

IMAGE RECONSTRUCTION FOR X-RAY COMPUTED TOMOGRAPHY
IN SECURITY SCREENING APPLICATIONS

A Dissertation

Submitted to the Faculty

of

Purdue University

by

Sherman Jordan Kisner

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

December 2013

Purdue University

West Lafayette, Indiana

ACKNOWLEDGMENTS

I would like to thank first and foremost my adviser Prof. Charles Bouman. The inspiration he has provided me over the years is immeasurable. Also the other members of my advisory committee, including Prof. Ken Sauer whom I'm honored to have had the opportunity to work with; and Professors Jan Allebach, Thomas Talavage and Peter Doerschuk who have all played important roles in my education from as far back as my undergraduate days. I would also like to thank lab-mates Singanallur Venkatakrishnan ("Venkat") for numerous stimulating discussions, and Pengchong Jin who wrote the initial version of the helical reconstruction code used in part of my work. I would like also to acknowledge the good people at Morpho Detection, in particular Dr. Sondre Skatter, who has been instrumental in facilitating our collaboration, and also the ALERT Center at Northeastern University and DHS who sponsored much of this work. And finally, special mention for my partner Eri Haneda who has given endless support in life, and inspiration toward completing my graduate work.

PREFACE

Most of the work presented in this dissertation was carried out in collaboration with vendors of security screening systems. Not surprisingly, in this field there are certain sensitive topics that cannot be presented in the public domain. While it was tempting to write a more comprehensive, but confidential, thesis, a strong effort was made to assemble components of the work that could be published in order to, I hope, benefit the community. For this reason, there are some aspects of the experimental analysis that are left intentionally vague, or altogether omitted. My hope is that this does not detract from the overall impact of the work.

TABLE OF CONTENTS

	Page
LIST OF TABLES	vi
LIST OF FIGURES	vii
ABBREVIATIONS	xi
ABSTRACT	xii
1 INTRODUCTION	1
2 BACKGROUND	4
2.1 Model-Based Iterative Reconstruction	4
2.2 Prior Modeling	6
2.3 Substitute Function ICD for the QGGMRF Prior	7
3 3D MBIR FOR A HELICAL-SCAN CT BAGGAGE SCANNER	11
3.1 Multi-slice Helical Scan CT	11
3.2 Reconstruction Parameters	12
3.3 Detector Afterglow	12
3.4 Fan Angle Offset	16
3.5 Data Weighting Matrix	17
3.6 Comparison of MBIR and DFM Reconstructions	21
3.7 Summary	26
4 SPARSE VIEW ANGLE ITERATIVE CT RECONSTRUCTION	27
4.1 Projection Model	27
4.1.1 Wide-beam projector	28
4.2 Experimental Results	30
4.2.1 Effect of sparse view angles	31
4.2.2 Effect of clutter on CT number accuracy	33
4.3 Discussion	37

	Page
4.4 Summary	38
4.5 Supplemental Analysis of Wide-Beam Projector Model	39
5 ADAPTIVE PRIORS FOR ITERATIVE IMAGE RECONSTRUCTION	41
5.1 Homogeneous Priors	42
5.2 Segmentation-Adapted Prior	42
5.3 Bilateral Adaptive Prior	44
5.4 Experimental Results	47
5.4.1 Results for Segmentation-adapted Prior	49
5.4.2 Results for Bilateral Adaptive Prior	57
6 ESTIMATION OF NOISE MODEL PARAMETERS	63
6.1 Noise Model for X-ray Projections	63
6.2 Estimation of Model Parameters	65
6.3 Experimental Parameter Estimation	66
6.4 Simulation of Projection Data	67
LIST OF REFERENCES	69
VITA	72

LIST OF TABLES

Table	Page
4.1 Image and Reconstruction Parameters	30
4.2 Root mean square error of reconstructions in Fig. 4.2. Units are offset Hounsfield (air=0).	33
4.3 Statistics for the reconstructed CT numbers for a randomly placed synthetic target. The values were computed by averaging over 60 trials of placing a round 1.7cm uniform target (1400 HU) at random locations within the bag perimeter in the ground truth images (See Fig. 4.3). The <i>Dev.</i> is the average deviation of reconstructed target voxels from the true value. Similarly, the <i>RMSE</i> is the root mean square deviation from the true target value. All values are in offset Hounsfield units (air=0). . .	36
5.1 Image and Reconstruction Parameters	48
5.2 Accuracy of segmentation-assisted reconstruction for a round 4 cm target using various clique weights at the edges and interior of segmented regions. The ordered triples contain the (<i>target deviation</i> , <i>target RMSE</i> , <i>image RMSE</i>), where the <i>target deviation</i> is the average deviation from 1400 HU within the target region, and the <i>target RMSE</i> and <i>image RMSE</i> are the root mean square error of the reconstruction within the target region only, and across the entire image respectively. The CT values for the original homogeneous-prior reconstruction are those indicated in the boxes (unity weighting coefficients). The <i>edge weight</i> was applied to pairwise cliques that fall across an edge of a segmented component boundary, and the <i>interior weight</i> was applied to cliques that fall in the interior of a segmented component.	56
5.3 Accuracy of bilateral adaptive prior reconstruction for the two target types. The ordered triples contain the (<i>target deviation</i> , <i>target RMSE</i> , <i>image RMSE</i>), where the <i>target deviation</i> is the average deviation from 1400 HU within the target region, and the <i>target RMSE</i> and <i>image RMSE</i> are the root mean square error of the reconstruction within the target region only, and across the entire image respectively. The CT values for the original homogeneous-prior reconstruction are listed in the top row. Subsequent rows list the results after k iterations of, (1) readjusting the clique weights, and (2) recomputing the MAP estimate with the new posterior.	62

LIST OF FIGURES

Figure	Page
2.1 Algorithm for substitute function ICD optimization of a cost containing the QGGMRF prior. The algorithm allows multiple local updates which can speed convergence (N_{sub} of 2 or 3 is typically sufficient). Also included is an optional over-relaxation step, which often speeds convergence in earlier iterations, however this may prevent full convergence to the global minimum. Two approaches are to (1) unset the flag after a certain number of iterations, or (2) step down the over-relaxation factor toward 1.0 as the iterations increase.	10
3.1 Illustration of the effect of afterglow correction on MBIR reconstructions.	15
3.2 Accounting for a small displacement in the relative mounting position of the detector array. The horizontal axis is in the z-direction (perpendicular to the plane of gantry rotation).	16
3.3 Effect of power law weighting in the data matrix entries. Target is a large uniform-density acetal cylinder (diam=15cm).	19
3.4 Effect of power law weighting in the data matrix entries. Target is a uniform-density acetal cylinder (diam=8cm) with four tungsten pins inset.	20
3.5 Comparison of DFM and MBIR reconstructions for a bag containing clothing and a can of nuts.	22
3.6 Comparison of DFM and MBIR reconstructions on a baggage scan. The region highlights a bulk uniform material susceptible to distortion from a nearby object.	23
3.7 Comparison of DFM and MBIR reconstructions on a baggage scan. The region highlights a bulk uniform material susceptible to distortion from a nearby object.	24
3.8 Comparison of DFM and MBIR reconstructions on a baggage scan. The region highlights the discrimination between two adjacent objects. . . .	25

Figure	Page
4.1 High resolution CT bag scan used as ground truth for the current experiment. Gray scale is in offset Hounsfield units (HU), where Air=0 HU. The original CT image contained reconstruction artifacts due the diversity in composition and morphology of items in the bag, which were removed by masking out the background (set to 0 HU), and retaining only the visible objects. A linear wide-beam projector was then used to generate sinogram data for the current analysis.	31
4.2 Image reconstructions from sparse view angle projection data. Four parallel projection data sets are considered, containing 64, 32, 16, and 8 view angles, uniformly spaced between 0 and 180 degrees. Reconstructions include filtered back projection (FBP) and MBIR using a Gaussian Markov random field prior (GMRF), and a q-generalized Gaussian MRF prior (QGGMRF). The gray scale range for all results shown is [0,2000] HU, as in the ground truth image of Fig. 4.1.	32
4.3 Ground truth images for investigating the effect of clutter on CT reconstruction accuracy. The bag contents from the ground truth image in Fig. 4.1 have been masked out to create a <i>low clutter</i> scene in (a). For evaluation, a synthetic target of uniform value (1400 HU) has been added, as highlighted by the box near the center of the low and high clutter scenes. A close-up of the target in (c) also shows a reference line highlighting voxels that will be examined in the experiment.	34
4.4 Reconstructions zoomed to the target area indicated in Figs. 4.3(a) and 4.3(b). All results are from 32-view angle parallel projection data, the top row generated from the low clutter scene of Fig. 4.3(a), and the bottom row from the high clutter scene of Fig. 4.3(b).	35
4.5 CT values from sparse view reconstructions of Fig. 4.4 for voxels along a line through the center of the target region (See Fig. 4.3(c)). Also shown are the true voxel values from ground truth.	35
5.1 Algorithm for model-based reconstruction using a segmentation-adapted prior. \tilde{x} is a label image from an initial reconstruction. (Here, a $\tilde{x}_s = 0$ label indicates a region not of interest.) The second optimization should use the first reconstruction as an initial condition to accelerate convergence. An ICD algorithm for the optimizations is described in detail in Chapter 2, including pseudocode in Figure 2.1.	43

Figure	Page
5.2 Algorithm for model-based reconstruction using a bilateral adaptive prior. The process of redefining the weights and reconstructing is repeated a fixed number of times. In each stage, the MAP estimate should use the previous reconstruction as an initial condition to accelerate convergence. For the MAP estimates, a detailed ICD algorithm was covered in Chapter 2, including pseudocode in Figure 2.1).	46
5.3 Targets used in simulated cargo containers for assessment of segmentation-adapted and bilateral prior reconstruction. The round target is 4 cm in diameter, the thin target is 4.4 mm wide, and both have uniform attenuation of 1400 HU.	47
5.4 Segmentation-adapted reconstruction for a target in <i>high clutter</i> level conditions. The target is 4 cm in diameter, with a 2.2 mm pixel size. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is a segmentation of the initial reconstruction, and (d) is the segmentation-adapted reconstruction (edge weight=0, interior weight=5). 1-D profiles along the indicated blue lines are provided in separate graphs.	50
5.5 Segmentation-adapted reconstruction for a 1400 HU target in <i>high clutter</i> level conditions. Vertical and horizontal pixel profiles through target.	51
5.6 Segmentation-adapted reconstruction for a target in <i>medium clutter</i> level conditions. The target is 4 cm in diameter, with a 2.2 mm pixel size. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is a segmentation of the initial reconstruction, and (d) is the segmentation-adapted reconstruction (edge weight=0, interior weight=5). 1-D profiles along the indicated blue lines are provided in separate graphs.	52
5.7 Segmentation-adapted reconstruction for a 1400 HU target in <i>medium clutter</i> level conditions. Vertical and horizontal pixel profiles through target.	53
5.8 Segmentation-adapted reconstruction for a target in <i>low clutter</i> level conditions. The target is 4 cm in diameter, with a 2.2 mm pixel size. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is a segmentation of the initial reconstruction, and (d) is the segmentation-adapted reconstruction (edge weight=0, interior weight=5). 1-D profiles along the indicated blue lines are provided in separate graphs.	54
5.9 Segmentation-adapted reconstruction for a 1400 HU target in <i>low clutter</i> level conditions. Vertical and horizontal pixel profiles through target.	55

Figure	Page
5.10 Bilateral adaptive reconstruction for the <i>round 4 cm</i> target. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is the bilateral adaptive reconstruction with coefficients determined from the result in (b), and (d) shows the result of the bilateral reconstruction after 16 iterations of determining the coefficients followed by re-optimization. 1-D profiles along the indicated blue lines are provided in separate graphs.	58
5.11 Bilateral adaptive reconstructions for the <i>round 4 cm</i> target. Profiles are shown for reconstructions after various iterations of prior readjustment.	59
5.12 Bilateral adaptive reconstruction for the <i>thin 4.4 mm</i> target. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is the bilateral adaptive reconstruction with coefficients determined from the result in (b), and (d) shows the result of the bilateral reconstruction after 16 iterations of determining the coefficients followed by re-optimization. 1-D profiles along the indicated blue lines are provided in separate graphs.	60
5.13 Bilateral adaptive reconstruction for the <i>thin 4.4 mm</i> target. Profiles are shown for reconstructions after various iterations of prior readjustment.	61

ABBREVIATIONS

2D	two-dimensional
3D	three-dimensional
ADC	analog-to-digital converter
ART	algebraic reconstruction technique
ATD	automated threat detection
ATR	automatic threat recognition
CT	computed tomography
DFM	direct Fourier method
FBP	filtered back projection
HU	Hounsfield units
ICD	iterative coordinate descent
MAP	maximum a posteriori
MBIR	model-based iterative reconstruction
MRF	Markov random field
GMRF	Gaussian Markov random field
GGMRF	generalized Gaussian Markov random field
QGGMRF	Q-generalized Gaussian Markov random field
RMSE	root mean-square error
TV	total variation

ABSTRACT

Kisner, Sherman J. Ph.D., Purdue University, December 2013. Image Reconstruction for X-ray Computed Tomography in Security Screening Applications. Major Professor: Charles A. Bouman.

X-ray tomographic systems have increasingly widespread use in security screening applications. For example, most major airports now utilize X-ray CT systems for efficient screening of baggage and cargo. While decades of research has benefited CT for medical diagnostics, a number of practical differences in the security application present a new set of challenges for the reconstruction problem. For example, the size and composition of the scan subjects, the throughput requirements, and the measure of image quality are all factors that lead to a different set of design considerations. This thesis investigates the application of model-based iterative reconstruction (MBIR) methods for X-ray CT systems in the security context. This reconstruction approach is demonstrated to produce high quality images that are much less susceptible to image distortion compared to direct reconstructions using direct Fourier methods or filtered back projection. Presented first is a mapping of a fully 3D MBIR algorithm to a multi-slice helical-scan CT system certified for baggage screening. Model enhancements for the system are discussed, and MBIR reconstructions are presented alongside direct reconstructions for a set of real baggage scans. Also investigated is the performance of model-based methods for reconstruction on sparse-view angle security CT systems. Advances in the forward and prior modeling are demonstrated to improve several aspects of the reconstructed images, including resolution, artifact suppression, and CT accuracy.

1. INTRODUCTION

While X-ray computed tomography (CT) is predominantly associated with medical diagnostics, it has also become an integral component of aviation security systems for efficient inspection of bags and containers prior to transport [1–3]. While the underlying principals are largely the same for medical and security CT, a different set of constraints are associated with security CT systems, such as the physical size and diversity of the scanned objects, the acceptable X-ray energy and dosage, the scan time requirements, and the sense of image quality. Such constraints present some new opportunities and challenges for the CT reconstruction problem [4, 5].

Security CT systems are tasked with producing a three dimensional image (or a series of 2D images) of the contents of a suitcase or cargo container that is suitable for both visual inspection and for analysis by automatic threat detection (ATD) algorithms. The quality of the CT images is a key factor in the performance of a screening system because any container that cannot be cleared based on the reconstruction must undergo additional inspection, requiring significantly more time and resources.

CT image quality is governed not only by the design of the physical system, but also by the image reconstruction methods employed, which transform the raw X-ray projection measurements into meaningful images [6–8]. Generally the choice of the reconstruction method involves a trade-off between the resilience to image noise and distortion, and the time and computing power required to perform the reconstruction. Direct analytical reconstruction such as *filtered back-projection* (FBP) and the *direct Fourier method* (DFM) have remained attractive due to their relatively light computational burden, as well as their long history of use in medical CT. However these methods can be highly susceptible to non-idealities in the measured data, which may produce significant image distortion. One key example is the pronounced effect

scanned metal can have on the reconstruction, which commonly shows up as streaks in the image.

Recently, there has been growing interest in the use of model-based reconstruction techniques in CT security systems [9–11]. The potential to produce high-quality reconstructions is facilitated by their ability to incorporate knowledge of the physical and statistical properties of both the scanner measurements and the targets. The model-based framework also allows a great deal of flexibility to tune the algorithm to a particular scanner and application. In helical-scan CT systems for example, the modeling explicitly accounts for the trajectory of the X-ray source, the cone angle of the detector array, and the detector point-spread function [12]. The forward model also accounts for measurement degradation due to photon quantum noise and electronic noise in the detectors [13]. Incorporating this information directly into the reconstruction allows less reliance on pre-correction and reformatting (interpolation) of the data to a standard set of tomographic measurements, which is often required for direct reconstruction methods.

One established class of model-based techniques includes a regularization of the solution through a *Markov random field* (MRF) prior model, which describes the conditional distribution of a voxel given its neighbors. Such methods formulate the reconstruction as a maximization of the *posterior* distribution (of the image, given the measurements), or a *MAP* estimate. The optimization usually requires an iterative strategy such as *iterative coordinate descent* (ICD) [14], or gradient based methods.

The choice of prior model has a strong influence on the characteristics of the solution. A Gaussian MRF (GMRF) prior, provides for fast convergence but tends to over-regularize the solution. The *generalized Gaussian MRF* (GGMRF) [15] provides for noise suppression while allowing for better edge preservation in the image. A further generalized family is the *Q-generalized Gaussian MRF* (QGGMRF) [12] which is highly controllable while providing for fast convergence [16].

This thesis investigates the application of model-based iterative reconstruction (MBIR) methods for X-ray CT systems in the security context. This reconstruction

approach is demonstrated to produce high quality images that are much less susceptible to image distortion compared to direct reconstructions using direct Fourier methods or filtered back projection. Presented in Chapter 3 is a mapping of a fully 3D MBIR algorithm to a multi-slice helical-scan CT system certified for baggage screening. Model enhancements for the system are discussed, and MBIR reconstructions are presented alongside direct reconstructions for a set of real baggage scans. Chapter 4 investigates the performance of model-based methods for reconstruction on sparse-view angle security CT systems, and Chapter 5 extends the sparse-view context to evaluate two proposed prior models that adapt to the image conditions. Overall, advances in the forward and prior modeling are demonstrated to improve several aspects of the reconstructed images, including resolution, artifact suppression, and CT accuracy.

2. BACKGROUND

2.1 Model-Based Iterative Reconstruction

Model-based iterative reconstruction (MBIR) is a statistical framework in which the image, x , and the projection measurements, y , are random vectors, and the reconstructed image is computed as the maximum *a posteriori* (MAP) estimate,

$$\hat{x} = \underset{x \geq 0}{\operatorname{argmin}} \{ -\log p(y|x) - \log p(x) \} . \quad (2.1)$$

The likelihood function $p(y|x)$ contains the characterization of the measurement data, including the system geometry, projection model, and measurement uncertainty. The prior distribution, $p(x)$, of the image plays the important role of regularizing the solution, which mitigates the effects of having a data set that is noisy, sparse, or that contains aberrations from the assumed data model.

The measured photon count, λ_i , corresponding to projection i can be modeled as a Poisson random variable with mean $\bar{\lambda}_i = \bar{\lambda}_{T,i} e^{-A_{i*}x}$, where $\bar{\lambda}_{T,i}$ is the source photon rate for projection i , and A_{i*} is the i^{th} row of the projection matrix. From the Poisson model, a Taylor expansion can be used to approximate the log likelihood term by the second degree polynomial in x [14],

$$\log p(y|x) \approx -\frac{1}{2}(y - Ax)^T D(y - Ax) + g(y) \quad (2.2)$$

where D is a diagonal weighting matrix with entries $D_{ii} = \lambda_i$, and $g(y)$ combines terms that are constant with respect to x . Note in this form, the photon count, λ_i , acts as a weighting coefficient for the cost term, $(y_i - A_{i*}x)^2$, associated with projection i . This has a simple interpretation that smaller photon counts are less reliable measurements, and hence are weighted less in the cost function. An advancement of this result is the weighting $D_{ii} = \lambda_i^2/(\lambda_i + \sigma_i^2)$, which accounts for both photon statistics and additive electronic measurement noise [13].

A common class of priors for image modeling are Markov random fields (MRF) [17–22], which model pixel interactions on a local scale. The density of an MRF can generally be written as a *Gibbs distribution*,

$$p(x) = \frac{1}{z} \exp \left\{ - \sum_{c \in \mathcal{C}} V_c(x_c) \right\} \quad (2.3)$$

where \mathcal{C} is the set of all image cliques, x_c is a subset of values from x in clique c , and z is a normalizing constant called the partition function. The present work considers only pairwise cliques, and a density that can be written more simply as,

$$p(x) = \frac{1}{z} \exp \left\{ - \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \right\} \quad (2.4)$$

where $\rho(\cdot)$ is a positive and symmetric *potential function*. The coefficients $b_{s,r}$ are typically used to adjust the penalty for different geometrical distances between neighbors. For example, if $d(s, r)$ is the Euclidean distance between neighboring voxels s and r , the weighting could be defined as,

$$b_{s,r} = \frac{C_{s,r}}{d(s, r)} \quad (2.5)$$

where the term, $C_{s,r}$, is chosen so that the coefficients associated with a given voxel approximately sum to one, $\sum_{r \in \partial s} b(s, r) \approx 1$.

A subtle but important detail to note is that the coefficients $b_{s,r}$ need to have the symmetry $b_{s,r} = b_{r,s}$ because the points s and r are unordered in the Gibbs distribution of Equation (2.4). The weights in Equation (2.5) satisfy this symmetry if $C_{s,r} = C_{r,s}$. On a periodic lattice, this is true with the simple normalizer

$$C_{s,r} = \left[\sum_{u \in \partial s} \frac{1}{d(s, u)} \right]^{-1} \quad (2.6)$$

which will be constant with respect to both s and r . More generally, one could first define a position-dependent weight for each point

$$C_u = \left[\sum_{w \in \partial u} \frac{1}{d(u, w)} \right]^{-1} \quad (2.7)$$

and then a symmetric clique coefficient,

$$b_{s,r} = \frac{C_s + C_r}{2d(s,r)} . \quad (2.8)$$

Substituting in the expressions for the log likelihood and log prior, the reconstruction is obtained by solving the following optimization,

$$\hat{x} = \underset{x \geq 0}{\operatorname{argmin}} \left\{ \frac{1}{2} \|y - Ax\|_D^2 + \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \right\} . \quad (2.9)$$

In the present work, the solution, \hat{x} , is computed by iterative coordinate descent (ICD) which solves the global optimization by a series of 1D minimizations with respect to each voxel [9, 14, 23].

2.2 Prior Modeling

Priors considered in this thesis fall within the family of potential functions defined by the *Q-generalized Gaussian Markov random field*, or QGGMRF [12], which has the form,

$$\rho(\Delta) = c^q \frac{|\Delta/c|^p}{1 + |\Delta/c|^{p-q}} . \quad (2.10)$$

This is a further generalization of a family of edge-preserving priors called the *generalized Gaussian MRF* [15]. The shape of the QGGMRF potential function is controlled by the parameters p , q , and c , with the approximate behavior,

$$\rho(\Delta) \approx \begin{cases} c^{q-p} |\Delta|^p & \text{for } |\Delta| \ll c \\ |\Delta|^q & \text{for } |\Delta| \gg c \end{cases} . \quad (2.11)$$

Of particular interest are the cases where $1 \leq q \leq p$, which ensures convexity of the potential function. Common values are $p = 2$ (quadratic near zero), and q close to 1, which is a generalization of a Huber prior [24]. This provides for both edge preservation and fast convergence properties [16]. Another special case ($p = q = 1$), is the absolute value prior, $\rho(\Delta) = |\Delta|$, which is closely related to *total variation* (TV) regularization [25, 26], often used in imaging problems.

2.3 Substitute Function ICD for the QGGMRF Prior

The cost function for the MAP estimation, using the approximated Poisson log likelihood and MRF prior, is given by the following,

$$f(x) = \frac{1}{2} \|y - Ax\|_D^2 + \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \quad (2.12)$$

where y is the vector of projection measurements, A is a linear projection operator, and ρ is the potential function associated with the prior model.

Iterative coordinate descent (ICD) solves the optimization by iteratively performing 1D minimizations with respect to each variable. Each update can be considered a line search $x \leftarrow x + \alpha d$, where d is a direction vector and α is a scalar step size. An ICD update of pixel x_s would correspond to a direction vector, d , having a “1” in position s and “0”s elsewhere. The key is to find the step size, α , for each update that minimizes the cost function along that particular direction, d . The update for pixel x_s results in the following new cost value.

$$f(x + \alpha d) = \frac{1}{2} \|y - Ax - \alpha Ad\|_D^2 + \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s + \alpha d_s - x_r - \alpha d_r) \quad (2.13)$$

$$= \frac{1}{2} \|y - Ax - \alpha A_{*,s}\|_D^2 + \sum_{r \in \partial s} b_{s,r} \rho(x_s - x_r + \alpha) \quad (2.14)$$

The optimal step size is then,

$$\alpha^* = \underset{\alpha \geq -x_s}{\operatorname{argmin}} f(x + \alpha d) . \quad (2.15)$$

Note that a positivity constraint has been specified for the updated pixel value. This is a simple 1D optimization over α but for most interesting priors, α^* cannot be written in closed form. One possible approach is to compute the derivative of f analytically and then numerically search for the root.

Another generally faster approach is *substitute function minimization*, in which a simpler function of α , commonly a quadratic, is determined that satisfies a couple key properties, and is much easier to optimize. The following properties guarantee

that as long as the update decreases the substitute function, q , it also decreases the cost function, f :

$$q(x; x') \geq f(x) \quad (2.16)$$

$$q(x'; x') = f(x') . \quad (2.17)$$

Here $q(x; x')$ is explicitly shown to depend on the image values, x' , before the update. These conditions are that the substitute function must be an upper bound to f , and it must equal f at the current image state x' .

In the optimization of (2.14) over α , a quadratic substitute function can be assigned to each term in the summation. The data term is quadratic to begin with, so the substitute function for this term is exact,

$$\frac{1}{2} \|y - Ax - \alpha A_{*,s}\|_D^2 \quad (2.18)$$

$$= \frac{1}{2} (y - Ax - \alpha A_{*,s})^t D (y - Ax - \alpha A_{*,s}) \quad (2.19)$$

$$= \frac{1}{2} (y - Ax)^t D (y - Ax) - (y - Ax)^t D (\alpha A_{*,s}) + \frac{1}{2} \alpha^2 A_{*,s}^t D A_{*,s} \quad (2.20)$$

$$= \text{const.} + \theta_1 \alpha + \frac{1}{2} \theta_2 \alpha^2 \quad (2.21)$$

where

$$\theta_1 = -(y - Ax)^t D A_{*,s} \quad (2.22)$$

$$\theta_2 = A_{*,s}^t D A_{*,s} . \quad (2.23)$$

For the terms associated with the prior, one can use a quadratic substitute function with the so called *symmetric bound* condition to ensure properties (2.16) and (2.17) hold. If we define $\Delta = x_s - x_r$, and $\Delta' = x'_s - x'_r$, the substitute function for $\rho(\Delta)$ is given by,

$$\rho(\Delta; \Delta') = \frac{\rho'(\Delta')}{2\Delta'} \Delta^2 . \quad (2.24)$$

The potential function for the QGGMRF prior is given by,

$$\rho(\Delta) = c^q \frac{|\Delta/c|^p}{1 + |\Delta/c|^{p-q}} \quad (2.25)$$

with derivative,

$$\rho'(\Delta) = \frac{p|\Delta|^{p-1}(1 + |\Delta/c|^{p-q}) - |\Delta|^p|\Delta|^{p-q-1}(1/c)^{p-q}(p-q)}{c^{p-q}(1 + |\Delta/c|^{p-q})^2} \text{sgn}(\Delta) \quad (2.26)$$

$$= \frac{c^{q-p}|\Delta|^{p-1}}{1 + |\Delta/c|^{p-q}} \left(p - \frac{(p-q)|\Delta/c|^{p-q}}{1 + |\Delta/c|^{p-q}} \right) \text{sgn}(\Delta) \quad (2.27)$$

which produces the substitute function,

$$\rho(\Delta; \Delta') = \frac{\rho'(\Delta')}{2\Delta'} \Delta^2 \quad (2.28)$$

$$= \frac{c^{q-p}|\Delta'|^{p-2}}{2(1 + |\Delta'/c|^{p-q})} \left(p - \frac{(p-q)|\Delta'/c|^{p-q}}{1 + |\Delta'/c|^{p-q}} \right) \Delta^2. \quad (2.29)$$

Minimizing the substitute function produces the following

$$\alpha^* = \underset{\alpha \geq -x_s}{\text{argmin}} \left\{ \theta_1 \alpha + \frac{1}{2} \theta_2 \alpha^2 + \sum_{r \in \partial s} b_{s,r} \rho(x_s - x_r + \alpha; x_s - x_r) \right\} \quad (2.30)$$

$$= \underset{\alpha \geq -x_s}{\text{argmin}} \left\{ \theta_1 \alpha + \frac{1}{2} \theta_2 \alpha^2 + \sum_{r \in \partial s} b_{s,r} w(x_s - x_r) (x_s - x_r + \alpha)^2 \right\} \quad (2.31)$$

$$= \max \left\{ \frac{-\theta_1 - 2 \sum_{r \in \partial s} b_{s,r} w(x_s - x_r) (x_s - x_r)}{\theta_2 + 2 \sum_{r \in \partial s} b_{s,r} w(x_s - x_r)}, -x_s \right\} \quad (2.32)$$

where

$$w(\Delta) = \frac{c^{q-p}|\Delta|^{p-2}}{2(1 + |\Delta/c|^{p-q})} \left(p - \frac{(p-q)|\Delta/c|^{p-q}}{1 + |\Delta/c|^{p-q}} \right). \quad (2.33)$$

Equation (2.32) gives a closed form solution to minimize the substitute function. A single substitute function update of x_s may not reach the minimum of the true cost, so multiple local updates of x_s may be performed to approach the true minimum. This can be computationally advantageous because θ_2 can be reused, and only a simple update of θ_1 is needed. The new value of θ_1 after a single update of x_s is the following.

$$\theta_1^{(k+1)} = -(y - A(x + \alpha^* d))^t D A_{*,s} \quad (2.34)$$

$$= -(y - Ax - \alpha^* A_{*,s})^t D A_{*,s} \quad (2.35)$$

$$= -(y - Ax)^t D A_{*,s} + \alpha^* A_{*,s}^t D A_{*,s} \quad (2.36)$$

$$= \theta_1^{(k)} + \alpha^* \theta_2 \quad (2.37)$$

ICD algorithm using multiple-update majorization of QGGMRF prior:

```

Initialize  $e \leftarrow y - Ax$ 
For  $K$  iterations {
  For each pixel  $s \in S$  {
     $\theta_1 \leftarrow -e^t D A_{*,s}$ 
     $\theta_2 \leftarrow A_{*,s}^t D A_{*,s}$ 
     $x_{old} \leftarrow x_s$ 
    For  $N_{sub}$  iterations {
       $\alpha^* \leftarrow \max \left\{ \frac{-\theta_1 - 2 \sum_{r \in \partial s} b_{s,r} w(x_s - x_r) (x_s - x_r)}{\theta_2 + 2 \sum_{r \in \partial s} b_{s,r} w(x_s - x_r)}, -x_s \right\}$ 
      where  $w(\Delta) = \frac{c^{q-p} |\Delta|^{p-2}}{2(1+|\Delta/c|^{p-q})} \left( p - \frac{(p-q) |\Delta/c|^{p-q}}{1+|\Delta/c|^{p-q}} \right)$ 
       $x_s \leftarrow x_s + \alpha^*$ 
       $\theta_1 \leftarrow \theta_1 + \alpha^* \theta_2$ 
    }
    if ( overrelaxation_flag ) {
       $x_s \leftarrow x_{old} + \text{overrelax\_factor} * (x_s - x_{old})$ 
    }
     $e \leftarrow e - A_{*,s} (x_s - x_{old})$ 
  }
}

```

Fig. 2.1. Algorithm for substitute function ICD optimization of a cost containing the QGGMRF prior. The algorithm allows multiple local updates which can speed convergence (N_{sub} of 2 or 3 is typically sufficient). Also included is an optional over-relaxation step, which often speeds convergence in earlier iterations, however this may prevent full convergence to the global minimum. Two approaches are to (1) unset the flag after a certain number of iterations, or (2) step down the over-relaxation factor toward 1.0 as the iterations increase.

3. 3D MBIR FOR A HELICAL-SCAN CT BAGGAGE SCANNER

X-ray computed tomography (CT) currently has widespread application in air travel security systems for the purpose of baggage screening. This chapter presents an implementation of fully 3D MBIR reconstruction for a multi-slice helical CT system that is certified for security screening. Innovations in the data model are introduced to enhance image quality for the security application. It is first demonstrated that the traditional MBIR noise weighting can substantially reduce metal artifacts, but under certain conditions can also contribute to irregular textures in uniform materials. A novel weighting function is then introduced that adaptively combines a traditional weighting with a power-law scaling, resulting in reconstructions with greatly improved texture and reduced metal streaking. Substantial improvements are also demonstrated from detector afterglow correction and calibration for fan beam offset. The image improvements afforded by MBIR can provide for better operator experience and potentially enable enhanced performance of automatic threat detection (ATD) software.

For assessment the MBIR reconstructions are compared to a native set of reconstructions supplied by the vendor of the CT system. The model based reconstructions using the adaptive noise weighting demonstrate higher resolution and lower noise images, with greater suppression of metal-induced artifacts [10].

3.1 Multi-slice Helical Scan CT

The CT system of interest in this study is a third generation cone beam CT scanner in which the X-ray source and detectors rotate about a fixed axis, and the scanned objects are continuously fed through on a conveyor belt as measurements

are acquired. It is common to define a coordinate system as fixed with respect to the image volume, with the source and detectors translating orthogonally to the rotation plane, producing a helical trajectory. A multi-slice system has a detector array consisting of many rows (slices) and typically hundreds of columns of detector cells arranged in an arc, providing a sufficient number of projections to recover a 3D image of the scanned target.

The data model employed here includes a linearized forward projection model in which the object attenuation image, $x \in \mathbb{R}^M$, and the projections, $y \in \mathbb{R}^N$, in the absence of noise are related by a sparse matrix operator A ,

$$y = Ax . \tag{3.1}$$

The matrix coefficient A_{ij} represents the contribution of voxel j in forming projection i . In this study the elements of A are calculated using a distance-driven projector [12, 27].

3.2 Reconstruction Parameters

For the 3D reconstructions in this study, the cliques are formed from a 26-point neighborhood, consisting of the nearest neighbors surrounding a given voxel in all three dimensions. The prior model is a Markov random field (MRF) with an absolute value potential function, $\rho(\Delta) = |\Delta|$, which is closely related to *total variation* (TV) regularizer [25, 26], often used in imaging problems. The reconstruction is computed using a parallelized implementation of iterative coordinate descent (ICD) described in Chapter 2.

3.3 Detector Afterglow

X-ray scintillation detectors exhibit an *afterglow* property, which is a residual signal in the scintillation crystal that remains after the incident X-rays are removed. This signal decays exponentially with multiple time constants associated with different

physical characteristics of the crystal. In a continuous scan mode, as occurs in a helical system, the afterglow effect results in a smoothing of the measurements since each given measurement is affected by the beam's position over its recent history. If the sampling rate of the system is on the order of the afterglow time constants or greater, this smoothing can be deconvolved to compensate for the afterglow effect.

In [28], the authors characterize the HiLightTM detector as having an impulse response,

$$h(t) = \sum_{n=1}^N \frac{\alpha_n}{\tau_n} e^{-\frac{t}{\tau_n}} u(t) \quad (3.2)$$

where α_n specifies the relative strength of the decay components (with time constants τ_n) of the detector. These are grouped as *primary speed*, with τ_n on the order of 1 ms, and *afterglow* components, with time constants on the order of hundreds of milliseconds.

If an input photon flux, $w(t)$, is applied to produce a detector output $x(t)$, then sampled at $t = k\Delta t$,

$$x_k = x(k\Delta t) = \int_{-\infty}^{k\Delta t} w(t') h(k\Delta t - t') dt' \quad (3.3)$$

$$= \sum_{j=-\infty}^k \int_{(j-1)\Delta t}^{j\Delta t} w(t') h(k\Delta t - t') dt' \quad (3.4)$$

$$\approx \sum_{j=-\infty}^k w(j\Delta t) \int_{(j-1)\Delta t}^{j\Delta t} h(k\Delta t - t') dt' \quad (3.5)$$

$$= \sum_{j=-\infty}^k w_j \sum_{n=1}^N \alpha_n (1 - e^{-\frac{\Delta t}{\tau_n}}) e^{-(k-j)\frac{\Delta t}{\tau_n}} \quad (3.6)$$

$$= w_k * h_k \quad (3.7)$$

where

$$h_k = \sum_{n=1}^N \beta_n e^{-k\frac{\Delta t}{\tau_n}} u(k), \text{ and} \quad (3.8)$$

$$\beta_n = \alpha_n (1 - e^{-\frac{\Delta t}{\tau_n}}) \quad (3.9)$$

The authors in [28] specify the deconvolution by the following recursion. Here, y_k is ideally the recovery of the input signal w_k , given measurements x_k .

$$y_k = \frac{x_k - \sum_{n=1}^N \beta_n \left\{ e^{-\frac{\Delta t}{\tau_n}} \left[y_{k-1} + e^{-\frac{\Delta t}{\tau_n}} \left[y_{k-2} + \cdots + e^{-\frac{\Delta t}{\tau_n}} \left(y_2 + e^{-\frac{\Delta t}{\tau_n}} y_1 \right) \right] \right] \right\}}{\sum_{n=1}^N \beta_n} \quad (3.10)$$

Or,

$$y_k = \frac{x_k - \sum_{n=1}^N \beta_n e^{-\frac{\Delta t}{\tau_n}} S_{nk}}{\sum_{n=1}^N \beta_n} \quad (3.11)$$

where,

$$S_{nk} = y_{k-1} + e^{-\frac{\Delta t}{\tau_n}} S_{n(k-1)} \quad (3.12)$$

$$S_{n1} = 0. \quad (3.13)$$

It is not immediately obvious how the above correction can be implemented as a finite difference equation. Note the afterglow impulse response $h(k)$ can be written as a sum of simple recursive filters,

$$h(k) = \sum_{n=1}^N h_n(k) \quad (3.14)$$

where

$$h_n(k) = \beta_n e^{-k \frac{\Delta t}{\tau_n}} u(k). \quad (3.15)$$

This yields a transfer function,

$$H(z) = X(z)/W(z) = \sum_{n=1}^N \frac{\beta_n}{1 - e^{-\frac{\Delta t}{\tau_n}} z^{-1}} \quad (3.16)$$


$$= \frac{b_0 + b_1 z^{-1} + b_2 z^{-2} + \cdots + b_{N-1} z^{-(N-1)}}{1 + a_1 z^{-1} + a_2 z^{-2} + \cdots + a_N z^{-N}} \quad (3.17)$$

which implies the difference equation,


$$x_k + a_1 x_{k-1} + \cdots + a_N x_{k-N} = b_0 w_k + b_1 w_{k-1} + \cdots + b_{N-1} w_{k-(N-1)}. \quad (3.18)$$

So the afterglow correction can be implemented as a recursive difference equation of the form

$$y_k = \frac{1}{b_0} x_k + \frac{a_1}{b_0} x_{k-1} + \cdots + \frac{a_N}{b_0} x_{k-N} - \frac{b_1}{b_0} y_{k-1} - \cdots - \frac{b_{N-1}}{b_0} y_{k-(N-1)}. \quad (3.19)$$




(a) Reconstruction from raw measurements



(b) Reconstruction from afterglow corrected measurements

Fig. 3.1. Illustration of the effect of afterglow correction on MBIR reconstructions.

Detector afterglow correction was performed using the above filter on the raw scanner measurements. Figure 3.1 illustrates the effect of afterglow correction on MBIR reconstructions. The example shows an axial slice with significantly improved resolution as a result of the correction.



(a) Original DFM reconstruction



(b) MBIR accounting for a fan offset

Fig. 3.2. Accounting for a small displacement in the relative mounting position of the detector array. The horizontal axis is in the z -direction (perpendicular to the plane of gantry rotation).

3.4 Fan Angle Offset

Image reconstruction algorithms, whether implicitly or explicitly, assume an exact geometry of the system. A slight displacement in the mounting position of the detector array relative to the source, even within manufacturing tolerances, was found to produce a clear periodic displacement artifact in the reconstructed target. Figures 3.2(a) and 3.2(b) illustrate the effect of accounting for a small offset in the fan angle of the detector array.

3.5 Data Weighting Matrix

From the form of the likelihood in Equation (2.2), the diagonal entries D_{ii} take the role of the approximate inverse variance of the projections y_i . This is true for $D_{ii} = \lambda_i$ when considering photon noise alone, and is the principle behind the result $D_{ii} = \lambda_i^2/(\lambda_i + \sigma_i^2)$ accounting for both photon and electronic noise. The inverse variance is a logical choice for the weighting of the sinogram entries, y_i , in contributing to the solution of Equation (2.9). However there is motivation for further exploration of these weights. For one, that D_{ii} represents an inverse variance presupposes that the projection y_i varies about a mean $A_{i*}x$. Inaccuracy in the linearized and discretized projection model, as well as other biasing influences (e.g. beam-hardening, non-linear partial volume effects, scatter) can also affect the reliability of each data term in forming the solution.

Experimentally, an advantage was found in generalizing the weighting D_{ii} as a function of the counts λ_i . One in particular is the family of power law functions of the form

$$D_{ii} = (\lambda_i/\lambda_{T,i})^r \quad (3.20)$$

where λ_i is the target scan count, $\lambda_{T,i}$ is the air scan count (a separate scan with no target present), and $0 \leq r \leq 1$. The case $r = 1$ is equivalent to the original Poisson-induced weighting, and the effect of decreasing r from 1 is primarily to increase the relative weights of the lower-count measurements.

Figures 3.3(a) and 3.3(b) show a slice from a DFM reconstruction and an MBIR reconstruction, using $D_{ii} = \lambda_i$ and $\rho(\Delta) = |\Delta|$, of a large uniform-density cylinder constituted of acetal copolymer. The irregular texture in the MBIR result is a behavior that was observed when using the traditional weightings in reconstructing relatively large objects that produce highly attenuated X-ray measurements. Figure 3.3(c) shows the effect of reducing r to 0.5, resulting in a substantial improvement in the variance and texture within the uniform object.

A negative side-effect of increasing the relative weights of the lower-count measurements is to worsen distortion due to highly dense materials such as metal. This effect is illustrated in Figure 3.4(b,c) for reconstructions of a smaller acetal cylinder containing four tungsten pins. Note the non-uniform texture for $r = 1$ is not as apparent compared to the result in Figure 3.3(b) due the cylinder's smaller size.


To retain both improvement in bulk object reconstruction and resilience to metal artifacts, a mixture of these weightings was produced that depends on the detected presence of metal. To this end, first define a function I_i indicating the likely presence of metal along projection ray i .

$$I_i = \begin{cases} 1 & , \text{ if for some } j, \text{ both } A_{ij} > 0 \text{ and } x_j^{(0)} > T \\ 0 & , \text{ otherwise} \end{cases} \quad (3.21)$$

Here $x_j^{(0)}$ is the CT number for voxel j from an initial reconstruction, and T is a threshold CT value that indicates metal. The weights are then mixed according to

$$D_{ii} = I_i (\lambda_i / \lambda_{T,i}) + (1 - I_i) (\lambda_i / \lambda_{T,i})^{0.5} \quad (3.22)$$

which selects a power of 1 or 0.5 for a given projection ray depending on the metal indicator. Figures 3.3(d) and 3.4(d) show the corresponding results of using the mixed weighting. In the current study, a threshold of $T = 3000$ Hounsfield units (HU) (offset so that air=0 HU, water=1000 HU) was used.



(a) DFM




(b) MBIR, $r=1$ (c) MBIR, $r=1/2$ (d) MBIR, $r=\{1,1/2\}$ mixture

Fig. 3.3. Effect of power law weighting in the data matrix entries. Target is a large uniform-density acetal cylinder (diam=15cm).



(a) DFM




(b) MBIR, $r=1$ (c) MBIR, $r=1/2$ (d) MBIR, $r=\{1,1/2\}$ mixture


Fig. 3.4. Effect of power law weighting in the data matrix entries. Target is a uniform-density acetal cylinder (diam=8cm) with four tungsten pins inset.

3.6 Comparison of MBIR and DFM Reconstructions


For further evaluation, model-based reconstructions were computed for several luggage scans using the new data weighting in Equation (3.22), Figures 3.5–3.8 present select regions from these results alongside direct Fourier reconstructions (DFM) for comparison. Afterglow correction was performed on these scans prior to both MBIR and DFM reconstruction. Figures 3.6 and 3.7 demonstrate a dramatic reduction in metal-induced streaking while reproducing uniform materials with little noise. Figures 3.5 and 3.8 highlight improvements with respect to resolution and object discrimination.

Image improvements afforded by MBIR provide for both better operator experience and enhanced performance of automatic threat detection (ATD). Operator experience is beyond the scope of this study, but it is evident that a cleaner image will help operators with fast and effective clearing of benign luggage.

The qualitative impact of the improved image reconstructions on the vendor’s proprietary ATD algorithms was also evaluated. The cleaner objects in Figure 3.5 can improve segmentation, leading to better object identification (and consequent classification). Noise reduction (Figures 3.6 and 3.7) also helps with classification, as well as with distinguishing containers from their contents. Figure 3.8 shows how MBIR could allow ATD to separate adjoining objects for further analysis. Each image improvement leads to a reduction in false alarms due to incorrect segmentation, processing, or classification. In turn, the reduction in false alarms opens an opportunity for increased detection on challenging configurations or new threats.



(a) DFM



(b) MBIR

Fig. 3.5. Comparison of DFM and MBIR reconstructions for a bag containing clothing and a can of nuts.




Fig. 3.6. Comparison of DFM and MBIR reconstructions on a baggage scan. The region highlights a bulk uniform material susceptible to distortion from a nearby object.




Fig. 3.7. Comparison of DFM and MBIR reconstructions on a baggage scan. The region highlights a bulk uniform material susceptible to distortion from a nearby object.




Fig. 3.8. Comparison of DFM and MBIR reconstructions on a baggage scan. The region highlights the discrimination between two adjacent objects.

3.7 Summary

This chapter presented an implementation of fully 3D MBIR reconstruction for a multi-slice helical CT scanner used for security screening. A novel data weighting was introduced in the MBIR model to enhance image quality for security applications. In comparing to DFM, the MBIR reconstructions demonstrated substantial improvement in resolution, noise reduction, and reduction of metal-induced streaking. These image improvements can be used to reduce false alarms in existing ATD algorithms or to reduce the cost of new detection schemes.

4. SPARSE VIEW ANGLE ITERATIVE CT RECONSTRUCTION

Certain driving factors in the security application of X-ray CT have led to non-standard system geometries [29]. These factors can include the need for high throughput, high workload over the scanner lifetime, and the accommodation of bulky scan targets. Some configurations result in far fewer view angles than typically found in medical CT scanners. The advantage of sparse-view architectures is the potential to reduce the cost and complexity of the system, but a disadvantage is that sparse-view data can lead to structured artifacts in the reconstructions due to the ill-posed nature of the inversion [30]. Several reconstruction approaches have been explored for sparse angle tomography [31–33]

This chapter investigates the application of model-based iterative reconstruction (MBIR) for sparse-view scanning architectures, and illustrates the potential value of these methods in the security context. To evaluate the model-based reconstructions, realistic models of baggage are used, with randomly inserted simple simulated objects. It is demonstrated that using this approach, the model-based reconstructions can substantially reduce artifacts and improve the accuracy of the estimated CT numbers.

4.1 Projection Model

A linear forward projection model is assumed in which, in the noiseless case, the object density image, $x \in \mathbb{R}^M$, and the projections, $y \in \mathbb{R}^N$, are related by a sparse matrix operator A ,

$$y = Ax . \tag{4.1}$$

The matrix coefficient A_{ij} represents the contribution of voxel j in forming projection i . In a line-beam model, A_{ij} is calculated as the length of ray i that intersects voxel

j . In a wide-beam model which accounts for the fact that photons are collected over a detector *area*, the coefficient A_{ij} is computed as the inner product of the projection of voxel j onto the face of sensor i , with a detector efficiency kernel which is typically a simple indicator function. All results presented in this study use the wide-beam model.

4.1.1 Wide-beam projector

An important component of model-based reconstruction is the forward projection model. A common model assumes a point X-ray source and point detectors. In other words, each measurement is associated with a simple line integral through the image space,

$$y_i = \int_{r \in S_i} \mu(r) dr \quad (4.2)$$

where $\mu(r)$ is the linear attenuation at location r , and S_i is the set of points in projection ray i . The measured photon count, λ_i , would then be in theory (neglecting noise),

$$\lambda_i = \lambda_0 e^{-y_i} \quad (4.3)$$

where λ_0 is the count from an empty calibration scan.

When the image is discretized into voxels, the projection equation becomes

$$y_i = \sum_j A_{ij} x_j \quad , \text{ or in matrix form,} \quad (4.4)$$

$$Y = AX \quad (4.5)$$

where x_j is the attenuation within voxel j , and A_{ij} is the length that projection ray i intersects voxel j .

Real detectors collect photons over the area of the detector face. This means the contribution from a particular voxel is determined not simply by the intersecting line segment but rather the volume of the voxel that projects from the source onto the detector face.

The X-ray attenuation associated with a particular material is generally quantified in the context of a narrow beam passing through a uniform mass of the material. In this respect, the total photons, λ_i , collected by a detector can be considered as a weighted sum of narrow X-ray beams impinging on the detector face. Dropping the detector index i for the moment,

$$\lambda = \int_{z \in C} \lambda_s e^{-\int_{r \in S(z)} \mu(r) dr} w(z) dz . \quad (4.6)$$

Here, λ_s is the source photon rate, C is the set of points on the detector face, and $w(z)$ is the sensitivity of the detector at location z on its face. Notice the set of points in the projection ray, $S(z)$, is now written as a function of the point where the ray terminates at the detector.

After incorporating the calibration scan measurements, λ_0 , the projection, $y_i = \log \frac{\lambda_0}{\lambda}$, can be approximated by,

$$\log \frac{\lambda_0}{\lambda} \approx \int_{z \in C} \left\{ \int_{r \in S(z)} \Delta \mu(r) dr \right\} p(z) dz \quad (4.7)$$

where now $\Delta \mu(r) = \mu(r) - \mu_0(r)$ is the change in attenuation between the calibration and target scans, and $p(z)$ is $w(z)$ normalized so that it integrates to 1 (see Section 4.5). This is essentially a weighted sum of the line integrals (relative to calibration) falling on the detector face. Note this is a very different expression from (4.6) which is a summation in the photon domain.

The reason for this approximation is principally computational, as this is simply a weighted sum of projections which can be incorporated directly into the projection matrix, A . Therefore the projection operation is exactly the same computationally as the point detector model in Equation (4.4), where the weighting coefficients in the A matrix have been adjusted to account for the wide beam. The only additional overhead is in the computation of the A matrix, which for small reconstruction volumes (e.g. 2D) can be precomputed and stored.

4.2 Experimental Results

This section presents both qualitative and quantitative analyses on the accuracy of model-based reconstruction. To evaluate the reconstructions with respect to ground truth, projection data was simulated in this study by applying the forward projector described in Section 4.1 to the clean high-resolution scan in Figure 4.1. The original CT image itself contained reconstruction artifacts due the diversity in composition and morphology of items in the bag, which were effectively removed by masking out the background, and retaining only the visible objects. Since the principal interest is examining the effect of a reduced number of view angles, and the effect of clutter in the image, no photon noise was incorporated in this study. The full set of image and reconstruction parameters is summarized in Table 4.1. For comparison, reconstructions were also computed using *filtered back projection*, the filter having a ramp frequency response multiplied by a Hamming window, and a cutoff frequency of 0.8 times the Nyquist rate.

Table 4.1
Image and Reconstruction Parameters

Sinogram parameters	no. of angles	{64,32,16,8}
	no. of translations	800
	Δ_θ	180/no_angles deg
	detector width	1 mm
Image geometry	xdim	800 pixels
	ydim	800 pixels
	field of view	80 cm
QGMRF parameters	p	2.0
	q	1.0
	c	15.0 HU



Fig. 4.1. High resolution CT bag scan used as ground truth for the current experiment. Gray scale is in offset Hounsfield units (HU), where Air=0 HU. The original CT image contained reconstruction artifacts due the diversity in composition and morphology of items in the bag, which were removed by masking out the background (set to 0 HU), and retaining only the visible objects. A linear wide-beam projector was then used to generate sinogram data for the current analysis.

4.2.1 Effect of sparse view angles

Figure 4.2 shows the reconstructions after forward projecting the bag scan of Fig. 4.1 at a limited number of equally spaced view angles between 0 and 180 degrees. Illustrated is the effect of reducing the number of view angles on reconstructions by (1) filtered back projection (FBP), (2) MBIR using a Gaussian Markov random field (GMRF) prior, and (3) a Q-generalized Gaussian MRF prior (QGGMRF). The corresponding *root mean square error* (RMSE) from ground truth for each of these reconstructions is listed in Table 4.2. The RMSE was computed from only those voxels having a density greater than air in the ground truth image.

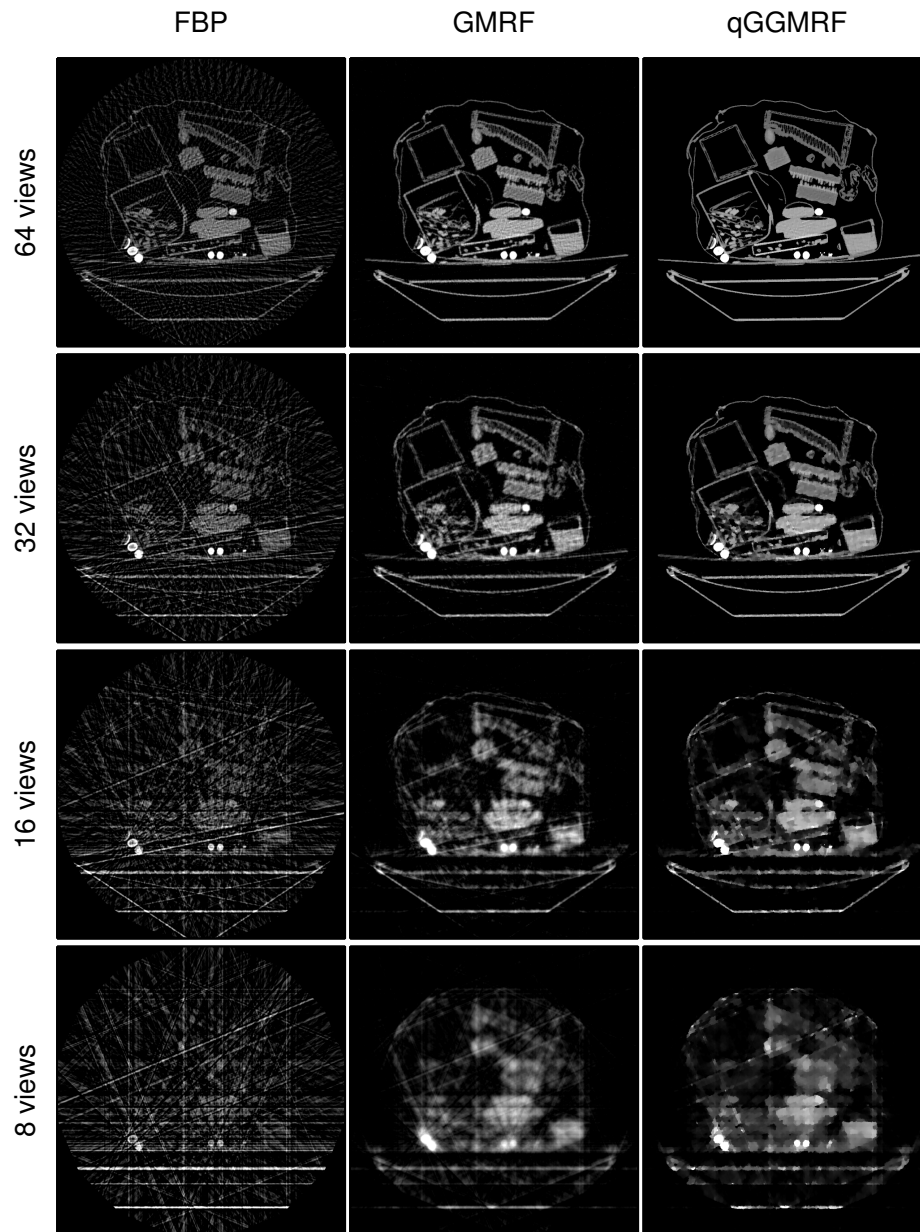


Fig. 4.2. Image reconstructions from sparse view angle projection data. Four parallel projection data sets are considered, containing 64, 32, 16, and 8 view angles, uniformly spaced between 0 and 180 degrees. Reconstructions include filtered back projection (FBP) and MBIR using a Gaussian Markov random field prior (GMRF), and a q-generalized Gaussian MRF prior (QGGMRF). The gray scale range for all results shown is $[0, 2000]$ HU, as in the ground truth image of Fig. 4.1.

Table 4.2
Root mean square error of reconstructions in Fig. 4.2. Units are offset
Hounsfield (air=0).

no. of views	FBP	GMRF	QGGMRF
64	481.0	237.8	112.8
32	628.4	361.1	277.1
16	746.2	498.9	453.8
8	854.4	607.1	598.5

4.2.2 Effect of clutter on CT number accuracy

To examine the effect of image “clutter” on the accuracy of CT number estimates, the ground truth image of Fig. 4.1 was modified in two respects. First, the contents of the bag scan were masked out to produce a *low clutter* scene (the original image was considered a *high clutter* scene). Second, a single round 1.7 cm diameter target of a known uniform CT value (1400 HU) was placed at some location inside the perimeter of the bag. Figures 4.3(a) and 4.3(b) illustrate this for the low and high clutter scenes, and 4.3(c) shows a close-up of the target.

Two experiments were performed. In the first, 32-view angle parallel projection data was produced from the images shown in Figures 4.3(a) and 4.3(b) (without the highlighting box). Reconstructions were computed using FBP, and MBIR using GMRF and QGGMRF priors. A close-up of the reconstructions around the target region are shown in Figure 4.4, which can be compared to the ground truth shown in Figure 4.3(c). Figure 4.5 shows the reconstructed CT numbers for voxels along the reference line shown in Figure 4.3(c).

The second experiment calculated the average accuracy of the reconstructed target voxels after placing the target at various locations in the bag. Specifically in each trial, (1) the synthetic target was placed at a random location inside the bag perimeter;




Fig. 4.3. Ground truth images for investigating the effect of clutter on CT reconstruction accuracy. The bag contents from the ground truth image in Fig. 4.1 have been masked out to create a *low clutter* scene in (a). For evaluation, a synthetic target of uniform value (1400 HU) has been added, as highlighted by the box near the center of the low and high clutter scenes. A close-up of the target in (c) also shows a reference line highlighting voxels that will be examined in the experiment.

(2) the image was forward projected to produce a 32-view angle sinogram; (3) the data was reconstructed; and (4) the target voxels' average deviation from the true value was calculated, as well as the root mean square of the deviations. Table 4.3 summarizes the results of this procedure averaged over 60 trials of random placement.




Fig. 4.4. Reconstructions zoomed to the target area indicated in Figs. 4.3(a) and 4.3(b). All results are from 32-view angle parallel projection data, the top row generated from the low clutter scene of Fig. 4.3(a), and the bottom row from the high clutter scene of Fig. 4.3(b).

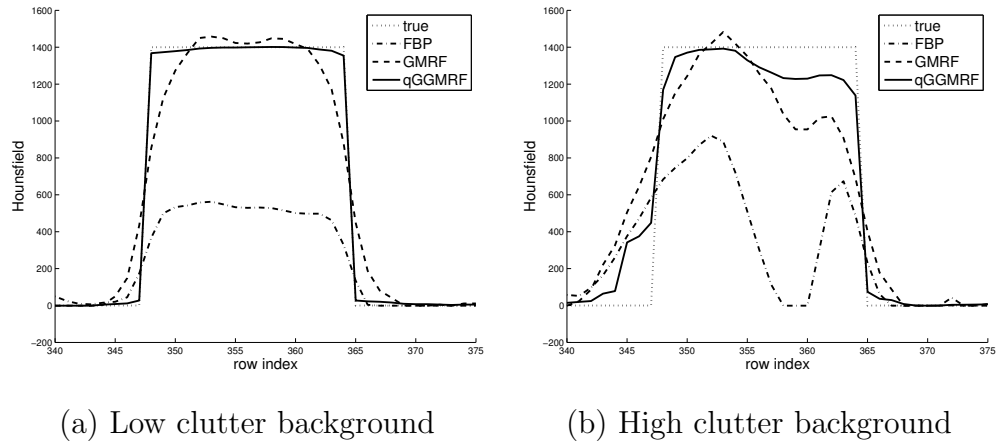


Fig. 4.5. CT values from sparse view reconstructions of Fig. 4.4 for voxels along a line through the center of the target region (See Fig. 4.3(c)). Also shown are the true voxel values from ground truth.

Table 4.3

Statistics for the reconstructed CT numbers for a randomly placed synthetic target. The values were computed by averaging over 60 trials of placing a round 1.7cm uniform target (1400 HU) at random locations within the bag perimeter in the ground truth images (See Fig. 4.3). The *Dev.* is the average deviation of reconstructed target voxels from the true value. Similarly, the *RMSE* is the root mean square deviation from the true target value. All values are in offset Hounsfield units (air=0).

	Low clutter		High clutter	
	<u>Dev.</u>	<u>RMSE</u>	<u>Dev.</u>	<u>RMSE</u>
FBP	-895.1	899.1	-647.8	702.7
GMRF	-157.2	280.4	-179.8	332.7
QGGMRF	-14.2	25.8	-87.3	209.2

4.3 Discussion

The most apparent advantage of model-based reconstruction from Figure 4.2 is a reduced susceptibility to streaking artifacts, whereas FBP quickly devolves into streaks as the number of views becomes small. While streaking patterns can be seen in dense regions of the MBIR-GMRF result at 64 views, the spatial extent of the streaks is much more localized than in FBP, and the regularization of the QGGMRF prior further reduces these dramatically.

In fact, for each data set in Figure 4.2 the QGGMRF prior produces a result with less structured error and a clearer edges than the GMRF prior. However, for extremely low view angles such as the 8-views case, the advantage of the QGGMRF over GMRF is minimal because the edge locations are not always accurate. These points are also reflected in the RMSE values listed in Table 4.2. In each case, the QGGMRF RMSE is smaller than the GMRF RMSE, with the difference becoming less significant as the number of views decreases.

Of note in Table 4.2 is the result that, in the mean square sense, MBIR-QGGMRF produces a more accurate reconstruction than FBP using only a quarter of the number of views. Specifically, the RMSE of QGGMRF at 16 views is smaller than that of FBP at 64 views, and QGGMRF at 8 views is smaller than FBP at 32 views. If this result generalizes, this is a particularly significant consideration since the number of views can have a direct correlation to system cost, scan time, and reconstruction time. Of course this marked difference in RMSE does not necessarily translate in the qualitative sense because visual inspection can somewhat compensate for the streaking in FBP. It should also be noted that one reason for the relatively high RMSE in the FBP reconstructions is a general underestimation of the CT numbers (which is apparent in Figure 4.2) partly due to image energy dispersal in the streaking. Presumably an appropriate image-dependent rescaling could be employed to provide a degree of compensation for this.

The 32-view synthetic target experiment of Figs. 4.4 and 4.5 reinforces several of the above observations. Namely, MBIR-QGGMRF produces a much more accurate reconstruction in terms of both CT numbers overall, and in terms of edge clarity. The low clutter scene results in effectively no visible streaking in any case, with QGGMRF producing very accurate CT numbers and very little blurring of the target boundary. In the high clutter scene, FBP fully splits the target into two disjoint objects, while the GMRF prior produces a recognizable object but with highly non-uniform CT numbers. The QGGMRF prior produces a significantly more uniform reconstruction of the target and reproduces the edges with notable accuracy by comparison.

The results of the random placement experiment summarized by Table 4.3 are a more general confirmation of the observations about the reconstructions in Fig. 4.4. Since the target position is allowed to vary, the results are not strongly dependent on any particularly strong streaking artifacts caused by the metallic objects in the image. Of note is the factor of 10 improvement in the accuracy in the low clutter scene, going from the GMRF to the QGGMRF prior. As observed in Fig. 4.5(a), this is due to the much more accurate edge reconstruction afforded by the QGGMRF model. Similar, while not as dramatic, improvements are produced for the high clutter scene.

4.4 Summary

This chapter presented the application of iterative model-based reconstruction on sparse view angle parallel projection data, generated from a typical bag scan. The MBIR reconstructions using two different prior models, a Gaussian Markov random field (GMRF) and a q-generalized Gaussian Markov random field (QGGMRF), were compared to a standard filtered back projection algorithm. Qualitative and quantitative measures demonstrated potentially great strengths in model-based reconstruction applied to transportation security, both in terms of reconstruction of form and in the CT number accuracy.

4.5 Supplemental Analysis of Wide-Beam Projector Model

First define the following quantities associated with a particular detector. For notational convenience, the detector index will be omitted here.

- λ_s : photon rate at detector in a vacuum (unoccluded)
- λ_0 : photon rate observed at detector in empty calibration scan
- λ : photon rate observed at detector during object scan
- $\mu_0(r)$: X-ray attenuation in calibration scan at location r
- $\mu(r)$: X-ray attenuation in object scan at location r
- $w(z)$: sensitivity of detector at location z on the detector face

At a given detector, the photon rates observed in the calibration and imaging scans are given by,

$$\lambda_0 = \int_{z \in C} \lambda_s e^{-\int_{r \in S(z)} \mu_0(r) dr} w(z) dz \quad (4.8)$$

$$\lambda = \int_{z \in C} \lambda_s e^{-\int_{r \in S(z)} \mu(r) dr} w(z) dz \quad (4.9)$$

where C is the set of points on the detector face, and $S(z)$ is the set of points on a line between the X-ray source and location z on the detector face. Manipulating equation (4.9),

$$\lambda = \int_{z \in C} e^{-\int_{r \in S(z)} \mu(r) - \mu_0(r) dr} \underbrace{\lambda_s e^{-\int_{r \in S(z)} \mu_0(r) dr} w(z)}_{u(z)} dz. \quad (4.10)$$

Now define a normalized version of $u(z)$,

$$p(z) = \frac{u(z)}{\int_{z' \in C} u(z') dz'} \quad (4.11)$$

so that $\int_{z \in C} p(z) dz = 1$. From the definition of $u(z)$ and λ_0 , we have $u(z) = \lambda_0 p(z)$, and

$$\lambda = \lambda_0 \int_{z \in C} e^{-\int_{r \in S(z)} \Delta \mu(r) dr} p(z) dz \quad (4.12)$$

where $\Delta\mu(r) \triangleq \mu(r) - \mu_0(r)$ is the change in attenuation relative to the calibration. Now we obtain the following result.

$$\log \frac{\lambda}{\lambda_0} = \log \int_{z \in C} e^{-\int_{r \in S(z)} \Delta\mu(r) dr} p(z) dz \quad (4.13)$$

$$\geq \int_{z \in C} \log \left\{ e^{-\int_{r \in S(z)} \Delta\mu(r) dr} \right\} p(z) dz \quad (4.14)$$

$$= \int_{z \in C} \left\{ - \int_{r \in S(z)} \Delta\mu(r) dr \right\} p(z) dz \quad (4.15)$$

The inequality in (4.14) is an application of Jensen's inequality since $\log(\cdot)$ is convex and $p(z)$ is a non-negative function that integrates to 1.

This leads to the result,

$$\log \frac{\lambda_0}{\lambda} \leq \int_{z \in C} \left\{ \int_{r \in S(z)} \Delta\mu(r) dr \right\} p(z) dz . \quad (4.16)$$

The left-hand side of (4.16) is used as an estimate of the projection through the target. This would be most accurately related to the target attenuation, μ , through equations (4.8) and (4.9). However the right-hand side of (4.16) is more convenient because it is simply a weighted sum of projections through the object which can be directly incorporated into the projection matrix in the discrete formulation. However the inequality in (4.16) means the reconstructed target attenuation $\Delta\mu$ will be underestimated as the MAP estimation penalizes the difference between the right and left-hand sides.

5. ADAPTIVE PRIORS FOR ITERATIVE IMAGE RECONSTRUCTION

The priors utilized in Chapters 3 and 4 could be described as *homogeneous* in that their characteristics do not change as a function of position in the image volume. There are multiple potential benefits to using non-homogeneous regularization. For example, the point spread function of an imaging system may be spatially dependent, so the regularization can be designed to achieve uniform resolution across the image [34].

From another standpoint, the imaged objects themselves, especially in the security application, are usually complex heterogeneous bodies with a wide dynamic range in density. The variability both within and between scans makes it difficult to characterize the scans with a “one-size-fits-all” prior. For this reason, there could be a strong potential to improve the reconstructions by adjusting the regularization given initial estimates of the underlying image.

This chapter explores adapting the prior model to suit the behavior of a given image. Two such techniques are proposed here, and both involve starting from an initial reconstruction, adjusting the weighting kernel in the prior according to the initial image estimate, and re-optimizing a posterior that includes the modified prior. The first is a *segmentation-adapted prior* in which a segmentation of an initial reconstruction steers the prior to increase the regularization within the interior of uniform-composition objects, and decrease near object boundaries. The second is a *bilateral adaptive prior* in which the clique weightings are determined similar to a bilateral filter kernel [35, 36]. Both approaches are demonstrated to smooth uniform objects while increasing accuracy around edges.

5.1 Homogeneous Priors

Recall the MBIR framework models the underlying attenuation image as a Gibbs distribution of the following form.

$$p(x) = \frac{1}{z} \exp \left\{ - \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \right\} \quad (5.1)$$

In the MAP cost function,

$$f(x) = \frac{1}{2} \|y - Ax\|_D^2 + \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \quad (5.2)$$

the potential function $\rho(x_s - x_r)$ effectively penalizes differences in value between neighboring voxels s and r . Typically, the weighting coefficients $b_{s,r}$ are only a function of the geometrical distance between neighbors. For example, using the expressions discussed in Section 2.1, one could first define the following weight for each point,

$$C_u = \left[\sum_{w \in \partial u} \frac{1}{d(u, w)} \right]^{-1} \quad (5.3)$$

and then a clique coefficient,

$$b_{s,r} = \frac{C_s + C_r}{2 d(s, r)} . \quad (5.4)$$

where $d(s, r)$ is the Euclidean distance between s and r . Assuming the reconstruction is performed on a lattice, the general characteristics of this prior would not change over the image volume, hence this prior could be described as *spatially homogeneous*.

5.2 Segmentation-Adapted Prior

This section introduces a type of non-homogeneous prior that adapts to the local image conditions, guided by a segmentation of an initial reconstruction. Depending on the expected or desired characteristics of the image, and the segmentation method used, there are several possible strategies for adapting the prior model. One scenario is the desire to more accurately reconstruct the CT values and morphology of uniform-composition objects. This could be achieved after segmenting objects of interest in an

initial reconstruction, and then, (1) increasing penalties for voxel differences within the objects, and (2) reducing penalties for voxel differences between pairs that fall across a region boundary. These penalty adjustments can be applied directly to the scaling coefficients $b_{s,r}$ in (5.2).

The method is described in Figure 5.1. Starting from some initial condition, $\hat{x}^{(0)}$, a MAP estimate is computed using a homogeneous prior. A segmentation of the result produces a label image, \tilde{x} , which is then used to adjust the prior coefficients $b_{s,r}$. Weights for voxel pairs in the interior of a segmented region are scaled up by $a_1 \geq 1$, and pairs across boundaries are scaled down by $0 \leq a_2 < 1$.

Model-based reconstruction using segmentation-adapted prior:

Initialize $\hat{x} \leftarrow \hat{x}^{(0)}$

For each voxel $u \in S$ {

$$C_u \leftarrow \left[\sum_{w \in \partial u} \frac{1}{d(u,w)} \right]^{-1} \quad // \ d(s,r) = \text{Euclidean distance between } s \text{ and } r$$

}

$$b_{s,r} \leftarrow \frac{C_s + C_r}{2d(s,r)}$$

$$\hat{x} \leftarrow \underset{x \geq 0}{\operatorname{argmin}} \left\{ \frac{1}{2} \|y - Ax\|_D^2 + \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \right\}$$

$\tilde{x} \leftarrow \text{SegmentationOfChoice}(\hat{x})$

$$b'_{s,r} \leftarrow \begin{cases} a_1 b_{s,r} & , \text{ if } \tilde{x}_s = \tilde{x}_r \neq 0 \\ a_2 b_{s,r} & , \text{ if } \tilde{x}_s \neq \tilde{x}_r \end{cases}$$

$$\hat{x} \leftarrow \underset{x \geq 0}{\operatorname{argmin}} \left\{ \frac{1}{2} \|y - Ax\|_D^2 + \sum_{\{s,r\} \in \mathcal{C}} b'_{s,r} \rho(x_s - x_r) \right\}$$

Fig. 5.1. Algorithm for model-based reconstruction using a segmentation-adapted prior. \tilde{x} is a label image from an initial reconstruction. (Here, a $\tilde{x}_s = 0$ label indicates a region not of interest.) The second optimization should use the first reconstruction as an initial condition to accelerate convergence. An ICD algorithm for the optimizations is described in detail in Chapter 2, including pseudocode in Figure 2.1.

5.3 Bilateral Adaptive Prior

Again recall the Gibbs distribution for the prior model.

$$p(x) = \frac{1}{z} \exp \left\{ - \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \right\} \quad (5.5)$$

In the *segmentation-adapted prior* described in the previous section, the clique weights, $b_{s,r}$, were determined from a segmentation of an initial homogeneous-prior reconstruction. While the experimental results show promise when the segmentation covers uniform regions, this approach has the drawback of being strongly dependent on the segmentation accuracy, particularly at the region boundaries.

Another technique proposed here is to base the coefficients on a *bilateral filter* kernel [35, 36], in which the clique weights are influenced by the local properties of a previous estimate of the image itself. Similar to the segmentation adapted prior, the goal here is to reduce the influence of neighboring voxels when they fall across an edge (large gradient), and increase the influence of neighbors within homogeneous regions. First define a smoothed version of a previous image estimate (which could be an initial reconstruction), $\tilde{x}_s = \sum_r x_r h_{s-r}$, where h is a Gaussian filter kernel. Then define the clique coefficients by the following,

$$b_{s,r} = C_{s,r} \exp \left(- \frac{(\tilde{x}_s - \tilde{x}_r)^2}{2\sigma_b^2} \right) \frac{1}{d(s,r)} . \quad (5.6)$$

The last term containing $d(s,r)$, the Euclidean distance between voxels s and r , is the homogeneous case described in Section 5.1. The new adaptive term containing the smoothed image \tilde{x} has the effect of decreasing the influence of voxel r on voxel s when the difference in their values is large. The parameter σ_b determines how strongly the neighbor differences influence the clique weight. Finally, the constant $C_{s,r}$ is a normalizer so that the coefficients associated with a given voxel sum to approximately one.

Note that because the cliques in Equation (5.5) are unordered voxel pairs, the coefficients must have the symmetry $b_{s,r} = b_{r,s}$ in order to define a consistent cost

function. Therefore the normalizing constants $C_{s,r}$ also need to have this symmetry. This can be achieved by first computing for each voxel, u ,

$$C_u = \left[\sum_{w \in \partial u} \exp \left(-\frac{(\tilde{x}_u - \tilde{x}_w)^2}{2\sigma_b^2} \right) \left(\frac{1}{d(u, w)} \right) \right]^{-1} . \quad (5.7)$$

Then a symmetric normalizing factor is computed as,

$$C_{s,r} = \frac{1}{2} (C_s + C_r) . \quad (5.8)$$

The overall method is described as pseudocode in Figure 5.2. Starting from some initial condition, a MAP estimate, $\hat{x}^{(0)}$, is computed using a homogeneous prior. A low-pass filtered version, \tilde{x} , of the initial estimate is then used to determine the bilateral prior weights for a subsequent reconstruction, $\hat{x}^{(1)}$. The cycle of redefining the weights and reconstructing can be repeated a number of times. Note that each optimization in the inner loop is convergent from the construction of the cost function and optimization method, but the sequence $\hat{x}^{(0)}$, $\hat{x}^{(1)}$, $\hat{x}^{(2)}$, \dots is not necessarily convergent. Note also that the calculation of the intermediate vector C_u allows the coefficients $b_{s,r}$ to be computed and symmetrized on the fly in the course of each optimization. The alternative would be to pre-compute the coefficients for all unique neighboring voxel pairs, which could require a large amount of memory.

Model-based reconstruction using bilateral adaptive prior:

Initialize \hat{x}

For each voxel $u \in S$ {

$$C_u \leftarrow \left[\sum_{w \in \partial u} \frac{1}{d(u,w)} \right]^{-1} \quad // \ d(s, r) = \text{Euclidean distance between } s \text{ and } r$$

}

$$b_{s,r} \leftarrow \frac{C_s + C_r}{2d(s,r)}$$

$$\hat{x}^{(0)} \leftarrow \underset{x \geq 0}{\operatorname{argmin}} \left\{ \frac{1}{2} \|y - Ax\|_D^2 + \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \right\}$$

$k \leftarrow 0$ // iteration variable; will readjust prior K times

While $k < K$ {

$$k \leftarrow k + 1$$

$$\tilde{x} \leftarrow \text{LowPassFilter}(\hat{x}^{(k-1)})$$

For each voxel $u \in S$ {

$$C_u \leftarrow \left[\sum_{w \in \partial u} \exp \left(-\frac{(\tilde{x}_u - \tilde{x}_w)^2}{2\sigma_b^2} \right) \frac{1}{d(u,w)} \right]^{-1} \quad // \text{ for coeff normalization}$$

}

$$b_{s,r} \leftarrow \frac{C_s + C_r}{2d(s,r)} \exp \left(-\frac{(\tilde{x}_s - \tilde{x}_r)^2}{2\sigma_b^2} \right)$$

$$\hat{x}^{(k)} \leftarrow \underset{x \geq 0}{\operatorname{argmin}} \left\{ \frac{1}{2} \|y - Ax\|_D^2 + \sum_{\{s,r\} \in \mathcal{C}} b_{s,r} \rho(x_s - x_r) \right\}$$

}

Fig. 5.2. Algorithm for model-based reconstruction using a bilateral adaptive prior. The process of redefining the weights and reconstructing is repeated a fixed number of times. In each stage, the MAP estimate should use the previous reconstruction as an initial condition to accelerate convergence. For the MAP estimates, a detailed ICD algorithm was covered in Chapter 2, including pseudocode in Figure 2.1).

5.4 Experimental Results

This section presents qualitative and quantitative experiments to assess the reconstructions using the adaptive priors introduced earlier in this chapter. To evaluate the reconstructions with respect to ground truth, sparse-view projection data (32 views) was simulated by applying a linear forward projector (see Sec. 4.1) to cargo container images having various levels of background “clutter”. Embedded within the container images is a synthetic target of uniform attenuation (1400 HU; water=1000 HU), shown in Figure 5.3. The embedded targets are the focus of the assessment, and the full background cargo images are not presented by request of the supplier. The full set of image and reconstruction parameters is summarized in Table 5.1.




Fig. 5.3. Targets used in simulated cargo containers for assessment of segmentation-adapted and bilateral prior reconstruction. The round target is 4 cm in diameter, the thin target is 4.4 mm wide, and both have uniform attenuation of 1400 HU.

Table 5.1
Image and Reconstruction Parameters

Sinogram parameters	no. of angles	32
	no. of translations	800
	Δ_θ	180/no angles deg
	detector width	1 mm
Image geometry	xdim	800 pixels
	ydim	800 pixels
	field of view	1.78 m
	voxel size	2.2 mm
QGGMRF parameters	p	2.0
	q	1.0
	c	15.0 HU

5.4.1 Results for Segmentation-adapted Prior

Reconstructions using the *segmentation-adapted prior* described in Section 5.2 were assessed using cargo images containing the target in Figure 5.3(a). Three container images were used having various levels of background clutter. Figures 5.4 and 5.5 show the result for the high-clutter conditions. A segmentation of an initial model-based reconstruction (homogeneous QGGMRF prior) is shown in 5.4(c). The final reconstruction after redefining the posterior is shown in 5.4(d). The boundary of the target has been reconstructed correctly in most places, except at a couple locations due to a mis-segmentation. The line profiles in Figure 5.5 illuminate the improvement not only the reduction in distortion, but also the improvement in the accuracy of the CT number estimates.

The less challenging cases in which the image contains less background objects are shown in Figures 5.6 through 5.9. In these cases, the line profiles demonstrate the improvement in CT number accuracy within the target objects, compared to the initial reconstructions. The medium clutter case again demonstrates a reduction in distortion, and the low clutter case demonstrates a significant reduction in bias of the CT number estimates.

Table 5.2 lists quantitative results for this experiment for different choices of edge and interior weighting in the modified prior. There is a general reduction in target bias with a decrease in the edge weight. In only the high clutter case, the target RMSE increases with a decrease in edge weight, likely due to the mis-segmentation in this case. Also in the high clutter results, the target RMSE generally decreases as the interior weight is increased due to a reduction in distortion.

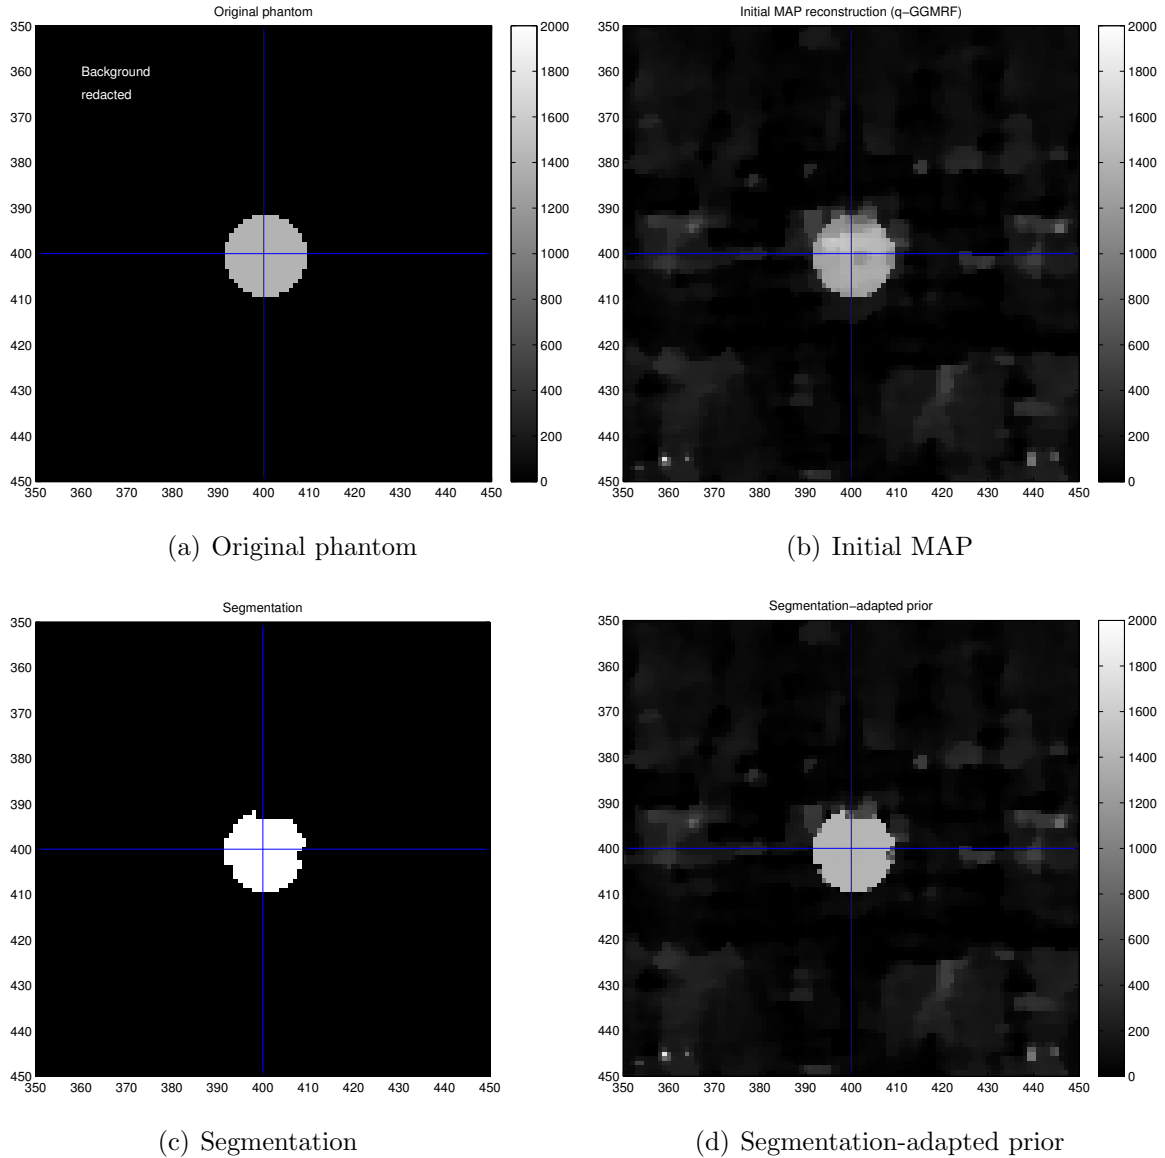




Fig. 5.4. Segmentation-adapted reconstruction for a target in *high clutter* level conditions. The target is 4 cm in diameter, with a 2.2 mm pixel size. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is a segmentation of the initial reconstruction, and (d) is the segmentation-adapted reconstruction (edge weight=0, interior weight=5). 1-D profiles along the indicated blue lines are provided in separate graphs.



(a) Vertical line profile



(b) Horizontal line profile

Fig. 5.5. Segmentation-adapted reconstruction for a 1400 HU target in *high clutter* level conditions. Vertical and horizontal pixel profiles through target.

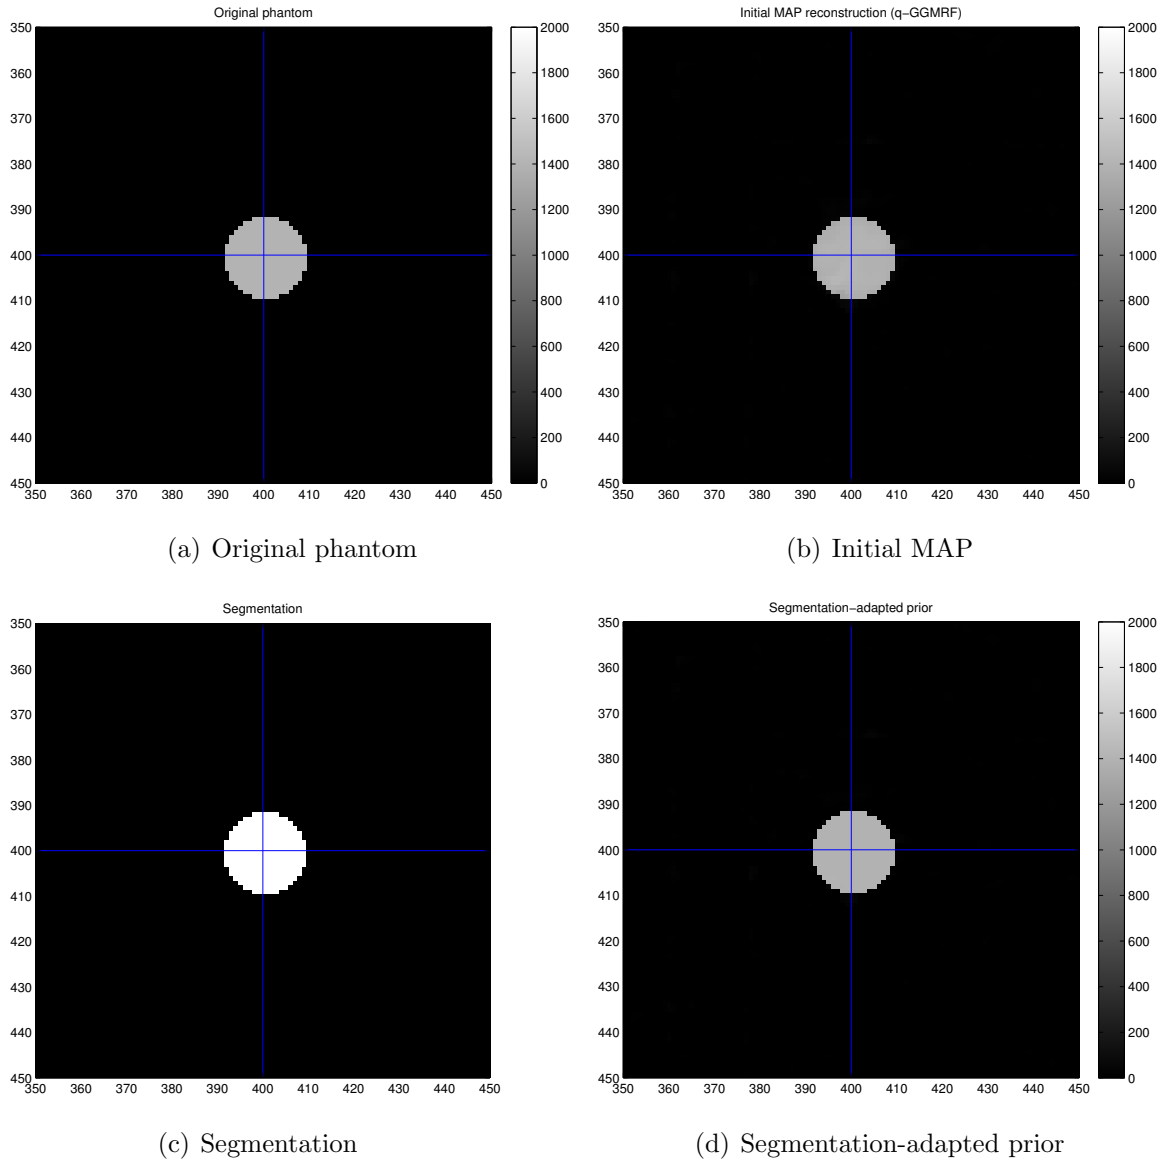

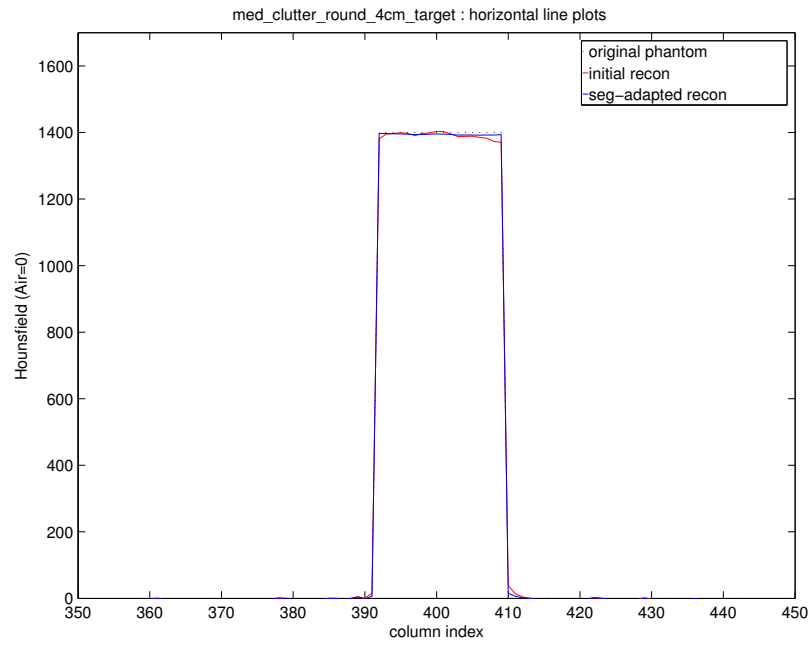


Fig. 5.6. Segmentation-adapted reconstruction for a target in *medium clutter* level conditions. The target is 4 cm in diameter, with a 2.2 mm pixel size. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is a segmentation of the initial reconstruction, and (d) is the segmentation-adapted reconstruction (edge weight=0, interior weight=5). 1-D profiles along the indicated blue lines are provided in separate graphs.



(a) Vertical line profile



(b) Horizontal line profile

Fig. 5.7. Segmentation-adapted reconstruction for a 1400 HU target in *medium clutter* level conditions. Vertical and horizontal pixel profiles through target.

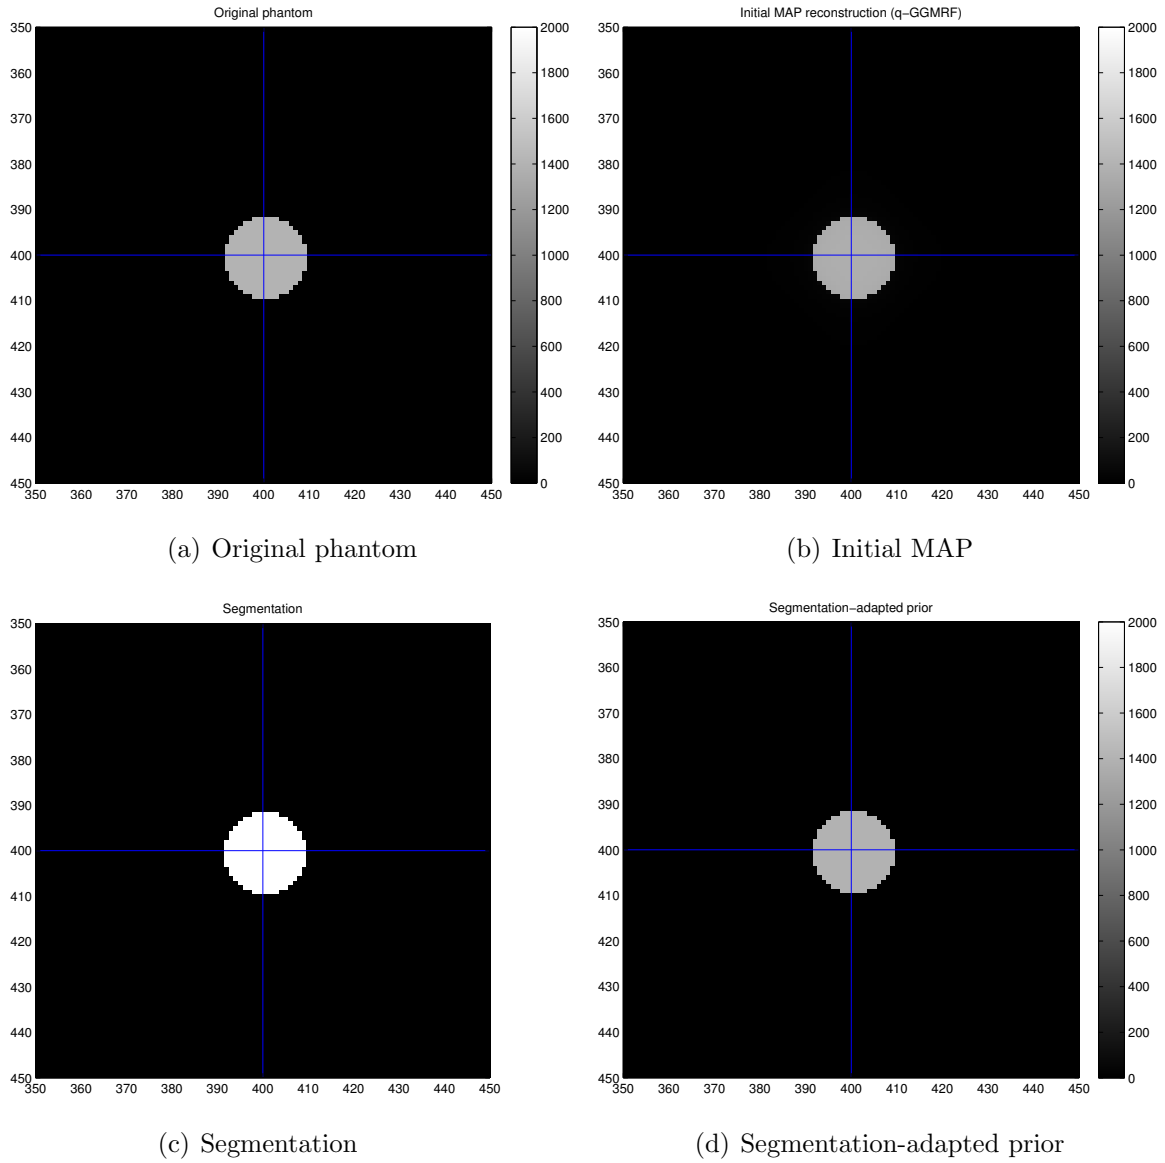

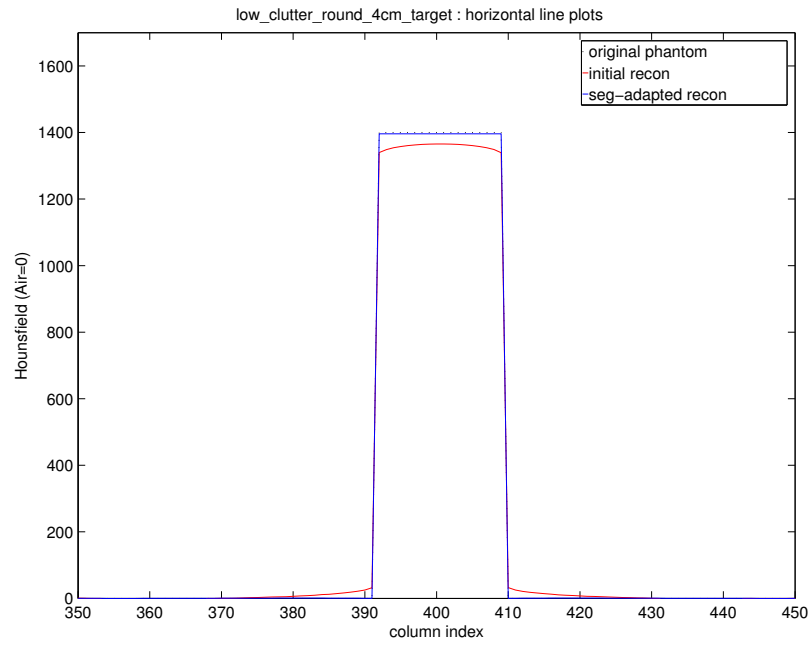


Fig. 5.8. Segmentation-adapted reconstruction for a target in *low clutter* level conditions. The target is 4 cm in diameter, with a 2.2 mm pixel size. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is a segmentation of the initial reconstruction, and (d) is the segmentation-adapted reconstruction (edge weight=0, interior weight=5). 1-D profiles along the indicated blue lines are provided in separate graphs.



(a) Vertical line profile



(b) Horizontal line profile

Fig. 5.9. Segmentation-adapted reconstruction for a 1400 HU target in *low clutter* level conditions. Vertical and horizontal pixel profiles through target.

Table 5.2

Accuracy of segmentation-assisted reconstruction for a round 4 cm target using various clique weights at the edges and interior of segmented regions. The ordered triples contain the (*target deviation*, *target RMSE*, *image RMSE*), where the *target deviation* is the average deviation from 1400 HU within the target region, and the *target RMSE* and *image RMSE* are the root mean square error of the reconstruction within the target region only, and across the entire image respectively. The CT values for the original homogeneous-prior reconstruction are those indicated in the boxes (unity weighting coefficients). The *edge weight* was applied to pairwise cliques that fall across an edge of a segmented component boundary, and the *interior weight* was applied to cliques that fall in the interior of a segmented component.

		<u>interior weight</u>		
		1.0	2.0	5.0
<u>Clutter</u>	<u>edge weight</u>	(target deviation, target RMSE, image RMSE)		
Low	1	(-48, 49, 49.1)	(-49, 49, 48.7)	(-49, 49, 48.8)
	0.5	(-29, 29, 28.9)	(-29, 29, 29.3)	(-30, 30, 30.4)
	0.2	(-12, 12, 11.9)	(-12, 12, 12.4)	(-15, 15, 14.7)
	0	(-0, 0, 0.2)	(-1, 1, 0.7)	(-4, 4, 3.9)
Med	1	(-14, 26, 159.7)	(-12, 16, 159.9)	(-11, 12, 159.7)
	0.5	(-9, 17, 157.7)	(-8, 10, 157.8)	(-8, 8, 157.8)
	0.2	(-7, 14, 158.3)	(-6, 8, 158.5)	(-6, 7, 158.5)
	0	(-6, 13, 159.1)	(-5, 7, 159.4)	(-5, 6, 159.3)
High	1	(-122, 203, 192.5)	(-121, 171, 192.4)	(-119, 170, 192.4)
	0.5	(-77, 214, 192.4)	(-79, 197, 192.4)	(-84, 196, 192.4)
	0.2	(-52, 240, 192.5)	(-55, 225, 192.4)	(-64, 222, 192.4)
	0	(-36, 263, 192.6)	(-41, 245, 192.5)	(-51, 240, 192.5)

5.4.2 Results for Bilateral Adaptive Prior

Model-based reconstructions using the *bilateral adaptive prior* described in Section 5.3 were assessed using cargo images containing the targets in Figure 5.3(a,b). For this analysis a single “high-clutter” background image was used (same background used for the results of Figure 5.4). After generating the sinogram data, initial model-based reconstructions were computed using a homogeneous QGGMRF prior.

Figure 5.10(c) shows the result after a single prior readjustment, and 5.10(d) after 16 readjustments. The edges in this case are reconstructed accurately, as opposed to the segmentation-adapted result in Figure 5.4. The profiles in Figure 5.11 show the behavior after various iterations of prior readjustment. The reconstruction accuracy around the edges is generally improved from the bilateral weighting, but the results indicate that multiple iterations do not always uniformly decrease the error. For example, the accuracy at the left edge in 5.11(a) is improved with each iteration, but the left edge in 5.11(b) is closest after only 2 passes. In addition, there are two outlier patches in the interior of the target whose error is significantly increased between 8 and 16 iterations.

The results for the thin target in Figures 5.12 and 5.13 show a similar result. The reconstruction accuracy is clearly improved from the bilateral weighting, but after one or two iterations it is not clear that additional passes benefit the reconstruction.

Table 5.3 lists quantitative results for these experiments. Note that multiple iterations are providing better average target CT values, but also increase both the image and target RMSE. In this case, two iterations produced the most accurate results under the condition that the overall image RMSE does not exceed that of the initial reconstruction.

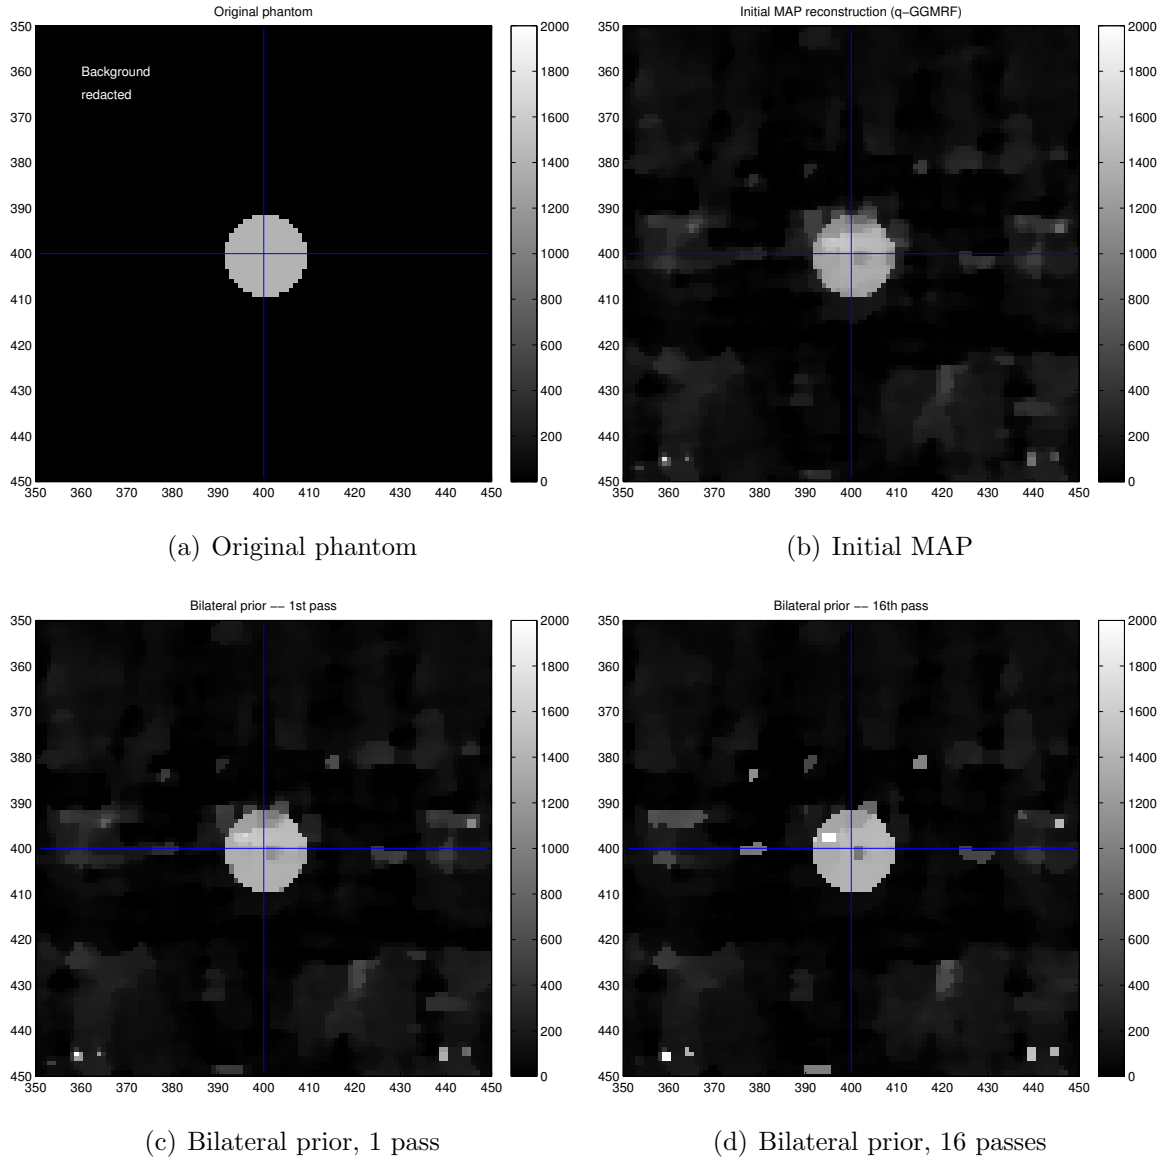
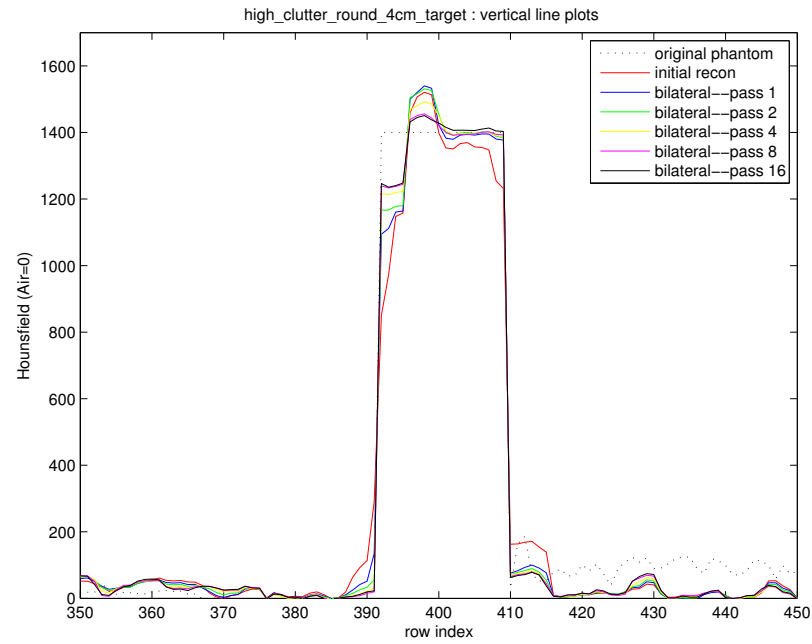
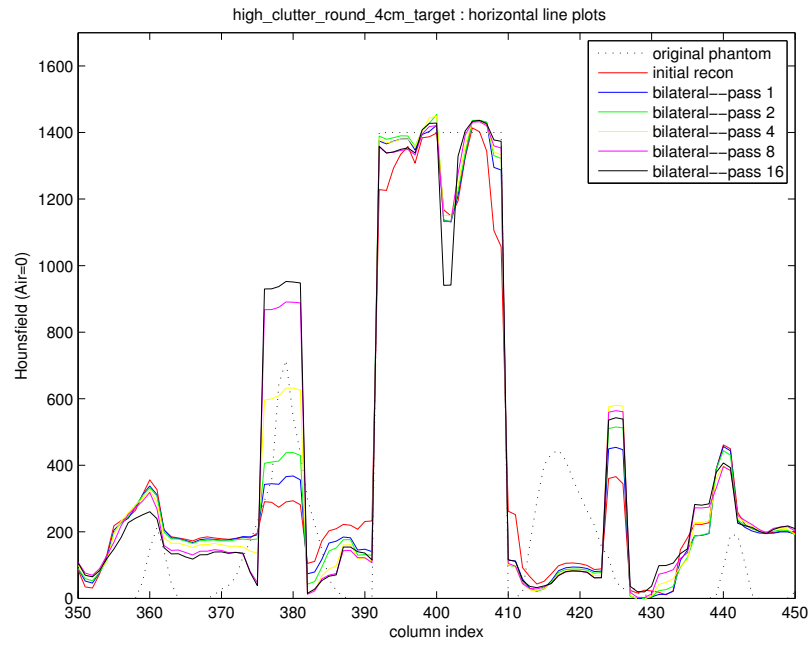


Fig. 5.10. Bilateral adaptive reconstruction for the *round 4 cm* target. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QGGMRF prior, (c) is the bilateral adaptive reconstruction with coefficients determined from the result in (b), and (d) shows the result of the bilateral reconstruction after 16 iterations of determining the coefficients followed by re-optimization. 1-D profiles along the indicated blue lines are provided in separate graphs.



(a) Vertical line profile



(b) Horizontal line profile

Fig. 5.11. Bilateral adaptive reconstructions for the *round 4 cm* target. Profiles are shown for reconstructions after various iterations of prior readjustment.

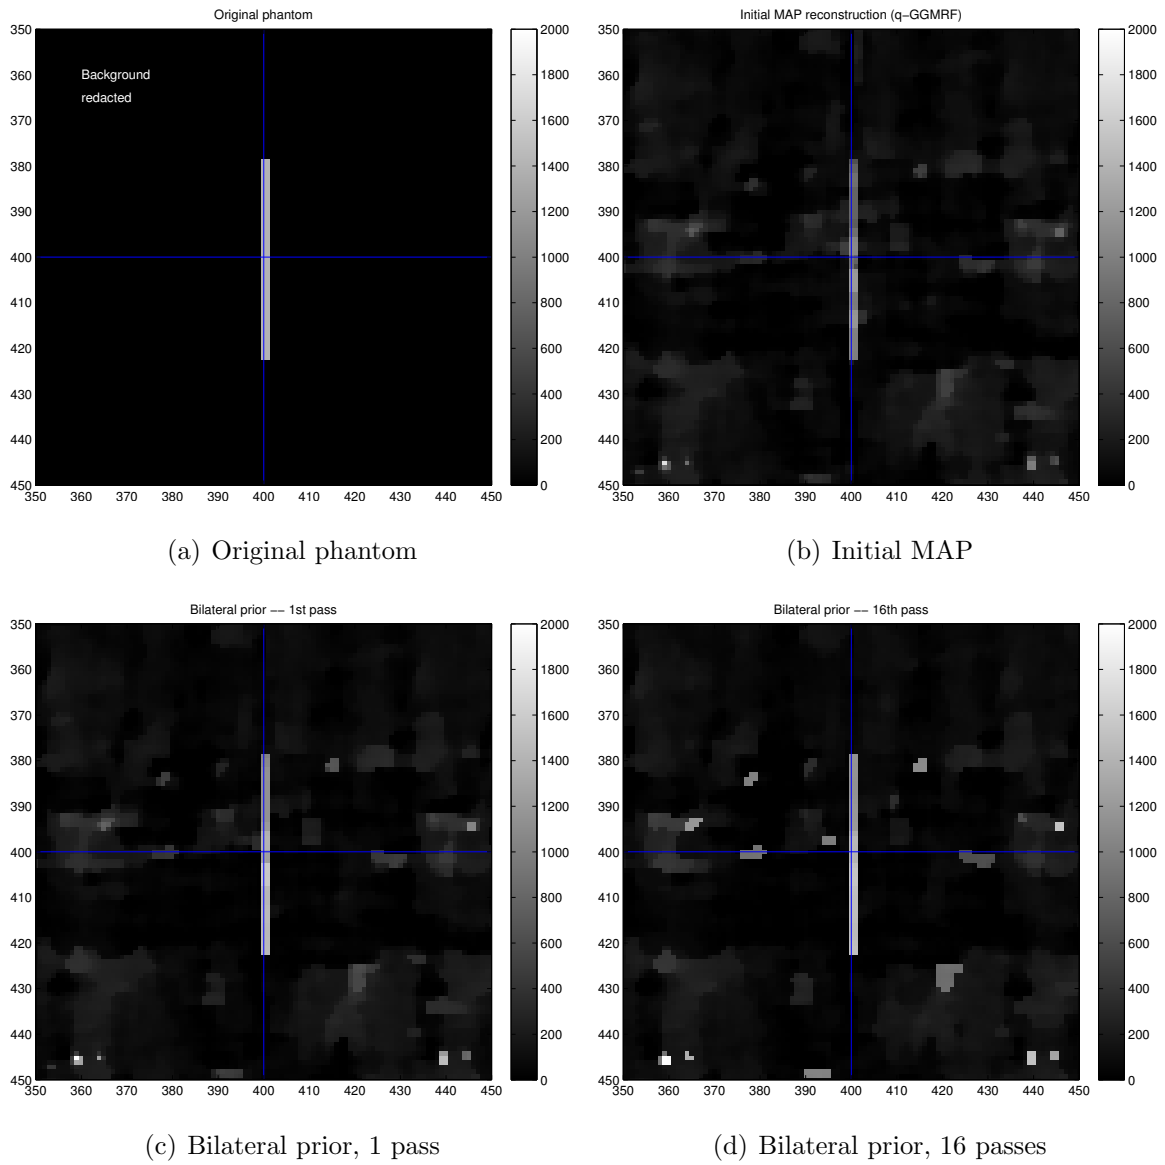




Fig. 5.12. Bilateral adaptive reconstruction for the *thin 4.4 mm* target. (a) is the original target phantom used to generate the 32-view sinogram, (b) is the initial reconstruction using a homogeneous QG-GMRF prior, (c) is the bilateral adaptive reconstruction with coefficients determined from the result in (b), and (d) shows the result of the bilateral reconstruction after 16 iterations of determining the coefficients followed by re-optimization. 1-D profiles along the indicated blue lines are provided in separate graphs.



(a) Vertical line profile



(b) Horizontal line profile

Fig. 5.13. Bilateral adaptive reconstruction for the *thin 4.4 mm* target. Profiles are shown for reconstructions after various iterations of prior readjustment.

Table 5.3

Accuracy of bilateral adaptive prior reconstruction for the two target types. The ordered triples contain the (*target deviation*, *target RMSE*, *image RMSE*), where the *target deviation* is the average deviation from 1400 HU within the target region, and the *target RMSE* and *image RMSE* are the root mean square error of the reconstruction within the target region only, and across the entire image respectively. The CT values for the original homogeneous-prior reconstruction are listed in the top row. Subsequent rows list the results after k iterations of, (1) readjusting the clique weights, and (2) recomputing the MAP estimate with the new posterior.

Iteration	4 cm diam	4.4 mm wide
	(target deviation, target RMSE, image RMSE)	
0 (initial recon)	(-122, 202, 192.5)	(-446, 474, 192.9)
1	(-49, 155, 191.1)	(-74, 182, 191.2)
2	(-30, 138, 192.3)	(-29, 135, 192.4)
4	(-21, 131, 194.9)	(-21, 135, 195.0)
8	(-15, 184, 197.8)	(-29, 145, 197.9)
16	(-13, 206, 199.7)	(-41, 160, 199.9)

6. ESTIMATION OF NOISE MODEL PARAMETERS

The noise characteristics of a CT system are important to model correctly, especially for model-based reconstruction methods. This not only benefits the reconstruction, but also allows better simulation of data from a real system.

The noise characteristics of a given system generally need to be measured experimentally. This chapter describes an approach for estimating parameters associated with photon statistics as well as noise in the sensors and electronics. The first section reviews a noise model often employed in MBIR. This is followed by parameter estimates based on X-ray transmission measurements on a step-wedge block, and the last section discusses simulation. The methods were applied to a real system, but the results are not included here for proprietary reasons.

6.1 Noise Model for X-ray Projections

Starting from a Poisson model for the photon counts, λ_i , the log likelihood of the measured projections, y , given the image, x , is approximated by the following [14],

$$\log p(y|x) \approx -\frac{1}{2}(y - Ax)^t D(y - Ax) + g(y) \quad (6.1)$$

where A is a forward projection operator, and

$$y_i = \log(\lambda_T/\lambda_i) \quad (6.2)$$

$$D = \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_n\} \quad (6.3)$$

where λ_i is the X-ray flux at detector i and λ_T is the source flux. In the MAP estimation of x , the cost function retains the term

$$\frac{1}{2}(y - Ax)^t D(y - Ax) . \quad (6.4)$$

In this context, the matrix D has an interpretation as the inverse covariance matrix of the projection vector, y . In other words, the variance of projection y_i is approximated by $1/\lambda_i$. Intuitively, in the cost framework the diagonal terms of D are weighting coefficients indicating the reliability of each projection measurement (e.g. lower counts are less reliable).

There are reasons for adjusting this data model in (6.4). One is that additional noise in the sensors and electronics will cause a deviation from the Poisson model of the measured counts. Another is that other transformations, in addition to a *log* operation, may be applied to the counts in estimating the projections, y , for example a beam hardening correction. The entries of D can therefore be adjusted to more accurately reflect the variance of the projection measurements.

Consider a practical X-ray measurement that contains electronic noise. Assume the measurement of photon count λ_i has an additive zero-mean white noise component, w_i .

$$\tilde{\lambda}_i = c (\lambda_i + w_i) \quad (6.5)$$

The constant c is implicit in the calibration of the acquisition system so that an air scan produces a measurement at the high end of the dynamic range of the analog-to-digital (A/D) converter (e.g. $c = 2^p/\lambda_T$, where p is the number of bits in the ADC). The mean and variance of the measurement, $\tilde{\lambda}_i$, are given by,

$$E[\tilde{\lambda}_i] = c E[\lambda_i] \quad (6.6)$$

$$Var(\tilde{\lambda}_i) = c^2 Var(\lambda_i) + c^2 \sigma_w^2 \quad (6.7)$$

$$= c^2 E[\lambda_i] + c^2 \sigma_w^2 \quad (6.8)$$

$$= c E[\tilde{\lambda}_i] + c^2 \sigma_w^2 \quad (6.9)$$

where we use the fact that for a Poisson random variable,

$$Var(\lambda_i) = E[\lambda_i] . \quad (6.10)$$

From the noisy calibrated measurements, the projections can be estimated by

$$y_i = \log(\tilde{\lambda}_T/\tilde{\lambda}_i) . \quad (6.11)$$

What is now needed is a new approximation of the variance of the projections, y_i , to substitute into the D matrix of the data term in (6.4). Consider a more general case of (6.11) where a well-behaved transformation f is used to estimate the projections, $y_i = f(\tilde{\lambda}_i)$. A first order approximation of the deviation of y_i from its mean is given by

$$(y_i - E[y_i]) \approx (\tilde{\lambda}_i - E[\tilde{\lambda}_i]) f'(E[\tilde{\lambda}_i]) . \quad (6.12)$$

Squaring each side and taking the expectation,

$$Var(y_i) \approx Var(\tilde{\lambda}_i) |f'(E[\tilde{\lambda}_i])|^2 . \quad (6.13)$$

If f consists of only the *log* operation as in (6.11), then

$$f(t) = \log(\tilde{\lambda}_T) - \log(t) \quad (6.14)$$

$$|f'(t)|^2 = 1/t^2 \quad (6.15)$$

and,

$$Var(y_i) \approx Var(\tilde{\lambda}_i) \frac{1}{E[\tilde{\lambda}_i]^2} \quad (6.16)$$

$$= \frac{c E[\tilde{\lambda}_i] + c^2 \sigma_w^2}{E[\tilde{\lambda}_i]^2} . \quad (6.17)$$

So the diagonal elements of D become,

$$D_{ii} = \frac{E[\tilde{\lambda}_i]^2}{c E[\tilde{\lambda}_i] + c^2 \sigma_w^2} . \quad (6.18)$$

Typically during a scan only one measurement is taken for each projection ray, so the mean $E[\tilde{\lambda}_i]$ must be estimated by the single measurement $\tilde{\lambda}_i$. The terms c and σ_w^2 are constants that can be estimated off-line.

6.2 Estimation of Model Parameters

To estimate the constants c and σ_w^2 associated with the (approximate) projection variances, start from the previous result,

$$Var(\tilde{\lambda}_i) = c E[\tilde{\lambda}_i] + c^2 \sigma_w^2 . \quad (6.19)$$

Having a set of X-ray transmission measurements of various density objects, from which $E[\tilde{\lambda}_i]$ and $Var(\tilde{\lambda}_i)$ can be estimated for several photon rates, one can perform a standard least squares regression to estimate c and $c^2\sigma_w^2$ in (6.19). Denoting the mean and variance estimates as $Var(\tilde{\lambda})_k$ and $E[\tilde{\lambda}]_k$ for $k = 1..n$, we can write the overdetermined system,

$$\begin{bmatrix} \vdots \\ Var(\tilde{\lambda})_k \\ \vdots \end{bmatrix} = \begin{bmatrix} \vdots & 1 \\ E[\tilde{\lambda}]_k & \vdots \\ \vdots & 1 \end{bmatrix} \begin{bmatrix} c \\ c^2\sigma_w^2 \end{bmatrix} \quad (6.20)$$

or, compacting notation,

$$V = H\Theta . \quad (6.21)$$

The least squares fit for the parameter vector, Θ , is then

$$\hat{\Theta} = (H^t H)^{-1} H^t V . \quad (6.22)$$

6.3 Experimental Parameter Estimation

A series of X-ray projections were acquired for a series of steel step-wedge blocks of varying thickness. Each region of a given thickness was used to estimate the mean and variance of the photon count measurements. Since perturbations in the mean over a given region can cause inaccuracies in the variance estimate, one can allow for a spatially varying mean in calculating the variance. First, a mean estimate, $\hat{\mu}_i$, was calculated for each pixel location in a given region by averaging over the surrounding 9×9 window. Then the variance estimate for the region was calculated as,

$$\hat{Var} = \frac{1}{n-1} \sum_{i=1}^n (x_i - \hat{\mu}_i)^2 \quad (6.23)$$

where x_i are the raw measurements for the given region.

After obtaining the mean and variance estimates for the various levels of attenuation from the step-wedge data, the parameters c , $c^2\sigma_w^2$, and hence σ_w , were estimated using Equation (6.22). The results have been omitted here for proprietary reasons.

6.4 Simulation of Projection Data

Using the noise model and parameter estimates calculated in the previous section, one can better simulate projection data for a given image scene. First, to determine the X-ray dosage in units of counts, recall the previous result relating the measurement, $\tilde{\lambda}_i$, to the actual photon count, λ_i .

$$E[\tilde{\lambda}_i] = c E[\lambda_i] \quad (6.24)$$

The average dosage as measured from the air scan would then be

$$E[\lambda_T] = E[\tilde{\lambda}_T]/c . \quad (6.25)$$

For the sake of example, assume the measured signal from the air scan is calibrated to $E[\tilde{\lambda}_T] = 2^p$ (the maximum value output from the ADC). The average dosage is then estimated as,

$$\bar{\lambda}_T = 2^p/c . \quad (6.26)$$

Now assuming the dosage is constant, we can determine the theoretical average count at detector i for a given image x using,

$$\bar{\lambda}_i = \bar{\lambda}_T \exp(-p_i) \quad (6.27)$$

$$= \bar{\lambda}_T \exp(-A_{i,*}x) . \quad (6.28)$$

Photon counting noise and electronic noise can then both be incorporated by replacing the theoretical count by the sum of a Poisson random variable (mean= $\bar{\lambda}_i$) and a Gaussian random variable (mean=0, var= σ_w), as

$$\lambda_i = \bar{\lambda}_i + \sqrt{\bar{\lambda}_i} Z_1 + \sigma_w Z_2 \quad (6.29)$$

where Z_1 and Z_2 are independent $\mathcal{N}(0, 1)$ random variables (note a Poisson random variable is well-approximated, in the mean-square sense, by a Gaussian if $\bar{\lambda}_i$ is not small). Given the possibility of obtaining a negative count in λ_i , its value should be floored.

If desired, the noisy count can be transformed back into a projection via,

$$y_i = \log \left(\frac{\bar{\lambda}_T}{\lambda_i} \right) , \quad (6.30)$$

however the model-based reconstructions generally benefit from having both the count, λ_i , and the dosage, λ_T .

LIST OF REFERENCES

LIST OF REFERENCES

- [1] R. C. Smith and J. M. Connelly, "Aspects of explosives detection," in *CT Technologies* (M. Marshall and J. Oxley, eds.), pp. 131–45, Boston: Elsevier, 2009.
- [2] S. Singh and M. Singh, "Explosives detection systems (EDS) for aviation security," *Signal Processing*, vol. 83, pp. 31–55, Jan. 2003.
- [3] G. Zentai, "X-ray imaging for homeland security," in *IEEE International Workshop on Imaging Systems and Techniques*, pp. 1–6, Sept. 2008.
- [4] Z. Ying, R. Naidu, K. Guilbert, D. Schafer, and C. R. Crawford, "Dual energy volumetric x-ray tomographic sensor for luggage screening," in *Sensors Applications Symposium, 2007. SAS '07. IEEE*, pp. 1–6, Feb. 2007.
- [5] M. Ellenbogen and R. Bijjani, "Liquids and homemade explosive detection," in *Proc. SPIE, Optics and Photonics in Global Homeland Security V*, vol. 7306, pp. 73060Y–1–7, May 2009.
- [6] T. Buzug, *Computed Tomography: From Photon Statistics to Modern Cone-Beam CT*. Springer, 2008.
- [7] J. Hsieh, *Computed Tomography: principles, design, artifacts, and recent advances*. Bellingham, WA: SPIE, 2nd ed., 2009.
- [8] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988.
- [9] P. Jin, E. Haneda, K. D. Sauer, and C. A. Bouman, "A model-based 3D multi-slice helical CT reconstruction algorithm for transportation security application," in *Proc. of The 2nd International Conference on Image Formation in X-Ray Computed Tomography*, June 2012.
- [10] S. J. Kisner, P. Jin, C. A. Bouman, K. Sauer, W. Garms, T. Gable, S. Oh, M. Merzbacher, and S. Skatter, "Innovative data weighting for iterative reconstruction in a helical CT security baggage scanner," in *47th IEEE International Carnahan Conference on Security Technology*, (Medellin, Colombia), October 2013.
- [11] S. J. Kisner, E. Haneda, C. A. Bouman, S. Skatter, M. Kourinny, and S. Bedford, "Limited view angle iterative CT reconstruction," in *Proceedings of SPIE-IS&T Electronic Imaging, Computational Imaging X*, vol. 8296, (Burlingame, CA), January 2012.
- [12] J.-B. Thibault, K. Sauer, C. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multi-slice helical CT," *Medical Physics*, vol. 34, no. 11, pp. 4526–4544, 2007.

- [13] J.-B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh, "A recursive filter for noise reduction in statistical iterative tomographic imaging," in *Proc. SPIE 6065, Computational Imaging IV*, Jan. 2006.
- [14] K. Sauer and C. Bouman, "A local update strategy for iterative reconstruction from projections," *IEEE Trans. on Signal Processing*, vol. 41, pp. 534–548, February 1993.
- [15] C. Bouman and K. Sauer, "A generalized Gaussian image model for edge-preserving MAP estimation," *IEEE Trans. on Image Processing*, vol. 2, pp. 296–310, July 1993.
- [16] Z. Yu, J. Thibault, C. Bouman, K. Sauer, and J. Hsieh, "Fast model-based X-Ray CT reconstruction using spatially non-homogeneous ICD optimization," *IEEE Trans. on Image Processing*, vol. 20, no. 1, pp. 161–175, 2011.
- [17] J. Besag, "Spatial interaction and the statistical analysis of lattice systems," *Journal of the Royal Statistical Society B*, vol. 36, no. 2, pp. 192–236, 1974.
- [18] J. Besag, "On the statistical analysis of dirty pictures," *Journal of the Royal Statistical Society B*, vol. 48, no. 3, pp. 259–302, 1986.
- [19] R. Kindermann and J. Snell, *Markov Random Fields and their Applications*. Providence: American Mathematical Society, 1980.
- [20] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, pp. 721–741, November 1984.
- [21] R. Dubes and A. Jain, "Random field models in image analysis," *Journal of Applied Statistics*, vol. 16, no. 2, pp. 131–164, 1989.
- [22] R. Chellappa and A. Jain, eds., *Markov Random Fields: Theory and Applications*. Boston: Academic Press, Inc., 1993.
- [23] C. A. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *IEEE Trans. on Image Processing*, vol. 5, pp. 480–492, March 1996.
- [24] R. Stevenson and E. Delp, "First international workshop on robust computer vision," in *Fitting curves with discontinuities*, pp. 127–136, October 1990.
- [25] J. Besag, "Towards Bayesian image analysis," *Journal of Applied Statistics*, vol. 16, no. 3, pp. 395–407, 1989.
- [26] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 14, pp. 259 – 268, 1992.
- [27] B. De Man and S. Basu, "Distance-driven projection and backprojection in three dimensions," *Physics in Medicine and Biology*, vol. 49, no. 11, pp. 2463–2475, 2004.
- [28] J. Hsieh, O. E. Gurmen, and K. F. King, "A recursive correction algorithm for detector decay characteristics in CT," in *Proc. SPIE, Medical Imaging 2000: Physics of Medical Imaging*, vol. 3977, pp. 298–305, 2000.

- [29] A. D. Foland, "Design of the discrete skew geometry and iterative reconstruction of the MV3D scanner," in *Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, 2011.
- [30] M. E. Davison, "The ill-conditioned nature of the limited angle tomography problem," *SIAM J. Appl. Math.*, vol. 43, pp. 428–448, 1983.
- [31] K. Hanson and G. Wecksung, "Bayesian approach to limited-angle reconstruction computed tomography," *J. Optical Society. America*, vol. 73, pp. 1501–1509, November 1983.
- [32] J. Prince and A. Willsky, "A hierarchical algorithm for limited-angle reconstruction," in *Proc. of IEEE Int'l Conf. on Acoust., Speech and Sig. Proc.*, (Glasgow, Scotland), pp. 1468–1471, May 23-26 1989.
- [33] B. Sahiner and A. Yagle, "Limited angle tomography using wavelets," in *IEEE Nuclear Science Symposium & Medical Imaging Conf.*, (San Francisco, CA.), pp. 1912–1916, 1993.
- [34] J. A. Fessler and L. Rogers, "Spatial resolution properties of penalized-likelihood image reconstruction: space-invariant tomographs," *IEEE Trans. on Image Processing*, vol. 5, pp. 1346–58, September 1996.
- [35] C. Tomasi and R. Manduchi, "Sixth international conference on computer vision," in *Bilateral filtering for gray and color images*, pp. 839–46, January 1998.
- [36] M. Elad, "On the origin of the bilateral filter and ways to improve it," *IEEE Trans. on Image Processing*, vol. 11, pp. 1141–1150, 2002.

VITA

VITA

Sherman J. Kisner was awarded a Ph.D. in Electrical Engineering in December 2013 by the School of Electrical and Computer Engineering at Purdue University in West Lafayette, Indiana. His research, conducted under the direction of Prof. Charles A. Bouman, involved model-based methods for X-ray computed tomography, with a focus on security screening systems. During his graduate education he taught full undergraduate courses in circuit analysis and digital signal processing, and also lectured and developed curricula for a graduate level image processing course. He was previously awarded degrees Bachelor of Science and Master of Science in Electrical Engineering, both from Purdue University.