LOW-DOSE X-RAY CT RECONSTRUCTION BASED ON JOINT SINOGRAM SMOOTHING AND LEARNED DICTIONARY-BASED REPRESENTATION

Ivana Stojanovic[†], *Homer Pien*[⋆], *Synho Do*[⋆], *W. Clem Karl*[†]

[†] Boston University, Boston, MA [⋆]Massachusetts General Hospital, Boston, MA

ABSTRACT

In this paper we propose two novel image reconstruction methods for low-dose X-ray CT data. Both approaches are based on anisotropic sinogram smoothing coupled with sparse local image representation with respect to a learned overcomplete dictionary. The redundant dictionary is learned from normal-dose CT training images and encodes artifact-free image behavior. The methods differ in the details of how the redundant dictionary information is included. Efficient solution approaches to the new formulations are provided. Comparative results on simulated low-dose imagery are given. Our approach is new in how it applies learning-based dictionary techniques to low-dose CT reconstruction, in its use of high quality training data in dictionary generation, and in its incorporation of anisotropic sinogram constraints together with the dictionary-based representation.

Index Terms— Low-dose CT, sparse representation, sinogram smoothing, learning, overcomplete dictionary representations

1. INTRODUCTION

X-ray computed tomography (CT) has established itself as a primary diagnostic imaging modality since its invention. Although it has become an invaluable diagnostic tool, its use of X-ray radiation with its associated increased risk of cancer has strongly motivated the reduction of dose during data acquisition.

CT dose reduction strategies involve the optimization of numerous system parameters, including exposure control, bowtie filter, prospective gating, among others. Additionally, two commonly used strategies include the reduction of source current (mA) and energy (kV). The reduction of current reduces the number of photons at the detector, thereby increases noise. The reduction of energy decreases photon penetration through tissue and dense structures, and thereby increases the presence of artifacts. As such, it is important to design low-dose imaging algorithms which are robust to higher noise and effectively suppress artifacts, in order to preserve the diagnostic quality of the reconstructed images as dose is reduced.

Many techniques have been developed to control noise. Some methods have focused on processing in the projection data domain prior to conventional image reconstruction [1, 2, 3, 4]. Others are based in the image domain using statistical reconstruction methods [4, 5, 3], in which a cost function is optimized with respect to an unknown image. Recently, the use of sparse image domain representations with respect to adaptive local dictionaries have been proposed for transmission CT [6, 7]. In [8] dictionary learning was proposed for denoising the projection sinogram.

In this paper we consider the problem of reconstruction of images from low-dose CT projection data. We propose two formulations for accomplishing this task that impose anisotropic sinogram smoothness coupled with sparse local image representations with respect to a learned overcomplete dictionary. In our approach the overcomplete dictionary is learned from artifact-free, normal-dose imagery. The two formulations differ in the specifics of how the dictionary information is included. We provide efficient iterative solution methods for both formulations. Results are provided on simulated low-dose imagery comparing our methods to the two related approaches - the sparse image-domain representation method in [7], and to the sinogram smoothing approach in [3].

2. STATISTICAL MODEL OF X-RAY CT

The 2-D Radon transform maps a function $s(\mathbf{x})$ from the spatial or image domain to the so called projection or sinogram domain $q(\theta, t)$ according to:

$$q(\theta, t) = \int_{\ell(\theta,t)} s(\mathbf{x})\, d\ell, \tag{1}$$

where $q(\theta, t)$ is the integral of $s(\mathbf{x})$ along the line ℓ at angle θ and offset t. In practice, we use a finite-dimensional representation of the unknown image on a discrete pixel grid with $s_j = s(\mathbf{x}_j)$ and assume a set projections along a finite set of lines $q_i = q(\theta_i, t_i)$. Then by collecting the sampled variables in vectors and discretizing the integral, the Radon transform (1) can be represented in discrete form as:

$$\mathbf{q} = \mathbf{\Phi s}, \tag{2}$$

where $\mathbf{q}$ is the vector of the q_i, $\mathbf{s}$ is the vector of the s_j, and $(\mathbf{\Phi})_{ij} = \int_{\ell(\theta_i,t_i)} b_j(\mathbf{x})\, d\ell$ with $b_j(\mathbf{x})$ the j-th square pixel basis function.

In X-ray transmission tomography the projections q_i are not measured directly. Instead, the raw data are the number of photons detected after passing through the material from the X-ray source. The incident photon flux y_T can be measured and thus assumed known. Under these assumptions the measured photon count y_i corresponding to projection i is a Poisson distributed random variable with mean and variance given by $y_T \exp[-q_i]$. The photon counts at different angles and offsets are independent and thus the overall data log-likelihood function can be reduced to:

$$L(\mathbf{y}|\mathbf{q}) = -\sum_{i=1}^{I} \left[y_T e^{-(\mathbf{q})_i} + y_i (\mathbf{q})_i + \log(y_i!) \right], \tag{3}$$

where $\mathbf{y}$ denotes the vector containing all the raw measured photon counts y_i. To simplify the subsequent development we instead will use the following quadratic approximation to the log-likelihood function in (3) obtained using a second order Taylor series approximation [9, 10] about the point $r_i = \log(y_T/y_i)$:

$$L(\mathbf{y}|\mathbf{q}) \approx -\frac{1}{2} \|\mathbf{r} - \mathbf{q}\|^2_{W_d} + c(\mathbf{y}), \tag{4}$$

where $\mathbf{r}$ is the vector of r_i and $W_d = \text{diag}[y_1, y_2, \cdots, y_I]$.

3. JOINT SINOGRAM AND DICTIONARY FORMULATIONS

We have observed image domain artifacts when applying existing dictionary-based methods to low-dose tomographic data. To overcome these artifacts we propose a new approach based on minimization of a cost functional which includes constraints on sinogram behavior as well as learned dictionary representations of the image. In addition, we propose creating the dictionary on artifact free, normal-dose imagery rather than attempting to learn it directly from the low-dose imagery at the time of inversion. The overall cost function structure we propose is the following:

$$J(\mathbf{s},\mathbf{q}) = J_D(\mathbf{y},\mathbf{q}) + J_R(\mathbf{q},\mathbf{s}) + J_I(\mathbf{s}), \tag{5}$$

where the three terms correspond to data domain, Radon model, and image domain costs respectively. We discuss each of these terms next. The first term $J_D(\mathbf{y},\mathbf{q})$ captures the statistical behavior of the observed data as well as constraints on sinogram behavior:

$$J_D(\mathbf{y},\mathbf{q}) = \|\mathbf{r}-\mathbf{q}\|^2_{W_d} + \alpha \mathbf{q}^T\mathbf{W}\mathbf{q}, \tag{6}$$

where the first term comes from the log-likelihood approximation (4) and the second term corresponds to the sinogram Markov random field model used in [3, 5], which serves to anisotropically smooth the sinogram. In particular, the term is given by:

$$\mathbf{q}^T\mathbf{W}\mathbf{q} = \sum_i \sum_{m\in\mathcal{N}_i} \frac{w_{im}}{2}(\mathbf{q}_i - \mathbf{q}_m)^2, \tag{7}$$

where $\mathcal{N}_i$ is the 2-D local neighborhood of the sinogram pixel i. Different weights can be chosen to match the correlation properties of the sinogram measurements along the angular and radial directions. We use the derivative penalty with weights chosen to smooth more in the angular than the radial direction, since we expect close angular projections to be similar. Overall this cost connects the observed data to the sinogram.

The next term $J_R(\mathbf{q},\mathbf{s})$ captures the geometric line integral model of the Radon transform in (2) and serves to connect the projection domain to the image domain:

$$J_R(\mathbf{q},\mathbf{s}) = \beta \left\|\mathbf{q}-\mathbf{\Phi}\mathbf{s}\right\|_2^2. \tag{8}$$

The last term $J_I(\mathbf{s})$ is the image domain regularization term and serves to capture prior information about the behavior of the image. For this term we use a sparse reconstruction constraint in an appropriate learned and redundant dictionary $\mathbf{D}$ [11]. First, a dictionary is trained on a set of overlapping patches obtained from similar images. In our method we obtain this dictionary from a normal-dose set of data using the KSVD algorithm of [11]. By using normal-dose images in creating the dictionary we aim to encode the behavior of artifact-free reconstructions. This learned dictionary is then used in two slightly different ways, which we describe next.

LDSR-SS Method

The first method we propose we term the local dictionary sparse representation with sinogram smoothing (LDSR-SS) method. In this method we construct $J_I(\mathbf{s})$ term of (5) by representing a patch of the attenuation image $\mathbf{s}$ as a sparse combination of some elements of the dictionary $\mathbf{D}$, i.e.

$$\mathbf{E}_j\mathbf{s} = \mathbf{D}\mathbf{x}_j, \qquad \|\mathbf{x}_j\|_0 \le T, \qquad j\in\mathcal{J},$$

where the representation coefficient vector $\mathbf{x}_j$ is assumed sparse as reflected in a constraint on its cardinality T. The matrix $\mathbf{E}_j$ extracts a 2-D local neighborhood of the pixel j and $\mathcal{J}$ denotes the set of pixels for which a local neighborhood is extracted. The definition of the set $\mathcal{J}$ controls the overlap between neighboring patches. With this image regularization term the overall inversion formulation can be written as:

$$\begin{aligned}\widehat{\mathbf{s}} = \underset{\mathbf{s},\,\{\mathbf{x}_j\},\,\mathbf{q}}{\arg\min} \quad & \|\mathbf{r}-\mathbf{q}\|^2_{W_d} + \alpha\mathbf{q}^T\mathbf{W}\mathbf{q} \\ & + \beta\|\mathbf{q}-\mathbf{\Phi}\mathbf{s}\|_2^2 \\ & + \gamma\sum_{j\in\mathcal{J}}\left(\|\mathbf{E}_j\mathbf{s}-\mathbf{D}\mathbf{x}_j\|_2^2 + \lambda_j\|\mathbf{x}_j\|_0\right).\end{aligned} \tag{9}$$

GDSR-SS Method

The minimization of (9) involves optimization over three variables and a local optimization process over each image patch. It is possible to combine the local patch optimizations together in a two step process of patch estimation and post-fitting image formation. This leads to our second formulation, which we term the global dictionary sparse representation with sinogram smoothing (GDSR-SS) method.

To derive this approach consider that each patch can be placed back on the large image canvas based on knowledge of its index j by simply performing the operation $\mathbf{E}_j^T(\mathbf{D}\mathbf{x}_j)$. Placing back each patch onto a large image canvas and summing we can write:

$$\sum_{j\in\mathcal{J}}\mathbf{E}_j^T(\mathbf{E}_j\mathbf{s}) = \sum_{j\in\mathcal{J}}\mathbf{E}_j^T(\mathbf{D}\mathbf{x}_j).$$

Note that $\sum_{j\in\mathcal{J}}\mathbf{E}_j^T\mathbf{E}_j$ is a diagonal matrix for which j-th diagonal element simply counts the number of patches each pixel is represented in. Assuming that each pixel is represented in at least one patch, this in an invertible matrix and can be moved to the right hand side of the equation:

$$\mathbf{s} = W_E\sum_{j\in\mathcal{J}}\mathbf{E}_j^T(\mathbf{D}\mathbf{x}_j), \qquad W_E = \left[\sum_{j\in\mathcal{J}}\mathbf{E}_j^T\mathbf{E}_j\right]^{-1}.$$

We can write this representation of $\mathbf{s}$ in a compact form as:

$$\begin{aligned}&\mathbf{s} = \widetilde{\mathbf{D}}\mathbf{x},\\ &\mathbf{x} = [\mathbf{x}_1^T, ..., \mathbf{x}_J^T]^T,\\ &\widetilde{\mathbf{D}} = W_E\,[\mathbf{E}_1^T\mathbf{D}, \mathbf{E}_2^T\mathbf{D}, ..., \mathbf{E}_N^T\mathbf{D}].\end{aligned}$$

Finally, we arrive at the following global dictionary approach:

$$\begin{aligned}\widehat{\mathbf{x}} = \underset{\mathbf{x},\,\mathbf{q}}{\arg\min} \quad & \|\mathbf{r}-\mathbf{q}\|^2_{W_d} + \alpha\mathbf{q}^T\mathbf{W}\mathbf{q} \\ & + \beta\|\mathbf{q}-\mathbf{\Phi}\widetilde{\mathbf{D}}\mathbf{x}\|_2^2 + \lambda\|\mathbf{x}\|_0, \\ \widehat{\mathbf{s}} = \widetilde{\mathbf{D}}\widehat{\mathbf{x}},\end{aligned} \tag{10}$$

where we have assumed equal weighting on all patch sparsities. Note that the last two terms in (10) are respectively the $J_R(\mathbf{q},\mathbf{s})$ term and the image domain regularization term $J_I(\mathbf{s})$ written in terms of the representation coefficients $\mathbf{x}$. By using this formulation, we effectively assumed that the local dictionary $\mathbf{D}$ learned on image patches of a similar normal-dose image sparsely and accurately represents each patch, and thus, there exists an accurate representation of the whole image $\mathbf{s}$ in the global dictionary $\widetilde{\mathbf{D}}$.

3.1. Optimization algorithms

We adopt a coordinate descent iterative minimization approach for finding the solutions of (9) and (10).

Optimization algorithm for LDSR-SS.

Initialization. Set $\mathbf{x} = 0, \mathbf{s} = 0$ and solve:

$$\widehat{\mathbf{q}} = \arg\min_{\mathbf{q}} \quad \|\mathbf{r} - \mathbf{q}\|^2_{W_d} + \alpha \mathbf{q}^T \mathbf{W} \mathbf{q},$$
$$\widehat{\mathbf{s}} = \text{FBP}(\widehat{\mathbf{q}}).$$

At each iteration solve:

Step 1. Sparse Representation. Given $\widehat{\mathbf{q}}$ and $\widehat{\mathbf{s}}$ find a sparse image representation:

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \quad \gamma \sum_{j \in \mathcal{J}} \left(\|\mathbf{E}_j \widehat{\mathbf{s}} - \mathbf{D}\mathbf{x}_j\|_2^2 + \lambda_j \|\mathbf{x}_j\|_0 \right).$$

Step 2. Image Update. Given $\widehat{\mathbf{x}}$ and $\widehat{\mathbf{q}}$ update image:

$$\widehat{\mathbf{s}} = \arg\min_{\mathbf{s}} \quad \beta \|\widehat{\mathbf{q}} - \mathbf{\Phi}\mathbf{s}\|_2^2 + \gamma \sum_{j \in \mathcal{J}} \|\mathbf{E}_j \mathbf{s} - \mathbf{D}\mathbf{x}_j\|_2^2.$$

Step 3. Sinogram Update. Given $\widehat{\mathbf{x}}$ and $\widehat{\mathbf{s}}$ update sinogram:

$$\widehat{\mathbf{q}} = \arg\min_{\mathbf{q}} \quad \|\mathbf{r} - \mathbf{q}\|^2_{W_d} + \alpha \mathbf{q}^T \mathbf{W} \mathbf{q} + \beta \|\mathbf{q} - \mathbf{\Phi}\widehat{\mathbf{s}}\|_2^2.$$

Optimization algorithm for GDSR-SS.

Initialization. Set $\mathbf{x} = 0$ and solve

$$\widehat{\mathbf{q}} = \arg\min_{\mathbf{q}} \quad \|\mathbf{r} - \mathbf{q}\|^2_{W_d} + \alpha \mathbf{q}^T \mathbf{W} \mathbf{q}.$$

At each iteration solve:

Step 1. Sparse Representation. Given $\widehat{\mathbf{q}}$ find a sparse image representation:

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \quad \beta \|\widehat{\mathbf{q}} - \mathbf{\Phi}\widetilde{\mathbf{D}}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_0.$$

Step 2. Sinogram update. For given $\mathbf{x}$ update sinogram:

$$\widehat{\mathbf{q}} = \arg\min_{\mathbf{q}} \quad \|\mathbf{r} - \mathbf{q}\|^2_{W_d} + \alpha \mathbf{q}^T \mathbf{W} \mathbf{q} + \beta \|\mathbf{q} - \mathbf{\Phi}\widetilde{\mathbf{D}}\widehat{\mathbf{x}}\|_2^2.$$

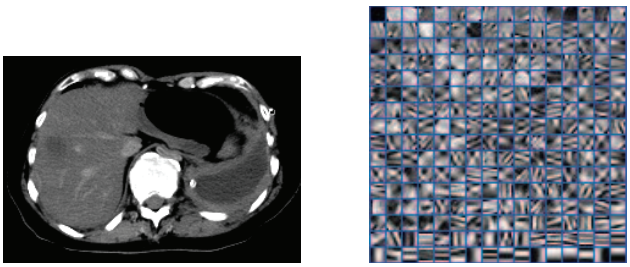
Output. The solution is: $\widehat{\mathbf{s}} = \widetilde{\mathbf{D}}\widehat{\mathbf{x}}$.

The optimization problems in the *Initialization* and *Step 3* of the LDSR-SS, as well as the problems in the *Initialization* and *Step 2* of the GDSR-SS method, are solved iteratively one pixel at the time, similarly to the approach taken in [3]. The *Step 1* of the LDSR-SS method is the sparse coding step which is solved by the Orthogonal Matching Pursuit (OMP) [11] independently for each patch. The solution of the quadratic problem in *Step 2* of the LDSR-SS is found via a conjugate gradient method. Finally, for the minimization problem in the *Step 1* of the GDSR-SS method we first replace the l_0-norm with the l_1-norm for computational purposes. The resulting l_1-least squares problem can be solved by many large-scale solvers and we use [12] in this work.

Finally, note that the *Initialization* step of the LDSR-SS optimization algorithm is the penalized reweighted least squares sinogram noise reduction method studied in [3]. Furthermore, when $\alpha = 0$ and $\mathbf{q} = \mathbf{\Phi}\mathbf{s}$, the LDSR-SS method reduces to the method studied in [7]. In the following section, along with reconstruction results of the LDSR-SS and GDSR-SS methods, we show reconstruction results for these two formulations as well.

4. EXPERIMENTAL RESULTS

We use simulated low-dose CT data to validate the proposed algorithms. We simulate dose reduction by varying the blank scan photon count parameter y_T. Our simulations are for the full-scan data when the low-dose is achieved by reducing the CT current. We assume a full scan data with regularly spaced projection angles in the range $[0, 180]$ deg with step of .25 deg in between the consecutive projection angles. The parameter w_{im} in (7) is chosen empirically such that it equals 1 for the two neighbors in the angular dimension and 0.25 for the two neighbors in the radial direction.



(a) (b)

Fig. 1. Dictionary training. (a) The FBP reconstruction from a normal data set. (b) The learned dictionary.

The global dictionary is learned from patches extracted from a normal data set image formed by FBP. The normal data set FBP image is shown in Fig.1(a), along with the learned dictionary Fig.1(b), where each patch representation was limited to less then 10 dictionary atoms.

In Fig.2 we show zoomed reconstructions from the low-dose CT scan. The parameter $\beta = 1$ for all reconstructions. The filtered backprojection (FBP) image from a noisy CT data is shown in Fig.2(a) and the LDSR-SS initialization image, i.e. the sinogram smoothing result followed by FBP, for $\alpha = y_T$ is shown in Fig.2(b). A similar formulation of the sinogram smoothing was studied in [3].

Fig.2(c) shows GDSR-SS reconstruction when the sinogram smoothing step is skipped. We modify the GDSR-SS cost function by setting $\alpha = 0$, $\mathbf{q} = \mathbf{\Phi}\widetilde{\mathbf{D}}\mathbf{x}$. The parameter $\lambda = 4.8e2$ was empirically determined to yield the best visual result. Similarly, Fig.2(d) shows LDSR-SS reconstruction when the sinogram smoothing step is skipped. We modify the LDSR-SS cost function by setting $\alpha = 0$, $\mathbf{q} = \mathbf{\Phi}\mathbf{s}$. The reduced cost function is then equivalent to the GD-SIR method studied in [7]. Step 1 is solved by limiting per patch representation to at most 5 atoms and the l_2-norm error of each representation to $1e^{-5}$.

Fig.2(e) shows GDSR-SS reconstruction for $\alpha = y_T$ and the l_2-norm error of global sparse representation $\sigma = 0.1$. The cost function of this step is solved in its BPDN formulation. Fig.2(f) shows LDSR-SS reconstruction for $\alpha = y_T$, with LDSR-SS sparse recovery Step 1 limited to at most 5 atoms and the l_2-norm error of every patch representation limited to $1e^{-5}$ for LDSR-SS. The parameter $\gamma = 1$ was empirically determined to yield the best visual result.

We observe that all methods suppress noise and improve contrast of detail features within the texture regions. However, the shape of some for small features is distorted in all reconstructions. Additionally, GDSR-SS and its modification change the appearance of the tissue the most significantly among all methods.

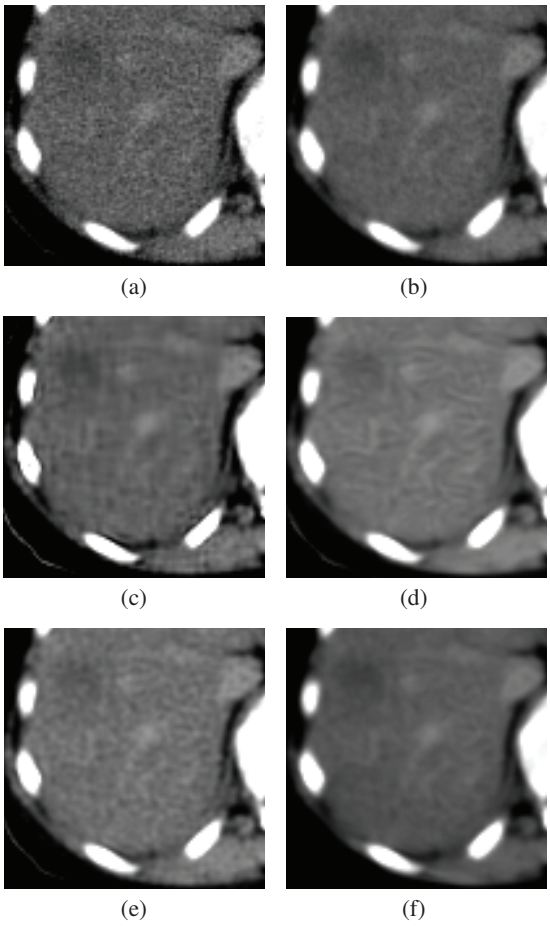


Fig. 2. Zoomed reconstruction from a low-dose data set. (a) FBP. (b) Sinogram smoothing method of [3]. (c) GDSR-SS, $\alpha = 0$, $\mathbf{q} = \mathbf{\Phi}\widetilde{\mathbf{D}}\mathbf{x}$. (d) GD-SIR method of [7]. (e) GDSR-SS. (f) LDSR-SS.

5. CONCLUSIONS AND DISCUSSION

We demonstrated several approaches to the low-dose tomography reconstruction based on a regularized inversion when prior terms are formulated both in the sinogram and the image domain. In the sinogram domain we imposed the sinogram smoothness penalty and in the image domain we utilized the sparse representation of local image patches in a dictionary learned on an image from a normal-dose dataset. Such dictionary is capable of representing fine structural information for a similar class of the low-dose CT images. We proposed the GDSR-SS and LDSR-SS formulations, where the LDSR-SS approach iterates between image and sinogram domain, in contrast to the GDSR-SS that iterates between the sinogram and the sparse representation space. The joint regularization approach produces promising results in terms of suppressing noise and preserving fine texture details. Furthermore, we contrasted the two proposed methods to the two existing approaches, [3] and [7], which are represented as the special cases of the LDSR-SS method. Our results indicate that the best reconstruction in terms of texture quality, fine details and noise suppression is obtained with the LDSR-SS.

A general concern with the learning method is that the dictionary trained on a normal-dose data set may be sensitive to the difference between the training image and the true image. The difficulty with learning from examples is that anomalies and special features may not be faithfully represented and recovered. Small image patches aim to reduce this concern as we encode local, not global behavior, which should be consistent between scans.

6. REFERENCES

[1] J. Hsieh, "Adaptive streak artifact reduction in computed tomography resulting from excessive X-ray photon noise," *Medical Physics*, vol. 25, no. 11, pp. 2139–2147, 1998.

[2] P.J. La Riviere, "Penalized-likelihood sinogram smoothing for low-dose CT," *Medical Physics*, vol. 32, no. 6, pp. 1676–1683, 2005.

[3] J. Wang, T. Li, H. Lu, and Z. Liang, "Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose X-ray computed tomography," *Medical Imaging, IEEE Transactions on*, vol. 25, no. 10, pp. 1272 –1283, oct. 2006.

[4] I.A. Elbakri and J.A. Fessler, "Statistical image reconstruction for polyenergetic X-ray computed tomography," *Medical Imaging, IEEE Transactions on*, vol. 21, no. 2, pp. 89 –99, feb. 2002.

[5] T. Li, X. Li, J. Wang, J. Wen, H. Lu, J. Hsieh, and Z. Liang, "Nonlinear sinogram smoothing for low-dose X-ray CT," *Nuclear Science, IEEE Transactions on*, vol. 51, no. 5, pp. 2505 – 2513, oct. 2004.

[6] H.Y. Liao and G. Sapiro, "Sparse representations for limited data tomography," in *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on*, may 2008, pp. 1375 –1378.

[7] Q. Xu, H. Yu, X. Mou, and G. Wang, "Dictionary learning based low-dose X-ray CT reconstruction," in *11th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, July 2011, pp. 259–263.

[8] J. Shtok, M. Elad, and M. Zibulevsky, "Sparsity-based sinogram denoising for low-dose computed tomography," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, may 2011, pp. 569 –572.

[9] K. Sauer and C. Bouman, "A local update strategy for iterative reconstruction from projections," *IEEE Trans. on Signal Proc.*, vol. 41, no. 2, pp. 534–548, Feb. 1993.

[10] C. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *IEEE Trans. on Image Proc.*, vol. 5, no. 3, pp. 480–492, Mar. 1996.

[11] M. Aharon, M. Elad, and A. Bruckstein, "K -SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *Signal Processing, IEEE Transactions on*, vol. 54, no. 11, pp. 4311 –4322, nov. 2006.

[12] E. van den Berg and M. P. Friedlander, "SPGL1: A solver for large-scale sparse reconstruction," June 2007, http://www.cs.ubc.ca/labs/scl/spgl1.