

Penalized-Likelihood Reconstruction for Sparse Data Acquisitions with Unregistered Prior Images and Compressed Sensing Penalties

J. W. Stayman^{*a}, W. Zbijewski^a, Y. Otake^b, A. Uneri^b, S. Schafer^a,
J. Lee^c, J. L. Prince^c, J. H. Siewerdsen^a

^aDept. of Biomedical Eng., Johns Hopkins University, Baltimore, MD USA 21205

^bDept. of Computer Science, Johns Hopkins University, Baltimore, MD USA 21218

^cDept. of Electrical and Computer Eng., Johns Hopkins University, Baltimore, MD USA 21218

ABSTRACT

This paper introduces a general reconstruction technique for using unregistered prior images within model-based penalized-likelihood reconstruction. The resulting estimator is implicitly defined as the maximizer of an objective composed of a likelihood term that enforces a fit to data measurements and that incorporates the heteroscedastic statistics of the tomographic problem; and a penalty term that penalizes differences from prior image. Compressed sensing (p-norm) penalties are used to allow for differences between the reconstruction and the prior. Moreover, the penalty is parameterized with registration terms that are jointly optimized as part of the reconstruction to allow for mismatched images. We apply this novel approach to synthetic data using a digital phantom as well as tomographic data derived from a cone-beam CT test bench. The test bench data includes sparse data acquisitions of a custom modifiable anthropomorphic lung phantom that can simulate lung nodule surveillance. Sparse reconstructions using this approach demonstrate the simultaneous incorporation of prior imagery and the necessary registration to utilize those priors.

INTRODUCTION

There are numerous clinical situations where repeated tomographic acquisitions are prescribed. For example in lung cancer treatment such scans are used both for diagnostics as well as image-guided procedures. Specifically, repeated surveillance scans are used diagnostically to monitor nodule size over the course of treatment. Similarly, CT may be used in a Cine mode to guide a biopsy needle. In both cases, there tends to be substantial similarity between images in the acquisition sequence. Such similarities have previously been exploited in reconstruction methods as image priors for subsequent reconstructions (most notably in PICCS reconstructions [1,2]) from sparse data acquisitions. These sparse acquisitions can be angularly undersampled, limited arcs, and/or highly truncated; providing the opportunity for significant dose reductions or decreased acquisition times.

PICCS reconstruction relies on compressed sensing norms (the L_1 norm being one particular choice) that are well-suited to ill-posed problems [3,4] due to their ability to enforce sparse reconstructions from ill-posed problems. Typically, this property is utilized by applying a so-called sparsifying transformation to the estimated imagery (like a spatial gradient), if the underlying image itself is not already a sparse entity. When prior images are utilized in a reconstruction, one would expect that the difference between a registered prior image and the new reconstruction is sparse – having significant values only in a region of change. (In some cases, additional sparsifying transforms are applied even though this difference may already be sparse.) If the prior image is not well-registered, one would expect there to be more significant differences and the prior image would have decreased utility. As such, PICCS approaches that include prior image registration have already been developed [5].

All of the approaches mentioned above seek to minimize an objective function composed of p-norm difference between the prior image and the estimate subject to a constraint that observed data matches the re-projected image estimate. This is typically applied as a linear constraint matching the log-transformed data with the re-projected estimate. While this approach is attractive since it strictly enforces the data match criterion, it does not recognize that different measurements may contain different information content. For example, it is common to presume that x-ray measurement follow a Poisson noise distribution. As such, the noise variance can be substantially space-variant with different “rays” possessing very different SNR. (Since PICCS treats all measurements equally it may be implicitly assuming a noise model that is homoscedastic in the log-transformed measurement space.)

In this paper we propose an alternate reconstruction approach that maintains a measurement model that includes the projection statistics; *and* incorporates an image prior (with possibly unregistered imagery). The specifics of this approach are outlined, and applied to simulated and real acquisitions of a lung nodule phantom in the following sections.

*web.stayman@jhu.edu; phone 1 410-955-1314; fax 410-955-1115;

METHODS

In [6], it was recognized that a total variation constraint could alternately be considered using a penalty method. In this case, one would perform an *unconstrained* optimization on an objective function that is composed of a balance of data fit terms and priors. These priors may enforce various smoothness criteria on the reconstructions, or they may incorporate previously acquired images. We utilize such a framework in the development of our penalized-likelihood approach. A brief mathematical introduction of this approach follows.

Consider a general (sparse) measurement model that relates a discretized object (μ) to mean measurements (\bar{y})

$$\bar{y} = I_0 \exp(-l) = I_0 \exp(-\mathbf{A}\mu)$$

where I_0 is the number of (unattenuated) photons for each projection ray, and \mathbf{A} is the so-called system matrix that is the discrete projection operator (for all angles and detector elements). One can presume an arbitrary noise model for the measurements; however, selecting the commonly applied Poisson model with independent measurements yields the following log-likelihood (dropping inconsequential constant terms)

$$\begin{aligned} y_i &\sim \text{Poisson} \rightarrow p(y_i|\mu) = \exp[-\bar{y}_i(\mu)] \frac{[\bar{y}_i(\mu)]^{y_i}}{y_i!} \\ L(y; \mu) &= p(y|\mu) = \prod_{i=1}^N p(y_i|\mu) = \prod_{i=1}^N \exp[-\bar{y}_i(\mu)] \frac{[\bar{y}_i(\mu)]^{y_i}}{y_i!} \\ \log L(y; \mu) &\cong \sum_{i=1}^N y_i \log[I_0 \exp(-\mathbf{A}\mu)]_i - [I_0 \exp(-\mathbf{A}\mu)]_i \end{aligned}$$

Maximization of the log-likelihood term alone would yield the maximum likelihood estimate whereas a general penalized-likelihood technique would maximize the following objective

$$\hat{\mu} = \underset{\mu}{\operatorname{argmax}} \log L(\mu; y) - R(\mu)$$

where $R(\cdot)$ is a general penalty to enforce desirable properties in the image (e.g. a roughness penalty to control noise in the reconstruction). Alternately, a PICCS-type objective can be expressed mathematically as follows

$$\hat{\mu} = \underset{\mu}{\operatorname{argmin}} [\alpha \|\Psi_1(\mu - \mu_T)\|_p + (1 - \alpha) \|\Psi_2\mu\|_p] \quad \text{s.t.} \quad \mathbf{A}\mu = \hat{l}(y)$$

where the objective is composed of two terms and one constraint. The first term enforces similarity between the attenuation estimate and a prior (or template) volume, μ_T . The use of a p-norm metric and a sparsifying operator, Ψ , allows for significant differences between the estimate and the prior (in the sparse domain). The second term is used to balance the similarity of the prior image with a general image property. Typically, Ψ is chosen to be a spatial gradient and the second term discourages noise or image roughness. Lastly, the constraint term enforces a relation of the attenuation back to the measurements using a linear equality constraint. Specifically, the projection of the attenuation estimate must match an estimate of the line integrals (typically found by normalizing and log-transforming the data). While this estimator represents a constrained optimization, we may turn this into an unconstrained optimization by replacing the constraint with a penalty function. Specifically,

$$\hat{\mu} = \underset{\mu}{\operatorname{argmin}} [\alpha \|\Psi_1(\mu - \mu_T)\|_p + (1 - \alpha) \|\Psi_2\mu\|_p + \gamma \|\mathbf{A}\mu - \hat{l}(y)\|].$$

In this case, one can select γ to be arbitrarily large to approximate the constrained objective. However, we now note that the data consistency term can be replaced by a likelihood-based data fidelity term. Specifically, we may now propose a penalized-likelihood solution (now a maximization because of the likelihood) using the following modified objective function

$$\hat{\mu} = \underset{\mu}{\operatorname{argmax}} \log L(\mu; y) - \beta_1 R_1(\mu; \mu_T) - \beta_2 R_2(\mu) \quad \text{where} \quad \begin{cases} R_1(\mu; \mu_T) = \|\Psi_1(\mu - \mu_T)\|_{p_1} \\ R_2(\mu) = \|\Psi_2\mu\|_{p_2} \end{cases}$$

where there are two penalties that may be chosen to match the analogous terms in the unconstrained PICCS objective. Two regularization parameters (β_1 , β_2) control the relative influence of each term. Similarly, we allow for different spar-

sifying operators and p-norms for each penalty. We will refer to the above reconstruction technique as PI-PLE (Prior Image Penalized Likelihood Estimation).

The above objective is suitable for image priors that are pre-registered; however, one may also jointly refine the image estimate and the registration by using the following objective that includes a registration transformation (\mathbf{W}) and optimization of the associated registration parameters (δ).

$$\{\hat{\mu}, \hat{\delta}\} = \underset{\mu, \delta}{\operatorname{argmax}} \log L(\mu; y) - \beta_1 R_1(\mu, \delta; \mu_T) - \beta_2 R_2(\mu) \quad \text{where} \quad \begin{cases} R_1(\mu, \delta; \mu_T) = \|\Psi_1(\mu - \mathbf{W}(\delta)\mu_T)\|_{p_1} \\ R_2(\mu) = \|\Psi_2\mu\|_{p_2} \end{cases}$$

We will refer to this technique as PIR-PLE (Prior Image with Registration, Penalized Likelihood Estimation). In general, the above objective requires nonconvex optimization for $p \leq 1$. Similarly, the registration operator is likely to induce local minima for general objects. The nonconvexity due to the p-norm may be handled by using the graduated nonconvexity approach [7], where one starts with $p > 1$ and reduces the value slowly during optimization. To simultaneously optimize for δ , we choose a space-alternating approach where we optimize over μ , then δ (and then repeat). While the optimization of δ is subject to local minima, we find that there is often a reasonable capture range given a good starting estimate.

The details of the optimization algorithm we utilize to solve the PIR-PLE objective are illustrated in Figure 1. Specifically for the optimization, we apply separable paraboloidal surrogates iterations [8,9] for fixed δ , followed by quasi-Newton (BFGS) or conjugate gradient iterations for fixed μ . Gradient-based optimization of δ requires differentiability of the PIR-PLE objective. Thus, one must be particularly careful in specifying various components of the PIR-PLE objective. Specifically, both the operator \mathbf{W} , and the p-norms must be handled with some care. We have chosen to implement \mathbf{W} as a kernel-based interpolator for which the gradient terms may be computed analytically. This presumes a differentiable kernel like b-splines or cubic kernels. The kernel-based approach is separable along each dimension, and highly parallelizable. We have focused on rigid transformations and have implemented both the transformation operator and the derivatives of the transformation operator with respect to the elements of δ in parallelized GPU code using nVidia's CUDA libraries. Because p-norms are not, in general, differentiable at zero, we modify the norm to be quadratic within a small neighborhood of zero and to be a shifted p-norm outside that neighborhood such that the function and derivative match at the boundary. This enforces differentiability of the entire objective.

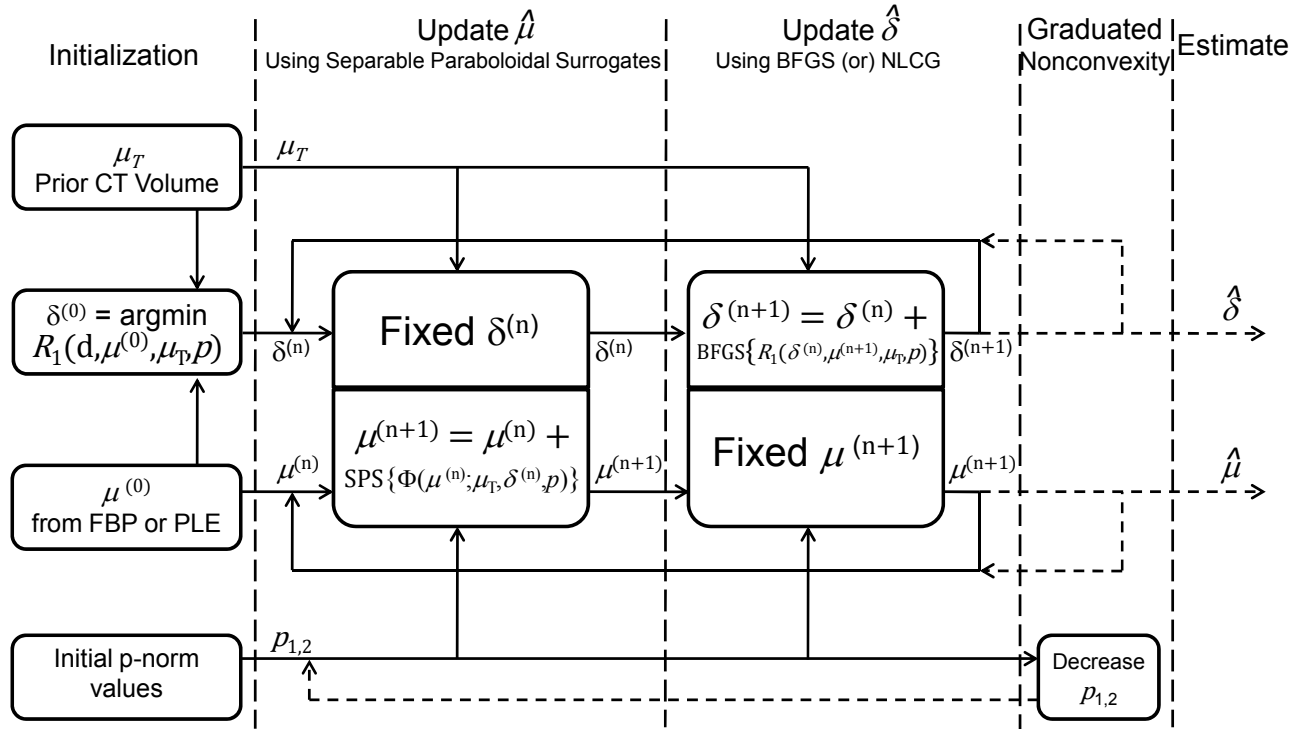


Figure 1. Flowchart representing the alternating maximization algorithm used to solve the PIR-PLE objective function for simultaneous estimation of the registration parameters (δ) and the attenuation volume (μ).

RESULTS

We applied the reconstruction approaches outlined in this paper to both simulated data and data collected from a cone-beam CT test bench. Various geometries have been investigated and presented here. We begin with a simple 2D case using simulated digital phantom data. We will show three acquisition cases: 1) 360° (360 angles) truncated tomographic data (the central 25% of the data is retained); 2) angularly undersampled data collected over 360° with 60 angles; and 3) limited-angle data over 90° (90 angles). In all cases, data were simulated with Poisson noise with 10^4 photons per detector element. The system geometry represented a system with a 55 cm source-to-detector distance, 43 cm source-to-axis distance, and 0.388 mm detector pixels (620 radial bins). Reconstructions were onto 0.25 mm square voxels (480 x 480 grid). The digital phantom is shown in Figure 2 along with an intentionally mismatched image prior, .



Figure 2. Left: True digital phantom. Right: Intentionally mismatched prior image. Note missing features including the truncated contrast ramp and line pairs, missing ellipses, and the differing diameter of the "bony" annulus on the right-hand side of the phantom.

Figure 3 compares reconstructions using penalized likelihood with no image priors (PLE); penalized likelihood with an unregistered (specifically an in-plane translation of (1.8,2.4) voxels) image prior (PI-PLE); and penalized likelihood with a jointly registered image prior (PIR-PLE). While the penalized-likelihood approaches experience significant artifacts (particularly in the region-of-interest and limited angles scans) we see greatly improved image quality when prior image penalties are included in the objective function. We note, however, that even small translation errors can potentially




Figure 3. Investigation of three imaging geometries (Truncated, Subsampled, Limited Angle) and three reconstruction methods. Left Column: Penalized-likelihood estimate with no prior image. Center Column: Penalized-likelihood with an unregistered prior image - (1.8,2.4) voxel translation. Right Column: Penalized-likelihood with an iteratively registered prior image. In this case $p_1=2$ and Ψ_1 equals the identity transform for the R_1 penalty; and $p_2=2$ and Ψ_2 equals the first-order spatial difference operator for the R_2 penalty.

degrade the image quality in the unregistered case (note the line pairs in PI-PLE vs. PIR-PLE, especially in the angular undersampling case). We note that in each of these three acquisition cases the PIR-PLE objective provides a highly accurate estimate of the in-plane translation with estimates being within a small fraction of a voxel. While the intentionally mismatched prior distorts the image data (note the oval shaped features on the right-hand side of the reconstructions that use prior image data); this is done in a predictable way - borrowing from the prior where data is missing and relying on the measurements otherwise. In particular, we see in the truncated data case that the annulus has been flattened due to the lack of horizontal frequencies in the acquisition in this region, and the reliance on the prior image whose annulus is of greater diameter. The opposite situation is found in the limited angle scenario where vertical frequencies are not obtained for this region with the consequence of reconstructing something that looks like a vertically oriented elliptical annulus.

A similar investigation was conducted using real data from a cone-beam computed tomography (CBCT) test bench. Figure 4 summarizes results from this investigation wherein an idealized lung nodule surveillance study was simulated. In this case a fully sampled CT scan of a phantom without a lung nodule is used as a prior image for reconstruction of a region-of-interest (ROI) short scan of the same (physically translated) phantom which now contains a lung nodule. This set of scans simulates the case of a patient having an initial scan with the finding of a suspicious nodule or region, who is then scheduled for a follow-up scan sometime in the future to assess any significant growth or change in that suspicious region of interest. The simplified in-plane translation of the phantom is a highly idealized version of the motion one would expect between the initial and follow-up scans.

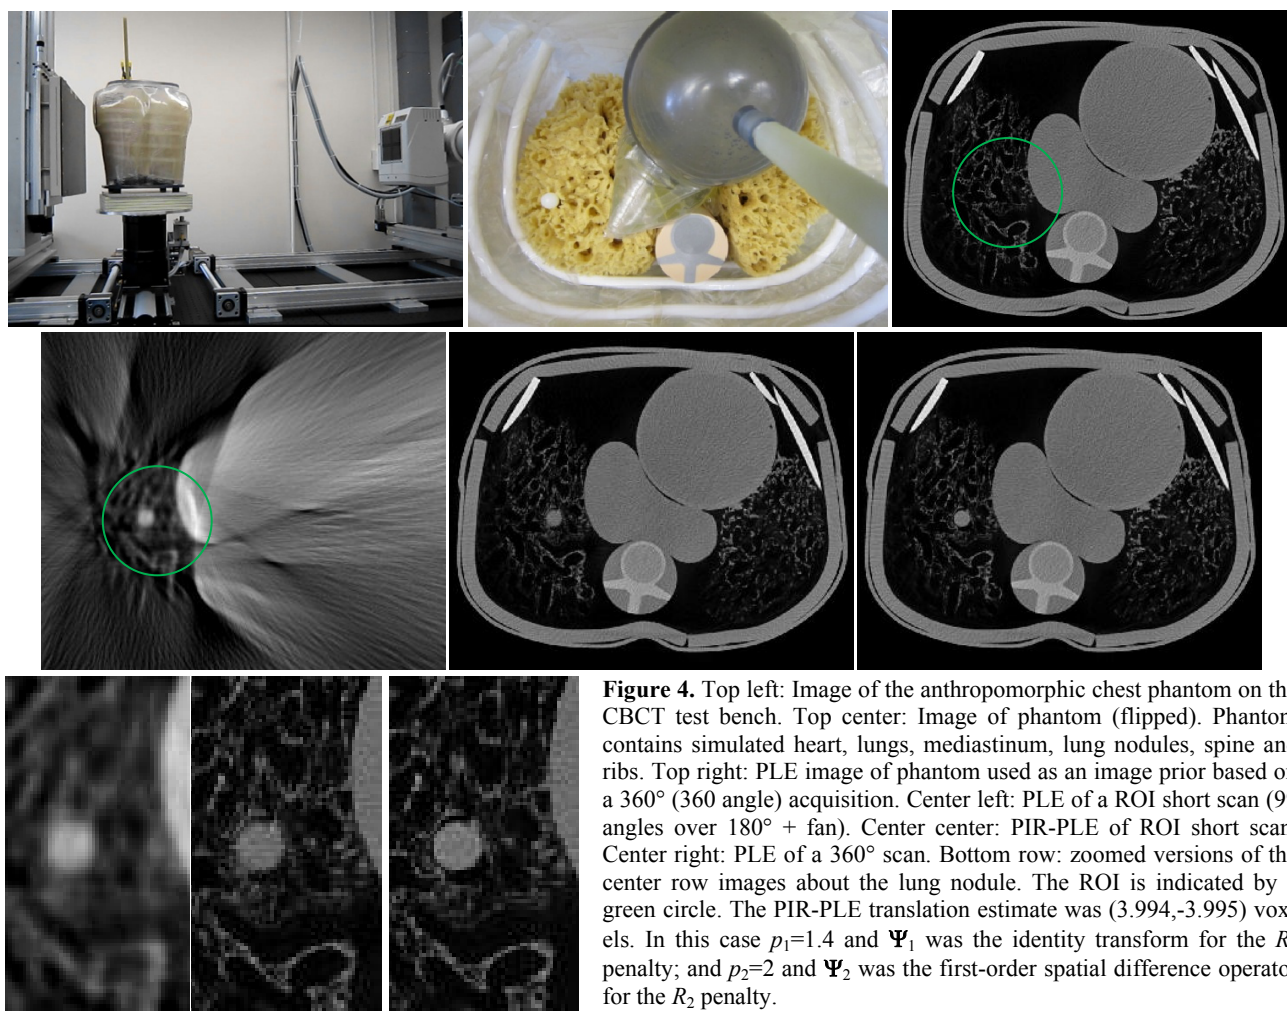


Figure 4. Top left: Image of the anthropomorphic chest phantom on the CBCT test bench. Top center: Image of phantom (flipped). Phantom contains simulated heart, lungs, mediastinum, lung nodules, spine and ribs. Top right: PLE image of phantom used as an image prior based on a 360° (360 angle) acquisition. Center left: PLE of a ROI short scan (99 angles over 180° + fan). Center center: PIR-PLE of ROI short scan. Center right: PLE of a 360° scan. Bottom row: zoomed versions of the center row images about the lung nodule. The ROI is indicated by a green circle. The PIR-PLE translation estimate was (3.994,-3.995) voxels. In this case $p_1=1.4$ and Ψ_1 was the identity transform for the R_1 penalty; and $p_2=2$ and Ψ_2 was the first-order spatial difference operator for the R_2 penalty.

For this study we used a chest phantom composed of a plastic torso shell that included tissue mimicking components for the chest wall, ribs, mediastinum, heart, and spine. Lungs were simulated using natural sponges wetted with water to approximate natural lung attenuation. A polyethylene sphere was used to simulate the lung nodule and was inserted into the interior of the natural sponge. The detector is a PaxScan 4030CB (Varian Medical Systems) with 1024, lateral, 0.388 mm detector pixels. The system geometry is 150 cm source-to-detector and 120 cm source-to-axis of rotation. Full 2D scan acquisitions were comprised of 360° rotations using 360 angles. The follow-up ROI scan was achieved by digitally subsetting complete data to be a short scan over 220° using 99 angles (1/2 the angular sampling and a short scan). A circular ROI was specified in the reconstruction volume and dynamic digital collimation was used to simulate exposure of the ROI only. Noise statistics of the unattenuated beam were used to approximate the noise equivalent number of photons (approximately 6×10^4 photons) for data scaling for the Poisson model. All data were reconstructed onto a 820 x 820, 0.4 mm voxel grid. Two test bench acquisitions were performed of the chest phantom: one with and one without the simulated lung nodule. The phantom was translated on the scanning platform by a few millimeters between scans.

The reconstruction of the initial scan using a standard PLE with quadratic penalty is shown in the top right of Figure 4. A circle indicates the ROI for the simulated follow-up scan. A reconstruction of the sparse ROI acquisition data is shown in the bottom left. We note that this scan is subject to significant artifacts outside the ROI, and within the ROI the spatial resolution is considerably reduced due to the angular undersampling. Reconstruction using the PIR-PLE objective is shown bottom center. In this case, the β_1 parameter was chosen to yield a good balance between the prior image and the data, and the β_2 parameter was chosen to yield qualitatively the same apparent spatial resolution of the reconstructed image. For comparison, a PLE reconstruction of a complete data set is also shown at the bottom right. We note excellent agreement between the PIR-PLE-based sparse data reconstruction and the fully sampled PLE reconstruction.

DISCUSSION

We have introduced a penalized-likelihood framework that allows for incorporation of prior images. Such a tool is particularly powerful in the case of sparse data acquisitions. We identified two specific objective functions that presume previously registered imagery (PI-PLE) or unregistered images that are registered simultaneously with reconstruction of sparse data (PIR-PLE). While this approach is motivated by compressed sensing methods like PICCS [1], this technique also allows for sophisticated measurement modeling that includes the statistical distribution of the data. Whereas traditional PICCS approaches enforce an equality constraint that the reconstruction match the data, the proposed approach allows for statistically motivated weightings between different measurements based on their information content. We have proposed an algorithm for solving the PIR-PLE objective using an alternating minimization approach. This algorithm relies on specific choices for the registration operator that ensures differentiability of the objective function.

Preliminary results with both simulated and test bench data suggest that the algorithm can simultaneously incorporate information from previously acquired images and estimate the registration parameters. Moreover, there appears to be a substantial improvement in image quality with a significant reduction in data acquired as compared with traditional approaches that do not incorporate the prior images. While we have identified the specific case of lung nodule surveillance as one particular application of this reconstruction framework, one might apply this approach to any scenario where sequences of (possibly misregistered) tomographic data are acquired.

REFERENCES

- [1] Chen, G.H., Tang, J. and Leng, S., "Prior image constrained compressed sensing (PICCS): A method to accurately reconstruct dynamic CT images from highly undersampled projection data sets," *Med. Phys.* **35**, 660–663 (2008)
- [2] Nett, B., Tang, J., Leng, S., and Chen, G.H., "Tomosynthesis via total variation minimization reconstruction and prior image constrained compressed sensing (PICCS) on a C-arm system," *Proc. SPIE*, **6913**, (2008)
- [3] Li, M., et al., "An accurate iterative reconstruction algorithm for sparse objects: application to 3D blood vessel reconstruction from a limited number of projections," *Phys. Med. Biol.*, **47**, 2599–2609 (2002)
- [4] Sidky, E. Y. and Pan, X., "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," *Phys. Med. Biol.*, **53**, 4777–4807 (2008)
- [5] Nett, B., Tang, J., Aagaard-Kienitz, B., Rowley, H., and Chen, G.H., "Low radiation dose C-arm cone-beam CT based on prior image constrained compressed sensing (PICCS): including compensation for image volume mismatch between multiple data acquisitions," *Proc. SPIE*, **7258**, 725–803 (2009)
- [6] Vogel, C.R., and Oman, M. E., "Iterative methods for total variation denoising," *SIAM J. Sci. Comput.*, **17**(1), 227–238 (1996)
- [7] Blake, A. and Zisserman, A., [*Visual Reconstruction*], MIT Press, Cambridge, (1987)
- [8] Erdogan, H., Fessler, J. A., "Ordered subsets algorithms for transmission tomography," *Phys. Med. Biol.*, **44**, 2835–2851, (1999)
- [9] Fessler, J.A., "Grouped coordinate descent algorithms for robust edge-preserving image restoration," *Proc SPIE Im. Recon. and Restor. II*, **3170**, 184–194 (1997)