

THE UNIVERSITY OF CHICAGO

SPECTRAL X-RAY IMAGING AND PHASE RETRIEVAL

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES
AND THE PRITZKER SCHOOL OF MEDICINE
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

COMMITTEE ON MEDICAL PHYSICS

BY
GIAVANNA JADICK

CHICAGO, ILLINOIS
AUGUST 2026

© 2026 by Giavanna Jadick

ORCID iD: <https://orcid.org/0000-0002-4894-2895>

For Ellen.

“A huge percentage of the stuff that I tend to be automatically certain of is, it turns out,
totally wrong and deluded.”

— David Foster Wallace

This Is Water

TABLE OF CONTENTS

LIST OF FIGURES	ix
LIST OF TABLES	xvi
ACKNOWLEDGMENTS	xvii
ABSTRACT	xix
1 INTRODUCTION	1
1.1 Phase contrast	4
1.1.1 Phase retrieval	6
1.2 Spectral imaging	7
1.2.1 The multiple measurement problem	9
1.2.2 Photon-counting detectors	10
1.3 Spectral XPCI and current challenges	11
1.4 Dissertation overview	13
2 X-RAY IMAGING FUNDAMENTALS	15
2.1 Conventional X-ray imaging model	15
2.2 X-ray sources	19
2.2.1 X-ray tubes	19
2.2.2 Synchrotron light sources	20
2.2.3 Beam geometry	21
2.3 X-ray detectors	21
2.3.1 Energy-integrating detectors	22
2.3.2 Photon-counting detectors	23
2.4 Computed tomography	26
2.5 Spectral imaging and basis decomposition	28
2.5.1 Interaction-probability decomposition	28
2.5.2 Material decomposition	31
2.5.3 K-edge imaging	32
2.5.4 Virtual monochromatic images	33
2.6 Summary	34
3 DUAL-ENERGY CT WITH MEGAVOLTAGE AND KILOVOLTAGE SPECTRA	35
3.1 Introduction	35
3.2 Background	36
3.3 Methods	38
3.3.1 Signal detection framework	39
3.3.2 Single-line integral through a two-material object	40
3.3.3 Fan-beam CT simulations of an anthropomorphic phantom	42
3.4 Results	45

3.4.1	Single-line integral	45
3.4.2	CT simulations	48
3.5	Discussion	53
3.6	Conclusion	61
3.7	Appendix: Simulated CT images	61
4	MV-KV DUAL-ENERGY CT WITH A PHOTON-COUNTING DETECTOR . .	70
4.1	Introduction	70
4.2	Background	71
4.3	Methods	74
4.3.1	Materials	75
4.3.2	Single-line integral	76
4.3.3	CT simulations	79
4.4	Results	80
4.4.1	Single-line integral	80
4.4.2	CT simulations	83
4.5	Discussion	84
4.6	Conclusion	89
4.7	Appendix: Spectral detectors with a single-MV spectrum	89
5	PHASE-CONTRAST IMAGING FUNDAMENTALS	94
5.1	Wave model	94
5.2	Free-space propagation	96
5.2.1	Fresnel diffraction	99
5.2.2	Transport-of-intensity equation (TIE)	100
5.3	Modulation through an object	102
5.3.1	Projection approximation	103
5.3.2	Multi-slice modulation	105
5.4	Imaging	106
5.4.1	Propagation-based imaging (PBI)	106
5.4.2	Other methods	107
5.5	Phase retrieval	107
5.5.1	Single-distance	108
5.5.2	Spectral	109
5.5.3	Material decomposition	110
5.6	Summary	111
5.7	Appendix: Derivation of the TIE	112
6	CRAMÉR–RAO LOWER BOUND FOR SPECTRAL PB-XPCI	114
6.1	Introduction	114
6.2	Background	114
6.3	Theory	117
6.3.1	Two-pixel model	118
6.3.2	Line-profile model	121

6.3.3	2D geometry model	125
6.4	Methods	126
6.5	Results and Discussion	129
6.5.1	Two-measurement models (absorption only)	129
6.5.2	Line-profile model (XPCI)	134
6.5.3	2D geometry model and MLE validation	140
6.6	Conclusion	142
6.7	Appendix: Partial derivatives of PB-XPCI measurements	146
7	ACCURACY OF THE PROJECTION APPROXIMATION	149
7.1	Introduction	149
7.2	Background	150
7.3	Methods	153
7.3.1	Rod phantom line profiles	154
7.3.2	QA phantom projections	155
7.3.3	Zebrafish micro-CT	157
7.4	Results and Discussion	159
7.4.1	Rod phantom line profiles	159
7.4.2	QA phantom projections	163
7.4.3	Zebrafish micro-CT	163
7.5	Conclusion	173
8	ITERATIVE MATERIAL DECOMPOSITION FOR DUAL-ENERGY PB-XPCI	175
8.1	Introduction	175
8.2	Background	176
8.3	Methods	178
8.3.1	Forward models	178
8.3.2	Material decomposition phase retrieval (MDPR)	180
8.3.3	Simulations	182
8.4	Results and Discussion	187
8.4.1	Spatial resolution	187
8.4.2	Propagation distance	188
8.4.3	Noise	190
8.4.4	Initialization	191
8.5	Conclusion	194
9	FULL SPECTRAL MATERIAL DECOMPOSITION PHASE RETRIEVAL	195
9.1	Introduction	195
9.2	Theory	197
9.2.1	Polychromatic forward model	197
9.2.2	Analytical inverse	198
9.2.3	Optimization-based solution	199
9.3	Methods	200
9.3.1	Materials	200

9.3.2	TIE-MD simulations	201
9.3.3	Iterative MDPR simulations	202
9.4	Results and Discussion	205
9.4.1	TIE-MD energy sensitivity	205
9.4.2	Iterative MDPR tuning	207
9.4.3	Model sensitivity	211
9.4.4	Calcification decomposition	213
9.5	Conclusion	216
10	CONCLUSIONS	219
10.1	Summary	219
10.2	Future Directions	222
	REFERENCES	224

LIST OF FIGURES

1.1	(Left) Portrait of Wilhelm Röntgen with his imaging apparatus. (Right) Radiograph of his wife Anna Bertha Ludwig's hand. Images from the Wellcome Collection, via World History Encyclopedia, licensed under Creative Commons (CC-BY and CC BY-NC) with grayscale conversion.	2
2.1	A simple model of X-ray attenuation along a single, monochromatic ray passing through a uniform object (Eq. 2.2).	16
2.2	A model of X-ray attenuation that accounts for the 3D object's spatial and energy-dependent $\mu(\mathbf{r}, E)$ and the polychromatic X-ray spectrum $N_0(E)$ (Eq. 2.3). . . .	17
2.3	Schematic of two common X-ray detectors: (left) the indirect-conversion scintillator-based EID used on most clinical CT scanners, and (right) direct-conversion semiconductor PCD that is being progressively implemented clinically.	22
2.4	Schematic of theoretical electrical readout for a spectral PCD channel (left) and the resulting detected energy spectra (right). The bin spectra were generated using an approximate non-ideal spectral threshold scheme where each interacting X-ray follows a Gaussian probability distribution of being detected at its true energy. More precisely, the bin spectra are distorted by other nonidealities such as charge sharing, K-escape, and pulse pile-up.	23
2.5	Two common CT acquisition geometries. (Left) Clinical fan-beam CT scanner using a source-detector array that rotates in tandem about the patient. (Right) Parallel-beam micro-CT with a stationary synchrotron X-ray source and rotating sample.	26
2.6	Linear attenuation coefficients in the X-ray energy range for some materials relevant to diagnostic imaging. Water has no apparent K-edge; calcium has a K-edge too low to be apparent in most images; iodine and gadolinium have sufficiently high K-edges to be candidate contrast agents for K-edge imaging. Also note that L-edges are visible for iodine and gadolinium around 4–8 keV. All materials were defined with density $\rho = 1 \text{ g/cm}^3$	32
3.1	The modeled detective efficiency function $\eta(E)$	40
3.2	The five spectra with magnitude scaled to deliver 1 μGy dose.	41
3.3	The two computational phantoms imaged (noiseless 80-keV monoenergetic images). The metal hip replacement is indicated by the arrow in panel (b). In the modeled hip prosthesis, the femoral head and outer shell were either titanium or stainless steel, and the inner lining was PMMA. CNR was computed using the delineated ROIs, and RMSE was evaluated within the phantom.	42
3.4	Linear attenuation coefficients of metals common in modern hip prostheses (cobalt-chromium alloys, titanium alloys, and surgical-grade stainless steel) in (a) the keV energy range and (b) the MeV energy range.	43
3.5	Heatmap of peak tissue SNR as a function of bone thickness for all dual-energy spectral combinations.	46
3.6	Heatmap of peak bone SNR as a function of bone thickness for all dual-energy spectral combinations.	47

3.7	With 1, 4, or 6 cm bone thickness, curve of bone SNR as a function of dose allocated to the high-energy spectrum for the detunedMV-80kV (a) and 140kV-80kV (b) spectral pairs.	47
3.8	Maximum SNR for tissue (solid lines, left scale) and bone (dashed lines, right scale) for the 140kV-80kV (kV-kV, square marker) and detunedMV-80kV (MV-kV, circle marker) spectral combinations as a function of bone thickness.	48
3.9	CNR in the detunedMV-80kV (MV-kV) and 140kV-80kV (kV-kV) VMIs as a function of energy, with horizontal lines at the CNR of each kV single-energy CT acquisition. The total dose of the dual-energy scans was equivalent to the dose of the single-energy scans (10 mGy).	49
3.10	Measured contrast (numerator of the CNR) in the detunedMV-80kV (MV-kV) and 140kV-80kV (kV-kV) VMIs as a function of energy, with horizontal lines at the value of each dose-matched kV single-energy CT acquisition.	50
3.11	Measured noise (denominator of the CNR) in the detunedMV-80kV (MV-kV) and 140kV-80kV (kV-kV) VMIs as a function of energy, with horizontal lines at the value of each dose-matched kV single-energy CT acquisition.	51
3.12	RMSE in the detunedMV-80kV (MV-kV) and 140kV-80kV (kV-kV) VMIs as a function of energy. RMSE was calculated relative to the monoenergetic ground truth XCAT.	51
3.13	Single-energy CT images (a-b), basis material images (c-d), and VMIs (e-f) for the 140kV-80kV DE-CT acquisition of the pelvis with no hip replacement. The VMIs correct the beam hardening visible in the single-energy CT images. . . .	63
3.14	Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the detunedMV-80kV DE-CT acquisition of the pelvis with no hip replacement. VMI synthesis corrects for the beam hardening in the 80kV single-energy image but results in increased noise from the detunedMV single-energy image.	64
3.15	Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the 140kV-80kV DE-CT acquisition with titanium hip replacement. In the VMIs, the beam hardening artifact of the single-energy images has been corrected, but there is residual streaking.	65
3.16	Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the detunedMV-80kV DE-CT acquisition with titanium hip replacement. The 300 keV VMI does not display streaking or beam hardening artifacts, but it shows a similar noise profile as the detunedMV single-energy image.	66
3.17	Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the 140kV-80kV DE-CT acquisition with stainless steel hip replacement. Both single-energy images and VMIs suffer from severe streaking artifacts.	67
3.18	Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the detunedMV-80kV DE-CT acquisition with stainless steel hip replacement. The 300 keV VMI highly alleviates the streaking artifact.	68
3.19	The dose-matched, beam-hardening corrected SE-CT images for each phantom (column) and kV spectrum (row). Images are shown with a window width of 500 HU and level of 50 HU.	69

4.1	The two phantoms used for the CT simulations. (a) An image quality phantom of tissue with seven bone inserts of varying densities. (b) A section of the XCAT phantom. Pixel values are in HU evaluated at 80 keV and shown at a window level of 0 and width of 1000. Each image field-of-view is 50 cm.	75
4.2	Flowchart summarizing the MV-kV simulation procedure. Single-energy MV and kV sinograms are simulated independently and then recombined for DE-CT (top row): sinogram-domain material decomposition is performed, then material images are reconstructed, and VMIs are generated using Eq. 3.6. Alternatively, the single-energy sinograms can be directly reconstructed (bottom row): the single-kV image (left) suffers from severe beam hardening, while the single-MV image (right) has poor contrast and high noise.	77
4.3	Basis material thickness line profiles through the IQ phantom used for approximate dose allocation optimization. Tissue thickness ranged between 37 and 40 cm, and bone thickness ranged from 1.5- to 3.5-cm equivalent for the seven different density inserts.	80
4.4	Basis material thickness line profiles through the XCAT phantom used for approximate dose allocation optimization. Tissue thickness ranged between 18 and 32 cm, and bone thickness varied up to 5 cm.	81
4.5	The dose allocation to the MV spectrum maximizing basis material SNR for (a) tissue and (b) bone as a function of bone thickness with fixed tissue thickness of 40 cm.	82
4.6	Dose-optimized tissue SNR using (a) photon-counting and (b) energy-integrating detector models for a variety of basis material thicknesses. Panel (c) shows the percent difference.	82
4.7	Dose-optimized bone SNR using (a) photon-counting and (b) energy-integrating detector models for a variety of basis material thicknesses. Panel (c) shows the percent difference.	83
4.8	Measured versus true bone density in the seven inserts of the image quality phantom. Error bars represent $k = 1$ standard deviation over all noisy simulations.	84
4.9	CNR measured in three bone inserts with density (a) 1.4 g/cm^3 , (b) 1.8 g/cm^3 , or (c) 2.2 g/cm^3 as a function of VMI energy. The error bars represent $k = 2$ standard deviations over all noisy simulations.	85
4.10	(a)–(b) 50-keV VMIs and (c)–(d) their noise profiles simulated with either PCD (left) or EID (right) schemes. Pixel units are HU. The noise profile is concentrated around the bone regions.	85
4.11	(a)–(b) 60-keV VMIs and (c)–(d) their noise profiles simulated with either PCD (left) or EID (right) schemes. Pixel units are HU. The noise profile is distributed around the bones and lateral edges of the phantom.	86
4.12	(a)–(b) 70-keV VMIs and (c)–(d) their noise profiles simulated with either PCD (left) or EID (right) schemes. Pixel units are HU. The noise profile has a more uniform texture across tissue regions.	86

4.13	The (a) kilovoltage and (b) megavoltage X-ray spectra with photon counts per 1-keV energy bin scaled to deliver approximately 1-μGy dose to the center of a 40-cm water cylinder.	92
4.14	Single-line integral basis material SNR measured as a function of spectral threshold imposed on the (a) kilovoltage and (b) megavoltage spectra for an object composed of 37-cm tissue and 3-cm bone.	93
5.1	Schematic of propagation-based imaging using multi-slice modulation through N object slices, each with thickness $\Delta z = T/N$, where $\exp[\dots]$ is shorthand for the projection approximation equation (Eq. 5.40) and $\mathcal{D}_z$ is the free-space propagation operator over distance z	105
6.1	Two-pixel CRLB toy model.	119
6.2	Multi-pixel line-profile toy model: a uniform plane wave is incident on a single-edge object. The PB-XPCI acquisition process is modeled, and then the measurements for the Fisher matrix are taken from a small region of pixels in a single row centered about the material edge.	122
6.3	2D geometry toy model for a system of two basis material cylinders.	125
6.4	Illustration of the detector response functions $D_a(E)$ modeled for the spectral PCD system given single energy thresholds of $E_t = 15, 20$, or 25 keV for a normalized 30-kVp spectrum $S(E)$ with a 5-keV FWHM spectral threshold blur.	128
6.5	Comparison of basis material CRLBs with either a two-pixel or two-energy system for 1-mm bone and 4-mm soft tissue background. The two-energy CRLB curve shows the minimum achieved for a given low energy, i.e. $\text{CRLB}_j^{(2E)}(E_1) = \min \text{CRLB}_j(E_1, E_2)$ for $E_2 > E_1$. The full two-energy $\text{CRLB}_j(E_1, E_2)$ heatmaps are shown in Fig. 6.6.	130
6.6	Two-energy/single-pixel CRLB heatmap for all energy pairs for 1-mm bone and 4-mm soft tissue. The white “+” indicates the minimum CRLB coordinates, at (11.8 keV, 23.0 keV) for both materials. The $E_1 = E_2$ diagonal is not defined for this system.	130
6.7	Two-energy/two-pixel CRLB for 1-mm bone and 4-mm soft tissue background. The white “+” indicates the minimum CRLB coordinates, at (15 keV, 15 keV) for bone and (10 keV, 10 keV) for tissue.	133
6.8	Two-energy/two-pixel CRLB for 1-mm fibroglandular tissue and 4-mm adipose tissue background. The white “+” indicates the minimum CRLB coordinates, at (8.8 keV, 8.8 keV) for gland and (8.0 keV, 8.0 keV) for adipose.	133
6.9	Bone/soft-tissue basis: line-profile dual-energy CRLB heatmaps for $R = 0$ mm (traditional) and 100 mm (XPCI), with line profiles for two example energies.	135
6.10	Fibroglandular/adipose-tissue basis: line-profile dual-energy CRLB heatmaps for $R = 0$ mm (traditional) and 100 mm (XPCI), with line profiles for two example energies.	136

6.11	Bone/soft tissue polychromatic line-profile CRLBs as a function of threshold energy. The solid black line is for a one-threshold (two-bin) system. The thin marked lines show different two-threshold (three-bin) systems, with color/marker indicating the lower energy threshold and the horizontal axis the higher threshold.	138
6.12	Gland/adipose tissue polychromatic line-profile CRLBs as a function of threshold energy. The solid black line is for a one-threshold (two-bin) system. The thin marked lines show different two-threshold (three-bin) systems, with color/marker indicating the lower energy threshold and the horizontal axis the higher threshold.	139
6.13	Example noisy simulated images used for the MLE validation study with a 20-keV energy threshold.	140
6.14	For the bone and soft tissue basis, comparison of CRLB and MLE simulation variance as a function of threshold energy for the polychromatic spectral PCD system. For each panel, the CRLB (solid line) uses the left vertical axis and the MLE (dashed line) uses the right vertical axis for scale. A ± 1 keV window moving average is shown over the MLE variances, which were computed for individual E_t in 0.1 keV intervals.	141
6.15	For the gland and adipose tissue basis, comparison of CRLB and MLE simulation variance as a function of threshold energy for the polychromatic spectral PCD system. For each panel, the CRLB (solid line) uses the left vertical axis and the MLE (dashed line) uses the right vertical axis for scale. A ± 1 keV window moving average is shown over the MLE variances, which were computed for individual E_t in 0.1 keV intervals.	141
6.16	For the bone and soft tissue basis, bias of the MLE thickness estimates.	141
6.17	For the gland and adipose tissue basis, bias of the MLE thickness estimates.	142
7.1	Fresnel number heatmaps in parameter ranges representing three X-ray imaging modalities: micro-CT, mammography, and clinical CT.	152
7.2	Schematic of the rod phantom.	154
7.3	Refractive index components for materials used in the rod phantom simulations.	155
7.4	Projected thickness of the two randomly generated PMMA-sphere clutter phantoms. On the left, the larger “mammo-clutter” phantom was generated over a full 5-cm ² FOV, and a 1-cm ² ROI (red box) was chosen for the simulation. On the right, the full 1-mm ² FOV was used for simulations of the smaller “micro-clutter” phantom.	156
7.5	Zebrafish phantom at the original 1.6- μ m resolution (left) and after the smooth upsampling operation to achieve 0.2- μ m resolution for accurate PB-XPCI micro-CT simulations.	157
7.6	Refractive index components for the zebrafish phantom simulations. The metal stains osmium tetroxide and uranyl acetate reflect a mixture of 1% stain in water.	158
7.7	Simulated aluminum rod phantom line profiles before (dashed line) and after (marked solid line) resampling at the detector pixel size $\Delta x = 0.5$ μ m.	160
7.8	Simulated PMMA rod phantom line profiles before (dashed line) and after (marked solid line) resampling at the detector pixel size $\Delta x = 0.5$ μ m.	160

7.9	RMSE between MS and PA rod phantom simulations as a function of detector pixel size Δx	161
7.10	RMSE between 100-slice MS line profiles and MS simulations with just N_z , illustrating convergence of MS simulations as $N_z \rightarrow \infty$	161
7.11	1-cm ² FOV mammo-clutter phantom simulations with the 55- μ m pixel detector.	164
7.12	1-cm ² FOV mammo-clutter phantom simulations with the 8- μ m pixel detector.	165
7.13	1-mm ² FOV micro-clutter phantom simulations with the 8- μ m pixel detector.	166
7.14	1-mm ² FOV micro-clutter phantom simulations with the 0.5- μ m pixel detector.	167
7.15	Ideal reconstructions of the stained and unstained zebrafish phantom to provide baseline reference for their native attenuation contrast. The ground truth phantom is shown on the left for reference with the same four-material colormap of Fig. 7.5. The full FOV is 2.5 mm with matrix size 4096 $\times$ 4096. The bottom row shows the detailed structures within the boxed 0.5-mm ROI.	168
7.16	Stained zebrafish contact image ($R = 0$), simulated with either PA or MS, with sub-micron detector pixel size (top row) or larger binned detector pixels (bottom row). 0.5-mm ROI of the full reconstructed matrix.	169
7.17	Unstained zebrafish contact image ($R = 0$), simulated with either PA or MS, with sub-micron detector pixel size (top row) or larger binned detector pixels (bottom row). 0.5-mm ROI of the full reconstructed matrix.	170
7.18	Stained zebrafish images with $R = 20$ mm, raw intensity measurements (top row) and post-phase retrieval (bottom row), simulated with either PA or MS. 0.5-mm ROI of the full reconstructions with sub-micron detector pixels.	171
7.19	Stained zebrafish images with $R = 20$ mm, raw intensity measurements (top row) and post-phase retrieval (bottom row), simulated with either PA or MS. 0.5-mm ROI of the full reconstructions with sub-micron detector pixels.	172
8.1	Schematic of the CT acquisition geometry and simulation phantom. The static coordinates (x, y) define the stationary plane of the phantom, while the rotating coordinates (ξ, ℓ) change with the view angle according to Eqs. 8.2–8.3.	183
8.2	SSIM and 1–RMSE in each basis image as a function of detector pixel size Δx using either PA- or MS-based MDPR.	185
8.3	PA, MS, and ground truth (GT) PMMA reconstructions at different spatial resolutions (phantom thickness = 205 μ m).	186
8.4	An example line profile through the 10-keV sinogram generated at different detector pixel sizes with simulation upsampling (top row; full MS model) and without upsampling (bottom row; optimization MS model).	186
8.5	SSIM and 1–RMSE in each basis image as a function of propagation distance scale relative to the FOV (R scale = R / FOV).	189
8.6	SSIM and 1–RMSE in each basis image as a function of incident flux.	190
8.7	Noisy data PA and MS reconstructions using either a constant initialization of 1s (left; “const init”) or a downsampled, blurred ground truth (right; “BGT”).	192
8.8	Ideal data PA and MS reconstructions using either a constant initialization of 1s (left; “const init”) or a downsampled, blurred ground truth (right; “BGT init”).	193

8.9	Loss history curves for the MDPR runs shown in Figs. 8.7 and 8.8, with end threshold $\Delta\mathcal{L}_t = 10^{-5}$	193
9.1	Polychromatic spectrum and energy-dependent detector efficiency function based on the absorption of a 1-mm thick silicon PCD.	201
9.2	Two-material phantom used for optimization tuning and sensitivity analysis. . .	203
9.3	Multi-material phantom used for robustness analysis.	203
9.4	Example monochromatic simulations and analytical TIE-MD output for three different energy pairs. The example intensity data in the left column is shown for the middle energy pair, (20 keV, 30 keV).	205
9.5	RMSE of TIE-MD output thickness maps from monochromatic simulated data or from polychromatic simulated data (requiring substitution of effective energies in the TIE-MD algorithm), with and without noise, at a range of energy thresholds. The right column shows the effective monoenergies calculated for each threshold.	206
9.6	Basis material CRLB (3-mm soft tissue background with variable 0.3 to 0.5-mm bone thickness) and TIE-MD RMSE as a function of energy threshold.	207
9.7	20-keV threshold reconstructions produced from different optimization initialization methods.	209
9.8	Central line profiles, averaged over the middle ten rows, in the reconstructed 20-keV images from the best initialization methods with either a constant or TIE-MD variant.	209
9.9	RMSE as a function of regularization weight γ for the constant (3,0) and TIE-blur50 initializations with 20-keV threshold data. Example thickness maps are shown for the initialization and γ at the starred point on the curve.	210
9.10	Thickness maps and differences relative to ground truth (GT), reconstructed using different optimizer forward-projection models. As reference, the bottom row shows the unoptimized TIE.	214
9.11	Thickness maps and differences relative to ground truth (GT), reconstructed using the full polychromatic forward projector in each iteration with different degrees of wave upsampling (1, 2, or $4\times$ the reconstructed resolution).	215
9.12	Raw data simulated at three noise levels with the multi-material phantom (threshold energy $E_t = 20$ keV).	216
9.13	MDPR-retrieved basis material thickness maps of the multi-material phantom data.	217
9.14	Line profiles through the bone-basis image in the central rows corresponding to the type I (top) and type II calcifications (bottom).	217

LIST OF TABLES

1.1	Summary of current spectral, single-distance PB-XPCI phase retrieval methods.	12
3.1	Peak coordinates ($r_{\max}$, $\text{SNR}_{\max}$) as a function of dose allocation for both basis materials and all spectral pairs with a bone thickness of 1 cm.	49
3.2	For each phantom, maximum CNR measured in the MV-kV and kV-kV VMIs, and single-energy CNRs with and without a beam-hardening correction (BHC). For each dual-energy acquisition, the VMI energy corresponding to the maximum CNR is indicated in parentheses. For each single-energy spectrum, the change in CNR with BHC is indicated in parentheses.	50
6.1	Density and elemental composition (% mass) used for the basis materials [1]. . .	126
7.1	Parameter ranges used to represent the three imaging modalities for Fresnel number heatmaps (Fig. 7.1).	151
7.2	Average and maximum difference ($ \text{PA} - \text{MS} $) in normalized raw intensity images ($R = 0$ or $10/100$ cm) and phase retrieval “PR” images ($R = 10/100$ cm) for the different detector and phantom sizes. The mammo-clutter phantom was 5-cm thick while the micro-clutter phantom was only 0.5 cm thick.	162
9.1	Different forward projections used for the optimization model sensitivity analysis.	204
9.2	RMSE [mm] of retrieved basis material thickness images from different initialization methods, with acquisition energy threshold optimized either for the TIE (20 keV) or CRLB (23 keV). The asterisk * indicates the minimum unrounded RMSE achieved for each material.	208
9.3	Model sensitivity analysis: RMSE [mm] for different optimization forward models, systematically excluding nonideal system elements such as energy dependencies and blur. Optimization was initialized with either TIE-blur50 or Const (3,0). The asterisk * indicates the minimum unrounded RMSE for each material. . .	211

ACKNOWLEDGMENTS

There are a number of people to whom I owe a serious solid for their support in making this dissertation at least partially coherent.¹

Before all is my advisor, Patrick La Rivière, who is frankly an unbeatable example of an excellent scientist and mentor. I am always amazed by Patrick's spot-on answer to seemingly any question, whether it is about physics, history, coffee, or the hot new restaurant in town. I will definitely be picking up a copy of the *Big Book of Amazing Facts*. I cannot adequately express my gratitude for his advice, patience, and friendship these five years.

Among the UChicago faculty, my dissertation committee has been my first line of support: Ingrid Reiser, Xiaochuan Pan, and Chien-Min Kao. Their questions and insights were imperative to many key elements of this project. I have also benefited from the wisdom of several more faculty mentors: Sam Armato, Maryellen Giger, Emil Sidky, Yulei Jiang, Zheng Feng Lu, and Emily Marshall, among others. Behind the scenes, I am grateful to Lili Gonzalez Hernandez and Elena Rizzo, who have always been supportive student advocates.

I am very fortunate to have met a number of great people among the students here. My first grad school friends were my Ph.D. cohort: Lucas Berens, Amy Renné, and Chisondi Warioba. I also want to thank my I-115 office buddies Hadley DeBrosse and Geneva Anderberg, along with many fantastic students I've worked with along the way.

To my family, Angie Moreschi, Chris Jadick, and Mica Jadick, thank you for long being my most unconditional supporters. I hope you enjoy reading some bits of this tome, and I apologize in advance for all the equations. I promise they're important!

Finally, to my lucky charm, Ellen Mines: I actually could not have done this without you. You never doubt for even a moment my greatest ambitions and ideas, and your quiet selflessness has gotten me through my most challenging moments. You probably know more of the basic science behind this thesis than some physicists. You are going to be a wonderful

1. This is the part of the dissertation where phase-contrast jokes are allowed, right?

doctor (the other kind).

This dissertation would not have been possible without the support of the U.S. National Science Foundation Graduate Research Fellowship Program, which gave me the great privilege of pursuing a research project of my choice under my faculty advisor. I also thank the Biological Sciences Graduate Fellowship and the AAPM/RSNA Graduate Fellowship for their support.

ABSTRACT

X-ray phase-contrast imaging (XPCI) is a promising modality sensitive to fine structures that are often invisible with traditional absorption-only X-ray imaging. However, since XPCI depends on two contrast mechanisms (absorption and phase), it in principle requires multiple independent measurements for accurate reconstruction. With emerging photon-counting detectors available to conveniently provide multiple energy measurements in a single shot, spectral XPCI has emerged as a promising candidate for quantitative phase retrieval.

Taken separately, spectral imaging and XPCI each aim to improve soft-tissue image quality, but they face distinct challenges: limited native soft-tissue sensitivity and the need for multiple measurements, respectively. The combined spectral XPCI approach could offer new advantages while resolving these same limitations. While previous spectral XPCI material decomposition approaches have demonstrated the technique’s potential to improve image quality, they typically rely on inverting relatively idealized and simplified forward models. This is likely due to the lack of an analytical solution for a full polychromatic system model and the computational challenges of developing a numerical or machine-learning based solution in the context of XPCI.

The purpose of this dissertation is to unite traditional spectral X-ray imaging with the emerging field of XPCI through the lens of material decomposition. In part one, we focus on absorption-contrast imaging, using the Cramér–Rao lower bound (CRLB) to assess and optimize dual-energy CT systems with megavoltage and kilovoltage spectra for improved radiotherapy imaging. Part two expands the scope to include propagation-based phase contrast, developing a spectral XPCI CRLB model and other simulation-based metrics for assessing basis material image quality. We utilize these to examine the common assumptions underlying existing spectral phase-retrieval and material-decomposition methods. Finally, in part three, we propose an improved material decomposition framework for spectral XPCI based on a more generalizable optimization-based approach. Using an iterative solution enables the

inversion of a more holistic forward model, bypassing common simplifying approximations required by existing closed-form inversions over parameter ranges typical of both current and novel sub-micron X-ray imaging systems. This work lays the foundation for future work aiming to more fully realize the benefits of spectral imaging through quantitative phase retrieval and material decomposition.

CHAPTER 1

INTRODUCTION

The term *X-ray* was coined in 1895 in Wilhelm Röntgen’s laboratory, upon observing the unexpected fluorescence of some crystals near a seemingly unrelated experiment with his cathode ray tube [2]. The name was meant as a temporary placeholder, using “X” as the mathematical designation for something unknown while he worked to understand the cause of this phenomenon [3]. More than a century later, we know much more about X-rays. A few languages have retitled the elusive rays to memorialize Röntgen’s discovery (in German, *röntgenstrahlen*), but most still refer to them by the mysterious moniker *X*. Though possibly an accident of time, the persistence of the name “X-ray” is perhaps appropriate, considering there are still many open scientific questions regarding their implementation.

Among these is the challenging task of how to leverage X-rays for soft-tissue imaging. Since their discovery, X-rays have been famous for picturing bones. In fact, the most prolific X-ray image of all time was a 15-minute exposure of the hand of Röntgen’s wife, Bertha, showing a detailed shadow of her finger bones and golden wedding ring (Figure 1.1). She is reported to have exclaimed, “Now I have seen my death!” upon seeing her own skeleton, as prior to the advent of X-ray imaging, it was only possible to see one’s bones post-mortem [4].

Today, X-ray images are not nearly so shocking. Standard radiographs are routinely used by radiologists for their excellent contrast between bones and surrounding tissues, offering conclusive diagnosis of fractures and dislocations. Certain lung conditions are also readily detected due to the high contrast between tissue and air-filled pulmonary structures. Broadly, there are roughly five appreciable levels of contrast in radiography (air, adipose tissue, soft tissue, bone, and metal) [5]. Other X-ray based modalities have been developed to further improve diagnostic power; for example, computed tomography (CT) acquires X-ray projections at different angles for three-dimensional (3D) image reconstruction to resolve




Figure 1.1: (Left) Portrait of Wilhelm Röntgen with his imaging apparatus. (Right) Radiograph of his wife Anna Bertha Ludwig’s hand. Images from the Wellcome Collection, via World History Encyclopedia, licensed under Creative Commons (CC-BY and CC BY-NC) with grayscale conversion.

overlapping structures, and fluoroscopy acquires a time series of low-dose projections for real-time interventions and diagnostic “videos” [6, 7]. The removal of overlapping structures in CT leads to better contrast resolution relative to single-projection radiography, roughly on the order of 5–10 Hounsfield units [6].

Yet, the native ability of X-rays to distinguish variations within soft tissues has not substantially improved in clinical applications. Many soft-tissue imaging tasks are more difficult using X-rays because, generally, sub-types of tissues (whether healthy or cancerous) are close to each other in density and atomic composition and consequently have only small differences in their X-ray absorption. Despite this, X-ray based modalities are frequently ordered for such imaging tasks anyway, owing to their rapid acquisition and relatively inexpensive implementation. The images can provide some initial evidence for various conditions but typically require follow-up imaging with a different modality like magnetic resonance imag-

ing (MRI) or ultrasound for physicians to make a definitive diagnosis, which can be time and resource exhaustive.

For example, consider digital mammography, a two-dimensional (2D) X-ray modality. As the primary breast cancer screening method, mammography is applied to an asymptomatic population to search for potentially malignant masses and microcalcifications—both of which have only small differences in X-ray absorption relative to surrounding breast tissue [8]. To be feasibly implemented at the population level, screening modalities must be rapid, reliable, cost-effective, and low-risk. Yet, due to the key challenge of detecting weakly absorptive lesions, to achieve sufficient sensitivity, mammography has notoriously low specificity. Approximately 1 in 10 patients will be called back for additional testing, and of those, only about 7% will receive a cancer diagnosis [9]. Within a decade of annual screening, over half of patients can expect at least one false-positive recall [9]. Despite this, even for women with dense breasts—the highest risk population—a systematic review of proposed supplemental screening modalities such as MRI and ultrasound found at best mixed benefits with economic modeling [10]. There would be great benefit to a better screening method pairing the economic benefits of X-rays with the sensitivity of supplemental modalities.

Several pre-clinical applications face similar challenges associated with high-resolution soft-tissue imaging. Many biological fields are exploring increasingly complex and small-scale phenomena, which contain untapped 3D information for treating and diagnosing disease [11–13]. One prominent example is in tumor research: work has shown that the 3D cellular and extracellular architecture of a tumor has critical implications for its progression and response to therapy [11]. Another is tissue phenotyping, which is fundamental to understanding the genetic and cellular origins of disease [12]. Data-driven analysis is a promising potential path for uncovering fundamental principles dictating these complexities, but this requires prohibitively massive amounts of high-quality images for model training, hence introducing the need for an experimentally convenient modality with suitable sensitiv-

ity to soft-tissue changes. The current conventional approaches of serial histology and serial block-face scanning electron microscopy require labor-intensive sample preparation—manual cutting into the serial slices, repeated metal staining, and fixation in resin [11]. An X-ray imaging approach like micro-CT has potential to be more convenient, since it is natively 3D and thus does not require serial slicing. However, the small size of the biological samples and fine tissue structures of interest have relatively weak X-ray absorption, so micro-CT still requires a labor-intensive step of applying metal stains [12, 14]. This is practically prohibitive the necessary high-throughput data acquisition for advanced understanding of micron-scale 3D architecture, which could enable tailored disease treatments and lead to better patient prognoses.

1.1 Phase contrast

For these imaging tasks and X-ray based modalities, differential X-ray absorption has been the dominant contrast mechanism dictating how well structures and diseases can be identified. For certain tasks, this is clearly sufficient; but for others like soft-tissue imaging tasks, X-ray modalities are more likely implemented for their practical advantages that make them more cost-effective but less informative. However, X-rays possess a second contrast mechanism that has been far underexplored relative to absorption until much more recent history: phase.

As a form of electromagnetic radiation, X-rays are fundamentally both absorbed and refracted, just like visible light. Their intensity and phase modulations are dictated by the complex refractive index of the materials through which they pass, commonly written as $n = 1 - \delta + i\beta$. The imaginary component β relates to absorption while the real decrement δ relates to phase shifts. As it turns out, in the X-ray energy range, δ tends to be roughly two or three orders of magnitude larger than β , particularly for low-atomic number materials most relevant to biomedical imaging. This suggests that phase could supply a better contrast

mechanism than absorption for traditionally challenging X-ray imaging tasks [4].

But there are a few reasons why X-ray refraction is not immediately noticeable in conventional medical images. Although an X-ray wave field accumulates phase shifts as it travels through an object, standard detectors measure intensity alone. So, the phase of the X-ray wave must somehow be translated into intensity variations in order to be measured. In the X-ray energy range, the refractive index typically has only a tiny deviation from unity, resulting in very small refraction angles and thus interference patterns that are too fine to be resolved with the spatial resolution of diagnostic imaging technology. Nonetheless, with the right equipment, methods have been developed to observe X-ray phase shifts, falling under the broad category of *X-ray phase-contrast imaging* (XPCI).

In the history of XPCI, one of the earliest published notes observing the bending of X-rays was from none other than Albert Einstein in 1918, over two decades after Röntgen’s discovery [4, 15]. Commenting on a radiograph, he described bright outlines along the edges of human limbs and noted that this suggests X-rays have a refractive index slightly smaller than 1, expressing it as $n = 1 - \epsilon$ where ϵ is about 10^{-6} (which is indeed correct) [15]. Unfortunately, neither the image nor its specific acquisition method were published, though a note in Einstein’s proof has been interpreted as suggesting it was a propagation-based image, meaning the bright edges were interference patterns or “Fresnel fringes” [4]. The same observation that $n < 1$ was again published by Compton a few years later in 1923 with some additional experimental evidence [16].

Another several decades passed before the wave nature of X-rays was successfully harnessed for imaging. In 1942, Frits Zernike developed phase-contrast microscopy in the visible light domain, showing how an otherwise transparent specimen can be rendered detectable by converting object-induced phase variations into measurable intensity contrast through interference with undiffracted light [17]. Similar ideas took root in the X-ray domain in 1965 with Bonse and Hart’s crystal interferometer, which established X-ray interferometry as the

first XPCI acquisition method [18]. It was not until the 1990s that phase-contrast imaging was demonstrated using no specialized optics equipment by Snigirev et al., relying only on the highly coherent and monochromatic radiation produced by the European Synchrotron Radiation Facility (ESRF) [19]. With a film detector and up to 200-cm object-to-detector distance, they directly observed edge enhancement effects at the bounds of cellulose and polypropylene fibers. This was soon extended to polychromatic radiation by Wilkins et al., who showed that a polychromatic phase-contrast image is essentially a weighted sum of monochromatic images, relaxing the X-ray source requirement to just high spatial coherence [20]. This method has become known as propagation-based imaging (PBI) and is possibly one of the most clinically feasible XPCI acquisition approaches, especially for XPCI-CT, due to its polychromatic compatibility and lack of extra optical elements like mirrors and gratings [21, 22].

1.1.1 Phase retrieval

By the 21st century, a number of XPCI acquisition methods were developed to encode object-induced phase shifts in intensity measurements: interferometry, analyzer-based imaging, grating-based methods, and PBI, among others [21]. These techniques induced qualitative image changes like bright and dark fringes along material interfaces in PBI to improve structural delineation. For quantitative imaging, however, an additional step was needed: solution of the phase-contrast inverse problem, known as *phase retrieval*.

The goal of phase retrieval is to reconstruct quantitative maps of the index of refraction $n(\mathbf{r})$, where $\mathbf{r}$ indicates an object's spatial dependency [23]. This is similar to how CT images are quantitative reconstructions of an object's X-ray attenuation coefficient, $\mu(\mathbf{r})$ [6, 7]. For projection imaging, this takes the form of quantitative images of the line integrals through the index of refraction, $\int d\ell n(\mathbf{r})$, just as traditional radiographs represent $\int d\ell \mu(\mathbf{r})$. The key challenge is that, unlike μ , the refractive index n is a complex value and thus has two

independent components: real and imaginary. Thus, the full phase retrieval problem requires the reconstruction of *two* separate images: the phase-shifting $\delta(\mathbf{r})$ and the absorptive $\beta(\mathbf{r})$. In principle, this requires either two independent measurements to solve or assumptions about the object’s composition or acquisition physics to reduce the number of unknowns.

This led to a practical conundrum. To realize the theorized signal-to-noise benefits owing to the several orders of magnitude advantage of the phase-shifting component, phase retrieval is paramount. Ideally these reconstructions would be quantitatively accurate, but the need for multiple images introduces experimental inconvenience as well as potential for human error.

Due to its simple experimental setup and parallels with diagnostic imaging, this dissertation focuses on the PBI or “PB-XPCI” acquisition technique. To date, the most prolific PBI phase retrieval algorithms utilize various physics assumptions and a single measurement or rely on measurements taken at multiple propagation distances [24–27]. A potential solution that has gained feasibility with recent advances in X-ray technology is *spectral* XPCI.

1.2 Spectral imaging

Spectral X-ray imaging has a decades-long history in the context of traditional absorption-contrast imaging [28]. The term “spectral” imaging broadly refers to any type of energy-resolved imaging, in which multiple acquisitions are acquired with different energy spectra, typically with the end goal of performing subsequent quantitative analysis. In clinical imaging, early applications typically used just two energies, so the method was called “dual-energy” (DE) imaging, which we now consider a sub-type of spectral imaging. More generally, spectral imaging might analogously be called “multi-energy” imaging. The use of *spectral* as an umbrella term is a relatively recent evolution, likely owing to new technology making the acquisition of more than two energy-resolved scans much more convenient [28]. Thus, these different terms are mixed throughout this dissertation but essentially the same.

The idea of using multiple-energy measurements for quantitative medical imaging was first proposed by Bertil Jacobson in 1953, who described a method for quantifying the concentration of iodine in water with two monochromatic radiographs [29]. However, Jacobson’s paper did not find immediate further traction, perhaps due to practical limitations associated with clinical translation. Years later, the 1970s saw a sudden revival in interest with the invention of CT imaging.

CT is a fundamentally quantitative imaging method in the sense that it is concerned with numerical reconstruction of 3D absorption coefficients, with the mathematical theory first formulated by Allan Cormack in 1963 [30, 31] and the concept independently re-derived and popularized by Godfrey Hounsfield in 1973 with the first clinical prototype [32]. With 3D reconstruction, CT can resolve overlapping structures that are obscured with single-projection radiography, improving the diagnostic capabilities of X-rays. Given these quantitative roots, spectral CT may have seemed like a natural extension; even Hounsfield’s original 1973 paper briefly noted the possibility of determining a material’s atomic number using two conventional polychromatic CT scans with tube voltages of 100 kV and 140 kV [32].

Just a few years later, in 1976, a thorough mathematical description of the spectral CT inverse problem was published by Alvarez and Macovski [33, 34]. In essence, they posited that the energy dependence of the X-ray attenuation coefficient $\mu(E)$ could be represented as a linear combination of a finite number of basis functions, and they claimed that the ultimate choice of those functions is an empirical problem. Their initial paper reported success using just two basis functions modeling the photoelectric and Compton scatter interaction probabilities [33], and they soon extended this to basis functions of “assumed materials” beginning with water and bone [34]. This introduced the idea of material decomposition, approximating a patient as comprising a finite set of basis materials. Several practical applications were soon to follow: atomic number and electron density decomposition [35], iodine quantification with K-edge analysis [36], bone mineral density analysis [37], among others

[28]. Dual-energy radiography was revisited as well, now using polychromatic spectra and the Alvarez-Macovski formalism [38].

1.2.1 The multiple measurement problem

Yet, spectral imaging encountered the same practical challenge as XPCI: how to acquire the multiple measurements? As a clinical modality, CT faced even more nonidealities, and patient movement between different energy acquisitions was a major issue, with even slight shifts violating the assumption of perfect pixel-by-pixel alignment. Thus, the sequential scanning approach faced significant barriers to be viable for clinical imaging [39].

The first prototype dual-energy CT system, the Siemens Somatom DR3, aimed to solve this with “fast kVp switching,” an approach in which the X-ray tube generator rapidly changes its voltage between pairs of pulses, so that two energy datasets may be acquired in a single scan [40]. It was implemented with the clinical goal of bone mineral quantification [41]. While this solved the patient motion issue, it introduced other complications associated with detector requirements and the generation of sufficient low-kVp flux [7]. In the late 1980s and the decades following, the clinical imaging industry once again focused their attention on other avenues of innovation.

In 2006, Siemens returned with a new prototype with the potential to solve the multiple measurement problem—dual-source CT [42]. The approach was actually developed to solve a non-spectral clinical problem: the need for sub-second time resolution with echocardiograph-controlled cardiac CT. Using two source-detector arrays mounted on the CT gantry with an angular offset of approximately 90° , the system achieved a mean temporal resolution of 62 ms. For spectral CT, this also introduced a more viable technology for acquiring multiple measurements without motion artifacts [43, 44]. Soon after, GE developed their own clinical dual-energy CT system based on fast-kVp switching, with advances in source and detector technology improving upon the aforementioned difficulties of this approach [45].

To this point, the multiple-measurement problem in spectral CT had mainly been tackled from a *source-side* perspective. Fuji and Philips were the first to introduce *detector-side* approaches in radiography and CT, respectively [28, 46, 47]. Detector-side methods utilize the intrinsic polychromaticity of a single X-ray beam and implement a specialized detector that can acquire energy-resolved measurements. Their initially proposed detectors used a “sandwich” idea of stacked scintillators, with the top layer designed to both (1) detect a portion of the incident beam and (2) allow the remaining, filtered portion of the beam to pass through to the bottom layer. This resulted in the top layer providing a relatively lower-energy measurement and the bottom a higher-energy measurement [39].

1.2.2 *Photon-counting detectors*

Broadly speaking, the detector-side philosophy may have been prescient in anticipating a larger spectral imaging paradigm shift that has emerged in the last decade with the advent of clinically viable photon-counting detectors (PCDs) [48–50]. The dual-source, fast-kVp switching, and dual-layer detector systems each faced their own practical challenges limiting broad clinical adoption. New PCDs offered an energy-discriminating detection approach as well as a number of other advantages over the standard energy-integrating detectors (EIDs) used by clinical CT scanners, such as improved geometric efficiency, elimination of electronic noise, better native contrast, and smaller pixel sizes.

The properties of both PCDs and EIDs will be described in detail in Section 2.3; but for the purpose of this introduction, the main point is that recent PCDs have largely solved the multi-measurement problem for spectral CT imaging and have thus resulted in a modern resurgence of interest in this modality, with the first clinical PCD-CT system (Siemens Naeotom Alpha) launched in 2022 [51]. Several other vendors are developing their own PCD-CT systems, with GE receiving FDA clearance for their proposed system just a few weeks prior to the time of writing (March 2026) [52]. A thorough review of the progress of

other current commercial PCD-CT systems is given in [50].

1.3 Spectral XPCI and current challenges

Returning to phase retrieval, the revitalized spectral CT industry has coincided with renewed interest in spectral XPCI as a more convenient acquisition technique to solve the multi-measurement problem for quantitative phase retrieval. In the early years after the introduction of PBI, the idea of spectral PB-XPCI as a path to quantitative phase retrieval had already been suggested in a simplified form [53, 54]. Now, with the recent development of modern PCDs for clinical X-ray imaging, this line of research is experiencing a surge of new attention [55–61].

It is worth noting that clinical CT scanners use PCDs with a significantly larger pixel size (roughly as small as ~ 0.2 mm) than PB-XPCI, which is pre-clinical and generally requires micron-scale resolution in order to resolve the phase-induced edge-enhancement effects. Nonetheless, motivated by the advantages of PCDs, there have been efforts to develop sufficiently high-resolution PCDs for XPCI [62–64], and they have been successfully demonstrated in benchtop spectral PB-XPCI experiments along with polychromatic micro-focus X-ray sources [55, 56, 61, 65]. This is a promising sign for the practical translatability of spectral XPCI, paralleling the recent advances in spectral PCD-CT.

The currently published spectral PB-XPCI phase retrieval methods are summarized in Table 1.1. These methods have shown promising results, with high-quality reconstructed images indicating proof of principle and the expected practical benefits of a spectral acquisition. Nonetheless, while many solutions claim to be quantitative, a number of common approximations are made that may limit their generalizable accuracy in ways that have not been sufficiently tested; widespread examples include the thin-object assumption, weak intensity and/or phase gradients, near-field propagation regime, and various material-composition assumptions. Some methods have only been presented for monochromatic spectra [58–60],

Reference	Method	Outputs
Gursoy & Das 2013 [55] / Vazquez et al. 2021 [56]	Analytical inverse solution in the frequency domain based on a linearized transport-of-intensity equation; proposed in 2013 and more robustly validated in 2021.	Absorption and phase images
Li et al. 2018 [57, 58]	“Linear iterative phase retrieval” (LIPR); iterative method based on the Alvarez–Macovski model with linearized attenuation update steps.	Atomic number and density [57]; basis material thicknesses [58]
Schaff et al. 2020 [59, 60]	Analytical inverse solution in the frequency domain similar to [55] with dual-energy imaging for direct material decomposition.	Basis material thickness images
Ghani et al. 2021 [61]	Iterative method using update steps based on the phase-attenuation duality approximation with Fresnel propagation.	Absorption and phase images

Table 1.1: Summary of current spectral, single-distance PB-XPCI phase retrieval methods.

which has its uses for early testing and biomedical research at synchrotron facilities, but for practical purposes is not clinically translatable and would still typically require sequential acquisitions, which can introduce source misalignment errors [57]. To fully capitalize on the purported practical translatability of spectral PCDs, a phase retrieval method should be compatible with polychromatic spectra and energy-discriminating detector model. Further, theoretical image quality metrics like the Cramér–Rao lower bound (CRLB) originally used for spectral CT imaging are paramount for system optimization and thorough evaluation of phase retrieval performance. Such metrics were developed early on for absorption-only spectral imaging [66], even presented in Alvarez and Macovski’s original work [33]. Similar theoretical metrics have emerged for some different XPCI techniques [67, 68], but still none developed for PB-XPCI, with or without a spectral element.

1.4 Dissertation overview

The objective of this dissertation is to unite traditional spectral imaging with the emerging field of XPCI through the lens of material decomposition, ultimately developing spectral XPCI as a better approach for quantitative phase retrieval. The primary aims are (1) to evaluate the theoretical performance of spectral PB-XPCI systems and (2) to develop a more general and accurate solution to the spectral PB-XPCI inverse problem by avoiding common approximations.

The dissertation is organized into three main parts. In the first, we focus on traditional absorption-contrast spectral imaging. Chapter 2 outlines the necessary background theory, presenting X-ray imaging fundamentals from an absorption-only foundation, while the following two chapters explore a potential application of dual-energy CT for improved soft-tissue contrast in radiotherapy using combined megavoltage (MV) and kilovoltage (kV) spectra. Chapter 4 introduces an estimation theory framework for assessing “MV-kV” dual-energy CT relative to traditional kV single- and dual-energy CT. Chapter 5 quantifies the potential improvements that could arise by using emerging PCDs with MV-kV imaging instead of traditional EIDs.

The second part expands upon the absorption-only theory to include propagation-based phase contrast from a forward modeling perspective. Chapter 6 parallels the first absorption-only fundamentals of Chapter 2, introducing phase-contrast theory with an eye to first principles and clear identification of simplifying assumptions. Chapter 6 extends the original spectral imaging estimation theory framework to evaluate the theoretical performance of spectral PB-XPCI systems prior to phase retrieval. Chapter 7 begins to explore specific approximations made in common phase retrieval methods using both analytical and simulation approaches.

Finally, the third part presents our solution to the spectral phase retrieval problem. Chapter 8 introduces our method, which uses an iterative solution and automatically dif-

ferentiable forward model to relax reliance on conventional physics approximations in the context of synchrotron micro-CT imaging. Chapter 9 extends this to the challenging task of inverting a full polychromatic imaging system model. The utility of our method is explored through comparison with other phase retrieval methods and an evaluation of which forward model elements are most essential for optimization performance.

We conclude in Chapter 10 with a summary of key results and proposed future directions.

CHAPTER 2

X-RAY IMAGING FUNDAMENTALS

We begin by orienting ourselves with an overview of X-ray imaging physics from the perspective of the diagnostic imaging physicist, in which absorption is the primary contrast mechanism. This provides a foundation for the following two chapters exploring clinical-scale dual-energy CT imaging. The more wholistic physics perspective including both absorptive and phase-shifting behavior will be covered beginning in Chapter 5.

2.1 Conventional X-ray imaging model

In diagnostic imaging, the X-ray is generally regarded as a high-energy photon that behaves as a discrete “packet” of energy [69]. In the diagnostic imaging realm, X-ray energies fall roughly within the range of 1 to 150 keV, while for radiation therapy treatments, they can span up to 25 MeV [6].

There are essentially two things that an X-ray can do when it passes through an object: either it (1) interacts with the object, or (2) it does not. The interacting portion of an X-ray beam determines the deposited radiation dose (and can thus be tailored for radiation therapy treatments) while the non-interacting portion passes through the object and can be measured using film or digital detectors, forming images. Consequently, the X-ray image characteristics hinge on the object’s interaction probability. The interaction probability per unit distance is called the linear attenuation coefficient μ ,

$$\mu = \mu_{\text{PE}} + \mu_{\text{CS}} + \mu_{\text{PP}}, \tag{2.1}$$

which is the sum of the contributions of three main photon interactions: the photoelectric effect (PE), Compton scatter (CS), and pair production (PP) [70]. Since the minimum energy for pair production to occur is 1.022 MeV (the resting mass of two electrons), μ_{PP} is dropped




Figure 2.1: A simple model of X-ray attenuation along a single, monochromatic ray passing through a uniform object (Eq. 2.2).

for most imaging applications. Then, if an ensemble of N_0 incident photons passes through an object with constant μ and thickness L (Figure 2.1), we expect the exiting number of non-interacting photons N to follow the Beer–Lambert law,

$$N = N_0 e^{-\mu L}. \quad (2.2)$$

In other words, the X-ray beam flux as a function of thickness follows an exponential decay function with coefficient μ [7].

The linear attenuation coefficient μ depends on the chemical composition and density of each material through which it passes as well as the energy the incident X-ray. Most objects are composed of multiple materials, and most X-ray beams are at least somewhat polyenergetic or “polychromatic,” i.e. comprise X-rays with a range of energies. A more complete model of an X-ray beam’s attenuation along a single ray is

$$N(E) = N_0(E) \exp \left[- \int d\ell \mu(\mathbf{r}, E) \right], \quad (2.3)$$

where E denotes energy, $N_0(E)$ is the incident energy spectrum (number of photons per energy), and $\mathbf{r}$ is the three-dimensional (3D) spatial coordinate (Figure 2.2). Now instead of characterizing the total number of photons exiting the object, we characterize $N(E)$, the exiting energy spectrum.

To actually acquire an image, the exiting X-ray spectrum must be converted into a




Figure 2.2: A model of X-ray attenuation that accounts for the 3D object’s spatial and energy-dependent $\mu(\mathbf{r}, E)$ and the polychromatic X-ray spectrum $N_0(E)$ (Eq. 2.3).

measurable electrical signal. This is often done either directly or via a scintillator that first converts the X-rays into visible light [6]. The full image is acquired by measuring signals along many rays through the object using a pixelated detector. This introduces additional factors to the detection model associated with the shape of the incident X-ray spectrum and specific detector properties. Here, we approximate a more wholistic measurement model as

$$m(\mathbf{r}_\perp) = \Delta x^2 \int dE D(E) \eta(E) I(\mathbf{r}_\perp, E) * \text{psf}(\mathbf{r}_\perp, E), \quad (2.4)$$

where $m(\mathbf{r}_\perp)$ is the detected image, Δx is the width of each detector element (assumed square), $D(E)$ is a detector response function, $\eta(E)$ is a detective efficiency function, $I(\mathbf{r}_\perp, E)$ is the X-ray flux (number of photons per unit area), $*$ indicates the 2D convolution with respect to $\mathbf{r}_\perp$, and $\text{psf}(\mathbf{r}_\perp, E)$ is the system point-spread function (PSF). The detector response characterizes how the detector converts a photon to a real measurement, for example $D(E) \sim E$ for a scintillator that emits visible light in proportion to the photon’s incident energy. The efficiency function characterizes the number of photons that are stopped by the detector and might be estimated using the Beer–Lambert law as $\eta(E) = 1 - \exp[-\mu_{\text{det}}(E)T_{\text{det}}]$. The system PSF can be characterized by its source and detector components,

$$\text{psf}(\mathbf{r}_\perp, E) = \text{psf}_{\text{src}}(\mathbf{r}_\perp, E) * \text{psf}_{\text{det}}(\mathbf{r}_\perp, E). \quad (2.5)$$

Note that while this model is more general, it still makes a number of approximations, and a more thorough approach is provided by Barrett and Myers [71]. For example, the PSF generally varies due to variations in source anode angle, obliquity factors, and magnification as a function of detector pixel location [72]. Additionally, $D(E)$ and $\eta(E)$ can depend on $\mathbf{r}_\perp$, and it is often necessary to implement detector-element-specific correction factors to avoid artifacts in detected images. This is often done via an air scan, which is a blank measurement acquired without an object. In the Eq. 2.4 model we implicitly assume that such correction factors are already applied and thus drop the $\mathbf{r}_\perp$ dependence.

The model additionally assumes that the system is linear and shift-invariant (LSI), specifically meaning that it has the following two properties [71]:

- **Linearity:** The image of two superimposed objects can be expressed as a linear combination of the images of each individual object, i.e. for linear system operator $\mathcal{S}$, object functions $f_1(\mathbf{r})$, $f_2(\mathbf{r})$, and scalars a, b :

$$\mathcal{S}[af_1(\mathbf{r}) + bf_2(\mathbf{r})] = a\mathcal{S}[f_1(\mathbf{r})] + b\mathcal{S}[f_2(\mathbf{r})]. \quad (2.6)$$

- **Shift-invariance:** A transverse spatial shift of the object results only in a corresponding shift of the output image, i.e. if $g(\mathbf{r}) = \mathcal{S}[f(\mathbf{r})]$, then:

$$\mathcal{S}[f(\mathbf{r} - \mathbf{r}_0)] = g(\mathbf{r} - \mathbf{r}_0). \quad (2.7)$$

While this cannot hold exactly for arbitrary imaging systems or arbitrarily large transverse shifts, the LSI framework is generally a reasonable approximation for radiological imaging systems over a limited field of view and therefore commonly implemented. Since LSI operators preserve vector operations, this introduces several convenient properties useful for later forward modeling.

2.2 X-ray sources

We now consider realistic energy spectra produced by common X-ray sources.

2.2.1 X-ray tubes

The most common X-ray source for medical imaging is an X-ray tube, composed of a cathode and anode contained within a glass envelope [73]. The cathode is usually a heated filament, which produces a cloud of electrons in proportion to the tube current or “mA” (within the effective operating range). Cold electron sources, such as carbon nanotubes (CNTs), have also been proposed [74]. The electrons are accelerated across the tube due to the potential difference between the cathode and anode. Typical tube voltages are around 50–150 kV in diagnostic imaging, with the specific value chosen depending on the task [6]. This potential difference is also called the peak tube voltage “kVp.” The accelerated electrons arrive at the anode with energy equal to $e \cdot \text{kVp}$, where $e \approx 1.602 \times 10^{-19}$ C is the elementary charge. X-rays are produced at a range of energies through bremsstrahlung and hard collisions with the anode [70]. X-rays exit the tube through a collimated window, which is directed to face the object for imaging applications.

One practical consideration is the anode target. The bombarding electrons are focused onto a small, finite-sized focal spot which affects the system’s spatial resolution via $\text{psf}_{\text{src}}(\mathbf{r}_{\perp})$. They deposit most of their energy as heat, and only a small portion (less than 1%) becomes X-rays [4]. For a good-quality image, the tube current and exposure time (mAs) must be high enough to generate sufficient X-ray flux but not so high as to damage the anode. Tungsten is a common anode material due to its high atomic number and high melting point, which makes it less prone to damage. Molybdenum and rhodium are also common choices in mammography [6]. Increasing the focal spot size will improve heat loading but degrade spatial resolution by broadening the PSF. This is somewhat alleviated by using an angled anode; along the central ray to the detector, the effective focal spot size is reduced

by a factor of $\sin(\theta)$, where θ is the anode angle [72]. Further, most anodes are designed as discs that rotate anywhere from 3,000 to 10,000 times per minute in order to spread the deposited energy over a much larger area of the anode [6]. Nonetheless, this results in some unavoidable trade-offs regarding high spatial resolution (small effective focal spot size) and physical limitations of the system. Diagnostic imaging typically uses focal spot sizes on the order of 1 mm, though smaller sizes are also used for high-resolution applications such as imaging breast calcifications.

Microfocus X-ray sources have also been developed, mainly for benchtop imaging applications. Using the same anode-impact principle, these sources can achieve focal spot sizes orders of magnitude smaller, around the range of 5–20 μm and down to single microns, with either conventional solid anodes or liquid-metal targets [4, 73]. Such sources are attractive for high-resolution imaging experiments, though their broad clinical implementation is limited by practical considerations including cost, flux, and heat loading. Novel alternatives continue to be explored, such as a recently proposed target made of fast-moving tungsten microparticles [75].

2.2.2 Synchrotron light sources

An alternative X-ray source is synchrotron radiation. In a synchrotron facility, electrons are accelerated to relativistic speeds with GeV energies and stored in a high-vacuum ring [4]. Magnetic fields deflect the electron beam to produce synchrotron radiation, which can branch off to individual beamlines for experiments. Relative to conventional X-ray tubes, synchrotron sources can provide much higher flux, smaller focal-spot size, low divergence, and high coherence [76]. With additional beamline optics such as crystal monochromators, the radiation can be filtered to produce nearly monochromatic beams with bandwidths as low as $\Delta E/E \sim 10^{-4}$. These properties make synchrotron light a good X-ray source choice for micro-CT and high-resolution phase-contrast imaging. A key disadvantage is accessibility,

since synchrotrons are large-scale user facilities with limited beamtime, making them better suited to scientific research applications.

2.2.3 Beam geometry

Another consideration is the geometry of how the X-rays propagate from the source and through the object [7]. Anode-based X-ray tubes are a type of point source, meaning X-rays radiate outward isotropically in a “cone beam” geometry. When only a single axial slice is considered, as in some early CT scanners, this is often called a “fan beam” geometry. With a point source, the object must be placed sufficiently far from the source so that the solid angle subtended by the X-ray beam entirely encompasses the object to be imaged. This also introduces a magnification factor M that increases the source PSF and minimizes the detector PSF/pixel size when projected onto the object plane. Clinical X-ray imaging methods like radiography and mammography usually use a contact-imaging approach, where the detector is placed immediately behind the patient, limiting undesirable magnification of the focal spot. If the point source is assumed to be infinitely far from the object, as for synchrotron sources, this becomes a roughly “parallel beam” geometry. This is only feasible when the source PSF is essentially negligible, otherwise the X-ray beam would be entirely defocused and the exit image uselessly blurry. A parallel beam geometry could also be theoretically achieved by an array of collimated point sources.

2.3 X-ray detectors

As described in the general measurement model (Eq. 2.4), the X-ray detector introduces a number of factors including a detector response function, efficiency function, and PSF. These vary depending on the type of detector used. The most common detectors for diagnostic imaging can be broadly categorized into energy-integrating detectors (EIDs) and photon-counting detectors (PCDs).




Figure 2.3: Schematic of two common X-ray detectors: (left) the indirect-conversion scintillator-based EID used on most clinical CT scanners, and (right) direct-conversion semiconductor PCD that is being progressively implemented clinically.

2.3.1 Energy-integrating detectors

By definition, an EID weights each incident photon by its energy and sums over these components to produce a measured signal with detective function $D(E) \sim E$ [48]. This is generally not a desirable property for an X-ray detector, as higher energies reduce image contrast, but rather a practical consequence of developing a detector that can be conveniently manufactured and implemented in commercial devices. Most current X-ray detectors for CT and radiography use an indirect-conversion EID scheme (Fig. 2.3), in which X-rays are first converted to visible light and then collected using a photodiode to measure an electrical signal [7]. The most common type is a solid-state indirect-conversion EID made of segmented scintillator blocks. X-rays interact via the photoelectric effect, producing a photoelectron that briefly excites other electrons in nearby atoms. When the excited electrons return to their ground state, visible and ultraviolet light is emitted in the form of characteristic radiation. The photodiodes capture a portion of this light, which becomes the detected signal. To avoid this light contaminating adjacent detector channels, it is necessary to coat the blocks with a reflective material. This can reduce the detector’s geometric efficiency.

Direct-conversion EIDs are also possible, bypassing the intermediate step of optical photon production. This approach was commonly used in the early days of CT, with detectors filled of high-pressure xenon gas that would be ionized by passing X-rays, allowing direct




Figure 2.4: Schematic of theoretical electrical readout for a spectral PCD channel (left) and the resulting detected energy spectra (right). The bin spectra were generated using an approximate non-ideal spectral threshold scheme where each interacting X-ray follows a Gaussian probability distribution of being detected at its true energy. More precisely, the bin spectra are distorted by other nonidealities such as charge sharing, K-escape, and pulse pile-up.

collection of the resulting charges [7]. While this xenon-gas detector was relatively low-cost, it struggled with low detective efficiency due to the inherently low density of gas and engineering challenges in designing a multi-row detector for 3D CT imaging, thus resulting in the shift to the conventional indirect-conversion EID used on most CT scanners today.

2.3.2 Photon-counting detectors

A PCD model equally weights photons, regardless of energy, with a simple counting scheme represented by $D(E) \sim 1$. This detector response function alone improves contrast relative to EIDs by no longer overweighting high-energy photons. The main practical implementation of this detector type, the direct-conversion semiconductor PCD (Fig. 2.3), also has a number of advantages over typical EIDs. In these detectors, X-rays interact with a semiconductor material, producing a charge cloud of electron-hole pairs that are attracted towards pixelated anodes and a monolithic cathode, respectively [7]. The charge collected by each anode is read out into separate detector channels. The combination of a monolithic semiconductor and pixelated anodes removes the need for reflective layers between scintillator blocks as

used for EIDs, improving geometric efficiency and allowing for much smaller pixel sizes. Each channel registers an electrical signal over time. X-ray interactions appear as spikes in this signal with pulse height proportional to their energy¹ (Fig. 2.4) [77]. The full electrical current–time readout is too large to be digitized and stored, so a processing circuit is designed to apply a pulse-height threshold for counting the X-ray interaction spikes in each detector channel (thus resulting in $D(E) \sim 1$). This pulse-height threshold can be selected to safely eliminate background electronic noise (typically around 10–20 keV for clinical CT), which is not possible with conventional EIDs [48, 49]. Further pulse-height thresholds may be applied to extract spectral information.

These PCDs have the advantages of smaller pixel sizes, no electronic noise, and native spectral information, but they contend with other challenges. When X-ray interactions occur near the boundary between two anodes, the charge cloud can be split into two smaller measurements in what is known as “charge sharing.” This causes two issues: multiple small pulses are registered across different detector channels for what should truly be a single large pulse, and spectral information is distorted when these pulses fall into an incorrect energy bin. This phenomenon can also be caused by secondary interactions of the electrons and escape of K-fluorescent photons. As pixel sizes get smaller, these distortions become more severe. Additionally, after each X-ray interaction, there is a small period of dead time before the detector is ready to count a new independent interaction. At high flux levels, multiple independent X-ray interactions are more likely to occur near-simultaneously, resulting in dead-time losses from “pulse pile-up.” However, since these effects cause somewhat predictable, there is extensive ongoing development of correction algorithms, such as coincidence detection to correct for charge sharing and empirical models to correct for spectral distortion [78–83]. In fact, the charge sharing phenomenon has even been leveraged advantageously to achieve micron-scale spatial resolution by fitting a Gaussian distribution to the pulses

1. The proportionality approximation $D(E) \sim 1$ is approximately true in the range 30–100 keV, which is most of the relevant diagnostic X-ray imaging range [48]. At higher energies, $D(E)$ slowly decreases.

measured in adjacent detector channels [62–64].

The most common semiconductor materials implemented in clinical and preclinical PCD-CT systems are cadmium telluride (CdTe), cadmium zinc telluride (CZT), and silicon (Si) [49]. CdTe and CZT have a higher atomic number ($Z_{\text{Cd}}=48$, $Z_{\text{Zn}}=30$, and $Z_{\text{Te}}=52$), which results in a higher stopping power such that they can be made just a few millimeters thick without sacrificing detective efficiency. Si has a lower atomic number ($Z_{\text{Si}} = 14$) and thus must be made a few centimeters thick to have comparable efficiency. The difference in atomic composition also affects the proportion of different photon interactions: the photoelectric effect is dominant in high- Z materials like CdTe and CZT, while Compton scatter is dominant in Si. This can cause greater count-multiplicity distortions in Si when Compton-scattered photons are detected more than once. On the other hand, high- Z detectors are more likely to experience K-fluorescence after photoelectric interactions, which causes spectral distortion. Another note is that Si is a well-understood semiconductor with a strong supply chain of silicon wafers due to their prevalence in the electronics industry, whereas the manufacturing infrastructure for CZT and CdTe is still in development [49]. Thus, even accounting for the advantages, one must content with these trade-offs and challenges in implementing semiconductor PCDs.

Alternatively, PCDs can be manufactured using an indirect-conversion scheme. Gamma cameras used in nuclear medicine imaging are one such example [69] These types of scintillator-based detectors had been assumed to be too slow for counting readout modes with the higher photon flux rates of X-ray imaging. However, recent work has indicated that certain materials such as lanthanum bromide doped with cerium ($\text{LaBr}_3\text{:Ce}$) can be just as fast as CZT and manufactured at 10–15% of the cost [84, 85]. This is an ongoing line of investigation, and direct-conversion PCDs are still largely dominant in current clinical implementations.




Figure 2.5: Two common CT acquisition geometries. (Left) Clinical fan-beam CT scanner using a source-detector array that rotates in tandem about the patient. (Right) Parallel-beam micro-CT with a stationary synchrotron X-ray source and rotating sample.

2.4 Computed tomography

Traditional X-ray radiography suffers from the problem of superposition, with 3D structural information projected onto a single 2D plane, resulting in indistinguishable overlapping structures and greatly reduced contrast [7]. This is especially true of weakly attenuating structures, which cause only subtle changes within a relatively constant background material. A famous example in neuroradiology is the task of visualizing the ventricles of the brain, which prior to the 1970’s had been considered impossible without pneumoencephalography—an uncomfortable procedure that required lumbar puncture to replace the cerebrospinal fluid (CSF) with air [86].

The invention of computed tomography (CT) largely solved such problems. CT scanners operate by acquiring several projection images at different angles θ around the object (Fig. 2.5). Simply having more angular views available visually allows one to begin resolving overlapping structures by their spatial shifts, and with sufficiently high sampling, one can mathematically reconstruct full images of the internal anatomy [30–32]. With brain CT, CSF

could be natively distinguished from surrounding gray and white matter, sparing patients further pneumoencephalography [86].

To formulate the CT reconstruction problem, we first note that radiography yields images that reflect the line integrals through the object’s linear attenuation coefficient, the negative-log intensity measurements $p(\mathbf{r}_\perp) \sim \int dz \mu(\mathbf{r}_\perp, z)$. CT scans acquire many of these line integral measurements at different angles, creating a sinogram $p(\mathbf{r}_\perp, \theta)$. The reconstruction operation produces spatially resolved maps of the linear attenuation coefficient, $\mu(\mathbf{r}_\perp)$. The art of the problem lies in performing said reconstruction. Analytical inverse solutions can offer near exact reconstructions under idealized conditions, such as the Fourier slice theorem, filtered back-projection (FBP), and Feldkamp–Davis–Kress (FDK) methods for different geometries [87, 88]. However, these can amplify noise from the raw images and ignore other realistic image model elements such as detector and source nonidealities. There has been extensive work developing better methods capable of accounting for these factors via iterative reconstruction (IR) and artificial intelligence (AI) approaches [89]. Iterative methods frame CT reconstruction as an optimization problem, which can be broadly posed as

$$\mathbf{x}^\star = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}), \quad (2.8)$$

where $\mathcal{L}$ is a loss function and $\mathbf{x}$ is the object to reconstruct [90]. The loss function is typically composed of a fidelity term, which compares the projected solution to the truly measured data, and a regularization term, which enforces other desirable properties such as smoothness and sparsity.

Once the loss function has been defined, the minimization problem is solved by providing an initial guess of the object and applying iterative update steps until a convergence criterion is met [90]. While iterative methods can reduce image noise and encourage more realistic image reconstruction through regularization, they can be computationally intensive [89]. This has motivated interest in machine learning methods, which are computationally

intensive to train but much more efficient in their actual implementation. Nonetheless, AI approaches must contend with the typical pitfalls of acquiring a sufficiently general training dataset, avoiding overfitting of the training data, and eliminating hallucinations. There is ongoing work exploring hybrid AI-IR methods as possible solutions to the drawbacks of each method alone; for example, using AI to optimize an iterative reconstruction approach or to design a regularization term [91–93].

2.5 Spectral imaging and basis decomposition

In Chapter 1, we outlined the history of spectral imaging and alluded to the many benefits that may arise from multiple energy-resolved measurements. The primary way to realize these benefits is to use the independent measurements to solve for a new set of images, which can provide quantitative insights or improve the detectability of subtle abnormalities. This process is called basis decomposition. Whereas raw intensity images simply represent the proportion of incident X-rays stopped by the detector, basis decomposition problems can be posed to yield other physically meaningful images such as interaction probabilities, electron density and atomic number distributions, or individual material thickness maps [33, 34].

2.5.1 *Interaction-probability decomposition*

The key idea enabling basis decomposition is the unique energy dependence of different materials’ linear attenuation coefficients $\mu(E)$. As introduced in Eq. 2.1, in the diagnostic imaging energy range, the attenuation coefficient can be represented as the sum of contributions from two types of photon interactions,² $\mu(E) = \mu_{\text{PE}}(E) + \mu_{\text{CS}}(E)$, where PE indicates the photoelectric effect and CS indicates Compton scatter. Each interaction has different

2. More precisely, note that this is valid for diagnostic X-ray imaging, with energies roughly in the range of 1–150 keV. In different energy ranges, the exponential coefficients for Eq. 2.9–2.10 change slightly, and pair-production is also possible at energies above 1.022 MeV [70].

functional dependence on X-ray energy and material atomic composition, with

$$\mu_{\text{PE}}(E) \propto \frac{Z^3}{E^3}, \quad (2.9)$$

$$\mu_{\text{CS}}(E) \propto \rho_e f_{\text{KN}}(E), \quad (2.10)$$

where Z is atomic number, ρ_e is electron density, and $f_{\text{KN}}(E)$ is the Klein–Nishina function [33, 70]. Then, it is valid to represent an inhomogeneous object in terms of two energy-dependent basis functions,

$$\mu(\mathbf{r}, E) = a_1(\mathbf{r})f_1(E) + a_2(\mathbf{r})f_2(E), \quad (2.11)$$

where $a_1(\mathbf{r}), a_2(\mathbf{r})$ account for the spatial distribution of the object and $f_1(E), f_2(E)$ are basis functions reflecting the energy dependence of the interaction probabilities,

$$f_1(E) = 1/E^3, \quad (2.12)$$

$$f_2(E) = f_{\text{KN}}(E). \quad (2.13)$$

The key insight of Alvarez and Macovski was to formulate the attenuation coefficient in this manner and successfully fit experimental measurements to these functions [33]. In the context of imaging, recall our ideal signal detection model (Eq. 2.3). Substituting our new expression for $\mu(\mathbf{r}, E)$ from Eq. 2.11 yields a two-component model

$$p(\mathbf{r}_\perp) = f_1(E) \int dz a_1(\mathbf{r}_\perp, z) + f_2(E) \int dz a_2(\mathbf{r}_\perp, z), \quad (2.14)$$

where

$$p(\mathbf{r}_\perp) \equiv \ln \left[\frac{I_0}{I(\mathbf{r}_\perp)} \right] \quad (2.15)$$

indicates the negative log-intensity measurement, and we have defined the ray to be directed along the z axis and thus separated the 3D spatial coordinate into its components $\mathbf{r} \equiv (\mathbf{r}_\perp, z)$.

Say we acquire two ideal monochromatic $p_i(\mathbf{r}_\perp)$, with incident X-ray spectra each comprising a single energy E_i ($i = 1, 2$). The two-measurement system can be represented in terms of its basis components in matrix form,

$$\begin{bmatrix} p_1(\mathbf{r}_\perp) \\ p_2(\mathbf{r}_\perp) \end{bmatrix} = \begin{bmatrix} f_1(E_1) & f_2(E_1) \\ f_1(E_2) & f_2(E_2) \end{bmatrix} \begin{bmatrix} A_1(\mathbf{r}_\perp) \\ A_2(\mathbf{r}_\perp) \end{bmatrix}, \quad (2.16)$$

where

$$A_j(\mathbf{r}_\perp) = \int dz a_j(\mathbf{r}_\perp, z). \quad (2.17)$$

Then the basis images $A_j(\mathbf{r}_\perp)$ can be immediately solved by taking the inverse of the 2×2 matrix of basis functions $f_j(E_i)$ on a pixel-wise basis for each $\mathbf{r}_\perp$. When these functions take the form of Eq. 2.13, the basis images $A_j(\mathbf{r}_\perp)$ represent the spatial distribution of the probability of photoelectric and Compton scatter interactions.

Note, this matrix-inversion solution technique is an idealization valid for true monochromatic measurements. With polychromatic incident spectra, the left-hand side measurement vector $\vec{p}$ could still in theory be taken as the raw intensity images, but construction of the 2×2 matrix on the right-hand side necessitates some single monoenergy-equivalent E_i for each acquisition. In practice, this is sometimes defined as an “effective energy” equal to the intensity-weighted sum energy of each acquisition’s detected spectrum. While this is a fast and convenient decomposition method, we show in Chapter 9 that such approximations can lead to undesirable inaccuracies in the retrieved basis images. More wholistically, iterative basis decomposition methods can be used to account for the full spectrum of polychromatic

acquisitions, with potential added benefits of regularized loss functions.

2.5.2 *Material decomposition*

The physical justification for two-image decomposition lies with the two primary X-ray interactions in the diagnostic imaging energy range. This same principle can justify decomposition into other forms of basis functions $f_j(E)$, thus giving rise to other physical meanings for the corresponding A_j . The most commonly chosen basis is one of “assumed materials” [34].

Assuming the object comprises just two materials, the line integrals through the object can be simply written as

$$\int dz \mu(\mathbf{r}_\perp, z, E) = \mu_1(E) T_1(\mathbf{r}_\perp) + \mu_2(E) T_2(\mathbf{r}_\perp), \quad (2.18)$$

where

$$T_j(\mathbf{r}_\perp) = \int dz t_j(\mathbf{r}_\perp, z) \quad (2.19)$$

is an image of the projected thickness of each material along the z -axis. This again separates the energy and spatial dependence of our object representation, such that the basis material thickness maps can be solved with two ideal monochromatic measurements, with the linear attenuation coefficients of each material as the basis functions, i.e. $f_j(E_i) \equiv \mu_j(E_i)$ and $A_j(\mathbf{r}_\perp) \equiv T_j(\mathbf{r}_\perp)$ in Eq. 2.16.

For dual-energy CT imaging, one can perform material decomposition on each set of raw dual-energy projection measurements to construct a material thickness sinogram $T_j(\mathbf{r}_\perp, \theta)$. When input to a typical CT reconstruction algorithm, this yields an image of $\rho_j(\mathbf{r}_\perp)$, a material density map.

The success of the material decomposition approach depends on how well the object can be approximated by the chosen materials. The original set of basis materials proposed by




Figure 2.6: Linear attenuation coefficients in the X-ray energy range for some materials relevant to diagnostic imaging. Water has no apparent K-edge; calcium has a K-edge too low to be apparent in most images; iodine and gadolinium have sufficiently high K-edges to be candidate contrast agents for K-edge imaging. Also note that L-edges are visible for iodine and gadolinium around 4–8 keV. All materials were defined with density $\rho = 1 \text{ g/cm}^3$.

Alvarez and Macovski was bone and water, which provides a good first approximation for humans [34]. Aluminum/PMMA and bone/soft tissue are other common choices. Traditionally chosen bases tend to comprise one highly attenuating material and one weakly attenuating material, as this has shown the best success in medical imaging. The main violation of the energy-dependence assumptions made so far involves the manual introduction of high- Z materials with significant K-edges.

2.5.3 *K-edge imaging*

The K-absorption edge refers to a material-specific energy value above which the photoelectric interaction probability increases suddenly. This is due to a threshold effect when the incident X-ray is just energetic enough to excite K-shell electrons, which had zero probability of interaction with lower-energy X-rays [70]. In fact, each electron shell has a threshold energy for photoelectric interactions. The K-shell simply has the highest such threshold since it is closest to the nucleus and therefore has the greatest binding energy. For natural

biological materials, the K-edge tends to be far below the X-ray energy range and thus irrelevant for diagnostic imaging, making the basis material decomposition sufficient for typical dual-energy imaging. However, the effect can also be leveraged advantageously in spectral imaging by introducing high- Z materials with significant K-edges as contrast agents (Fig. 2.6). One can acquire an image with energy just above and just below the K-edge, then subtract the two to produce an image isolating the contrast agent. This can be in theory perfectly performed using monochromatic sources but is more challenging with conventional polychromatic spectra. Spectral PCDs offer an appealing solution, since their energy bin thresholds can be precisely set to match the known K-edge of the contrast agent [7].

2.5.4 *Virtual monochromatic images*

A downside of basis decomposition is that it can amplify noise. This can be ameliorated by recombining the basis images into a single virtual monochromatic image (VMI) using the original two-component model of Eq. 2.11. To solve the original decomposition problem, the basis functions $f_j(E_i)$ were taken as known and p_i as given measurements. To compute a VMI at some energy of interest E_0 , we instead take the newly-solved basis images $a_j(\mathbf{r}_\perp)$ as inputs and calculate $f_j(E_0)$ using their defined functional form, such that

$$\text{VMI}(\mathbf{r}_\perp, E_0) = a_1(\mathbf{r}_\perp)f_1(E_0) + a_2(\mathbf{r}_\perp)f_2(E_0). \quad (2.20)$$

This is essentially just a restatement of Eq. 2.11. Since the noise of the two basis images tends to be correlated, combining them in this manner can cancel out individually affected pixels and to improve overall CNR. Since VMIs can be synthesized at arbitrary energies, it is also possible to identify an optimal E_0 for maximizing CNR. For polychromatic spectral imaging, VMIs also have the benefit of acting as a beam-hardening correction.

2.6 Summary

To conclude, this chapter has laid out the necessary absorption-contrast imaging theory for the following sections. We have outlined two prominent approaches to improving contrast in diagnostic imaging: spectral imaging for basis material decomposition, and CT imaging for distinguishing overlapping structures. Furthermore, we have made explicit conventional assumptions in the imaging model and the limitations of typical source and detector technology. With this foundation, Chapters 3 and 4 present a possible approach to improving soft-tissue imaging for radiotherapy: spectral CT imaging with a megavoltage spectrum.

CHAPTER 3

DUAL-ENERGY CT WITH MEGAVOLTAGE AND KILOVOLTAGE SPECTRA

3.1 Introduction

When imaging for radiation therapy, high soft-tissue contrast is essential. Tumors must be accurately imaged at all stages: before treatment for dose calculations, during treatment for patient positioning, and after treatment for cancer monitoring [94]. The current imaging standard in radiation therapy is single-energy (SE) CT. While SE-CT provides clear contrast between bone and soft tissue, it has difficulty distinguishing variations within tissue structures. To achieve a sufficient contrast-to-noise ratio (CNR) to identify a tumor on a soft-tissue background, a higher patient dose may be necessary. With the advent of modern image-guided radiation therapy (IGRT) and its daily imaging requirements, individual image doses are in danger of accumulating over the course of treatment up to the scale of a therapy dose fraction [95]. Thus, there is interest in finding methods that achieve sufficient image quality for radiation therapy with reduced dose.

One potential solution could be on-board dual-energy (DE) CT. As outlined in Chapters 1 and 2, DE-CT is routinely implemented in diagnostic imaging to enhance image quality. With carefully chosen X-ray spectra and possible contrast agents, there are a number of spectral techniques that can be implemented to improve image quality [96–99]. A typical downside is the practical step of acquiring specialized DE-CT hardware for acquiring simultaneous energy-resolved measurements, whether that be a dual-source system, fast-kVp switching device, or spectral photon-counting detector. However, modern radiation therapy treatment systems are often already equipped with this hardware: two X-ray sources that rotate in tandem, a megavoltage (MV) source for treatment and a kilovoltage (kV) source for on-board imaging. While the primary purpose of the MV beam is radiation therapy treatment,

it can also be detuned to operate with a softer energy spectrum (peak photon counts at ≈ 1 MeV) more appropriate for imaging. Given this readily available equipment, there exists an unexplored potential for enhancing image quality in the context of radiation therapy using MV-kV DE-CT [98–101].

The purpose of this chapter is to assess and optimize a proposed MV-kV DE-CT acquisition scheme, laying the foundation for a more accessible pathway to enhanced soft-tissue contrast for radiotherapy imaging.¹

3.2 Background

A number of methods have been explored to improve image quality in radiation therapy settings without delivering excess patient dose. CT is currently the only accepted imaging modality for 3D treatment planning dose calculations, as it provides empirical information on electron density and atomic composition. Unfortunately, X-ray imaging modalities like CT require increased dose to yield sufficient soft-tissue contrast for this purpose and additionally suffer from a lack of functional information and, in some cases, metal artifacts. To overcome these drawbacks, there has been recent interest in magnetic resonance (MR) and positron emission tomography (PET) imaging at different stages of the treatment planning process [94, 101]. MR- and PET-linac systems are currently being developed and even introduced in some clinics [103, 104]. Though these systems show promise, they are emerging and expensive, and CT imaging still remains necessary for dose calculation. Combined MV-kV imaging has the unique advantage of utilizing existing hardware, resulting in a much lower barrier to implementation and the potential to be realized in a shorter time frame. Moreover, MV images have the potential to be acquired during patient treatment, providing dual energy information to augment kV images without necessitating additional dose [100, 105, 106].

1. This chapter is based on work published by SPIE in [102] under a Creative Commons Attribution 4.0 International License. [DOI: 10.1117/1.JMI.11.2.023501]

While some prior work has explored ways to combine MV and kV information for various applications, work specifically exploring MV-kV DE-CT is limited. This may be due to the drawbacks of MV images, which can potentially contaminate kV images and reduce image quality if the two are combined naïvely. Since MeV photons are generally more penetrating than keV photons and have greater dose deposition per photon, MV images typically have lower contrast and greater noise relative to kV images with the same dose [106, 107]. There are limited situations in which MV imaging alone is sufficient. For example, MV localizers or CT images have been shown to be sufficient for radiotherapy setup verification, which is useful for linacs lacking a kV X-ray source [108]. The superior quality of kV imaging is still necessary for initial treatment planning.

When strategically implemented, MV information can be combined with kV information to achieve better soft-tissue contrast than could be achieved with either source alone and an equivalent total dose. Various techniques have been described previously [99, 106, 107, 109–115]. For example, the greater penetrability of MV photons becomes an advantage when imaging highly attenuating objects; thus, MV data may be synthesized with kV images for metal artifact correction [112, 115]. Similar methods have been implemented for target tracking during radiotherapy [110, 111].

Previous work has shown promise for MV-kV DE-CT [99, 106, 107, 109, 113, 114]. With only partial angular information, combined MV-kV CT images can achieve superior image quality relative to kV or MV alone by inheriting the advantages of each single-energy image: higher contrast from kV images and reduction in streak artifacts from MV images. These techniques utilize 90–110° of data from each spectrum, possibly with a small amount (10–15°) of overlap, then implement either a linear gray scale conversion or histogram mapping of pixels to reconstruct a single kV or MV image with complete angular information. Though these methods improve efficiency by reducing rotation time, the imperfect spectral mapping can cause artifacts and reduce image quality. To our knowledge, previous work focusing

specifically on MV-kV DE-CT has not considered datasets with complete angular information in both MV and kV domains.

The precise conditions for MV-kV DE-CT to yield improved image quality with equivalent dose have not been fully characterized. The experimental nature of past work, utilizing image quality phantoms, limits the number of data points feasible for robust analysis and optimization of MV-kV DE-CT. These methods typically acknowledge the drawback of greater dose deposition by MeV photons but do not consider the effect of dose distribution between MV and kV spectra. They observe that image quality is best when MV beams penetrate more highly attenuating material inserts, but they image only a small number of material inserts within image quality phantoms. A more complete analysis should assess image quality for a continuum of material attenuations and spectral dose distributions.

We implement the analytical method proposed by Roessl et al., which uses estimation theory to characterize basis material signal-to-noise ratio (SNR) along a single-line integral with dual-energy X-rays [66]. This method has not yet been applied to MV-kV imaging. With this technique, we are able to survey a wide range of parameters without experimental measurements, facilitating more robust characterization of situations in which MV-kV imaging provides superior image quality and quantification of the degree of improvement. This similarly allows us to assess whether it is possible to achieve equivalent image quality and reduced dose with MV-kV imaging. We consider three variables: spectral combination, dose allocation between the two spectra, and material composition. To gauge utility of MV-kV DE-CT in a more clinically realistic setup, we also extend this model to a CT raytracing simulation with parameters informed by the single-line optimization.

3.3 Methods

Two models were developed for quantifying basis material image quality: a toy model using a single-line integral with a two-material object and a fan-beam CT simulation with

a computational anthropomorphic phantom. The single-line model was used to maximize basis material SNR as a function of spectral dose allocation for an object with various bone thicknesses. These results were used to inform the simulated CT imaging task.

3.3.1 Signal detection framework

The detected signal measurements m_i ($i = 1, 2$) were calculated using a discretized form of the signal detection model,

$$m_i = \sum_E D(E) \text{Poi} \left\{ I_i(E) \exp \left[- \int d\ell \mu(\mathbf{r}, E) \right] \eta(E) \Delta E \right\} \quad (3.1)$$

where $I_i(E)$ is the incident X-ray spectrum (photons per energy), ΔE is the discrete energy bin width, and the notation $\text{Poi}\{\bar{x}\}$ indicates generation of a realization of a Poisson random variable with mean $\bar{x}$. In Eq. 3.1, the argument of the Poisson noise is the mean number of photons of energy E stopped by the detector. This indicates that the signal noise model is compound Poissonian weighted by $D(E)$.

An energy-integrating detector was modeled ($D(E) = E$) with detective efficiency $\eta(E)$ as shown in Fig. 3.1. The detective efficiency function was computed to yield performance consistent with that of a previously described high-detective quantum efficiency (DQE) xenon gas detector [116, 117]. Such detectors have been implemented on commercial tomotherapy units (Accuray Inc., Sunnyvale, CA) for fan-beam MV CT imaging [118, 119].

Input spectra

Five polychromatic spectra $I_i(E)$ were modeled (Fig. 3.2). Three kV spectra were chosen to represent common diagnostic CT options (80kV, 120kV, and 140kV). The two MV spectra were modeled after a typical treatment beam (6MV) and a treatment beam with energy detuned to below 3 MV for imaging (detunedMV) [120, 121]. To ensure a valid basis for




Figure 3.1: The modeled detective efficiency function $\eta(E)$.

comparison, the flux of each spectrum was scaled to deliver the same dose to the center of a 40-cm diameter water cylinder (depth $d_w = 20$ cm) under the condition of charged particle equilibrium,

$$\text{Dose}(d_w) = \int_{E_{\min}}^{E_{\max}} dE I_i(E) \exp[-\mu_w(E)d_w] \left(\frac{\mu_{\text{en}}(E)}{\rho} \right)_w E, \quad (3.2)$$

where $\mu_w(E)$ is the linear attenuation coefficient of water and $(\mu_{\text{en}}(E)/\rho)_w$ is the mass energy absorption coefficient of water [70].

3.3.2 Single-line integral through a two-material object

Basis material SNR was computed using an estimation theory framework for a single ray incident on a two-material object of ICRU tissue and bone with densities $\rho_{\text{tissue}} = 1.00$ g/cm³ and $\rho_{\text{bone}} = 1.85$ g/cm³ [1, 66]. The tissue thickness was fixed at $t_{\text{tissue}} = 40$ cm and the bone thickness was varied from $t_{\text{bone}} = 1$ to 10 cm. All MV-kV and kV-kV spectral pairs were considered, yielding nine DE combinations. The total single-line dose allocated to both spectra was set to 1 μGy . For a typical CT acquisition with near 1000 projection views, this




Figure 3.2: The five spectra with magnitude scaled to deliver 1 μGy dose.

would sum to a dose of 1 mGy. SNR was characterized as a function of dose allocation r in 1% increments (from 1% to 99%), where the high-energy spectrum was rescaled by r and the low-energy spectrum by $1 - r$.

The SNR for each basis material j was defined as the ratio of the true mass thickness ($A_j \equiv \rho_j t_j$) to the square root of the Cramér-Rao Lower Bound (CRLB) on variance,

$$\text{SNR}_j = \frac{A_j}{\sigma_{A_j}}. \quad (3.3)$$

The CRLB was found from the Fisher information $\mathcal{F}$ using the known relation $\sigma_{A_j}^2 \geq \mathcal{F}_{jj}^{-1}$ [122]. The signal noise model, which was energy-weighted compound Poisson, was approximated as a Gaussian with mean and variance matching the first two moments of the true distribution [97]. The mean of each measurement is simply m_i as in Eq. 3.1, and the corresponding variance σ_i^2 is

$$\sigma_i^2 = \sum_E D^2(E) I_i(E) \exp \left[- \int d\ell \mu(\mathbf{r}, E) \right] \eta(E) \Delta E. \quad (3.4)$$

Thus, each DE acquisition ($i = 1, 2$) yields a Fisher information matrix with elements that




Figure 3.3: The two computational phantoms imaged (noiseless 80-keV monoenergetic images). The metal hip replacement is indicated by the arrow in panel (b). In the modeled hip prosthesis, the femoral head and outer shell were either titanium or stainless steel, and the inner lining was PMMA. CNR was computed using the delineated ROIs, and RMSE was evaluated within the phantom.

can be written in terms of m_i and σ_i^2 and their partial derivatives with respect to the true mass thicknesses A_j :

$$\mathcal{F}_{pq} = \sum_{i=1,2} \frac{1}{\sigma_i^2} \frac{\partial m_i}{\partial A_u} \frac{\partial m_i}{\partial A_v} + \frac{1}{2} \sum_{i=1,2} \frac{1}{(\sigma_i^2)^2} \frac{\partial \sigma_i^2}{\partial A_u} \frac{\partial \sigma_i^2}{\partial A_v}, \quad (3.5)$$

where i is the spectral index and u, v are the basis material indices [66].

3.3.3 Fan-beam CT simulations of an anthropomorphic phantom

To assess whether MV-kV DE-CT may provide advantages in clinical imaging scenarios, the signal detection framework was extended from the single-line model to a fan-beam CT geometry with 1200 views, 800 detector channels, and a fan angle of 47° . A single 360° rotation was simulated for each acquisition. Beam transmission through a computational anthropomorphic phantom, the extended cardiac torso (XCAT), was calculated [123]. The phantom had dimensions of 512×512 with 1 mm^2 pixels. Path lengths through each pixel were determined using Siddon's algorithm for the exact radiological path through a CT array [124].




Figure 3.4: Linear attenuation coefficients of metals common in modern hip prostheses (cobalt-chromium alloys, titanium alloys, and surgical-grade stainless steel) in (a) the keV energy range and (b) the MeV energy range.

Since MeV photons are generally more penetrating than keV photons, the effect of high attenuation was considered by imaging the pelvis region with and without a metal hip replacement (Fig. 3.3) [125]. A hip prosthesis comprises three primary components: a femoral head, a cap or lining, and an outer shell. Each part must be wear-resistant, biocompatible, and capable of bearing high static and dynamic loads; tradeoffs among these factors inform the proper choice of materials for each patient’s unique presentation [126]. Metals have traditionally been and continue to be the central component of implants for total hip arthroplasty [125–127]. Contemporary metals of choice are stainless steel, titanium alloys, or cobalt-chromium (Co-Cr) alloys [126]. Figure 3.4 shows the linear attenuation coefficients of these materials, indicating that titanium alloys are generally less attenuating than Co-Cr alloys and stainless steel. To represent these typical attenuation categories, we modeled two types of metal implants: (i) commercially pure titanium and (ii) surgical grade stainless steel for both the femoral head and shell with a PMMA lining.

The total dose of each DE-CT acquisition was set to 10 mGy. Sinograms were generated

for two spectral pairs, “MV-kV” (9 mGy detunedMV and 1 mGy 80kV) and “kV-kV” (5 mGy 140kV and 5 mGy 80kV). Each sinogram pair was decomposed into ICRU tissue and bone basis materials using an iterative Gauss–Newton algorithm [128]. While basis material decomposition may occur in either the sinogram or image domain, the sinogram-domain decomposition was chosen because of its advantage of ameliorating beam hardening artifacts. The basis material sinograms were then reconstructed into basis material images (BMIs) using fan-beam filtered back projection (FFBP) including a general sinc window with cutoff frequency at 80% of the Nyquist frequency [7]. These BMIs correspond to the densities of the materials ρ_i . The reconstructed images had a matrix size of 512×512 and field-of-view of 50 cm. Virtual monoenergetic images (VMIs) were generated at various energies E_0 as a linear combination of the BMIs,

$$\text{VMI}(E_0) = \rho_1 \left(\frac{\mu(E_0)}{\rho} \right)_1 + \rho_2 \left(\frac{\mu(E_0)}{\rho} \right)_2, \quad (3.6)$$

where $(\mu(E_0)/\rho)_j$ is the known mass attenuation coefficient of basis material j at the energy E_0 .

For comparison, three 10 mGy single-energy (SE) acquisitions were also generated (80kV, 120kV, and 140kV) and reconstructed using the same FFBP algorithm. As the dual-energy VMIs have the advantage of ameliorating beam-hardening effects, a two-material beam-hardening correction (BHC) was applied to the SE images. A water correction was first applied in the sinogram domain using a fourth-degree polynomial remapping based on known attenuation characteristics [7]. A bone correction was then applied in the image domain by generating an artifact-only image from a thresholded “bone” image, which was linearly combined with the original image to create the final corrected image [7]. This allowed us to compare dose-matched and artifact-corrected SE- and DE-CT, which is more fair than comparison of each 10 mGy DE-CT acquisition with its noisier constituent <10 mGy SE-CT simulations. Better image quality with equivalent dose also indicates that dose can be

reduced while maintaining equivalent image quality.

To evaluate image quality, CNR was computed in each 10 mGy SE-CT image and DE-CT VMI using measurements from the ROIs delineated in Fig. 3.3. CNR was defined as

$$\text{CNR} = \frac{|\text{Avg}[\text{ROI}_1] - \text{Avg}[\text{ROI}_2]|}{\sqrt{\text{Var}[\text{ROI}_1] + \text{Var}[\text{ROI}_2]}}, \quad (3.7)$$

where ROI_1 is the signal and ROI_2 is the background. We also examined contrast and noise individually, defined as the numerator and denominator of Eq. 3.7, respectively.

To evaluate accuracy, the VMIs were registered to the input phantom, and root-mean-square-error (RMSE) was computed relative to the monoenergetic ground truth,

$$\text{RMSE} = \sqrt{\text{Avg}[(\text{VMI} - \text{XCAT})^2]}. \quad (3.8)$$

Note that SE-CT measurements are not energy-dependent, whereas CNR, RMSE, VMI, and XCAT include implicit energy dependence from each monoenergetic evaluation.

3.4 Results

3.4.1 Single-line integral

To identify the most promising spectral pairs, Fig. 3.5 and 3.6 present heatmaps of peak tissue and bone SNR, respectively, for each spectral combination and the ten bone thicknesses. Peak SNR was found from the curve of SNR as a function of dose allocation r . For both basis materials, the 140kV-80kV pair yields the highest SNRs overall. The detunedMV-80kV pair yields the highest SNRs of the MV-kV pairs. These two spectral pairs will be the focus of further analysis (“kV-kV” and “MV-kV,” respectively). Tissue SNR is maximized for both pairs with 1 cm bone, and bone SNR is maximized for kV-kV at 4 cm bone and MV-kV at 6 cm bone.




Figure 3.5: Heatmap of peak tissue SNR as a function of bone thickness for all dual-energy spectral combinations.

Looking at optimal spectral dose distribution, Fig. 3.7 shows tissue basis material SNR as a function of high-energy dose allocation r for three different bone thicknesses. Based on the peaks in the two heatmaps, bone thicknesses of 1, 4, and 6 cm were chosen. Table 3.1 lists the coordinates of the peak basis material SNR for all spectral combinations with 1 cm bone. The MV-kV curve is skewed toward allocating a greater proportion of dose to the MV spectrum, peaking at $r = 0.92, 0.83$, and 0.75 for 1, 4, and 6 cm, respectively. The kV-kV curve favors a more equal dose distribution, peaking at $r = 0.51, 0.44, 0.40$. As bone thickness increases, SNR is maximized by increasing the dose allocated to the low-energy spectrum.

To assess the effect of increasing object attenuation, Fig. 3.8 shows peak SNR as a function of bone thickness for the MV-kV and kV-kV pairs. At low bone thicknesses, kV-kV SNR is higher than MV-kV SNR for both basis materials. The tissue SNR monotonically




Figure 3.6: Heatmap of peak bone SNR as a function of bone thickness for all dual-energy spectral combinations.




Figure 3.7: With 1, 4, or 6 cm bone thickness, curve of bone SNR as a function of dose allocated to the high-energy spectrum for the detunedMV-80kV (a) and 140kV-80kV (b) spectral pairs.




Figure 3.8: Maximum SNR for tissue (solid lines, left scale) and bone (dashed lines, right scale) for the 140kV-80kV (kV-kV, square marker) and detunedMV-80kV (MV-kV, circle marker) spectral combinations as a function of bone thickness.

decreases with bone thickness, while the bone SNR reaches a maximum at a thickness of 4 cm (kV-kV) or 6 cm (MV-kV). As bone thickness increases, the difference between the two curves decreases, and at 8 cm bone, the MV-kV curves intersect the kV-kV curves and begin to yield higher SNRs. This is due to the more rapid SNR drop-off of kV-kV imaging at high bone thickness.

3.4.2 CT simulations

To quantify image quality in the simulated CT images, Fig. 3.9 shows CNR in the VMIs of the pelvis phantom with and without the metal hip replacement (titanium or steel). CNR was measured as a function of VMI energy, and horizontal lines were drawn at the fixed CNR for each dose-matched, beam-hardening corrected kV SE-CT image. These represent thresholds above which DE-CT may yield improved image quality for a given imaging task. The DE-CT peak CNRs and SE-CT CNR thresholds with and without BHC are listed in Table 3.2. In

Table 3.1: Peak coordinates ($r_{\max}$, $\text{SNR}_{\max}$) as a function of dose allocation for both basis materials and all spectral pairs with a bone thickness of 1 cm.

Spectral pair	Tissue		Bone	
	$r_{\max}$	$\text{SNR}(r_{\max})$	$r_{\max}$	$\text{SNR}(r_{\max})$
6MV-80kV	0.93	47.99	0.89	3.36
6MV-120kV	0.93	31.79	0.91	1.86
6MV-140kV	0.93	26.79	0.91	1.49
detunedMV-80kV	0.92	61.23	0.87	4.24
detunedMV-120kV	0.92	41.11	0.89	2.39
detunedMV-140kV	0.91	34.88	0.89	1.93
140kV-80kV	0.51	154.02	0.45	9.89
140kV-120kV	0.51	35.72	0.50	2.09
120kV-80kV	0.50	120.38	0.45	7.94




Figure 3.9: CNR in the detunedMV-80kV (MV-kV) and 140kV-80kV (kV-kV) VMIs as a function of energy, with horizontal lines at the CNR of each kV single-energy CT acquisition. The total dose of the dual-energy scans was equivalent to the dose of the single-energy scans (10 mGy).

Table 3.2: For each phantom, maximum CNR measured in the MV-kV and kV-kV VMIs, and single-energy CNRs with and without a beam-hardening correction (BHC). For each dual-energy acquisition, the VMI energy corresponding to the maximum CNR is indicated in parentheses. For each single-energy spectrum, the change in CNR with BHC is indicated in parentheses.

Acquisition	CNR		
	Pelvis	Pelvis with titanium	Pelvis with steel
MV-kV	5.30 (61 keV)	1.78 (239 keV)	1.35 (315 keV)
kV-kV	7.16 (93 keV)	2.82 (116 keV)	0.94 (122 keV)
80kV	4.16	0.97	0.02
80kV (BHC)	5.59 (+1.43)	1.33 (+0.36)	0.14 (+0.12)
120kV	4.79	1.47	0.46
120kV (BHC)	6.00 (+1.21)	2.42 (+0.95)	0.41 (−0.05)
140kV	4.94	1.63	0.73
140kV (BHC)	6.17 (+1.23)	2.63 (+1.00)	1.05 (+0.32)




Figure 3.10: Measured contrast (numerator of the CNR) in the detunedMV-80kV (MV-kV) and 140kV-80kV (kV-kV) VMIs as a function of energy, with horizontal lines at the value of each dose-matched kV single-energy CT acquisition.




Figure 3.11: Measured noise (denominator of the CNR) in the detuned MV-80kV (MV-kV) and 140kV-80kV (kV-kV) VMIs as a function of energy, with horizontal lines at the value of each dose-matched kV single-energy CT acquisition.




Figure 3.12: RMSE in the detuned MV-80kV (MV-kV) and 140kV-80kV (kV-kV) VMIs as a function of energy. RMSE was calculated relative to the monoenergetic ground truth XCAT.

the pelvis and pelvis with titanium, single-energy CNR improves considerably with BHC. With steel, the single-energy images are too noisy for the BHC algorithm implemented to sufficiently reduce streaking artifacts. With BHC, single-energy CT is able to outperform DE-CT in certain cases. For the pelvis and pelvis with titanium, kV-kV DE-CT yields the best CNR, and the peak MV-kV CNR is lower than the SE-CT threshold. For the more highly attenuating steel, MV-kV DE-CT yields the best CNR, and the peak kV-kV CNR is below the SE-CT threshold. Without BHC, the SE-CT threshold is surpassed by the peak CNR of both MV-kV and kV-kV acquisitions for all three phantoms.

To better understand the CNR results, Fig. 3.10 and Fig. 3.11 show the separate contrast and noise values within the overall ratio. At low monoenergies, both MV-kV and kV-kV VMIs yield the highest native contrast, as expected. This is accompanied by a high noise level, especially in the MV-kV VMIs. This is likely due to the relative up-weighting of the 80kV acquisition with only 10% dose allocation at low VMI monoenergies. At high monoenergies, DE-CT VMIs generally have less noise and lower contrast than the dose-matched SE-CT images. The relative rate of change of the two factors differs for each phantom, as indicated in the conglomerate ratios shown in Fig. 3.9.

To quantify accuracy of the simulated CT images relative to the known ground truth, Fig. 3.12 shows RMSE for each DE-CT VMI. Generally, RMSE appears to decrease with greater VMI energy. Depending on the imaging setup, this lower error may come at the cost of higher CNR (Fig. 3.9). For the pelvis without metal, the minimum RMSE is 44.9 HU for kV-kV and 51.1 HU for MV-kV DE-CT. The smallest error overall is achieved with kV-kV DE-CT, though the difference relative to MV-kV is small (+6.2 HU). With metal, the RMSE measurements are much larger, likely due to the overall higher noise and possible streaking artifacts. For the titanium case, the scale of the error is similar, with a minimum RMSE of 66.9 HU for kV-kV and 69.0 HU for MV-kV DE-CT (+2.1 HU). For the steel case, MV-kV yields the lowest error; the minimum RMSE is 222.2 HU for kV-kV and 130.8 HU

for MV-kV DE-CT (-91.4 HU). The discrepancy in accuracy is also reflected in the contrast measurements shown in Fig. 3.10. Ideally, MV-kV and kV-kV acquisitions would yield VMIs with identical contrast values that match the known ground truth for each monoenergy. The contrast curves match well for the pelvis alone but diverge with a titanium or steel implant.

The simulated CT images are shown in an appendix (Section 3.7). Each DE-CT acquisition includes the constituent single-energy images, basis material images, and example VMIs. The dose-matched, beam-hardening corrected SE-CT images are also shown.

3.5 Discussion

This work presents an evaluation of MV-kV DE-CT imaging, which could be realized using existing hardware in radiotherapy settings. We approached this task beginning with an estimation theory framework for calculating SNR along a single ray and then expanding to a simulated fan-beam CT acquisition of an anthropomorphic phantom. By comparing MV-kV DE-CT images with dose-matched single-energy kV CT images and diagnostic kV-kV DE-CT images, we were able to assess whether MV-kV imaging could be used to achieve improved image quality or reduced dose relative to current clinical standards. Even small dose reductions could potentially compound over many IGRT fractions to result in a significantly lower total dose from imaging over a patient’s course of treatment.

Our single-line integral estimation theory method requires at least two input X-ray spectra to compute a theoretical upper bound on the achievable basis material SNR [66]. We use this method to compare MV-kV DE-CT to diagnostic kV-kV DE-CT. However, in a realistic radiation therapy scenario, the available imaging option would likely be a single-energy scan (either MV or kV) rather than kV-kV DE-CT, since existing treatment systems lack dual kV X-ray sources. The availability of dual MV and kV X-ray sources is a unique practical advantage of MV-kV DE-CT. While the kV-kV single-line integral SNR calculations provide useful context for the MV-kV SNRs, our CT simulation allows us to make a more clini-

cally relevant comparison of single-energy kV CT with dual-energy MV-kV DE-CT. We also simulated kV-kV DE-CT images for context with the single-line integral SNRs.

Of the nine spectral pairs considered (six MV-kV and three kV-kV), the best basis material SNRs were found using the 140kV-80kV spectral pair. This is expected, as this spectral pair is commonly used for diagnostic DE-CT imaging, since it maximizes the energy separation between high- and low-energy spectra given the peak voltages conventionally available with modern X-ray tubes [96]. Of the MV-kV DE pairs, the best SNRs were found using the detunedMV-80kV pair and the second-best with the 6MV-80kV pair. As the detuned beam has a lower effective energy than the treatment beam, it is expected to yield better image quality due to its higher detective efficiency and native contrast. However, the 6MV treatment beam has the potential to be used for acquiring images simultaneously with radiation therapy treatment, while the detuned beam could only be used for imaging before or after treatment [105]. In this work, we chose to focus on the detuned beam. This choice provides a metric of the best achievable MV-kV CT image quality using beams currently available on therapy treatment systems, which is useful for an assessment of clinical viability.

One unique consideration of this study was the effect of dose allocation between the two spectra in a dual-energy acquisition. We found that spectral dose distribution has a considerable effect on MV-kV image quality. As kV X-ray spectra have relatively similar effective energy and detective efficiency, dose allocation may not be a typical consideration for diagnostic DE-CT. Single-ray basis material SNR peaked with approximately 90% dose allocated to the detunedMV spectrum or 50% dose to the 140kV spectrum when paired with the 80kV low-energy spectrum. The exact optimal dose allocation varied depending on the basis material and bone thickness. The asymmetric SNR-versus-dose allocation curve is a unique aspect of MV-kV DE-CT, likely owing to the greater dose deposition per photon with energy in the MeV range. In general, the optimal r was slightly higher for tissue SNR than for bone SNR. For both spectral pairs, r also decreased with increasing bone thickness.

This relation was steeper for the MV-kV DE-CT pair. These results suggest spectral dose allocation is an important consideration for clinical implementation of MV-kV CT. Though past work assessing combined MV-kV image quality often draws comparisons with single-energy acquisitions, there is lack of a consideration of dose variations [106, 109, 113–115]. A fair comparison of MV-kV DE-CT with single-energy CT should utilize dose-matched acquisitions with optimal spectral dose distributions depending on object composition. Our method provides one such way of carrying out this optimization.

For the fan-beam CT simulations, we chose the dose distribution between spectra using the single-ray approximations, with a 90:10 MV-kV distribution and 50:50 kV-kV distribution in each ray. Optimal dose allocation is more complex for CT than it is for a single ray. A CT acquisition comprises thousands of line integrals over many views and detector channels, each passing through a distinct section of the anatomy with varying attenuation. For rays including greater bone thickness, our results indicate that more dose should be allocated to the low-energy spectrum. We expect that image quality in our DE-CT simulations could be improved with more thorough dose allocation optimization. Traditional CT imaging uses bowtie filtration to reduce the dose allocated to more peripheral detector angles, producing more similar spectrum magnitudes after attenuation and, consequently, more uniform noise in each channel [129]. Tube current modulation (TCM) may additionally be implemented to modulate the dose delivered at each view angle, equalizing the noise in each projection [130]. The single-line model could be used to inform a DE-CT TCM algorithm that incorporates a method for optimizing dose allocation to each spectrum at each view. Since the optimal dose allocation as a function of bone thickness varied more quickly for the detunedMV-80kV pair than it did for the diagnostic 140kV-80kV pair, MV-kV DE-CT image quality may especially benefit from such an algorithm, and it is worth considering for future applications.

In the single-line integral model, tissue SNR monotonically decreased as a function of bone thickness. Bone SNR peaked at 4 cm for the 140kV-80kV pair and 6 cm for the detunedMV-

80kV pair. At low bone thicknesses, the SNR achieved with the kV-kV pair was higher than that with the MV-kV pair. This is more relevant for most imaging scenarios, especially at antero-posterior or postero-anterior CT view angles. However, at 8 cm bone thickness and greater, the MV-kV pair yields the highest SNRs. With the greater attenuation due to high bone thickness, the higher penetrability of the MeV photons becomes more beneficial as kV images begin to suffer from photon starvation. This effect has been utilized in other work for artifact correction around highly attenuating objects, namely metal implants [112, 115]. Our findings corroborate this effect. For this reason, we explored DE-CT basis material image quality with and without metal implants in the XCAT phantom [123].

The simulated CT images show an advantage of MV-kV DE-CT over dose-matched single-energy kV CT and kV-kV DE-CT when imaging the pelvis. As expected, when with a steel hip prosthesis. This suggests potential utility of the technique for metal artifact correction. When imaging the pelvis alone or with the less attenuating titanium implant, kV-kV DE-CT yielded the best image quality. Though both dual-energy acquisitions could be used to generate low-energy VMIs with greater native contrast, the higher noise in the MV-kV acquisitions translated into a lower CNR than the dose-matched, artifact-corrected single-energy CT images of the pelvis without metal. With either titanium or steel hip prosthesis, as expected, image quality suffered for all SE-CT acquisitions. The severity of metal artifacts was especially apparent in the constituent kV images with the more highly attenuating steel implant, and MV-kV VMIs yielded the best CNR and RMSE. For the relatively less attenuating titanium implant, the kV-kV and MV-kV DE-CT produced a higher CNR than CNRs converge above 200 keV. However, visual inspection of the 300 keV VMIs (Section 3.7) indicates that the images have very different textures; the kV-kV image (Fig. 3.15f) has better contrast but a residual streak artifact, while the MV-kV image (Fig. 3.16f) is noisier but with more uniform texture. Likewise, the dose-matched, artifact-corrected SE-CT images produce a higher CNR than the MV-kV VMIs in the pelvis alone or with titanium, but there

are some residual beam-hardening artifacts that affect the noise texture (Fig. 3.19). Our metric of image quality, CNR, is unable to capture this difference in texture, and it is not clear which image would be universally preferable. Further work might explore whether such texture differences are clinically significant.

Dual-energy acquisitions can bypass the need for a separate beam-hardening correction (BHC) by using a material decomposition algorithm in the sinogram domain. We found that BHC is necessary for single-energy CNRs to exceed the peak MV-kV CNRs. We did not apply any post-processing to the DE-CT VMIs. Looking at the MV-kV pelvis-only images (Fig. 3.14), it appears that the main contributor to VMI noise is the MV acquisition. Even with 90% of the total dose allocated to the MV spectrum, it has a lower incident flux than the kV spectrum due to the higher dose deposition per photon in the MeV energy range. DE-CT has the unique potential to mitigate the high noise in the MV acquisitions with a multi-channel denoising technique that utilizes the shared edge structure of images acquired with different spectra [131, 132]. Though not implemented in our work, such a method could be particularly useful for MV-kV DE-CT, as kV images have much lower noise and sharper edges to inform denoising in the MV images. Commercial imaging equipment may implement a proprietary denoising algorithm, as well. Given the potential for MV-kV imaging to generate low-energy VMIs with high native contrast, such methods might be explored to test whether MV noise could be sufficiently reduced to produce a higher CNR than that of single-energy CT for more general imaging scenarios. Nevertheless, these simulation results demonstrate the value of MV information for metal artifact correction and VMI generation.

Other work has similarly drawn comparisons between combined MV-kV images and single-energy kV images [106, 109]. Li et al. observed that MV-kV VMIs of an image quality phantom can yield improved CNR relative to single-energy kV images with proper selection of virtual monoenergy [109]. They found low monoenergy is preferable for low-density material inserts, and conversely, high monoenergy is preferable for high-density inserts. This

trend matches our findings of VMI CNR as a function of monoenergy in the pelvis with and without metal. Similarly, Yin et al. measured comparable or better contrast in aggregate MV-kV reconstructions compared to kV alone, depending on the material [106]. It is relevant to note that these studies utilized a partial-angle acquisition technique and did not account for dose distribution, which is distinct from our method.

We chose to model a fan-beam, high-DQE, energy-integrating detector (EID) with an efficiency near 20% in the MeV energy range (Fig. 3.1). This emulates the xenon gas detectors that have been used for MV CT on commercial tomotherapy units (Accuray Inc., Sunnyvale, CA) [116–119]. Prior research indicates that a DQE(0) of at least 20% is needed for sufficient image quality without delivering excessive dose to the patient, with non-negligible dose concerns emerging due to daily imaging of patients for modern IGRT [95, 133]. However, conventional linear accelerators typically acquire MV images using a cone-beam CT geometry and flat-panel electronic portal imaging device (EPID) with a much lower detective efficiency of 1 to 2% [134]. Thus, it is of interest to consider the applicability of our work in the context of conventional cone-beam CT. MV imaging technology is rapidly evolving, and new EPIDs are being introduced with higher detective efficiency, improved spatial resolution, and reduced quantum noise. A variety of promising techniques are being explored, including multi-layered scintillators, novel scintillating materials, and direct conversion detectors [133–137]. New EPIDs have been developed with measured DQE(0) of 6.7% [134], 8.0% [135], 9.7% [136], and 22% [133]. One application of that research demonstrated that single-energy MV CT images acquired with a four-layer EPID have superior image quality relative to MV and kV images acquired with a single-layer detector in the presence of titanium and aluminum implants [138]. Our results support these findings in the context of dual-energy MV-kV CT for steel artifact correction. Furthermore, as noted, MV-kV DE-CT has the unique potential to compensate for low MV counting efficiency with a multi-channel denoising technique. Thus, we expect our work is also relevant for cone-beam MV-kV DE-CT with either an

emerging high-DQE EPID or multi-channel denoising.

Diagnostic kV imaging technology is also advancing. Photon-counting detectors (PCDs) have recently debuted in clinical CT scanners, offering advantages such as higher spatial resolution, lower quantum noise, and dose reduction with image quality comparable to conventional EIDs [48, 51]. Spectral PCDs mounted on radiotherapy treatment systems could offer a new avenue for single-shot MV-kV DE-CT by using signal pulse height thresholds to “bin” detected photons into sinograms with different energy spectra using a single X-ray beam. Another ongoing area of research is data-driven DE-CT imaging techniques. Deep learning models have been introduced for the generation of dual-energy images from single-energy data and material decomposition [139, 140]. This work could potentially be utilized in MV-kV imaging, for example with improved material decomposition algorithms that are trained to account for the large contrast and noise differences.

This work was a simplified theoretical analysis of MV-kV DE-CT, and many limitations could be more realistically modeled in future work. Our compound Poisson noise model neglected X-ray scatter, patient motion, and electronic noise. A real CT acquisition will be affected by these complexities, and a DE-CT system using simultaneous acquisitions will also experience cross-scatter from the two beams. Though the simultaneous acquisition method introduces this additional scatter, it has the advantage of reducing motion artifacts relative to a sequential acquisition method. A Monte Carlo simulation could be additionally implemented to approximate both single-source scatter and dual-source cross-scatter, in order to better weigh the costs and benefits of each technique [141, 142]. Other work has presented new methods for scatter reduction between MV and kV sources, which could also be considered [143]. Though we did not account for electronic noise in our detector model, new photon-counting detectors are able to threshold out this noise. This could also be an avenue for future work. For image reconstruction, we implemented a standard filtered back-projection algorithm [7]. More modern iterative and deep learning methods could be

implemented, which may include more advanced noise reduction. Such algorithms likely especially benefit MV CT, since MV images tend to be noisier than kV images when dose-matched.

One potential extension of this work would be a more realistic differentiation of the two source-detector arrays. While traditional CT is acquired with a full field-of-view, some work has considered MV images that are truncated by multi-leaf collimators (MLCs) [105]. Future work could simulate or acquire MLC leakage images, especially with a simulated dose plan that opens the MLCs around the tumor. These images would use a higher incident flux, and attenuation through the MLCs would be calculated. Such images may suffer from a limited field-of-view and ring artifacts, for which corrections should be explored. On a linac, imaging systems likely utilize cone-beam CT geometry, which can further suffer from greater scatter and smaller field-of-view relative to fan-beam CT [144]. However, it is important to assess these practical constraints in relation to the standard single-energy CT images. Different detector materials and geometries should be considered for the MV and kV systems. In this work, we assumed the same detector was used for all acquisitions, with a fan shape and a fixed number of channels. This matched detector acquisition allows for material decomposition in the sinogram domain, which has the advantage of alleviating beam hardening artifacts. In current clinical settings, the kV source would likely have a larger fan angle and more detector channels. These different geometries will produce sinograms with different angular extents and other properties, likely resulting in different reconstructed image qualities and necessitating implementation of an image-domain material decomposition algorithm. This may affect the resulting image quality. Even so, the results of this initial investigation provide sufficient motivation for future studies with more faithful geometric modeling.

3.6 Conclusion

This work presents an analysis of MV-kV DE-CT imaging. We estimate that basis material SNR is maximized with 90% dose allocated to the MV spectrum. For bone thicknesses greater than 8 cm, SNR is maximized using MV-kV DE-CT. In a simulated CT scan of the pelvis with a stainless steel hip prosthesis, MV-kV VMIs can produce higher CNR and lower RMSE than diagnostic kV-kV VMIs, indicating potential utility of this technique for metal artifact correction. With pelvis alone or a less attenuating titanium prosthesis, MV-kV DE-CT can generate low-energy VMIs with higher native contrast, but dose-matched single-energy CT yields a better CNR due to lower noise. For MV-kV to outperform single-energy CT in these cases, multi-channel denoising methods, multi-view dose allocation optimization, and improved megavoltage detectors might be explored. These results demonstrate the potential utility of MV-kV DE-CT, more robustly quantify the parameters for optimal implementation, and motivate future experimental investigation into clinical applications.

3.7 Appendix: Simulated CT images

For qualitative analysis, Fig. 3.13–3.18 show a sampling of simulated SE-CT images, BMIs, and VMIs used for our analysis, providing complementary visual information for the measurements in Fig. 3.9 and 3.12. Figures 3.13 and 3.14 show the kV-kV and MV-kV DE-CT acquisitions of the pelvis without metal; Fig. 3.15 and 3.16 show the images with titanium hip replacement; and Fig. 3.17 and 3.18 show the images with stainless steel hip replacement. Figure 3.19 shows the dose-matched, beam-hardening corrected SE-CT images used for comparison.

Qualitatively, in each constituent kV SE-CT image, there is noticeable beam hardening in the pelvis and severe streaking artifacts with either titanium or stainless steel metal hip replacement. These effects are more severe for the more highly attenuating metal, stainless

steel. The detuned MV SE-CT images display less beam hardening, and although the metal hip replacements increase noise, they do not cause as severe streaking.

In the dose-matched, artifact-corrected SE-CT images used for comparison with DE-CT VMIs (Fig. 3.19), there is a considerable reduction in noise and beam hardening. In images of the pelvis alone, the artifact is almost entirely corrected. With a titanium implant, there is noticeable improvement but some residual streaking that affects noise texture. With steel, although the cupping artifact is mostly corrected, the streaking artifacts are too severe to be eliminated.

Similarly, with either metal, after basis material decomposition, the kV-kV and MV-kV BMIs both show considerable streaking. In the MV-kV case, this is likely due to artifact contamination from the 80kV SE-CT image. This contamination improves after synthesis into virtual monoenergetic images. Without metal, both the MV-kV and kV-kV VMIs have visibly good contrast and no beam hardening. With titanium, the 300 keV VMIs for both MV-kV and kV-kV acquisitions improve upon the SE-CT images; kV-kV has less noise but a moderate streaking artifact, while MV-kV is slightly noisier but devoid of streaking. With the more attenuating stainless steel, the 300 keV MV-kV VMI has greatly reduced metal artifact contamination compared to the kV-kV VMIs.




Figure 3.13: Single-energy CT images (a-b), basis material images (c-d), and VMIs (e-f) for the 140kV-80kV DE-CT acquisition of the pelvis with no hip replacement. The VMIs correct the beam hardening visible in the single-energy CT images.




Figure 3.14: Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the detunedMV-80kV DE-CT acquisition of the pelvis with no hip replacement. VMI synthesis corrects for the beam hardening in the 80kV single-energy image but results in increased noise from the detunedMV single-energy image.




Figure 3.15: Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the 140kV-80kV DE-CT acquisition with titanium hip replacement. In the VMIs, the beam hardening artifact of the single-energy images has been corrected, but there is residual streaking.




Figure 3.16: Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the detunedMV-80kV DE-CT acquisition with titanium hip replacement. The 300 keV VMI does not display streaking or beam hardening artifacts, but it shows a similar noise profile as the detunedMV single-energy image.




Figure 3.17: Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the 140kV-80kV DE-CT acquisition with stainless steel hip replacement. Both single-energy images and VMIs suffer from severe streaking artifacts.




Figure 3.18: Single-energy images (a-b), basis material images (c-d), and VMIs (e-f) for the detunedMV-80kV DE-CT acquisition with stainless steel hip replacement. The 300 keV VMI highly alleviates the streaking artifact.




Figure 3.19: The dose-matched, beam-hardening corrected SE-CT images for each phantom (column) and kV spectrum (row). Images are shown with a window width of 500 HU and level of 50 HU.

CHAPTER 4

MV-KV DUAL-ENERGY CT WITH A PHOTON-COUNTING DETECTOR

4.1 Introduction

In the previous chapter, we demonstrated the ability of MV-kV DE-CT to generate images with higher native contrast than can be achieved with dose-matched single-kV images of the pelvis [102]. The technique also shows promise for metal artifact correction due to the high penetrability of megavoltage-energy photons. However, the limitations of MV imaging can curb the benefits of MV-kV DE-CT in common imaging scenarios. Though fewer MV X-rays are attenuated by the body relative to kV X-rays, they deposit much greater dose per stopped photon. To deliver the same dose as a kV spectrum, a MV spectrum's flux must be considerably reduced, resulting in much noisier images. Furthermore, MV detective efficiency is generally much lower than kV efficiency when using currently available X-ray detectors [116, 117, 134]. When implementing MV-kV DE-CT imaging, the greater noise of the MV data can contaminate the kV data and result in VMIs with an overall lower CNR. Clinical translation of MV-kV DE-CT will likely require a technique that reduces the noise of MV images.

In this chapter, we explore photon-counting detectors (PCDs) as a potential solution to this challenge for MV-kV DE-CT. The advancing implementation of PCDs in diagnostic radiology makes this topic of particular interest for MV imaging, which faces unique challenges relative to conventional kV imaging. Yet, there has been limited investigation of PCDs beyond the diagnostic CT imaging energy range, with tube potentials generally less than 150 kV. Using CT simulations and the optimization framework described in the previous chapter, we compare the performance of MV-kV DE-CT with non-spectral PCDs versus conventional

energy-integrating detectors (EIDs) and assess the degree of potential improvement.¹

4.2 Background

Though image noise might always be decreased by increasing incident flux, this results in a higher imaging dose, which is especially undesirable in radiotherapy imaging settings. Unlike in diagnostic imaging, radiation therapy patients must be imaged repeatedly in a relatively short period of time to track their body morphology and tumor progression for precise and accurate treatment planning. The total dose of this required imaging can sum to the scale of a treatment fraction, motivating the search for accessible, high-contrast, low-dose imaging for radiotherapy [133]. Ideally, dose in the context of radiotherapy imaging would be equivalent to or lower than that in diagnostic settings to account for the patient’s repeat imaging. A recent literature review reported a diagnostic reference dose level range of 10 to 15 mGy for chest-abdominopelvic CT tumor imaging [146]. Historically, MV CT for radiotherapy has been limited by the much higher dose it requires for sufficient image quality [147].

An appealing avenue for improved MV image quality might be found on the detector side [148]. Conventional MV imaging utilizes EIDs, with recent advances predominantly focused on improving detective efficiency. Most MV imaging uses flat-panel electronic portal imaging devices (EPIDs) with efficiencies as low as 1–2% [134]. MV CT with these detectors requires a dose on the order of 300 mGy for sufficient image quality [149]. Novel EPIDs are being developed with improved efficiency in the range of 5–20% through a variety of strategies, such as layering of multiple scintillators or development of new scintillating materials [133–135, 138]. With these emerging detectors, doses of 40 mGy can yield acceptable soft-tissue contrast [150]. Modern tomotherapy systems are also capable of relatively high-efficiency MV CT imaging, with a fan-beam xenon gas detector that stops approximately 20% of MV

1. This chapter is based on work published by SPIE in [145] under a Creative Commons Attribution 4.0 International License. [DOI: 10.1117/1.JMI.11.S1.S12811]

photons [116–119]. These systems can achieve acceptable soft-tissue contrast with doses as low as 10 mGy [151].

However, EIDs have certain downsides aside from detective efficiency. The uniting characteristic of EIDs is that they weight detected photons by their incident energy and sum them into the final signal. This results in low-energy photons, which have greater soft-tissue contrast than high-energy photons, contributing less information to the final image. Though this reduces both kV and MV image contrast, it is especially unfortunate for MV imaging, in which the highest energy photons are weighted a full two orders of magnitude more than the most useful low-energy photons. A second downside of EIDs is that they suffer from electronic noise, which is independent of the unavoidable quantum noise associated with photon counting statistics. Since the flux of an MV treatment beam must be greatly decreased to deliver an acceptably low imaging dose to the patient, detected photon counts are generally much lower, and electronic noise can become a more significant component of the final image. These two effects result in both lower contrast and higher noise, doubly detrimental to contrast-to-noise ratio.

As an emerging alternative, PCDs intrinsically avoid the energy weighting issue of EIDs and can threshold out the electronic noise background [48]. While PCDs also face downsides that can degrade spatial and energy resolution, including K-escape, charge sharing, and pulse pile-up, correction techniques are being developed to account for these factors and even apply them advantageously; for example, charge sharing measurements have been used to achieve sub-micron resolution in a PCD for CT imaging [62, 79, 81]. Further, PCDs have recently debuted in diagnostic CT scanners and show great potential for improving image quality relative to conventional EIDs [51].

A key promise of PCDs is that, if desired, one can extract spectral information by defining multiple pulse-height thresholds. These “spectral PCDs” enable multi-energy imaging without the historical need for repeat acquisitions, dual source-detector arrays, multi-layer

detectors, fast-kVp switching sources, or other specialized equipment [152]. With a spectral PCD, multi-energy imaging could be implemented on a radiation therapy treatment system using a single MV spectrum. However, the imposition of spectral thresholds significantly reduces the photon counts per energy bin, exacerbating the key challenge of MV CT—the high noise due to low flux at acceptable dose levels. We address the case of single-MV imaging with a spectral PCD in an appendix (Section 4.7)

The benefits of PCDs might still be realized in the context of radiation therapy imaging by MV-kV DE-CT with “non-spectral PCDs.” These detectors essentially operate as a spectral PCD with one very low energy threshold to remove electronic noise. Even without any energy resolution, non-spectral PCDs have great potential for improving MV imaging by alleviating the high-energy photon up-weighting issue and reducing low-dose image noise—two key issues of EIDs. To gain multi-energy information with non-spectral PCDs, one requires a traditional dual-source setup, as in MV-kV DE-CT on a radiation therapy treatment system.

From a practical standpoint, the potential implementation of MV-kV DE-CT with non-spectral PCDs depends on the availability of a PCD with sufficient MV counting efficiency. The PCDs implemented in clinical and preclinical CT systems are typically made of cadmium telluride (CdTe), cadmium zinc telluride (CZT), or silicon (Si) semiconductors [49]. CdTe and CZT have a higher atomic number and therefore can be made relatively thin (mm-scale) for diagnostic imaging. These detectors are likely too thin for MV imaging. Using an estimate based on linear attenuation coefficient, a 3-mm CdTe slab will stop only 10% of 3-MeV photons. This is better than most commercial EPIDs but only half the reported 20% efficiency needed for sufficiently low-dose MV imaging [133]. Si PCDs are a more affordable option due to the broader availability of silicon wafers, but for sufficient detective efficiency, they require centimeter-scale thickness due to their lower atomic number [153]. For this reason, Si detectors have been developed using an “edge-on” approach. Perhaps surprisingly, this could be more advantageous for MV imaging: a 3-cm Si slab is expected

to attenuate 20% of 3-MeV photons, which is comparable to the high-efficiency MV EID we modeled. Unfortunately, MV photons will produce larger charge clouds and thus are likely to suffer from more charge sharing. This would be especially prominent in Si detectors, where there is a greater proportion of Compton events contributing to overall photon attenuation [153]. Though these PCDs may offer sufficient MV efficiency, they may suffer losses in spatial resolution. In an emerging application, some work has explored novel edge-on CZT detectors for positron emission tomography (PET) imaging [154–156]. These detectors are especially promising for MV imaging, as a 3-cm CdTe slab can stop 50% of 3-MeV photons. The drawback of edge-on CZT detectors relative to Si is likely manufacturing cost. These available and emerging diagnostic CT PCDs show promise for MV imaging and indicate the feasibility of MV-kV DE-CT using PCDs in place of the traditional EIDs.

4.3 Methods

We implemented three approaches for comparing MV-kV image quality achieved with non-spectral photon-counting and energy-integrating detector models for MV imaging on advanced radiotherapy treatment systems. Spectral PCDs are discussed separately in the appendix (Section 4.7).

The single-line integral toy model was used to assess a theoretical limit on achievable image quality and to optimize dose allocation between the MV and kV spectra (Section 4.3.2). Simulated MV-kV DE-CT images of an image quality phantom were used to assess material decomposition accuracy with the two detector models (Section 4.3.3). MV-kV DE-CT images of an anthropomorphic phantom were also simulated for qualitative noise analysis (Section 4.3.3). The two phantoms used for the CT simulations are shown in Fig. 4.1. We first describe the materials and parameters common to these approaches (Section 4.3.1).




Figure 4.1: The two phantoms used for the CT simulations. (a) An image quality phantom of tissue with seven bone inserts of varying densities. (b) A section of the XCAT phantom. Pixel values are in HU evaluated at 80 keV and shown at a window level of 0 and width of 1000. Each image field-of-view is 50 cm.

4.3.1 Materials

The modeled MV-kV dual-energy system utilized a detuned MV X-ray spectrum and 80-kVp spectrum typical for diagnostic imaging. This MV source is based on a typical 6 MV radiotherapy treatment beam with its spectrum softened such that the average photon energy is below 3 MeV and the maximum flux is at approximately 1 MeV [120, 121]. The detuned beam is achievable on existing treatment systems and more favorable for MV imaging by mitigating some drawbacks of high-energy X-rays, but it would not be used for radiotherapy treatment.

Both EID and PCD models used the same energy-dependent detective efficiency function $\eta(E)$. This allowed us to isolate the effect of the different photon weighting schemes of the two detector models, which is of particular interest in the high-energy case of MV imaging. The detective efficiency function was modeled based on existing xenon gas EIDs with relatively high detective quantum efficiency in the MV energy range. The function is near unity in the keV energy range and decreases to 20% at energies up to 6 MeV [102]. Such detectors are

implemented in existing tomotherapy systems (Accuray Inc., Sunnyvale, CA) for fan-beam MV imaging [118, 119]. For our theoretical comparison, we applied this same efficiency function to the PCD model, though such detectors do not exist in reality. Other high-efficiency MV detector options are also emerging, for example utilizing stacks of electronic portal imaging devices (EPIDs) as EIDs or edge-on PCDs [133, 153–156].

The raytracing CT simulation utilized a 55-cm source-to-isocenter distance and 100-cm source-to-detector distance. The fan-beam detector had 800 channels and a total fan angle of 50 deg. For all simulations, the two input spectra magnitudes were rescaled such that they would deliver a total dose of 10 mGy to the center of a 40-cm water cylinder [102]. A total of 1200 projection views were acquired over a 360° rotation. For each dual-energy acquisition, sinogram-domain material decomposition into tissue and bone basis materials was performed using a Gauss–Newton algorithm [128]. Images were reconstructed using fan-beam filtered back-projection with a matrix size of 512, field-of-view of 50 cm, and ramp filter with cutoff frequency at 80% of the Nyquist limit. VMIs were generated at a continuum of energies ($E_0 = 20\text{--}120$ keV) as a linear combination of the basis material images using Eq. 3.6.

The MV-kV DE-CT simulation workflow is summarized as a flowchart in Fig. 4.2, which includes example images demonstrating how VMIs can alleviate single-energy image artifacts.

4.3.2 *Single-line integral*

In the previous chapter, we described an estimation theory approach for computing the Cramèr–Rao lower bound (CRLB) on variance in the context of basis material decomposition along a single ray through two materials in the context of EIDs only and no electronic noise [100, 102]. This allowed us to compare MV-kV DE imaging with traditional diagnostic (kV-kV) DE imaging. In this chapter, we build upon this approach to compare two types of detectors.

The detector response function is distinct for integrating measurements [$D(E) = E$] and




Figure 4.2: Flowchart summarizing the MV-kV simulation procedure. Single-energy MV and kV sinograms are simulated independently and then recombined for DE-CT (top row): sinogram-domain material decomposition is performed, then material images are reconstructed, and VMIs are generated using Eq. 3.6. Alternatively, the single-energy sinograms can be directly reconstructed (bottom row): the single-kV image (left) suffers from severe beam hardening, while the single-MV image (right) has poor contrast and high noise.

counting measurements $[D(E) = 1]$, resulting in different noise distributions and thus different functional forms of the CRLB for each detector. The PCD follows a standard Poisson distribution, while the EID follows an energy-weighted compound Poisson distribution, approximated as a Gaussian yielding the Fisher matrix elements of Eq. 3.5. Using the same notation, the Fisher matrix elements for the Poissonian PCD measurements are

$$\mathcal{F}_{pq} = \sum_{i=1,2} \frac{1}{m_i} \frac{\partial m_i}{\partial A_p} \frac{\partial m_i}{\partial A_q}, \quad (4.1)$$

where to reiterate, m_i are the intensity measurements and A_j are the mass thicknesses of the basis materials $A_j \equiv \rho_j t_j$.

Furthermore, one major advantage of PCDs is their potential to eliminate electronic noise by applying some minimum pulse-height threshold greater than the noisy background. We assess the utility of this strategy by extending the EID CRLB model to include electronic

noise by adding a factor of $\sigma_e \times \overline{E}_i$ to the EID noise model, where σ_e is the standard deviation in counts due to the electronic noise and $\overline{E}_i$ is the intensity-weighted average energy of spectrum I_i . We chose $\sigma_e = 10$ photons. In contrast, the PCD model used $\sigma_e = 0$, simulating a low pulse-height threshold that eliminates background noise but does not yield other multi-energy information.

Interpreting the true mass thickness as the ideal signal, we defined SNR_j as the ratio of A_j to the square root of its CRLB σ_{A_j} (Eq. 3.3) as a figure of merit. We computed SNR_j for a range of tissue thicknesses ($t_1 = 10\text{--}50$ cm, $\rho_1 = 1.06$ g/cm³) and bone thicknesses ($t_2 = 0.1\text{--}10$ cm, $\rho_2 = 1.85$ g/cm³). We also considered the effect of dose allocation between the MV and kV spectra. The magnitude of each spectrum was initially scaled such that the total dose delivered by a single polychromatic ray to the center of a 40-cm water cylinder would be 1 μGy [102]. For a CT scan with 1000 views, this corresponds to a 1-mGy dose. To ensure the total dose remained fixed for each acquisition, the magnitude of the MV spectrum was rescaled by a factor r (ranging from 0.01 to 0.99), and the kV spectrum was rescaled by the remaining $1 - r$. We compared the dose-optimized SNRs achieved with either an EID or non-spectral PCD model, the dose allocation factor r needed to achieve this optimal SNR, and the effect of basis material thickness.

The resulting SNRs are scalable to other dose levels. For the PCD case, which has a pure Poisson counting noise model, noise is proportional to the inverse square root of flux, so SNR increases with the square root of dose per ray [μGy] at any dose allocation factor. Though this square root rule-of-thumb is not exact for the EID case due to photon energy weighting and electronic noise, these nonidealities are mainly relevant at very low doses, so the rule is still widely applied [157–161].

4.3.3 CT simulations

As an extension of the single-line integral model, we simulated MV-kV DE-CT acquisitions of two phantoms. To assess basis material decomposition with the two detector schemes, we created a computational image quality (IQ) phantom comprising a 40-cm tissue cylinder with seven 3-cm bone inserts with densities varying from 1.0 to 2.2 g/cm³. Twenty MV-kV DE-CT acquisitions were simulated using both EIDs and non-spectral PCDs as described above, providing several noisy realizations for statistical analysis. In the bone basis material image, the mean pixel value in each insert was measured and compared to the known ground truth. Contrast-to-noise ratio (CNR) was measured for each bone insert (Eq. 3.7), using ground-truth masks to define ROIs in the region of the bone insert at the density of interest and the uniform tissue background.

For qualitative analysis of image noise profiles using either detector scheme, we simulated MV-kV DE-CT acquisitions of an extended cardiac torso (XCAT) phantom [123]. We generated twenty noisy CT images using both detector models and synthesized the basis material images into VMIs using Eq. 3.6. Noise profile images were generated by taking the pixel-by-pixel standard deviation over the multiple noisy realizations.

For both phantoms, the total DE-CT dose was fixed at 10 mGy. The dose allocation between the MV and kV spectra were approximately optimized using the results of the single-line integral. An initial 180° noiseless simulation with just one central detector channel was conducted to yield basis material thickness line profiles through the phantom, and a single dose allocation factor r was chosen depending on the varying the material thickness pairs over the several projection view angles. The thickness line profiles for the two phantoms are shown in Figs. 4.3 and 4.4. Though this method assumed *a priori* knowledge of the phantom composition, it was highly approximate. Similar if not better methods can likely be implemented in clinical scenarios using anterior-posterior and lateral scout scans, as is commonly used for tube current modulation. Based on these thickness profiles, we




Figure 4.3: Basis material thickness line profiles through the IQ phantom used for approximate dose allocation optimization. Tissue thickness ranged between 37 and 40 cm, and bone thickness ranged from 1.5- to 3.5-cm equivalent for the seven different density inserts.

utilized 85% and 90% dose allocated to the MV spectrum for the IQ and XCAT simulations, respectively.

4.4 Results

4.4.1 Single-line integral

Figure 4.5 shows the dose allocation factor r_{opt} that optimizes basis material SNR using 40-cm tissue and a variety of bone thicknesses. Relative to the EID SNR, the PCD SNR is optimized with a slightly lower dose allocated to the MV spectrum, though the difference decreases as bone thickness increases. The optimal dose allocation is distinct for bone and tissue, with bone generally having a slightly lower r_{opt} than tissue. The largest changes in r_{opt} occur with respect to bone thickness: small bone thicknesses (<1 cm) are best with near 90% dose allocated to the MV spectrum, whereas higher bone thicknesses are optimized with a more equal 50% dose to the MV spectrum. Based on this, we infer the optimal dose allocation for a given image acquisition depends primarily on object composition, with




Figure 4.4: Basis material thickness line profiles through the XCAT phantom used for approximate dose allocation optimization. Tissue thickness ranged between 18 and 32 cm, and bone thickness varied up to 5 cm.

additional but less pronounced dependence on the material of interest (tissue or bone) and detector (EID or PCD).

Figures 4.6 and 4.7 show heatmaps of dose-optimized basis material SNR at all bone and tissue thickness pairs assessed. Across all material compositions, PCDs produce between 15 and 45% improvement in SNR. The greatest percent improvements are seen for very thin objects, which also have the lowest overall SNRs due to their lower attenuation. The optimal dose factors that generated the maximum SNRs shown in these heatmaps were used to inform the approximate dose allocation optimization of the CT scans. For example, Fig. 4.5 shows r_{opt} as a function of bone thickness for 40-cm tissue. For each phantom, characteristic average tissue thickness and non-zero bone thickness were approximated from the the angle-dependent thickness profiles (Figs. 4.3 and 4.4). Average non-zero bone thickness was generally low (≈ 2 cm), where the $r_{\text{opt}}(t_{\text{bone}})$ curve is slowly varying (Fig. 4.5) and only 2–3% different for EIDs and PCDs. From this we selected the constant values of $r_{\text{opt}} = 0.85$ and 0.90 for the IQ phantom and XCAT phantom, respectively.




Figure 4.5: The dose allocation to the MV spectrum maximizing basis material SNR for (a) tissue and (b) bone as a function of bone thickness with fixed tissue thickness of 40 cm.




Figure 4.6: Dose-optimized tissue SNR using (a) photon-counting and (b) energy-integrating detector models for a variety of basis material thicknesses. Panel (c) shows the percent difference.




Figure 4.7: Dose-optimized bone SNR using (a) photon-counting and (b) energy-integrating detector models for a variety of basis material thicknesses. Panel (c) shows the percent difference.

4.4.2 CT simulations

To assess basis material decomposition accuracy for the two detector schemes, Fig. 4.8 shows bone basis material measurements in the seven IQ phantom inserts. The measurements are similar using either PCD or EID at all densities, and both detector schemes result in a slight underestimation of true bone density. However, the basis material images of PCDs and EIDs produce distinct noise levels in VMIs, as illustrated in Fig. 4.9, which shows CNR in three bone density inserts at a continuum of monoenergies. PCDs consistently yield higher VMI CNR than EIDs at all monoenergies, with peak CNR generally occurring at a slightly lower monoenergy. The improvement in CNR with PCDs is especially pronounced at lower (<50 keV) VMI energies.

Figures 4.10–4.12 show VMIs and their corresponding noise profiles using either the PCD or EID model. Broadly, EIDs result in noisier images. In the lower energy VMIs (50 keV), the noise profile is concentrated around bony regions and in streaks between bones. While this is apparent for both detector models, it is more pronounced with EIDs. This results in a less isotropic noise texture in the EID images. As VMI energy increases to 60 and 70 keV (Figs. 4.11–4.12), the noise texture shifts to be more isotropic across the tissue regions of the phantom. While the EID images still have greater noise magnitude than the PCD




Figure 4.8: Measured versus true bone density in the seven inserts of the image quality phantom. Error bars represent $k = 1$ standard deviation over all noisy simulations.

images, the two noise textures appear to be more similar.

4.5 Discussion

MV-kV DE-CT imaging is a proposed modality that could be implemented on existing radiotherapy treatment systems, possibly bringing the advantages of multi-energy imaging to routine radiotherapy imaging over the course of a patient’s cancer treatment. In this work, we explored the implications of exchanging the EIDs in a previously studied MV-kV DE-CT system with non-spectral PCDs of identical counting efficiency [102]. This allowed us to assess the effect of the energy-weighting detection scheme and electronic noise. We expected the energy-weighting effect of EIDs is particularly detrimental for MV imaging due to the multiple order-of-magnitude increase in weighting for photons carrying the least contrast information. A 3 MeV photon would be weighted $100\times$ greater than a 30 keV photon. Furthermore, since the MV spectrum flux must be greatly reduced relative to the kV spectrum to reach acceptable dose levels, we expected the effect of a constant electronic noise background may be more pronounced. Our findings corroborate these expectations.




Figure 4.9: CNR measured in three bone inserts with density (a) 1.4 g/cm^3 , (b) 1.8 g/cm^3 , or (c) 2.2 g/cm^3 as a function of VMI energy. The error bars represent $k = 2$ standard deviations over all noisy simulations.




Figure 4.10: (a)–(b) 50-keV VMIs and (c)–(d) their noise profiles simulated with either PCD (left) or EID (right) schemes. Pixel units are HU. The noise profile is concentrated around the bone regions.




Figure 4.11: (a)–(b) 60-keV VMIs and (c)–(d) their noise profiles simulated with either PCD (left) or EID (right) schemes. Pixel units are HU. The noise profile is distributed around the bones and lateral edges of the phantom.




Figure 4.12: (a)–(b) 70-keV VMIs and (c)–(d) their noise profiles simulated with either PCD (left) or EID (right) schemes. Pixel units are HU. The noise profile has a more uniform texture across tissue regions.

In the single-line integral model, for each basis material, dose-optimized SNR initially increases with thickness as the true mass thickness $\rho \times t$ (numerator) increases. Each SNR eventually peaks due to increases in the CRLB (denominator), likely due to reduction in detected photon counts as they are attenuated by the object. This general relationship is true for the two detector models. Comparing PCD SNRs to EID SNRs, we observe a saddle surface shape in the percent improvement heatmap for all material thickness pairs (panel (c) of Fig. 4.6–4.7). The greatest improvement (approx. 45%) in SNR was seen at low thicknesses, where the signal is the weakest but the electronic noise background is less significant relative to the large number of detected photons. PCDs yield a smaller but still observable (approx. 15%) improvement at more moderate thickness pairs. At very high thicknesses, the percent improvement increases to near 30%, likely due to the high attenuation of the incident beams resulting in a more dominant electronic noise component in the EID signal.

The dose optimization factor r_{opt} was slightly different depending on the detector model, and 0–5% less of the total dose should be allocated to the MV spectrum when using a PCD. Tables of r_{opt} values for all material thickness pairs were used to choose the dose allocations of 0.85 and 0.90 for the IQ phantom and XCAT phantom CT simulations, respectively. More specifically, r_{opt} varies depending on the specific thickness pair of each view angle as in Figs. 4.3 and 4.4. In a separate experiment, we simulated CT scans with r_{opt} modulated as a function of view angle according to the thickness profiles for each phantom and detector type. We found that CNR did not significantly improve relative to the CNRs for the constant values of 0.85 and 0.90, indicating that these constant r_{opt} choices are sufficient for this study. The scope of this paper focuses on fairly comparing EID and PCD performance, not necessarily optimizing performance as a function of dose allocation. This subject might be further investigated in future work. Since such “ r_{opt} modulation” would require additional software for MV-kV DE-CT on existing radiotherapy machines, and it did not largely affect

the results,

In the MV-kV DE-CT simulations, the Gauss–Newton material decomposition method implemented had similar accuracy for either EID or PCD detector model [128]. By combining the two basis material images into a VMI, greater CNR can be achieved at a range of monoenergies. As expected, CNR increases with bone density due to the greater attenuation of the signal relative to the background, and the monoenergy maximizing CNR increases with bone density due to beam hardening. For each individual bone density, PCD CNR was greater than EID CNR at all VMI energies examined (20–120 keV). The peak PCD CNR occurred at a slightly lower monoenergy than the peak EID CNR. This might be attributed to the improvement in PCD basis material image noise. In principle, the two material images’ noise is correlated such that when they are recombined using Eq. 3.6, the resulting VMI will have less noise. This depends on the energy. Since the mass attenuation coefficients in the VMI calculation decrease as energy increases, low-energy VMIs can amplify the noise in the material images. This effect is not as severe for the less noisy PCD images, so the peak CNR is at a lower monoenergy with higher native contrast. The greatest CNR improvements were observed at low monoenergies (<50 keV), while only slight improvement was seen after peak PCD CNR. For a more clinically relevant assessment, we also simulated MV-kV VMIs of an anthropomorphic phantom. Qualitatively, for both detectors, lower monoenergies produced higher soft-tissue contrast but greater noise, especially around regions with bone and along the lateral patient axis (Fig. 4.10). The noise texture and magnitude is more prominent in the EID images. For both detector models, soft-tissue contrast decreased and noise texture became more uniform at higher monoenergies, but the noise magnitude was still higher using the EID (Fig. 4.11–4.12). The greatest improvements in image quality with the PCD appear to be due to improvements in noise magnitude and texture, while contrast is subjectively more similar.

The results of this computational assessment of non-spectral PCDs for MV-kV DE-CT

motivate future experimental work. The CT simulation study provides an initial estimate of the noise reduction benefits achievable with PCDs. An experimental study would more realistically include other factors such as photon scatter, spatial resolution, and available PCDs. These results could be compared to simulated findings for a better understanding of the clinical translatability of our results. Simulations and estimation theory calculations allow us to conveniently explore a range of parameters relevant for radiotherapy imaging scenarios, complementing these experimental investigations. These computational methods could also be used to optimize image acquisition parameters to ensure fair comparison of different detector models.

4.6 Conclusion

PCDs show promise for improving MV-kV dual-energy CT imaging by improving MV image noise characteristics. In terms of single-line integral basis material SNR, we found that non-spectral PCDs yielded 15–45% improvement over EIDs. In CT simulations, material decomposition accuracy was similar for both detectors, while non-spectral PCDs resulted in higher CNR and more uniform noise texture in VMIs. Relative to existing high-efficiency MV EIDs, similar or better MV detective efficiency might be achieved by emerging diagnostic PCDs if utilizing an edge-on setup. This computational study motivates experimental assessment of PCDs for MV imaging and potential clinical translation of PCDs for MV-kV DE-CT imaging on radiotherapy systems.

4.7 Appendix: Spectral detectors with a single-MV spectrum

An additional area of interest is spectral imaging with a single MV X-ray source. While the main text focuses on non-spectral detectors, PCDs can also possess the spectral capability to resolve the energy of incident photons by using a detection scheme with more than one

pulse-height threshold, enabling the possibility of material decomposition without need for multiple polychromatic spectra. This subject has traditionally been explored in the context of diagnostic kV spectra. If feasible with MV spectra, the technique might be a great advantage for radiotherapy imaging on systems lacking a kV source entirely. Radiotherapy systems must have some on-board imaging device for patient positioning verification at the time of treatment. Since MV localizers can be sufficient for this purpose and there are practical costs associated with an additional kV X-ray source, not all systems are equipped with dual MV-kV sources [108]. Generally, kV images are more desirable for their higher contrast and lower noise relative to MV images acquired with the same dose. Energy-discriminating PCDs could, in theory, be used with a MV beam for on-board, single-shot spectral imaging and subsequent reconstruction of higher contrast images without need for a kV X-ray source.

This might offer advantages over single-MV imaging. Though the imposition of spectral thresholds would reduce the photon counts of each detected energy bin, this might be alleviated relative to a single-MV acquisition by recombining spectral data in VMIs. Like single-MV imaging, the utility of spectral MV imaging will likely depend on having sufficiently high detective efficiency. It might also benefit from more extensive treatment beam detuning, so that flux could be increased without higher dose and that more low-energy information would be available to complement high-energy information for material decomposition. However, in comparison to MV-kV DE imaging, dose-matched spectral MV imaging will always have a much lower total flux due to the greater dose deposition per attenuated MV photon. As concluded in the main text, MV-kV DE imaging is optimized with most incident flux allocated to the MV beam, depending on object composition. Energy bins in spectral MV imaging could not be dose-optimized due to the fixed shape of the X-ray spectrum. For these reasons, in most scenarios, we do not expect spectral MV imaging to outperform MV-kV DE imaging. In this section, we quantitatively investigate MV imaging with a spectral PCD.

We modeled a spectral PCD with one, idealized energy threshold to sort detected photons into two energy bins. We applied the single-line integral model to probe achievable basis-material SNR as a function of this threshold. The same detective efficiency model as described in Section 4.3.1 was implemented. The object was modeled as 37-cm of tissue and 3-cm of bone, as in the IQ phantom used for the dual-energy simulations. We considered either a megavoltage spectrum (detunedMV) or a standard diagnostic kilovoltage spectrum (120kV) applied independently in combination with the spectral PCD. Figure 4.13 shows the shape and magnitude of the spectra, which were scaled to deliver the same 1- μ Gy dose to the center of a 40-cm water cylinder. As kV photons generally deposit less dose than a MV photon, the magnitude of counts in the kV spectrum is much higher than it is in the MV spectrum.

Figure 4.14 shows the estimated basis material SNR as a function of spectral threshold for the kV and MV acquisitions. The location of the peak of each curve indicates the optimal spectral threshold for imaging this material pair. The peak SNRs achieved with the kV spectrum are far higher than those achieved with the MV spectrum, which reaches a maximum SNR of approximately 1.2 for tissue. For both basis materials, the 120kV SNR is maximized with an ideal spectral threshold of 70 keV. The maximum SNRs are 153 for bone and 31 for tissue, which are greater than the non-spectral, dose-optimized MV-kV SNRs shown in Fig. 4.6 and 4.7 (67 and 14, respectively). This likely owes to the large difference in total counts between the kV and MV spectra, the order-of-magnitude difference in the modeled detective efficiency function at kV and MV energy ranges, and lower native contrast at higher X-ray energies. The difference may also be attributed to the material thicknesses that were assessed (37 cm of tissue and 3 cm of bone). At greater object thicknesses, spectral MV or non-spectral MV-kV systems have greater potential to outperform single-kV systems [102]. Furthermore, this initial simple experiment utilized a highly idealized spectral response, granting a large advantage to the kV spectral PCD relative to the MV-




Figure 4.13: The (a) kilovoltage and (b) megavoltage X-ray spectra with photon counts per 1-keV energy bin scaled to deliver approximately 1- μ Gy dose to the center of a 40-cm water cylinder.

kV non-spectral PCD single-line SNRs. In generating the spectral signals, we assumed the electrical signal corresponding to each stopped photon is perfectly proportional to its energy. In reality, the amplitude recorded by the detector is distorted by K-fluorescence, Compton scatter, and charge sharing [49]. These non-idealities are material-specific, with CdTe- and CZT-based PCDs more affected by K-fluorescence and Si-based PCDs measuring more low-energy Compton interactions. A more realistic spectral response might be modeled by a Gaussian photopeak centered at the incident photon energy with a low-energy tail [82, 162].

The primary outcome of this experiment is motivation for our focus on MV-kV dual-energy imaging with a non-spectral PCD rather than MV-only imaging with a spectral detector. In future work, it may be of interest to explore a dual-energy MV-kV system that uses a spectral PCD, which could utilize a more advanced basis material decomposition approach to benefit from more than two measurements. Work focusing on this subject should take care to implement a realistic PCD spectral response. We expect that megavoltage-only imaging with a spectral PCD will require much higher dose or greater detective efficiency to be of clinical utility. With a non-spectral PCD, a combination of megavoltage and kilovoltage




Figure 4.14: Single-line integral basis material SNR measured as a function of spectral threshold imposed on the (a) kilovoltage and (b) megavoltage spectra for an object composed of 37-cm tissue and 3-cm bone.

spectra is likely a more feasible approach to material decomposition on a radiotherapy system.

CHAPTER 5

PHASE-CONTRAST IMAGING FUNDAMENTALS

In the previous three chapters, we only considered the principle of X-ray absorption, adopting a ray-driven perspective and taking for granted that μ is the per-distance attenuation probability of a photon. This is the conventional presentation of X-ray physics in medical imaging contexts, and it is used for both forward model simulations and inverse problem solutions. However, like any other photon, X-rays have both particle and wave behavior; more holistically, it is essential to consider both absorption and refraction [23]. As it turns out, in the energy range and spatial resolution range typical of modern diagnostic imaging technology, it is not unreasonable that the traditional view has neglected X-ray refraction. The high energy of X-rays results in smaller refraction angles relative to the spatial resolution of most clinical imaging systems, rendering phase shifts immeasurable with conventional equipment. With advancing source and detector technology, these potential benefits may be within reach.

In this section, we lay out the necessary phase-contrast imaging theory, beginning with a general X-ray wave model, then describing possible paths of approximations to simplify the model for practical implementation.

5.1 Wave model

As X-rays are simply electromagnetic radiation, we can set up our framework starting with the same Maxwell equations applied in the visible light regime or any other energy range. Certain simplifications will be made that are appropriate for our specific imaging regime; these will be clearly noted.

In the absorption-only framework, attenuation was modeled along a single-line integral described by the Beer–Lambert law (Eq. 2.3). Now consider a full 3D, time-varying, complex

X-ray wave field $\Psi(\mathbf{r}, t)$, where $\mathbf{r} = (x, y, z)$ is a point in space and t is a moment in time. As a complex value, Ψ has both magnitude and phase, so it may be written as

$$\Psi(\mathbf{r}, t) = \sqrt{I(\mathbf{r}, t)} e^{i\phi(\mathbf{r}, t)}, \quad (5.1)$$

where $I = |\Psi|^2$ is the intensity and ϕ is the phase. If the field is coherent and monochromatic with energy E , then its time and spatial dependence are separable [23],

$$\Psi(\mathbf{r}, E, t) = \psi(\mathbf{r}, E) e^{-i\omega(E)t}, \quad (5.2)$$

where $\omega(E) = E/\hbar$ is the angular frequency.

We also assume the wave is paraxial, so it can be described as a series of x - y planes advancing perpendicularly to the axis of propagation z . Without loss of generality, it is also useful to write the spatial field as the product of an unscattered component and a perturbed envelope,

$$\psi(\mathbf{r}, E) = \tilde{\psi}(\mathbf{r}, E) e^{ik(E)z}. \quad (5.3)$$

The intensity of the unscattered component $|e^{ik(E)z}|^2 = 1$, so the full complex wave $\psi(\mathbf{r}, E)$ and the envelope $\tilde{\psi}(\mathbf{r}, E)$ always have the same intensity. This is useful for imaging applications, where X-ray detectors are only capable of measuring intensity rather than the full complex wave field, so it can be convenient to focus just on the envelope.

We consider two cases of wave evolution along the z axis: propagation through free space, and modulation through perturbing media.

5.2 Free-space propagation

The scalar vacuum wave equation dictates how the wave evolves through space and time,

$$\left(\frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \nabla^2\right) \Psi(\mathbf{r}, t) = 0, \quad (5.4)$$

where c is the speed of light in a vacuum and ∇^2 is the Laplacian operator,

$$\nabla^2 \equiv \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}. \quad (5.5)$$

Because the wave equation is linear, a general polychromatic X-ray wave field can be decomposed into a sum over its monochromatic components [23],

$$\Psi(\mathbf{r}, t) = \int dE \psi(\mathbf{r}, E) e^{-i\omega(E)t}, \quad (5.6)$$

each of which evolves independently according to the corresponding monochromatic wave equation.¹ So, while the following derivations are written for a single monochromatic component with energy E , the same principles are applicable to polychromatic radiation by substitution into Eq. 5.6.

Since the monochromatic field has separable space and time components, substituting Eq. 5.2 into Eq. 5.4 yields the Helmholtz equation,

$$\left(\nabla^2 + k^2(E)\right) \psi(\mathbf{r}, E) = 0, \quad (5.7)$$

which is essentially a time-independent wave equation. The separation reflects the fact that the wave's time development at any point is dictated by oscillations with fixed frequency $\omega(E) = ck(E)$, where $k(E) = E/\hbar c$ is the wave number. By direct substitution, it is apparent

1. This linearity does not model inelastic scatter, since it assumes that different energy components are independent [23].

that the elementary plane waves are solutions to the Helmholtz equation,

$$\psi^{(\text{PW})}(\mathbf{r}, E) = \exp [i(k_x x + k_y y + k_z z)] , \quad (5.8)$$

since

$$\nabla^2 \psi^{(\text{PW})}(\mathbf{r}, E) = -k^2(E) \psi^{(\text{PW})}(\mathbf{r}, E), \quad (5.9)$$

and we have expanded the components of the wavenumber,

$$k^2 = k_x^2 + k_y^2 + k_z^2. \quad (5.10)$$

For forward propagation in the z direction, we take the positive root²

$$k_z = \sqrt{k^2 - k_x^2 - k_y^2}. \quad (5.11)$$

Substituting this relation into Eq. 5.8, the plane wave can be written to isolate its z -dependence as

$$\psi^{(\text{PW})}(\mathbf{r}, E) = \exp [i(k_x x + k_y y)] \exp \left[i z \sqrt{k^2(E) - k_x^2 - k_y^2} \right]. \quad (5.12)$$

When $z = 0$, this gives the expression for the unpropagated plane wave,

$$\psi^{(\text{PW})}(\mathbf{r}_\perp, z = 0, E) = \exp [i(k_x x + k_y y)] , \quad (5.13)$$

where $\mathbf{r}_\perp = (x, y)$ denotes the coordinates in the transverse plane. Then, Eq. 5.12 can be

2. Note that this assumes k_z is real, or that $k_x^2 + k_y^2 \leq k^2$, which equivalently neglects evanescent waves [23].

written as the product of a propagation factor and the initial plane wave,

$$\psi^{(\text{PW})}(\mathbf{r}_\perp, z = R, E) = \psi^{(\text{PW})}(\mathbf{r}_\perp, z = 0, E) \exp \left[iR \sqrt{k^2(E) - k_x^2 - k_y^2} \right], \quad (5.14)$$

where R is the propagation distance.

To calculate the full propagated wave, we sum the contributions of the k_x - k_y plane wave components. In other words, we first decompose the field in the plane $z = 0$ into transverse plane waves,

$$\psi(\mathbf{r}_\perp, z = 0, E) = \iint dk_x dk_y \hat{\psi}(\mathbf{k}_\perp, z = 0, E) \exp [i(k_x x + k_y y)], \quad (5.15)$$

where $\mathbf{k}_\perp = (k_x, k_y)$ denotes the transverse plane frequency coordinates, $\hat{\psi}(\mathbf{k}_\perp, z = 0, E)$ is the 2D Fourier transform (FT) of $\psi(\mathbf{r}_\perp, z = 0, E)$ with respect to x and y , also called the angular spectrum. Each component propagates independently according to Eq. 5.14, so the full expression for the downstream wave is

$$\begin{aligned} \psi(\mathbf{r}_\perp, z = R, E) = \\ \iint dk_x dk_y \hat{\psi}(\mathbf{k}_\perp, z = 0, E) \exp \left[iR \sqrt{k^2(E) - k_x^2 - k_y^2} \right] \exp [i(k_x x + k_y y)]. \end{aligned} \quad (5.16)$$

More compactly, we can write this using an operator formalism,

$$\psi(\mathbf{r}_\perp, z = R, E) = \mathcal{D}_R \psi(\mathbf{r}_\perp, z = 0, E), \quad (5.17)$$

where $\mathcal{D}_R$ is the free-space propagation operator

$$\mathcal{D}_R = \mathcal{F}^{-1} \exp \left[iR \sqrt{k^2(E) - k_x^2 - k_y^2} \right] \mathcal{F}, \quad (5.18)$$

which uses cascaded operator notation acting from right to left. The FT is always taken

with respect to x and y (or their corresponding frequency components k_x, k_y for the inverse FT).

This point provides a general forward model for X-ray wave propagation in free space. There are a few possible options for simplifying further. Models implementing more assumptions and approximations typically lend themselves to a more convenient inverse problem solution at the cost of quantitative accuracy when their conditions are violated. Here, we focus on two common implementations: the Fresnel diffraction operator and the transport-of-intensity equation (TIE).

5.2.1 Fresnel diffraction

Assuming our wave is paraxial, the transverse spatial frequencies are small compared to the wavenumber, i.e. $k_x^2 + k_y^2 \ll k^2(E)$. Then,

$$\begin{aligned} \sqrt{k^2(E) - k_x^2 - k_y^2} &= k(E) \sqrt{1 - \frac{k_x^2 + k_y^2}{k^2(E)}} \\ &\approx k(E) - \frac{k_x^2 + k_y^2}{2k(E)}, \end{aligned} \quad (5.19)$$

after a binomial approximation. Substituting Eq. 5.19 into Eq. 5.16 gives the Fresnel approximation to the free-space propagator,

$$\psi(\mathbf{r}_\perp, z = R, E) \approx e^{ik(E)R} \mathcal{F}^{-1} \left\{ \exp \left[-\frac{iR}{2k(E)} (k_x^2 + k_y^2) \right] \mathcal{F} \{ \psi(\mathbf{r}_\perp, z = 0, E) \} \right\}. \quad (5.20)$$

Recall Eq. 5.3, in which we separated the wave field into an envelope $\tilde{\psi}(\mathbf{r}_\perp, z, E)$ and an unscattered factor $e^{ik(E)z}$, as we see in the leading term of Eq. 5.20 here. This factor does not affect the intensity of the wave field. Turning our attention to the perturbed envelope,

we can again characterize its evolution through free space using operator notation,

$$\tilde{\psi}(\mathbf{r}_\perp, z = R, E) = \mathcal{D}_R^{(\text{F})} \tilde{\psi}(\mathbf{r}_\perp, z = 0, E) \quad (5.21)$$

where $\mathcal{D}_R^{(\text{F})}$ is the Fresnel free-space propagation operator,

$$\mathcal{D}_R^{(\text{F})} = \mathcal{F}^{-1} H_R(\mathbf{k}_\perp, E) \mathcal{F}, \quad (5.22)$$

$$H_R(\mathbf{k}_\perp, E) = \exp \left[-\frac{iR}{2k(E)} (k_x^2 + k_y^2) \right]. \quad (5.23)$$

Note that by the convolution-multiplication theorem, this Fourier-space multiplication is also equivalent to spatial convolution of the field at $z = 0$ with the kernel $h_R(\mathbf{r}_\perp, E) \equiv \mathcal{F}^{-1} \{H_R(\mathbf{k}_\perp, E)\}$ [71].

5.2.2 Transport-of-intensity equation (TIE)

If we substitute our envelope expression for the wave (Eq. 5.3) into the Helmholtz equation (Eq. 5.7), apply the paraxial approximation, and make some lengthy simplifications,³ we arrive at the free-space paraxial equation,

$$\left(2ik(E) \frac{\partial}{\partial z} + \nabla_\perp^2 \right) \tilde{\psi}(\mathbf{r}, E) = 0, \quad (5.24)$$

where $\nabla_\perp^2$ is the transverse Laplacian,

$$\nabla_\perp^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}. \quad (5.25)$$

3. A more detailed description can be found in [23].

Further substituting in the original expression for Ψ (Eq. 5.1) and expanding yields the general TIE,

$$-\nabla_{\perp} \cdot [I(\mathbf{r}_{\perp}, z) \nabla_{\perp} \phi(\mathbf{r}_{\perp}, z)] = k(E) \frac{\partial I(\mathbf{r}_{\perp}, z)}{\partial z}, \quad (5.26)$$

with the full derivation given in an appendix (section 5.7) [24]. The operators on the left-hand side dictate how the wave components evolve in the transverse plane ($\nabla_{\perp}$), while the right-hand side dictates how the intensity changes along the optical axis ($\partial/\partial z$), providing some useful intuition in separating out the dependencies.

To further simplify, one may apply a finite difference approximation to the intensity planes along the propagation axis, i.e.,

$$\frac{\partial I(\mathbf{r}_{\perp}, z)}{\partial z} \approx \frac{I(\mathbf{r}_{\perp}, z + \Delta z) - I(\mathbf{r}_{\perp}, z)}{\Delta z}. \quad (5.27)$$

Assuming small propagation distance $R = \Delta z$, the intensity portion of the TIE under this finite difference approximation becomes

$$I(\mathbf{r}_{\perp}, z = R) = I(\mathbf{r}_{\perp}, z = 0) - \frac{R}{k(E)} \nabla_{\perp} \cdot [I(\mathbf{r}_{\perp}, z = 0) \nabla_{\perp} \phi(\mathbf{r}_{\perp}, z = 0)]. \quad (5.28)$$

Expanding the last term,

$$\begin{aligned} & \nabla_{\perp} \cdot [I(\mathbf{r}_{\perp}, z = 0) \nabla_{\perp} \phi(\mathbf{r}_{\perp}, z = 0)] \\ &= I(\mathbf{r}_{\perp}, z = 0) \nabla_{\perp}^2 \phi(\mathbf{r}_{\perp}) + \nabla_{\perp} I(\mathbf{r}_{\perp}, z = 0) \cdot \nabla_{\perp} \phi(\mathbf{r}_{\perp}, z = 0). \end{aligned} \quad (5.29)$$

Applying the assumption that $I(\mathbf{r}_{\perp}, z = 0)$ has weak transverse intensity gradients, i.e. $\nabla_{\perp} I(\mathbf{r}_{\perp}, z = 0) \approx 0$, then only the Laplacian of the phase term remains and the finite-

difference-approximated TIE simplifies to:

$$I(\mathbf{r}_\perp, z = R) \approx I(\mathbf{r}_\perp, z = 0) \left(1 - \frac{R}{k(E)} \nabla_\perp^2 \phi(\mathbf{r}_\perp, z = 0) \right). \quad (5.30)$$

5.3 Modulation through an object

Now with the understanding of how X-ray wave fields evolve through free space, let us consider the case of a perturbing object. Assume the object is exclusively composed of materials that are static, non-magnetic, zero charge and current density, and elastically scattering [23]. The object can be characterized by its spatial dependence of its energy-dependent, complex refractive index,⁴

$$n(\mathbf{r}, E) = 1 - \delta(\mathbf{r}, E) + i\beta(\mathbf{r}, E). \quad (5.31)$$

The wavenumber in media is simply the vacuum wavenumber scaled by the refractive index [24], allowing us to revise the vacuum Helmholtz equation to account for the object,

$$\left(\nabla^2 + k^2(E) n^2(\mathbf{r}, E) \right) \psi(\mathbf{r}, E) = 0, \quad (5.32)$$

giving the inhomogeneous Helmholtz equation.

We can then substitute in our expression for ψ in terms of its envelope $\tilde{\psi}$ (Eq. 5.3) and make use of the identity for the Laplacian of a product,

$$\nabla^2 \left[\tilde{\psi}(\mathbf{r}, E) e^{ik(E)z} \right] = \tilde{\psi}(\mathbf{r}, E) \nabla^2 e^{ik(E)z} + e^{ik(E)z} \nabla^2 \tilde{\psi}(\mathbf{r}, E) + 2 \nabla \tilde{\psi}(\mathbf{r}, E) \nabla e^{ik(E)z}. \quad (5.33)$$

If we factor out the term $e^{ik(E)z}$, substitute the expression for the Laplacian in terms of the

4. Some texts adopt the sign convention $n = 1 - \delta - i\beta$ [4] while others use $n = 1 - \delta + i\beta$ [23] (as we do here). The choice is arbitrary as long as the wave equation is still satisfied and subsequent equations are internally consistent.

transverse Laplacian,

$$\nabla^2 = \nabla_{\perp}^2 + \frac{\partial^2}{\partial z^2}, \quad (5.34)$$

and make the paraxial approximation, we arrive at the inhomogeneous paraxial equation

$$\left(2ik(E) \frac{\partial}{\partial z} + \nabla_{\perp}^2 + k^2(E) \left(n^2(\mathbf{r}, E) - 1 \right) \right) \tilde{\psi}(\mathbf{r}, E) = 0. \quad (5.35)$$

This form of the paraxial approximation neglects the second derivative of the wave field in the z direction, which holds when the wave field varies much more strongly in the transverse plane than it does along the propagation axis.

Say we have a known wave incident on an object at $z = 0$. To model the wave's modulation through the object, the task is then to use Eq. 5.35 to solve for the wave exiting the object at $z = z_0$ in terms of the incident wave. This generally involves further approximations.

5.3.1 Projection approximation

If we assume that $\nabla_{\perp}^2 \psi(\mathbf{r})$ is negligible within the object and rearrange Eq. 5.35 to isolate the z -derivative and substitute $\mathbf{r} = (\mathbf{r}_{\perp}, z)$, we arrive at a partial differential equation,

$$\frac{\partial}{\partial z} \tilde{\psi}(\mathbf{r}_{\perp}, z, E) \approx \frac{k}{2i} \left(1 - n^2(\mathbf{r}, E) \right) \tilde{\psi}(\mathbf{r}_{\perp}, z, E), \quad (5.36)$$

which has known solution [23],

$$\tilde{\psi}(\mathbf{r}_{\perp}, z = z_0, E) \approx \exp \left[\frac{k}{2i} \int_0^{z_0} dz \left(1 - n^2(\mathbf{r}_{\perp}, z, E) \right) \right] \tilde{\psi}(\mathbf{r}_{\perp}, z = 0, E). \quad (5.37)$$

This is known as the *projection approximation* (PA), also called the thin-sample approximation or geometric optics approximation [163]. Practically, this amounts to assuming that

internal propagation-based phase effects are negligible, which is generally true when the object is thin. A rough metric of PA validity is the Fresnel number,

$$N_F = \frac{\Delta^2}{\lambda(E)z}, \quad (5.38)$$

where Δ is the smallest feature size of interest and z is the distance under consideration [4]. The condition $N_F \gg 1$ indicates that refraction at the spatial resolution Δ is roughly negligible relative to the length scale z . For evaluating PA specifically, one could substitute the detector pixel size and object thickness as proxies for these parameters, respectively [164].

Furthermore, recall Eq. 5.31's expression for $n(\mathbf{r}, E)$ in terms of refractive and absorptive components $\delta(\mathbf{r}, E)$, $\beta(\mathbf{r}, E)$. In the X-ray energy range, these are much smaller than 1, roughly on the order of 10^{-6} and 10^{-9} , respectively. So, we can make the following approximation,

$$\begin{aligned} 1 - n^2(\mathbf{r}, E) &= 1 - (1 - \delta(\mathbf{r}, E) + i\beta(\mathbf{r}, E))^2 \\ &\approx 2(\delta(\mathbf{r}, E) - i\beta(\mathbf{r}, E)) \end{aligned} \quad (5.39)$$

which drops higher order δ and β terms in the expansion of n^2 , roughly 10^{-12} to 10^{-18} for X-ray energies. This provides the final expression for the object-modulated X-ray wave field according to PA,

$$\tilde{\psi}(\mathbf{r}_\perp, z = z_0, E) \approx \exp \left[-ik(E) \int_0^{z_0} dz (\delta(\mathbf{r}, E) - i\beta(\mathbf{r}, E)) \right] \tilde{\psi}(\mathbf{r}_\perp, z = 0, E). \quad (5.40)$$

Note that when we take the intensity of the exit wave, $I(\mathbf{r}_\perp) = |\psi(\mathbf{r}_\perp)|^2$, this simplifies to the Beer–Lambert law for attenuation-based imaging, where β is related to the linear




Figure 5.1: Schematic of propagation-based imaging using multi-slice modulation through N object slices, each with thickness $\Delta z = T/N$, where $\exp[\dots]$ is shorthand for the projection approximation equation (Eq. 5.40) and $\mathcal{D}_z$ is the free-space propagation operator over distance z .

attenuation coefficient,

$$\mu(\mathbf{r}, E) = 2k(E)\beta(\mathbf{r}, E). \quad (5.41)$$

5.3.2 Multi-slice modulation

If refractive effects within the object are not negligible, the initial PA assumption neglecting $\nabla_{\perp}^2 \psi(\mathbf{r})$ is not valid. This generally occurs with thick objects containing small features, such that the Fresnel number condition (Eq. 5.38) is not satisfied. In this case, a better model is the multi-slice (MS) method [4]. In this model, the object is divided into N slices that are each thin enough to safely satisfy PA. Then, the wave field's modulation through the object is modeled by sequentially applying PA (Eq. 5.40) and a free-space propagation operator over each slice thickness, as illustrated in Fig. 5.1.

While this is still an approximation of the inhomogeneous paraxial equation (Eq. 5.35), it is more general than PA and can be especially beneficial for forward models of high-resolution imaging systems [165, 166].

5.4 Imaging

We now have the formalism and a variety of forward models describing how complex X-ray wave fields propagate through different media. For imaging, we require an additional method for detecting the wave. We lack a method for directly measuring both real and imaginary parts of the full complex wave field, as X-ray detectors are only capable of measuring intensity. Thus, a variety of acquisition methods have been developed for converting the phase of the X-ray wave field into detectable intensity variations.

5.4.1 *Propagation-based imaging (PBI)*

Propagation-based imaging (PBI) is possibly the most conceptually simple XPCI acquisition technique, as it involves no specialized optical equipment. If the source is sufficiently coherent and the system has high enough spatial resolution, then phase-contrast effects can be apparent in the raw intensity images as high- and low-intensity fringes at material edges due to interference between the refracted and unrefracted portions of the wave field [23]. In the X-ray energy range, this generally requires a system spatial resolution no larger than tens of microns. This is much smaller than what is achieved with conventional clinical X-ray sources and detectors used for absorption-based imaging, but there is work developing sufficient equipment for medical translation [62, 64]. Currently, PBI is widely implemented at synchrotron facilities for micro-CT and has also been demonstrated with benchtop x-ray sources using sufficiently small focal-spot sizes [20, 53, 55, 167–169].

PBI forward modeling involves chaining together operations for complex object modulation, free-space propagation to the detector, and intensity detection. For the latter, the standard monochromatic component summation, PSF blurring, and detective efficiency models can be applied as used for absorption-only imaging (described in Chapter 2). However, even when the detector has relatively low resolution, for discrete modeling of wave evolution, it is important to upsample the wave field and free-space propagator to accurately account

for sub-pixel-resolution interference patterns.

5.4.2 *Other methods*

Despite the conceptual simplicity of PBI, it can be challenging to satisfy its strict spatial-coherence and detector-resolution requirements. Although this dissertation focuses on PBI, other XPCI acquisition strategies can encode phase effects into raw intensity images using additional optical elements and more complex experimental setups [22, 170]. For example, in grating-based interferometry, sample-induced refraction is measured from shifts in a periodic interference, which can enable simultaneous absorption, differential phase, and dark-field imaging. Grating-based methods are also being developed for translational CT applications, including recent demonstrations of human-scale dark-field CT prototypes [171]. Another choice is analyzer-based imaging, in which an analyzing crystal converts small angular deviations of the transmitted beam into intensity variations, providing sensitivity to refraction and ultra-small-angle scattering. Similarly, with edge-illumination and related speckle-based methods, masks or random structured illumination patterns are used to infer refraction and scattering from sample-induced distortions of the detected intensity pattern. Together, these methods illustrate a range of experimental strategies available for XPCI, while PBI remains especially attractive for this dissertation for its comparatively simple optics and forward modeling.

5.5 Phase retrieval

The goal of quantitative XPCI is to leverage raw intensity images to solve the phase retrieval inverse problem: recovery of both attenuation and phase-shifting maps of the object. Many phase retrieval methods have been developed and refined for a variety of imaging applications. In this section we give a brief overview of the three analytical methods most relevant to this dissertation. For simplicity, we assume a parallel beam geometry. If desired, there

are straightforward modifications to each method that account for a magnification factor, detailed in the original manuscripts [24, 55].

5.5.1 *Single-distance*

Perhaps the most ubiquitous approach is Paganin's method for phase retrieval from a single propagation-based image [24]. It is very experimentally convenient to acquire just one image, so a variety of methods using this acquisition approach have been developed. However, any single-measurement method must apply certain approximations to reduce the two independent unknowns (phase and attenuation) to just one. The key additional assumption of Paganin's method is that the object is made of a single material. In a comparison of seven single-distance PB-XPCI phase retrieval algorithms, this single-material method showed strong performance and ease of calculation [22].

To derive Paganin's algorithm, adopt a finite-difference approximated TIE and PA-based forward model. By assuming the object is homogeneous, we can drop the $\mathbf{r}$ -dependence of the refractive index $n(\mathbf{r}, E) \rightarrow n(E)$ and take it as given. For biological samples, a common choice of assumed material is water. Then, the PA expression simplifies to

$$\psi(\mathbf{r}_\perp, z = T, E) = \exp \left[-ik(E)T(\mathbf{r}_\perp)(\delta(E) - i\beta(E)) \right] \psi(\mathbf{r}_\perp, z = 0, E), \quad (5.42)$$

where $T(\mathbf{r}_\perp)$ is the projected thickness map of the assumed homogeneous object. If the wave is then propagated to a detector over distance R , $T(\mathbf{r}_\perp)$ can be solved for in Fourier space in terms of the measured intensity image $I_{\text{meas}}(\mathbf{r}_\perp)$ and other known parameters,

$$T(\mathbf{r}_\perp) = -\frac{1}{\mu(E)} \ln \left[\mathcal{F}^{-1} \left\{ \mu(E) \frac{\mathcal{F}\{I_{\text{meas}}(\mathbf{r}_\perp)\} / I_{\text{in}}}{R\delta(E)|\mathbf{k}_\perp|^2 + \mu(E)} \right\} \right], \quad (5.43)$$

where I_{in} is the uniform incident intensity. It is then trivial to generate phase and attenuation images from this thickness map as $\delta(E)T(\mathbf{r}_\perp)$ and $\mu(E)T(\mathbf{r}_\perp)$, respectively. Note that this

is essentially a matched low-pass filter that suppresses the high-frequency phase-contrast fringes as well as image noise, which can improve contrast-to-noise ratio when the algorithm's assumptions are satisfied.

5.5.2 Spectral

To avoid the homogeneous object approximation without necessitating multiple distance measurements, a potential single-shot solution could utilize energy-resolved measurements from a spectral photon-counting detector as proposed by Gürsoy and Das [55]. Like the single-material method, it implements a finite difference-approximated TIE and PA-based forward model, yielding monochromatic log-measurements

$$p_R(\mathbf{r}_\perp, E) = \int dz \mu(\mathbf{r}, E) + \frac{R}{k(E)} \nabla^2 \phi(\mathbf{r}_\perp, E). \quad (5.44)$$

Recall $\mu(E)$ can be written as the sum of its Compton scatter and photoelectric effect dependencies (Eq. 2.13),

$$\mu(E) = NK \frac{Z^5}{E^3} + NZ f_{\text{KN}}(E), \quad (5.45)$$

where N is concentration and K is a dimensionless constant. An analogous expression for the phase $\phi(E)$ utilizes its proportionality to electron density ρ_e ,

$$\phi(\mathbf{r}_\perp, E) = -\frac{2\pi r_e}{k(E)} \int dz \rho_e(\mathbf{r}), \quad (5.46)$$

where r_e is the classical electron radius. Then, a multi-energy phase retrieval solution can be found using an Alvarez–Macovski style decomposition,

$$p_R(\mathbf{r}_\perp, E) = \frac{K}{E^3} a_1(\mathbf{r}_\perp) + \left(f_{\text{KN}}(E) + \frac{2\pi r_e R}{k^2(E)} \nabla^2 \right) a_2(\mathbf{r}_\perp), \quad (5.47)$$

where $a_1(\mathbf{r}_\perp), a_2(\mathbf{r}_\perp)$ are the unknown basis projection images that can be solved using standard spectral imaging approaches [33]. It is then trivial to recombine the basis projections into phase and absorption projections using Eqs. 5.45–5.46.

The manuscript also discusses a strategy for using polychromatic incident radiation and multiple intensity images from different spectral PCD energy bins, using the median energy of each bin and a weighted-least squares solution [55].

5.5.3 *Material decomposition*

A similar spectral phase retrieval method was proposed by Schaff et al. to directly perform material decomposition from PB-XPCI measurements [59]. The method adopts the same linearized TIE model starting point for Gürsoy and Das (Eq. 5.44), written with the phase projection $\phi(\mathbf{r}_\perp)$ in terms of the object's phase-shifting component $\delta(\mathbf{r})$,

$$p_R(\mathbf{r}_\perp, E) = \int dz \mu(\mathbf{r}, E) - R \nabla^2 \int dz \delta(\mathbf{r}, E). \quad (5.48)$$

Leveraging the Fourier differentiation theorem and taking the FT of this equation, the Laplacian operator can be expressed as a multiplication with 2D frequency components,

$$\hat{p}_R(\mathbf{k}_\perp, E) = \mathcal{F} \left\{ \int dz \mu(\mathbf{r}, E) \right\} + R \mathbf{k}_\perp^2 \mathcal{F} \left\{ \int dz \delta(\mathbf{r}, E) \right\}. \quad (5.49)$$

Assuming a finite number of basis materials ($j = 1, \dots, B$), the monochromatic intensity projections can be expressed in terms of a linear sum of basis thickness projections,

$$\hat{p}_R(\mathbf{k}_\perp, E) = \sum_j^B A_j(E) \hat{T}_j(\mathbf{k}_\perp), \quad (5.50)$$

$$A_j(E) \equiv \mu_j(E) + R \mathbf{k}_\perp^2 \delta_j(E), \quad (5.51)$$

where $\widehat{T}_j(\mathbf{k}_\perp)$ is the FT of the thickness map $T_j(\mathbf{r}_\perp)$. Just as was described in Chapter 2 for traditional dual-energy absorption-only material decomposition, a system of equations can be set up using the FT-log projections with different acquisition energies, and the FT-thickness map is solved for each $\mathbf{k}_\perp$ component using a matrix inversion. The real-space thickness maps are recovered by taking the inverse FT.

5.6 Summary

This concludes the overview of essential phase-contrast imaging fundamental concepts for this dissertation. By deriving the four forward models (two for free-space propagation and two for object modulation), we now have a foundational understanding of each one’s underlying approximations. This provides context for the potential drawbacks of the most common phase retrieval algorithms.

The forward models presented here assume the incident wave is paraxial and in the X-ray energy range, and that the object is static, non-magnetic, and elastically scattering [23]. The most prolific XPCI inverse problem solutions tend to rely on the finite-difference approximated TIE and projection approximation [22, 24, 55, 59]. These can break down when propagation distances become too large, the object is thick, or the resolution is high [4]. Fresnel diffraction propagation and multi-slice modulation may yield a more accurate model at the cost of a more challenging or computationally intensive inverse solution [165, 166]. In the next chapters, we explore the regime of approximation validity and establish an inverse-agnostic estimation theory metric for assessing and optimizing best-achievable PB-XPCI performance.

5.7 Appendix: Derivation of the TIE

First substitute the original definition of the wave Ψ in terms of intensity I and phase ϕ (Eq. 5.1) into the free-space paraxial equation (Eq. 5.24):

$$\left(2ik(E)\frac{\partial}{\partial z} + \nabla_{\perp}^2\right) \sqrt{I(\mathbf{r}, t)} \exp[i\phi(\mathbf{r}, t)] = 0. \quad (5.52)$$

To expand, compute the partial derivative with respect to z and the transverse Laplacian separately,

$$\frac{\partial \Psi(\mathbf{r}, t)}{\partial z} = \exp[i\phi(\mathbf{r}, t)] \left(\frac{\partial \sqrt{I(\mathbf{r}, t)}}{\partial z} + i\sqrt{I(\mathbf{r}, t)} \frac{\partial \phi(\mathbf{r}, t)}{\partial z} \right), \quad (5.53)$$

$$\begin{aligned} \nabla_{\perp}^2 \Psi(\mathbf{r}, t) = \exp[i\phi(\mathbf{r}, t)] & \left(\nabla_{\perp}^2 \sqrt{I(\mathbf{r}, t)} - \sqrt{I(\mathbf{r}, t)} |\nabla_{\perp} \phi(\mathbf{r}, t)|^2 \right. \\ & \left. + 2i \nabla_{\perp} \sqrt{I(\mathbf{r}, t)} \cdot \nabla_{\perp} \phi(\mathbf{r}, t) + i\sqrt{I(\mathbf{r}, t)} \nabla_{\perp}^2 \phi(\mathbf{r}, t) \right). \end{aligned} \quad (5.54)$$

Substituting these into the paraxial wave equation and canceling the common factor $\exp[i\phi(\mathbf{r}, t)]$ gives:

$$\begin{aligned} 0 = \nabla_{\perp}^2 \sqrt{I(\mathbf{r}, t)} - \sqrt{I(\mathbf{r}, t)} |\nabla_{\perp} \phi(\mathbf{r}, t)|^2 - 2k\sqrt{I(\mathbf{r}, t)} \frac{\partial \phi(\mathbf{r}, t)}{\partial z} \\ + i \left(2k \frac{\partial \sqrt{I(\mathbf{r}, t)}}{\partial z} + 2\nabla_{\perp} \sqrt{I(\mathbf{r}, t)} \cdot \nabla_{\perp} \phi(\mathbf{r}, t) + \sqrt{I(\mathbf{r}, t)} \nabla_{\perp}^2 \phi(\mathbf{r}, t) \right). \end{aligned} \quad (5.55)$$

Taking the imaginary part yields

$$0 = 2k \frac{\partial \sqrt{I(\mathbf{r}, t)}}{\partial z} + 2\nabla_{\perp} \sqrt{I(\mathbf{r}, t)} \cdot \nabla_{\perp} \phi(\mathbf{r}, t) + \sqrt{I(\mathbf{r}, t)} \nabla_{\perp}^2 \phi(\mathbf{r}, t). \quad (5.56)$$

Multiplying through by $\sqrt{I(\mathbf{r}, t)}$ and using

$$\frac{\partial I(\mathbf{r}, t)}{\partial z} = 2\sqrt{I(\mathbf{r}, t)} \frac{\partial \sqrt{I(\mathbf{r}, t)}}{\partial z}, \quad \nabla_{\perp} I(\mathbf{r}, t) = 2\sqrt{I(\mathbf{r}, t)} \nabla_{\perp} \sqrt{I(\mathbf{r}, t)}, \quad (5.57)$$

we obtain

$$0 = k \frac{\partial I(\mathbf{r}, t)}{\partial z} + \nabla_{\perp} I(\mathbf{r}, t) \cdot \nabla_{\perp} \phi(\mathbf{r}, t) + I(\mathbf{r}, t) \nabla_{\perp}^2 \phi(\mathbf{r}, t). \quad (5.58)$$

Finally, recognizing that

$$\nabla_{\perp} \cdot (I(\mathbf{r}, t) \nabla_{\perp} \phi(\mathbf{r}, t)) = \nabla_{\perp} I(\mathbf{r}, t) \cdot \nabla_{\perp} \phi(\mathbf{r}, t) + I(\mathbf{r}, t) \nabla_{\perp}^2 \phi(\mathbf{r}, t), \quad (5.59)$$

we arrive at the general TIE:

$$-\nabla_{\perp} \cdot [I(\mathbf{r}_{\perp}, z) \nabla_{\perp} \phi(\mathbf{r}_{\perp}, z)] = k(E) \frac{\partial I(\mathbf{r}_{\perp}, z)}{\partial z}. \quad (5.60)$$

CHAPTER 6

CRAMÉR–RAO LOWER BOUND FOR SPECTRAL PB-XPCI

6.1 Introduction

As we have introduced, propagation-based X-ray phase-contrast imaging (PB-XPCI) is a valuable emerging technique that has enabled high-quality imaging of weakly attenuating structures [4, 163]. Traditional X-ray imaging relies solely on absorption contrast, which provides excellent discrimination between bones and soft tissue but often struggles to resolve small features and subtle variations within tissue. XPCI builds upon absorption-contrast imaging by aiming to measure phase shifts in the X-ray wave field, which can reveal changes within the object even when intensity shifts are small. However, phase retrieval suffers from a multi-measurement problem in order to recover the two independent components of the object: phase and absorption. Spectral PB-XPCI is a compelling practical solution for its potential to acquire multiple energy-resolved measurements in a single shot, and specific spectral phase retrieval solutions are still emerging. In this chapter, we introduce estimation theory toy models for PB-XPCI that can be used as inverse solution-agnostic benchmarks for achievable image quality and tools for system optimization [172, 173].

6.2 Background

X-ray detectors are sensitive to intensity only, so measurement of phase shifts requires specialized acquisition methods to convert them into detectable intensity variations [22, 24, 26]. In PB-XPCI, free-space propagation between the object and detector causes phase shifts to manifest visually as an edge-enhancement effect [23]. The method is predominantly used for synchrotron micro-CT and benchtop XPCI with a micro-focus X-ray source [20, 53, 55, 167–169]. While raw PB-XPCI images can qualitatively enhance the delineation of similarly absorptive structures, quantitative imaging analogous to attenuation coefficient reconstruc-

tion requires *phase retrieval*—the conversion of raw intensity measurements into maps of phase-shifting coefficients. Practical implementations often make various simplifying assumptions, such as treating the object as comprising a single material or being a pure phase object [24]. Even when these methods produce qualitatively appealing images, they may not yield accurate reconstructions. More generally, due to the two contrast mechanisms (attenuation and phase), if these other approximations are not valid, at least two independent acquisitions are needed for quantitative phase retrieval [174]. These are typically acquired at different propagation distances, which can be experimentally inconvenient and error-prone when one considers practicalities such as changes in magnification, alignment errors, and motion artifacts.

A promising alternative is *multi-energy* or “spectral” XPCI, in which the multiple independent measurements are obtained by using different X-ray spectra rather than by varying propagation distance [55, 59]. Multi-energy imaging has long been of interest in absorption-contrast X-ray imaging, where the multiple measurements are useful for performing material decomposition and generating high-contrast virtual monochromatic images [33]. For diagnostic CT imaging, multi-energy measurements have typically been acquired using some specialized acquisition setup, such as repeat scans with different peak tube voltages (kVp), fast-kVp switching, or multiple detector-source arrays. With the recent introduction of energy-discriminating spectral photon-counting detectors (PCDs) in clinical CT scanners, these measurements can now be conveniently acquired in a single shot [48, 49, 175]. The increasing accessibility of PCDs in diagnostic CT has generated further motivation for developing multi-energy imaging techniques. In the context of PB-XPCI, multi-energy imaging enables quantitative phase retrieval while minimizing experimental complexity; and with a spectral PCD, it may similarly be performed in a single shot, the gold standard for convenience [55, 57, 61, 176]. A recently proposed micron-resolution PCD design meeting the spatial resolution requirements of PB-XPCI further extends the potential of this approach

[62, 177].

Material decomposition from spectral PB-XPCI can be regarded as a phase retrieval problem, as it must convert raw edge-enhanced images into quantitative basis images. Schaff et al. developed an analytical method by inverting the linearized transport-of-intensity equation (TIE) and directly solving for basis material thickness images in Fourier space [59]. This is effectively equivalent to applying a matching low-pass filter that suppresses both edge enhancement and high-frequency noise [178]. It is conceptually similar to the Alvarez–Macovski model for attenuation-based decomposition [33]. The work demonstrated that with proper selection of the incident X-ray energies, both basis material images can benefit from phase information in addition to absorption contrast [178]. Though developed for monochromatic sources, this method could be extended to polychromatic spectra with prior knowledge of the imaging system and its noise characteristics [55, 176]. Other recent approaches to multi-energy PB-XPCI have explored novel acquisition techniques to ease imaging workflow and optimization-based reconstruction methods that may directly account for polychromatic source spectra [57, 61, 179–181]. This introduces many promising options for spectral XPCI yet also trade-offs in convenience, computational complexity, and possible accuracy.

For any imaging modality, theoretical image quality metrics such as the Cramér–Rao lower bound (CRLB) are indispensable tools for assessing and optimizing system performance. The CRLB is an estimate of the minimum achievable variance of an unbiased estimator and has been widely applied in spectral X-ray imaging, most notably with the basis material CRLB model detailed by Roessl and Herrmann, which has guided applications optimizing image acquisition protocols and noise-weighted image reconstruction [66, 131]. Likewise for XPCI broadly, analytical CRLB models have been developed for grating-based and analyzer-based acquisition techniques and can also indicate whether specific phase retrieval methods approach their theoretical performance limit [67, 68]. These existing methods are

each computed per a single detector pixel measurement.

A CRLB framework for spectral PB-XPCI would enable systematic evaluation of proposed material decomposition and phase retrieval approaches. Unlike the attenuation-only CRLB, a PB-XPCI CRLB must account for the spatial nature of the edge enhancement effect, with crucial information distributed over multiple detector pixels and influenced by the object geometry. This increases the model complexity with both a PB-physics forward model and several spatially-dependent acquisition parameters. We previously developed a numerical estimate of the spectral PB-XPCI CRLB using a finite-difference approximation, showing improved signal-to-noise ratio (SNR) of the basis material thickness estimate with increasing propagation distance [172]. In this work, we develop a CRLB framework for spectral PB-XPCI that incorporates realistic detector models and object geometry, and we demonstrate its utility for optimizing spectral imaging performance with additional phase-contrast information [173].

6.3 Theory

This section outlines the logical development of a spectral PB-XPCI CRLB toy model, building up from the absorption-only dual-energy CRLB to ultimately include spectral propagation-based phase information. The first claim we make is, since propagation-based phase contrast effects manifest in the raw intensity measurements as an edge-enhancement effect, it is essential to include multiple spatially resolved measurements in building the PB-XPCI Fisher information. To understand the implications of this, we first rehash the conventional absorption-only single-pixel model and illustrate why two basis materials cannot be estimated from a single scalar measurement; then, we introduce a two-pixel model to see how the problem is affected by additional spatial information with and without multiple energies. From there, we expand the forward model to include PB phase contrast and realistic imaging system effects, and we evaluate the Fisher information over one- and

two-dimensional toy models accounting for different types of geometric edge information.

6.3.1 Two-pixel model

Recall the general form of the Fisher information $\mathcal{I}$ for counting measurements that follow a Poisson distribution [122]:

$$\mathcal{I}_{uv} = \sum_k \frac{1}{m_k} \frac{\partial m_k}{\partial \theta_u} \frac{\partial m_k}{\partial \theta_v}, \quad (6.1)$$

where k is the independent measurement index, $\vec{\theta}$ is the vector of unknown parameters to be estimated, and m_k are the expected measurements. In the context of imaging, this is a reasonable model for any monochromatic measurements or for polychromatic measurements using a PCD.¹

In a simple version of the absorption-only CRLB, m_k are ideal measurements through a two-material object,

$$m_k \equiv m(E_k) = I(E_k) \exp[-\mu_1(E_k)T_1 - \mu_2(E_k)T_2], \quad (6.2)$$

where $I(E_k)$ is the incident intensity of the monochromatic beam with energy E_k , $\mu_j(E)$ is the linear attenuation coefficient of the j^{th} basis material, and T_j is the corresponding thickness. With this formulation, $\vec{\theta} = (T_1, T_2)$, and the partial derivatives for Eq. 6.1 are

$$\begin{aligned} \frac{\partial m(E_k)}{\partial T_j} &= -\mu_j(E_k) I(E_k) \exp[-\mu_1(E_k)T_1 - \mu_2(E_k)T_2,] \\ &= -\mu_j(E_k) m(E_k). \end{aligned} \quad (6.3)$$

1. In this chapter, we only examine situations with Poisson statistics. A polychromatic acquisition with an energy-integrating detector is better represented by an energy-weighted compound Poisson distribution, which has a different form for the Fisher matrix [66].




Figure 6.1: Two-pixel CRLB toy model.

If there is only a single energy measurement, substituting Eq. 6.3 into Eq. 6.1, we see the single-measurement 2×2 Fisher matrix is

$$\mathcal{I}(E_k) = \begin{bmatrix} \mu_1^2(E_k) m(E_k) & \mu_1(E_k) \mu_2(E_k) m(E_k) \\ \mu_1(E_k) \mu_2(E_k) m(E_k) & \mu_2^2(E_k) m(E_k) \end{bmatrix}, \quad (6.4)$$

which has determinant

$$\det \mathcal{I}(E_k) = \mathcal{I}_{uu}(E_k) \mathcal{I}_{vv}(E_k) - \mathcal{I}_{uv}(E_k) \mathcal{I}_{vu}(E_k) = 0, \quad (6.5)$$

and thus has no inverse.

This is the origin point for requiring multiple independent measurements to improve the conditioning of the material decomposition problem, traditionally taking the form of multiple energies E_k . With multiple energy measurements, so long as $\mu_j(E_k)$ are sufficiently different, $\mathcal{I}$ is invertible.

Instead, consider a system with a single energy measurement and two unique spatially-resolved measurements as in Fig. 6.1, which adopts the notation $m_i(E_k)$ with index i for the spatial coordinate. The first measurement $m_1(E_k)$ matches the form of the traditional intensity measurement, passing through both materials. The second measurement $m_2(E_k)$

passes through just the second “background” material and thus provides new information. By Eq. 6.1, the total Fisher matrix can be written as the sum of Fisher matrices for each independent measurement. The first independent measurement’s Fisher matrix $\mathcal{I}^{(1)}$ takes the form of Eq. 6.4, while the second pixel’s is

$$\mathcal{I}^{(2)} = \begin{bmatrix} 0 & 0 \\ 0 & \mu_2^2(E_k) m_2(E_k) \end{bmatrix}. \quad (6.6)$$

Then, the determinant of the full two-pixel system $\mathcal{I}^{(2\text{px})} \equiv \mathcal{I}^{(1)} + \mathcal{I}^{(2)}$ is

$$\det \mathcal{I}^{(2\text{px})} = \mu_1^2(E_k) \mu_2^2(E_k) m_1(E_k) m_2(E_k). \quad (6.7)$$

So long as each intensity measurement is greater than 0, this determinant is nonzero and the Fisher matrix is invertible. Recall the CRLBs are the diagonal elements of the inverse Fisher matrix, so the single-energy/two-pixel absorption-only CRLBs each have simple closed-form expressions,

$$\text{CRLB}_1^{(2\text{px})}(E_k) = \frac{1}{m_1(E_k) \mu_1^2(E_k)} + \frac{1}{m_2(E_k) \mu_1^2(E_k)}, \quad (6.8)$$

$$\text{CRLB}_2^{(2\text{px})}(E_k) = \frac{1}{m_2(E_k) \mu_2^2(E_k)}. \quad (6.9)$$

Likewise, the two-energy/single-pixel Fisher matrix can be written as the sum $\mathcal{I}^{(2E)} \equiv \mathcal{I}^{(1)}(E_1) + \mathcal{I}^{(1)}(E_2)$, which has its own CRLB solution described in Chapters 3 and 4 and is

commonly used as a toy model for absorption-contrast dual-energy imaging [66],

$$\text{CRLB}_1^{(2E)}(E_1, E_2) = \frac{m_1(E_1) \mu_{21}^2 + m_1(E_2) \mu_{22}^2}{m_1(E_1) m_1(E_2) (\mu_{11}\mu_{22} - \mu_{12}\mu_{21})^2} \quad (6.10)$$

$$\text{CRLB}_2^{(2E)}(E_1, E_2) = \frac{m_1(E_1) \mu_{11}^2 + m_1(E_2) \mu_{12}^2}{m_1(E_1) m_1(E_2) (\mu_{11}\mu_{22} - \mu_{12}\mu_{21})^2} \quad (6.11)$$

where for compactness we use notation $\mu_{jk} \equiv \mu_j(E_k)$ in these expressions.

To evaluate combined spectral PB-XPCI, we must consider a system with both edge information (multiple pixels) and spectral information (multiple energies). The simplest toy model is a two-energy/two-pixel system, which has Fisher information summing over both energy and spatial contributions,

$$\mathcal{I}^{(2E/2\text{px})} = \sum_i \sum_k \mathcal{I}^{(i)}(E_k), \quad (6.12)$$

where $i = 1, 2$ is the pixel index and $k = 1, 2$ is the energy index. Then, the CRLBs are calculated in the same way, as the diagonal elements of the inverse Fisher information matrix.

6.3.2 Line-profile model

The two-pixel toy model provides a simple form of spatial information, but it is not sufficient for resolving multiple propagation-based fringes. Instead, it serves as the foundation for a multi-pixel line-profile toy model that can be used to evaluate PB-XPCI as well as traditional imaging (Fig. 6.2). Building upon the absorption-only case, consider a full PB-XPCI forward model accounting for both real and imaginary parts of the object's refractive index and free-space propagation to a detector. For convenience, we define a single complex material factor $A(\mathbf{r}, E) \equiv \delta(\mathbf{r}, E) - i\beta(\mathbf{r}, E)$, which will be useful for compactly writing the forward model. Making the projection approximation (PA), the normalized intensity of an ideal monochromatic beam after passing through the object and propagating to the detector can




Figure 6.2: Multi-pixel line-profile toy model: a uniform plane wave is incident on a single-edge object. The PB-XPCI acquisition process is modeled, and then the measurements for the Fisher matrix are taken from a small region of pixels in a single row centered about the material edge.

be written as

$$I(\mathbf{r}_\perp, E) = \left| \mathcal{D}_R(E) \exp \left[-ik(E) \int dz A(\mathbf{r}, E) \right] \right|^2, \quad (6.13)$$

where $\mathbf{r}_\perp$ is the two-dimensional (2D) spatial coordinate (x, y) , R is the distance between the object and the detector, and $\mathcal{D}_R(E)$ is an energy-dependent free-space propagation operator. The propagation operator is the key element of this forward model translating the object's complex exit wave into measurable intensity fringes in the detector plane. We assume a linear form of $\mathcal{D}_R(E)$, the Fresnel diffraction integral (Eq. 5.18). This formulation is more general than the finite-difference approximated TIE widely applied for analytical phase retrieval algorithms [23, 24, 55, 59].

The normalized monochromatic components of Eq. 6.13 are used to define a more general equation for a non-ideal, polychromatic imaging setup. According to spectral decomposition, the polychromatic signal can be expressed as the sum of its individual energy components. So for acquisition a , we model the nonideal measured image in terms of monochromatic

components as

$$m_a(\mathbf{r}_\perp) = \Delta x^2 \int dE \left(\text{psf}_a(\mathbf{r}_\perp, E) * \eta_a(E) D_a(E) S_a(E) I(\mathbf{r}_\perp, E) \right), \quad (6.14)$$

where Δx is the width of each finite detector element (assumed square), $S_a(E)$ is the incident X-ray source spectrum (photon counts per energy per area), $\eta_a(E)$ is the detector efficiency function, $D_a(E)$ is the detector response function, $\text{psf}_a(\mathbf{r}_\perp, E)$ is the system point-spread function (PSF), and $*$ indicates the 2D convolution operation with respect to the spatial coordinate $\mathbf{r}_\perp$.

For the context of material decomposition, it is necessary to simplify Eq. 6.13 by defining a finite number of basis materials $j = 1, 2, \dots, B$ with spatially independent material factors $A(\mathbf{r}, E) \rightarrow A_j(E)$ and spatially varying projected thickness profiles $t_j(\mathbf{r}_\perp)$. The aim of material decomposition is to reconstruct images of each $t_j(\mathbf{r}_\perp)$ given known $A_j(E)$. Having discretized the object into a finite number of materials, the integral of Eq. 6.13 simplifies to a sum,

$$\int dz A(\mathbf{r}, E) = \sum_j^B t_j(\mathbf{r}_\perp) A_j(E). \quad (6.15)$$

For a line-profile single-edge model as shown in Fig. 6.2, we can write the 2D thickness profiles in terms of a single characteristic thickness T_j ($j = 1, 2$),

$$\sum_j^B t_j(\mathbf{r}_\perp) A_j(E) = T_1 A_1(E) \text{HS}(x - x_0) + T_2 A_2(E), \quad (6.16)$$

where $\text{HS}(x)$ is the Heaviside step function,

$$\text{HS}(x) = \begin{cases} 1, & x \geq 0 \\ 0, & x < 0, \end{cases} \quad (6.17)$$

and x_0 is the center of the object edge.

To calculate the basis material CRLBs for this model, we can take a small region-of-interest (ROI) of pixels in a single detector row centered about x_0 . Then the Fisher information matrix elements are

$$\mathcal{I}_{uv} = \sum_a \sum_{\mathbf{r}_\perp \in \text{ROI}} \frac{1}{m_a(\mathbf{r}_\perp)} \frac{\partial m_a(\mathbf{r}_\perp)}{\partial T_u} \frac{\partial m_a(\mathbf{r}_\perp)}{\partial T_v}. \quad (6.18)$$

This assumes the detector follows Poisson counting statistics and that the intensity measurements in each acquisition a and pixel $\mathbf{r}_\perp$ are independent.

To calculate the partial derivatives of the PB-XPCI measurements with respect to each thickness for substitution into Eq. 6.18, we can exploit the linearity of the convolution and free-space propagation operators along with a string of chain rule substitutions (detailed in the appendix, section 6.7) to see that for each toy model, we only need to compute the partial derivatives of the basis functions,

$$f_1(\vec{\theta}; \mathbf{r}_\perp) = T_1 \text{HS}(x - x_0), \quad f_2(\vec{\theta}; \mathbf{r}_\perp) = T_2, \quad (6.19)$$

where $\vec{\theta} = (T_1, T_2)$. The partials are

$$\frac{\partial f_1(\vec{\theta}; \mathbf{r}_\perp)}{\partial T_1} = \text{HS}(x - x_0), \quad \frac{\partial f_1(\vec{\theta}; \mathbf{r}_\perp)}{\partial T_2} = 0, \quad (6.20)$$

$$\frac{\partial f_2(\vec{\theta}; \mathbf{r}_\perp)}{\partial T_1} = 0, \quad \frac{\partial f_2(\vec{\theta}; \mathbf{r}_\perp)}{\partial T_2} = 1. \quad (6.21)$$




Figure 6.3: 2D geometry toy model for a system of two basis material cylinders.

Note that for accurate discrete calculation of pixel values within the ROI, the wave field should be upsampled and padded for free-space propagation [4].

6.3.3 2D geometry model

While the line-profile toy model is a logical starting point for incorporating XPCI information in the CRLB, it encodes edge enhancement for only a single material along a single axis. More generally, if more system geometry information is available, a 2D toy model can be built by defining $t_j(\mathbf{r}_\perp)$ as a separable function comprising a constant, characteristic thickness T_j and geometry factor $g_j(\mathbf{r}_\perp)$ describing the material's shape in the transverse plane,

$$t_j(\mathbf{r}_\perp) = T_j g_j(\mathbf{r}_\perp). \quad (6.22)$$

The constant thickness gives a single location-independent parameter for which we can estimate the basis material CRLB, while the geometry factor encodes edge information necessary to account for PB edge enhancement. Figure 6.3 shows a schematic of this acquisition setup for $B = 2$ basis material cylinders.

To compute the partial derivatives for the Fisher information, we use the following 2D

Table 6.1: Density and elemental composition (% mass) used for the basis materials [1].

Material	ρ [g/cm ³]	H	C	N	O	Na	Mg	P	S	Cl	K	Ca
Soft Tissue	1.06	10.2	14.3	3.4	70.8	0.2	–	0.3	0.3	0.2	0.3	–
Bone	1.92	3.4	15.5	4.2	43.5	0.1	0.2	10.3	0.3	–	–	22.5
Adipose	0.90	11.2	61.9	1.7	25.1	–	–	–	–	–	–	–
Gland	1.10	10.2	18.4	3.2	67.6	–	–	–	–	–	–	–

geometry basis functions

$$f_1(\vec{\theta}; \mathbf{r}_\perp) = T_1 g_1(\mathbf{r}_\perp), \quad f_2(\vec{\theta}; \mathbf{r}_\perp) = T_2 g_2(\mathbf{r}_\perp), \quad (6.23)$$

where $\vec{\theta} = (T_1, T_2)$ and $g_j(\mathbf{r}_\perp)$ is assumed to be fixed by the known object geometry. Then the partial derivatives are

$$\frac{\partial f_1(\vec{\theta}; \mathbf{r}_\perp)}{\partial T_1} = g_1(\mathbf{r}_\perp), \quad \frac{\partial f_1(\vec{\theta}; \mathbf{r}_\perp)}{\partial T_2} = 0, \quad (6.24)$$

$$\frac{\partial f_2(\vec{\theta}; \mathbf{r}_\perp)}{\partial T_1} = 0, \quad \frac{\partial f_2(\vec{\theta}; \mathbf{r}_\perp)}{\partial T_2} = g_2(\mathbf{r}_\perp). \quad (6.25)$$

6.4 Methods

The proposed toy model CRLBs were explored using two basis material pairs: 1-mm bone with 4-mm soft tissue, which presents a typical high-absorption contrast decomposition task often performed with conventional dual-energy imaging; and 1-mm fibroglandular (“gland”) tissue with 4-mm adipose tissue, a weakly attenuating material pair providing a relatively more challenging task. The elemental compositions used to compute the refractive indices of the basis materials is reported in Table 6.1.

First, to contextualize the additional information provided purely from multiple spatial measurements, both of the two-measurement toy models were compared (two-pixel/single-

energy versus two-energy/single-pixel) in the absence of phase information for an ideal monochromatic imaging system with 10^3 total incident photons. CRLB was computed at a range of energies between 8 and 32 keV. The combined two-energy/two-pixel model was also considered to assess the relative contributions of spectral and spatial information, respectively.

Second, the potential advantage of PB-XPCI was evaluated by expanding to a multi-pixel line-profile toy model, which is capable of resolving edge-enhancement effects for $R > 0$ with sufficiently high spatial resolution. For each material pair, a 20-pixel ROI was defined in a single detector row centered about the 1-mm material's edge (bone or gland). An ideal PCD was modeled with pixel width $\Delta x = 1 \mu\text{m}$. The incident number of photons per pixel was scaled so that the total counts summed over the entire ROI matched the earlier two-measurement systems. Basis material CRLBs were then computed for idealized monochromatic acquisitions at a range of energy pairs for either $R = 0 \text{ mm}$ (traditional absorption-only imaging) or $R = 100 \text{ mm}$ (PB-XPCI). The model was validated through comparison with the two-energy/two-pixel model for the absorption-only case.

The same line-profile model was also used to assess a nonideal, polychromatic PB-XPCI system using a spectral PCD for multiple energy-resolved intensity measurements. The source spectrum was generated based on a tungsten anode with 30-kV peak tube voltage and 0.5-mm aluminum filtering, as might be used for mammographic imaging [182]. To accurately simulate the continuous integral over all energy components of the full polychromatic wave field, the spectrum was discretized into monochromatic component bins of width $\Delta E = 0.05 \text{ keV}$. The total flux of the spectrum was matched to the monochromatic system, such that $\sum_i S(E_i)\Delta E = 10^3 \text{ per } \mu\text{m}^2$, with single-micron pixels.

A spectral detection scheme was defined in terms of thresholds energies E_t , which corresponded to energy “bin” detector response functions $D_a(E)$ used to generate the multiple energy-resolved acquisitions. Given a single threshold, an ideal PCD would simply sort each




Figure 6.4: Illustration of the detector response functions $D_a(E)$ modeled for the spectral PCD system given single energy thresholds of $E_t = 15, 20$, or 25 keV for a normalized 30-kVp spectrum $S(E)$ with a 5-keV FWHM spectral threshold blur.

detected photon into either the low or high bin based on whether its true energy E_0 is lower or higher than E_t , respectively. We introduced a more realistic detection model that assumes each photon has finite energy resolution with a Gaussian probability distribution of being detected at energy E , with mean $= E_0$ and full-width at half-maximum (FWHM) $= 0, 2$, or 5 keV. Figure 6.4 shows the full polychromatic spectrum filtered by example detector response functions for different thresholds.

With this polychromatic imaging model, the line-profile CRLB was computed as a function of threshold energy for both a single-threshold (two-bin) system and a two-threshold (three-bin) system, probing the utility of additional energy-resolved information and the optimal thresholds for each basis material.

Finally, to assess generalizability, the line-profile model was extended to a 2D geometry toy model for the single-threshold polychromatic imaging system. The basis object geometry functions $g_j(\mathbf{r}_\perp)$ were defined as cylinders with diameter $w = 30$ μm . The centers of the cylinders were offset by a small shift (20% of the radius along both x and y axes) to disentangle a portion of the PB edge-enhancement effect for each basis material while maintaining contrast overlap, which is fairly representative of a typical imaging task. The detector was modeled as a 64×64 array of 1- μm pixels with a Gaussian spatial PSF (FWHM $= 2$ μm). The spectral energy threshold model incorporated a nonideal Gaussian FWHM

of 5 keV. CRLB was evaluated as a function of threshold energy for both bone/tissue and gland/adipose material pairs.

To check how well the CRLB predictions would correspond with experimental imaging scenarios, spectral XPCI simulations were performed for the same input parameters. Variance and bias of the subsequent thickness estimates was measured over $N = 10^5$ unique noisy realizations for each. The material decomposition was performed using a maximum-likelihood estimation (MLE) approach; for each set of simulated images $\mathbf{y}$, the estimated T_1^* and T_2^* were found as the arguments minimizing the Poisson negative log-likelihood function,

$$\mathcal{L}(T_1, T_2) = \sum_i [\boldsymbol{\lambda}_i - \mathbf{y}_i \log(\boldsymbol{\lambda}_i) + \log(\mathbf{y}_i!)] , \quad (6.26)$$

where $\boldsymbol{\lambda} = f(T_1, T_2)$ is the array of forward projected images for a given thickness pair and i is the array index. The forward model f was written using the wave-optics library `chromatix` and used for both simulation and optimization [183]. We used this to assess how well MLE material decomposition could saturate the CRLB on variance and, secondarily, to identify whether the optimal acquisition energy parameters were the same according to both the CRLB and MLE simulations.

6.5 Results and Discussion

6.5.1 Two-measurement models (*absorption only*)

To understand the individual information contributions of spatial (multi-pixel) and spectral (multi-energy) information, Fig. 6.5 directly compares the magnitude of the CRLB for each two-measurement system. At all energies, the two-pixel system yields a CRLB much smaller than the two-energy system. This is not particularly surprising, since the two-pixel system includes additional assumptions about the object composition that significantly improve the conditioning of the problem. Recall that the ideal dual-monoenergy material decomposition




Figure 6.5: Comparison of basis material CRLBs with either a two-pixel or two-energy system for 1-mm bone and 4-mm soft tissue background. The two-energy CRLB curve shows the minimum achieved for a given low energy, i.e. $\text{CRLB}_j^{(2E)}(E_1) = \min \text{CRLB}_j(E_1, E_2)$ for $E_2 > E_1$. The full two-energy $\text{CRLB}_j(E_1, E_2)$ heatmaps are shown in Fig. 6.6.




Figure 6.6: Two-energy/single-pixel CRLB heatmap for all energy pairs for 1-mm bone and 4-mm soft tissue. The white “+” indicates the minimum CRLB coordinates, at (11.8 keV, 23.0 keV) for both materials. The $E_1 = E_2$ diagonal is not defined for this system.

task can be solved by inverting a simple 2×2 matrix of linear attenuation coefficients as described in Chapter 2,

$$\begin{bmatrix} p_1(E_1) \\ p_1(E_2) \end{bmatrix} = \begin{bmatrix} \mu_1(E_1) & \mu_2(E_1) \\ \mu_1(E_2) & \mu_2(E_2) \end{bmatrix} \begin{bmatrix} T_1 \\ T_2 \end{bmatrix}, \quad (6.27)$$

where p_i is the log-measurement of the i^{th} pixel, and the two-energy system only considers the pixel including both materials ($i = 1$). Analogously, the two-pixel system with a single energy E_k can be written as

$$\begin{bmatrix} p_1(E_k) \\ p_2(E_k) \end{bmatrix} = \begin{bmatrix} \mu_1(E_k) & \mu_2(E_k) \\ 0 & \mu_2(E_k) \end{bmatrix} \begin{bmatrix} T_1 \\ T_2 \end{bmatrix}, \quad (6.28)$$

which is upper triangular and thus straightforward to solve by back substitution. Note that for polychromatic acquisitions, these solution methods generally require effective monoenergies to generate the 2×2 matrices of linear attenuation coefficients.

For our context of spectral imaging, we are primarily interested in system performance as a function of energy. For each basis material j , the energy E^* that optimizes system performance is the one minimizing the CRLB, i.e.

$$E_j^* \equiv \arg \min_E \sigma_{T_j}^2(E). \quad (6.29)$$

The optima of each CRLB are located at the zeros of the first derivative with respect to E . These optima are minima if the second derivative is greater than 0. As a stronger condition, if the CRLBs are strictly convex as a function of E , there is guaranteed to be a single minimizing optimum energy. Due to the elaborate energy-dependence of the linear attenuation

coefficient in terms of contributions from Compton scatter (Klein–Nishina function), it is quite tedious to write closed form first and second derivatives, so we instead note that this is true by observation and with numerical verification for both systems (Fig. 6.5).

In fact, the two-pixel solution appears to be algebraically equivalent to an idealized log-subtraction decomposition method used for digital subtraction angiography (DSA) [184–186]. DSA is a fluoroscopic technique used to image vasculature by acquiring images both with and without a contrast agent, then subtracting the paired log-projections to remove obfuscating background structures [184, 185]. The non-contrast projections are analogous to the background-only pixel in our two-pixel case. In the context of DSA, some existing work has compared conventional subtraction with proposed energy-resolved angiography (ERA) methods, much like the comparison of two-pixel and two-energy toy models here [184–188]. A key motivation for ERA is to reduce artifacts associated with multiple acquisitions needed for DSA, for example motion and misregistration [186, 188, 189]. In terms of SNR, traditional DSA still generally outperforms ERA due to noise amplification associated with spectral material decomposition, although system optimization is improving its future prospects [184–186, 188]. Therefore, it is not surprising to see in Fig. 6.5 that the two-pixel method outperforms the two-energy method. While we continue to focus on spectral XPCI here, these two toy-model CRLB frameworks may also be useful for future work in angiography.

Fig. 6.7 shows the two-energy/two-pixel CRLB heatmap as a function of energy pair for bone/tissue. Unlike the traditional two-energy single-line integral CRLB (Fig. 6.6), the two-pixel CRLB is minimized at a point along the diagonal $E_1 = E_2$, 15 keV for bone and 10 keV for soft tissue. This indicates that, when adding a background-only measurement to a dual-energy imaging problem, it is actually better to acquire a single energy measurement at the optimal value rather than to divide the dose budget between distinct energies. Further, returning to Fig. 6.5, the optimal energy “pair” for the two-energy/two-pixel CRLB appears




Figure 6.7: Two-energy/two-pixel CRLB for 1-mm bone and 4-mm soft tissue background. The white “+” indicates the minimum CRLB coordinates, at (15 keV, 15 keV) for bone and (10 keV, 10 keV) for tissue.




Figure 6.8: Two-energy/two-pixel CRLB for 1-mm fibroglandular tissue and 4-mm adipose tissue background. The white “+” indicates the minimum CRLB coordinates, at (8.8 keV, 8.8 keV) for gland and (8.0 keV, 8.0 keV) for adipose.

to be the same as the one optimal energy for the single-energy/two-pixel CRLB. This matches expectation. Along the diagonal $E_1 = E_2$, the two-energy/two-pixel Fisher information (Eq. 6.12) simplifies to

$$\mathcal{I}^{(2E/2\text{px})}(E_1, E_2 = E_1) = \sum_i \mathcal{I}^{(i)}(E_1), \quad (6.30)$$

which is identical to the single-energy/two-pixel Fisher information. Therefore, if the total photons are matched, the diagonal of the two-energy/two-pixel CRLB matrix will match the shape of the single-energy/two-pixel CRLB curve. If there exists a unique minimizer E^* for the single-energy/two-pixel CRLB, the value at (E^*, E^*) must be either a minimum or a saddle point for the two-energy/two-pixel case. This does not necessarily exclude the existence of other local minimizers; however, none are apparent for the bone/tissue system examined in Fig. 6.7. Similarly, Fig. 6.8 shows the two-energy/two-pixel CRLB for a more weakly attenuating material pair of fibroglandular tissue on an adipose tissue background, which also has a single global minimizer along the $E_1 = E_2$ diagonal.

6.5.2 Line-profile model (XPCI)

As an initial check, we expect the multi-pixel line-profile CRLBs to match the two-pixel CRLBs when $R = 0$, the edge is centered in the ROI at $x_0 = 0$, and the total photon counts are matched. Figs. 6.9 and 6.10 confirm this, showing the line-profile CRLBs as a function of dual-energy pair for both bone/soft-tissue and gland/adipose-tissue objects, dose matched to their corresponding two-pixel CRLB heatmaps of Figs. 6.7 and 6.8. Example line profiles are also presented to contextualize the raw edge enhancement. The $R = 0$ case perfectly matches the two-pixel CRLBs at all energies. With $R = 100$ mm, the optimal signal-material CRLB improves by several orders of magnitude (bone or gland), while the background CRLB changes only slightly (soft tissue or adipose tissue). This reflects the new




Figure 6.9: Bone/soft-tissue basis: line-profile dual-energy CRLB heatmaps for $R = 0$ mm (traditional) and 100 mm (XPCI), with line profiles for two example energies.




Figure 6.10: Fibroglandular/adipose-tissue basis: line-profile dual-energy CRLB heatmaps for $R = 0$ mm (traditional) and 100 mm (XPCI), with line profiles for two example energies.

presence of edge-enhancement information for the first material alone. Furthermore, while the 0-mm heatmaps are smooth and convex, the 100-mm signal-material CRLB heatmap has several additional local minima. This is likely a result of the partial averaging effects of sub-pixel interference patterns displayed at the finite detector resolution, which appears to give rise to a periodicity in the optimal energies. The global minimum is still along the $E_1 = E_2$ diagonal, although the optimizing value depends on the propagation distance, increasing slightly from 15.0 to 17.6 keV for bone and from 8.8 to 10.2 keV for gland.

Figs. 6.11 and 6.12 extend the line-profile CRLB to the nonideal polychromatic system with an energy-discriminating PCD using spectral uncertainty FWHM of 0, 2, or 5 keV. The solid black line shows the CRLB as a function of threshold energy for the one-threshold (two-bin) system, while the thin marked lines show the results for different two-threshold (three-bin) systems. The first observation is that, unlike the ideal monochromatic system, the minimizing thresholds are always at intermediate values, affirming the utility of spectral information with a polychromatic source. If the minimizing threshold were instead at the minimum or maximum possible energy, this would imply that spectral information does not provide useful additional information. Second, the one-threshold curve serves as an upper envelope for the two-threshold curves, indicating that further energy-resolved information can only improve CRLB. Third, like the monochromatic system, the ideal threshold systems (FWHM = 0 keV) show rapid oscillations in local minima, but these are gradually smoothed out as FWHM increases to 5 keV. This may reflect the particularly informative nature of specific intermediate energies, which gradually reduce their relative contribution to either bin when the energy threshold is swept across. This effect is not as prominent in the background material (soft tissue or adipose tissue), likely due to the lack of edge enhancement as a uniform slab.




Figure 6.11: Bone/soft tissue polychromatic line-profile CRLBs as a function of threshold energy. The solid black line is for a one-threshold (two-bin) system. The thin marked lines show different two-threshold (three-bin) systems, with color/marker indicating the lower energy threshold and the horizontal axis the higher threshold.




Figure 6.12: Gland/adipose tissue polychromatic line-profile CRLBs as a function of threshold energy. The solid black line is for a one-threshold (two-bin) system. The thin marked lines show different two-threshold (three-bin) systems, with color/marker indicating the lower energy threshold and the horizontal axis the higher threshold.




Figure 6.13: Example noisy simulated images used for the MLE validation study with a 20-keV energy threshold.

6.5.3 2D geometry model and MLE validation

Figure 6.13 shows example noisy polychromatic simulation images used for the MLE validation, and Figs. 6.14 and 6.15 show the CRLB and MLE variance curves for the bone/soft tissue basis and gland/adipose basis, respectively. The trends in MLE as a function of threshold energy closely matched the CRLB, with the same energies minimizing both MLE variance and CRLB for each material. Further, in comparison with the line-profile case, the shape of the first material (bone or gland) CRLB curve closely matched the corresponding line-profile case (single-threshold and spectral FWHM = 5 keV), with minimizers at 21 keV for bone and 16 keV for gland. This reflects the shared information content at the edge of bone/gland in both models. The shape of the background material curve (soft tissue or adipose tissue) differed, with the introduction of multiple local minima, a result of the newly introduced edge information. The optimal energy threshold was unique for each material, with a slightly higher value for the denser material of each basis relative to the less dense material (16 keV for soft tissue and 15 keV for adipose). All optimal energy thresholds were at intermediate energy values, again indicating the benefits of spectral separation in context of material decomposition with a polychromatic source.

The MLE variance was slightly greater than the CRLB, with average MLE:CRLB ratios of 1.11 ± 0.04 for bone, 2.31 ± 0.08 for soft tissue, 2.43 ± 0.14 for gland, and 1.77 ± 0.56 for




Figure 6.14: For the bone and soft tissue basis, comparison of CRLB and MLE simulation variance as a function of threshold energy for the polychromatic spectral PCD system. For each panel, the CRLB (solid line) uses the left vertical axis and the MLE (dashed line) uses the right vertical axis for scale. A ± 1 keV window moving average is shown over the MLE variances, which were computed for individual E_t in 0.1 keV intervals.




Figure 6.15: For the gland and adipose tissue basis, comparison of CRLB and MLE simulation variance as a function of threshold energy for the polychromatic spectral PCD system. For each panel, the CRLB (solid line) uses the left vertical axis and the MLE (dashed line) uses the right vertical axis for scale. A ± 1 keV window moving average is shown over the MLE variances, which were computed for individual E_t in 0.1 keV intervals.




Figure 6.16: For the bone and soft tissue basis, bias of the MLE thickness estimates.




Figure 6.17: For the gland and adipose tissue basis, bias of the MLE thickness estimates.

adipose. This may be related to a bias in the simulation thickness estimates. Figs. 6.16 and 6.17 show the bias as a function of threshold energy for bone/tissue and gland/adipose, respectively. All materials had a slight negative bias, though it was on the order of a few microns and less than 1% of the true thickness at all thresholds. Bias was generally the smallest near the variance-minimizing threshold energies of each material. Since the CRLB is a bound on the lowest achievable variance of unbiased estimator, it may not be fully predictive in the presence of significant bias.

6.6 Conclusion

In this work, we propose a CRLB framework for PB-XPCI and use it to examine multi-energy XPCI for material decomposition. Spectral XPCI is a promising emerging approach to phase-contrast imaging, owing to its ability to acquire multiple independent measurements for quantitative phase retrieval or material decomposition in a single-shot and thus bypassing inconvenient multi-shot acquisitions or potentially inaccurate approximations. A generalizable optimization framework is essential to realize its full potential.

The CRLB provides a valuable, phase retrieval-agnostic metric for assessing and optimizing system performance. This has been demonstrated previously for other XPCI acquisition systems and for material decomposition in traditional absorption-only X-ray imaging,

yet a gap has existed for spectral PB-XPCI specifically [66–68]. This may be because the PB-XPCI CRLB must accommodate multiple detector pixel measurements, since key information resides in the high- and low-intensity fringes that develop at material edges in the image, introducing a new layer of complexity to traditional theoretical toy models [172]. We accounted for this by treating spatially-resolved pixel measurements as independent and summing over them in the Fisher information calculation, a reasonable first approximation similar to the conventional assumption of independent measurements in different spectral energy bins [66]. By progressively building up from just two pixels to a full 2D geometry toy model, we are able to elucidate the information contributions of spectral, spatial, and phase-contrast information.

Interestingly, the spatially resolved measurements in the CRLB model parallel the temporally resolved measurements acquired for angiography [184–186]. Past work has indicated promise for joint energy-resolved and conventional time-resolved DSA with a spectral PCD [190]. This would indicate similar possible benefits in the proposed CRLB models even in the absorption-only domain. Though the multi-energy benefit did not necessarily translate for the ideal monochromatic imaging system, we saw improvements with a polychromatic system and spectral PCD, as was used in the prior work on energy-resolved angiography.

The different apparent spectral benefit for monochromatic versus polychromatic systems is possibly related to the fact that the simple 2×2 matrix inversion material decomposition strategy (Eqs. 6.27 and 6.28) is only valid for monochromatic acquisitions. With polychromatic measurements, this simple decomposition method necessitates collapsing the full spectrum to a single monoenergy for each acquisition. In reality, a polychromatic acquisition contains more complicated information due to the many energy components summed together, especially when including propagation-based edge enhancement. Further, the effective energies of the spectrum cannot be tuned at will as they can for a monoenergetic acquisition, since they are constrained to the range of possible pairs as a function of thresh-

old energy. Thus, the monochromatic analysis suggesting that there is a single optimal monoenergy is not applicable to the spectral polychromatic system. The polychromatic CRLB reflects this, finding an optimal spectral threshold for splitting a single polychromatic image into two energy “bin” images.

An additional note regarding the dual-monoenergy system concerns the effect of the incident spectra $S_a(E)$ as defined for the general measurement model (Eq. 6.14). While the polychromatic analysis utilized a spectrum function generated based on a tungsten-anode source, the monochromatic dual-energy system modeled here simply fixed $S(E) = 1$ for each acquisition. With monochromatic synchrotron radiation, this is not an unreasonable simplification in the context of sequential dual-energy scans. While a synchrotron’s flux does change as a function of energy, with sequential scans, the acquisition time can be modulated such that the total intensity per acquisition is equal. More generally, it could be possible to modulate acquisition time and find some optimal intensity distribution across monoenergies; i.e., while we show that when monoenergetic $S_i \equiv S(E_i) = 1$ the optimal energy pair converges to a single energy, we do not consider the case of $S_1 \neq S_2$. In Chapter 3, we showed that MV-kV dual-energy imaging could outperform single-kV imaging, but only with optimal intensity allocation where roughly 10% of the total dose was delivered from the MV spectrum. It is possible that a dual-monoenergetic system could see similar benefits if the spectral allocation is optimized rather than fixed equal to 1. The CRLB framework developed here enables such analysis, which we leave to future work.

Despite the additional model complexity, the 2D component has the advantage of allowing for assessment of several other system parameters simultaneously, including detector pixel size and object width, providing a more holistic view of PB-XPCI performance. This also opens up potential for exploring a variety of toy model geometries beyond the two-cylinder model examined in this work. Uniquely, it can account for certain realistic 2D effects such as source and detector PSF, which opens up the potential for propagation distance opti-

mization in addition to multi-energy analysis when using a cone-beam source that includes a magnification factor.

There is some prior work considering alternative methods for quantifying theoretical material decomposition performance for DE PB-XPCI using phase retrieval methods that analytically invert the linearized TIE [24, 59, 178]. This has used a simpler model, but the overall findings are consistent with our CRLB trends. Notably, the CRLB offers a more general performance assessment framework that is independent of any specific material decomposition or phase retrieval algorithm. Because TIE-based approaches rely on approximations that may not hold in regimes where our CRLB indicates improved best-achievable SNR (e.g. small Δx and sufficiently large R), material decomposition methods that relax TIE assumptions may yield performance gains [57, 61, 176, 179, 180]. Future work could compare these approaches systematically across conditions where CRLB and TIE validity diverge, highlighting the most promising avenues for improving material decomposition performance. A new element of our estimation theory framework is the inclusion of a polychromatic spectrum and energy-dependent detector response.

As with any theoretical image quality metric, certain limiting simplifications are unavoidable. Here, we focused on Poisson counting statistics and did not model effects such as scatter or more complex PCD effects such as pulse pile-up and dead time. Nonetheless, we were able to include other realistic elements such as the system PSF, which cannot be modeled in a typical single-line integral CRLB [66]. Extending the forward model and derivative expressions to include more effects could be an interesting direction for future work building upon this XPCI-CRLB.

In conclusion, our proposed XPCI-CRLB framework establishes a useful foundation for optimizing spectral PB-XPCI systems and advancing their translation toward practical biomedical applications. Our results affirm the utility of multi-energy information with a polychromatic spectral PCD setup, an XPCI system of particular interest for its convenient

single-shot acquisition implementable in benchtop systems. We present a logical progression of models, beginning with the traditional single-pixel CRLB and gradually including more spatial and phase-contrast information. We validated the 2D toy model against a simulation study to confirm that CRLB optimization predictions translate to practical imaging scenarios. Future experiments might implement the CRLB with specific toy model material configurations emulating the imaging task at hand, allowing early optimization of acquisition parameters to achieve the best material decomposition and phase retrieval performance.

6.7 Appendix: Partial derivatives of PB-XPCI measurements

For a generic material decomposition toy model, the projected thickness term can be parameterized in terms of basis functions f_j ,

$$\sum_j^B t_j(\mathbf{r}_\perp) A_j(E) = \sum_j^B f_j(\vec{\theta}; \mathbf{r}_\perp) A_j(E), \quad (6.31)$$

where $\vec{\theta}$ is the vector of parameters to be estimated. For example, for a standard two-material decomposition problem, $\vec{\theta} = (T_1, T_2)$.

Calculation of the Fisher matrix requires the partial derivatives of the intensity measurements with respect to each parameter, which is

$$\frac{\partial}{\partial \theta_u} m_a(\mathbf{r}_\perp; \vec{\theta}) = \Delta x^2 \int dE \left(\text{psf}_a(\mathbf{r}_\perp, E) * \eta_a(E) D_a(E) S_a(E) \frac{\partial I(\mathbf{r}_\perp, E; \vec{\theta})}{\partial \theta_u} \right), \quad (6.32)$$

since the monochromatic components $I(\mathbf{r}_\perp, E)$ are the only constituent terms that depend on $\vec{\theta}$, and the integral and convolution are linear operators. To simplify, each monochromatic

intensity can be written in terms of the complex wave field incident on the detector,

$$I(\mathbf{r}_\perp, E; \vec{\theta}) = \left| \psi_R(\mathbf{r}_\perp, E; \vec{\theta}) \right|^2, \quad (6.33)$$

$$\psi_R(\mathbf{r}_\perp, E; \vec{\theta}) \equiv \mathcal{D}_R(E) \exp \left[-ik(E) \sum_j^B f_j(\vec{\theta}; \mathbf{r}_\perp) A_j(E) \right]. \quad (6.34)$$

The identity for the derivative of the squared modulus of a complex function $|a(x)|^2 \equiv a(x)\overline{a(x)}$ is

$$\frac{\partial}{\partial x} |a(x)|^2 = 2 \operatorname{Re} \left[\overline{a(x)} a'(x) \right], \quad (6.35)$$

where the overbar indicates the complex conjugate. Since the propagation operator is linear and independent of $\vec{\theta}$,

$$\frac{\partial}{\partial \theta_u} \psi_R(\mathbf{r}_\perp, E; \vec{\theta}) = \mathcal{D}_R(E) \frac{\partial}{\partial \theta_u} \exp \left[-ik(E) \sum_j^B f_j(\vec{\theta}; \mathbf{r}_\perp) A_j(E) \right]. \quad (6.36)$$

This simply requires evaluation of the partial derivative of the argument of the exponent, since the derivative of an exponentiated function is

$$\frac{\partial}{\partial x} e^{a(x)} = a'(x) e^{a(x)}. \quad (6.37)$$

Applying this to Eq. 6.36,

$$\frac{\partial}{\partial \theta_u} \left(-ik(E) \sum_j^B f_j(\vec{\theta}; \mathbf{r}_\perp) A_j(E) \right) = -ik(E) \sum_j^B A_j(E) \frac{\partial f_j(\vec{\theta}; \mathbf{r}_\perp)}{\partial \theta_u}. \quad (6.38)$$

This demonstrates that the only model-specific derivative needed is the partial derivative of each basis function f_j with respect to the estimated parameters. Once these derivatives are specified, the remaining operations follow from the common PB-XPCI forward model,

as these derivatives are substituted back into the chain of equations above.

$$\frac{\partial}{\partial \theta_u} \psi_R(\mathbf{r}_\perp, E; \vec{\theta}) = \mathcal{D}_R(E) \left\{ \left[-ik(E) \sum_j^B A_j(E) \frac{\partial f_j(\vec{\theta}; \mathbf{r}_\perp)}{\partial \theta_u} \right] \exp \left[-ik(E) \sum_j^B f_j(\vec{\theta}; \mathbf{r}_\perp) A_j(E) \right] \right\}. \quad (6.39)$$

CHAPTER 7

ACCURACY OF THE PROJECTION APPROXIMATION

7.1 Introduction

PB-XPCI faces several hurdles before clinical realization and optimal implementation. A key challenge is the need for both accurate and computationally efficient discrete forward modeling. Free-space propagation simulation generally requires sub-micron sampling, even for detectors with relatively low resolution, in order to account for partial averaging effects over the high-frequency interference patterns. With increasingly large fields-of-view (FOVs), 3D volumes, and tomography rotation angles, the memory requirements of a PB-XPCI simulation can compound rapidly. Further, to solve the phase retrieval problem, one must invert the chosen forward model, which may not have an analytical solution and thus can magnify the computational challenge. For this reason, several approximations are commonly implemented in XPCI forward models.

One of the most ubiquitous simplifications is the projection approximation (PA), which neglects diffraction and refraction within the object by treating the object as a single projected complex transmission function [23]. This is typically valid for thin objects, high X-ray energies, and large transverse feature sizes/lower spatial resolution [4, 163]. Under PA, the X-ray wave field exiting an object is modeled according to the line integrals through the object's complex index of refraction. If one can recover the full complex exit wave from the raw data, under PA, it is straightforward to use the real and imaginary line integral component sinograms as input to a conventional CT reconstruction method such as filtered back-projection (FBP). But if the conditions of PA are not satisfied, conventional phase retrieval methods will not necessarily recover exact line integrals, jeopardizing the accuracy of both recovered projections and tomographic reconstructions.

It is essential to understand the validity of approximations under various conditions to

make informed decisions in imaging simulations and inverse solutions. Given the apparent direction of advances in micro-CT technology, large FOVs and high spatial resolution, PA is an approximation of particular interest [62, 169]. Note that in CT imaging, FOV is the limiting factor for both the maximum width and maximum thickness of the object, as it is rotated and imaged over many view angles. Thus, the purpose of this chapter is to assess the accuracy of PA and to compare it to a more general model in the context of PB-XPCI forward modeling and existing inverse problem solutions, including conventional phase retrieval and CT reconstruction.¹

7.2 Background

A detailed derivation of the PA model and its founding assumptions can be found in Chapter 5. In short, given a wave field $\psi(\mathbf{r})$ and object with refractive index $n(\mathbf{r}, E)$, the PA forward model follows from the inhomogeneous paraxial equation (Eq. 5.35) when one assumes that $\nabla_{\perp}^2 \psi(\mathbf{r}) \approx 0$ within the object. This is valid in practice when refractive effects within the object are negligible, which is generally true of thin objects with relatively large features [4].

A simple unitless metric for PA validity uses the Fresnel number N_F ,

$$N_F = \frac{\Delta x^2}{\lambda(E)T}, \quad (7.1)$$

where Δx is the smallest feature size of interest (assigned here as detector pixel size), $\lambda(E)$ is the X-ray wavelength, and T is the object thickness [4]. PA is typically valid when $N_F \gg 1$, with a hard cut-off of $N_F > 10$ proposed [163]. Considering the context of advancing imaging technology, Fig. 7.1 shows heatmaps of the Fresnel number for ranges of parameters corresponding to conventional and emerging micro-CT, mammography, and clinical CT (parameters listed in Table 7.1). The condition $N_F \gg 1$ is well satisfied for clinical CT even

1. Parts of this chapter are based on conference proceedings papers available on ArXiv [191] and [164] under the CC BY-NC-SA 4.0 licence.

Table 7.1: Parameter ranges used to represent the three imaging modalities for Fresnel number heatmaps (Fig. 7.1).

Imaging modality	Energy E [keV]	Resolution Δx [μm]	Thickness z [cm]
Micro CT	14	0.1 – 5.0	0.1 – 2.0
Mammography	18	5 – 50	2 – 10
Clinical CT	50	50 – 400	10 – 50

with potential $\sim 100\text{-}\mu\text{m}$ high-resolution PCD pixel sizes at a maximum 50-cm FOV. Similarly, the condition is well satisfied for mammography at realistic clinical detector pixel sizes, $\sim 10\text{ }\mu\text{m}$, though it may drop below the $N_F = 10$ threshold for single-micron scales and thick samples. The condition is not clearly satisfied for micro-CT, especially with sub-micron pixels or centimeter-thick samples.

Though a useful first estimate, the Fresnel number does not capture the full picture. Eq. 7.1 provides a simple criterion based on a single characteristic feature size and length scale, but real objects contain a range of material interfaces and structure sizes that may be larger than the pixel size. Beyond the object, the exit wave evolves further as it propagates through free space, and detector nonidealities such as blur and noise influence whether internal-object refractive effects actually convert to measured intensity changes. Post-processing algorithms such as phase retrieval and tomographic reconstruction can also affect the image properties. While the Fresnel number offers useful orientation, direct comparisons between PA and more general wave optics models are still necessary for robust assessment of different samples.

Previous work from Morgan et al. investigated PA versus an exact solution for PB-XPCI of a cylinder edge as a proxy for biological specimens [192], additionally developing another analytical metric for PA validity that accounts for additional acquisition parameters such as the cylinder’s phase-shifting component δ and propagation distance [193]. A key finding was the significance of propagation distance: experimentally measured 25-keV synchrotron




Figure 7.1: Fresnel number heatmaps in parameter ranges representing three X-ray imaging modalities: micro-CT, mammography, and clinical CT.

images with half-micron pixel-size detector were found to more closely match PA simulations as distance increased up to 100 cm. More precisely, they found PA underestimated fringe magnitude over distances up to the scale of the sample thickness. Nonetheless, especially at more typical PB-XPCI scales using larger propagation distances, PA very closely matched experimental observation and improved simulation computation time relative to the exact solution from hours to under a second.

Other models may offer an improved balance of generality and computational efficiency. One potential alternative is the multi-slice (MS) method, in which the object is first divided into a series of sufficiently thin axial slices, then the incident wave field's evolution is modeled using sequential thin-object transmission functions and free-space propagation operators over each slice width [4]. The MS model is also an approximation, neglecting backscattered waves and making the paraxial assumption, but it is more general than PA because it does not collapse the entire object into a single projected transmission function. Still, it requires additional propagation operations scaling with the number of slices, increasing computation time and inverse problem complexity.

Yet, MS-based inverse solutions are already established for X-ray ptychography, a computational microscopy technique that utilizes overlapping coherent diffraction pattern mea-

surements to perform phase retrieval [194–197]. Like PB-XPCI, conventional ptychographic reconstruction often models the object using a single projected transmission function. Since ptychography is often performed at nanometer-scale resolution, the depth-of-field can quickly become short relative to the sample thickness, motivating earlier development of more accurate forward models and corresponding inverse solutions [4]. The MS approximation offered a practical solution accounting for depth-dependent refractive effects [165]. Recently, an MS-based phase retrieval method has been proposed compatible with both ptychography and propagation-based nanotomography, using automatic differentiation to perform iterative optimization while accounting for internal refractive effects without requiring an analytical inverse [198].

Though this past work focuses on smaller scales than micro-CT or high-resolution clinical imaging, it motivates benchmarking PA performance against MS forward models when the assumptions of PA may be violated. Thus, the specific goals of this chapter are to explore the regime of PA accuracy, to identify when depth-dependent refractive effects become apparent in the MS model, and to assess whether these differences meaningfully affect conventional phase retrieval and CT reconstruction. This lays the foundation for a parallel comparison of PA- and MS-based iterative phase retrieval methods in the following chapter. We build upon prior work by directly comparing the two models and assessing them at different size scales.

7.3 Methods

Noiseless monochromatic PB-XPCI simulations were performed using both PA and MS forward models. The simulations imaged a variety of simple and detailed phantoms generated at sub-micron spatial resolution with either high- or low-attenuation contrast materials. All simulations were defined using the wave-optics library `chromatix` and assumed a monochromatic source with an ideal photon-counting detector [183]. Three primary simulation studies




Figure 7.2: Schematic of the rod phantom.

were performed: (1) line profiles through a simple rod phantom, (2) full-FOV projection images of an image-quality phantom, and (3) micro-CT scans of stained and unstained zebrafish phantoms.

7.3.1 Rod phantom line profiles

The first phantom was a simple image quality phantom comprising a 1-cm thick water slab with ten embedded rods (Fig. 7.2). Each rod had a 10- μm radius, and they were equally distributed throughout the slab's depth in the direction of wave propagation. This geometry was chosen as a simple form of depth information that might induce detectable internal phase shifts neglected by the PA model.

To consider the effect of material composition, rods were defined as either aluminum or PMMA. Fig. 7.3 shows the refractive index components of the phantom materials as a function of energy. Aluminum has the highest absorptive and phase-shifting components of all materials, likely owing to its relatively high density. Compared to water, PMMA is more weakly absorptive (smaller β) but more strongly phase-shifting (larger δ).

Single-projection PB-XPCI simulations were performed using a 14-keV incident beam and a range of spatial resolutions ($\Delta x = 0.05\text{--}5.0\ \mu\text{m}$), propagation distances ($R = 0$ or 20 mm), and number of MS slices ($N_z = 2\text{--}100$). Central line profiles were compared in the raw images produced by the two forward models.




Figure 7.3: Refractive index components for materials used in the rod phantom simulations.

7.3.2 QA phantom projections

To provide more transverse spatial detail, mammography quality assurance (QA)-inspired “clutter phantoms” of PMMA spheres were generated at two image scales (Fig. 7.4) [199]. The first scale, “mammo-clutter,” was a 5-cm³ cube randomly populated with non-overlapping spheres of radii between 0.3 and 1.5 cm, targeting a fill ratio of 60% of the total volume. A 1-cm ROI was chosen for efficient imaging simulations. The second scale, “micro-clutter,” was a 1-mm² transverse area with half-cm depth, generated using the same random fill algorithm with sphere radii 10 μ m–1 mm and a target fill ratio of 50%. The anisotropic volume shape and slightly lower fill target were chosen for the micro-clutter phantom due to the large number of non-overlapping micron-sized spheres needed to achieve the target, which induced significant computation time. For phantom generation, the spheres were allowed to spill outside of the imaged volumes.

Three detectors were modeled representing real detectors used at different imaging scales. First, a 55- μ m pixel detector was based on the Medipix family of energy-discriminating PCDs, a promising detector for spectral XPCI, with a 21- μ m FWHM Gaussian PSF [56, 61, 200, 201]. The second detector had 8- μ m pixel width, inspired by the inCiTe 3D X-ray microscope




Figure 7.4: Projected thickness of the two randomly generated PMMA-sphere clutter phantoms. On the left, the larger “mammo-clutter” phantom was generated over a full 5-cm² FOV, and a 1-cm² ROI (red box) was chosen for the simulation. On the right, the full 1-mm² FOV was used for simulations of the smaller “micro-clutter” phantom.

(trademark KA Imaging), with a 6- μ m FWHM Gaussian PSF. Last, a 0.5- μ m pixel detector was modeled after a micron-resolution detector used for half-centimeter FOV micro-CT imaging [169]. This detector model used a Lorentzian PSF with 1.1- μ m FWHM, reflecting the lens optics system used to achieve its sub-micron effective pixel pitch [202].

The 55- μ m and 8- μ m detectors were used to simulate images of the larger mammo-clutter phantom at $R = 0$ or 100 cm, while the 8- μ m and 0.5- μ m detectors were used to simulate images of the smaller micro-clutter phantom at $R = 0$ or 10 cm. The nonzero propagation distance was chosen in proportion to the object size ($20 \times$ max phantom thickness). All simulations were noiseless with monochromatic 10-keV beam. The MS model used $N_z = 50$ slices, with slice width $\Delta z = 1$ mm for the larger phantom and $\Delta z = 0.1$ mm for the smaller phantom.

Single-distance phase retrieval (PR) was performed on the $R > 0$ images to see whether the intensity differences in the raw MS- and PA-simulated data actually translated to appreciable differences. The PR algorithm used was the widely-implemented homogeneous object method, with an assumed material of PMMA [24].




Figure 7.5: Zebrafish phantom at the original 1.6- μm resolution (left) and after the smooth upsampling operation to achieve 0.2- μm resolution for accurate PB-XPCI micro-CT simulations.

7.3.3 Zebrafish micro-CT

As an experimental analogue, micro-CT scans were simulated using a structurally rich zebrafish phantom. The original phantom was created by segmenting a reconstructed synchrotron micro-CT scan of a metal-stained zebrafish with 1.6- μm voxel size. The two metals, osmium tetroxide and uranyl acetate, approximately targeted cell lipids and nuclei, respectively, allowing segmentation into four materials (air, water, tissue, adipose).

Since the native 1.6- μm voxel size was too coarse to simulate sub-micron wave interference effects, a custom smooth upsampling algorithm was developed. Each material mask was first upsampled individually to 0.2- μm pixels ($8\times$) using nearest-neighbor interpolation. A small




Figure 7.6: Refractive index components for the zebrafish phantom simulations. The metal stains osmium tetroxide and uranyl acetate reflect a mixture of 1% stain in water.

Gaussian blur was applied ($\sigma = 8$ pixels), and the result was re-binarized into 1s (material present) and 0s (material absent) using a normalized threshold of 0.35, which was manually selected to preserve fine structures without excess blooming. Overlapping material masks were assigned to a single material by choosing the material with the highest blurred intensity, with any remaining empty space assigned to vacuum, and the individual material masks were recombined into a single upsampled phantom with smooth material interfaces (Fig. 7.5).

Simulations were performed for both metal-stained (24 keV) and unstained (14 keV) material compositions, with refractive index components in Fig. 7.6. For the stained simulations, osmium tetroxide and uranyl acetate mixtures replaced adipose and tissue, respectively. The stained simulations also used a more energetic source (24 keV instead of 14 keV) to mitigate photon starvation from the higher attenuation.

To eliminate edge artifacts, the phantom was padded with 1000 empty vacuum pixels on all sides, yielding a final 13130×13130 zebrafish array of $0.2\text{-}\mu\text{m}$ pixels. Projection images were taken at 6874 rotation angles equally spaced over 180 degrees to satisfy Nyquist sampling criterion for parallel-beam CT [71]. The free-space propagation distance was either 0 or 20 mm, and the detector was defined with pixel size $\Delta x = 0.6\text{ }\mu\text{m}$ and ideal focal spot.

Lower-resolution images were also simulated by block-averaging the ideal sinograms to a pixel size $\Delta x = 2.4 \text{ } \mu\text{m}$.

Single-distance, homogeneous material PR was performed in the sinogram domain with an assumed material of water, yielding effective thickness sinograms [24]. For the $R = 20$ mm sinograms, this is a true form of phase retrieval, while for the $R = 0$ contact sinograms, this is equivalent to taking the log-intensity measurements and normalizing by the linear attenuation coefficient of water. The final images were reconstructed from the PR sinograms using FBP with a matrix size of 4096 and FOV of 2.5 mm [7].

7.4 Results and Discussion

7.4.1 Rod phantom line profiles

Figs. 7.7 and 7.8 show a region of the PA- and MS-simulated rod phantom line profiles for aluminum and PMMA rod materials, respectively. The line profile region shown is centered on the rod-to-water background material interface. In the contact image ($R = 0$ mm), PA shows no fringes due to the lack of any propagation, resulting in a large difference relative to MS. After a short propagation downstream ($R = 20$ mm, twice the phantom thickness), there is still a subtle difference in the line profiles. This convergence matches expectation from previous work examining PA in modeling PB-XPCI of cylinders, with PA converging to exact solutions and experimental measurements as R increases beyond the scale of the object thickness [192, 193]. Interestingly, finite detector sampling effects may overemphasize differences between PA and MS at certain distances due to apparent transverse shifts in the fringe peak and trough locations. This effect is apparent in the 20-mm aluminum rod line profiles around detector coordinate 11 μm and 12.5 μm .

Fig. 7.9 explores this finite detector effect by looking at pixel-wise root-mean-square-error (RMSE) between MS- and PA-simulated line profiles as a function of detector pixel size. As




Figure 7.7: Simulated aluminum rod phantom line profiles before (dashed line) and after (marked solid line) resampling at the detector pixel size $\Delta x = 0.5 \mu\text{m}$.




Figure 7.8: Simulated PMMA rod phantom line profiles before (dashed line) and after (marked solid line) resampling at the detector pixel size $\Delta x = 0.5 \mu\text{m}$.




Figure 7.9: RMSE between MS and PA rod phantom simulations as a function of detector pixel size Δx .




Figure 7.10: RMSE between 100-slice MS line profiles and MS simulations with just N_z , illustrating convergence of MS simulations as $N_z \rightarrow \infty$.

expected according to the Fresnel number criterion (Eq. 7.1), PA shows better agreement with MS at larger pixel sizes. The difference approaches a plateau above $\Delta x \approx 2 \mu\text{m}$ for both PMMA and aluminum. The downward trend was not monotonic, with small bumps possibly reflecting the finite sampling effect when MS and PA fringes fall near the edge of differently spatially distributed pixels, an inevitability when sweeping over a range of pixel sizes. Additionally, RMSE was roughly an order of magnitude larger for aluminum relative to PMMA, likely due to its larger δ increasing its induced phase shifts.

Fig. 7.10 assesses the number of slices needed for MS simulations of the rod phantom to

Table 7.2: Average and maximum difference ($|PA - MS|$) in normalized raw intensity images ($R = 0$ or $10/100$ cm) and phase retrieval “PR” images ($R = 10/100$ cm) for the different detector and phantom sizes. The mammo-clutter phantom was 5-cm thick while the micro-clutter phantom was only 0.5 cm thick.

Phantom	Detector Δx	Images	Avg. diff.	Max. diff.
mammo-clutter	55 μm	$R = 0$ cm	0.00021	0.0056
		$R = 100$ cm	0.00017	0.0047
		PR	3.7 μm	36.8 μm
	8 μm	$R = 0$ cm	0.00051	0.093
		$R = 100$ cm	0.00044	0.075
		PR	3.9 μm	43.2 μm
micro-clutter	8 μm	$R = 0$ cm	0.0022	0.015
		$R = 10$ cm	0.0017	0.016
		PR	4.2 μm	15.8 μm
	0.5 μm	$R = 0$ cm	0.010	0.16
		$R = 10$ cm	0.011	0.27
		PR	4.5 μm	19.2 μm

converge with 0.1- μm pixel sampling, comparing MS as a function of number of object slices N_z to MS_{100} (the 100-slice MS simulation taken as an estimate of the true line profile). For both materials, RMSE drops below 10% of the maximum around $N_z = 20$ (21 for PMMA and 18 for aluminum), after which slowly plateaus. This corresponds to a slice width of 0.5 mm, suggesting this is a sufficient thickness for satisfying PA at 1-cm thickness with 10- μm cylinders. Notably, the Fresnel number is very different depending on whether the cylinder size or detector resolution is taken as the feature size of interest ($N_F = 2260$ for cylinder size versus $N_F = 0.226$ for detector pixel over 0.5-mm thickness). A more adequate prediction of PA sufficiency might use a few detector pixels as the feature size.

7.4.2 QA phantom projections

Table 7.2 summarizes the average and maximum differences in the PA and MS images of the clutter phantoms imaged with the different detector scales. For both phantoms, the difference between forward models was much larger for the smaller detector pixel size. In the larger mammo-clutter phantom, the PA – MS difference was also slightly larger in the $R = 0$ mm exit waves relative to the $R = 100$ cm propagated waves. This difference with propagation distances was not obvious in the micro-clutter phantom images. Overall, differences in the PR thickness estimates were very small, on the order of ~ 10 μm , which is less than 1% of the both phantom thicknesses (5 or 0.5 cm). However, the difference was relatively larger for the micro-clutter phantom, since it had a smaller total thickness.

The full raw intensity, retrieved thickness, and PA minus MS difference images for both phantoms and detectors can be seen in Figs. 7.11–7.14. In the intensity images, differences between the MS and PA simulations are concentrated at the object edges, as expected due to PA underestimating fringe amplitude. In the mammo-clutter PR images, this translated to slightly lower thickness estimates from PA on the “thick” side of material interfaces and higher estimates on the other side. This also seems true in the micro-clutter PR images, though the greater detail of edge structures makes it more challenging to identify individual sphere edges.

7.4.3 Zebrafish micro-CT

Fig. 7.15 shows the ideal stained and unstained absorption-only reconstructions. A small ROI (25% of the full FOV) is also shown to illustrate finer details, which will be the focus of the remaining images examined. The metal stains considerably improve the contrast of cell lipids and nuclei on the water background, reflecting their higher attenuation.

Figs. 7.16 and 7.17 show the reconstructed contact image ROIs ($R = 0$) for the stained and unstained zebrafish, respectively, comparing MS and PA at the sub-micron and larger




Figure 7.11: 1-cm² FOV mammo-clutter phantom simulations with the 55-μm pixel detector.




Figure 7.12: 1-cm² FOV mammo-clutter phantom simulations with the 8-μm pixel detector.




Figure 7.13: 1-mm² FOV micro-clutter phantom simulations with the 8- μm pixel detector.




Figure 7.14: 1-mm² FOV micro-clutter phantom simulations with the 0.5- μ m pixel detector.




Figure 7.15: Ideal reconstructions of the stained and unstained zebrafish phantom to provide baseline reference for their native attenuation contrast. The ground truth phantom is shown on the left for reference with the same four-material colormap of Fig. 7.5. The full FOV is 2.5 mm with matrix size 4096×4096 . The bottom row shows the detailed structures within the boxed 0.5-mm ROI.




Figure 7.16: Stained zebrafish contact image ($R = 0$), simulated with either PA or MS, with sub-micron detector pixel size (top row) or larger binned detector pixels (bottom row). 0.5-mm ROI of the full reconstructed matrix.




Figure 7.17: Unstained zebrafish contact image ($R = 0$), simulated with either PA or MS, with sub-micron detector pixel size (top row) or larger binned detector pixels (bottom row). 0.5-mm ROI of the full reconstructed matrix.




Figure 7.18: Stained zebrafish images with $R = 20$ mm, raw intensity measurements (top row) and post-phase retrieval (bottom row), simulated with either PA or MS. 0.5-mm ROI of the full reconstructions with sub-micron detector pixels.

binned detector pixel sizes. In the stained images, the MS simulation reveals subtle edge enhancement of the water/air interfaces, which is invisible to PA and lost with the large detector pixel size. The overall PA–MS difference is much larger with the sub-micron detector pixels. Stain-to-stain material interfaces did not induce detectable edge enhancement differences between PA and MS forward models at either detector resolution. In the unstained images, the maximum magnitude of the PA–MS difference images is roughly three times as large as it was for the stained difference images. Further, a close look at the sub-micron detector images reveals observable differences between PA and MS simulations at the tissue-to-adipose material interfaces (in place of the stain-to-stain interfaces). This is a consequence of the metal stains' smaller δ relative to water, tissue, and adipose (Fig. 7.6).

Figs. 7.18 and 7.19 show the subsequent propagated $R = 20$ mm images for the stained




Figure 7.19: Stained zebrafish images with $R = 20$ mm, raw intensity measurements (top row) and post-phase retrieval (bottom row), simulated with either PA or MS. 0.5-mm ROI of the full reconstructions with sub-micron detector pixels.

and unstained zebrafish, respectively, reconstructed using the sub-micron detector pixel data. Like in the contact images, the magnitude of the PA–MS difference is much larger for the unstained image. This is especially pronounced in the PR images considering the relative difference; the PR-estimated thickness of the metal stains is over $10\times$ larger than the estimated thickness of the soft tissues, but the PA–MS difference is smaller for the metal stains. Also in the PR images, PA simulations yield a lower recovered thickness on the “thick” side of material edges, matching the observation in the projected image domain with the QA phantom simulations.

Comparing stained and unstained PR images, the unstained images offer improved differentiation of cell nuclei. The addition of metal stains causes large blurring artifacts in the PR images, resulting in nuclei and lipids appearing as large indistinguishable clusters. This may be due to the homogeneous material assumption made in the phase retrieval algorithm. An assumed single material of water is generally not an accurate characterization of the heterogeneous zebrafish, but it is more reasonable for an object comprising adipose and tissue than it is for highly-attenuating metals. Thus, while the raw metal-stained $R=20$ mm images show prominent material-interface delineation, even between metal stains, this is not necessarily translated to a better PR image.

7.5 Conclusion

In sum, this PB-XPCI simulation comparison provides a more nuanced analysis of the accuracy of the projection approximation (PA). The results indicate that object thickness and feature size alone are not sufficient to determine whether PA is an adequate forward model. Although the Fresnel number provides a useful first indicator, PA–MS differences also depend on material properties, propagation distance, detector resolution, and any specific post-processing methods applied to the raw data. This is consistent with previous work, which we extend to include varied phantoms, phase retrieval, and CT reconstruction

[192, 193].

In the rod phantom simulations, PA–MS differences were most apparent for sub-micron pixel sizes and higher-contrast materials, where depth-dependent refraction within the object produced measurable changes in the intensity profiles. The QA phantom simulations showed that these raw-data differences can be substantially reduced by large detector pixels and blur, as well as approximate single-distance phase retrieval. Finally, the zebrafish CT simulations reflected similar findings in context of a structurally complex biological phantom, especially through comparison of the stained and unstained material compositions.

Together, these results affirm that PA is typically a reasonable forward model at current X-ray imaging scales. However, for emerging high-resolution/wide-FOV micro-CT, PA errors can become apparent in the raw data domain even when they are less apparent after phase retrieval, suggesting it may be important for quantitative PR methods to account for internal refractive effects. This motivates continued development and benchmarking of more general forward models and inverse solutions in regimes where internal phase shifts might be non-negligible. We explore this in the next chapter with an iterative phase retrieval framework.

CHAPTER 8

ITERATIVE MATERIAL DECOMPOSITION FOR DUAL-ENERGY PB-XPCI

8.1 Introduction

Most PB-XPCI phase retrieval methods rely on the projection approximation (PA): the assumption that refractive effects within the object are negligible [22, 24, 26, 60, 203, 204]. As concluded in the previous chapter, PA offers an appealing forward model due to its conceptual simplicity, computational convenience, and accuracy at most conventional X-ray imaging scales. However, for thicker objects, higher-resolution imaging, and strongly phase-shifting materials, wave interference within the object can become a significant factor contributing to image quality, violating the key assumption of PA. This is not necessarily bad for imaging—the high resolution and increased edge enhancement in such scenarios can provide useful information about the sample. In these regimes, a multi-slice (MS) model can provide a more complete description of the relevant imaging physics while remaining more computationally tractable than a full wave equation solution [165].

For quantitative phase retrieval, accurate forward modeling is essential. If the chosen model does not capture the relevant imaging physics, the inverse problem being solved may differ from the true physics underlying the measured data, and resulting images may suffer from artifacts or reduced accuracy. This issue is increasingly important for wide-field micro-CT and nanotomography, where advances in high-resolution detectors and coherent X-ray sources render internal refractive effects more relevant considerations [166, 169].

Here, we present an MS-based phase retrieval approach for dual-energy CT material decomposition with PB-XPCI. Using automatic differentiation (AD), the solution can be efficiently implemented without need to derive closed-form gradients of the forward model. This also allows the PA and MS models to be interchanged within the same optimization

framework, isolating the effect of the forward model on convergence and reconstruction quality. We use this framework to assess when MS-based phase retrieval can improve material decomposition performance relative to PA-based phase retrieval, and whether any improvement justifies the additional computational complexity.¹

8.2 Background

Iterative reconstruction (IR) methods have long been established in the clinical CT reconstruction literature [7, 90, 205]. In fact, the first CT prototypes from Cormack and Hounsfield utilized a type of iterative reconstruction, algebraic reconstruction technique (ART), years before Shepp and Logan’s introduction of the efficient filtered back-projection (FBP) method [30, 32, 87, 90]. Yet, iterative methods suffer from a high computational cost and longer run time to produce the final images.

It was not until the 2010s that IR became practically implemented, largely spurred by interest in achieving low-dose CT and thus needing clinically acceptable reconstruction methods capable of handling noisier raw data [7, 206]. By incorporating known statistical models, hybrid IR/analytical update schemes, and improved computational power, IR became compatible with routine clinical workflow [207]. While FBP remains an important analytical standard, modern clinical CT reconstruction commonly incorporates some form of IR, especially through vendor-specific algorithms [208].

A similar interest in iterative methods has arisen in the context of PB-XPCI to achieve quantitative phase retrieval. Analytical phase retrieval methods, such as the widely used Paganin method [24], provide rapid solutions under simplifying approximations, most notably that the object is homogeneous or has a known relationship between its phase-shifting and absorptive components [22]. When the approximations are valid, these methods can be both accurate and convenient; but if the assumptions are violated, the images can suffer from

1. This chapter is related to work also shared as conference proceedings [179, 180].

artifacts and inaccuracies. An example of this was seen in the previous chapter, applying Paganin’s method to metal-stained and unstained zebrafish micro-CT scans, when the presence of highly attenuating stains on the soft tissue background violated the single-material assumption and resulted in severe blurring of fine structures. The necessary assumptions for phase retrieval can be relaxed somewhat with multiple independent measurements, which in principle provide the necessary information to solve for the two independent contrast mechanisms (phase and absorption). Even so, a general forward model for PB-XPCI does not necessarily have an analytical inverse solution, especially accounting for more realistic physics elements such as repeat diffraction events within the object.

Iterative phase retrieval is thus an interesting direction for quantitative PB-XPCI, opening the possibility to include more model-based elements when closed-form inversion is insufficient. Model-based IR methods have been proposed for PB-XPCI CT, incorporating a simplified TIE-based forward model into the CT reconstruction problem [209, 210]. Other optimization approaches have been developed implementing more accurate Fresnel propagation forward models and multiple distance measurements to further relax the model assumptions [204]. For spectral PB-XPCI specifically, linear iterative phase retrieval (LIPR) has been proposed for single-distance, near-field, dual-energy imaging, based on a simplified Alvarez–Macovski decomposition model [57]. Other methods have explored techniques appropriate for energy-resolved PCD measurements [61]. These approaches broadly indicate promise for iterative phase retrieval but also highlight the continued implementation of simplifying approximations in forward projections and update steps, improving algorithm stability and speed at the cost of generalizability and possible accuracy.

Today, computational advances such as GPU acceleration and AD make it more feasible to efficiently implement iterative phase retrieval with more general forward models [183]. In particular, AD makes it possible to calculate gradients for numerical models by applying the chain rule through the elementary operations making up each computer function, bypassing

the need for a closed-form inverse or update rule. This is especially useful for PB-XPCI, with forward model operations that are challenging to invert analytically, including Fresnel propagation, nonideal detector response, and repeated diffraction events within the sample. Recently, an AD-based phase retrieval method was successfully developed for nanotomography using an MS model for wave evolution through the object, demonstrating the feasibility of iterative phase retrieval with a more general model [166]. However, as we saw in the previous chapter, the relative benefit of MS forward modeling over the simpler PA model is largely application-dependent, and it remains open to explore whether the additional computational expense and potential instability is necessary at larger imaging scales.

To this end, this chapter compares AD-enabled PA- and MS-based iterative phase retrieval for dual-energy PB-XPCI material decomposition. By posing phase retrieval as an optimization problem, our framework bypasses the need for a closed-form inverse solution and minimizes reliance on model approximations like the linearized TIE. With AD, the update steps also do not need to be manually derived for each forward model. This allows the effect of MS modeling to be isolated within the same reconstruction framework, with the goal of identifying when MS improves performance enough to justify its computational costs.

8.3 Methods

8.3.1 Forward models

With the context of dual-energy imaging and two-material decomposition, let $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2)$ denote the object, where $\mathbf{x}_j$ is the voxel-wise volume fraction of the j^{th} basis material. Then

for energy E , the spatially-varying refractive index components can be parameterized as

$$\begin{aligned} A(\mathbf{r}, E) &= \delta(\mathbf{r}, E) - i\beta(\mathbf{r}, E) \\ \Rightarrow A(\mathbf{x}, E) &= \sum_{j=1}^2 \mathbf{x}_j [\delta_j(E) - i\beta_j(E)], \end{aligned} \quad (8.1)$$

where $\delta_j(E)$ and $\beta_j(E)$ are the phase-shifting and absorptive components of the complex refractive index of basis material j , respectively.

The CT acquisition geometry is shown in Fig. 8.1. For each CT view angle θ , the coordinates of the rotated transverse plane are

$$\xi = x \cos \theta + y \sin \theta, \quad (8.2)$$

$$\ell = -x \sin \theta + y \cos \theta, \quad (8.3)$$

where ℓ is the direction of wave propagation through the object and ξ is the detector coordinate. The detector was defined as a plane perpendicular to the axis of propagation with finite element size $\Delta\xi$.

Using the PA model, the wave's modulation through the object is defined according to the line integrals through its refractive index components, yielding monochromatic sinogram measurements

$$f^{(\text{PA})}(\mathbf{x})(E, \xi, \theta) = S(E) \left| \mathcal{D}_R(E) \exp \left[-ik(E) \int_0^T d\ell A(\mathbf{x}, E) \right] \right|^2, \quad (8.4)$$

where $S(E)$ is the incident spectrum in units of number of photons per detector element, $\mathcal{D}_R(E)$ is the energy-dependent free-space propagator, R is the propagation distance, $k(E)$ is the wavenumber, and T is the object thickness in the direction of propagation. The free-space propagator was defined as the Fresnel operator convolution (Eq. 5.21). For the dual-monoenergetic simulations in this chapter, $S(E)$ was set equal to a constant, such that

the total incident intensity of each scan was equal.

For the multi-slice (MS) model, the object was divided into N_s slices of thickness $\Delta\ell = T/N_s$. The transmission function for each slice i was defined as

$$\mathcal{T}_i(\mathbf{x}, E, \xi, \theta) = \exp \left[-ik(E) \int_{(i-1)\Delta\ell}^{i\Delta\ell} A(\mathbf{x}, E) d\ell \right]. \quad (8.5)$$

Starting with a unit incident plane wave $\psi_0(\xi, \theta, E) = 1$, for each CT angle, the wave field was propagated through the object slices according to

$$\psi_i(\mathbf{x}, E, \xi, \theta) = \mathcal{D}_{\Delta\ell}(E) \mathcal{T}_i(\mathbf{x}, E, \xi, \theta) \psi_{i-1}(\mathbf{x}, E, \xi, \theta), \quad i = 1, \dots, N_s. \quad (8.6)$$

The final sinogram intensity measurements are then

$$f^{(\text{MS})}(\mathbf{x})(E, \xi, \theta) = S(E) |\mathcal{D}_R(E) \psi_{N_s}(\mathbf{x}, E, \xi, \theta)|^2. \quad (8.7)$$

8.3.2 *Material decomposition phase retrieval (MDPR)*

For spectral XPCI, material decomposition might be considered a type of phase retrieval problem, since it requires converting raw intensity images into a new quantitative basis. Once the new basis is acquired, it is trivial to convert between assumed material images and phase/absorption images using known basis function relations [55], much like synthesis of virtual monochromatic images. Thus, we term the problem *material decomposition phase retrieval* (MDPR).

MDPR was posed as an optimization problem,

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}), \quad (8.8)$$

with loss function

$$\mathcal{L}(\mathbf{x}) = \|f(\mathbf{x}) - \mathbf{y}\|_2^2 + \gamma r(\mathbf{x}) \quad (8.9)$$

comprising an L2 norm fidelity term and an additional regularization weighted by γ , where $\mathbf{y}$ is the measured data.

Common regularization terms in CT imaging include smoothness and sparsity constraints, such as total variation (TV) and the L1 norm, respectively [90]. While these can improve image quality, they generally require time-consuming manual tuning of regularization weights. Here, with the primary goal of comparing PA and MS performance, we chose a minimalist regularization strategy to reduce need for hand-tuning, defining $r(\mathbf{x})$ as a material exclusion penalty equal to the L1 norm of the material-product image,

$$r(\mathbf{x}) = \|\mathbf{x}_1 \mathbf{x}_2\|_1. \quad (8.10)$$

This is simply the pixel-wise product of the two basis material images, which applies a penalty only when both materials are nonzero and therefore encourages a single material per pixel. The term was found to be necessary to encourage material differentiation in the absence of an informed initial guess of the object. The penalty is a fair prior for CT imaging in the context of basis material decomposition, which spatially separates individual materials through tomographic reconstruction, so long as partial volume averaging effects are minimal. The regularization weight γ was set to either 0 for idealized simulations or tuned to the task using a coarse parameter sweep.

The loss function was minimized using the Adaptive Moment Estimation (ADAM) optimizer, implemented with `optax` [211]. ADAM is a first-order optimizer that scales voxel-wise updates using running estimates of the gradient and squared gradient magnitude. This can improve stability when different basis material images have different gradient scales due

to their varying attenuation and phase sensitivities. Here, ADAM was implemented with fixed learning rate $\alpha = 10^{-2}$, first- and second-moment scaling parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and small constant to avoid dividing by zero $\epsilon = 10^{-8}$. The convergence criterion was defined as when the change in loss over 50 iterations was less than some threshold $\Delta\mathcal{L}_t$, which was set depending on the task.

The solution space was constrained to $\mathbf{x} \geq 0$. A finite support constraint was also applied given an assumed known external object structure, in this case a cylinder. This constraint was reported as a key element of successful MS-based phase retrieval method for PB-XPCI nanotomography [166]. It is also straightforward to compute the finite support constraint function empirically using an approximate analytical reconstruction and small intensity threshold.

All reconstructions were performed using a single NVIDIA Tesla V100 32-GB GPU, making memory the primary constraint on phantom size, angular sampling, detector resolution, and MS slice count in this initial study. Extension to thicker samples, larger FOVs, or more densely sampled CT acquisitions at sub-micron resolution will require either larger-memory GPUs or multi-GPU computation strategies. Prior work has explored distributed-memory MS propagation using array tiling [212], and conventional tomographic strategies such as batching or subdividing projection angles may also help reduce memory demand [7]. Here, we instead focused on smaller-scale simulations that could be reconstructed on a single GPU, enabling a broad parameter sweep over acquisition and reconstruction conditions. Future work can use these findings to guide larger-scale MS reconstruction studies, including high-resolution wide-field micro-CT [169].

8.3.3 Simulations

A simple image quality phantom was generated (Fig. 8.1) comprising a PMMA cylinder spanning 90% of the FOV with spherical aluminum inserts. The phantom was discretized




Figure 8.1: Schematic of the CT acquisition geometry and simulation phantom. The static coordinates (x, y) define the stationary plane of the phantom, while the rotating coordinates (ξ, ℓ) change with the view angle according to Eqs. 8.2–8.3.

into a 4096×4096 array of pixels with width Δx_{sim} . Measured sinogram data was generated using a “full” MS forward model, which applied intermediate transmission steps to all 4096 slices at each CT view angle. To sufficiently model sub-pixel wave interference, the detector pixel size was downsampled by s pixels relative to the desired simulation resolution using block averaging, i.e. $\Delta \xi = s \Delta x_{\text{sim}}$. The specific scale and pixel size varied depending on the specific experiment. Noise was optionally added to each monochromatic sinogram using a Poisson counting distribution based on the incident intensity $S(E)$.

The dual-energy pair (6 keV, 10 keV) was selected using the line-profile CRLB detailed in Ch. 6 for a 1-mm PMMA/0.1-mm Aluminum basis with 0.2- μm pixels and 1-mm propagation distance to the detector. The energies separately minimized the CRLB of the basis materials. For consistency, the same energies were fixed across all simulations, regardless of phantom scaling and system resolution. A separate experiment might explore individually-optimized energy pairs.

MDPR and CT reconstruction were performed simultaneously with the optimization

framework detailed above. The reconstruction matrix was defined as a square array of pixels with width matching the size of the detector elements, $\Delta x = \Delta \xi$. MDPR was independently repeated with either the PA or MS forward model used to compute the fidelity term of the loss function. For efficient optimization, the MS model used $N_s = 8$, and neither model upsampled the phantom during the forward projection step.

Simulation parameters were chosen to represent the orders of magnitude most relevant to high-resolution micro-CT and nanotomography. All CT simulations acquired 180 view angles equally spaced 1° apart. While this is relatively coarsely sampled angularly for the high-resolution detectors, it was mainly chosen to reduce the total memory required while maintaining larger reconstruction matrix sizes.

The first set of experiments considered MDPR performance under ideal conditions: noiseless data, unregularized optimization ($\gamma = 0$), and highly informed initialization, with $\mathbf{x}$ equal to the ground truth (GT) phantom perturbed by a small 2D Gaussian blur (FWHM = 2% of the FOV), called the “blurred ground truth” (BGT). MS- and PA-based MDPR performance was measured as a function of detector resolution and object-to-detector propagation distance using a small convergence threshold $\Delta \mathcal{L}_t = 10^{-6}$. For the resolution study, detector pixel size ranged from $\Delta x = 20$ nm to 2 μm , with the phantom scaled to a $205\text{-}\mu\text{m}$ FOV and a fixed propagation distance of half the FOV. For the object-to-detector distance study, the pixel size was fixed at $0.2\text{-}\mu\text{m}$, and the propagation distance was scaled from 0 to $4\times$ the FOV.

Idealized MDPR performance was then assessed as a function of noise level. Noise was added to the ideal $0.2\text{-}\mu\text{m}$ detector pixel, half-FOV propagation distance sinograms by varying incident number of photons $S(E)$ from 10^2 to 10^7 , with fewer photons corresponding to noisier images. Higher noise may obfuscate appreciable differences in the MS and PA forward projections. MDPR was performed using either model on each set of noisy data.

Finally, MDPR performance and convergence were examined for a more challenging opti-




Figure 8.2: SSIM and $1 - \text{RMSE}$ in each basis image as a function of detector pixel size Δx using either PA- or MS-based MDPR.

mization task. The previously generated noisy sinogram with $S(E) = 10^3$ was reconstructed using an uninformed initial guess of the object, $\mathbf{x}$ equal to a uniform array of 1s for both objects. A regularization weight $\gamma = 0.02$ was chosen to minimize the fidelity term after a coarse parameter search. The loss value was logged for each iteration, and MDPR performance was compared between forward models and initializations.

For quantitative comparison, MDPR accuracy was benchmarked against the GT phantom downsampled to match the reconstruction matrix resolution (pixel size Δx). Two main figures of merit were used for each basis material image: the structural-similarity index measure (SSIM) to capture perceived resemblance of the GT, and the root-mean-square error (RMSE) to quantify numerical accuracy relative to the GT [213].




Figure 8.3: PA, MS, and ground truth (GT) PMMA reconstructions at different spatial resolutions (phantom thickness = 205 μm).




Figure 8.4: An example line profile through the 10-keV sinogram generated at different detector pixel sizes with simulation upsampling (top row; full MS model) and without upsampling (bottom row; optimization MS model).

8.4 Results and Discussion

8.4.1 Spatial resolution

Fig. 8.2 shows MDPR performance as a function of pixel size. RMSE is plotted as $1 - \text{RMSE}$ so that trends match SSIM (higher indicates better performance). For both PA and MS models, performance generally improves with smaller pixel sizes, reflecting the greater phase-contrast information content at material edges in the raw sinogram data. MS performed better than PA at the two smallest pixel sizes ($\Delta x \leq 0.4 \mu\text{m}$), with performance converging above $1 \mu\text{m}$. This matches the expectation that high resolution is necessary for internal refractive effects to significantly contribute to detected intensity measurements, making an MS model beneficial. As pixel size increased, the unregularized reconstructions suffered from an increase in streaking artifacts, despite a lack of Poisson noise in the raw sinogram data (Fig. 8.3). The effect was more severe for PA, which is especially noticeable at the highest resolution ($\Delta x = 0.2 \mu\text{m}$), where the MS image shows very little noise or streaking in the PMMA reconstruction.

There are a few possible reasons. The sinogram data was relatively coarsely sampled angularly, and Fig. 8.4 shows a close look at sinogram line profiles generated with the full MS model versus the more efficient optimization MS model used only for forward projections in the loss function. The key difference between models is that full MS is given a high resolution phantom ($\Delta x_{\text{sim}} = 0.05 \mu\text{m}$) and thus simulates the sub-detector pixel interference effects before block averaging to the lower resolution of the sinogram. This creates a smoothing effect and is more realistic. In contrast, the optimization MS model did not upsample due to the lack of a high-resolution ground truth, setting the reconstructed image pixel size, detector sampling, and simulation sampling all equal. The line profiles of Fig. 8.4 show the consequences of this. At high resolutions, the fringes are still approximately in the correct spatial location, though failing to upsample leads to an underestimate in their magnitude,

and small additional fringes are visible in regions that should be smooth. At micron pixel sizes and greater, the fringe location shifts slightly, and the smooth curves take on a more step-like shape.

These model inaccuracies suggest that an upsampling routine within the optimization forward model might be necessary for accurate MDPR. This introduces an interesting question of how to best implement the upsampling, since it inherently requires sub-pixel information about the object not available in the context of an inverse problem. Natural first directions could be nearest-neighbor upsampling and linear interpolation, though these may have the drawbacks of creating false structural edges and over-smoothing true edges, respectively. One could also attempt MDPR at resolutions higher than the detector pixel size, though this could become poorly conditioned due to the lack of measured sub-pixel information. Either method will require additional computer memory to perform the reconstruction. A regularization strategy might also help reduce XPCI-specific artifacts [210, 214]. Optimization upsampling remains an important future direction, and the highest resolution ($\Delta x = 0.2 \text{ }\mu\text{m}$) was chosen as the focus of the remaining experiments due to its reduced artifacts.

8.4.2 *Propagation distance*

Fixing pixel size, Fig. 8.5 instead shows the effect of increasing the propagation distance between the object and the detector. Each FOM is shown as a function of “ R scale”, defined as the ratio of the propagation distance to the FOV. Previous work has indicated that MS more accurately characterizes the detected image at short propagation distances on the order of the object thickness, when the relative contribution of internal refractive effects is high [193]. We observe this effect here, with MS mostly outperforming PA at R scales up to 2, depending on the basis material and FOM referenced.

PA performance generally improved as a function of R , reflecting the greater phase information available due to larger Fresnel fringe magnitude at structural edges. However, there




Figure 8.5: SSIM and 1-RMSE in each basis image as a function of propagation distance scale relative to the FOV ($R \text{ scale} = R / \text{FOV}$).

are two interesting exceptions. First, in the PMMA image, PA yielded an initial high SSIM value at $R = 0$ before dropping and then smoothly increasing as a function of R . This may reflect the numerical sensitivity of phase-contrast forward models in the context of optimization; at $R = 0$, the PA model lacks any phase model in its update step, which could produce qualitatively smoother images. However, MS still outperforms slightly at $R = 0$, indicating it is still better to account for phase shifts accumulating within the object up to the exit plane. Second, at R scales around 2, PA image quality improvements plateaued or even slightly deteriorated in the case of aluminum RMSE. A possibility is that the material basis already has appreciable attenuation contrast, reducing the relative importance of phase information for differentiating aluminum and PMMA, and the nonlinearity of the XPCI models may require additional upsampling or regularization to realize its benefits without introducing artifacts [210, 214]. This could also explain why the MS model performance tended to decline or plateau at earlier R scales than the PA model. Weakly attenuating basis pairs and improved regularization strategies are a continued direction of investigation [179, 215]. A more uniform trend may therefore arise under more stable conditions or over a larger range




Figure 8.6: SSIM and 1-RMSE in each basis image as a function of incident flux.

of R .

8.4.3 Noise

Fig. 8.6 begins to relax the ideal reconstruction scenario by adding Poisson noise to the sinograms prior to performing MDPR. As expected, for both PA and MS models, the FOMs strictly improve as I_0 increases (and noise correspondingly decreases). SSIM appears to plateau around 10^5 photons per pixel while RMSE plateaus earlier near 10^4 . This offers useful experimental guidance on the minimum radiation dose or exposure time to achieve optimal image quality. Further, the image quality advantage of MS over PA is most significant at low noise levels, with MS converging or even falling below PA at the highest noise level examined ($I_0 = 10^2$). The one exception is that MS plateaus at a slightly worse RMSE than PA in the aluminum basis image, though we note that both values are very high ($1 - \text{RMSE} \geq 0.95$) and the difference is small. These exceptions may reflect the numerical sensitivity of phase-contrast forward models discussed above.

8.4.4 Initialization

Wave-based optimization problems are known to be challenging due to their frequency oscillations and sensitivity to initialization [166, 216]. To isolate the effect of PA and MS on their optimal solution and to accelerate convergence, the previous three experiments used a blurred-GT (BGT) initialization. More realistically, the GT is unknown, and the initial estimate must be chosen without additional information. A common method is to use an approximate analytical solution. In our case, with dual-energy PB-XPCI CT and sub-micron spatial resolution, suitable options are still emerging. Attempting FBP reconstruction with standard absorption-contrast material decomposition (inverting the 2×2 matrix of linear attenuation coefficients as described in Ch. 2) led to severe noise and streaking artifacts related to the significant phase-contrast fringes. Initial blurring compromised structural details and absorption-contrast decomposition by modulating the true detected intensity. Thus, we explored a simple uninformed initialization strategy using a constant full field of 1s.

The MS- and PA-based reconstruction images using both initialization strategies are shown in Figs. 8.7 and 8.8, with and without noise, respectively. The ideal BGT images show the same streaking artifacts noted in the detector resolution sweep. These are also partially visible in the constant initialization images, though there are additional artifacts. The MS model is able to more accurately recover the uniform circular shape of the aluminum inserts, whereas PA introduces some irregular straight blurred edges. The same effect can be seen in the PMMA images to a weaker extent. Both PA and MS return a false fringe around the PMMA cylinder, likely related to the optimization model underestimating fringe magnitude due to lack of wave upsampling during forward projection. The artifact is more irregular with PA. In terms of contrast, both constant and BGT initializations accurately recovered the aluminum volume fractions (truth = 1), whereas the constant initialization led to a negatively biased PMMA volume fraction. This was similar but slightly more severe for PA, with average pixel value in the GT PMMA mask of 0.38 for PA and 0.40 for MS.




Figure 8.7: Noisy data PA and MS reconstructions using either a constant initialization of 1s (left; “const init”) or a downsampled, blurred ground truth (right; “BGT”).

This may be related to the regularization parameter, which was found to be necessary to break the symmetry of PMMA and aluminum image reconstructions, but as a type of L_1 norm may also encourage sparsity. Future work will explore alternative regularization and improved uninformed initialization strategies, ideally finding a suitable analytical estimate that introduces spatial separation of materials without need for a priori information or aggressive material overlap penalties.

In terms of optimization performance, MS tended to converge in slightly fewer iterations than PA, for both initializations and noise levels (Fig. 8.9). This was especially true with BGT initialization, with MS requiring roughly 50 fewer iterations. On average, each PA iteration required 0.24 s to run, while each MS iteration required 0.44 s. This resulted in an increased run time of 12% to 83% for MS-based MDPR relative to PA. MS speed will be highly dependent on the number of slices used in the optimization. Here, $N_s = 8$ was selected using test simulations as the minimum slices required to achieve forward-projection convergence. The minimum number of slices will likely be larger for thicker samples and/or more detailed internal structures, requiring new evaluation to balance algorithm run time




Figure 8.8: Ideal data PA and MS reconstructions using either a constant initialization of 1s (left; “const init”) or a downsampled, blurred ground truth (right; “BGT init”).




Figure 8.9: Loss history curves for the MDPR runs shown in Figs. 8.7 and 8.8, with end threshold $\Delta\mathcal{L}_t = 10^{-5}$.

with model accuracy.

8.5 Conclusion

This chapter demonstrated an iterative MDPR framework for directly comparing PA- and MS-based forward models in dual-energy PB-XPCI material decomposition. Using an automatically differentiable forward model and loss function, the framework bypasses the need for simplified model-specific inverses or manually derived update steps. Across the conditions explored here, MS-based MDPR provided the clearest benefit at sub-micron resolution, short propagation distances, and low noise levels—realizing the benefits of MS in the relevant situations highlighted in the previous chapter. Nonetheless, its advantages were limited by sensitivity to initialization, regularization, and run time. Practical implementation will require an approach to improve optimization stability and handle large memory requirements.

Since the MS forward model uniquely captures sample depth-dependent diffraction effects, it is most relevant for 3D imaging modalities such as CT, where the goal is to spatially separate internal structures through tomographic reconstruction. For 2D projection-based modalities, phase retrieval usually aims to recover projected phase and absorption images, or projected thickness maps in the context of material decomposition. Even when samples are thick or detailed enough to produce significant depth-dependent interference effects, it is not obvious how to implement an MS model for projection-based phase retrieval, which implicitly assumes that the parameter of interest can be represented as a projected quantity. Even without MS modeling, single-projection XPCI could still benefit from PA-based iterative phase retrieval methods that include acquisition-model features that are difficult to capture analytically. In the next chapter, we explore iterative MDPR for quantitative spectral XPCI with a polychromatic spectrum and energy-discriminating detector.

CHAPTER 9

FULL SPECTRAL MATERIAL DECOMPOSITION PHASE RETRIEVAL

9.1 Introduction

Thus far, we have presented evidence for the promise of propagation-based X-ray phase-contrast imaging (PB-XPCI) as an emerging modality with greater sensitivity to fine structures that are often invisible with traditional absorption-only X-ray imaging. However, since XPCI depends on two contrast mechanisms (absorption and phase), it in principle requires multiple independent measurements for generalized accurate reconstruction and therefore suffers from a multiple measurement problem. The traditional solution relies on multiple propagation distance measurements, which is experimentally inconvenient and risks misalignment artifacts [217].

With emerging photon-counting detectors (PCDs) available to conveniently provide multiple energy measurements in a single shot, *spectral* XPCI has become a compelling direction for development. In this spirit, in Chapter 8, we proposed an iterative material decomposition phase retrieval (MDPR) method leveraging multiple monochromatic measurements from high-resolution, synchrotron micro-CT to reconstruct basis material volumes. However, this does not fully solve the multi-measurement problem, as multiple monochromatic energy measurements often still require sequential synchrotron micro-CT scans. In theory, these multiple monoenergetic scans are an improvement upon the typical multiple distance acquisitions, since they do not require movement of the sample. However, changing synchrotron energy is inconvenient, requiring a significant wait time for beam retuning, during which the detector may drift or other unforeseen complications may arise. The gold standard solution to the multi-measurement problem is not necessarily multi-energy: broadly, it must be *single-exposure*.

Single-shot multi-energy imaging can be conveniently achieved via a polychromatic source coupled with a detector capable of filtering measurements into energy-resolved data bins, such as a spectral PCD [48, 49]. If monochromatic measurements are necessary, single-exposure multi-energy imaging has been achieved at synchrotron facilities by leveraging higher-harmonic radiation of the crystal monochromator [218]. A more accessible solution for benchtop PB-XPCI systems is to use a micro-focus X-ray source, avoiding the need for beamtime at a synchrotron facility but necessitating algorithms capable of handling the polychromatic measurements. Yet most analytical spectral phase retrieval solutions are generally derived for more idealized monochromatic measurements [57–59]. If polychromaticity is considered, it is typically with weighted-average effective monoenergies or least-squares solutions [55, 56, 61, 176]. A truly quantitative solution to the spectral phase retrieval problem should be compatible with both polychromatic source spectra and other non-ideal system effects. While iterative phase retrieval methods have been developed with this in mind, they have not yet been adapted to general forward models or multiple polychromatic measurements [209, 210].

The purpose of this chapter is to expand our proposed iterative MDPR framework to full spectral imaging, natively accounting for multiple polychromatic measurements and other energy-dependent system elements.¹ Our approach also benefits from the additional modeling features described in the previous chapter, namely more accurate propagation modeling that relaxes typical assumptions such as a linearized transport-of-intensity equation [57, 59]. We operate at a spatial resolution sufficiently large for accurate forward modeling without the computationally expensive multi-slice model, focusing on a single 2D projection and the task of reconstructing basis material thickness maps.

1. Parts of this chapter are adapted from previously published conference proceedings [215].

9.2 Theory

9.2.1 Polychromatic forward model

When the projection approximation is valid, the normalized, monochromatic wave field intensity at the detector plane is

$$I_R(\mathbf{r}_\perp, E) = \left| \mathcal{D}_R(E) \exp \left[-ik(E) \int_0^T dz A(\mathbf{r}, E) \right] \right|^2, \quad (9.1)$$

where $\mathcal{D}_R(E)$ is the free-space propagation operator, T is the maximum object thickness, and A characterizes the 3D distribution of the object's refractive index in terms of its real and imaginary components,

$$A(\mathbf{r}, E) = \delta(\mathbf{r}, E) - i\beta(\mathbf{r}, E). \quad (9.2)$$

For single-projection imaging, it is convenient to rewrite the integral of Eq. 9.1 as a discrete sum over individual material thickness maps, such that

$$\int_0^T dz A(\mathbf{r}, E) = \sum_{j=1}^B t_j(\mathbf{r}_\perp) A_j(E), \quad (9.3)$$

for a B -material object.

For material decomposition, this often serves as a starting point with $B = 2$ assumed basis materials [34, 59]. In our optimization framework notation, we define the object $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_B)$ as the array of material thickness maps $\mathbf{x} \equiv t_j(\mathbf{r}_\perp)$. Then, the full polychromatic forward model can be written as the integral over monochromatic components, including the energy-dependent nonideal system elements,

$$f(\mathbf{x}) = \int dE [D(E) \eta(E) S(E) I_R(\mathbf{x}, E) * \text{psf}(\mathbf{r}_\perp, E)] \quad (9.4)$$

where $D(E)$ is a detector response function, $\eta(E)$ is a detective efficiency function, $S(E)$ is the polychromatic source spectrum in photon counts per detector pixel, $*$ indicates the 2D convolution operator with respect to $\mathbf{r}_\perp$, psf is the total system point-spread function (PSF), and $I_R(\mathbf{x}, E)$ is simply Eq. 9.1 notated as a function of the object.

For single-shot imaging with an energy-discriminating detector, such as a spectral PCD, multiple energy-resolved measurements can be modeled by defining per-bin detector response functions $D_i(E)$, where i is the energy bin index and the spectral forward model becomes $f(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_N(\mathbf{x}))$, where N is the number of bins. Using a full energy-dependent function rather than a simple binary energy threshold allows the possibility of incorporating binning nonidealities, such as the finite energy photopeak width or spectral distortions from K-escape and charge sharing.

9.2.2 Analytical inverse

An analytical inverse solution to the spectral MDPR problem has been developed using a simplified forward model based on the linearized transport-of-intensity equation (TIE), termed here “TIE-MD” [55, 59]. The algorithm and its derivation are detailed in Section 5.5.3. It has shown promising results in the context of monochromatic synchrotron radiation [60]. For use with polychromatic measurements, however, one must estimate spectral contributions in terms of refractive indices evaluated at monochromatic effective energies [56]. These are generally defined as

$$E_i^{(\text{eff})} = \frac{\int dE E w_i(E)}{\int dE w_i(E)}, \quad (9.5)$$

where the energy weight w_i is the expected number of photons at each energy E measured by the detector,

$$w_i(E) = D_i(E) \eta(E) S(E). \quad (9.6)$$

9.2.3 Optimization-based solution

We apply the automatic differentiation (AD)-enabled iterative MDPR method introduced in Chapter 8 to solve the optimization problem,

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}), \quad (9.7)$$

with loss function based on the Poisson negative log-likelihood function,

$$\mathcal{L}(\mathbf{x}) = \frac{1}{N_d} \sum_{i,d} [\bar{y}_{i,d} - y_{i,d} \ln(\bar{y}_{i,d})] + \gamma r(\mathbf{x}), \quad (9.8)$$

where i is the energy bin index, d is the detector pixel index, N_d is the total number of bin-pixel measurements, $\mathbf{y}$ is the measured data, $\bar{\mathbf{y}} = f(\mathbf{x})$ is the predicted data, and $r(\mathbf{x})$ is a regularization function.

For 2D projection imaging, it is not appropriate to choose $r(\mathbf{x})$ as a material overlap penalty as was used with the 3D CT optimization problem solved in Chapter 8, since the significance of the task lies in separating overlapping thickness projections. Here, the regularization penalty was chosen as a scaled, weighted TV penalty over the 8-pixel neighborhood $\mathcal{N}$ around each pixel,

$$r(\mathbf{x}) = \sum_{j,k} \sum_{n \in \mathcal{N}} \frac{d_n}{s_j} [x_{j,k} - x_{j,n}], \quad (9.9)$$

where j is the basis material index, k is the image pixel index, n is the neighbor pixel index,

s_j are material-dependent scale factors, and d_n are neighbor pixel-distance scale factors $d_n \in \{1, 1/2\}$. The material scales were chosen based on the relative thickness magnitude of each basis material, which was key to avoid oversmoothing when one material was significantly more sparse than the other. Similar neighborhood-based, edge-preserving regularization terms have been successfully implemented in previous iterative PB-XPCI reconstruction algorithms [209, 210].

9.3 Methods

9.3.1 Materials

Simulation studies were performed using a 50-kVp spectrum generated based on a tungsten anode with 2-mm aluminum filtration [182]. The detector was modeled as a PCD with efficiency $\eta(E)$ based on the absorption of a 1-mm thick silicon sensor, similar to a detector previously used for spectral PB-XPCI with a 55- μm pixel pitch Medipix detector [56, 200, 219]. For direct-conversion PCDs, the two most common materials are silicon or cadmium telluride (CdTe)/cadmium-zinc-telluride (CZT) [48]. Silicon offers improved robustness to low-energy distortions and energy resolution, with FWHM in the range of 3.5–5.4 keV versus 5–10 keV for CdTe/CZT [49]. Therefore, energy-resolved bin response functions $D_i(E)$ were generated for given threshold energies E_t using the Gaussian probability method described in Section 6.4 with FWHM = 5 keV, which is the higher end for silicon (see Fig. 6.4 for example $D_i(E)$ at different thresholds). This is a more realistic model than an ideal step-function threshold model, incorporating reduced energy separation of the bin images.

The 50-kVp spectrum $S(E)$ and detective efficiency function $\eta(E)$ are shown in Fig. 9.1. More generally, we note that these functions may vary across detector pixels as a function of $\mathbf{r}_\perp$. Since it is straightforward to correct these variations with an initial air scan, we did not model any $\mathbf{r}_\perp$ -dependent changes in these two functions. Additionally, the total system PSF




Figure 9.1: Polychromatic spectrum and energy-dependent detector efficiency function based on the absorption of a 1-mm thick silicon PCD.

was modeled as a 2D Gaussian kernel with constant FWHM independent of detector pixel location and photon energy. The PSF function includes contributions from both detector and source blur, and more generally it will broaden at wider cone beam angles and higher energies [71]. The constant PSF here serves as a first approximation appropriate for this initial study.

9.3.2 TIE-MD simulations

To assess the analytical TIE-MD sensitivity to polychromatic/monochromatic system mismatches, initial exploratory simulations were performed using energy, spatial resolution, and propagation distance parameters similar to those previously demonstrated [59, 60]. The detector comprised a 256×256 array of 6.5- μm pixels, positioned a distance $R = 20$ cm downstream from the object. The phantom comprised overlapping ICRU bone and soft-tissue ellipsoids with true maximum thicknesses of 0.5 mm and 1 mm, respectively. It was generated at $8\times$ higher resolution than the detector ($\Delta x_{\text{sim}} = 0.8125 \mu\text{m}$). The Gaussian system PSF was modeled with FWHM equal to the detector pixel size.

Simulations were performed over a range of energy thresholds, 16 to 26 keV. Effective energy pairs (E_1, E_2) were calculated for each threshold using Eq. 9.5. Two forms of synthetic

data were generated: true polychromatic and effective monochromatic. The former represents realistic data acquired with the full polychromatic energy-dependent imaging system, while the latter matches the idealized effective monoenergy model assumed by the TIE-MD algorithm. TIE-MD was performed on both datasets with and without Poisson noise added, and root-mean-square error (RMSE) was computed in the resulting thickness images relative to the ground truth (GT).

9.3.3 Iterative MDPR simulations

A slightly larger bone/soft-tissue phantom was designed to assess and tune the optimization-based MDPR algorithm (Fig. 9.2), with detector pixel size $\Delta x = 20 \text{ }\mu\text{m}$, system PSF of $\text{FWHM} = 2\Delta x$, and propagation distance $R = 300 \text{ cm}$. The phantom spanned a 5-mm field-of-view (FOV), with a 3-mm thick ellipsoid-cylinder with four superimposed bone ellipsoids of maximum thickness ranging from 0.2 to 0.5 mm. It was generated at $8\times$ higher spatial resolution ($\Delta x_{\text{sim}} = 2.5 \text{ }\mu\text{m}$), which was found to be sufficiently small to capture sub-detector pixel size wave interference effects. This relatively larger scale was chosen to emulate parameters more common for benchtop PB-XPCI, rather than synchrotron micro-CT used for previous TIE-MD experiments [59]. This scale is equivalent to the effective pixel size of the 55- μm pixel pitch Medipix detector with a cone-beam micro-focus source and magnification factor of 2.75, which is within the typical range of applications using magnification around 2–3 [55, 56, 61, 176].

To identify the optimal energy threshold for the imaging system, polychromatic simulations were performed as a function of threshold with subsequent TIE-MD to generate initial basis material images. Noise was added to the raw intensity data based on a Poisson distribution with incident spectrum scaled to have total photons per pixel $I = \int dE S(E) = 10^5$. The simulation error was defined as the RMSE of TIE-MD thickness estimates relative to the GT thickness maps. This was compared to the line-profile CRLB (Section 6.3.2) with a




Figure 9.2: Two-material phantom used for optimization tuning and sensitivity analysis.




Figure 9.3: Multi-material phantom used for robustness analysis.

3-mm soft-tissue background and a range of 0.3, 0.4, or 0.5-mm bone thickness signal. The thresholds minimizing simulation error and CRLB were used for the remaining simulations.

The initialization and regularization parameters for the optimization-based MDPR were then tuned. Two initialization strategies were considered: (1) a structure-naïve constant initialization with the bone-basis image set to 0 mm and the tissue-basis image set to either 0, 1, or 3 mm; or (2) an initial analytical estimate of the raw TIE-MD estimates or Gaussian-blurred TIE-MD with a 25- or 50-pixel FWHM. A finite support constraint mask was applied to the initialized thickness maps, estimated using blurred and thresholded intensity data.

The loss function was minimized with the ADAM optimizer and a learning rate of 10^{-4} [211]. The regularization weight was chosen as $\gamma = 0.02$ based on a parameter sweep over $\gamma = 0.001$ to 0.3. The MDPR convergence criterion was defined as $\Delta\mathcal{L} \leq 10^{-5}$ over 50 iterations, which was found to reliably stop iterations only after true plateaus.

After tuning the best energy threshold, initialization, and regularization, a model sensitivity analysis was performed. Realistic, noisy simulations were performed using the identified

Optimization $f(\mathbf{x})$	Model description
Full polychromatic	Eq. 9.4 with all nonideal elements.
Full + 2 $\times$ upsample	Full poly + intermediate $\Delta x_{\text{sim}} = \Delta x/2$ upsampling during wave propagation.
Full + 4 $\times$ upsample	Full poly + intermediate $\Delta x_{\text{sim}} = \Delta x/4$ upsampling during wave propagation.
Effective monochromatic	$S(E)$ and $\eta(E)$ replaced by equivalent monochromatic values corresponding to the effective energies (Eq. 9.5).
Ideal threshold	No energy threshold blur; step function detector responses $D_1(E) = 1 - \text{HS}(E - E_t)$ and $D_2(E) = \text{HS}(E - E_t)$.
Ideal PSF	No source or detector blur; ideal system PSF function $\text{psf}(\mathbf{r}_\perp) = \delta_{\text{Dirac}}(\mathbf{r}_\perp)$.

Table 9.1: Different forward projections used for the optimization model sensitivity analysis.

parameters. Optimization-based MDPR was performed using different forward projectors $f(\mathbf{x})$ within the loss function to assess which model elements are most important for accurate material decomposition. Each model is summarized in Table 9.1.

Finally, model robustness to multiple materials and low doses was assessed using a multi-material phantom (Fig. 9.3). The same simulation scale was used, with 20- μm pixels and 5-mm FOV. The phantom materials were chosen to emulate mammography, with a background ellipsoid-cylinder of fibroglandular tissue, three superimposed adipose ellipsoids, and two rows of microcalcifications with maximum thickness between 80 and 200 μm . The first row of calcifications was defined as type I (calcium oxalate dihydrate), usually associated with benign lesions, while the second row was defined with the more likely malignant type II composition (calcium hydroxyapatite) [220, 221]. Simulations were performed with $R = 300$




Figure 9.4: Example monochromatic simulations and analytical TIE-MD output for three different energy pairs. The example intensity data in the left column is shown for the middle energy pair, (20 keV, 30 keV).

cm (XPCI) or 0 cm (absorption-only imaging) at three different noise levels corresponding to total incident photons per pixel $I = 200, 10^3$, and 10^5 . Iterative MDPR was performed using a basis of soft tissue and bone. This is more similar to a realistic material decomposition task in biomedical imaging, where the assumed basis materials will be similar but not identical to the true broader set of biological materials. The adipose and gland materials are expected to primarily contribute to the tissue-basis image while the calcifications contribute to the bone-basis image [215].

9.4 Results and Discussion

9.4.1 TIE-MD energy sensitivity

Fig. 9.4 shows the exploratory monochromatic dual-energy simulations performed to assess TIE-MD in ideal conditions. In all cases, residual edge-enhancement rings are left in the bone-basis image. The tissue-basis images more successfully correct the edge fringes but instead suffer from periodic artifacts due to the Fourier-space inversion approach of the TIE-MD method. The intermediate energy pair (20 keV, 30 keV) visually appears to offer the best




Figure 9.5: RMSE of TIE-MD output thickness maps from monochromatic simulated data or from polychromatic simulated data (requiring substitution of effective energies in the TIE-MD algorithm), with and without noise, at a range of energy thresholds. The right column shows the effective monoenergies calculated for each threshold.

balance of artifact minimization and MD performance, although noise is still considerably amplified relative to the raw data. Shifting the monoenergy pairs just 5 keV lower or higher results in significant artifact increase in the tissue-basis image, while the bone-basis image appears to be more robust to energy changes despite residual edges. This indicates a high sensitivity of the method to acquisition energy selection even in the idealized monochromatic case.

The implications of this for full polychromatic imaging are probed by Fig. 9.5, which shows the RMSE of TIE-MD estimated thickness images produced from true polychromatic data or equivalent effective-energy monochromatic simulations. The difference between the polychromatic and monochromatic curves reflects the practical consequence of the MDPR model mismatch. For both materials, the monochromatic curve is more flat as a function of threshold, with lower thresholds generally producing better RMSE. Instead using the true polychromatic data, the curve shows a much more pronounced minimum in the neighborhood of $E_t = 20$ keV. The polychromatic case also appears to be more noise-sensitive, especially in the bone RMSE, where the difference between ideal and noisy curves is noticeably larger with the polychromatic data.

These initial TIE-MD results serve as motivation for the proposed iterative MDPR algo-




Figure 9.6: Basis material CRLB (3-mm soft tissue background with variable 0.3 to 0.5-mm bone thickness) and TIE-MD RMSE as a function of energy threshold.

rithm. By accounting for the realistic energy-dependent detector elements, we aim to close the gap between the monochromatic and polychromatic mismatch, offering a solution more appropriate for polychromatic spectral PB-XPCI likely to be implemented with benchtop imaging systems.

9.4.2 Iterative MDPR tuning

To choose the best energy threshold to use for the spectral polychromatic simulations, Fig. 9.6 shows the basis material CRLB and TIE-MD simulation error as a function of threshold energy. The thresholds minimizing either parameter were similar but not identical, with 23 keV minimizing the CRLBs and 20 keV minimizing the simulation error. To choose one, iterative MDPR was performed using data from both energy thresholds using the different proposed initialization strategies. Table 9.2 lists the RMSEs produced for each threshold/initialization pair, while Fig. 9.7 shows the 20-keV reconstructions for qualitative reference. The constant initialization strategy reliably reconstructed the bone ellipsoids but struggled to recover the soft-tissue background, with retrieved thickness converging within roughly 0.5-mm of the initialized thickness. This resulted in especially large soft-tissue RMSE for the 0-mm and 1-mm tissue initializations. The 3-mm constant tissue initialization “Const (3,0)” produced the best tissue images among the const-init methods, but it failed to recover the cylindrical curve

Init method	Soft tissue		Bone	
	20 keV	23 keV	20 keV	23 keV
Const (0, 0)	2.098	2.088	0.179	0.178
Const (1, 0)	1.362	1.363	0.118	0.118
Const (3, 0)	0.668	0.656	0.018	0.018*
TIE	0.343	1.172	0.037	0.095
TIE-blur25	0.324	1.059	0.036	0.088
TIE-blur50	0.307*	0.853	0.036	0.074

Table 9.2: RMSE [mm] of retrieved basis material thickness images from different initialization methods, with acquisition energy threshold optimized either for the TIE (20 keV) or CRLB (23 keV). The asterisk * indicates the minimum unrounded RMSE achieved for each material.

of the thickness map, instead producing a more flat thickness distribution. This is especially apparent in the central line profiles shown in Fig. 9.8. In contrast, the TIE-initialization methods yielded tissue maps with more accurate, curved thickness profiles. The TIE instead suffered from the low-frequency artifacts identified in the initial TIE-MD sensitivity analysis. These were somewhat ameliorated by applying a Gaussian blur to the TIE-MD estimates prior to optimization. While the TIE artifacts were not entirely removed by the iterative MDPR process, for example in the increased background thickness around the smallest bone ellipsoid, the overall RMSE was minimized using the 20-keV simulations with TIE-blur50 initialization. Therefore, these parameters were adopted for the model sensitivity analysis.

An additional observation is that the TIE-initialized RMSEs were much more sensitive to the 3-keV change in energy threshold, while the constant-initialized RMSEs changed by <0.01 mm. Looking at the 23-keV threshold only, the Const (3,0) initialization produced the best RMSEs in both material images. Further, both constant and TIE-based initialization




Figure 9.7: 20-keV threshold reconstructions produced from different optimization initialization methods.




Figure 9.8: Central line profiles, averaged over the middle ten rows, in the reconstructed 20-keV images from the best initialization methods with either a constant or TIE-MD variant.




Figure 9.9: RMSE as a function of regularization weight γ for the constant (3,0) and TIE-blur50 initializations with 20-keV threshold data. Example thickness maps are shown for the initialization and γ at the starred point on the curve.

left residual bone ellipsoid artifacts in the tissue image, likely due to its high attenuation. This suggests constant initialization can still be a valid strategy when the conditions of the TIE are not optimal, especially in terms of qualitative performance.

A noteworthy success of all initialization methods is the elimination of residual Fresnel fringes in the material images. While there was some material bleed into incorrect images (i.e. residual tissue background in the bone image and residual bone ellipsoids in the tissue image), this largely indicates that the iterative MDPR process successfully converts PB edge-enhancement information into the retrieved thicknesses.

As the final step of the optimization tuning, the regularization weight was chosen as $\gamma = 0.02$. Fig. 9.9 shows the basis-image RMSEs as a function of γ for the 20-keV threshold simulation with the two best initializations, const (3,0) and TIE-blur50. The value of $\gamma = 0.02$ was chosen to balance RMSE and visual appearance of each material image produced by the different methods, aiming to recover the smoothest tissue-thickness profile without blurring the bone ellipsoid edges.

Optimizer $f(\mathbf{x})$	Soft tissue		Bone	
	const init	TIE init	const init	TIE init
Full polychromatic	0.676	0.303	0.018*	0.036
Full + $2\times$ upsample	0.586	0.299*	0.027	0.035
Full + $4\times$ upsample	0.586	0.307	0.041	0.043
Effective monochromatic	0.826	0.354	0.039	0.071
Ideal threshold	0.447	0.363	0.029	0.018
Ideal PSF	0.652	0.303	0.018	0.035
TIE-MD (no optimization)	—	0.464	—	0.089

Table 9.3: Model sensitivity analysis: RMSE [mm] for different optimization forward models, systematically excluding nonideal system elements such as energy dependencies and blur. Optimization was initialized with either TIE-blur50 or Const (3,0). The asterisk * indicates the minimum unrounded RMSE for each material.

9.4.3 Model sensitivity

Using the tuned simulation and optimization parameters, the model sensitivity analysis was performed by solving for the thickness maps with the iterative MDPR algorithm using forward projectors with varying levels of realism. Table 9.3 lists the basis-material RMSEs of the final estimated thickness maps for the different models.

Overall, the full polychromatic model produced among the most accurate reconstructions, but the results also show that increased forward model realism does not guarantee improved performance without adequate initialization. TIE initialization substantially reduced tissue RMSE relative to constant initialization, with the best result obtained using the full model with $2\times$ upsampling. For bone, the lowest RMSE was obtained using the full polychromatic model with constant initialization, although the bone images were generally robust to model

variations. Several models gave similar bone-basis performance, with only sub- μm differences in RMSE. The effective-monochromatic approximation produced the largest performance degradation, especially for constant initialization. The ideal threshold bin-response functions were also important with TIE initialization. In contrast, removing the PSF or increasing the intermediate wave-propagation upsampling in the full model produced only modest RMSE changes, indicating that these factors were not as dominant in the optimization performance.

The images shown in Fig. 9.10 complement the RMSE values listed in the table, showing the final basis-material estimates produced by each model and their difference images relative to the GT. All models successfully smooth out most of the low-frequency artifacts and residual fringes present in the raw TIE-MD initialization image (bottom row). Among the iterative MDPR images, the effective-monochromatic results are the most noticeably different, with higher noise, stronger remnant fringes, and the largest overestimate in bone thickness. This reflects the importance of fully modeling the polychromatic spectrum, since each monoenergetic component contributes to the wave-interference patterns that determine the final detected edge enhancement.

Likewise, Fig. 9.11 shows the basis images and GT differences for the full polychromatic model with different intermediate wave upsampling factors. Although the RMSE values differ by only a few μm , the images are noticeably smoother and less noisy with higher upsampling. However, the upsampled images also suffer from artifacts, such as interspersed jagged edges with $2\times$ upsampling and isolated bright and dark speckles with $4\times$ upsampling. This may be associated with the nearest-neighbor upsampling method used to match the estimated thickness maps to the wave-propagation sampling during the forward projection step. Additionally, while upsampling helped smooth the soft tissue image, it worsened the underestimation of bone thickness. Given these trade-offs, matched wave sampling is likely sufficient for most implementations. To realize the benefits of upsampling without the instability, a smoother phantom upsampling method such as 2D linear or quadratic interpolation

could be attempted, although this would introduce considerably greater computational expense at each iteration.

9.4.4 *Calcification decomposition*

Fig. 9.12 shows the raw intensity images of the multi-material phantom simulated at the three noise levels, with and without phase contrast ($R = 0$ or 300 cm). At the highest noise level, the largest type II calcifications are visible both with and without phase contrast, while the type I calcifications are challenging to detect at all sizes. The larger adipose ellipsoids are also barely visible without the additional edge enhancement provided by PB-XPCI.

In the basis-material domain, Fig. 9.13 shows the estimated thickness maps retrieved using iterative MDPR with constant $(3, 0)$ initialization, which was found to be more stable for the multi-material phantom. As a closer look at calcification recovery, Fig. 9.14 shows the central line profiles through the type I and type II calcifications in the bone-basis image. At the highest noise level, the 300-cm images successfully recover the four largest type I and type II calcifications, with only the smallest 80- μm calcifications, spanning just four detector pixels, possibly lost to noise. At lower noise levels, the size and material differences among all calcification types become apparent in the bone-basis image. The $R = 300$ -cm images also capture the adipose ellipsoids, although their thickness contribution is split between the tissue and bone basis images.

In the 0-cm images, the lack of edge enhancement leads to smoother thickness maps with considerably less noise than the 300-cm thickness estimates. However, this relative improvement comes at the cost of calcification detectability. While the $>100\text{-}\mu\text{m}$ type II calcifications are detectable in all images, the type I calcifications are blurred and nearly entirely lost except in the highest-dose images. Although this might be possibly mitigated by acquisition-specific tuning of the regularization weight γ , there may be insufficient contrast information in the raw intensity images to reliably recover these smaller and inherently




Figure 9.10: Thickness maps and differences relative to ground truth (GT), reconstructed using different optimizer forward-projection models. As reference, the bottom row shows the unoptimized TIE.




Figure 9.11: Thickness maps and differences relative to ground truth (GT), reconstructed using the full polychromatic forward projector in each iteration with different degrees of wave upsampling (1, 2, or 4 $\times$ the reconstructed resolution).




Figure 9.12: Raw data simulated at three noise levels with the multi-material phantom (threshold energy $E_t = 20$ keV).

lower-contrast calcifications. Overall, these simulations indicate that iterative MDPR for full polychromatic spectral XPCI remains robust with low doses and, more realistically, imperfectly matched basis materials.

9.5 Conclusion

To close, this chapter presents a full spectral MDPR solution, in the sense that the iterative optimization-based approach accounts for the complete polychromatic source spectrum and energy-dependent system response at each update step. After tuning the optimization parameters, the approach successfully recovered basis material thickness maps under challenging conditions, including high noise, small calcifications, and imperfectly matched basis materials. The method is still imperfect, leaving some residual material cross talk between basis material images, but it offers large improvement over the analytical reference algorithm, TIE-MD. Further, by systematically varying the elements included in the optimizer forward model, we demonstrated the value of the full spectral approach, with the polychromatic model outperforming both analytical TIE-MD and optimization-based MDPR using effective monoenergies. These results support full spectral MDPR as a promising reconstruction




Figure 9.13: M DPR-retrieved basis material thickness maps of the multi-material phantom data.




Figure 9.14: Line profiles through the bone-basis image in the central rows corresponding to the type I (top) and type II calcifications (bottom).

approach for future benchtop micro-focus PB-XPCI systems, with polychromatic sources and energy-discriminating detector schemes the most promising candidates for single-shot quantitative phase retrieval.

CHAPTER 10

CONCLUSIONS

Altogether, this dissertation has united traditional spectral X-ray imaging with the emerging field of X-ray phase-contrast imaging (XPCI) through the lens of material decomposition. Taken separately, spectral imaging and XPCI each aim to improve soft-tissue image quality, but they face distinct challenges: limited native soft-tissue sensitivity and the need for multiple measurements, respectively. Thus, spectral XPCI has emerged as a promising candidate for quantitative phase retrieval, enabled by multi-measurement, single-shot acquisition with energy-resolving photon-counting detectors (PCDs). We have thoroughly examined the assumptions underlying existing spectral phase-retrieval and material-decomposition methods, developed original theoretical and simulation-based metrics for assessing retrieved image quality, and presented a more general optimization-based solution. This iterative solution enables the inversion of a more general forward model, bypassing common simplifying approximations required by existing closed-form inversions over parameter ranges typical of both current and novel sub-micron X-ray imaging systems. This lays the foundation for future work aiming to more fully realize the benefits of spectral imaging through quantitative phase retrieval and material decomposition.

10.1 Summary

The first part of the dissertation focused on absorption-contrast spectral X-ray imaging. Chapter 2 presented an overview of the relevant X-ray imaging fundamentals. This included a review of common X-ray sources and detectors used throughout this work, from micron-scale synchrotron imaging to new clinical PCD-CT systems [48, 49, 51, 62], and the basic theory underlying two-energy basis decomposition methods [33, 34].

In Chapters 3 and 4, we extended this absorption-contrast theory to examine and opti-

mize dual-energy CT systems with the Cramér–Rao lower bound (CRLB) in the context of combined megavoltage (MV) and kilovoltage (kV) spectra [102, 145]. The MV-kV spectral combination is conveniently attainable using the two X-ray sources already available on most radiotherapy machines: MV for treatment, and kV for on-board imaging. In Chapter 3, our results demonstrated the benefits of MV-kV dual-energy CT over dose-matched single-kV CT, especially for metal artifact correction. However, the results also highlighted potential pitfalls when MV image noise becomes dominant. Chapter 4 explored how this limitation might be mitigated by using a PCD to improve native MV image contrast and eliminate electronic noise, which can dominate at the low photon-count levels needed to reduce MV imaging dose to diagnostic imaging levels. Based on the single-line integral CRLB, we predicted 15–45% improvement in MV-kV basis-material image quality using PCDs instead of traditional detectors. In CT simulations, we found PCDs yielded more uniform noise texture and improved contrast-to-noise ratio (CNR) in virtual monoenergetic images (VMIs). These simulation studies motivate future experimental work to validate our findings and provide a tool for spectral dose allocation optimization prior to patient scanning.

In the subsequent chapters, the dissertation scope expanded to include phase contrast, with the goal of developing a more generalizable spectral phase retrieval approach. Chapter 5 began this transition by reviewing propagation-based (PB) XPCI theory, detailing each assumption made to arrive at conventional imaging forward models. The chapter also provided derivations of the analytical phase retrieval methods implemented in later sections: the widely used single-distance method for a homogeneous object and the multi-energy methods for phase retrieval and material decomposition [24, 55, 59]. This overview highlighted both advantages and potential quantitative weaknesses when approximations were not valid.

In Chapter 6, we developed new CRLB toy models suitable for assessing material decomposition performance with PB-XPCI [172, 173]. The unique challenge was in designing a model to account for the spatially distributed edge-enhancement information without

trivializing the material decomposition problem by introducing excess absorption-contrast information. Building upon the traditional single-line integral CRLB, we systematically developed three models of increasing detail: a spatially resolved two-pixel system, a line-profile single-edge CRLB, and a 2D toy geometry model. We used these models to assess optimal monochromatic energy pairs and polychromatic energy thresholds for spectral XPCI, finding similar trends between the 1D line-profile and 2D geometry CRLBs and validating the optimal thresholds with a maximum-likelihood estimation (MLE) simulation study. These models establish a useful foundation for optimizing spectral PB-XPCI acquisition and advancing its translation toward practical applications.

Chapter 7 began the task of assessing PB-XPCI forward models with the context of quantitative phase retrieval, specifically focusing on the widely applied projection approximation (PA). While PA is fairly accurate at conventional imaging scales, it may not fully capture phase-contrast effects emerging from refraction within the object itself when the detector resolution is high, feature sizes are small, and propagation distances short [165, 192, 193]. We verified these predictions of previous work by comparing the PA model to a more general multi-slice (MS) model [164, 191, 222]. We found that meaningfully large differences between MS and PA simulations do not necessarily translate to significant differences after applying common phase retrieval methods, such as the single-distance, homogeneous-object method [24]. Micro-CT simulations of metal-stained and unstained zebrafish also highlighted artifacts arising when the homogeneous-material assumption was violated. These findings motivated the development of a more general phase retrieval solution, aiming to capture the additional information of object depth-dependent phase shifts and larger attenuation for improved quantitative performance.

To this end, Chapters 8 and 9 presented an optimization-based material decomposition phase retrieval (MDPR) solution and demonstrated its advantages over methods that invert approximated forward models [179, 180, 215]. Using automatically differentiable (AD) for-

ward models, our iterative MDPR solution reduced its reliance on simplifying assumptions and more fully saturated achievable basis material image quality. Chapter 8 focused on dual-energy synchrotron micro-CT with the MS-based forward model, while Chapter 9 focused on polychromatic benchtop XPCI with a nonideal energy-discriminating PCD. In both chapters, the AD-enabled optimization framework allowed different forward-model components to be swapped directly within the reconstruction, lending itself to assessment of which model elements are most essential for accurate MDPR. This provided evidence for when the additional computational expense of a more general forward model is warranted, beyond a simple comparison of iterative and analytical solutions. We ultimately demonstrate a key improvement to quantitative MDPR lies in accounting for the full polychromatic X-ray spectrum and energy-dependent system response.

10.2 Future Directions

There are two natural future directions. The first is to expand the computational simulations to larger scales using multiple GPUs coupled with distributed-memory strategies [212]. Since accurate PB-XPCI simulation typically requires substantial wave field upsampling, large-scale high-resolution reconstruction remains computationally challenging, particularly for 3D CT imaging [166]. Improved computational strategies will therefore be necessary for practical implementation of the optimization-based theory presented here. The second direction is validation with real spectral XPCI data. Although the simulation studies in this dissertation were designed to minimize forward-model approximations, practical imaging systems introduce additional nonidealities that are difficult to fully capture in simulation, including photon scatter, source and detector calibration, PCD spectral distortions, and other imperfect knowledge of the imaging system. Such effects are best addressed through experimental studies, spanning simple image quality phantoms as well as biological samples. These experiments would test both the accuracy of the proposed full spectral forward model

and the robustness of the optimization-based MDPR framework for nonideal acquisitions.

REFERENCES

- [1] J. H. Hubbell and S. M. Seltzer, “X-ray mass attenuation coefficients, NIST standard reference database 126,” National Institute of Standards and Technology, NIST Standard Reference Database 126, 2004.
- [2] O. Glasser, *Wilhelm Conrad Roentgen and the Early History of Roentgen Rays*. Springfield, IL: Charles C. Thomas, 1934.
- [3] W. C. Röntgen, “Ueber eine neue art von strahlen,” *Sitzungsber. Phys.-Med. Ges. Würzburg*, vol. 29, pp. 132–141, 1895.
- [4] C. Jacobsen, *X-ray Microscopy*. Cambridge University Press, 2019. ISBN 978-1-107-07657-0
- [5] W. E. Brant and C. A. Helms, *Fundamentals of Diagnostic Radiology*, 4th ed. Lippincott Williams & Wilkins, 2012. ISBN 9781451171396
- [6] J. T. Bushberg, J. A. Seibert, E. M. Leidholdt, and J. M. Boone, *The Essential Physics of Medical Imaging*, 3rd ed. Lippincott Williams & Wilkins, 2011. ISBN 9780781780575
- [7] J. Hsieh, *Computed Tomography Principles, Design, Artifacts, and Recent Advances*, 4th ed. SPIE press, 2022, vol. PM344. ISBN 9781510646872
- [8] L. Heck and J. Herzen, “Recent advances in X-ray imaging of breast tissue: From two- to three-dimensional imaging,” *Phys. Med.*, vol. 79, pp. 69–79, 2020. doi: 10.1016/j.ejmp.2020.10.025
- [9] D. L. Miglioretti, R. Smith-Bindman, T. Onega, L. M. Henderson, J. Alford-Teaster, D. L. Weaver, B. L. Sprague, S. D. Herschorn, A. N. A. Tosteson, and K. Kerlikowske, “Association between false-positive results and return to screening mammography in the breast cancer surveillance consortium cohort,” *Ann. Intern. Med.*, 2024. doi: 10.7326/M24-0123
- [10] S. N. Duggan, M. Azharuddin, R. Hernández, C. Robertson, D. Cooper, E. McCall, P. Manson, G. Virgili, M. Clarke, S. Treweek, and M. Brazzelli, “Supplemental imaging modalities for breast cancer screening in women with dense breasts: A systematic review with economic considerations,” *Breast*, vol. 85, p. 104668, 2026. doi: 10.1016/j.breast.2025.104668
- [11] B. D. de Senneville, F. Z. Khoubai, M. Bevilacqua, A. Labedade, K. Flosseau, C. Chardot, S. Branchereau, J. Ripoché, S. Cairo, E. Gontier, and C. Grosset, “Deciphering tumour tissue organization by 3D electron microscopy and machine learning,” *Communications biology*, vol. 4, no. 1, p. 1390, 2021. doi: 10.1038/s42003-021-02919-z

- [12] M. A. Yakovlev, K. Liang, C. R. Zaino, D. J. Vanselow, A. L. Sugarman, A. Y. Lin, P. J. La Rivière, Y. Zheng, J. D. Silverman, J. C. Leichty, S. X. Huang, and K. C. Cheng, “Quantitative geometric modeling of blood cells from x-ray histotomograms of whole zebrafish larvae,” *eLife*, vol. 12, p. RP89432, 2023. doi: 10.7554/eLife.89432
- [13] M. S. Ngu, D. J. Vanselow, A. L. Sugarman, R. A. Saint-Fort, C. R. Zaino, M. A. Yakovlev, K. C. Cheng, and K. C. Ang, “Staining and resin embedding of whole *Daphnia magna* samples for micro-CT imaging enabling 3d visualization of cells, tissues, and organs,” *PLoS ONE*, vol. 19, no. 11, p. e0313389, 2024. doi: 10.1371/journal.pone.0313389
- [14] M. S. Ngu, D. J. Vanselow, C. R. Zaino, A. Q. Sugarman, M. A. Yakovlev, K. C. Cheng, and K. C. Ang, “Staining and resin embedding of whole daphnia samples for micro-CT imaging,” *bioRxiv*, pp. 2023–05, 2023.
- [15] A. Einstein, “Lassen sich brechungsexponenten der körper für röntgenstrahlen experimentell ermitteln?” *Verh. Dtsch. Phys. Ges.*, vol. 20, pp. 86–87, 1918.
- [16] A. H. Compton, “The total reflexion of x-rays,” *Phil. Mag.*, vol. 45, no. 270, pp. 1121–1131, 1923.
- [17] F. Zernike, “Phase contrast, a new method for the microscopic observation of transparent objects,” *Physica*, vol. 9, no. 7, pp. 686–698, 1942. doi: 10.1016/S0031-8914(42)80035-X
- [18] U. Bonse and M. Hart, “An X-ray interferometer,” *Appl. Phys. Lett.*, vol. 6, no. 8, pp. 155–156, 1965. doi: 10.1063/1.1754212
- [19] A. Snigirev, I. Snigireva, V. Kohn, S. Kuznetsov, and I. Schelokov, “On the possibilities of X-ray phase contrast microimaging by coherent high-energy synchrotron radiation,” *Rev. Sci. Instrum.*, vol. 66, no. 12, pp. 5486–5492, 1995. doi: 10.1063/1.1146073
- [20] S. W. Wilkins, T. E. Gureyev, D. Gao, A. Pogany, and A. W. Stevenson, “Phase-contrast imaging using polychromatic hard x-rays,” *Nature*, vol. 384, no. 6607, pp. 335–338, 1996. doi: 10.1038/384335a0
- [21] T. E. Gureyev, S. C. Mayo, D. E. Myers, Y. Nesterets, D. Paganin, A. Pogany, A. W. Stevenson, and S. Wilkins, “Refracting Röntgen’s rays: propagation-based X-ray phase contrast for biomedical imaging,” *J. Appl. Phys.*, vol. 105, no. 10, 2009. doi: 10.1063/1.3115402
- [22] A. Burvall, U. Lundström, P. A. Takman, D. H. Larsson, and H. M. Hertz, “Phase retrieval in X-ray phase-contrast imaging suitable for tomography,” *Opt. Express*, vol. 19, no. 11, pp. 10 359–10 376, 2011. doi: 10.1364/OE.19.010359
- [23] D. Paganin, *Coherent X-Ray Optics*, ser. Oxford Series on Synchrotron Radiation. Oxford University Press, 2006, no. 6. ISBN 9780198567288

- [24] D. Paganin, S. C. 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,” *J. Microsc.*, vol. 206, no. 1, pp. 33–40, 2002. doi: 10.1046/j.1365-2818.2002.01010.x
- [25] S. C. Mayo, T. J. Davis, T. E. Gureyev, P. R. Miller, D. Paganin, A. Pogany, A. W. Stevenson, and S. W. Wilkins, “X-ray phase-contrast microscopy and microtomography,” *Opt. Express*, vol. 11, no. 19, pp. 2289–2302, 2003. doi: 10.1364/OE.11.002289
- [26] M. A. Beltran, D. M. Paganin, K. Uesugi, and M. J. Kitchen, “2D and 3D X-ray phase retrieval of multi-material objects using a single defocus distance,” *Opt. Express*, vol. 18, no. 7, pp. 6423–6436, 2010. doi: 10.1364/OE.18.006423
- [27] T. M. Gureyev, T. J. Davis, A. Pogany, S. C. Mayo, and S. W. Wilkins, “Optical phase retrieval by use of first Born- and Rytov-type approximations,” *Appl. Opt.*, vol. 43, no. 12, pp. 2418–2430, 2004. doi: 10.1364/AO.43.002418
- [28] E. Fredenberg, “Spectral and dual-energy X-ray imaging for medical applications,” *Nucl. Instrum. Methods Phys. Res., Sect. A*, vol. 878, pp. 74–87, 2018. doi: 10.1016/j.nima.2017.07.044
- [29] B. Jacobson, “Dichromatic absorption radiography. Dichromography,” *Acta Radiol.*, vol. 39, no. 6, pp. 437–452, 1953. doi: 10.3109/00016925309136730
- [30] A. M. Cormack, “Representation of a function by its line integrals, with some radiological applications,” *J. Appl. Phys.*, vol. 34, no. 9, pp. 2722–2727, 1963. doi: 10.1063/1.1729798
- [31] A. M. Cormack, “Representation of a function by its line integrals, with some radiological applications. II,” *J. Appl. Phys.*, vol. 35, no. 10, pp. 2908–2913, 1964. doi: 10.1063/1.1713127
- [32] G. N. Hounsfield, “Computerized transverse axial scanning (tomography): Part 1. description of system,” *The British journal of radiology*, vol. 46, no. 552, pp. 1016–1022, 1973. doi: 10.1259/0007-1285-46-552-1016
- [33] R. E. Alvarez and A. Macovski, “Energy-selective reconstructions in X-ray computerised tomography,” *Phys. Med. Biol.*, vol. 21, no. 5, p. 733, 1976. doi: 10.1088/0031-9155/21/5/002
- [34] A. Macovski, R. Alvarez, J.-H. Chan, J. Stonestrom, and L. Zatz, “Energy dependent reconstruction in X-ray computerized tomography,” *Comput. Biol. Med.*, vol. 6, no. 4, pp. 325–336, 1976. doi: 10.1016/0010-4825(76)90069-X
- [35] R. A. Rutherford, B. R. Pullan, and I. Isherwood, “Measurement of effective atomic number and electron density using an emi scanner,” *Neuroradiology*, vol. 11, no. 1, pp. 15–21, 1976. doi: 10.1007/BF00327253

- [36] S. J. Riederer and C. A. Mistretta, “Selective iodine imaging using K-edge energies in computerized X-ray tomography,” *Med. Phys.*, vol. 4, no. 6, pp. 474–481, 1977. doi: 10.1118/1.594357
- [37] H. K. Genant and D. Boyd, “Quantitative bone mineral analysis using dual energy computed tomography,” *Invest. Radiol.*, vol. 12, no. 6, pp. 545–551, 1977. doi: 10.1097/00004424-197711000-00015
- [38] L. A. Lehmann, R. E. Alvarez, A. Macovski, W. R. Brody, N. J. Pelc, S. J. Riederer, and A. L. Hall, “Generalized image combinations in dual KVp digital radiography,” *Med. Phys.*, vol. 8, no. 5, pp. 659–667, 1981. doi: 10.1118/1.595025
- [39] K. Taguchi, I. Blevis, and K. Iniewski, Eds., *Spectral, Photon Counting Computed Tomography: Technology and Applications*, 1st ed. Boca Raton, FL: CRC Press, 2020. ISBN 9781138598126
- [40] W. A. Kalender, W. H. Perman, J. R. Vetter, and E. Klotz, “Evaluation of a prototype dual-energy computed tomographic apparatus. I. phantom studies,” *Med. Phys.*, vol. 13, no. 3, pp. 334–339, 1986. doi: 10.1118/1.595958
- [41] J. R. Vetter, W. H. Perman, W. A. Kalender, R. B. Mazess, and J. E. Holden, “Evaluation of a prototype dual-energy computed tomographic apparatus. II. determination of vertebral bone mineral content,” *Med. Phys.*, vol. 13, no. 3, pp. 340–343, 1986. doi: 10.1118/1.595951
- [42] T. G. Flohr, C. H. McCollough, H. Bruder, M. Petersilka, K. Gruber, C. Süß, M. Grasruck, K. Stierstorfer, B. Krauss, R. Raupach, A. N. Primak, A. Küttner, S. Achenbach, C. Becker, A. Kopp, and B. M. Ohnesorge, “First performance evaluation of a dual-source CT (DSCT) system,” *Eur. Radiol.*, vol. 16, no. 2, pp. 256–268, 2006. doi: 10.1007/s00330-005-2919-2
- [43] T. R. C. Johnson, B. Krauss, M. Sedlmair, M. Grasruck, H. Bruder, D. Morhard, C. Fink, S. Weckbach, M. Lenhard, B. Schmidt, T. Flohr, M. F. Reiser, and C. R. Becker, “Material differentiation by dual energy CT: initial experience,” *Eur. Radiol.*, vol. 17, no. 6, pp. 1510–1517, 2007. doi: 10.1007/s00330-006-0517-6
- [44] B. Ruzsics, H. Lee, P. L. Zwerner, M. Gebregziabher, P. Costello, and U. J. Schoepf, “Dual-energy CT of the heart for diagnosing coronary artery stenosis and myocardial ischemia-initial experience,” *Eur. Radiol.*, vol. 18, no. 11, pp. 2414–2424, 2008. doi: 10.1007/s00330-008-1022-x
- [45] D. Zhang, X. Li, and B. Liu, “Objective characterization of GE discovery CT750 HD scanner: Gemstone spectral imaging mode,” *Med. Phys.*, vol. 38, no. 3, pp. 1178–1188, 2011. doi: 10.1118/1.3551999
- [46] A. Altman and R. Carmi, “TU-E-210A-03: A double-layer detector, dual-energy CT — principles, advantages and applications,” *Med. Phys.*, vol. 36, no. 6Part24, p. 2750, 2009. doi: 10.1118/1.3182434

- [47] S. Kido, H. Nakamura, W. Ito, K. Shimura, and H. Kato, “Computerized detection of pulmonary nodules by single-exposure dual-energy computed radiography of the chest (part 1),” *Eur. J. Radiol.*, vol. 44, no. 3, pp. 198–204, 2002. doi: 10.1016/S0720-048X(02)00268-1
- [48] T. Flohr, M. Petersilka, A. Henning, S. Ulzheimer, J. Ferda, and B. Schmidt, “Photon-counting CT review,” *Physica Medica*, vol. 79, pp. 126–136, 2020. doi: 10.1016/j.ejmp.2020.10.030
- [49] M. Danielsson, M. Persson, and M. Sjölin, “Photon-counting X-ray detectors for CT,” *Phys. Med. Biol.*, vol. 66, no. 3, p. 03TR01, 2021. doi: 10.1088/1361-6560/abc5a5
- [50] F. R. Schwartz, A. D. Sodickson, P. J. Pickhardt, D. V. Sahani, M. H. Lev, and R. Gupta, “Photon-counting CT: Technology, current and potential future clinical applications, and overview of approved systems and those in various stages of research and development,” *Radiology*, vol. 314, no. 3, p. e240662, 2025. doi: 10.1148/radiol.240662
- [51] K. Rajendran, M. Petersilka, A. Henning, E. R. Shanblatt, B. Schmidt, T. G. Flohr, A. Ferrero, F. Baffour, F. E. Diehn, L. Yu, P. Rajiah, J. G. Fletcher, S. Leng, and C. H. McCollough, “First clinical photon-counting detector CT system: technical evaluation,” *Radiology*, vol. 303, no. 1, pp. 130–138, 2022. doi: 10.1148/RADIOL.212579
- [52] U.S. Food and Drug Administration, “510(k) premarket notification: K253520, Photonova Spectra; Photonova Spectra Select,” 2026, decision date: March 20, 2026. [Online]. Available: <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmn.cfm?ID=K253520>
- [53] T. E. Gureyev, S. Mayo, S. Wilkins, D. Paganin, and A. W. Stevenson, “Quantitative in-line phase-contrast imaging with multienergy X rays,” *Phys. Rev. Lett.*, vol. 86, no. 25, p. 5827, 2001. doi: 10.1103/PhysRevLett.86.5827
- [54] T. E. Gureyev, A. W. Stevenson, D. M. Paganin, T. Weitkamp, A. Snigirev, I. Snigireva, and S. W. Wilkins, “Quantitative analysis of two-component samples using in-line hard X-ray images,” *J. Synchrotron Radiat.*, vol. 9, no. 3, pp. 148–153, 2002. doi: 10.1107/S0909049502004971
- [55] D. Gürsoy and M. Das, “Single-step absorption and phase retrieval with polychromatic X-rays using a spectral detector,” *Opt. Lett.*, vol. 38, no. 9, pp. 1461–1463, 2013. doi: 10.1364/OL.38.001461
- [56] I. Vazquez, I. E. Harmon, J. C. Rodriguez Luna, and M. Das, “Quantitative phase retrieval with low photon counts using an energy resolving quantum detector,” *J. Opt. Soc. Am. A*, vol. 38, no. 1, pp. 71–79, 2021. doi: 10.1364/JOSAA.396717
- [57] H. T. Li, A. M. Kingston, G. R. Myers, L. Beeching, and A. P. Sheppard, “Linear iterative near-field phase retrieval (LIPR) for dual-energy X-ray imaging and material discrimination,” *J. Opt. Soc. Am. A*, vol. 35, no. 1, pp. A30–A39, 2018. doi: 10.1364/JOSAA.35.000A30

- [58] H. T. Li, F. Schaff, L. C. P. Croton, K. S. Morgan, and M. J. Kitchen, “Quantitative material decomposition using linear iterative near-field phase retrieval dual-energy X-ray imaging,” *Phys. Med. Biol.*, vol. 65, no. 18, p. 185014, 2020. doi: 10.1088/1361-6560/ab9558
- [59] F. Schaff, K. S. Morgan, J. A. Pollock, L. C. Croton, S. B. Hooper, and M. J. Kitchen, “Material decomposition using spectral propagation-based phase-contrast X-ray imaging,” *IEEE Trans. Med. Imaging*, vol. 39, no. 12, pp. 3891–3899, 2020. doi: 10.1109/TMI.2020.3006815
- [60] F. Schaff, J. A. Pollock, K. S. Morgan, and M. J. Kitchen, “Spectral propagation-based X-ray phase-contrast computed tomography,” *J. Med. Imaging*, vol. 9, no. 3, p. 031506, 2022. doi: 10.1117/1.JMI.9.3.031506
- [61] M. U. Ghani, A. Yan, L. L. Fajardo, X. Wu, and H. Liu, “Dual-energy phase retrieval algorithm for inline phase sensitive X-ray imaging system,” *Opt. Express*, vol. 29, no. 17, pp. 26 538–26 552, 2021. doi: 10.1364/OE.431623
- [62] C. Sundberg, M. Persson, J. J. Wikner, and M. Danielsson, “1- μ m spatial resolution in silicon photon-counting CT detectors,” *J. Med. Imaging*, vol. 8, no. 6, pp. 063 501–063 501, 2021. doi: 10.1117/1.JMI.8.6.063501
- [63] D. W. O’Connell, K. S. Morgan, G. Ruben, F. Schaff, L. C. P. Croton, G. A. Buckley, D. M. Paganin, K. Uesugi, and M. J. Kitchen, “Photon-counting, energy-resolving and super-resolution phase contrast X-ray imaging using an integrating detector,” *Opt. Express*, vol. 28, no. 5, pp. 7080–7094, 2020. doi: 10.1364/OE.384928
- [64] R. Brunskog, M. Persson, Z. Jin, and M. Danielsson, “First experimental evaluation of a high-resolution deep silicon photon-counting sensor,” *J. Med. Imaging*, vol. 11, no. 1, p. 013503, 2024. doi: 10.1117/1.JMI.11.1.013503
- [65] P. Bartl, T. Michel, F. Nachtrab, N. Uhlmann, and G. Anton, “Evaluation of X-ray phase-contrast imaging with the medipix,” *Nucl. Instrum. Methods Phys. Res. Sect. A*, vol. 633, no. Supplement 1, pp. S143–S147, 2011. doi: 10.1016/j.nima.2010.06.150
- [66] E. Roessl and C. Herrmann, “Cramér–Rao lower bound of basis image noise in multiple-energy X-ray imaging,” *Phys. Med. Biol.*, vol. 54, no. 5, pp. 1307–1318, 2009. doi: 10.1088/0031-9155/54/5/014
- [67] Y. Ge, K. Li, and G.-H. Chen, “Cramér–Rao lower bound in differential phase contrast imaging and its application in the optimization of data acquisition systems,” in *Medical Imaging 2014: Physics of Medical Imaging*, ser. Proc. SPIE, vol. 9033, 2014. doi: 10.1117/12.2043930 p. 90330F.
- [68] K. Majidi, J. Li, C. Muehleman, and J. G. Brankov, “Noise and analyzer-crystal angular position analysis for analyzer-based phase-contrast imaging,” *Phys. Med. Biol.*, vol. 59, no. 8, p. 1877, 2014. doi: 10.1088/0031-9155/59/8/1877

- [69] S. R. Cherry, J. A. Sorenson, and M. E. Phelps, *Physics in Nuclear Medicine*, 4th ed. Elsevier Saunders, 2012. ISBN 9781416051985
- [70] F. H. Attix, *Introduction to Radiological Physics and Radiation Dosimetry*. John Wiley & Sons, Ltd, 1986. ISBN 0471011460
- [71] H. H. Barrett and K. J. Myers, *Foundations of Image Science*. Wiley, 2004. ISBN 9780471153009
- [72] P. J. La Rivière and P. Vargas, “Correction for resolution nonuniformities caused by anode angulation in computed tomography,” *IEEE Trans. Med. Imaging*, vol. 27, no. 9, pp. 1333–1341, 2008. doi: 10.1109/TMI.2008.923639
- [73] R. Behling, “X-ray sources: 125 years of developments of this intriguing technology,” *Physica Medica*, vol. 79, pp. 162–187, 2020. doi: 10.1016/j.ejmp.2020.07.021
- [74] H. Sugie, M. Tanemura, V. Filip, K. Iwata, K. Takahashi, and F. Okuyama, “Carbon nanotubes as electron source in an x-ray tube,” *Applied Physics Letters*, vol. 78, no. 17, pp. 2578–2580, 2001. doi: 10.1063/1.1367278
- [75] R. Behling, C. Hulme, G. Poludniowski, P. Tolias, and M. Danielsson, “Tungsten microparticle streams enable simultaneous reflective and transmissive x-ray emission and high central image resolution,” in *Medical Imaging 2026: Physics of Medical Imaging*, vol. 13924. SPIE, 2026. doi: 10.1117/12.3084831 p. 139240R.
- [76] S. Shin, “New era of synchrotron radiation: fourth-generation storage ring,” *AAPPS Bull.*, vol. 31, no. 1, p. 21, 2021. doi: 10.1007/s43673-021-00021-4
- [77] S. S. Hsieh, S. Leng, K. Rajendran, S. Tao, and C. H. McCollough, “Photon counting CT: Clinical applications and future developments,” *IEEE Transactions on Radiation and Plasma Medical Sciences*, vol. 5, no. 4, pp. 441–452, 2021. doi: 10.1109/TRPMS.2020.3020212
- [78] S. S. Hsieh and M. Sjolín, “Digital count summing vs analog charge summing for photon counting detectors: A performance simulation study,” *Med. Phys.*, vol. 45, no. 9, pp. 4085–4093, 2018. doi: 10.1002/mp.13098
- [79] M. Touch, D. P. Clark, W. Barber, and C. T. Badea, “A neural network-based method for spectral distortion correction in photon counting X-ray CT,” *Phys. Med. Biol.*, vol. 61, no. 16, p. 6132, 2016. doi: 10.1088/0031-9155/61/16/6132
- [80] S. S. Hsieh, “Coincidence counters for charge sharing compensation in spectroscopic photon counting detectors,” *IEEE Trans. Med. Imaging*, vol. 39, no. 3, pp. 678–687, 2020. doi: 10.1109/TMI.2019.2933986
- [81] A. S. Wang, D. Harrison, V. Lobastov, and J. E. Tkaczyk, “Pulse pileup statistics for energy discriminating photon counting X-ray detectors,” *Med. Phys.*, vol. 38, no. 7, pp. 4265–4275, 2011. doi: 10.1118/1.3592932

- [82] 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,” *Med. Phys.*, vol. 41, no. 4, p. 041905, 2014. doi: 10.1118/1.4866890
- [83] S. Marupudi, J. A. Manus, M. U. Ghani, S. J. Glick, and B. Ghammraoui, “Evaluation of charge summing correction in CdTe-based photon-counting detectors for breast CT: performance metrics and image quality,” *J. Med. Imaging*, vol. 12, no. 1, pp. 013501–013501, 2025. doi: 10.1117/1.JMI.12.1.013501
- [84] K. Taguchi, D. R. Schaart, M. C. Goorden, and S. S. Hsieh, “Imaging performance of a LaBr₃:Ce scintillation detector for photon counting x-ray computed tomography: Simulation study,” *Med. Phys.*, vol. 52, no. 1, pp. 158–170, 2025. doi: 10.1002/mp.17436
- [85] S. S. Hsieh, K. Taguchi, M. C. Goorden, and D. R. Schaart, “The potential of scintillator-based photon counting detectors: Evaluation using Monte Carlo simulations,” in *Medical Imaging 2025: Physics of Medical Imaging*, ser. Progress in Biomedical Optics and Imaging - Proceedings of SPIE, J. M. Sabol, K. Li, and S. Abbaszadeh, Eds., vol. 13405. SPIE, 2025. doi: 10.1117/12.3045837 p. 1340537.
- [86] W. P. Dillon, “50th anniversary of computed tomography: past and future applications in clinical neuroscience,” *J. Med. Imaging*, vol. 8, no. 5, p. 052112, 2021. doi: 10.1117/1.JMI.8.5.052112
- [87] L. A. Shepp and B. F. Logan, “The Fourier reconstruction of a head section,” *IEEE Trans. Nucl. Sci.*, vol. 21, no. 3, pp. 21–43, Jun. 1974. doi: 10.1109/TNS.1974.6499235
- [88] L. A. Feldkamp, L. C. Davis, and J. W. Kress, “Practical cone-beam algorithm,” *J. Opt. Soc. Am. A*, vol. 1, no. 6, pp. 612–619, 1984. doi: 10.1364/JOSAA.1.000612
- [89] P. J. La Rivière and C. R. Crawford, “From EMI to AI: a brief history of commercial CT reconstruction algorithms,” *J. Med. Imaging*, vol. 8, no. 5, p. 052111, 2021. doi: 10.1117/1.JMI.8.5.052111
- [90] B. De Man and J. A. Fessler, “Statistical iterative reconstruction for X-Ray computed tomography,” in *Biomedical Mathematics: Promising Directions in Imaging, Therapy Planning and Inverse Problems*, Y. Censor, M. Jiang, and G. Wang, Eds. Madison, WI: Medical Physics Publishing, 2010, pp. 113–140. ISBN 978-1-930524-48-4
- [91] M. J. Willemink and P. B. Noël, “The evolution of image reconstruction for CT—from filtered back projection to artificial intelligence,” *European Radiology*, vol. 29, no. 5, pp. 2185–2195, 2019. doi: 10.1007/s00330-018-5810-7
- [92] F. K. Kopp, M. Catalano, D. Pfeiffer, E. J. Rummeny, and P. B. Noël, “Evaluation of a machine learning based model observer for x-ray CT,” in *Medical Imaging 2018: Image Perception, Observer Performance, and Technology Assessment*, ser. Proceedings of SPIE, C. K. Abbey and C. R. Mello-Thoms, Eds., vol. 10577, 2018. doi: 10.1117/12.2293582 p. 105770S.

- [93] H. Zhang, G. J. Gang, H. Dang, and J. W. Stayman, “Regularization analysis and design for prior-image-based x-ray CT reconstruction,” *IEEE Trans. Med. Imaging*, vol. 37, no. 12, pp. 2675–2686, 2018. doi: 10.1109/TMI.2018.2847250
- [94] G. C. Pereira, M. Traughber, and R. F. Muzic, “The role of imaging in radiation therapy planning: Past, present, and future,” *Biomed Res Int*, vol. 2014, p. 231090, 2014. doi: 10.1155/2014/231090
- [95] P. Alaei and E. Spezi, “Imaging dose from cone beam computed tomography in radiation therapy,” *Physica Medica*, vol. 31, no. 7, pp. 647–658, Nov. 2015. doi: 10.1016/j.ejmp.2015.06.003
- [96] D. Marin, D. T. Boll, A. Mileto, and R. C. Nelson, “State of the art: Dual-energy CT of the abdomen,” *Radiology*, vol. 271, no. 2, pp. 327–342, May 2014. doi: 10.1148/radiol.14131480 Publisher: Radiological Society of North America.
- [97] D. S. Rigie and P. J. La Rivière, “Optimizing spectral CT parameters for material classification tasks,” *Physics in Medicine and Biology*, vol. 61, no. 12, pp. 4599–4622, May 2016. doi: 10.1088/0031-9155/61/12/4599 Publisher: IOP Publishing.
- [98] M. F. Kruis, “Improving radiation physics, tumor visualisation, and treatment quantification in radiotherapy with spectral or dual-energy CT,” *J. Appl. Clin. Med. Phys.*, vol. 23, no. 1, p. e13468, 2022. doi: 10.1002/acm2.13468 Publisher: Wiley Online Library.
- [99] S. Sajja, Y. Lee, M. Eriksson, H. Nordström, A. Sahgal, M. Hashemi, J. G. Mainprize, and M. Ruschin, “Technical principles of dual-energy cone beam computed tomography and clinical applications for radiation therapy,” *Advances in Radiation Oncology*, vol. 5, no. 1, pp. 1–16, Jan. 2020. doi: 10.1016/j.adro.2019.07.013
- [100] G. Jadick and P. La Rivière, “Optimization of MV-kV dual-energy CT imaging for tomographic therapy,” in *Medical Imaging 2023: Physics of Medical Imaging*, vol. 12463, SPIE. SPIE, apr 2023. doi: 10.1117/12.2653674 pp. 557–566.
- [101] W. van Elmpt, G. Landry, M. Das, and F. Verhaegen, “Dual energy CT in radiotherapy: Current applications and future outlook,” *Radiotherapy and Oncology*, vol. 119, no. 1, pp. 137–144, Apr. 2016. doi: 10.1016/j.radonc.2016.02.026
- [102] G. Jadick, G. Schlafly, and P. J. La Rivière, “Dual-energy computed tomography imaging with megavoltage and kilovoltage X-ray spectra,” *J. Med. Imaging*, vol. 11, no. 2, pp. 023 501–023 501, 2024. doi: 10.1117/1.JMI.11.2.023501
- [103] O. M. Oderinde, S. M. Shirvani, P. D. Olcott, G. Kuduvalli, S. Mazin, and D. Larkin, “The technical design and concept of a PET/CT linac for biology-guided radiotherapy,” *Clinical and Translational Radiation Oncology*, vol. 29, pp. 106–112, Jul. 2021. doi: 10.1016/j.ctro.2021.04.003

- [104] D. Winkel, G. H. Bol, P. S. Kroon, B. van Asselen, S. S. Hackett, A. M. Werensteijn-Honingh, M. P. Intven, W. S. Eppinga, R. H. Tijssen, L. G. Kerkmeijer, H. C. de Boer, S. Mook, G. J. Meijer, J. Hes, M. Willemsen-Bosman, E. N. de Groot-van Breugel, I. M. Jürgenliemk-Schulz, and B. W. Raaymakers, “Adaptive radiotherapy: The Elekta Unity MR-linac concept,” *Clinical and Translational Radiation Oncology*, vol. 18, pp. 54–59, Sep. 2019. doi: 10.1016/j.ctro.2019.04.001
- [105] K. J. Ruchala, G. H. Olivera, J. M. Kapatoes, E. A. Schloesser, P. J. Reckwerdt, and T. R. Mackie, “Megavoltage CT imaging as a by-product of multileaf collimator leakage,” *Phys. Med. Biol.*, vol. 45, no. 7, p. N61, Jul. 2000. doi: 10.1088/0031-9155/45/7/401
- [106] F.-F. Yin, H. Guan, and W. Lu, “A technique for on-board CT reconstruction using both kilovoltage and megavoltage beam projections for 3D treatment verification,” *Med. Phys.*, vol. 32, no. 9, pp. 2819–2826, 2005. doi: 10.1118/1.1997307
- [107] J. Zhang and F.-F. Yin, “Minimizing image noise in on-board CT reconstruction using both kilovoltage and megavoltage beam projections,” *Med. Phys.*, vol. 34, no. 9, pp. 3665–3673, 2007. doi: 10.1118/1.2768862
- [108] T. R. Mackie, J. Balog, K. Ruchala, D. Shepard, S. Aldridge, E. Fitchard, P. Reckwerdt, G. Olivera, T. McNutt, and M. Mehta, “Tomotherapy,” *Seminars in Radiation Oncology*, vol. 9, no. 1, pp. 108–117, Jan. 1999. doi: 10.1016/S1053-4296(99)80058-7
Radiation Therapy Treatment Optimization.
- [109] H. Li, B. Liu, and F.-F. Yin, “Generation of virtual monochromatic CBCT from dual kV/MV beam projections,” *Med. Phys.*, vol. 40, no. 12, p. 121910, 2013. doi: 10.1118/1.4824324
- [110] W. Luo, S. Yoo, Q. J. Wu, Z. Wang, and F.-F. Yin, “Analysis of image quality for real-time target tracking using simultaneous kV-MV imaging,” *Med. Phys.*, vol. 35, no. 12, pp. 5501–5509, 2008. doi: 10.1118/1.3002313
- [111] W. Mao, N. Riaz, L. Lee, R. Wiersma, and L. Xing, “A fiducial detection algorithm for real-time image guided IMRT based on simultaneous MV and kV imaging,” *Med. Phys.*, vol. 35, no. 8, pp. 3554–3564, 2008. doi: 10.1118/1.2953563
- [112] E. Pearson, X. Pan, and C. Pelizzari, “Dual-energy (MV/kV) CT with probabilistic attenuation mapping for IGRT applications,” in *Medical Imaging 2015: Physics of Medical Imaging*, vol. 9412. SPIE, Mar. 2015. doi: 10.1117/12.2082068 pp. 1391–1398.
- [113] H. Wertz, D. Stsepankou, M. Blessing, M. Rossi, C. Knox, K. Brown, U. Gros, J. Boda-Heggemann, C. Walter, J. Hesser, F. Lohr, and F. Wenz, “Fast kilovoltage/megavoltage (kVMV) breathhold cone-beam CT for image-guided radiotherapy of lung cancer,” *Phys. Med. Biol.*, vol. 55, no. 15, pp. 4203–4217, 2010. doi: 10.1088/0031-9155/55/15/001

- [114] M. Blessing, D. Stsepankou, H. Wertz, A. Arns, F. Lohr, J. Hesser, and F. Wenz, “Breath-hold target localization with simultaneous kilovoltage/megavoltage cone-beam computed tomography and fast reconstruction,” *Int. J. Radiat. Oncol. Biol. Phys.*, vol. 78, no. 4, pp. 1219–1226, 2010. doi: 10.1016/j.ijrobp.2010.01.030
- [115] H. Li, W. Giles, L. Ren, J. Bowsher, and F.-F. Yin, “Implementation of dual-energy technique for virtual monochromatic and linearly mixed CBCTs,” *Med. Phys.*, vol. 39, no. 10, pp. 6056–6064, 2012. doi: 10.1118/1.4752212
- [116] H. Keller, M. Glass, R. Hinderer, K. Ruchala, R. Jeraj, G. Olivera, and T. R. Mackie, “Monte Carlo study of a highly efficient gas ionization detector for megavoltage imaging and image-guided radiotherapy,” *Med. Phys.*, vol. 29, no. 2, pp. 165–175, 2002. doi: 10.1118/1.1445414
- [117] K. J. Ruchala, G. H. Olivera, E. A. Schloesser, and T. R. Mackie, “Megavoltage CT on a tomotherapy system,” *Physics in Medicine and Biology*, vol. 44, no. 10, pp. 2597–2621, Oct. 1999. doi: 10.1088/0031-9155/44/10/316
- [118] S. L. Meeks, J. F. Harmon Jr., K. M. Langen, T. R. Willoughby, T. H. Wagner, and P. A. Kupelian, “Performance characterization of megavoltage computed tomography imaging on a helical tomotherapy unit,” *Med. Phys.*, vol. 32, pp. 2673–2681, 2005. doi: 10.1118/1.1990289
- [119] K. Schombourg, F. Bochud, and R. Moeckli, “Stability of the Helical TomoTherapy Hi-Art II detector for treatment beam irradiations,” *J. Appl. Clin. Med. Phys.*, vol. 15, no. 6, pp. 119–127, 2014. doi: 10.1120/jacmp.v15i6.4897
- [120] R. Jeraj, T. R. Mackie, J. Balog, G. Olivera, D. Pearson, J. Kapatoes, K. Ruchala, and P. Reckwerdt, “Radiation characteristics of helical tomotherapy,” *Med. Phys.*, vol. 31, no. 2, pp. 396–404, 2004. doi: 10.1118/1.1639148
- [121] T. R. Mackie, “History of tomotherapy,” *Phys. Med. Biol.*, vol. 51, no. 13, pp. R427–R453, 2006. doi: 10.1088/0031-9155/51/13/R24
- [122] S. M. Kay, *Fundamentals of statistical signal processing: estimation theory*. Prentice-Hall, Inc., 1993.
- [123] W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. W. Tsui, “4D XCAT phantom for multimodality imaging research,” *Med. Phys.*, vol. 37, no. 9, pp. 4902–4915, 2010. doi: 10.1118/1.3480985
- [124] R. L. Siddon, “Fast calculation of the exact radiological path for a three-dimensional CT array,” *Med. Phys.*, vol. 12, no. 2, pp. 252–255, 1985. doi: 10.1118/1.595715
- [125] A. Nouri and C. Wen, “Stainless steels in orthopedics,” in *Structural Biomaterials*, ser. Woodhead Publishing Series in Biomaterials, C. Wen, Ed. Woodhead Publishing, 2021, pp. 67–101. ISBN 978-0-12-818831-6

- [126] F. Zivic, S. Affatato, M. Trajanovic, M. Schnabelrauch, N. Grujovic, and K. L. Choy, *Biomaterials in clinical practice: advances in clinical research and medical devices*. Springer, 2017.
- [127] A. Goharian and M. Abdullah, “Bioinert metals (stainless steel, titanium, cobalt chromium),” *Trauma plating systems*, vol. 115, 2017.
- [128] D. S. Rigie and P. La Rivière, “An efficient material decomposition method using the Gauss-Newton algorithm,” in *2015 IEEE Medical Imaging Conference*. IEEE, 2015. doi: 10.13140/RG.2.1.1680.3289
- [129] A. Wunderlich and F. Noo, “Achieving uniform noise in direct fan-beam CT reconstruction through bowtie filter design,” in *2007 IEEE Nuclear Science Symposium Conference Record*, vol. 6, 2007. doi: 10.1109/NSSMIC.2007.4437083 pp. 4379–4382.
- [130] G. Jadick, E. Abadi, B. Harrawood, S. Sharma, W. P. Segars, and E. Samei, “A scanner-specific framework for simulating CT images with tube current modulation,” *Phys. Med. Biol.*, vol. 66, no. 18, p. 185010, 2021. doi: 10.1088/1361-6560/ac2269
- [131] D. S. Rigie and P. J. La Rivière, “Joint reconstruction of multi-channel, spectral CT data via constrained total nuclear variation minimization,” *Phys. Med. Biol.*, vol. 60, no. 5, pp. 1741–1762, 2015. doi: 10.1088/0031-9155/60/5/1741
- [132] D. Rigie, P. La Rivière, and A. Petschke, “Spectral X-ray computed tomography reconstruction using a vectorial total variation,” Jun. 6 2017, uS Patent 9,672,638.
- [133] J. Star-Lack, D. Shedlock, D. Swahn, D. Humber, A. Wang, H. Hirsh, G. Zentai, D. Sawkey, I. Kruger, M. Sun, E. Abel, G. Virshup, M. Shin, and R. Fahrig, “A piecewise-focused high DQE detector for MV imaging,” *Med. Phys.*, vol. 42, no. 9, pp. 5084–5099, 2015. doi: 10.1118/1.4927786
- [134] J. Rottmann, D. Morf, R. Fueglistaller, G. Zentai, J. Star-Lack, and R. Berbeco, “A novel EPID design for enhanced contrast and detective quantum efficiency,” *Phys. Med. Biol.*, vol. 61, no. 17, p. 6297, 2016. doi: 10.1088/0031-9155/61/17/6297
- [135] Y.-H. Hu, D. Shedlock, A. Wang, J. Rottmann, P. Baturin, M. Myronakis, P. Huber, R. Fueglistaller, M. Shi, D. Morf, J. Star-Lack, and R. I. Berbeco, “Characterizing a novel scintillating glass for application to megavoltage cone-beam computed tomography,” *Med. Phys.*, vol. 46, no. 3, pp. 1323–1330, 2019. doi: 10.1002/mp.13355
- [136] T. C. Harris, J. Seco, D. Ferguson, M. Lehmann, P. Huber, M. Shi, M. W. Jacobson, I. Valencia Lozano, M. Myronakis, P. Baturin, R. Fueglistaller, D. Morf, and R. I. Berbeco, “Clinical translation of a new flat-panel detector for beam’s-eye-view imaging,” *Phys. Med. Biol.*, vol. 65, no. 22, p. 225004, 2020. doi: 10.1088/1361-6560/abb571
- [137] T. Harris, J. Seco, D. Ferguson, M. Jacobson, M. Myronakis, I. V. Lozano, M. Lehmann, P. Huber, R. Fueglistaller, D. Morf, J. D. Mamon, H J abd Mancias,

- N. E. Martin, and R. I. Berbeco, “Improvements in beam’s eye view fiducial tracking using a novel multilayer imager,” *Phys. Med. Biol.*, vol. 66, no. 15, p. 155007, 2021. doi: 10.1088/1361-6560/ac1246
- [138] T. C. Harris, M. W. Jacobson, M. E. Myronakis, M. Lehmann, P. Huber, D. Morf, I. Ozoemelum, Y.-H. Hu, D. Ferguson, R. Fueglistaller, P. Corral Arroyo, and R. I. Berbeco, “Impact of a novel multilayer imager on metal artifacts in MV-CBCT,” *Physics in Medicine and Biology*, vol. 68, no. 14, p. 145009, 2023. doi: 10.1088/1361-6560/ace09a
- [139] J. Zhu, T. Su, X. Zhang, J. Yang, D. Mi, Y. Zhang, X. Gao, H. Zheng, D. Liang, and Y. Ge, “Feasibility study of three-material decomposition in dual-energy cone-beam CT imaging with deep learning,” *Phys. Med. Biol.*, vol. 67, no. 14, p. 145012, 2022. doi: 10.1088/1361-6560/ac7b09
- [140] T. Lyu, W. Zhao, Y. Zhu, Z. Wu, Y. Zhang, Y. Chen, L. Luo, S. Li, and L. Xing, “Estimating dual-energy CT imaging from single-energy CT data with material decomposition convolutional neural network,” *Medical image analysis*, vol. 70, p. 102001, 2021. doi: 10.1016/j.media.2021.102001
- [141] E. Abadi, B. Harrawood, S. Sharma, A. Kapadia, W. P. Segars, and E. Samei, “DukeSim: a realistic, rapid, and scanner-specific simulation framework in computed tomography,” *IEEE Trans. Med. Imaging*, vol. 38, no. 6, pp. 1457–1465, 2019. doi: 10.1109/TMI.2018.2886530
- [142] S. Sharma, E. Abadi, A. Kapadia, W. P. Segars, and E. Samei, “A GPU-accelerated framework for rapid estimation of scanner-specific scatter in CT for virtual imaging trials,” *Phys. Med. Biol.*, vol. 66, no. 7, p. 075004, 2021. doi: 10.1088/1361-6560/abeb32
- [143] M. Kaur, J. Luce, M. Lehmann, D. Morf, L. Zhu, H. Kang, M. Walczak, M. M. Harkenrider, and J. C. Roeske, “Effect of scattered megavoltage x-rays on markerless tumor tracking using dual energy kilovoltage imaging,” *J. Appl. Clin. Med. Phys.*, p. e13993, 2023. doi: 10.1002/acm2.13993
- [144] K. Y. Seet, A. Barghi, S. Yartsev, and J. Van Dyk, “The effects of field-of-view and patient size on CT numbers from cone-beam computed tomography,” *Phys. Med. Biol.*, vol. 54, no. 20, p. 6251, 2009. doi: 10.1088/0031-9155/54/20/014
- [145] G. Jadick, M. Ventura, and P. J. La Rivière, “Utility of photon-counting detectors for MV-kV dual-energy computed tomography imaging,” *J. Med. Imaging*, vol. 11, no. S1, pp. S12811–S12811, 2024. doi: 10.1117/1.JMI.11.S1.S12811
- [146] G. Paulo, J. Damilakis, V. Tsapaki, A. A. Schegerer, J. Repussard, W. Jaschke, G. Frija, M. Hierath, J. Clark, and European Society of Radiology, “Diagnostic reference levels based on clinical indications in computed tomography: a literature review,” *Insights into imaging*, vol. 11, pp. 1–9, 2020. doi: 10.1186/s13244-020-00899-y

- [147] J. Sillanpaa, J. Chang, G. Mageras, E. Yorke, F. De Arruda, K. E. Rosenzweig, P. Munro, E. Seppi, J. Pavkovich, and H. Amols, “Low-dose megavoltage cone-beam computed tomography for lung tumors using a high-efficiency image receptor,” *Med. Phys.*, vol. 33, no. 9, pp. 3489–3497, 2006. doi: 10.1118/1.2222075
- [148] M. Ventura, G. Jadick, and P. J. La Rivière, “Comparison of energy-integrating detectors and photon-counting detectors for MV-kV dual-energy imaging on a tomographic therapy system,” in *Medical Imaging 2024: Physics of Medical Imaging*, vol. 12925, SPIE. SPIE, 2024. doi: 10.1117/12.3006854 pp. 542–546.
- [149] B. A. Groh, J. H. Siewerdsen, D. G. Drake, J. W. Wong, and D. A. Jaffray, “A performance comparison of flat-panel imager-based MV and kV cone-beam CT,” *Med. Phys.*, vol. 29, no. 6, pp. 967–975, 2002. doi: 10.1118/1.1477234
- [150] Y. El-Mohri, L. E. Antonuk, Q. Zhao, R. B. Choroszuca, H. Jiang, and L. Liu, “Low-dose megavoltage cone-beam CT imaging using thick, segmented scintillators,” *Phys. Med. Biol.*, vol. 56, no. 6, p. 1509, 2011. doi: 10.1088/0031-9155/56/6/001
- [151] A. P. Shah, K. M. Langen, K. J. Ruchala, A. Cox, P. A. Kupelian, and S. L. Meeks, “Patient dose from megavoltage computed tomography imaging,” *Int. J. Radiat. Oncol. Biol. Phys.*, vol. 70, no. 5, pp. 1579–1587, 2008. doi: 10.1016/j.ijrobp.2007.11.048
- [152] C. H. McCollough, K. Rajendran, S. Leng, L. Yu, J. G. Fletcher, K. Stierstorfer, and T. G. Flohr, “The technical development of photon-counting detector CT,” *European radiology*, vol. 33, no. 8, pp. 5321–5330, 2023. doi: 10.1007/s00330-023-09545-9
- [153] C. Sundberg, M. Danielsson, and M. Persson, “Compton coincidence in silicon photon-counting CT detectors,” *J. Med. Imaging*, vol. 9, no. 1, p. 013501, Jan. 2022. doi: 10.1117/1.JMI.9.1.013501
- [154] Y. Gu, J. L. Matteson, R. T. Skelton, A. C. Deal, E. A. Stephan, F. Duttweiler, T. M. Gasaway, and C. S. Levin, “Study of a high resolution, 3-D positioning cross-strip Cadmium Zinc Telluride detector for PET,” in *IEEE*, Oct. 2008. doi: 10.1109/NSS-MIC.2008.4774098 pp. 3596–3603.
- [155] S. Abbaszadeh, Y. Gu, P. D. Reynolds, and C. S. Levin, “Characterization of a sub-assembly of 3D position sensitive cadmium zinc telluride detectors and electronics from a sub-millimeter resolution PET system,” *Phys. Med. Biol.*, vol. 61, no. 18, p. 6733, Aug. 2016. doi: 10.1088/0031-9155/61/18/6733
- [156] Y. Wang and S. Abbaszadeh, “Reduction of the low energy threshold through new modular data acquisition electronics for cross-strip cadmium zinc telluride (CZT) based PET system,” in *2021 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, Oct. 2021. doi: 10.1109/NSS/MIC44867.2021.9875680 pp. 1–3.
- [157] M. Diwakar and M. Kumar, “A review on CT image noise and its denoising,” *Biomedical Signal Processing and Control*, vol. 42, pp. 73–88, 2018. doi: 10.1016/j.bspc.2018.01.010

- [158] H. Lu, T. Hsiao, X. Li, and Z. Liang, “Noise properties of low-dose CT projections and noise treatment by scale transformations,” in *2001 IEEE Nuclear Science Symposium Conference Record (Cat. No. 01CH37310)*, vol. 3. IEEE, 2001, pp. 1662–1666.
- [159] J. B. Solomon, X. Li, and E. Samei, “Relating noise to image quality indicators in CT examinations with tube current modulation,” *Am. J. Roentgenol*, vol. 200, no. 3, pp. 592–600, 2013. doi: 10.2214/AJR.12.8580
- [160] S. Leng, L. Yu, J. Wang, J. G. Fletcher, C. A. Mistretta, and C. H. McCollough, “Noise reduction in spectral CT: Reducing dose and breaking the trade-off between image noise and energy bin selection,” *Med. Phys.*, vol. 38, no. 9, pp. 4946–4957, 2011. doi: 10.1118/1.3609097
- [161] J. Ma, Z. Liang, Y. Fan, Y. Liu, J. Huang, W. Chen, and H. Lu, “Variance analysis of X-ray CT sinograms in the presence of electronic noise background,” *Med. Phys.*, vol. 39, no. 7Part1, pp. 4051–4065, 2012. doi: 10.1118/1.4722751
- [162] A. S. Wang and N. J. Pelc, “Impact of photon counting detector spectral response on dual energy techniques,” in *Medical Imaging 2010: Physics of Medical Imaging*, vol. 7622. SPIE, 2010, pp. 1148–1156.
- [163] D. M. Paganin and D. Pelliccia, “X-ray phase-contrast imaging: a broad overview of some fundamentals,” *Adv. Imaging Electron Phys.*, vol. 218, pp. 63–158, 2021. doi: 10.1016/bs.aiep.2021.04.002
- [164] G. Jadick and P. La Rivière, “Modeling propagation-based X-ray phase-contrast imaging: validity of the projection approximation,” in *Proc. Virtual Imaging Trials in Medicine*, 2024. doi: 10.48550/arXiv.2405.05359 pp. 68–72.
- [165] K. Li, M. Wojcik, and C. Jacobsen, “Multislice does it all—calculating the performance of nanofocusing X-ray optics,” *Opt. Express*, vol. 25, no. 3, pp. 1831–1846, 2017. doi: 10.1364/OE.25.001831
- [166] M. Du, S. Kandel, J. Deng, X. Huang, A. Demortiere, T. T. Nguyen, R. Tucoulou, V. De Andrade, Q. Jin, and C. Jacobsen, “Adorym: A multi-platform generic X-ray image reconstruction framework based on automatic differentiation,” *Opt. Express*, vol. 29, no. 7, pp. 10 000–10 035, 2021. doi: 10.1364/OE.418296
- [167] R. P. Xian, C. L. Walsh, S. E. Verleden, W. L. Wagner, A. Bellier, S. Marussi, M. Ackermann, D. D. Jonigk, J. Jacob, P. D. Lee, and P. Tafforeau, “A multiscale X-ray phase-contrast tomography dataset of a whole human left lung,” *Sci. Data*, vol. 9, no. 1, p. 264, 2022. doi: 10.1038/s41597-022-01353-y
- [168] S. Zabler, P. Cloetens, J.-P. Guigay, J. Baruchel, and M. Schlenker, “Optimization of phase contrast imaging using hard X rays,” *Rev. Sci. Instrum.*, vol. 76, no. 7, 2005. doi: 10.1063/1.1960797

- [169] M. A. Yakovlev, D. J. Vanselow, M. S. Ngu, C. R. Zaino, S. R. Katz, Y. Ding, D. Parkinson, S. Y. Wang, K. C. Ang, P. La Rivière, and K. Cheng, “A wide-field micro-computed tomography detector: micron resolution at half-centimetre scale,” *J. Synchrotron Rad.*, vol. 29, no. 2, pp. 505–514, 2022. doi: 10.1107/S160057752101287X
- [170] S. C. Mayo, P. R. Miller, S. Wilkins, T. J. Davis, D. Gao, T. E. Gureyev, D. Paganin, D. Parry, A. Pogany, and A. W. Stevenson, “Quantitative X-ray projection microscopy: phase-contrast and multi-spectral imaging,” *J. Microsc.*, vol. 207, no. 2, pp. 79–96, 2002. doi: 10.1046/j.1365-2818.2002.01046.x
- [171] M. Viermetz, N. Gustschin, C. Schmid, J. Haeusele, M. von Teuffenbach, P. Meyer, F. Bergner, T. Lasser, R. Proksa, T. Koehler, and F. Pfeiffer, “Dark-field computed tomography reaches the human scale,” *Proceedings of the National Academy of Sciences*, vol. 119, no. 8, p. e2118799119, 2022. doi: 10.1073/pnas.2118799119
- [172] G. Jadick and P. La Rivière, “Cramér-Rao lower bound in the context of spectral X-ray imaging with propagation-based phase contrast,” in *Medical Imaging 2024: Physics of Medical Imaging*, vol. 12925. SPIE, 2024. doi: 10.1117/12.3006282 pp. 86–93.
- [173] G. Jadick and P. L. Rivière, “Basis material Cramér–Rao lower bound: the theoretical utility of propagation-based phase contrast,” in *Medical Imaging 2026: Physics of Medical Imaging*, ser. Proceedings of SPIE, K. Li, S. Abbaszadeh, and A. Ganguly, Eds., vol. 13924, 2026. doi: 10.1117/12.3087671 pp. 114–119.
- [174] K. A. Nugent, “X-ray noninterferometric phase imaging: a unified picture,” *J. Opt. Soc. Am. A*, vol. 24, no. 2, pp. 536–547, 2007. doi: 10.1364/JOSAA.24.000536
- [175] K. Rangarajan, A. Gupta, S. Dasgupta, U. Marri, A. K. Gupta, S. Hari, S. Banerjee, and C. Arora, “Ultra-high resolution, multi-scale, context-aware approach for detection of small cancers on mammography,” *Scientific reports*, vol. 12, no. 1, p. 11622, 2022. doi: 10.1038/s41598-022-15259-7
- [176] 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*, vol. 10948. SPIE, 2019. doi: 10.1117/12.2513513 pp. 1157–1164.
- [177] R. Brunskog, M. Persson, Z. Jin, and M. Danielsson, “Experimental evaluation of a micron-resolution CT detector,” in *Medical Imaging 2024: Physics of Medical Imaging*, vol. 12925. SPIE, 2024. doi: 10.1117/12.2692858 p. 129250B.
- [178] F. Schaff, K. S. Morgan, D. M. Paganin, and M. J. Kitchen, “Spectral X-ray imaging: Conditions under which propagation-based phase-contrast is beneficial,” *Phys. Med. Biol.*, vol. 65, no. 20, p. 205006, 2020. doi: 10.1088/1361-6560/aba318
- [179] G. Jadick and P. J. La Rivière, “Material decomposition of weakly absorptive structures with spectral X-ray phase-contrast CT,” in *2025 IEEE Nuclear Science Symposium (NSS), Medical Imaging Conference (MIC) and Room*

- Temperature Semiconductor Detector Conference (RTSD)*. IEEE, 2025. doi: 10.1109/NSS/MIC/RTSD57106.2025.11287842 pp. 1–1.
- [180] G. Jadick and P. La Rivière, “Optimization-based phase retrieval for material decomposition with multi-energy computed tomography,” in *18th International Meeting on Fully 3D Image Reconstruction in Radiology and Nuclear Medicine*, May 2025. doi: 10.48550/arXiv.2508.12509
 - [181] S. Vrbaški, S. Donato, E. Longo, P. Segars, E. Samei, R. Longo, and A. Contillo, “Synchrotron split-beam method for clinical photon-counting computed tomography,” *Physica Med.*, vol. 142, p. 105734, 2026. doi: 10.1016/j.ejmp.2026.105734
 - [182] A. M. Hernandez and J. M. Boone, “Tungsten anode spectral model using interpolating cubic splines: unfiltered X-ray spectra from 20 kV to 640 kV,” *Med. Phys.*, vol. 41, no. 4, p. 042101, 2014. doi: 10.1118/1.4866216
 - [183] D. Deb, G.-J. Both, E. Bezzam, A. Kohli, S. Yang, A. Chaware, C. Allier, C. Cai, G. Anderberg, M. H. Eybposh, M. C. Schneider, R. Heintzmann, F. A. Rivera-Sanchez, C. Simmerer, G. Meng, J. Tormes-Vaquero, S. Han, S. C. Shanmugavel, T. Maruvada, X. Yang, Y. Kim, B. Diederich, C. Joo, L. Waller, N. J. Durr, N. C. Pégard, P. J. La Rivière, R. Horstmeyer, S. Chowdhury, and S. C. Turaga, “Chromatix: a differentiable, GPU-accelerated wave-optics library,” *bioRxiv*, may 2026. doi: 10.1101/2025.04.29.651152
 - [184] J. Tanguay, H. K. Kim, and I. A. Cunningham, “A theoretical comparison of x-ray angiographic image quality using energy-dependent and conventional subtraction methods,” *Med. Phys.*, vol. 39, no. 1, pp. 132–142, 2012. doi: 10.1118/1.3658728
 - [185] C. S. Burton, J. R. Mayo, and I. A. Cunningham, “Energy subtraction angiography is comparable to digital subtraction angiography in terms of iodine Rose SNR,” *Med. Phys.*, vol. 43, no. 11, pp. 5925–5933, 2016. doi: 10.1118/1.4962651
 - [186] S. Aubert, I. A. Cunningham, and J. Tanguay, “Theoretical comparison of energy-resolved and digital-subtraction angiography,” *Med. Phys.*, vol. 49, no. 11, pp. 6885–6902, 2022. doi: 10.1002/mp.15973
 - [187] C. S. Burton, J. R. Mayo, and I. A. Cunningham, “Theoretical and experimental comparison of image signal and noise for dual-energy subtraction angiography and conventional x-ray angiography,” in *Medical Imaging 2015: Physics of Medical Imaging*, ser. Proceedings of SPIE, C. Hoeschen and D. Kontos, Eds., vol. 9412. SPIE, 2015. doi: 10.1117/12.2082724 p. 941219.
 - [188] S. Aubert and J. Tanguay, “Signal-difference-to-noise comparison of temporal subtraction, kV-switching dual-energy and photon-counting dual-energy x-ray angiography,” *Med. Phys.*, vol. 50, no. 12, pp. 7400–7414, 2023. doi: 10.1002/mp.16800

- [189] S. Aubert and J. Tanguay, “Experimental optimization of single-exposure dual-energy angiography with photon-counting x-ray detectors,” *Med. Phys.*, vol. 50, no. 2, pp. 763–777, 2023. doi: 10.1002/mp.16079
- [190] S. Liao, X. Zhang, X. Jiang, M. Tivnan, J. W. Stayman, and G. J. Gang, “Joint temporal and spectral processing for improved digital subtraction angiography using photon-counting detectors,” *IEEE Transactions on Biomedical Engineering*, vol. 72, no. 12, pp. 3501–3514, 2025. doi: 10.1109/TBME.2025.3570925
- [191] G. Jadick and P. La Rivière, “Accuracy of propagation-based phase-contrast CT under the projection approximation,” in *8th International Conference on Image Formation in X-Ray Computed Tomography*, 2024. doi: 10.48550/arXiv.2508.12505
- [192] K. S. Morgan, K. K. W. Siu, and D. M. Paganin, “The projection approximation versus an exact solution for X-ray phase contrast imaging, with a plane wave scattered by a dielectric cylinder,” *Opt. Commun.*, vol. 283, no. 23, pp. 4601–4608, 2010. doi: 10.1016/j.optcom.2010.07.012
- [193] K. S. Morgan, K. K. W. Siu, and D. M. Paganin, “The projection approximation and edge contrast for X-ray propagation-based phase contrast imaging of a cylindrical edge,” *Opt. Express*, vol. 18, no. 10, pp. 9865–9878, 2010. doi: 10.1364/OE.18.009865
- [194] F. Pfeiffer, “X-ray ptychography,” *Nat. Photon.*, vol. 12, no. 1, pp. 9–17, 2018. doi: 10.1038/s41566-017-0072-5
- [195] A. M. Maiden, M. J. Humphry, and J. M. Rodenburg, “Ptychographic transmission microscopy in three dimensions using a multi-slice approach,” *J. Opt. Soc. Am. A*, vol. 29, no. 8, pp. 1606–1614, 2012. doi: 10.1364/JOSAA.29.001606
- [196] E. H. R. Tsai, I. Usov, A. Diaz, A. Menzel, and M. Guizar-Sicairos, “X-ray ptychography with extended depth of field,” *Opt. Express*, vol. 24, no. 25, pp. 29 089–29 108, 2016. doi: 10.1364/OE.24.029089
- [197] P. Li and A. Maiden, “Multi-slice ptychographic tomography,” *Sci. Rep.*, vol. 8, p. 2049, 2018. doi: 10.1038/s41598-018-20530-x
- [198] M. Du, Y. S. Nashed, S. Kandel, D. Gürsoy, and C. Jacobsen, “Three dimensions, two microscopes, one code: Automatic differentiation for X-ray nanotomography beyond the depth of focus limit,” *Sci. Adv.*, vol. 6, no. 13, p. eaay3700, 2020. doi: 10.1126/sci-adv.aay3700
- [199] D. Luzhbin, C.-C. Chang, S.-J. Chang, and J. Wu, “Dual-energy tissue cancellation in mammography for improved detection of microcalcifications and neoplasms: A phantom study,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1025, p. 166062, 2022. doi: 10.1016/j.nima.2021.166062

- [200] R. Ballabriga, M. Campbell, E. Heijne, X. Llopart, L. Tlustos, and W. Wong, “Medipix3: A 64k pixel detector readout chip working in single photon counting mode with improved spectrometric performance,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 633, pp. S15–S18, 2011. doi: 10.1016/j.nima.2010.06.108 11th International Workshop on Radiation Imaging Detectors (IWORID).
- [201] M. Das and Z. Liang, “Spectral X-ray phase contrast imaging for single-shot retrieval of absorption, phase, and differential-phase imagery,” *Opt. Lett.*, vol. 39, no. 21, pp. 6343–6346, 2014. doi: 10.1364/OL.39.006343
- [202] A. Koch, C. Raven, P. Spanne, and A. Snigirev, “X-ray imaging with submicrometer resolution employing transparent luminescent screens,” *J. Opt. Soc. Am. A*, vol. 15, no. 7, pp. 1940–1951, 1998. doi: 10.1364/JOSAA.15.001940
- [203] 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,” *Opt. Express*, vol. 25, no. 26, pp. 33 543–33 558, 2017. doi: 10.1364/OE.25.033543
- [204] K. A. Mohan, J.-B. Forien, V. Sridhar, J. Cuadra, and D. Parkinson, “Maximum likelihood based phase-retrieval using fresnel propagation forward models with optional constraints,” *IEEE Transactions on Computational Imaging*, vol. 9, pp. 1111–1125, 2023. doi: 10.1109/TCI.2023.3332007
- [205] J. A. Fessler, “Statistical image reconstruction methods for transmission tomography,” in *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis*, M. Sonka and J. M. Fitzpatrick, Eds. Bellingham, WA: SPIE, 2000, pp. 1–70.
- [206] A. Padole, R. D. A. Khawaja, M. K. Kalra, and S. Singh, “Ct radiation dose and iterative reconstruction techniques,” *AJR Am J Roentgenol*, vol. 204, no. 4, pp. W384–W392, 2015. doi: 10.2214/AJR.14.13241
- [207] L. L. Geyer, U. J. Schoepf, F. G. Meinel, J. W. Nance, G. Bastarrika, J. A. Leipsic, N. S. Paul, M. Rengo, A. Laghi, and C. N. De Cecco, “State of the art: Iterative ct reconstruction techniques,” *Radiology*, vol. 276, no. 2, pp. 339–357, 2015. doi: 10.1148/radiol.2015132766
- [208] W. Stiller, “Basics of iterative reconstruction methods in computed tomography: A vendor-independent overview,” *Eur J Radiol*, vol. 109, pp. 147–154, 2018. doi: 10.1016/j.ejrad.2018.10.025
- [209] L. Hehn, K. Morgan, P. Bidola, W. Noichl, R. Gradl, M. Dierolf, P. B. Noël, and F. Pfeiffer, “Nonlinear statistical iterative reconstruction for propagation-based phase-contrast tomography,” *APL Bioengineering*, vol. 2, no. 1, p. 016105, 2018. doi: 10.1063/1.4990387

- [210] L. Hehn, R. Gradl, M. Dierolf, K. S. Morgan, D. M. Paganin, and F. Pfeiffer, “Model-based iterative reconstruction for propagation-based phase-contrast X-ray CT including models for the source and the detector,” *IEEE Transactions on Medical Imaging*, vol. 39, no. 6, pp. 1975–1987, 2020. doi: 10.1109/TMI.2019.2962615
- [211] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations*, 2015. doi: 10.48550/arXiv.1412.6980 ArXiv:1412.6980.
- [212] S. Ali, M. Du, M. F. Adams, B. Smith, and C. Jacobsen, “Comparison of distributed memory algorithms for X-ray wave propagation in inhomogeneous media,” *Opt. Express*, vol. 28, no. 20, pp. 29 590–29 618, 2020. doi: 10.1364/OE.400240
- [213] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: From error visibility to structural similarity,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, Apr. 2004. doi: 10.1109/TIP.2003.819861
- [214] S. Huhn, L. M. Lohse, J. Lucht, and T. Salditt, “Fast algorithms for nonlinear and constrained phase retrieval in near-field X-ray holography based on Tikhonov regularization,” *Optics Express*, vol. 30, no. 18, pp. 32 871–32 886, 2022. doi: 10.1364/OE.462368
- [215] G. Jadick and P. La Rivière, “Spectral X-ray phase-contrast imaging of microcalcifications via automatic-differentiation material decomposition,” in *Medical Imaging 2026: Physics of Medical Imaging*, ser. Proceedings of SPIE, K. Li, S. Abbaszadeh, and A. Ganguly, Eds., vol. 13924, 2026. doi: 10.1117/12.3088061 pp. 135–140.
- [216] O. G. Ernst and M. J. Gander, “Why it is difficult to solve helmholtz problems with classical iterative methods,” in *Numerical Analysis of Multiscale Problems*, ser. Lecture Notes in Computational Science and Engineering, I. G. Graham, T. Y. Hou, O. Lakkis, and R. Scheichl, Eds. Berlin, Heidelberg: Springer, 2012, vol. 83, pp. 325–363. ISBN 978-3-642-22061-6
- [217] K. A. Mohan, D. Y. Parkinson, and J. A. Cuadra, “Constrained non-linear phase retrieval for single distance X-ray phase contrast tomography,” *arXiv*, 2020, arXiv preprint; DOI not found.
- [218] J. N. Ahlers, K. M. Pavlov, M. J. Kitchen, S. A. Harker, E. J. Pryor, J. A. Pollock, M. K. Croughan, Y. Y. How, M.-C. Zdora, L. F. Costello, D. W. O’Connell, C. Hall, and K. S. Morgan, “Single-exposure x-ray dark-field imaging via a dual-energy propagation-based setup,” *Optics Letters*, vol. 50, no. 7, pp. 2171–2174, 2025. doi: 10.1364/OL.553310
- [219] 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, no. 3, pp. 1413–1421, 2015. doi: 10.1109/TNS.2015.2425227

- [220] C. N. Luna, I. S. Reiser, and G. J. Gang, “Detection and classification of breast cancer and calcification using multi-contrast X-ray phase tomosynthesis and a deep learning classifier,” in *Medical Imaging 2025: Physics of Medical Imaging*, vol. 13401, SPIE. SPIE, feb 2025. doi: 10.1117/12.3043359 pp. 451–454.
- [221] L. Frappart, I. Remy, H. C. Lin, A. Bremond, D. Raudrant, B. Grousseau, and J. Vauzelle, “Different types of microcalcifications observed in breast pathology: correlations with histopathological diagnosis and radiological examination of operative specimens,” *Virchows Archiv A*, vol. 410, no. 3, pp. 179–187, 1987. doi: 10.1007/BF00710823
- [222] G. Jadick, C. Riggs, and P. La Rivière, “Quantitative forward modeling of propagation-based x-ray phase-contrast imaging at clinical scale,” *Med. Phys.*, vol. 52, no. 10, p. e70059538, 2025, 2025 Annual Meeting.