Applications of Statistical Modeling in Iterative CT Image Reconstruction

David Simon Perlmutter

A thesis
submitted in partial fulfillment of the
requirements for the degree of

Masters of Science

University of Washington

2015

Reading Committee:

Adam Alessio, Chair

Paul Kinahan

Les Atlas

Program Authorized to Offer Degree:
Electrical Engineering

University of Washington

Abstract

Applications of Statistical Modeling in Iterative CT Image Reconstruction

David Simon Perlmutter

Chair of the Supervisory Committee:
Associate Professor Adam Alessio
Radiology

Traditionally, x-ray CT images are produced by an algorithm called filtered back projection, or FBP. FBP is an analytical solution to the idealized CT image reconstruction problem, the inverse problem of turning raw x-ray measurements into a full 3-dimensional (3D) image, and is derived assuming a continuous set of noiseless measurements. However real CT data are noisy and biased, especially so if the scans are performed at low x-ray dose, and advanced statistical estimation techniques have been shown to produce higher quality images than FBP. This work presents two applications of statistical modeling in CT image reconstruction. The first application discusses the statistics of CT data noise, and compares the performance of several common models for estimation in a simplified 1D experiment. The second application concerns modeling temporal CT data, in which the measured data typically contain redundancies. It proposes an estimation method that exploits these redundancies to address two key challenges in CT image reconstruction: reducing noise and lowering computation time. We demonstrate this noise reduction analytically and through experimental simulations. In addition, a third study validates the use of the statistical models used in this work by comparing them to measured data from a clinical CT scanner. Overall, these methods contribute to the methodology of statistical CT image reconstruction to enable ultra-low dose x-ray CT imaging.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I'd like to chiefly thank Profs. Adam Alessio and Paul Kinahan for giving me an opportunity in their lab and for their help, ideas and support throughout my Masters work. Also much thanks to Soo Mee Kim for sharing ideas and code, the entire IRL lab for being great to work with the last two years, as well as to Dr Bruno De Man and his colleagues at GE Global Research and Patrick La Riviere and Dimpil Modgil at U Chicago for helpful conversations and Prof. Les Atlas for serving on my thesis committee. And finally, thanks to my family for their support throughout grad school. The work in this thesis was supported by the National Institutes of Health under grants R01-HL109327 and R01- CA115870.

Chapter 1

INTRODUCTION

This thesis investigates new statistical modeling techniques for x-ray computed tomography (CT) image reconstruction, the process of turning raw x-ray measurements into a full 3D image. The goal of this work, as in most advanced CT modeling research, is to produce the highest quality images possible from a given set of data, or equivalently, to enable a sufficient level of image quality at the lowest possible radiation dose. This intrinsic trade off between image quality (or more specifically, image noise) and radiation dose derives from the fundamental statistics of x-ray detection (see 1.1). Understanding and exploiting these statistics are crucial to advancing the quality of CT imaging. The drawback to statistical methods, however, is that they typically produce algorithms with significantly longer computation time. Statistical image reconstruction has been used in other tomographic modalities, such as Positron Emission Tomography (PET), since the late 90s. However CT has been slow to adopt statistical methods because the size of CT data and images, typically 1-2 orders of magnitude larger than PET, lengthens computation time beyond practical feasibility in the clinic. Efforts to increase the speed of reconstruction algorithms are crucial to advancing statistical modeling research so that they can overcome this clinical hurdle.

The thesis is organized into five sections. The remainder of Chapter 1 introduces the basics of CT imaging, with particular emphasis on statistical reconstruction methods. It also provides a short mathematical background on maximum likelihood estimation (MLE), which is the theoretical framework for methods used in this work. Chapter 2 compares simulated CT data, which uses a Compound Poisson plus Gaussian (CPG) signal model, to experimental data from a clinical CT scanner. It examines the validity of the CPG model and estimates model parameters to improve the fidelity of the simulation. Chapter 3 presents a study

comparing the performance of various advanced statistical models using a simplified 1D experiment. This experiment allows one to analyze the relative benefits and trade offs between models of different fidelity. Chapter 4 presents a new algorithm for dynamic CT imaging in which assumptions about redundancy in the CT data enable both improved image quality and reduced computation time over standard methods. The final chapter concludes the work and provides areas of future work.

1.1 X-ray CT Tomography

X-ray Computed Tomography (CT) is a medical imaging modality used to produce 3D, voxelized, anatomical maps of the human body. A CT image shows the ability of a particular region of the body to attenuate x-ray photons, which corresponds to its material composition and density. Modern scanners can acquire full body images in a few minutes, with a resolution under 1mm^3, and are used in a wide variety of clinical procedures such as cancer screening, diagnosis of cardiac and lung disease and surgical planning. A simple way to understand the underlying technology behind CT is through the etymology of the term "computed tomography" itself. Tomography, which derives from the Greek "tomo" meaning slice, and "graphy" meaning sketch, is the process of forming an image of an object by collecting projections, or slices, through that object. In x-ray CT, the projections are simple radiographs, conventionally known as 2D x-ray images. However tomographic modalities such as PET and Optical Coherence Tomography (OCT) acquire projections through other sources of penetrating radiation. In a CT scan, an x-ray source and detector rotate around a subject, collecting multiple projections from different angles, or views, of the body (see fig. 1.1). Stacked together, these projections are called a sinogram. Because the measured data are just a set of projections, the final CT image must be "computed", or reconstructed, via an algorithm. The choice of reconstruction algorithm has a crucial impact on the presentation and quality of the resulting image. The classic reconstruction algorithm for CT is filtered backprojection (FBP), which is used in virtually all commercial CT scanners because of its simplicity and computational speed (for a good review of FBP, see [16]). However more

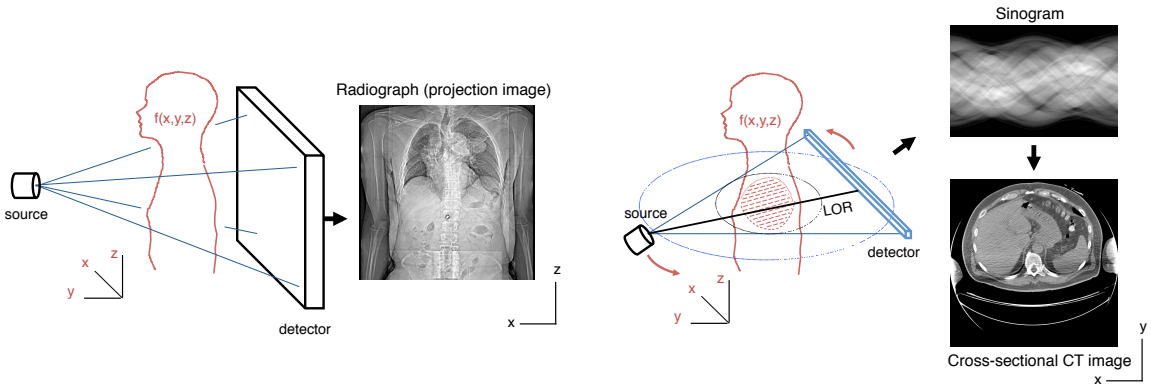

Figure 1.1: Left: Schematic of conventional 2D radiography ("x-ray imaging") with a static x-ray source and flat panel detector. It produces a projection image, which is a summing or overlay of 3D information onto a 2D plane. Right: A CT image acquisition with source and detector mounted on a rotating gantry. Projections (1D in this example) from multiple angles are measured to form a sinogram, from which a true cross-sectional image is produced.

advanced algorithms, particularly statistically based methods, have been shown to produce higher quality images than FBP and can flexibly accommodate different scanner geometries. These methods achieve higher performance by incorporating knowledge of the statistics of the data into their algorithm. Before addressing reconstruction methods we will briefly discuss the primary sources of randomness in CT measurements.

The detection of monochromatic x-ray photons is primarily governed by Beer's law of attenuation which states that the attenuation of light through a scattering material is exponentially proportional to the material's thickness.

$$y = I_0 e^{-\mu p} \tag{1.1}$$

where I_0 is the initial flux, y is the attenuated x-ray flux, p is the length or density of material, and μ is an attenuation coefficient. In the context of CT, p is the sum of material density across multiple voxels along a particular ray, or line of response (LOR), through the body

(fig. 1.1),

$$p_j = \sum_i a_{ij} x_i \tag{1.2}$$

where the index j defines an LOR, x_i represent an image voxel, and a_{ij} are elements of the projection matrix A (see below), representing the relative contribution of voxel i to LOR j (for simplicity one can think of a_{ij} as an indicator variable). The task of reconstruction is to produce an estimate of each x in the image given a set of measurements y. Unfortunately, Beer's law only predicts the average detector response, and real photon measurements are intrinsically noisy. The primary source of randomness is due to the quantum nature of light, which causes ambiguity in the arrival of monochromatic photons that has been empirically shown to follow a Poisson distribution. Poisson random variables have a signal to noise ratio (SNR) linearly proportional to their mean, so high-count x-ray measurements contain less noise than low-count ones. Thus a high intensity the x-ray source, with higher the patient dose, leads to higher the image quality. Most statistical methods model x-ray data as either a Poisson or Gaussian distribution, which is known to be a good approximation of Poisson random variables with large means. However these models are only approximations because of a variety of complicating factors (see [21] for a good review on the topic of CT data modeling). First, most CT sources produce a broad spectrum of x-ray photons, and the attenuation coefficient μ is energy dependent and varies with different materials in the body. This follows a more complicated, multi-energetic form of Beer's Law,

$$y = \int_{\mathcal{E}} I_0(\mathcal{E}) e^{-\sum_k \mu_k(\mathcal{E})\, p_k} \tag{1.3}$$

where $\mathcal{E}$ is the x-ray energy and k indexes different material types along the LOR. Detection efficiency is also energy dependent and can vary from detector to detector. Also, most commercial CT scanners use scintillating detectors, which first convert an x-ray photon into optical light before detection. The detected optical signal depends on the energy and trajectory of the incoming x-ray, its depth of interaction in the scintillating material, and the spectrum of optical photons produced [21]. These scintillator interactions are quite complicated and are typically poorly modeled, although they likely cause significant randomness

in the data. Finally, the detector electronics themselves produce some thermal noise even in the absence of x-rays. While these are the primary sources of statistical noise in the data, a wide variety of undesirable deterministic effects such as x-ray source spot size (which affects resolution), x-ray scatter, and patient motion also degrade CT data. One challenge of statistical reconstruction lies in the ability to choose which effects to model so as to balance image quality with computational speed.

1.2 Statistical Reconstruction

Statistical reconstruction is a subset of (though often conflated with) iterative reconstruction, a broad research area of tomographic reconstruction techniques based in the fields of estimation and inverse problems. The common theme of iterative methods is that they produce a series of image estimates via a feedback mechanism in which the current image is used to simulate a hypothetical set of measurements. This operation is known as image projection. This synthetic data are then compared to the true measured data, and the discrepancy used to produce the next image estimate through a "backprojection". Ideally, these estimates will converge to a stable final image. Analytic reconstruction methods such as FBP, in contrast, form a single image directly from the data. The algebraic reconstruction technique (ART) and its generalizations [3, 13, 15], some of the first iterative reconstruction algorithms proposed for CT, are based on a general method for solving linear systems of equations of the form $Ax = b$. While ART allows for more flexibility in modeling the geometry of the data-generating process (the projection operator, "A"), it is not fully statistical because it does not incorporate knowledge of data noise. Fully statistical methods [14, 26] treat the measured data as a random variable drawn from a probability distribution and attempt to reconstruct an image which is "most likely" to have produced it. The most well-known of these is the maximum likelihood expectation maximization algorithm (MLEM) for emission tomography [28]. MLEM, which is the basis for most clinical PET and SPECT reconstruction software today, treats the data as Poisson distributed. While MLEM was

adapted for transmission tomography [20] it converges slowly for transmission problems, and several alternative Poisson likelihood algorithms have been proposed to speed up reconstruction [4, 9, 19]. Ordered subsets (OS) is a technique used in many algorithms [17] to speed up image projection, which is typically the computational bottleneck in iterative algorithms. Instead of projecting a complete set of data in each iteration, OS algorithms break data into N equally-sized subsets and perform a projection and image update sequentially for each one. This increases the speed of each iteration by roughly an order of N, but can lead to unstable convergence. Finally, many algorithms model CT data as Gaussian instead of Poisson distributed. While potentially less accurate leads to a simpler least squares estimation problem [27] and thus faster convergence.

Beyond algorithmic speed, a variety of other modeling factors have been studied in the last few decades that address the quality of CT reconstruction. Much of this work falls under the category of geometric modeling, which focuses on designing projection operators that balance accuracy and speed. Other techniques seek to avoid the cost of projection altogether and perform statistical restoration of the data itself [18], followed by simple FBP. Several groups have used more sophisticated statistical models than Poisson and Gaussian in their reconstruction algorithms that incorporate, among other things, polyenergetic x-ray spectra [7, 8], noise in the detector electronics [35, 37] and scattering. Finally, Bayesian methods [12], which rely on a priori image knowledge, have had a significant impact on improving image quality. The power of image priors lies in their ability to incorporate empirical side information about the image itself, not just the data acquisition process, directly into reconstruction. This is done via a penalty, or regularization, function on the image. For example, the classic quadratic Markov random field prior (QMRF) [4] favors smoothness in the reconstructed image by penalizing the difference between neighboring pixels, squared. This is especially crucial for low dose data, in which reconstructions are dominated by high frequency noise. A major benefit of the Bayesian framework is it allows the user to design penalties flexibly to promote all kinds of image behavior. An enormous number of image priors have been proposed to encourage image properties such as smoothness, sparse edges [29], uniform spatial

resolution [30], and consistency with previously acquired scans of the same patient [6], and many more.

1.3 Maximum Likelihood Estimation

This section provides a brief background on maximum likelihood estimation (MLE), the mathematical framework upon which most statistical reconstruction is based. Maximum likelihood codifies the intuition that, when estimating the parameters of a statistical model, one should choose those that make the probability of observing the measured data "most likely" according to the model. Formally, we define a statistical model as a parameterized distribution $p(y|x)$ over a set of observed data, y, with a single parameter x (the extension to multiple parameters is straightforward). To choose a CT related example, y could be the outcome of a series of N x-ray detector measurements, and x the true x-ray source intensity. Then, the maximum likelihood estimate is the parameter value that maximizes $p(y|x)$ over the observed data *as a function of* x. When viewed in this way, $p(y|x)$ is know as a likelihood function.

$$\hat{x}_{MLE} \;=\; \arg\max_{x} \; p(y|x) \tag{1.4}$$

The beauty (and perhaps drawback) of MLE is that inference on x relies solely on a model of the data, $p(y|x)$, which is often well understood; it does not model the parameters themselves. In our example, if we assume that each x-ray measurement y_j is Poisson distributed with a mean equal to the source intensity, x, then,

$$\hat{x}_{MLE} \;=\; \arg\max_{x} \; p(y|x) \;=\; \prod_{j=1}^{N} p(y_j|x) \;=\; \prod_{j=1}^{N} \frac{x^{y_j}\, e^{-x}}{y_j!}. \tag{1.5}$$

Likelihood functions are often messy expressions that are difficult to optimize. It is common to take their log before optimizing, which since log is monotonic does not affect the position

of its maximum. Thus, denoting the log-likelihood function as $L(x)$,

$$\hat{x}_{MLE} \;=\; \arg\max_x \; L(x) \;=\; \sum_{j=1}^{N} y_j \log(x) - x - \log(y_j!) \tag{1.6}$$

$$\frac{\partial L}{\partial x} \;=\; \sum_{j=1}^{N} \frac{y_j}{x} - 1 \;=\; 0, \qquad\qquad \hat{x}_{MLE} \;=\; \frac{1}{N}\sum_{j=1}^{N} y_j = \bar{y} \tag{1.7}$$

the MLE of x is the measurement mean.

The downside of not modeling parameters probabilistically is that one can become over reliant on the data and ignore potentially useful side information about the parameters directly. This is especially true when the data are very noisy and should not be fully trusted. Bayesian statistics offers an elegant modification to MLE called maximum a posteriori (MAP) estimation. Intuitively, MAP estimation seeks the *most likely parameters* given the data, not those that make the *data* most likely as in MLE. Mathematically, MAP seeks to maximize the posterior function, $p(x|y)$, instead of the likelihood function, $p(y|x)$. The mathematical connection between the likelihood and posteriori functions is known as Bayes' Rule, which states that the posterior is equal to the likelihood times a prior term, which is a probability function on x itself,

$$p(x|y) \;=\; p(y|x)\, p(x). \tag{1.8}$$

This assumption that the parameters themselves are probabilistic is intrinsic to Bayesian thinking. For example we may believe, before seeing the measured data, that our true x-ray source intensity is equally likely to be between 0 and 100 photons, or Gaussian distributed around 60 photons, or always odd numbered. The power of MAP estimation lies in the ability of the user to design informative, flexible and computationally practical priors like these. Taking the log of the posterior function for simplicity of optimization gives,

$$\hat{x}_{MAP} \;=\; \arg\max_x \; L(x) \;+\; R(x) \tag{1.9}$$

where $R(x)$ is the log of the prior term, known as a regularization function. In this form we can see that the MAP estimate is a balance between two factors: the likelihood function,

which measures how well the parameters fit the data, and the regularization function, which measures how closely the parameters conform to our initial expectations.

Chapter 2

MODEL ANALYSIS FOR ULTRA LOW DOSE EXPERIMENTAL CT DETECTOR DATA

In this work we use experimental data from a clinical GE VCT scanner to validate and fit a standard CT detector noise model, the Compound Poisson plus Gaussian (CPG) model. We are particularly interested in conditions of low x-ray flux, and will show the model is most well suited for that regime. First, we describe the CPG model and our testing procedure. We then present experimental results and evaluate model performance. Finally, we compare our experimental data to simulations from catsim, a GE-developed CT simulation software tool, using experimentally-determined parameters settings.

2.1 Model

The Compound Poisson plus Gaussian (CPG) model describes a CT detector signal as a sum of two terms: a Compound Poisson random variable that describes the arrival of x-ray photons of different energies, and a Gaussian electronic noise term. Specifically the detector signal, Y, is given as,

$$Y = G\sum_{i=1}^{\mathcal{E}_m} \mathcal{E}_i X_i \; + \; N(i_d, \sigma_e^2) \qquad\qquad X_i \sim Poisson(\lambda_i) \tag{2.1}$$

where X_i is the Poisson distributed number of x-ray photon arrivals of energy $\mathcal{E}_i$ with mean λ_i, i_d and σ_e^2 are the mean and variance of the electronic noise, respectively, and G is a constant gain term that converts photon energy into detector signal units. Note that since G is constant, the detected signal from a photon is linearly proportional to its energy. A goal of this study is to find experimental values for G, i_d, and σ_e.

We can further characterize the model by calculating its basic statistics. Consider a com-

pound Poisson random variable as a weighted sum of a finite number of independent Poisson random variables, X_i, with means λ_i, and weights $\mathcal{E}_i$. That is,

$$S = \sum_{i=1}^{m} \mathcal{E}_i X_i \qquad X_i \sim Poisson(\lambda_i) \tag{2.2}$$

and, define $\lambda_0 = \sum_{i=1}^{m} \lambda_i$. Also, define the function, $\rho(\mathcal{E})$, such that $\rho(\mathcal{E}_i) = \lambda_i$. Physically, S is the detector signal (minus electronic noise) and $\rho(\mathcal{E})$ is the incident photon spectrum, made of energies, $\mathcal{E}_i$, and intensities λ_i. X_i is then the number of detected photons of energy $\mathcal{E}_i$, with mean λ_i. Then the mean of S is,

$$\mu_S = E[S] = E\left[\sum_{i=1}^{m} \mathcal{E}_i X_i\right] \tag{2.3}$$

$$= \sum_{i=1}^{m} \mathcal{E}_i E[X_i] = \sum_{i=1}^{m} \mathcal{E}_i \lambda_i = \lambda_0 \sum_{i=1}^{m} \frac{\lambda_i}{\lambda_0} \mathcal{E}_i \tag{2.4}$$

$$= \lambda_0 \mu_{\mathcal{E}} \tag{2.5}$$

where μ_E is the mean energy of the *normalized* spectrum, $\frac{\rho(\mathcal{E})}{\lambda_0}$. The variance of S is,

$$\sigma_S^2 = Var[S] = Var\left[\sum_{i=1}^{m} \mathcal{E}_i X_i\right] \tag{2.6}$$

$$= \sum_{i=1}^{m} \mathcal{E}_i^2 Var[X_i] = \sum_{i=1}^{m} \mathcal{E}_i^2 \lambda_i \tag{2.7}$$

$$= \lambda_0 \sum_{i=1}^{m} \frac{\lambda_i}{\lambda_0} \mathcal{E}_i^2 \tag{2.8}$$

$$= \lambda_0 \rho_{\mathcal{E}} \tag{2.9}$$

where $\rho_{\mathcal{E}}$ is the mean of the *squared* spectrum, or the second moment of $\frac{\rho(\mathcal{E})}{\lambda_0}$. Hence, the mean and variance of the CPG variable Y are given as,

$$\mu_Y = G\lambda_0 \mu_{\mathcal{E}} + i_d \tag{2.10}$$

$$\sigma_Y^2 = G^2 \lambda_0 \rho_{\mathcal{E}} + \sigma_e^2 \,. \tag{2.11}$$

Note that both mean and variance are linearly proportional to λ_0, up to additive constants. Thus, when comparing a change in σ_Y^2 vs μ_Y, *as a function of* λ_0,

$$\frac{\Delta\sigma_Y^2}{\Delta\mu_Y}(\lambda_0) = \frac{G^2\lambda_0\rho_{\mathcal{E}}}{G\lambda_0\mu_{\mathcal{E}}} = G\frac{\rho_{\mathcal{E}}}{\mu_{\mathcal{E}}} \approx G\mu_{\mathcal{E}} \tag{2.12}$$

By knowing properties of the incident x-ray spectrum (namely its mean and second moment) we can back out the detector gain factor, G, by plotting signal variance vs mean and examining the slope, without knowing the true experimental flux. The last approximation in Eq 2.12 is only meant to build intuition that slope increases with higher spectrum mean energy – all calculations use the true expression.

2.2 Experiments

2.2.1 Dark current data

Our first experiments characterized the electronic noise by collecting data with the x-ray source off. Table 2.1 lists our standard experimental parameters for data collection, using the VCT Service Mode tool (all data presented used these settings unless otherwise noted). It should be noted that the units of raw signal reported by the VCT by the Service Tool Mode are unknown and treated here as arbitrary. We assume the goal of catsim is to produce data in these units. The only data available to us in Service Mode were the mean and standard deviation of each detector over the full scan; we do not have access to the full data set. Table 2.1 also shows the experimental electronic noise has mean 897.3 and standard deviation 14.99. We use these experimental values for i_d and σ_e, respectively, in later catsim simulations, but should mention two caveats. First, not every detector shares the same electronic noise statistics. There is a consistent, detector-to-detector variation in dark mean, with a standard deviation of 28.1 (σ_e likewise varies across detectors). We have not explored the structure of this variation. Second, since we are currently only able to access the mean and standard deviation of the detector data, we have no knowledge of higher order statistics of the electronic noise – it is only *assumed* to be Gaussian.

X-ray Rotation	Off
Scan Time	3 sec
Inter Scan Time	1 sec
Trig Rate	984 Hz
Detector Size	0.625 mm
Filter	Large
Calibration Vector	Full
Mean	14.88
Standard Deviation	14.99

Table 2.1: Standard VCT scan parameters.

2.2.2 X-ray on data

We now turn to estimating G experimentally via variance vs mean plots using Eq. 2.12. First, we check if the experimental variance is linear with mean as predicted by the CPG model. We hold the tube voltage at 120 kVp and change tube current to vary detector flux. In addition the bowtie filter, a standard CT dose-reduction filter that preferentially attenuates x-rays along LORs at the edge of the body, provides a large spread of fluxes within each scan. To explore the low flux regime, we collected scans with about 2 mm lead between the tube and detector at 10-650 mA tube current, as well as air scans with 10-300 mA current. Figure 2.1 shows 3 views of this variance vs mean data for a single row of detectors, with lead scans in blue and air scans in red. While the low flux regime (far right) is linear, detectors with signal beyond about 50,000 have variance that grows super-linearly. This heightened variance is not predicted by the CPG model, and requires further investigation. We will restrict our analysis to data in the low flux regime.

After determining an appropriate flux level, we collected a larger amount of variance-mean data to directly calculate G. For a given tube voltage, we collected scans at multiple current values, as before. We used 4.5 cm of aluminum as attenuating material to lower flux levels

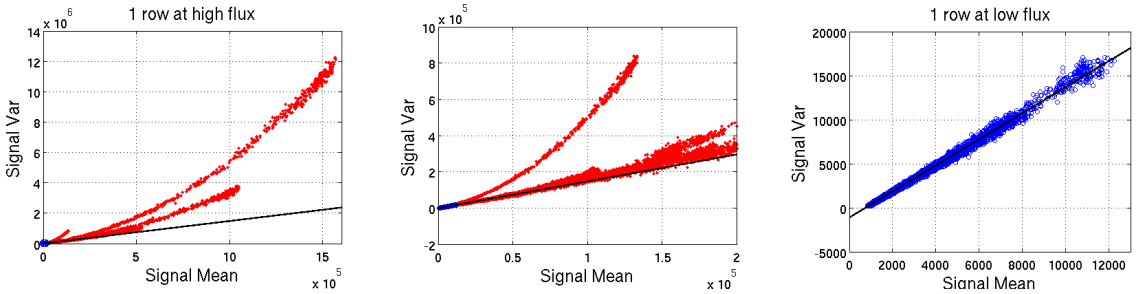

Figure 2.1: Signal variance vs mean for a single row of detectors under a wide range of photon fluxes at 120 kVp. The left figure shows the entire data set, while the middle and right figures are progressively more zoomed in. Scans were taken with both a lead target (blue data) and no target (red).

to an appropriate level. We then plotted the average variance vs mean of the central 16 detectors in each row (over 6 rows), and found the slope of the best fit line. We repeated this experiment at 80, 100, 120, and 140 kVp, which each have different expected slopes. Figure 2.2 shows these four data sets. Notice the increasing slope with higher kVp, which agrees with the intuition that slope is proportional to mean energy. Lastly, we removed about half the aluminum from the target and recollected 80 kVp data to verify that we could measure a change in slope due to the beam hardening from the aluminum. This data is also shown in figure 2.2. Using spectrum data in catsim, we calculated the incident spectrum mean and second moment to convert each slope into a G value. Since the CPG model gain is constant across energy, each experiment should produce the same G value. Table 2.2 shows the slopes and G values for each data set, along with their standard deviation. The mean value is $G = 0.0195$. The higher G values at 120 and 140 kVp may be due to slightly super-linear variance at signal above 30,000. Regardless, the fit is quite good.

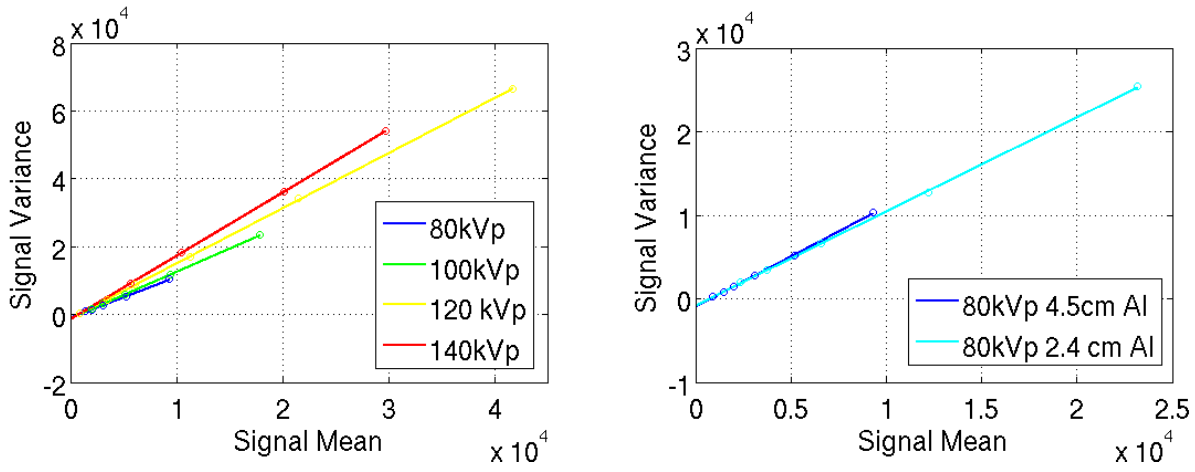

Figure 2.2: (Left) Low flux detector variance vs mean for 4 different kVp settings and 4.5 cm Al target, with best fit slopes in legend. (Right) Same analysis for 80 kVp with 4.5 cm Al target (blue) and 2.4 cm Al target (light blue).

2.3 Simulations

With new values for i_d, σ_e^2, and G, we used these parameters to create catsim simulations of our experiment. We used a 4.5 cm (and 2.4 cm) thick aluminum analytic box phantom, as well as the appropriate kVp, mA, and bowtie filter. The simulated sinograms were then processed identically to the experimental data, first by finding detector means and variances, then plotting the results and calculating their slope. A comparison of simulation and experimental results is shown in figure 2.3. Looking at the upper left image, there appears to be good agreement between simulated and experimental data; both data sets follow almost exactly the same linear trend, except for a slightly higher experimental variance at high signal. To prove the importance of an accurate G, the 140 kVp simulations were redone with $G = 0.039$, twice its experimental value. This data, plotted in black in the upper left figure, does not match any experimental data. While the mean-variance trends match, the catsim data has less signal than experimental for the same mA (by about a factor of 4), which implies an underestimate of total x-ray photons detected. This is seen by looking at the ratio of

	80 kVp, 2.4 cm	80 kVp, 4.5 cm	100 kVp	120 kVp	140 kVp
Slope	1.13	1.20	1.37	1.62	1.86
$\frac{\rho_\mathcal{E}}{\mu_\mathcal{E}}$	59.0	62.3	72.0	82.1	91.9
G	0.0193	0.0191	0.0190	0.0198	0.0203
σ_G	2.4e-12	3.8e-12	1.2e-12	1.9e-12	0.9e-12

Table 2.2: Experimentally calculate G values. Note, $\frac{\rho_\mathcal{E}}{\mu_\mathcal{E}}$ is calculated after beam hardening.

simulated to experimental signal, plotted in the lower left of figure 2.3. In catsim, spectra are stored in units of photon flux per mA, and are linearly scaled by mA, acquisition time, and detector area to produce total photons. The simulated acquisition time matches experiment and should not be an issue. We used catsim's default detector size and fill factor which to the best of our knowledge accurately represent the 0.625 mm detection bin at iso-center used in experiment. However an error here could easily account for our discrepancy. Finally, while simulation mA matched experiment, the experimental signal mean is slightly sub-linear with mA at high current (above about 300 mA, see figure 2.4). This probably means the tube flux itself droops at high mA. Note that this is a relationship between input mA and flux, a separate issue from the increased variance at high detected signal. Ignoring this issue for now, we conclude that the catsim spectra themselves underestimate VCT tube flux. Figure 2.3 includes new catsim data with each catsim spectra rescaled to better fit experiment. Observe that the 80 kVp simulations still appear to be mismatched at higher signal. This is because the 80 kVp scans required higher mA (up to 320 mA), and so catsim overestimates flux. The 140 kVp data only used up to 60 mA, and is not affected by this problem. The exact rescaling factors are listed in table 2.3.

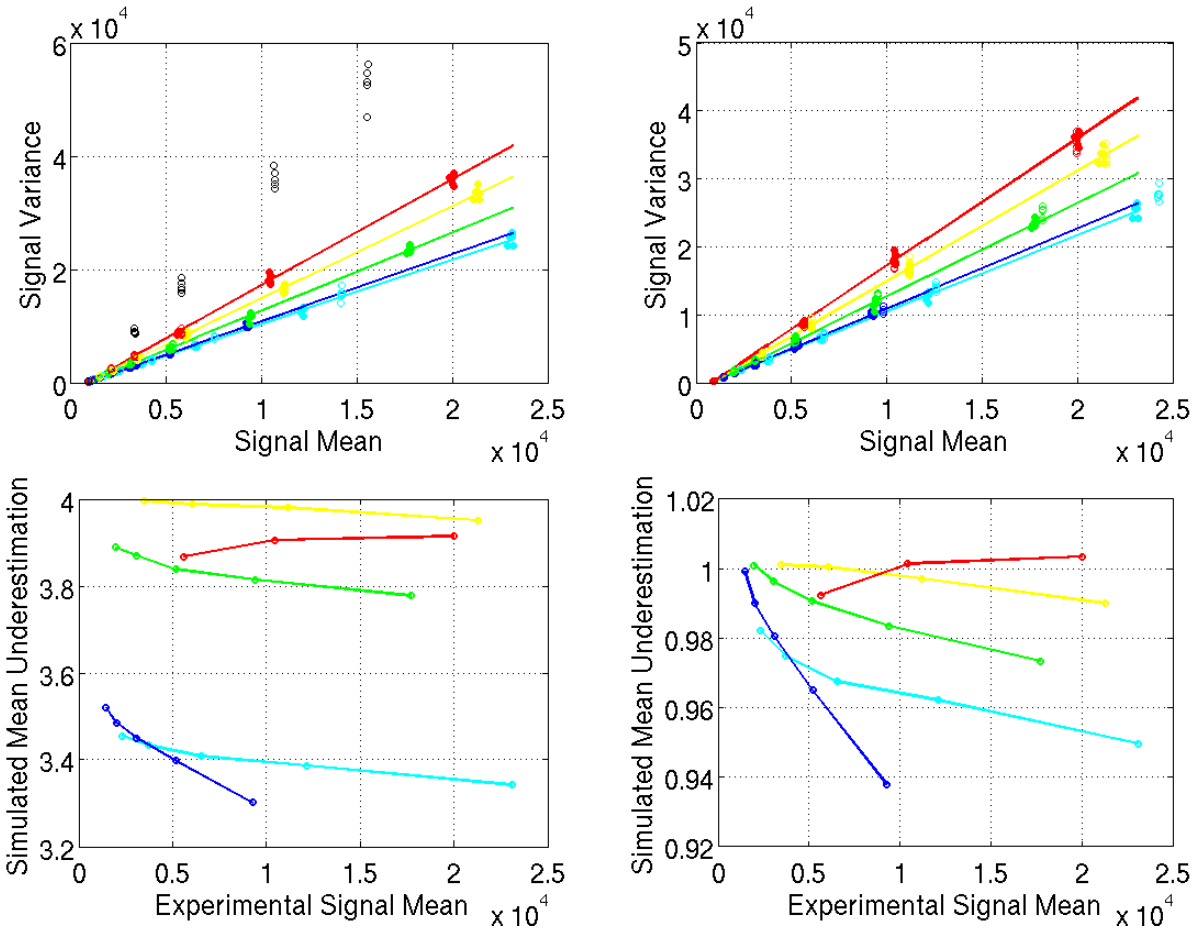

Figure 2.3: (First row, left) Experimental data (filled circles) with equivalent catsim simulations (open circles). The black data points are 140 kVp simulations using $G = 0.039$, twice the correct value. Lines are best fit to experimental data. (Right) Simulations with rescaled spectra. (Second row, left) Ratio of experimental to simulated mean, as a function of experimental mean. (Right) Ratio with rescaled spectra.

2.4 Conclusion

Our experimental results show that the CPG detector model is a valid noise model for a clinical VCT scanner in the low-dose regime. Both detector variance and mean are linear with mA (and by assumption, flux), and their ratio increases with spectra mean energy, as

G	i_d	σ_e	C_{80kVp}	C_{100kVp}	C_{120kVp}	C_{140kVp}
0.195 I.u.	897.2 I.u.	14.99 I.u.	3.52	3.88	3.99	3.90

Table 2.3: Experimentally determined catsim parameters

predicted by the model. Furthermore, the photon gain value, G, predicted by multiple spectra were consistent, showing that G is well-modeled as constant across energy. However at higher flux, the model begins to break down in two ways; first, signal mean becomes slightly sub-linear with mA, which is not modeled in catsim. More importantly, signal variance becomes super-linear with mean. Possible causes of this behavior are an unexpected change in x-ray spectra at high mA, or non-ideal detector behavior at high counts. More investigation is needed to understand this phenomenon. The real experimental values of i_d, σ_e^2, and G were then put into catsim to simulate variance vs mean data. These simulation results agreed well with experiment, up to an intensity scaling factor for each spectra, which we believe to be a real source of error in the catsim spectrum data.

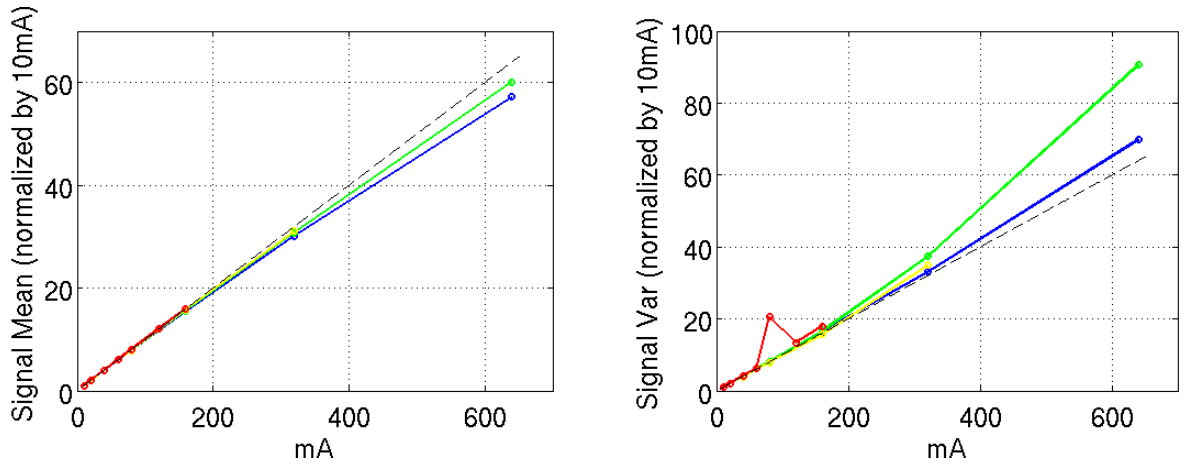

Figure 2.4: (Left) Signal mean as a function of mA, normalized by the mean at 10mA, using the data set for figure 2.1. For all kVp, signal starts to drop below linear at high mA. (Right) Variance vs mA.

Chapter 3

STATISTICAL COMPARISON OF LIKELIHOOD MODELS FOR LOW DOSE X-RAY CT

Iterative reconstruction in x-ray CT using maximum likelihood (ML) estimation seeks to improve image quality over analytic techniques by accurately modeling the x-ray statistics [22]. However there are a variety of statistical models to choose from, each providing a different trade off between data fidelity and computational simplicity. For example, some complex models incorporate the multi-energetic x-ray spectrum, energy-integrating detectors, or readout noise of the detector electronics present in real CT systems. Unfortunately, image reconstruction from sophisticated models is often intractable, and approximations are required. Following the work of Fessler on monoenergetic transmission tomography [11], we simplify the reconstruction task to estimating the depth of material along a single axis, which makes estimation tractable for any model. We use this methodology to evaluate the potential benefit of ML estimation using 1) the Compound Poisson model of quantum noise [34] and 2) additive Gaussian electronic noise models.

3.1 Methodology

The Compound Poisson (CP) model is considered a fairly realistic noise model for x-ray CT data. Unfortunately, there is no closed-form CP likelihood, making ML estimation difficult. However the CP probability density function (PDF) can be directly computed by repeated convolution of Poisson density functions [34]. We simulate an experiment in which a single x-ray attenuation measurement is taken from a monolithic block of material. The material composition and input x-ray spectra are fixed and known, so changes in the measured signal depend only on the unknown material depth, θ. By numerically computing the CP PDF

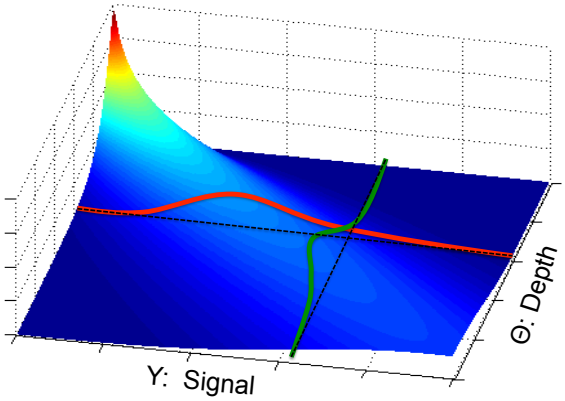

Figure 3.1: A notional 2D likelihood surface showing the relationship between the density and likelihood function of a random variable. A probability density function (shown in red) is a function of the measurement, y, for a fixed parameter, θ. A likelihood function (green) is a function of the parameter, for a fixed measurement. We build the exact Compound Poisson likelihood for a single parameter through repeated calculation of its densities.

for all (reasonable) values of θ, we build an exact, 2D likelihood surface, $L(y, \theta)$, for a grid of possible measurements y and depths θ. Thus, the ML estimate (MLE) for a particular measurement is simply the argmax of a 1D slice of $L(y, \theta)$. A notional visualization of $L(y, \theta)$ is shown in fig. 3.1.

Once the MLE, $\hat{\theta}$, has been determined, its bias and variance can be calculated using,

$$\text{bias}\{\hat{\theta}\} = \mathbb{E}[\hat{\theta} - \theta] = \sum_y (\hat{\theta}(y) - \theta)\, p(y) \tag{3.1}$$

$$\text{var}\{\hat{\theta}\} = \mathbb{E}[(\hat{\theta} - \mathbb{E}[\hat{\theta}])^2] = \sum_y (\hat{\theta}(y) - \mathbb{E}[\hat{\theta}(y)])^2\, p(y) \tag{3.2}$$

where $p(y)$ is the probability distribution of the "true" data-generating model. Note that these are functions of the true depth, θ. Using this approach, we compare bias and variance of several likelihood models in single parameter estimation.

We also consider the case of multiple, iid measurements of a single parameter. In some ways this is overly optimistic, since clinical CT scans typically take only one measurement

from each LOR. However combining multiple measurements more accurately reflects the real tomography problem. In this case, the likelihood is a sum of terms, one for each measurement, and we can find the MLE in the same manner as before. Equations (3.1) and (3.2) become difficult, multi-dimensional summations, though, so we instead use Monte Carlo simulation to calculate MLE bias and variance.

3.2 Quantum Noise

We first compare the effects of quantum noise modeling for ML estimation. We select the following likelihood models for comparison:

1. Poisson: $Y = G\,E_0 X(\lambda_0)$

2. Multi-Energy Poisson: $Y = G\,E_0 X(\sum_j \frac{E_j}{E_0}\lambda_j)$

3. Compound Poisson: $Y = G\,\sum_j E_j X(\lambda_j)$

4. Gaussian: $Y = G\,E_0\,N(\lambda_0, \lambda_0)$

where, Y is a detector measurement, $\lambda_j = I_j e^{-\mu_j \theta}$ are noise-free photon arrivals with energy E_j and initial flux I_j, μ_j is an energy-dependent attenuation coefficient, and θ is the unknown depth of material. Here, the x-ray spectrum, initial intensity, and material composition are all known. λ_0 and E_0 are the average total photon arrivals and average photon energy, respectively. The notation $X(\lambda)$ denotes a Poisson random variable with mean λ, and $N(\mu, \sigma^2)$ a Gaussian random variable with mean μ and variance σ^2. G is a constant gain factor which was experimentally-determined from a GE VCT scanner.

Model 1 has been well studied, leading to fast algorithms [27]. Algorithms also exist for model 2 [7], which accounts for the poly-energetic nature of the x-ray spectrum. The $\frac{E_j}{E_0}$ term is a linear weight used in [36] for energy-integrating detectors. In effect, it modifies the spectrum to match the signal mean of a Compound Poisson. Model 4 approximates the Poisson distribution with a Gaussian with signal-dependent variance.

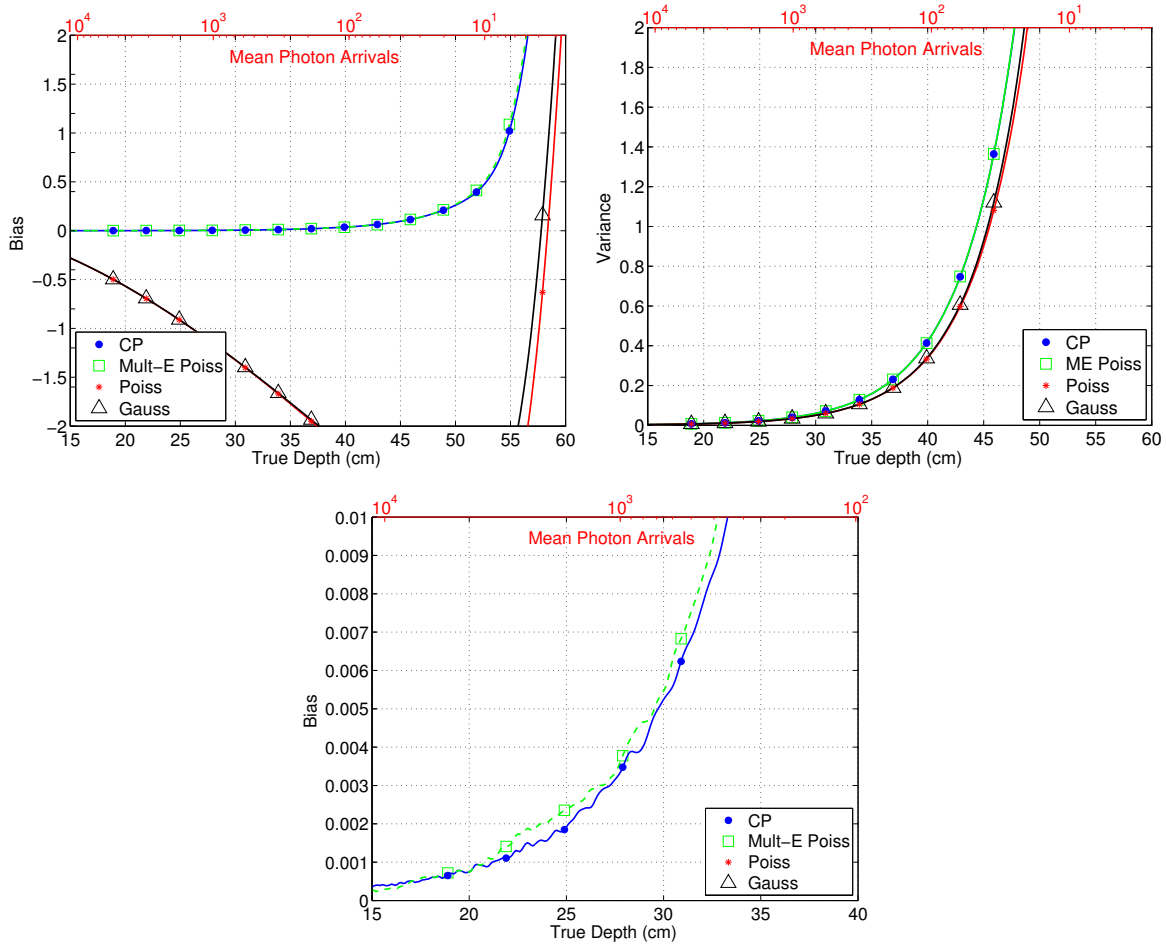

Figure 3.2: Single measurement MLE bias for several likelihood functions, assuming a "true" Compound Poisson model with no electronic noise. The mean photon arrivals, λ_0, for each depth is shown above the figure. The data was calculated analytically from eqs. (3.1) and (3.2). The Poisson model has a strong negative bias due to beam hardening. The multi-energy Poisson accounts well for beam-hardening, and gives nearly identical performance to the Compound Poisson. The variance of each estimator is nearly identical. A zoomed-in plot of bias shows that CP likelihood achieves a slightly lower bias than multi-energy Poisson.

Figure 3.2 shows the MLE bias and variance for each model for the single measurement experiment. Results for the multi-measurement experiment are nearly identical. The input

spectrum was 120kVp and 1mA. The material is 100% water. In calculating bias and variance, the CP is assumed to be the "truth" model. The data was calculated analytically from eqs. (3.1) and (3.2). The Poisson model has a strong negative bias due to beam hardening. The multi-energy Poisson accounts well for beam-hardening, and gives nearly identical performance to the Compound Poisson. Both the Poisson and Gaussian models have slightly lower variance than the poly-energetic models.

3.3 Electronic Noise

We next compare the effects of electronic noise on MLE performance, which we model as additive Gaussian noise with known variance, σ^2, and mean, i_d, (assumed zero). The variance was determined experimentally from a GE VCT scanner. In general, the addition of Gaussian noise can make finding the MLE difficult because the likelihood function has no closed form expression. In our single parameter case, we can construct a new likelihood function by simply convolving the original PDF with a Gaussian PDF and continue the method described above. In order to focus purely on the contribution of electronic noise, we remove the effects of beam hardening in the following way; we assume a material whose attenuation coefficient, $\bar{\mu}$, is constant and independent of energy, which we set equal to the attenuation coefficient of water at 60 keV. Now, the attenuated x-rays always retain the same spectral shape – there is no beam hardening – but are still poly-energetic. In this scenario, the Poisson and Multi-Energy Poisson models are equivalent. We now consider the following likelihood models:

1. Poisson + Gaussian: $Y = G\,E_0 X(\lambda_0) + N(i_d, \sigma^2)$

2. CP + Gaussian: $Y = G\,\sum_j E_j X(\lambda_j) + N(i_d, \sigma^2)$

3. Gaussian: $Y = G\,E_0\,N(\lambda_0, \lambda_0 + \sigma^2)$

4. Poisson

5. Compound Poisson

6. Post-log Least Squares: $P = X(p)$

7. Post-log Weighted LS: $P = N(p, y_i + \sigma^2)$

where $p = e^{-\bar{\mu}\theta}$ is the attenuation coefficient produced by depth θ. These post-log models require a log conversion of the data: $p_i = -\log(y_i/I_0)$. Models in which the Gaussian noise is not explicitly included (#3-6) involve a thresholding step to remove negative (or non-positive) signal values. We again use a 120 kVp, 1mA x-ray source, but now use the Compound Poisson plus Gaussian (CPG) model as "truth".

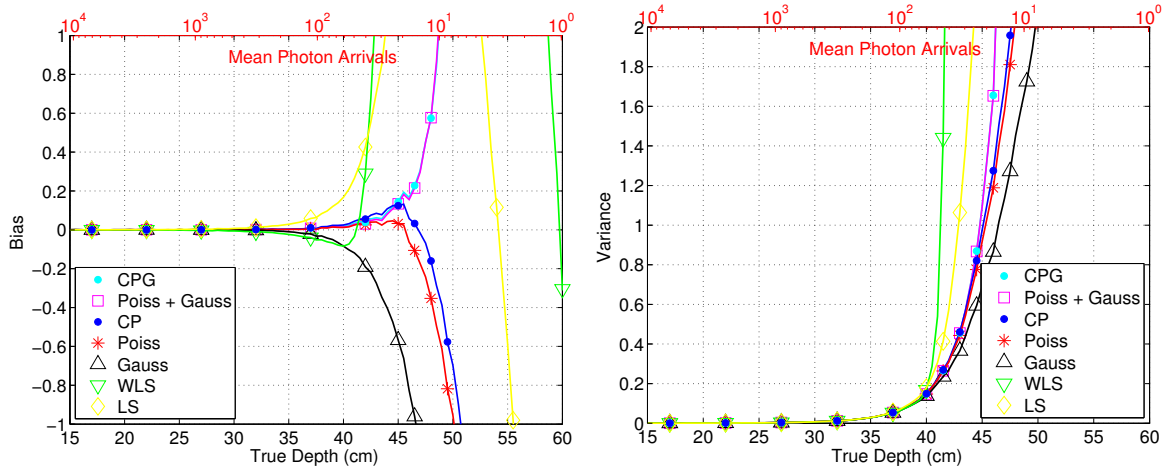

Figure 3.3: MLE bias for several likelihood functions, assuming a "true" Compound Poisson plus Gaussian model. $N = 10$ measurements are used for each MLE, and there is no beam hardening. The data was calculated via Monte Carlo simulation. In the absence of beam hardening, the Poisson + Gaussian and CPG models are nearly identical. Poisson with thresholding has lower bias than CP for several hundred photon arrivals, but is otherwise similar. Post-log models begin to show bias at photon arrivals less than an order of magnitude above pre-log models.

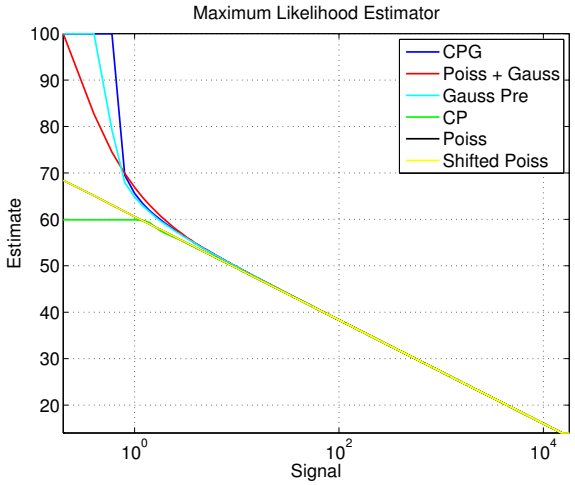

Figure 3.4: Maximum likelihood estimate vs signal level in the single parameter, single measurement model with Gaussian electronic noise and no beam hardening. As the measured signal approaches zero, the MLE of models that explicitly include Gaussian noise in the likelihood tend to increase rapidly, while models that use thresholding then to level off. This results in a positive bias at low signal for models that include electronic noise, and a negative bias for threshold models. At zero, all MLEs go to infinity.

Figure 3.3 shows the MLE bias for each model. In the absence of beam hardening, Poisson and CP models are nearly identical. At signal levels significantly greater than the electronic noise standard deviation (>20 detected photons), thresholding is nearly identical to models that explicitly include electronic noise. At lower signal levels, both types of models begin to fail, but in different ways; threshold models have a systematic negative bias, while explicit electronic noise models tend to have a positive bias. This is because, as the measured signal approaches zero, the MLE of models that include a Gaussian in the likelihood tend to increase rapidly, while models that use thresholding level off (see fig. 3.4). Both biases are of comparable magnitude, and both estimators have similar variance. To partially summarize these results, table 3.1 ranks all models according to the depth of material at which each model fails, where failure is defined as a RMSE of 1 cm.

Model	Mean Detected Photons at RMSE = 1 cm	Flux Required at 50 cm
WLS	56.4 cm	4.3 mA
Poisson + Gaussian	17.4 cm	1.10 mA
Shifted Poisson	16.4 cm	1.05 mA
ME Poisson + Gaussian	14.5 cm	0.93 mA
Comp. Poisson + Gaussian	14.5 cm	0.93 mA
Poisson	13.7 cm	0.87 mA
ME Poisson	12.3 cm	0.79 mA
Comp. Poisson	11.8 cm	0.76 mA

Table 3.1: Ranking of all models simulated according to the depth of material at which each model fails, where failure is defined as a RMSE of 1 cm. Choosing a different failure rate might change the rank order. The second column show the initial flux required to achieve a 1 cm RMSE at 50 cm, which is proportional to the dose savings one might achieve through advanced modeling in this experiment.

3.4 Discussion and Conclusion

We have used a single parameter system to numerically compute the performance of MLEs for several x-ray CT likelihood models, some of which are typically intractable in full tomography. While our single parameter, known material system is a simplification of full x-ray tomography, we believe it reasonably captures the effects of noise modeling. Using this methodology, we compared the relative benefit of advanced statistical models for two particular physical noise sources: quantum noise and electronic noise. For quantum noise, we found that the more accurate Compound Poisson model alters the MLE slightly for detected signals equivalent to $\sim$10 photon counts, and significantly below $\sim$1 count (not shown). These

differences negligibly affected estimator bias and variance. This is because ML estimation in the <1 photon signal regime is dominated by zero signal data, for which Compound Poisson and Poisson MLEs agree. Simply accounting for beam hardening effects with a multi-energy Poisson model is sufficient. For electronic noise, we found that thresholding negative data values results in a systematic negative bias at signal levels at or below the electronic noise standard deviation, while models that explicitly incorporate Gaussian electronic noise have a systematic positive bias, and exhibit worse RMSE.

Chapter 4

MIXED CONFIDENCE ESTIMATION FOR DYNAMIC ITERATIVE CT RECONSTRUCTION

Often in CT imaging, local regions of interest require enhanced resolution or noise properties, depending on the application. For example in targeted reconstruction, the goal is to produce a high-resolution image in a small region of interest without fully reconstructing the entire scanner field of view. 4D CT procedures often wish to reduce noise preferentially in a dynamic region over time. Finally, CT attenuation correction for PET/SPECT requires clearly capturing the boundary between objects, but is not impacted by subtle variations in tissue density.

Previous reconstruction algorithms have taken advantage of spatially-varying confidence in different ways. In targeted reconstruction, several methods use low-resolution analytic reconstructions outside the ROI, then iteratively modify within it [31], [25]. In 4D CT, a common approach is to use a low-resolution composite image, reconstructed from a time-averaged sinogram, to aid reconstruction. Composite images have been used as a weight applied to filtered back projection [32] images, and as a prior term in a total variation minimization algorithm [5]. In our previous work [1, 2] we proposed the Constrain-Static Target-Kinetic (CSTK) reconstruction algorithm as a method to reduce computation time in 4D CT image reconstruction by devoting full computational resources to only the dynamic region of interest. This work extends that effort by presenting an analytic argument, based on an estimator variance analysis [10], that CSTK additionally improves noise levels throughout the image, including the region of interest. This technique could be characterized as having side-information to increase the confidence of the static region, while desiring an optimal image of the lower-confidence, kinetic region. This class of methods is advantageous for

both substantially reducing reconstruction time and reducing noise in the lower-confidence region. We feel this analysis can be extended to the situations above, where locally varying performance can be leveraged. We verify our analytic argument with simulation studies.

4.0.1 CSTK Algorithm

Constrain-Static Target-Kinetic reconstruction (CSTK) is a method to reduce computation time of most iterative 4D CT reconstruction algorithms. It comprises the following steps; 1) classify each image pixel as either *static* or *kinetic* across frames, perhaps using a high-noise estimate of each frame, 2) form a low-noise, low-resolution "composite image" to initialize all frames, and 3) update only the kinetic pixels in each frame. The resulting computation reduction scales linearly with the percentage of dynamic pixels, minus the time to form the composite image. Previous work [2] showed two applications, Retrospective Gated CT Angiography and Dynamic Perfusion CT, in which CSTK provided similar image quality to conventional OSEM reconstruction with 50% dynamic pixels, and therefore 50% compute time.

4.1 Statistical Formulation

4.1.1 Static Model

We start by adopting the standard quadratic approximation to the static transmission tomography problem. We wish to estimate the x-ray attenuation coefficients of each pixel in an image, $\theta = [\theta_1, ..., \theta_m]^T \in \mathbb{R}^m$, from observations $Y = [y_1, ..., y_n]^T \in \mathbb{R}^n$, where $y_i = -\log(p_i/I_0)$ are the post-log-corrected, measured sinogram bins. By taking the second order Taylor series expansion of the Poisson likelihood [27], the system model can by approximated as,

$$P(Y|\,\theta) \approx N(A\theta\,,\,Q) \tag{4.1}$$

where A is the tomographic model, a forward projection operator, and $Q = \text{diag}(1/p_i)$. Recognizing this as a weighted least squares problem, the maximum likelihood estimator

$\hat{\theta}_{ML}$ can be written explicitly as the least squares solution,

$$\hat{\theta}_{ML} = \arg\max_{\theta} P(Y|\theta) = \left(A^T W A\right)^{-1} A^T W Y \tag{4.2}$$

Where $W = Q^{-1}$. It is easy to show that $\hat{\theta}_{ML}$ is unbiased ($E[\hat{\theta}_{ML}] = \mu_{\hat{\theta}} = \theta$) and that $Cov[\hat{\theta}_{ML}] = \Sigma_{\hat{\theta}} = (A^T W A)^{-1}$. If we choose to add a quadratic prior term to control noise amplification, the solution becomes,

$$\hat{\theta}_{MAP} = \arg\max_{\theta} P(\theta|Y) = \left(A^T W A + R\right)^{-1} A^T W Y \tag{4.3}$$

where R describes the prior term. In the following we will assume no prior for simplicity, although this analysis could be extended to the quadratic prior case.

4.1.2 Dynamic Model

We now consider the 4D extension to the static problem in which we want to estimate multiple images over K time frames, or $\boldsymbol{\theta} = \left[\theta^1, \ldots, \theta^K\right]^T$ from $\boldsymbol{Y} = \left[Y^1, \ldots, Y^K\right]^T$, where each $\theta^j \in \mathbb{R}^m$, $Y^j \in \mathbb{R}^n$. A straightforward approach is to treat each frame as a separate static estimation problem, i.e.,

$$\hat{\theta}^j_{ML} = \left(A^T W^j A\right)^{-1} A^T W^j Y^j \tag{4.4a}$$

$$\mu^j_{\hat{\theta}_{ML}} = \theta^j \quad \Sigma^j_{\hat{\theta}_{ML}} = (A^T W^j A)^{-1} \tag{4.4b}$$

but this solution ignores potentially useful information about θ^j in $Y^{\neg j}$, where the $\neg j$ superscript indicates all frames besides j. To improve on this, we separate the parameter space into a *static* and *dynamic* component, namely, $\theta^j = [\theta_s, \theta^j_d]^T$ (see fig. 4.1). Previous work on CSTK [2] suggests practical ways to do this partitioning. It is assumed that θ_s is constant across all frames, and θ^j_d constant only across frame j. This assumption allows us to factor the posterior distribution as,

$$P(\theta^j | \boldsymbol{Y}) = P(\theta_s, \theta^j_d | Y^j, Y^{\neg j}) \tag{4.5}$$

$$= P(\theta^j_d | Y^j, \theta_s)\, P(\theta_s | \boldsymbol{Y}). \tag{4.6}$$

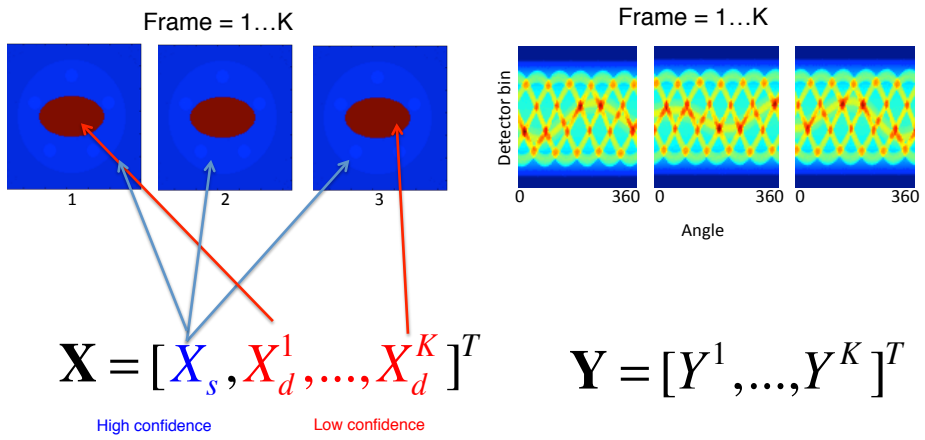

Figure 4.1: Decomposition of image into static and dynamic regions.

The first term in eq. (4.6) is the conditional of eq. (4.1), and so is itself multivariate Gaussian variate. The second term is a marginal of the composite image, which is Gaussian distributed if the composite is Gaussian distributed. Although this may not be true for arbitrary composite images, we restrict the composite image to unbiased linear estimators of eq. (4.1), which we call $\hat{\theta}^c$ with covariance $\Sigma_{\hat{\theta}}^c$. Thus, if:

$$\Sigma_\theta = \begin{bmatrix} \Sigma_{ss} & \Sigma_{sd} \\ \Sigma_{ds} & \Sigma_{dd} \end{bmatrix}$$

then,

$$P\left(\theta_s|\boldsymbol{Y}\right) = N\left(\theta_s,\, \Sigma^c_{ss}\right) \tag{4.7a}$$

$$P\left(\theta^j_d|\, Y^j, \theta_s\right) = N\left(\mu^j_{d|s},\, \Sigma^j_{d|s}\right) \tag{4.7b}$$

$$\mu^j_{d|s} = \mu^j_{\theta_d} + \Sigma^j_{ds}(\Sigma^j_{ss})^{-1}(\hat{\theta}^c - \theta_s) \tag{4.7c}$$

$$\Sigma^j_{d|s} = \Sigma^j_{dd} - \Sigma^j_{ds}(\Sigma^j_{ss})^{-1}\Sigma^j_{sd} \tag{4.7d}$$

where eq. (4.7a) come directly from the marginal distribution and eqs. (4.7b) to (4.7d) from the conditional distribution of a multivariate Gaussian. Notice that eq. (4.7b) depends on θ_s only linearly through its mean. This ensures that the product of eq. (4.7a) and eq. (4.7b) is also Gaussian, parameterized by

$$\mu^j_\theta = \left[\theta_s,\, \theta^j_d\right]^T \tag{4.8a}$$

$$\Sigma^j_\theta = \tag{4.8b}$$

$$\begin{bmatrix} \Sigma^c_{ss} & \Sigma^j_{ds}(\Sigma^j_{ss})^{-1}\Sigma^c_{ss} \\ \Sigma^c_{ss}(\Sigma^j_{ss})^{-1}\Sigma^j_{sd} & \Sigma^j_{d|s} + \Sigma^j_{ds}(\Sigma^j_{ss})^{-1}\Sigma^c_{ss}(\Sigma^j_{ss})^{-1}\Sigma^j_{sd} \end{bmatrix}.$$

Equation (4.8) is the main result of this paper. Notice that $\hat{\theta}_{CSTK} = \mu^j_\theta$ is unbiased, as might be expected (this would not be true if we included a prior term). Also, the covariance can be directly compared with eq. (4.4), the straightforward reconstruction approach. First, the static pixel variance under CSTK is simply Σ^c_{ss}. A reasonable choice for the composite image would be the average of each frame estimate, θ^j, in which case $\Sigma^c_\theta \approx \frac{1}{K}\Sigma^j_{\hat{\theta}_{ML}}$. More interestingly, the dynamic pixel variance (lower right term of 4.8b), is the sum of two terms; Σ^j_{dd}, and a correction factor, Σ^*, where,

$$\Sigma^* = \Sigma^j_{ds}(\Sigma^j_{ss})^{-1}\Sigma^c_{ss}(\Sigma^j_{ss})^{-1}\Sigma^j_{sd} \;-\; \Sigma^j_{ds}(\Sigma^j_{ss})^{-1}\Sigma^j_{sd}.$$

To better understand 4.8b, consider two cases; when $\Sigma^c_{ss} = \mathbf{0}$ and $\Sigma^c_{ss} = \Sigma^j_{ss}$. In the first case, the full dataset $\boldsymbol{Y}$ provides perfect information about θ_s, and the covariance of the dynamic portion of the image, θ^j_d, reduces to the conditional covariance $\Sigma^j_{d|s}$, which is

guaranteed to be smaller than Σ_{dd}^{j}. In the second case, the covariance of the static portion of the composite image equals the covariance of the static portion of a single frame; i.e. the composite image did not improve the estimate of θ_s. Then $\Sigma^* = \mathbf{0}$, and the covariance is simply Σ_{dd}^{j}. In practice, we expect to lie somewhere in between these two bounds; the composite provides some extra, but not perfect, information about θ_s, which helps lower the covariance of *both* dynamic and static pixel estimates.

4.1.3 Numerical Validation

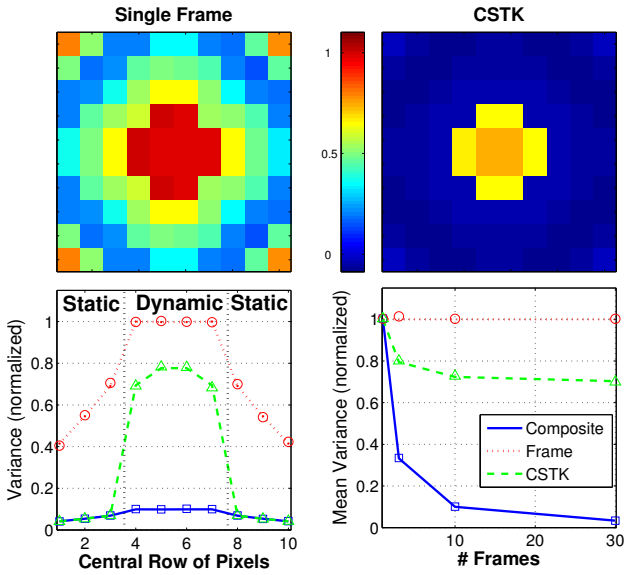

Figure 4.2: First row presents theoretical variance image for single frame reconstruction and CSTK. The CSTK estimate achieves lower variance everywhere, including the central, dynamic region (yellow region, in upper right). Lower left: Profile through central row of variance image for 3 reconstruction methods, with simulated values as datum and theoretical values as lines. Lower right: Average pixel variance of the dynamic region as a function of frames acquired. Simulated data are datum and theoretical values are lines.

To validate our analysis we compared our predicted covariance with those measured from simulations of a small, 10x10 pixel image. We used a sinogram of 15 detectors x 16 views to make explicit computation of the pseudoinverse tractable. We added Gaussian noise according to the approximate signal model in eq. (4.1). Figure 4.2 shows the predicted variance of each pixel using simple framed recon and CSTK, from eq. (4.4) and eq. (4.8), respectively. While the true image is static, we nonetheless treat the inner pixels (red pixels in fig. 4.2, upper left) as dynamic, and collected 10 frames of data. The CSTK image shows dramatic variance reduction in the static region, but also significant reduction in the inner, dynamic region. We then computed $\hat{\theta}_c$, $\hat{\theta}^1_{ML}$, and $\hat{\theta}_{CSTK}$ using $\boldsymbol{Y}$, Y^1, and the CSTK method. Figure 4.2 shows that the predicted, theoretical variance of a single row of pixels agrees well with the sample variance over 1000 noise realizations. Figure 4.2 also shows the average variance of over all dynamic pixels as the number of frames increases. The framed recon is constant, since it doesn't share information across frames, while the composite image shows the typical $\frac{1}{N}$ variance reduction. The CSTK variance initially decreases rapidly, then plateaus to the average variance of Σ^j_{dd}.

4.2 Experimental Results

To demonstrate the utility of these results in a more realistic dynamic CT scenario, we simulated data from a 128x128 dynamic target, pictured in the top row of fig. 4.3 (the central object is moving continuously at a rate of 8 pixels per frame, where a frame is a full 360 degree revolution of the detector). The dynamic pixels were chosen a priori as an ellipse, pictured in fig. 4.4. Note, the dynamic pixels are not truly constant within a frame, as assumed in the statistical analysis. The measurement sinogram was 200 detectors x 150 views per frame for 5 frames, with Poisson distributed noise; $p_i \sim \text{Poisson}(\lambda_i)$, where $\lambda_i = I_0\, e^{-\sum_j c_{i,j}\theta_j}$ across line of response i. I_0 was chosen as low as possible such that $p_i > 0$ for all i. For reconstruction, we use iterative coordinate descent (ICD), but only update dynamic pixels. The choice of composite image is important, and must balance low-noise and computation time. We chose our composite image to be the ICD reconstruction of

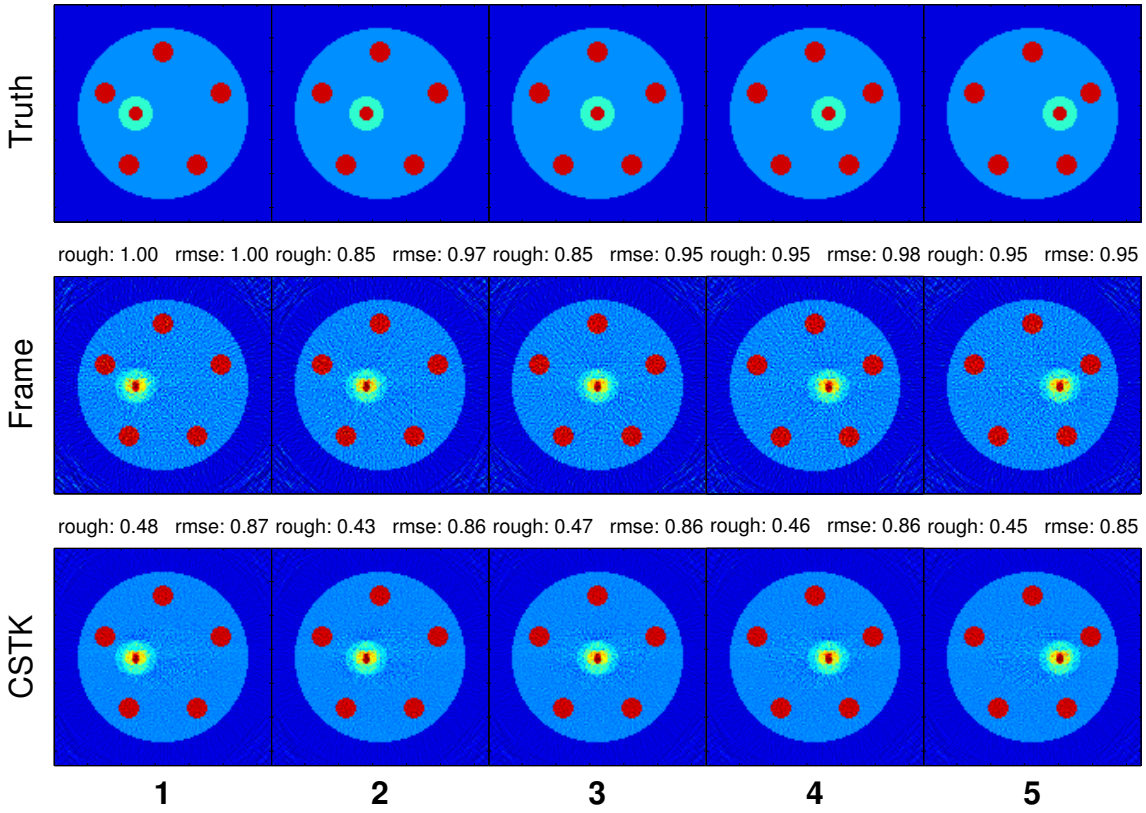

Figure 4.3: Top row: Truth image in the center of each time frame. Row 2-3: Framed ML, and CSTK estimate for each frame. Above each image are image quality metrics of dynamic pixels only. RMSE is the total root mean squared error of the reconstruction. 'rough' is the standard deviation of pixels in a flat patch in the dynamic region. Both metrics are normalized.

the average sinogram across all frames. While $\theta^c = \frac{1}{K}\sum_j \hat{\theta}^j_{ML}$ was suggested previously, it requires estimation of the entire image for each frame, which we would like to avoid. As an approximation, we use $\hat{\theta}_{ML}$ on the entire data set, $\boldsymbol{Y}$, computed with ICD.

Figure 4.3 compares the CSTK reconstruction of each frame with that of simple framed recon. The CSTK noise appears lowest, particularly in the static region. However we are primarily concerned with performance in the dynamic region. We put a circular mask around

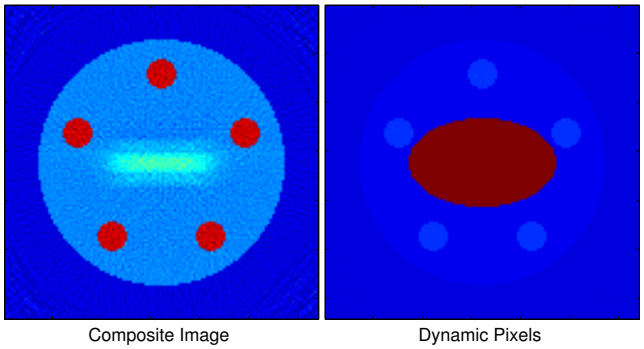

Figure 4.4: Left: Composite image using ICD on the full data set, $\boldsymbol{Y}$. Right: Mask for elliptical region of dynamic pixels.

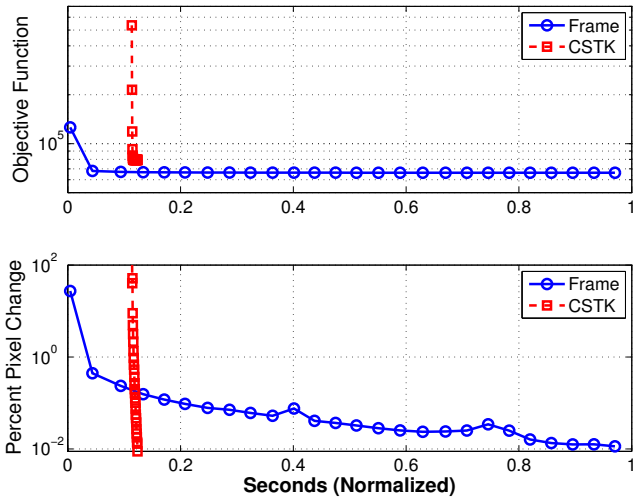

Figure 4.5: Convergence metrics for reconstruction of first frame. Each data point is a full iteration of ICD. Top: Negative log-likelihood value vs. computation time for simple framed (blue) and CSTK (red) recon. Bottom: Mean percent pixel change vs. computation time. Algorithm was terminated when mean pixel change was $< 1\%$.

the moving ball in each frame. The remaining dynamic pixels should be constant, and so we take their spatial variance as a measure of image variance. We present the image roughness

(pixel-to-pixel variance) in the uniform portion of the dynamic region and total RMSE in the dynamic region. Both metrics are normalized and shown above each image. Figure 4.5 compares the convergence of both methods. To fairly account for the composite compute time, T_c, in CSTK, the CSTK data are offset by T_c/K, the time to compute reconstruction from averaged sinogram. From fig. 4.5 and fig. 4.3, the CSTK method achieves about 15% less noise in the dynamic region in 90% less compute time than simple frame recon. Both factors should increase with number of frames and the percentage of static pixels, as long as the static/dynamic pixel segmentation remains accurate.

4.3 Discussion and Conclusion

We have shown that CSTK reconstruction, applied to ML estimation of cardiac gated CT imaging, can both save computation time *and* lower noise throughout the image. The results confirm the intuition that side information, in the form of increased confidence of particular parameters, can decrease noise. As the strength of the side information increases, i.e. more frames or a higher percentage of static pixels in CSTK and the correlation between side information and pixels of interest increases, the noise reduction is greater. Future work is needed to better understand the effect of CSTK in MAP reconstruction, when a prior term is added to the cost function.

Chapter 5

CONCLUSION AND FUTURE WORK

The work presented here presents examples of the role statistical analysis and modeling can play in the trade off between image quality, dose, and reconstruction time in low-dose CT imaging. The first work established the validity of compound Poisson plus Gaussian (CPG) model with real experimental data, and suggests improvements to the catsim simulation tool. In the second work, a statistical methodology was developed to compare the relative importance of various statistical effects in low-dose CT modeling. It suggested that incorporating beam hardening effects into the statistical model improves estimator bias, although a pre-processing beam hardening correction still achieved reasonable performance. However capturing the full compound Poisson nature of the x-ray spectrum was not significant above a few photons of signal. Electronic noise caused a more significant deterioration of performance at low dose. However explicitly modeling this noise did not improve performance in the low flux regime, when the mean measured photon signal approximately equaled the electronic noise standard deviation. These results were presented at the IEEE Medical Imaging Conference (MIC), 2014 [23]. Future work will compare these models to experimental data from a CT scanner, and propose alternative noise variance reduction techniques to combat electronic noise degradation at ultra low flux. The third work gave a statistical argument for information sharing in redundant dynamic CT imaging that both reduces noise and saves computation. This method is an alternative to more complex motion correction techniques, and is best suited for applications in which quantitative accuracy is desired in a dynamic region of interest. The work was presented at the Conference on Image Formation in X-ray CT, 2014 [24]. Future work will develop a fast variance prediction method to estimate the Mixed Confidence Estimation performance gain, and apply this method on a clinical data

set.

BIBLIOGRAPHY

[1] A. Alessio and P. Kinahan. Statistical reconstruction of targeted regions with application to CT and PET cardiac imaging. *American Journal of Roentgenology (supplement)*, 186(4):A46, 2006.

[2] A. Alessio and P. La Riviere. Constrain static target kinetic iterative image reconstruction for 4D cardiac CT imaging. *Proc. SPIE*, 7873:78730S–78730S–7, 2011.

[3] A.H. Andersen and A.C. Kak. Simultaneous algebraic reconstruction technique (sart): A superior implementation of the art algorithm. *Ultrasonic Imaging*, 6(1):81–94, 1984. cited By 477.

[4] C.A. Bouman and K. Sauer. A unified approach to statistical tomography using coordinate descent optimization. *IEEE Transactions on Image Processing*, 5(3):480–492, 1996. cited By 203.

[5] G.H Chen, J. Tang, and S. Leng. Prior image constrained compressed sensing (piccs): a method to accurately reconstruct dynamic ct images from highly undersampled projection data sets. *Med Phys*, 35(2):660–3, Feb 2008.

[6] H Dang, A S Wang, Z Zhao, M S Sussman, J H Siewerdsen, and J W Stayman. Joint estimation of deformation and penalized-likelihood ct reconstruction using previously acquired images. In *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med*, pages 424–427, 2013.

[7] B. De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. An iterative maximum-likelihood polychromatic algorithm for CT. *Medical Imaging, IEEE Transactions on*, 20(10):999–1008, Oct 2001.

[8] I.A. Elbakri and J.A. Fessler. Statistical image reconstruction for polyenergetic x-ray computed tomography. *IEEE Transactions on Medical Imaging*, 21(2):89–99, 2002. cited By 198.

[9] H Erdogan and J A Fessler. Accelerated monotonic algorithms for transmission tomography. In *Image Processing, 1998. ICIP 98. Proceedings. 1998 International Conference on*, pages 680–684, 1998.

[10] J. Fessler. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): applications to tomography. *Image Processing, IEEE Transactions on*, 5(3):493–506, Mar 1996.

[11] J.A. Fessler. Hybrid poisson/polynomial objective functions for tomographic image reconstruction from transmission scans. *Image Processing, IEEE Transactions on*, 4(10):1439–1450, Oct 1995.

[12] S Geman and D McClure. Bayesian image analysis: an application to single photon emission tomography. *Amer. Statist. Assoc*, pages 12–18, 1985.

[13] P. Gilbert. Iterative methods for the three-dimensional reconstruction of an object from projections. *Journal of Theoretical Biology*, 36(1):105–117, 1972. cited By 459.

[14] M. Goitein. Three-dimensional density reconstruction from a series of two-dimensional projections. *Nuclear Instruments and Methods*, 101(3):509 – 518, 1972.

[15] R Gordon, R Bender, and G T Herman. Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray photography. *Journal of Theoretical Biology*, 29(3):471 – 481, 1970.

[16] J Hsieh. *Computed Tomography: Principles, Design, Artifacts, and Recent Advances*. SPIE Press monograph. SPIE Press, 2003.

[17] H. M. Hudson and R. S. Larkin. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Transactions on Medical Imaging*, 13(4):601–609, 1994. cited By 1919.

[18] P.J. La Riviere, Junguo Bian, and P.A. Vargas. Penalized-likelihood sinogram restoration for computed tomography. *Medical Imaging, IEEE Transactions on*, 25(8):1022–1036, Aug 2006.

[19] K Lange. Convergence of em image reconstruction algorithms with gibbs smoothing. *IEEE Transactions on Medical Imaging*, 9(4):439–446, 1990. cited By 189.

[20] K. Lange and R. Carson. Em reconstruction algorithms for emission and transmission tomography. *Journal of Computer Assisted Tomography*, 8(2):306–316, 1984. cited By 852.

[21] J Nuyts, B De Man, J A Fessler, W Zbijewski, and F J Beekman. Modelling the physics in iterative reconstruction for transmission computed tomography. *Nuclear Science Symposium and Medical Imaging Conf. (NSS/MIC), IEEE*, June 2014.

[22] J. Nuyts, B. De Man, J.A. Fessler, W. Zbijewski, and F.J. Beekman. Modelling the physics in iterative reconstruction for transmission computed tomography. *Phys Med Biol.*, 58(12):R63–R96, Jun 2013.

[23] D S Perlmutter, S M Kim, P E Kinahan, and A M Alessio. Comparison of Maximum Likelihood Estimators for Low Flux X-ray CT. In *Nuclear Science Symposium and Medical Imaging Conf. (NSS/MIC), IEEE*, 2014.

[24] D S Perlmutter, S M Kim, P E Kinahan, and A M Alessio. Mixed Confidence Estimation for Iterative CT Reconstruction. In *The 3rd Intl. Conf. on Image Formation in X-ray Computed Tomography*, pages 29–32, 2014.

[25] E. Rashed and et al. Practical statistical models for region-of-interest tomographic reconstruction and long object problem. In *Nuclear Science Symposium Conference Record, 2007. NSS '07. IEEE*, volume 5, pages 3505–3511, Oct 2007.

[26] A J Rockmore and A Macovski. A maximum likelihood approach to emission image reconstruction from projections. *Nuclear Science, IEEE Transactions on*, 23(4):1428–1432, 1976.

[27] K. Sauer and C. Bouman. A local update strategy for iterative reconstruction from projections. *Signal Processing, IEEE Transactions on*, 41(2):534–548, Feb 1993.

[28] L.A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *Medical Imaging, IEEE Transactions on*, 1(2):113–122, Oct 1982.

[29] E Y Sidky and X Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in medicine and biology*, 53(17):4777, 2008.

[30] J.W. Stayman and J.A. Fessler. Compensation for nonuniform resolution using penalized-likelihood reconstruction in space-variant imaging systems. *Medical Imaging, IEEE Transactions on*, 23(3):269–284, March 2004.

[31] C.W. Stearns, R.M. Manjeshwar, and S.D. Wollenweber. An efficient algorithm for targeted reconstruction of tomographic data. In *Nuclear Science Symposium Conference Record, 2006. IEEE*, volume 5, pages 2808–2811, Oct 2006.

[32] M. Supanuch and et al. Radiation dose reduction in time-resolved CT angiography using highly constrained back projection reconstruction. *Phys Med Biol*, 54(14):4575–93, Jul 2009.

[33] J.-B. Thibault, K.D. Sauer, C.A. Bouman, and J. Hsieh. A three-dimensional statistical approach to improved image quality for multislice helical ct. *Medical Physics*, 34(11):4526–4544, 2007. cited By 363.

[34] Bruce R. Whiting. Signal statistics in x-ray computed tomography, 2002.

[35] Jingyan Xu and B M W Tsui. Electronic Noise Modeling in Statistical Iterative Reconstruction. *Image Processing, IEEE Transactions on*, 18(6):1228–1238, 2009.

[36] Jingyan Xu and B.M.W. Tsui. Quantifying the importance of the statistical assumption in statistical X-ray CT image reconstruction. *Medical Imaging, IEEE Transactions on*, 33(1):61–73, Jan 2014.

[37] Mehmet Yavuz and Jeffrey A. Fessler. Statistical image reconstruction methods for randoms-precorrected {PET} scans. *Medical Image Analysis*, 2(4):369 – 378, 1998.