

THE UNIVERSITY OF CHICAGO

ENHANCED MEGAVOLTAGE IMAGING FOR RADIOTHERAPY BY LIGHT-FIELD
IMAGING OF SCINTILLATORS

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

COMMITTEE ON MEDICAL PHYSICS

BY

BENJAMIN MARTIN PREUSSER

CHICAGO, ILLINOIS

DECEMBER 2023

Copyright © 2023 by Benjamin Martin Preusser
All Rights Reserved

To my parents and sisters.

Comparison is the thief of joy. — Unknown

TABLE OF CONTENTS

LIST OF FIGURES	vii
LIST OF TABLES	x
ACKNOWLEDGMENTS	xi
ABSTRACT	xii
1 INTRODUCTION	1
1.1 Megavoltage Imaging	1
1.1.1 Detective Quantum Efficiency	3
1.1.2 The history of the electronic portal imaging detector	5
1.2 Light-field Imaging	10
1.3 Light-field Scintillator Imaging	15
2 FORWARD MODEL	17
2.1 TOPAS Monte Carlo	18
2.1.1 Monte Carlo Models	21
2.2 Optics	27
2.2.1 Camera Geometric Relations	27
2.2.2 Optical Effects	35
2.3 Reconstruction of volumetric scene	38
2.3.1 Limited Angle Tomography	39
2.3.2 Reconstruction Algorithms	42
2.4 Simulation Experiments	43
2.5 Results	44
2.6 Discussion & Conclusions	54
3 MATERIAL DECOMPOSITION & SPECTRAL ANALYSIS OF IMAGING WITH A LIGHT-FIELD CAMERA	56
3.1 Introduction	56
3.2 LKH-5's Material Decomposition Performance	57
3.2.1 Light-field imager axial resolution	58
3.2.2 Modeling the optical photon production in scintillator	59
3.2.3 Material Decomposition Results	64
3.2.4 Light-field Material Decomposition Results	66
3.3 Cramér-Rao Lower Bound Analysis	68
3.3.1 CRLB Results	72
3.4 Conclusion	74

4	LIGHT-FIELD IMAGER EXPERIMENTS	75
4.1	Equipment	75
4.2	DQE experiments	82
4.2.1	DQE procedure	82
4.2.2	DQE Results	85
4.3	Spectral Extraction Experiments	91
4.4	General Experiment Observations	94
4.5	Conclusion	97
5	SUMMARY & CONCLUSIONS	99
5.1	DQE Summary & Conclusions	99
5.2	Spectral Performance Summary & Conclusions	100
5.3	Future Directions	101
	REFERENCES	104

LIST OF FIGURES

1.1	A 3D printed phantom is imaged at increasing X-ray energy.	3
1.2	A cross-section of a camera-based EPID	6
1.3	The optical train of a conventional camera vs a standard plenoptic camera.	11
1.4	Angular information collected by a conventional camera vs angular information collected by a standard light-field camera. The yellow and green rays of light are focused onto two different pixels behind the MLA array.	12
1.5	Standard plenoptic camera data structure.	13
2.1	An illustration of the forward and back projection operations. ML refers to the main lens. MLA refers to the microlens array.	17
2.2	The X-ray photon fluences used in the TOPAS simulations.	22
2.3	Pulse-Height-Spectrum for LKH-5 at 6 MV. The total light collected represents the sum of all pixel values in a single gamma event image.	27
2.4	The optical train inside the Lytro Illum – a standard model light-field camera.	28
2.5	Point P_0 will be out of focus if moved off the z_0 plane. z'_1 refers to the out-of-focus image planes, where C indicates the resulting magnitude of blur in image space.	31
2.6	The transverse PSF displayed as a function of distance from the camera. The amount of transverse blur is minimum at the focal plane. The cameras have identical apertures and focal lengths. The middle images display the transverse blur axially throughout the focal stack. The edge images display the 2D transverse blur for the top and bottom focal stack slices. For the conventional camera, the top and bottom slices have different transverse blurs due to the asymmetry in the 3D PSF. In this example, there is no noticeable transverse plane blur for the Lytro Illum sub-aperture images.	34
2.7	Vignetting for a medium sensor	36
2.8	High-level ASTRA geometry for light-field cameras.	41
2.9	A 3D phantom modeled and 3D printed to compare the simulation and experiment for initial verification.	45
2.10	A view of a rectified microlens array image created by the Monte Carlo ray-tracing model	46
2.11	A view of a hexagonal microlens array image created by the Monte Carlo ray-tracing model	46
2.12	DQE for lens matched to ideal Lytro Illum	47
2.13	MTF for lens matched to ideal Lytro Illum	48
2.14	The normalized noise power spectrum for 6 MV	48
2.15	Example of gamma event images in the 2.2 LKH-5 scintillator	49
2.16	Approximated DQE for the experimental camera vs the Lytro Illum.	50
2.17	Lens matched to an F/0.95 lens.	52
2.18	MTF comparison between F/0.95 and F/2 LF cameras	52
2.19	DQE from one meter away for a LF camera.	53
2.20	MTF comparison between Monte Carlo ray-tracing and volumetric approach	54

3.1	PSF arrangement	58
3.2	Shift-variant PSF for light-field camera	59
3.3	LKH-5 optical response to irradiation at different energies. Images for 2.5 MV and 6 MV are exposure matched with 4 mm of copper buildup. The 225 kV image was taken at a longer exposure, hence its appearance is brighter.	61
3.4	Comparing the theoretical optical photon production with the LKH-5 scintillator at 6 MV to measured data. The continuous analytical model was used.	61
3.5	TOPAS material phantom. Two 5 cm of solid water and squares of 2 mm of Al, Au, and Cu.	62
3.6	TOPAS model of dose absorbed in the LKH-5 scintillator	63
3.7	The solid water material phantom’s contrast as different energies. The 2.5 MV and 6 MV images are window-level matched to better reflect the decrease in contrast at 6 MV compared to 2.5 MV. Materials are solid water, gold, copper, and aluminum.	65
3.8	Material removal using coarse axial resolution. The material removal for 225 kV, 2.5 MV, and 6 MV are the top, middle, and bottom rows, respectively.	67
3.9	Reconstructed focal stack slices of the material phantom at 225 kV, 2.5 MV, and 6 MV separated by row. Focus 3 is the focus at the top of the scintillator, then it goes focus 2, then 1 with increasing depths.	69
3.10	Axial response for the Lytro Illum vs F/0.95 light-field camera	70
3.11	The CRLB setup. A delta X-ray beam passes through soft tissue and bone and is detected by a light-field camera.	71
3.12	A sample of the viewpoint images used to compute the CRLB.	73
4.1	Cameras used during experiment	76
4.2	Raw image from Lytro Illum sensor. The MLA array found in the Lytro is hexagonally arranged and there is an RGB Bayer color filter above the sensor pixels. Image is of a 3D printed copper-doped phantom being imaged at 6 MV.	79
4.3	The experimental setup. The camera is placed inside a lead housing out of the X-ray beam. The X-ray beam is incident on the LKH-5 suspended above a front surface mirror, reflecting the optical photons to the camera. The light-field imager was placed in a lightbox for the LINAC experiments. The top of the scintillator was covered with black construction paper. On top of the black paper the phantoms would be placed.	80
4.4	Resolution at the focal plane. The ruler in the lower righthand of the images indicates a distance of 3 mm.	81
4.5	Tungsten Edge used for MTF calculations	83
4.6	NPS image processing	84
4.7	MTFs for the Lytro Illum and for the conventional camera	86
4.8	LKH-5 resolution differences based on the color channel on the Lytro Illum	87
4.9	2D NNPS for Lytro and CoolSNAP cameras. X and Y directional NNPS.	88
4.10	Noise banding from dark-frame subtraction analysis	89
4.11	Electronic noise.	89
4.12	6 MV DQE	90

4.13	Refocusing on Lytro Illum data with ML-EM	91
4.14	Axial signal dependence on frequency in the transverse plane. Bar pattern image is from a focal stack slice from the Lytro Illum at 225 kV.	92
4.15	Measured axial point spread function for a 5 μm point source. The 3D deconvolution result with the measured PSF is shown with the original bar pattern image.	93
4.16	Material axial responses imaged on the Lytro Illum. It is important to note that LKH-5 has a higher index of refraction than air, therefore light will appear 36.7% closer to the camera than it really is. These axial responses are not corrected for refraction, therefore the light appears to come from the middle of the scintillator rather than the top.	94
4.17	Potential thermal variance dependence for Lytro and conventional camera	96
4.18	LKH-5 scintillation response to one minute of exposure of 6 MV at LINAC max dose rate.	97
5.1	Preliminary proton beam simulation and reconstruction using the developed forward model.	102

LIST OF TABLES

2.1	EJ-260 Specifications	20
2.2	LKH-5 Specifications	20
2.3	Camera Parameters at $Z = 20$ cm	44
2.4	Camera Parameters at $Z = 1$ m	45
3.1	Material removal signal to background ratio at 2.5 MV	66
3.2	Material removal signal to background ratio at 6 MV	66
3.3	CRLB figures of merit $A/\sqrt{\sigma^2}$ for Lytro Illum (F/2).	73
3.4	CRLB figures of merit $A/\sqrt{\sigma^2}$ for an F/0.95 light-field camera.	74
4.1	Lytro Illum Specifications	77
4.2	CoolSNAP Camera Specifications	77

ACKNOWLEDGMENTS

Many people have played a role in this project. I want to thank:

- PATRICK LA RIVIÈRE for the continuous support and patience during my PhD study. Your breadth of knowledge is inspirational and gives me something to strive for in my future endeavors.
- JORDAN SLAGOWSKI for serving on my committee, for your enthusiasm, and the many hours spent aiding me in the clinic for data collection.
- CHIEN-MIN KAO & RODNEY WIERSMA for serving on my committee and giving valuable suggestions.
- ERIK PEARSON for your time helping me get set up on the XRAD and the many hours you spent with me during the initial light-field experiments.
- BAIYANG DAI for your friendship, support, and the many adventures we had in Chicagoland.
- PHIL VARGAS & COREY SMITH & NIKOLAJ REISER for the helpful suggestions, Python support, and office sharing. I enjoyed the many conversations we have had over the years.
- TALON CHANDLER for the early guidance and insight on light-field imaging and the world of optics. Your enthusiasm for the sciences is contagious.
- LUCAS BERENS, DYMARKCO DAVIS & CHRISTOPHER VALDES for rotating with me and providing helpful suggestions and support to this project.

ABSTRACT

Megavoltage imaging during radiotherapy provides a beam’s eye view of patient anatomy, allowing for portal dosimetry and real-time tumor position tracking without additional dose. However, current electronic portal imaging devices (EPIDs) suffer poor contrast-to-noise ratios because of their limited (1-2%) detective quantum efficiency (DQE). Conventional EPIDs use a thin layer (<1mm) of scintillator to convert x-rays to visible light that is then detected by an array of photodetectors. The scintillator must be thin in order to reduce blurring caused by light spread. In addition, conventional EPIDs do not provide any form of spectral decomposition, and the megavoltage detectors that do are expensive. Light-field imaging offers a way to image a volumetric light source like a scintillator by recording the emitted 4D light field and computationally refocusing from a single exposure. Computational reconstruction allows for out-of-focus light to be removed, allowing for an order of magnitude thicker scintillators to be imaged while preserving resolution and allowing for spectral extraction. A forward model was developed to simulate light-field cameras and scintillators to explore the potential of light-field cameras in scintillator imaging. With the forward model, DQEs and spectral performance were evaluated. A prosumer light-field camera was used experimentally and evaluated in terms of DQE and spectral performance to evaluate the simulation model. In addition, a conventional scientific camera was used experimentally. The simulations revealed that light-field cameras provide an advantage over conventional cameras for especially thick scintillators; however, the noise from reconstruction is the limiting factor preventing the light-field camera from being competitive with multi-layered EPIDs for thinner scintillators. In addition, the spectral performance of light-field cameras for X-ray-based imaging was marginal at 225 kV and negligible in megavoltage images. The scientific camera results were validated by simulation, and the prosumer light-field camera was shown to be unsuitable for scientific studies.

CHAPTER 1

INTRODUCTION

Megavoltage imaging is a technique used in radiotherapy cancer treatment to ensure patient alignment and track the tumor field. Megavoltage imagers can provide an on-demand beam's eye view of the tumor during radiation treatment that would otherwise be difficult to do with the onboard kilovoltage imagers, which are orthogonal to the treatment beam. Current megavoltage imagers suffer from low contrast and low DQE. Improving megavoltage detectors has been an active area of research due to the potential payoffs for patient imaging. Thanks to new developments in scintillators, camera optics, and computationally efficient algorithms, progress has been made on a number of fronts.

Light-field imaging of a thick monolithic scintillator is a novel technique that uses the scintillator to indirectly convert megavoltage X-rays into optical photons that can then be imaged by a light-field camera. Depending on the thickness of the scintillator and the light-field camera being used this could improve the low quantum efficiency of megavoltage detectors while offering higher resolution.

This chapter begins with a brief explanation of megavoltage imaging and why it is difficult to get a high-quality image. It provides a literature review of creating more efficient megavoltage detectors. We then discuss the development of light-field imaging and review the literature on current light-field attempts at volumetric imaging. Finally, the chapter concludes with an overview of our research on light-field imaging of scintillators.

1.1 Megavoltage Imaging

Accurate patient placement for external radiation treatment is paramount for accurate radiation delivery. As a result, clinicians use a variety of methods to align the patient to the treatment field, including lasers located throughout the room to provide an isocenter and an

optical surface monitoring system (OSMS), which tracks patient movement in real-time and is often used in lieu of lasers. In addition, a variety of imagers are used to assist with alignment: kilovoltage CT imager, magnetic resonance imaging (MRI), and the electronic portal imaging device (EPID) [1]. The American Association of Physicists in Medicine (AAPM) Task Group 142 recommends that the laser alignment to isocenter should be within ± 2 mm or even ± 1 mm depending on the radiation application [2]. However, despite the best efforts to align a patient, humans are unpredictable and susceptible to slight movements that can occur post-alignment and even during the treatment itself. Therefore it would be ideal to be able to track patient movement in real-time, not just the patient’s body at large, but even internal movement as well such as monitoring the heart and lungs and their movement cycles. Tumor tracking with the treatment beam itself is desired as it would provide a beam’s eye view of the tumor and would be able to tell when the tumor leaves the treatment field. In addition, it would spare the patient additional imaging dose that is delivered by tracking the patient with the kilovoltage imager system that images perpendicular to the treatment beam [3].

In diagnostic X-ray energies, images taken with the standard 30-140 kV range have significantly better contrast and material separability due to materials having greater differential absorption from the photoelectric effect than compared to megavoltage imaging. Figure 1.1 showcases an example of a 3D-printed copper-based filament being imaged at increasing X-ray energies and the associated loss of contrast.

The photoelectric effect has a cross-sectional area that is inversely proportional to the energy cubed ($1/E^3$) [4]. Defining contrast C as

$$C(\%) = \frac{I_1 - I_2}{(I_1 + I_2)/2} \cdot 100\%, \quad (1.1)$$

where I_1 is the intensity measurement of the object of interest and I_2 is the intensity measurement of the background [5]. For 1 cm of bone in a 20 cm soft tissue phantom, the




Figure 1.1: A 3D printed phantom is imaged at increasing X-ray energy.

contrast decreases from 50% to 4% from diagnostic to megavoltage imaging. Similarly, 1 cm of lung goes from 17% to 4% with the change to megavoltage imaging.

Being able to detect an object out of the background in an image depends on the imager's resolution and noise properties. The concept of having a metric to define how well a system can convert an incident signal to a useful image signal is called detective quantum efficiency. It is commonly used to describe and compare different megavoltage imagers.

1.1.1 *Detective Quantum Efficiency*

When consumers buy a new cellphone, they want to know how well the camera compares to other phone cameras. They may look at metrics like the total number of megapixels, the low-light performance, and the color science the phone uses to colorize the final images. Similarly, when it comes to X-ray imaging systems, doctors and researchers want to know how efficient the imager is regarding X-ray dose, signal-to-noise ratio (SNR), resolution, and field-of-view of the imager. A commonly used metric to define how efficient an X-ray imager is in converting the input SNR_{in} to an output image signal SNR_{out} is known as the DQE [6]

$$DQE(f) = \frac{SNR_{out}^2}{SNR_{in}^2}. \quad (1.2)$$

The f represents frequency in terms of lines per mm (lp/mm). An ideal imaging detector will have a DQE of 1, meaning all incoming X-ray photons will be captured and contribute to the final image. Another DQE formalism that is commonly used to relate the resolution of the imaging system in terms of the modulation transfer function (MTF), the normalized noise power spectrum (NNPS), and the incident photon fluence q by

$$DQE(f) = \frac{|MTF(f)|^2}{q \cdot nNPS(f)}. \quad (1.3)$$

The MTF represents the transmission of spatial frequencies through the imaging system. For both X-ray and optical systems, the MTF can be determined by imaging a point source and summing the point source in either the X-Y directions to generate point-spread functions denoted PSF_X and PSF_Y for the X and Y axes. The directional MTF can be calculated by taking the magnitude of the discrete Fourier transform of the PSF,

$$MTF(f) = |\mathbf{FT}[PSF(x)]|. \quad (1.4)$$

For comparison between different systems, the MTF is normalized to 1 on the zeroth frequency. For X-ray systems, the X-ray source size, X-ray source energy, and the pixel resolution of the detector all affect the maximum resolution achievable by the X-ray imaging system [7]. Depending if the DQE is being calculated by experiment or by simulation, the formalism of NNPS changes. Regardless the NNPS captures the noise of the system over spatial frequencies by taking the Fourier transform of the noise autocorrelation. In an experiment, the noise autocorrelation will be calculated by taking flood-field images and analyzing them, and in simulation by looking at the noise from individual gamma photon events to save on computation time [6, 8].

Clearly, high resolution and low noise are ideal for maximizing the DQE. At lower energies, DQE can be in the 50-70% range [9], due to the X-rays being easier to attenuate.

Whereas the detectors used in megavoltage imaging typically have a DQE of 1% with some experimental detectors having a higher DQE but compromises are made on resolution to achieve the higher DQE [10].

1.1.2 The history of the electronic portal imaging detector

The first portal imaging devices used to assess patient setup for radiotherapy in the mid-1970s relied on films. The structure of the film-based EPID contained a metal buildup plate, the film, and a backing material to hold the film. The metal buildup plate was placed in front of the film to increase the detector's quantum efficiency. The metals used in buildup plates are generally copper, steel, or brass. The buildup layer serves as a way to generate electrons which will then interact with the film. In addition, the layer of metal absorbs low-energy X-rays that would have otherwise interacted with the film blurring the image. The backing material for the film was sometimes a metal plate as well, which can provide back-scattered electrons. In other cases, glass or plastic is used to provide general structural support.

The use of film portal imagers quickly showed the importance of performing routine quality assessments of patient alignment; however, given the manual nature of using films: setup, replacing, and developing, there was a limit to how many films could be used. Therefore, quality assurance at the time was done on a weekly basis with a margin of error of 10 mm [5]. As treatment precision improved, so did the need for accurate setups. As a result, over the years, the EPID has undergone several advancements.

The first-generation EPID technology can be divided into camera-based EPIDs, photo-stimulable phosphors, ionization chamber matrix devices, scanning arrays, and other systems [11]. Of the first-generation EPIDS, only the camera-based and ionization chamber matrix designs became commercially available. While each technology has advantages and disadvantages, the EPID designs we are most interested in competing with is the camera-based EPIDs and the second-generation EPIDs based on active matrix flat-panel technology, which

serves as the basis for EPIDs available commercially today [11].

Camera-based EPIDs

Camera-based EPIDs, which followed the film-based portal imagers, had a similar design structure as shown in Figure 1.2, comprising metal buildup plate of, traditionally, 1 mm of copper and thin scintillator, usually gadolinium oxysulphide, with thicknesses ranging from 100-400 mg/cm⁻¹. Unlike the film method, where the image was formed from direct interactions with the gamma photons and electrons, the camera-based EPIDs used indirect detection by scintillation, which converts the energy deposited from gammas and electrons into optical photon. As shown in Figure 1.2, there is no backplate as the optical photons need to exit the scintillator to reach the optical camera. The number of scintillation photons that are produced from an event can be calculated by




Figure 1.2: A cross-section of a camera-based EPID

$$N = E_s \left(\frac{\eta_{scint}}{E_{scint}} \right) . \quad (1.5)$$

Where E_s is the X-ray source energy, η_{scint} is the maximum energy conversion efficiency of the scintillator, and E_{scint} is the energy of the fluoresced optical photons [12]. Therefore, the better the energy conversion of a scintillator, the more optical photons will be produced. While equation 1.5 is written in monoenergetic formalism, most scintillators and X-ray sources operate in polychromatic fashion.

The brightness of a scintillator as perceived by a camera depends on the number of photons hitting the sensor. While camera optics will be more thoroughly discussed in the forward model of chapter 2, the two largest factors determining the number of photons that a lens will collect are the distance the camera is away from the light source and the diameter of the main lens. A lens' light collection efficiency, also known as optical coupling (OC) efficiency to a light source, can be expressed as a percentage by

$$g_{oc} = \frac{T}{4(F/\#)^2(1+m)^2 + 1} , \quad (1.6)$$

where T represents the lens' transmission factor for light (70-80%) for a multi-piece photographic lens, m the demagnification ratio, and $F/\#$ representing the F-number of the main lens, which is defined as the ratio of the lens' focal length over the lens' diameter [12]. F-numbers for photographic lenses generally range from F/0.95 up to F/32 or higher. While custom-made lenses can have a specific F-number, generally, consumer-based lenses list F-numbers based on the power of $\sqrt{2}^p$, therefore consumer lenses go by an F-number scale of (1,1.41,2,2.8,...). Comparing the light collection efficiency of two sequential lenses F/2 and F/2.8 set to the same focal length, the diameter decreases by a factor of $1/\sqrt{2}$ from F/2 to F/2.8, thus reducing the flux by half. Taking a picture with an F/2 lens with a shutter speed of 1/100th of a second is the same as taking a picture with an F/2.8 lens at 1/50th of

a second [13]. The demagnification term m is defined as the object size to image size. Macro lenses that allow for near-camera focusing may produce one-to-one magnification. However, most photographic lenses have an object-to-image ratio greater than 4. For example, a lens may have a magnification ratio of 1:10 reported on its barrel, which means the image of the object will be a tenth of its true real-world size.

Analyzing equation 1.6, a camera needs to be as close as possible to make the magnification close to 1 with the largest F-number possible to collect the most light to maximize the lens collection efficiency. However, imaging both close and at wide apertures leads to images with the region in focus being only a few microns wide, therefore a lot of macro photography images tend to be dark due to increasing the F-number to have a large enough depth-of-field to capture the entire subject. This leads to a competing trade-off of aperture vs distance [14]. To combat this tradeoff, many macro photographers have turned to focal stacking. Focal stacking is the process of taking multiple images in rapid sequence. Each image has a slightly different focus. The images are then stored in a 3D array, where an algorithm can go pixel by pixel or by local neighborhoods and compare which pixel axially in the stack should be used in the final macro image to create an image with an extended depth-of-field (EDOF). While focal stacking works well in many situations, most focal stacking algorithms rely on the assumptions 1) that in-focus features have their highest contrast when they are in focus and 2) that the region is brightest when it is in focus compared to when it is out of focus [15]. Focal stacking algorithms for microscopy have since moved on from those assumptions and now use processes like complex wavelet techniques, which assume that in-focus components will have higher spatial frequency features than when they are out-of-focus and the complex wavelet techniques allow for the localization of frequencies in the spatial domain in order to compute the extended depth-of-field image [16]. Despite these advancements in computationally expanding the depth-of-field in microscopy and macro photography, when it comes to creating extended depth-of-field images for transparent scenes that have little frequency

variation axially in the focal stack or in the transverse planes, many of these algorithms will return noisy pixels as the most in-focus pixel. This is problematic when imaging something like a scintillator that can be transparent. Also as scintillators grow thicker the resolution suffers due to light-spread within the scintillator [17]. As a result, the scintillator thickness used in the camera-based EPIDs has been limited because it must match the camera's depth-of-field thickness and it must minimize the scintillator's light spread.

A few camera-based EPIDs were commercially available such as Siemens "Beamview", Infimed's "Theraview", and Philips/Elekta "SRI-100/iView" [11]. The major draw of these EPIDs was the real-time performance. The Theraview system employed a cooled CCD sensor that had a readout of 4 frames per second [5]. The field-of-view for the Theraview was 28 by 28 cm at ISO center, and the pixels could be binned to 512 by 512 for a fast readout, or for high-resolution photos a max resolution of 1024x1024 pixels.

Despite having a decent resolution, the camera-based EPID was bulky to work around in patient setup, and the DQE of this generation of camera-based EPIDs was less than 1%. The camera-based EPID have numerous inefficiencies. Depending on the scintillator only a fraction of the light escapes the scintillator (10-20%) and from there depending on the aperture size only a fraction of that light enters the camera. The camera readout introduces additional noise, further increasing the system's NPS and lowering the DQE by several percent [18].

Recently, camera-based scintillator imagers have been making a comeback. Detecting the Bragg peak of a proton beam can be tricky with a thin scintillator if the beam was not attenuated enough. Therefore, researchers have been developing large liquid scintillator camera setups to capture the Bragg peak and assess its depth in the scintillator [19]. Using large scintillators also has an application in traditional photon/electron machine QA, where a large plastic scintillator can be imaged by multiple cameras to perform 3D dose assessment [20]. In addition, camera technology has improved. Larger format sensors with more pixels

and computational photography being the two largest improvements.

Active matrix flat-panel EPIDs

The second generation and most prevalent EPID technology is based on active matrix flat panel imagers (AMFPIs) [11]. The construction of AMFPIs is similar to the previous EPID designs: X-ray converter, then detection device. The X-rays are converted either by the in-direct detection (metal/phosphor) or by direct detection (metal/photoconductor). In the photoconductor, a voltage is generated by electron-hole pairs, which is later read out. The novel component is the readout device, the active matrix. The active matrix is similar to a camera sensor. Both feature pixels arranged in a 2D grid. In AMFPIs, during X-ray irradiation, each pixel accumulates a charge, which is later read out row by row and digitally processed to form an image.

Benefits of AMFPI EPIDs include their speed, field-of-view, and bit-depth. AMFPI EPIDs can have up to a 41 by 41 cm field-of-view with 1024x1024 pixels. They offer up to 16-bit depth allowing for high dynamic range in an image. However, there are still drawbacks to AMFPI. Their DQE is quantum limited due to their thin X-ray conversion layer, meaning the only way to increase the quantum efficiency is by increasing the thickness of the X-ray converter layer. If the X-ray converter is a scintillator, increasing the scintillator thickness there is a sharp falloff in resolution due to the isotropic light spread [17]. Additional problems like image lag/ghosting and high maintenance and replacement costs also plague AMFPI EPIDs. Despite these drawbacks, AMFPI EPIDs are the most commercially used EPID design at the time of writing.

1.2 Light-field Imaging

A conventional photographic camera records the integral light intensity distribution captured by the camera's main lens. A plenoptic camera is similar to conventional cameras as it also

records the integral light intensity distribution, however, it introduces a microlens array (MLA) between the camera’s sensor and the main lens, which preserves the direction of the incoming light. The recording of both the light-intensity distribution and its direction is referred to as recording the light-field as a light ray can be parameterized as a 4D vector based on the coordinates of intersection in the main lens and sensor planes [21].

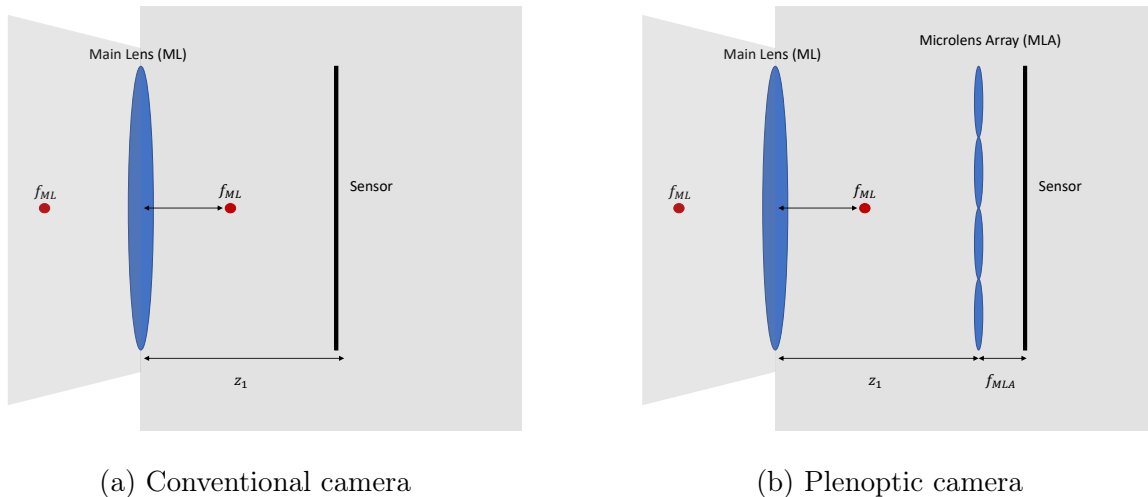


Figure 1.3: The optical train of a conventional camera vs a standard plenoptic camera.

In the past decade, two types of plenoptic cameras have become available for consumers to purchase: unfocused plenoptic cameras, also known as the standard plenoptic camera and focused plenoptic cameras. The placement of the microlens array in relation to the image sensor and the main lens largely identifies the type of plenoptic camera. In a standard plenoptic camera, the MLA is placed at the back focal plane of the main lens, and the image sensor is placed at the back focal plane of the MLA. In a focused light-field camera an intermediate image is formed either in the front or behind the microlens array [22]. Each design has pros and cons. The standard plenoptic model maximizes the angular information as each pixel behind a microlens is recording a different angle. If the same pixel is picked behind each microlens, a perspective image can be formed. Figure 1.4 shows how the angular information is recorded to different pixels behind a microlens.

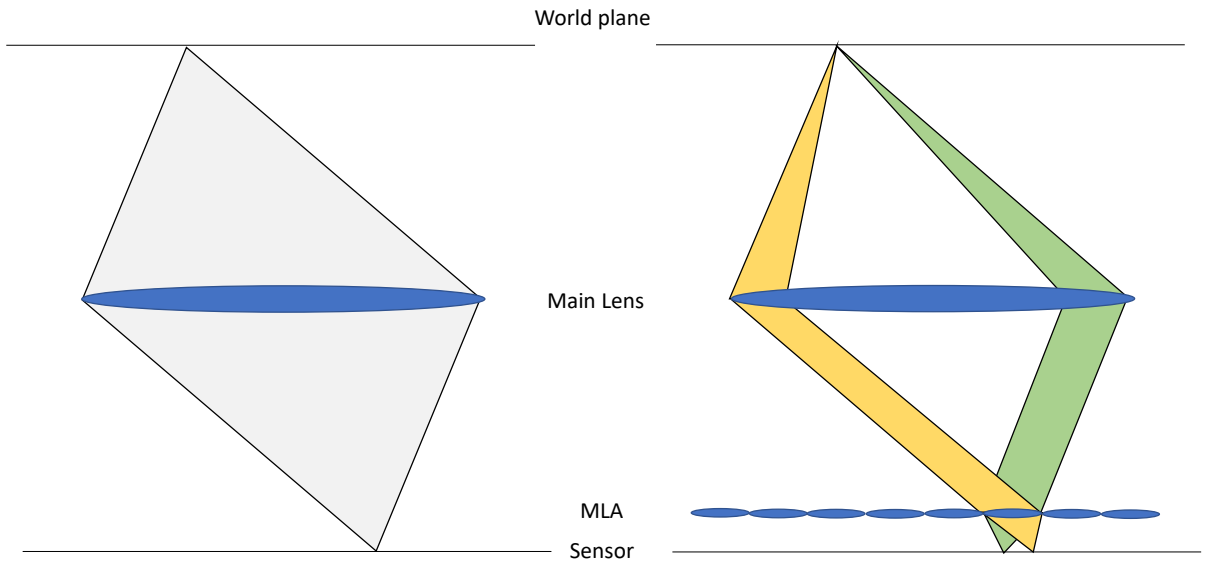


Figure 1.4: Angular information collected by a conventional camera vs angular information collected by a standard light-field camera. The yellow and green rays of light are focused onto two different pixels behind the MLA array.


Figure 1.5 shows example data collected from a standard plenoptic camera in practice. In Figure 1.5, approximately two hundred thousand microlenses are in front of the camera sensor, each $20 \mu\text{m}$ in diameter. Two example viewpoint images, also known as sub-aperture images as the width of the main lens is the limiting factor in how much parallax can be observed for an image scene [21].

The major drawback to the standard plenoptic camera is the spatial resolution is limited to the number of microlenses; even two hundred thousand microlenses result in a final image of only 540×375 pixels. The focused plenoptic camera trades some of the angular information for increased spatial resolution, however, that increases computation, and given our imaging application, the angular information is more novel, hence we limited our focus to standard plenoptic cameras.


The novelty of a light-field camera is its refocusing ability using both the recorded light



(a) 40 megapixel microlens array image



(b) Zoomed section of the MLA image.



(c) Viewpoint 1



(d) Viewpoint 2

Figure 1.5: Standard plenoptic camera data structure.

intensity and angular information. There are several ways to refocus an image computation-ally, however, they can be broken down into analytic and iterative. Analytical refocusing include methods like Ren Ng's integration approach, expressed as

$$E_{\alpha f}(x', y') = \frac{1}{\alpha^2 f^2} \int \int L_f^{u,v}(u(1 - 1/\alpha) + x'/\alpha, v(1 - 1/\alpha) + y'/\alpha) du dv, \quad (1.7)$$

where $E_{\alpha f}(x', y')$ represents a refocused image, x, y refers to a specific microlens on the MLA, u, v refers to a point on the main lens plane, $L_f^{u,v}$ is the sub-aperture image which is dilated by a factor of α and shifted by $(u(1 - 1/\alpha) + x'/\alpha, v(1 - 1/\alpha) + y'/\alpha)$. f corresponds to the focal length and α represents the plane of the refocused image relative to the nominal

focus plane [21]. Essentially by the integration method a refocused image can be constructed by adding and shifting the sub-aperture images. The number of sub-aperture images used in the refocusing process directly relates to the depth-of-field (DoF) of the refocused image. Refocused images that use only a few sub-aperture images will give an image with a much wider DoF than when all the sub-aperture images are used because adding more sub-aperture images in the integration process increases the angular information, which increases blur. This flexibility of sub-aperture image usage allows a light-field camera to simulate what a photographed scene would look like if it were imaged by conventional cameras with different f/numbers.

Iterative refocusing treats the refocusing problem as a limited angle tomography problem or as a 3D deconvolution problem, [23, 24]. In the iterative refocusing methods, attempts are made to remove out-of-focus light. In general, refocusing a camera does not remove the out-of-focus light; similarly refocusing with the integration method does not remove the out-of-focus light. The issue with iterative methods is the camera has to be modeled well enough to model the imager’s system geometry transfer matrices. Both 3D deconvolution and iterative tomographic reconstruction require system matrices. 3D deconvolutions are difficult to do with a light-field camera due to a light-field camera being shift-variant both in the transverse and axial planes [24]. Modeling the full volume space to the sensor image is also quite difficult as storing that matrix can take hundreds of gigabytes, as for example, the Lytro Illum has a 40 million pixelated sensor array if we wish to reconstruct a 100x100x50 voxelized scene that would create a system matrix with 2^{13} elements. As a result, many of the iterative approaches use sparse matrices to reduce computation significantly and in many cases, only parts of the system matrix are computed on demand.

Light-field imaging has been emerging in different areas of manufacturing. Light-field cameras have been used to image gas injectors to flames [24]. In terms of light-field camera availability, the Lytro company was a pioneer in making light-field cameras inexpensive to

prosumer photography enthusiasts with their first generation light-field camera and their 2nd generation light-field camera referred to as the Lytro Illum which was released in 2014 [25]. Since then Lytro has been dissolved. However, light-field cameras are still available from Raytrix, which develops light-field cameras for scientific applications [26].

1.3 Light-field Scintillator Imaging

The application of applying a light-field camera to image a thick monolithic scintillator is very novel. Archambault *et al.* was the first group to publish the idea of using a light-field camera to perform 3D dose verification for radiotherapy [20, 27]. Their application was solely quality assessment based, which leaves a large application for light-field cameras for active patient imaging unexplored. In addition, their research always utilized multiple views to reconstruct the 3D volume, either by using a mirror to get multiple views to the camera or by using the LINAC's EPID to provide the view orthogonal to the camera's. Their results were encouraging for being able to reconstruct the 3D volume with multiple camera projections, however, it does not indicate how efficient their light-field imager system is in terms of DQE.

The light-field scintillator imager we present here is entirely novel. We use a glass-based scintillator with a noticeably higher density than the scintillator used in previous light-field research. In addition, we explore the potential of performing material decomposition with a light-field camera. To explore the material decomposition potential we develop an open-source forward model that extends from the Monte Carlo scintillation simulation to the computational volume reconstruction. We explore theoretically how efficient in terms of DQE a light-field camera imager can be and whether is it competitive with current approaches.

Two major approaches exist to increase the DQE of EPID detectors: columnated scintillators and multi-layered detectors. The first approach is to use thicker scintillators; however, instead of a thick monolithic scintillator, the scintillator is broken into thousands of long thin

columns, which are then optically shielded from each other and glued together to make a large scintillator [28]. The DQE(0) of the piecewise-focused detector developed by Star-Lack *et al.* is reported to be around 24%, which is significantly higher than traditional EPIDs. However, the construction of a columnated scintillator is expensive and the column's width limits the resolution.

The multi-layered detector approach uses multiple thinner scintillators sandwiched with a copper buildup and photodetector for each scintillator. This approach is also expensive and introduces electronics that can lead to scatter into the lower layers. The prototype developed by Rottmann *et al.* used four layers of scintillator, which increased the DQE(0) to 6.7% [29]. However, since the first prototype, they have further optimized the layers. In more recent prototypes, they have a thin scintillator layer for resolution, and a thick scintillator layer for detection [30]. In the optimized for DQE multi-layered detector at 6 MV they had a DQE(0) of 7% to 15% where the better DQE was achieved by using a significantly thicker, but columnated scintillator for the detection layer [30]. In addition, using a multi-layered detector they can perform a spectral evaluation of the incident X-rays due to lower energy X-rays being preferentially attenuated in the first layers. This makes this EPID model the greatest competition to our model where we propose to image a thick monolithic scintillator to increase the DQE but also to use a light-field camera to refocus and potentially perform material decomposition.

CHAPTER 2

FORWARD MODEL

Performing experiments can be quite expensive, especially when the outcome of the experiment is uncertain. In the case of medical physics, scintillators, cameras, and access to an X-ray source can cost tens of thousands of dollars, especially if the hardware is customized for a particular application. An alternative approach to doing physical experiments is to perform simulations. In our case, we did perform several experiments with a light-field camera which are presented in chapter 4. However, the camera we used was a prosumer-level camera that has several significant limitations on the quality of the data that it could capture. To gain a better understanding of the promise and limitations of light-field scintillator imaging, we develop here a forward model that is used to validate our experimental results and also to simulate more optimal light-field cameras.

The order of this chapter follows the process of simulating data in Monte Carlo to projecting it through the camera, then reconstructing the volume through iterative methods. Figure 2.1 visually illustrates the forward and back projection operations and the substeps in the operations.

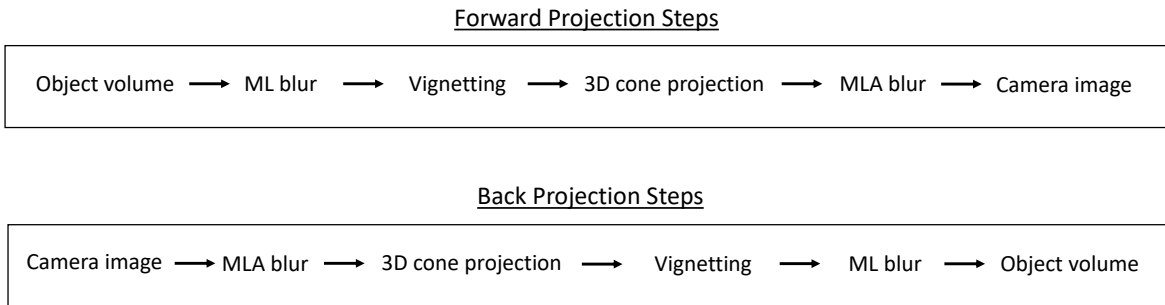


Figure 2.1: An illustration of the forward and back projection operations. ML refers to the main lens. MLA refers to the microlens array.

2.1 TOPAS Monte Carlo

A Monte Carlo simulation is a powerful tool used to predict the probability of an outcome when random variables are present. While Monte Carlo had its beginning in the early 18th century, it was not until the 19th and 20th centuries that simulations took off as a way to verify experiments, analyze data, and provide insight into mathematical statistics.

Early applications of Monte Carlo were used to study what can be considered traditional deterministic problems. For an example of a deterministic problem, Comte de Buffon (1707-1788) would drop needles onto an image with horizontal lines. He would count the number of times X the needles intersected the lines and he would throw the needles N times. Doing this, he was able to show that he could estimate π by using $\pi = \frac{2N}{X}$. While π isn't random, it was one of the first instances where randomness was used to calculate a deterministic parameter [31, 32].

Nowadays, Monte Carlo can still be used to calculate deterministic parameters, however, it is largely used to simulate events and parameters that are entirely probabilistic. In the case of our study, X-ray interactions with matter are a probabilistic event. There are several popular Monte Carlo frameworks used in medical physics: Geant4, which was developed in the 90s for high energy space applications, EGSnrc for modeling photons and electrons, and TOPAS which was initially developed for proton research based on Geant4. However, TOPAS has since been expanded to include X-rays and optical photons and various other medical physics-related applications [33, 34, 35]. As a result of TOPAS being designed for medical physics applications and for medical physicists as the intended users, we chose to use that interface for our Monte Carlo simulations. In addition, TOPAS is open-source and free to use if you are a researcher and/or a student.

Monte Carlo simulations can be computationally expensive depending on the scale and the amount of parallelization acceleration the user wants to do. Historically the central processing units (CPU) of the computer were responsible for performing the calculations.

However, recently graphical processing units (GPU) have made a comeback in the scientific community due to their role in deep learning applications. While a CPU on a personal computer may only have 8 CPU cores compared to the hundreds of thousands of processing cores found on a GPU, CPU cores tend to operate a factor of 10 or faster than a GPU core [36]. Therefore, where GPUs tend to excel is in highly parallelizable problems like matrix element-wise multiplication where each core can compute a part of it. Unfortunately, TOPAS being based on the older Geant4 Monte Carlo framework is not parallelizable for GPU, forcing the user to rely on multiple CPUs to complete a simulation efficiently. Therefore our Monte Carlo simulations along with our iterative reconstructions were performed on a scientific computing server that features 512 Intel Xeon Gold 6130 processor cores, and 16 Nvidia Tesla V100 32 GB GPUs. TOPAS simulations were done using 300 of the 512 CPUs in parallel and only one GPU was required for the iterative reconstructions.

Our primary goal in using Monte Carlo is to simulate a thick transparent scintillator being struck by X-rays or, potentially, protons. As discussed in Chapter 1, our scintillator needs to be transparent to its own optical fluorescence, and be thick enough to be considered novel compared to the more traditional EPID scintillators.

In our simulations, we simulate an LKH-5 scintillator and an EJ-260 scintillator. We chose to simulate the LKH-5 scintillator as we have a small sample of it that we used in our experiments presented in Chapter 4. The LKH-5 scintillator parameters are shown in Table 2.2. While there are some disagreements in the literature about the true LKH-5 parameters due to LKH-5 being a proprietary experimental scintillator developed by Collimated Holes Inc (Campbell, CA, USA) all reported values are roughly the same [30, 37, 38]. The other scintillator EJ-260, produced by Eljen Technology (Sweetwater, TX, USA) we did not perform any physical experiments with. However, the EJ-260 scintillator has been previously used in light-field scintillator research and more recently in 3D proton beam QA. Therefore it is of interest for its potential use as a large 3D volumetric scintillator [20, 39]. The

Table 2.1: EJ-260 Specifications

EJ-260 Scintillator Specifications	
Composition	H(0.085), C(0.915)
Scintillator thickness	2.3, 5 cm
Peak λ	490 nm
Density	1.023 g/cm ³
Refractive Index (λ_{490})	1.58
Scintillation Yield (ph/MV)	9200
Absorption Length (λ_{545})	350 cm

Table 2.2: LKH-5 Specifications

LKH-5 Scintillator Specifications	
Composition	O(0.3), Ba(0.4), Si(0.2), Tb(0.1)
Scintillator thickness	2.3, 5 cm
Peak λ	545 nm
Density	3.83 g/cm ³
Refractive Index (λ_{545})	1.58
Scintillation Yield (ph/MV)	4000
Absorption Length (λ_{545})	160 mm

specifications for EJ-260 are shown in Table 2.1 [40].

The most notable differences between LKH-5 and the EJ-260 scintillators are the density, optical clarity, optical output, and emitted wavelength. LKH-5 is a glass-based scintillator with a density of 3.83 g/cm³, whereas EJ-260 is plastic-based scintillator with a density of 1.023 g/cm³. In literature, other groups that have used a thick (> 1 cm) piece of LKH-5 scintillator in experiments often left off a copper conversion layer due to its density. In our experiments, we also observed at megavoltage energies that increasing the copper thickness leads to a lower resolution, which will be shown in Chapter 4 [38]. On its own, LKH-5 has an effective atomic number of $Z = 35$, which corresponds to bromine.

Regarding optical clarity, the EJ-260 permits optical photons to travel an order of magnitude further than the LKH-5 at 350 cm and 160 mm absorption values, respectively. The optical absorption of a scintillator is defined as the distance where the beam reaches a $1/e$ intensity. In addition to having a larger absorption length, EJ-260 has twice the optical

output of 9000 photons/MV versus LKH-5's 4000 photons/MV. Others have also noted how dim LKH-5 is compared to the more commonly used GOS scintillation material used in EPIDs, which can have optical yields of 60k photons/MV [41].

2.1.1 Monte Carlo Models

Our scintillator Monte Carlo research employs two approaches: a volumetric dose approach and an optical ray tracing approach. A research group from the Centre de Recherche at Chude Quebec, the only other group that we know has researched light-field scintillator imager applications for 3D dose verification, generated their simulated data by using treatment planning software to simulate the beam passing through the large scintillator. They assumed that the number of optical photons emitted by a voxel would be proportional to the dose absorbed. From there, rays were cast from each voxel to the camera using commercial optical software Zemax [20, 27, 42, 43]. This approach does not capture specular and diffuse scatter from scintillator coating or nearby boundary objects. In addition, it does not capture any internal reflections. While Zemax can simulate some of these effects, it is behind a commercial license.

An alternate approach entails simulating the scintillation process within the scintillator and then ray-tracing individual photons to the camera sensor. This captures the optical effects that would have been missed using the volumetric approach. Additional drawbacks to this ray-tracing approach caused us to primarily focus on the volumetric approach for our material decomposition tasks; however, we implemented the ray-tracing approach to simulate the DQE for various configurations.

The incident X-ray fluences used in the Monte Carlo simulations are shown in Figure 2.2. The 225 kV spectrum for the small animal irradiator was previously computed by Monte Carlo by colleagues at the University of Chicago. Colleagues also provided the 2.5 MV and 6 MV spectrums for the LINACs used in the clinic at the University of Chicago.




Figure 2.2: The X-ray photon fluences used in the TOPAS simulations.

Monte Carlo Volumetric Model

In our volumetric dose calculation approach, the scintillator was broken down into a number of voxels depending on the application and the base resolution of the camera. The primary role of the volumetric models was to determine the spectral component of scintillator imaging with a light-field camera, largely because the scintillator could be broken into larger voxels in the transverse plane, which reduces the computational memory requirement. In the volumetric implementation, TOPAS does not continuously compute and write the calculated dose to an output file, meaning the simulation is lost if a simulation is interrupted. As a way to combat the risk of losing a simulation, smaller batches can be simulated; however, the overhead memory allocations and the time to reserve resources can be significant, as a typical densely voxelized scintillator can require tens of gigabytes. The default transverse size for voxels was set to 50 by 50 μm , which is slightly smaller than the best resolution achievable with a 20- μm diameter MLA which, when factoring in the geometric properties, would have a resolution of 60 μm in our experimental configuration. The axial voxel resolution was picked to be sub-millimeter for the 2.3 cm thick scintillators and to be 1 mm for the 5 cm scintillators.

To assess the LF imager’s MTF for the volumetric approach, a 4 cm long line source was

simulated at 6 MV. The line source was angled at 2.5 degrees, and 3 million incident photons were used. An important factor affecting an X-ray imager’s resolution is the focal spot size of the X-ray source. The scintillator was placed 150 cm away in the simulations to minimize the focal spot blurring. The focal spot sizes were assumed to be symmetric in the X and Y plane, and we assumed the largest FWHM direction size from the literature for the sources’ diameter. Therefore, the source size diameter for the LINAC was assumed to be 1.5 mm [44, 45], and 3 mm for the kilovoltage experiments [46]. In the kilovoltage experiments, the diameter of the small animal irradiator source rotation is limited to 60 cm, and therefore in those simulations, the scintillator was adjusted to 55 cm from the source.

Monte Carlo Ray-tracing Model

To determine the DQE of the system, we need to track the incident fluence and the resulting optical fluence on the camera. TOPAS is able to simulate a scintillation event, where optical photons are emitted as rays. In TOPAS, only ray physics is simulated for optical photons, which includes: reflection, refraction, absorption, and transmission. More advanced things like polarization can also be simulated. The downside of TOPAS’ optical implementation is the computational time requirement.

Tables 2.1 and 2.2 list the optical properties used for the scintillators in the simulations. In TOPAS, if a material does not have any optical properties listed, TOPAS assumes the material is opaque to light. Therefore optical properties were added for air and we added black paper to the top surface of the scintillator to generate a scatter similar to what would have been observed in the experiments. In retrospect, black paper should not have been used as a material for shielding the scintillator due to its reflectivity of 10%. The black paper was modeled as a rough surface that reflected light diffusely while absorbing 90% of incident photons [47].

TOPAS records all the produced scintillation photons in a phase-space file that logs the

photon's X, Y, Z coordinates and several of the directional cosine angles allowing for the calculation of the momentum vector of the photon when it crosses the phase space plane. Therefore, simple ray-tracing transfer matrices can propagate a ray through space and a thin lens using the phase space output [48]. Equation 2.1 represents the transmission of a ray through a thin lens with a magnification factor of $M = -\frac{z_1}{z_0}$ where z_1 is the image distance, and z_0 is the object distance,

$$\begin{pmatrix} x' \\ \theta' \end{pmatrix} = \begin{pmatrix} M & 0 \\ -1/f & 1/M \end{pmatrix} \begin{pmatrix} x \\ \theta \end{pmatrix}. \quad (2.1)$$

Equation 2.1 is expressed in 1D, where x and x' represent the initial and final offset of the ray off the optical axis of the system. Assuming a symmetric lens, this can be replicated in 2D. In the case of the light-field camera, the rays are traced to the MLA plane by equation 2.1 and then transmitted to the sensor by

$$\begin{pmatrix} x' \\ \theta' \end{pmatrix} = \begin{pmatrix} 0 & f_{MLA} \\ -1/f_{MLA} & 1 \end{pmatrix} \begin{pmatrix} x \\ \theta \end{pmatrix}. \quad (2.2)$$

The sensor was assumed to be ideal and capable of photon counting, and the counts were binned into $2 \mu\text{m}$ pixels.

DQE method for simulation

In the introduction chapter 1, the concept of DQE was introduced to quantify the efficiency of an imaging system. As will be seen in the experiment chapter, an approach exists to measure the DQE experimentally by imaging an edge or line source to get the MTF and by also imaging a set of empty flood-field images with the X-ray source on and off to quantify the quantum noise of the imaging system. While simulating a line source is not very computationally expensive – simulating hundreds of flood-field images is. The fluence used in

clinics for a flood field usually ranges $> 10^8$ photons. As a result, we simulate the DQE using the fast simulation approach developed by Star-Lack *et al.* known as the "FLS" method [6] recognizing scientists Fujita [49], Lubberts [50], and Swank [51] for their contributions to the DQE calculation.

Recall that the DQE is defined as

$$DQE(f) = \frac{|MTF(f)|^2}{q \cdot nNPS(f)}. \quad (2.3)$$

The MTF was generated by simulating a 4 cm line source angled at 2.5 degrees. In the FLS method the denominator of the DQE is calculated by

$$qNNPS(f) = qNNPS(0) \cdot \frac{NPS(f)}{NPS(0)}. \quad (2.4)$$

Where $NNPS$ is the normalized noise power spectrum, and NPS is the noise power spectrum calculated from measuring single gamma events. In reality, billions of photons are hitting the scintillator at once. However, we can simulate a flux of one X-ray per simulation. A large fraction of single gamma events go unrecorded though, due to the low light-collection efficiency of the optical camera. Therefore, the quantum efficiency (QE) is defined as the ratio of detected events to the number of simulated gamma events,

$$\eta = N_{detect}/N_{\gamma}. \quad (2.5)$$

Where N_{detect} is the number of detected gamma events as recorded by the camera. Thirty-thousand gamma events were simulated to calculate QE. To calculate the noise power spectrum without biasing it, the optical output of the scintillators was artificially increased to 1 million photons per MV. Three thousand new gamma events were recorded with this brighter

response. The 1D NPS was calculated from the gamma events by

$$NPS_{detect}(f) = \left| DFT_i \left(\sum_j PSF_{detect}(i, j) \right) \right|^2. \quad (2.6)$$

Where i, j indicate pixels, and DFT is the 1D discrete Fourier transform. PSF_{detect} refers to the 2D point-source response function for a detected gamma event. Summing along i or j will result in the X or Y directional NPS. The average NPS for all the detected gamma events was done by

$$NPS_{\bar{S}}(f) \equiv \frac{1}{N_{detect}} \sum NPS_{detect}(f). \quad (2.7)$$

The Swank factor was used to scale $NPS_{\bar{S}}(f)$ to the appropriate height. In an ideal photon counting system, the recorded value would be one-to-one based on the number of incident X-ray photons, however in the case where the image is formed by integrating scintillation pulses like in a camera scintillator setup, the signal-to-noise ratio is reduced by the Swank factor which is based on the shape of the pulse-height spectrum (PHS) response of the scintillator from the gamma events [51]. The pulse height spectrum is determined by measuring the total light output received by the camera for a gamma event and recording that output frequency to generate the light output distribution. The light output distribution is normalized, and from the distribution, the Swank factor can be determined by

$$qNNPS(0) = \frac{M_2 M_0}{\eta M_1^2}. \quad (2.8)$$

Where M_i represents the i th moment of the PHS distribution, which is defined as

$$M_i = \int x^i P(x) dx = \frac{\sum_x x^i \cdot counts_x}{\sum_x counts_x}, \quad (2.9)$$

where x represents the total light recorded from a single gamma event and $counts_x$ represents the number of times that particular light total was observed. Figure 2.3 shows an example

PHS for the LKH-5 scintillator at 6 MV. Combining the Swank factor and the NPS leads to equation 2.4, the denominator in the DQE calculation, and is significantly faster than simulating traditional flood-field images.




Figure 2.3: Pulse-Height-Spectrum for LKH-5 at 6 MV. The total light collected represents the sum of all pixel values in a single gamma event image.

2.2 Optics

2.2.1 Camera Geometric Relations

Modeling a standard model light-field camera, where the microlens array is at a fixed distance from the sensor, requires knowledge of the focal length of the main lens f_1 , the focal length of the microlens f_2 , and the distance between the micro-lens array and the main lens z_1 . Figure 2.4 showcases the optical train of a standard LF camera. In practice, the selection of an MLA and the placement of the array is coupled to the properties of the main lens; specifically, the MLA is chosen to match the image-side F-number ($F/\#$) of the main lens.

The $F/\#$ of a lens is defined by

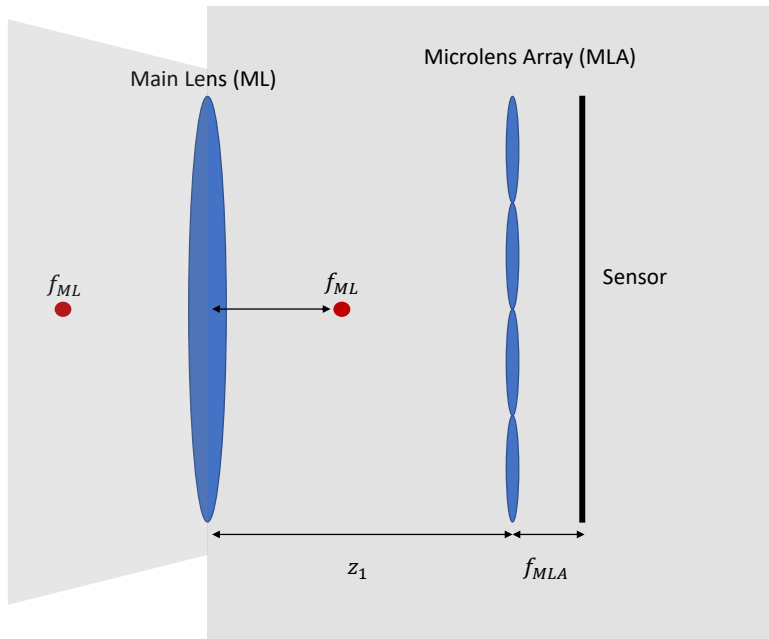


Figure 2.4: The optical train inside the Lytro Illum – a standard model light-field camera.

$$F/\# = \frac{f}{D}, \quad (2.10)$$

where D is the diameter of the lens and f is the focal length. It is important to note that the focal length of a lens is only defined for infinity focus. In applications with a finite focus distance, the sensor has to be further than the focal length away to form an image in a camera. Therefore, the image-side $F/\#$ uses the actual image distance z_1 and is defined as

$$(F/\#)_{image} = \frac{z_1}{D}. \quad (2.11)$$

The focus distance of a camera can be calculated with the thin-lens equation assuming the object (z_0) is sufficiently large, typically four times the focal length for non-telephoto lenses,

and for telephoto lenses, 7-9 times the focal length of the lens.

$$\frac{1}{z_0} + \frac{1}{z_1} = \frac{1}{f}. \quad (2.12)$$

Solving for z_0 yields,

$$z_0 = \left| \frac{z_1 f}{z_1 - f} \right|. \quad (2.13)$$

The system's magnification is defined as $M = z_0/z_1$, which is counter-intuitive as most definitions of magnification are flipped. Using this definition of magnification, the size of the image in the microlens (s_i, t_i) plane is related to that in object space (s_o, t_o) by

$$\begin{aligned} s_o &= M s_i, \\ t_o &= M t_i. \end{aligned} \quad (2.14)$$

Camera lenses are typically sold based on their focal length and the sensor size the lens can be used on. Therefore, f_1 and z_0 are the constraints in the camera optimization problem. Similar to how z_0 was calculated by equation 2.13, z_1 can be calculated if z_0 and f_1 are given. Ideally, we would want to use the highest focal length possible to get the best possible resolution on the scintillator, however, there are trade-offs in using the highest focal-length lens, with one of the largest drawbacks being the resulting small field-of-view (FOV). The FOV of a camera can be calculated by

$$FOV = M * d_s, \quad (2.15)$$

where M is the magnification and d_s is the physical size of the sensor. Smaller sensors introduce a lower max focal length to meet a FOV requirement. Sensor sizes have grown considerably since the first camera-based EPIDs, with the so-called 'full-frame' sensor, which has a diagonal measurement of 35 mm, being the reference size for sensors today. There are

larger sensor formats than the full-frame, such as the medium sensors which can go up to mid-50 mm in diagonal; however, the price of medium sensors can vary significantly based on the pixel count or the size.

A competing factor that has a role in deciding whether to max out the focal length or use a smaller focal length but move the camera closer is the depth-of-field (DoF) effect. Decreasing the focal length increases the amount of the image that is in-focus unless the camera is moved closer to the subject. Using a higher focal length lens tends to collapse the image scene, meaning that further objects appear closer to the near objects and this leads to a shallower DoF. The depth-of-field (DoF) of a conventional camera can be defined as

$$DoF = \frac{2(F/\#)Cz_0^2f_1^2}{f_1^4 - (F/\#)^2C^2z_0^2}, \quad (2.16)$$

where all terms have been previously defined except C , which represents a flexible value that depends on how the image is being viewed. If the image were being viewed on a printed piece of paper, C would represent the printer's resolution. In our case, C could be thought of as the width of a microlens or smaller as a point source in the focus plane. Equation 2.16 implies that as an object moves further away, higher focal lengths are required to counteract increasing DoF. The DoF equation 2.16 is quite comprehensive as it is looking at higher-order effects on the DoF; however, its utility in programming a forward model is limited. If we assume perfect focus at the focus plane z_0 and only look at cases where $z > z_0$, where z represents a different plane in object space, we can rederive the depth-of-field with respect to a front-of-focus DoF and a rear-of-focus DoF [52]. Figure 2.5 shows the off-focus blur experienced on the sensor if point P_0 is moved off the focus plane z_0 .

The change in the image plane distance z'_1 can be expressed as

$$\begin{aligned} z'_1 &= \frac{z_1(D+C)}{D} : z'_1 > z_1, \\ z'_1 &= \frac{z_1(D-C)}{D} : z'_1 < z_1, \end{aligned} \quad (2.17)$$




Figure 2.5: Point P_0 will be out of focus if moved off the z_0 plane. z'_1 refers to the out-of-focus image planes, where C indicates the resulting magnitude of blur in image space.

Here D is the diameter of the main lens. Relating the variables focal length f , diameter D , and blur C , by substituting z'_1 into the thin-lens equation, the resulting near z_n and far z_f distances that meet blur criterion C on the sensor are expressed by

$$\begin{aligned} z_n &= \frac{z_0(D + C)}{D + \frac{Cz_0}{f}}, \\ z_f &= \frac{z_0(D - C)}{D - \frac{Cz_0}{f}}. \end{aligned} \tag{2.18}$$

The total DoF is given by the difference $z_f - z_n$. Given that z_0 is usually greater than f in most standard imaging applications, it can be seen that for a fixed focus and blur C , increasing the focal length compresses the scene and decreases the DoF. Similarly, increasing the aperture decreases DoF.

Implementing the DoF effect into the forward model requires translating the DoF blur into object space. This can be done by solving for C in equations 2.18 and multiplying it by the magnification factor.

$$\begin{aligned} C_{no} &= \frac{Df(z - z_0) z_0}{z_0(f - z) z_1}, \\ C_{fo} &= \frac{Df(z_0 - z) z_0}{z_0(f - z) z_1}. \end{aligned} \tag{2.19}$$

Here the subscript o indicates object space blur, and z indicates distance of object plane. As can be seen with this formulation, this model assumes zero blur at the plane $z = z_0$, which is inaccurate as there is Fraunhofer diffraction, which will be discussed more in-depth in section 2.2.2. As a result, the magnified Fraunhofer diffraction can be substituted for the PSF at the focus plane. Given that the camera simulations in this research use large apertures, the Fraunhofer diffraction in object space will appear as a point given the high resolution of the lens. However, if a small aperture lens is used, resulting in a noticeable diffraction pattern, the blurs calculated with equations 2.19 can be convolved with the Fraunhofer diffraction pattern.

A further adjustment is needed in the object-space blur formulations to replicate the DoF effect in a light-field camera. Unlike in a conventional camera, where each pixel is exposed to the entire aperture, a light-field camera divides the aperture with the MLA, hence giving each sub-aperture image a very large DoF. Given the large DoF in the sub-aperture images of an LF camera, many implementations of light-field refocusing choose to ignore the DoF effect, however, in our case of trying to maximize scintillator thickness, it is important to know to what extent the camera will be able to refocus. If the object is blurred in the sub-aperture images, the light-field camera cannot refocus on that plane without using iterative reconstruction methods.


The conventional camera DOF/blur equations show that the aperture plays a significant

role in the DoF. The effective aperture for each sub-aperture image in a standard plenoptic camera is given by


$$D_{subaperture} = z_1 \frac{\Delta x}{f_2}, \quad (2.20)$$

where f_2 is the MLA focal length and Δx is the sensor's pixel pitch. A common alteration to equation 2.20 is to replace Δx with $s_e = \max(\Delta x, r_{psf})$ as a sensor with a higher pixel count does not bring more angular information if the diffraction from the MLA spans multiple pixels. For example, while the Lytro Illum has 40 megapixels with a pixel pitch of $1.4 \mu\text{m}$, the PSF from the MLA is between $2.5\text{-}3 \mu\text{m}$, effectively reducing the angular information to 20 million rays. Factoring in the fill factor of the MLA and the vignetting of the signal, several million more pixels are wasted.

To illustrate the depth-of-field difference between a light-field camera and a conventional camera, Figure 2.6 shows each respective camera's 3D point spread function and the amount of blur that would occur for each transverse plane in a focal stack. The light-field 3D PSF in Figure 2.6b is for the sub-aperture images taken by the Lytro Illum measured for $197 \geq Z \leq 223 \text{ mm}$, which corresponds to the scintillator thickness and centered on the standard distance that the scintillator was used in our experimental measurements. As can be seen given the smaller aperture sampling from the MLA, the LF camera has zero noticeable blur from main lens' DoF effects within this refocusing range. The conventional camera has a considerable amount of blur for planes off the focal plane shown in Figure 2.6a. In the conventional camera's PSF axial slice image, it can be seen that there is an asymmetry across the focal plane. Objects that sit closer to the camera go out of focus faster than objects that sit behind the focal plane. A similar asymmetry would be seen for the LF PSF if the range was expanded.



(a) Conventional Camera



(b) Light-Field Camera (Lytro Illum)

Figure 2.6: The transverse PSF displayed as a function of distance from the camera. The amount of transverse blur is minimum at the focal plane. The cameras have identical apertures and focal lengths. The middle images display the transverse blur axially throughout the focal stack. The edge images display the 2D transverse blur for the top and bottom focal stack slices. For the conventional camera, the top and bottom slices have different transverse blurs due to the asymmetry in the 3D PSF. In this example, there is no noticeable transverse plane blur for the Lytro Illum sub-aperture images.

2.2.2 Optical Effects

Vignetting

Vignetting is the reduction of light towards the periphery of an image compared to the center. Vignetting has an effect on how much light we can collect for each part of an image. There are different types of vignetting: natural, optical, and mechanical [53]. Mechanical vignetting is from obstructions on the exterior of the lens barrel, blocking the edge light, so this can be effectively avoided with careful design. Optical vignetting occurs with very wide apertures and is from the falloff that very oblique beams experience entering the lens compared to beams entering the lens on the optical axis. The effect is caused by the oblique beam experiencing a smaller oval aperture than the standard circular aperture. However, we can ignore this form of vignetting because of the distance of the scintillator. None of the rays will be obliquely entering the lens. Natural vignetting refers to the light collection falloff and should be modeled. Natural vignetting can be calculated for each point in the object volume by

$$flux = \frac{LSA \cos^4 \theta}{z^2} \quad (2.21)$$

where z is the axial distance of the voxel from the main lens, L is luminance, A is the lens area, S is the voxel area, and θ is the angle the voxel makes with the optical axis. Natural vignetting at high focal lengths can be ignored; however, at small focal lengths, there can be up to a 40% scene intensity falloff towards the edge of the image. Another thing of interest is the z^2 falloff in the denominator. A z^2 falloff is expected for point sources; however, the light falloff from a light source that occupies a fixed area may fall off at a more linear rate depending on the distance from the camera. As a general rule, for a circular light source, if the camera is further than 10x the diameter of the light source, the falloff can be estimated as z^2 .

While not expressed in the vignetting equation, the sensor size is another factor that is important to how much vignetting will be visible. The vignetting equation is written in object space so we can utilize the volumetric simulation approach; however, it is important to note a larger sensor will experience more vignetting from the fact that it has a larger FOV compared to a smaller sensor. To illustrate the vignetting that would occur for a medium format sensor found on high-end cameras is shown in Figure 2.7 for our simulations at $Z = 1$ m.




Figure 2.7: Vignetting for a medium sensor

Diffraction

In the previous section 2.2.1, a great deal of detail went into highlighting the different optical/hardware components of the light-field camera and discussing their impact on the final captured image. When it comes to determining the maximum resolution of a camera, the variables of greatest interest are the camera's F-number and the pixel size of the camera. In our forward model, we will assume that the lenses in the camera are ideal lenses, free of aberrations. In that case, the maximum resolution will be limited by diffraction if the diffraction spans over multiple pixels.

There are two types of diffraction that are prevalent in optics, Fresnel and Fraunhofer diffraction [54]. Determining which diffraction is being observed on an imaging screen depends on the viewing distance from the aperture, the size of the aperture, and the wavelength of light being measured [54]. In optical imaging with a lens, the observed diffraction pattern is Fraunhofer diffraction. Fraunhofer diffraction can be expressed analytically by the expression

$$I_{psf}(r) = \left(\frac{A}{\lambda z}\right)^2 \left[\frac{J_1\left(\frac{kw r}{z}\right)}{\frac{kw r}{z}}\right]^2. \quad (2.22)$$

Here z is the viewing distance of the screen from an aperture with a radius of w , J_1 is the first-order Bessel function, k is the wave number, which is given by the expression $k = \frac{2\pi}{\lambda}$, and r is the radial distance from the PSF center $r = \sqrt{x^2 + y^2}$. I_{psf} in its current form, all terms being squared, accounts for the scintillator light being incoherent. The PSF of a polychromatic light source paired with a monochromatic camera with no Bayesian color filter can be determined by the summation of each wavelength's PSF $I(r) = \sum_{\lambda} I_{psf}(r, \lambda)$. However, in practice, most scintillators have a primary wavelength emission and unless the camera is a scientific camera, a color filter is often on the sensor by default allowing for a single PSF computation using eqn 2.22 for the primary wavelength.

In our simulation model, we forward project $y = Ax$ and back project the data $x = A^T y$ iteratively. Expanding the forward and back projection operators A and A^T , respectively, we can incorporate the point-spread blurring into our model by using the PSF and its adjoint in each respective direction. In our case, the PSF is equal to its adjoint meaning for both forward and back projections the data can be blurred by the same PSF.

We are simulating standard plenoptic cameras and as discussed in the introduction, the standard light-field camera's resolution is determined by the size of the MLA. As a result, the blur from the primary lens can be ignored as the spatial sampling by the microlens is an order of magnitude larger than the measured blur. However, if an upsampling operation is

being performed to increase the resolution, which effectively cuts the microlens' diameters by the upsampling factor, then incorporating the blurring from the main lens is critical.

2.3 Reconstruction of volumetric scene

In computed tomography, multiple projections of a patient are used to reconstruct the 3D slices of a patient. In order to provide slices with minimal blurring and relatively isotropic spatial resolution, these projections have to be over a wide range of angles around the patient. In our application, we wish to perform a volumetric reconstruction of a scintillator from a very limited angular set that spans the width of the main lens. As a result, we expect significant blurring axially. As a result, our primary motivation and metric of success for simulation will be improving axial resolution.

Previous groups that have explored light-field scintillator imaging have used either one light-field camera projection paired with an EPID projection for axial resolution optimization or two orthogonal light-field projections, which allow for imaging independent of the presence of an EPID [20]. It has been shown that the light-field camera's axial resolution is dependent on the reconstruction algorithm used, but also the scene. Liu *et al.*, assessed a light-field camera's performance for both discrete and continuous luminous fields, where the former is an image scene consisting of scattered point sources in a 3D volume, whereas the latter is more akin to scintillator imaging where a point source occupies every voxel in the image volume. In their assessment, they concluded that a single plenoptic camera was not sufficient to reconstruct a continuous luminous source; however, the axial resolution was not null indicating the potential for partial spectral extraction. [55].

We define the reconstruction as the standard linear model of

$$y = Ax, \tag{2.23}$$

where A represents the forward projection matrix of the scintillation volume x to the camera’s sensor image y .

The construction of the projection matrix is performed by Astra Toolbox [56]. Light-field cameras can be modeled as a limited-angle tomography problem with existing tomographic software. However, conventional tomographic software does not model the optical photon production in the scintillator or the optical blurring of the camera system, so there exists no out-of-box solutions to simulate a light-field scintillator imager. As a result, the forward model developed here combines several open-source tools. The model consists of several components: scintillator simulation, limited angle tomography projection, optical blurring, and lastly, image reconstruction.

2.3.1 Limited Angle Tomography

Light-field imaging is a limited-angle tomography problem akin to tomosynthesis imaging for breast cancer detection. In tomosynthesis, high-resolution 2D projections are collected as the X-ray source moves across in an arc trajectory above the breast. A greater arc angle allows for better Z-axis resolution. A greater number of projections increases the in-plane (X-Y) resolution. Ideally, one would want to take as many projections with as wide of an arc angle as possible to give the best volumetric resolution; however, given dose constraints, there is an optimization between the number of projections vs the arc angle size to meet in-plane resolution demands while providing some additional Z-resolution [57].

Unlike tomosynthesis, where each projection image requires additional dose, light-field cameras collect their angular information in a single shot. The amount of angular information a light-field camera can capture is limited by its lens diameter and the number of pixels behind each microlens. Akin to tomosynthesis, the wider the main lens, the better the Z-axis resolution, and the more pixels behind each microlens, the better the in-plane resolution is from removing out-of-plane blurring effects. Given the similarities between tomosynthesis

and light-field geometry, Viganò *et al.* developed a light-field camera framework utilizing an existing open-source tomographic software, ASTRA toolbox [23, 56]. ASTRA toolbox is a flexible tomography framework to model unique source, volume, and detector configurations while supporting standard tomographic configurations like CT. The coding foundation of ASTRA toolbox is written in C++, with CUDA for GPU acceleration in the projection operations. To expand to a wider research audience, ASTRA toolbox also features a Python wrapper interface.

Viganò’s light-field framework for ASTRA is known as PLENOPTOMOS (Plenoptic imaging reconstruction tools) and is open-source and based on Python [58]. The light-field camera frameworks available in PLENOPTOMOS are the standard plenoptic camera and plenoptic 2.0 cameras. While more recent cameras released by Raytrix feature hybrid MLAs that use microlenses of different focal lengths interlaced to expand the DoF further, those cameras cannot be readily simulated due to the underlying architecture of ASTRA toolbox assuming identical geometry effects for each microlens and sensor pixel [26]. While the multi-focus light-field cameras offer a significantly larger DoF, in our study, we consider scintillator thicknesses that are well within the DoF for a standard plenoptic camera.

The ASTRA toolbox is versatile in simulating unconventional X-ray setups because it allows for the geometry to be specified in vector formalism. For each projection image, a 12D vector is specified $(srcX, srcY, srcZ, dX, dY, dZ, uX, uY, uZ, vX, vY, vZ)$, where *src* refers to the source location in 3D, *d* refers to the center of the detector, and *u* and *v* represent the detector basis functions. For example, a detector on the X,Y plane would have *u, v* vectors of $u = (1, 0, 0), v = (0, 1, 0)$. With this vector-based formalism, it is possible to simulate geometries with multiple light-field cameras placed adjacent to each other, geometries with two light-field cameras looking at the scintillator orthogonally, or to simulate a combination of light-field and conventional camera geometries. While we did not study these setups, they represent possible directions for future research.

Figure 2.8 displays the high-level geometry used to simulate a light-field camera. Unlike a real light-field camera where the angular information is captured by pixels behind a microlens, the angular information is simulated here as shifting the X-ray source by $\Delta u, \Delta v$, which can be calculated by the equation

$$\Delta u = \frac{f_1 \Delta x}{f_2}. \quad (2.24)$$

Where Δx is the pixel size, f_2 is the MLA focal length, and f_1 is the main lens focal length. Δv can be computed similarly by using the pixel dimension in the y dimension.

In ASTRA toolbox, the center of the volume is the origin of the coordinate system. ASTRA is rather unconventional with its source placement. In Figure 2.8, the source is displayed as in a conventional tomographic acquisition. In the light-field application, the source is the light being emitted from the volume; therefore, the source is placed at the center axial slice of the volume, and the MLA is placed distance Z_0 away. The magnification effect is simulated in ASTRA for a thin lens by scaling the MLA pixel grid basis vectors accordingly.

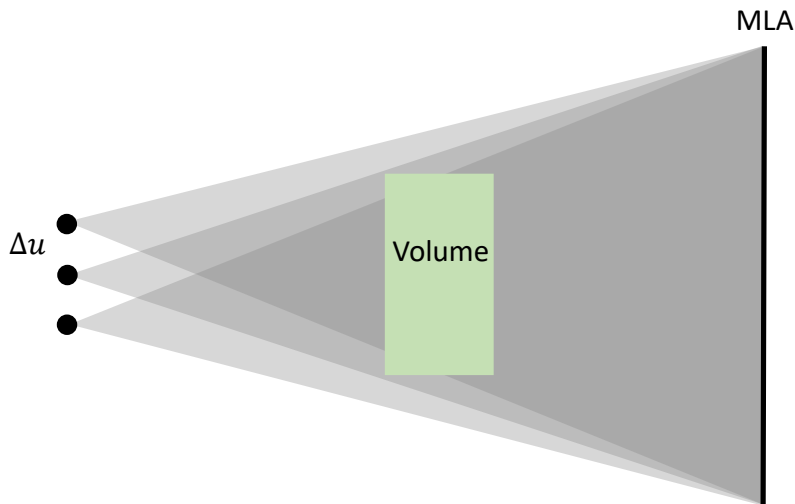


Figure 2.8: High-level ASTRA geometry for light-field cameras.

2.3.2 Reconstruction Algorithms

There exists a variety of algorithms from the tomography literature that could be applied to light-field imaging reconstruction. One such algorithm that has been used in similar transparent, luminous imaging scenes with a light-field camera is the maximum likelihood estimation algorithm (ML-EM), [20]. In ML-EM, the k th iteration step is given by:

$$x^{(k)} = \frac{x^{(k-1)}}{A^T \mathbf{1}} A^T \left(\frac{y_m}{Ax^{(k)}} \right), \quad (2.25)$$

where y_m is the measured light-field projection, x is the image volume which varies with each iteration k , and A and A^T represent the forward and back projection matrices, respectively. ML-EM has several useful qualities, including enforced data positivity, useful statistical properties, and ease of implementation. In opaque imaging scenes, x can be denoted as x_z where x_z is just a plane some distance z away instead of reconstructing over the entire image volume; however, in our application, we will be doing 3D refocusing to prevent noise amplification which occurs with optimization per x_z in transparent scenes.

The stopping criterion for the ML-EM is when the maximum log-likelihood function reaches its maximum. The following expression gives the log-likelihood function:

$$\mathcal{L}(y_m, \hat{y}^{(k)}) = \sum_{p=1}^N \left(y_m \ln \hat{y}^{(k)} - \hat{y}^{(k)} \right), \quad (2.26)$$

where \hat{y} is the light-field projection data from the k th iteration of ML-EM, and the summation is over all the pixels. A stopping criteria of $|\mathcal{L}^{(k)} - \mathcal{L}^{(k-1)}| < \epsilon$ with $\epsilon = 0.001$. In addition, the root-mean-squared error was calculated using

$$RMSE = \sqrt{\frac{\sum^N (x^{(k)} - x^{(k-1)})^2}{N}}, \quad (2.27)$$

where N and the summation are the total number of voxels in the imaging scene.

2.4 Simulation Experiments

To compare camera models, we simulated the required MTFs and NNPSs to calculate the DQE for the studied energies. We simulated the physical cameras used in our experiments to verify the simulation’s validity. The physical cameras used in the experiments were the Lytro Illum and a scientific monochromatic CoolSNAP camera [59, 60]. While it was impossible to fully verify all the camera parameters for the Lytro and CoolSNAP due to propriety software or lack of software altogether, we performed educated estimates on the few unknown parameters. We made high-level comparisons between the simulated data vs. the experimental. In addition to simulating the physical cameras, the scene geometry was replicated, including scintillator to source distance, scintillator material, and scintillator thickness.

More ideal light-field cameras and scintillator parameters were also simulated. More ideal light-field cameras were simulated by using cameras with faster F-numbers (F/0.95) and comparing them against their equivalent conventional cameras. The F/0.95 aperture was selected as it is quite unusual for a consumer lens to have a larger aperture than that on a 30 mm focal length lens, which was approximately the focal length for the lens on the Lytro Illum. However, for simulations where the camera was far away, the F-number limit was increased to F/1.4 as zoom lenses typically do not exceed 100 mm in diameter. For example, a 115 mm focal length zoom lens at F/1.4 would have a physical aperture diameter of 82 mm. As a result, an F/1.4 115 mm zoom lens is large and typically costs a thousand USD or more potentially.

The main scintillator parameters we looked at were scintillator thickness and scintillator material. As previously mentioned, we used LKH-5 experimentally, but EJ-260 is an alternative. Both scintillators were simulated at a thickness of 5 cm, and LKH-5 was additionally simulated at 2.2 cm to match experiments. The 5 cm thickness was arbitrarily selected, as the thickness could have been increased further. However, increasing scintillator thickness requires decreasing the lens zoom, which reduces the resolution. At a certain thickness,

the resolution will become non-competitive with an EPID, and the scintillator cost will be prohibitive.

2.5 Results

Tables 2.3 and 2.4 show the camera specifications used in the simulations. In the simulations, the scintillator was modeled to be only 8 by 8 cm, as the NPS images were generated by a delta ray in the center of the field of view. The resolution reported is the resolution at the scintillator plane factoring in camera resolution/magnification. To remain competitive with an EPID, the scintillator plane resolution needs to be $<340 \mu\text{m}$, which requires the MLAs to be considerably smaller. The tables show that even with $20 \mu\text{m}$ diameter microlens, the object resolution is already at $125 \mu\text{m}$. The field-of-view (FOV) for the simulated 28.9 mm, 40 mm, and 115 mm focal length lenses, assuming a medium sensor with dimensions of 33 x 44 mm, are 20.7 x 27.5 cm, 14 x 18.8 cm, and 25.4 x 33.8 cm, respectively. Note, the 115 mm focal length lens was used at 1 meter away, resulting in the larger FOV despite being a stronger lens.

Table 2.3: Camera Parameters at $Z = 20 \text{ cm}$

Camera Model	MLA Diameter	f_1	$(F/\#)_{ML}$	$(F/\#)_{MLA}$	DOF	Resolution
Lytro	20 μm	28.9 mm	2	2.3	37 mm	125 μm
LF ideal	20 μm	28.9 mm	0.95	1.10	17 mm	125 μm
Experiment camera	NA	40 mm	2.8	NA	1.8 mm	22.3 μm

Figure 2.9 shows a 3D-modeled phantom that consists of pillars at different depths. The first two pillars represent the bottom of the scintillator, and the last thicker pillar represents the top of the scintillator. To verify the refocusing ability of the forward model, the 3D phantom was also 3D printed and imaged with the Lytro Illum. Shown are the center focal stack slice with the middle pillar in focus. The amount of blur experienced by the simulated

Table 2.4: Camera Parameters at $Z = 1$ m

Camera Model	MLA Diameter	f_1	$(F/\#)_{ML}$	$(F/\#)_{MLA}$	DOF	Resolution
LF ideal	20 μm	115 mm	1.4	1.58	37.5 mm	153 μm

model agrees well with the Lytro Illum. In addition, both the experimental data and modeled data can be refocused using the same parameters indicating further agreement.




Figure 2.9: A 3D phantom modeled and 3D printed to compare the simulation and experiment for initial verification.

Figure 2.10 displays a cropped sensor image of a rectified MLA array of the line spread source inside a 5 cm thick scintillator. While almost all LF camera processing software rectifies the MLA image to have the MLAs appear in a rectangular arrangement, we can also simulate MLAs in a hexagonal pattern, which is what the Lytro Illum has. Figure 2.11 shows a Monte Carlo ray-tracing example of a tungsten edge where the hexagonal arrangement is the main focus. Depending on the sensor shape, the circle packing increase from rectangular packing to hexagonal packing increases the coverage from 78% up to 91%, which leads to a better quantum efficiency for the camera.

The $DQE(0)$ of an ideal Lytro Illum light-field equivalent camera in our experimental setup is found to be 0.065 when assuming the MLA's F-number is matched to the main lens' image-side F-number; however, in application, the MLA is fixed, and the Lytro Illum has a zoom lens, which prevents the MLA from being F-number matched for all zooms. Factoring




Figure 2.10: A view of a rectified microlens array image created by the Monte Carlo ray-tracing model




Figure 2.11: A view of a hexagonal microlens array image created by the Monte Carlo ray-tracing model

in the MLA mismatch, a more accurate $DQE(0)$ for the Lytro is 0.053. An equivalent conventional camera with the same main lens ($F/2$ 28.9 mm) has a $DQE(0)$ of 0.068. The ideal Lytro Illum supports an $F/2$ aperture and when compared to an $F/2$ conventional

camera, the conventional camera has an overall higher DQE across spatial frequencies as shown in Figure 2.12. The source of this unexpected result stems from the noise amplified during the ML-EM reconstruction. The LF camera does provide equal to or better resolution compared to the conventional camera shown in Figure 2.13. However, as seen in the noise-power spectrums, the ideal Lytro Illum’s NPS flatlines at a significantly higher base of 0.3 compared to the <0.05 for the conventional camera.

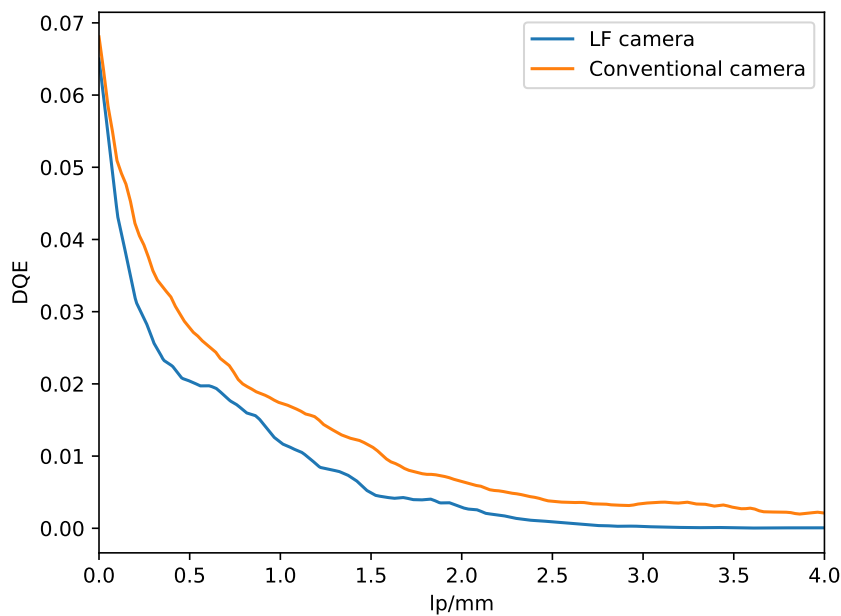


Figure 2.12: DQE for lens matched to ideal Lytro Illum

Recall that the noise power spectrum in the rapid Monte Carlo DQE approach calculates the NPS from single gamma events. Given that the LKH-5 scintillator is several centimeters thick and the CSDA range of an electron at 6 MV is on the order of a centimeter, many gamma events will appear trail-like, where the path of the electron or x-ray is illuminated by optical photons. Several examples of gamma events are shown in Figure 2.15, which shows two gamma events recorded as seen by the conventional camera and as seen by a light-field camera. The gamma events shown for the light-field camera are the reconstructed focal stack

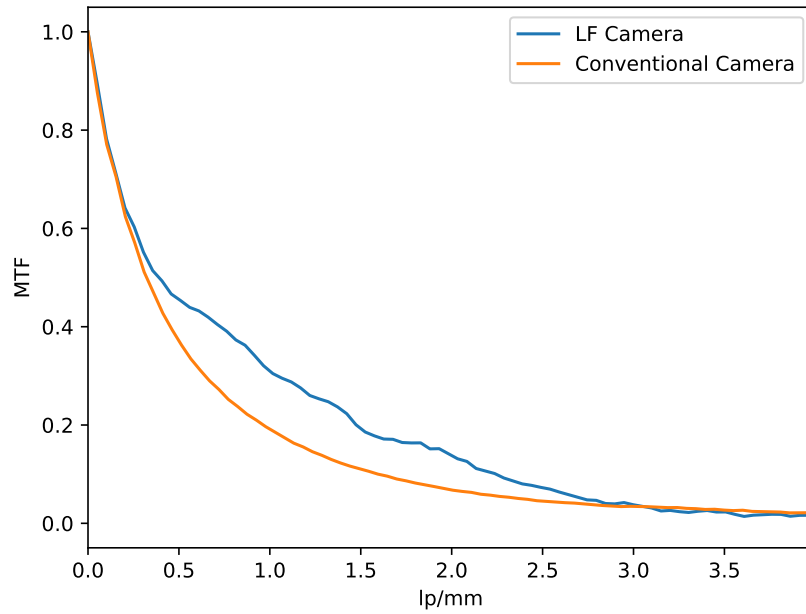


Figure 2.13: MTF for lens matched to ideal Lytro Illum

images summed into one image. Given that both cameras are focused to the center of the scintillator, it can be seen from the first gamma event example that the light-field camera was able to resolve the '3' like pattern, whereas the conventional camera's is out-of-focus.







Figure 2.14: The normalized noise power spectrum for 6 MV




(a) Conventional camera



(b) LF camera



(c) Conventional camera



(d) LF camera

Figure 2.15: Example of gamma event images in the 2.2 LKH-5 scintillator

The selling point of using a light-field camera is its ability to refocus and perform volumetric reconstructions. However, given the noise amplification from reconstruction if a wide aperture lens is not used, the gains of using a light-field camera will be negligible. A potential area of future research would be to better model and suppress the noise in the iterative algorithm.

Figure 2.12 showed the DQE for a lens-matched scenario. As will be shown in chapter 4, we were not able to lens match the conventional camera to the Lytro. What ended up being used in the experiment was a Nikon 40mm f/2.8G AF-S DX Micro Nikkor Lens [61] and unfortunately, with the camera that was being used the lens configurations could not be exported as a file with an image. Therefore, using educated approximations the DQE for the experimental camera was estimated and shown in Figure 2.16. The DQE(0) from experiment for the conventional camera was 0.159 vs 0.161 predicted. The measured vs predicted DQE(0) for the Lytro was 0.004 vs 0.053, respectively. While these simulations lack electronic noise and assume perfect photon absorption on the sensor, it was able to model the conventional camera indicating successful forward modeling. It will be shown in chapter 4 that electronic noise is one of the sources of the DQE degradation for the Lytro and that the Lytro cannot be modeled without additional specifications on the lens configuration.




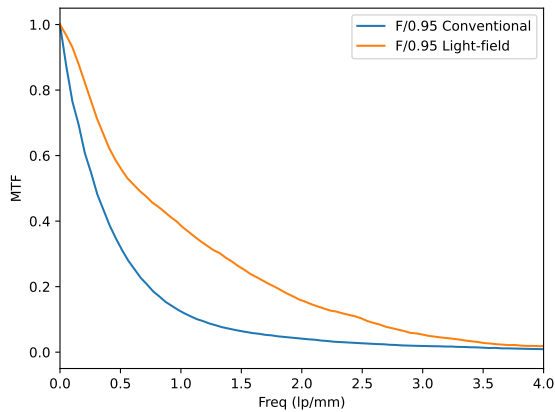
Figure 2.16: Approximated DQE for the experimental camera vs the Lytro Illum.

As previously discussed, as the F-number is increased the depth-of-field increases resulting


in a less out-of-focus light blur as everything is in focus. The only source of blur for an all-in-focus image would be from scintillator blur. While significant, it is not significant enough for the LKH-5 to significantly degrade the resolution with higher F-number lenses. In addition, imaging with higher F-number lenses is not ideal for scintillators such as LKH-5, which is very dim at an optical output of 4000 optical photons per MV resulting in poor optical coupling.

The simulated 2.2 cm LKH-5 scintillator at 6 MV has a quantum efficiency of 33%. If the LKH-5 scintillator was used in an EPID, the EPID would have a DQE(0) of 19% however, the EPID's DQE would exponentially fall to zero from the light scatter within the scintillator. Using an F/2 lens on a traditional camera sensor results in a QE 12.75% and for the light-field camera a QE of 11.96% with the difference stemming from a small percentage of photons being lost from the MLA and not landing in an MLA image. At 2.5 MV, the QE of the LF camera falls to 10.95%, which mirrors experiments where exposure times had to be increased to accommodate the lower energy X-ray fluence.

Given the rather poor QE despite the scintillator interacting with a third of all incident X-rays, there are effectively three ways to improve the optical coupling between the scintillator and the camera: use a lower F/# lens, move the camera closer, and lastly and harder to do, use a brighter scintillator. In our proof of concept design, we already have the camera as close as possible, therefore, simulating a brighter lens with an F/0.95 leads to a significantly improved DQE as shown in Figure 2.17. In this case, the brighter lens led to better optical coupling, which resulted in the light-field camera outperforming the conventional camera at all but higher frequencies. It is important to note that assuming ray-optics and basic refocusing techniques such as the integration technique this configuration of using an F/0.95 lens with 20 μm diameter micro-lenses leads to a DoF in a sub-aperture image of only 17.34 mm, slightly smaller than the scintillator. However, ML-EM was able to effectively remove the small DoF effect from the sub-aperture images leading to better resolution across the



(a) MTF comparison



(b) DQE comparison

Figure 2.17: Lens matched to an F/0.95 lens.

spatial frequencies compared to the F/2 light-field camera shown in Figure 2.18. From a photography standpoint, the F/0.95 lens does offer a higher maximum resolution than an F/2.




Figure 2.18: MTF comparison between F/0.95 and F/2 LF cameras

The previous DQE results were all based on our proof-of-concept geometry, where the camera was only 20 cm away from the scintillator. In a clinical setting, that would not be the case. The distance from the scintillator to the camera would be closer to a meter. In Figure 2.19 the DQE is shown for a 5 cm slab of EJ-260 and a 5 cm slab of LKH-5 viewed with a F/1.4 115 mm lens. As can be seen in Figure 2.19, there are still additional gains to increasing the scintillator thickness, in particular for the LKH-5 scintillator. The EJ-260 scintillator, while being twice as bright as the LKH-5 scintillator, lacks the density to attenuate the beam to be competitive with LKH-5. While this is encouraging for further investigation, the price of using a large 40x40x5cm slab of LKH-5 may be a large cost barrier for an EPID with the price of a LF camera on top.




Figure 2.19: DQE from one meter away for a LF camera.

Recall that the DQE was computed from the ray-tracing model. Despite ray tracing being extremely computationally expensive, it does capture higher-order effects like light scatter, reflection, and absorption that the volumetric approach does not. Comparing the MTFs

for the line source simulated by both forward models there are some noticeable differences. In Figure 2.20, the ray-tracing’s MTF has a lower resolution at higher frequencies than the volumetric approach, possibly from light scatter off the top of the scintillator which is believed to cause long-ranged blurring in scintillator images [62].




Figure 2.20: MTF comparison between Monte Carlo ray-tracing and volumetric approach

2.6 Discussion & Conclusions

The results from the simulations indicate there is potential for light-field cameras in some scintillator imaging configurations. More research on noise suppression during the reconstruction phase could reduce the NPS of light-field cameras making them more competitive with conventional cameras when the DoF effect is minimal, which suggests that light-field cameras are best for very thick scintillator imaging. While light-field cameras do have a refocusing range limit that is dependent on the optics and also on the algorithm used, that is another area of potential exploration. It was observed that our ML-EM algorithm was

able to refocus and also effectively remove out-of-focus light. To what extent ML-EM could be used to deconvolve out-of-focus light was not thoroughly investigated; However, in the material decomposition chapter 3, the axial resolution of ML-EM is explored in spectral decomposition tasks.

Our light-field scintillator model improved upon existing tools such as ASTRA-toolbox and PLENOPTOMOS to offer arguably one of the most versatile light-field reconstruction tools. There are several additional algorithms like Chambolle-Pock that could be easily employed; however, that requires searching that algorithm’s parameter space [63, 64]. While our forward model does require the user to process their light-field data into an acceptable rectified format, the only other major light-field repository is the Light field toolbox for MATLAB, which does offer good tools to get access to the raw camera images. However, many of its refocusing abilities are limited to opaque scenes [65]. The small resource base for light-field camera users and the proprietary nature of camera manufacturers further emphasizes the need for open-source and accessible code to expand the use of light-field cameras.

CHAPTER 3

MATERIAL DECOMPOSITION & SPECTRAL ANALYSIS OF IMAGING WITH A LIGHT-FIELD CAMERA

3.1 Introduction

In conventional optics, spectral imaging refers to the measurement of different wavelengths of light. A cellphone camera may be limited to only detecting visible photons, while NASA has many instruments, such as the Airborne Visible / Infrared Imaging Spectrometer (AVIRIS) used in earth remote sensing, that can detect photons from a much broader range from the visible to infrared [66]. Using a combination of wavelengths, AVIRIS can do things such as see through smoke to determine a future fire's fuel source or to detect soil moisture levels.

Spectral imaging in relation to X-ray imaging means being able to resolve and separate X-rays based on their energy E . X-rays are exponentially attenuated and can be modeled as

$$I_{out} = I_{in}e^{\mu(E)x} , \quad (3.1)$$

where a beam consisting of I_{in} photons/area passes through a material with linear attenuation coefficient $\mu(E)$ with a thickness x . I_{out} is the exiting beam fluence [4]. Lower energy X-rays are preferentially absorbed in a material than higher energy which has led to various detector designs that try to utilize this effect. However, the challenge is separating the signal from the low vs high energy interactions in the detector [30]. The multi-layer EPID achieves its energy separation by using multiple detection layers, with each layer being additionally hardened by a metal plate [29]. In the light-field imager design, the energy separation will ideally appear by reconstructing the scintillation volume and the recorded light in each slice would be representative of the energy absorbed in that layer.

Spectral imaging of X-rays can allow for an increase in the contrast-to-noise ratio, or to

perform material decomposition. Material decomposition is the process of identifying the amount and type of material that an X-ray has passed through. Determining the material can allow for material removal from an image. For instance, bone is typically removed to improve the contrast of tissue sitting behind the bone.

There are two ways to generate spectral information, X-ray source modifications or detector design. Dual-energy CT is a commonly used imaging modality that uses spectral information to perform material decomposition. X-ray source switching between high and low energies or having two sources in the CT is one approach dual-energy CTs will use to generate spectral information. However, this approach requires additional imaging by the two energies leading to an increase in dose. An alternative approach is to make use of the polychromatic X-ray emission from an X-ray source and alter the detector design to filter from low to high-energy X-rays [67]. This approach has the potential to limit the imaging dose requirement for spectral imaging by gathering all the information in a single acquisition. In megavoltage imaging, improving contrast is very important, given the lower contrast inherent in megavoltage images. Unlike dual-energy CT using two imaging acquisitions at different energies, we use the polychromatic nature of X-ray generation as the spectral information.

This chapter is broken up into two sections: exploring the contrast-to-noise enhancement of using a thick monolithic scintillator paired with a light-field camera to perform material removal from images, and a Cramér-Rao lower bound analysis on the material decomposition performance of an ideal light-field camera,

3.2 LKH-5's Material Decomposition Performance

The monolithic scintillator being simulated and used is axially 2.2-2.3 cm thick. In terms of scintillators, this is a novel thickness even for LKH-5, which has only been used in thicknesses up to 1.5 cm in previous studies [29]. Since the camera is the limiting factor in how thick a scintillator can be, we first explore the theoretical performance of a light-field camera

imaging a point source. Even though a scintillator is a continuous light source, imaging a point source will give an upper bound on the axial resolution.

Using a thick monolithic scintillator for material decomposition is novel, and given that existing megavoltage multi-layered detectors utilize multiple metal buildup layers, it is of interest to see if a monolithic scintillator can perform the same level of material removal before factoring in the degradation that occurs by imaging the scintillator with a light-field camera.

3.2.1 *Light-field imager axial resolution*

In the first simulation, an array of point sources were placed on the camera’s focal plane shown in Figure 3.1 to assess the axial resolution of the light-field cameras. The points were arranged as the dots found on the 5 side of a dice. One point was placed in the center of the field-of-view and the edge points were placed at 20% and 80% of the field-of-view distance. In addition, the PSF array was shifted closer by half the scintillator thickness and further by half the scintillator thickness from the camera to simulate point sources off the focus plane.

A camera replicating the Lytro Illum with its F/2 lens and an ideal light-field camera with an F/0.95 lens were simulated. The scene was reconstructed using 500 iterations of ML-EM.

Figure 3.2 for the Lytro Illum and the F/0.95 light-field camera shows the axial resolution for the center point source. The full-width-half-max (FWHM) value for the two cameras was calculated to be 345 μm . Notably, this 345 μm is reported for a camera-to-scintillator distance of 210 mm. However,





Figure 3.1: PSF arrangement

using a sufficient zoom lens can preserve this axial resolution for a light-field camera a meter

away.

The asymmetric shape of the axial PSF was expected due to conventional cameras having asymmetric PSFs, as discussed in chapter 2. As a result of this asymmetry, a point source that occurs in front of the focal plane at a distance equal to the front surface of the scintillator will have an approximately three times wider FWHM than the point source located in the focal plane. On the far side of the focal plane, the PSF has an FWHM approximately two times wider than the focal plane's. Figure 3.2 shows relative sizes and shapes for off-the-focal plane point sources relative to a point source on the center. In addition to axial point source variation, a point source on the periphery of the image will have a slightly different PSF than the center PSF, but the transverse effect is significantly less than the axial effect.



(a) The axial PSF for a light-field camera. (b) Relative shapes for point sources off the focal plane.

Figure 3.2: Shift-variant PSF for light-field camera

3.2.2 Modeling the optical photon production in scintillator

Analytically

Since we are reconstructing the scintillator volume into a focal stack, each layer in the focal stack is going to correspond to a set thickness. The probability that an X-ray interacted in

layer i can be expressed as

$$M_i(A) = \eta_i^{tot} \int N_i(E; A) E dE, \quad (3.2)$$

where η_i^{tot} is an efficiency factor given by the product of the lens' collection efficiency, scintillator energy conversion efficiency, and the quantum efficiency of the camera at the emission wavelength of the scintillator [68]. $N_i(E; A)$ represents the number of X-rays absorbed by layer i of the scintillator. $N_i(E; A)$ can be expressed as $N_i(E; A) = I(E; A)D_i(E)$ where $I(E; A)$ is line integral of attenuation through the object and $D_i(E)$ is the probability of photon of energy E being detected in layer i and which can be expressed by

$$D_i(E) = (1 - e^{-\mu(E)t_i}) \prod_{i'=1}^{i-1} e^{-\mu(E)t_{i'}}, \quad (3.3)$$

where $\mu_i(E)$ is the linear attenuation coefficient and t is the layer thickness [68]. Alternatively, we can also express the light production in continuous space by

$$N(z) = I_0 \int Z(E)\psi(E)\mu_s(E) \left[e^{-\mu_s(E)(z+z_{max})} \otimes h(z - z_{max}) \right] dz, \quad (3.4)$$

where I_0 represents the initial X-ray fluence per area, $Z(E)$ is the normalized energy distribution of the incident photons in units of keV^{-1} , $\psi(E)$ represents an energy-integrating model where $\psi(E) = kE$ with k as a proportionality constant with units of optical photons per unit keV, μ_s represents the linear attenuation coefficient of the scintillator, and $h(z - z_{max})$ represents the electrons and their spread, where z_{max} is the reference to scintillator width. At 225 kV, the electron travel is not significant. However, at 6 MV, the electrons can travel on the order of a cm generating additional optical photons and lowering the scintillator's resolution [4]. Figure 3.3 showcases LKH-5's optical response to 225 kV, 2.5 MV, and 6 MV. At 225 kV, LKH-5 attenuates the entire X-ray beam, however, at 6 MV, the attenuation

is 33%. Figure 3.4 shows the difference between modeled optical photons using equation 3.4, TOPAS Monte Carlo, and the measured optical light for 6 MV for 4 mm of copper buildup. Both models agree with the measured data, which is slightly lower than expected. The TOPAS model experiences a drop in absorbed dose at the bottom of the scintillator because the scintillator is simulated suspended in air. The measured signal experiences an increase in light output at the bottom of the scintillator due to reflection.





Figure 3.3: LKH-5 optical response to irradiation at different energies. Images for 2.5 MV and 6 MV are exposure matched with 4 mm of copper buildup. The 225 kV image was taken at a longer exposure, hence its appearance is brighter.




(a) Model optical photon production vs measured with camera.

Figure 3.4: Comparing the theoretical optical photon production with the LKH-5 scintillator at 6 MV to measured data. The continuous analytical model was used.

Monte Carlo

The alternate approach is to simulate the light production within the scintillator by Monte Carlo. Using the volumetric Monte Carlo approach discussed in chapter 2 we simulated flood field images at 225 kV, 2.5 MV, and 6 MV to check the attenuation in the scintillator. The model created for spectral evaluation features an object that consists of solid water, gold, copper, and aluminum. Each material attenuates the beam before reaching the scintillator to allow assessment of material identification. Figure 3.5 shows a wireframe model of the material phantom used in the simulation. The origins of this phantom design can be found in Myronakis' *et al.* work on megavoltage imaging with a multi-layer detector [69]. As shown in Figure 3.5, the gold, copper, and aluminum are 2 mm squares sandwiched between two 5 cm solid water slabs. The 2 mm thickness was an arbitrary selection as previous work on multi-layer megavoltage detectors do not provide their phantom's material thicknesses [69].




(a) TOPAS material phantom geometry

Figure 3.5: TOPAS material phantom. Two 5 cm of solid water and squares of 2 mm of Al, Au, and Cu.

Figure 3.6 displays the simulated dose absorbed curves for LKH-5, assuming no copper

buildup. What is important about Figure 3.6 is the presence of a buildup of dose which would correspond to an increase in optical photons. Unfortunately, experimental images that looked at the scintillator orthogonally at 6 MV consisted of buildup and additional material. Verifying the presence of buildup would be a good model verification for future work.



(a) TOPAS material phantom geometry

Figure 3.6: TOPAS model of dose absorbed in the LKH-5 scintillator

Material Removal

We followed the methodology Rottmann *et al.* used for their multi-layer detector material removal images [29]. The monolithic scintillator signal was broken up into a focal stack of 30 slices, each slice corresponding to a thickness of $733 \mu\text{m}$. Each slice image I was corrected for flood field I_F , dark current $I_D = 0$, and gain G which was the average signal from the flood field image by

$$I = \frac{I - I_D}{I_F - I_D} \cdot G. \quad (3.5)$$

Recall the material phantom being used here has various materials. Therefore, the material removal performance is material dependent. To remove certain material from an image, a weight factor must be determined to create a weighted log-subtraction image. The process of weighted log-subtraction takes the natural log of one focal slice image I_1 and subtracts it from another natural log slice image I_2 multiplied by a weight w , to form a new image I_w which can be expressed as

$$\ln(I_w) = \ln(I_1) - w \cdot \ln(I_2). \quad (3.6)$$

I_1 and I_2 can be any combination of slices or multiple slices combined into one. The optimal weighting factor w is determined based on minimizing,

$$\Delta S = \bar{I}_m - \bar{I}_b, \quad (3.7)$$

where \bar{I}_m and \bar{I}_b are the mean signal from an ROI drawn over the metal of interest and the background material, respectively. Ideally, ΔS should equal zero if the metal is fully removed while leaving the other materials detectable.

3.2.3 Material Decomposition Results

The material phantom at 225 kV, 2.5 MV, and 6 MV is shown in Figure 3.7. Going around the phantom clockwise and starting at the top left, the materials are solid water (background), gold, copper, and aluminum (bone substitute). As expected, there is a noticeable decrease in contrast with increasing energy. We have 30 slices each with an axial thickness of 733 μm , and we explored two approaches to merging the slices into two images to be used to form a weighted log subtracted image. Note these two methods use the voxelized dose array from Monte Carlo and are not true focal stacks reconstructed by a light-field camera. In the first method, only the top and bottom slices are considered, inspired by Rottmann *et al.* who use

only the first and bottom layers to form the weighted image. Given the axial resolution of a light-field camera is on the order of $300 \mu\text{m}$, and from work by Liu *et al.* has shown the axial resolution for a light-field camera is worse for a continuous light source than a discrete light source, it is likely that two reasonable slices, one from the top and one from the bottom half of the scintillator will have the largest difference [55]. Therefore, the 2nd method of merging slices together tries to capture the light-field effect of decreased axial resolution. We summed the first 15 slices to form a top image and the bottom 15 to form a bottom image.




Figure 3.7: The solid water material phantom’s contrast as different energies. The 2.5 MV and 6 MV images are window-level matched to better reflect the decrease in contrast at 6 MV compared to 2.5 MV. Materials are solid water, gold, copper, and aluminum.

Tables 3.1 and 3.2 contain the resulting material removing performance by showing the ratio of metal to the background. The scintillator layer image formation method is expressed as fine or coarse, corresponding to using single or summing slices together. Perfect removal would result in a ratio of 1. Comparing the results to the multi-layered EPID, using a thicker scintillator offers a clear advantage for both 2.5 MV and 6 MV energies. All the multi-layered EPID material to background ratios were reported to be 1 ± 0.005 , whereas we achieved greater material preservation, especially at 2.5 MV. At 6 MV, the improvements are slightly better than a multi-layered EPID, however, in the process of removing a material, the other materials are significantly affected. Based on the performance difference between using single slices versus summing together multiple slices to form the two images used in material

removal, the coarse axial resolution approach offered better performance. Incorporating the additional slices increased the amount of spectral information used. A better weighting scheme could be applied instead of summing the slices to further improve spectral utility.

Table 3.1: Material removal signal to background ratio at 2.5 MV

Removed Material (Resolution)	Aluminum	Copper	Gold
Aluminum (Fine)	1.0	1.017	1.023
Copper (Fine)	0.994	1.0	0.9756
Gold (Fine)	0.9976	1.008	0.999
Aluminum (Coarse)	1.01	1.016	0.791
Copper (Coarse)	0.9995	1.0	0.852
Gold (Coarse)	0.9883	0.96	0.992

Table 3.2: Material removal signal to background ratio at 6 MV

Removed Material (Resolution)	Aluminum	Copper	Gold
Aluminum (Fine)	1.0	1.006	1.001
Copper (Fine)	0.999	1.001	0.986
Gold (Fine)	1.00	1.006	1.001
Aluminum (Coarse)	1.0	0.994	0.949
Copper (Coarse)	1.002	0.9995	0.952
Gold (Coarse)	1.038	1.0864	0.9985

Figure 3.8 shows the associated weighted log-subtracted images, each with a different metal removed. While 2.5 MV and 6 MV are the energies of interest, at 225 kV, due to the majority of the beam being attenuated, aluminum can cleanly be removed, with the other two metals remaining visible. Gold, however, being a high-density metal can not be removed from the images due to significant beam attenuation.

3.2.4 Light-field Material Decomposition Results

In chapter 2, we went into great detail about how a light-field camera can remove out-of-focus light and improve the MTF over a conventional camera. However, having a high resolution in the axial direction is quite challenging due to the limited parallax information




Figure 3.8: Material removal using coarse axial resolution. The material removal for 225 kV, 2.5 MV, and 6 MV are the top, middle, and bottom rows, respectively.

that the camera can collect and the continuous light source. The material phantom from the previous section was computationally refocused with a simulated Lytro Illum where the entire camera configuration can be set and another light-field camera with an $F/0.95$ aperture. Five hundred iterations of ML-EM were performed to reconstruct the phantom at

225 kV, 2.5 MV, and 6 MV.

Figure 3.9 shows slices from a reconstructed focal stack of the material phantom. As can be seen from the slices, the light-field camera can refocus the slices. However, in regions of low variation, such as a section over the solid water, the observable axial variation will only be a couple of percent, shown in Figure 3.10. While only a couple of percent, the axial curves still captured the exponential response of the scintillator at 225 kV. As will be seen from the physical experiments in chapter 4, where the integration approach was used to generate focal stacks, those also only experienced a couple of percent variation axially, which is a good agreement with simulation. In the comparison between the simulated Lytro Illum and the F/0.95 camera, increasing the aperture does improve the axial resolution near the top of the scintillator. Towards the bottom of the scintillator, the larger aperture camera amplifies the noise. In addition, not all regions experience little axial variation. For example, the edges of the scintillator experience significant light changes throughout the focal stack at 225 kV due to the edge being a high-frequency object.

3.3 Cramér-Rao Lower Bound Analysis

Previously, we looked at material removal based on the availability of spectral information. Recall that the material phantom we used to show material removal was structured in layers: solid water, metal, then solid water. The incident X-ray spectrum on the scintillator contains information that captures how much material and what types of material the beam went through. If perfect detection existed, a close estimate of the materials and thicknesses could be made. However, X-ray interactions are probabilistic. Given that we are imaging with a light-field camera that can only take a finite set of measurements, there will always be a variance in predicting deterministic quantities such as the material thickness. Since the thickness of a material is a deterministic value, when we try to estimate it, the precision of an unbiased estimator is, at most, the Fisher information \mathcal{F} . This statement can also be




Figure 3.9: Reconstructed focal stack slices of the material phantom at 225 kV, 2.5 MV, and 6 MV separated by row. Focus 3 is the focus at the top of the scintillator, then it goes focus 2, then 1 with increasing depths.

stated as the lower bound on the variance is the inverse of the Fisher information $\sigma^2 \geq \mathcal{F}^{-1}$. Finding the lower bound on an unbiased estimator's variance is called Cramér-Rao Lower Bound Analysis (CRLB) [68].




Figure 3.10: Axial response for the Lytro Illum vs F/0.95 light-field camera

The setup we wish to find the CRLB variance for is an X-ray beam that passes through soft tissue and bone with thicknesses t_1 and t_2 and is detected by the scintillator imaged by a light-field camera, shown in Figure 3.11. Here z_{max} refers to the distance from the focus plane of the camera, which is set to the center of the scintillator. Recall refocusing can be expressed as an integral over the main lens (u, v) and microlens (s, t) planes by

$$L(u_i, v_i, s_i, t_i) = \int_{-z_{max}}^{z_{max}} \delta \left[\left(1 - \frac{z}{z_0}\right) u_i - M s_i \left(1 - \frac{z}{z_0}\right) v_i - M t_i \right] N(z; A_1, A_2) dz, \quad (3.8)$$

where L is light-field samples from an image scene, M is the camera's magnification, z_0 is the camera's focus distance, and $N(z; A_1, A_2)$ is from equation 3.4, where we modeled the light output from the scintillator as a continuous function. However, now the beam has passed through materials A_1 and A_2 , which indicate soft tissue and bone, respectively. Let $A = \rho t$, where ρ is a material's density, and t is a material's thickness. This expression allows for the mass absorption coefficient to be used instead when modeling the attenuation through




Figure 3.11: The CRLB setup. A delta X-ray beam passes through soft tissue and bone and is detected by a light-field camera.

materials.

Only a delta ray will be modeled to make the computation easier, reducing the integral complexity. The measured light field depends on the size of the X-ray beam w^2 . In a single image acquisition, each pixel is a sample i , therefore the total measurement can be written as

$$\bar{y}_i = L(u_i, v_i, s_i, t_i)w^2. \quad (3.9)$$

The elements in the Fisher information matrix for an energy-integrating detector, assuming a Gaussian noise model, can be calculated from

$$\mathcal{F}_{jk} = E \left[-\frac{\partial^2 \mathcal{L}}{\partial A_j \partial A_k} \right] = \sum_{i=1}^M \frac{1}{\sigma^2} \frac{\partial \bar{y}_i}{\partial A_j} \frac{\partial \bar{y}_i}{\partial A_k} + \frac{1}{2} \sum_{i=1}^M \frac{1}{(\sigma_i^2)^2} \frac{\partial \sigma_i^2}{\partial A_j} \frac{\partial \sigma_i^2}{\partial A_k}. \quad (3.10)$$

Breaking equation 3.10 down into parts [70]. The Fisher elements are equal to taking the partial differentials in respect to the materials A_1, A_2 , which refer to soft tissue and bone,

for both the measurement \bar{y}_i and the variance of the measurement σ_i^2 , which is expressed by

$$\sigma_i^2 = w_b^2 I_0 k^2 \int_{-z_{max}}^{z_{max}} p_i^2(z) \int E^2 Z(E) \mu_s(E) T(E; A_1, A_2) e^{-\mu_s(E)(z+z_{max})} dE dz + \sigma_e^2. \quad (3.11)$$

σ_i^2 is almost identical to the \bar{y}_i expression, however in σ_i^2 the terms of $k, E, p_i(z)$ as scalars get squared in a variance calculation, and the remaining non-squared terms are the mean of a Poisson process. $p_i(z)$ represents the light-field integration integral that maps the volume to the sensor image. $p_i(z)$ effectively captures each voxel's weight in the forward projection. We assumed zero electronic noise, therefore $\sigma_e^2 = 0$.

With the Fisher matrix computed, the standard deviation for A_1, A_2 can be computed by the inverse of the Fisher. A signal-to-noise ratio $A_j/\sqrt{\sigma_j^2}$ is used to assess the prediction ability of estimating the thickness of soft tissue and bone.

3.3.1 CRLB Results

The CRLB was analyzed for 225 kV, 2.5 MV, and 6 MV for the Lytro Illum, and a more optimal light-field camera with an F/0.95 aperture placed in the same configuration as the Lytro Illum. To determine the fluence I_0 for each energy an imaging dose of 1 cGy was assumed. Therefore, the incident fluences were 7.3×10^{10} , 2.38×10^9 , and 2×10^9 for 225 kV, 2.5 MV, and 6 MV, respectively. The beam area w_b was selected to match the highest resolution achievable given the microlens diameter and magnification, which resulted in $w_b = .125$ mm. Figure 3.12 displays a couple of viewpoint images used in the CRLB analysis. Even though the beam has an area of w_b^2 , the beam appears as a delta beam based on the light-field camera's resolution.

Table 3.3 has the resulting figures of merit $A_j/\sqrt{\sigma_j^2}$ for 5 cm of bone, and 40 cm of soft tissue for the Lytro Illum. A density of 1 g/cm^3 was assumed for soft tissue, and 1.85 g/cm^3




Figure 3.12: A sample of the viewpoint images used to compute the CRLB.

for bone. As expected, as energy increases the ability to identify material decreases. The dropoff in the figure of merit from 225 kV where the entire X-ray beam is attenuated to 2.5 MV is significant. Table 3.4 has the figures of merit for a better light-field camera. From the previous section, it was shown that the axial resolution increase due to the F/0.95 is marginal, likewise, the gains in CRLB with a better light-field camera are marginal. The best way to improve the CRLB is by increasing the scintillator thickness to capture more of the beam.

Table 3.3: CRLB figures of merit $A/\sqrt{\sigma^2}$ for Lytro Illum (F/2).

Material	225 kV	2.5 MV	6 MV
Soft tissue	4792	17.49	16.36
Bone	1165	3.879	3.663

Table 3.4: CRLB figures of merit $A/\sqrt{\sigma^2}$ for an F/0.95 light-field camera.

Material	225 kV	2.5 MV	6 MV
Soft tissue	5680	19.11	18.09
Bone	1383	4.23	4.05

3.4 Conclusion

Imaging a thick monolithic scintillator and extracting spectral information with a light-field camera is problematic in most instances, with some areas of hope for future research. While the evidence that a light-field camera will be able to fully resolve a 3D volumetric scintillator being irradiated at 2.5 MV and 6 MV at this point is weak, ML-EM was the only iterative algorithm fully explored. More advanced algorithms may provide better noise regulations as controlling noise is imperative when the light variation through the scintillator is less than 20% at 6 MV. In addition, if the spectral information could be extracted with a light-field camera, from simulations, it is clear that there is useful information within the scintillator that would allow for better material removal than the multi-layer EPID.

CHAPTER 4

LIGHT-FIELD IMAGER EXPERIMENTS

We performed a variety of experiments to assess whether an inexpensive prosumer light-field camera can image a thick monolithic scintillator. Currently, in the literature there is little to no data about scintillator light-field imaging other than from a research group in Canada that used a Raytrix camera to image a large 20 x 20 x 20 cm plastic scintillator for 3D dose verification [20]. However, what does exist in the literature suggests that imaging a continuous light source like a scintillator with a single light-field camera will be challenging [55].

Our experiments can be broken into three categories: exploratory, DQE imaging metrics, and spectral extraction. The light-field camera we used was the Lytro Illum; however, Lytro was absorbed by Google [71]. Resulting in very little information being available on detailed camera specifications. Therefore, many experiments were done just to gain insight into the internal camera operations. With a basic understanding of the camera workings, we moved on to assessing the resolutions (spectral & image) of the light-field imager at 225 kV, 2.5 MV, and 6 MV energies. In addition, we took flood-field and dark current images to calculate the noise power spectrum to calculate the DQE. The last thing we looked at with the Lytro Illum is the potential to extract spectral information from the reconstructed focal stack.

4.1 Equipment

This study used a Lytro Illum standard plenoptic camera as it is relatively inexpensive compared to the available focused plenoptic cameras. Table 4.1 provides an overview of the important camera specifications for the Lytro Illum [24, 59]. The conventional camera used was a CoolSNAP HQ², and important camera specifications for it can be found in Table 4.2 [60]. The camera lens used on the CoolSNAP for imaging was an AF-S DX Micro

NIKKOR 40mm f/2.8G [61]. The aperture of the lens was placed wide open to the F/2.8 setting. Figure 4.1 shows the cameras with the NIKKOR lens attached to the CoolSNAP. The NIKKOR lens is designed for an APS-C sensor size, which is smaller than a full-frame sensor. As a result, the listed 40 mm focal length on the lens has an equivalent focal length of 60 mm on a full-frame sensor and the equivalent of a 15.53 mm focal length lens on the CoolSNAP which has an even smaller sensor than the APS-C format.



(a) Conventional Camera (CoolSNAP with 40mm F/2.8 lens)



(b) Light-Field Camera (Lytro Illum)

Figure 4.1: Cameras used during experiment

Both cameras can be controlled remotely via a computer in another room, which is important to avoid exposure to scattered X-rays. The CoolSNAP, a rather dated camera, can be controlled by an IEEE 1394a FireWire connection to a Windows XP/Vista 32 desktop

Table 4.1: Lytro Illum Specifications

Main Lens type	Built-in
MLA arrangement	Hexagonal
MLA diameter (pixels)	14
f/number	2
focal length	30-250 mm (35 mm equivalent)
Sensor Type	CMOS
Chroma	RGB
Sensor size	1/2 in.
Sensor dimension	7728 x 5368 pixels
Pixel size	1.4 μm
ISO	80-3200
Max shutter speed	1/4000 sec
Shutter type	Rolling
Max Resolution using Lytro software	2450 x 1634
Standard Resolution	570 x 375

Table 4.2: CoolSNAP Camera Specifications

Main Lens type	None
Sensor type	CCD
Chroma	None
Readout	90 frames/ms
Resolution using Lytro software	1392 x 1040
Pixel size	6.45 μm

running Micro-Manager, an ImageJ-based software that allows remote camera control and immediate image processing through ImageJ [72]. The Illum can be remote-controlled via either a Mac or Windows computer. The Illum can be run as an Android device and controlled through Lytro’s Power Tools beta python package [73]. As mentioned, Lytro is rather a black box camera and while the power tools outputs images of the MLA, the images have already been colorized and the MLA rectified into a rectangular grid. While the rectified MLA images are easier to work with, the recorded light intensity has been gamma corrected to represent brightness in the way humans perceive brightness. In photography, the linear light intensity as recorded by the camera is converted to an image that we observe on the camera screen by

$$I_{out} = cI_{int}^{\gamma}, \quad (4.1)$$

where I represents the light intensity, c a constant, and γ which is largely based on the manufacturer's color science and may or may not be recorded by the camera. The Lytro Illum does not report the used gamma value. Therefore, a way to semi-convert the data back to linear intensity is by $I_{out} = 2^{Dx}$ where D is the bit-depth of the sensor and x is the gamma-affected data normalized to 0 to 1. Understanding the color science employed by the Illum is important, as shown in Figure 4.2 depending on the scintillator's optical light wavelength, it may only be detected by half the pixels if it is green light or a quarter of the pixels if it is blue or red. In the case of LKH-5, its primary emission is 545-550 nm, allowing it to be detected by the green pixels in the Bayer filter. We used the gamma-affected images colorized using Lytro's in-house color science due to the superior rectification of converting the hexagonal MLA array to a rectangular arranged array.

By default, the Illum's output pixel resolution is determined by the number of microlens, which results in the pictures having a resolution of 570 x 375 pixels. However, depending on the software, for exemplifying Lytro's in-house tools, the images can be outputted as 2450 x 1634 pixels. Light-field focal stacks were reconstructed with Ren Ng's integration method [21], which was implemented in Python by the PLENOPTOMOS package [23]. With the focal stacks, an extended-depth-of-field (EDOF) image was computed by using a Sobel operator to select the most in-focus pixels throughout stack [16] to compare against the most in-focus focal stack slice.

Images were taken in an XRAD 225Cx small animal irradiator and a Varian TrueBeam LINAC. On the XRAD, the scintillator was primarily imaged at 225 kV. On the LINAC, 2.5 MV and 6 MV were the primary energies explored. The dose rate was never calibrated for the XRAD, however, at 225 kV, we used the highest current of 13 mA, which based on machine specifications, likely resulted in a dose rate greater than 400 cGy/min. On the other






Figure 4.2: Raw image from Lytro Illum sensor. The MLA array found in the Lytro is hexagonally arranged and there is an RGB Bayer color filter above the sensor pixels. Image is of a 3D printed copper-doped phantom being imaged at 6 MV.

hand, the LINAC was calibrated to deliver 0.85 cGy per monitor unit (MU) to water at a depth of 10 cm in a 10 by 10 cm field. The imager prototype was placed on the treatment couch with a source-to-scintillator surface distance of 100 cm. Field sizes for the scintillator measurements were reduced to the scintillator's surface area of (3.4 x 2 cm). The maximum dose rate was used for all imaging situations. On the LINAC, the maximum dose rates were 60 MU/min and 600 MU/min for 2.5 MV and 6 MV, respectively

Figure 4.3 shows the setup components: the scintillator, mirror, lead box, and camera. In the small animal irradiator, external light can easily be shielded with tin foil. However, for the LINAC experiments, a large cardboard box was placed over the imager system to help shield it against treatment room lights that could not be covered like the treatment room emergency indicator lights. The scintillator itself was covered with black construction paper to prevent objects on top of the scintillator from being visible to the camera.




(a) Physical setup. Arrows indicate the placement of the setup inside a light-shield box and the position of the scintillator.



(b) Setup cross-section

Figure 4.3: The experimental setup. The camera is placed inside a lead housing out of the X-ray beam. The X-ray beam is incident on the LKH-5 suspended above a front surface mirror, reflecting the optical photons to the camera. The light-field imager was placed in a lightbox for the LINAC experiments. The top of the scintillator was covered with black construction paper. On top of the black paper the phantoms would be placed.

The cameras were always focused to the center of the scintillator. The CoolSNAP lacked motorized control of the focus and had to be set by hand. The Lytro Illum did have focus control that the user could specify; however, it was easier to allow the Illum to auto-focus itself to the center. Since the scintillator is transparent, we used a 3D printed phantom that consisted of pillars that were 3 mm in diameter and placed at different depths. Figure 4.4 showcase the 3D-printed phantom with both cameras focused on the center pillar as much as possible. It can be seen in Figure 4.4 that the conventional camera had a wide depth-of-field that encompasses most of the pillars to some degree. The 3D-printed plastic has a marble rock texture that gives the plastic a texture to observe focus effects more clearly. The closest pillar is located on the left side of the images in Figure 4.4 and corresponds to the base of the scintillator. Similarly, the thicker rightmost pillar corresponds to the top of the scintillator. The axial distance between the nearest and furthest pillar was 2.3 cm. At the focus plane, it was measured that the conventional camera had a resolution of $36 \mu\text{m}$ and the Lytro had a resolution of $136 \mu\text{m}$.



(a) CoolSNAP focal plane

(b) Lytro camera focal plane

Figure 4.4: Resolution at the focal plane. The ruler in the lower righthand of the images indicates a distance of 3 mm.

4.2 DQE experiments

4.2.1 DQE procedure

The definitions of detective quantum efficiency (DQE), modulation transfer function (MTF), photon fluence per unit area q , noise power spectrum (NPS), and the normalized noise power spectrum (NNPS) were presented in chapter 1 as taken from Star-Lack [6]; however, they are repeated here for clarity,

$$DQE(f) = \frac{|MTF(f)|^2}{q \cdot nNPS(f)}. \quad (4.2)$$

At kV energies, the DQE measurement protocol has been standardized by the International Electrotechnical Commission [8]. Given that our light-field imager is a prototype, the procedure had to be adjusted somewhat, especially at megavoltage energies. At megavoltage energies, calculating the MTF is not as well established. In kilovoltage energies, a simple slab of tungsten, as shown in Figure 4.5 can be used to calculate the MTF. Figure 4.5 is the edge we used for kV and MV MTF calculations. The edge itself is 2 cm of tungsten. In megavoltage energies, doing an MTF by a slit would be ideal; however, that is incredibly difficult to do experimentally over such a small scintillator. Using a 2 cm tungsten will result in some error, especially in the megavoltage due to insufficient attenuation. From the literature, 4 cm of tungsten is the ideal compromise between sufficient attenuation and not having an edge too thick that it causes geometric effects. Therefore, for a 2 cm piece of tungsten, we expect the MTF to be overestimated at 6 MV. For 225 kV, the MTF error should be less than a percent [74].

To calculate the NPS, 110 flood-field and 110 dark-field images were taken. Flood-field images refer to images where the X-ray source is on and nothing is placed on the scintillator. Dark-field images refer to background noise images with the X-ray source off. Ideally, dark-field images should be captured before and after the flood-field images; however, whether





Figure 4.5: Tungsten Edge used for MTF calculations

that is done depends on if the camera sensor is stable or if the scintillator has an afterglow.

The linearized flood-field images $I(x, y)$ were dark-current corrected by subtracting out the average dark-field image. Depending on the approach, the flood-field images can be zero-meaned by subtracting two flood-field frames together. The other zero-mean approach uses a 2nd-order polynomial to remove any flood-field beam gradients, followed by calculating the zero mean by subtracting each detrended frame by the average detrended image. $S_g(x, y)$ represents the nonstochastic signal removed by detrending and averaging. We used the 2nd-order polynomial detrend approach. Figure 4.6 shows a flood field image and an ROI undergoing zero meaning for the NPS calculation.

Due to the limited size of the Lytro images, only two partially overlapping regions of interest (ROIs) of 128x128 pixels could be placed over the scintillator for analysis. For the CoolSNAP images, ROIs were 256x256 pixels wide and were broken up into 12 non-overlapping ROIs; however, the NPS calculation allows up to 50% overlap to maximize the pixels used in the analysis. From the square modulus of the Fourier transform of the ROIs, the NPS was calculated using equation 4.3 where N_{fl} refers to the number of ROI images,




(a) Flood-field image




(b) Flood-field ROI



(c) Detrended ROI



(d) Detrended Mean ROI



(e) Final zero-mean ROI

Figure 4.6: NPS image processing

$N_x = N_y$ the dimensions of the ROIs, and $\Delta x = \Delta y$ the pixel resolution at the focal plane:

$$NPS(f_x, f_y) = \frac{\Delta x \Delta y}{N_x N_y N_{fl}} \sum_{g=1}^{N_{fl}} \left| \sum_{n=1}^{N_x} \sum_{m=1}^{N_y} [I_g(n, m) - S_g(n, m)] e^{\frac{-i2\pi(nf_x + mf_y)}{N_x N_y}} \right|^2. \quad (4.3)$$

To verify that the NPS was calculated correctly, the NPS was compared against the variance of the zero-mean ROIs via

$$\sigma^2 = \int \int NPS(f_x, f_y) df_x df_y. \quad (4.4)$$

The NPS was normalized by dividing it by the square of the mean signal found in the linearized flood-field images before zero-meaning. The flood-field mean signal is sometimes called the 'large area' signal in DQE literature [6]. To convert the 2D-NNPS to directional 1D NNPS, the frequencies within ten rows or columns of each axis were binned such that for a frequency (f), $f = \sqrt{f_x^2 + f_y^2}$, it would be averaged into a bin if it fell in the range of $f_{bin} = f \pm \frac{0.01}{\Delta pixel}$ where $\Delta pixel$ is the pixel resolution at the focal plane.

In experiment, the NPS being measured is the cumulative sum of different noise sources: quantum noise, electronic noise, thermal effects, and fixed pattern noise [75]. The CoolSNAP camera is a scientific camera with active cooling. In contrast, the Lytro is a prosumer camera that lacks the sensor systems required to provide consistent repeated measurements without significant heat up. To quantify the electronic noise on the Lytro Illum equation 4.3 was used to calculate the NPS for the dark-field images. Pairs of dark-field images were subtracted and the NPS was calculated for each pair. Since electronic readout between dark-field frames is uncorrelated when dark-field frames are subtracted the NPS will yield a value double what would be observed in a single frame, therefore the result was divided by two [38].

4.2.2 DQE Results

The measured MTFs for the CoolSNAP and Lytro Illum are presented in Figure 4.7. As expected, with increasing X-ray energy, the resolution decreased. The MTFs at 225 kV were compared against a diagnostic bar phantom, shown in Figure 4.8, where the 2.00 lp/mm is visible on the Lytro Illum. There is some uncertainty with the MTF calculation as the resolution at the focus plane has to be manually calculated for both cameras and can result

in a shift in the MTF if a pixel or two off.




Figure 4.7: MTFs for the Lytro Illum and for the conventional camera

The conventional camera could read significantly higher resolutions at 225 kV; however, it depended on focus. If the focus was slightly closer to the top, then the resolution dramatically improved, but if it was slightly below the center of the scintillator, the resolution was similar to the Lytro’s. Recall that Lytro’s images were being refocused with the integration method, which does not remove out-of-focused light; therefore, each focal slice would have an out-of-focus light blur of an $F/2$ camera. The conventional camera had an $F/2.8$ lens meaning the out-of-focus light blur would be smaller than the Lytro Illum. The conventional camera also has 6.45 micron pixels compared to the Lytro’s 20 micron microlenses giving the conventional camera a higher resolution. Another factor that may have helped the conventional camera at 225 kV is the lack of a Bayer filter. LKH-5 primarily emits at 550 nm, however, it does emit red optical photons, which seem to vary with the incident X-ray energy. In particular, at 225 kV the bar pattern is more visible using only the red channel on the Lytro vs when using the green channel shown in Figure 4.8. The green channel is the primary channel used in this study as the red and blue channels are generally quite noisy.

The 2D NNPS for the Lytro Illum and the CoolSNAP are shown in 4.9 for 6 MV.





Figure 4.8: LKH-5 resolution differences based on the color channel on the Lytro Illum

Recall that in the NNPS calculation, the axes frequencies are ignored as the 2D detrending removes those frequencies, leading to the observed lack of signal on the axes. The Lytro Illum experiences noise equally in both the X and Y directions, whereas the CoolSNAP experiences more noise in the X direction than the Y direction. The 1D directional NNPS graphs for 6 MV are shown in Figure 4.9. Unsurprisingly, from the NNPS results, it is clear that the sensors read row by row vertically, as shown in Figure 4.10 where banding is observed from the subtraction of two dark-field frames.


The reported NNPS results are from one-second exposures. In a quantum-noise limited system, the NNPS when multiplied by the photon fluence q should not change with increasing exposures. Future research should be done to explore if one-second exposures were quantum-limited; however, for one-second exposures, 10 MU of dose would have been delivered, and from literature, many megavoltage imagers become quantum-noise limited past 2 MU [38]. In addition, at one-second and longer exposures, electronic noise becomes a significant concern. Shown in Figure 4.11, the Lytro Illum has an electronic NPS one hundred times noisier than the CoolSNAP. Despite being one hundred times noisier than the CoolSNAP, the effect of electronic noise is less than 1% from a dark-field analysis before the camera heats up. The Lytro's and CoolSNAP's electronic noise shapes were as expected. The Lytro Illum's electronic noise was from a reconstructed focal stack slice corresponding to the most in-focus

MTF measurement. This results in the electronic noise having a structure. In contrast, the CoolSNAP's electronic noise had the standard white noise shape.




(a) Lytro 2D NNPS

(b) CoolSNAP 2D NNPS




(c) Lytro 1D NNPS



(d) CoolSNAP 1D NNPS

Figure 4.9: 2D NNPS for Lytro and CoolSNAP cameras. X and Y directional NNPS.


Without knowing the fluence of the XRAD or LINAC, it is tricky to calculate the actual DQE. The large difference in DQE for the CoolSNAP and the Lytro Illum in Figure 4.12 is the result of the camera distances being different and the cameras not being lens matched. It was thought that both cameras were placed 20 cm away from the scintillator based on the camera body position, however, for the CoolSNAP the optics and lens ended up being




(a) Lytro noise banding

(b) CoolSNAP noise banding

Figure 4.10: Noise banding from dark-frame subtraction analysis



(a) Lytro



(b) CoolSNAP

Figure 4.11: Electronic noise.

several centimeters closer leading to a higher optical coupling. In simulations, it was shown that for F/2 matched lenses a conventional camera would outperform the LF camera and here the CoolSNAP had an F/2.8 lens. To approximate what the DQE(0) could be for the LINAC at 6 MV requires approximations from the literature. Star-Lack *et al.* report a KERMA dose equivalent factor of 1.42×10^7 photons/(cGy-mm²) for a 10 by 10 cm² field.

Our machine had a 0.85 cGy per MU calibration. Assuming an output factor reduction in the dose observed in our small field of 0.88 [76] we can reduce the 0.85 cGy/MU by that factor. Putting it together, the photon fluence per mm^2 is estimated to be

$$600 \frac{\text{MU}}{\text{min}} \times 1/60 \frac{\text{min}}{\text{sec}} \times 1.42 * 10^7 \frac{\text{cGy}}{\text{MU} - \text{mm}^2} \times 0.85 \frac{\text{cGy}}{\text{MU}} \times 0.88 \text{output} = 10.62 * 10^7 \frac{\text{ph}}{\text{mm}^2} .$$

From simulation, we expected to see a DQE(0) for the conventional camera around 16.1% and we measured a DQE(0) of 15.9% shown in Figure 4.12. For the Lytro Illum, we expected a DQE(0) of 5.3%, and we measured a DQE(0) of 0.4%. The gap between expected and measured for the Lytro Illum stems partly from the Bayer filter blocking 50% of the pixels as LKH-5's primary emission is green. This would reduce the expected DQE(0) to 2.65%. The remaining gap likely comes from poor lenses since the Lytro is only a prosumer camera and unquantified electronic noise. The CoolSNAP's DQE agrees well at DQE(0) and between 0.5-1 lp/mm. The noise in the CoolSNAP's DQE comes from the MTF measurement as the scintillator has several scratches that cannot be filtered out, and if there is a scratch on either the top or bottom surface, it will be present in the MTF calculation.

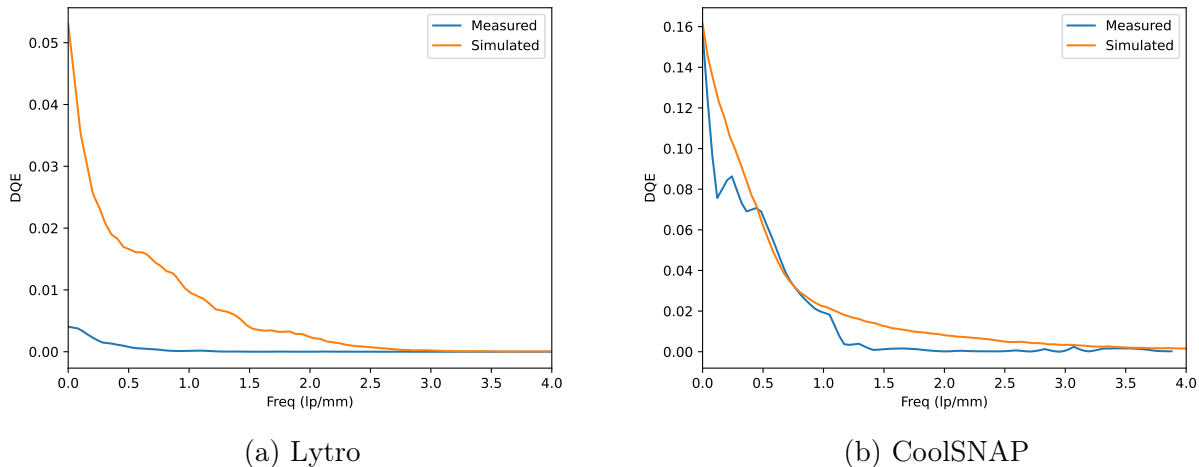


Figure 4.12: 6 MV DQE

4.3 Spectral Extraction Experiments

The integration approach works well for refocusing with little knowledge of the light-field camera’s internal design. If the focus plane can be estimated, which it can be from the MLA pattern, refocusing just requires knowing the relative difference between the focus plane and the refocus plane as input to the integration refocus algorithm. However, in order to accurately remove the out-of-focus light, an iterative algorithm is required. In our simulations we implemented ML-EM to reconstruct the 3D scenes, so we tried to apply it to the Lytro Illum data. Figure 4.13 showcases the rapid increase in noise using ML-EM on the Lytro Illum data.



Figure 4.13: Refocusing on Lytro Illum data with ML-EM

The fact that the Lytro Illum is largely still a black box is problematic, especially when using more advanced algorithms such as ML-EM that rely on the accuracy of the imaging system matrices to improve the resulting reconstructions. It is possible to extract enough information from the Lytro images and their JSON files to get the camera to refocus over the image volume without having to do a guess and check. However, parameters which have an effect on the system matrices, like the aperture width, which is calculated from the F-number, and the lens focal length, are still somewhat ambiguous on the Lytro Illum. The camera advertises a constant F/2 aperture over the zoom range, however, the MLA is fixed. Therefore, it is unclear if they are preserving the aperture in the traditional sense $F/\# = f/D$ where f is the focal length and D is the lens diameter or if they are preserving the perceived aperture to the MLA array. An F/2 lens may be F/2 at infinity focus; however,

for finite focus there would be a higher F-number. The aperture dilemma is further convolved by the reported value in the image JSON files being $F/2.24$, not the advertised $F/2$. Given the poor performance of ML-EM, our spectral extraction results are reported by using the integrated approach.

Depth estimation tasks require that there be some frequency changes that occur within the image scene. For example, if a sheet of paper was uniformly lit and angled slightly, a light-field camera would not be able to tell which end of the paper is closer or farther. In the case of the scintillator, if a blank scintillator with nothing on it is imaged and then computationally refocused over the volume, there would be little to no axial variation seen in the focal stack. Figure 4.14 displays a kilovoltage bar pattern being imaged at 225 kV. The image shown is a slice from the focal stack with the bar pattern in the sharpest focus. At 225 kV, LKH-5 is able to completely attenuate the X-ray beam; therefore most of the light is coming from the top half of the scintillator. In Figure 4.14 different regions of interest are highlighted and their mean pixel value is calculated per slice, which is shown in the graph. For the regions of interest with little variation, the axial resolution is basically zero. For the region that contains high frequency, computationally refocusing via the integration method will blur high frequencies resulting in a noticeable variation in the axial signal. However, it is not a true signal representation of the light production within the scintillator.

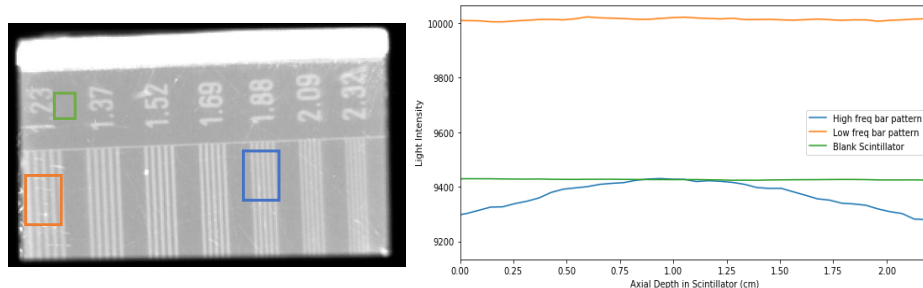


Figure 4.14: Axial signal dependence on frequency in the transverse plane. Bar pattern image is from a focal stack slice from the Lytro Illum at 225 kV.

One approach was made to attempt a better 3D deconvolution of the Lytro Illum data.

The approach relied on imaging a $5\ \mu\text{m}$ pinhole that was illuminated by a green LED. The pinhole was imaged in a replicated scintillator imaging geometry. The pinhole was then computationally refocused to form a 3D PSF focal stack. The axial blur of the 3D PSF can be seen in Figure 4.15. The benefit of imaging a 3D point source is that any unknown camera parameters are also captured in the PSF images. The 3D deconvolution was done with DeconvolutionLab2, an open-source plugin available for ImageJ and developed by EPFL in Switzerland [77]. An example of the improved contrast after factoring in image window adjustments is shown in Figure 4.15.





Figure 4.15: Measured axial point spread function for a $5\ \mu\text{m}$ point source. The 3D deconvolution result with the measured PSF is shown with the original bar pattern image.

We performed an additional experiment where different materials were placed on top of the bar pattern. As discussed, high-frequency objects tend to have the best axial resolution. We found that $1.88\ \text{lp/mm}$ bars had the greatest axial variation out of the different line pairs. As a result, we compared the different material responses on a $1.88\ \text{lp/mm}$ line. Shown in Figure 4.16 is the effect of deconvolution on the materials' axial response. While not a true

representation of the light produced within the scintillator, each signal does have a different profile shape and amplitude that could be used in a rudimentary material identification task.



(a) Axial slice through focal stack before deconvolution. Normalized to blank scintillator. (b) Axial slice through focal stack after deconvolution. Normalized to blank scintillator.

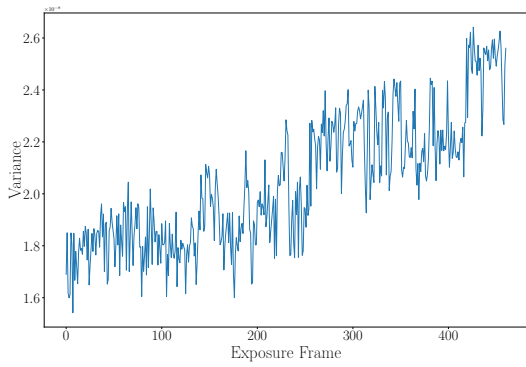
Figure 4.16: Material axial responses imaged on the Lytro Illum. It is important to note that LKH-5 has a higher index of refraction than air, therefore light will appear 36.7% closer to the camera than it really is. These axial responses are not corrected for refraction, therefore the light appears to come from the middle of the scintillator rather than the top.

4.4 General Experiment Observations

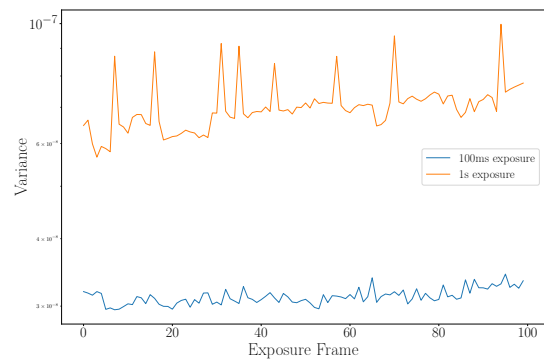
Given the black-box nature of the Lytro, it is hard to quantify the noise precisely. The camera being a novelty device lacks many of the required features for low-noise performance imaging. The most notable missing feature is sensor cooling. In the DQE experiments where the camera sits in a lead box, after finishing the Lytro is hot to the touch, suggesting temperatures greater than 100 degrees Fahrenheit. The issue is further compounded by the small pixel sizes, which require a second or more of exposure time per image, exacerbating the heating problem. Rottmann *et al.* [29] noted a 70% saturation level was used to make the

electronic noise negligible compared to the photon-shot noise, which was orders of magnitude higher than what was done here. To explore this thermal effect further, a couple hundred dark-current images were taken at 500 ms exposure with both cameras on an open-air table. The noise variance responses for the open-air experiment are shown in Figure 4.17 (a) and (c) for the Illum and CoolSNAP, respectively. The Lytro was warm to the touch after completing the couple hundred images; however not as hot as in the shielding. Figure 4.17 (b) is the Lytro's noise variance which occurred during the background readings from the DQE experiment where the Illum would have been in the lead shielding. The 100 ms exposure was the first set of background images, followed by the 1s background images. The flood-field images for 100ms/1s were taken after the background exposure measurements. The CoolSNAP in open air could cool itself, effectively lowering the electronic noise. The order of image acquisitions in Figure 4.17 (d) occurred in the same order as the Lytro: 100 ms, then 1 s exposures, then the flood-field images for DQE calculation.


Besides the variability of one of the cameras being a prosumer non-scientific camera, the scintillator itself is an experimental scintillator. Figure 4.18 showcases how LKH-5 responds to one minute of exposure at the maximum dose rate at 6 MV from a TrueBeam LINAC. The light output of the scintillator is quite stable over the exposure time, however, the afterglow is significant. We didn't measure the full decay time of the scintillator's glow, however, even after several seconds the scintillator was still emitting light a couple of times higher than its base signal. This afterglow effect is not entirely surprising as it is well documented for terbium-doped scintillators [78]. Besides the afterglow, throughout this project, the scintillator has noticeably yellowed compared to its original clear state. Montague *et al.* have shown that LKH-5 can be annealed to help preserve the longevity of the scintillator and reduce radiation damage. [79]. One final discovery about LKH-5 is the post-irradiation activation in beams greater than 6 MV [79]. While we only irradiated LKH-5 once above 15 MV, it has been noted by others that LKH-5 becomes activated enough that it should




(a) Open-air Lytro Illum



(b) Lead Shielding Lytro Illum



(c) Open-air CoolSNAP



(d) Lead Shielding CoolSNAP

Figure 4.17: Potential thermal variance dependence for Lytro and conventional camera

not be handled for 15 minutes, which is concerning in terms of patient and clinician dose exposure [79]




Figure 4.18: LKH-5 scintillation response to one minute of exposure of 6 MV at LINAC max dose rate.

4.5 Conclusion

Performing the light-field imager experiments was important to distinguish the differences between simulation and reality. We were able to accurately model the CoolSNAP camera indicating a successful forward model; however, the Lytro was unable to be effectively modeled. The Lytro Illum did not offer much improvement in terms of resolution given the difficult nature of accurately modeling the camera’s system matrices. In addition, the poor modeling made the spectral ability of the Lytro Illum marginal. However signals can be generated for different materials, as seen in Figure 4.16, but whether the signals can be utilized for accurate material identification is left for future work.

An additional area of research that was ignored but may improve experimental data is having a better data processing pipeline. Currently, images are taken with the Illum, and

then an MLA image is extracted using Lytro's power tools. The power tools do not correct for hot pixels, dead pixels, vignetting, or even using accurate color science. While we attempted to do hot pixel correction by looking at the surrounding pixels and replacing them with the average using the micro-lens image, it might be better to reconstruct the scene with the hot pixels and replace them based on the scene's context rather than using the neighboring pixels [80]. However, this is left up to future work due to the potential of having a better light-field camera where access to the manufacturer may be possible for better data processing.

CHAPTER 5

SUMMARY & CONCLUSIONS

Light-field imaging of a thick monolithic scintillator is a novel application. What we have shown through simulation and experiments is that using a single light-field camera to capture a continuous light field is an inherently ill-posed problem that requires further research to translate it into a clinical setting.

5.1 DQE Summary & Conclusions

The simulations we did and experiments we performed with the Lytro Illum showed that the conventional camera would have a better DQE(0) than this light-field camera. There are still several aspects of the simulations that should be refined. As we saw experimentally, the Lytro Illum performed a tenth of what was predicted when it came to DQE. While some of the factors that limited the Lytro Illum could be improved, such as removing the color filter, and ensuring the sensor is cooled, things like the attenuation from the microlens cannot be easily solved, and the light-field camera will have to compensate by ensuring noise management is incorporated into the reconstruction algorithm.

The role of having good imaging processing software cannot be understated with light-field photography. Besides Lytro's beta power tools to interface with the camera, all the software used in this study is open-source for researchers, students, and educators. Additional noise correction could have been done that may have improved the Lytro slightly. Hot pixel correction and color correction were all rudimentarily implemented in our studies; however, lacking camera specifics minimized the corrections' full potential. In addition, experimentally, both cameras likely had higher noise readings due to excess scatter coming through the lead glass from scatter events from the scintillator or from the buildup material.

The agreement between simulation and experiment for the CoolSNAP camera is very

encouraging that the forward model works. The percent difference between $DQE(0)$ predicted and measured for the CoolSNAP was 1.24%. As expected from the literature, when the magnification of a camera is less than 5, the optical coupling is greatly improved. For the CoolSNAP, using the 40 mm focal length lens the quantum efficiency of imaging single gamma photon events was 29%. The scintillator’s quantum efficiency was 33%. If the 2.2 cm LKH-5 scintillator was used in an EPID, the $DQE(0)$ would be 19%, which does not leave much improvement for the CoolSNAP which was at 15.9%. However, this was a proof-of-concept experiment design, and in reality, the camera would not be sitting less than 20 cm away from the scintillator with a macro-photography lens that allows for very close focusing.

We predicted in simulation that if we used a thicker scintillator with a medium to high-end lens we could get a $DQE(0)$ of almost 18% with a light-field camera, which would be considered novel and competitive in the current EPID market. However, factors like price still come into play and the DQE value of 18% has some heavy assumptions in it, like perfect photon absorption, and no electronic noise. According to the literature, LKH-5 had a quoted price of 1 USD/cm³. For a 40 x 40 x 2 cm scintillator, which would give a field-of-view competitive with current active-matrix EPIDs, the scintillator would cost 3,200 [79]. However, that price is from literature and not quoted from the manufacturer. Based on simulations, the alternative scintillator EJ-260, while likely costing half the price lacks the required density to be competitive with LKH-5. A major drawback to LKH-5 was how easy it was to scratch and chip, which definitely had a negative impact on the experimental noise measurements.

5.2 Spectral Performance Summary & Conclusions


The spectral performance of imaging a thick monolithic scintillator with only one camera was poor in all instances other than at 225 kV in certain simulations. As previously mentioned, noise needs to be better managed, through actions like TV regularization potentially. It

was surprising that ML-EM did not perform as well as hoped. ML-EM has been used by the other light-field group researching light-field imaging for dose assessment purposes. However, in all their reported results they used multiple projections to reconstruct so it is hard to say how they would have performed with a single projection, which is a more ill-posed scenario. In addition, ML-EM was used in an assessment of using plenoptic cameras to image both discrete and continuous light-fields and it seemed to fair well [55]. However, even in the assessment of plenoptic cameras for volumetric imaging, they were reconstructing a very small volume of 64x64x64 voxels and also concluded that two cameras are best for continuous light-fields.


5.3 Future Directions

Given how much effect the light field type being imaged (discrete points vs continuous) has on the overall reconstruction and spectral availability, applications like proton beam imaging may be a better alternative for light-field imaging. In particular proton beams that have a narrow width will fare better, as the higher the transverse frequency a beam has the better the spatial assessment will be axially. In addition, the Bragg peak is a rapid change axially, which will potentially help with axial localization. In a preliminary simulation using our developed forward model, two non-pencil proton beams of 100 MeV and 120 MeV were simulated hitting a 20 x 20 x 20 cm cube of EJ-260 scintillation material. The scintillation volume was reconstructed with ML-EM and the resulting axial responses for each beam are shown in Figure 5.1. The light-field camera used to reconstruct used an F/2 70 mm focal length lens placed at a meter away. With this configuration, the native refocusing range based on purely ray optics is 15.8 cm, and the resolution is 265 μm . As shown the depths of the Bragg peaks differed from the modeled dose depths until the refraction was considered, which brought the peaks within 5 mm.


Determining the location of the Bragg peak is important for improving proton beam



(a) Proton beams as seen from a conventional camera.



(b) Reconstructed axial response vs actual placement.



(c) Reconstructed axial response vs actual placement corrected for refraction

Figure 5.1: Preliminary proton beam simulation and reconstruction using the developed forward model.

treatment. Despite the differences between protons and photons, proton treatments utilize a calibration function to relate x-ray CT values to proton stopping power ratios, which in turn can be used to calculate the effective range of protons. However, this indirect approach can be a significant source of uncertainty [81]. Proton CT (pCT) has been an active area of research over the last decade. Proton CT would provide a more accurate 3D mapping of

proton-stopping power.

More recently, Darne *et al.* have shown a pCT setup using a volumetric liquid scintillator [82]. A volumetric scintillator detector allows for the comprehensive capture of a wide distribution of residual proton beam energies in a single acquisition and can be used to generate proton radiographs. Solid scintillators have been used as well to generate radiographs, however at quite low proton beam energies not suitable for treatment [83]. Historically, these volumetric scintillators have been imaged with conventional cameras and while some optical effects can be corrected for at the focus plane via deconvolution, LF cameras potentially provide a way to do 3D reconstruction, further correcting for optical effects while also providing spectral information on the Bragg peak depth as shown in our preliminary study. While additional simulations should be performed with higher fluence and with pencil beams, the initial results are encouraging.

However, the major concern with proton imaging and general light-field imaging is the speed of the camera. Additional lenses in a light-field camera paired with small pixels results in exposure times in excess of a second, even at megavoltage energies, which may leave light-field cameras in the realm of QA. The future of light-field imaging comes down to improving the resources in the light-field imaging community, bringing down the cost of scientific cameras and experiment equipment, and pioneering additional novel applications for a novel technology.

REFERENCES

- [1] Ui-Jung Hwang, Kwanghyun Jo, Young Kyung Lim, Jung Won Kwak, Sang Hyoun Choi, Chiyoung Jeong, Mi Young Kim, Jong Hwi Jeong, Dongho Shin, Se Byeong Lee, Jeong Hoon Park, Sung Yong Park, and Siyong Kim. A new method and device of aligning patient setup lasers in radiation therapy. *Journal of Applied Clinical Medical Physics*, 17:49, 2016.
- [2] Task group 142 report: quality assurance of medical accelerators. *Medical Physics*, 36:4197–4212, 2009.
- [3] Michael G Herman, James M Balter, David A Jaffray, Kiarin P Mcgee, Peter Munro, Marcel Van Herk, and John W Wong. Clinical use of electronic portal imaging: Report of aapm radiation therapy committee task group 58. *Medical Physics*, pages 712–737, 2001.
- [4] Frank Herbert Attix. *Introduction to Radiological Physics and Radiation Dosimetry*. WILEY-VCH Verlag GmbH & Co. KGaA, 1986.
- [5] John W. Wong. *Electronic Portal Imaging Devices (EPID)*, pages 207–213. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [6] Josh Star-Lack, Mingshan Sun, Andre Meyer, Daniel Morf, Dragos Constantin, Rebecca Fahrig, and Eric Abel. Rapid monte carlo simulation of detector DQE(f). *Medical Physics*, 41, 2014.
- [7] Harrison H. Barrett and William Swindell. *Radiological Imaging: The Theory of Image Formation, Detection, and Processing*. Academic Press, 1996.
- [8] International Electrotechnical Commission. *Medical electrical equipment: characteristics of digital x-ray imaging devices—part 1: determination of the detective quantum efficiency. Document no. 62220-1*. International Electrotechnical Commission, Geneva, Switzerland, 2003.
- [9] Ehsan Samei and Michael J. Flynn. An experimental comparison of detector performance for direct and indirect digital radiography systems. *Medical Physics*, 30:608–622, 4 2003.
- [10] Y. El-Mohri, K. W. Jee, L. E. Antonuk, M. Maolinbay, and Q. Zhao. Determination of the detective quantum efficiency of a prototype, megavoltage indirect detection, active matrix flat-panel imager. *Medical Physics*, 28:2538–2550, 12 2001.
- [11] M. C. Kirby and A. G. Glendinning. Developments in electronic portal imaging systems. *The British Journal of Radiology*, 79 Spec No 1, 2006.
- [12] Swathi Lakshmi Balasubramanian and Ganapathy Krishnamurthi. X-ray scintillator lens-coupled with cmos camera for pre-clinical cardiac vascular imaging—a feasibility study. *PLOS ONE*, 17:e0262913, 2 2022.

- [13] Eugene Hecht. *Optics*. Pearson, 5th edition, 2015.
- [14] Spencer Cox. What is magnification in photography? <https://photographylife.com/what-is-magnification>, 2021.
- [15] Focusall: Focal stacking of microscopic images using modified harris corner response measure. *IEEE/ACM transactions on computational biology and bioinformatics / IEEE, ACM*, 13:326, 3 2016.
- [16] Brigitte Forster, Dimitri Van De Ville, Jesse Berent, Daniel Sage, and Michael Unser. Extended depth-of-focus for multi-channel microscopy images: a complex wavelet approach. pages 660–663, 2004.
- [17] Jean Pierre Bissonnette, I. A. Cunningham, and P. Munro. Optimal phosphor thickness for portal imaging. *Medical Physics*, 24:803–814, 1997.
- [18] G. Pang and J. A. Rowlands. Electronic portal imaging with an avalanche-multiplication-based video camera. *Medical Physics*, 27:676–684, 4 2000.
- [19] Chinmay D. Darne, Fahed Alsanea, Daniel G. Robertson, Fada Guan, Tinsu Pan, David Grosshans, and Sam Beddar. A proton imaging system using a volumetric liquid scintillator: A preliminary study. *Biomedical Physics and Engineering Express*, 5, 2019.
- [20] Mathieu Goulet, Madison Rilling, Luc Gingras, Sam Beddar, Luc Beaulieu, and Louis Archambault. Novel, full 3d scintillation dosimetry using a static plenoptic camera. *Medical Physics*, 41, 2014.
- [21] Ren Ng. *Digital light field photography*. PhD thesis, Stanford University, 2006.
- [22] On the fundamental comparison between unfocused and focused light field cameras. *Applied Optics, Vol. 57, Issue 1, pp. A1-A11*, 57:A1–A11, 1 2018.
- [23] Tomographic approach for the quantitative scene reconstruction from light field images. *Optics Express*, 26:22574, 2018.
- [24] Martin Eberhart and Stefan Loehle. Light-field imaging for plasma wind-tunnel application. *Journal of Thermophysics and Heat Transfer*, 33:407–415, 2019.
- [25] Liam McCabe. Lytro announces illum light field camera: Digital photography review. 2014.
- [26] Raytrix 3d light-field vision. <https://raytrix.de/>.
- [27] Madison Rilling, Guillaume Allain, Simon Thibault, and Louis Archambault. Tomographic-based 3d scintillation dosimetry using a three-view plenoptic imaging system. *Medical Physics*, 47:3636–3646, 8 2020.

- [28] Josh Star-Lack, Daniel Shedlock, Dennis Swahn, Dave Humber, Adam Wang, Hayley Hirsh, George Zentai, Daren Sawkey, Isaac Kruger, Mingshan Sun, Eric Abel, Gary Virshup, Mihye Shin, and Rebecca Fahrig. A piecewise-focused high dqe detector for mv imaging. *Medical Physics*, 42:5084–5099, 9 2015.
- [29] Joerg Rottmann, Daniel Morf, Rony Fueglistaller, George Zentai, Josh Star-Lack, and Ross Berbeco. A novel epid design for enhanced contrast and detective quantum efficiency. *Physics in Medicine & Biology*, 61:6297, 8 2016.
- [30] Ingrid Valencia Lozano, Mengying Shi, Marios Myronakis, Paul Baturin, Rony Fueglistaller, Pascal Huber, Mathias Lehmann, Daniel Morf, Dianne Ferguson, Matthew W. Jacobson, Thomas Harris, Ross I. Berbeco, and Christopher L. Williams. Frequency-dependent optimal weighting approach for megavoltage multilayer imagers. *Physics in Medicine & Biology*, 66:085012, 4 2021.
- [31] Virginia Polytechnic Institute. Monte carlo simulation. <https://sites.google.com/a/vt.edu/monte-carlo-simulation/history>.
- [32] Robert L. Harrison. Introduction to monte carlo simulation. *AIP conference proceedings*, 1204:17, 1 2010.
- [33] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytracsek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell’Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani, D. Gibin, J. J. Gomez Cadenas, I. Gonzalez, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F. W. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata, M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampen, V. Lara, V. Lefebure, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. Mora de Freitas, Y. Morita, K. Murakami, M. Nagamatu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O’Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M. G. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. di Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. Safai Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J. P. Wellisch, T. Wenaus, D. C. Williams, D. Wright, T. Yamada, H. Yoshida, and D. Zschesche. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506:250–303, 7 2003.
- [34] I Kawrakow, DWO Rogers, E Mainegra-Hing, F Tessier, RW Townson, and BRB Walters. Egsnrc toolkit for monte carlo simulation of ionizing radiation transport, 2000.

- [35] J. Perl, J. Shin, J. Schümann, B. Faddegon, and H. Paganetti. Topas: an innovative proton monte carlo platform for research and clinical applications. *Medical Physics*, 39:6818–6837, 2012.
- [36] Ebubekir Buber and Banu Diri. Performance analysis and cpu vs gpu comparison for deep learning. *2018 6th International Conference on Control Engineering and Information Technology, CEIT 2018*, 10 2018.
- [37] Collimated Holes Inc. <http://www.collimatedholes.com/>.
- [38] Yue Houng Hu, Daniel Shedlock, Adam Wang, Joerg Rottmann, Paul Baturin, Marios Myronakis, Pascal Huber, Rony Fueglistaller, Mengying Shi, Daniel Morf, Josh Star-Lack, and Ross I. Berbeco. Characterizing a novel scintillating glass for application to megavoltage cone-beam computed tomography. *Medical Physics*, 46:1323–1330, 2019.
- [39] Chinmay D. Darne, Daniel G. Robertson, Fahed Alsanea, Charles Antoine Collins-Fekete, and Sam Beddar. A novel proton-integrating radiography system design using a monolithic scintillator detector: Experimental studies. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1027:166077, 3 2022.
- [40] Eljen Technology. Green emitting ej-260, ej-262. <https://eljentechnology.com/products/plastic-scintillators/ej-260-ej-262>.
- [41] Physics considerations in mv-cbct multi-layer imager design. *Physics in Medicine & Biology*, 63:125016, 6 2018.
- [42] Madison Rilling, Guillaume Allain, Simon Thibault, and Louis Archambault. Tomographic-based 3d scintillation dosimetry using a three-view plenoptic imaging system. *Medical Physics*, 47:3636–3646, 8 2020.
- [43] Zemax software. <https://www.zemax.com/>.
- [44] P. Munro, J. A. Rawlinson, and A. Fenster. Therapy imaging: Source sizes of radiotherapy beams. *Medical Physics*, 15:517–524, 7 1988.
- [45] A. Fenster, J. J. Battista, and P. Munro. X-ray sources of medical linear accelerators: focal and extra-focal radiation. *Medical Physics*, 20:1417–1427, 1993.
- [46] Xrad specifications. https://lfic.dfci.harvard.edu/sites/default/files/pdfs/IGRT_04.pdf.
- [47] Jennifer L. Marshall, Patrick Williams, Jean-Philippe Rheault, Travis Prochaska, Richard D. Allen, and D. L. DePoy. Characterization of the reflectivity of various black materials. In Suzanne K. Ramsay, Ian S. McLean, and Hideki Takami, editors, *Ground-based and Airborne Instrumentation for Astronomy V*, volume 9147, page 91474F. International Society for Optics and Photonics, SPIE, 2014.

- [48] Ivan Bazarov. Geometric optics. <https://www.classe.cornell.edu/~ib38/teaching/p330/Lecture5.pdf>, 2016.
- [49] A simple method for determining the modulation transfer function in digital radiography. *IEEE transactions on medical imaging*, 11:34–39, 1992.
- [50] G. Lubberts. Random noise produced by x-ray fluorescent screens*. *JOSA*, Vol. 58, Issue 11, pp. 1475-1483, 58:1475–1483, 11 1968.
- [51] Robert K Swank. Absorption and noise in x-ray phosphors. *Journal of Applied Physics*, 44:4199–4203, 1973.
- [52] Linda Shapiro. *Computer Vision*. Pearson, 1st edition, 2 2001.
- [53] Sidney Ray. *Applied Photographic Optics*. Focal Press, 3rd edition, 2002.
- [54] Joseph Goodman. *Introduction to Fourier Optics*. McGraw-Hill, 2 edition, 1996.
- [55] Hecong Liu, Qianlong Wang, and Weiwei Cai. Assessment of plenoptic imaging for reconstruction of 3d discrete and continuous luminous fields. *Journal of the Optical Society of America A*, 36:149, 2 2019.
- [56] Wim van Aarle, Willem Jan Palenstijn, Jeroen Cant, Eline Janssens, Folkert Bleichrodt, Andrei Dabrovolski, Jan De Beenhouwer, K. Joost Batenburg, and Jan Sijbers. Fast and flexible x-ray tomography using the astra toolbox. *Optics Express*, Vol. 24, Issue 22, pp. 25129-25147, 24:25129–25147, 10 2016.
- [57] Nikki Tirada, Guang Li, David Dreizin, Luke Robinson, Gauri Khorjekar, Sergio Dromi, and Thomas Ernst. Digital breast tomosynthesis: Physics, artifacts, and quality control considerations. *Radiographics*, 39:413–426, 3 2019.
- [58] Nicola Viganò, Francesco Brun, Pablo Martinez Gil, and Charlotte Herzog. Plenoptomos. <https://github.com/cicwi/plenoptomos>.
- [59] Lytro illum 40 megaray light field camera overview. https://www.dpreview.com/products/lytro/compacts/lytro_illum.
- [60] Photometrics coolsnap hq². <https://www.photometrics.com/wp-content/uploads/2019/10/HQ2-Datasheet.pdf>.
- [61] Af-s dx micro nikkor 40mm f/2.8g lens specifications. <https://www.nikonusa.com/en/nikon-products/product/camera-lenses/af-s-dx-micro-nikkor-40mm-f%252f2.8g.html>.
- [62] Duane Smalley, Stuart Baker, Brandon Baldonado, Jesus Castaneda, Andrew Corredor, Jessica H. Clayton, Logan Fegenbush, Cort Gautier, Amanda Gehring, Todd Haines, James Lucero, John Stearns, and Katie Walters. Image restoration of high-energy x-ray radiography with a scintillator blur model. *Nuclear Instruments and Methods in Physics*

Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 968:163910, 7 2020.

- [63] Emil Y. Sidky, Jakob H. Jørgensen, and Xiaochuan Pan. Convex optimization problem prototyping for image reconstruction in computed tomography with the chambolle-pock algorithm. *Physics in Medicine & Biology*, 57:3065–3091, 11 2011.
- [64] Nicola Viganò, Pablo Martínez Gil, Charlotte Herzog, Ombeline de la Rochefoucauld, Robert van Liere, and Kees Joost Batenburg. Advanced light-field refocusing through tomographic modeling of the photographed scene. *Optics Express*, 27:7834, 2019.
- [65] Decoding, calibration and rectification for lenselet-based plenoptic cameras. *Computer Vision and Pattern Recognition (CVPR)*, pages 1027–1034, 2013.
- [66] Gregg Vane, Robert O. Green, Thomas G. Chrien, Harry T. Enmark, Earl G. Hansen, and Wallace M. Porter. The airborne visible/infrared imaging spectrometer (aviris). *Remote Sensing of Environment*, 44:127–143, 5 1993.
- [67] Dual- and multi-energy ct: Principles, technical approaches, and clinical applications. *Radiology*, 276:637–653, 9 2015.
- [68] Dimple Modgil, David S. Rigie, Yuxin Wang, Xianghui Xiao, Phillip A. Vargas, and Patrick J. La Rivière. Material identification in x-ray microscopy and micro CT using multi-layer, multi-color scintillation detectors. *Physics in Medicine & Biology*, 60:8025–8045, 2015.
- [69] Marios Myronakis, Rony Fueglistaller, Joerg Rottmann, Yue-Houng Hu, Adam Wang, Paul Baturin, Pascal Huber, Daniel Morf, Josh Star-Lack, and Ross Berbeco. Spectral imaging using clinical megavoltage beams and a novel multi-layer imager. *Physics in Medicine & Biology*, 62:9127–9139, 2017.
- [70] E. Roessl and C. Herrmann. Cramér-rao lower bound of basis image noise in multiple-energy x-ray imaging. *Physics in Medicine & Biology*, 54:1307–1318, 2009.
- [71] Scott Hayden. Lytro website goes to life support mode as company winds down after google acquisition – road to vr. <https://www.roadtovr.com/report-google-acquires-light-field-company-lytro-move-hire-employees/>, 4 2018.
- [72] Advanced methods of microscope control using manager software. *Journal of biological methods*, 1:e10, 11 2014.
- [73] Lytro Inc. Lytro power tools beta. <https://github.com/kmader/lytro-power-tools>, 2015.
- [74] Technical note: Accuracy of mtf measurements with an edge phantom at megavoltage x-ray energies. *Medical Physics*, 46:5685–5689, 12 2019.

- [75] P. Monnin, H. Bosmans, F. R. Verdun, and N. W. Marshall. A comprehensive model for quantum noise characterization in digital mammography. *Physics in Medicine & Biology*, 61:2083, 2 2016.
- [76] David M. Klein, Ramesh C. Tailor, Louis Archambault, Lilie Wang, Francois Therriault-Proulx, and A. Sam Beddar. Measuring output factors of small fields formed by collimator jaws and multileaf collimator using plastic scintillation detectors. *Medical Physics*, 37:5541, 2010.
- [77] Daniel Sage, Lauréne Donati, Ferréol Soulez, Denis Fortun, Guillaume Schmit, Arne Seitz, Romain Guiet, Cédric Vonesch, and Michael Unser. Deconvolutionlab2: An open-source software for deconvolution microscopy. *Methods*, 115:28 – 41, 2017.
- [78] Michael Stuart West. *Radiation-induced luminescence in terbium-doped silicate glasses*. PhD thesis, William & Mary, 1997.
- [79] Mairead E. Montague, Cordell Delzer, Xianfei Wen, Kathleen C. Goetz, Daniel Shedlock, and Jason P. Hayward. Radiation hardness characterization of lkh-5 scintillating glass. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 982:164551, 12 2020.
- [80] High quality light field extraction and post-processing for raw plenoptic data. *IEEE Transactions on Image Processing*, 29:4188–4203, 2020.
- [81] Ming Yang, X Ronald Zhu, Peter C Park, Uwe Titt, Radhe Mohan, Gary Virshup, James E Clayton, and Lei Dong. Comprehensive analysis of proton range uncertainties related to patient stopping-power-ratio estimation using the stoichiometric calibration. *Physics in Medicine & Biology*, 57:4095–4115, 2012.
- [82] Chinmay D. Darne, Fahed Alsanea, Daniel G. Robertson, Fada Guan, Tinsu Pan, David Grosshans, and Sam Beddar. A proton imaging system using a volumetric liquid scintillator: A preliminary study. *Biomedical Physics and Engineering Express*, 5, 2019.
- [83] Sodai Tanaka, Teiji Nishio, Masato Tsuneda, Keiichiro Matsushita, Shigeto Kabuki, and Mitsuru Uesaka. Improved proton ct imaging using a bismuth germanium oxide scintillator. *Physics in Medicine & Biology*, 63, 2018.