TECHNISCHE UNIVERSITÄT MÜNCHEN
TUM School of Natural Sciences

Processing Algorithms for Grating-Based X-Ray Phase Contrast with a Clinical CT Prototype

Clemens Maria Schmid

Vollständiger Abdruck der von der TUM School of Natural Sciences der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

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

Die Dissertation wurde am 25.09.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Natural Sciences am 06.11.2023 angenommen.

TECHNICAL UNIVERSITY OF MUNICH

DISSERTATION

Processing Algorithms for Grating-Based X-Ray Phase Contrast with a Clinical CT Prototype

Author:
Clemens Maria Schmid

Supervisor:
Prof. Dr. Franz Pfeiffer

Chair of Biomedical Physics
TUM School of Natural Sciences

September 2023

Abstract

Computed tomography (CT) is a highly valuable imaging technique used in non-invasive medical diagnosis. However, for soft tissue in the human body the differences in attenuation and therefore image contrast are inherently small. Grating-based X-ray phase-contrast imaging is a relatively new method that exploits additional interaction mechanisms between photons and matter, namely refraction and small-angle scattering. This technique generates supplementary images with distinct contrast properties. Previous challenges arose due to the susceptibility of the Talbot-Lau interferometer as the primary component of the experimental setup to mechanical vibrations. Consequently, acquiring suitable data for clinical routine has been hindered in the past. In this thesis we propose a comprehensive processing pipeline designed to detect and analyze spatially and temporally variable fluctuations occurring in an interferometer installed on a continuously rotating clinical CT gantry. By leveraging correlations in the vibrations of the modular grating setup we identify a small number of relevant fluctuation modes. This enables the reconstruction of samples without the presence of vibration artifacts. Furthermore, we introduce an extended signal-retrieval method that accounts for higher-order terms of the interference pattern's Fourier series, as well as various system movements. By combining these innovative approaches we establish a processing pipeline that facilitates accurate and rapid tomographic reconstructions of the linear attenuation coefficient, linear diffusion coefficient, and refractive index decrement in clinical CT using a Talbot-Lau interferometer. The advancements presented in this thesis have the potential to significantly improve the detection and diagnosis of lung-related conditions, thereby offering valuable insights for enhanced patient care.

Zusammenfassung

Die Röntgen-Computertomographie (CT) ist ein wertvolles bildgebendes Verfahren für die nichtinvasive medizinische Diagnose. Bei weichem Gewebe im menschlichen Körper sind die Unterschiede in der Abschwächung und damit der Bildkontrast jedoch von Natur aus gering. Die gitterbasierte Röntgen-Phasenkontrast-Bildgebung ist eine relativ neue Methode, die sich zusätzliche Wechselwirkungsmechanismen zwischen Photonen und Materie zunutze macht: Brechung und Kleinwinkelstreuung. Diese Technik erzeugt zusätzliche Bilder mit eigenständigen Kontrasteigenschaften. Bisherige Herausforderungen ergaben sich aus der Anfälligkeit des Talbot-Lau-Interferometers als Hauptkomponente des Versuchsaufbaus für mechanische Vibrationen. Folglich wurde die Gewinnung geeigneter Daten für die klinische Routine in der Vergangenheit behindert. In dieser Arbeit wird eine umfassende Verarbeitungspipeline zur Erkennung und Analyse räumlich und zeitlich variabler Fluktuationen vorgeschlagen, die in einem Interferometer auftreten, das auf einer kontinuierlich rotierenden klinischen CT-Gantry installiert ist. Durch die Ausnutzung von Korrelationen in den Schwingungen des modularen Gitteraufbaus wird eine kleine Anzahl relevanter Fluktuationsmoden identifiziert. Dies ermöglicht die Rekonstruktion von Proben ohne das Vorhandensein von Vibrationsartefakten. Darüber hinaus wird eine erweiterte Methode zur Signalgewinnung eingeführt, die Terme höherer Ordnung in der Fourier-Reihe des Interferenzmusters sowie verschiedene Systembewegungen berücksichtigt. Durch die Kombination dieser Ansätze wird eine Verarbeitungspipeline geschaffen, die genaue und schnelle tomographische Rekonstruktionen des linearen Absorptionskoeffizienten, des linearen Diffusionskoeffizienten und der Brechungsindexabnahme in der klinischen CT mit einem Talbot-Lau-Interferometer ermöglicht. Die in dieser Arbeit vorgestellten Fortschritte haben das Potenzial, die Erkennung und Diagnose von Lungenerkrankungen erheblich zu verbessern und bieten damit wertvolle Erkenntnisse für eine verbesserte Patientenversorgung.

“It has captivated our imaginations
Haunted our dreams
And here it is
Spinning before our eyes.”

—Bas Oskam

Contents

Chapter 1

Introduction

Medical imaging using X-rays is widely utilized in contemporary clinical practice due to its effectiveness. It offers numerous advantages such as affordability, speed, and the ability to provide quantitative results, distinguishing it from magnetic resonance imaging (MRI). Nonetheless, there are certain drawbacks associated with X-ray imaging, including radiation exposure and limited contrast when imaging soft tissues. To overcome these limitations, researchers have explored the wave properties of X-ray radiation and the resulting refraction phenomena to enhance soft-tissue contrast and capture functional signals from porous structures. However, it is important to note that the refraction angles observed in this context are typically minute, necessitating modifications to conventional X-ray imaging setups to detect them accurately. One such approach involves utilizing a Talbot-Lau interferometer, which employs a grating-based phase-contrast imaging technique.

Context Grating-based X-ray phase contrast [1–3] uses the Talbot effect to retrieve additional information about the sample from the X-ray wavefront. Besides the conventional attenuation coefficient, the refractive index decrement and small-angle scattering as the linear diffusion coefficient [4] can be obtained. It has been shown that the latter provides signals from structures smaller than the system's direct resolution [5–8]. A periodic intensity pattern is generated by a modulation grating G_1, creating self-images at specific distances. The "intensity" is the local mean of the pattern, the relative magnitude of the modulation is called "visibility", and the position of the pattern is the "phase". These quantities are altered by the presence of a sample, where attenuation leads to an overall intensity reduction of the pattern, refraction shifts its lateral position, and coherent small-angle scattering (diffusion) reduces the visibility. As the interference pattern is usually too small to be resolved directly an analyzer grating G_2 is placed in front of the detector to subsample the wavefront [3]. One of the gratings is moved in small increments to obtain the convolution of the G_2 modulation with the interference pattern at multiple positions. From these data points the three signals transmission, dark-field contrast, and differential phase shift can be retrieved. This procedure is called "phase stepping" [3, 9]. The method was developed with highly coherent synchrotron radiation and brought to laboratory setups by including a third grating G_0 in the interferometer [9]. Placed between G_1 and a conventional X-ray source, it transforms the latter into many small slit sources which are mutually incoherent but produce individual G_1 interference patterns adding up constructively at the detector plane. The combination of G_0, G_1, and G_2 gratings is called a Talbot-Lau interferometer. Fig. 3.1 shows a sketch of the experimental setup of such an interferometer with inverse geometry [10] in which the sample is placed between G_1 and G_2.

Studies involving in-vivo pigs [5, 11] and ex-vivo human subjects [7, 12] provided compelling evidence of the potential and feasibility of dark-field chest radiography in humans. The translation of this technology to human lung imaging posed challenges primarily due to the size of the human thorax and the production difficulties associated with the reacquired gratings. The Klinikum rechts der Isar (the university hospital of the Technical University of Munich) successfully installed and commissioned the first system for clinical dark-field chest radiography [13]. This achievement marks a significant milestone in bringing the potential of dark-field X-ray imaging to the forefront of clinical practice, offering new avenues for improved diagnosis, monitoring, and management of lung-related conditions [14–18].

The application of dark-field imaging extends beyond radiography to include dark-field computed tomography [19]. This advanced technique enables three-dimensional visualization and analysis of the lung, offering valuable insights into various pathologies. The clinical significance of dark-field imaging lies in its ability to enhance the detection of subtle tissue changes, such as early signs of lung diseases or the assessment of treatment response. Furthermore, sophisticated image processing algorithms play a crucial role in extracting meaningful information from dark-field images, aiding in diagnosis and treatment planning.

Research project The first Talbot-Lau interferometer mounted in a continuously rotating clinical gantry is presented in [20]. It is a modified commercial CT platform with 70 cm bore size operated in a standard clinical scan protocol and with sufficient field-of-view

to capture a human. Its interferometer design [21] and characteristics [22] have been investigated in the literature.
Unlike conventional lab-based setups, this scanner presents a highly unstable environment for interferometry due to its continuous rotation. The stepping curve is sampled based on inherent system vibrations, eliminating the need for explicit phase stepping. These vibrations introduce spatial fluctuations in intensity, visibility, and fringe phase, independent of the sample's properties. The movement of the gratings caused by vibrations and centrifugal forces from rotation results in complex spatial phase variations. Moreover, the interferometer's visibility is reduced due to the scanner's movement during detector exposures, effectively blurring the modulation pattern. To cover the entire detector the G_2 grating is constructed by combining multiple smaller gratings with each potentially vibrating independently. Finally, slight movements of the X-ray focal spot on the rotating anode cause defocusing in the interferometer, leading to spatial fluctuations in visibility and intensity. The triangular G_1 profile combined with unsteady sampling of the stepping curve in a small angular window leads to unprecedented relevance of higher-order Fourier terms of the wavefront sampled by the G_2.

Contribution This work presents several advances of processing algorithms used in grating-based dark-field and phase imaging to facilitate dark-field CT with in clinical practice.
We propose a method to model the spatio-temporal fluctuations in a Talbot-Lau interferometer. The fluctuations in visibility and phase are modeled per G_2 tile as linear combinations of two-dimensional polynomials. To enable reconstruction of a sample scan, these are reduced to a joint vibration model over all G_2 tiles by employing principal component analysis (PCA) and keeping only few components after factor analysis. The joint model is translated to a sample scan and enables the estimation of tile-wise vibrations behind a sample by determining the coefficients of the PCA model outside of the sample.
We also introduce an extended signal-retrieval method to account for higher-order terms of the interference pattern's Fourier series, the macroscopic movement of the gratings during a gantry rotation, slow thermal deformation, and continuous sample rotation during acquisition.
The combination of the presented methods forms a processing pipeline for fast, accurate tomographic reconstructions of linear attenuation coefficient, linear diffusion coefficient, and refractive index decrement in the context of clinical CT with a Talbot-Lau interferometer.

Chapter 2

Background and previous work

This chapter presents the theoretical foundations of X-ray phase and dark-field imaging with a Talbot-Lau interferometer. The chapter begins with an overview of conventional clinical X-ray imaging techniques, including physical processes, radiography, and CT. These techniques have been the cornerstone of medical imaging for many years, providing valuable insights into the internal structures of the human body. However, they are limited in their ability to differentiate soft tissues with similar attenuation coefficients and may expose patients to higher doses of ionizing radiation.

The chapter then discusses the principles of X-ray grating interferometry with clinical sources, focusing on the Talbot-Lau interferometer. This innovative imaging setup leverages the wave nature of X-rays and the interference phenomena to extract additional information beyond conventional absorption-based imaging. The Fourier series and stepping curve analysis are discussed, providing a mathematical framework to understand the interference patterns generated by the grating interferometer.

The next section explores the various types of image projections that can be obtained from the Talbot-Lau interferometer. These projections include visibility, dark-field, and differential phase projections, each capturing distinct physical phenomena and providing complementary information about the imaged object. The mathematical formulations and interpretations of these projections are presented, highlighting their potential for enhancing image contrast and enabling the visualization of subtle tissue variations.

Finally, the chapter discusses the application of tomography to extract quantitative parameters from the acquired projections, specifically focusing on the estimation of the diffusion coefficient and refractive index decrement. Tomography plays a crucial role in reconstructing cross-sectional images of the imaged object and enables the three-dimensional visualization of internal structures. The principles of tomographic reconstruction, including line integrals and the filtered back-projection algorithm, are elaborated.

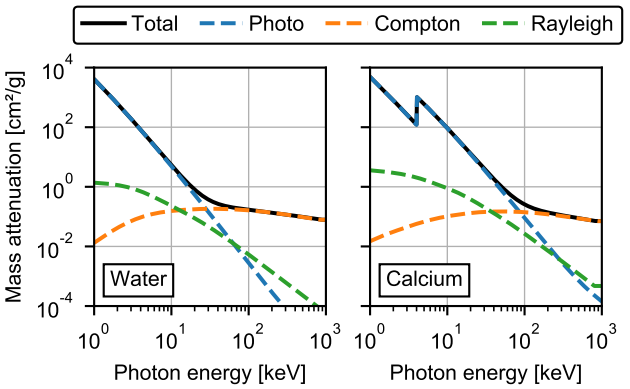

Figure 2.1: Energy dependence of the mass attenuation coefficient of water and calcium. The total attenuation is the sum of photoelectric absorption, Compton, and Rayleigh scatting, respectively. The sharp edge for calcium corresponds to its K-edge at 4 keV.

2.1 Conventional clinical X-ray imaging

In the field of medical imaging, X-ray imaging plays a crucial role in the diagnosis and treatment of various diseases and conditions. Conventional clinical X-ray imaging techniques, such as radiography and CT, have been widely used for decades. In this section, we will explain the fundamental principles underlying these techniques as well as their limitations. It is largely based on the comprehensive discussions in [23] and [24].

2.1.1 Interactions and image contrast

X-ray radiation consists of high-energy photons, i.e. electromagnetic radiation, in the approximate energy range of 1 keV to 1000 keV per photon. The effect relevant to conventional clinical imaging is their absorption on the path from X-Ray source through patient (or object) to the individual detector pixel. One single photon is assumed to either reach the pixel unaffected or being completely absorbed such that it does not reach the detector pixel—or the detector at all, to be precise. The attenuation is therefore not an effect of a single photon, but only defined as the reduction of an initial beam intensity comprised of many photons.

The attenuation of X-rays in the specified energy range is a compound effect of multiple physical processes where the most important three are photoelectric absorption, Compton scattering, and Rayleigh scattering. All of these processes depend on the X-Ray energy E and the atomic number Z of the material in a different way. In clinical imaging they are combined into one macroscopic quantity, the linear attenuation coefficient μ. It can be expressed as the sum of the scattering cross sections of the individual processes,

$$\begin{aligned}\frac{\mu(E,Z)}{\rho_\text{a}} &= \sigma_\text{tot}(E,Z) \qquad (2.1)\\ &= \sigma_\text{Photo}(E,Z) + \sigma_\text{Comp}(E,Z) + \sigma_\text{Rayl}(E,Z)\,,\end{aligned}$$

with the atomic number density ρ_a, the total cross section $\sigma_\text{tot}(E,Z)$, and the cross sections σ_Photo, σ_Comp, and σ_Rayl of photoelectric absorption, Compton scattering, and Rayleigh scattering, respectively.

The linear attenuation coefficient μ can also be motivated with the complex refractive index n of a material. It is defined as

$$n = 1 - \delta + i\beta\,, \qquad (2.2)$$

with the real part $1-\delta$ causing a phase shift of the electromagnetic wave and the imaginary part β being associated with its attenuation.

We define a plane wave ψ with amplitude ψ_0 traveling in direction $\boldsymbol{r}$ with wave vector $\boldsymbol{k}$ and frequency ω,

$$\begin{aligned}\psi(\boldsymbol{r},t) &= \psi_0 \exp\left(i\left(\boldsymbol{k}\cdot\boldsymbol{r} - \omega t\right)\right) \qquad (2.3)\\ &= \psi_0 \exp\left(i\omega\left(\left(1-\delta+i\beta\right)r/c - t\right)\right)\,, \qquad (2.4)\end{aligned}$$

where the second line uses the dispersion relation $\omega = kc/n$ and the assumption $\boldsymbol{k}\cdot\boldsymbol{r} = kr \Leftrightarrow \boldsymbol{k}\|\boldsymbol{r}$. We compare (2.4) with the wave in a vacuum,

$$\psi_\text{vac}(r,t) = \psi_0 \exp\left(i\omega\left(r/c - t\right)\right)\,, \qquad (2.5)$$

where the refractive index is defined as $n_\text{vac} = 1$ and therefore $\delta = \beta = 0$. We express (2.4) in terms of ψ_vac,

$$\psi(r,t) = \psi_\text{vac}(r,t)\exp\left(-i\delta\omega r/c\right)\exp\left(-\beta\omega r/c\right)\,, \qquad (2.6)$$

where it becomes apparent that the wave in a medium is influenced via a phase shift $-\delta\omega r/c$ (first exponential term) and exponential attenuation $\exp\left(-\beta\omega r/c\right)$ (second exponential term). The phase shift is used later in this work and not relevant to this discussion of conventional clinical X-ray imaging.

Generally, a detector can only measure the intensity of a signal. The intensity $I(r,t)$ of the wave is computed by its squared absolute value,

$$I(r,t) = |\psi(r,t)|^2\,. \qquad (2.7)$$

To compute the intensity of a wave traveling through a homogeneous medium, we square (2.6) and use the thickness d instead of r,

$$\begin{aligned}I(d,t) &= \psi_0^2 \exp\left(-\left(4\pi\beta/\lambda\right)d\right) \qquad (2.8)\\ &= I_0 \exp\left(-\mu d\right)\,, \qquad (2.9)\end{aligned}$$

with the second line using the relation $\omega = 2\pi c/\lambda$ and the definition of the linear attenuation coefficient,

$$\mu = 4\pi\beta/\lambda\,. \qquad (2.10)$$

(2.9) is commonly called the Lambert-Beer law. It describes the exponential attenuation of X-rays and therefore governs the image formation in clinical X-ray imaging. The unit of the linear attenuation coefficient is cm^{-1}. It is also common to provide the mass attenuation coefficient $\mu/\rho(E)$ given in units of $\text{cm}^2\,\text{g}^{-1}$ where ρ is the density of the material, compound, or mixture in $\text{g}\,\text{cm}^{-3}$.

For this thesis, the exact physical processes underlying the attenuation coefficient are not directly relevant. We use it mainly as a macroscopic quantity describing the material-specific attenuation of an X-ray beam and thereby providing visual contrast of the human body.

2.1.2 Radiography

Radiography is a widely used X-ray imaging technique that provides two-dimensional (2D) images of the internal structures of an object. It involves the transmission of X-rays through the object onto an X-ray detector, which captures the attenuated X-rays. The detector converts the X-rays into an electrical signal, which is then processed to produce a radiographic image. Corresponding to (2.9), the measured intensity $I(u,v)$ at pixel (u,v) is basically the shadow of the human body in the X-ray spectrum.

A clinical X-ray source emits a broad spectrum of X-rays that pass through the patient's body. The X-rays are attenuated to different extents as they interact with the various tissues and structures within the body. Dense structures, such as bones, attenuate X-rays more than soft tissues, resulting in differences in the transmitted X-ray intensities. These intensity variations are captured by the X-ray detector, forming the basis for the contrast in the radiographic image.

The acquired radiographic image represents a 2D projection of the internal structures. The intensity of each pixel in the image corresponds to the X-ray attenuation along the path from the X-ray source to the detector. The technique is comparatively simple, affordable, and requires only little radiation dose. However, due to the superposition of structures along the X-ray path, the resulting image may lack depth information

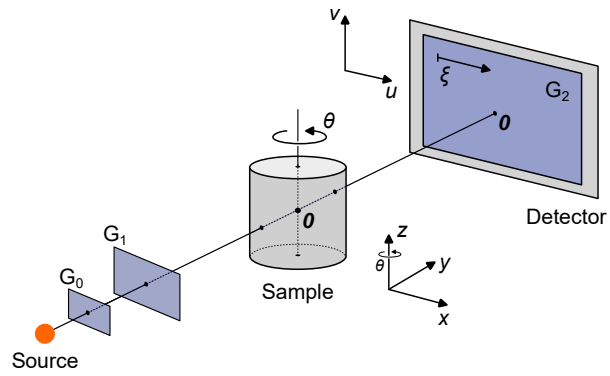

Figure 2.2: Coordinate system of a cone-beam setup with a Talbot-Lau interferometer in inverse geometry. The location in the sample is described with (x, y, z) and on the projection with (u, v). The X-ray beam travels along positive y. $(0, 0, 0)$ is in the center of the sample and $(0, 0)$ in the center of the projection. The stepping position of G_2 is described by ξ. The bars in each grating are along z or v, respectively. The sample coordinates rotate with θ around z.

and suffer from overlapping structures, limiting its diagnostic capabilities.

2.1.3 Tomographic reconstruction

Tomography is a powerful imaging technique that allows for the visualization of internal structures in a cross-sectional manner. It reconstructs a 2D or 3D representation of the object based on its projection data acquired from multiple angles. This section discusses the canonical case of parallel X-ray illumination, a 2D object, and a flat detector. The coordinate systems for a general cone-beam setup are illustrated in Fig. 2.2. It also depicts the three gratings of a Talbot-Lau interferometer which is discussed in section 2.2.

Line integrals At the heart of tomographic imaging lies the concept of line integrals, which capture the attenuation of X-rays as they traverse the imaged object along different paths. The line integral represents the total attenuation coefficient along a line connecting the X-ray source and the detector. Mathematically, it is expressed as an integral of the attenuation coefficient μ,

$$l(\theta, u) = \iint \mu(x, y)\, \delta(x\cos(\theta) + y\sin(\theta) - u)\, \mathrm{d}x\, \mathrm{d}y\,, \tag{2.11}$$

where $l(\theta, u)$ is the line integral at angle θ and distance u from the center of rotation, $\delta(\cdot)$ is the Dirac delta function, and $\mu(x, y)$ is the linear attenuation coefficient at position (x, y) within the object. This translation from $\mu(x, y)$ to $l(\theta, u)$ is called the "Radon transform". In this pure form it is only valid for parallel-beam geometry which is the canonical case for introducing tomographic reconstruction. In general it is called a "tomographic forward-projection" or just "projection".

By acquiring a set of line integrals at different angles a projection data set is obtained which serves as the input for the tomographic reconstruction process. The process therefore aims to obtain $\mu(x, y)$ from projection data $l(\theta, u)$. The latter is commonly called the "sinogram".

As stated before, the physically measured signal is always an intensity I. We measure $I(\theta, u)$ with the object in the beam path, $I_0(\theta, u)$ without the object and compute the line integral by taking the negative logarithm of the transmission,

$$\begin{aligned} l(\theta, u) &= -\ln \frac{I(\theta, u)}{I_0(\theta, u)} \\ &= -\ln T(\theta, u)\,, \end{aligned} \tag{2.12}$$

with $T(\theta, u)$ shorthand for the transmission in angle θ and distance u. (2.12) is a inversion of the Lambert-Beer law (2.9).

Filtered back-projection Performing the inversion of (2.11), i.e. obtaining $\mu(x, y)$ from $l(\theta, u)$, involves the "Fourier slice theorem". It states that the one-dimensional (1D) Fourier transform $\mathcal{F}_u$ of the line integral along u is equivalent to a radial slice through the 2D Fourier transform $\mathcal{F}_{x,y}$ of the original function $\mu(x, y)$,

$$\mathcal{F}_u \left[l(\theta, u)\right](\omega, \gamma) = \mathcal{F}_{x,y} \left[\mu(x, y)\right](\omega\cos\gamma, \omega\sin\gamma)\,, \tag{2.13}$$

where ω and γ are radius and angle in polar coordinates in Fourier space of the sample, respectively. By measuring line integrals $l(\theta, u)$ of an unknown object $\mu(x, y)$, we therefore radially sample the 2D Fourier transform of μ. Instead of rebinning this polar sampling into cartesian sampling and performing a 2D inverse Fourier transform (which is problematic in real-world, discretized systems) to obtain $\mu(x, y)$, we can express the inverse 2D Fourier transform in polar coordinates and combine it with (2.13),

$$\mu(x, y) = \frac{1}{2\pi} \int_0^\pi \mathcal{F}_u^{-1} \left[|\omega|\, \mathcal{F}_u \left[l(\theta, u)\right]\right] \mathrm{d}\theta\,, \tag{2.14}$$

where $\mathcal{F}_u^{-1}$ is the 1D inverse Fourier transform along u and the integral over θ from 0 to π is the "tomographic backwards-projection" or simply "back-projection". It can be seen as a smearing over the (x, y) image plane under angle θ. The term $|\omega|$ originates from the transform to polar coordinates. It is commonly seen as a high-pass filter, enhancing high-frequency features in the sinogram before back-projection. Other filter terms can be used instead to change resolution and noise of the reconstructed image $\mu(x, y)$.

Discretized notation In a real-world measurement all quantities are discretized into pixels and voxels respectively. The image is described by a vector $\boldsymbol{\mu} = (\mu_i) \in \mathbb{R}^N$ with N the number of voxels of the 2D or 3D image. The sinogram is described by the vector $\boldsymbol{l} = (l_j) \in \mathbb{R}^M$ with M the total number of data points, i.e. the number of projection angles times the number of detector pixels. The tomographic projection is written as a matrix multiplication,

$$\boldsymbol{l} = \boldsymbol{A}\boldsymbol{\mu}\,, \tag{2.15}$$

$$l_j = \sum_i a_{ij}\mu_i\,, \tag{2.16}$$

with $\boldsymbol{A} = (a_{ij}) \in \mathbb{R}^{N\times M}$ the tomographic system matrix. It describes the contribution of voxel i to line integral j. The tomographic back-projection is expressed by the transpose operation,

$$\boldsymbol{z} = \boldsymbol{A}^\top\boldsymbol{l}\,, \tag{2.17}$$

$$z_i = \sum_j a_{ij}l_j\,, \tag{2.18}$$

with $\boldsymbol{z} = (z_i) \in \mathbb{R}^M$. It is apparent that $\boldsymbol{z}$ is not generally identical to $\boldsymbol{\mu}$, i.e. the transpose of $\boldsymbol{A}$ is not its inverse. Still, many introductions to tomographic reconstruction first show the back-projection of a sinogram, somehow expecting it to be the reconstructed image.

The system matrix $\boldsymbol{A}$ is very sparse and far too large to be computed and stored at once. Instead, its elements are computed on-the-fly to perform the matrix multiplications in (2.15) and (2.17). This is typically done on graphics processing units (GPUs), leveraging their "single instruction, multiple data" (SIMD) architecture and texture memory for fast bilinear interpolation. There are multiple techniques how to compute $\boldsymbol{A}$ in a given geometry, all involving a tradeoff between accuracy and computational cost [25].

2.2 X-ray grating interferometry with clinical sources

As mentioned in the previous section, X-Rays are electromagnetic radiation and therefore governed by the material-specific refractive index n. While conventional clinical X-ray imaging uses its imaginary part β and associated absorption to generate image contrast, the real part δ also influences the radiation. More concretely, it causes a phase shift of the incoming wave, leading to a refraction by a certain angle α. The reason why this is typically ignored in clinical imaging is the fact that the angular change is minuscule for X-ray energies and materials commonly found in the human body. X-ray beams being refracted by the object reach the same detector pixel and the refraction

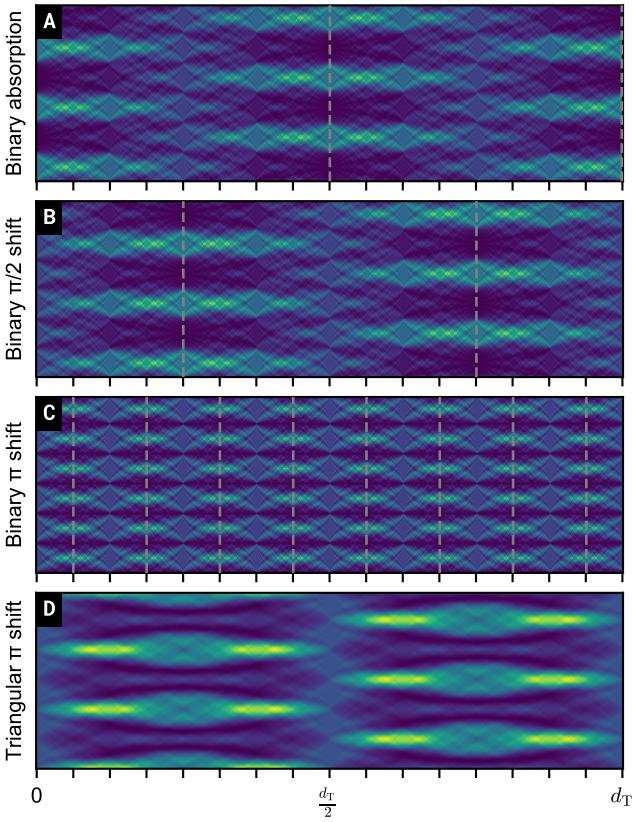

Figure 2.3: Simulated Talbot carpets for **A**-**C** the three common grating types and **D** the novel triangular design. The simulation assumes monochromatic X-Ray radiation and parallel illumination traveling from left to right. **A** The binary absorption grating creates a self-image at d_T and a laterally shifted self-image at $d_T/2$. **B** A phase-shifting G_1 with $\pi/2$ has self-images at $d_T/4$ and $3\,d_T/4$. **C** For a π phase shift there are self-images at $(2k+1)/16\,d_T$ with $k \in \mathbb{N}$ and the lateral pattern period is $p_1/2$. **D** A triangular phase shift with a maximum of π creates a complex pattern with more localized peaks but no distinct self-image.

therefore does not generate image contrast. One type of optical setup to make the refraction and associated effects visible in the context of clinical X-ray imaging is the Talbot-Lau interferometer, presented in the following.

2.2.1 Talbot effect

The core principle of the interferometer is the "Talbot effect" [26]. It describes the phenomenon that a periodic structure illuminated by a plane wave creates a diffraction pattern similar to this structure in specific distances. Here, we use a 1D grating as a periodic structure. The specific distances d_T in which self-images occur a called "Talbot distances" d_T,

$$d_T = n\frac{2p_1^2}{\lambda}\,, \tag{2.19}$$

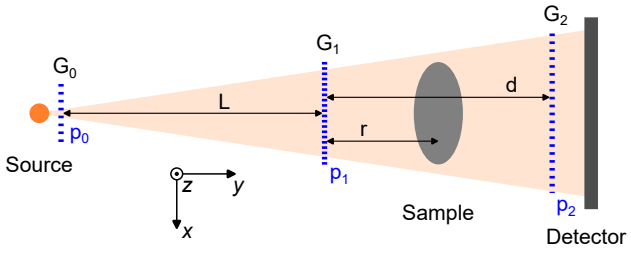

Figure 2.4: Geometry of a Talbot-Lau interferometer in a fan-beam setup. The geometric magnification is governed by the G_0-G_1 distance L and G_1-G_2 distance d. The grating periods p_0, p_1, and p_2 are chosen accordingly. The G_1-object distance r determines the effective angular sensitivity.

with $n \in \mathbb{Q}$ a rational factor depending on the grating, p_1 the grating period, and λ the wavelength of the illuminating radiation. Conventionally, for binary absorption gratings, $n \in \mathbb{N}$. More relevant to clinical application are phase-shifting gratings which introduce a periodic phase shift but have zero or little absorption themselves to maximize dose efficiency.

This grating introducing the diffraction pattern is commonly called "interference grating", "beam-splitter grating", or simply G_1. The visualization of the interference pattern over propagation distance is called "Talbot carpet". Fig. 2.3 shows Talbot carpets simulated with Fresnel propagation for different grating types and monochromatic X-ray radiation. The binary attenuation G_1 in **A** has the first self-image at $n = 1$, and a shifted self-image at $n = 1/2$. For the $\pi/2$ phase-shifting G_1 in **B**, self-images occur at $n = 1/4$ and $n = 3/4$. The π phase-shifting G_1 in **C** produces a self-image already at $n = 1/16$. Additionally, the period of the lateral intensity modulation is half the period of the G_1 itself. Finally, **D** shows a novel design with a triangular phase shift, discussed in [21] and used in the dark-field computed tomography system (DFCT). It offers an improved overall visibility for a polychromatic setup and imperfect G_0 and G_2 gratings (which is not visible in the ideal, monochromatic simulation of Fig. 2.3 **D**.)

2.2.2 Analyzer grating

To describe the geometry of the interferometer setup, we define the distance from a point-like X-Ray source to grating G_1 as L, and the distance from G_1 to detector as d. The period p_2 of the lateral intensity modulation at the detector follows from geometric magnification,

$$p_2 = p_1 \frac{L+d}{L} \,. \tag{2.20}$$

For clinical X-Ray energies the wavelength is in the order of picometers ($\lambda \approx 10^{-12}$ m). To implement a setup with a propagation distance of meters ($d_T \approx 10^{0}$ m) a grating period in the order of micrometers is required ($p_1 \approx 10^{-6}$ m) according to (2.19). Even for $d = 9L$ in (2.20) the pattern period p_2 would still be $p_2 = 10p_1 = 10^{-5}$ m.

Common clinical X-ray detectors currently do not offer pixel sizes small enough to resolve the interference pattern directly. Instead, we can use an additional grating in front of the detector to sub-sample the wavefront. The grating is commonly called "analyzer grating" or G_2. It is a binary absorption grating with period p_2 from (2.20). For a G_1 with π phase shift, the intensity pattern has period $p_1/2$ and the analyzer period p_2 is halved accordingly.

The intensity recorded by a detector pixel behind the G_2 grating is the convolution of the interference pattern in front of the G_2 with the binary absorption profile of the G_2 itself. It follows that (a) the measured intensity is not the interference pattern itself (unless G_2 has infinitesimal slit width, i.e. we convolve with the Delta function $\delta(\cdot)$), and (b) multiple measurements have to be acquired with different lateral positions in order to completely sample the modulation pattern. (a) poses no problem for imaging because we can still retrieve the relevant changes to the wavefront caused by the object. (b) is called "stepping": the gratings are moved in small increments relative to each other for each measurement, resulting in a "stepping curve" of pixel intensities depending on lateral grating shift.

2.2.3 Source grating

In reality, the X-ray source always has some finite extent. It results in a blurred interference pattern computed by projecting the spatial focal spot intensity profile through G_1 onto the detector and performing a convolution. Typical spot sizes of clinical X-Ray sources are larger than the interference pattern produced by the Talbot effect and would wash it out completely.

We introduce a third grating G_0 into the beam path placed between X-ray source and G_1 [9]. It is a binary absorption grating and transforms the large continuous focal spot into many narrow slit sources with each source as wide as one G_0 grating slit. It is therefore called "source grating". The slit sources are mutually incoherent but create constructive interference patterns at the detector if the period p_0 is chosen according to the setup geometry,

$$p_0 = p_2 \frac{L}{d} \,, \tag{2.21}$$

which is called "Lau effect" [27]. The combination of G_0, G_1, and G_2 gratings is therefore called "Talbot-Lau interferometer".

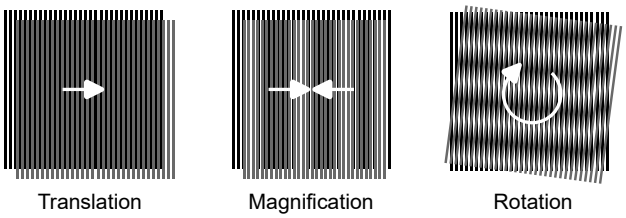

Figure 2.5: Origin of Moire fringes. Left shows a translation without fringes. The center illustrates periodicity mismatch from longitudinal displacement and resulting fringes parallel to the grating bars. Right shows rotated gratings leading to fringes perpendicular to the grating bars. (Viewing this figure on a computer screen may alter the fringes if the pixel size is similar to the rendered lines. Zoom in to avoid this problem.)

2.2.4 Stepping curve

The final intensity signal measured by a detector pixel is a combination of the wave propagation of the G_1 modulation, the convolution with the analyzer grating G_2 profile, and the convolution with the source grating G_0 profile projected through G_1 onto G_2. A detailed derivation of the resulting analytical equation can be found in [28]. In this work, it is sufficient to state that the final pixel intensity $I(\xi)$ depending on the lateral grating position ξ is described by a Fourier series,

$$I(\xi) = a_0 \left[1 + \sum_{n=1}^{\infty} a_n \cos\left(n \left(2\pi \frac{\xi}{p_2} + \phi \right) \right) \right], \tag{2.22}$$

with a_0 mean magnitude of the pattern, a_n the Fourier coefficients of the modulation, and ϕ a constant macroscopic reference phase which depends on the grating position. $I(\xi)$ in (2.22) is commonly called the "stepping curve".

The "visibility" V of a modulated signal is generally defined as the ratio between peak-to-peak and mean intensity,

$$V = \frac{I_{\text{max}} - I_{\text{min}}}{I_{\text{max}} + I_{\text{min}}}. \tag{2.23}$$

In the context of (2.22), the visibility of the stepping curve is dominated by the first Fourier coefficient a_1. We can therefore approximate the visibility with this coefficient and drop the higher terms,

$$I(\xi) \approx a_0 \left[1 + a_1 \cos\left(2\pi \frac{\xi}{p_2} + \phi \right) \right], \tag{2.24}$$

giving the canonical model in X-Ray imaging with a Talbot-Lau interferometer and clinical sources. Parameters a_0, a_1, and ϕ are setup-specific, vary over the interferometer area, and are determined empirically.

Unless the gratings are perfectly aligned with respect to each other the final macroscopic intensity pattern over the detector will be governed by Moire fringes. In general they result from the superposition of two periodic signals with slightly different periods, illustrated in Fig. 2.5. In a Talbot-Lau interferometer the periodic signals are the propagated interference pattern, the G_2 absorption profile, and the G_0 absorption profile at the location of the G_2. For a given setup with p_0, p_1, and p_2 slight deviations from (2.20) and (2.21) cause a mismatch. In terms of (2.24) this results in the phase ϕ varying with the detector pixel. Moving one of the gratings changes ξ in (2.24) globally and "moves" the Moire fringes laterally. The apparent lateral movement of the macroscopic fringes is much larger than the actual microscopic grating displacement ξ.

2.2.5 Image contrast

The image contrast in a stepped Talbot-Lau interferometer is associated with three separate physical effects shown in Fig. 2.6. Each of them influences (2.24) in a different way. We formulate the stepping curve $I^{\text{s}}(\xi)$ with a sample in the beam path,

$$I^{\text{s}}(\xi) = a_0^{\text{s}} \left[1 + a_1^{\text{s}} \cos\left(2\pi \frac{\xi}{p_2} + \phi^{\text{s}} \right) \right] \tag{2.25}$$

$$= a_0 T \left[1 + a_1 D \cos\left(2\pi \frac{\xi}{p_2} + \phi + \Phi \right) \right] \tag{2.26}$$

with T the sample transmission, Φ the differential phase shift, and D the so-called "dark-field" of the sample.

Although this model is motivated by a spatial interference pattern introduced in section 2.2.1, the contrast mechanisms are only corresponding to changes of the stepping curve $I(\xi)$ *within one detector pixel*. The lateral modulation of the signal (and its modification by the sample) is assumed local to a pixel. The variable ξ is the *microscopic* stepping position of the grating, not the *macroscopic* position on the detector (u, v) in Fig. 2.2. It follows that the image resolution does not depend on the grating periods but on the detector resolution.

Absorption The image contrast via absorption in the sample is used in conventional clinical X-ray imaging and discussed in section 2.1. In short, it is the combined effect of photoelectric absorption, Compton scattering, and Rayleigh scattering in the sample. As displayed in Fig. 2.6 **A**, it reduces the overall magnitude of the pattern, i.e. mean *and* modulation amplitude. The transmission is computed as the ratio of a_0^{s} and a_0,

$$T = \frac{a_0^{\text{s}}}{a_0}. \tag{2.27}$$

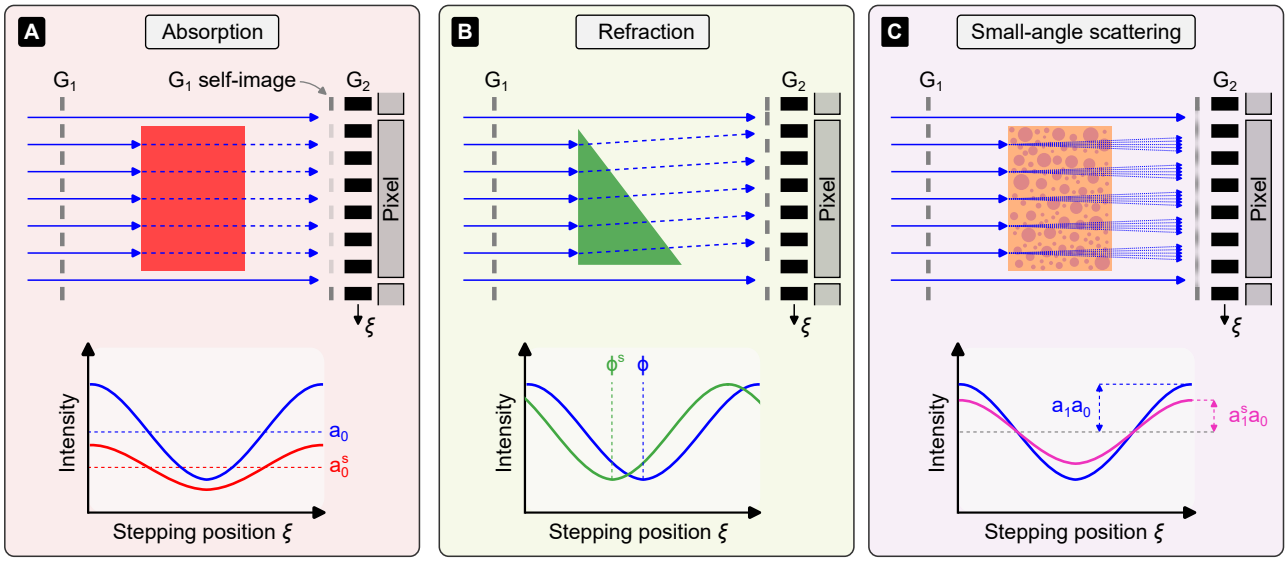

Figure 2.6: Contrast mechanisms in a Talbot-Lau interferometer. **A** Attenuation in the object leads to a reduced mean of the local stepping curve. **B** Refraction shifts the local interference pattern laterally. **C** Small-angle scattering reduces the local modulation amplitude. Note that the horizontal axis in the lineplots corresponds to lateral position ξ of the G_2, not the longitudinal beam position.

According to (2.12) we can compute the projected attenuation $l(\theta, u)$ from transmission T and perform tomographic reconstruction, giving volumetric slices of the linear attenuation coefficient μ.

Refraction As discussed in section 2.1.1, X-Ray radiation is influenced by the real part δ of the refractive index n, causing a material-specific refraction and lateral phase shift. This effect is visualized in Fig. 2.6 **B**. More precisely, the refraction angle α is proportional to the spatial gradient of the projected electron density ρ_e,

$$\alpha \propto \frac{\partial}{\partial u} \int \rho_e(x, y, z)\, \mathrm{d}y\,, \tag{2.28}$$

where we assume the X-ray beam traveling in y direction according to Fig. 2.2 for ease of notation. The lateral interference pattern shift Δu is proportional to α,

$$\Delta u = d \tan\alpha \approx d\alpha\,, \tag{2.29}$$

where we use the small-angle approximation for α and G_1-G_2 distance d from (2.20) and Fig. 2.4. The resulting phase shift Φ of the stepping curve (2.26) is computed using the G_2 period p_2,

$$\Phi = 2\pi \frac{\Delta u}{p_2} = 2\pi \frac{d}{p_2}\alpha = S_\alpha \alpha\,. \tag{2.30}$$

The term S_α is the "angular sensitivity" of the interferometer. It connects the physically intrinsic refraction angle α with the measured lateral phase shift Φ.

(2.30) assumes that the refraction occurs at the position of G_1. In real setups the sample is placed somewhere between G_1 and G_2 leading to a reduced sensitivity $\widetilde{S}_\alpha$,

$$\widetilde{S}_\alpha(r) = S_\alpha \left(1 - \frac{r}{d}\right)\,, \tag{2.31}$$

with r the distance from G_1 to sample. For large samples with extent similar to r and d, the sensitivity varies within the sample and makes tomographic reconstruction more complicated. For small samples, we can assume a constant factor r/d.

Combining (2.28) and (2.30) we see that integrating the phase shift Φ yields the projected electron density ρ_e,

$$\int \Phi(u)\, \mathrm{d}u \propto \int \rho_e(x, y, z)\, \mathrm{d}y\,, \tag{2.32}$$

with u the pixel coordinate from (2.14) and the detector axis perpendicular to the grating bars, i.e. the direction sensitive for refraction. Integrating $\Phi(u)$ is commonly combined with filtered back-projection and more specifically the Fourier filter acting along u. Generally, the derivative of a function $f(u)$ can be expressed via its Fourier transform,

$$\mathcal{F}_u\left[\frac{\partial}{\partial u} f(u)\right](\omega) = 2\pi i\omega \mathcal{F}_u\left[f(u)\right](\omega)\,, \tag{2.33}$$

with ω the coordinate in Fourier space corresponding so u in real space. Similarly, the integral of $f(u)$ can be expressed via its Fourier transform,

$$\mathcal{F}_u\left[\int f(u)\, \mathrm{d}u\right](\omega) = \frac{\mathcal{F}_u\left[f(u)\right](\omega)}{2\pi i\omega}\,, \tag{2.34}$$

omitting some arbitrary constant of integration. We combine (2.34) with the conventional equation for filtered back-projection (2.14),

$$\begin{aligned}\rho_{\mathrm{e}}(x,y) &= \\ &= \frac{1}{2\pi}\int_0^\pi \mathcal{F}_u^{-1}\left[|\omega|\,\mathcal{F}_u\left[\int \Phi(\theta,u)\,\mathrm{d}u\right]\right]\mathrm{d}\theta \\ &= \frac{1}{2\pi}\int_0^\pi \mathcal{F}_u^{-1}\left[\frac{\operatorname{sgn}\omega}{2\pi i}\,\mathcal{F}_u\left[\Phi(\theta,u)\right]\right]\mathrm{d}\theta\,, \end{aligned} \tag{2.35}$$

fusing the integration of $\Phi(u)$ and filter $|\omega|$ into the "Hilbert filter" $\operatorname{sgn}\omega/2\pi i$. We can therefore obtain volumetric slices of the electron density ρ_{e} in the object by performing filtered back-projection with adapted filter weights directly on the differential phase projections Φ.

Small-angle scattering The third contrast mechanism is small-angle scattering at material interfaces displayed in Fig. 2.6 **C**. The scattered X-rays still reach the same detector pixel but the lateral coherence and according modulation amplitude of the stepping curve is reduced. Mathematically, a_0 of (2.24) is unaffected while a_1 is lower, i.e. the visibility is reduced. The ratio between a_1^{s} and a_1 is the dark-field signal D in (2.26),

$$D = \frac{a_1^{\mathrm{s}}}{a_1}\,. \tag{2.36}$$

"Scattering" in this context means small deflections of the X-ray beams at material intersections, i.e. microscopic changes of the electron density ρ_{e}. The exact definition varies in the literature. In this work we define dark-field and Rayleigh scattering to be separate effects. While the strongly forward-directed Rayleigh scattering in microscopically homogeneous matter can cause a loss in lateral coherence of the stepping curve, we would regard this effect as an artifact rather than a valid source of dark-field signal.

In contrast, we do regard microscopic structures as the cause of dark-field contrast. The X-ray beam is refracted at each material interface with the normal vectors of the interfaces assumed isotropic. After many interactions the beam is "widened" and reaches the detector pixel. The small-angle scattering in dark-field contrast can therefore be interpreted as unresolved refraction. The microscopic structures are not resolved directly but generate a signal indicating their presence.

While the dark-field projections already offer diagnostic value [14–18] we can also compute a volumetric representation of the scattering power via tomographic reconstruction: the "linear diffusion coefficient" ε. It is motivated by considering a sharp X-ray beam being widened into a Gaussian angular distribution $A_1(\alpha)$ by a thin slice of scattering material,

$$A_1(\alpha) = \frac{1}{\sigma_1\sqrt{2\pi}}\exp\left(-\frac{\alpha^2}{2\sigma_1^2}\right)\,, \tag{2.37}$$

with azimuthal angle α from the initial beam direction and σ_1 the width of the distribution. After many such slices, the final distribution A is a convolution of all distributions A_i,

$$A(\alpha) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{\alpha^2}{2\sigma^2}\right)\,, \tag{2.38}$$

with total width $\sigma^2 = \sum_i \sigma_i^2$. Going from discrete slices to a continuous distribution $\sigma(x)$,

$$\sigma^2 = \int \frac{\mathrm{d}\sigma^2(x)}{\mathrm{d}x}\,\mathrm{d}x = \int \varepsilon(x)\,\mathrm{d}x\,, \tag{2.39}$$

gives the scattering power density $\varepsilon(x)$. We assume small angles α and transform the angular distribution $A(\alpha)$ into a spatial distribution $\widetilde{A}(\xi)$ on the detector,

$$\widetilde{A}(\xi) = \frac{1}{\sigma d\sqrt{2\pi}}\exp\left(-\frac{\xi^2}{2d^2\sigma^2}\right)\,, \tag{2.40}$$

with d the distance between object and detector and ξ the stepping coordinate within one pixel. The stepping curve $I(\xi)$ with period $2\pi/p_2$ is convolved with this lateral distribution $\widetilde{A}(\xi)$ resulting in a new stepping curve. Its modulation amplitude is reduced by D,

$$\begin{aligned} D &= \exp\left(-\frac{2\pi^2 d^2}{p_2^2}\sigma^2\right) \\ &= \exp\left(-\frac{S_\alpha^2}{2}\int \varepsilon(x,y)\,\mathrm{d}y\right)\,, \end{aligned} \tag{2.41}$$

connecting the measured visibility reduction D from (2.36) with the line integral of the linear diffusion coefficient ε over y via the angular sensitivity S_α from (2.30). The volumetric distribution $\varepsilon(x,y)$ can therefore be obtained from D with tomographic reconstruction analogous to the linear attenuation coefficient μ in section 2.1.3.

Coming back to the discussion about the term "scattering", (2.41) highlights that we define valid dark-field signal as generated by a mechanism subject to Gaussian scattering profiles and the line integral thereof, instead of merely reducing the observed ratio a_1^{s}/a_1.

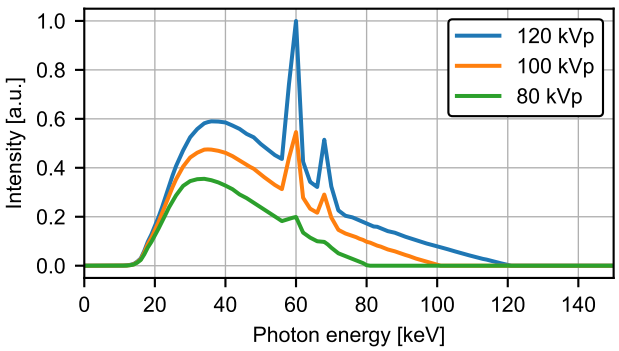

Figure 2.7: Simulated polychromatic X-ray spectra for peak photon energies 80 keV, 100 keV, and 120 keV. The spectrum is a combination of the continuous bremsstrahlung and the distinct K_α and K_β emission lines of the tungsten anode.

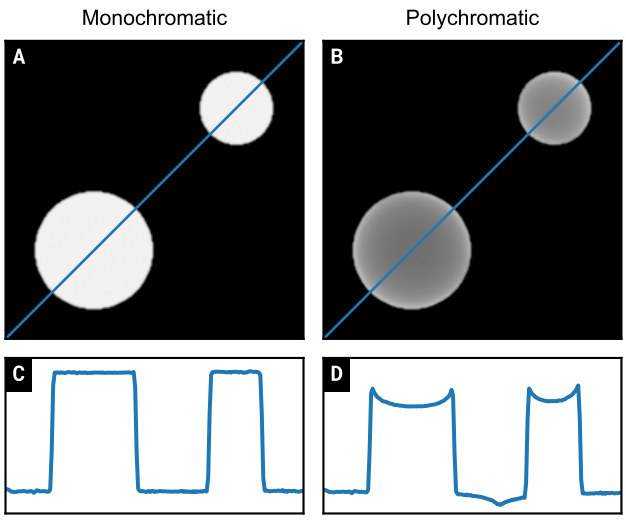

Figure 2.8: Beam-hardening artifacts in the reconstructed linear attenuation coefficient μ of two aluminum cylinders. **A** shows a slice of the monochromatic simulation. **B** shows the reconstructed image of a polychromatic simulation. **C**,**D** are line profiles through the cylinders. The underestimation of the attenuation coefficient manifests as cupping.

2.3 Image artifacts

The discussion so far assumed monochromatic X-ray radiation and ideal interferometer gratings. This section goes into image artifacts resulting from polychromatic X-Ray radiation and realistic gratings.

2.3.1 Spectral effects

Polychromatic X-ray spectrum The spectrum of clinical X-ray tubes is not monochromatic. Broadly speaking, they use a high voltage V_{peak} to accelerate electrons onto a target anode made from tungsten or molybdenum. The electrons are decelerated or absorbed in the material which generates X-ray photons. Their energy spectrum is a combination of the continuous "bremsstrahlung" and the sharp K_α and K_β lines of the anode material. Simulated spectra for a tungsten anode and different V_{peak} are shown in Fig. 2.7.

Beam-hardening As shown in Fig. 2.1 the linear attenuation coefficient $\mu(E)$ depends on the X-ray energy E and generally decreases with increasing energy in the range relevant to clinical imaging. Combined with the polychromatic X-ray spectrum this leads to "beam-hardening": the spectrum of the X-ray beam changes while traveling through the sample because lower energies are absorbed more. The mean attenuation coefficient is lower for the hardened beam. Overall, this results in an underestimation of the reconstructed μ as illustrated in Fig. 2.8.

Imperfect gratings, polychromatic visibility The interferometer gratings are finite physical objects governed by the same physics as the measured object. Both absorption and phase shift induced on the wavefront depend on the X-ray energy and the gratings' thickness, materials, and microscopic structures.
G_0 and G_2 are binary absorption gratings. Ideally their transmission profile should be either 1 or 0. For real gratings, the absorption in the grating bars is finite and decreases with increasing photon energy. Furthermore, the absorption of grating slits is non-zero because they contain "bridges" for mechanical support. The G_1 grating is either also a binary absorption or a phase-shifting grating. In the first case, the same limitations apply as for G_0 and G_2. For the phase-shifting case, the phase shift is not constant but inversely proportional to the X-ray energy. Additionally the thin G_1 absorbs some of the radiation in its bars.

These effects lead to a interferometer visibility $V(E)$ that depends on the photon energy E. Simulated visibility spectra are shown in Fig. 2.9 for different G_1 designs. For ideal G_0 and G_2 gratings, the visibility would be largest for the "design energy" and decreases for increasing X-ray energies. With realistic G_0 and G_2 the optimal energy is lowered. At 80 keV the visibility increases abruptly due to the K-edge of the grating material gold.

Beam-hardening-induced visibility reduction The spectral visibility in Fig. 2.9 decreases with increasing X-ray energy. If the polychromatic beam is hardened by the energy-dependent attenuation coefficient $\mu(E)$ (as discussed before), the effective visibility decreases even if no small-angle scattering has occurred. This leads to an overestimation of the diffusion coefficient ε generally in all attenuating materials but most noticeably in homogeneous materials like water. The visibility loss $\Delta V(T)$ can be modeled as a function of the sample transmission T for some homogeneous, absorbing calibration material like polyoxymethylene (POM) [22]. After

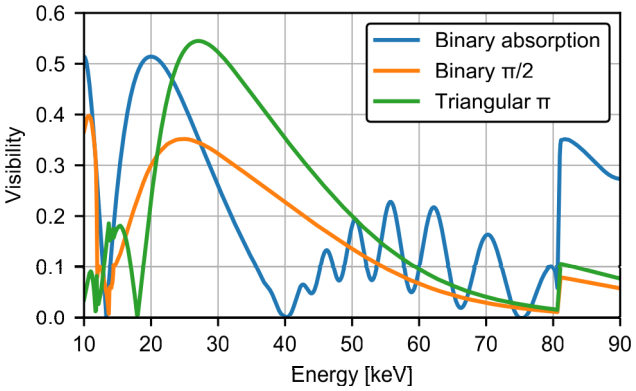

Figure 2.9: Simulation of polychromatic visibility spectra for different G_1 designs. All gratings are simulated realistically. For phase-shifting G_1's the design energy is 35 keV. G_0 and G_2 are 200 and 300 μm high and have bridge fractions of 10 % and 1 %, respectively. The absorption G_1 is 200 μm high.

signal retrieval on the sample scan the acquired dark-field D is corrected by $\Delta V(T)$ before tomographic reconstruction.

2.3.2 Interferometer vibrations

The Talbot-Lau interferometer described in section 2.2 is susceptible to mechanical instabilities due to its fine grating structures. These instabilities can lead to artifacts in the reconstructed images, which can be mitigated by incorporating a parameterization of vibrations into the image reconstruction process. This approach typically formulates the problem as an optimization problem, where the goal is to estimate the position of the stepped grating. The grating shift is modeled as a global shift of the observed phase on the detector. A widely used algorithm for this approach was presented in [29] and has been applied to grating-based X-ray differential phase contrast in [30] and [31]. This algorithm alternates between optimizing the images and the grating shift based on the model likelihood. Another alternative is to maximize image smoothness, as discussed in [32] and [33]. The vibration parameters are determined through a nested optimization process, which uses a linearized least-squares algorithm for obtaining the projections during optimization. Additionally, errors resulting from incorrect stepping positions can be mitigated through post-processing techniques outlined in [34]. These techniques predict artifacts based on the reference phase image, assuming that the errors are uncorrelated with the sample images.
Mechanical instabilities can cause gratings to vibrate, which can lead to spatial variations in the phase of the image. [35] proposed an alternating algorithm that optimizes the model likelihood with respect to grating positions and pixel-wise flat-fields. This algorithm can determine the vibration coefficients in a single step if only transverse shifts are considered. However, for movements along the optical axis or rotation of the gratings, the optimization process becomes iterative.
[36] presented a different algorithm that focuses on correcting flux variations (i.e., global intensity offsets) along with stepping errors. This algorithm is computationally very efficient, but it only accounts for global offsets in intensity and phase.
[13] described a chest radiography scanning setup that exhibits visibility fluctuations caused by the movement of the interferometer during a single exposure. To address this issue, [37] presented a processing algorithm that incorporates vibrations both during and between exposures, accounting for variations in visibility and phase.
In contrast, [38] utilized a stationary tiled interferometer that does not consider tile-wise vibrations.
[39] and [40] introduced a small animal scanner equipped with rotating gantries and gratings. The system operates in a "step-and-shoot" mode where only one grating position is measured per rotation angle. The gantry moves precisely to each angular position and remains stationary for phase stepping, rather than rotating continuously. Due to the limited scale and rotation speed of this setup, the authors observe phase drifts dependent only on the gantry's orientation.
[41] presented a larger phase-contrast CT system incorporating clinical components. This tabletop setup involves rotating the sample, but the authors do not discuss any form of grating vibrations.

In summary, various methods have been proposed to address spatial phase fluctuations, combinations of phase with visibility or flux offsets, or a combination of these factors individually. However, none of the literature mentioned simultaneously considers spatial fluctuations in all channels. Furthermore, to our knowledge, no previous work has focused on modeling the coupled vibration of a tiled G_2 grating.

2.4 Prototype for dark-field computed tomography

The measurements in this work were conducted with the dark-field computed tomography system (DFCT) presented in [20]. Both its design considerations [21] and characteristics [22] have been discussed in the literature and most detailed in [42]. This section gives only a brief overview of the setup.

The base unit is a commercial CT system "Brilliance iCT SP" manufactured by "Koninklijke Philips N.V.". It is a state-of-the-art clinical CT system with air bearings. A Talbot-Lau interferometer is mounted on the gantry with custom holders designed by our group. An annotated photograph of the setup is shown in Fig. 2.10 and a schematic drawing in Fig. 2.11. The

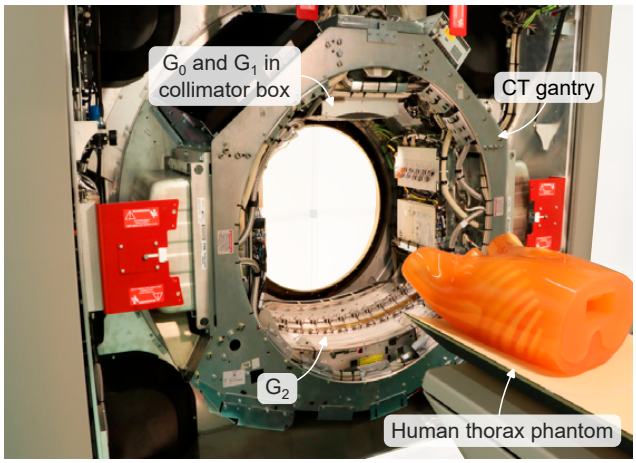

Figure 2.10: Annotated photograph of the DFCT with opened maintenance cover. It shows the gantry, the collimator box containing G_0 and G_1, the tiled G_2 in front of the detector, and a human thorax phantom on the movable table. Figure adapted from [20].

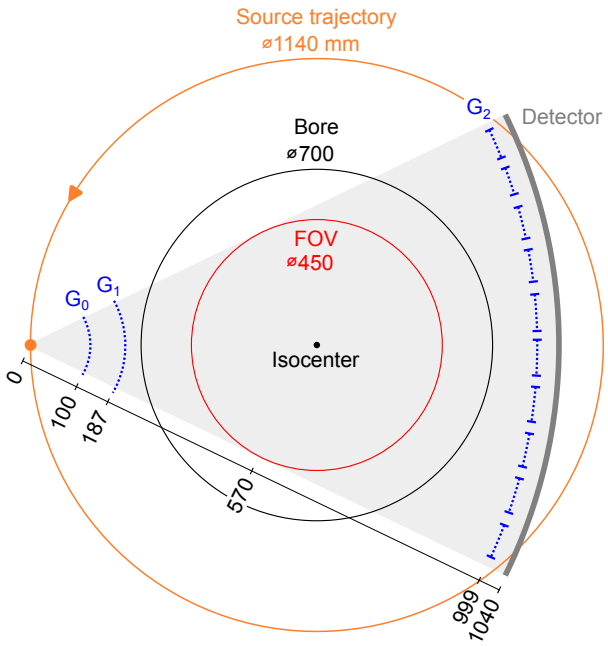

Figure 2.11: Schematic drawing of the DFCT. G_0 and G_1 are placed close to the X-ray source in "inverse" geometry. All gratings and the detector are curved to focus on the focal spot. The effective FOV is 450 mm with a fan angle of 52.5°. The G_2 is implemented by multiple small gratings tiled along detector width. The gaps between tiles are exaggerated here. All distances are given in mm.

interferometer is in inverse geometry, meaning G_0 and G_1 are placed close to the X-ray source and the object is between G_1 and G_2. G_0 and G_1 both consist of a single grating respectively and are mounted inside the collimator box ("visor") of the CT, replacing the original bowtie filter. The G_2 grating is mounted in front of the detector and covers a large width of approximately 80 cm. It consists of multiple small gratings tiled horizontally on a shared holder. Each tile can be adjusted individually to form a consistent fringe pattern over the full detector width.

Both X-ray source and detector are from the original base unit. The source is a rotating anode with up to 120 kVp and 550 mA. Its focal spot can be moved between consecutive exposures to improve spatial resolution by emulating a half-pixel detector shift [43]. More relevant in our context is the possibility of doing virtual stepping without moving the gratings [44] which is not explored in this work.
The detector is an integrating flat-panel type with a framerate of up to 4800 exposures per second. It has 64×672 pixels (rows×columns) which are 1.04 mm high and 1.41 mm wide. Due to limitations in grating manufacturing mainly concerning the G_0 height we only use the central 32 detector rows.

The Talbot-Lau interferometer is implemented in inverse geometry as shown in Fig. 2.11. G_0 and G_1 are placed very close to the X-ray source to maximize:

- dose efficiency: only G_2 blocks radiation after the patient,
- G_2 period p_2 (for a given p_0): large $d/L \Rightarrow$ large p_2 as of (2.21),

where p_0 determines the overall lateral coherence and, for a fixed total length $L + d$, the maximum angular sensitivity of the setup.

G_0 and G_2 are binary absorption gratings. Most notable is the small pitch of G_0 given its relatively large area, in combination representing the current state-of-the-art in grating manufacturing [45, 46]. The G_2 is implemented as multiple horizontally tiled gratings. They sit on a shared holder in which each tile can be adjusted to form a consistent overall fringe pattern. The G_1 is a phase grating with a triangular profile. This novel design offers improved polychromatic visibility and is easy to manufacture [21].

The purpose of the DFCT is tomographic dark-field imaging of human subjects in a clinical context. We operate the system solely with its original protocols to ensure dose compatibility. The gantry is continuously rotating to enable fast tomographic scans. As a consequence, we can not perform stepping as with conventional lab-based interferometer setups which involves multiple exposures for every gantry angle. Instead, we use movement of the gratings induced mainly by the rotating anode and the cooling pump to sample the stepping curve. This has significant implications for the signal processing which are discussed in the following chapters.

Chapter 3

Air-scan processing

This chapter introduces the model for the interferometer in the dynamic environment of a continuously rotating clinical CT gantry. The origins and parameterization of fluctuations in all image modalities are discussed. An optimization method for determining reference flat-fields terms and vibration parameterization coefficients from raw scanner data is presented. Two additional extensions to the conventional model are introduced, namely higher-order visibility terms, and macroscopic lateral movement of the gratings.

The vibration model discussed in this chapter has been published as [47].

3.1 Grating vibrations

The canonical model of a stepped Talbot-Lau interferometer without a sample is [9]

$$y_{pt}^{\text{simple}} = I_p \left(1 + V_p \cos\left(\phi_p + \gamma_t\right)\right) , \qquad (3.1)$$

with the expected intensity y_{pt}^{simple} in detector pixel $p \in \{1, \ldots, P\}$ at stepping position index $t \in \{1, \ldots, T\}$, flat-field intensity I_p, flat-field visibility V_p, flat-field phase ϕ_p, and the global phase shift γ_t induced by moving one of the gratings perpendicular to the grating bars. The flat-fields (I_p, V_p, ϕ_p) are intrinsic to the interferometer setup and usually assumed to be constant during a scan and between scans. The only parameter which is changed between exposures is the global phase offset γ_t by controlling one of the grating positions. In laboratory setups γ_t can be deliberately chosen. We use the term "explicit" phase stepping for this conventional method. The experimental setup of interest in this work is a commercial clinical gantry platform (Brilliance iCT, Philips) which has been retrofitted with gratings to enable human CT scans giving dark-field contrast. The system design is presented in [20]. The gantry is operated in a continuously rotating manner and, given the fact that the tube is not pulsed and the detector is read out at 2 to 4 kHz, it is impractical to implement explicit phase stepping. Instead, the acquisition relies on the vibrations intrinsic to the system which generate sufficient sampling of γ_t to perform phase retrieval both for air and sample-scans. Therefore the index

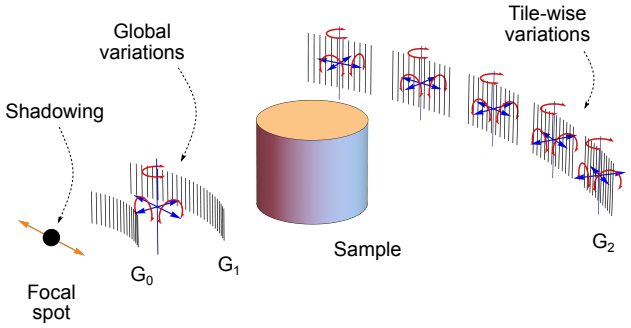

Figure 3.1: Schematic depiction of the vibrations in the setup [47]. The G_0-G_1 combination vibration creates global phase fluctuations on the detector. The individual G_2 tile movements lead to tile-wise phase and visibility variations. The focal spot movement creates global intensity and visibility fluctuations due to shadowing. The interference pattern amplitude is further reduced by grating movement during an exposure.

t refers to both the stepping position as well as the gantry angle as they are changing simultaneously. We use the term "implicit" phase stepping for exploiting the vibrations to sample the stepping curve at initially unknown points.

A schematic of the interferometer and the expected modes of vibrations is shown in Fig. 3.1. We do not use a single grating for sampling of the stepping curve and all gratings are moving simultaneously. The observed phase variation therefore is a compound effect from all gratings.

As indicated in Fig. 3.1 we expect various vibrations which will lead to a pixel- and time-dependent change of the intensity, visibility, and phase of the interference pattern. The model (3.1) is extended to

$$\begin{aligned} y_{pt} =& I_p \left(1 + I_{pt}^{\text{vib}}\right) \times \\ & \times \left[1 + V_p \left(1 + V_{pt}^{\text{vib}}\right) \cos\left(\phi_p + \phi_{pt}^{\text{vib}}\right)\right] , \end{aligned} \qquad (3.2)$$

with I_{pt}^{vib}, V_{pt}^{vib}, and ϕ_{pt}^{vib} representing spatial and temporal fluctuations over p and t.

It is obviously infeasible to fit the unknowns on the right-hand side from the measurements y_{pt} directly without any parameterization. We have $P \times T$ data

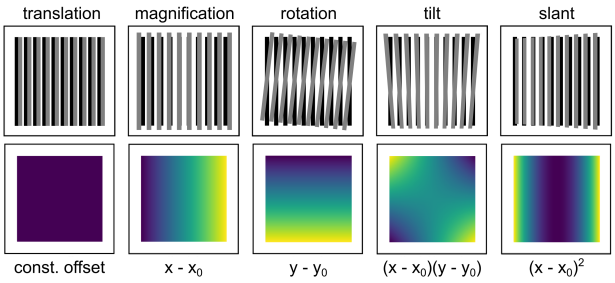

Figure 3.2: Polynomial phase shifts form grating movement and rotation. Different deviations from the desired position result in changing observed phase, modeled by 2D polynomials. The maximum order is two perpendicular to the grating bars (x here) and one parallel to the grating bars (y here).

points versus $3 \times P + 3 \times P \times T \gg P \times T$ unknowns (three flat-fields + three vibrations terms). Therefore we need some parameterization for the unknown terms.

In principle, we could come up with some parameterization for the flat-fields I_p, V_p, ϕ_p. They have very "pixel-wise" characteristics however, which are not trivial to capture with low-dimensional functions over p.

It is much more feasible to find a model for the vibrations $I_{pt}^{\text{vib}}, V_{pt}^{\text{vib}}, \phi_{pt}^{\text{vib}}$ over p, as we will show in the following. Furthermore it is more profitable in terms of decreased degrees-of-freedom (DOF) as the number of vibrations unknowns $P \times T$ per channel is much larger than the number of unknowns P per flat-field.

We initially discuss the vibration model for phase and visibility, which can be parameterized by 2D polynomials. First, we will first discuss the vibration model behind one G_2 tile, which is possible because the tiles can be treated as separate interferometers.

3.1.1 Phase fluctuations

As indicated in Fig. 3.1 we assume translations and rotations of the G_0-G_1 assembly, the G_2 carrier, and the individual G_2 tiles. According to [35] the resulting changes of the phase can be accurately modeled by low-order two-dimensional polynomials over the detector behind a G_2 tile. The effects are illustrated in Fig. 3.2. We will introduce our notation for this phenomenon in the following.

We define a two-dimensional polynomial term $\boldsymbol{\mathcal{P}}_{ij} = (\mathcal{P}_{ijp})$ of order i along the width of the detector and order j along the height of the detector in pixel p as

$$\mathcal{P}_{ijp} = u(p)^i v(p)^j \,, \tag{3.3}$$

with $u(p)$ the u coordinate along the width and $v(p)$ the v coordinate along the height of the detector in pixel p as defined in Fig. 2.2. Here, the pixel index p refers to the area behind a single G_2 tile. Each term of the polynomial is multiplied with a coefficient γ_{ijt} to calculate the amplitude of the local phase offset per exposure t, leading to the phase vibration

$$\phi_{pt}^{\text{vib}} = \sum_{i,j} \gamma_{ijt} \mathcal{P}_{ijp} \,. \tag{3.4}$$

As motivated in [35], the maximum polynomial order along the grating bars (j here) is one, whereas the maximum perpendicular (i here) is two. This equates to $(2+3) \times T$ unknowns to describe the phase vibrations for one scan (behind one G_2 tile). The global phase shift γ_t from (3.1) is represented in component $\mathcal{P}_{00p}$ and coefficient γ_{00t}. This model is therefore a superset of the conventional method of regarding only a global phase change per exposure.

3.1.2 Visibility fluctuations

Additional to the phase fluctuations we experimentally observe spatial variations in visibility. These originate from (a) movement of the gratings during one exposure and (b) movement of the focal spot. (b) causes defocusing, or more precisely a non-perpendicular incident angle of the radiation onto the grating, referred to as "shadowing".

Finite exposure time

We will first discuss (a): a change in total phase (i.e. the terms inside the cosine) during an exposure generally leads to a drop in visibility for that exposure [48] because the moving fringe pattern is spatially averaged out over the exposure time. This is the case for our scanner as it is continuously rotating and the detector is read out continuously. The amount of visibility reduction depends on the ratio between exposure time and amount of phase change during that exposure. This in turn depends on the strength of the vibration and its frequency.

[48] models visibility loss with long exposure times where one exposure integrates over many vibration cycles. We use the model introduced in [37] due to comparatively short exposure times at our scanner.

The coefficients γ_{ijt} to describe the phase vibrations are discretized in time t in our model (3.4). For the context of finite exposure times, they can be regarded as continuous variables in time τ and the vibration linearized around the timepoints t as

$$\phi_p^{\text{vib}}(\tau) \approx \phi_{pt}^{\text{vib}} + \dot{\phi}_{pt}^{\text{vib}} \tau \,, \tag{3.5}$$

where

$$\dot{\phi}_{pt}^{\text{vib}} = \left. \frac{\partial \phi_p^{\text{vib}}(\tau)}{\partial \tau} \right|_{\tau = t} \,. \tag{3.6}$$

We therefore treat the values γ_{ijt} as infinitesimal "snapshots" of the in reality continuous phase $\phi_p^{\text{vib}}(\tau)$.

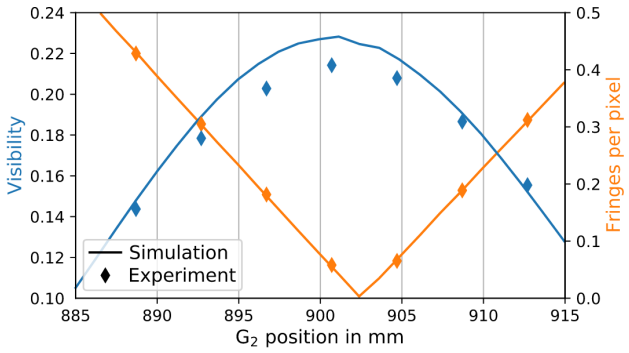

Figure 3.3: Visibility depending on the G_2 position [21]. We observe maximum visibility at one specific G_2 distance. Deviations from that position cause a noticeable reduction in visibility.

Let Δt be the exposure time during which the phase changes. To compute the observed visibility drop, we average the cosine of (3.2) over the exposure time Δt:

$$\frac{1}{\Delta t}\int_{t-\frac{\Delta t}{2}}^{t+\frac{\Delta t}{2}} \cos\left(\phi_p + \phi_p^{\text{vib}}(\tau)\right)\,\mathrm{d}\tau \approx \left(1-\frac{\Delta t^2}{24}\left(\dot{\phi}_{pt}^{\text{vib}}\right)^2\right)\cos\left(\phi_p+\phi_{pt}^{\text{vib}}\right)\,. \tag{3.7}$$

To get from the integral to the approximation, we perform a Taylor expansion in τ of the cosine to second order. The factor $\Delta t^2/24$ originates from the integration of the cosine.

Apparently, the space of possible visibility reductions caused by the finite exposure time is given as a sum of pair-wise products of the phase displacements, when using (3.4) for $\dot{\phi}_{pt}^{\text{vib}}$:

$$1-\frac{\Delta t^2}{24}\left(\dot{\phi}_{pt}^{\text{vib}}\right)^2 = 1-\frac{\Delta t^2}{24}\left(\sum_{i,j}\dot{\gamma}_{ijt}\mathcal{P}_{ijp}\right)^2\,. \tag{3.8}$$

The spatial characteristics of the visibility drop are therefore also described by 2D polynomials, but their effective maximum order is doubled compared to the phase shift of (3.4). This equation underlines our statement in the beginning of this section: the amount of visibility loss depends on the exposure time (squared) and the amount of phase change per exposure, quantified by its "velocity" $\dot{\phi}$.

Defocusing

As mentioned as (b) before, changes of the X-ray focal spot and/or grating location can also cause visibility fluctuations. Ideally, the gratings are bent to a cylindrical surface and carefully focused onto the intended source position. Due to their vibration they move out of the ideal location, which leads to a reduction of visibility [21]. The same is true for movement of the focal spot. It lies on a rotating anode typical for commercial clinical CTs which need high flux. Despite it being engineered to high precision, due to the material wearing out from electron irradiation the exact location of X-ray generation varies slightly. This is not noticeable in conventional absorption imaging. In the Talbot-Lau interferometer in our prototype, even the slightest movement changes the interference pattern and ultimately reduces the system visibility.

We assume that the observed reduction can be also approximated by 2D polynomials. This effect is additive to the aforementioned finite exposure time model. We can therefore formulate a general model for the visibility fluctuation:

$$V_{pt}^{\text{vib}} = \sum_{i,j}\beta_{ijt}\mathcal{P}_{ijp}\,. \tag{3.9}$$

The maximum polynomial order in i and j is doubled in comparison to (3.4), such that the squared phase change in (3.7) can be modeled. The coefficients β_{ijt} in (3.9) do not have a physical representation on their own. They do merely empirically model the observed total visibility fluctuation as a combination of the finite exposure time and interferometer defocusing.

This leads to $2\times(2+3)\times T$ unknowns β_{ijt} to describe the visibility fluctuation per G_2 tile.

We want to note that the reduction of visibility varying over exposures leads to variations of the noise level and potentially artifacts in reconstructions of dark-field and differential phase, as they are sensitive to the system visibility [49]. This effect is not captured by the presented model.

3.1.3 Principal vibration components

With the proposed tile-wise polynomial parameterization, we can model the individual tiles for an air-scan. This enables us to obtain satisfactory flat-fields terms I_p, V_p, and ϕ_p of (3.2). They are shown and discussed in section 3.4.1.

However, the overall goal is to measure and reconstruct an unknown sample. A detailed discussion on the processing pipeline itself is presented in chapter 4, but it also motivates the air-scan processing to produce a vibration model with as little DOF as possible. Furthermore, if the vibration over G_2 tiles is correlated, we can find a model which has joint components over all tiles. We can approximate vibrations in the center (behind the sample) from regions on the left- and right-hand side of the detector (outside the sample).

It is therefore desirable to approximate the fluctuations with few dominant, shared modes across G_2 tiles, instead of modeling each G_2 tile separately. The tile

movement is expected to be correlated because of the shared G_2 carrier. This leads to similar coefficients γ_{ijt} over some or all G_2 tiles. It also implies correlations in the visibility drop caused by the finite exposure time as we expect $\dot{\gamma}_{ijt}$ to be similar over G_2 tiles. It is beneficial to exploit these correlations and approximate the fluctuations in visibility and phase with a reduced model.

We can use principal component analysis (PCA) to find this new vibration representation. Let $\boldsymbol{X} \in \mathbb{R}^{P\times T}$ be a data matrix with P variables in the rows and T observations in the columns. The singular value decomposition (SVD) on $\boldsymbol{X}$ is given as

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\top , \tag{3.10}$$

with the orthogonal matrix $\boldsymbol{U} \in \mathbb{R}^{P\times P}$, the diagonal matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{P\times T}$, and the orthogonal matrix $\boldsymbol{V} \in \mathbb{R}^{T\times T}$. $\boldsymbol{\Sigma}$ contains the singular values of $\boldsymbol{X}$ on its diagonal which we define to be in descending order so that $\boldsymbol{\Sigma}$ is unique. The rows of $\boldsymbol{\Sigma}\boldsymbol{V}^\top$ are called the "principal components" of $\boldsymbol{X}$ and the columns of $\boldsymbol{U}$ are called the "principal directions" of $\boldsymbol{X}$, i.e. the magnitude of each principal component per observation. The principal components are used as a basis to express the data $\boldsymbol{X}$, where the basis components are sorted in order of descending "importance". We can therefore use only the first couple basis components to express $\boldsymbol{X}$ to a high degree of accuracy while drastically reducing the amount of variables. In practice, we use the SVD implementation in `numpy.linalg.svd`.

We translate our definition of PCA to be used on the vibration terms $\boldsymbol{V}^{\text{vib}} = (V_{pt}^{\text{vib}})$ and $\boldsymbol{\phi}^{\text{vib}} = (\phi_{pt}^{\text{vib}})$. Let the function PCA be defined as acting on a matrix $\boldsymbol{X}$ and returning $\boldsymbol{\Sigma}\boldsymbol{V}^\top$ and $\boldsymbol{U}^\top$:

$$\text{PCA}(\boldsymbol{X}) \to (\boldsymbol{\Sigma}\boldsymbol{V}^\top, \boldsymbol{U}^\top) . \tag{3.11}$$

To apply PCA on the combined fluctuations from all G_2 tiles, their terms V_{pt}^{vib} and ϕ_{pt}^{vib} are concatenated along the width of the interferometer and the index p is changed from locally on a G_2 tile to globally on the detector. In practice, this simply means to use independent polynomials per G_2 tile when processing the raw data, combining the resulting 2D vibrations to an array with the same shape as the raw data, and applying PCA on the combined array.

It has proven useful in practice to exclude the global fluctuations $\boldsymbol{\beta}_{00}$ and $\boldsymbol{\gamma}_{00}$ (with $\boldsymbol{\mathcal{P}}_{00} = \mathbf{1}$) from PCA and use them as the "0th" component of the PCA models. We simply add them again later to the PCA approximation in (3.14) and (3.15). This allows for an intuitive meaning of the first term as a global visibility drop and phase shift, respectively.

With this definition of PCA, the joint principal components $\boldsymbol{\mathcal{B}} = (\mathcal{B}_{kp})$ and $\boldsymbol{\mathcal{C}} = (\mathcal{C}_{kp})$ (both $\in \mathbb{R}^{P\times T}$) over all G_2 tiles of the visibility and phase fluctuations are determined via

$$\text{PCA}\left(\boldsymbol{V}^{\text{vib}} - \boldsymbol{\beta}_{00}\right) \to (\boldsymbol{\mathcal{B}}, \boldsymbol{\beta}^\star) , \tag{3.12}$$

$$\text{PCA}\left(\boldsymbol{\phi}^{\text{vib}} - \boldsymbol{\gamma}_{00}\right) \to (\boldsymbol{\mathcal{C}}, \boldsymbol{\gamma}^\star) . \tag{3.13}$$

The coefficients $\beta_{kt}^\star, \gamma_{kt}^\star \in \mathbb{R}^{T\times T}$ correspond only to the magnitude of principal component k at exposure t, not to the original polynomial model. They express the magnitude of the vibration component k in exposure t, respectively. The joint polynomial vibration across all G_2 tiles are therefore expressed in $\boldsymbol{\mathcal{B}}$ and $\boldsymbol{\mathcal{C}}$.

The fluctuations are approximated with a reduced number $B^\star$ and $C^\star$ of modes:

$$V_{pt}^{\text{vib}} \approx \beta_{00t} + \sum_{k=1}^{B^\star} \beta_{kt}^\star \mathcal{B}_{kp} , \tag{3.14}$$

$$\phi_{pt}^{\text{vib}} \approx \gamma_{00t} + \sum_{k=1}^{C^\star} \gamma_{kt}^\star \mathcal{C}_{kp} . \tag{3.15}$$

The reduced set of vibration components is used for reconstruction in a sample-scan. How to determine $B^\star$ and $C^\star$ empirically during air-scan processing is discussed in section 3.4.2. The most dominant components of $\boldsymbol{\mathcal{B}}$ and $\boldsymbol{\mathcal{C}}$ in our prototype are shown in Fig. 3.14.

3.1.4 Why no direct vibrations?

We briefly want to motivate why we use a parameterized model for V_{pt}^{vib} and ϕ_{pt}^{vib} instead of determining them directly from rearranging the model (3.2).

We rearrange (3.2) to isolate the visibility terms

$$V_p\left(1 + V_{pt}^{\text{vib}}\right) = \left(\frac{\hat{y}_{pt}}{I_p\left(1 + I_{pt}^{\text{vib}}\right)} - 1\right) \times \times \left(\cos\left(\phi_p + \phi_{pt}^{\text{vib}}\right)\right)^{-1} , \tag{3.16}$$

with $\hat{y}_{pt}$ the measured intensities. Evaluating the right-hand side of this equation leads to numerical instability as the cosine term can be very close to 0 and its reciprocal very large. It is therefore not feasible in practice.

We rearrange (3.2) in a similar manner to isolate the phase terms

$$\cos\left(\phi_p + \phi_{pt}^{\text{vib}}\right) = \left(\frac{\hat{y}_{pt}}{I_p\left(1 + I_{pt}^{\text{vib}}\right)} - 1\right) \times \times \left(V_p\left(1 + V_{pt}^{\text{vib}}\right)\right)^{-1} . \tag{3.17}$$

Evaluating this equation is feasible, as the visibility terms are not close to 0 and their reciprocal is numerically stable. However, to obtain the phase vibration we have to compute the $\arccos$ of the right-hand side. $\arccos(x)$ is only defined for $x \in (-1, 1)$. Due to noise and labile intermediate values, the right-hand side of (3.17) can be $\notin (-1, 1)$. It is therefore also not possible to directly compute ϕ^{vib}_{pt} from (3.17) in practice.

3.1.5 Intensity fluctuations

The defocusing of the interferometer caused by displacements of the focal spot and grating movement, introduced in section 3.1.2, also leads to modulations in intensity. As shown in [21] there is a shadowing effect when grating lamella are not perfectly aligned to the focal spot, i.e. deviate from a cylindrical surface around the focal spot. Furthermore, imperfections in a grating (especially cracks from bending) cause small intensity variations, which change their projected position on the detector when the position of the focal spot changes. In contrast to visibility and phase vibrations, we do not use an explicit model for these shadowing effects in intensity. The structure of the imperfections is too fine to be captured by (tile-wise) 2D polynomials. Luckily, also in contrast to visibility and phase, we can access the intensity vibrations directly:

$$I^{\text{vib}}_{pt} = \frac{\hat{y}_{pt}}{y_{pt}} - 1\,. \tag{3.18}$$

Here, $\hat{y}_{pt}$ denotes the measured values and y_{pt} the predicted signal including only visibility and phase fluctuations. In other words, we use the normalized residuum at some point in the processing pipeline—after introducing phase and visibility vibrations—to determine the main vibration modes in intensity. We compute the global flux $\boldsymbol{\alpha}_{00} = (\alpha_{00t}) \in \mathbb{R}^T$ as

$$\alpha_{00t} = \frac{\langle \hat{y}_{pt} \rangle_p}{\langle \hat{y}_{pt} \rangle_{pt}}\,. \tag{3.19}$$

As with visibility and phase, we apply PCA on the determined vibrations $\boldsymbol{I}^{\text{vib}}$ without the global fluctuation $\boldsymbol{\alpha}_{00}$:

$$\text{PCA}\left(\boldsymbol{I}^{\text{vib}} - \boldsymbol{\alpha}_{00}\right) \to (\boldsymbol{\mathcal{A}}, \boldsymbol{\alpha}^\star)\,, \tag{3.20}$$

We obtain the model $\boldsymbol{\mathcal{A}} = (\mathcal{A}_{kp}) \in \mathbb{R}^{P\times T}$ with coefficients $\boldsymbol{\alpha}^\star = (\alpha^\star_{kt}) \in \mathbb{R}^{T\times T}$. Again we approximate I^{vib}_{pt} with a reduced number $A^\star$ of modes:

$$I^{\text{vib}}_{pt} \approx \alpha_{00t} + \sum_{k=1}^{A^\star} \alpha^\star_{kt} \mathcal{A}_{kp}\,. \tag{3.21}$$

How to determine $A^\star$ empirically is discussed in section 3.4.2. The most dominant components of $\boldsymbol{\mathcal{A}}$ in our prototype are shown in Fig. 3.14.

As will be discussed later in section 3.2.3, we use 2D polynomials spanning the whole detector as an intermediate model for $\boldsymbol{\mathcal{A}}$ during the processing pipeline. Only at a later stage, we can calculate a good approximation for $\boldsymbol{I}^{\text{vib}}$ and determine the final model from that via (3.18) and (3.20).

3.2 Likelihood optimization of vibrations and flat-fields

This section first defines (linearized) statistical phase retrieval and extends it to include spatial variations in intensity, visibility, and phase. Then a minimization around SPR is presented to determine the temporal coefficients of the vibration model.

3.2.1 Linearized statistical phase retrieval

We first discuss phase retrieval regarding (3.1), meaning no per-exposure changes except the global phase offset γ_t. This is the established method used in conventional lab-based interferometer setups. The flat-field terms $(\hat{I}_p, \hat{V}_p, \hat{\phi}_p)$ are determined by minimizing the pixel-wise weighted least-squares problem

$$\mathcal{L}_p(y_{pt}, \hat{y}_{pt}) = \sum_{t=1}^{T} \frac{\left(y_{pt} - \hat{y}_{pt}\right)^2}{y_{pt}} \tag{3.22}$$

$$\approx \sum_{t=1}^{T} \frac{\left(y_{pt} - \hat{y}_{pt}\right)^2}{\hat{y}_{pt}}\,, \tag{3.23}$$

$$(\hat{I}_p, \hat{V}_p, \hat{\phi}_p) = \underset{I_p, V_p, \phi_p}{\arg\min}\, \mathcal{L}_p(y_{pt}, \hat{y}_{pt})\,, \tag{3.24}$$

with the predicted signal y_{pt} and the measured values $\hat{y}_{pt}$ in pixel p at exposure t. The weights $1/\hat{y}_{pt}$ are motivated by the measurement being a Poisson process and the variance of the signal $\hat{y}_{pt}$ therefore approximated by the signal itself.

The minimization in (3.24) can be performed directly by linearizing (3.1) with respect to the flat-field terms (I_p, V_p, ϕ_p) [49]. A trigonometric identity and change in variables to a_p and b_p is applied,

$$\begin{aligned} y_{pt} &= I_p + I_p V_p \left(\cos\phi_p \cos\gamma_t - \sin\phi_p \sin\gamma_t\right) \\ &= I_p + a_p \cos\gamma_t + b_p \sin\gamma_t\,, \end{aligned} \tag{3.25}$$

and (3.24) is solved for $(\hat{I}_p, \hat{a}_p, \hat{b}_p)$ as an overdetermined linear system with $T > 3$. The initial flat-field terms of (3.1) are obtained via

$$\hat{V}_p = \frac{\sqrt{\hat{a}_p^2 + \hat{b}_p^2}}{\hat{I}_p}\,, \tag{3.26}$$

$$\hat{\phi}_p = \arctan2(-\hat{b}_p, \hat{a}_p)\,. \tag{3.27}$$

For ease of notation the SPR can be written in matrix form. Let $\hat{\boldsymbol{y}}_p = (\hat{y}_{pt}) \in \mathbb{R}^T$ be the vector of T measurements in pixel p, $\boldsymbol{w}_p = (w_{pt}) \in \mathbb{R}^T$ the weight vector with $w_{pt} = 1/\hat{y}_{pt}$, and $\boldsymbol{\gamma} = (\gamma_t) \in \mathbb{R}^T$ the vector of global phase offsets. The system matrix $\boldsymbol{A}$ is defined as

$$\boldsymbol{A} = \begin{pmatrix} 1 & \cos\gamma_0 & \sin\gamma_0 \\ \vdots & \vdots & \vdots \\ 1 & \cos\gamma_T & \sin\gamma_T \end{pmatrix} \in \mathbb{R}^{T\times 3}\,. \tag{3.28}$$

Note that $\boldsymbol{A}$ does not depend on the pixel p. The weighted least squares version of the linearized model from (3.25) is formulated in matrix notation as

$$\left(\boldsymbol{A}^\top \boldsymbol{W}_p \boldsymbol{A}\right) \boldsymbol{x}_p = \boldsymbol{A}^\top \boldsymbol{W}_p \hat{\boldsymbol{y}}_p\,, \tag{3.29}$$

with the diagonal weights matrix $\boldsymbol{W}_p \in \mathbb{R}^{T\times T}$ given as $\boldsymbol{W}_p = \operatorname{diag} \boldsymbol{w}_p$ and the solution vector $\boldsymbol{x}_p = (I_p, a_p, b_p) \in \mathbb{R}^3$. Solving (3.29) for $\boldsymbol{x}_p$ is performed via matrix factorization. We implement the SPR with `JAX` [50] and use `jax.numpy.linalg.solve`, which in turn uses LU decomposition.

We have to extend the SPR to include vibrations in all image channels as in (3.2) in order to use it for the DFCT. Luckily, also with the addition of vibration terms, the cost function is still linear with respect to (wrt.) the flat-field terms I_p, V_p, and ϕ_p. Only the system matrix has to be adapted. We define the columns of the system matrix $\boldsymbol{A}_p \in \mathbb{R}^{T\times 3}$ in pixel p as

$$\left[\boldsymbol{A}_p\right]_{t,1} = 1 + I_{pt}^{\text{vib}}\,, \tag{3.30}$$

$$\left[\boldsymbol{A}_p\right]_{t,2} = \left(1 + I_{pt}^{\text{vib}}\right)\left(1 + V_{pt}^{\text{vib}}\right)\cos\phi_{pt}^{\text{vib}}\,, \tag{3.31}$$

$$\left[\boldsymbol{A}_p\right]_{t,3} = \left(1 + I_{pt}^{\text{vib}}\right)\left(1 + V_{pt}^{\text{vib}}\right)\sin\phi_{pt}^{\text{vib}}\,, \tag{3.32}$$

where $[\,\cdot\,]_{t,i}$ with $i = 1, 2, 3$ denotes row t and the three columns i of $\boldsymbol{A}_p$. The weighted least-squares problem is formulated the same as in (3.29) except that $\boldsymbol{A}_p$ now also depends on the pixel p:

$$\left(\boldsymbol{A}_p^\top \boldsymbol{W}_p \boldsymbol{A}_p\right) \boldsymbol{x}_p = \boldsymbol{A}_p^\top \boldsymbol{W}_p \hat{\boldsymbol{y}}_p\,. \tag{3.33}$$

For a given set of vibrations I_{pt}^{vib}, V_{pt}^{vib}, and ϕ_{pt}^{vib} we can now compute the resulting flat-field terms $\hat{I}_p$, $\hat{V}_p$, and $\hat{\phi}_p$ in an efficient, parallized manner. Constructing $\boldsymbol{A}_p$ and solving (3.33) for $\boldsymbol{x}_p$ is implemented in `em_jax.spr.lsq_all`.

3.2.2 Nested optimization of flat-fields and vibrations

When we process an air-scan, we initially set the vibration models $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ to 2D (tile-wise) polynomials as discussed before. But we neither know the flat-fields $\boldsymbol{I}$, $\boldsymbol{V}$, and $\boldsymbol{\phi}$, nor the vibration coefficients $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$. In principle, we could optimize some likelihood cost function wrt. all unknowns at once. Due to the model (3.2) being non-linear and non-convex we would most probably not find the global minimum this way—also known as (aka.) the physically correct solution. Instead, we split the unknowns into two groups, flat-fields $(\boldsymbol{I}, \boldsymbol{V}, \boldsymbol{\phi})$ and vibration coefficients $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$, and treat them separately during optimization.

"Separately" can mean different things in this context. One common method is alternating optimization: one minimizes the cost function only wrt. to one set of unknowns to convergence, while holding the other set(s) fixed. Then the sets are switched and the cost function is optimized wrt. the other set. While this method is generally stable and widely used, it has been shown that the convergence can be slow in our context [37]. Instead, we use "nested" optimization: we put in the vibration coefficients $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$, build $\boldsymbol{A}_p$ to apply SPR as in (3.33), and compute the residuum of the linearized solve for SPR. The optimization then minimizes this residuum wrt. the vibration coefficients, while the intermediate flat-fields are re-computed every iteration. A flowchart of this optimization principle is shown in Fig. 3.4.

More formally, for a given vibration model $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ the pixel-wise cost function L_p is defined as a weighted sum over the squared residuum after phase retrieval

$$L_p = (\boldsymbol{A}_p \boldsymbol{x}_p - \hat{\boldsymbol{y}}_p)^\top \boldsymbol{W}_p (\boldsymbol{A}_p \boldsymbol{x}_p - \hat{\boldsymbol{y}}_p)\,, \tag{3.34}$$

and the sum over all pixels p gives the total cost C

$$C(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}) = \sum_{p=1}^{P} m_p L_p(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})\,, \tag{3.35}$$

with m_p the grating mask. It is 0 outside of the active grating area and between G_2 tiles (see Fig. 3.9), and 1 otherwise. For a given fluctuation model $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ the function $C(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ is minimized with respect to the vibration parameters

$$(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) = \underset{\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma}}{\arg\min}\, C(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})\big|_{(\boldsymbol{\mathcal{A}},\boldsymbol{\mathcal{B}},\boldsymbol{\mathcal{C}})}\,, \tag{3.36}$$

giving the optimal coefficients $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$.

For the first-order gradient descent algorithm used for optimization (L-BFGS [51]), the gradient of C with respect to the temporal vibration parameters has to be computed. Of special interest is the derivative of L_p through the SPR. The gradient of L_p with respect to α_{kt} is computed by

$$\frac{\mathrm{d}L_p}{\mathrm{d}\alpha_{kt}} = 2\boldsymbol{x}_p^\top \boldsymbol{A}_p^\top \boldsymbol{W}_p \frac{\partial \boldsymbol{A}_p}{\partial \alpha_{kt}} \boldsymbol{x}_p - 2\hat{\boldsymbol{y}}_p^\top \boldsymbol{W}_p \frac{\partial \boldsymbol{A}_p}{\partial \alpha_{kt}} \boldsymbol{x}_p \tag{3.37}$$

$$= 2\left(\boldsymbol{A}_p \boldsymbol{x}_p - \hat{\boldsymbol{y}}_p\right)^\top \boldsymbol{W}_p \frac{\partial \boldsymbol{A}_p}{\partial \alpha_{kt}} \boldsymbol{x}_p\,, \tag{3.38}$$

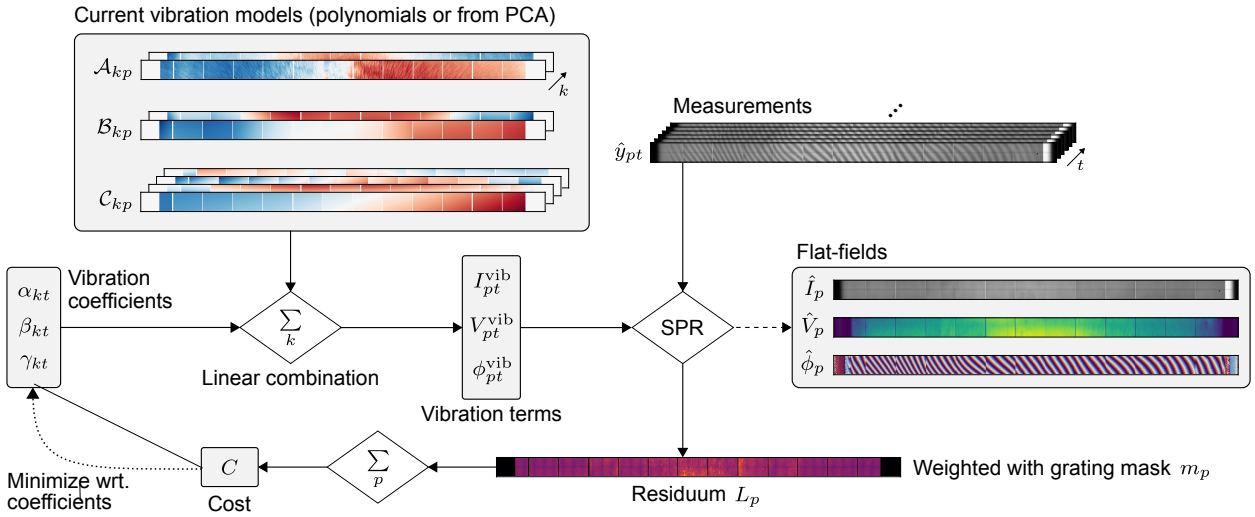

Figure 3.4: Flowchart of the nested optimization of vibration coefficients and flat-fields. The coefficients $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ are combined with the vibrations models $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ to form vibration terms $\boldsymbol{I}^{\text{vib}}$, $\boldsymbol{V}^{\text{vib}}$, and $\boldsymbol{\phi}^{\text{vib}}$. They are in turn combined with the measurements $\hat{\boldsymbol{y}}$ in SPR. The pixel-wise residuum L_p of the pixel-wise linearized solve is summed over pixels p, giving the cost function C. It it minimized wrt. the coefficients. The flat-fields $\hat{\boldsymbol{I}}$, $\hat{\boldsymbol{V}}$, and $\hat{\boldsymbol{\phi}}$ are not used during optimization. The vibration models $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ are not altered during optimization (but between optimizations).

as the solution vector $\boldsymbol{x}_p$ minimizes L_p by definition and $\partial L_p / \partial \boldsymbol{x}_p = 0$. Therefore the derivative $\partial \boldsymbol{x}_p / \partial \alpha_{kt}$ through the linearized solve for $\boldsymbol{x}_p$ is not necessary. The remaining derivative $\partial \boldsymbol{A}_p / \partial \alpha_{kt}$ is straightforward. This concept works equivalently for β_{kt} and γ_{kt}. It can be seen as an application of the implicit function theorem [52].

We implement the nested optimization around SPR in `JAX`. The fact that $\partial L_p / \partial \boldsymbol{x}_p = 0$ is expressed by using `jax.lax.stop_gradient` on the solution vector $\boldsymbol{x}_p$ after the linear solve:

```
@jax.jit
def get_cost_after_spr(A, b):
    x = jnp.linalg.solve(
        jnp.dot(A.T, A),
        jnp.dot(A.T, b)
    )
    x = jax.lax.stop_gradient(x)
    L = jnp.square(jnp.dot(A, x) - b)
    C = jnp.sum(L)
    return C
```

Here, `A` already contains the statistical weights $\boldsymbol{W}_p$. This code nicely illustrates that we only need to compute $\boldsymbol{x}_p$ for $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$, but do not need to differentiate through the process of computing $\boldsymbol{x}_p$. It results in a speedup of roughly 30 %.

3.2.3 Optimization strategy

To avoid running into local minima during the optimization, the cost function $C(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ is not optimized with respect to all vibrations immediately: we ignore most

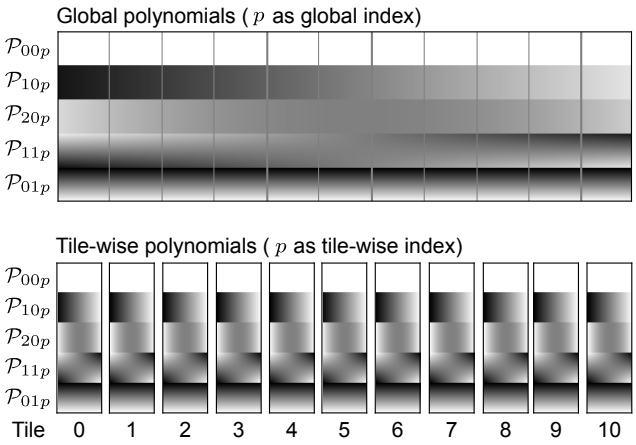

Figure 3.5: Illustration of global and tile-wise polynomials $\mathcal{P}_{ijp}$. The optimization pipeline starts with treating the whole G_2 as one interferometer and using global polynomials as vibration models. Later we allow tile-wise differences and use tile-wise polynomials.

vibration terms in the beginning and introduce the parameters successively as outlined below. All vibration coefficients are initialized with $\boldsymbol{\alpha} = \boldsymbol{\beta} = \boldsymbol{\gamma} = \boldsymbol{0}$. Before being introduced, a parameter is not included in the model.

We found the SPR to be most sensitive to a global offset in intensity and phase, as modeled by $\mathcal{P}_{00p}$ with $\mathcal{P}_{ijp}$ defined in (3.3). A convenient method to find approximations $\widetilde{\alpha}_{00t}$ and $\widetilde{\gamma}_{00t}$ for the corresponding coefficients is described in [53]. We briefly outline the method with our notation here.

We first approximate the flat-field $\boldsymbol{I}$ via the pixel-wise mean over all exposures,

$$\widetilde{I}_p = \left\langle \hat{y}_{pt} \right\rangle_t , \tag{3.39}$$

where $\left\langle \cdot \right\rangle_t$ is the pixel-wise mean over all exposures t. This implicitly assumes randomly distributed phase over t, such that

$$\left\langle \cos\left(\phi_p + \phi_{pt}^{\text{vib}}\right) \right\rangle_t \approx \mathbf{0} . \tag{3.40}$$

We approximate the flux $\widetilde{\alpha}_{00t}$ by the mean of the normalized intensities $\hat{y}_{pt}/\widetilde{I}_p$ over p:

$$\widetilde{\alpha}_{00t} = \left\langle \frac{\hat{y}_{pt}}{\widetilde{I}_p} \right\rangle_p . \tag{3.41}$$

This in turn requires that the fringes are sufficiently small, such that

$$\left\langle \cos\left(\phi_p + \phi_{pt}^{\text{vib}}\right) \right\rangle_p \approx \mathbf{0} . \tag{3.42}$$

We compute the global phase offsets $\widetilde{\gamma}_{00t}$ by PCA on the normalized intensities

$$\text{PCA}\left(\frac{\hat{\boldsymbol{y}}}{\widetilde{\boldsymbol{\alpha}}_{00}\widetilde{\boldsymbol{I}}}\right) \to (\boldsymbol{Y}, \boldsymbol{U}) , \tag{3.43}$$

where we could compute flat-fields V_p and ϕ_p from the principal components $\boldsymbol{Y}$, but don't because we employ nested optimization later which gives us the flat-fields (I_p, V_p, ϕ_p) anyway. The important output is $\boldsymbol{U} = (U_{kt})$ with which we calculate $\widetilde{\alpha}_{00t}$ via

$$\widetilde{\alpha}_{00t} = \arctan\left(\frac{U_{2,t}}{U_{1,t}}\right) , \tag{3.44}$$

using the first and second components of $\boldsymbol{U}$ per exposure t.

Afterward we start with the main processing pipeline. We initially approximate the whole G_2 as one single interferometer and introduce tile-wise terms only in the end. We first find global offsets (including global fluctuations in intensity), then global polynomials, and finally tile-wise polynomials in visibility and phase. After global visibility and phase polynomials, we compute an intermediate approximation of the intensity PCA model $\boldsymbol{\mathcal{A}}$. It contains slight fringe artifacts and has to be blurred, but is required to determine accurate tile-wise vibrations in visibility and phase afterward. In the final steps the PCA components on all channels are computed, including an un-blurred intensity PCA model.

In detail, the steps of the optimization pipeline are:

1. Compute $\widetilde{\alpha}_{00t}$ and $\widetilde{\gamma}_{00t}$ as described in [53].
2. Introduce global phase offset $(\gamma_{00t}, \mathcal{P}_{00p})$, initialized with $\widetilde{\gamma}_{00t}$. $\widetilde{\alpha}_{00t}$ is kept constant.
3. Introduce global visibility offset $(\beta_{00t}, \mathcal{P}_{00p})$.
4. Introduce global intensity offset $(\alpha_{00t}, \mathcal{P}_{00p})$ where α_{00t} is initialized with $\widetilde{\alpha}_{00t}$ from step 1.
5. Introduce global phase polynomials $(\gamma_{ijt}, \mathcal{P}_{ijp})$ with $i = 0, 1, 2$ and $j = 0, 1$ as in (3.3).
6. Introduce global visibility polynomials $(\beta_{ijt}, \mathcal{P}_{ijp})$ with $i = 0, \ldots, 4$ and $j = 0, 1, 2$.
7. Determine PCA model for spatial intensity fluctuations $(\boldsymbol{\mathcal{A}}, \boldsymbol{\alpha}^\star)$ as described in (3.20) and (3.21). Blur $\boldsymbol{\mathcal{A}}$ with a 2D Gaussian filter over the detector.
8. Introduce tile-wise phase polynomials $(\gamma_{ijt}, \mathcal{P}_{ijp})$ with $i = 0, 1, 2$ and $j = 0, 1$.
9. Introduce tile-wise visibility polynomials $(\beta_{ijt}, \mathcal{P}_{ijp})$ with $i = 0, 1, 2, 3, 4$ and $j = 0, 1, 2$.
10. Determine dominant vibrations in visibility $(\boldsymbol{\mathcal{B}}, \boldsymbol{\beta}^\star)$ and phase $(\boldsymbol{\mathcal{C}}, \boldsymbol{\gamma}^\star)$ via PCA as described in (3.13), (3.14), and (3.15).
11. Recompute PCA model for spatial intensity fluctuations $(\boldsymbol{\mathcal{A}}, \boldsymbol{\alpha}^\star)$.

Each of the steps (except 1, 7, 10 and 11) is a minimization of (3.36) with the L-BFGS algorithm [51] of the current fluctuation model coefficients to a certain small threshold $\varepsilon > 0$.

The optimization of the vibration coefficients $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ is performed via gradient-based minimization of $C(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ with the L-BFGS algorithm. The L-BFGS implementation used is from `scipy`. It requires double-precision floating point arrays on CPU, so we copy $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ from GPU to CPU and cast to `float64` for every iteration. Compared to the computational cost of computing the gradient in the first place this is cheap, because the number of elements to move around is small—we optimize *coefficients*, not *images*.

Orthogonality

Looking at (3.2) it becomes obvious that the three groups of coefficients do not contribute equally to C. Put differently, the Hessian $\boldsymbol{\mathcal{H}}_C$ of C is not a diagonal matrix, nor are its diagonal elements equally scaled. Ideally, this is covered by the approximated (inverse) Hessian in the L-BFGS solver. But we can make its task easier by (a) making sure $\boldsymbol{\mathcal{H}}_C$ is diagonal and (b) rescaling its diagonal.

(a) is achieved by using a vibration model with pairwise orthogonal (and therefore completely orthogonal) model components per channel. This makes the Hessian block-wise diagonal, because the vibration coefficients for one exposure are not correlated. For the models $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ from PCA, the components are

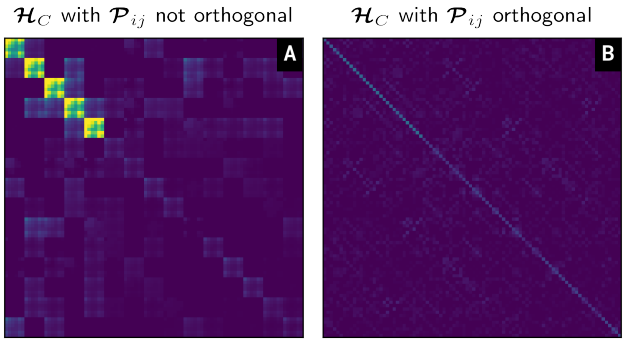

Figure 3.6: Hessian matrix $\boldsymbol{\mathcal{H}}_C \in \mathbb{R}^{60\times 60}$ wrt. the vibration coefficients of five exposures for 2D polynomials as vibration models in intensity, visibility, and phase. **A** The left-hand Hessian corresponds to non-orthogonal polynomials as defined in (3.3). One bright square represents the intensity coefficients α_{ijt} of one exposure t. The five squares in the upper-left therefore correspond to the intensity coefficients over five exposures, which are clearly correlated within an exposure and over exposures, and furthermore dominate in magnitude compared to the visibility and phase coefficients in the center and lower-right, respectively. **B** The right-hand Hessian corresponds to pair-wise orthogonal polynomials, which are also scaled per-channel with flat-fields $\hat{\boldsymbol{I}}$ and $\hat{\boldsymbol{V}}$, respectively. It is approximately diagonal, which is beneficial for convergence of C.

already pair-wise orthogonal by the definition of PCA. In contrast, the 2D polynomials $\boldsymbol{\mathcal{P}}_{ij}$ from (3.3) are not pair-wise orthogonal in their original definition. We can make them orthogonal over i, j with the Gram-Schmidt process [54],

$$\widetilde{\mathcal{P}}_{ijp} = \text{Gram-Schmidt}_{ij}\left(\mathcal{P}_{ijp}\right) , \tag{3.45}$$

such that $\widetilde{\boldsymbol{\mathcal{P}}}_{ij}$ are pair-wise orthogonal polynomials with $\langle \widetilde{\boldsymbol{\mathcal{P}}}_{ij}, \widetilde{\boldsymbol{\mathcal{P}}}_{kl} \rangle_p = 0$ if $i \neq k$ and $j \neq l$.

(b) is analogous to the simple pre-conditioning method of using the pre-computed inverse of the Hessian diagonal as a pre-conditioner. We approximate this by scaling the vibrations models per channel by the L2 norm over pixels of the respective component $\boldsymbol{\mathcal{P}}_{ij}$ or $\boldsymbol{\mathcal{A}}_k$. More formally, we define weighted L2-norms $\|\cdot\|_2^{\mathsf{I}}$ and $\|\cdot\|_2^{\mathsf{V}}$ for intensity and visibility,

$$\|\boldsymbol{x}\|_2^{\mathsf{I}} = \sqrt{\sum_{p=1}^{P} \left(m_p I_p x_p\right)^2} , \tag{3.46}$$

$$\|\boldsymbol{x}\|_2^{\mathsf{V}} = \sqrt{\sum_{p=1}^{P} \left(m_p I_p V_p x_p\right)^2} , \tag{3.47}$$

with placeholder variable $\boldsymbol{x} \in \mathbb{R}^P$ and grating mask $\boldsymbol{m}$. We compute scaled polynomials $\widetilde{\boldsymbol{\mathcal{P}}}_{ij}^{\mathsf{I}}$, $\widetilde{\boldsymbol{\mathcal{P}}}_{ij}^{\mathsf{V}}$, and $\widetilde{\boldsymbol{\mathcal{P}}}_{ij}^{\phi}$ in intensity, visibility, and phase,

$$\widetilde{\boldsymbol{\mathcal{P}}}_{ij}^{\mathsf{I}} = \frac{\boldsymbol{\mathcal{P}}_{ij}}{\|\boldsymbol{\mathcal{P}}_{ij}\|_2^{\mathsf{I}}} , \tag{3.48}$$

$$\widetilde{\boldsymbol{\mathcal{P}}}_{ij}^{\mathsf{V}} = \widetilde{\boldsymbol{\mathcal{P}}}_{ij}^{\phi} = \frac{\boldsymbol{\mathcal{P}}_{ij}}{\|\boldsymbol{\mathcal{P}}_{ij}\|_2^{\mathsf{V}}} , \tag{3.49}$$

and use them instead of $\boldsymbol{\mathcal{P}}_{ij}$ from (3.3) during air-scan processing. The application of PCA is not altered by these polynomials. They merely improve convergence of the cost function C, but if anything converge to the same vibration terms $\hat{\boldsymbol{V}}^{\text{vib}}$ and $\hat{\boldsymbol{\phi}}^{\text{vib}}$, which are then put into PCA. Strictly speaking, we only use $\boldsymbol{\mathcal{P}}_{00}$ in the specified air-scan processing pipeline.

The vibration models $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ from PCA are already pair-wise orthogonal in k, but not scaled wrt. the flat-fields. We define scaled PCA models,

$$\widetilde{\boldsymbol{\mathcal{A}}}_k = \frac{\boldsymbol{\mathcal{A}}_k}{\|\boldsymbol{\mathcal{A}}_k\|_2^{\mathsf{I}}} , \tag{3.50}$$

$$\widetilde{\boldsymbol{\mathcal{B}}}_k = \frac{\boldsymbol{\mathcal{B}}_k}{\|\boldsymbol{\mathcal{B}}_k\|_2^{\mathsf{V}}} , \tag{3.51}$$

$$\widetilde{\boldsymbol{\mathcal{C}}}_k = \frac{\boldsymbol{\mathcal{C}}_k}{\|\boldsymbol{\mathcal{C}}_k\|_2^{\mathsf{V}}} , \tag{3.52}$$

and use them mainly in sample-scan processing with the final estimation for flat-fields $\hat{\boldsymbol{I}}$ and $\hat{\boldsymbol{V}}$ from the air-scan.

3.3 Further perturbations

Additionally to the modeled vibrations, we identified additional deviations in the DFCT data from the conventional model of a stepped Talbot-Lau interferometer. We employ the main processing pipeline discussed until now and use the flat-fields $(\hat{\boldsymbol{I}}, \hat{\boldsymbol{V}}, \hat{\boldsymbol{\phi}})$ and vibration terms $(\hat{\boldsymbol{I}}^{\text{vib}}, \hat{\boldsymbol{V}}^{\text{vib}}, \hat{\boldsymbol{\phi}}^{\text{vib}})$ to determine the remaining perturbations discussed in the following.

We therefore implicitly assume that we can correctly determine the flat-fields and vibrations terms *before* identifying the perturbations discussed here. The general motivation is that these effects have zero mean over a full axial scan and do not negatively influence the air-scan processing. They do however negatively influence the image quality when processing a sample-scan, which involves phase-retrieval in small angular windows, in which the effects have non-zero mean. Simply put, we have to determine the perturbations in the air-scan, even though we need them only for sample-scan image quality.

3.3.1 Higher orders of visibility

As discussed in section 2.2.4 the model of a Talbot-Lau interferometer is a Fourier series, where only the first two terms a_0 and a_1 are used in practice. The

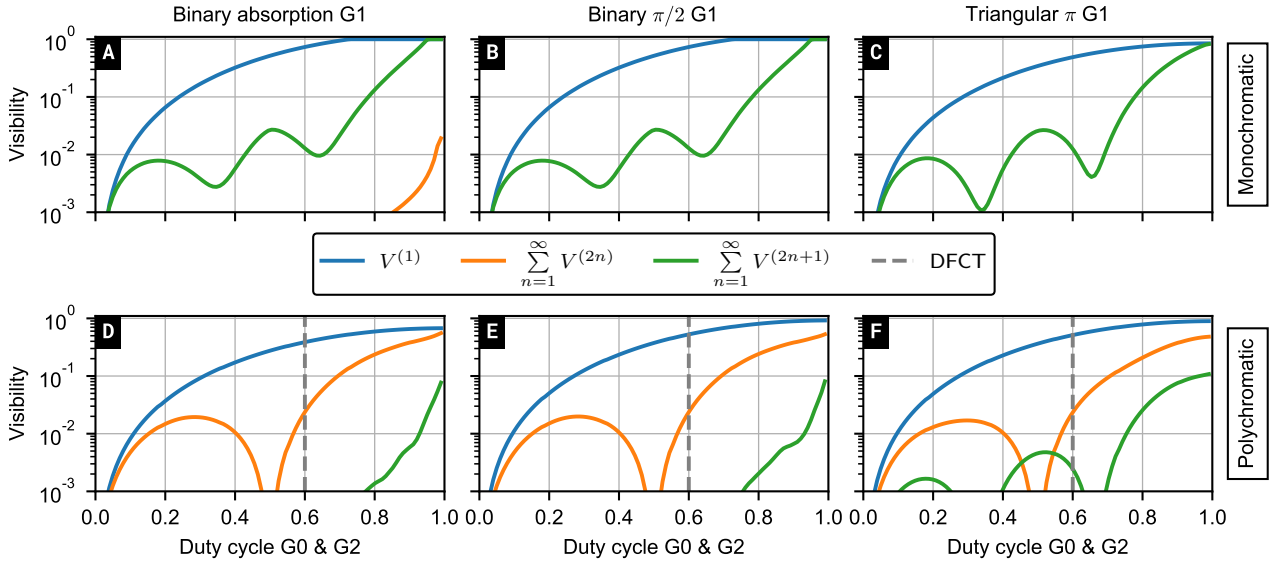

Figure 3.7: Simulation of visibility orders depending on the G_0 and G_2 duty cycle for different G_1 configurations. The G_1 duty cycle is 50 % in all simulations. Orange lines represent the sum of higher visibilities with even order, and green lines the sum of odd orders. **A**-**C** The top row depicts monochromatic simulations with the gratings in the respective optimal fractional Talbot distance. Even orders do not contribute to the visibility. **D**-**F** The bottom row shows polychromatic simulations with a 80 kVp Tungsten spectrum and 40 keV design energy. Even orders do contribute to the visibility, but not at duty cycle 50 %. **F** The triangular G_1 has the strongest odd higher orders. The gray dashed line represents the DFCT duty cycle of approximately 60 %.

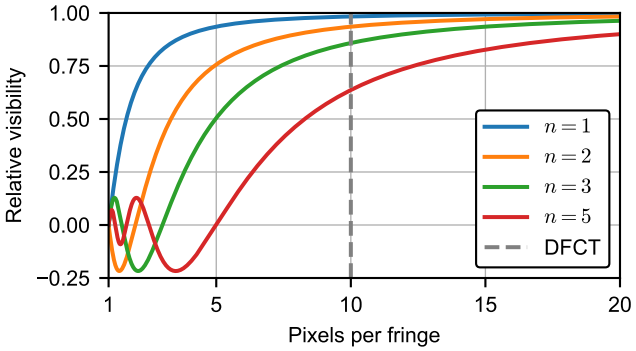

Figure 3.8: Relative visibility decrease depending on number of pixels per fringe and visibility order n. The visibility generally decreases with finer fringes. Higher orders are affected much more by this effect, because their fringes are smaller. The DFCT is tuned to roughly ten pixels per fringe which barely affects the first-order visibility. Negative visibilities imply aliasing and a lateral phase shift of π.

higher orders a_2, a_3, etc. get smaller very quickly and are typically ignored as they contribute very little to the overall visibility contrast.

The DFCT has a phase-shifting G_1 with a triangular profile. It has been shown [21] that such a triangular G_1 has stronger higher visibility (i.e. Fourier series) orders than a G_1 with a binary profile. These are not included in our model (3.2) so far.

Bech [55] showed that the stepping curve has only odd terms if the analyzer grating G_2 has 50 % duty cycle. More generally, the even order Fourier terms of any arbitrary function f are removed by a convolution with a symmetric G_0 or G_2, whose even order Fourier series terms are zero.

The magnitude of the terms a_n decreases rapidly with $n \geq 2$ mainly depending on the size of the G_0 and G_2 slits [28]. For symmetric G_0 and G_2, i.e. a duty cycle of 50 %, $a_n = 0$ if n is even. A simulation comparing the visibility orders for different duty cycles is shown in Fig. 3.7. The DFCT has G_0 and G_2 with duty cycle of approximately 60 % [21], leading to non-zero even visibility terms. Note that the derivation in [28] is valid only for binary G_1 profiles, not for triangular phase-shifting gratings presented in [21] and used in the DFCT.

The DFCT has a G_2 with 60 % and G_0 with 56 % duty cycle [21]. This leads to non-zero (albeit small) visibilities of even order as seen in Fig. 3.7. Experimentally, we observe and use the second order $V^{(2)}$ and third order $V^{(3)}$.

Due to the additive structure of the model including higher-order visibilities (3.55), visibilities $V^{(n)}$ of order $n \geq 2$ do not influence the result of first-order SPR (as in (3.33)) if the sampling points of the stepping curve, i.e. γ_t in the conventional model (3.1) or ϕ_{pt}^{vib} in (3.2), cover at least $(0, 2\pi)$. Instead, they end up in the residuum after the (linear) fit of the model to the data.

This means that the air-scan processing does not have to take the higher-order visibilities into account to

obtain accurate results for flat-fields $(\boldsymbol{I}, \boldsymbol{V}, \boldsymbol{\phi})$ and vibration coefficients $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$. It applies per-pixel SPR on $T = 2400$ sampling points covering $(0, 10\pi)$ (see section 3.4.2) leading to a robust fit of the first-order model. Furthermore, even with our triangular G_1, the higher orders are still comparatively small (see Fig. 3.10). Only for small window sizes in sliding-window SPR they become noticeable as artifacts in the dark-field D. We discuss these artifacts in section 3.4 and chapter 4. Our method of determining the higher visibility orders of the DFCT in the air-scan processing is discussed here.

The visibility fluctuation $\boldsymbol{V}^{\text{vib}}$ also influences higher-order visibility terms. To quantify the effect, we extend the notion of continuously changing phase $\boldsymbol{\phi}^{\text{vib}}(\tau)$ as in (3.7) to higher orders n of the phase term inside the cosine,

$$\begin{aligned}\frac{1}{\Delta t}\int_{t-\frac{\Delta t}{2}}^{t+\frac{\Delta t}{2}} \cos\left(n\left(\phi_p + \phi_p^{\text{vib}}(\tau)\right)\right)\,\mathrm{d}\tau \\ \approx \left(1 - n^2\frac{\Delta t^2}{24}\left(\dot{\phi}_{pt}^{\text{vib}}\right)^2\right)\cos\left(n\left(\phi_p + \phi_{pt}^{\text{vib}}\right)\right)\,,\end{aligned} \tag{3.53}$$

where we again approximate the cosine by its Taylor expansion to second order in τ. The structure of (3.53) is very similar to (3.7), except for the added factor n^2 for the visibility drop. We therefore assume the spatio-temporal fluctuations of the higher-order visibility terms $\boldsymbol{V}^{\text{vib},(n)}$ to be approximated by the weighted first-order fluctuations $\boldsymbol{V}^{\text{vib}}$,

$$V_{pt}^{\text{vib},(n)} = n^2 V_{pt}^{\text{vib}}\,, \tag{3.54}$$

and extend the air-scan model (3.2) to include higher-order visibility terms,

$$\begin{aligned}y_{pt}^{\text{higher}} = I_p\left(1 + I_{pt}^{\text{vib}}\right)\Bigg[1 + \\ + \sum_{n=1}^{N} V_p^{(n)}\left(1 + n^2 V_{pt}^{\text{vib}}\right)\times \\ \times\cos\left(n\left(\phi_p^{(n)} + \phi_{pt}^{\text{vib}}\right)\right)\Bigg]\,,\end{aligned} \tag{3.55}$$

with the number of visibility orders N and the a priori unknown higher-order visibilities $\boldsymbol{V}^{(n)}$ and phases $\boldsymbol{\phi}^{(n)}$. This model has the same additive structure as (3.2) and we can perform a linearized least-squares solve for $\boldsymbol{I}$, $\boldsymbol{V}^{(n)}$, and $\boldsymbol{\phi}^{(n)}$ in a very similar manner. We define a pixel-wise system matrix $\boldsymbol{H}_p \in \mathbb{R}^{T\times(1+2N)}$ as

$$\begin{aligned}\left[\boldsymbol{H}_p\right]_{t,1} &= 1 + I_{pt}^{\text{vib}}\,, \\ \left[\boldsymbol{H}_p\right]_{t,2n} &= \left(1 + I_{pt}^{\text{vib}}\right)\left(1 + n^2 V_{pt}^{\text{vib}}\right)\cos\left(n\phi_{pt}^{\text{vib}}\right)\,, \\ \left[\boldsymbol{H}_p\right]_{t,2n+1} &= \left(1 + I_{pt}^{\text{vib}}\right)\left(1 + n^2 V_{pt}^{\text{vib}}\right)\sin\left(n\phi_{pt}^{\text{vib}}\right)\,,\end{aligned} \tag{3.56}$$

for row t and $n \in \{1, ..., N\}$. Therefore $\boldsymbol{H}_p$ has $1 + 2N$ columns: one for the intensity term and two per visibility order. The pixel-wise weighted least-squares problem is formulated as

$$\left(\boldsymbol{H}_p^\top \boldsymbol{W}_p \boldsymbol{H}_p\right)\boldsymbol{x}_p = \boldsymbol{H}_p^\top \boldsymbol{W}_p \hat{\boldsymbol{y}}_p\,, \tag{3.57}$$

with the statistical weights $\boldsymbol{W}_p$ from (3.33). (3.57) is solved for $\boldsymbol{x}_p$ via matrix decomposition, giving the solution vector $\hat{\boldsymbol{x}}_p \in \mathbb{R}^{1+2N}$. The solution terms $\hat{I}_p$, $\hat{V}_p^{(n)}$, and $\hat{\phi}_p^{(n)}$ are computed from the solution vector $\hat{\boldsymbol{x}}_p$,

$$\hat{I}_p = \left[\hat{\boldsymbol{x}}_p\right]_1\,, \tag{3.58}$$

$$\hat{V}_p^{(n)} = {}^{1}\!/_{\hat{I}_p}\sqrt{\left[\hat{\boldsymbol{x}}_p\right]_{2n}^2 + \left[\hat{\boldsymbol{x}}_p\right]_{2n+1}^2}\,, \tag{3.59}$$

$$\hat{\phi}_p^{(n)} = \operatorname{arctan2}\left(-\left[\hat{\boldsymbol{x}}_p\right]_{2n+1}, \left[\hat{\boldsymbol{x}}_p\right]_{2n}\right)\,. \tag{3.60}$$

We formulate the model (3.55) with independent phases $\boldsymbol{\phi}^{(n)}$. The (linear) solve of the model for $\hat{\boldsymbol{\phi}}^{(n)}$ therefore gives numerically independent phases. Physically, we do expect a relation between them,

$$\phi_p^{(n)} = n\phi_p^{(1)}\,, \tag{3.61}$$

because the model (2.2.4) is motivated from the Fourier series description of the stepping curve in section 2.2.4 and all orders work on the same "macroscopic" interferometer phase. However, we can perform the described linear solve (3.57) only if we treat $\boldsymbol{\phi}^{(n)}$ separately. Even though this is suboptimal from a physical point of view, it gives us the opportunity to verify the physical sensibility of our model: we can compare the obtained phase orders $\hat{\boldsymbol{\phi}}^{(n)}$ and check whether they satisfy (3.61). The results for the DFCT are presented in section 3.4.1.

Concerning the empirical processing routine, we first perform the air-scan processing pipeline as discussed in section 3.2.3 and ignore higher-order terms. Then we find the higher-order terms $\boldsymbol{V}^{(n)}$ and $\boldsymbol{\phi}^{(n)}$ as described here. The higher-order model (3.55) is never used "online" in the nested optimization discussed in section 3.2.2, but only once in the very end to determine the higher-order terms, given the previously determined vibration terms $\hat{\boldsymbol{I}}^{\text{vib}}$, $\hat{\boldsymbol{V}}^{\text{vib}}$, and $\hat{\boldsymbol{\phi}}^{\text{vib}}$.

3.3.2 Macroscopic grating movement

The principle of image formation with a Talbot-Lau interferometer is to detect angular deviations much

smaller than a detector pixel. In explicit phase stepping, one of the gratings is moved in small increments over one period of that grating. The grating movement is therefore by design much smaller than the spatial resolution of the detector. This is the reason why we can usually assume the flat-fields $\boldsymbol{I}$, $\boldsymbol{V}$, and $\boldsymbol{\phi}$ as constant, because the projection of the *macroscopic* grating structures onto the detector does not change during a measurement. In contrast, the vibration model from the previous sections (and the phase modulation in conventional setups) is assumed to be caused by *microscopic* movements of the gratings.

As we will see in section 3.4.4 and Fig. 3.15 **C**, the global phase γ_{00} moves over approximately 10π or five grating periods during one rotation of the gantry in the DFCT. The grating period geometrically magnified onto the detector is the same for all gratings and we can use the period p_2 of the analyzer grating G_2 for further discussion. In the DFCT we have $p_2 = 45\,\mu$m, equating to a lateral shift of the macroscopic grating structure of $5 \times 45 = 255\,\mu$m. The width of the (non-square) detector pixels is approximately 1.41 mm. The projection of the gratings therefore changes by $0.255/1.41 \approx 15\,\%$ of a pixel. The usual assumption that the grating movement is much smaller than a detector pixel does not hold in the DFCT.

As long as the expected projected lateral movement is smaller than one pixel, the change of the flat-fields can be modeled by adding their respective spatial derivatives,

$$\widetilde{\boldsymbol{I}} \approx \boldsymbol{I} + a_u \left(\partial_u \boldsymbol{I}\right) + a_v \left(\partial_v \boldsymbol{I}\right) , \tag{3.62}$$

$$\widetilde{\boldsymbol{V}} \approx \boldsymbol{V} + a_u \left(\partial_u \boldsymbol{V}\right) + a_v \left(\partial_v \boldsymbol{V}\right) , \tag{3.63}$$

$$\widetilde{\boldsymbol{\phi}} \approx \boldsymbol{\phi} + a_u \left(\partial_u \boldsymbol{\phi}\right) + a_v \left(\partial_v \boldsymbol{\phi}\right) , \tag{3.64}$$

where $\widetilde{\boldsymbol{I}}$, $\widetilde{\boldsymbol{V}}$, and $\widetilde{\boldsymbol{\phi}}$ are the effective projected flat-field terms, $a_u, a_v \in \mathbb{R}$ with $-1 < a_u, a_v < 1$ are the horizontal and vertical macroscopic shifts in pixel fractions, and ∂_u, ∂_v denote the horizontal and vertical spatial gradient, respectively. We omit index t in (3.62) to (3.64) for readability, but the factors a_u and a_v do change per exposure. Equations (3.62) to (3.64) can be interpreted as performing a Taylor expansion to first order, where we approximate the first derivative as the spatial gradient. We combine these displacement terms into movement terms I_{pt}^{mov}, V_{pt}^{mov}, and ϕ_{pt}^{mov} (all $\in \mathbb{R}^{P\times T}$) and formulate an adapted model,

$$\begin{aligned} y_{pt}^{\text{mov}} = \left(I_p + I_{pt}^{\text{mov}}\right)\left(1 + I_{pt}^{\text{vib}}\right)\Big[1 + \\ + \left(V_p + V_{pt}^{\text{mov}}\right)\left(1 + V_{pt}^{\text{vib}}\right) \times \\ \times \cos\left(\phi_p + \phi_{pt}^{\text{mov}} + \phi_{pt}^{\text{vib}}\right)\Big] , \end{aligned} \tag{3.65}$$

with vibration terms I_{pt}^{vib}, V_{pt}^{vib}, and ϕ_{pt}^{vib} from (3.21), (3.14), and (3.15).

Modeling $\boldsymbol{I}^{\text{mov}}$, $\boldsymbol{V}^{\text{mov}}$, and $\boldsymbol{\phi}^{\text{mov}}$ directly over the subpixel shifts a_u and a_v is impractical. The observed lateral shift of the flat-fields is a compound effect of the individual grating movements. The macroscopic structures of the gratings contribute differently to the several flat-fields and it is therefore not guaranteed that the *observed* lateral shift is identical in intensity, visibility, and phase. Put differently, the modalities might each have their own a_u and a_v.

Instead, we empirically determine the observed effect from measured air-scan data. We perform SPR on angular subsets of the full axial air-scan with the previously determined vibration terms $(I_{pt}^{\text{vib}}, V_{pt}^{\text{vib}}, \phi_{pt}^{\text{vib}})$ incorporated into a system matrix $\boldsymbol{M}_{pt}$, such that the resulting terms constitute the effective flat-field terms from (3.65). We define the system matrix $\boldsymbol{M}_{pt} \in \mathbb{R}^{\mathcal{M}\times 3}$,

$$\begin{aligned} \left[\boldsymbol{M}_{pt}\right]_{\tau,1} &= 1 + I_{p\tau}^{\text{vib}} , \\ \left[\boldsymbol{M}_{pt}\right]_{\tau,2} &= \left(1 + I_{p\tau}^{\text{vib}}\right)\left(1 + V_{p\tau}^{\text{vib}}\right)\cos\phi_{p\tau}^{\text{vib}} , \\ \left[\boldsymbol{M}_{pt}\right]_{\tau,3} &= \left(1 + I_{p\tau}^{\text{vib}}\right)\left(1 + V_{p\tau}^{\text{vib}}\right)\sin\phi_{p\tau}^{\text{vib}} , \end{aligned} \tag{3.66}$$

with $\mathcal{M}$ the number of adjacent exposures in one angular window around exposure t and $\tau \in \text{window}(t)$ the index inside that window. Compared to $\boldsymbol{A}_p$ discussed in section 3.2.1, the system matrix $\boldsymbol{M}_{pt}$ also depends on the exposure t, because its index τ is always centered around it.

For a more elaborate discussion of this so-called "sliding-window phase-retrieval", see section 4.2.1. In short, it is a method of determining approximations $\hat{\boldsymbol{I}}^{\text{mov}}$, $\hat{\boldsymbol{V}}^{\text{mov}}$, and $\hat{\boldsymbol{\phi}}^{\text{mov}}$ which have high spatial resolution (pixel-wise) but lower temporal resolution (determined by the angular window size $\mathcal{M}$).

We formulate the according least-squares problem similarly to section 3.2.1 as

$$\left(\boldsymbol{M}_{pt}^{\top}\boldsymbol{W}_{pt}\boldsymbol{M}_{pt}\right)\boldsymbol{x}_{pt} = \boldsymbol{M}_{pt}^{\top}\boldsymbol{W}_{pt}\hat{\boldsymbol{y}}_{pt}^{(1)} . \tag{3.67}$$

The data vector $\hat{\boldsymbol{y}}_{pt}^{(1)} \in \mathbb{R}^{\mathcal{M}}$ represents the measured data $\hat{\boldsymbol{y}}$ around exposure t with the higher-order visibility terms $n > 1$ subtracted, such that only the first-order visibility remains. The weights $\boldsymbol{W}_{pt}$ are computed accordingly. The movement terms $\boldsymbol{I}^{\text{mov}}$, $\boldsymbol{V}^{\text{mov}}$, and $\boldsymbol{\phi}^{\text{mov}}$ are computed from the solution vector $\hat{\boldsymbol{x}}_{pt}$,

$$\hat{I}_{pt}^{\text{mov}} = \left[\hat{\boldsymbol{x}}_{pt}\right]_1 - \hat{I}_p , \tag{3.68}$$

$$\hat{V}_{pt}^{\text{mov}} = \frac{1}{\hat{I}_{pt}^{\text{mov}}}\sqrt{\left[\hat{\boldsymbol{x}}_{pt}\right]_2^2 + \left[\hat{\boldsymbol{x}}_{pt}\right]_3^2} - \hat{V}_p , \tag{3.69}$$

$$\hat{\phi}_{pt}^{\text{mov}} = \arctan2\left(-\left[\hat{\boldsymbol{x}}_{pt}\right]_3, \left[\hat{\boldsymbol{x}}_{pt}\right]_2\right) - \hat{\phi}_p , \tag{3.70}$$

with the previously determined flat-fields $\hat{\boldsymbol{I}}$, $\hat{\boldsymbol{V}}$, and $\hat{\boldsymbol{\phi}}$. We determine the dominant components of $\boldsymbol{I}^{\text{mov}}$,

$\boldsymbol{V}^{\text{mov}}$, and $\boldsymbol{\phi}^{\text{mov}}$ with PCA,

$$\text{PCA}\left(\boldsymbol{I}^{\text{mov}}\right) \rightarrow \left(\boldsymbol{R}^{\text{I}}, \boldsymbol{r}^{\text{I}}\right) , \tag{3.71}$$

$$\text{PCA}\left(\boldsymbol{V}^{\text{mov}}\right) \rightarrow \left(\boldsymbol{R}^{\text{V}}, \boldsymbol{r}^{\text{V}}\right) , \tag{3.72}$$

$$\text{PCA}\left(\boldsymbol{\phi}^{\text{mov}}\right) \rightarrow \left(\boldsymbol{R}^{\phi}, \boldsymbol{r}^{\phi}\right) , \tag{3.73}$$

and approximate the movement terms in (3.65) with the first two components respectively,

$$\hat{I}^{\text{mov}}_{pt} \approx \sum_{k=1}^{2} r^{\text{I}}_{kt} R^{\text{I}}_{kp} , \tag{3.74}$$

$$\hat{V}^{\text{mov}}_{pt} \approx \sum_{k=1}^{2} r^{\text{V}}_{kt} R^{\text{V}}_{kp} , \tag{3.75}$$

$$\hat{\phi}^{\text{mov}}_{pt} \approx \sum_{k=1}^{2} r^{\phi}_{kt} R^{\phi}_{kp} , \tag{3.76}$$

as we expect the change being caused by lateral movement along height and width of the detector, modeled by (3.62) to (3.64).

In sample-scan processing, we use $\boldsymbol{R}^{\text{I}}$, $\boldsymbol{R}^{\text{V}}$, and $\boldsymbol{R}^{\phi}$ determined as described above. The coefficients $\boldsymbol{r}^{\text{I}}$, $\boldsymbol{r}^{\text{V}}$, and $\boldsymbol{r}^{\phi}$ have to be re-determined. We can either perform a likelihood fit similar to the vibration coefficients, or just use the low-frequency component from the air-scan coefficients, because we expect the movement to be correlated with the gantry movement and therefore comparatively slow anyway. Results for $\boldsymbol{R}^{\text{I}}$, $\boldsymbol{R}^{\text{V}}$, and $\boldsymbol{R}^{\phi}$ and the influence of movement correction on reconstructions is shown in section 3.4.3.

We can estimate the effect of the flat-field movement terms on the sample projections after sliding-window phase-retrieval in transmission T_{pt}, D_{pt}, and Φ_{pt}. For that we assume that the erroneous projections $\widetilde{T}_{pt}$, $\widetilde{D}_{pt}$, and $\widetilde{\Phi}_{pt}$ obtained *without* the movement constitute the same total signal as the corrected projections,

$$\widetilde{T}_{pt} I_p \stackrel{!}{=} T_{pt}\left(I_p + I^{\text{mov}}_{pt}\right) , \tag{3.77}$$

$$\widetilde{D}_{pt} V_p \stackrel{!}{=} D_{pt}\left(V_p + V^{\text{mov}}_{pt}\right) , \tag{3.78}$$

$$\widetilde{\Phi}_{pt} + \phi_p \stackrel{!}{=} \Phi_{pt} + \phi_p + \phi^{\text{mov}}_{pt} , \tag{3.79}$$

and obtain the following relations,

$$\widetilde{T}_{pt} = T_{pt}\left(1 + \frac{I^{\text{mov}}_{pt}}{I_p}\right) , \tag{3.80}$$

$$\widetilde{D}_{pt} = D_{pt}\left(1 + \frac{V^{\text{mov}}_{pt}}{V_p}\right) , \tag{3.81}$$

$$\widetilde{\Phi}_{pt} = \Phi_{pt} + \phi^{\text{mov}}_{pt} . \tag{3.82}$$

In phase, the error is simply additive. In transmission and dark-field, the macroscopic movement creates an error in the projections proportional to the movement itself, inversely weighted with the respective flat-field. The error is therefore larger in areas with small reference intensity or visibility, respectively. The intensity flat-field $\hat{\boldsymbol{I}}$ in the DFCT (see Fig. 3.9) is relatively homogeneous. This makes the expected error in transmission small for two reasons: a) The intensity has no falloff at the borders of the interferometer, and b) the spatial gradients $\partial_u \hat{\boldsymbol{I}}$ and $\partial_v \hat{\boldsymbol{I}}$ are small as well, which determine the magnitude of the movement error as of (3.62). In visibility, we do expect a larger error at the far left- and right-hand side of the interferometer field of view (FOV), because the visibility magnitude tapers off there (also see Fig. 3.9), and at the same time the spatial gradient along the width of the detector $\partial_u \hat{\boldsymbol{V}}$ is large. This can be circumvented by generously cropping the projections used for reconstruction, if the sample is not too large. Alternatively, optimizing the grating setup for homogeneous visibility lessens this problem.

3.4 Experimental results

We discuss the results determined with the optimization method discussed in the previous sections. An axial air-scan is used as measurement data $\hat{\boldsymbol{y}}$. First, we show the flat-field terms $\hat{\boldsymbol{I}}$, $\hat{\boldsymbol{V}}^{(n)}$, and $\hat{\boldsymbol{\phi}}^{(n)}$. They characterize the interferometer and are of primary interest in conventional grating-based phase-contrast imaging, at least for $n = 1$. Then, we discuss the fluctuation models $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$. They exhibit different global and tile-wise behavior depending on the image channel.

To show the importance of this adapted model, we employ a simple sample-processing method using a global polynomial model (spanning all G_2 tiles) on the one hand, and our determined PCA model $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ on the other hand. We process both another air-scan and a measurement of a test phantom. In the same comparison we show the impact of previously discussed higher-order visibility terms and the macroscopic movement of flat-fields.

3.4.1 Flat-fields

We first show the resulting flat-fields $\hat{\boldsymbol{I}}$, $\hat{\boldsymbol{V}}^{(n)}$, and $\hat{\boldsymbol{\phi}}^{(n)}$ of an air-scan in the DFCT. They are presented in Fig. 3.9 together with an exposure of the raw data $\hat{\boldsymbol{y}}$ and the grating mask $\boldsymbol{m}$ indicating where a pixel represents usable interferometer area and therefore contributes to the likelihood C in (3.35).

The intensity $\hat{\boldsymbol{I}}$ is relatively homogeneous over the width of the detector. Only between G_2 tiles and on the far left- and right-hand side of the detector (outside the active grating area) it is noticeably larger. This homogeneity is an effect of the bent gratings, avoiding shadowing artifacts as discussed in [21]. A remaining

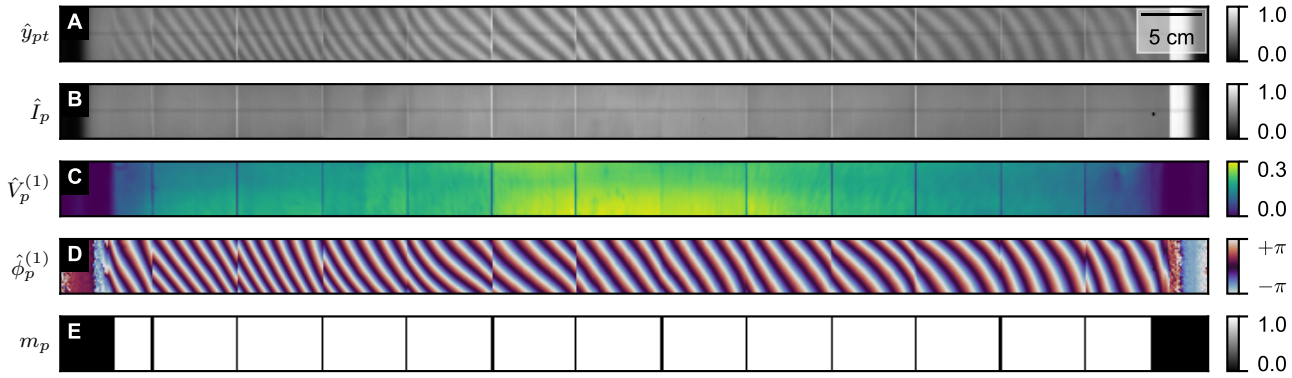

Figure 3.9: DFCT flat-fields after processing an air-scan. **A** shows one exposure of the raw data $\hat{\boldsymbol{y}}$. The fringes and G_2 tiles are clearly visible. The tile borders manifest as bright vertical lines. **B** shows the intensity profile $\hat{\boldsymbol{I}}$. It is relatively homogeneous. Only the tile borders and especially the grating-free area on the far right are brighter than the G_2 area. **C** depicts the (first-order) visibility $\hat{\boldsymbol{V}}^{(1)}$. Similarly to the intensity it peaks in the center with values around 30 %. It is very low at tile borders and obviously 0 outside the grating area at the far left and right. **D** shows the (first-order) phase pattern $\hat{\boldsymbol{\phi}}^{(1)}$. It matches the fringes and is tuned to 4-6 vertical fringes per tile. **E** depicts the grating mask $\boldsymbol{m}$. It is 0 outside the grating area and at tile borders, and 1 otherwise.

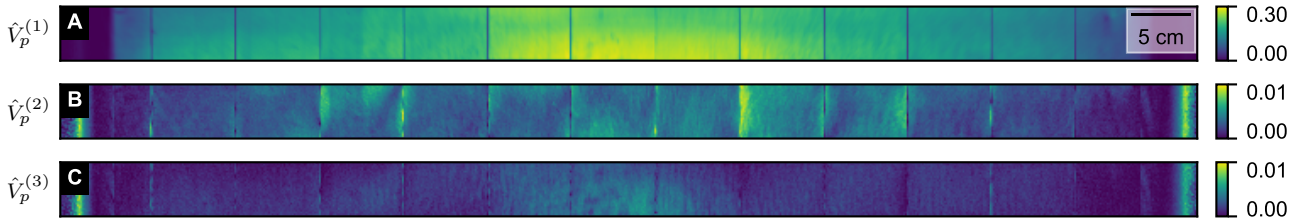

Figure 3.10: DFCT visibility in **A** first, **B** second, and **C** third order. The first-order visibility $\hat{\boldsymbol{V}}^{(1)}$ reaches a maximum of 30 % in the center of the detector. Both the second and third order visibilities $\hat{\boldsymbol{V}}^{(2)}$ and $\hat{\boldsymbol{V}}^{(3)}$ are much smaller in overall magnitude and are below 1 % almost everywhere. $\hat{\boldsymbol{V}}^{(2)}$ is largest at G_2 borders. $\hat{\boldsymbol{V}}^{(3)}$ has a similar trend as $\hat{\boldsymbol{V}}^{(1)}$ and is largest in the center.

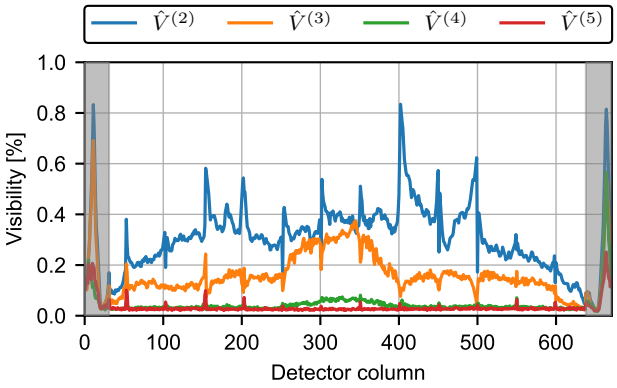

Figure 3.11: Lineplots of DFCT higher-order visibilities along the width of the detector. The visibility magnitude decreases with increasing order. Gray areas indicate invalid values from outside the interferometer FOV.

feature is a horizontal line with slightly less intensity. It is an effect of the detector and does not influence our processing further.

The first-order visibility $\hat{\boldsymbol{V}}^{(1)}$ peaks at 30 % and shows a stronger horizontal gradient, being large in the center and smaller towards the edges. Even still, compared to unbent gratings the profile is homogeneous [21]. Between G_2 tiles and outside of the active grating area on the far left- and right-hand side it is close to zero.

The pattern of the first-order phase $\hat{\boldsymbol{\phi}}^{(1)}$ resembles the fringes on the raw data $\hat{\boldsymbol{y}}$. The interferometer is tuned to four to six vertical fringes per G_2 tile. This amounts to approximately ten pixels per fringe. According to Fig. 3.8 this fringe size has only a very small negative effect on the visibility.

The grating mask $\boldsymbol{m}$ is zero on the far left- and right-hand side of the detector, and between G_2 tiles. Theoretically, our model should also be valid outside of the grating area, as the visibility is simply zero there. But with very low visibility comes very strong sensitivity to noise, and it has proven useful to explicitly exclude such areas from the vibration coefficient optimization by masking them. As the adjustment of G_2 tiles is an elaborate process, some grating borders consist of two-pixel columns which show up as thick black lines in $\boldsymbol{m}$. These lead to circular artifacts in the reconstructed volumes because we do not employ any kind of typewriter shift (moving detector, sample, or focal

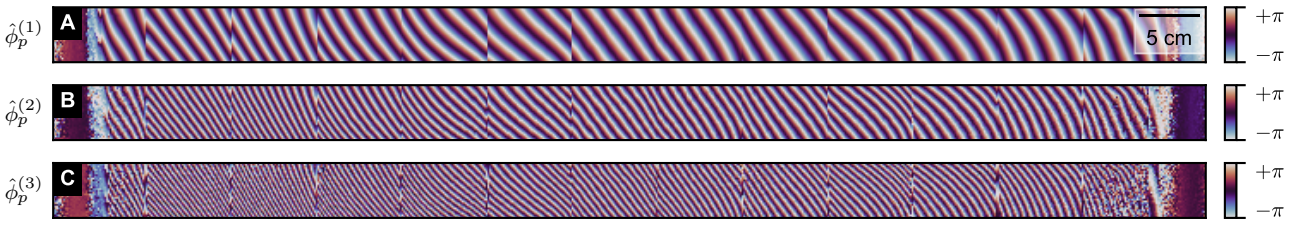

Figure 3.12: DFCT phase in **A** first, **B** second, and **C** third order. The first-order phase $\hat{\phi}^{(1)}$ visually matches the fringes of the raw data $\hat{y}$ seen in Fig. 3.9. The higher orders $\hat{\phi}^{(2)}$ and $\hat{\phi}^{(3)}$ have double and triple the spatial frequency, respectively.

spot between exposures). We discuss them in chapter 4 and see that these artifacts are not noticeable after ring correction [56].

We show higher orders of visibility $\hat{\boldsymbol{V}}^{(n)}$ and phase $\hat{\boldsymbol{\phi}}^{(n)}$ in Fig. 3.10 and Fig. 3.12 for $n = 1, 2, 3$, respectively. The higher-order visibilities $\hat{\boldsymbol{V}}^{(2)}$ and $\hat{\boldsymbol{V}}^{(3)}$ in Fig. 3.10 are much smaller than the first order $\hat{\boldsymbol{V}}^{(1)}$ and below 1 % almost everywhere. $\hat{\boldsymbol{V}}^{(2)}$ seems to be largest at borders of G_2 tiles. This could be caused by the mounting brackets deforming the gratings slightly and changing their effective duty cycle at the edges. $\hat{\boldsymbol{V}}^{(3)}$ is even smaller overall, and has a similar horizontal trend as $\hat{\boldsymbol{V}}^{(1)}$ in that it is largest in the center. According lineplots for the horizontal visibility profiles are shown in Fig. 3.11. Higher-order visibilities up to $n = 10$ are shown in Fig. A.1 and are basically zero.

The phase patterns $\hat{\boldsymbol{\phi}}^{(2)}$ and $\hat{\boldsymbol{\phi}}^{(3)}$ in Fig. 3.12 have double and triple the spatial frequency as $\hat{\boldsymbol{\phi}}^{(1)}$, respectively. This matches our expectation from (3.61). As discussed in section 3.3.1, we expect decreasing visibility magnitude with increasing order both from the convolution of the Talbot carpet with G_0 and G_2, as well as because of the fringe period becoming smaller than one detector pixel. We see in Fig. 3.12 that the third-order fringes are already barely larger than one pixel. It is therefore not surprising that for even higher orders the determined phase patterns $\hat{\boldsymbol{\phi}}^{(n)}$ are only noise, as seen in Fig. A.2.

As stated earlier, we only use $\hat{\boldsymbol{V}}^{(2)}$ and $\hat{\boldsymbol{V}}^{(3)}$ for the higher-order correction during sample-scan processing. It is discussed in section 4.3.3.

3.4.2 Vibration model

We discuss the dominant spatial fluctuations $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ determined by PCA. The number of components $A^\star$ and $B^\star$ in intensity and visibility (as introduced in (3.21) and (3.14)) to keep for processing a sample-scan is chosen by plotting the singular values determined during PCA sorted by their magnitude for each channel. These so-called "scree plots" [57] are shown in Fig. 3.13. The magnitude of the singular values corresponds to the magnitude of the principal directions

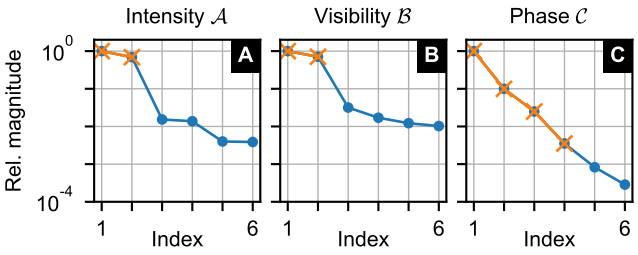

Figure 3.13: Scree plots of PCA model components for **A** intensity, **B** visibility, and **C** phase. In **A** and **B** there are two dominant components, after which the singular values drop to almost zero. Therefore we choose $A^\star = B^\star = 2$. In **C** the decrease is more gradual and the number $C^\star = 4$ is determined manually based on the quality of the reconstructed images. The magnitudes of the singular values correspond to the magnitudes of model coefficients shown in Fig. 3.15.

(i.e. vibration coefficients) shown in Fig. 3.15. The factors before the first "knee" of the scree plots are considered important and $A^\star$ and $B^\star$ are assigned to this number in the respective channel. With this method we determine $A^\star = B^\star = 2$. The principal components $k > 2$ diminish very quickly and therefore contribute only very little to $\boldsymbol{I}^{\text{vib}}$ and $\boldsymbol{V}^{\text{vib}}$, respectively. The scree plot for the phase channel does not show a pronounced first knee. $C^\star$ is empirically chosen by first determining $A^\star$ and $B^\star$ and then increasing $C^\star$ until no visual difference in the reconstructions is perceived. Here, we determine $C^\star = 4$.

In principle, the number of model components $A^\star$, $B^\star$, and $C^\star$ can be determined and updated after each processed air-scan. While we do update the models $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ after each air-scan, the number of model components in the DFCT has been constant for the duration of this thesis.

Principal components

The dominant principal components forming the vibration models in intensity, visibility, and phase are shown in Fig. 3.14. In the intensity channel **A**,**B**, the first two principal components are dominated by large gradients showing a distinctive left-right and outer-inner trend, respectively. We assume they originate from the shadowing caused by the drifting of the focal spot [21].

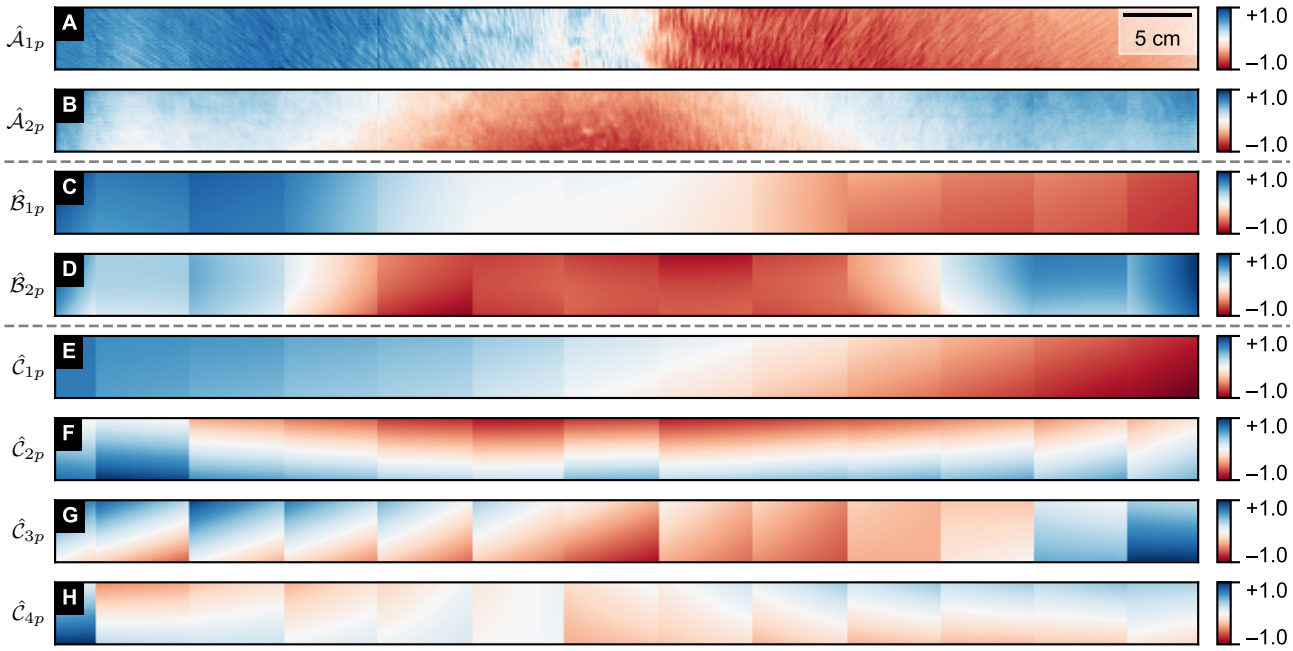

Figure 3.14: The most dominant principal components of the spatial fluctuations in **A**,**B** intensity, **C**,**D** visibility, and **E**-**H** phase. All components are scaled to [-1, 1] only for visualization. The vibrations models are determined only in the active grating area. **A**,**B** The intensity components are a combination of large gradients over the whole detector and small interferometer structures, respectively. Both **C**,**D** visibility and **E**-**H** phase components emanate from a tile-wise polynomial model and can not express small structures. While the visibility vibration consists of mainly global gradients, the phase components also show distinct tile-wise properties. The images shown are cut to the valid interferometer area.

Additionally, the components show fine details which are vertical in the center and more diagonal on the outside of the field of view (FOV). They may originate from macroscopic structures in the gratings from bending them onto the cylindrical geometry [58].

The two visibility components **C**,**D** are also dominated by a large trend over the width of the detector. They resemble the intensity vibrations in that the first factor is a horizontal ramp with a flat region in the center and the second factor has an inner-outer trend. As the visibility fluctuation is modeled with tile-wise polynomials, the components do not show fine details like the intensity fluctuations. One can still identify tile-to-tile differences, like the bright tile second from the right or the details in tile corners.

The vibrations in the phase channel **E**-**H** exhibit the most distinct tile-wise behavior. The first two principal components **E**,**F** are smooth ramps over the whole detector with only the right-most tile slightly brighter. They capture the movement of the G_2 carrier. The third and fourth components **G**,**H** however have only tile-wise characteristics and show no discernible global trend.

Model coefficients

The determined "principal directions" aka. model coefficients $\hat{\alpha}^\star$, $\hat{\beta}^\star$, and $\hat{\gamma}^\star$ are shown in Fig. 3.15. The global terms $\hat{\alpha}_{00}$, $\hat{\beta}_{00}$, and $\hat{\gamma}_{00}$ dominate in all modalities. We note again that these global terms could be incorporated directly in the PCA models, but it has proven useful in practice to discern between an artificial "global" fluctuation and the spatial vibrations determined in PCA.

All coefficients in all modalities are comprised of a low-frequency component (correlated to the gantry position) and few high-frequency components originating from vibrating subsystems in the DFCT, like the rotating anode and the cooling pump [22]. This becomes apparent in the temporal Fourier spectra shown in Fig. 3.16. They are very sparse with high-frequency peaks at 176 Hz, 189 Hz, and 2×176 Hz, depending on the channel. Apart from the global phase $\hat{\gamma}_{00}$ at 176 Hz and corresponding global visibility drop $\hat{\beta}_{00}$ at 2×176 Hz, all other coefficients are mainly at 189 Hz.

This work does little more than acknowledge the existence of these few dominant frequencies. Our team has shown that the amplitude and frequency of the vibrations are essential to the quality of the images after signal retrieval [59]. In contrast, here, we assume the DFCT as-is with its given vibrations, and aim to optimize the processing and reconstruction with them.

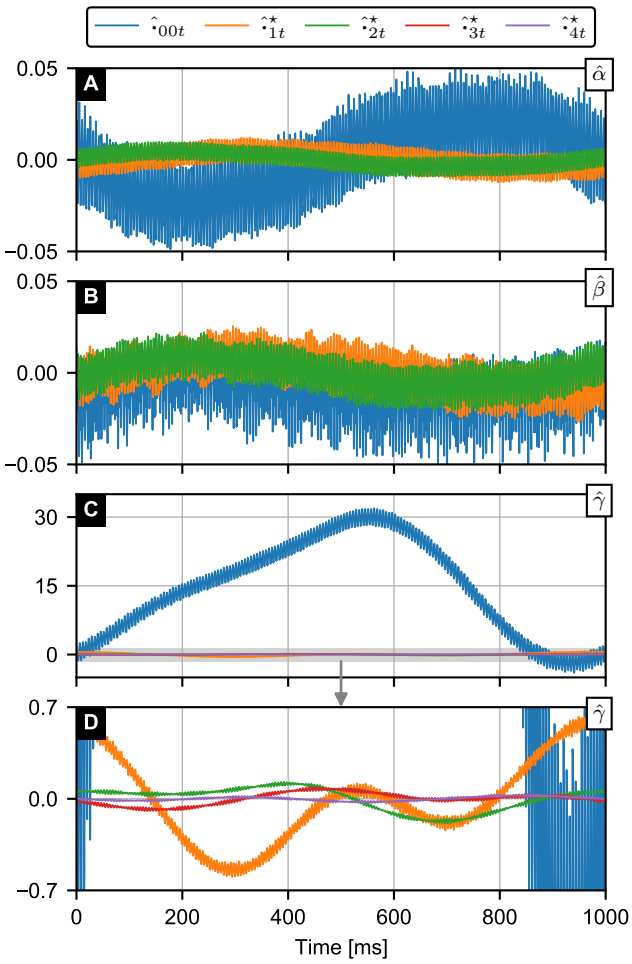

Figure 3.15: Plot of per-shot coefficients of PCA model in **A** intensity $\hat{\boldsymbol{\alpha}}^{\star}$, **B** visibility $\hat{\boldsymbol{\beta}}^{\star}$, and **C**,**D** phase $\hat{\boldsymbol{\gamma}}^{\star}$. **A** The intensity is mostly influenced by the global flux variation $\hat{\alpha}_{00}$. **B** The visibility reduction from the varying phase per exposure is captured mostly by the global term $\hat{\beta}_{00}$. **C** The global phase $\hat{\gamma}_{00}$ varies over $10\,\pi$ per rotation, i.e five fringe periods. The influence of the PCA model components is much smaller. The high-frequency oscillation of $\hat{\gamma}_{00}$ has a range of approximately $1.5\,\pi$.

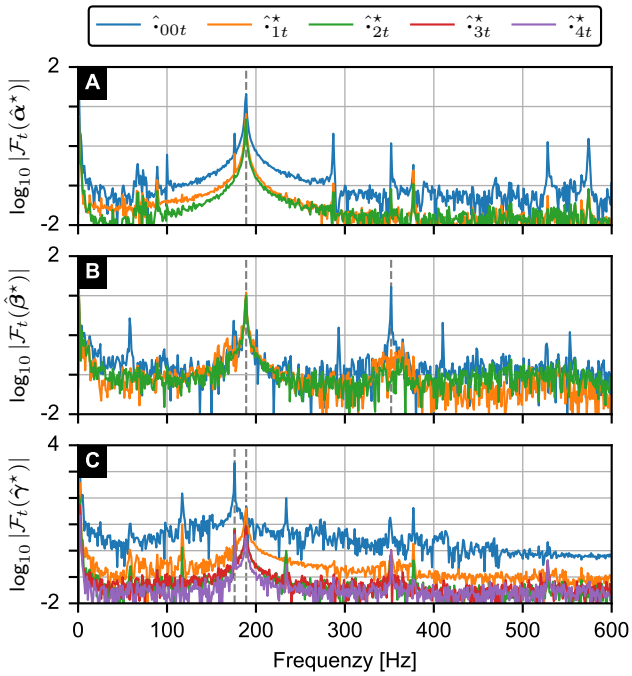

Figure 3.16: Plot of Fourier spectra of per-shot coefficients of PCA model in **A** intensity $\hat{\boldsymbol{\alpha}}^{\star}$, **B** visibility $\hat{\boldsymbol{\beta}}^{\star}$, and **C** phase $\hat{\boldsymbol{\gamma}}^{\star}$. It is apparent that all vibration coefficients have sparse frequency spectra, i.e. consist of a small number of oscillations. The dominant frequencies are 189 Hz in intensity, 189 Hz and 2×176 Hz in visibility, and 176 Hz and 189 Hz in phase.

3.4.3 Movement correction

The components $\hat{\boldsymbol{R}}^{\mathsf{I}}$, $\hat{\boldsymbol{R}}^{\mathsf{V}}$, and $\hat{\boldsymbol{R}}^{\phi}$ for the macroscopic movement model of section 3.3.2 are determined on the same axial air-scan in the DFCT and presented in Fig. 3.17. For visualization, they are min/max-scaled to [-1, 1]. During processing, they are instead scaled to have the same L2 norm (analogously to the PCA vibration model components) for better convergence, as described in section 3.2.3. We use two model components per channel, because we motivate the perturbance from macroscopic grating movement to be similar to the spatial gradients of the respective flat-fields. The components in Fig. 3.17 do indeed mimic the spatial gradients of the flat-fields of the respective channel. All components show vertical grating structures mainly in the center, which are probably caused by cracks in the substrate after bending [22]. The second components in visibility $\hat{\boldsymbol{R}}^{\mathsf{V}}_2$ and phase $\hat{\boldsymbol{R}}^{\phi}_2$ contain a distinct blob-like feature on the far right. It matches the local drop in visibility in the first-order visibility flat-field $\hat{\boldsymbol{V}}^{(1)}$ shown in Fig. 3.9. The intensity components $\hat{\boldsymbol{R}}^{\mathsf{I}}_2$ are modulated with slight, high-frequency fringe artifacts. This is likely due to the movement being modeled in part by the vibration model in intensity $\boldsymbol{\mathcal{A}}$, as it is determined directly via PCA on the intermediate relative residuum during the processing pipeline. The remaining movement terms should be fairly small and apparently approach the contrast of very slight remaining fringe artifacts.

This assumption is validated by the visualization of the error from movement terms in Fig. 3.18. It depicts the mean of the absolute value of the deviation from ideal sample projections, motivated in section 3.3.2. We compute the mean of the absolute value (as opposed to the value itself), because the terms have approximately zero mean over all exposures from an air-scan. The images in Fig. 3.18 naturally resemble the model components shown in Fig. 3.17.
The error in transmission $\langle|\hat{\boldsymbol{I}}^{\mathsf{mov}}/\hat{\boldsymbol{I}}|\rangle$ is roughly one order of magnitude smaller than in dark-field and phase. This is not surprising, as a) the intensity flat-field $\hat{\boldsymbol{I}}$ is comparatively homogeneous and its spatial gradient

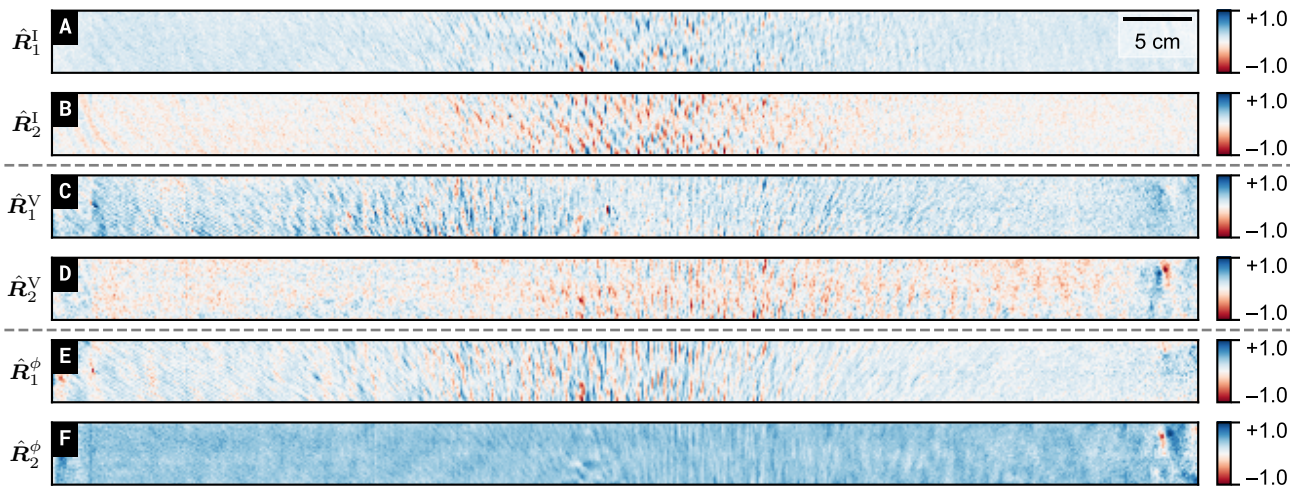

Figure 3.17: Components of macroscopic flat-field movement in **A**,**B** intensity, **C**,**D** visibility, and **E**,**F** phase. All components are min/max-scaled to [-1, 1] for visualization. Although empirically obtained by sliding-window SPR on an air-scan, they do resemble the spatial gradients of the respective flat-fields and in that show macroscopic grating features. **A**,**B** In intensity, slight residual fringes are visible. The images shown are cut to the valid interferometer area.

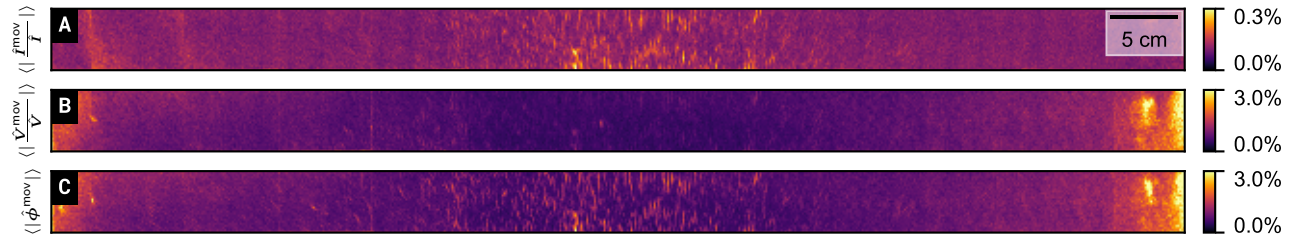

Figure 3.18: Absolute value of projection error from flat-field movement, averaged over exposures t, in **A** transmission, **B** dark-field, and **C** differential phase. All terms show a similar crack structure in the center. As expected, it mimics the spatial gradient of the respective flat-field $\hat{\boldsymbol{I}}$, $\hat{\boldsymbol{V}}$, or $\hat{\phi}$. **A** The magnitude of the transmission error is smaller by roughly one order of magnitude. **B** The dark-field is affected mostly on the far left and right because of the small visibility. **C** The phase error has stronger features in the center than dark-field, but similar large values at the borders.

therefore small, and b) the perturbation that does occur is in part covered by the vibration model in intensity, which is computed by PCA on the intermediate model residuum.
The dark-field error is dominated by large values at the far left- and right-hand side of the inner FOV, increasing from the center. These are explained by the low first-order visibility $\hat{\boldsymbol{V}}^{(1)}$ at the interferometer borders, as seen in Fig. 3.9. The distinct blob-like structure near the right border also appears in the dark-field error.
The error in the phase channel is generally the largest. Similar to the dark-field error, its magnitude increases at the borders of the interferometer. But in contrast to the dark-field, it also shows crack-like features in the center. Ultimately, the projection error caused by macroscopic flat-field movement is largest in the phase channel.

3.4.4 Impact on reconstructed images

We show the relevance of the proposed reference processing pipeline by comparing reconstruction results using the PCA model $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ from Fig. 3.14 to global polynomial fluctuations $\boldsymbol{\mathcal{P}}$ (as defined in (3.3)) in all channels. The processing pipeline used for the sample-scans is described in section 4.2.2. The processing pipeline for the polynomial basis is effectively steps 1-6 from section 3.2.3 and the PCA model for intensity in step 7 replaced with global polynomials $\mathcal{P}_{ijp}$ with $i, j = 0, 1, 2$.
In addition, we show the effect of including higher-order visibility terms and macroscopic flat-field movement as discussed in section 3.3.

Even though this chapter is concerned with processing air-scans, we have to look at reconstuctions of sample-scans to appreciate the relevance of an accurate vibration model. To obtain accurate flat-fields $(\hat{\boldsymbol{I}}, \hat{\boldsymbol{V}}, \hat{\phi})$, we would not need tile-wise polynomials or high-resolution intensity fluctuations from PCA. Over one axial scan of 2400 projections, the mean of the

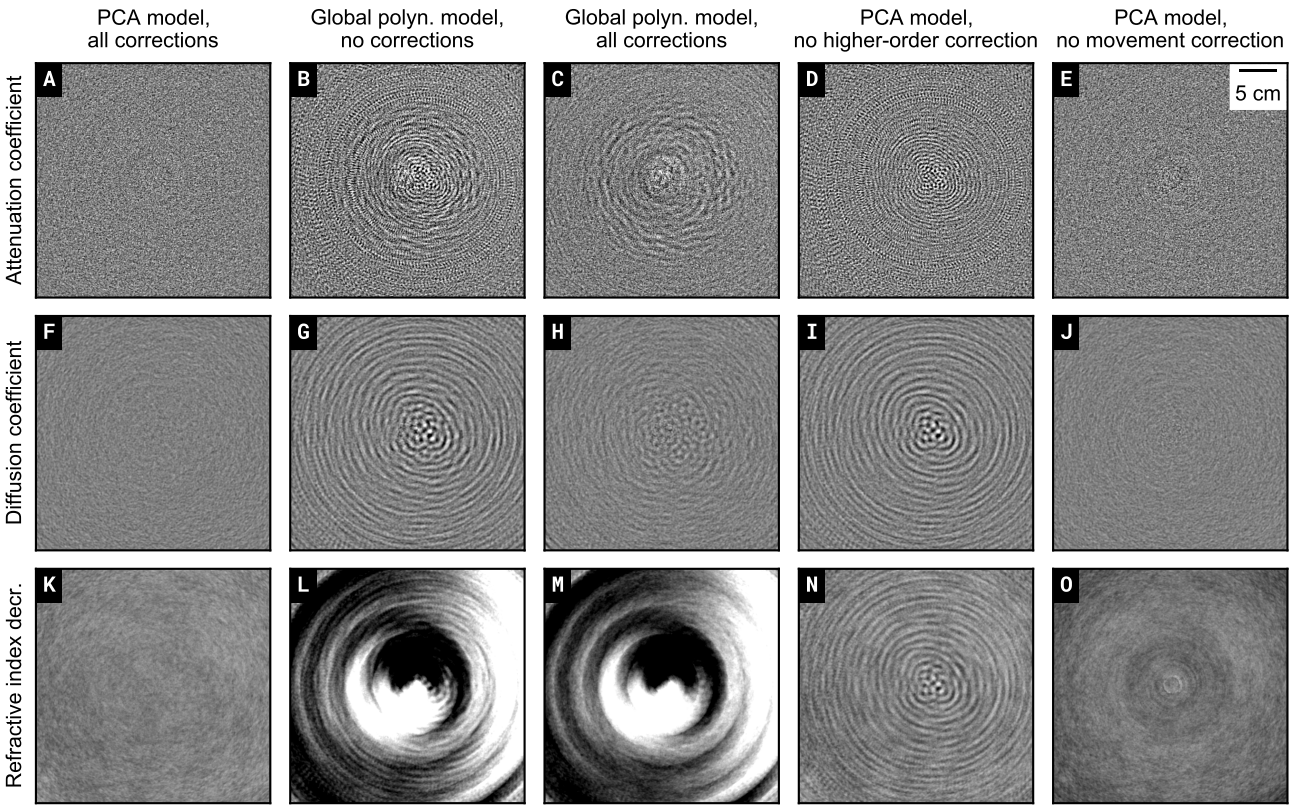

Figure 3.19: Comparison of air-scan reconstructions for different vibration models, higher-order correction and flat-field movement correction in attenuation coefficient (top), diffusion coefficient (middle) and refractive index decrement (bottom). The left-most column **A**,**F**,**K** shows reconstructed slices with the PCA model and both higher-order and macroscopic movement correction applied, i.e. the presented method. The second column **B**,**G**,**L** shows the same slices with the most conventional methods of a global polynomial vibration model and neither corrections. Columns three to five show the slices with one of the presented techniques omitted, respectively. The windowing is **A**-**E**: [-3, 3] $\times$ 10^{-3} cm^{-1}; **F**-**J**: [-3, 3] $\times$ 10^{-2} cm^{-1}; **K**-**O**: [-2, 2] $\times$ 10^{-1} cm^{-1}.

vibrations is roughly 1 or 0, respectively, and they can therefore be ignored in (3.2).

In contrast, accurate vibrations are important when we apply SPR on a small set of consecutive exposures. This is the case for sample-scan reconstruction, where we perform sliding-window phase-retrieval. Therefore, it is appropriate to compare sample-scan reconstructions to assess the air-scan processing.

Air-scan

We look at the reconstruction of an air-scan as a homogeneous, known "sample". The full detector area can be used to determine vibration parameters $(\boldsymbol{\alpha}^{\mathrm{s}}, \boldsymbol{\beta}^{\mathrm{s}}, \boldsymbol{\gamma}^{\mathrm{s}})$ at the time of the "sample"-scan. This allows us to analyze the importance of an accurate vibration model $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ without the need for a sophisticated sample-scan processing pipeline. Two directly subsequent scans are performed with no sample in the beam path. One scan consists of 2400 exposures over a full 360° gantry rotation which takes 1 s. The X-ray tube is operated at 80 kVp and 550 mA.

The second scan is used with the proposed reference processing pipeline from section 3.2.3 to extract the PCA vibration model shown in Fig. 3.14, the higher-order visibility terms $\hat{\boldsymbol{V}}^{(2)}$ and $\hat{\boldsymbol{V}}^{(3)}$, and the movement models $\hat{\boldsymbol{R}}^{\mathrm{I}}$, $\hat{\boldsymbol{R}}^{\mathrm{V}}$, and $\hat{\boldsymbol{R}}^{\phi}$. The first scan is processed and reconstructed with the pipeline in chapter 4 as if it was a sample-scan. We use eleven subsequent angular views per sliding-window. The duration between scans was 30 seconds.

The results are compared with a simplified pipeline in which fluctuations in all channels are only modeled with polynomials $\mathcal{P}_{ijp}$ as in (3.3) up to second order along the width and height of the whole detector. In terms of the air-scan processing in section 3.2.3, this equates to performing steps 1-6 with $i = j = 0, 1, 2$ in visibility and phase. Also in intensity, polynomials $\boldsymbol{\mathcal{P}}$ replace the dominant components $\boldsymbol{\mathcal{A}}$ from PCA. For sample-scan processing and specifically equation (4.5), the same global polynomial vibration model is used in intensity, visibility, and phase. Therefore,

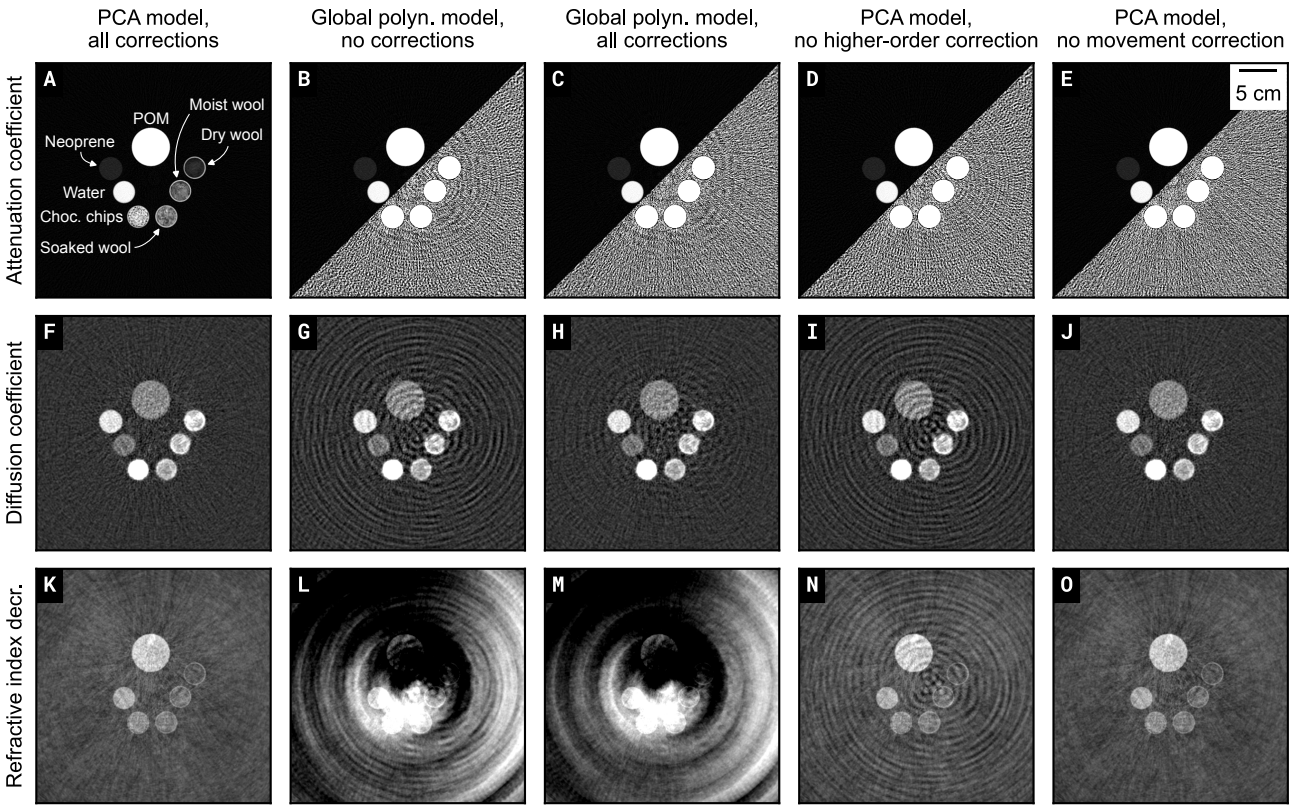

Figure 3.20: Comparison of sample-scan reconstructions for different vibration models, higher-order correction and flat-field movement correction in attenuation coefficient (top), diffusion coefficient (middle) and refractive index decrement (bottom). The left-most column **A**,**F**,**K** shows reconstructed slices with the PCA model and both higher-order and macroscopic movement correction applied, i.e. the presented method. The second column **B**,**G**,**L** shows the same slices with the most conventional methods of a global polynomial vibration model and neither corrections. Columns three to five show the slices with one of the presented techniques omitted, respectively. The windowing is **A**-**E**: [-0.05, 2.2] $\times$ 10^{-1} cm^{-1}; **F**-**J**: [-2, 8] $\times$ 10^{-2} cm^{-1}; **K**-**O**: [-2, 4] $\times$ 10^{-1} cm^{-1}. Narrow window in **B**-**E**: [-5, 5] $\times$ 10^{-3} cm^{-1}.

this simplified pipeline does not handle tile-wise vibrations nor the intricate details of the intensity fluctuations. Even still, it covers all established methods discussed in section 2.3.2.

We show the central slices of the resulting volumes in linear attenuation coefficient $\hat{\mu}$, linear diffusion coefficient $\hat{\varepsilon}$, and refractive index decrement $\hat{\delta}$ in Fig. 3.19. The left-most column **A**,**F**,**K** represents the complete presented pipeline, i.e. using the PCA vibration model, higher-order visibility terms, and macroscopic flat-field movement. All image modalities are homogeneous, showing no distinct artifacts. The noise pattern is different however, with the attenuation **A** showing high-frequency (white) noise, while diffusion **F** and refractive index **K** show rather patchy, low-frequency noise.
The second-from-left column **B**,**G**,**L** is the result of the most conventional processing pipeline. We apply the simple vibration model of global 2D polynomials in all image modalities, and regard neither higher-order visibilities nor flat-field movement. The reconstructions therefore suffer from strong, somewhat circular artifacts in all channels. Again, the attenuation **B** exhibits features with spatially rather high frequency. The diffusion **G** looks similar, albeit with features of approximately half the spatial frequency. The refractive index **L** is dominated by broad ring-like structures, while there are some finer bands visible.
The following columns isolate the omission of one novel technique presented in this work, respectively.
The central column **C**,**H**,**M** shows the artifacts generated by using a naive vibration model. In attenuation **C** only lower-frequency, circular artifacts remain. Their apparent magnitude is slightly lower than the total error in **B**. The diffusion **H** is least affected by the naive vibration model, exhibiting similar features to **C** but smaller in magnitude. In the refractive index **M** the dominating ring structures remain and are therefore apparently caused by incorrect vibration terms, namely tile-wise phase fluctuations.
The second-from-right column **D**,**I**,**N** isolates the error from the higher-order visibility terms. Somewhat unexpectedly, but as explained before, the small higher-order terms create significant artifacts in sliding-window phase retrieval and the corresponding reconstructed images. In attenuation **D** they are responsible for the high-frequency circular features. Apparently,

the error in transmission has double the frequency of the fringe pattern. In contrast, these "fringe artifacts" in diffusion **I** and refractive index **N** have once the fringe frequency, and look very overall very similar to each other. In diffusion **I**, they pose a larger error than vibrations in **H**, while in refractive index the vibrations in **M** are much more significant.
Finally, the right-most column **E**,**J**,**O** visualizes the impact of macroscopic flat-field movement. In both attenuation **E** and diffusion **J** the effect is barely noticeable despite the narrow window. Only the refractive index **O** exhibits thin circular structures mainly in the center. They correspond to the expected projection error in **C**.

Overall, the artifacts resulting from naive processing of the DFCT data are different in magnitude and structure, depending on image modality and effect. The attenuation coefficient is most affected by both the vibration model and the higher-order visibility terms. In diffusion, the higher-order terms dominate. The refractive index decrement shows the largest error overall from all modalities, which in turn is generated mostly by an incorrect vibration model, resulting in broad tile-wise bands.

As mentioned previously, we do employ a ring removal [56] on the reconstructed images. This is necessary to avoid sharp concentric ring artifacts in all channels originating from invalid sample projections at G_2 tile borders (as seen in Fig. 3.9). No such artifacts are visible in Fig. 3.19.

Rod phantom

To assess the impact of the discussed image artifacts compared to the sample contrast, we conduct the same comparisons for the scan of a test object. It consists of a POM cylinder of 5 cm diameter and six Falcon tubes with 3 cm diameter each, filled with wool at three different levels of dampness, chocolate chips, water, and neoprene, respectively. Again a subsequent air-scan is performed to obtain the flat-fields $(\hat{\boldsymbol{I}}, \hat{\boldsymbol{V}}, \hat{\phi})$, the dominant vibration components $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$, the higher-order visibility terms $\hat{\boldsymbol{V}}^{(2)}$ and $\hat{\boldsymbol{V}}^{(3)}$, and the movement models $\hat{\boldsymbol{R}}^{\mathsf{I}}$, $\hat{\boldsymbol{R}}^{\mathsf{V}}$, and $\hat{\boldsymbol{R}}^{\phi}$. The duration between scans was 30 seconds. All scan parameters are identical to the air-scan reconstruction. The vibration coefficients are determined by fitting the model $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ in the sample-free area on the far left- and right-hand side of the interferometer.

The central slices of the reconstructions of linear attenuation coefficient $\hat{\mu}$, linear diffusion coefficient $\hat{\varepsilon}$, and refractive index decrement $\hat{\delta}$ are shown in Fig. 3.20. The artifacts in all channels are similar to the air-scan reconstruction in Fig. 3.19. The windowing is adapted to the full sample contrast in each modality, and therefore enables the comparison of noise and artifact magnitudes.

The left-most column **A**,**F**,**K** shows the reconstructed images with all novel techniques employed and represents the state-of-the-art image quality in the DFCT at the time of this thesis. The attenuation **A** qualitatively has the largest contrast-to-noise ratio, compared to diffusion **F** and refractive index **K**. The POM and water cylinders exhibit the strongest attenuation signal. While the neoprene contrast is very low, the attenuation of the three wool cylinders increases with the dampness and corresponding relative water content. Overall, each material is well separated from the noise floor. In contrast, the diffusion coefficient **F** shows a much higher noise floor compared to the sample signal. Similar to the air-scan image, no distinct artifacts are visible. The neoprene, dry wool, and chocolate chips exhibit the strongest scattering signal, due to their porous structure and resulting high number of uni-directional air-material interfaces. In this comparison, we ignore the effect of beam-hardening-induced visibility reduction [22]. It results in POM and water having a non-zero diffusion coefficient, despite the materials being homogeneous on the length scales relevant for dark-field signal. While important for medical imaging, this has no consequence for the discussion of artifacts in the context of the presented methods in this chapter. The complementary contrast between attenuation **A** and diffusion **F** is still apparent. The refractive index **K** qualitatively looks very similar to the attenuation **A**, albeit with a much higher noise floor and decreased contrast-to-noise ratio (CNR).
The second-from-left column **B**,**G**,**L** finally depicts the relevance of the new techniques presented in this chapter by relating the image errors with the sample contrast. The artifacts themselves are unsurprisingly very similar to Fig. 3.19. However, in attenuation **B** they are only visible if a narrow window is chosen, completely saturated from sample signal. The sample's diffusion signal **G** is approximately of the same magnitude as the image artifacts. The tubes and different materials in them are still discernible, but significantly modulated by the concentric features corresponding to the back-projected fringe pattern. The refractive index **L** is affected strongest by the artifacts. Their contrast is stronger than the sample contrast, resulting in barely visible tubes and their contents.
The middle, second-from-right, and right columns serve the same purpose as in Fig. 3.19, isolating the impact of a single novel method, respectively. Again, the attenuation is affected similarly by the vibration model **C** and the higher-order correction **D**. In diffusion, the higher-order visibilities **I** are most detrimental. But also an naive vibration model **H** lead to significant image errors. In contrast, the refractive index is completely unusable with incorrect vibrations **M**. Uncorrected higher-order terms **N** are less severe, but still comparable to the total error in diffusion and significant on their own. The macroscopic flat-field movement in the right-most column **E**,**J**,**O** is only discernible

in refractive index **O**. In these examples, it has the lowest overall impact of the discussed effects.

3.5 Discussion

The main finding in this chapter is that by using the proposed PCA model, we can significantly reduce the number of parameters to model the vibration state of each projection compared to previous methods. The approach is quite universal and may be applied to other setups, especially if they use a tiled G_2 grating. The number of parameters $A^\star$, $B^\star$, and $C^\star$ per projection might be adjusted according to the scree plots. If the setup is reasonably stable, they have to be determined only once—in contrast to the model components $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ themselves, which should be updated somewhat regularly.

We further propose estimating the vibration coefficients in a sample-scan on a per-projection basis from sample-free detector area. This simple method is used in order to validate the model and to investigate the number of PCA components needed for a reconstruction without vibration artifacts. If one is only interested in flat-field terms $\boldsymbol{I}$, $\boldsymbol{V}$, and ϕ, a simpler polynomial model (or even only global flux and phase offset) might be sufficient. But for an adequate reconstruction of an unknown sample, a more accurate vibration model is required, as shown by the comparisons of reconstructed images.

The proposed correction works independently of the attenuation or phase shift of the sample. We determine the inherent interferometer vibrations, which are independent of the sample properties in the model. The tiles of G_2 have to exhibit some joint vibration such that the number of model parameters can be reduced by PCA. Theoretically, the tiles could be completely independent. However, as the fluctuation is induced by an external vibration and the tiles are mounted on a common holder, similar vibrations can be expected in practice. Lacking a comparable system to the DFCT, we assume that our vibration model is universally applicable. The air-scan processing in section 3.2.3, including tile-wise vibrations and the PCA method, is robust to noise and shortened rotation times: we can choose a large tube current and if the statistics are still limited by short rotation time, we can use several gantry rotations to increase the number of data points. The sample-scan (potentially with a patient) is governed by dose limitations, which is a problem orthogonal to the vibration model.

The exact order of the optimization steps in section 3.2.3 is empirically found and may be different for another setup. We show the detailed algorithm due to the novelty of the experimental platform and the lack of comparison in the literature. It is generally advisable to begin with the terms which influence the cost function the most, namely global intensity $\boldsymbol{\alpha}_{00}$ and phase $\boldsymbol{\gamma}_{00}$. Then we introduce global spatial fluctuations, and then tile-wise spatial fluctuations. We do expect a very similar pipeline to be successful in an alternative implementation of a DFCT.

Reducing the spatial fluctuations in all channels—for which we formulate the model and PCA method—while keeping the global phase variation for sampling may not be feasible, as the grating movement would have to be strictly transversal. Removing one grating may reduce system vibrations and their impact on image quality. Recent work uses structured anodes [60, 61] and could make the G_0 grating obsolete. However, the fact that the vibration is not the same on all G_2 tiles clearly shows that at least some vibrations cannot be avoided by omitting G_0. Furthermore, the structured anodes can only be used for very small fan-angles. Reducing the number of G_2 tiles by implementing larger tiles would naturally reduce the complexity and should be a motivation for further advances in the grating fabrication field.

We motivate, determine, and correct for higher-order terms of the Fourier model of the stepping curve. Due to the phase-shifting G_1 with its triangular profile and the G_0 and G_2 gratings with $\approx 60\,\%$ duty cycle the observed stepping curve has stronger terms $n > 1$ than a conventional setup. Furthermore, these terms cause image artifacts in our pipeline because the phase stepping is not guaranteed to be evenly sampled over 2π in every angular window. We determine them with the proposed extended SPR algorithm, assuming an interferometer vibration state found with the first-order model. The higher-order model is only used for SPR but not in an "online" optimization of the vibrations. Although not completely rigorous this procedure has proven accurate enough in practice. However, a closer examination of the interplay of vibrations and higher Fourier orders is necessary to fully understand the imaging system.

The *macroscopic* movement of the flat-fields is potentially in part modeled by the vibration terms $\boldsymbol{I}^{\text{vib}}$, $\boldsymbol{V}^{\text{vib}}$, and ϕ^{vib}—which originally are motivated by spatio-temporal fluctuations caused by *microscopic* movement of the gratings. In visibility and phase, we model these fluctuations with 2D polynomials. In intensity, we compute the vibration term $\boldsymbol{I}^{\text{vib}}$ directly from the relative residuum of model and measurement. This residual probably is a combination of vibration- and movement effects.

All in all, the corrections employed on the DFCT data are somewhat empirical, even though motivated by physical effects. We have no independent guarantee that the terms we determine in air-scan processing describe the underlying processes accurately and are

phenomenologically separate. We do however determine a model which describes the data with a small number of parameters, enabling tomographic reconstruction of an unknown sample.

In summary, we propose a processing scheme to identify and correct for vibrations of a Talbot-Lau interferometer mounted inside a rotating clinical CT gantry. All channels show spatial fluctuations changing between exposures which have a specific shape over the detector. Polynomials in visibility and phase over the whole detector and per G_2 tile are used in the data model and an implicit optimization scheme with linearized phase retrieval as an internal layer is employed to determine their individual magnitude per exposure. The resulting tile-wise vibrations are coupled by applying PCA and keeping only the first few dominant components for processing a sample-scan. In the intensity channel, dominant fluctuation components are identified by PCA on the normalized residuum. The vibration model for the intricate fluctuations allows an artifact-free reconstruction in the presence of a sample with only a small number of parameters. A comparison with a vibration model using global polynomials shows that the latter is not sufficient to capture the system's dynamics and leads to artifacts in the reconstruction of a sample (see Fig. 3.19 and 3.20). The proposed algorithm can be used to identify setup-specific vibrations in a clinically relevant platform where mechanical instabilities can not be avoided.

Chapter 4

Sample-scan processing

This chapter discusses challenges and techniques concerning determining and reconstructing the three contrast modalities transmission, dark-field, and differential phase of a measured sample. The main challenge lies in determining the interferometer vibrations during the time of the sample scan. It is not trivial as neither vibration coefficients nor the sample projections are known initially, although we generally keep the vibration models from PCA as discussed in chapter 3. We have to consider a large sample covering the whole detector and the temporal stability of the setup.

First, we define the model with a sample in the beam path. Then we motivate processing pipelines and optimization schemes for small objects where there is sample-free detector area, and for large objects with no sample-free area. To obtain satisfactory image quality of the tomographic reconstructions with sliding-window SPR we present and employ several novel post-processing techniques.

4.1 Model notation

The coefficients $\boldsymbol{\alpha}^{\mathsf{s}}$, $\boldsymbol{\beta}^{\mathsf{s}}$, and $\boldsymbol{\gamma}^{\mathsf{s}}$ describe the state of the interferometer vibrations at the time of the sample scan. For ease of notation we add the DC component $\boldsymbol{\mathcal{P}}_{00}$ at index $k = 0$ to the PCA models from chapter 3,

$$\mathcal{A}_{0,p} := \mathcal{B}_{0,p} := \mathcal{C}_{0,p} := \mathbf{1}\,, \tag{4.1}$$

and use coefficients $\alpha^{\mathsf{s}}_{0,t}$, $\beta^{\mathsf{s}}_{0,t}$, and $\gamma^{\mathsf{s}}_{0,t}$ instead of $\alpha^{\mathsf{s}}_{00t}$, β^{s}_{00t}, and $\gamma^{\mathsf{s}}_{00t}$ accordingly. We also drop the $\cdot^{\star}$ identifier for the sample-scan coefficients because we do not need to differentiate between polynomials and PCA model in this chapter.

We use helper variables I^{tot}_{pt}, V^{tot}_{pt}, and ϕ^{tot}_{pt} to describe the complete interferometer state at the time of the sample scan combining flat-fields and vibrations,

$$I^{\mathsf{tot}}_{pt} = I_p \left(1 + \sum_{k=0}^{A^\star} \mathcal{A}_{kp}\alpha^{\mathsf{s}}_{kt}\right)\,, \tag{4.2}$$

$$V^{\mathsf{tot}}_{pt} = V_p \left(1 + \sum_{k=0}^{B^\star} \mathcal{B}_{kp}\beta^{\mathsf{s}}_{kt}\right)\,, \tag{4.3}$$

$$\phi^{\mathsf{tot}}_{pt} = \phi_p + \sum_{k=0}^{C^\star} \mathcal{C}_{kp}\gamma^{\mathsf{s}}_{kt}\,, \tag{4.4}$$

and adapt the model (3.2) with vibrations but without other perturbations accordingly,

$$y^{\mathsf{s}}_{pt} = T_{pt} I^{\mathsf{tot}}_{pt} \left[1 + D_{pt} V^{\mathsf{tot}}_{pt} \cos\left(\Phi_{pt} + \phi^{\mathsf{tot}}_{pt}\right)\right]\,, \tag{4.5}$$

with the sample's projected transmission $\boldsymbol{T}$, dark-field $\boldsymbol{D}$, and differential phase $\boldsymbol{\Phi}$ under angle t in pixel p. While the PCA models $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ are inherited unchanged from the air-scan the coefficients $\boldsymbol{\alpha}^{\mathsf{s}}$, $\boldsymbol{\beta}^{\mathsf{s}}$, and $\boldsymbol{\gamma}^{\mathsf{s}}$ are in general different from the air-scan and have to be re-determined.

It should be emphasized that even though the vibration state in a sample-scan is initially unknown it is assumed to be independent of the sample. Put differently, a sample-scan has the same characteristics in terms of the interferometer as an air-scan, just that a sample is also in the beam and modulates the respective parts of the model. Only because this modulation is initially unknown the process of determining $\boldsymbol{\alpha}^{\mathsf{s}}$, $\boldsymbol{\beta}^{\mathsf{s}}$, and $\boldsymbol{\gamma}^{\mathsf{s}}$ is different and arguably more involved. The interferometer itself is not changed by the presence of the sample.

We formulate the likelihood function C^{s} using (4.5),

$$C^{\mathsf{s}} = \sum_{t=1}^{T}\sum_{p=1}^{P} w_{pt}\left(y^{\mathsf{s}}_{pt} - \hat{y}^{\mathsf{s}}_{pt}\right)^2\,, \tag{4.6}$$

with $\hat{y}^{\mathsf{s}}_{pt}$ the measured intensities and w_{pt} the statistical weights. Due to the measurement being a Poisson process, the weights can be approximated by $w_{pt} \approx 1/\hat{y}^{\mathsf{s}}_{pt}$. The optimization problem is formulated

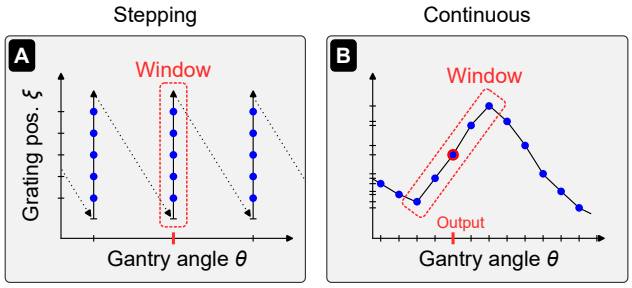

Figure 4.1: Acquisition schemes for **A** stepping and **B** continuous grating movement. **A** Stepping requires separate gantry and grating movement. The input window for signal retrieval consists of one gantry angle θ over which the stepping position ξ is changed. **B** In continuous acquisition both θ and ξ change simultaneously. The signal retrieval works on a sliding window of exposures.

as a minimization of C^{s},

$$\left(\hat{\boldsymbol{T}}, \hat{\boldsymbol{D}}, \hat{\boldsymbol{\Phi}}, \hat{\boldsymbol{\alpha}}^{\mathrm{s}}, \hat{\boldsymbol{\beta}}^{\mathrm{s}}, \hat{\boldsymbol{\gamma}}^{\mathrm{s}}\right) = \underset{\boldsymbol{T},\boldsymbol{D},\boldsymbol{\Phi},\boldsymbol{\alpha}^{\mathrm{s}},\boldsymbol{\beta}^{\mathrm{s}},\boldsymbol{\gamma}^{\mathrm{s}}}{\arg\min} C^{\mathrm{s}} \,. \tag{4.7}$$

In practice, it is not feasible to fit all unknowns of (4.7) immediately. The likelihood function is highly non-convex and it can not be expected that the determined minimum is physically sensible (i.e. it is not the global minimum). We can incorporate prior knowledge in the optimization process by including only some parameters in the beginning with sensible initial values and introducing more parameters step by step. The specific methods of optimizing the model are explained in the following sections.

4.2 Alternating optimization in projection domain

For an unstable, non-convex optimization problem it is a common technique to split the set of parameters into two or more groups. Each parameter set is optimized with respect to the objective function, while the other sets are kept constant. After convergence of one set, this procedure is repeated with the sets exchanged. In general this technique has slower overall convergence than joint optimization, but is more robust against running into local minima [37]. For the optimization in the context of (4.7) we split the parameters into two groups: projections $(\boldsymbol{T}, \boldsymbol{D}, \boldsymbol{\Phi})$ and vibration coefficients $(\boldsymbol{\alpha}^{\mathrm{s}}, \boldsymbol{\beta}^{\mathrm{s}}, \boldsymbol{\gamma}^{\mathrm{s}})$. The parameters in each set are optimized jointly, respectively.

4.2.1 Projection optimization

For the projections, the optimization is performed via sliding-window SPR [62]. It applies SPR on a small angular window of consecutive exposure to determine the sample projections of the exposure central to that window. The current estimate for the vibration state $(\hat{\boldsymbol{\alpha}}^{\mathrm{s}}, \hat{\boldsymbol{\beta}}^{\mathrm{s}}, \hat{\boldsymbol{\gamma}}^{\mathrm{s}})$ is assumed constant and the model (4.5) is linearized with respect to the sample projections $\boldsymbol{T}$, $\boldsymbol{D}$, and $\boldsymbol{\Phi}$. The system matrix $\boldsymbol{A}^{\mathrm{s}}_{pt} \in \mathbb{R}^{\mathcal{M}\times 3}$ is defined in accordance to section 3.2.1,

$$\left[\boldsymbol{A}^{\mathrm{s}}_{pt}\right]_{\tau,1} = I^{\mathrm{tot}}_{p\tau} \,, \tag{4.8}$$

$$\left[\boldsymbol{A}^{\mathrm{s}}_{pt}\right]_{\tau,2} = I^{\mathrm{tot}}_{p\tau} V^{\mathrm{tot}}_{p\tau} \cos\left(\phi^{\mathrm{tot}}_{p\tau}\right) \,, \tag{4.9}$$

$$\left[\boldsymbol{A}^{\mathrm{s}}_{pt}\right]_{\tau,3} = I^{\mathrm{tot}}_{p\tau} V^{\mathrm{tot}}_{p\tau} \sin\left(\phi^{\mathrm{tot}}_{p\tau}\right) \,, \tag{4.10}$$

with

$$\tau \in \mathrm{window}(t) \,, \tag{4.11}$$

$$\mathrm{window}(t) = \left\{t - \lfloor \mathcal{M}/2 \rfloor, \ldots, t + \lfloor \mathcal{M}/2 \rfloor\right\} \,, \tag{4.12}$$

and $\mathcal{M}$ the user-defined number of consecutive exposures used for a window. It is usually an odd number so that the index t really is in the center of the window. The indexing is wrapped around the scan's end in the case of an axial scan. We formulate the weighted least-squares problem per output pixel p and output angle t,

$$\left(\left(\boldsymbol{A}^{\mathrm{s}}_{pt}\right)^{\top} \boldsymbol{W}_{pt} \boldsymbol{A}^{\mathrm{s}}_{pt}\right) \boldsymbol{x}^{\mathrm{s}}_{pt} = \left(\boldsymbol{A}^{\mathrm{s}}_{pt}\right)^{\top} \boldsymbol{W}_{pt} \hat{\boldsymbol{y}}^{\mathrm{s}}_{pt} \,. \tag{4.13}$$

We solve (4.13) for $\hat{\boldsymbol{x}}^{\mathrm{s}}_{pt}$ with LU decomposition in JAX and compute the sample projections,

$$\hat{T}_{pt} = \left[\hat{\boldsymbol{x}}^{\mathrm{s}}_{pt}\right]_1 \,, \tag{4.14}$$

$$\hat{D}_{pt} = {}^{1}\!/_{\hat{T}_{pt}} \sqrt{\left[\hat{\boldsymbol{x}}^{\mathrm{s}}_{pt}\right]_2^2 + \left[\hat{\boldsymbol{x}}^{\mathrm{s}}_{pt}\right]_3^2} \,, \tag{4.15}$$

$$\hat{\Phi}_{pt} = \arctan2\left(-\left[\hat{\boldsymbol{x}}^{\mathrm{s}}_{pt}\right]_3, \left[\hat{\boldsymbol{x}}^{\mathrm{s}}_{pt}\right]_2\right) \,. \tag{4.16}$$

Although this mathematical notation is cumbersome, the actual computation is quite simple: to obtain the sample's projections under angle t, we use SPR on the adjacent projections around t with flat-fields and vibrations as input. In comparison to chapter 3, the system matrix $\boldsymbol{A}^{\mathrm{s}}_{pt}$ depends on output pixel p and output angle t. The larger number of indices ($\boldsymbol{A}^{\mathrm{s}}_{pt}$ depends on p and t, and is then indexed with τ) is explained by every output projection t having its own input window of exposures $\tau \in \mathrm{window}(t)$.

Sliding-window SPR involves the implicit assumption that the sample does not rotate and its projections do not change over the input angle τ, illustrated in Fig. 4.1. This is obviously not the case in a continuously rotating system. The discrepancy results in "rotation" artifacts which are especially strong in regions where the sample's projections change quickly. This is the case at sample edges and in the outer regions of the FOV with large projected movement. The

resulting artifacts are in general small enough to not negatively affect the overall optimization. We discuss a correction method in section 4.3.2 and image quality in section 4.4.

4.2.2 Vibration optimization

Solving (4.7) for the vibration coefficients $(\boldsymbol{\alpha}^{\mathrm{s}}, \boldsymbol{\beta}^{\mathrm{s}}, \boldsymbol{\gamma}^{\mathrm{s}})$ is performed with the L-BFGS algorithm [51]. The coefficient vectors are concatenated into one combined parameter vector while the current estimate for the sample projections $(\hat{\boldsymbol{T}}^{\mathrm{s}}, \hat{\boldsymbol{D}}^{\mathrm{s}}, \hat{\boldsymbol{\Phi}}^{\mathrm{s}})$ is assumed constant.

We implement the likelihood C^{s} and its gradient wrt. the vibrations in `JAX`. The results are transferred from the GPU to host memory and passed to the L-BFGS implementation in `scipy`. Unlike the nested optimization in section 3.2.2 which computes the non-trivial vibration gradient through SPR, here we fix the projections $(\hat{\boldsymbol{T}}^{\mathrm{s}}, \hat{\boldsymbol{D}}^{\mathrm{s}}, \hat{\boldsymbol{\Phi}}^{\mathrm{s}})$ and the cost function gradient is rather straightforward. Nevertheless it is convenient in practice to use the automatic differentiation (autodiff) features of `JAX` instead of implementing it by hand.

As mentioned before the temporal progression of $(\boldsymbol{\alpha}^{\mathrm{s}}, \boldsymbol{\beta}^{\mathrm{s}}, \boldsymbol{\gamma}^{\mathrm{s}})$ is similar to the air-scan coefficients $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$. Primarily the influence of the gantry rotation is reproducible from scan to scan because the angular position is encoded with each measured exposure. In contrast, the exact state of the various high-frequency oscillations influencing the gratings is not known a priori. A good initialization for the sample-scan vibrations $(\boldsymbol{\alpha}^{\mathrm{s}}, \boldsymbol{\beta}^{\mathrm{s}}, \boldsymbol{\gamma}^{\mathrm{s}})$ is therefore a low-pass-filtered version of the air-scan parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$.

Small samples For small samples which do not cover the whole active interferometer we can use the sample-free area to determine $(\hat{\boldsymbol{\alpha}}^{\mathrm{s}}, \hat{\boldsymbol{\beta}}^{\mathrm{s}}, \hat{\boldsymbol{\gamma}}^{\mathrm{s}})$ before any signal retrieval of the sample projections. The pipeline is illustrated in Fig. 4.2. In sample-free area, $T = D = 1$ and $\Phi = 0$. The model (4.5) simplifies to

$$\widetilde{y}^{\mathrm{s}}_{pt} = I^{\mathrm{tot}}_{pt} \left(1 + V^{\mathrm{tot}}_{pt} \cos \phi^{\mathrm{tot}}_{pt}\right) . \tag{4.17}$$

Eq. (4.17) looks similar to the canonical model (3.1) and we could linearize it accordingly. However, performing sliding-window SPR on $\widetilde{\boldsymbol{y}}^{\mathrm{s}}$ to obtain $\boldsymbol{I}^{\mathrm{tot}}$, $\boldsymbol{V}^{\mathrm{tot}}$, and $\boldsymbol{\phi}^{\mathrm{tot}}$ is not feasible: the vibration terms change much more quickly than the rotating sample's projections $\boldsymbol{T}$, $\boldsymbol{D}$, and $\boldsymbol{\Phi}$ and the assumption of window-wise constancy is strongly violated. Instead we find $(\hat{\boldsymbol{\alpha}}^{\mathrm{s}}, \hat{\boldsymbol{\beta}}^{\mathrm{s}}, \hat{\boldsymbol{\gamma}}^{\mathrm{s}})$ with the L-BFGS solver as explained earlier. The advantage of the vibration model determined with PCA as discussed in chapter 3 is the low number of coefficients per exposure. The results in section 3.4.4 were obtained with this small-sample pipeline. In this chapter, we use it to analyze the impact of sample-free area size on reconstruction quality and the temporal stability of the PCA model in section 4.4.2.

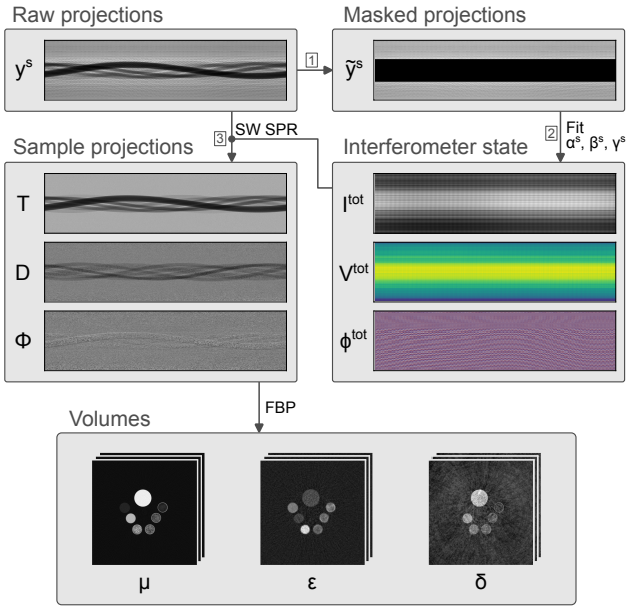

Figure 4.2: Processing of a small sample. We use the sample-free area to fit the PCA vibration model coefficients. Afterward the sample projections are determined with sliding-window SPR. They are reconstructed with filtered back-projection (FBP).

Large samples We investigate the influence of the relative size of the sample-free detector area on the quality of the tomographic reconstruction in the "small sample pipeline" in section 4.4.2 and Fig. 4.8. Without jumping ahead, it should not be surprising that the quality of the vibration fit decreases for small (or no) sample-free detector area. This scenario is relevant for the DFCT because we aim at measuring human subjects which may fill the whole FOV.

When there is insufficient sample-free interferometer area we employ an alternating optimization of C^{s} in (4.7) wrt. projection and vibration quantities. In general we do not optimize all quantities immediately but introduce them in succession.
All vibrations coefficients $(\boldsymbol{\alpha}^{\mathrm{s}}, \boldsymbol{\beta}^{\mathrm{s}}, \boldsymbol{\gamma}^{\mathrm{s}})$ are initialized with low-pass-filtered versions of the air-scan values $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$.
Concerning initial values of sample projections, the transmission $\boldsymbol{T}$ is most detrimental in the cost function but also can be initialized with a good approximation: we perform tomographic reconstruction with filtered back-projection (FBP) on the unprocessed projections $\hat{\boldsymbol{y}}^{\mathrm{s}}$ normalized by the reference intensity $\hat{\boldsymbol{I}}$,

$$\hat{\boldsymbol{\mu}} = \mathcal{R}^{-1} \left[-\ln \left(\frac{\hat{\boldsymbol{y}}^{\mathrm{s}}}{\hat{\boldsymbol{I}}} \right) \right] , \tag{4.18}$$

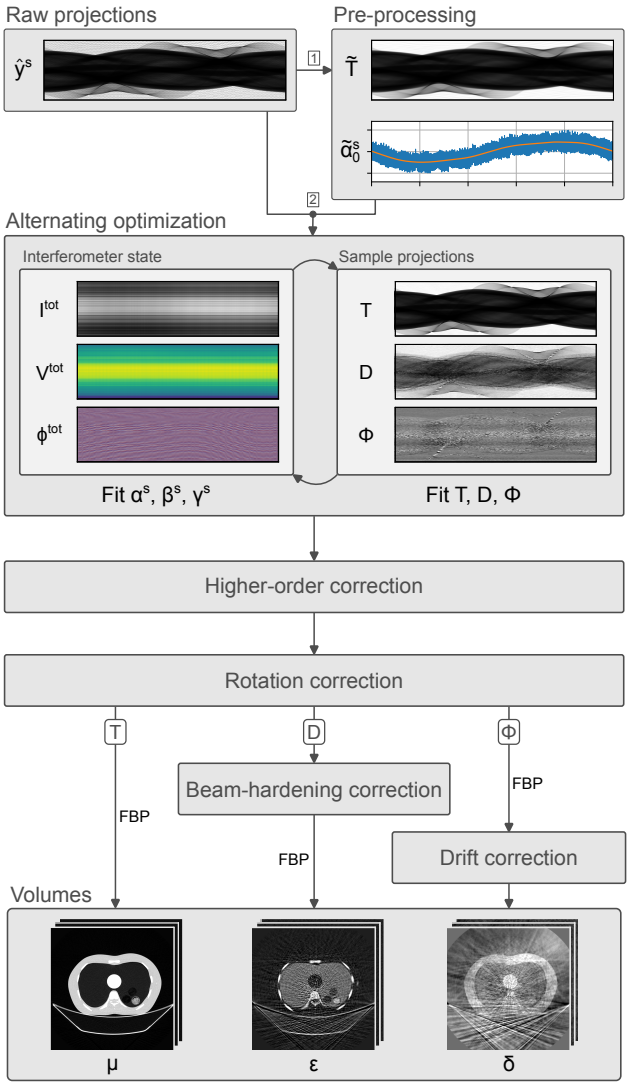

Figure 4.3: Flowchart of the processing pipeline for the scan of a large sample. We compute pre-processing estimates $\widetilde{\boldsymbol{T}}$ and $\widetilde{\boldsymbol{\alpha}}_0^{\mathrm{s}}$ from the raw projection data $\hat{\boldsymbol{y}}^{\mathrm{s}}$. We use them to initialize the alternating optimization of vibration coefficients $(\hat{\boldsymbol{\alpha}}^{\mathrm{s}}, \hat{\boldsymbol{\beta}}^{\mathrm{s}}, \hat{\boldsymbol{\gamma}}^{\mathrm{s}})$ and sample projections $(\hat{\boldsymbol{T}}, \hat{\boldsymbol{D}}, \hat{\boldsymbol{\Phi}})$. After convergence we employ corrections of higher-order Fourier terms, sample rotation, beam-hardening-induced visibility reduction (only on dark-field $\hat{\boldsymbol{D}}$), and slow phase drift (only on diffusion coefficient $\hat{\boldsymbol{\delta}}$). The tomographic reconstruction into attenuation coefficient $\hat{\boldsymbol{\mu}}$, diffusion coefficient $\hat{\boldsymbol{\varepsilon}}$, and refractive index decrement $\hat{\boldsymbol{\delta}}$ are performed with FBP, respectively.

to obtain an initial guess of the attenuation coefficient $\hat{\boldsymbol{\mu}}$. $\mathcal{R}^{-1}$ is the inverse Radon transform and shorthand for the FBP. $\hat{\boldsymbol{\mu}}$ is modulated by the fringe pattern but has comparatively good quality. We then perform tomographic forward-projection,

$$\widetilde{\boldsymbol{T}} = \exp\left(-\mathcal{R}\left[\hat{\boldsymbol{\mu}}\right]\right) , \tag{4.19}$$

and obtain an approximation of the sample transmission $\widetilde{\boldsymbol{T}}$ before doing any phase-retrieval or fitting of the vibration coefficients. $\mathcal{R}$ represents the Radon transform and is shorthand for tomographic foward-projection. Due to the double enforcement of tomographic consistency through (4.18) and (4.19) there are hardly any fringe artifacts in $\widetilde{\boldsymbol{T}}$.

The global flux variation $\boldsymbol{\alpha}_0^{\mathrm{s}}$ is also estimated before phase-retrieval. We use a high-pass filter on the exposure-wise mean of the normalized measured data $\hat{\boldsymbol{y}}^{\mathrm{s}}$ and combine it with the low-pass-filtered air-scan values $\hat{\boldsymbol{\alpha}}_{00}$,

$$\widetilde{\boldsymbol{\alpha}}_0^{\mathrm{s}} = \mathrm{HF}_t\left[\left\langle \frac{\hat{\boldsymbol{y}}^{\mathrm{s}}}{\hat{\boldsymbol{I}}\,\widetilde{\boldsymbol{T}}} \right\rangle_p\right] + \mathrm{LF}_t\left[\hat{\boldsymbol{\alpha}}_{00}\right] \tag{4.20}$$

with HF_t and LF_t the temporal high- and low-pass filter, respectively. We implement LF_t with a Gaussian filter and define $\mathrm{HF}_t\left[\boldsymbol{x}\right] = \boldsymbol{x} - \mathrm{LF}_t\left[x\right]$.

After those initializations we begin with the main optimization loop consisting of repeated minimization of C^{s} in (4.7). All parameters are held fixed at their initial values until explicitly listed here, i.e. added as an actively optimized parameter. The individual stages of the alternating optimization are listed in table 4.1 by the parameters they optimize.

Step	Vibrations	Projections
1	γ_0^{s}	$\boldsymbol{T}$
2	γ^{s}	$\boldsymbol{T}$
3	$\boldsymbol{\alpha}^{\mathrm{s}}$,$\gamma^{\mathrm{s}}$	$\boldsymbol{T}$, $\boldsymbol{D}$
4	$\boldsymbol{\alpha}^{\mathrm{s}}$, $\boldsymbol{\beta}^{\mathrm{s}}$, γ^{s}	$\boldsymbol{T}$, $\boldsymbol{D}$, $\boldsymbol{\Phi}$

Table 4.1: Optimization steps in the processing pipeline for a large sample. For each step we first minimize C^{s} wrt. the vibrations and then wrt. the projections. The results are used as initializations in the next step, respectively. Until a variable is listed, it is held at its initialization and not included in the cost function gradient.

4.3 Image enhancement methods

We found several methods to improve the image quality specifically in processing the scan of a large sample. We explain each of them in the following. The impact on experimental results is shown in section 4.4.

4.3.1 Tomographic consistency

Enforcing tomographic consistency in this projection-focused optimization pipeline is useful not only to compute an initial transmission estimate $\widetilde{\boldsymbol{T}}$ as mentioned before, but also to reduce artifacts in intermediate sample projection estimates of $\boldsymbol{T}$, $\boldsymbol{D}$, and $\boldsymbol{\Phi}$ listed in table 4.1. We formulate the general method,

$$\boldsymbol{T}^{\mathrm{FP}} = \exp\left(-\mathcal{R}\left[\mathcal{R}^{-1}\left[-\ln\left(\boldsymbol{T}\right)\right]\right]\right) , \tag{4.21}$$

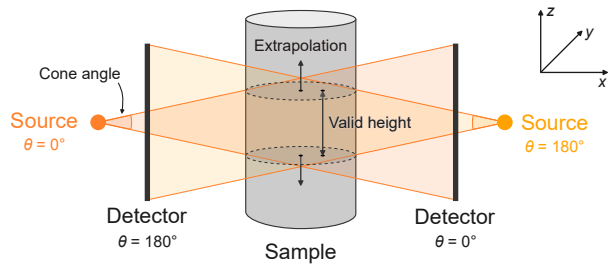

Figure 4.4: Enforcing tomographic consistency by forward-projecting the FBP. Due to cone-beam geometry, the top and bottom of the sample are not estimated correctly. We extrapolate the top- and bottom-most axial slice before forward-projection.

with $\boldsymbol{T}^{\text{FP}}$ the projections with enforced tomographic consistency. It works analogously for dark-field $\boldsymbol{D}$ and differential phase $\boldsymbol{\Phi}$, reducing noise and fringe artifacts from incorrect intermediate vibration estimates.

This method can be seen as a middle-ground between a strict "projections-then-FBP" pipeline, and a full-fledged intensity-based statistical iterative reconstruction (IBSIR) going directly from raw sample projections to all sample volumes. The former requires phase-stepping for each projection angle which the DFCT does not allow. The latter correctly incorporates the continuous rotation into the model, but (unnecessarily) introduces a sophisticated statistical model for the projection data. If we deem the fit of the vibration coefficients to be the main challenge in our processing pipeline (in contrast to low photon counts), then the projection-based processing combined with intermediate tomographic consistency enforcement is a sensible compromise.
In cone-beam geometry the FBP can not reconstruct the full sample FOV along the rotation axis leading to under-estimation at the top and bottom of the forward-projected volumes. We mitigate this error by extrapolating the intermediate reconstructed volume vertically along the rotation axis before forward-projection, illustrated in Fig. 4.4. This approximation is obviously only valid for samples that are sufficiently invariant along the rotation axis.

4.3.2 Rotation artifacts

As mentioned before the sliding-window SPR assumes that the sample projections $(\boldsymbol{T}, \boldsymbol{D}, \boldsymbol{\Phi})$ do not change in the angular window of exposures. This assumption is violated in the DFCT because the gantry is continuously rotating and each exposure corresponds to a different projection angle. The resulting projection artifacts are different depending on the signal modality. In general their severity depends on the projected lateral movement of the sample which in turn is larger for voxels far away from the rotation axis.

The transmission $\hat{\boldsymbol{T}}$ is simply blurred out over the angular window and the reconstructed attenuation coefficient $\hat{\boldsymbol{\mu}}$ has lower spatial resolution with increased distance from the isocenter. The dark-field $\hat{\boldsymbol{D}}$ and differential phase $\hat{\boldsymbol{\Phi}}$ are strongly correlated and both show artifacts resembling the reference phase pattern $\hat{\phi}$. They are discussed in [59] which calls them "movement-induced crosstalk-artifacts". The reconstructed diffusion coefficient $\hat{\varepsilon}$ and refractive index $\hat{\delta}$ are modulated with streaks between high-absorbing features outside the central FOV.

We have found that in the DFCT the non-constant transmission $\boldsymbol{T}$ in an angular window is the main source of rotation artifacts. Luckily it is also least affected by movement and easiest to approximate both in terms of noise as well as sampling ϕ of the stepping curve. We therefore aim to find a good estimate $\hat{\boldsymbol{T}}$ and perform signal retrieval again on $\hat{\boldsymbol{y}}^{\text{s}}/\hat{\boldsymbol{T}}$ without the transmission influence to obtain improved estimates $\hat{\boldsymbol{D}}$ and $\hat{\boldsymbol{\Phi}}$. [59] discusses multiple ways to achieve this and reduce rotation artifacts in the context of large samples in the DFCT. Here we present one method because it is used for the experimental results in section 4.4.

It relies on the reduction of projection artifacts after reconstructing the projections with FBP and then projecting them again with tomographic foward-projection (FP) as outlined in section 4.3.1. The outline of the algorithm is:

1. Compute $\hat{\boldsymbol{T}}$ with sliding-window SPR on $\hat{\boldsymbol{y}}^{\text{s}}$.
2. Correction loop:
   (a) Compute $\hat{\boldsymbol{\mu}} = \mathcal{R}^{-1}\left[-\ln \hat{\boldsymbol{T}}\right]$.
   (b) Compute $\hat{\boldsymbol{T}}^{\text{FP}} = \exp\left(-\mathcal{R}\left[\hat{\boldsymbol{\mu}}\right]\right)$.
   (c) Compute $\hat{\boldsymbol{T}}^{\star}$, $\hat{\boldsymbol{D}}$, and $\hat{\boldsymbol{\Phi}}$ with sliding-window SPR on $\hat{\boldsymbol{y}}^{\text{s}}/\hat{\boldsymbol{T}}^{\text{FP}}$.
   (d) Update $\hat{\boldsymbol{T}} \leftarrow \hat{\boldsymbol{T}} \times \hat{\boldsymbol{T}}^{\star}$

First we obtain a transmission estimate $\hat{\boldsymbol{T}}$ with normal sliding-window SPR, preferably with a slightly reduced window size $\mathcal{M}^{-}$. We compute the attenuation coefficient with FBP on $\hat{\boldsymbol{T}}$ and forward-project it again to obtain an improved transmission estimate $\hat{\boldsymbol{T}}^{\text{FP}}$. A virtual transmission-corrected measurement $\hat{\boldsymbol{y}}^{\text{s}}/\hat{\boldsymbol{T}}^{\text{FP}}$ is computed on which we perform sliding-window SPR, preferably with a slightly larger window $\mathcal{M}^{+}$. The obtained estimates $\hat{\boldsymbol{D}}$ and $\hat{\boldsymbol{\Phi}}$ (and by extension their reconstructions $\hat{\varepsilon}$ and $\hat{\delta}$) already show less rotation artifacts. We sharpen the transmission estimate $\hat{\boldsymbol{T}}$ with the residual $\hat{\boldsymbol{T}}^{\star}$ originating by the slight blurring inherent to the FBP and FP. The steps can be repeated multiple times to further decrease rotation artifacts while the noise increases with each iteration.

4.3.3 Corrections for higher visibility orders

The significance of higher-order terms in the Fourier series description of the stepping curve has been discussed in section 3.3.1 already. In short, we determine the higher-order visibilities $\boldsymbol{V}^{(n)}$ and phases $\boldsymbol{\phi}^{(n)}$ in air-scan processing after iterative optimization of the interferometer vibrations $\boldsymbol{I}^{\text{vib}}$, $\boldsymbol{V}^{\text{vib}}$, and $\boldsymbol{\phi}^{\text{vib}}$. Equations (3.54) and (3.55) describe the interaction with interferometer vibrations terms. For the sample-scan reconstructions shown in section 3.4.4 we remove the terms $n \geq 2$ and perform sliding-window SPR. This involves the implicit assumption that the sample does not influence the higher-order terms which is accurate enough for small samples as seen in Fig. 3.20. In contrast, for large samples we have to include the sample influence to obtain accurate reconstructions.

Dark-field Section 2.2.5 motivates dark-field signal via small-angle scattering in the sample and a widening of an X-ray beam into a Gaussian intensity profile, reducing lateral coherence. As of (2.41) the relative visibility reduction D is computed by performing a convolution of the projected intensity distribution $\widetilde{A}(x)$ with a stepping curve $I(x)$ with period p_2 (commonly the G_2 period). We assume a period $p^{(n)}$ of the higher-order terms,

$$p^{(n)} = \frac{p_2}{n}\,, \tag{4.22}$$

and obtain the higher-order dark-field signal $D^{(n)}$ analogously to (2.41),

$$D^{(n)} = \exp\left(-\frac{2\pi^2 d^2}{(p_2/n)^2}\sigma^2\right) = D^{n^2}\,. \tag{4.23}$$

The relative reduction of the higher-order visibilities is apparently the first-order dark-field D to the n^2-th power. Qualitatively, because $D \leq 1$, the reduction is stronger with higher orders n.

Differential phase We introduce the phase shift $\Phi^{(n)}$ acting on the higher-order terms of the Fourier series. Intuitively, according to the higher-order stepping curve period in (4.22),

$$\Phi^{(n)} = n\Phi\,, \tag{4.24}$$

with the first-order differential phase Φ induced by the sample.

During sample-scan processing, we first determine the interferometer vibration state $(\hat{\boldsymbol{I}}^{\text{tot}}, \hat{\boldsymbol{V}}^{\text{tot}}, \hat{\boldsymbol{\phi}}^{\text{tot}})$ and sample projections $(\hat{\boldsymbol{T}}, \hat{\boldsymbol{D}}, \hat{\boldsymbol{\Phi}})$ using only the first-order model (4.5). Then we compute higher-order dark-field $\hat{\boldsymbol{D}}^{(n)}$ and differential phase $\hat{\boldsymbol{\Phi}}^{(n)}$ from the current first-order estimates $\hat{\boldsymbol{D}}$ and $\hat{\boldsymbol{\Phi}}$ to approximate the first-order projections $\hat{\boldsymbol{y}}^{\text{s}}$. We again perform sliding-window SPR on the first-order estimate and obtain sample projections without higher-order artifacts.

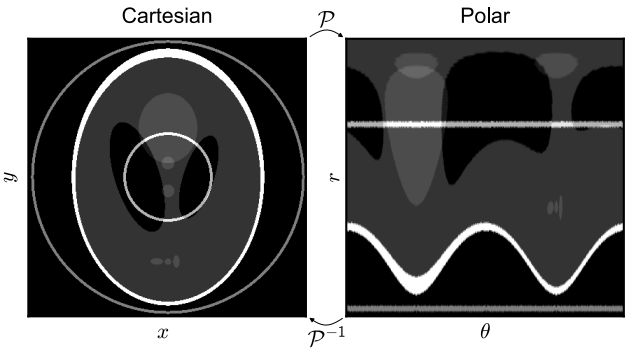

Figure 4.5: Illustration of the polar transform of an axial slice. Concentric features in cartesian coordinates are straight lines along θ in polar coordinates. We use the polar transform and its mean over θ to approximate the artifacts from slow phase drift in the tomographic reconstruction of the refractive index decrement δ.

4.3.4 Slow phase drift

The DFCT is based on a commercial CT system and is thermally comparatively stable due to on-board liquid cooling and being located in an air-conditioned room. As already mentioned the Talbot-Lau interferometer is sensitive to the microscopic movements of each grating and we do observe the remaining small thermal effects as a slow drift of the reference phase ϕ from slight thermal deformation of the gratings. "Drift" means slow movement or deformation of the gratings. The resulting change of ϕ is similar to the low-order 2D lateral polynomials in Fig. 3.2 which are also used for the vibration model in chapter 3. However, we found that the slow phase drift differs from the interferometer vibration and is different from the PCA vibration model $\hat{\boldsymbol{\mathcal{C}}}$. We show in section 4.4 that the result in the tomographic reconstruction of the refractive index decrement $\boldsymbol{\delta}$ are concentric bands. This is expected since the drift is much slower ($\approx 10^3$ s) than the typical duration of an axial scan ($\approx 10^0$ s) and therefore constant over the scan.

We use the similarity of the attenuation $\hat{\boldsymbol{\mu}}$ with the refractive index $\hat{\boldsymbol{\delta}}$ image and the concentric nature of the drift to approximate and remove it from $\hat{\boldsymbol{\delta}}$. We approximate the refractive index decrement $\hat{\boldsymbol{\delta}}$ with the scaled attenuation coefficient $\hat{\boldsymbol{\mu}}$,

$$\hat{\boldsymbol{\delta}} \approx \lambda\hat{\boldsymbol{\mu}}\,, \tag{4.25}$$

with scalar λ an empirically determined scaling parameter [63, 64]. Equation (4.25) is an approximation for a single material in a given interferometer with a fixed relationship between attenuation coefficient μ and refractive index decrement δ, combined

with the setup-specific sensitivity S_α from (2.30). It has proven sufficiently accurate for this correction, but using piecewise-linear relationships for multiple materials would also be possible [64].
The image error $\hat{\boldsymbol{\xi}}$ from the slow phase drift is estimated by their difference,

$$\hat{\boldsymbol{\xi}} = \hat{\boldsymbol{\delta}} - \lambda\hat{\boldsymbol{\mu}}\,. \tag{4.26}$$

We transform $\hat{\boldsymbol{\xi}}$ into cylindrical coordinates and subtract its mean over polar angle θ from the transformed refractive index $\hat{\boldsymbol{\delta}}$,

$$\hat{\boldsymbol{\delta}}^{\text{corr}} = \mathcal{P}^{-1}\left[\mathcal{P}\left[\hat{\boldsymbol{\delta}}\right] - \left\langle \mathcal{P}\left[\hat{\boldsymbol{\xi}}\right]\right\rangle_\theta\right], \tag{4.27}$$

resulting in the corrected refractive index image $\hat{\boldsymbol{\delta}}^{\text{corr}}$. $\mathcal{P}$ and $\mathcal{P}^{-1}$ represent the transform into polar coordinates and its inverse, respectively. We define the polar plane (constant z) to be the axial plane and accordingly the z axis to be the rotation axis.
Experimental results are presented in section 4.4.

4.3.5 Cross-channel bilateral filter

The dark-field projections and diffusion coefficient reconstructions inherently have increased noise compared to transmission and attenuation coefficient, respectively [49]. The similarity of object edges in both channels can be exploited with various methods to improve image quality [65, 66]. The tool most commonly used in our pipeline is a cross-channel bilateral filter due to its speed, preservation of edges, and comparatively low number of parameters [67, 68].

The common bilateral filter is a moving-window filter, similar to a Gaussian or median filter. In that, it works on a patch of the input image and computes the value of the output image central to the patch. In contrast to the mentioned filters the bilateral filter "identifies" and preserves object edges. It works by applying both index-wise and value-wise weights to the elements of the patch, such that data points further away from the patch center both in distance and value are considered less. Formally, we define the bilateral filter $\mathcal{B}_{f,g}(\boldsymbol{x})$ on $\boldsymbol{x} = (x_i) \in \mathbb{R}^N$,

$$\left(\mathcal{B}_{f,g}(\boldsymbol{x})\right)_i = \frac{\sum_{j\in\Omega_i} w_j x_j}{\sum_{j\in\Omega_i} w_j} \quad \text{with} \quad w_j = f(x_j, x_i)\, g(j,i)\,, \tag{4.28}$$

with Ω_i the patch around index i, f the value-wise weighting function, and g the index-wise (or "spatial") weighting function. Typical functions are a Gaussian potential for f and either a Gaussian potential or constant 1 for g. For the latter case $g(j,i) = 1$, the strength of the spatial filter is controlled with the size of the patch Ω_i.

This single-channel filter is already powerful and widely used in signal processing. Its evaluation is computationally somewhat expensive because $f(x_j, x_i)$ can not be precomputed in contrast to $g(j,i)$. It is therefore commonly computed on a GPU.

We extend the bilateral filter to two input channels $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^N$. Input $\boldsymbol{x}$ is assumed to be less noisy and/or of higher resolution than $\boldsymbol{y}$. The output of the cross-channel filter is a filtered version of $\boldsymbol{y}$ with the patch weights w_j determined on $\boldsymbol{x}$. We therefore expect that both channels have similar features (albeit with different contrast), which is generally the case in attenuation and diffusion coefficient. Formally, we define the cross-channel bilateral filter $\mathcal{B}^{\text{cc}}_{f,g}(\boldsymbol{x}, \boldsymbol{y})$,

$$\left(\mathcal{B}^{\text{cc}}_{f,g}(\boldsymbol{x}, \boldsymbol{y})\right)_i = \frac{\sum_{j\in\Omega_i} w_j y_j}{\sum_{j\in\Omega_i} w_j} \quad \text{with} \quad w_j = f(x_j, x_i)\, g(j,i)\,. \tag{4.29}$$

$\mathcal{B}^{\text{cc}}_{f,g}$ is identical to $\mathcal{B}_{f,g}$ except for y_j in (4.29) instead of x_j in (4.28). We implement the cross-channel bilateral filter in `JAX`. Experimental results are shown in section 4.4.

4.4 Experimental results

We show experimental results of the processing pipeline presented in the previous sections of this chapter.
First we concentrate on the "small sample" pipeline using sample-free detector area and analyze the temporal stability of the method, especially the PCA model. The presented correction for slow drift of the reference phase is applied on reconstructions of the refractive index decrement both to show its validity and to isolate potential artifacts from a varying PCA vibration model. Next, we investigate the robustness with respect to the size of the sample-free detector area to determine interferometer vibrations.
The results of the processing pipeline for large samples (i.e. no sample-free area) are presented and discussed on measurements of a human-sized thorax phantom filled with various inserts. We show the impact of the presented rotation correction, phase drift correction, and cross-channel bilateral filter.

4.4.1 Temporal stability & phase drift

Temporal stability We analyze the temporal stability of the PCA model by performing the sample processing on several air-scans with increasing time spacing to a scan of the rod phantom. The resulting reconstructions are shown in Fig. 4.6. Five air-scans are recorded and processed with the proposed pipeline to obtain respective PCA vibration models. The temporal distance to the sample-scan is one, five,

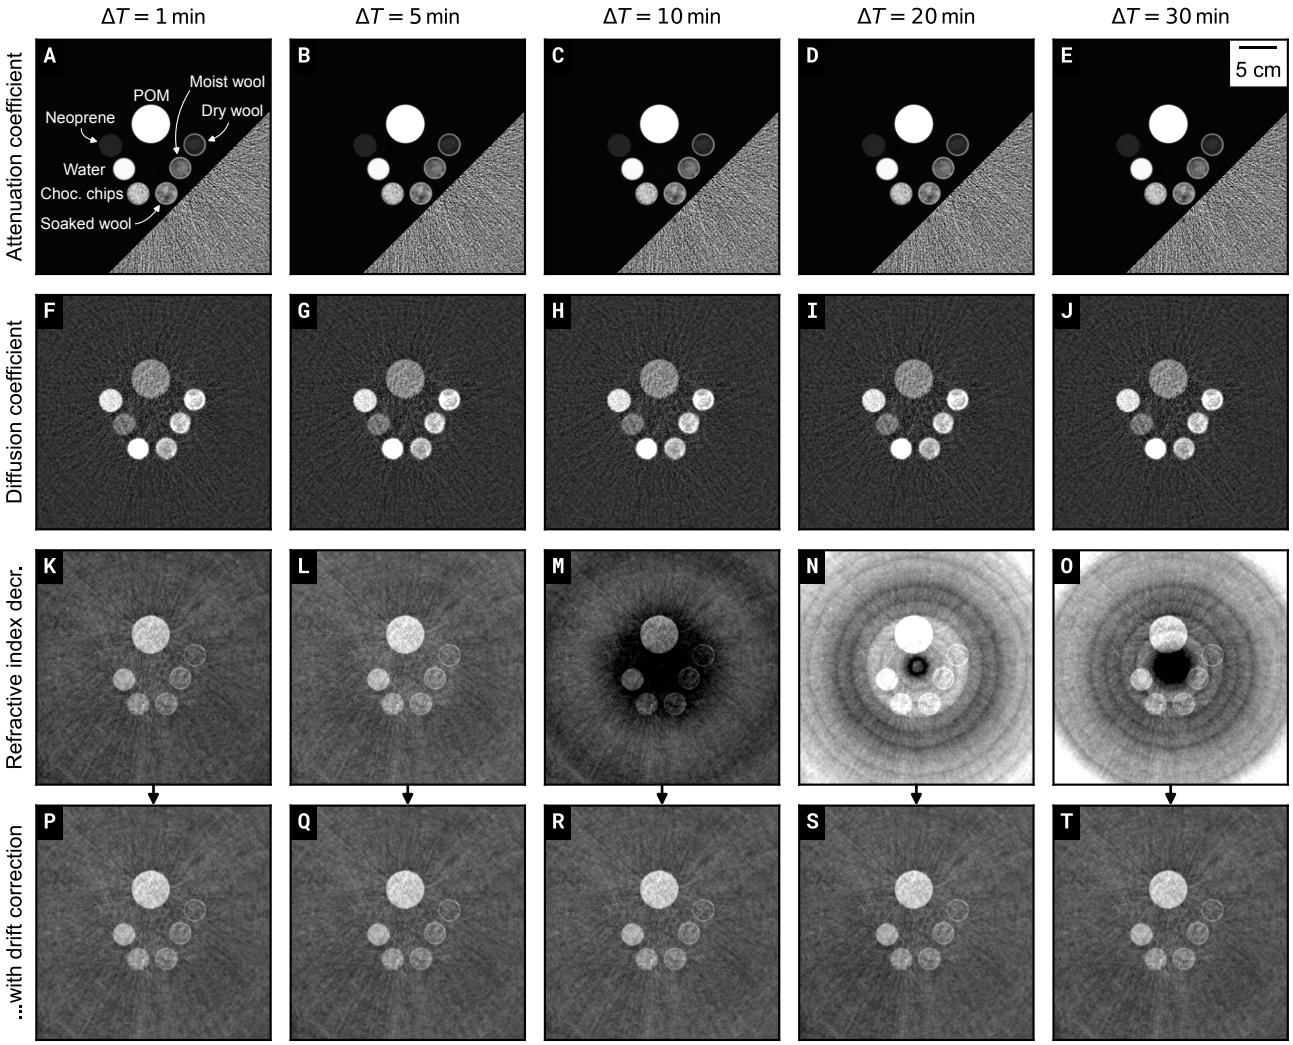

Figure 4.6: Reconstructions of the rod phantom using PCA models from different air-scans with increasing time distance ΔT to test the temporal stability of the model and the phase drift correction. None of the reconstructions shows vibrations artifacts (as seen in Fig. 3.20). The refractive index slices in **M**-**O** show circular bands stemming from slow thermal drift of the reference phase flat-field. Because the drift is much slower ($\approx 10^3$ s) than the duration of a scan ($\approx 10^0$ s) the effect is concentric. We apply the correction method from section 4.3.4 indicated by the small black arrows and obtain slices **P**-**T**. They all look identical and show neither drift nor vibration artifacts. The windowing is **A**-**E**: [-0.05, 2.2] $\times$ 10^{-1} cm^{-1}; **F**-**J**: [-2, 8] $\times$ 10^{-2} cm^{-1}; **K**-**O**: [-2, 4] $\times$ 10^{-1} cm^{-1}. Narrow window in **A**-**E**: [-5, 5] $\times$ 10^{-3} cm^{-1}.

ten, 20, and 30 minute(s), respectively. The sample-scan is processed with the "small sample pipeline" discussed in section 4.2.2, i.e. by fitting the per-shot vibration coefficients $(\alpha^s, \beta^s, \gamma^s)$ in the sample-free detector area at the left- and right-hand side of the detector. The only difference between the reconstructions in Fig. 4.6 are the PCA vibration models $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ obtained from each air-scan. If the vibration model was not valid for the sample-scan data due to thermal drift effects we expect "vibration" artifacts similar to Fig. 3.20.

None of the axial slices shown in Fig. 4.6 show such artifacts. The attenuation coefficient $\hat{\mu}$ in **A**-**E** and diffusion coefficient $\hat{\varepsilon}$ in **F**-**J** look identical for all air-scans and time differences. If there was a slight change in the vibration characteristics they appear to still be modeled by an air-scan 30 minutes apart.

The slices **K**-**O** of the refractive index decrement $\hat{\delta}$ do not look identical: especially **M**-**O** are overlaid with strong concentric features whose contrast is larger than the sample contrast. They originate from slow thermal drift and deformation of the interferometer which results in slow changes of the reference phase $\hat{\phi}$. This change apparently can not be modeled by the PCA vibration model $\mathcal{C}$ which we fit to the sample-free detector area. Therefore the difference to $\hat{\phi}$ ends up in the projections $\hat{\Phi}$ and ultimately in the reconstructed refractive index $\hat{\delta}$. The drift appears to be noticeable after 10 minutes in this comparison, or roughly 10^3 seconds in general. In any case, this timeframe is much longer than the typical duration of an axial scan at one second or even a helical scan at $\approx 10^1$ seconds.

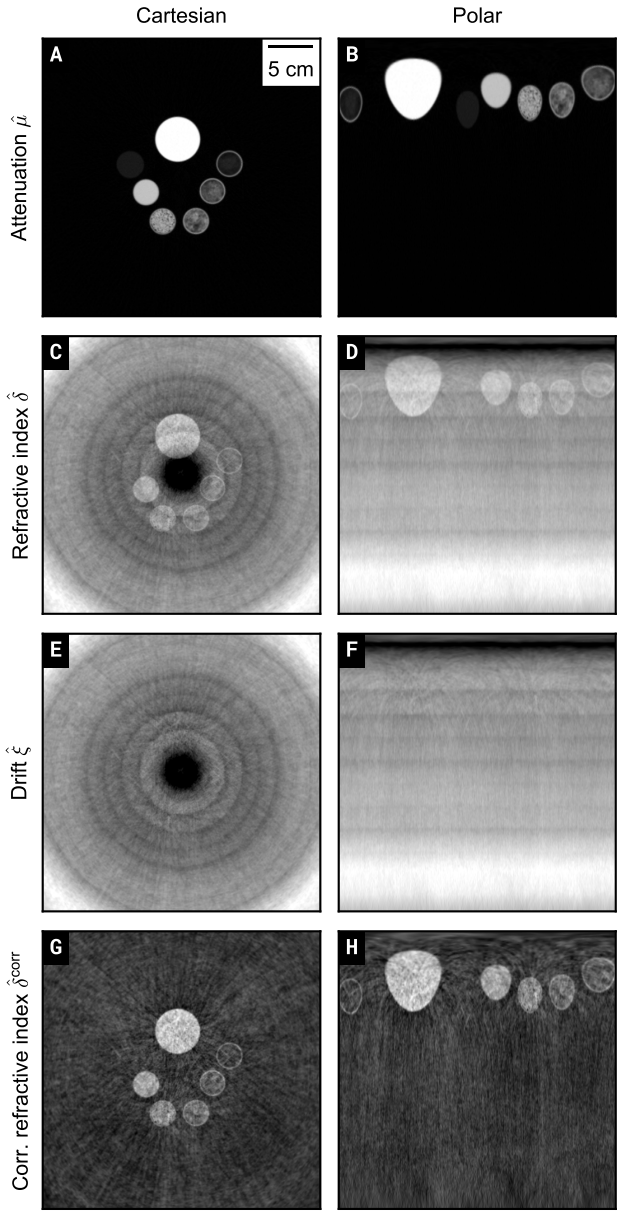

Figure 4.7: Axial slices in cartesian and polar coordinates during drift correction. **A**,**B** show the attenuation coefficient $\hat{\mu}$ in cartesian and polar corrdinates, respectively. **C**,**D** show the refractive index decrement $\hat{\delta}$ with strong concentric bands before drift correction. **E**,**F** represent the difference $\hat{\xi}$ between refractive index $\hat{\delta}$ and scaled attenuation $\lambda\hat{\mu}$. **G**,**H** show the corrected refractive index $\hat{\delta}^{\text{corr}}$ obtained by $\hat{\delta} - \langle\hat{\xi}\rangle_\theta$.

The effective interferometer phase $\hat{\phi}$ therefore is almost constant during the scan and the artifacts in the tomographic reconstruction $\hat{\delta}$ are perfectly concentric.

Drift correction As discussed in section 4.3.4 we use this condition to approximate the effect and compute corrected images $\hat{\delta}^{\text{corr}}$. Fig. 4.7 illustrates the method: we subtract the scaled reconstructed attenuation coefficient $\hat{\mu}$ from the refractive index decrement $\hat{\delta}$ and mean over the polar angle θ to obtain the drift estimate $\langle\hat{\xi}\rangle_\theta$. It is subtracted from $\hat{\delta}$ giving a corrected refractive index $\hat{\delta}^{\text{corr}}$ without concentric bands. The axial slices **P**-**T** in Fig. 4.6 look identical. The method apparently works well for small samples of which we can already compute a tomographic reconstruction of $\boldsymbol{\mu}$ and $\boldsymbol{\delta}$ apart from the drift. The approximation $\hat{\boldsymbol{\delta}} \approx \lambda\hat{\boldsymbol{\mu}}$ discussed in section 4.3.4 seems to be accurate enough in the DFCT.

4.4.2 Impact of sample-free detector area size

While the "small-sample pipeline" is useful for analyzing the isolated effect of the vibration model, phase drift correction, and other effects discussed in section 3.4.4, it is not viable for reconstructing large samples using the whole FOV of the DFCT.

Image quality We analyze the impact of the size of the sample-free detector area used to fit $(\boldsymbol{\alpha}^{\text{s}}, \boldsymbol{\beta}^{\text{s}}, \boldsymbol{\gamma}^{\text{s}})$ on the quality of the tomographic reconstructions of the rod phantom. The air- and sample-scan are identical to section 3.4.4. The full detector width along all G_2 tiles as seen in Fig. 3.9 is 545 pixels and 77.2 cm, corresponding to an isocenter FOV of 42.4 cm. We employ the "small sample pipeline" and manually restrict the sample-free detector area to 50 %, 30 %, 20 %, 10 %, and 5 % of the detector width, evenly distributed to the left- and right side of the sample. The resulting tomographic reconstructions of the attenuation coefficient $\hat{\mu}$, the diffusion coefficient $\hat{\varepsilon}$, and the refractive index decrement $\hat{\delta}$ are shown in Fig. 4.8.
The left-most column **A**,**F**,**K** represents the high-quality images with plenty of area to fit the vibrations $(\hat{\boldsymbol{\alpha}}^{\text{s}}, \hat{\boldsymbol{\beta}}^{\text{s}}, \hat{\boldsymbol{\gamma}}^{\text{s}})$. While the image quality varies between modalities each image, is free of vibration-specific artifacts as previously in section 3.4.4. The relative width of the sample-free detector area decreases with the right-hand columns. In the attenuation coefficient $\hat{\mu}$, slight artifacts begin to appear for 10 % relative width in **D** and are stronger for 5 % relative width in **E**. In both cases they are only discernible in the narrow value window, not in the window capturing the whole sample. In **F**-**J** diffusion coefficient $\hat{\varepsilon}$, slight circular features occur only at 5 % relative width in **J**. Note that they are larger in relative contrast than in **E** due to the full value window. The **K**-**O** refractive index decrement $\hat{\delta}$ generally is most susceptible to decreasing sample-free area. Vignetting at the FOV border starts at 20 % in **M** and significant artifacts are visible for 5 % relative width in **O**.

Orthogonality We use the respective "orthogonality" of the PCA model components $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ depending on the sample-free detector area size as a

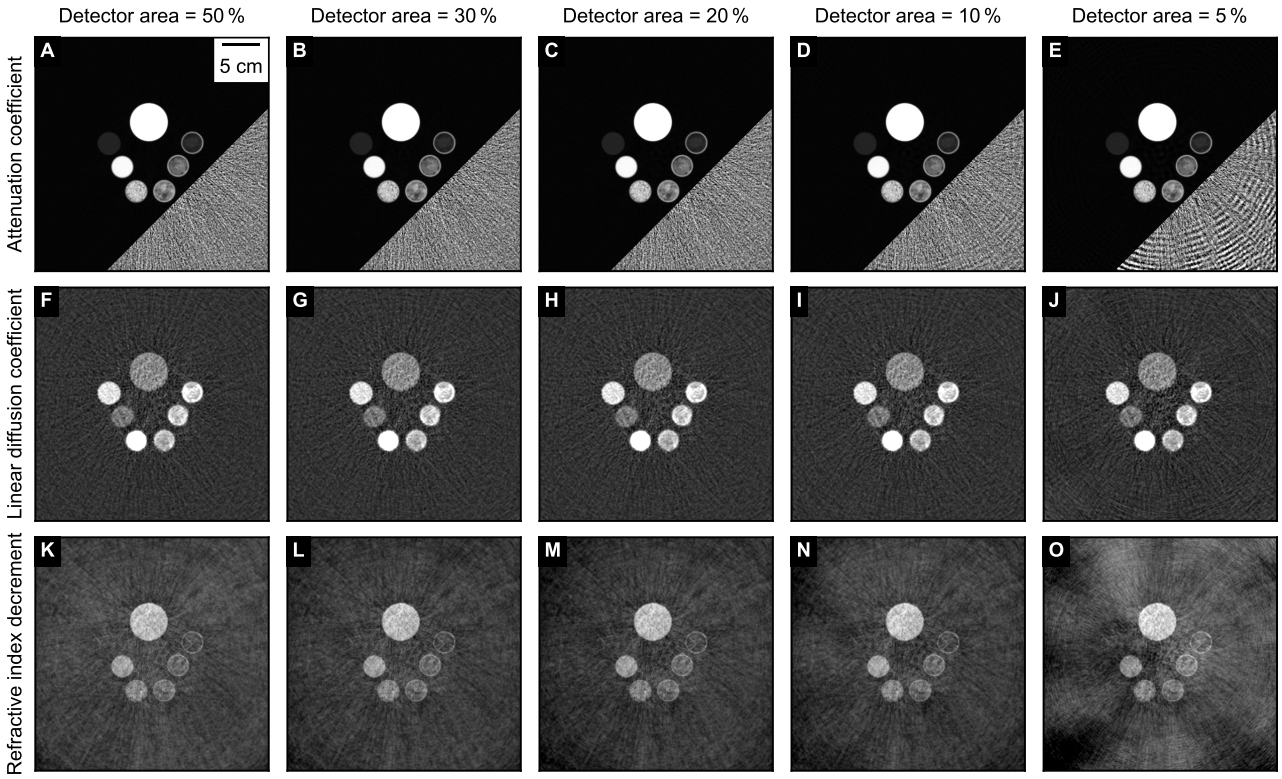

Figure 4.8: Reconstructions of a phantom using different amounts of the detector width in the "small sample pipeline" to determine the vibrations coefficients per exposure $\boldsymbol{\alpha}^{\mathrm{s}}$, $\boldsymbol{\beta}^{\mathrm{s}}$, and $\boldsymbol{\gamma}^{\mathrm{s}}$ [47]. The percentage above each column indicates the relative amount of detector width used for the fit. The image quality is constant down to 10 % detector width in attenuation **A**-**D** and diffusion coefficient **F**-**I**. The refractive index decrement is more sensitive due to the higher amount of components in the PCA model. It shows deteriorating quality for 20 % in the outer regions **M** and 10 % close to the sample **N**. For 5 % detector width **E**,**J**,**O**, all channels suffer in image quality. The windowing is **A**-**E**: [-0.05, 2.2] $\times$ 10^{-1} cm^{-1}; **F**-**J**: [-2, 8] $\times$ 10^{-2} cm^{-1}; **K**-**O**: [-1, 3] $\times$ 10^{-1} cm^{-1}. Narrow window in **A**-**E**: [-5, 5] $\times$ 10^{-3} cm^{-1}.

measure of the stability of the coefficient fit in the latter. As discussed in section 3.2.3 the vibration models

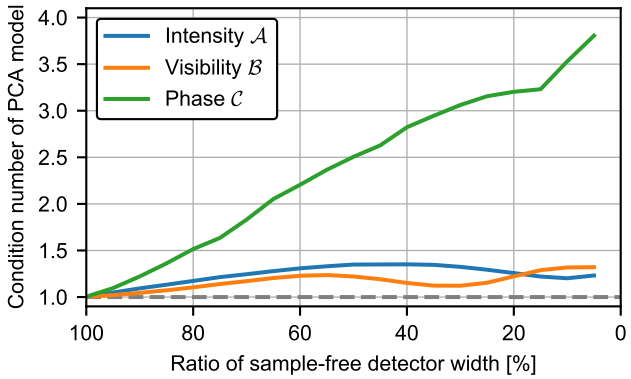

Figure 4.9: Condition number as a measure for orthogonality of the PCA model $(\boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{C}})$ inside the sample-free detector area used for estimating vibration coefficients $\boldsymbol{\alpha}^{\mathrm{s}}$, $\boldsymbol{\beta}^{\mathrm{s}}$, and $\boldsymbol{\gamma}^{\mathrm{s}}$ during a sample-scan. The components are orthogonal (condition number 1) only for full width. When using less detector area the condition number increases, decreasing the stability of the coefficient fit. Figure adapted from [47].

are orthogonalized and scaled wrt. the reference intensity $\hat{\boldsymbol{I}}$ and visibility $\hat{\boldsymbol{V}}$, respectively, such that the Hessian matrix of the cost function is as diagonal and equally scaled as possible, which is beneficial for convergence.

All this is performed on the full detector area. However, we use only a fraction of the detector area for determining the sample-scan vibration coefficients $\hat{\boldsymbol{\alpha}}^{\mathrm{s}}$, $\hat{\boldsymbol{\beta}}^{\mathrm{s}}$, and $\hat{\boldsymbol{\gamma}}^{\mathrm{s}}$. The vibration models restricted to the sample-free detector area are not orthogonal anymore. We compute the condition number κ for each vibration model $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ inside a fraction of the detector area to assess their "orthogonality" with a single number, respectively. The condition number $\kappa\left(\boldsymbol{X}\right)$ is given as the ratio $\sigma_{\mathrm{max}}\left(\boldsymbol{X}\right)/\sigma_{\mathrm{min}}\left(\boldsymbol{X}\right)$ of the maximal and minimal singular values of $\boldsymbol{X}$. The condition numbers of the vibration models depending on the relative detector area are shown in Fig. 4.9. If $\kappa = 1$ the models are orthogonal. This is apparently only the case for 100 % detector area. When less area is used the condition number increases. For intensity $\boldsymbol{\mathcal{A}}$ and visibility $\boldsymbol{\mathcal{B}}$ this increase is pretty small and $\kappa \leq 1.5$ down to 5 % width. The phase model $\boldsymbol{\mathcal{C}}$ is more sensitive to

the area decrease. Its condition number increases almost fourfold for 5 % relative width.
The prominent sensitivity of the phase vibration model $\mathcal{C}$ is in accordance with the comparison in Fig. 4.8, in which the axial slices of the refractive index decrement $\hat{\boldsymbol{\delta}}$ in **K**-**O** exhibit the strongest quality reduction with decreasing relative detector area for the coefficient fit.

4.4.3 Pre-processing estimations

We present the two pre-processing steps of the pipeline for large samples discussed in section 4.2.2. The estimation of the sample transmission $\widetilde{\boldsymbol{T}}$ is shown in Fig. 4.10 and the approximation of the global flux variation $\widetilde{\alpha}_0^{\mathrm{s}}$ in Fig. 4.11. Both are crucial for the alternating optimization of the model cost function (4.7) as a good initialization.

Both are conducted on an axial scan of the commercially available "Lungman" phantom ("Multipurpose Chest Phantom N1" from Kyoto Kagaku). This object, visible in the annotated photograph Fig. 2.10, is a life-sized model of a human thorax including surrogates for spine and rib bones. It can be filled with various inserts and has been used extensively for the investigations of dark-field radiography and computed tomography [20, 69, 70]. In the example presented here, the thorax is filled with a neoprene insert, a POM cylinder, and three Falcon tubes containing water, air, or chocolate chips. The neoprene insert has similar porosity as healthy lung tissue and produces dark-field contrast accordingly, but otherwise has low attenuation. POM and water on the other hand should exhibit only attenuation contrast because they are homogeneous on the length scales relevant to small-angle scattering. Chocolate sprinkles create both attenuation and dark-field signal.

Transmission estimate We illustrate the transmission estimation in Fig. 4.10. **A** shows the sinogram of the raw data $\hat{\boldsymbol{y}}^{\mathrm{s}}$ normalized by the reference intensity $\hat{\boldsymbol{I}}$. The overall strongly attenuating thorax phantom fills approximately two thirds of the detector width. Even more is taken by the patient couch. The data is modulated with the interferometer fringe pattern, especially visible in the remaining sample-free detector area. As shown in **B** we perform tomographic reconstruction via FBP on the normalized projections to obtain an approximation of the attenuation coefficient $\widetilde{\boldsymbol{\mu}}$. It clearly resembles the thorax phantom but is superimposed with circular structures ("fringe artifacts") especially in the center. We compute the tomographic forward-projection $\exp\left(-\mathcal{R}\left[\widetilde{\boldsymbol{\mu}}\right]\right)$ to obtain the transmission estimate $\widetilde{\boldsymbol{T}}$ in **C**. It has slightly reduced lateral resolution but also a strongly reduced interferometer fringe pattern.

Flux estimate Afterward we find an approximation for the per-exposure flux variation $\widetilde{\boldsymbol{\alpha}}_0^{\mathrm{s}}$ illustrated in Fig. 4.11. The starting point is the raw data $\hat{\boldsymbol{y}}^{\mathrm{s}}$ normalized with the reference intensity $\hat{\boldsymbol{I}}$ and the just acquired transmission estimate $\widetilde{\boldsymbol{T}}$, shown in **A**. Visually, it almost resembles "flat" air-scan data $\hat{\boldsymbol{y}}$ apart from the remaining edges of the sample. The latter originate from the reduced spatial resolution of $\mathcal{R}\left[\mathcal{R}^{-1}\left[\,\cdot\,\right]\right]$ mentioned before. The mean of each exposure is plotted in **B** and already shows the characteristic high-frequency modulation of the flux combined with the low-frequency gantry influence. Its low-frequency component contains some residual sample features compared to the flux $\hat{\boldsymbol{\alpha}}_{00}$ from an air-scan in **C**. We therefore combine the high-frequency component of **B** (blue) with the low-frequency component of **C** (orange) to construct the flux estimate $\widetilde{\boldsymbol{\alpha}}_0^{\mathrm{s}}$ in **D**. It has the correct slow gantry movement (adapted to the encoded starting angle of the scan) modulated with the fast vibrations of the on-board components.

4.4.4 Large sample pipeline

We present tomographic reconstructions of the thorax phantom in Fig. 4.12 at different stages of the processing pipeline discussed in section 4.2.2.

Pipeline stages Fig. 4.12 **A**-**D** show the linear attenuation coefficient $\hat{\boldsymbol{\mu}}$. It has the highest image quality overall of the three contrast channels. With only pre-processing estimates and optimized global phase $\hat{\gamma}_0^{\mathrm{s}}$ in **A** the image quality is already quite good apart from fringe features in the center. The latter are gone in **B** after fitting all interferometer vibrations and the correction for higher-orders in **C** looks identical. The rotation correction in **D** does sharpen the image slightly at the cost of increased high-frequency noise especially below spine bone and table. In all slices it is apparent that the spatial resolution decreases with increased distance from the rotation axis in the center. The tips of the patient table are blurred in rotation direction.
The linear diffusion coefficient $\hat{\varepsilon}$ in **E**-**I** has overall increased noise levels and stronger fringe/vibration artifacts than the attenuation coefficient $\hat{\boldsymbol{\mu}}$. After pre-processing and global phase in **E** the thorax phantom and neoprene insert are discernible but strongly modulated by artifacts. They are progressively reduced by **F** including all vibrations and **G** removing higher-order visibility components. The impact of the rotation correction in **H** is more subtle. It mostly reduces streak artifacts originating from the spine in the center and the table's tips. Finally, we employ a projection-based correction of beam-hardening-induced visibility reduction in **I** as mentioned in section 2.3.1. It is calibrated to a single material (typically POM) and is discussed in more detail in [22]. In the thorax phantom it

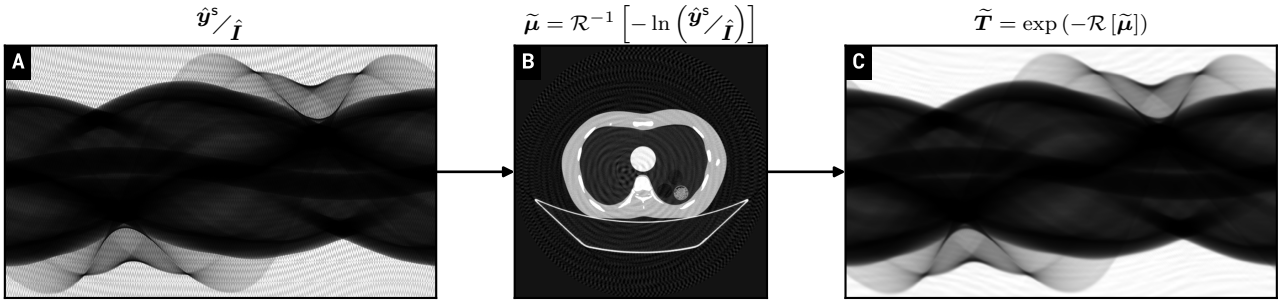

Figure 4.10: Estimating sample transmission by forward-projecting the FBP of normalized raw projections. **A** shows the unprocessed measurement $\hat{\boldsymbol{y}}^{\mathrm{s}}$ normalized with the reference intensity $\hat{\boldsymbol{I}}$. It is reconstructed to an estimate of the attenuation coefficient $\widetilde{\boldsymbol{\mu}}$ in **B** with FBP. This estimate is in turn forward-projected and gives the transmission estimate $\widetilde{\boldsymbol{T}}$ in **C**, in which fringes (and vibrations, although not discernible here) are strongly suppressed. This method is also used on intermediate sample projections $(\hat{\boldsymbol{T}}, \hat{\boldsymbol{D}}, \hat{\boldsymbol{\Phi}})$ obtained with sliding-window SPR during alternating optimization of large samples.

reduces erroneous diffusion signal in the tissue surrogate and central POM cylinder. The bone surrogates still show significant diffusion contrast, which is caused in part by incorrect calibration as well as the surrogate material probably exhibiting some dark-field scattering itself. Even still, the high scattering power of the large neoprene insert and the small tube with chocolate sprinkles is clearly discernible in **I**.

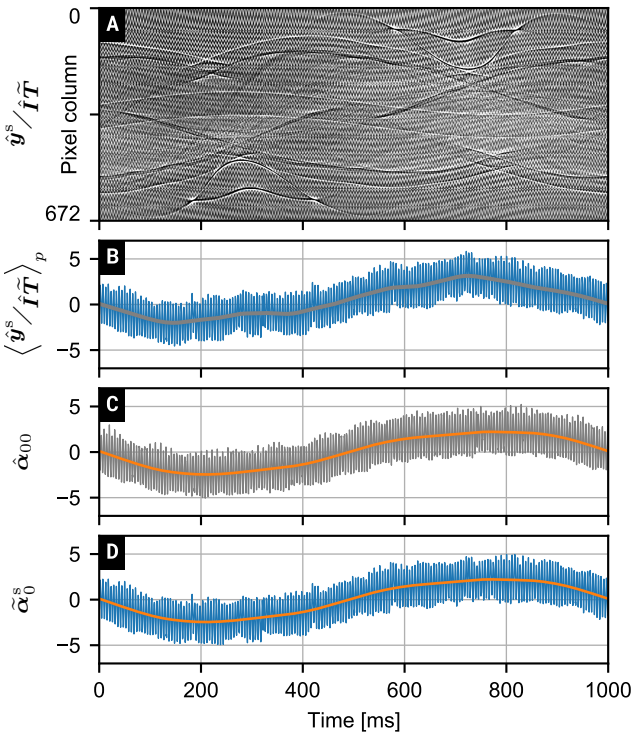

Figure 4.11: Estimating flux of sample-scan data before phase retrieval. **A** shows the sample-scan sinogram $\hat{\boldsymbol{y}}^{\mathrm{s}}$ normalized with the reference intensity $\hat{\boldsymbol{I}}$ and pre-processing transmission estimate $\widetilde{\boldsymbol{T}}$. We combine the **B** high-frequency mean of this sinogram with the **C** low-frequency mean of the air-scan flux $\hat{\boldsymbol{\alpha}}_{00}$ to estimate the **D** flux $\widetilde{\boldsymbol{\alpha}}_0^{\mathrm{s}}$ of the sample-scan measurement. The windowing in **A** is [0.8, 1.2]. The plots in **B**,**C**,**D** are given in percent %.

The refractive index decrement $\hat{\boldsymbol{\delta}}$ is depicted in **J**-**N**. It has the worst image quality overall concerning noise and structural artifacts like streaks and concentric bands. After pre-processing and optimizing the global phase in **J** the thorax phantom is barely discernible. Large ramps overlay the tomographic reconstruction and have higher contrast than the sample itself. They are removed in **K** after optimizing all interferometer vibrations. The noise level is still quite high and the broad concentric bands from slow thermal drift of the reference phase $\hat{\phi}$ dominate. Due to low statistics (and a somewhat naive Gaussian noise model) the noise around the central spine is very high. The reduced rotation artifacts in **M** are therefore hardly appreciable. Only after estimating and removing the thermal drift in **N** the thorax phantom really becomes discernible. Especially the patient table creates strong streak artifacts.

Cross-channel bilateral filter The final images of the diffusion coefficient $\hat{\varepsilon}$ and the refractive index decrement $\hat{\boldsymbol{\delta}}$ in Fig. 4.12 **I**,**N** can be further improved by employing the cross-channel bilateral filter $\mathcal{B}^{\mathrm{cc}}$ discussed in section 4.3.5. We use the attenuation coefficient $\hat{\boldsymbol{\mu}}$ from Fig. 4.12 **D** as the reference image to compute value-wise weights and in that determine object edges. Axial slices of filtered images $\mathcal{B}^{\mathrm{cc}}_{f,g}(\hat{\boldsymbol{\mu}}, \hat{\varepsilon})$ and $\mathcal{B}^{\mathrm{cc}}_{f,g}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\delta}})$ are shown in Fig. 4.13. We use a Gaussian potential for the value-wise function f and a box-like kernel $g(j,i) = 1$ for the spatial filter.
The filtered diffusion coefficient $\mathcal{B}^{\mathrm{cc}}_{f,g}(\hat{\boldsymbol{\mu}}, \hat{\varepsilon})$ in **B** has significantly reduced noise and streaks, especially in the center around the spine. The edges of the spine bone are correctly preserved. Next to it, the thin gap between rib bones and neoprene insert becomes visible. The underestimation of $\hat{\varepsilon}$ as a dark streak is not removed by the filter but seems more dominant due to the reduced surrounding high-frequency noise. The neoprene insert does look more homogeneous after filtering. However there is still a medium-frequency

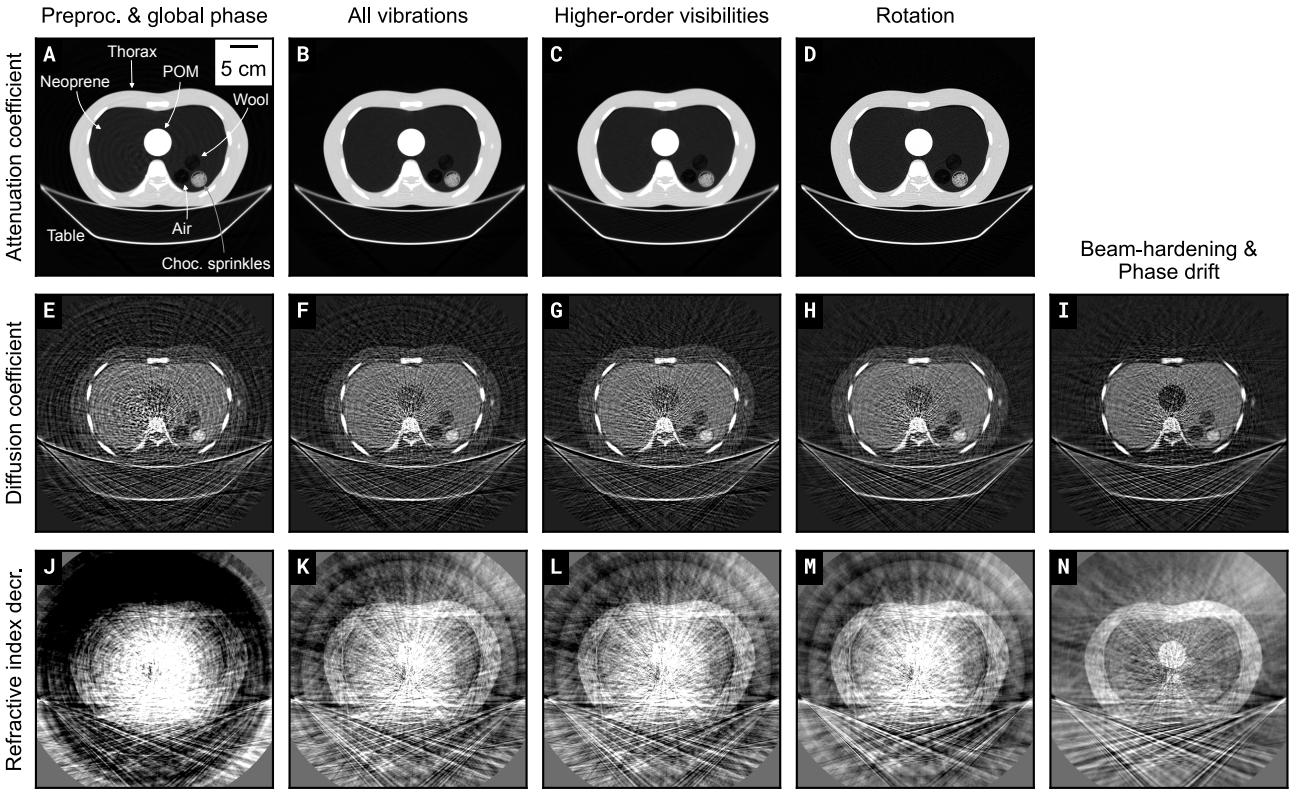

Figure 4.12: Tomographic reconstructions of the thorax phantom at different stages of the processing pipeline. **A**-**D** The first row depicts the attenuation coefficient $\hat{\boldsymbol{\mu}}$. Its quality is already reasonable in the beginning after pre-processing estimates $(\widetilde{\boldsymbol{T}}, \widetilde{\boldsymbol{\alpha}}_0^{\mathrm{s}})$ and optimized global phase $\hat{\gamma}_0^{\mathrm{s}}$. **E**-**I** The second row shows the diffusion coefficient $\hat{\boldsymbol{\varepsilon}}$. It contains fringe artifacts until the higher-order visibility terms are estimated and removed in **G**. **I** The right-most image illustrates the beam-hardening correction which is only applied on $\hat{\boldsymbol{D}}$. **J**-**N** The third show depicts the refractive index decrement $\hat{\boldsymbol{\delta}}$. Its quality is generally the worst of all contrast modalities. The noise level is high and the slow phase drift creates concentric bands whose contrast is stronger than the sample. Only in the last step **N** which estimates and removes the drift, the sample becomes visible. The dominant remaining artifacts in **I** and **N** are streaks from the table's tips. They are far away from the rotation axis and their projected movement is therefore very large. The windowing is **A**-**E**: [-0.05, 2.5] $\times$ 10^{-1} cm^{-1}; **F**-**J**: [-0.2, 1.5] $\times$ 10^{-1} cm^{-1}; **K**-**O**: [-3, 4] $\times$ 10^{-1} cm^{-1}.

modulation visible. It is a combination of rotation artifacts, low statistics near the spine, and potentially slightly incorrect interferometer vibrations. At the border of the central POM cylinder is a bright spot that originates from the voxel sitting right on the sharp, high-contrast edge in the attenuation channel. The value-wise weights f make the filter kernel effectively very small at the border which makes such artifacts possible. Overall, the structure of the POM edge is sharp and has no overshoots otherwise. The streaks inside the patient table are reduced by the filter,
The overall image quality of the refractive index decrement $\mathcal{B}^{\mathrm{cc}}_{f,g}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\delta}})$ in **D** is also significantly improved. The high-frequency noise in the center near the POM cylinder and the spine is reduced and the surrounding tissue surrogate is more homogeneous. The streaks in the patient table are greatly reduced. However, the artifacts that do remain seem more prominent. There is a distinct streak originating from the right-hand table tip through the center of the thorax, and multiple streaks parallel to the straight table undersides. Furthermore, despite the drift correction (and vibration optimization), there are some large gradients spanning the axial slice.

4.5 Discussion

In this chapter we presented processing and reconstruction pipelines for phase-retrieval and tomographic reconstructions of sample measurements in the DFCT.

The main challenge is determining the interferometer state at the time of the scan in order to isolate the sample's projections in transmission, dark-field, and differential phase. We assume that the spatial fluctuations in intensity, visibility, and phase caused by system vibrations can be described as a linear combination of the low-parameter models acquired with PCA on an air-scan as discussed in chapter 3. The coefficients of the linear combination per exposure have

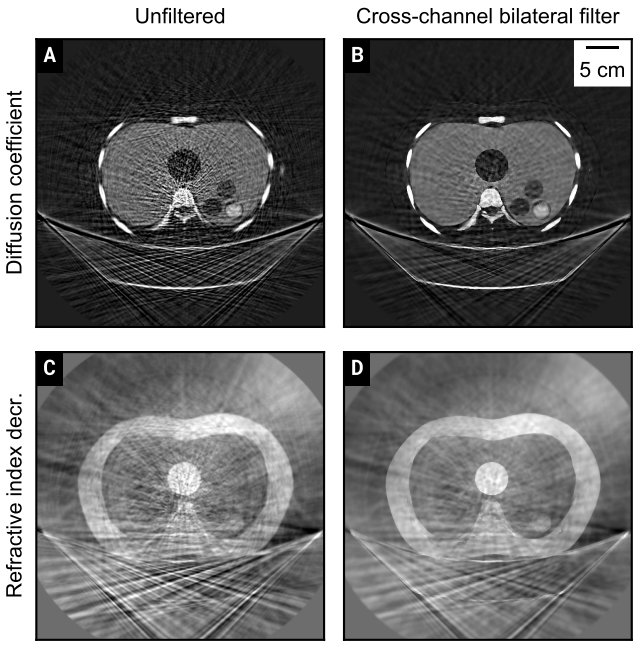

Figure 4.13: Cross-channel bilateral filter $\mathcal{B}^{\mathrm{cc}}(\hat{\mu}, \cdot)$ on thorax phantom with reference $\hat{\mu}$ in **A**,**B** diffusion coefficient $\hat{\varepsilon}$ and **C**,**D** refractive index decrement $\hat{\delta}$. In both modalities the high-frequency noise is strongly suppressed while sharp edges are preserved. The remaining artifacts become more visible, such as medium-frequency "blobs" in **B** or a streak from the table's tip in **D**. The windowing is identical to Fig. 4.12.

to be updated. We formulate a data model with the sample projections and interferometer state, and optimize the resulting likelihood function wrt. the vibration coefficients and ultimately the sample projections.

If the sample is reasonably small we can estimate the coefficients of the vibration model from the sample-free interferometer area before doing signal retrieval of the sample. The fit of the coefficients is iterative, starting with few dominant terms and incrementally going to the full model from the air-scan. Due to good initialization and the small number of parameters the fit process is stable. We can therefore estimate the interferometer state within the sample area from data around the sample.

With an estimation of the interferometer state we perform phase retrieval on the sample measurement. The DFCT is rotating continuously and the acquired measurement of one axial scan contains only one exposure per projection angle and stepping position. Strictly speaking this is not enough information to correctly separate the three contrast channels, as we only measure one projection per gantry angle but want to determine three projections from it. The modulation of the stepping curve phase via the system vibrations is sufficiently fast compared to the gantry rotation, such that we can assume the sample's projection constant over a small set of subsequent exposures. With this assumption and the interferometer state determined with sample-free area, we perform SPR on each angular window separately and obtain the sample's projections in transmission, dark-field, and differential phase. Tomographic reconstructions via FBP are performed on each of them separately giving volumetric images of linear attenuation coefficient, linear diffusion coefficient, and refractive index decrement.

The purpose of this research project and the DFCT is ultimately dark-field CT on human patients. Given the geometry of the DFCT and limitations of grating fabrication we probably will not be able to use sample-free interferometer area to determine the interferometer state when imaging a human. We therefore employ an alternating minimization scheme of the likelihood cost function wrt. the vibration model coefficients and sample projections. The optimization of the vibration coefficients is still performed with an iterative first-order solver. However, the model does include the current estimate of the sample projections. After convergence we perform sliding-window SPR with the current interferometer state and update the sample projections. These two steps are repeated until satisfactory overall convergence. We compute estimations for the transmission and the per-exposure intensity fluctuation in a pre-processing step to initialize the alternating optimization.
The obtained transmission, dark-field, and differential phase projections of the large sample are improved in several post-processing steps. We estimate and remove the higher-order terms of the Fourier series with the sample's dark-field correctly adapted, suppress artifacts from sample rotation, and correct for beam-hardening reduced visibility reduction and slow thermal phase drift, respectively. These effects were ignored for a small sample but are relevant for a larger object.
The image quality of the diffusion coefficient and the refractive index decrement is further enhanced by employing the presented cross-channel bilateral filter with the attenuation coefficient as a reference. It requires object features to be similar between contrast channels.

Overall, the results in this chapter demonstrate that we can compute tomographic reconstructions for attenuation, diffusion, and refractive index decrement for small and large samples in the DFCT. The attenuation objectively has the best image quality. Artifacts from an incorrect interferometer state, higher-order Fourier terms, and sample rotation are small and the noise level is lowest. The diffusion coefficient is more sensitive to vibrations and higher-order terms but provides a novel and functional image contrast. The refractive index decrement generally has the poorest image quality. It is also impacted by slow drift of the reference phase caused by thermal deformation of the gratings, for which we only have an empirical post-processing

correction. It works well for small samples but is limited by systematic artifacts with large samples, such that the drift can not be completely removed.

There are several simplifications and limitations of the current approach.
Most notably, our projection-based pipeline does not use a valid model for continuous acquisition without stepping. The resulting artifacts are suppressed by the proposed rotation correction but not removed. Instead, the IBSIR approach correctly models the acquisition and does not require any approximations, going from raw measurement directly to tomographic reconstruction by including the Radon transform in the model. It is computationally expensive and slow to converge, but could be initialized with the results from the presented pipeline (both in interferometer state and sample reconstructions) and serve as a post-processing filter, removing rotation artifacts and enforcing prior knowledge about the sample (local smoothness and similar object edges).
Once we apply this valid model and obtain improved image quality, other limitations of the current projection-based pipeline will likely become more noticeable. We estimate the higher-order terms of the Fourier series based on the optimization of the first-order model. The full model with higher orders is therefore never used "online" in an iterative optimization. Residual influences might become visible with reduced overall image noise.
The correction for beam-hardening-induced visibility reduction can also be improved. At the moment we determine a global calibration for a single material. With volumetric information, we could perform a segmentation and calculate a more physically sensible estimation of the effect. Furthermore, we do not yet incorporate visibility hardening at all [71].

A logical next step for the sample acquisition overall is moving to a helical trajectory and improved z coverage. This requires a more involved tomographic reconstruction in virtual "wedge" geometry which has been done in the literature for conventional absorption contrast but not in the context of dark-field and phase contrast. We can either perform signal retrieval and rebinning separately or combine them in one step. Another possibility to reconstruct helical scans without the need for other geometries is to use iterative reconstruction, either in the form of statistical iterative reconstruction (SIR) on already separated sample projections or in the form of IBSIR directly on the raw helical measurement. However, the difficulties of the latter concerning convergence and computational complexity will be increased by the larger datasets.

Lastly, it would be beneficial for the research of clinical dark-field CT in general to describe the information encoding and signal retrieval process in a more formalized framework. [59] motivates phase retrieval as a signal demodulation process but does not connect it with tomographic reconstruction and sampling in Radon space. Derivating quantitative measures for sampling quality and resulting image properties such as noise characteristics and undersampling artifacts (which rotation artifacts basically are) would open the possibility for objectively optimal setup design and signal retrieval.

Chapter 5

Conclusion

This thesis develops improved data processing algorithms relevant to clinical dark-field computed tomography with X-rays. It introduces several additions to the conventional model motivated by the physical processes governing the image formation. They are evaluated with measurements from a prototype setup based on a commercial CT gantry. This chapter summarizes the findings and outlines further steps on the path to clinical dark-field computed tomography with X-rays.

Interferometer vibrations We first motivate and develop a model for the spatio-temporal fluctuations we observe on air-scan data acquired with the DFCT. It is a commercial CT system from "Philips N.V". which our group retrofitted with a Talbot-Lau interferometer. The scanner is operated in the original clinical protocols in which the gantry rotation and detector readout are continuous. Starting from the conventional model used for X-ray imaging with a stepped Talbot-Lau interferometer, we introduce 2D polynomial terms for each G_2 tile in visibility and phase which vary for each detector exposure. Lateral movement and rotation of the interferometer gratings lead to variations in the effective phase at the G_2 grating. Due to the finite exposure time the phase changes during detector readout which reduces the system visibility according to our analytically derived model. Defosing of the interferometer from movement of the gratings and the focal spot further reduces the visibility.

Nested linearized phase retrieval We extend the conventional SPR method with these vibration terms and implement a nested optimization scheme in which the parameterization coefficients are optimized wrt. the residuum after the linear solve in SPR. The algorithm is executed on a GPU via `JAX`. In using air-scan data we use projections from a full axial gantry rotation in which the stepping phase on the interferometer varies over 10π and is sufficient for stable phase-retrieval. The method gives us reference intensity, visibility, and phase flat-fields, as well as the estimated spatio-temporal tile-wise fluctuations of visibility and phase.

Principal component analysis These fluctuations are analyzed in PCA to determine their dominant components and reduce the number of parameterization coefficients per exposure while still accurately modeling the vibration. In this manner, for the DFCT we identify two components in visibility and four in phase. The PCA method is then used to determine dominant fluctuations in the intensity contrast of the interferometer. By re-arranging the model with current estimates for reference flat-fields and aforementioned vibrations in visibility and phase we can apply PCA directly and obtain a intensity fluctuation parameterization, which in our case consists of two major components.

Higher orders of visibility Another novelty of the DFCT processing is the consideration of higher-order terms of the Fourier series describing the stepping curve measured behind the G_2 grating. While the conventional model only considers the mean and first-order modulation (effectively assuming the curve to be an offset cosine) we observe significant second- and third-order terms which become noticeable as artifacts if the phase-retrieval is performed over less than one stepping period. In the DFCT they are especially relevant because the triangular G_1 generates a pattern with inherently larger higher orders, and the phase-retrieval for sample reconstruction works on small angular windows in which the phase varies only slightly and often $< 2\pi$. We adapt the vibration models and the first-order SPR to include these higher orders of visibility and use axial air-scan data to determine them, in order to remove them after the vibration fit and compute virtual first-order projections.

Macroscopic grating movement The final addition to the model without a sample is the macroscopic lateral movement of the gratings. In contrast to a lab-based setup where one grating is moved over microscopic distances much smaller than a detector pixel, in the DFCT the gantry rotation induces macroscopic grating displacement. We mathematically motivate the resulting effects from the spatial gradients of the flat-fields and determine them with PCA on the residuum from performing sliding-window SPR on an air-scan. Considering this effect avoids concentric ring artifacts in the tomographic reconstructions especially of the

refractive index decrement, due to the spatial gradient of the phase flat-field with fringes being so large.

Sliding-window phase retrieval In accordance with our projection-based pipeline we use sliding-window SPR to obtain projections from a measured sample in transmission, dark-field, and differential phase. This method uses small angular windows of subsequent exposures to estimate the sample projections in the exposure central to the window, given the interferometer state as a combination of reference flat-fields and vibration model. Due to the continuous rotation of the scanner this method violates the model because it assumes a constant sample in the angular window. However, implemented on GPU it is much faster than iterative methods such as IBSIR and is shown to produce satisfactory results if we employ a set of image enhancement methods.

Processing a large sample The main challenge in processing a scan of a sample is determining the interferometer state at the time of the scan. We assume that the low-dimensional vibration model from the latter is still valid (and show this to be the case for up to 30 minutes between scans) and only update the per-exposure coefficients for each model component. The main optimization loop is a minimization of a likelihood cost function alternating between vibration coefficients and sample projections. In principle, we could employ the same nested scheme as in air-scan processing, performing "inner" sliding-window SPR at each iteration of the "outer" optimization of the vibration coefficients. This has proven both unstable and computationally expensive in practice and we opted for the robust alternating scheme instead.

Additionally, we present two novel image enhancement methods for post-processing: the reduction of rotation artifacts and the estimation of slow thermal drift of the reference phase flat-field, both leading to artifacts in the tomographic reconstruction of the sample projections.

The rotation artifacts stem from the sliding-window SPR implicitly assuming a static sample. We enforce tomographic consistency by forward-projecting the FBP of the sample transmission and performing sliding-window SPR on virtual projections with the transmission estimate removed. This strongly reduces streak artifacts in the final reconstructions of all image modalities.

Despite active thermal regulation of the setup (both the room and on-board components) there is small deformation of the gratings over time, resulting in strong artifacts in the reconstructed refractive index decrement. We estimate the drift as the polar mean of the difference between refractive index decrement and scaled attenuation coefficient, effectively approximating the phase shift with the scaled transmission.

Combined with the temporal stability of the PCA vibration model this gives us a processing pipeline that is stable over time and handles large samples using the full interferometer FOV. Its validity is demonstrated with tomographic reconstructions of both a small rod phantom, and a human-sized thorax phantom filled with various inserts. Employing the presented cross-channel bilateral filter on the final results of the diffusion coefficient and refractive index decrement further increases their contrast-to-noise ratio by using the high quality of the attenuation image.

Formalization of the signal retrieval Despite the satisfactory experimental results of the presented projection-based reconstruction pipeline, the method lacks a formal framework that quantifies the information density in the context of continuous rotation, phase oscillation, and sliding-window phase retrieval. The meta-parameters used by our method such as oscillation frequency and angular window size are empirically found. At the moment we can't predict the magnitude of image errors caused by higher-order visibility terms, insufficient phase sampling in an angular window, or projected lateral sample movement (rotation). There has been work from our group to formalize the signal retrieval [59] but there is not an "end-to-end" model for image quality yet.

Even still, there exist more empirical methods for improved signal retrieval which have not been discussed in this work. According to [59] the sliding-window SPR can be extended from only angular windows to spatial windows, effectively implementing patch-wise SPR. Combined with small, horizontal fringes (instead of vertical fringes in this work) the image quality in axial slices of tomographic reconstructions can be greatly improved. This method (and the others presented in the cited work) introduce even more meta-parameters, further increasing the need for a formal framework of information density in this context.

As an alternative to the presented sliding-window SPR and all its modifications, we could use IBSIR for tomographic reconstruction [44]. It is an iterative reconstruction method that incorporates the Radon transform in its model and therefore correctly expresses the continuous rotation and phase sampling implicitly. Unfortunately, a direct consequence is that this method is computationally very expensive, requiring multiple tomographic back- and forward-projections per optimizer iteration. With the insights from this thesis and the results from tomographic reconstruction (and filtering) as an initialization of IBSIR, it could prove viable to further increase image quality without the need for manually tuned meta-parameters mentioned above.

Helical acquisition A crucial step in the direction of clinical viability of the setup is the transition from axial to helical acquisition, covering the whole human thorax for lung imaging. While the core principles of our pipeline are still valid for helical acquisitions, it does introduce the need for several adjustments. FBP in helical geometry with the focus-centered detector in the DFCT is more complicated. The most viable solution is a rebinning to virtual wedge geometry [72, 73] which has only been explored with conventional absorption contrast in the literature. To obtain valid projections of the dark-field and especially the differential phase, we have to combine the rebinning with sliding-window SPR, fusing them into one bilinear interpolation.

Furthermore, due to dose limitations and the prolonged scan time, the effective dose per axial slice will be significantly lower than for the tomographic reconstructions shown in this work. This increases image noise especially in the diffusion coefficient. Further post-processing filters or an IBSIR post-processing step will be necessary. The processing time will be increased due to the large dataset.

Dual-phase G_1 An orthogonal possibility to increase dose effectiveness are dual-pitch G_1 gratings [74, 75]. This novel design features a macroscopic modulation of the microscopic grating structure. The latter governs the correlation length and corresponding length scales to whose dark-field signal the setup is sensitive, but the former creates a macroscopic intensity modulation at the detector, eliminating the need for the analyzer grating G_2. The macroscopic modulation indeed is decreased by small-angle scattering of the Talbot carpet from the microscopic structure. Without the G_2 absorbing approximately half the radiation, the dose is effectively doubled. However, first investigations in the literature have shown a lower visibility compared to conventional designs. As the noise in the dark-field image is proportional to $V^{-2}I^{-1}$, the visibility of the dual-pitch G_1 has to be at least 70 % of the conventional design in order to be beneficial.

Spectral imaging Spectral imaging with dual-energy or photon-counting devices is a heavily researched topic with increasing commercial adoption in recent years [76] although proposed already in 1976 [77]. It considers the energy dependency of the attenuation coefficient to generate virtual mono-energetic images, material decomposition, and contrast-enhanced images in general. Using energy-resolved measurements in dark-field and phase contrast has been investigated in recent work but is still at a pre-clinical stage [78, 79]. If the DFCT would be combined with a dual-layer or photon-counting detector, we could exploit the improved noise characteristics and correlation between attenuation and phase shift via the projected electron density, ultimately improving image quality.

Clinical studies Finally, the ultimate goal is to provide diagnostic information in clinical practice. To evaluate the potential of the DFCT reader studies are necessary, comparing it to conventional CT with absorption contrast. Furthermore, there exists a prototype for dark-field radiography that has shown great promise for early detection of various lung diseases at low radiation dose [13–17]. The added value of volumetric diffusion information will be measured against this established method.

Bibliography

[1] A. Lohmann and D. Silva, "An interferometer based on the Talbot effect," *Opt. Commun.*, vol. 2, no. 9, 1971.

[2] S. Yokozeki and T. Suzuki, "Shearing interferometer using the grating as the beam splitter," *Appl. Opt.*, vol. 10, no. 7, 1971.

[3] A. Momose *et al.*, "Demonstration of X-ray Talbot interferometry," *Jpn. J. Appl. Phys., Part 2*, vol. 42, no. 7B, 2003.

[4] F. Pfeiffer *et al.*, "Hard-X-ray dark-field imaging using a grating interferometer," *Nat. Mater.*, vol. 7, no. 2, 2008.

[5] L. B. Gromann *et al.*, "In-vivo X-ray dark-field chest radiography of a pig," *Sci. Rep.*, vol. 7, 4807, 2017.

[6] J. Tanaka *et al.*, "Cadaveric and in vivo human joint imaging based on differential phase contrast by X-ray Talbot-Lau interferometry," *Z. Med. Phys.*, vol. 23, no. 3, 2013.

[7] K. Willer *et al.*, "X-ray dark-field imaging of the human lung—a feasibility study on a deceased body," *PLOS ONE*, vol. 13, no. 9, e0204565, 2018.

[8] W. Yashiro, Y. Terui, K. Kawabata, and A. Momose, "On the origin of visibility contrast in x-ray Talbot interferometry," *Optics express*, vol. 18, no. 16, 2010.

[9] F. Pfeiffer, T. Weitkamp, O. Bunk, and C. David, "Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources," *Nat. Phys.*, vol. 2, no. 4, 2006.

[10] T. Donath *et al.*, "Inverse geometry for grating-based x-ray phase-contrast imaging," *Journal of Applied Physics*, vol. 106, no. 5, 2009.

[11] A. A. Fingerle *et al.*, "Imaging features in post-mortem x-ray dark-field chest radiographs and correlation with conventional x-ray and CT," *European radiology experimental*, vol. 3, 2019.

[12] J. Andrejewski *et al.*, "Whole-body x-ray dark-field radiography of a human cadaver," *European Radiology Experimental*, vol. 5, no. 1, 2021.

[13] K. Willer *et al.*, "X-ray dark-field chest imaging for detection and quantification of emphysema in patients with chronic obstructive pulmonary disease: A diagnostic accuracy study," *The Lancet Digital Health*, vol. 3, no. 11, 2021.

[14] T. Urban *et al.*, "Qualitative and quantitative assessment of emphysema using dark-field chest radiography," *Radiology*, 2022.

[15] F. T. Gassert *et al.*, "Assessment of inflation in a human cadaveric lung with dark-field chest radiography," *Radiology: Cardiothoracic Imaging*, vol. 4, no. 6, 2022.

[16] M. Frank *et al.*, "Dark-field chest X-ray imaging for the assessment of covid-19-pneumonia," *Communications Medicine*, vol. 2, no. 1, 2022.

[17] T. Urban *et al.*, "Dark-field chest radiography outperforms conventional chest radiography for the diagnosis and staging of pulmonary emphysema," *Investigative Radiology*, vol. Publish Ahead of Print, 2023.

[18] M. Kattau *et al.*, "X-ray dark-field chest radiography: A reader study to evaluate the diagnostic quality of attenuation chest X-rays from a dual-contrast scanning prototype," *European Radiology*, 2023.

[19] M. Bech *et al.*, "Quantitative x-ray dark-field computed tomography," *Physics in Medicine & Biology*, vol. 55, no. 18, 2010.

[20] M. Viermetz *et al.*, "Dark-field computed tomography reaches the human scale," *Proceedings of the National Academy of Sciences*, vol. 119, no. 8, 2022.

[21] M. Viermetz *et al.*, "Technical design considerations of a human-scale Talbot-Lau interferometer for dark-field CT," *IEEE Transactions on Medical Imaging*, 2022.

[22] M. Viermetz *et al.*, "Initial characterization of Dark-Field CT on a clinical gantry," *IEEE Transactions on Medical Imaging*, 2022.

[23] T. Buzug, *Computed Tomography*. Springer, Berlin, 2008.

[24] A. C. Kak and M. Slaney, *Principles of computerized tomographic imaging*. SIAM, 2001.

[25] H. Turbell, "Cone-beam reconstruction using filtered backprojection," Dissertation, Linköpings Universitet, 2001.

[26] H. Talbot, "Facts relating to optical science. no. IV," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 9, no. 56, 1836.

[27] E. Lau, "Beugungserscheinungen an Doppelrastern," *Annalen der Physik*, vol. 437, no. 7-8, 1948.

[28] M. Chabior, "Contributions to the characterization of grating-based x-ray phase-contrast imaging," Dissertation, Technische Universität Dresden, 2011.

[29] K. Okada, A. Sato, and J. Tsujiuchi, “Simultaneous calculation of phase distribution and scanning phase shift in phase shifting interferometry,” *Opt. Commun.*, vol. 84, no. 3-4, 1991.

[30] M. Seifert *et al.*, “Optimisation of image reconstruction for phase-contrast X-ray Talbot-Lau imaging with regard to mechanical robustness,” *Phys. Med. Biol.*, vol. 61, no. 17, 2016.

[31] M. Marschner *et al.*, “Helical X-ray phase-contrast computed tomography without phase stepping,” *Sci. Rep.*, vol. 6, 23953, 2016.

[32] S. Kaeppler *et al.*, “Improved reconstruction of phase-stepping data for Talbot-Lau X-ray imaging,” *J. Med. Imaging*, vol. 4, no. 3, 034005, 2017.

[33] C. Hauke *et al.*, “Enhanced reconstruction algorithm for moiré artifact suppression in Talbot-Lau X-ray imaging,” *Phys. Med. Biol.*, vol. 63, no. 13, 135018, 2018.

[34] F. De Marco *et al.*, “Analysis and correction of bias induced by phase stepping jitter in grating-based X-ray phase-contrast imaging,” *Opt. Express*, vol. 26, no. 10, 2018.

[35] J. Dittmann, A. Balles, and S. Zabler, “Optimization based evaluation of grating interferometric phase stepping series and analysis of mechanical setup instabilities,” *J. Imaging*, vol. 4, no. 6, 77, 2018.

[36] K. Hashimoto, H. Takano, and A. Momose, “Improved reconstruction method for phase stepping data with stepping errors and dose fluctuations,” *Opt. Express*, vol. 28, no. 11, 2020.

[37] W. Noichl *et al.*, “Correction for mechanical inaccuracies in a scanning Talbot-Lau interferometer,” *IEEE Transactions on Medical Imaging*, 2023.

[38] M. Seifert *et al.*, “Talbot-Lau x-ray phase-contrast setup for fast scanning of large samples,” *Scientific Reports*, vol. 9, no. 1, 2019.

[39] A. Tapfer *et al.*, “Experimental results from a preclinical X-ray phase-contrast CT scanner,” *Proceedings of the National Academy of Sciences*, vol. 109, no. 39, 2012.

[40] A. Yaroshenko *et al.*, “Pulmonary emphysema diagnosis with a preclinical small-animal x-ray dark-field scatter-contrast scanner,” *Radiology*, vol. 269, no. 2, 2013.

[41] Z. Wu *et al.*, “Prototype system of noninterferometric phase-contrast computed tomography utilizing medical imaging components,” *Journal of Applied Physics*, vol. 129, no. 7, 2021.

[42] M. Viermetz, “Development of the first Human-scale Dark-field Computed Tomography System,” Dissertation, Technical University of Munich, 2022.

[43] M. Kachelrieß, M. Knaup, C. Penßel, and W. A. Kalender, “Flying focal spot (ffs) in cone-beam CT,” *IEEE transactions on nuclear science*, vol. 53, no. 3, 2006.

[44] M. v. Teuffenbach *et al.*, “Grating-based phase-contrast and dark-field computed tomography: A single-shot method,” *Scientific reports*, vol. 7, no. 1, 2017.

[45] J. Mohr *et al.*, “High aspect ratio gratings for X-ray phase contrast imaging,” in *AIP Conference proceedings*, American Institute of Physics, vol. 1466, 2012.

[46] T. J. Schröter *et al.*, “Large field-of-view tiled grating structures for X-ray phase-contrast imaging,” *Review of Scientific Instruments*, vol. 88, no. 1, 2017.

[47] C. Schmid *et al.*, “Modeling vibrations of a tiled Talbot-Lau interferometer on a clinical CT,” *IEEE Transactions on Medical Imaging*, 2022.

[48] F. Horn *et al.*, “Implementation of a Talbot-Lau interferometer in a clinical-like c-arm setup: A feasibility study,” *Scientific Reports*, vol. 8, no. 1, 2018.

[49] T. Weber *et al.*, “Noise in x-ray grating-based phase-contrast imaging,” *Med. Phys.*, vol. 38, no. 7, 2011.

[50] J. Bradbury *et al.*, *JAX: Composable transformations of Python+NumPy programs*, version 0.4.2, 2018.

[51] D. C. Liu and J. Nocedal, “On the limited memory bfgs method for large scale optimization,” *Mathematical programming*, vol. 45, no. 1, 1989.

[52] S. G. Krantz and H. R. Parks, *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media, 2002.

[53] G. Pelzer *et al.*, “Reconstruction method for grating-based x-ray phase-contrast images without knowledge of the grating positions,” *Journal of Instrumentation*, vol. 10, no. 12, 2015.

[54] E. Cheney and D. Kincaid, *Linear Algebra: Theory and Applications*. Jones and Bartlett Publishers, 2009.

[55] M. Bech, “X-ray imaging with a grating interferometer,” Dissertation, University of Copenhagen, 2009.

[56] N. T. Vo, R. C. Atwood, and M. Drakopoulos, “Superior techniques for eliminating ring artifacts in X-ray micro-tomography,” *Opt. Express*, vol. 26, no. 22, 2018.

[57] G. Lewith, W. B. Jonas, and H. Walach, *Clinical Research in Complementary Therapies*. Elsevier, 2002.

[58] N. Gustschin *et al.*, “Quality and parameter control of X-ray absorption gratings by angular X-ray transmission,” *Optics express*, vol. 27, no. 11, 2019.

[59] J. Haeusele *et al.*, “Advanced phase-retrieval for stepping-free X-ray dark-field computed tomography,” *IEEE Transactions on Medical Imaging*, 2023.

[60] K. Vegso *et al.*, "Recent advance in grating-based x-ray phase tomography," in *Developments in X-Ray Tomography XI*, B. Müller and G. Wang, Eds., SPIE, 2017.

[61] G. Zan *et al.*, "High-resolution multicontrast tomography with an X-ray microarray anode–structured target source," *Proceedings of the National Academy of Sciences*, vol. 118, no. 25, 2021.

[62] I. Zanette *et al.*, "Trimodal low-dose X-ray tomography," *Proceedings of the National Academy of Sciences*, vol. 109, no. 26, 2012.

[63] M. Willner *et al.*, "Quantitative X-ray phase-contrast computed tomography at 82 kev," *Optics express*, vol. 21, no. 4, 2013.

[64] S. Van Gogh *et al.*, "Towards clinical-dose grating interferometry breast CT with fused intensity-based iterative reconstruction," *Optics Express*, vol. 31, no. 5, 2023.

[65] S. Gu, L. Zhang, W. Zuo, and X. Feng, "Weighted nuclear norm minimization with application to image denoising," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014.

[66] M. J. Ehrhardt and S. R. Arridge, "Vector-valued image processing by parallel level sets," *IEEE Transactions on Image Processing*, vol. 23, no. 1, 2013.

[67] T. Koehler and E. Roessl, "Simultaneous denoising in phase contrast tomography," in *AIP Conference Proceedings*, American Institute of Physics, vol. 1466, 2012.

[68] S. Allner *et al.*, "Bilateral filtering using the full noise covariance matrix applied to x-ray phase-contrast computed tomography," *Physics in Medicine & Biology*, vol. 61, no. 10, 2016.

[69] M. Frank *et al.*, "Dosimetry on first clinical dark-field chest radiography," *Medical physics*, vol. 48, no. 10, 2021.

[70] N. Gustschin *et al.*, "Dark-field imaging on a clinical CT system: Performance and potential based on first results," in *7th International Conference on Image Formation in X-Ray Computed Tomography*, SPIE, vol. 12304, 2022.

[71] F. De Marco *et al.*, "X-ray dark-field signal reduction due to hardening of the visibility spectrum," *arXiv preprint arXiv:2011.03542*, 2020.

[72] C. Bontus and T. Kohler, "Reconstruction algorithms for computed tomography," *Advances in Imaging and Electron Physics*, vol. 151, 2008.

[73] T. Kohler *et al.*, "Evaluation of helical cone-beam CT reconstruction algorithms," in *2002 IEEE Nuclear Science Symposium Conference Record*, IEEE, vol. 2, 2002.

[74] A. Pandeshwar, M. Kagias, Z. Shi, and M. Stampanoni, "Envelope modulated x-ray grating interferometry," *Applied Physics Letters*, vol. 120, no. 19, 2022.

[75] R. Tang *et al.*, "Detailed analysis of the interference patterns measured in lab-based X-ray dual-phase grating interferometry through wave propagation simulation," *Optics Express*, vol. 31, no. 2, 2023.

[76] M. Patino *et al.*, "Material separation using dual-energy CT: Current and emerging applications," *Radiographics*, vol. 36, no. 4, 2016.

[77] R. E. Alvarez, "Energy-selective reconstructions in x-ray computerised tomography," *Physics in Medicine and Biology*, vol. 21, no. 5, 1976.

[78] K. Mechlem *et al.*, "Spectral differential phase contrast x-ray radiography," *IEEE transactions on medical imaging*, vol. 39, no. 3, 2019.

[79] K. Mechlem *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*, vol. 65, no. 6, 2020.

List of Figures

List of Tables

List of abbreviations

1D one-dimensional

2D two-dimensional

3D three-dimensional

aka. also known as

autodiff automatic differentiation

CNR contrast-to-noise ratio

CT computed tomography

DFCT dark-field computed tomography system

DOF degrees-of-freedom

FBP filtered back-projection

FOV field of view

FP tomographic foward-projection

GPU graphics processing unit

IBSIR intensity-based statistical iterative reconstruction

MRI magnetic resonance imaging

PCA principal component analysis

POM polyoxymethylene

SIMD “single instruction, multiple data”

SIR statistical iterative reconstruction

SPR (linearized) statistical phase retrieval

SVD singular value decomposition

wrt. with respect to

Appendix A

Auxiliary data

A.1 Higher-order reference visibility and phase flat-fields

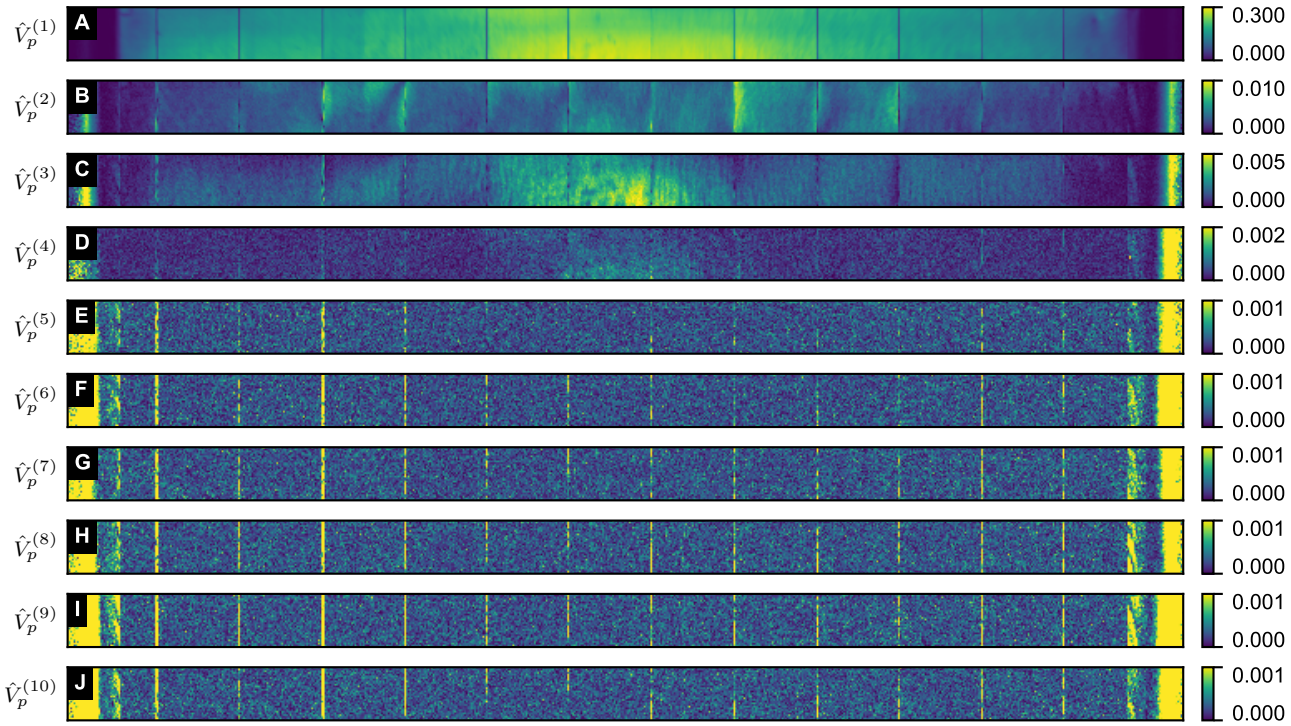

Figure A.1: DFCT reference visibility $\hat{\boldsymbol{V}}^{(n)}$ for $n = 1...10$. It is determined by performing the presented processing pipeline in chapter 3 for an axial air-scan and using the generalized linearized phase retrieval for higher-order Fourier terms. Please note the different value windows in the colorbars. **A** The first-order visibility $\hat{\boldsymbol{V}}^{(1)}$ is by far the largest in magnitude, reaching 30 % in the detector center. **B**,**C** For $n = 2$ and $n = 3$ the magnitude is below 1 %. These are the terms we use in practice. **D** $\hat{\boldsymbol{V}}^{(4)}$ is still discernible over the noise floor but below 0.1 % in magnitude. **E**-**J** The higher orders $n = 5...10$ are not discernible and only show noise despite the high number of projections from a full axial scan used for the signal retrieval. Their minuteness is a combination of the higher-order Fourier series terms inherently decreasing with order n, and the effective visibility generally decreasing with smaller fringe size.

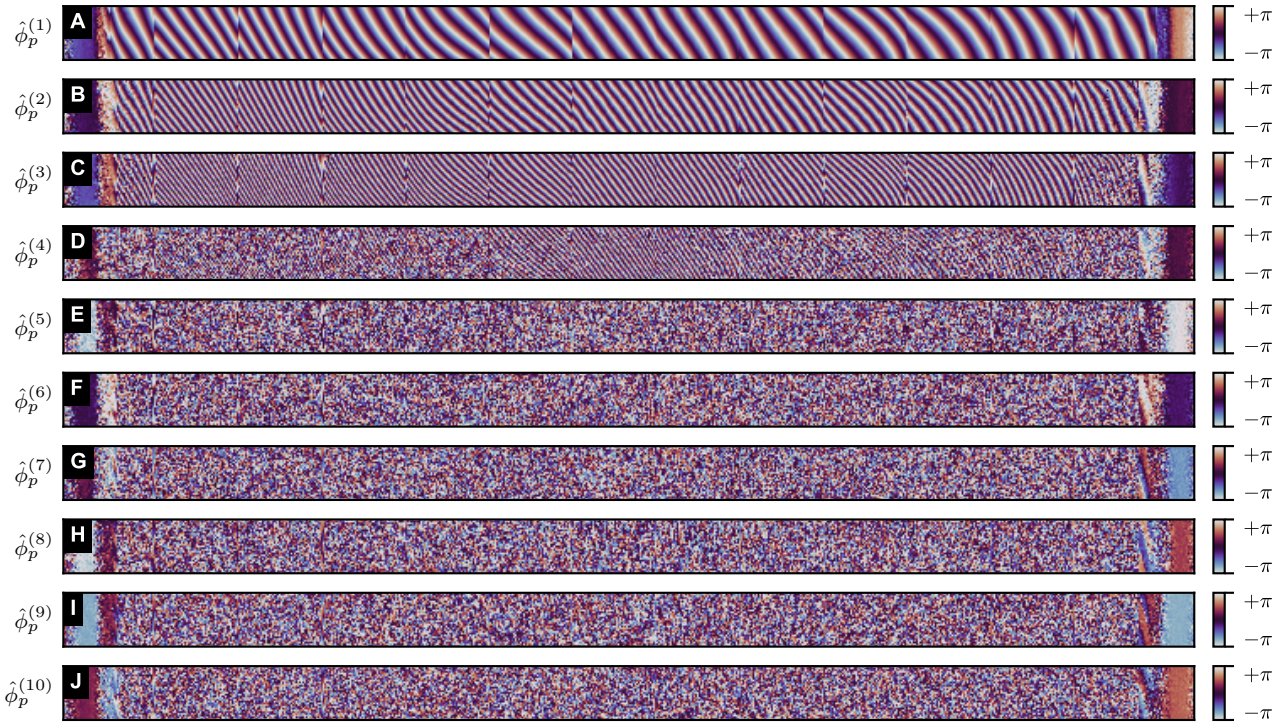

Figure A.2: DFCT reference phase $\hat{\boldsymbol{\phi}}^{(n)}$ for $n = 1...10$. It is determined by performing the presented processing pipeline in chapter 3 for an axial air-scan and using the generalized linearized phase retrieval for higher-order Fourier terms. **A** The first-order phase $\hat{\boldsymbol{\phi}}^{(1)}$ visually matches the fringe pattern of the raw intensity measurements because the first-order visibility $\hat{\boldsymbol{V}}^{(1)}$ dominates the overall visibility contrast. **B**,**C** For $n = 2$ and $n = 3$ the fringe pattern has double and triple the spatial frequency, respectively. These are the terms we use in practice. The doubled and tripled lateral frequency is expected from the physical model but not enforced in the signal retrieval. Its occurrence therefore is evidence of the physical presence of the higher orders in the data. **D** $\hat{\boldsymbol{\phi}}^{(4)}$ is only discernible in the center where $\hat{\boldsymbol{V}}^{(4)}$ is larger than the noise floor. **E**-**J** The higher orders $n = 5...10$ are not discernible and only show noise despite the high number of projections from a full axial scan used for the signal retrieval. Their minuteness is a combination of the higher-order Fourier series terms inherently decreasing with order n, and the effective visibility generally decreasing with smaller fringe size.

Scientific contributions

Publications with first authorship

- **Clemens Schmid**, Manuel Viermetz, Nikolai Gustschin, Wolfgang Noichl, Jakob Haeusele, Tobias Lasser, Thomas Koehler, and Franz Pfeiffer. "*Modeling vibrations of a tiled Talbot-Lau interferometer on a clinical CT system*". IEEE Transactions on Medical Imaging, 2022.

Publications with co-authorship

- Manuel Viermetz*, Nikolai Gustschin*, **Clemens Schmid**, Jakob Haeusele, Maximilian von Teuffenbach, Pascal Meyer, Frank Bergner, Tobias Lasser, Roland Proksa, Thomas Koehler, and Franz Pfeiffer. "*Dark-field computed tomography reaches the human scale*". Proceedings of the National Academy of Sciences, 2022.
- Manuel Viermetz, Nikolai Gustschin, **Clemens Schmid**, Jakob Haeusele, Peter Noël, Roland Proksa, Stefan Loescher, Thomas Koehler, and Franz Pfeiffer. "*Technical design considerations of a human-scale Talbot-Lau interferometer for dark-field CT*". IEEE Transactions on Medical Imaging, 2022.
- Manuel Viermetz, Nikolai Gustschin, **Clemens Schmid**, Jakob Haeusele, Bernhard Gleich, Bernhard Renger, Thomas Koehler, and Franz Pfeiffer. "*Initial Characterization of Dark-Field CT on a Clinical Gantry*". IEEE Transactions on Medical Imaging, 2022.
- Jakob Haeusele, **Clemens Schmid**, Manuel Viermetz, Nikolai Gustschin, Tobias Lasser, Thomas Koehler, and Franz Pfeiffer. "*Advanced Sliding Window Signal Extraction for Rapidly Acquired Dark-field CT*". IEEE Transactions on Medical Imaging, 2023.

* Shared first authorship.

Presentations

- **Clemens Schmid**, Manuel Viermetz, Nikolai Gustschin, Jakob Haeusele, Tobias Lasser, Thomas Koehler, and Franz Pfeiffer. "*Dark-Field Imaging on a Clinical CT System: Modeling of Interferometer Vibrations*". 7th International Conference on Image Formation in X-Ray Computed Tomography, 2022.

Acknowledgments

This thesis stands on the proverbial “shoulders of giants”. It would not have been possible without the help and previous work of many people whom I want to thank here.

- **Franz Pfeiffer** for giving me the chance to do research on this topic and using the group’s facilities, for your insights into physics and medical application, for supporting my research stay in Boston, and for allowing me to be a part of the international science community through conferences and collaborations.
- **Thomas Koehler** for your expertise in X-ray imaging and physics, for supporting this project with industry experience, and for your cool and collected manner of doing and discussing research.
- **Tobias Lasser** for being the mentor of my PhD research, and for your support concerning both algorithms as well as good scientific practice and communication.
- **Manuel, Niko, and Jakob** for being my close colleagues and good friends in this research project, for the many discussions, explanations, questions, and shared experiences both in fun and frustration.
- **Wolfgang** for being my go-to senior colleague about all things phase retrieval, computational physics, and good scientific practice.
- **Max, Amanda, Catalin, Robert, Nikola, Tom, and Florian** for your contributions to this research project and, by extension, this thesis.
- **Martin, Wolfgang, and Johannes M.** for your continuous support of the hard- and software infrastructure of the chair, for sharing your GNU/Linux wisdom, and for entrusting me with admin responsibilities.
- **Julia Herzen** for our discussions and your insights, both on- and off-topic.
- **All past and present members of the chair** for the collaborative atmosphere and good times in general.
- **My family and friends** for supporting me throughout these years, and for reminding me that there indeed are more things to life than science.