Title

Sparse Primary Sampling: A Novel Scatter Estimation and Correction Strategy for Cone-beam Computed Tomography

ISBN

9798189297939

Author

Li, Alan R

Publication Date

2026-09-03

Peer reviewed|Thesis/dissertation

Sparse Primary Sampling: A Novel Scatter Estimation and Correction Strategy for Cone-beam Computed Tomography

by

Alan R Li

DISSERTATION

Submitted in partial satisfaction of the requirements for degree of

DOCTOR OF PHILOSOPHY

in

Bioengineering

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO
AND
UNIVERSITY OF CALIFORNIA, BERKELEY

Approved:

Signed by:
Ke Sheng
73766CDB522A473...

Ke Sheng

Chair

DocuSigned by:
Youngho Seo
Signed by:
Steven Conolly
DocuSigned by:
Galateia Kazakia
Signed by:
Qihui Lyu
8999354EC122454...

Youngho Seo

Steven Conolly

Galateia Kazakia

Qihui Lyu

Committee Members

Acknowledgments

I am deeply grateful to my advisor, Ke Sheng, for his guidance, mentorship, and support throughout my work and doctorate.

I am grateful to my collaborators Dan Ruan, Ahmad Sakaamini, and Qihui Lyu, for their cooperation and assistance throughout my Ph.D. journey.

I am grateful to the members of my thesis committee, Youngho Seo, Steven Conolly, and Galateia Kazakia, for their feedback and guidance.

I also thank my fellow lab members for their camaraderie throughout my time in the lab. Finally, I am profoundly thankful to my parents and friends, whose encouragement and support made this work possible.

Sparse Primary Sampling: A Novel Scatter Estimation and Correction Strategy for Cone-beam Computed Tomography

Alan Li

ABSTRACT

Cone-beam computed tomography (CBCT) has become integral to image-guided interventions and radiotherapy owing to its compact geometry and capacity for in-room volumetric imaging. Its broad-beam acquisition, however, generates substantial scattered radiation that violates the line-integral assumption underlying reconstruction, biasing attenuation estimates and degrading Hounsfield-unit (HU) accuracy, contrast-to-noise ratio (CNR), and spatial resolution. Existing scatter-correction strategies each impose characteristic penalties: hardware rejection methods such as anti-scatter grids attenuate primary fluence and raise dose; beam-stop arrays sacrifice primary signal and often require dual scans; and software approaches based on analytical kernels, Monte Carlo simulation, or machine learning trade computational cost, prior-knowledge requirements, or robustness for accuracy. This thesis introduces sparse primary sampling (SPS), a hybrid scatter estimation and correction strategy designed to recover scatter while preserving the primary signal in a single, dose-neutral acquisition. The SPS grid comprises a radiolucent holder populated with small focused tungsten collimators that expose a sparse lattice of "smart pixels" to near-primary-only signal, while the surrounding detector records the full primary-plus-scatter signal. Because Compton scatter in the diagnostic energy range is spatially smooth and low-frequency dominant, the scatter sampled at these points can be interpolated across the detector and refined through an alternating-minimization reconstruction that jointly estimates the image and the scatter distribution under edge-preserving and smoothness regularization. The method was validated computationally using GPU-accelerated Monte Carlo simulations of a cylindrical CT phantom and patient-derived head-and-neck and pelvic

phantoms spanning low and high scatter-to-primary regimes, and benchmarked against a three-dimensional Richardson–Lucy denoising approach. A sampling density of approximately 0.08% of detector pixels was identified as optimal, balancing scatter sampling against interpolation error and yielding up to a 94.5% reduction in HU root-mean-square error for the CT phantom, alongside substantial cupping-artifact and CNR recovery, without dose overhead or prior anatomical knowledge. A prototype SPS grid was further evaluated experimentally on an in-house benchtop CBCT system using a Catphan phantom, comparing filtered backprojection and total-variation-regularized iterative reconstruction with and without SPS correction. Despite geometric misalignment and prototype limitations that constrained absolute performance, SPS correction reduced cupping non-uniformity and improved HU linearity and CNR, with iterative reconstruction showing the strongest recovery, demonstrating proof of concept under experimental conditions. Collectively, these results aim to establish SPS as a dose-efficient, single-scan scatter-correction technique that requires no patient-specific priors and integrates measurement, estimation, and reconstruction within a unified workflow. The thesis concludes by identifying the engineering refinements and expanded phantom studies needed to close the gap between the demonstrated experimental performance and the technique's computationally established potential.

Contents

List of Figures

List of Tables

List of Formulas

List of Abbreviations

ASG	Anti-scatter grid
BSA	Beam stop array
CBCT	Cone-beam computed tomography
CNR	Contrast-to-noise ratio
CT	Computed tomography
DQE	Detective quantum efficiency
ESF	Edge spread function
FBCT	Fan-beam computed tomography
FBP	Filtered back-projection
FDK	Feldkamp–Davis–Kress (reconstruction algorithm)
FFT	Fast Fourier transform
FISTA	Fast iterative shrinkage-thresholding algorithm
FOV	Field of view
HN	Head and neck
HU	Hounsfield unit
IGRT	Image-guided radiotherapy
IR	Iterative reconstruction
LSF	Line spread function
MC	Monte Carlo
ML	Machine learning

MTF	Modulation transfer function
NPS	Noise power spectrum
PCB	Printed circuit board
PSF	Point spread function
RL	Richardson–Lucy
RMSE	Root-mean-square error
ROI	Region of interest
SD	Sampling density
SDD	Source-to-detector distance
SID	Source-to-isocenter distance
SPR	Scatter-to-primary ratio
SPS	Sparse primary sampling
SR-POLY	Polyenergetic statistical reconstruction
TV	Total variation
UI	Uniformity index

List of Symbols

Roman letters

D	Time-averaged blank-field measurement
E_0	Incident photon energy
E'	Scattered photon energy
f	Spatial frequency
$f(x, y)$	Object attenuation function
g	Total detected signal (primary + scatter)
h	Planck's constant
$h'(x, y)$	Scatter convolution kernel
H_μ	Huber penalty function
HU_i	Hounsfield value of the i-th pixel in an ROI
HU_μ	Mean Hounsfield value within an ROI
I	X-ray intensity measured per detector pixel
I_0	Unattenuated (incident) X-ray intensity
I_b	Blank-field detector scan
l	Photon path length
N	Number of photons
p	Primary signal
$p(t, \theta)$	Projection (Radon transform of the object)
$p_0(x, y)$	Primary photon signal

r_e	Classical electron radius
s	Scatter signal
s_i	Scaling factor
x, y	Cartesian coordinates within a projection
y_c	Corrected line integral
Z	Atomic number

Greek letters

α	Reduced photon energy, $h\nu/m_0c^2$
Δx	Pixel sampling interval
θ	Scattering angle
λ	Virtual scatter distribution (3-D RL fitting)
μ	Linear attenuation coefficient
ν	Photon frequency
σ	Standard deviation within an ROI

Chapter 1

Introduction

Since its discovery in the late nineteenth century, radiography has undergone substantial methodological development, culminating in the introduction of computed tomography (CT) in the latter half of the twentieth century.

CT reconstructs three-dimensional volumes from two-dimensional projection data, eliminating object superposition and improving volumetric spatial resolution relative to projection radiography. Tomographic slices can be reconstructed following principles derived from the Radon transform, which requires a source trajectory satisfying Tuy's data-sufficiency condition[1] for geometric completeness, and sufficiently angularly dense sampling that satisfies the Nyquist criterion.

Many CT acquisition geometries have been developed to satisfy this sampling requirement, the most prominent being fan-beam CT (FBCT) and cone-beam CT (CBCT). FBCT samples a narrow, collimated fan of radiation across one or a few detector rows, building up a volume slice by slice, while CBCT spreads a divergent beam across a two-dimensional detector to capture an entire volume in a single rotation. FBCT remains the clinical standard for diagnostic imaging owing to its acquisition speed, spatial resolution, and cost-effectiveness. CBCT, by virtue of its open geometry and volumetric acquisition, has been adopted across a wide range of settings, including dental and head-and-neck imaging, fixed and mobile C-arm systems for interventional radiology and surgery,

preclinical small-animal scanners, and on-board imaging systems integrated into radiation therapy linacs.

Of these applications, it is the role of CBCT in image-guided radiation therapy (IGRT) that motivates the present work. Kilovoltage and megavoltage CBCT systems mounted on C-arm linear accelerators provide in-room, volumetric assessment of patient anatomy immediately prior to or during treatment, supplementing conventional two-dimensional portal imaging and enabling near-real-time setup verification and adaptive guidance. The compact footprint and open geometry of CBCT are well suited to this environment, where the closed-ring gantry and confined bore required by FBCT for rapid helical source–detector trajectories are impractical.

Despite these advantages, CBCT image quality remains markedly inferior to that of FBCT, and the cone-beam geometry that enables volumetric acquisition is also its principal liability. Exposing a large two-dimensional detector to a broad beam generates substantial scattered radiation, whereas FBCT intrinsically rejects scatter by collimating to a narrow fan. Compounded by lower detective quantum efficiency (DQE), narrower dynamic range, and motion artifacts arising from slow gantry rotation, scatter is the dominant source of degradation. The detected scatter signal in cone-beam projections grows with subject size and frequently exceeds the primary signal, biasing the measured attenuation coefficients, introducing nonlinearities, and degrading Hounsfield-unit accuracy, image contrast, and spatial resolution.

Scatter correction has therefore been central to CBCT's clinical maturation, with successive refinements driving broader clinical adoption. Numerous CBCT-specific correction techniques have been reported, but few have been incorporated into commercial linacs; those omitted typically increase patient dose, prolong acquisition, or impose system-specific complexity that limits portability. An optimal correction technique should minimize dose and acquisition time, admit a simple yet robust implementation, and effectively suppress scatter-induced artifacts.

This thesis proposes a novel scatter-correction technique that preserves the primary signal while concurrently estimating scatter from sparse measurements, enabling accurate projection-domain

correction. The method employs a sparse primary sampling (SPS) grid that is robust, adaptable, and capable of simultaneously sampling primary and scattered signals.

The objectives of this thesis are to:

- Characterize scatter radiation in diagnostic x-ray imaging.
- Review the development of CBCT and the role of scatter in shaping it.
- Establish a baseline from historical scatter-correction methods.
- Survey current state-of-the-art correction techniques.
- Present the SPS grid technique and quantify the achievable image-quality improvements.

1.1 Thesis Overview

Chapter 2 reviews the historical development of CBCT and its initial motivations. The principal hardware components are described alongside reconstruction theory, and a summary of x-ray physics is introduced to provide a foundation for scatter physics and its consequences for CBCT reconstruction.

Chapter 3 surveys CBCT scatter-correction methods in the literature, broadly classified as software-based, hardware-based, or hybrid approaches. State-of-the-art techniques are reviewed and serve as baselines for evaluating the clinical role of the SPS grid. The limitations of prior and existing methods are then assessed with respect to scatter measurement, scatter estimation, imaging dose, acquisition time, complexity, and robustness.

Chapter 4 introduces the SPS method as a response to the limitations identified in Chapter 3 and states the specific aims of the work. The SPS grid concept and prototype design are presented, followed by the computational validation framework — the Monte Carlo simulation setup, digital phantoms, smart-pixel modeling, the alternating-minimization reconstruction algorithm, and the comparative 3-D Richardson–Lucy method. The computational study is then presented and analyzed

across the CT, head-and-neck, and pelvic phantoms, evaluating the theoretical performance and feasibility of the SPS grid across the relevant scatter-to-primary-ratio regimes.

Chapter 5 presents and analyzes the experimental hardware study using the prototype SPS grid. The benchtop CBCT system, prototype grid, beam stop array, and Catphan phantom are described, the hardware validation methodology is detailed, and the reconstructed image quality is assessed and benchmarked against reference reconstructions.

Chapter 6 examines the sources of error and limitations of the experimental study, organized across the benchtop system and its geometric alignment, the prototype SPS grid, the projection preprocessing, the reconstruction framework, and the scope of the phantom validation. It distinguishes the non-idealities that suppress the measured performance below the method's potential from the use of a grid-free correction target that flatters it, and consolidates the corrective actions into a path forward for the next iteration.

Chapter 2

Background

Diagnostic radiography is subject to a variety of image artifacts, including scatter radiation, beam hardening, and hardware inhomogeneities. To provide a theoretical foundation for understanding image artifacts and scatter artifact correction techniques, this chapter will first provide an overview of relevant X-ray mechanics and interactions.

- X-ray interactions with matter will be discussed to describe how an image signal is detected and how scattered radiation is generated.
- How X-ray photons are generated is described with a summary of the clinical implementation of an X-ray source.

The next section introduces the technology of Cone beam Computed Tomography to provide context and significance for scatter radiation.

- The CBCT hardware components will be described to build upon the prior theoretical foundation and provide the practical implementation of how medical systems generate and detect X-rays.
- CBCT reconstruction will be described to show how the measured X-ray projection data is converted into cross-sectional images, building from conventional CT reconstruction theory toward the cone-beam geometry.

- The artifacts that degrade CBCT image quality — beam hardening, streak and partial-volume effects, and truncation — will be reviewed, with particular focus on scatter as the dominant artifact source motivating this work.

2.1 Physics of Diagnostic X-ray Imaging

Diagnostic X-ray imaging uses photons with energies typically ranging from 10 keV to 100 keV. X-rays in this energy range are typically generated by accelerating electrons between two electrodes and bombarding the anode in a vacuum-sealed x-ray tube. The electrons are generated via thermionic emission, whereby heating certain materials to high temperatures provides sufficient kinetic energy to release electrons. The photons produced by the electron bombardment follow a broad energy spectrum, as shown by **Figure 2.1**. This characteristic X-ray curve is due to the contributions of two radiative processes, Bremsstrahlung and characteristic radiation.

2.1.1 Bremsstrahlung

Bremsstrahlung, German for "braking radiation", is electromagnetic radiation produced from a charged particle changing velocity within an external Coulomb field, with the energy loss conserved as a radiated photon. In the context of X-ray imaging, the electrons impacting the anode, also known as the target, will produce Bremsstrahlung in a continuous spectrum of energy due to the ensemble statistical cross section with target's atomic nuclei. An electron grazing the atomical orbitals may exhibit a slight decrease in energy while an electron directly hitting the target material nucleus will lose all its kinetic energy (**Figure 2.2**). In classical mechanics, the radiated photon will have the energy E = hv, where h is Planck's constant, v is the frequency, and E is the energy equivalent to the loss of kinetic energy. The energy spectrum exhibits predictable behavior, such as an energy peak equal to the electrons' peak energy, and can be modeled via Kramer's Law. With some approximation, the photon production is dependent on the photon energy produced in a linearly

decreasing relationship, although due to the internal target filtering and external X-ray window filter, most of the lower energy photons are absorbed.

2.1.2 Characteristic Radiation

Characteristic radiation is the discrete electromagnetic radiation emitted when electrons bound within atomic orbitals transition between energy levels. The difference in the binding energies of electron shells determines the frequency/energy of the emitted photon, resulting in multiple discrete lines in the EM spectrum of an X-ray generator corresponding to different electron-shell transitions. The innermost shell, K-shell, has the highest binding energy, with L, M, and N shells exhibiting lower binding energy, respectively, as the distance from the nucleus increases.

In the context of X-ray generators, accelerated electrons interact with atomic orbitals and eject bound electrons, resulting in a vacancy within an orbital. Depending on which energy transition, characteristic X-rays are denoted by the resulting shell of the transitioning electron and a subscript of α, β, depending on whether an adjacent shell or a nonadjacent shell donated the electron, respectively (**Figure 2.3**).

2.1.3 Interactions with Matter

Although not as interacting as charged particles such as the previous electron, X-ray photons interact with matter via various processes (**Figure 2.4**). X-ray photons exceed the energy threshold required to liberate electrons from atomic orbitals and are termed ionizing radiation. Similar to the charged-particle case, X-ray photons undergoing interactions will undergo energy and/or momentum changes, which can be accompanied by energy deposition, the generation of secondary particles, or annihilation. The processes are as follows, Rayleigh scattering, Compton scattering, photoelectric absorption, and pair production. Pair production occurs above 1.02 MeV, which is far beyond the diagnostic X-ray energy range and will not be discussed for our purposes. The previous three processes constitute the aggregate attenuation behavior of photons and will be discussed after the description of each individual process.

2.1.3.1 Rayleigh Scattering

Rayleigh scattering involves a photon interacting with the atom and exciting all bound electrons (**Figure 2.5**).

The excitation occurs from the photon's oscillating EM field expending energy to oscillate the atom's electrons in phase, which will then radiate the expended energy back out as another photon of the same energy as the incident radiation. Since the energy is equivalent from before the interaction,

$$E' = E_0 \tag{2.1}$$

Rayleigh scattering is also known as elastic scattering. The electron oscillation is a collective behavior of the atomic electrons with the superposition of individual electron waves producing the emitted scattered radiation, hence Rayleigh scattering, also known as coherent scattering. Rayleigh scattering is dominant within lower energy regimes within the range from <30 keV, as the binding energy of most K-shells has not been reached.

The emitted photon's scattering angle probability can be calculated using the classical Thomson differential cross-section of a free electron with added terms such as an atomic form factor to approximate the effective charge.

$$\frac{d\sigma_0}{d\Omega} = \frac{r_0^2}{2}(1 + \cos^2\theta) \tag{2.2}$$

where r_0^2 is the classical electron radius, θ is the scattering angle with respect to the incident axis of the beam, and $\frac{d\sigma_0}{d\Omega}$ is the differential Thomson interaction cross section describing the fraction of incident energy scattered into a given unit solid angle. Despite the lower incidence at higher energies, with at most 10% of interactions at 30 keV, Rayleigh scattering is strongly biased around an angular deflection of 0 degrees and is a significant contributor to scatter, especially for lower energy X-ray imaging.

2.1.3.2 Compton Scatter

In contrast to Rayleigh scatter, Compton scatter involves photons interacting and exciting individual bound electrons, hence the additional name of incoherent scatter (**Figure 2.6**).

In the context of X-ray imaging, Compton scattering is the dominant form of scatter in the diagnostic energy range of 10 keV to 100 keV and remains dominant at therapeutic energies up to 30 MeV. After 1.022 MeV, pair production occurs, but for the purposes of scatter discussion in diagnostic imaging, the process will be ignored. During Compton scatter, an incoming ionizing photon interacts with a bound electron, which probabilistically is assumed to be an outer shell electron due to the low binding energy, and ejects the electron while also losing energy in the process. From energy and momentum conservation, the incident photon's energy is equal to the summed scattered photon's energy and the ejected electron's kinetic energy, as shown by:

$$E_0 = E' + E_{e^-} \tag{2.3}$$

For the purposes of modeling energy loss and scattering angles, the binding energy of the outer valence electron is ignored as it is negligible. Due to the loss of the photon's energy, Compton scattering is also known as inelastic scattering. The ejected electron does not travel far in the surrounding material and is promptly absorbed following potential interactions with charged particles, such as further excitations, ionizations, or bremsstrahlung. These secondary interactions can generate additional secondary particles, depositing energy in the local area.

The scattering angle of the photon can be derived following momentum and energy conservation, with the resulting formulation shown below

$$E' = \frac{E_0}{1 + \frac{E_0}{511\ keV}(1 - \cos\theta)} \tag{2.4}$$

Where E' is the scattered photon's energy, E_0 is the incident photon energy, and θ is the deflection angle of the scattered photon with respect to the incident trajectory. The 511 keV term is derived

from the electron's rest mass and refers to the maximum energy transferred to the electron by the incident photon.

The probabilistic angular distribution of scattering Compton photons can be described by the Klein-Nishina formula with the differential angular scattering cross section as follows

$$\frac{d\sigma}{d\Omega}(\theta) = r_e^2 \left(\frac{1}{1+\alpha(1-cos\theta)}\right)^2 \left(\frac{1+\cos^2\theta}{2}\right)\left(1+\frac{\alpha^2(1-cos\theta)^2}{(1+\cos^2\theta)[1+\alpha(1-\cos\theta)]}\right) \quad (2.5)$$

Where $\alpha = \frac{h\nu}{m_0c^2}$ and r_e is the classical electron radius. As the above formulation assumes a free electron, bound electron Compton scatter angular profiles can be plotted using the incoherent scattering function as a correction factor. As seen in **Figure 2.7**, higher energy photons show a forward bias in the scattering angle which is relevant to X-ray imaging as more scatter photons will be detected by the detector [2].

Additionally, Compton scattering at diagnostic energies results in less energy loss in the scattered photon, while higher-energy incident photons result in more energy transfer to the scattered electron (more energy loss in the scattered photon). Consequently, Compton-scattered photons are a significant contributor to image quality in diagnostic X-ray imaging, as they retain much of their penetrability and can reach the detector.

As the incident photon energy increases for X-ray imaging, the relative probability of Compton Scatter to Rayleigh, or, to be discussed, Photoelectric absorption increases as the difference between the electron binding energy and photon energy increases. As the photon interacts with individual electrons, Compton scattering is dependent on the electron density of the atom, which varies minimally among biological elements aside from hydrogen. Hydrogen contains no neutrons and has double the electron density of other materials in tissues, which equates to double the interaction cross section.

2.1.3.3 Photoelectric Effect

In the photoelectric effect, a photon is absorbed by a bound electron, and the excitation surpasses the electron's binding energy, resulting in electron emission. All the incident photon's energy (E_o) is absorbed by the electron, and the emitted electron, dubbed a photoelectron, has kinetic energy (E_{pe}) equal to the energy surplus after the binding energy (E_b).

$$E_{pe} = E_o - E_b \tag{2.6}$$

As seen in the equation, the incident photon must have energy larger than the binding energy of the photoelectron. The emitted photoelectron results in an ionized atom, and the vacancy in inner electron orbitals is filled by outer shell electrons that can result in an electron cascade. The difference in binding energy between the orbitals is either released as characteristic X-rays or Auger electrons, which are electrons emitted by absorbing the energy difference.

The precise binding energy differs between materials, but in general, the outer valence shell electrons have the least binding energy, which increases in the inner L and K shells, respectively. Incident photons have higher interaction probabilities the closer their energy is to the various shell binding energies, resulting in characteristic attenuation curves shown in the figure (**Figure 2.8**). The closer the photon energy is to the binding energy of an electron shell, the higher the chance of photoelectric interaction with an electron of the shell. A consequence of this interaction cross-section is sharp discontinuities known as absorption edges that occur when the photon energy is just above the binding energy. Apart from these discontinuities, the interaction cross-section for photoelectric absorption decreases as the photon energy (E_o) increases. The probability of photoelectric interaction per unit mass with the atomic number and energy is approximately proportional as follows:

$$\sigma \propto \frac{Z^3}{E_o^3} \tag{2.7}$$

Where Z is the atomic number and E_o is the energy of the incident photon.

The photoelectric effect is relevant to X-ray imaging, as the absorbed photons do not produce scattered photons as in previous interactions. The presence of the absorption edges leads to material-specific contrast between bone and soft tissue, which is highly relevant for X-ray imaging. The complete absorption of the photons is relevant for the detection of diagnostic X-rays with detector pixels composed of higher Z materials to increase detection efficiency. High Z materials also work well as shielding for general X-ray shielding or blocking scatter via anti-scatter grids as discussed in Chapter 3.

2.1.4 X-ray Attenuation

The various photon interactions compose the total material linear attenuation coefficient μ as follows:

$$\mu = \mu_{Rayleigh} + \mu_{photoelectric} + \mu_{Compton} + \mu_{pair\ production} \tag{2.8}$$

The linear attenuation coefficient describes the fraction of photons removed traveling through a medium with units of photons $(unitless)$ per distance (m^{-1}). The linear attenuation coefficient of a material depends on the density of the material, as the interaction probability scales with the number of atoms per unit distance. When considering X-ray imaging, variations in tissue density directly affect measured grey values in the produced images.

Figure 2.8 shows the interaction contributions of the individual photon interactions for μ of water and lead [3] depending on photon energy. As mentioned previously, the relevant interaction for scatter production is Compton scattering, which is dominant in the diagnostic energy range of 10 keV to 100 keV. Photoelectric contribution is significant at lower energies, and tissues such as bone exhibit higher photoelectric absorption at the electron shell absorption edges.

The total attenuation of a photon beam can be described by the number of interacting photons N_{int} removed from the beam is proportional to the number of incident photons N_o traveling through an infinitesimal material thickness interval dx such that the number of transmitted photons $dN(x)$ is given by:

$$dN(x) = -N_{int}(x) = -\mu(x) * N_o(x)dx \tag{2.9}$$

Integration of 2.8, assuming a homogeneous material medium, results in the Lambert-Beers equation of

$$N = N_o e^{-\mu x} \tag{2.10}$$

For diagnostic X-ray imaging, this formula provides the means of deriving the effective linear attenuation coefficient of an X-ray source passing through an object using a projection image, which measures N. The following form of the Lambert-Beers equation is used for CT reconstruction and describes the image formed at the detector, as shown as follows

$$I = I_o * e^{\int_{-\infty}^{\infty} -\mu(l)\ dl} \tag{2.11}$$

Where I is the X-ray intensity measured by the detector per pixel and I_0 is the unattenuated X-ray intensity from the source. l is the parameter describing the photon path, and is the line integral calculated to obtain the effective total linear attenuation coefficient. The derived total linear attenuation coefficient contains the beam path material composition; by extension, it also contains the image information.

2.2 Diagnostic X-ray Production

The classical method of X-ray production involves accelerating electrons from a cathode towards an anode within an evacuated chamber, also known as the X-ray tube. A brief overview of components in the X-ray tube is given in **Figure 2.9**

The major components are listed as follows: cathode, anode, rotor/stator, glass or metal envelope, tube port, cable sockets, and tube housing, but focus will be placed on the X-ray generation components such as the cathode, anode, and tube housing/envelope.

The cathode is the negative terminal of the X-ray tube, comprising one or two focal-spot filaments and a focusing cup. The filament is heated to produce electrons via thermionic emission, and the number of electrons leaving the filament surface is proportional to the circuit current. The focusing cup is maintained at the same potential difference relative to the anode as the filament, shaping the emitted electrons towards the anode. Under an applied voltage difference, the electrons flow towards and are concentrated at the anode focal spot, which is the geometric source of the X-ray photons.

The anode is the positive terminal of the X-ray tube, where thermal electrons emitted from the cathode strike. The sudden impact from the accelerated electrons with the anode, or target, converts the electron kinetic energy into either heat or radiated particles, including X-ray photons. The target material is a high-Z material, normally tungsten, to maximize bremsstrahlung efficiency and handle heat deposition from the electrons. Additionally, the target/anode is angled with respect to the central axis in the X-ray field geometry. Effectively, the projected focal length by the photons is smaller than the actual focal length produced from the anode. The geometric foreshortening of the focal spot length at the anode and along the central ray is called the line focus principle and allows for increased power delivery with smaller effective focal spot sizes without overheating the anode. A consequence of the line focus principle is the heel effect or reduction in the x-ray beam intensity distribution in the anode direction due to "self-attenuation" of photons in the target.

At the target, X-ray photons are produced from both bremsstrahlung and characteristic photon emission, and the combined interactions produce the broad bremsstrahlung spectrum and sharp characteristic X-ray peaks of a typical X-ray curve as seen in **Figure 2.10** [4]:

The peak tube voltage determines the maximum kinetic energy of the photons, assuming complete kinetic energy conversion of the electron into radiation via bremsstrahlung. For diagnostic imaging energies, the bremsstrahlung conversion efficiency for a tungsten target is approximately 1%, with 99% of the incident electron energy being wasted as heat or low-energy radiation that is filtered. The low energy X-rays are removed by the inherent materials of the X-ray tube and intentional filters, often made with aluminum, are used to further filter and shape the beam.

2.3 Computed Tomography

X-ray projection radiography provided the first glimpses into the human body without an invasive procedure, but the images produced were limited in the information provided. The projection of 3-Dimensional information onto a 2-D film resulted in anatomical features being overlaid and obscured. The development of X-ray Computed Tomography overcame the limitations of 2-D projections and enabled visualization of 3-D information by reconstructing crosssectional images of an object from angular radiographic projections.

The following sections provide an introduction to conventional CT, focusing on core principles that intersect with CBCT, such as image acquisition and reconstruction. Next, the CBCT system will be detailed by highlighting shared characteristics with conventional CT and CBCT-specific features. Finally, a summary of the various CBCT image artifacts and associated sources will be provided to provide context and differentiate the focus of the paper, scattered radiation-induced artifacts.

2.3.1 Conventional CT

The FBCT is the most widely adopted development of multiple decades and generations of CT scanner designs, improving upon image quality, image acquisition speed, and imaging dose efficiency with a multi-slice CT capable of finishing a scan in the order of seconds. A modern CT scanner consists of a gantry containing the rotating source and detector used to capture angular projections, as shown in **Figure 2.11** [5].

For third-generation CT systems, the source and detector are rigidly fixed with respect to each other, and the detectors used are curved scintillating solid-state detectors.

The fan beam geometry of conventional CT refers to the narrow X-ray beam with minimal divergence along the patient length (z-axis) of the CT scanner, as shown in **Figure 2.11**.

In comparison, a broad-beam geometry exhibits significant beam divergence, with a cone angle approaching the fan angle (**Figure 2.12**). In practice, FBCT utilizes a narrow cone beam geometry wider than a true fan beam to expose multiple detector arrays rather than a single row, allowing for

larger volumes to be covered in a single gantry scan. To image the patient anatomy, the CT gantry rotates along the patient's length in a helical path to capture the full volume of interest.

2.3.2 CT Reconstruction

The CT scanning geometry in **Figure 2.11** shows the fan beam, which consists of rays representing the beam path taken by the X-ray photons from the source to the detector. Each ray is attenuated according to the Lambert-Beer's law with an exponential decay term composed of the attenuation coefficients as seen in Equation 2.10. It is assumed that the object of interest is within the imaging boundaries determined by the detector and beam angles to extend the integral to infinity. By measuring the incident intensity I_0 beforehand, the line integral p, given by equation 2.11, along the X-ray path l can be isolated from the Lambert-Beer's equation and provide the linear summation of attenuation coefficients which can be extracted via reconstruction algorithms.

$$p = -\ln \frac{I}{I_0} = \int_{-\infty}^{+\infty} \mu(l) dl \tag{2.12}$$

Where I is the transmitted intensity through an object. In practice, the line integral measured by the projections is inaccurate or distorted, yet the derived algorithms assume the equation is ideal. These errors lead to incorrect data and appear as artifacts in the reconstructed CT images, such as scatter, beam hardening, partial volume, or truncation artifacts.

2.3.3 Image Reconstruction

To provide an intuitive understanding of CBCT reconstruction, the basic principles of fan-beam reconstruction will be introduced. For simplicity, assume the CT projections are obtained using parallel-beam geometry, where attenuation paths are parallel to each other. We assume ideal line integrals defined by a projection at an angle θ and distance from the origin t, $p(t, \theta)$, of the following form

$$p(t,\theta) = -\ln\frac{I(t,\theta)}{I_0} = \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(x,y)\,\delta(t - x\cos\theta - y\sin\theta)\,dx\,dy \tag{2.13}$$

Where $f(x,y)$ is the mathematical function of the object of interest. The set of line integrals formed by the projection, $p(t,\theta)$, is also known as the Radon transform of the function $f(x,y)$. Using the Radon transform, the object information $f(x,y)$ can be obtained following the Fourier Slice Theorem which states that the Fourier transform of a function's projection is equivalent to a radial line at constant θ in the 2-D Fourier Transform of the function

$$F_{1D}(k,\ \theta) = F_{2D}(k\cos\theta, k\sin\theta) \tag{2.14}$$

In practice, the set of angular projections and line integrals obtained from CT scans is finite and as such the resulting object $f(x,y)$ is discrete in the reconstruction corresponding to the spatial resolution of the image. The Radon transform approximation obtained from discrete projections data is sufficient for diagnostic medical imaging and enables modern CT imaging. As such, there are numerous approaches to the CT reconstruction problem with the simplest method performing a filtered backprojection (FBP):

$$f(x,y) = \int_0^{\pi} F^{-1}(F_{1D}\ (p(t,\ \theta)) * |\omega|\ )d\theta \tag{2.15}$$

The $|\omega|$ density compensation filter is often known as the ramp filter for its characteristic ramp shape in the Fourier domain. Using FBP, the reconstructed image does not exhibit the blurring caused by simple backprojection smearing and provides a tomographic slice of the object.

The values of the reconstructed object $f(x,y)$ correspond to the attenuation coefficients of the materials in the object. To scale the attenuation coefficients in a visible range, the attenuation coefficients are converted to Hounsfield Units (HU) with the following equation

$$f(x,y)_{HU} = \frac{f(x,y)_{\mu} - \mu_{water}}{\mu_{water}} * 1000 \tag{2.16}$$

The resulting Hounsfield scale is centered at 0 HU, which corresponds to water. Air (or vacuum, representing minimal attenuation) is assigned -1000 HU, while bone typically ranges from about 600 to 1000 HU. Larger HU values correspond to metallic materials such as those found in metal implants and scale until 2000 HU. Consequently, the full 4096 HU range can be stored in 12 bits and visualized with 12-bit color depth.

2.4 Cone-beam Computed Tomography

Instead of utilizing narrow fan beam geometry, CBCT utilizes broad cone beams to capture projection images of an object with beam cone angles nearly equivalent to the fan angles (**Figure 2.13**).

CBCT projections are captured using true cone beam geometry with 2-D flat panel detectors rather than narrow cone beam geometry with curved detector arrays found in MDCTs. Both CBCT and FBCT use a third-generation rotate-rotate source and detector configuration, but CBCT does not require the helical scan trajectory used in FBCT for volumetric imaging. For CBCT reconstruction, a circular trajectory is sufficient for collecting projection data. As a result, CBCT supports more flexible scanning geometries, and its compact form factor enables integration with a wide range of existing flat panel systems, including C-arms and linac-mounted platforms. Integrating CBCT into existing systems enables online or onboard volumetric imaging, which is essential during radiological interventions or radiation therapy for on-demand guidance and anatomical verification. Cone-beam image reconstruction shares underlying principles with conventional FBCT FBP but differs in implementation due to beam divergence and flat-panel detector geometry. To account for these differences, the basic cone-beam reconstruction technique, the Feldkamp-Davis-Kress (FDK) algorithm, was developed to reconstruct the full volume scanned, which will be described next.

2.4.1 Cone-beam Reconstruction

The FDK method is a derived FBP technique that accounts for increased cone angle in projections and reconstructs the total 3D volume instead of sequential individual slices as in FBCT FBP. The

total density at a point within the cone beam is taken to be the sum of the contributions from all projections that intersect the point. Like FBP, FDK relies on a ramp filter and additionally uses a pre-weighting factor, dependent on the fan and cone angles, to weigh the contribution at each point. Unlike FBP, the FDK method is approximate, with the result always deviating from the ground truth.

An important consequence of note is that from the circular trajectory, according to the Tuy sufficiency condition [1], CBCT reconstruction is unable to provide data for exact reconstruction of an object outside of the central axial plane that the source trajectory resides in, as a complete intersection from all projections is impossible. As a result, cone beam artifacts can arise in reconstructed non-central axial slices due to data insufficiency, and this will be discussed in the following section.

2.4.2 Cone-beam CT Artifacts

Cone-beam images suffer from the same artifacts that conventional FBCT exhibits in addition to cone-beam artifacts caused by the previously mentioned data insufficiency. In general, image artifacts are inaccurate HU values caused by deviations from the assumptions of image reconstruction, such as nonlinearities caused by scatter contribution or beam hardening effects. A common consequence of CBCT reconstruction errors is the cupping artifact, where normally homogenous regions have a darkening HU gradient towards the center of an object. A non-exhaustive list of cone-beam artifacts is described in the following subsections.

2.4.2.1 Beam Hardening

As shown in **Figure 2.10**, the x-ray source produces a polyenergetic photon spectrum with a broad, uneven distribution across energy bins. The linear attenuation coefficient μ varies with photon energy, leading to lower-energy photons being preferentially absorbed relative to higher-energy photons, causing the mean photon energy of the beam to increase with depth and resulting in beam hardening. As higher-energy photons have lower μ projections, they underestimate the central regions of an object with longer path lengths, while shorter-path-length regions, such as the object periphery,

exhibit less beam hardening. Denser structures, such as bones, also exhibit preferential absorption of lower-energy photons and contribute to beam hardening, with extremely dense structures causing streak artifacts, as discussed next. Beam hardening artifacts are mitigated by filtering the outgoing X-ray beam with metallic filters, such as aluminum, to remove most low-energy photons.

2.4.2.2 Streak Artifacts

When excessive photon attenuation occurs in a region near a high-density object, such as a metal implant, the signal received by the detector is insufficient for reconstruction. The resulting HU values of beam paths behind the dense object are error-prone due to the lack of photon statistics in combination with beam hardening. Techniques developed to mitigate the streak artifacts are known as metal artifact reduction techniques, but are not the focus of this discussion. The metal implants also induce increased scatter in the object, resulting in additional scatter artifacts that follow a streak pattern.

2.4.2.3 Partial Volume

When the resolution of the reconstruction is limited, e.g., from choosing a large focal spot size resulting in a larger point spread function and reduced ability to discriminate between specific materials, partial volume artifacts occur. The assumption that each voxel is a homogeneous material breaks down, resulting in a voxel with mixed materials and an average attenuation coefficient. Partial volume artifacts typically occur when thick slices are reconstructed along the patient's length and can be overcome using higher resolution images or correcting with a thinner slice reconstruction.

2.4.2.4 CBCT Truncation

Truncation artifacts occur when the beam's field of view does not fully encompass the object, leading to incomplete projection data and inconsistencies during reconstruction. Although truncated anatomy still contributes to attenuation along the beam path, it is not captured in all projections, leading to missing information and artifacts such as shading, cupping, or bright bands near the edges

of the field of view. While truncation artifacts are not unique to cone-beam CT, they are often more pronounced due to the large cone angle and limited detector size. In circular CBCT acquisition, reconstruction is most accurate near the central plane, and artifacts generally increase in severity for slices farther from this plane.

2.4.2.5 Scatter

Scatter is a major contributor to HU inaccuracies in CBCT, leading to artifacts such as cupping and overall image degradation. These artifacts arise from violations of the assumption that only primary photons are detected, which underlies the line-integral model used to estimate attenuation coefficients. In reality, scattered photons generated from interactions within the object also reach the detector, leading to an overestimation of detected intensity and an apparent reduction in attenuation.

This results in shading and blurring in the reconstructed image, with underestimation of HU values, reduced soft-tissue contrast, and compromised quantitative accuracy. The problem is exacerbated in cone-beam geometries, where the large radiation field increases both scatter production and detection. Path integrals with high scatter-to-primary ratio detector measurements, particularly those passing through highly attenuating materials such as bone or metal, can produce pronounced streaking artifacts in the reconstruction, similar in appearance to beam hardening effects.

The focus of this thesis is the development of a novel technique to mitigate scatter artifacts in cone-beam CT and correct the resulting image degradation while preserving signal information. The next chapter provides context and an overview of the current landscape of scatter correction techniques in CBCT, laying a foundation for evaluating the proposed method.

2.5 Figures

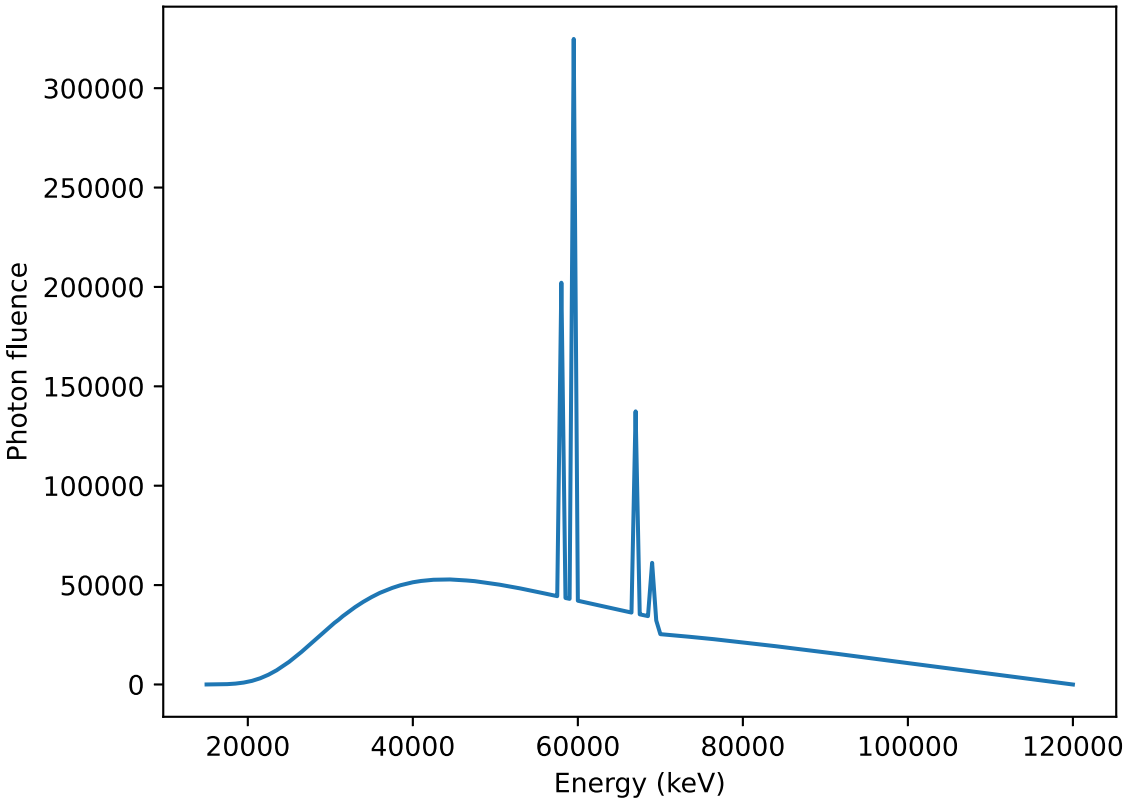

Figure 2.1: Filtered polyenergetic X-ray spectrum, showing the continuous bremsstrahlung distribution with the superimposed characteristic radiation peaks.

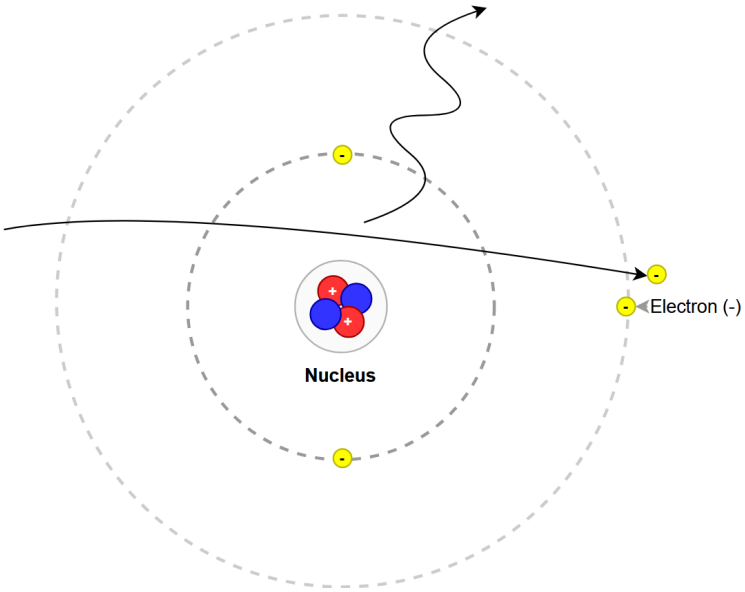

Figure 2.2: Production of bremsstrahlung radiation: an incident electron is deflected and decelerated by the Coulomb field of the nucleus, emitting a photon equal to the kinetic energy lost.

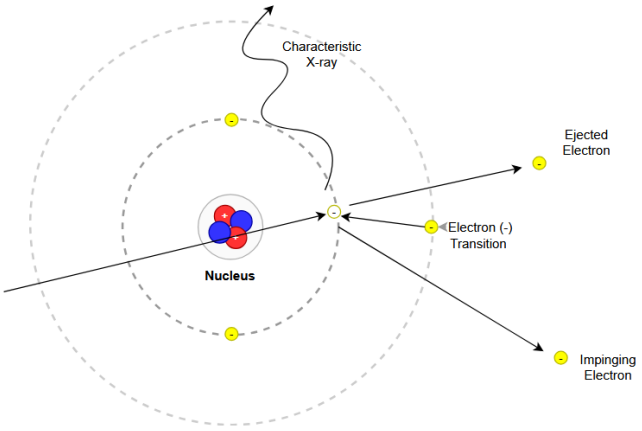

Figure 2.3: Production of characteristic radiation: an impinging electron ejects an inner-shell electron, and an outer-shell electron transitions to fill the vacancy, emitting a characteristic X-ray photon equal to the difference in the shell binding energies.

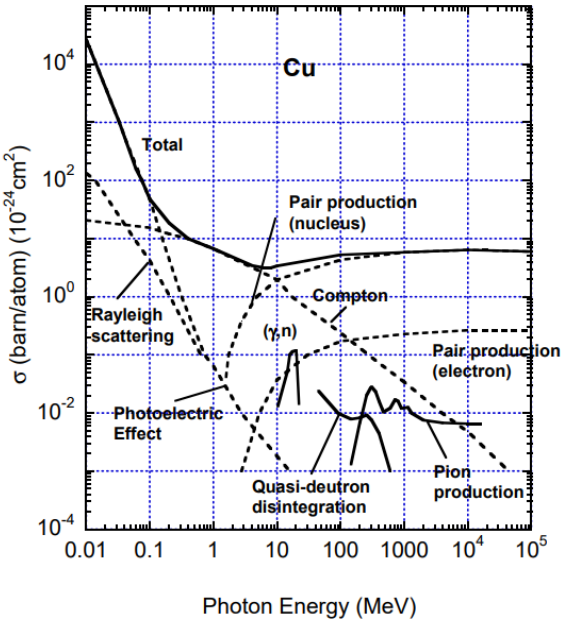

Figure 2.4: Photon interaction cross sections for copper as a function of photon energy, showing the relative contributions of Rayleigh scattering, the photoelectric effect, Compton scattering, and pair production to the total cross section. Adapted from KEK (rcwww.kek.jp/research/shield/photon_r.pdf).

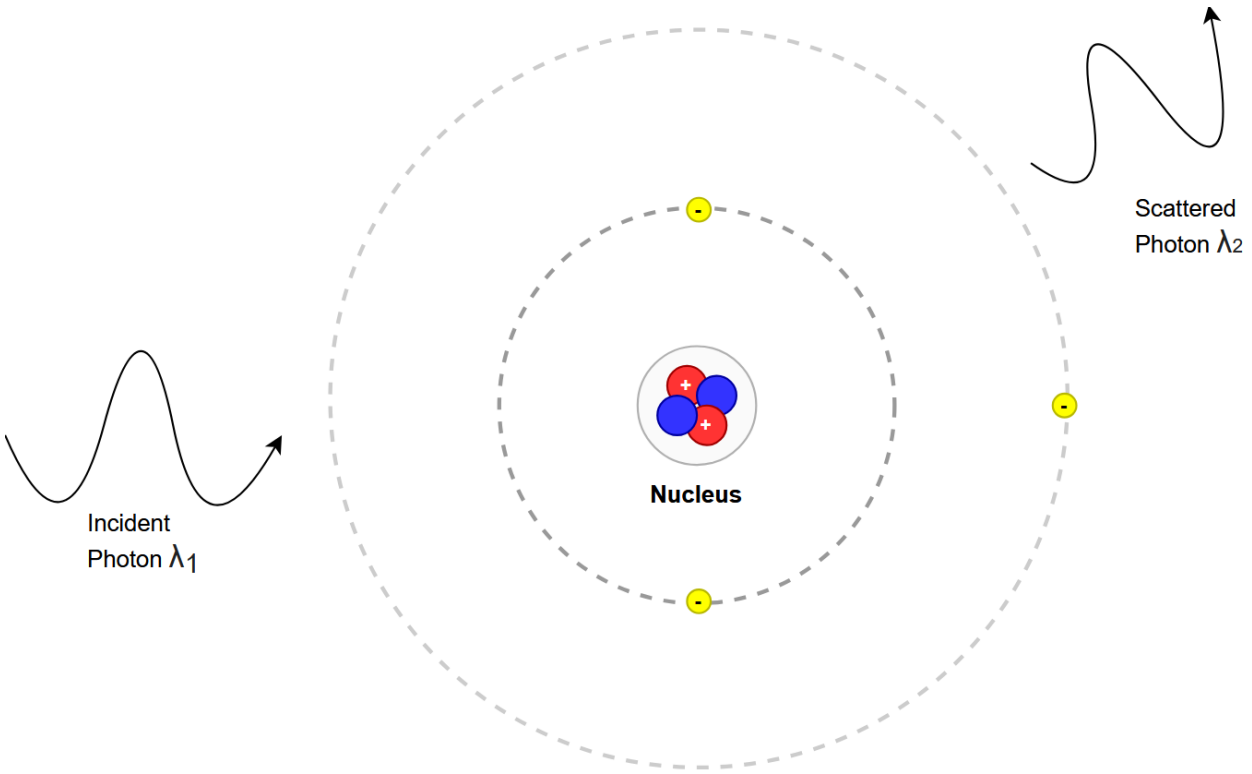

Figure 2.5: Rayleigh (elastic) scattering: an incident photon of wavelength λ_1 interacts with a bound atomic electron and is redirected as a scattered photon of essentially equal wavelength λ_2, changing direction without losing energy.

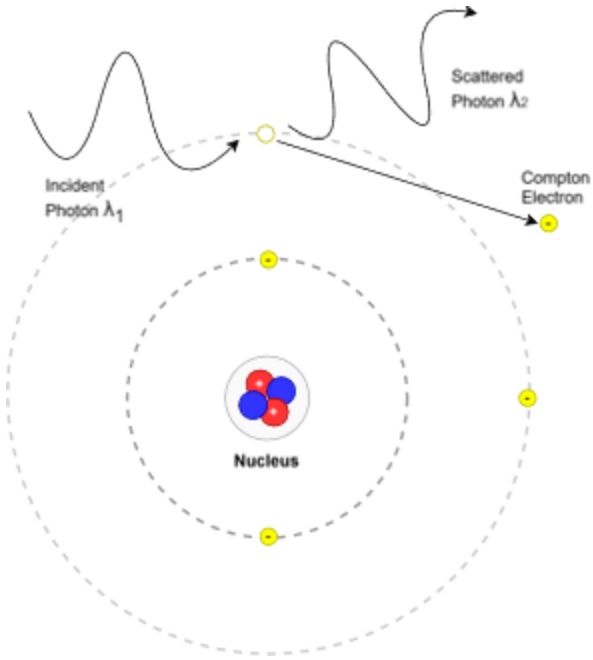

Figure 2.6: Compton (incoherent) scattering: an incident photon of wavelength λ_1 transfers part of its energy to a loosely bound electron, ejecting it as a Compton electron and producing a scattered photon of longer wavelength λ_2.

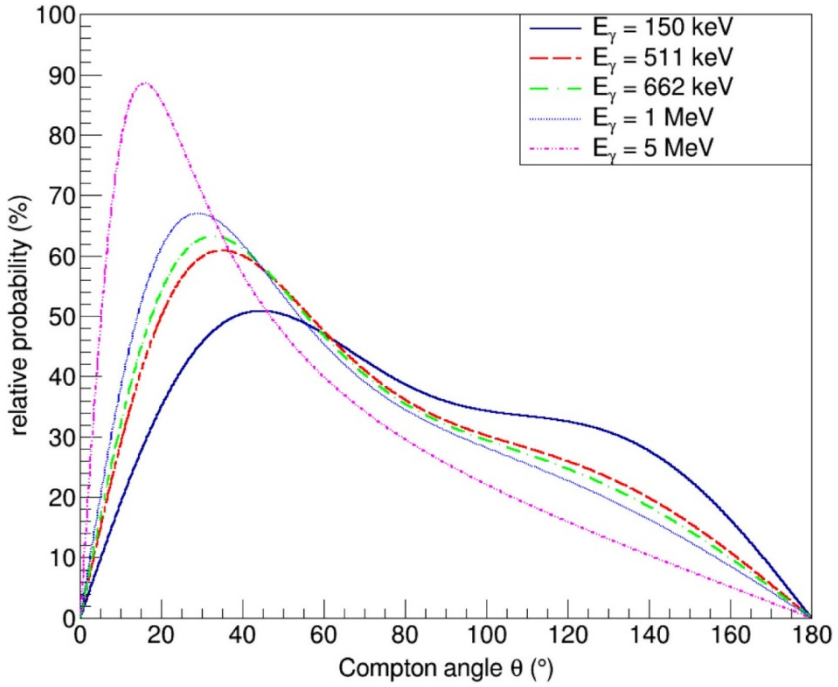

Figure 2.7: Relative probability of Compton scattering as a function of scattering angle θ for several incident photon energies, showing the increasing forward bias of scattered photons at higher energies.

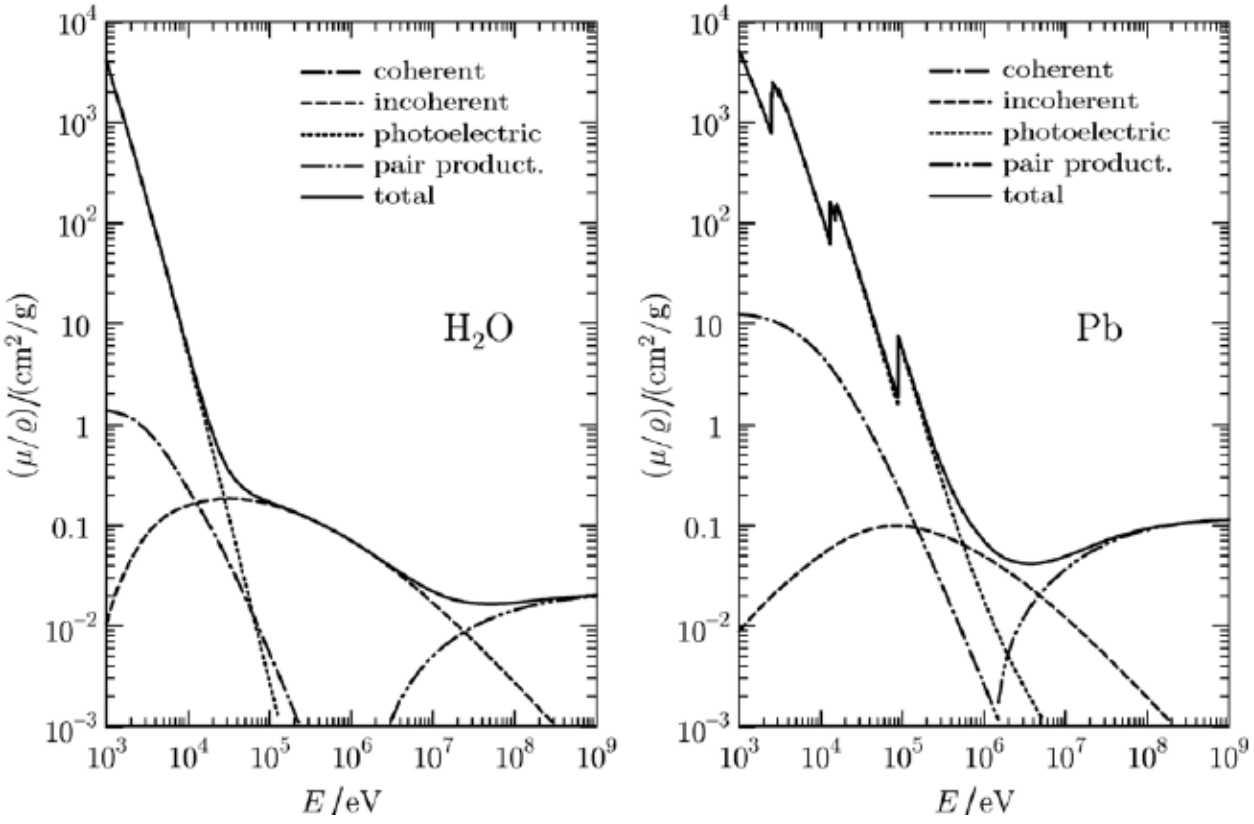

Figure 2.8: Mass attenuation coefficient as a function of photon energy for water (left) and lead (right), separated into the coherent (Rayleigh), incoherent (Compton), photoelectric, and pair-production contributions to the total.

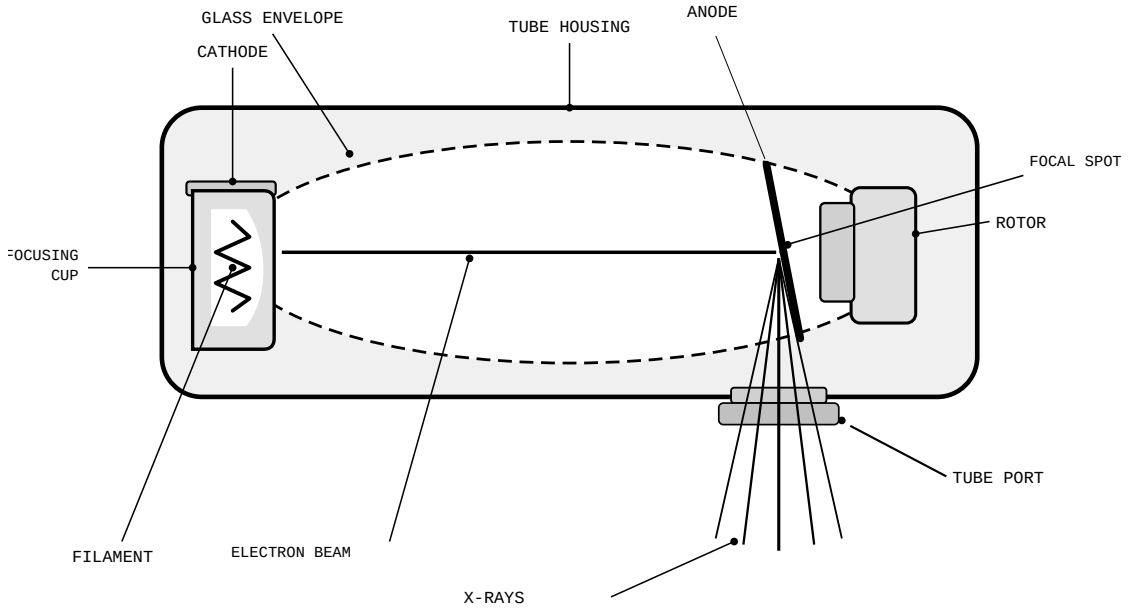

Figure 2.9: Schematic of a rotating-anode X-ray tube and its principal components, including the cathode, filament, focusing cup, glass envelope, anode, focal spot, rotor, and tube port.

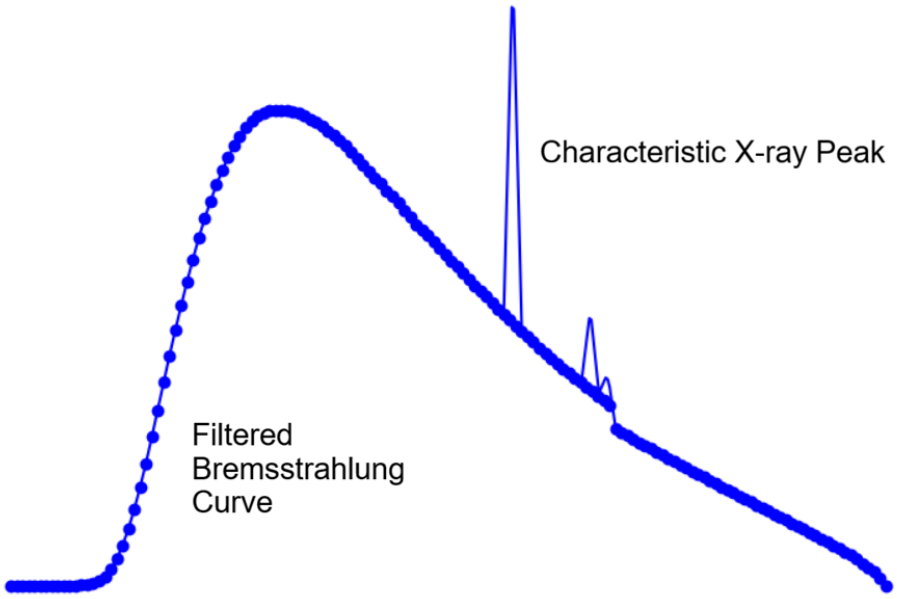

Figure 2.10: Representative X-ray output spectrum, showing the filtered bremsstrahlung continuum with the superimposed characteristic X-ray peaks of the anode material.

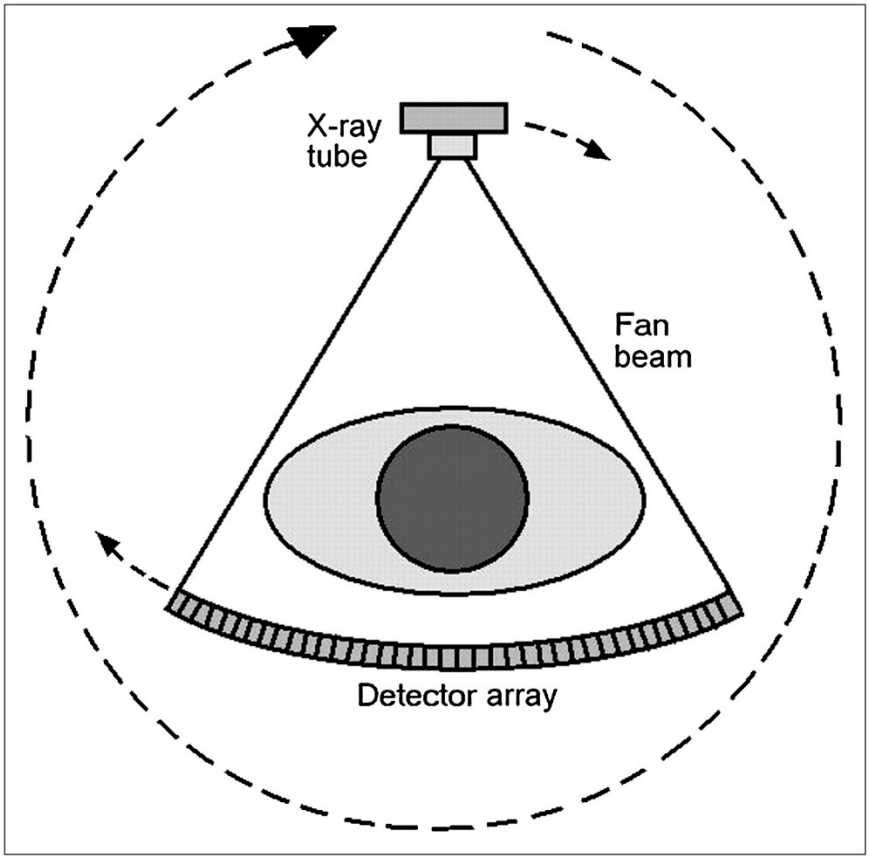

Figure 2.11: Third-generation fan-beam CT geometry, in which the X-ray tube and curved detector array rotate together about the patient.

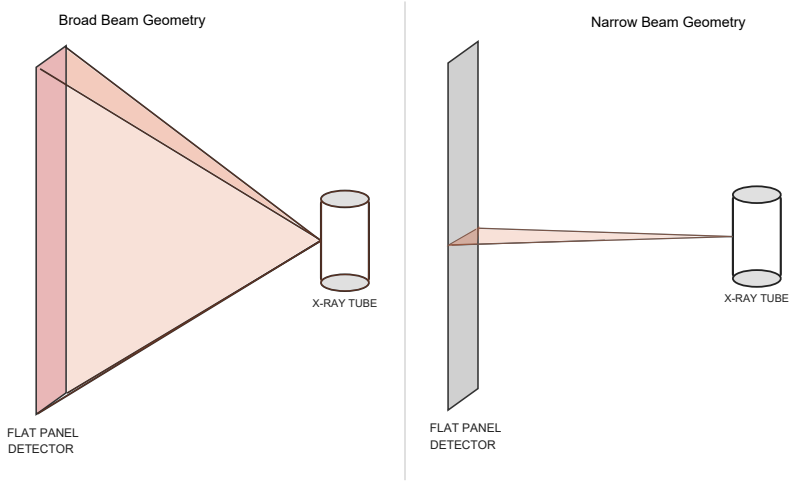

Figure 2.12: Comparison of broad-beam geometry (left), which exposes the full two-dimensional flat-panel detector, and narrow-beam geometry (right), which collimates the beam to a thin profile to limit scatter at the detector.

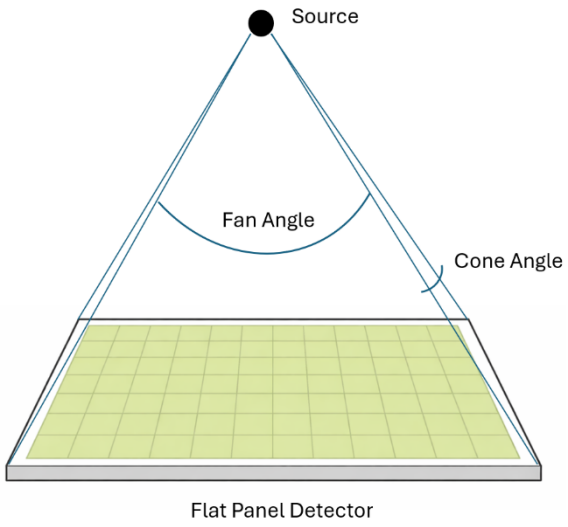

Figure 2.13: Cone-beam CT geometry, in which a point source illuminates a two-dimensional flat-panel detector, defined by the in-plane fan angle and the out-of-plane cone angle.

Chapter 3

Scatter Correction Techniques for Cone-beam CT Projections

The previous chapter presented the various image artifacts that hamper image quality in CBCT reconstruction, including reductions in HU accuracy due to inhomogeneities, shading artifacts, and blurring. Scatter was identified as a major contributor to reduced image quality in CBCT, whereas it is less of an issue in conventional CT, which is more widely utilized.

As mentioned previously, the focus of this chapter will be to provide an overview of scatter correction techniques and build a foundation for evaluating the novel proposed scatter correction technique of this thesis. A nonexhaustive overview of relevant scatter techniques will be categorized and broken down, with both the benefits and limitations described for each. Although not a strict organization, scatter correction techniques can be considered as hardware-based, software-based, or a hybrid of the two. For hardware-based methods, techniques can be further broken down into pre- and post-patient methods.

For note during the evaluation of CBCT scatter techniques, it is known that scatter-to-primary ratios vary between anatomical regions, with larger anatomies such as the torso and pelvis having more difficult regions to correct for, as the SPR is much larger. Each technique was evaluated using

different phantoms and anatomies, and direct comparison is not the focus of this chapter, but rather the analysis of the pros and cons of each technique.

The flow of the chapter is as follows:

- Software-based methods such as analytical scatter kernels or statistical estimation via Monte-Carlo simulations or Machine learning.
- Hardware-based methods are categorized according to pre-patient devices, such as scatter measurement devices like BSAs, or post-patient devices that block scatter, such as Antiscatter Grids.

3.1 Software-based correction methods

Software-based scatter correction methods aim to estimate or model the scatter distribution to correct measured projections used for reconstruction. These techniques generally involve solving a constrained optimization problem that iteratively estimates the scatter components and the reconstructed images. Other software-based approaches include Monte Carlo methods to statically approximate the scatter distribution and machine-learning methods that either estimate scatter or correct for it during or after image reconstruction. The approaches differ in their methods for estimating the scatter component and how prior information is leveraged and computation resources are optimized.

3.1.1 Kernel-based Scatter Correction

Scatter correction via analytical kernels assumes that the scatter distribution is a blurred version of the primary photon distribution, related by a point-spread function (PSF) or scatter kernel. By performing a deconvolution with the scatter kernel, the primary photon distribution can be extracted from the measured photon signal [6]. The measured photon distribution at the detector can be modeled as:

$$p(x, y) = p_0(x, y) * h'(x, y) \tag{3.1}$$

where $p(x, y)$ is the measured pixel value, $p_0(x, y)$ is the primary photon signal, and $h'(x, y)$ is the convolution kernel which can be expressed as:

$$h'(x, y) = h'_p(x, y) + h'_s(x, y) \tag{3.2}$$

where $h'_p(x, y)$ is the primary distribution component, normally a $\delta(x, y)$ function at one detector pixel, while $h'_s(x, y)$ is the scatter kernel from a single photon, assumed to be spatially invariant and circularly symmetrical. The X-ray beam can be conceptualized as a combination of pencil beams, and the resulting scatter at a pixel is the superposition of multiple convolved primary pencil beams. The scatter can then be removed from measured projections through iterative deconvolution to obtain corrected projections for reconstruction [6], [7], [8].

However, analytical kernel-based methods, although computationally efficient with no dose overhead, suffer from limited scatter correction due to their lack of anatomical specificity, which relies on assumptions such as spatially invariant scatter kernels and simplified scatter attenuation [7]. The scatter kernel accuracy can be improved for a specific system and geometry by including parameters such as object size, airgap, exposure parameters, and anti-scatter hardware, but it remains limited because patient parameters cannot be feasibly determined experimentally. Ultimately, the simplicity and speed of kernel-based methods must be traded off for increased accuracy and improved scatter correction effectiveness.

3.1.2 Monte Carlo Simulation Methods

Assuming the scatter distribution from a projection is properly modeled, the measurements can be corrected by simply subtracting the scatter contribution. Monte Carlo (MC) methods can be considered the 'gold standard' for scatter modeling, as they directly simulate the physical transport of photons through matter, generating both primary and scattered signals [9]. However,

naive MC simulations are computationally expensive, a significant limitation that is further exacerbated in iterative reconstruction frameworks that require repeated forward MC projections [10]. Computational limitations of Monte Carlo simulations include the need for long particle histories to generate a low-noise estimate of the scatter distribution and memory constraints that depend on the simulated geometry fidelity. Methods exist to accelerate Monte Carlo-based approaches, such as variance reduction techniques or GPU acceleration [11], [12] but careful assumptions are required to simplify and accelerate MC simulations. Modifying statistical interactions can bias results and yield nonphysical scatter estimates. One such method by Zbijewski [13] uses low-count MC simulations to obtain a noisy estimate of the scatter, which is then denoised using 3-D Richardson-Lucy (3-D RL) deconvolution.

Although Monte Carlo simulations provide robust, fundamentally consistent estimates of scatter distributions, the computational cost required for accuracy makes clinical application prohibitive. The number of photons scales linearly with computation time, but noise decreases at a rate of $\frac{1}{\sqrt{N}}$ where N is the number of particles. When simulating a full set of CBCT projections, the computation time is compounded by multiple orders of magnitude, scaling from hours to days. A caveat of noisy Monte Carlo data is that the simulated scatter-estimation noise is combined with the measured data noise, and this combination may require postprocessing to mitigate errors [14]. For high-fidelity scatter modeling, the MC simulation requires a priori information from a previous scan to produce accurate patient geometry. As patient scans are not ground-truth measurements, the scanned data must be processed to remove existing scatter contamination, registration errors, and FOV truncation, unless a typically infeasible total-body scan is utilized [15].

An alternative software-based scatter correction technique, similar to Monte Carlo-based methods for modeling physical interactions, is Acuros CTS, which solves the Boltzmann diffusion equation with accuracy approaching MC methods while being computationally less expensive [16], [17]. However, the technique depends on a high-fidelity planning CT to accurately model patient geometry and consequently shares the same practical difficulties as Monte Carlo-based approaches. As a computational solver, numerical approximations and discretization are used to make the problem

computationally tractable, with memory usage and time efficient, but these compromises come at the cost of modeling accuracy. Parameter tuning via Pareto optimization is vital for achieving good performance and a balance between accuracy and runtime.

3.1.3 Machine Learning

Recently, machine-learning (ML) methods have been adopted to approach the scatter migration problem [18], [19]. The application of machine learning to CBCT scatter correction involves training neural networks to recognize the complex, non-linear relationships between scatter-corrupted measurements and the underlying primary signal. With a fully trained model, ML-based scatter correction shows excellent results, with speeds far exceeding those of existing software-based methods [20].

However, preparing and training an ML model is a significant bottleneck for ML methods, as carefully curated and properly labeled data for a given system is required for good scatter correction performance [21]. One approach utilizes synthetic CTs generated with Monte Carlo simulations as part of the deep learning workflow, but this method requires considerations as stated in the previous section[22]. The biggest drawback is the lack of robustness in ML scatter correction with hallucinations, both false positives and false negatives, impacting the confidence and explainability of medical images with no direct connection to ground truth measurements.

In brief, software methods have achieved varying levels of success. However, without physically detecting and differentiating the scatter components, the robustness of software-based methods can be challenged, such as during motion-induced artifacts.

3.2 Hardware-based correction methods

Scatter compensation techniques utilizing hardware can be categorized into pre- and post-patient devices depending on their placement in the source trajectory relative to the patient and detector. Hardware methods aim to correct for scatter either through scatter rejection or scatter measurement

followed by subtraction in the projection. To keep things focused, hardware-based techniques with relevance to CBCT scatter correction will be mainly discussed as techniques such as a bowtie filter or airgap are widely utilized and complimentary to novel developments.

3.2.1 Pre-patient Hardware-based Scatter Correction Techniques

Pre-patient hardware strategies are designed to modify or attenuate the X-ray beam before it enters the subject. These methods focus on reducing the source of scatter by limiting the irradiated volume or by encoding information into the primary beam to facilitate downstream computational separation of the scatter component [23], [24], [25].

One class of techniques involves scatter modeling and estimation using beam blockers and beam stop arrays (BSAs) [26], [27], [28], [29], [30], [31], [32]. These methods aim to sample the scatter directly by blocking primary transmission and isolating the scatter signal in the attenuated primary's shadows, thereby interpolating the full scatter distribution. An alternative to direct scatter measurement via BSAs is primary fluence modulation [25]. By utilizing semitransparent blockers, the high-frequency components of the primary distribution can be modulated without perturbing the low-frequency dominant scatter distribution. The low-frequency primary information is then extracted via filtering and demodulation to estimate the scatter signal.

Despite their utility, scatter-blocking-based methods have significant limitations. BSAs inherently cause primary signal loss in blocked regions, necessitating dedicated recovery strategies. The conventional two-scan mode addresses this by employing a secondary open-field acquisition [29], though at the cost of doubled scan time, increased patient dose, and greater susceptibility to motion artifacts. Single-scan alternatives have been explored to mitigate these penalties, including moving blocking arrays that shift with each angular projection and interpolating data from adjacent projections [31], [32], as well as simpler stationary configurations such as the optimized or half-beam blocker [26], [30]. However, moving arrays introduce mechanical complexity and require precise synchronization among components, thereby increasing susceptibility to detector lag. Stationary designs, while simpler, necessitate full scans and rely on data redundancy assumptions that break

down in the presence of moving anatomies. Primary fluence modulation, though it avoids outright signal blockage, introduces its own technical demands: accurate spatial encoding is critical for signal recovery, and the technique therefore requires precise mechanical alignment and carefully tuned demodulation parameters to reconstruct reliable image data. An intrinsic difficulty with pre-patient devices is susceptibility to gantry sag, where the C-arm gantry that aligns the source and detector deviates from alignment due to weight. For pre-patient devices, misalignments in positioning relative to the source are magnified at the detector, leading to inaccuracies in scatter correction. Additional concerns with pre-patient devices include scatter from imperfect blocking, reduced detector FOV, and, when the treatment source serves as the signal source, incompatibility with real-time treatment guidance.

3.2.2 Post-patient

Post-patient hardware is located between the subject and the detector. Their role is to provide a physical filter that discriminates between ballistic primary photons and off-axis scattered photons by absorbing scattered photons.

Anti-scatter grids (ASGs) are the most clinically prevalent implementation of post-patient hardware, and operate by rejecting scattered radiation through aligned septa and spacers while permitting the passage of primary photons [33]. However, 1-D ASGs are intrinsically limited to rejecting CBCT scatter due to the 1-D septa being unable to perfectly align with the broad angular spread of scatter from a 3-D irradiated volume. Two-dimensional ASGs, arranged in a cellular pattern, overcome this limitation by rejecting scatter from more potential trajectories in the object geometry, offering superior performance over conventional one-dimensional linear grids in CBCT [34].

Despite their widespread adoption, ASGs carry significant limitations. The most consequential is the dose penalty imposed by the Bucky factor, the ratio of radiation rejected, because grid septa and spacers inevitably attenuate a portion of the primary beam alongside scattered photons; the exposure must be increased to maintain equivalent primary fluence at the detector and preserve

image quality. Although 2-D ASGs are a significant improvement in imaging dose utilization, there is still attenuation up to 7% of primary fluence [34]. In general, ASGs demand precise mechanical alignment during deployment. Misalignment, whether off-center, off-focus, or off-level, produces grid cutoff artifacts characterized by underexposure at image edges and structured intensity non-uniformities, which must be calibrated and corrected for. However, the practicality of 2-D ASGs is constrained by their design and the high-precision tungsten fabrication they require, driving up manufacturing costs, increasing weight for potential gantry sag, and compounding the challenge of widespread clinical implementation[35]. Additionally, each 2-D ASG is tailored to a specific scanning protocol and lacks adaptability, further constrained by high fabrication costs, substantial weight, and fragile septa.

3.3 Hybrid correction methods

Hybrid methods that combine hardware and software strategies to address the scatter problem have also been explored. However, core problems, such as reducing primary transmission, remain [36], [37], [38], [39]. Scatter correction hardware, such as beam blockers, is used to sample scatter in shadow regions, and the measurements are used in an iterative scheme to reconstruct the full scatter. Limitations to these methods include narrowing the detector field of view (FOV) via beam blocker edge-based scatter sampling [37], multiple scan sequences for strip-based lead collimators [38], or the previously mentioned difficulties with moving beam blockers [39].

3.4 Conclusions

Across hardware-based and software-based approaches, a central tension exists between scatter rejection and scatter measurement as competing correction philosophies. Hardware methods that physically reject scatter, such as ASGs, avoid the need for explicit scatter modeling but impose inherent primary signal attenuation, requiring increased exposure to compensate, which is a dose penalty that scales with the number of projections required for a technique. Measurement-based

approaches, whether via beam stop arrays, primary modulation, or software estimation, attempt to recover scatter without outright rejection but introduce their own burdens: blocked regions must be restored, modulated signals must be demodulated, and Monte Carlo or iterative methods demand substantial computational resources, rendering real-time clinical application impractical. Complexity, whether mechanical, computational, or procedural, emerges as a persistent barrier across all categories. Pre-patient devices require precise source alignment and are susceptible to gantry sag; post-patient grids demand high-precision fabrication and careful calibration; and software methods entail significant training overhead or simulation costs. Hybrid approaches, while promising in principle, largely inherit the limitations of their constituent components rather than resolving them. Collectively, these trade-offs highlight the absence of a unified solution that simultaneously achieves effective scatter correction, dose efficiency, and practical clinical deployability. The following section proposes a novel scatter correction method that seeks to overcome the limitations of previous techniques.

Chapter 4

The Sparse Primary Sampling Grid: Concept and Computational Validation

This chapter presents our novel scatter correction technique designed to address the practical limitations of current hardware- and software-based methods identified in Chapter 3, and establishes its theoretical performance through a computational validation study. The technique uses a novel sparse primary sampling device to sample the primary signal sparsely while minimizing signal loss. Unlike previous hardware scatter estimation techniques, the primary signal is mostly retained at the sampling points and can be recovered in a single scan without significant mechanical complexity. The new method, dubbed SPS scatter correction, benefits from robust measurement data for estimation while requiring no dose overhead.

The chapter is organized as follows:

- The theoretical background for the SPS grid method and the initial prototype SPS grid design are introduced.
- The computational validation framework is described, including the Monte Carlo simulation setup, digital phantoms, smart pixel modeling, the SPS alternating-minimization reconstruction algorithm, and the comparative 3-D Richardson-Lucy (RL) fitting method.

- Quantitative image-quality metrics of CNR, MTF, HU RMSE, and uniformity index (UI) are defined with reference to the phantom ROIs.
- Reconstruction results for the CT, head-and-neck, and pelvic phantoms are presented, benchmarking the SPS grid against primary-only, uncorrected total-signal, and 3-D RL fitting reconstructions.
- SPS technique performance is discussed across sampling densities and compared against existing hardware- and software-based scatter correction methods, and the assumptions and limitations of the computational study are examined, with implications for the experimental hardware validation presented in Chapter 5.

4.1 Sparse Primary Sampling Grid Theory

Scatter, as discussed in Chapter 2, is the result of various charged-particle and photon interactions that scatter primary photons and produce secondary photons in the imaged geometry. The scattered photons contribute to the pixel values measured by the detector but do not reflect the true radiologic path of the primary photons, leading to inaccuracies and deviations in the reconstruction. For the CBCT energy range of 10 keV to 120 keV, the dominant interaction contributing to measured scatter is Compton scatter [40], as referenced from **Figure 2.4**, which can be modeled by the Klein-Nishina formula to predict the angular distribution for a given energy range.

As seen in **Figure 4.1**, Compton scatter at the energy ranges associated with CBCT is broadly distributed in scatter angles. As a consequence, the Compton scatter distribution can be assumed locally smooth [41], which translates to a low-frequency-dominant signal in the Fourier domain. **Figure 4.2** confirms these findings using pelvis data obtained from a Monte Carlo simulation, where most of the information is concentrated in the lower frequencies and along the Cartesian directions.

As most of the scatter information is encoded as low-frequency signals, a sparse sampling approach can be utilized to reconstruct and estimate the scatter distribution for scatter correction. By sparsely sampling the scatter in K-space using a design such as a Cartesian grid, most of the signal

in **Figure 4.2** can be captured by the detector, and the full scatter distribution can be estimated by interpolating to fill in signal voids. However, isolating the scatter from the detector will result in the loss of primary photons required for image reconstruction. By utilizing the linear relationship of the detector signal, the scatter can be obtained by knowing both the primary signal and total signal, as seen below:

$$s = g - p \tag{4.1}$$

where p is the primary signal, g is the total signal, and s is the scatter signal. The proposed SPS grid enables measurement of the primary signal while maintaining the total signal, and is described in the following section.

4.2 Sparse Primary Sampling Grid Design

The proposed SPS grid is designed to sparsely sample the primary signal while minimizing signal loss. To accomplish this, the SPS grid uses hollow tungsten inserts machined to only a few pixels in diameter, with a thickness sufficient for scatter rejection, as shown in **Figure 4.3**. To enable flexible sampling of the primary signal at varying sampling densities and locations, a modular design was conceived, composed of a holder plate and individual cylindrical tungsten collimators. Holes are machined in the holder plate and aligned to focus on the source. The holder plate is radiolucent and minimally attenuates incoming X-rays. The cylindrical tungsten collimators are inexpensively sourced by cutting hollow tungsten tubes into multiple segments. By inserting the tungsten tube segments into holes oriented toward the source, pixels covered by the tubes are exposed only to primary photons, with minimal scatter contamination. These pixels are termed smart pixels. The total signal from the smart pixels covered by the tubes is recovered using nearest-neighbor interpolation.

The optimal smart pixel dimensions were determined through heuristic decision using NIST attenuation values [42] and results from the computational study introduced in the following sections.

Considerations for tube dimensions include the number of pixels blocked and scatter contamination on the central pixel. Increasing scatter contamination will limit the improvement in image quality from the SPS scatter correction method, while blocking more pixels results in larger interpolation errors that lead to downstream deviations in reconstruction accuracy and potential artifacts.

Figure 4.4 shows preliminary investigations of smart-pixel designs conducted via Monte Carlo simulations in TOPAS. Tubes were modeled as tungsten hollow cylinders pointed straight at the source. Anatomical scatter was simulated with a 25.6×25.6×25.6 cm^3 cube of water in the beam path. Finite source size and penumbral effects were simulated through a translational shift in the x-y plane, while tube misalignment was simulated through an angular tilt of the tube.

Benefits of the modular holder design include rapid prototyping, flexible manufacturing within design tolerances, light weight, and economy. Current SPS grid prototypes are 3-D printed and simple to mount to the detector, as discussed further in Chapter 5. Regarding the design's feasibility, aligning the individual smart pixel tubes to the source requires high precision during holder fabrication. For a smart pixel tube with an inner diameter of 1 mm and a detector-to-source distance of 150 cm, the angular accuracy of the holes in the plate must be within 0.04° of the ideal angles towards the source. This machining tolerance is achievable with existing industrial standards [43].

4.3 Computational Validation Methodology

To determine the feasibility of the SPS scatter correction technique, a computational study was conducted using Monte Carlo simulations to generate projections and model phantoms. The study establishes the theoretical upper bound on SPS correction performance before any hardware factors are introduced.

The workflow for CBCT scatter correction is pictured in **Figure 4.5** and involves obtaining raw projections with an SPS grid to sample the isolated primary signal at smart pixels while maintaining signal elsewhere. The smart pixel scatter signal is recovered using nearest-neighbor interpolation, and an initial scatter distribution is estimated via 2-D cubic spline interpolation. A dedicated

reconstruction algorithm that uses an iterative alternating minimization loop simultaneously improves the scatter model and corrects the image for scatter artifacts.

For the computational study workflow, the raw projections are simulated via Monte Carlo simulations with a modeled SPS grid, as described in greater detail in the next section.

4.3.1 CBCT Monte Carlo Simulation Setup

The MC-GPU v1.3 simulation code [12] is used to calculate the CBCT projections at sufficient histories within a reasonable computation time via GPU acceleration. The X-ray transport interactions and voxel materials are adapted from the MC particle simulation toolkit PENELOPE 2006. MC-GPU allows for the analysis of photon counts at the detector by separating the scatter and primary signal. Scatter photons are further categorized into Compton, Rayleigh, and multiple-scatter photons, each tracked separately at the detector. A 120 kVp X-ray source from a tungsten target with a 4.3 mm aluminum filter is utilized for the simulations. The CT scan protocols place the source at 150 cm from the detector with varying isocenter distances based on the center of the object geometry, as shown in **Figure 4.6**. Both CT and HN phantom simulations utilized the same particle statistics of 5×10^{11} histories per projection, while the pelvic phantom utilized 1×10^{12} histories per projection to increase particle statistics while remaining computationally feasible.

Each scan consists of 360 projections with 1° angular separation. For smaller phantom simulations, the simulated detector is a 400×400 flat panel with 1 mm² detector elements for a total size of 40×40 cm² and a total detector element count of 160,000. For larger phantom simulations, the detector simulated is a 600×600 flat panel with 1 mm² detector elements for a total size of 60×60 cm² and a total detector element count of 360,000. The detector approximates a clinical FPD within MC-GPU memory and speed constraints. The simulated detector measures the raw photon count and associated photon energy at each detector pixel.

4.3.2 Digital Phantoms

For the computational study, three phantoms were simulated following PENELOPE 2006 material definitions to be compatible with MC-GPU. For quantitative image quality evaluation, a custom cylindrical CT phantom (**Figure 4.7**) was simulated to benchmark image improvement in spatial resolution, contrast, and scatter artifact recovery via cupping artifact reduction.

The CT phantom measures 25.6 cm in diameter and 25.6 cm thick with 1 mm³ voxels, and is placed at 100 cm source-to-isocenter distance (SID). Within the CT phantom there are two halves, a line pair slice and an embedded contrast rod slice, which are shown in the digital phantom. The phantom material is water, and the line-pair test material is bone. The contrast rods consist of lung, cartilage, bone, glandular, brain, and adipose inserts. The images were reconstructed at 1×1×1 mm³ voxel sizes to match the digital phantom dimensions and provide a ground truth comparison with the synthetic matrix defining the CT phantom.

For clinical anatomy, a head and neck (HN) patient phantom and pelvic patient phantom were simulated to represent clinical anatomy in low and high scatter-to-primary ratio (SPR) regimes, respectively. Both phantoms are depicted in **Figure 4.8**. For the head and pelvis phantom simulations, a patient DICOM was converted to PENELOPE-compatible material definitions using a continuous tissue model [11]. The head phantom is approximately 25×19×25 cm³ with 1×1×2.5 mm³ voxels. The pelvis phantom was converted following the same protocol and is approximately 34×39×39 cm³ with 1×1×3 mm³ voxels. The four materials simulated were adipose, air, soft tissue, and bone of varying densities based on the converted patient DICOM HU values.

The HN phantom operates in a low SPR regime due to its relatively compact anatomy and the predominance of bony structures and air cavities, which limit the volume of scattering material. Scatter-induced CNR degradation is moderate, and the primary-only reconstruction provides HU values close to the nominal tissue definitions, making it a reliable benchmark. The pelvic phantom operates in a high SPR regime owing to its larger anatomical volume and the substantial soft tissue mass surrounding the central bony structures of the pelvis, which generates significantly more Compton scatter than the HN anatomy. This elevated scatter produces more severe HU depression

and stronger cupping gradients, compressing tissue contrast more aggressively and making scatter artifact correction more demanding. The high SPR also means that residual scatter contamination from imperfect correction has a larger absolute impact on image quality and secondary scatter artifacts. The pelvic phantom therefore motivates the dedicated artifact analysis presented later in this chapter.

4.3.3 Smart Pixel Modeling

For the simulation study, the SPS grid was modeled with the attenuating effects of the SPS holder plate omitted and with the holder plate itself modeled as a perfect radiolucent material. To simplify the simulation geometry, the smart pixels are modeled based on the detector output, assumed to measure the primary signal with tunable contamination from simulated scatter, while the non-smart pixels measure the total signal. For example, with 1.5 mm-thick tungsten collimators, smart-pixel scatter contamination is modeled by adding 4% of the scatter fluence to the simulated primary signal. Values for scatter contamination were obtained with NIST reference data [42] and verified using Monte Carlo simulation data as seen in **Figure 4.4**. This simplification allows for arbitrary grid schemes and comparisons using the same simulation particle statistics.

For the MC-GPU simulations, smart pixels were placed uniformly in a grid pattern, as shown in **Figure 4.9**.

From preliminary simulation comparisons of various tube geometries, a tube with an inner diameter of 1 mm, an outer diameter of 4 mm, and a length of 2 cm was selected, corresponding to a 5×5-pixel dead zone surrounding a single smart pixel for a detector with 1×1 mm² pixels. Tube dimensions were determined by matching a detector pixel's dimensions and achieving sufficient attenuation while minimizing detector coverage by the tube shadows. Scatter attenuation in the smart pixels ranged from 96–98%, with values obtained by comparing with an identical simulation omitting the tubes. Potential penumbral effects from the tube geometry are ignored in the simulation; concerns regarding the tube penumbra are discussed in the limitations section of this chapter. The

percentage of smart pixels relative to total detector pixels gives rise to a sampling density metric where

$$SD = \frac{\#\ of\ Smart\ Pixels}{\#\ of\ Detector\ Pixels} \tag{4.2}$$

Various sampling densities (SD) of the total detector pixels were tested, with 1.05%, 0.275%, 0.0756%, 0.0506%, 0.0156%, and 0.0056% SD selected for determining the efficacy of the SPS grid and scatter reconstruction algorithm in removing scatter artifacts in the reconstructed image without affecting the total photon fluence.

4.3.4 SPS Reconstruction Algorithm

An alternating minimization reconstruction is utilized to achieve an optimal solution for both image reconstruction and scatter estimation. The CBCT sparse primary sampling reconstruction problem is formulated as follows:

$$\mathcal{L}_{tot}(x, s; g) = \mathcal{L}_{fid}(Ax, g - s) + \lambda_1 \mathcal{L}_{sm,x}(Dx) + \lambda_2 \mathcal{L}_{sm,s}(Ds). \tag{4.3}$$

where $\mathcal{L}_{tot}$ is the total cost function, $\mathcal{L}_{fid}$ is the measurement fidelity term and is a function of the forward projected image, Ax, and the corrected projection, $g - s$, with g as the total signal measurement and s as the estimated scatter measurement. $\mathcal{L}_{sm,x}$ and $\mathcal{L}_{sm,s}$ are regularization terms that assume smooth behavior for the image, x, and scatter, s, with $\lambda_1,\ \lambda_2$ as the respective hyperparameter weightings. λ_1 and λ_2 were selected using root mean squared error (RMSE) comparisons, described later, of the CT phantom reconstruction to score performance. The values selected for $\lambda_1,\ \lambda_2$ are insensitive to changes, with minimal image performance differences between patient anatomy and scanning geometry. For physical intuition, λ_1 models assumptions in the image domain, such as edge preservation, while λ_2 models assumptions in the scatter domain, such as being low-frequency dominant.

The estimates for the image and scatter, x and s, are related to the measurement g as follows:

$$p + s + n = g, p = I_b e^{-Ax} \tag{4.4}$$

where p is the primary signal, g is the total signal measurement, s is Poisson-distributed measurement noise, and I_b is the blank field scan of the detector. For convenience, an intermediate line integral variable y is introduced:

$$y = \log I_b - \log(g - s) \tag{4.5}$$

As y contains noise from the total signal measurement g, the fidelity term L_{fid} reflects the relationship $y \cong Ax$. Edge-preserving piecewise smoothness in the image domain is encoded with the Huber function [44], $L_{sm}, x(\cdot) = H_\mu$, and smoothness of the scatter is regularized with a Tikhonov term, $L_{sm}, s(\cdot) = ||\cdot|\,|_2^2$. Both Huber and Tikhonov regularization are continuous and differentiable, allowing efficient solving through FISTA. The Huber function is defined as follows:

$$H_\mu(t) = \begin{cases} \frac{1}{2\mu}t^2\,, & |t| \le \mu \\ |t| - \frac{\mu}{2}\,, & |t| > \mu \end{cases} \tag{4.6}$$

The primary image and the scatter projection signal are solved alternately based on the best estimate of the other optimization variable with the following two objective functions:

$$\min_x \frac{1}{2}\|Ax - y\|_2^2 + \lambda_1 H_\mu(Dx), \quad \text{s.t. } x \ge 0 \tag{4.7}$$

$$\min_s \frac{1}{2}\left\|I_b e^{-Ax} + s - g\right\|_2^2 + \lambda_2 \|Ds\|_2^2, \quad \text{s.t. } 0 \le s \le g \tag{4.8}$$

The schema in **Figure 4.10** summarizes the reconstruction pipeline minimizing L_{tot}. The reconstruction pipeline begins with a blank field scan followed by a phantom scan. The sparsely distributed smart pixels collect only the primary signals, which are subtracted from the nearest neighbor interpolation of the total signal of surrounding pixels to obtain the scatter signal at the

smart pixel. Another interpolation is performed to estimate the full detector-space scatter s_0 and form the intermediate line integral. The line integral produces an initial scatter-corrected image x_k , which is used as a static variable to estimate a better scatter estimate s_{k+1} at all non-smart pixels. The loop continues until a set number of iterations. Forward and back projection operators were provided by the Michigan Image Reconstruction Toolbox [45]. Each submodule was solved utilizing the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [46] for its speed and familiarity in the implementation. A previous study that implemented FISTA served as the reference [47].

To compare with the SPS reconstruction results, the image was reconstructed using the primary-only signal and the uncorrected total signal, representing best- and worst-case reconstruction results, respectively. The primary-only image and the total signal image were generated using the following cost functions and line integrals $y_p,\ y_t$:

$$\begin{aligned} &\min_x \frac{1}{2} \left\| Ax - y_p \right\|_2^2 + \lambda_1 H_\mu(Dx), \quad \text{s.t. } x \geq 0,\ y_p = \log\left(I_b\right) - \log(p) \\ &\min_x \frac{1}{2} \left\| Ax - y_t \right\|_2^2 + \lambda_1 H_\mu(Dx), \quad \text{s.t. } x \geq 0,\ y_t = \log\left(I_b\right) - \log(g) \end{aligned} \tag{4.9}$$

4.3.5 Comparative Method: 3-D Richardson-Lucy Fitting

For the evaluation of the computational SPS simulation study, the comparative 3-D Richardson-Lucy fitting method was chosen based on the work of Zbijewski and Beekman [13]. Zbijewski proposed an iterative MC-based scatter artifact reduction scheme for cone-beam micro-CT that forms the basis of the 3-D RL fitting approach adopted here. Their method begins by reconstructing an initial image from scatter-contaminated projections, then uses MC simulation on that reconstruction to obtain a first-pass scatter estimation. The scatter estimation can be used to correct initial projections and produce images that can be used either to improve the scatter estimation by repeating the simulation correction step or to provide sufficiently corrected images for visualization. The paper proposes a two-stage acceleration strategy: a rapid, low-count MC simulation is first executed, and the inherently noisy result is then de-noised by fitting 3-D Gaussian basis functions via the RL algorithm. Extending the RL fit into the neighboring angular projections exploits the known smoothness of

scatter distributions as a function of projection angle, yielding an additional order-of-magnitude speedup over the earlier 2-D projection-plane-only formulation. Compared to plain MC simulation, the combined approach reduces required computation time by three to four orders of magnitude. The corrected scatter estimate is fed back as a background term in a poly-energetic statistical reconstruction algorithm (SR-POLY), which simultaneously corrects for beam hardening. Validation on simulated rat-abdomen phantom data showed that a single correction cycle nearly eliminates scatter-induced cupping, and quantitative errors in a water phantom were reduced from approximately 12% to 1% after two correction cycles on experimentally acquired data. The method requires no mechanical modifications to the scanner and adds only a modest overhead to total reconstruction time.

As a computational study, the SPS grid method is benchmarked against the 3-D RL method with one correction cycle. Runtime, image quality recovery, and factors such as parameter tuning and robustness are evaluated between the two techniques. For the initial low-count MC simulations, the respective histories were reduced by a magnitude of 4; e.g., the CT and HN phantoms were simulated with 5×10^7 histories and the pelvic phantom with 1×10^8 histories per angular projection.

Scatter from low-count MC projections is isolated from the primary for scatter estimation. The scatter is denoised via an iterative RL algorithm. The smooth scatter is estimated by the following relationship:

$$\widetilde{p}(x, y, \theta) = G(x, y, \theta) * \lambda(x, y, \theta) \tag{4.10}$$

where $p(x,\ y,\ \theta)$ is the scatter estimate obtained from the convolution of $G(x,\ y,\ \theta)$, a 3-D Gaussian point spread function, with $\lambda(x,\ y,\ \theta)$, the virtual scatter distribution. The input parameters $x,\ y$ are the Cartesian coordinates within each projection, while θ is the corresponding angular slice. λ is determined iteratively following the RL algorithm:

$$\lambda^{(k+1)}(x, y, \theta) = \lambda^{(k)}(x, y, \theta) \frac{G(x, y, \theta) * (\frac{p(x,y,\theta)}{\widetilde{p}^{(k)}(x,y,\theta)})}{\sum_{x,y,\theta} G(x, y, \theta)} \tag{4.11}$$

where $p(x,\ y,\ \theta)$ is the low-count MC scatter distribution and $\widetilde{p}^{(k)}(x, y, \theta)$ is the estimated scatter distribution after k iterations of the 3-D RL fitting. The optimal $\sigma_{x,y}$, σ_{θ}, for $G(x, y, \theta)$ and optimum number of iterations were determined by minimizing the error with the high-count MC scatter distribution. A more thorough analysis and derivation of the optimal parameters can be found in [13], [48], [49]. The total signal is corrected for scatter using the estimated denoised scatter to produce an estimate of the primary measurement for image reconstruction.

4.4 Quantitative Evaluation Metrics

4.4.1 Synthetic CT Phantom Image Metrics

The synthetic cylindrical CT phantom serves as the primary computational benchmark for evaluating the SPS grid technique and provides a controlled reference against which scatter correction performance can be precisely quantified. As the water-equivalent material was defined virtually, the CT phantom's HU value serves as ground truth for comparison. The phantom body is composed of water-equivalent material and is divided into two functional halves along the axial direction: a line-pair slice and a contrast-comparison slice, each targeting a distinct dimension of image-quality degradation caused by scatter.

4.4.2 Clinical Patient Phantom Image Metrics

Unlike a standardized CT phantom's precisely machined cylindrical inserts arranged at known radial positions within a uniform background, the HN phantom and the pelvic phantom pose fundamentally different ROI selection problems. Both phantoms are derived from patient DICOMs converted to PENELOPE 2006-compatible material definitions using a continuous tissue model, resulting in realistic anatomical geometries composed of four simulated materials: adipose, air, soft tissue, and bone of continuously varying densities mapped from the original patient HU values. Because the tissue boundaries are irregular and the background is heterogeneous rather than uniform

water-equivalent material, the CNR measurement strategy was omitted for the clinical phantoms, and HU accuracy was used as a surrogate for evaluating image-quality recovery.

4.4.3 Contrast-to-Noise Ratio (CNR)

As CNR is crucial for soft-tissue contrast and anatomical discrimination, which scatter is known to reduce, a contrast comparison slice was used to measure the SPS grid CNR recovery. The contrast comparison slice (**Figure 4.11**) contains six cylindrical material inserts embedded in the water-equivalent phantom body at equivalent radial positions. The insert materials of lung, cartilage, bone, glandular tissue, brain, and adipose span a wide range of HU values from approximately −1000 HU (lung/air) to +1000 HU (bone), mimicking the tissue diversity encountered in clinical anatomy. The background ROI used for CNR calculation is a water-equivalent region at the same radial distance as the inserts, ensuring that any cupping gradient affects both ROI and background consistently and isolates contrast effects from uniformity effects. CNR for the contrast slice was calculated as follows using the material insert ROIs and background ROI shown in **Figure 4.11**:

$$CNR = \frac{|\mu_{ROI} - \mu_{BG}|}{\sqrt{\sigma_{ROI}^2 + \sigma_{BG}^2}} \tag{4.12}$$

where μ is the mean pixel value of the corresponding material region of interest (ROI) or background and σ is the standard deviation calculated from the corresponding ROI or background.

4.4.4 HU Accuracy

For the CT phantom, the contrast-slice material inserts were used to determine the HU recovery of the SPS grid technique by calculating the RMSE. RMSE quantifies both the systematic offset (bias) and the pixel-to-pixel variability (noise) within each insert, making it a composite metric of accuracy. For the material insert HU RMSE calculations, the following definition is used:

$$RMSE = \frac{\sqrt{\sum_{i=1}^{N} \left\| HU_i - HU_\mu \right\|_2^2}}{N} \tag{4.13}$$

where N is the number of pixels in the ROI, HU_i is the i-th pixel within the ROI, and HU_μ is the mean pixel value in the ROI. As the CT phantom material inserts were well characterized, an aggregate RMSE was calculated from all ROI data points.

For the clinical phantoms, ROIs are placed within regions that are spatially homogeneous in the ground-truth primary-only reconstruction, avoiding tissue boundaries, partial-volume regions, and areas of known anatomical complexity, such as trabecular bone interfaces or adipose-muscle boundaries. The same ROI positions defined on the primary-only reconstruction are applied identically to the uncorrected, SPS-corrected, and 3-D RL corrected reconstructions. This direct comparison ensures that any HU differences between reconstruction conditions are attributable solely to the scatter correction performance rather than to ROI placement variability. RMSE between ROIs was evaluated differently based on the tissue, as homogeneous ROIs could not be selected.

4.4.5 Modulation Transfer Function (MTF)

For spatial image resolution, the line pair slice is used to benchmark and measure the MTF of the SPS grid technique. The line pair slice contains alternating bars of bone and water-equivalent material arranged at progressively increasing spatial frequencies, forming a bar pattern test gauge analogous to the CTP528 high-resolution module of the Catphan 503. The bone material was chosen for its high contrast against the water background, providing high signal modulation even in the presence of scatter-induced HU depression and ensuring the spatial frequency response is measurable across a wide range of scatter contamination levels. The MTF curve was estimated by determining the contrast in the line pairs and curve fitting to the sampled points with the following definition:

$$MTF(f) = \left| \int LSF(x) e^{-2\pi i f x} dx \right| = 100\% * \frac{C(0)}{C(f)}$$
$$C(f) = \frac{HU(f)_{\max} - HU(f)_{\min}}{HU(f)_{\max} + HU(f)_{\min}} \tag{4.14}$$

where $HU(f)_{max}$ is the HU value of the bright peak of the corresponding line pair frequency f and $HU(f)_{\min}$ is the HU value of the dark valley of the corresponding line pair frequency. $C(0)$ is approximately unity when calculated using the background water. For converting the linear attenuation coefficients to CT HU numbers via $HU = (\mu_{material} - \mu_{water})/\mu_{water} \times 1000$, μ_{water} was obtained by averaging the background of the ideal primary-only reconstruction.

4.4.6 HU Uniformity

To evaluate cupping artifact recovery, the UI was measured using peripheral ROIs and a central ROI in a uniform water-equivalent region (**Figure 4.12**). UI directly quantifies cupping artifact severity as a single scalar metric, insensitive to absolute HU levels. The UI is determined by:

$$UI = \frac{1}{N} \sum_{i=1}^{N} \left(\frac{HU_{\mu,i} - HU_{\mu,c}}{HU_{\mu,c} + 1000} \right) x100\% \tag{4.15}$$

where N is the number of peripheral ROIs, $HU_{\mu,i}$ is the mean HU value of the number *i* peripheral ROI, and $HU_{\mu,c}$ is the mean HU value of the central ROI. A UI closer to 0 indicates a smaller cupping artifact effect. Additional cupping artifact evaluation for all simulated phantoms was performed using a central line profile to provide further visualization of the cupping artifact recovery and a quantitative assessment of the HU artifact error.

4.4.7 Algorithm Speed

Although not relevant in evaluating image quality improvements, the runtime of a scatter correction technique is an important practical consideration for clinical translation. All scatter correction

methods were implemented in MATLAB code, running in the same environment and on the same computer. Runtime was evaluated using MATLAB's timer functionality and recorded once each reconstruction technique finished. For the 3-D RL scatter estimation technique, the kernel parameters were precalculated and the smoothing operation of the low-count Monte Carlo was timed. The 3-D RL reconstruction runtime does not include the time for kernel tuning (approximately 6 days) required for optimal image reconstruction quality. For the SPS grid scatter correction technique, a single-pass alternating minimization was utilized with optimized preselected FISTA iterations. For final image reconstruction, both techniques used the same iterative reconstruction cost function:

$$\min_x \frac{1}{2} \|Ax - y_c\|_2^2 + \lambda_1 H_\mu(Dx), \quad \text{s.t. } x \geq 0, \; y_c = \log(I_b) - \log(p) \tag{4.16}$$

where y_c is the corrected line integral, run until convergence. The shared system ran Ubuntu 18.04.06 as the operating system with an Intel Core i7-9700k CPU and GeForce RTX 2070 with 64 GiB of RAM.

4.5 Results

As the phantom voxels were anisotropic along the patient's superior-inferior axis, image analysis was focused on transverse reconstructions. Transverse slice SPS grid reconstruction, primary-only reconstruction, total signal reconstruction, and 3-D RL scatter-corrected reconstruction results for the CT phantom, HN phantom, and pelvic phantom were used to obtain the quantitative image metrics defined above. For patient phantoms, lower sampling densities (0.0506%, 0.0156%, 0.0056%) are omitted due to severe reconstruction artifacts caused by smart pixel interpolation errors. Additional scatter-recovery analysis was performed on the pelvic phantom to profile anatomy-specific secondary artifacts.

4.5.1 CT Phantom Results

As discussed previously, the phantom body is water-equivalent, with well-defined embedded material inserts and a controlled-scatter regime. All four image-quality metrics of UI, MTF, HU RMSE, and CNR can be evaluated simultaneously and compared against a ground-truth primary-only reconstruction, uncorrected reconstruction, and 3-D RL fitting corrected reconstruction (**Figure 4.13**). The chosen grid densities of 1.05% and 0.08% are shown, representing high and low sampling densities, respectively, as reconstruction results for lower sampling densities are unusable in a practical clinical setting. To facilitate evaluation in different clinical niches, soft tissue and bone windows were utilized for all reconstruction visualization results. Visually, the scatter reconstruction techniques exhibit adequate scatter artifact recovery but suffer from residual artifacts due to incomplete correction, with the 1.05% SPS reconstruction showing pronounced artifacts.

4.5.2 Cupping Artifact Evaluation

For qualitative evaluation, a line profile along the central horizontal axis of the line pair phantom (**Figure 4.14**) was plotted, showing the cupping artifact and comparative recovery. Using the primary-only reconstruction as a benchmark, the SPS grid reconstruction results are observed to deviate more from the ground truth, the sparser the SD ($\Delta HU|_{\mu}^{1.05\%}$: 68.7, $\Delta HU|_{\mu}^{0.275\%}$: 32.18, $\Delta HU|_{\mu}^{0.0756\%}$: 32.08, $\Delta HU|_{\mu}^{0.0506\%}$: 32.52, $\Delta HU|_{\mu}^{0.0156\%}$: 34.1, $\Delta HU|_{\mu}^{0.0056\%}$: 29.67). For all SDs tested, omitting 1.05% SD, SPS grid scatter reconstruction results maintained cupping-artifact-free line profiles. For comparison, the 3-D RL method resulted in $\Delta HU|_{\mu}^{3DRL}$: 3.93. UI results are shown in **Table 4.1**, and SPS reconstruction results reveal a decreasing relationship between SD and UI.

4.5.3 Spatial Resolution Evaluation

For evaluating the spatial resolution recovery, the modulation transfer function (**Figure 4.15**) was fitted using data points obtained from the CT phantom line pair test and compared to evaluate

changes in the image spatial resolution of the scatter-corrected results. For SPS reconstruction, the spatial resolution drastically improves from 1.05% to 0.08% SD, while lower SDs (0.05%, 0.0156%, 0.0056%) have minimal improvements in spatial resolution. At 0.08% sampling density, the SPS reconstruction method shows good agreement with the best-case primary signal MTF and 3-D RL method MTF, with a 12.6% deviation from the best-case MTF. The fitted data for SPS corrected reconstruction does not intersect the MTF_{10} threshold, whereas the uncorrected reconstruction fitted data intersects at around 0.436 lp/mm.

4.5.4 HU Accuracy and CNR

In **Figure 4.16**, the contrast rod slice compares the HU values and CNR of inserts composed of lung, bone, cartilage, adipose, glandular, and brain tissue across the sampling densities tested. The background ROI for both CNR and RMSE calculations is placed in the water-equivalent phantom body at the same radial distance as the inserts.

Table 4.1: CT phantom HU accuracy, uniformity, RMSE, and runtime.

Material	**Primary**	**Total**	**3-D RL**	**1.05%**	**0.28%**	**0.08%**	**0.05%**	**0.016%**	**0.006%**
Lung	−318.73 ±0.65	−282.87 ±6.60	−317.90 ±0.62	−303.01 ±10.13	−310.93 ±1.81	−312.44 ±1.02	−312.19 ±0.57	−317.19 ±0.66	−319.70 ±0.69
Cartilage	115.95 ±0.60	−22.61 ±7.38	118.34 ±0.57	108.20 ±0.67	107.25 ±0.61	108.15 ±0.61	108.60 ±0.60	98.57 ±0.70	98.87 ±0.68
Adipose	−90.92 ±0.49	−146.40 ±3.53	−90.81 ±0.47	−87.49 ±0.84	−90.66 ±0.54	−91.59 ±0.57	−91.57 ±0.44	−99.67 ±0.53	−100.17 ±0.52
Bone	1610.41 ±10.39	801.09 ±51.64	1614.39 ±9.95	1522.89 ±33.37	1534.15 ±22.43	1544.60 ±16.58	1548.35 ±11.42	1516.98 ±14.94	1531.32 ±13.77
Brain	42.85 ±0.65	−68.74 ±10.66	42.01 ±0.66	37.50 ±0.89	37.18 ±0.69	37.51 ±0.78	38.16 ±0.69	28.65 ±0.79	27.98 ±0.78
Gland	60.80 ±0.48	−58.16 ±8.01	62.02 ±0.53	55.25 ±1.81	54.27 ±0.62	53.51 ±0.54	55.23 ±0.45	45.66 ±0.58	45.45 ±0.57
Water	20.33 ±1.87	−104.4 ±10.42	21.46 ±2.00	14.34 ±4.38	14.18 ±2.85	14.65 ±2.65	15.22 ±2.67	6.12 ±2.00	5.89 ±1.89
UI	0.59	7.92	0.12	1.36	1.45	1.54	1.88	4.77	5.27
RMSE	0.00	29.41	0.44	2.37	1.72	1.63	1.51	3.53	3.54
Run time			671 s			280 s			

Table 4.1 presents the mean HU value and standard deviation for each insert ROI together with the total HU RMSE across all non-bone ROIs, with the primary-only reconstruction serving as the ideal benchmark. The SPS grid reduces HU RMSE across all tested sampling densities. RMSE improves monotonically from 0.006% to 0.08% SD, where a 94.5% reduction is achieved relative to the uncorrected reconstruction. Beyond 0.08% SD, increasing sampling density causes a sharp increase in RMSE as interpolation errors from denser tube coverage begin to dominate over the gains from improved scatter sampling.

Table 4.2: CNR for the CT phantom inserts.

Material	Primary	Total	3-D RL	1.05%	0.28%	0.08%	0.05%	0.016%	0.006%
Lung	134.66	10.49	129.40	21.87	69.85	89.13	101.00	121.53	126.23
Cartilage	38.72	4.60	37.63	18.57	26.96	28.65	28.60	34.20	36.14
Adipose	47.12	3.01	45.44	19.51	31.00	32.94	34.38	41.87	43.92
Bone	129.73	14.59	133.27	39.96	60.14	79.55	108.85	89.20	97.38
Brain	8.93	1.69	7.71	4.39	6.50	6.66	6.84	8.07	8.25
Gland	17.19	2.51	16.05	6.61	11.56	12.18	12.82	15.34	16.06

Table 4.2 presents the CNR results obtained from the ROIs. In contrast to the RMSE trend, decreasing sampling density yields better CNR performance, with the 0.006% SD result showing CNR comparable to the 3-D RL method. For HU accuracy and CNR performance, the 3-D RL method achieves superior results, although the SPS grid reaches similar CNR for the lower SD configurations.

4.5.5 Head and Neck Phantom

The HN phantom represents clinical anatomy in a low scatter-to-primary ratio regime. The phantom is approximately 25×19×25 cm³ with 1×1×2.5 mm³ voxels, composed of adipose, air, soft tissue, and bone of continuously varying densities derived from the converted patient DICOM HU values. Reconstruction results for the HN phantom are shown in **Figure 4.17**. Similar to the CT phantom results, visual evaluation shows scatter recovery in all scatter-corrected reconstructions, but the secondary artifacts in the 1.05% reconstruction are exacerbated with residual streaks, although less severe in the bone window.

To evaluate the cupping artifact recovery in the HN phantom, line profiles were plotted along sagittal and coronal planes (**Figure 4.18**). Both plots show good agreement with the visual evaluation, with recovery of the cupping artifacts in all scatter-corrected reconstructions along with the residual error in the 1.05% reconstruction. The SPS-corrected reconstruction at 0.08% SD closely aligns with the primary-only reference profile across the full lateral extent of the phantom. The low SPR of the HN anatomy means that the absolute magnitude of the cupping depression is more modest than in larger anatomy such as the pelvic case presented next. Nonetheless, the SPS correction

required to restore HU accuracy is acceptable, particularly in the central soft-tissue regions where the cumulative scatter from multiple projections and heterogeneous anatomy of bone and air produces the greatest error.

HU accuracy was evaluated using ROIs placed within bone, adipose, and soft-tissue regions (**Figure 4.19**). Results confirm the same relationship between sampling density and HU accuracy observed in the CT phantom. The best-performing SD for the HN phantom is 0.08%, achieving RMSE reductions relative to the uncorrected reconstruction of 92.8% for bone, 89.2% for adipose, and 93.7% for soft tissue. For comparison, the 3-D RL method achieves RMSE reductions of 98.5% for bone, 93.5% for adipose, and 99.6% for soft tissue.

Table 4.3: HN phantom HU accuracy and runtime.

Material	Primary	Total	3-D RL	SPS (1.05%)	SPS (0.28%)	SPS (0.08%)
Bone	1871.64 ±119.45	1275.00 ±151.77	1879.47 ±119.00	1801.18 ±160.01	1833.75 ±120.58	1829.54 ±121.71
Adipose	−55.74 ±4.68	−46.02 ±6.99	−55.90 ±5.01	−78.47 ±23.15	−57.35 ±5.98	−55.46 ±4.85
Soft Tissue	76.77 ±1.61	−50.92 ±1.92	76.65 ±1.50	56.74 ±8.62	69.43 ±1.93	68.82 ±2.13
Run time			597 s			224 s

Despite this gap, SPS performance in the low-SPR HN phantom is sufficient that residual HU error after correction falls within acceptable ranges for qualitative identification, as seen in **Table 4.3**, demonstrating suitability for anatomical regions with moderate scatter burdens.

4.5.6 Pelvic Phantom Results

The pelvic phantom represents clinical anatomy in a high SPR regime and is of particular interest for CBCT reconstruction due to the challenging geometry. The phantom is approximately 34×39×39 cm³ with 1×1×3 mm³ voxels, composed of the same four material classes as the HN phantom but with significantly greater soft-tissue volume generating substantially more Compton scatter (**Figure 4.20**). For the pelvic phantom, the 1.05% SPS reconstruction lacks the distinct artifacts the smaller CT and HN phantoms presented, while 3-D RL exhibits large streaking artifacts.

The pelvic phantom line profiles (**Figure 4.21**) confirm recovery of the more severe cupping artifacts characteristic of the high-SPR pelvic anatomy. The SPS-corrected reconstruction broadly recovers the HU depression seen in the uncorrected total signal reconstruction across most of the anatomical cross-section. However, the corrected pelvic profiles exhibit localized deviations from the primary-only reference, manifested as shading artifacts not present in the HN phantom results. These shading artifacts arise from scatter estimation errors in regions where the scatter distribution exhibits higher spatial frequency content than the SPS reconstruction algorithm can resolve at the chosen SD, compounded by reduced particle statistics at the simulated dose level. The SPS-corrected

results nonetheless confirm the method's capacity to correct scatter artifacts in high-SPR anatomical settings, whereas the 3-D RL method is susceptible to noise-induced errors in larger anatomy.

HU accuracy was evaluated using ROIs placed within bone, adipose, and soft-tissue regions (**Figure 4.22**). The best-performing SD for the pelvic phantom is again 0.08%, achieving RMSE reductions relative to the uncorrected reconstruction of 87.5% for bone, 91.7% for adipose, and 90.9% for soft tissue. For comparison, the 3-D RL method achieves RMSE reductions of 94.6% for bone, 93.5% for adipose, and 97.6% for soft tissue.

Table 4.4: Pelvic phantom HU accuracy and runtime.

Material	Primary	Total	3-D RL	SPS (1.05%)	SPS (0.28%)	SPS (0.08%)
Bone	874.93 ±0.51	−75.21 ±8.25	925.60 ±6.17	729.98 ±0.11	741.08 ±2.04	755.82 ±1.77
Adipose	−108.50 ±0.57	−402.15 ±1.45	−89.90 ±4.43	−135.63 ±0.49	−138.17 ±0.45	−132.84 ±0.93
Soft Tissue	38.93 ±2.94	−339.08 ±0.68	30.42 ±0.31	−.65 ±0.95	.79 ±1.31	4.45 ±0.62
Run time			1222 s			694 s

Both scatter correction techniques show worse HU accuracy recovery, with 3-D RL performing slightly better (**Table 4.4**). This decreased performance reflects the greater challenge posed by the high-SPR pelvic scatter distribution, as well as additional factors discussed next. The absolute HU error after SPS correction remains larger in the pelvic phantom than in the HN phantom for equivalent SDs, a direct consequence of the higher SPR producing a larger residual scatter contribution when the correction is imperfect.

4.5.7 Pelvic Scatter Artifact Analysis

As an additional consideration for the results presented in this section, the high susceptibility of the pelvic anatomy to scatter artifacts motivated further work to isolate artifacts and errors caused by the smart pixel tubes. The scan configurations tested varied the scatter contamination percentage and the presence of a pixel dead zone. Worst-case, best-case, and ideal-case scatter contamination were tested at 4%, 2%, and 0%, respectively. Reconstructions with and without smart-pixel dead zones were performed for each scatter contamination variation; for an experimental equivalent, the two-scan protocol would remove the smart-pixel dead zones. To diagnose particle-count artifacts, an additional ideal smart-pixel reconstruction was tested using smoothed projection data, smoothed with a median filter to approximate increased particle statistics. The presented results provide deeper insight into the theoretical limits of the SPS grid and illustrate the effects of practical constraints, such as scatter contamination and interpolation errors, on the technique's efficacy.

Additionally, due to constraints on computational time and hardware, the simulation's pelvic phantom projections have a lower imaging dose than conventional clinical scans. An MC-GPU simulation of 1×10^{12} particles corresponds to about 2.9 cGy for 10% DQE, which is a conservative estimate compared to clinical imaging doses [50], with higher DQE of up to 70% for kV FPDs [51]. Sufficient histories were run to isolate the pelvic shading artifacts as smart pixel dead zone artifacts (**Figure 4.23**), and more work was conducted to distinguish the dead zone artifacts from Monte Carlo noise-induced artifacts. The simulation's insufficient particle statistics resulted in increased noise in projections, compounded interpolation errors from the SPS reconstruction method and decreased fitting accuracy for the 3-D RL scatter correction method. This noise error is exhibited as shaded regions within the central anatomy and low HU depressions in the sagittal line plots (**Figure 4.24**). From the 0% scatter contamination figure, a portion of the HU error from the SPS grid reconstruction is an offset that scales with the scatter contamination.

Table 4.5: Pelvic ROI HU accuracy across scatter-contamination and scan-mode protocols.

ROI	Primary	Total	3-D RL	SPS (1.05%)	SPS (0.28%)	SPS (0.08%)
4%-1scan						
Bone	874.93 ± 0.51	−75.21 ± 8.25	925.60 ± 6.17	729.98 ± 0.11	741.08 ± 2.04	755.82 ± 1.77
Adipose	−108.50 ± 0.57	−402.15 ± 1.45	−89.90 ± 4.43	−135.63 ± 0.49	−138.17 ± 0.45	−132.84 ± 0.93
Soft	38.93 ± 2.94	−339.08 ± 0.68	30.42 ± 0.31	−.65 ± 0.95	.79 ± 1.31	4.45 ± 0.62
4%-2scan						
Bone	874.93 ± 0.51	−75.21 ± 8.25	925.60 ± 6.17	747.15 ± 0.48	747.60 ± 0.35	768.01 ± 0.55
Adipose	−108.50 ± 0.57	−402.15 ± 1.45	−89.90 ± 4.43	−137.83 ± 3.50	−140.50 ± 2.38	−131.24 ± 0.68
Soft	38.93 ± 2.94	−339.08 ± 0.68	30.42 ± 0.31	2.85 ± 2.47	9.16 ± 1.53	4.48 ± 2.56
2%-1scan						
Bone	874.93 ± 0.51	−75.21 ± 8.25	925.60 ± 6.17	792.41 ± 0.12	803.85 ± 0.40	820.54 ± 0.36
Adipose	−108.50 ± 0.57	−402.15 ± 1.45	−89.90 ± 4.43	−121.84 ± 0.41	−124.24 ± 0.37	−118.04 ± 0.70
Soft	38.93 ± 2.94	−339.08 ± 0.68	30.42 ± 0.31	18.86 ± 0.65	20.47 ± 0.65	24.37 ± 0.77
2%-2scan						
Bone	874.93 ± 0.51	−75.21 ± 8.25	925.60 ± 6.17	802.89 ± 0.97	798.19 ± 0.24	816.66 ± 0.37
Adipose	−108.50 ± 0.57	−402.15 ± 1.45	−89.90 ± 4.43	−124.43 ± 0.49	−123.93 ± 0.74	−118.51 ± 0.73
Soft	38.93 ± 2.94	−339.08 ± 0.68	30.42 ± 0.31	18.78 ± 0.58	20.83 ± 0.43	23.88 ± 1.00
0%-1scan						
Bone	874.93 ± 0.51	−75.21 ± 8.25	925.60 ± 6.17	874.77 ± 9.89	890.74 ± 9.74	909.18 ± 10.43
Adipose	−108.50 ± 0.57	−402.15 ± 1.45	−89.90 ± 4.43	−106.37 ± 3.46	−111.71 ± 1.67	−100.18 ± 1.11
Soft	38.93 ± 2.94	−339.08 ± 0.68	30.42 ± 0.31	33.78 ± 3.26	48.68 ± 3.63	46.10 ± 2.21
0%-2scan						
Bone	874.93 ± 0.51	−75.21 ± 8.25	925.60 ± 6.17	882.51 ± 5.08	882.57 ± 7.38	908.58 ± 10.09
Adipose	−108.50 ± 0.57	−402.15 ± 1.45	−89.90 ± 4.43	−109.20 ± 3.31	−111.88 ± 1.94	−100.15 ± 1.15
Soft	38.93 ± 2.94	−339.08 ± 0.68	30.42 ± 0.31	43.88 ± 4.98	51.49 ± 5.48	46.23 ± 2.04
Smooth 0%-2scan						
Bone	874.93 ± 0.51	−75.21 ± 8.25	925.60 ± 6.17	864.15 ± 0.33	870.36 ± 0.29	895.54 ± 0.39
Adipose	−108.50 ± 0.57	−402.15 ± 1.45	−89.90 ± 4.43	−108.98 ± 0.31	−109.66 ± 0.23	−101.66 ± 1.29
Soft	38.93 ± 2.94	−339.08 ± 0.68	30.42 ± 0.31	43.50 ± 0.95	42.91 ± 1.15	45.83 ± 0.84

Table 4.6: Pelvic ROI RMSE across scatter-contamination and scan-mode protocols.

ROI	Total	3-D RL	SPS (1.05%)	SPS (0.28%)	SPS (0.08%)
4%-1scan					
Bone	950.16	51.00	144.95	133.86	119.11
Adipose	293.66	19.07	27.14	29.68	24.36
Soft	378.01	9.02	39.64	38.21	34.57
4%-2scan					
Bone	950.16	51.00	127.78	127.33	106.92
Adipose	293.66	19.07	29.53	32.08	22.75
Soft	378.01	9.02	36.34	29.96	34.61
2%-1scan					
Bone	950.16	51.00	82.52	71.08	54.39
Adipose	293.66	19.07	13.35	15.76	9.57
Soft	378.01	9.02	20.22	18.64	14.79
2%-2scan					
Bone	950.16	51.00	72.04	76.74	58.27
Adipose	293.66	19.07	15.95	15.45	10.04
Soft	378.01	9.02	20.33	18.31	15.27
0%-1scan					
Bone	950.16	51.00	9.46	18.22	35.58
Adipose	293.66	19.07	3.93	3.64	8.37
Soft	378.01	9.02	7.32	11.04	7.93
0%-2scan					
Bone	950.16	51.00	8.95	10.29	34.93
Adipose	293.66	19.07	3.32	3.88	8.41
Soft	378.01	9.02	8.22	14.36	8.02
Smooth 0%-2scan					
Bone	950.16	51.00	10.79	4.58	20.62
Adipose	293.66	19.07	0.70	1.28	6.92
Soft	378.01	9.02	5.52	5.13	7.45

Further analysis introduces seven different scan protocols in **Tables 4.5 and 4.6**, with the scatter contamination percentage (#%) included and whether smart-pixel dead zones were simulated, denoted as 1- or 2-scans as previously introduced. As observed from **Tables 4.5 and 4.6**, together with **Figure 4.24**, there is a consistent HU error that is linearly related to scatter contamination and decreases with decreasing SD. However, omitting smart pixel dead zones, in both 4% and 2% scatter contamination results, the higher SDs of 1.05% and 0.28% do not show significant differences. At

0% scatter contamination, regardless of smart-pixel dead zones, the expected behavior of higher SD resulting in improved imaging is recovered. Smoothing the MC projection data improved SPS reconstruction performance, with both lower RMSE and a tighter spread in HU values due to lower standard deviation.

4.5.8 Runtime Comparison with the 3-D RL Fitting Method

For computational benchmarking, the 0.08% SPS grid reconstruction was compared with the 3-D RL denoising using low-count MC scatter projections, with time comparisons in **Tables 4.1, 4.3, and 4.4**. From the tables, SPS reconstruction is faster than 3-D RL reconstruction. For further consideration, the SPS reconstruction method is independent of prior patient information, whereas the 3-D RL method requires a low-dose preliminary CBCT scan to model via Monte Carlo simulation and lengthy kernel tuning. Additionally, large anatomy requires more histories to simulate for sufficient statistics even when downsampling scatter statistics for the low-count projections, leading to order-of-magnitude decreases in speed.

4.6 Discussion

In this study, a novel scatter artifact reduction method for CBCT using a hybrid scatter correction approach is presented. A flexible detector-mounted primary signal sampling device was designed that allows for varying sampling densities and patterns. The SPS design leverages the low spatial frequency nature of the scatter signal in CBCT to reconstruct the scatter from sparsely sampled smart pixels [52], [53]. The SPS grid simultaneously measures the total signal and the smart-pixel primary signal while attenuating total photon transmission by only a small amount. An alternating iterative scheme is presented to reconstruct the scatter signal from the sparsely sampled smart pixels. Given the full signal measurement, the complete scatter signal is recovered for image reconstruction. The feasibility and performance of the sparse primary reconstruction strategy were assessed via Monte Carlo simulations, and the method was benchmarked against software-based scatter correction using

3-D RL fitting to evaluate its applicability. The CT phantom results from the SPS reconstruction demonstrate the technique's ideal performance with extremely sparse primary sampling and effective removal of scatter artifacts that can hinder spatial and contrast resolution.

In this study, SPS sampling densities of 1.05%, 0.275%, 0.0756%, 0.0506%, 0.0156%, and 0.0056% were selected to investigate the relationship between sampling density and reconstruction performance, as well as to determine any potential sampling limits of the method. The 0.08% SD provided the best results, optimizing UI, RMSE, and CNR, resulting in a 94.5% reduction in total image HU RMSE for the CT phantom. For comparison, the 3-D RL method achieves a 98.51% reduction in RMSE. For the SPS reconstruction technique, there is a tradeoff between lower and higher SD, as depicted in **Figure 4.25**, with lower SD providing smoother scatter profiles but less sensitivity to small anatomical features, while higher SD provides better estimates of higher frequency features in the scatter profile at the cost of incurring more interpolation errors from increased tube coverage. As a result, there is a limit to the improvement in image quality from decreasing SD, as the SPS reconstruction results' uniformity and HU accuracy improve to 0.08% before diminishing. Metrics sensitive to tube interpolation errors, such as CNR and MTF, do not show an optimal SD, with both metrics improving as the SD decreases.

For the HN and pelvic phantoms, **Tables 4.3 and 4.4** show the best-case RMSE reduction for the 0.08% SPS grid reconstruction. A roughly 90% RMSE reduction was seen in both HN and pelvic phantoms, although the absolute HU error was significantly larger in the pelvic phantom from the higher SPR. For both HN and pelvic phantoms, extreme streak artifacts from interpolation errors at sampling densities at or below 0.08% render lower sampling densities unviable. From **Figure 4.23**, higher sampling densities resulted in patterned HU errors due to data insufficiency and interpolation errors caused by smart-pixel dead zones. The errors are further exacerbated in the soft-tissue window and significantly degrade image quality. In **Figure 4.20**, the higher sampling density reconstruction for the pelvic phantom did not have severe artifacts, suggesting that larger anatomies are less susceptible to interpolation errors. These findings suggest the sampling limit for the SPS method is around 0.08% of detector pixels. Additionally, the shading behavior in the

pelvic phantom can be understood and analyzed using the data from **Figure 4.25** and **Tables 4.5 and 4.6**, respectively. Although the CT phantom scatter profile is a simplification of the pelvic scatter profile, the 0.08% SD can be understood as a composite of both 1.05% and 0.006% scatter behavior, with some regions drastically underestimating and overestimating the scatter profile. Scatter contamination introduces an offset in the scatter contribution, which is added to the primary fluence, leading to a consistent shift in HU error. In **Figures 4.20 and 4.24**, the combined HU errors for peaks and depressions, along with the HU underestimation, result in dark-shaded regions. The SPS method is also noise-sensitive due to the nature of downsampling detector measurements, and the improvement in image quality from more particle statistics is evident in **Tables 4.5 and 4.6** for the smoothed results.

When comparing image improvements achieved using conventional 1-D anti-scatter grids (ASGs) in CBCT, HU inaccuracy can be reduced from approximately 40% to 12%, depending on the grid ratio. However, increased grid ratios reduce CNR by up to 65% and require an imaging dose approximately 1.6 to 2.5 times higher to restore soft-tissue CNR [54]. The SPS method at 0.08% SD achieves better HU accuracy, comparable worst-case HU accuracy to 1-D ASG, and improves worst-case CNR by 60% without increasing the imaging dose.

4.6.1 Evaluating the SPS Grid Against the Comparative 3-D RL Method

The comparative 3-D RL fitting method circumvents the limitations of slow MC-based scatter correction and dataset-dependent learning models by employing rapid MC acceleration via low-count simulations derived from prior patient scans. In an *in-silico* environment, the 3-D RL method demonstrates superior performance compared to the SPS method. However, in clinical implementation, the SPS approach does not require prior knowledge of the scanned anatomy, whereas the 3-D RL fitting method necessitates a preliminary scan to generate an accurate Monte Carlo geometry for low-photon-count simulations. Additionally, the 3-D RL fitting method requires extensive iterative kernel tuning for optimal scatter-fitting results, with each kernel requiring adjustment for individual anatomical variations. This requirement reduces the practical feasibility of

the 3-D RL method and similar software-based approaches. As seen from the pelvic data, the 3-D RL method struggles with high SPR anatomy, as more particle statistics are required for accurate fitting, thereby increasing simulation time by orders of magnitude. In these experiments, the tuning step typically required 24 hours or more and scaled with the parameter space and projection dimensions, which significantly reduces the practicality of this method in the clinical setting. This SPR limitation is not encountered in the original publication [13], as the original technique assumed small animal anatomy for micro-CT. In contrast, the SPS method seamlessly integrates projection acquisition, scatter correction, and image reconstruction into a unified workflow.

4.7 Assumptions and Limitations of the Computational Study

It is assumed that the scatter contamination is consistent within each smart pixel regardless of placement and that the pixels covered by the tungsten tubes are completely blocked. The scatter contamination within each smart pixel may not be constant due to grid placement and can be improved. A different simulation toolkit, such as FASTCAT, may provide better results by explicitly simulating the tungsten tubes in the scanning geometry, which is not possible with MC-GPU's strict geometry definitions. In the current iteration of the SPS method, pixels covered by the smart pixel tubes are discarded; however, if the attenuation is consistent, the total signal may be recoverable and used to improve interpolation.

Penumbral effects were ignored, as preliminary simulations show minimal shadows with misalignments, but errors during gantry rotation and tube alignment may compound, resulting in reconstruction artifacts. Future work is needed to account for the error accumulation, with experimental studies to determine the severity and potential artifacts. Potential solutions to address error accumulation include increasing the tube alignment error tolerance by adjusting the inner and outer diameters.

Further improvements in the SPS technique, particularly concerning sampling patterns, are possible. While this study employed a uniform smart-pixel sampling pattern, a non-uniform

approach tailored to the scatter distribution could improve reconstruction accuracy. For example, higher sampling densities toward the detector center may yield improved reconstructions without increasing the total number of collimators. Optimization of smart pixel density and pattern remains a topic for future investigation. In this study, straightforward regularized reconstruction and ray tracing simulation were performed to focus on the gains from SPS sparse sensing; however, the reconstruction results could conceivably be improved by combining with Monte Carlo or Boltzmann forward projection.

These computational findings establish the theoretical performance envelope of the SPS grid technique under idealized geometric and statistical conditions. The next chapter translates the technique into hardware, presenting the benchtop CBCT system, the prototype SPS grid, and the early experimental validation results, against which the simulated upper bound established here serves as the reference for interpreting the experimental performance gap.

4.8 Figures

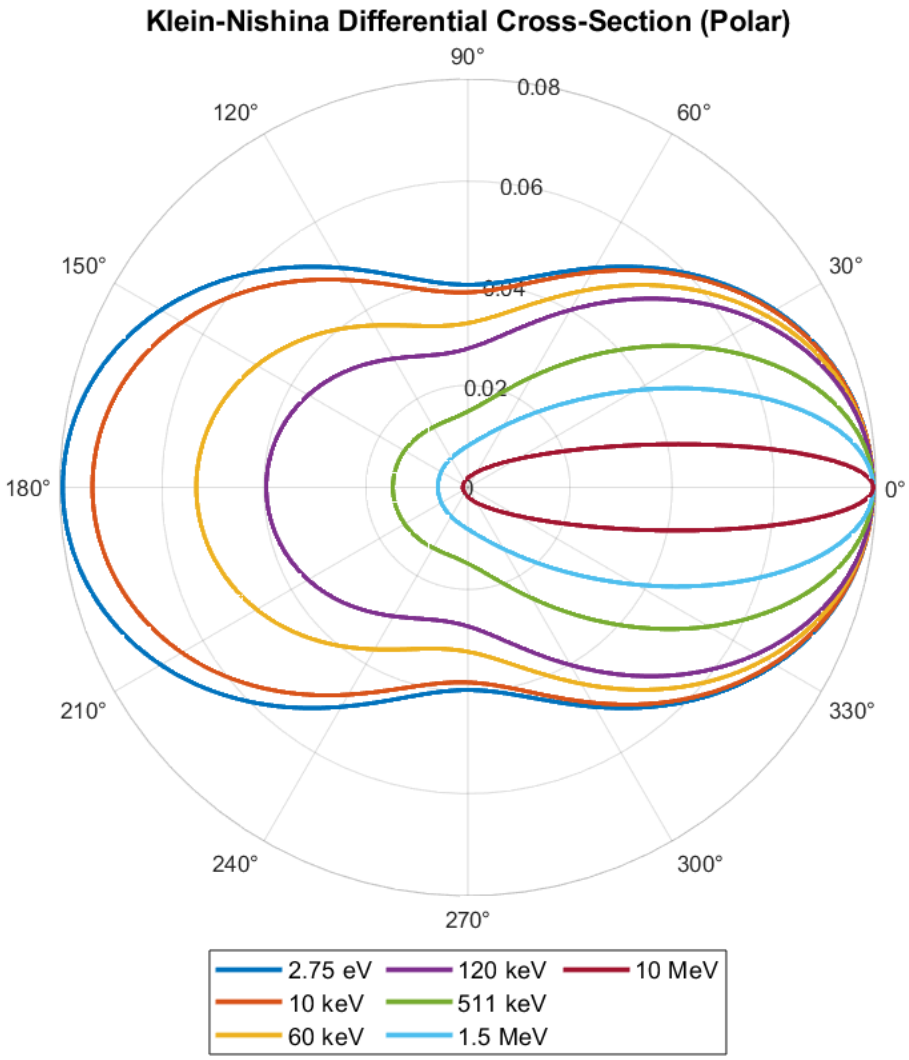

Figure 4.1: Klein-Nishina angular distribution of Compton scatter for the CBCT energy range (10–120 keV).

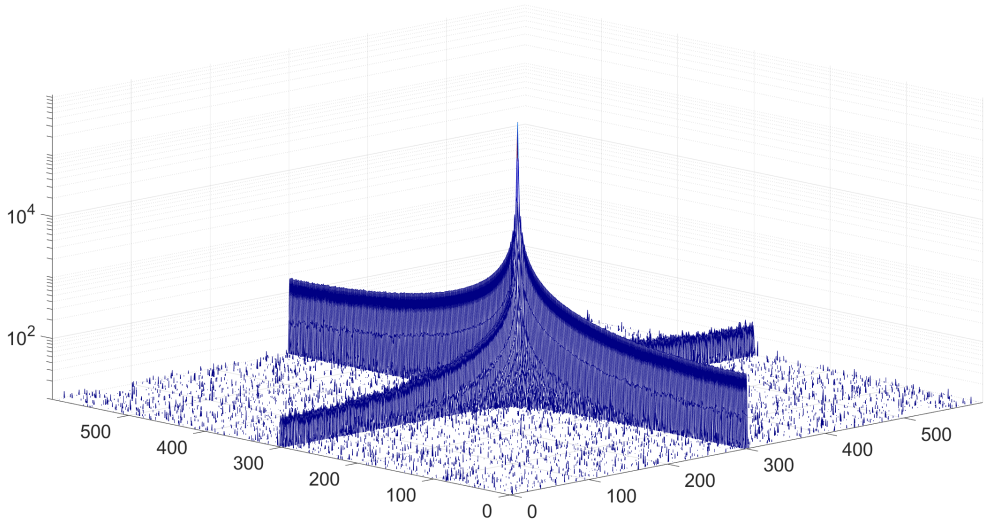

Figure 4.2: Fourier-domain (Kx, Ky) representation of the scatter distribution from Monte Carlo pelvis data, showing low-frequency dominance.

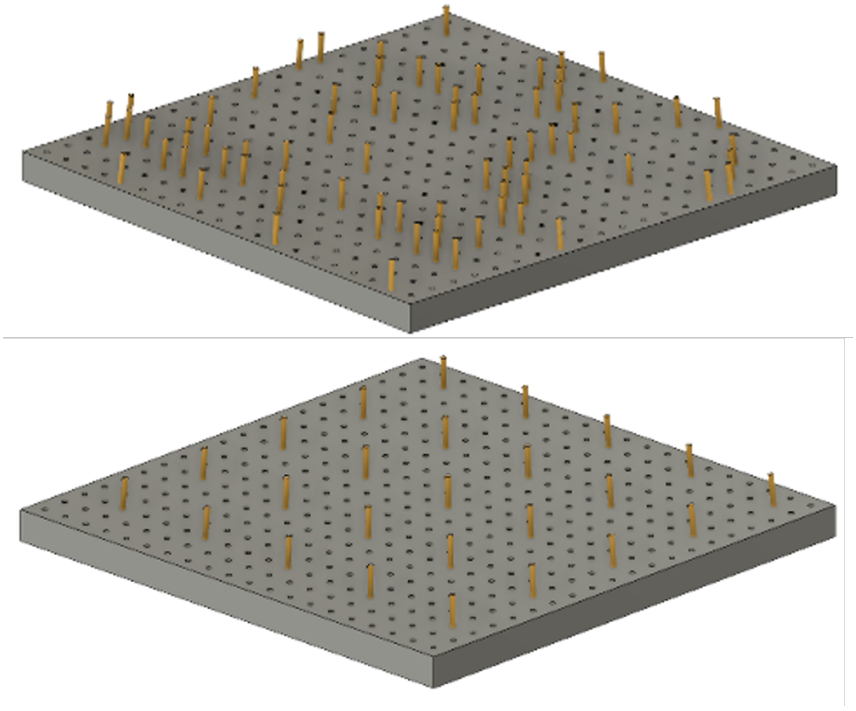

(a)

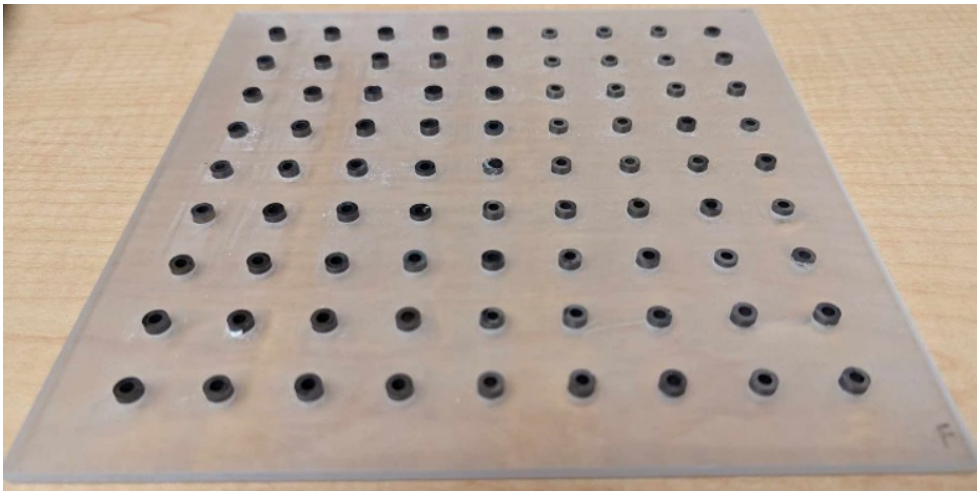

(b)

Figure 4.3: (a) 3-D model of the SPS grid designed for tungsten tube inserts, shown aligned for 20 cm SSD in regular and random sampling configurations. (b) Prototype sparse primary sampling grid holder plate with tungsten tube inserts.

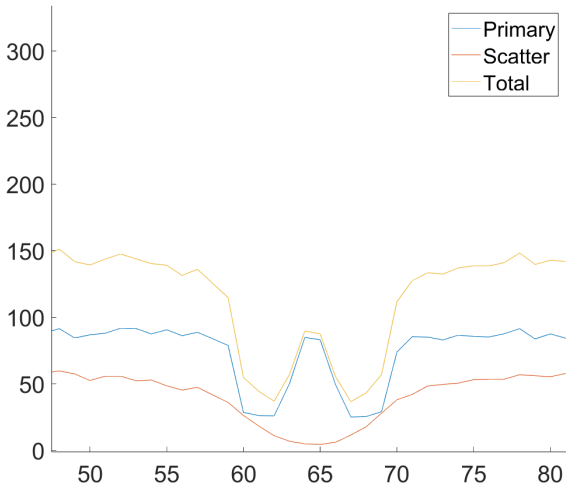

Figure 4.4: Energy deposition surrounding one individual tungsten tube (1 mm ID, 4 mm OD, 2 cm tall) modeled via TOPAS with 1×10^{12} histories. Primary, scatter, and total photon energy deposition are shown.

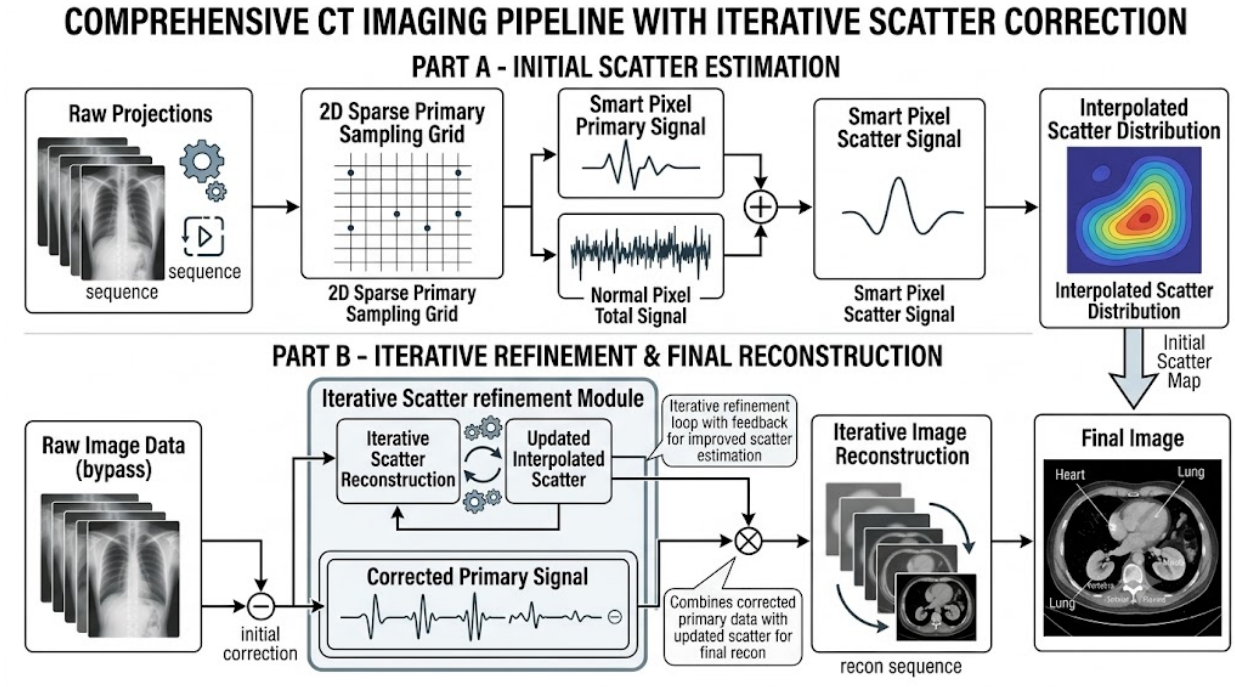

Figure 4.5: Pipeline for scatter reconstruction and correction via the SPS grid.

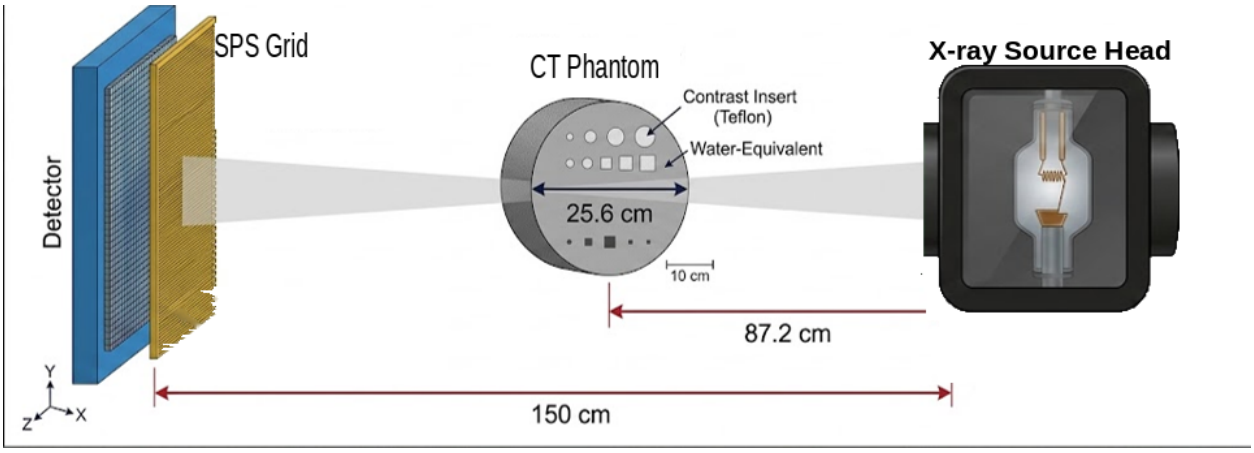

Figure 4.6: Monte Carlo CBCT scanning geometry with source-to-detector distance of 150 cm and phantom-dependent isocenter distances.

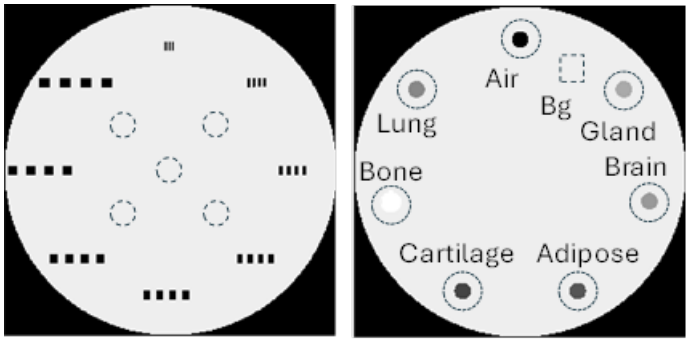

Figure 4.7: Digital CT phantom for use in Monte Carlo. (a) Line pair comparison slice with ROI used for HU uniformity. (b) Contrast comparison slice with ROI and water background used for HU comparisons and CNR calculations. Pixel values shown are not correlated to density and are purely visual.

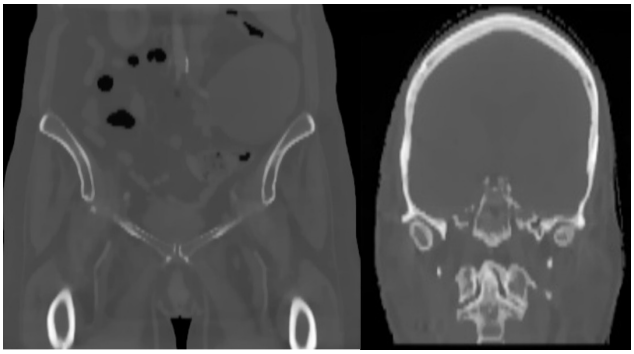

Figure 4.8: Sagittal view of the clinical head-and-neck and pelvic phantoms. Note: phantom voxels are anisotropic, and the image is stretched along the superior-inferior direction to match clinical dimensions.

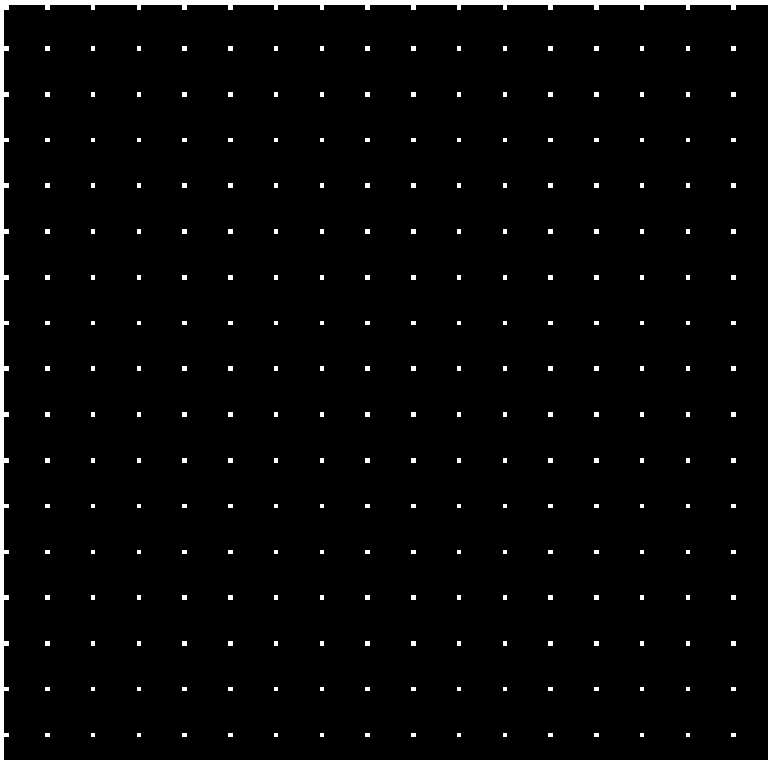

Figure 4.9: Uniform smart pixel grid pattern used in the MC-GPU simulations.

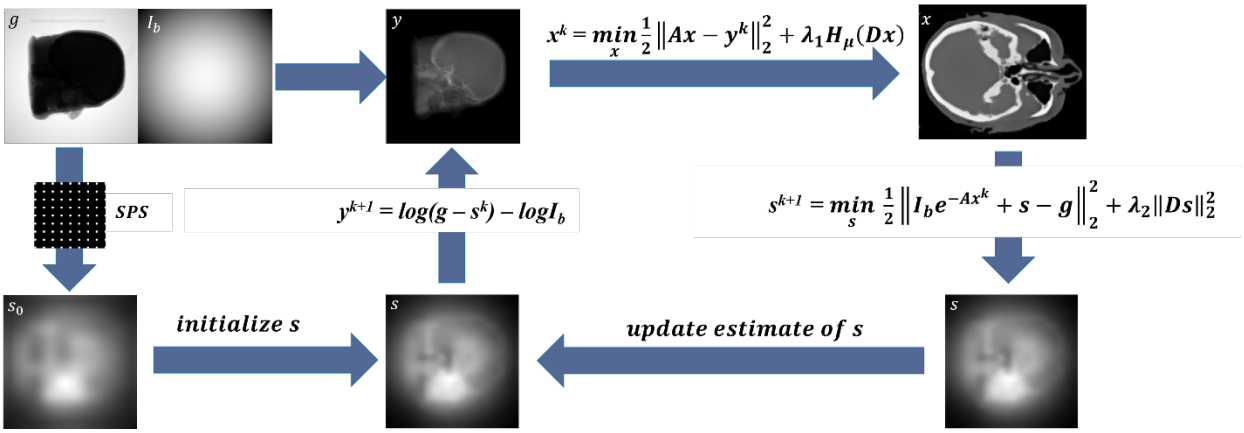

Figure 4.10: Schema of the alternating minimization algorithm, showing the conversion between the line integral and photon count along with the alternating minimizations. g is the total signal measurement, I_b is the blank field measurement, s_0 is the initial scatter interpolation obtained from the SPS grid, s is the reconstructed scatter signal, y is the corrected line integral, H_μ is the Huber penalty function, and x is the reconstructed image.

Figure 4.11: Contrast comparison slice with material insert ROIs and water-equivalent background ROI used for CNR and HU accuracy calculations.

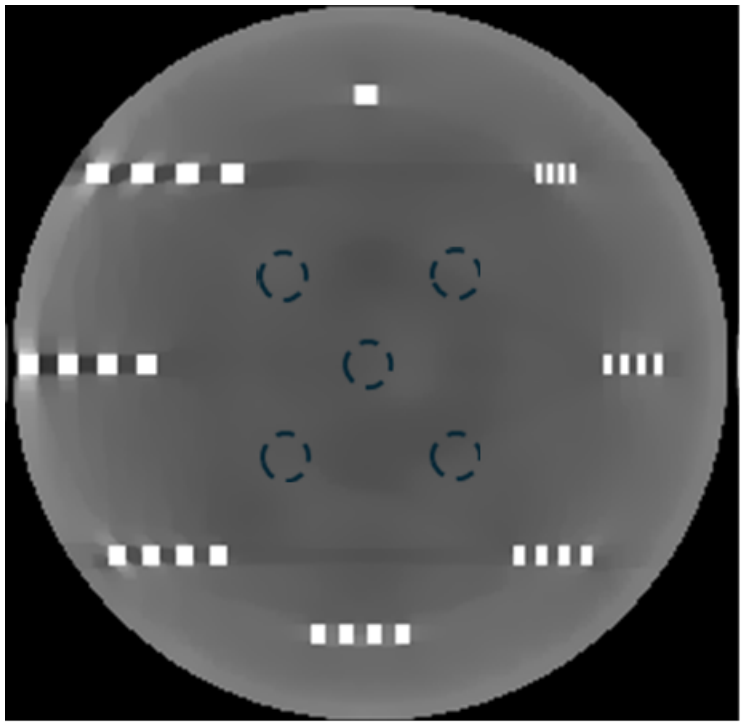

Figure 4.12: Peripheral and central ROIs in the uniform water-equivalent region used for the uniformity index measurement.

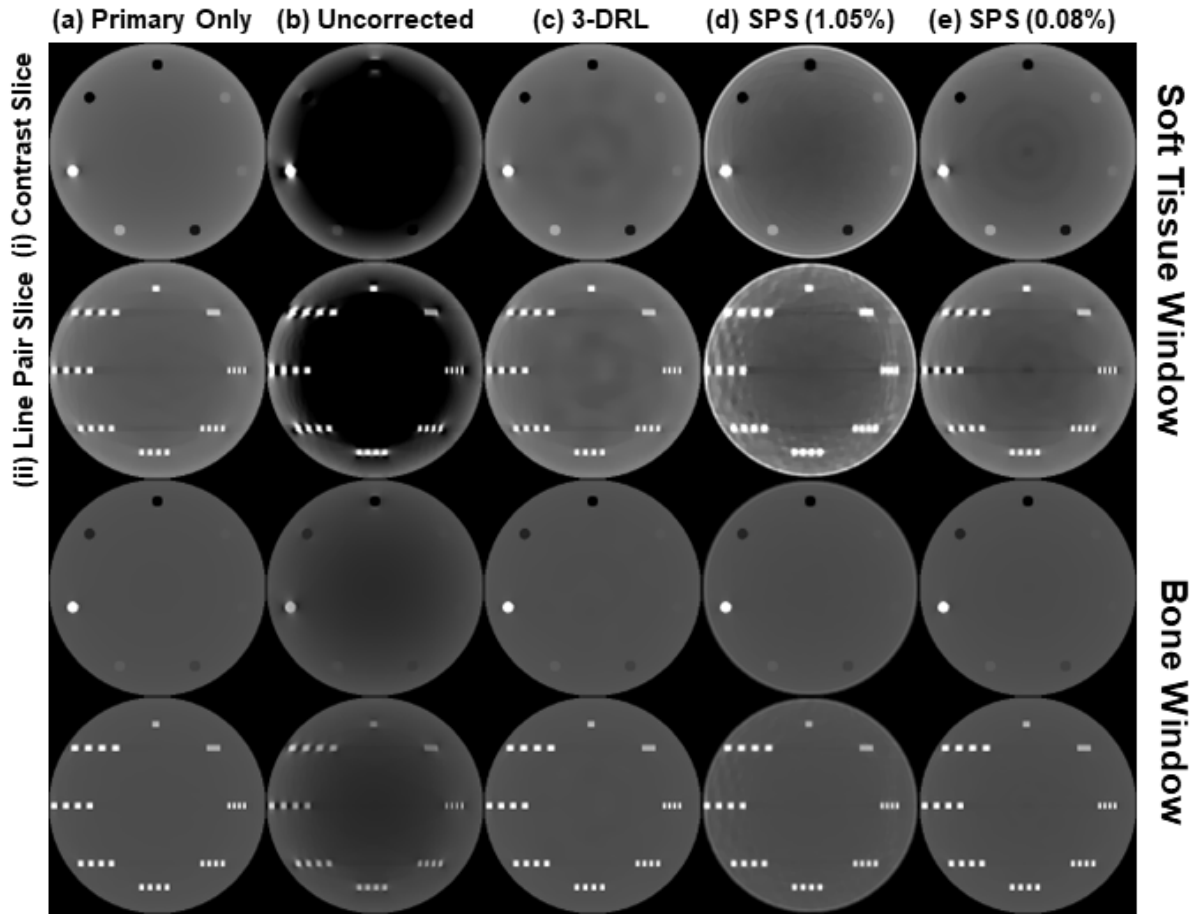

Figure 4.13: CT phantom reconstruction results showing both line pair and contrast slices comparing (a) primary-only reconstruction, (b) uncorrected total-signal reconstruction, (c) 3-D RL fitting corrected reconstruction, and (d-e) SPS grid reconstruction at representative sampling densities, displayed in clinical soft-tissue ([−120, 240] HU) and bone ([−600, 1400] HU) windows.

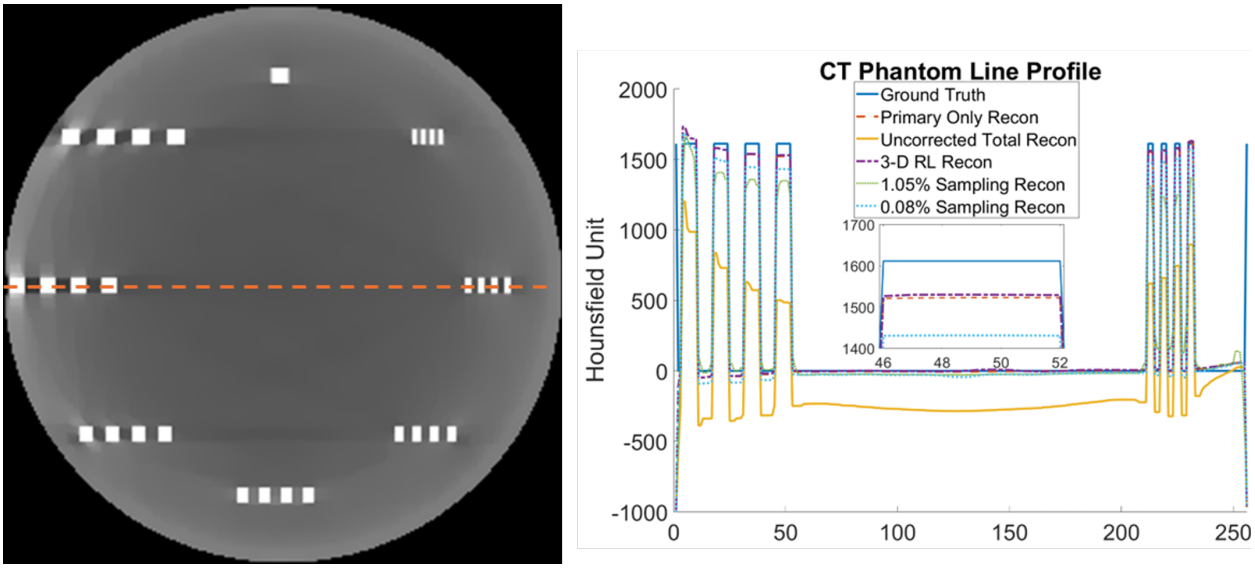

Figure 4.14: Line profile taken from the CT phantom showing the scatter recovery performance of each case: ground truth phantom using absolute HU values, primary-only signal reconstruction, SPS corrected signal reconstruction, 3-D RL scatter corrected reconstruction, and uncorrected total signal reconstruction.

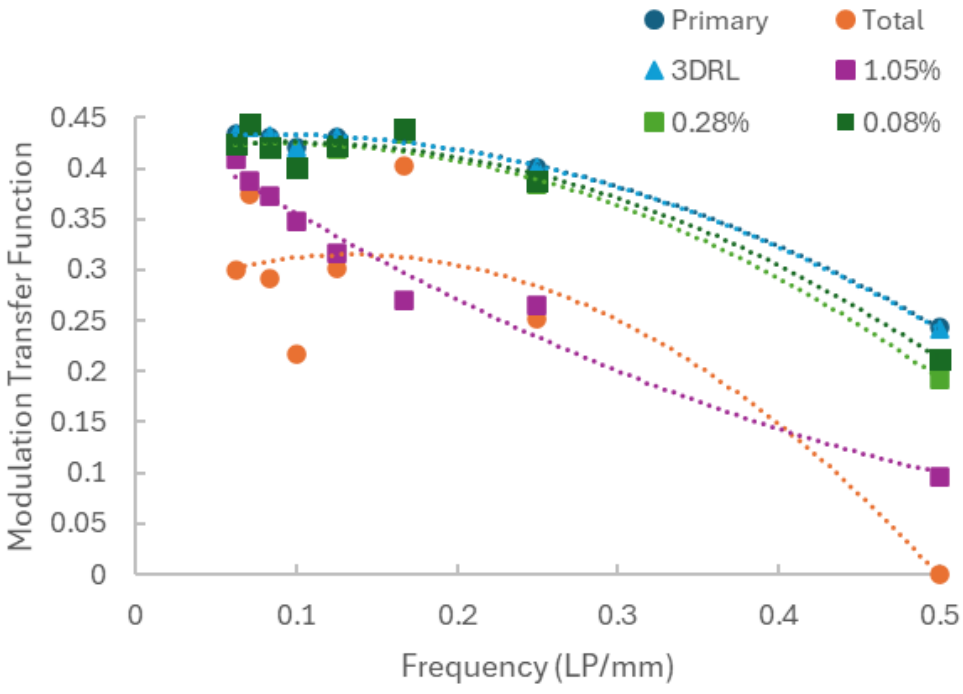

Figure 4.15: Fitted MTF curves from the CT phantom line pair test comparing SPS corrected reconstructions across sampling densities against primary-only, uncorrected, and 3-D RL reconstructions.

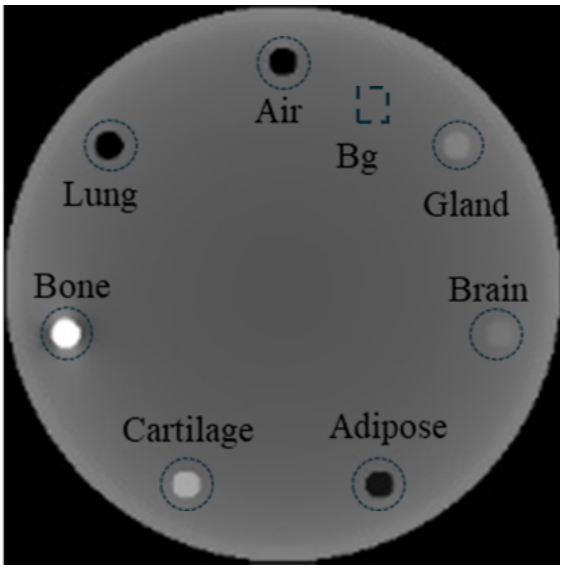

Figure 4.16: Contrast rod slice comparing HU values and CNR of inserts composed of lung, bone, cartilage, adipose, glandular, and brain tissue across the sampling densities tested.

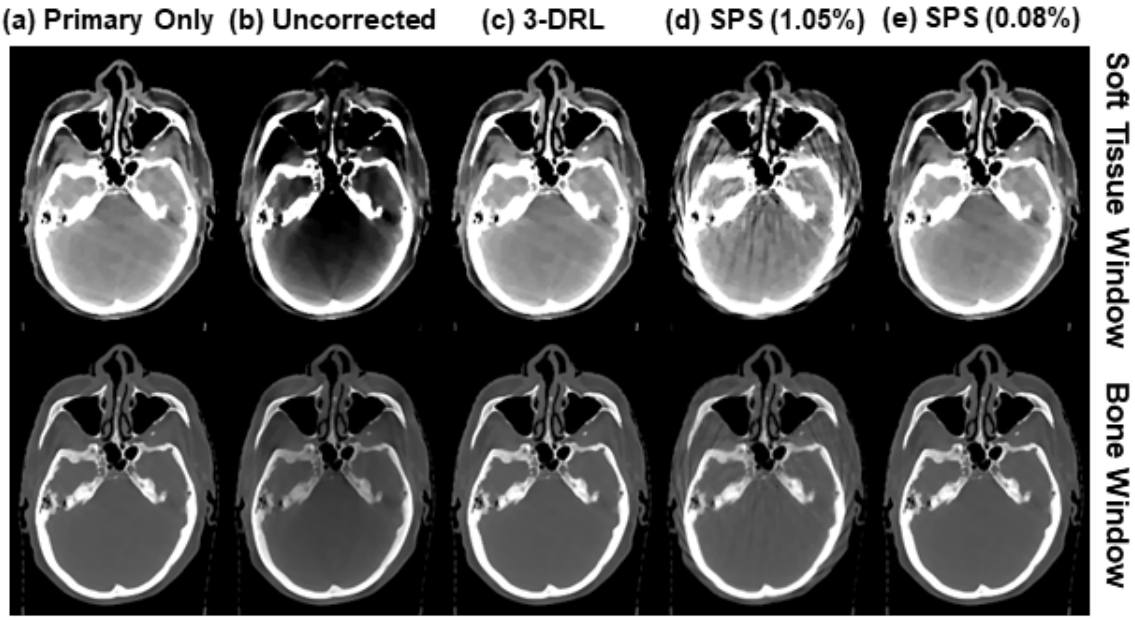

Figure 4.17: HN phantom reconstruction results across reconstruction conditions, displayed in clinical soft-tissue ([−120, 240] HU) and bone ([−600, 1400] HU) windows. The (a) primary-only reconstruction serves as the ideal reconstruction benchmark; the (b) uncorrected total-signal reconstruction is the worst case, exhibiting a large HU cupping artifact; the (c) 3-D RL reconstruction is the computational benchmark reference method; the (d) 1.05% SPS reconstruction shows the worst SPS performance, with additional streak artifacts; and the (e) 0.08% SPS reconstruction shows the best scatter recovery of the SPS method.

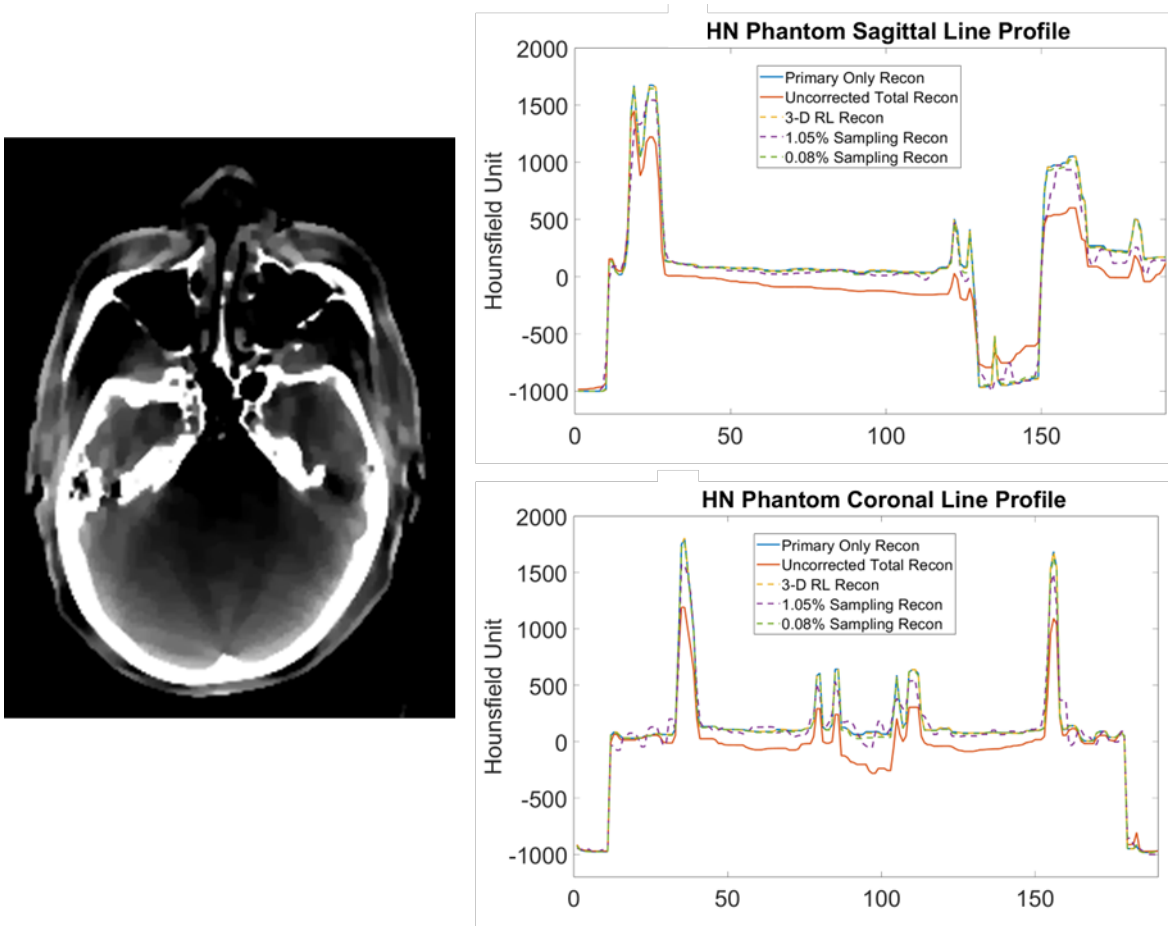

Figure 4.18: Line profiles along sagittal and coronal planes of the HN phantom showing cupping artifact recovery along the central pixels of figure and corresponding line plots.

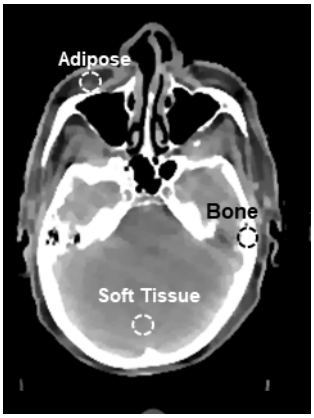

Figure 4.19: ROI-based HU accuracy evaluation for the HN phantom (bone, adipose, and soft-tissue regions).

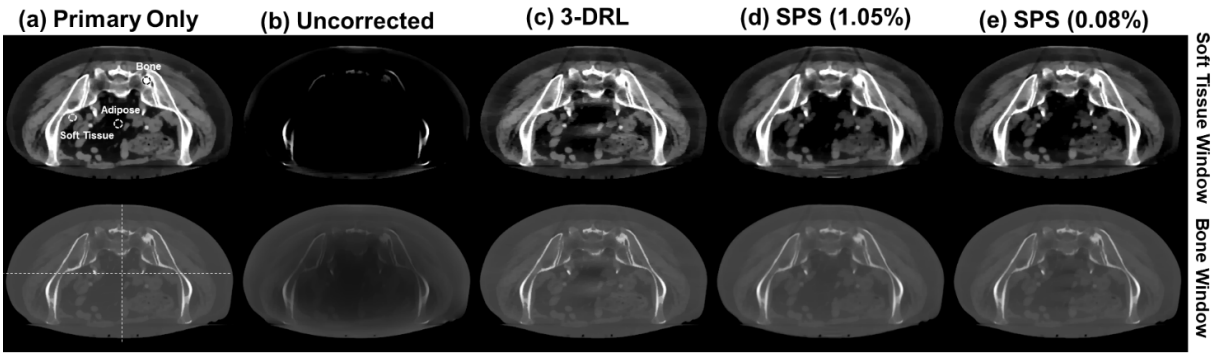

Figure 4.20: Pelvic phantom reconstruction results across reconstruction conditions, displayed in clinical soft-tissue ([−120, 240] HU) and bone ([−600, 1400] HU) windows. The (a) primary-only reconstruction serves as the ideal reconstruction benchmark; the (b) uncorrected total-signal reconstruction is the worst case, exhibiting a large HU cupping artifact; the (c) 3-D RL reconstruction is the computational benchmark reference method and exhibits streaking artifacts; both (d) 1.05% and (e) 0.08% exhibit cupping artifact recovery.

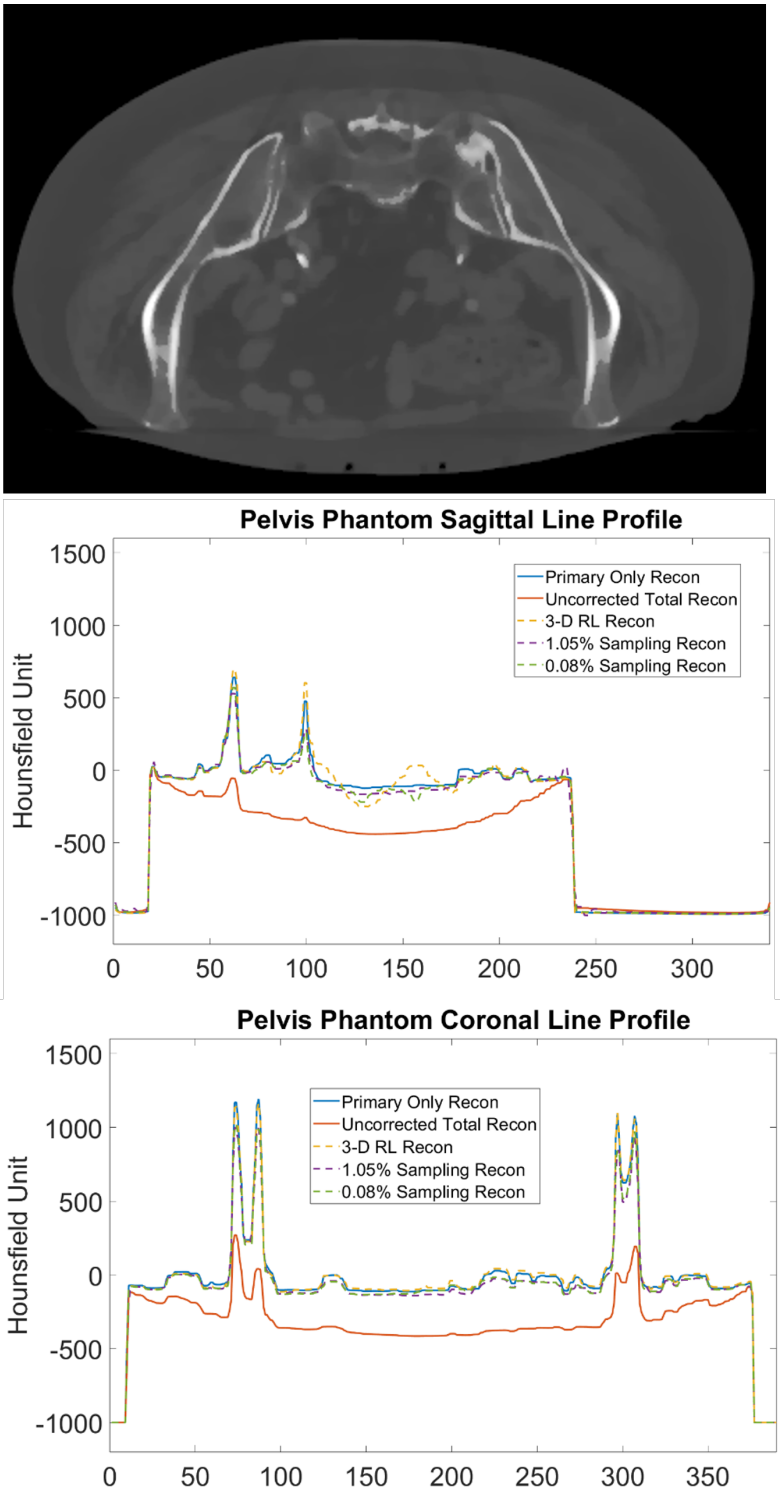

Figure 4.21: Pelvic phantom line profiles confirming recovery of the severe cupping artifacts characteristic of high-SPR anatomy.

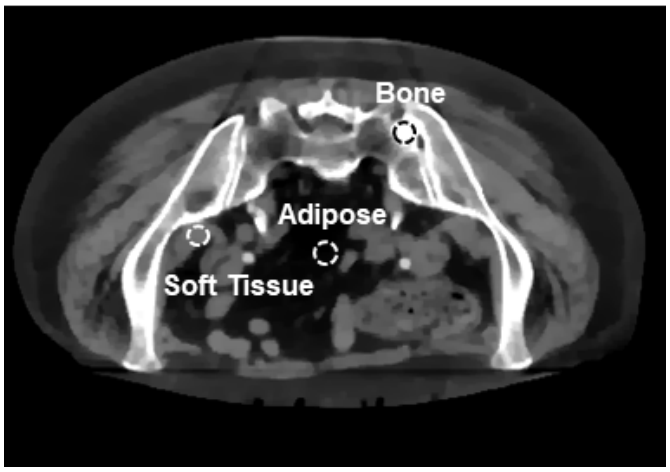

Figure 4.22: ROI-based HU accuracy evaluation for the pelvic phantom (bone, adipose, and soft-tissue regions).

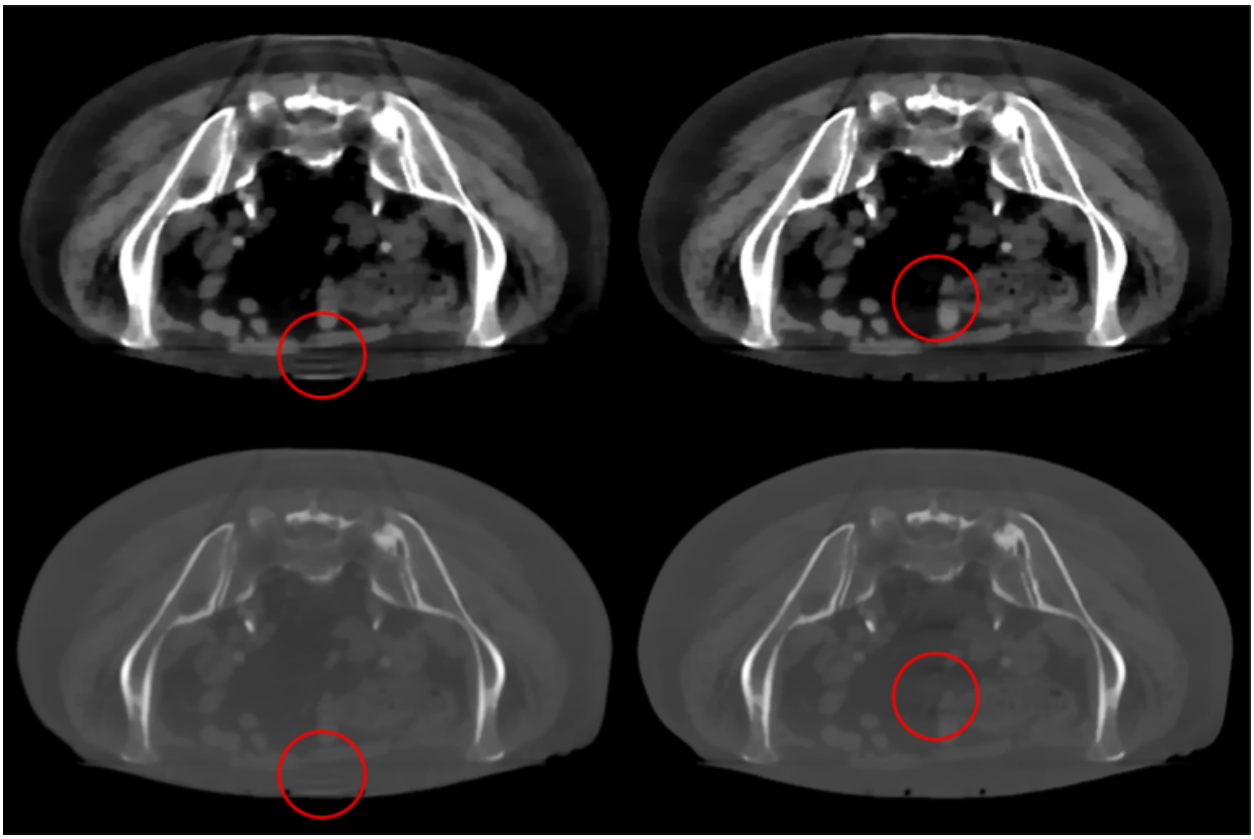

Figure 4.23: Isolation of pelvic shading artifacts as smart pixel dead zone artifacts.

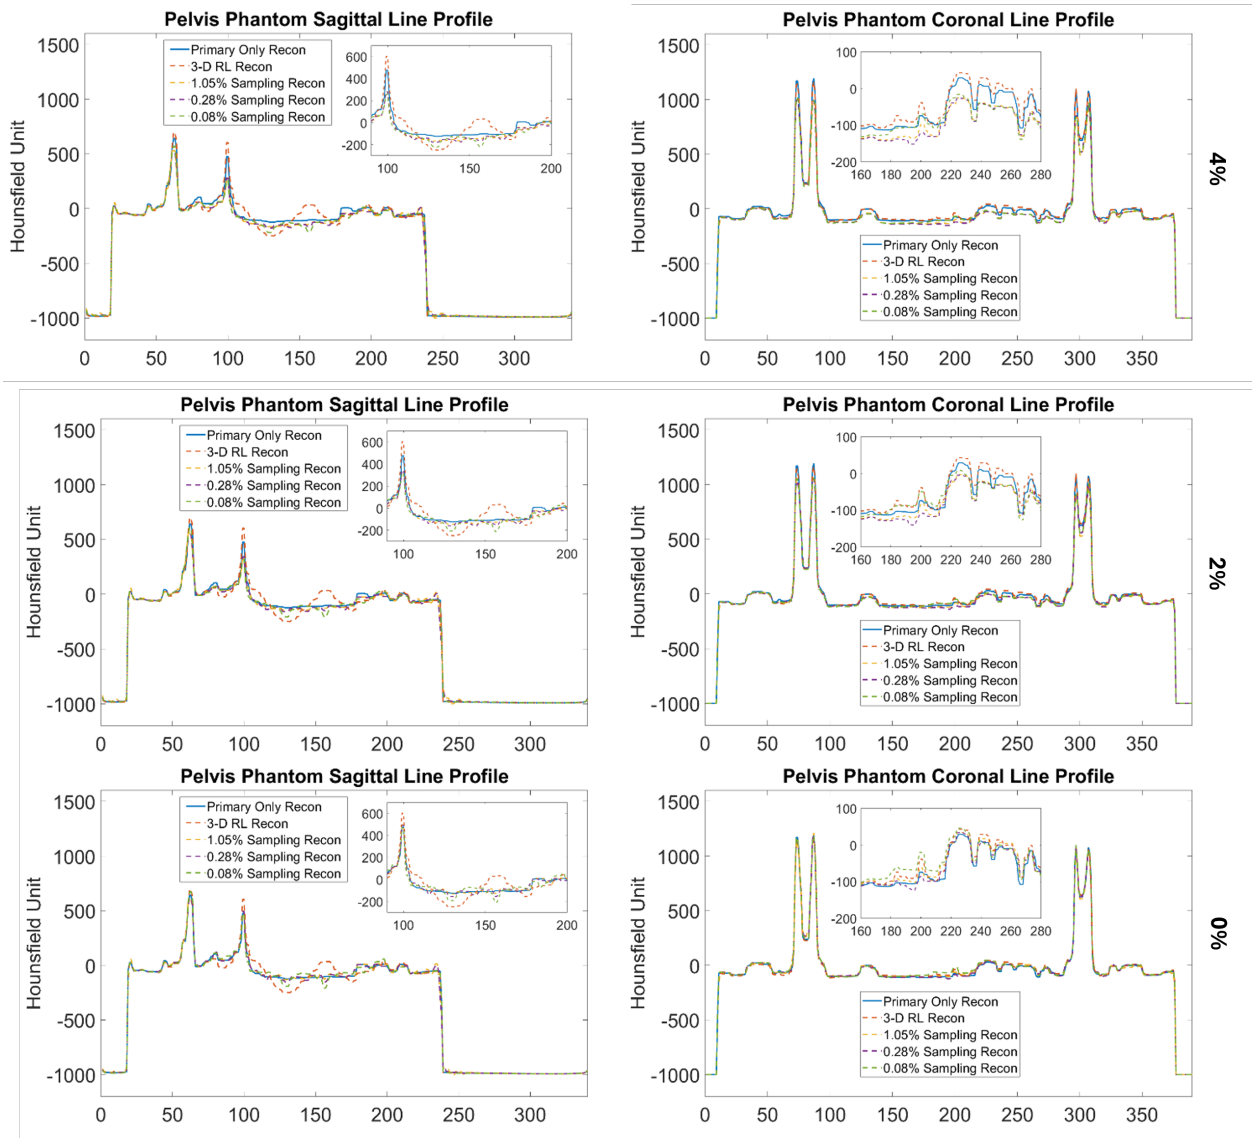

Figure 4.24: Line profiles taken from the pelvic phantom comparing the relative scatter recovery of different scatter contamination cases: 4%, 2%, and 0% scatter contamination.

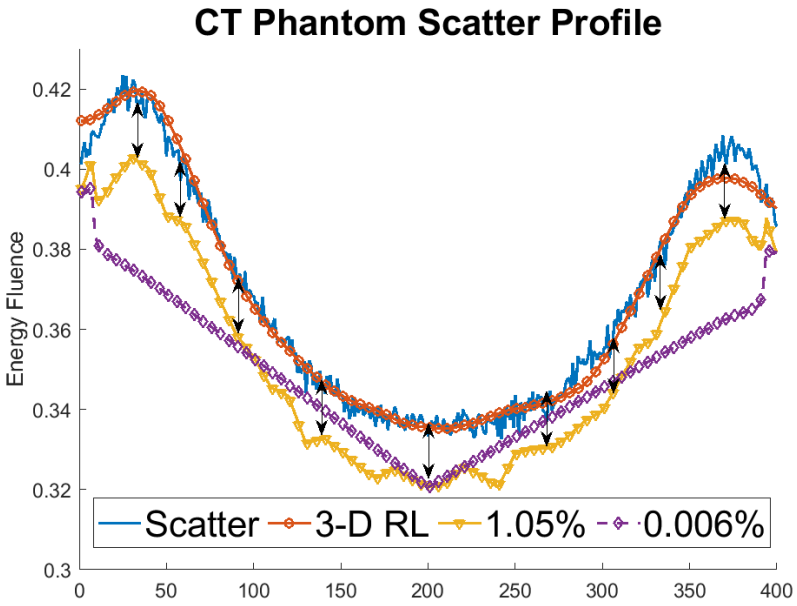

Figure 4.25: CT phantom comparison of MC ground-truth scatter, 3-D RL fitted scatter, and reconstructed scatter from the SPS method at SD of 1.05% and 0.006%. Black arrows depict the constant HU error gap from scatter contamination.

Chapter 5

Experimental Hardware Implementation and Early Validation of the SPS Grid

Following validation of the theoretical SPS grid via the computational Monte Carlo study in Chapter 4, an experimental study was conducted to further develop the SPS grid technique. This chapter presents the experimental apparatus, methodology, and the initial results and discussion of the experimental implementation of the SPS grid. The chapter is organized as follows:

- The benchtop CBCT system utilized for the experiments is introduced, including components, scanning geometry, system design, and associated firmware and software.
- Experimental hardware, including the prototype SPS grid, the beam stop array (BSA), and the Catphan phantom, is described.
- The hardware validation study design is outlined, using the BSA to provide a best-case scatter measurement reference.
- Quantitative image-quality metrics are introduced using the standard CT phantom modules.
- The imaging and preprocessing workflow for the experimental benchtop CT is detailed.

- Reconstruction methods and results for the CT phantom using the prototype SPS grid against the uncorrected signal are presented for proof-of-concept performance and discussed.
- Assumptions and limitations of the experimental study are discussed, and areas for improvement in future studies and SPS iterations are presented.

5.1 Benchtop CBCT System

The initial hardware feasibility study was conducted using a custom benchtop system modeling the scanning sequence of a third-generation CBCT, using a rotating stage as the moving reference frame. **Figures 5.1 and 5.2** show the initial conception and latest iteration of the system, respectively.

The core components of the benchtop system are seen in **Figure 5.3**, with the imaging system flow labeled. The central computer interfaces with the detector, stage controller, and Arduino microcontroller to synchronize the imaging system. At the computer, the user operates software dedicated to controlling the detector, source, and stage, with the user interfaces shown in **Figure 5.4**. Detector and source software for readout, acquisition, and parameter input were provided by iRay and Spellman, respectively. The Python interface for synchronizing beam output and detector acquisition via the Arduino microcontroller, and for stage rotation via a Sercomm terminal, was originally provided by Celestial Oncology, with the source code updated and modified to suit the new system configuration.

The Arduino microcontroller serves as the interface between the computer and the X-ray controller, while also synchronizing the detector acquisition trigger with the X-ray pulse window. For stage synchronization, the Arduino exchanges ready states with the Python interface to trigger exposure start based on the set angle for each projection.

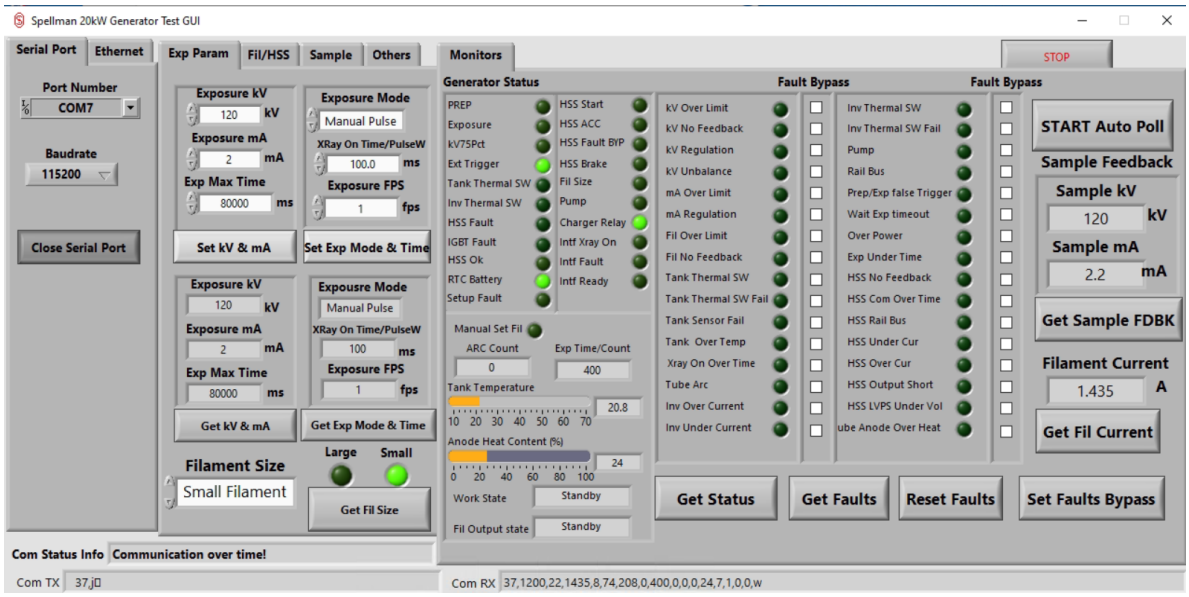

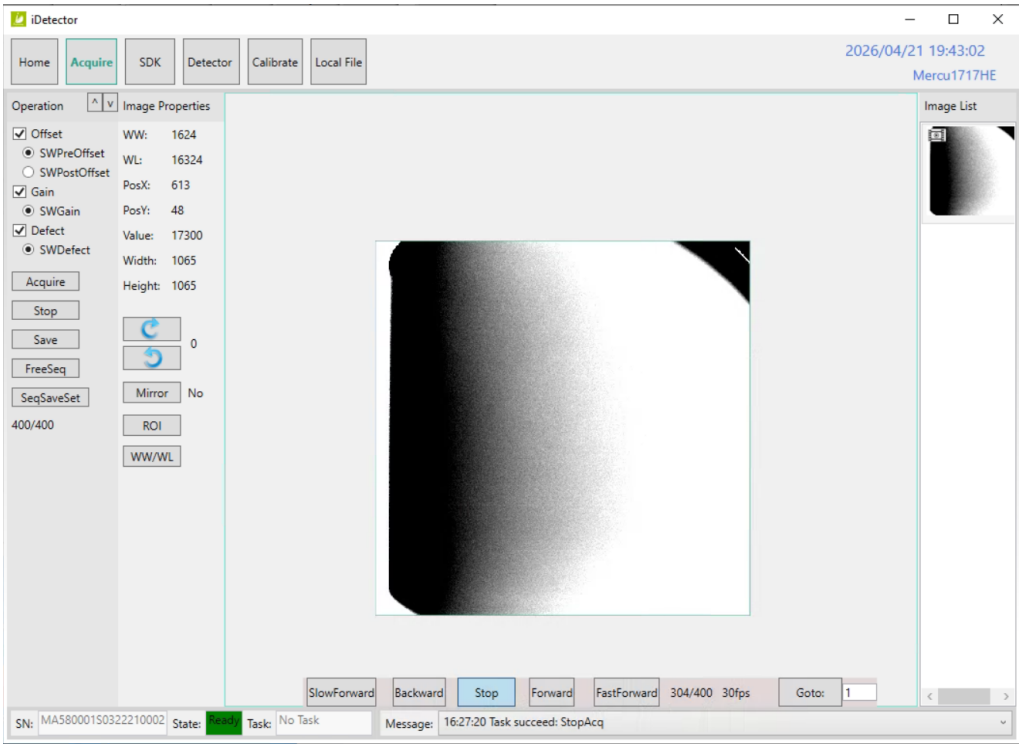

To summarize the capabilities of the CBCT system, the relevant operating values for each imaging component are given in **Table 5.1**.

Table 5.1. Operating specifications of the benchtop CBCT imaging components.

X-ray Source	Detector	Rotary Stage
kV Range: 40–125 kV	Scintillator: GOS/CsI	Travel Range: 360°
mA: 0.5–200 mA	Pixel Size: 100 μm	Accuracy: 0.020°
Peak output: 20 kW	Pixel Dimensions: 4260 × 4260	Precision: 0.0002°
Filtration: 0.7 mm Al	Resolution: 5 lp/mm	Min Motion: 0.001°
Beam Angle: 10°	Energy Range: 40 keV – 15 MeV	Speed: 20°/s
Min Pulse Time: 10 ms	Frame Rate: 10 fps (1×1), 40 fps (4×4)	Load Capacity: 100 N

Specific output energy configurations are detailed in the methodology of the hardware study. Detector binning and FOV are also provided when analyzing the hardware study results. For a step-and-shoot CBCT imaging sequence, the practical angular speed and framerate is 2 FPS for 1° spaced projections.

5.2 Experimental Hardware

5.2.1 SPS Grid Prototype

An initial SPS 3-D model was created in Fusion 360 (**Figure 5.5**) using a flexible Python script that allows variable parameters based on the source-to-detector distance, the grid dimensions, the number of smart pixel collimators, and the collimator diameter. The initial SPS prototypes were fabricated using third-party manufacturers, allowing for material selection, with initial fabrications using PVC (**Figure 5.6**). Manufacturing angular tolerances are within the margin of error for the theoretically calculated requirement of 0.04°, but for the hardware study, conventional tungsten carbide rods with larger tolerances are used as they can be cut to size onsite.

Additionally, to facilitate alignment of the grid to the beam source, a matching alignment grid was developed for the SPS grid. Based on the same generation Python script, the alignment grid

can be easily modeled and fabricated using personal 3-D printers, while more robust designs can similarly be manufactured by third-party companies.

5.2.2 Beam Stop Array

For the beam stop array, a product (**Figure 5.7**) was obtained from QRM [55] consisting of a regular array of small lead discs embedded in a low-attenuation supporting plate. The BSA is mounted between the X-ray source collimator and the imaged object such that each lead disc fully attenuates the primary signal along its line of sight, leaving only scattered photons to reach the shadowed detector pixels [29]. The intensity recorded at each shadowed pixel thus constitutes a direct sample of the scatter distribution at that location, and a complete two-dimensional scatter map for the projection is recovered by spatial cubic spline interpolation across the sparse sample lattice.

To obtain ground-truth scatter measurements at every projection angle, two complete scans of the object are acquired under matched geometry and X-ray technique: one with the BSA in place and one without (**Figure 5.8**). The first scan yields a direct sparse measurement of the scatter distribution at each angle, from which the full two-dimensional scatter map is recovered by spatial cubic spline interpolation. The second scan provides the corresponding total-signal projection, free of the primary loss imposed by the lead discs. Subtracting the interpolated scatter projection from the matched total-signal projection at each view yields a scatter-corrected primary projection across the full angular range. The doubled exposure required by this acquisition strategy is impractical for clinical imaging but acceptable for the present benchmark study, where it isolates the achievable image quality under direct scatter measurement from errors introduced by angular interpolation or model-based estimation.

Within the present hardware study, the BSA-corrected reconstruction serves as the best-case scatter reference against which the proposed SPS grid technique is benchmarked. Because the BSA samples the scatter distribution directly rather than inferring it from the primary signal, the resulting reconstruction approximates the image quality achievable under near-ideal scatter knowledge for

the benchtop geometry and provides an upper bound on the artifact reduction expected from a single-scan correction method.

5.2.3 CT Phantom

The Catphan 503 [56] (**Figure 5.9**) was selected as the primary phantom for the hardware validation study, as it bridges the gap between the simulation study and the imaging metric validation. The Catphan family is the industry standard for evaluating physical image quality in clinical CT and CBCT, and its use across recent scanner-comparison and mobile CBCT studies [57], [58] provides a well-characterized basis for benchmarking the SPS grid technique against published results. The metrics enabled by the Catphan modules align directly with the image degradations introduced by scatter, such as cupping/shading non-uniformity, HU bias, reduced contrast-to-noise ratio, and altered noise texture, outlined by the QA framework established by AAPM Task Group 66 for CT scanners used in radiation oncology [59].

The Catphan 503 contains three modules relevant to this evaluation: the CTP404 sensitometry module, with cylindrical inserts of seven materials (air, PMP, LDPE, polystyrene, acrylic, Delrin, and Teflon) spanning approximately −1000 HU to +1000 HU; the CTP528 module, with a 1–21 lp/cm aluminum bar gauge and a sub-pixel tungsten-carbide bead source for visual resolution evaluation and MTF measurement; and the CTP486 module, cast from a single uniform material within ±2% of water density for image noise, uniformity, and noise power spectrum (NPS) analyses. Notably, the 503 configuration omits the dedicated low-contrast module included in higher-tier Catphan models, so low-contrast performance is instead characterized through CNR measurements on the CTP404 inserts.

5.3 Hardware Validation Study Design

The hardware validation study implements the benchtop CBCT system, the SPS grid, the beam stop array, and the phantoms in an experimental methodology to quantitatively assess whether the SPS

grid technique recovers image quality lost to scatter under realistic acquisition conditions. Three full CBCT scans of each phantom are acquired, each under a distinct geometric imaging configuration corresponding to the different hardware arrangement in the beam path. The shared parameters across all three configurations are 105 cm source-to-detector distance (SDD), 55 cm SID, a 30 cm air gap between phantom and detector, and a 20 cm diameter rotation stage supporting the 20 cm diameter Catphan 503. X-ray technique and scan protocol are matched across configurations so that the only experimental variable is the hardware introduced into the beam path. The shared geometry and the three configurations are illustrated in **Figure 5.10**.

The first configuration is acquired with neither the BSA nor the SPS grid in the beam path. The resulting projections preserve the full scatter contamination present in the benchtop geometry and are reconstructed directly to provide the uncorrected control, establishing the worst-case lower bound against which any scatter-correction improvement is measured. The second configuration is acquired with the BSA mounted in front of the X-ray source between the collimator and the phantom, with the SPS grid removed from the detector. Combined with the uncorrected projections from the first configuration, this provides the matched total-signal and BSA-blocked projection pair required for the BSA correction algorithm, which subtracts the interpolated scatter image from the matched total-signal projection at each view to yield the BSA-referenced reconstruction. Because the BSA samples are measured directly at each view rather than inferred from the primary signal, this configuration provides a close approximation of ground-truth scatter and serves as the best-case reference. The third configuration is acquired with the SPS grid mounted at the detector face, with smart-pixel collimators oriented toward the focal spot per the alignment procedure described above, and the BSA removed from the source side. The resulting projections are reconstructed to produce the SPS-corrected reconstruction. The three reconstructions are compared on identically positioned phantom slices to allow per-voxel and per-ROI comparisons.

A comparative arm using a clinical anti-scatter grid (ASG) was considered during the design of this study, as ASGs are the dominant hardware-based scatter-rejection technique in clinical CBCT and would provide a useful benchmark against an existing standard-of-care alternative. However,

the manufacturing tolerances required to align an ASG to the source-to-detector geometry of the benchtop system, combined with the limited availability of focused grids matched to the 105 cm SDD and the cost of custom fabrication, made an ASG arm impractical within the scope of this study. The BSA reference arm partially addresses this gap by providing an upper-bound image-quality benchmark independent of any specific hardware-based rejection technique, and a direct ASG comparison is identified as future work in Chapter 6.

5.4 Quantitative Evaluation Metrics

Unlike the computationally defined phantoms of Chapter 4, in which material HU values are known exactly and a ground-truth primary-only reconstruction is directly accessible, the experimental study relies on a physical Catphan phantom whose insert compositions and nominal HU values are specified by the manufacturer. Ground-truth primary-only images are therefore unavailable, and reconstruction accuracy is instead assessed relative to published nominal values and through physically meaningful ratios between regions. Image quality is characterized by the same four metrics introduced in Section 4.4 — HU accuracy (here as HU linearity), spatial resolution via the MTF, image uniformity, and CNR — following the phantom evaluation methodology of Okazaki et al. [57]. The metric definitions established in Section 4.4 carry over and are not repeated here; this section presents only the modifications required by the physical Catphan phantom and the corresponding experimental measurements.

Because the experimental prototype SPS grid and benchtop setup introduce geometric distortions and miscalibration effects absent from the simulation study, reconstructed images are normalized prior to metric computation to decouple scatter-correction performance from residual hardware calibration errors. This normalization strategy is described alongside the experimental reconstruction pipeline in Section 5.5.

5.4.1 Catphan Phantom Module Overview

The Catphan phantom (The Phantom Laboratory, Salem, NY) is the standard quality assurance phantom for CBCT imaging systems and serves as the primary image quality benchmark throughout the experimental study. The relevant modules are CTP404 sensitometry, CTP528 high-resolution, and CTP486 image uniformity.

The CTP404 module (**Figure 5.11**) contains seven cylindrical material inserts as referenced in section 5.2.3 embedded in a water-equivalent epoxy background at known radial positions. The inserts span a wide range of nominal HU values from approximately −1000 HU (air) to approximately +990 HU (dense bone/Teflon), providing a well-characterized multi-material contrast target for HU accuracy and CNR evaluation.

The CTP528 module (**Figure 5.12**) contains a high-contrast bar pattern gauge with line pair frequencies from 1 to 21 lp/cm. The alternating high-density and low-density bars at progressively increasing spatial frequencies allow the MTF to be sampled across the relevant frequency range for clinical CBCT imaging.

The CTP486 module (**Figure 5.13**) consists of a homogeneous water-equivalent disk designed specifically for image noise and uniformity assessment. Five square ROIs, one central and four peripheral, are placed within the uniform disk to quantify the spatial HU variation induced by scatter-driven cupping artifacts.

5.4.2 HU Linearity

Because no ground-truth primary-only reconstruction is available, the per-ROI HU RMSE used in Section 4.4.4 is replaced by a HU-linearity analysis against the manufacturer's nominal values. HU linearity is assessed by placing circular ROIs within each CTP404 material insert and performing a linear regression of measured mean HU values against the corresponding nominal HU values specified by the phantom manufacturer. The coefficient of determination (R^2) of this regression quantifies the fidelity of the measured-to-nominal relationship; $R^2 = 1$ indicates perfect linear correspondence.

Deviations from linearity reflect residual scatter-induced HU error, beam-hardening effects, or normalization imperfections in the reconstruction pipeline. The linear regression is defined as:

$$HU_{measured} = a \cdot HU_{nominal} + b \tag{5.1}$$

where a is the fitted slope and b is the fitted intercept. The coefficient of determination is:

$$R^2 = 1 - \frac{\sum_i \left(HU_{meas,i} - \widehat{HU}_{meas,i} \right)^2}{\sum_i \left(HU_{meas,i} - \overline{HU}_{meas} \right)^2} \tag{5.2}$$

where the summation runs over all CTP404 material inserts i, $HU_{meas,i}$ is the mean pixel value within the circular ROI corresponding to insert i, $\widehat{HU}_{meas,i}$ is the regression-predicted value for insert i, and $\overline{HU}_{meas}$ is the mean of all measured insert HU values. A slope a close to unity and intercept b close to zero indicates both accurate scaling and offset, while R^2 approaching 1 indicates low scatter in the measured-to-nominal relationship across the full dynamic range of insert materials.

5.4.3 MTF

Spatial resolution is quantified from the CTP528 module. Unlike the bar-pattern contrast method used in the computational study (Section 4.4.5), the experimental MTF is derived by the edge-spread-function (ESF) method, which estimates the frequency response from a sharp physical edge rather than from discrete line-pair modulations, giving a continuous, densely sampled response that is less susceptible to reconstruction artifacts from setup errors. The ESF is sampled perpendicular to a high-contrast edge in the CTP528 module, and the line-spread function (LSF) is obtained by central-difference differentiation: The ESF is extracted by sampling a line of pixel values perpendicular to a high-contrast edge within the CTP528 module. The LSF is then obtained by differentiating the ESF with respect to position using a central-difference approximation:

$$LSF(x) = \frac{dESF(x)}{dx} \approx \frac{ESF(x + \Delta x) - ESF(x - \Delta x)}{2\Delta x} \tag{5.3}$$

where Δx is the pixel sampling interval in mm. To reduce spectral leakage introduced by the finite measurement length of the edge profile, the LSF is multiplied element-wise by a Hann window $w(x)$ prior to transformation:

$$\begin{aligned} LSF_w(x) &= LSF(x) \cdot w(x) \\ w(x) &= \frac{1}{2} \times \left(1 - cos\left(\frac{2\pi x}{N - 1}\right)\right) \end{aligned} \tag{5.4}$$

where N is the total number of samples in the LSF. The windowed LSF is then zero-padded to 1024 samples to interpolate the frequency spectrum more finely without adding information, and the MTF is computed as the magnitude of the fast Fourier transform (FFT) of the windowed, padded LSF, normalized so that $MTF(0) = 1$:

$$MTF(f) = \frac{|FFT\,(LSF_w(x))|}{|FFT\,(LSF_w(x))|_{f=0}} \tag{5.5}$$

where f is the spatial frequency in 1/mm, determined by the pixel sample spacing Δx and the padded FFT length. The MTF is reported at 10% and 50% modulation thresholds (MTF_{10} and MTF_{50}), corresponding to the limiting resolution and the half-power spatial frequency, respectively. Because the experimental prototype introduces geometric reconstruction errors that may broaden the ESF independently of scatter, MTF comparisons between SPS-corrected and uncorrected reconstructions are interpreted as relative improvements in scatter-limited resolution recovery rather than absolute system resolution characterizations.

5.4.4 Image Uniformity

Image uniformity is quantified with the UI defined in Section 4.4.6 (Eq. 4.15). For the experimental study, the UI is computed from five ROIs in the homogeneous CTP486 module — one central and four peripheral (top, bottom, left, right). A central line profile across the module is additionally reported to visualize the cupping-artifact shape and the spatial extent of scatter-correction recovery.

5.4.5 CNR

CNR follows the combined-noise definition given in Section 4.4.3 (Eq. 4.12). For the experimental study, CNR is computed for each CTP404 insert material, with the background ROI placed at the same radial distance from the phantom center as the corresponding insert so that any residual cupping gradient affects signal and background equally and the metric isolates contrast recovery from uniformity recovery. CNR is reported separately per insert to capture the material-dependent variation in contrast recovery between high-density (e.g., Teflon, bone) and low-density (e.g., LDPE, air) inserts.

5.5 Experimental Imaging Workflow

5.5.1 Projection Acquisition

Phantom projections are acquired by taking 360 projections in a step-and-shoot sequence with stage rotation and exposure. The phantom was placed at 30 cm distance from the detector to the rotation axis, with a 75 cm distance to the X-ray source. The radiographic exposure technique utilized was 120 kV, 3 mA, and 100 ms pulses with a large focal spot of 0.6 mm. The detector was calibrated for 4×4 binning with 0.4×0.4 mm² pixel sizes for a total of 1065×1065 pixels. Darkfield projections and blankfield projections utilized the same exposure technique.

The initial scatter estimation was obtained using an SPS grid projection (**Figure 5.14**) and performing an SPS smart pixel boundary interpolation to recover the missing scatter signal from each of the smart pixels, which contain only the primary signal, by subtracting the interpolated uncorrected signal (**Figure 5.15**).

The full-field scatter distribution was then estimated by interpolating between each smart-pixel sampling point, with averaged scatter values padding the boundary pixels to interpolate external SPS values (**Figure 5.16**).

Attempts were made to perform SPS scatter correction on the same projections used to obtain the interpolated scatter (**Figure 5.14**), but streaking artifacts from interpolation errors caused by nonideal measurements and geometric misalignment in the SPS smart-pixel collimators rendered analysis unfeasible (**Figure 5.17**). To emulate ideal conditions in which interpolation errors are minimized with an optimal SPS grid setup, the uncorrected Catphan projections without the SPS grid were used as the scatter-correction target.

5.5.2 Projection Processing

The detector center and beam center from the source are misaligned due to various complications during setup, and as a result, the field size at the detector is truncated. Pixels on both sides of the detector were truncated in the projection to remove unexposed pixels, resulting in a 2-D pixel resolution of 965×1065.

Blankfield and darkfield projections were used for projection correction and line integral calculations, respectively. Projections were time-averaged into one 2-D reference matrix for subtraction and line integral calculation:

$$\widetilde{B}(x,y) = \frac{1}{N_B}\sum_k B_k(x,y), \quad \widetilde{D}(x,y) = \frac{1}{N_D}\sum_k D_k(x,y), \tag{5.6}$$

where B and D are the time-averaged blankfield and darkfield projections, respectively. Darkfield scans were subtracted from both the blankfield and object projections.

The experimental benchtop CT contains multiple setup deviations and exposure inconsistencies that must be accounted for in the projections. The exposure between each projection fluctuates, with some angular projections detecting more overall fluence at the detector than others. To compensate for the exposure variation, a shared air ROI was selected in all angular projections and used to normalize the overall fluence with a scaling factor. The scaling factor was derived using the ratio of the mean air value in the same ROI location in both the blankfield and projection after darkfield correction:

$$s_i = \frac{\overline{P^{\text{air}}}}{\overline{B^{\text{air}}}} \tag{5.7}$$

where s_i is the scaling factor and P_{air} is the mean value of the projection ROI over the mean value of the blank projection ROI B_{air}.

The next corrected error was mechanical misalignment between the rotation stage center and the actual radiation field center. This misalignment can normally be corrected in the reconstruction input parameters for a conventional CBCT rotation frame with a stationary object, but the moving rotation frame in the experimental benchtop system is swapped with the object. As such, the object's misalignment on the rotation stage results in a geocentric orbit in the imaging reference frame, which cannot be corrected using geometrical parameters. For proof-of-concept and initial analysis, a preprocessing step for the CT phantom projections was implemented to both center and align the angular projections for CBCT reconstruction. As the CT phantom is cylindrically symmetric, the centroid can be determined via cross-correlation and used to align the phantom in each projection. The cross-correlation describes the degree of overlap between the phantom's projected intensity profile and its flipped mirror image as a function of relative shift along the detector axis. Subpixel shifts were handled by upsampling to apply the shift and then downsampling to the original resolution.

5.6 Reconstruction Results and Discussion

The original intent of this experimental study was to evaluate the performance of the SPS grid scatter correction technique and benchmark the scatter estimation accuracy with the BSA method described in Section 5.3. However, the BSA reconstruction is omitted from the current analysis: difficulties with the beam stop array setup and the inability to obtain robust scatter-to-primary-ratio measurements under the present geometry prevented its realization, as detailed among the experimental limitations in Section 6.5. Instead, the reconstruction methods compared for the experimental study are FDK and TV-regularized iterative reconstruction (IR TV), with and without SPS grid scatter correction.

Instead of the intended BSA reconstruction, a clinical CT scan was used as the benchmark for ideal scatter correction which should be considered when comparing between the results. It must be emphasized that, owing to the smart-pixel interpolation streaking described in Section 5.5.1, the SPS grid scatter correction was applied to the uncorrected, grid-free projections rather than to the SPS grid projections themselves. This substitution circumvents the dead-zone interpolation error rather than resolving it, so although the results validate the grid's scatter-measurement concept, they do not yet constitute the single-scan correction that defines the method, and its implications are examined in Section 6.3.

Furthermore, the computational study's proposed alternating minimization was not implemented, as the geometrical distortions from imperfect alignment and mechanical deviations are propagated and amplified at each alternating iteration, resulting in circular artifacts in the final reconstruction. As an additional consideration, the prototype SPS grid was devised and used with a design focused solely on testing implementation, rather than on specific sampling densities or optimal dimensions. The consequences of the prototype SPS grid are discussed in Chapter 6.

5.6.1 HU Accuracy

To evaluate the HU accuracy of the reconstruction techniques, the line integral must be converted into HU values, as the clinical CT system outputs only HU values. The HU conversion was calibrated separately for the FDK, IR, and clinical reconstructions using an air and water ROI selected in module CTP486 for the corresponding scatter-corrected CT and clinical control slice (**Figure 5.18**).

For the HU linearity comparisons using module CTP404 (**Figure 5.19**), ROIs were selected corresponding to known materials, and **Table 5.2** describes the manufacturer's nominal values compared with the mean measured ROI values across the various reconstruction methods. The corresponding regression lines for each reconstruction were plotted and fitted using the mean HU values of the material insert ROIs (**Figure 5.20**).

Table 5.2. Nominal CTP404 insert HU values compared with the mean measured ROI values for each reconstruction method.

Material Insert	Nominal HU	Uncorrected FDK	SPS Corrected FDK	Uncorrected IR TV	SPS Corrected IR TV	Clinical CT
Air	−1000	−561.53	−705.83	−703.07	−900	−977.88
PMP	−200	−162.83	−140.59	−204.64	−262.51	−246.93
LDPE	−100	−119.11	−75.65	−156.49	−195.86	−162.82
Polystyrene	−35	−100.45	−45.51	−139.56	−169.5	−114.33
Acrylic	120	−46.35	41.6	−82.32	−89.24	24.64
Delrin	340	79.61	206.95	82.94	77.36	224.13
Teflon	990	337.04	591.91	397.62	506.04	768.52
Water	0	−18.05	62.99	−30.59	−45.72	3.23
R^2	1	0.451	0.651	0.550	0.699	0.877

Compared with the uncorrected reconstructions, the SPS-corrected reconstructions show partial scatter recovery in HU linearity, with the IR TV results showing superior recovery compared with the FDK results. The linear regression error for the IR TV results was reduced from 37.3% to 20.3% when SPS correction was applied.

5.6.2 Spatial Resolution

Using the CTP528 module and the line pairs present, the spatial resolution of each dataset and reconstruction method was measured via the modulation transfer function. The MTF was derived from the ESF of a line ROI (**Figure 5.21**) and describes the fidelity with which the reconstruction retains object contrast at each spatial frequency, decreasing from unity at zero frequency to zero at the system's resolution limit; the frequency at 50% modulation (MTF_{50}) was used as the summary metric of spatial resolution (**Figure 5.22**). For further data, MTF_{10} of **Figure 5.22** was as follows, in order of appearance: MTF10 = 1.344 1/mm, MTF10 = 1.276 1/mm, MTF10 = 0.595 1/mm, MTF10 = 0.696 1/mm, and MTF10 = 0.688 1/mm. Spatial resolution improved minimally with SPS correction, with the largest differences arising from the reconstruction technique. Due to the geometric distortions remaining in the reconstruction and upsampling in the clinical CT reconstruction, the spatial resolution cannot be evaluated confidently.

5.6.3 Image Uniformity

The CTP486 module (**Figure 5.23**) was analyzed to determine the cupping artifact recovery of the SPS scatter correction, with quantitative results shown in **Table 5.3**.

Table 5.3. Uniformity index and maximum HU deviation for each reconstruction method.

		SPS		SPS	
	Uncorrected	Corrected	Uncorrected	Corrected IR	Clinical
Metric	FDK	FDK	IR TV	TV	CT
UI	16.932%	4.208%	15.215%	3.356%	0.236%
Max Δ HU	135.5	39.5	147.3	33.3	2.4

The SPS corrected reconstructions showed an improvement of approximately 13% in uniformity, but remain an order of magnitude away from the clinical CT uniformity.

5.6.4 Contrast Resolution

Contrast resolution was benchmarked using the CTP404 module and the various material inserts (**Figure 5.24**). For visualization in the bar plot, LDPE, polystyrene, and Delrin were compared among the various reconstruction techniques. **Table 5.4** presents the full CNR results for all analyzed ROIs.

Table 5.4. CNR results for all CTP404 material insert ROIs across reconstruction methods.

		SPS		SPS	
Material	Uncorrected	Corrected	Uncorrected	Corrected IR	Clinical
Insert	FDK	FDK	IR TV	TV	CT
Air	18.787	20.778	22.841	28.926	127.222
PMP	3.775	4.646	3.972	6.646	29.752
LDPE	2.598	3.229	2.838	5.601	18.901

Material Insert	Uncorrected FDK	SPS Corrected FDK	Uncorrected IR TV	SPS Corrected IR TV	Clinical CT
Polystyrene	1.941	2.422	2.424	4.89	15.151
Acrylic	0.198	0.163	0.303	0.847	2.806
Delrin	4.497	4.704	5.589	7.914	24.56
Teflon	7.779	7.911	8.395	8.626	47.03
Water	1.251	0.811	2.278	3.582	0.049

Overall, the SPS-corrected IR TV reconstruction performed best at improving CNR across all material inserts. However, compared to clinical CT performance, the SPS method remains far from the desired results, being a factor of 3–5 off in CNR.

5.6.5 Discussion

Figure 5.25 presents the preliminary results of a prospective experimental study comparing clinical CT reconstructions with a prototype SPS grid paired with an ad hoc FDK or iterative reconstruction (IR) algorithm utilizing total variation (TV) regularization. Qualitatively, the SPS scatter correction mitigates cupping artifacts by an absolute margin of approximately 100 HU, though uniformity index measurements indicate some residual cupping remains. Quantitatively, the current SPS implementation enhances HU accuracy, spatial resolution, and CNR. In optimal scenarios, such as with the LDPE inserts, the CNR is doubled compared to uncorrected CBCT. Although the overall CNR performance still falls short of diagnostic CT benchmarks, the SPS reconstruction results strongly align with prior theoretical modeling.

By including both FDK and IR with TV reconstruction techniques, the image quality improvements from SPS scatter correction can be compared across algorithms. Across all imaging metrics, SPS scatter-corrected IR with TV performed significantly better than the other reconstruction methods, suggesting that the choice of reconstruction algorithm improves the SPS technique's

scatter-correction recovery. This suggests that a more finely tuned reconstruction algorithm, such as the alternating minimization algorithm from the computational study, can improve the performance of SPS scatter correction. Due to residual geometric alignment imperfections and miscalibrations during reconstruction, the image quality improvements from the experimental SPS technique have not been fully realized.

5.7 Figures

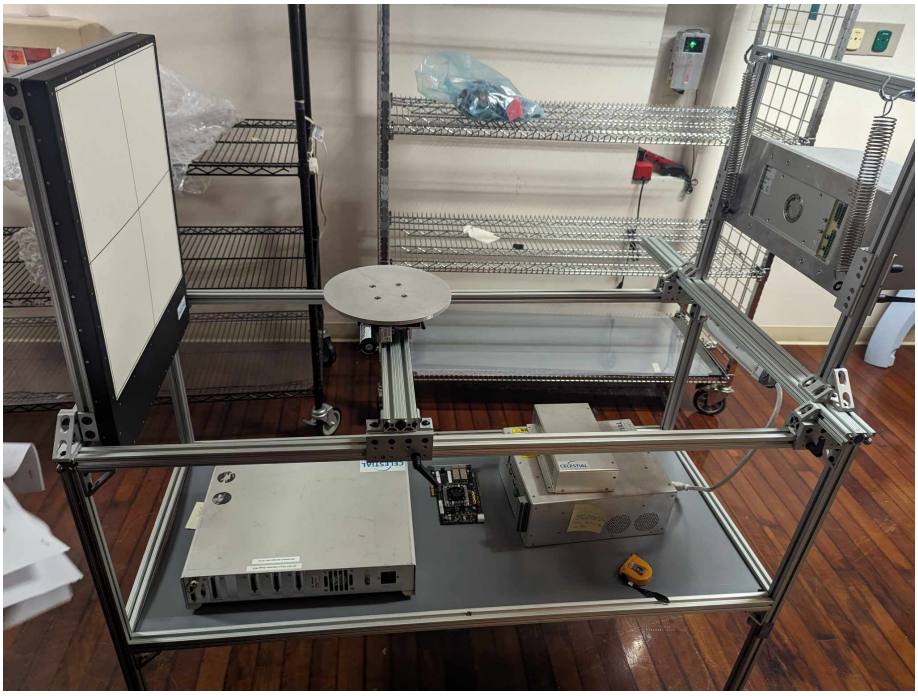

Figure 5.1: Initial conception of the benchtop CBCT system, inherited detector, deprecated X-ray source-controller and rotary stage, stage controller, and arduino microcontroller.

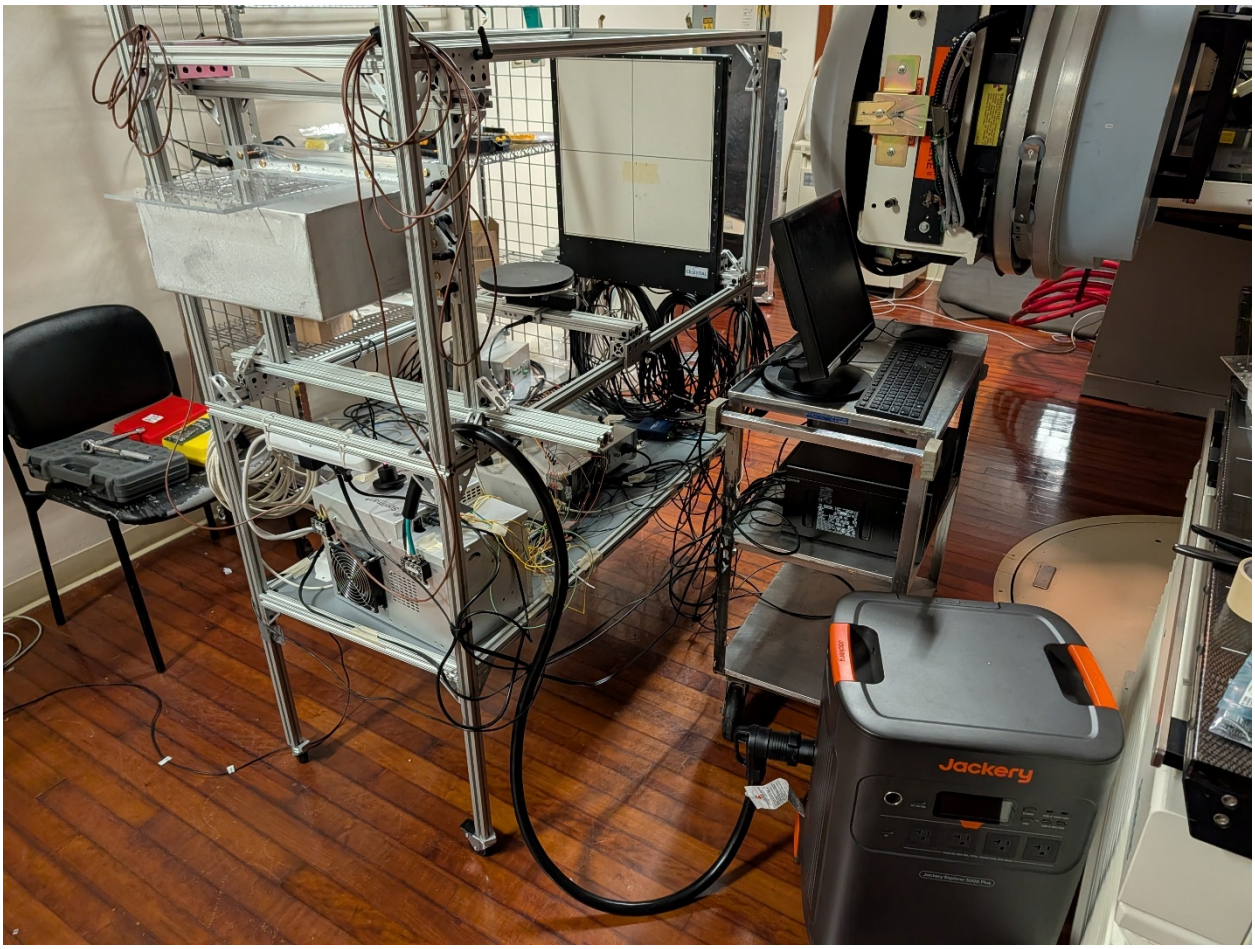

Figure 5.2: Latest iteration of the benchtop CBCT system showing the redesigned benchtop cage, portable high power supply, new X-ray source-controller, interfacing computer, inherited detector, and new rotary stage.

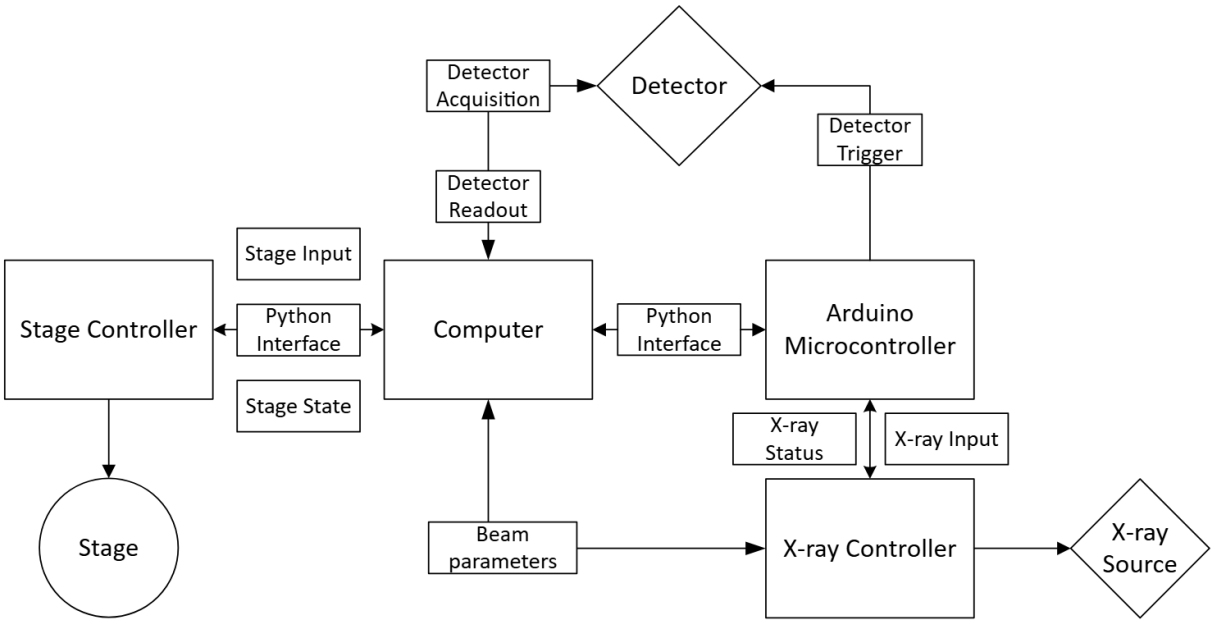

Figure 5.3: Core components of the benchtop system with the imaging system flow labeled. The computer serves as the master containing all three GUIs required to operate the stage, source, and detector with synchronization performed by the arduino. The computer also receives the measured projection data and listens to the status of each component.

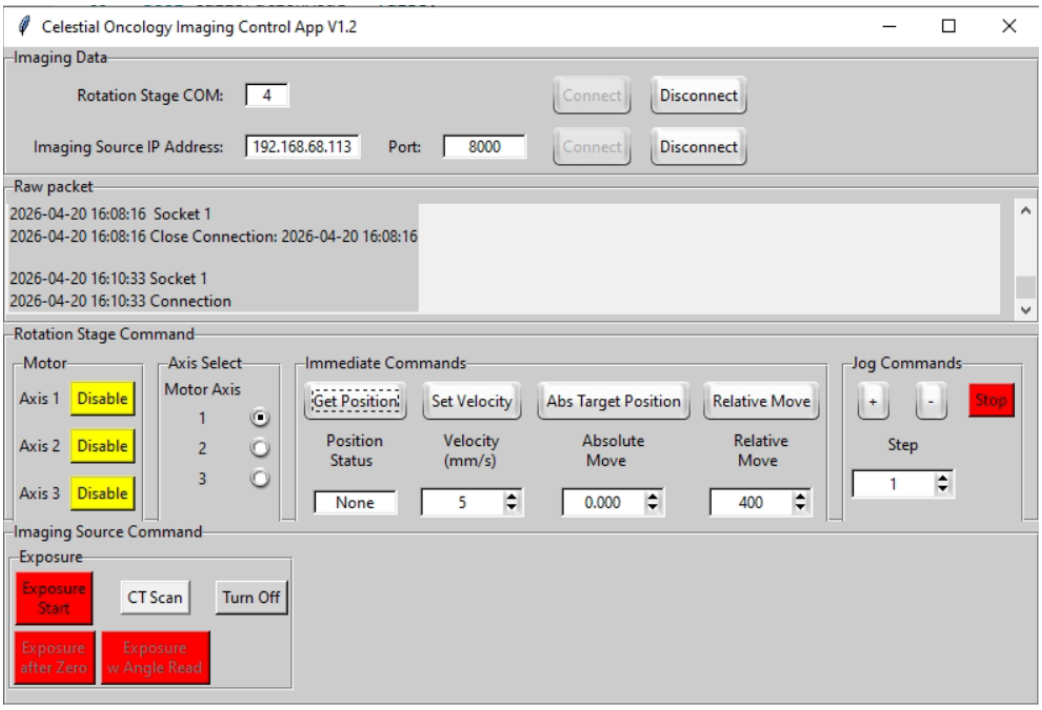

Figure 5.4: User interfaces for the X-ray source configuration, detector, and stage-exposure control software, respectively from top to bottom.

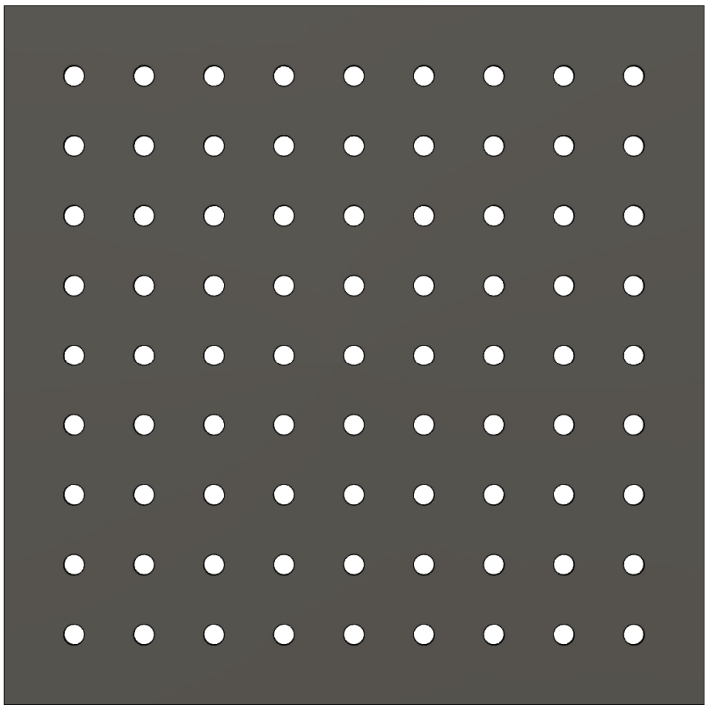

Figure 5.5: Initial SPS 3-D model created in Fusion 360 with parameterized Python generation script.

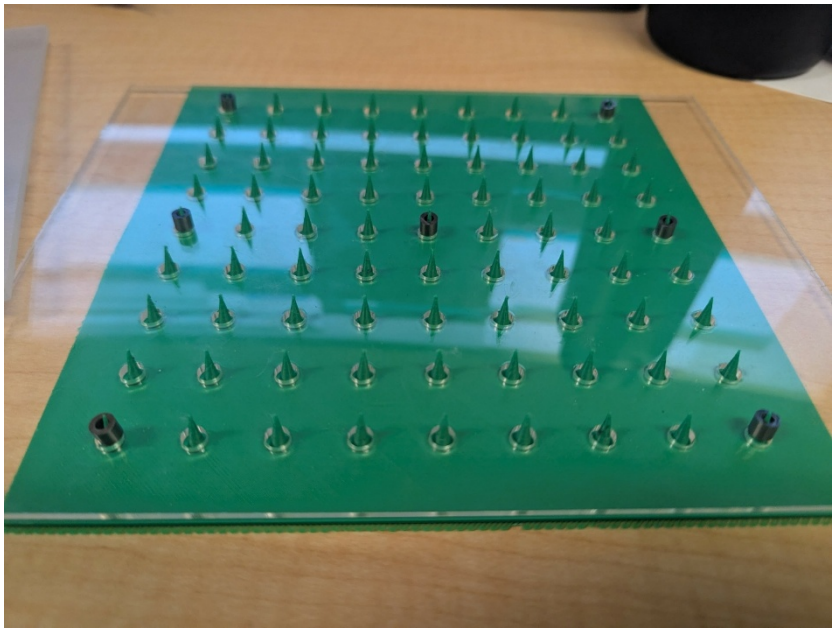

Figure 5.6: Initial PVC-fabricated SPS grid prototype along with alignment grid prototype.

Figure 5.7: Beam stop array consisting of a regular array of small lead discs embedded in a low-attenuation supporting plate.

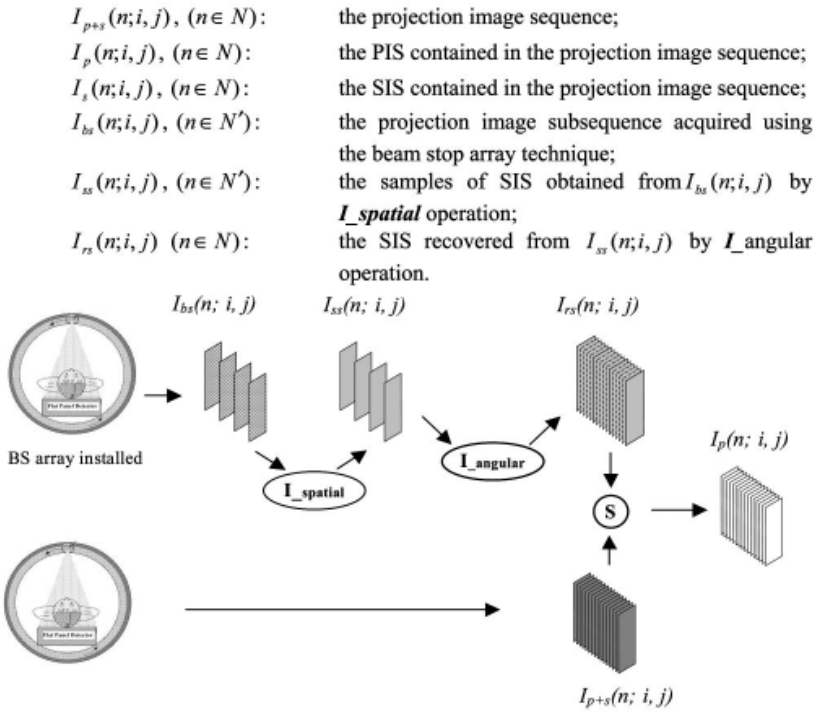

Figure 5.8: Two-scan BSA acquisition strategy: a BSA scan yielding sparse scatter samples and a matched open-field scan providing the total signal, combined to produce scatter-corrected projections.

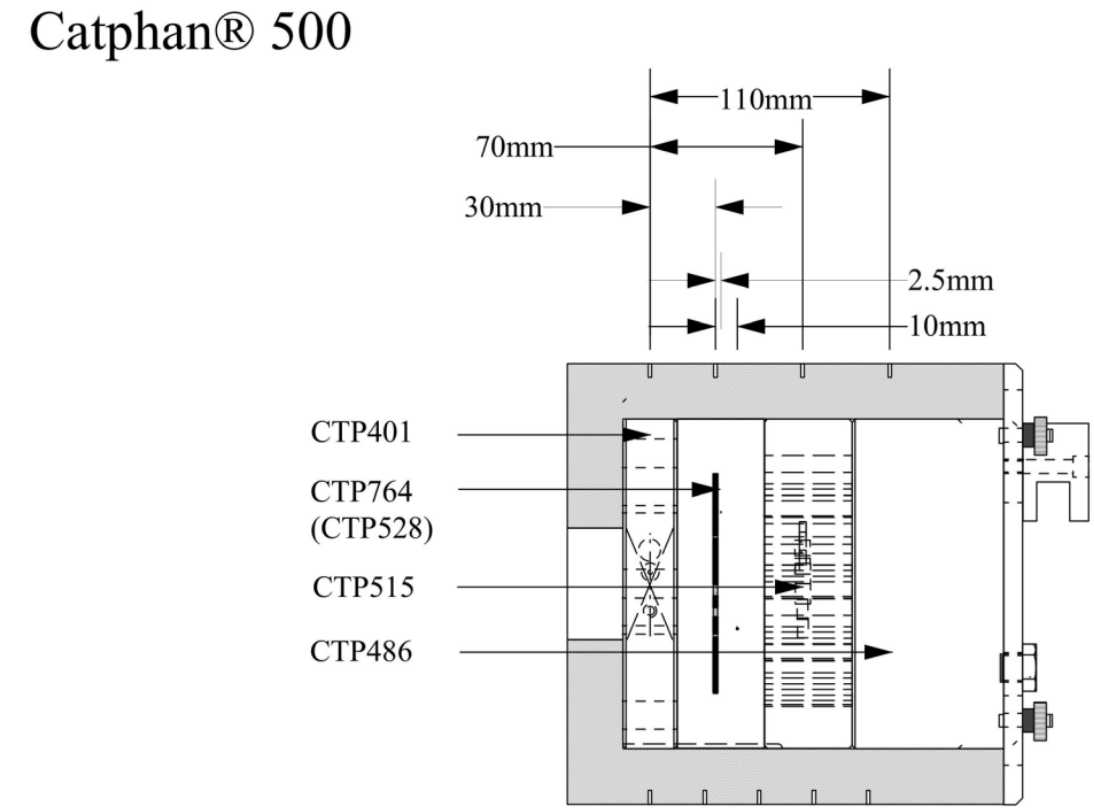

Figure 5.9: Catphan 503 phantom used as the primary phantom for hardware validation.

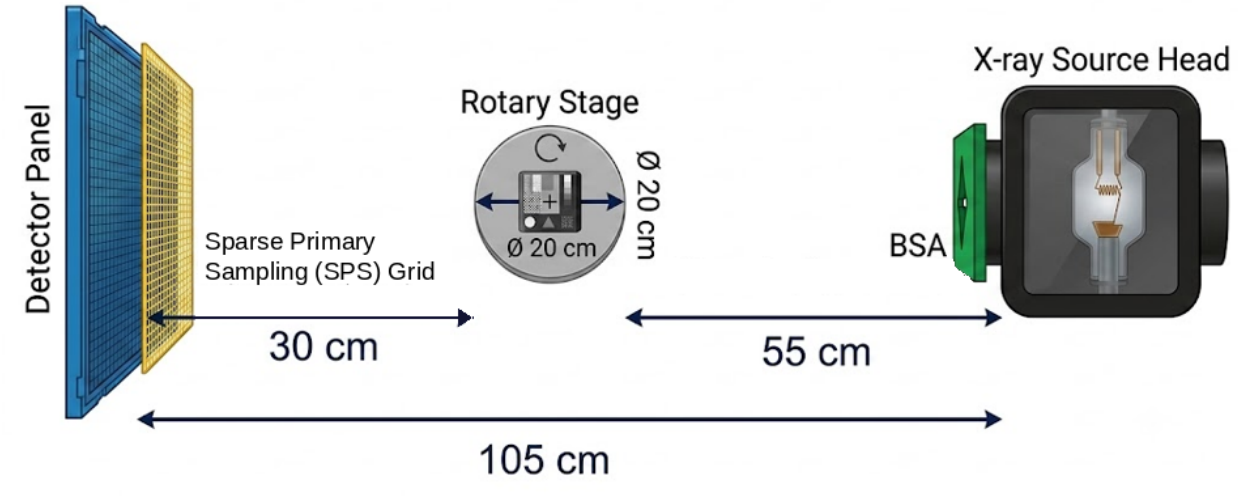

Figure 5.10: Shared experimental geometry for the three acquisition configurations: uncorrected control (no BSA or SPS grid), BSA-referenced configuration, and SPS grid configuration. BSA and SPS grid omitted as required.

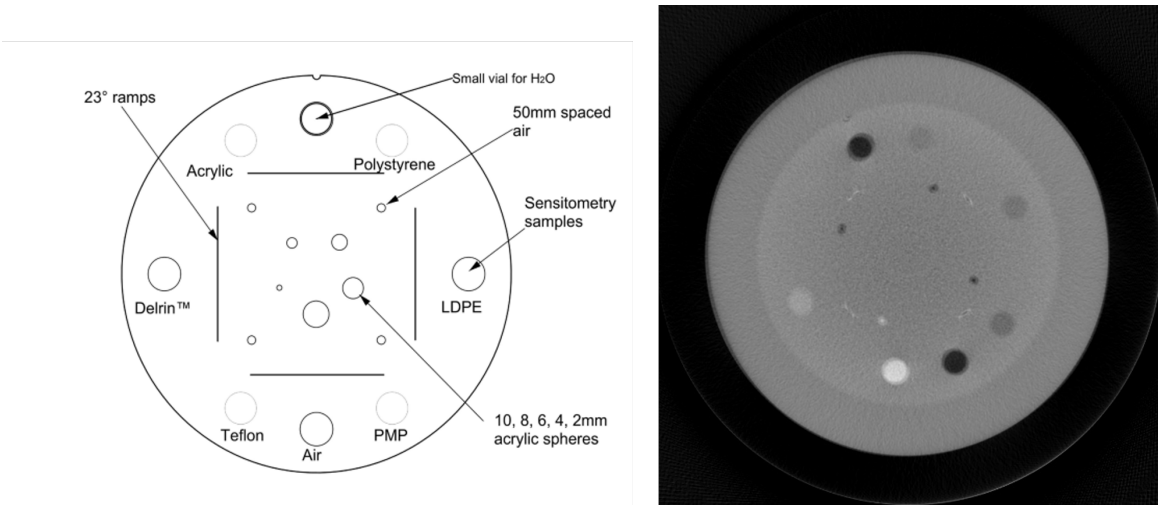

Figure 5.11: CTP404 sensitometry module with material inserts with the left showing the module as depicted by the Catphan manual and the right showing the reconstructed slice.

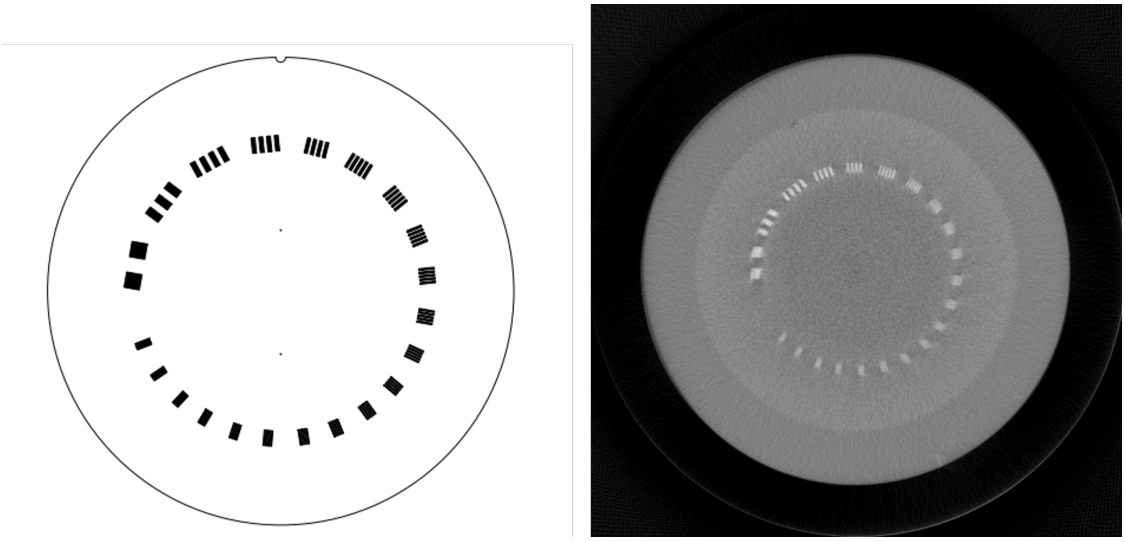

Figure 5.12: CTP528 high-resolution module with the left showing the module as depicted by the Catphan manual and the right showing the reconstructed slice.

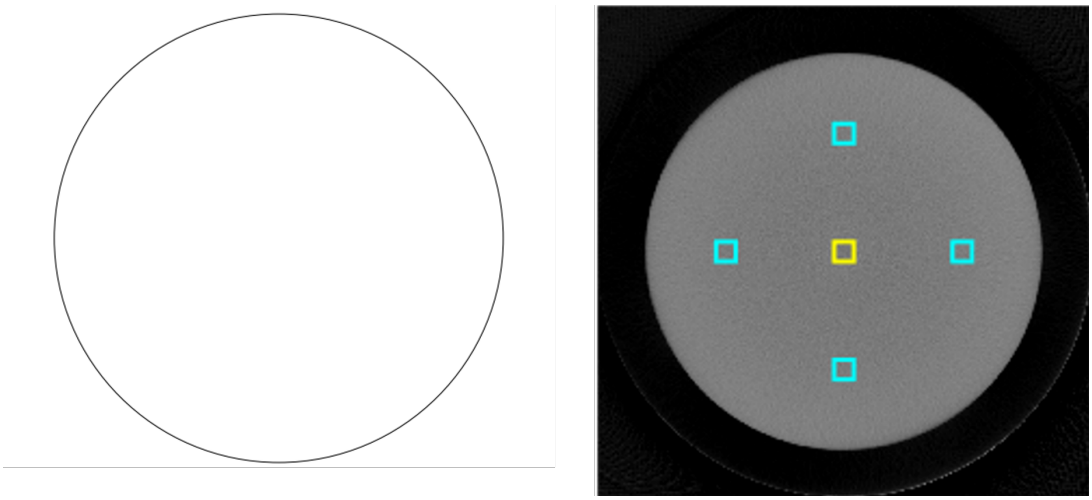

Figure 5.13: CTP486 image uniformity module with the left showing the module as depicted by the Catphan manual and the right showing the reconstructed slice.

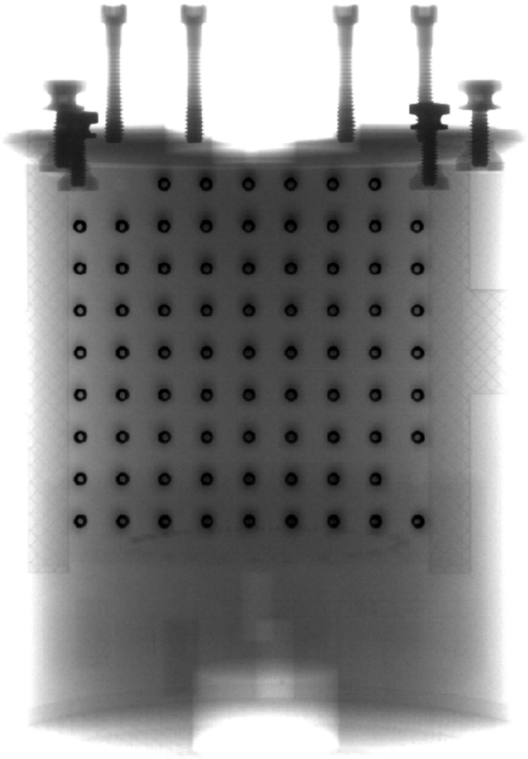

Figure 5.14: SPS grid projection used for the initial scatter estimation. Note the shadows of the 3-D printed in the projection.

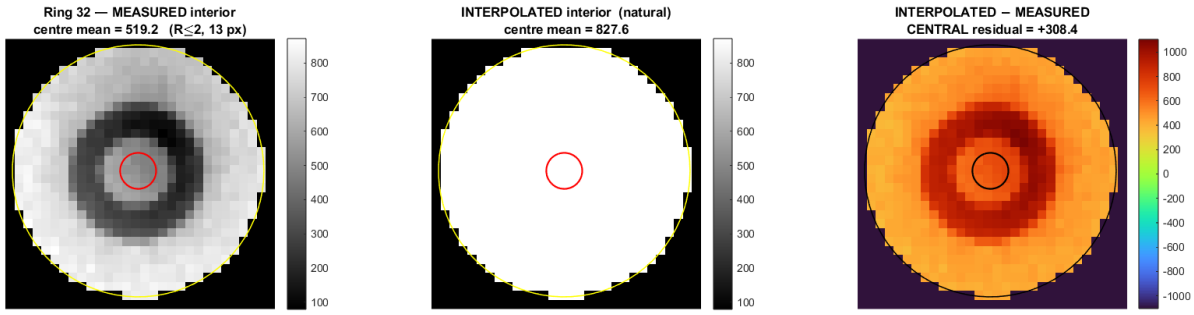

Figure 5.15: SPS smart pixel boundary interpolation recovering the missing scatter signal at each smart pixel. Shown here is the evaluation of the scatter at smart pixel 32 of the grid array.

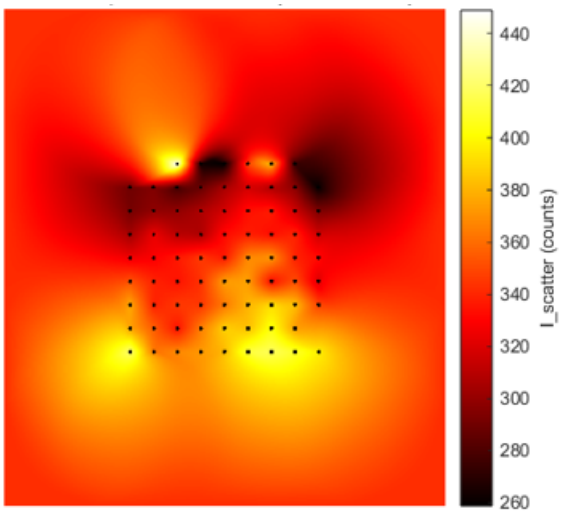

Figure 5.16: The full-field scatter distribution estimated by interpolating between smart pixel sampling points with extrapolation using averaged scatter as the value for exterior points.

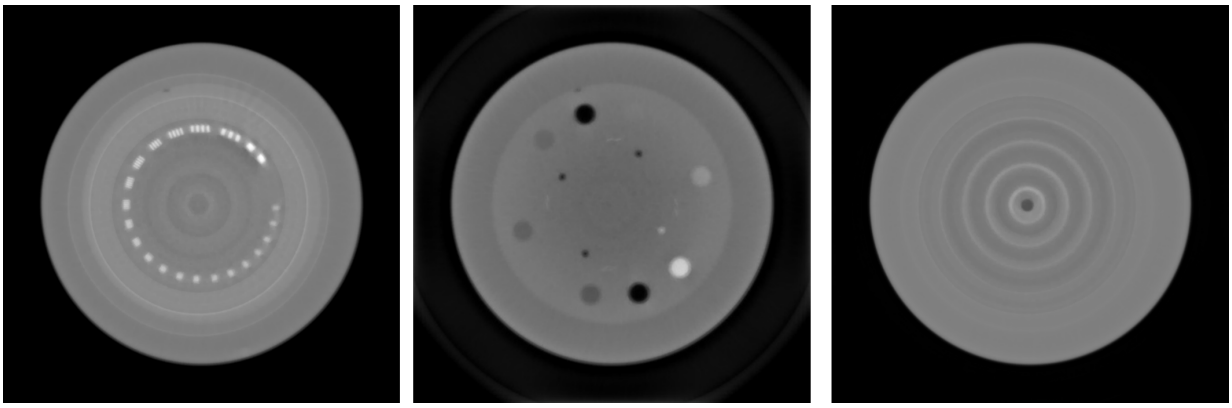

Figure 5.17: Streaking artifacts from interpolation errors caused by nonideal measurements and geometric misalignment in the SPS smart pixel collimators.

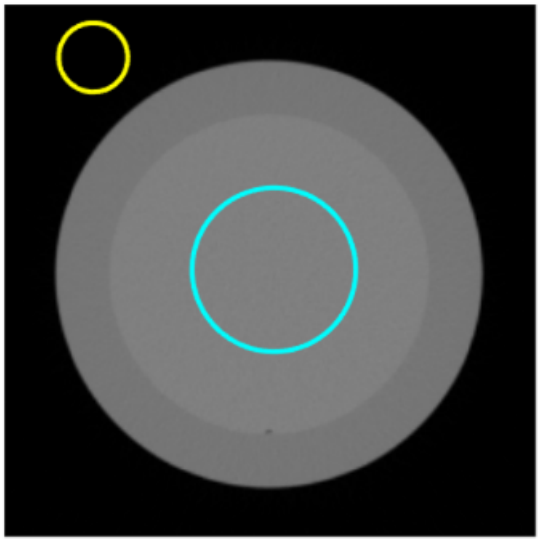

Figure 5.18: Air (yellow) and water (teal) ROIs in module CTP486 used for HU calibration of each reconstruction.

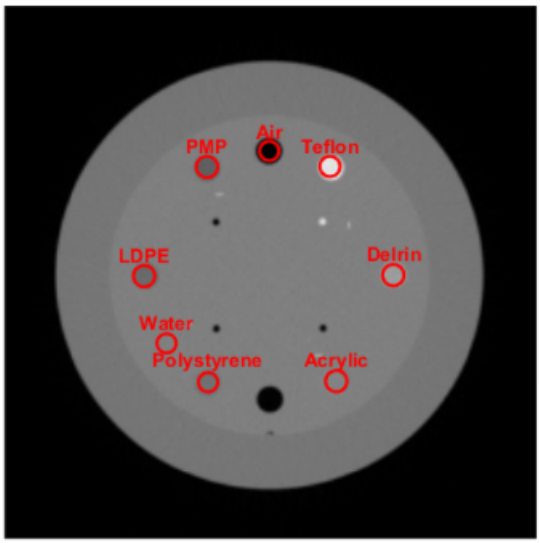

Figure 5.19: Reconstructed slice used for HU accuracy evaluation various material ROIs labeled.

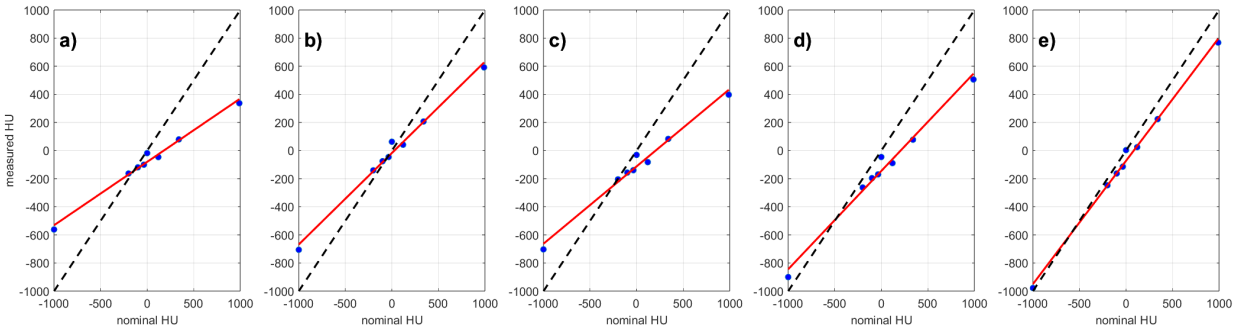

Figure 5.20: HU linearity plots for (a) uncorrected FDK, (b) SPS corrected FDK, (c) uncorrected IR TV, (d) SPS corrected IR TV, and (e) clinical CT control comparison.

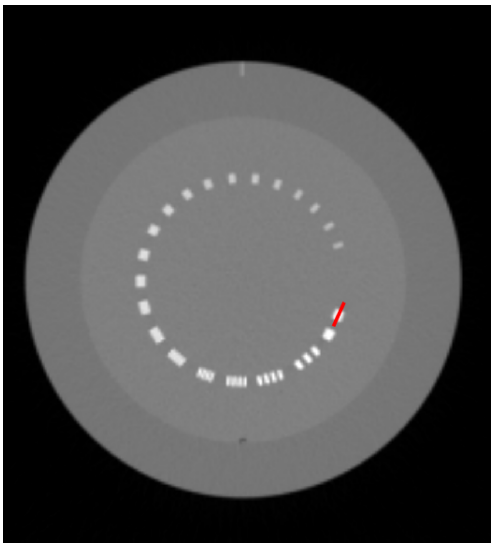

Figure 5.21: CTP528 module and line ROI used for ESF-based MTF measurement.

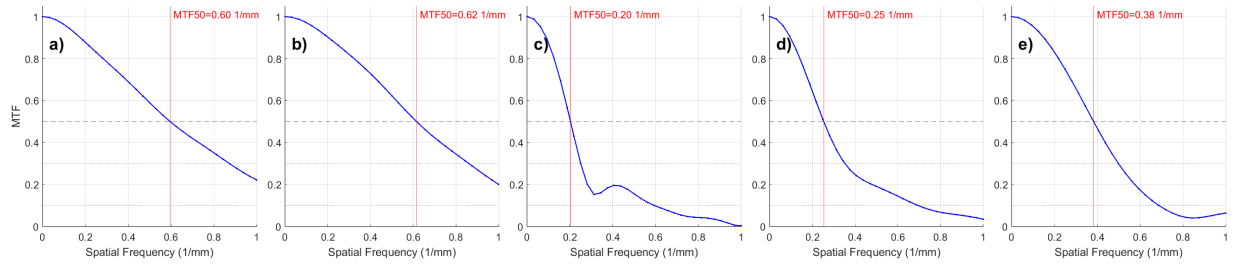

Figure 5.22: MTF comparison across reconstruction methods with MTF_{50} as the summary metric for (a) uncorrected FDK, (b) SPS corrected FDK, (c) uncorrected IR TV, (d) SPS corrected IR TV, and (e) clinical CT control comparison.

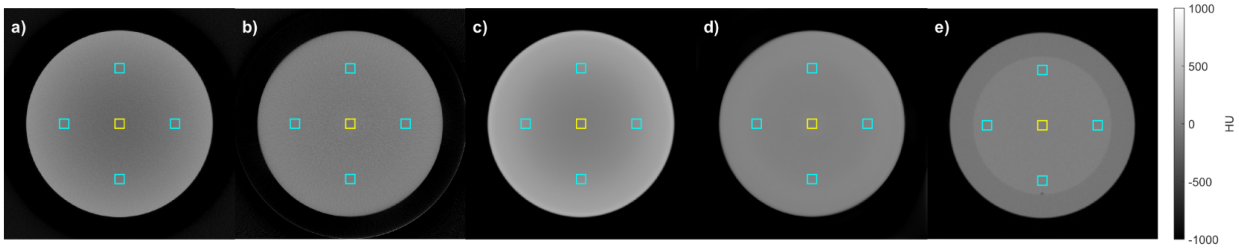

Figure 5.23: CTP486 uniformity analysis and central line profiles showing cupping artifact recovery for (a) uncorrected FDK, (b) SPS corrected FDK, (c) uncorrected IR TV, (d) SPS corrected IR TV, and (e) clinical CT control comparison.

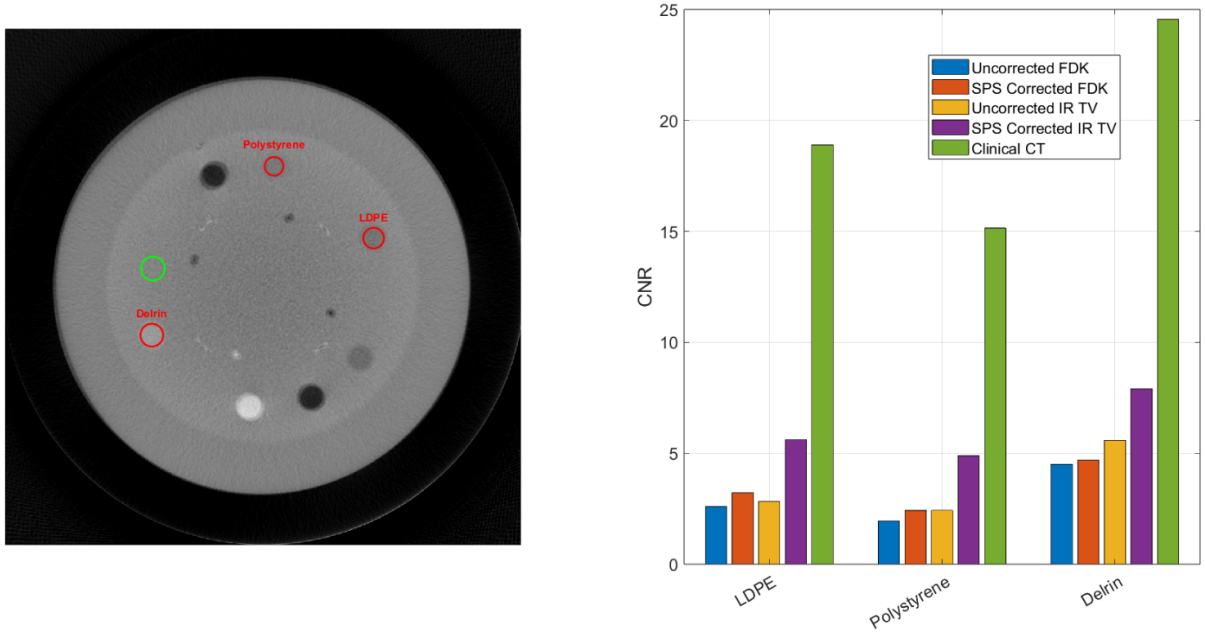

Figure 5.24: CTP404 module with corresponding material inserts ROI's mean values visualized in the bar plot with the various reconstruction techniques CNR compared

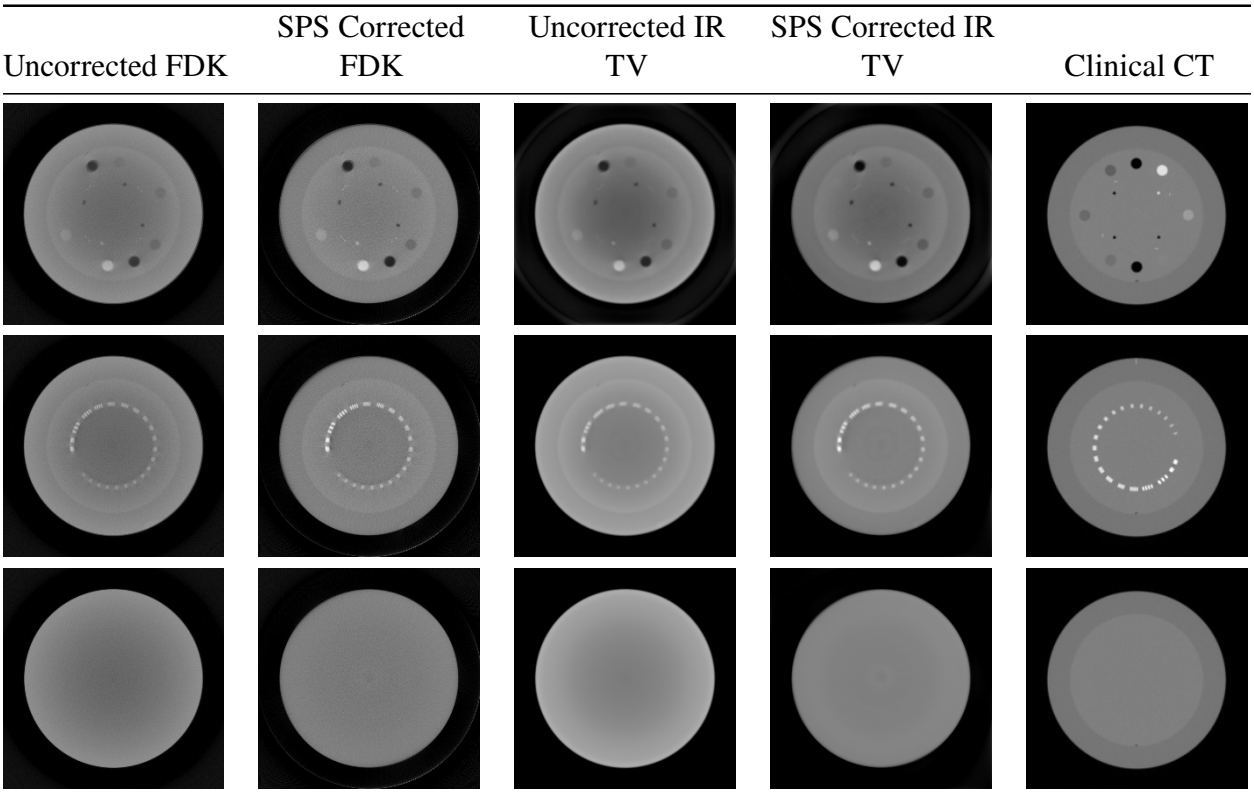

Figure 5.25: Reconstruction comparison across uncorrected FDK, SPS corrected FDK, uncorrected IR TV, SPS corrected IR TV, and clinical CT.

Chapter 6

Sources of Error and Limitations of the Experimental Study

The experimental results presented in Chapter 5 establish the SPS grid as a working proof of concept, visibly recovering the scatter-induced cupping artifact and improving HU linearity, uniformity, and contrast-to-noise ratio relative to the uncorrected reconstruction. These results were obtained, however, under a preliminary and largely un-optimized configuration: the source-to-object distance and tube output were selected pragmatically to yield usable images for an initial study rather than tuned for performance, and the apparatus departs substantially from the idealized geometry and photon statistics assumed in the computational study of Chapter 4. This chapter examines the principal sources of error and the limitations they impose, organized into five domains: the benchtop system and its geometric alignment, the prototype SPS grid, the preprocessing and correction applied to the raw projections, the reconstruction framework, and the scope of the phantom validation. Considered together, these factors indicate that the reported image-quality recovery represents a conservative lower bound on the achievable performance of the technique rather than its ceiling, with the gap relative to the computational predictions attributable in large part to identified and correctable non-idealities rather than to the SPS concept itself. One limitation is an important exception to this framing and works in the opposite direction — the use of a grid-free projection as

the scatter-correction target, discussed in Section 6.3, inflates the reported image quality relative to the true single-scan grid pipeline and is therefore examined in detail there. Each domain is examined in turn, and each points toward a specific corrective action that is consolidated into the path forward in Section 6.6.

6.1 Benchtop System and Geometric Alignment

The benchtop CBCT system models the third-generation scanning sequence using a moving reference frame in which the object is mounted on a rotating stage while the X-ray source and flat-panel detector remain stationary, in contrast to the clinical configuration in which a rigid source–detector assembly rotates about a stationary patient. This inversion is convenient for a benchtop apparatus but shifts the dominant geometric error from gantry sag to the alignment of the object, source, detector, and stage relative to a common beam axis, with consequences that propagate through the entire reconstruction. Owing to engineering difficulties, vendor delays, and the rudimentary state of the prototype platform, the system could not be brought to ideal geometric alignment for this study, and several distinct error sources combine to limit the achievable reconstruction quality.

A first limitation arises from incomplete characterization of the X-ray field. The beam dimensions specified by the source manufacturer (Spellman) were not originally incorporated into the system model, and the field was assumed to be circular for the purpose of geometry definition. In practice the field is not the assumed full circle: it is clipped by a sharp vertical cutoff on one side, so that the exposed region forms a dome-shaped circular segment and a vertical strip of the detector on the cutoff side receives no primary signal (**Figure 6.1**).

This truncation is addressed in preprocessing (Section 6.3) by cropping the detector to a balanced field of view, at the cost of usable area. A related limitation is the absence of source-output calibration: characterizing and stabilizing the tube output requires a suitable power supply operating from 240 V AC mains, which was not readily available in the clinical environment. As a consequence

the tube output was observed to fluctuate between projections without an absolute reference, and the residual fluctuation is addressed only approximately downstream (Section 6.3).

The dominant geometric errors are mechanical. The beam central axis is laterally shifted from the detector center, and the rotation axis of the rotary stage is offset from, and not co-linear with, the beam path. In addition to these in-plane offsets, the detector and the source each carry some residual angular misalignment: the detector is tilted — a pitch — from imperfect leveling, while the source is rotated about the vertical axis — a yaw — so that the beam field is canted horizontally rather than projecting squarely onto the detector. The source yaw is attributable in large part to the steel mounting frame, which was modeled slightly oversized and therefore does not seat flush against the benchtop cage. This frame-induced yaw is plausibly the compound result of several tolerances stacking unfavorably: the tolerances built into the 3-D model, the fabrication tolerances of the third-party manufacturer, measurement error incurred when the physical dimensions were taken, and the construction tolerance of the benchtop assembly itself.

These offsets and tilts persist because the platform lacks the precision actuation needed to correct them. There are no micrometer-precision stages for detector translation orthogonal to the beam path, nor for the rotary-stage position along the beam path or orthogonal to it in either the vertical or horizontal direction. Source positioning is especially coarse: the source assembly is heavy, its vertical adjustment consists of loosening the mount and physically lifting it to the approximate height, and its horizontal motion orthogonal to the beam requires loosening brackets and sliding the entire cage. Laser alignment qualitatively reveals the residual offsets, but in the absence of motorized micrometer control the residual error cannot be reduced below the precision of manual positioning. This precision floor is decisive: the per-collimator angular tolerance required by the smart-pixel design — within approximately 0.04° of the ideal source-focused angle for a 1 mm collimator, as established in Chapter 4 — is well below what manual alignment of the grid and stage can reliably achieve. Ideally the detector would be mounted within the same cage as the source to eliminate the relative leveling error directly, but this would require a lengthy and unanticipated redesign and rebuild of the platform.

A final system-level limitation concerns the acquisition electronics and timing. The benchtop is presently controlled through a rudimentary control board that is prone to intermittent connection errors; a more robust printed circuit board (PCB) is needed both to improve reliability and to make subsequent modification and debugging tractable. The source–detector exposure-acquisition synchronization was calibrated from the component documentation supplemented by a brute-force parameter search, and while this reduced the per-projection exposure variation, the timing is not fully optimized and likely contributes residual variation that propagates into the projections.

Collectively, these geometric and system errors do more than degrade the images directly: they introduce artifacts independent of scatter that contaminate the quantitative metrics, and, critically, they preclude the use of the alternating-minimization reconstruction of Chapter 4, whose per-iteration behavior amplifies rather than tolerates geometric inconsistency (Section 6.4).

6.2 Prototype SPS Grid

The prototype SPS grid was devised to test implementation feasibility rather than to realize a specific sampling density or optimized collimator geometry, and it departs from the computational design of Chapter 4 in several respects that together account for a substantial portion of the experimental-to-computational performance gap.

The prototype is undersized relative to the irradiated field, so its smart-pixel collimators sample only a small central portion of the projection. The scatter distribution must therefore be recovered largely by extrapolation beyond the sampled region rather than by interpolation within it, which is intrinsically less reliable and reduces the demonstrated robustness of the correction (the validation consequences are discussed in Section 6.5). A direct artifact of the prototype's small size is that the 3-D-printed collimator holders appear within the projections; in a full-size grid spanning the field these holders would lie outside the imaged region and would not contribute to the measured signal.

The prototype collimators are also larger than optimal, with an inner diameter of 3 mm and an outer diameter of 6 mm (**Figure 6.2**), against the optimized 1 mm inner and 4 mm outer diameters

identified computationally (**Figure 6.3**). The oversized smart pixels occlude a substantially larger fraction of the detector and produce more pronounced penumbral effects under misalignment, and the resulting interpolation errors manifest as the severe streaking artifacts that rendered direct SPS-grid reconstruction unfeasible (**Figure 5.17**).

Compounding the dimensional non-idealities, the smart-pixel collimators are imperfectly aligned to the source. The manufacturer documentation did not specify the true beamline center, so although laser alignment of the grid was implemented, it was referenced to an approximate rather than the true focal-spot-defined center, leaving the collimators not uniformly focused toward the focal spot. The dedicated alignment grid developed alongside the prototype showed promise as an alignment aid, but realizing its benefit requires a fully mechanically aligned system and more development time than was available; on the present platform, with the geometric errors of Section 6.1 unresolved, the alignment grid could not be fully exploited.

6.3 Projection Preprocessing and Correction

The preprocessing pipeline described in Section 5.5.2 was necessary to render the misaligned benchtop projections usable, but each correction carries an assumption whose validity limits the interpretation of the results, and the pipeline as a whole works without being perfect, leaving residual alignment error in the corrected projections.

Although the physical beam field is truncated on only one side, leaving the dome-shaped exposed region traced to Section 6.1, pixels were cropped from both sides of the detector to keep the reconstruction field of view balanced about the center, reducing the array to 965×1065. This balancing necessarily sacrifices exposed pixels on the untruncated side as well, so the crop reduces the effective FOV for reconstruction. Darkfield and blankfield references were time-averaged and applied for offset subtraction and line-integral calculation in the standard manner, under the assumption of consistent output that the fluctuating source output does not fully satisfy.

The per-projection source-output fluctuation, together with the need to compare reconstructions acquired under physically different beam-path configurations (with and without the SPS grid, and the intended beam stop array arm), was compensated using a shared air-ROI scaling factor (Equation 5.7) that renormalizes each projection's overall fluence to the blank-field reference. This correction is necessarily ad hoc: it assumes the inter-projection and inter-configuration variation is a spatially uniform multiplicative gain change, correcting only a global offset and not any spatial or spectral redistribution of the output. Whether this single-ROI normalization is sufficient or whether it introduces additional error of its own when reconciling the different scan configurations remains to be investigated, and any uncorrected component propagates into both the scatter estimate and the final reconstruction.

The geocentric orbit of the phantom, which results from the object's misalignment on the rotation stage and cannot be removed through the standard geometric reconstruction parameters, was corrected by a cross-correlation centroid alignment that exploits the cylindrical symmetry of the Catphan, with subpixel shifts handled by upsampling and downsampling. This method is effective for the symmetric Catphan but is fundamentally limited: it relies on a symmetry that clinical anatomy does not possess, and the subpixel resampling introduces a small interpolation error. The dependence on phantom symmetry is therefore a barrier to applying the present preprocessing pipeline to asymmetric or anthropomorphic geometries without a different alignment strategy, such as an explicit geometric calibration phantom.

The most consequential preprocessing limitation concerns the SPS projections themselves, and it warrants emphasis because it bears directly on what the experiment does and does not establish. As reported in Section 5.5.1, attempts to perform SPS scatter correction directly on the grid projections produced streaking artifacts from interpolation errors — driven by the oversized, imperfectly aligned smart pixels of Section 6.2 — severe enough to render the reconstruction unfeasible for analysis. To proceed, the scatter distribution was still estimated from the grid scan, but the projection that was corrected and reconstructed was the separate, grid-free uncorrected acquisition rather than the grid projection itself. Because the grid-free projection contains no smart-pixel dead zones,

this substitution removes the dominant source of interpolation error and is the principal reason the reconstructed image quality is as high as reported; the reported metrics should accordingly be read as the performance obtainable given a good scatter estimate under near-ideal interpolation, not as the performance of the full grid-based acquisition in its present prototype form.

Two consequences follow, pulling in opposite directions. On one hand, the scatter-measurement concept at the heart of the technique is not diminished: the scatter estimate was genuinely derived from the grid's sparse primary samples, so the experiment does validate the grid's ability to sample the primary signal and infer the scatter distribution, and the recovered image quality confirms that this estimate is accurate enough to support effective correction. On the other hand, the demonstration is no longer the single-scan, dose-neutral correction that defines the SPS method and distinguishes it from two-scan beam-stop approaches: applying a scatter map measured in one acquisition to a separate grid-free acquisition is effectively a two-scan procedure. The experiment therefore does not yet demonstrate the dose- and scan-reduction advantage that motivates the method, and it confounds a like-for-like comparison with the computational study of Chapter 4, in which a single grid projection supplied both the smart-pixel primary samples and the total signal to be corrected. Unlike the geometric and grid non-idealities catalogued elsewhere in this chapter, which suppress the reported performance below the method's potential, this substitution inflates image recovery relative to the current single-scan grid pipeline; the two effects must be kept distinct when interpreting the results. Recovering acceptable image quality directly from the grid projection — closing the dead-zone interpolation error rather than circumventing it — is therefore the central preprocessing objective for the next iteration and the prerequisite for an experimental result that is genuinely comparable to the computational prediction and that substantiates the single-scan feasibility claim.

6.4 Reconstruction Framework

The reconstruction algorithms used in this study, FDK and total-variation–regularized iterative reconstruction (IR TV), were adopted out of necessity rather than designed for the SPS sampling

scheme, and their parameters were not finely tuned to its requirements. The alternating-minimization algorithm developed and validated computationally in Chapter 4 which jointly estimates the scatter field and the image using total-variation regularization in the image domain and L2-norm (Tikhonov) regularization in the scatter domain was not employed experimentally. Its joint regularization is precisely the mechanism that smooths over the interpolation errors arising from occluded smart-pixel regions, and would in principle have been the most appropriate reconstruction for SPS data. However, the residual geometric misalignments described in Section 6.1 are amplified at each alternating iteration, accumulating into circular artifacts that dominate the final image. The alternating-minimization reconstruction is therefore contingent on a level of geometric calibration the present apparatus cannot yet provide.

Two further mismatches limit the FDK and IR TV reconstructions. First, the forward model implicit in any SPS-aware reconstruction assumes collimators uniformly focused toward the focal spot and a well-characterized circular acquisition geometry; the experimental geometry violates both, so the reconstruction operates on data inconsistent with its model. FDK in particular is an approximate algorithm that assumes an ideal circular trajectory and is not robust to the residual orbit and offset present here. Second, the regularization weights were selected against the noise characteristics of the simulation rather than the experimental projections, and no spatially varying data-fidelity weighting was applied at the occluded or penumbral pixels surrounding each smart pixel. The observation in Section 5.6.5 that SPS-corrected IR TV outperformed SPS-corrected FDK across every metric is consistent with this analysis: the more capable reconstruction recovers more of the available scatter correction, implying that a reconstruction tailored to the SPS sampling pattern — once the geometry permits it — would recover substantially more.

6.5 Validation Scope and Generalizability

The experimental validation was restricted to a single Catphan 503 phantom, which bounds what the results can demonstrate in several respects. Most directly, the undersized prototype grid samples

only a small central portion of the irradiated field, so the directly measured scatter samples cover predominantly the central region of the scatter distribution and the correction relies on extrapolating the distribution beyond the sampled region — extrapolation that is less reliable than interpolation and that reduces the demonstrated robustness of the method. The truncation described in Section 6.3 narrows the sampled region further.

The choice of phantom also limits generalizability. The Catphan is near-cylindrically symmetric and largely homogeneous; while ideal for standardized image-quality metrics, it does not exercise the high SPR, heterogeneous regime in which scatter correction is most demanding and the SPS technique is most clinically motivated. No anatomically realistic or anthropomorphic phantom was evaluated, and no large-volume phantom approximating the torso or pelvis, geometries that generate the most challenging scatter and that the computational study of Chapter 4 identified as the hardest cases, was included. A further consequence of the geometric errors of Section 6.1 is that the quantitative metrics extracted from the Catphan may themselves not be fully robust: residual geometric distortion in the reconstruction broadens edges and perturbs uniformity independently of scatter, so the spatial-resolution metrics in particular (Section 5.6.2) should be read as relative indicators rather than absolute characterizations.

The study also lacks the reference arms against which the SPS reconstruction was intended to be benchmarked. The BSA comparison, designed in Section 5.3 to provide a near-ideal direct measurement of scatter, was not implemented owing to several practical difficulties with the BSA setup that together prevented robust scatter-to-primary-ratio measurements under the present geometry. First, the array could not be reliably centered in the beam path, leaving its lead discs laterally offset from their intended positions relative to the field. Second, the available mounting placed the BSA too close to the source, so the projected shadows of the lead discs were magnified to the point of heavily reducing the overall exposure reaching the detector. Third, as a consequence of this geometry, only a few of the array's lead inserts fell within the detector field of view, leaving too few scatter samples to interpolate a reliable full-field scatter map. Properly mounting the BSA at an appropriate source-to-array distance and in correct lateral alignment would require either a

redesign of the benchtop cage or a dedicated 3-D-printed mounting fixture, neither of which are possible without more time. In its place, a clinical CT-simulator scan of an equivalent Catphan was used as the ideal-correction reference. This substitution carries two caveats. First, the clinical scans were not acquired at the same reconstruction resolution as the benchtop data and had to be upscaled to compute and compare metrics, which limits the fidelity of the comparison, particularly for spatial resolution. Second, the clinical CT simulator is in some sense an unfair benchmark: it is unrealistic to expect a CBCT system relying on hardware-based scatter correction alone to approach the image quality of a diagnostic CT simulator, so the clinical scan establishes an aspirational rather than a representative upper bound. A more reasonable benchmark — a clinical CBCT equipped with an anti-scatter grid, representing the current standard of care for hardware-based scatter rejection — was not included in this study. Taken together, the present results are bounded only by the uncorrected case below and an idealized CT reference above, and should be regarded as investigatory rather than conclusive; broader phantom studies are required to establish clinical utility. Most importantly, because the correction was applied to a grid-free acquisition rather than the grid projection itself (Section 6.3), the single-scan, dose-neutral advantage that distinguishes the SPS method from two-scan techniques — and that is the principal claim the experimental study was intended to support — remains to be demonstrated on hardware.

Table 6.1. Summary of the experimental sources of error, their impact on reconstructed image quality, and recommended remedies for future iterations.

Source of error	Impact on image quality	Recommended Solution
Benchtop system and geometric alignment		
Beam-axis, stage-axis, detector-pitch, and source-yaw misalignment	Geometric distortion and edge blurring (reduced spatial resolution); scatter-independent artifacts; corrupted quantitative metrics; precludes the alternating-minimization reconstruction	Motorized micrometer stages on all axes with explicit geometric calibration; level the detector pitch and square the source-mounting frame to remove the yaw; co-mount the detector in the source cage
Source-output fluctuation (no output calibration; 240 V supply unavailable)	Projection-to-projection gain errors producing ring and shading artifacts and HU inaccuracy	240 V power supply enabling source-output calibration; per-projection output monitoring
Acquisition electronics and timing (rudimentary control board; imperfect exposure–acquisition synchronization)	Intermittent dropouts and residual exposure variation, contributing artifacts and HU inaccuracy	Robust PCB controller and refined exposure–acquisition synchronization
Un-optimized acquisition technique (source-to-object distance and exposure chosen pragmatically)	Suboptimal scatter-to-primary ratio and noise, reducing contrast/CNR and HU accuracy	Optimize geometry and exposure once the platform is mechanically stable
Prototype SPS grid		
Undersized grid sampling only the central field, with 3-D-printed holders falling within the imaged region	Scatter recovered by extrapolation, plus holder occlusion, producing residual cupping/shading and streak artifacts and peripheral HU inaccuracy	Refabricate at full field-spanning size so holders lie outside the field; non-uniform sampling tailored to the scatter distribution
Oversized smart pixels (3 mm ID / 6 mm OD vs optimized 1 mm / 4 mm)	Large dead zones and penumbra producing severe streaking artifacts and degraded resolution and contrast	Refabricate at the optimized collimator dimensions (1 mm ID, 4 mm OD)
Collimators not focused to the true focal spot (incomplete documentation; laser referenced to an approximate center)	Smart-pixel scatter contamination and interpolation error, producing streak artifacts and HU inaccuracy	Align to the true beamline center using the alignment grid on a calibrated system
Projection preprocessing and correction		
One-sided field truncation with symmetric crop	Reduced field of view and loss of peripheral scatter samples, producing truncation and extrapolation artifacts	Correct the field characterization and beam centering to remove the truncation
Ad-hoc air-ROI flux normalization across scan configurations	Imperfect inter-projection and inter-configuration gain matching, producing shading artifacts and HU inaccuracy	Ground normalization in measured source output via per-projection monitoring
Symmetry-dependent cross-correlation centering; residual mis-centering and subpixel resampling	Residual geometric distortion and edge blurring (reduced resolution); does not generalize to asymmetric anatomy	Explicit geometric calibration independent of phantom symmetry
Grid-free projection used as the correction target in place of the grid projection	Removes dead-zone interpolation error, inflating apparent resolution, contrast, and HU recovery; not a single-scan demonstration and confounds comparison with the computational study	Recover image quality directly from the grid projection: total-signal recovery in covered pixels and an SPS-tailored reconstruction
Reconstruction framework		
Alternating-minimization reconstruction unusable under residual geometry	Per-iteration amplification of geometric error, producing circular artifacts	Re-enable after geometric calibration; add spatially varying data-fidelity weighting at occluded and penumbral pixels
Borrowed, untuned reconstruction (FDK / IR-TV; regularization tuned on simulation noise)	Suboptimal scatter-correction recovery: residual cupping and limited contrast and resolution gains	Reconstruction tailored to the SPS sampling pattern with experiment-matched regularization
Validation scope and generalizability		
Single symmetric Catphan; no anthropomorphic or large high-SPR phantom	Untested in the high-SPR, heterogeneous regime where artifacts and HU errors are most severe; limited generalizability	Extend to larger, heterogeneous, and anthropomorphic phantoms across the SPR regimes
Missing BSA and clinical-CBCT/ASG benchmarks; clinical-CT reference upscaled	No matched upper-bound benchmark and geometry-driven uncertainty in resolution, contrast, and HU metrics	Add a BSA reference arm and a clinical-CBCT-with-anti-scatter-grid benchmark; match reconstruction resolution

6.6 Implications for Future Iterations

The limitations above as summarized by **Table 6.1** define a clear and largely correctable path for the next iteration of the experimental program. On the platform, the priority is to replace coarse manual positioning with motorized micrometer-controlled stages on every degree of freedom that currently lacks one — detector translation orthogonal to the beam, rotary-stage position along and orthogonal to the beam in both vertical and horizontal directions, and source elevation and lateral translation — supported by an explicit geometric calibration and by correction of the detector pitch and the source yaw. Squaring or remodeling the oversized source-mounting frame to seat flush would remove the frame-induced yaw of the beam field, and co-mounting the detector within the source cage, though it entails a redesign, would eliminate the relative leveling error at its source. A robust PCB-based control board should replace the present rudimentary board to remove intermittent connection errors and to ease modification and debugging, and acquiring a suitable 240 V AC power supply would enable source-output calibration, potentially removing the residual fluctuation that the air-ROI normalization currently approximates. Together these changes would suppress the dominant system-level errors and, in doing so, re-enable the alternating-minimization reconstruction by removing the geometric inconsistency that currently causes its iterations to diverge.

On the hardware and algorithmic side, the grid should be refabricated at the optimized collimator dimensions (1 mm inner, 4 mm outer diameter), at full field-spanning size so that the holders fall outside the imaged region, and beam-centered with the aid of the alignment grid on a properly aligned system. The reconstruction should be adapted to the SPS sampling pattern through spatially varying data-fidelity weighting at occluded and penumbral pixels, non-uniform sampling tailored to the scatter distribution, and recovery of the total signal in tube-covered pixels. As an intermediate validation and debugging aid, modeling the benchtop system itself in Monte Carlo — including its measured misalignments — would provide a controlled sanity check against which the experimental reconstructions could be compared and through which specific alignment errors could be diagnosed and attributed. Finally, the validation must be extended beyond the single symmetric phantom to larger, heterogeneous, and anthropomorphic phantoms spanning the low- and high-SPR regimes

examined computationally, and should incorporate both the BSA reference arm and a clinical CBCT-with-anti-scatter-grid benchmark, so that the achievable image quality can be assessed against both a direct scatter measurement and a representative standard of care. Addressing these items is expected to close much of the gap between the experimental performance reported in Chapter 5 and the theoretical performance established in Chapter 4, positioning the technique for evaluation on a clinical CBCT platform where its single-scan, dose-neutral, and mechanically static character offers a direct integration path.

6.7 Figures

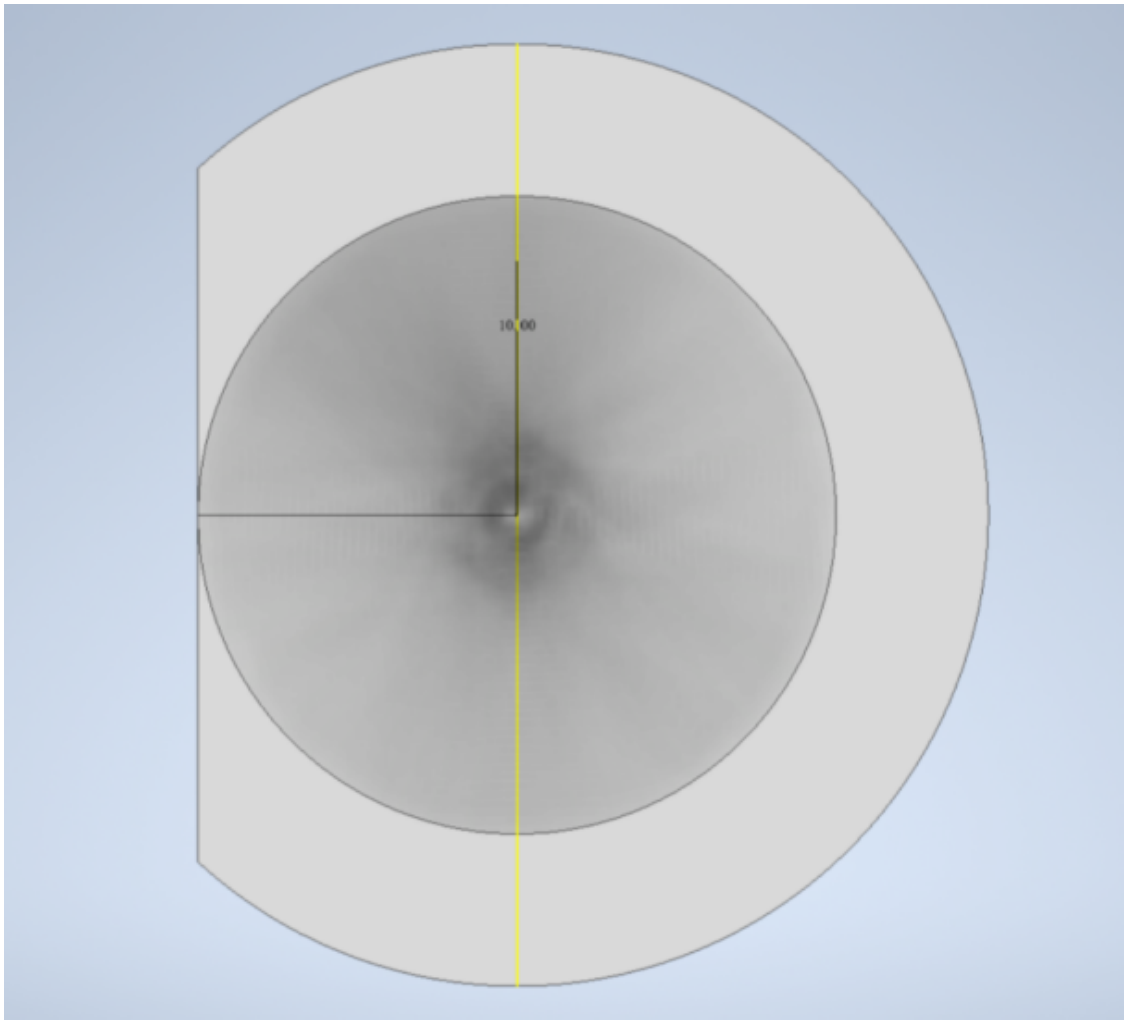

Figure 6.1: Circular truncated field geometry of the beam provided by Spellman. For reference, 10 cm is demarcated by the black line.

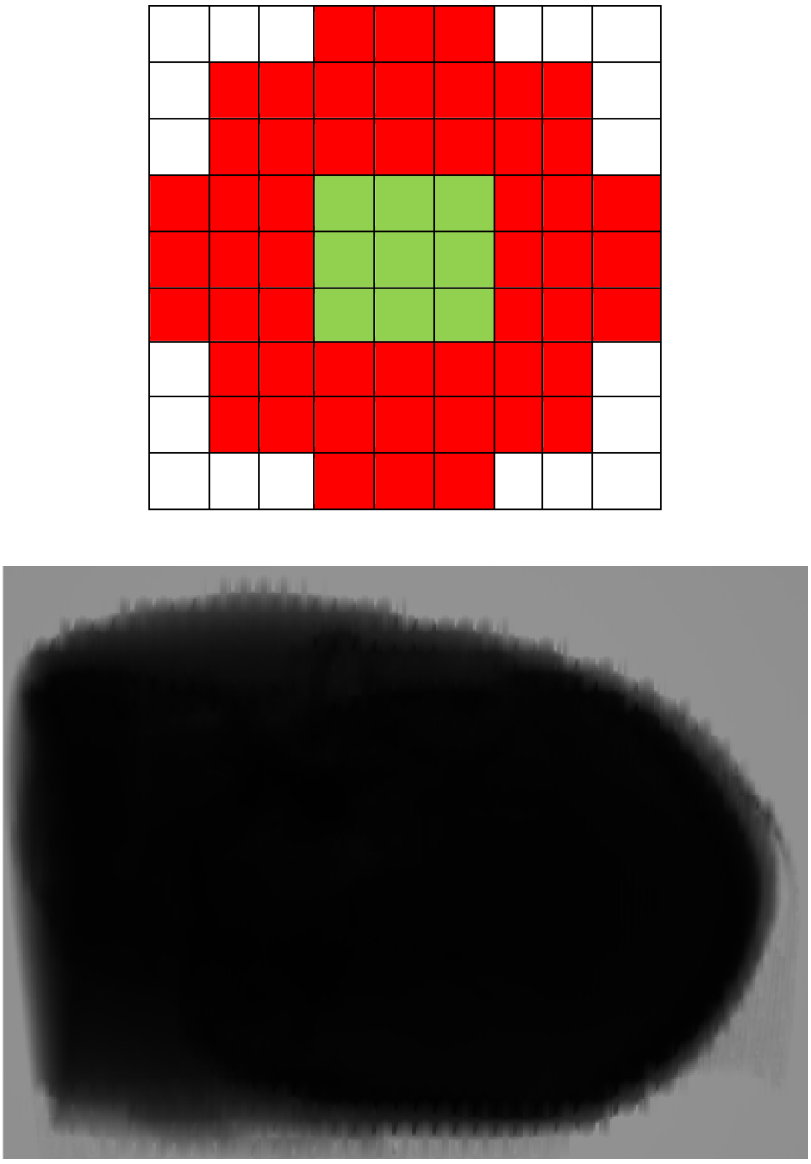

Figure 6.2: Simulated 3 mm ID and 6 mm OD tubes previously used in the initial computation study conception which resulted in streak artifacts from interpolation errors in the projection. Red pixels represent blocked pixels while green pixels represent primary only smart pixels.

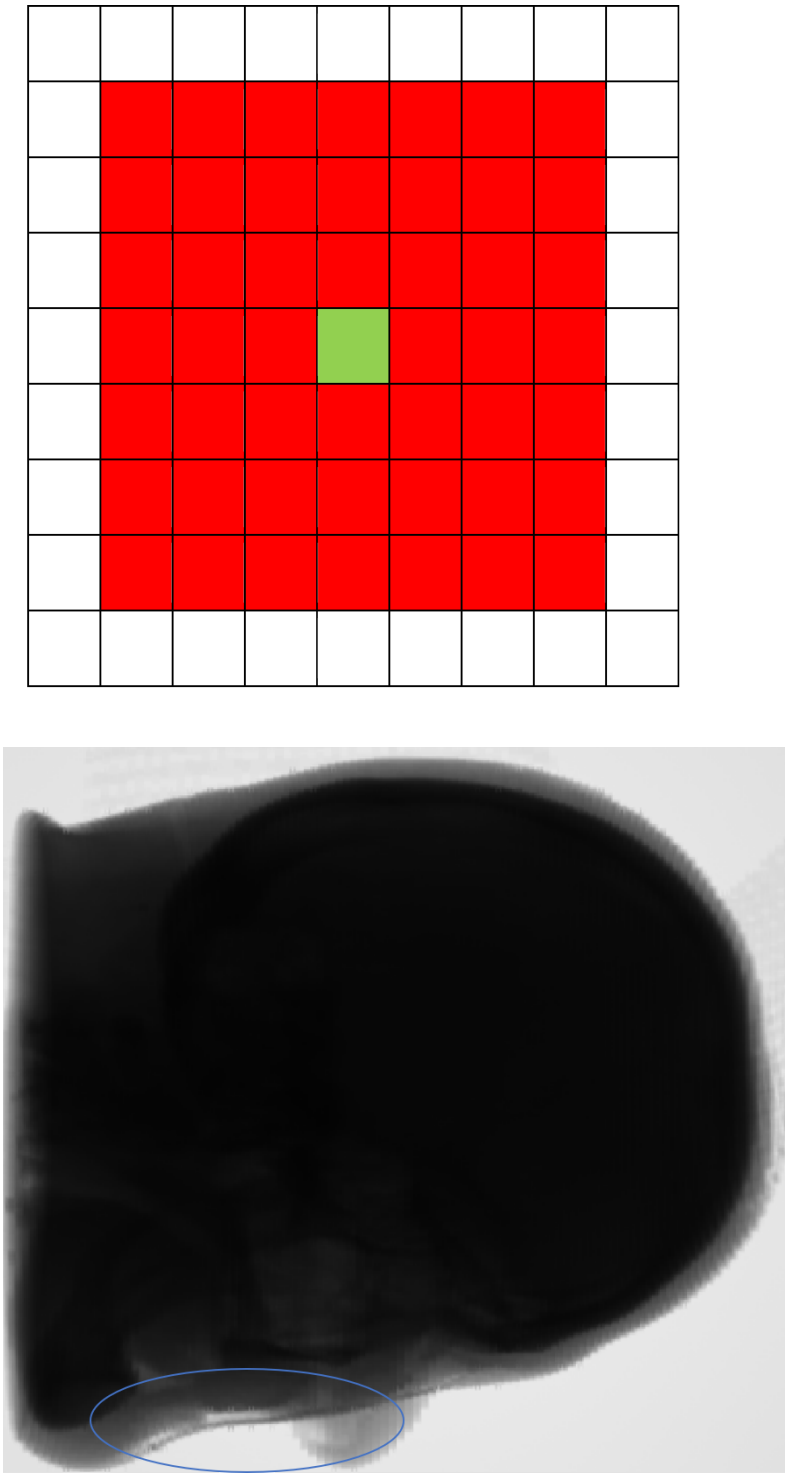

Figure 6.3: Simulated 1 mm ID and 4 mm OD tubes previously used in the initial computation study conception mitigated streak artifacts by minimizing detector footprint and interpolation error. Red pixels represent blocked pixels while green pixels represent primary only smart pixels.

References

[1] H. K. Tuy. “An inversion formula for cone-beam reconstruction.” *SIAM Journal on Applied Mathematics*, **43**(3), 1983, pp. 546–552. DOI: 10.1137/0143035.

[2] V. Babiano, J. Balibrea, L. Caballero, D. Calvo, I. Ladarescu, J. Lerendegui, S. Mira Prats, and C. Domingo-Pardo. “First i-TED demonstrator: A Compton imager with Dynamic Electronic Collimation.” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **953**, Feb. 2020, p. 163228. ISSN: 0168-9002. DOI: 10.1016/j.nima.2019.163228.

[3] F. Salvat and J. M. Fernández-Varea. “Overview of physical interaction models for photon and electron transport used in Monte Carlo codes.” *Metrologia*, **46**(2), Mar. 2009, S112. ISSN: 0026-1394. DOI: 10.1088/0026-1394/46/2/S08.

[4] G. Hernández and F. Fernández. “A model of tungsten anode x-ray spectra.” *Medical Physics*, **43**(8Part1), 2016, pp. 4655–4664. ISSN: 2473-4209. DOI: 10.1118/1.4955120.

[5] L. W. Goldman. “Principles of CT and CT Technology.” *Journal of Nuclear Medicine Technology*, **35**(3), Sept. 2007, pp. 115–128. ISSN: 0091-4916, 1535-5675. DOI: 10.2967/jnmt.107.042978.

[6] H. Li, R. Mohan, and X. R. Zhu. “Derivation of Scatter Kernel in CBCT Imaging System.” In: *2006 Fortieth Asilomar Conference on Signals, Systems and Computers*. Oct. 2006, pp. 1851–1854. DOI: 10.1109/ACSSC.2006.355082. URL: https://ieeexplore.ieee.org/document/4176892.

[7] H. Li, R. Mohan, and X. R. Zhu. "Scatter kernel estimation with an edge-spread function method for cone-beam computed tomography imaging." *Physics in Medicine & Biology*, **53**(23), Nov. 2008, p. 6729. ISSN: 0031-9155. DOI: 10.1088/0031-9155/53/23/006.

[8] J. M. Boone and J. A. Seibert. "An analytical model of the scattered radiation distribution in diagnostic radiology." *Medical Physics*, **15**(5), 1988, pp. 721–725. ISSN: 2473-4209. DOI: 10.1118/1.596186.

[9] G. Mettivier, N. Lanconelli, S. L. Meo, and P. Russo. "Scatter Correction in Cone-Beam Breast Computed Tomography: Simulations and Experiments." *IEEE Transactions on Nuclear Science*, **59**(5), Oct. 2012, pp. 2008–2019. ISSN: 1558-1578. DOI: 10.1109/TNS.2012.2201955.

[10] G. Poludniowski, P. M. Evans, V. N. Hansen, and S. Webb. "An efficient Monte Carlo-based algorithm for scatter correction in keV cone-beam CT." *Physics in Medicine & Biology*, **54**(12), June 2009, p. 3847. ISSN: 0031-9155. DOI: 10.1088/0031-9155/54/12/016.

[11] A. Sisniega, W. Zbijewski, J. Xu, H. Dang, J. W. Stayman, J. Yorkston, N. Aygun, V. Koliatsos, and J. H. Siewerdsen. "High-fidelity artifact correction for cone-beam CT imaging of the brain." *Physics in Medicine & Biology*, **60**(4), Jan. 2015, p. 1415. ISSN: 0031-9155. DOI: 10.1088/0031-9155/60/4/1415.

[12] A. Badal and A. Badano. "Accelerating Monte Carlo simulations of photon transport in a voxelized geometry using a massively parallel graphics processing unit." *Medical Physics*, **36**(11), Nov. 2009, pp. 4878–4880. ISSN: 0094-2405. DOI: 10.1118/1.3231824.

[13] W. Zbijewski and F. Beekman. "Efficient Monte Carlo based scatter artifact reduction in cone-beam micro-CT." *IEEE Transactions on Medical Imaging*, **25**(7), July 2006, pp. 817–827. ISSN: 1558-254X. DOI: 10.1109/TMI.2006.872328.

[14] L. Zhu, Y. Xie, J. Wang, and L. Xing. "Scatter correction for cone-beam CT in radiation therapy." *Medical Physics*, **36**(6Part1), 2009, pp. 2258–2268. ISSN: 2473-4209. DOI: 10.1118/1.3130047.

[15] Y. Xu, T. Bai, H. Yan, L. Ouyang, A. Pompos, J. Wang, L. Zhou, S. B. Jiang, and X. Jia. "A Practical Cone-beam CT Scatter Correction Method with Optimized Monte Carlo Simulations for Image-Guided Radiation Therapy." *Physics in medicine and biology*, **60**(9), May 2015, pp. 3567–3587. ISSN: 0031-9155. DOI: 10.1088/0031-9155/60/9/3567.

[16] A. Maslowski, A. Wang, M. Sun, T. Wareing, I. Davis, and J. Star-Lack. "Acuros CTS: A fast, linear Boltzmann transport equation solver for computed tomography scatter - Part I: Core algorithms and validation." *Medical Physics*, **45**(5), May 2018, pp. 1899–1913. ISSN: 2473-4209. DOI: 10.1002/mp.12850.

[17] A. Wang et al. "Acuros CTS: A fast, linear Boltzmann transport equation solver for computed tomography scatter - Part II: System modeling, scatter correction, and optimization." *Medical Physics*, **45**(5), May 2018, pp. 1914–1925. ISSN: 2473-4209. DOI: 10.1002/mp.12849.

[18] Y. Nomura, Q. Xu, H. Shirato, S. Shimizu, and L. Xing. "Projectiondomain scatter correction for cone beam computed tomography using a residual convolutional neural network." *Medical Physics*, **46**(7), July 2019, pp. 3142–3155. ISSN: 0094-2405. DOI: 10.1002/mp.13583.

[19] S. Xie, C. Yang, Z. Zhang, and H. Li. "Scatter Artifacts Removal Using Learning-Based Method for CBCT in IGRT System." *IEEE Access*, **6**, 2018, pp. 78031–78037. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2018.2884704.

[20] J. Maier, E. Eulig, T. Vöth, M. Knaup, J. Kuntz, S. Sawall, and M. Kachelrieß. "Real-time scatter estimation for medical CT using the deep scatter estimation: Method and robustness analysis with respect to different anatomies, dose levels, tube voltages, and data truncation." *Medical Physics*, **46**(1), 2019, pp. 238–249. ISSN: 2473-4209. DOI: 10.1002/mp.13274.

[21] M. Rossi and P. Cerveri. "Comparison of Supervised and Unsupervised Approaches for the Generation of Synthetic CT from Cone-Beam CT." *Diagnostics*, **11**(8), Aug. 2021, p. 1435. ISSN: 2075-4418. DOI: 10.3390/diagnostics11081435.

[22] F. Madesta, T. Sentker, C. Rohling, T. Gauer, R. Schmitz, and R. Werner. "Monte Carlo-based simulation of virtual 3 and 4-dimensional cone-beam computed tomography from computed

tomography images: An end-to-end framework and a deep learning-based speedup strategy." *Physics and Imaging in Radiation Oncology*, **32**, Sept. 2024, p. 100644. ISSN: 2405-6316. DOI: 10.1016/j.phro.2024.100644.

[23] E.-P. Rührnschopf and K. Klingenbeck. "A general framework and review of scatter correction methods in x-ray cone-beam computerized tomography. Part 1: Scatter compensation approaches." *Medical Physics*, **38**(7), 2011, pp. 4296–4311. ISSN: 2473-4209. DOI: 10.1118/1.3599033.

[24] L. Ren, Y. Chen, Y. Zhang, W. Giles, J. Jin, and F.-F. Yin. "Scatter Reduction and Correction for Dual Source Cone-beam CT using Pre-patient Grids." *Technology in cancer research & treatment*, **15**(3), June 2016, pp. 416–427. ISSN: 1533-0346. DOI: 10.1177/1533034615587615.

[25] L. Zhu, N. R. Bennett, and R. Fahrig. "Scatter Correction Method for X-Ray CT Using Primary Modulation: Theory and Preliminary Results." *IEEE Transactions on Medical Imaging*, **25**(12), Dec. 2006, pp. 1573–1587. ISSN: 1558-254X. DOI: 10.1109/TMI.2006.884636.

[26] H. Lee, L. Xing, R. Lee, and B. P. Fahimian. "Scatter correction in cone-beam CT via a half beam blocker technique allowing simultaneous acquisition of scatter and image information." *Medical Physics*, **39**(5), 2012, pp. 2386–2395. ISSN: 2473-4209. DOI: 10.1118/1.3691901.

[27] T. Niu and L. Zhu. "Scatter correction for full-fan volumetric CT using a stationary beam blocker in a single full scan." *Medical Physics*, **38**(11), 2011, pp. 6027–6038. ISSN: 2473-4209. DOI: 10.1118/1.3651619.

[28] J. S. Maltz et al. "Focused beam-stop array for the measurement of scatter in megavoltage portal and cone beam CT imaging." *Medical Physics*, **35**(6Part1), 2008, pp. 2452–2462. ISSN: 2473-4209. DOI: 10.1118/1.2924220.

[29] R. Ning, X. Tang, and D. Conover. "X-ray scatter correction algorithm for cone beam CT imaging." *Medical Physics*, **31**(5), 2004, pp. 1195–1202. ISSN: 2473-4209. DOI: 10.1118/1.1711475.

[30] X. Liang et al. “Scatter correction for a clinical cone-beam CT system using an optimized stationary beam blocker in a single scan.” *Medical Physics*, **46**(7), 2019, pp. 3165–3179. ISSN: 2473-4209. DOI: 10.1002/mp.13568.

[31] X. Liu, C. C. Shaw, T. Wang, L. Chen, M. C. Altunbas, and S. C. Kappadath. “An Accurate Scatter Measurement and Correction Technique for Cone Beam Breast CT Imaging Using Scanning Sampled Measurement (SSM) Technique.” *Proceedings of SPIE–the International Society for Optical Engineering*, **6142**, Feb. 2006, pp. 614234–1-614234-7. ISSN: 0277-786X. DOI: 10.1117/12.656655.

[32] C. Zhao, Y. Zhong, J. Wang, and M. Jin. “Simultaneous Dose Reduction and Scatter Correction for 4D Cone-Beam Computed Tomography.” In: *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. Oct. 2017, pp. 1–3. DOI: 10.1109/NSSMIC.2017.8532916. URL: https://ieeexplore.ieee.org/document/8532916.

[33] M. Uffmann and C. Schaefer-Prokop. “Digital radiography: The balance between image quality and required radiation dose.” *European Journal of Radiology*. Digital Radiography, **72**(2), Nov. 2009, pp. 202–208. ISSN: 0720-048X. DOI: 10.1016/j.ejrad.2009.05.060.

[34] C. Altunbas, B. Kavanagh, T. Alexeev, and M. Miften. “Transmission characteristics of a two dimensional antiscatter grid prototype for CBCT.” *Medical Physics*, **44**(8), 2017, pp. 3952–3964. ISSN: 2473-4209. DOI: 10.1002/mp.12346.

[35] C. Altunbas, Y. Park, Z. Yu, and A. Gopal. “A unified scatter rejection and correction method for cone beam computed tomography.” *Medical Physics*, **48**(3), 2021, pp. 1211–1225. ISSN: 2473-4209. DOI: 10.1002/mp.14681.

[36] G. Wu, C. R. Inscoe, J. Calliste, J. Shan, Y. Z. Lee, O. Zhou, and J. Lu. “Estimating scatter from sparsely measured primary signal.” *Journal of Medical Imaging*, **4**(1), Jan. 2017, p. 013508. ISSN: 2329-4302. DOI: 10.1117/1.JMI.4.1.013508.

[37] B. Meng, H. Lee, L. Xing, and B. P. Fahimian. "Single-scan patient-specific scatter correction in computed tomography using peripheral detection of scatter and compressed sensing scatter retrieval." *Medical Physics*, **40**(1), 2013, p. 011907. ISSN: 2473-4209. DOI: 10.1118/1.4769421.

[38] F. Yang, D. Zhang, K. Huang, W. Shi, and X. Wang. "Scattering Estimation for Cone-Beam CT Using Local Measurement Based on Compressed Sensing." *IEEE Transactions on Nuclear Science*, **65**(3), Mar. 2018, pp. 941–949. ISSN: 1558-1578. DOI: 10.1109/TNS.2018.2803739.

[39] C. Zhao, L. Ouyang, J. Wang, and M. Jin. "Blind deconvolution for blocker-based scatter correction of CBCT." In: *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. Oct. 2015, pp. 1–3. DOI: 10.1109/NSSMIC.2015.7582043. URL: https://ieeexplore.ieee.org/document/7582043.

[40] J. Puente Torres, H. Crespo Sariol, J. Yperman, Á. Brito Sauvanell, R. Carleer, and J. Navarro Campa. "A novel X-ray radiography approach for the characterization of granular activated carbons used in the rum production." *Journal of Analytical Science and Technology*, **9**(1), Jan. 2018, p. 1. ISSN: 2093-3371. DOI: 10.1186/s40543-017-0133-x.

[41] G. Jarry, S. A. Graham, D. J. Moseley, D. J. Jaffray, J. H. Siewerdsen, and F. Verhaegen. "Characterization of scattered radiation in kV CBCT images using Monte Carlo simulations." *Medical Physics*, **33**(11), 2006, pp. 4320–4329. ISSN: 2473-4209. DOI: 10.1118/1.2358324.

[42] J. Hubbell and S. Seltzer. *Tables of X-Ray Mass Attenuation Coefficients and Mass Energy-Absorption Coefficients (version 1.4)*. http://physics.nist.gov/xaamdi. 2004. URL: http://physics.nist.gov/xaamdi.

[43] M. P. Jahan, M. Rahman, and Y. S. Wong. "A review on the conventional and micro-electrodischarge machining of tungsten carbide." *International Journal of Machine Tools and Manufacture*, **51**(12), Dec. 2011, pp. 837–858. DOI: 10.1016/j.ijmachtools.2011.08.016.

[44] N. Parikh and S. Boyd. "Proximal Algorithms." *Foundations and Trends in Optimization*, **1**(3), Jan. 2014, pp. 127–239. ISSN: 2167-3888. DOI: 10.1561/2400000003.

[45] J. A. Fessler. *Michigan Image Reconstruction Toolbox*. https://github.com/JeffFessler/mirt. URL: https://github.com/JeffFessler/mirt.

[46] A. Beck and M. Teboulle. "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems." *SIAM Journal on Imaging Sciences*, **2**(1), Jan. 2009, pp. 183–202. DOI: 10.1137/080716542.

[47] Q. Lyu, C. Yang, H. Gao, Y. Xue, D. O'Connor, T. Niu, and K. Sheng. "Technical Note: Iterative megavoltage CT (MVCT) reconstruction using block-matching 3D-transform (BM3D) regularization." *Medical Physics*, **45**(6), 2018, pp. 2603–2610. ISSN: 2473-4209. DOI: 10.1002/mp.12916.

[48] W. Zbijewski and F. Beekman. "Fast scatter estimation for cone-beam X-ray CT by combined Monte Carlo tracking and Richardson-Lucy fitting." In: *IEEE Symposium Conference Record Nuclear Science 2004.* Vol. 5. Oct. 2004, pp. 2774–2777. DOI: 10.1109/NSSMIC.2004.1466264. URL: https://ieeexplore.ieee.org/document/1466264.

[49] A. Colijn and F. Beekman. "Accelerated simulation of cone beam X-ray scatter projections." *IEEE Transactions on Medical Imaging*, **23**(5), May 2004, pp. 584–590. ISSN: 1558-254X. DOI: 10.1109/TMI.2004.825600.

[50] C. Ordóñez-Sanz, M. Cowen, N. Shiravand, and N. D. MacDougall. "CBCT imaging: a simple approach for optimising and evaluating concomitant imaging doses, based on patient-specific attenuation, during radiotherapy pelvis treatment." *The British Journal of Radiology*, **94**(1124), Aug. 2021, p. 20210068. ISSN: 0007-1285. DOI: 10.1259/bjr.20210068.

[51] C. Lindsay, M. Bazalova-Carter, A. Wang, D. Shedlock, M. Wu, M. Newson, L. Xing, W. Ansbacher, R. Fahrig, and J. Star-Lack. "Investigation of combined kV/MV CBCT imaging with a high-DQE MV detector." *Medical Physics*, **46**(2), 2019, pp. 563–575. ISSN: 2473-4209. DOI: 10.1002/mp.13291.

[52] Y. Chen, B. Liu, J. M. O'Connor, C. S. Didier, and S. J. Glick. "Characterization of scatter in cone-beam CT breast imaging: Comparison of experimental measurements and

Monte Carlo simulation." *Medical Physics*, **36**(3), 2009, pp. 857–869. ISSN: 2473-4209. DOI: 10.1118/1.3077122.

[53] A. Malusek, M. P. Sandborg, and G. A. Carlsson. "Simulation of scatter in cone beam CT: effects on projection image quality." In: *Medical Imaging 2003: Physics of Medical Imaging*. Vol. 5030. SPIE. June 2003, pp. 740–751. DOI: 10.1117/12.479940. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/5030/0000/Simulation-of-scatter-in-cone-beam-CT--effects-on/10.1117/12.479940.full.

[54] S. Schafer, J. W. Stayman, W. Zbijewski, C. Schmidgunst, G. Kleinszig, and J. H. Siewerdsen. "Antiscatter grids in mobile C-arm cone-beam CT: Effect on image quality and dose." *Medical Physics*, **39**(1), Jan. 2012, pp. 153–159. ISSN: 0094-2405. DOI: 10.1118/1.3666947.

[55] *Beam Stop Phantom | A PTW Company*. https://www.qrm.de/en/products/beam-stop-phantom. URL: https://www.qrm.de/en/products/beam-stop-phantom.

[56] *Cone-Beam Phantom | A PTW Company*. https://www.qrm.de/en/cone-beam-phantom. URL: https://www.qrm.de/en/cone-beam-phantom.

[57] K. Okazaki, W. Cao, R. Taleei, F. Mourtada, J. Li, K. Mooney, P. R. Anne, and Y. Chen. "Evaluation of image quality in mobile conebeam CT with dose modulation using automatic exposure control: A phantom study." *Journal of Applied Clinical Medical Physics*, **26**(10), Sept. 2025, e70248. ISSN: 1526-9914. DOI: 10.1002/acm2.70248.

[58] Y. Li, Y. Jiang, H. Liu, X. Yu, S. Chen, D. Ma, J. Gao, and Y. Wu. "A phantom study comparing low-dose CT physical image quality from five different CT scanners." *Quantitative Imaging in Medicine and Surgery*, **12**(1), Jan. 2022, pp. 766–780. ISSN: 2223-4292. DOI: 10.21037/qims-21-245.

[59] S. Mutic, J. R. Palta, E. K. Butker, I. J. Das, M. S. Huq, L.-N. D. Loo, B. J. Salter, C. H. McCollough, and J. Van Dyk. "Quality assurance for computed-tomography simulators and the computed-tomography-simulation process: Report of the AAPM Radiation Therapy

Committee Task Group No. 66." *Medical Physics*, **30**(10), 2003, pp. 2762–2792. ISSN: 2473-4209. DOI: 10.1118/1.1609271.

Publishing Agreement

It is the policy of the University to encourage open access and broad distribution of all theses, dissertations, and manuscripts. The Graduate Division will facilitate the distribution of UCSF theses, dissertations, and manuscripts to the UCSF Library for open access and distribution. UCSF will make such theses, dissertations, and manuscripts accessible to the public and will take reasonable steps to preserve these works in perpetuity.

I hereby grant the non-exclusive, perpetual right to The Regents of the University of California to reproduce, publicly display, distribute, preserve, and publish copies of my thesis, dissertation, or manuscript in any form or media, now existing or later derived, including access online for teaching, research, and public service purposes.

Signed by:
Alan R Li
45D2392F718F4C0...
Author Signature

06/25/2026
Date