

Modeling and Pre-Treatment of Photon-Starved CT Data for Iterative Reconstruction

Zhiqian Chang*, *Student Member, IEEE*, Ruoqiao Zhang, *Member, IEEE*,
Jean-Baptiste Thibault, *Member, IEEE*, Debashish Pal, *Member, IEEE*,
Lin Fu, *Member, IEEE*, Ken Sauer, *Member, IEEE*, and Charles Bouman, *Fellow, IEEE*

Abstract—An increasing number of X-ray CT procedures are being conducted with drastically reduced dosage, due at least in part to advances in statistical reconstruction methods that can deal more effectively with noise than can traditional techniques. As data become photon-limited, more detailed models are necessary to deal with count rates that drop to the levels of system electronic noise. We present two options for sinogram pre-treatment that can improve the performance of photon-starved measurements, with the intent of following with model-based image reconstruction. Both the local linear minimum mean-squared error (LLMMSE) filter and pointwise Bayesian restoration (PBR) show promise in extracting useful, quantitative information from very low-count data by reducing local bias while maintaining the lower noise variance of statistical methods. Results from clinical data demonstrate the potential of both techniques.

Index Terms—Adaptive filtering, Bayesian inference, computed tomography, iterative reconstruction, low-dose CT, photon starvation, statistical reconstruction.

I. INTRODUCTION

CONCERNS over long-term health effects of radiation exposure in CT motivate diagnostic imaging with minimal X-ray dosages. The improved image quality offered by model-based, iterative reconstruction (MBIR) methods in the presence of noisy data [1]–[3] has created expectations that lower exposure levels can be used to extract sufficient information to make diagnoses and/or furnish partial information for subsequent procedures. These improved reconstruction capabilities have led to experimentation with dramatically reduced exposures in the clinic [4]–[6]. Dealing with the resulting

noise characteristics of model-based methods is critical to clinical utility when applying MBIR to extremely low signal-to-noise ratio (SNR) data. Improving quantitative accuracy in these settings may be important in applications such as lung assessment near high-attenuation shoulder regions in screening for emphysema [7], [8].

Photon starvation in CT detectors has long been an issue, particularly in such low-flux-rate applications as attenuation compensation scanning systems in emission tomography [9], where extremely low SNR is routine. In conventional diagnostic CT, near-total photon starvation has generally been treated as an outlier event, with ad hoc treatments to allow such measurements to be treated in the conventional processing chain. Particularly in cases of metal or other nearly radio-opaque objects in the field of view, correction of “bad” data with various interpolation or data replacement approaches [10]–[12] reduces severe artifacts from highly localized X-ray blockage. Dramatically reduced overall dosage however, may cause a significant fraction of data to be photon-starved; to this situation the replacement methods are not particularly well-suited.

As indicated by its name, MBIR could allow us to model the data in such a way that detectors registering no photon energy, or even a negative number due to additive electronic noise, do not pose a serious statistical problem. Effects of polychromaticity in beam hardening, scatter, Poisson statistics, arbitrary detector response shapes, etc. can be included in a forward model that allows all readings to contribute appropriately, statistically, to the reconstructed imagery. However, for the sake of limiting the computational cost of iterative methods, simplifications are necessary, and we will seek to treat photon-starved data such that reconstruction can be equally simple to the normal-dosage case. That is, we will consider approaches that can be applied as pre-treatment of data for application of conventional MBIR routines. A variety of sinogram preprocessing techniques has improved performance in generic reduced dosage settings [13]–[22], primarily as a prelude to reconstruction by deterministic methods such as filtered back projection (FBP). The fact that we plan to reconstruct via MBIR changes the requirements for sinogram data restoration. The aforementioned noise preprocessing methods are aimed primarily at reducing noise variance to enhance the limited control of FBP over that attribute of reconstructions. MBIR has the capability of suppressing noise according to locally varying statistics and spatially adaptive stochastic image models; therefore it is less sensitive to data variance and can tolerate

Manuscript received July 5, 2016; revised August 26, 2016; accepted August 30, 2016. Date of publication September 7, 2016; date of current version December 29, 2016. This research was supported by GE Healthcare. Asterisk indicates corresponding author.

*Z. Chang is with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556 USA (e-mail: zchang1@nd.edu).

R. Zhang is with the Department of Radiology, University of Washington, Seattle, WA 98195 USA (e-mail: zhangrq@uw.edu).

J.-B. Thibault and D. Pal are with Applied Science Laboratory, GE Healthcare, Waukesha, WI 53188 USA (e-mail: jean-baptiste.thibault@med.ge.com; debashish.pal@ge.com).

L. Fu is with Image Reconstruction Laboratory, GE Global Research, Niskayuna, NY 12309 USA (e-mail: fulin@ge.com).

K. Sauer is with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556 USA (e-mail: sauer@nd.edu).

C. Bouman is with the Department of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: bouman@purdue.edu).

Digital Object Identifier 10.1109/TMI.2016.2606338

sinogram correction methods aimed at eliminating bias at the expense of potentially increased sinogram variance.

We propose two methods for X-ray CT sinogram restoration under very low photon counts. The first is in the same vein as [14], for recovering local means in photon counts through spatially adaptive filtering. Our algorithm is based on non-stationary, signal-dependent noise and modeled after that of Kuan, et. al [23], called the local linear minimum mean-squared error estimator (LLMMSE). Local estimates of statistics parameterize filter coefficients. The second alternative is designed to better preserve conditional independence among sinogram photon count data, and operates point-wise. It relies on a Bayesian scalar estimation model, with a weak *a priori* density on the Poisson mean of each detector output, and will be referred to as pointwise Bayesian restoration (PBR).

Preliminary results on our photon-starvation methods were presented in two conference papers [24], [25]. Further studies included in this paper have improved the methods' robustness.

The paper is organized as follows. In Section II, conventional noise modeling for statistical methods is described. Section III discusses the issues with photon-starved measurements at extremely low dose and then presents two proposed data correction methods. Section IV shows results of sinogram correction as well as reconstructed images from both phantom and clinical scans. Finally, Section V discusses the algorithms' comparison and potential future work.

II. STATISTICAL MODELING OF SINOGRAM MEASUREMENTS

We begin with treatment of modeling issues that are especially pronounced at very low transmission rates, but merit inclusion in a large fraction of CT scans. Transmission tomographic data is initially captured as detector outputs approximately proportional to numbers of photons arriving at the detector surface. If a monoenergetic input Poisson rate of λ_T is attenuated by material having line integral density p_j between detector and source, corresponding to projection j , the random received count is, by the Beer-Lambert law [18],

$$\Lambda_j \sim \text{Poisson}\{\lambda_T e^{-p_j}\}. \quad (1)$$

This can easily be generalized to the case of spatially and/or temporally varying input rates. Conditioned on the set of integral projections $\{p_j\}$, counts may be modeled as mutually independent, resulting in the log-likelihood

$$L(\lambda|p) = \sum_{j=1}^M (-\lambda_T e^{-p_j} - \lambda_j p_j + \text{constant}(\lambda_j)), \quad (2)$$

where M denotes the total number of projection measurements. While (2) is not highly complex, greater simplicity in both computation and intuitive understanding of the estimation problem is available in a collection of one-dimensional Taylor series expansions [26] about the point y in data space,

$$L_j(\lambda_j|p_j) \approx L_j(\lambda_j|y_j) - d_{1,j}(y_j - p_j) + \frac{d_{2,j}}{2}(y_j - p_j)^2, \quad (3)$$

in which $d_{k,j}$ is the k -th derivative of $L_j(\lambda_j|p)$ with respect to p_j . Probably the most useful and common choice of




Fig. 1. Log-likelihood function $L_j(\lambda_j|p_j)$ in (2) (smooth curves) and corresponding quadratic approximation (curves with markers) for varying photon counts. Input flux $\lambda_T = 2 \times 10^5$.

expansion point for the approximation of (3) is the maximum-likelihood value of the integral density,

$$y_j = -\log \frac{\lambda_j}{\lambda_T}. \quad (4)$$

This transformation is also the standard mapping of counts to projections in X-ray CT.

The end goal of this work is not estimation of sinogram values, but rather the discretized image, x . We model each unknown, true integral projection as $p_j = A_{j*}x$, which requires that the matrix A capture the geometry of the voxelized 3D image, the detector sensitivity profiles and a linearization of the response of materials to the generally polychromatic X-rays. Choosing a diagonal data weighting matrix W populated by entries $w_j = -d_{2,j}$, the approximate ML estimation of x is given by

$$\hat{x} = \arg \min_{x \in \Omega} \left\{ \frac{1}{2} (y - Ax)^T W (y - Ax) \right\}, \quad (5)$$

where Ω is the feasible region of x .

This approximation has provided a basis for a range of useful statistical approaches to tomographic estimation formulations and their optimizations [27], [28]. The entries in the matrix W may vary over several orders of magnitude in common scans, representing inverse variance of individual measurements. Exploitation of this varying data reliability is the primary advantage of statistical methods over deterministic methods. The Hessian of this formulation, $A^T W A$, is an approximation to the Fisher information matrix [29] for estimating voxel values from the sinogram data. The diagonal entries therefore provide a measure of the information available for the value of a voxel when the remainder of the image is fixed.

Fig. 1 shows scalar log-likelihoods in p_j and their quadratic approximations for four values of the received photon count λ_j . The higher counts on the left yield log-likelihoods that are well-matched by the Taylor series approximations. As counts decrease, representing greater attenuation due to

greater densities, the deviation from quadratic becomes dramatic. Simultaneously, the curvature is reduced proportionally to λ_j . In a vector estimation problem such as the one at hand for x , the compromise made to minimize (5) will tend to accept approximately equal weighted costs among measurements. Thus low-count projection estimates in the vector optimal estimate for x may be forced far from their scalar maximum-likelihood values of (4), into the region where the simple quadratic approximation is quite inaccurate. When counts drop to starvation levels (4) and (5) may need adjustment from their simplest, high-dosage versions to allow useful interpretation of the full set of measurements.

A. Additive Poisson Noise

For clinical CT X-rays, the primary attenuation mechanism is Compton scatter [18]. This deflection of photons from their original path is accompanied by energy loss, but a large fraction remains within the sensitivity range and the physical capture surface of detectors. Most detectors are not capable of energy discrimination and cannot distinguish between scattered and unscattered photons. Scatter forms an additive offset to Λ , sometimes constituting the majority of captured energy in uncollimated, or poorly collimated, systems. Because angular distribution is wide relative to detector size, this bias in counts is often modeled as independent of local image content, which allows it to be treated as an additive, independent Poisson, whose rate we denote γ_j .

Eq.(1) may then be modified as

$$\Lambda_j \sim \text{Poisson}\{\lambda_T e^{-p_j} + \gamma_j\}. \quad (6)$$

Our log-likelihood becomes

$$L_j(\lambda_j|x) = -(\lambda_T e^{-p_j} + \gamma_j) + \lambda_j \log(\lambda_T e^{-p_j} + \gamma_j). \quad (7)$$

The parameter γ_j of these additive Poisson variates possess slowly varying spatial features, which can often be estimated with sufficient accuracy to be subtracted in preprocessing. Additional variance from the higher counts, however, remains. The background rate may also make the log-likelihood non-convex [27], [30]. In cases where $\lambda_j > \gamma_j$ the log-likelihood has its maximum at

$$y_j = -\log \frac{(\lambda_j - \gamma_j)}{\lambda_T}, \quad (8)$$

the natural point for the Taylor expansion of (7) similarly to (3). Second derivatives result in data weighting coefficients

$$w_j = \frac{(\lambda_j - \gamma_j)^2}{\lambda_j}, \quad (9)$$

demonstrating information loss in the background counts [27].

Fig. 2 shows the effects of γ_j on the log-likelihood. In cases of very small numbers of unscattered, unattenuated photons such as the problems of interest here, scattered energy may be a major noise source. Incorrect scatter compensation may lead to bias in projection values, as illustrated in Fig. 2. Failing to compensate for the 20% fraction of radiation received as scatter leads to a shift of 0.223 for the ML line integral estimates.




Fig. 2. Log-likelihood function $L_j(\lambda_j|p_j)$ for different noise models (with additive Poisson noise γ_j and electronic noise σ_j of Section II.B) at $\lambda_j = 20$, and $\lambda_T = 2 \times 10^5$.

B. Electronic Noise

Solid state detectors utilize scintillating material to convert X-ray photons to light, followed by photodiodes transforming light into electrical current. The data acquisition system (DAS) converts these analog signals into signals intended to be proportional to photon accumulations, before mathematical operations to form projection estimates. Photodiodes and electronic components of the DAS add corruption to the signal in addition to Poisson characteristics of photon counting. We refer to this corruption as electronic noise due to its appearing in the analog signal of the DAS. It is modeled as additive and zero-mean Gaussian, independent of the information-laden Poisson variates, similarly to previous work in photon-limited optical imaging [31]. In well-designed systems, the electronic noise variance is sufficiently low compared to the Poisson mean of counts to allow it to be neglected in computing variance of projection measurements. With very limited dosage and penetration of X-rays, photon numbers may fall to the level of electronic noise variance and both noise sources are essential in modeling the statistics of the received λ_j .

We map the electronic noise into the photon domain, including it as the Gaussian component of the combined Poisson-Gaussian model,

$$\Lambda_j = K_j + N_j, \quad (10)$$

where

$$K_j \sim \text{Poisson}\{\lambda_T e^{-p_j} + \gamma_j\} \\ N_j \sim \mathcal{N}(0, \sigma_e^2),$$

and σ_e^2 is the electronic noise variance, here assumed homogeneous for all detectors. Define $\theta_j := \lambda_T e^{-p_j} + \gamma_j$, the mean count at the j^{th} ray. The probability density function (PDF) of Λ_j is formed by the convolution of PDFs of K_j and N_j , as

$$f_{\Lambda_j}(\lambda_j|x) = (f_{K_j} * f_{N_j})(\lambda_j|x) \\ = \sum_{k=0}^{\infty} (2\pi\sigma_e^2)^{-\frac{1}{2}} \exp\left\{-\frac{(\lambda_j - k)^2}{2\sigma_e^2}\right\} \frac{\theta_j^k e^{-\theta_j}}{k!}. \quad (11)$$

The joint PDF of all received count measurements Λ is

$$f_{\Lambda}(\lambda|x) = \prod_{j=1}^M f_{\Lambda_j}(\lambda_j|x). \quad (12)$$

Fig. 2 shows the effect of this second noise source on the log-likelihood, which again fails to be concave. The nature of the PDF of (11) presents complex computation for derivatives of the log-likelihood. Closed-form evaluation of the ML projection value is therefore not practical; we employ a lookup table indexed by λ_j to record maxima from (11) in the variable p_j . Given the same measurement λ_j , the ML integral density estimates with different electronic noise models are practically unchanged, with a shift of 0.005 and 0.014 for $\sigma_e = 2$ and $\sigma_e = 4$ respectively.

The primary effect of electronic noise is increasing variance of photon counts and projection estimates. For the weights in (5), we use an approximation of variance,

$$\text{Var}[Y_j] \approx \left| \frac{dy_j}{d\lambda_j} \right|^2 \text{Var}[\Lambda_j] \approx \frac{\lambda_j + \sigma_j^2}{(\lambda_j - \gamma_j)^2}. \quad (13)$$

So the statistical weight of y_j is given by

$$w_j = \frac{(\lambda_j - \gamma_j)^2}{\lambda_j + \sigma_j^2}. \quad (14)$$

This shows a further decrease in weighting due to this second noise source. For reconstruction, we pursue a MAP estimation for image x with an *a priori* model $U(x)$. Thus, the total objective function is defined as

$$\Phi(x|\lambda) = \frac{1}{2}(y - Ax)^T W(y - Ax) + U(x). \quad (15)$$

III. PROCESSING OF PHOTON-STARVED SINOGRAMS

We consider cases of extremely low photon counts, in which the Poisson mean is comparable to the standard deviation of electronic noise in terms of counts. Negative data outputs are then common occurrences due to Gaussian error, provided the system allows direct access to raw measurements without truncation such that underlying information may be recovered. The key counts-projection conversion of (8) then is invalid since the $\log()$ does not admit negative arguments. In standard pre-processing, this means no valid measurement exists for these points, and for ultra-low-dose cases, a significant number of data may be missing. However, these raw data hold information, and finding a valid way to extract it may be important in extending the range of statistical methods to unprecedentedly low exposure rates.

Non-positive measurements are commonly replaced by somewhat arbitrary, small positive numbers or simply discarded, which is equivalent to setting the corresponding w_j to zero. Alternatively, values may be interpolated from neighboring detectors with positive values [10]. When negative and zero measurements are common, these replacement methods bias overall local count values positively and local projection values negatively. The result may be low-spatial-frequency shading reminiscent of beam-hardening effects.

Error diffusion methods offer an attractive option for preserving local means while correcting individual channels' negative values. Floyd and Steinberg proposed a recursive filter that minimizes quantization errors [32], and that has a counterpart for processing of low-count regions of sinograms [33]. The filter's output is a strictly positive set of sinogram counts, with error from negative-valued channels spatially dithered. However, a positive value, ϵ , for the output floor must be chosen, and we have found local bias to be relatively sensitive to ϵ in the presence of large numbers of negatives.

A. Denoising in Counts

Integral projections A_{j*x} are exponentiated to form the mean of the Poisson photon counts. Thus accurate estimation of the mean will yield accurate estimates of the projections through the $\log()$ transformation. Any averaging is best done among photon counts rather than projection values, since per Jensen's inequality, averaging values after the convex $-\log()$ operation will bias measurements upward. This averaging is primarily intended for correcting estimates of means, but there may also be merit in smoothed maps of data weights [27] in these cases. A simple approach would be repeated linear filtering that is terminated at the point all negative count measurements have disappeared. Such filtering could well establish local count means, but will be costly in terms of resolution [15], [20], [21].

Our goal is to adapt continuously and locally to the statistics of low-quality count data, similarly to the adaptive filters that have found success in FBP preprocessing for noise suppression [13], and suffering minimal resolution loss. Somewhat similarly to recursive, least-squares linear filtering, we model the counts data as a non-stationary process, with low-frequency variation in local means for larger structure, and higher-frequency variation representing edges and texture. The adaptation of the filter will depend on local estimates of statistics for both signal and noise. The local linear minimum mean-squared error (LLMMSE) filter [23] of Kuan, et. al., appears to capture the characteristics appropriate to our denoising problem. A continuously varying convex combination of original signal values and local averages, directed by local statistics, has shown itself quite efficient in removal of signal-dependent noise [34], [35].

Operation of the LLMMSE filter is summarized by

$$\hat{\lambda}_j = \bar{\lambda}_j + \frac{\sigma_{s,j}^2}{\sigma_{s,j}^2 + \sigma_{n,j}^2}(\lambda_j - \bar{\lambda}_j), \quad (16)$$

where $\bar{\lambda}_j$ is the local average count value, intended to provide the mean of the Poisson-Gaussian variates. Each raw measurement λ_j is decomposed into the local mean component and the residual prediction error component $(\lambda_j - \bar{\lambda}_j)$. The parameters $\sigma_{s,j}^2$ and $\sigma_{n,j}^2$ are signal and electronic noise variances, respectively. The former is estimated directly from the counts, yielding the approximation $\sigma_{s,j}^2 \approx \bar{\lambda}_j$. The principal supplemental noise $\sigma_{n,j}^2$ will be the electronics component, modeled as Gaussian, typically at the level of a handful of photons. The weighting of the residual in (16) is the approximate ratio of signal variance to observation variance.




Fig. 3. Low-signal processing of photon-starved measurements (λ_j) using the LMMSE filter ($\hat{\lambda}_j$) and pointwise Bayesian recovery (PBR) ($\hat{\theta}_j$) of Section III.B. (λ_j are realizations of a Poisson-Gaussian process with $\theta_j = 2$ and $\sigma_e = 4$, and $\hat{\lambda}_j$ and $\hat{\theta}_j$ are computed using a 9-point sliding window.)

The defining equation (16) may be rewritten as

$$\hat{\lambda}_j = \eta_j \lambda_j + (1 - \eta_j) \bar{\lambda}_j, \quad (17)$$

where $\eta_j \in [0, 1]$ is

$$\eta_j = \frac{\sigma_{s,j}^2}{\sigma_{s,j}^2 + \sigma_{n,j}^2} \approx \frac{\bar{\lambda}_j}{\bar{\lambda}_j + \sigma_e^2}$$

$$1 - \eta_j = \frac{\sigma_{n,j}^2}{\sigma_{s,j}^2 + \sigma_{n,j}^2} \approx \frac{\sigma_e^2}{\bar{\lambda}_j + \sigma_e^2}.$$

When LLMMSE is applied as pre-processing for MBIR reconstruction in the formulation in (15), regularization strength in the log *a priori* $U(x)$ may need adjustment to preserve resolution. Alternatively, the operation of the filter can be restricted to very low SNR regions by thresholding η_j with

$$\hat{\eta}_j = \begin{cases} \max\{\eta_j, 0\} & \bar{\lambda}_j \leq \lambda_{th} \\ 1 & \bar{\lambda}_j > \lambda_{th} \end{cases} \quad (18)$$

This adaptation allows MBIR to be applied without change to overall regularization levels. Should we reach the case $\eta_j \leq 0$, then $\hat{\lambda}_j$ is fixed by $\bar{\lambda}_j$. The LLMMSE filter does not guarantee positivity in its output, but in practice, the small number of remaining negatives may be effectively zero-weighted in reconstruction. After filtering, the counts $\{\hat{\lambda}_j\}$ are fed into the processes in computing integral projections and corresponding weights described in Section II.

Here the estimate of $\bar{\lambda}_j$ is computed as an average over a small neighborhood, which is sufficient to provide a good approximation of local mean for the LLMMSE filter without excessively degrading spatial resolution. To our knowledge, this type of filter has not previously been applied in CT, and may be desirable only in conjunction with statistical reconstruction, where the data weighting term W in (14) helps control the residual noise not removed by the filter. Fig. 3 shows an example of an LLMMSE denoised signal,

with significantly reduced noise level and well-preserved local means. Compared to raw signal λ_j , few counts remain negative.

B. Bayesian Inference for Means of Counts

LLMMSE sinogram preprocessing implicitly places a prior distribution on sinogram data that introduces correlation among neighboring channels. Most MAP formulations of the reconstruction problem rely on the independence of sinogram data (when conditioned on the image x) and applies prior beliefs for noise suppression only in the image domain. A preprocessing that could more effectively preserve the independence of sinogram data would better fit existing MAP reconstruction methods. To that end, we introduce an alternative approach that allows use of negative-valued measurements without sinogram filtering. This technique may also better prevent positive bias in photon counts than does the LLMMSE. The process results from a Bayesian view of sinograms, but its principal effects are administered pointwise to the data; we therefore entitle it pointwise Bayesian restoration (PBR).

The basis for our MBIR formulation has been to use the best quadratic approximation of the log-likelihood we can produce, such that it can be inserted into the many available optimization routines suited to that form. We may generate, perhaps numerically, a log likelihood from (11), and would normally choose the ML value of the projection for our expansion point, but for negative counts the ML value is unbounded, as the log-likelihood has no finite maximum point. We therefore augment this procedure with the application of an *a priori* density to the mean of the Poisson process. It is hoped that a weak prior model may stabilize the log-likelihood in ultra-low-dose cases sufficiently to utilize a much greater fraction of low SNR data by using the scalar MAP value as input to the remainder of the data preparation.

1) Conjugate Prior: The conjugate prior approach [36], in which the posterior is of the same family as the prior distribution, is natural for our problem, providing a relatively simple parameterization. The prior and posterior are called “conjugate” distributions in this case. If a conditionally Poisson distribution has rate $\theta > 0$,

$$f_K(k|\theta) = \frac{\theta^k e^{-\theta}}{k!}, \quad (19)$$

its conjugate is the gamma distribution, characterized by the shape parameter $\alpha > 0$ and the rate parameter $\beta > 0$,

$$f_\Theta(\theta; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta}, \quad \theta > 0. \quad (20)$$

If the measurement is $K = k$ from the Poisson counts, θ will have the *a posteriori* PDF

$$\begin{aligned} f_\Theta(\theta|k) &= \frac{f_K(k|\theta) f_\Theta(\theta)}{f_K(k)} \\ &= \frac{(1 + \beta)^{k+\alpha}}{\Gamma(k + \alpha)} \theta^{k+\alpha-1} e^{-(1+\beta)\theta}. \end{aligned} \quad (21)$$

The posterior for the Poisson rate has the form of the gamma PDF, with shape parameter of $k + \alpha$ and rate parameter

of $1 + \beta$. ML estimation from observation of the Poisson variate provides the unbiased

$$\hat{\theta}_{\text{MLE}} = k \quad (22)$$

for a single realization. The PBR approach gives

$$\hat{\theta}_{\text{MAP}} = E[\Theta | K = k] = \frac{k + \alpha}{1 + \beta}. \quad (23)$$

2) Parameter Selection:: For non-positive measurements, $\hat{\theta}_{\text{MLE},j} = \lambda_j$ poses a problem for the $\log()$ operation in converting counts to line-integral densities. Thus, informative prior knowledge is applied to provide a positivity constraint. The true scalar *a priori* distribution of θ_j may be derived from a model of the random image X . For the present purpose, though, such a distribution for all possible X would likely be far less useful than one that adapts to some partial knowledge of the local characteristics of the projections. Similarly to the LLMMSE method in this regard, we use local statistics to select parameters for the density of a single detector count. Assuming an *a priori* distribution for θ_j in (11) is given to follow $\text{Gamma}(\alpha_j, \beta_j)$ as in (20), we choose *a priori* expectation and variance based on the local data statistics so that

$$\begin{aligned} E[\Theta_j] &= \frac{\alpha_j}{\beta_j} := \bar{\theta}_j \\ \text{Var}[\Theta_j] &= \frac{\alpha_j}{\beta_j^2} := \bar{\theta}_j + \sigma_e^2 \end{aligned}$$

where $\bar{\theta}_j$ is an approximation of the mean of Λ_j using its local average, similar to $\bar{\lambda}_j$ in (16), but with a required positive minimum $\epsilon_\theta > 0$,

$$\bar{\theta}_j = \max\{\bar{\lambda}_j, \epsilon_\theta\}. \quad (24)$$

Then, the parameters of the conjugate prior are calculated as

$$\alpha_j = \frac{\bar{\theta}_j^2}{\bar{\theta}_j + \sigma_e^2} \quad (25)$$

$$\beta_j = \frac{\bar{\theta}_j}{\bar{\theta}_j + \sigma_e^2} \quad (26)$$

Inserting α_j and β_j into (20) and evaluating the *a posteriori* mean of combined Poisson-Gaussian model of PDF in (11), the PBR estimation of Poisson mean is given by

$$\begin{aligned} \hat{\theta}_{\text{MAP},j} &= \frac{\int \theta_j f_{\Lambda_j}(\lambda_j | \theta_j) f_{\Theta_j}(\theta_j) d\theta_j}{\int f_{\Lambda_j}(\lambda_j | \theta_j) f_{\Theta_j}(\theta_j) d\theta_j} \\ &= \frac{\sum_{k=0}^{\infty} h(k; \alpha_j + 1, \beta_j, \lambda_j)}{\sum_{k=0}^{\infty} h(k; \alpha_j, \beta_j, \lambda_j)} > 0, \end{aligned} \quad (27)$$

where

$$\begin{aligned} h(k; \alpha_j, \beta_j, \lambda_j) &= \exp\left\{-\frac{(\lambda_j - k)^2}{2\sigma_e^2}\right\} \Gamma^{-1}(k+1) \\ &\quad \Gamma(k + \alpha_j)(1 + \beta_j)^{-(k + \alpha_j)}. \end{aligned}$$

(See derivations in Appendix A.)

Compared to the LLMMSE denoising method, PBR has potential advantages in retaining greater channel independence and enforcing positivity, as suggested by Fig. 3. Fig. 4 shows that relatively high-count measurements associated with small local means ($\bar{\lambda}_j$) are less suppressed with PBR than with LLMMSE denoising, illustrating the potential of the Bayesian




Fig. 4. Mapping of photon-starved measurements ranging $[-5\sigma_e, 5\sigma_e]$ using LLMMSE filter (λ_j) and PBR estimation (θ_j) respectively, with parameter $\sigma_e = 4$ and two different values of parameter λ_j .

method to preserve resolution better than LLMMSE by retaining more of these local deviations. Also, the Bayesian method appears more robust to inaccurate parameter selection in the design of the local sinogram transformation. As shown in Fig. 4, with the same perturbation on the model parameter $\bar{\lambda}_j$, the resultant PBR estimates vary less than do the LLMMSE estimates. For cases in which the observed photon count value λ_j is close to the chosen parameter $\bar{\lambda}_j$ ($\lambda_j \approx 2.5$ in Fig. 4), the respective shifts in the two methods' curves due to changing $\bar{\lambda}_j$ are approximately equivalent. However, for $\lambda_j \approx 15$, for example, the two curves for the LLMMSE method have a greater distance between them, while the PBR curves show significantly less difference. To fix the regularization strength outside photon-starved regions, we apply PBR only to low-count regions when $\bar{\lambda}_j \leq \lambda_{\text{th}}$, similar to the treatment used in LLMMSE denoising.

IV. RESULTS

We applied the proposed methods to helical CT scan data acquired on a Discovery CT750 HD scanner (GE Healthcare, Waukesha, WI). All data were composed of 984 views per rotation, with a detector array made up of 888×64 sub-detectors, having maximum collimation of 40 mm at isocenter. A large bowtie filter was used. The electronic noise level was pre-calibrated from a scan with no tube current, yielding a standard deviation of $\sigma_e = 3.9$ photon counts, assumed constant across channels. All data were water-corrected for beam hardening, but not otherwise treated for polychromatic effects. Subsequently, model-based iterative reconstruction (MBIR) results were achieved by applying the algorithm in [3], [37] with the *q*-GGMRF as *a priori* image model and non-homogeneous iterative coordinated descent [38] for optimization. The reconstruction field of view was 50 cm. Resolution of reconstructed images was 512×512 pixels, with slice thickness of 0.625 mm.

A. Phantom Study

The "low-signal phantom" (LSP) provides a highly varying attenuation profile with sufficiently low minimum counts




Fig. 5. High dosage comparison of FBP standard kernel (a) and MBIR (b) with LSP phantom. Display window width 400 HU.




Fig. 6. Low-dose LSP reconstructions with different algorithms applied on photon-starved regions. Input photon count λ_T at center of bowtie was approximately 2.4×10^4 . (a) FBP with adaptive smoothing; (b) MBIR with no low-signal processing and zero-weights for negative measurements; (c) FBP with LLMMSE before adaptive smoothing; (d) MBIR with LLMMSE and zero-weights for residual negative channels; (e) FBP with PBR before adaptive smoothing; (f) MBIR with PBR. Display window width 400 HU.

in some projection paths to test the techniques discussed in Section III. The phantom is oval-shaped in cross-section with major axis of 45 cm and minor axis of 25 cm. Two Teflon rods were inserted parallel to the z -axis in imitation of large bones. The phantom was scanned at both high (120 kVp, 835 mA) and low (120 kVp, 20 mA) dosage settings, with speed of 1.0 sec/rotation, 20 mm collimation at normalized pitch of 31/32. Fig. 5 shows comparison between FBP and MBIR at high dosage, when low-count measurements are not encountered. The same pretreatment in converting counts to line integral projections has been applied for both algorithms. As shown, MBIR produces higher quality image than FBP with apparent reduction of noise, homogeneous texture across the plane and fine details.

The low dosage experiment begins in Fig. 6 (a), where the FBP with standard kernel, employing an adaptive low-pass filter to fill negative measurements [13], is affected by both low SNR and significant bias in the horizontal,

high-attenuation paths. The high dose reference Fig. 5 (b) is essentially free of beam-hardening type artifacts using standard water-correction treatment, suggesting that local biases in low-dosage examples derive from low-signal photon-counting phenomena, when some detectors reach photon starvation level and register negative measurements. General interpolation or smoothing type approaches negatively bias projection estimates, as observed in the accompanying low-frequency dark shading artifacts along the horizontal rays of large attenuation. Statistical iterative methods, on the other hand, have the convenience of discarding negative measurements by putting zero weights on them. As shown in Fig. 6 (b), drastic reduction of noise and artifacts is achieved, compared to 6 (a), but there are remaining residual negative biases in horizontal high attenuation paths.

To improve projection estimation in low-count regions, we first considered the use of the LLMMSE filter as proposed in (17) to denoise low SNR signals. $\bar{\lambda}_j$ was computed as a local 3×3 average in order to provide only the minimum amount of necessary local smoothing, and the threshold $\lambda_{th} = 3\sigma_e$ was used across the sinogram. Thus only areas with local average counts on the order of electronic noise standard deviation were affected by the adaptive filter. As a low-signal treatment, the LLMMSE filter can be applied for both FBP and MBIR preprocessing steps. For FBP, LLMMSE was applied before adaptive smoothing in counts domain; while for MBIR, LLMMSE was followed by the $\log()$ operation, and few remaining negative channels were zero-weighted in reconstruction. As for this low-dose LSP scan, 80% of the originally registered negative measurements became positive after denoising. In Figs. 6 (c) and 6 (d), both FBP and MBIR results show significant improvement in terms of biases. For the FBP reconstruction approach, all measurements were equally weighted. Fig. 6 (c) is considerably noisier than 6 (a), especially in the areas where high density projections passed, because the adaptive smoothing was less aggressive and thus those projections have higher variances. However, due to dynamic statistical weighting in (14) and regularization model $U(x)$ in [37], MBIR obtained results with further reduced biases and relatively uniform noise behavior.

With the Bayesian inference method, the same threshold of $\lambda_{th} = 3\sigma_e$ was used to apply the MAP estimate for low-count regions. A minimum value $\epsilon_\theta = 0.01$ was assumed for the floor of the *a priori* means. The result in Fig. 6 (f) shows the advantage of using all measurements for reconstruction. Both visual and quantitative assessments favor Figs. 6 (d) and 6 (f) over Fig. 6 (b). The statistical weights play an important role in balancing the information and noise carried by each projection. On the other hand, plugging in PBR to conventional FBP preprocessing led to an increased non-homogenous noise pattern, shown in Fig. 6 (e).

The chart in Fig. 7 measures quantitative improvements of results in Figs. 6 (b), (d) and (f) at selected regions of interest (ROI's). Both proposed methods consistently outperformed the conventional zero-weighting model. Reference values were measured from Fig. 5 (b).

Line integral projection estimations at view angles with highest attenuations are plotted in Fig. 8. Of the central

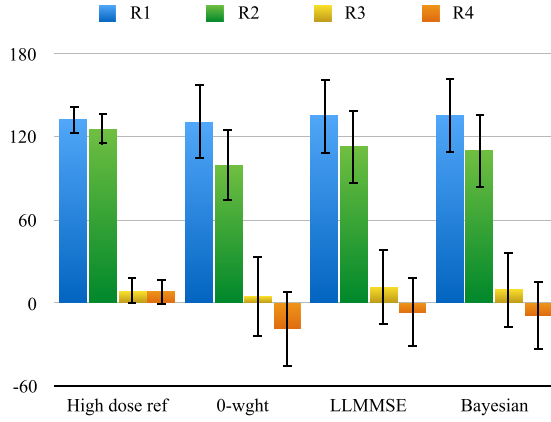


Fig. 7. Low-dose LSP chart of means (HU) and standard deviations (error bars) at selected ROI's in the images shown in Fig. 6.

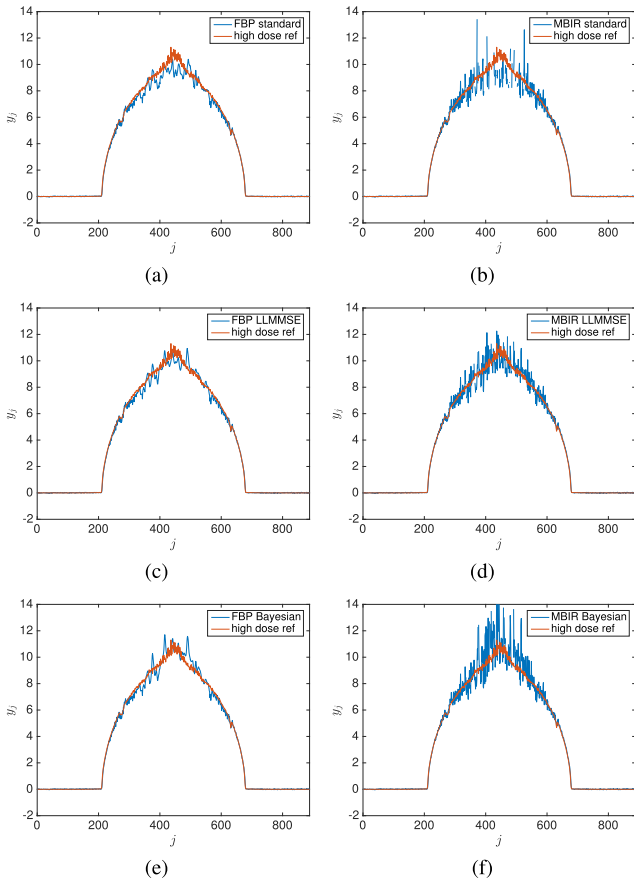


Fig. 8. Low-dose LSP line integral projection estimates with different algorithms applied on photon-starved regions. (a) Adaptive smoothing; (b) no low-signal processing and discarding negative measurements; (c) concatenating LLMMSE and adaptive smoothing; (d) LLMMSE with residual negative channels discarded; (e) concatenating Bayesian inference and adaptive smoothing; (f) Bayesian inference.

200 channels of measurements at this angle, 33% of registered counts are negative, shown as absent data points in Fig. 8 (b). The adaptive smoothing filter used in FBP pretreatment is very aggressive and caused negative bias and loss of resolution. Both LLMMSE and PBR methods, however, better preserved local means as well as channel-to-channel variations. Combining LLMMSE or PBR with FBP's adaptive smoothing




Fig. 9. Reconstruction comparison of a low-dose abdomen scan. Input photon count λ_T at center of bowtie was approximately 1.1×10^4 . (a) high-dose MBIR as reference; (b) low-dose MBIR with zero-weights on negatives; (c) low-dose MBIR with LLMMSE filter for low-signal denoising; (d) low-dose MBIR with Bayesian inference for low-signal processing. Display window width 400 HU. (Data courtesy of M.K. Kalra and A. Padole of Massachusetts General Hospital.)

is plausible, but needs further adjustment to control noise amplification.

B. Low-Dose Clinical Application

We also applied the proposed methods to clinical data, scanned at 120 kVp, 20 mA, with 40 mm collimation, at pitch ratio of 63/64. A high dosage scan (120 kVp, 400 mA) was taken of the same patient for reference in Fig. 9 (a). In this circumstance, the object was more rounded, and there is no particular path that has apparently higher attenuation than others. However, the input flux was low enough that there was a small portion of measurements falling into negative territory. In Fig. 9 (b), although there is no localized bias, the overall image suffers from loss of density. LLMMSE denoising and PBR inference achieved similar performance in preserving local means from otherwise corrupted signals. Results in Figs. 9 (c) and 9 (d) are substantially improved and closer to their true HU values, compared to the high dose reference in Fig. 9 (a). Measured means and standard deviations are shown in Fig. 10. To compute $\bar{\lambda}_j$ and $\bar{\theta}_j$ for the two methods, a 3×3 neighborhood was selected and processing was applied when $\bar{\lambda}_j \leq 5\sigma_e$.

C. Ultra-Low-Dose Clinical Application

We extended the low-signal study with an ultra-low dose clinical scan, taken at 80 kVp, 10 mA, with 40 mm

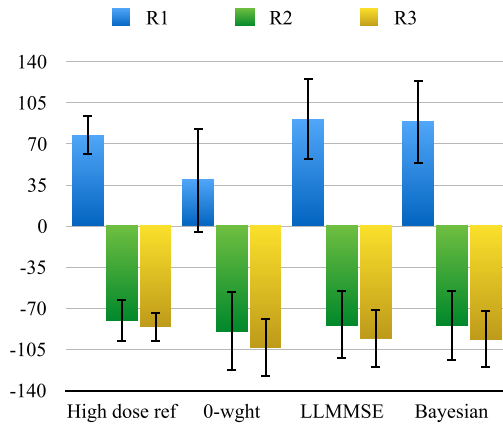


Fig. 10. Low-dose abdomen chart of means (HU) and standard deviations (error bars) at selected ROI's in the images shown in Fig. 9.

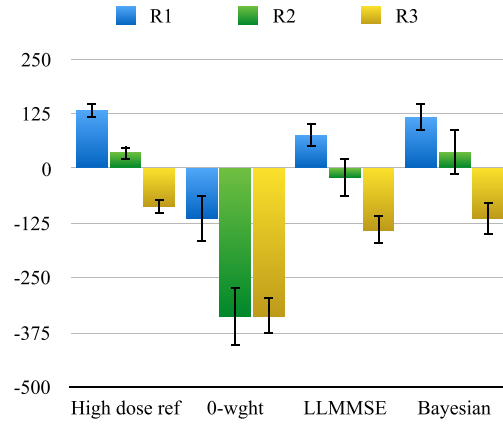


Fig. 12. Ultra-low-dose chest chart of means (HU) and standard deviations (error bars) at selected ROI's in the images shown in Fig. 11.




Fig. 11. Reconstruction comparison of a ultra-low-dose chest scan. Input photon count λ_T at center of bowtie was approximately 1.4×10^3 . (a) high-dose MBIR as reference; (b) low-dose MBIR with zero-weights on negatives; (c) low-dose MBIR with LLMMSE filter for low-signal denoising; (d) low-dose MBIR with Bayesian inference for low-signal processing. Display window width 600 HU. (Data courtesy of M.K. Kalra and A. Padole of Massachusetts General Hospital.)




Fig. 13. Zoomed-in comparison of Fig. 11 for ultra-low-dose chest scan. (a) high-dose MBIR as reference; (b) low-dose MBIR with zero-weights on negatives; (c) low-dose MBIR with LLMMSE filter for low-signal denoising; (d) low-dose MBIR with Bayesian inference for low-signal processing. Display window width 600 HU.

collimation, at pitch ratio of 33/64. A high dosage (120 kVp, 310 mA) reference scan of the same patient was taken in Fig. 11 (a) with the illustrated image nearly co-located with low-signal reconstructions. Fig. 11 (b) shows a heavily negatively biased image, suggesting that a significant number of negative measurements was registered across all view angles and zero-weighting negatives was no longer feasible to achieve informative reconstruction. Figs. 11 (c) and 11 (d) used a 5×5 neighborhood to compute $\bar{\lambda}_j$ and $\bar{\theta}_j$, and the threshold for low-count region was determined when $\bar{\lambda}_j \leq 5\sigma_e$. In spite of

the nearly complete photon starvation in some projections, the ensemble of measurements still contained significant information from the chest anatomy. The PBR result in Fig. 11 (d) outperformed LLMMSE in 11 (c) marginally, by preserving higher definition boundaries and less bias. Quantitative improvements of means and standard deviations at selected ROI's are shown in Fig. 12. A zoomed-in visual comparison is presented in Fig. 13. In practice, such extremely low-count data may occur primarily in attenuation scans for PET/CT systems, in which case image quality may be secondary in importance to appropriate local mean values.

V. CONCLUSION

The two methods presented in this paper have shown promise in dealing with very limited-count CT data. From preliminary results of handling non-positive measurements, it appears that LLMMSE denoising achieves comparable image quality to the more sophisticated Bayesian inference at moderately low-dose scenarios under statistical reconstruction framework. But in extreme cases of photon starvation, PBR showed more capability in avoiding biases. As a pretreatment in projection estimation, both methods have the potential of being applied for analytical reconstruction algorithms, but will require further adjustment. Our subsequent work with these innovations will include more extensive, quantitative phantom studies and other clinical applications.

This work has proceeded from the point of view that maintaining the simplicity of a quadratic model is important for computational savings in evaluation of derivatives of the log-likelihood. Therefore, single-pass sinogram treatments were designed, assuming that the log() operation would be performed on counts-domain data before initiation of reconstruction. The degree of importance of the quadratic objective form may vary substantially with the type of optimization followed in achieving the MBIR reconstruction. We have oriented these techniques toward sequential update methods, in which pre-log data operations could be quite costly due to non-linear operations in computation of the derivative at each voxel update. For global update methods, such as are typically used in ordered-subsets, gradient-type algorithms, the relative cost of using data retaining the pre-log form in place of least squares is potentially much lower than in the sequential case. There is likely additional benefit to statistical modeling in MBIR in exploiting this pre-log formulation. Subsequent publications will address the comparison between the methods in this paper and more accurate, more costly pre-log modeling.

APPENDIX A

The Bayesian estimation with conjugate prior in (27) can be derived as follows:

$$\begin{aligned}
 N_j &= \int_0^\infty \theta_j f_{\Lambda_j}(\lambda_j | \theta_j) f_{\Theta_j}(\theta_j) d\theta_j \\
 &= \int_0^\infty \theta_j \left(\sum_{k=0}^\infty (2\pi\sigma_e^2)^{-\frac{1}{2}} \exp\left\{-\frac{(\lambda_j - k)^2}{2\sigma_e^2}\right\} \frac{\theta_j^k e^{-\theta_j}}{k!} \right) \\
 &\quad \left(\frac{\beta_j^{\alpha_j}}{\Gamma(\alpha_j)} \theta_j^{\alpha_j-1} e^{-\beta_j\theta_j} \right) d\theta_j \\
 &= (2\pi\sigma_e^2)^{-\frac{1}{2}} \frac{\beta_j^{\alpha_j}}{\Gamma(\alpha_j)} \sum_{k=0}^\infty \left(\exp\left\{-\frac{(\lambda_j - k)^2}{2\sigma_e^2}\right\} (k!)^{-1} \right. \\
 &\quad \left. \int_0^\infty \theta_j^{k+\alpha_j-1} e^{-(1+\beta_j)\theta_j} d\theta_j \right) \\
 &= (2\pi\sigma_e^2)^{-\frac{1}{2}} \frac{\beta_j^{\alpha_j}}{\Gamma(\alpha_j)} \sum_{k=0}^\infty \left(\exp\left\{-\frac{1}{2\sigma_e^2}(\lambda_j - k)^2\right\} \Gamma^{-1}(k+1) \right. \\
 &\quad \left. \Gamma(k + \alpha_j + 1)(1 + \beta_j)^{-(k+\alpha_j+1)} \right) \\
 &= (2\pi\sigma_e^2)^{-\frac{1}{2}} \frac{\beta_j^{\alpha_j}}{\Gamma(\alpha_j)} \sum_{k=0}^\infty h(k; \alpha_j + 1, \beta_j, \lambda_j)
 \end{aligned}$$

$$\begin{aligned}
 D_j &= \int_0^\infty f_{\Lambda_j}(\lambda_j | \theta_j) f_{\Theta_j}(\theta_j) d\theta_j \\
 &= \int_0^\infty \left(\sum_{k=0}^\infty (2\pi\sigma_e^2)^{-\frac{1}{2}} \exp\left\{-\frac{(\lambda_j - k)^2}{2\sigma_e^2}\right\} \frac{\theta_j^k e^{-\theta_j}}{k!} \right) \\
 &\quad \left(\frac{\beta_j^{\alpha_j}}{\Gamma(\alpha_j)} \theta_j^{\alpha_j-1} e^{-\beta_j\theta_j} \right) d\theta_j \\
 &= (2\pi\sigma_e^2)^{-\frac{1}{2}} \frac{\beta_j^{\alpha_j}}{\Gamma(\alpha_j)} \sum_{k=0}^\infty \left(\exp\left\{-\frac{(\lambda_j - k)^2}{2\sigma_e^2}\right\} (k!)^{-1} \right. \\
 &\quad \left. \int_0^\infty \theta_j^{k+\alpha_j-1} e^{-(1+\beta_j)\theta_j} d\theta_j \right) \\
 &= (2\pi\sigma_e^2)^{-\frac{1}{2}} \frac{\beta_j^{\alpha_j}}{\Gamma(\alpha_j)} \sum_{k=0}^\infty \left(\exp\left\{-\frac{1}{2\sigma_e^2}(\lambda_j - k)^2\right\} \Gamma^{-1}(k+1) \right. \\
 &\quad \left. \Gamma(k + \alpha_j)(1 + \beta_j)^{-(k+\alpha_j)} \right) \\
 &= (2\pi\sigma_e^2)^{-\frac{1}{2}} \frac{\beta_j^{\alpha_j}}{\Gamma(\alpha_j)} \sum_{k=0}^\infty h(k; \alpha_j, \beta_j, \lambda_j)
 \end{aligned}$$

ACKNOWLEDGMENTS

The authors thank Dr. Mannudeep Kalra and Dr. Atul Padole of Massachusetts General Hospital for supplying clinical data sets.

REFERENCES

- [1] D. Polite, S. Yan, J. O'Sullivan, D. Snyder, and B. Whiting, "Implementation of alternating minimization algorithms for fully 3D CT imaging," in *Proc. SPIE*, San Jose, CA, USA, Jan. 2005, pp. 362–373.
- [2] A. Ziegler, T. Köhler, and R. Proksa, "Noise and resolution in images reconstructed with FBP and OSC algorithms for CT," *Med. Phys.*, vol. 34, pp. 585–598, Feb. 2007.
- [3] J.-B. Thibault, K. Sauer, C. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multi-slice helical CT," *Med. Phys.*, vol. 34, pp. 4526–4544, Nov. 2007.
- [4] R. C. Nelson, S. Feuerlein, and D. T. Boll, "New iterative reconstruction techniques for cardiovascular computed tomography: How do they work, and what are the advantages and disadvantages?" *J. Cardiovasc. Comput. Tomogr.*, vol. 5, no. 5, pp. 286–292, 2011.
- [5] P. J. Pickhardt *et al.*, "Abdominal CT with model-based iterative reconstruction (MBIR): Initial results of a prospective trial comparing ultralow-dose with standard-dose imaging," *Am. J. Roentgenol.*, vol. 199, no. 6, pp. 1–9, 2012.
- [6] Y. Yamada *et al.*, "Model-based iterative reconstruction technique for ultralow-dose computed tomography of the lung," *Invest. Radiol.*, vol. 47, pp. 482–489, Aug. 2012.
- [7] E. J. Stern and M. S. Frank, "CT of the lung in patients with pulmonary emphysema: Diagnosis, quantification, and correlation with pathologic and physiologic findings," *Am. J. Roentgenol.*, vol. 162, no. 4, pp. 791–798, 1994.
- [8] K. Satoh *et al.*, "CT assessment of subtypes of pulmonary emphysema in smokers," *Chest*, vol. 120, no. 3, pp. 725–729, 2001.
- [9] T. Xia, A. Alessio, B. D. Man, R. Manjeshwar, E. Asma, and P. Kinahan, "Ultra-low dose CT attenuation correction for PET/CT," *Phys. Med. Biol.*, vol. 57, no. 2, pp. 309–328, 2012.
- [10] W. A. Kalender, R. Hobel, and J. Ebersberger, "Reduction of CT artifacts caused by metallic implants," *Radiology*, vol. 164, pp. 576–577, Aug. 1987.
- [11] E. Meyer, R. Raupach, M. Lell, B. Schmidt, and M. Kachelrieß, "Normalized metal artifact reduction (NMAR) in computed tomography," *Med. Phys.*, vol. 37, pp. 5482–5493, Oct. 2010.
- [12] C. Lemmens, D. Faul, and J. Nuyts, "Suppression of metal artifacts in CT using a reconstruction procedure that combines MAP and projection completion," *IEEE Trans. Med. Imag.*, vol. 28, no. 2, pp. 250–260, Feb. 2009.

- [13] J. Hsieh, "Adaptive streak artifact reduction in computed tomography resulting from excessive X-ray photon noise," *Med. Phys.*, vol. 25, pp. 2139–2147, Nov. 1998.
- [14] M. Kachelriess, O. Watzke, and W. A. Kalender, "Generalized multi-dimensional adaptive filtering for conventional and spiral single-slice, multi-slice, and cone-beam CT," *Med. Phys.*, vol. 28, pp. 475–490, Apr. 2001.
- [15] T. Li *et al.*, "Nonlinear sinogram smoothing for low-dose X-ray CT," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 10, pp. 2505–2513, Oct. 2004.
- [16] P. J. La Rivière, "Penalized-likelihood sinogram smoothing for low-dose CT," *Med. Phys.*, vol. 32, pp. 1676–1683, Jun. 2005.
- [17] J. Wang, T. Lu, H. Lu, and Z. Liang, "Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose X-ray computed tomography," *IEEE Trans. Med. Imag.*, vol. 25, no. 10, pp. 1272–1283, Oct. 2006.
- [18] J. Hsieh, *Computed Tomography, Principles, Design, Artifacts, and New Development*, 2nd ed. Bellingham, WA, USA: SPIE Press, 2009.
- [19] I. Mori, Y. Machida, M. Osanai, and K. Iinuma, "Photon starvation artifacts of X-ray CT: their true cause and a solution," *Radiol. Phys. Technol.*, vol. 6, pp. 130–141, Jan. 2013.
- [20] O. Demirkaya, "Reduction of noise and image artifacts in computed tomography by nonlinear filtration of the projection images," *Proc. SPIE*, vol. 4322, pp. 917–923, Feb. 2001.
- [21] G.-F. Rust, V. Aurich, and M. Reiser, "Noise/dose reduction and image improvements in screening virtual colonoscopy with tube currents of 20 mAs with nonlinear Gaussian filter chains," *Proc. SPIE*, vol. 4683, pp. 1–12, Feb. 2002.
- [22] S. M. Kim, A. M. Alessio, D. S. Perlmutter, J.-B. Thibault, B. De Man, and P. E. Kinahan, "Analysis of statistical models for iterative reconstruction of extremely low-dose CT data," in *Proc. IEEE Nucl. Sci. Symp. Med. Imag. Conf.*, Seattle, WA, USA, Nov. 2014, pp. 1–4.
- [23] D. T. Kuan, A. A. Sawchuk, T. C. Strand, and P. Chavel, "Adaptive noise smoothing filter for images with signal-dependent noise," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 7, no. 3, pp. 165–177, Mar. 1985.
- [24] R. Zhang, Z. Chang, J.-B. Thibault, K. Sauer, and C. Bouman, "Statistical modeling challenges in model-based reconstruction for X-ray CT," in *Proc. 6th Comput. Imag. Conf.*, vol. 8657. Burlingame, CA, USA Feb. 2013.
- [25] Z. Chang, R. Zhang, J.-B. Thibault, K. Sauer, and C. Bouman, "Statistical X-ray computed tomography imaging from photon-starved measurements," *Proc. SPIE*, vol. 9020, p. 90200G, Mar. 2014.
- [26] K. Sauer and C. A. Bouman, "A local update strategy for iterative reconstruction from projections," *IEEE Trans. Signal Process.*, vol. 41, no. 2, pp. 534–548, Feb. 1993.
- [27] J. Fessler, "Hybrid poisson/polynomial objective functions for tomographic image reconstruction from transmission scans," *IEEE Trans. Image Process.*, vol. 4, no. 10, pp. 1439–1450, Oct. 1995.
- [28] C. A. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *IEEE Trans. Image Process.*, vol. 5, no. 3, pp. 480–492, Mar. 1996.
- [29] J. Fessler, "Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography," *IEEE Trans. Image Process.*, vol. 5, no. 3, pp. 493–506, Mar. 1996.
- [30] K. Sauer and J.-B. Thibault, "Imaging from low-intensity data," in *Bayesian Approach to Inverse Problems*, J. Idier, Ed. Hoboken, NJ, USA: Wiley, 2008, pp. 343–362.
- [31] D. L. Snyder, A. M. Hammoud, and R. L. White, "Image recovery from data acquired with a charge-coupled-device camera," *J. Opt. Soc. Am. A*, vol. 10, no. 5, pp. 1014–1023, 1993.
- [32] R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial grey scale," *Proc. Soc. Inf. Display*, vol. 17, no. 2, pp. 75–77, 1976.
- [33] J.-B. Thibault, C. A. Bouman, K. Sauer, and J. Hsieh, "A recursive filter for noise reduction in statistical iterative tomographic imaging," *Proc. SPIE*, vol. 6065, p. 60650X, Jan. 2006.
- [34] K. Krissian, C.-F. Westin, R. Kikinis, and K. G. Vosburgh, "Oriented speckle reducing anisotropic diffusion," *IEEE Trans. Image Process.*, vol. 16, no. 5, pp. 1412–1424, May 2007.
- [35] F. Argenti, G. Torricelli, and L. Alparone, "MMSE filtering of generalised signal-dependent noise in spatial and shift-invariant wavelet domains," *Signal Process.*, vol. 86, pp. 2056–2066, Aug. 2006.
- [36] H. Raiffa and R. Schlaifer, *Applied Statistical Decision Theory*, Boston, MA, USA: Clinton Press, Inc., 1961.
- [37] Z. Chang *et al.*, "Adaptive regularization for uniform noise covariance in iterative 3D CT," in *Proc. Int. Meeting Fully Three-Dimensional Image Reconstruction Radiol. Nucl. Med.*, Newport, RI, USA, May/Jun. 2015, pp. 256–259.
- [38] Z. Yu, J. B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh, "Fast model-based X-ray CT reconstruction using spatially nonhomogeneous ICD optimization," *IEEE Trans. Image Process.*, vol. 20, no. 1, pp. 161–175, Jan. 2011.