Lehrstuhl für Biomedizinische Physik
Technische Universität München

Towards Low-Dose Spectral Phase-Contrast X-ray Imaging

September 2020

Korbinian Mechlem
PhD thesis

Supervisor:
Prof. Dr. Franz Pfeiffer

Technische Universität München

Technische Universität München
Physik Department
Lehrstuhl für Biomedizinische Physik

Towards low-dose spectral phase-contrast X-ray imaging

Korbinian Maximilian Mechlem

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

Die Dissertation wurde am 30.09.2020 bei der Technischen Universität München eingereicht und durch die Fakultät für Physik am 20.11.2020 angenommen.

Abstract

Spectral X-ray imaging with energy-resolving photon-counting detectors has gained increasing interest for medical imaging applications in recent years. By exploiting information about the energy-dependent attenuation, different materials can be distinguished and information about the chemical composition of an object is obtained. However, the high noise level of the material selective images is still a fundamental problem, especially for medical imaging applications where a low radiation dose is desired. In this work, we investigate how photon-counting detectors can be used to combine spectral X-ray imaging with the emerging grating-based differential phase-contrast X-ray imaging technology. In addition to the attenuation, this technique measures the phase shift that an X-ray wave exhibits when traversing an object. Numerical simulations of radiography and computed tomography applications demonstrate that the combination of these two imaging techniques (spectral differential phase-contrast imaging) benefits from the strengths of the individual methods while the weaknesses are mitigated. Quantitatively accurate material selective images are obtained and the noise level is strongly reduced, compared to spectral X-ray imaging. Moreover, we develop several processing and image reconstruction techniques for spectral and spectral differential phase-contrast imaging. Our results demonstrate that the image quality can be strongly improved by accounting for photon-counting detector imperfections, using statistical noise models and incorporating prior knowledge in a Bayesian image reconstruction framework. Although many results need be confirmed by additional experiments, we believe that this work shows the potential of spectral differential phase-contrast imaging with photon-counting detectors in the fields of medical imaging and non-destructive testing, especially for low-dose applications.

Zusammenfassung

Die spektrale Röntgenbildgebung mit energieauflösenden photonenzählenden Detektoren hat in den letzten Jahren zunehmend an Interesse für medizinische Bildgebungsanwendungen gewonnen. Durch die Ausnutzung von Informationen über die energieabhängige Abschwächung der einfallenden Röntgenstrahlung können verschiedene Materialien unterschieden und Informationen über die chemische Zusammensetzung eines Objekts gewonnen werden. Der hohe Rauschpegel der materialselektiven Bilder ist jedoch nach wie vor ein grundlegendes Problem, insbesondere für medizinische Bildgebungsanwendungen, bei denen eine niedrige Strahlendosis erwünscht ist. In dieser Arbeit untersuchen wir, wie photonenzählende Detektoren verwendet werden können, um die spektrale Röntgenbildgebung mit der aufkommenden gitterbasierten differenziellen Phasenkontrast-Röntgenbildgebungstechnologie zu kombinieren. Zusätzlich zur Abschwächung der Intensität misst diese Technik die Phasenverschiebung, die eine Röntgenwelle beim Durchqueren eines Objekts erfährt. Numerische Simulationen von Radiographie- und Computertomographie-Anwendungen zeigen, dass die Kombination dieser beiden Bildgebungstechnologien (spektrale differenzielle Phasenkontrast-Bildgebung) von den Stärken der einzelnen Methoden profitiert, während die Schwächen ausgeglichen werden. Es können quantitativ genaue materialselektive Bilder berechnet werden und der Rauschpegel ist im Vergleich zur spektralen Röntgenbildgebung stark reduziert. Darüber hinaus entwickeln wir verschiedene Verarbeitungs- und Bildrekonstruktionstechniken für die spektrale und spektrale differenzielle Phasenkontrast-Bildgebung. Unsere Ergebnisse zeigen, dass die Bildqualität stark verbessert werden kann, indem man die Unvollkommenheiten aktueller photonenzählender Detektoren berücksichtigt, statistische Rauschmodelle verwendet und Vorwissen in ein bayesschen Bildrekonstruktionsansatz einfließen lässt. Obwohl viele Ergebnisse durch zusätzliche Experimente bestätigt werden müssen, glauben wir, dass diese Arbeit das Potenzial der spektralen differenziellen Phasenkontrast-Bildgebung mit photonenzählenden Detektoren in den Bereichen medizinische Bildgebung und zerstörungsfreie Prüfung zeigt. Dies gilt insbesondere für Anwendungen mit niedriger Dosis.

Contents

1 Introduction **1**

2 Fundamentals of spectral and phase-contrast X-ray imaging **5**
2.1 X-ray generation . . . 5
2.2 Free space wave propagation . . . 7
2.2.1 Angular spectrum propagation . . . 8
2.2.2 Fresnel propagation . . . 8
2.2.3 Incoherent sources . . . 11
2.3 X-ray interaction with matter . . . 12
2.3.1 Complex refractive index . . . 12
2.3.2 Projection approximation . . . 13
2.3.3 Fresnel scaling theorem . . . 14
2.3.4 Eikonal equation (geometrical optics) . . . 15
2.3.5 X-ray attenuation . . . 16
2.4 X-ray detection . . . 22
2.4.1 Quantum efficiency and photon statistics . . . 22
2.4.2 Classification of X-ray detectors . . . 23
2.4.3 Photon-counting detectors . . . 25
2.5 Statistical signal processing . . . 29
2.5.1 Maximum a posteriori and maximum-likelihood estimators . . . 31
2.5.2 The Cramér-Rao lower bound . . . 33
2.6 Reconstruction techniques for Computed Tomography . . . 33
2.6.1 Filtered backprojection . . . 33
2.6.2 Statistical iterative reconstruction . . . 36
2.7 Spectral X-ray imaging . . . 51
2.7.1 Technical realizations . . . 51
2.7.2 Basis material approximation . . . 51
2.7.3 Material decomposition algorithms . . . 54
2.8 Grating-based differential phase-contrast imaging . . . 60
2.8.1 Talbot-Lau interferometer . . . 62
2.8.2 Signal extraction . . . 66
2.8.3 Image artifacts . . . 70

3 Numerical simulation framework **74**
3.1 Spectral X-ray imaging . . . 74
3.2 Spectral differential phase-contrast X-ray imaging . . . 77

4 A post-processing algorithm for spectral CT material selective images using learned dictionaries **80**
4.1 Introduction . . . 80
4.2 Methods . . . 81
4.2.1 Dictionary denoising . . . 81
4.2.2 Modifications for denoising of basis material images . . . 82
4.2.3 Bilateral Filtering . . . 84
4.2.4 Numerical Simulation . . . 85

4.2.5 Experimental measurement . . . 85
4.3 Results . . . 86
4.4 Discussion . . . 86
4.5 Conclusion . . . 90

5 Spectral angiography material decomposition using an empirical forward model and a dictionary-based regularization 92
5.1 Introduction . . . 93
5.2 Methods . . . 94
5.2.1 Empirical model for material decomposition . . . 94
5.2.2 Regularized material decomposition . . . 96
5.2.3 Numerical simulation . . . 99
5.2.4 Experimental measurement . . . 101
5.3 Results . . . 103
5.3.1 Numerical simulation . . . 103
5.3.2 Experimental measurement . . . 105
5.4 Discussion . . . 107
5.5 Conclusion . . . 109

6 Joint statistical iterative material image reconstruction for spectral CT using a semi-empirical forward model 110
6.1 Introduction . . . 110
6.2 Methods . . . 111
6.2.1 Spectral estimator . . . 111
6.2.2 Penalized likelihood function for SIR . . . 112
6.2.3 Optimization algorithm . . . 113
6.2.4 SIR of decomposed line integrals . . . 117
6.2.5 Numerical simulation . . . 117
6.2.6 Experimental measurement . . . 118
6.3 Results . . . 119
6.4 Discussion . . . 124
6.5 Conclusion . . . 128

7 Spectral differential phase-contrast X-ray radiography 130
7.1 Introduction . . . 130
7.2 Methods . . . 131
7.2.1 Spectral phase-contrast model . . . 131
7.2.2 Numerical Experiment . . . 135
7.3 Results . . . 136
7.4 Conclusions . . . 141

8 A theoretical framework for comparing noise characteristics of spectral, differential phase-contrast and spectral differential phase-contrast X-ray imaging 144
8.1 Methods . . . 145
8.1.1 Forward models and signal extraction . . . 145
8.1.2 Noise analysis with the Cramér Rao lower bound (CRLB) . . . 148
8.1.3 Prediction of covariances and noise power spectra . . . 150
8.1.4 Numerical simulation . . . 151
8.2 Results . . . 152
8.3 Discussion . . . 157
8.4 Conclusion . . . 159

8.5 Appendix . . . 159
8.5.1 Noise propagation for the integration step in DPC imaging . . . 159
8.5.2 Calculating the noise power spectrum from the covariances . . . 160

9 Statistical iterative reconstruction for spectral differential phase-contrast CT 161
9.1 Methods . . . 161
9.1.1 Forward model and CT reconstruction . . . 161
9.1.2 Numerical simulation . . . 164
9.2 Results . . . 165
9.3 Discussion and Outlook . . . 165

10 Conclusions and outlook 168

11 Appendix 172
11.1 Wave propagation . . . 172
11.1.1 Extended sources and the projection approximation . . . 172
11.1.2 Talbot effect in combination with a phase-shifting object . . . 173
11.1.3 Fresnel scaling theorem and the multi-slice projection approximation . . . 174

Abbreviations 178

List of publications and scientific presentations 179

Bibliography 182

1 Introduction

With their ability to penetrate matter, X-rays enable the non-invasive acquisition of information about the inside of an object. In the past decades, X-ray imaging techniques have been steadily improved and have become an invaluable tool for clinical diagnosis and non-destructive testing applications. Compared to the first prototypes, modern X-ray imaging systems offer improved image quality, contrast and resolution while dramatically reducing the scanning time and the radiation dose. The development of X-ray computed tomography (CT) in the early 1970s was a particularly important milestone. By obtaining projection measurements at various angles around an object, CT allows to reconstruct a three-dimensional map of object properties (such as the attenuation coefficient). Advances in X-ray imaging technology are still driven by improvements in three different fields that are strongly interlinked:

- data acquisition and processing hardware (e.g. X-ray source, detector, CT gantry and computer hardware),
- novel imaging methods (e.g. spectral X-ray imaging, phase-contrast X-ray imaging, sparse sampling CT), and
- data processing algorithms (e.g. statistical iterative reconstruction, deep learning based image analysis, denoising algorithms).

In the last years, spectral X-ray imaging techniques have gained increasing interest for medical imaging and non-destructive testing applications. Compared to conventional attenuation imaging, spectral X-ray imaging provides additional information about the composition of an object by conducting measurements with two or more distinct photon energy spectra. The important special case of using two different spectra is often called dual energy imaging. By exploiting information about the energy-dependent attenuation, basis material decomposition algorithms allow to distinguish between different materials and to obtain information about the chemical composition of an object. Moreover, image artifacts (referred to as "beam hardening") related to the polychromatic X-ray spectrum can be efficiently suppressed. These unique properties have proven to be beneficial for clinical diagnosis in many different application [1]. The principle and mathematical concept of dual energy X-ray imaging was already proposed in 1976 by Alvarez and Macovski [2]. However, dual energy imaging has only gained popularity after more recent hardware advances (which will be discussed in more detail in section 2.7.1) have made it possible to use the technique in clinical CT scanners. Material decomposition algorithms typically lead to a degradation of the signal-to-noise ratio and noise amplification compared to the unprocessed spectral images [2, 3]. This is still a fundamental problem of spectral CT since it limits the usability of basis material images. Noise amplification is particularly problematic for low-dose medical imaging because the increased quantum noise already leads to higher noise levels in the unprocessed spectral images.

As a result of technological progress, the number of X-ray imaging procedures has increased steadily, so that they now represent a major source of radiation exposure for the population [4]. Due to concerns about radiation-induced cancer (especially for CT scans [5, 6]), reducing the dose delivered to the patient while preserving image quality is an important challenge for spectral and conventional X-ray imaging techniques. Recent progress in computer hardware has enabled the application of advanced denoising and reconstruction algorithms that have the

potential to substantially increase the image quality. In this context, reducing the computational time by performing calculations in parallel on graphics processing units (GPUs) has been a key concept. For conventional CT, it has been demonstrated that statistical iterative reconstruction techniques can significantly reduce the radiation dose compared to analytical image reconstruction methods [7, 8, 9, 10]. As will be discussed later (see section 2.6.2), one of the largest advantages of SIR techniques is their flexibility. In contrast to analytical reconstruction methods, it is possible to incorporate a noise model, a (simplified) simulation of the physical processes and prior knowledge about the reconstructed images.

Spectral X-ray imaging with photon-counting detectors (PCDs) has gained increasing popularity in the last years. Originally developed at CERN in the late 1980s for the analysis of particle collider experiments [11, 12], PCD technology has been transferred to X-ray imaging applications. In contrast to standard integrating detectors, PCDs directly convert the incident X-rays to electron-hole pairs in a semiconductor sensor. The generated electrical signals are analyzed separately for each detector pixel with dedicated fast readout circuits, which allows individual photons to be counted. These properties of PCDs have proven promising for spectral X-ray imaging applications. By comparing the generated electrical signal to a set of predefined thresholds, PCDs are capable of (partially) resolving the x-ray photon energy. This enables the simultaneous acquisition of several energy selective images that are perfectly registered, both in the spatial and temporal domain. Other established dual energy imaging techniques (see section 2.7.1) are difficult to extend to three or more energy resolved measurements. Contrary to conventional X-ray detectors, electronic readout noise is efficiently suppressed, which is particularly important for low-dose applications. In addition, PCDs can achieve higher spatial resolution because they do not require an additional conversion step to visible light photons in a scintillator. There are several reasons why PCDs are not yet established in clinical X-ray imaging applications. In addition to high manufacturing costs, current PCD technology suffers from various performance degrading effects that will be discussed in section 2.4.3. Moreover, manufacturing imperfections and inhomogeneities of the semiconductor sensor material lead to spatial variations of the detector response, which degrade the image quality.

Besides spectral imaging, phase-contrast X-ray imaging is another emerging technology that can provide improved image quality and additional information compared with conventional X-ray imaging. Among the many phase-contrast imaging methods, grating-based differential phase-contrast (DPC) imaging is particularly promising for medical imaging and non-destructive testing applications since it places the lowest demands on spatial and temporal coherence [13]. Because of the low source coherence requirements, DPC imaging can be performed with standard X-ray tubes. Grating-based differential phase-contrast imaging exploits an entirely different contrast generating mechanisms in addition to the conventional attenuation contrast: The phase-contrast image is obtained by indirectly measuring the (differential) phase shift that an X-ray wave exhibits when traversing an object [13, 14] with the help of a three-grating interferometer. This differential phase shift can be directly related to the electron density distribution of the object. Depending on the imaging parameters, DPC imaging can achieve a highly improved contrast-to-noise ratio (CNR) compared to attenuation-based imaging [15, 16, 17, 18, 19]. This applies particularly to materials with low atomic numbers (e.g. tissues), which have low absorption contrast. It has been recently demonstrated that - similarly to spectral imaging - material selective images can in principle be calculated from DPC measurements by using the attenuation and phase shift information [20, 21]. Moreover, the so-called dark-field signal can be extracted from the interferometric measurements [22]. It is related to ultra-small angle scattering and provides information about the microstructure of an object far below the detector resolution. Contrary to ultra high resolution X-ray imaging, dark-field imaging is performed on large objects, using conventional X-ray tubes and standard detectors with a large field of view.

This unique feature has lead to a variety of promising future medical imaging applications (see section 2.8.2). However, similarly to conventional attenuation imaging, the polychromatic X-ray spectrum causes beam hardening artifacts in both the phase and the dark-field channel [23, 24]. Furthermore, the phase shift cannot be unambiguously determined. This phase-wrapping problem is another important source of artifacts for grating-based phase-contrast imaging [25, 26].

In summary, spectral and grating-based differential phase-contrast (DPC) X-ray imaging are two emerging technologies with great potential for medical imaging applications. They could help to reduce the radiation dose and offer valuable additional information (basis material images, dark-field image) compared with conventional attenuation-based X-ray imaging. Both X-ray imaging techniques could be further enhanced by combining them with PCD technology. In the case of spectral imaging, potential advantages of PCDs compared to standard dual energy imaging have already been demonstrated in many pre-clinical and clinical studies [27, 28, 29, 30, 31, 32, 33, 34]. However, both imaging techniques still face experimental and algorithmic challenges that must be overcome on the way to routine use in clinical and non-destructive testing applications. In this work, we devise and investigate processing and reconstruction algorithms for spectral and DPC X-ray imaging with PCDs. Moreover, we develop a mathematical model and processing algorithms for combining both X-ray imaging methods. From a hardware perspective, this novel method, which we call spectral differential phase-contrast (SDPC) imaging, is realized by combining a grating interferometer with an energy-resolving PCD. These algorithmic developments aim at increasing the diagnostic or informative value of spectral and SDPC imaging, especially for low-dose applications. This is achieved by reducing the noise level, eliminating image artifacts and increasing the quantitative accuracy of the obtained images.

Outline The thesis is organized as follows. Chapter 2 discusses the fundamentals of spectral and DPC X-ray imaging, on which the results of this thesis are based. Besides the fundamentals of X-ray physics, the working principle of the data acquisition hardware as well as various signal processing and image reconstruction techniques are explained. We also review the state of the art and important applications for both X-ray imaging techniques. In chapter 3, we present the numerical simulation framework for spectral and DPC imaging that was partly developed in this work and discuss its limitations. Throughout this work, numerical simulations were used to study the performance and characteristics of the developed reconstruction and denoising algorithms. The first part of the results of this thesis (chapters 4 - 6) focuses on spectral X-ray imaging. In chapters 5 and 6, we address the challenges of photon-counting detector (PCD) imperfections and their adverse effect on the image quality by combining material decomposition algorithms with (semi-) empirical forward models for the measurement process. By tuning the model parameters with calibration measurements, pixel-to-pixel variations of the detector response can be taken into account. To suppress noise amplification during material decomposition, we investigate multi-channel post-processing (chapter 4) and regularization (chapter 5) techniques as well as a novel statistical iterative reconstruction algorithm (chapter 6) that combines material decomposition and image reconstruction in a single processing step.
In the second part of this thesis, a new way of combining spectral and grating-based DPC imaging is discussed. Numerical simulations of radiography (chapter 7) and CT applications (chapter 9) show that the combination of these two imaging techniques benefits from the strengths of the individual methods while the weaknesses are mitigated. Chapter 8 analyzes the unique noise characteristics of spectral spectral differential phase-contrast (SDPC) imaging. Moreover, we theoretically compare the performance of spectral, DPC and SDPC imaging as a function of the setup parameters.

As already partly demonstrated in this thesis, most of the processing techniques for spectral

imaging that were discussed in chapters 4 - 6 can also be transferred to SDPC imaging to further improve the image quality or reduce the radiation dose. Although many results need be confirmed by additional numerical simulations and experimental measurements, we believe that this work shows the potential of SDPC imaging with PCDs for future clinical and non-destructive testing applications. Our results indicate that by combining SDPC imaging with advanced processing algorithms and PCD technology, valuable additional information compared to the individual imaging techniques can be obtained while improving the image quality, especially for low-dose applications.

2 Fundamentals of spectral and phase-contrast X-ray imaging

This chapter discusses the theoretical foundation of spectral and differential phase-contrast imaging and thus serves as a basis for the following chapters. We start by explaining the physical principles behind X-ray imaging, including X-ray generation, wave propagation, interaction of X-rays with matter and X-ray detection. The second part of this chapter focuses on on two algorithmic techniques that are particularly relevant for X-ray imaging: statistical signal processing and CT reconstruction. Finally, we discuss the basic principles of spectral and differential phase-contrast imaging and review the state of the art in these two research fields.

2.1 X-ray generation

In this section, we focus on the physics of X-ray generation using conventional X-ray tubes. In particular, the physical as well as statistical properties that are relevant for spectral and differential phase-contrast imaging will be discussed. Although other X-ray sources such as synchrotron radiation ([35], chapters 2.2 - 2.5) or free electron lasers [36] outperform X-ray tubes in many important aspects (e.g. photon flux, spatial and temporal coherence), almost all medical imaging as well as most non-destructive testing applications rely on X-ray tubes. High costs and complexity as well as large space requirements often render synchrotrons and other high-brilliance X-ray sources impractical for routine applications. The recent development of compact synchrotron sources [37, 38, 39] could provide an interesting compromise between high-brilliance X-ray sources and applicability in future medical imaging tasks.
Figure 2.1 depicts a simplified schematic drawing of a rotating anode X-ray tube. The reader is referred to [40] for a comprehensive treatment of modern X-ray tube technology. Electrons are emitted from the cathode by heating a filament (typically made of tungsten) to approximately 2000 K [40]. The thermal energy allows the electrons to overcome their binding energy to the filament. After emission from the filament, the electrons are accelerated towards the anode by the applied high bias voltage. At the anode, the interaction with the Coulomb fields of the atoms in the anode material cause a rapid deceleration of the incident electrons. In a classical electrodynamics picture, the deceleration of the charged electrons generates electromagnetic radiation (*bremsstrahlung*). Usually, the deceleration of an electron is a multi-step process and several photons are emitted during the deceleration process. However, it is also possible that the whole kinetic energy of an electron is transferred to one photon. The maximum photon energy is determined by energy conservation and is thus equal to the kinetic energy of the electron. In summary, due to the multi-step deceleration process, *bremsstrahlung* produces a continuous spectrum of photon energies with a maximum energy equal to the electron energy.
The continuous spectrum is superimposed by a line spectrum that is characteristic of the anode material. These sharp emission lines (width of $\approx$ 40 eV for tungsten [40]) are generated by collisions of the incident electrons with electrons in the inner shells of the anode atoms. The electron in the inner shell is kicked out of the atom and the vacant position is filled by an electron from one of the outer shells. In this process, a photon with characteristic energy determined by the energy difference between the participating electron shells is emitted. For most anode materials, only the K_α and K_β emission lines are in the diagnostics energy range ($\approx 15-150$ keV). They represent transitions from the L- and M-shell to the K-shell, respectively. A typical X-ray spectrum of a tungsten anode for an acceleration voltage of 120 kVp is shown

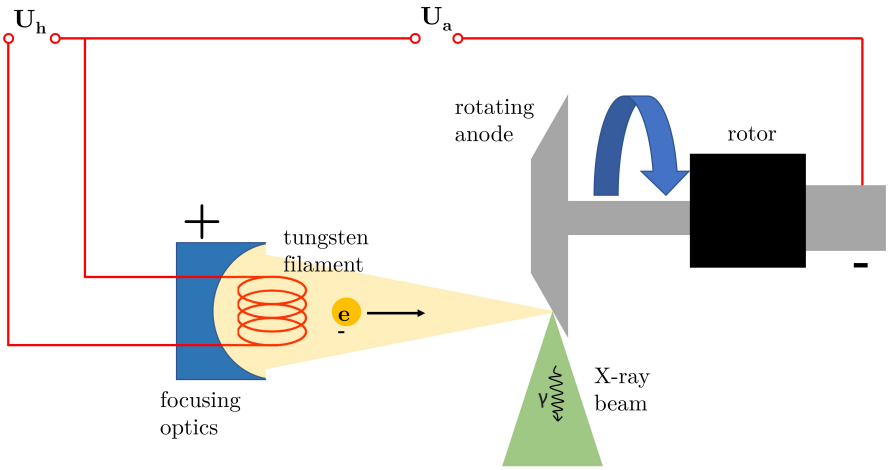

Figure 2.1: Schematic drawing of a rotating anode. Electrons are emitted from an electrically heated tungsten filament and subsequently accelerated towards the rotating anode by the high bias voltage (U_a). Additional focusing optics shape the electron beam to a small spot on the anode. The rapid deceleration of the electrons upon impact on the anode causes the emission of X-rays. To distribute the heat load over a larger area, the anode target is rotated continuously.

in figure 2.2. Although the peaks caused by the characteristic emission lines are far higher than the smooth *bremsstrahlung* background, they only contribute about 7% to the total intensity [40] due to their small width.
The aforementioned mechanisms for generating X-rays by interactions of fast electrons with the anode material are rather inefficient: Only approximately 1% of the primary electron energy is converted to X-ray radiation [41]. The vast majority of the kinetic energy is converted to heat or carried away by scattered electrons. Since the X-ray emission is nearly isotropic, only about 0.03 % of the primary electron energy is used for X-ray imaging [40]. As there are no efficient X-ray lenses in the diagnostic energy range, the X-ray beam is delimited by apertures to extinguish unwanted radiation.
To achieve a sufficiently high photon flux, modern X-ray tubes have an electric input power of up to approximately 100 kW (for CT applications). This leads to a significant heating problem in the X-ray tube and consequently, some kind of cooling mechanism has to be implemented. The most common approach is to use a rotating anode where the material hit by the electron beam is constantly exchanged by the rotation. In this way, the energy of the electron beam and therewith the heat load is spread out over a (circular) line instead of a small spot and can be dissipated more easily. To avoid blurring of the acquired X-ray images, it is desirable that the X-rays are only emitted from a very small spot on the anode. Additional electron focusing optics that direct the electron beam onto a small spot on the anode are therefore typically placed around the filament.
The generation of a photon by an incident electron that interacts with the anode material can be viewed as a random event with a certain probability p. Under this assumption, the probability that q photons are emitted in the time interval $[0, T]$ is given by a binomial distribution:

$$P(Q = q) = \binom{N}{q} p^q (1-p)^{N-q}, \tag{2.1}$$

where N is the number of electrons that hit the anode in the time interval $[0, T]$. One can show (see reference [41], chapter 2.6.1 for a derivation of this result) that in the limit $N \to \infty,\ p \to 0$, the binomial distribution converges to the Poisson distribution:

$$P(Q = q) = \frac{\lambda^q}{q!} e^{-\lambda}, \quad \lambda = Np. \tag{2.2}$$

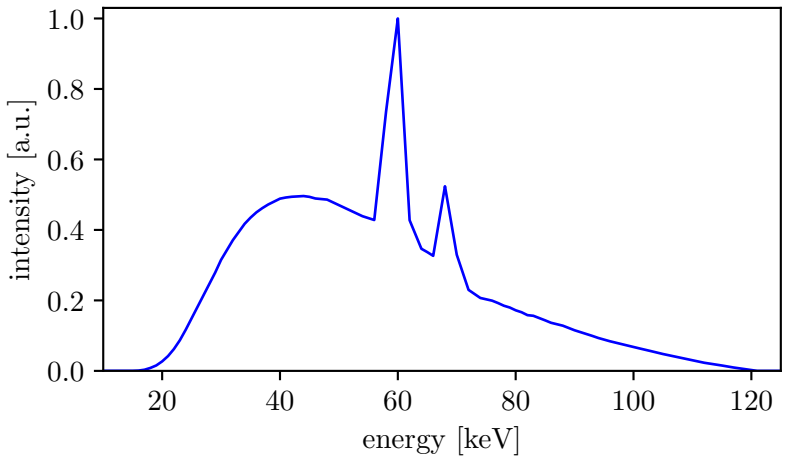

Figure 2.2: Simulated X-ray spectrum for a tungsten anode with an acceleration voltage of 120 kVp. The TASMIP algorithm [42] was used to generate the simulated spectrum.

For a typical X-ray tube, approximately 10^{15} electrons arrive at the anode per millisecond and the probability that a particular electron generates a photon is small. The statistical distribution of photons that are emitted in a certain time interval is thus well approximated by a Poisson distribution.
Some of the aforementioned properties of typical X-ray tubes are particularly relevant for mathematically modeling and simulating spectral as well as differential phase-contrast imaging:

- Conventional X-ray sources emit a polychromatic spectrum with a broad range of photon energies.
- The number of emitted photons in a certain time interval follows a Poisson distribution.
- X-rays are emitted from an extended source spot. Considering the process of photon generation, the emitted radiation is spatially incoherent. This aspect will be discussed in more detail in section 2.2.3.

2.2 Free space wave propagation

Since the wavelength of X-rays is typically much smaller than the investigated structures, conventional X-ray imaging is well described by geometrical optics. However, a wave-optical approach is needed to describe the interference effects that grating-based phase-contrast imaging relies on. The following sections provide an introduction to this approach, starting with the propagation of electromagnetic waves in a vacuum. The derivations are mainly based on reference [43], where additional information can be found. In the absence of matter, the components of the electric and magnetic field are decoupled in Maxwell's equations. The electromagnetic disturbance can thus be described by a single scalar field:

$$\phi(x, y, z, t) = \phi_\omega(x, y, z)e^{-i\omega t}, \tag{2.3}$$

where $w = 2\pi f$ is the angular frequency of the monochromatic wave. For simplicity, this section focuses on monochromatic wave functions, however a generalization to polychromatic wave functions is easily obtained by spectral decomposition [43]. Starting from Maxwell's equations, one can show that the spatial components of the wave function $\phi_\omega(x, y, z)$ obey the Helmholtz equation:

$$(\nabla^2 + k^2)\phi_\omega(x, y, z) = 0, \quad k = \frac{\omega}{c} = \frac{2\pi}{\lambda}, \tag{2.4}$$

where λ is the wavelength and c is the speed of light. Elementary plane waves of the form:

$$\phi_\omega^{(\mathrm{PW})}(x,y,z) = e^{i(k_x x + k_y y + k_z z)} \tag{2.5}$$

are solutions to the Helmholtz equation, provided that:

$$k^2 = k_x^2 + k_y^2 + k_z^2. \tag{2.6}$$

Assuming plane waves that are forward propagating with respect to the optical axis z, $\phi_\omega^{(\mathrm{PW})}(x,y,z)$ can be rewritten as:

$$\phi_\omega^{(\mathrm{PW})}(x,y,z) = e^{i(k_x x + k_y y)} e^{iz\sqrt{k^2-k_x^2-k_y^2}} = \phi_\omega^{(\mathrm{PW})}(x,y,0) e^{iz\sqrt{k^2-k_x^2-k_y^2}}. \tag{2.7}$$

The second equality in eq. 2.7 means that the value of the propagated wave function over the plane $z = z_0$ ($z_0 > 0$) can be obtained by multiplying $\phi_\omega^{(\mathrm{PW})}(x,y,0)$ with the free-space propagator $e^{iz_0\sqrt{k^2-k_x^2-k_y^2}}$.

2.2.1 Angular spectrum propagation

The propagated value of an arbitrary wave function can be calculated by decomposing it into elementary plane waves using a two-dimensional Fourier representation:

$$\phi_\omega(x,y,z=0) = \frac{1}{2\pi} \int\int \tilde{\phi}_\omega(k_x,k_y,z=0) e^{i(k_x x + k_y y)} dk_x dk_y. \tag{2.8}$$

The value of the wave function at the plane $z = z_0$ is calculated by multiplying the plane waves with the corresponding free space propagator:

$$\phi_\omega(x,y,z_0) = \frac{1}{2\pi} \int\int \tilde{\phi}_\omega(k_x,k_y,0) e^{iz_0\sqrt{k^2-k_x^2-k_y^2}} e^{i(k_x x + k_y y)} dk_x dk_y. \tag{2.9}$$

This equation is known as the angular spectrum representation of the propagated wave field [43]. In operator notation, eq. 2.9 can be expressed as:

$$\phi_\omega(x,y,z_0) = \mathcal{F}^{-1} e^{iz_0\sqrt{k^2-k_x^2-k_y^2}} \mathcal{F}\ \phi_\omega(x,y,0), \tag{2.10}$$

i.e. the wave function at $z = z_0$ is calculated by first taking the Fourier transform of the wave function at $z = 0$, multiplying it with the free-space propagator and then transforming the result back to real space.

2.2.2 Fresnel propagation

Fresnel propagation is a useful approximation to angular spectrum propagation for a paraxial incident wave $\phi_\omega(x,y,z=0)$. In mathematical terms, the condition of paraxiality is expressed as:

$$|k_x| \ll k_z, \quad |k_y| \ll k_z \tag{2.11}$$

for all non-negligible components of $\phi_\omega(x,y,z=0)$ (compare eq. 2.8). In a geometrical optics picture, this corresponds to the assumption that all X-rays make small angles with respect to the optical axis. As will be discussed in section 2.3, this assumption is normally fulfilled for X-ray imaging because the refraction angles for X-rays are very small. Moreover, the Fresnel scaling theorem (see section 2.3.3) allows to use the Fresnel approximation in combination with point sources. Using a binomial approximation for the exponent in eq. 2.10:

$$\sqrt{k^2 - k_x^2 - k_y^2} \approx k - \frac{k_x^2 + k_y^2}{2k}, \tag{2.12}$$

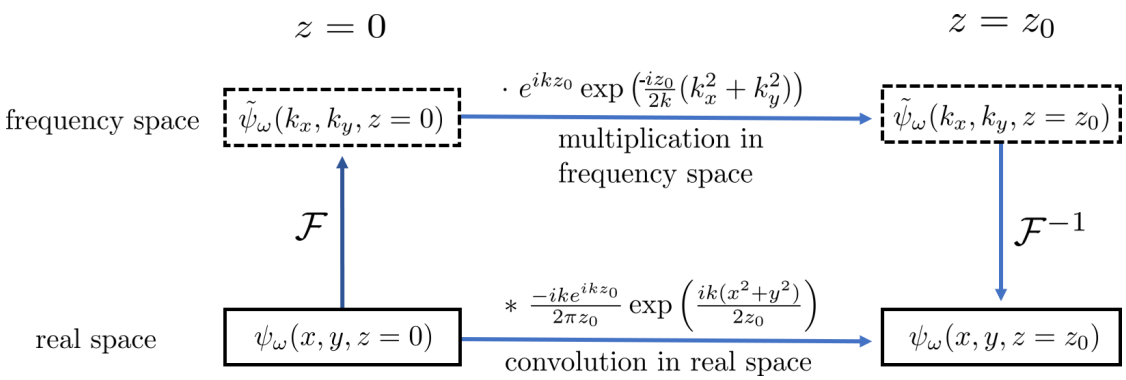

Figure 2.3: Visualization of two different strategies to calculate the propagated wave function in Fresnel approximation. The propagated wave field can either be calculated by a convolution in real space or by transforming into frequency space, multiplying by the Fresnel propagator and then transforming back into real space.

the propagated wave field is given by:

$$\phi_\omega(x, y, z_0) = \mathcal{F}^{-1} e^{ikz_0} e^{\frac{-iz_0}{2k}(k_x^2+k_y^2)} \mathcal{F}\ \phi_\omega(x, y, 0). \tag{2.13}$$

According to equation 2.13, the propagated wave field in Fresnel approximation is calculated by Fourier-transforming the initial wave field, multiplying with the Fresnel propagator $e^{ikz_0} e^{\frac{iz_0}{2k}(k_x^2+k_y^2)}$ in frequency space and then transforming back into real space. Note that the phase factor e^{ikz_0} can be pulled out of the inverse Fourier transform. Since multiplication in frequency space corresponds to convolution in real space, $\phi_\omega(x, y, z_0)$ can also be calculated in real space (see reference [43] for a derivation of this result):

$$\phi_\omega(x, y, z_0) = \phi_\omega(x, y, z = 0) * H(x, y, z_0), \tag{2.14}$$

where $*$ denotes a convolution and $H(x, y, z_0)$ is the real space form of the Fresnel propagator:

$$H(x, y, z_0) = \mathcal{F}^{-1}\left(e^{ikz_0} e^{\frac{-iz_0}{2k}(k_x^2+k_y^2)}\right) = \frac{-ike^{ikz_0}}{2\pi z_0} e^{\frac{ik(x^2+y^2)}{2z_0}}. \tag{2.15}$$

In summary, the propagated wave function at a plane $z = z_0$ can be calculated from the wave function at the plane $z = 0$ either by multiplication with the Fresnel-propagator in frequency space or by convolution with the real-space form of the Fresnel propagator in real space. Figure 2.3 illustrates the two different ways for calculating the propagated wave function in free space with the Fresnel approximation. From a numerical perspective, multiplying with the Fresnel propagator in frequency space is often preferred because the computational complexity of the convolution is much higher (compared to an element-wise multiplication and two fast Fourier transforms).

Huygens-Fresnel principle: The convolution formulation of Fresnel propagation (eq. 2.14) allows an interpretation in terms of the Huygens-Fresnel principle. Suppose there is a point source in the z=0 plane, i.e. :

$$\phi_\omega(x, y, z = 0) = \phi_\omega^{(0)} \delta(x - x_0)\delta(y - y_0). \tag{2.16}$$

The wave field over the plane $z = z_0$ is then given by:

$$\phi_\omega(x, y, z = z_0) = \phi_\omega^{(0)} \frac{-ike^{ikz_0}}{2\pi z_0} \exp\left(\frac{ik((x - x_0)^2 + (y - y_0)^2)}{2z_0}\right). \tag{2.17}$$

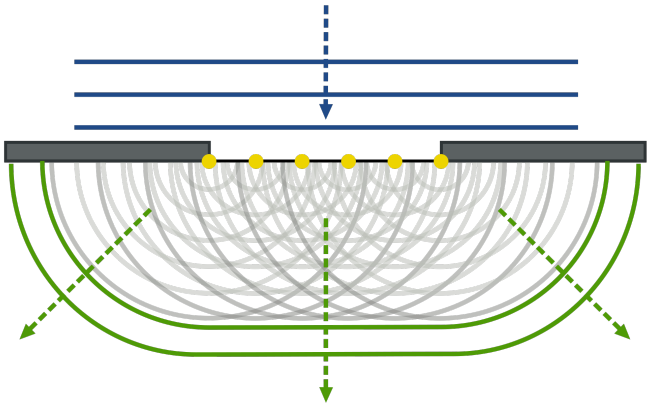

Figure 2.4: Illustration of the Huygens-Fresnel principle: Every point on the aperture can be interpreted as the source of a spherical wave. The propagated wave function is given by the superposition of these spherical waves. This figure was taken from reference [45].

This corresponds to a spherical wave (in parabolic approximation) emanating from the point $(x_0, y_0, 0)$. Moreover, eq. 2.17 represents the impulse response of the linear optical system [44]. Due to the linearity of the convolution in eq. 2.14, the propagated wave function $\phi_\omega(x, y, z = z_0)$ for an arbitrary wave function $\phi_\omega(x, y, z = 0)$ in the $z = 0$ plane can be interpreted as a superposition of spherical waves. Figure 2.4 shows an illustration of the Huygens-Fresnel principle.

Local interference for X-ray imaging The aim of this paragraph is to demonstrate that Fresnel propagation is essentially a local interference phenomenon within the parameter range that is typically used for medical X-ray imaging and non-destructive testing. This means that the majority of the contributions to the value of the propagated wave function $\phi_\omega(x_0, y_0, z_0)$ at a certain point (x_0, y_0) come from a small neighborhood around this point at the $z = 0$ plane $(\phi_\omega(x_0 \pm \Delta x, y_0 \pm \Delta y, z = 0))$ [44]. Using the formula for Fresnel propagation in real space (see eq. 2.14), $\phi_\omega(x_0, y_0, z_0)$ is given by:

$$\phi_\omega(x_0, y_0, z_0) = \frac{-ike^{ikz_0}}{2\pi z_0} \int\int \phi_\omega(x, y, 0) \exp\left(\frac{ik((x_0 - x)^2 + (y_0 - y)^2)}{2z_0}\right) dxdy. \tag{2.18}$$

Figure 2.5 (a) shows a plot of the real part of the complex exponential term in the integral of eq. 2.18 as a function of the distance between x_0 and x ($d_x = x_0 - x$) for a propagation distance of $z_0 = 1$ m and an X-ray energy of 40 keV (corresponds to $k = 2.0 \cdot 10^{11}$ m^{-1}). These values provide a reasonable order-of magnitude estimate for various medical imaging and non-destructive testing applications. One can see that the oscillation period p_x decreases with increasing d_x. If $d_x \gg p_x$, p_x can be approximated as

$$p_x = \frac{2\pi z_0}{k d_x} \approx 3 \cdot 10^{-11} \text{ m}^2 \frac{1}{d_x}. \tag{2.19}$$

For $d_x = 100$ μm, the oscillation period is already approximately 0.3 μm which is much shorter than the period of any regular structure (e.g. gratings for phase-contrast imaging) that typically occurs in an X-ray imaging experiment. In this case, $\phi_\omega(x_0, y_0, z_0)$ varies slowly in comparison to the complex exponential term in eq. 2.18 and the contributions to the integral in eq. 2.18 for $|d_x| = |x_0 - x| > 100$ μm tend to cancel out because of the rapid oscillations. This is visualized in figure 2.5 (b) where the intensity of the propagated wave function $|\phi_\omega(x_0, y_0, z_0)|^2$ for a homogeneous incident wave ($\phi_\omega(x_0, y_0, z = 0) = 1$) is calculated by only considering a small neighborhood $x_0 \pm \Delta x$ in the convolution integral of eq. 2.18. For $\Delta x > 10$ μm, the

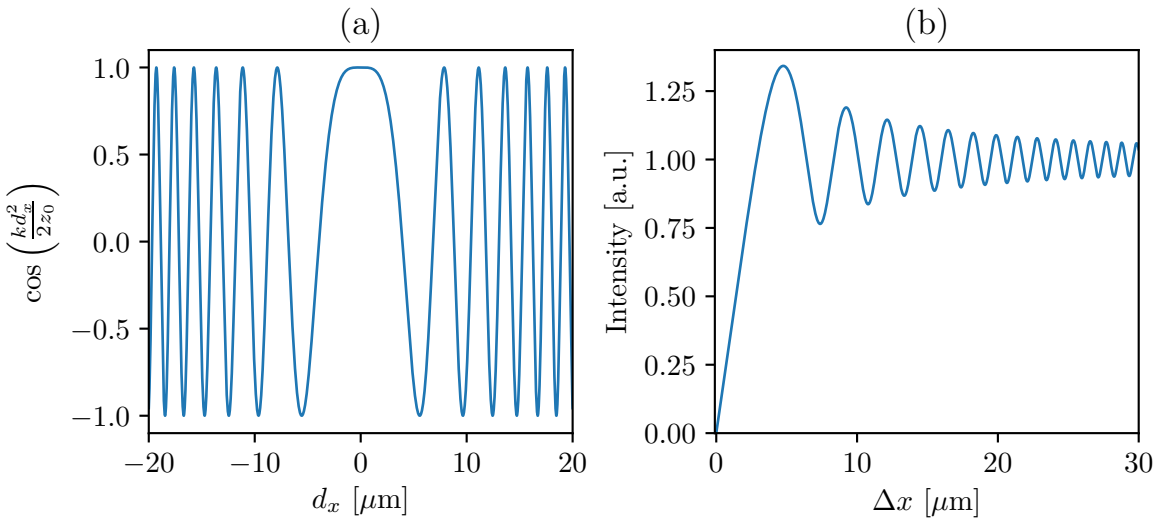

Figure 2.5: (a): Plot of the real part of the complex exponential factor in eq. 2.18 as a function of $d_x = x_0 - x$ for a propagation distance of $z_0 = 1$ m and an X-ray energy of 40 keV. The frequency of the oscillations increases with increasing distance between x_0 and x. (b): Intensity of the propagated wave function as a function of the size of the neighborhood (Δx) around x_0 that was considered for the convolution integral in eq. 2.18.

intensity of the propagated wave field starts to oscillate rapidly around the ground-truth value and the amplitude of the oscillations decreases as Δx is increased. Consequently, the majority of the contributions to the propagated wave function come from a small area around the point of interest (x_0, y_0) in the $z = 0$ plane. As a rule of thumb, for a sufficiently slowly varying wave field $\phi_\omega(x_0, y_0, z = 0)$, most contributions come from a square with width $4\sqrt{\lambda z_0}$ centered around the point of interest (see reference [44], chapter 4.2.2 for more details). This motivates the use of "ray-based" (forward) models for the measured intensities both in the case of conventional as well as grating-based phase phase-contrast X-ray imaging (see sections 2.6.2).

2.2.3 Incoherent sources

As discussed in chapter 2.1, generating photons by electron-electron collisions is an independent stochastic process, i.e. the interaction of an impinging electron with the anode material is normally independent of other impinging electrons. An X-ray tube thus emits incoherent radiation from an extended area (the source spot). This effect has to be considered in order to obtain an accurate physical model of the X-ray imaging process. Towards this end, we consider the wave field $\phi(x, y, z_d, t)$ at a point $(x, y, z = z_d)$ in the observation plane that is produced by two monochromatic point sources located at $(x_1, y_1, z = 0)$ and $(x_2, y_2, z = 0)$, respectively. Using the Huygens-Fresnel principle (eq. 2.17), $\phi(x, y, z_d, t)$ can be expressed as:

$$\phi(x, y, z_d, t) = K_1\phi(x_1, y_1, 0, t) + K_2\phi(x_2, y_2, 0, t), \tag{2.20}$$

where K_1 and K_2 are complex constants. Since the integration time is typically extremely long compared to the oscillation period $T = 2\pi/\omega$, an X-ray detector placed at (x, y, z_d) measures

the time-averaged intensity $I(x, y, z_d)$ of the wave field:

$$\begin{aligned} &I(x,y,z_d) = \langle|\phi(x,y,z_d,t)|^2\rangle_t = \langle\phi(x,y,z_d,t)\phi^*(x,y,z_d,t)\rangle_t \\ &= \langle|K_1\phi(x_1,y_1,0,t)|^2\rangle_t + \langle|K_2\phi(x_2,y_2,0,t)|^2\rangle_t + \\ &K_1K_2^*\langle\phi(x_0,y_0,0,t)\phi^*(x_2,y_2,0,t)\rangle_t \ + K_1^*K_2\langle\phi^*(x_1,y_1,0,t)\phi(x_2,y_2,0,t)\rangle_t \\ &= |K_1|^2 I(x_1,y_1,0) + |K_2|^2 I(x_2,y_2,0) + 2\mathrm{Re}\left[K_1K_2^*J_{12}\right], \\ &J_{12} = \langle\phi(x_1,y_1,0,t)\phi^*(x_2,y_2,0,t)\rangle_t. \end{aligned} \tag{2.21}$$

The last line of eq. 2.21 defines the mutual intensity J_{12} [46]. By normalizing with the intensity of the two point sources, the complex coherence factor u_{12} is obtained [47]:

$$u_{12} = \frac{J_{12}}{\left[I(x_1,y_1,0)I(x_2,y_2,0)\right]^{1/2}}. \tag{2.22}$$

The complex coherence factor characterizes the degree of spatial coherence and is directly related to the visibility of interference phenomena (e.g. interference fringes in Young's double slit experiment). Perfect coherence corresponds to $u_{12} = 1$, whereas incoherent sources are characterized by $u_{12} = 0$ because there is no defined phase relationship between the two point sources. For incoherent sources, the observed intensity thus corresponds to the sum of the intensities that the two point sources produce individually:

$$I(x,y,z_d) = |K_1|^2 I(x_1,y_1,0) + |K_2|^2 I(x_2,y_2,0). \tag{2.23}$$

Equation 2.23 can be generalized to a spatially extended source by modeling it with an infinite number of point sources.

2.3 X-ray interaction with matter

This section discussed the interaction of X-rays with matter. Starting from a wave-optical approach that considers Maxwell's equations, the influence of matter on the wave function is examined and the notion of a complex refractive index is developed. The physical origin of the different contributions to the refractive index is more conveniently discussed in the particle picture.

2.3.1 Complex refractive index

Assuming a linear, isotropic and non-magnetic material (which is normally the case for X-ray imaging applications), a similar equation to the free-space Helmholtz equation (see eq. 2.4) can be derived [43]:

$$\left[\nabla^2 + k^2 n_\omega(x,y,z)\right]\psi_\omega(x,y,z) = 0, \tag{2.24}$$

where

$$n_\omega(x,y,z) = \sqrt{\frac{\epsilon_\omega(x,y,z)}{\epsilon_0}} \tag{2.25}$$

is the complex refractive index. The permittivity of the material and the permittivity of free space are denoted by $\epsilon_\omega(x, y, z)$ and ϵ_0, respectively. The derivation of eq. 2.24 also assumes that the material properties (i.e. $n_\omega(x, y, z)$) vary slowly compared to the wavelength of the X-ray radiation. Similar to free-space propagation, the more general case of a polychromatic wave function is obtained by a superposition of the corresponding monochromatic wave functions. The complex refractive index is often expressed in the form:

$$n_\omega = 1 - \delta_w + i\beta_w, \tag{2.26}$$

where $\delta_\omega \ll 1$ and $\beta_\omega \ll 1$ are real, positive numbers. The physical meaning of these two quantities will be discussed in more detail in the following sections. The 'inhomogeneous' Helmholtz equation (eq. 2.24) is difficult to solve exactly. Therefore, various approximations such as the first Born approximation and the projection approximation have been developed. The latter will be discussed in more detail because (in combination with Fresnel propagation) it is well suited for numerical simulations of X-ray imaging.

2.3.2 Projection approximation

As will be discussed in section 2.3.4, the refraction angles for X-rays are usually very small since the real part of the complex refractive index is close to unity. In this case, the projection approximation provides a useful approximate solution to the 'inhomogeneous' Helmholtz equation (eq. 2.24) for X-ray imaging of weakly scattering objects. Form a geometrical optics perspective, this corresponds to the assumption that the perturbation to the ray paths which would have existed in the absence of the object are negligible [43]. Therefore, the phase and amplitude shift of the wave function that is caused by the object can be considered separately from the effect of free space propagation. This corresponds to reducing the object to a plane (with surface normal parallel to the propagation direction) that causes the same phase shift and attenuation. One can show [43] that the influence of the 'object plane' on the phase and amplitude of the incident wave function is approximately given by:

$$\psi_\omega^{(d)}(x,y,z=0) = \exp\left(-ik\int_{z=0}^{z=z_0}\left[\delta_\omega(x,y,z) - i\beta_\omega(x,y,z)\right]dz\right)\psi_\omega^{(u)}(x,y,z=0), \qquad (2.27)$$

where $\psi_\omega^{(d)}$ and $\psi_\omega^{(u)}$ are the wave functions downstream and upstream of the object plane, respectively. The object is assumed to extend from $z = 0$ to $z = z_0$ (compare figure 2.6) and the object plane is located at $z = 0$. Combining eq. 2.27 with the Fresnel approximation to free-space propagation (eq. 2.13) yields the wave field at $z = z_0$ calculated under the projection approximation:

$$\psi_\omega(x,y,z=z_0) = \mathcal{F}^{-1}e^{ikz_0}e^{\frac{-iz_0}{2k}(k_x^2+k_y^2)}\mathcal{F}\ \phi_\omega^{(d)}(x,y,0). \qquad (2.28)$$

As long as the propagation distance from the object to the detector is much larger than the extent of the object in beam direction, the projection approximation provides accurate results [48, 49]. However, if the extent of the object is comparable to the propagation distance, it is necessary to switch to a multi-slice approach. For this approach, the object is split into several adjacent slabs. In a second step, the slabs are approximated by a series of equidistant planes. The phase and amplitude shift of each plane is applied to the wave function before propagating it to the next plane. This iterative procedure for obtaining $\psi_\omega(x,y,z=z_0)$ is illustrated in figure 2.6.

By comparing the wave field obtained with the projection approximation to the results obtained with the first Born approximation (see reference [43] for a detailed derivation), the refractive index decrement δ_ω can be related to the electron density $\rho_e(x,y,z)$ inside the object:

$$\delta_\omega(x,y,z) = \frac{2\pi r_e\rho_e(x,y,z)}{k^2} = \frac{r_e\rho_e(x,y,z)(hc)^2}{2\pi E^2}, \qquad (2.29)$$

where r_e is the classical electron radius, c is the speed of light, h is Planck's constant and E is the photon energy. This comparison allows to establish a connection between the wave picture and the particle picture. Under the first Born approximation, the phase shift caused by the object can be identified with coherent Thomson scattering at quasi-free electrons inside the object. In a classical description of Thomson scattering, the incoming electromagnetic wave sets quasi-free electrons in oscillations which then radiate spherical waves. Eq. 2.29 is only valid for quasi-free electrons which means that the binding energy of the electrons is small compared to the photon

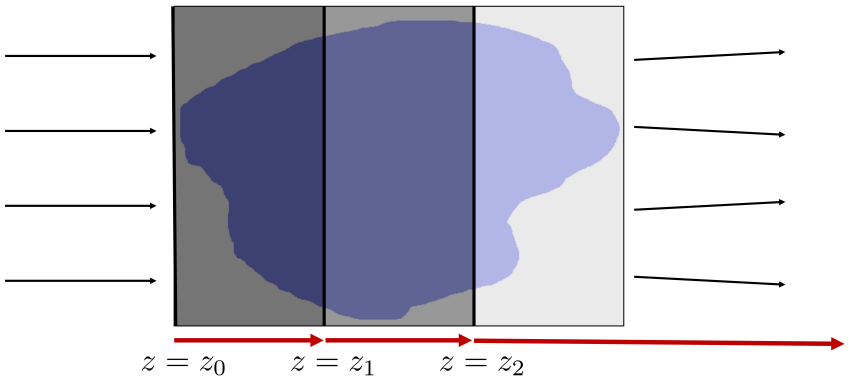

Figure 2.6: Illustration of the multi-slice projection approximation. A plane wave is incident on an object that is split into three slabs. The key idea of the projection approximation is to represent the three slabs by three planes (located at z_0, z_1 and z_2) with the same phase- and amplitude shifting properties. Free space propagation (represented by red arrows) is used to calculate the evolution of the wave field between the planes.

energy. For X-ray energies close to atomic absorption edges (see section 2.3.5), an additional dispersion correction term is introduced. In the diagnostic energy range, δ_ω is in the order of 10^{-6}.

2.3.3 Fresnel scaling theorem

The Fresnel scaling theorem relates the diffracted intensity pattern of an object that is produced by point-source illumination to the intensity pattern that is obtained by plane-wave illumination. In combination with the projection approximation, this theorem is particularly useful for numerical simulations of cone-beam geometries with the Fresnel approximation. Directly using Fresnel propagation (eq. 2.13) for numerical simulations with a point source is difficult because the Fourier transformed point source is not bandlimited [50]. Based on the convolution formulation of Fresnel propagation, the following relationship between the point source ($I_\omega^{(P)}$) and the plane wave ($I_\omega^{(\infty)}$) intensity pattern can be derived [43]:

$$I_\omega^{(P)}(x, y, z = l + d) = \frac{1}{M^2} I_\omega^{(\infty)} \left(\frac{x}{M}, \frac{y}{M}, z = l + \frac{d}{M} \right), \tag{2.30}$$

where l is the distance between the point source and the object, d is the distance between the object and the detector plane and $M = (l+d)/l$ is the geometric magnification. According to eq. 2.30, the point source intensity pattern at a distance d downstream of the object is calculated by:

- (i) taking the plane wave intensity pattern at a distance of d/M downstream of the object,
- (ii) magnifying it by a factor of M and
- (iii) dividing the obtained intensities by M^2.

Figure 2.7 illustrates the connection between the plane wave and the point source intensity patterns.

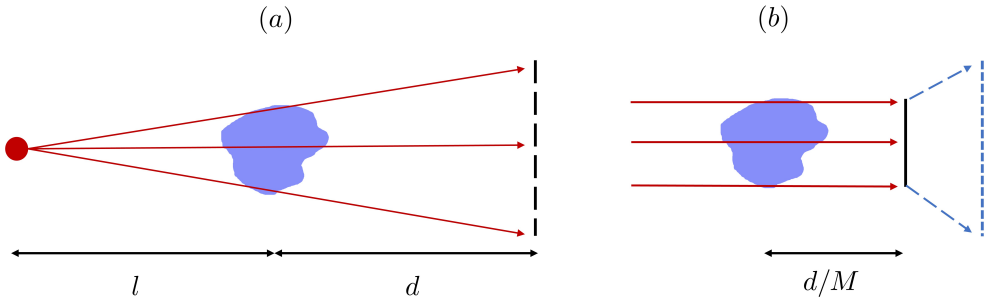

Figure 2.7: Illustration of the Fresnel scaling theorem. The point-source intensity pattern (a) at a propagation distance d from the object is directly proportional to the magnified plane wave intensity pattern (b) at the scaled propagation distance $d/M = dl/(l+d)$.

2.3.4 Eikonal equation (geometrical optics)

The eikonal equation establishes a connection between the wave theory and the ray theory of geometrical optics. The ray picture emerges from the wave theory in the limit of small wavelengths (compared to the other length scales). Writing $\phi_\omega(x,y,z)$ as

$$\phi_\omega(x,y,z) = \sqrt{I_\omega(x,y,z)}e^{i\psi_\omega(x,y,z)}, \tag{2.31}$$

the eikonal equation can be derived from the inhomogeneous Helmholtz equation (eq. 2.24, see reference [43] or [51] for details):

$$|\nabla\psi_\omega(x,y,z)|^2 = k^2 n_\omega(x,y,z)^2 + \frac{\nabla^2\sqrt{I_\omega(x,y,z)}}{\sqrt{I_\omega(x,y,z)}} \tag{2.32}$$

In the limit of geometrical optics ($\lambda \to 0$ or equivalently $k \to \infty$), the intensity $I_\omega(x,y,z)$ is approximately constant over length scales comparable to the wavelength λ. In this case, the second term in eq. 2.32 can be neglected. This corresponds to neglecting diffraction effects and yields the eikonal equation of geometrical optics:

$$|\nabla\psi_\omega(x,y,z)| = k n_\omega(x,y,z) \tag{2.33}$$

By defining rays as orthogonal trajectories to the wavefronts, a link between wave optics and geometrical optics is established. It is also possible to extend the standard theory of geometrical optics and consider diffraction effects by introducing diffraction rays [52]. Since wavefronts represent surfaces of constant phase ($\psi_\omega(x,y,z) = \text{const.}$), the local direction of a ray $\vec{t}(x,y,z)$ is given by the normalized phase gradient:

$$\vec{t}(x,y,z) = \nabla\psi_\omega(x,y,z)/|\nabla\psi_\omega(x,y,z)|. \tag{2.34}$$

With the eikonal equation of geometrical optics (eq. 2.33), the trajectory of rays through an inhomogeneous medium can be calculated.

X-ray refraction and the projection approximation The definition of rays as orthogonal trajectories to the wavefronts allows to connect the phenomenon of X-ray refraction to the projection approximation. More specifically, a relation between the refraction angle α and the gradient of the projected electron density can be established. Towards this end, we consider the situation depicted in figure 2.8 where a monochromatic plane wave $\phi_\omega(x,y,z) = e^{ikz}$ is incident on a homogeneous, purely phase shifting prism.

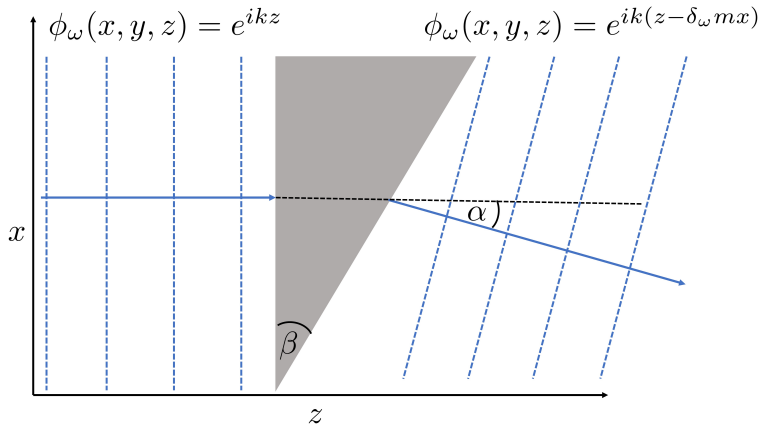

Figure 2.8: Illustration of the connection between the geometrical refraction angle α and the projection approximation. Wavefronts are indicated by dashed lines whereas the solid lines represent rays in the geometrical optics picture.

According to the projection approximation, the wave function immediately downstream of the prism is given by:

$$\phi_\omega(x,y,z) = e^{-ik\int \delta_\omega(x,y,z)dz} e^{ikz} = e^{ik(z-\delta_\omega mx)}, \tag{2.35}$$

where the slope $m = \tan\beta$ is related to the inclination angle β of the prism. Considering the definition of the ray direction (eq. 2.34), the refraction angle α (see figure 2.8) is calculated as:

$$\alpha \approx \tan\alpha = \left(\frac{\frac{\partial}{\partial x}\psi_\omega(x,y,z)}{\frac{\partial}{\partial z}\psi_\omega(x,y,z)}\right) = -\frac{\partial}{\partial x}\int \delta_\omega(x,y,z)dz = -\delta_\omega m \tag{2.36}$$

Since $\delta_\omega \ll 1$, the paraxial approximation ($\alpha \approx \sin\alpha \approx \tan\alpha$) can be used. Typical refraction angles for X-ray are in the order of 10^{-6} rad or less. Contrary to visible light, the real part of the refractive index is less than unity. X-ray are thus deflected towards regions of lower electron density. Combining eq. 2.36 with eq. 2.29, the refraction angle can be expressed in terms of the gradient of the projected electron density:

$$\alpha = -\frac{r_e(hc)^2}{2\pi E^2}\frac{\partial}{\partial x}\int \rho_e(x,y,z)dz. \tag{2.37}$$

2.3.5 X-ray attenuation

Various interaction mechanisms between X-rays and matter cause X-ray attenuation. In the wave-optical picture, this corresponds to a reduction of the intensity of an X-ray wave when transversing an object. By taking the squared modulus of the eq. 2.27, the ratio of the intensities upstream ($I_\omega^{(u)}$) and downstream ($I_\omega^{(d)}$) of the object can be calculated with the projection approximation:

$$\frac{I_\omega^{(d)}(x,z,y)}{I_\omega^{(u)(x,y,z)}} = \exp\left(-2k\int_{z=0}^{z=z_0}\beta_\omega(x,y,z)dz\right). \tag{2.38}$$

With the definition of the linear attenuation coefficient:

$$\mu_\omega = 2k\beta\omega \quad \text{or equivalently} \quad \mu(E) = \frac{2\pi hc}{E}\beta(E), \tag{2.39}$$

a generalized version of the Lambert-Beer law is derived:

$$\frac{I^{(d)}(E,x,y,z)}{I^{(u)}(E,x,y,z)} = \exp\left(-\int \mu(E,x,y,z)ds\right), \tag{2.40}$$

where E is the photon energy. Compared to eq. 2.38, the integral over the z-coordinate has been generalized by a line integral along the propagation direction of the wave (which corresponds to the X-ray path in the geometrical optics limit). In the particle picture, X-ray attenuation is identified with absorption or scattering of X-ray photons, which is a statistical process. The right-hand side of eq. 2.40 is thus interpreted as the probability that an incident photon is transmitted through the object. Assuming that N photons with energy E are incident on the object, the probability that q photons are transmitted follows a Bernoulli distribution [41]:

$$P(Q=q) = \binom{N}{q} p^q (1-p)^{N-q}, \quad p = \exp\left(-\int \mu(E,x,y,z)ds\right). \tag{2.41}$$

Eq. 2.41 is valid for a fixed number of incident photons N. As illustrated in section 2.1 , N is Poisson distributed. It can be shown [53] that a Bernoulli process with probability p acting on a Poisson process with mean λ results in a Poisson process with mean λp. Consequently, the number of transmitted photons q with a certain energy E is again Poisson distributed:

$$P(Q=q) = \frac{(\lambda p)^q}{q!} e^{-\lambda p}. \tag{2.42}$$

The attenuation coefficient μ is related to the interaction cross-section per atom $\sigma^a(E)$:

$$\mu(E,x,y,z) = \sigma^a(E) n_{\mathrm{p}}(x,y,z), \tag{2.43}$$

where n_{p} is the particle density. Instead of the attenuation coefficient μ, the mass attenuation coefficient is often considered, since it is independent of the particle density:

$$\frac{\mu(E)}{\rho} = \frac{\sigma^a(E)}{M_a}, \tag{2.44}$$

where ρ is the density and M_a the atomic mass of the element under consideration. The typical photon energy range for biomedical applications is between $15 - 150$ keV. In this energy range, there are three relevant interaction mechanisms between photons and matter: photoelectric absorption, Compton scattering and Rayleigh scattering. These three interaction mechanisms will be discussed in more detail in the following.

Photoelectric absorption Photoelectric absorption is characterized by the interaction of a bound electron with an incoming X-ray photon. The photon is absorbed and the photon energy is used to overcome the binding energy of the electron. Any excess energy is converted into kinetic energy of the expelled electron and the ionized atom. The vacancy created by the expelled electron is filled by an electron from an outer shell. This recombination process can lead to the emission of X-rays with characteristic energy (fluorescent lines). Provided that the photon energy exceeds the binding energy of the K-shell, most photoelectric absorption events occur with the K-shell electrons of the absorbing atom [54]. A quantum-mechanical description of photon-atom interactions is necessary to calculate the absorption cross-section for the photoelectric effect. Using first order perturbation theory, an approximate expression for the absorption cross-section σ_{ph} far away from absorption edges can be derived [35]:

$$\sigma_{\mathrm{ph}}(Z,E) = C_{\mathrm{ph}} Z^5 E^{-3.5}, \tag{2.45}$$

where C_{ph} is a proportionality constant and Z is the atomic number. To get more accurate results, various correction factors such as relativistic effects and screening by the atomic electrons

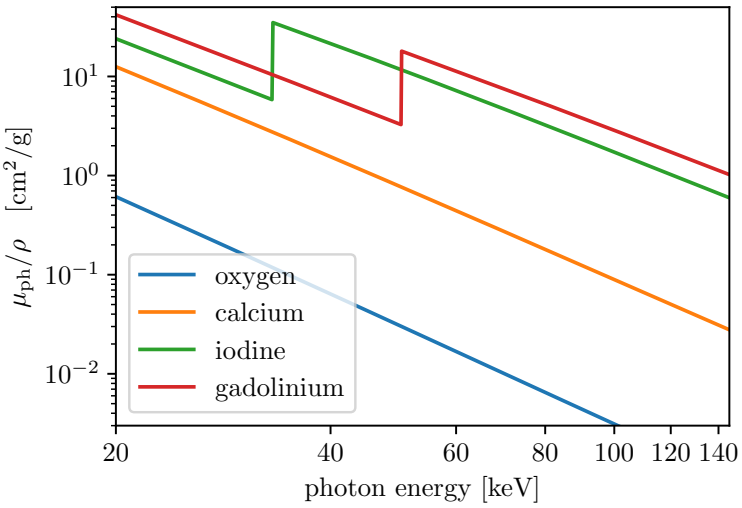

Figure 2.9: Photoelectric contribution to the mass attenuation coefficient for oxygen, calcium, iodine and gadolinium as a function of the photon energy. The attenuation data was taken from the NIST database [57, 58].

need to be considered [55]. The dependencies of the cross-sections on the atomic number and the photon energy can thus in general not be separated:

$$\sigma_{\mathrm{ph}}(Z, E) \neq C_{\mathrm{ph}} Z^{\alpha} E^{-\beta}. \tag{2.46}$$

Nevertheless, eq. 2.46 is often used to parameterize the photoelectric absorption cross-section for light elements ($Z \leq 20$). In this case, the absorption edges are located below the energy range of interest and the exponents α and β vary depending on the specific element or mixture to be parameterized [56] ($\alpha \in [4.0, 5.0]$, $\beta \in [3.0, 3.5]$). Figure 2.9 shows the photoelectric contribution to the mass attenuation coefficient as a function of the photon energy for four different elements: oxygen ($Z = 8$), calcium ($Z = 20$), iodine ($Z = 53$) and gadolinium ($Z = 64$). In this double-logarithmic plot, the graphs for oxygen and calcium closely resemble straight lines, which confirms that the photoelectric attenuation cross-section follows a power law. Moreover, as predicted by eq. 2.45, the overall attenuation cross-section increases strongly with increasing atomic number. For iodine and gadolinium, the binding energy of the K-shell electrons is in the diagnostic energy range. Once the photon energy increases above the binding energy, the K-shell electrons contribute to the photoelectric absorption which explains the sharp rise of the attenuation coefficients for iodine and gadolinium at $E = 33.2$ keV and $E = 50.2$ keV, respectively. As will be discussed in section 2.7.1, spectral X-ray imaging with contrast agents relies on the characteristic photoelectric attenuation pattern of heavier elements (such as iodine and gadolinium) to distinguish between contrast agent and other materials.

Compton scattering In the particle picture, Compton scattering represents a collision of a photon with a weakly bound orbital electron. The photon is deflected from its original path and due to the transferred momentum, the collision electron is ejected from the atom. Compton scattering is an incoherent scattering process because the incident photon loses a significant fraction of its energy. The energy E' of the X-ray photon after the collision can be derived from the conservation of momentum and energy:

$$E'(E, \theta) = \frac{E}{1 - \frac{E}{m_e c^2}(1 - \cos\theta)}, \tag{2.47}$$

where m_e is the electron mass and θ is the scattering angle. Based on quantum mechanical calculations, Klein and Nishina [59] derived an analytical expression for the differential scattering

cross section (DSCS) per unit solid angle for a free electron:

$$\frac{d\sigma_{\mathrm{C}}^{\mathrm{KN}}(E,\theta)}{d\Omega} = \frac{r_e^2}{2}\left(\frac{E'(E,\theta)}{E}\right)^2\left(\frac{E'(E,\theta)}{E} + \frac{E}{E'(E,\theta)} - \sin^2\theta\right), \tag{2.48}$$

where r_e is the classical electron radius. The total scattering cross section per electron is obtained by integrating over all possible scattering angles [60]:

$$\sigma_{\mathrm{C}}^{\mathrm{KN}}(E) = 2\pi r_e\left(\frac{1+\alpha}{\alpha^3}\left[\frac{2\alpha(1+\alpha)}{1+2\alpha} - \ln(1+2\alpha)\right] + \frac{\ln(1+2\alpha)}{2\alpha} - \frac{1+3\alpha}{(1+2\alpha)^2}\right), \quad \alpha = \frac{E}{m_e c^2}. \tag{2.49}$$

Since eq. 2.49 is valid for a single electron, the Compton-scattering contribution μ_{Compton} to the attenuation coefficient is proportional to the electron density n_e of the absorber [41]:

$$\mu_{\mathrm{Compton}}(E) = n_e \sigma_{\mathrm{C}}^{\mathrm{KN}}(E). \tag{2.50}$$

Within the Klein-Nishina approximation of scattering on free electrons, μ_{Compton} is thus independent of the atomic number of the attenuating material.
For elements with high atomic numbers and lower X-ray energies, electronic binding effects become more important and the Klein-Nishina approximation becomes increasingly inaccurate. The binding energy can be considered by multiplying the differential Klein-Nishina cross section with a correction function $S(x,Z)$ that depend on the momentum transfer variable $x(E,\theta) = \frac{E}{hc}\sin\left(\frac{\theta}{2}\right)$ and the atomic number of the atom under consideration [54, 61]:

$$\frac{d\sigma_{\mathrm{C}}^{\mathrm{a}}(E,Z,\theta)}{d\Omega} = \frac{d\sigma_{\mathrm{C}}^{\mathrm{KN}}(E,\theta)}{d\Omega} S(x(E,\theta),Z), \tag{2.51}$$

where $d\sigma_{\mathrm{C}}^{\mathrm{a}}(E,Z,\theta)/d\Omega$ is the Compton DSCS for one atom of the scattering material. Figure 2.10 (a) shows the (average) DSCS per electron for carbon (calculated according to eq. 2.51) as a function of the scattering angle at a photon energy of 80 keV together with the Klein-Nishina approximation (eq. 2.48) of the DSCS. Whereas the Klein-Nishina approximation is accurate for large scattering angles, it overestimates the DSCS in forward direction. This effect becomes more pronounced for lower energies and heavier elements. In figure 2.10 (b), the total Compton scattering cross-sections per electron for carbon and iodine are plotted as a function of the photon energy together with the Klein-Nishina approximation (see eq. 2.49). For carbon and high photon energies, the Klein-Nishina formula approximates the Compton scattering cross-section quite accurately. Since the orbital electrons are more tightly bound to the iodine atom, the approximation of scattering on free electrons breaks down in this case and the Klein-Nishina formula overestimates the total scattering cross-section.
Photons that are scattered in forward direction can reach the detector and deteriorate the image quality for both attenuation-based imaging and phase-contrast imaging. Moreover, these scattered photons are difficult to include in a physical model of the measurement acquisition process.

Rayleigh scattering Rayleigh scattering is a quasi-elastic interaction process between a photon and an atom. The photon is scattered on bound atomic electrons and deflected from its original path. In contrast to Compton scattering, the energy loss of the photon is negligible and no electrons are ejected from the atom. From the perspective of classical electrodynamics, this process is characterized by forced oscillations of all atomic electrons with the frequency of the incident X-ray radiation. Since the X-ray frequency is much larger than the resonant frequency of the electrons, all electrons radiate in antiphase with the incident radiation [60]. The scattered contributions from individual electrons are coherent with each other which leads to interference

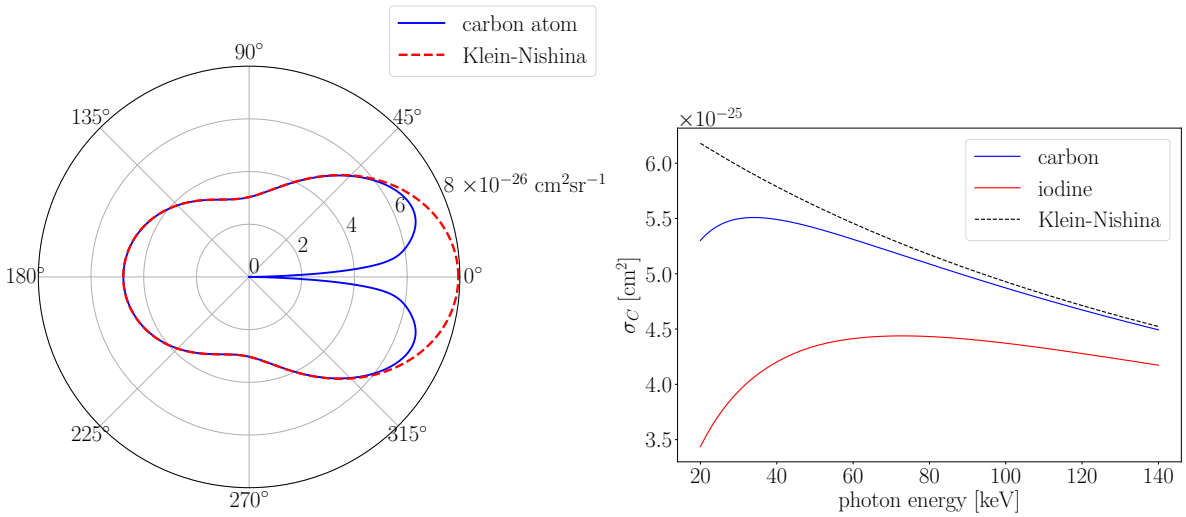

Figure 2.10: (a): Differential Compton scattering cross-section per electron as a function of the scattering angle for a carbon atom and 80 keV photon energy. The Klein-Nishina approximation to the differential scattering cross section is also shown for comparison. (b): Total Compton scattering cross-section per electron for carbon and iodine as a function of the photon energy in comparison with the Klein-Nishina approximation of scattering on free electrons.

effects between the individual scattered waves [60]. Consequently, the DSCS per electron is non-additive for Rayleigh scattering. It is thus more reasonable to consider the DSCS per atom:

$$\frac{d\sigma_{\mathrm{R}}^{\mathrm{a}}(E,Z,\theta)}{d\Omega} = \frac{r_e^2}{2}(1+\cos^2\theta)[F(x(E,\theta),Z)]^2, \tag{2.52}$$

where $F(x,Z)$ is the so-called atomic form factor for Rayleigh scattering and $x(E,\theta) = \frac{E}{hc}\sin(\frac{\theta}{2})$ is the momentum transfer variable. Similarly to Compton scattering, $F(x,Z)$ can be interpreted as a multiplicative correction factor to the DSCS for Thomson scattering on free electrons. Calculations of $F(x,Z)$ are based on the self-consistent Hartree field model [61, 60]. For large scattering angles, the individual scattering contributions interfere destructively, whereas there is constructive interference for small scattering angles:

$$\lim_{x\to 0} F(x,Z) = Z \tag{2.53}$$

Consequently, the probability of Rayleigh scattering is strongly increased in forward direction. This effect can be seen in figure 2.11 (a), where the Rayleigh DSCS for a carbon atom is plotted for photon energies of 20 keV and 70 keV. With increasing photon energy, the angular distribution of Rayleigh scattered photons becomes increasingly forward directed. The total Rayleigh scattering cross-section per atom is obtained by integrating the DSCS (eq. 2.52) over all angles. Although the exact results depend on the atomic number, the total Rayleigh scattering cross-section is approximately proportional to:

$$\sigma_{\mathrm{R}}^{\mathrm{a}}(E,Z) \propto \frac{Z^{\frac{8}{3}}}{E^2}. \tag{2.54}$$

The accuracy of this approximation increases for light elements and high photon energies. The double-logarithmic plot in figure 2.11 (b) shows the total Rayleigh scattering cross-section per atom as a function of the photon energy for carbon (Z= 6), calcium (Z=20) and gadolinium (Z=64). The three graphs deviate slightly from straight lines which corresponds to small inaccuracies in the proportionality law of eq. 2.54, especially for low energies. As expected from eq.

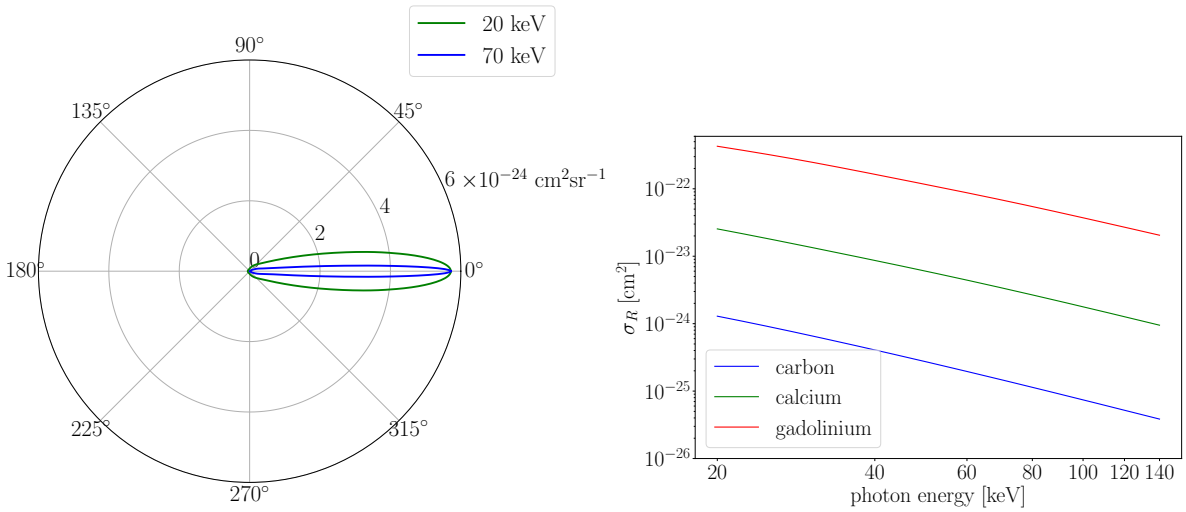

(a) (b)

Figure 2.11: (a): Rayleigh DSCS for a carbon atom at photon energies of 20 keV and 70 keV. (b): Total Rayleigh scattering cross section per atom as a function of the photon energy for carbon, calcium and gadolinium.

2.54, the scattering cross-section per atom rises strongly with increasing atomic number. Rayleigh scattering is an important source of undesired scattered radiation that reaches the detector because of its characteristic angular distribution that strongly favors scattering in forward direction.

Attenuation coefficient The total attenuation coefficient μ is the sum of the attenuation coefficients of the individual interaction mechanisms:

$$\mu = \mu_{\text{photoelectric}} + \mu_{\text{Compton}} + \mu_{\text{Rayleigh}}. \tag{2.55}$$

The relative importance of these three interaction mechanisms strongly depends on the photon energy and the atomic number of the attenuating element. This point is illustrated in figure 2.12, where the contributions of the three interaction mechanisms to the mass attenuation coefficient μ/ρ of carbon, calcium and gadolinium are plotted as a function of the photon energy. In figure 2.12, the Compton scattering contribution to the mass attenuation coefficient varies only slightly with the atomic number Z of the attenuating element. This observation is made plausible with the Klein-Nishina approximation, in which μ_{Compton} is proportional to the electron density and independent of Z (compare eq. 2.50 and eq. 2.49). As predicted by the corresponding proportionality laws (see eq. 2.54 and eq. 2.46), the mass attenuation coefficients for photoelectric absorption and Rayleigh scattering increase strongly for heavier elements. In the case of carbon, Compton scattering is the dominant contribution to the mass attenuation coefficient except for low photon energies ($E < 20$ keV), where photoelectric attenuation dominates due to the quickly increasing absorption probability for lower energies (see eq. 2.46). The most important contribution to the mass attenuation coefficient of Calcium is photoelectric absorption. However, for high photon energies, ($E > 85$ keV), the photoelectric absorption falls below the slowly varying Compton scattering contribution. In the low energy range, Rayleigh scattering becomes more important than Compton scattering but both scattering mechanisms remain an order of magnitude smaller than photoelectric absorption. The mass attenuation coefficient of heavier elements such as Gadolinium is clearly dominated by photoelectric absorption over the entire diagnostic energy range.

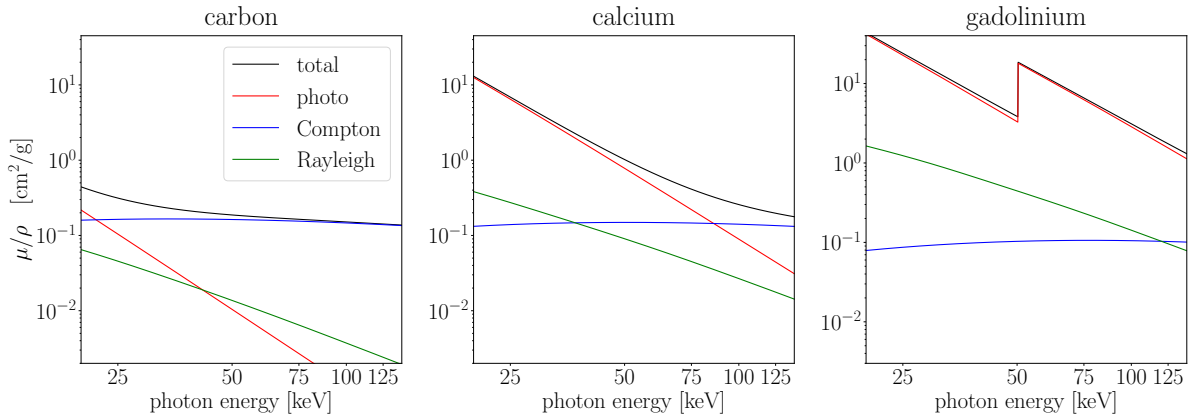

Figure 2.12: Contributions of photoelectric absorption, Compton scattering and Rayleigh scattering to the mass attenuation coefficients of carbon, calcium and gadolinium as a function of the photon energy.

Mixture rule The mixture rule states that the X-ray attenuation of a mixture or compound is simply given by a weighted sum of the individual elements:

$$\mu(E) = \sum_i \sigma_i^a(E) n_\mathrm{p}^i, \tag{2.56}$$

where $\sigma_i^a(E)$ is the total attenuation cross-section per atom of element i and n_p^i is the corresponding particle density. Equation 2.56 can be rewritten in terms of the mass attenuation coefficients:

$$\frac{\mu_c}{\rho_c} = \sum_i w_i \left(\frac{\mu}{\rho}\right)_i, \tag{2.57}$$

where μ_c and ρ_c are the attenuation coefficient and the density of the compound, respectively. The weight fraction of element i is denoted by w_i. The mixture rule is valid if the effects of the chemical or crystalline environment can be neglected [62]. This is typically the case for medical imaging and non-destructive testing since the X-ray energies are orders of magnitude higher than chemical binding energies ($\sim$ eV). As will be discussed in section 2.7.1, the mixture rule is an important assumption for spectral X-ray imaging algorithms that represent the energy-dependent attenuation by two or more basis materials.

2.4 X-ray detection

Modern X-ray detectors feature two essential components: an X-ray sensor that interacts with the incident photons and converts them into electrical signals and processing electronics that convert the electric signals into digital images. From a wave-optical point of view, only the intensity of the incident wave field can be detected. The phase of the wave field is not directly measurable. In the photon energy range between 15 and 150 keV, the three attenuation mechanisms discussed in section 2.3.5 (photoeffect, Compton and Rayleigh scattering) are relevant for describing the interactions between the incident photons and the sensor material.

2.4.1 Quantum efficiency and photon statistics

An ideal X-ray detector would convert all photons that interact with the sensor material into electrical signals. In this case, the detection efficiency is only limited by the quantum efficiency $\eta(E)$ of the sensor which represents the interaction probability of an incident photon with the

sensor material. The interaction probability is equivalent to the probability that the photon is not transmitted through the sensor:

$$\eta(E) = 1 - e^{-\mu(E)d_s}, \tag{2.58}$$

where d_s is the thickness of the sensor. Following the argumentation in section 2.3.5 (eq. 2.41 and 2.42), the number of X-ray photons with a certain energy E that are transmitted through the object and subsequently detected is again Poisson distributed:

$$P(Q = q) = \frac{\lambda^q}{q!}e^{-\lambda}, \quad \lambda = S(E)\exp\left(-\int \mu(E,x,y,z)ds\right)\eta(E), \tag{2.59}$$

where $S(E)$ is the average number of photons emitted from the source in direction of the detector (pixel). Since the sum of two independent Poisson processes with mean values λ_1 and λ_2 is a Poisson process with mean $\lambda_1 + \lambda_2$, the total number of detected photons also follows a Poisson distribution:

$$P(Q_\mathrm{t} = q) = \frac{\lambda_\mathrm{t}^q}{q!}e^{-\lambda_\mathrm{t}}, \quad \lambda_\mathrm{t} = \int S(E)\exp\left(-\int \mu(E,x,y,z)ds\right)\eta(E)dE. \tag{2.60}$$

Equation 2.60 is only valid for ideal detectors. In a real detector, various undesired effects can alter the photon statistics and introduce noise correlations between neighboring pixels.

2.4.2 Classification of X-ray detectors

X-ray detectors can be classified by the working principle of the radiation sensor and the processing electronics. For both of these essential components of an X-ray detector, there are two different design principles [63]:

Indirect vs. direct conversion sensors Figure 2.13 illustrates the difference between indirect and direct conversion sensors. In the case of indirect conversion, an incident X-ray photon is first converted into visible light photons in a scintillation crystal. Photoelectric absorption and Compton scattering generate electron-hole pairs in the scintillating material and visible light is emitted when they recombine [64]. The intensity of the visible light is proportional to the energy of the incident X-ray photon. Typical scintillator materials include cesium iodine (CsI), cadmium tungstate ($CdWO_4$) and gadolinium oxysulphide (Gd_2O_2S) [65, 41]. Since the most commonly used scintillator materials contain elements with high atomic numbers, indirect conversion sensor normally have a high quantum efficiency. In a second step, the visible light photons generate electron-hole pairs in a photodiode, which enables the conversion to an electrical signal. Some of the scintillation light does not reach the photodiode or spreads to neighboring pixel (scintillator blur). The scintillator blur causes noise correlations between neighboring pixels and limits the spatial resolution of the imaging system.
In a direct conversion detector, the radiation sensor consists of a layer of semiconductor material to which a high bias voltage is applied. The most commonly used materials are Si, Ge, GaAs, CdTe and CdZnTe. Similar to the indirect conversion mechanism, an incident X-ray photon creates electron-hole pairs inside the sensor material. The generated charges are accelerated towards readout electrodes by the electric field before they can recombine. The direct conversion mechanism is more efficient because losses in the scintillator material are avoided. To be more precise, an impinging X-ray photon creates more electron-hole pairs in a direct conversion detector compared to the indirect conversion scheme. Consequently, a direct conversion sensor can achieve a higher signal-to noise ratio. This is particularly relevant in combination with energy-integrating readout electronics, which will be discussed in the next paragraph. In some materials, however, the high bias voltage can generate dark currents [66] (i.e. an electrical current without X-ray illumination).

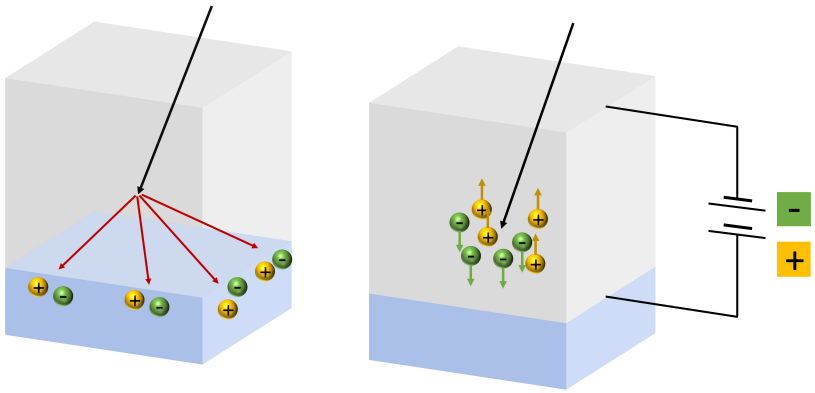

Figure 2.13: Working principle of direct and indirect conversion sensors. In the case of indirect conversion, an incident X-ray photon is first converted into visible light photons in a scintillation crystal. In a second step, the visible light photons generate electron-hole pairs in the semiconductor layer of a photodiode, which enables the conversion to electrical signals. In a direct conversion detector, a high bias voltage is applied to the semiconductor sensor. An incident X-ray photon generates free charges that are accelerated towards the electrodes by the electric field.

Energy-integrating vs. single pulse processing electronics Energy-integrating detectors use storage capacitors to integrate the electrical charges from all impinging photons in a certain time interval. After the integration time has passed, the discharge current of the capacitor is measured and digitized. Due to the analog integration step, energy-integrating detectors can measure comparatively high photon fluxes. The measured signal is proportional to the total photon energy deposited in the sensor. Besides quantum noise (see eq. 2.60), electronic noise, which is generated by bias currents and thermal noise in the analog parts of the detection electronics, represents an additional contribution to the total noise level [67].
Single pulse-processing detectors feature dedicated application-specific integrated circuits (ASICs) that enable an analysis of the individual signal pulses generated by impinging photons. Ideal single pulse-processing detectors thus count the number of photons that interact with the sensor material in a certain time interval. By only counting signal pulses that exceed a certain threshold, electronic readout noise can be avoided. This is particularly relevant for ultra low-dose imaging applications, where the signal level is comparable to the electronic noise level [68, 69, 70]. In contrast to energy-integrating detectors, the contribution of each photon to the total signal strength is independent of the photon energy. This can be advantageous for some X-ray imaging applications because single pulse-processing detectors assign more weight to low energy photons which typically carry more contrast information [71, 72]. Some ASICs implement several discriminators (i.e. several signal threshold levels) which allows the acquisition of energy resolved measurements. The additional energy information enables spectral X-ray imaging and material differentiation (see section 2.7.1). Moreover, the energy information can be used to improve the contrast-to-noise ration (CNR) of conventional attenuation images by an energy-weighted addition [73, 74, 75, 76, 77]. Compared to energy-integrating detectors, single pulse processing electronics are expensive and difficult to manufacture because of the relatively complicated processing electronics that are required for each pixel. Moreover, the densely packed electronics have a high power consumption which can cause heat dissipation problems [78].

Modern detector systems for medical radiography and CT applications mostly rely on indirect conversion sensors in combination with energy-integrating processing electronics. Driven by

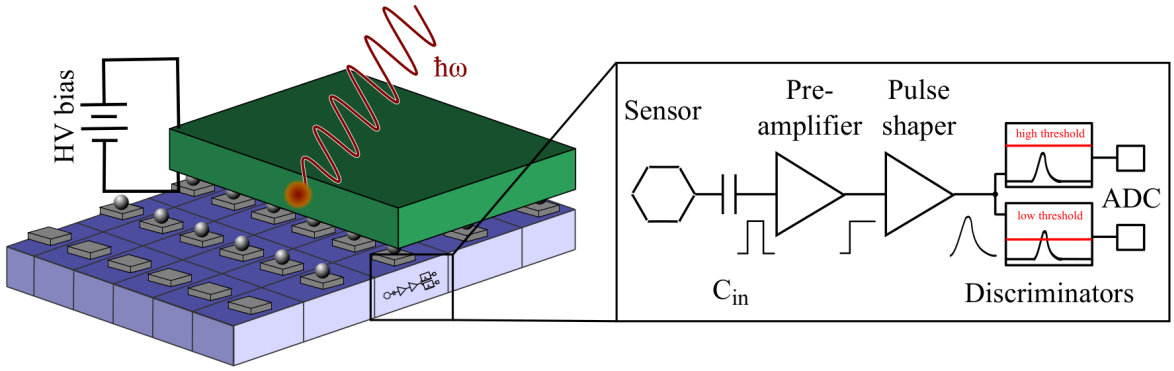

Figure 2.14: Architecure of a photon-counting detector. The detection system consists of two main components: A semiconductor sensor and readout ASICs which are connected to the sensor via micro-bumps. The readout ASICs (see zoom-in) contain components for pulse-processing, energy discrimination and analog-to-digital conversion. (Reprinted, with permission, from reference [84].)

recent advances in hybrid photon-counting detector technology [79, 80, 81], the combination of direct conversion sensors and single pulse processing (or single photon-counting) electronics has gained increasing popularity over the last years. Although applications in clinical routine are still rare, many pre-clinical and clinical studies [27, 28, 29, 30, 31, 32, 33, 34] have investigated the potential benefits of photon-counting detector technology for medical imaging. Since photon-counting detectors (PCDs) offer many advantages compared to other spectral X-ray imaging technologies (see section 2.7.1), they will be discussed in more detail in the following.

2.4.3 Photon-counting detectors

The hybrid pixel PCD technology was developed at CERN in the late 1980's for the analysis of particle collider experiments [11, 12]. By modifying the readout ASICs, the PCD technology has been transferred to other imaging applications [79], such as crystallography [82] , adaptive optics [83] and X-ray imaging.
Figure 2.14 shows the architecture of a PCD. The detection system features two main parts: A direct conversion semiconductor sensor and an array of readout ASICs (one ASIC for each detector pixel). These two components are manufactured individually and then connected via a bump-bonding process. The hybrid architecture allows an individual optimization of the sensor and readout ASICs. Moreover, due to the modular concept, it is possible to connect different sensor materials to the same readout ASIC. The signal processing chain of the readout ASICs consists of a pre-amplifier, a pulse shaper, one or more discriminators and an analog-to-digital converter. Since the induced charge by a single X-ray photon is very small ($\sim 10^{-15}$ C), the electronic signal is amplified by the pre-amplifier before further processing of the signal pulses. In the absence of undesired detector effects (see next paragraph), the number of electron-hole pairs generated in the semiconductor sensor and thus also the pulse height of the amplified signal is directly proportional to the energy of the impinging X-ray photon. The pulse shaper is essentially a band-pass filter that improves the signal-to-noise ratio (SNR) of the pulse generated by the pre-amplifier by amplifying the signal and attenuating electronic noise. The band-pass filter also changes the shape of the pulse in the temporal domain and is thus referred to as pulse shaper. The shaped signal is then fed to the discriminators which represent the interface between the analog and digital part of the signal processing chain. The discriminator detects a signal if the pulse height exceeds the corresponding pre-specified threshold value. In other words, each discriminator only detects photons with energies that exceed the energy associated with the corresponding threshold. Finally, the signal is converted from analog to digital and for

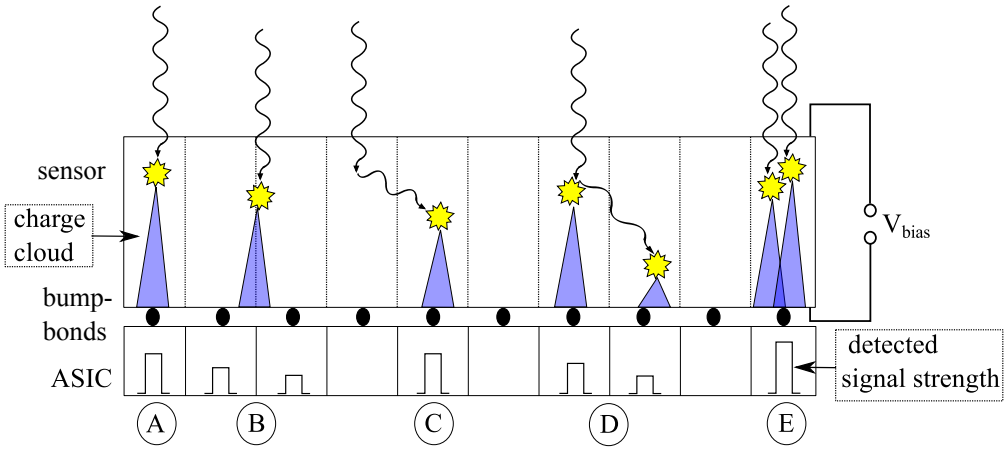

Figure 2.15: Sensor effects in photon-counting detectors that affect the spectral and/or the spatial resolution. (A): The incident photon interacts by photoelectric absorption and the entire charge cloud is detected by the pixel in which the interaction took place. This represents the ideal scenario with no performance degrading effect. (B): The photon interacts at the boundary between two adjacent pixels and the charge cloud spreads over two adjacent pixels which leads to charge-sharing. (C): Rayleigh scattering in the sensor causes the photon to be detected in a different pixel. (D): Similar to charge-sharing, the photon energy can be distributed to different pixels by Compton scattering in the sensor material. (E): Pulse pile-up is caused by two photons hitting the same detector pixel in a time interval that is shorter than the temporal resolution of the detector. (Reprinted, with permission, from reference [85].)

each signal that exceeds the threshold value, a counter is incremented.

Performance degrading effects in PCDs For an ideal detector with N discriminators it is possible to detect photons in N sharply defined energy windows (also called energy bins) by subtracting adjacent counter values. However, as depicted in figure 2.15, several performance degrading effects influence both the spatial and spectral resolution of a real PCD. Moreover, they can cause deviations from the idealized assumption of uncorrelated Poisson statistics (compare eq. 2.60).

Charge-sharing After generation of a charge cloud in the semiconductor material by an incident photon, the free charge carriers drift to the electrodes. Diffusion effects and the Coulomb force cause the charge cloud to spread out laterally. If the charges are created near a pixel boundary, parts of the charge cloud may be detected in several neighboring pixels (see figure 2.15 (B)). If the charge fractions exceed the discriminator threshold, the detector falsely registers several photons in adjacent pixels with a fraction of the original photon energy. Charge-sharing thus causes distortions in the energy response of the PCD and produces a low energy tail in the observed spectrum [79]. Furthermore, the spatial resolution of the detector is reduced since the spreading charge cloud effectively acts as a blurring mechanism in the spatial domain. The impact of the charge-sharing distortions depend on the sensor material, the pixel size, the thickness of the semiconductor sensor and the bias voltage. Geometrical considerations lead to the conclusion that the probability of charge sharing decreases with increasing detector pixels size and decreasing thickness of the semiconductor sensor. The sensor thickness of the PCD thus defines the tradeoff between quantum efficiency and charge sharing distortions. This tradeoff can be partially mitigated by using high-Z semiconductor materials that provide high quantum efficiencies with comparatively thin sensor layers. A higher bias voltage leads to faster drift times and consequently, the charge cloud has less time to spread out which reduces the amount of charge sharing. Some PCDs implement an on-chip charge-sharing correction scheme [86, 87].

The ASICs feature an additional network of charge summing circuits that allows neighboring pixels to communicate with each other and detect coincident pulses. The detected photon is then only assigned to the pixel with the highest charge sharing fraction, which represents the location where the photon most likely hit the sensor. A disadvantages of the additional charge-sharing correction electronics is the increased processing time which potentially leads to undesired effects due to pulse pile-up (see next but one paragraph).

Scattering and fluorescent emission X-ray photons can be deflected from their original path by Rayleigh and Compton scattering interactions with the sensor material which potentially results in a detection event in a neighboring pixel. Whereas Rayleigh scattering (see figure 2.15 (C)) only degrades the spatial resolution, the effect of Compton scattering is similar to charge-sharing (see figure 2.15 (D)). A fraction of the photon energy is transferred to the scattered electron before the scattered photon is detected in a neighboring pixel. Compton scattering thus causes a relatively homogeneous background in the measured photon energy spectra [63]. As discussed in section 2.3.5, the relative probability of photoelectric absorption compared to scattering interactions increases with increasing atomic number of the sensor material. Furthermore, the increased total absorption probability for high-Z materials leads to a shorter mean free path length of the scattered photons. Consequently, performance degrading effects due to scattering interactions can be partly suppressed by using high-Z semiconductor sensors.
Photoelectric absorption as well as Compton scattering processes eject an electron from the interacting atom. The vacancy is filled by one of the electrons from the outer shells. This rearrangement can lead to the emission of a characteristic fluorescence photon that might reach another detector pixel or escape from the sensor. Fluorescent emission thus influences both the spatial and spectral resolution of the PCD. Despite the larger attenuation coefficient for high-Z materials, the mean free path length of fluorescence photons increases with the (effective) atomic number of the sensor material because of the increased energy of the K-alpha emission lines. In silicon, the mean free path length is only 12 μm, whereas in a CdTe sensor the mean free path lengths are 110 μm and 58 μm for photons emitted from Cd and Te atoms, respectively [79].

Pulse pile-up Collecting the charges and processing the signal after an x-ray photon has hit a detector pixel needs a certain amount of time, the so-called deadtime of the detector. Given the stochastic nature of the photon arrival times, it is possible that the signal pulses of two or more photons (that interact with the same detector pixel) overlap. The probability for such a pulse pile-up event depends on the photon flux and the detector deadtime. Moreover, the detector pixel size plays an important role because it strongly influences the effective photon flux per detector pixel. For state-of-the-art PCDs, pile-up effects become relevant at a flux of approximately $10^6 - 10^8$ photon counts per second and square millimeter [78]. Depending on the time delay between the signal pulses generated by two incident photons, two types of pulse-pile effects can be differentiated. Peak pile-up occurs if the two pulses strongly overlap, i.e. the difference in photon arrival time is much smaller than the deadtime. As illustrated in figure 2.16 (a), the individual pulses cannot be resolved and the ASIC only registers one pulse with an amplitude that is approximately given by the sum of the amplitudes of the individual pulses. The so-called tail pile-up scenario (see figure 2.16 (b)) is characterized by a larger time delay between the two pulses. The two pulses can thus be separated, however the tail of the first pulse overlaps with the peak of the second pulse. In the example of figure 2.16 (b), the second pulse is falsely registered in the high energy threshold.
In summary, pulse pile-up effects lead to losses in the number of registered photon and distortions of the spectral detector response. In the last years, analytical pulse pile-up model have been developed [88, 89, 90, 91] that can correct or mitigate the pulse pile-up induced bias. Nevertheless, the spectral distortions as well as the lost photon counts have an adverse effect on the noise level for spectral X-ray imaging applications.

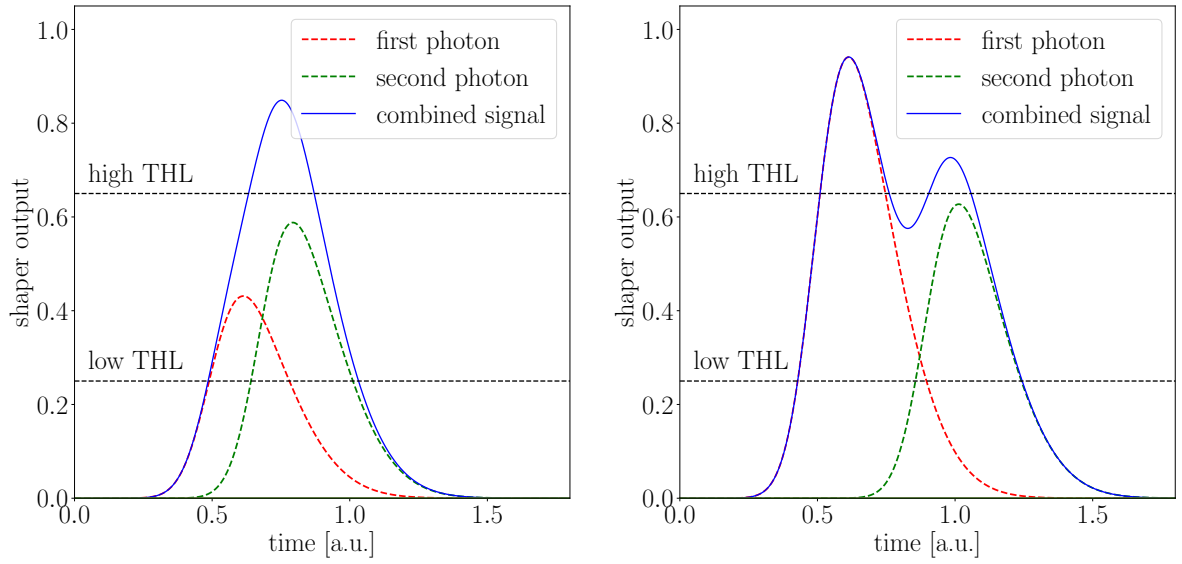

Figure 2.16: Illustration of the difference between the peak pile-up effect (a) and the tail pile-up effect (b) in photon counting detectors. In case of the peak pile-up effect, the two signal pulses that were generated by two incident photons overlap so closely that the individual pulses cannot be resolved. Instead of registering two photon counts in the low energy threshold, the readout ASIC interprets the combined signal as one photon count in the high energy threshold. In the tail pile-up scenario, the two pulses can be separated, however the tail of the first pulse overlaps with the peak of the second pulse. Due to this overlap, the second pulse is falsely registered in the high energy threshold.

Sensor polarization Crystal defects can temporarily trap electrons and holes that have been generated by incicident X-ray photons. The trapped charges locally distort the electric field that is induced by the bias voltage. This results in a decreased charge collection efficiency and a temporarily and spatially varying spectral response of the PCD. Compared to silicon, this effect is much more pronounced in high-Z semiconductor sensors (e.g. CdTe) where the defect concentration and the material quality is much more difficult to control. Polarization effects can be minimized by using high bias voltages and operating the PCD at low temperatures [79]. Polarization effects an also be mitigated by periodically resetting the bias voltage [92, 93].

Electronic noise Electronic noise arising in the analog processing electronics influences the signal that is fed to the discriminators and thus degrades the energy resolution of the PCD. However, contrary to energy-integrating detectors, the threshold-based processing architecture of PCDs generally prevents that electronic noise influences the total number of registered photon counts.

Many design choices for PCDs involve a tradeoff between different performance degrading effects. For example, a smaller pixel size reduces pulse-pile up effects, but increases the amount of charge-sharing. Compared to silicon, high-Z semiconductor sensor like CdTe increase the quantum efficiency and reduce charge-sharing effects (because the sensor thickness can be smaller). However, polarization effects and manufacturing imperfections play a larger role since CdTe and other high-Z semiconductors are much more difficult to process. This can lead to a spatially varying spectral response as well as temporal drifts. Both effects have to be considered in the signal processing pipeline to obtain artifact-free images.

2.5 Statistical signal processing

As discussed in the previous sections, X-ray generation, attenuation and detection are stochastic processes and any radiography or CT measurement is thus contaminated by noise. In order to obtain high quality X-ray images, the stochastic nature of these processes needs to be taken into account, in particular for low-dose measurements. This section considers a selection of statistical signal processing techniques that are particularly relevant for X-ray imaging and will be used extensively in this thesis: maximum a posteriori estimators, maximum likelihood estimators and the Cramér-Rao lower bound. We start by explaining some fundamental mathematical concepts that form the basis for the aforementioned statistical signal processing techniques.

Inverse Problems Inverse problems are characterized by determining system parameters from a set of observations that were caused by the parameters to be determined. They frequently occur in physics or related fields and are particularly relevant for X-ray imaging. Many of the results presented in this work are based on solving inverse problems that occur in spectral imaging, (spectral) differential phase-contrast imaging and CT reconstruction. Given a set of parameters, the expected observations can be calculated by constructing a *forward model* which is based on the physical laws governing the measurement process. The forward model f is a function that maps the parameter values to the expected observations:

$$f : \mathbb{R}^M \mapsto \mathbb{R}^N, \ \hat{\vec{x}} = f(\vec{\theta}) \tag{2.61}$$

This forward problem typically has a unique solution that is comparatively easy to calculate. The associated inverse problem, however, is typically much more difficult to solve. First of all, an analytical inversion of the forward model is not possible in most cases. For example, directly inverting the forward model is impossible if the inverse problem is underdetermined, i.e. the number of model parameters is larger than the number of observations. In this case, one

iteratively minimizes the difference between the forward model and the measured data:

$$\vec{\theta} = \arg\min_{\vec{\theta}} G(f(\vec{\theta}), \vec{x}), \tag{2.62}$$

where the function G is a similarity measure between the forward model and the measured data. A common choice is the mean-squared error:

$$G(f(\vec{\theta}), \vec{x}) = \frac{1}{N}(\hat{\vec{x}} - \vec{x}) \cdot (\hat{\vec{x}} - \vec{x}). \tag{2.63}$$

Moreover, inverse problems are often *ill-posed* or *ill-conditioned*. An ill-posed problem violates one of the following three conditions [94]:

- a solution exists
- the solution is unique
- the solution depends continuously on the observed variables.

For an ill-conditioned inverse problem, a small change in the observed data causes a large change of the calculated model parameters. This means that the solution of the inverse problem is non-robust because experimental measurements inevitably contain some amount of noise. A common strategy to stabilize the solution of ill-conditioned or ill-posed inverse problem is the introduction of additional assumptions about the model parameters via a so-called *regularization* term $R(\vec{\theta})$. The corresponding modified optimization problem can be expressed as:

$$\vec{\theta} = \arg\min_{\vec{\theta}} \left[G(f(\vec{\theta}), \vec{x}) + \lambda R\left(\vec{\theta}\right) \right], \tag{2.64}$$

where λ is a regularization parameter that controls the tradeoff between data fidelity and the assumptions about $\vec{\theta}$. The connection between the regularization term and Bayesian statistics will be discussed in section 2.5.1.

Estimator An estimator $\hat{\theta}$ is a rule for calculating unknown parameters $\vec{\theta}$ based on observed data $\vec{x}$. It can be viewed as a function that maps the observed data to an estimate of the desired parameters:

$$\hat{\theta} : \mathbb{R}^N \mapsto \mathbb{R}^M, \quad \vec{\theta} = \hat{\theta}(\vec{x}), \tag{2.65}$$

where N is the number of measurements and M is the number of parameters. The restriction to real coordinate space is sufficient for X-ray imaging applications. In many cases (in particular for inverse problems), $\hat{\theta}(\vec{x})$ can only be implicitly defined. Since the measured data $\vec{x}$ are random variables and the estimator is a function of the data, it can also be viewed as a random variable. A simple X-ray imaging example is the estimation of the line integrals of the linear attenuation coefficient from intensity measurement. Assuming a monochromatic source, the Lambert-Beer law can be written as (compare eq. 2.40):

$$y_i = b_i e^{-l_i}, \quad l_i = \int \mu(x, y, z) ds_i, \tag{2.66}$$

where y_i is the measured intensity for detector pixel i, $l_i = \int \mu(x, y, z) ds_i$ denotes the line integral along the path from the source to detector pixel i and b_i is the reference intensity measured without the sample in the beam path. Following the notation of eq. 2.65, $\vec{\theta} = (l_1, ..., l_N)^T$, $\vec{x} = (y_1, ..., y_N, b_1, ..b_N)^T$ and $\hat{\theta}$ is given by:

$$\hat{\theta} : \mathbb{R}^{2N} \mapsto \mathbb{R}^N, \quad l_i = -\ln\left(\frac{y_i}{b_i}\right), \tag{2.67}$$

where N is the number of detector pixel.
There are three commonly used metrics to characterize the performance of an estimator. The bias of an estimator is the expectation value of the difference between the estimated parameters $\hat{\theta}(\vec{x})$ and the ground truth parameter values $\vec{\theta}^{\mathrm{gt}}$:

$$\mathrm{bias}\left(\hat{\theta}\right) = \mathrm{E}\left[\hat{\theta}(\vec{x}) - \vec{\theta}^{\mathrm{gt}}\right] = \mathrm{E}\left[\hat{\theta}(\vec{x})\right] - \vec{\theta}^{\mathrm{gt}}. \tag{2.68}$$

With the probability density function (pdf) of the measured data given the ground truth parameter values $p(\vec{x}|\vec{\theta}^{\mathrm{gt}})$, the expectation value in eq. 2.68 can be calculated explicitly for each component by:

$$\mathrm{E}\left[\hat{\theta}_i\right] = \int p\left(\vec{x}|\vec{\theta}^{\mathrm{gt}}\right)\hat{\theta}_i(\vec{x})dx_1, ..., dx_M. \tag{2.69}$$

The mean-squared error (MSE) and the variance σ^2 are defined component-wise as:

$$\begin{aligned} \mathrm{MSE}(\hat{\theta}_i) &= E\left[\left(\hat{\theta}_i(\vec{x}) - \theta_i^{\mathrm{gt}}\right)^2\right] \\ \sigma^2(\hat{\theta}_i) &= E\left[\left(\mathrm{E}\left[\hat{\theta}_i(\vec{x})\right] - \hat{\theta}_i(\vec{x})\right)^2\right]. \end{aligned} \tag{2.70}$$

Noise-induced bias of an estimator In the case of a nonlinear estimator, i.e. a nonlinear mapping $\hat{\theta}(\vec{x})$ from the observed data $\vec{x}$ to parameter space, noise-induced bias can occur. This effect occurs even if the estimator is unbiased for noise-free data (i.e. $\vec{\theta}^{\mathrm{gt}} = \hat{\theta}(E\left[\vec{x}\right])$) because taking the expectation value and applying a nonlinear function are not interchangeable:

$$E\left[\hat{\theta}(\vec{x})\right] \neq \hat{\theta}(E\left[\vec{x}\right]) \tag{2.71}$$

Figure 2.17 illustrates the effect of noise-induced bias for a simple one-dimensional estimation problem. A well-known X-ray imaging example for noise-induced or statistical bias is the estimation of the line integrals that was discussed in the last paragraph (see eq. 2.66). With the help of Jensen's inequality, one can show that the line integrals (and thus the reconstructed attenuation coefficients in a CT scan) are systematically overestimated for low photon statistics [53, 95].

2.5.1 Maximum a posteriori and maximum-likelihood estimators

Maximum a posteriori (MAP) estimation is an approach for solving inverse problems that is based on Bayesian statistics. The idea is to find the parameter values $\vec{\theta}_{\mathrm{MAP}}$ that maximize the probability of the parameters given the measured data:

$$\vec{\theta}_{\mathrm{MAP}} = \arg\max_{\vec{\theta}} P\left(\vec{\theta}\,|\,\vec{x}\right) \tag{2.72}$$

Since $P\left(\vec{\theta}\,|\,\vec{x}\right)$ cannot be accessed directly for inverse problems, eq. 2.72 is reformulated with Bayes's theorem:

$$\vec{\theta}_{\mathrm{MAP}} = \arg\max_{\vec{\theta}} P\left(\vec{\theta}\,|\,\vec{x}\right) = \arg\max_{\vec{\theta}} \frac{P\left(\vec{x}\,|\,\vec{\theta}\right)P\left(\vec{\theta}\right)}{P(\vec{x})} = \arg\max_{\vec{\theta}} P\left(\vec{x}\,|\,\vec{\theta}\right)P\left(\vec{\theta}\right), \tag{2.73}$$

where the last equality holds because $P(\vec{x})$ is independent of $\vec{\theta}$. The function $\vec{\theta} \mapsto P\left(\vec{x}|\vec{\theta}\right)$ is known as *likelihood function* since it maps a parameter vector to the likelihood of the measured data given the parameter vector. The forward model allows a direct calculation of the likelihood function if the statistical distribution of the measured data is known or can be approximated.

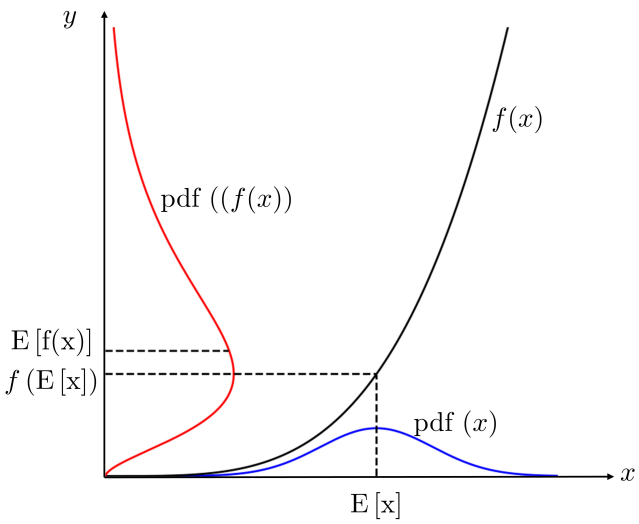

Figure 2.17: Noise-induced bias of a one-dimensional nonlinear estimator. The nonlinear mapping function $\hat{\theta}(x) = x^2$ skews the Gaussian distribution of the measured data after mapping to parameter space. In this case, taking the expectation value and applying $\hat{\theta}(\cdot)$ are not interchangeable.

A common example is the assumption of independent Gaussian noise. In this case, $P\left(\vec{x} \mid \vec{\theta}\right)$ is given by:

$$P\left(\vec{x} \mid \vec{\theta}\right) = \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(\hat{x}_i - x_i)^2}{2\sigma_i^2}\right), \tag{2.74}$$

where σ_i is the variance of x_i. Contrary to frequentist statistics, Bayesian statistics interprets probabilities as a degree of belief in an event. The parameters $\vec{\theta}$ are thus also treated as random variables and $P(\vec{\theta})$ represents the prior (i.e. before seeing the data) probability distribution of $\vec{\theta}$. Prior knowledge and assumptions about $\vec{\theta}$ that are incorporated in $P(\vec{\theta})$ can be used to stabilize the (possibly ill-conditioned) optimization problem and reduce the variance of the estimator. However, the incorporation of prior knowledge typically leads to biased estimates. In order to increase the numerical stability of the optimization problem, one optimizes the logarithm of the *penalized likelihood function* $P\left(\vec{x} \mid \vec{\theta}\right) P\left(\vec{\theta}\right)$. Since the logarithm is a strictly monotonic function, this transformation will not change the position of the maximum. By convention, the negative penalized log-likelihood function is minimized, which is mathematically equivalent to maximizing the penalized log-likelihood function:

$$\vec{\theta}_{\mathrm{MAP}} = \arg\min_{\vec{\theta}} \left[-L(\vec{\theta}) + \lambda R(\vec{\theta})\right], \tag{2.75}$$

where $L(\vec{\theta}) = \ln\left(P(\vec{x}|\vec{\theta})\right)$ is the log-likelihood function and $\lambda R(\vec{\theta}) = -\ln\left(P(\vec{\theta})\right)$ is the regularization term. By writing the prior probability of the parameters as:

$$P(\vec{\theta}) = e^{-\lambda R(\vec{\theta})}, \tag{2.76}$$

the regularization term that was introduced in the last section to stabilize the solution of inverse problems (compare eq. 2.64) can be interpreted as prior knowledge about the optimization parameters within the framework of Bayesian statistics.
In the case of a uniform prior that assigns the same prior probability to all parameter configurations $\vec{\theta}$, MAP estimation reduces to maximum-likelihood (ML) estimation:

$$\vec{\theta}_{\mathrm{ML}} = \arg\min_{\vec{\theta}} \left[-L(\vec{\theta})\right]. \tag{2.77}$$

The ML estimator is a common choice for inverse and (other) problems because it can be easily constructed if a forward model and knowledge about the measurement statistics are available. It naturally incorporates the measurement statistics by giving more weight to more reliable measurements. Furthermore, the ML estimator is asymptotically unbiased and asymptotically achieves the minimum variance for an unbiased estimator (Cramér-Rao lower bound, see next section) in the limit of low noise levels [96].

2.5.2 The Cramér-Rao lower bound

The Cramér-Rao lower bound (CRLB) is a powerful tool from estimation theory that predicts a lower bound for the variance of an unbiased estimator. Given a parameter vector $\vec{\theta}$, it can be shown that [97, 96]:

$$\boldsymbol{C}(\vec{\theta}) - [\boldsymbol{F}(\vec{\theta})]^{-1} \geq 0, \tag{2.78}$$

i.e. the matrix $\boldsymbol{C}(\vec{\theta}) - [\boldsymbol{F}(\vec{\theta})]^{-1}$ is positive semidefinite. Here, $\boldsymbol{C}(\vec{\theta})$ is the covariance matrix of $\vec{\theta}$:

$$\boldsymbol{C}_{uv} = \mathrm{E}\left[(\theta_u - E(\theta_u))(\theta_v - E(\theta_v))\right]. \tag{2.79}$$

The Fisher information matrix $\boldsymbol{F}(\vec{\theta})$ is the expectation value of the curvature of the negative log-likelihood function:

$$\boldsymbol{F}_{uv} = \mathrm{E}\left[-\frac{\partial^2 L(\vec{\theta})}{\partial\theta_u \partial\theta_v}\right]. \tag{2.80}$$

From equation 2.78, a lower bound for the variance of the estimated parameters can be deduced:

$$\boldsymbol{C}_{uu} = \sigma^2(\theta_u) \geq (\boldsymbol{F}^{-1})_{uu}. \tag{2.81}$$

An unbiased estimator is called *efficient* if its covariance matrix meets the CRLB:

$$\boldsymbol{C}(\vec{\theta}) = \boldsymbol{F}(\vec{\theta}). \tag{2.82}$$

In this work, the CRLB is used to predict the covariance matrices of different estimators for spectral, differential phase-contrast and spectral differential phase-contrast X-ray imaging. This allows the optimization of various image acquisition parameters as well as a performance comparison between different X-ray imaging methods (see section/chapter 8).

2.6 Reconstruction techniques for Computed Tomography

CT reconstruction is the process of inferring the three-dimensional distribution of object properties from a set of projection measurements obtained at various angles around the object. In this section, we will focus on CT reconstruction of the linear attenuation coefficient. The basic principles of CT reconstruction for attenuation-based X-ray imaging also apply to spectral and differential phase-contrast X-ray imaging. These imaging techniques enable a three-dimensional reconstruction of other object properties, such as basis material volume fraction or the electron density (see section 2.7.1 and 2.8). In the last decades, CT reconstruction was almost exclusively performed with analytical reconstruction algorithms such as filtered backprojection. Recently, statistical iterative reconstruction (SIR) algorithms which can be viewed as MAP estimators for the inverse problem of CT reconstruction have become an active field of research.

2.6.1 Filtered backprojection

Although other methods exist [98, 99], the filtered backprojection (FBP) algorithm is by far the most common analytical CT reconstruction method. For simplicity, the FBP algorithm will be derived in two-dimensional parallel beam geometry, similar to the discussion in reference [41]. A

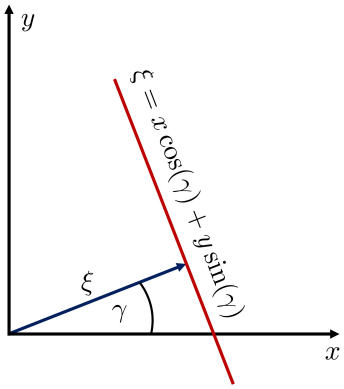

Figure 2.18: Graphical visualization of the coordinate systems ((x, y) and (γ, ξ)) as well as the line of projection $\xi = x\cos(\gamma) + y\sin(\gamma)$ for the FBP algorithm in parallel beam geometry. The red line represents all points that are projected onto $p_\gamma(\xi)$, where ξ is the distance from the rotation center and γ is the rotation angle. (This figure was inspired by a similar figure in reference [41].)

generalization to the three-dimensional case with parallel beam geometry is straightforward. By modifying the basic reconstruction principle for the parallel beam case, the FBP algorithm can be applied to fan beam [100, 101], cone beam [102, 103] and helical [104, 105, 106] acquisition geometries.

Mathematically, the measurement acquisition for a CT scan in parallel beam geometry can be described by a *Radon transform*. The Radon transform [107] maps a two-dimensional function $g(x, y)$ to its projections $p(\gamma, \xi) = p_\gamma(\xi)$. The function $g(x, y)$ represents the property of the object that is reconstructed (e.g. the attenuation coefficient), γ is the projection angle and ξ denotes the distance of the ray from the point of origin (see figure 2.18). The Randon transform is given by:

$$p_\gamma(\xi) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} g(x, y)\delta(x\cos(\gamma) + y\sin(\gamma) - \xi)dxdy, \tag{2.83}$$

where the argument of the δ-distribution parameterizes the line of projection ($x\cos(\gamma)+y\sin(\gamma) = \xi$, see figure 2.18). Inverting the Radon transform and calculating $g(x, y)$ from the *sinogram* $p(\gamma, \xi)$ is based on the Fourier slice theorem, which states: The one-dimensional Fourier transform $P_\gamma(\omega)$ of the projection $p_\gamma(\xi)$ is equivalent to a radial slice of the two-dimensional Fourier transform $G(u, v)$ of $g(x, y)$ taken at the angle γ:

$$P_\gamma(\omega) = G(\omega\cos(\gamma), \omega\sin(\gamma)). \tag{2.84}$$

The Fourier slice theorem and its connection to image reconstruction is illustrated in figure 2.19. We start the derivation of the FBP algorithm by expressing $g(x, y)$ as the inverse Fourier transform of $G(u, v)$:

$$g(x, y) = \frac{1}{2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} G(u, v)e^{2\pi i(xu+yv)}dudv. \tag{2.85}$$

By performing a coordinate transformation to polar coordinates, the Fourier slice theorem (eq. 2.84) can be combined with the Fourier representation of $g(x, y)$:

$$\begin{aligned} g(x, y) &= \frac{1}{2\pi}\int_0^{2\pi}\int_{-\infty}^{\infty} G(\omega\cos\gamma, \omega\sin\gamma)e^{2\pi i\omega(x\cos(\gamma)+y\sin(\gamma))}\omega d\omega d\gamma \\ &= \frac{1}{2\pi}\int_0^{2\pi}\int_{-\infty}^{\infty} P_\gamma(\omega)e^{2\pi i\omega(x\cos(\gamma)+y\sin(\gamma))}\omega d\omega d\gamma \\ &= \frac{1}{2\pi}\int_0^{\pi}\int_{-\infty}^{\infty} P_\gamma(\omega)e^{2\pi i\omega\xi}|\omega| d\omega d\gamma, \end{aligned} \tag{2.86}$$

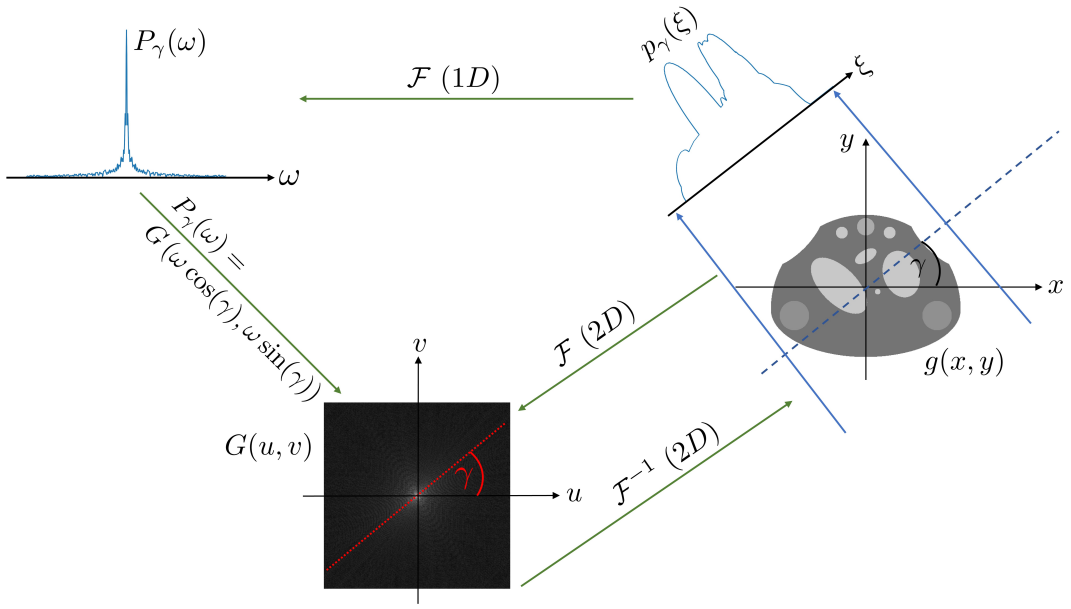

Figure 2.19: Visualization of the Fourier slice theorem and its connection to image reconstruction. The Fourier slice theorem allows to obtain the two-dimensional Fourier transform $G(u, v)$ of the image $g(x, y)$ from the projections $p_\gamma(\xi)$. (This figure was inspired by a similar figure in reference [108].)

where the symmetry properties of the Fourier transform and the definition of $\xi = x\cos(\gamma) + y\sin(\gamma)$ have been used to arrive at the last equality (see reference [41] pg. 181 for the derivation). The inner integral in the last line of eq. 2.86:

$$h_\gamma(\xi) = \int_{-\infty}^{\infty} P_\gamma(\omega) e^{2\pi i\omega\xi} |\omega| d\omega \tag{2.87}$$

can be interpreted as a high-pass filter that is implemented in Fourier space. Since a multiplication in Fourier space is equivalent to a convolution in real space, the filtered projections $h_\gamma(\xi)$ could also be obtained by convolving $p_\gamma(\xi)$ with $\mathcal{F}^{-1}(|\omega|)$. The outer integral in the last line of eq. 2.86 represents a backprojection of the filtered projections:

$$g(x, y) = \frac{1}{2\pi} \int_0^\pi h_\gamma(\xi) d\gamma. \tag{2.88}$$

For each angle γ, the backprojection smears the filtered projections $h_\gamma(\xi)$ over the image plane along the ray directions given by ξ. In summary, the FBP algorithms comprises two major steps:

- 1. high-pass filtering the projections $p_\gamma(\xi)$ by:
  a) multiplying with $|\omega|$ in frequency space
  b) convolving with $\mathcal{F}^{-1}(|\omega|)$ in real space
- 2. backprojecting the filtered projections $h_\gamma(\xi)$

Equation 2.86 cannot be applied directly to reconstruct attenuation images from a CT scan since only a discrete sampling of $p_\gamma(\xi)$ (both in the spatial and angular dimension) is measured. Moreover, $g(x, y)$ has to be discretized in order to perform numerical computations. A discretized version of the Randon transform (compare eq. 2.83) is given by:

$$-\ln\left(\frac{y_i}{b_i}\right) = l_i = \sum_j a_{ij}\mu_j, \tag{2.89}$$

where the variables γ and ξ have been summarized in the sinogram index i. The line integrals l_i of the attenuation coefficient are approximated by a forward projection of the image $\vec{\mu}$. The elements a_{ij} of the projection matrix $\boldsymbol{a}$ represents the contribution of voxel j to the line integral for sinogram index i and μ_j denotes the attenuation coefficient for voxel j. Calculating the measured line integrals as $l_i = y_i/b_i$ implies the assumption of a monochromatic X-ray source or an effective energy which describes the polychromatic source spectrum sufficiently well. In the discrete case, a sufficient angular sampling is necessary to avoid aliasing artifacts. A common rule of thumb is [109]:

$$N_{\text{proj}} = \frac{\pi}{2} N_{\text{px}}, \tag{2.90}$$

where N_{proj} is the number of angular projections and N_{px} is the number of pixels in one detector row.

2.6.2 Statistical iterative reconstruction

From a Bayesian perspective, statistical iterative reconstruction (SIR) can be viewed as the task of finding the image $\vec{\mu}$ that maximizes the probability of the measured data $\vec{y}$:

$$\vec{\mu}_{\text{r}} = \arg\max_{\vec{\mu}} P(\vec{\mu}|\vec{y}). \tag{2.91}$$

This optimization problem corresponds to a MAP estimator for the inverse problem of CT reconstruction. As explained in section 2.5.1, the optimization problem can be reformulated in terms of the penalized log-likelihood function $\phi(\vec{\mu})$:

$$\vec{\mu}_{\text{r}} = \arg\min_{\vec{\mu}} \phi(\vec{\mu}) = \arg\min_{\vec{\mu}} \left[-L(\vec{\mu}) + \lambda R(\vec{\mu})\right]. \tag{2.92}$$

SIR algorithms offers many advantages compared to analytical CT reconstruction methods. The possibility to incorporate an accurate noise model is one of the reasons for the improved image quality of SIR methods compared to FBP reconstruction. Moreover, the MAP estimation framework allows to include prior knowledge in the form of a regularization term, which suppresses noise and enhances the image quality compared to analytical methods. Analytical reconstruction methods that rely on the Fourier slice theorem impose severe restrictions on the acquisition geometry and the physical model of the measurement acquisition. SIR algorithms, on the contrary, can handle undersampling, missing data and provide the opportunity to include a sophisticated simulation of the CT physics. Minimizing the penalized log-likelihood function requires iterative optimization algorithms. The associated high computational complexity leads to long reconstruction times compared to FBP reconstruction. This is the most important disadvantage of SIR, especially for clinical applications. In the last years, the advent of highly efficient optimization algorithms as well as advances in computer hardware have mitigated this shortcoming.
SIR is a modular concept with many different components which often can be customized and optimized independently of each other:

- **discretization:** To reconstruct the image from the measurements, it is necessary to discretize the distribution of attenuation values $\mu(x, y, z)$. There are different choices for this discretization, leading to different reconstruction results.

- **regularization:** Physically implausible images are assigned with a low prior probability. This stabilizes the ill-conditioned problem of CT reconstruction.

- **forward model:** The forward model is based on a physical simulation of the CT scan and calculates the expected measurement values based on the current image estimate.

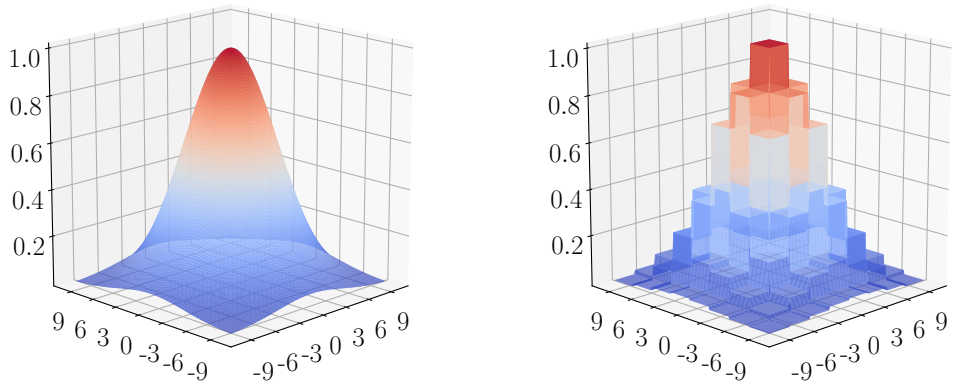

(a) (b)

Figure 2.20: Discrete representation of a two-dimensional continuous function (a) using a pixel basis (b).

- **data model:** The log-likelihood of the measurement data given the current image estimate is computed. This computation requires comparing the expected measurements $\hat{\vec{y}}$ (as determined by the forward model) with the actual measurement data $\vec{y}$ according a statistical model of the measurement acquisition process.
- **optimization algorithm:** Iterative optimization algorithms are necessary to find the minimum of the penalized log-likelihood function $\phi(\vec{\mu})$.

The modular design of SIR techniques allows to tailor the CT reconstruction to the requirements of a specific imaging task by adjusting (some of) the aforementioned components. The flexibility gained by this approach is one of the great strengths of SIR. In the following, the individual components of SIR will be explained in more detail.

Discretization The numerical evaluation and minimization of the penalized log-likelihood function requires a finite paramererization $\vec{\mu}$ of the spatial distribution of attenuation coefficients $\mu(x, y, z)$. A general approach for parameterizing $\mu(x, y, z)$ is to represent it in terms of a basis expansion [53]:

$$\mu(x, y, z) = \sum_{j=1}^{N} \mu_j \chi_j(x, y, z), \tag{2.93}$$

where $\chi_j(x, y, z)$ are the basis functions and μ_j are the coefficients of the basis expansion. The basis functions should have compact support in order to reduce the computational complexity of the forward model. [53]. The most common choice of $\chi_j(x, y, z)$, which was also used in this work, is the so-called voxel basis:

$$\chi_j(x, y, z) = \text{rect}\left(\frac{x - x_j}{\Delta}\right)\left(\frac{y - y_j}{\Delta}\right)\left(\frac{z - z_j}{\Delta}\right), \tag{2.94}$$

where $\text{rect}(\cdot)$ is the rectangular function, (x_j, y_j, z_j) is the center of the jth voxel and Δ is the pixel size. Figure 2.20 shows the effect of discretizing a two-dimensional function by using a pixel basis. Spherically symmetric Kaiser-Bessel functions ("blobs") have been investigated as an alternative to the voxel basis because of their mathematically appealing properties [110, 111, 112]. Other, less frequently used discretization choices include wavelets [113] and B-splines [114]. In the following, the voxel basis will be assumed since it is the most common choice and

was used throughout this work. The fact that the object is discretized into a finite number of basis functions inherently leads to discrepancies between the measured projection data and the forward model [115, 116]. The mismatch between the real measurement values and the simulated projections may cause artifacts such as edge overshoots and aliasing patterns [117]. Choosing a larger number of basis functions (i.e. smaller voxels) reduces the severity of these artifacts.

Forward model An important part of the forward model is the forward projection, which is a discretized version of the Radon transform (compare eq. 2.89):

$$\vec{l} = \boldsymbol{a}\vec{\mu}; \quad l_i = \sum_j a_{ij}\mu_j, \tag{2.95}$$

where the matrix elements a_{ij} represent an approximation of the (average) intersection length for sinogram index i and voxel j. Even though $\boldsymbol{a}$ is sparse, pre-computing the matrix elements $\boldsymbol{a}$ and executing a matrix-vector multiplication to calculate the forward projection is unfeasible because of the large size of $\boldsymbol{a}$. For typical CT applications, the memory requirements to store $\boldsymbol{a}$ (in uncompressed form) would be in the exabyte range [118]. Consequently, the matrix elements a_{ij} have to be computed on the fly. The backprojection, which is the transpose of the forward projection is needed to calculate the gradient of the penalized log-likelihood function. In may cases, the gradient $\vec{g}$ of the penalized log-likelihood function can be expressed as a backprojection of a gradient $\vec{g}_\mathrm{p}$ in projection space:

$$\vec{g} = \boldsymbol{a}^T\vec{g}_\mathrm{p}; \quad g_j = \sum_i a_{ij}g_{\mathrm{p},i}. \tag{2.96}$$

Since the gradient is required for efficiently minimizing the penalized log-likelihood function, the backprojection also plays an important role in SIR frameworks. The forward and backprojection operations are computationally expensive and thus they are often one of the most important bottlenecks in SIR methods. Various approaches for calculating forward and backprojections with different tradeoffs between computational complexity and accuracy of the projection operations have been developed.

We will discuss pixel-driven and ray-driven projectors in more detail because the in-house tomographic projector [119, 118] that was used in this work relies on these concepts. Both of these projector types can be used for arbitrary 3D cone-beam geometries. Moreover, they are particularly suited for highly parallelized implementations on graphics processing units (GPUs) [119, 118, 120, 121]. The recent progress in GPU technology is one of the driving forces behind the increasing popularity of SIR techniques for medical imaging and non-destructive testing [122]. Ray-driven projectors work by tracing rays through the image and calculating the line integrals as a weighted sum of the image voxels that are close to the corresponding rays [123]. For illustration purposes, figure 2.21 (a) shows the working principle of a ray-driven projector in 2D parallel beam geometry. For each sinogram index i, a ray connects the source and the center of the corresponding detector pixel. The line integrals are obtained by calculating the intersection length with each row of voxels and then interpolating the attenuation coefficient from the two nearest voxels within that row. This approach is generalized to 3D by calculating the slab intersection length combined with bi-linear interpolation between the four nearest voxels within that slab [124, 125]. Ray-driven projectors are typically used for forward projections because the ray-driven backprojection can introduce Moiré artifacts [126]. Moreover, the forward projection is easier to parallelize because of the highly non-sequential memory access of the ray-driven backprojection [123]. Pixel-driven projectors [127, 128] (see figure 2.21 (b)) consider lines that connect the source and the center of each image pixel (or voxel in 3D). The location where the line intersects with the detector surface determines the relative weight of the adjacent detector pixels for the subsequent linear interpolation. Since this method introduces high-frequency artifacts and is difficult to parallelize for the forward projection [123], it is most

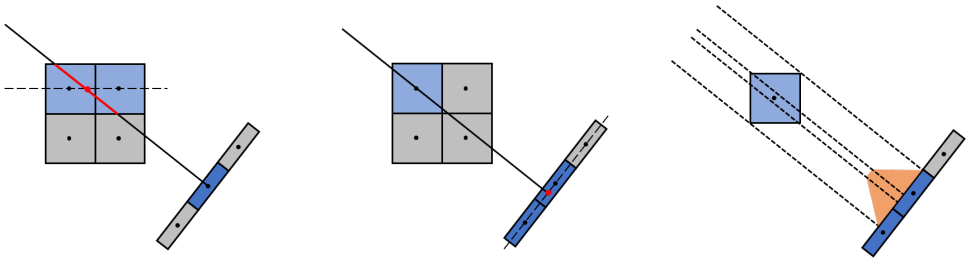

(a) (b) (c)

Figure 2.21: Working principle of ray-driven (a), pixel-driven (b) and separable footprint projectors (c). Ray-driven projectors interpolate between different voxels in image space whereas pixel-driven projectors interpolate between adjacent detector pixels. Footprint projectors achieve the highest accuracy by considering the area of overlap between the image voxels and the detector pixels.

commonly used for implementing the backprojection. The tomographic projector that was used in this work thus combines a ray-driven forward projection with a pixel-driven backprojection. This strategy allows an efficient parallelization and optimization of both operations on GPUs, however the resulting projector pair is *unmatched*. Due to the different interpolation strategies, the backprojection is not consistent with the forward projection. More precisely, the backprojection matrix (which is calculated on the fly) is not the transpose of the forward projection matrix. The implications of using a unmatched projector pair concerning convergence and image quality are still under investigation [129, 130, 131]. Ray-driven forward- and pixel-driven backprojectors sacrifice some accuracy for computational efficiency and speed of computation. More accurate methods such as distance driven [123] or separable footprint [132] (see figure 2.21 (c)) projectors consider the overlap between the image voxels and the detector pixels to calculate approximations of the average intersection lengths. These methods implement matched forward and backprojector pairs. In the last years, parallelized GPU implementations have reduced the computational time for these methods [133, 134, 135, 136]. In its simplest form, the forward model is given by eq. 2.95. Combining the forward projector with the Lambert-Beer law provides a simple model for the expected number of photon counts:

$$\hat{y}_i = b_i \exp\left(-\sum_j a_{ij}\mu_j\right), \tag{2.97}$$

where $\hat{y}_i$ is the expected number of photon counts for sinogram index i and b_i is the corresponding reference intensity (without the sample in the beam path). Directly modeling the expected intensities avoids the statistical bias that is caused by calculating the measured line integrals (see section 2.5.1). Within the limits of computational feasibility, the SIR framework can be combined with more sophisticated forward models that simulate various aspects of the measurement aquisition process. Examples include modeling the polychromatic source spectrum [137, 138, 139], the effect of scattered radiation reaching the detector [140, 53], an extended X-ray source spot [141, 142, 143], detector crosstalk [143, 144, 145] as well as object motion [146, 147] .

Data model In order to calculate the data term $-L(\vec{\mu})$ of the penalized log-likelihood function $\phi(\vec{\mu})$, the probability $P(\vec{y}|\vec{\mu}) = P\left(\vec{y}|\hat{\vec{y}}(\vec{\mu})\right)$ of the measured data $\vec{y}$ given the current image estimate $\vec{\mu}$ needs to be calculated. Calculating $P(\vec{y}|\vec{\mu})$ involves comparing the measured data to the expected intensities predicted by the forward model $\hat{\vec{y}}$ according to a statistical model of the measured signals. As derived in section 2.4.1, the number of detected X-ray quanta follows

a Poisson distribution for an ideal detector. Depending on the detector type, the statistical distribution of the signals measured with a real detector differs from this idealized approximation. For an energy-integrating detector, the energy-dependent weighting leads to a compound Poisson process. By including electronic noise of the detector, $P(y_i|\hat{y}_i)$ can be expressed as [115, 148]:

$$P(y_i|\hat{y}_i) \sim \text{compound Poisson}\left(\hat{y}_i, \Phi(E)\right) + \text{Gaussian}\left(M_\text{e}, \sigma_\text{e}^2\right), \tag{2.98}$$

where $\Phi(E)$ is the effective photon energy spectrum and M_e and σ_e^2 denote the mean and variance of the electronic noise, respectively. Although still an approximation, Whiting et al. [67] have shown that this statistical model for energy-integrating detectors agrees well with experimental data. Strictly speaking, this model is only valid for a detector with a small point spread function since charge-sharing or scintillator blurring effects introduce noise correlations between neighboring pixels. In practice, the noise model of eq. 2.98 is often simplified to a Poisson distribution:

$$P\left(y_i|\hat{y}_i(\vec{\mu})\right) = \frac{\hat{y}_i(\vec{\mu})^{y_i}}{y_i!} e^{-\hat{y}_i(\vec{\mu})}. \tag{2.99}$$

Although the compound Poisson noise model has been used in a SIR framework [149], the increased computational complexity normally outweighs the incremental benefit of an exact noise model [124]. Since photon-counting detectors intrinsically suppress electronic noise and weigh each photon equally, the Poisson model is considered to be a good approximation in this case [78, 150]. Assuming that all measurements are statistically independent, the likelihood of the measured data for the Poisson model is expressed as:

$$P(\vec{y}|\vec{\mu}) = \prod_{i=1}^{M} \frac{\hat{y}_i(\vec{\mu})^{y_i}}{y_i!} e^{-\hat{y}_i(\vec{\mu})}. \tag{2.100}$$

Consequently, the negative log-likelihood is given by:

$$-L(\vec{\mu}) = \sum_{i=1}^{M} \hat{y}_i(\vec{\mu}) - y_i \ln(\hat{y}_i(\vec{\mu})), \tag{2.101}$$

where terms independent of $\vec{\mu}$ have been dropped. For a sufficiently large number of photons ($\hat{y}_i \geq 10$), the Poisson distribution is well approximated by a Gaussian distribution:

$$P\left(y_i|\hat{y}_i(\vec{\mu})\right) = \frac{1}{\sqrt{2\pi\sigma^2(\hat{y}_i)}} \exp\left(-\frac{(y_i - \hat{y}_i)^2}{2\sigma^2(\hat{y}_i)}\right) \tag{2.102}$$

If electronic noise can be neglected, the variance $\sigma^2(\hat{y}_i)$ is equal to the expected photon counts $\hat{y}_i$, which in turn are often approximated by the measured photon counts y_i:

$$\sigma^2(\hat{y}_i) = \hat{y}_i \approx y_i. \tag{2.103}$$

Another common approach is to consider the statistical distribution of the line integrals $l_i = y_i/b_i$. Using error propagation and neglecting the variance of the reference measurements b_i, one can derive $P(l_i|\hat{l}_i)$ from eq. 2.102:

$$P\left(l_i|\hat{l}_i(\vec{\mu})\right) = \frac{1}{\sqrt{2\pi\sigma^2(l_i)}} \exp\left(-\frac{(l_i - \hat{l}_i)^2}{2\sigma^2(l_i)}\right), \tag{2.104}$$

where $\hat{l}_i = \hat{y}_i/b_i$. By propagating the variance through the logarithm, one can show that:

$$\sigma^2(l_i) \approx \frac{1}{\hat{y}_i} \approx \frac{1}{y_i}. \tag{2.105}$$

The key idea of the second approximation in eq. 2.105 is to simplify $-L(\vec{\mu})$ to a quadratic function which is much easier to minimize:

$$-L(\vec{\mu}) = \sum_{i=1}^{M} w_i \Big(\hat{l}_i - l_i\Big)^2, \tag{2.106}$$

where w_i is a precomputed weight. The standard choice $w_i = y_i$ follows from eq. 2.105. Thibault et al. [151] have proposed the following modified weights which include the electronic noise σ_e^2:

$$w_i = \frac{y_i^2}{y_i + \sigma_e^2}. \tag{2.107}$$

As explained before, working with the line integrals l_i instead of the intensities y_i leads to statistical bias for low photon statistics. In the case of photon-counting detectors, zero-count measurements can provide valuable information [70]. Due to the logarithmic conversion, this information is difficult to include in line integral based data models.

Regularization In terms of Bayesian statistics, regularization for attenuation-based imaging can be understood as the process of assigning prior probabilities to different attenuation maps $\vec{\mu}$ in order to stabilize the ill-conditioned problem of CT reconstruction. A low prior probability is assigned to images that fit the measurement data well but are physically implausible. Researchers have developed a plethora of different regularization strategies based on certain prior assumptions about the images (for a comprehensive overview see reference [152]). Choosing the right regularizer for a particular imaging task is still an active field of research. This section will mainly focus on the most commonly used approach that is based on the assumption that the attenuation map $\vec{\mu}$ is locally smooth, i.e. neighboring voxels tend to have similar values. Noise suppression is thus achieved by assigning a low prior probability to attenuation maps $\vec{\mu}$ with high local variance. The regularization term for this family of nearest-neighbor based regularizers can be written in the following general form [53]:

$$R(\vec{\mu}) = \sum_{j=1}^{N} \sum_{k=1}^{N_p} w_{jk} \psi(|\mu_j - \mu_k|), \tag{2.108}$$

where N is the number of voxels and N_p is the number of nearest neighbors that are taken into account. The potential function $\psi(\cdot)$ penalizes the absolute difference between neighboring voxels and w_{jk} denotes a weight given to this penalty. In most applications, the weights are simply chosen inversely proportional to the distance between the two voxels under consideration. However, regularizers that aim at uniform spatial resolution [153, 154] or noise [155, 156] modify the local (i.e. as a function of the voxel index j) and the directional (i.e. as function of the neighbor index k) regularization strength. The overall regularization strength (controlled by the regularization parameter λ) determines the tradeoff between noise suppression and regularization-induced bias. The roughness penalty of nearest-neighbor based regularizers can cause blurring and distortion of edges and small structures. Over-regularization can thus decrease the spatial resolution and interfere with the detection of small or low-contrast features.
The choice of the potential function $\psi(\cdot)$ also strongly influences the properties of the reconstructed image $\vec{\mu}$. Possibly the most natural choice is a quadratic penalty:

$$\psi(\Delta) = \Delta^2, \tag{2.109}$$

where Δ denotes the difference between two neighboring voxel values. While the quadratic penalty has good noise suppression properties, it tends to oversmooth boundaries between different materials and small structures. The Huber potential [157] mitigates this effect by linearly penalizing larger value differences:

$$\psi(\Delta) = \begin{cases} \Delta^2, & \text{if } |\Delta| \le \gamma \\ 2\gamma|\Delta| - \gamma^2, & \text{if } |\Delta| > \gamma \end{cases}. \tag{2.110}$$

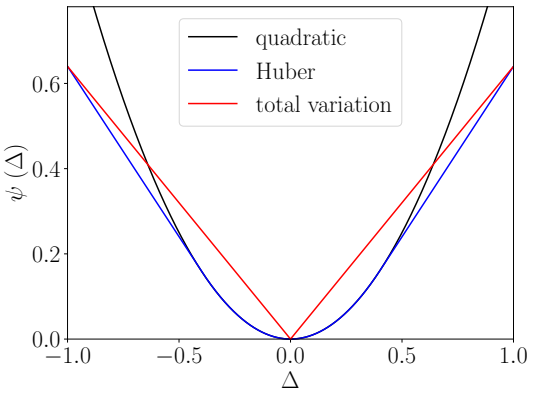

Figure 2.22: Comparison between three commonly used potential functions (quadratic, Huber and total variation) for nearest-neighbor based regularizers.

Up to the threshold γ, the Huber potential coincides with the quadratic potential but a linear continuous continuation is used for $|\Delta| > \gamma$. If the value of γ is chosen correctly, noise is penalized quadratically while the penalty for edges only increases linearly with the edge height. Anisotropic total-variation (TV) regularization is another popular regularization strategy [158, 159, 160] that uses a linear penalty function:

$$\psi(\Delta) = |\Delta|. \tag{2.111}$$

TV regularizers have good edge preserving properties, however they potentially introduce staircase artifacts in low dose CT reconstructions due to the restoration of artifical edges that were introduced by noise [152, 161]. Figure 2.22 shows a comparison between the convex potential functions for quadratic, Huber and TV regularizers. Convex potential functions are normally preferred because they guarantee (in combination with a convex data term) that the optimization algorithm converges to a unique global minimum. Additionally, specialized optimization algorithms that rely on convex optimization theory can be used.
Another important class of regularization strategies is based on compressed sensing theory. Compressed sensing is a signal processing technique that allows to recover signals from undersampled measurements by assuming sparsity of the signal in a certain domain. TV regularization can be viewed as a compressed sensing method because the TV regularizer enforces sparsity of the gradient of the image $\vec{\mu}$. Dictionary-based regularizers assume that small image patches can be sparsely represented by a linear combination of so-called dictionary atoms. The dictionary atoms are learned from high-quality training images and contain typical structures and image features that frequently occur in these images. In terms of Bayesian statistics, dictionary-based regularization corresponds to the prior assumption that small patches of the attenuation image can be sparsely represented in the dictionary basis (which consists of the dictionary atoms). Dictionary-based image denoising and regularization will be explained in more detail in chapter 4 and 5, respectively.
Besides nearest-neighbor and compressed sensing based regularizers, other regularization techniques have been developed. Example include regularizers based on nonlinear neighborhood filters [162, 163], prior images [164, 165, 166] (e.g. from an earlier CT scan of the same patient/object) and wavelets [167].

Optimization algorithm The optimization algorithm is an essential part of SIR techniques since closed-form solutions for minimizing the penalized log-likelihood function $\phi(\vec{\mu})$ do not exist in most cases. A non-iterative approach to minimizing could be based on solving the following

system of equations:

$$\frac{\partial \phi(\vec{\mu})}{\partial \mu_j} = 0 \quad \forall j = 1, ..., N, \tag{2.112}$$

However, even if eq. 2.112 represents a system of linear equations, direct inversion is impractical because of the large size of the system matrix. Efficient iterative optimization algorithms are therefore needed to minimize $\phi(\vec{\mu})$ reliably and quickly. Due to the larger number of optimization variables, iterative minimization algorithms for SIR are almost exclusively gradient-based. Consequently, convergence to a local optimum cannot be ruled out if $\phi(\vec{\mu})$ has several local optima. In this case, the initialization of the algorithm (initial guess) decides if the global optimum is reached. It is therefore desirable that $\phi(\vec{\mu})$ is convex, which implies that a unique minimum exists. In the last decades, a large variety of optimization algorithms with different characteristics have been developed and used for SIR [53]. The spectrum ranges from general purpose algorithms for large-scale convex [168] and non-convex [169, 170] optimization problems to specialized solvers that exploit the specific structure of the penalized log-likelihood function. The compatibility with different modular components of the SIR framework (e.g. different forward models or regularizers) is the main advantage of general purpose solvers. However, they typically converge much slower than specialized solvers that have been tailored to a particular SIR model.
In this section, we will focus on two optimization algorithms that have been used in this work. The nonlinear conjugate gradient (NLCG) algorithm is a general purpose solver that was used for SIR of spectral phase-contrast data (see chapter 9) as well as a variety of other MAP estimation problems in projection space. The ordered subsets separable quadratic surrogate (OS-SQS) algorithm is a highly efficient optimization algorithm that was originally developed for attenuation-based SIR. In this work, it was adapted to one-step SIR for spectral CT (see chapter 6).

Nonlinear Conjugated Gradient (NLCG) The NLCG algorithm is based on the conjugate gradient (CG) method which is used to solve large systems of linear equations:

$$\boldsymbol{A}x = b, \quad x, b \in \mathbb{R}^N, \quad \boldsymbol{A} \in \mathbb{R}^{N \times N}. \tag{2.113}$$

Solving this system of equations is equivalent to minimizing the following quadratic function of N variables:

$$f(x) = \frac{1}{2} x^T \boldsymbol{A} x - b^T x + c, \quad c \in \mathbb{R}. \tag{2.114}$$

Based on the tutorial by Shewchuk [171], we first introduce the CG method before giving a brief overview of the modifications that enable the generalization to the NLCG algorithm, which can be applied to non-quadratic functions. More details about the NLCG algorithm can be found in the optimization literature [172, 173, 174, 175].
The CG algorithm performs repeated line searches along certain search directions $d_{(n)} \in \mathbb{R}^N$ where the index n represents the n-th iteration of the algorithm. This means that the new guess $x_{(n+1)}$ for the optimization variables lies on the line described by:

$$x_{(n+1)} = x_{(n)} + \alpha_{(n)} d_{(n)}, \ \alpha_{(n)} \in \mathbb{R}, \tag{2.115}$$

where $x_{(n)}$ is the current guess. Minimizing eq. (2.115) with respect to $\alpha_{(n)}$ yields:

$$\alpha^{(n)} = \frac{r_{(n)}^T r_{(n)}}{d_{(n)}^T \boldsymbol{A} d_{(n)}}, \tag{2.116}$$

where $r_{(n)} = -\nabla f(x_{(n)}) = b - A x_{(n)}$ is the residual. The key difference to the gradient descent algorithm is the selection of subsequent search directions. Whereas for gradient descent the search direction coincides with the gradient at the current position:

$$d_{(n)} = -\nabla f\left(x_{(n)}\right), \tag{2.117}$$

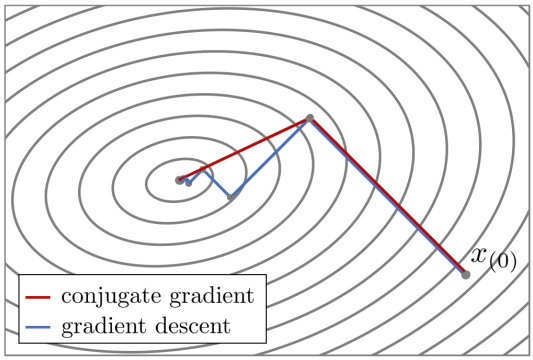

Figure 2.23: Comparsion between conjugate gradient and gradient descent optimization for a two-dimensional problem. Whereas the gradient descent algorithm takes several steps along similar search directions, the conjugate gradrient algorithm converges in two iterations (figure modified from reference [176]).

the search directions d are chosen to be conjugate with respect to $\boldsymbol{A}$ for the CG method:

$$d_{(q)}^T \boldsymbol{A} d_{(j)} = 0 \quad \forall\ q \neq j. \tag{2.118}$$

Depending on the shape of $f(x)$, new search directions are often close to previous ones for the gradient descent algorithm. The CG method avoids this by choosing conjugate search directions, which greatly accelerates convergence. In contrary to steepest descent, the CG algorithm theoretically converges in N iterations (with N being the number of optimization variables). An example of the benefit of conjugate search directions for a two-dimensional optimization problem is shown in figure 2.23. Conjugate search directions are efficiently constructed by Gram-Schmidt conjugation:

$$\begin{aligned} &\text{Initialize: } d_{(0)} = r_{(0)} = -\nabla f\left(x_{(0)}\right), \text{ then construct recursively:} \\ &\beta_{(n+1)} = \frac{r_{(n+1)}^T r_{(n+1)}}{r_{(n)}^T r_{(n)}} \\ &d_{(n+1)} = r_{(n+1)} + \beta_{(n+1)} d_{(n)}. \end{aligned} \tag{2.119}$$

Algorithm 1 presents the complete CG method, which combines Gram-Schmidt conjugation with a linesearch along the new search direction (compare eq. 2.115 and 2.116).

Initialize $d^{(0)} = r^{(0)} = -\nabla f(x_{(0)})$

for *i=0,...,N* **do**

$\alpha_{(n)} = \frac{r_{(n)}^T r_{(n)}}{d_{(n)}^T \boldsymbol{A} d_{(n)}}$

$x_{(n+1)} = x_{(n)} + \alpha_{(n)} d_{(n)}$

$r_{(n+1)} = -\nabla f(x_{(n+1)}) = r_{(n)} - \alpha_{(n)} \boldsymbol{A} d_{(n)}$

$\beta_{(n+1)} = \frac{r_{(n+1)}^T r_{(n+1)}}{r_{(n)}^T r_{(n)}}$

$d_{(n+1)} = r_{(n+1)} + \beta_{(n+1)} d_{(n)}$

end

Algorithm 1: Outline of the CG method.

A detailed analysis [171] shows that the convergence properties of the CG algorithm are mainly determined by the eigenvalues of the matrix $\boldsymbol{A}$. An upper bound for the error after n iterations

is given by:

$$\frac{f(x_{(n)}) - f(x_{(N)})}{f(x_{(0)}) - f(x_{(N)})} \leq 4\left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^{2n}, \tag{2.120}$$

where $x^{(N)}$ indicates the solution of eq. (2.114) because the algorithm converges after N iterations. The spectral condition number κ is defined as the ratio between the largest and smallest eigenvalue of $\boldsymbol{A}$: $\kappa = \lambda_{\max}/\lambda_{\min}$. From eq. (2.120), it follows that a smaller spectral condition number leads to faster convergence. One can also show that convergence is accelerated if the eigenvalues are not evenly distributed between $\lambda_{\min}$ and $\lambda_{\max}$ but clustered together. Performing N CG iterations to guarantee convergence is computationally infeasible for large-scale systems. Moreover, floating point errors prevent an exact convergence for numerical calculations. In practice, the CG algorithms is thus terminated once the solution is approximated with sufficient accuracy.

Preconditioning The goal of preconditioning is to induce a coordinate transformation that decreases the spectral condition number of the matrix $\boldsymbol{A}$ or leads to a better clustering of the eigenvalues and thus accelerates convergence [177]. Suppose that the system matrix $\boldsymbol{A}$ can be approximated by a matrix $\boldsymbol{M}$ which is easier to invert than $\boldsymbol{A}$ and $\boldsymbol{M}$ can be written as $\boldsymbol{M} = \boldsymbol{E}\boldsymbol{E}^T$, $\boldsymbol{E} \in \mathbb{R}^{N\times N}$. The preconditioned CG method (see algorithm 2) essentially minimizes $f(\bar{x})$ after the coordinate transformation:

$$\bar{x} = \boldsymbol{E}^T x \tag{2.121}$$

By performing a few variable substitutions, the transformation matrix $\boldsymbol{E}^T$ is eliminated entirely from the algorithm (see reference [171] for the derivation) and the coordinate transform does not have to be carried out explicitly.

Initialize $r_{(0)} = -\nabla f(x_{(0)}), d_{(0)} = \boldsymbol{M}^{-1} r_{(0)}$

for *i=0,...,N* **do**

$$\alpha_{(n)} = \frac{r_{(n)}^T \boldsymbol{M}^{-1} r_{(n)}}{d_{(n)}^T \boldsymbol{A} d_{(n)}}$$
$$x_{(n+1)} = x_{(n)} + \alpha_{(n)} d_{(n)}$$
$$r_{(n+1)} = -\nabla f(x_{(n+1)}) = r_{(n)} - \alpha_{(n)} \boldsymbol{A} d_{(n)}$$
$$\beta_{(n+1)} = \frac{r_{(n+1)}^T \boldsymbol{M}^{-1} r_{(n+1)}}{r_{(n)}^T \boldsymbol{M}^{-1} r_{(n)}}$$
$$d_{(n+1)} = \boldsymbol{M}^{-1} r_{(n+1)} + \beta_{(n+1)} d_{(n)}$$

end

Algorithm 2: Outline of the preconditioned CG method.

The convergence properties of the preconditioned CG algorithm are determined by the spectral condition number and the clustering of eigenvalues of the matrix $\boldsymbol{M}^{-1}\boldsymbol{A}$ instead of $\boldsymbol{A}$. In general, there is a tradeoff between the computational complexity of determining $\boldsymbol{M}^{-1}$ and the convergence acceleration that the preconditioner provides. The ideal preconditioner $\boldsymbol{M} = \boldsymbol{A}$ would lead to convergence within one iteration, however determining $\boldsymbol{M}^{-1}$ is as difficult as solving the original system of equations (or equivalently minimizing $f(x)$). The diagonal preconditioner $\boldsymbol{M} = \text{diag}\ (A_{11}, ..., A_{NN})$ scales $f(x)$ along the coordinate axes. It is trivial to invert:

$$M_{jj}^{-1} = \frac{1}{A_{jj}},\ j = 1, ..., N\ , \tag{2.122}$$

but large convergence improvements can only be expected if the matrix $\boldsymbol{A}$ is dominated by its diagonal entries.

In the general case of non-quadratic functions, a Taylor expansion can be used to calculate a local quadratic approximation of $f(x)$. This strategy allows to transfer the basic principles of the CG method to the NLCG algorithm. In the NLCG algorithm, the Hessian matrix $\boldsymbol{H}$ evaluated at the current guess $x_{(n)}$ plays the role of the matrix $\boldsymbol{A}$ in the CG algorithm:

$$H_{ij}(x) = \frac{\partial^2 f(x)}{\partial x_i \partial x_j}. \tag{2.123}$$

Since the Hessian matrix is a function of the optimization variables x, the meaning of "conjugate search directions" (compare eq. 2.118) changes from one iteration to the next. The more $f(x)$ deviates from a quadratic function, the more quickly the search directions lose conjugacy. Besides a periodic restart, several other empirical modifications of the Gram-Schmidt process for constructing the search directions (see eq. 2.119) are used in the NLCG algorithm [178, 179, 170]. Furthermore, an exact linesearch (see eq. 2.116) is difficult to perform due to the non-quadratic nature of $f(x)$. Instead of finding the exact minimum along the current search direction, one settles for a stepsize α_n that fulfills the Wolfe conditions [180, 181]:

$$\begin{aligned} &(1):\ f\left(x_{(n)} + \alpha_{(n)} d_{(n)}\right) \leq f\left(x_{(n)}\right) + c_1 \alpha_{(n)} d_{(n)}^T \nabla f(x_{(n)}) \\ &(2):\ d_{(n)}^T \nabla f\left(x_{(n)} + \alpha_{(n)} d_{(n)}\right) \geq c_2 d_{(n)}^T \nabla f(x_{(n)}). \end{aligned} \tag{2.124}$$

A common empirical choice for the value of the constants is $c_1 = 0.1$ and $c_2 = 10^{-4}$ [172]. The Wolfe conditions ensure that the optimizer makes sufficient progress towards the minimum in each iteration.
Preconditioning can also be combined with the NLCG algorithm to accelerate convergence. An important difference to the CG method is that the Hessian $\boldsymbol{H}$ changes during optimization. Consequently, the preconditioner might become less effective if the curvature of $f(x)$ varies strongly. In many practical cases, however, the initial guess is already sufficiently close to the minimum so that the Hessian only varies slightly. In the case of SIR for attenuation-based CT for example, a sufficient initial guess can often be obtained by FBP reconstruction.

Ordered subset separable quadratic surrogate (OSSQS) The OSSQS algorithm is a specialized solver for CT reconstruction. It exploits the structure of the penalized log-likelihood function for standard attenuation-based SIR and is thus less flexible than general purpose solvers. More specifically, the OSSQS algorithm requires a nearest-neighbor based regularization term $R(\vec{\mu})$ with a convex potential function $\psi(\Delta)$ and a log-likelihood function of the form:

$$-L(\vec{\mu}) = \sum_{i=1}^{M} h_i\left(\hat{l}_i\right), \tag{2.125}$$

i.e. the log-likelihood function is separable with respect to the sinogram indices and for each sinogram index i, the log-likelihood can be expressed as a function of the forward model for the line integrals $\hat{l}_i$. The standard data models (compare eq. 2.101 and 2.106) fulfill the requirement of eq. 2.125. For example, in the case of a Gaussian data model, $h_i(\hat{l}_i) = w_i(\hat{l}_i - l_i)^2$. The key idea of the OSSQS algorithm is to replace the penalized log-likelihood function $\phi(\vec{\mu}) = -L(\vec{\mu}) + \lambda R(\vec{\mu})$ at each iteration step by a surrogate function $\varphi(\vec{\mu}; \vec{\mu}^{(n)})$ which is easier to minimize. The vector $\vec{\mu}^{(n)}$ denotes the estimate of the attenuation coefficients at the n-th iteration. Minimizing the surrogate functions monotonically decreases $\phi(\vec{\mu})$ as long as the following condition for the surrogate functions are fulfilled [182]:

$$\begin{aligned} &1.\quad \varphi\left(\vec{\mu}^{(n)}; \vec{\mu}^{(n)}\right) = \phi\left(\vec{\mu}^{(n)}\right) \\ &2.\quad \frac{\partial \varphi}{\partial \mu_j}\left(\vec{\mu}^{(n)}; \vec{\mu}^{(n)}\right) = \frac{\partial \phi}{\mu_j}\left(\vec{\mu}^{(n)}\right) \quad \forall j = 1, ..., N \\ &3.\quad \varphi\left(\vec{\mu}; \vec{\mu}^{(n)}\right) \geq \phi\left(\vec{\mu}\right). \end{aligned} \tag{2.126}$$

The OSSQS algorithm replaces the penalized log-likelihood function by a separable surrogate function at each iteration step. This means that the surrogate function (at each iteration) can be written as:

$$\varphi\left(\vec{\mu};\vec{\mu}^{(n)}\right)=\sum_{i=1}^{N} f_i(\mu_i). \tag{2.127}$$

If the f_i are also quadratic, the minimization step is reduces to finding the minimum of N one-dimensional quadratic functions, which can be performed analytically and in parallel. Figure 2.24 illustrates the concept of minimizing $\phi(\vec{\mu})$ by the separable surrogate function principle for a one-dimensional objective function.

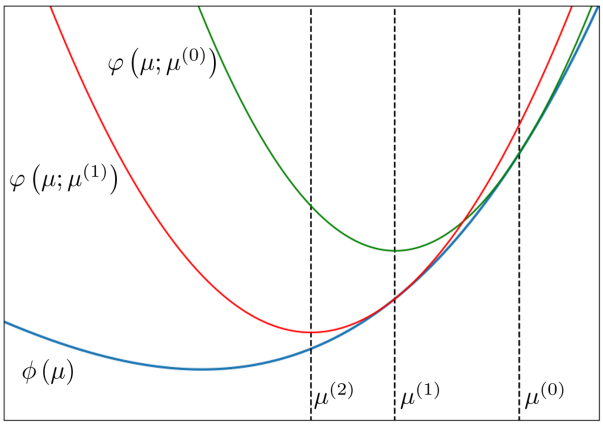

Figure 2.24: Illustration of the quadratic surrogate principle in 1D. Starting with an initial guess $\mu^{(0)}$, at each iteration step a quadratic surrogate function for the more complicated penalized log-likelihood function $\phi(\mu)$ is constructed (at the position of the current guess) and then analytically minimized.

In order to find a separable surrogate function for $\phi(\vec{\mu})$, surrogate functions for the likelihood and the penalty term are constructed separately. In case of the likelihood term, the separable quadratic surrogate function is constructed in two steps. First, a quadratic (non-separable) surrogate function for the likelihood term is created. This surrogate function is then replaced by a second surrogate function which is quadratic and separable. A quadratic surrogate function for the likelihood term which fulfills the surrogate conditions (eq. 2.126) is given by:

$$Q_1\left(\vec{\mu};\vec{\mu}^{(n)}\right)=\sum_{i=1}^{M} q_i\left(\hat{l}_i,\hat{l}_i^{(n)}\right)\geq -L(\mu)=\sum_{i=1}^{M} h_i(\hat{l}_i), \tag{2.128}$$

with

$$q_i\left(\hat{l}_i;\hat{l}_i^{(n)}\right)=h_i\left(\hat{l}_i^{(n)}\right)+\frac{\partial h_i}{\partial \hat{l}_i}\left(\hat{l}_i^{(n)}\right)\left(\hat{l}_i-\hat{l}_i^{(n)}\right)+\frac{1}{2}c_i\left(\hat{l}_i^{(n)}\right)\left(\hat{l}_i-\hat{l}_i^{(n)}\right)^2, \tag{2.129}$$

where $\hat{l}_i^{(n)}=\sum_j a_{ij}\mu_j^{(n)}$ is the estimate of the i-th line integral at the n-th iteration. and $c_i(\hat{l}_i^{(n)})$ is the curvature of the parabola q_i. The curvature $c_i(\hat{l}_i^{(n)})$ of the quadratic function $q_i\left(\hat{l}_i;\hat{l}_i^{(n)}\right)$ should be chosen as small as possible (within the constraints of the third condition in eq. 2.126) to allow for larger update steps (compare figure 2.24) and thus accelerate convergence of the algorithm. It can be shown [182] that the optimum curvature is given by:

$$c_i(\hat{l}_i^{(n)})=\begin{cases} 2\dfrac{h_i(0)-h_i\left(\hat{l}_i^{(n)}\right)+\frac{\partial}{\partial \hat{l}_i}h_i\left(\hat{l}_i^{(n)}\right)\,\hat{l}_i^{(n)}}{\left(\hat{l}_i^{(n)}\right)^2} & ,\hat{l}_i^{(n)}>0 \\ \dfrac{\partial^2}{\partial \hat{l}_i^2}h_i(0), \quad \hat{l}_i^{(n)}=0 & \end{cases}. \tag{2.130}$$

However, the optimum curvature is only valid if a non-negativity constraint on $\vec{\mu}$ is enforced. Finding the second, separable surrogate function $Q_2\left(\vec{\mu};\vec{\mu}^{(n)}\right)$ to the quadratic surrogate function $Q_1\left(\vec{\mu};\vec{\mu}^{(n)}\right)$ is based on rewriting the sum:

$$\hat{l}_i = \sum_{j=1}^{N} a_{ij}\mu_j = \sum_{j=1}^{N} \alpha_{ij}\left(\frac{a_{ij}}{\alpha_{ij}}\left(\left(\mu_j - \mu_j^{(n)}\right) + \sum_{r=1}^{N} a_{ir}\mu_r^{(n)}\right)\right), \tag{2.131}$$

where

$$\sum_{j=1}^{N} \alpha_{ij} = 1, \quad \forall\ i \text{ and } \alpha_{ij} \geq 0\ \forall i,j. \tag{2.132}$$

Applying Jensen's inequality to the convex function $q_i\left(\hat{l}_i;\hat{l}_i^{(n)}\right)$ yields:

$$q_i\left(\hat{l}_i;\hat{l}_i^{(n)}\right) \leq \sum_{j=1}^{N} \alpha_{ij} q_i\left(\frac{a_{ij}}{\alpha_{ij}}\left(\mu_j - \mu_j^{(n)}\right) + \hat{l}_i^{(n)};\hat{l}_i^{(n)}\right), \quad \hat{l}_i^{(n)} = \sum_{r=1}^{N} a_{ir}\mu_r^{(n)}. \tag{2.133}$$

A popular choice for the α_{ij} providing fast convergence is:

$$\alpha_{ij} = \frac{a_{ij}}{\gamma_i}, \tag{2.134}$$

where $\gamma_i = \sum_{k=1}^{N} a_{ik}$ is the forward projection of the image $\vec{\mu} = 1$. The resulting separable and quadratic surrogate function is then given by:

$$Q_2\left(\vec{\mu};\vec{\mu}^{(n)}\right) = \sum_{i=1}^{M}\sum_{j=1}^{N} \frac{a_{ij}}{\gamma_i} q_i\left(\gamma_i\left(\mu_j - \mu_j^{(n)}\right) + \hat{l}_i^{(n)};\hat{l}_i^{(n)}\right). \tag{2.135}$$

One can verify that $Q_2\left(\vec{\mu};\vec{\mu}^{(n)}\right)$ is tangent to $Q_1\left(\vec{\mu};\vec{\mu}^{(n)}\right)$ at the current estimate $\vec{\mu}^{(n)}$. Thus the conditions of eq. 2.126 hold for $Q_2\left(\vec{\mu};\vec{\mu}^{(n)}\right)$ with respect to $Q_1\left(\vec{\mu};\vec{\mu}^{(n)}\right)$ and also with respect to the original likelihood term.
A similar procedure for finding a surrogate function is also applied to the penalty term:

$$R(\vec{\mu}) = \sum_{j=1}^{N}\sum_{k\in N_j} w_{jk}\psi(\mu_j - \mu_k), \tag{2.136}$$

where N_j denotes the neighborhood of voxel j. Exploiting the convexity and symmetry of the potential function $\psi(\mu_j - \mu_k)$, Jensen's inequality is used to construct a separable surrogate function $\hat{\psi}_{jk}(\vec{\mu};\vec{\mu}^{(n)})$ for $\psi(\mu_j - \mu_k)$ [183]:

$$\begin{aligned}\psi(\mu_j - \mu_k) &= \psi\left(\frac{1}{2}\left[2\mu_j - \mu_j^{(n)} - \mu_k^{(n)}\right] + \frac{1}{2}\left[-2\mu_k + \mu_j^{(n)} + \mu_k^{(n)}\right]\right)\\ \leq \hat{\psi}_{jk}(\vec{\mu};\vec{\mu}^{(n)}) &= \frac{1}{2}\psi\left(2\mu_j - \mu_j^{(n)} - \mu_k^{(n)}\right) + \frac{1}{2}\psi\left(-2\mu_k + \mu_j^{(n)} + \mu_k^{(n)}\right)\\ &= \frac{1}{2}\psi\left(2\mu_j - \mu_j^{(n)} - \mu_k^{(n)}\right) + \frac{1}{2}\psi\left(2\mu_k - \mu_j^{(n)} - \mu_k^{(n)}\right).\end{aligned} \tag{2.137}$$

Combined with equation 2.136, this leads to a separable surrogate function for the penalty term:

$$S\left(\vec{\mu};\vec{\mu}^{(n)}\right) = \sum_{j=1}^{N}\sum_{k\in N_j} w_{jk}\hat{\psi}_{jk}\left(\vec{\mu};\vec{\mu}^{(n)}\right). \tag{2.138}$$

The final separable surrogate function for the penalized log-likelihood function is defined as:

$$\varphi\left(\vec{\mu};\vec{\mu}^{(n)}\right) = Q_2\left(\vec{\mu};\vec{\mu}^{(n)}\right) + \lambda S\left(\vec{\mu};\vec{\mu}^{(n)}\right). \tag{2.139}$$

An update step simply consists of finding the minimum of N functions depending on only one variable. For quadratic data terms and penalty terms $\psi(\mu_j - \mu_k) = (\mu_j - \mu_k)^2$, the minimum can be found analytically since $\varphi\left(\vec{\mu};\vec{\mu}^{(n)}\right)$ is also quadratic in this case. If the penalty term or data term is non-quadratic, one can approximately minimize the surrogate functions, e.g. by applying a Newton-Raphson line search:

$$\mu_j^{(n+1)} = \mu_j^{(n)} - \left(\frac{\partial^2\varphi\left(\vec{\mu};\vec{\mu}^{(n)}\right)}{\partial\mu_j^2}\right)^{-1}\frac{\partial\varphi\left(\vec{\mu};\vec{\mu}^{(n)}\right)}{\partial\mu_j}. \tag{2.140}$$

In the following, several techniques that accelerate convergence and reduce the computational complexity of the separable quadratic surrogate (SQS) algorithm that was derived in the last paragraph will be presented. Without these acceleration techniques, the SQS algorithm essentially corresponds to diagonally preconditioned gradient descent optimization.

Ordered subsets The idea behind ordered subsets is to use only a subset of the available projection data in each update step. This saves computational time since the computationally expensive forward projection and backprojection is only applied to a subset of the data. The key approximation in this extension to the SQS algorithm is:

$$W\left(\sum_{i\in s_m} h_i(\hat{l}_i)\right) + \lambda R(\vec{\mu}) = \phi_m(\vec{\mu}) \approx \phi(\vec{\mu}) = \left(\sum_{i=1}^{M} h_i(\hat{l}_i)\right) + \lambda R(\vec{\mu}). \tag{2.141}$$

Here, the sum over all projections is replaced by the sum over a subset (s_m) of all projections, scaled by the factor W which corresponds to the number of subsets. The scaling factor W thus ensures that the effective regularization strength is independent of the number of subsets. One full iteration is completed by cycling through all subsets and updating the current guess on the basis of $\phi_m(\vec{\mu})$ instead of $\phi(\vec{\mu})$. Ideally, one subiteration would take approximately $1/W$ of the time of one full iteration but still decrease $\phi(\vec{\mu})$ almost as much as one full iteration. The ordered subset extension of the SQS-algorithm is not guaranteed to monotonically decrease $\phi(\vec{\mu})$ since an approximation (compare eq. 2.141) is involved. It was reported [184] that using too many subsets results in a limit cycle of the algorithm because the approximation of eq. 2.141 becomes worse with a larger number of subsets, especially close to the minimum of the penalized log-likelihood function. One possible strategy to overcome this issue might be to decrease the number of subsets as the optimization progresses towards the minimum of $\phi(\vec{\mu})$.

Precomputed curvature Computing the update (eq. 2.140) involves computing the curvature of the surrogate function $Q_2\left(\vec{\mu};\vec{\mu}^{(n)}\right)$. In most cases, an initial guess that is sufficiently close to the optimum of $\phi(\vec{\mu})$ can be found by analytical reconstruction methods. As a result, the curvature only changes slighlty during the optimization. Furthermore, monotonicity of the optimization algorithm cannot be guaranteed anymore with the ordered subsets acceleration. Consequently, the optimum curvature principle (eq. 2.130) that guarantees a monotonically decreasing penalized log-likelihood function can be abandoned. Instead of re-computing the optimum curvature at each iteration, an approximate curvature is precomputed and used for all subsequent iterations. Sufficiently close to the optimum of $\phi(\vec{\mu})$ (i.e. $\hat{l}_i \approx l_i$), it can be shown that the curvature of $Q_2\left(\vec{\mu};\vec{\mu}^{(n)}\right)$ is well approximated by [185]:

$$\frac{\partial^2 Q_2\left(\vec{\mu};\vec{\mu}^{(n)}\right)}{\partial\mu_j^2} \approx \sum_{i=1}^{M} a_{ij}\gamma_i\frac{\partial^2 h_i}{\partial\hat{l}_i^2}(l_i). \tag{2.142}$$

Precomputing the contribution of the data term to the curvature of the surrogate function $\varphi(\vec{\mu};\vec{\mu}^n)$ saves computational time since one forward and one backprojection are saved at each (sub)iteration. A potential downside of this approach is that the approximation of eq. 2.142 is only valid if a good initial guess can be obtained by analytical reconstruction methods.

Nesterov acceleration The (OS-)SQS algorithm can be combined with momentum techniques to further accelerate convergence. Momentum techniques use the gradient information from previous iterations to modify the image update step and give more weight to directions of persistent reduction in the objective function [186]. In 1983, Y. Nesterov published a momentum technique [187] that uses the last two descent directions to better inform the current update and provide a convergence rate of $O(1/n^2)$, where n denotes the number of iterations. This technique was extended in 2005 [188] to use information from all previous descent directions. Nesterov also proved that the convergence rate of $O(1/n^2)$ is optimal for any first order optimization method. Ordinary gradient descent methods like NLCG or OSSQS only provide a convergence rate of $O(1/n)$. Combining Nesterov's momentum technique with OSSQS is particularly appealing because it could in theory provide a convergence rate of $O(1/(nW)^2)$ [189] in early iterations, where W denotes the number of subsets. However, combining algorithms which cannot guarantee a monotonical descent of the penalized log-likelihood function with Nesterov's momentum technique might lead to unstable behaviour. It nevertheless turned out to be successful in practice [190, 189] . Instabilities of the algorithm were reported if the number of subsets was large ($W > 11$ in [190]) and it was observed that the 2005 version of the momentum technique is more stable for CT reconstruction [189]. Algorithm 3 and algorithm 4 present Nesterov's momentum technique (from 1983 and 2005, respectively) in combination with the OSSQS algorithm [189]. In this context, $\vec{\Delta}$ is the image update:

$$\vec{\Delta}_j(\vec{\mu}) = -\left(\frac{\partial^2 \varphi_m\left(\vec{\mu};\vec{\mu}^{(n)}\right)}{\partial \mu_j^2}\right)^{-1} \frac{\partial \varphi_m\left(\vec{\mu};\vec{\mu}^{(n)}\right)}{\partial \mu_j}, \tag{2.143}$$

where m denotes the current subset index (see eq. 2.140). The current image estimate is denoted by $\vec{\mu}$, the quantity $\vec{v}$ is the cumulative momentum from all image updates and $\vec{z}$ is a state variable that linearly combines the current image estimate with the cumulative momentum. To differentiate between full iterations (using all angular projections exactly once) and subiterations, fractional iteration indices are introduced (e.g. $\vec{\mu}^{(2+3/W)}$ corresponds to the second full iteration and the third subiteration).

Initialize $\vec{z}^{(0)} = \vec{\mu}^{(0)}, d_0 = 1$
for *t=1,2,...* **do**
  for *m=1,2,...,W* **do**
    $k = tW + m$
    $d_{k+1} = \frac{1}{2}\left(1 + \sqrt{1 + 4d_k^2}\right)$
    $\vec{\mu}^{\left(\frac{k+1}{W}\right)} = \vec{z}^{\left(\frac{k}{W}\right)} - \Delta\left(\vec{z}^{\left(\frac{k}{W}\right)}\right)$
    $\vec{z}^{\left(\frac{k+1}{W}\right)} =$
      $\vec{\mu}^{\left(\frac{k+1}{W}\right)} + \frac{d_{k-1}}{d_{k+1}}\left(\vec{\mu}^{\left(\frac{k+1}{W}\right)} - \vec{\mu}^{\left(\frac{k}{W}\right)}\right)$
  end
end

Algorithm 3: Nesterov's Momentum technique (1983) combined with the OS-SQS algorithm.

Initialize $\vec{z}^{(0)} = \vec{v}^{(0)} = \vec{\mu}^{(0)}, d_0 = 1$
for *t=1,2,...* **do**
  for *m=1,2,...,W* **do**
    $k = tW + m$
    $d_{k+1} = \frac{1}{2}\left(1 + \sqrt{1 + 4d_k^2}\right)$
    $\vec{\mu}^{\left(\frac{k+1}{W}\right)} = \vec{z}^{\left(\frac{k}{W}\right)} - \Delta\left(\vec{z}^{\left(\frac{k}{W}\right)}\right)$
    $\vec{v}^{\left(\frac{k+1}{W}\right)} =$
      $\vec{z}^{(0)} - \sum_{l=0}^{nW+m} d_l \Delta\left(\vec{z}^{\left(\frac{l}{W}\right)}\right)$
    $\vec{z}^{\left(\frac{k+1}{W}\right)} = \vec{\mu}^{\left(\frac{k+1}{W}\right)} +$
      $\frac{d_{k+1}}{\sum_{l=0}^{k+1} d_l}\left(\vec{v}^{\left(\frac{k+1}{W}\right)} - \vec{\mu}^{\left(\frac{k+1}{W}\right)}\right)$
  end
end

Algorithm 4: Nesterov's Momentum technique (2005) combined with the OSSQS algorithm.

2.7 Spectral X-ray imaging

Compared to conventional attenuation imaging, spectral X-ray imaging techniques provide additional information about the composition of an object by conducting measurements with two or more distinct photon energy spectra. The important special case of using two different spectra is often called dual energy imaging. Sometimes the term spectral imaging is exclusively used for imaging systems that acquire measurements with three or more distinct photon energy spectra. Sampling the energy-dependent attenuation allows to distinguish between different materials and to obtain information about the chemical composition. Basis material decomposition algorithms calculate material-specific images which can be used for quantitative imaging, reducing beam hardening artifacts and improving the contrast-to-noise ratio. This section gives an overview of different technical realizations of spectral X-ray imaging, basis material decomposition algorithms and their theoretical foundation as well as applications of spectral X-ray imaging in the fields of clinical diagnosis and non-destructive testing.

2.7.1 Technical realizations

Figure 2.25 shows an overview of the most common approaches for acquiring energy-resolved X-ray images. Although illustrated for CT, all of these approaches could also be adapted for radiography applications. One can differentiate between spectral imaging techniques that operate with different source spectra (fast kVp-switching, dual source CT, twin-beam CT) and techniques that use an energy-resolved detection mechanism (dual layer detectors, photon-counting detectors). Fast kVp-switching techniques [191, 192, 193] (see figure 2.25 (a)) change the X-ray spectrum by rapidly switching the acceleration voltage of the X-ray tube while roating the gantry. Dual source systems [194, 195] (figure 2.25 (b)) feature two X-ray tubes and two detectors that are aligned perpendicular to each other. This technique enables the use of different X-ray filters to increase the spectral separation between the two spectra. Twin-beam CT systems [196] (figure 2.25 (c)) rely entirely on two different X-ray filters to split the source spectrum into two different X-ray spectra, that are incident on the left and right half of the X-ray detector. Dual layer CT scanners [197, 198] and radiography systems [199] (figure 2.25 (d)) use two scintillation detectors that are stacked on top of each other. The inner layer preferably detects low energy photons. It transmits a large fraction of the high energy photons, which subsequently interact with the scintillation layer of the outer detector. The possibility to acquire energy-resolved measurements with PCDs has already been discussed in detail in section 2.4.3. Unlike the other technical realizations, PCD technology enables spectral imaging with three or more well separated photon energy spectra. Moreover, PCDs provide the possibility to acquire energy-resolved measurements with perfect spatial and temporal registration. Among the other systems, only dual layer CT yields well registered spectral projection measurements. Despite the theoretical advantages of PCDs for spectral imaging, current clinical dual energy CT systems are exclusively based on one of the other technical approaches. High costs as well as performance degrading effects of current PCDs (see section 2.4.3) have so far prevented the application of photon counting spectral CT for routine clinical diagnosis. However, advancements in PCD technology and recent pre-clinical [72, 30, 200, 31, 201, 202] studies indicate that clinical photon-counting spectral CT scanners will be feasible in the near future.

2.7.2 Basis material approximation

As was discussed in section 2.3.5, the attenuation coefficient depends on the X-ray energy, the density and the atomic number of the element under investigation. The mixture rule (eq. 2.56 and eq. 2.57) is used to calculate the attenuation of compounds. However, in the diagnostic energy range and for light elements ($Z < 20$), the energy-dependent attenuation of any material

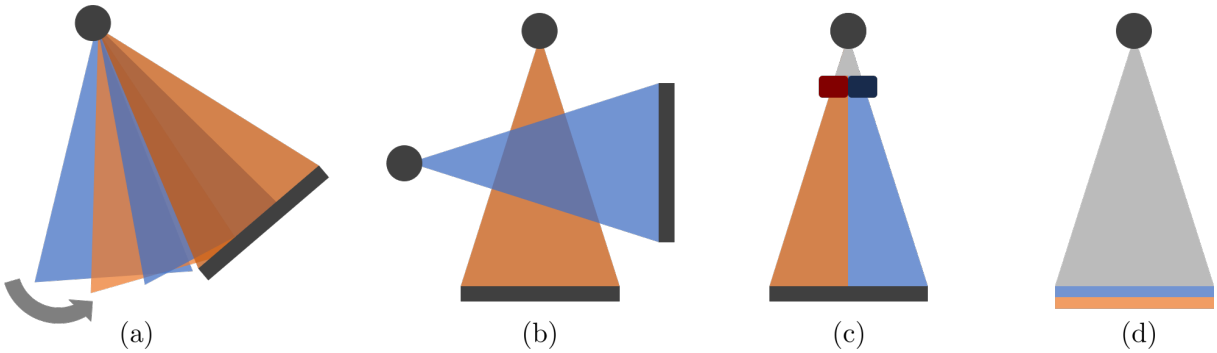

Figure 2.25: Different technical realizations of spectral CT. (a): The fast kVp-switching technique is based on rapidly changing the acceleration voltage of the X-ray tube during the CT scan. (b): Dual source systems use two X-ray tubes and two detectors oriented perpendicular to each other. (c): Twin-beam CT systems use two different X-ray filters to split the source spectrum into two different X-ray spectra. (d): Dual layer (or photon-counting) CT systems feature an energy-resolved detection mechanism.

can be approximately parameterized by two basis materials [203]:

$$\frac{\mu(E)}{\rho} = \sum_{b=1}^{2} w_b \, \frac{\mu_b(E)}{\rho_b} \tag{2.144}$$

or equivalently:

$$\mu(E) = \sum_{b=1}^{2} \alpha_b \, \mu_b(E), \tag{2.145}$$

where b is the basis material index. By comparing with the mixture rule (eq. 2.57), the coefficients w_b and α_b can be interpreted as the basis material weight fractions and volume fractions, respectively. The basis material approximation of eq. 2.144 and eq. 2.145 is motivated by the dominance of two attenuation mechanisms in the diagnostic energy range [2]: photoelectric absorption and Compton scattering. Although the energy dependency of both effects is a function of the atomic number, the variance is comparatively small for light elements. Consequently, the energy dependency of the attenuation coefficient can be approximately separated from the dependency on the atomic number [204, 205]. Figure 2.26 investigates the accuracy of the basis material approximation in more detail. In this example, the energy-dependent attenuation of a range of light elements (from $Z = 6$ to $Z = 20$) was modeled by using carbon ($Z = 6$) and calcium ($Z = 20$) as basis materials. Figure 2.26 (a) shows the relative error of modeling the energy-dependent attenuation of aluminum ($Z = 13$) with the two basis materials. In this case, the relative error e_r was defined as:

$$e_\mathrm{r}(E) = \frac{\mu_\mathrm{Al}(E) - \sum \alpha_b \, \mu_b(E)}{\mu_\mathrm{Al}(E)}. \tag{2.146}$$

For energies above 30 keV, the relative error is below 1%. In most applications, the larger relative error for lower energies (up to 3%) is less relevant because the X-ray source only emits a small fraction of the total flux in this energy range. Moreover, most of the low energy photons are attenuated by the object and thus do not contribute to the measured signal. In figure 2.26 (b), the root-mean-squared relative error (RMSRE) of approximating the attenuation coefficients with the two basis materials is plotted as a function of the atomic number. The RMSRE reaches its maximum for $Z = 13$ and decreases for elements that are closer to one of the basis materials. This behavior is reasonable since the energy dependency of the attenuation mechanisms is effectively interpolated from the two basis materials. This interpolation works

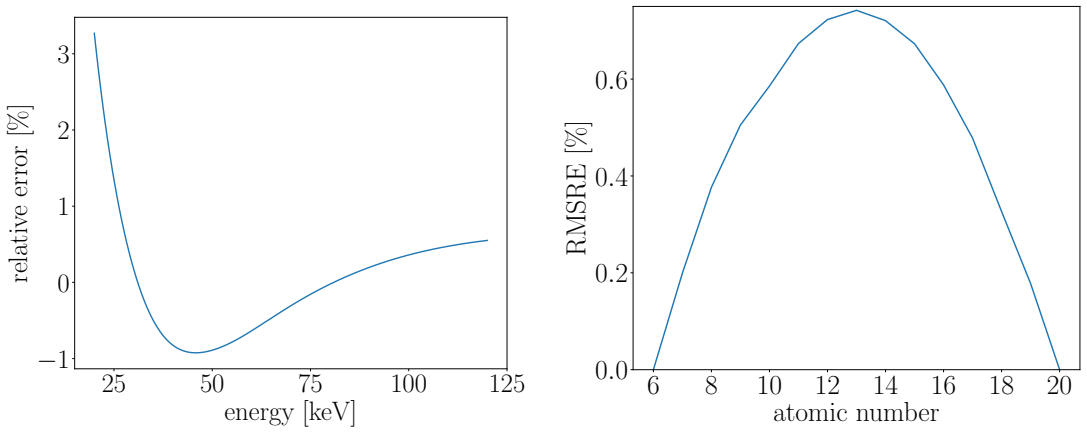

(a) (b)

Figure 2.26: Accuracy of the basis material approximation. In this example, the energy-dependent attenuation of various light elements ($Z = 6-20$) was modeled using carbon ($Z = 6$) and calcium ($Z = 20$) as basis materials. (a): Relative error of modeling the energy-dependent attenuation of aluminum. (b): Root-mean-squared relative error (RMSRE) of approximating the attenuation coefficients with the two basis materials as a function of the atomic number.

better for elements with atomic numbers close to $Z = 6$ (carbon) or $Z = 20$ (calcium).

The basis material approximation (eq. 2.144 and eq. 2.145) allows to accurately parameterize $\mu(E)$ from a limited number of energy-resolved measurements, to which current spectral X-ray imaging techniques are restricted. It is particularly well suited for medical imaging applications because the human body consists almost exclusively of light elements. If heavier elements (e.g. contrast agents such as iodine or gadolinium) with characteristic absorption edges in the diagnostic energy range are present, the parameterization can be extended to include their characteristic energy-dependent attenuation:

$$\mu(E) = \sum_{b=1}^{2} \alpha_b \ \mu_b(E) + \sum_{b=1}^{B} \alpha_b^{\mathrm{H}} \ \mu_b^{\mathrm{H}}(E), \tag{2.147}$$

where the first term denotes the usual approximation for light elements and the second term represents the contributions from B elements (or compounds) with absorption edges in the relevant energy range. In this case, $2 + B$ spectral measurements need to be conduced to unambiguously determine the volume fractions of all basis materials with the help of material decomposition algorithms (see section 2.7.3). In addition to the material-specific information that material decomposition algorithms provide, accurately parameterizing $\mu(E)$ from a few energy-resolved measurements enables the correction of *beam hardening artifacts*.

Beam hardening Beam hardening is one of the most important causes of image artifacts in conventional attenuation-based CT. In radiography applications, beam hardening can lead to incorrect quantitative values but it does not cause visual image artifacts. Beam hardening artifacts arise from a mismatch between the physics of the measurement process and the mathematical model that is used for image reconstruction. The standard Lambert-Beer model ($y_i = b_i \ \exp(-l_i)$, compare eq. 2.97) implicitly assumes that the polychromatic X-ray spectrum can be modeled by an effective energy. A more exact model could be based on a polychromatic version of the Lambert-Beer law (compare eq. 2.60):

$$y_i = \int S(E)\eta(E) \ \exp\left(-\int \mu(E,x,y,z)ds_i\right)dE \tag{2.148}$$

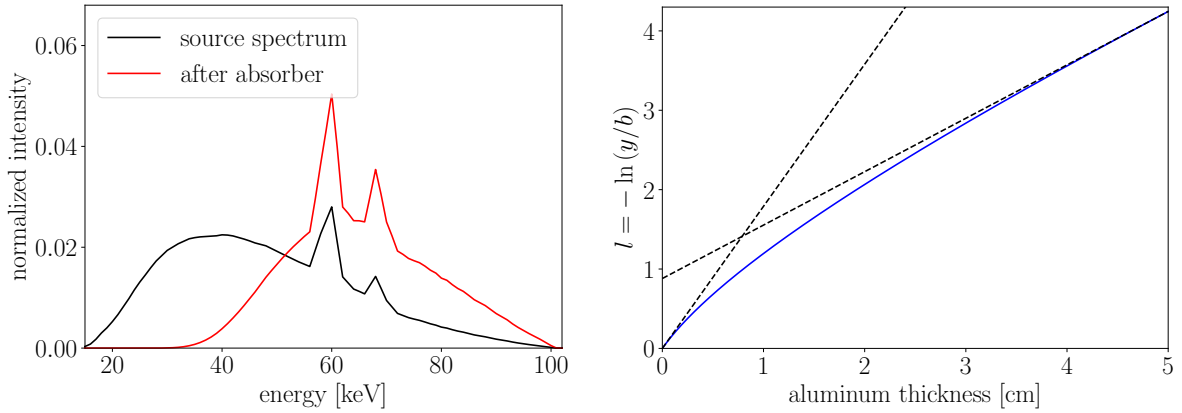

(a) (b)

Figure 2.27: Beam hardening artifacts in projection space. (a): Comparison of normalized X-ray spectra before and after transmittance through 3 cm of aluminum. (b): Beam hardening causes a nonlinear relationship between the measured line integrals and the absorber thickness. The dashed lines represent two linear functions that would be obtained with monochromatic spectra. The associated energies (which determine the slopes) correspond to the mean energy of the spectra before and after beam hardening, respectively.

However, $\mu(E)$ cannot be determined from standard attenuation measurements that only use a single X-ray spectrum. Many empirical beam hardening correction algorithms assume that the object predominately consists of one material (which has to be known a-priori) [206, 207, 208, 209, 210, 211]. This allows the application of the polychromatic Lambert-Beer law, but a complete suppression of beam hardening artifacts is hardly possible if the object consists of two or more materials with different attenuation properties. As discussed in section 2.3.5, the attenuation coefficient generally decreases with increasing photon energy. X-rays with lower energies are thus preferably attenuated and the X-ray spectrum behind an absorber is "hardened". As an example, figure 2.27 compares the normalized X-ray spectra before and after passing an absorber (3 cm of aluminum). The incident spectrum was generated by a tungsten anode operated at an acceleration voltage of 100 kVp. Beam hardening increases the effective energy of the X-ray spectrum, which causes a decrease of the apparent attenuation coefficient behind an absorber. It thus leads to a nonlinear relationship between the measured line integrals l_i and the true absorber thickness. This effect is visualized in figure 2.27 (b), where the measured line integrals l_i are plotted as a function of the true absorber thickness.
In a CT scan, beam hardening causes inconsistencies of the measured line integrals for different angular views. As a consequence, the reconstructed images are corrupted by cupping and streak artifacts. Figure 2.28 illustrates typical beam hardening artifacts by analyzing a CT reconstruction of two aluminum spheres. The axial slice in figure 2.28 (a) was reconstructed from a simulated CT scan with a polychromatic spectrum generated by a tungsten anode operated at 100 kVp. Besides streak artifacts, cupping (i.e. the apparent attenuation coefficient decreases towards the center of the object) is visible. Figure 2.28 (b) shows a lineplot through the center of the bottom sphere. The comparison between polychromatic and monochromatic reconstructions clearly shows the effect of beam hardening induced cupping.

2.7.3 Material decomposition algorithms

Material decomposition algorithms exploit the material-specific characteristics of the energy-dependent attenuation to calculate basis material images. As mentioned before, the basis material approximation plays a key role in accomplishing this task from a limited number of spectral

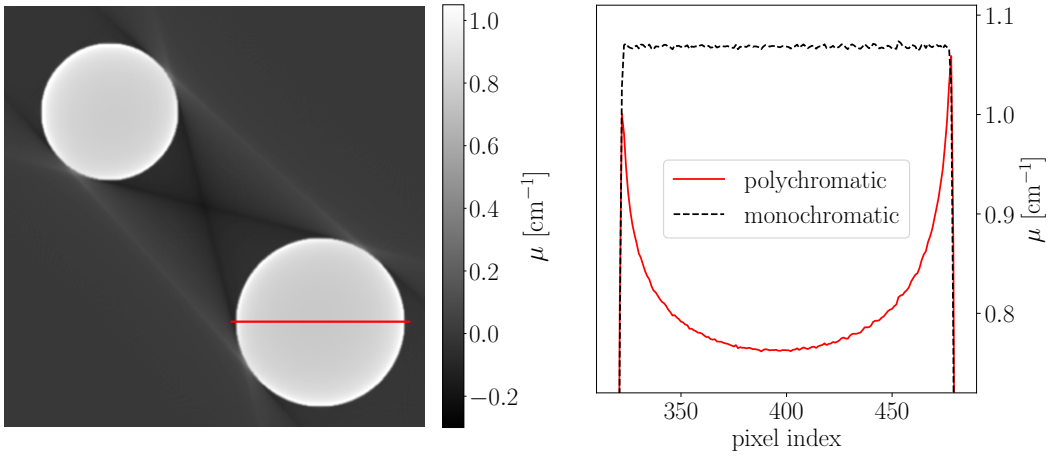

(a) (b)

Figure 2.28: Illustration of typical beam hardening artifacts using the example of a CT reconstruction of two aluminum spheres. (a): An axial slice of the reconstructed image shows streak artifacts and cupping. (b): The lineplot through the center of the bottom sphere analyzes the effect of cupping artifacts in more detail.

measurements. On the other hand, however, this means that for light elements, the basis material images do not necessarily directly provide information about the chemical composition of an object. They rather represent surrogate materials with the same attenuation properties. Figure 2.29 exemplarily illustrates a three-material decomposition task (using bone, tissue and gadolinium contrast agent as basis materials) with an idealized PCD that divides the spectrum into four non-overlapping energy bins.

The additional information provided by the basis material images has proven to be beneficial for clinical diagnosis in many different application cases (see reference [1] for a comprehensive overview). Iodine basis material images, which can be calculated from contrast-enhanced spectral CT scans, allow to distinguish between different causes of intracranial hemorrhage [212, 213]. Moreover, iodine images are used for tumor recognition [214] and staging [215] as well as detecting perfusion defects [216]. Bone and calcium basis material images facilitate the evaluation of bone mineral density [217] and the assessment of atherosclerosis [218], respectively. Other applications of material decomposition include kidney stone characterization [219], assessment of pulmonary nodules [220] and virtual non-enhanced imaging [221]. From a mathematical perspective, basis material decomposition can be viewed as an estimator that maps the energy-resolved intensity measurements to a set of basis material images. In the case of spectral CT, the energy-resolved sinogram $\vec{y}$ is mapped to the basis material volume fractions $\vec{\alpha}$ in image space:

$$\hat{\theta} : \mathbb{R}^{M \times S} \mapsto \mathbb{R}^{N \times B}, \quad \vec{\alpha} = \hat{\theta}(\vec{y}), \tag{2.149}$$

where B is the number of basis materials, N and M are the number of image voxels, and sinogram entries, respectively and $S \geq B$ represents the number of spectral measurements. Recently, so-called one-step SIR algorithms have been developed that allow a direct estimation of $\vec{\alpha}$ according to eq. 2.149. These methods and their potential advantages will be discussed in more detail in chapters 6 and 9. As of yet, one-step SIR algorithms are rarely used in practice. In order to simplify the estimation problem, material decomposition and image reconstruction are normally performed in two separate steps. Depending on the order of these two steps, two classes of spectral CT material decomposition algorithms can be differentiated: Images-based material decomposition algorithms [222, 223, 224] first reconstruct conventional attenuation images corresponding to different photon energy spectra before performing material decomposition

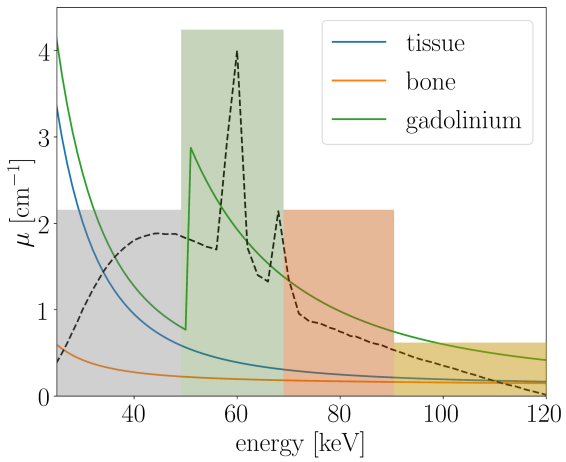

Figure 2.29: Illustration of a three-material decomposition task. Spectral X-ray imaging allows to differentiate between tissue, bone and gadolinium contrast by exploiting differences in their energy-dependent attenuation. In this example, energy-resolved measurements were acquired with an idealized PCD that features four non-overlapping energy bins.

in image space. Projection-based material decomposition algorithms [225, 226, 227, 228] first calculate the basis material line integrals $\vec{A}$ in projection space:

$$\hat{\theta} : \mathbb{R}^{M\times S} \mapsto \mathbb{R}^{M\times B}, \quad \vec{A} = \hat{\theta}(\vec{y}). \tag{2.150}$$

Projection-based material decomposition can also be used for spectral radiography applications. Under the assumption of statistically independent measurements for each sinogram index i, the estimation problem is separable:

$$\left(A_i^1, ..., A_i^B\right)^T = \hat{\theta}\left(\left(y_i^1, ..., y_i^S\right)^T\right) \quad \forall i = 1, ..., M, \tag{2.151}$$

where A_i^b is the line integral for basis material b and sinogram index i. The basis material line integrals represent the equivalent pathlenghts through the basis materials that explain the spectral projection measurements $\left(y_i^1, ..., y_i^S\right)^T$ for sinogram index i. In a second step, the basis material images are reconstructed from the calculated basis material line integrals. Figure 2.30 summarizes the three aforementioned methods to obtain basis material images.

In contrast to image-based decomposition, projection-based decomposition algorithms require spatially registered projection measurements and are therefore not compatible with some technical realizations of spectral CT (e.g. dual source CT). However, they enable an almost perfect correction of beam hardening artifacts, which is difficult to achieve with image-based methods. Furthermore, projection-based decomposition algorithms allow to incorporate an accurate statistical noise model and they form the basis for one-step SIR techniques. In this work, novel projection-based material decomposition algorithms for PCD-based spectral CT and spectral radiography applications have been developed. For these reasons, we will focus on projection-based material decomposition in the following.

Noise amplification Both image-based and projection-based material decomposition algorithms are affected by noise amplification and a degradation of the signal-to-noise ratio (SNR) compared to the unprocessed spectral images. This is still a fundamental problem of spectral X-ray imaging techniques since it limits the usability of material selective images. Noise amplification is particularly problematic for low-dose medical imaging because the increased quantum noise

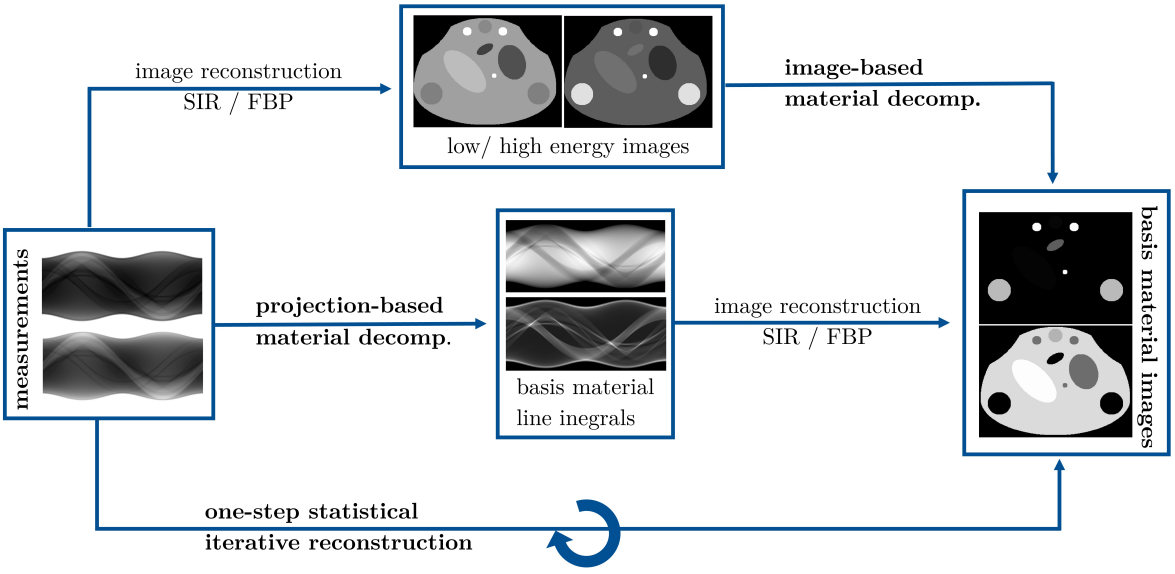

Figure 2.30: Overview of different methods for obtaining basis material images in spectral CT. One-step statistical iterative reconstruction algorithms combine image reconstruction and material decomposition in a joint framework. In most cases, however, material decomposition and image reconstruction are performed separately. Depending on the order of these two steps, image-based and projection-based material decomposition algorithms can be differentiated.

already leads to higher noise levels in the unprocessed spectral images. Consequently, advanced denoising algorithms that are specifically tailored to spectral imaging play an important role in obtaining high quality basis material images. In addition to image post-processing methods [229, 230, 231, 232, 233], MAP estimators in the form of regularized material decomposition [234, 235, 236] as well as image reconstruction algorithms [237, 238] have shown promising results.

The degradation of the SNR is most easily explained for a simplified image-based material decomposition task. Suppose that a spectral CT scan with two monochromatic X-ray sources (operated at energies E_l and E_h, respectively) has been conducted. In this case, image-based two-material decomposition reduces to solving the following system of equations for each voxel j:

$$\begin{aligned} &\text{(I)} \quad \mu_j^l = \alpha_j^1 \mu_1(E_l) + \alpha_j^2 \mu_2(E_l) \\ &\text{(II)} \quad \mu_j^h = \alpha_j^1 \mu_1(E_h) + \alpha_j^2 \mu_2(E_h), \end{aligned} \tag{2.152}$$

where $\mu_1(E)$ and $\mu_2(E)$ represent the energy-dependent attenuation coefficients of the two basis materials. For simplicity, we assume that both basis materials do not have K-edge discontinuities in the relevant energy range. The reconstructed attenuation coefficients for the low and high energy image are denoted by μ_j^l and μ_j^h, respectively. Solving for the volume fractions of the first basis material α_j^1 yields:

$$\alpha_j^1 = \frac{\mu_j^l - \mu_j^h \beta}{\mu_1(E_l) - \mu_1(E_h)\beta}, \quad \beta = \frac{\mu_2(E_l)}{\mu_2(E_h)} > 1, \tag{2.153}$$

with a similar result for the second basis material. The SNR for the basis material image is given by:

$$\mathrm{SNR}\left(\alpha_j^1\right) = \frac{\mu_j^l - \mu_j^h \beta}{\sqrt{\sigma^2(\mu_j^l) + \sigma^2(\mu_j^h)\beta^2}}. \tag{2.154}$$

As can be seen from the numerator in eq. 2.154, the signal decreases whereas the denominator in eq. 2.154 shows that the noise level of the basis material images increases compared to the attenuation images (noise amplification), resulting in a degradatio of the SNR. In the more general case of basis material decomposition with polychromatic spectra, the CRLB (see section 2.5.2) can be used to arrive at a similiar conclusion [227]. The prediction of (co-)variances with the CRLB will be discussed in chapter 8. A more detailed analysis shows that the noise in the basis material images is anticorrelated [3, 229, 2]. The degree of anticorrelation and the overall noise level increase with increasing spectral overlap of the energy-resolved measurements.
Noise amplification is also the reason why a material decomposition into more than two low-Z materials is unfeasible for almost all practical applications. In theory, based on small differences in the energy-dependent attenuation, more than two low-Z materials could be differentiated by conducting three or more spectral measurements [239]. However, as detailed in section 2.7.2, the energy-dependent attenuation of a third basis material is already closely modeled by the other two basis materials. This leads to highly increased noise levels (by a factor of $\approx 10^3 - 10^4$) [240] if a material decomposition into three or more materials without absorption edges in the relevant energy range is attempted.

Projection-based material decomposition In the following, we will assume that the spectral projection measurements have been acquired with a PCD. However, a generalization to other spectral X-ray imaging methods is straightforward. Projection-based material decomposition algorithms can be divided into two categories: algorithms based on a physical model of the measurement acquisition and empirical methods. The former use a polychromatic version of the Lambert-Beer law to construct a forward model which connects the expected spectral intensity measurements to the basis material line integrals:

$$\hat{y}_i^s = \int \phi_{\text{eff}}(E) \exp\left(-\sum_{b=1}^{B} A_i^b \, \mu_b(E)\right) dE, \tag{2.155}$$

where $\hat{y}_i^s$ is the expected intensity for detector pixel i and energy bin s of the PCD and $\mu_b(E)$ represents the energy-dependent attenuation coefficient of the b-th basis material. The effective spectrum $\phi_{\text{eff}}(E)$ includes the source spectrum and detector effects (quantum efficiency, detector response). Scattered radiation that reaches the detector is not considered in this forward model, which can lead to biased decomposition results for some applications [241, 242, 243]. Using the forward model of eq. 2.155, projection-based material decomposition is formulated as the inverse problem of determining the basis material line integrals from the spectral projection measurements. This inverse problem is most commonly solved with an ML estimator:

$$\vec{A}_{\text{ML}} = \arg\min_{\vec{A}} \; -L(\vec{A}) \tag{2.156}$$

Assuming independent Poisson statistics for the spectral measurements y_i^s, the log-likelihood function $-L(\vec{A})$ is given by:

$$-L(\vec{A}) = \sum_i \sum_s \hat{y}_i^s(\vec{A}_i) - y_i^s \ln\left(\hat{y}_i^s(\vec{A}_i)\right). \tag{2.157}$$

Due to the separable structure of $-L(\vec{A})$, the basis material line integrals can be determined independently for each detector pixel:

$$\vec{A}_i = \left(A_i^1, ..., A_i^B\right)^T = \arg\min_{\vec{A}_i} \sum_s \hat{y}_i^s(\vec{A}_i) - y_i^s \ln\left(\hat{y}_i^s(\vec{A}_i)\right). \tag{2.158}$$

Accurately determining the effective spectrum (including the source spectrum and the detector response) is often difficult in practice. This applies in particular to performing spectral X-ray

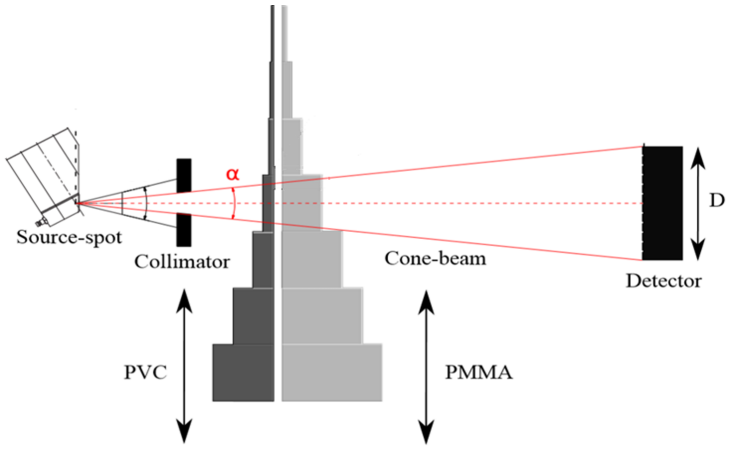

Figure 2.31: Possible technical realization of a basis material calibration measurement. The setup features two calibration phantoms (made of PMMA and PVC), which comprise several homogeneous blocks of varying thickness. By independently moving the calibration phantoms in vertical direction, transmission measurements with different basis material line integral combinations can be acquired quickly. (Modified, with permission, from reference [85].)

measurements with PCDs because manufacturing imperfections lead to inter-pixel variations in the detector response. Therefore, the spectral detector response would have to be simulated or measured individually for each pixel. Furthermore, various effects (e.g. material aging or thermal drifts) change the system parameters over time. Ignoring inter-pixel variations leads to an artificial noise pattern in projection images, even if a reference intensity (flatfield) correction is applied. Due to beam hardening effects combined with spectral distortions of the PCD response, a simple flatfield correction cannot fully remove these artifacts. Since the artifact pattern is consistent over different angular views, it causes ring artifacts in the CT reconstruction. Empirical material decomposition approaches therefore rely on calibration measurements to circumvent the problem of determining the effective spectrum in experimental measurements. The calibration measurements typically consist of a set of transmission measurements through various thicknesses combinations of two or more calibration materials. A possible technical realization for performing these transmission measurements is depicted in figure 2.31. Ideally, the calibration materials coincide with the basis materials that are used for the subsequent basis material decomposition. For low-Z materials, one can use eq. 2.145 to perform a basis material transformation [203] after material decomposition. This means representing $\mu(E)$ with a pair of basis materials that are different from the calibration materials. The calibration measurements are used to determine a number of fit parameters $\vec{c}_i$ of an empirical function that maps the spectral intensity measurements to the basis material thicknesses:

$$\left(A_i^1, ..., A_i^B\right)^T = \theta\left(\left(y_i^1, ..., y_i^S\right)^T;\ \vec{c}_i\right) \tag{2.159}$$

By determining the fit parameters $\vec{c}_i$ individually for each pixel, inter-pixel variations of the PCD response can be taken into account. Mapping the spectral measurements to the basis material thicknesses can either be performed directly, as suggest by eq. 2.159, or indirectly by combining an empirical forward model:

$$\left(\hat{y}_i^1, ..., \hat{y}_i^S\right)^T = f\left(\left(A_i^1, ..., A_i^B\right)^T;\ \vec{C}_i\right) \tag{2.160}$$

with a MAP or ML estimator. The latter approach has the advantage of efficient estimation (i.e. achieving the CRLB) in cases where the number of spectral measurements exceeds the number of basis materials ($S > B$). As could be expected intuitively, increasing the number of

energy-resolved measurements (while keeping the dose constant) generally decreases the noise level of the basis material images, although there are diminishing returns [240]. Various empirical models based on polynomial functions [244, 245, 2, 246], neural networks [247, 248] surrogate spectra [249, 250] and lookup tables [225, 251] have been proposed.

Virtual monochromatic images After reconstruction of the basis material images, virtual monochromatic images (VMIs) can be synthesized by applying eq. 2.147 individually to each voxel. VMIs simulate conventional attenuation images that would have been obtained from a CT scan with a monochromatic X-ray source. By tuning the virtual X-ray energy, the contrast and noise properties of the VMIs can be manipulated. Figure 2.32 exemplarily depicts the CNR and the absolute noise level of the VMIs as a function of the virtual X-ray energy. Although the exact shape of the graphs depends on the object and the image acquisition parameters [252], the general trends apply to a wide range of spectral X-ray imaging tasks. In clinical applications, VMIs synthesized at an X-ray energy between $40 - 70$ keV [253] have been used to enhance the image contrast and reduce the noise level, which facilitates various low-contrast detection tasks [254, 255, 256]. High energy VMIs (≈ 120 keV) enable the suppression of beam hardening and metal artifact. This is relevant for medical imaging [257, 258, 259] as well as non-destructive testing applications [260]. Using micro-CT for non-destructive testing applications as an example, figure 2.33 demonstrates how VMIs can successfully eliminate beam hardening artifacts.

Electron density and effective atomic number images Spectral X-ray imaging can provide spatially resolved information about the electron density and the effective atomic number of an object. The basis material volume fractions obtained from material decomposition allow to approximately calculate the local electron density ρ_e^j:

$$\rho_e^j = \sum_{b=1}^{B} \alpha_j^b \rho_e(b), \tag{2.161}$$

where $\rho_e(b)$ represents the electron density of the b-th basis material and j is the voxel index. There are several heuristic approaches [261] for calculating the effective atomic number Z_{eff}^j within an image voxel. Glasser [262] established an expression that is directly compatible with basis material decomposition:

$$Z_{\text{eff}}^j = \sqrt[3]{\sum_{b=1}^{B} w_b^j Z_b^3}, \tag{2.162}$$

where w_b^j and Z_b are the weight fraction and the (effective) atomic number of the b-th basis material, respectively. An important application of effective atomic number (Z-effective) images is particle beam therapy planning, where they can improve the dose calculation accuracy [263, 264, 265]. Z-effective images are also employed for some specialized tasks in clinical diagnosis [266].

2.8 Grating-based differential phase-contrast imaging

Contrary to attenuation-based methods, phase-contrast X-ray imaging techniques use the real part of the complex refractive index (compare eq. 2.26) as contrast generation mechanism. There are various approaches for translating the phase shift that an X-ray wave exhibits when traversing an object into a detectable intensity signal [267, 268]. Among these, grating-based differential phase-contrast (DPC) imaging is particularly promising for medical and non-destructive testing applications, since it places the lowest demands on spatial and temporal coherence [13]. Because

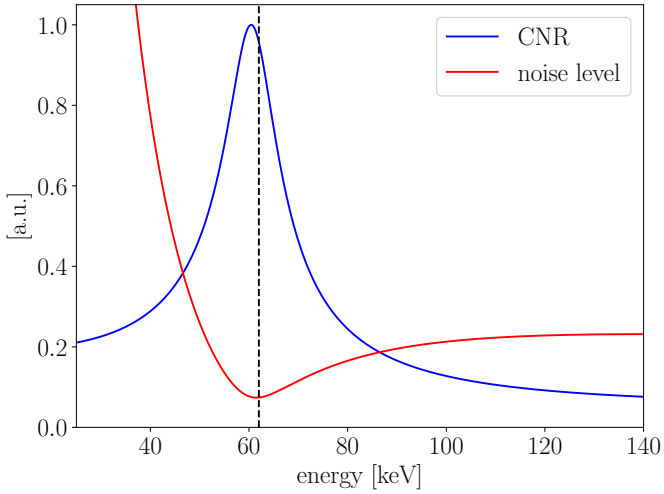

Figure 2.32: Contrast-to-noise ratio (CNR) and noise level (noise standard deviation) of virtual monochromatic images as a function of the (virtual) X-ray energy. In this example, the noise level reaches a minimum for a photon energy of 62 keV. At this point, the anti-correlated noise of the basis material images cancels out maximally. Note that the maximum of the CNR curve does not coincide perfectly with the minimum noise level.

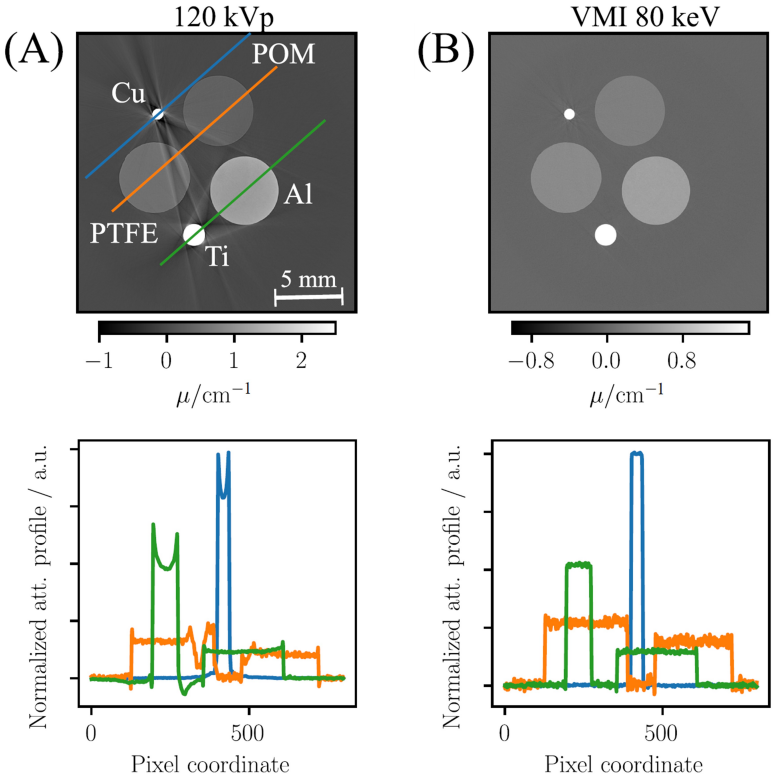

Figure 2.33: Reduction of beam hardening artifacts by calculating virtual monochromatic images (VMIs) in the field of non-destructive testing. The comparison between the conventional attenuation image (A) and the VMI at 80 keV (B) demonstrates the suppression of streak artifacts. The lineplots at the bottom show that, in contrast to the conventional image, the VMI is not affected by cupping artifacts. (Modified from reference [260].)

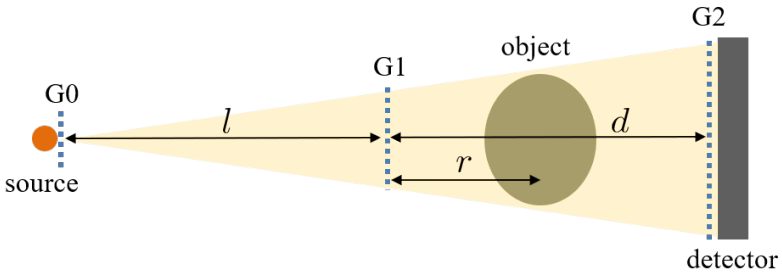

Figure 2.34: Setup for DPC imaging with standard X-ray tubes. Compared to conventional X-ray imaging, an additional Talbot-Lau interferometer comprised of three gratings (G0, G1, G2) is inserted into the beam path. The G0-G1 and G1-G2 distances are denoted by l and d, respectively and r describes the distance between the beam splitter grating (G1) and (the center of) the object. In this setup, the object is placed downstream of the G1 grating, however configurations with the object positioned upstream of the G1 grating are also possible.

of the low source coherence requirements, DPC imaging can be performed with standard X-ray tubes. Moreover, the so-called dark-field signal can be extracted from the interferometric measurements [22]. It is related to ultra-small angle scattering and provides information about the microstructure of an object far below the detector resolution. In this work, a novel model for combining grating-based DPC imaging with spectral X-ray imaging has been developed. For this reason, we will focus exclusively on grating-based phase-contrast imaging in the following.

2.8.1 Talbot-Lau interferometer

DPC imaging techniques employ a Talbot-Lau interferometer to extract the differential phase shift and dark-field signals (in addition to the attenuation information) from so-called stepping curve measurements. Figure 2.34 depicts a typical DPC imaging setup. A source grating (G0) is placed right in front of the X-ray source. The Talbot-Lau interferometer also features a beam-splitter grating (G1) and a detector grating (G2) which is placed just upstream of the detector. The G1 grating is positioned between the other two gratings. The object can be either placed downstream or upstream of the G1 grating. For simplicity, we will first consider a scenario in which a monochromatic plane wave is incident on the G1 grating before turning to polychromatic, spatially incoherent X-ray sources. This idealized model is applicable to DPC imaging experiments with synchrotron sources, where the source grating can be omitted [269, 270, 271].

Talbot effect Grating interferometry relies on the Talbot self-imaging effect which describes the phenomenon that the image of any periodic structure illuminated by a plane wave is repeated at regular distances away from the structure. The pattern of the G1 grating will thus be reproduced at certain discrete distances downstream of the grating, the so-called Talbot distances:

$$d_{\mathrm{T}}^{m} = m\frac{2p_1^2}{\lambda}, \quad m \in \mathbb{N}, \tag{2.163}$$

where p_1 is the period of the G1 grating. The Talbot effect can be derived from Fresnel diffraction theory (see for example reference [43]). In the appendix (section 11.1.2), we perform a similar (although less general) derivation by considering a cosine grating upon which an additional phase ramp is imposed. The significance of the phase ramp will become clear in the next section, where the influence of a phase shifting object on the Talbot self-images will be discussed. Although self-images can be generated with arbitrary periodic structures, we will focus on binary, rectangular gratings in the following. Binary gratings are almost exclusively used in experimental

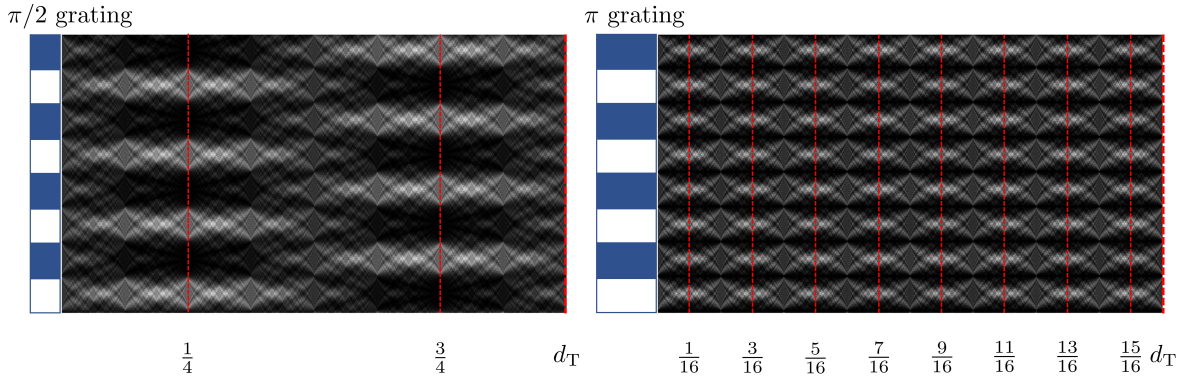

Figure 2.35: Simulated intensity distribution behind a $\pi/2$ (left) and π (right) phase shifting grating (see chapter 3 for details about the numerical simulation). The imprinted phase modulations are converted into a binary intensity patterns at the fractional Talbot distances of $d_F^m = m/4\ d_T^m$ and $d_F^m = m/16\ d_T^m$ for the $\pi/2$ and π phase shifting grating, respectively (m being an odd integer). Whereas the period of the intensity pattern is preserved for the $\pi/2$ grating, it is halved compared to the grating period for the π phase shifting grating.

DPC imaging setups because they are easiest to produce in high quality. Depending on the characteristics of the beam splitter grating, self-images can also be found at certain fractional Talbot distances d_F^m:

$$d_F^m = \frac{1}{n} d_T^m = \frac{m}{n} \frac{2p_1^2}{\lambda}, \quad m, n \in \mathbb{N}, \tag{2.164}$$

where m is the fractional Talbot order an $1/n$ describes the scaling between the fractional (d_F^m) and the full (d_T^m) Talbot distances. An important special case that is exploited for grating-based DPC imaging is the conversion of a pure phase modulation into a detectable intensity modulation for odd fractional Talbot orders (i.e. m is an odd integer in eq. 2.164). The fractional Talbot effect thus enables the application of phase shifting G1 gratings with negligible attenuation that do not reduce the X-ray flux. The shape of the intensity pattern and the location of the fractional Talbot distances (i.e. the value of n in eq. 2.164) depend on the induced phase shift and the duty cycle of the G1 grating. Suleski [272] has compiled a comprehensive overview of the different possibilities to use the fractional Talbot effect with phase shifting binary gratings. The most commonly used phase shifting gratings have a duty cycle (i.e. ratio of gap width to grating period) of 0.5 and induce a phase shift of π or $\pi/2$ at the design energy. They are typically made of low-Z materials like silicon or nickel (to reduce the X-ray attenuation) and the height of the grating bars is on the order of a few microns. By combining the projection approximation (eq. 2.27) and the formula for the refractive index decrement (eq. 2.29), it follows that the induced phase shift is inversely proportional to the photon energy. This proportionality will become relevant for polychromatic X-ray spectra. Figure 2.35 shows numerical simulations of the intensity pattern downstream of π and $\pi/2$ phase shifting gratings. As can be seen from figure 2.35, the fractional Talbot distances for π and $\pi/2$ phase shifting gratings are located at $d_F^m = m/16\ d_T^m$ and $d_F^m = m/4\ d_T^m$, respectively ($m \in 2\ \mathbb{N} + 1$). Whereas the period of the intensity pattern is preserved for a $\pi/2$ grating, it is halved compared to the grating period for a π phase shifting grating.

Analyzer grating Signal processing techniques for DPC imaging extract information about an object by analyzing the distortions of the Talbot self-image. These distortions can be resolved directly if the spatial resolution of the detector is comparable to the grating periods. However, in order to obtain a compact setup and to achieve a sufficiently high sensitivity (see section 2.8.2),

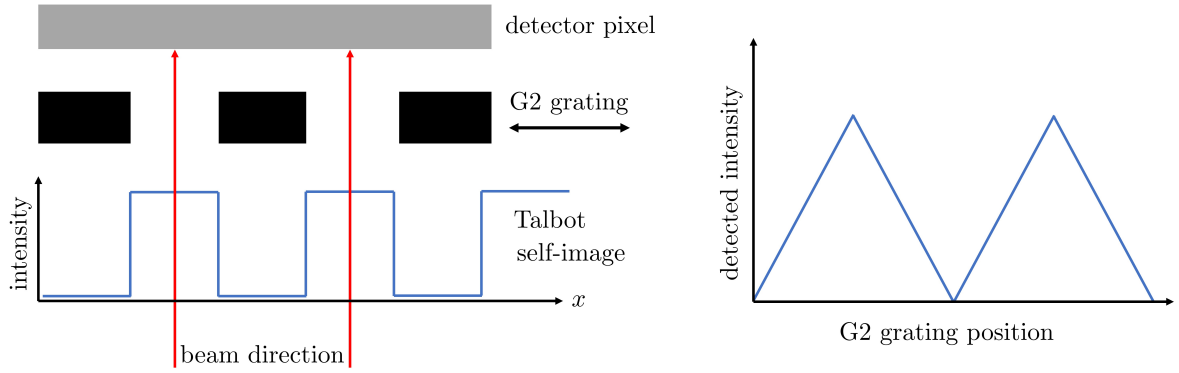

Figure 2.36: Acquisition of stepping curve measurements (in parallel beam geometry) by scanning the G2 grating across the detector. For each detector pixel, the measured intensity varies depending on the relative position of the self-image and the G2 grating. Using the stepping curve measurements, sample-induced distortions of the self-image can be detected without directly resolving the self-image pattern (see figure 2.37).

grating periods on the order of 10 μm are necessary, which is at least one order of magnitude below the typical pixel sizes of X-ray detectors used for medical imaging and non-destructive testing applications. As an example, the first fractional Talbot distance for a $\pi/2$ phase shifting grating with a period of 10 μm and a design energy of 45 keV is already at $d_F^1 = 1.8$ m. Therefore, an additional analyzer grating (G2) is placed right in front of the detector. The analyzer grating allows to convert distortions of the Talbot self image into detectable intensity modulations, eliminating the need to directly resolve the self-image pattern. It is an absorption grating that has the same period as the Talbot self-image. By moving the analyzer grating in the direction perpendicular to the grating bars, a so-called stepping curve can be recorded. The process of acquiring stepping curve measurements is illustrated in figure 2.36. Under the assumption of illuminating the phase shifting grating with a monochromatic plane wave, the measured intensity is a periodic triangular function of the analyzer grating position [273]. This result is obtained by convolving the intensity profile of the self-image with the transmission function of the analyzer grating (both of which are rectangular functions). The stepping curve measurements sample one period of the triangular function at a few equally spaced analyzer grating positions. Instead of directly resolving the Talbot self-image, signal extraction can be performed by analyzing the distortions of the stepping curve that are caused by the object. The analyzer grating thus effectively decouples the spatial resolution requirements from the grating periods.

Cone beam geometry The assumption of plane wave illumination (which corresponds to parallel beam geometry in the ray picture) is unrealistic for experimental DPC setups that use conventional X-ray tubes. In this case, a point source is a more appropriate model. However, in the next section, we will show that the spatial extent of the source spot also needs to be considered. The Fresnel scaling theorem (see section 2.3.3) allows us to transfer the results that have been obtained for parallel beam geometry to cone beam geometry [1]. Compared to parallel beam geometry, the fractional Talbot distances as well as the self-image in the detector plane are scaled by the magnification factor $M_G = (l+d)/l$, where l is the distance between the point source and the G1 grating and d is the distance between G1 and G2 (compare figure 2.34).

[1] Strictly speaking, Fresnel propagation cannot be used for large fan or cone angles, since the paraxial approximation becomes inaccurate. However, the refraction angles for X-rays are extremely small and wave propagation effects only have a local influence in the near field (see section 2.2.2). This allows the definition of local optical axes to which the paraxial approximation is applicable.

Consequently, the period of the analyzer grating p_2 also has to be scaled by M_G in order to match to the period of the self-image:

$$p_2 = M_\mathrm{G} p_1. \tag{2.165}$$

In fan- and cone beam geometries, the angle between the X-rays and the grating bars depends on the location. Outside the central beam, the incident X-rays no longer hit the grating surface at right angles. To assure that the imprinted phase shift and attenuation patterns are spatially homogeneous, the gratings have to be bent. For each of the three gratings, the radius of curvature is equal to the distance between the source and the corresponding grating. This ensures that the diverging X-rays are always perpendicular to the grating surfaces.

Source grating Conventional X-ray tubes generate an extended, spatially incoherent source spot. As was discussed in section 2.2.3, the total intensity is the sum of the intensities produced by individual incoherent (point) sources. In the appendix (section 11.1.1), we combine Fresnel propagation with the projection approximation to derive the intensity distribution for an extended, incoherent source. It can be calculated by a convolution of the point source intensity pattern with a scaled version (scaling factor $s = -d/l$) of the spatial intensity distribution in the source plane. For conventional X-ray tubes, the size of the source spot is typically much larger than the G1 grating period and the Talbot self-image would thus be completely blurred out. To prevent this, Pfeiffer et al. [13] proposed to place an additional absorption grating (G0) close to the source. This source grating divides the source spot into an array of mutually incoherent line sources. By adjusting the period of the source grating (p_0) to the period of the Talbot self-image:

$$p_0 = \frac{l}{d} M_\mathrm{G} p_1 = \frac{l}{d} p_2, \tag{2.166}$$

the intensity patterns produced by the individual line sources add constructively. This effect was first described for visible light by Lau [274]. The three-grating *Talbot-Lau interferometer* thus enables DPC imaging experiments with incoherent sources [13], which is a key step to transferring the technique to medical imaging and non-destructive testing applications.
There are different possibilities to realize a Talbot-Lau interferometer with a given total setup length. In theory, the G1 grating can be placed anywhere between the other two gratings by adjusting the grating periods of G0 and G2 according to eq. 2.166 and eq. 2.165. Symmetric ($l = d$), conventional ($l > d$) and inverse ($l < d$) geometry setups have been investigated (compare figure 2.34 for the definition of l and d). The three setup geometries have different advantages and disadvantages concerning technical aspects such as grating fabrication, as well interferometer sensitivity [275, 276].
With the additional source grating, the measured intensity $I(x)$ as a function of the analyzer grating position x can be described by a convolution of three rectangular functions. Together with the aforementioned scaling factor $s = -d/l$, the effective period of the G0 grating also matches the period of the self-image. By performing the convolutions in Fourier space [276], the observed intensity $I(x)$ can be written as a cosine series:

$$I(x) = b \left[1 + \sum_{q=1}^{\infty} v_q \cos\left(2\pi x \frac{q}{p_2} + q\phi_\mathrm{R} \right) \right], \tag{2.167}$$

where ϕ_R is a reference phase that depends on the initial positions of the three gratings. The magnitude of the Fourier coefficients v_q decrease rapidly with increasing q [276]. This motivates the standard stepping curve model, which is obtained by truncating the Fourier series after the first term:

$$I(x) = b \left[1 + V \cos\left(2\pi \frac{x}{p_2} + \phi_\mathrm{R} \right) \right], \tag{2.168}$$

where b and V describe the average value and the visibility of the stepping curve, respectively. The visibility of an intensity modulation is generally defined as:

$$V = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \tag{2.169}$$

where $I_{\max}$ and $I_{\min}$ represent the maximum and the minimum of the intensity modulation, respectively.

Polychromatic spectrum To simplify the analysis, a monochromatic X-ray source has been assumed in the previous paragraphs. Since different monochromatic components of a polychromatic wave field are mutually incoherent [277], the detected intensity is obtain by a weighted sum of the individual monochromatic intensities. At the location of the analyzer grating, the pattern of the G1 grating will only be perfectly reproduced for the design energy of the interferometer. An analysis of the Fresnel propagator in real space (eq. 2.15) shows that (for a fixed wave field in the source plane) the propagated intensity only depends on the ratio k/z_0. Scaling the X-ray energy by a factor of s is thus equivalent to scaling the propagation distance by $1/s$. As can be seen from figure 2.35, the self-images for different X-ray energies approximately resemble the G1 pattern provided that the difference to the design energy is not too large. However, the phase shift that is induced by the G1 grating also depends on the X-ray energy (compare eq. 2.29). Nevertheless, the resulting stepping curves for different X-ray energies can still be approximated by eq. 2.168 provided that the spectral bandwidth is not too large. By adjusting the visibility parameter V, the imperfect self-images can be taken into account.
Moreover, only ideal absorption gratings (G0, G2) have been considered in the last paragraphs. Current grating fabrication techniques achieve a maximum aspect ratio of ≈ 20. Although highly attenuating materials (e.g. gold) are used, a fraction of the X-ray photons is transmitted through the grating bars. This also leads to an energy-dependent reduction of the visibility of the stepping curve.

2.8.2 Signal extraction

An object that is placed in the beam path distorts the Talbot self-image and thus the measured stepping curve. By comparing the mean value, the visibility and the phase of the stepping curves to a reference scan, spatially resolved information about the attenuation, the phase shift and the dark-field signal caused by the sample is obtained. The influence of the object on the stepping curve can be expressed as [278]:

$$I(x_u) = b e^{-\mu d} \left[1 + V e^{-\epsilon d} \cos\left(\phi_u + \Delta\phi\right)\right], \tag{2.170}$$

where d is the thickness of the sample and the parameter ϕ_u denotes the reference phase for stepping position x_u of the analyzer grating:

$$\phi_u = 2\pi \frac{x_u}{p_2} + \phi_{\mathrm{R}}. \tag{2.171}$$

The most commonly used methods to determine the three signal channels (μ, $\Delta\phi$ and ϵ) from the measurements are Fourier processing [22] and least-squares fitting [279]. In chapter 8, we will discuss an ML estimator for DPC signal extraction. Similar to conventional attenuation imaging, the stepping curve model ignores the polychromatic X-ray spectrum and implicitly assumes that the stepping curve measurements can be described by an effective energy. Figure 2.37 depicts the three physical mechanisms with which an object can influence the stepping curve. The following paragraphs discuss these effects in more detail.

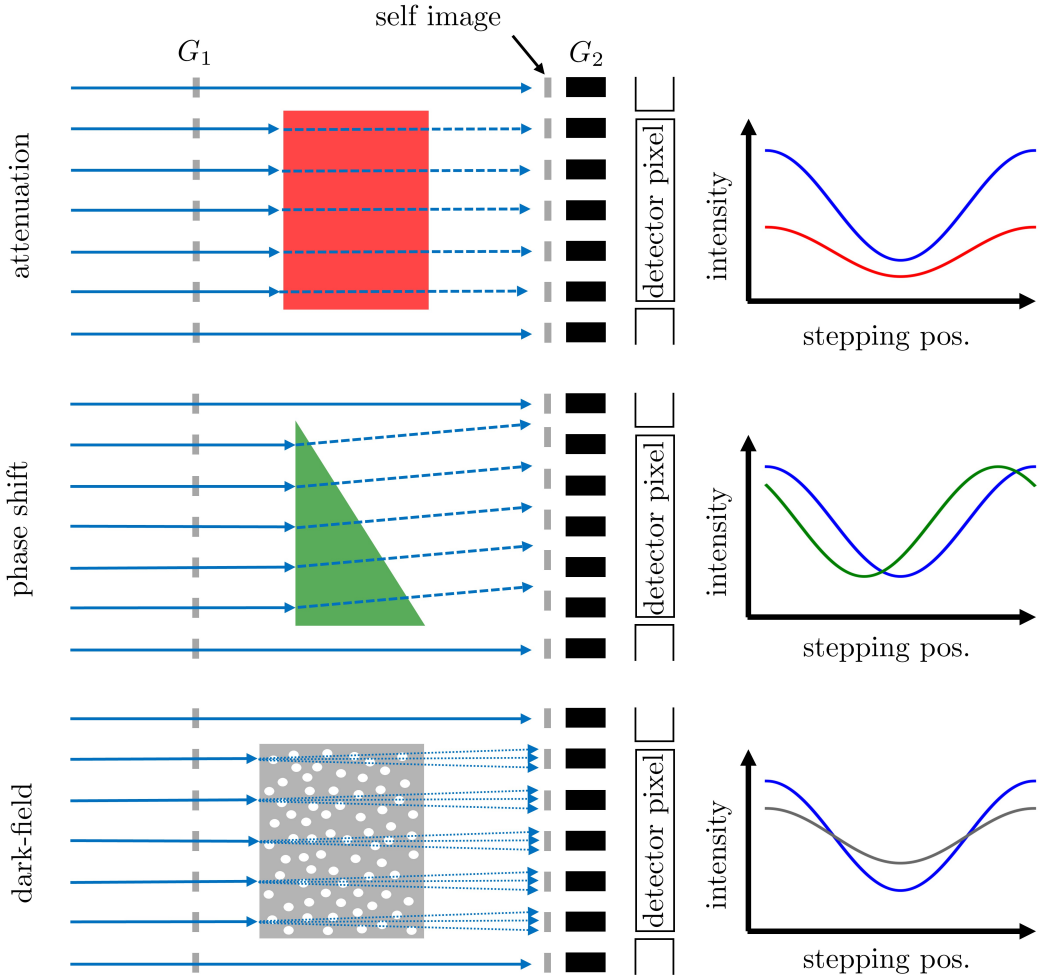

Figure 2.37: The three signal channels that can be extracted from grating-based DPC measurements and their influence on the stepping curve. An attenuating object (top row) causes an overall intensity reduction of the Talbot self-image and thus also the stepping curve. A purely phase shifting object (center row) refracts the incident X-rays which leads to a lateral displacement of the self-image. Consequently, the measured stepping curve is shifted laterally compared to the reference measurement. In the bottom row, the incident photons are scattered by a microstructured object, which leads to local distortions and blurring of the Talbot self-image. Although small angle scattering preserves the average intensity, the visibility of the stepping curve is reduced.

Attenuation X-ray attenuation leads to an overall intensity reduction of the Talbot self-image and thus also the stepping curve. This effect is modeled by the multiplicative term $e^{-\mu d}$, which reflects the Lambert-Beer law. By simply calculating the mean intensity values of the stepping curves for each pixels, the conventional attenuation image is reproduced. However, compared to a setup for conventional X-ray imaging, the grating interferometer reduces the X-ray flux and hardens the source spectrum. Consequently, attenuation images obtained from conventional and DPC X-ray imaging setups differ slightly although the same contrast generation mechanism is used.

Phase shift As was discussed in section 2.3.4, an object can cause refraction of the incident X-rays. The refraction angle is related to the gradient of the projected electron density of the object (compare eq. 2.37):

$$\alpha = -\frac{r_e (hc)^2}{2\pi E^2} \frac{\partial}{\partial x} \int \rho_e(x,y,z) dz. \tag{2.172}$$

The hybrid model of wave propagation (Talbot effect) and geometrical optics (ray refraction) immediately explains the phase shift $\Delta\phi$ of the stepping curve that is caused by a constant gradient of the projected electron density. Refraction leads to a lateral displacement of the Talbot self-image. In parallel beam geometry and assuming that the object location coincides with the G1 position, the self-image is displaced by $\Delta x = \alpha d$ which causes a phase shift of the stepping curve:

$$\Delta\phi = -2\pi \frac{\Delta x}{p_2} = -2\pi \frac{d}{p_2} \alpha \equiv \mathcal{S}_0^\alpha \alpha, \tag{2.173}$$

where the angular sensitivity $S_0^\alpha = \Delta\phi/\alpha$ has been introduced. It determines the proportionality factor between the phase shift $\Delta\phi$ of the stepping curve at a given refraction angle α and thus has a large influence on the CNR of the reconstructed images. In the appendix (section 11.1.2), we derive the result of eq. 2.173 by combining Fresnel propagation and the projection approximation, i.e. without the hybrid wave/ray picture. In cone beam geometry, eq. 2.173 has to be modified. Assuming that the object is placed between G1 and G2, the angular sensitivity is given by [275]:

$$\mathcal{S}^\alpha = \mathcal{S}_0^\alpha \left(1 - \frac{r}{d}\right), \tag{2.174}$$

where r is the distance between the G1 grating and the object. By combining eq. 2.172 with eq. 2.174, the phase shift of the stepping curve can be directly related to the gradient of the projected electron density:

$$\begin{aligned} \Delta\phi_i &= \frac{r_e d}{p_2} \left(1 - \frac{r}{d}\right) \frac{(hc)^2}{E^2} \frac{\partial}{\partial x} \int_{C_i} \rho_e(x,y,z) ds \\ &= \frac{M}{a} \frac{r_e d}{p_2} \left(1 - \frac{r}{d}\right) \frac{(hc)^2}{E^2} \left(\tilde{\rho}_e^{(i+1)} - \tilde{\rho}_e^{(i)}\right) \equiv \mathcal{S}^e \left(\tilde{\rho}_e^{(i+1)} - \tilde{\rho}_e^{(i)}\right), \end{aligned} \tag{2.175}$$

where i is the pixel (or sinogram) index, C_i is the line that connects the source and detector pixel i and $\tilde{\rho}_e^i = \int_{C_i} \rho_e(x,y,z) ds$ is the projected electron density for detector pixel i. In the second line of eq. 2.175, the gradient of the projected electron density was approximated by a forward difference:

$$\frac{\partial}{\partial x} \int_{C_i} \rho_e(x,y,z) ds = \frac{M}{a} \left(\tilde{\rho}_e^{(i+1)} - \tilde{\rho}_e^{(i)}\right), \tag{2.176}$$

where a is the detector pixel size. The object magnification M has to be taken into account because the projected electron density profile that is measured at the detector plane is magnified by a factor of M compared to the object plane. Equation 2.175 directly connects the phase shift of the stepping curve to the gradient of the projected electron density. By integrating (or rather summing) the phase shifts $\Delta\phi_i$, the projected electron density profile is obtained. For CT

applications, the electron density line integrals can then be used to reconstruct a 3D image of the electron density distribution. With the Hilbert filter, integration and image reconstruction (FBP) can be performed together in one step [14]. The integration step and the effect of the sensitivity parameter $\mathcal{S}^e$ on the noise level of the electron density images will be discussed in chapter 8. Laboratory experiments have demonstrated that the electron density images obtained by DPC imaging can provide a strongly improved contrast-to-noise ratio compared to attenuation-based imaging, particularly for low-Z, tissue-like materials [15, 16, 17, 19, 18]. As will be discussed in chapter 8, the CNR depends strongly on the imaging parameters. Examples of potential clinical applications include mammography [280, 281, 282], assessment of atherosclerotic plaque [283, 284], cartilage imaging [285, 286] and histological examinations [287, 288]. Similar to dual energy techniques, DPC imaging offers two complementary image channels (the attenuation coefficient and the electron density) that can be used to calculate the effective atomic number [289, 290, 291] and to obtain basis material images [21, 20].

Dark-field signal Grating-based DPC imaging provides a third contrast channel that is inaccessible to spectral or conventional attenuation-based X-ray imaging methods [22]. The dark-field signal is related to unresolvable microstructures with typical length scales far below the detector resolution. According to the projection approximation (see section 2.3.2), a microstructured object induces small local phase shifts to an incident wave function. Using the eikonal equation of geometrical optics (section 2.3.4), rays can be defined as orthogonal trajectories to the wavefront. In the geometrical optics picture, local distortions of the wavefront can thus be interpreted as ultra-small angle scattering on the microstructures. The small phase variations (or the scattered X-rays) locally distort the Talbot self-images. The same effect was used to explain the phase shift of the stepping curve, however the key difference is that a constant phase shift (i.e. a constant electron density gradient over the length scale of a detector pixel) was assumed for the differential phase-contrast signal. Under these conditions, X-rays are refracted which causes a uniform shift of the self-image pattern. In case of the dark-field signal, the distortions of the wavefront and thus the Talbot self-image vary significantly over small length scales. This leads to a reduced visibility of the stepping curve. With the help of wave-optical calculations [292, 293] and small angle neutron scattering theory [294], the visibility reduction can be quantified and connected to the autocorrelation function of the electron density:

$$\frac{V_s(\xi)}{V(\xi)} = e^{\int_{\text{path}} \sigma(E)(G(\xi)-1)dz}, \tag{2.177}$$

where ξ is the correlation length, V_s/V is the ratio of the visibilities with and without the sample in the beam path, $\sigma(E) \propto 1/E^2$ is the macroscopic scattering cross section and $G(\cdot)$ is the projected normalized autocorrelation function of the scattering length density ρ_{sl} [295, 296] along the beam direction:

$$G(x) = \int \gamma(x,0,0)dz, \quad \gamma(\vec{r}) = \frac{\int_V \rho_{\text{sl}}(\vec{s})\rho_{\text{sl}}(\vec{s}+\vec{r})d\vec{s}}{\int_V \rho_{\text{sl}}(\vec{s})\rho_{\text{sl}}(\vec{s})d\vec{s}}. \tag{2.178}$$

Assuming standard one-dimensional gratings, wavefront fluctuations in the direction parallel to the grating bars (y-direction) have no influence on the detected intensity for the stepping curve measurements. Consequently, the grating interferometer is only sensitive to microstructures in the direction perpendicular to the grating bars (x-direction). The scattering length density is directly proportional to the electron density far away from the absorption edges (which is normally the case for medical imaging). If the sample is placed between the phase shifting grating (G1) and the detector grating (G2), the correlation length ξ is calculated as:

$$\xi = \frac{hc}{E}\frac{d_{S,G2}}{p_2}, \tag{2.179}$$

where $d_{S,G2}$ is the distance between the sample and the G2 grating. Neglecting the spatial extent of the object along the projection direction (i.e., a constant $d_{S,G2}$ is assumed for the whole object) and assuming a homogeneous dark-field generating material, eq. 2.177 can be simplified to:

$$\frac{V_s(E)}{V(E)} = e^{d_\epsilon \sigma(E)(G(E)-1)}, \tag{2.180}$$

where d_ϵ represents the pathlength through the microstructured material. Modeling the visibility reduction by the term $e^{-\epsilon d}$ in the standard stepping curve model of eq. 2.170 is motivated by identifying the linear diffusion coefficient ϵ with:

$$\epsilon = \sigma(E_{\text{eff}})(1 - G(E_{\text{eff}})), \tag{2.181}$$

where E_{eff} is the effective energy of the polychromatic spectrum.
The dark-field signal provides structural information on the micrometer scale without the need to directly resolve these structures. In contrary to ultra high resolution X-ray imaging, dark-field imaging can be performed on large objects, using conventional X-ray tubes and standard detectors with a large field of view. This unique feature has lead to a variety of promising non-destructive testing and medical imaging applications of X-ray dark-field imaging. In the field of non-destructive testing, the dark-field signal has been used for studying fiber orientation [297, 298], revealing small cracks and fractures [299] and investigating water transport processes in porous materials [300, 301]. Potential clinical applications of X-ray dark-field imaging include foreign body detection [302], mammography [280, 303] and classification of kidney stones [304]. Many pre-clinical studies have focused on lung imaging since the microstructured air-tissue interfaces generates a strong dark-field signal. Among other things, dark-field radiography has the potential to improve the detection of emphysema [305], pulmonary fibrosis [306] and chronic obstructive pulmonary disease [307].

2.8.3 Image artifacts

In this section, a number of image artifacts that potentially arise in DPC imaging will be discussed. We will focus on image artifacts that can be removed or mitigated by combining DPC imaging with spectral imaging (see chapter 7). Most of these artifacts are caused by the fact that the standard stepping curve model ignores the polychromatic spectrum. However, without spectral measurements, the energy-dependent attenuation cannot be determined unambiguously (see section 2.7.2) and thus polychromatic effects cannot be eliminated in most cases.

Phase-wrapping Due to the periodic nature of the gratings in a Talbot-Lau interferometer, the spatial shift of the self-image and thus the phase shift of the stepping curve cannot be determined unambiguously. In other words, one can add integer multiples of 2π to the phase shift $\Delta\phi$ without changing the stepping curve model (see eq. 2.170):

$$I(x_u, \Delta\phi) = I(x_u, \Delta\phi + n2\pi), \quad n \in \mathbb{Z}. \tag{2.182}$$

If the object causes a phase shift of more than half the period of the self-image pattern, the phase will be wrapped back into the interval $]-\pi, \pi]$ [276]. Consequently, an incorrect projected electron density gradient will be determined (compare eq. 2.175), which leads to incorrect values for the projected electron density after the integration step. In the case of a CT scan, the tomographic inconsistencies subsequently cause strong artifacts in the reconstructed images. Although phase-unwrapping strategies for DPC imaging based on prior knowledge [26, 308] and the attenuation image [309, 25] have been published, a complete elimination of phase-wrapping artifacts for arbitrary objects is difficult.

Beam hardening As discussed in the previous section, both the phase shift $\Delta\phi$ and the linear diffusion coefficient ϵ depend on the X-ray energy. Similar to the attenuation coefficient μ, both quantities decrease with increasing X-ray energy. In combination with polychromatic X-ray spectra, all three DPC imaging channels are thus susceptible to beam hardening artifacts. Analogous to conventional X-ray imaging (see section 2.7.2), beam hardening artifacts are caused by the mismatch between the (implicitly) monochromatic stepping curve model and a more exact, polychromatic model:

$$I(x_u) = \int S(E) e^{-\mu(E)d} \left[1 + V(E) e^{-\epsilon(E)d} \cos\left(\phi_u + \Delta\phi(E)\right)\right] \, dE. \tag{2.183}$$

Under the assumption that the attenuation coefficient $\mu(E)$ is known a-priori and spatially uniform, Yan et al. [310] have demonstrated that a model similar to eq. 2.183 can be used to correct beam hardening artifacts in the phase shift channel. In chapter 7, we will discuss how this result can be translated to a larger class of samples by combining DPC imaging with spectral imaging. To obtain more insight into beam hardening phenomena for DPC imaging, it is useful to analyze the DPC signals that are obtained with the standard stepping curve model when the measurements are simulated with the polychromatic model of eq. 2.183. For example, since $I = I_0 e^{-\mu d}$ (according to standard stepping curve model) the attenuation coefficient μ is given by:

$$\mu = -\frac{1}{d} \ln\left(\frac{I}{I_0}\right), \tag{2.184}$$

where I and I_0 can be determined by simply averaging the sample and reference stepping curves, respectively. Describing the measured intensities I and I_0 with the polychromatic model leads to:

$$\mu = -\frac{1}{d} \ln\left(\frac{\int S(E) e^{-\mu(E)d} dE}{\int S(E) dE}\right). \tag{2.185}$$

For small sample thicknesses d, one can use a first order Taylor expansion [276, 109] to arrive at an expression for the effective linear attenuation coefficient $\bar{\mu}$:

$$\bar{\mu} = \int S^{\mathrm{N}}(E) \mu(E) dE, \quad S^{\mathrm{N}}(E) = \frac{S(E)}{\int S(E) dE} \tag{2.186}$$

The effective attenuation coefficient is thus given by a weighted average of the energy dependent attenuation coefficient $\mu(E)$. The weights are determined by the normalized spectrum $S^{\mathrm{N}}(E)$. The analysis can be extended to larger thicknesses by considering a Taylor expansion in a small interval around the mean object thickness $\bar{d}$. In this case, the weights are given by the hardened normalized spectrum $S^{\mathrm{N}}\left(E, \bar{d}\right)$:

$$S^{\mathrm{N}}\left(E, \bar{d}\,\right) = \frac{S(E) \exp\left(-\mu(E)\bar{d}\,\right)}{\int S(E) \exp\left(-\mu(E)\bar{d}\,\right) \, dE}. \tag{2.187}$$

A mathematical explanation for beam hardening artifacts in the attenuation channels can be found by considering the effective attenuation coefficient $\bar{\mu}$ as a function of the mean object thickness $\bar{d}$. With increasing object thickness, low energies are more and more attenuated and $S^{\mathrm{N}}\left(E, \bar{d}\right)$ is thus distorted in favor of higher energies. This means that more weight is given to higher X-ray energies when calculating $\bar{\mu}$ according to eq. 2.186. Since $\mu(E)$ decreases with increasing energy, $\bar{\mu}(\bar{d})$ decreases with increasing object thickness.
In analogy to the effective attenuation coefficient, the effective phase shift $\Delta\bar{\phi}$ and the effective linear diffusion coefficients $\bar{\epsilon}$ can be derived [276]. For simplicity, we will assume that the reference phase ϕ_u is independent of the X-ray energy in the following. With the additional

assumption of small electron density gradients (and thus phase shifts), the effective phase shift can be approximated as [276]:

$$\Delta\bar{\phi} = \int a^{\mathrm{N}}(E,\bar{d})\Delta\phi(E)dE, \tag{2.188}$$

where $a^{\mathrm{N}}(E,\bar{d})$ is the normalized amplitude of the stepping curve:

$$a^{\mathrm{N}}(E,\bar{d}) = \frac{S(E)\exp\left(-\mu(E)\bar{d}\right)V(E)\exp\left(-\epsilon(E)\bar{d}\right)}{\int S(E)\exp\left(-\mu(E)\bar{d}\right)V(E)\exp\left(-\epsilon(E)\bar{d}\right)dE}\ . \tag{2.189}$$

As can be seen from eq. 2.188, both the effective visibility $V(E)\exp\left(-\epsilon(E)\bar{d}\right)$ and the effective spectrum $S(E)\exp\left(-\mu(E)\bar{d}\right)$ influence the normalized amplitude spectrum $a^{\mathrm{N}}(E,\bar{d})$ which in turn influences the effective phase shift. With the standard stepping curve model, there are thus cross-talk effects between the different signals. This means that attenuation or visibility reduction effects can have an impact on the extracted phase shift. In case of the dark-field signal, it is easier to first consider the effective visibility $\bar{V}$ [276]:

$$\bar{V} = \left|\int S^{\mathrm{N}}(E,\bar{d})V(E)e^{-\epsilon(E)\bar{d}}e^{i\Delta\phi(E)}dE\right|. \tag{2.190}$$

The effective linear diffusion coefficient can then be defined as:

$$\bar{\epsilon} = -\frac{1}{d}\ln\left(\bar{V}/\bar{V}_0\right), \tag{2.191}$$

where $\bar{V}_0$ is the effective visibility of the reference stepping curve:

$$\bar{V}_0 = \int S^{\mathrm{N}}(E)V(E)dE. \tag{2.192}$$

The visibility reduction for both the polychromatic and the standard stepping curve model is described analogously to the intensity reduction (μ and $\mu(E)$ are replaced by ϵ and $\epsilon(E)$ respectively). Consequently, the dark-field signal is susceptible to the same kind of beam hardening artifacts as the attenuation signal. Moreover, equation 2.190 shows that the dark-field signal for the standard DPC model is both influenced by changes in the effective spectrum as well as the phase shift. These cross talk effects can falsely produce a dark-field signal even tough no microstructure is present. We will discuss these effects (which represent important special cases of eq. 2.190) in more detail in the following.

Beam hardening induced dark-field Distortions in the effective spectrum induce a dark-field signal that is not related to small angle scattering [24, 311, 312]. For simplicity, we assume that the object has no microstructure ($\epsilon(E) = 0$) and induces no phase shift ($\Delta\phi(E) = 0$). In this case, eq. 2.190 can be simplified to:

$$\bar{V} = \int S^{\mathrm{N}}(E,\bar{d})V(E)dE, \tag{2.193}$$

i.e. the effective visibility is calculated as the weighted average of the energy dependent visibility $V(E)$. Distortions of the normalized spectrum $S^{\mathrm{N}}(E,\bar{d})$ due to beam hardening change the weighting function for $V(E)$ and thus cause a dark-field signal. This effect is illustrated in figure 2.38.

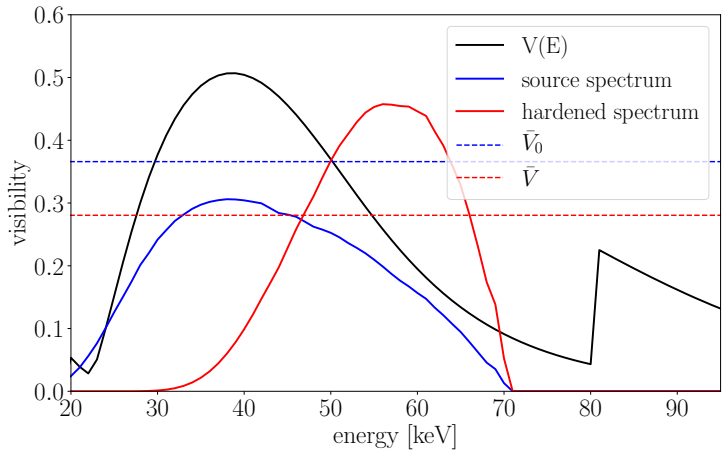

Figure 2.38: Crosstalk between beam hardening in the attenuation channel and the dark-field signal. The effective visibility $\bar{V}$ is calculated as the weighted integral of the energy-dependent visibility $V(E)$ of the interferometer (compare eq. 2.193). Beam hardening in the attenuation channel distorts the normalized spectrum (and thus the weighting function for $V(E)$) towards higher energies, which causes an artificial dark-field signal.

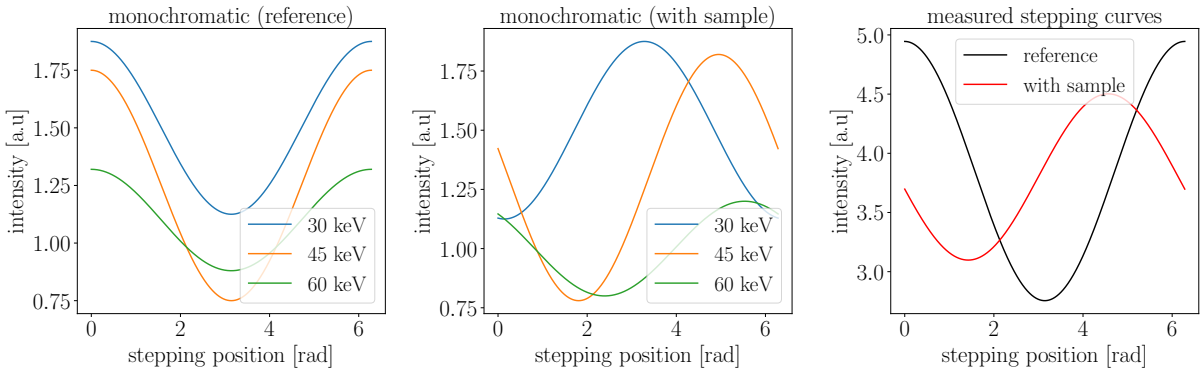

Figure 2.39: Visualization of the phase shift dispersion effect. In this simplified example, the source spectrum consists of three monochromatic energies. Since $\Delta\phi(E) \propto 1/E^2$, a purely phase shifting sample causes different phase shifts for each of the three energies (compare the graphs on the left and in the middle). As a consequence, the amplitude of the measured stepping curve (displayed on the right) is reduced compared to the reference stepping curve.

Phase shift dispersion A large electron density gradient causes cross-talk between the phase shift and the dark-field signal [311]. Since $\Delta\phi(E) \propto 1/E^2$, the phase shift for different X-ray energies varies strongly for large electron density gradients. The measured stepping curve can be viewed as the sum of individual monochromatic stepping curves (compare eq. 2.183). Adding shifted cosine functions reduces the amplitude of the measured stepping curve compared to the reference stepping curve (where $\Delta\phi(E) = 0$). Figure 2.39 demonstrates this effect for a simplified X-ray spectrum consisting of three monochromatic energies. The phase shift dispersion effect can also be deduced from eq. 2.190 by setting $\mu(E) = \epsilon(E) = 0$:

$$\bar{V} = \left| \int S^{\mathrm{N}}(E) V(E) e^{i\Delta\phi(E)} dE \right| < \int S^{\mathrm{N}}(E) V(E) dE = V_0. \tag{2.194}$$

3 Numerical simulation framework

Throughout this work, various numerical simulations of spectral- and spectral phase-contrast X-ray measurements with PCDs have been conducted. Numerical simulations allow to rapidly generate realistic input data with perfect control and knowledge of the measurement acquisition process. Moreover, the imaging parameters (such as acceleration voltage, sample properties, noise level etc.) can be varied quickly, which enables systematic studies of their influence on the image quality. Numerical simulations thus provide an invaluable tool for investigating the performance and characteristics of the reconstruction and denoising algorithms that were developed in the course of this work. There is a tradeoff between a fast, simple and traceable simulation and an accurate model of the measurement acquisition process. Since the main focus of this work is devising and testing novel algorithms for spectral- and spectral phase-contrast X-ray imaging, we employed a simulation frameworks that focuses on the most essential physical characteristics of the measurement process and neglects several effects that are of minor importance for most of the experimental measurements conduced in this work. However, depending on the application and the imaging parameters, there can be a significant mismatch between the simulated data and a real experiment. The potential influence of neglecting different effects will be discussed more thoroughly in the following sections. We first explain the simulation pipeline for spectral X-ray imaging before describing the extension to SDPC X-ray imaging.

3.1 Spectral X-ray imaging

Figure 3.1 shows the simulation pipeline for spectral and SDPC X-ray imaging. Most of the steps for simulating spectral measurements with PCDs are identical to the ones for simulating SDPC imaging. In our framework, the key difference between numerical simulations of spectral and SDPC measurements is how the influence of the sample (and the additional interferometer for SDPC imaging) is calculated. As a first step, the source spectrum $S(E)$ is simulated with the TASMIP algorithm [42] which interpolates between experimentally measured X-ray spectra. In the second step, the attenuation of the sample is considered and the transmitted intensity (as a function of the photon energy) is calculated. This central part of the simulation is based on the combination of the polychromatic Lambert-Beer law (eq. 2.148) and a forward projector (see section 2.6.2). For each projected ray (connecting the source spot and detector pixel i), the transmitted spectrum $S_i^{\mathrm{t}}(E)$ is calculated as:

$$S_i^{\mathrm{t}}(E) = S(E) \exp\left(- \sum_{r=1}^{R} \sum_{j=1}^{N} a_{ij} \alpha_j^r \mu_r(E) \right), \tag{3.1}$$

where α_j^r and $\mu_r(E)$ represent the volume fraction (in voxel j) and the energy-dependent attenuation of the $r-$th material, respectively. Prior to the simulation, a discretized 3D representation of the simulated object containing the volume fractions of all R materials found in the object has to be defined. The energy-dependent attenuation is interpolated from the NIST database [58]. For compounds, the mixture rule (eq. 2.57) is applied. Using a forward projector implies adopting the ray picture of geometrical optics. Since spectral X-ray imaging is normally performed with an extended, incoherent X-ray source, wave propagation effects can be neglected. On the other hand, the effect of an extended X-ray source is not modeled by the forward projector. An extended source spot could be included by convolving the simulated projections with a scaled

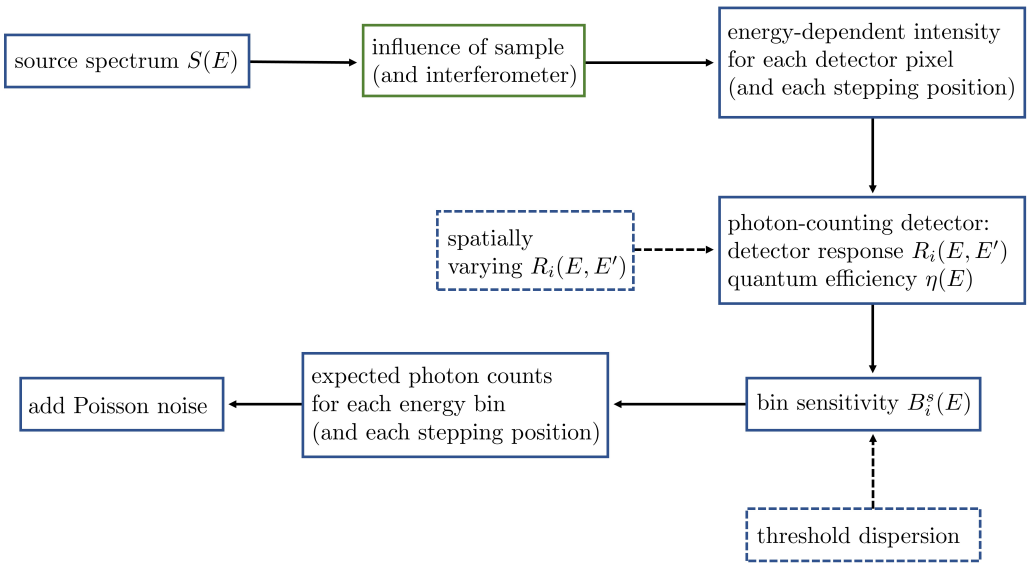

Figure 3.1: Flowchart of the algorithm for numerical simulations of spectral X-ray measurements.

version of the intensity distribution in the source plane (compare section 11.1.1). The experimental measurements in this work were performed with a microfocus X-ray tube. Geometrical considerations (pixel size, magnification, source spot size) lead to the conclusion that source blur can be neglected for this particular experimental setup. The simulation model of eq. 3.1 does not consider scattered radiation (Rayleigh- and Compton scattering) that reaches the detector. The impact of scattered radiation mainly depends on geometrical parameters, such as the illuminated volume, the object size, and the distance between the object and the detector. It has already been demonstrated that neglecting scattered radiation leads to biased decomposition results in many applications (e.g. clinical spectral CT) [313, 242]. However, computationally involved Monte-Carlo simulations have to be performed to obtain accurate scatter estimates [314, 315]. The experimental measurements in this work were performed with a small field of view (and thus a small illuminated volume) and a large air gap between the sample and the detector. Consequently, the impact of scattered X-rays is low, even though no anti-scatter grid was used. Due to the non-ideal quantum efficiency of the PCD sensor and various performance degrading effects that were discussed in section 2.4.3, the effective spectrum that is "seen" by the PCD does not coincide with the transmitted spectrum $S_i^{\mathrm{t}}(E)$. The effective spectrum $\phi_i^s(E)$ for pixel i and energy bin s of the PCD is modeled by:

$$\phi_i^s(E) = S(E)\eta(E)B_i^s(E)\exp\left(-\sum_{r=1}^{R}\sum_{j=1}^{N} a_{ij}\alpha_j^r\mu_r(E)\right), \tag{3.2}$$

where $\eta(E)$ denotes the quantum efficiency of the sensor material. The bin sensitivity $B_i^s(E)$ describes the probability that an incident photon with energy E is detected in energy bin s of the PCD. It is calculated according to an empirical PCD response model [316], which considers several performance degrading effects. We used a python simulation framework developed at our institute that implements a modified version of the aforementioned empirical response model. A detailed description of the simulation framework can be found in reference [84]. In order to calculate the bin sensitivity, $B_i^s(E)$, the energy response function $R_i(E, E')$ is considered. It describes the probability that an incident photon with energy E is counted at a (fictive) detector

threshold which corresponds to the energy E'. The energy response function is given by:

$$\begin{aligned} R_i(E,E') = c_p(E)\left(-\frac{(E-E')^2}{2\sigma_p^2(E)}\right) + c_e(E)\exp\left(-\frac{(E-E')^2}{2\sigma_e^2(E)}\right) \\ + c_f(E)\exp\left(-\frac{(E-E')^2}{2\sigma_f^2(E)}\right) + Z(E,E'). \end{aligned} \tag{3.3}$$

The first term in eq. 3.3 represents photoelectric absorption of the incident photon and the corresponding Gaussian peak is thus centered around the photon energy E. Due to electronic noise and sensor effects, the peak has a finite width $\sigma_p(E)$. By analogy, the second and third term refer to K-escape and fluorescence peaks, respectively. The normalization constants $c_p(E), c_e(E), c_f(E)$ control the relative frequency of photoelectric absorption, K-escape events and fluorescent emission and the term $Z(E,E')$ describes a slowly varying charge sharing background. For each incident energy E, the response function is normalized such that $\int (R(E,E')dE' = 1$. The bin sensitivity can be calculated from the response function by integrating between the two thresholds that form the energy bin under consideration:

$$B_i^s(E) = \int_{E_s}^{E_{(s+1)}} R_i(E,E')dE'. \tag{3.4}$$

The empirical model considers the degradation of the spectral response through various performace degrading effects. This causes an overlap of the effective spectra for different energy bins which leads to increased noise levels compared to an ideal PCD. However, the spatial effect of charge sharing and other interactions that can cause noise correlations between different detector pixels is neglected. The degree of noise correlation depends on the detector pixel size, the threshold setting and the application of a charge-sharing correction scheme. Cascaded system models and experimental measurements show that (in contrast to scintillation detectors) neighboring pixels are only weakly correlated for PCDs [317, 318]. The empirical model also neglects pulse pile-up effects, which become relevant at high flux levels. In our experimental measurements, the flux of the X-ray sources was orders of magnitudes lower than the X-ray flux of a clinical CT. In this case, pulse-pile up effects can be safely neglected. To obtain more realistic simulations for clinical imaging tasks, the simulation framework could be combined with an analytical pulse pile-up model [319]. Within the simulation framework, a non-homogeneous detector response can be simulated by either varying the parameters of the energy response $R_i(E,E')$ or slightly modifying the threshold levels from pixel to pixel. This so-called threshold dispersion leads to a spatially varying bin sensitivity (compare eq. 3.4). The expected number of photon counts $\hat{y}_i^s$ for each detector pixel i and energy bin s of the PCD is calculated by integrating the effective spectrum $\phi_i^s(E)$ (see eq. 3.2) over all energies. Numerically, this is realized by summing over a discretized effective spectrum:

$$\hat{y}_i^s = \sum_{q=1}^{Q} S(E_q)\eta(E_q)B_i^s(E_q)\exp\left(-\sum_{r=1}^{R}\sum_{j=1}^{N} a_{ij}\alpha_j^r\mu_r(E_q)\right). \tag{3.5}$$

As discussed in section 2.6.2, the Poisson noise model is a good approximation for PCDs. The measured number of photon counts y_i^s is thus simulated by drawing a random number from a Poisson distribution with expectation value $\hat{y}_i^s$.

$$y_i^s = \text{Poisson}\left(\hat{y}_i^s\right). \tag{3.6}$$

It is often advisable to perform the numerical simulations on finer grids, i.e. decreasing both to the voxel and the detector pixel spacing compared to the grids that are used for reconstruction. In a second step, the simulated detector data can be rebinned to the desired pixel size. Having

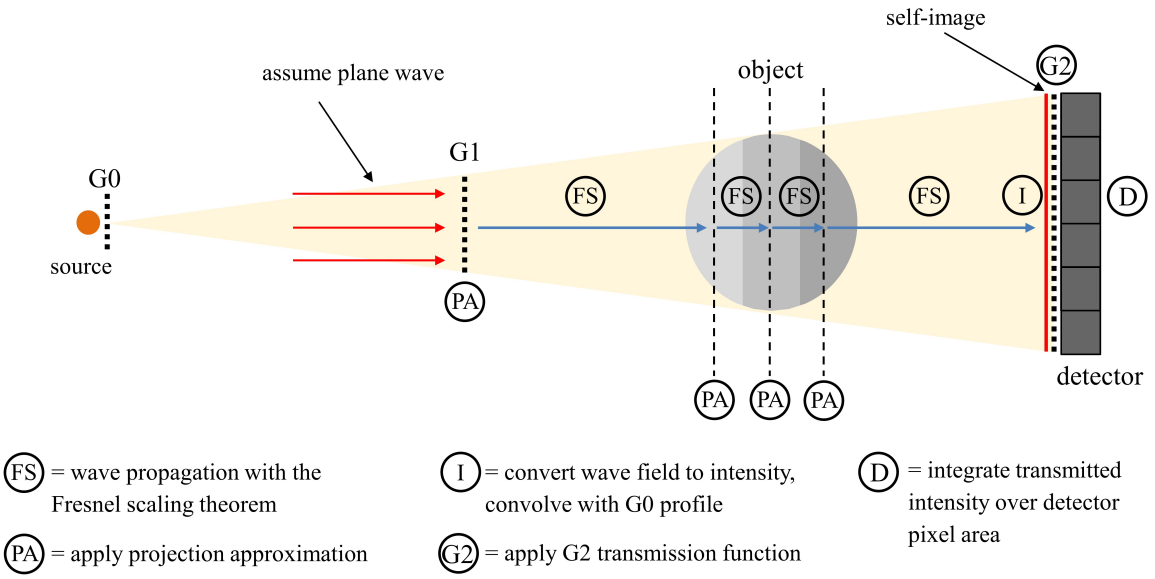

Figure 3.2: Illustration of the different algorithmic steps required to simulate grating-based DPC imaging by combining Fresnel propagation and the (multi-slice) projection approximation.

a perfect match between the simulation and the reconstruction is sometimes referred to as the "inverse crime" [115]. This type of simulation is generally sufficient for investigating the stability and other fundamental properties of a reconstruction algorithm, but the performance is likely overestimated compared to real data, i.e. an upper performance bound is obtained [115, 320].

3.2 Spectral differential phase-contrast X-ray imaging

Due to the three-grating interferometer that is inserted into the setup for SDPC (and DPC) imaging, wave propagation effects cannot be neglected anymore. In particular, the G0 grating creates an array of thin line sources that interfere constructively. Consequently, in contrast to spectral imaging, local wave propagation effects in the near field (compare section 2.2.2) are not blurred out by an extended, incoherent source. This means that a wave-optical approach is necessary to simulate the influence of the sample and the grating interferometer on the detected intensity. Towards this end, we combine Fresnel propagation with the (multi-slice) projection approximation. The simulation framework for SDPC imaging extends an existing software tool for grating simulations (gXsim) that was developed at our institute. Similarly to spectral imaging, the source spectrum is discretized and the simulation is conducted separately for each energy (since different monochromatic components of a polychromatic wave field are mutually incoherent). After the transmitted spectrum $S^{\mathrm{t}}_{i,r}(E)$ for each detector pixel and stepping position r has been obtained, the subsequent simulation steps are analogous to spectral imaging. The output of the simulation framework for SDPC imaging is the number of photon counts y_i^{rs} for each detector pixel, energy bin of the PCD and stepping position of the G2 grating.
For the wave-optical calculations, we assume that the G0 grating is placed very close to the X-ray source, i.e. propagation effect from the source spot to the G0 grating can be neglected. This means that the G0 grating can be regarded as an array of mutually incoherent line sources. As explained in section 2.2.3, the intensity in the detector plane for extended, incoherent sources can be calculated by first performing wave-optical simulations with a perfect point source and then convolving the obtained intensity profile with a scaled version of the intensity distribution in the G0 plane. We use the Fourier space version of Fresnel propagation (eq. 2.13) instead of the real space version (eq. 2.14) because of its much higher computational efficiency. However,

a point source is difficult to simulate with this approach because its Fourier transform is not bandlimited, which is incompatible with the discrete and finite-sized grid that is used in all computer simulations [50]. This problem is circumvented by assuming plane wave illumination and then converting the simulated wave fields to cone beam geometry by applying the Fresnel scaling theorem. The simulation grid has to be much finer than the grating periods ($\gtrsim 100$ samples per grating period) in order to model the interferometric effects correctly. Consequently, a large number of grid points is needed for realistic detector sizes, which makes wave propagation simulations computationally expensive. Since the X-ray source is incoherent in the direction parallel to the grating bars, the wavefront can be propagated separately for each detector row. Splitting up the 2D wavefront into a series of 1D wavefronts and individually propagating them drastically reduces the memory requirements as well as the computational complexity [321]. Considering the local character of Fresnel propagation in the near field (compare section 2.2.2), the simulation could be further accelerated by an additional splitting in the direction perpendicular to the grating bars.

Figure 3.2 shows the individual steps to calculate the intensity pattern in the detector plane. The G1 grating is assumed to be illuminated by a plane wave. In the first step, the projection approximation is used to determine the wave field immediately downstream of the G1 grating. Next, we apply the Fresnel scaling theorem (eq. 2.30) to calculate the propagated wave field at the location of the sample under point source illumination. In the appendix (section 11.1.3), we derive a slightly modified version of the Fresnel scaling theorem that can be applied to the wave fields (instead of the intensity distributions). Depending on the geometry and the sample thickness, the sample is split into one or more slabs. The slabs are subsequently approximated by a series of equidistant planes that cause the same attenuation and phase shift as the corresponding slabs. The phase shift and attenuation of each plane is applied to the wave function before using the modified Fresnel scaling theorem to propagate the wave field to the next plane. In our numerical implementation, the grid spacing for the wave field is increased with increasing propagation distance to account for the magnification of the wave field (compared to plane wave illumination) according to the Fresnel scaling theorem. However, the grid spacing for the sample is independent of the propagation distance. Consequently, the discrete representation of the sample planes has to be interpolated onto the current wave field grid. After the exit-surface wave function for the last sample plane has been calculated, the Fresnel scaling theorem is applied one more time to calculate the intensity distribution of the self-image in the analyzer grating plane. As mentioned in the beginning of this section, the obtained point source intensity pattern is then convolved with a scaled version of the intensity distribution in the G0 plane to consider the mutually incoherent line sources. Finally, the intensity immediately downstream of the G2 grating is obtained by multiplying the self-image with the spatial transmission function of the G2 grating for each stepping position. We calculate the transmitted spectrum $S^{\mathrm{t}}_{i,r}(E)$ for each detector pixel by integrating the intensity pattern over the corresponding detector pixel area.

In theory, the dark-field signal could be simulated by considering the microstructure of the object when applying the multi-slice projection approximation [322]. However, to get accurate results, the object would have to be divided into a large number of planes (> 100), which strongly increases the computational complexity, especially for large samples. Moreover, the microstructure of the sample would have to be known with sub-micrometer resolution for a realistic simulation. We therefore use a simpler approach to include the effect of ultra-small angle scattering on microstructures: The amplitudes of the energy-resolved stepping curves $S^{\mathrm{t}}_{i,r}(E)$ are scaled by the factor $\exp\left(-d_\epsilon f_\epsilon(E)\right)$, where d_ϵ is the pathlength through the microstructured material and $f_\epsilon(E)$ describes the visibility reduction per unit length as a function of the photon energy (which could for example be determined by experimental measurements). This approximation neglects the position dependency of the dark-field signal via the correlation length. In chapter 7, we will discuss a more exact heuristic approach for simulating the dark-field signal that is based on eq. 2.177.

The limitations of the simulation framework as well as the neglected physical effects that were discussed for spectral imaging (see last section) also apply to SDPC imaging. Additionally, the simulation framework for SDPC imaging is restricted to samples for which the inhomogeneous Helmholtz equation (eq. 2.24) can be used. For most medical imaging and non-destructive testing applications cases, the assumption of a linear, isotropic and non-magnetic sample is fulfilled. The assumption that the complex refractive index varies slowly compared with the wavelength of the X-ray radiation excludes certain objects (e.g. crystals) from the simulation. Since the deflection angles for X-rays are small, the Fresnel approximation to wave propagation is generally valid.

4 A post-processing algorithm for spectral CT material selective images using learned dictionaries

As was discussed in section 2.7.3, material-selective images suffer from noise amplification and a degradation of the SNR. Combining spectral X-ray imaging with advanced denoising algorithms is therefore essential in order to enhance the image quality and usability of basis material images for various tasks such as artifact reduction, quantitative imaging and clinical diagnosis. In this work, several denoising techniques for spectral X-ray imaging have been explored. Regularized material decomposition and one-step statistical iterative reconstruction techniques for spectral CT will be discussed in chapter 5 and 6, respectively. In this chapter, an image post-processing technique for spectral CT based on learned dictionaries will be presented. Dictionary denoising is a powerful noise reduction technique which separates image features from noise by modeling small image patches as a sparse linear combination of dictionary atoms. These dictionary atoms are learned from training images prior to the denoising process. An important advantage of image post-processing techniques is their flexibility. Since they only require the decomposed basis material images as input, post-processing techniques can be combined with different technical realizations of spectral CT and image-based as well as projection-based material decomposition algorithms. Material selective images typically show strong structural correlations, i.e. edges and other small structures tend to be aligned in different image channels. We have adapted the standard dictionary denoising technique to make use of the structural correlation as well as the anti-correlated noise which is typically present in material selective images. Numerical simulations and an experimental measurement show that our algorithm achieves improved image quality compared to two other post-processing methods, namely conventional dictionary denoising and bilateral filtering. Moreover, the novel multi-channel post-processing algorithm serves as a basis for the development of dictionary-based multi-channel regularization strategies, which will be discussed in chapter 5.

The text and the figures in this chapter are taken from the following publication of the author:

- Mechlem, Korbinian, et al. "A post-processing algorithm for spectral CT material selective images using learned dictionaries." Biomedical Physics & Engineering Express 3.2 (2017): 025009.

Compared to the original publication, the text has been slightly modified to achieve a better integration into this doctoral thesis.

4.1 Introduction

The importance of suppressing noise and improving the image quality is reflected by the development of several denoising algorithms for material decomposition in spectral CT [230, 223, 229, 231]. Another approach is to incorporate the structural correlations between different spectral images prior to material decomposition. This can be realized by creating a reference image (e.g. by summing up all spectral channels) which is then used for noise reduction in the individual spectral images. Different strategies based on similarity matrices [323], total-variation regularization [324] and the correlation coefficient [325] have been investigated. In the last years,

dictionary-based denoising has received a lot of attention in the scientific community. The key idea is to model small image patches as a sparse linear combination of learned dictionary atoms. Motivated by the success of the method in fields like image/video denoising [326, 327] and magnetic resonance imaging [328], this technique has also been applied to CT image reconstruction. Dictionary-based denoising was used for artifact suppression [329] and few-view [330] as well as low-dose [331] image reconstruction in conventional CT. The unique capability of learned dictionaries to differentiate between noise and features leads to strong noise suppression. Recently, a tensor-based dictionary learning approach for spectral CT reconstruction has been proposed [332]. By generalizing the image patches to tensors which include the spectral dimension, the structural correlation between spectral images is taken into account. However, the strong noise correlations makes it difficult to translate this approach directly to material selective images. In this chapter, we present a new dictionary-based algorithm for denoising material selective images. It exploits the structural correlation between basis material images as well as the fact that noise in material selective images is typically highly anti-correlated [229]. Since the algorithm is applied to basis material images as a post-processing step, it is applicable to both image-based and projection-based material decomposition. We demonstrate, on simulated as well as real measurement data, that the denoising algorithm leads to strongly improved image quality. Moreover, our the novel algorithm achieves superior image quality compared to two other post-processing methods, namely conventional dictionary denoising and bilateral filtering [333].

4.2 Methods

Our algorithm for denoising of basis material images is a modification of the well established dictionary denoising process aimed at exploiting the structural correlation between material selective images as well as the anti-correlated noise. First we describe dictionary denoising, which was used as a basis for our new algorithm, before explaining the modifications made to achieve enhanced performance for basis material images.

4.2.1 Dictionary denoising

Dictionary denoising techniques divide the image into small overlapping patches which are processed individually. The final denoised image is then compiled from the individually processed image patches. The degree of overlap of the patches can be described by the sliding distance, which is defined as the distance (measured by the number of voxels) between the centers of adjacent image patches. For illustration purposes, figure 4.1 shows the extraction of overlapping patches (size 3×3 pixels) from a 2D-image, using a sliding distance of two pixels. However, we chose to use 3D cubic patches for our algorithm in order to avoid horizontal streak artifacts in sagittal or coronal views [334]. The key assumption of dictionary denoising is that natural image patches have a sparse representation in a suitable basis. This means that they can be modeled as a linear combination of a small number n of basis functions, with $n \ll N$, where N denotes the number of voxels in an image patch. To ensure an optimum sparse representation, the basis functions (also called dictionary atoms) are obtained from application-specific training images by a dictionary learning algorithm. The final dictionary represents an over-complete basis comprised of K ($K \gg N$) dictionary atoms which reflect typical structures occurring in the training images. In general, learning the dictionary atoms from application-specific training images allows for a more sparse representation of the image patches compared to using a generic basis. Contrary to structures and image features, noise cannot be sparsely represented with the dictionary basis. By modeling the image patches as a linear combination of a few dictionary atoms, noise is therefore effectively suppressed. Since dictionary denoising operates on image patches instead of individual voxels and is able to 'recognize' image features, the tradeoff be-

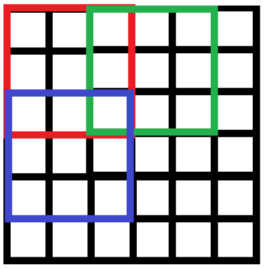

Figure 4.1: Graphical illustration of covering a 2D-image with overlapping patches (size 3×3 pixels), using a sliding distance of two pixels.

tween noise and resolution can be partially mitigated. This tradeoff is typical for many denoising methods which operate on individual voxels (e.g. filtering, nearest-neighbor-based regularization). In the following, the mathematical description of dictionary denoising will be introduced. An image patch of $d \times d \times d$ voxels can be expressed as a N-dimensional vector $x \in \mathbb{R}^N$, $N = d^3$. The dictionary comprised of K atoms is written as a matrix $D \in \mathbb{R}^{N \times K}$ where the columns of the matrix represent the dictionary atoms. The denoising problem for one patch can then be formulated as:

$$\min_{\alpha} \| \alpha \|_0 \text{ s.t. } \| x - D\alpha \|_2^2 \leq \epsilon, \tag{4.1}$$

where $\alpha \in \mathbb{R}^K$ is a vector with few nonzero entries, $\epsilon > 0$ is a small error tolerance and $\| \bullet \|_0$ is the l_0-norm. Because solving eq. (4.1) is NP-hard, greedy algorithms are employed to compute an approximate solution in acceptable time. We chose to use the orthogonal matching pursuit algorithm [335]. The dictionary atoms were learned from training images before the denoising process by solving the following optimization problem:

$$\min_{\alpha, D} \sum_{s=1}^{S} \| \alpha_s \|_0 \text{ s.t. } \sum_{s=1}^{S} \| x_s - D\alpha_s \|_2^2 \leq S\epsilon, \tag{4.2}$$

with the help of the fast online learning method [336]. The index s indicates the different patches which were extracted from the training images. We chose to employ a "global dictionary" approach. This means that the dictionary is fixed prior to the denoising process and not dynamically adapted during denoising ("adaptive dictionary"). Using a fixed dictionary saves computational time and the resulting image quality is comparable to adaptive dictionary denoising for conventional CT images [331].

4.2.2 Modifications for denoising of basis material images

Although our algorithm theoretically works with more than two basis material images, we will focus on two-material decomposition in the following. The performance of the algorithm for three-material decomposition will be investigated in the future. Material decomposition algorithms for spectral CT typically produce two basis material images with highly anti-correlated noise. Suitable linear combinations of the basis material images yield virtual monochromatic images which represent the attenuation at a certain reference energy. The associated coefficients of the linear combination are positive for both basis materials and therefore anti-correlated noise is decreased. By choosing the reference energy at which the anti-correlated noise maximally cancels out, a virtual monochromatic image with substantially improved signal-to-noise ratio (SNR) compared to the basis material images is obtained. In general, it is much easier to separate image features from noise in this virtual monochromatic image (called minimum noise image in the following) than performing the same task for the basis material images. Furthermore, the minimum noise image shares the same structures and edges with the basis material images while

the voxel values in all three images are different. The image patches x_s can be assumed to be the sum of their ground truth values (x_s^t) and noise (x_s^n):

$$x_s = x_s^t + x_s^n, \tag{4.3}$$

with s being the image patch index. A key assumption for the following steps is that the ground truth values of the minimum noise image patches and the basis material image patches are related by linear transformations:

$$\begin{aligned} x_s^{a,t} &= m_s^a + \beta_s^a x^{o,t}, \\ x_s^{b,t} &= m_s^b + \beta_s^b x^{o,t}, \end{aligned} \tag{4.4}$$

where $x_s^{a,t}$, $x_s^{b,t}$ and $x^{o,t}$ indicate the ground truth values of the s-th image patch of basis material image a, b and the minimum noise image, respectively. The fit coefficients for the linear transformation from the minimum noise image patches to the image patches of basis material a and b are denoted by m_s^a, β_s^a and m_s^b, β_s^b, respectively. In the following, the noisy image patches will be denoted by x_s^a, x_s^b and x_s^o.

The first step of our algorithm is to subtract the mean values m_s^o from the minimum noise image patches x_s^o in order to identify edges and structures independent of a constant offset. From now on, offset-corrected patches will be marked with a tilde (e.g. $\tilde{x}_s^o$). After the offset correction, we perform dictionary denoising on the image patches $\tilde{x}_s^o$:

$$\begin{aligned} \forall s: \quad \alpha_s = \min_{\alpha} \| \alpha \|_0 \text{ s.t. } \| \tilde{x}_s^o - D\alpha \|_2^2 \leq \epsilon, \\ \tilde{x}_s^{o,d} = D\alpha_s, \end{aligned} \tag{4.5}$$

where $\tilde{x}_s^o$ denotes the s-th patch from the minimum noise image and $\tilde{x}_s^{o,d}$ denotes the corresponding denoised patch. Using the assumption of eq. 4.4, we will compute the denoised basis material image patches $x_s^{a,d}$ and $x_s^{b,d}$ by applying linear transformations to the processed minimum noise image patches $\tilde{x}_s^{o,d}$. Similarly to the minimum noise image, the mean values m_s^a and m_s^b are subtracted from the image patches x_s^a and x_s^b in order to obtain the offset-corrected image patches $\tilde{x}_s^a$ and $\tilde{x}_s^b$. Adopting the same notation as in eq. 4.4 implies that the mean values of the image patches will be used to approximate the corresponding coefficients of the linear transformations. The fit coefficients β_s^a and β_s^b are approximated by projecting the denoised and offset-corrected patches of the minimum noise image onto the offset-corrected basis material image patches:

$$\beta_s^r = \frac{\langle \tilde{x}_s^r, \tilde{x}_s^{o,d} \rangle}{\| \tilde{x}_s^{o,d} \|_2^2}, \quad r = \{a, b\}, \tag{4.6}$$

where $\langle \bullet \rangle$ signifies the scalar product. We obtain the final denoised basis material image patches $x_s^{a,d}$ and $x_s^{b,d}$ by performing the following linear transformations with the denoised and offset-corrected minimum noise image patches $\tilde{x}_s^{o,d}$:

$$\begin{aligned} x_s^{a,d} &= m_s^a + \beta_s^a \tilde{x}_s^{o,d}, \\ x_s^{b,d} &= m_s^b + \beta_s^b \tilde{x}_s^{o,d}. \end{aligned} \tag{4.7}$$

In summary, dictionary denoising is used to identify structures and edges in the minimum noise image. This process works much more efficiently and reliably on the minimum noise image compared to the basis material images because of the improved image quality and the reduced noise level. Exploiting the structural correlations of the basis material images and the minimum noise image, the denoised basis material images are calculated by applying local linear transformations to the processed minimum noise image. Figure 4.2 shows a graphical representation of the aforementioned joint dictionary denoising algorithm and summarizes the most important steps.

The size of the image patches plays an important role for the performance of the algorithm. On the one hand, larger patches lead to better dictionary denoising results and the fit coefficients for the linear transformations can be determined more reliably. On the other hand, the assumption of eq. 4.4 might not be fulfilled anymore if the patch size is chosen too large. Furthermore, the computational complexity increases drastically with increasing patch size. We chose a patch size of $8 \times 8 \times 8$ voxels for our experiments to balance the aforementioned tradeoffs. Additionally, the computational complexity of the algorithm can be reduced by increasing the sliding distance d_s of the image patches, since the total number of image patches is proportional to $1/d_s^3$. However, we found that the image quality decreases slightly with increasing sliding distance. We chose to use a sliding distance of two voxels for our experiments.

In order to reduce blocking artifacts [326], which become more prevalent with increasing sliding distance, we introduce a non-uniform weighting scheme for the contributions of different image patches to a certain voxel:

$$v^d = \frac{\sum_{p=1}^{P} w_p v_p^d}{\sum_{p=1}^{P} w_p}, \quad w_p = e^{-r_p^2/\delta}, \tag{4.8}$$

where v^d indicates the value of the voxel v after denoising. The summation index p includes all P image patches which contain the voxel v. The weights w_p decrease exponentially with the squared distance r_p^2 between the voxel v and the center of the image patch p. The quantity v_p^d denotes the value of voxel v in image patch p after denoising and δ is an adjustable parameter. With this weighting scheme, voxels close to the patch centers are given more weight. This reflects the idea that structures at the centers of the image patches can be detected more reliably than structures at the boundaries of the patches. Heuristically, we found that $\delta = 8.5$ gives the best results for a patch size of $8 \times 8 \times 8$ voxels and a sliding distance of two voxels. We use the FORBILD head phantom [337] (size 800^3 voxels) as training image. The final dictionary contains 2048 atoms with 512 ($=8^3$) voxels.

4.2.3 Bilateral Filtering

Bilateral filtering is a well known noise reducing and edge preserving post-processing method in the field of image processing. Similarly to our algorithm, an adapted version of the bilateral filter is able to exploit structural correlations between several aligned images. We therefore chose to compare our algorithm to this adapted version of the bilateral filter. Conventional bilateral filtering is a generalization of Gaussian smoothing and can be written as:

$$v_i = \frac{1}{N_i} \sum_{j \in N_i} v_j f(r_{ij}) e^{-(v_i - v_j)^2/2\sigma^2}, \tag{4.9}$$

where N_i is a geometrical neighborhood of pixel v_i and $f(r_{ij})$ is a distance dependent weighting factor (r_{ij} is the geometrical distance between the voxels v_i and v_j). The weighting factor $e^{-(v_i - v_j)^2/2\sigma^2}$ controls the degree of averaging in dependence of the difference between the voxel values. The idea is to suppress the averaging process across edges and encourage averaging if the voxel values are close to each other. To ensure this, the tuning parameter σ is normally chosen to be comparable to the image noise level. A generalization of this idea is to suppress the averaging process if an edge is detected in either of the two images a and b:

$$\begin{aligned} a_i &= \frac{1}{N_i} \sum_{j \in N_i} a_j f(r_{ij}) e^{-(a_i - a_j)^2/2\sigma_a^2} e^{-(b_i - b_j)^2/2\sigma_b^2}, \\ b_i &= \frac{1}{N_i} \sum_{j \in N_i} b_j f(r_{ij}) e^{-(a_i - a_j)^2/2\sigma_a^2} e^{-(b_i - b_j)^2/2\sigma_b^2}. \end{aligned} \tag{4.10}$$

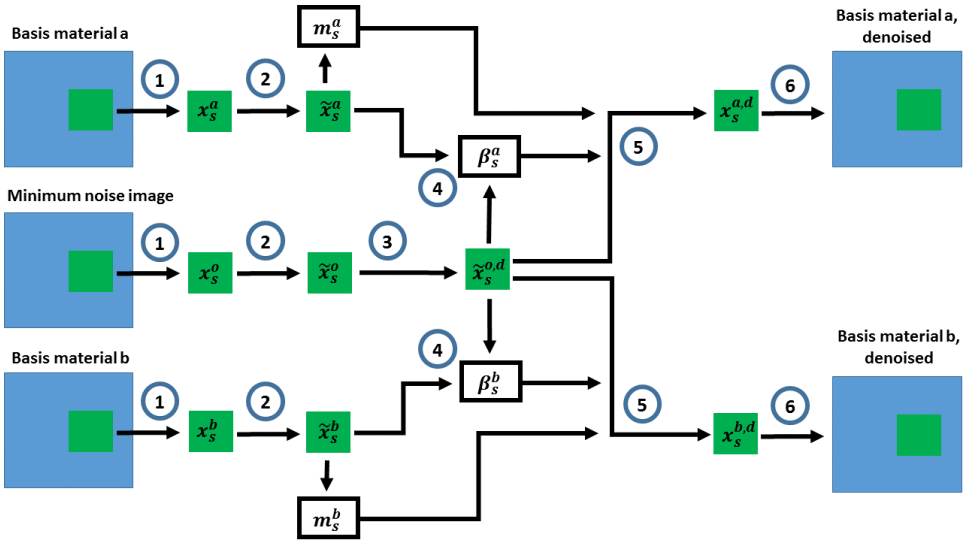

Figure 4.2: Graphical representation of the dictionary-based algorithm for denoising of basis material images, using one image patch as an example. The key steps of the algorithm are: (1) Extract image patches from the same locations in the basis material images and the minimum noise image. (2): Subtract the mean values of the image patches. (3): Apply dictionary denoising to the offset-corrected minimum noise patch. (4): Calculate β_s^a and β_s^b using eq. 4.6. (5): Linearly transform the processed minimum noise image patch to obtain the corresponding basis material image patches according to eq. 4.7. (6): The final basis material images are compiled using all the denoised basis material image patches, compare eq. 4.8.

In order to make bilateral filtering and dictionary denoising more comparable, we used a 3D version of the adapted bilateral filter with a geometrical neighborhood of the same size as the dictionary patches ($8 \times 8 \times 8$ voxels).

4.2.4 Numerical Simulation

A spectral CT scan of the FORBILD thorax phantom [337] (size 768^3 voxels) was simulated. We assumed acceleration voltages of 100 kVp and 140 kVp for the low and high energy scan, respectively. In both cases, the spectrum was filtered with 0.2 mm of copper and an ideal energy-integrating detector with a CsI-based scintillation layer was assumed (1 mm thickness). After reconstruction of the low and high energy images via filtered backprojection (FBP) an image-based material decomposition into a bone and soft tissue image was conducted using direct matrix inversion.

4.2.5 Experimental measurement

A CT scan of a human knee was conducted at an experimental setup. The usage of the knee specimen for research projects was approved by the institutional review board. The donor had dedicated its body for educational and research purposes, and provided written informed consent prior to death, in compliance with local institutional and legislative requirements. The ex-vivo human knee specimen was fixed in formalin. The tube was operated at an acceleration voltage of 110 kVp and a CdTe-based photon-counting detector (XC-Flite FX1, XCounter AB, pixel size 200 μm$\times$200 μm) with thresholds set to 27 and 52 keV was used. In total, 1201 projections were taken and the tube loading was 131 mAs. A projection-based material decomposition algorithm [249] was applied in order to obtain basis material images representing Compton scattering and photoelectric absorption.

4.3 Results

To investigate the performance of the new algorithm, we applied it as a post-processing method to basis material images obtained from projection-based as well as image-based material decomposition techniques. Furthermore, we compare the performance of our new joint dictionary denoising algorithm to adapted bilateral filtering as well as conventional dictionary denoising. Before turning to experimental measurements, we first present the results of a numerical simulation since this provides the possibility to compare the denoised images with a ground truth image.
The bottom row of figure 4.3 shows the simulated bone (d) and soft tissue image (e) as well as the minimum noise image (f). In the top row of figure 4.3 the ground truth values for the bone (a) and soft tissue image (b) as well as the minimum noise image (c) are displayed. The top row of figure 4.4 shows the results of applying joint dictionary denoising (a), conventional dictionary denoising (b) and bilateral filtering (c) to the bone image. Similarly, soft tissue images processed with our new algorithm (d), conventional dictionary denoising (e) and bilateral filtering (f) are displayed in the bottom row of figure 4.4. The tuning parameters of the algorithms (ϵ for dictionary-based denoising and σ for bilateral filtering) were optimized by maximizing the structural similarity index (SSI), using the ground truth images as references. Since only one parameter was tuned for all denoising algorithms, this approach is not prone to overfitting. In table 4.1, the mean squared error (MSE) and SSI compared to the ground truth are given for various denoising methods.
Figure 4.5 compares the influence of the dictionary patch size and the sliding distance on the image quality for joint dictionary denoising. The same parameter ϵ as in figure 4.4 (scaled by the dictionary patch size) was used for all experiments.
Figure 4.6 shows images of a human knee in transverse slice orientation at the level of the patellofemoral joint. Figures 4.6 (a) and (b) show the basis images representing Compton scattering and photoelectric absorption, respectively. In figure 4.6 (c), the corresponding minimum noise image is displayed. Figure 4.7 compares the results of applying different denoising methods (bilateral filtering, conventional and joint dictionary denoising) to the Compton and photo images displayed in figure 4.6 (a) and (b). Since there was no reference image available, the denoising parameters were visually tuned to achieve a comparable edge sharpness for all methods. The following denoising parameters were used: Joint dictionary denoising: $\epsilon = 0.1807\ \mathrm{cm}^{-2}$, conventional dictionary denoising: $\epsilon_{\mathrm{Compton}} = 0.3315\ \mathrm{cm}^{-2}$, $\epsilon_{\mathrm{photo}} = 2.885 \times 10^{10}\ \mathrm{keV}^6/\mathrm{cm}^2$, bilateral filtering: $\sigma^2_{\mathrm{Compton}} = 1.004 \times 10^{-3}\ \mathrm{cm}^{-2}$, $\sigma^2_{\mathrm{photo}} = 8.084 \times 10^{7}\ \mathrm{keV}^6/\mathrm{cm}^2$.

4.4 Discussion

For the numerical simulation, material decomposition via matrix inversion leads to a strong degradation of the SNR and highly anti-correlated noise in the material selective images. Therefore, the noise level can be greatly reduced by calculating the virtual monochromatic image at which the anti-correlated noise cancels out maximally (see figure 4.3 (f)). Joint dictionary denoising leads to basis material images with strongly improved image quality compared to the unprocessed images. The processed images look similar to the ground truth images. The most apparent differences are slightly blurred edges and the presence of a small amount of low frequency noise in the denoised images. Numerical experiments show that this low frequency noise is mostly caused by the uncertainties in determining the correct mean values (m_s^a, m_s^b) of the image patches. However, since the variation (due to noise) of the mean value of an image patch is much smaller than the variation of an individual voxel, the noise level is reduced compared to the unprocessed images. In order to achieve a similar noise reduction with conventional dictionary denoising, spatial resolution and edge sharpness have to be sacrificed (compare the zoomed region of figure 4.4 (b) and (e)). The loss of edge sharpness only occurs for the lower contrast

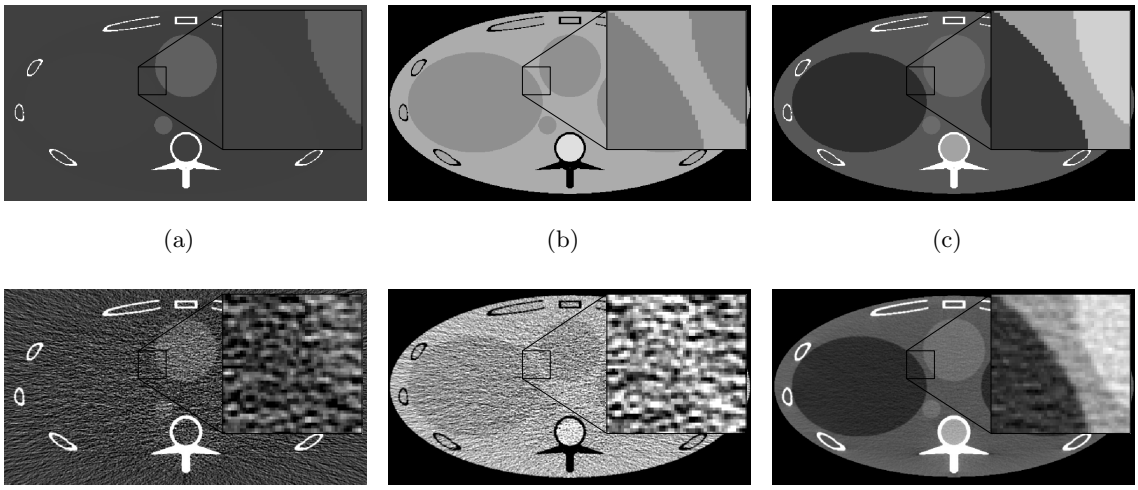

(a) (b) (c)

(d) (e) (f)

Figure 4.3: Bone (a, d), soft tissue (b, e) and minimum noise (c, f) images for the numerical simulation. The top row shows the ground truth while the bottom row shows the noisy images obtained from a simulated low-dose scan of the phantom. The range of the windows is [-0.19 g/cm^3, 0.58 g/cm^3], [0.42 g/cm^3, 1.38 g/cm^3] and [0.19 cm^{-1}, 0.33 cm^{-1}] for the bone, soft tissue and minimum noise images, respectively.

Table 4.1: Mean squared error (MSE) and structural similarity index (SSI) for the bone and soft tissue image compared to the ground truth for various denoising methods. The MSE was normalized with respect to the unprocessed basis material images.

	No denoising	Bilateral filtering	Conventional dictionary denoising	Joint dictionary denoising
MSE, bone image	1	0.746	0.337	0.296
MSE, soft tissue image	1	0.414	0.222	0.184
SSI, bone image	0.332	0.888	0.946	0.953
SSI, soft tissue image	0.249	0.807	0.906	0.942

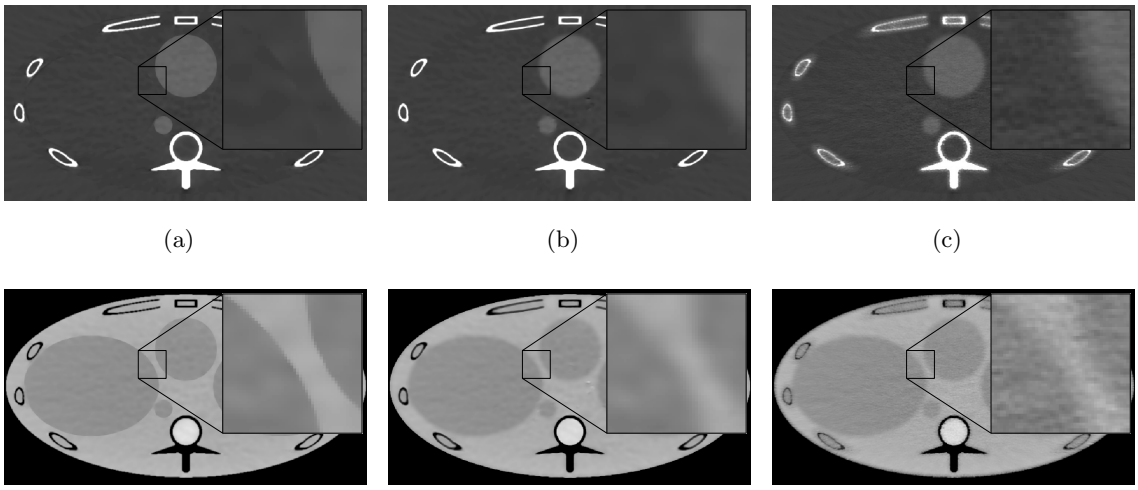

(a) (b) (c)

(d) (e) (f)

Figure 4.4: Comparison of different denoising methods for the numerical simulation. The bone and soft tissue images are shown in the top and bottom row, respectively. The columns of the figure represent different denoising methods: Joint dictionary denoising (a, d), conventional dictionary denoising (b, e) and bilateral filtering (c, f). The range of the windows is [-0.19 g/cm^3, 0.58 g/cm^3] and [0.42 g/cm^3, 1.38 g/cm^3] for the bone and soft tissue images, respectively.

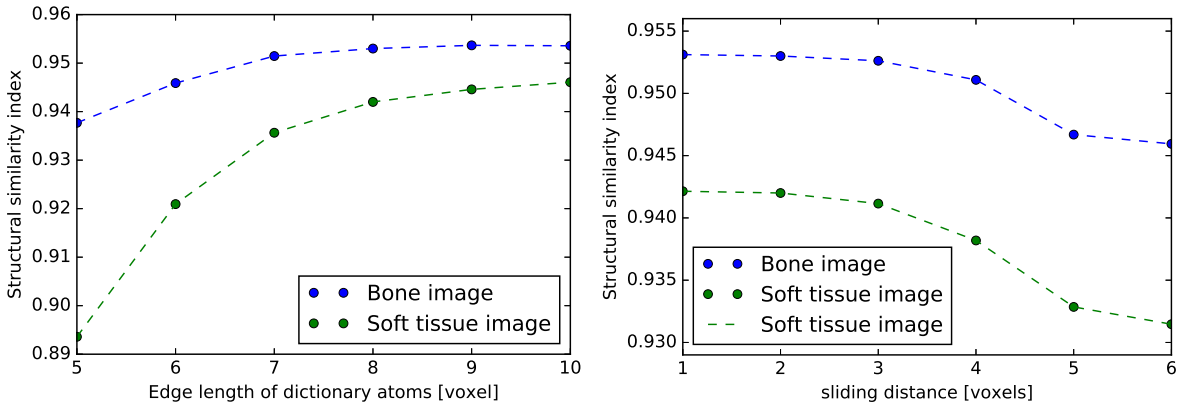

(a) (b)

Figure 4.5: Influence of the dictionary patch size (a) and the sliding distance (b) on the image quality for joint dictionary denoising. The image quality was measured by calculating the structural similarity compared to the ground truth.

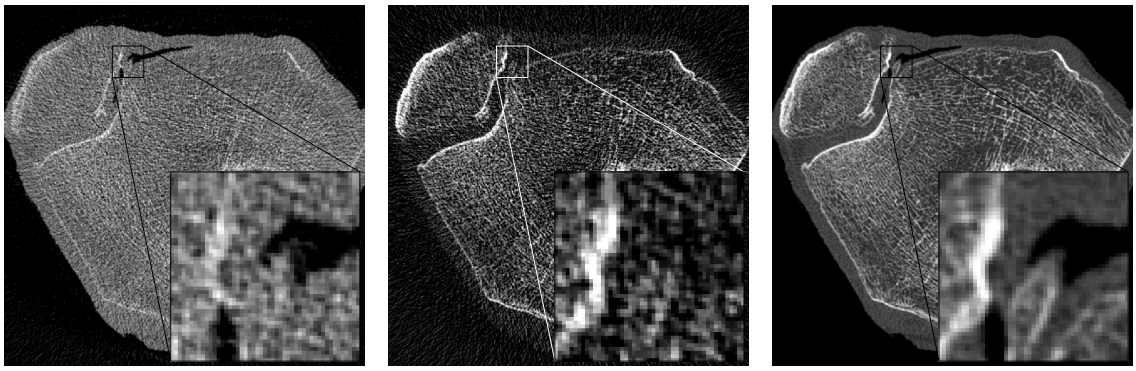

(a) (b) (c)

Figure 4.6: Unprocessed Compton (a), photo (b) and minimum noise (c) image for the experimental measurement of a human knee. The range of the windows is [0.0 keV^3/cm, 68400 keV^3/cm], [0.035 cm^{-1}, 0.35 cm^{-1}] and [0.09 cm^{-1}, 0.59 cm^{-1}] for the photo, Compton and minimum noise image, respectively.

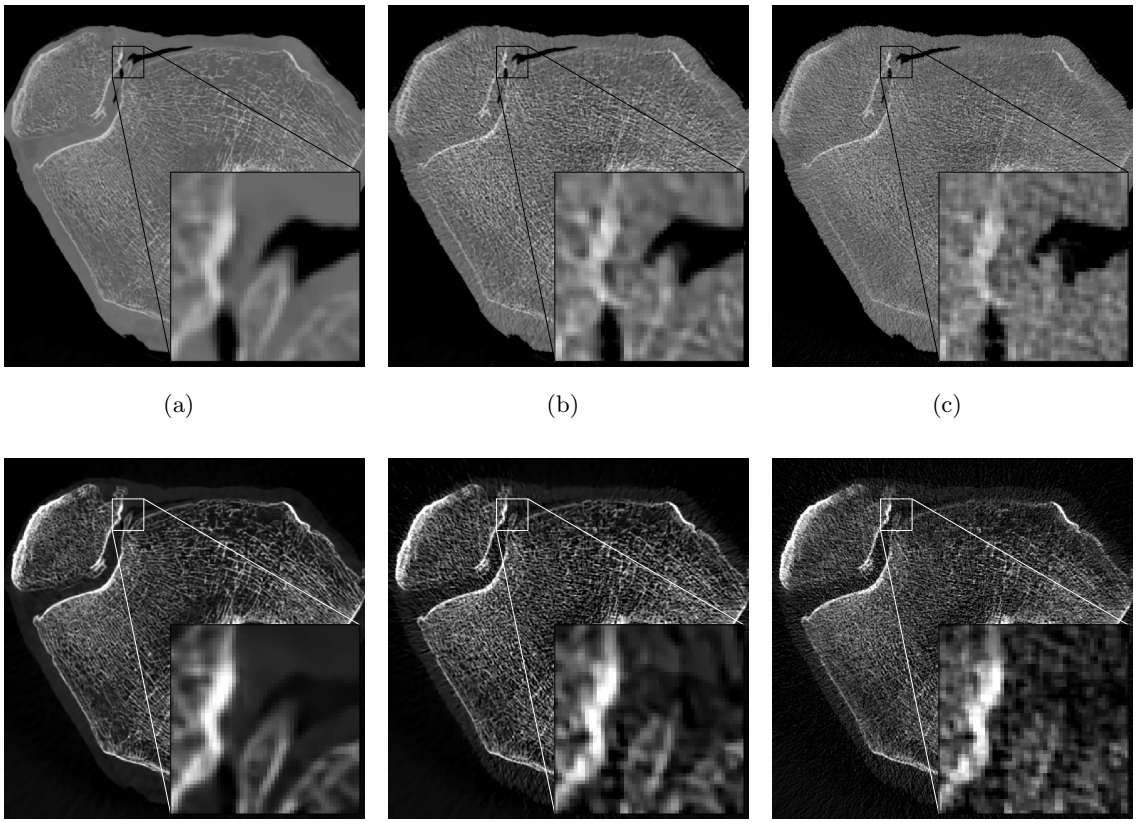

(a) (b) (c)

(d) (e) (f)

Figure 4.7: Comparison of bilateral filtering (c,f), joint (a,d) and conventional (b,e) dictionary denoising for the experimental measurement of a human knee. The top row shows Compton images whereas the bottom row shows photo images. The range of the windows is [0.0 keV^3/cm, 68400 keV^3/cm] and [0.035 cm^{-1}, 0.35 cm^{-1}] for the photo and Compton images, respectively.

features in the soft tissue region. Conventional dictionary denoising is still able to accurately distinguish the edges between soft tissue and bone from noise. In the case of bilateral filtering, reduced edge sharpness in the soft tissue region and an increased noise level compared to conventional dictionary denoising can be observed. The aforementioned qualitative statements are supported by quantitative image quality measurements. Measured in terms of MSE and SSI, the images produced by our algorithm are notably closer to the ground truth compared to bilateral filtering and conventional dictionary denoising.
Depending on the experimental parameters, one could imagine a scenario where structures of the basis material images cancel (almost) completely in the minimum noise image. In this case, structures and edges could get lost in the basis material images because they are compiled from linearly transformed minimum noise image patches. This is a limitation of the joint dictionary denoising algorithm in its current form. A possible extension of the algorithm could compare several virtual monochromatic images and locally choose the one with the best SNR. However, this would require a method to locally estimate the image signal.
As figure 4.5 shows, the image quality for both basis material images becomes worse if the dictionary patch size is reduced. This effect is more pronounced for the soft tissue image. Since the soft tissue image has a higher noise level, the uncertainties in determining the fit parameters for the linear transformations (compare eq. 4.4) grow faster with decreasing patch size. The image quality also deteriorates slightly with increasing sliding distance, especially if the sliding distance is larger than three voxels. Up to a sliding distance of three voxels, the image quality is only marginally reduced. It is therefore reasonable to use a sliding distance of two or three voxels to save computational time.
The goal of the experimental measurement was to demonstrate that our algorithm achieves a strong improvement in image quality in case of a clinically relevant image with complicated structures. Compared to the minimum noise image (figure 4.6 (c)), the basis material images (figure 4.6 (a) and (b)) show a decreased SNR and the noise is anti-correlated. In contrary to the numerical simulation, where maximizing the SSI leads to a comparable noise level for all denoising methods, the denoising parameters for the experimental measurements were visually tuned to achieve comparable edge sharpness for all denoising methods. As can be seen from figure 4.7 (a) and (d), joint dictionary denoising efficiently removes noise from the basis material images, while fine structures and features (visible in the minimum noise image) can be clearly identified. In the unprocessed basis material images, most of these structures vanish in the noise. Consequently, the image quality of the basis material images can be greatly improved by applying joint dictionary denoising. Bilateral filtering and conventional dictionary denoising leads to higher noise levels and some image features are lost (see for example the top-right region of the images in figure 4.7).

4.5 Conclusion

We have developed a new method for denoising basis material images in spectral CT. As a post-processing method, it can be used for image-based as well as projection-based material decomposition techniques. The algorithm is based on the capability of learned dictionaries to preserve image features while suppressing noise. We have introduced several modifications of the conventional dictionary denoising algorithm in order to exploit the structural correlations of basis material images as well as the anti-correlated noise. Dictionary denoising is applied to the virtual monochromatic image at which the anti-correlated noise maximally cancels out. Suppressing noise and identifying image features is in general much more efficient and reliable for this minimum noise image compared to performing the same task directly on the basis material images. An exception to this occurs if structures cancel out in the minimum noise image, which is a potential limitation of the algorithm. The denoised basis material images are subsequently calculated by applying linear transformations [338] to the processed minimum

noise image patches. We demonstrated that post-processing basis material images with the proposed algorithm leads to highly improved image quality. Noise is strongly suppressed while almost no blurring of edges and structures occurs. Furthermore, our joint dictionary denoising algorithm leads to superior image quality compared to conventional dictionary denoising and bilateral filtering. Improving the image quality of basis material images is an important goal because noise amplification on material decomposed images is a fundamental problem of spectral CT. This applies in particular to medical imaging where the desire for low radiation exposure of the patient leads to increased noise levels. Our algorithm therefore has the potential to improve the usability of basis material images for various tasks such as artifact reduction, quantitative imaging and clinical diagnosis.

5 Spectral angiography material decomposition using an empirical forward model and a dictionary-based regularization

In this chapter, we present a novel algorithm for projection-based material decomposition. It features an empirical polynomial forward model that is tuned by calibration measurements. An existing direct estimator [244] (compare eq. 2.159) is reformulated as MAP estimator for the inverse problem of determining the basis material line integrals. The modified estimator has the advantage of achieving the CRLB (see section 2.5.2) if the number of spectral measurements exceeds the number of basis materials. Moreover, the applied Bayesian approach of MAP estimation allows to incorporate a noise-suppressing regularization term into the material decomposition algorithm. In contrast to the post-processing method discussed in the last chapter, material decomposition and image denoising are performed jointly. We devise an efficient regularization strategy by generalizing the Bayesian formulation of dictionary denoising to simultaneously consider several correlated image channels. The commonly used nearest-neighbor-based regularization strategies, such as the Huber or TV penalty, assume a piece-wise constant image. In many cases, this assumption is justified for 3D CT images, however, typically it does not hold for projection images. Furthermore, dictionary-based regularizers mitigate the problem of regularization crosstalk, which will be discussed in more detail in chapter 6. In many regards, the novel regularization strategy can be viewed as an advancement of the dictionary-based post-processing algorithm that was discussed in the previous chapter. The post-processing algorithm relies on the minimum noise image for identifying edges and structures. This potentially carries the risk of missing structures in the basis material images that cancel each other in the minimum noise image. The dictionary-based regularization strategy that will be discussed in this chapter combines local whitening transformations with the simultaneous orthogonal matching pursuit (SOMP) algorithm to eliminates this disadvantage. In addition, the local whitening transformations offer an elegant way of considering a spatially varying noise level in all image channels. We focused on spectral coronary angiography as a potential clinical application of projection-based material decomposition with photon counting detectors. Numerical and real experiments show that spectral angiography with realistic dose levels and gadolinium contrast agent concentrations is feasible using the proposed decomposition algorithm and currently available photon-counting detector technology.
The text and the figures in this chapter are taken from the following publication of the author:

- Mechlem, Korbinian, et al. "Spectral angiography material decomposition using an empirical forward model and a dictionary-based regularization." IEEE Transactions on Medical Imaging 37.10 (2018): 2298-2309.

Compared to the original publication, the text has been slightly modified to achieve a better integration into this doctoral thesis.

We have also investigated an extension of the proposed decomposition algorithm to spectral CT. With minor modifications, the presented empirical forward model can be incorporated into a one-step SIR algorithm that performs material decomposition and image reconstruction simultaneously. In this case, the one-step SIR algorithm was combined with a nearest-neighbor based regularization term. The key findings coincide with the results that will be presented in chapter 6, where a different empirical forward model is combined with a one-step SIR algorithm.

For this reason, extending the empirical decomposition algorithm that was presented in this chapter to one-step SIR will not be discussed in detail in this work. More information can be found in the following publication of the author:

- Mechlem, Korbinian, et al. "Statistical iterative reconstruction for spectral CT using ratios of polynomial functions." 14th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine (2017).

Generalizing the dictionary-based regularization approach to 3D and combining it with one-step SIR algorithms would be an interesting future project. Possible advantages compared to nearest-neighbor based regularizers include improved image quality and reduced regularization crosstalk.

5.1 Introduction

In the last decades, invasive coronary angiography (ICA) has been the gold standard for the imaging of coronary arteries and the diagnosis of coronary heart disease. The examination is usually performed with a C-arm unit, while contrast agent is administered to the coronary vessels by a catheter. Despite the method's indisputable clinical worth, the overlapping of coronary vessels and the surrounding soft tissue and bone structures results in reduced image contrast and poor visibility of smaller vessels. For cardiac [339] imaging, it was demonstrated that this "anatomical noise" reduces the probability of detecting stenotic lesions. Digital subtraction angiography (DSA) aims to overcome the aforementioned limitation by acquiring two projection images before and after contrast agent injection. The difference image removes the anatomical background and thus allows for the visualization of the coronary vessels without interfering anatomical structures. However, motion artifacts are a major problem for DSA. Cardiac and respiratory movements as well as other involuntary motions can compromise the image quality of DSA [340]. Furthermore, the chemical composition of coronary artery calcifications cannot be assessed directly in both ICA and DSA. Driven by recent advances in photon-counting detector (PCD) technology [79, 80, 81], spectral imaging has obtained increasing attention concerning its use in medical diagnostics. The ability to (partially) resolve the incident photon energy in combination with the spatial and temporal registration of the acquired spectral data allows for projection-based material decomposition. Projection-based material decomposition algorithms exploit the additional information about the energy-dependent attenuation to generate basis material images. Using contrast agents with K-edge discontinuities in the relevant energy range, three and more basis materials can be differentiated [28, 341, 342, 343, 344, 345]. Spectral coronary angiography could be a promising advancement of ICA, as the ability to differentiate between bone, soft tissue and contrast agent allows for the generation of DSA-like images without the risk of motion artifacts. Furthermore, it is possible to distinguish between calcified and non-calcified coronary plaques, as the latter are associated with certain risk factors [346]. However, noise amplification is particularly problematic for multi-material decomposition tasks, such as the aforementioned decomposition into bone, soft tissue and contrast agent. The reduced SNR limits the diagnostic value of basis material images, especially for low-dose medical imaging. Although many sophisticated denoising algorithms exploiting the noise-correlation between material-specific images have been developed for spectral CT [229, 347, 232], there are only a few publications considering projection data [235, 348, 349].
Another major challenge for material decomposition is constructing an accurate model for the measured spectral signals as a function of the basis material thicknesses. In this chapter, we present a new empirical forward model based on rational functions that is tuned by calibration measurements. Contrary to the existing decomposition algorithm based on rational functions [244], the proposed algorithm achieves the theoretical optimum for the variance even if the number of spectral measurements exceeds the number of basis materials. In order to suppress noise

introduced in the decomposition and improve image quality, we combine a statistical model of the measurement process with a dictionary-based regularization in a Bayesian framework. We generalize the dictionary-based regularization approach to simultaneously consider several structurally correlated basis material images. The performance of the presented decomposition algorithm and its merit for applications in spectral coronary angiography are evaluated in computer simulations as well as experimental measurements. In both cases, we investigated a thorax phantom with coronary vessels filled by a gadolinium contrast agent and added artificial coronary calcifications.

5.2 Methods

5.2.1 Empirical model for material decomposition

The empirical model for material decomposition is based on a publication by Cardinal and Fenster [244]. In the context of dual energy CT, they demonstrated that the measured log-signals are well approximated by polynomial functions of the basis material line integrals:

$$l_i^l = P(A_i^0, A_i^1; \vec{c}_l), \quad l_i^h = P(A_i^0, A_i^1; \vec{c}_h), \tag{5.1}$$

where A_i^k is the basis material path length for basis material k (here $k \in (0,1)$) and detector pixel i and the vectors $\vec{c}_l$ and $\vec{c}_h$ describe a set of fit parameters. The log-signals are defined as:

$$l_i^l = -\ln(\hat{y}_i^l / b_i^l), \quad l_i^h = -\ln(\hat{y}_i^h / b_i^h), \tag{5.2}$$

where $\hat{y}_i^l$ and $\hat{y}_i^h$ denote the expected intensities for the low and high-energy measurement, respectively. Similarly, b_i^l and b_i^h describe the reference intensities without the sample in the beam path for the low and high-energy measurement. Quantities with a hat denote an expected value, whereas the same quantity without a hat represents a measured value. For simplicity, the reference intensities are assumed to be noise-free, i.e. $\hat{b} = b$. The functions $P(A_i^0, A_i^1; \vec{c}_l)$ and $P(A_i^0, A_i^1; \vec{c}_h)$ are given by the ratio of two low-order polynomials:

$$\begin{aligned} &P(A_i^0, A_i^1; \vec{c}_l) = \mathcal{N}(A_i^0, A_i^1; \vec{c}_l) / \mathcal{D}(A_i^0, A_i^1; \vec{c}_l) \\ &\mathcal{N} = c_{l0} + c_{l1} A_i^0 + c_{l2} A_i^1 + c_{l3} \left(A_i^0\right)^2 + c_{l4} \left(A_i^1\right)^2 + c_{l5} A_i^0 A_i^1 \\ &\mathcal{D}(A_i^0, A_i^1; \vec{c}_l) = 1 + c_{l6} A_i^0 + c_{l7} A_i^1 \end{aligned} \tag{5.3}$$

with a similar equation for $P(A_i^0, A_i^1; \vec{c}_h)$. The fit coefficients $\vec{c}_l$ and $\vec{c}_h$ depend on the source spectrum and the detector response. Modeling the log-signals with the aforementioned rational functions is motivated by comparing their Taylor expansions with the Taylor expansion of the analytically calculated log signal function [350] (i.e. using the source spectrum and the detector response). Cardinal and Fenster inverted the relation given in eq. 5.3 to derive a dual energy material decomposition algorithm:

$$A_i^0 = \tilde{P}(l_i^l, l_i^h; \vec{d}_0), \quad A_i^1 = \tilde{P}(l_i^l, l_i^h; \vec{d}_1). \tag{5.4}$$

After determining the parameters $\vec{d}_0$ and $\vec{d}_1$ for the rational functions $\tilde{P}(l_i^l, l_i^h; \vec{d}_0)$ and $\tilde{P}(l_i^l, l_i^h; \vec{d}_1)$ with the help of calibration measurements, the basis material line integrals are obtained by simply evaluating the calibrated rational functions at the measured low and high-energy log-signals. This material decomposition approach has many desirable properties such as computational efficiency, robustness to noise, and high accuracy. Furthermore, a small number of calibration measurements (9-16) is sufficient to robustly determine the fit coefficients $\vec{d}$ and the algorithm has good extrapolation capabilities beyond the calibration range. However, it was shown that a generalization of this decomposition algorithm to the case of having more spectral measurements than basis materials leads to excessive noise [225]. The obtained variances are up to two orders

of magnitude above the Cramer Rao lower bound (CRLB), which describes the theoretical lower limit for the variance of an unbiased estimator. The increased variance is a large disadvantage for material decomposition with multi-bin PCDs, which allow for spectral measurements with more energy bins than basis materials. In general, increasing the number of spectral measurements leads to reduced noise levels for the basis material images, as the additional spectral information can be exploited by the decomposition algorithm [240]. To overcome this limitation of the decomposition algorithm, we propose using the non-inverted relationship between the log-signals (see eq. 5.1) to obtain an empirical forward model for the measured photon counts. By modeling the statistics of the spectral measurements, the empirical model can be combined with a maximum-likelihood (ML) estimator that is asymptotically (in the limit of large photon counts) unbiased and achieves the CRLB. In the following, the proposed modifications will be explained in more detail. The expected number of photon counts $\hat{y}_i^s$ for detector pixel i and spectral measurement index s is given by:

$$\hat{y}_i^s = b_i^s e^{-\hat{l}_i^s},\ \hat{l}_i^s = P_i^s(\vec{A}_i;\vec{c}_{is}),\ \vec{A}_i = (A_i^0,...,A_i^n)^T \tag{5.5}$$

where b_i^s denotes the reference intensity for detector pixel i and spectral index s. Compared to the original model (cf. eq. 5.1), $P_i^s(\vec{A}_i;\vec{c}_{is})$ is a function of n different basis material line integrals, and the fit parameters $\vec{c}_{is}$ are determined individually for each spectral measurement and each pixel. Because of the manufacturing imperfections of current PCDs, the energy response varies between different pixels. Thus, a global fit (averaging all pixels) would lead to systematic errors for the decomposed basis material thicknesses. For a three-material decomposition task, which is the focus of this paper, $P(\vec{A}_i) = P(A_i^0, A_i^1, A_i^2)$ is given by:

$$\begin{aligned} &P(A_i^0, A_i^1, A_i^2) = \mathcal{N}(A_i^0, A_i^1, A_i^2)/\mathcal{D}(A_i^0, A_i^1, A_i^2), \\ &\mathcal{N} = c_0 + \sum_{r=1}^{3} c_r A_i^r + \sum_{r=1}^{3} c_{(3+r)}(A_i^r)^2 + c_7 A_i^0 A_i^1 + c_8 A_i^0 A_i^2 + c_9 A_i^1 A_i^2 \\ &\mathcal{D} = 1 + \sum_{r=1}^{3} c_{(9+r)} A_i^r \end{aligned} \tag{5.6}$$

To simplify the index notation, the indices i and s for the fit coefficients $\vec{c}_{is}$ have been dropped in eq. 5.6. The fit parameters $\vec{c}_{is}$ are determined by a least-squares fit to the calibration measurements:

$$\vec{c}_{is} = \arg\min \sum_{k=1}^{K} w_k \Big(l_{ik}^s - P\left(\vec{A}_{ik};\vec{c}_{is}\right)\Big)^2. \tag{5.7}$$

The measured log-signals for calibration measurement k are denoted by l_{ik}^s, and $\vec{A}_{ik}$ represents the corresponding basis material thicknesses. The weights w_k are chosen according to the statistical uncertainty of the corresponding calibration point. We observed that inconsistencies in the calibration measurements (e.g. noise in simulations or detector drift in an experimental measurement) could sometimes lead to outliers for the fit coefficients (e.g. extremely large values for one fit coefficient). To avoid numerical problems during the subsequent decomposition, the fit coefficients $\vec{c}_{is}$ were constrained to predefined, heuristically chosen intervals. We determined the intervals $[c_{rs}^{\min}, c_{rs}^{\max}]$ (r being the fit coefficient index and s the spectral measurement index) by first performing an unconstrained fit for all detector pixels and then choosing $c_{rs}^{\min}$ and $c_{rs}^{\max}$ such that 99% of the unconstrained fit coefficients lie within the interval $[c_{rs}^{\min}, c_{rs}^{\max}]$. Based on the assumption that the energy-dependent attenuation of all basis materials monotonically decreases, Cardinal and Fenster showed that all fit coefficients should be positive [350]. As this assumption is not valid for K-edge materials, we allowed for negative fit coefficients. After the empirical forward model has been calibrated (cf. eq. 5.5 and eq. 5.7), ML decomposition of a spectral measurement is performed by finding the set of basis material line integrals $\vec{A}_i$

that maximizes the probability of the spectral measurements. Assuming independent Poisson statistics, this corresponds to minimizing the following negative log-likelihood function for each pixel:

$$\vec{\mathcal{A}}_i = \arg\min - \mathcal{L}_i(\vec{A}_i) = \sum_{s=1}^{S} \hat{y}_i^s(\vec{A}_i) - y_i^s \ln(\hat{y}_i^s(\vec{A}_i)). \tag{5.8}$$

In summary, the modified algorithm models the log-signals as a function of the basis material thicknesses instead of modeling the basis material thicknesses as a function of the log-signals. Equation 5.5, allows one to express the expected number of photon counts as a function of the basis material line integrals. This allows for a statistical treatment of the measured photon counts y_i^s in terms of a maximum-likelihood estimator. However, the modifications increase the computational complexity, as an iterative optimization algorithm is needed to find the minimum of $-\mathcal{L}_i(\vec{A}_i)$.

5.2.2 Regularized material decomposition

Since material decomposition suffers from noise amplification and the degradation of the SNR compared to the unprocessed images, the direct output of a material decomposition algorithm is often unsatisfactory in terms of image quality. In order to keep the radiation dose to the patient as low as possible, it is necessary to combine the material decomposition algorithm with an efficient denoising strategy. For conventional CT, the Bayesian approach of combining a statistical model of the projection measurements with prior knowledge in the form of a regularization term has been shown to strongly improve the image quality and reduce the noise levels compared to analytical reconstruction methods [7]. While this approach has also been successfully adapted to CT reconstruction of basis material images [351, 352, 353], to the best of our knowledge, the potential of regularization strategies for material decomposition in the projection domain was only recently investigated [235]. Most regularization approaches for material decomposition penalize the difference between neighboring pixels (or voxels) individually for each basis material image (e.g. using a quadratic or Huber penalty). An entirely different regularization strategy that has gained increasing interest in the last years is based on a Bayesian formulation of dictionary denoising [331, 354, 355]. This approach assumes that small image patches can be modeled as a sparse linear combination of dictionary patches. The dictionary patches are pre-trained on high-quality images and reflect typical structures occurring in these images. Consequently, noise cannot be sparsely represented in the dictionary basis. It has been demonstrated that a dictionary-based regularization can outperform total variation (TV) regularization [331, 354]. Moreover, edge-preserving nearest-neighbor-based regularization strategies, such as the Huber or TV penalty, assume a piece-wise constant image (or equivalently a sparse image gradient). In many cases, this assumption is justified for 3D CT images, however, typically it does not hold for projection images. We therefore chose to combine the basis material decomposition algorithm with a dictionary-based regularization term. We introduce several modifications to the dictionary denoising algorithm in order to tailor it for the simultaneous denoising of several correlated imaging channels (in this case, basis material images). The following paragraphs explain the mathematical formulation of dictionary denoising as well as the aforementioned modifications.

Dictionary denoising

Dictionary denoising operates on small overlapping quadratic image patches that are processed individually. The degree of overlap is described by the sliding distance, which is defined as the distance between the centers of adjacent image patches. An image patch of $\sqrt{N} \times \sqrt{N}$ pixels can be described as a vector $\vec{x} \in \mathbb{R}^N$. The dictionary is represented by a matrix $\boldsymbol{D} \in \mathbb{R}^{N \times K}$ whose columns are the dictionary patches. In general, an over-complete dictionary basis is chosen, i.e. $K >> N$ because this allows for a more sparse representation of the image patches. The task

of finding a suitable linear combination of dictionary patches to model an image patch $\vec{x}$ can be written as a constrained optimization problem:

$$\min_{\vec{\alpha}} ||\vec{\alpha}||_0 \text{ s.t. } ||\vec{x} - \boldsymbol{D}\vec{\alpha}||_2^2 \leq \epsilon, \tag{5.9}$$

where $\vec{\alpha} \in \mathbb{R}^K$ is a vector containing the coefficients for the linear combination and $\epsilon > 0$ is a small error tolerance. The l_0-norm $||\bullet||_0$ counts the number of nonzero entries in the vector $\vec{\alpha}$. Equation 5.9 is interpreted as representing an image patch $\vec{x}$ as a linear combination of dictionary patches with a maximum error of ϵ using as few dictionary patches as possible. Because solving this optimization problem is NP-hard, greedy algorithms are employed to compute an approximate solution in acceptable time. A popular choice is the orthogonal matching pursuit algorithm (OMP) [335]. With the help of the method of Lagrange multipliers, eq. 5.9 can be rewritten in an unconstrained form:

$$\min_{\vec{\alpha}} ||\vec{x} - \boldsymbol{D}\vec{\alpha}||_2^2 + v||\vec{\alpha}||_0, \tag{5.10}$$

where v is the Lagrange multiplier. In our implementation, the dictionary $\boldsymbol{D}$ is trained using the online dictionary learning method [336] and is kept fixed during the optimization ("global dictionary"). The alternative is to use an adaptive dictionary that is dynamically re-trained on the current image estimates [331, 355]. Since our goal is to develop a dictionary-based prior in a Bayesian framework, using an adaptive dictionary would violate the fundamental principle of a data-independent prior.

Multi-channel generalization

A straightforward generalization of eq. 5.9 to several imaging channels is to apply dictionary denoising individually to each image. In order to write the resulting optimization problem compactly, the vectors $\vec{x}$ and $\vec{\alpha}$ are generalized to matrices $\boldsymbol{x} \in \mathbb{R}^{N\times C}$, $\boldsymbol{\alpha} \in \mathbb{R}^{K\times C}$, where C is the number of channels. The matrices $\boldsymbol{x}$ and $\boldsymbol{\alpha}$ can be written as $\boldsymbol{x} = (\vec{x}_0, ..., \vec{x}_C)$ and $\boldsymbol{\alpha} = (\vec{\alpha}_0, ..., \vec{\alpha}_C)$, i.e. each column of $\boldsymbol{x}$ (or $\boldsymbol{\alpha}$) contains the image patch (or coefficient vector) for one specific channel. With these definitions, the unconstrained optimization problem can be written as:

$$\min_{\boldsymbol{\alpha}} ||\boldsymbol{x} - \boldsymbol{D}\boldsymbol{\alpha}||_F^2 + v||\boldsymbol{\alpha}||_0, \tag{5.11}$$

where $||\bullet||_F$ is the Frobenius norm. Similarly, the l_0-'norm' is generalized to counting the number of nonzero elements of the matrix $\boldsymbol{\alpha}$. By treating each image channel independently from the others, the formulation of the optimization problem in eq. 5.11 does not account for the strong correlations between the channels. In our case, the basis material images often share the same edges and structures and thus the image patches for different channels can be modeled by using the same set of dictionary patches (but different coefficient vectors). To exploit structural correlations between different channels and perform joint dictionary denoising, the sparsity assumption is modified. Instead of assuming that the image patches for each channel can be modeled by a small number of dictionary patches, it is assumed that image patches from different channels can be modeled by a common small set of dictionary patches. Mathematically, this assumption is expressed by penalizing the number of nonzero rows of the matrix $\boldsymbol{\alpha}$ instead of penalizing the number of nonzero elements. The optimization problem is therefore rewritten as:

$$\min_{\boldsymbol{\alpha}} ||\boldsymbol{x} - \boldsymbol{D}\boldsymbol{\alpha}||_F^2 + v||\boldsymbol{\alpha}||_{\text{row}-0}, \tag{5.12}$$

with the row-l_0 norm:

$$||\boldsymbol{\alpha}||_{\text{row}-0} = ||\vec{\beta}||_0, \ \beta_i = \max_j (|\boldsymbol{\alpha}_{ij}|) \tag{5.13}$$

The new optimization problem (eq. 5.12) can be solved by the simultaneous orthogonal matching pursuit algorithm (SOMP) [356], which is a modified version of the OMP algorithm. Theoretical [357] as well as experimental studies [358] show that the SOMP algorithm improves the reliability

of signal recovery compared to the application of the OMP algorithm separately to each channel. We used the implementation of the SOMP algorithm in the "SPAMS" software package [359] for our experiments.

Whitening transformation

The SOMP algorithm works best if the noise is independent and identically distributed for all channels, since all channels are treated equally in the corresponding optimization problem (see eq. 5.12). However, basis material decomposition typically leads to noise correlations and different noise levels for the basis material images. A well-known example is the highly anti-correlated noise structure for dual energy basis material decomposition into photo and Compton images. The implicit assumption of a strong correlation between different channels in the SOMP algorithm can result in correlated noise being modeled by dictionary patches. Separating structures and noise works best for correlated structures and uncorrelated noise. We therefore apply a basis transformation to the image patches before applying the SOMP algorithm. In the new basis, the noise is uncorrelated and has the same magnitude in all channels. Given an estimate of the covariance between different channels (with covariance matrix $\boldsymbol{R}$), Alvarez [360] demonstrated that the corresponding transformation matrix $\boldsymbol{\varphi}$ has the form:

$$\boldsymbol{\varphi} = \left[\frac{\vec{\varphi}_1}{\sqrt{\eta_1}} \; ... \; \frac{\vec{\varphi}_C}{\sqrt{\eta_C}} \right]^T, \tag{5.14}$$

where $\vec{\varphi}_1, ..., \vec{\varphi}_C$ and $\eta_1, ..., \eta_C$ are the eigenvectors and eigenvalues of $\boldsymbol{R}$, respectively.

Penalized log-likelihood function

Having introduced maximum-likelihood basis material decomposition (section 5.2.1) and the concepts for joint dictionary denoising, we are now in the position to compile the penalized log-likelihood function for regularized material decomposition. Similar to the single-channel strategy for conventional CT, we assume that image patches that are poorly modeled by a sparse linear combination of dictionary patches have a low prior probability. Contrary to maximum-likelihood decomposition, the optimization problem for regularized decomposition is not separable (with respect to different pixels) because the dictionary regularization term couples neighboring pixels. We therefore define the matrix $\boldsymbol{A} \in \mathbb{R}^{N_t \times C}$, which holds all basis material images, with N_t representing the total number of pixels in a basis material image. Similar to section 5.2.1, $\vec{A}_i$ represents the different basis material thicknesses for one specific pixel (with index i). With these definitions, the penalized log-likelihood function $\phi(\boldsymbol{A})$ is given by:

$$\begin{aligned} \phi(\boldsymbol{A}, \tilde{\alpha}) &= -\mathcal{L}(\boldsymbol{A}) + \lambda \mathcal{R}(\boldsymbol{A}, \tilde{\alpha}) \\ -\mathcal{L}(A) &= \sum_{i=1}^{N_t} \sum_{s=1}^{S} \hat{y}_i^s(\vec{A}_i) - y_i^s \ln(\hat{y}_i^s(\vec{A}_i)) \\ \mathcal{R}(\boldsymbol{A}) &= \sum_{q=1}^{N_p} ||(E_q \boldsymbol{A}) {\boldsymbol{\varphi}_q}^T - \boldsymbol{D}\boldsymbol{\alpha}_q||_F^2 + v||\boldsymbol{\alpha}_q||_{\text{row}-0}. \end{aligned} \tag{5.15}$$

For the regularization term, N_p denotes the number of overlapping image patches that are extracted for each channel and q is the patch index. The coefficient matrix for the q-th patch is denoted by $\boldsymbol{\alpha}_q$ and $\tilde{\alpha} = \{\boldsymbol{\alpha}_1, ..., \boldsymbol{\alpha}_q\}$ is the set of all coefficient matrices. Applying the patch extraction operator E_q to $\boldsymbol{A}$ results in a matrix $E_q \boldsymbol{A} \in \mathbb{R}^{N \times C}$, i.e. an image patch with N pixels is extracted for each channel and written into a matrix. For each patch index q, the transformation matrix $\boldsymbol{\varphi}_q$ is determined individually, because the noise correlations and noise levels vary depending on the local basis material thickness combinations. The covariance matrix

for the i-th pixel $\boldsymbol{R}^{(i)}$ can be estimated by calculating the inverse of the fisher information matrix $\boldsymbol{\mathcal{I}}^{(i)} \in \mathbb{R}^{C\times C}$:

$$\boldsymbol{R}^{(i)}_{\alpha\beta} \geq \left(\boldsymbol{\mathcal{I}}^{(i)}\right)^{-1}_{\alpha\beta};\ \boldsymbol{\mathcal{I}}^{(i)}_{\alpha\beta}(\vec{A}_i) = -E\left[\frac{\partial^2 \mathcal{L}_i(\vec{A}_i)}{\partial A_i^\alpha \partial A_i^\beta}\right], \tag{5.16}$$

see eq. 5.8 for the definition of $\mathcal{L}_i(\vec{A}_i)$. Equation 5.16 is known as Cramér-Rao inequality and defines a lower bound (the Cramér-Rao lower bound, CRLB) for the (co-)variance of an unbiased estimator. Since the ground truth values for the basis material thicknesses $\vec{A}_i$ are unknown, we use the result of a maximum-likelihood decomposition ($\vec{\mathcal{A}}_i$, compare eq. 5.8) as an approximation. For each patch index q, we calculate the average covariance matrix $\boldsymbol{R}_q \in \mathbb{R}^{C\times C}$ by averaging the covariance matrices of each pixel contained in the image patch. The transformation matrix $\boldsymbol{\varphi}_q$ is then calculated from $\boldsymbol{R}_q$ as described in section 5.2.2. As suggested for the single-channel case [331], the penalized log-likelihood function can be minimized by alternatingly optimizing $\boldsymbol{A}$ and $\tilde{\alpha}$ while keeping the other optimization variable fixed. First, the basis material images are kept fixed and $\tilde{\alpha}$ is updated:

$$\min_{\tilde{\alpha}} \sum_{q=1}^{N_p} ||(E_q \boldsymbol{A})\boldsymbol{\varphi}_q{}^T - \boldsymbol{D}\boldsymbol{\alpha}_q||_F^2 + v||\boldsymbol{\alpha}_q||_{\text{row}-0} \tag{5.17}$$

This optimization problem can be solved individually for each patch by the SOMP algorithm. In a second step, $\tilde{\alpha}$ is kept fixed and the basis material images (represented by $\boldsymbol{A}$) are optimized:

$$\min_{\boldsymbol{A}} -\mathcal{L}(\boldsymbol{A}) + \lambda \left(\sum_{q=1}^{N_p} ||(E_q \boldsymbol{A})\boldsymbol{\varphi}_q{}^T - \boldsymbol{D}\boldsymbol{\alpha}_q||_F^2\right) \tag{5.18}$$

If $\tilde{\alpha}$ is kept fixed, this optimization problem is separable with respect to each pixel because $\mathcal{L}(\boldsymbol{A})$ is separable and the second term in eq. 5.18 can be rewritten as:

$$\sum_{q=1}^{N_p} ||(E_q \boldsymbol{A})\boldsymbol{\varphi}_q{}^T - \boldsymbol{D}\boldsymbol{\alpha}_q||_F^2 = \sum_{i=1}^{N_t} \sum_{q\in\mathcal{P}_i} ||\boldsymbol{\varphi}_q \vec{A}_i - \vec{\rho}_q||_2^2, \tag{5.19}$$

where $\mathcal{P}_i$ is the set of all patch indices that contain pixel i and $\vec{\rho}_q$ is the row of $(\boldsymbol{D}\boldsymbol{\alpha}_q)$ that corresponds to pixel i (written as vector). Minimizing eq. 5.18 thus reduces to solving N_t subproblems with C optimization variables.
Compared to ML decomposition, the computational complexity for penalized likelihood (PL) decomposition is strongly increased. On the workstation used for our experiments (Intel Xeon processor E5-2690 v4), the ML decomposition for 400×400 pixels takes less than one second and the PL decomposition using a patch size of $10 \times 10 = 100$ pixels and a sliding distance of 3 pixels takes about 100 seconds (dominated by the computational time for the SOMP algorithm). With slight losses in image quality, the computational time for the PL decomposition could be reduced by either increasing the sliding distance or decreasing the patch size, as the total number of patches is inversely proportional to the squared sliding distance and the computational complexity of the SOMP algorithm increases at least proportionally to the patch size.

5.2.3 Numerical simulation

We simulated a coronary angiography projection measurement with a PCD (1mm thick cadmium telluride (CdTe) sensor, 400×400 pixels, 5 thresholds per pixel) and a tungsten anode with aluminum filtration. A segmented coronary artery from a clinical CT scan was placed in the center of the FORBILD thorax phantom [337]. The artery contained a mixture of gadolinium (Gd) contrast agent and blood with an effective Gd concentration of 50 mg/ml. Furthermore, two

Table 5.1: Basis material thicknesses for the simulated calibration measurement.

basis material	thicknesses
PVC	[0.3, 1.0, 1.7] cm
PMMA	[9.5, 17.25, 25.0] cm
Gd	[0.0, 37.97, 75.95] μm

Table 5.2: Acquisition parameters for the numerical simulation.

parameter	kVp	Al filter	threshold positions
optimization range	60-140 kVp	0-5 mm	20-138 keV
optimum parameters	140 kVp	0.7 mm	[20,36,50,66,80] keV

calcifications were placed in the artery. The chemical composition of the calcifications was chosen according to data published by Schmid [361]. The remaining parts of the digital phantom were filled with soft tissue and bone. Figure 5.2(a) shows a slice of the modified FORBILD phantom and an overview of the acquisition geometry. The projection measurement was taken at a 40-degree angle (40° LAO [362]). This projection angle is often chosen for conventional angiography because the spine does not block the view of the coronary vessels. A spatially varying detector response was simulated by incorporating a Gaussian distribution of the pixel-wise thresholds ($\sigma = 1\,\text{keV}$) centered around the corresponding global values. We optimized the acquisition parameters (filter thickness, acceleration voltage and threshold positions) by minimizing the noise level in the gadolinium basis material image while keeping the absorbed dose constant. The noise level was determined by calculating the CRLB (cf. eq. 5.16) for 15 cm of soft tissue, 0.5 cm of bone and 12.5 μm of pure gadolinium (equivalent to 2 mm of gadolinium solution with a concentration of 50 mg/ml) in the beam path. This combination reflects typical path lengths for the simulated projection measurement. As will be shown later, the proposed ML decomposition is practically unbiased and achieves the CRLB. Consequently, the CRLB predicts the noise level in the decomposed basis material images accurately. We calculated the absorbed dose according to the method presented in [363]. The absorbed dose for the simulated projection of the phantom was 1 mGy, which corresponds to an air kerma of 2.77 mGy and $1.3 \cdot 10^6$ photon counts registered by the detector without the phantom in the beam path (flatfield). The resulting acquisition parameters can be found in table 5.2 and figure 5.1 shows the corresponding effective spectra for the five energy bins of the simulated PCD. Prior to the simulated projection measurement of the phantom, we simulated a calibration measurement by placing homogeneous blocks of PMMA, PVC and gadolinium with three different thicknesses (cf. table 5.1) in the beam path. In total, 27 calibration points were acquired by measuring all possible combinations of basis material thicknesses (using exactly one block of each basis material). The thickness ranges for the different basis materials were chosen such that they cover the relevant pathlengths in the projection measurement of the phantom. As there are no dose considerations, the calibration measurement was simulated with significantly better statistics than the projection measurement ($\approx$ 100 times more photons) and the exposure time was adjusted to have the same photon statistics for each calibration point. In addition, a similar "test-calibration" measurement was simulated to evaluate the bias of the proposed decomposition algorithm. In this case, we used 8 evenly spaced thicknesses for each basis material (resulting in $8^3 = 512$ combinations) and the thickness ranges were increased by 10 % compared to the original calibration measurement. The dictionary for regularized material decomposition was trained on the conventional projection image of the thorax phantom. We used a patch size of 10×10 pixels and a sliding distance of 3 pixels. These parameters were chosen heuristically as a compromise between image quality and computational complexity.

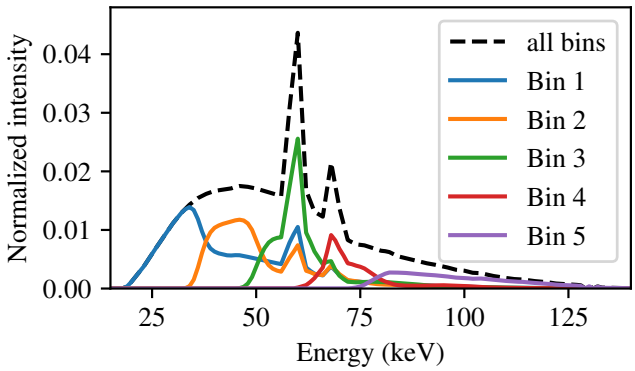

Figure 5.1: Effective spectra (including quantum efficiency and detector response) for the simulated photon-counting detector.

5.2.4 Experimental measurement

Sample preparation The sample used for the experimental measurement was a semi-anthropomorphic thorax phantom (QRM-GmbH, Forchheim, Germany). The thorax phantom contains a cylindrical notch in the center that was equipped with a water container. To emulate a coronary angiography scenario as realistically as possible, an ex-vivo porcine heart was put into this water container. Figure 5.2(b) shows an overview of the sample and the acquisition geometry. Instead of injecting contrast agent directly into the coronary arteries of the porcine heart, we used artificial arteries made of polyethylene tubes and fixed them on top of the heart's biological main arteries. Thereby, a uniform and stable filling of the artificial vessels was provided during the measurement. According to the natural size of human coronary arteries [364], the tubes had an inner diameter of 1 and 3.5 mm, whereby the latter were equipped with artificial coronary calcifications fixated to the internal wall of the tubes. These calcifications were made of hydroxylapatite powder ($Ca_5(PO_4)_3(OH)$), which was alloyed with tissue glue (Histoacryl, B. Braun Surgical S.A., Rubí, Spain) to solidify the mixture and to resemble the chemical composition of calcified plaques according to [361]. The prepared tubes were filled with clinical gadolinium contrast agent (DOTAREM®, Guerbert, Sulzbach, Germany), commonly used in magnetic resonance imaging. Approximately 15 ml of the contrast agent were injected in undiluted form, corresponding to a gadolinium concentration of 78.6 mg/ml.

Experimental setup The experimental measurements were performed on a static setup (see figure 5.3) consisting of an X-ray tube, a photon-counting detector system and several positioning devices mounted on an optical table. The statically mounted X-ray source (XWT-160-CT, X-RayWorX, Garbsen, Germany) is a micro-focus tube with a tungsten reflection target. Data acquisition was done with the hybrid pixel detector prototype SANTIS 0804 (Dectris AG, Baden, Switzerland) with a 1 mm thick CdTe sensor and an active area of 515 x 256 pixels. The detector pixel size is 150 μm and each pixel has four separate energy thresholds. The setup is equipped with a sample holder movable in x-, y- and z-direction. Between the X-ray tube and the sample holder, the calibration phantoms are mounted on linear stages. The phantoms consist of an aluminum frame that can be equipped with cuboid blocks of various thicknesses and materials.

Acquisition parameters Prior to the acquisition of the phantom, we performed calibration measurements in order to determine the fit parameters of the proposed forward model (cf. eq. 5.5 - 5.6) for each detector pixel. For the calibration, a grid of 4x4x4 different thickness combinations of titanium, POM (polyoxymethylen) and gadolinium contrast agent were used (cf. table 5.3), resulting in a total of 64 calibration points. The acquisition time was increased

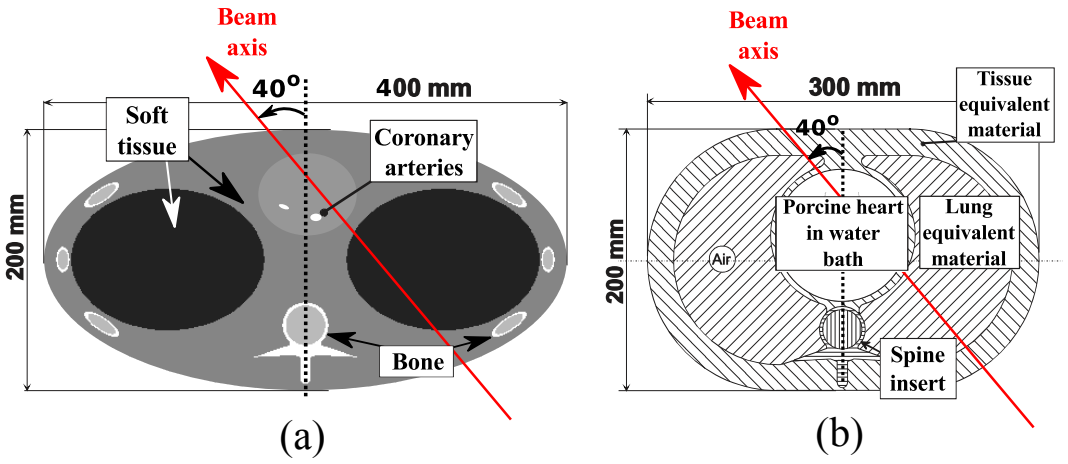

Figure 5.2: Overview of phantoms and acquisition geometry for the numerical simulation (a) and the experimental measurement (b). The numerical simulations were performed with a modified FORBILD thorax phantom [337]. For the experimental measurements, a semi-anthropomorphic thorax phantom was used. The projection angle for both the numerical simulation and the experimental measurement was 40° relative to the sagittal plane.

Table 5.3: Basis material thicknesses for the calibration routine in the experimental measurement.

basis material	thicknesses
Ti	$[0, 0.01, 0.034, 0.066]$ cm
POM	$[0, 3.0, 6.0, 12.4]$ cm
Gd (39.3 mg/ml)	$[0.0, 0.45, 0.90, 1.48]$ cm

for calibration points with higher line integral values to provide a constant photon statistic of approximately $2.5 \cdot 10^5$ photons per pixel. Similar to the numerical simulation, the experimental measurement of the thorax phantom was done at a 40-degree angle (40° LAO [362]). Due to the limited field of view of the detector system, the sample stage was used to extend the imaged area to 23x8 cm^2 by stitching a series of acquisitions with the beam being collimated to the detector area (16cm^2) in order to suppress scattered radiation. Analogous to the numerical simulation, the acquisition parameters were optimized to yield minimum noise in the decomposed basis material images. The setup geometry and the optimized acquisition parameters for the measurement are listed in table 5.4. The radiation dose was assessed by means of the incident air kerma at the position of the sample. The kerma was measured by a NOMEX® T11049 Multimeter (PTW, Freiburg, Germany).

Table 5.4: Summary of acquisition and geometry parameters for the experimental measurement.

acceleration voltage	120 kVp	tube current	2.3 mA
aluminum filtration	3.5 mm	thresholds	20, 48, 62, 74 keV
flatfield photon counts	$2 \cdot 10^6$	incident air kerma	9.36 mGy
source-detector dist.	1200 mm	source-sample dist.	800 mm
effective pixel size	100 μm	field of view	230 × 77 mm

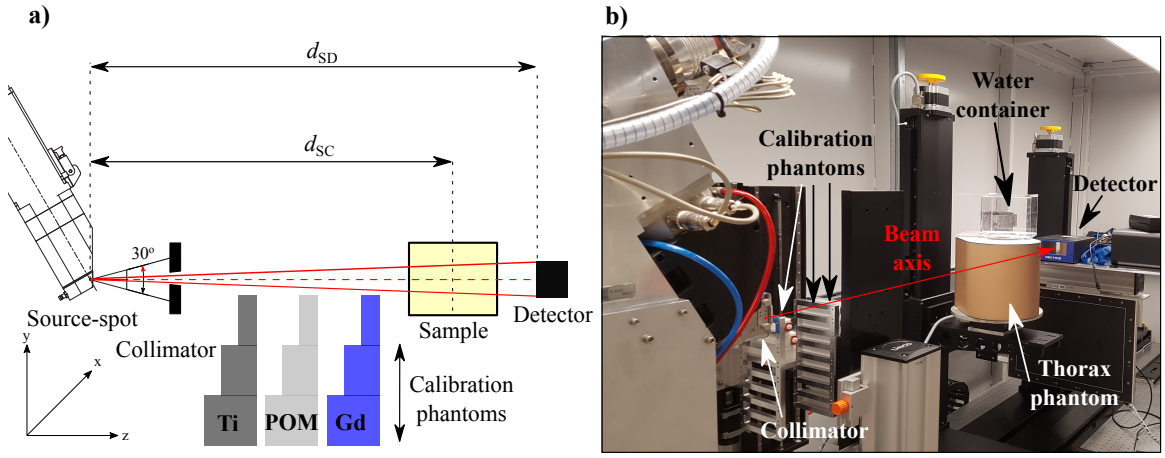

Figure 5.3: Sketch (left) and picture (right) of the experimental setup used for the experimental measurements. The source-to-detector and source-to-center distance are denoted by d_{SD} and d_{SC}, respectively. The calibration phantoms (Ti, POM and Gd) consist of cuboid blocks of various thicknesses mounted to an aluminum frame and can be moved into the x-ray beam by linear stages. The x-ray beam is collimated to the active detector area. The sample comprises a porcine heart that was put into a water container and surrounded by a thorax phantom.

5.3 Results

5.3.1 Numerical simulation

After calibrating the polynomial forward model with the simulated calibration measurement (cf. eq. 5.7), we performed an ML decomposition of the "test calibration" measurement to evaluate the bias of the proposed model. Table 5.5 shows the maximum and root-mean squared error (RMSE) of the decomposed basis material line integrals. More precisely, the mean error of all detector pixels was calculated and then the maximum error and the RMSE were calculated by considering the aforementioned mean error for all calibration points. In addition to the decomposition error for gadolinium, we calculated the decomposition error for gadolinium with a reduced density of 50 mg/ml (abbreviated Gd_{50}). This density corresponds to the gadolinium concentration for the simulated contrast agent and thus simplifies a comparison with the decomposition errors for the other basis materials. The maximum and RMSE decomposition errors inside the calibration range are in the order of $20 - 120\,\mu$m. These errors are slightly larger compared to the results of Cardinal and Fenster [244] for the original rational polynomial model in the dual energy case. However, the proposed model has been tested for a more complex decomposition task - using 5 energy bins and three basis materials instead of two basis materials and two spectra in the dual energy case. Outside the calibration range, the maximum error rises by a factor of ≈ 3, which is consistent with the results from [244]. Moreover, we investigated the variance of the proposed decomposition algorithm. Figure 5.4(a) shows the variance of the basis material line integrals for one specific decomposition task (12.5 μm of Gd, 0.93 cm of PVC and 14.5 cm of PMMA) as a function of the photon statistics, whereas fig. 5.4(b) shows the systematic bias for the same task. Transformed into a Gd, bone and soft tissue basis, these thicknesses reflect typical pathlengths for the projection measurement of the phantom. Furthermore, fig. 5.4(a) shows the theoretical minimum for the variance (CRLB, cf. eq. 5.16). For this experiment, the threshold dispersion was disabled (for both the calibration and the decomposition) as it represents an additional contribution to the variance and thus complicates a comparison with the CRLB.

The proposed algorithm achieves the CRLB over a large range of photon statistics ($\approx 10^8 - 10^3$ flatfield photon counts). Similarly, the bias remains small and constant over this range. For

Table 5.5: Maximum and root-mean squared (RMSE) errors for the decomposition of the "test-calibration" measurement. Gd_{50} corresponds to gadolinium with a reduced density of 50 mg/ml (compare section 5.3.1). The maximum error was calculated separately inside and outside of the calibration range.

basis material	max. error [μm] inside cal. range	max. error [μm] outside cal. range	RMSE [μm]
Gd	0.266	0.710	0.127
Gd_{50}	42.0	112	20.1
PVC	71.6	207	33.3
PMMA	123	328	53.9

very high photon counts ($> 10^8$), the algorithm does not achieve the CRLB anymore. This is because the number of photons of the test measurement exceeds the number of photons for the calibration measurement. Therefore, the noise of the calibration measurement cannot be neglected anymore and represents a significant contribution to the experimental variance. In the case of very low photon counts ($< 10^3$), deviations from the CRLB are also observable. In this range of photon counts, the decomposition becomes increasingly biased. The reason for this statistical bias is the propagation of noise through the nonlinear decomposition algorithm [353]. Figure 5.5 shows the basis material images for the simulated projection angiography. We compare the ground truth images with the results from ML decomposition (cf. eq. 5.8) and PL decomposition using a dictionary prior (cf. eq. 5.15). The stopping parameter for the SOMP algorithm ϵ (cf. eq. 5.17 and [356] for a detailed explanation of the SOMP algorithm) was chosen heuristically according to visual appearance and the regularization strength λ was tuned to achieve 15 % of the noise level for ML decomposition. The ground truth images were calculated by performing a basis transformation of the materials in the thorax phantom into equivalent volume fractions of PVC, PMMA, Gd and then projecting the resulting 3D basis material images onto the detector. Although one could transform the output of the decomposition algorithms (equivalent thicknesses of PVC, PMMA and Gd) into a bone, soft tissue and Gd basis, we chose to work with the calibration material because the energy-dependent attenuation of PVC and PMMA is already similar to bone and soft tissue, respectively. Furthermore, we rescaled the gadolinium line integrals to an equivalent gadolinium density of 50 mg/ml (abbreviated by Gd_{50}), which corresponds to the gadolinium density of the contrast agent. For comparison, fig. 5.6(a) shows the conventional projection image, which was generated by calculating the log signals l_i using all energy bins:

$$l_i = -\ln\left(\frac{\sum_{s=1}^{5} y_i^s}{\sum_{s=1}^{5} b_i^s}\right). \tag{5.20}$$

The blue arrows in the conventional angiography image (fig. 5.6(a)) show areas where the ribs overlap with the coronary arteries. In this region, the "anatomical noise" impedes the assessment of the coronary vessels. In the basis material images (see fig. 5.5), however, contributions from bone and the contrast agent can be clearly separated. As indicated by the green arrow, calcified tissue and contrast agent cannot be distinguished reliably in the conventional image. This is because the energy-dependent attenuation averaged over the effective spectrum is quite similar for calcified tissue and the mixture of blood and gadolinium contrast agent. However, the calcifications can be identified as bright spots and vacanies in the PVC and Gd basis material image, respectively (see bottom left zoom in fig. 5.5). No systematic deviations or artifacts are observable between the basis material images generated by ML decomposition and the ground truth. But compared to the conventional image, the noise level of the ML decomposed images is strongly increased. PL decomposition successfully reduces the noise level compared to ML decomposition while retaining the anatomical features.

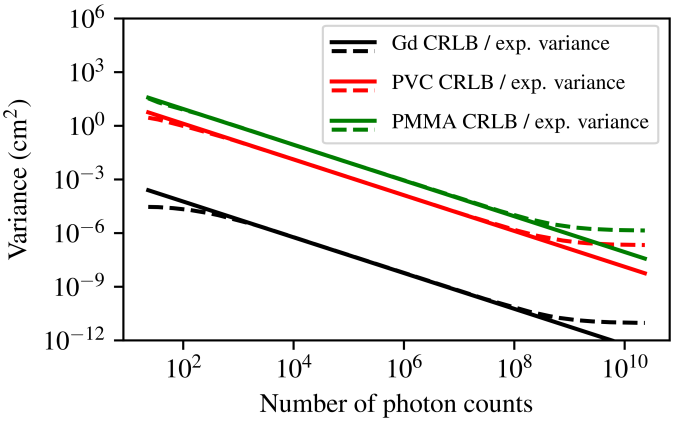

(a)

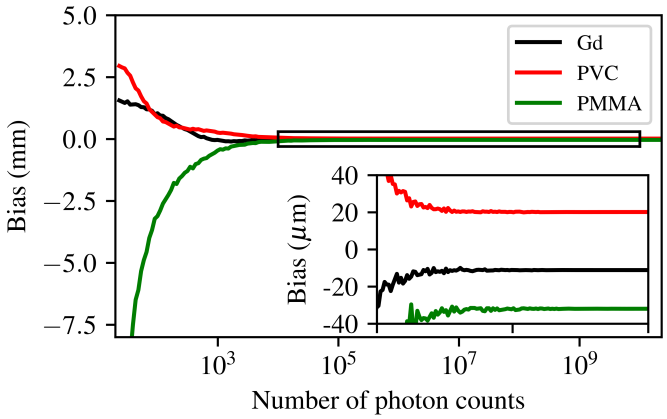

(b)

Figure 5.4: Experimental variance (a) and bias (b) of ML decomposition for a decomposition task of 12.5 μm of Gd, 0.93 cm of PVC and 14.5 cm PMMA. The variance and bias are plotted as a function of the photon statistics in the flatfield, defined as the number of photons registered by the detector without the basis materials in the beam path. The variance is compared with the theoretical lower limit for an unbiased estimator (CRLB).

5.3.2 Experimental measurement

The data acquired in the experimental measurement of the thorax phantom (cf. section 5.2.4) was decomposed into basis material line integrals of gadolinium solution, titanium and POM. After decomposition into the calibration materials, we performed a basis transformation into a (pure) gadolinium, soft tissue and calcium basis. Similarly to the numerical simulation, the obtained gadolinium line integrals were rescaled to equivalent densities of the used contrast agent (78.6 mg/ml Gd, abbreviated by Gd_{80}). Figure 5.7 shows the resulting basis material images for ML and PL decomposition, whereas fig. 5.6(b) shows the conventional image calculated by eq. (5.20). Similar to the numerical simulation, the stopping parameter ϵ for the SOMP algorithm was chosen according to visual appearance and the regularization strength λ was set to obtain 15 % of the noise level of an ML decomposition of the acquired data. In the conventional image, the calcifications overlap with the gadolinium filled arteries and the soft tissue background arising from the porcine heart and the thorax phantom. This makes it difficult to distinguish between calcifications (cf. fig. 5.6(b) arrow 2), trapped air inclusions (cf. fig. 5.6(b) arrow 1) and "anatomical noise" caused by the structural features of the porcine heart (cf. fig. 5.6(b) arrow 3). In contrast, these features can be differentiated by combining the information of all three basis material images. As already demonstrated in the numerical simulation, intensively calcified

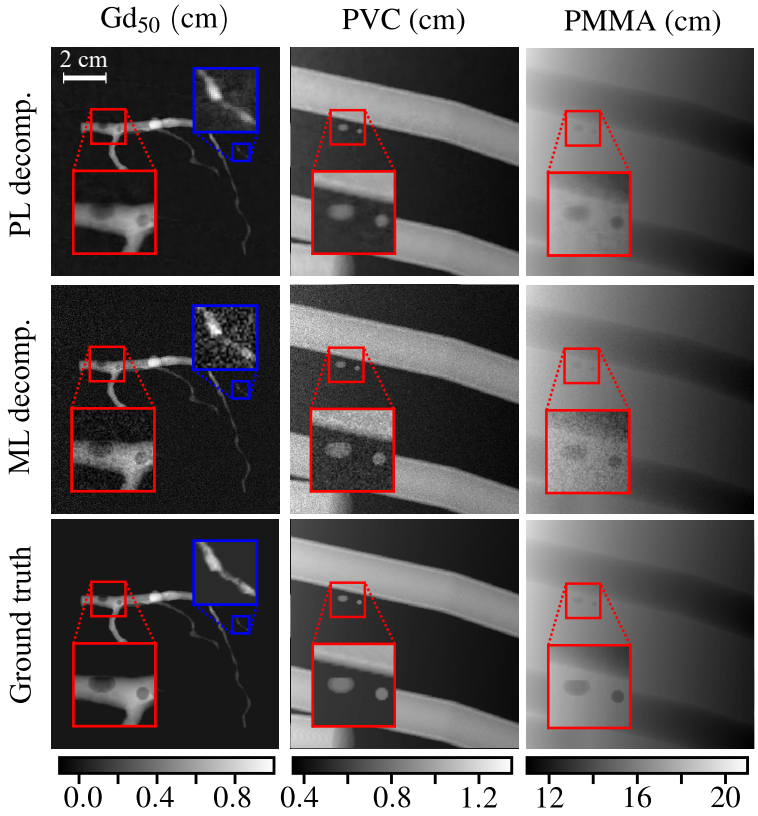

Figure 5.5: Basis material images for the simulated angiography measurement. The first and second row show the results of a PL and ML decomposition, respectively. The ground truth images are displayed in the third row. The bottom left zoomed inset shows the calcifications of the coronary artery. In the case of the Gd image, an additional inset on the top right shows a magnification of a small coronary vessel.

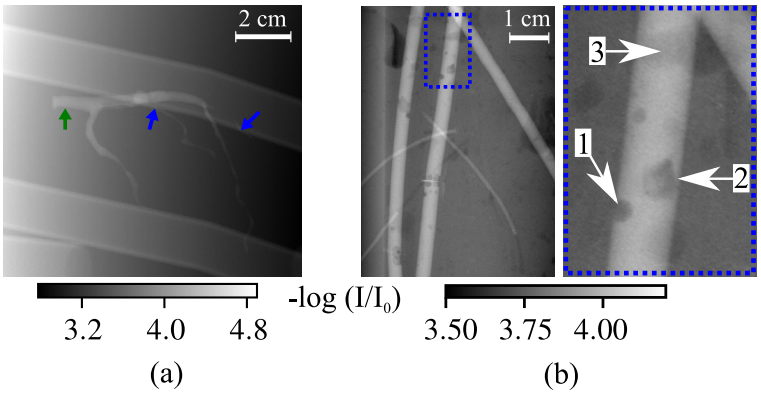

Figure 5.6: Conventional image for the simulated (a) and experimental (b) angiography measurement. Equation 5.20 explains the calculation of the conventional images. The colored arrows in (a) mark areas where the basis material images (see fig. 5.5) offer additional diagnostic information. The blue inlay in (b) highlights an air inclusion (1), a calcification (2) and "anatomical noise" (3).

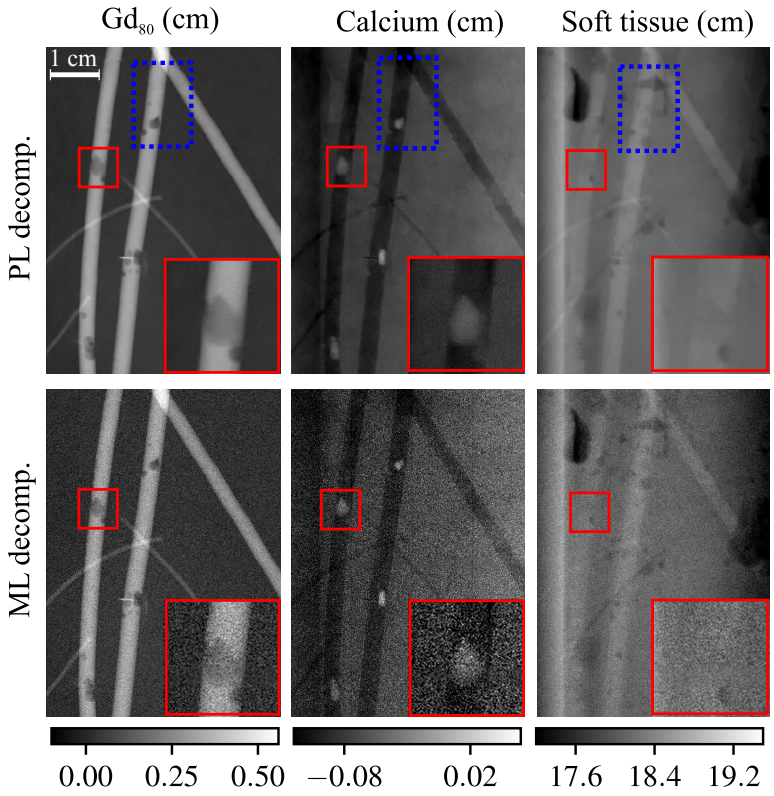

Figure 5.7: Basis material images for the experimental measurement calculated with PL (top row) and ML (bottom row) decomposition. The red inset shows a zoomed view of a calcification and the blue rectangle highlights a scenario where calcifications, air inclusion and "anatomical noise" can be distinguished (compare zoomed view in fig. 5.6).

structures and the contrast agent can have similar average attenuation values, which results in poor detectability of the introduced calcifications in the conventional image. In contrast, the basis material images provide a clear separation of vessels and the introduced calcifications (cf. red inlays fig. 5.7).

5.4 Discussion

The proposed polynomial forward model is able to accurately fit the spectral transmission as a function of the basis material line integrals with a small number of calibration measurements. In combination with ML estimation, this results in a practical, minimally biased decomposition algorithm that does not require knowledge about the source spectrum and the detector response. In real experiments, other effects such as source and detector drifts will most likely dominate over the bias of the empirical model. Moreover, the proposed estimator is efficient (i.e. achieves the CRLB) over a large range of photon counts. The range of ultra low photon statistics ($< 10^3$ photon counts, which is less than 0.1 % of the dose for the numerical and real experiment), where deviations from the CRLB and statistical bias are observable, is not relevant for projection imaging because the image quality would not be satisfactory. An important advantage of our decomposition method compared to the approach by Cardinal and Fenster is achieving the CRLB even if the number of thresholds exceeds the number of basis materials. We therefore investigated two different scenarios (using 4 thresholds for the experiment and 5 thresholds for the simulation) to demonstrate that the method is not limited to the case in which the number of thresholds exceeds the number of basis materials by one. Furthermore, an accurate simulation of the experimental measurement was not possible, as the response of the photon counting

detector prototype was unknown.
We focused on spectral angiography as a promising application of projection-based multi-material decomposition. We have designed the experiments with clinically realistic conditions, i.e. using human phantom sizes and realistic amounts and concentrations of the gadolinium contrast agent as well as reasonable dose levels.
Although clinical angiography is, in general, performed with iodinated contrast agents, the use of iodine in spectral imaging is disadvantageous, as its K-edge is located at the lower end of the energy range typically used in clinical applications. The photon spectrum behind the patient hardly contains energies below the iodine K-edge because the human body absorbs a large fraction of the low energy photons, especially when the X-rays pass through bone or other strongly absorbing materials. As the energy dependent attenuation of iodine above the K-edge (33 keV) can be accurately reproduced by a suitable linear combination of bone and soft tissue, a three-material decomposition leads to excessive noise [240]. We have chosen gadolinium as a contrast agent because of the higher K-edge (50 keV) and because it is approved as a contrast agent for magnetic resonance imaging. Gadolinium-based contrast agents have already been used for conventional [365, 366] as well as digital subtraction angiography [367, 368] in cases where the application of iodine-based contrast agents was contraindicated. With the help of a pressure injector [366], a considerable dilution of the contrast agent (reducing the CNR) can be prevented. Concerning the maximum amount of Gd-contrast agent, a limit of $0.3 - 0.4$ mmol gadolinium per kilogram can be found in the literature [369, 370, 366]. For a patient weighing 75 kg, this corresponds to a limit of $23 - 30$ ml of the contrast agent used in this study. We injected 15 ml of the contrast agent in the experimental measurement, which is below the recommended limit for angiography and in line with the recommended dose for magnetic resonance imaging [371]. The possibility of various side effects [372] is a disadvantage of gadolinium-based contrast agents. However, the presented algorithm also works with different contrast agents exhibiting the K-edge in the diagnostic energy range. For example, recent studies have identified tantalum-based contrast agents (K-edge at 67 keV) as a promising candidate for X-ray imaging [373].
The results of numerical and real experiments show that contributions from contrast agent, soft tissue and calcium can be clearly separated, which provides valuable additional information. The calcium and soft tissue images allow one to identify plaques and determine their chemical composition, whereas the "anatomical noise" is removed in the contrast agent image, resulting in a DSA-like image without the risk of motion artifacts. However, compared to the conventional images, the ML decomposed basis material images show a strongly increased noise level. Noise amplification and a degradation of the SNR compared to conventional images are well-known problems limiting the utility of basis material decomposition for spectral CT or spectral projection measurements. By combining the decomposition algorithm with an efficient denoising strategy, the effect of noise amplification can be mitigated, increasing the image quality and thus the clinical value of the basis material images. In our experiments, the PL decomposition in combination with a dictionary prior reduced the noise level by a factor of ≈ 6.7 while preserving even small and low contrast structures (cf. zoomed insets in fig. 5.5 and fig. 5.7). In this work, we chose to use a dictionary-based prior instead of a more common nearest-neighbor-based regularization strategy because the assumption of a piece-wise constant image is better suited for CT than for projection imaging. However, regularized material decomposition has recently also been successfully performed with nearest-neighbor-based priors [235, 374]. A comparison between different regularization strategies would be interesting, but is considered out of the scope of this work.
To achieve a comparable image quality to the numerical simulation, the dose level for the experimental measurement was increased considerably (cf. incident air kerma 9.36 mGy vs 2.77 mGy). However, this increased demand for radiation dose can be explained by taking into account the different conditions for the experiment. Although the numerical simulation considered charge-sharing and the resulting overlap of the energy bins as a major cause of amplified noise in

spectral imaging, the overlap of the energy bins in the experimental measurement was possibly larger than assumed in the simulation. Apart from that, the higher number of available energy bins in the simulation (5 vs 4), adversely affects the noise level in the measurement. Lastly, the experimental outcome might be influenced by imperfections and sensor effects of current photon-counting detectors. The presented work did not consider scattered radiation and its influence on the proposed method, as it was neglected in the simulations and suppressed in the experimental measurements by collimation of the X-ray beam to a small area of 16 cm^2. In a clinical large field of view scenario, scattered radiation would add a non-negligible low-frequency background to the acquired data resulting in biased basis material images. This low-frequency bias cannot be removed by the dictionary denoising approach. In this case, the algorithm would have to be combined with either hardware or software-based scatter suppression strategies.

5.5 Conclusion

We have developed a basis material decomposition algorithm that features two key novelties. First, an empirical forward model based on rational functions. The model uses a comparatively small number of fit parameters, which allows for a pixel-wise accurate tuning of the model with a small number of calibration measurements. By inverting the originally proposed polynomial functions and combining the resulting forward model with a statistical model of the measurement, the algorithm achieves the CRLB even if the number of acquired energy bins exceeds the number of basis materials. The second novelty is a joint regularization strategy for projection images based on a dictionary prior. Compared to single-channel dictionary denoising, the SOMP algorithm encourages common structures and edges in the reconstructed basis material images. We demonstrated that this denoising strategy leads to highly improved image quality compared to maximum-likelihood decomposition. In this work, we focused on spectral angiography as a potential clinical application of projection-based material decomposition with photon counting detectors. Numerical and real experiments show that spectral angiography with clinically relevant dose levels and concentrations of gadolinium contrast agent is feasible using the proposed decomposition algorithm and currently available photon counting detector technology. The ability to discriminate between calcium-containing structures, soft tissue and contrast agent as well as the removal of “anatomical noise” are important advantages compared to conventional angiography. We therefore believe that spectral angiography can offer additional value for clinical diagnosis.

6 Joint statistical iterative material image reconstruction for spectral computed tomography using a semi-empirical forward model

One-step SIR algorithms allow to directly estimate the basis material volume fractions from the spectral projection measurements (compare eq. 2.149). This is accomplished by a forward model which connects the (expected) spectral projection measurements and the material selective images. It has been demonstrated that performing image reconstruction and material decomposition jointly leads to superior image quality compared to two-step methods (projection-based and image-based material decomposition). Furthermore, one-step SIR algorithms enable new acquisition methods for spectral CT by allowing the reconstruction of basis material images from sparse spectral measurements [375, 376]. Previous approaches have focused on a physical forward model (compare eq. 2.155). In this chapter, we present a new algorithm for joint statistical iterative material image reconstruction that relies on a semi-empirical forward model which is tuned by calibration measurements. This strategy allows to model spatially varying properties of the imaging system without requiring detailed prior knowledge of the system parameters. We have developed an efficient optimization algorithm by modifying an existing approach based on separable surrogate functions [377] (see also section 2.6.2 in the theory part). Our results demonstrate highly accelerated convergence and reduced reconstruction times. Mory et al. have later independently confirmed our findings by comparing five optimization algorithms for one-step spectral CT [378]. In their case study, our algorithm converged one to three orders of magnitude faster than the other optimization algorithms. Moreover, we demonstrate how joint SIR techniques can mitigate statistical bias (cf. section 2.5). Although post-processing techniques for two-step material decomposition algorithms can also improve the image quality, they typically cannot correct systematically biased images. Suppressing statistical bias is particularly relevant for low-dose quantitative spectral CT. Our findings concerning the benefits of one-step SIR algorithms as well as some of the methodology devised in this chapter have motivated the development of a one-step SIR algorithm for spectral phase-contrast CT, which will be discussed in chapter 9.
The text and the figures in this chapter are taken from the following publication of the author:

- Mechlem, Korbinian, et al. "Joint statistical iterative material image reconstruction for spectral computed tomography using a semi-empirical forward model." IEEE Transactions on Medical Imaging 37.1 (2017): 68-80.

Compared to the original publication, the text has been slightly modified to achieve a better integration into this doctoral thesis.

6.1 Introduction

Most spectral CT algorithms separate the process of material decomposition and image reconstruction. Material decomposition is either performed in projection space prior to image reconstruction [225, 226, 227, 228] or an image based decomposition algorithm is applied after reconstructing the images corresponding to different photon energy spectra [222, 223, 224].

Separating these steps is suboptimal because the full statistical information contained in the spectral tomographic measurements cannot be exploited. Statistical iterative reconstruction (SIR) techniques provide a third, mathematically elegant approach to obtaining material selective images. By including a statistical model of the projection measurements and incorporating prior knowledge about the reconstructed images, improved image quality compared to analytical reconstruction methods has been reported for conventional CT. Consequently, SIR enables significant dose reductions for conventional clinical CT [7, 8, 9, 10].
By adapting SIR for spectral CT, image reconstruction and material decomposition can be performed jointly. This is accomplished by a forward model which directly connects the (expected) spectral projection measurements and the material selective images. To obtain this forward model, detailed knowledge of the different photon energy spectra was assumed in previous work [352, 379, 380, 381]. However, as discussed in section 2.7.3, accurately determining the effective spectrum (including the source spectrum and the detector response) is often difficult in practice. This applies in particular to performing spectral CT measurements with photon counting detectors (PCDs). Many projection-based decomposition algorithms therefore rely on calibration measurements to circumvent the problem of determining the spectrum in experimental measurements [244, 225].
In this chapter, a new algorithm for statistical iterative material decomposition is presented. It uses a semi-empirical forward model which is fine-tuned by calibration measurements and allows to model spatially varying source spectra as well as variations of the detector response. Moreover, an efficient optimization algorithm based on separable surrogate functions [185] is employed. This partially negates the major shortcoming of SIR techniques, namely high computational cost and long reconstruction times. Numerical as well as real experiments show that our new algorithm leads to improved results compared to projection-based material decomposition by overcoming limits given by the classical separation of decomposition and reconstruction.

6.2 Methods

6.2.1 Spectral estimator

Basis material decomposition relies on expressing the energy-dependent attenuation coefficient $\mu(E)$ of an object as a linear combination of a few so-called basis materials b:

$$\mu(E) = \sum_{b=1}^{B} A^b f^b(E), \tag{6.1}$$

where A^b is the line integral of material b and $f^b(E)$ describes the energy-dependent attenuation of basis material b. This approximation is valid because the energy-dependent x-ray attenuation can be described as a superposition of only a few physical effects (e.g. Compton- and photoelectric effect, K-edges) in the photon energy range typically used for CT. The expected number of photon counts $\hat{y}_i$ registered at detector pixel i after penetrating the object is given by:

$$\hat{y}_i = \int_0^\infty \phi_{\text{eff},i}(E)\, \exp\left(- \sum_{b=1}^{B} A_i^b f^b(E) \right) dE, \tag{6.2}$$

where E denotes the photon energy and $\phi_{\text{eff},i}(E)$ is an effective x-ray spectrum which includes all source and detector effects. The quantity A_i^b describes the line integral of basis material b along the path between the source and detector pixel i. Based on a recently proposed estimator [249] for projection-based material decomposition, we approximate $\hat{y}_i$ as the sum of a small number ($\sim$ 2-4) of exponential functions:

$$\hat{y}_i \left(A_i^0, ..., A_i^B \right) = \sum_{r=1}^{R} C_{ri}\, \exp\left(- \sum_{b=1}^{B} A_i^b f^b(\mathcal{E}_{ri}) \right), \tag{6.3}$$

with fit parameters C_{ri} and $\mathcal{E}_{ri}$. This can be interpreted as fitting a surrogate spectrum for $\phi_{\text{eff},i}(E)$ in eq. 6.2 which is described as a sum of δ-functions. The approximation is motivated by previous work concerning beam hardening corrections. It was demonstrated and physically motivated that the thickness-dependent attenuation of a homogeneous material can be modeled by considering two dominant effective energies [382, 383]. Depending on the degree of beam hardening, we found that the accuracy can be further improved by adding another energy. The coefficients C_{ri} and $\mathcal{E}_{ri}$ are determined individually for each detector pixel by fitting $\hat{y}_i$ to calibration measurements with known A_i^b. Assuming Poisson statistics for the calibration measurements, this is done by minimizing the following negative log-likelihood function with respect to the fit parameters $\vec{C}_i = (C_{0i}, ..., C_{Ri})^T$ and $\vec{\mathcal{E}}_i = (\mathcal{E}_{0i}, .., \mathcal{E}_{Ri})^T$:

$$\begin{aligned} -L_i(\vec{C}_i, \vec{\mathcal{E}}_i) &= \sum_{k=1}^{M} \hat{y}_{ik}(\vec{C}_i, \vec{\mathcal{E}}_i) - y_{ik} \ln\left(\hat{y}_{ik}(\vec{C}_i, \vec{\mathcal{E}}_i)\right), \\ \hat{y}_{ik}(\vec{C}_i, \vec{\mathcal{E}}_i) &= \sum_{r=1}^{R} C_{ri} \exp\left(-\sum_{b=1}^{B} A_{ik}^b f^b(\mathcal{E}_{ri})\right). \end{aligned} \tag{6.4}$$

The quantity A_{ik}^b represents the basis material line integrals for calibration measurement k and y_{ik} denotes the measured photon count for detector pixel i and calibration measurement k. For PCDs, the statistical distribution of photon counts is usually well described by a Poisson distribution as long as pulse pile-up effects are moderate [78]. Compared to a discretized version of eq. 6.2, the number of fit parameters is significantly reduced and is typically smaller than ten. A low dimensional parameter space allows an efficient and fast pixel-wise minimization of $-L_i(\vec{C}_i, \vec{\mathcal{E}}_i)$ and stabilizes the ill-conditioned problem of fitting a sum of exponential functions [384]. This becomes particularly important when the number of calibration points is reduced, which is desirable from an experimental perspective. We chose to use the Nelder-Mead algorithm to minimize eq. 6.4. The fit parameters $\vec{C}_i$ were constraint to be positive and the energy parameters $\vec{\mathcal{E}}_i$ were constrained to a predefined interval.
Fitting the parameters $\vec{C}_i$ and $\vec{\mathcal{E}}_i$ with the help of calibration measurements is repeated for all effective spectra used for the spectral CT measurement. We therefore introduce the index s to denote the fit coefficients $\vec{C}_i^s$ and $\vec{\mathcal{E}}_i^s$ of different spectra. Similarly, the expected and measured photon counts for different spectra are indicated by $\hat{y}_i^s$ and y_i^s, respectively.
After the surrogate spectra have been determined, projection-based material decomposition can be performed by minimizing the negative log-likelihood of the spectral measurements y_i^s with respect to the basis material line integrals A_i^b:

$$\begin{aligned} -L(A_i^0, ..., A_i^B) &= \sum_{s=1}^{S} \hat{y}_i^s - y_i^s \ln(\hat{y}_i^s), \\ \hat{y}_i^s &= \sum_{r=1}^{R} C_{ri}^s \exp\left(-\sum_{b=1}^{B} A_i^b f^b(\mathcal{E}_{ri}^s)\right). \end{aligned} \tag{6.5}$$

It is possible to directly decompose into a set of basis materials different from the calibration materials by replacing the energy-dependent attenuation functions $f^b(E)$. The calibrated parameters $\vec{C}_i^s$ and $\vec{\mathcal{E}}_i^s$ can still be used.

6.2.2 Penalized likelihood function for SIR

We will restrict ourselves to two basis materials ($B = 2$) in the following, however a generalization to more than two basis materials is possible and will be investigated in the future. Let α_j^b denote the contribution of basis material b to image voxel j ($j = 1, ..., N$). The basis material images

are then compactly written as $\vec{\alpha}^b$ and $\vec{\alpha} = (\alpha_1^1, ..., \alpha_N^1, \alpha_1^2, ..., \alpha_N^2)^T$ denotes the vector of all optimization variables. The detector pixel index i is generalized to a sinogram index, which includes the different projection angles of a CT measurement. Eq. 6.3 represents a continuous and differentiable expression for the expected number of photons counts in dependence of the basis material line integrals and can therefore be interpreted as a semi-empirical forward model (adjusted with the help of calibration measurements) for SIR. By assuming Poisson statistics for the spectral CT measurements y_i^s, we derive the following log-likelihood of the measurement data:

$$\begin{aligned} -L(\vec{\alpha}) &= \sum_{s=1}^{S} \sum_{i=1}^{P} \hat{y}_i^s - y_i^s \ln(\hat{y}_i^s), \\ \hat{y}_i^s(\vec{\alpha}) &= \sum_r C_{ri}^s \exp\left(- \sum_{b=1}^{2} f^b(\mathcal{E}_{ri}^s) \sum_{j=1}^{N} a_{ij} \alpha_j^b \right), \end{aligned} \tag{6.6}$$

where $\sum_{j=1}^{N} a_{ij}\alpha_j^b$ represents a forward projection with projection matrix a_{ij}.

Since the problem of CT reconstruction is ill-conditioned, direct minimization of $-L(\vec{\alpha})$ leads to very noisy basis material image estimates. By incorporating prior knowledge about the basis material images in the form of a regularization term $R(\vec{\alpha})$, noise can be efficiently suppressed. We chose to use the Huber regularization separately for both material images:

$$\begin{aligned} R(\vec{\alpha}) &= \sum_{b=1}^{2} \sum_{j=1}^{N} \sum_{k \in N_j} \lambda_b \psi(\alpha_j^b - \alpha_k^b), \\ \psi(\Delta) &= \begin{cases} \frac{\Delta^2}{2} & \text{if } |\Delta| \leq \gamma \\ \gamma|\Delta| - \frac{\gamma^2}{2} & \text{if } |\Delta| \geq \gamma \end{cases}, \end{aligned} \tag{6.7}$$

where N_j is a geometrical neighborhood of voxel j and γ is a tuning parameter. The idea is to penalize small differences between neighboring voxels quadratically, whereas larger differences, which are more likely to represent image features, are penalized linearly. The parameters λ_b control the strength of the regularization for the basis material image $\vec{\alpha}^b$. The final penalized likelihood function which is minimized during the SIR is given by the sum of the data term $-L(\vec{\alpha})$ and the regularization term $R(\vec{\alpha})$:

$$\vec{\alpha}^{(\text{opt})} = \min_{\vec{\alpha}} \theta(\vec{\alpha}) = \min_{\vec{\alpha}} \left(-L(\vec{\alpha}) + R(\vec{\alpha}) \right). \tag{6.8}$$

6.2.3 Optimization algorithm

The optimization algorithm is based on minimizing separable quadratic surrogate functions $\varphi(\vec{\alpha}; \vec{\alpha}^{(n)})$ of the penalized log-likelihood function $\theta(\vec{\alpha}) = -L(\vec{\alpha}) + R(\vec{\alpha})$ at each iteration step (n). If the following conditions for the surrogate functions are fulfilled, this procedure also minimizes $\theta(\vec{\alpha})$ [185]:

$$\begin{aligned} &1. \quad \varphi(\vec{\alpha}^{(n)}; \vec{\alpha}^{(n)}) = \theta(\vec{\alpha}^{(n)}) \\ &2. \quad \left. \frac{\partial \varphi(\vec{\alpha}; \vec{\alpha}^{(n)})}{\partial \alpha_j^b} \right|_{\vec{\alpha} = \vec{\alpha}^{(n)}} = \left. \frac{\partial \theta(\vec{\alpha})}{\partial \alpha_j^b} \right|_{\vec{\alpha} = \vec{\alpha}^{(n)}} \\ &3. \quad \varphi(\vec{\alpha}; \vec{\alpha}^{(n)}) \geq \theta(\vec{\alpha}). \end{aligned} \tag{6.9}$$

In contrary to the original penalized likelihood function, the surrogate functions can be minimized analytically and in parallel for different voxels (since they are separable). Weidinger et al. [377] have derived a separable surrogate function for the data term $-L(\vec{\alpha})$, assuming that

the effective spectra of the measurements are known and the same for all detector pixels (i.e. using the forward model of eq. 6.2 with $\phi_{\text{eff},i}(E) = \phi_{\text{eff}}(E)\ \forall i$). Since the semi-empirical forward model of eq. 6.3 has a similar mathematical form compared to a discretized version of the spectral forward model given in eq. 6.2, a separable surrogate function for the data term can be derived by following the approach presented in [377]. In order to simplify the derivation of the surrogate function, we define the following abbreviations:

$$
\begin{aligned}
&h_i^s(\hat{y}_i^s) = \hat{y}_i^s - y_i^s\ \ln(\hat{y}_i^s), \quad -L(\vec{\alpha}) = \sum_{s=1}^{S}\sum_{i=1}^{P} h_i^s(\hat{y}_i^s) \\
&l_{ri}^{sb}(\vec{\alpha}^b) = \sum_{j=1}^{N} a_{ij}\alpha_j^b f^b(\mathcal{E}_{ri}^s), \quad l_{ri}^s(\vec{\alpha}) = \sum_{b=1}^{2} l_{ri}^{sb}(\vec{\alpha}^b) \\
&t_{ri}^s(\vec{\alpha}) = e^{-l_{ri}^s(\vec{\alpha})}, \quad \beta_{ri}^{s,(n)} = \frac{\hat{y}_i^s(\vec{\alpha}^{(n)})}{t_{ri}^s(\vec{\alpha}^{(n)})}.
\end{aligned}
\tag{6.10}
$$

In a first step, the summation over the fit coefficients is moved out of the convex functions $h_i^s(\hat{y}_i^s)$. Using the definitions of eq. 6.10, the forward model in eq. 6.6 can be rewritten as:

$$
\hat{y}_i^s = \sum_{r=1}^{R} \frac{C_{ri}^s}{\beta_{ri}^{s,(n)}} t_{ri}^s(\vec{\alpha})\beta_{ri}^{s,(n)}. \tag{6.11}
$$

Since it holds that:

$$
\sum_{r=1}^{R} \frac{C_{ri}^s}{\beta_{ri}^{s,(n)}} = 1, \tag{6.12}
$$

applying Jensen's inequality to $h_i^s(\hat{y}_i^s)$ yields the following surrogate function:

$$
-L(\vec{\alpha}) \leq \sum_{s=1}^{S}\sum_{i=1}^{P}\sum_{r=1}^{R} \frac{C_{ri}^s}{\beta_{ri}^{s,(n)}} h_i^s\left(t_{ri}^s(\vec{\alpha})\beta_{ri}^{s,(n)}\right) \equiv Q_1\left(\vec{\alpha};\vec{\alpha}^{(n)}\right). \tag{6.13}
$$

The next step is to approximate $Q_1\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$ by another surrogate function which is quadratic with respect to l_{ri}^{sb}. This can be accomplished by a second-order expansion of $g_{ri}^s(l_{ri}^{s1}, l_{ri}^{s2}) \equiv h_i^s\left(t_{ri}^s\left(\vec{\alpha}\right)\beta_{ri}^{s,(n)}\right) = h_i^s\left(\exp\left(-\sum_{b=1}^{2} l_{ri}^{sb}\right)\beta_{ri}^{s,(n)}\right)$ about the current line integrals $l_{ri}^{sb}(\vec{\alpha}^{b,(n)})$:

$$
\begin{aligned}
&g_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2}) \approx q_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2}) = g_{ri}^{s,(n)}(l_{ri}^{s1,(n)}, l_{ri}^{s2,(n)}) \\
&+ \sum_{b=1}^{2} \left.\frac{\partial g_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2})}{\partial l_{ri}^{sb}}\right|_{l_{ri}^{sb}=l_{ri}^{sb,(n)}} \left(l_{ri}^{sb} - l_{ri}^{sb,(n)}\right) \\
&+ \frac{1}{2}\sum_{k=1}^{2}\sum_{m=1}^{2} T_{ri}^{s,(n)}\left(l_{ri}^{sk} - l_{ri}^{sk,(n)}\right)\left(l_{ri}^{sm} - l_{ri}^{sm,(n)}\right),
\end{aligned}
\tag{6.14}
$$

where $l_{ri}^{sb,(n)}$ is an abbreviation for $l_{ri}^{sb}(\vec{\alpha}^{b,(n)})$.
To guarantee monotonic convergence of the algorithm, the curvature $T_{ri}^{s,(n)}$ must be chosen such that $g_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2}) \leq q_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2})$. However, similar to conventional CT, the curvature can be determined such that $g_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2}) \approx q_{ri}^{s,(n)}$. Following the arguments in [185], a suitable approximation is $T_{ri}^{s,(n)} \approx y_i^s$. With this choice of curvature the monotonicity of the algorithm is compromised. However, it is difficult to find a curvature which satisfies $g_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2}) \leq q_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2})$ [377]. Furthermore, we chose to use a precomputed approximation to reduce the

computational complexity of the algorithm. Combining eq. 6.13 and eq. 6.14 yields the new surrogate function $Q_2\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$:

$$Q_2\left(\vec{\alpha};\vec{\alpha}^{(n)}\right) = \sum_{s=1}^{S}\sum_{i=1}^{P}\sum_{r=1}^{R}\frac{C_{ri}^{s}}{\beta_{ri}^{s,(n)}}q_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2}). \tag{6.15}$$

Lastly, $Q_2\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$ is replaced by another surrogate function which is separable with respect to the image voxels. By rewriting the line integrals l_{ri}^{sb} :

$$\begin{aligned} l_{ri}^{sb}(\vec{\alpha}^{b}) &= \sum_{j=1}^{N} a_{ij}\alpha_j^b f^b(\mathcal{E}_{ri}^{s}) \equiv \sum_{j=1}^{N} w_{ij}\lambda_{ij}^{srb}(\alpha_j^b) = \\ &\sum_{j=1}^{N} w_{ij}\left(\frac{a_{ij}f^b(\mathcal{E}_{ri}^{s})}{w_{ij}}\left(\alpha_j^b - \alpha_j^{b,(n)}\right) + l_{ri}^{sb,(n)}\right), \end{aligned} \tag{6.16}$$

with:

$$w_{ij} = \frac{a_{ij}}{\sum_{j=1}^{N} a_{ij}}, \quad \sum_{j=1}^{N} w_{ij} = 1, \tag{6.17}$$

Jensen's inequality can be applied to pull the sum over the image voxels out of the functions $q_{ri}^{s,(n)}(l_{ri}^{s1}, l_{ri}^{s2})$, which yields the final surrogate function $Q_3\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$:

$$\sum_{s=1}^{S}\sum_{i=1}^{P}\sum_{r=1}^{R}\sum_{j=1}^{N} w_{ij}\frac{C_{ri}^{s}}{\beta_{ri}^{s,(n)}}q_{ri}^{s,(n)}\left(\lambda_{ij}^{sr1}(\alpha_j^1), \lambda_{ij}^{sr2}(\alpha_j^2)\right). \tag{6.18}$$

This surrogate function is separable with respect to different voxels, but not separable with respect to the basis material contributions (α_j^1,α_j^2) in a certain voxel.
A separable surrogate function for the regularization term $\mathcal{R}\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$ can be computed analogously to conventional CT (see e.g. [185]):

$$\mathcal{R}\left(\vec{\alpha};\vec{\alpha}^{(n)}\right) = \sum_{b=1}^{2}\sum_{j=1}^{N}\sum_{k\in N_j}\lambda_b\left[\frac{1}{2}\psi\left(2\alpha_j^b - \alpha_j^{b,(n)} - \alpha_k^{b,(n)}\right) + \frac{1}{2}\psi\left(2\alpha_k^b - \alpha_j^{b,(n)} - \alpha_k^{b,(n)}\right)\right]. \tag{6.19}$$

The surrogate function for the penalized log-likelihood function $\varphi(\vec{\alpha},\vec{\alpha}^{(n)})$ is then determined by the sum of the surrogate functions for the data term $Q_3\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$ and the regularization term $\mathcal{R}\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$. At each iteration step, the surrogate functions are minimized analytically:

$$\begin{aligned} \vec{\alpha}^{(n+1)} &= \vec{\alpha}^{(n)} - \left(H_{\varphi}^{(n)}\right)^{-1}\cdot\nabla\left(\varphi\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)\right)\Big|_{\vec{\alpha}=\vec{\alpha}^{(n)}}, \\ \left(H_{\varphi}^{(n)}\right)^{-1} &= -\left(H_{Q_3}^{(n)} + H_{\mathcal{R}}^{(n)}\right)^{-1}, \\ \nabla\left(\varphi\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)\right) &= \nabla\left(Q_3\left(\vec{\alpha};\vec{\alpha}^{(n)}\right) + \left(\mathcal{R}(\vec{\alpha};\vec{\alpha}^{(n)}\right)\right) = \nabla\theta\left(\vec{\alpha}^{(n)}\right), \end{aligned} \tag{6.20}$$

where $H_{Q_3}^{(n)}$, $H_{\mathcal{R}}^{(n)}$ and $H_{\varphi}^{(n)}$ are the Hessian matrices of $Q_3\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$, $\mathcal{R}\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$ and $\varphi(\vec{\alpha},\vec{\alpha}^{(n)})$, respectively. The last equality in eq. 6.20 holds because the surrogate function $\varphi\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$ fulfills conditions (1) and (2) of eq. 6.9. Motivated by a similar strategy for conventional CT ("precomputed curvature [185]"), we precompute the Hessian matrices $H_{\varphi}^{(n)}$ based on the initial basis material image estimates:

$$H_{\varphi}^{(n)} \approx H_{\varphi}^{(0)} \equiv H_{\varphi}. \tag{6.21}$$

Table 6.1: Overview of the ordered subsets separable quadratic surrogate algorithm without (top) and with Nesterov acceleration (bottom, compare the table in [385]). The subscript S_m for the gradient indicates that an approximation of the gradient is calculated based on the current angular subset.

Ordered subsets algorithm
For n=0,1,...
For m=0,1,...,M-1
$k = nM + m$
$\vec{\alpha}^{(k+1)} = \vec{\alpha}^{(k)} - (H_\varphi)^{-1} \cdot \nabla_{S_m} \theta \left(\vec{\alpha}^{(k)}\right)$
Ordered subsets algorithm with Nesterov acceleration
Initialize $\vec{z}^{(0)} = \vec{v}^{(0)} = \vec{\alpha}^{(0)},\ t_0 = 1$
For n=0,1,...
For m=0,1,...,M-1
$k = nM + m$
$t_{k+1} = \frac{1}{2}\left(1 + \sqrt{1 + 4t_k^2}\right)$
$\vec{\alpha}^{(k+1)} = \vec{z}^{(k)} - (H_\varphi)^{-1} \cdot \nabla_{S_m} \theta \left(\vec{z}^{(k)}\right)$
$\vec{v}^{(k+1)} = \vec{z}^{(0)} - \sum_{l=0}^{nM+m} t_l \, (H_\varphi)^{-1} \cdot \nabla_{S_m} \theta \left(\vec{z}^{(l)}\right)$
$\vec{z}^{(k+1)} = \vec{\alpha}^{(k+1)} + \frac{t_{k+1}}{\sum_{l=0}^{k+1} t_l} \left(\vec{v}^{(k+1)} - \vec{\alpha}^{(k+1)}\right)$

This approach strongly reduces the computational complexity of the algorithm. Exploiting the separability of $\varphi\left(\vec{\alpha}; \vec{\alpha}^{(n)}\right)$, N 2×2 matrices and the corresponding inverse matrices (to precompute $(H_\varphi)^{-1}$) have to be computed. The approximation in eq. 6.21 is valid if the initial guess $\vec{\alpha}^{(0)}$ is sufficiently close to the optimum. A suitable initial guess can be found by applying projection-based material decomposition (compare eq. 6.5) and then reconstructing the basis material images with filtered backprojection (FBP). Computing the gradient of $Q_3\left(\vec{\alpha}; \vec{\alpha}^{(n)}\right)$ is typically slower than computing the gradient of $\mathcal{R}\left(\vec{\alpha}; \vec{\alpha}^{(n)}\right)$ because it involves computationally expensive forward- and backprojections. Similar to conventional CT, the computational complexity of one iteration can thus be further reduced by using the ordered subsets principle. The projection angles are divided into M subsets $S_1, ..., S_M$ and at each subiteration, an approximation for the gradient is calculated based on a subset S_m of the projection angles. At subiteration m, this corresponds to approximating $-L(\vec{\alpha})$ by:

$$-L(\vec{\alpha}) = \sum_{s=1}^{S} \sum_{i=1}^{P} \hat{y}_i^s(\vec{\alpha}) - y_i^s \, \ln\left(\hat{y}_i^s(\vec{\alpha})\right) \approx M \sum_{s=1}^{S} \sum_{i \in S_m} \hat{y}_i^s(\vec{\alpha}) - y_i^s \, \ln\left(\hat{y}_i^s(\vec{\alpha})\right), \tag{6.22}$$

and then constructing $Q_3\left(\vec{\alpha}; \vec{\alpha}^{(n)}\right)$. One iteration is completed after cycling through all subsets. Ideally, the computation of the gradient of the data term is accelerated by a factor of M while the convergence remains comparable to the original algorithm. Finally, the optimization algorithm can be further accelerated by combining it with Nesterov's momentum technique [188]. This method accelerates convergence by using previously calculated gradients as momentum terms. Combining ordered subsets and Nesterov's momentum technique has led to highly improved convergence for conventional CT [385]. Table 6.1 gives an overview of the ordered subset algorithm and the modified version which incorporates momentum terms.

6.2.4 SIR of decomposed line integrals

Statistical iterative basis material image reconstruction can also be realized by separating the process of material decomposition and image reconstruction. First, a projection based material decomposition algorithm is applied, yielding basis material line integrals in projection space. In a second step, a SIR with the decomposed line integrals is performed. This approach is very flexible because it can be combined with different empirical material decomposition algorithms without requiring a continuous and differentiable forward model. The most straightforward approach is to reconstruct the basis material images separately and independently of each other [238]. In this case, the data term is given by:

$$-L(\vec{\alpha}^b) = \sum_{i=1}^{P} \frac{1}{\sigma_{i,b}^2} \left(A_i^b - \hat{A}_i^b \right)^2, \tag{6.23}$$

where A_i^b is the basis material line integral determined by projection-based decomposition and $\hat{A}_i^b$ represents a forward projection of basis material image $\vec{\alpha}^b$:

$$\hat{A}_i^b = \sum_{j=1}^{N} a_{ij} \alpha_j^b. \tag{6.24}$$

The quantity $\sigma_{i,b}^2$ is an estimate of the variance of the basis material line integral A_i^b according to the Cramér-Rao lower bound (CRLB), which describes a theoretical lower bound for the variance of an unbiased estimator (see [227] for the calculation).
By jointly reconstructing the basis material images, the noise correlations of the decomposed basis material line integrals are taken into account [386]. Introducing the vectors $\vec{A}_i = (A_i^0,\ A_i^1)$ and $\vec{\hat{A}}_i = (\hat{A}_i^0,\ \hat{A}_i^1)$, the data term can be written as:

$$-L(\vec{\alpha}^b) = \sum_i \left(\vec{A}_i - \vec{\hat{A}}_i \right) B_i \left(\vec{A}_i - \vec{\hat{A}}_i \right)^T, \tag{6.25}$$

where B_i is a 2×2 symmetric matrix approximating the inverse covariance matrix of the decomposed basis material line integrals. We refer the reader to [386] for the derivation of B_i and a more detailed description of the model.
For both joint and separate SIR of decomposed line integrals, we applied the Huber regularization penalty individually to each basis material image (compare eq. 6.7).

6.2.5 Numerical simulation

A low-dose spectral CT scan of a phantom consisting of an elliptic cylinder containing soft-tissue which is filled with several smaller bone cylinders was simulated. We assumed a density of 0.8 g/cm^3 for the bone cylinders, corresponding to cancellous bone. Figure 6.1 shows the corresponding soft tissue and bone basis material images. As can be seen from figure 6.1, we simulated several low contrast objects in the center of the phantom and in the center of the bottom-right bone cylinder by varying the soft tissue and bone density, respectively. We assumed a tungsten anode with an acceleration voltage of 120 kVp, a copper filter (thickness 0.2 mm) and a CdTe-based photon-counting detector (size 600×100 pixels) with thresholds set to 25 and 60 keV. Furthermore, pulse pile-up and scattering towards the detector were not considered. To simulate a spatially varying detector response, the thresholds for each pixel were drawn from a Gaussian distribution ($\sigma = 2$ keV) centered around the corresponding mean values (25 and 60 keV). We simulated 1200 projection measurements equally spaced between 0 and 360°. The tube loading was chosen such that $2 \cdot 10^4$ photons were registered on average per detector pixel and projection measurement without the phantom in the beam path (flatfield

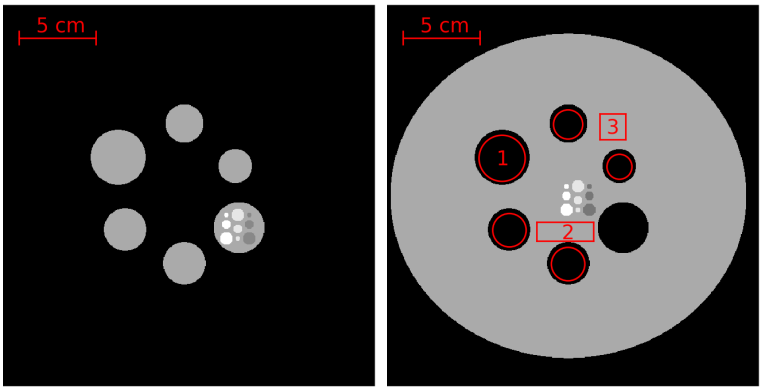

(a) (b)

Figure 6.1: Bone (a) and soft tissue (b) basis material image for the cylindrical phantom used for the numerical simulation. The regions (1), (2) and (3) were used for quantitative evaluations (see section 6.3).

intensity). Prior to the CT scan, we simulated a calibration measurement by placing homogeneous blocks of PMMA and PVC with three different thicknesses ($6.4, 35.2, 64.0$ mm for PVC and $25.0, 140.5, 256.0$ mm for PMMA) in the beam path. In total, 9 calibration measurements were acquired by taking projection measurements of all possible combinations of blocks (using one block of each basis material). The flatfield intensity was increased to $1 \cdot 10^9$ photons for the calibration measurement in order to suppress the effect of fitting errors due to noise in the calibration measurements. In real experiments, the calibration measurement is typically performed with strongly increased photon statistics compared to the CT scan because there are no dose considerations and the number of calibration measurements is much lower than the number of projections in the CT scan. To evaluate the accuracy of the calibrated model, we simulated a set of test measurements using the same photon statistics. In total, 64 test measurements were simulated which correspond to varying pathlengths for cortical bone (density 1.92 g/cm^3, 0-24 mm) and soft tissue ($25 - 370$ mm). Figure 6.2 shows an overview of the calibration measurements and the test measurements (converted to corresponding PVC/PMMA thicknesses). The arrows connecting the test measurements indicate the definition of the measurement index, which will be used for evaluation purposes.

6.2.6 Experimental measurement

A CT scan of a human knee was conducted with an experimental setup. The usage of the knee specimen for research projects was approved by the institutional review board. The donor had dedicated its body for educational and research purposes, and provided written informed consent prior to death, in compliance with local institutional and legislative requirements. The tube was operated at an acceleration voltage of 110 kVp and a CdTe-based photon-counting detector with 1536×128 pixels (XC-Flite FX1, XCounter AB, pixel size 100 μm $\times$ 100 μm) and thresholds set to 27 and 52 keV was used. In total, 1201 projections were taken and the tube loading was 131 mAs. Between the projection measurements, the specimen was randomly shifted by a few pixels, which is a common micro-CT technique for ring artifact suppression [387]. Before the CT scan, 9 calibration measurements were acquired by placing homogeneous blocks of PMMA (thicknesses $2, 32, 64$ mm) and PVC (thicknesses $1, 16, 32$ mm) in the beam path, using at most one block of PVC and PMMA for each calibration measurement.

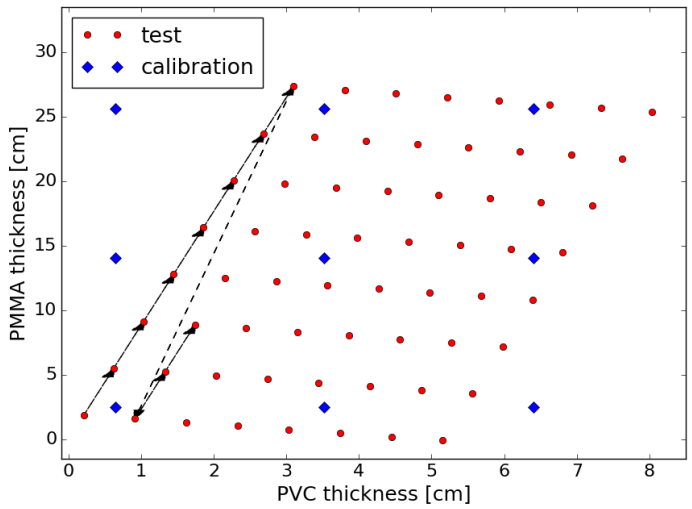

Figure 6.2: Overview of (equivalent) PMMA and PVC thicknesses for the simulated calibration and test measurements. The arrows connecting the test measurement points indicate the definition of the measurement index which is used for evaluation purposes.

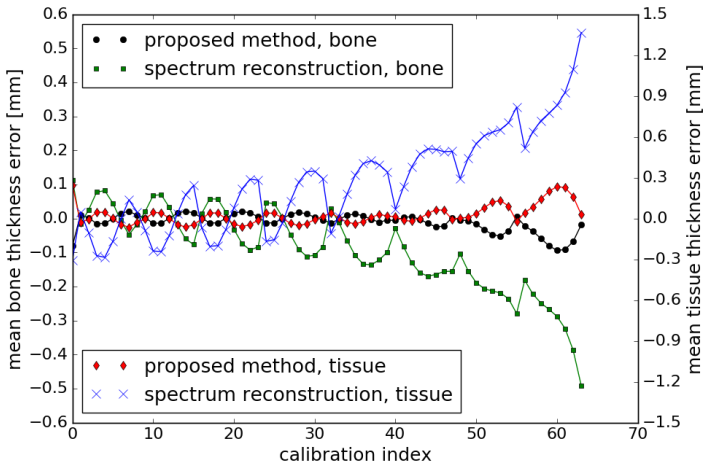

Figure 6.3: Comparison of the accuracy of the decomposed line integrals for the test measurements described in section 6.2.5. For both the proposed method and the spectrum reconstruction the same calibration measurement and the same number of fitting parameters (6) was used.

6.3 Results

To determine the accuracy of the projection-based material decomposition approach (compare eq. 6.5), the test measurements described in section 6.2.5 were decomposed into bone and soft tissue line integrals. Figure 6.3 shows the deviations from the correct line integrals averaged over all detector pixels.

We used three exponential terms (i.e. 6 fit parameters) to fit the semi-empirical forward model according to eq. 6.4 for all detector pixels and both energy bins. In general, the number of exponential terms used for fitting depends on the desired accuracy and the expected degree of beam hardening in the measurements (which in turn depends on the shape of the spectrum and the largest basis material line integrals in the sample). This dependency is illustrated in figure 6.4 where the average accuracy of the forward model (for the low energy bin) is plotted in dependence of the test measurement index. The accuracy of the forward model $\mathrm{acc}_{\mathrm{fwd}}$ was

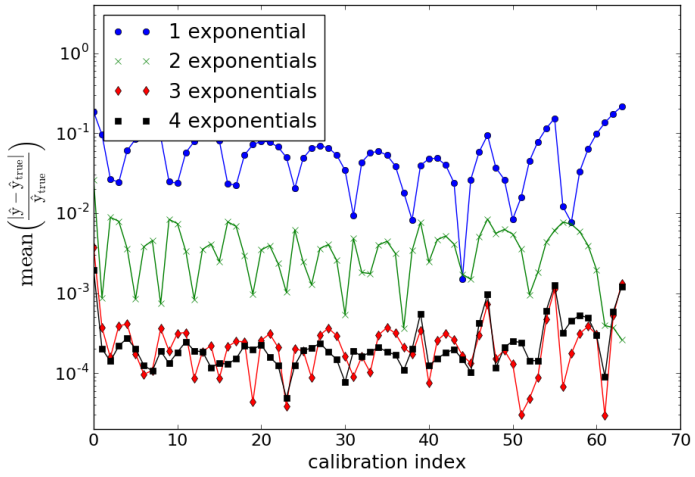

Figure 6.4: Accuracy of the fitted forward model for the low energy bin using different numbers of exponential terms. The accuracy was determined by calculating the mean relative difference between the photon counts predicted by the forward model and the expected photon counts calculated with eq. 6.2.

calculated as follows:

$$\mathrm{acc}_{\mathrm{fwd}} = \mathrm{mean}\left(\frac{|\hat{y} - \hat{y}_{true}|}{\hat{y}_{true}}\right), \tag{6.26}$$

where $\hat{y}$ indicates the photon counts predicted by the forward model and $\hat{y}_{true}$ represents the expected counts calculated according to eq. 6.2. The mean is taken over all detector pixels.

The accuracy of the proposed semi-empirical forward model is compared to the method of 'spectrum reconstruction', which has already been investigated by different authors [388, 389]. Mathematically, the corresponding forward model looks similar to eq. 6.3 but the spectrum is divided into uniform energy bins, i.e. the energies $\mathcal{E}_{ri}$ have the same predefined values for each detector pixel and only the coefficients C_{ri} are fitted with the help of calibration measurements. The same calibration measurement and the same number of fitting parameters (6) were used for estimating the spectrum. We chose to use the expectation-maximization algorithm for reconstructing the spectrum because of its robustness and the possibility to incorporate constraints such as non-negativity of the coefficients [388]. After spectrum estimation, a maximum-likelihood decomposition algorithm (see eq. 6.5) was applied to the simulated test measurements. Figure 6.3 compares the resulting average bias of the decomposed line integrals to the bias of the proposed semi-empirical forward model. The average absolute decomposition error (of all calibration points) was 117 μm and 353 μm for the bone and soft tissue line integrals, respectively. The proposed semi-empirical model reduced the average error to 20 μm and 52 μm for the bone and soft tissue line integrals, respectively.

Figure 6.5 (a) and (b) show the influence of the photon statistics on the bias and standard deviation of our projection-based material decomposition approach. To eliminate the effect of threshold dispersion, a single detector pixel with thresholds set to 25 and 60 keV was used for this experiment. After fitting the semi-empirical forward model with the help of the simulated calibration measurement (see section 6.2.5), spectral CT measurements with 10.2 mm of cortical bone and 172.9 mm of soft tissue placed in the beam path were simulated. The flatfield intensity was varied from 10^8 to 10^4 photons and 10^6 different noise realizations were calculated for each flatfield intensity. Figure 6.5 (a) shows the mean error of the projection-based material decomposition, whereas figure 6.5 (b) compares the experimental standard deviation for both basis materials with the CRLB.

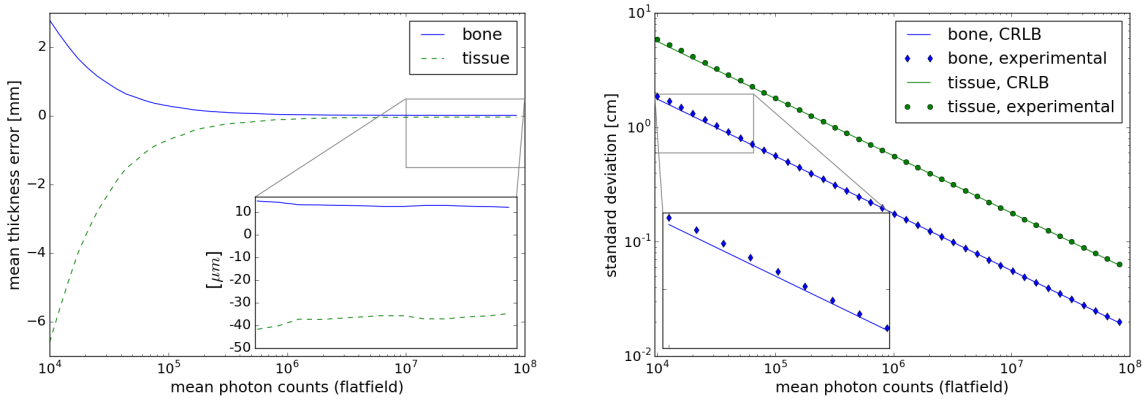

(a) (b)

Figure 6.5: Bias (a) and standard deviation (b) of the projection-based material decomposition approach in dependence of the photon statistics. For each flatfield intensity, the bias and variance were computed by decomposing 10^6 noise realizations of a simulated spectral measurement with 10.2 mm of cortical bone and 172.9 mm of soft tissue placed in the beam path. The Cramér-Rao lower bound (CRLB) was calculated according to the deviation given in [227].

The left column of figure 6.6 shows the soft tissue (a) and bone image (d) image obtained with projection-based decomposition followed by a filtered backprojection (FBP) for the simulated spectral CT scan. Basis material images reconstructed with the SP-SIR,EW algorithm are displayed in the middle column of figure 6.6. The right column shows the same basis material images reconstructed with the proposed JI-SIR algorithm. The regularization strengths λ_b were tuned to achieve 20 % of the FBP noise level (in region (3), compare figure 6.1) for both basis material images and both SIR algorithms. Furthermore, the same γ parameter for the Huber regularization was used for both iterative reconstruction methods.

In figure 6.7, the bias of different reconstruction methods is investigated. The images show the difference between the mean of all 100 horizontal slices of the reconstructed basis material images and the horizontal slices of the corresponding ground truth images. Since the phantom is constant in vertical direction, all horizontal ground truth slices are equal. The mean of all slices was taken to suppress noise which might conceal the systematic bias in the reconstructed images. The first two columns of figure 6.7 show the bias of projection-based material decomposition followed by FBP reconstruction. A noise-free spectral CT scan was used for the first column, whereas the noisy CT scan described in section 6.2.5 was used to generate the difference images in the second column. The bias of five different statistical iterative reconstruction algorithms was investigated. Besides the proposed joint reconstruction algorithm which works directly on the spectral count data (abbreviated by "JI-SIR" in the following), we study the bias of joint (JP-SIR) and separate SIR (SP-SIR) algorithms working on decomposed basis material line integrals (compare section 6.2.4). Here we differentiate between two different strategies to determine the statistical weights for SIR of decomposed line integrals ($1/\sigma_{i,b}$ for SP-SIR and B_i for JP-SIR, see section 6.2.4). "Estimated weights" (abbreviated EW) refers to approximating the variances (for SP-SIR) or the covariance matrices (for JP-SIR) with the noisy simulated CT scan, whereas as noise-free simulation is used to calculate the "true weights" (TW). The third and fourth column of figure 6.7 show the difference images for SP-SIR, EW and JP-SIR, EW, respectively. The corresponding difference images using true weights are not shown because it turned out that they are very similar to the difference images of the noisy FBP reconstructions (second column). The last column shows the bias of the JI-SIR algorithm. The same γ parameter for the Huber regularization (see eq. 6.7) was used for all iterative reconstructions. In order to make all difference images comparable and avoid regularization induced bias, the regularization strengths

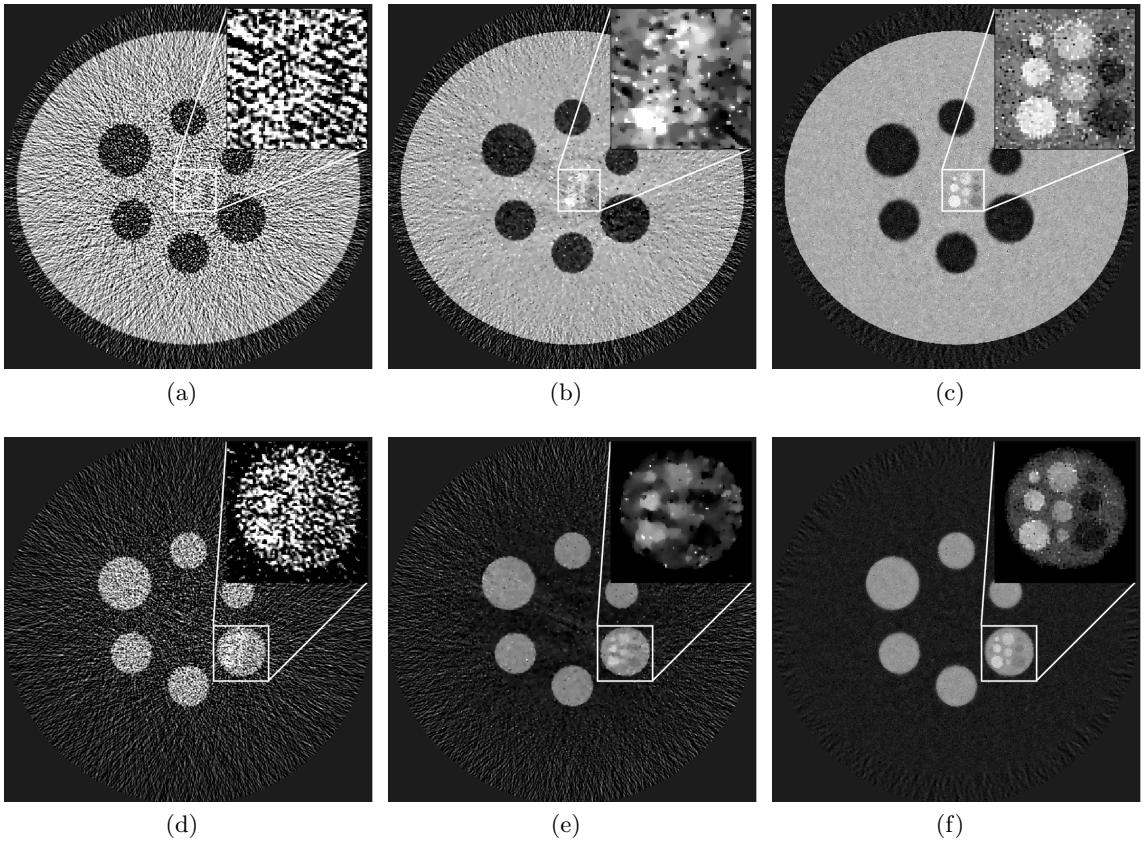

Figure 6.6: Reconstructed basis material images (soft tissue, top row and bone, bottom row) for the simulated spectral CT scan. The left and middle column show the soft tissue and bone images obtained with projection-based decomposition followed by FBP (a,d) and separate SIR (b,e), whereas the right column shows the corresponding basis material images obtained with the proposed joint SIR algorithm (c,f). For each basis material, the same windows were chosen for all reconstruction methods.

were tuned to achieve the FBP noise level for all iterative reconstructions. We measured the noise level by calculating the standard deviation in a homogeneous region of the basis material images (region 3 in figure 6.1). Table 6.2 presents the average basis material density bias for the aforementioned reconstruction methods. The densities were determined in two different regions corresponding to the centers of the bone cylinders and the center of the soft tissue cylinder (compare regions 1 and 2 in figure 6.1).

In figure 6.8, we compare the convergence of two different versions of the optimization algorithm presented in section 6.2.3 for the simulated spectral CT scan. The conventional version of the algorithm does not used ordered subsets and Nesterov acceleration. Furthermore, the Hessian matrix of the surrogate function $\varphi\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$ is computed at each iteration step. The accelerated version of the algorithm uses six subsets (each containing 200 angles), Nesterov's acceleration technique and the Hessian matrix of $\varphi\left(\vec{\alpha};\vec{\alpha}^{(n)}\right)$ is precomputed (compare eq. 6.21). The number of subsets was chosen such that the computational time for the gradient of the data term and the regularization term were approximately equal. Both algorithms were initialized with FBP images and subsequently run for 120 (sub)iterations. The graph shows the normalized l_2-distance to a converged reference (obtained by performing 3000 iterations of the conventional algorithm) in dependence of the runtime. We calculated the normalized l_2-distances for the basis material

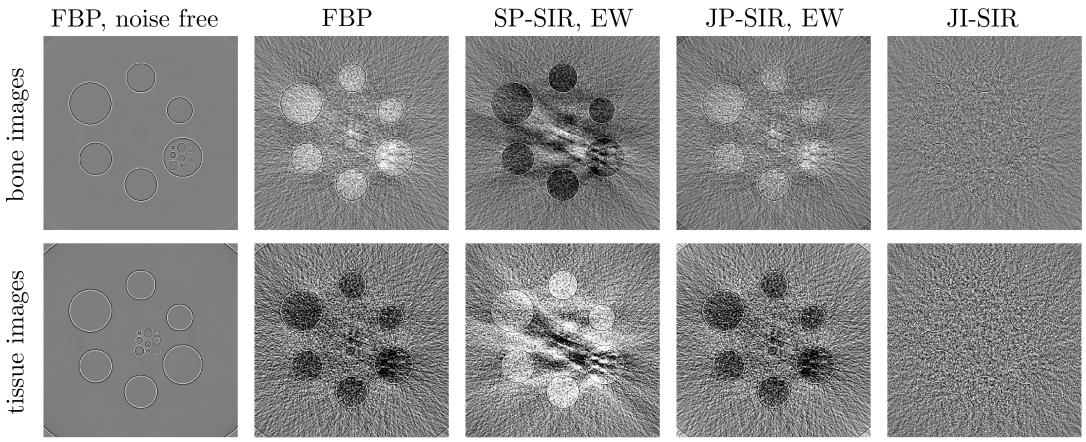

Figure 6.7: Difference of reconstructed basis material images to the ground truth for various reconstruction methods (zoomed into the center of the phantom). In order to suppress noise, the images were generated by averaging all reconstructed horizontal slices and then subtracting the ground truth. Besides projection-based material decomposition followed by FBP, the difference images for separate and joint SIR of decomposed line integrals using estimated statistical weights (SP-SIR, EW and JP-SIR, EW, respectively) are shown. The last column displays the bias of the proposed joint SIR algorithm (JI-SIR).

Table 6.2: Evaluation of the bias of different reconstruction methods for the numerical simulation. The table shows the difference to the ground truth bone and soft tissue densities for two different regions corresponding to the bone cylinders (region 1) and an area slightly below the center of the simulated sample (region 2, see figure 6.1).

density bias [mg/cm^3]	bone, region 1	tissue, region 1	bone, region 2	tissue, region 2
FBP, noise-free	+0.61	-0.94	+0.09	-0.18
FBP	+59.44	-81.25	19.90	-26.81
SP-SIR, EW	-76.17	+84.73	-6.16	-6.32
SP-SIR, TW	+55.31	-75.71	+17.06	-25.72
JP-SIR, EW	+30.02	-74.78	+6.85	-16.90
JP-SIR, TW	+56.81	-82.21	+20.32	-26.83
JI-SIR	-0.25	+0.86	+0.23	+0.33
ground truth density [mg/cm^3]	800	0	0	1060

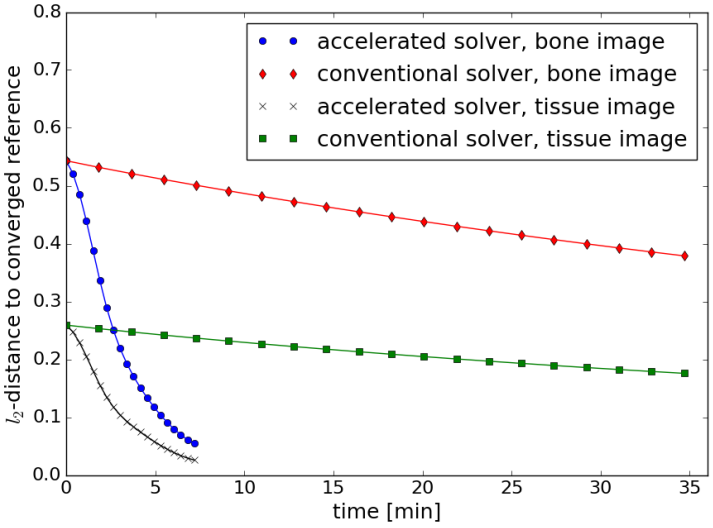

Figure 6.8: Convergence of two different version of the optimization algorithm presented in section 6.2.3. The graph shows the normalized l_2-distance to a converged reference in dependence of the runtime.

images ($l_2^{b,(n)}$ iteration n) as:

$$l_2^{b,(n)} = \sqrt{\frac{\sum_{j=1}^{N}\left(\alpha_j^{b,(n)} - \alpha_j^{b,(\mathrm{ref})}\right)^2}{\sum_{j=1}^{N}\left(\alpha_j^{b,(ref)}\right)^2}}, \tag{6.27}$$

where $\alpha_j^{b,(\mathrm{ref})}$ indicates the reference images.

The results for the experimental measurement are shown in figure 6.9. In the left column row of figure 6.9, photo- (a) and Compton effect (d) basis material images obtained with projection-based material decomposition followed by FBP are presented. Basis material images reconstructed with the SP-SIR,EW algorithm are displayed in the middle column (b,e). The right column of figure 6.9 shows the basis material images reconstructed with our joint SIR algorithm (c,f). The regularization parameters were tuned according to visual appearance. We used three exponential terms to fit the semi-empirical forward model for both energy bins with the help of the calibration measurements.

6.4 Discussion

Figure 6.3 (a) shows that the proposed projection-based decomposition algorithm is slightly biased ($<$ 100 μm average error of the decomposed line integrals). This bias results from the semi-empirical approximation of the true spectral forward model (compare eq. 6.3 and 6.2). The small number of fit parameters does not allow to fit the spectral forward model perfectly. Comparable systematic errors have been reported for other projection-based material decomposition methods which rely on calibration measurements [244, 225]. In an experimental measurement, other bias inducing effects (e.g. scattering, detector drifts) will typically dominate over the bias of the semi-empirical forward model.

A low dimensional representation of the spectral forward model is advantageous because the number of calibration measurements can be reduced. Furthermore, a small number of fit parameters minimizes the risk of overfitting due to inconsistencies in the calibration measurements (e.g. noise). From figure 6.3 and the calculation of the mean absolute errors, it is evident that the method of spectrum reconstruction (combined with maximum likelihood decomposition)

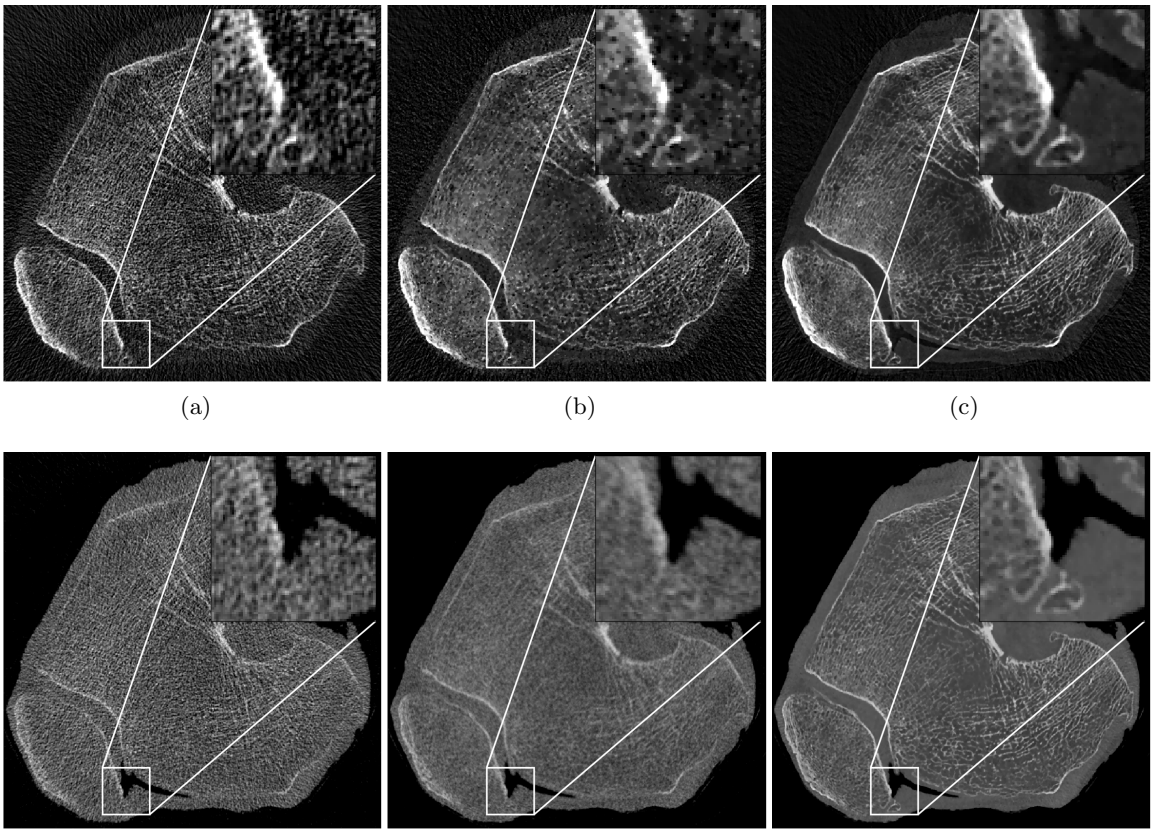

Figure 6.9: Axial slices of the reconstructed basis material images (photoeffect, top row and Compton effect, bottom row) for the experimental measurement of a human knee. The left and middle column of figure 6.9 show basis material images obtained with projection-based material decomposition followed by FBP (a,d) and separate SIR (b,e), respectively. In the right column, photo- (c) and Compton effect (f) basis material images reconstructed with our joint SIR algorithm are presented. For each basis material, the same windows were chosen for all reconstruction methods.

leads to a significantly larger decomposition bias if the same number of fitting parameters and the same calibration measurements are used. This becomes particularly obvious for the test measurements with the largest basis material thicknesses and thus the most beam hardening.
In case of the numerical simulation, we found that the best results are obtained if the number of calibration measurements is slightly larger than the number of fit parameters. Adding more calibration measurements does not considerably improve the accuracy of the fitted model. Figure 6.4 illustrates the influence of the number of exponential terms on the accuracy of the semi-empirical forward model. With two exponential functions, the error of the predicted photons counts can already be reduced below 1 %. Adding another term leads to errors of approximately 0.02 %, whereas using four exponential terms does not reduce the average absolute error considerably. A possible reason is overfitting to noise due to the increased number of fitting parameters. Furthermore, adding more exponential terms leads to diminishing returns in this case, since the spectral transmission can already be accurately predicted with three exponential terms.
In conclusion, the proposed semi-empirical forward model approximates the spectral transmission very accurately while using a small number of fit parameters and calibration measurements. These properties are desirable from an experimental perspective since they enable an accurate material decomposition with low experimental overhead.
Although the simulation includes detector imperfections (charge-sharing, electronic noise, escape and fluorescence peaks) to model a realistic spectral response of the photon-counting detector, the spatial effect of these imperfections has been neglected. In practice, some of these effects could lead to an inferior performance of the proposed semi-empirical forward model, since it only includes photon counts directly related to the path of the primary radiation and neglects spatial noise correlations. However, most of the aforementioned detector imperfections (e.g. charge sharing) only distort the expected number of photon counts locally, i.e. affecting a few neighboring pixels. For the calibration, the distortions average out if the calibration phantom is homogeneous over the detector area. If the measured object is not homogeneous, effects like charge sharing or fluorescence peaks should lead to a blurring of the decomposed line integrals. Furthermore, pulse pile-up and scattering from the sample towards the detector are not included in the proposed model. If these effects cannot be neglected, the semi-empirical model could be combined with correction algorithms for pile-up and scattering.
The analysis of the bias and variance of the projection-based material decomposition approach (figure 6.5) shows that the decomposed line integrals are negligibly biased at high and moderate photon statistics. However they become considerably biased if the flatfield intensity is decreased below 10^5 photon counts. The bone line integrals are more and more overestimated while the soft tissue line integrals are underestimated. In the case of low photon statistics, this noise-induced bias is more than two orders of magnitudes larger than the bias of the semi-empirical forward model. The reason for this statistical bias is the propagation of noise through a nonlinear material decomposition operation (compare also the discussion in section 2.5). In general, calculating the expectation value and applying a nonlinear function $g(\cdot)$ are not interchangeable:

$$E(g(X)) \neq g(E(X)), \tag{6.28}$$

where X is a random variable and $E(\cdot)$ denotes the expectation value. Figure 6.5 (a) indicates that statistical bias might be a major limiting factor for the quantitative accuracy of material decomposition algorithms in the case of low-dose spectral CT measurements. Biased decomposition results at low photons statistics have also been reported for a polynomial-based material decomposition methods which does not rely on maximum-likelihood estimation [246]. Figure 6.5 (b) shows that the standard deviation of the decomposed line integrals is essentially equal to the theoretical optimum (CRLB) for a large range of flatfield intensities. At very low photon counts ($< 10^5$), the experimental standard deviation lies slightly above the CRLB.
Similar to the projection-based analysis, figure 6.7 and table 6.2 present an evaluation of the

noise-induced bias for the simulated spectral CT scan. In case of the noise-free FBP, the reconstructed basis material images only noticeably differ from the ground truth at the edges of the phantom. These differences can be attributed to the backprojection process. As expected from the projection-based analysis, the quantitative values in table 6.7 support the hypothesis that the FBP reconstructions from noise-free data are almost unbiased. The FBP reconstructions of the simulated low-dose scan overestimate the bone density, while the soft tissue density is underestimated. This effect is particularly apparent in regions where mostly large basis material line integrals contribute to the reconstructed voxel values (i.e. the bone cylinders and the center of the soft tissue cylinder). Since the corresponding spectral CT measurements have low photon statistics, the statistical bias is larger. The bias for the SP-SIR, TW and the JP-SIR, TW algorithm is comparable to the bias of the FBP reconstructions (see table 6.2). This is reasonable because these reconstruction algorithms work on basis material line integrals with the same bias. However, the FBP reconstruction uses a uniform weighting scheme while line integrals are weighted inversely proportional to the variance predicted by the CRLB and the inverse covariance matrix for the SP-SIR and the JP-SIR algorithm, respectively. Calculating the statistical weights (which is a nonlinear operation) with noisy projection data results in systematically biased weights. Consequently, the bias of the weight estimation interacts with the bias of the projection-based material decomposition for the SP-SIR,EW and the JP-SIR,EW algorithm. Assuming the same expected intensities, calculating the statistical weights with noisy data falsely gives more weight to measurements with higher photon counts. In the case of the SP-SIR,EW algorithm, the bias of the projection-based decomposition is overcompensated, leading to underestimated bone densities and overestimated soft tissue densities in region (1) (see figure 6.7). Using correlated statistical weights seems to mitigate this effect, resulting in reduced bias for the JP-SIR,EW algorithm. Figure 6.7 and table 6.2 show that the bias of the proposed JI-SIR algorithm is comparable to the bias of the noise-free FBP reconstruction and is therefore almost two orders of magnitude smaller than the bias of the other reconstruction algorithms. The reason for this is that, unlike all other methods, material decomposition and image reconstruction are performed jointly. By combining image reconstruction and material decomposition, the basis material contributions for a certain voxel are determined by using the statistical information from all projection measurements involving this voxel. Furthermore, the regularization term couples neighboring voxels, which effectively increases the number of spectral projection measurements involved in determining the basis material contributions for a certain voxel. Projection-based material decomposition methods do not use information from other projections to calculate the basis material line integrals, which leads to more bias. In a second step, information from different basis material line integrals is used for image reconstruction, but the bias introduced in the decomposition step propagates into the basis material images.

In summary, two-step methods which separate the process of material decomposition and image reconstruction, are susceptible to two different sources of statistical bias (projection-based decomposition and calculating the statistical weights). Our results show that the proposed joint statistical iterative reconstruction algorithm improves the quantitative accuracy for low-dose CT scans by avoiding these sources of statistical bias.

Figures 6.6 and 6.9 compare the image quality of basis material images reconstructed with projection-based decomposition followed by FBP or SIR (i.e. SP-SIR,EW) and the proposed joint reconstruction (JI-SIR). In case of the numerical simulation, the low-contrast density variations of the basis material images can be recovered with the JI-SIR algorithm (see figure 6.6 (c,f)) while they vanish in the noise for the FBP reconstructions. To achieve the same noise level as the JI-SIR algorithm (20 % of the FBP noise level), the regularization strength has to be chosen comparatively high for the SP-SIR,EW reconstruction. This leads to artifacts and blurring of the low contrast insets (see figure 6.6 (b,e)).

Figure 6.9 demonstrates that the joint SIR algorithm achieves highly improved image quality compared to projection-based material decomposition combined with FBP or SIR in the case

of a clinically relevant sample. Compared to the FBP reconstructions (see figure 6.9 (a,d)) the JI-SIR algorithm strongly reduces the noise level in the basis material images while preserving small structural details (see figure 6.9 (c,f)). Despite increased noise levels compared to the JI-SIR algorithm, some small structures are lost in the basis material images reconstructed with the SP-SIR,EW algorithm (see figure 6.9 (b,e)).
The comparison of basis material images reconstructed with projection-based material decomposition followed by FBP or SIR and the JI-SIR algorithm shows that better tradeoffs between noise and resolution can be achieved by accounting for the noise correlations between basis material images with joint statistical iterative reconstruction techniques. Similar findings have been reported in the reference for the JP-SIR algorithm [386].
A potential problem with applying Huber regularization separately to both images is regularization crosstalk. Suppose regularization is applied to basis material image A while no regularization is applied to basis material B. If the regularization smooths out edges (or reduces the edge height) in image A, the corresponding edges will reappear with different amplitude in image B, which leads to undesired image appearance and regularization induced bias. This effect is caused by the coupling of the basis material images through the spectral forward model and the data term. The edges in image B are adjusted in order to optimize the agreement between the measured data and the spectral forward model, which corresponds to minimizing the log-likelihood of the data term. Experimentally, we found that the crosstalk effect becomes noticeable if the strength of the regularization is chosen comparatively high to achieve a strong noise suppression. Recently, local response prediction techniques have been generalized to spectral CT [390]. This allows the suppression of crosstalk artifacts by locally tuning the regularization strength for the basis material images.
A future project will be the combination of joint regularization techniques and the data model presented in this work. We believe that the regularization crosstalk can be suppressed and the image quality can be further improved by using a regularization term which encourages the same structures and edges in the basis material images. This could for example be accomplished by adapting the dictionary based regularization strategy that was presented in chaper 5 to 3D CT reconstruction.
The comparison of the conventional and the accelerated optimization algorithm in figure 6.8 shows that the reconstruction time can be strongly reduced by employing the acceleration techniques presented in section 6.2.3. The total runtime for 120 (sub)iterations was 36.2 and 7.5 minutes for the conventional and accelerated solver, respectively. By using ordered subsets and precomputing the Hessian matrices, the computational time for one subiteration was reduced to approximately one fifth of the conventional algorithm. The final normalized l_2-distances for the bone and soft-tissue images were (0.37, 0.17) and (0.049, 0.024) for the conventional and accelerated algorithm, respectively. Consequently, the convergence per iteration is highly accelerated by employing Nesterov's momentum technique. Reducing the reconstruction time of iterative algorithms is an important goal for practical applications since long reconstruction times remain a major challenge of SIR methods.

6.5 Conclusion

We have developed a new statistical iterative basis material image reconstruction algorithm which is based on a semi-empirical forward model. Since the spectral forward model is adjusted with the help of calibration measurements, detailed knowledge about the source spectrum and the detector response is not required. Furthermore, inter-pixel variations of the detector response can be incorporated, which is crucial for spectral CT in combination with photon counting detectors. The small number of fitting parameters allows to reduce the number of calibration measurements, which is desirable from an experimental perspective. By employing an efficient

optimization algorithm based on separable surrogate functions, we were able to reduce the runtime of the reconstruction algorithm and accelerate convergence. Reducing the computational cost of SIR techniques is an important goal for practical applications, especially if large datasets are acquired. The low-dose simulation of a spectral CT scan indicates that statistical bias can be a major limiting factor for the quantitative accuracy of basis material image reconstruction. Contrary to reconstruction techniques which separate material decomposition and image reconstruction, our joint SIR approach shows almost no statistical bias. Moreover, the joint SIR algorithm leads to improved image quality compared to separate SIR because the of the joint statistical noise model which includes noise correlations between the basis material images. This was demonstrated by comparing reconstructed basis material images from a simulated CT scan and an experimental measurement of a human knee. However, if the regularization strength is comparatively high, which is particularly relevant for low-dose measurements, regularization crosstalk might lead to undesired image appearance. Regularizing the basis material images separately is probably suboptimal in combination with the joint SIR approach, which couples these images through the spectral forward model. In the future, we plan to combine our algorithm with joint regularization methods, which encourage common structures in the basis material images.

7 Spectral differential phase-contrast X-ray radiography

In this chapter, we investigate the combination of spectral imaging and grating-based DPC imaging in projection space. We propose to combine these two emerging X-ray imaging technologies by measuring stepping curves in a standard grating-based differential phase contrast setup at several different effective spectra. Although spectral differential phase-contrast (SDPC) imaging could in theory be realized by all common dual energy techniques (compare section 2.7.1), PCD technology appear to be the most promising approach. This is because PCDs provide spatially and temporally registered energy-resolved stepping curve measurements, which simplifies data analysis and processing. As both spectral and DPC imaging methods can determine the (projected) electron density, we devise a novel material decomposition algorithm that uses the spectral and the phase contrast information simultaneously. Numerical experiments show that the combination of these two imaging techniques benefits from the strengths of the individual methods while the weaknesses are mitigated. Due to the additional phase shift information, the noise level of the basis material images is strongly reduced. Compared with conventional DPC imaging, the additional spectral information can eliminate beam hardening as well as phase wrapping artifacts (see section 2.8.3) and suppresses long-range noise correlations. The noise characteristics of SDPC imaging (in comparison with spectral and DPC X-ray imaging) are analyzed in more detail in chapter 8. In chapter 9, we extend the model for projection-based SDPC imaging to 3D CT reconstruction. Motivated by the results for spectral imaging (see chapter 6), we investigate the benefits of a one-step SIR algorithm for SDPC CT imaging.
The text and the figures in this chapter are taken from the following publication of the author:

- Mechlem, Korbinian, et al. "Spectral differential phase contrast x-ray radiography." IEEE Transactions on Medical Imaging (2019).

Compared to the original publication, the text has been slightly modified to achieve a better integration into this doctoral thesis.

7.1 Introduction

It has been recently demonstrated that material selective images can in principle be calculated from DPC measurements by using the attenuation and phase shift information [20, 21]. A possible advantage of this approach compared with spectral X-ray imaging is the additional information provided by the dark-field image. Moreover, depending on the setup parameters, DPC imaging can achieve a highly improved contrast-to-noise ratio (CNR) compared to attenuation-based imaging [15, 16, 17, 18, 19]. However, partly due to the noise characeristics of DPC imaging (see chapter 8), it is unclear whether the observed CNR advantage in laboratory experiments could be tranferred to medical imaging applications [391]. Obtaining quantitaitvely accurate material selective images is another challenge for DPC imaging. As discussed in section 2.8.3, the polychromatic X-ray spectrum causes beam hardening artifacts in both the phase and the dark-field channel [23, 24]. Furthermore, the phase shift cannot be unambiguously determined. This phase-wrapping problem is another important source of artifacts for grating-based phase-contrast imaging [25, 26]. Spectral X-ray imaging techniques provide basis material images with high quantitative accuracy, but noise amplification during material decomposition remains a

major challenge. The combination of spectral and DPC imaging could therefore mitigate some of the aforementioned shortcomings of the individual imaging techniques. Several approaches for combining spectral and differential phase-contrast imaging have already been published. The combination of these two emerging X-ray imaging technologies can be realized by measuring stepping curves in a standard grating-based differential phase contrast setup at several different effective spectra (e.g., by using a photon counting detector). Epple et. al used spectral phase-contrast measurements to mitigate phase-wrapping artifacts [392, 393]. Moreover, it has been demonstrated that a weighted addition of phase- [394] or dark-field images [395] extracted from different energy bins of a photon counting detector can improve the CNR. In this chapter, we investigate a new algorithm for combining grating-based DPC radiography and spectral radiography. We extract two basis material images and a dark-field image by simultaneously using the spectral and the phase contrast information. The connection between spectral and DPC imaging is established via the projected electron density that can be determined with both imaging techniques. The fact that both the (energy-dependent) attenuation and the phase shift depend on the projected electron density has already been used for coded-aperture imaging [396]. Numerical simulations show that the combination of these two imaging techniques benefits from the strengths of the individual methods while the weaknesses are mitigated. Quantitatively accurate basis material images without beam hardening or phase-wrapping artifacts are obtained. Furthermore, the additional phase shift information leads to highly reduced basis material image noise levels compared to conventional spectral X-ray imaging. In this chapter, we concentrate on human thorax radiography as a potential medical application of spectral differential phase-contrast (SDPC) radiography. Spectral radiography has two important advantages compared to conventional radiography for chest imaging: First, the decomposition into basis material images facilitates the assessment of small lung nodules by eliminating "anatomical noise" (i.e., anatomical structures such as the ribs superimposed on the projection image) [397]. Second, basis material decomposition allows the detection of calcifications within a lung nodule, which is an important indicator of benignancy [398]. However, due to the reduced signal-to-noise ratio of the basis material images, quantum noise (instead of "anatomical noise") is often a major limiting factor for spectral chest radiography [397]. We demonstrate how the reduced noise levels and the additional information about the microstructure (dark-field image) provided by SDPC radiography could be beneficial for clinical diagnosis.

7.2 Methods

7.2.1 Spectral phase-contrast model

We begin by explaining maximum-likelihood projection-based material decomposition for spectral imaging before extending the model to incorporate energy-resolved phase-contrast measurements. Disregarding materials with K-edges in the relevant energy range, the energy-dependent attenuation of any material is well described by a linear combination of two basis materials:

$$\mu(E) = A_1 f_1(E) + A_2 f_2(E), \tag{7.1}$$

where A_α is the line integral of material α and $f_\alpha(E)$ describes the the energy-dependent attenuation of basis material α. This is a consequence of the fact that there are only three interaction mechanisms (photoeffect, Compton effect and Rayleigh-scattering) in the energy range typically used for medical imaging ($\approx 20 - 140\,\mathrm{keV}$) and that the attenuation cross section is dominated by Compton-scattering and photoelectric absorption. Therefore, we focus on two basis materials in the following, however a generalization to more basis materials is straightforward. With the approximation of eq. 7.1, the expected number of photon counts $\hat{y}_i^s$ for detector pixel i and spectrum s is given by:

$$\hat{y}_i^s = \int_0^\infty S(E) R^s(E) e^{-A_1^i f_1(E) - A_2^i f_2(E)} dE, \tag{7.2}$$

where $S(E)$ is the source spectrum and $R^s(E)$ describes the detector response, i.e., the probability that a photon with energy E is detected by energy bin s of the photon-counting detector. Maximum-likelihood decomposition is performed by finding the basis material line integrals $\vec{A}_1 = (A_1^1, ..., A_1^N)^T$, $\vec{A}_2 = (A_2^1, ..., A_2^N)^T$ that maximize the probability of the measured data. Assuming Poisson statistics, this corresponds to minimizing the following negative log-likelihood function [399, 352, 236]:

$$-L(\vec{A}_1, \vec{A}_2) = \sum_i \sum_s \hat{y}_i^s - y_i^s \ln(\hat{y}_i^s), \tag{7.3}$$

where y_i^s denotes the measured photon counts for detector pixel i and effective spectrum s. This optimization problem can be solved separately for each detector pixel.
For conventional phase-contrast imaging, a stepping curve is acquired by shifting one of the gratings. The polychromatic spectrum is neglected and the expected intensity for the stepping curve measurements is approximated as [278]:

$$\hat{y}_i^r = be^{-\mu_i}\left(1 + Ve^{-\epsilon_i}\cos(\phi^r + \Delta\phi_i)\right), \tag{7.4}$$

where $\hat{y}_i^r$ and ϕ^r are the expected number of photons and the reference phase of the stepping curve (for step r), respectively. The quantities μ_i, $\Delta\phi_i$ and ϵ_i describe the attenuation of the object, the phase shift of the stepping curve and the visibility reduction (dark-field signal), respectively. The reference intensity measured with an open beam is given by b and V is the visibility of the reference stepping curve.

Taking eq. 7.2 and eq. 7.4 into account, we propose the following forward model for spectral differential phase-contrast (SDPC) measurements:

$$\hat{y}_i^{rs} = \int_0^\infty S(E)R^s(E)e^{-A_1^i f_1(E) - A_2^i f_2(E)}\left[1 + V(E)e^{-d_\epsilon^i f_\epsilon(E)}\cos(\phi^r(E) + \Delta\phi_i(E))\right]dE, \tag{7.5}$$

where $\hat{y}_i^{rs}$ is the expected number of photon counts for detector pixel i, stepping position r and energy bin s. Compared to eq. 7.2, we have introduced a third, "dark-field basis material" that is responsible for the visibility reduction of the stepping curve, i.e., for the dark-field signal. According to Strobl [294], the visibility reduction is given by:

$$\frac{V_s(\xi)}{V(\xi)} = e^{\int_{\text{path}} \sigma(E)(G(\xi)-1)dz}, \tag{7.6}$$

where V_s/V is the ratio of the visibilities with and without the sample in the beam path, $\sigma(E) \propto \frac{1}{E^2}$ is the macroscopic scattering cross section and $G(\cdot)$ is the autocorrelation function of the scattering length density [295]. The scattering length density is directly proportional to the electron density far away from the absorption edges (which is normally the case for medical imaging). If the sample is placed between the phase shifting grating (G1) and the detector grating (G2), the correlation length ξ is calculated as:

$$\xi = \frac{hc}{E}\frac{d_{S,G2}}{p_2}, \tag{7.7}$$

where h is Planck's constant, c is the speed of light, $d_{S,G2}$ is the distance between the sample and the G2 grating and p_2 is the corresponding grating period. Neglecting the spatial extend of the object along the projection direction (i.e., a constant $d_{S,G2}$ is assumed for the whole object) and assuming a homogeneous dark-field generating material, eq. 7.6 can be simplified to:

$$\frac{V_s(E)}{V(E)} = e^{d_\epsilon \sigma(E)(G(E)-1)}, \tag{7.8}$$

where d_ϵ represents the pathlength through the microstructured material. Comparing eq. 7.6 with the SDPC forward model (eq. 7.5), $f_\epsilon(E)$ can be identified with $\sigma(E)(1 - G(\xi))$. Writing

$\sigma(E)(G(\xi)-1)$ only as a function of the energy corresponds to neglecting the spatial extend of the object in the projection direction since a constant $d_{S,G2}$ has to be assumed for the whole object. With this identification, f_ϵ(E) describes the visibility reduction per unit length as a function of the photon energy.

The phase shift $\Delta\phi(E)$ is proportional to the gradient of the projected electron density $\tilde{\rho}_e(x,y)$ perpendicular to the grating bars [291, 290]:

$$\Delta\phi_i(E) = \frac{q}{E^2}\frac{\partial}{\partial x}\tilde{\rho}_e(x_i,y_i), \tag{7.9}$$

where q is a proportionality factor depending on natural constants and the interferometer properties. Assuming that the object is placed between G1 and G2, q is given by [275]:

$$q = \frac{M r_e h^2 c^2 d_{G1,G2}}{p_2}\left(1-\frac{d_{S,G1}}{d_{G1,G2}}\right), \tag{7.10}$$

where r_e is the classical electron radius. The magnification M has to be taken into account because the projected electron density profile that is measured at the detector plane is magnified by a factor of M compared to the sample plane. Similar to the dark-field model, the spatial extend of the object in beam direction has to be neglected and a fixed mean value for $d_{S,G1}$ is assumed in the SDPC forward model. It is worthwhile to note that in a possible extension of the method to 3D statistical iterative CT reconstruction, the position-dependency of the dark-field and the phase shift along the optical axis could be taken into account [400]. In eq. 7.9, the spatial coordinates of detector pixel i in the detector plane are given by x_i and y_i. The projected electron density $\tilde{\rho}_e^i = \tilde{\rho}_e(x_i,y_i)$ is obtained by integrating the three-dimensional electron density distribution $\rho_e(x,y,z)$ along the projection direction:

$$\tilde{\rho}_e(x,y) = \int \rho_e(x,y,z)dz. \tag{7.11}$$

The key idea for connecting spectral and phase-contrast X-ray imaging, and for eliminating the projected electron density as optimization variable, is expressing the projected electron density as the sum of the projected basis material electron densities:

$$\tilde{\rho}_e^i = A_1^i\rho_e(M_1) + A_2^i\rho_e(M_2), \tag{7.12}$$

where $\rho_e(M_1)$ and $\rho_e(M_2)$ represent the electron densities of the basis materials. Similarly to the standard spectral imaging approximation of modeling the energy-dependent attenuation of any material by a linear combination of the basis materials, the approximation of eq. 7.12 works well in the relevant energy range $(20-140\,\mathrm{keV})$.

In analogy to eq. 7.3, the basis material line integrals can be reconstructed by minimizing the negative log-likelihood function:

$$-\tilde{L}(\vec{A}_1,\vec{A}_2,\vec{d}_\epsilon) = \sum_i\sum_s\sum_r \hat{y}_i^{rs} - y_i^{rs}\ln\left(\hat{y}_i^{rs}\right). \tag{7.13}$$

In a real experiment, all setup-dependent quantities $(S(E), R^s(E), V(E), \phi^r(E))$ can depend on the spatial position and therefore on the detector pixel index i. For simplicity and clarity, we have suppressed this possible dependency in this section. Figure 7.1 shows an overview of the workflow for both SDPC imaging and conventional spectral imaging. While SDPC material decomposition takes energy-resolved stepping curve measurements as an input, standard material decomposition takes energy-resolved intensity images as an input. In a setup with a grating interferometer, the intensity images are generated by summing up all stepping curve measurements. In a non-interferometric setup, the output images of a photon counting detector can be used directly. Both imaging methods calculate bone and soft tissue basis material images. SDPC material decomposition generates an additional dark-field image as an output.

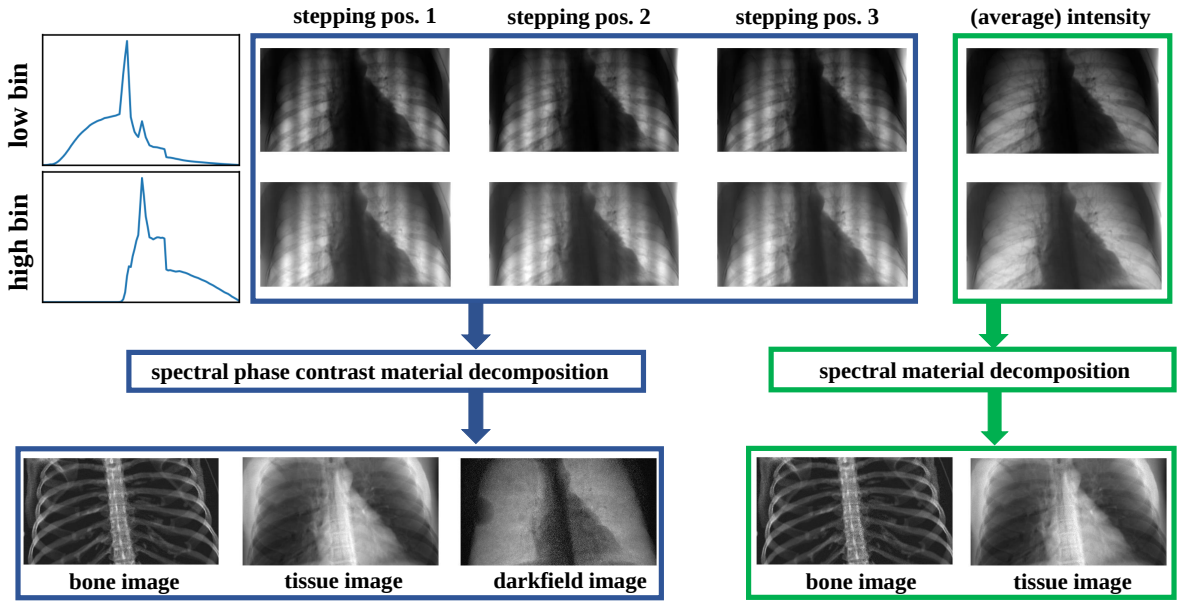

Figure 7.1: Overview of the workflow for spectral phase-contrast material decomposition and conventional spectral material decomposition. While spectral phase-contrast material decomposition takes energy-resolved stepping curve measurements as an input, standard material decomposition uses energy-resolved intensity images as an input. Both methods calculate bone and soft tissue basis material images. Spectral phase-contrast material decomposition generates an additional dark-field image as an output.

In contrast to standard spectral imaging, the SDPC optimization problem is not separable with respect to the detector pixels because the gradient of the electron density couples the pixel values in the direction perpendicular to the grating bars. The optimization can therefore be performed separately for each detector row, but not for each detector pixel. Moreover, the cosine term in the forward model (see eq. 7.5) introduces local optima into the log-likelihood function of eq. 7.13. In the results section, we analyze the occurrence of these local optima in more detail. The optimization result therefore depends on the initial guess, unless a global optimization strategy is used. In our implementation, we used a nonlinear conjugate gradient algorithm (NLCG) to minimize the log-likelihood for SDPC imaging. Like most gradient descent algorithms, the NLCG algorithms converges to a local optimum. We did not use a global optimization strategy due to the drastically increased computational cost. We obtained an initial guess for $\vec{A}_0$ and $\vec{A}_1$ by smoothing the result of the conventional spectral material decomposition with a Gaussian filter ($\sigma = 3$ pixels). The visibility extinction coefficients ϵ_i can be determined for each detector pixel using the standard signal extraction methods for conventional phase-contrast imaging. With the help of the approximation:

$$\epsilon_i \approx d_\epsilon^i f_\epsilon(E_{\text{eff}}), \tag{7.14}$$

where E_{eff} is an effective energy as defined in [23], an initial guess for $\vec{d}_\epsilon$ can be obtained. The initial guess for $\vec{d}_\epsilon$ is also smoothed with a Gaussian filter. Using this initialization strategy, the NLCG algorithm converged to the global optimum in our simulation study. To ensure that $\hat{y}_i^{rs}(\vec{A}_1, \vec{A}_2, \vec{d}_\epsilon) > 0$ (otherwise numerical problems can arise during the optimization because of the logarithm in the log-likelihood function, see eq. 7.13), we constrain the dark-field line integrals $\vec{d}_\epsilon$ to a predefined lower bound (in this case -1 cm). This prevents the term $V(E)e^{-d_\epsilon^i f_\epsilon(E)}$ in the SDPC forward model (eq. 7.5) from becoming larger than one (averaged over the spectrum).

There are different possibilities to calculate the gradient of the projected electron density (compare eq. 7.9) numerically. Using the forward or backward difference directly couples neighboring

pixels, but it has the disadvantage of introducing a slight systematic inconsistency between the forward model and the measured phase shift. The measured phase shift in a particular pixel is determined by the average of the projected electron density gradient in this pixel, but by using the forward or backward difference, the gradient of the forward model is shifted by $\pm\frac{1}{2}$ pixels. Using a central difference approximation eliminates this inconsistency but it has the disadvantage of generating a "checkerboard-like" noise pattern because the electron density gradient only couples optimization variables that are two pixels apart (e.g., using basis material M_1 as an example, the gradient for pixel i couples $A_1^{(i+1)}$ and $A_1^{(i-1)}$. To circumvent these problems, we minimize $-\tilde{L}(\vec{A}_1, \vec{A}_2, \vec{d}_\epsilon)$ twice, using both the forward and backward difference approximation and then average the resulting basis material images.

7.2.2 Numerical Experiment

Spectral phase-contrast projection measurements of a human thorax were simulated with the wave-optical simulation package that was presented in chapter 3. We assumed a tungsten X-ray source operated at an acceleration voltage of 120 kVp and a photon-counting detector with 830×2000 pixels (pixel size 300 μm), a 2 mm thick cadmium telluride (CdTe) sensor and two thresholds per pixel. The threshold positions (15 keV and 63 keV) were optimized for spectral X-ray imaging. The simulated setup features a symmetric grating interferometer with a design energy of 50 keV and a total length of 2.4 m. Figure 7.2 (a) shows an overview of the setup geometry and the grating specifications for the simulated thorax radiography. The G0 and G2 grating are made of gold with a grating bar height of 200 μm. The G1 grating is made of nickel and induces a phase shift of $\pi/2$ for the design energy of 50 keV. The patient is assumed to stand 40 cm away from G1, between G1 and G2. Figure 7.2 (b) shows a plot of the energy-dependent visibility of the interferometer together with the effective spectra for the low and high energy bin of the photon counting detector. As pointed out by Kottler et. al [401], a three-grating interferometer with a $\pi/2$ phase shifting grating exhibits a second visibility peak at twice the design energy (here 100 keV). The second important factor determining the visibility is the attenuation of the gold gratings. In the range from $60-80$ keV, the gold gratings become increasingly transparent which explains the decreasing visibility in this energy range. For energies above the K-edge of gold (80.7 keV), the attenuation of the gratings, and thus also the visibility, rises sharply. The reference visibilities in the low and high energy bin were 27.5% and 14.7%, respectively. The thorax phantom is based on a publicly available CT scan [402] that was segmented into bone, soft tissue and lung tissue. Two artificial lung nodules (one of them calcified) were placed in the lung. They were simulated by replacing a spherical region of the lung tissue (diameter of 8 mm) by tissue and calcified tissue (with a calcium volume fraction of 3.5 %), respectively. Furthermore, a pneumothorax was simulated by replacing a region of the lung with air. Due to the large size of the thorax phantom, the projection approximation was used in our simulation. To analyze the bias and the noise properties of the proposed material decomposition approach, we conducted an additional simulation of the central detector line. In this simulation, the thorax phantom was divided into 64 different planes (with surface normal vectors oriented parallel to the beam direction) and we alternatingly applied the projection approximation before propagating the wavefront to the location of the next plane. This simulation strategy accounts for the position-dependent phase shift that is caused by a specific electron density gradient (compare eq. 7.9 and eq. 7.10).

In theory, it is possible to simulate the visibility reduction by including the microstructure of the lung in the wave-optical simulation [322]. However, in our case this approach is infeasible because of the comparatively large size of the thorax phantom. Moreover, 3D images of the lung with sub-micrometer resolution would be required for a realistic simulation. Therefore, we simulated the dark-field signal by scaling the amplitudes of the simulated stepping curves by the factor $e^{-d_\epsilon f_\epsilon(E)}$, where d_ϵ is the pathlength through the lung tissue. At the design energy of 50 keV, the simulated interferometer has a correlation length of $\xi = 1.8$ μm, which is much

smaller than the typical structure size of the lung (the alveoli have a diameter of ≈ 200 μm). We therefore approximated $G(\xi)$ by a first order Taylor expansion to obtain an expression for $f_\epsilon(E)$ that does not rely on detailed knowledge about the microstructure of the lung:

$$G(\xi) \approx G(0) + G'(0)\xi = 1 + G'(0)\xi. \tag{7.15}$$

Combining this approximation with eq. 7.6 and using $\sigma(E) \propto \frac{1}{E^2}$, the visibility reduction can be written as:

$$\frac{V_s}{V} = e^{-d_\epsilon \frac{C}{E^2}\xi}, \tag{7.16}$$

where C is a proportionality constant. Assuming a fixed distance between the object and G2 ($d_{S,G2}$), eq. 7.16, is simplified to:

$$\frac{V_s}{V} = e^{-d_\epsilon \frac{C}{E^3}}. \tag{7.17}$$

Neglecting higher order terms of the Taylor expansion in eq. 7.15 might not always be justified. The most notable counterexample is a (diluted) dispersion of microspheres, for which $G(\xi)$ can be approximated by [295]:

$$G(\xi) = G\left(\xi / r_\mathrm{s}\right) \approx e^{-\frac{9}{8}\left(\frac{\xi}{r_\mathrm{s}}\right)^2}, \tag{7.18}$$

where r_s is the radius of the spheres. In this case, $G'(0) = 0$ and a second-order Taylor expansion leads to an E^{-4} dependency of the dark-field signal for $\xi \ll r_s$. Nevertheless, the E^{-3} dependency of the dark-field signal that was obtained in eq. 7.17 coincides with preliminary dark-field measurements of a porcine lung that were conducted at an experimental setup at our Institute. For more complicated and less symmetric structures than the microspheres, it can be expected that $G'(0)$ does not vanish.

Equation 7.16 considers the fact that the dark-field signal depends on the position along the optical axis and was therefore used for the additional simulation of the central detector line, whereas eq. 7.17 was used for the simulation of the whole image.

For the spectral phase-contrast simulations, a stepping curve with 5 steps homogeneously distributed between 0 and 2π was measured for each of the two energy bins of the photon counting detector. We conducted an additional simulation of a conventional spectral measurement in order to compare the performance of SDPC decomposition to conventional spectral material decomposition. For this simulation, the three gratings were removed, but otherwise the same imaging parameters were used. The dose for all simulated measurement was 0.3 mGy (estimated with the method presented in [363]).

7.3 Results

In order to look at the origin of local optima more closely, we investigate a simplified version of $-\tilde{L}(\vec{A}_1, \vec{A}_2, \vec{d}_\epsilon)$: $\vec{d}_\epsilon$ is eliminated by assuming that there are no dark-field generating materials in the object. Moreover, only the basis material line integrals for the central pixel of a detector row (A_1^c, A_2^c, c being the index of the central pixel) are considered as optimization variables. The basis material line integrals for all other pixels are kept fixed at their ground truth values (5 cm of tissue and 0.75 cm of bone). Figures 7.3 (a) and (b) show contour plots of the log-likelihood function for SDPC and conventional material decomposition, respectively (using bone and soft tissue as basis materials). To better separate the effects of phase shift and attenuation, we consider two artificial basis materials M_1, M_2 with the following properties:

$$\begin{aligned} \rho_e(M_1) &= 0\,\mathrm{cm}^{-3}, & f_1(E) &\propto \frac{1}{E^3}, \\ \rho_e(M_2) &= 10^{23}\,\mathrm{cm}^{-3}, & f_2(E) &\propto \rho_e(M_2) f_\mathrm{kn}(E), \end{aligned} \tag{7.19}$$

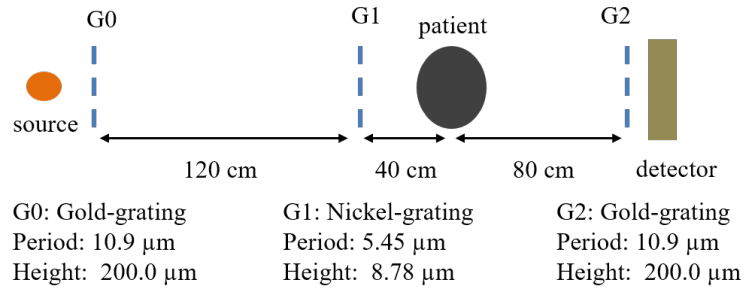

(a)

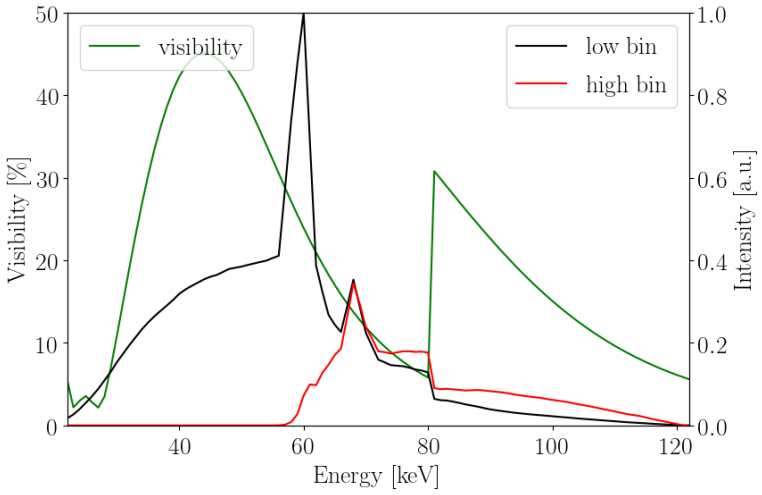

(b)

Figure 7.2: (a): Overview of the setup geometry and the grating specifications for the numerical experiment. (b): Energy-dependent visibility of the interferometer together with the effective spectra for the two bins of the simulated photon-counting detector.

where $f_{\mathrm{kn}}(E)$ is the Klein-Nishina function [2]. As far as the energy-dependent attenuation is concerned, M_1 can be interpreted as photoeffect while M_2 represents Compton-scattering. However, due to the definition of the electron densities for M_1 and M_2, this comparison cannot be taken literally. Figures 7.3 (c) and (d) show contour plots of the log-likelihood function in the photo-Compton basis for SDPC and conventional material decomposition, respectively. For both the bone-tissue as well as the photo-Compton basis, the log-likelihood function for conventional material decomposition has one global optimum and the contour lines form ellipses with an inclination of approximately 45° with respect to the coordinate axes. The shape of the log-likelihood surface therefore reflects the typical anti-correlated noise of the decomposed basis material images. The log-likelihood surfaces for SDPC material decomposition have the same overall shape. However, close to the minimum, the surface shows deviations compared to conventional material decomposition. In the following, we focus on the photo-Compton basis because the deviations to conventional decomposition are easier to explain in this basis. The phase shift term in the forward model for SDPC imaging, which in turn depends on the gradient of the projected electron density, modifies the log-likelihood functions predominately in the direction of basis material M_2 (i.e., the Compton effect). If the electron density deviates from its ground truth value, the phase shift predicted by the forward model no longer matches the physical phase shift. Consequently, the predicted stepping curve $\hat{y}_c^{rs}$ deviates from the measured values y_c^{rs} and the value of the log-likelihood function increases. For conventional phase-contrast imaging, the value of the log-likelihood function decreases again once one of the following two

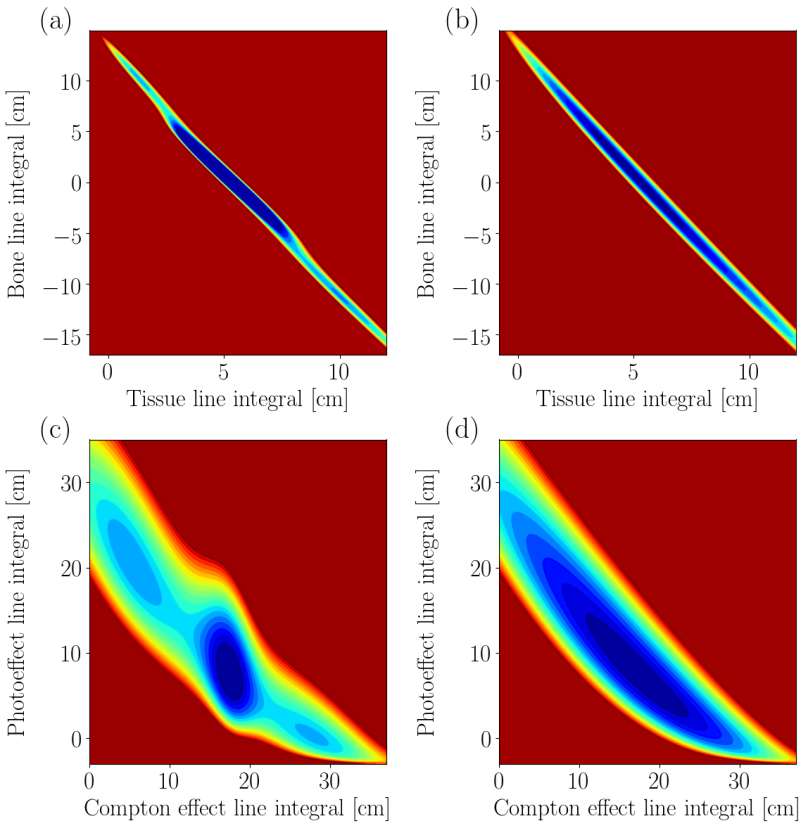

Figure 7.3: Contour plots of the simplified log-likelihood functions for spectral phase-contrast material decomposition (a,c) and conventional spectral material decomposition (b,d). The top row shows the log-likelihood functions in the bone-tissue basis whereas the bottom row shows the photo-Compton basis.

conditions is met:

$$\begin{aligned} \Delta\phi_c &> \Delta\phi_c^{\mathrm{gt}} + \pi \\ \Delta\phi_c &< \Delta\phi_c^{\mathrm{gt}} - \pi, \end{aligned} \tag{7.20}$$

where $\Delta\phi_c$ and $\Delta\phi_c^{\mathrm{gt}}$ are the phase shift predicted by the forward model and the ground truth phase shift, respectively. Furthermore, if $\Delta\phi_c = \Delta\phi_c^{\mathrm{gt}} + m \cdot 2\pi,\ m \in \mathbb{Z}$ the log-likelihood function has the same value again. This corresponds to the phase-wrapping problem for conventional phase-contrast imaging. For SDPC imaging, the analysis is complicated by the energy-dependency of the phase shift $\Delta\phi(E)$, however the aforementioned arguments qualitatively explain the occurrence of the two "barriers" close to the global minimum of the log-likelihood function and their orientation perpendicular to the Compton effect axis. Moreover, the appearance of the two local minima can be explained by the analogy to the phase-wrapping problem. In contrast to conventional phase-contrast imaging, the correct phase shift can be unambiguously determined by comparing the value of the log-likelihood function at the minima. At the position of the two local minima, the forward model does not correctly predict the average intensities of the spectral stepping curve measurements, and thus the value of the log-likelihood function is higher compared to the global minimum. Consequently, SDPC imaging eliminates phase-wrapping artifacts, provided one can find the global optimum of the log-likelihood function. However, reliably finding the global optimum of the (non-simplified) log-likelihood function (compare eq. 7.13) in acceptable computational time remains an open question. For the simulated setup, the local optima are well separated from the global optimum (by ±5 cm of tissue thickness and ± 10 cm of bone thickness). With the simulated dose level of

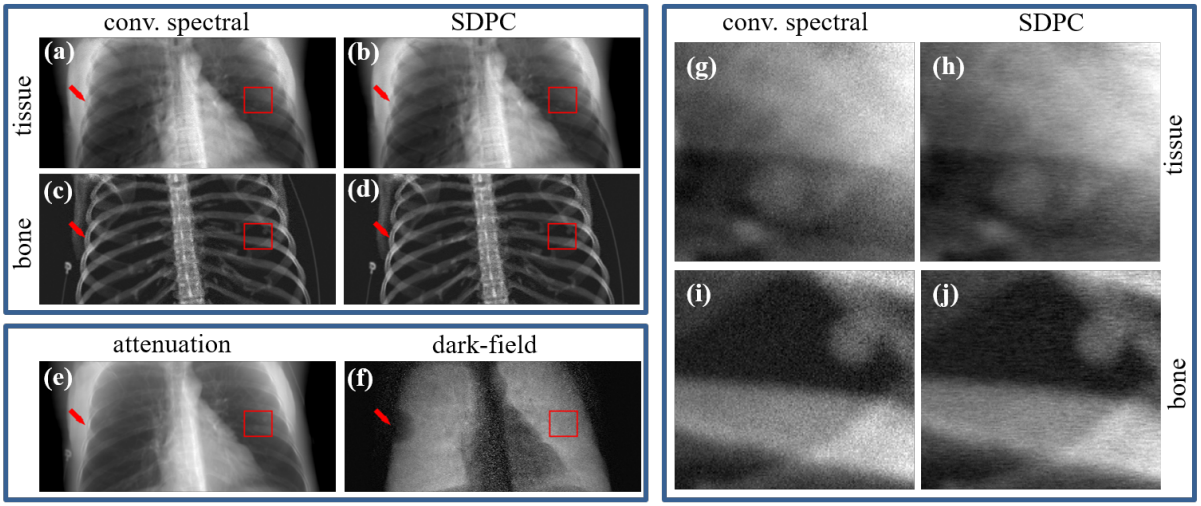

Figure 7.4: Numerical simulation of a human thorax radiography. The left part shows a comparison between conventional spectral material decomposition (a),(c) and spectral differential phase-contrast (SDPC) material decomposition ((b),(d)) for both the tissue and bone images. The red rectangle in the images highlights the area where the artificial lung nodules are placed. Moreover, the conventional attenuation image (e) as well as the dark-field image (f) that was obtained by the proposed SDPC decomposition algorithm are displayed. The red arrows highlight the location of a simulated pneumothorax. The right part of the figure shows a zoom of the area where the artificial lung nodules are placed. The zoomed images facilitate a visual comparison of conventional spectral decomposition ((g),(i)) and SDPC decomposition ((h),(j)) regarding image quality and noise levels.

0.3 mGy, it is thus unlikely that the initial guess obtained from a conventional spectral decomposition is close to one of the local optima. By smoothing the initial guess (compare Section 7.2.1), we further reduce the probability of convergence to a local optimum. Phase unwrapping would be more difficult for a more sensitive grating interferometer (because the local optima are spaced more closely) or reduced photon statistics (because the quality of the initial guess decreases). In these cases, regularized material decomposition techniques [235, 236] could possibly improve the stability of the phase-wrapping correction.

Figure 7.4 shows a comparison of tissue and bone basis material images obtained by conventional material decomposition ((a),(c)) and SDPC decomposition ((b),(d)). The red arrows highlight the location of the simulated pneumothorax and the red rectangles highlight the area where the two artificial lung nodules are located. Furthermore, the conventional attenuation image (e) (calculated by summing up the two spectral measurements) and the dark-field image calculated by SDPC decomposition (f) are displayed. Compared to the attenuation image, the soft tissue images facilitate the detection of the lung nodule that is overlayed by a rib. Furthermore, the calcification of the upper nodule can be detected in the bone images. In this simulation study, the artificial pneumothorax is easily detected in the dark-field image because the missing lung tissue causes a reduction of the dark-field signal in the corresponding area. In the conventional attenuation image, the pneumothorax is much harder to spot because of the low attenuation of the lung and the superimposed soft tissue and bone structures.
The right part of figure 7.4 shows soft tissue ((g),(h)) and bone ((i),(j)) images for both decomposition methods zoomed into the area where the lung nodules are located (compare red rectangles in the left part of the figure). The basis material images generated with the proposed SDPC decomposition algorithm ((h) and (j)) show a reduced noise level as well as a different noise texture compared to conventional material decomposition ((g) and (i)). The different noise texture is caused by the phase shift term which couples neighboring pixels in direction perpen-

dicular to the grating bars. A quantitative noise analysis (compare Table 7.1) in the zoomed region shows that SDPC imaging reduces the variance by a factor of 2.5 and 1.9 for the tissue and bone image, respectively. Consequently, the dose delivered to the patient can be reduced by a factor of 2.5 and 1.9 while maintaining the same noise level for the tissue and bone image, respectively. The additional phase shift information for SDPC decomposition overcompensates the reduced photon statistics (because of the G2 attenuation) in comparison to conventional spectral decomposition. The G0 and G1-grating only influence the flux and not the dose, because they are placed between the source and the patient, but they cause beam hardening which decreases the spectral separation between the low and high energy bin.
In order to assess the stability of the SDPC decomposition algorithm, the simulation was repeated for 30 different dose levels ranging from 0.003 mGy to 3 mGy. Figure 7.5 exemplarily shows the variance of the tissue image in the same zoomed region as a function of the dose level for both decomposition methods. The variance of the bone image follows a similar pattern. For both decomposition methods, the experimentally obtained variances (dots) are inversely proportional to the photon statistics (indicated by solid lines) for a large range of dose levels. The SDPC decomposition only exceeds the predicted variances for ultra low dose levels (< 0.01 mGy). In conclusion, the variance reduction reported in Table 7.1 remains valid over a large range of dose levels.
We conducted an additional analysis of the variance and bias for conventional and SDPC material decomposition by looking at the decomposition results for the central detector line with 30,000 different noise realizations. As mentioned in Section 7.2.2, the position dependency of the phase shift as well as the visibility reduction was considered in this simulation. Figure 7.6 (a) shows the conventional attenuation image of the simulated thorax radiography. The red line in figure 7.6 (a) indicates the central detector line for which the analysis was conducted. Moreover, markers for three different areas ((1), (2) and (3)) are displayed. Figure 7.6 (c) displays the variance reduction of SDPC imaging in comparison to conventional spectral imaging for the soft tissue and bone image in dependence of the pixel index. In area (1), where the open beam hits the detector, SDPC decomposition shows the largest variance reduction (by a factor of 5.5 for the tissue image and a factor of 3 for the bone image). In area (2), the beam is strongly attenuated as the projected bone and soft tissue thicknesses are comparatively large. Beam hardening leads to an increased average energy of the spectrum and thus the phase shifts become smaller ($\Delta\phi(E) \propto \frac{1}{E^2}$). Furthermore, the effective visibility of the stepping curves is reduced because the gold gratings become increasingly transparent at higher energies and because the effective energy of the hardened spectrum is shifted away from the design energy of the interferometer (compare figure 7.2 (b)). These effects explain the smaller advantage of SDPC imaging compared to conventional spectral imaging in area (2). In the area of the lungs (3), the attenuation is comparatively weak, but the dark-field signal of the lungs strongly reduces the effective visibility of the stepping curve. The noise level for SDPC decomposition increases with decreasing visibility, similar to conventional phase-contrast imaging, where the image noise is inversely proportional to the effective visibility squared [279]. This explains why the smallest variance reduction factors occur in the area of the lungs.
Figure 7.6 (b) exemplarily shows the decomposition bias of the soft tissue image, whereas figure 7.6 (d) shows the decomposition bias for the dark-field image. The decomposition bias for the bone image is anticorrelated to the soft tissue image and smaller in magnitude. The total bias b_{total} can be viewed as the sum of three bias contributions:

$$b_{\text{total}} = b_{\text{model}} + b_{\text{proj}} + b_{\text{stat}}. \tag{7.21}$$

The first term b_{model} is the inherent bias of the SDPC model that remains even when an infinitely thin object (along the projection direction) is measured. It is determined by performing an SDPC decomposition from noise-free data simulated with the projection approximation. The most important cause of this bias is the fact that the material decomposition is performed on a

Table 7.1: Quantitative noise analysis for spectral differential phase-contrast (SDPC) and conventional spectral imaging (conv. spectral) in the region of the thorax phantom where the artificial lung nodules are located (compare right part of figure 7.4). Additionally, we evaluate how much the dose for SDPC imaging can be reduced compared to conventional spectral imaging while achieving the same noise levels.

	Variance [mm^2]		Dose reduction
	conv. spectral	SPDC	factor
Tissue image	3.5	1.4	250 %
Bone image	0.40	0.21	190 %

much coarser grid than the simulation. This gives rise to (nonlinear) partial volume effects (for example, a changing electron density gradient within one pixel cannot be modeled). However, the magnitude of these effects is comparatively small (maximum bias < 200 μm).
The second term b_{proj} describes the bias of implicitly using the projection approximation in the SDPC forward model. Since the position along the optical axis cannot be recovered from a radiography measurement, the position dependency of the phase shift and visibility reduction cannot be taken into account. This contribution to the bias could in principle be eliminated in a possible future extension of the SDPC algorithm to statistical iterative CT reconstruction. It is determined by calculating the bias of an SDPC decomposition simulated with the multi-step propagation approach and subtracting b_{model}. For both the soft tissue and dark-field image, b_{proj} is approximately an order of magnitude larger than b_{model}. As can be seen from eq. 7.10 and eq. 7.7, the relative magnitude of b_{proj} for a small volume element is proportional to the ratio of the distance to the center of the object (in projection direction) to the distance between the gratings G1 and G2. For the numerical simulation, the maximum ratio is approximately 10 %. Nevertheless, the maximum thickness bias of ≈ 1 mm (relative error < 1 %) is much smaller than the aforementioned geometrical considerations suggest. Since b_{proj} depends linearly on the distance to the center of the object, a large part of the bias contributions of volume elements upstream and downstream of the center of the object cancel out. In case of the soft tissue image, b_{proj} is correlated to large electron density gradients (e.g., ribs) that occur outside of the center of the thorax phantom (with respect to the optical axis). For the dark-field image, b_{proj} is caused by an asymmetric distribution of lung tissue with respect to the center of the thorax phantom.
The term b_{stat} describes a noise-induced contribution to the total bias. This effect has already been investigated for conventional phase-contrast imaging[403] and spectral imaging [353]. The statistical bias was determined by subtracting b_{model} and b_{proj} from the total bias that was obtained by analyzing the 30,000 different noise realizations. The strong attenuation in area (2) leads to decreased photons statistics and consequently to an increase of the statistical bias for both the soft-tissue and the dark-field image. At a dose level of 0.3 mGy the noise-induced bias is small (< 200 μm), except for three large spikes in the dark-field image. The two spikes at the edges of the phantom can be correlated with large electron density gradients. As the visibility in this area is strongly reduced (via phase shift dispersion [311]), the dark-field noise and thus the statistical bias increases. In the area of the spine, the strong attenuation increases the dark-field noise level far enough that the constraint of $d_\epsilon > -1$ cm becomes active, which explains the large bias spike in the center of the phantom.

7.4 Conclusions

We have developed a novel model that combines energy-resolved and phase-contrast X-ray imaging. The connection between the two modalities is made via the projected electron density, which can be determined with both methods. In our experiments, the combination of these two imaging techniques shows promising results for medical imaging applications. The simulation of a

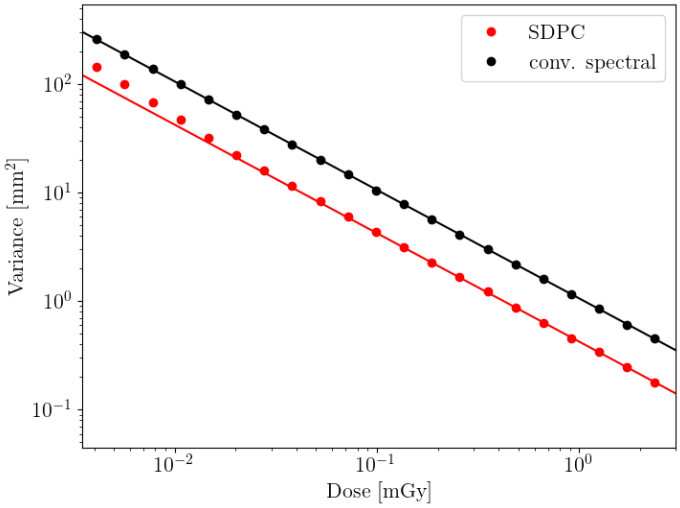

Figure 7.5: Variance of the tissue image for SDPC and spectral decomposition in the zoomed region of figure 7.4 as a function of the dose level. The experimentally obtained variances are represented by dots and the solid lines indicate the expected inverse proportionality to the dose level.

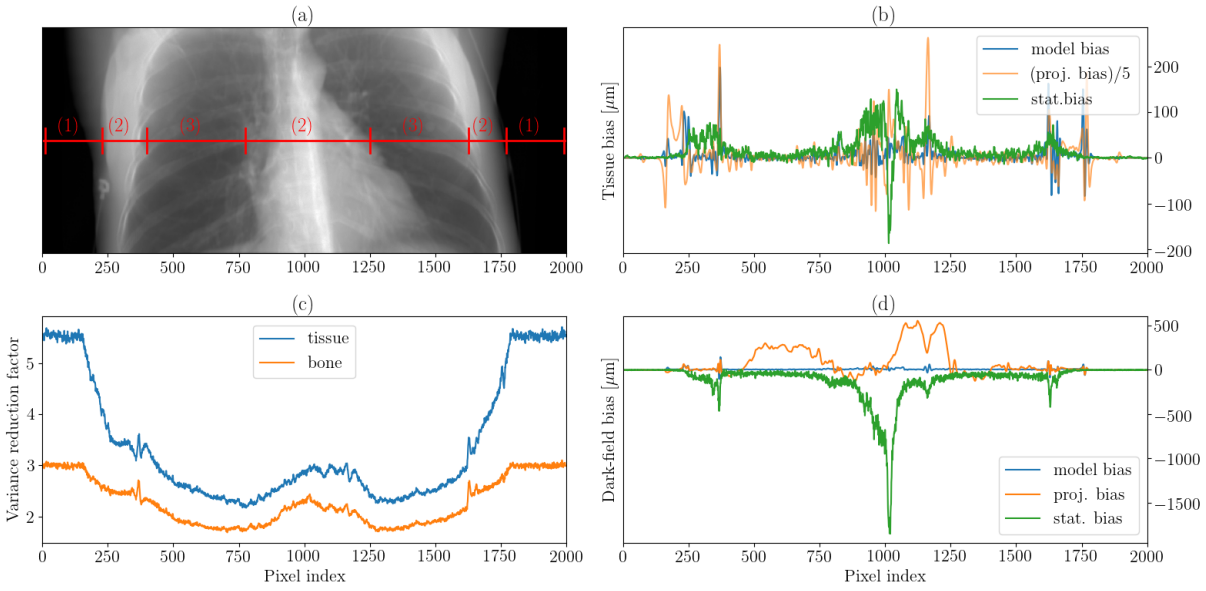

Figure 7.6: Analysis of variance and bias for SDPC and conventional spectral material decomposition. (a): Conventional attenuation image with the central detector line marked in red. (b): Lineplots of the decomposition bias for the tissue image. The decomposition bias is composed of three terms: The inherent bias of the model (model bias), the noise-induced bias (stat. bias) and a bias term that is caused by the position dependency of the dark-field signal and the phase shift (proj. bias) The last term was divided by 5 to fit on the same scale. (c): Variance reduction factor of SDPC decomposition compared to conventional spectral decomposition as a function of the detector pixel index. (d): Dark-field decomposition bias for SDPC imaging as a function of the detector pixel index.

human thorax radiography demonstrates that quantitatively correct basis material images with highly reduced noise levels compared to conventional spectral imaging can be obtained under clinically realistic conditions (e.g., object size, spatial resolution, dose). Compared to spectral imaging, the additional information on the phase shift determines the projected electron density more precisely. This additional constraint leads to a strong noise reduction for the basis material line integrals. Within this proof-of-concept study, we cannot definitely answer the question if this noise reduction directly relates to improved clinical diagnostics (e.g., for the detection of lung nodules). An important point to consider in this context is that (unlike spectral imaging) SDPC imaging leads to correlated noise in the direction perpendicular to the grating bars. The simulations give an upper bound for the noise advantage of SDPC imaging in a real experiment because various unfavourable experimental effects (e.g., imperfect gratings) were not considered. However, the large dose reduction factors suggest that reduced noise levels could also be achieved in real experiments. In relation to spectral imaging, the optimization of the log-likelihood function is complicated by the occurrence of local optima. In our numerical experiments, the convergence to local optima could be avoided by a suitable initial guess. Since this cannot be guaranteed for arbitrary objects or setup parameters, choosing the right optimization algorithm for SDPC imaging remains an open question. Another important advantage of the proposed spectral phase-contrast model is the additional information provided by the dark-field image. This contrast mechanism gives information about the object's microstructure that is inaccessible to conventional spectral imaging. As the dark-field signal decreases when the microstructure of the lung is destroyed, dark-field X-ray imaging is a promising candidate for the early detection of lung diseases [404, 306, 405].

8 A theoretical framework for comparing noise characteristics of spectral, differential phase-contrast and spectral differential phase-contrast X-ray imaging

In the last chapter, we have investigated the combination of grating-based DPC radiography and spectral radiography [406] and have developed a basis material decomposition algorithm that uses the spectral and the phase contrast information simultaneously. Our numerical experiments have demonstrated that SDPC radiography yields quantitatively correct basis material images with strongly reduced noise levels compared to conventional spectral imaging. Analogous to DPC imaging, we have observed correlated noise patterns for SDPC imaging. Since the phase shift of the stepping curve is proportional to the gradient of the projected electron density, the phase shift information couples neighboring pixels in the direction perpendicular to the grating bars. In the case of DPC imaging, the integration step that is necessary to calculate electron density images causes long-range noise correlations and excessive low frequency noise. Based on these noise characteristics, Raupach and Flohr have argued that DPC imaging will produce inferior image quality compared with standard attenuation-based imaging for clinical CT applications [391]. However, it could be expected that SDPC imaging produces less low frequency noise, since it can be viewed as a mixture of spectral and DPC imaging. In the case of spectral imaging, there are no noise correlations between neighboring pixels, but the overall noise level is typically relatively high due to noise amplification during material decomposition.
Taking the characteristics of all three aforementioned imaging methods into account, the question of the optimal method for a particular imaging task arises. In this chapter, we investigate one important aspect of this question by conducting an in-depth noise analysis of all three imaging methods. Towards this end, we develop a noise analysis framework based on the Cramér Rao lower bound (CRLB) [96]. This framework allows the determination of the optimum imaging parameters for all three methods and the prediction of noise correlations and noise power spectra. Since all three aforementioned imaging methods can determine the projected electron density (PED), we focus on PED images for our comparison. Although only projection imaging is considered in this work, a generalization to 3D computed tomography is possible. Using an imaging task that can be viewed as a simplified version of human X-ray radiography, we demonstrate that the combination of spectral and phase-contrast imaging has the potential to generate PED images with strongly reduced noise levels compared with the individual methods. Our theoretical analysis suggests that SDPC imaging outperforms spectral and DPC imaging for a large range of clinically relevant pixel sizes. Finally, we discuss the noise characteristics of SDPC imaging in the context of the comparison between attenuation-based imaging and DPC imaging by Raupach and Flohr [391].
The text and the figures in this chapter are taken from the following publication of the author:

- Mechlem, Korbinian, et al. "A theoretical framework for comparing noise characteristics of spectral, differential phase-contrast and spectral differential phase-contrast X-ray imaging." Physics in Medicine & Biology (2020).

Compared to the original publication, the text has been slightly modified to achieve a better integration into this doctoral thesis.

8.1 Methods

Before explaining the derivation of the noise analysis framework based on the CRLB, we give a brief overview of the physical models (forward models) of the measurement acquisition processes for all three imaging methods under consideration (spectral, SDPC and DPC imaging). Furthermore, we formulate the signal extraction process as a maximum-likelihood (ML) estimation problem for all three imaging methods. ML estimation has already been used successfully for projection-based material decomposition in spectral imaging [316, 249]. In this case, the ML estimator has many desirable properties such as the capability to handle overdetermined systems (i.e. having more energy bins than basis materials). Moreover, it is unbiased and efficient (i.e. it achieves the minimum variance for an unbiased estimator) in the limit of low noise levels. We also successfully applied an ML-based decomposition algorithm to SDPC imaging [406]. In the case of DPC imaging, Fourier processing [22] is the most commonly used method of signal extraction due to the comparatively low computational complexity. Using error propagation, the noise properties of DPC imaging in combination with Fourier processing have already been investigated for CT reconstruction [18] and in the projection domain [407, 408, 409]. However, it has been demonstrated that a Fourier-based estimator for DPC imaging is not efficient [410]. For this reason, and to make the signal extraction for all three methods more comparable, we use an ML estimator for DPC imaging. This estimator is similar to weighted least-squares signal extraction methods for DPC imaging [411, 412].
In the following, we assume that a photon-counting detector (PCD) is used to acquire energy-resolved measurements for all three imaging methods. The number of registered photon counts is therefore modeled by a Poisson distribution. Moreover, we assume that there is no correlation between the photon counts registered in different detector pixels or energy bins. These assumptions are valid for an ideal PCD, but current real PCDs exhibit various undesired sensor effects (e.g. pulse pile-up, charge sharing) that could cause a violation of these assumptions [413, 91]. In the limit of negligible electronic noise and detector blur, the noise analysis presented in this work could be extended to other spectral imaging methods, such as dual source CT or kVp-switching. As will be explained later, the differential phase shift for SDPC and DPC imaging couples pixels in the direction perpendicular to the orientation of the grating bars. However, the pixels remain uncoupled in the direction parallel to the grating bars. For simplicity, we will therefore consider a one-dimensional PCD in the following because each detector row can be treated separately for all three imaging methods.

8.1.1 Forward models and signal extraction

Basis material decomposition for spectral imaging relies on the assumption that the energy-dependent attenuation of any material can be modeled by a linear combination of a few basis materials. Neglecting materials with K-edges in the relevant energy range for medical imaging ($\approx 20 - 140$ keV), only two basis materials are needed. This is a consequence of the fact that there are only three interaction mechanisms (photoelectric effect, Compton-scattering and Rayleigh-scattering) in this energy range. Furthermore, the contribution of Rayleigh-scattering to the total attenuation cross-section is typically small compared with the other two interaction mechanisms. For simplicity, we will focus on two basis materials in the following. For spectral imaging, the expected number of photon counts $\hat{y}_i^s$ registered in energy bin s and detector pixel i is thus modeled by [316]:

$$\hat{y}_i^s = \int_0^\infty t(\mathcal{E})\mathcal{R}^s(\mathcal{E})e^{-A_1^i f_1(\mathcal{E}) - A_2^i f_2(\mathcal{E})}d\mathcal{E}, \tag{8.1}$$

where $t(\mathcal{E})$ is the source spectrum and $\mathcal{R}^s(\mathcal{E})$ describes the detector response, i.e. the probability that a photon with energy $\mathcal{E}$ is detected by energy bin s of the PCD. The functions $f_1(\mathcal{E})$ and $f_2(\mathcal{E})$ represent the energy-dependent attenuation of the two basis materials and the

corresponding basis material line integrals (for detector pixel i) are denoted by A_1^i and A_2^i. Assuming uncorrelated Poisson statistics, the negative log-likelihood function for spectral imaging is given by:

$$-L(\vec{A}_1, \vec{A}_2) = \sum_{i=1}^{N} \sum_{s=1}^{S} \hat{y}_i^s - y_i^s \ln\left(\hat{y}_i^s\right), \tag{8.2}$$

where N and S represent the number of detector pixels and energy bins of the PCD, respectively. The quantity y_i^s denotes the measured number of photon counts. ML decomposition is performed by finding the basis material line integrals $\vec{A}_1 = \left(A_1^1, ..., A_1^N\right)$, $\vec{A}_2 = \left(A_2^1, ..., A_2^N\right)$ that minimize $-L(\vec{A}_1, \vec{A}_2)$. This optimization problem can be solved separately for each detector pixel. The PED (ρ_e) can be calculated by a linear combination of the basis material line integrals:

$$\rho_\mathrm{e}^i = A_1^i \rho_\mathrm{e}^\mathrm{V}(\mathrm{M}_1) + A_2^i \rho_\mathrm{e}^\mathrm{V}(\mathrm{M}_2), \tag{8.3}$$

where $\rho_\mathrm{e}^\mathrm{V}(\mathrm{M}_1)$ and $\rho_\mathrm{e}^\mathrm{V}(\mathrm{M}_2)$ are the volume electron densities of the two basis materials. The idea of approximating the electron density of any material by a linear combination of the two basis materials is conceptually very similar to the standard dual energy assumption of modeling the energy-dependent attenuation with two basis materials. Consequently, equation 8.3 is also only valid for the dual energy parameter range (low-Z materials, energy range $\approx 20 - 140$ keV). For DPC imaging, the three contrast modalities (attenuation, differential phase shift and dark-field) are extracted from stepping curve measurements that are generated by shifting one of the gratings [13, 22]. In contrast to spectral and SDPC imaging, no energy-resolved measurements are acquired. We therefore assume a PCD with just one threshold for DPC imaging. The stepping curve is typically modeled by a cosine (or sine) function [22] and beam hardening effects caused by the polychromatic spectrum are neglected. The expected intensity $\hat{y}_i^r$ for stepping position r and detector pixel i can thus be written as: [278]:

$$\hat{y}_i^r = b e^{-\mu_i} \left(1 + V e^{-\epsilon_i} \cos\left(\phi_r + \Delta\phi_i\right)\right), \tag{8.4}$$

where ϕ_r and V are the reference phase (for step r) and the reference visibility of the stepping curve, respectively. The quantities μ_i, $\Delta\phi_i$ and ϵ_i describe the attenuation of the object, the phase shift of the stepping curve and the visibility reduction (dark-field signal), respectively. The reference intensity measured with an open beam is given by the parameter b. The standard DPC stepping curve model (equation 8.4) implicitly represents the polychromatic spectrum by an effective X-ray energy. This approximation is reasonably accurate for weakly attenuating samples that only slightly distort the incident X-ray spectrum. Since enhancing soft tissue contrast is one of the main application cases for DPC imaging, this assumption is often justified. The attenuation, differential phase and dark-field images ($\vec{\mu}, \Delta\vec{\phi}, \vec{\epsilon}$) are calculated by minimizing the negative log-likelihood of the measured data:

$$-L\left(\vec{\mu}, \Delta\vec{\phi}, \vec{\epsilon}\right) = \sum_{i=1}^{N} \sum_{r=1}^{R} \hat{y}_i^r - y_i^r \ln\left(\hat{y}_i^r\right), \tag{8.5}$$

where R is the number of stepping positions and y_i^r denotes the number of photon counts measured for detector pixel i and stepping position r. Similarly to spectral imaging, this optimization problem is separable with respect to the detector pixels. From the differential phase shift, the gradient of the PED can be calculated:

$$\Delta\phi_i = \mathcal{S}\left(\rho_\mathrm{e}^{i+1} - \rho_\mathrm{e}^i\right), \tag{8.6}$$

where the sensitivity $\mathcal{S}$ represents the conversion factor between a PED difference and the corresponding phase shift of the stepping curve. The higher the sensitivity, the larger the phase shift for a given electron density difference between two neighboring pixels. Assuming that the

sample is placed between the G1 and G2 grating (compare figure 8.1), the sensitivity is given by [275]:

$$\mathcal{S} = \frac{r_e d}{p_2 a}\left(1 - \frac{l}{d}\right)\left(\frac{hc}{\mathcal{E}_{\text{eff}}}\right)^2, \tag{8.7}$$

where r_e is the classical electron radius, d is the distance between the G1 and G2 grating, l is the distance between the G1 grating and the object and p_2 is the period of the G2 grating. The parameter a represents the effective pixel size (i.e. the detector pixel size divided by the geometrical magnification of the setup) and $\mathcal{E}_{\text{eff}}$ is the effective energy of the setup. In analogy to attenuation-based imaging, it can be defined as [23]:

$$\mathcal{E}_{\text{eff}} = \left(\frac{\int_0^\infty t(\mathcal{E})\mathcal{R}(\mathcal{E})e^{-\mu(\mathcal{E})}V(\mathcal{E})\mathcal{E}^{-2}d\mathcal{E}}{\int_0^\infty t(\mathcal{E})\mathcal{R}(\mathcal{E})e^{-\mu(\mathcal{E})}V(\mathcal{E})d\mathcal{E}}\right)^{-\frac{1}{2}}, \tag{8.8}$$

where $\mu(\mathcal{E})$ and $V(\mathcal{E})$ represent the energy-dependent attenuation of the object and the energy-dependent visibility of the interferometer, respectively. Assuming a known value of the PED at the left and right edges of the detector (for simplicity we assume that the open beam hits the edges of the detector, i.e. $\rho_\text{e}^1 = \rho_\text{e}^N = 0$), the PED for an arbitrary pixel can be calculated by integration of the differential phase shifts:

$$\rho_\text{e}^i = \frac{1}{\mathcal{S}}\sum_{q=1}^{i-1}\Delta\phi_q. \tag{8.9}$$

Alternatively, it would be possible to integrate starting from the right side of the detector:

$$\rho_\text{e}^i = -\frac{1}{\mathcal{S}}\sum_{q=i}^{N-1}\Delta\phi_q. \tag{8.10}$$

However, these integration strategies are suboptimal for our goal of providing an in-depth noise analysis of all three imaging methods. The summation introduces a strong spatial dependency of the PED variance, even for a homogeneous sample (i.e. a sample that produces the same expected stepping curve measurements for each detector pixel). In this case, the variance of the PED increases linearly from left to right or right to left (if using equation (8.9) or equation (8.10), respectively). However, error propagation calculations show (see appendix) that the variance remains constant (even for nonhomogeneous samples) if the electron density is calculated as the average of both summation strategies:

$$\rho_\text{e}^i = \frac{1}{2\mathcal{S}}\left(\sum_{q=1}^{i-1}\Delta\phi_q - \sum_{q=i}^{N-1}\Delta\phi_q\right). \tag{8.11}$$

In the appendix, we show that the electron density variance is reduced by a factor of two when compared to variance obtained using equation (8.9) or (8.10).
SDPC imaging can be viewed as a combination of spectral and DPC imaging. The spectral and phase contrast information is used simultaneously by acquiring energy-resolved stepping curve measurements. The expected number of photon counts $\hat{y}_i^{rs}$ for detector pixel i, stepping position r and energy bin s can be modeled as:

$$\hat{y}_i^{rs} = \int_0^\infty t(\mathcal{E})\mathcal{R}^s(\mathcal{E})e^{-A_1^i f_1(\mathcal{E}) - A_2^i f_2(\mathcal{E})}\left[1 + V(\mathcal{E})e^{-d_\epsilon^i f_\epsilon(\mathcal{E})}\cos\left(\phi_r(\mathcal{E}) + \Delta\phi_i(\mathcal{E})\right)\right]d\mathcal{E}, \tag{8.12}$$

where d_ϵ^i is the line integral of an artificial dark-field basis material (see [406] for a more detailed explanation) and $f_\epsilon(\mathcal{E})$ is the corresponding energy-dependency of the dark-field signal. Compared to the forward model of DPC imaging (equation 8.4), the polychromatic spectrum is taken into account and thus the visibility $V(\mathcal{E})$, reference phase $\phi_r(\mathcal{E})$ and the phase

shift $\Delta\phi_i(\mathcal{E})$ become energy-dependent. In a real experiment, all setup-dependent quantities $(t(\mathcal{E}), \mathcal{R}^s(\mathcal{E}), V(\mathcal{E}), \phi_r(\mathcal{E}))$ can depend on the spatial position and therefore on the detector pixel index i. For simplicity and clarity, we have omitted this possible dependency in this section. Similarly to DPC imaging, the phase shift depends on the gradient of the PED:

$$\Delta\phi_i(\mathcal{E}) = \frac{r_\mathrm{e} d}{p_2 a}\left(1-\frac{l}{d}\right)\left(\frac{hc}{\mathcal{E}}\right)^2\left(\rho_\mathrm{e}^{i+1}-\rho_\mathrm{e}^{i}\right) \equiv \mathcal{S}(\mathcal{E})\left(\rho_\mathrm{e}^{i+1}-\rho_\mathrm{e}^{i}\right), \tag{8.13}$$

where $\mathcal{S}(\mathcal{E})$ is the energy-dependent sensitivity of the setup. The key idea behind connecting spectral and DPC imaging and eliminating the PED as an additional optimization variable is expressing the PED as the sum of the projected basis material electron densities:

$$\rho_e^i = A_1^i \rho_\mathrm{e}^\mathrm{V}(\mathrm{M}_1) + A_2^i \rho_\mathrm{e}^\mathrm{V}(\mathrm{M}_2). \tag{8.14}$$

Two basis materials images $(\vec{A}_1, \vec{A}_2)$ and a dark-field $(\vec{d}_\epsilon)$ image can be reconstructed by minimizing the following negative log-likelihood function:

$$-L(\vec{A}_1, \vec{A}_2, \vec{d}_\epsilon) = \sum_{i=1}^{N}\sum_{s=1}^{S}\sum_{r=1}^{R} \hat{y}_i^{rs} - y_i^{rs}\ln\left(\hat{y}_i^{rs}\right). \tag{8.15}$$

As the forward model depends on the basis material thicknesses and their spatial gradients (via the gradient of the PED), the log-likelihood cannot be optimized separately for each detector pixel.

8.1.2 Noise analysis with the Cramér Rao lower bound (CRLB)

The CRLB is a powerful tool from estimation theory that predicts a lower bound for the variance of an unbiased estimator. Given a parameter vector $\vec{a}$, it can be shown that [97]:

$$\boldsymbol{C}(\vec{a}) - [\boldsymbol{F}(\vec{a})]^{-1} \geq 0, \tag{8.16}$$

i.e. the matrix $\boldsymbol{C}(\vec{a}) - [\boldsymbol{F}(\vec{a})]^{-1}$ is positive semidefinite. Here, $\boldsymbol{C}(\vec{a})$ is the covariance matrix of $\vec{a}$:

$$\boldsymbol{C}_{uv} = \mathrm{E}\left[(a_u - E(a_u))(a_v - E(a_v))\right], \tag{8.17}$$

where $E(\cdot)$ denotes the expectation value. The Fisher information matrix $\boldsymbol{F}(\vec{a})$ is the expectation value of the curvature of the negative log-likelihood function:

$$\boldsymbol{F}_{uv} = \mathrm{E}\left[-\frac{\partial^2 L(\vec{a})}{\partial a_u \partial a_v}\right]. \tag{8.18}$$

From equation 8.16, a lower bound for the variance of the estimated parameters can be deduced:

$$\boldsymbol{C}_{uu} = \sigma^2(a_u) \geq (\boldsymbol{F}^{-1})_{uu}. \tag{8.19}$$

Since the ML estimator is unbiased and achieves the CRLB in the limit of low noise levels, the CRLB can be used to predict the noise levels for all three imaging methods under consideration. As will be shown later, the CRLB is a good predictor of the noise levels for clinically realistic photon statistics.

In the case of spectral imaging, the optimization problem is separable with respect to different pixels. We therefore obtain N 2×2 (inverse) Fisher matrices. For simplicity, the dependency on the pixel index i is suppressed in the following. As derived by Roessel and Herrmann [227], the Fisher matrix is given by:

$$\boldsymbol{F}_{uv} = \sum_{s=1}^{S} \frac{1}{\hat{y}^s}\frac{\partial \hat{y}^s}{\partial A_u}\frac{\partial \hat{y}^s}{\partial A_v}, \tag{8.20}$$

where $u, v \in (1, 2)$ are the basis material indices and

$$\frac{\partial \hat{y}^s}{\partial A_u} = -\int_0^\infty t(\mathcal{E})\mathcal{R}^s(\mathcal{E})e^{-A_1 f_1(\mathcal{E})-A_2 f_2(\mathcal{E})} f_u(\mathcal{E})d\mathcal{E}. \tag{8.21}$$

The elements of the Fisher matrix can be rewritten in a slightly more intuitive form:

$$\boldsymbol{F}_{uv} = \sum_{s=1}^{S} \hat{y}^s \bar{f}_u^s \bar{f}_v^s, \quad \bar{f}_u^s = \frac{1}{\hat{y}^s}\frac{\partial \hat{y}^s}{\partial A_u}, \tag{8.22}$$

where $\bar{f}_u^s$ is the weighted average attenuation caused by basis material u in energy bin s. The weights for each energy bin are determined by the effective spectra (including object attenuation). The lower bound for the variances of the basis material line integrals (see equation 8.19) is then calculated using the analytical inversion formula for 2×2 matrices. From the variances of the basis material thicknesses, the variance of the PED can be calculated by standard error propagation:

$$\sigma^2(\rho_\mathrm{e}) = \rho_\mathrm{e}^\mathrm{V}(\mathrm{M}_1)^2\sigma^2(A_1) + \rho_\mathrm{e}^\mathrm{V}(\mathrm{M}_2)^2\sigma^2(A_2) + 2\rho_\mathrm{e}^\mathrm{V}(M_1)\rho_\mathrm{e}^\mathrm{V}(M_2)\mathrm{Cov}(A_1, A_2), \tag{8.23}$$

where the covariance of the two basis material thicknesses $\mathrm{Cov}(A_1, A_2)$ is estimated by $(\boldsymbol{F}^{-1})_{12}$. Similarly to the derivation for spectral imaging [227], The Fisher matrix for DPC imaging can be calculated as:

$$\boldsymbol{F}_{uv}(\vec{a}) = \sum_{r=1}^{R} \frac{1}{\hat{y}^r}\frac{\partial \hat{y}^r}{\partial a_u}\frac{\partial \hat{y}^r}{\partial a_v}, \quad \vec{a} = (\mu, \Delta\phi, \epsilon)^T. \tag{8.24}$$

Explicitly calculating the gradients leads to:

$$\boldsymbol{F} = \sum_{r=1}^{R} \frac{1}{\hat{y}^r} \begin{pmatrix} (\hat{y}_r)^2 & \hat{y}^r Q\sin(\phi_r^\mathrm{eff}) & \hat{y}^r Q\cos(\phi_r^\mathrm{eff}) \\ \hat{y}^r Q\sin(\phi_r^\mathrm{eff}) & Q^2\sin^2(\phi_r^\mathrm{eff}) & Q^2\sin(\phi_r^\mathrm{eff})\cos(\phi_r^\mathrm{eff}) \\ \hat{y}^r Q\cos(\phi_r^\mathrm{eff}) & Q^2\sin(\phi_r^\mathrm{eff})\cos(\phi_r^\mathrm{eff}) & Q^2\cos^2(\phi_r^\mathrm{eff}) \end{pmatrix}, \tag{8.25}$$

where $\phi_r^\mathrm{eff} = \phi_r - \Delta\phi$ is the effective phase and $Q = be^{-\mu}Ve^{-\epsilon}$ is the expected intensity ($be^{-\mu}$) multiplied by the effective visibility ($Ve^{-\epsilon}$) of the stepping curve. We are particularly interested in the lower bound for the variance of the differential phase shift $\sigma^2(\Delta\phi)$, as the PED is calculated by integration of the differential phase shifts (compare equation 8.11). It is given by:

$$\sigma^2(\Delta\phi) \geq \left(\boldsymbol{F}^{-1}\right)_{22}. \tag{8.26}$$

For standard phase stepping with $R > 3$ equidistantly distributed steps, the off-diagonal elements $\boldsymbol{F}_{12}, \boldsymbol{F}_{21}, \boldsymbol{F}_{13}, \boldsymbol{F}_{31}$ vanish [279]. Moreover, numerical evaluations show that $\boldsymbol{F}_{23}$ and $\boldsymbol{F}_{32}$ are small compared with the corresponding diagonal entries and vanish in the limit of $R \to \infty$. In this case, a simple interpretation of the lower bound for $\sigma^2(\Delta\phi)$ is possible:

$$\sigma^2(\Delta\phi) \geq \frac{1}{\boldsymbol{F}_{22}} \propto \left[Rbe^{-\mu}\left(Ve^{-\epsilon}\right)^2\right]^{-1}, \tag{8.27}$$

i.e. $\sigma^2(\Delta\phi)$ is inversely proportional to the number of phase steps, the average number of photon counts per step and the effective visibility squared. The same result was obtained for a least-squares estimator (instead of an ML estimator) [279]. By applying standard error propagation techniques to equation 8.11, the variance of the PED can be calculated from the variance of the differential phase shift:

$$\sigma^2(\rho_\mathrm{e}^i) = \frac{1}{4\mathcal{S}^2}\sum_{q=1}^{N-1}\sigma^2(\Delta\phi_q). \tag{8.28}$$

For SDPC imaging, the optimization problem cannot be solved separately for each detector pixel. Consequently, there are 3N optimization variables, which we summarize in the vector $\vec{a} =$

$\left(A_1^1, ..., A_1^N, A_2^1, ..., A_2^N, d_\epsilon^1, ..., d_\epsilon^N\right)^T$ The Fisher matrix has $3N \times 3N$ entries that are calculated similarly to the other two imaging methods:

$$\boldsymbol{F}_{uv} = \sum_{i=1}^{N} \sum_{s=1}^{S} \sum_{r=1}^{R} \frac{1}{\hat{y}_i^{rs}} \frac{\partial \hat{y}_i^{rs}}{\partial a_u} \frac{\partial \hat{y}_i^{rs}}{\partial a_v} \tag{8.29}$$

However, most of the elements of the Fisher matrix are zero because the differential phase shift only couples neighboring pixels. To write the partial derivatives of the forward model more compactly, we define the following abbreviations:

$$\begin{aligned} \alpha_i^s(\mathcal{E}) &= t(\mathcal{E})\mathcal{R}^s(\mathcal{E})e^{-A_1^i f_1(\mathcal{E}) - A_2^i f_2(\mathcal{E})} \\ \beta_i^r(\mathcal{E}) &= V(\mathcal{E})e^{-d_\epsilon^i f_\epsilon(\mathcal{E})} \cos\left(\phi_r(\mathcal{E}) + \Delta\phi_i(\mathcal{E})\right) \\ \gamma_i^r(E) &= V(\mathcal{E})e^{-d_\epsilon^i f_\epsilon(\mathcal{E})} \sin\left(\phi_r(\mathcal{E}) + \Delta\phi_i(\mathcal{E})\right) \end{aligned} \tag{8.30}$$

The non-zero partial derivatives of the forward model ($\hat{y}_i^{rs}$) are given by:

$$\begin{aligned} \frac{\partial \hat{y}_i^{rs}}{\partial A_u^i} &= \int_0^\infty -\alpha_i^s(\mathcal{E})\left(1 + \beta_i^r(\mathcal{E})\right) f_u(\mathcal{E}) + \alpha_i^s(\mathcal{E})\gamma_i^r(\mathcal{E})\mathcal{S}(\mathcal{E})\rho_\mathrm{e}(\mathrm{M}_u)d\mathcal{E}, \quad u \in (1,2) \\ \frac{\partial \hat{y}_i^{rs}}{\partial A_u^{(i+1)}} &= -\int_0^\infty \alpha_i^s(\mathcal{E})\gamma_i^r(\mathcal{E})\mathcal{S}(\mathcal{E})\rho_\mathrm{e}(\mathrm{M}_u)d\mathcal{E}, \quad u \in (1,2) \\ \frac{\partial \hat{y}_i^{rs}}{\partial d_\epsilon^i} &= -\int_0^\infty \alpha_i^s(\mathcal{E})\beta_i^r(\mathcal{E})f_\epsilon(\mathcal{E})d\mathcal{E}. \end{aligned} \tag{8.31}$$

The Fisher matrix for SDPC imaging must be inverted numerically to calculate the CRLB.

8.1.3 Prediction of covariances and noise power spectra

In this section, we explain our approach for predicting covariances and noise power spectra for all three imaging methods. Given an estimate of the covariances of the projected electron densities between different pixels, the noise power spectrum (NPS) can be estimated as:

$$\mathrm{NPS}(k) = \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} e^{2\pi j \frac{v-u}{N}} \mathrm{Cov}(\rho_\mathrm{e}^{(u+1)}, \rho_\mathrm{e}^{(v+1)}), \tag{8.32}$$

where k is the spatial frequency and j is the imaginary unit. In the appendix, we give a quick derivation of this result. For spectral imaging, the material decomposition is conducted separately for each pixel. This means that there are no noise correlations between different pixels:

$$\mathrm{Cov}\left(\rho_\mathrm{e}^u, \tilde{\rho}_\mathrm{e}^v\right) = 0, u \neq v, \quad \mathrm{Cov}\left(\rho_\mathrm{e}^u, \rho_\mathrm{e}^u\right) = \sigma^2(\rho_\mathrm{e}^u). \tag{8.33}$$

For DPC imaging, the differential phase shifts are also uncorrelated. The covariance matrix of the differential phase shifts $\boldsymbol{C}^{\Delta\phi}$ is thus a diagonal matrix with diagonal elements $\boldsymbol{C}_{ii}^{\Delta\phi} = \sigma^2(\Delta\phi_i)$. However, the integration step that is necessary to obtain the projected electron densities (compare equation 8.11) introduces noise correlations. Using error propagation, the covariance matrix C^{ρ_e} for the projected electron densities can be calculated from the covariance of the differential phase shifts:

$$\boldsymbol{C}^{\rho_\mathrm{e}} = \boldsymbol{B}\boldsymbol{C}^{\Delta\phi}\boldsymbol{B}^T, \tag{8.34}$$

where $\boldsymbol{B}$ is the transformation matrix between phase shifts ($\Delta\vec{\phi}$) and projected electron densities ($\vec{\rho}_e$):

$$\vec{\rho}_e = \boldsymbol{B}\vec{\phi}. \tag{8.35}$$

The entries of the matrix $\boldsymbol{B}$ can be deduced form equation 8.11. In the case of SDPC imaging, we use the entries of the inverse Fisher matrix as an estimate for the covariances:

$$\mathrm{Cov}(a^u, a^v) \approx \boldsymbol{F}_{uv}^{-1}, \quad \vec{a} = \left(A_1^1, ..., A_1^N, A_2^1, ..., A_2^N, d_\epsilon^1, ..., d_\epsilon^N\right)^T. \tag{8.36}$$

Similarly to DPC imaging, the covariance matrix for the projected electron densities is calculated as:

$$\boldsymbol{C}^{\rho_e} = \tilde{\boldsymbol{B}}\boldsymbol{F}^{-1}\tilde{\boldsymbol{B}}^T, \quad \vec{\rho}_e = \tilde{\boldsymbol{B}}\vec{a}. \tag{8.37}$$

The entries of the transformation matrix $\tilde{\boldsymbol{B}}$ can be deduced from equation 8.14.

8.1.4 Numerical simulation

To compare the noise level of the projected electron densities for all three imaging methods and to test the predictions of the noise analysis framework, we simulated a radiography measurement of an homogeneous object. The X-ray beam had to penetrate 12 cm of soft tissue and 1 cm of cortical bone, which corresponds to a PED of $12\ \mathrm{cm} \cdot 3.52 \cdot 10^{23}\ \mathrm{cm}^{-3} + 1\ \mathrm{cm} \cdot 5.95 \cdot 10^{23}\ \mathrm{cm}^{-3} = 4.82 \cdot 10^{24}\ \mathrm{cm}^{-2}$. The aforementioned thicknesses could reflect typical path lengths for medical imaging tasks (e.g. thorax radiography or head CT) . The simulated object had no internal microstructure, i.e. it did not generate a dark-field signal. We assumed a tungsten X-ray source and a PCD with a pixel size of 200 μm, a 2 mm thick cadmium telluride sensor and two thresholds per pixel. As explained in the methods section, only one detector row (containing 400 pixels) was simulated. For DPC and SDPC imaging, a symmetric three-grating interferometer (operated at the first fractional Talbot order) was inserted into the beam path. The attenuation gratings (G0 and G2) were assumed to be made of gold with a grating height of 200 μm. The G1 grating was assumed to be made of nickel and to induce a phase shift of $\pi/2$ for the design energy. Although the duty cycle of the gratings influences the noise level [414], it was kept fixed at its standard value (0.5) in this simulation study in order to reduce the number of possible acquisition parameter combinations. The total length of the simulated setup was 2.4 m and the object was placed between G1 and G2 (40 cm away from G1). Figure 8.1 shows an overview of the setup geometry that was kept fixed for all three imaging methods. All other imaging

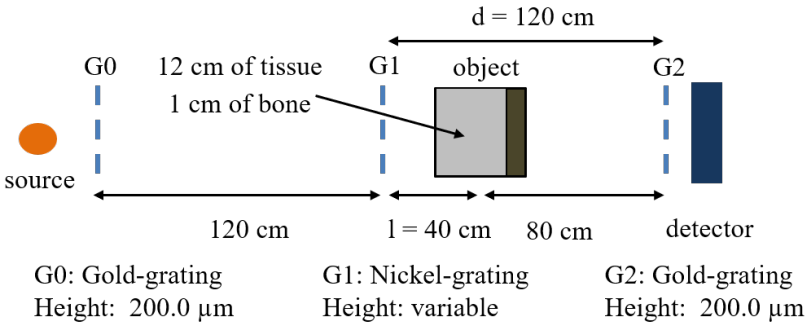

Figure 8.1: Overview of the setup geometry that is kept fixed for all three imaging methods. The three gratings are removed for spectral imaging.

parameters (acceleration voltage, threshold positions and design energy of the interferometer) were optimized individually for each imaging method (if applicable). For all three imaging methods, the low threshold of the PCD was kept fixed at 15 keV. The position of the high threshold was optimized for SDPC and spectral imaging while only one threshold (at 15 keV) was used for DPC imaging. For DPC and SDPC imaging, the design energy was tuned by changing the height of the phase shifting grating as well as the grating periods. Changing the grating periods is necessary to keep the first fractional Talbot distance at the detector position. Stepping curves (using 5 steps equally distributed between 0 and 2π) were simulated with a

wave-optical simulation package based on Fresnel propagation and the projection approximation (see chapter 3). For all three imaging methods, the variable setup parameters were optimized by minimizing the variance of the PED predicted by the CRLB while keeping the dose to the object constant (1 mGy). The dose was estimated based on an empirical model that was fitted to Monte-Carlo simulations [363]. Since the G0 and G1 gratings are placed between the source and the object, they only influence the (spectral) photon flux incident on the object. The G2 grating, however, attenuates photons that have (potentially) contributed to the dose delivered to the object because it is placed between the object and the detector.

8.2 Results

The method for finding the optimum setup parameters is illustrated in figure 8.2 (a) and (b). Figure 8.2 (a) shows a contour plot of the predicted PED variance for SDPC imaging (at the optimum threshold position) as a function of the acceleration voltage and the design energy. As demonstrated in the last section, the PED for the simulated decomposition task is on the order of 10^{24} cm^{-2}, while the corresponding standard deviations are approximately two orders of magnitude lower. This explains the larger numerical values for the PED variances of up to 10^{46} cm^{-4}. The minimum variance is achieved for an acceleration voltage of 140 kVp and a design energy of 60 keV. As the acceleration voltage is decreased, the optimum design energy also decreases. Figure 8.2 (b) shows a similar contour plot for DPC imaging. In comparison with SDPC imaging, the optimum acceleration voltage is much lower (for this particular decomposition task). However, similarly to SDPC imaging, the optimum design energy is positively correlated with the acceleration voltage. This behavior is reasonable because a large part of the spectrum should be concentrated around the design energy of the interferometer for optimum performance. Table 8.1 shows the optimum setup parameters as well as the minimum PED

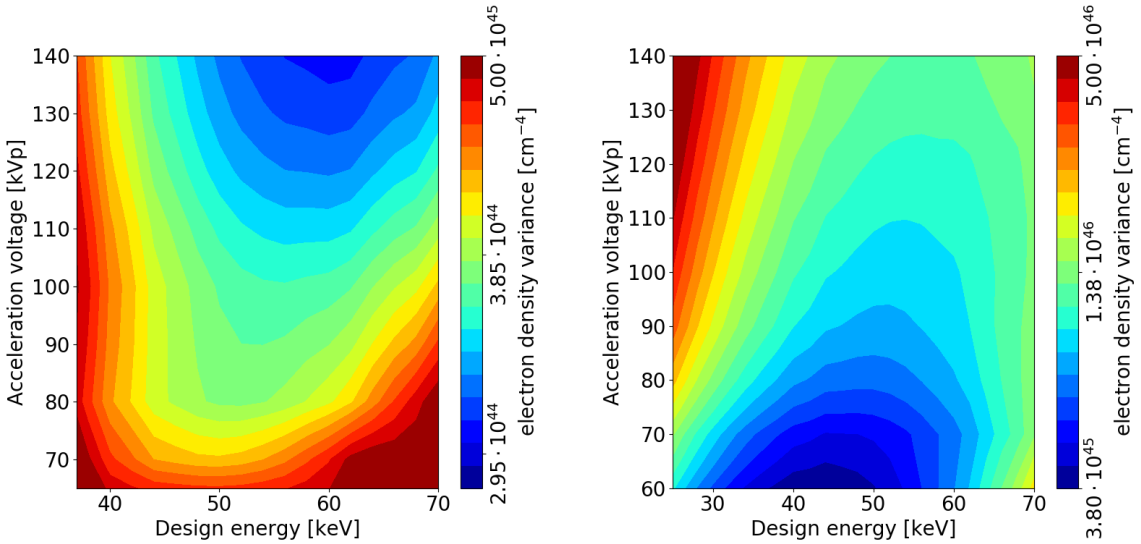

(a) (b)

Figure 8.2: (a): Contour plot of the predicted projected electron density (PED) variance (at the optimum threshold position) for SDPC imaging as a function of the acceleration voltage and the design energy. (b): Predicted PED variance for DPC imaging as a function of the acceleration voltage and the design energy.

variances for all three imaging methods. Interestingly, the optimum setup parameters for SDPC imaging are very similar to those for spectral imaging, while DPC imaging favors a considerably lower acceleration voltage (60 kVp) and design energy (44 keV). A possible explanation for

Table 8.1: Optimum imaging parameters and PED variances for spectral, DPC and SDPC imaging. The optimum parameters were determined by minimizing the PED variance predicted by the CRLB while keeping the dose to the object constant (1 mGy).

	optimization range	optimum parameters		
		spectral	DPC	SDPC
acceleration voltage [kVp]	$40-140$	140	60	140
high threshold [keV]	$15-140$	63	-	65
design energy [keV]	$20-70$	-	44	60
G2 grating period [μm]	$9.22-17.25$	-	11.63	9.96
min. PED variance [cm^{-4}]		$1.77 \cdot 10^{45}$	$3.82 \cdot 10^{45}$	$3.04 \cdot 10^{44}$

these results can be found by looking at figure 8.3, which shows the energy-dependent visibility as well as the effective spectra for DPC and SDPC imaging with optimum setup parameters. The effective spectrum includes the source spectrum, the attenuation of the gratings and the detector response. Kottler et. al [401] demonstrated that a grating interferometer with a $\pi/2$ phase shifting grating (G1) exhibits a second visibility peak at twice the design energy. The second important factor that influences the visibility is the attenuation of the gold gratings. In the range from $60-80\,\mathrm{keV}$, the gold gratings become increasingly transparent, but for energies above the K-edge of gold (80.7 keV), the attenuation of the gratings and thus also the visibility rises sharply. In the case of DPC imaging, a lower effective energy (see equation 8.8) is preferable

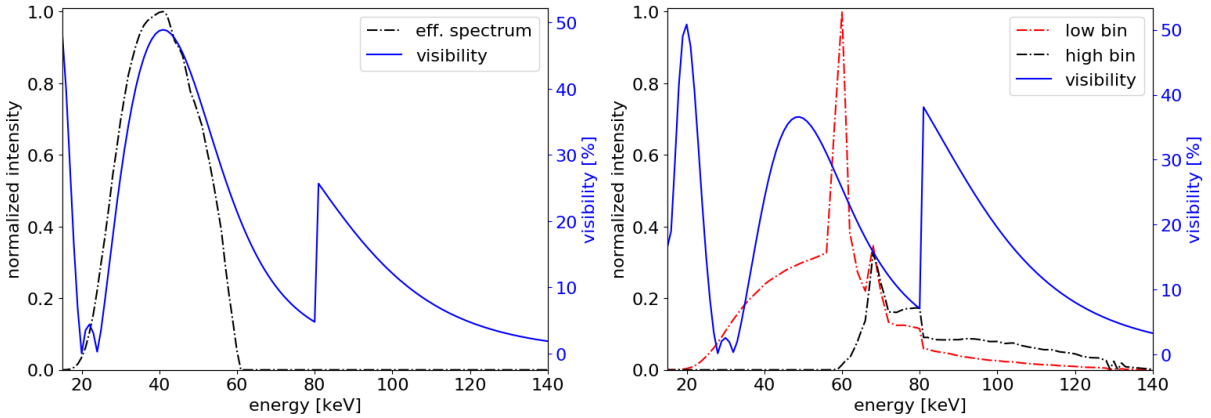

(a) (b)

Figure 8.3: Energy-dependent visibilities and effective spectra for DPC (a) and SDPC (b) imaging using the optimum acquisition parameters given in table 8.1.

because it increases the sensitivity of the setup (compare equation 8.6). In other words, a fixed gradient of the PED causes a larger phase shift for lower effective energies. Since the spatial resolution and the setup geometry are fixed for our simulation, the sensitivity can only be tuned via the design energy and the effective energy. The acceleration voltage (60 kVp) is matched to the design energy (44 keV) so that a large fraction of the effective spectrum is concentrated around the visibility peak. Achieving a high visibility is important because for DPC imaging, the variance of the PED is proportional to the visibility squared (compare equation 8.27). The comparatively high acceleration voltage (140 kVp) for SDPC imaging leads to a decreased sensitivity compared with DPC imaging. However, this decrease in sensitivity is compensated by the improved performance of spectral imaging at higher acceleration voltages. Moreover, the decrease in sensitivity and visibility compared with DPC imaging is less pronounced for the low energy bin. Due to the higher acceleration voltage, the optimum design energy for SDPC

imaging is considerably larger (60 keV). Consequently, the visibility peak for low energies is reduced compared with DPC imaging. On the other hand, the visibility for energies larger than the K-edge of Gold (80.7 keV) is increased.
The predicted PED variances show that for this particular imaging task, SDPC imaging can achieve considerably lower noise levels compared with spectral and DPC imaging (variance reduction by a factor of 5.8 and 12.6. respectively). It is important to note that for DPC imaging, the variance depends approximately linearly on the number of pixels in one detector row. This behavior is caused by the integration that is required to convert the differential phase shifts to a PED profile (see equation 8.11). For SDPC imaging, the PED variance only depends very weakly on the number of pixels. The reason for this weak dependence will be explained later when we analyze the noise correlations for all three imaging methods.

We conducted numerical experiments with 100 different dose levels ranging from 0.001 to

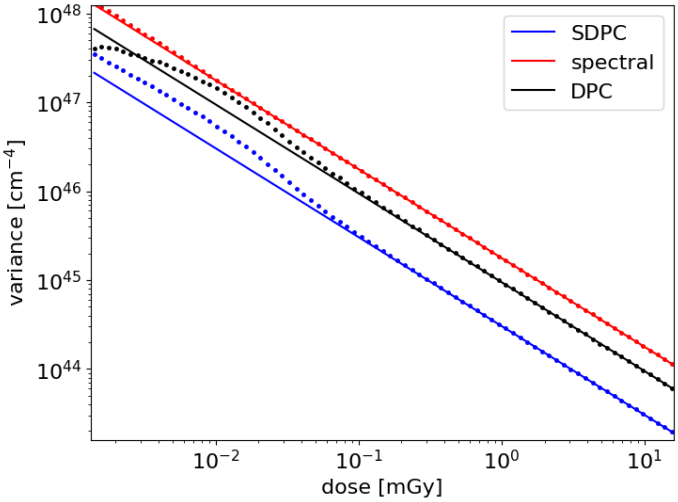

Figure 8.4: Experimentally obtained PED variances (dots) and theoretical predictions (solid lines) for all three imaging methods as a function of the dose level.

10 mGy in order to test if the ML-based estimators achieves the predicted PED variances. For these experiments, we reduced the detector size to 100 pixels to reduce the computational time. For each dose level and imaging method, 10,000 different noise realizations were simulated and the ML-based estimators (compare equations 8.2,8.5,8.15), were used to calculate the PED. Figure 8.4 shows the experimentally achieved PED variances (dots) together with the variances predicted by the CRLB (solid lines) for all three imaging methods. All estimators achieve the CRLB for a large range of dose levels ($\approx 0.1 - 10$ mGy). For lower dose levels, deviations from the CRLB are observable, especially for DPC and SDPC imaging. In the case of DPC imaging, a similar dependency of the variance on the dose level (or photon statistics) has been reported for a Fourier-based estimator [391]. In this simulation, DPC imaging achieves a lower variance than spectral imaging because of the reduced detector size. However, the main goal of this simulation was to test the predictions of the noise analysis framework rather than comparing the performance of the three imaging methods.
It has already been pointed out that the integration step for DPC imaging causes long-range correlations between the PED values for one detector row. Consequently, the NPS for DPC imaging is dominated by low frequencies. In the case of spectral imaging, the basis material line integrals can be determined separately for each detector pixel and thus the PED values determined by spectral imaging are uncorrelated (at least for an ideal PCD). Therefore, the noise power should be equally distributed between all frequencies ("white noise"). Since SDPC imaging is a combination of both aforementioned imaging methods, it could be expected that the NPS for SDPC imaging is a mixture between the noise power spectra for the other two

imaging methods. In section 8.1.3, we have presented a framework for predicting the PED covariances as well as the noise power spectra for all three imaging methods. We conducted another numerical experiment to verify the predicted covariances and noise power spectra. In this simulation, the dose was fixed at 1 mGy and the original detector size (400 pixels) was used. For each of the three imaging methods, 50,000 different noise realizations were simulated and the PEDs were calculated separately for each noise realization. Figure 8.5 (a) shows the experimentally calculated normalized covariances (dots) together with the theoretical predictions (solid lines) for all three imaging methods. More precisely, the covariances between the PED for the central detector pixel (i.e. pixel index 200) and the PEDs of all other pixels is shown. In figure 8.5 (b), the experimentally obtained noise power spectra are plotted together with the theoretical predictions. We used equation 8.42 for the experimental calculation of the NPS. For both the covariances and the noise power spectra, the theoretical predictions are in good

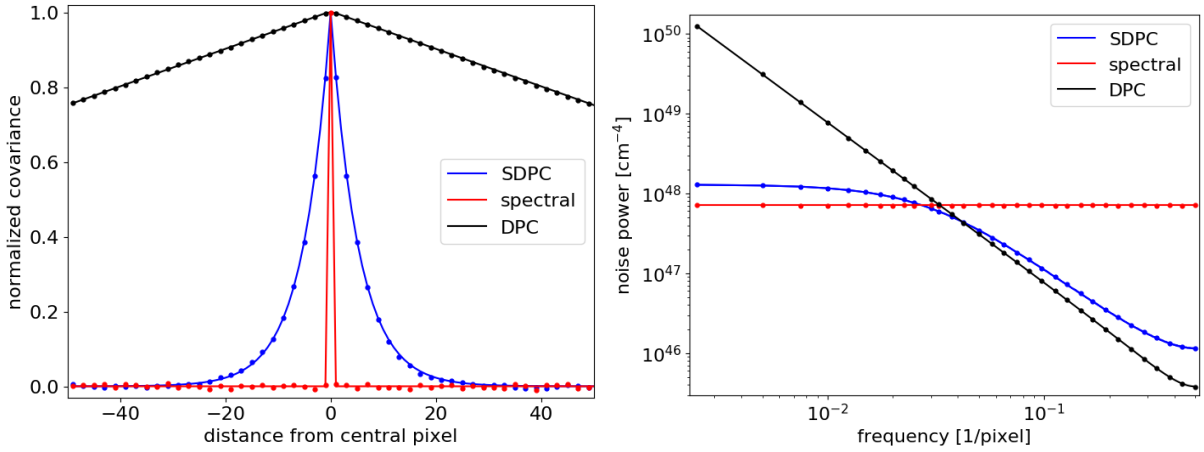

(a) (b)

Figure 8.5: (a): Experimentally calculated normalized covariances between the central detector pixel and the other pixels (dots) together with the theoretical predictions (solid lines). (b): Noise power spectra for all three imaging methods obtained in the numerical experiment (dots) and the corresponding theoretical predictions (solid lines).

agreement with the experimentally obtained values. As expected, the covariance of the PED between two different detector pixels is zero for spectral imaging. For DPC imaging, the covariance decreases linearly with the distance from the central detector pixel until it drops to zero at the edges of the detector (not shown in figure 8.5 (a)). This reflects the well-known long-range noise correlations introduced by the integration step. In the case of SDPC imaging, however, the covariance rapidly drops to zero with increasing distance from the central pixel. For the imaging parameters given in table 8.1, almost no noise correlations are observable if the distance to the central pixel is larger than approximately 25 pixels. Compared to DPC imaging, the additional spectral information eliminates long-range noise correlations. On the other hand, it could be argued that compared with spectral imaging, the additional phase shift term for SDPC imaging couples the PED values of neighboring pixels and thus introduces local noise correlations. The covariance graph for SDPC imaging explains the aforementioned weak dependence of the PED variance on the number of detector pixels. Contrary to DPC imaging, the covariance rapidly decreases with the distance between two pixels and thus the assumption of fixed, known PED values at the edges of the detector does not influence the noise level (except for pixels close to the edges). The different noise correlations for the three imaging methods are also reflected in the different noise power spectra. As expected, the noise power is independent of the frequency for spectral imaging, whereas the noise power for DPC imaging increases drastically for lower spatial frequencies. For high frequencies, the noise power spectrum for SDPC imaging is similar

to DPC imaging. For low frequencies, however, the noise power does not increase but converges to a constant value that lies slightly above the noise power graph for spectral imaging. The higher noise power for SDPC imaging in the low frequency area can be explained by the attenuation of the G2-grating, which removes a part of the X-ray beam that has contributed to the dose delivered to the object.
In our simulation study, we assumed a fixed total length of the setup and a symmetric grating interferometer that operates in the first fractional Talbot distance. Given these conditions, the effective pixel size is the most important tuning factor for the sensitivity of a DPC or SDPC imaging setup (compare equation 8.7 and 8.13). Since the PED variance strongly depends on the sensitivity, we investigated the theoretically predicted performance of all three imaging methods as a function of the effective pixel size. For each pixel size, the optimum imaging parameters were determined with the noise analysis framework presented in the methods section. Figure 8.6 (a) shows the predicted PED variances as a function of the effective pixel size. For visualization purposes, the number of photon counts per pixel (instead of the dose delivered to the object) was kept constant. This means that the dose delivered to the object was chosen inversely proportional to the squared effective pixel size (i.e. the effective pixel area). If the dose was kept constant, the PED variances would have to be multiplied by an additional factor that is inversely proportional to the effective pixel area for all three imaging methods. Under these

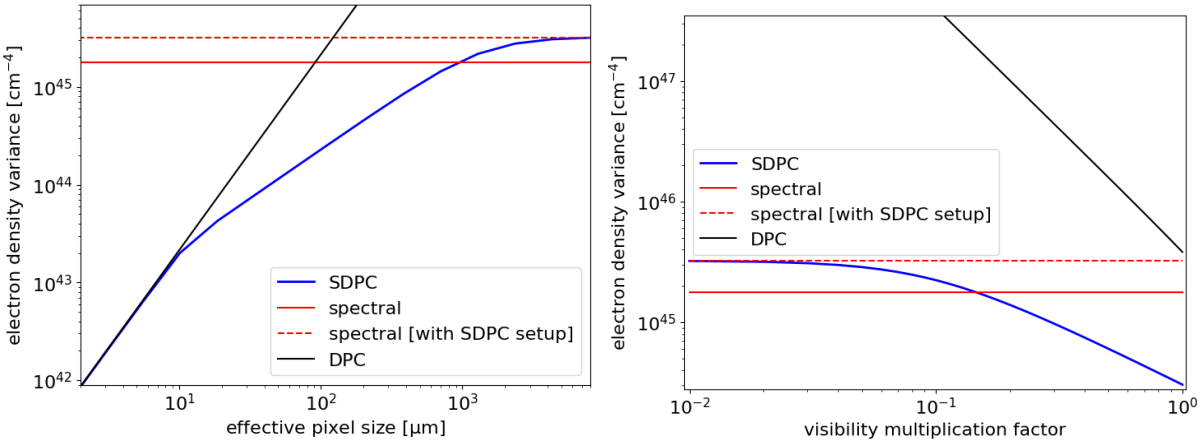

(a) (b)

Figure 8.6: (a): Minimum PED variances for all three imaging methods as a function of the effective pixel size. For visualization purposes, the number of photon counts per pixel (instead of the dose delivered to the object) was kept constant.
(b): PED variances with standard resolution and imaging parameters (see table 8.1) as a function of the visibility. For this experiment, the energy-dependent visibilities (see figure 8.3) were multiplied by a visibility reduction factor m ($0 < m < 1$) to simulate a decreased visibility due to experimental effects.

conditions, the PED variance for spectral imaging is independent of the pixel size. In the case of DPC imaging, the variance is proportional to the squared effective pixel size, which reflects the well-known sensitivity-dependent noise level for DPC imaging (compare also equation 8.28). For large effective pixel sizes ($a > 4$ mm), the PED variance for SDPC imaging is almost constant and approximately twice as large as the variance for spectral imaging. It coincides with the PED variance one would obtain if spectral imaging was performed with the "unnecessary" grating interferometer inserted into the beam path (see dashed line in figure 8.6 (a)). It follows that for larger pixel sizes (or low sensitivities), the additional information on the basis material line integrals provided by the phase shift of the stepping curve can be neglected compared to the spectral information. Consequently, SDPC imaging reduces to spectral imaging in the limit of

large effective pixel sizes. For this particular imaging task, SDPC imaging and spectral imaging have equal PED variances for an effective pixel size of $a = 950$ μm. At this point, the benefit of the additional phase shift information is counterbalanced by the reduced photon statistics caused by the attenuation of the G2 grating. For very small pixel sizes ($a < 10$ μm), the SDPC variances almost coincide with the DPC variances. In this case, SDPC imaging reduces to DPC imaging because the spectral information provides little additional value for determining the PED in comparison to the phase shift of the stepping curve. In the area between these limiting cases (10 μm $< a < 950$ μm), the SDPC variance lies below both the spectral and the DPC variance because both the spectral and the phase shift information are relevant for determining the PED. The largest variance reduction (by a factor of 8.8) compared to both DPC and spectral imaging is achieved for the point of equal variance of spectral and DPC imaging ($a = 90$ μm). At this point, the spectral and the phase shift information are of equal importance for determining the PED.

It is important to note that the predicted PED variance only depends on the geometrical imaging parameters (pixel size, object position, grating distances, grating periods) via the sensitivity. Consequently, similar trends for the PED variances could have been observed by varying geometrical parameters other than the pixel size. However, increasing the sensitivity via these geometrical parameters also increases the phase shift of the stepping curve that a given sample causes (via the gradient of the PED). For DPC imaging, this might lead to additional phase wrapping artifacts. As will be discussed in the next section, phase wrapping artifacts can in theory be avoided for SDPC imaging.

The number of detector pixels ($N = 400$) was kept constant in this analysis of the noise characteristics. This means that the field of view decreases with decreasing effective pixel size a. If the field of view was kept constant, the total number of detector pixels would have to be scaled by $1/a$. As discussed in the results section, increasing the number of detector pixels would have an adverse effect on the PED variance for DPC imaging, but the influence on SDPC imaging would be negligible as long as the spectral information prevents long-range noise correlations.

In a real experiment, various undesirable effects (e.g. vibrations, grating imperfections) could cause a reduced visibility of the stepping curves compared with the simulations. Moreover, we have assumed that the object does not generate a dark-field signal which would also reduce the effective visibility. We therefore investigated the influence of a visibility reduction on the PED variance for SDPC and DPC imaging. For this study, the imaging parameters were kept fixed at their optimum values (compare table 8.1). We simulated a visibility reduction by multiplying the energy-dependent visibilities (compare figure 8.3) with a visibility reduction parameter m ($0 < m < 1$). Figure 8.6 (b) shows the predicted PED variances for SDPC and DPC imaging as a function of the visibility reduction parameter. For comparison, the variances for spectral imaging with and without the interferometer in the beam path are also plotted as constant lines. As expected, the variance for DPC imaging rises rapidly with decreasing visibility ($\sigma^2(\rho_\mathrm{e}) \propto (Ve^{-\epsilon})^{-2}$, compare equation 8.27). In the case of SDPC imaging, however, the variances rises much more slowly with decreasing visibility (approximately proportionally to $(Ve^{-\epsilon})^{-1}$) before it approaches a constant value for low visibilities. This is reasonable because in the limit of low visibilities, SDPC imaging effectively reduces to spectral imaging (with the gratings in the beam path). For $m = 0.15$, the variances for SDPC imaging and spectral imaging are equal. This means that for this particular reconstruction task, the visibility can be reduced by a factor of $1/m = 6.7$ compared with the idealized simulation before the attenuation of the G2 grating outweighs the variance reduction achieved by the additional phase shift information.

8.3 Discussion

A quantitative comparison between conventional attenuation imaging and DPC imaging is difficult because of the different underlying contrast generating mechanisms for these two imaging

methods. For example, a contrast-to-noise ratio comparison [15, 16] is complicated because the differences in contrast for the two imaging methods depends on the choice of the investigated materials. Moreover, depending on the setup parameters, the contrast between different materials can vanish completely for both imaging methods. For this reason, we concentrated on electron density images, which can be calculated with all three imaging methods (spectral, DPC and SDPC imaging) under consideration in this comparison study. Nevertheless, the presented noise analysis framework can also be used to determine the optimum acquisition parameters for other imaging tasks, such as dark-field imaging or basis material decomposition. Although we have focused on projection imaging, the noise analysis framework could be generalized to computed tomography imaging (e.g. by using error propagation for the filtered backprojection [415]).

For the investigated imaging task, our analysis predicts highly reduced noise levels for SDPC imaging in comparison to both DPC imaging and spectral imaging. To achieve a fair comparison, the imaging parameters were optimized individually for each of the three methods. For a large range of clinically realistic dose levels, the ML-based estimators achieve the predicted noise levels for all three imaging methods. The validity of the predicted covariances and noise power spectra was confirmed by additional numerical simulations. Our analysis focused on the projected electron density, but earlier studies have shown that the noise advantage of SDPC imaging compared to spectral imaging is also transferred to the basis material images [406]. In our study, DPC imaging only outperforms spectral imaging for comparatively high spatial resolutions ($a < 90$ μm), however this conclusion does not apply to SDPC imaging. For a large range of clinically relevant effective pixel sizes (up to 950 μm), SDPC imaging theoretically outperforms the other two imaging methods by simultaneously using both the spectral and phase contrast information. Although we have only considered one imaging scenario in this work, the general trends should apply to a large range of imaging tasks. Nevertheless, the location of the break-even points between the three imaging methods will vary depending on the object size and the chemical composition. In a real experiment, both grating imperfections and ultra-small angle scattering by the object (i.e. nonzero dark-field signals) could reduce the visibility of the stepping curves compared to our simulations. The noise reduction for SDPC imaging compared to spectral imaging will thus be smaller than theoretically predicted. Nonetheless, because of the large theoretical noise reduction factors, we believe that SDPC imaging can also outperform the other two imaging methods in real experiments. Moreover, the noise level for SDPC imaging increases less rapidly with decreasing visibility compared with DPC imaging (see figure 8.6 (b)). Regardless of the potential for noise reduction, the combination of spectral and phase-contrast imaging provides additional information that is inaccessible with the individual imaging methods. Compared with spectral imaging, the dark-field image yields additional information about the object's microstructure and, compared with DPC imaging, two basis material images can be calculated and beam hardening artifacts in all three imaging channels can be corrected.

Raupach and Flohr [416] state that the noise correlations for DPC imaging are disadvantageous for clinical diagnosis. They argue that the dominance of low frequencies in the NPS is unfavourable for recognizing structures. Furthermore, because of the low-frequency noise, reducing the noise level by pixel binning is less effective compared with conventional attenuation-based imaging. In general, these arguments do not apply to SDPC imaging because of the fundamentally different noise correlations compared to DPC imaging. An exception is the limit of very small pixel sizes (or high sensitivities), where SDPC imaging effectively converges to DPC imaging. In our study, the influence of the spectral information causes a rapid decrease of the covariance between two pixels with increasing distance. Consequently, there is no excessive low frequency noise. Furthermore, Raupach and Flohr [391] argue that - contrary to conventional attenuation imaging - there is a dose limit for DPC imaging below which no meaningful signals can be extracted. The interferometric measurement acquisition process only allows unique determination of the phase shift of the stepping curve across a 2π interval. Larger phase shifts are

wrapped back into this interval. According to Raupach and Flohr, statistical phase wrapping leads to a loss of the phase shift information for low dose levels. More precisely, the average extracted phase shift is biased to zero, regardless of the underlying physical phase shift if a standard averaging procedure is used. As a result of the additional spectral information when compared to DPC imaging, the calculated phase shift is not restricted to a 2π interval for SDPC imaging. Qualitatively speaking, the spectral information determines in which 2π interval the phase shift lies (via the PED profile) and the exact value is fine-tuned by the stepping curve information. In principle, there is thus no information loss through statistical phase wrapping. However, it has been shown that the log-likelihood function for SDPC imaging has local optima that can be explained by an analogy to the phase wrapping problem for DPC imaging [406]. Although we could avoid the convergence to local optima in previous simulation studies by choosing a suitable initial guess, this strategy is likely to fail at extremely low dose levels. Nevertheless, the convergence to local optima is an optimization problem rather than a fundamental restriction of SDPC imaging because it could in principle be avoided with a global optimization strategy (or by incorporating prior information in the form of a regularization term).

8.4 Conclusion

In this work, we have developed a noise analysis framework that allows the calculation of (co-) variances and noise power spectra for spectral, DPC and SDPC imaging. An important practical application of this framework is finding the optimum imaging parameters for all three methods. Our analysis shows that the combination of spectral and phase-contrast imaging is a promising imaging technique with various advantages compared with the individual methods. SDPC imaging provides additional information compared with both spectral imaging (dark-field image) and DPC imaging (basis material images). Moreover, we demonstrated that SDPC imaging enables a strong noise (or dose) reduction compared with the other two imaging methods for a large range of clinically relevant pixel sizes. Finally, the additional spectral information compared to DPC imaging eliminates excessive low-frequency noise, which can be a major drawback of DPC imaging, especially in projection space.

8.5 Appendix

8.5.1 Noise propagation for the integration step in DPC imaging

Applying standard error propagation techniques to equation 8.9, the variance of the PED for a certain pixel ($\sigma^2(\tilde{\rho}_e^i)$) can be calculated:

$$\sigma^2(\rho_\mathrm{e}^i) = \frac{1}{\mathcal{S}^2} \sum_{q=1}^{i-1} \sigma^2(\Delta\phi_q), \tag{8.38}$$

where $\sigma^2(\Delta\phi_q)$ is the variance of the differential phase shift (which can be calculated using the CRLB). In the special case of a homogeneous sample (i.e. $\sigma^2(\Delta\phi_q) = \text{const.} = \sigma^2(\Delta\phi)\ \forall q$), the average variance of the electron density $\sigma^2_\mathrm{avg}(\rho_\mathrm{e})$ can be calculated as:

$$\begin{aligned}\sigma^2_\mathrm{avg}(\rho_\mathrm{e}) &= \frac{1}{N-2} \sum_{i=2}^{N-1} \sigma^2(\rho_\mathrm{e}^i) = \frac{1}{\mathcal{S}^2(N-2)} \sum_{i=2}^{N-1} \sum_{q=1}^{i-1} \sigma^2(\Delta\phi) \\ &= \frac{1}{\mathcal{S}^2(N-2)} \sum_{i=2}^{N-1} (i-1)\sigma^2(\Delta\phi) = \frac{N-1}{2\mathcal{S}^2}\sigma^2(\Delta\phi).\end{aligned} \tag{8.39}$$

The pixels with index 1 and N are excluded from the average variance calculation because it was assumed that the PED is known with certainty at the edges of the detector. The same result is obtained if the integration from right to left (equation 8.10) is considered. If the left- and right-integrated PED profiles are averaged (see equation 8.11), the variance of the PED is given by:

$$\sigma^2(\rho_{\mathrm{e}}^i) = \frac{1}{4\mathcal{S}^2}\left(\sum_{q=1}^{i-1}\sigma^2(\Delta\phi_q) + \sum_{q=i}^{N-1}\sigma^2(\Delta\phi_q)\right) = \frac{1}{4\mathcal{S}^2}\sum_{q=1}^{N-1}\sigma^2(\Delta\phi_q). \tag{8.40}$$

As can be seen from equation 8.40, the variance of the electron density is independent of the detector pixel index. Consequently, the variance is spatially constant, even for a nonhomogeneous sample. In the special case of a homogeneous sample, $\sigma^2_{\mathrm{avg}}(\rho_{\mathrm{e}})$ is given by:

$$\sigma^2_{\mathrm{avg}}(\rho_{\mathrm{e}}) = \frac{1}{N-2}\sum_{i=2}^{N-1}\sigma^2(\rho_{\mathrm{e}}^i) = \frac{1}{4\mathcal{S}^2}(N-1)\sigma^2(\Delta\phi). \tag{8.41}$$

Comparing equation (8.40) and (8.39), it follows that (in the case of a homogeneous sample) the average variance is reduced by a factor of two when averaging left- and right-integrated electron density profiles.

8.5.2 Calculating the noise power spectrum from the covariances

In this section, we derive how the noise power spectrum (NPS) can be calculated from an estimate of the covariances (compare equation 8.32). For a homogeneous object, the NPS can be calculated as [417]:

$$\mathrm{NPS}(k) = E\left[\left|\mathcal{F}(\bar{\rho}_{\mathrm{e}}(x)\right|^2\right], \quad \bar{\rho}_{\mathrm{e}}(x) = \rho_{\mathrm{e}}(x) - E\left[\rho_{\mathrm{e}}(x)\right], \tag{8.42}$$

where k is the spatial frequency and $\bar{\rho}_e(x)$ is the offset-corrected PED as a function of the spatial coordinate x. In discrete form, the NPS is given by:

$$\begin{aligned}\mathrm{NPS}(k) &= E\left[\left|\sum_{q=0}^{N-1} e^{-2\pi j\frac{qk}{N}}\bar{\rho}_{\mathrm{e}}^{(q+1)}\right|^2\right] = E\left[\left(\sum_{q=0}^{N-1} e^{-2\pi j\frac{qk}{N}}\bar{\rho}_{\mathrm{e}}^{(q+1)}\right)\left(\sum_{u=0}^{N-1} e^{+2\pi j\frac{uk}{N}}\bar{\rho}_{\mathrm{e}}^{(u+1)}\right)\right] \\ &= E\left[\sum_{q=0}^{N-1}\sum_{u=0}^{N-1} e^{2\pi j\frac{(u-q)k}{N}}\bar{\rho}_{\mathrm{e}}^{(q+1)}\bar{\rho}_{\mathrm{e}}^{(u+1)}\right] = \sum_{q=0}^{N-1}\sum_{u=0}^{N-1} e^{2\pi j\frac{(u-q)k}{N}} E\left[\bar{\rho}_{\mathrm{e}}^{(q+1)}\bar{\rho}_{\mathrm{e}}^{(u+1)}\right].\end{aligned} \tag{8.43}$$

Since $\mathrm{E}\left[\bar{\rho}_{\mathrm{e}}^{(q+1)}\bar{\rho}_{\mathrm{e}}^{(u+1)}\right] = \mathrm{Cov}\left(\rho_{\mathrm{e}}^{(q+1)}, \rho_{\mathrm{e}}^{(u+1)}\right)$, the result of equation 8.32 is obtained.

9 Statistical iterative reconstruction for spectral differential phase-contrast computed tomography

The theoretical noise analysis that was conducted in the last chapter has demonstrated that SDPC radiography potentially enables a strong noise (or dose) reduction for a large range of clinically relevant parameters. Since image reconstruction via FBP is a linear operation, it could be expected that the same findings apply to 3D CT reconstruction. In this chapter, we extend our approach for projection-based SDPC imaging to CT reconstruction. Motivated by the benefits for spectral imaging (see chapter 6), we focus on the development of a one-step SIR algorithm for SDPC CT. A major obstacle for the translation of grating-based DPC imaging to clinical CT applications is the phase stepping procedure. In the case of a conventional FBP reconstruction, at least three different phase steps (i.e. positions on the stepping curve) would have to be recorded for each angle by shifting one of the gratings. This procedure is both time consuming and incompatible with a continuously rotating gantry. Numerical simulations demonstrate the possibility of eliminating the phase stepping procedure by using the one-step SIR approach for SDPC imaging (SDPC-SIR). Moreover, our investigations indicate that the SDPC-SIR algorithm enables CT acquisitions with strongly reduced dose levels compared with both conventional spectral imaging and SDPC-FBP reconstruction. However, more realistic numerical simulations as well as experimental measurements will be necessary to confirm these preliminary findings.
The text in this chapter is partly based on the following publication of the author:

- Mechlem, Korbinian, et al. "Statistical iterative reconstruction for spectral phase contrast CT." 15th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine. Vol. 11072. International Society for Optics and Photonics, 2019.

9.1 Methods

9.1.1 Forward model and CT reconstruction

The most straightforward approach for extending the projection-based SDPC imaging method that was presented in chapter 7 to CT reconstruction is depicted in figure 9.1. Each angular view is processed individually with the proposed projection-based decomposition algorithm and then the three volumes are reconstructed separately (via FBP or SIR):

$$\vec{m}_1 = \mathrm{RECO}\left(\vec{A}_1\right), \quad \vec{m}_2 = \mathrm{RECO}\left(\vec{A}_1\right), \quad \vec{\epsilon} = \mathrm{RECO}\left(\vec{d}_\epsilon^{\,i}\right), \tag{9.1}$$

where $\vec{m}_1$, $\vec{m}_2$ and $\vec{\epsilon}$ denote the volume fractions for the two attenuation basis materials and the dark-field basis material, respectively. The quantities $\vec{A}_1$, $\vec{A}_2$ and $\vec{d}_\epsilon$ represent the corresponding basis material sinograms and $\mathrm{RECO}(\cdot)$ is an image reconstruction operator. Alternatively, image reconstruction and material decomposition can be performed jointly in a one-step SIR framework. This is accomplished by a forward model that directly connects the (expected) spectral stepping curve measurements and the three basis material volumes. As discussed in

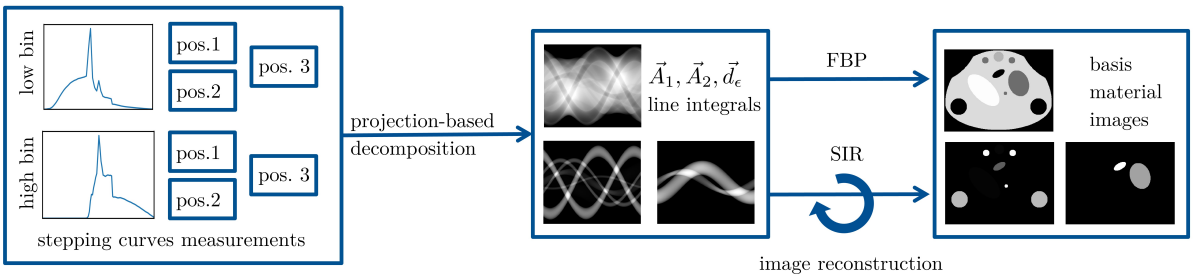

Figure 9.1: Two-step approach for SDPC CT reconstruction. First, the three basis material sinograms are calculated by processing each angular view independently with the projection-based decomposition algorithm that was discussed in chapter 7. In a second step, the three basis material volumes are reconstructed by applying standard image reconstruction techniques (FBP, SIR).

chapter 7, the expected number of photon counts $\hat{y}_i^{rs}$ for sinogram index i, stepping position r and energy bin s of the photon-counting detector can be modeled by:

$$\hat{y}_i^{rs} = \int_0^\infty S(E)R^s(E)e^{-A_1^i f_1(E) - A_2^i f_2(E)} \left[1 + V(E)e^{-d_\epsilon^i f_\epsilon(E)} \cos\left(\phi_i^r(E) + \Delta\phi_i(E)\right)\right] dE. \quad (9.2)$$

The key modification compared to the projection-based approach is calculating the basis material line integrals by a forward projection with system matrix $\boldsymbol{a}$:

$$A_1^i = \sum_j a_{ij} m_1^j, \quad A_2^i = \sum_j a_{ij} m_2^j, \quad d_\epsilon^i = \sum_j a_{ij} \epsilon^j, \quad (9.3)$$

where j is the voxel index. The phase shift $\Delta\phi(E)$ is proportional to the gradient of the projected electron density $\tilde{\rho}_e(x, y)$ perpendicular to the grating bars:

$$\Delta\phi_i(E) = \frac{q}{E^2} \frac{\partial}{\partial x} \tilde{\rho}_e^i(x_i, y_i), \quad (9.4)$$

where q is a proportionality factor depending on natural constants and the interferometer properties (compare eq. 7.10). The spatial coordinates for sinogram index i in the detector plane are given by x_i and y_i. Similar to the projection-based approach, we model the gradient of the projected electron density by a forward difference in the direction perpendicular to the grating bars:

$$\frac{\partial}{\partial x} \tilde{\rho}_e^i(x_i, y_i) = \frac{1}{t} \left(\tilde{\rho}_e^{(i+1)}(x_{(i+1)}, y_i) - \tilde{\rho}_e^i(x_i, y_i) \right), \quad (9.5)$$

where t is the effective pixel size (including the geometrical magnification of the object). The projected electron density $\tilde{\rho}_e^i = \tilde{\rho}_e^i(x_i, y_i)$ is obtained by forward projecting the three-dimensional electron density distribution $\vec{\rho}_e$ along the projection direction: $\tilde{\rho}_e^i = \sum_j a_{ij}\ \rho_e^j$. In contrast to chapter 7, we express the volume electron density (instead of the projected electron density) as:

$$\rho_e^j = m_1^j \rho_e(M_1) + m_2^j \rho_e(M_2), \quad (9.6)$$

where $\rho_e(M_1)$ and $\rho_e(M_2)$ are the volume electron densities of the basis materials. Figure 9.2 illustrates the iterative optimization algorithm that is used for solving the inverse problem of basis material image reconstruction for SDPC CT. Starting with an initial guess for the basis material volumes, the expected basis material sinograms are determined by forward projecting the corresponding volumes. Next, the expected spectral stepping curve measurements $\hat{y}_i^{rs}$ are calculated with the forward model of eq. 9.2. By comparing the expected stepping curves to the experimental measurements y_i^{rs} and incorporating prior knowledge in the form of a regularization term $R(\vec{m}_1, \vec{m}_2, \vec{\epsilon})$, an image update is calculated. Mathematically, this iterative process corresponds

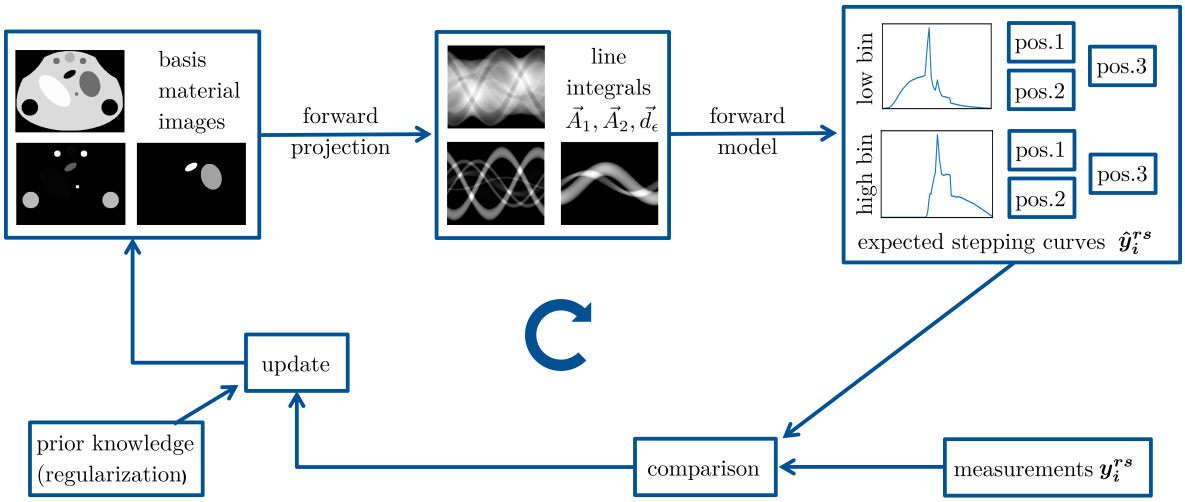

Figure 9.2: One-step statistical iterative reconstruction of spectral differential phase-contrast CT data. Material decomposition and image reconstruction are performed jointly by using a forward model that directly connects the basis material image to the expected spectral stepping curve measurements. By comparing the forward model with the measured data and including prior knowledge in the form of a regularization term, an image update is calculated. This iterative process is repeated until convergence.

to minimizing the penalized log-likelihood function $\phi(\vec{m}_1, \vec{m}_2, \vec{\epsilon}) = -L(\vec{m}_1, \vec{m}_2, \vec{\epsilon}) + R(\vec{m}_1, \vec{m}_2, \vec{\epsilon})$ (compare section 2.6.2). Assuming Poisson statistics for the measured photon counts y_i^{rs}, the negative log-likelihood function is given by:

$$-L(\vec{m}_1, \vec{m}_2, \vec{\epsilon}) = \sum_i \sum_s \sum_r \hat{y}_i^{rs} - y_i^{rs} \ln\left(\hat{y}_i^{rs}\right). \tag{9.7}$$

In this work, we use a Huber regularizer that is applied separately to each basis material image. The penalized log-likelihood function is minimized with a NLCG solver.
The forward model for one-step SDPC CT imaging neglect that the phase shift and the dark-field signal depend on the position along the optical axis. This approximation is reasonably accurate if the object size is small compared to the inter-grating distances.
In the case of standard DPC CT reconstruction, three or more stepping curve measurements per angular view are necessary to uniquely determine the attenuation, phase shift and dark-field images. Consequently, the gantry has to be halted at each projection angle to acquire the stepping curve measurements. This procedure is both time consuming and incompatible with a continuous CT acquisition (e.g. helical CT). Moreover, the mechanical movement of the grating makes the interferometer susceptible to drift artifacts. Recently, it has been demonstrated that a one-step SIR reconstruction algorithm for DPC CT can eliminate the phase stepping procedure. Teuffenbach et al. [418] accomplished this by varying the reference phase of the stepping curve as a function of the projection angle and the spatial position on the detector. This allows to elegantly interpolate the full stepping curve information from adjacent angular views and neighboring detector pixels. Figure 9.3 shows Radon space images that illustrate the acquisition scheme for one-step DPC CT. In figure 9.3 (a), a conventional phase stepping is displayed. For each angle and each detector pixel (in one detector row), a stepping curve with three different stepping positions (visualized by colored dots) is acquired. In figure 9.3 (b), only one stepping position is measured for each projection. However, the reference phase varies as a function of both the projection angle and the spatial position on the detector. In this work, we investigate whether the same technique can be used to eliminate the phase stepping procedure for SDPC CT. The novel acquisition scheme is incorporated into the reconstruction

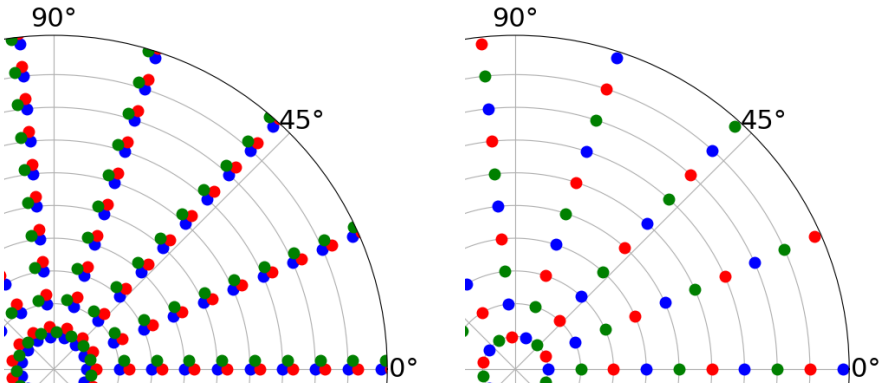

(a) (b)

Figure 9.3: Illustration of two acquisition schemes for (spectral) differential phase-contrast CT in Radon space. (a): Standard acquisition scheme: A full stepping curve with $R \geq 3$ phase steps (visualized by colored dots) is measured for each projection. (b): Only one stepping position is measured for each projection angle, but the reference phase varies with the projection angle and the spatial position on the detector.

framework by eliminating the sum over the phase steps in the log-likelihood function (eq. 9.7) and incorporating the varying reference phases $\phi_i^r(E)$ (as a function of the sinogram index i) into the forward model of eq. 9.2. In a real measurement, the reference phase could be varied as a function of the projection angle by focal spot deflection. A spatially varying reference phase could be achieved by slightly detuning the interferometer to generate fringes in the detector plane.

Besides comparing the two-step approaches for spectral and SDPC imaging, we compare the proposed one-step SIR algorithm for SDPC CT to the one-step SIR algorithm for spectral CT that was presented in chapter 6. Since, the source spectrum and the detector response are known in our simulation studies, the calibration step for the semi-empirical forward model that was discussed in chapter 6 can be omitted. This corresponds to setting $V(E) = 0$ in the forward model for SDPC imaging (eq. 9.2). Furthermore, the volume fractions $\vec{\epsilon}$ of the microstructured material are eliminated as optimization variables because the dark-field image cannot be determined with spectral CT.

9.1.2 Numerical simulation

Simulating a full CT scan with the wave-optical approach presented in chapter 3 would be a very time consuming task. To reduce the computational complexity and focus on the fundamental algorithmic properties of SDPC CT reconstructions, we use the forward model of eq. 9.2 and eq. 9.3 for our numerical simulations. While investigations with a perfect match between the simulation and the reconstruction ("inverse crime", see also chapter 3) are generally sufficient for analyzing the stability and other fundamental properties of a reconstruction algorithm, the performance is likely overestimated compared to real data. In particular, the position dependency of the phase shift and the dark-field signal are neglected. An analysis of the accuracy of the SDPC forward model (as was conducted in projection space, see chapter 7) is thus not possible with this type of simulation.

We simulated a CT scan of a scaled version of the Forbild microCT phantom [337] (see scalebar in the top left image of figure 9.4). The phantom was assumed to be made of PMMA and the inserts were filled with different materials (PVC, PMMA, water, aluminum). In two of the ellipsoids close to the center of the phantom, porous filling materials (PMMA and PVC) that generate a dark-field signal ($f_\epsilon(E) \propto 1/E^3$) were assumed. We simulated a tungsten X-ray

source with an acceleration voltage of 140 kVp and a photon-counting detector with a pixel size of 500 μm, a 2 mm thick CdTe sensor and two thresholds per pixel. The threshold positions (15 and 63 keV) were optimized for spectral imaging. The symmetric grating interferometer was comprised of two attenuation gratings (G0 and G2) made of gold and a $\pi/2$ phase-shifting grating (G1) with a design energy of 62 keV. We assumed a total interferometer length of 2.4 m and the phantom was placed 20 cm downstream of the G1 grating. The energy-dependent visibility $V(E)$ of the interferometer was simulated with the wave-optical approach described in section 3.2. Combining the simulated visibility with the forward model (eq. 9.2), stepping curves (using 5 equally distributed steps per angle) with the phantom in the beam path were simulated. These stepping curves were used for projection-based SDPC material decomposition followed by FBP image reconstruction (FBP-SDPC). In case of the one-step SIR algorithm for SDPC CT (OS-SIR-SDPC), a measurement without phase stepping was simulated by selecting only one stepping positions for each angle. This means that only one fifth of the dose was used compared to the two-step reconstruction approach. Different stepping positions were selected for consecutive angular views in order to improve the sampling of the stepping curves. Input data for spectral imaging was generated by averaging the full stepping curves to calculate conventional low and high energy sinograms. These sinograms were used to generate basis material images with conventional spectral material decomposition in projection space followed by FBP reconstruction (FBP-spectral) as well as spectral one-step SIR (OS-SIR-spectral).

9.2 Results

Figure 9.4 shows axial slices of the obtained PMMA and PVC basis material images for all four image reconstruction methods (FBP-spectral, FBP-SDPC, OS-SIR-spectral and OS-SIR-SDPC). As could be expected from SDPC radiography simulations (see chapters 7 and 8), the basis material images calculated with the FBP-SDPC algorithm show strongly reduced noise levels compared with the FBP-spectral reconstruction without compromising the quantitative accuracy. Compared to two-step spectral CT, the noise level (i.e variance of the basis material images in a homogeneous region close to the center of the phantom) is reduced by a factor of 5.9 and 4.2 for the PMMA and PVC image, respectively. Moreover, in agreement with the analysis that was conducted in chapter 8, the noise power spectrum for the FBP-SDPC reconstructions is shifted towards lower spatial frequencies. The regularization parameters for the OS-SIR-spectral and OS-SIR-SDPC reconstruction were tuned to achieve the same noise levels as the FBP-SDPC algorithm for all basis material images. The image quality and edge sharpness of both one-step reconstruction algorithms is very similar, despite the reduced dose level for the OS-SIR-SDPC reconstruction. Upon closer inspection, regularization crosstalk artifacts in the PVC images are observable for both one-step reconstruction methods.

In figure 9.5, the dark-field images calculated by the FBP-SDPC and OS-SIR-SDPC algorithm are shown. There are no regularization crosstalk artifacts observable for the dark-field image generated by the OS-SIR-SDPC algorithm because the dark-field image is only weakly coupled to the other two basis material images by the data term. A quantitatively correct dark-field image is reconstructed with the OS-SIR-SDPC algorithm despite the missing information (and reduced photon statistics) compared to a full phase-stepping scan. The image quality for both reconstruction methods is comparable, however some salt-and-pepper noise artifacts are observable for the OS-SIR-SDPC reconstruction.

9.3 Discussion and Outlook

As could be expected from the noise analysis in projection space, SDPC CT provides basis material images with strongly reduced noise levels compared to spectral CT. In the future, a more detailed noise analysis study for SDPC CT could be conducted by extending the noise

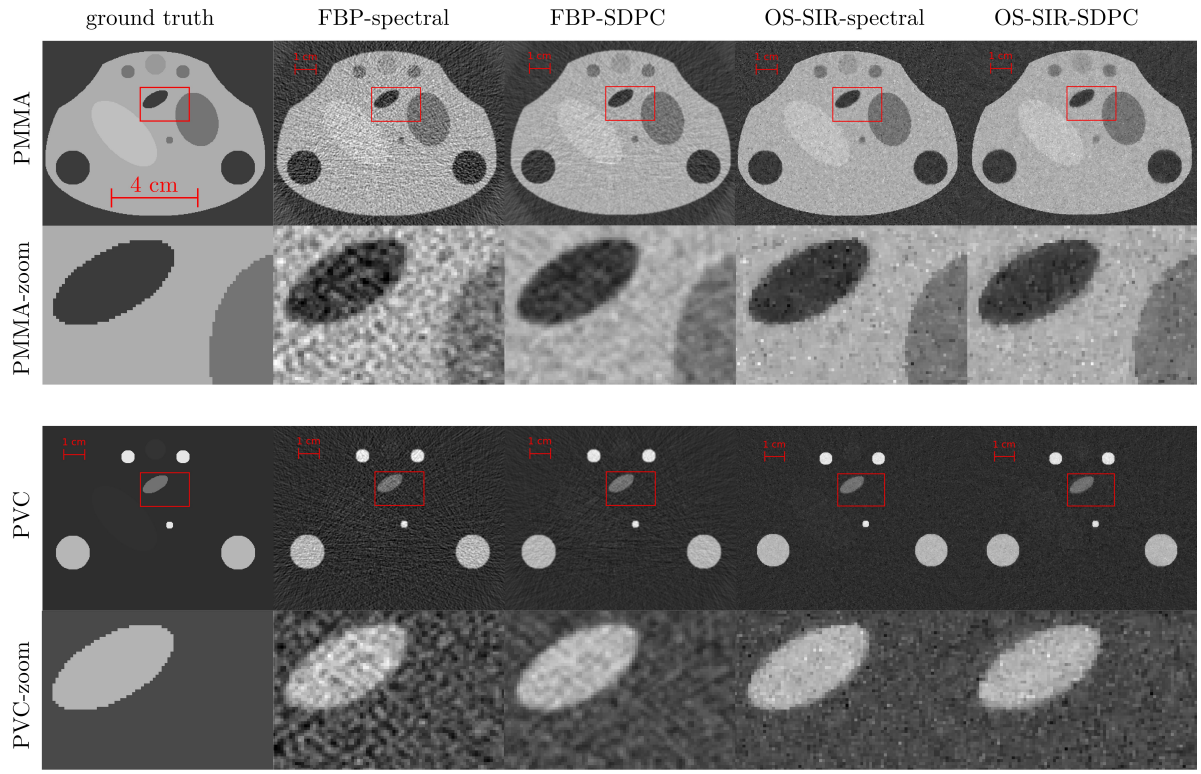

Figure 9.4: Comparison of the PMMA and PVC basis material images that were obtained with the four investigated image reconstruction methods. As could be expected from the results in the previous chapters, both the additional information for SDPC imaging as well as the one-step SIR techniques strongly improve the image quality compared with spectral imaging and two-step reconstruction algorithms, respectively.

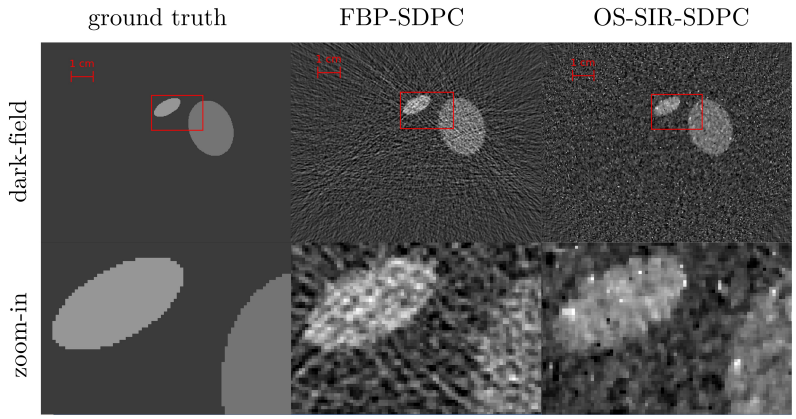

Figure 9.5: Comparison dark-field images that were reconstructed with the FBP-SDPC and the OS-SIR-SDPC algorithm. In both cases, the microstructured objects are correctly reconstructed. While this result is not surprising for the two-step method (since the full stepping curve information is available for the FBP-SDPC reconstruction), it demonstrates that the one-step SIR algorithm is able to infer information about the dark-field signal from incomplete data.

analysis framework (see chapter 8) to 3D CT reconstruction. In this simulation study, spectral CT was performed with the "unnecessary" grating interferometer in the beam path. Since the G2 grating attenuates approximately 43% of the photons that were transmitted through the object, the variance reduction factors achieved with the FBP-SDPC algorithm (compared to the FBP-spectral algorithm) of 5.9 and 4.2 correspond to potential dose reduction factors of 3.4 and 2.4 for the PMMA and PVC basis material images, respectively. In this simulation study, a comparatively large detector pixel size of 500 μm was chosen. A larger noise (or dose) advantage of SDPC imaging could be expected for CT applications that require higher spatial resolution.
Our preliminary results suggest that SDPC CT without phase stepping is feasible in combination with a one-step SIR algorithm. In agreement with the results for DPC CT by Teuffenbach et al. [418], the one-step SIR approach allows to reconstruct three quantitatively correct basis material images from incomplete stepping curve data. In our simulation study, the image quality of the one-step spectral CT reconstruction could be maintained for the one-step SDPC reconstruction despite using only one fifth of the available stepping curve measurements. Considering the attenuation of the G2 grating, this corresponds to a potential dose reduction of 65%. However, a more detailed analysis of noise-resolution tradeoff curves for both reconstruction approaches is necessary to draw definitive conclusions.
We believe that the image quality of the proposed one-step SIR technique for SDPC CT could be further improved by combining it with a more sophisticated regularization strategy. Advanced multi-channel regularizers (such as the dictionary-based regularizer that was presented in chapter 5) would most likely also mitigate the observed regularization crosstalk artifacts. Another important future project is developing a more efficient optimization algorithm for the inverse problem of one-step SDPC CT reconstruction. Unfortunately, deriving a fast solver based on surrogate functions (similar to one-step SIR for spectral CT, see chapter 6) has turned out to be difficult because of the the additional phase shift term in the forward model that couples neighboring detector pixels.
A limitation of this proof-of-concept simulation study is the perfect match between the simulation and the reconstruction model. More realistic simulations and experimental measurements are needed to investigate the influence of position-dependent phase shift and dark-field signals. This will be particularly relevant for possible clinical CT applications, where the size of the patient is comparable to the inter-grating distances. In the future, we plan to adjust the SDPC CT forward model to account for the position dependency. This could be achieved by modifying the forward projector for the electron density and the volume fractions of the microstructured material.
In summary, this work demonstrates that SDPC CT can offer two advantages compared to conventional spectral CT: Noise (or dose) reduction and additional information about the microstructure (dark-field image) that is inaccessible with attenuation-based imaging methods. With the OS-SIR-SDPC algorithm presented in this work, it is possible to use these advantages without relying on a time-consuming phase stepping scan. However, more detailed investigations as well as more realistic numerical simulations and experimental measurements are necessary to confirm these preliminary results.

10 Conclusions and outlook

Spectral and grating-based differential phase-contrast (DPC) X-ray imaging are two emerging technologies that offer additional information compared with conventional attenuation-based X-ray imaging. However, both imaging techniques still face experimental and algorithmic challenges that must be overcome on the way to routine use in clinical and non-destructive testing applications. In the first part of the results of this thesis, various algorithmic processing techniques for spectral X-ray imaging have been devised and investigated. To address the challenges of photon-counting detector (PCD) imperfections and noise amplification during material decomposition, we have developed calibration-based empirical forward models (chapters 5 and 6), multi-channel post-processing (chapter 4) and regularization (chapter 5) techniques as well as one-step statistical iterative reconstruction (SIR) algorithms (chapter 6). In the second part of this thesis, a new way of combining spectral and grating-based DPC imaging is discussed. Numerical simulations of radiography (chapter 7) and CT applications (chapter 9) show that the combination of these two imaging techniques benefits from the strengths of the individual methods while the weaknesses are mitigated. Compared with DPC imaging, the additional spectral information allows to mitigate beam hardening and phase wrapping artifacts that can compromise the quantitative accuracy. Compared with spectral imaging, a theoretical analysis (chapter 8) shows that the noise (or dose) level can be strongly reduced by exploiting the additional phase shift information.

These contributions highlight the potential and the possible benefits of spectral differential phase-contrast (SDPC) X-ray imaging techniques for future clinical and non-destructive testing applications. However, thorough experimental studies will be necessary to confirm the findings that were obtained with numerical simulations. Many of the processing techniques for spectral imaging that were discussed in chapters 4 - 6 can be transferred to SDPC imaging to further improve the image quality or reduce the dose level.

Figure 10.1 illustrates the connection between the individual chapters in this work and their contribution towards potential future low-dose SDPC applications in the fields of medical imaging and non-destructive testing. Dashed lines (and red text) indicate possible future work towards the experimental realization of low-dose SDPC imaging.

In chapter 4, we developed a dictionary-based multi-channel denoising algorithm for spectral CT. Performing dictionary denoising of the minimum noise image in combination with local linear transformations to the basis material images allowed us to exploit the structural as well as the noise correlations between different basis material images. The novel denoising algorithm outperformed two other state-of-the-art post-processing techniques. An advancement of this algorithm for radiography applications was presented in chapter 5. By reformulating the dictionary denoising process as prior knowledge in a Bayesian framework, the denoising technique could be combined with a maximum-likelihood estimator for material decomposition that considers the measurement statistics in projection space. The calculation of the minimum noise image was replaced by local whitening transformations, which eliminates the potential risk of missing structures in the basis material images. In addition, the local whitening transformations offer an elegant way of considering a spatially varying noise level in all image channels. To obtain high quality basis material images, inter-pixel differences in the detector response of current PCDs have to be taken into account by the material decomposition algorithm. Towards this end, we developed an empirical polynomial forward model that is tuned by calibration measure-

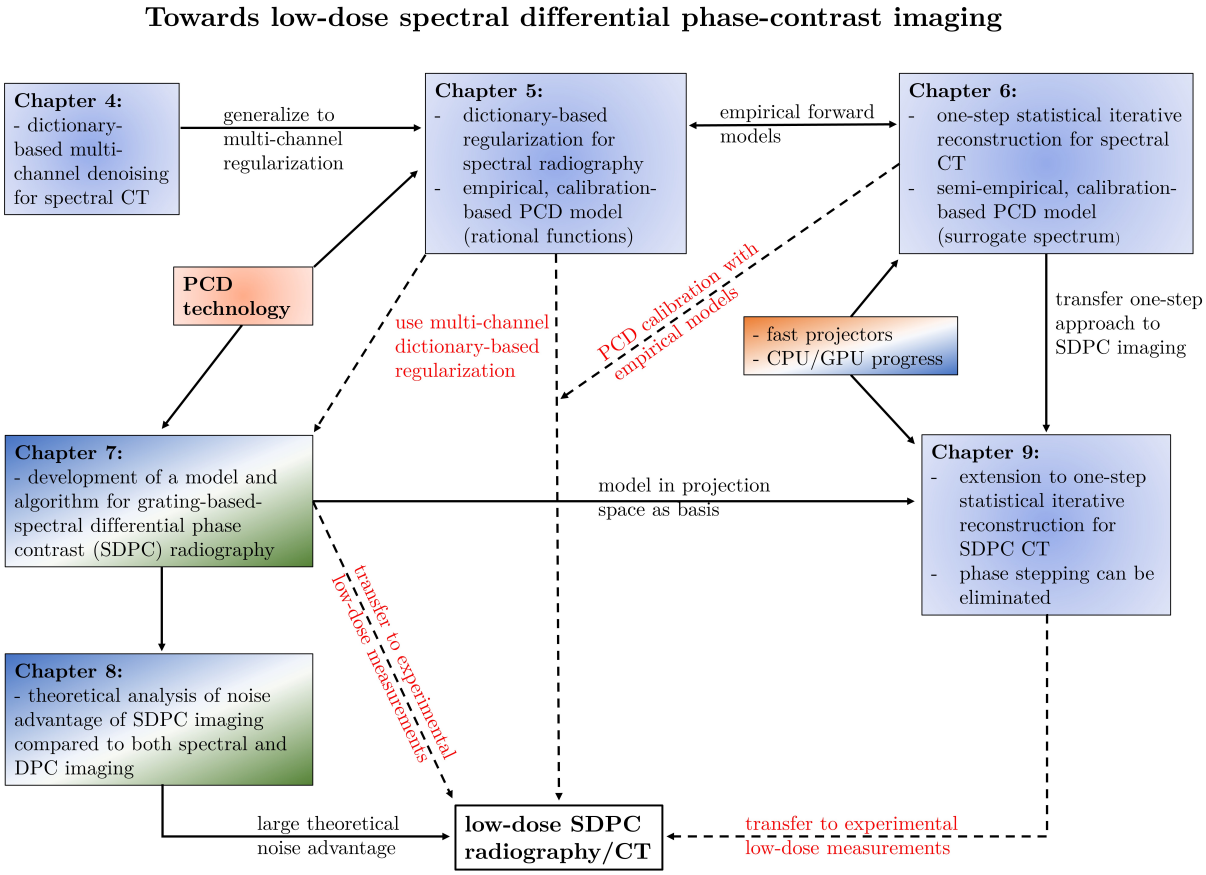

Figure 10.1: Graphical summary of the connection between the individual chapters in this work and their contribution towards potential future low-dose spectral differential phase-contrast applications in the fields of medical imaging and non-destructive testing. Dashed lines (and red text) indicate possible future work towards the experimental realization of low-dose spectral differential phase-contrast imaging. A blue background indicates algorithmic developments, whereas orange and green background colors represent hardware advances and novel physical models, respectively.

ments. This strategy allows to model spatially varying properties of the imaging system without requiring detailed prior knowledge of the system parameters. We focused on spectral coronary angiography as a potential clinical application of projection-based material decomposition with photon counting detectors. Numerical and real experiments show that spectral angiography with realistic dose levels and gadolinium contrast agent concentrations is feasible using the proposed decomposition algorithm and currently available PCD technology. A different empirical estimator was investigated in chapter 6. The corresponding forward model based on surrogate spectra is again tuned by calibration measurements. Both empirical estimators are practically unbiased and achieve the theoretical lower limit for the variance of an unbiased estimator (CRLB), even if the number of spectral measurements is larger than the number of basis materials. In chaper 6 we extended the projection-based estimator to 3D CT reconstruction by developing a one-step SIR algorithm which performs image reconstruction and basis material decomposition jointly in one step. The empirical forward model was adapted to CT reconstruction and combined with a high efficient optimization algorithm that converges one to three orders of magnitude faster than other existing approaches [378]. This partially negates the major shortcoming of SIR techniques, namely high computational cost and long reconstruction times, and thus represents a key step towards using advanced one-step SIR techniques in practice. Besides the improved im-

age quality compared with two-step approaches, we demonstrated how one-step SIR algorithms can mitigate statistical bias, which is particularly relevant for low-dose quantitative spectral CT.

In chapter 7, a new algorithm for combining grating-based DPC radiography and spectral radiography was investigated. Two basis material images and a dark-field image are extracted by simultaneously using the spectral and the phase contrast information. In contrast to conventional DPC imaging, quantitatively correct values for the pathlengths through a microstructured material can be retrieved. The connection between spectral and DPC imaging was established via the projected electron density that can be determined with both imaging techniques. We focused on human thorax radiography as a promising medical application of SDPC radiography and demonstrated potential benefits for clinical diagnosis. In addition to the reduced noise level compared with spectral radiography, the dark-field image could provide valuable additional diagnostic information about the lungs. We concluded that phase wrapping artifacts can in principle be eliminated with SDPC imaging, provided that the global optimum of the log-likelihood function is found. Although convergence to local optima could be avoided by a suitable initial guess in our studies, it remains an open question how convergence to the global optimum can be guaranteed with an efficient gradient-based optimization algorithm. Nevertheless, the existence of a unique global optimum is an important advantage compared to DPC imaging. An interesting future project would be combining the SDPC material decomposition algorithm with a multi-channel regularization technique (e.g. the dictionary-based regularizer discussed in chapter 5) to further improve the image quality and reduce the probability of convergence to local optima.
The noise analysis framework that was developed in chapter 8 allows the calculation of (co-) variances and noise power spectra for spectral, DPC and SDPC imaging in projection space. An important practical application of this framework is finding the optimum acquisition parameters for a particular imaging task. In the future, the predicted covariances could be used to optimize the performance of regularized material decomposition algorithms (e.g. to achieve a uniform resolution or noise level). Our noise analysis indicated that SDPC imaging outperforms both spectral and DPC imaging for a large range of imaging parameters that are relevant for medical imaging applications. Compared with DPC imaging, the additional spectral information mitigates long-range noise correlations and excessive low frequency noise, which can be a major drawback of DPC imaging, especially in projection space. Nevertheless, depending on the interferometer sensitivity, short to medium-range noise correlations remain. Whether these correlations can impede the recognition of small image features (compared with spectral imaging) needs to be evaluated in future reader studies.
In chapter 9, the projection-based SDPC algorithm was extended to 3D CT reconstruction. In particular, we transferred the one-step SIR techniques that were originally developed for spectral CT (see chapter 6) and DPC CT [418] to SDPC CT. The ability of one-step SIR algorithms to handle non-standard acquisition schemes and incomplete information allowed us to eliminate the time consuming phase stepping procedure. Numerical simulations indicated that the image quality of one-step SDPC CT reconstructions is superior to both two-step SDPC CT and one-step spectral CT reconstructions. However, more detailed investigations as well as more realistic numerical simulations and experimental measurements are necessary to confirm these preliminary results. Furthermore, a more thorough analysis of the noise characteristics for SDPC CT could be conducted by extending the noise analysis framework to 3D CT reconstruction. In the future, we plan to adjust the SDPC CT forward model to account for the position dependencies of the dark-field and phase shift signals. This will be particularly relevant for possible clinical CT applications, where the size of the patient is comparable to the inter-grating distances.

We believe that the contributions of this work could provide the foundation for further research towards low-dose SDPC radiography and CT applications in the field of medical imaging.

Due to the steady progress of computing hardware, the use of the processing and reconstruction algorithms that have been discussed in this thesis will probably be feasible for most application cases. Given the recent advances in PCD technology, it is likely that some of the current performance degrading effects of PCDs will be mitigated in the future. This would lead to further improvements in the image quality for spectral and SDPC X-ray imaging applications. The theoretical noise analysis in chapter 8 shows that the largest noise advantages compared to existing techniques could be expected for imaging tasks with comparably high spatial resolution and moderate attenuation (e.g. mammography). In cases where the expected noise advantage compared to spectral imaging is smaller (e.g. most human CT applications), the dark-field image could still provide valuable additional diagnostic information.

Although a first proof-of-concept study has demonstrated the experimental feasibility of SDPC imaging [419], there are still some challenges left that need to be overcome on the way to potential medical imaging applications. To obtain quantitatively accurate and artifact-free basis material images in experimental measurements, both the effective spectrum and the energy-dependent visibility of the interferometer have to be known precisely. We believe that a combination of the physical SDPC forward model with a calibration-based approach (such as the semi-empirical forward model discussed in chapter 6) could be a viable solution for considering grating inhomogeneities and spatial variations of the detector response. Another important future project is combining the SDPC decomposition algorithm with advanced multi-channel regularizers that consider the unique noise correlations for SDPC imaging (both between different channels as well as different detector pixels). The improved image quality would most likely allow to further decrease the radiation dose for many applications. Moreover, novel multi-channel regularizers (such as the dictionary-based regularizer that was presented in chapter 5) would most likely also mitigate the observed regularization crosstalk artifacts. Grating drifts (e.g. due vibrations or thermal expansion) have been neglected in the numerical simulations. Depending on the setup stability and the acquisition time, the SDPC model may have to be combined with a reference phase estimation algorithm. Because of the longer measurement times, this is mainly relevant for experimental CT measurements, especially for acquisition schemes without phase stepping.

11 Appendix

11.1 Wave propagation

11.1.1 Extended sources and the projection approximation

Under the projection approximation and the paraxial approximation (Fresnel propagation), the effect of an extended, incoherent source can be considered by a convolution of the point source intensity pattern with a scaled version of the spatial intensity distribution of the source. This result will be derived by comparing the intensity pattern of two shifted point sources under the projection approximation. For simplicity, we reduce the problem to two dimensions and only consider shifts in the x-direction. Figure 11.1 shows two point sources P and Q in the $z = 0$ plane located at $x = 0$ and $x = x_0$, respectively. A weakly scattering object is located at $z = l$ and the detector is placed at $z = l + d$. According to the Huygens-Fresnel principle (eq. 2.17), the incident wavefronts in the $z = l$ plane are given by:

$$\begin{aligned} \phi_\omega^{(Q)}(x, z = l) &= \sqrt{I_0} \exp\left(\frac{ik(x - x_0)^2}{2l}\right) \\ \phi_\omega^{(P)}(x, z = l) &= \sqrt{I_0} \exp\left(\frac{ikx^2}{2l}\right) , \end{aligned} \tag{11.1}$$

where I_0 is the intensity of the wave field in the $z = l$ plane. Since we are eventually interested in the intensity in the detector plane, pure phase factors can be ignored in this derivation. Applying the projection approximation and using the convolution formulation of Fresnel propagation (the one-dimensional version of eq. 2.14), the wave-functions in the detector plane are given by:

$$\begin{aligned} \phi_\omega^{(Q)}(x, z = l + d) &= \sqrt{I_0}\sqrt{\frac{k}{2\pi d}} \int \exp\left(\frac{ik(x' - x_0)^2}{2l}\right) f(x') \exp\left(\frac{ik(x - x')^2}{2d}\right) dx' \\ \phi_\omega^{(P)}(x, z = l + d) &= \sqrt{I_0}\sqrt{\frac{k}{2\pi d}} \int \exp\left(\frac{ikx'^2}{2l}\right) f(x') \exp\left(\frac{ik(x - x')^2}{2d}\right) dx' , \end{aligned} \tag{11.2}$$

where $f(x)$ represents the influence of the object in the projection approximation. By expanding the quadratic terms in the integral and eliminating phase factors, the intensity of the wave fields

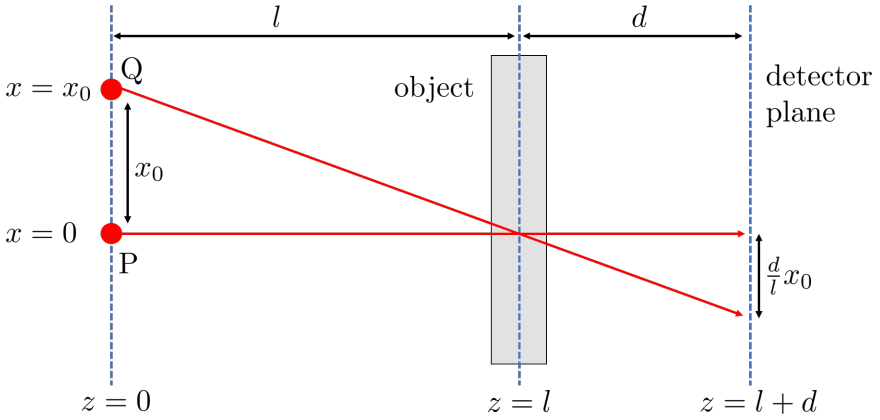

Figure 11.1: Illustration of the shift of the intensity pattern in the detector plane when moving a point source in the source plane ($z = 0$) from $x = 0$ to $x = x_0$.

in the detector plane can be calculated as:

$$\begin{aligned} I^{(Q)}(x, z = l + d) &= I_0 \frac{k}{2\pi d} \left| \int \exp\left(\frac{ik}{2} \left[x'^2 \left(\frac{1}{l} + \frac{1}{d} \right) - \frac{2x'}{d} \left(\frac{x_0 d}{l} + x \right) \right] \right) f(x') dx' \right|^2 \\ I^{(P)}(x, z = l + d) &= I_0 \frac{k}{2\pi d} \left| \int \exp\left(\frac{ik}{2} \left[x'^2 \left(\frac{1}{l} + \frac{1}{d} \right) - \frac{2x'}{d} x \right] \right) f(x') dx' \right|^2 \end{aligned} \tag{11.3}$$

From the above pair of equations, it can be seen that:

$$I^{(Q)}\left(x - \frac{x_0 d}{l}, z = l + d \right) = I^{(P)}\left(x, z = l + d \right), \tag{11.4}$$

i.e. shifting the point source by x_0 in the source plane is equivalent to shifting the intensity pattern by $-x_0 d/l$ in the detector plane. Since the individual intensities for incoherent sources simply add up (see section 2.2.3), the effect of an extended source can be considered by (spatially) scaling the source intensity profile with the factor $-d/l$ and then convolving the point source intensity profile in the detector plane with the scaled intensity distribution of the extended source. The two rays in figure 11.1 indicate how the same result can be obtained in the simpler (and less general) geometrical optics picture.

11.1.2 Talbot effect in combination with a phase-shifting object

In this section, we derive the displacement of the Talbot self-image that is caused by a purely phase-shifting object in the wave picture using Fresnel propagation and the projection approximation. The same result is more easily obtained by combining wave propagation (for the Talbot effect, see section 2.8) and geometrical optics (to explain the shift of the self-image with deflected rays, compare sections 2.3.4 and 2.8). However, the following derivation is more rigorous in the sense that the hybrid wave/ray picture is not necessary. For simplicity, we reduce the problem to two dimensions and assume a purely phase shifting prism, identical to the one presented in section 2.3.4. Moreover, we assume an attenuating cosine grating at $z = 0$, that transmits the following wave field:

$$\phi_g(x, z = 0) = \left(1 + \cos\left(\frac{2\pi}{p} x \right) \right) = \left(1 + \frac{1}{2} \left(e^{i\frac{2\pi}{p}x} + e^{-i\frac{2\pi}{p}x} \right) \right) \tag{11.5}$$

The derivation could be generalized to arbitrary periodic gratings by Fourier decomposition (see for example reference [43], chapter 4.2.1 for a more general derivation of the Talbot effect). According to the projection approximation, the influence of the prism that is placed directly behind the grating is modeled by (compare eq. 2.35):

$$\phi(x, z = 0) = e^{iqx} \phi_{\mathrm{g}}(x, z = 0) = \left(1 + \frac{1}{2} \left(e^{i\frac{2\pi}{p}x} + e^{-i\frac{2\pi}{p}x} \right) \right) e^{iqx}, \quad q = -k\delta_\omega m, \tag{11.6}$$

where m is the slope of the prism. The propagated wave field at $z = z_0$ can be calculated by Fresnel propagation:

$$\phi_\omega(x, z_0) = e^{ikz_0} \mathcal{F}^{-1} e^{\frac{-iz_0}{2k} k_x^2} \mathcal{F}\, \phi_\omega(x, z = 0) \tag{11.7}$$

As a first step, the Fourier transform of $\phi_\omega(x, z = 0)$ is calculated:

$$\mathcal{F}\left(\phi_\omega(x, z = 0) \right) = \sqrt{2\pi} \left(\delta\left(q - k_x \right) + \frac{1}{2} \delta\left(\frac{2\pi}{p} + q - k_x \right) + \frac{1}{2} \delta\left(-\frac{2\pi}{p} + q - k_x \right) \right) \tag{11.8}$$

Combining eq. 11.7 and eq. 11.8 yields:

$$\begin{aligned}
\phi_\omega(x,z_0) &= e^{ikz_0}\int e^{\frac{-iz_0}{2k}k_x^2}\left(\delta\left(q-k_x\right)+\frac{1}{2}\delta\left(\frac{2\pi}{p}+q-k_x\right)+\frac{1}{2}\delta\left(-\frac{2\pi}{p}+q-k_x\right)\right)e^{ik_xx}dk_x \\
&= e^{ikz_0}\left[e^{iqx}e^{-i\frac{z_0}{2k}q^2}+\frac{1}{2}e^{-i\frac{z_0}{2k}\left(\frac{2\pi}{p}+q\right)^2}e^{i\left(\frac{2\pi}{p}+q\right)x}+\frac{1}{2}e^{-i\frac{z_0}{2k}\left(-\frac{2\pi}{p}+q\right)^2}e^{i\left(-\frac{2\pi}{p}+q\right)x}\right] \\
&= e^{ikz_0}e^{iqx}e^{-i\frac{z_0}{2k}q^2}\left[1+\frac{1}{2}e^{-i\frac{z_0}{2k}\left(\frac{2\pi}{p}\right)^2}\left(e^{i\left(\frac{2\pi}{p}x-\frac{z_0}{2k}\frac{4\pi q}{p}\right)}+e^{-i\left(\frac{2\pi}{p}x-\frac{z_0}{2k}\frac{4\pi q}{p}\right)}\right)\right] \\
&= e^{ikz_0}e^{iqx}e^{-i\frac{z_0}{2k}q^2}\left[1+e^{-i\frac{z_0}{2k}\left(\frac{2\pi}{p}\right)^2}\cos\left(\frac{2\pi}{p}\left[x-\frac{z_0q}{k}\right]\right)\right]
\end{aligned} \tag{11.9}$$

A detector placed in the $z = z_0$ plane will measure the intensity pattern $I(x, z_0) = \phi_\omega(x, z_0)\phi^*_\omega(x, z_0)$, which is given by [44]:

$$I(x,z_0) = 1+2\cos\left(\frac{z_0}{2k}\left(\frac{2\pi}{p}\right)^2\right)\cos\left(\frac{2\pi}{p}\left[x-\frac{z_0q}{k}\right]\right)+\cos^2\left(\frac{2\pi}{p}\left[x-\frac{z_0q}{k}\right]\right). \tag{11.10}$$

To produce a Talbot self image in the $z = z_0$ plane, the propagation distance must be equal to a Talbot distance:

$$z_0 = m\frac{2p^2}{\lambda} = m\frac{kp^2}{\pi} \implies \frac{z_0}{2k}\left(\frac{2\pi}{p}\right)^2 = 2\pi m, \quad m\in\mathbb{N} \tag{11.11}$$

At a Talbot distance, the expression for the intensity (eq. 11.10) reduces to:

$$I(x,z_0) = \left[1+\cos\left(\frac{2\pi}{p}\left[x-\frac{z_0q}{k}\right]\right)\right]^2, \tag{11.12}$$

which can be interpreted as a perfect self-image of the intensity pattern directly behind the grating (compare eq. 11.5) that is shifted by:

$$\Delta x = \frac{z_0q}{k} = -\delta_\omega m z_0 = \alpha z_0. \tag{11.13}$$

This result coincides with the geometrical optics picture in which the shift of the self-image is explained by deflected X-rays (compare section 2.3.4).

11.1.3 Fresnel scaling theorem and the multi-slice projection approximation

In this section, we derive how the Fresnel scaling theorem can be combined with the multi-slice projection approximation. This combination is particularly relevant for fast and accurate simulations of grating-based DPC imaging (see chapter 3). However, it is not immediately obvious how to combine these two simulation techniques. The multi-slice projection approximation requires the wave function $\phi(x, y, z = z_0)$ in different planes perpendicular to the optical axis, whereas the Fresnel scaling theorem relates the intensity in a plane $z = z_0$ that is obtained with point source illumination to the one that would have been obtained with plane wave illumination. We therefore re-examine the derivation of the Fresnel scaling theorem that is given in reference [43] to establish a similar relation for the wave functions. In the following, we assume a symmetric DPC setup with a thin sample placed between the G1 and the G2 grating. The corresponding geometrical parameters are defined in figure 11.2. For simplicity, the G0 grating is replaced by a perfect point source. However, as discussed in sections 2.2.3 and 11.1.1, the effect of the G0 grating can be included by a convolution with the intensity profile in the G0 plane. Our goal is to calculate the intensity distribution in the G2 plane ($z = l + r_1 + r_2$). Towards this

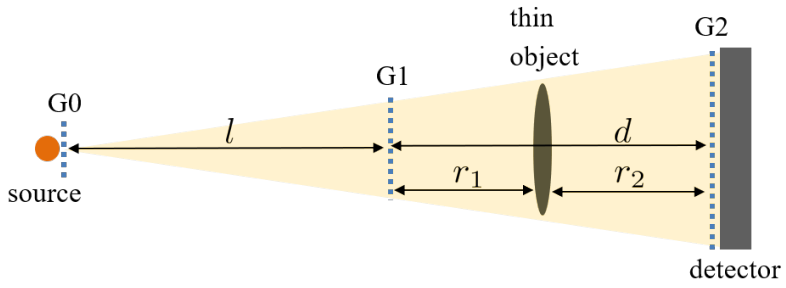

Figure 11.2: Three-grating interferometer for DPC imaging with the sample placed between the G1 and G2 grating. The definition of the geometry parameters in this figure will be used for the derivations in this section.

end, the projection approximation needs to be applied twice: At the location of the G1 grating ($z = l$) and at the location of the sample ($z = l + r_1$). Given the wave field $\phi_\omega(x, y, z = l)$, the wave field in the plane $z = l + r_1$ can be calculated with the convolution formulation of Fresnel propagation:

$$\phi_\omega(x, y, z = l + r_1) = \frac{-ike^{ikr_1}}{2\pi r_1} \int\int \phi_\omega(x', y', z = l) \exp\left(\frac{ik}{2r_1}(x - x')^2 + (y - y')^2\right) dx'dy' \tag{11.14}$$

From now on, we ignore multiplicative phase factors that are independent of x and y since they are irrelevant for the intensity in the G2 plane. Expanding the quadratic terms in eq. 11.14 yields:

$$\begin{aligned}\phi_\omega(x, y, z = l + r_1) = \frac{k}{2\pi r_1} \exp\left(\frac{ik}{2r_1}(x^2 + y^2)\right) \int\int \phi_\omega(x', y', z = l) \times \\ \exp\left(\frac{ik}{2r_1}(x'^2 + y'^2)\right) \exp\left(\frac{-ik}{r_1}(xx' + yy')\right) dx'dy'.\end{aligned} \tag{11.15}$$

Under the paraxial and the projection approximation, the wave field for point-source illumination $\phi_\omega(x, y, z = l)$ is related to the wave field for plane wave illumination $\phi_\omega^{(\mathrm{PW})}(x, y, z = l)$ via (compare eq. 2.17):

$$\phi_\omega(x, y, z = l) = \phi_\omega^{(\mathrm{PW})}(x, y, z = l) \exp\left(\frac{ik}{2l}(x^2 + y^2)\right) \tag{11.16}$$

Combining eq. 11.15 and eq. 11.16, we obtain an expression for $\phi_\omega(x, y, z = l + r_1)$ in terms of the exit-surface wave field at $z = l$ for the case of plane wave illumination:

$$\begin{aligned}\phi_\omega(x, y, z = l + r_1) = \frac{k}{2\pi r_1} \exp\left(\frac{ik}{2r_1}(x^2 + y^2)\right) \int\int \phi_\omega^{(\mathrm{PW})}(x', y', z = l) \times \\ \exp\left(\frac{ik}{2}(x'^2 + y'^2)\left(\frac{1}{r_1} + \frac{1}{l}\right)\right) \exp\left(\frac{-ik}{r_1}(xx' + yy')\right) dx'dy'.\end{aligned} \tag{11.17}$$

Since:

$$\frac{1}{r_1} + \frac{1}{l} = \frac{1}{r_1}\left(\frac{l + r_1}{l}\right) = \frac{M_1}{r_1}, \tag{11.18}$$

where M_1 is the magnification of the G1 grating, eq. 11.17 can be rewritten as:

$$\begin{aligned}\phi_\omega(x, y, z = l + r_1) = \frac{k}{2\pi r_1} \exp\left(\frac{ik}{2r_1}(x^2 + y^2)\right) \int\int \phi_\omega^{(\mathrm{PW})}(x', y', z = l) \times \\ \exp\left(\frac{ikM_1}{2r_1}(x'^2 + y'^2)\right) \exp\left(\frac{-ik}{r_1}(xx' + yy')\right) dx'dy'.\end{aligned} \tag{11.19}$$

The wave field at $z = l + r_1$ in parallel beam geometry can be obtained from eq. 11.19 by taking the limit $l \to \infty$ (i.e. $M_1 \to 1$):

$$\begin{aligned}\phi_\omega^{(\mathrm{PW})}(x, y, z = l + r_1) = \frac{k}{2\pi r_1} \exp\left(\frac{ik}{2r_1}(x^2 + y^2)\right) \int\int \phi_\omega^{(\mathrm{PW})}(x', y', z = l) \times \\ \exp\left(\frac{ik}{2r_1}(x'^2 + y'^2)\right) \exp\left(\frac{-ik}{r_1}(xx' + yy')\right) dx'dy'.\end{aligned} \tag{11.20}$$

As can be seen by direct substitution, the wave fields at $z = l + r_1$ for parallel beam and cone beam geometry are related by:

$$\phi_\omega(x, y, z = l + r_1) = M_1^{-1} \phi_\omega^{(\mathrm{PW})}\left(\frac{x}{M_1}, \frac{y}{M_1}, z = l + \frac{r_1}{M_1}\right) \exp\left(\frac{ik(M_1 - 1)}{2r_1 M_1}(x^2 + y^2)\right). \tag{11.21}$$

This relationship is very similar to the Fresnel scaling theorem (compare eq. 2.30), except for the additional phase factor.
Since:

$$\frac{(M_1 - 1)}{r_1 M_1} = \frac{l + r_1 - l}{r_1(l + r_1)} = \frac{1}{l + r_1}, \tag{11.22}$$

the phase factor can be rewritten as:

$$\exp\left(\frac{ik(M_1 - 1)}{2r_1 M_1}(x^2 + y^2)\right) = \exp\left(\frac{ik}{2(l + r_1)}(x^2 + y^2)\right). \tag{11.23}$$

The phase factor in eq. 11.21 thus corresponds to the wave field that a point source located at $(x, y, z) = (0, 0, 0)$ produces in the plane $z = l + r_1$. We define an intermediate wave field $\phi_\omega^{(\mathrm{IM})}(x, y, z = l + r_1)$ such that:

$$\phi_\omega(x, y, z = l + r_1) = \phi_\omega^{(\mathrm{IM})}(x, y, z = l + r_1) \exp\left(\frac{ik}{2(l + r_1)}(x^2 + y^2)\right) \tag{11.24}$$

By combining eq. 11.21 with eq. 11.24, we obtain the relation:

$$\phi_\omega^{(\mathrm{IM})}(x, y, z = l + r_1) = M_1^{-1} \phi_\omega^{(\mathrm{PW})}\left(\frac{x}{M_1}, \frac{y}{M_1}, z = l + \frac{r_1}{M_1}\right). \tag{11.25}$$

Repeating the derivation of eqs. 11.17 - 11.21 (except that l and r_1 are replaced by $l + r_1$ and r_2, respectively) leads to:

$$\begin{aligned}\phi_\omega(x, y, z = l + r_1 + r_2) = & M_2^{-1} \phi_\omega^{(\mathrm{IM})}\left(\frac{x}{M_2}, \frac{y}{M_2}, z = l + r_1 + \frac{r_2}{M_2}\right) \\ & \exp\left(\frac{ik(M_2 - 1)}{2r_2 M_2}(x^2 + y^2)\right),\end{aligned} \tag{11.26}$$

where $M_2 = (l + r_1 + r_2)/(l + r_1)$ is the magnification of the sample in the detector plane. The intensity in the detector plane is given by:

$$I_\omega(x, y, z = l + r_1 + r_2) = M_2^{-2} \left| \phi_\omega^{(\mathrm{IM})}\left(\frac{x}{M_2}, \frac{y}{M_2}, z = l + r_1 + \frac{r_2}{M_2}\right) \right|^2 \tag{11.27}$$

Considering eqs. 11.25 - 11.27, the following algorithm can be used for simulating the intensity in the detector plane in a DPC setup:

- 1.) Initialize a plane wave at $z = l$.
- 2.) Apply the projection approximation in the G1 plane to obtain the exit surface wave immediately downstream of the G1 grating.

- 3.) Apply the Fresnel scaling theorem to the wave function (instead of the intensity) to obtain an intermediate wave field at $z = l + r_1$:

$$\phi_\omega^{((\mathrm{IM})}(x, y, z = l + r_1) = M_1^{-1} \phi_\omega^{(\mathrm{PW})} \left(\frac{x}{M_1}, \frac{y}{M_1}, z = l + \frac{r_1}{M_1} \right). \tag{11.28}$$

- 4.) Apply the projection approximation to calculate the influence of the sample on the wave field.
- If necessary, repeat steps 3.) and 4.) to simulate an extended sample.
- 5.) Apply the Fresnel scaling theorem for the intensities (eq. 11.27) to calculate the intensity in the detector plane.

Abbreviations

3D ..... three-dimensional
ASIC ..... application-specific integrated circuit
CdTe ..... cadmium telluride
CG ..... conjugate gradient
CNR ..... contrast-to-noise ratio
CRLB ..... Cramér-Rao lower bound
CT ..... computed tomography
DPC ..... differential phase-contrast
DSA ..... digital subtraction angiography
DSCS ..... differential scattering cross section
FBP ..... filtered backprojection
GPU ..... graphics processing unit
ICA ..... invasive coronary angiography
MAP ..... maximum a posterior
ML ..... maximum likelihood
MSE ..... mean-squared error
NLCG ..... nonlinear conjugate gradient
NPS ..... noise power spectrum
OMP ..... orthogonal matching pursuit
OS ..... ordered subsets
OSSQS ..... ordered subsets separable quadratic surrogate
PCD ..... photon-counting detector
PED ..... projected electron density
PL ..... penalized likelihood
RMSE ..... root-mean squared error
RMSRE ..... root-mean squared relative error
SDPC ..... spectral differential phase-contrast
SIR ..... statistical iterative reconstruction
SNR ..... signal-to-noise ratio
SOMP ..... simultaneous orthogonal matching pursuit
SQS ..... separable quadratic surrogate
SSI ..... structural similarity index
TV ..... total variation
VMI ..... virtual monochromatic image

List of publications and scientific presentations

First-authored publications (peer-reviewed)

Mechlem, K., Allner, S., Ehn, S., Mei, K., Braig, E., Münzel, D., Pfeiffer, F., Noël, P. B. (2017). A post-processing algorithm for spectral CT material selective images using learned dictionaries. *Biomedical Physics & Engineering Express*, 3(2), 025009.

Mechlem, K., Ehn, S., Sellerer, T., Braig, E., Münzel, D., Pfeiffer, F., Noël, P. B. (2017). Joint statistical iterative material image reconstruction for spectral computed tomography using a semi-empirical forward model. *IEEE Transactions on Medical Imaging*, 37(1), 68-80.

Mechlem, K., Sellerer, T., Ehn, S., Münzel, D., Braig, E., Herzen, J., Noël, P. B., Pfeiffer, F. (2018). Spectral angiography material decomposition using an empirical forward model and a dictionary-based regularization. *IEEE Transactions on Medical Imaging*, 37(10), 2298-2309.

Mechlem, K., Sellerer, T., Viermetz, M., Herzen, J., Pfeiffer, F. (2020). Spectral differential phase contrast X-ray radiography. *IEEE Transactions on Medical Imaging*, 39(3), 578-587.

Mechlem, K., Sellerer, T., Viermetz, M., Herzen, J., Pfeiffer, F. (2020). A theoretical framework for comparing noise characteristics of spectral, differential phase-contrast and spectral differential phase-contrast X-ray imaging. *Physics in Medicine & Biology*.

Co-authored publications (peer-reviewed)

Marschner, M., Birnbacher, L., Mechlem, K., Noichl, W., Fehringer, A., Willner, M., Scherer, K., Herzen, J., Noël, P. B., Pfeiffer, F. (2016). Two-shot x-ray dark-field imaging. *Optics express*, 24(23), 27032-27045.

Ehn, S., Sellerer, T., Mechlem, K., Fehringer, A., Epple, M., Herzen, J., Pfeiffer F., Noël, P. B. (2016). Basis material decomposition in spectral CT using a semi-empirical, polychromatic adaption of the Beer–Lambert model. *Physics in Medicine & Biology*, 62(1), N1.

Eggl, E., Mechlem, K., Braig, E., Kulpe, S., Dierolf, M., Güunther, B., Achterhold, K., Herzen J., Gleich, B., Rummeny, E., Noël, P. B., Pfeiffer, F., Münzel, D. (2017). Mono-energy coronary angiography with a compact synchrotron source. *Scientific reports*, 7(1), 1-7.

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. (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.

Braig, E., Böhm, J., Dierolf, M., Jud, C., Günther, B., Mechlem, K., Allner, S., Sellerer, T., Achterhold, K., Gleich, B., Noël, P. B., Pfeiffer, D., Rummeny, E., Herzen, J., Pfeiffer, F. (2018). Direct quantitative material decomposition employing grating-based X-ray phase-contrast CT. *Scientific reports*, 8(1), 1-10.

Sellerer, T., Ehn, S., Mechlem, K., Duda, M., Epple, M., Noël, P. B., Pfeiffer, F. (2019). Quantitative dual-energy micro-CT with a photon-counting detector for material science and non-destructive testing. *PloS one*, 14(7).

Heck, L., Dierolf, M., Jud, C., Eggl, E., Sellerer, T., Mechlem, K., Günther, B., Achterhold, K., Gleich, B., Metz, S., Pfeiffer, D., Kröninger, K., Herzen, J. (2019). Contrast-enhanced spectral mammography with a compact synchrotron source. *PloS one*, 14(10).

Scherer, K., Hammel, J., Sellerer, T., Mechlem, K., Renger, B., Bähr, A., Kupatt, C., Hinkel, R., Herzen, J., Pfeiffer, F., Rummeny, E., Pfeiffer, D. (2019). Dynamic quantitative iodine myocardial perfusion imaging with dual-layer CT using a porcine model. *Scientific reports*, 9(1), 1-9.

Presentations at international conferences

Mechlem, K., Ehn, S., Sellerer, T., Pfeiffer, F., Noël, P. B. (2017). Statistical iterative material image reconstruction for spectral CT using a semi-empirical forward model. *SPIE Medical Imaging 2017: Physics of Medical Imaging* (poster).

Mechlem, K., Ehn, S., Sellerer, T., Pfeiffer, F., Noël, P. B. (2017). Statistical iterative reconstruction for spectral CT using ratios of polynomial functions. *14th International Meeting on FullyThree-Dimensional Image Reconstruction in Radiology and Nuclear Medicine* (talk).

Mechlem, K., Sellerer, T., Viermetz, M., Herzen, J., Pfeiffer, F. (2019). Spectral differential phase contrast X-ray radiography. *International Symposium on Biomedical Applications of X-Ray Phase Contrast Imaging* (talk).

Mechlem, K., Sellerer, T., Herzen, J., Pfeiffer, F. (2019). Statistical iterative reconstruction for spectral phase contrast CT. *15th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine* (talk).

Mechlem, K., Sellerer, T., Braig, E., Ehn, S., Pfeiffer, D., Herzen, J., Pfeiffer, F. (2019). Reconstruction and processing techniques for photon-counting data. *1st International Symposium on Photon Counting Technologies and Applications* (talk).

Acknowledgments

Throughout the entire period of this doctoral thesis, I have experienced great support, guidance and assistance from many different people. I would like to thank:

- **Franz Pfeiffer** for giving me the opportunity to conduct my research at the Chair of Biomedical Physics. Thank you for providing an excellent working and research environment and for your scientific guidance, especially concerning the combination of spectral and phase-contrast X-ray imaging.
- **Thorsten Sellerer** for his support with experimental measurements, countless insightful discussions and for the positive and productive working atmosphere. You could be described as my "experimental counterpart". I think we worked together very well as a team and managed to advance research in the fields of spectral and spectral phase-contrast X-ray imaging. Furthermore I would like to thank **Sebastian Ehn**, who was my second "experimental counterpart" during the first half of my time as a PhD student. Thank you for the enjoyable and productive time as a part of the "spectral island" team.
- **Julia Herzen and Peter Nöel** for their guidance, insightful discussions and for correcting my paper drafts.
- my office colleagues throughout the years: **Thorsten, Eva, Sebastian E., Manuel, Kirsten, Josef** for the friendly and positive atmosphere. You have made and important contribution to the fact that I thoroughly enjoyed my time at the Chair of Biomedical Physics.
- the members of the "reconstruction office" **Lorenz H., Wolfgang, Andreas, Sebastian A., Lorenz B.** for the interesting discussions and their help with algorithmic and IT-related problems. A special thanks goes to **Sebastian Allner** for his support with CT reconstructions and his helpful an supportive attitude during my time as a master student at the Chair of Biomedical Physics. You certainly contributed to my decision to do a PhD thesis in the same group.
- the E17 IT team (**Martin, Wolfgang, Sebastian A.**) for keeping the software and computing infrastructure running smoothly for the rest of us. You have saved us a tremendous amount of time with your work!
- **Lorenz H., Eva, Wolfgang and Thorsten** for proofreading my PhD thesis.
- **Nelly and Klaus** for their help with managing administrative work and other issues.
- all members of the group for the positive and welcoming atmosphere throughout the years.
- my family for their support during my time as a PhD student.

Bibliography

[1] Manuel Patino, Andrea Prochowski, Mukta D Agrawal, Frank J Simeone, Rajiv Gupta, Peter F Hahn, and Dushyant V Sahani. Material separation using dual-energy ct: current and emerging applications. *Radiographics*, 36(4):1087–1105, 2016.

[2] Robert E Alvarez and Albert Macovski. Energy-selective reconstructions in x-ray computerised tomography. *Physics in Medicine & Biology*, 21(5):733, 1976.

[3] Frederick Kelcz, Peter M Joseph, and Sadek K Hilal. Noise considerations in dual energy ct scanning. *Medical physics*, 6(5):418–425, 1979.

[4] D A Schauer and O W Linton. Ncrp report no. 160, ionizing radiation exposure of the population of the united states, medical exposure—are we doing less with more, and is there a role for health physicists? *Health physics*, 97(1):1–5, 2009.

[5] David J Brenner and Eric J Hall. Computed tomography—an increasing source of radiation exposure. *New England Journal of Medicine*, 357(22):2277–2284, 2007.

[6] Amy Berrington De González, Mahadevappa Mahesh, Kwang-Pyo Kim, Mythreyi Bhargavan, Rebecca Lewis, Fred Mettler, and Charles Land. Projected cancer risks from computed tomographic scans performed in the united states in 2007. *Archives of internal medicine*, 169(22):2071–2077, 2009.

[7] Peter B Noël, Bernhard Renger, Martin Fiebich, Daniela Münzel, Alexander A Fingerle, Ernst J Rummeny, and Martin Dobritz. Does iterative reconstruction lower ct radiation dose: evaluation of 15,000 examinations. *PloS one*, 8(11):e81141, 2013.

[8] Peter B Noël, Alexander A Fingerle, Bernhard Renger, Daniela Münzel, Ernst J Rummeny, and Martin Dobritz. Initial performance characterization of a clinical noise–suppressing reconstruction algorithm for mdct. *American Journal of Roentgenology*, 197(6):1404–1409, 2011.

[9] Amy K Hara, Robert G Paden, Alvin C Silva, Jennifer L Kujak, Holly J Lawder, and William Pavlicek. Iterative reconstruction technique for reducing body radiation dose at ct: feasibility study. *American Journal of Roentgenology*, 193(3):764–771, 2009.

[10] Andreas Sauter, Thomas Koehler, Alexander A Fingerle, Bernhard Brendel, Vivien Richter, Michael Rasper, Ernst J Rummeny, Peter B Noël, and Daniela Münzel. Ultra low dose ct pulmonary angiography with iterative reconstruction. *PloS one*, 11(9):e0162716, 2016.

[11] Francis Anghinolfi, P Aspell, K Bass, Werner Beusch, L Bosisio, C Boutonnet, P Burger, M Campbell, E Chesi, C Claeys, et al. A 1006 element hybrid silicon pixel detector with strobed binary output. *IEEE transactions on nuclear science*, 39(4):654–661, 1992.

[12] F Antinori, G Alexeev, and E Andersen. First results from the 1994 lead beam run of wa97. *Nuclear Physics. A*, 590(1-2):139c–146c, 1995.

[13] Franz Pfeiffer, Timm Weitkamp, Oliver Bunk, and Christian David. Phase retrieval and differential phase-contrast imaging with low-brilliance x-ray sources. *Nature physics*, 2(4):258, 2006.

[14] Franz Pfeiffer, Christian Kottler, O Bunk, and C David. Hard x-ray phase tomography with low-brilliance sources. *Physical review letters*, 98(10):108105, 2007.

[15] Julia Herzen, Tilman Donath, Franz Pfeiffer, Oliver Bunk, Celestino Padeste, Felix Beckmann, Andreas Schreyer, and Christian David. Quantitative phase-contrast tomography of a liquid phantom using a conventional x-ray tube source. *Optics express*, 17(12):10010–10018, 2009.

[16] Joseph Zambelli, Nicholas Bevins, Zhihua Qi, and Guang-Hong Chen. Radiation dose efficiency comparison between differential phase contrast CT and conventional absorption CT. *Medical physics*, 37(6Part1):2473–2479, 2010.

[17] Tilman Donath, Franz Pfeiffer, Oliver Bunk, Christian Grünzweig, Eckhard Hempel, Stefan Popescu, Peter Vock, and Christian David. Toward clinical x-ray phase-contrast CT: demonstration of enhanced soft-tissue contrast in human specimen. *Investigative radiology*, 45(7):445–452, 2010.

[18] Guang-Hong Chen, Joseph Zambelli, Ke Li, Nicholas Bevins, and Zhihua Qi. Scaling law for noise variance and spatial resolution in differential phase contrast computed tomography. *Medical physics*, 38(2):584–588, 2011.

[19] Klaus J Engel, D Geller, T Köhler, G Martens, S Schusser, G Vogtmeier, and E Rössl. Contrast-to-noise in x-ray differential phase contrast imaging. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 648:S202–S207, 2011.

[20] Huajie Han, Renfang Hu, Faiz Wali, Zhao Wu, Kun Gao, Shenghao Wang, Yonggang Gu, Yi Jin, and Chao Zhai. Phase-contrast imaging for body composition measurement. *Physica Medica*, 43:25–33, 2017.

[21] Eva Braig, Jessica Böhm, Martin Dierolf, Christoph Jud, Benedikt Günther, Korbinian Mechlem, Sebastian Allner, Thorsten Sellerer, Klaus Achterhold, Bernhard Gleich, et al. Direct quantitative material decomposition employing grating-based x-ray phase-contrast CT. *Scientific reports*, 8(1):16394, 2018.

[22] Franz Pfeiffer, Martin Bech, Oliver Bunk, Philipp Kraft, Eric F Eikenberry, Ch Brönnimann, Christian Grünzweig, and Christian David. Hard-x-ray dark-field imaging using a grating interferometer. *Nature materials*, 7(2):134, 2008.

[23] Michael Chabior, Tilman Donath, Christian David, Oliver Bunk, Manfred Schuster, Christian Schroer, and Franz Pfeiffer. Beam hardening effects in grating-based x-ray phase-contrast imaging. *Medical physics*, 38(3):1189–1195, 2011.

[24] Wataru Yashiro, Patrik Vagovič, and Atsushi Momose. Effect of beam hardening on a visibility-contrast image obtained by x-ray grating interferometry. *Optics express*, 23(18):23462–23471, 2015.

[25] Wilhelm Haas, Martin Bech, Peter Bartl, Florian Bayer, André Ritter, Thomas Weber, Georg Pelzer, Marian Willner, Klaus Achterhold, Jürgen Durst, et al. Phase-unwrapping of differential phase-contrast data using attenuation information. In *Medical Imaging 2011: Image Processing*, volume 7962, page 79624R. International Society for Optics and Photonics, 2011.

[26] I Zanette, T Weitkamp, S Lang, M Langer, J Mohr, C David, and J Baruchel. Quantitative phase and absorption tomography with an x-ray grating interferometer and synchrotron radiation. *physica status solidi (a)*, 208(11):2526–2532, 2011.

[27] Salim Si-Mohamed, Daniel Bar-Ness, Monica Sigovan, David P Cormode, Philippe Coulon, Emmanuel Coche, Alain Vlassenbroek, Gabrielle Normand, Loic Boussel, and Philippe Douek. Review of an initial experience with an experimental spectral photon-counting computed tomography system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 873:27–35, 2017.

[28] Julia Dangelmaier, Daniel Bar-Ness, Heiner Daerr, Daniela Muenzel, Salim Si-Mohamed, Sebastian Ehn, Alexander A Fingerle, Melanie A Kimm, Felix K Kopp, Loic Boussel, et al. Experimental feasibility of spectral photon-counting computed tomography with two contrast agents for the detection of endoleaks following endovascular aortic repair. *European radiology*, 28(8):3318–3325, 2018.

[29] Felix K Kopp, Heiner Daerr, Salim Si-Mohamed, Andreas P Sauter, Sebastian Ehn, Alexander A Fingerle, Bernhard Brendel, Franz Pfeiffer, Ewald Roessl, Ernst J Rummeny, et al. Evaluation of a preclinical photon-counting ct prototype for pulmonary imaging. *Scientific reports*, 8(1):17386, 2018.

[30] Amir Pourmorteza, Rolf Symons, Veit Sandfort, Marissa Mallek, Matthew K Fuld, Gregory Henderson, Elizabeth C Jones, Ashkan A Malayeri, Les R Folio, and David A Bluemke. Abdominal imaging with contrast-enhanced photon-counting ct: first human experience. *Radiology*, 279(1):239–245, 2016.

[31] Rolf Symons, Daniel S Reich, Mohammadhadi Bagheri, Tyler E Cork, Bernhard Krauss, Stefan Ulzheimer, Steffen Kappler, David A Bluemke, and Amir Pourmorteza. Photon-counting computed tomography for vascular imaging of the head and neck: First in vivo human results. *Investigative radiology*, 53(3):135–142, 2018.

[32] Rolf Symons, Amir Pourmorteza, Veit Sandfort, Mark A Ahlman, Tracy Cropper, Marissa Mallek, Steffen Kappler, Stefan Ulzheimer, Mahadevappa Mahesh, Elizabeth C Jones, et al. Feasibility of dose-reduced chest ct with photon-counting detectors: initial results in humans. *Radiology*, 285(3):980–989, 2017.

[33] Shuai Leng, Wei Zhou, Zhicong Yu, Ahmed Halaweish, Bernhard Krauss, Bernhard Schmidt, Lifeng Yu, Steffen Kappler, and Cynthia McCollough. Spectral performance of a whole-body research photon counting detector ct: quantitative accuracy in derived image sets. *Physics in Medicine & Biology*, 62(17):7216, 2017.

[34] Shuai Leng, Kishore Rajendran, Hao Gong, Wei Zhou, Ahmed F Halaweish, Andre Henning, Steffen Kappler, Matthias Baer, Joel G Fletcher, and Cynthia H McCollough. 150-μm spatial resolution using photon-counting detector computed tomography technology: technical performance and first patient images. *Investigative radiology*, 53(11):655–662, 2018.

[35] Jens Als-Nielsen and Des McMorrow. *Elements of modern X-ray physics*. John Wiley & Sons, 2011.

[36] Paul Emma, R Akre, J Arthur, R Bionta, C Bostedt, J Bozek, A Brachmann, P Bucksbaum, Ryan Coffee, F-J Decker, et al. First lasing and operation of an ångstrom-wavelength free-electron laser. *nature photonics*, 4(9):641, 2010.

[37] Zhirong Huang and Ronald D Ruth. Laser-electron storage ring. *Physical review letters*, 80(5):976, 1998.

[38] Martin Bech, Oliver Bunk, Christian David, Ronald Ruth, Jeff Rifkin, Rod Loewen, Robert Feidenhans'l, and Franz Pfeiffer. Hard x-ray phase-contrast imaging with the compact light source based on inverse compton x-rays. *Journal of synchrotron radiation*, 16(1):43–47, 2009.

[39] Elena Eggl, Martin Dierolf, Klaus Achterhold, Christoph Jud, Benedikt Günther, Eva Braig, Bernhard Gleich, and Franz Pfeiffer. The munich compact light source: initial performance measures. *Journal of synchrotron radiation*, 23(5):1137–1142, 2016.

[40] Rolf Behling. *Modern diagnostic x-ray sources: technology, manufacturing, reliability*. CRC Press, 2015.

[41] Thorsten M Buzug. Computed tomography. In *Springer Handbook of Medical Technology*, pages 311–342. Springer, 2011.

[42] John M Boone and J Anthony Seibert. An accurate method for computer-generating tungsten anode x-ray spectra from 30 to 140 kv. *Medical physics*, 24(11):1661–1670, 1997.

[43] David Paganin. *Coherent X-ray optics*. Number 6. Oxford University Press on Demand, 2006.

[44] Joseph W Goodman. *Introduction to Fourier optics*. Roberts and Company Publishers, 2005.

[45] Arne Nordmann. Wave diffraction in the manner of huygens and fresnel. `http://en.wikipedia.org/wiki/Huygens-Fresnel_principle#/media/File:Refraction_on_an_aperture_-_Huygens-Fresnel_principle.svg`, 2007. [Online; accessed 11-November-2019].

[46] Frederik Zernike. The concept of degree of coherence and its application to optical problems. *Physica*, 5(8):785–795, 1938.

[47] Joseph W Goodman. *Statistical optics*. John Wiley & Sons, 2015.

[48] Kaye Susannah Morgan, Karen Kit Wan Siu, and DM Paganin. The projection approximation and edge contrast for x-ray propagation-based phase contrast imaging of a cylindrical edge. *Optics express*, 18(10):9865–9878, 2010.

[49] Kaye S Morgan, Karen KW Siu, and David M Paganin. The projection approximation versus an exact solution for x-ray phase contrast imaging, with a plane wave scattered by a dielectric cylinder. *Optics Communications*, 283(23):4601–4608, 2010.

[50] Jason Daniel Schmidt. Numerical simulation of optical wave propagation with examples in matlab. SPIE Bellingham, Washington, USA, 2010.

[51] Max Born and Emil Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.

[52] Joseph B Keller. Geometrical theory of diffraction. *Josa*, 52(2):116–130, 1962.

[53] Jeffrey A Fessler, M Sonka, and J Michael Fitzpatrick. Statistical image reconstruction methods for transmission tomography. *Handbook of medical imaging*, 2:1–70, 2000.

[54] Ervin B Podgoršak et al. *Radiation physics for medical physicists*. Springer, 2006.

[55] DJ Hawkes and Daphne F Jackson. An accurate parametrisation of the x-ray attenuation coefficient. *Physics in Medicine & Biology*, 25(6):1167, 1980.

[56] DR White. An analysis of the z-dependence of photon and electron interactions. *Physics in Medicine & Biology*, 22(2):219, 1977.

[57] JH Hubbell. Photon mass attenuation and energy-absorption coefficients. *The International Journal of Applied Radiation and Isotopes*, 33(11):1269–1290, 1982.

[58] Nist x-ray attenuation database. `https://www.nist.gov/pml/x-ray-mass-attenuation-coefficients`. Online: accessed: 11-September-2019.

[59] Oskar Klein and Yoshio Nishina. Über die streuung von strahlung durch freie elektronen nach der neuen relativistischen quantendynamik von dirac. *Zeitschrift für Physik*, 52(11-12):853–868, 1929.

[60] Norman Allen Dyson. *X-rays in Atomic and Nuclear Physics*. Cambridge University Press, 2005.

[61] JH Hubbell, Wm J Veigele, EA Briggs, RT Brown, DT Cromer, and d RJ Howerton. Atomic form factors, incoherent scattering functions, and photon scattering cross sections. *Journal of physical and chemical reference data*, 4(3):471–538, 1975.

[62] Daphne F Jackson. Chemical effects in x-ray transmission measurements. *Nuclear Instruments and Methods in Physics Research*, 193(1-2):387–389, 1982.

[63] Glenn F Knoll. *Radiation detection and measurement*. John Wiley & Sons, 2010.

[64] Martin Nikl. Scintillation detectors for x-rays. *Measurement Science and Technology*, 17(4):R37, 2006.

[65] Efrat Shefer, Ami Altman, Rolf Behling, Raffy Goshen, Lev Gregorian, Yalon Roterman, Igor Uman, Naor Wainer, Yoad Yagil, and Oren Zarchin. State of the art of ct detectors and sources: a literature review. *Current Radiology Reports*, 1(1):76–91, 2013.

[66] Michael Overdick, Christian Baumer, Klaus Juergen Engel, Johannes Fink, Christoph Herrmann, Hans Kruger, Matthias Simon, Roger Steadman, and Gnter Zeitler. Status of direct conversion detectors for medical imaging with x-rays. *IEEE Transactions on Nuclear Science*, 56(4):1800–1809, 2009.

[67] Bruce R Whiting, Parinaz Massoumzadeh, Orville A Earl, Joseph A O'Sullivan, Donald L Snyder, and Jeffrey F Williamson. Properties of preprocessed sinogram data in x-ray computed tomography. *Medical physics*, 33(9):3290–3303, 2006.

[68] Xiaolan Wang, Alexander Zamyatin, and Daxin Shi. Dose reduction potential with photon counting computed tomography. In *Medical Imaging 2012: Physics of Medical Imaging*, volume 8313, page 831349. International Society for Optics and Photonics, 2012.

[69] Rolf Symons, Tyler E Cork, Pooyan Sahbaee, Matthew K Fuld, Steffen Kappler, Les R Folio, David A Bluemke, and Amir Pourmorteza. Low-dose lung cancer screening with photon-counting ct: a feasibility study. *Physics in Medicine & Biology*, 62(1):202, 2016.

[70] Korbinian Mechlem. Advanced statistical iterative reconstruction for x-ray computed tomography. Master's thesis, Technical University of Munich, 2016.

[71] Zhicong Yu, Shuai Leng, Steven M Jorgensen, Zhoubo Li, Ralf Gutjahr, Baiyu Chen, Ahmed F Halaweish, Steffen Kappler, Lifeng Yu, Erik L Ritman, et al. Evaluation of conventional imaging performance in a research whole-body ct system with a photon-counting detector array. *Physics in Medicine & Biology*, 61(4):1572, 2016.

[72] Ralf Gutjahr, Ahmed F Halaweish, Zhicong Yu, Shuai Leng, Lifeng Yu, Zhoubo Li, Steven M Jorgensen, Erik L Ritman, Steffen Kappler, and Cynthia H McCollough. Human imaging with photon-counting-based ct at clinical dose levels: Contrast-to-noise ratio and cadaver studies. *Investigative radiology*, 51(7):421, 2016.

[73] RN Cahn, B Cederström, Mats Danielsson, A Hall, M Lundqvist, and D Nygren. Detective quantum efficiency dependence on x-ray energy weighting in mammography. *Medical Physics*, 26(12):2680–2683, 1999.

[74] Jürgen Giersch, Daniel Niederlöhner, and Gisela Anton. The influence of energy weighting on x-ray imaging quality. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 531(1-2):68–74, 2004.

[75] Taly Gilat Schmidt. Optimal "image-based" weighting for energy-resolved ct. *Medical physics*, 36(7):3018–3027, 2009.

[76] Kesava S Kalluri, Mufeed Mahd, and Stephen J Glick. Investigation of energy weighting using an energy discriminating photon counting detector for breast ct. *Medical physics*, 40(8):081923, 2013.

[77] Polad M Shikhaliev. Projection x-ray imaging with photon energy weighting: experimental evaluation with a prototype detector. *Physics in Medicine & Biology*, 54(16):4971, 2009.

[78] Katsuyuki Taguchi and Jan S Iwanczyk. Vision 20/20: Single photon counting x-ray detectors in medical imaging. *Medical physics*, 40(10), 2013.

[79] R Ballabriga, J Alozy, M Campbell, E Frojdh, E.H.M. Heijne, T Koenig, X Llopart, J Marchal, D. Pennicard, T Poikela, L Tlustos, P Valerio, W Wong, and M Zuber. Review of hybrid pixel detector readout ASICs for spectroscopic X-ray imaging. *Journal of Instrumentation*, (01):P01007–P01007, jan 2016.

[80] S. Kappler, A. Henning, B. Kreisler, F. Schoeck, K. Stierstorfer, and T. Flohr. Photon counting ct at elevated x-ray tube currents: contrast stability, image noise and multi-energy performance. *Proc.SPIE*, 9033:9033 – 9033 – 8, 2014.

[81] Peter Trueb, Pietro Zambon, and Christian Broennimann. Assessment of the spectral performance of hybrid photon counting x-ray detectors. *Medical Physics*, 44(9):207–214, 2017.

[82] Gregor Hülsen, Christian Broennimann, Eric F Eikenberry, and Armin Wagner. Protein crystallography with a novel large-area pixel detector. *Journal of applied crystallography*, 39(4):550–557, 2006.

[83] John Vallerga, Jason McPhate, Anton Tremsin, Oswald Siegmund, Bettina Mikulec, and Allan Clark. Optically sensitive medipix2 detector for adaptive optics wavefront sensing. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 546(1-2):263–269, 2005.

[84] Sebastian Ehn. *Photon-counting hybrid-pixel detectors for spectral X-ray imaging applications.* PhD thesis, Technical University of Munich, 2017.

[85] Thorsten Sellerer. Projection based dual-energy micro-ct with a photon-counting detector. Master's thesis, Technical University of Munich, 2016.

[86] Christer Ullberg, Mattias Urech, Niclas Weber, Anders Engman, Anna Redz, and Fredrik Henckel. Measurements of a dual-energy fast photon counting cdte detector with integrated charge sharing correction. In *Medical Imaging 2013: Physics of Medical Imaging*, volume 8668, page 86680P. International Society for Optics and Photonics, 2013.

[87] R Ballabriga, G Blaj, M Campbell, M Fiederle, D Greiffenberg, EHM Heijne, X Llopart, R Plackett, S Procz, L Tlustos, et al. Characterization of the medipix3 pixel readout chip. *Journal of Instrumentation*, 6(01):C01052, 2011.

[88] Katsuyuki Taguchi, Eric C Frey, Xiaolan Wang, Jan S Iwanczyk, and William C Barber. An analytical model of the effects of pulse pileup on the energy spectrum recorded by energy resolved photon counting x-ray detectors. *Medical physics*, 37(8):3957–3969, 2010.

[89] Ewald Roessl, Heiner Daerr, and Roland Proksa. A fourier approach to pulse pile-up in photon-counting x-ray detectors. *Medical physics*, 43(3):1295–1298, 2016.

[90] Jochen Cammin, Jennifer Xu, William C Barber, Jan S Iwanczyk, Neal E Hartsough, and Katsuyuki Taguchi. A cascaded model of spectral distortions due to spectral response effects and pulse pileup effects in a photon-counting x-ray detector for ct. *Medical physics*, 41(4):041905, 2014.

[91] Adam S Wang, Daniel Harrison, Vladimir Lobastov, and J Eric Tkaczyk. Pulse pileup statistics for energy discriminating photon counting x-ray detectors. *Medical physics*, 38(7):4265–4275, 2011.

[92] RO Bell, G Entine, and HB Serreze. Time-dependent polarization of cdte gamma-ray detectors. *Nuclear Instruments and Methods*, 117(1):267–271, 1974.

[93] Tomoyuki Seino, Shinya Kominami, Yuichiro Ueno, and Kensuke Amemiya. Pulsed bias voltage shutdown to suppress the polarization effect for a cdte radiation detector. *IEEE Transactions on Nuclear Science*, 55(5):2770–2774, 2008.

[94] Jacques Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton university bulletin*, pages 49–52, 1902.

[95] Jeffrey A Fessler. Hybrid poisson/polynomial objective functions for tomographic image reconstruction from transmission scans. *IEEE Transactions on Image Processing*, 4(10):1439–1450, 1995.

[96] Steven M Kay. *Fundamentals of statistical signal processing*. Prentice Hall PTR, 1993.

[97] Louis L Scharf. *Statistical signal processing*, volume 98. Addison-Wesley Reading, MA, 1991.

[98] Gabor T Herman. Two direct methods for reconstructing pictures from their projections: A comparative study. *Computer Graphics and Image Processing*, 1(2):123–144, 1972.

[99] Maria Magnusson. Linogram and other direct fourier methods for tomographic reconstruction. 1993.

[100] Henrik Turbell. *Cone-beam reconstruction using filtered backprojection*. PhD thesis, Linköping University Electronic Press, 2001.

[101] Gengsheng Lawrence Zeng. *Medical image reconstruction: a conceptual tutorial*. Springer, 2010.

[102] Yangbo Ye and Ge Wang. Filtered backprojection formula for exact image reconstruction from cone-beam data along a general scanning curve. *Medical Physics*, 32(1):42–48, 2005.

[103] Lee A Feldkamp, LC Davis, and James W Kress. Practical cone-beam algorithm. *Josa a*, 1(6):612–619, 1984.

[104] Alexander Katsevich. Theoretically exact filtered backprojection-type inversion algorithm for spiral ct. *SIAM Journal on Applied Mathematics*, 62(6):2012–2026, 2002.

[105] Stefan Schaller, Thomas Flohr, Klaus Klingenbeck, Jens Krause, Theobald Fuchs, and Willi A Kalender. Spiral interpolation algorithm for multislice spiral ct. i. theory. *IEEE transactions on medical imaging*, 19(9):822–834, 2000.

[106] Theobald Fuchs, Jens Krause, Stefan Schaller, Thomas Flohr, and Willi A Kalender. 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, 2000.

[107] Johann Radon. On the determination of functions from their integral values along certain manifolds. *IEEE transactions on medical imaging*, 5(4):170–176, 1986.

[108] Mahsa Noori Asl and Alireza Sadremomtaz. Analytical image reconstruction methods in emission tomography. 2013.

[109] Avinash C Kak, Malcolm Slaney, and Ge Wang. Principles of computerized tomographic imaging. *Medical Physics*, 29(1):107–107, 2002.

[110] Robert M Lewitt. Alternatives to voxels for image representation in iterative reconstruction algorithms. *Physics in Medicine & Biology*, 37(3):705, 1992.

[111] Andy Ziegler, Thomas Köhler, Tim Nielsen, and Roland Proksa. Efficient projection and backprojection scheme for spherically symmetric basis functions in divergent beam geometry. *Medical physics*, 33(12):4653–4663, 2006.

[112] Rolf-Dieter Bippus, Thomas Koehler, Frank Bergner, Bernhard Brendel, Eberhard Hansis, and Roland Proksa. Projector and backprojector for iterative ct reconstruction with blobs using cuda. In *Fully 3D 2011: 11th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine, Potsdam, Germany, 11-15 July 2011*, 2011.

[113] Pavel Kisilev, Michael Zibulevsky, and Yehoshua Y Zeevi. Wavelet representation and total variation regularization in emission tomography. In *Proceedings 2001 International Conference on Image Processing (Cat. No. 01CH37205)*, volume 1, pages 702–705. IEEE, 2001.

[114] Kenneth M Hanson and George W Wecksung. Local basis-function approach to computed tomography. *Applied Optics*, 24(23):4028–4039, 1985.

[115] Johan Nuyts, Bruno De Man, Jeffrey A Fessler, Wojciech Zbijewski, and Freek J Beekman. Modelling the physics in the iterative reconstruction for transmission computed tomography. *Physics in Medicine & Biology*, 58(12):R63, 2013.

[116] Xiaochuan Pan, Emil Y Sidky, and Michael Vannier. Why do commercial ct scanners still employ traditional, filtered back-projection for image reconstruction? *Inverse problems*, 25(12):123009, 2009.

[117] Wojciech Zbijewski and Freek J Beekman. Comparison of methods for suppressing edge and aliasing artefacts in iterative x-ray ct reconstruction. *Physics in Medicine & Biology*, 51(7):1877, 2006.

[118] Andreas Fehringer. *Real-time iterative reconstruction for x-ray computed tomography*. PhD thesis, Technical University of Munich, 2019.

[119] Andreas Fehringer, Tobias Lasser, Irene Zanette, Peter B Noël, and Franz Pfeiffer. 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, 2014.

[120] Van-Giang Nguyen and Soo-Jin Lee. Parallelizing a matched pair of ray-tracing projector and backprojector for iterative cone-beam ct reconstruction. *IEEE Transactions on Nuclear Science*, 62(1):171–181, 2015.

[121] Yi Du, Gongyi Yu, Xincheng Xiang, and Xiangang Wang. Gpu accelerated voxel-driven forward projection for iterative reconstruction of cone-beam ct. *Biomedical engineering online*, 16(1):2, 2017.

[122] Philippe Despres and Xun Jia. A review of gpu-based medical image reconstruction. *Physica Medica*, 42:76–92, 2017.

[123] Bruno De Man and Samit Basu. Distance-driven projection and backprojection in three dimensions. *Physics in Medicine & Biology*, 49(11):2463, 2004.

[124] B De Man and JA Fessler. Statistical iterative reconstruction for x-ray computed tomography. *Biomedical Mathematics: Promising Directions in Imaging, Therapy Planning and Inverse Problems*, pages 113–40, 2010.

[125] Peter M Joseph. An improved algorithm for reprojecting rays through pixel images. *IEEE transactions on medical imaging*, 1(3):192–196, 1982.

[126] Bruno De Man and Samit Basu. Distance-driven projection and backprojection. In *2002 IEEE Nuclear Science Symposium Conference Record*, volume 3, pages 1477–1480. IEEE, 2002.

[127] TM Peters. Algorithms for fast back-and re-projection in computed tomography. *IEEE transactions on nuclear science*, 28(4):3641–3647, 1981.

[128] W Zhuang, SS Gopal, and TJ Hebert. Numerical evaluation of methods for computing tomographic projections. *IEEE transactions on nuclear science*, 41(4):1660–1665, 1994.

[129] Gengsheng Zeng. Counter examples for unmatched projector/backprojector in an iterative algorithm. *Chinese Journal of Academic Radiology*, 04 2019.

[130] Gengsheng L Zeng and Grant T Gullberg. Unmatched projector/backprojector pairs in an iterative reconstruction algorithm. *IEEE transactions on medical imaging*, 19(5):548, 2000.

[131] Filippo Arcadu, Marco Stampanoni, and Federica Marone. On the crucial impact of the coupling projector-backprojector in iterative tomographic reconstruction. *arXiv preprint arXiv:1612.05515*, 2016.

[132] Yong Long, Jeffrey A Fessler, and James M Balter. 3d forward and back-projection for x-ray ct using separable footprints. *IEEE transactions on medical imaging*, 29(11):1839–1850, 2010.

[133] Xiaobin Xie, Madison G McGaffin, Yong Long, Jeffrey A Fessler, Minhua Wen, and James Lin. 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, 2017.

[134] Camille Chapdelaine, Nicolas Gac, Ali-Mohammad Djafari, and Estelle Parra-Denis. New gpu implementation of separable footprint (sf) projector and backprojector: first results. In *The Fifth International Conference on Image Formation in X-Ray Computed Tomography*, pages 314–317, 2018.

[135] Rui Liu, Lin Fu, Bruno De Man, and Hengyong Yu. Gpu-based branchless distance-driven projection and backprojection. *IEEE transactions on computational imaging*, 3(4):617–632, 2017.

[136] Daniel Schlifske and Henry Medeiros. A fast gpu-based approach to branchless distance-driven projection and back-projection in cone beam ct. In *Medical Imaging 2016: Physics of Medical Imaging*, volume 9783, page 97832W. International Society for Optics and Photonics, 2016.

[137] Bruno De Man, Johan Nuyts, Patrick Dupont, Guy Marchal, and Paul Suetens. An iterative maximum-likelihood polychromatic algorithm for ct. *IEEE transactions on medical imaging*, 20(10):999–1008, 2001.

[138] Idris A Elbakri and Jeffrey A Fessler. Segmentation-free statistical image reconstruction for polyenergetic x-ray computed tomography. In *Proceedings IEEE International Symposium on Biomedical Imaging*, pages 828–831. IEEE, 2002.

[139] Idris A Elbakri and Jeffrey A Fessler. Segmentation-free statistical image reconstruction for polyenergetic x-ray computed tomography with experimental validation. *Physics in Medicine & Biology*, 48(15):2453, 2003.

[140] Jing Wang, Weihua Mao, and Timothy Solberg. Scatter correction for cone-beam computed tomography using moving blocker strips: A preliminary study. *Medical physics*, 37(11):5792–5800, 2010.

[141] Kevin J Little and Patrick J La Rivière. An algorithm for modeling non-linear system effects in iterative ct reconstruction. In *IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, pages 2174–2177. IEEE, 2012.

[142] Bing Feng, Jeffrey A Fessler, and Michael A King. 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, 2006.

[143] Steven Tilley II, Jeffrey H Siewerdsen, and J Webster Stayman. 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, 2015.

[144] Lorenz Hehn, Steven Tilley, Franz Pfeiffer, and J Webster Stayman. Blind deconvolution in model-based iterative reconstruction for ct using a normalized sparsity measure. *Physics in Medicine & Biology*, 64(21):215010, 2019.

[145] Steven Tilley II, Jeffrey H Siewerdsen, Wojciech Zbijewski, and J Webster Stayman. 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, 2016.

[146] AA Isola, Andreas Ziegler, T Koehler, WJ Niessen, and Michael Grass. Motion-compensated iterative cone-beam ct image reconstruction with adapted blobs as basis functions. *Physics in Medicine & Biology*, 53(23):6777, 2008.

[147] Qiulin Tang, Jochen Cammin, Somesh Srivastava, and Katsuyuki Taguchi. A fully four-dimensional, iterative motion estimation and compensation method for cardiac ct. *Medical physics*, 39(7Part1):4291–4305, 2012.

[148] Jianhua Ma, Zhengrong Liang, Yi Fan, Yan Liu, Jing Huang, Wufan Chen, and Hongbing Lu. Variance analysis of x-ray ct sinograms in the presence of electronic noise background. *Medical physics*, 39(7Part1):4051–4065, 2012.

[149] Idris A Elbakri and Jeffrey A Fessler. Efficient and accurate likelihood for iterative image reconstruction in x-ray computed tomography. In *Medical Imaging 2003: Image Processing*, volume 5032, pages 1839–1850. International Society for Optics and Photonics, 2003.

[150] F Yu Daniel and Jeffrey A Fessler. Mean and variance of single photon counting with deadtime. *Physics in Medicine & Biology*, 45(7):2043, 2000.

[151] Jean-Baptiste Thibault, Charles A Bouman, Ken D Sauer, and Jiang Hsieh. A recursive filter for noise reduction in statistical iterative tomographic imaging. In *Electronic Imaging 2006*, pages 60650X–60650X. International Society for Optics and Photonics, 2006.

[152] Hao Zhang, Jing Wang, Dong Zeng, Xi Tao, and Jianhua Ma. Regularization strategies in statistical image reconstruction of low-dose x-ray ct: A review. *Medical physics*, 45(10):e886–e907, 2018.

[153] J Webster Stayman and Jeffrey A Fessler. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE transactions on medical imaging*, 19(6):601–615, 2000.

[154] Jang Hwan Cho and Jeffrey A Fessler. Regularization designs for uniform spatial resolution and noise properties in statistical image reconstruction for 3-d x-ray ct. *IEEE transactions on medical imaging*, 34(2):678–689, 2014.

[155] Jang Hwan Cho and Jeffrey A Fessler. Quadratic regularization design for 3d axial ct: Towards isotropic noise. In *2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC)*, pages 1–5. IEEE, 2013.

[156] Zhiqian Chang, Ruoqiao Zhang, Jean-Baptiste Thibault, Debashish Pal, Lin Fu, Ken Sauer, and Charles Bouman. Adaptive regularization for uniform noise covariance in iterative 3d ct. In *Proc. 13th Int. Meeting Fully Three-Dimensional Image Reconstruction Radiol. Nucl. Med.*, pages 256–259, 2015.

[157] Peter J Huber. *Robust statistics*. Springer, 2011.

[158] Zhen Tian, Xun Jia, Kehong Yuan, Tinsu Pan, and Steve B Jiang. Low-dose ct reconstruction via edge-preserving total variation regularization. *Physics in Medicine & Biology*, 56(18):5949, 2011.

[159] Emil Y Sidky and Xiaochuan Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology*, 53(17):4777, 2008.

[160] Junguo Bian, Jeffrey H Siewerdsen, Xiao Han, Emil Y Sidky, Jerry L Prince, Charles A Pelizzari, and Xiaochuan Pan. Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam ct. *Physics in Medicine & Biology*, 55(22):6575, 2010.

[161] Maurice Debatin. *Accurate Low-Dose Iterative CT Reconstruction from Few Projections using Sparse and Non-Local Regularization Functions*. PhD thesis, Ruprecht Karl University of Heidelberg, 2016.

[162] Michael Elad. On the origin of the bilateral filter and ways to improve it. *IEEE Transactions on image processing*, 11(10):1141–1151, 2002.

[163] Wei Xu and Klaus Mueller. Evaluating popular non-linear image processing filters for their use in regularized iterative ct. In *IEEE Nuclear Science Symposuim & Medical Imaging Conference*, pages 2864–2865. IEEE, 2010.

[164] Hua Zhang, Jing Huang, Jianhua Ma, Zhaoying Bian, Qianjin Feng, Hongbing Lu, Zhengrong Liang, and Wufan Chen. Iterative reconstruction for x-ray computed tomography using prior-image induced nonlocal regularization. *IEEE Transactions on Biomedical Engineering*, 61(9):2367–2378, 2013.

[165] J Webster Stayman, Hao Dang, Yifu Ding, and Jeffrey H Siewerdsen. Pirple: a penalized-likelihood framework for incorporation of prior images in ct reconstruction. *Physics in Medicine & Biology*, 58(21):7563, 2013.

[166] Pascal Thériault Lauzier and Guang-Hong Chen. Characterization of statistical prior image constrained compressed sensing. i. applications to time-resolved contrast-enhanced ct. *Medical physics*, 39(10):5930–5948, 2012.

[167] Abolfazl Mehranian, Mohammad Reza Ay, Arman Rahmim, and Habib Zaidi. X-ray ct metal artifact reduction using wavelet domain $l\_\{0\}$ sparse regularization. *IEEE transactions on medical imaging*, 32(9):1707–1722, 2013.

[168] Sathish Ramani and Jeffrey A Fessler. A splitting-based iterative algorithm for accelerated statistical x-ray ct reconstruction. *IEEE transactions on medical imaging*, 31(3):677–688, 2011.

[169] Jorge Nocedal. Updating quasi-newton matrices with limited storage. *Mathematics of computation*, 35(151):773–782, 1980.

[170] Yu-Hong Dai and Yaxiang Yuan. A nonlinear conjugate gradient method with a strong global convergence property. *SIAM Journal on optimization*, 10(1):177–182, 1999.

[171] Jonathan Richard Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. `http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf`, 1994. [Online; accessed 25-March-2015].

[172] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.

[173] Jorge Nocedal. Conjugate gradient methods and nonlinear optimization. *Linear and nonlinear conjugate gradient-related methods*, pages 9–23, 1996.

[174] William W Hager and Hongchao Zhang. A survey of nonlinear conjugate gradient methods. *Pacific journal of Optimization*, 2(1):35–58, 2006.

[175] Yu-Hong Dai. Nonlinear conjugate gradient methods. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.

[176] Oleg Alexandrov. Illustration of the conjugate gradient method. `http://en.wikipedia.org/wiki/Conjugate_gradient_method#/media/File:Conjugate_gradient_illustration.svg`, 2007. [Online; accessed 11-November-2019].

[177] Jeffrey A Fessler and Scott D Booth. Conjugate-gradient preconditioning methods for shift-variant pet image reconstruction. *IEEE transactions on image processing*, 8(5):688–699, 1999.

[178] Magnus Rudolph Hestenes and Eduard Stiefel. *Methods of conjugate gradients for solving linear systems*, volume 49. NBS Washington, DC, 1952.

[179] Elijah Polak and Gerard Ribiere. Note sur la convergence de méthodes de directions conjuguées. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique*, 3(R1):35–43, 1969.

[180] Philip Wolfe. Convergence conditions for ascent methods. *SIAM review*, 11(2):226–235, 1969.

[181] Philip Wolfe. Convergence conditions for ascent methods. ii: Some corrections. *SIAM review*, 13(2):185–188, 1971.

[182] Hakan Erdogan. *Statistical image reconstruction algorithms using paraboloidal surrogates for PET transmission scans*. PhD thesis, The University of Michigan, 1999.

[183] Kenneth Lange and Jeffrey A Fessler. Globally convergent algorithms for maximum a posteriori transmission tomography. *IEEE Transactions on Image Processing*, 4(10):1430–1438, 1995.

[184] Donghwan Kim, Debashish Pal, Jean-Baptiste Thibault, and Jeffrey A Fessler. Improved ordered subsets algorithm for 3D x-ray CT image reconstruction. *Proc. 2nd Intl. Mtg. on image formation in X-ray CT*, pages 378–81, 2012.

[185] Hakan Erdogan and Jeffrey A Fessler. Ordered subsets algorithms for transmission tomography. *Physics in Medicine & Biology*, 44(11):2835, 1999.

[186] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147, 2013.

[187] Yurii Nesterov. A method of solving a convex programming problem with convergence rate O (1/k2). In *Soviet Mathematics Doklady*, volume 27, pages 372–376, 1983.

[188] Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.

[189] Donghwan Kim, Sathish Ramani, and Jeffrey A Fessler. Combining ordered subsets and momentum for accelerated x-ray ct image reconstruction. *IEEE transactions on medical imaging*, 34(1):167–178, 2014.

[190] Adam S Wang, J Webster Stayman, Yoshito Otake, Sebastian Vogt, Gerhard Kleinszig, and Jeffrey H Siewerdsen. Accelerated statistical reconstruction for c-arm cone-beam ct using nesterov's method. *Medical physics*, 42(5):2699–2708, 2015.

[191] Naveen Chandra and David A Langan. Gemstone detector: dual energy imaging via fast kvp switching. In *Dual energy CT in clinical practice*, pages 35–41. Springer, 2011.

[192] Mitchell M Goodsitt, Emmanuel G Christodoulou, and Sandra C Larson. Accuracies of the synthesized monochromatic ct numbers and effective atomic numbers obtained with a rapid kvp switching dual energy ct scanner. *Medical physics*, 38(4):2222–2232, 2011.

[193] Willi A Kalender, WH Perman, JR Vetter, and Ernst Klotz. Evaluation of a prototype dual-energy computed tomographic apparatus. i. phantom studies. *Medical physics*, 13(3):334–339, 1986.

[194] Thomas G Flohr, Cynthia H McCollough, Herbert Bruder, Martin Petersilka, Klaus Gruber, Christoph Süβ, Michael Grasruck, Karl Stierstorfer, Bernhard Krauss, Rainer Raupach, et al. First performance evaluation of a dual-source ct (dsct) system. *European radiology*, 16(2):256–268, 2006.

[195] Martin Petersilka, Herbert Bruder, Bernhard Krauss, Karl Stierstorfer, and Thomas G Flohr. Technical principles of dual source ct. *European journal of radiology*, 68(3):362–368, 2008.

[196] André Euler, Anushri Parakh, Anna L Falkowski, Sebastian Manneck, David Dashti, Bernhard Krauss, Zsolt Szucs-Farkas, and Sebastian T Schindera. Initial results of a single-source dual-energy computed tomography technique using a split-filter: assessment of image quality, radiation dose, and accuracy of dual-energy applications in an in vitro and in vivo study. *Investigative radiology*, 51(8):491–498, 2016.

[197] Raz Carmi, Galit Naveh, and Ami Altman. Material separation with dual-layer ct. In *IEEE Nuclear Science Symposium Conference Record, 2005*, volume 4, pages 3–pp. IEEE, 2005.

[198] Gert Jan Pelgrim, Robbert W van Hamersvelt, Martin J Willemink, Bernhard T Schmidt, Thomas Flohr, Arnold Schilham, Julien Milles, Matthijs Oudkerk, Tim Leiner, and Rozemarijn Vliegenthart. Accuracy of iodine quantification using dual energy ct in latest generation dual source and dual layer ct. *European radiology*, 27(9):3904–3912, 2017.

[199] David L Ergun, Charles A Mistretta, David E Brown, RT Bystrianyk, WK Sze, F Kelcz, and DP Naidich. Single-exposure dual-energy computed radiography: improved detection and processing. *Radiology*, 174(1):243–249, 1990.

[200] Shengzhen Tao, Kishore Rajendran, Cynthia H McCollough, and Shuai Leng. Feasibility of multi-contrast imaging on dual-source photon counting detector (pcd) ct: An initial phantom study. *Medical Physics*, 2019.

[201] David P Cormode, Salim Si-Mohamed, Daniel Bar-Ness, Monica Sigovan, Pratap C Naha, Joelle Balegamire, Franck Lavenne, Philippe Coulon, Ewald Roessl, Matthias Bartels, et al. Multicolor spectral photon-counting computed tomography: in vivo dual contrast imaging with a high count rate scanner. *Scientific reports*, 7(1):4784, 2017.

[202] Salim Si-Mohamed, Valérie Tatard-Leitman, Alexis Laugerette, Monica Sigovan, Daniela Pfeiffer, Ernst J Rummeny, Philippe Coulon, Yoad Yagil, Philippe Douek, Loic Boussel, et al. Spectral photon-counting computed tomography (spcct): in-vivo single-acquisition multi-phase liver imaging with a dual contrast agent protocol. *Scientific reports*, 9(1):8458, 2019.

[203] LA Lehmann, RE Alvarez, Aetal Macovski, WR Brody, NJ Pelc, Stephen J Riederer, and AL Hall. Generalized image combinations in dual kvp digital radiography. *Medical physics*, 8(5):659–667, 1981.

[204] Jeffrey F Williamson, Sicong Li, Slobodan Devic, Bruce R Whiting, and Fritz A Lerma. On two-parameter models of photon cross sections: Application to dual-energy ct imaging. *Medical physics*, 33(11):4115–4129, 2006.

[205] Eric L Gingold and Bruce H Hasegawa. Systematic bias in basis material decomposition applied to quantitative dual-energy x-ray imaging. *Medical physics*, 19(1):25–33, 1992.

[206] Peter Hammersberg and Måns Mångård. Correction for beam hardening artefacts in computerised tomography. *Journal of X-ray Science and Technology*, 8(1):75–93, 1998.

[207] Jan Jakubek. Data processing and image reconstruction methods for pixel detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 576(1):223–234, 2007.

[208] Graham Davis, Nitin Jain, and James Elliott. A modelling approach to beam hardening correction. In *Developments in X-ray Tomography VI*, volume 7078, page 70781E. International Society for Optics and Photonics, 2008.

[209] Marc Kachelrieß, Katia Sourbelle, and Willi A Kalender. Empirical cupping correction: A first-order raw data precorrection for cone-beam computed tomography. *Medical physics*, 33(5):1269–1274, 2006.

[210] Gabor T Herman. Demonstration of beam hardening correction in computed tomography of the head. *Journal of computer assisted tomography*, 3(3):373–378, 1979.

[211] JM Meagher, CD Mote, and HB Skinner. Ct image correction for beam hardening using simulated projection data. *IEEE Transactions on Nuclear Science*, 37(4):1520–1524, 1990.

[212] SJ Kim, HK Lim, HY Lee, CG Choi, DH Lee, DC Suh, SM Kim, JK Kim, and B Krauss. Dual-energy ct in the evaluation of intracerebral hemorrhage of unknown origin: differentiation between tumor bleeding and pure hemorrhage. *American Journal of Neuroradiology*, 33(5):865–872, 2012.

[213] Rajiv Gupta, Catherine M Phan, Christianne Leidecker, Thomas J Brady, Joshua A Hirsch, Raul G Nogueira, and Albert J Yoo. Evaluation of dual-energy ct for differentiating intracerebral hemorrhage from iodinated contrast material staining. *Radiology*, 257(1):205–211, 2010.

[214] Michelle M McNamara, Mark D Little, Lauren F Alexander, L Van Carroll, T Mark Beasley, and Desiree E Morgan. Multireader evaluation of lesion conspicuity in small pancreatic adenocarcinomas: complimentary value of iodine material density and low kev simulated monoenergetic images using multiphasic rapid kvp-switching dual energy ct. *Abdominal imaging*, 40(5):1230–1240, 2015.

[215] Zilai Pan, Lifang Pang, Bei Ding, Chao Yan, Huan Zhang, Lianjun Du, Baisong Wang, Qi Song, Kemin Chen, and Fuhua Yan. Gastric cancer staging with dual energy spectral ct imaging. *PLoS One*, 8(2):e53651, 2013.

[216] Eun Jin Chae, Joon Beom Seo, Yu Mi Jang, Bernhard Krauss, Choong Wook Lee, Hyun Joo Lee, and Koun-Sik Song. Dual-energy ct for assessment of the severity of acute pulmonary embolism: pulmonary perfusion defect score compared with ct angiographic obstruction score and right ventricular/left ventricular diameter ratio. *American Journal of Roentgenology*, 194(3):604–610, 2010.

[217] Alexis Laugerette, Benedikt J Schwaiger, Kevin Brown, Lena C Frerking, Felix K Kopp, Kai Mei, Thorsten Sellerer, Jan Kirschke, Thomas Baum, Alexandra S Gersing, et al. Dxa-equivalent quantification of bone mineral density using dual-layer spectral ct scout scans. *European radiology*, pages 1–11, 2019.

[218] Bryan M Zweig, Meetkumar Sheth, Steve Simpson, and Mouaz H Al-Mallah. Association of abdominal aortic calcium with coronary artery calcium and obstructive coronary artery disease: a pilot study. *The International Journal of Cardiovascular Imaging*, 28(2):399–404, 2012.

[219] Matthias Eiber, Konstantin Holzapfel, Markus Frimberger, Michael Straub, Heike Schneider, Ernst J Rummeny, Martin Dobritz, and Armin Huber. Targeted dual-energy single-source ct for characterisation of urinary calculi: experimental and clinical experience. *European radiology*, 22(1):251–258, 2012.

[220] Eun Jin Chae, Jae-Woo Song, Joon Beom Seo, Bernhard Krauss, Yu Mi Jang, and Koun-Sik Song. Clinical utility of dual-energy ct in the evaluation of solitary pulmonary nodules: initial experience. *Radiology*, 249(2):671–681, 2008.

[221] Long-Jiang Zhang, Jin Peng, Sheng-Yong Wu, Z Jane Wang, Xin-Sheng Wu, Chang-Sheng Zhou, Xue-Man Ji, and Guang-Ming Lu. Liver virtual non-enhanced ct with dual-source, dual-energy ct: a preliminary study. *European radiology*, 20(9):2257–2264, 2010.

[222] Darin P Clark and Cristian T Badea. Spectral diffusion: an algorithm for robust material decomposition of spectral ct data. *Physics in medicine and biology*, 59(21):6445, 2014.

[223] Tianye Niu, Xue Dong, Michael Petrongolo, and Lei Zhu. Iterative image-domain decomposition for dual-energy ct. *Medical physics*, 41(4), 2014.

[224] Clemens Maaß, Matthias Baer, and Marc Kachelrieß. Image-based dual energy ct using optimized precorrection functions: A practical new approach of material decomposition in image domain. *Medical physics*, 36(8):3818–3829, 2009.

[225] Robert E Alvarez. Estimator for photon counting energy selective x-ray imaging with multibin pulse height analysis. *Medical physics*, 38(5):2324–2334, 2011.

[226] Joonki Noh, Jeffrey A Fessler, and Paul E Kinahan. Statistical sinogram restoration in dual-energy ct for pet attenuation correction. *IEEE transactions on medical imaging*, 28(11):1688–1702, 2009.

[227] E Roessl and C Herrmann. Cramér–rao lower bound of basis image noise in multiple-energy x-ray imaging. *Physics in medicine and biology*, 54(5):1307, 2009.

[228] Philip Stenner, Timo Berkus, and Marc Kachelriess. Empirical dual energy calibration (edec) for cone-beam computed tomography. *Medical physics*, 34(9):3630–3641, 2007.

[229] Willi A Kalender, Ernst Klotz, and Lena Kostaridou. An algorithm for noise suppression in dual energy ct material density images. *IEEE transactions on medical imaging*, 7(3):218–224, 1988.

[230] Michael Petrongolo and Lei Zhu. Noise suppression for dual-energy ct through entropy minimization. *IEEE transactions on medical imaging*, 34(11):2286–2297, 2015.

[231] Richard J Warp and James T Dobbins III. Quantitative evaluation of noise reduction strategies in dual-energy imaging. *Medical physics*, 30(2):190–198, 2003.

[232] Korbinian Mechlem, Sebastian Allner, Sebastian Ehn, Kai Mei, Eva Braig, Daniela Münzel, Franz Pfeiffer, and Peter B Noël. A post-processing algorithm for spectral ct material selective images using learned dictionaries. *Biomedical Physics & Engineering Express*, 3(2):025009, 2017.

[233] Shengzhen Tao, Kishore Rajendran, Cynthia H McCollough, and Shuai Leng. Material decomposition with prior knowledge aware iterative denoising (md-pkaid). *Physics in Medicine & Biology*, 63(19):195003, 2018.

[234] Tom Hohweiller, Nicolas Ducros, Françoise Peyrin, and Bruno Sixou. Spectral ct material decomposition in the presence of poisson noise: a kullback–leibler approach. *IRBM*, 38(4):214–218, 2017.

[235] Nicolas Ducros, Juan Felipe Perez-Juste Abascal, Bruno Sixou, Simon Rit, and Françoise Peyrin. Regularization of nonlinear decomposition of spectral x-ray projection images. *Medical physics*, 44(9):e174–e187, 2017.

[236] Korbinian Mechlem, Thorsten Sellerer, Sebastian Ehn, Daniela Münzel, Eva Braig, Julia Herzen, Peter B Noël, and Franz Pfeiffer. Spectral angiography material decomposition using an empirical forward model and a dictionary-based regularization. *IEEE transactions on medical imaging*, 37(10):2298–2309, 2018.

[237] Wonseok Huh and Jeffrey A Fessler. Iterative image reconstruction for dual-energy x-ray ct using regularized material sinogram estimates. In *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 1512–1515. IEEE, 2011.

[238] Carsten O Schirra, Ewald Roessl, Thomas Koehler, Bernhard Brendel, Axel Thran, Dipanjan Pan, Mark A Anastasio, and Roland Proksa. Statistical reconstruction of material decomposed data in spectral ct. *IEEE transactions on medical imaging*, 32(7):1249–1257, 2013.

[239] Hans Bornefalk. Xcom intrinsic dimensionality for low-z elements at diagnostic energies. *Medical physics*, 39(2):654–657, 2012.

[240] Robert E Alvarez. Dimensionality and noise in energy selective x-ray imaging. *Medical physics*, 40(11):111909, 2013.

[241] D Clark, SM Johnston, GA Johnson, and CT Badea. The effect of scatter correction on dual energy micro-ct. In *Proc. of SPIE Vol*, volume 8668, pages 86684T–1, 2013.

[242] Jens Wiegert, Klaus Jürgen Engel, and Christoph Herrmann. Impact of scattered radiation on spectral ct. In *Medical Imaging 2009: Physics of Medical Imaging*, volume 7258, page 72583X. International Society for Optics and Photonics, 2009.

[243] Klaus J Engel, Christian Bäumer, Jens Wiegert, and Günter Zeitler. Spectral analysis of scattered radiation in ct. In *Medical Imaging 2008: Physics of Medical Imaging*, volume 6913, page 69131R. International Society for Optics and Photonics, 2008.

[244] H Neale Cardinal and Aaron Fenster. An accurate method for direct dual-energy calibration and decomposition. *Medical physics*, 17(3):327–341, 1990.

[245] Okkyun Lee, Steffen Kappler, Christoph Polster, and Katsuyuki Taguchi. Estimation of basis line-integrals in a spectral distortion-modeled photon counting detector using low-order polynomial approximation of x-ray transmittance. *IEEE transactions on medical imaging*, 36(2):560–573, 2016.

[246] Dufan Wu, Li Zhang, Xiaohua Zhu, Xiaofei Xu, and Sen Wang. A weighted polynomial based material decomposition method for spectral x-ray ct imaging. *Physics in Medicine & Biology*, 61(10):3749, 2016.

[247] Robert Alvarez. Near optimal neural network estimator for spectral x-ray photon counting data with pileup. *arXiv preprint arXiv:1702.01006*, 2017.

[248] Mengheng Touch, Darin P Clark, William Barber, and Cristian T Badea. A neural network-based method for spectral distortion correction in photon counting x-ray ct. *Physics in Medicine & Biology*, 61(16):6132, 2016.

[249] S Ehn, T Sellerer, K Mechlem, A Fehringer, M Epple, J Herzen, F Pfeiffer, and PB Noël. Basis material decomposition in spectral ct using a semi-empirical, polychromatic adaption of the beer–lambert model. *Physics in Medicine and Biology*, 62:N1–N17, 2017.

[250] Korbinian Mechlem, Sebastian Ehn, Thorsten Sellerer, Eva Braig, Daniela Münzel, Franz Pfeiffer, and Peter B Noël. Joint statistical iterative material image reconstruction for spectral computed tomography using a semi-empirical forward model. *IEEE transactions on medical imaging*, 37(1):68–80, 2017.

[251] Robert E Alvarez. Efficient, non-iterative estimator for imaging contrast agents with spectral x-ray detectors. *IEEE transactions on medical imaging*, 35(4):1138–1146, 2015.

[252] Lifeng Yu, Jodie A Christner, Shuai Leng, Jia Wang, Joel G Fletcher, and Cynthia H McCollough. Virtual monochromatic imaging in dual-source dual-energy ct: radiation dose and image quality. *Medical physics*, 38(12):6371–6379, 2011.

[253] Cynthia H McCollough, Shuai Leng, Lifeng Yu, and Joel G Fletcher. Dual-and multi-energy ct: principles, technical approaches, and clinical applications. *Radiology*, 276(3):637–653, 2015.

[254] Moritz H Albrecht, Thomas J Vogl, Simon S Martin, John W Nance, Taylor M Duguay, Julian L Wichmann, Carlo N De Cecco, Akos Varga-Szemes, Marly van Assen, Christian Tesche, et al. Review of clinical applications for virtual monoenergetic dual-energy ct. *Radiology*, 293(2):260–271, 2019.

[255] Lifeng Yu, Shuai Leng, and Cynthia H McCollough. Dual-energy ct–based monochromatic imaging. *American journal of Roentgenology*, 199(5_supplement):S9–S15, 2012.

[256] Paul Apfaltrer, Sonja Sudarski, David Schneider, John W Nance Jr, Holger Haubenreisser, Christian Fink, Stefan O Schoenberg, and Thomas Henzler. Value of monoenergetic low-kv dual energy ct datasets for improved image quality of ct pulmonary angiography. *European journal of radiology*, 83(2):322–328, 2014.

[257] Eric Pessis, Raphaël Campagna, Jean-Michel Sverzut, Fabienne Bach, Mathieu Rodallec, Henri Guerini, Antoine Feydy, and Jean-Luc Drapé. Virtual monochromatic spectral imaging with fast kilovoltage switching: reduction of metal artifacts at ct. *Radiographics*, 33(2):573–583, 2013.

[258] Julia Dangelmaier, Benedikt J Schwaiger, Alexandra S Gersing, Felix F Kopp, Andreas Sauter, Martin Renz, Isabelle Riederer, Rickmer Braren, Daniela Pfeiffer, Alexander Fingerle, et al. Dual layer computed tomography: reduction of metal artefacts from posterior spinal fusion using virtual monoenergetic imaging. *European journal of radiology*, 105:195–203, 2018.

[259] Paul Stolzmann, Sebastian Winklhofer, Nicole Schwendener, Hatem Alkadhi, Michael J Thali, and Thomas D Ruder. Monoenergetic computed tomography reconstructions reduce beam hardening artifacts from dental restorations. *Forensic science, medicine, and pathology*, 9(3):327–332, 2013.

[260] Thorsten Sellerer, Sebastian Ehn, Korbinian Mechlem, Manuela Duda, Michael Epple, Peter B Noël, and Franz Pfeiffer. Quantitative dual-energy micro-ct with a photon-counting detector for material science and non-destructive testing. *PloS one*, 14(7), 2019.

[261] Anne Bonnin, Philippe Duvauchelle, Valérie Kaftandjian, and Pascal Ponard. Concept of effective atomic number and effective mass density in dual-energy x-ray computed tomography. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 318:223–231, 2014.

[262] Otto Glasser, Edith H Quimby, Lauriston S Taylor, JL Weatherwax, and Russell H Morgan. Physical foundations of radiology. *Academic Medicine*, 37(3):253, 1962.

[263] Masatoshi Saito and Shota Sagara. A simple formulation for deriving effective atomic numbers via electron density calibration from dual-energy ct data in the human body. *Medical physics*, 44(6):2293–2303, 2017.

[264] Nora Hünemohr, Bernhard Krauss, Christoph Tremmel, Benjamin Ackermann, Oliver Jäkel, and Steffen Greilich. Experimental verification of ion stopping power prediction from dual energy ct data in tissue surrogates. *Physics in Medicine & Biology*, 59(1):83, 2013.

[265] Esther Bär, Arthur Lalonde, Gary Royle, Hsiao-Ming Lu, and Hugo Bouchard. The potential of dual-energy ct to reduce proton beam range uncertainties. *Medical physics*, 44(6):2332–2344, 2017.

[266] Etienne Danse et al. Spectral ct detection of entrapped gallstone based on z-effective map. *Journal of the Belgian Society of Radiology*, 102(1), 2018.

[267] Marco Endrizzi. X-ray phase-contrast imaging. *Nuclear instruments and methods in physics research section A: Accelerators, spectrometers, detectors and associated equipment*, 878:88–98, 2018.

[268] SW Wilkins, Ya I Nesterets, TE Gureyev, SC Mayo, A Pogany, and AW Stevenson. On the evolution and relative merits of hard x-ray phase-contrast imaging methods. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2010):20130021, 2014.

[269] Atsushi Momose. Recent advances in x-ray phase imaging. *Japanese Journal of Applied Physics*, 44(9R):6355, 2005.

[270] Atsushi Momose, Shinya Kawamoto, Ichiro Koyama, Yoshitaka Hamaishi, Kengo Takai, and Yoshio Suzuki. Demonstration of x-ray talbot interferometry. *Japanese journal of applied physics*, 42(7B):L866, 2003.

[271] Timm Weitkamp, Ana Diaz, Christian David, Franz Pfeiffer, Marco Stampanoni, Peter Cloetens, and Eric Ziegler. X-ray phase imaging with a grating interferometer. *Optics express*, 13(16):6296–6304, 2005.

[272] Thomas J Suleski. Generation of lohmann images from binary-phase talbot array illuminators. *Applied optics*, 36(20):4686–4691, 1997.

[273] Martin Bech. *X-ray imaging with a grating interferometer*. PhD thesis, University of Copenhagen, 2009.

[274] Ernst Lau. Beugungserscheinungen an doppelrastern. *Annalen der Physik*, 437(7-8):417–423, 1948.

[275] Tilman Donath, Michael Chabior, Franz Pfeiffer, Oliver Bunk, Elena Reznikova, Juergen Mohr, Eckhard Hempel, Stefan Popescu, Martin Hoheisel, Manfred Schuster, et al. Inverse geometry for grating-based x-ray phase-contrast imaging. *Journal of Applied Physics*, 106(5):054703, 2009.

[276] Michael Chabior. *Contributions to the characterization of grating-based x-ray phase-contrast imaging*. PhD thesis, Technische Universität Dresden, 2011.

[277] William H Carter. Coherence theory, 1993.

[278] Bernhard Brendel, Maximilian Teuffenbach, Peter B Noël, Franz Pfeiffer, and Thomas Koehler. Penalized maximum likelihood reconstruction for x-ray differential phase-contrast tomography. *Medical physics*, 43(1):188–194, 2016.

[279] Thomas Weber, Peter Bartl, Florian Bayer, Jürgen Durst, Wilhelm Haas, Thilo Michel, Andre Ritter, and Gisela Anton. Noise in x-ray grating-based phase-contrast imaging. *Medical physics*, 38(7):4133–4140, 2011.

[280] Susanne Grandl, Kai Scherer, Anikó Sztrókay-Gaul, Lorenz Birnbacher, Konstantin Willer, Michael Chabior, Julia Herzen, Doris Mayr, Sigrid D Auweter, Franz Pfeiffer, et al. Improved visualization of breast cancer features in multifocal carcinoma using phase-contrast and dark-field mammography: an ex vivo study. *European radiology*, 25(12):3659–3668, 2015.

[281] E Eggl, S Schleede, Martin Bech, K Achterhold, S Grandl, A Sztrókay, K Hellerhoff, D Mayr, R Loewen, RD Ruth, et al. X-ray phase-contrast tomosynthesis of a human ex vivo breast slice with an inverse compton x-ray source. *EPL (Europhysics Letters)*, 116(6):68003, 2017.

[282] Thomas Koehler, Heiner Daerr, Gerhard Martens, Norbert Kuhn, Stefan Löscher, Udo van Stevendaal, and Ewald Roessl. Slit-scanning differential x-ray phase-contrast mammography: Proof-of-concept experimental studies. *Medical physics*, 42(4):1959–1965, 2015.

[283] Christopher Habbel, Holger Hetterich, Marian Willner, Julia Herzen, Kristin Steigerwald, Sigrid Auweter, Ulrich Schüller, Jörg Hausleiter, Steffen Massberg, Maximilian Reiser, et al. Ex vivo assessment of coronary atherosclerotic plaque by grating-based phase-contrast computed tomography: correlation with optical coherence tomography. *Investigative radiology*, 52(4):223–231, 2017.

[284] Holger Hetterich, Marian Willner, Sandra Fill, Julia Herzen, Fabian Bamberg, Alexander Hipp, Ulrich Schüller, Silvia Adam-Neumair, Stefan Wirth, Maximilian Reiser, et al. Phase-contrast ct: qualitative and quantitative evaluation of atherosclerotic carotid artery plaque. *Radiology*, 271(3):870–878, 2014.

[285] Atsushi Momose, Wataru Yashiro, Kazuhiro Kido, Junko Kiyohara, Chiho Makifuchi, Tsukasa Ito, Sumiya Nagatsuka, Chika Honda, Daiji Noda, Tadashi Hattori, et al. X-ray phase imaging: from synchrotron to hospital. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2010):20130023, 2014.

[286] Julia Herzen, Dimitrios C Karampinos, Peter Foehr, Lorenz Birnbacher, Manuel Viermetz, Rainer Burgkart, Thomas Baum, Fabian Lohoefer, Moritz Wildgruber, Franz Schilling, et al. 3d grating-based x-ray phase-contrast computed tomography for high-resolution quantitative assessment of cartilage: An experimental feasibility study with 3t mri, 7t mri and biomechanical correlation. *PloS one*, 14(2):e0212106, 2019.

[287] Irene Zanette, Timm Weitkamp, Geraldine Le Duc, and Franz Pfeiffer. X-ray grating-based phase tomography for 3d histology. *RSC Advances*, 3(43):19816–19819, 2013.

[288] Joan Vila-Comamala, Carolina Arboleda, Lucia Romano, Willy Kuo, Kristina Lang, Konstantins Jefimovs, Zhentian Wang, Gad Singer, David Vine, Wenbing Yun, et al. Development of laboratory grating-based x-ray phase contrast microtomography for improved pathology. *Microscopy and Microanalysis*, 24(S2):192–193, 2018.

[289] Marian Willner, Martin Bech, Julia Herzen, Irene Zanette, Dieter Hahn, Johannes Kenntner, Juergen Mohr, Alexander Rack, Timm Weitkamp, and Franz Pfeiffer. Quantitative x-ray phase-contrast computed tomography at 82 kev. *Optics express*, 21(4):4155–4166, 2013.

[290] Zhihua Qi, Joseph Zambelli, Nicholas Bevins, and Guang-Hong Chen. Quantitative imaging of electron density and effective atomic number using phase contrast CT. *Physics in Medicine & Biology*, 55(9):2669, 2010.

[291] Lorenz Birnbacher, Marian Willner, Mathias Marschner, Daniela Pfeiffer, Franz Pfeiffer, and Julia Herzen. Accurate effective atomic number determination with polychromatic grating-based phase-contrast computed tomography. *Optics express*, 26(12):15153–15166, 2018.

[292] W Yashiro, Y Terui, K Kawabata, and A Momose. On the origin of visibility contrast in x-ray talbot interferometry. *Optics express*, 18(16):16890–16901, 2010.

[293] Susanna K Lynch, Vinay Pai, Julie Auxier, Ashley F Stein, Eric E Bennett, Camille K Kemble, Xianghui Xiao, Wah-Keat Lee, Nicole Y Morgan, and Han Harold Wen. Interpretation of dark-field contrast and particle-size selectivity in grating interferometers. *Applied optics*, 50(22):4310–4319, 2011.

[294] Markus Strobl. General solution for quantitative dark-field contrast imaging with grating interferometers. *Scientific reports*, 4:7243, 2014.

[295] F Prade, A Yaroshenko, J Herzen, and F Pfeiffer. Short-range order in mesoscale systems probed by x-ray grating interferometry. *EPL (Europhysics Letters)*, 112(6):68002, 2016.

[296] Robert Andersson, Léon F Van Heijkamp, Ignatz M De Schepper, and Wim G Bouwman. Analysis of spin-echo small-angle neutron scattering measurements. *Journal of Applied Crystallography*, 41(5):868–885, 2008.

[297] Vincent Revol, Bernhard Plank, Rolf Kaufmann, Johann Kastner, Christian Kottler, and Antonia Neels. Laminate fibre structure characterisation of carbon fibre-reinforced polymers by x-ray scatter dark field imaging with a grating interferometer. *NDT & E International*, 58:64–71, 2013.

[298] Andreas Malecki, Guillaume Potdevin, Thomas Biernath, Elena Eggl, Konstantin Willer, Tobias Lasser, Jens Maisenbacher, Jens Gibmeier, Alexander Wanner, and Franz Pfeiffer. X-ray tensor tomography. *EPL (Europhysics Letters)*, 105(3):38002, 2014.

[299] Torsten Lauridsen, Marian Willner, Martin Bech, Franz Pfeiffer, and Robert Feidenhans. Detection of sub-pixel fractures in x-ray dark-field tomography. *Applied Physics A*, 121(3):1243–1250, 2015.

[300] Friedrich Prade, Kai Fischer, Detlef Heinz, Pascal Meyer, Jürgen Mohr, and Franz Pfeiffer. Time resolved x-ray dark-field tomography revealing water transport in a fresh cement sample. *Scientific reports*, 6:29108, 2016.

[301] F Yang, F Prade, M Griffa, I Jerjen, C Di Bella, J Herzen, A Sarapata, F Pfeiffer, and Pietro Lura. Dark-field x-ray imaging of unsaturated water transport in porous materials. *Applied Physics Letters*, 105(15):154105, 2014.

[302] Eva-Maria Braig, Lorenz Birnbacher, Florian Schaff, Lukas Gromann, Alexander Fingerle, Julia Herzen, Ernst Rummeny, Peter Noël, Franz Pfeiffer, and Daniela Muenzel. Simultaneous wood and metal particle detection on dark-field radiography. *European radiology experimental*, 2(1):1, 2018.

[303] Kai Scherer, Eva Braig, Sebastian Ehn, Jonathan Schock, Johannes Wolf, Lorenz Birnbacher, Michael Chabior, Julia Herzen, Doris Mayr, Susanne Grandl, et al. Improved diagnostics by assessing the micromorphology of breast calcifications via x-ray dark-field radiography. *Scientific reports*, 6:36991, 2016.

[304] Kai Scherer, Eva Braig, Konstantin Willer, Marian Willner, Alexander A Fingerle, Michael Chabior, Julia Herzen, Matthias Eiber, Bernhard Haller, Michael Straub, et al. Non-invasive differentiation of kidney stone types using x-ray dark-field radiography. *Scientific reports*, 5:9527, 2015.

[305] Katharina Hellbach, Andre Yaroshenko, Felix G Meinel, Ali Ö Yildirim, Thomas M Conlon, Martin Bech, Mark Mueller, Astrid Velroyen, Mike Notohamiprodjo, Fabian Bamberg, et al. In vivo dark-field radiography for early diagnosis and staging of pulmonary emphysema. *Investigative radiology*, 50(7):430–435, 2015.

[306] Andre Yaroshenko, Katharina Hellbach, Ali Önder Yildirim, Thomas M Conlon, Isis Enlil Fernandez, Martin Bech, Astrid Velroyen, Felix G Meinel, Sigrid Auweter, Maximilian Reiser, et al. Improved in vivo assessment of pulmonary fibrosis in mice using x-ray dark-field radiography. *Scientific Reports*, 5:17492, 2015.

[307] Konstantin Willer, Alexander A Fingerle, Lukas B Gromann, Fabio De Marco, Julia Herzen, Klaus Achterhold, Bernhard Gleich, Daniela Muenzel, Kai Scherer, Martin Renz, et al. X-ray dark-field imaging of the human lung—a feasibility study on a deceased body. *PloS one*, 13(9):e0204565, 2018.

[308] Griffin Rodgers, Georg Schulz, Hans Deyhle, Shashidhara Marathe, Christos Bikis, Timm Weitkamp, and Bert Müller. A quantitative correction for phase wrapping artifacts in hard x-ray grating interferometry. *Applied Physics Letters*, 113(9):093702, 2018.

[309] Iwan Jerjen, Vincent Revol, Philipp Schuetz, Christian Kottler, Rolf Kaufmann, Thomas Luethi, Konstantins Jefimovs, Claus Urban, and Urs Sennhauser. Reduction of phase artifacts in differential phase contrast computed tomography. *Optics express*, 19(14):13604–13611, 2011.

[310] Aimin Yan, Xizeng Wu, and Hong Liu. Polychromatic x-ray effects on fringe phase shifts in grating interferometry. *Optics express*, 25(6):6053–6068, 2017.

[311] Georg Pelzer, Gisela Anton, Florian Horn, Jens Rieger, André Ritter, Johannes Wandner, Thomas Weber, and Thilo Michel. A beam hardening and dispersion correction for x-ray dark-field radiography. *Medical physics*, 43(6Part1):2774–2779, 2016.

[312] Nicholas Bevins, Ke Li, Joseph Zambelli, and Guang-Hong Chen. Type ii beam hardening artifacts in phase contrast imaging. In *Medical Imaging 2013: Physics of Medical Imaging*, volume 8668, page 866816. International Society for Optics and Photonics, 2013.

[313] A Sossin, V Rebuffel, J Tabary, JM Létang, N Freud, and L Verger. Influence of scattering on material quantification using multi-energy x-ray imaging. In *2014 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pages 1–5. IEEE, 2014.

[314] Wojciech Zbijewski and Freek J Beekman. Efficient monte carlo based scatter artifact reduction in cone-beam micro-ct. *IEEE transactions on medical imaging*, 25(7):817–827, 2006.

[315] G Poludniowski, PM Evans, VN Hansen, and S Webb. An efficient monte carlo-based algorithm for scatter correction in kev cone-beam ct. *Physics in Medicine & Biology*, 54(12):3847, 2009.

[316] JPea Schlomka, E Roessl, R Dorscheid, S Dill, G Martens, T Istel, C Bäumer, C Herrmann, R Steadman, G Zeitler, et al. Experimental feasibility of multi-energy photon-counting k-edge imaging in pre-clinical computed tomography. *Physics in medicine and biology*, 53(15):4031, 2008.

[317] J Xu, Wojciech Zbijewski, G Gang, Joseph Webster Stayman, Katsuyuki Taguchi, M Lundqvist, E Fredenberg, JA Carrino, and JH Siewerdsen. Cascaded systems analysis of photon counting detectors. *Medical physics*, 41(10):101907, 2014.

[318] Xu Ji, Ran Zhang, Yongshuai Ge, Guang-Hong Chen, and Ke Li. Signal and noise characteristics of a cdte-based photon counting detector: cascaded systems analysis and experimental studies. In *Medical Imaging 2017: Physics of Medical Imaging*, volume 10132, page 1013207. International Society for Optics and Photonics, 2017.

[319] Ewald Roessl, Heiner Daerr, Klaus Jürgen Engel, Axel Thran, Carsten Schirra, and Roland Proksa. Combined effects of pulse pile-up and energy response in energy-resolved, photon-counting computed tomography. In *2011 IEEE Nuclear Science Symposium Conference Record*, pages 2309–2313. IEEE, 2011.

[320] Jari Kaipio and Erkki Somersalo. Statistical inverse problems: discretization, model reduction and inverse crimes. *Journal of computational and applied mathematics*, 198(2):493–504, 2007.

[321] Johannes Wolf, Andreas Malecki, Jonathan Sperl, Michael Chabior, Markus Schüttler, Dirk Bequé, Cristina Cozzini, and Franz Pfeiffer. Fast one-dimensional wave-front propagation for x-ray differential phase-contrast imaging. *Biomedical optics express*, 5(10):3739–3747, 2014.

[322] A Malecki, G Potdevin, and F Pfeiffer. Quantitative wave-optical numerical analysis of the dark-field signal in grating-based x-ray interferometry. *EPL (Europhysics Letters)*, 99(4):48001, 2012.

[323] Tonghe Wang and Lei Zhu. Dual energy ct with one full scan and a second sparse-view scan using structure preserving iterative reconstruction (spir). *Physics in Medicine and Biology*, 61(18):6684, 2016.

[324] Zhicong Yu, Shuai Leng, Zhoubo Li, and Cynthia H McCollough. Spectral prior image constrained compressed sensing (spectral piccs) for photon-counting computed tomography. *Physics in Medicine and Biology*, 61(18):6707, 2016.

[325] Miaoshi Wang, Yanbo Zhang, Rui Liu, Shuxu Guo, and Hengyong Yu. An adaptive reconstruction algorithm for spectral ct regularized by a reference image. *Physics in Medicine and Biology*, 61(24):8699, 2016.

[326] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006.

[327] Matan Protter and Michael Elad. Image sequence denoising via sparse and redundant representations. *IEEE Transactions on Image Processing*, 18(1):27–35, 2009.

[328] Saiprasad Ravishankar and Yoram Bresler. Mr image reconstruction from highly undersampled k-space data by dictionary learning. *IEEE transactions on medical imaging*, 30(5):1028–1041, 2011.

[329] Yang Chen, Luyao Shi, Qianjing Feng, Jian Yang, Huazhong Shu, Limin Luo, Jean-Louis Coatrieux, and Wufan Chen. Artifact suppressed dictionary learning for low-dose ct image processing. *IEEE transactions on medical imaging*, 33(12):2271–2292, 2014.

[330] Yang Lu, Jun Zhao, and Ge Wang. Few-view image reconstruction with dual dictionaries. *Physics in Medicine & Biology*, 57(1):173, 2011.

[331] Qiong Xu, Hengyong Yu, Xuanqin Mou, Lei Zhang, Jiang Hsieh, and Ge Wang. Low-dose x-ray ct reconstruction via dictionary learning. *IEEE transactions on medical imaging*, 31(9):1682–1697, 2012.

[332] Yanbo Zhang, Xuanqin Mou, Ge Wang, and Hengyong Yu. Tensor-based dictionary learning for spectral ct reconstruction. *IEEE Transactions on Medical Imaging*, 2016.

[333] Thomas Koehler and Ewald Roessl. Simultaneous de-noising in phase contrast tomography. In *International Workshop on X-ray and Neutron Phase Imaging with Gratings*, volume 1466, pages 78–83. AIP Publishing, 2012.

[334] Kai Mei, Alexander Valentinitsch, Felix Kopp, Andreas Fehringer, Franz Pfeiffer, Ernst Rummeny, Jan Bauer, and Peter Noël. Iterative ct image reconstruction using 3d dictionary learning. *Proceedings of Fully 3D Image Reconstruction in Radiology and Nuclear Medicine*, pages 130–133, 2015.

[335] Yagyensh Chandra Pati, Ramin Rezaiifar, and Perinkulam Sambamurthy Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proceedings of 27th Asilomar conference on signals, systems and computers*, pages 40–44. IEEE, 1993.

[336] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11(Jan):19–60, 2010.

[337] H. Bruder, T. Flohr, T. Fuchs, T. von der Haar, M. Karolczak, M. Kachelrieß, G. Lauritsch, S. Schaller, U. Seibert, K. Sourbelle, K. Stierhofer, O. Watzke, and P. Stenner. Forbild phantoms. `http://www.imp.uni-erlangen.de/phantoms/`. Accessed: 2016-08-17.

[338] Korbinian Mechlem, Sebastian Allner, Kai Mei, Franz Pfeiffer, and Peter B Noël. Dictionary-based image denoising for dual energy computed tomography. In *SPIE Medical Imaging*, pages 97830E–97830E. International Society for Optics and Photonics, 2016.

[339] Ken Ohara, Charles E Metz, Maryellen Lissak Giger, et al. Investigation of basic imaging properties in digital radiography. 13. effect of simple structured noise on the detectability of simulated stenotic lesions. *Medical physics*, 16(1):14–21, 1989.

[340] Jesse Tanguay, Ho Kyung Kim, and Ian A Cunningham. A theoretical comparison of x-ray angiographic image quality using energy-dependent and conventional subtraction methods. *Medical physics*, 39(1):132–142, 2012.

[341] RK Panta, ST Bell, JL Healy, R Aamir, CJ Bateman, M Moghiseh, APH Butler, and NG Anderson. Element-specific spectral imaging of multiple contrast agents: a phantom study. *Journal of Instrumentation*, 13(02):T02001, 2018.

[342] Rolf Symons, Bernhard Krauss, Pooyan Sahbaee, Tyler E Cork, Manu N Lakshmanan, David A Bluemke, and Amir Pourmorteza. Photon-counting ct for simultaneous imaging of multiple contrast agents in the abdomen: an in vivo study. *Medical physics*, 2017.

[343] Rolf Symons, Tyler E Cork, Manu N Lakshmanan, Robert Evers, Cynthia Davies-Venn, Kelly A Rice, Marvin L Thomas, Chia-Ying Liu, Steffen Kappler, Stefan Ulzheimer, Veit Sandfort, David A. Bluemke, and Amir Pourmorteza. Dual-contrast agent photon-counting computed tomography of the heart: initial experience. *The international journal of cardiovascular imaging*, 33(8):1253–1261, 2017.

[344] Daniela Muenzel, Heiner Daerr, Roland Proksa, Alexander A Fingerle, Felix K Kopp, Philippe Douek, Julia Herzen, Franz Pfeiffer, Ernst J Rummeny, and Peter B Noël. Simultaneous dual-contrast multi-phase liver imaging using spectral photon-counting computed tomography: a proof-of-concept study. *European Radiology Experimental*, 1(1):25, 2017.

[345] Daniela Muenzel, Daniel Bar-Ness, Ewald Roessl, Ira Blevis, Matthias Bartels, Alexander A Fingerle, Stefan Ruschke, Philippe Coulon, Heiner Daerr, Felix K Kopp, et al. Spectral photon-counting ct: Initial experience with dual–contrast agent k-edge colonography. *Radiology*, 283(3):723–728, 2016.

[346] Khurram Nasir, Juan J Rivera, Hyuk-Jae Chang, Matthew J Budoff, and Roger S Blumenthal. Calcified versus noncalcified atherosclerosis: Implications for evaluating cardiovascular risk. *Current Cardiovascular Risk Reports*, 3(2):150–155, mar 2009.

[347] Michael Petrongolo, Xue Dong, and Lei Zhu. A general framework of noise suppression in material decomposition for dual-energy ct. *Medical physics*, 42(8):4848–4862, 2015.

[348] B Brendel, F Bergner, K Brown, and T Koehler. Penalized likelihood decomposition for dual layer spectral ct. In *Proceeding of the fourth International Meeting in X-ray Computed Tomography*, pages 41–44, 2016.

[349] Cynthia H McCollough, Michael S Van Lysel, Walter W Peppler, and Charles A Mistretta. A correlated noise reduction algorithm for dual-energy digital subtraction angiography. *Medical physics*, 16(6):873–880, 1989.

[350] H Neale Cardinal and Aaron Fenster. Analytic approximation of the log-signal and log-variance functions of x-ray imaging systems, with application to dual-energy imaging. *Medical physics*, 18(5):867–879, 1991.

[351] David S Rigie and Patrick J La Rivière. Joint reconstruction of multi-channel, spectral ct data via constrained total nuclear variation minimization. *Physics in medicine and biology*, 60(5):1741, 2015.

[352] Yong Long and Jeffrey A Fessler. Multi-material decomposition using statistical image reconstruction for spectral ct. *IEEE transactions on medical imaging*, 33(8):1614–1626, 2014.

[353] Korbinian Mechlem, Sebastian Ehn, Thorsten Sellerer, Eva Braig, Daniela Münzel, Franz Pfeiffer, and Peter B Noël. Joint statistical iterative material image reconstruction for spectral computed tomography using a semi-empirical forward model. *IEEE transactions on medical imaging*, 37(1):68–80, 2018.

[354] Sara Soltani, Martin S Andersen, and Per Christian Hansen. Tomographic image reconstruction using training images. *Journal of Computational and Applied Mathematics*, 313:243–258, 2017.

[355] Cheng Zhang, Tao Zhang, Ming Li, Chengtao Peng, Zhaobang Liu, and Jian Zheng. Low-dose ct reconstruction via l1 dictionary learning regularization using iteratively reweighted least-squares. *Biomedical engineering online*, 15(1):66, 2016.

[356] Joel A Tropp, Anna C Gilbert, and Martin J Strauss. Algorithms for simultaneous sparse approximation. part i: Greedy pursuit. *Signal Processing*, 86(3):572–588, 2006.

[357] Jie Chen and Xiaoming Huo. Theoretical results on sparse representations of multiple-measurement vectors. *IEEE Transactions on Signal Processing*, 54(12):4634–4643, 2006.

[358] NV Aravind, K Abhinandan, Vineeth V Acharya, and S Sumam David. Comparison of omp and somp in the reconstruction of compressively sensed hyperspectral images. In *Communications and Signal Processing (ICCSP), 2011 International Conference on*, pages 188–192. IEEE, 2011.

[359] J Mairal. Spams - sparse modeling software. `http://spams-devel.gforge.inria.fr/`. Accessed: 2018-04-24.

[360] Robert E Alvarez. Near optimal energy selective x-ray imaging system performance with simple detectors. *Medical physics*, 37(2):822–841, 2010.

[361] Karl Schmid, William O. McSharry, Cornelis H. Pameijer, and J.Paul Binette. Chemical and physicochemical studies on the mineral deposits of the human atherosclerotic aorta. *Atherosclerosis*, 37(2):199–210, oct 1980.

[362] Carlo Di Mario and Nilesh Sutaria. Coronary angiography in the angioplasty era: projections with a meaning. *Heart*, 91(7):968–976, 2005.

[363] John M Boone. Parametrized x-ray absorption in diagnostic radiology from monte carlo calculations: Implications for x-ray detector design. *Medical physics*, 19(6):1467–1473, 1992.

[364] J. T. Dodge, B. G. Brown, E. L. Bolson, and H. T. Dodge. Lumen diameter of normal human coronary arteries. Influence of age, sex, anatomic variation, and left ventricular hypertrophy or dilation. *Circulation*, 86(1):232–246, jul 1992.

[365] Keika Ose, Tomoki Doue, Kan Zen, Mitsuyoshi Hadase, Takahisa Sawada, Akihiro Azuma, and Hiroaki Matsubara. Gadolinium'as an alternative to iodinated contrast media for x-ray angiography in patients with severe allergy. *Circulation Journal*, 69(4):507–509, 2005.

[366] Elizabeth Juneman, Layth Saleh, Hoang Thai, Steven Goldman, and Mohammad Reza Movahed. Successful coronary angiography with adequate image acquisition using a combination of gadolinium and a power injector in a patient with severe iodine contrast allergy. *Experimental & Clinical Cardiology*, 17(1):17, 2012.

[367] FD Hammer, PP Goffette, Jacques Malaise, and P Mathurin. Gadolinium dimeglumine: an alternative contrast agent for digital subtraction angiography. *European radiology*, 9(1):128–136, 1999.

[368] Christiane M Erley, Birgit D Bader, Elke D Berger, Nurdan Tuncel, Sabine Winkler, Gunnar Tepe, Teut Risler, and Stephan Duda. Gadolinium-based contrast media compared with iodinated media for digital subtraction angiography in azotaemic patients. *Nephrology Dialysis Transplantation*, 19(10):2526–2531, 2004.

[369] David J Spinosa, John A Kaufmann, and Gary D Hartwell. Gadolinium chelates in angiography and interventional radiology: a useful alternative to iodinated contrast media for angiography. *Radiology*, 223(2):319–325, 2002.

[370] M Giozzet, E Cavagna, M Dea De, G Tarroni, D Casol, L Silvestro De, C Tessarin, and E Paoli Vitali De. Gadolinium for dsa in two patients with azotemia: Images of suitable quality and risk of acute renal failure. *Giornale italiano di nefrologia: organo ufficiale della Societa italiana di nefrologia*, 20(3):298–301, 2003.

[371] Dotarem prescribing information. `http://www.guerbet-us.com/fileadmin/user_upload/usa_home/products/2017_0825_Dotarem_PI_FINAL_FROM_FDA.pdf`. Accessed: 2018-04-10.

[372] Gregory J Nadolski and S William Stavropoulos. Contrast alternatives for iodinated contrast allergy and renal dysfunction: options and limitations. *Journal of vascular surgery*, 57(2):593–598, 2013.

[373] Paul F FitzGerald, Matthew D Butts, Jeannette C Roberts, Robert E Colborn, Andrew S Torres, Brian D Lee, Benjamin M Yeh, and Peter J Bonitatibus Jr. A proposed ct contrast agent using carboxybetaine zwitterionic tantalum oxide nanoparticles: Imaging, biological, and physicochemical performance. *Investigative radiology*, 51(12):786, 2016.

[374] Mai Quyen Pham, Nicolas Ducros, and Barbara Nicolas. A bvmf-b algorithm for nonconvex nonlinear regularized decomposition of spectral x-ray projection images. In *Medical Imaging 2017: Physics of Medical Imaging*, volume 10132, page 101323K. International Society for Optics and Photonics, 2017.

[375] Matthew Tivnan, Wenying Wang, Steven Tilley II, Jeffrey H Siewerdsen, and J Webster Stayman. Optimized spatial-spectral ct for multi-material decomposition. In *15th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, volume 11072, page 1107211. International Society for Optics and Photonics, 2019.

[376] Steven Wayne Tilley II, Wojciech Zbijewski, and Joseph Webster Stayman. Model-based material decomposition with a penalized nonlinear least-squares ct reconstruction algorithm. *Physics in medicine and biology*, 2018.

[377] Thomas Weidinger, Thorsten M Buzug, Thomas Flohr, Steffen Kappler, and Karl Stierstorfer. Polychromatic iterative statistical material image reconstruction for photon-counting computed tomography. *International journal of biomedical imaging*, 2016, 2016.

[378] Cyril Mory, Bruno Sixou, Salim Si-Mohamed, Loic Boussel, and Simon Rit. Comparison of five one-step reconstruction algorithms for spectral ct. *Physics in Medicine & Biology*, 63(23):235001, 2018.

[379] Alex Sawatzky, Qiaofeng Xu, Carsten O Schirra, and Mark A Anastasio. Proximal admm for multi-channel image reconstruction in spectral x-ray ct. *IEEE transactions on medical imaging*, 33(8):1657–1668, 2014.

[380] C Cai, Thomas Rodet, Samuel Legoupil, and Ali Mohammad-Djafari. A full-spectral bayesian reconstruction approach based on the material decomposition model applied in dual-energy computed tomography. *Medical physics*, 40(11):111916, 2013.

[381] Rina Foygel Barber, Emil Y Sidky, Taly Gilat Schmidt, and Xiaochuan Pan. An algorithm for constrained one-step inversion of spectral ct data. *Physics in medicine and biology*, 61(10):3784, 2016.

[382] E Van de Casteele, D Van Dyck, J Sijbers, and E Raman. An energy-based beam hardening model in tomography. *Physics in Medicine and Biology*, 47(23):4181, 2002.

[383] J Alles and RF Mudde. Beam hardening: Analytical considerations of the effective attenuation coefficient of x-ray tomography. *Medical physics*, 34(7):2882–2889, 2007.

[384] Andrei A Istratov and Oleg F Vyvenko. Exponential analysis in physical phenomena. *Review of Scientific Instruments*, 70(2):1233–1257, 1999.

[385] Donghwan Kim, Sathish Ramani, and Jeffrey A Fessler. Combining ordered subsets and momentum for accelerated x-ray ct image reconstruction. *IEEE transactions on medical imaging*, 34(1):167–178, 2015.

[386] Ruoqiao Zhang, Jean-Baptiste Thibault, Charles A Bouman, Ken D Sauer, and Jiang Hsieh. Model-based iterative reconstruction for dual-energy x-ray ct using a joint quadratic likelihood model. *IEEE transactions on medical imaging*, 33(1):117–134, 2014.

[387] GR Davis and JC Elliott. X-ray microtomography scanner using time-delay integration for elimination of ring artefacts in the reconstructed image. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 394(1):157–162, 1997.

[388] Emil Y Sidky, Lifeng Yu, Xiaochuan Pan, Yu Zou, and Michael Vannier. A robust method of x-ray source spectrum estimation from transmission measurements: Demonstrated on computer simulated, scatter-free transmission data. *Journal of applied physics*, 97(12):124701, 2005.

[389] Li Zhang, Guowei Zhang, Zhiqiang Chen, Yuxiang Xing, Jianping Cheng, and Yongshun Xiao. X-ray spectrum estimation from transmission measurements using the expectation maximization method. In *Nuclear Science Symposium Conference Record, 2007. NSS'07. IEEE*, volume 4, pages 3089–3093. IEEE, 2007.

[390] Wenying Wang, Steven Tilley II, Matthew Tivnan, and J Webster Stayman. Local response prediction in model-based ct material decomposition. In *15th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, volume 11072, page 110720Z. International Society for Optics and Photonics, 2019.

[391] Rainer Raupach and Thomas G Flohr. Analytical evaluation of the signal and noise propagation in x-ray differential phase-contrast computed tomography. *Physics in Medicine & Biology*, 56(7):2219, 2011.

[392] FM Epple, G Potdevin, P Thibault, S Ehn, J Herzen, A Hipp, F Beckmann, and F Pfeiffer. Unwrapping differential x-ray phase-contrast images through phase estimation from multiple energy data. *Optics express*, 21(24):29101–29108, 2013.

[393] Franz M Epple, Sebastian Ehn, Pierre Thibault, Thomas Koehler, Guillaume Potdevin, Julia Herzen, David Pennicard, Heinz Graafsma, Peter B Noël, and Franz Pfeiffer. Phase unwrapping in spectral x-ray differential phase-contrast imaging with an energy-resolving photon-counting pixel detector. *IEEE transactions on medical imaging*, 34(3):816–823, 2015.

[394] Yongshuai Ge, Ran Zhang, Ke Li, and Guang-Hong Chen. Improving radiation dose efficiency of x-ray differential phase contrast imaging using an energy-resolving grating interferometer and a novel rank constraint. *Optics express*, 24(12):12955–12968, 2016.

[395] Georg Pelzer, Andrea Zang, Gisela Anton, Florian Bayer, Florian Horn, Manuel Kraus, Jens Rieger, Andre Ritter, Johannes Wandner, Thomas Weber, et al. Energy weighted x-ray dark-field imaging. *Optics express*, 22(20):24507–24515, 2014.

[396] Mini Das and Zhihua Liang. Spectral x-ray phase contrast imaging for single-shot retrieval of absorption, phase, and differential-phase imagery. *Optics letters*, 39(21):6343–6346, 2014.

[397] F Kelcz, FE Zink, WW Peppler, DG Kruger, DL Ergun, and CA Mistretta. Conventional chest radiography vs dual-energy computed radiography in the detection and characterization of pulmonary nodules. *AJR. American journal of roentgenology*, 162(2):271–278, 1994.

[398] Hany Kashani, Jianan Grace Gang, Nicholas A Shkumat, Carlos A Varon, John Yorkston, Richard Van Metter, Narinder S Paul, and Jeffrey H Siewerdsen. Development of a high-performance dual-energy chest imaging system: initial investigation of diagnostic performance. *Academic radiology*, 16(4):464–476, 2009.

[399] E. Roessl and R. Proksa. K-edge imaging in x-ray computed tomography using multi-bin photon counting detectors. *Physics in Medicine and Biology*, (15):4679–4696, 2007.

[400] Udo van Stevendaal, Zhentian Wang, Thomas Köhler, Gerhard Martens, Marco Stampanoni, and Ewald Roessl. Reconstruction method incorporating the object-position dependence of visibility loss in dark-field imaging. In *Medical Imaging 2013: Physics of Medical Imaging*, volume 8668, page 86680Z. International Society for Optics and Photonics, 2013.

[401] Christian Kottler, Vincent Revol, Rolf Kaufmann, and Claus Urban. Dual energy phase contrast x-ray imaging with talbot-lau interferometer. *Journal of Applied Physics*, 108(11):114906, 2010.

[402] Simba chest health analysis system public lung image database. `http://www.via.cornell.edu/databases/simbadb.html`. Accessed: 2018-07-31.

[403] Xu Ji, Yongshuai Ge, Ran Zhang, Ke Li, and Guang-Hong Chen. Studies of signal estimation bias in grating-based x-ray multicontrast imaging. *Medical physics*, 44(6):2453–2465, 2017.

[404] Simone Schleede, Felix G Meinel, Martin Bech, Julia Herzen, Klaus Achterhold, Guillaume Potdevin, Andreas Malecki, Silvia Adam-Neumair, Sven F Thieme, Fabian Bamberg, et al. Emphysema diagnosis using x-ray dark-field imaging at a laser-driven compact synchrotron light source. *Proceedings of the National Academy of Sciences*, 109(44):17880–17885, 2012.

[405] Katharina Hellbach, Andre Yaroshenko, Konstantin Willer, Tina Pritzke, Alena Baumann, Nina Hesse, Sigrid Auweter, Maximilian F Reiser, Oliver Eickelberg, Franz Pfeiffer, et al. Facilitated diagnosis of pneumothoraces in newborn mice using x-ray dark-field radiography. *Investigative radiology*, 51(10):597–601, 2016.

[406] Korbinian Mechlem, Thorsten Sellerer, Manuel Viermetz, Julia Herzen, and Franz Pfeiffer. Spectral differential phase contrast x-ray radiography. *IEEE Transactions on Medical Imaging*, 2019.

[407] Michael Chabior, Tilman Donath, Christian David, Manfred Schuster, Christian Schroer, and Franz Pfeiffer. Signal-to-noise ratio in x ray dark-field imaging using a grating interferometer. *Journal of applied physics*, 110(5):053105, 2011.

[408] Vincent Revol, Christian Kottler, Rolf Kaufmann, Ulrich Straumann, and Claus Urban. Noise analysis of grating-based x-ray differential phase contrast imaging. *Review of Scientific Instruments*, 81(7):073709, 2010.

[409] Wataru Yashiro, Yoshihiro Takeda, and Atsushi Momose. Efficiency of capturing a phase image using cone-beam x-ray talbot interferometry. *JOSA A*, 25(8):2025–2039, 2008.

[410] Yongshuai Ge, Ke Li, and Guang-Hong Chen. Cramér-rao lower bound in differential phase contrast imaging and its application in the optimization of data acquisition systems. In *Medical Imaging 2014: Physics of Medical Imaging*, volume 9033, page 90330F. International Society for Optics and Photonics, 2014.

[411] M Marschner, M Willner, G Potdevin, A Fehringer, PB Noël, F Pfeiffer, and J Herzen. Helical x-ray phase-contrast computed tomography without phase stepping. *Scientific reports*, 6:23953, 2016.

[412] Sebastian Kaeppler, Jens Rieger, Georg Pelzer, Florian Horn, Thilo Michel, Andreas Maier, Gisela Anton, and Christian Riess. Improved reconstruction of phase-stepping data for talbot–lau x-ray imaging. *Journal of Medical Imaging*, 4(3):034005, 2017.

[413] Paurakh L Rajbhandary, Scott S Hsieh, and Norbert J Pelc. Effect of spatio-energy correlation in pcd due to charge sharing, scatter, and secondary photons. In *Medical Imaging 2017: Physics of Medical Imaging*, volume 10132, page 101320V. International Society for Optics and Photonics, 2017.

[414] T Thuering and M Stampanoni. Performance and optimization of x-ray grating interferometry. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2010):20130027, 2014.

[415] Avinash C. Kak and Malcolm Slaney. *Principles of computerized tomographic imaging*. IEEE press New York, 1988.

[416] Rainer Raupach and Thomas Flohr. Performance evaluation of x-ray differential phase contrast computed tomography (pct) with respect to medical imaging. *Medical physics*, 39(8):4761–4774, 2012.

[417] JH Siewerdsen, IA Cunningham, and DA Jaffray. A framework for noise-power spectrum analysis of multidimensional images. *Medical physics*, 29(11):2655–2671, 2002.

[418] Maximilian von Teuffenbach, Thomas Koehler, Andreas Fehringer, Manuel Viermetz, Bernhard Brendel, Julia Herzen, Roland Proksa, Ernst J Rummeny, Franz Pfeiffer, and Peter B Noël. Grating-based phase-contrast and dark-field computed tomography: a single-shot method. *Scientific reports*, 7(1):1–8, 2017.

[419] Thorsten Sellerer, Korbinian Mechlem, Christer Ullberg, Julia Herzen, and Franz Pfeiffer. Experimental feasibility of photon-counting spectral x-ray phase-contrast computed tomography. Annual Meeting of the Radiological Society of North America, 2019.