with high noise levels and ROI reconstructions. In addition, the generalization to estimate multiple geometry parameters and to complex datasets could not be demonstrated. Investigating the limitations of this approach would be desireable.

For applications at the NanoCT, the focus for both methods was on the horizontal detector offset, which was shown to have a significant influence on the image quality. However, estimating all geometry parameters at once is not feasible, due to the computational cost of the first method or the fact that the optimization problem of the second method is highly unconstrained. Thus, the properties of the experimental setup have to be understood reasonably well to determine which geometry parameters are not reliable due to random shifts and need to be estimated and which geometry parameters can be assumed to be known such that their influence of the image quality can be neglected. In addition, it should be assessed if it is sufficient to estimate the global geometry parameters (for instance tip and tilt of the rotation axis), if one can parameterize the dependency of the geometry parameters with the views, or if it is necessary to estimate the parameters for each view individually. Estimating only global parameters reduces the complexity of the optimization problems significantly.

In conclusion, for most scenarios, the first method using the normalized sparsity measure for CT will be more suitable as it is more robust. In general, the exact implementation of any of the described methods varies between different setups and needs to be tailored individually to the respective uncertainties of the geometry parameters.

9. Conclusion and outlook

This final chapter presents a short summary of the main results followed by an outlook regarding future developments, showing how the key algorithms designed here can be combined and utilized for different imaging scenarios.

9.1. Summary of results

In the following sections, the main algorithmic advances to improve image quality in different X-ray imaging modalities are summarized for each chapter.

Non-linear SIR framework for PBI and PB-CT

In Chapter 5, a versatile non-linear SIR framework was proposed for PBI and PB-CT using solely a single propagation distance. This framework directly incorporates into the reconstruction a model for the propagation of the X-rays and a statistical description of the measurements. For PBI, if a pure phase object is assumed, the analytical solution of the proposed framework coincides with Bronnikov's phase retrieval algorithm. Using the homogeneity assumption, it was demonstrated experimentally that the results are in good agreement with the results obtained by the single-material phase-retrieval algorithm by Paganin. For PB-CT, in the limit of an incoherent source, the proposed framework coincides with a SIR algorithm for conventional attenuation-based CT. The main result of this chapter is a SIR algorithm for PB-CT, which employs the homogeneity assumption. Experiments performed with a laboratory X-ray source showed how materials that are not described by the physical mean model can be excluded from the reconstruction. In further experiments using synchrotron radiation, the proposed approach was able to remove artifacts arising from a cochlear implant, artifacts that appear because the implant violates the homogeneity assumption. This allows for better assessment of whether the insertion of the implant has resulted in subtle damages, which would naturally be close to the implant itself, hence in the region of the image where artifacts are otherwise found.

Modeling the source and the detector in homogeneous PBI and PB-CT

Building upon the results of Chapter 5, in Chapter 6 reconstruction algorithms for homogeneous PBI and PB-CT were proposed, which additionally model the effects of the spatial coherence of the X-ray source and the response of the X-ray detector including noise correlations. The model for the interference effects was extended to match the formulation of the single-material phase-retrieval algorithm. For homogeneous PBI, an analytical solution for a special case of the proposed algorithm could be derived, which further generalizes the single-material phase-retrieval algorithm and its extensions. Moreover, a relation between the single-material phase-retrieval algorithm and regularized image denoising was outlined. For PB-CT, a simulation study as well as an experimental study at the MuCLS were performed. The proposed algorithms were compared to each other and to several analytical

reconstruction approaches. As a main result, the proposed algorithms improve image quality significantly compared to conventional reconstruction approaches. In particular, modeling noise correlations has a crucial impact on resolution and the suppression of overshoots at sample edges, which are characteristic to reconstructions in PB-CT.

Joint deconvolution CT reconstruction

In the context of conventional attenuation-based CT, Chapter 7 presented an optimization approach, which allows for joint estimation of low-dimensional parameterized models of the system blur during tomographic reconstruction. In detail, a new regularization term was devised based on the normalized sparsity metric, which was introduced in the context of blind deconvolution in optical imaging. An extensive simulation study demonstrated that the devised objective function has a global optimum for the true blur parameters and that the proposed optimization approach is able to recover these parameters. Furthermore, this approach was successfully verified experimentally at a CT test-bench setup. As a consequence, not only laborious characterizations of all components that add to the system blur can be omitted, but this approach has furthermore the potential to include more subtle contributions that are difficult to characterize, such as interpolations in the projection operations.

Optimization-based geometry estimation for CT at the nanometer scale

Two optimization-based approaches were investigated in Chapter 8, which estimate the geometry of CT components directly from the measurements. The first method is based on the metric introduced in Chapter 7 for joint deconvolution CT reconstruction. The versatility of this approach allows, in principle, for the estimation of arbitrary geometry parameters for every view individually. The significant improvements in image quality were demonstrated for a complex biological sample which exceeded the FOV. Thereby, the horizontal drifts of the respective components could be estimated for each view. Furthermore, an approach which jointly estimates the geometry parameter during tomographic reconstruction was investigated. In detail, an optimization algorithm, which is independent of the model for the projection operations, was introduced and analyzed using a simulation study. This approach was also verified experimentally by estimating the horizontal detector offset for every view individually.

9.2. Outlook

Finally, this section presents ideas to further refine and enhance the results shown in this work. It concludes with two case studies demonstrating how the results outlined in this work could be combined and utilized for two different experimental setups.

9.2.1. Further development

Building upon the SIR framework for PB-CT presented in Chapter 5, a challenging task would be to examine means to more independently reconstruct the attenuation and phase-shifting properties of the sample by including additional priors into the reconstruction framework. Independently, a more detailed comparison of the proposed algorithm, which employs the homogeneity assumption, to the conventional two-step approaches could be performed.

On the one hand, the bias introduced due to Jensen's inequality should be reduced and, on the other hand, one could propagate the noise of the measurements through the phase-retrieval step and then utilize the highly correlated noise in a subsequent tomographic reconstruction using for instance means addressed in Chapter 6.

Concerning the results in Chapter 6, the differences between the proposed generalized single-material phase-retrieval algorithm and the conventional single-material phase-retrieval algorithm and its extension should be evaluated on experimental data. In PB-CT, choosing the best propagation distance to be most sensitive to the phase information is crucial. As the proposed algorithm can account for the source blur and the detector blur differently, it would be worthwhile to examine how the reconstruction algorithm influences the optimal choice of the propagation distance. This would then directly impact the choice of the experimental parameters. Moreover, means to include polychromatic effects at least approximately by an additional linear operator could be investigated. Moreover, the proposed concept could be transferred to other phase-contrast techniques such as GB-CT, where higher resolution is also limited by the finite source size and the detector. This has the potential to further increase resolution in laboratory environments.

For blind deconvolution CT reconstruction introduced in Chapter 7, the estimation of higher-dimensional blur parameters would be desirable, as in most cases the system blur cannot be described accurately by a single parameter. In addition, further means to solve the optimization problem would be of interest, as with a truly gradient-based approach the restrictions on the dimensionality of the blur model would be reduced dramatically. For instance, one could investigate analytical approaches to approximate the dependency of the volume on the blur parameters. Further, the influence of the likelihood term could be alleviated by including noise correlations in the reconstruction framework in a similar manner to Chapter 6. Thus, the high-frequency components, which are essential when estimating the blur, would be emphasized. Moreover, the effects of noise should be further investigated, as the proposed approach is prone to high noise levels. For instance, using the Huber penalty as a basis for devising an improved regularization term seems reasonable. Another promising approach would be to use this framework to jointly estimate the material specific constant $z\delta/\mu$, parameterizing the interference term in PB-CT within the reconstruction approach introduced in Chapter 6. Similar to blur, a change of this value alters the response to noise. This seems to be particularly promising for several reasons. Firstly, the high-frequency components are automatically enhanced on the measured projections, which are most important when estimating the blur. Secondly, the high frequencies in the volume are reduced due to the integrated phase retrieval. Thus, the response of the regularization term should be less corrupted by noise. Eventually, optimizing this factor would potentially also include the effects of the source, detector, and noise as these are closely related to this factor as outlined in Chapter 6.

Finally, the metric-guided geometry optimization algorithm investigated in Chapter 8 could be applied to additional geometry parameters to evaluate if the image quality can be further improved. In addition, a comparison to other metrics proposed in the literature would be of interest as well as evaluating if this method can be transferred to other setups. Regarding the joint geometry estimation, the challenges arising when applying this approach to more complex samples should be investigated further. Most beneficial for both methods would be on the one hand, more precise experimental investigations to determine, which of the respective CT components are most prone to misalignment, and on the other hand, the development of parameterized models describing the drifts of the individual components, as this would reduce the number of unknown parameters significantly.

9.2.2. Application to different setups

The first setup is the NanoCT described in Subsection 8.3.1. The following text outlines how the results presented here can be combined to improve image quality. Due to the small focal spot and large propagation distance at this setup, interference effects are present, which can be used to improve contrast and reduce the noise level. In addition, this setup uses high magnifications to improve the spatial resolution. However, the resolution is ultimately limited by the extent of the source. Thus, the algorithm for PB-CT introduced in Chapter 6 would be an ideal candidate, as it includes a model for the interference effects and can account for the source blur. The blur and noise correlations induced by the photon-counting detector do not necessarily have to be modeled as the influence of the detector is comparably small. One main challenge at this setup is mechanical stability. As demonstrated in Chapter 8, means exist to estimate additional geometry parameters during tomographic reconstruction. However, the geometry parameters prone to drift should be identified beforehand and should be parameterized accordingly to reduce the number of unknown parameters which need to be jointly estimated as much as possible.

The second setup presented within the range of this work is the first experimental hutch of the MuCLS, which was described in Subsection 6.5.1. In contrast to the NanoCT, the geometry can be assumed to be known, thus no additional geometry parameters need to be estimated. Due to the small divergence of the beam, efficient detectors with small pixel-sizes are used to increase resolution, which however spread out the signal more. Imaging techniques applied at this setup include conventional attenuation-based CT as well as PB-CT. For conventional CT, modeling the detector response and the attenuation of the X-rays should be sufficient in general, because by placing the detector close to the sample the influence of the source is reduced. This results in the algorithm discussed in Chapter 6 in the limiting case of no phase effects. If the detector response is not known, the approach detailed in Chapter 7 can be used provided the noise level is moderate. For PB-CT, the distance between the sample and the detector is increased to be more sensitive to interference effects. The distance is ultimately limited by the influence of the source. Thus, the propagation distance, the spatial coherence of the source and the detector response have to be incorporated in order to be most sensitive to the phase-shifting properties of the sample. All of these are accounted for in the algorithm introduced in Chapter 6. As discussed above, using the framework derived in Chapter 7 seems to be a promising approach to jointly estimate the material specific constant $z\delta/\mu$ parameterizing the interference term.

Overall, including additional parameters, which have to be estimated, and more complex models describing the image formation can improve the image quality significantly but also make the optimization problem more difficult. Appropriate algorithms need to be tailored to each setup individually. Thereby, the most prominent factors that diminish the image quality have to be identified and included such that the resulting optimization problem remains solvable.

Bibliography

[Achterhold et al., 2013] Achterhold, K., Bech, M., Schleede, S., Potdevin, G., Ruth, R., Loewen, R., and Pfeiffer, F. (2013). Monochromatic computed tomography with a compact laser-driven X-ray source. *Scientific Reports*, 3:1313.

[Allner, 2019] Allner, S. (2019). PhD thesis, Technical University of Munich, Germany. In preparation.

[Als-Nielsen and McMorrow, 2011] Als-Nielsen, J. and McMorrow, D. (2011). *Elements of modern X-ray physics*. John Wiley & Sons.

[Andersen and Kak, 1984] Andersen, A. H. and Kak, A. C. (1984). Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm. *Ultrasonic Imaging*, 6(1):81–94.

[Arcadu et al., 2016] Arcadu, F., Stampanoni, M., and Marone, F. (2016). On the crucial impact of the coupling projector-backprojector in iterative tomographic reconstruction. *arXiv preprint arXiv:1612.05515*.

[Arhatari et al., 2010] Arhatari, B. D., Gates, W. P., Eshtiaghi, N., and Peele, A. G. (2010). Phase retrieval tomography in the presence of noise. *Journal of Applied Physics*, 107(3):034904.

[Attwood and Sakdinawat, 2017] Attwood, D. and Sakdinawat, A. (2017). *X-rays and extreme ultraviolet radiation: principles and applications*. Cambridge University Press.

[Beckmann, 2006] Beckmann, E. C. (2006). CT scanning the early days. *The British Journal of Radiology*, 79(937):5–8.

[Behling, 2015] Behling, R. (2015). *Modern Diagnostic X-ray Sources: Technology, Manufacturing, Reliability*. CRC Press.

[Beltran et al., 2018] Beltran, M. A., Paganin, D. M., and Pelliccia, D. (2018). Phase-and-amplitude recovery from a single phase-contrast image using partially spatially coherent X-ray radiation. *Journal of Optics*, 20(5):055605.

[Beltran et al., 2011] Beltran, M. A., Paganin, D. M., Siu, K. K. W., Fouras, A., Hooper, S. B., Reser, D. H., and Kitchen, M. J. (2011). Interface-specific X-ray phase retrieval tomography of complex biological organs. *Physics in Medicine & Biology*, 56(23):7353.

[Beltran et al., 2010] Beltran, M. A., Paganin, D. M., Uesugi, K., and Kitchen, M. J. (2010). 2D and 3D X-ray phase retrieval of multi-material objects using a single defocus distance. *Optics Express*, 18(7):6423–6436.

[Bidola, 2017] Bidola, P. (2017). *Characterization and application of high resolution phase-contrast laboratory micro-CT setups*. PhD thesis, Technical University of Munich, Germany.

[Bidola et al., 2015] Bidola, P. M., Zanette, I., Achterhold, K., Holzner, C., and Pfeiffer, F. (2015). Optimization of propagation-based phase-contrast imaging at a laboratory setup. *Optics Express*, 23(23):30000–30013.

[Birnbacher, 2018] Birnbacher, L. (2018). *High-sensitivity grating-based phase-contrast computed tomography with incoherent sources*. PhD thesis, Technical University of Munich, Germany.

[Boas and Fleischmann, 2011] Boas, F. E. and Fleischmann, D. (2011). Evaluation of two iterative techniques for reducing metal artifacts in computed tomography. *Radiology*, 259(3):894–902. PMID: 21357521.

[Bonse and Hart, 1965] Bonse, U. and Hart, M. (1965). An X-ray interferometer. *Applied Physics Letters*, 6(8):155–156.

[Bravin et al., 2013] Bravin, A., Coan, P., and Suortti, P. (2013). X-ray phase-contrast imaging: from pre-clinical applications towards clinics. *Physics in Medicine & Biology*, 58(1):R1–R35.

[Bremmer, 1952] Bremmer, H. (1952). On the asymptotic evaluation of diffraction integrals with a special view to the theory of defocusing and optical contrast. *Physica*, 18(6-7):469–485.

[Brendel et al., 2016] Brendel, B., von Teuffenbach, M., Noël, P. B., Pfeiffer, F., and Koehler, T. (2016). Penalized maximum likelihood reconstruction for X-ray differential phase-contrast tomography. *Medical Physics*, 43(1):188–194.

[Brent, 1971] Brent, R. P. (1971). An algorithm with guaranteed convergence for finding a zero of a function. *The Computer Journal*, 14(4):422–425.

[Bronnikov, 1999] Bronnikov, A. V. (1999). Reconstruction formulas in phase-contrast tomography. *Optics Communications*, 171(4–6):239–244.

[Bronnikov, 2002] Bronnikov, A. V. (2002). Theory of quantitative phase-contrast computed tomography. *Journal of the Optical Society of America*, 19(3):472–480.

[Bui et al., 2017] Bui, M., Albarqouni, S., Schrapp, M., Navab, N., and Ilic, S. (2017). X-ray posenet: 6 DoF pose estimation for mobile X-ray devices. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1036–1044.

[Burvall et al., 2011] Burvall, A., Lundström, U., Takman, P. A. C., Larsson, D. H., and Hertz, H. M. (2011). Phase retrieval in X-ray phase-contrast imaging suitable for tomography. *Optics Express*, 19(11):10359–10376.

[Buzug, 2008] Buzug, T. M. (2008). *Computed tomography: from photon statistics to modern cone-beam CT*. Springer Science & Business Media.

[Carter, 1993] Carter, W. H. (1993). Coherence theory.

[Chan and Wong, 1998] Chan, T. F. and Wong, C.-K. (1998). Total variation blind deconvolution. *IEEE Transactions on Image Processing*, 7(3):370–375.

[Cho et al., 2005] Cho, Y., Moseley, D. J., Siewerdsen, J. H., and Jaffray, D. A. (2005). Accurate technique for complete geometric calibration of cone-beam computed tomography systems. *Medical Physics*, 32(4):968–983.

[Cloetens et al., 1996] Cloetens, P., Barrett, R., Baruchel, J., Guigay, J.-P., and Schlenker, M. (1996). Phase objects in synchrotron radiation hard X-ray imaging. *Journal of Physics D: Applied Physics*, 29(1):133–146.

[Commission, 2008] Commission, I. E. (2008). *Radiotherapy equipment: coordinates, movements and scales*. IEC publication.

[Cont, 2016] Cont, D. (2016). X-ray beam characterization at the Munich Compact Light Source. Master's thesis, Technical University of Munich, Germany.

[Cowley, 1995] Cowley, J. M. (1995). *Diffraction physics*. Elsevier.

[David et al., 2002] David, C., Nöhammer, B., Solak, H. H., and Ziegler, E. (2002). Differential X-ray phase contrast imaging using a shearing interferometer. *Applied Physics Letters*, 81(17):3287–3289.

[Davis et al., 1995] Davis, T. J., Gao, D., Gureyev, T. E., Stevenson, A. W., and Wilkins, S. W. (1995). Phase-contrast imaging of weakly absorbing materials using hard X-rays. *Nature*, 373(6515):595–598.

[De Man and Basu, 2004] De Man, B. and Basu, S. (2004). Distance-driven projection and backprojection in three dimensions. *Physics in Medicine & Biology*, 49(11):2463.

[De Rosier and Klug, 1968] De Rosier, D. and Klug, A. (1968). Reconstruction of three dimensional structures from electron micrographs. *Nature*, 217(5124):130.

[De Witte et al., 2009] De Witte, Y., Boone, M., Vlassenbroeck, J., Dierick, M., and Van Hoorebeke, L. (2009). Bronnikov-aided correction for X-ray computed tomography. *Journal of the Optical Society of America*, 26(4):890–894.

[Dittmann, 2018] Dittmann, J. (2018). Cone beam geometry calibration for micro computed tomography systems using arbitrary fiducial markers and the relation between projection matrices and real space geometry. *arXiv preprint arXiv:1808.03337*.

[Donath et al., 2006] Donath, T., Beckmann, F., and Schreyer, A. (2006). Automated determination of the center of rotation in tomography data. *Journal of the Optical Society of America*, 23(5):1048–1057.

[Eggl et al., 2016] Eggl, E., Dierolf, M., Achterhold, K., Jud, C., Günther, B., Braig, E., Gleich, B., and Pfeiffer, F. (2016). The Munich compact light source: initial performance measures. *Journal of Synchrotron Radiation*, 23(5):1137–1142.

[Eisebitt et al., 2004] Eisebitt, S., Lüning, J., Schlotter, W., Lörgen, M., Hellwig, O., Eberhardt, W., and Stöhr, J. (2004). Lensless imaging of magnetic nanostructures by X-ray spectro-holography. *Nature*, 432(7019):885.

[Elad, 2002] Elad, M. (2002). On the origin of the bilateral filter and ways to improve it. *IEEE Transactions on Image Processing*, 11(10):1141–1151.

[Elbakri and Fessler, 2003] Elbakri, I. A. and Fessler, J. A. (2003). Efficient and accurate likelihood for iterative image reconstruction in X-ray computed tomography. In *Medical Imaging 2003: Image Processing*, volume 5032, pages 1839–1851. International Society for Optics and Photonics.

[Endrizzi, 2018] Endrizzi, M. (2018). X-ray phase-contrast imaging. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 878:88–98.

[Erdogan and Fessler, 1998] Erdogan, H. and Fessler, J. A. (1998). Accelerated monotonic algorithms for transmission tomography. In *International Conference on Image Processing*, volume 2, pages 680–684. IEEE.

[Erdogan and Fessler, 1999a] Erdogan, H. and Fessler, J. A. (1999a). Monotonic algorithms for transmission tomography. *IEEE Transactions on Medical Imaging*, 18(9):801–814.

[Erdogan and Fessler, 1999b] Erdogan, H. and Fessler, J. A. (1999b). Ordered subsets algorithms for transmission tomography. *Physics in Medicine & Biology*, 44(11):2835.

[Fehringer, 2019] Fehringer, A. (2019). PhD thesis, Technical University of Munich, Germany. In preparation.

[Fehringer et al., 2014] Fehringer, A., Lasser, T., Zanette, I., Noël, P. B., and Pfeiffer, F. (2014). A versatile tomographic forward-and back-projection approach on multi-GPUs. In *Medical Imaging 2014: Image Processing*, volume 9034, page 90344F. International Society for Optics and Photonics.

[Feldkamp et al., 1984] Feldkamp, L. A., Davis, L. C., and Kress, J. W. (1984). Practical cone-beam algorithm. *Journal of the Optical Society of America*, 1(6):612–619.

[Feng et al., 2006] Feng, B., Fessler, J. A., and King, M. A. (2006). Incorporation of system resolution compensation (RC) in the ordered-subset transmission (OSTR) algorithm for transmission imaging in SPECT. *IEEE Transactions on Medical Imaging*, 25(7):941–949.

[Ferrucci et al., 2015] Ferrucci, M., Leach, R. K., Giusca, C., Carmignato, S., and Dewulf, W. (2015). Towards geometrical calibration of X-ray computed tomography systems—a review. *Measurement Science and Technology*, 26(9):092003.

[Ferstl et al., 2019] Ferstl, S., Schwaha, T., Ruthensteiner, B., Hehn, L., Allner, S., Müller, M., Dierolf, M., Achterhold, K., and Pfeiffer, F. (2019). Nanoscopic X-ray tomography for correlative microscopy of a small meiofaunal sea-cucumber. In preparation.

[Fessler, 2009] Fessler, J. (2009). Analytical tomographic image reconstruction methods. *Image Reconstruction: Algorithms and Analysis*, 66:67.

[Fessler, 2000] Fessler, J. A. (2000). Statistical image reconstruction methods for transmission tomography. In Fitzpatrick, J. M. and Sonka, M., editors, *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis*, pages 1–70. SPIE Press, Bellingham.

[Fessler, 2003] Fessler, J. A. (2003). Analytical approach to regularization design for isotropic spatial resolution. In *Nuclear Science Symposium Conference Record, 2003 IEEE*, volume 3, pages 2022–2026. IEEE.

[FORBILD phantoms, 1999] FORBILD phantoms (1999). Forbild phantoms. http://www.imp.uni-erlangen.de/phantoms/.

[Förster et al., 1980] Förster, E., Goetz, K., and Zaumseil, P. (1980). Double crystal diffractometry for the characterization of targets for laser fusion experiments. *Kristall und Technik*, 15(8):937–945.

[Fourmont, 2003] Fourmont, K. (2003). Non-equispaced fast Fourier transforms with applications to tomography. *Journal of Fourier Analysis and Applications*, 9(5):431–450.

[Fuchs et al., 2000] Fuchs, T., Krause, J., Schaller, S., Flohr, T., and Kalender, W. A. (2000). Spiral interpolation algorithms for multislice spiral CT. II. Measurement and evaluation of slice sensitivity profiles and noise at a clinical multislice system. *IEEE Transactions on Medical Imaging*, 19(9):835–847.

[Gang et al., 2014] Gang, G. J., Stayman, J. W., Zbijewski, W., and Siewerdsen, J. H. (2014). Task-based detectability in CT image reconstruction by filtered backprojection and penalized likelihood estimation. *Medical Physics*, 41(8Part1):081902.

[Geman and Geman, 1984] Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741.

[Gilbert, 1972] Gilbert, P. (1972). Iterative methods for the three-dimensional reconstruction of an object from projections. *Journal of Theoretical Biology*, 36(1):105–117.

[Gonzalez and Woods, 2001] Gonzalez, R. C. and Woods, R. E. (2001). *Digital Image Processing*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2nd edition.

[Goodman, 2005] Goodman, J. W. (2005). Introduction to Fourier optics. *Introduction to Fourier optics, 3rd ed., by JW Goodman. Englewood, CO: Roberts & Co. Publishers, 2005*, 1.

[Gordon et al., 1970] Gordon, R., Bender, R., and Herman, G. T. (1970). Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and X-ray photography. *Journal of Theoretical Biology*, 29(3):471–481.

[Gradl et al., 2017] Gradl, R., Dierolf, M., Hehn, L., Günther, B., Yildirim, A. Ö., Gleich, B., Achterhold, K., Pfeiffer, F., and Morgan, K. S. (2017). Propagation-based phase-contrast X-ray imaging at a compact light source. *Scientific Reports*, 7(1):4908.

[Groso et al., 2006] Groso, A., Abela, R., and Stampanoni, M. (2006). Implementation of a fast method for high resolution phase contrast tomography. *Optics Express*, 14(18):8103–8110.

[Gullberg et al., 1990] Gullberg, G. T., Tsui, B. M., Crawford, C. R., Ballard, J. G., and Hagius, J. T. (1990). Estimation of geometrical parameters and collimator evaluation for cone beam tomography. *Medical Physics*, 17(2):264–272.

[Gureyev et al., 2009] Gureyev, T., Mayo, S., Myers, D., Nesterets, Y., Paganin, D., Pogany, A., Stevenson, A., and Wilkins, S. (2009). Refracting Röntgen's rays: propagation-based X-ray phase contrast for biomedical imaging. *Journal of Applied Physics*, 105(10):102005.

[Gureyev et al., 2014a] Gureyev, T., Mayo, S., Nesterets, Y. I., Mohammadi, S., Lockie, D., Menk, R., Arfelli, F., Pavlov, K., Kitchen, M., Zanconati, F., et al. (2014a). Investigation of the imaging quality of synchrotron-based phase-contrast mammographic tomography. *Journal of Physics D: Applied Physics*, 47(36):365401.

[Gureyev et al., 2004] Gureyev, T., Stevenson, A., Nesterets, Y. I., and Wilkins, S. (2004). Image deblurring by means of defocus. *Optics Communications*, 240(1-3):81–88.

[Gureyev et al., 2014b] Gureyev, T. E., Nesterets, Y. I., de Hoog, F., Schmalz, G., Mayo, S. C., Mohammadi, S., and Tromba, G. (2014b). Duality between noise and spatial resolution in linear systems. *Optics Express*, 22(8):9087–9094.

[Gureyev et al., 2017] Gureyev, T. E., Nesterets, Y. I., Kozlov, A., Paganin, D. M., and Quiney, H. M. (2017). On the "unreasonable" effectiveness of transport of intensity imaging and optical deconvolution. *Journal of the Optical Society of America*, 34(12):2251–2260.

[Gureyev and Nugent, 1996] Gureyev, T. E. and Nugent, K. A. (1996). Phase retrieval with the transport-of-intensity equation. ii. orthogonal series solution for nonuniform illumination. *Journal of the Optical Society of America A*, 13(8):1670–1682.

[Hager and Zhang, 2006] Hager, W. W. and Zhang, H. (2006). A survey of nonlinear conjugate gradient methods. *Pacific Journal of Optimization*, 2(1):35–58.

[Hartley and Zisserman, 2003] Hartley, R. and Zisserman, A. (2003). *Multiple View Geometry in Computer Vision*. Cambridge University Press, New York, NY, USA, 2 edition.

[Hashemi et al., 2017] Hashemi, S., Song, W. Y., Sahgal, A., Lee, Y., Huynh, C., Grouza, V., Nordström, H., Eriksson, M., Dorenlot, A., Régis, J. M., et al. (2017). Simultaneous deblurring and iterative reconstruction of CBCT for image guided brain radiosurgery. *Physics in Medicine & Biology*, 62(7):2521.

[Hehn, 2015] Hehn, L. (2015). High-performance algorithms for improved reconstruction of X-ray computed tomography. Master's thesis, Technical University of Munich, Germany.

[Hehn et al., 2019a] Hehn, L., Gradl, R., Dierolf, M., Morgan, K. S., Paganin, D., and Pfeiffer, F. (2019a). Model-based iterative reconstruction for propagation-based phase-contrast X-ray CT including models for the source and the detector. Submitted.

[Hehn et al., 2018a] Hehn, L., Gradl, R., Voss, A., Günther, B., Dierolf, M., Jud, C., Willer, K., Allner, S., Hammel, J. U., Hessler, R., et al. (2018a). Propagation-based phase-contrast tomography of a guinea pig inner ear with cochlear implant using a model-based iterative reconstruction algorithm. *Biomedical Optics Express*, 9(11):5330–5339.

[Hehn et al., 2018b] Hehn, L., Morgan, K., Bidola, P., Noichl, W., Gradl, R., Dierolf, M., Noël, P. B., and Pfeiffer, F. (2018b). Nonlinear statistical iterative reconstruction for propagation-based phase-contrast tomography. *APL Bioengineering*, 2(1):016105.

[Hehn et al., 2019b] Hehn, L., Tilley, S., Pfeiffer, F., and Stayman, J. W. (2019b). Blind deconvolution in model-based iterative reconstruction for CT using a normalized sparsity measure. *Physics in Medicine & Biology*, 64(21):215010.

[Hemberg et al., 2003] Hemberg, O., Otendal, M., and Hertz, H. (2003). Liquid-metal-jet anode electron-impact X-ray source. *Applied Physics Letters*, 83(7):1483–1485.

[Hofmann et al., 2014] Hofmann, C., Knaup, M., and Kachelrieß, M. (2014). Effects of ray profile modeling on resolution recovery in clinical CT. *Medical Physics*, 41(2):021907.

[Hu et al., 2011] Hu, Y., Xie, L., Luo, L., Nunes, J. C., and Toumoulin, C. (2011). L0 constrained sparse reconstruction for multi-slice helical CT reconstruction. *Physics in Medicine & Biology*, 56(4):1173.

[Huber, 2011] Huber, P. J. (2011). Robust statistics. In *International Encyclopedia of Statistical Science*, pages 1248–1251. Springer.

[Huh and Fessler, 2011] Huh, W. and Fessler, J. A. (2011). Iterative image reconstruction for dual-energy X-ray CT using regularized material sinogram estimates. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1512–1515.

[Ingal and Beliaevskaya, 1995] Ingal, V. N. and Beliaevskaya, E. A. (1995). X-ray plane-wave topography observation of the phase contrast from a non-crystalline object. *Journal of Physics D: Applied Physics*, 28(11):2314.

[Irvine et al., 2014] Irvine, S., Mokso, R., Modregger, P., Wang, Z., Marone, F., and Stampanoni, M. (2014). Simple merging technique for improving resolution in qualitative single image phase contrast tomography. *Optics Express*, 22(22):27257–27269.

[Joseph, 1982] Joseph, P. M. (1982). An improved algorithm for reprojecting rays through pixel images. *IEEE Transactions on Medical Imaging*, 1(3):192–196.

[Kak and Slaney, 1988] Kak, A. C. and Slaney, M. (1988). *Principles of computerized tomographic imaging*. IEEE press New York.

[Katsevich, 2002] Katsevich, A. (2002). Theoretically exact filtered backprojection-type inversion algorithm for spiral CT. *SIAM Journal on Applied Mathematics*, 62(6):2012–2026.

[Keller, 1962] Keller, J. B. (1962). Geometrical theory of diffraction. *Journal of the Optical Society of America A*, 52(2):116–130.

[Kim et al., 2015] Kim, D., Ramani, S., and Fessler, J. A. (2015). Combining ordered subsets and momentum for accelerated X-ray CT image reconstruction. *IEEE Transactions on Medical Imaging*, 34(1):167–178.

[Kingma and Ba, 2014] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

[Kingston et al., 2011] Kingston, A., Sakellariou, A., Varslot, T., Myers, G., and Sheppard, A. (2011). Reliable automatic alignment of tomographic projection data by passive auto-focus. *Medical Physics*, 38(9):4934–4945.

[Kitchen et al., 2017] Kitchen, M. J., Buckley, G. A., Gureyev, T. E., Wallace, M. J., Andres-Thio, N., Uesugi, K., Yagi, N., and Hooper, S. B. (2017). CT dose reduction factors in the thousands using X-ray phase contrast. *Scientific Reports*, 7(1):15953.

[Koehler et al., 2015] Koehler, T., Daerr, H., Martens, G., Kuhn, N., Löscher, S., van Stevendaal, U., and Roessl, E. (2015). Slit-scanning differential X-ray phase-contrast mammography: Proof-of-concept experimental studies. *Medical Physics*, 42(4):1959–1965.

[Kostenko et al., 2013] Kostenko, A., Batenburg, K. J., King, A., Offerman, S. E., and van Vliet, L. J. (2013). Total variation minimization approach in in-line X-ray phase-contrast tomography. *Optics Express*, 21(10):12185–12196.

[Krishnan et al., 2011] Krishnan, D., Tay, T., and Fergus, R. (2011). Blind deconvolution using a normalized sparsity measure. In *CVPR 2011*, pages 233–240.

[Kyriakou et al., 2008] Kyriakou, Y., Lapp, R., Hillebrand, L., Ertel, D., and Kalender, W. (2008). Simultaneous misalignment correction for approximate circular cone-beam computed tomography. *Physics in Medicine & Biology*, 53(22):6267.

[La Rivière et al., 2006] La Rivière, P. J., Bian, J., and Vargas, P. A. (2006). Penalized-likelihood sinogram restoration for computed tomography. *IEEE Transactions on Medical Imaging*, 25(8):1022–1036.

[Langer et al., 2008] Langer, M., Cloetens, P., Guigay, J.-P., and Peyrin, F. (2008). Quantitative comparison of direct phase retrieval algorithms in in-line phase tomography. *Medical Physics*, 35(10):4556–4566.

[Langer et al., 2014] Langer, M., Cloetens, P., Hesse, B., Suhonen, H., Pacureanu, A., Raum, K., and Peyrin, F. (2014). Priors for X-ray in-line phase tomography of heterogeneous objects. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 372(2010).

[Langer et al., 2012] Langer, M., Cloetens, P., Pacureanu, A., and Peyrin, F. (2012). X-ray in-line phase tomography of multimaterial objects. *Optics Letters*, 37(11):2151–2153.

[Lehmann and Casella, 2006] Lehmann, E. L. and Casella, G. (2006). *Theory of point estimation*. Springer Science & Business Media.

[Levin et al., 2011] Levin, A., Weiss, Y., Durand, F., and Freeman, W. T. (2011). Understanding blind deconvolution algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2354–2367.

[Liu and Nocedal, 1989] Liu, D. C. and Nocedal, J. (1989). On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528.

[Long et al., 2010] Long, Y., Fessler, J. A., and Balter, J. M. (2010). 3D forward and back-projection for X-ray CT using separable footprints. *IEEE Transactions on Medical Imaging*, 29(11):1839–1850.

[Longo et al., 2016] Longo, R., Arfelli, F., Bellazzini, R., Bottigli, U., Brez, A., Brun, F., Brunetti, A., Delogu, P., Di Lillo, F., Dreossi, D., et al. (2016). Towards breast tomography with synchrotron radiation at elettra: first images. *Physics in Medicine & Biology*, 61(4):1634.

[Man et al., 1999] Man, B. D., Nuyts, J., Dupont, P., Marchal, G., and Suetens, P. (1999). Metal streak artifacts in X-ray computed tomography: a simulation study. *IEEE Transactions on Nuclear Science*, 46(3):691–696.

[Mayo et al., 2003] Mayo, S., Davis, T., Gureyev, T., Miller, P., Paganin, D., Pogany, A., Stevenson, A., and Wilkins, S. (2003). X-ray phase-contrast microscopy and microtomography. *Optics Express*, 11(19):2289–2302.

[Mayo et al., 2007] Mayo, S., Miller, P., Gao, D., and Sheffield-Parker, J. (2007). Software image alignment for X-ray microtomography with submicrometre resolution using a sem-based X-ray microscope. *Journal of Microscopy*, 228(3):257–263.

[Mayo and Sexton, 2004] Mayo, S. C. and Sexton, B. (2004). Refractive microlens array for wave-front analysis in the medium to hard X-ray range. *Optics Letters*, 29(8):866–868.

[McJunkin et al., 2018] McJunkin, J. L., Durakovic, N., Herzog, J., and Buchman, C. A. (2018). Early outcomes with a slim, modiolar cochlear implant electrode array. *Otology & Neurotology*, 39(1):e28–e33.

[Miao et al., 1999] Miao, J., Charalambous, P., Kirz, J., and Sayre, D. (1999). Extending the methodology of X-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature*, 400(6742):342.

[Momose, 1995] Momose, A. (1995). Demonstration of phase-contrast X-ray computed tomography using an X-ray interferometer. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 352(3):622–628.

[Momose et al., 2003] Momose, A., Kawamoto, S., Koyama, I., Hamaishi, Y., Takai, K., and Suzuki, Y. (2003). Demonstration of X-ray talbot interferometry. *Japanese Journal of Applied Physics*, 42(7B):L866.

[Moré and Thuente, 1994] Moré, J. J. and Thuente, D. J. (1994). Line search algorithms with guaranteed sufficient decrease. *ACM Transactions on Mathematical Software (TOMS)*, 20(3):286–307.

[Morgan et al., 2010] Morgan, K., Siu, K. K., and Paganin, D. (2010). The projection approximation and edge contrast for X-ray propagation-based phase contrast imaging of a cylindrical edge. *Optics Express*, 18(10):9865–9878.

[Morgan et al., 2011] Morgan, K. S., Paganin, D. M., and Siu, K. K. W. (2011). Quantitative single-exposure X-ray phase contrast imaging using a single attenuation grid. *Optics Express*, 19(20):19781–19789.

[Morgan et al., 2012] Morgan, K. S., Paganin, D. M., and Siu, K. K. W. (2012). X-ray phase imaging with a paper analyzer. *Applied Physics Letters*, 100(12):124102.

[Muders and Hesser, 2014] Muders, J. and Hesser, J. (2014). Stable and robust geometric self-calibration for cone-beam CT using mutual information. *IEEE Transactions on Nuclear Science*, 61(1):202–217.

[Müller et al., 2017] Müller, M., de Sena Oliveira, I., Allner, S., Ferstl, S., Bidola, P., Mechlem, K., Fehringer, A., Hehn, L., Dierolf, M., Achterhold, K., et al. (2017). Myoanatomy of the velvet worm leg revealed by laboratory-based nanofocus X-ray source tomography. *Proceedings of the National Academy of Sciences*, 114(47):12378–12383.

[Natterer, 1986] Natterer, F. (1986). *The mathematics of computerized tomography*, volume 32. Siam.

[Nelder and Mead, 1965] Nelder, J. A. and Mead, R. (1965). A simplex method for function minimization. *The Computer Journal*, 7(4):308–313.

[Nesterets and Gureyev, 2014] Nesterets, Y. I. and Gureyev, T. E. (2014). Noise propagation in X-ray phase-contrast imaging and computed tomography. *Journal of Physics D: Applied Physics*, 47(10):105402.

[Nikl, 2006] Nikl, M. (2006). Scintillation detectors for X-rays. *Measurement Science and Technology*, 17(4):R37–R54.

[Nocedal, 1980] Nocedal, J. (1980). Updating quasi-newton matrices with limited storage. *Mathematics of Computation*, 35(151):773–782.

[Nocedal and Wright, 2006] Nocedal, J. and Wright, S. (2006). *Numerical optimization*. Springer Science & Business Media.

[Noël et al., 2013] Noël, P. B., Renger, B., Fiebich, M., Münzel, D., Fingerle, A. A., Rummeny, E. J., and Dobritz, M. (2013). Does iterative reconstruction lower CT radiation dose: evaluation of 15,000 examinations. *PLoS ONE*, 8(11):e81141.

[Noo et al., 2000] Noo, F., Clackdoyle, R., Mennessier, C., White, T. A., and Roney, T. J. (2000). Analytic method based on identification of ellipse parameters for scanner calibration in cone-beam tomography. *Physics in Medicine & Biology*, 45(11):3489.

[Nunez and Llacer, 1990] Nunez, J. and Llacer, J. (1990). A fast Bayesian reconstruction algorithm for emission tomography with entropy prior converging to feasible images. *IEEE Transactions on Medical Imaging*, 9(2):159–171.

[Nuyts et al., 2013] Nuyts, J., De Man, B., Fessler, J. A., Zbijewski, W., and Beekman, F. J. (2013). Modelling the physics in the iterative reconstruction for transmission computed tomography. *Physics in Medicine & Biology*, 58(12):R63.

[O'Connell et al., 2017] O'Connell, B. P., Hunter, J. B., Haynes, D. S., Holder, J. T., Dedmon, M. M., Noble, J. H., Dawant, B. M., and Wanna, G. B. (2017). Insertion depth impacts speech perception and hearing preservation for lateral wall electrodes. *The Laryngoscope*, 127(10):2352–2357.

[Oliphant, 2006] Oliphant, T. E. (2006). *A guide to NumPy*, volume 1. Trelgol Publishing USA.

[Olivo et al., 2001] Olivo, A., Arfelli, F., Cantatore, G., Longo, R., Menk, R. H., Pani, S., Prest, M., Poropat, P., Rigon, L., Tromba, G., Vallazza, E., and Castelli, E. (2001). An innovative digital imaging set-up allowing a low-dose approach to phase contrast applications in the medical field. *Medical Physics*, 28(8):1610–1619.

[Olivo and Speller, 2007] Olivo, A. and Speller, R. (2007). A coded-aperture technique allowing X-ray phase contrast imaging with conventional sources. *Applied Physics Letters*, 91(7):074106.

[O'Sullivan, 1985] O'Sullivan, J. (1985). A fast sinc function gridding algorithm for Fourier inversion in computer tomography. *IEEE Transactions on Medical Imaging*, 4(4):200–207.

[Pack et al., 2004] Pack, J. D., Noo, F., and Kudo, H. (2004). Investigation of saddle trajectories for cardiac CT imaging in cone-beam geometry. *Physics in Medicine & Biology*, 49(11):2317.

[Paganin, 2006] Paganin, D. (2006). *Coherent X-ray optics*. 6. Oxford University Press on Demand.

[Paganin et al., 2002] Paganin, D., Mayo, S., Gureyev, T. E., Miller, P. R., and Wilkins, S. W. (2002). Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object. *Journal of Microscopy*, 206(1):33–40.

[Paganin and Morgan, 2019] Paganin, D. M. and Morgan, K. S. (2019). X-ray fokker–planck equation for paraxial imaging. *arXiv preprint arXiv:1908.01473*.

[Panetta et al., 2008] Panetta, D., Belcari, N., Del Guerra, A., and Moehrs, S. (2008). An optimization-based method for geometrical calibration in cone-beam CT without dedicated phantoms. *Physics in Medicine & Biology*, 53(14):3841.

[Patel et al., 2009] Patel, V., Chityala, R., Hoffmann, K., Ionita, C., Bednarek, D., and Rudin, S. (2009). Self-calibration of a cone-beam micro-CT system. *Medical Physics*, 36(1):48–58.

[Peatross and Ware, 2015] Peatross, J. and Ware, M. (2015). *Physics of light and optics*. Brigham Young University, Department of Physics.

[Perrone and Favaro, 2014] Perrone, D. and Favaro, P. (2014). Total variation blind deconvolution: The devil is in the details. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2909–2916.

[Peters, 1981] Peters, T. (1981). Algorithms for fast back-and re-projection in computed tomography. *IEEE Transactions on Nuclear Science*, 28(4):3641–3647.

[Pfeiffer et al., 2005] Pfeiffer, F., Bunk, O., Schulze-Briese, C., Diaz, A., Weitkamp, T., David, C., Van Der Veen, J., Vartanyants, I., and Robinson, I. (2005). Shearing interferometer for quantifying the coherence of hard X-ray beams. *Physical Review Letters*, 94(16):164801.

[Pfeiffer et al., 2007] Pfeiffer, F., Kottler, C., Bunk, O., and David, C. (2007). Hard X-ray phase tomography with low-brilliance sources. *Physical Review Letters*, 98(10):108105.

[Pfeiffer et al., 2006] Pfeiffer, F., Weitkamp, T., Bunk, O., and David, C. (2006). Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources. *Nature Physics*, 2(4):258.

[Qi and Leahy, 2006] Qi, J. and Leahy, R. M. (2006). Iterative reconstruction techniques in emission computed tomography. *Physics in Medicine & Biology*, 51(15):R541.

[Radon, 1917] Radon, J. (1917). Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Verh. Sachs. Akad. Wiss. Leipzig, Math Phys Klass*, 69.

[Rinkel et al., 2008] Rinkel, J., Dillon, W. P., Funk, T., Gould, R., and Prevrhal, S. (2008). Computed tomographic metal artifact reduction for the detection and quantitation of small features near large metallic implants: a comparison of published methods. *Journal of Computer Assisted Tomography*, 32(4):621–629.

[Rit et al., 2014] Rit, S., Oliva, M. V., Brousmiche, S., Labarbe, R., Sarrut, D., and Sharp, G. C. (2014). The reconstruction toolkit (RTK), an open-source cone-beam CT reconstruction toolkit based on the insight toolkit (ITK). In *Journal of Physics: Conference Series*, volume 489, page 012079. IOP Publishing.

[Riviere and Vargas, 2008] Riviere, P. J. L. and Vargas, P. (2008). Correction for resolution nonuniformities caused by anode angulation in computed tomography. *IEEE Transactions on Medical Imaging*, 27(9):1333–1341.

[Rodenburg et al., 2007] Rodenburg, J., Hurst, A., Cullis, A., Dobson, B., Pfeiffer, F., Bunk, O., David, C., Jefimovs, K., and Johnson, I. (2007). Hard-X-ray lensless imaging of extended objects. *Physical Review Letters*, 98(3):034801.

[Röntgen, 1895] Röntgen, W. C. (1895). Über eine neue Art von Strahlen. *Aus den Sitzungsberichten der Würzburger Physik.-medic. Gesellschaft. Würzburg*, pages 137–147.

[Rottman et al., 2015a] Rottman, C., Hinkle, J., Cheryauka, A., Whitaker, R., and Joshi, S. (2015a). Joint cone-beam reconstruction and geometry estimation for mobile c-arms. In *The 13th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, pages 697–700.

[Rottman et al., 2015b] Rottman, C., McBride, L., Cheryauka, A., Whitaker, R., and Joshi, S. (2015b). Mobile C-arm 3D reconstruction in the presence of uncertain geometry. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 692–699. Springer.

[Rudin et al., 1992] Rudin, L. I., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268.

[Ruhlandt et al., 2014] Ruhlandt, A., Krenkel, M., Bartels, M., and Salditt, T. (2014). Three-dimensional phase retrieval in propagation-based phase-contrast imaging. *Physical Review A*, 89(3):033847.

[Schaff, 2017] Schaff, F. (2017). *Directional small-angle X-ray scattering computed tomography*. PhD thesis, Technical University of Munich, Germany.

[Schaller et al., 2000] Schaller, S., Flohr, T., Klingenbeck, K., Krause, J., Fuchs, T., and Kalender, W. A. (2000). Spiral interpolation algorithm for multislice spiral CT. I. Theory. *IEEE Transactions on Medical Imaging*, 19(9):822–834.

[Scherer et al., 2015] Scherer, K., Willer, K., Gromann, L., Birnbacher, L., Braig, E., Grandl, S., Sztrókay-Gaul, A., Herzen, J., Mayr, D., Hellerhoff, K., et al. (2015). Toward clinically compatible phase-contrast mammography. *PLoS ONE*, 10(6):e0130776.

[Schomberg and Timmer, 1995] Schomberg, H. and Timmer, J. (1995). The gridding method for image reconstruction by Fourier transformation. *IEEE Transactions on Medical Imaging*, 14(3):596–607.

[Schoonjans et al., 2011] Schoonjans, T., Brunetti, A., Golosio, B., del Rio, M. S., Solé, V. A., Ferrero, C., and Vincze, L. (2011). The xraylib library for X-ray–matter interactions. Recent developments. *Spectrochimica Acta Part B: Atomic Spectroscopy*, 66(11–12):776–784.

[Seibt, 2006] Seibt, P. (2006). *Algorithmic Information Theory*. Springer.

[Shepp and Logan, 1974a] Shepp, L. and Logan, B. (1974a). Reconstructing interior head tissue from X-ray transmissions. *IEEE Transactions on Nuclear Science*, 21(1):228–236.

[Shepp and Logan, 1974b] Shepp, L. A. and Logan, B. F. (1974b). The Fourier reconstruction of a head section. *IEEE Transactions on Nuclear Science*, 21(3):21–43.

[Shewchuk et al., 1994] Shewchuk, J. R. et al. (1994). An introduction to the conjugate gradient method without the agonizing pain.

[Sidky et al., 2010] Sidky, E. Y., Anastasio, M. A., and Pan, X. (2010). Image reconstruction exploiting object sparsity in boundary-enhanced X-ray phase-contrast tomography. *Optics Express*, 18(10):10404–10422.

[Sidky et al., 2006] Sidky, E. Y., Kao, C.-M., and Pan, X. (2006). Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT. *Journal of X-ray Science and Technology*, 14(2):119–139.

[Sidky and Pan, 2008] Sidky, E. Y. and Pan, X. (2008). Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology*, 53(17):4777.

[Snigirev et al., 1995] Snigirev, A., Snigireva, I., Kohn, V., Kuznetsov, S., and Schelokov, I. (1995). On the possibilities of X-ray phase contrast microimaging by coherent high-energy synchrotron radiation. *Review of Scientific Instruments*, 66(12):5486–5492.

[Stayman et al., 2013] Stayman, J. W., Dang, H., Otake, Y., Zbijewski, W., Noble, J., Dawant, B., Labadie, R., Carey, J. P., and Siewerdsen, J. H. (2013). Overcoming nonlinear partial volume effects in known-component reconstruction of cochlear implants. In *Proc. SPIE 6886, Medical Imaging 2013: Physics of Medical Imaging*, page 86681L.

[Stayman and Fessler, 2000] Stayman, J. W. and Fessler, J. A. (2000). Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Transactions on Medical Imaging*, 19(6):601–615.

[Sziklas and Siegman, 1974] Sziklas, E. A. and Siegman, A. (1974). Diffraction calculations using fast fourier transform methods. *Proceedings of the IEEE*, 62(3):410–412.

[Taguchi and Iwanczyk, 2013] Taguchi, K. and Iwanczyk, J. S. (2013). Vision 20/20: Single photon counting X-ray detectors in medical imaging. *Medical Physics*, 40(10):100901.

[Teague, 1983] Teague, M. R. (1983). Deterministic phase retrieval: a green's function solution. *Journal of the Optical Society of America*, 73(11):1434–1441.

[Thibault et al., 2007] Thibault, J.-B., Sauer, K. D., Bouman, C. A., and Hsieh, J. (2007). A three-dimensional statistical approach to improved image quality for multislice helical CT. *Medical Physics*, 34(11):4526–4544.

[Thompson et al., 2019] Thompson, D. A., Nesterets, Y. I., Pavlov, K. M., and Gureyev, T. E. (2019). Fast three-dimensional phase retrieval in propagation-based X-ray tomography. *Journal of Synchrotron Radiation*, 26(3):825–838.

[Tilley et al., 2018a] Tilley, S., Jacobson, M., Cao, Q., Brehler, M., Sisniega, A., Zbijewski, W., and Stayman, J. W. (2018a). Penalized-likelihood reconstruction with high-fidelity measurement models for high-resolution cone-beam imaging. *IEEE Transactions on Medical Imaging*, 37(4):988–999.

[Tilley et al., 2016a] Tilley, S., Siewerdsen, J. H., and Stayman, J. W. (2016a). Model-based iterative reconstruction for flat-panel cone-beam CT with focal spot blur, detector blur, and correlated noise. *Physics in Medicine & Biology*, 61(1):296–319.

[Tilley et al., 2016b] Tilley, S., Siewerdsen, J. H., Zbijewski, W., and Stayman, J. W. (2016b). Nonlinear statistical reconstruction for flat-panel cone-beam CT with blur and correlated noise models. In *Medical Imaging 2016: Physics of Medical Imaging*, volume 9783, page 97830R. International Society for Optics and Photonics.

[Tilley et al., 2016c] Tilley, S., Zbijewski, W., Siewerdsen, J. H., and Stayman, J. W. (2016c). Modeling shift-variant X-ray focal spot blur for high-resolution flat-panel cone-beam CT. In *Conference proceedings. International Conference on Image Formation in X-ray Computed Tomography*, volume 2016, page 463. NIH Public Access.

[Tilley et al., 2018b] Tilley, S. W., Zbijewski, W., and Stayman, J. W. (2018b). Model-based material decomposition with a penalized nonlinear least-squares CT reconstruction algorithm. *Physics in Medicine & Biology*, page 035005.

[Tomasi and Manduchi, 1998] Tomasi, C. and Manduchi, R. (1998). Bilateral filtering for gray and color images. In *Sixth International Conference on Computer Vision*, pages 839–846. IEEE.

[Turbell, 2001] Turbell, H. (2001). *Cone-beam reconstruction using filtered backprojection*. PhD thesis, Linköping University Electronic Press.

[Tuy, 1983] Tuy, H. K. (1983). An inversion formula for cone-beam reconstruction. *SIAM Journal on Applied Mathematics*, 43(3):546–552.

[Tward and Siewerdsen, 2008] Tward, D. J. and Siewerdsen, J. H. (2008). Cascaded systems analysis of the 3D noise transfer characteristics of flat-panel cone-beam CT. *Medical Physics*, 35(12):5510–5529.

[Ullherr and Zabler, 2015] Ullherr, M. and Zabler, S. (2015). Correcting multi material artifacts from single material phase retrieved holo-tomograms with a simple 3D Fourier method. *Optics Express*, 23(25):32718–32727.

[Vidal et al., 2005] Vidal, F. P., Létang, J. M., Peix, G., and Cloetens, P. (2005). Investigation of artefact sources in synchrotron microtomography via virtual X-ray imaging. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 234(3):333–348.

[Von Smekal et al., 2004] Von Smekal, L., Kachelrieß, M., Stepina, E., and Kalender, W. A. (2004). Geometric misalignment and calibration in cone-beam tomography: Geometric misalignment and calibration in cone-beam tomography. *Medical Physics*, 31(12):3242–3266.

[Wang and Qi, 2012] Wang, G. and Qi, J. (2012). Penalized likelihood pet image reconstruction using patch-based edge-preserving regularization. *IEEE Transactions on Medical Imaging*, 31(12):2194–2204.

[Wanna et al., 2014] Wanna, G. B., Noble, J. H., Carlson, M. L., Gifford, R. H., Dietrich, M. S., Haynes, D. S., Dawant, B. M., and Labadie, R. F. (2014). Impact of electrode design and surgical approach on scalar location and cochlear implant outcomes. *The Laryngoscope*, 124(S6).

[Weber et al., 2017] Weber, L., Hänsch, A., Wolfram, U., Pacureanu, A., Cloetens, P., Peyrin, F., Rit, S., and Langer, M. (2017). Registration of phase-contrast images in propagation-based X-ray phase tomography. *Journal of Microscopy*, pages 36–47.

[Weitkamp et al., 2005] Weitkamp, T., Diaz, A., David, C., Pfeiffer, F., Stampanoni, M., Cloetens, P., and Ziegler, E. (2005). X-ray phase imaging with a grating interferometer. *Optics Express*, 13(16):6296–6304.

[Wilde et al., 2016] Wilde, F., Ogurreck, M., Greving, I., Hammel, J. U., Beckmann, F., Hipp, A., Lottermoser, L., Khokhriakov, I., Lytaev, P., Dose, T., Burmester, H., Müller, M., and Schreyer, A. (2016). Micro-CT at the imaging beamline P05 at PETRA III. *AIP Conference Proceedings*, 1741(1):030035.

[Wilkins et al., 1996] Wilkins, S., Gureyev, T. E., Gao, D., Pogany, A., and Stevenson, A. (1996). Phase-contrast imaging using polychromatic hard X-rays. *Nature*, 384(6607):335.

[Wilkins et al., 2014] Wilkins, S. W., Nesterets, Y. I., Gureyev, T. E., Mayo, S. C., Pogany, A., and Stevenson, A. W. (2014). On the evolution and relative merits of hard X-ray phase-contrast imaging methods. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 372(2010):20130021.

[Willmott, 2011] Willmott, P. (2011). *An introduction to synchrotron radiation: techniques and applications*. John Wiley & Sons.

[Wolf, 1982] Wolf, E. (1982). New theory of partial coherence in the space–frequency domain. part i: spectra and cross spectra of steady-state sources. *Journal of the Optical Society of America A*, 72(3):343–351.

[Wolf, 2007] Wolf, E. (2007). *Introduction to the theory of coherence and polarization of light*. Cambridge University Press.

[Wu and Fessler, 2011] Wu, M. and Fessler, J. A. (2011). GPU acceleration of 3D forward and backward projection using separable footprints for X-ray CT image reconstruction. In *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med*, volume 6, page 021911. Citeseer.

[Xie et al., 2017] Xie, X., McGaffin, M. G., Long, Y., Fessler, J. A., Wen, M., and Lin, J. (2017). Accelerating separable footprint (SF) forward and back projection on GPU. In *Medical Imaging 2017: Physics of Medical Imaging*, volume 10132, page 101322S. International Society for Optics and Photonics.

[Xu et al., 2012] Xu, Q., Yu, H., Mou, X., Zhang, L., Hsieh, J., and Wang, G. (2012). Low-dose X-ray CT reconstruction via dictionary learning. *IEEE Transactions on Medical Imaging*, 31(9):1682–1697.

[Ye and Wang, 2005] Ye, Y. and Wang, G. (2005). Filtered backprojection formula for exact image reconstruction from cone-beam data along a general scanning curve. *Medical Physics*, 32(1):42–48.

[Yu et al., 2000] Yu, D. F., Fessler, J. A., and Ficaro, E. P. (2000). Maximum-likelihood transmission image reconstruction for overlapping transmission beams. *IEEE Transactions on Medical Imaging*, 19(11):1094–1105.

[Zeng, 2010] Zeng, G. L. (2010). *Medical image reconstruction: a conceptual tutorial.* Springer.

[Zeng and Gullberg, 2000] Zeng, G. L. and Gullberg, G. T. (2000). Unmatched projector/backprojector pairs in an iterative reconstruction algorithm. *IEEE Transactions on Medical Imaging*, 19(5):548–555.

[Zernike, 1942] Zernike, F. (1942). Phase contrast, a new method for the microscopic observation of transparent objects. *Physica*, 9(7):686–698.

[Zernike, 1955] Zernike, F. (1955). How I discovered phase contrast. *Science*, 121(3141):345–349.

[Zhang et al., 2014a] Zhang, H., Han, H., Wang, J., Ma, J., Liu, Y., Moore, W., and Liang, Z. (2014a). Deriving adaptive MRF coefficients from previous normal-dose CT scan for low-dose image reconstruction via penalized weighted least-squares minimization. *Medical Physics*, 41(4):041916.

[Zhang et al., 2014b] Zhang, H., Ouyang, L., Ma, J., Huang, J., Chen, W., and Wang, J. (2014b). Noise correlation in CBCT projection data and its application for noise reduction in low-dose CBCT. *Medical Physics*, 41(3):031906.

[Zhang et al., 2018] Zhang, H., Wang, J., Zeng, D., Tao, X., and Ma, J. (2018). Regularization strategies in statistical image reconstruction of low-dose X-ray CT: A review. *Medical Physics*, 45(10):e886–e907.

[Zheng et al., 2018] Zheng, J., Fessler, J. A., and Chan, H.-P. (2018). Detector blur and correlated noise modeling for digital breast tomosynthesis reconstruction. *IEEE Transactions on Medical Imaging*, 37(1):116–127.

List of abbreviations

BAC	Bronnikov-aided correction
BFGS	Broyden–Fletcher–Goldfarb–Shanno
CCD	charged-coupled device
CG	conjugate gradient
CT	computed tomography
CTF	contrast transfer function
DFT	discrete Fourier transform
FBP	filtered backprojection
FDK	Feldkamp, Davis, Kress
FFT	fast Fourier transform
FOV	field-of-view
FT	Fourier transform
FWHM	full width at half maximum
GB-CT	grating-based phase-contrast computed tomography
GBI	grating-based phase-contrast imaging
IFT	inverse Fourier transform
L-BFGS	limited-memory BFGS
MAP	maximum a posteriori
MBA	modified Bronnikov algorithm
MBIR	model-based iterative reconstruction
MRF	Markov random field
MSE	mean squared error
MuCLS	Munich Compact Light Source
NLPV	nonlinear partial volume effects
NPS	noise power spectrum
PAG	Paganin's single-material phase-retrieval algorithm
PB-CT	propagation-based phase-contrast computed tomography
PBI	propagation-based phase-contrast imaging

PCI	phase-contrast imaging
PE	polyethylene
PMMA	poly(methyl methacrylate)
PSF	point spread function
RMSE	root MSE
ROI	region of interest
SAD	source-to-axis distance
SDD	source-to-detector distance
SIR	statistical iterative reconstruction
SPECT	single-photon emission computed tomography
SR	synchrotron radiation
STIR	statistical TIE-based iterative reconstruction
TIE	transport-of-intensity equation
TPE	transport-of-phase equation
TV	total-variation

List of algorithms

List of figures

List of publications and scientific presentations

First-authored publications (peer-reviewed)

- **Hehn, L.; Tilley, S.; Pfeiffer, F.; Stayman, J. W.:** Blind deconvolution in model-based iterative reconstruction for CT using a normalized sparsity measure. Physics in Medicine & Biology **64** (21), 2019, 215010

- **Hehn, L.; Gradl, R.; Voss, A.; Günther, B.; Dierolf, M.; Jud, C.; Willer, K.; Allner, S.; Hammel, J. U.; Hessler, R.; Morgan, K. S.; Herzen, J.; Hemmert, W.; Pfeiffer, F.:** Propagation-based phase-contrast tomography of a guinea pig inner ear with cochlear implant using a model-based iterative reconstruction algorithm. Biomedical Optics Express **9** (11), 2018, 5330

- **Hehn, L.; Morgan, K. S.; Bidola, P.; Noichl, W.; Gradl, R.; Dierolf, M.; Noël, P. B.; Pfeiffer, F.:** Nonlinear statistical iterative reconstruction for propagation-based phase-contrast tomography. APL Bioengineering **2** (1), 2018, 016105

Submitted manuscript:

- **Hehn, L.; Gradl, R.; Dierolf, M.; Morgan, K. S.; Paganin, D. M.; Pfeiffer, F.:** Model-based iterative reconstruction for propagation-based phase-contrast X-ray CT including models for the source and the detector. IEEE Transactions on Medical Imaging, *submitted*

Co-authored publications (peer-reviewed)

- **Yang, L.; Gradl, R.; Dierolf, M.; Möller, W.; Kutschke, D.; Feuchtinger, A.; Hehn, L.; Donnelley, M.; Günther, B.; Achterhold, K.; Walch, A.; Stoeger, T.; Razansky, D.; Pfeiffer, F.; Morgan, K. S.; Schmid, O.**: Multimodal Precision Imaging of Pulmonary Nanoparticle Delivery in Mice: Dynamics of Application, Spatial Distribution, and Dosimetry. Small, 2019, 1904112

- **Gradl, R.; Dierolf, M.; Yang, L.; Hehn, L.; Günther, B.; Möller, W.; Kutschke, D.; Stoeger, T.; Gleich, B.; Achterhold, K.; Donnelley, M.; Pfeiffer, F.; Schmid, O.; Morgan, K. S.**: Visualizing treatment delivery and deposition in mouse lungs using in vivo X-ray imaging. Journal of Controlled Release **307**, 2019, 282-291

- **Günther, B.; Hehn, L.; Jud, C.; Hipp, A.; Dierolf, M.; Pfeiffer, F.:** Full-field structured-illumination super-resolution X-ray transmission microscopy. Nature Communications **10** (1), 2019, 2494

- **Gross, V.; Müller, M.; Hehn, L.; Ferstl, S.; Allner, S.; Dierolf, M.; Achterhold, K.; Mayer, G.; Pfeiffer, F.:** X-ray imaging of a water bear offers a new look at tardigrade internal anatomy. Zoological Letters **5** (1), 2019, 14

- **Gradl, R.; Morgan, K. S.; Dierolf, M.; Jud, C.; Hehn, L.; Günther, B.; Moller, W.; Kutschke, D.; Yang, L.; Stoeger, T.; Pfeiffer, D.; Gleich, B.; Achterhold, K.; Schmid, O.; Pfeiffer, F.:** Dynamic in vivo chest X-ray dark-field imaging in mice. IEEE Transactions on Medical Imaging **38** (2), 2019, 649-656

- **Gradl, R.; Dierolf, M.; Günther, B.; Hehn, L.; Möller, W.; Kutschke, D.; Yang, L.; Donnelley, M.; Murrie, R.; Erl, A.; Stoeger, T.; Gleich, B.; Achterhold, K.; Schmid, O,; Pfeiffer, F.; Morgan, K. S.:** In vivo Dynamic Phase-Contrast X-ray Imaging using a Compact Light Source. Scientific Reports **8** (1), 2018, 6788

- **Götzfried, J.; Döpp, A.; Gilljohann, M.; Ding, H.; Schindler, S.; Wenz, J.; Hehn, L.; Pfeiffer, F.; Karsch, S.:** Research towards high-repetition rate laser-driven X-ray sources for imaging applications. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **909**, 2018, 286-289

- **Döpp, A.; Hehn, L.; Götzfried, J.; Wenz, J.; Gilljohann, M.; Ding, H.; Schindler, S.; Pfeiffer, F.; Karsch, S.:** Quick X-ray microtomography using a laser-driven betatron source. Optica **5** (2), 2018, 199-203

- **Müller, M.; de Sena Oliveira, I.; Allner, S.; Ferstl, S.; Bidola, P.; Mechlem, K.; Fehringer, A.; Hehn, L.; Dierolf, M.; Achterhold, K.; Gleich, B.; Hammel, J. U.; Jahn, H.; Mayer, G.; Pfeiffer, F.:** Myoanatomy of the velvet worm leg revealed by laboratory-based nanofocus X-ray source tomography. Proceedings of the National Academy of Sciences **114** (47), 2017, 12378-12383

- **Gradl, R.; Dierolf, M.; Hehn, L.; Günther, B.; Yildirim, A. Ö.; Gleich, B.; Achterhold, K.; Pfeiffer, F.; Morgan, K. S.:** Propagation-based phase-contrast X-ray imaging at a compact light source. Scientific Reports **7** (1), 2017, 4908

First-authored conference proceedings

- **Hehn, L.; Morgan, K. S.; Dierolf, M.; Gradl, R.; Noël, P. B.; Pfeiffer, F.:** Model-based iterative reconstruction for propagation-based phase-contrast computed tomography, suitable for laboratory sources. The 5th International Conference on Image Formation in X-Ray Computed Tomography (CT Meeting), 2018

- **Hehn, L.; Morgan, K. S.; Bidola, P.; Noichl, W.; Gradl, R.; Dierolf, M.; Pfeiffer, F.; Noël, P. B.:** Statistical iterative reconstruction for propagation-based phase-contrast tomography in the near-field regime. The 14th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine (Fully3D), 2017

Oral presentations at international conferences

- **June 2017:** Statistical iterative reconstruction for propagation-based phase-contrast tomography in the near-field regime. The 14^{th} International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine (Fully3D), Xi'an, China

Poster presentations at international conferences

- **May, 2018:** Model-based iterative reconstruction for propagation-based phase-contrast computed tomography, suitable for laboratory sources. The 5^{th} International Conference on Image Formation in X-Ray Computed Tomography (CT Meeting), Salt Lake City, UT, USA

Acknowledgments

Last but not least, I would like to acknowledge the help and support of my supervisors, collaborators, colleagues and friends, without whom this endeavor would not have been possible.

First of all, I would like to thank my supervisor **Franz Pfeiffer** most sincerely for the opportunity to pursue my doctorate at the Chair of Biomedical Physics. Thank you for providing a positive working environment and an excellent infrastructure, which includes various exceptional imaging setups. I would also like to thank you for supporting my work, my research exchange to Johns Hopkins University and various other collaborations.

I am extremely thankful to my mentor **Martin Dierolf** for always taking the time to discuss any problems encountered throughout my doctoral studies and your clear thoughts on solving these problems. I also highly valued your in-depth feedback on various manuscripts and your help during experiments.

Thank you, **Kaye Morgan**, for being very supportive and for your invaluable advice on my work and all your help on the sometimes laborious peer-review process. Thank you for always being positive and motivating. I would also like to thank **David Paganin** for your time discussing various topics in greatest detail. Your broad and in-depth knowledge on these topics and feedback on my work have been invaluable.

I would like to thank **Web Stayman** for giving me the opportunity to work at an exceptional lab at Johns Hopkins University. You have been very supportive and I have learned a lot from you. I would also like to thank your student **Steven Tilley** who was extremely helpful throughout this research exchange. I'm thankful to the whole team for providing a positive and productive working environment, including **Grace Gang**, **Hao Zhang**, **Wenying Wang**, **Gabriela Rodal** and **Andrew Mao**.

Moreover, I am very thankful to **Peter Noël**. You have been a very supportive mentor. In particular, thank you for encouraging me to go to various conferences and to undertake the research exchange, all of which has contributed significantly to this work.

Thank you, **Regine Gradl**, for all the time we have spent together at the lab and in front of a computer. I am now confident to come near non-simulated physics experiments. And thank you for always being a step ahead in writing your thesis, which kept me motivated. I would like to thank **Benedikt Günther** for many fun and valuable discussions as well as exciting excursions to synchrotrons. Thank you, **Christoph Jud**, for fun beamtimes and the rest of the MuCLS team including **Juanjuan Huang**, **Stephanie Kulpe**, **Elena Eggl**, **Karin Burger**, **Johannes Brantl**, **David Cont**, **Johannes Melcher** and **Lisa Heck**.

An integral part of the pleasant working environment is due to the colleagues and friends at our office. I would like to thank **Sebastian Allner** for inexhaustible advice for all kinds of processing and reconstruction problems, which have been invaluable for this work. I am very thankful to **Wolfgang Noichl** for extremely fun and in-depth discussions or figuring out names of our colleagues. Working with you on various optimization problems has been very productive and enjoyable and contributed significantly to this work. Thank you, **Lorenz Birnbacher**, **Max von Teuffenbach**, **Clemens Schmid**, **Andreas Fehringer**, **Mathias Marschner** and **Fabio De Marco** for all of your help and this positive atmosphere.

I would like to thank **Simone Ferstl** and **Mark Müller** for providing fascinating datasets and for many very enjoyable discussions and collaborations. I would also like to thank **Katharina Scheidt** for measuring samples used within this work as well as **Madleen Busse** for sample preparation.

Thank you **Klaus Achterhold** for always taking the time to discuss numerous topics and mathematical derivations. In particular, thank you for your thorough feedback on this thesis. I am also thankful to **Julia Herzen** who gave valuable advice on various topics throughout my doctorate and for organizing experiments at the synchrotron. Thank you both for making the teaching part of my doctorate a very pleasant experience.

I am very thankful to the many more collaborators and colleagues I had the chance to work with. I am thankful to **Pidassa Bidola** for being very helpful in particular regarding measuring samples. Thank you, **Alen Begic**, for helping me to create various samples. I would like to thank **Andreas Döpp** for a very pleasant collaboration, **Konstantin Willer** and **Jörg Hammel** for your help at the synchrotron. I would further like to thank **Werner Hemmert**, his student **Andrej Voss** and his collaborator **Roland Hessler** for their support and for an extremely pleasant collaboration. Thank you, **Tobias Lasser**, for valuable and motivating discussions and your mathematical rigor on various occasions throughout my doctorate.

I would like to thank all colleagues from the neighboring office, including **Eva Prexl**, **Thorsten Sellerer**, **Korbinian Mechlem**, **Manuel Schultheiß**, **Josef Scholz** as well as **Kirsten Taphorn**, in particular for fun times during coffee breaks and lunch.

Thank you, **Nelly De Leiris** for your help regarding administrative work. I am also thankful to **Veronica Bodeck** for being very helpful dealing with administrative work particularly regarding my research exchange and business trips.

I would also like to thank the rest of the chair for contributing to this very positive working environment.

Lastly, I would like to thank my family and Sarah for their continuous support in every aspect of my life.