Task-Driven Trajectory Design for Endovascular Embolization

by

Gabriela M. Rodal

A thesis submitted to The Johns Hopkins University
in conformity with the requirements for the degree of
Master of Science in Engineering

Baltimore, Maryland
February, 2019

Abstract

Computed Tomography (CT) is one of the most useful and widely applied imaging modalities, employed in both diagnostic and treatment planning purposes in the medical field. Circular and spiral acquisition trajectories are traditionally employed and work well in many cases. The advent of technologies such as robotic C-arms in interventional imaging allow for more complex data acquisitions, which enables potential improvements in image quality, increased field of view, and sampling. This capability has particular potential crucial in interventional cases where images may be compromised by complex anatomy or surgical tools. In this work, we present a paradigm that uses custom non-circular orbits and prior patient information along with segmentation and registration techniques to account for surgical tools and/or implants, to improve image quality. The framework leverages the anatomical model to optimize a parameterized source-detector trajectory for a variety of specific imaging tasks. We propose an overall workflow for orbit customization with investigations of the various workflow stages as well as the overall performance of the framework.

Thesis Committee Members

J. Webster Stayman, Ph.D. (Primary Advisor)

Assistant Professor, Department of Biomedical Engineering

Johns Hopkins University

Wojciech B. Zbijewski, Ph.D.

Assistant Professor, Department of Biomedical Engineering

Johns Hopkins University

Grace J. Gang, Ph. D.

Research Associate, Department of Biomedical Engineering

Johns Hopkins University

Acknowledgements

This thesis represents the culmination of my time at Hopkins, where I found not only a passion for medicine and technology, but also lasting respect and friendship for the people I met. When I arrived at Hopkins in 2013, I never would have imagined it would be my home for the next five and a half years. Everything was new and intimidating, from the campus, to the classes, to Baltimore itself. But with time, I and all my peers grew into the big shoes Hopkins expected us to fill with the help of great mentors and professors. Dr. Stayman stands out as one such outstanding professor who genuinely cares for the wellbeing and growth of his students both academically and in life. His cheer and encouragement made my thesis year a delight, even when experiments failed, or results were frustrating. My time learning from him, both in his Imaging Instrumentation class and in lab, has been invaluable to me, and I deeply appreciate all that he has done as an advisor and as a person.

I feel incredibly blessed to have worked with such brilliant, talented, and kind people at the AIAI lab. During my time with AIAI, I've learned about medical imaging and designing experiments (and where to get the best food in Baltimore).I want to express my sincerest gratitude to the other members of my thesis committee, Dr. Wojciech Zbijewski and Dr. Grace Gang, for their advice over the past few months and for taking the time to read my thesis despite their busy schedules. Thank you to everyone in the AIAI Lab for the support and kindness, including Dr. Grace Gang, Dr. Hao Zhang, Dr. Steven Tilley, Wenying Wang, Matthew Tivnan, Jessica Flores, Amalie Shi, Stephen Liu, Xueqi Guo, Kailun Cheng, Andrew Mao, Lorenz Hehn, Nischita Kaza, and Jalen Cooke. I must give special thanks to Sam Bourne, who has advised and

supported me since beginning the MSE program at Hopkins. He has truly gone above and beyond for his advisees and has made the program what it is. My time at Hopkins and with the AIAI lab has been a unique and precious experience, filled with memories that I will treasure for years to come. I am more grateful than I can convey with words here in these acknowledgements—so thank you to everyone who's been there along the way. This thesis is dedicated to you.

Table of Contents

List of Figures

Chapter 1

Introduction to Optimizing Orbits to Improve Imaging with Metal Implants

1.1 Context and Objective of Work

Medical imaging is an important component of diagnostic medicine as well as of the modern surgical workflow. Computed Tomography (CT) has emerged as one of the most common imaging modalities, with tens of millions of scans taken each year within the medical sector alone (Brenner and Hall, 2007). The popularity of CT is largely due to its fast acquisition times compared to other imaging modalities as well as the high resolution of the 3D images produced.

A special computed tomography modality with a flat panel detector, cone beam computed tomography (CBCT), has seen increasing applications. Unlike fan-beam CT, CBCT exposure incorporates the entire field of view (FOV) so only one rotation about the gantry is necessary to acquire enough data for 3D reconstruction rather than a helical progression requiring stacking separate 2D reconstructions. This allows for more rapid acquisition of a data set for the entire FOV and requires a less expensive detector, contributing to the low cost and small foot print of CBCT. Additional advantages of CBCT include increased x-ray tube efficiency and reduced image distortion due to internal patient movement. However, CBCT

reconstructions frequently require scatter correction as large amounts of scattered radiation are detected, affecting the image quality due to noise and degraded contrast resolution. (Yan and Leahy, 2000) Despite this limitation, CBCT has found a wide range of medical applications in radiotherapy guidance, diagnosis, and treatment planning. For example, early detection via advanced imaging systems, including CBCT, allows doctors to identify tumors, hemorrhages, and other abnormalities in a timely matter to produce better patient outcomes. (Lee et al., 2016)

1.2 Motivation for Work

Importantly, imaging modalities can be used for surgical guidance, such as C-arm based CBCT which allows for additional mobility of the source and detector. This mobility provides additional flexibility in acquisition design, making CBCT invaluable in the operating room for image guided procedures. (Stayman and Siewerdsen, 2013) The presence of surgical tools and the necessity of highly attenuating metal implants often introduce severe metal artifacts that degrade image quality and obscure visualization of important anatomical features–frequently at the surgical site itself where physicians most need clear images.

Metal artifacts often arise because satisfactory images cannot be calculated from projections with missing or distorted data, which occurs when all or most photons are attenuated by the metal for all projections. To correct for metal artifacts, many approaches have been proposed. Traditionally, solutions are based on iterative or algebraic reconstruction techniques. However, these require a large number of iterations to converge. Kalender proposed a simpler solution which uses linear interpolation to “fill in” the missing data and found that metal artifacts were reduced. (Kalender et al., 1987) Iterative methods, such as the Metal Deletion Technique

(MDT), focus on the principle that projection data near the metal is less accurate due to the factors such as beam hardening and under-sampling. Metal pixels are deleted from the reconstruction with each iteration and forward projected values from previous iterations replace the distorted data. (Boas and Fleischmann, 2011; Boas and Fleischmann, 2012). Ultimately, the core of each solution is to "fill in" missing data – this can also be accomplished by leveraging the mobility of robotic C-arms to navigate closely around the metal coil to avoid passing many photons through it.

The tasks chosen for this work are focused upon areas with degraded image quality due to a metal implant or areas where photons are highly attenuated by bone. Traditionally, circular and spiral orbits are used for CT scans and have been the standard data acquisition approach for decades and are the most common trajectory for CBCT scans as well. However, non-circular scans have also been investigated for a variety of applications in which an expanded FOV or improved sampling is desired. This thesis explores the impact of different orbits on image quality, given a specific task.

A change in orbit as simple as tilting a circular orbit can positively impact image quality and localization. This has been demonstrated in head imaging (Menzel et al., 1999). These non-standard orbits have been proposed to expand field of view but we believe there are additional opportunities to, for example, improve image quality surrounding a metal implants, and to reduce radiation dose. To perform such a design, a performance metric must be chosen. Task-based acquisition design has been introduced in recent years to optimize these non-circular orbits for specific tools, procedures, and patient anatomy. We propose task-based detectability, which accommodates many system dependences including the imaging task, patient anatomy, anatomical location, and reconstruction parameters. (Gang et al., 2014)

1.3 Utility of Non-Standard Orbits

1.3.1 "Complete" Orbits

Orlov's condition is based on the set of acquisition angles observed by parallel-beam sampling. Tuy modifies the parallel-beam geometry to describe the limitations of cone-beam geometry. (Metzler, 2003) Tuy's condition is based on the relationship of a point to the curve of the cone-beam focal-source point. (Tuy et al., 1983) A traditional cone-beam CT circular scan does not yield a "complete" orbit, excluding vital 3D information about the object of interest as it violates Tuy's condition. Such "incompleteness" manifests as cone-beam artifacts including a loss of z-axis frequencies and streaking. Such artifacts can be reduced by improving 3D sampling and data completeness using non-circular trajectories. However, "complete" orbits may not be necessary for specific tasks. For example, non-circular orbits can reposition missing frequency information to frequencies less relevant to the given imaging task. In mammotomography, physical limitations limit the field of view for cameras with parallel collimators such that the tissue immediately anterior to the chest wall cannot be imaged with a simple circular orbit. Various complex 3D trajectories have been proposed to avoid physical limitations imposed by the patient or bed and compensate for the missing data. Investigations have demonstrated that these more complex 3D orbits yield reconstructions with less distortion and more complete sampling. (Madhav et al., 2009; Brzymialkiewicz et al., 2006)

1.3.2 Tilted Circular Orbits

Tilting the source-detector orbit relative to the longitudinal axis of the object can produce significant improvement in some imaging scenarios. Such improvements range from reduced

radiation to the patient's eyes (Nikupaavo et al., 2015) to improved image quality adjacent to the skull base (Menzel et al., 1999) where a large amount of highly-attenuating bone drastically reduces the number of photons reaching the detector. Additional flexibility in trajectory to move around complexly shaped/positioned highly attenuating objects yields further improvements to these images.

1.3.3 Extended Field of View and Sampling

Non-circular orbits have been used to provide extended axial (Yu et al., 2016) and elliptical (Herbst et al., 2015) fields of view. Such orbits can be used to address the "long body problem". The long body problem presents itself in helical orbits for multi-slice CT. Incomplete data is created at each end of the spiral orbits which is impossible to reconstruct without artifacts. Several solutions have been proposed such as spirals with circular "end caps" and a line plus partial helical scan. (Yan and Leahy, 2000)

1.4 Past Work Focused on Task-Driven Optimization

In recent years, there has been an effort to drive data acquisitions for particular tasks. We describe a few approaches to demonstrate the utility of task-driven optimization for imaging system design, acquisition, and reconstruction.

1.4.1 Tube Current Modulation

Tube current modulation can be altered such that the radiation dose to the patient is reduced while maintaining the image quality. This is typically achieved by decreasing current for

less attenuating areas or increasing current for more attenuating areas. The current is modulated based upon measured attenuation or a sinusoidal-type function corresponding to the transaxial shape of the patient. Tube modulation can either be preprogrammed or implemented in real time using a feedback mechanism or some dynamic combination of the two.

Task-based optimization has been used to efficiently and effectively design CT tube current modulation and orbital tilt in filtered back projection (FBP) reconstructions. By leveraging a theoretical model based upon the implicit function theorem and Fourier approximations, detectability of a sphere in a head phantom was improved by 30% compared to the unmodulated case and detectability of a line-pair pattern was improved by 80%. (Gang et al., 2015)

1.4.2 Fluence-Field Modulation

Fluence field modulation (FFM) is an emerging technology wherein the x-ray beam is dynamically customized. This is an extension of tube current modulation and provides additional opportunity to reduce radiation dose to patients. Strategies for FFM include inverse geometry CT and dynamic beam filtration using a number of approaches including multiple aperture devices (Stayman et al., 2016), liquid-filled bowtie filters (Shunhavanich et al., 2015; Szezykutowicz et al, 2015), and split bowtie filters (Mail et al., 2009).

Prior anatomical knowledge from a low-dose scout CT scan can be used to predict image quality, for example, spatial resolution and noise as a function of FFM. This information may be used to determine detectability in a volume-of-interest given the FFM acquisition parameters, which then allows for an ideal FFM to be estimated. Gang et al. used a covariance

matrix adaptation evolutionary strategy to maximize the minimum detectability over a volume-of-interest when using a model-based iterative reconstruction (MBIR) for image formation. (Gang et al., 2016)

1.4.3 Regularization with Penalized Weighted Least Squares Reconstruction

Regularization is used to promote smoothing and/or edge sharpness in an image. This can be achieved using any number of penalty functions, which will be discussed in depth in future chapters. Dang et al. created a spatially-varying regularization term $R(\mu)$ which penalizes intensity differences between neighboring voxels to promote local smoothness but also to maximize task-based detectability of small hemorrhages. (Dang et al., 2017)

The same detectability metric used in Gang's work described previously, was implemented for this study to optimize the penalty function such that the image quality was maximized for a specific volume of interest. Dang determined that the local MTF and NPS exhibit shift-variance and anisotropy and that the measured MTF and NPS were in accordance with those approximated for the optimization. It also concluded that image was most improved in areas of high attenuation such as near the skull base when using the task-based penalty; these areas are otherwise prone to over-smoothing data. (Dang et al., 2016)

Chapter 2

Basic Principles of Computed Tomography

2.1 Physical Principles

The groundbreaking discovery of X-rays in 1895 by Wilhelm Roentgen served as a catalyst for modern diagnostic radiology. The versatility of x-rays contributed to its ubiquity in the medical field, as it can be used to image nearly all tissues and organs through projection radiography, or more recently with tomography.

Computed Tomography, or CT, refers to a computerized x-ray imaging procedure in which a beam of x-rays is projected through an object. A motorized x-ray source rotates the beam about the object of interest. (Prince and Links, 2015) Once the photons are attenuated by the object, the information provided by the digital x-ray detectors located opposite the source is used to reconstruct cross-sectional “slices”, or tomographic images, that reveal the composition of the object. Digitally stacking these slices, or projecting a wider beam of photons, such as a fan or cone, allows for a three dimensional reconstruction of the object.

2.1.1 X-ray Generation and Detection

X-rays are a form of electromagnetic (EM) radiation, with energies between 10-150 keV when used for medical diagnostic purposes. This radiation is produced in a cathode ray tube by accelerating electrons in a vacuum between the cathode and the anode. At the cathode, a tungsten filament is heated causing electrons to be emitted from the metal. The free electrons are then accelerated by an applied voltage to collide with the metal target anode. The collision instantly decelerates the fast-moving electrons, and 99% of their kinetic energy is transformed into heat. Only a tiny percentage of the electron's energy is converted into useful X-ray radiation. (Behling, 2015)

X-rays are generated primarily by Bremsstrahlung radiation, or "braking radiation", but also by characteristic radiation. In the Bremsstrahlung process, which generates approximately 80% of the observed photons, the accelerated electrons travelling in a material are deflected, slowed or completely stopped by the attractive force from the positively charged nucleus that it encounters. (Behling, 2015) Kinetic energy is lost as the electron slows. By conservation of energy, this lost energy must be either absorbed by the atom or converted to another form of energy. When the electron is sufficiently close to the nucleus, the electrostatic interaction and braking effect increases and the resulting "braking" photon possesses more energy. Therefore, the energy produced by this process produce photons of a variety of wavelengths within the X-ray spectrum.

The second method of producing these photons is a high-energy collision between high-speed electrons accelerated through the vacuum tube and the anode metal's atoms. This collision produces characteristic radiation. Upon colliding, a lower-orbital electron is ejected leaving a vacancy that is filled by higher-orbital electrons. In dropping to a lower energy orbital, an X-ray

photon is produced from the energy difference between the higher and lower orbitals. These energies are dependent upon the anode material, as the probability for filling a K-shell vacancy differs by orbital level and thus the number of generated photons at each characteristic energy is different. (Prince and Links, 2015) It is important to note that this process will only generate an X-ray photon if the kinetic energy of the accelerated electron is greater than the energy binding the electron to the nucleus, otherwise the lower orbital electron will not be removed, and the reaction will not occur.

The emitted photon intensity can be controlled by modifying the anode current and the accelerating voltage. Higher anode current increases the cathode's heating, which increases the thermionic emission of electrons from the tungsten which will in turn increase the number of emitted X-ray photons. (Prince and Links, 2015) Increasing the accelerating voltage increases the kinetic energy of the electrons thus increasing the energy of the emitted X-ray photons once the electrons collide with the anode.

An increase in photon energy corresponds with an increase in penetrating power. In imaging, this is vital as low energy photons may not fully penetrate the patient in order to reach the detector and be used to produce an image. Absorbed photons may mutate DNA, causing abnormal base pairings as well as double-strand breaks. The downstream effects of these radiation-induced errors include somatic mutations and possibly cancer. Due to the increased usage of CT for diagnostic imaging, a large body of research has been dedicated to reducing this ionizing dose while maintaining image quality.

2.1.2 X-ray Interactions with Matter and Beer-Lambert Law

Upon emission from the source, the x-ray photons interact with matter primarily through photoelectric absorption and Compton scatter. An x-ray undergoing photoelectric absorption will eject a K-shell electron from an atom with the same energy of the incident x-ray. Photoelectric interactions usually occur with electrons that are firmly bound to the atom, with a relatively high binding energy. The interactions are most probable when the electron binding energy is only slightly less than the energy of the photon, such that the electron can be ejected. The probability of the interaction is proportional to $\left(\frac{Z}{E}\right)^3$, where Z is the atomic number and E is the energy of the photon.

The second manner in which x-rays typically interact with matter is through Compton scatter, in which the incident photon loses only a portion of its energy to a valence electron. This electron is then ejected while a lower-energy photon is deflected, or "scattered", by the material at an angle. The energy of the scattered photon is dependent upon the scattering angle. This can be undesirable as the scattered photons that continue in a forward direction become a secondary radiation source and reduce the quality of the image. (Behling, 2015)

The surviving photon information collected by the digital detector is dependent upon the density of the materials within the object. X-ray attenuation is a function of initial photon count, I_0 and the density of the materials which compose the scanned object.

The basic tenant of the Beer-Lambert law is that photons are attenuated exponentially as they travel through an object. Mathematically, this means that for a monochromatic X-ray beam of energy E, the law states that the remaining photons after the beam traverses distance ℓ within any present material is as follows

$$I = I_0 e^{-\int_0^{\ell} \mu(x)dx}$$

where the linear attenuation coefficient μ can vary along the beam's path based upon the specific material being traversed at position x. In general, μ is also energy-dependent which is important due to the polychromatic x-ray beam. Using the monochromatic model above ignores these spectral effects and can lead to artifacts such as beam hardening. In this work, we ignore such effects. The general goal of X-ray CT is to recover the attenuation coefficient, μ, from the information at the detector, I, though mathematical methods. (Mahesh, 2013)

2.2 Data Acquisition Trajectories

Circular trajectories are popular in cone-beam CT imaging. However, 3D image reconstruction based on circular trajectories does not satisfy Tuy's data sufficiency condition for an exact three-dimensional reconstruction, which states that any plane passing through the FOV of the object must intersect the source trajectory in at least one point. (Metzler, 2003) Despite this, circular acquisitions are widely used despite the potential for cone-beam artifacts. The appropriateness of a trajectory is necessarily dependent on the task at hand, and whether the sampling provided by a particular acquisition is sufficient to properly image a region of interest. Investigation of custom non-circular orbits is central to the work presented here.

2.3 Reconstruction Considerations

2.3.1 Forward Model with Mono-Energetic X-rays

The mono-energetic forward model, written in matrix-vector form is:

$$\bar{y} = \boldsymbol{D}\{g\}\exp(-l)$$

Where $l = \boldsymbol{A}\mu$, where $\boldsymbol{D}$ is a diagonal operator, g are measurement-dependent gains (detector sensitivities, for example), l denotes a vector of all line integrals with μ denoting the vector of attenuation values of the object, **A** is the linear projection operator describing system geometry. Poisson noise models are often applied to describe expected variations in y.

2.3.2 Mathematical Fundamentals

To obtain 3D images from projection data requires an inversion of the model. Classic reconstruction models rely on the "Projection-Slice Theorem", or the "Fourier Slice Theorem", which relates the one-dimensional Fourier transform of the projection data and the two-dimensional Fourier transform of the object. It states that the one-dimensional Fourier transform of the projection equals a line passing through the origin of the two-dimensional Fourier transform of the object at the same an angular orientation (θ) as the angle θ at which the projection was acquired. (Prince and Links, 2015)

If $F\ (u, v) = \ \mathcal{F}_{2D}\{f(x, y)\}$ is the Fourier transform of the object, the Projection-Slice Theorem states that

$$F(\rho\ , \theta) = \mathcal{F}_{1D}\{l(u, \theta)\} = P(\rho\ , \theta)$$

Where u is the horizontal positional index on the detector, P is the Fourier transform of the line integral (projection) data l, and ρ , θ are the polar coordinates in the Fourier domain. (Prince and Links, 2015)

2.3.3 Filtered Back Projection Reconstruction

The Projection-Slice Theorem allows derivation of the filtered back projection (FBP) algorithm. The classic derivation relies on several fundamental assumptions about the scanner geometry, including that it be a parallel-beam scanner and have a circular source-detector trajectory. In the parallel beam case, the object image is $f(x, y)$:

$$f(x,y) = \int_0^{\pi} \left[\int_{-\infty}^{\infty} |\varrho|\, G(\varrho,\theta) e^{j2\pi\varrho\ell}\, d\varrho \right]_{\ell = x\cos\theta + y\,\sin\theta}$$

Where the 1-D Fourier transform of a projection with respect to ℓ (line integral projectection data) is:

$$G(\varrho,\theta) = \mathcal{F}_{1D}\{g(\ell,\theta)\} = \int_{-\infty}^{\infty} g(\ell,\theta) e^{-j2\pi\varrho\ell}\, d\ell$$

In which ϱ denotes spatial frequency and $|\varrho|$ is a "ramp" filter in the frequency domain. G is the Fourier transform of data g, and ϱ, θ are polar coordinates in the Fourier domain.

The ramp filter is frequently employed with a cutoff frequency at or below the Nyquist sampling limit in order to avoid amplifying high frequency noise. Generally, the direct application of the Fourier method leads to sampling that is inversely proportional to ϱ. Additional "weights" are required to compensate for the sparser sampling at higher frequencies. (Prince and Links, 2015)

In summary, this straightforward FBP reconstruction is the implementation of a ramp filter in the frequency domain followed by a backprojection over the line ℓ , defined as $x\cos\theta + y\,sin\theta$, summed over filtered projections at each rotation angle. This approach is the most widely used reconstruction algorithm in clinical CT scanners.

2.3.4 Penalized Likelihood Image Reconstruction

More modern reconstruction approaches include iterative methods which have demonstrated an ability to reduce noise over classic FBP images. While such statistical iterative reconstruction may result in longer reconstruction times compared to FBP reconstructions using the same raw projection data, the resulting images are often substantially less noisy with preserved spatial resolution and border definition. This means that image quality can be maintained or even improved despite radiation dose reduction.

The objective function of penalized-likelihood (PL) estimation is as follows:

$$\hat{\mu} = \arg\max_{\mu} \log L(y;\mu) - R(\mu;\delta)$$

where $R(\mu;\delta)$ represents a penalty term.

In PL image reconstruction, the regularization parameter β controls the tradeoff between resolution and noise. However, this parameter can be difficult to choose and has no explicit direct relation to spatial resolution and noise levels. Moreover, these properties are object dependent, varying between objects or within different areas of the same object.

The additional δ parameter is included to allow for more complex penalties, such as the Huber penalty or a hyperbola function. For a Huber penalty, the δ term specifies the value at

which the loss function becomes linear. This property allows the Huber penalty to combine the sensitivity of the quadratic loss function (L2) and the robustness of the absolute value function (L1) . While the quadratic function allows for simpler optimization and global smoothing, the nonquadratic functions such as the Huber function preserves edges. Edge preservation is often desired when iteratively reconstructing an image, but it also has the potential to emphasize streak artifacts and distort the image. (Fessler and Hero, 1995) It should be noted that the analysis and prediction of resolution and noise properties can be difficult for nonquadratic penalty functions.

The use of iterative reconstruction is vital when employing non-traditional acquisitions. Iterative methods are not limited by challenges such as sparse data sampling or noncircular trajectories, such as those discussed in this work. These methods will adapt to the data acquired and improve the image quality wherever possible given the information provided.

2.4 Challenges Posed by Limited Data

Interventional CT systems often are subject to a variety of kinds of incomplete data. Limited data can include truncated data, limited angle data, and few-view data:

1. Truncated data results from a reduction in the FOV, which truncates the projections during the acquisition and produces severe artifacts during reconstruction. This results in complicated artifacts which are not easily remedied. When the data is truncated axially, a helical scan may be employed but specialized methods must be employed to handle transverse truncation. (Clackdoyle and Defrise, 2010)
2. Limited angle data refers to CT scans in which projection data is acquired in such a manner that the source-detector do not turn the required 180 degrees plus the angular

width of fan beam. These limited-angle scans significantly degrade the quality of the reconstructed image. Conventional FBP and iterative methods do not perform well in these cases and certain artifacts present themselves regardless of the reconstruction method. (Clackdoyle and Defrise, 2010)

3. Few-view CT refers to a CT scan in which the number of acquired projections are significantly fewer than the conventional hundreds or thousands of views over 360 degrees. This type of sparse scan is favored from an economic view in that it dramatically reduces the acquisition time and increases equipment life cycle. (Kudo et al, 2002) These sparse-view CTs contain insufficient sampling, where the number of projections does not satisfy the Nyquist sampling theorem and the resulting FBP reconstructions contain streak artifacts. (Sidky et al., 2006) Iterative reconstructions based on an assumption of image sparsity in the image gradient domain have demonstrated an ability to minimize the artifacts resulting from such sparse sampling.

Chapter 3

Formulation of the Task Driven Design Problem

3.1 Resolution

Resolution is a measurement of an imaging system's ability to reproduce an object's detail, and can be influenced by an array of different factors. As features of interest become smaller, higher resolution is required to be able to resolve details. Spatial resolution can be assessed qualitatively using line pair targets. While acquiring such image data is straightforward, the technique is biased by observer subjectivity and only provides information on the limiting value of the imaging system.

The modulation transfer function (MTF) is a metric describing resolution of linear and shift-invariant imaging systems in the Fourier domain. The MTF is related to the spatial-domain point spread function (PSF). The PSF describes the normalized intensity distribution of a point-source image that has passed through the entire imaging system. The 3D $MTF(f_x, f_y, f_z)$ is defined in terms of the 3D system point-spread function, $p(x, y, z)$, which has been normalized to unity volume:

$$MTF(f_x, f_y, f_z) = \frac{|\mathcal{F}\{p(x, y, z)\}|}{\|\mathcal{F}\{p(x, y, z)\}\|}$$

where f_x, f_y, f_z are spatial-frequency variables corresponding to spatial-frequency variables $x, y,$ and z respectively.

This metric allows investigators to explore information across all spatial frequencies. We note that while the system MTF only applies to a shift-invariant system, a locally shift-invariant system can be analyzed in a specific region of interest using the *local* MTF. This local MTF denotes the frequency information in a small neighborhood surrounding a j^{th} voxel for a targeted investigation.

3.2 Noise

Noise in computed tomography is random variations in voxel values due to stochastic effects in the imaging system. Image noise can be influenced by a variety of factors including kVp, mA, exposure time, slice thickness, reconstruction filter, etc. For example, an increase in exposure time when acquiring data can have a positive impact on the Signal to Noise ratio (SNR). A limited description of noise that has gained widespread appeal is the use of standard deviation of voxel values as an easy and quick estimate of CT image noise. However, any number of noise textures may yield the same standard deviations.

The Noise Power Spectrum (NPS) describes noise in the Fourier domain for linear, shift-invariant systems with wide-sense stationary statistics. (Friedman et al., 2013). The three-dimensional NPS is defined as the three-dimensional Fourier transform of the system's autocovariance function, K:

$$NPS(f_x, f_y, f_z) = \mathcal{F}\{K(\tau_x, \tau_y, \tau_z)\}$$

where $\tau_x, \tau_y,$ and τ_z are the distances between voxels in each corresponding spatial direction. Similarly to local MTF, the local NPS can be obtained by investigating a small neighborhood surrounding voxel j of interest for locally stationary imaging systems.

3.3 Computing Detectability as a Performance Metric

Imaging performance is defined with respect to an imaging task. (Vennart, 1996) We choose a task-based detectability index that integrates noise, spatial resolution, a task function, and an observer model. (Gang 2014) The observer is the entity performing the imaging task. One such model is the nonprewhitening observer model which was successfully employed in various works and over a broad range of imaging conditions and tasks (Gang et al., 2011). For a non-pre-whitening observer, the detectability index relies upon the MTF and NPS as follows:

$$d_j'^2(\Omega) = \frac{\left\{\iiint \left[MTF_j(\Omega) \cdot W_{Task}\right]^2 df_x df_y df_z\right\}^2}{\iiint NPS_j(\Omega)\left[MTF_j(\Omega) \cdot W_{Task}\right]^2 df_x df_y df_z}$$

where W_{Task} is the detection task function chosen, defined by the Fourier transform of the difference between stimulus present and stimulus absent. This is described in section 4.3.

Note that this is an expression of local detectability with the local MTF and NPS estimates (denoted MTF_j and NPS_j). Also note that the local MTF and NPS change based upon acquisition parameters Ω including the orbit which is defined based on a system geometry specified in section 3.6. While MTF and NPS are typically measured using specialized

experiments, analytic expressions can be derived in some circumstances including for PL reconstruction (though local expression are required). This is discussed in Section 4.2.

3.4 Objective Functions Based on Task

The detectability index provides an objective for the design of an optimal source-detector orbit. The most straightforward approach to obtain a task-driven non-circular orbit is to assess the detectability at a single location within the specified region of interest. The detectability at this point would then the maximized:

$$\widehat{\Omega} = \underset{\Omega \in \Omega_{feasible}}{\operatorname{argmax}} \; d_j'^2(\Omega)$$

where the estimated parameter set $\widehat{\Omega}$ defines the optimal trajectory and is chosen from a range of possible parameters $\Omega_{feasible}$, limited by factors such as hardware or time restrictions.

While a single location optimized for maximum detectability is straightforward, it runs the risk of highly optimizing at only that location while sacrificing image quality in the larger region of interest. Exploring the detectabilities across various locations and maximizing the minimum detectability allows us to measure performance at a number of different locations in the field of view.

As such, the aim is to optimize detectability over a number of sampled locations in the region of interest, where the best orbit would yield the maximum minimum detectability over a range of locations, where:

$$\widehat{\Omega} = \underset{\Omega \in \Omega_{feasible}}{\operatorname{argmax}} \min_{j \in j_{ROI}} \left\{ d_j'^2(\Omega) \right\}$$

This stipulates that the location of minimum performance is what drives the design, as the worst part of the image is most often what needs to be improved most. Thus, the chosen orbit cannot be one that sacrifices image quality in one location for the improvement of another region. (Stayman and Siewerdsen, 2013) Doing so would mean that there is a "new" location of minimum performance that would need to be maximized.

3.5 Optimization Algorithms and Influences

The covariance matrix adaptation evolutionary strategy (CMA-ES) algorithm (Hansen and Ostermeier, 2001) employed in this work to estimate the solution to the maximum-minimum detectability objective function. To estimate the solution to the objective function, a population size of 40 was applied.

3.6 Parameterization of Prospective Orbit

The trajectory being investigated is parameterized using 9 periodic basis functions, constrained to tilts within a range limited by the C-Arm or available equipment. This work constrained the tilts to range within $\phi = -30° \; to \; 30°$, with a sparse sampling of gantry angles (for design purposes) every 10 degrees from $\theta = 1° \; to \; 360°$.

Chapter 4

Efficient Implementation of Task Driven Optimization

4.1 Computational Barriers to Practical Applications

In theory, one could spend the time and computational power to create an accurate MTF and NPS for every possible orbit of interest in order to determine which yields the best image with maximum detectability. This requires MTF/NPS computation for each of several locations around the region of interest for *each* potential orbit. Even for efficient optimization algorithms, thousands of MTF/NPS evaluations are required. This efficient computation is critical to implementation.

In order to overcome these inefficiencies, Fourier approximations for resolution and covariance can be employed to reduce computation times and effectively integrate task-based optimization of orbits with surgical workflow (discussed in Chapter 5). The previously-derived resolution and covariance predictors for penalized likelihood estimators can provide accurate approximations to the local resolution properties and covariance functions for tomographic systems given a good estimate of the mean measurements.

4.2 Expression for Local MTF and NPS

The local impulse response, or point spread function (PSF), measures the change in the mean reconstructed image due to perturbation of a particular pixel in a noiseless object. (Fessler, 1994) This gives a reasonable predictor of the local image resolution, while similarly, the local covariance can give a reliable predictor of local noise.

The local impulse response, l_j, and local covariance, c_j, can be approximated as

$$l_j \approx [\boldsymbol{A}(\Omega)^T \boldsymbol{D}\{\bar{y}\}\boldsymbol{A}(\Omega) + \beta \boldsymbol{R}]^{-1} \boldsymbol{A}(\Omega)^T \boldsymbol{D}\{\bar{y}\}\boldsymbol{A}(\Omega) e_j$$

$$c_j \approx [\boldsymbol{A}(\Omega)^T \boldsymbol{D}\{\bar{y}\}\boldsymbol{A}(\Omega) + \beta \boldsymbol{R}]^{-1} [\boldsymbol{A}(\Omega)^T cov\{y\}\boldsymbol{A}(\Omega)] \, [\boldsymbol{A}(\Omega)^T \boldsymbol{D}\{\bar{y}\}\boldsymbol{A}(\Omega) + \beta \boldsymbol{R}]^{-1} \, e_j$$

where e_j is the j^{th} unit vector of all zeros excluding the j^{th} location which is one, and $\boldsymbol{D}\{\cdot\}$ is a diagonal operator that places the vector argument on the diagonal of a matrix. These predictors demonstrate that the reconstructed image is dependent upon the system geometry ($\boldsymbol{A}(\Omega)$), the regularization ($\beta \boldsymbol{R}$), location (e_j) and the prior projection data itself ($\bar{y}$). (Fessler and Rogers, 1996; Fessler, 1996) These predictors allow us to predict local resolution and noise properties respectively given knowledge of patient anatomy through a preliminary CT scan (such as a preoperative CT). However, these calculations require inversions of large matrices that can be quite computationally taxing.

4.2.1 Fourier Approximation

It is worth noting that the projection of a single point $\boldsymbol{A}(\Omega)e_j$ is incredibly sparse. Intuitively, only the intensity and position of non-zero values within the Fourier space would be of interest. With approximately only one non-zero value per projection, it would make sense to

limit the size of the matrix by precomputing a much smaller linear operator, $\boldsymbol{L}_j$, which is only $N^3 \times N_{views}$ for each sampled location j.

For a more efficient alternative, the Fourier approximation may be applied as described by Qi and Leahy in 1999 and Stayman and Fessler in 2000, under a presumption of local shift-invariance and local stationarity. The approximate forms for local MTF and local NPS are as follows:

$$MTF_j = \mathcal{F}\{l_j\} \approx \frac{\mathcal{F}\{\boldsymbol{A}(\Omega)^T \boldsymbol{D}\{\bar{y}\}\boldsymbol{A}(\Omega)e_j\}}{\mathcal{F}\{\boldsymbol{A}(\Omega)^T \boldsymbol{D}\{\bar{y}\}\boldsymbol{A}(\Omega)e_j + \beta \boldsymbol{R}e_j\}}$$

$$NPS_j = \mathcal{F}\{c_j\} \approx \frac{\mathcal{F}\{\boldsymbol{A}(\Omega)^T \boldsymbol{D}\{\bar{y}\}\boldsymbol{A}(\Omega)e_j\}}{\mathcal{F}\{\boldsymbol{A}(\Omega)^T \boldsymbol{D}\{\bar{y}\}\boldsymbol{A}(\Omega)e_j + \beta \boldsymbol{R}e_j\}^2}$$

where $\mathcal{F}$ denotes the 3D discrete Fourier transform and relies upon an element-by-element vector division. The repeated forward and back projections as well as the Fourier transforms represent a significant computational burden, which proves limiting given the desired workflow for this work.

Recognizing the limitation of the above approximations, Stayman and Fessler observed that the projection-backprojection term is linear in the diagonal weighting such that:

$$\mathcal{F}\left\{\boldsymbol{A}(\Omega)^T \boldsymbol{D}\{w\}\boldsymbol{A}(\Omega)e_j\right\} = \boldsymbol{L}_j w$$

Where the linear operator $\boldsymbol{L}_j$ can be precomputed, stored, and later repeatedly called upon for fast calculation of MTF and NPS. (Stayman and Fessler, 2004)

4.3 Weighting Functions for Probe Locations in Area of Interest

Weighting functions are evaluated for each sampled location in an area of interest, describing the degree of attenuation that photons experience for a given source position (based upon gantry angle, θ and tilt angle, ϕ). The optimization algorithm will design a trajectory such that the photons rarely pass through the most highly attenuating path. The weighting functions will highly weigh the least attenuating beam paths to favor them during optimization. These weights assist in maximizing the minimum detectability across several locations.

4.4 Efficient Sampling

4.4.1 Volume Subsampling

Because the local PSF and local covariance are relatively compact spatially for reasonable imaging properties, one can constrain the volume used to calculate the local MTF and NPS. For this work, we focus on a PSF and covariance ROI of 50 x 50 x 50 voxels.

4.4.2 Angular Subsampling

Relevant frequency information is conveyed by the shape and orientation of each Fourier plane. Estimating these planes for each projection is computationally expensive, and a finely sampled MTF and NPS is not required for design purposes. The number of Fourier planes calculated can be limited by creating a sparser MTF/NPS by angular under-sampling. (Stayman and Fessler, 2004) Where before hundreds or thousands of Fourier planes were calculated, a mere 30 to 50 planes can convey sufficient relevant frequency information.

While a sparse MTF/NPS may be used in trajectory design, they can be further improved by interpolating the calculated planes to produce a smoother representation of the remaining Fourier space. In doing so, we artificially create a predictor that approximates the MTF that would have resulted from calculating planes for each projection across full 360-degree scan.

4.5 Thresholding of $\boldsymbol{L_j}$

As described in Section 4.2.1, $\boldsymbol{L_j}$ is a linear operator that quickly forms the $\mathcal{F}\{\boldsymbol{AWA}e_j\}$ required in MTF and NPS calculations. We note that there are many values of $\mathcal{F}\{\boldsymbol{AWA}e_j\}$ that are small or have little impact on a particular d' evaluation. This is attributed to some frequencies in the task function being near zero. Thus, we may select a threshold and not evaluating all frequencies, making the Fourier space approximation even more efficient by the elimination of data points that have little to no contribution to the detectability (d') calculation.

Chapter 5

Workflow

5.1 Scout Views and Preliminary Reconstruction

For interventional procedures, an initial CT image of the area of interest is often acquired to assess the operative site and create a surgical plan. This CT image contains vital prior anatomical information that can be used in our proposed interventional workflow. While useful, the preoperative CT can significantly differ from the intraoperative state. For example, positioning is different and anatomical damage is possible. Notably, hardware is often introduced and imaging around such hardware is typically the goal. To accommodate such changes, a scout scan of only two perpendicular views is proposed to approximate hardware locations/volume as well as patient registration. The anteroposterior and lateral views were chosen for simplicity. In this work, we focus on determining the localization of an embolization coil when treating brain aneurysms, but it is possible to expand this work to encompass various other implant models as well. This can include pedicle screws and rods in spinal surgery or prostatic artery embolization.

To determine the rough appearance of the embolization coil, we segment the dense metal coil from a two-projection reconstruction, refine the model, and place the model within a registered patient volume from the preoperative CT. The two-projection reconstruction is an incredibly

rough approximation of the patient anatomy and coil structure, and will yield a streaky, "boxier" representation of the coil. This preliminary image is produced via simple FDK reconstruction.

(a)

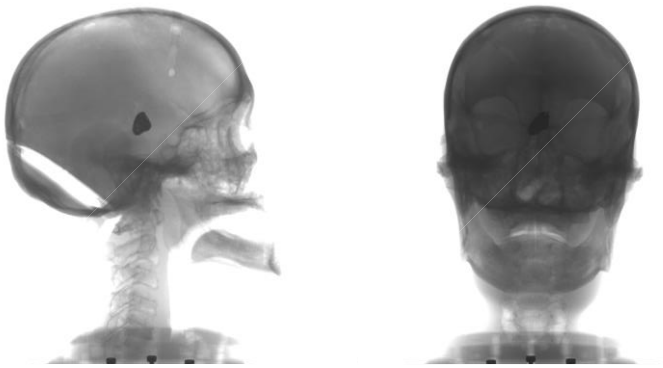

(b)

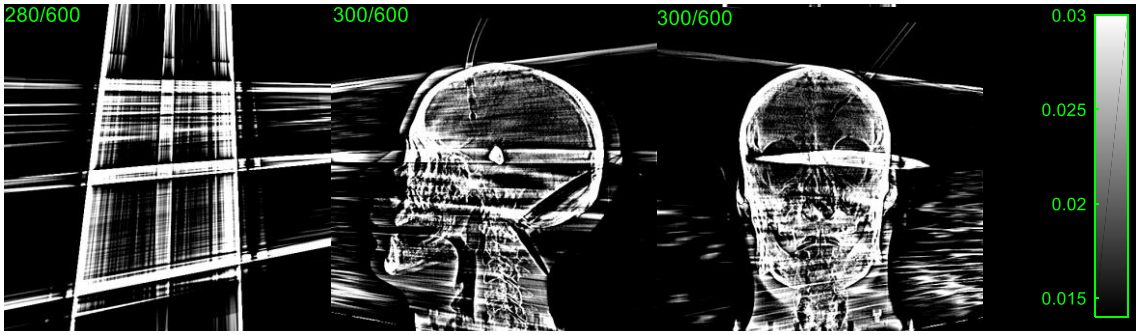

Figure 1: (a) AP/Lateral views of benchtop projection acquisition of head phantom with implanted coil (b) FBP reconstruction using only the two projections in Figure 1(a). This image illustrates the streaks prevalent in the reconstruction and gives an idea of the potential difficulties in segmenting the coil out from the head.

5.2 Segmentation of Coil

We focus on cases where the significant change between the preoperative CT scan and the intraoperative CT is the addition of highly attenuating metal. Other changes such as changes in hardware and changes in anatomy from the operation are usually relatively negligible in comparison; the metal hardware tends to dominate image quality and imaging performance

issues. As such, the segmentation and positioning of the embolization coil is sufficient to model the artifacts likely to occur in the intraoperative image. The highly-attenuating metal implants have a significant impact on the statistics of the data (weighting of the projections $\bar{y}$ in the predictors). The workflow for the four-step segmentation is described below.

1. Segmentation of the coil was accomplished first by custom thresholding to create an image with only the metal coil, its artifacts, and the denser bone within the skull. We selected a value of 2% above of the maximum intensity in the two-view reconstruction.

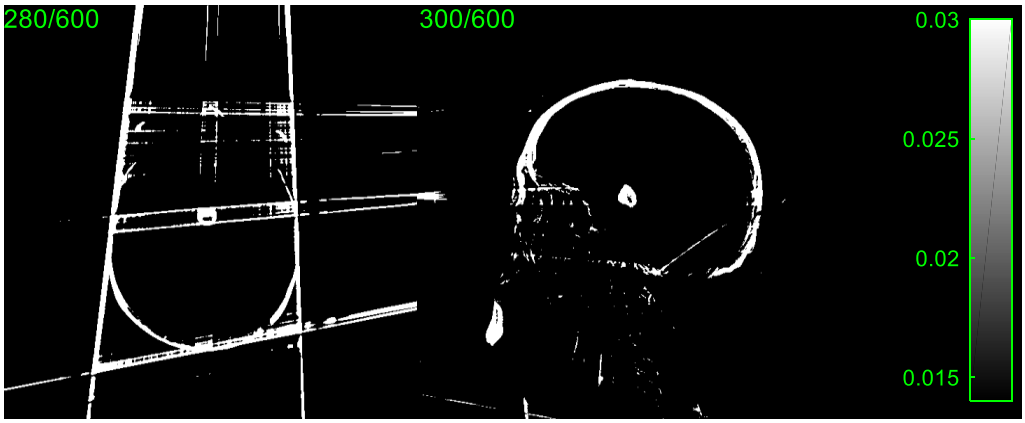

Figure 2: Result of initial thresholding of two-projection reconstruction. The axial view demonstrates that a simple thresholding preserves unwanted edges and streaks, while the shape and size of the coil are well preserved. Higher thresholding may eliminate some streaks but would affect the integrity of the coil model.

2. A diamond structuring element (MATLAB's *strel* function) is used to better define the structure of the coil by disconnecting the streaks from the coil and if thin enough, eliminate them altogether. For best results, MATLAB's eroding function, *imerode,* was applied twice once with structure size 2 and subsequently with structure size 1. The two steps ensure that the coil shape and size remain relatively intact.

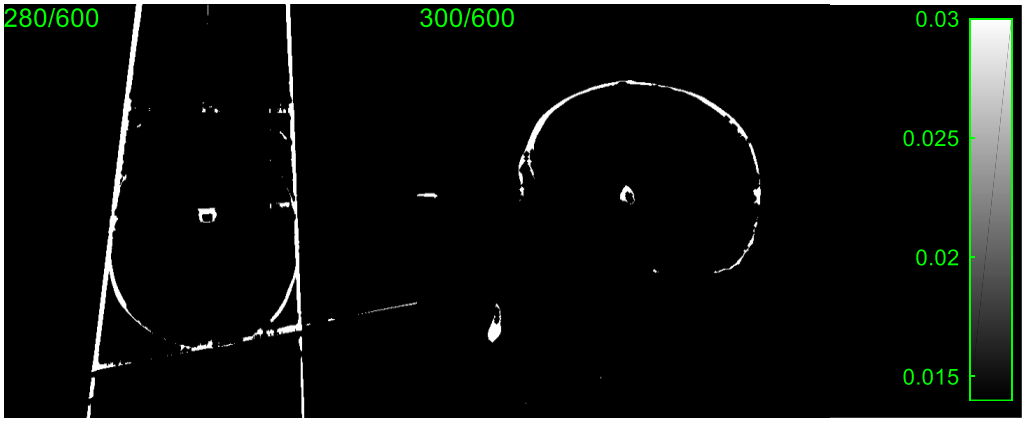

Figure 3: Result of two erosion steps. The second step in the segmentation workflow largely removes unwanted artifacts.

3. The coil was then isolated using two methods. Knowing the relative structure of a human head, and assuming that the patient was centered reasonably well, an ellipsoid centered within the skull can be extracted such that the majority of the bone in the segmentation is removed. Following this, we can make use of MATLAB's connected components function (*bwconncomp*) to extract the largest cluster of connected voxels – our coil, while removing all other extraneous segments leftover.

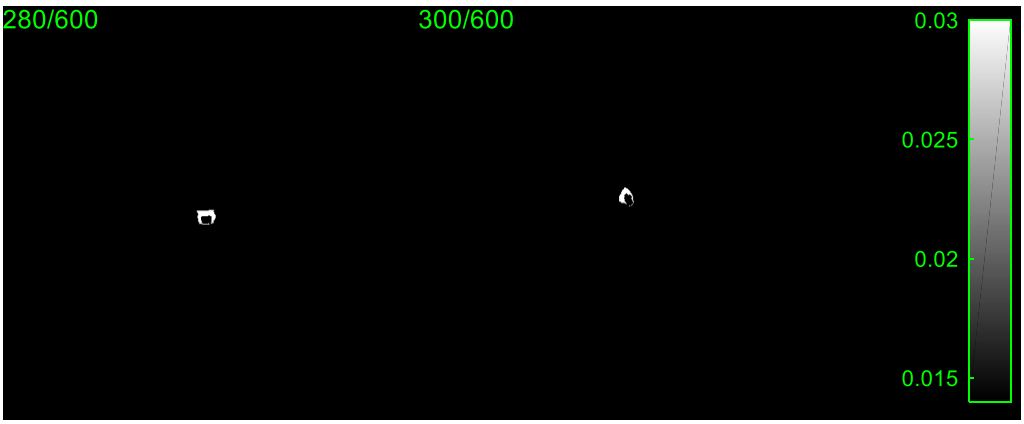

Figure 4: Result of connected component method. The coil is isolated, but now appears hollow and incomplete.

4. MATLAB's, *imdilate*, function can then be used twice to restore the size and shape of the coil. Once with a diamond and then again with a sphere to round off any sharp edges left from the two-projection reconstruction.

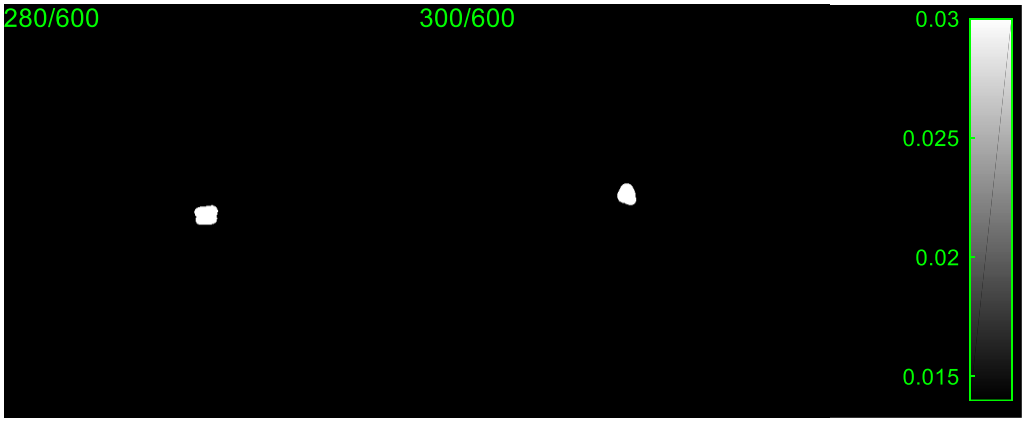

Figure 5: Result of two dilation steps. The coil is still boxy but has filled in and remains true to its original shape and size.

The shape and size of the coil are an adequate approximation of the true coil, as discussed in the verification chapter, in section 6.2. While the axial view of the coil is somewhat square as a result of only using two projections to generate the coil model, the segmentation is sufficient for our purposes. Following this segmentation, the desired metal coil's attenuation can then easily be substituted in the binary image created.

5.3 Registration of Patient Coordinates

Following coil segmentation from the preoperative CT scan, a rigid registration of the patient is performed. This permits specification of the designed trajectory so that it is transformed into the intraoperative patient's coordinates. This is accomplished in this work using a 3D-3D registration program, Elastix. (Klein et al., 2010; Shamonin et al. 2014)

5.4 Optimization Workflow

5.4.1 Task Function Construction

For the intracranial embolization case, we choose a task function emulating a hemorrhage detection task. The task function was constructed by creating a Gaussian stimulus the size of a 3 mm hemorrhages. A Fourier transform of the bleed specifies the frequency information relevant to the detection task.

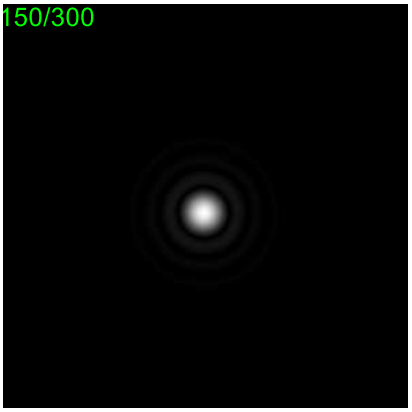

Figure 6: Frequency domain task function for Gaussian detection.

5.4.2 Projection Generation

Projection data was generated using a separable footprint projector. (Long et al., 2010) Poisson noise was then artificially added to simulation data to emulate physical data with quantum noise.

5.4.3 Calculate Weighting Functions

Locations of interest are sampled from the area where improved detectability is desired. For the embolization scenario, we select a sphere of 30 sample points around the coil. The placement of these points of interest are depicted in the following figure by purple asterisks.

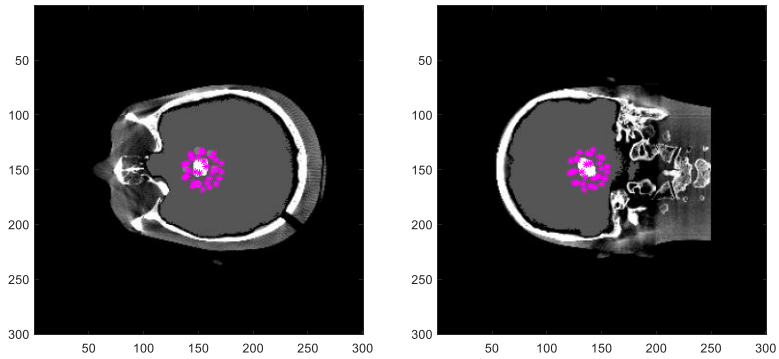

Figure 7: Sphere of sampled points surrounding coil. This illustrates the ideal position and size of the sphere, encompassing the area with bleeds where improved detectability is most desired.

For each sampled location around the region of interest, statistical weighting functions must be generated from the raw projection data. These functions depict the statistical weights for all source locations at each task location (sampled within the region of interest), and how noisy each of those positions are. Precomputation of these weights permits efficient optimization but also illustrates projections and location combinations that lead to low weights (source positions leading to poor statistics). For example, if an area is highly attenuating, the weight map will display a dark area and the optimized orbit will likely avoid those positions for all possible sampled locations and strive to pass through the lighter areas of each weight map.

The weight maps below display the attenuations through a specific point of interest given a location of the source, defined by the gantry angle θ (from 0 to 359 degrees, sampled every 10 degrees) and the out of plane tilt angle ϕ (from -30 to 30 degrees, sampled every 2 degrees). The

optimal orbit is plotted on these weight maps to illustrate how the orbits with the best detectability indices avoid the highly attenuating, darker areas of the weight map such as the metal coil (black circle) and the areas where the source position passes photons through a lot of dense bone (dark grey). Instead, the orbit attempts to choose source positions where the photons are less attenuated (the lighter grey/white areas).

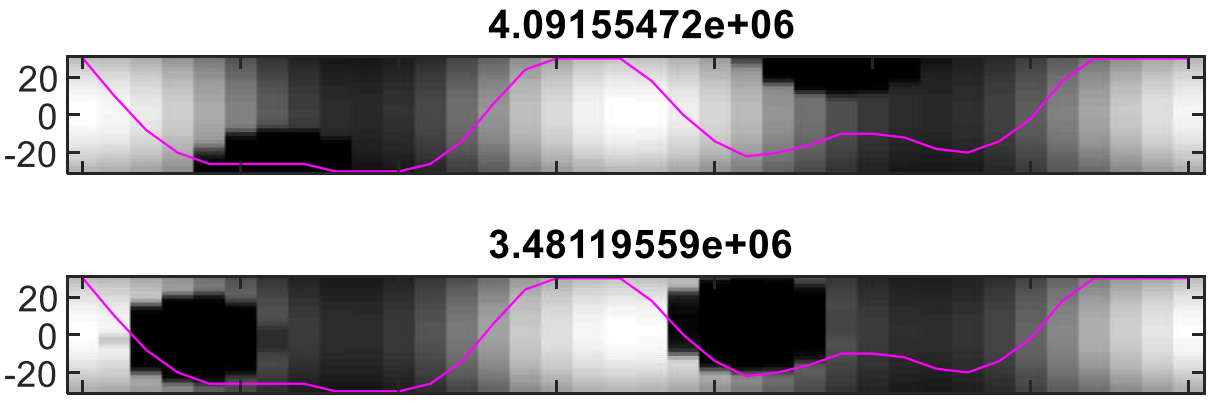

Figure 8: Weight maps for two sampled locations displaying the attenuations across an array of source-detector positions, defined by the gantry angle (x axis) and out of plane angle (the tilt angle). The orbit optimized across 30 sampled points for maximum minimum detectability is plotted atop the two weight maps. The detectability indices of the orbit for the sampled point are displayed above their respective weight maps.

5.4.4 CMA-ES Optimization

Using a population size of 40, the CMA-ES algorithm uses the multi-location objective to maximize the minimum detectability. The resultant vector of the optimization represents the 9 coefficients of the 9 periodic basis functions.

At least 40 optimizations were conducted in order to ensure that the best possible solution was reached. This was achieved by initializing each optimization with a different seed, and choosing the orbit with the highest minimum detectability. For more complex cases, up to 100 optimizations were executed.

5.5 Reconstruction Using Parameterized Orbit

The parameterized orbit obtained from the detectability optimization is easily converted into a series of gantry and tilt angles that can be used to calculate projection matrices based upon the source-detector positions. Once these projection matrices have been obtained, continuing with a standard PL reconstruction is fairly straight forward. The reconstruction was generated by iteratively using ordered subsets; the voxels were all updated simultaneously and projected forward and back once to compute the likelihood gradient. The penalty gradient and curvature were computed for every iteration (no precomputation of the curvature). Both FBP and zero image initializations were explored for this work.

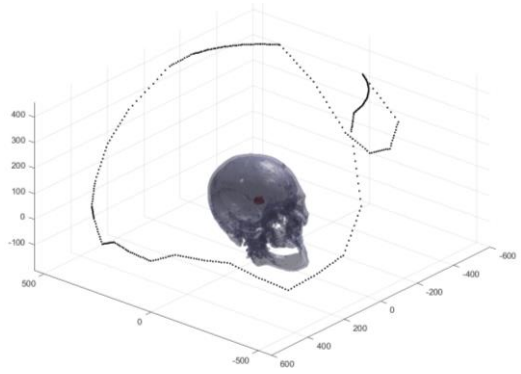

Figure 9: An example of a designed trajectory, with all 360 source positions plotted as black points surrounding a head phantom with a red coil embedded.

Chapter 6

Validation

This chapter details the investigations conducted of the various workflow stages and overall performance of the framework proposed in Chapter 5. We believe the contents of this chapter validate Chapter 7 results for the specific tasks investigated.

6.1 MTF Considerations

The optimization and detectability index rely upon the accuracy of the MTF approximation used to describe the object and task. As such, it is important that the estimated MTF's orientation and magnitude be compared against the "proper" MTF, constructed by determining the point spread function (measures the change in the mean reconstructed image due to perturbation of a particular pixel in the noiseless object).

This comparison determines whether there are any inconsistencies in geometry or reconstruction. In order to validate the accuracy of the MTF approximation, "true" MTFs were calculated using a head phantom with the metal coil artificially implanted. A stimulus point was placed within the area of interest surrounding the metal at one of the sampled location points used for approximations of MTF and NPS. A phantom with and without a stimulus were forward projected and subsequently iteratively reconstructed using the same parameters for both PL reconstructions. The Fourier transform of the difference between the two images yields the MTF,

in other words, the Fourier transform of the PSF. This is referred to as the “long method.” As previously discussed, subsampling Fourier space allows for more efficient computation of an ideal trajectory. Therefore, we crop the reconstructed images to the desired subsampled dimensions of the approximation before the Fourier transform of the “long method” cropped PSF.

To approximate the MTF, the same single voxel used for the “long” method was chosen for the MTF approximation. The estimated MTF was calculated using the procedure described in Section 4.2.

As shown below, the estimated MTF and the true MTF contain the same frequency information for a variety of stimulus locations surrounding a coil, with null cones of roughly the same size and shape within the Fourier space. It should be noted that the magnitudes of the MTFs are different but were easily normalized on a case-by-case basis throughout this work to produce roughly identical MTFs as the intensity distribution across the non-zero values were approximately equal. Furthermore, these results were generated without subsampling the MTF for the long case, but would yield similar results, simply at a lesser resolution – similar to those seen in the estimated case. For reference, the locations (as PSFs for easier viewing) referred to in figure 10(a) for each MTF are labelled in figure 10(b).

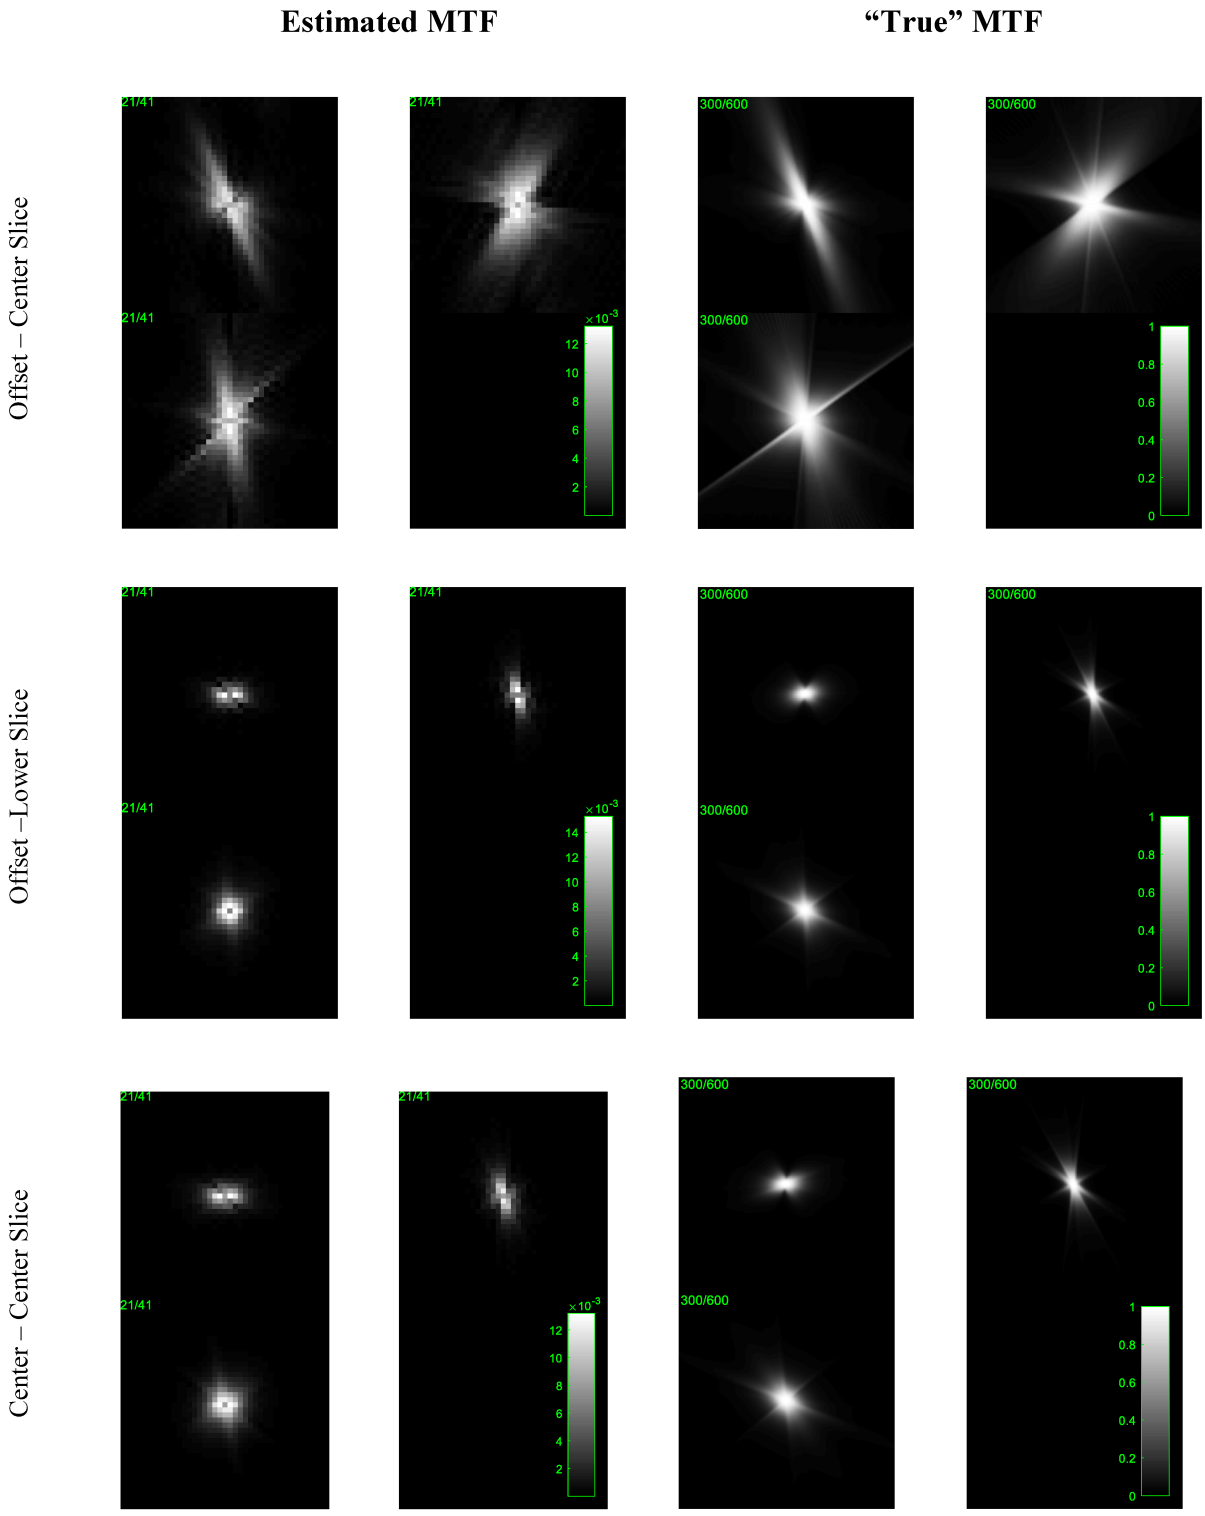
Estimated MTF
"True" MTF
Offset – Center Slice
Offset –Lower Slice
Center – Center Slice
21/41
300/600
×10^-3

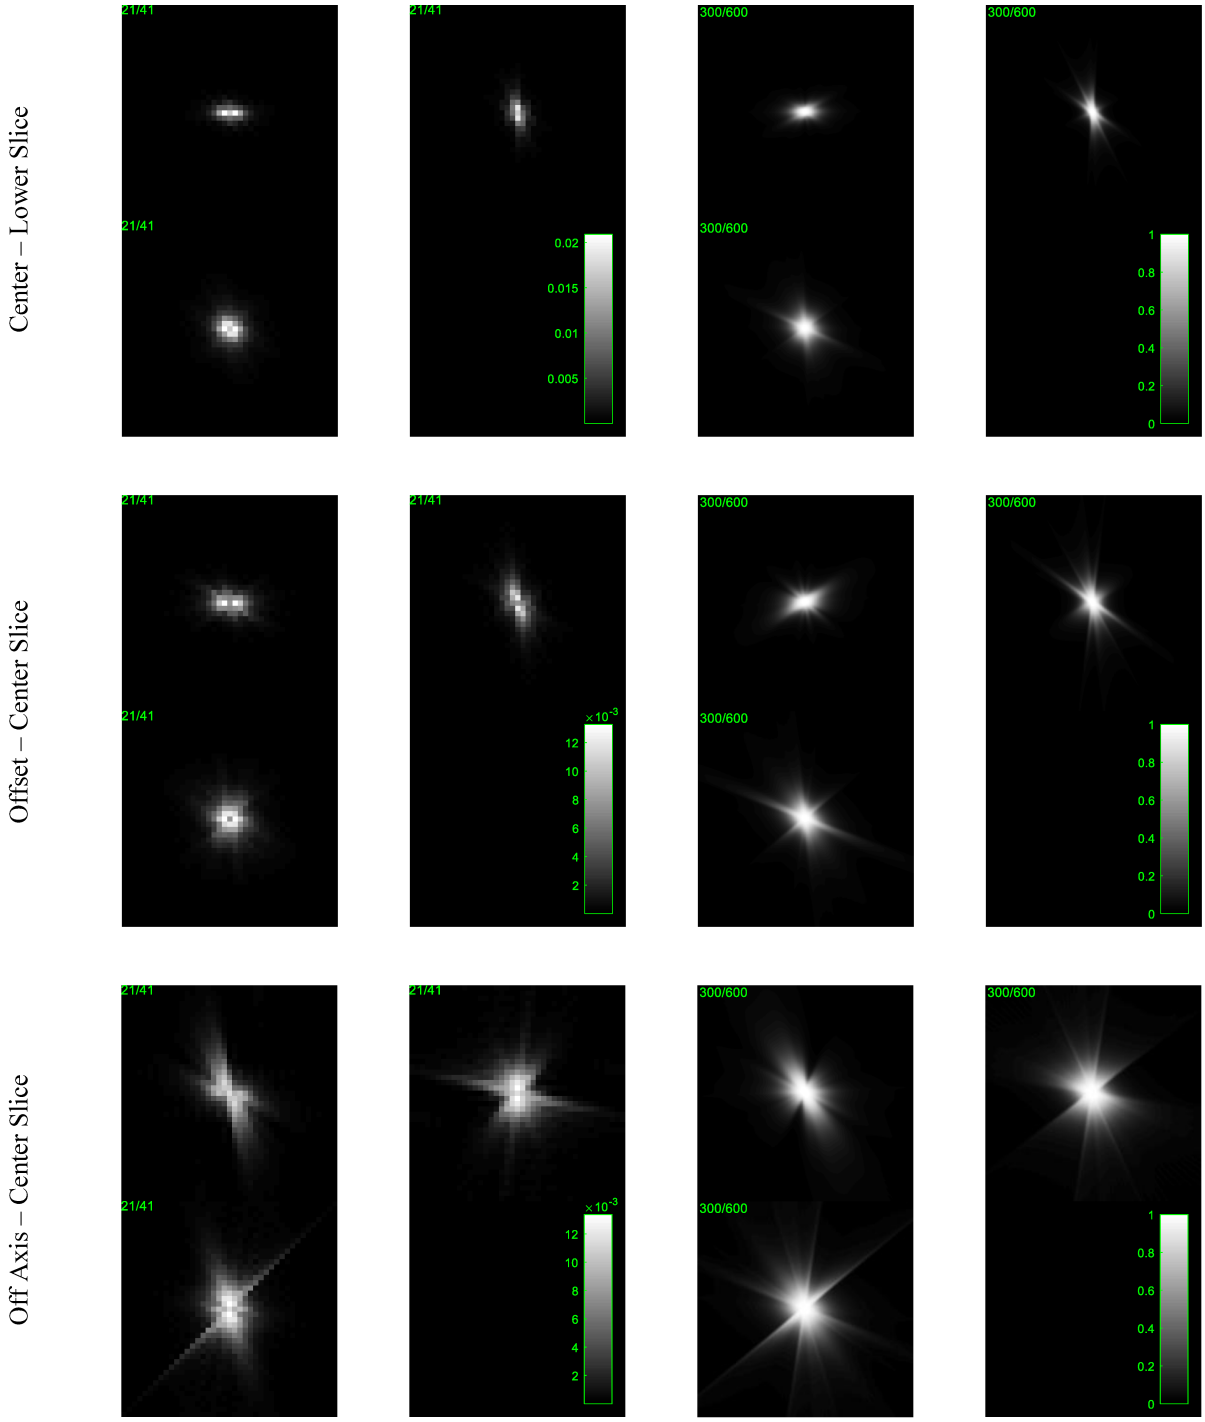
Center – Lower Slice
Offset – Center Slice
Off Axis – Center Slice
21/41
300/600
0.02
0.015
0.01
0.005
×10⁻³
12
10
8
6
4
2
1
0.8
0.6
0.4
0.2
0

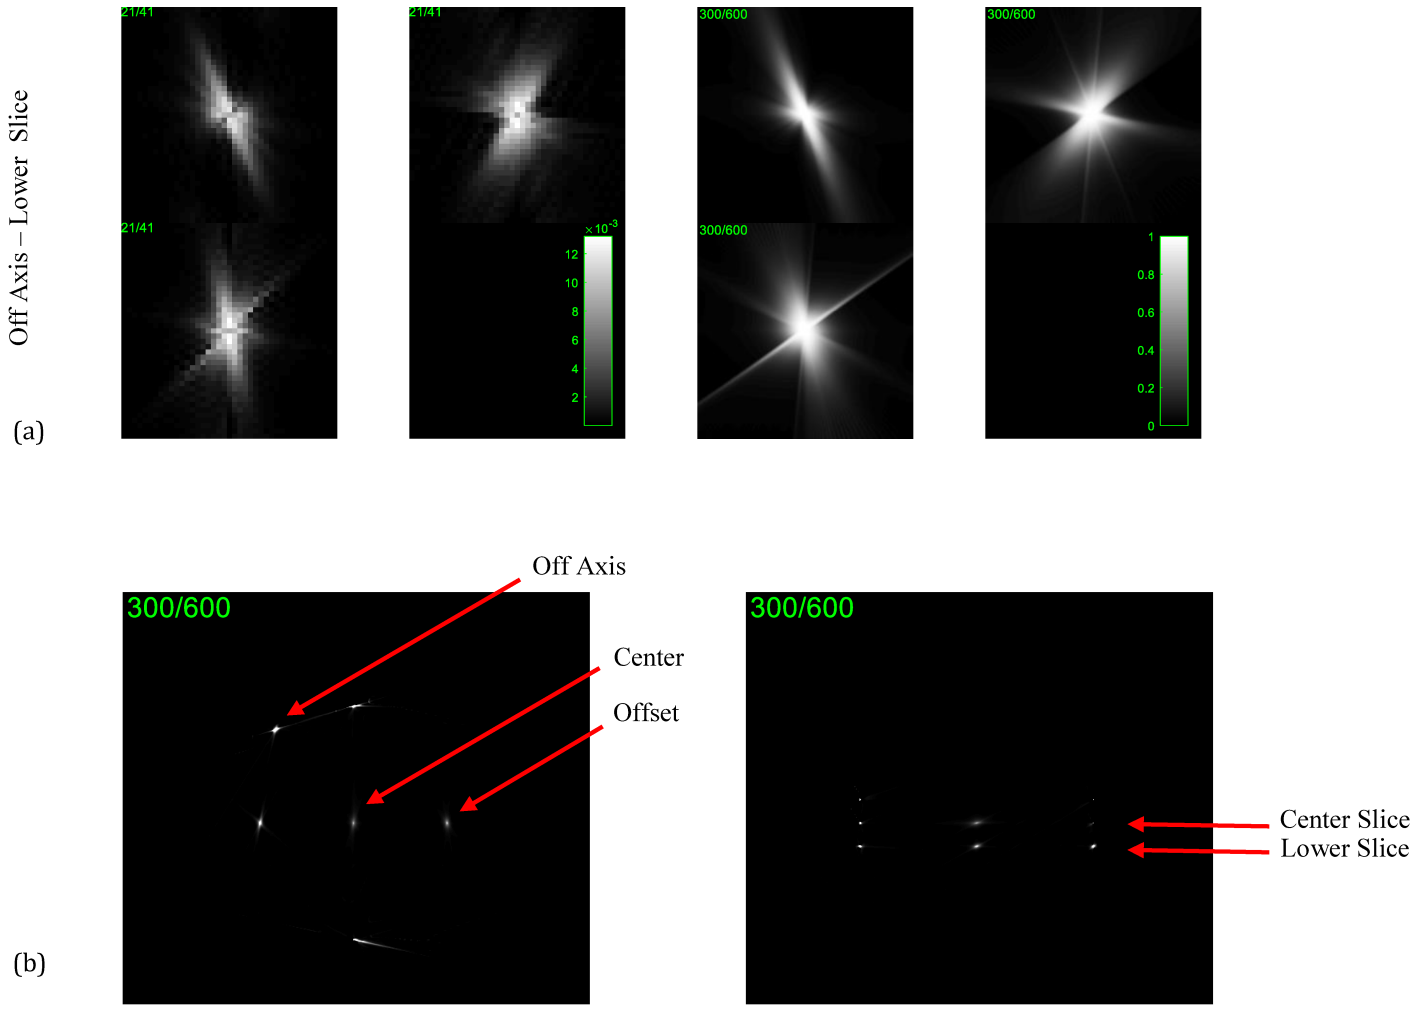

Figure 10: (a) Several stimulus locations on two separate slices were investigated to match estimated MTFs to true MTFs. (b) The stimulus locations are illustrated as the PSFs.

6.2 Verifying Coil Segmentation & 3D-3D Registration

For the purposes of testing the validity of the segmentation method, a full reconstruction (with 360 projections) of the skull and metal coil (or any implant) can be acquired. The segmented coil in a zero volume can then be subtracted from the "true" reconstruction to qualitatively determine whether too much or too little of the metal has been eroded away, and whether the shape of the segmented coil properly represents the profile of the original coil.

To verify that the coil has been properly placed, the coil-less phantom that had the coil segmentation artificially implanted can be compared to the full reconstruction of the bench data. The transformation will have been determined by a rigid registration reliant upon the skull, so we expect to see approximately zero values around the skull and neck, while the coil should line up relatively nicely excluding the slight differences in shape between the true coil and the segmentation. This is merely for verification purposes. As detailed in Chapter 5, the true workflow requires that the images be registered using a basic FBP reconstruction of the two scout views collected and the prior information available (the image of the head excluding the coil). As the skull information is mostly present, an adequate registration is possible.

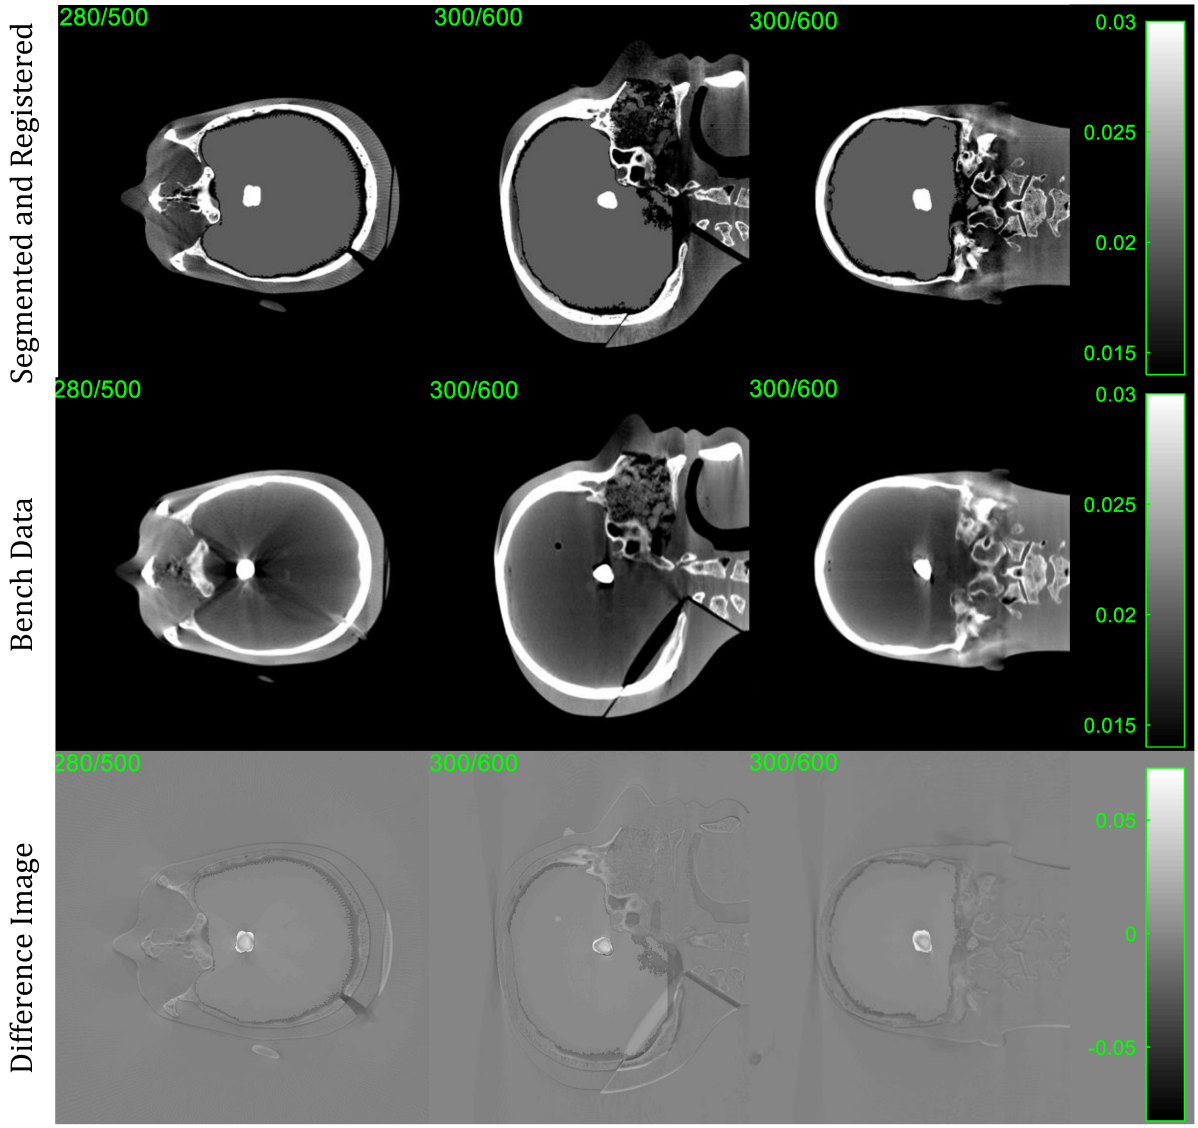

Figure 11: Registration verification using a difference image. The images demonstrate proper alignment of the skull and coil.

Zooming into the coil itself as seen in the figure below, it is clear that the segmentation of the coil mostly represents the shape and size of the metal such that we may approximate its effects on a reconstruction. The difference in attenuation between the coil and the model can be mostly attributed to the reconstruction method – FBP would never be able to estimate the

attenuation of the coil whereas a method like PL or PWLS would fare far better as the errors in attenuation are due to beam hardening.

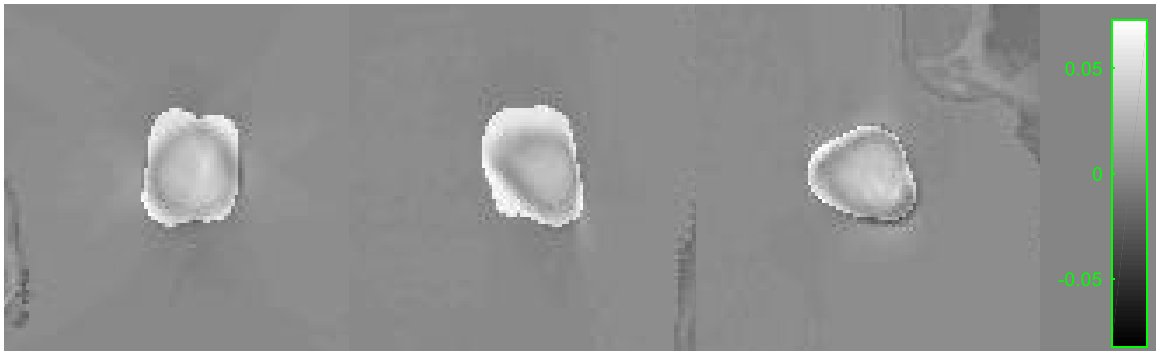

Figure 12: A magnified, cropped view of the difference image centered on the metal coil. The images demonstrate an acceptable approximation of the coil shape and size after segmentation in all three views.

The images demonstrate that the skull and coil are adequately aligned, with some error primarily in the artificial segmentation of the brain for the simulation data.

In assessing and validating these factors, we can be confident that the results presented in the following chapter are accurate.

Chapter 7

Results and Discussion

7.1 Optimized Orbits

Following multiple optimizations, the trajectory with the maximum minimum detectability was chosen to reconstruct the data, as previously described in section 5.4.4.

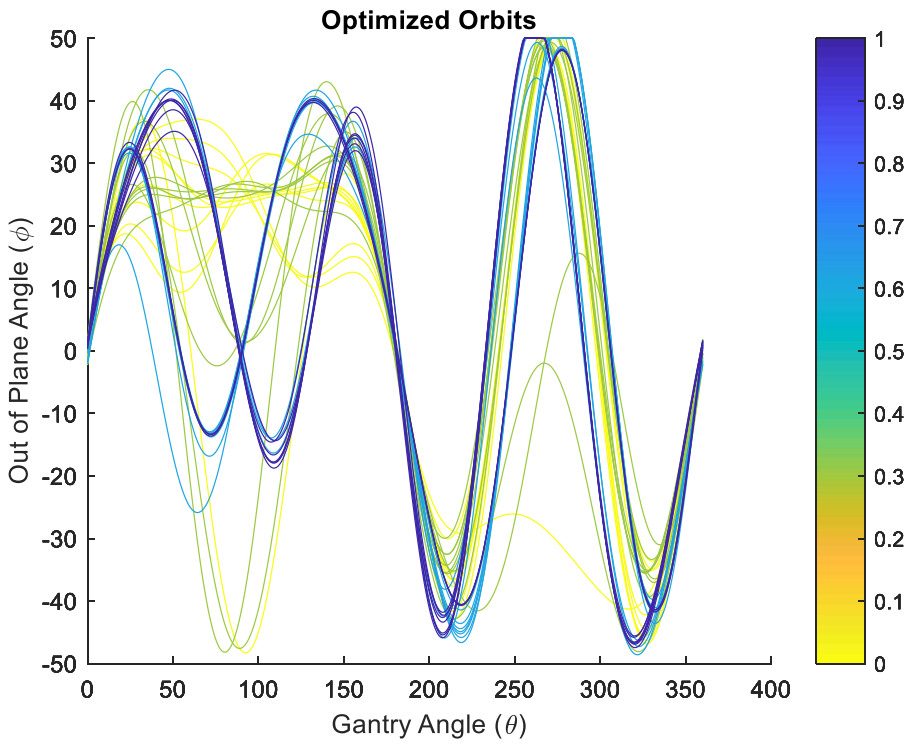

Figure 13: A plot of 40 designed orbits, with the detectabilities encoded by color – darker colors assigned to orbits with higher detectability indices. The figure illustrates the similarity between the trajectories with the highest detectabilities.

The figure above (designed for the pelvis phantom) demonstrates that the dark blue orbits with the highest (top 20%) detectabilities are nearly identical, exhibiting far more clustering than orbits with lower detectabilities. We conclude that the trajectory with the highest detectability is a relatively robust solution.

7.2 Reconstruction for Optimized Orbit

Once an ideal orbit has been converged upon, we can begin the reconstruction process. This process requires fine tuning several factors and techniques, described in this section. In order to efficiently explore the effects of regularization parameters and their effects upon image quality for the given phantom, a simple elliptic cylindrical phantom was generated containing a highly-attenuating ellipsoid and surrounding lower contrast spheres, as shown below.

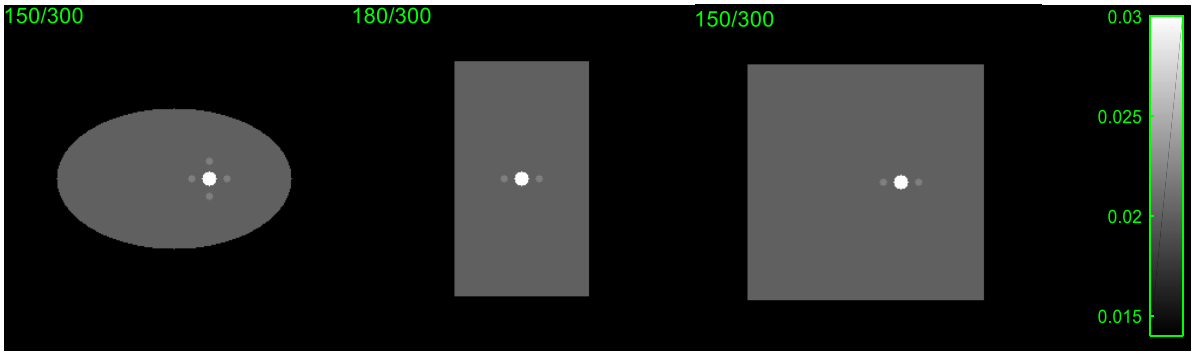

Figure 14: Three views of simple phantom, center sliced through the metal coil.

This phantom is roughly the shape and volume of a human head, with the ellipsoid modelling an embolization coil and the low contrast spheres representing brain bleeds that can be difficult to visualize following embolization procedures.

7.2.1 Number of Iterations

In order to confirm that the penalized likelihood reconstruction has converged upon a solution, a reconstruction with successful parameters was set to run for many iterations to confirm that the difference between the reconstructions is negligible or insignificant enough when compared to the time dedicated to the reconstruction.

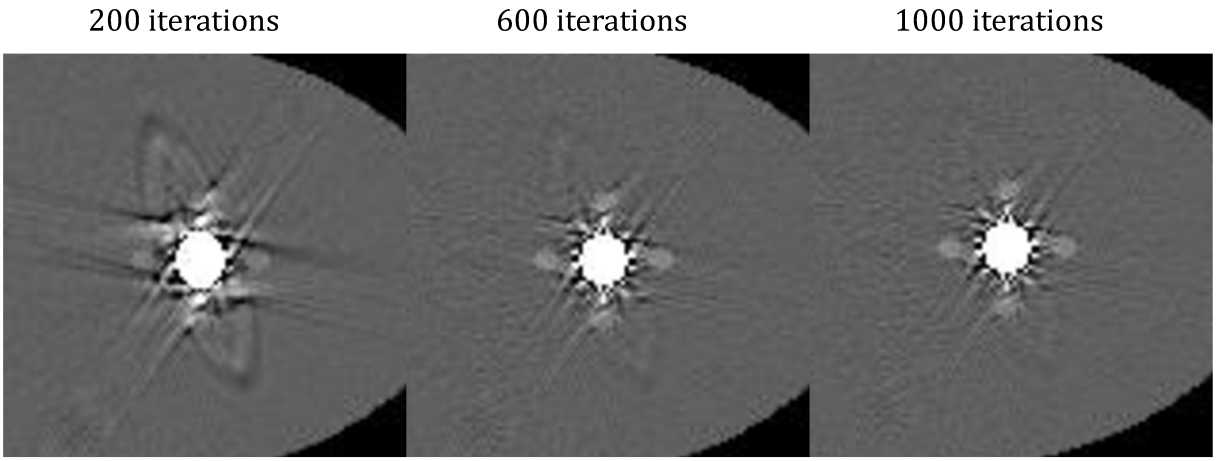

Figure 15: Multiple maximum iterations were explored, with 1000 iterations yielding the best results but at the cost of lengthy reconstruction times. 600 iterations was found to be an acceptable compromise between reduced artifacts and efficiency.

These results demonstrate a huge reduction in artifacts between iterations 200 and 600, but with diminishing returns regarding image quality between 600 and 1000 iterations. Therefore, it can be concluded that 600 iterations is sufficient to assess improvements in image quality of a designed orbit over that of a circular orbit in PL reconstructions of a similar model, such as that of a human head.

7.2.2 Noise Realizations

Different Poisson noise realizations were evaluated to confirm that improved visibility of the bleeds was true across various noise realizations. The results demonstrated that the results produced by the optimized trajectory were not coincidentally effective for an image with a specific noise level but for any noise realization.

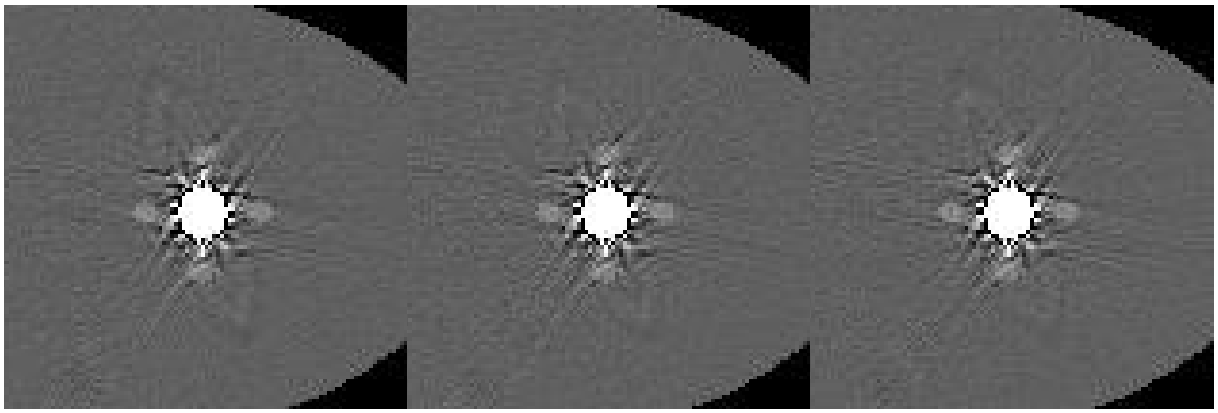

Figure 16: Various noise realizations were explored to ensure the results of the designed trajectory carried true across each. The figure illustrates that this is the case.

7.2.3 Initializations

A good initial guess helps to speed up convergence for penalized likelihood reconstruction. In general, the closer the starting point is to the final image estimation, the fewer iterations are required. However, if the algorithm is initialized with artificial edges or streaks, then the reconstruction may take many iterations to eliminate these false objects. In order to ensure that the algorithm is provided with a good first guess, three initializations were tested: truth, zeros, and an FBP reconstruction.

Similarly, to verify sufficient convergence and to ensure image features are a product of a converged reconstruction process, rather than being intermediate errors from insufficient convergence, we initialized the PL reconstructions with the “true” phantom volume. “True”

initialization was chosen for this comparison because the PL reconstruction will converge fastest from the truth, and the artifacts likely to diminish or be eliminated should be clear despite the limited number of iterations used for this experiment. As such, the initial "guess" is the best guess investigated and any deviations from this are likely attributed to the reconstruction workflow.

The following results were created without noise and with identical reconstruction parameters, including quadratic penalties where β = 1 E 6. This β value was empirically found to be a good regularization parameter for this data set when employing quadratic penalties. Section 7.2.1 provides evidence supporting the decision to reconstruct images using 600 iterations. This is the case for all of the following reconstructions unless otherwise stated.

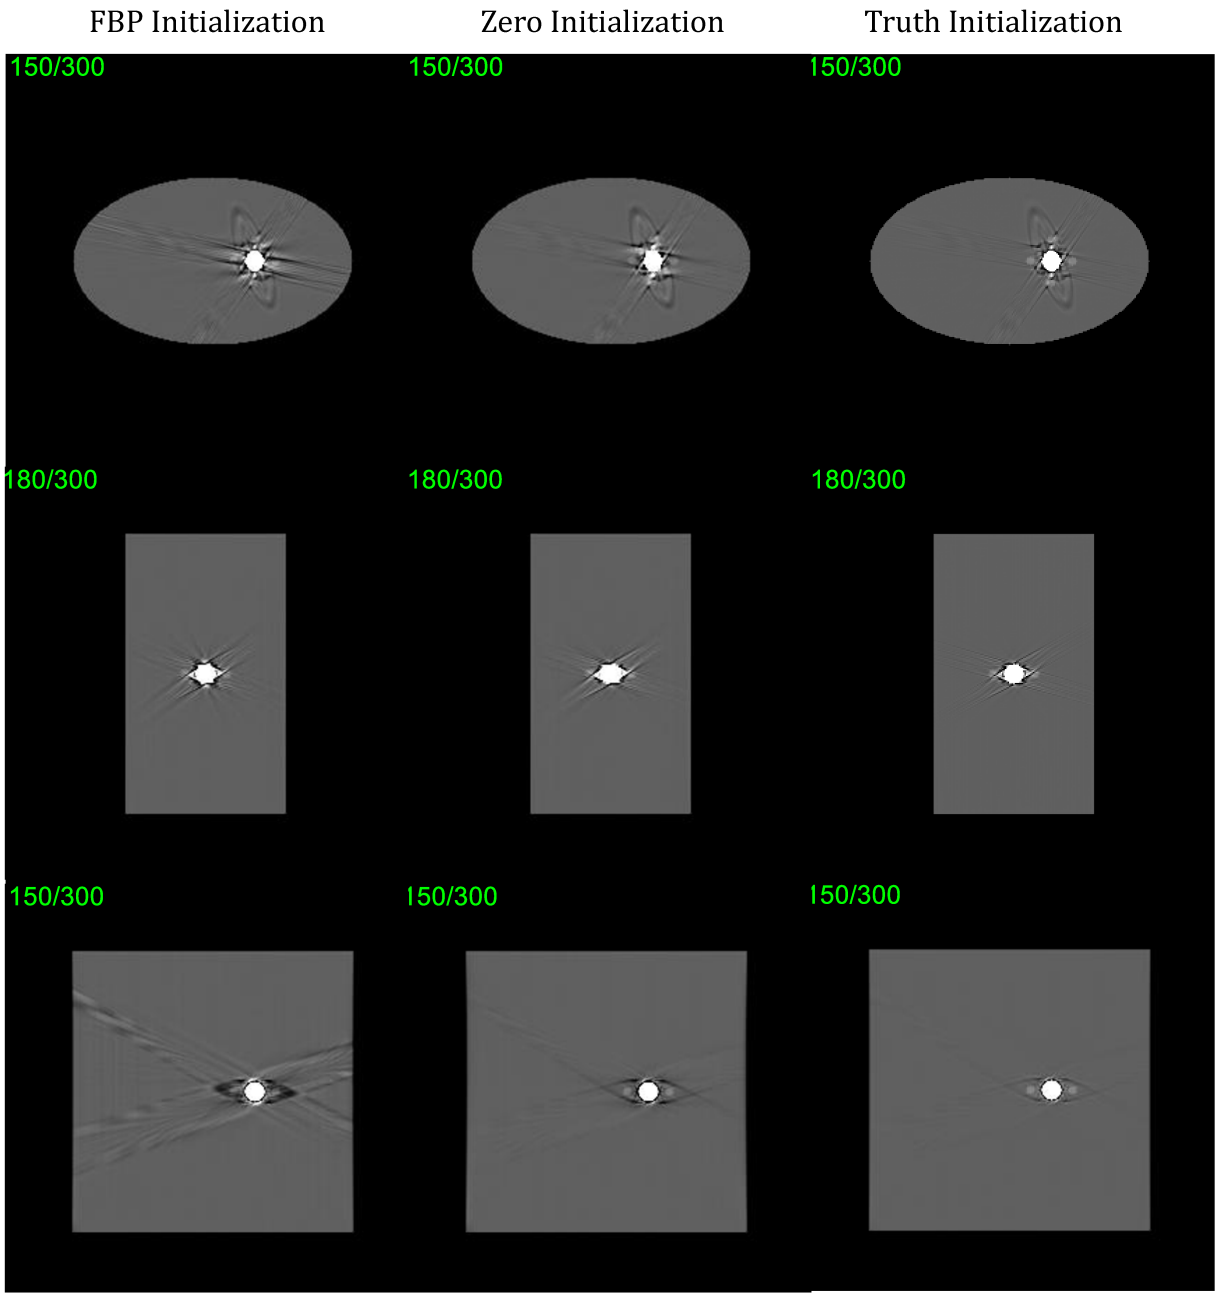

Figure 17: Three initializations were investigated to determine the effect of the initializations on the artifacts present, with results improving from left to right.

FBP, zero, and truth initializations were explored to assess their impact upon an image of the same object, with the same acquisition, reconstruction parameters, and noise kernels. Importantly, significant differences in image quality were evident despite reconstruction with the same number of iterations. It was found that the PL reconstruction initialized with FBP produces

an image with significant streak artifacts and with very poor visibility for the low contrast bleeds near the metal coil. These streaks can be easily avoided by initializing with a matrix of zeros. This initialization most closely resembles the PL reconstruction initialized with the true phantom, which results in the least severe artifact of the three initializations investigated. While some streak artifacts are still present, they have been drastically reduced, especially given that the images were reconstructed using only 200 iterations.

7.2.4 Quadratic vs Nonquadratic Penalties

In general, a roughness-penalized, likelihood-based estimator can be expressed as:

$$\hat{\mu} = \arg\max_{\mu} \log L(y;\mu) - R(\mu) = \arg\max_{\mu} \log L(y;\mu) - \beta_R \|\boldsymbol{\Psi}_R \mu\|^{p_R}$$

Where $R(\mu)$ denotes an arbitrary regularizer, and we have chosen a specific roughness penalty which incorporates the image gradient operator $\boldsymbol{\Psi}_{\mathrm{R}}$, a p-norm operator. As previously discussed, the parameter β_R is a scalar that allows for a manual tuning of the noise-resolution tradeoff.

The differences between the different p-norm operators are significant enough to merit discussion. The p-norm penalties explored include total variation (TV) style regularization where $p_R = 1$, quadratic regularization where $p_R = 2$, and Huber penalty. The quadratic regularizer applies blur relatively uniformly across the volume, while the TV regularization allows larger differences between voxels and can result in highly nonuniform smoothing with strong edge preservation. (Vogel and Oman, 1996; Stayman 2014).

This work primarily explored the effects of Quadratic vs. Huber penalty in a phantom with low-contrast bleeds located near a highly attenuating ellipsoid. The following figure

presents the best reconstructions achieved for each penalty studied over an investigation of a range of $\beta = 1\ E\ 3\ to\ 1\ E\ 8$ for quadratic penalties and $\beta = 1\ E\ 3\ to\ 1\ E\ 8\ , \delta = 1\ E - 5\ to\ 1$ for Huber penalties. As previously stated, both images were created over the course of 600 iterations. The quadratic data set performed best with $\beta = 1\ E\ 6$, while the Huber data performed best with $\beta = 1\ E\ 5\ , \delta = 1\ E - 2$.

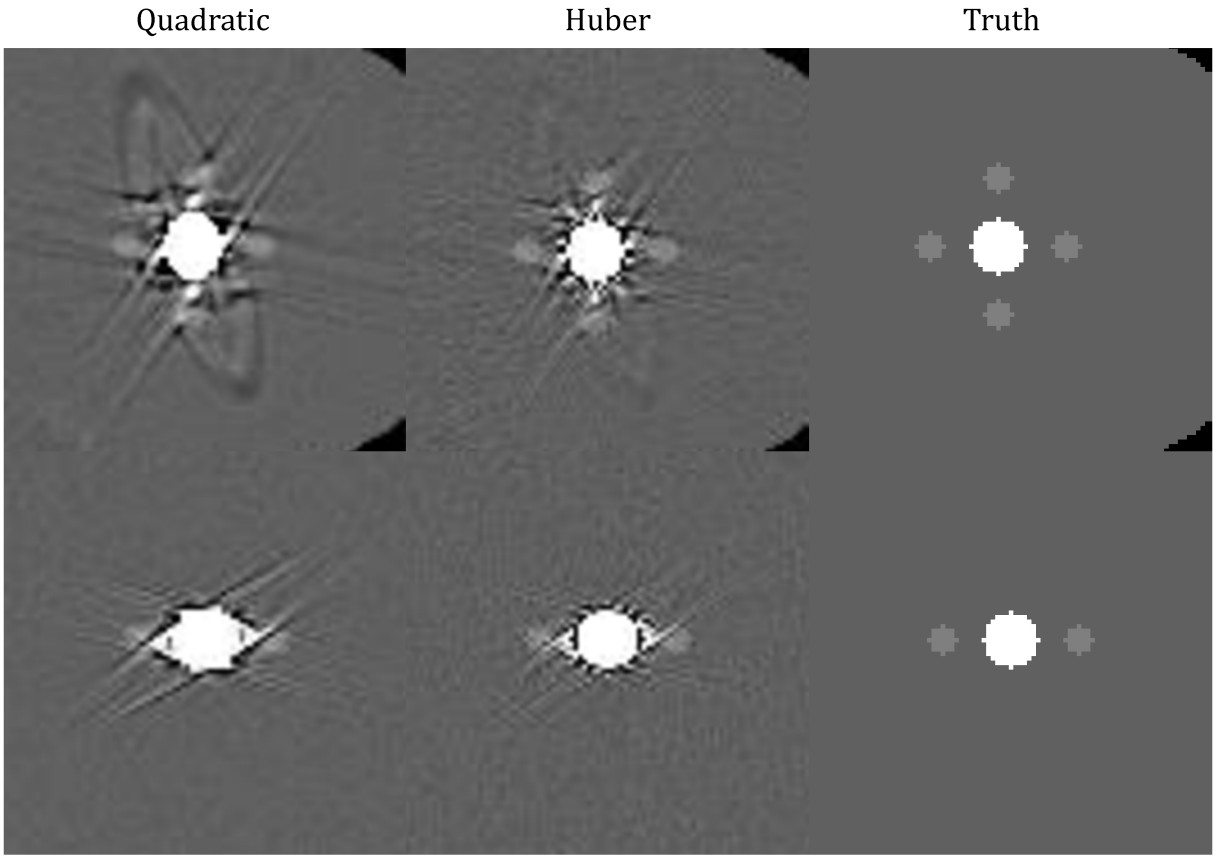

Figure 18: Effects of two different penalty functions on visibility of low-contrast bleeds. The images demonstrate a significant reduction in metal artifacts in the Huber-penalty reconstruction.

The simple case presented above demonstrates the artifact reduction when the Huber penalty is implemented, in which the top and bottom bleeds in the axial views can be seen. As expected, the Huber penalty has a better resolution-to-noise tradeoff. While the quadratic image

contains less noise due to its smoothing the entire image, the Huber penalty preserved edges of interest such that the bleeds are far more visible despite identical projection acquisition.

7.3 Bleed Visibility around High-Attenuators

Various phantoms were studied, all with the same task in mind – to improve image quality in the area surrounding or near a highly attenuating object. The results are summarized in the following subsections.

7.3.1 Bleed Visibility in Simple Phantom

Using the simple phantom previously described, we were able to achieve a significant improvement in image quality surrounding the platinum sphere. This is a clear distinction between the metal coil and the low contrast bleed in the image created using the designed trajectory, whereas the bleeds lack definition in the circular trajectory. It is impossible to tell where the bleeds start and end with the traditional acquisition technique. Though both images were created with the same parameters in a Huber penalized iterative reconstruction, the more designed orbit visibly outperforms the circular one.

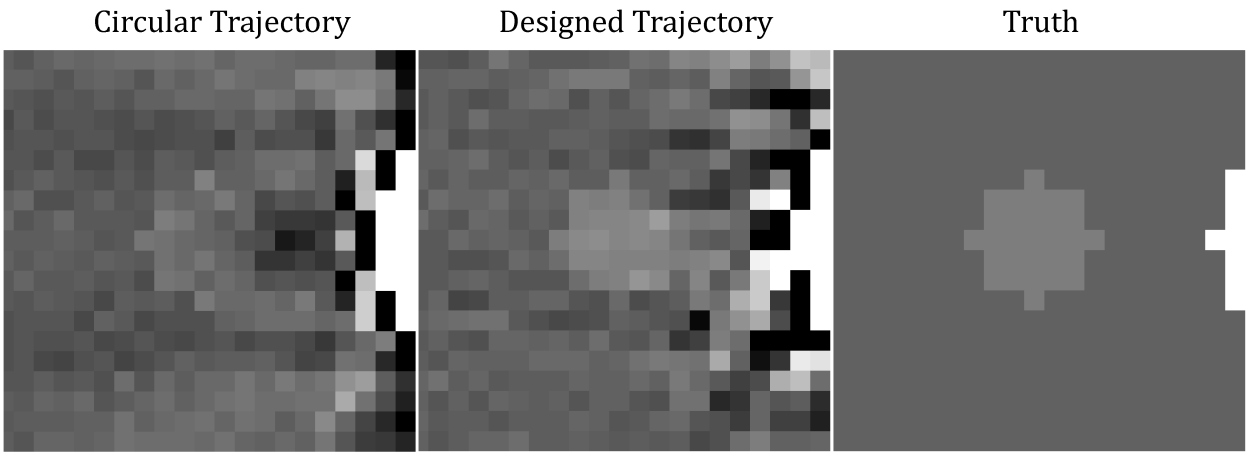

Figure 19: A bleed's visibility in the axial view. The shape and position of the bleed is more distinct in the designed trajectory reconstruction. The circular trajectory image displays less contrast between the bleed and the surrounding brain.

A simple root mean squared error (RMSE) calculation reveals that the error for the designed trajectory when compared to truth was $0.0283\ mm^{-1}$ but $0.0317\ mm^{-1}$ for the circular results. The structural similarity index metric (SSIM) was also calculated, yielding 0.9859 and 0.9839 for designed and circular results respectively.

7.3.2 Bleed Visibility in Head Phantom

Similarly, results were produced for low-contrast bleeds within a head phantom. The results in this section demonstrate a complete workflow (though the designed orbit acquisition was simulated due to limited movement of the available hexapod). The figure below demonstrates a dramatic improvement in the axial view in bleed visibility for all five bleeds within the center slice, which is typically the area most affected by metal implant artifacts.

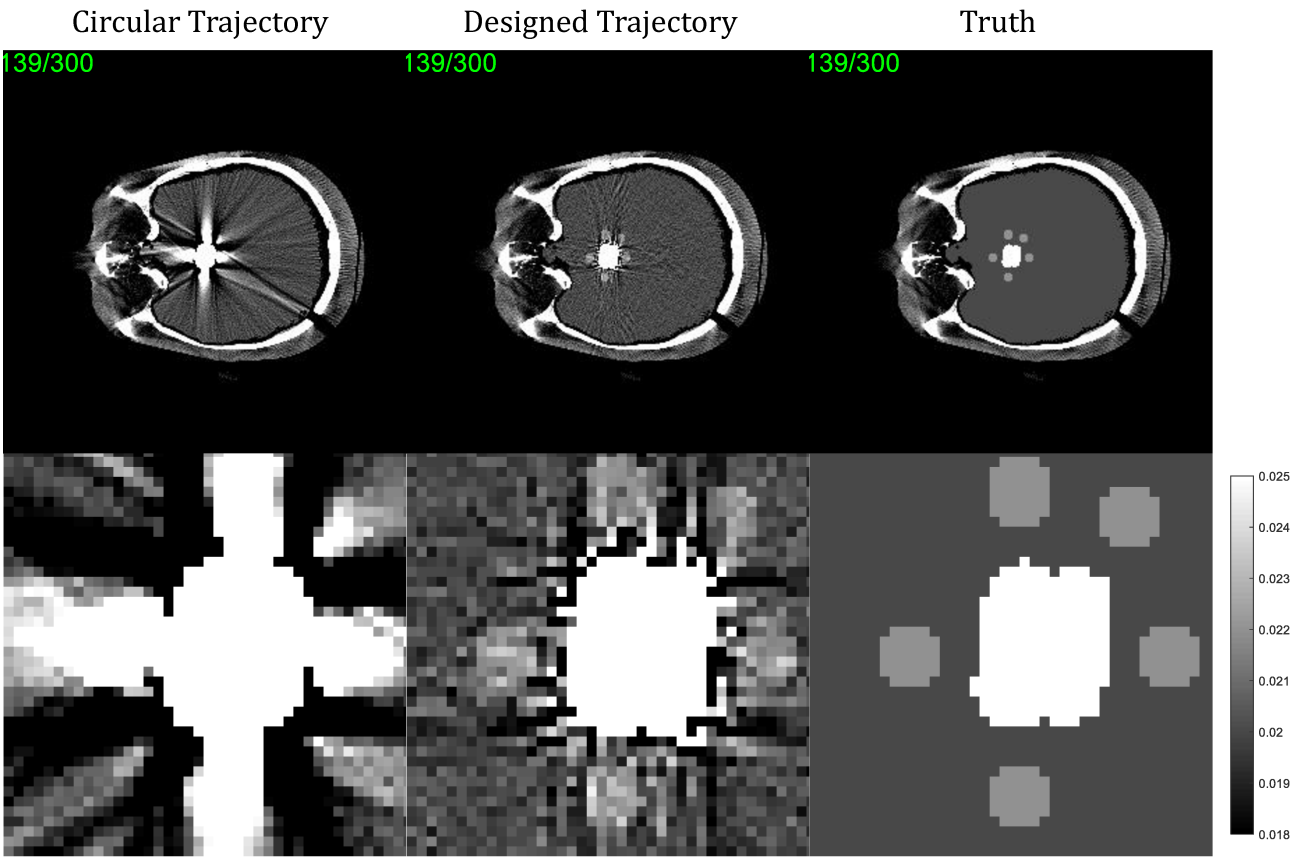

Figure 20: Results of head trials. Axial view displayed for brevity. The circular results illustrate all five bleeds are completely obstructed, while the designed trajectory allows for the position, shape, and size of the bleeds to be clearly seen.

Quantitatively, we found that the RMSE error for the cropped view of the traditional circular orbit was 0.1366 mm^{-1} whereas the error for the designed orbit was 0.0491 mm^{-1} – a decrease of approximately 64%. As further assurance of the improvement, the SSIM was also calculated for this set of results yielding 0.8702 and 0.9735 for circular and designed reconstructions, respectively. The SSIM maps below of the cropped center slice depict the inconsistencies in reconstructing the metal coil itself for both cases. However, the similarity index of the surrounding area containing the bleeds is far better matched to the phantom in the designed orbit than in the circular orbit.

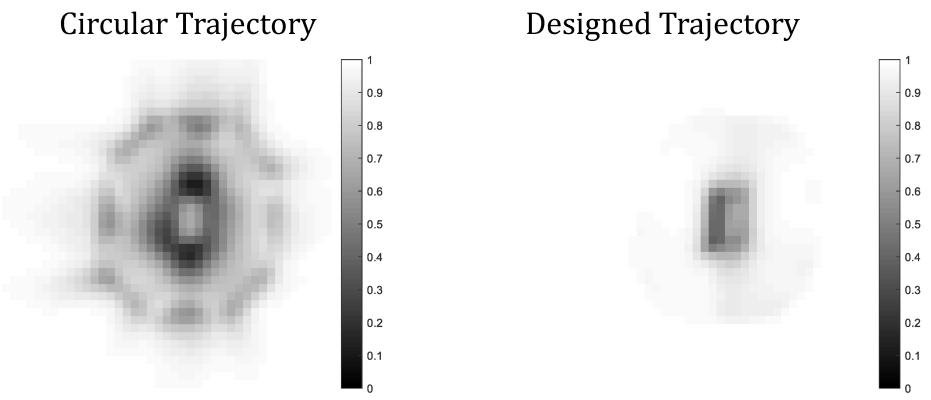

Figure 21: SSIM maps for the circular and designed trajectories when compared to the truth. The similarity between the designed results and truth is evident, especially in the area containing the bleeds.

7.3.3 Extension to Bleed Visibilities in a Prostate Phantom

The usefulness of task-based orbit design is not limited to embolization coils in head-like phantoms. CT imaging to guide and evaluate prostate artery embolization (PAE) is a valuable tool in the treatment of benign prostatic hyperplasia (BPH). Unfortunately, the position of the prostate poses a problem as the surrounding pelvic bone heavily attenuates the photons, leading to poor images of the prostate at lower doses. This work demonstrates the efficacy of task-based orbit design for imaging the prostate and higher frequency information within it at low doses, while maintaining image quality. The generated orbits, optimized by leveraging known pelvis geometry, can be used to improve the quality of prostate images in both average and obese patients, and across extremes of prostate size and position. One such orbit is shown below, with the prostate of interest colored red, followed by the cropped reconstructions of the prostatic bleed.

(a)

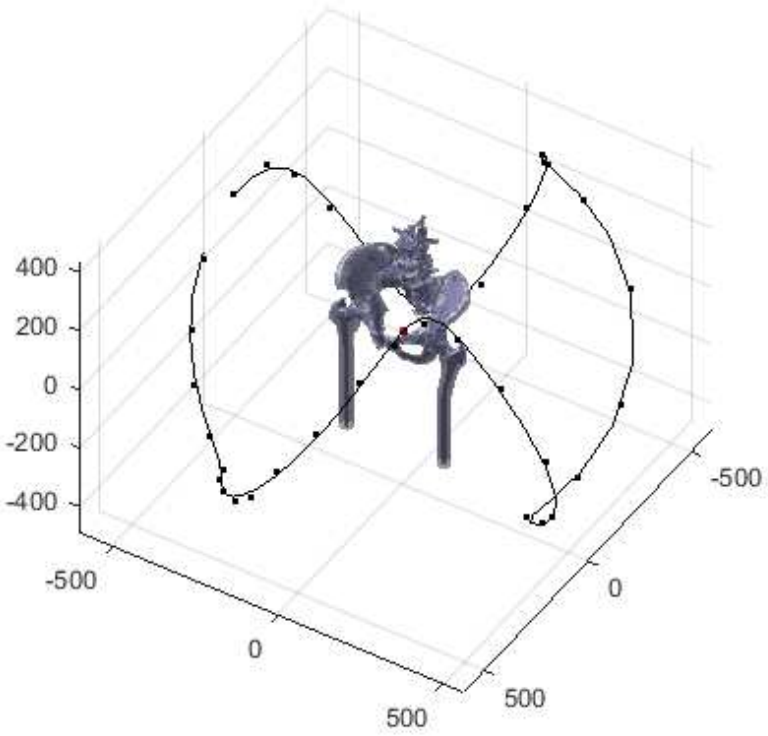

(b)

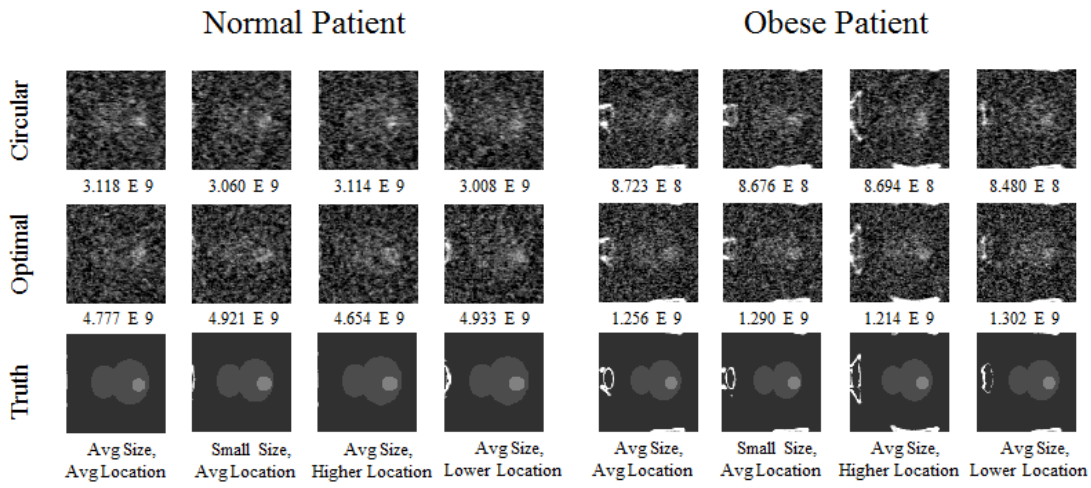

Figure 22: (a) General trajectory found successful for depicted normal patient, with prostate in average position and size. (b) Results for circular and designed (optimal) trajectories in both normal and obese patients. The bleed is visibly clearer and truer to shape in the optimal results in both normal and obese patients, across a variety of prostate sizes and locations.

With photon counts as low as 1000, FBP reconstructions are incapable of producing a useful image in this region. We compare our penalized-likelihood reconstructions using the

customized orbits to standard circular orbits for two small bleeds in an average prostate. In both cases, we find there is a clear increase in the calculated detectability,d', as well as an improvement in the embolization's appearance and its similarity to the phantom. Moreover, as seen below, smaller bleeds are resolved using the designed trajectory shown in the figure below.

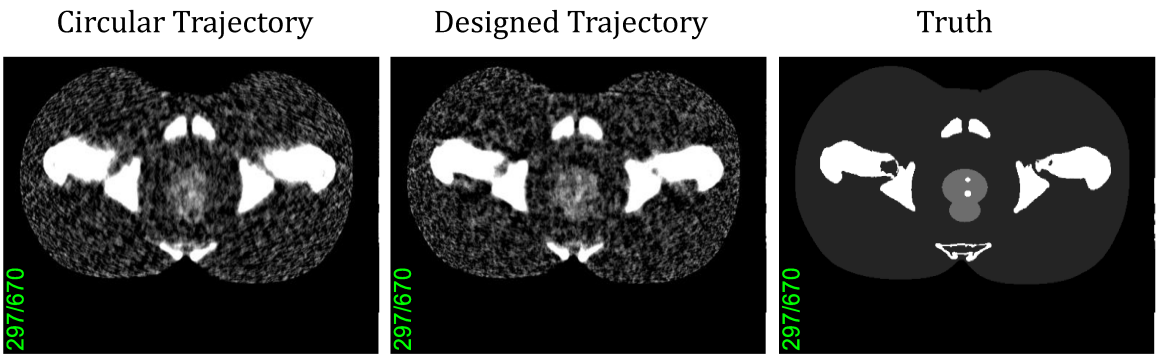

Figure 23: Axial views illustrate the improved bleed visibility for the larger prostatic bleed as well as resolving the smaller bleed.

Chapter 8

Conclusion

This work was based upon the hypothesis that designing an acquisition trajectory passing through the least attenuating areas of the imaged object will yield better image quality. To achieve this, an anatomical model using prior scans as well as segmentation and registration of changes due to surgical intervention were leveraged to optimize the CT orbit such that detectabilities are maximized across the area of interest. We have demonstrated the ability of a properly designed acquisition to dramatically improve the visibility of low contrast objects within simple phantoms as well as within a human body.

These effects were especially pronounced in cases where a highly attenuating object, such as metals within the imaged object, were present. As previously discussed, medical procedures such as embolization in the head can benefit greatly from task-based orbit optimization. These improvements are not limited to the applications explicitly discussed in this thesis. Intraoperative procedures involving metal may all benefit from this technique; these procedures may include pedicle screws for spinal fusion, knee replacement implants, and titanium plate fixation for sternal dehiscence.

The optimization framework described in this work could be performed before surgery based on the prior CT scan and planning information regarding the surgical tools (e.g., embolization). However, real-time optimization may be required under certain circumstances

during surgery. In that case, more efficient optimization and reconstruction routines are necessary. Future work will focus on leveraging the GPU for accelerated CMA-ES optimization and model-based reconstruction.

It is also important to note that we have demonstrated the use of designed orbits in cases without metal artifacts as well. Areas with abundant bone, such as the pelvis, can be difficult to image properly with high image quality and low radiation dose. We have shown that the low-contrast prostatic bleeds were far more visible in the designed trajectory results while using the same initial photon count. Other areas of the body such as the torso and the skull can also benefit from designed orbits, regardless of the presence of metal tools or implants.

Moreover, the validity of each stage of the workflow was investigated and verified. In doing so, the importance and impact of each of the parameters and factors was assessed and detailed, giving insight into the algorithm itself. This work expands upon previous work (Stayman and Siewerdsen, 2013) which focused upon quadratically penalized model-based image reconstruction, finding that Huber penalty is more effective for reducing streak artifacts without blurring out the entire coil area.

Future studies can further explore these medical applications, potentially with cadavers and later with patients. The optimization problem can be further sophisticated by changing the parameterization of the orbit. Rather than employing the gantry (θ) and tilt angles (ϕ) alone, both Euler angles and a translation (θ, ϕ, z) can be employed to create a more dynamic trajectory that may provide additional benefits to image quality by navigating the x-rays through more complex anatomical areas. Additionally, the orbit parameterization can include magnification to increase the field of view if beneficial to the task. As an added benefit, initial photon count can also be included as part of the optimization to reduce radiation to areas where

less photons are required to acquire a proper projection image. Additional optimization parameters can be included in the same task driven optimization framework.

This work has demonstrated an effective method to significantly improve image quality across a variety of tasks and scenarios. Future work can address the limitations discussed within this chapter to further improve the workflow for intraoperative applications.

References

Behling R. *Modern Diagnostic X-Ray Sources: Technology, Manufacturing, Reliability*. 10.1201/b18655 (2015).

Boas FE, Fleischmann D. Computed tomography artifacts: Causes and reduction techniques. Imaging in Medicine. 4(2):229–240. doi: 10.2217/iim.12.13 (2012).

Boas FE, Fleischmann D. Evaluation of two iterative techniques for reducing metal artifacts in computed tomography. Radiology 259(3), 894-902 (2011).

Brenner DJ, Hall, EJ. Computed tomography--an increasing source of radiation exposure. N Engl J Med. 357(22):2277-2284. doi:10.1056/NEJMra072149 (2007).

Brzymialkiewicz CN, Tornai MP, Mckinley RL, Cutler SJ, Bowsher J. Performance of dedicated emission mammotomography for various breast shapes and sizes. Physics in medicine and biology. 51. 5051-64. 10.1088/0031-9155/51/19/021 (2006).

Clackdoyle R and Defrise M. Tomographic reconstruction in the 21st century. Signal Processing Magazine, IEEE 27.4, pp. 60–80 (2010).

Dang H, Stayman JW, Xu J et al. Task-based statistical image reconstruction for high-quality cone-beam CT. Phys Med Biol. 62:8693 (2017).

Dang H, Stayman JW, Xu J, Sisniega A, Zbijewski W, Wang X, Foos DH, Aygun N, Koliatsos VE, Siewerdsen JH. Regularization design for high-quality cone-beam CT of intracranial hemorrhage using statistical reconstruction. SPIE Medical Imaging, San Diego, CA, Vol. 9783, 97832Y (2016).

Fessler JA and Hero AO. Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms. IEEE Tr. Im. Proc., 4(10):1417–29 (1995).

Fessler JA and Rogers WL. Spatial resolution properties of penalized-likelihood image reconstruction: space-invariant tomographs. *IEEE Trans. Image Process.* 5 1346–58 (1996).

Fessler JA. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. IEEE Tr. Im. Proc., 5(3):493–506 (1996).

Friedman SN, Fung GS, Siewerdsen JH, Tsui BM. A simple approach to measure computed tomography (CT) modulation transfer function (MTF) and noise-power spectrum (NPS) using the American College of Radiology (ACR) accreditation phantom. *Medical physics*, *40*(5), 051907 (2013).

Gang G, Siewerdsen JH, Stayman JW. Task-Based Design of Fluence Field Modulation in CT for Model-Based Iterative Reconstruction. 4th International Conference on Image Formation in X-Ray Computed Tomography, pp. 407–410, Bamberg, Germany (2016).

Gang G, Siewerdsen JH, Stayman JW. Task-driven tube current modulation and regularization design in computed tomography with penalized-likelihood reconstruction. SPIE Medical Imaging, pp. 978324, International Society for Optics and Photonics (2016).

Gang G, Stayman JW, Ehtiati T, Siewerdsen JH. Task-driven image acquisition and reconstruction in cone-beam CT. Physics in medicine and biology, 60 (8), pp. 3129–50, ISSN: 1361-6560 (2015).

Gang G, Stayman JW, Ehtiati T, Siewerdsen JH. Task-driven image acquisition and reconstruction in cone-beam CT. Physics in medicine and biology60(8), 3129–50 (2015).

Hansen N and Ostermeier A. Completely derandomized self-adaptation in evolution strategies. Evolutionary Computation 9(2): 159–195 (2001).

Herbst M, Schebesch F, Berger M, Choi J-H, Rahrig R, Hornegger J, Maier AK. Dynamic detector offsets for field of view extension in C-arm computed tomography with application to weight-bearing imaging Med. Phys. 42 2718–29 (2015).

Kalender WA, Hebel R, Ebersberger J. Reduction of CT artifacts caused by metallic implants. Radiology 164(2), 576–577 (1987).

Klein S, Staring M, Murphey K, Viergeer MA, Pluim JPW. Elastix: a toolbox for intensity based medical image registration. IEEE Transactions on Medical Imaging, vol. 29, no. 1, pp. 196 – 205 (2010).

Kudo H et al. New super-short-scan algorithms for fan-beam and cone beam reconstruction. Nuclear Science Symposium Conference Record, IEEE. Vol. 2. IEEE. pp. 902–906 (2002).

Lee EW, So N, Chapman R, McWilliams JP, Loh CT, Busuttil RW, Kee ST. Usefulness of intra-procedural cone-beam computed tomography in modified balloon-occluded retrograde transvenous obliteration of gastric varices. World J Radiol, 8(4): 390-396 (2016).

Long Y, Fessler JA, Balter JM. 3D forward and back-projection for X-ray CT using separable footprints. IEEE Trans. Med. Imag., vol. 29, no. 11, pp. 1839–1850 (2010).

Madhav P, Crotty D, Mckinley RL, Tornai MP. Evaluation of tilted cone-beam CT orbits in the development of a dedicated hybrid mammotomograph. Physics in medicine and biology. 54. 3659-76. 10.1088/0031-9155/54/12/004. (2009).

Mahesh M. The essential physics of medical imaging, third edition. Med Phys, 40, p. 077301 (2013).

Mail N, Moseley DJ, Siewerdsen JH, Jaffray DA. The influence of bowtie filtration on cone-beam CT image quality. *Med Phys*. 36:22–32 (2009).

Menzel HG, Schibilla H, Teunen D. European Guidelines on Quality Criteria for Computed Tomography European Guidelines on Quality Criteria, 1–71. (1999).

Metzler SD, Bowsher J, Jaszczak RJ. Geometrical Similarities of the Orlov and Tuy Sampling Criteria and a Numerical Algorithm for Assessing Sampling Completeness. IEEE Transactions on Nuclear Science - IEEE TRANS NUCL SCI. 50. 1241 - 1245 vol.2. 10.1109/NSSMIC.2002.1239544 (2003).

Nikupaavo U, Kaasalainen T, Reijonen V, Ahonen S-M, Kortesniemi M. Lens dose in routine head CT: comparison of different optimization methods with anthropomorphic phantoms. *AJR. Am. J. Roentgenol.* 204 117–23 (2015).

Ouadah S, Jacobson M, Stayman JW, Weiss C, Siewerdsen JH, Ehtiati T. Task-driven orbit design and implementation on a robotic C-arm system for cone-beam CT, Progress in Biomedical Optics and Imaging - Proceedings of SPIE, vol. 10132, ISSN 16057422 (print), Conference Paper (English), DOI: 10.1117/12.2255646 (2017).

Prince JL and Links JM. Medical Imaging Signals and Systems. Upper Saddle River, NJ: Pearson (2015).

Qi J and Leahy RM. Fast computation of the covariance of MAP reconstructions of PET images ed K M Hanson. International Society for Optics and Photonics, pp 344–55 (1999).

Shamonin DP, Bron E, Lelieveldt BPF, Smits M, Klein S, Staring M. Fast Parallel Image Registration on CPU and GPU for Diagnostic Classification of Alzheimer's Disease. Fronteirs in Neuroinformatics, vol. 7, no. 50, pp. 1-15 (2014).

Shunhavanich P, Hsieh SS, Pelc NJ. Fluid-filled dynamic bowtie filter: a feasibility study. Proc SPIE.9412:94121L. (2015).

Sidky EY, Kao C, Pan X. Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT. Journal of X-ray Science and Technology 14.2, pp. 119–139 (2006).

Stayman JW and Fessler JA. Efficient calculation of resolution and covariance for penalized-likelihood reconstruction in fully 3-D SPECT. *IEEE Trans. Med. Imaging* 23 1543–56 (2004).

Stayman JW and Fessler JA. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Trans. Med. Imaging* 19 601–15 (2000).

Stayman JW and Siewerdsen JH. Task-Based Trajectories in Iteratively Reconstructed Interventional Cone-Beam CT. Proceedings of the International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine, pp. 257–260 (2013).

Szczykutowicz TP and Hermus J. Fluid dynamic bowtie attenuators. Proc SPIE. 9412:94120X (2015).

Tuy HK. An Inversion Formula for Cone-Beam Reconstruction. *SIAM J. Appl. Math.*, vol. 43, no. 3, pp. 546–552 (1983).

Vennart, W. ICRU Report 54: medical imaging – the assessment of image quality. Radiography, 3, pp. 243-244, 10.1016/S1078-8174(97)90038-9 (1996).

Vogel CR and Oman ME. Iterative Methods for total variation denoising. *Sci Comput*. 17(1):227-238 (1996).

Yan XH and Leahy RM. Cone beam tomography with circular, elliptical and spiral orbits. Physics in Medicine and Biology. 37. 493. 10.1088/0031-9155/37/3/001 (2000).

Yu Z, Lauritsch G, Dennerlein F, Mao Y, Hornegger J, Noo F. Extended ellipse-line-ellipse trajectory for long-object cone-beam imaging with a mounted C-arm system. Phys. Med. Biol. 61 1829–51 (2016).

GABRIELA M. RODAL

240.421.1917 grodal1@jhu.edu

EDUCATION

JOHNS HOPKINS UNIVERSITY, Whiting School of Engineering, Baltimore, MD May 2019
Master of Science in Biomedical Engineering, GPA: 3.97/4.0

- Specialization in Imaging (Computed Tomography)
- *Relevant Courses:* Imaging Instrumentation, Computer Integrated Surgery, Computer Vision, Mathematical Image Analysis, Integrated Photonics, Magnetic Resonance, Advanced Topics in Optical Medical Imaging

Bachelor of Science in Biomedical Engineering, GPA: 3.63/4.0, Dean's List – All Semesters December 2016

- Specialization in Instrumentation, Sensors, and Nanotechnology
- Minor in Entrepreneurship and Management with focus in Business Law
- *Relevant Coursework:* Biophotonics Course & Lab, Medical Imaging Systems, Business Law I & II, Business Creation and Contracts, Patent Law, Business Ethics

PROJECTS

JOHNS HOPKINS UNIVERSITY, Biomedical Engineering Department, Design Team Program, Baltimore, MD
Project: "Early Detection of Neonatal Sepsis in Low Resource Settings"

- Designed and constructed a detection device and adapted device to process in situ and transmit data to basic cellular device via an audio jack and alert community healthcare workers of at-risk neonates via SMS
- Illustrated all aspects of design using Creo software
- Co-authored Design History File, project proposal, business plan, and IRB application conforming to all patent and intellectual property laws
- Presented final proposal at the Johns Hopkins Center for Bioengineering Innovation and Design Annual Design Day competition and finalist in Business Plan Competition and Bay Area Social Innovation Challenge

JOHNS HOPKINS UNIVERSITY, Computer Science Department, Baltimore, MD

- Constructed a neural network using TensorFlow in a variety of structures, including Xception, AlexNet, VGG, ResNet50 and Inception, to identify locations of lung nodules in patient CT scans to compare efficacy of the different methods, where all methods achieved accuracies of above 90 percent.
- Pre-processed images to create proper patches for network learning and analysis.

JOHNS HOPKINS UNIVERSITY, Electrical Engineering Department Design Day Program, Baltimore, MD

- Designed and constructed a Near Field Scanning Optical Microscope with near commercial-grade resolution
- Presented at the Johns Hopkins Electrical Engineering Annual Design Day

JOHNS HOPKINS UNIVERSITY, Department of Mechanical Engineering, Baltimore, MD

- Proposed an Optical Coherence Tomography catheter probe to detect misfiring and inactive neurons within the gastrointestinal tract. Authored R21 research grant proposal
- Researched and authored a critical analysis of commercial products of liposomal carriers of doxorubicin. Quantified and modelled distribution system to test validity of the analysis and conclusions

EXPERIENCE

JOHNS HOPKINS UNIVERSITY, Baltimore, MD August 2017 – February 2019
Graduate Student Investigator, Advanced Imaging Algorithms & Instrumentation Laboratory
Master's Thesis: "Task-Driven Trajectory Design for Endovascular Embolization"

- Leverage detectability index of local embolization coil to drive customized acquisitions and image formation for statistical reconstruction methods of head and pelvis data to improve image quality in operative area
- Employ advanced image processing techniques to segment and manipulate volumetric data

JOHNS HOPKINS UNIVERSITY, Baltimore, MD May 2018 - Present
Professional Development Chair, Graduate Representative Organization

- Collaborate with senior university administration to organize and facilitate professional development events and activities for the Homewood graduate community
- Improve student employment outcomes and engage alumni in networking events across departments

JOHNS HOPKINS UNIVERSITY, Baltimore, MD August 2017 – January 2018
Head Teaching Assistant, Systems Bioengineering 3 Course, Department of Biomedical Engineering

- Lead 150+ engineering seniors in 2 weekly discussion sections to improve student's comprehension of novel course content and the mathematical background necessary to succeed in the course
- Held office hours each week for additional problem set and coding help; guest lectured when necessary
- Revised exam questions and weekly assigned problem sets. Wrote detailed rubrics for all assignments and exams. Managed 6 graduate student graders to ensure timely and accurate feedback on student work

JOHNS HOPKINS UNIVERSITY, Baltimore, MD January 2016 – May 2017
Chemistry Lab Teaching Assistant, Department of Chemistry

- Lead 20 students in weekly experiments, graded exams, performed administrative tasks
- Communicated and clarified key concepts and proper lab techniques to students

JOHNS HOPKINS UNIVERSITY, Baltimore, MD June 2014 – November 2015
Undergraduate Researcher, Enteric Neuromuscular Disorders and Pain Laboratory

- Constructed an experimental pump to test and analyze the in-vitro peristaltic activity in the small and large intestines of mice to develop a consistent model of gut movement
- Investigated effects of drug and optogenetically modified conditions
- Presented findings to Principal Investigator and 10 senior faculty every month
- Executed experiments requiring reagent preparation and delicate surgeries and dissections

SKILLS

Design: Math Modeling, Optical System Design, Circuit and Device Design
Computer Applications: MATLAB, Python, Creo Parametric, AutoCAD, Java, ImageJ
Business and Logistics: Business Plans, Patent Research, IRB Applications, FDA Regulations, R21 Grant Proposals
Languages: Spanish, French