Spectral Phase Contrast Imaging and Phase Retrieval

by
Ivan Vazquez

A dissertation submitted to the Department of Physics,

College of Natural Science and Mathematics

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in Physics

Chair of Committee: Mini Das

Committee Member: Howard C. Gifford

Committee Member: Liming Li

Committee Member: Mark A. Meier

University of Houston
December 2020

DEDICATION/EPIGRAPH

To Mima and to my loving parents, Ivan and Dagmaris.

ACKNOWLEDGMENTS

This work would have been impossible without the help of Maria Pesek and my family. Maria helped me stay motivated through the many challenges I encountered during the four years I worked at the University of Houston. I am especially thankful for her help with meals and for advising me on many personal and professional decisions. My family's support was also fundamental. They encouraged me and helped me tackle many obstacles. I want to commend my father and Lilia for their efforts in creating a healthy and quiet environment to work.

I want to thank Amareswararao (Amar) Kavuri and Nathaniel (Nate) R. Fredette for our many enriching discussions in the lab. They are both excellent researchers and great friends. I also want to thank the other members of the lab for helping with data collection and with brainstorming new ideas.

ABSTRACT

Phase-contrast X-ray imaging (PCI) can significantly improve the contrast of small and weakly-absorbing materials. This is because of the sensitivity of PCI measurements to changes in both the attenuation and phase of an X-ray wavefield due to interactions with matter. Yet, quantitative estimates of material properties from PCI measurements require the use of phase retrieval (PR) algorithms. Our team developed a capable PR strategy for propagation-based PCI that estimates a pair of material properties from two or more energy-resolved measurements. Part of this work is dedicated to broadening our understanding of the key advantages and practical limitations of our method. To this end, a comprehensive analysis of the underlying theory and the sensitivity to factors such as noise are provided. Based on our findings, we proposed a set of strategies that can help enhance the quality and accuracy of estimated values. We also examined the performance of the PR method with laboratory measurements of a heterogeneous sample. To simultaneously obtain multiple energy-resolved measurements, we used a polychromatic source in combination with high-resolution photon-counting detectors. Our findings demonstrate that the properties of multiple (more than two) unique materials can be estimated with accuracies above 90%. Additionally, our strategy produced results with satisfactory accuracy and image quality even when the detected photon counts were as low as 250 photons per pixel. The underpinning theory in our approach relies on the weak object approximation to simplify the transport-of-intensity equation (TIE). Thus, a section of this work is dedicated to examining the errors related to enforcing the weak object approximation when deriving a TIE-based PR algorithm. We also tested other prevalent methods to simplify the TIE. Experimental results revealed that errors introduced by the weak object approximation were below 5% even for the case of aluminum, which has a relatively high atomic number.

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGMENTS iii

ABSTRACT iv

LIST OF FIGURES x

LIST OF TABLES xi

NOMENCLATURE xii

1 Introduction 1
- 1.1 Fundamentals of medical X-ray imaging . . . 2
  - 1.1.1 X-ray interaction mechanisms . . . 3
- 1.2 Phase contrast imaging . . . 5
- 1.3 Spectral X-ray imaging . . . 7
  - 1.3.1 Polychromatic radiation and its consequences . . . 7
  - 1.3.2 Photon-counting detectors . . . 9
- 1.4 Outline . . . 10

2 Propagation-based phase contrast imaging 11
- 2.1 Experimental geometry and practical requirements . . . 11
- 2.2 Physics of propagation-based imaging . . . 14
  - 2.2.1 Weakly interacting object approximation . . . 16
  - 2.2.2 Phase and amplitude contrast transfer functions . . . 17
  - 2.2.3 The transport of intensity equation . . . 19
  - 2.2.4 Validity of the near-field and projection approximations . . . 20

3 Spectral phase retrieval assuming weak X-ray interactions 23
- 3.1 Theory and derivation . . . 24
- 3.2 On the impact of data noise in spectral phase retrieval . . . 28
- 3.3 Analytic frequency modulation functions of the retrieved quantities . . . 30
- 3.4 Noise mitigation strategies . . . 35
  - 3.4.1 Apodization: frequency-domain filtering . . . 38
  - 3.4.2 Spatial denoising: block-matching and three-dimensional filtering . . . 40
  - 3.4.3 Smoothing by feature scaling . . . 43
  - 3.4.4 Discussion . . . 46
- 3.5 Errors from misrepresenting the energy . . . 46

4 Phase Retrieval of Soft Materials at Low Photon Counts 50
- 4.1 Materials and methods . . . 50
  - 4.1.1 Multi-Material Phantom and Biological Sample . . . 50
  - 4.1.2 Source and Detectors . . . 52
  - 4.1.3 Spectroscopic Measurements . . . 52

4.1.4 Quality and Accuracy Metrics . . . 54
4.1.5 Phase-retrieval workflow . . . 55
4.2 Results . . . 55
4.2.1 Phase retrieval of multi-material phantom . . . 55
4.2.2 Phase Retrieval with Low Photon Counts . . . 57
4.2.3 Quality Improvements on a Bio-sample Image . . . 59
4.3 Discussion . . . 60

5 Assessments of the Error Introduced by Approximations to the Transport of Intensity Equation 62
5.1 Analytical form of the weakly interacting object approximation . . . 63
5.1.1 Evaluating the approximation with an analytical phantom . . . 64
5.2 Influence of the weak-object approximation on phase retrieval . . . 66
5.3 Overview of the homogeneous and phase-attenuation duality approximations . . . 69
5.4 Derivation of a phase retrieval to assess the weak-object approximation . . . 70
5.5 Phase retrieval with experimental data . . . 71
5.5.1 Correcting spectral distortions . . . 73
5.5.2 Accurate estimates of the effective energy . . . 74
5.5.3 Analyzing the influence of the weak-object approximation . . . 76
5.5.4 Analysis of the phase-attenuation duality and homogeneous approximations . . . 77
5.6 Summary and conclusions . . . 80

6 CONCLUDING REMARKS AND FUTURE WORK 83

A Two-material Phase Retrieval: Derivation for Point-Sources 85
A.1 Homogeneous phase retrieval equation . . . 87

B Hybrid Phase Retrieval 88
B.1 Phase-attenuation duality algorithm . . . 90

C Simulation Framework for PCI Measurements 92
C.1 Simulating noisy data . . . 93

D Spectral correction 95

E Noise Power Spectrum 97

F Strategy to Estimate the Effective Energy 98

BIBLIOGRAPHY 100

LIST OF FIGURES

1 (a) Mass attenuation curves for water. (b) Relative contribution of Rayleigh scattering (σ_r), Compton scattering (σ_c), and photoelectric absorption (τ) to the total linear attenuation coefficient (μ). 4

2 Real (δ) and imaginary (β) parts of the complex refractive index of aluminum and water. 6

3 Two enlarged pixels of a PCD illustrating the basic steps that enable spectroscopic imaging. Photons deposit a charge proportional to their energy. The charge is converted to a voltage that is later amplified, compared to a user-defined threshold (TH), and counted if it exceeds the threshold. 9

4 (a) Depiction of a propagation-based phase contrast experiment geometry. The X-ray source and the sample are separated by a distance R_1 while the distance from the exit-plane of the sample to the detector is R_2. In conventional X-ray imaging, $R_2 \approx 0$. (b) Sketch showing the entrance and exit planes. The former is defined as the plane immediately upstream of the object while the latter is the plane immediately downstream of the object. In PCI, X-rays propagate from the exit plane towards the detector plane. 12

5 (a) Radially symmetric two-dimensional contrast transfer functions for the phase and amplitude. (b) One dimensional profiles along $\nu_y = 0$. The direct imaging regime $\sqrt{\lambda z}\nu_x < 1$ is highlighted. This region includes the optimum-phase contrast where the phase CTF is at a maximum. It also includes the near-field regime, where $\sin\chi \approx \chi$. (c) Shows how the smoothing effect of the point-spread-function of a system can impact the phase and amplitude CTF, labeled as pCTF and aCTF. . . 18

6 (a) Plot of the Fresnel number, as defined in Eq. 29, for compressed breast thicknesses ranging up to 8 cm and microcalcifications below 10 microns. A white line shows the $N_F = 10$ datum, which marks the limit of validity for the projection approximation. (b) Plot showing the Fresnel number as defined in Eq. 30 and a white line marking the near-field regime. An energy of 20 keV was used for both plots. 21

7 Retrieved and tabulated properties of polymethyl methacrylate (PMMA). Phase retrieval was performed on experimental measurements of a 3-mm PMMA rod obtained with a photon-counting detector and a polychromatic X-ray source. 27

8 Cascaded framework used to test the sensitivity of the SPR algorithm to noise and other relevant sources of error in the data and the phase retrieval workflow. Here F represents the forward modeling process that generates measurements at defined energies and N represents the introduction of noise into the simulated measurements. The operator F^{-1} represents the phase retrieval step. 28

9 Noise power spectrum of the simulated (noisy) measurements and two attenuation properties related to the retrieved quantities (a_1 and a_2). 29

10 Depiction of the hypothetical experiment used to learn about the frequency response of the retrieved quantities, a_1 and a_2. A sample is measured twice at different X-ray energies ensuring that the radiation is collected by a single pixel. This results in two sparse measurements, d_{E_i} with constant Fourier representations. 31

11 Normalized amplitude response of the retrieved terms, i.e. $|\tilde{A}_1|/\max|\tilde{A}_1|$ and $|\tilde{A}_2|/\max|\tilde{A}_2|$, for different energy pairs. The energy pairs used to analyze Eq. 43 and 42 are listed on the perimeter. 33
12 Real part (left) and amplitude response (right) of the FT of the photoelectric absorption term $\tilde{A}_1$ for the 15-50 keV energy pair. 34
13 (a) Amplitude response of the $\tilde{A}_1$ term using the energies 15-and-20 keV with two-times (2X) and three-times (3X) geometrical magnification. (b) Amplitude response of $\tilde{A}_1$ for the energies of 15-and-50 keV and two different materials, PMMA and PTFE. 34
14 (a) Workflow showing the approach used to investigate the three noise mitigation strategies. Here W represents an apodization strategy using a Tukey windowing kernel, D represents a spatial denoising step preceding phase retrieval, and scaling the known terms (feature scaling) is represented by the phase retrieval operator F^{-1} itself. The last choice illustrates that feature scaling is only a modification of how F^{-1} is implemented and does not amount to an added step. (b) Shows the regions used to estimate the percent error and SNR of the resulting images. The red, blue, and green boxes enclose the water, POM, and PMMA rods, respectively. The dashed white box shows the region used to estimate the standard deviation needed to compute the SNR. 36
15 Shows the effect of increasing the amount of high frequency content that is removed by the Tukey windowing kernel is shown. As more high frequency content is removed (going from left to right), the visual appearance of the image seems to improve. This, however, comes at the price of image resolution loss. 39
16 Outcome of applying the BM3D filtering scheme, represented by D, to a noisy PCI measurement. Comparing the average profiles below the noisy and filtered images shows how this routine reduces noise without substantial damage to sharp edges. . . 41
17 Results of the PE and CS terms when the inputs to PR are filtered by the BM3D scheme using four different values for the standard deviation of noise. 42
18 Increasing the scaling term introduces a sharp radial cut-off in the frequency response of the estimated terms. When the scale is set to unit y, the characteristic behavior of the two terms described in Sec. 3.3.The electron density rolls smoothly to zero, while the PE term increases with frequency. A line profile for the scale of 100 shows the sharp transition that results from feature scaling, which can lead to ringing artifacts and deterioration of the measurement quality. 44
19 The effect of different scaling terms for simulated data. The last term is smooth and appears free of artifacts, but is highly inaccurate. 45
20 (a)Results of phase retrieval using two (noise-free) monochromatic input measurements simulated with energies of 23.33 and 40 keV. The low-energy measurement (23.33 keV) was assigned an incorrect energy value during phase retrieval. The energy values used for the low energy bin are listed below the resulting PE and CS, which were estimated for an energy of 30 keV, c.f. Eq. 34. Each plot also shows the expected PE and CS and ignores negative values for clarity. (b) Average percent error along the rod profile. 47
21 (a) Two monochromatic measurements were simulated with energies of 30 and 60 keV. To induce an error, the 60 keV bin was mislabeled during PR as 62, 64 and 66 keV. (b) Average percent error along the rod profile. 48

22 Experimental setup used to image the multi-material phantom. The process of energy binning is also depicted. The Medipix3RX detector captures specific sections of the incident spectrum, and a set of energy bins corresponding to narrow sections of the incident spectrum are formed using energy binning. 51
23 Three steps involved in our spectral phase retrieval workflow. First, bad pixel correction (BPC) is implemented to remove dead-pixels and outliers from the measurements. Next, energy binning (EB) is implemented to create pseudo-monochromatic measurements. Finally, the phase retrieval (PR) is applied to the energy bins, and the material properties a_1 and a_2 are obtained (Sec. 3.1). Energy-dependent properties, like CS and PE, using a_1 and a_2 can be estimated. 55
24 Phase retrieval results for the multi-material phantom. (a) PCI measurements of multi-material phantom showing the expected edge enhancement. (b) pLAC map with the average retrieved profile (green) and expected (red) values superimposed. (c) The retrieved phase difference map with the retrieved (green) and expected (red) average profiles. The retrieved quantities correspond to an effective energy of 40 keV. 56
25 (a) Average profile of the retrieved and expected photoelectric term for the four material phantom. (b) Normalized contrast for the three retrieved mechanisms. The contrast estimated for each material was normalized by the maximum value for each of the attenuation mechanisms. 57
26 In all plots, 2X indicates the results for the measurements taken at a two-times magnification and the 3X the configuration with a magnification of three. (a) Signal difference-to-noise ratio (SNR) of the retrieved projected linear attenuation coefficient (pLAC), retrieved phase, and negative log of the measured phase contrast images (PCIs). An effective energy of 38 keV was estimated for the PCI measurement, and the retrieved quantities were computed for this energy (b) Gain in SNR found as the ratio between the SNR of the pLAC and PCI images as well as the phase and PCI images. (c) Average percent error computed from the estimated and expected electron density of POM. 58
27 (a) The negative log of the normalized intensity, $-\ln[I/I^{\text{in}}]$, from measurements taken at the contact plane, that is $M = 1$. (b) The same quantity as (a) but using measurements taken in a phase-enhancing configuration, that is $M = 3$. (c) Shows the same quantity as (a) and (b) but with intensity values reconstructed from the retrieved material properties. The insets show the same ROI in all three images, highlighting the quality and detail in each. 59
28 (a) Two-material phantom composed of two homogeneous objects with basic shapes. The thickness A is constant and the field-of-view (FOV) is limited to a small region that encapsulates the ends of the rod but is far from the ends of the rectangular slab. (b) Geometrical configuration of the PCI measurements used for this section. 65
29 (a) Plot of the left side of the inequality described by Eq. 60 using three radii and $\Delta\mu_E = 1$. (b) Plot of the maximum value of the left side of Eq. 60 for different combinations of $\Delta\mu$ and r. 66
30 Plot of the percent error between the intensity predicted by Eq. 61 with and without enforcing the weak-interaction approximation. The intensities were estimated for PMMA rods of different thicknesses. The horizontal axis represents the ratio between the distance away from the center of each rod and the corresponding radius. 68

31 Percent error between the intensities predicted Eq. 61 with and without enforcing the WOA in Eq. 63. One horizontal axis represent different contrast values ($\Delta\mu$) between the target and background materials while the other shows the ratio between the distance from the center of the target to the radius of the target. The ratio was kept to 98 % of the way to the edge of the target. 69
35 Average percent error in the projected thickness of the Al rod after using the two-material PR equation (Eq. 81 in Appx. A) on our data with water in the background. 79
36 Hybrid phase retrieval workflow. In the first step a_2 is estimated. Next a numerical Laplacian operation is used to obtain approximate $\boldsymbol{\nabla}^2 a_2$. Finally, the a_1 term is estimated from one or multiple bins. A series of optional denoising steps as well as a final smoothing of a_1 can be applied to improve the result. 88
37 Projected linear attenuation coefficients of graphite (left) and POM (right) rod estimated from outputs of the HyPR workflow. The first row shows the result of not using any denoising or smoothing. The second row shows the result of applying denoising and smoothing based on the standard deviation of the noise in the input data at each filtering stage. The images in the last row were found using user-defined values for the standard deviations. 90
38 Three flat-field corrected monochromatic measurements simulated at different energies and displaying unique amounts of noise. The noise is Poisson distributed. 94
39 Workflow used to correct spectral distortions. 95
40 (a) Original and corrected profile of a PCI measurement of an aluminum rod. The measurement corresponds to an effective energy of 24 keV. A dashed line shows the expected response for a bin bounded at the same energy values. (b) Original and corrected profiles of a PMMA bin having an effective energy of 26 keV. 96
41 Workflow used to estimate the noise power spectrum. 97
43 Estimating the effective energy iteratively. The error at each iteration is indicates with ε. 99

LIST OF TABLES

1 Average percent error and SNR for the quantitative assessment of the apodization scheme using different choices of cut-off frequency, ν_c. 40

2 Average percent error and SNR for the quantitative assessment of spatial denoising the inputs to the phase retrieval algorithm using a BM3D filter. The first column lists the different standard deviation values used, which the filtering algorithm interprets as the standard deviation of the noise in the input images. 43

3 Average percent error and SNR to assess the effect of feature scaling. For these results, four values of γ during phase retrieval were used (Eq. 48). The scale of the energy-functions was determine by the choice of units, namely keV. 46

4 The properties listed are effective atomic number (Z_{eff}), electron density (ρ_e), mass density (ρ_m), total linear attenuation (μ), photoelectric absorption (τ_{pe}), Compton scattering (σ_{cs}), and rod diameter ($\varnothing$). 51

5 Sets of measurements used to study the error in phase retrieval due to approximations in the TIE. The sets were defined by grouping energy bins with similar mean counts per pixel in the flat-field measurements and, therefore, similar SNR. The same bin boundaries were used for air and water as background materials, which resulted in different SNR values since the acquisition time was 60 s in both cases. The contrast in LAC ($\Delta\mu$) between Al and the background materials is also listed for every E_{eff}. For air as background, the change in LAC is equal to the LAC of Al. 72

6 Constants for the PE model found by minimizing Eq. 90 along with tabulated material properties. All properties marked with an asterisk correspond to tabulated properties. 93

NOMENCLATURE

TIE	Transport of intensity equation
PR	Phase retrieval
CT	Computed tomography
PCI	Phase-contrast imaging
PB-PCI	Propagation-based phase contrast imaging
SNR	Signal-to-noise ratio
CNVR	Contrast-to-noise-variance ratio
MMD	Multi-material decomposition
SPR	Spectral phase retrieval
FT	Fourier transform
CS	Compton Scattering
PE	Photoelectric absorption
LAC	Linear attenuation coefficient
pLAC	Projected linear attenuation coefficient
FOV	Field-of-view
ROI	Region of interest
Z	Effective atomic number
PCD	Photon-counting detector
CTF	Contrast transfer function

1 Introduction

From detecting the onset of cancer in a patient to identifying an explosive device in a package, X-ray imaging is used routinely to visualize internal structures in opaque objects. This highly effective and non-invasive technique has found a broad range of industrial, academic, and clinical applications. The history of X-ray imaging can be divided into three epochs. The first begins with the formal discovery of X-rays by Röntgen in 1895. In the years that followed, a vigorous race unfolded in the scientific community. The primary objectives of many scientists were to uncover the nature of these then mysterious (dubbed "X") rays and to identify new applications [1].

The widespread popularity of X-rays in medicine eventually led to the revolutionary invention of computed tomography (CT) imaging in 1972 by Godfrey Hounsfield. This marked the start of the second epoch, which saw the creation and subsequent development of a tool to address the limitation of 2D, or projection, imaging. With this, imaging scientists could render an accurate representation of internal 3D structures. In medicine, CT imaging led to novel approaches in diagnosing diseases and complications with greater accuracy.

Unlike the first two, the third epoch lacks a well-defined start. Instead, it is defined here to include important (and somewhat recent) technological advancements that have enabled the significant improvement of existing techniques and the emergence of promising ones. Three notable advancements are the introduction of highly coherent X-ray sources like synchrotron facilities, the development of high-resolution photon-counting detectors, and the arrival of powerful personal computing. All of these areas remain under active development today. X-ray phase-contrast imaging, which is the primary topic in this work, falls in the final stage.

In this section, elemental concepts related to medical X-ray imaging are discussed. This includes a brief description of how X-rays interact with matter to produce conventional (absorption-based) images. A short introduction to phase-contrast imaging and, subsequently, a description of spectral imaging with photon-counting detectors are also included. The former is kept to an overview of the technique since later parts of this work elaborate on the physics and show applications.

1.1 Fundamentals of medical X-ray imaging

The prevailing method for generating X-rays involves the controlled bombardment of a dense metallic target (commonly tungsten) with electrons. This process can be understood as a statistical phenomena, where each electron has a probability p of interacting with the target to produce an X-ray quanta [2]. For an arbitrarily large ensemble of N electrons, the probability, P, that the random variable, $\mathcal{N}^{\text{in}}$, describing the number of generated X-ray quanta will take the value n is given by the Binomial distribution

$$P(\mathcal{N}^{\text{in}} = n) = \binom{N}{n} p^n (1-p)^{N-n}. \tag{1}$$

With the expectation value of $\mathcal{N}^{\text{in}}$ defined as $E\left[\mathcal{N}^{\text{in}}\right] = \bar{n}$, it can be shown that the production of X-rays is a Poisson process. In other words, it satisfies

$$P(\mathcal{N}^{\text{in}} = n) = \frac{\bar{n}^n}{n!} e^{-\bar{n}}. \tag{2}$$

If we infer that the process of detecting photon also satisfies Eq. 1, then the sequence of production and detection of X-ray quanta is, as a whole, also a Poisson process [2].

The Poisson-distributed nature of measurements discussed here is essential to understand the statistical character of data collected with photon-counting detectors (PCDs). These are the class of detectors used for all experimental results included in this work. One implication of this statistical viewpoint is the treatment of X-rays as "particles". This will later be reconsidered to properly account for observations related to phase-contrast imaging.

Another defining characteristic of X-ray radiation is energy. Typical X-ray sources, like those in most clinical settings, emit radiation within a broad (polychromatic) spectrum. The upper limit of the spectrum can be set by regulating the kinetic energy of the bombarding electrons while the lower limit can be controlled by applying beam filters [3, 4]. In medical imaging, X-ray energies can range between 10 and 150 keV. The interested reader can find additional details about the

production of X-rays in introductory textbooks like [4]. For details on the topic of *X-ray photon statistics*, two pertinent sources are [2] and [5].

1.1.1 X-ray interaction mechanisms

So far, the interaction of X-rays with matter directly has not been considered. This section covers the processes by which the intensity of X-rays penetrating an object is attenuated. Attenuation refers to the process by which X-rays are either absorbed or scattered [3]. The quantity that describes the probability that X-ray quanta will be attenuated is the linear attenuation coefficient (LAC) μ, which has units of cm^{-1} and can be expanded in term of five interaction coefficients as:

$$\mu = \sigma_r + \sigma_c + \tau + \pi + \gamma, \tag{3}$$

where σ_r is the Rayleigh (or coherent) scattering component, σ_c is the Compton (or incoherent) scattering component, τ is the photoelectric absorption component, π is the pair-production component, and γ is the triplet-production component [3].

All of the quantities in Eq. 3 are energy-dependent and are typically given as spatial distributions when used to describe an object. For example, $\mu(x, y, z; E)$ defines the LAC at location (x, y, z) for energy E. The LAC depends on the elemental composition and mass density. However, the energy dependence of μ is characterized by the former while the mass density can be interpreted as a scaling factor. A related quantity that depends only on the elemental composition is the *mass attenuation coefficient*, which is found by dividing the LAC by the nominal density of a material.

Of the five interaction mechanisms represented on the right-hand side of Eq. 3, only the first three are relevant in medical X-ray imaging since π and γ are feasible only above 1.02 and 2.04 MeV, respectively. Furthermore, in the imaging of soft tissues, contributions to the LAC from Rayleigh scattering do not exceed 12% for energies above 30 keV [4]. This leads to the approximate yet common expression [6]

$$\mu(E) = \sigma_r(E) + \sigma_c(E) + \tau(E) \approx \sigma_c(E) + \tau(E), \tag{4}$$

where $E \in [10, 150]$ keV represents energies typical to biomedical applications. Fig. 1a shows the component attenuation curves of water, which closely resemble the properties of soft tissue, between 20 and 80 keV. The values for this plot, and for many more elements, compounds, and common materials can be found in the National Institute of Standards and Technology (NIST) database [7]. The relative contribution of σ_c, σ_r, and τ to the total attenuation are plotted in Fig. 1b. For all energies, σ_r remains near or below 10% while σ_c becomes the dominant term after 30 keV and accounts for over 80% of μ above 50 keV.

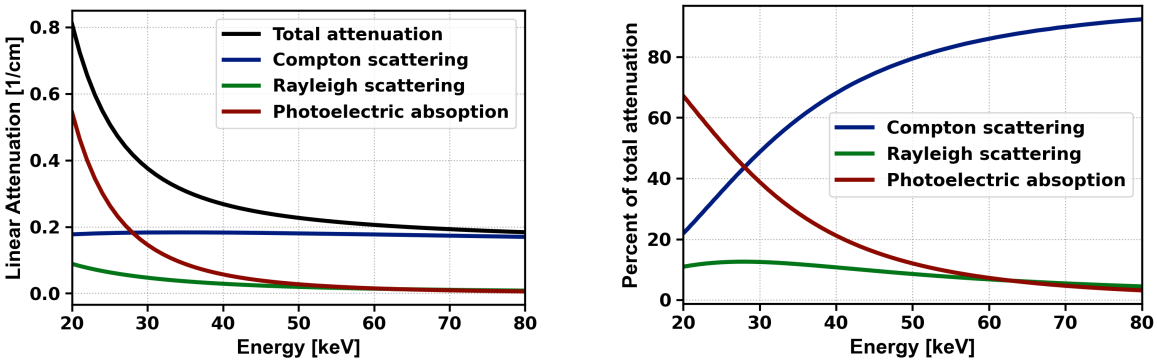

(a) (b)

Figure 1: (a) Mass attenuation curves for water. (b) Relative contribution of Rayleigh scattering (σ_r), Compton scattering (σ_c), and photoelectric absorption (τ) to the total linear attenuation coefficient (μ).

In *spectral imaging* applications, which involve multiple measurements of well-defined energies, it can be beneficial to express the LAC as a linear combination of energy-dependent basis functions. For example, we can write [6, 8, 9]:

$$\begin{aligned} \mu(E) &= a_0 \frac{1}{E^{n_r}} + a_1 \frac{1}{E^{n_p}} + a_2 \Phi_{KN}(E) \\ &= a_0 f_r(E) + a_1 f_p(E) + a_2 f_c(E), \end{aligned} \tag{5}$$

where $f_r(E)[= E^{-n_r}]$ and $f_p(E)[= E^{-n_p}]$ describe the energy functions of Rayleigh scattering and the photoelectric absorption, respectively. The energy dependence of Compton scattering, Φ_{KN}, is defined by the Klein-Nishina function. The quantities a_0, a_1, and a_2 describe energy-independent material properties. Although a_2 is understood to represent the electron density of a material, a

general consensus on the definitions of a_0 and a_1 is absent in the literature. Likewise, the values of the exponents n_r and n_p are material dependent [9] and therefore have application-specific definitions. It is common to see n_r and n_p defined as sets of acceptable values rather than a pair of specific numbers. Nevertheless, the usefulness of Eq. 5 for tasks like identifying material from X-ray images [6, 10] and retrieving the phase of X-rays [11] is recognized.

The intensity of X-rays at a plane coinciding with the detector (e.g. $z = 0$) can be related to the attenuating properties of a material by the *Beer-Lambert law*. For the case of monochromatic radiation traveling along the z-axis we have that

$$I(x, y, 0; E) = I^{\text{in}}(E) \exp\Big\{ - \int \mu(\boldsymbol{r}, E) \mathrm{d}z \Big\}. \tag{6}$$

Here $I^{\text{in}}(E)$ represents the intensity of X-rays incident on an object and $\boldsymbol{r}[= (x, y, z)]$ are spatial coordinates. Eq. 6 relies on particle-like interactions and, therefore, cannot be used to learn about the phase of X-rays.

1.2 Phase contrast imaging

Variations in the properties of a medium can result in changes to the amplitude and phase of electromagnetic waves like X-rays. Consequently, waves are refracted (or deviated) relative to their original direction. This can alter the intensity in ways that encode information about the phase-shifting properties of the medium. The material property that describes this type of interaction between X-ray and matter is the complex index of refraction

$$n = 1 - \delta + i\beta = 1 - \delta + i\frac{\mu}{2k} \tag{7}$$

where δ regulates changes to the phase (or refractions) of the radiation, $\beta[= \mu/2k]$ is proportional to the LAC, and k represents the wave-number. Like the attenuation properties described in Sec. 1.1.1, every quantity here is energy-dependent and can be expressed as a spatial distribution.

In general, changes to the phase of X-rays (and the ensuing path deviations) are extremely

small, since $\delta, \beta \ll 1$, and require specialized equipment and geometries to be detected [12, 13]. Phase-contrast imaging (PCI) refers to techniques where the measured intensities are sensitive to phase shifts. For many materials relevant to biomedical imaging, δ varies as E^{-2} while β varies as E^{-4} away from absorption edges [14, 15]. Thus, δ becomes relatively more important with increasing energy. Fig. 2 shows a comparison between the values of δ and β for aluminum and water, which highlights that δ is significantly larger for both materials. A large and growing number of PCI methods have been introduced over the years. They can be categorized based on the instrumentation or complexity required as well as the property of the phase (e.g. first or second derivative) they are most sensitive to [15].

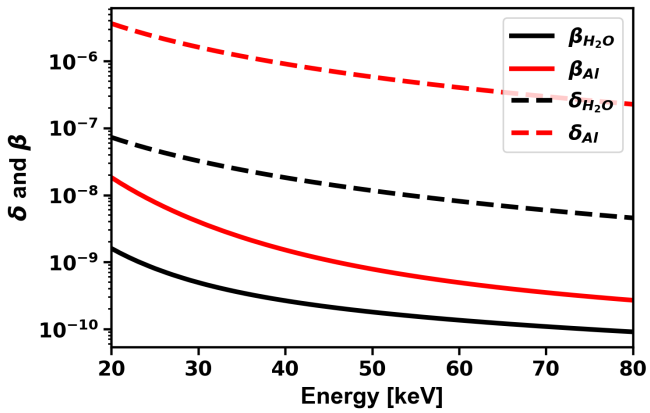

Figure 2: Real (δ) and imaginary (β) parts of the complex refractive index of aluminum and water.

Three promising PCI prospects for widespread application in biomedicine, listed in decreasing order of complexity, are grating-based interferometry, coded aperture, and propagation-based. Of these, the first two require careful alignment of specialized X-ray optics and are most sensitive to the first derivative of the phase. Both the phase and its derivative refer to the quantities along planes transverse to the propagation direction, i.e. the xy-plane for waves traveling along the z-axis. Propagation-based (PB) PCI is perhaps the simplest and most common technique [16]. The characteristic of the phase that is most relevant for PB-PCI is the second derivative or curvature. Several reviews about these and other methods have been published [12, 15–17]. Wilkins et al. (2014) provide a comprehensive overview highlighting some important shortcomings and advantages

of the techniques discussed here.

1.3 Spectral X-ray imaging

As we saw in Sec. 1.1.1, the intensity of X-rays reaching a detector placed immediately after an object is regulated by the LAC. Thus, in principle, Eq. 6 can give a method to identify materials from a single energy measurement when the X-ray energy is known. In practice, however, this is complicated by the dependence of the LAC on material density, which is often unknown. As a consequence, two materials with different chemical compositions can lead to identical (and indistinguishable) intensities [18]. A way to circumvent this limitation is with the use of multiple energy measurements to help uncover the energy dependence of a material and, with that, identify the related composition. This approach is fundamental to all spectral X-ray techniques.

Energy-specific X-ray measurements can be obtained in a variety of ways, each having its merits and disadvantages. Herein we focus on the use of photon-counting detectors in combination with polychromatic X-ray sources. A comprehensive description of this and other spectral X-ray imaging techniques relevant to biomedical imaging is given in [18].

1.3.1 Polychromatic radiation and its consequences

Since typical X-ray sources are polychromatic, the intensity measured at a detector element can be written in a more general and applicable form (compared to Eq. 6) as:

$$I(x,y,0;E_{\text{eff}}) = \int_{E_{\min}}^{E_{\max}} N^{\text{in}}(E) w_d(E) Q(E) \exp\left\{ -\int \mu(x,y,z;E)\mathrm{d}z \right\} \mathrm{d}E. \tag{8}$$

Here $N^{\text{in}}(E)$ is the number of incident photons with energy E propagating along the z-axis. The position of the detector is assumed to coincide with the plane $z = 0$. The functions $Q(E)$ and $w_d(E)$, respectively, describe the detection efficiency for photons with energy E and the weight given to them by the detector. The *effective energy*, E_{eff}, is defined as the energy of a monochromatic beam that would yield the same intensity as the right-hand side of Eq. 8 [19]. It is sometimes possible to

obtain an accurate estimate of E_{eff} experimentally [20].

When the bandwidth $\Delta E[= E_{\max} - E_{\min}]$ of the spectrum is not negligible, several undesirable effects must be carefully considered. An example arises when comparing Eq. 8 for monochromatic and polychromatic radiation. For the former, Eq. 8 affords a simple relationship between the projected LAC ($\int \mu(x,y,z;E)\mathrm{d}z$) and the measured intensity. In particular,

$$-\ln\left[\frac{I(x,y,0;E)}{N^{\text{in}}(E)w_d(E)Q(E)}\right] = -\ln\left[\frac{I^{\text{in}}(E)\exp\{-\int \mu(x,y,z;E)\mathrm{d}z\}}{I^{\text{in}}(E)}\right] = \int \mu(x,y,z;E)\mathrm{d}z, \quad (9)$$

where $I^{\text{in}}(E)[= N^{\text{in}}(E)w_d(E)Q(E)]$ is the intensity of the photons detected in the absence of an object along the beam path. Eq. 9 shows that the projected LAC is not influenced by the w_d and Q for the case of monochromatic radiation. On the other hand, a more complicated relationship is used to estimate the projected LAC [21], that is

$$-\ln\left[\frac{I(x,y,0;E_{\text{eff}})}{\int_{E_{\min}}^{E_{\max}} I^{\text{in}}(E)\mathrm{d}E}\right] = -\ln\left[\frac{\int_{E_{\min}}^{E_{\max}} I^{\text{in}}(E)\exp\{-\int \mu(\boldsymbol{r},E)\mathrm{d}z\}\mathrm{d}E}{\int_{E_{\min}}^{E_{\max}} I^{\text{in}}(E)\mathrm{d}E}\right] \approx \left\langle \int \mu(x,y,z;E)\mathrm{d}z \right\rangle_{I^{\text{in}}(E)}. \quad (10)$$

The right-hand side describes an ensemble average of the projected LAC with respect to the incident spectrum, i.e. $I^{\text{in}}(E)$, which can be used to identify a material when $I^{\text{in}}(E)$ is known for all energies. Yet, the validity of the approximation in Eq. 10 deteriorates (rapidly) as the bandwidth of the incident spectrum increases [21]. Also, values predicted by Eq. 10 are sensitive to w_d and Q.

Another negative consequence of imaging with a broad polychromatic spectrum is the possibility for beam hardening artifacts. These result from the preferential attenuation of lower energy photons, which leads to an X-ray spectrum that is thickness dependent. Put differently, the spectrum within an object can be significantly different (especially in the lower energy region) than the one that first interacted with the object's surface. The detrimental consequences of beam hardening, especially in tomographic imaging, are well documented in the literature, and its mitigation remains an area of active research [21–23].

1.3.2 Photon-counting detectors

Energy integrating detectors, the kind most commonly used in diagnostic imaging, combine all impinging photons into one intensity measurement. In addition, photons are weighted proportionally to their energy. This favors high-energy photons, which convey less contrast compare to the low-energy counterparts. In term of Eq. 8, this corresponds to defining the limits of the integral to encompass the full extent of the incident spectrum and ensuring that $w_d(E) \propto E$. Thus, measurements taken by these devices are highly susceptible to the complications outlined above. By contrast, energy resolving PCDs offer a way to accept photons confined to a narrow user-defined energy window. The major advantages of modern PCDs include the ability to obtain multiple spectroscopic measurements simultaneously, with high detector efficiency, high resolution, and insensitivity to electronic noise.

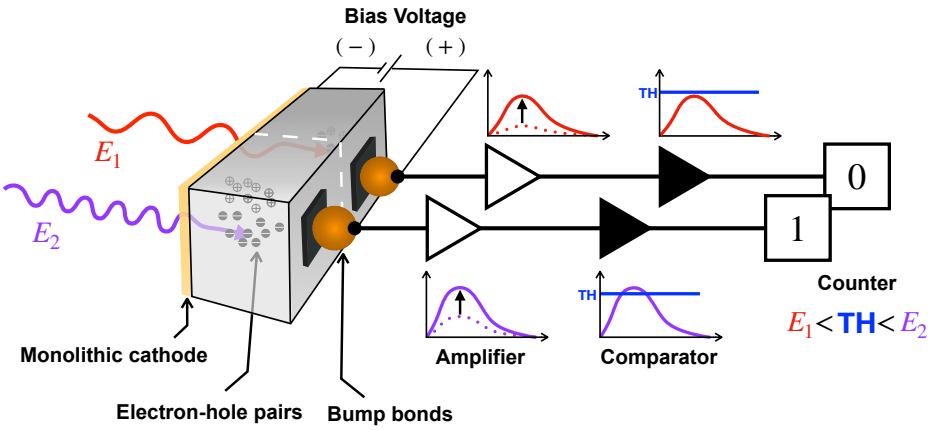

Figure 3: Two enlarged pixels of a PCD illustrating the basic steps that enable spectroscopic imaging. Photons deposit a charge proportional to their energy. The charge is converted to a voltage that is later amplified, compared to a user-defined threshold (TH), and counted if it exceeds the threshold.

For this work we used high-resolution (pixel pitch of 55 μm) PCDs from the Medipix3 collaboration, which were developed at CERN (Geneva). Even though the Medipix class of detectors was originally designed for use in high energy and particle physics, their utility for biomedical applications is widely recognized today [18, 24, 25]. Aside from a high spatial resolution, Medipix3 PCDs offer high detection efficiency, virtually zero electronic noise, energy resolutions of a few keV, and

the ability to record up to eight (different) spectroscopic measurements simultaneously [26].

Fig. 3 provides a simplified representation of how PCDs discriminate photons based on their energy. Namely, the incident radiation interacts with the sensor material resulting in a voltage pulse. The pulse is, in turn, amplified, compared to a user-defined threshold (TH) value, and counted if it exceeds the threshold. Only photons that produce signals within predefined energy levels are thus counted, allowing for spectroscopic imaging.

1.4 Outline

Section 2 of this dissertation explores the theory of propagation-based phase-contrast imaging. To this end, an overview of the problem within the context of the Fresnel-Kirchhoff diffraction theory is given. This section also introduces the transport of intensity equation. A breakdown of the experimental arrangement and conditions required by this method is also presented.

In Section 3 a phase retrieval algorithm developed by our group is introduced along with an in-depth analysis of its benefits and limitations. This includes a close look at the material properties that can be accurately estimated. A detailed analysis of how the phase retrieval algorithm handles noise in the input data is presented in this section.

A case study using the spectral phase retrieval of Section 3 is covered in Section 4. This report demonstrates that judicious implementation of our method can lead to significant improvements in quality as well as accurate recovery of material properties. We tested the algorithm with biologically relevant samples and, thus, believe the results to be of significance to clinical application.

In Section 5, the main focus is to provide an in-depth analysis of the consequences related to simplifying the transport of intensity equation with the weakly interacting object approximation. This is a key assumption in the phase retrieval of Section 3. To conduct our analysis, a framework that linked material properties (i.e. thickness and linear attenuation) to the breakdown of this approximation was derived. Through simulation, we were able to uncover details about the influence of this approximation in phase retrieval. The theoretical predictions were tested with experimental data, which lead to insightful and encouraging results.

2 Propagation-based phase contrast imaging

Propagation-based phase-contrast imaging (PB-PCI) is possibly the simplest and most common method to obtain phase-sensitive measurements in X-ray imaging. A distinguishing characteristic of PB-PCI that relates to its simplicity is the lack of optics along the beam path. In principle, phase-related signals in this imaging mode form through the free-space propagation of the X-ray wavefront perturbed by the sample. The only difference from a conventional imaging geometry is a suitable–and relatively large–separation between the sample and the detector. Although this basic description ignores other important differences, such as stricter requirements on the minimal lateral source coherence and detector resolution, it helps illustrate what is perhaps the most attractive aspect of PB-PCI.

Today, propagation-based imaging is a mature technique with some examples in X-ray imaging dating back to the early 1960's [15]. Yet, it was not until the mid 1990's [27] that the arrival of third generation synchrotron facilities aided much of the early efforts to understand the experimental aspects of this technique using hard X-rays ($\lambda < 2$ Å). Around the same time, in 1996, a seminal work by Wilkins et al. [28] showed convincing proof that propagation-based experiments could be performed with microfocus polychromatic (laboratory) sources. The widespread use of this technique has led to many promising results in laboratory and preclinical applications [29–34].

This section provides a description of the underlying physics that governs the signal formation in PB-PCI. The treatment here is by no means exhaustive, since its aim is to help in understanding the results shown in other parts of this work. Thorough accounts of the material covered here can be found in [12, 14, 17, 35].

2.1 Experimental geometry and practical requirements

In propagation-based imaging experiments, a key requirement is a non-negligible distance between the detector and the sample. Interference patterns, which encode the refractive properties of the sample, intensify as they propagate over this distance. Similarly to conventional (contact)

radiography, the sample is placed some distance R_1 downstream of the source and a distance R_2 upstream of the detector while no optics are added along the beam path. For contact imaging, however, R_2 is set as close to zero as practically possible. On the other hand, the choice of R_1 and R_2 for PB-PCI must be made judiciously to achieve sufficient lateral spatial coherence and a desirable signal level while remaining in the *near-field regime*, c.f. Sec 2.2.4.

Fig. 4a shows a schematic of the typical geometry used for PB-PCI measurements. Throughout this work R_1 will be used to indicate the source-to-object distance and R_2 will represent the object-to-detector distance. The optic axis will be assumed to coincide with the z-axis of a three-dimensional Cartesian coordinate system. Fig. 4b depicts the location and terminology used to describe the planes immediately before and after the object. Here the former is called the *entrance plane* while the latter is defined to coincide with the $z = 0$ plane and called the *object* or *exit* plane.

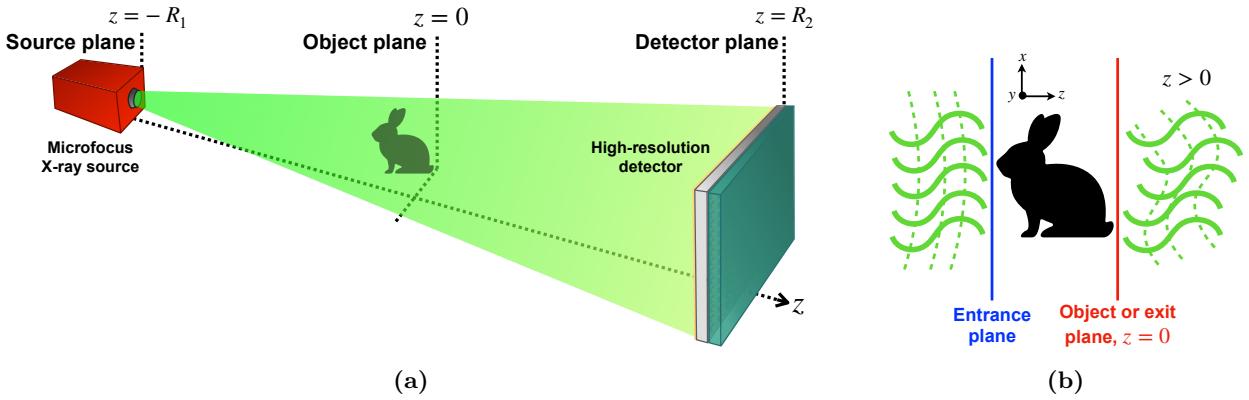

Figure 4: (a) Depiction of a propagation-based phase contrast experiment geometry. The X-ray source and the sample are separated by a distance R_1 while the distance from the exit-plane of the sample to the detector is R_2. In conventional X-ray imaging, $R_2 \approx 0$. (b) Sketch showing the entrance and exit planes. The former is defined as the plane immediately upstream of the object while the latter is the plane immediately downstream of the object. In PCI, X-rays propagate from the exit plane towards the detector plane.

Another critical component in a PB-PCI system is a source that provides sufficient lateral spatial coherence. Practical sources have a finite spatial extent, or *focal spot size*, which hinders the coherence of the X-ray waves. To illustrate this concept, one can consider a finite source represented as a spatial distribution of point sources with varying intensities. For simplicity, lets

assume that the source can be properly represented as a two-dimensional Gaussian distribution of intensities along the (x, y)-plane having a standard deviation $\sigma_{x,y}$ defined in units of cm. The lateral spatial coherence of the radiation impinging on the object can be estimated as [15, 36]

$$\ell_{x,y} = \frac{R_1 \lambda}{d} = 5.265 \times 10^{-8} [\text{keV} \cdot \text{cm}] \frac{R_1}{E \sigma_{x,y}}, \tag{11}$$

where d represents the full width at half maximum ($\sim 2.355\sigma_{x,y}$) of the source intensity distribution and λ is the wavelength of the X-ray radiation. The final result was obtained using $\lambda = hc/E$, where $h[\sim 4.136 \times 10^{-18}$ keV $\cdot$ s] and $c[\sim 2.998 \times 10^{10}$ cm/s] are Plank's constant and the speed of light in free-space. Here E is the energy of the X-rays in keV and R_1 has units of cm. As the ratio $R_1/E\sigma_{x,y}$ indicates, perfect coherence (i.e. $\ell_{x,y} \to \infty$) is only achieved for the case of an infinitesimally small source (i.e. $\sigma_{x,y} \to 0$). Furthermore, if the source-to-object distance and the focal spot size remain unchanged, the coherence deteriorates with increasing energy. On the other hand, the coherence improves when the source-to-object distance is increased. To ensure a significant phase contrast $\ell_{x,y}$ needs to satisfy $\ell_{x,y} \geq \nu_f^{-1}$, where ν_f is the spatial frequency (in cm^{-1}) corresponding the feature size of interest [15].

Unlike with spatial coherence, the requirements on temporal (or chromatic) coherence in near-field propagation-based imaging are less stringent. In fact, as reinforced in the next section, the character of the signals is essentially insensitive to the energy (and energy spread) of the incident beam [14, 35]. The energy does influence the signal amplitude, which can be important to achieve a desirable signal-to-noise level.

Detectors used in propagation-based methods must be able to resolve, and thus benefit from, interference (or diffraction) patterns in the signal. For a source of quasi-spherical waves, e.g. a microfocus source, the necessary detector resolution must be smaller than $\sqrt{MR_2\lambda}$, where $M[= 1 + R_2/R_1]$ is the geometrical magnification factor [16]. Therefore, building an effective PB-PCI system requires careful consideration of the source, detector, and geometrical configuration.

2.2 Physics of propagation-based imaging

Suppose a scattering object is composed of a linear and nonmagnetic medium that can be properly described using the complex index of refraction, $n(\boldsymbol{r}, \omega)$, introduced in Sec. 1.2. Furthermore, let the source of X-ray wave-fields be monochromatic and sufficiently far away so that the radiation incident on the object can be aptly treated as plane waves traveling down the z-axis. If the refractive index of the object is essentially constant (i.e. slowly varying) over length-scales of the order of λ (the wavelength of the X-rays), it follows that the Cartesian components $\Psi(\boldsymbol{r}, \omega)$ of the X-ray wave-field satisfy [37]:

$$\left(\boldsymbol{\nabla}^2 + k^2\right) \Psi(\boldsymbol{r}, \omega) = -4\pi F(\boldsymbol{r}, \omega)\Psi(\boldsymbol{r}, \omega). \tag{12}$$

Here $k[= 2\pi/\lambda]$ is the wavenumber, $\boldsymbol{r}[= (x, y, z)]$ are spatial coordinates, and $\boldsymbol{\nabla}^2[= \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}]$ is the Laplacian operator in three dimensions. The components $\Psi(\boldsymbol{r}, \omega)[= \psi(\boldsymbol{r})e^{-i\omega t}]$, commonly referred to as *wave functions*, are complex-valued and represent monochromatic propagating X-rays. The quantity $F(\boldsymbol{r}, \omega)$ is the *scattering potential* of the medium and is given by

$$F(\boldsymbol{r}, \omega) = \frac{1}{4\pi} k^2 [n^2(\boldsymbol{r}, \omega) - 1]. \tag{13}$$

Eq. 12, which is sometimes called the *inhomogeneous Helmholtz equation*, is a central result in scattering theory [37]. As Paganin and Pelliccia elegantly show in [17], Eq. 12 can be use to obtain a representation of the field at the exit plane, c.f. Sec. 2.1 and Fig. 4a, when *paraxial* propagation through the object and the *projection approximation* are enforced. The former requires the direction of the X-ray wave-field perturbed by the medium to remain effectively parallel to the propagation axis of the incident (unperturbed) field as it propagates through the object. If valid, the wave function can be written as the product of a "perturbation envelope" $u(\boldsymbol{r})$ and the unperturbed field $\exp\{ikz\}$ [12, p. 72]. Therefore, $\psi(\boldsymbol{r}) = u(\boldsymbol{r}) \exp\{ikz\}$, where the dependence on frequency

was omitted. Additionally, we can write Eq. 12 as

$$\left(2ik\frac{\partial}{\partial z} + \boldsymbol{\nabla}_\perp^2\right) u(\boldsymbol{r}) = -4\pi F(\boldsymbol{r})u(\boldsymbol{r}), \tag{14}$$

where $\boldsymbol{\nabla}_\perp^2 [= \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}]$ are the transverse components of the Laplacian operator.

When the projection approximation is valid at the exit plane, the contributions from $\boldsymbol{\nabla}_\perp^2$ in Eq. 12 can be ignored to show that [12, p. 73–74]

$$\begin{aligned} u_{\mathrm{o}}(x,y) &\approx u^{\mathrm{in}}(x,y)\tau(x,y) \\ &= u^{\mathrm{in}}(x,y)\exp\left\{i\varphi(x,y) - \frac{\mu_p(x,y)}{2}\right\}, \end{aligned} \tag{15}$$

where the subscript "o" indicates values at the exit (or object) plane, $u^{\mathrm{in}}(x,y)$ represents the field entering the object, and $\tau(x,y)[= \exp\{i\varphi(x,y) - \mu_p(x,y)/2\}]$ is the *transmission function* that encapsulates the properties of the object. In Eq. 15, the phase shift induced by the object is represented by $\varphi(x,y)$ while $\mu_p(x,y)[= \int \mu(\boldsymbol{r})\mathrm{d}z]$ gives the projected linear attenuation coefficient, c.f. Sec. 1.1.1. Another important consequence of the projection approximation is that the phase shifts can be written in terms of the projected electron density, $\int \rho_e(\boldsymbol{r})\mathrm{d}z$, that is:

$$\varphi(x,y) = -\lambda r_e \int \rho_e(\boldsymbol{r})\mathrm{d}z, \tag{16}$$

where r_e is the classic electron radius. From Eq. 15, the wave-field at the exit plane can be related to the properties of the object.

Once the field exits the object and propagates through (essentially) free-space, it evolves according to Fresnel diffraction theory. In free-space $n(\boldsymbol{r}) = 1$ so that $F(\boldsymbol{r}) = 0$ and Eq. 14 can be used to arrive at the *free-space paraxial wave equation*:

$$\left(2ik\frac{\partial}{\partial z} + \boldsymbol{\nabla}_\perp^2\right) u(\boldsymbol{r}) = 0. \tag{17}$$

If the paraxial approximation is valid on the half-space $z > 0$ and we define $u^{\mathrm{in}} = \exp\{ikz\}$ in

Eq. 15, then using the Fresnel-Kirchhoff integral we can write the field at z that satisfies Eq. 17 as [14, 35]

$$u_z(x,y) = \frac{i\exp\{ikz\}}{\lambda z}\int \tau(X,Y)\exp\left\{\frac{i\pi}{\lambda z}\left[(x-X)^2+(y-Y)^2\right]\right\}\mathrm{d}X\mathrm{d}Y, \tag{18}$$

where the notation $u_z(x,y) = u(x,y,z)$ was used for conciseness. In Fourier-space, this integral is represented more practically as

$$\tilde{U}_z(\boldsymbol{\nu}) = \exp\{ikz\}\tilde{T}(\boldsymbol{\nu})\exp\left\{i\pi\lambda z|\boldsymbol{\nu}|^2\right\}, \tag{19}$$

where $\tilde{U}$ and $\tilde{T}$ are the Fourier transforms (FT) of u and τ, respectively. Here the Fourier-domain vector $\boldsymbol{\nu}[=(\nu_x,\nu_y)]$ is the conjugate of (x,y) and $|\cdot|$ is the modulus operator. An important use of Eq. 19 involves simulating the response of propagation-based imaging systems [38, 39]. Two examples are a study of the performance of CO_2 as a contrast agent in PB-PCI [40] and the development of a simulation packages for phase-contrast X-ray breast imaging [41]. Eq. 19 also helps uncover some fundamental properties of propagation-based imaging, as demonstrated in the following sections and by [14, 35].

2.2.1 Weakly interacting object approximation

The subsequent analysis relies on the *weakly interacting object approximation* and similar accounts can be found in [14, 35, 42]. Under this approximation, the values of the φ and the μ_p are assumed to be sufficiently small so that the transmission function, to first order, can be written as

$$\tau(x,y) \approx 1 + i\varphi(x,y) - \mu^*(x,y), \tag{20}$$

where $\mu^* = \mu_p/2$. Upon substituting the FT of Eq. 20 into Eq. 19, the result is

$$\tilde{U}_z(\boldsymbol{\nu}) \approx \left[\delta(\boldsymbol{\nu}) + i\tilde{P}(\boldsymbol{\nu}) - \tilde{M}(\boldsymbol{\nu})\right]\exp\{i\chi(\boldsymbol{\nu})\}, \tag{21}$$

where $\delta(\boldsymbol{\nu})$ is a two-dimensional Dirac delta function in Fourier-domain and $\chi(\boldsymbol{\nu}) = \pi\lambda z|\boldsymbol{\nu}|^2$. The FT of φ and μ^* are represented by $\tilde{P}$ and $\tilde{M}$, respectively. It follows from Eq. 21 that:

$$u_z \approx 1 - \varphi\mathcal{F}\{\sin\chi\} - \mu^*\mathcal{F}\{\cos\chi\} + i\left[\varphi\mathcal{F}\{\cos\chi\} - \mu^*\mathcal{F}\{\sin\chi\}\right] \tag{22}$$

where the explicit dependence of u_z, μ^*, and φ on (x, y) as well as the dependence of χ on $\boldsymbol{\nu}$ was omitted. The operator $\mathcal{F}\{\cdot\}$ represents a two-dimensional Fourier transform.

The FT of the intensity, $I_z(x, y)[= |u_z(x, y)|^2]$, corresponding to Eq. 22 and containing only first-order terms is:

$$\tilde{I}_z \approx \delta + 2\tilde{P}\sin\chi - 2\tilde{M}\cos\chi, \tag{23}$$

where $\tilde{I}_z$, $\tilde{P}$, $\tilde{M}$, δ, and χ all depend on $\boldsymbol{\nu}$. Much can be learned from Eq. 23. For example, if the *near-field* regime is defined to correspond with values of χ sufficiently small so that $\sin\chi \approx \chi$, it follows that for materials with negligible attenuation ($\tilde{M} \approx 0$) but non-negligible phase-shifts the measured intensity is [35]

$$\begin{aligned} I_z(x, y) &= 1 - \frac{\lambda z}{2\pi}\boldsymbol{\nabla}_\perp^2\varphi(x, y) \\ &= 1 + \frac{1}{E^2}\frac{zr_e h^2c^2}{2\pi}\boldsymbol{\nabla}_\perp^2\int \rho_e(\boldsymbol{r})\mathrm{d}z \end{aligned}. \tag{24}$$

This result highlights how in the near-field regime, and when Eq. 16 is valid, the character of signals caused by phase shifts is controlled by variations in the electron density of the sample, which is independent of energy. The energy, in this case, acts only as a scalar multiplier. As explained further in the next section, Eq. 23 is also useful to learn about the factors that influence the transfer of contrast in propagation-based systems [35].

2.2.2 Phase and amplitude contrast transfer functions

The transfer of phase and amplitude contrast predicted by Eq. 23 is regulated by $\sin\chi$ and $\cos\chi$, respectively. Plotting these two contrast transfer functions (CTFs) in a coordinate system defined

by $(\sqrt{\lambda z}\nu_x, \sqrt{\lambda z}\nu_y)$ is instructive, since it reveals how they are affected by energy ($\propto \lambda^{-1}$), propagation distance (z), and the frequency content of the object ($\boldsymbol{\nu}$). Fig. 5 shows the two-dimensional phase and amplitude CTFs (a) as well as one-dimensional cross-sections along $\nu_y = 0$ for a system with perfect resolution (b) and after the influence of an arbitrary point-spread-function (PSF) is considered (c).

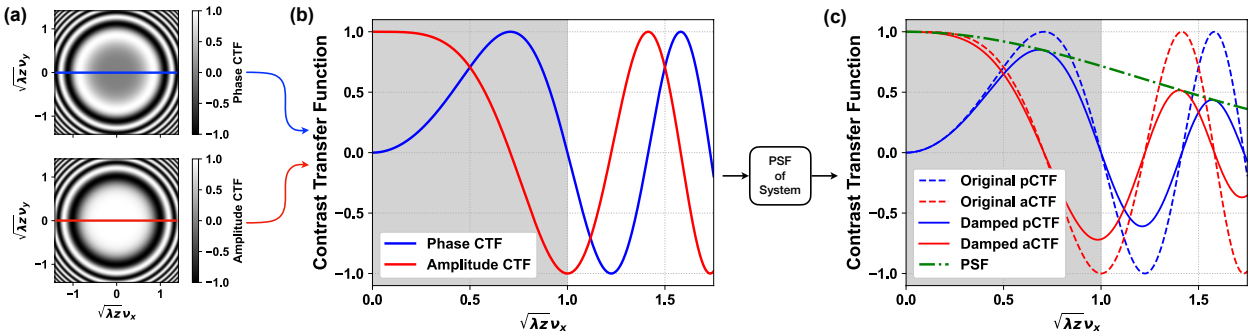

Figure 5: (a) Radially symmetric two-dimensional contrast transfer functions for the phase and amplitude. (b) One dimensional profiles along $\nu_y = 0$. The direct imaging regime $\sqrt{\lambda z}\nu_x < 1$ is highlighted. This region includes the optimum-phase contrast where the phase CTF is at a maximum. It also includes the near-field regime, where $\sin\chi \approx \chi$. (c) Shows how the smoothing effect of the point-spread-function of a system can impact the phase and amplitude CTF, labeled as pCTF and aCTF.

In Fig. 5a, the oscillation rate of the CTFs increases away from the origin. To understand the consequence of this, consider how two periodic features of close but different frequencies would be treated if the energy of X-rays is held constant. At small propagation distances, such that $\chi < \pi$, the two features will likely receive a similar gain factor from each CTF since they will be in close proximity and in a region where the CTFs oscillate gently. On the other hand, for sufficiently large propagation distances, the handling of each feature can vary significantly due to the rapidly changing CTFs. For example, the CTFs can introduce differences in polarity between the pair of features. These vastly different treatments of the phase and amplitude of features with similar frequency content can lead to the visually unrecognizable images that are known to emerge in the holographic regime [14]. A similar argument can be made if one instead considers decreasing the energy while maintaining a fixed propagation distance.

A gray band in Fig. 5a and 5b highlights the "direct-imaging" regime [14], which encapsulates

the near-field and the optimum phase-contrast regions. It corresponds to $\chi < \pi$. In this region, the two contrasts are complementary. For instance, phase contrast favours high-frequency features while amplitude contrast performs best for negligible propagation distances. In Fig. 5b, we see how the PSF of a system dampens the CTFs, thus leading to lower contrast. Since the PSF decreases in amplitude with increasing frequency, its influence can be especially detrimental to phase contrast.

2.2.3 The transport of intensity equation

The *transport of intensity equation* (TIE) provides a functional relationship between the longitudinal evolution of the intensity and variations of the phase and intensity of the wave-field along transverse planes. In other words, it can be used to relate the intensity and phase of X-rays on a plane (e.g. the object plane) to the intensity measured some perpendicular distance, e.g. R_2, away. For this reason, this equation is widely used in the development of algorithms that obtain quantitative estimates of the phase, called phase-retrieval algorithms. The TIE can be derived by first defining $\psi_z(x, y)$ as [43]

$$\psi_z(x, y) = \sqrt{I_z(x, y)} \exp\{i\varphi_z(x, y)\}, \tag{25}$$

such that $|\psi_z(x, y)|^2 \equiv I_z(x, y)$. The imaginary part of the equation that results from inserting Eq. 25 into Eq. 14 is the TIE [12] and is given by:

$$-k\frac{\partial}{\partial z} I_z(x, y) = \boldsymbol{\nabla}_\perp \cdot I_z(x, y)\boldsymbol{\nabla}_\perp \varphi_z(x, y). \tag{26}$$

In a strict sense, Eq. 26 assumes monochromatic and planar radiation. For this reason, it is of practical convenience to use modern synchrotron sources, which can provide quasi-monochromatic and quasi-planar X-rays. On the other hand, the more easily accessible polychromatic microfocus sources emit quasi-spherical radiation. Fortunately, this is readily addressed by accounting for geometrical magnification [42]. Also, the assumption of monochromaticity is met by properly estimating an effective energy for the polychromatic measurement.

It is common to analyze Eq. 26 at the plane $z = 0$, especially for X-ray techniques where the aim is to estimate the properties at $z = 0$ and the values of R_1 and R_2 are kept constant (e.g. single-shot propagation-based methods). In the near-field regime, it is possible to expand the z-derivative in Eq. 26 as a finite difference so that

$$-k\frac{I_{R_2}(x,y) - I_0(x,y)}{R_2} = \boldsymbol{\nabla}_\perp \cdot I_z(x,y)\boldsymbol{\nabla}_\perp \varphi_z(x,y)\big|_{z=0}, \tag{27}$$

or similarly,

$$I_{R_2}(x,y) = I_0(x,y) - \frac{R_2}{k}\boldsymbol{\nabla}_\perp \cdot I_z(x,y)\boldsymbol{\nabla}_\perp \varphi_z(x,y)\big|_{z=0}. \tag{28}$$

Eq. 28 explicitly shows how the TIE relates the intensity at the detector plane, I_{R_2}, to the intensity at the exit surface, I_0, and a term proportional to the propagation distance, R_2, and variations in the phase and intensity.

2.2.4 Validity of the near-field and projection approximations

Since much of the theory in this section is valid in the near-field regime and depends on the projection approximation, the requirements for these conditions to hold in X-ray imaging is important to understand. The scope will be limited to "soft" biological materials, i.e. those with low ($\lesssim 13$) effective atomic numbers like body parts. For the projection approximation to be accurate, the size of the finest feature of interest must be significantly larger than $\sqrt{\lambda\mathcal{T}}$, where $\mathcal{T}$ is the object thickness [44]. We also request for the Fresnel number, N_F, to satisfy the inequality

$$N_F \equiv \frac{a^2}{\lambda\mathcal{T}} \gg 1, \tag{29}$$

where a is the size of the smallest feature of interest [12].

One can show that this condition is satisfied for resolutions down to 10 μm at 20 keV in objects as thick as 80 cm [44]. As a practical example, consider the case of mammography, which poses the

highest resolution requirement in radiology [4]. In approximately 30 % of patients, microcalcifications in the breast are accompanied by malignant lesions [45]. If we establish that the inequality in Eq. 29 is satisfied whenever $N_F \geq 10$ [17], then the limits of the projection approximation can be mapped for compressed breast thicknesses up to 8 cm at an energy of 20 keV ($\lambda \approx 62$ Å). Fig. 6a shows the Fresnel number that would result for microcalcifications ranging between 1 to 10 microns. A white line indicates the datum where $N_F = 10$ so that all configurations to the right of it satisfy the projection approximation. Note that this condition is met for microcalcifications of a few microns in relatively thick breasts.

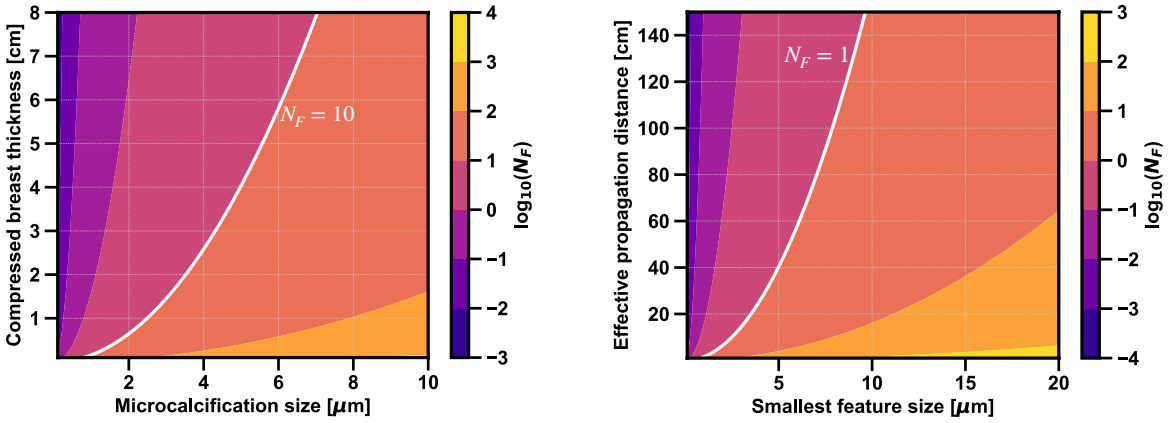

(a) **(b)**

Figure 6: (a) Plot of the Fresnel number, as defined in Eq. 29, for compressed breast thicknesses ranging up to 8 cm and microcalcifications below 10 microns. A white line shows the $N_F = 10$ datum, which marks the limit of validity for the projection approximation. (b) Plot showing the Fresnel number as defined in Eq. 30 and a white line marking the near-field regime. An energy of 20 keV was used for both plots.

The near-field (or Fresnel) region can be delineated by replacing $\mathcal{T}$ in Eq. 29 with the effective propagation distance R_{eff}, i.e. $N_F = a^2/\lambda R_{\text{eff}}$, and defining the boundary [13, p. 307]

$$\text{Near-field region:} \qquad R_{\text{eff}} \approx \frac{a^2}{\lambda} \quad \text{or} \quad N_F \approx 1. \tag{30}$$

Fig. 6b shows the Fresnel number for features up to 20 microns in size and propagation distances up to 150 cm. The white contour marks the boundary where $N_F \approx 1$. Values much greater than this approach the contact regime while values much smaller correspond to the Fraunhofer (or far-field)

regime. From the plot, the near-field approximation for features of 10 microns can be satisfied. Additionally, the contact and near-field regime for distances up to about 150 cm can be accessed.

3 Spectral phase retrieval assuming weak X-ray interactions

Improvements in signal strength combined with greater sensitivity to subtle fluctuations in material properties are characteristics of PCI that can benefit many important clinical X-ray imaging applications. Examples include breast imaging, lung imaging, and bone mineral-density estimation [29, 46, 47]. Such tasks often involve distinguishing between materials with differences in their attenuating properties that fall beneath those detectable by conventional (absorption-based) imaging. The intensities measured in PCI can encode information about both the phase-shifting and attenuating properties of matter. However, these methods do not represent a direct measurement of either property. For this reason, PCI measurements alone are not sufficient to identify materials and quantify their properties.

The task of recovering (or retrieving) material properties depends on how accurately one can describe and model the physical mechanisms and link these to observations. Therefore, models using a mathematical treatment of X-ray radiation that account for their wave nature are required. The process of using such models to retrieve material properties is known as phase retrieval (PR). Even within the narrow scope of propagation-based imaging, a wide variety of PR algorithms have been proposed in the literature [48, 49]. The two most popular mathematical frameworks used to derive PR algorithms for PB-PCI are the transport of intensity equation (TIE) and the contrast transfer function [48]. The following discussion will focus solely on TIE-based techniques where the necessary information for PR can be obtained in a single X-ray exposure. Therefore, methods that require measurements at multiple positions are not included since this may be unattainable or impractical in key biomedical applications like clinical computed tomography.

Two widely used phase retrieval techniques are the single-material (SM) approximation method [50] and the "phase-attenuation duality" (PAD) method [51]. With implementations that resemble the use of a low-pass frequency filter, these methods give an efficient way to estimate material properties from a single X-ray measurement. In particular, the SM technique can aid in estimating the projected thickness, while the PAD method can help find the projected electron density. Due to

their efficiency, the two techniques can be readily adapted to computationally demanding tasks like tomographic imaging [49]. However, the SM method [50] requires that samples be homogeneous and for their refractive properties to be known *a priori*. On the other hand, the "duality" condition, where the phase and the attenuation are both proportional to the electron density, is valid (with error below 10%) for biological materials above X-ray energies of 60 keV [51]. Even when these conditions are met, each technique estimates a single material property, which may not be adequate to differentiate closely-related materials.

Recently, our group proposed a phase retrieval approach capable of estimating the projected photoelectric (PE) absorption and electron density of a sample, which can be used to obtain the X-ray attenuation and phase delay caused by the sample [11, 52]. This proposed method requires a set of independent spectroscopic measurements. Unlike several alternatives [49], our phase retrieval approach is accurate for a broad range of X-ray energies (ranging from 20 keV to 100 keV) and presents fewer constraints on the heterogeneity of samples. The following section presents a comprehensive derivation of this technique. Similar accounts can be found in [11] and [53].

3.1 Theory and derivation

The phase retrieval scheme in [11] stems from the TIE, which was introduced in Sec. 2.2.3. If we express the dependence on energy, E, in Eq. 26 and expand the div(grad()) operation, then

$$-\frac{2\pi}{\lambda}\frac{\partial}{\partial z}I(\boldsymbol{r}_\perp, z; E) = \boldsymbol{\nabla}_\perp I(\boldsymbol{r}_\perp, z; E)\cdot\boldsymbol{\nabla}_\perp\varphi(\boldsymbol{r}_\perp, z; E) + I(\boldsymbol{r}_\perp, z; E)\boldsymbol{\nabla}_\perp^2\varphi(\boldsymbol{r}_\perp, z; E). \tag{31}$$

Here, $\lambda[\approx 12.4(\text{Å}\cdot\text{keV})/E(\text{keV})]$ represent the wavelength of the X-ray radiation while $\boldsymbol{r}_\perp[= x\hat{\boldsymbol{x}}+y\hat{\boldsymbol{y}}]$ and $\boldsymbol{\nabla}_\perp[= (\partial/\partial x)\hat{\boldsymbol{x}} + (\partial/\partial y)\hat{\boldsymbol{y}}]$ are the position vector and differential operator on the plane transverse to the X-ray's propagation direction. The transmitted intensity and phase at position (x, y, z) are represented by $I(\boldsymbol{r}_\perp, z; E)$ and $\varphi(\boldsymbol{r}_\perp, z; E)$, respectively. Equation 31 is non-linear and, therefore, difficult to implement directly. Instead, solutions are often found only after a number of simplifying approximations are made.

A common step in simplifying Eq. 31 involves assuming that the material(s) of interest introduce only weak variations to the phase and intensity of the impinging radiation. This approximation is sometimes called the *weakly interacting object approximation*. As a consequence, the propagation direction of the wave-field after interacting with the object remains nearly parallel to the original (unperturbed) direction and the projection approximation can be enforced [12]. Many materials, including biological samples, satisfy this approximation within a permissible tolerance [54]. As mentioned in Sec. 2.2.3, the TIE is further simplified when measurements are taken in the near-field regime, where the left-hand-side of Eq. 31 may be approximated by its final difference counterparts. With the simplifications above, and assuming plane wave illumination, the TIE can be modified to:

$$I(\boldsymbol{r}_\perp, R_2; E) = I(\boldsymbol{r}_\perp, 0; E)\left[1 - \frac{\lambda R_2}{2\pi}\boldsymbol{\nabla}_\perp^2 \varphi(\boldsymbol{r}_\perp, 0; E)\right]. \tag{32}$$

Under the projection approximation, the Beer-Lambert law can be used to write $I(\boldsymbol{r}_\perp, 0; E)$ in terms of the linear attenuation coefficient $\mu(\boldsymbol{r}_\perp, z; E)$, c.f. Eq. 6 in Sec. 1.1.1. By invoking a Rytov-like approximation [11] and replacing the plane wave illumination assumption with the field due to a point source some distance R_1 away from the sample, the following equation is derived:

$$\ln\left[\frac{I^{\text{in}}(M\boldsymbol{r}_\perp; E)}{I(M\boldsymbol{r}_\perp, R_2; E)}\right] = \int \mu(\boldsymbol{r}_\perp, z; E)\mathrm{d}z + \frac{\lambda}{2\pi}\frac{R_2}{M}\boldsymbol{\nabla}_\perp^2 \varphi(\boldsymbol{r}_\perp, 0; E), \tag{33}$$

where $M[= 1 + R_2/R_1]$ is the geometrical magnification factor and $\int \mu(\boldsymbol{r}_\perp, z; E)\mathrm{d}z$ represents the projected linear attenuation coefficient (pLAC). In past work by our team [11], we showed that the (non-linear) energy dependence of the TIE form given in Eq. 33 can be linearized by using a well-known decomposition of the pLAC [6]. After this process, the TIE gives an effective way to estimate the refractive properties of a sample (phase and absorption) from a set of two or more spectroscopic measurements, which can be acquired in a single exposure using photon counting

detectors. In particular, we write the pLAC as

$$\begin{aligned}\int \mu(\boldsymbol{r}_\perp, z; E)\mathrm{d}z &= \frac{1}{E^m}\int kZ^n(\boldsymbol{r}_\perp, z)\rho_e(\boldsymbol{r}_\perp, z)\mathrm{d}z + \Phi_{\mathrm{KN}}(E)\int \rho_e(\boldsymbol{r}_\perp, z)\mathrm{d}z \\ &= \frac{1}{E^m}a_1(\boldsymbol{r}_\perp) + \Phi_{\mathrm{KN}}(E)a_2(\boldsymbol{r}_\perp),\end{aligned} \tag{34}$$

where k is a constant related to the PE absorption, Z represents the effective atomic number, $\Phi_{\mathrm{KN}}(E)$ is the Klein-Nishina function, and $\rho_e(\boldsymbol{r}_\perp, z)$ is the electron density. The exponents n and m in Eq. 34 are typically defined so that n and m are both approximately 3 for biological samples [13,55]. A numerical fitting of tabulated PE absorption values [7] for a set of carbon-based plastics resulted in a nearly constant value for m of 3.30 ± 0.02, c.f. Appx. C. Interestingly, an essentially identical result was found by [9] for materials of similar composition.

The transmitted phase in Eq. 33 can be approximated as

$$\varphi(\boldsymbol{r}_\perp, 0; E) = -r_e\lambda a_2(\boldsymbol{r}_\perp), \tag{35}$$

where r_e is the classical electron radius. If we let $d(M\boldsymbol{r}_\perp; E) = \ln[I^{\mathrm{in}}(M\boldsymbol{r}_\perp; E)/I(M\boldsymbol{r}_\perp, R_2; E)]$, we get that:

$$d(M\boldsymbol{r}_\perp; E) = \left[\frac{1}{E^m}\right]a_1(\boldsymbol{r}_\perp) + \left[\Phi_{\mathrm{KN}}(E) - \frac{\lambda^2 r_e R_2}{2\pi M}\boldsymbol{\nabla}_\perp^2\right]a_2(\boldsymbol{r}_\perp), \tag{36}$$

where the terms in squared brackets are constants defined by the energy and geometry of the spectroscopic measurements. This forward model is the central result of this derivation.

A solution to Eq. 36 can be found in Fourier-space, where it can be readily adapted to a least squares, or a similarly efficient, solver. In particular, in Fourier-space Eq. 36 becomes

$$\mathcal{F}\{d(\boldsymbol{\xi}; E)\} = \left[\frac{1}{E^m}\right]\mathcal{F}\{a_1(\boldsymbol{\xi}/M)\} + \left[\Phi_{\mathrm{KN}}(E) + M\frac{\lambda^2 r_e R_2}{2\pi}|\boldsymbol{\nu}|^2\right]\mathcal{F}\{a_2(\boldsymbol{\xi}/M)\} \tag{37}$$

where we substituted $\boldsymbol{\xi}[= M\boldsymbol{r}_\perp]$ to represent the magnified coordinates at the detector plane. The operation $\mathcal{F}\{\cdot\}$ represents a 2D FT along the transverse plane, $\boldsymbol{\nu}[= \nu_x\hat{\boldsymbol{\nu}}_x + \nu_y\hat{\boldsymbol{\nu}}_y]$ is the Fourier-domain vector conjugate to $\boldsymbol{\xi}$, and $|\cdot|$ is the modulus operator. In this work we solved Eq. 37 using

weighted least squares with weights defined by the contrast-to-noise-variance ratio (CNVR) of each spectral measurement [56].

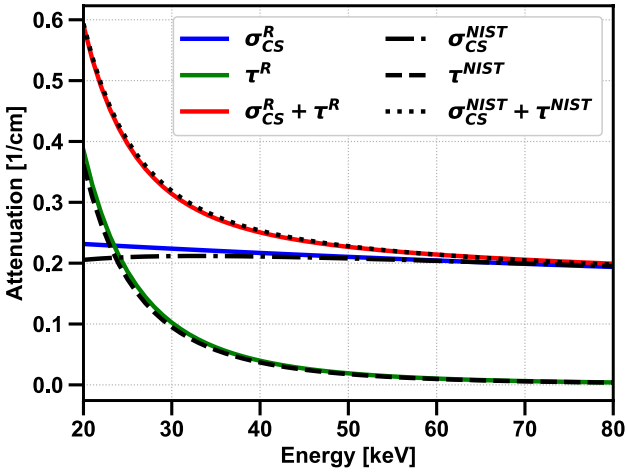

Figure 7: Retrieved and tabulated properties of polymethyl methacrylate (PMMA). Phase retrieval was performed on experimental measurements of a 3-mm PMMA rod obtained with a photon-counting detector and a polychromatic X-ray source.

After phase retrieval, a map of the projected electron density, $a_2(\boldsymbol{r}_\perp)$, and the energy independent part of PE absorption, $a_1(\boldsymbol{r}_\perp)$, are obtained. From these energy independent material properties, the projected values for CS and PE can be estimated at any energy using the first and second terms on the right-hand side of Eq. 34, [6]. Furthermore, the pLAC and transmitted phase can be found using Eq. 34 and 35, respectively. Finally, the laplacian of the transmitted phase can also be estimated using:

$$\boldsymbol{\nabla}_\perp^2 \varphi(\boldsymbol{\xi}/M; E) = \mathcal{F}^{-1}\left\{ -\left[M^2 \lambda r_e |\boldsymbol{\nu}|^2\right] \mathcal{F}\left\{a_2(\boldsymbol{\xi}/M)\right\}\right\}. \tag{38}$$

Figure 7 shows a comparison between tabulated [7] and retrieved values of Compton scattering, σ_{CS}, and photoelectric abortion, τ, over energies ranging between 20 to 80 keV. The retrieved values, which are indicated by the superscript R, were obtained using our spectral phase retrieval algorithm. The inputs were spectroscopic measurements of a 3-mm polymethyl methacrylate (PMMA) rod acquired in our lab with a polychromatic source and a photon-counting detector. As the plot shows, it is possible to obtain highly accurate estimates of material properties with the algorithm

presented in this section. In general, however, the accuracy will depend on factors like the effective energy of the measurements, the influence of experimental parameters on the approximations built into Eq. 36, and noise.

3.2 On the impact of data noise in spectral phase retrieval

In the absence of a noise mitigation strategy, the two retrieved properties, i.e. a_1 and a_2 in Eq. 36, display notable differences in quality when dealing with noisy data. This is illustrated by the cascaded framework shown in Fig. 8. Here the material properties of a phantom are simulated and used to generate the intensity expected by the phase retrieval algorithm, c.f. Appx. C. In other words, we use the forward model in Eq. 36 to generate synthetic data from a pair of projected material properties (a_1 and a_2) at specific energy values.

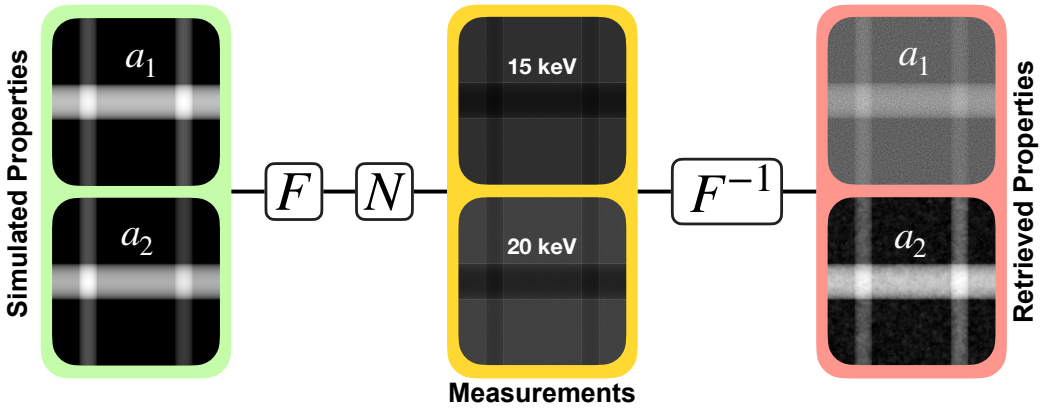

Figure 8: Cascaded framework used to test the sensitivity of the SPR algorithm to noise and other relevant sources of error in the data and the phase retrieval workflow. Here F represents the forward modeling process that generates measurements at defined energies and N represents the introduction of noise into the simulated measurements. The operator F^{-1} represents the phase retrieval step.

In Fig. 8, the forward simulation step is expressed by F while the retrieval process is represented as F^{-1}. If we describe a phantom with properties a_1 and a_2 as $P(a_1, a_2)$, then we can use the notation of operators and write the simulated measurements as $F[P(a_1, a_2)]$. Similarly, the phase retrieval process becomes $F^{-1}[F[P(a_1, a_2)]] = P(a_1, a_2)$. When Poisson noise is introduced in the simulated measurements, which corresponds to the application of the N operator in Fig. 8,

inaccuracies in the estimated terms after phase retrieval reflect the impact of noise. A qualitative inspection of the retrieved properties in Fig. 8 shows how the noise magnitude and the texture are different between the retrieved a_1 and a_2 terms.

A thorough understanding of the observed differences in the image-noise of the two terms can be obtained by investigating how noise is treated across the frequency spectrum of the images. One way to obtain this information is by estimating the noise power spectrum (NPS), which describes the frequency-dependent character of noise [4, p. 86]. In particular, the NPS can be estimated and compared for inputs and outputs of the system F^{-1}. A short but comprehensive description of the steps involved in estimating the NPS is given in Appx. E. For the inputs to F^{-1}, we can simulate a pair of independent (noisy) flat-field measurements for each of the energies used in this example, i.e. 15 and 20 keV, and divide (or flat-field correct) one by the other. This, in turn, produces a dataset that captures the noise in the flat-field corrected inputs to F^{-1}. Fig. 9 shows the NPS of the simulated input measurements and the corresponding retrieved properties. Here the negative log of the simulated measurements was used during the estimation of the NPS to ensure that the expected values of the inputs and retrieved quantities were the same, that is zero.

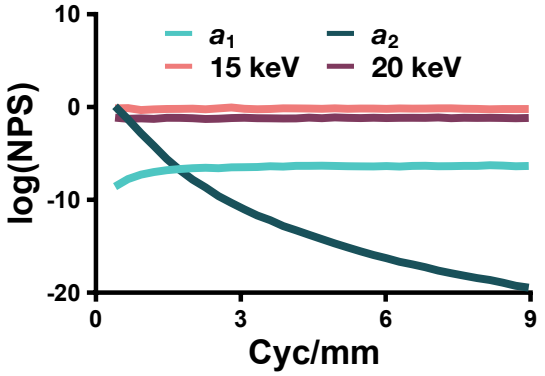

Figure 9: Noise power spectrum of the simulated (noisy) measurements and two attenuation properties related to the retrieved quantities (a_1 and a_2).

To draw some meaningful conclusions from the result in Fig 9, one can compare the NPS of the input measurements to those of the retrieved properties. One can also compare the NPS of the two properties with each other. To start, note that the NPS values are highest for the 15 keV

measurement. This trend reflects the properties of the spectrum used to simulate the measurements, for which counts at 20 keV were higher than those at 15 keV and led to "higher" noise (and noise power) for the latter. Next, except for the values at the lowest of frequencies shown, the noise power in both of the retrieved properties is significantly lower than that of the inputs. Additionally, the noise of the input images is uncorrelated (or "flat") while the same is not true for a_1 and a_2. By contrast, the noise in a_2 is correlated [4, p. 87] while that of a_1 shows a slight increase at low frequencies preceding a nearly flat profile.

The observed differences in the noise texture between the retrieved terms in Fig. 8 are readily explained by contrasting their respective NPS. Namely, the NPS of a_2 shows a significant and rapid reduction in noise power with increasing frequency and, thus, resembles what would be seen after a smoothing operation. The results of a_1 are strikingly different. The noise power of the a_1 term seems to increase, although slightly, with frequency. As a consequence, high-frequency content, where SNR is typically low, can be preferentially amplified relative to the lower frequency band. A similar result is achieved by sharpening (the opposite of smoothing) operations. The apparent smoothing of a_2 and a sharpening of a_1 will be addressed in more detail in later sections.

The strategy used in this section is useful for more than only understanding how noise affects phase retrieval. Two other important sources of error that can be easily studied with this approach are the influence of mischaracterizing the effective energy of the measurements and ignoring the blurring induced by the source and detector. Insight into these two problems is also given at the end of this section.

3.3 Analytic frequency modulation functions of the retrieved quantities

Another powerful method to understand the frequency response of a system involves obtaining and interpreting its response to an impulse. The result after such an analysis is sometimes called the point spread (or impulse response) function (PSF) of the system and paints a clear picture of how the system modifies the information in all resolvable frequencies. This section explores a way to obtain a pair of analytical functions analogous to the PSF of the spectral phase retrieval technique.

The main objectives are to obtain a suitable description of the frequency response of a_1 and a_2 and to relate the observations to material properties and experimental variables. By understanding how undesirable behaviors emerge, we can propose strategies to potentially mitigate them.

To keep the notation concise, lets rewrite the Fourier-domain version of the forward model, that is Eq. 37, as

$$\tilde{D}_{E_i}(\boldsymbol{\nu}) = f_{1,E_i}\tilde{A}_1(\boldsymbol{\nu}) + f_{2,E_i}(\boldsymbol{\nu})\tilde{A}_2(\boldsymbol{\nu}), \tag{39}$$

where the notation $\tilde{G}(\boldsymbol{\nu})[= \mathcal{F}\{g(\boldsymbol{r}_\perp)\}]$ represents the FT of any arbitrary function $g(\boldsymbol{r}_\perp)$ while $\boldsymbol{\nu}[= (\nu_x, \nu_y)]$ are the Fourier-domain coordinates conjugate to $\boldsymbol{r}_\perp[= (x, y)]$. The function f_{1,E_i} represents the energy dependence of PE absorption, i.e. E_i^{-n} with $n \sim 3$, and f_{2,E_i} is defined as

$$f_{2,E_i} = \Phi_{E_i} + \left(R_2 \frac{h^2c^2}{2\pi E_i^2}\right)|\boldsymbol{\nu}|^2 = \Phi_{E_i} + \Gamma_{E_i}|\boldsymbol{\nu}|^2. \tag{40}$$

In Eq. 39 and 40, the energy dependence of each variable is indicated with a subscript, e.g $\tilde{D}(\boldsymbol{\nu}; E_i) = \tilde{D}_{E_i}(\boldsymbol{\nu})$. The magnification is assumed to be unity, $M = 1$, for simplicity.

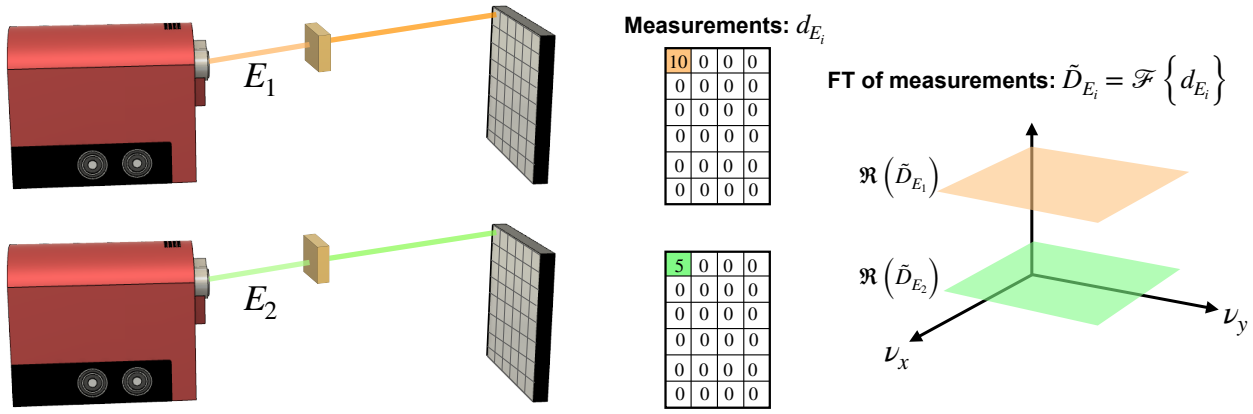

Figure 10: Depiction of the hypothetical experiment used to learn about the frequency response of the retrieved quantities, a_1 and a_2. A sample is measured twice at different X-ray energies ensuring that the radiation is collected by a single pixel. This results in two sparse measurements, d_{E_i} with constant Fourier representations.

To understand how each term behaves in frequency domain, one can consider how the algorithm reacts to a pair of spatial impulses. These types of signals have constant Fourier-domain representations. An experimental realization of this concept is depicted in Fig. 10. Namely, a sample

illuminated by two X-ray beams having energies E_1 and E_2 is considered. The transmitted X-rays are collected by a single pixel, resulting in a pair of sparse measurements like the ones shown in Fig. 10. The FT of the measurements are purely real and constant in frequency ($\boldsymbol{\nu}$) allowing:

$$\tilde{D}_{E_1} = f_{1,E_1}\tilde{A}_1(\boldsymbol{\nu}) + f_{2,E_1}(\boldsymbol{\nu})\tilde{A}_2(\boldsymbol{\nu}) \tag{41a}$$

$$\tilde{D}_{E_2} = f_{1,E_2}\tilde{A}_1(\boldsymbol{\nu}) + f_{2,E_2}(\boldsymbol{\nu})\tilde{A}_2(\boldsymbol{\nu}), \tag{41b}$$

where the values of $\{\tilde{D}_{E_i} : i = 1, 2\}$ are constant. Solving for $\tilde{A}_2(\boldsymbol{\nu})$ we get

$$\begin{aligned}\tilde{A}_2(\boldsymbol{\nu}) &= \frac{\tilde{D}_{E_1} f_{1,E_2} - \tilde{D}_{E_2} f_{1,E_1}}{f_{1,E_2} f_{2,E_1}(\boldsymbol{\nu}) - f_{1,E_1} f_{2,E_2}(\boldsymbol{\nu})} \\ &= \frac{\tilde{D}_{E_1} f_{1,E_2} - \tilde{D}_{E_2} f_{1,E_1}}{(f_{1,E_2}\Phi_{E_1} - f_{1,E_1}\Phi_{E_2}) + (f_{1,E_2}\Gamma_{E_1} - f_{1,E_1}\Gamma_{E_2})\,|\boldsymbol{\nu}|^2},\end{aligned} \tag{42}$$

where the definition in Eq. 40 is used for the second expression. With the exception of the last term in the denominator, every term in Eq. 42 is constant in $\boldsymbol{\nu}$. Furthermore, if the energies are defined so that $E_1 < E_2$, it can be shown that the numerator as well as the two expressions enclosed in parenthesis in the denominator are all negative. It follows that the shape of the amplitude spectrum of $\tilde{A}_2(\boldsymbol{\nu})$, i.e. $\left|\tilde{A}_2(\boldsymbol{\nu})\right|$, will resemble the transfer function of a smoothing (or low-pass) filter. This agrees with the observations made in Sec 3.2. This result still holds when $E_1 > E_2$.

After substituting Eq. 42 into Eq. 41a and rearranging the result, we get that:

$$\tilde{A}_1(\boldsymbol{\nu}) = \frac{\tilde{D}_{E_1}}{f_{1,E_1}} - \frac{\tilde{D}_{E_1} f_{1,E_2} - \tilde{D}_{E_2} f_{1,E_1}}{f_{1,E_1} f_{1,E_2} - f_{1,E_1}^2\,[f_{2,E_2}(\boldsymbol{\nu})/f_{2,E_1}(\boldsymbol{\nu})]}. \tag{43}$$

As was the case for $\tilde{A}_2$, in Eq. 43 almost every term is constant in $\boldsymbol{\nu}$, with the only exception being the ratio inside the square brackets. Unlike $\tilde{A}_2$, however, the dependence on $\boldsymbol{\nu}$ is more complicated. Yet, the shape of $\left|\tilde{A}_1(\boldsymbol{\nu})\right|$ can be studied by specifying a pair of energy values and modeling the absolute value of Eq. 43. Fig. 11 shows the amplitude response of the two terms modelled for a variety of energy pairs. The material properties used to define $\tilde{D}_{E_i}$ were those of polymethyl

methacrylate (PMMA). Since the material was assumed to be a uniform slab, the only property needed was the linear attenuation coefficient. In addition, the radial symmetry of equations 42 and 43 in frequency-domain allow their representation with a single axial profile.

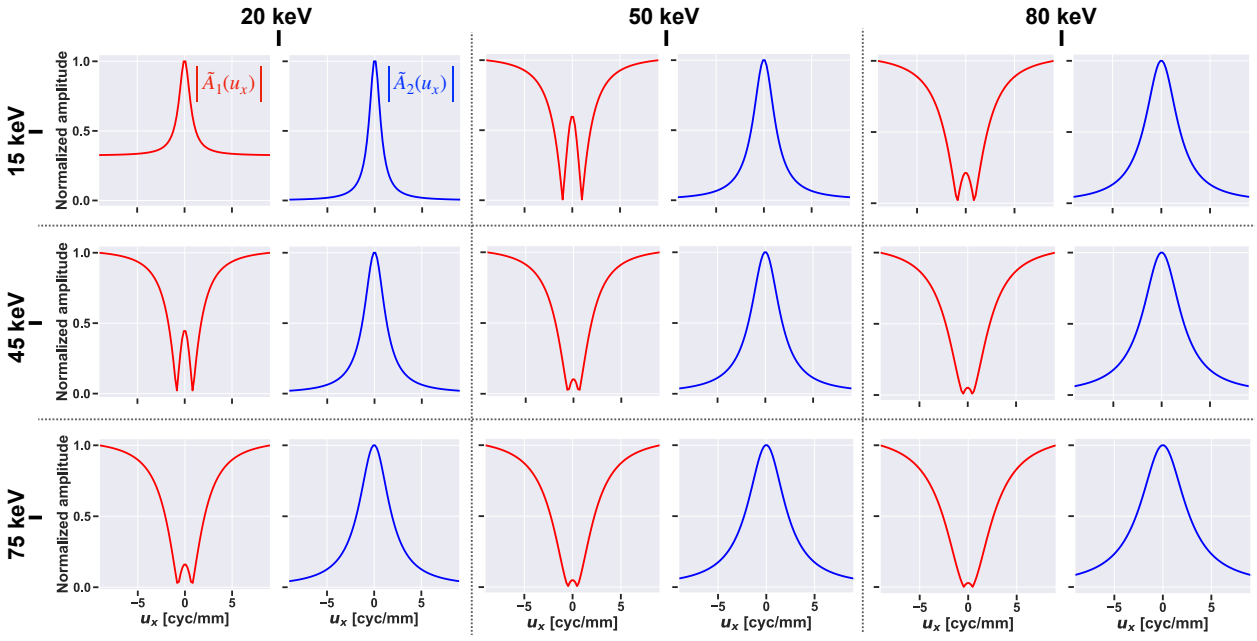

Figure 11: Normalized amplitude response of the retrieved terms, i.e. $|\tilde{A}_1|/\max|\tilde{A}_1|$ and $|\tilde{A}_2|/\max|\tilde{A}_2|$, for different energy pairs. The energy pairs used to analyze Eq. 43 and 42 are listed on the perimeter.

Fig. 11 reveals some important details about the frequency response of each term. First, we see that the amplitude response of $\tilde{A}_2$ resembles a smoothing kernel for all energy combinations as expected. Additionally, we see that the width of the profile broadens as either (or both) of the energies increase. This results in more high-frequency content being accepted by this term, which can influence the resulting image quality.

As expected from the comparison between the Eq. 43 and Eq. 42, the behavior of the $\tilde{A}_1$ term is more complex. The NPS analysis in Sec 3.2 gave some evidence to the amplification of power with increasing frequency, which resembles the effect of a sharpening operation. For the chosen energy pairs here, this is the dominant behavior and it becomes more prominent with increasing energy. An exception occurs when the energies are 15 and 20 keV, in which both terms favor lower frequencies. However, the amplitude response of the $\tilde{A}_1$ term does not roll down to zero as would

be the case for a high-cut or low-pass filter.

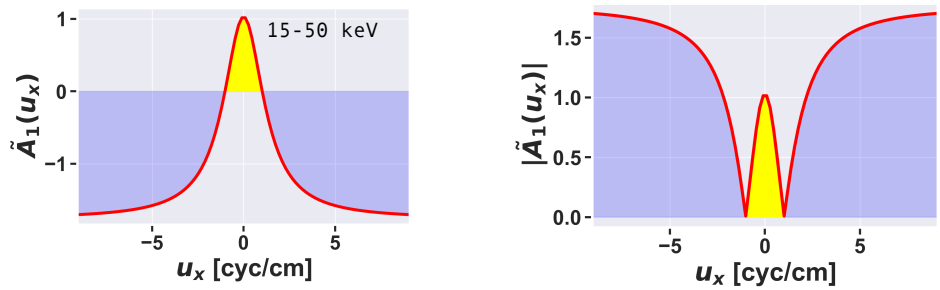

Figure 12: Real part (left) and amplitude response (right) of the FT of the photoelectric absorption term $\tilde{A}_1$ for the 15-50 keV energy pair.

Considering Fig. 12, the amplification of amplitudes with increasing frequencies in the $\tilde{A}_1$ term can be explained by revisiting Eq. 43. In particular, the $\tilde{A}_1$ is allowed to become negative. Since the amplitude response is given as the modulus of the Fourier transform, it is strictly positive and displays a switch in the behavior, from "smoothing" to "sharpening", at the zero-crossing points of $\tilde{A}_1$. Fig. 12 illustrates this for the case of the 15-50 keV pair in Fig. 11. Namely, $\tilde{A}_1$ becomes negative after about 1 cyc/mm where the amplitude spectrum $|\tilde{A}_1|$ shows a kink. The sign reversal of the negative branches results in amplification of the high-frequency content.

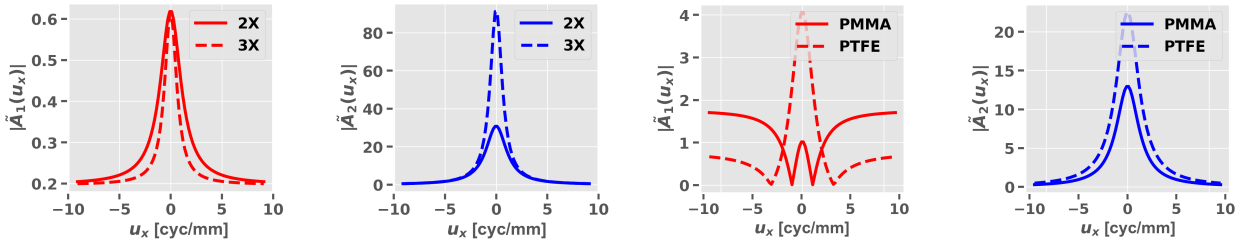

(a) 15 and 20 keV **(b)** 15 and 50 keV

Figure 13: (a) Amplitude response of the $\tilde{A}_1$ term using the energies 15-and-20 keV with two-times (2X) and three-times (3X) geometrical magnification. (b) Amplitude response of $\tilde{A}_1$ for the energies of 15-and-50 keV and two different materials, PMMA and PTFE.

The method in this section also provides a way to explore the influence of different experimental and material properties in the frequency response of the retrieved terms. This is achieved by modifying the f_{2,E_i} terms to reflect the desired parameters and, subsequently, estimating the frequency response of $\tilde{A}_1$ and $\tilde{A}_2$. Fig. 13a shows the effect of varying the geometrical magnification, M, from

two-times (2X) to three-times (3X) for the 15-20 keV energy pair. The reason for selecting this pair of inputs was the similar tapering behavior of the retrieved terms. For $\tilde{A}_1$, the 3X case shows stronger attenuation of higher frequencies (greater smoothing). On the other hand, the $\tilde{A}_2$ term has a comparable profile width for both magnifications but the 3X case has a higher amplification at low frequencies compared to the 2X case.

Fig. 13b shows the effect of using materials with different LACs for the 15-50 keV energy pair. For this example, comparing the changes going from PMMA to polytetrafluoroethylene (PTFE) gives insight into the effect of increasing the LAC. Increasing the attenuation of the material results in a slight broadening of the $\tilde{A}_2$ term and strengthens its low-frequency response. Using the common definition of the cut-off of a low-pass filter, i.e. the frequency corresponding to 50% of the maximum gain, it appears increasing the attenuation leads to a lower "cut-off frequency" for $|\tilde{A}_2|$. The $\tilde{A}_1$ term of both materials show a significant difference. Not only is the amplitude response approximately four times higher at low frequencies, but the overall behavior may also be better described as smoothing rather than sharpening going from PMMA to PTFE. This change can lead to noticeable differences in the quality of the resulting images.

3.4 Noise mitigation strategies

From the results in the previous sections, it is evident that many of the undesirable features in the retrieved images of the photoelectric (or the related a_1) term are connected to how the high-frequency content of the input is modified. In particular, an inherent increase in amplitudes with increasing frequency was observed. For this reason, two justifiable approaches to improve the quality of the retrieved images are filtering some of the high-frequency content and denoising the input. Filtering by "rolling-off" high-frequencies counteracts the intrinsic behavior of the a_1 term. On the other hand, denoised input images will have higher SNR across all frequencies and, therefore, should yield retrieved images that are less affected by the amplification of high-frequencies.

The following sections present our current methods to implement each of these two noise mitigation techniques. However, we acknowledge that a myriad of alternatives, some of which may lead

to better results, exist. A third approach, involving an unintended smoothing of the output caused by scaling the known terms in Eq. 36, will also be discussed. This last method currently lacks a solid theoretical foundation, is prone to artifacts, and it might be a byproduct of the numerical solver. For these reasons, this method is proposed as a way to obtain a glimpse of what a more proper noise mitigation strategy can produce.

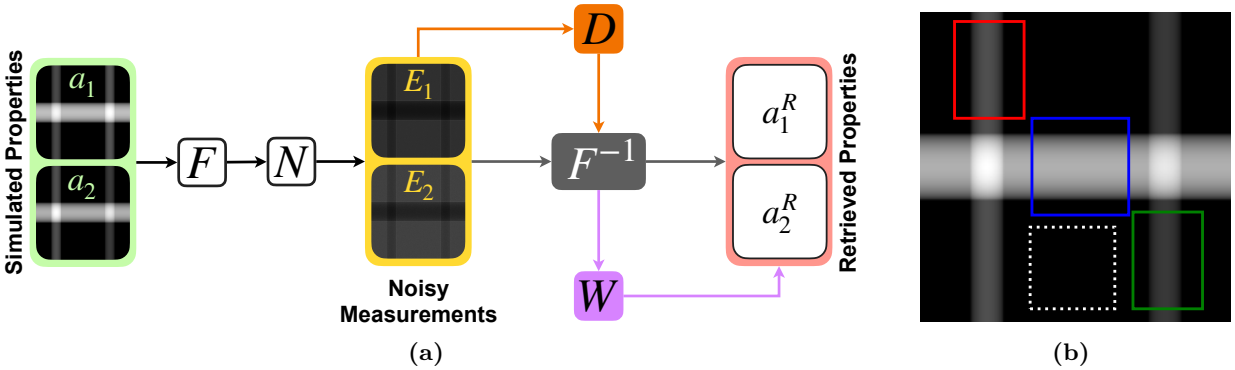

Figure 14: (a) Workflow showing the approach used to investigate the three noise mitigation strategies. Here W represents an apodization strategy using a Tukey windowing kernel, D represents a spatial denoising step preceding phase retrieval, and scaling the known terms (feature scaling) is represented by the phase retrieval operator F^{-1} itself. The last choice illustrates that feature scaling is only a modification of how F^{-1} is implemented and does not amount to an added step. (b) Shows the regions used to estimate the percent error and SNR of the resulting images. The red, blue, and green boxes enclose the water, POM, and PMMA rods, respectively. The dashed white box shows the region used to estimate the standard deviation needed to compute the SNR.

The performance of the three noise mitigation strategies was analyzed quantitatively by estimating the average percent error and SNR of the projected photoelectric (PE) and Compton scattering (CS) maps found after phase retrieval. As described in Sec. 3.1, these two quantities are found from the material properties a_1 and a_2, respectively. In our analysis, we use the simulation framework described in Sec. 3.2 with modifications to include the denoising strategies. Fig. 14 shows the resulting workflow. Only two measurements were simulated (15 and 20 keV). Like before, the operator F represents the forward simulation step, which operates on the numerical phantom with simulated material properties a_1 and a_2, i.e. $P(a_1, a_2)$. The operation $N[F[P(a_1, a_2)]]$ introduces Poisson noise into the noise-free simulated measurements, $F[P(a_1, a_2)]$, and F^{-1} operates on the

noisy measurements to recover the material properties. The last step represents phase retrieval and gives the estimated properties a_1^R and a_2^R so that:

$$P(a_1^R, a_2^R) = F^{-1}[N[F[P(a_1, a_2)]]]. \tag{44}$$

The aim is then to compare $P(a_1^R, a_2^R)$ with $P(a_1, a_2)$, or related material properties, using some metric to assess how well the influence N was mitigated.

The percent error used in this comparison, for the case of Compton scattering, can be given by:

$$\epsilon_{\text{cs}}(x^*, y^*) = \frac{\left|\sigma_{a_2^R}(x^*, y^*) - \sigma_{a_2}(x^*, y^*)\right|}{\sigma_{a_2}(x^*, y^*)} \cdot 100\%, \tag{45}$$

where $\sigma_{a_2^R}$ and σ_{a_2} are the retrieved and noise-free CS maps, respectively, while the superscript "*" indicates coordinates within each region-of-interest (ROI) where the material properties are not zero. A similar equation can be obtained for the PE term. Fig. 14 shows the ROIs used for the water (red), POM (blue), and PMMA (green) rods in our phantom. The SNR was calculated using the differential signal (peak-to-background) of each rod divided by the standard deviation of the ROI enclosed by the dashed white box.

For the simulated measurements, we used a pixel pitch of 55 microns, a 512×512 pixel array, and a monochromatic source. The two vertical rods (water and PMMA) were modeled with a diameter of 1 mm while the horizontal (POM) rod had a diameter of 2 mm. A geometrical magnification factor of three was used so that the dimensions of the rods at the detector plane correspond to either 3 or 6 mm. One benefit of working with rods is that their projections resemble the tapered lobe of a sinusoidal function and, thus, represent half of a cycle of a sinusoid. This perspective allows one to simplify the interpretation of results in frequency-domain, which can help inform the choices of cut-off frequencies.

3.4.1 Apodization: frequency-domain filtering

The removal of high-frequency content (or apodization) is a common practice in medical imaging that can mitigate artifacts and improve the signal-to-noise ratio [57]. Apodization is generally implemented as a frequency filter. Consequently, this technique is most effective when the frequency-domain representation of a measurement displays a clear separation between the "mostly-signal" and "mostly-noise" regions. This, however, is rarely the case. This section will discuss the use of a radial cosine-tapered, or Tukey windowing kernel, applied in frequency-domain following phase-retrieval. Since a virtually unlimited variety of filters with varying levels of sophistication exist [58], a more optimal apodization strategy is possible to find. Nevertheless, this relatively simple method is suitable for medical imaging [57].

The radial Tukey windowing kernel in frequency- or $\boldsymbol{\nu}$-domain is defined as:

$$\mathcal{W}(\boldsymbol{\nu}) = \begin{cases} 1 & |\boldsymbol{\nu}| < \nu_c \\ \cos^2\left(\dfrac{\left(\pi\left(|\boldsymbol{\nu}| - \nu_c\right)\right)}{2w}\right) & \nu_c \leq |\boldsymbol{\nu}| < \nu_c + w \\ 0 & \nu_c + w \leq |\boldsymbol{\nu}| \end{cases} \tag{46}$$

where $|\cdot|$ is the modulus operator and ν_c is a user-defined radial cut-off frequency. A second user-defined parameter is the transition-band width, w, which connects the pass-band region, where $\mathcal{W}(\boldsymbol{\nu}) = 1$, with the stop-band region, where $\mathcal{W}(\boldsymbol{\nu}) = 0$. The choices of ν_c and w regulate the trade-off between unwanted features (like noise and artifacts) and the resolution in the treated images.

Fig. 15 shows the results of applying Eq. 46 to the retrieved PE and CS maps. A constant transition width of approximately 0.25 cyc/mm was used along with four different cut-off frequencies of 8.2, 1.6, 0.5, and 0.2 cyc/mm. The kernels estimated for these choices of w and ν_c are shown on the last row of Fig. 15. Since the rods are either 3 or 6 mm in diameter at the detector plane, one can expect the horizontal and vertical rods to be close approximations of a sinusoid oscillating 1/6 and 1/12 of a cycle every 1 mm. Therefore, frequencies far greater than approximately 0.2 cycles

per mm will have negligible contributions to the information in the phase retrieved images. This observation is corroborated by the last column in Fig. 15, where the expected shape of the material properties is obvious.

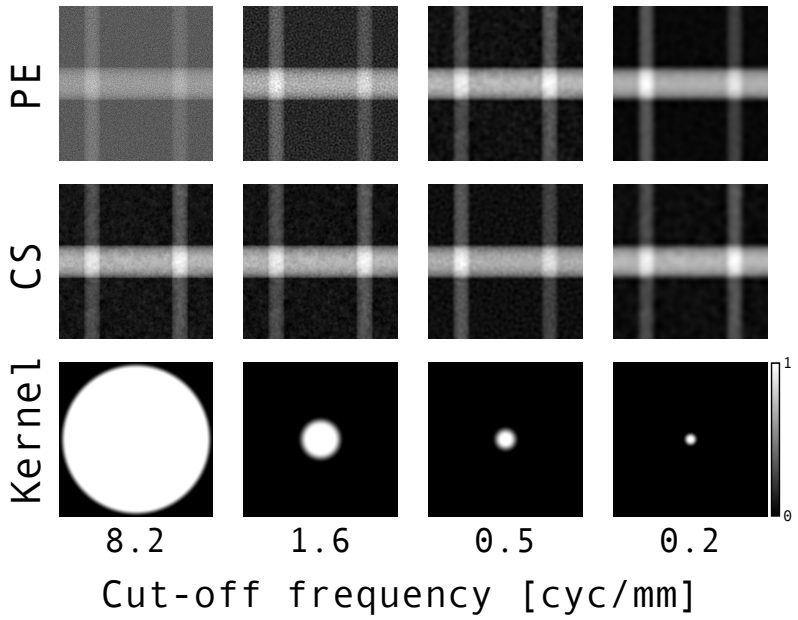

Figure 15: Shows the effect of increasing the amount of high frequency content that is removed by the Tukey windowing kernel is shown. As more high frequency content is removed (going from left to right), the visual appearance of the image seems to improve. This, however, comes at the price of image resolution loss.

Table 1 summarizes the average percent error and SNR values for each of the materials and seleted cut-off frequencies. By decreasing ν_c from 8.2 down to 0.2 cyc/mm, the SNR increased by nearly 38 folds for PE, but only by two-folds for CS in all materials. The average percent error did not show the same consistency across the materials. For the case of CS, only marginal gains are shown and the cut-off frequency decreases from 0.5 to 0.2 cyc/mm for POM. The improvements are more significant for PE, where the error improves by factors of 12, 9, and 16 for the water, POM, and PMMA rods, respectively.

This section concludes with noting that although the results in Table 1 are encouraging, the corresponding apodization scheme depends on the judicious choice of w and ν_c. In particular, selecting a somewhat unbiased value of ν_c, i.e. a frequency near or at the Nyquist-frequency like

		Water		POM		PMMA	
ν_c	Property	% Error	SNR	% Error	SNR	% Error	SNR
8.2	CS	11.7	10.5	4.8	14.8	9.8	12.0
	PE	80.4	1.4	39.6	1.7	113.7	1.0
1.6	CS	11.3	11.2	4.5	15.7	9.5	12.8
	PE	16.2	6.6	8.4	7.9	23.3	4.7
0.5	CS	9.5	14.2	4.2	20.0	7.7	16.2
	PE	7.3	18.0	3.9	21.7	9.9	12.9
0.2	CS	8.3	23.0	4.4	32.2	7.7	26.2
	PE	6.9	53.5	4.3	63.1	7.1	38.6

Table 1: Average percent error and SNR for the quantitative assessment of the apodization scheme using different choices of cut-off frequency, ν_c.

8.2 cyc/mm in this case, may not lead to satisfactory results. Instead, the best results were obtained by placing the cut-off close to where we expected most of the relevant content to end. Nevertheless, apodization seems to provide a valuable way to improve the quality and accuracy of the results.

3.4.2 Spatial denoising: block-matching and three-dimensional filtering

Section 3.3 described how the PR scheme preferentially amplifies the high-frequency content of input measurements when estimating a_1. The high-frequency regions of the Fourier-representation measurements typically suffer from low SNR, and amplifying these regions can lead to poor image quality. Applying spatial denoising on the measurements can mitigate this effect by increasing the SNR across all frequencies. This strategy is shown in Fig. 14 as the path over D.

An ideal denoising strategy is one that operates on noisy data to ensure that $P(a_1^R, a_2^R)$ is exactly equal to $P(a_1, a_2)$. Thus, using the operator notation from Sec. 3.2, we expect that

$$F^{-1}[D[N[F[P(a_1, a_2)]]]] = P(a_1^R, a_2^R) = P(a_1, a_2) \tag{47}$$

where the last equality is true if and only if $D[N[\cdot]]$ equals the identity operator. The denoising operation chosen for this section is the block-matching and three-dimensional filtering (BM3D) [59,60]. This decision was motivated by a recent and encouraging applications of BM3D in medical

imaging [61, 62] and by the impressive edge-preserving characteristics of this filter. The latter is highlighted by Fig. 16 and is a key requirement for PB-PCI, in which edge enhancement is often cited as a key benefit.

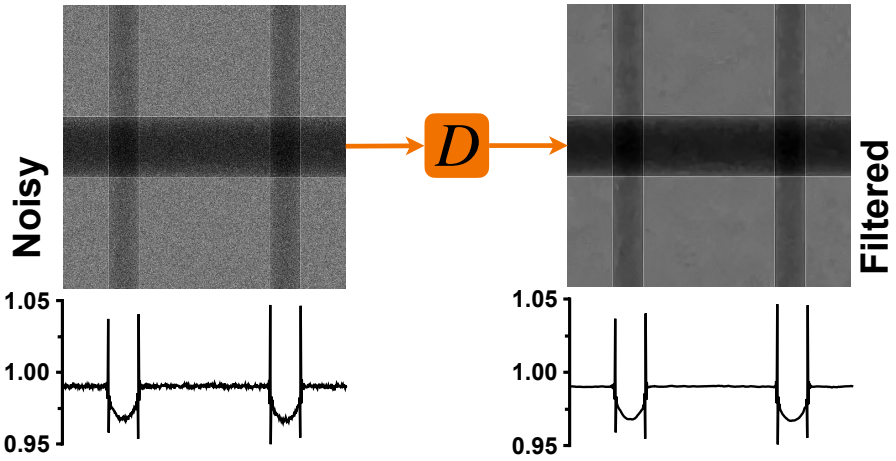

Figure 16: Outcome of applying the BM3D filtering scheme, represented by D, to a noisy PCI measurement. Comparing the average profiles below the noisy and filtered images shows how this routine reduces noise without substantial damage to sharp edges.

The BM3D algorithm can be broadly described in five steps. The following description will consider how BM3D works at a particular "reference" location (or patch). First, the entire image is denoised, which typically involves a Wiener-denoising strategy. Next, a group of patches similar to the reference patch are found across the image and subsequently stacked to form a 3D group. These patches are selected by ensuring that a measure of dissimilarity, with respect to the reference patch, falls beneath a predefined threshold. Once the patches are found, the next step is to denoise the patches. At this stage, the denoising routine involves a transformation of the 3D group into a domain where noise is assumed to correspond to a sparse set of small coefficients that can be easily removed by thresholding. In other words, a clear distinction between signal and noise is assumed. This step is commonly implemented using wavelet-based hard thresholding. For the fourth step, the patches in the denoised 3D stack are aggregated to produce a 2D patch. The final step involves repeating all previous steps, but a Wiener filter is used to denoise the 3D patches rather than hard thresholding. A detailed description of BM3D is well beyond the scope of this work, but can be found in [59, 60].

Fig. 17 shows the PE and CS terms obtained by treating the noisy inputs to the PR algorithm with a freely available BM3D Python module. The module was created, and is actively maintained, by the group that developed the original BM3D algorithm. The only input required to use the BM3D filter is an estimate of the standard deviation of the noise (nSTD) in the images. Yet, a motivated user can modify the algorithm at its multiple steps to optimize the performance for a given task. For our measurements, the standard deviation of the noise was found to be between 0.015 and 0.020.

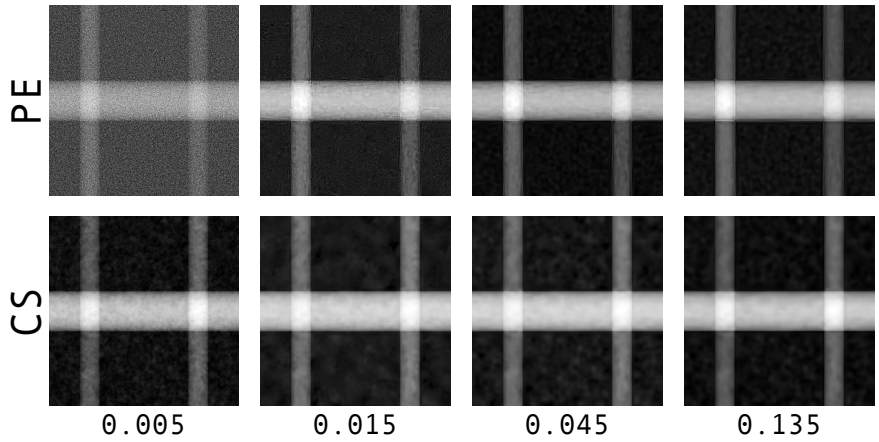

Figure 17: Results of the PE and CS terms when the inputs to PR are filtered by the BM3D scheme using four different values for the standard deviation of noise.

A qualitative analysis of Fig. 17 shows that although the CS term is not noticeably altered by applying the BM3D filter to the input data, the PE term is greatly improved. This is corroborated by the results in Table 2. Interestingly, the SNR of the CS term was highest when the nSTD was closest to that of the input images. The SNR value then decreases, but only by a few units. The error of both properties, however, appears to benefit by the increasing nSTD. The PE shows a different SNR trend. The SNR value seems to nearly saturate at around the value obtained when the nSTD was 0.045. For this term, the error at the maximum nSTD is 9 to 12 times smaller than at 0.005.

It should be noted that a different and perhaps equally valid approach is to implement BM3D after, rather than before, PR. This workflow could lead to better results. Another important point

		Water		POM		PMMA	
nSTD	Property	% Error	SNR	% Error	SNR	% Error	SNR
0.005	CS	11.5	10.9	4.7	15.3	9.6	12.4
	PE	93.7	1.2	45.7	1.5	130.3	0.9
0.015	CS	8.3	19.5	3.6	27.7	6.9	22.3
	PE	18.9	6.1	7.9	7.3	30.0	4.3
0.045	CS	7.4	15.3	3.3	21.7	6.3	17.5
	PE	11.6	22.8	4.3	27.5	16.4	16.3
0.135	CS	8.2	15.4	3.4	22.1	5.9	17.7
	PE	8.8	23.2	4.9	28.1	10.2	16.7

Table 2: Average percent error and SNR for the quantitative assessment of spatial denoising the inputs to the phase retrieval algorithm using a BM3D filter. The first column lists the different standard deviation values used, which the filtering algorithm interprets as the standard deviation of the noise in the input images.

is that the Python-based BM3D module assumes that the images noise is additive and spatially correlated stationary Gaussian noise [60]. As discussed in Sec. 1.1, the noise in X-ray imaging (especially in acquisitions with low exposure times from laboratory sources) can display statistical properties that are significantly different than the ones assumed by this filter. For this reason, we present this data as preliminary results and plan to explore the usefulness of this capable filter in the near future.

3.4.3 Smoothing by feature scaling

The final noise mitigation strategy is arguably the simplest to use and it can undoubtedly improve the quality and accuracy of the retrieved images. This technique is applied by changing the scaling between the energy-dependent functions of the a_1 and a_2 terms in Eq. 36. The new scaling must be accounted for by modifying the unknown terms (a_1 and a_2) to ensure that the mathematical identity of the original model is preserved. This strategy can be demonstrated with the notation used in Eq. 39, with $M = 1$ for simplicity, such that:

$$\tilde{D}_{E_i}(\boldsymbol{\nu}) = \frac{\gamma}{\gamma} f_{1,E_i} \tilde{A}_1(\boldsymbol{\nu}) + \frac{N_a}{N_a} f_{2,E_i}(\boldsymbol{\nu}) \tilde{A}_2(\boldsymbol{\nu}) = \gamma f_{1,E_i} \tilde{A}'_1(\boldsymbol{\nu}) + N_a f_{2,E_i}(\boldsymbol{\nu}) \tilde{A}'_2(\boldsymbol{\nu}). \tag{48}$$

Here N_a is Avogadro's number and γ is an arbitrary constant while the unknowns are redefined so that $\tilde{A}_1' = \tilde{A}_1/\gamma$ and $\tilde{A}_2' = \tilde{A}_2/N_a$. Furthermore, the energy is assumed to be specified in keV, which dictates the scales of the f_{1,E_i} and f_{2,E_i} terms. As Eq. 48 shows, this scheme is mathematically equivalent to multiplying by one. Yet, notable improvements in quality can be obtained by regulating the value of γ. The formulas do not justify the outcomes, which may be a byproduct of our choice of numerical solver.

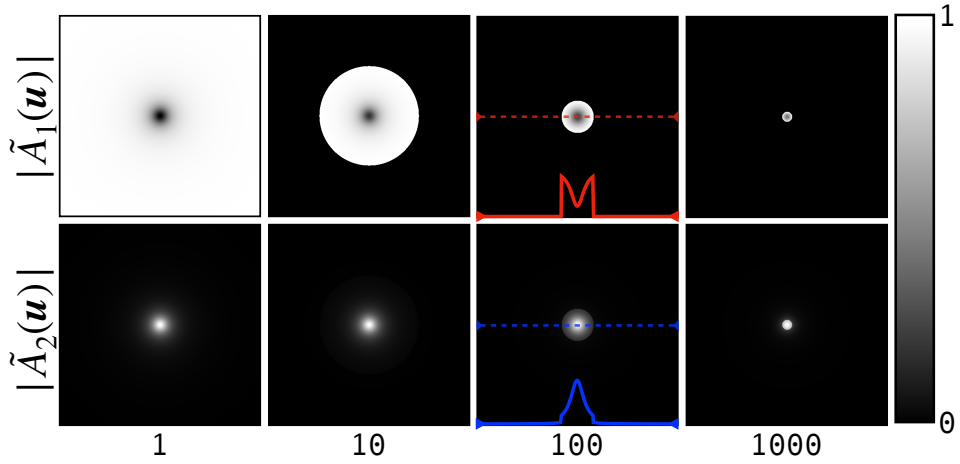

Figure 18: Increasing the scaling term introduces a sharp radial cut-off in the frequency response of the estimated terms. When the scale is set to unit y, the characteristic behavior of the two terms described in Sec. 3.3.The electron density rolls smoothly to zero, while the PE term increases with frequency. A line profile for the scale of 100 shows the sharp transition that results from feature scaling, which can lead to ringing artifacts and deterioration of the measurement quality.

Using an approach similar to the one in Sec 3.3, the results of this study revealed how changing the value of γ can produce improved results. Two artificial measurements of a single illuminated pixel were used as inputs, and the results provide insight into how the PR algorithm modulates frequencies. Fig. 18 shows the outcome of four different choices of γ: 1, 10, 100, and 1000. If no scaling is applied, i.e. $\gamma = 1$, then the amplitude spectrum of a_1 and a_2 agree with our previous finding. Namely, $|\tilde{A}_1|$ resembles a sharpening kernel, while $|\tilde{A}_2|$ resembles a smoothing kernel. Furthermore, as γ increases, the amplitude spectrum of both terms appears to multiply with a narrowing binary mask that eliminates high frequency content. In other words, values inside some radius in Fourier-domain were left unchanged (multiplied by one), while those outside were

completely removed (multiplied by zero).

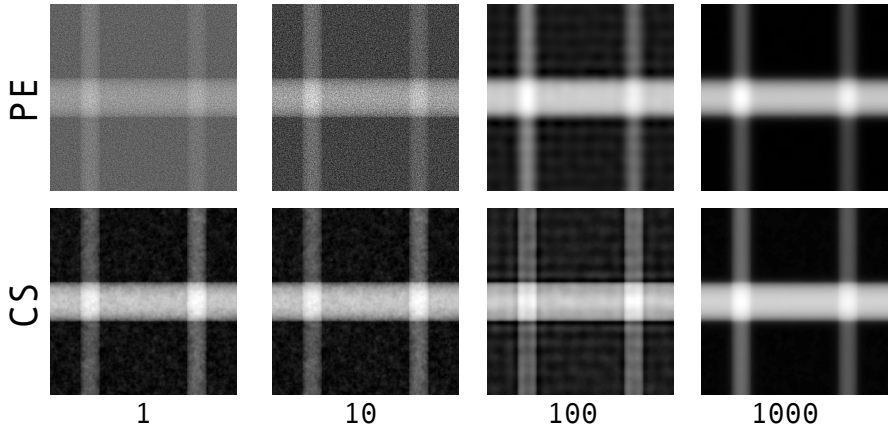

Figure 19: The effect of different scaling terms for simulated data. The last term is smooth and appears free of artifacts, but is highly inaccurate.

This approach can be used to retain relevant portions of Fourier-domain while avoiding "mostly-noise" regions because it resembles a radial windowing. This, in turn, can produce accurate solutions with improved quality. However, the frequency content, as shown in Fig. 18, is truncated abruptly, which is known to cause ringing artifacts [63]. Fig. 19 shows the estimated CS and PE terms for the same four values of γ used in Fig. 18. Note that when $\gamma = 100$, visible ringing appears in the two terms. Although the images corresponding to $\gamma = 1000$ do not display obvious ringing artifacts, the material properties there are highly inaccurate as shown in Table 3.

Table 3 lists the SNR and average percent errors from the analysis of this noise mitigation technique. From $\gamma = 1$ to $\gamma = 100$, both the SNR and the error improve significantly for the PE term. In particular, the SNR was between 27 and 30 times larger, while the error was 9 to 11 times smaller in all cases. For $\gamma = 1000$, a drastic increase in SNR is counterbalanced by a similarly sharp increase in the error. This observation suggests that for some γ, the trade-off between accuracy and quality should reach an optimum value.

Unlike the PE term, the error for CS increases above $\gamma = 10$ while the SNR improves only slightly between $\gamma = 1$ and $\gamma = 100$. Nevertheless, the loss in accuracy is not larger than two folds in all cases, which is about five times smaller than the improvements to the PE term. From the results in Fig. 19 and Table 3 we conclude that this method can be effective, but it may suffer from

artifacts and poor accuracy.

		Water		POM		PMMA	
γ	Property	% Error	SNR	% Error	SNR	% Error	SNR
1	CS	11.7	10.5	4.8	14.7	9.8	12.0
	PE	97.6	1.2	48.0	1.4	135.6	0.8
10	CS	11.6	10.6	4.7	14.9	9.7	12.1
	PE	39.4	2.7	20.0	3.3	56.6	2.0
100	CS	20.0	10.7	7.2	18.2	13.8	13.1
	PE	11.4	34.0	5.2	37.5	11.8	24.5
1000	CS	446.6	200.9	384.0	262.3	283.0	160.9
	PE	100.0	342.9	100.0	481.7	100.0	275.4

Table 3: Average percent error and SNR to assess the effect of feature scaling. For these results, four values of γ during phase retrieval were used (Eq. 48). The scale of the energy-functions was determine by the choice of units, namely keV.

3.4.4 Discussion

Since the Tukey kernel of Sec. 3.4.1 and the mitigation scheme in this section work in a similar way, there is worth in comparing the two techniques. The former allows for an easier and more conventional way to control the section of frequency-domain to exclude. Additionally, the Tukey kernel transitions smoothly from pass-band to stop-band, and it is therefore less prone to ringing artifacts. For these reasons, we suggest the method in Sec. 3.4.1 for better results, but remark that the technique here is useful to obtain quick improvements. In our experience using Eq. 48, values of γ ranging between 100 and 1000 can produce acceptable results in most cases.

3.5 Errors from misrepresenting the energy

The energy-dependent, but otherwise constant, terms in Eq. 36 must be properly defined to obtain an accurate estimates of the material basis a_1 and a_2. The terms are defined with values that correspond to the energies of the input measurements, which are not strictly monochromatic but instead cover a (narrow) band of energies. In practice, effective energies are estimated for each

measurement using one of several strategies, c.f. Sec. F. Thus, it is natural to wonder how misrepresenting the effective energies of the input bins can influence the accuracy of the results. This issue is explored here using the simulation strategy introduced in Sec. 3.2, see also Sec. C. In particular, the energy information required for PR was purposefully misrepresented relative to the values used to simulate the input measurements.

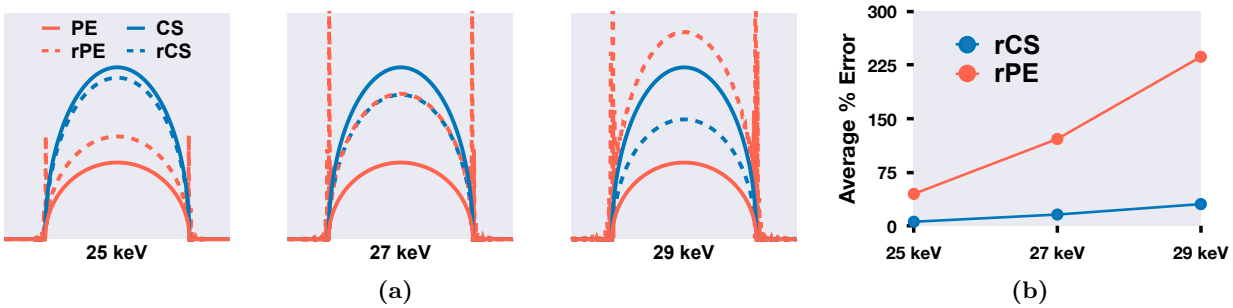

Figure 20: (a)Results of phase retrieval using two (noise-free) monochromatic input measurements simulated with energies of 23.33 and 40 keV. The low-energy measurement (23.33 keV) was assigned an incorrect energy value during phase retrieval. The energy values used for the low energy bin are listed below the resulting PE and CS, which were estimated for an energy of 30 keV, c.f. Eq. 34. Each plot also shows the expected PE and CS and ignores negative values for clarity. (b) Average percent error along the rod profile.

Fig. 20a shows the PE and CS contributions to the LAC estimated for 30 keV after PR. Two monochromatic measurements of a 3 mm PMMA rod with energies of 23.33 and 40 keV were used as inputs. The low-energy measurement was purposefully assigned values of 25, 27, and 29 keV during PR. In every case, the discrepancy between the actual and specified energies led to an underestimation of the CS term and overestimation of the PE term. Also, the magnitude of the average percent error increased monotonically with the error in the specified energy (Fig. 20).

The error in the PE term can be explained by comparing the mathematical form of this term with both the actual and the incorrectly specified energy values. Ideally, when $E_{\text{specified}} = E_{\text{actual}}$ the retrieved (a_1^R) and simulated (a_1) energy-independent properties should satisfy

$$\text{PE} = \frac{a_1}{E_{\text{actual}}^3} = \frac{a_1^R}{E_{\text{specified}}^3}. \tag{49}$$

If instead $E^{3.3}_{\text{specified}} > E^{3.3}_{\text{actual}}$, the PR algorithm preserves the equality by compensating the a_1^R term.

The energy-dependent function of the CS term (i.e. the Klein-Nishina function Φ_E) also decreases with increasing energy. Therefore, based on the logic above one could expect for the CS term to be overcompensated. However, unlike the $E^{-3.3}$ dependence of the PE term, the Klein-Nishina function varies gradually with increasing energy. In particular, the largest percent error in the expected value of Φ_E, which was found as $100 \times (\Phi_{23.3} - \Phi_{29})/\Phi_{23.3}$, is 1.9 % compared to 51.2 % for the energy dependence of the PE term. Furthermore, since the LAC is a combination of PE and CS, an overcompensated PE is likely to be accounted for with a reduction of the CS to minimize the overall error.

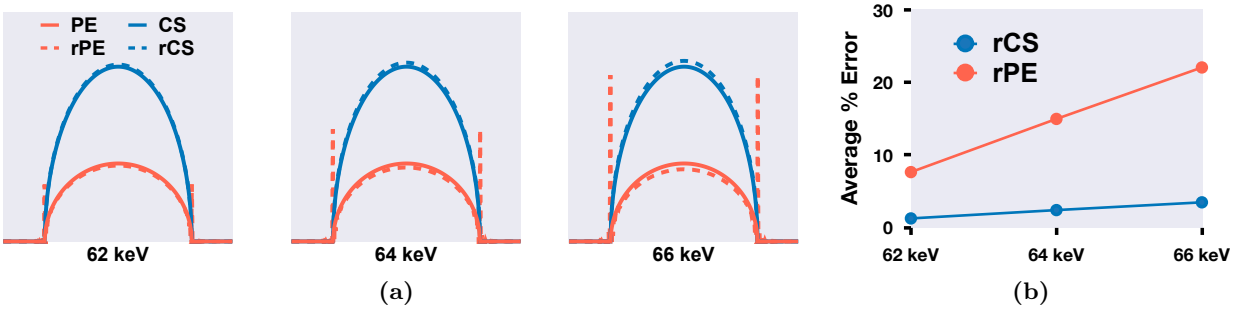

Figure 21: (a) Two monochromatic measurements were simulated with energies of 30 and 60 keV. To induce an error, the 60 keV bin was mislabeled during PR as 62, 64 and 66 keV. (b) Average percent error along the rod profile.

The same study was repeated using two mono-energetic measurements of 30 and 60 keV. The results are shown in Fig. 21. In this case, the high-energy measurement was assigned incorrect values of 62, 64, and 66 keV. Compared to the results in Fig. 20, the errors in Fig. 21 are about 10 times smaller and remain below 25 % in all cases. One reason for this difference is the sensitivity of energy-dependent material properties in the high- and low-energy regimes. For example, the LAC changes rapidly with increasing energy in the low (< 40 keV) regime. On the other hand, changes between two high energy values are small. In terms of the energy functions, those corresponding to the low energy measurement are properly characterized and those for the high energy are close to the actual values since above 60 keV the two functions change gradually. This results in relatively

small errors.

The results in this section highlight the importance of correctly estimating the effective energy of the input measurement. Additionally, errors in the low-energy regime are more detrimental and an effort should be made to address them. This concept should be carefully considered when using photon-counting detectors with high-Z sensor materials like CdTe. Spectral distortions in these detectors, which are more severe at low energies, can yield difficulties in estimating the effective energy of a measurement.

4 Phase Retrieval of Soft Materials at Low Photon Counts

The imaging of poorly-attenuating materials presents an important challenge in clinical and industrial X-ray imaging. Phase contrast imaging, along with phase retrieval, can improve the visibility and identification of these materials. Here we present experimental results for the phase retrieval technique described in Sec. 3.1, which is capable of correctly estimating the photoelectric absorption and Compton scattering due to a sample over a wide range of energies (Fig. 7). Typical methods to identify materials and quantify their abundance, like material decomposition [6, 64, 65], rely on correct estimates of energy-dependent properties. Therefore, these results are evidence that our method can help link phase contrast imaging with these spectral techniques.

The measurements used were obtained in a simple PB-PCI configuration, c.f. Sec. 2.1, using energy resolving photon counting detectors (PCDs). The use of PCDs allows one to obtain the necessary spectroscopic measurements in a single step when a polychromatic laboratory X-ray source is used. Our findings demonstrate the accurate estimate of multiple attenuation mechanisms for biologically relevant materials. In particular, errors in the estimates fall well below 10%. The signal-to-noise ratio of our estimated properties is substantially (around 10 times) greater than that of the corresponding PCI measurements. We highlight how the achieved accuracy in conjunction with improvements in image quality can help identify materials that could otherwise be indistinguishable with conventional techniques. Finally, insight into how reductions in the number of detected photons affects the accuracy of our estimates will be given.

4.1 Materials and methods

4.1.1 Multi-Material Phantom and Biological Sample

A key objective was to investigate the performance of our phase retrieval method with materials of low (< 10) but similar effective atomic numbers (Z_{eff}). To this end, we built a multi-material phantom from four homogeneous rods, each having a different composition, separated by a small

Material	Z_{eff}	ρ_e [$\times 10^{23}$ e/cc]	ρ_m [g/cc]	μ [cm^{-1}]	τ_{pe} [cm^{-1}]	σ_{cs} [cm^{-1}]	$\varnothing$[cm]
PS	5.70	3.40	1.050	0.229	0.0212	0.188	0.296
PMMA	6.47	3.83	1.180	0.277	0.0396	0.211	0.296
POM	6.95	4.52	1.410	0.344	0.0604	0.248	0.318
Graphite	6.00	5.53	1.839	0.382	0.0403	0.304	0.310

Table 4: The properties listed are effective atomic number (Z_{eff}), electron density (ρ_e), mass density (ρ_m), total linear attenuation (μ), photoelectric absorption (τ_{pe}), Compton scattering (σ_{cs}), and rod diameter ($\varnothing$).

(∼1 mm) air gap as shown in Fig. 22. The materials selected were polystyrene (PS), polymethylmethacrylate (PMMA), polyoxymethylene (POM), and graphite. A list of material properties relevant to this study is given in Table 4. Except for the rod diameters and the mass density of graphite, both of which were measured in our laboratory, all of the properties used for comparison were obtained from the NIST database [66].

Although relatively simple, this phantom provides heterogeneity in the horizontal direction and the constituents have characteristics similar to those of biological tissues [67, 68]. Furthermore, all of the Z_{eff} values are well below 10, ranging between 5.70 and 6.95. Materials with these properties are difficult to image and distinguish with absorption-based techniques.

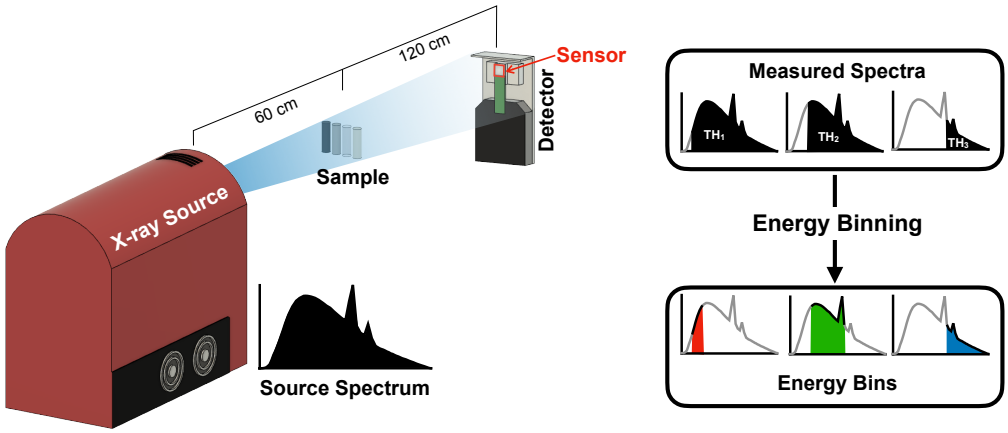

Figure 22: Experimental setup used to image the multi-material phantom. The process of energy binning is also depicted. The Medipix3RX detector captures specific sections of the incident spectrum, and a set of energy bins corresponding to narrow sections of the incident spectrum are formed using energy binning.

In this study of how differences in detected photon-counts influence the quality and accuracy

of our results (c.f. Sec. 4.2.2), the POM rod was used as the sample. To demonstrate the benefits of the proposed phase retrieval algorithm on a biological sample, an Emory oak leaf was imaged. This sample was selected because it contained features in a rich variety of sizes, which allowed for resolving characteristics of our approach in a practical application.

4.1.2 Source and Detectors

The spectroscopic measurements of the rod phantoms were collected with a single-chip Medipix3RX detector. This detector used a 1 mm thick cadmium telluride (CdTe) sensor and a 256×256 pixel array separated at a uniform pitch of 55 μm. The high quantum efficiency of the hybrid CdTe sensor allowed the use of a high source voltage of 90 keV for the multi-material phantom.

The Emory oak leaf in Sec. 4.2.3 was imaged using a single chip Medipix3RX with a 675 μm Si sensor and a pixel pitch is 55 μm. This detector was selected because of its robustness to spectral distortions at the low energies (between 10 and 40 keV) used to image the leaf. All of the measurements for this article were made in *charge summing mode*, which mitigates spectral distortions due to charge sharing [69].

A polychromatic microfocus X-ray tube (Hamamatsu 150 kV L12161-07) was used for the source. In all measurements, the *small focus mode* was selected. This configuration achieves a focal spot of about 7 μm and offers sufficient spatial coherence for PB-PCI.

4.1.3 Spectroscopic Measurements

Spectral data was generated in two steps: threshold (TH) scanning and energy binning (Fig. 22). In the first step, a series of independent scans were collected, each using different low-energy cut-off (or threshold) values. Each TH scan consists of the counts due to photons with energies above the threshold value. The second step involved subtracting two scans with the desired separation in low energy cut-offs. By obtaining spectroscopic data in this way, further insight into the sensitivity of our results to factors like different choices of energy bins and spectral distortions can be gained [70].

The multi-material phantom was scanned 60 cm downstream of the source and 120 cm upstream

of the detector (Fig. 22). Counts were recorded for 60 seconds using a peak voltage of 90 keV and a current of 110 micro-Amps. From the TH measurements, three energy bins bounded between 18–32, 32–45, and 45–58 keV were created. The effective energy of each bin was found as the weighted average of the energies it enclosed. In particular, the effective energy of the j-th bin was found as

$$E_{\text{eff}}^{j} = \sum_{i=1}^{M} c_i E_i^j \Big/ \sum_{i=1}^{M} c_i, \tag{50}$$

where c_i is the weight (or recorded counts) for the i-th energy within the bin and M represents the number of energies in the j-th bin. From Eq. 50 we estimated values of 25, 36, and 50 keV for the three bins, respectively.

Two additional sets of measurements of the POM rod were recorded for the accuracy and quality analysis described in Sec. 4.2.2. The measurements were taken with different source-to-object distances to investigate how different magnifications influenced the results. For the first measurements, the rod was positioned 60 cm from the detector and 60 cm from the source, resulting in a magnification factor of two (2X). In the second configuration, the distance from the detector was increased to 120 cm for a magnification factor of three (3X). In both cases, the source peak voltage and current were set to 90 keV and 110 μA, respectively.

Three energy bins were constructed from the TH scans of the POM rod. Bounded between 20–40, 40–60, and 60–90 keV, the effective energies of the bins were 28, 48, and 68 keV, respectively. To obtain measurements corresponding to different acquisition times, 60 consecutive (1 s) frames were recorded while maintaining a constant flux at the detector plane of approximately 1.86×10^5 counts/mm^2/s for the 2X case and 8.26×10^4 counts/mm^2/s for the 3X case. The frames were combined (or integrated) to produce results equivalent to acquisitions lasting 1, 3, 7, 15, and 30 seconds.

The leaf sample was imaged with three thresholds scans (with TH values of 10, 20, and 30 keV) in a PCI geometry. The source-to-object distance was 60 cm, and the object-to-detector distance was 120 cm. A single scan with a TH of 10 keV in a conventional (object-plane) geometry was also

recorded. In both cases, the peak source voltage was 45 keV, and the current was set to 220 μA for the PCI measurements and 60 μA for measurements at the object-plane. From the PCI scans, three consecutive bins (10–20, 20–30, and 30–45 keV) were created with effective energies of 15, 25, and 35 keV. All of the scans were recorded for 60 seconds.

4.1.4 Quality and Accuracy Metrics

For the accuracy and quality analysis, scans of the POM rod were taken with different geometries and different exposure times. To quantify how reductions in exposure time (and detected photons) influence the quality of the estimates, we compared the differential signal-to-noise ratio (SNR) [71] of the retrieved pLAC and phase to that of the lowest energy TH scans (of 20 keV). The TH scans are similar to measurements from an energy-integrating detector and captured the counts between 20 keV and the kVp. However, data acquired with PCDs, as in this case, benefit from zero dark noise and the equal (energy independent) weighting of detected photons [72]. In preparation to compute the SNR, all TH scans were first flat-field corrected, and the negative natural log of the results was found.

Before calculating the SNR, the 2D data was "projected" onto an axis perpendicular to the edge of the POM rod [73]. This produced a 1D representation of our data containing the same number of points as the 2D counterpart. The signal level was then estimated by averaging all of the values enclosed in a narrow region near the center of the rod. The background level and the noise were both estimated from the average and the standard deviation of a broad air region next to the rod. Since the effective energy of the TH scans was 38 keV, the pLAC and phase values were estimated for the same energy. The label "PCI" in the results is used to indicate outcomes linked to the TH scans.

To examine the accuracy, the average percent error between the retrieved and expected profiles of the projected electron density were compared. In particular, we used

$$\epsilon = \frac{1}{P}\sum_{i=1}^{P}\left|\frac{r_i - e_i}{e_i}\right|, \tag{51}$$

where ϵ represents the average percent error and P is the total number of points along the profile. Here r_i and e_i are the retrieved and expected values at the i-th position of along the profile of a rod.

4.1.5 Phase-retrieval workflow

The phase retrieval workflow is composed of three simple steps (Fig. 23). The first step involves identifying and correcting defective pixels in all the measurements. The next two steps in the workflow involve energy binning and flat-field correction. The effective energy of each bin is estimated during energy binning. The final step is to implement the retrieval algorithm. For this, only the effective energies of the binned measurements that serve as inputs need to be specified. To mitigate the noise in the a_1 term, the strategy described in Sec. 3.4.3 was applied.

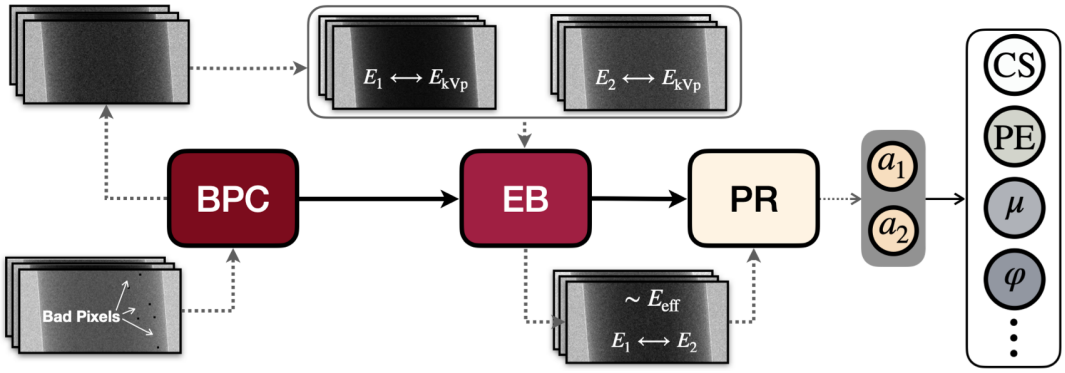

Figure 23: Three steps involved in our spectral phase retrieval workflow. First, bad pixel correction (BPC) is implemented to remove dead-pixels and outliers from the measurements. Next, energy binning (EB) is implemented to create pseudo-monochromatic measurements. Finally, the phase retrieval (PR) is applied to the energy bins, and the material properties a_1 and a_2 are obtained (Sec. 3.1). Energy-dependent properties, like CS and PE, using a_1 and a_2 can be estimated.

4.2 Results

4.2.1 Phase retrieval of multi-material phantom

Fig. 24a shows a PCI measurement of the multi-material phantom that corresponds to an effective X-ray energy of 38 keV. Edge enhancement, a form of phase contrast, can be identified as bright

vertical fringes running along material boundaries, and it is also visible in the average profile (blue curve). The retrieved pLAC and phase of our phantom, estimated for 40 keV, are shown in Fig. 24b and 24c. To obtain these results, we used three (non-overlapping) spectroscopic measurements having effective energies of 25, 36, and 50 keV. As the expected (red) and the average retrieved (green) profiles indicate, the quantities estimated by our phase retrieval display good accuracy. Namely, the average errors between the expected and retrieved profiles of the pLAC and phase were 7.8 % and 3.3 %, respectively.

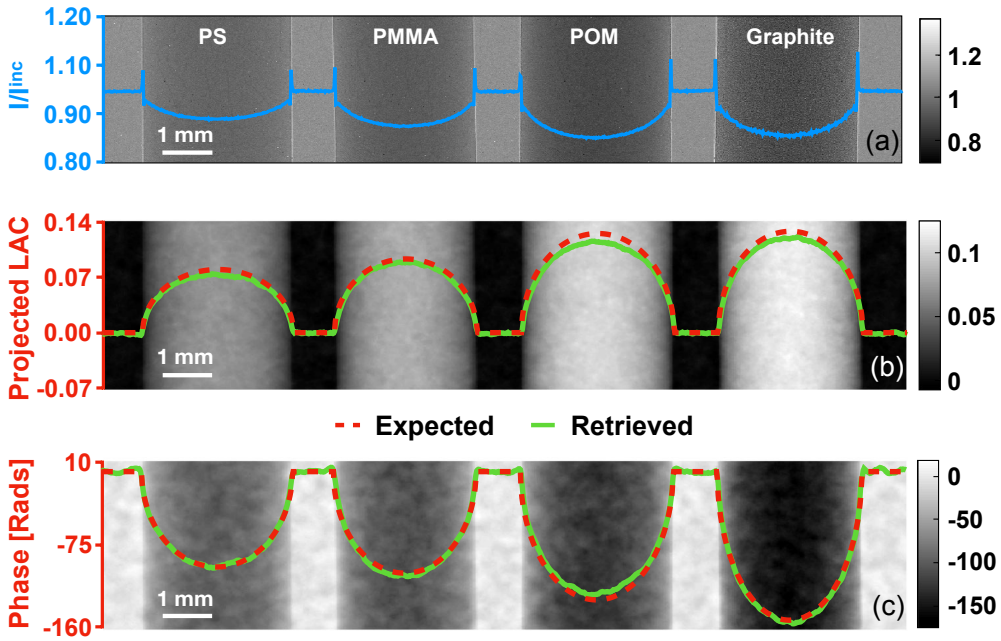

Figure 24: Phase retrieval results for the multi-material phantom. (a) PCI measurements of multi-material phantom showing the expected edge enhancement. (b) pLAC map with the average retrieved profile (green) and expected (red) values superimposed. (c) The retrieved phase difference map with the retrieved (green) and expected (red) average profiles. The retrieved quantities correspond to an effective energy of 40 keV.

A comparison between Fig. 24a and 24c highlights the limitation of relying on attenuation alone to differentiate and properly identify materials. POM and graphite may not be easily discernible from an absorption-based radiograph at 40 keV due to the closeness of their attenuation properties at this energy (Fig. 24b). By contrast, Fig. 24c shows a better separation between these two materials.

Similar to the phase, the projected PE absorption proves as another useful contrast mechanism

for identifying some of the materials in the phantom. The average PE absorption profile is in good agreement (6.3% error) with the expected response (Fig. 25a). Fig. 25b shows the normalized PE contrast along with the other mechanisms. The three curves are unique in shape and show different amounts of distinguishability between pairs of materials, which can be judged by comparing the separation between any two contrast values. For example, one can see from Fig. 25b that the difference between the contrast of graphite and POM due to the phase are greater than the difference due to attenuation. To reiterate, phase retrieval methods which require only one phase enhanced measurement estimate a single contrast mechanism. This method could prove insufficient for proper material differentiation. Based on this evidence, this method can be effective for imaging tasks like quantitative imaging of soft tissue heterogeneity.

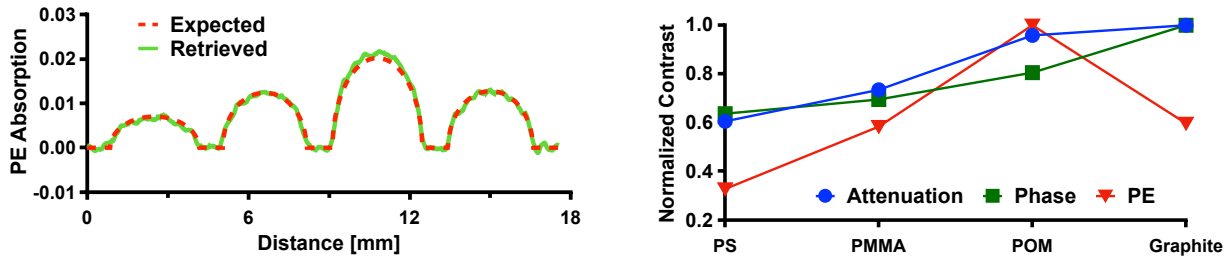

(a) (b)

Figure 25: (a) Average profile of the retrieved and expected photoelectric term for the four material phantom. (b) Normalized contrast for the three retrieved mechanisms. The contrast estimated for each material was normalized by the maximum value for each of the attenuation mechanisms.

4.2.2 Phase Retrieval with Low Photon Counts

The results of the quality analysis for varying photon counts are summarized in Fig. 26a and 26b. For both sets of images, 2X and 3X, the pLAC images gave the highest SNR values (Fig. 26a). Interestingly, the SNR for the 3X images were higher than those of the 2X images, although the opposite was true for their PCI measurements. This observation indicates that gains in signal strength with propagation distance combined with the reduction in effective pitch outweigh the effect of increased noise during phase retrieval. A similar result was observed in a recent publication [74].

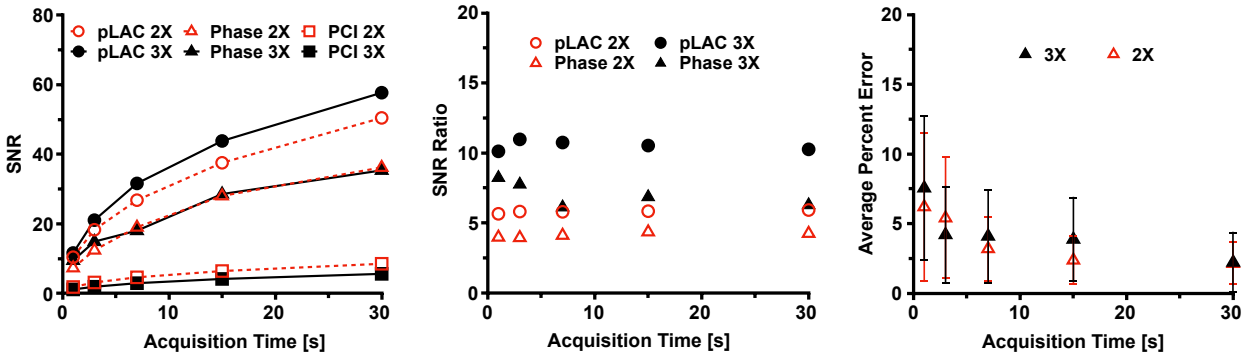

(a) (b) (c)

Figure 26: In all plots, 2X indicates the results for the measurements taken at a two-times magnification and the 3X the configuration with a magnification of three. (a) Signal difference-to-noise ratio (SNR) of the retrieved projected linear attenuation coefficient (pLAC), retrieved phase, and negative log of the measured phase contrast images (PCIs). An effective energy of 38 keV was estimated for the PCI measurement, and the retrieved quantities were computed for this energy (b) Gain in SNR found as the ratio between the SNR of the pLAC and PCI images as well as the phase and PCI images. (c) Average percent error computed from the estimated and expected electron density of POM.

The phase images also showed substantial improvements relative to the PCI images (Fig. 26b). Although the SNR values are nearly equal for both experimental configurations in Fig. 26a, the gains (see Fig 26b) for 3X are greater than those of 2X. It is worth noting that the pLAC is estimated as a linear combination of the two material properties obtained during PR, c.f. Eq. 34 in Sec. 3.1. By contrast, only one property (the electron density) is used to calculate the phase. As illustrated by the differences in SNR between the pLAC and phase, the process of combining the two estimates improves the quality of the resulting images.

The ratio between the SNR values (Fig. 26b) of the pLAC and PCI images, for all acquisition times, revealed nearly constant gains of 10.52 ± 0.35 folds for 3X and 5.80 ± 0.01 folds for 2X. Similarly flat gain factors of 7.06 ± 0.92 and 4.13 ± 0.17 were obtained between the phase and PCI images for 3X and 2X, respectively. The larger variability in the SNR gains of the 3X configuration is likely due to the quality of the measurements at one and three seconds. The counts for these two acquisition times were exceptionally low since an average of only 250 ± 31 counts/pixel/s were measured for the 3X configuration.

The accuracy of the retrieval for different acquisition times was examined using the average percent error between the retrieved and expected values of the projected electron density. Fig. 26c summarizes the results for both configurations. The error bars at each point represent standard deviation of the error values. The results show errors below 10% for all acquisition times and below 5% for acquisitions lasting three or more seconds. Additionally, the accuracy of the electron density estimated in the 2X configuration is slightly improved for most acquisition times lower than 30 seconds, which could be related to the higher quality of the inputs to the PR algorithm resulting from a higher flux on the detector in 2X compared to 3X.

4.2.3 Quality Improvements on a Bio-sample Image

To further demonstrate the benefits of our phase retrieval method, and specifically for a biological sample, we imaged a small Emory Oak leaf (Fig. 27a). This sample was selected because of its abundance of high spatial-frequency structures, which are enhanced by PCI geometries. Fig. 27a and 27b show the negative log of the normalized intensities, $-\ln[I/I^{\text{in}}]$, for a standard (contact) and a PCI measurement of the leaf section. While this quantity is approximately equal to the pLAC in the case of the contact measurement, the same is not true for the PCI measurement.

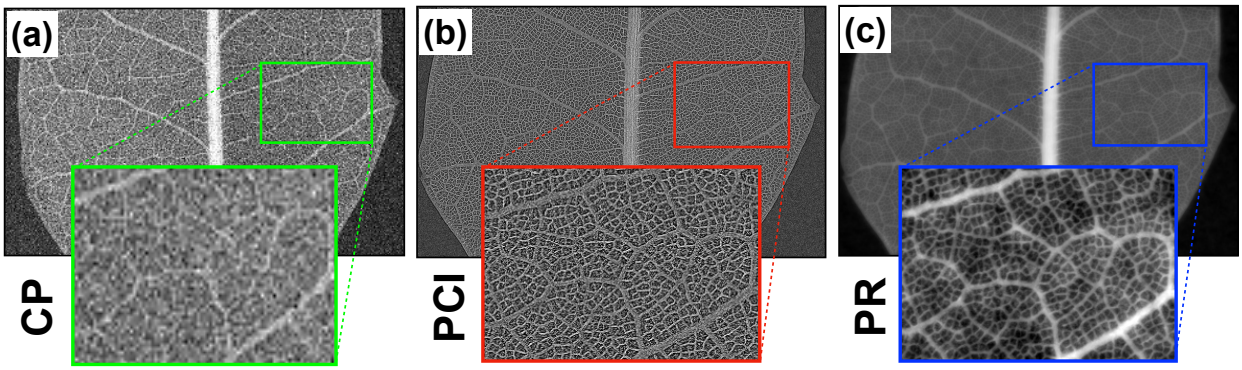

Figure 27: (a) The negative log of the normalized intensity, $-\ln[I/I^{\text{in}}]$, from measurements taken at the contact plane, that is $M = 1$. (b) The same quantity as (a) but using measurements taken in a phase-enhancing configuration, that is $M = 3$. (c) Shows the same quantity as (a) and (b) but with intensity values reconstructed from the retrieved material properties. The insets show the same ROI in all three images, highlighting the quality and detail in each.

To ensure similar counting statistics, the flux for the contact and PCI measurements was,

for both, 480 photons/pixel/s. For comparison, the negative log of the normalized intensity was estimated using the retrieved properties (Fig. 27c). The intensities in all the figures correspond to the same effective energy of 20 keV. Three insets provide an enlarged view of a vein network on the right blade and serve to compare the vein visibility at various sizes and thickness.

In the contact image (Fig. 27a), only relatively large veins on either side of the midrib can be easily discerned. By contrast, the PCI image (Fig. 27b) shows an impressive enhancement in the number and quality of small features that can be resolved. The retrieved image (Fig. 27c) also shows a clearly visible vein network as well as improvements in SNR. The enhanced quality allows for resolving contrast variations between veins (due to differences in size) and at different parts of the veins, possibly due to changes in composition.

4.3 Discussion

In this study, we demonstrated that a phase retrieval method proposed by our group [11] can accurately estimate the material properties of heterogeneous samples with significant improvements to image quality. As indicated by our findings, the complete refractive properties (absorption and phase) can be simultaneously estimated from two (or more) spectroscopic PCI measurements. The spectral PCI images needed for the results were obtained using photon-counting quantum detectors and a polychromatic source.

Unlike some popular phase retrieval algorithms, this method does not restrict the number of materials in a sample. Furthermore, two unique physical properties can be estimated without prior knowledge of the sample. This can prove advantageous for material identification tasks, which may not be possible with phase retrieval methods where a single property is estimated. We sought to highlight this advantage by retrieving the properties of weakly attenuating materials with small differences in effective atomic numbers. Such materials are often indiscernible in conventional radiography.

An important finding in this work relates to the stability of the proposed method with increasing quantum noise. We found that gains in SNR of ten folds and errors below 10% could be obtained

even at low acquisition times, with as few as 250 counts per pixel. This gain in SNR can translate to significant dose reduction in PCI mammography and tomographic imaging like PCI CT.

To date, few attempts have been made to understand the impact that object thickness has in PCI and phase retrieval applications. Our recent efforts in this front show that even for a 5 cm thick soft material, up to 70% of the phase signal is retained in the intensity measurement [75, 76]. Our team is currently expanding on these findings to determine the viability of our method for clinical applications.

Although demonstrated here for propagation-based PCI, the foundation of our approach has been successfully adapted to other PCI geometries such as coded-aperture [77]. In the future, we plan to explore if these geometries offer higher SNR gains, as they are known to improve phase enhancement in comparison to propagation geometry.

Finally, gains in SNR may not adequately represent the benefit for human interpretation of phase retrieved images. Previous work by our team dealt with task-based image quality assessments with human- or model-observers as a reliable metric for evaluating imaging technology [78, 79]. We plan to pursue this direction for phase contrast imaging and phase retrieval. Deciphering the combined impact of absorption and phase contrast in human-observer perception presents new opportunities and challenges in task-based image quality assessment research.

5 Assessments of the Error Introduced by Approximations to the Transport of Intensity Equation

The transport of intensity equation (TIE), c.f. Sec. 2.2.3 and Sec. 3.1, has become the most common mathematical framework to derive PR algorithms for PB-PCI [80,81]. This equation describes how the intensity of a propagating field relates to variations in material properties perpendicular to the direction of propagation. Due to its convoluted and non-linear form, the TIE is generally used only after applying several simplifying assumptions.

A common and important strategy to simplify the TIE is to apply the *weakly-interacting object* (or "*weak object*") approximation [11, 14, 28, 35, 48]. This approach assumes a sufficiently weak interaction between X-rays and the sample so that second-order variations in material properties can be ignored. Analogously, the dot product of the two gradient terms in Eq. 31 can be neglected and the TIE can be written as an inhomogeneous Poisson equation for which solutions can be readily found. This section presents a strategy to understand the errors related to enforcing the weak object approximation in PR algorithms. We also show how the validity of this approximation can depend on the sample thickness as well as the differences between the LAC of a sample and its background.

Two other common simplifications are to assume that the sample is homogeneous [50] or that the energies of the radiation are high enough so that a "duality" between the phase and attenuation of X-rays emerges [51]. For the latter, the phase and attenuation become primarily dependent on and can be written in terms of the electron density. We studied the error due to these two assumptions with experimental data using existing PR techniques that enforce them. In particular, we tested the homogeneity assumption with the scheme proposed in [50], while the phase-attenuation duality (PAD) condition was tested with the method given in [51]. Appx. A and B give derivation of these two techniques in forms well suited for "point-like" laboratory sources. To test the weak object approximation, a new PR scheme was derived in Sec. 5.4. In all three cases the retrieved property was the *projected thickness*. We will refer to PR methods based on the homogeneous, PAD, and

weak object approximations as hPR, PAD-PR, and wPR respectively.

5.1 Analytical form of the weakly interacting object approximation

For this section, it is convenient to write the TIE in the form given by [82], that is:

$$-k(E)\frac{\partial}{\partial z}\ln I(\boldsymbol{r}_\perp,z;E)=\boldsymbol{\nabla}_\perp\ln I(\boldsymbol{r}_\perp,z;E)\cdot\boldsymbol{\nabla}_\perp\varphi(\boldsymbol{r}_\perp,z;E)+\boldsymbol{\nabla}_\perp^2\varphi(\boldsymbol{r}_\perp,z;E). \tag{52}$$

This results is equivalent to Eq. 31 since $g^{-1}(\boldsymbol{r}_\perp)[\partial g(x)/\partial x]=\partial\ln g(x)/\partial x$ for some arbitrary function $g(x)$. All variables are defined in Sec. 3.1. Due to a nonlinear coupling between the phase and the intensity (i.e. $\boldsymbol{\nabla}_\perp\ln I\cdot\boldsymbol{\nabla}_\perp\varphi$), an implementation of Eq. 52 without simplifications is often prohibitively complex.

A common way to simplify Eq. 52 involves assuming that changes to the phase and intensity of the X-rays due to interactions with the sample are sufficiently weak so that (at $z=0$) the first term can be neglected on the right-hand [83]. Equivalently, we require that

$$\frac{\boldsymbol{\nabla}_\perp\ln I_E(\boldsymbol{r}_\perp,0)\cdot\boldsymbol{\nabla}_\perp\varphi_E(\boldsymbol{r}_\perp,0)}{\boldsymbol{\nabla}_\perp^2\varphi_E(\boldsymbol{r}_\perp,0)}\ll 1, \tag{53}$$

where the energy dependence is written as a subscript, e.g. $I(\boldsymbol{r}_\perp,z;E)=I_E(\boldsymbol{r}_\perp,z)$. To better understand the validity of this condition, an analytical representation can be developed with the projection approximation. This simplifying assumption is generally valid in biomedical applications [44].

We start by considering a sample composed of two homogeneous materials with projected thicknesses given by $T^{(m)}(\boldsymbol{r}_\perp)$, where the superscript $m[=1,2]$ identifies each material, see Fig. 28b. In addition, their combined thickness is defined as $A(\boldsymbol{r}_\perp)[=T^{(1)}(\boldsymbol{r}_\perp)+T^{(2)}(\boldsymbol{r}_\perp)]$. If the projection approximation holds at the $z=0$ plane, then the intensity and phase of the X-ray wavefield due to the sample can be written as [50, 84]:

$$I_E(\boldsymbol{r}_\perp)=I_E^{\text{in}}\exp\left(-\Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp)-\mu_E^{(1)}A(\boldsymbol{r}_\perp)\right) \tag{54}$$

and

$$\varphi_E(\boldsymbol{r}_\perp) = -\frac{2\pi}{k_E} r_e \left[\Delta\rho_e T^{(2)}(\boldsymbol{r}_\perp) + \rho_e^{(1)} A(\boldsymbol{r}_\perp)\right], \tag{55}$$

where as before r_e is the classical electron radius, ρ_e is the electron density, μ_E is the linear attenuation coefficient (LAC) at energy E, $\Delta\mu_E = \mu_E^{(2)} - \mu_E^{(1)}$, and $\Delta\rho_e = \rho_e^{(2)} - \rho_e^{(1)}$. With these results, the differential terms in Eq. 53 can be readily defined. Without lost of generality, we can focus on a 1-D derivative (e.g. x) and arrive at:

$$\frac{\partial}{\partial x} \ln I_E(\boldsymbol{r}_\perp) = -\Delta\mu_E \frac{\partial}{\partial x} T^{(2)}(\boldsymbol{r}_\perp) - \mu_E^{(1)} \frac{\partial}{\partial x} A(\boldsymbol{r}_\perp) \tag{56}$$

and

$$\frac{\partial}{\partial x}\varphi_E(\boldsymbol{r}_\perp) = -\frac{2\pi}{k_E} r_e \left[\Delta\rho_e \frac{\partial}{\partial x} T^{(2)}(\boldsymbol{r}_\perp) + \rho_e^{(1)} \frac{\partial}{\partial x} A(\boldsymbol{r}_\perp)\right]. \tag{57}$$

Similarly, the second derivative becomes

$$\frac{\partial^2}{\partial x^2}\varphi_E(\boldsymbol{r}_\perp) = -\frac{2\pi}{k_E} r_e \left[\Delta\rho_e \frac{\partial^2}{\partial x^2} T^{(2)}(\boldsymbol{r}_\perp) + \rho_e^{(1)} \frac{\partial^2}{\partial x^2} A(\boldsymbol{r}_\perp)\right]. \tag{58}$$

An analogous set of equations would result for the y-component of the gradient operation in Eq. 53. The y-component equations, in combination with Eq. 56 to 58, allow for an analytic expression of Eq. 53 when the first and second order differentials of $T^{(2)}(\boldsymbol{r}_\perp)$ and $A(\boldsymbol{r}_\perp)$ can be defined in tractable functional forms. Although these results are limited to a two-material sample, they give a powerful tool to study the TIE with analytical techniques.

5.1.1 Evaluating the approximation with an analytical phantom

A simple yet insightful example is that of a rod embedded inside a rectangular slab, each composed of a different material. Fig. 28a shows a rendering of an object with these characteristics. The ends of the slab are assumed to be found sufficiently far away from the field-of-view so that they can be ignored. In [84], such an object is called "ternary" since it can account for up to three unique materials (including free-space). Fortunately, a physical phantom with characteristics analogous

to Fig. 28a can be constructed with ease. The setup depicted in Fig. 28b shows how experimental measurements were collected in this study to test the findings described in this section.

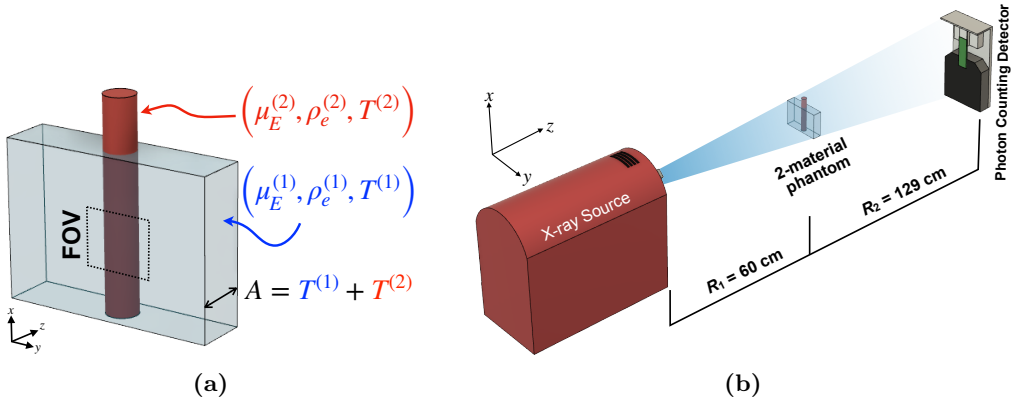

Figure 28: (a) Two-material phantom composed of two homogeneous objects with basic shapes. The thickness A is constant and the field-of-view (FOV) is limited to a small region that encapsulates the ends of the rod but is far from the ends of the rectangular slab. (b) Geometrical configuration of the PCI measurements used for this section.

The basic geometrical properties of the object lead to additional simplifications. First, the projected thickness of a rod can be properly represented with a 1-D piece-wise function [85]:

$$T^{(2)}(x) = \begin{cases} 2\sqrt{r^2 - x^2} & |x| \leq r \\ 0 & \text{otherwise} \end{cases}. \tag{59}$$

This, in turn, allows the use of the 1-D derivatives in Eq. 56 to 58. Next, since the combined thickness $A(\boldsymbol{r}_\perp)[= A]$ is a constant, we have that $\partial_x A(\boldsymbol{r}_\perp) = \partial_{xx} A(\boldsymbol{r}_\perp) = 0$ and Eq. 53 can be written as:

$$2\Delta\mu_E r \left[(x/r)^2 \sqrt{1 - (x/r)^2}\right] \ll 1, \qquad |x| \leq r. \tag{60}$$

An analysis of Eq. 60 leads to important points. First, the expression inside the square brackets controls the shape, and its values are bounded between 0 and ~0.385 irrespective of the choice of r. For this reason, the maximum value of the left side of the inequality is approximately $0.77 r \Delta\mu_E$, and the validity of Eq. 60 is therefore regulated by the product of r and $\Delta\mu_E$. Fig. 29a shows three curves corresponding to different radii with $\Delta\mu_E$ set to unity for all cases. By holding $\Delta\mu_E$

constant, the effect of changing r is isolated to see how the condition deteriorates as the value of r increases. The distance between the peak value and one shrinks with increasing radius. This result indicates that the validity of the condition in Eq. 53 for regions of an image near material boundaries can be less acceptable in thicker samples.

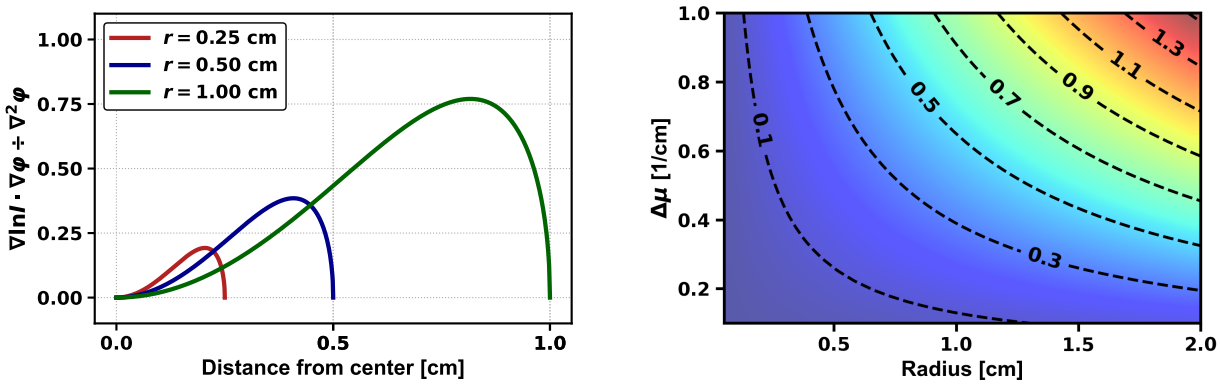

(a) (b)

Figure 29: (a) Plot of the left side of the inequality described by Eq. 60 using three radii and $\Delta\mu_E = 1$. (b) Plot of the maximum value of the left side of Eq. 60 for different combinations of $\Delta\mu$ and r.

Fig. 29b shows the maximum value of the left-hand side of Eq. 60 for different combinations of r and $\Delta\mu_E$. A set of the contour can be understood as delineations between different regions of validity for the inequality. One limitation of this viewpoint is that, as shown in Fig. 29a, the values in Fig. 29b pertain to a single point (the maximum) along the profile. Furthermore, all other points predicted by Eq. 60 for a given choice of r and $\Delta\mu_E$ fall below the values used for Fig. 29b. Nonetheless, the information conveyed by the figure can help inform where the inequality is most at risk of breakdown. We also note that a significant portion near the maximum of each curve, in Fig. 29a, shows relatively high amplitudes and the values are close to those used to create Fig. 29b.

5.2 Influence of the weak-object approximation on phase retrieval

Phase retrieval (PR) can be viewed as the inverse to a mathematical model that describes the physics of phase contrast image formation with sufficient accuracy (c.f. Sec. 3.2). For this reason, a key aspect of a PR algorithm is the accuracy with which the measured intensity is predicted by

its conjugate forward model. Here, the primary concern is understanding the influence of Eq. 53 in PR schemes that use this approximation to simplify the TIE. To achieve this, we can investigate how the intensity predicted by an approximation-free or "full" TIE differs from that predicted after Eq. 53 is enforced. For the case of the two-material phantom in Fig. 28a, the intensity obtained from the full TIE can be written as:

$$I_E(\boldsymbol{r}_\perp, R_2) = I^{\text{in}} e^{-(\text{AC}+\text{PC})} \tag{61}$$

where

$$\text{AC} = \Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp) + \mu_E^{(1)} A \tag{62}$$

represents the part of the exponent that provides the absorption contrast, while the phase contrast is described by the term

$$\text{PC} = \frac{2\pi R_2}{k_E^2} r_e \Delta\rho_e \left\{ \Delta\mu_E \left[\boldsymbol{\nabla}_\perp T^{(2)}(\boldsymbol{r}_\perp) \right]^2 - \boldsymbol{\nabla}_\perp^2 T^{(2)}(\boldsymbol{r}_\perp) \right\}. \tag{63}$$

To arrive at Eq. 61, the z-derivative on the left-hand side of Eq. 52 was expanded as a finite difference. This is understood to be valid in the near-field regime where R_2 is sufficiently small, see Sec. 2.2.4. A comprehensive analysis of this approximation can be found elsewhere [86]. Note that when PC is negligible, e.g. $R_2 \approx 0$, Eq. 61 yields the intensity of the sample predicted by Beer's law.

Enforcing the condition in Eq. 53 is equivalent to neglecting the quadratic term inside the set brackets in Eq. 63. Therefore, by comparing the response of Eq. 61 with and without neglecting the quadratic term in PC (Eq. 63), we can estimate the error in the intensities and gain insight about the influence that enforcing Eq. 53 can have in PR. Furthermore, we can learn how material properties can influence the error.

Fig. 30 describes the percent error for five polymethyl methacrylate (PMMA) rods of different thicknesses in the absence of a background material (i.e. $\mu_E^{(1)} = 0$). The horizontal axis corresponds

to the ratio between the radius and the distance away from the center of each rod. This choice facilitates a comparison between the different radii. The values were limited in the horizontal axis to 96% of the way to the edge, since the errors converged for higher values. Interestingly, the plot shows error that becomes less severe as the thickness of the sample increases. This may seem to contradict the results in Fig. 29a. However, in here we are concerned with the errors in the intensity as a whole. As the material thickness increases, the AC term, which is directly proportional to the thickness, dominates over the PC term. Therefore, errors related to approximations in the PC term become less significant even if the approximations become less valid.

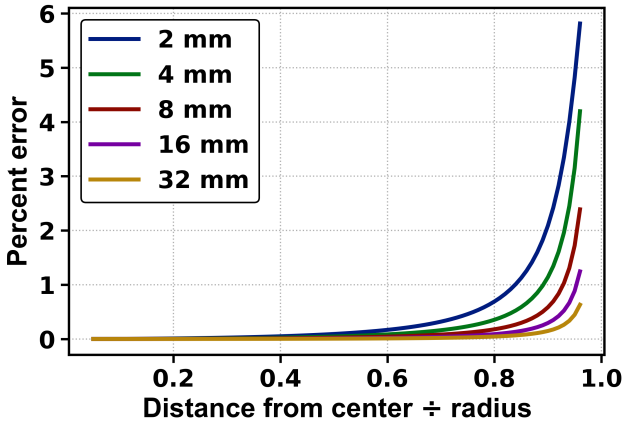

Figure 30: Plot of the percent error between the intensity predicted by Eq. 61 with and without enforcing the weak-interaction approximation. The intensities were estimated for PMMA rods of different thicknesses. The horizontal axis represents the ratio between the distance away from the center of each rod and the corresponding radius.

The contrast between the properties of the rod and the background can also influence the error in the expected intensity. Fig. 31 shows the changes in error along the profile of the rod for a set of $\Delta\mu$ values. To obtain these results, we simulated a 3 mm Al rod embedded in a 10 mm block of polyvinyl chloride (PVC), polytetrafluoroethylene (PTFE), polyoxymethylene (POM), PMMA, polystyrene (PS), as well as no background material (or free-space). At an energy of 30 keV, this yielded $\Delta\mu$ values of 0.99, 2.16, 2.58, 2.69, 2.77, and 3.05 respectively. The distance from the center of the rod was allowed to reach 98% of the way to the edge.

Similar to the results found when evaluating the inequality in Eq. 53 (Fig. 29), reducing the

contrast between the linear attenuation of the two materials is favorable for the weak object approximation and leads to lower error. Therefore, the approximation is more valid when dealing with similar materials, as it may be the case when imaging different tissue types. Since the same variation in material properties can be obtained for different material combinations and different energies, the results shown here generalize beyond the choices we made.

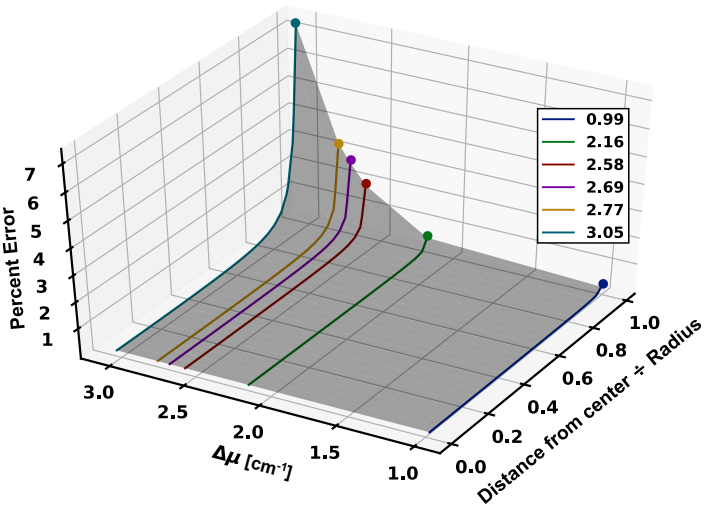

Figure 31: Percent error between the intensities predicted Eq. 61 with and without enforcing the WOA in Eq. 63. One horizontal axis represent different contrast values ($\Delta\mu$) between the target and background materials while the other shows the ratio between the distance from the center of the target to the radius of the target. The ratio was kept to 98 % of the way to the edge of the target.

5.3 Overview of the homogeneous and phase-attenuation duality approximations

The homogeneous approximation can be enforced in Eq. 61 by assuming that the sample is composed of a single material so that $\mu_E^{(1)} = 0$, $\Delta\mu_E = \mu_E^{(2)}$, and $\Delta\rho_e = \rho_e^{(2)}$. The resulting intensity can be used to obtain a PR equation that solves for $T^{(2)}(\boldsymbol{r}_\perp)$ as demonstrated by [50]. One way to approximate a heterogeneous sample as a homogeneous object is to describe its properties with average values. The LAC and electron density can be written as:

$$\bar{\mu}_E = \frac{1}{N}\sum_{i=1}^{N}\mu_E^{(i)} \qquad \text{and} \qquad \bar{\rho}_e = \frac{1}{N}\sum_{i=1}^{N}\rho_e^{(i)}. \tag{64}$$

We emphasize that a PR equation based on the homogeneous approximation is an exact inverse to Eq. 61 when the object is composed of a single material. Therefore, the error this approximation introduces will be dictated by how well the properties of the object can be approximated with single values of μ and ρ_e, as it is assumed in Eq. 64.

The PAD condition assumes that the LAC can be properly represented by only considering loss of photons due to Compton scattering (i.e. $\mu \approx \sigma_{cs}$). When the projection approximation is valid, we can write

$$\mu_E^{(i)} \approx \Phi_{\mathrm{KN}}(E)\rho_e^{(i)}T^{(i)}(\boldsymbol{r}_\perp), \tag{65}$$

where $\Phi_{\mathrm{KN}}(E)$ is the Klein-Nishina function. For Eq. 65 to be valid (with around 90 % accuracy), one should limit the use of this approximation to materials with effective atomic number falling below 10 and X-ray energies above 60 keV [51]. Materials that can satisfy this condition include soft tissues, water, and several carbon-based polymers. Sec. 1.1.1 discussed the case of water, for which Compton scattering accounts for over 80% of the total attenuation above 50 keV and rises above 90% for energies above 60 keV (Fig. 1).

5.4 Derivation of a phase retrieval to assess the weak-object approximation

When the gradient term can be neglected, Eq. 52 simplifies to:

$$-k_E\frac{\partial}{\partial z}\ln I_E(\boldsymbol{r}_\perp, z) = \boldsymbol{\nabla}_\perp^2\varphi(\boldsymbol{r}_\perp). \tag{66}$$

From the finite difference expansion of the left-hand side we can write

$$\ln\frac{I_E(\boldsymbol{r}_\perp, R_2)}{I^{\mathrm{in}}} = -\mu_E^{(1)}A - \Delta\mu_E T(\boldsymbol{r}_\perp) + \frac{2\pi R_2}{k_E^2}r_e\Delta\rho_e\boldsymbol{\nabla}_\perp^2 T(\boldsymbol{r}_\perp). \tag{67}$$

where T is understood to be the projected thickness of the target material, i.e. $m = 2$. This equation assumes two important characteristics about the incident radiation. First, the source is assumed to have an infinitesimally small dimension or to be sufficiently far away from the object so perfect

spatial coherence is achieved. Second, the incident radiation is assumed to be monochromatic. Combined, these two assumptions mean that the object is illuminated by a fully coherent wavefield [12]. The interested reader can find more information about these two assumptions in [12,14,35,87].

The geometrical magnification, M, that results when the source is a finite distance away from the illuminated object can be easily introduced into Eq. 67 following the steps in [53]. With this addition, an estimate of the projected thickness of the second material can be readily obtained in Fourier-domain as:

$$T(\boldsymbol{r}_\perp) = -\frac{1}{\Delta\mu_E}\mathcal{F}^{-1}\left\{\frac{\mathcal{F}\left\{\ln\left[\frac{I_E(M\boldsymbol{r}_\perp, R_2)}{I^{\text{in}}}e^{\mu_E^{(1)}A}\right]\right\}}{1+\Gamma_E\frac{\Delta\rho_e}{\Delta\mu_E}|\boldsymbol{\nu}|^2}\right\} \tag{68}$$

where $\Gamma_E[=2\pi r_e M R_2/k_E^2]$ depends on the energy and experimental geometry.

Similar equations for the homogeneous approximation and the phase-attenuation duality are derived in Appx. A and B, respectively. Equation 68 shares many similarities with the PR formula in [84], c.f. Eq. 81 in Appx. A. An important difference pertains to the placement of the natural log. The PR formula in [84] can be shown to predict the same intensity as Eq. 61. Therefore, this equation is used in Sec. 5.5.4 as an "approximation-free" PR implementation.

5.5 Phase retrieval with experimental data

The error in $T(\boldsymbol{r}_\perp)$ estimated with the three PR algorithms was analyzed using measurements of a physical phantom that closely resembled Fig. 28a. To construct the phantom, a rectangular plastic container with a constant inner-thickness of 15.30 ± 0.02 mm was used as well as an aluminum rod measuring 3.20 ± 0.02 mm in diameter. Using a 3-D printer, we manufactured a lid to facilitate the placement and removal of the rod from the middle of the container. The rod was imaged in the container with and without water. This allowed the testing of two material contrasts: aluminum in air and aluminum in water. To ignore the effect of the plastic walls, we used measurements of the empty box without the aluminum rod as flat-fields.

To satisfy the assumption of high temporal and spatial coherence, a microfocus source was

used in conjunction with a high-resolution photon counting detector. In particular, a single-chip CdTe Medipix3RX detector was used in *charge-summing mode*, which mitigates the effect of *charge sharing* [69]. The Medipix3RX detector has a 256×256 pixel array separated at a uniform pitch of 55 μm. For the source, we used a polychromatic microfocus X-ray tube (Hamamatsu 150 kV L12161-07). The source was operated in "small focus mode" at a peak voltage of 90 keV and at a current of 110 μA. The sample was placed 60 cm away from the source and 129 cm away from the detector (Fig. 28b). A set of threshold measurements were recorded with 22 low-energy cut-offs ranging from 17 to 72 keV and uniformly spaced at intervals of 2.5 keV. Each measurement lasted 60 seconds.

				Background Material			
				Air		Water	
Set	Bin Bounds [keV]	E_{eff} [keV]	Mean Counts	SNR	$\Delta\mu$ [cm^{-1}]	SNR	$\Delta\mu$ [cm^{-1}]
1	22.0 → 27.0	24.5	5473	11.1	5.3	7.0	4.7
	27.0 → 37.0	31.5	5780	11.8	2.7	8.2	2.3
	32.0 → 47.0	38.2	5652	10.2	1.7	7.5	1.4
2	29.5 → 37.0	33.0	3798	7.6	2.4	5.5	2.1
	39.5 → 52.0	44.9	3669	6.7	1.2	4.7	1.0
	47.0 → 62.0	53.6	3515	6.3	0.9	4.4	0.7
3	32.0 → 37.0	34.4	2292	4.5	2.2	3.3	1.9
	49.5 → 59.5	54.4	2344	4.2	0.9	3.0	0.6
	57.0 → 72.0	62.8	2284	4.9	0.7	3.3	0.5

Table 5: Sets of measurements used to study the error in phase retrieval due to approximations in the TIE. The sets were defined by grouping energy bins with similar mean counts per pixel in the flat-field measurements and, therefore, similar SNR. The same bin boundaries were used for air and water as background materials, which resulted in different SNR values since the acquisition time was 60 s in both cases. The contrast in LAC ($\Delta\mu$) between Al and the background materials is also listed for every E_{eff}. For air as background, the change in LAC is equal to the LAC of Al.

From the measurements, three sets of energy bins were created for each of the background materials, which yielded a total of 18 unique energy bins. The criteria used to define each set was a similar number of average counts per pixel in the flat-fields. Thus, the elements of each set have similar SNR values. To estimate the SNR, a narrow region (of about 10 pixels wide by

240 pixels long) extending along the middle of the rod was used to obtain the "signal" level. The "background" level and "noise variance" were both extracted from the intensity values in a broader region (around 20 pixels wide by 240 pixels long) on the air gap to the left of rod, see Fig. 32b. The SNR was found as (signal − background)/$\sqrt{\text{noise variance}}$.

Table 5 lists the properties of the elements in each set of energy bins, including the SNR and $\Delta\mu$ for air and water as background materials. To obtain $\Delta\mu$ between air and Al, air was considered to be non-attenuating (i.e. an LAC of zero). Note that substituting air for water as the background component led to a lower SNR. This is because the acquisition time was 60 seconds for all threshold measurements. When water surrounded the rod, the resulting images had lower counts per pixel because water has a higher attenuation than air.

5.5.1 Correcting spectral distortions

A key goal in this study was to check if any of the error trends described in Sec. 5.2 could be observed after using Eq. 68 with experimental data. After analyzing the average percent error between the expected $T(\boldsymbol{r}_\perp)$ and that retrieved from the bins listed in Table 5, we observed that energy-dependent (or spectral) distortions dominated the trend. Furthermore, we assumed that trends due to the weak object approximation would be completely obscured by the high magnitude of the errors due to spectral distortions. Figure 32a shows a plot of the average percent error obtained after applying Eq. 68 without mitigating spectral distortions. The errors increase with decreasing energy rising above 20% below 40-keV and reaching $\sim 50\%$ at 24.5 keV.

We associated the trend in Fig. 32a primarily to effects introduced by the hybrid PCD used. These types of detectors can introduce errors that are highly correlated with the energy of the incident beam [25, 88, 89]. In some cases, detector-related distortions can be so severe that they counterbalance the benefits of using PCDs [88]. For the CdTe Medipix3RX, spectral distortions are strongest at low energies near the fluorescence peak of Cd ($\sim$ 27 keV) and Te ($\sim$ 31 keV) and become less severe with increasing energies. Another possible source of energy-dependent distortions influencing the measurements is beam hardening, which is also more detrimental at

lower energies.

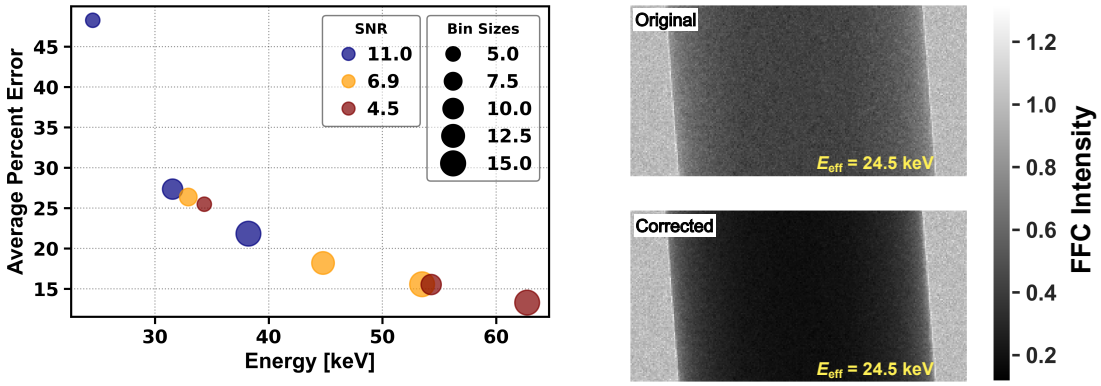

(a) Al in Air **(b)**

Figure 32: (a) Error in the retrieved thickness found by applying Eq. 68 with the three sets of energy bins in Table 5 corresponding to the Al rod in air. Since the elements in each set have similar SNRs, the average SNR of the measurements and a unique color were used to differentiate each set. The size of each point is related to the size of the corresponding energy bin in keV. (b) Flat-field corrected (FFC) intensity before and after spectral correction for the Al rod surrounded by air at the lowest effective energy in Table 5.

Spectral distortions were mitigated with a strategy similar to the one proposed by [88]. In particular, we calibrated the intensity values in the images to ensure a close match with the values predicted by Eq. 61 for the phantom. Figure 32b shows the flat-field corrected (FFC) intensity of the lowest energy bin ($E_{\text{eff}} = 24.5$ keV) in Table 5 before and after spectral correction, where the background material is air. As the obvious change in the intensities of the sample indicates, the original bin was severely distorted. In principle, the intensity at every pixel can be remapped to a different value, and therefore the corrected images can be significantly different from the originals. An effort was made to minimize changes to the noise "texture" and signals at (and near) material boundaries. Additional details of our spectral correction strategy can be found in Appx. D.

5.5.2 Accurate estimates of the effective energy

Since each bin in Table 5 corresponds to a (relatively narrow) band of energies, to represent a monochromatic measurement with these bins requires assigning an *effective energy*, E_{eff}, to each. In phase retrieval algorithms like the ones considered in this section, the effective energy of a

measurement is an important parameter that influences the accuracy and quality of the results [84]. As Eq. 68 implies, the estimated thickness depends on the proper definition of three energy-dependent parameters, i.e. Γ_E, $\mu_E^{(1)}$, and $\mu_E^{(2)}$.

We estimated E_{eff} for our energy bins in three steps. First, the intensity of a numerical phantom was simulated. The phantom was designed to have the dimensions and properties of the physical analog. Because of its symmetry, the simulated intensity at energy E could be represented with a 1-D profile:

$$I_E(x) = S_E e^{-(\text{AC}+\text{PC})} \tag{69}$$

where S_E are the photons emitted by the source at energy E, while AC and PC are the amplitude and phase contrast terms defined by Eq. 62 and 63, respectively. This step involved estimating an X-ray spectrum, S_E, that properly represented the one emitted by our source. For this we used xpecgen, a freely available Python program developed after [90]. The parameters required by the program (e.g. target material and peak voltage) were selected to reflect our experimental choices.

Next we simulated the FFC intensities corresponding to our bins. For a bin bounded between E_i and E_j, the FFC intensity was found using

$$I_{E_i \to E_j}(x) = \sum_{n=E_i}^{E_j} I_n(x) \Bigg/ \sum_{k=E_i}^{E_j} S_k. \tag{70}$$

In the final step we searched for a monochromatic FFC intensity, $I_{E_{\text{eff}}}(x)$, that closely represented $I_{E_i \to E_j}(x)$. We achieved this with an iterative scheme that minimized the cumulative absolute difference between the intensities, that is

$$E_{\text{eff}} = \arg\min_{E_{\text{eff}}} \int |I_{E_i \to E_j}(x) - I_{E_{\text{eff}}}(x)| \mathrm{d}x, \tag{71}$$

where the limits of the integral are defined to include the profile of the sample. In practice, the integral was applied as a sum that ran over the discrete elements of the profile. The value of E_{eff} was updated until the cumulative difference reflected a 99.5 % match between the monochromatic

and polychromatic profiles.

5.5.3 Analyzing the influence of the weak-object approximation

Figures 33a shows the average percent error obtained from the measurements of the Al rod surrounded by air after correcting spectral distortions. The errors were between ~ 3 to ~ 15 times lower than in Fig. 32a, where spectral distortions were not corrected. Furthermore, the errors remain below ~ 5 % in all cases, averaging 4.4 ± 0.5 %. Although a decrease with increasing energy was anticipated, the opposite was true in our results. This unexpected trend could have resulted from the interplay between variations in data quality and some unintentional consequences of the processing described in Sec. 5.5.1 and 5.5.2.

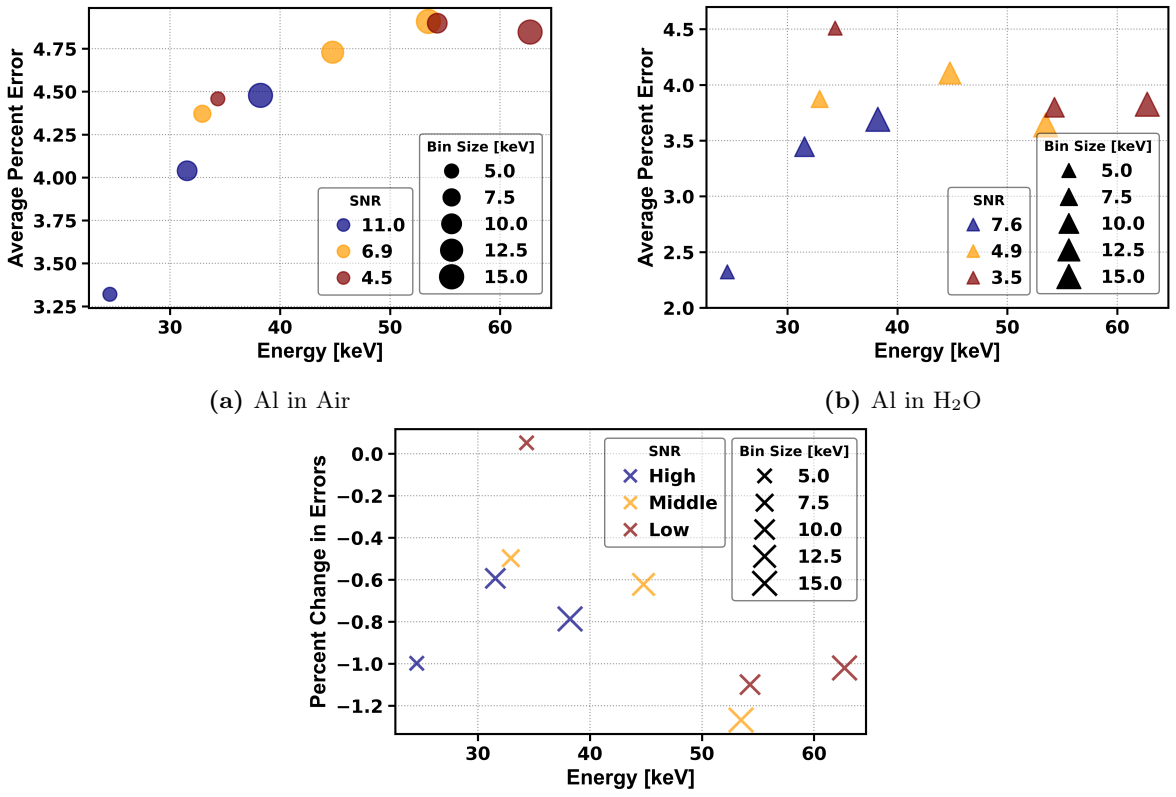

(a) Al in Air

(b) Al in H_2O

(c)

Figure 33: (a) Error in the retrieved thickness of the Al rod surrounded by air after spectral correction was applied to the data. (b) Error in the retrieved thickness of the Al rod submerged in water after spectral correction. (c) Relative change between the average percent errors found after applying wPR to the bins with air in the background, (a), and the errors found after substituting the background with water, (b).

To elaborate on our interpretation of the trend in Fig. 33a, first recall that our strategy to mitigate spectral distortions ensures a close agreement between the corrected intensity values and those predicted by Eq. 70. Also, each intensity in the numerator of Eq. 70 is found with the full TIE (Eq. 61). Our scheme to estimate E_{eff} aims to ensure that one can treat the corrected data as a monochromatic measurements with an accuracy of at least 99.5 %. Ideally, after these two steps, the errors from the experimental data should conform with the analytical results in Sec. 5.2. In practice, however, factors like noise and source-induced blurring will introduce additional errors that are not accounted for in this study.

The results for the Al rod submerged in water, shown in Fig. 33b, also fail to show a trend that decreases with increasing energy. Interestingly, a comparison with Fig. 33a reveals improvements in the accuracy, which agrees with our analytical prediction. This encouraging outcome emerged despite the data with water in the background having lower SNR. Figure 33c shows the change in the error after replacing air with water as the background material, in which a negative value indicates a reduction in the error. Aside from the bin at 34.4 keV, which shows a slight (< 0.1 %) increase in the error, every bin shows more than 0.4 % improvement. Also, higher energy bins appear to benefit the most.

5.5.4 Analysis of the phase-attenuation duality and homogeneous approximations

Since the PAD-PR equation estimates the projected electron density, $\rho_{\text{p}}[= \int \rho_e(\boldsymbol{r}_\perp, z)\mathrm{d}z]$, the thickness of the rod was found using

$$\rho_{\text{p}} = \rho_e^{(1)} A + \Delta\rho_e T^{(2)}(\boldsymbol{r}_\perp) \longrightarrow T^{(2)} = \frac{1}{\Delta\rho_e}\left[\rho_{\text{p}} - \rho_e^{(1)} A\right]. \tag{72}$$

This PR scheme is intended to work with materials having effective atomic numbers below 10, which is not met in our case since Al has an atomic number of 13. Even when the condition for the atomic number is satisfied, the energy of the X-rays is expected to be higher than 60-keV [51]. A more specific, but related, requirement is to use an energy where Compton scattering accounts

for more than 90 % of the sample's total LAC. For Al, Compton scattering constitutes ~ 80 % of the LAC at 100 keV, which is well above the highest bin in Table 5.

Perhaps as expected from the PAD-PR requirements, the errors in Fig. 34a and 34b are significantly higher than those found using Eq. 68. The trends in both plots resemble the one influenced by spectral distortions (Fig. 32a). This results because the assumptions in PAD-PR become more valid with increasing energy and thus the estimated thickness becomes more accurate. Additionally, unlike the results in Fig. 33, using water as the background material led to higher errors.

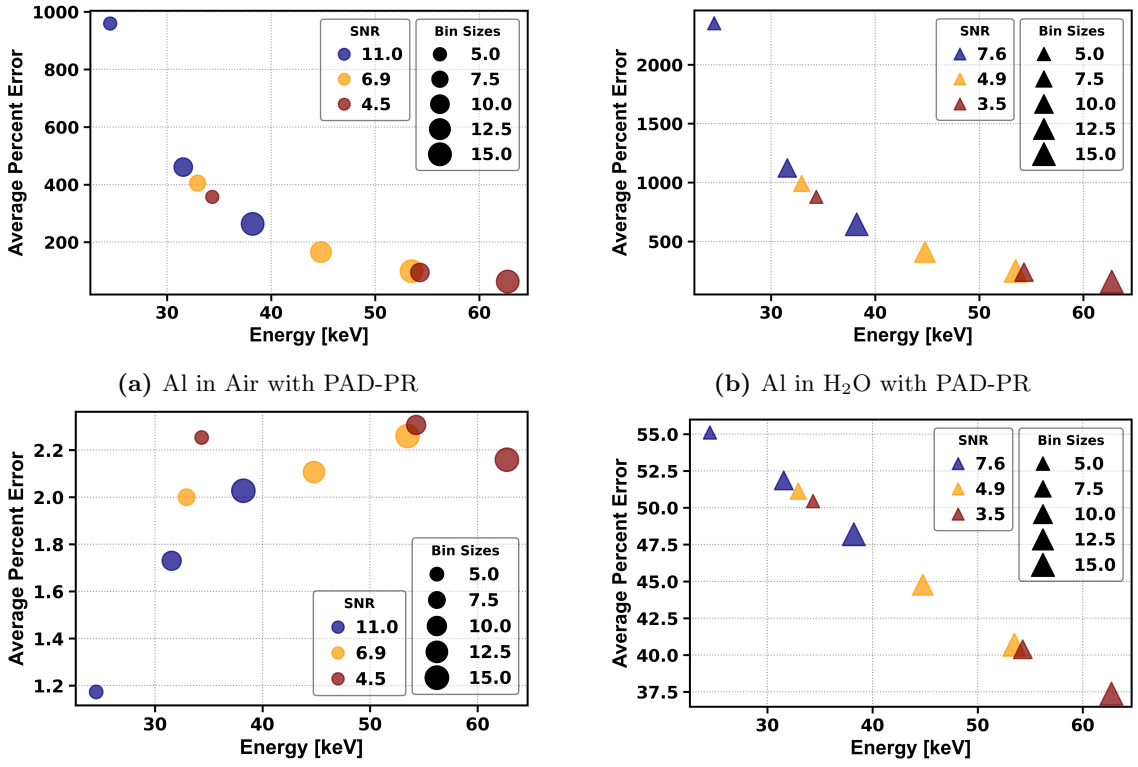

(a) Al in Air with PAD-PR

(b) Al in H_2O with PAD-PR

(c) Al in Air with hPR

(d) Al in H_2O with hPR

Figure 34: (a) Error in the retrieved thickness of three sets of spectroscopic measurements. The measurements in each set were chosen to have nearly equal SNR. The size of each point is related to the size of the energy bin corresponding to the measurement. (b) Error in the retrieved thickness after implementing spectral correction.

Figure 34c and 34d show the results corresponding to the hPR equation for data with air and water as the background material, respectively. The former shows the most accurate estimates of

$T(\boldsymbol{r}_\perp)$ in this study. This is likely due to the fact that the hPR equation and Eq. 68 are Fourier-transform analogs if one assumes a single-material object in the latter. Therefore, the hPR equation can be viewed as an exact inverse to the Al rod in air data processed with the steps described in Sec. 5.5.1 and 5.5.2.

The hPR was far less successful for the case of two materials (Fig. 34d). As described in Sec. 5.3, the average properties of the two materials were used to recover the projected thickness with hPR. The sample was treated as a slab with a constant thickness of 1.53 cm and the average percent error was calculated over the area of the image containing the projection of the rod. When the properties of the rod approach those of the background material (i.e. a homogeneous phantom), the error found in this way should decrease to a minimum.

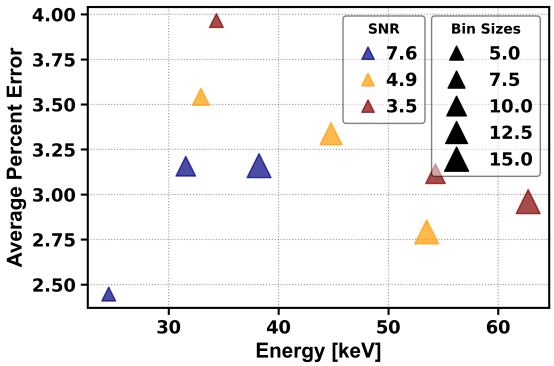

Figure 35: Average percent error in the projected thickness of the Al rod after using the two-material PR equation (Eq. 81 in Appx. A) on our data with water in the background.

Highly accurate results for the data of the rod surrounded by water can be obtained with the PR algorithm proposed by [84]. Similar to how the homogeneous PR in [50] represents the "exact" solution to a single-material phantom (i.e. Al surrounded by air), the two-material PR in [84] is the exact inverse of the experimental two-material phantom (e.g. Al submerged in water). In other words, the two-material phase retrieval in [84] is the Fourier-transform analog of Eq. 63 without any further simplifications. Fig. 35 shows the errors obtained after applying the two-material PR equation to our data with Al surrounded by water. These results show the lowest errors achieved for this phantom.

To conclude, both Fig. 33a and 34c show trends with energies that are clearly similar. Since the algorithms are fundamentally different, the similarities in the two error trends may support the view that they are caused by characteristics of the data, like differences in SNR and others not considered here. We also acknowledge that the results in Fig. 34d are somewhat arbitrary since the material properties could have been defined in a way different than Eq. 64. Finally, errors in Fig. 34c are lower than those in Fig. 35, even though the two were processed with PR algorithms based on "approximation-free" versions of the TIE. The opposite was observed for the data processed with the wPR (Fig. 33a and 33b). These trends could be evidence that reducing the material contrast benefits the weak object approximation.

5.6 Summary and conclusions

This section discussed how the weak object approximation can influence PR algorithms that rely on its accuracy. A good example is the algorithm presented in Sec. 3.1. Unlike the PR techniques tested here, the algorithm in Sec. 3.1 can handle more than three materials and does not require prior knowledge of the sample properties. In Sec. 5.1, an analytical representation for the weak object approximation was derived and analyzed for its validity. To achieve this, we assumed a two-material object built with basic shapes and invoked the projection approximation. Our analysis produced an explicit relationship between the approximation and the properties of the object, namely the thickness and LAC. Our results demonstrate how this approximation becomes less valid when the thickness of the object and/or the contrast between the LAC of the object and the background material increases.

When the TIE is modified with a simplifying assumption, the predicted intensity can deviate from that of the "full" or "approximation-free" TIE. If we consider the full TIE as an accurate description of the physics governing the formation of phase contrast images, then learning how approximations modify the intensity can aid in understanding some of the errors in PR because PR algorithms are generally derived as inverses to simplified-versions of the TIE. In Sec. 5.2, this logic was used to conduct a numerical analysis of the intensities predicted by the TIE with and without

the weak object approximation. The results showed that reducing the difference between the LAC of the object and the background was beneficial and led to a smaller discrepancy between the intensities. A less intuitive result indicated that increasing the object's thickness can lead to lower errors in PR because attenuation contrast, which becomes increasingly dominant as the thickness grows, is properly described by the TIE even after enforcing the weak object approximation. These results corroborate the fact that PB-PCI is most effective for small and poorly-absorbing objects. In these cases, phase-related effects can be more relevant, and differences between the measured intensity and the intensity expected after enforcing the weak object approximation will become more influential.

In this work, the impacts of homogeneous [50] and phase-attenuation duality [51] assumptions on PR were also considered. These approximations are used to derive the two most popular PR techniques today and thus have been extensively scrutinized in the literature [49]. Therefore, our findings are viewed as additions to the existing body of work. We tested these two approximations using existing algorithms with experimental data collected in our laboratory. For the weak object approximation, we derived a new PR algorithm with it as the only simplification to what we considered the full TIE (Eq. 61).

Our measurements were collected using an object built to closely resemble the phantom in Sec. 5.2. The object was imaged in a PB-PCI geometry using a microfocus X-ray source and a photon counting detector. Because of the similarities between the physical and numerical phantom, evidence of the trends was expected to be obtained in our numerical analysis. However, the results showed that errors introduced by spectral distortions dominated and were likely masking other trends. We mitigated spectral distortions with the steps described in Sec. 5.5.1 and 5.5.2, which ensured that our measurements conformed with the full TIE (Eq. 61).

Upon applying the wPR equation to the corrected data and analyzing the error in $T(\boldsymbol{r}_\perp)$, the expected improvements with increasing energies were absent. A possible explanation may be that errors introduced by unaccounted factors, including source-induced blurring and differences in data quality, were more significant. Interestingly, the errors decreased after reducing the contrast

between the target and background material. In particular, improvements were seen after changing the background from air to water, which agrees with our analytical prediction. In general, our findings show that PR algorithms enforcing the weak object approximation can recover material properties with high ($> 95\%$) accuracy.

The errors from the PAD-PR equation were more predictable. The assumption that the LAC can be properly described by only considering Compton scattering becomes more valid with increasing energy, which leads to smaller errors. The relatively large magnitude of the errors can be explained by the properties of our sample. We used a target sample composed of Al, which has an atomic number of 13 and falls above the value of 10 recommended for PAD-PR. Also, the effective energies of our measurements are all below ~ 60 keV, where Compton scattering accounts for only ~ 50 % of the total LAC.

Not surprisingly, the hPR algorithm produced exceptionally low errors for the single-material variation of our phantom (Al surrounded by air). Although we worked with experimental data, our processing forced measurements of our single-material phantom into a form for which hPR was close to an exact inverse. This was confirmed by processing our data with water in the background using the algorithm proposed by [84]. This, in principle, can be tuned to represent the exact inverse of the two-material phantom data. In both cases, the lowest errors were obtained for their respective phantoms. Not surprisingly, the hPR equation gave large errors for the case of the two-material phantom.

In future studies an effort should be made to ensure a nearly constant quality across all of the measurements. Furthermore, the noise should be pushed to the lowest practical level by acquiring for longer times. Source-related blur may be accounted for with a strategy similar to the one presented in [91]. Finally, more variations of target-background materials can help support our findings.

6 Conclusions

As demonstrated in this work, X-ray phase-contrast (PC) images can be collected with micro-focus polychromatic sources and energy-resolving detectors. The enhanced contrast inherent to these images can benefit tasks where the goal is to distinguish materials with similar properties. Such materials are generally difficult, or impossible, to differentiate with conventional (absorption-based) techniques.

A capable phase retrieval (PR) algorithm is needed to characterize materials using PC images since PC images provide limited quantitative information. To this end, our team proposed a spectral phase retrieval (SPR) technique, c.f. Sec. 3.1. This work demonstrates how this PR method can help discriminate materials in heterogeneous samples. Our results show exceptional quantitative accuracy. For this reason, we believe that a possible application for our method would be the quantitative imaging of soft tissue.

One of our primary objectives was to understand how the SPR technique is influenced by factors such as noise, errors in user-inputs, and the breakdown of approximations made during its derivation. With theoretical and computational analyses, we uncovered how noise was amplified in the retrieved quantities. This allowed us to propose, and subsequently investigate, strategies to mitigate noise. We also discovered the effect of misrepresenting the energy of measurements during PR and saw evidence of the impact that detector-related spectral distortions can have on our estimates.

A major characteristic of our PR method is the ability to decompose spectral PC measurements into a pair of material-basis [6]. Material-basis decomposition is not new. It was first proposed in the late 1970s for conventional contact imaging, and its advantages are well documented in the literature. These include the ability to mitigate beam-hardening artifact and virtual removal of contrast agents. What our method effectively achieves is an extension of this powerful technique to propagation-based phase-contrast imaging (PB-PCI). In doing so, and as validated here, it gives a way to handle materials with higher degrees of heterogeneity than other popular alternatives.

Moving forward we will systematically increase the complexity of our samples. Our first focus will be the treatment of highly heterogeneous samples, that is samples with some constituent materials covering a wide range of effective atomic numbers. We will also concentrate on expanding our findings into phase-contrast computed tomography.

A Two-material Phase Retrieval: Derivation for Point-Sources

An extension of the popular of the popular homogeneous PR equation [50] was proposed by [84] and is given here for completeness. The reader may view this appendix as complementary to the material in Sec. 5. We also show how the two-material PR equation reduces to the homogeneous PR under the assumption of a single material, which means that the derivation applies to both results. The original works, that is [50] and [84], contain details beyond those presented here, including mentions of all simplifying assumptions. Another related work is [92], where the two-material algorithm is carefully derived.

Consider a three-component object for which the complex index of refraction can take one of three values, namely

$$n = \begin{cases} 1 - \delta_E^{(1)} + i\beta_E^{(1)} & \text{inside material–1} \\ 1 - \delta_E^{(2)} + i\beta_E^{(2)} & \text{inside material–2} \\ 1 & \text{otherwise,} \end{cases} \tag{73}$$

where δ and β are defined in Sec. 1.2. The subscript E is used to imply an energy dependence while the superscript "(j)", with $j \in \{1, 2\}$, is used to differentiate between the two materials. As in Sec. 5.1, let the combined projected thickness, $T(\boldsymbol{r}_\perp)$, of material–1 and material–2 satisfy $A(\boldsymbol{r}_\perp) = T_1(\boldsymbol{r}_\perp) + T_2(\boldsymbol{r}_\perp) \approx$ constant. Under the validity of the projection approximation we can write the intensity and phase at the exit plane ($z = 0$) as:

$$I_E(\boldsymbol{r}_\perp, 0) = I_E^{\text{in}} \exp\Big\{-\mu_E^{(1)} A - \Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp)\Big\} \tag{74a}$$

and

$$\varphi_E(\boldsymbol{r}_\perp, 0) = -k_E \left[\delta_E^{(1)} A + \Delta\delta_E T^{(2)}(\boldsymbol{r}_\perp)\right], \tag{74b}$$

respectively [12]. Here we have that $\Delta\delta_E = \delta_E^{(2)} - \delta_E^{(1)}$ and $\Delta\mu_E = \mu_E^{(2)} - \mu_E^{(1)}$ while k_E represents the wavenumber. Note that, since $A(\boldsymbol{r}_\perp) \approx$ constant, the spatial dependence was excluded. Inserting

these definitions into the TIE (c.f. Sec. 2.2.3) and rearranging the terms we get

$$-k_E \frac{\partial}{\partial z} I_E(\boldsymbol{r}_\perp, z)\bigg|_{z=0} = I_E^{\text{in}} \exp\Big\{-\mu_E^{(1)} A\Big\} \Delta\delta_E \boldsymbol{\nabla}_\perp \cdot \Big[\exp\Big\{-\Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp)\Big\} \boldsymbol{\nabla}_\perp T^{(2)}(\boldsymbol{r}_\perp)\Big] \tag{75}$$

From the identity $(-\Delta\mu_E)^{-1}\boldsymbol{\nabla}_\perp \exp\{-\Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp)\} = \exp\{-\Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp)\}\boldsymbol{\nabla}_\perp T^{(2)}(\boldsymbol{r}_\perp)$ and using $\boldsymbol{\nabla}_\perp \cdot \boldsymbol{\nabla}_\perp = \boldsymbol{\nabla}_\perp^2$ it follows that Eq. 75 can be written as:

$$-k_E \frac{\partial}{\partial z} I_E(\boldsymbol{r}_\perp, z)\bigg|_{z=0} = -I_E^{\text{in}} \exp\Big\{-\mu_E^{(1)} A\Big\} \frac{\Delta\delta_E}{\Delta\mu} \boldsymbol{\nabla}_\perp^2 \exp\Big\{-\Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp)\Big\}. \tag{76}$$

The z-derivative on the left-hand side of Eq. 76 can be expanded using a finite difference approximation, which is appropriate in the near-field regime, so that:

$$\frac{I_E(\boldsymbol{r}_\perp, R_2) - I_E(\boldsymbol{r}_\perp, 0)}{R_2} = -I_E^{\text{in}} \exp\Big\{-\mu_E^{(1)} A\Big\} \frac{\Delta\delta_E}{\Delta\mu} \boldsymbol{\nabla}_\perp^2 \exp\Big\{-\Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp)\Big\}, \tag{77}$$

where R_2 is the separation (measured along the z-axis) between the illuminated object and the detector. Inserting the definition of $I_E(\boldsymbol{r}_\perp, 0)$ given by Eq. 74a and rearranging the results, we get that

$$\exp\Big\{\mu_E^{(1)} A\Big\} \frac{I_E(\boldsymbol{r}_\perp, R_2)}{I_E^{\text{in}}} = \left[1 - R_2 \frac{\Delta\delta_E}{\Delta\mu_E} \boldsymbol{\nabla}_\perp^2\right] \exp\Big\{\Delta\mu_E T^{(2)}(\boldsymbol{r}_\perp)\Big\}. \tag{78}$$

Until now, the formalism only considers planar radiation, for example X-rays due to a point source an infinite distance away from the illuminated object. Luckily, some trivial substitutions [12] are all that is needed to properly account for a point source some finite distance R_1 away from the object. Namely, if we let $\boldsymbol{\xi}[= M\boldsymbol{r}_\perp]$ represents the coordinates at the detector-plane, then:

$$\exp\Big\{\mu_E^{(1)} A\Big\} \frac{I_E(\boldsymbol{\xi}, R_2)}{I_E^{\text{in}}} = \left[1 - M R_2 \frac{\Delta\delta_E}{\Delta\mu_E} \boldsymbol{\nabla}_\perp^2\right] \exp\Big\{\Delta\mu_E T^{(2)}(\boldsymbol{\xi}/M)\Big\}. \tag{79}$$

where $M[= 1 + R_2/R_2]$ is the geometrical magnification factor [53]. Please note that here we assume $I_E(\boldsymbol{\xi}, R_2)$ and I_E^{in} to be measured at the detector-plane.

From the properties of the Fourier transform, we can represent Eq. 79 in $\boldsymbol{\nu}$-domain as [50]

$$\mathcal{F}\left\{\exp\left\{\mu_E^{(1)}A\right\}\frac{I_E(\boldsymbol{\xi},R_2)}{I_E^{\text{in}}}\right\}=\left[1+MR_2\frac{\Delta\delta_E}{\Delta\mu_E}|\boldsymbol{\nu}|^2\right]\mathcal{F}\left\{\exp\left\{\Delta\mu_E T^{(2)}(\boldsymbol{\xi}/M)\right\}\right\}, \tag{80}$$

where $\boldsymbol{\nu}$ are the frequency coordinates conjugate to $\boldsymbol{\xi}$. After rearranging and applying the inverse Fourier transform, the thickness of material–2 is given by:

$$T^{(2)}(\boldsymbol{\xi}/M)=-\frac{1}{\Delta\mu_E}\ln\mathcal{F}^{-1}\left\{\frac{\mathcal{F}\left\{\frac{I(\boldsymbol{\xi},R_2)}{I_E^{\text{in}}}\exp\left\{\mu_E^{(1)}A\right\}\right\}}{MR_2|\boldsymbol{\nu}|^2\frac{\Delta\delta_E}{\Delta\mu_E}+1}\right\} \tag{81}$$

A.1 Homogeneous phase retrieval equation

The steps outlined above lead to the homogeneous PR equation [50] if all quantities related to material–1, that is $\delta^{(1)}$, $\mu^{(1)}$, and $T^{(1)}$, are zero. Namely, if we substitute $\mu^{(1)}=0$ along with $\Delta\delta_E=\delta_E^{(2)}$ and $\Delta\mu_E=\mu_E^{(2)}$ into Eq. 81, then

$$T(\boldsymbol{\xi}/M)=-\frac{1}{\mu_E}\ln\mathcal{F}^{-1}\left\{\frac{\mathcal{F}\left\{\frac{I(\boldsymbol{\xi},R_2)}{I_E^{\text{in}}}\right\}}{MR_2|\boldsymbol{\nu}|^2\frac{\delta_E}{\mu_E}+1}\right\}, \tag{82}$$

where the superscript "(2)" was dropped. Both equations, 81 and 82, require prior knowledge of the complex refractive index of the sample to give accurate results.

B Hybrid Phase Retrieval

The aim of this appendix is to describe a two-step (or hybrid) phase-retrieval approach, dubbed hyPR, to estimate the projected electron density (a_2) and the energy-independent part of the photoelectric absorption (a_1). Although these quantities can be recovered with the algorithm presented in Sec. 3.1, the workflow here gives results with better quality. A drawback of hyPR, however, is the potential for results that are highly subjective. Namely, the workflow allows for four stages of filtering and the choices of implementation can vary substantially between different users. Yet, it is possible to implement the workflow without any user-inputs, but this may only lead to marginal improvements in some cases. Another option to minimize user-bias is a data-driven approach where each filtering stage is only influenced by the standard deviation of the noise in the input data.

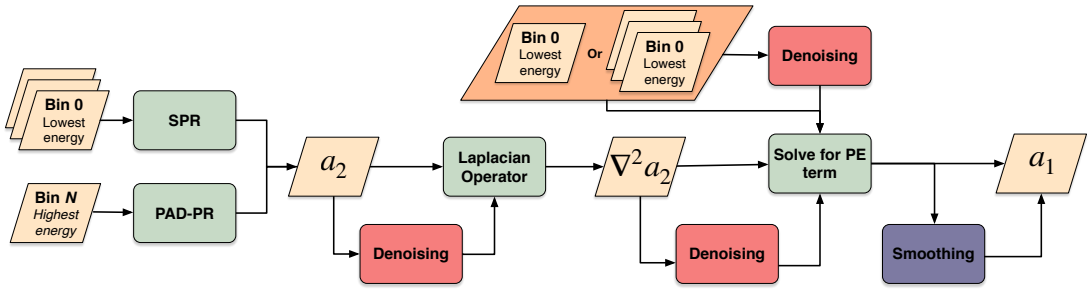

Figure 36: Hybrid phase retrieval workflow. In the first step a_2 is estimated. Next a numerical Laplacian operation is used to obtain approximate $\boldsymbol{\nabla}^2 a_2$. Finally, the a_1 term is estimated from one or multiple bins. A series of optional denoising steps as well as a final smoothing of a_1 can be applied to improve the result.

A breakdown of the hyPR workflow is shown in Fig. 36. The first step involves using the spectral phase retrieval (SPR) algorithm of Sec. 3.1 or the phase-attenuation duality (PAD-PR) technique proposed by [51], see Sec. B.1. An advantage of the latter is higher resilience to noise when only a pair of spectroscopic measurements is available. This is related to the fact that applying PAD-PR is analogous to applying a low-pass filter. As a consequence, hyPR with PAD-PR may be better suited for dual-energy applications. We will elaborate on this claim in a later part of this appendix.

The second step of the hyPR workflow involves the calculation of the Laplacian of a_2, which

can be done in space- or frequency-domain. The space domain approach involves a convolution with the kernel

$$k_L = \Delta^{-2} \begin{bmatrix} 0 & 1 & 0 \\ 1 & 4 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \tag{83}$$

where Δ is the detector pitch. This is equivalent to using a *mid-point finite difference approximation* to the second derivative twice, once for each coordinate direction.

The Laplacian operator can be preceded (and followed) by the application of edge-preserving denoising. The current method of choice is the BM3D routine, which is briefly described in Sec. 3.4.2. As previously shown, see Sec. 3.4.2, this sophisticated routine provides a capable denoising strategy with excellent preservation of edges and sharp signals [60]. Also, the only input required by this algorithm is the standard variation of the noise in the data, which is estimated inside of an ROI defined by the user.

The third non-optional step involves solving for a_1 and it requires two inputs: the Laplacian of a_2 and one (or more) energy bin(s). It is possible, at least in principle, to estimate a_1 with the same bin used to estimate a_2 with PAD-PR. Although this approach was not tested, we expect the resulting a_1 term to display an unacceptable level of noise. In space-domain, the a_1 term is found using:

$$a_1(\boldsymbol{\xi}/M) = \frac{1}{f_p(E_i)} \left[d(\boldsymbol{\xi}, E_i) - f_c(E_i) a_2(\boldsymbol{\xi}/M) + f_{pc}(E_i) \boldsymbol{\nabla}^2 a_2(\boldsymbol{\xi}/M) \right], \tag{84}$$

where $f_{pc}[= M r_e R_2 h^2 c^2 / 2\pi E^2]$ is the energy dependence of the phase contrast term, see Eq. 61–63. The term $d(\boldsymbol{\xi}/M; E_i)$ represents the negative natural log of the flat-field corrected input bin with effective energy E_i and the functions $f_c(E_i)$ and $f_p(E_i)$ are defined in Sec. 1.1.1. Note that it is possible to denoise $d(\boldsymbol{\xi}, E_i)$ prior to solving for a_1.

After applying Eq. 84 we obtain an estimate of a_1. If the quality of the resulting a_1 term is not satisfactory, then the term can be convolved with a smoothing kernel. Currently, the only option for this is *Gaussian smoothing*. Like the BM3D denoising filter, the Gaussian kernel only needs the specification of a standard deviation, which can be defined to be equal to the the standard

deviation of the noise in the data.

Figures 37 show the pLAC reconstructed with the a_1 and a_2 terms, see Eq. 34, found using the HyPR workflow. Each image shows a graphite (left) and a POM (right) rod. The pLAC were calculated for energies of 20, 50, and 80 keV. This allows us to visualize a region where PE (or a_1) dominates the pLAC, i.e. at 20 keV, as well as a region where CS (or a_2) is the dominant term. The first row corresponds to the path in the workflow without denoising or smoothing, see Fig. 36. For the second row, every filter was used and the standard deviation of the noise was found automatically using an ROI on the air region to the left of the first rod. The last row shows the result of making user-defined choices at every filtering stage. While the first two rows are virtually independent of use-input, the last row represents one of many possible outcomes.

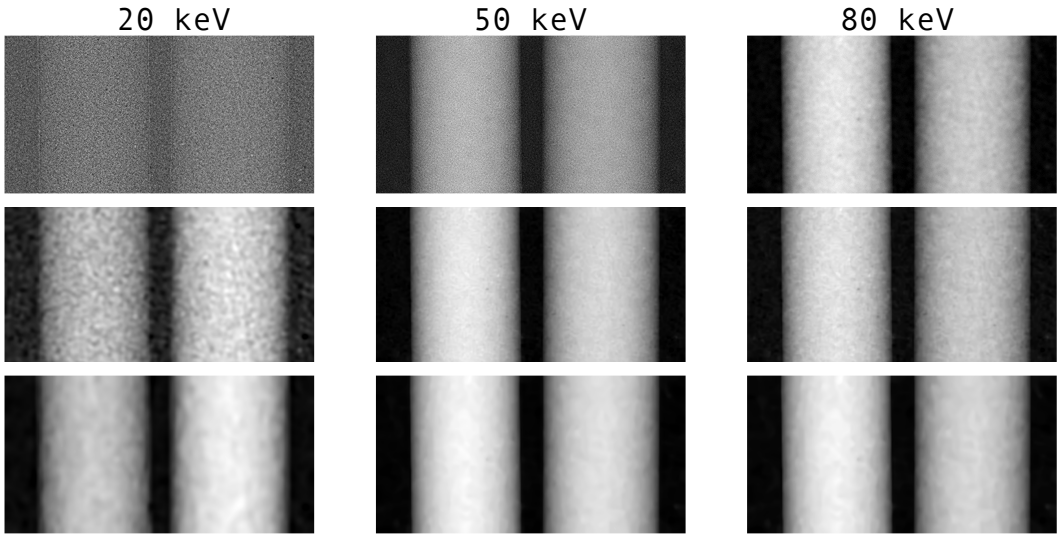

Figure 37: Projected linear attenuation coefficients of graphite (left) and POM (right) rod estimated from outputs of the HyPR workflow. The first row shows the result of not using any denoising or smoothing. The second row shows the result of applying denoising and smoothing based on the standard deviation of the noise in the input data at each filtering stage. The images in the last row were found using user-defined values for the standard deviations.

B.1 Phase-attenuation duality algorithm

The main assumption in the phase-attenuation duality algorithm is that the phase and the linear attenuation can both be written in term of the projected electron density at the object plane [51].

In particular, this means that the intensity can be written as:

$$I_{E_h}(\boldsymbol{r}_\perp, 0) \approx I_{E_h}^{\text{in}} \exp\{-\Phi_{E_h} a_2(\boldsymbol{r}_\perp)\} \tag{85}$$

where Φ_{E_h} is the Klein-Nishina function and E_h is an energy above 60 keV. If the phase is defined as $\varphi_{E_h}(\boldsymbol{r}_\perp) = -r_e \lambda_{E_h} a_2(\boldsymbol{r}_\perp)$ and the steps outlined in Appx. A are followed, it can be shown that

$$a_2(\boldsymbol{\xi}/M) = -\frac{1}{\Phi_{E_h}} \ln \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\left\{ \dfrac{I_{E_h}(\boldsymbol{\xi}, R_2)}{I_{E_h}^{\text{in}}} \right\}}{1 + \dfrac{M R_2 h^2 r_e c^2}{2\pi E_h^2 \Phi_{E_h}} |\boldsymbol{\nu}|^2} \right\}. \tag{86}$$

Here r_e is the classical electron radius, h is Plank's constant, and c is the speed of light in vacuum. All other variables have definitions given in Appx. A. For Eq. 86 to be accurate, it is expected that the components of the sample of interest will have effective atomic numbers below 9.

C Simulation Framework for PCI Measurements

In this work, PCI measurements were simulated using two approaches. For measurements that exactly matched the values expected by the spectral PR method of Sec. 3.1, we used its conjugate forward model. In other words, we simulated the intensity as

$$I_E(M\boldsymbol{r}_\perp) = I_E^{\text{in}} \exp\left\{ \frac{k}{E^m} Z^n \rho_e T(\boldsymbol{r}_\perp) + \Phi_E \rho_e T(\boldsymbol{r}_\perp) - \frac{h^2 c^2 r_e R_2}{2\pi M E^2} \boldsymbol{\nabla}_\perp^2 \rho_e T(\boldsymbol{r}_\perp). \right\} \tag{87}$$

Note that this equation assumes a homogeneous object with a projected thickness given by $T(\boldsymbol{r}_\perp)$ and requires knowledge of several material-dependent constants. In particular, we had to specify ρ_e, Z, k, n, and m.

While the values of Z and ρ_e can be found in standard tables, the same is not true for the constants in the PE term, i.e. k, n, and m. However, it is possible to estimate these constants by minimizing the difference between the PE model in Eq. 87, i.e.

$$\tau_{_E} = \frac{k}{E^m} Z^n \rho_e, \tag{88}$$

and the tabulated values for the PE contribution to the pLAC in the NIST (or similar) database. The model in Eq. 88 can be linearized by computing the natural log of both sides such that

$$\ln \tau_{_E} - \ln \rho_e = \ln k - m \ln E + n \ln Z. \tag{89}$$

It follows that estimates of k, n, and m can be found by minimizing the cost function

$$J(m,n,k) = \frac{1}{2m} \sum_{i=1}^{m} \left[(\ln k - m \ln E_i + n \ln Z) - \left(\ln \tau_{E_i}^{\text{ST}} - \ln \rho_e^{\text{ST}} \right) \right]^2, \tag{90}$$

where the superscript ST is used to indicate values from standard tables.

We minimized Eq. 90 for PS, PMMA, POM, water, and PTFE. Table 6 lists the resulting constants along with tabulated values for ρ_e, Z, and the mass density ρ_m. The material are listed

in increasing order relative to their Z value. The exponents of the energy, m, are in excellent agreement with the findings in [9], where $m = 3.30$ for Carbon-based substances. To test our simulation approach, we implemented SPR with two or more simulated noise-free measurements of different energies. In every case, and irrespective of our particular choices of energies, we found virtually zero error between the expected and retrieved properties.

Material	k	m	n	ρ_e^* [e/cm^3]	Z^*	ρ_m^* [g/cm^3]
PS	2.75×10^{-22}	3.33	2.23	3.40×10^{23}	5.70	1.050
PMMA	5.89×10^{-23}	3.30	3.12	3.83×10^{23}	6.47	1.180
POM	2.28×10^{-23}	3.29	3.61	4.52×10^{23}	6.95	1.410
Water	9.17×10^{-24}	3.29	4.05	3.34×10^{23}	7.42	0.997
PTFE	1.39×10^{-24}	3.27	4.87	6.36×10^{23}	8.43	2.200

Table 6: Constants for the PE model found by minimizing Eq. 90 along with tabulated material properties. All properties marked with an asterisk correspond to tabulated properties.

The second type of simulations used in this work is described in Sec. 5. Unlike Eq. 87, which was found from a simplified version of the TIE, the simulation equation in Sec. 5, i.e. Eq. 61, is obtained from the full TIE in the near-field regime. Nevertheless, the properties of the objects simulated with Eq. 61 must satisfy several restrictive conditions. For example, the sample can be composed of at most two different materials and each component must have a simple shape with a simple representation.

C.1 Simulating noisy data

Noise was added to our simulated data using a pseudo-random number generator with values distributed according to a Poisson distribution. A Poisson distribution is completely described with a single value corresponding to its mean, which also defines the variance. Thus, we used the intensity at each pixel as the mean of a Poisson probability density function and realized a new (random) value. To make the noise in our simulations reflect the energy-dependence observed in experimental data, we simulated a realistic X-ray spectrum [90] and used the counts at each energy to define the counts of the incident (flat) fields.

Figure 38 shows three noisy mono-energetic PCI measurements simulated with the approach outlined above. The X-ray spectrum used to estimate the counts in the flat-field is also shown. Polychromatic measurements, i.e. those corresponding to bins wider than one keV, were simulated as the weighted sum of monochromatic measurements contained in the bin, where the weights were defined using the counts in the flat-field.

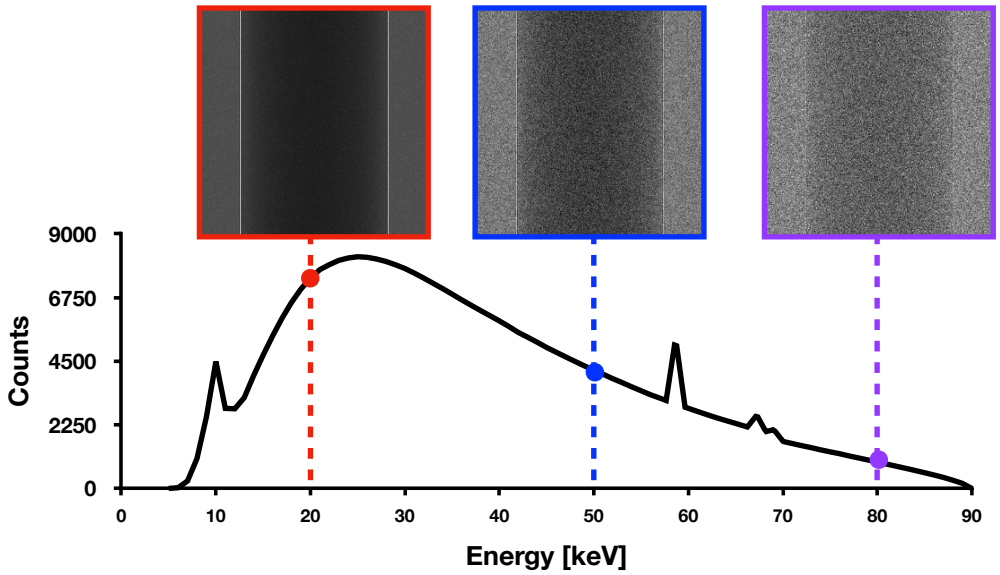

Figure 38: Three flat-field corrected monochromatic measurements simulated at different energies and displaying unique amounts of noise. The noise is Poisson distributed.

D Spectral correction

We implemented a simulation-based approach to improve the accuracy of the recorded intensity values, which were affected by spectral distortions and other sources of error. In our case, spectral distortions were primarily caused by detector-related effects like charge sharing, K-escape, and pulse-pileup [70, 89]. Evidence of beam-hardening was also observed. Like detector-related distortions, beam-hardening is also an energy-dependent phenomena.

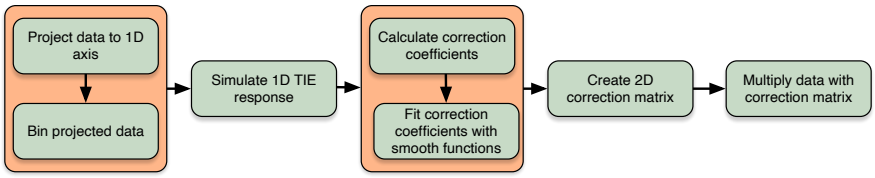

Figure 39: Workflow used to correct spectral distortions.

The "spectral correction" approach is outlined in Fig. 39. The first step is to obtain an accurate 1-D representation of the 2-D data, which is achieved with the "reprojection" method proposed by [93]. Next, a 1-D PCI response is simulated in space-domain with the exact TIE equation (Eq. 61–63). This is possible because of the simple shape of the phantoms. A related, but more general, spectral correction approach is proposed in [88].

After the simulation, a set of coefficients characterizing the mismatch between the 1D representation of the data and the simulated PCI response are computed as the ratio of the two. Since the coefficients reflect the noise in the data, we fit them with a smooth function. The choice of fitting function depends on several factors. For example, some low-Z materials have nearly constant correction coefficients that can be properly represented with lines. On the other hand, materials like aluminum can require piecewise functions to properly account for the differences in the correction coefficients of the air and sample regions. A possible explanation for this is beam-hardening. Namely, thin materials with low-Z will not introduce significant distortions so that the errors in the air and material regions are similar. By contrast, the relatively high attenuation of Al makes it more prone to beam-hardening and the variations in flux between the air and material region can be substantial.

The final two steps involve creating a correction matrix and, subsequently, applying the correction matrix to the data. The information obtained during reprojection is used to achieve this. In particular, the fitted coefficients are analyzed at the (non-uniform) position of every pixel along the 1-D reprojection axis. The coefficients corresponding to every pixel position are then multiplied by the original intensities to obtain a better match of the simulated data. The use of smooth coefficients minimizes the influence of our approach to the noise texture in the original data. We also ensured that changes at material boundaries were minimal by disallowing significant corrections at, and in close proximity to, the edges of the rods.

Figure 40 shows two low-energy (24 and 26 keV) bins of an Al and PMMA rod before and after spectral correction. The workflow in Fig. 39 was used to correct the original. As the appreciable difference between the original and corrected profile in Fig. 40a shows, the Al rod at this energy was subject to severe spectral distortions. Interestingly, distortions along the profile of the PMMA rod were minimal.

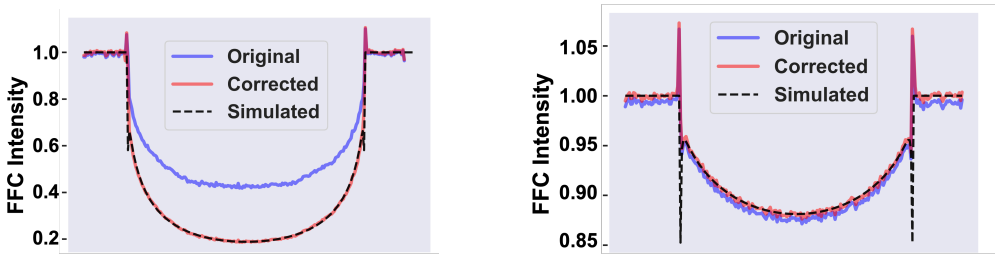

(a) Al at $\sim$ 24 keV **(b)** PMMA at $\sim$ 26 keV

Figure 40: (a) Original and corrected profile of a PCI measurement of an aluminum rod. The measurement corresponds to an effective energy of 24 keV. A dashed line shows the expected response for a bin bounded at the same energy values. (b) Original and corrected profiles of a PMMA bin having an effective energy of 26 keV.

The method here is specific to the samples used we used. Nevertheless, the underlying approach achieves the same objective as the popular method in [88]. Namely, distorted intensity values are transformed into more accurate representations agreeing with those predicted by the relevant physical model. In their work, [88] use Beer's law to predict accurate intensity values. Here we used the TIE, which is more proper for PCI measurements. The results in Sec. 5.5.2 support the effectiveness of the spectral correction scheme described here.

E Noise Power Spectrum

The noise power spectrum (NPS) gives the spectral decomposition of the noise in an image, which conveys the noise power at each frequency [4]. Figure 41 illustrates the workflow used to obtain the NPS results shown in this work. This method of estimating the NPS is a 2-D extension of Welch's method for spectral density estimation [94].

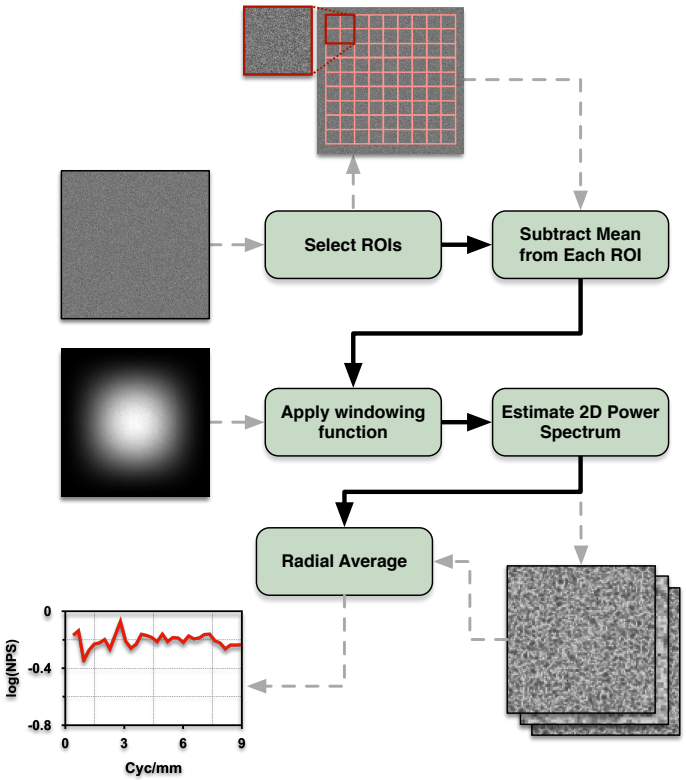

Figure 41: Workflow used to estimate the noise power spectrum.

In the first step, a set of overlapping ROIs are defined. For this, the user must define three parameters: the number, the size, and the overlap percentage of each ROI. Next, the mean value of each ROI is found and subtracted from the enclosed pixels. Spatial tapering windows are applied to each ROI in the third step. A common choice is the Hanning window, which is used in this work. Finally, the 2-D power spectrum of each tapered ROI is estimated and, in turn, the radial average of the combined results is found. This end product is a 1-D representation of the NPS.

F Strategy to Estimate the Effective Energy

The aim here is to describe a strategy that gives a highly accurate estimate of the effective energy of non-monochromatic measurements (or bins). Non-monochromatic refers to measurements corresponding to X-ray energies ranging over regions of more than one keV. A common way to define the effective energy of a bin is with the average energies that correspond to the bin. Another common and, as we will show, more accurate method is to use the weighted average. In practice, however, this requires knowledge of the source's spectrum to define the weights. The method advocated here is a simulation-based method applicable to PCI measurements and it produces results more accurate than the weighted-mean approach.

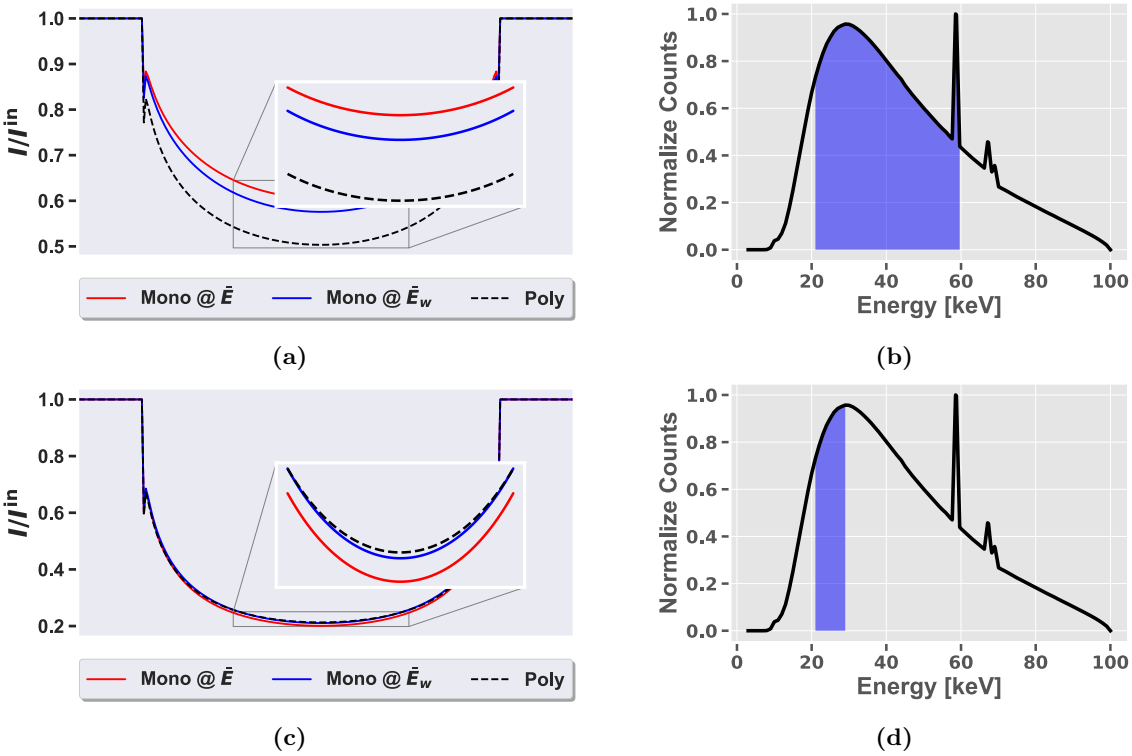

Figure 42: (a) Polychromatic measurement (dashed line) simulated for an energy bin covering the shaded region in (b). A red and blue curve indicate the monochromatic measurements found with the average and the weighted-average energy of the bin, respectively. (b) Simulated source spectrum and the region used to estimate (a). (c) Measurements for a narrower polychromatic bin. (d) Bin used to estimate the polychromatic measurement in (c).

Figures 42a and 42c show the flat-field corrected intensities simulated using the full TIE (see

Sec. 5.1) for a broad and a narrow bin, respectively. We indicated the simulated measurement using a dashed line in conjunction with the label "Poly". A clear depiction of the simulated source spectrum used as well as the regions matching each of the polychromatic measurements is given in Fig. 42b and 42d. From Fig. 42a and 42c we see that the effective energy found as a weighted average was more accurate in both cases. Perhaps not surprisingly, the average and weighted-average strategies are less accurate when the bin is wider. On the other hand, the monochromatic response for the weighted average is almost equal to the polychromatic measurement when the bin width was 10 keV. The results support the use of narrow bins for ether strategy, especially if the weighted-average can be used.

To improve the accuracy of the effective energy, a monochromatic measurement can be simulated and compared to the polychromatic response. Through an iterative approach, the mismatch between the polychromatic and monoenergetic profiles can be minimized until it falls below a desirable value. Fig. 43 shows this strategy in action. In this case, the initial prediction of the effective energy was 72 keV, which corresponded to 30 % error. The algorithm updated the energy five times until the error was 1 %.

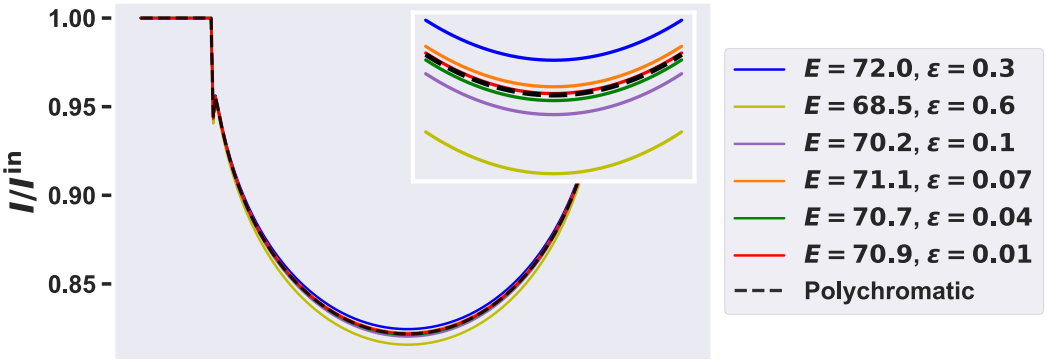

Figure 43: Estimating the effective energy iteratively. The error at each iteration is indicates with ε.

Bibliography

[1] A. Assmus, "Early history of x-rays," *Beam Line*, vol. 25, no. 2, pp. 10–24, 1995.

[2] T. M. Buzug, *Computed Tomography: From Photon Statistics to Modern Cone-Beam CT*. Berlin Heidelberg: Springer-Verlag, 2008.

[3] J. Beutel, H. L. Kundel, and R. L. Van Metter, *Handbook of Medical Imaging: Physics and Psychophysics*, vol. 1. Spie Press, 2000.

[4] J. T. Bushberg and J. M. Boone, *The Essential Physics of Medical Imaging*. Lippincott Williams & Wilkins, Dec. 2011.

[5] J. A. Fessler, M. Sonka, and J. M. Fitzpatrick, "Statistical image reconstruction methods for transmission tomography.," in *Handbook of medical imaging*, vol. 2, pp. 1–70, Spie Press, 2000.

[6] R. E. Alvarez and A. Macovski, "Energy-selective reconstructions in x-ray computerised tomography," *Physics in medicine and biology*, vol. 21, no. 5, pp. 733–744, 1976.

[7] J. Hubbell, "Photon cross sections, attenuation coefficients and energy absorption coefficients," *National Bureau of Standards Report NSRDS-NBS29, Washington DC*, 1969.

[8] J. Hale, *Fundamentals of Radiological Science*. Thomas, 1974.

[9] R. A. Rutherford, B. R. Pullan, and I. Isherwood, "Measurement of effective atomic number and electron density using an EMI scanner," *Neuroradiology*, vol. 11, pp. 15–21, May 1976.

[10] A. M. Alessio and L. R. MacDonald, "Quantitative material characterization from multi-energy photon counting CT," *Medical physics*, vol. 40, no. 3, pp. 031108–031108, 2013.

[11] D. Gürsoy and M. Das, "Single-step absorption and phase retrieval with polychromatic x-rays using a spectral detector," *Opt. Lett.*, vol. 38, pp. 1461–1463, May 2013.

[12] D. M. Paganin, *Coherent X-Ray Optics*. Oxford University Press, Jan. 2006.

[13] J. Als-Nielsen and D. McMorrow, *Elements of Modern X-ray Physics*. John Wiley & Sons, Apr. 2011.

[14] S. Mayo, P. Miller, S. Wilkins, T. Davis, D. Gao, T. Gureyev, D. Paganin, D. Parry, A. Pogany, and A. Stevenson, "Quantitative x-ray projection microscopy: Phase-contrast and multi-spectral imaging," *Journal of microscopy*, vol. 207, no. 2, pp. 79–96, 2002.

[15] S. W. Wilkins, Y. Nesterets, T. E. Gureyev, S. C. Mayo, A. Pogany, and A. W. Stevenson, "On the evolution and relative merits of hard x-ray phase-contrast imaging methods," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 372, pp. 20130021–20130021, Mar. 2014.

[16] H. Wen, "Biomedical x-ray phase-contrast imaging and tomography," in *Springer Handbook of Microscopy* (P. W. Hawkes and J. C. H. Spence, eds.), Springer Handbooks, pp. 1451–1468, Cham: Springer International Publishing, 2019.

[17] D. M. Paganin and D. Pelliccia, "Tutorials on x-ray phase contrast imaging: Some fundamentals and some conjectures on future developments," *arXiv preprint arXiv:1902.00364*, Jan. 2019.

[18] C. H. McCollough, S. Leng, L. Yu, and J. G. Fletcher, "Dual- and multi-energy CT: Principles, technical approaches, and clinical applications," *Radiology*, vol. 276, no. 3, pp. 637–653, 2015.

[19] E. C. McCullough, "Photon attenuation in computed tomography," *Medical Physics*, vol. 2, no. 6, pp. 307–320, 1975.

[20] M. R. Millner, W. H. Payne, R. G. Waggener, W. D. McDavid, M. J. Dennis, and V. J. Sank, "Determination of effective energies in CT calibration," *Medical Physics*, vol. 5, no. 6, pp. 543–545, 1978.

[21] J. Alles and R. F. Mudde, "Beam hardening: Analytical considerations of the effective attenuation coefficient of x-ray tomography," *Medical Physics*, vol. 34, no. 7, pp. 2882–2889, 2007.

[22] P. M. Shikhaliev, "Beam hardening artefacts in computed tomography with photon counting, charge integrating and energy weighting detectors: a simulation study," *Physics in Medicine and Biology*, vol. 50, pp. 5813–5827, Dec. 2005.

[23] D. Vavrik and J. Jakubek, "Radiogram enhancement and linearization using the beam hardening correction method," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 607, pp. 212–214, Aug. 2009.

[24] J. Jakůbek, "Semiconductor pixel detectors and their applications in life sciences," *Journal of Instrumentation*, vol. 4, pp. P03013–P03013, Mar. 2009.

[25] K. Taguchi and J. S. Iwanczyk, "Vision 20/20: Single photon counting x-ray detectors in medical imaging," *Medical Physics*, vol. 40, no. 10, pp. 100901–100901, 2013.

[26] R. Ballabriga, J. Alozy, G. Blaj, M. Campbell, M. Fiederle, E. Frojdh, E. H. M. Heijne, X. Llopart, M. Pichotka, S Procz, L. Tlustos, and W. Wong, "The Medipix3RX: a high resolution, zero dead-time pixel detector readout chip allowing spectroscopic imaging," *Journal of Instrumentation*, vol. 8, no. 02, pp. C02016–C02016, 2013.

[27] P. Cloetens, R. Barrett, J. Baruchel, J.-P. Guigay, and M. Schlenker, "Phase objects in synchrotron radiation hard x-ray imaging," *Journal of Physics D: Applied Physics*, vol. 29, pp. 133–146, Jan. 1996.

[28] S. W. Wilkins, T. E. Gureyev, D. Gao, A. Pogany, and A. Stevenson, "Phase-contrast imaging using polychromatic hard x-rays," *Nature*, vol. 384, no. 6607, p. 335–338, 1996.

[29] R. A. Lewis, "Medical phase contrast x-ray imaging: current status and future prospects," *Physics in Medicine and Biology*, vol. 49, pp. 3573–3583, July 2004.

[30] R. A. Lewis, N. Yagi, M. J. Kitchen, M. J. Morgan, D. Paganin, K. K. W. Siu, K. Pavlov, I. Williams, K. Uesugi, M. J. Wallace, C. J. Hall, J. Whitley, and S. B. Hooper, "Dynamic

imaging of the lungs using x-ray phase contrast," *Physics in Medicine and Biology*, vol. 50, pp. 5031–5040, Oct. 2005.

[31] M. J. Kitchen, R. A. Lewis, M. J. Morgan, M. J. Wallace, M. L. Siew, K. K. W. Siu, A. Habib, A. Fouras, N. Yagi, K. Uesugi, and S. B. Hooper, "Dynamic measures of regional lung air volume using phase contrast x-ray imaging," *Physics in Medicine and Biology*, vol. 53, pp. 6065–6077, Oct. 2008.

[32] A. Bravin, P. Coan, and P. Suortti, "X-ray phase-contrast imaging: from pre-clinical applications towards clinics," *Physics in Medicine and Biology*, vol. 58, pp. R1–R35, Dec. 2012.

[33] S. C. Mayo, A. W. Stevenson, and S. W. Wilkins, "In-line phase-contrast x-ray imaging and tomography for materials science," *Materials*, vol. 5, pp. 937–965, May 2012.

[34] L. C. P. Croton, K. S. Morgan, D. M. Paganin, L. T. Kerr, M. J. Wallace, K. J. Crossley, S. L. Miller, N. Yagi, K. Uesugi, S. B. Hooper, and M. J. Kitchen, "In situ phase contrast x-ray brain CT," *Scientific Reports*, vol. 8, pp. 11412–11412, July 2018.

[35] A. Pogany, D. Gao, and S. Wilkins, "Contrast and resolution in imaging with a microfocus x-ray source," *Review of Scientific Instruments*, vol. 68, no. 7, pp. 2774–2782, 1997.

[36] J. W. Goodman, *Statistical Optics*. John Wiley & Sons, May 2015.

[37] M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, Feb. 2000.

[38] A. Peterzol, J. Berthier, P. Duvauchelle, C. Ferrero, and D. Babot, "X-ray phase contrast image simulation," *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 254, pp. 307–318, Jan. 2007.

[39] E. R. Shanblatt, Y. Sung, R. Gupta, B. J. Nelson, S. Leng, W. S. Graves, and C. H. McCollough, "Forward model for propagation-based x-ray phase contrast imaging in parallel- and cone-beam geometry," *Optics Express*, vol. 27, pp. 4504–4521, Feb. 2019.

[40] U. Lundström, D. H. Larsson, A. Burvall, P. A. C. Takman, L. Scott, H. Brismar, and H. M. Hertz, "X-ray phase contrast for CO2 microangiography," *Physics in Medicine and Biology*, vol. 57, pp. 2603–2617, Apr. 2012.

[41] K. Bliznakova, P. Russo, G. Mettivier, H. Requardt, P. Popov, A. Bravin, and I. Buliev, "A software platform for phase contrast x-ray breast imaging research," *Computers in Biology and Medicine*, vol. 61, pp. 62–74, June 2015.

[42] J. M. Cowley, *Diffraction Physics*. Elsevier, Dec. 1995.

[43] M. R. Teague, "Deterministic phase retrieval: A Green's function solution," *JOSA*, vol. 73, no. 11, pp. 1434–1441, 1983.

[44] T. Vo-Dinh, *Biomedical Photonics Handbook*. CRC Press, Mar. 2003.

[45] P. Henrot, A. Leroux, C. Barlier, and P. Génin, "Breast microcalcifications: The lesions in anatomical pathology," *Diagnostic and Interventional Imaging*, vol. 95, pp. 141–152, Feb. 2014.

[46] L. Brombal, F. Arfelli, P. Delogu, S. Donato, G. Mettivier, K. Michielsen, P. Oliva, A. Taibi, I. Sechopoulos, R. Longo, and C. Fedon, "Image quality comparison between a phase-contrast synchrotron radiation breast CT and a clinical breast CT: a phantom based study," *Scientific Reports*, vol. 9, pp. 1–12, Nov. 2019.

[47] M. J. Kitchen, R. A. Lewis, M. J. Morgan, M. J. Wallace, M. L. Siew, K. K. W. Siu, A. Habib, A. Fouras, N. Yagi, K. Uesugi, and S. B. Hooper, "Dynamic measures of regional lung air volume using phase contrast x-ray imaging," *Physics in Medicine and Biology*, vol. 53, pp. 6065–6077, Oct. 2008.

[48] M. Langer, P. Cloetens, J.-P. Guigay, and F. Peyrin, "Quantitative comparison of direct phase retrieval algorithms in in-line phase tomography," *Medical physics*, vol. 35, no. 10, pp. 4556–4566, 2008.

[49] A. Burvall, U. Lundström, P. A. C. Takman, D. H. Larsson, and H. M. Hertz, "Phase retrieval in x-ray phase-contrast imaging suitable for tomography," *Optics Express*, vol. 19, pp. 10359–10376, May 2011.

[50] D. Paganin, S. Mayo, T. E. Gureyev, P. R. Miller, and S. W. Wilkins, "Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object," *Journal of microscopy*, vol. 206, no. 1, pp. 33–40, 2002.

[51] X. Wu, H. Liu, and A. Yan, "X-ray phase-attenuation duality and phase retrieval," *Optics letters*, vol. 30, no. 4, pp. 379–381, 2005.

[52] M. Das and Z. Liang, "Approximated transport-of-intensity equation for coded-aperture x-ray phase-contrast imaging," *Optics letters*, vol. 39, no. 18, pp. 5395–5398, 2014.

[53] I. Vazquez, N. R. Fredette, and M. Das, "Quantitative phase retrieval of heterogeneous samples from spectral x-ray measurements," in *Medical Imaging 2019: Physics of Medical Imaging* (T. G. Schmidt, G.-H. Chen, and H. Bosmans, eds.), vol. 10948, pp. 1157 – 1164, International Society for Optics and Photonics, SPIE, Apr. 2019.

[54] X. Wu and H. Liu, "Clinical implementation of x-ray phase-contrast imaging: Theoretical foundations and design considerations," *Medical physics*, vol. 30, no. 8, pp. 2169–2179, 2003.

[55] F. H. Attix, *Introduction to Radiological Physics and Radiation Dosimetry*. John Wiley & Sons, 2008.

[56] T. G. Schmidt, "Optimal "image-based" weighting for energy-resolved ct," *Medical Physics*, vol. 36, no. 7, pp. 3018–3027, 2009.

[57] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI Pulse Sequences*. Elsevier, Sept. 2004.

[58] S. W. Smith, *The scientist and engineer's guide to digital signal processing*. California Technical Pub. San Diego, 1997.

[59] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Transactions on Image Processing*, vol. 16, pp. 2080–2095, Aug. 2007.

[60] Y. Mäkinen, L. Azzari, and A. Foi, "Exact transform-domain noise variance for collaborative filtering of stationary correlated noise," in *2019 IEEE International Conference on Image Processing (ICIP)*, pp. 185–189, Sept. 2019.

[61] T. Zhao, J. Hoffman, M. McNitt-Gray, and D. Ruan, "Ultra-low-dose CT image denoising using modified BM3D scheme tailored to data statistics," *Medical Physics*, vol. 46, no. 1, pp. 190–198, 2019.

[62] V. Hanchate and K. Joshi, "MRI denoising using BM3D equipped with noise invalidation denoising technique and VST for improved contrast," *SN Applied Sciences*, vol. 2, pp. 1–8, Jan. 2020.

[63] F. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proceedings of the IEEE*, vol. 66, pp. 51–83, Jan. 1978.

[64] P. R. Mendonça, P. Lamb, and D. V. Sahani, "A flexible method for multi-material decomposition of dual-energy CT images," *IEEE transactions on medical imaging*, vol. 33, no. 1, pp. 99–116, 2014.

[65] N. R. Fredette, A. Kavuri, and M. Das, "Multi-step material decomposition for spectral computed tomography," *Physics in Medicine & Biology*, vol. 64, pp. 145001–145001, July 2019.

[66] V. Shen, D. Siderius, W. Krekelberg, and H. Hatch, "NIST standard reference simulation website," *NIST Standard Reference Database*, no. 173, pp. 2014–2017, 2017.

[67] M. E. Phelps, E. J. Hoffman, and M. M. Ter-Pogossian, "Attenuation coefficients of various body tissues, fluids, and lesions at photon energies of 18 to 136 keV," *Radiology*, vol. 117, pp. 573–583, Dec. 1975.

[68] P. C. Shrimpton, "Electron density values of various human tissues: in vitro compton scatter measurements and calculated ranges," *Physics in Medicine and Biology*, vol. 26, pp. 907–911, Sept. 1981.

[69] E. N. Gimenez, R. Ballabriga, G. Blaj, M. Campbell, I. Dolbnya, E. Frodjh, I. Horswell, X. Llopart, J. Marchal, J. McGrath, D. Omar, R. Plackett, K. Sawhney, and N. Tartoni, "Medipix3RX: Characterizing the Medipix3 redesign with synchrotron radiation," *IEEE Transactions on Nuclear Science*, vol. 62, pp. 1413–1421, June 2015.

[70] J. Cammin, J. Xu, W. C. Barber, J. S. Iwanczyk, N. E. Hartsough, and K. Taguchi, "A cascaded model of spectral distortions due to spectral response effects and pulse pileup effects in a photon-counting x-ray detector for CT," *Medical Physics*, vol. 41, no. 4, pp. 041905–041905, 2014.

[71] J. L. Prince and J. M. Links, *Medical Imaging Signals and Systems*. Pearson Prentice Hall Upper Saddle River, 2006.

[72] R. Ballabriga, J. Alozy, M. Campbell, E. Frojdh, E. H. M. Heijne, T. Koenig, X. Llopart, J. Marchal, D. Pennicard, T. Poikela, L. Tlustos, P. Valerio, W. Wong, and M. Zuber, "Review of hybrid pixel detector readout ASICs for spectroscopic x-ray imaging," *Journal of Instrumentation*, vol. 11, pp. P01007–P01007, Jan. 2016.

[73] C. R. Crawford, "Reprojection using a parallel backprojector," *Medical Physics*, vol. 13, no. 4, pp. 480–483, 1986.

[74] M. J. Kitchen, G. A. Buckley, T. E. Gureyev, M. J. Wallace, N. Andres-Thio, K. Uesugi, N. Yagi, and S. B. Hooper, "CT dose reduction factors in the thousands using x-ray phase contrast," *Scientific reports*, vol. 7, no. 1, pp. 1–9, 2017.

[75] S. Vespucci, C. Lewis, C. S. Park, and M. Das, "Examining phase contrast sensitivity to signal location and tissue thickness in breast imaging," in *Medical Imaging 2018: Physics of Medical Imaging*, vol. 10573, pp. 533 – 541, International Society for Optics and Photonics, SPIE, Mar. 2018.

[76] C. Lewis, I. Vazquez, S. Vespucci, and M. Das, "Contribution of scatter and beam hardening to phase contrast imaging," in *Medical Imaging 2019: Physics of Medical Imaging* (T. G. Schmidt, G.-H. Chen, and H. Bosmans, eds.), vol. 10948, pp. 1274 – 1280, International Society for Optics and Photonics, SPIE, Mar. 2019.

[77] M. Das and Z. Liang, "Spectral x-ray phase contrast imaging for single-shot retrieval of absorption, phase, and differential-phase imagery," *Optics Letters*, vol. 39, pp. 6343–6346, Nov. 2014.

[78] M. Das, H. C. Gifford, J. M. O'Connor, and S. J. Glick, "Penalized maximum likelihood reconstruction for improved microcalcification detection in breast tomosynthesis," *IEEE Transactions on Medical Imaging*, vol. 30, pp. 904–914, Apr. 2011.

[79] H. C. Gifford, Z. Liang, and M. Das, "Visual-search observers for assessing tomographic x-ray image quality," *Medical Physics*, vol. 43, pp. 1563–1575, Mar. 2016.

[80] T. E. Gureyev and S. W. Wilkins, "On x-ray phase retrieval from polychromatic images," *Optics Communications*, vol. 147, pp. 229–232, Feb. 1998.

[81] T. E. Gureyev, Y. I. Nesterets, A. Kozlov, D. M. Paganin, and H. M. Quiney, "On the "unreasonable" effectiveness of transport of intensity imaging and optical deconvolution," *JOSA A*, vol. 34, no. 12, pp. 2251–2260, 2017.

[82] N. Streibl, "Phase imaging by the transport equation of intensity," *Optics Communications*, vol. 49, pp. 6–10, Feb. 1984.

[83] K. M. Pavlov, T. E. Gureyev, D. Paganin, Y. I. Nesterets, M. J. Morgan, and R. A. Lewis, "Linear systems with slowly varying transfer functions and their application to x-ray phase-contrast imaging," *Journal of Physics D: Applied Physics*, vol. 37, pp. 2746–2750, Sept. 2004.

[84] M. Beltran, D. Paganin, K. Uesugi, and M. Kitchen, "2D and 3D x-ray phase retrieval of multi-material objects using a single defocus distance," *Optics Express*, vol. 18, no. 7, pp. 6423–6436, 2010.

[85] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. SIAM, 2001.

[86] L. Waller, L. Tian, and G. Barbastathis, "Transport of intensity phase-amplitude imaging with higher order intensity derivatives," *Optics Express*, vol. 18, pp. 12552–12561, June 2010.

[87] T. E. Gureyev, A. Roberts, and K. A. Nugent, "Partially coherent fields, the transport-of-intensity equation, and phase uniqueness," *JOSA A*, vol. 12, pp. 1942–1946, Sept. 1995.

[88] H. Ding and S. Molloi, "Image-based spectral distortion correction for photon-counting x-ray detectors," *Medical Physics*, vol. 39, no. 4, pp. 1864–1876, 2012.

[89] K. Taguchi, E. C. Frey, X. Wang, J. S. Iwanczyk, and W. C. Barber, "An analytical model of the effects of pulse pileup on the energy spectrum recorded by energy resolved photon counting x-ray detectors," *Medical Physics*, vol. 37, no. 8, pp. 3957–3969, 2010.

[90] G. Hernández and F. Fernández, "A model of tungsten anode x-ray spectra," *Medical Physics*, vol. 43, no. 8Part1, pp. 4655–4664, 2016.

[91] M. A. Beltran, D. M. Paganin, and D. Pelliccia, "Phase-and-amplitude recovery from a single phase-contrast image using partially spatially coherent x-ray radiation," *Journal of Optics*, vol. 20, no. 5, pp. 055605–055605, 2018.

[92] I. Häggmark, W. Vågberg, H. M. Hertz, and A. Burvall, "Comparison of quantitative multi-material phase-retrieval algorithms in propagation-based phase-contrast x-ray tomography," *Optics Express*, vol. 25, no. 26, pp. 33543–33558, 2017.

[93] E. Samei, M. J. Flynn, and D. A. Reimann, "A method for measuring the presampled MTF of digital radiographic systems using an edge test device," *Medical physics*, vol. 25, no. 1, pp. 102–113, 1998.

[94] P. Welch, "The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms," *IEEE Transactions on Audio and Electroacoustics*, vol. 15, pp. 70–73, June 1967.