

Prospective Regularization Analysis and Design for Prior-Image-Based Reconstruction of X-ray CT

Hao Zhang, Hao Dang, Grace J. Gang, and J. Webster Stayman

Abstract— Prior-image-based reconstruction (PIBR) methods, which incorporate a high-quality patient-specific prior image into the reconstruction of subsequent low-dose CT acquisitions, have demonstrated great potential to dramatically reduce data fidelity requirements while maintaining or improving image quality. However, one challenge with the PIBR methods is in the selection of the prior image regularization parameter which controls the balance between information from current measurements and information from the prior image. Too little prior information yields few improvements for PIBR, and too much prior information can lead to PIBR results too similar to the prior image obscuring or misrepresenting features in the reconstruction. While exhaustive parameter searches can be used to establish prior image regularization strength, this process can be time consuming (involving a series of iterative reconstructions) and particular settings may not generalize for different acquisition protocols, anatomical sites, patient sizes, etc. Moreover, optimal regularization strategies can be dependent on the location within the object further complicating selection.

In this work, we propose a novel approach for prospective analysis of PIBR. The methodology can be used to determine prior image regularization strength to admit specific anatomical changes without the need to perform iterative reconstructions in advance. The same basic methodology can also be used to prescribe uniform (shift-invariant) admission of change throughout the entire imaging field of view. The proposed predictive analytical approach was investigated in two phantom studies, and compared with the results from exhaustive search based on numerous iterative reconstructions. The experimental results show that the proposed analytical approach has high accuracy in predicting the admission of specific anatomical features, allowing for prospective determination of the prior image regularization parameter.

I. INTRODUCTION

SEQUENTIAL CT studies are common in many clinical applications such as disease monitoring, image-guided radiotherapy, and image-guided surgeries. Prior-image-based reconstruction (PIBR) has been used in sequential imaging to reduce radiation dose. Specifically, a high-quality CT scan may be first performed as reference, followed by a series of much lower exposure scans. Several studies have suggested that PIBR, through incorporation of a high-quality patient-specific prior image with subsequent low-dose acquisitions, there is great potential to improve image quality or further reduce data fidelity requirements in X-ray CT [1]–[13]. For instance, Nett et al. [2] integrated a fully-sampled prior image into the prior image constrained compressed sensing (PICCS) framework to improve subsequent image reconstruction from under-sampled acquisitions. Stayman et al. [3], [8] proposed a prior image registration,

penalized likelihood estimation (PIRPLE) approach which employs a penalized-likelihood framework and incorporates patient-specific prior image through a regularization term. Ma et al. [4], Xu and Muller [6], Xu and Tsui [7], and Zhang et al. [9] explored prior image induced nonlocal means (NLM) methods to improve the subsequent low-dose CT image reconstruction. Zhang et al. [12] extracted image textures of muscle, fat, bone, lung, etc. from high-quality prior image as *a priori* knowledge for texture-preserving Bayesian reconstruction of follow-up low-dose CT images. Additionally, Pourmorteza et al. [13] developed a penalized-likelihood method that reconstructs the difference image (with respect to a prior image) directly from the measurements.

While these prior-image-based reconstruction (PIBR) methods enjoy certain success, one common challenge with them is how to control the amount of information from the prior image in PIBR. Too little prior information yields little benefit, while too much prior information can lead the PIBR may obscure important anatomical changes. In the extreme, it is possible to force some approaches to produce a near-exact copy of the prior image despite true anatomical change. Typically, the balance between data fitting and integration of prior image information is controlled by regularization parameters. Examples include the parameter α in PICCS, β_P in PIRPLE, β and h in NLM-based methods, β in texture-preserving reconstruction, etc.

In practice, one can perform a series of image reconstructions with different regularization parameter values (i.e., parameter sweep) and choose the value whose corresponding reconstruction optimizing a certain image quality metric. However, there are problems with such an exhaustive search. First, such a parameter sweep can be extremely time-consuming due to a large number of iterative reconstructions. (Though recent strategies developed by Wu et al. [14] might be applied to reduce this burden by obtaining a sweep of parameters with a single iterative reconstruction.) Secondly, such parameter optimizations do not necessarily generalize to other scans. It is well-known that the performance of model-based iterative reconstruction (MBIR) approaches can be dependent on the x-ray technique, anatomical site, patient size, etc. Thus, optimization in one scenario does not necessarily extend to all imaging situations. Additionally, MBIR methods including PIBR approaches can exhibit significant variation in image quality within the field-of-view of a single study due to inherent shift-variances in the reconstruction algorithm. Such nonuniformities within a single image volume complicates not only the optimization of regularization

parameters but also the interpretation of individual images. A better understanding of imaging properties and predictive analytic expressions for imaging performance for PIBR methods would have great potential to provide consistent behavior across imaging conditions.

Prior work in the analysis of PIBR methods includes information source mapping [15] where reconstructions are (retrospectively) decomposed into two categories: features derived from measurement data, and features found in the prior image. Other work has included approximate analytic forms [16] for computationally efficient parameter optimization.

In this work, we define a bias metric associated with PIBR and use a new analysis of a specific PIBR approach – PIRPLE – to find closed-form analytic expressions for prior image regularization strength for specific levels of bias. The approach is applied in simulation studies and the ability to accurately and prospectively set regularization for specific goals including uniformity is demonstrated.

II. METHODS AND MATERIALS

A. PIRPLE framework

The PIRPLE approach [8] is a PIBR method that combines both a statistical measurement model and integration of prior images. In this work, we consider the PIRPLE objective function without registration (e.g., presuming the prior image is already registered to the current data), which may be written as:

$$\hat{\mu}_{\text{PIRPLE}} = \arg \max L(y; \mu) - \beta_R \|\Psi_R \mu\|_{p_R}^{p_R} - \beta_P \|\Psi_P (\mu - \mu_P)\|_{p_P}^{p_P} \quad (1)$$

where the first term, L , denotes the data fit term based on the log-likelihood, followed by two regularization terms. The first regularizer is a standard roughness penalty where Ψ_R denotes a pairwise voxel difference operator, and the second is a prior image penalty term that encourages similarity with a prior image, μ_P . Traditionally, one uses $p_R = p_P = 1$ because a l_1 norm penalty function encourages edge preservation with the standard penalty and encourages similarity but allows for sparse differences in the prior image penalty. In many cases Ψ_P is selected to be the identity matrix since the anatomical changes between the prior image and the current anatomy are already sparse. To ensure a differentiable objective function, we approximate the l_1 norm using a Huber function with a small δ value (10^{-4} mm^{-1} in these investigations). The two regularization parameters, β_R and β_P control the relative strengths of the image roughness and prior image penalties, respectively. With the above assumptions, the simplified objective function of PIRPLE is written as:

$$\hat{\mu}_{\text{PIRPLE}} = \arg \max L(y; \mu) - \beta_R \|\Psi_R \mu\|_1 - \beta_P \|\mu - \mu_P\|_1 \quad (2)$$

This objective function can be solved iteratively with an optimization algorithm, such as ordered subsets separable paraboloidal surrogates (OS-SPS) [17].

B. Closed-form approximation of the PIRPLE objective

The implicitly stated estimator in (2) is difficult to analyze directly. Previous analysis efforts have found approximate closed-form expressions that have helped to facilitate understanding of such implicitly defined estimators. For example, the

data fidelity term in Eq. (2) can be approximated by a weighted least-squares term using a second-order Taylor expansion of the log-likelihood function [18]. Similarly, the modified l_1 norm can also be approximated by a quadratic function as in [15], [16]. Specifically,

$$\|x\|_1^l = \sum_i f(x_i) \approx \sum_i \kappa_i(\tau_i) x_i^2 = (x)^T \mathbf{D} \{ \kappa(\tau) \} (x) \quad (3)$$

where $f(x_i)$ denotes the Huber function and $\kappa(\tau_i) x_i^2$ is a weighted quadratic approximation to f that is matched in value at $x_i = \tau_i$. These functions are defined as:

$$f(x_i) = \begin{cases} \frac{1}{2\delta} x_i^2 & |x_i| < \delta \\ |x_i| - \frac{\delta}{2} & |x_i| \geq \delta \end{cases}, \quad \kappa_i(\tau_i) = \begin{cases} \frac{1}{2\delta} & |\tau_i| < \delta \\ \frac{|\tau_i| - \delta/2}{\tau_i^2} & |\tau_i| \geq \delta \end{cases}$$

The last term in (3) shows a vector form of this approximation, where the diagonal weighting is a function of the vector, τ , which denotes the points (i.e. the image volume) about which the approximation is made.

Thus, the objective function in Eq. (2) can be approximated using a sum of weighted 2-norms:

$$\hat{\mu}_{\text{PIRPLE}} \approx \arg \min \|\mathbf{A}\mu - l\|_{\mathbf{W}}^2 + \beta_R \|\Psi_R \mu\|_{\mathbf{D}_R}^2 + \beta_P \|\mu - \mu_P\|_{\mathbf{D}_P}^2 \quad (4)$$

where l denotes the vector of line integral measurements (formed from the data y). Each 2-norm is weighted differently with $\mathbf{W} = \mathbf{D}\{y\}$, a weighting matrix with the measurements y along the diagonal; $\mathbf{D}_R = \mathbf{D}\{\kappa(\Psi_R \tilde{\mu})\}$ a weighting associated with the standard regularizer and $\mathbf{D}_P = \mathbf{D}\{\kappa(\tilde{\mu} - \mu_P)\}$ a weighting associated with the prior image penalty. For the approximation in (4) to be accurate an image estimate $\tilde{\mu}$ must chosen to be relatively close to the PIRPLE solution $\hat{\mu}_{\text{PIRPLE}}$.

Since the objective function in Eq. (4) is a quadratic form, it has the following closed-form solution:

$$\hat{\mu}_{\text{PIRPLE}} \approx (\mathbf{A}^T \mathbf{W} \mathbf{A} + \beta_R \Psi_R^T \mathbf{D}_R \Psi_R + \beta_P \mathbf{D}_P)^{-1} (\mathbf{A}^T \mathbf{W} l + \beta_P \mathbf{D}_P \mu_P) \quad (5)$$

This approximation will be critical to the following analytic development.

C. Determination of β_P in a PIRPLE reconstruction

The parameter β_P controls amount of information from the prior image to include in the reconstruction estimate. A smaller β_P restricts the amount of information from the prior image, and a larger β_P allows the use of more information from the prior image. Finding an optimal β_P is important because too little information yields no improvement and too much prior information can obscure specific features. In previous work, we have noted that the transition between different regions of imaging performance can be abrupt.

Consider an example problem where PIRPLE was applied to lung nodule surveillance. Given a prior image without a lung nodule, a low-dose sparse acquisition of patient anatomy with a nodule is acquired and reconstructed. Under the assumption of perfect registration, the only anatomical change between the two scans is the nodule, thus PIBR methods should work well. When β_P is varied, one can see reliable reconstruction of the change for low β_P , followed by a quick transition where the contrast of the change drops quickly, and then for high β_P the




Figure 1. The reconstructed change intensity ($\Delta\hat{\mu}$) versus prior image penalty strength. In previous work, we have observed that the contrast is reliably reproduced up to a certain β_P (first plateau, green region), then the change abruptly disappears (orange region), not to appear with higher β_P (red region).

nodule has completely disappeared. This abrupt and nonlinear behavior is illustrated in Figure 1. This behavior suggests a few things. First, the bias associated with a prior image penalty can be described as the fraction of the actual contrast that appears in the reconstruction. Second, analysis should likely be focused on the transition region since it will be difficult to relate accurate reconstructions of change to specific β_P values in plateau regions.

Following this idea, we can consider the abrupt transition region where a reconstructed change ($\Delta\hat{\mu}$) is estimated with half the actual contrast, e.g. $\hat{\mu} = \mu_P + \Delta\mu/2$ to attempt to find the β_P^* that would achieve that contrast. Or, more generally, we can consider where the reconstruction achieves a fractional contrast with $\hat{\mu} = \mu_P + \gamma\Delta\mu$ for $0 < \gamma < 1$. Returning to (5), and considering the case where regularization is dominated by the prior image data (e.g., $\beta_R = 0$) this implies:

$$\mu_P + \gamma \Delta \mu = (\mathbf{A}^T \mathbf{W} \mathbf{A} + \beta_P \mathbf{D}_P)^{-1} (\mathbf{A}^T \mathbf{W} l + \beta_P \mathbf{D}_P \mu_P). \quad (6)$$

Since the line integral estimates should be $l \approx \mathbf{A}(\mu_P + \Delta \mu)$, one can substitute for l and make the following manipulations

$$\begin{aligned} (\mathbf{A}^T \mathbf{W} \mathbf{A} + \beta_P \mathbf{D}_P)(\mu_P + \gamma \Delta \mu) &= \mathbf{A}^T \mathbf{W} l + \beta_P \mathbf{D}_P \mu_P \\ \gamma \mathbf{A}^T \mathbf{W} \mathbf{A} \Delta \mu + \gamma \beta_P \mathbf{D}_P \Delta \mu &= \mathbf{A}^T \mathbf{W} \mathbf{A} \Delta \mu \\ \beta_P \mathbf{D}_P \Delta \mu &= \frac{(1-\gamma)}{\gamma} \mathbf{A}^T \mathbf{W} \mathbf{A} \Delta \mu \end{aligned} \quad (7)$$

Returning to (3) and presuming that $\delta \approx 0$ and that the approximation point $\tilde{\mu} = \mu_P + \gamma \Delta \mu$, the same as the reconstruction target, we find that

$$\mathbf{D}_P \approx \mathbf{D} \left\{ \frac{1}{|\tilde{\mu} - \mu_P|} \right\} = \mathbf{D} \left\{ \frac{1}{|\gamma \Delta \mu|} \right\}.$$

Thus, (7) becomes

$$\beta_P \bar{\mathbf{I}} = (1 - \gamma) \mathbf{D}^{-1} \{ \text{sign}(\Delta \mu) \} \mathbf{A}^T \mathbf{W} \mathbf{A} \Delta \mu \quad (8)$$

If one considers a particular $\Delta \mu$ to guarantee in the reconstruction, (8) suggests that we can compute an ideal β_P^* that achieves the fractional contrast given by γ . Since (8) is a vector equation,

we expect that any design will only hold locally near the change $\Delta \mu$. Focusing only on nonnegative changes simplifies the above expression. If we consider a nonnegative change $\Delta \mu(j)$ to be centered at location j , we may use the following as a regularization design objective:

$$\beta_P^* = (1 - \gamma) [\mathbf{A}^T \mathbf{W} \mathbf{A} \Delta \mu(j)]_j \quad (9)$$

where $[\cdot]_j$ returns the j^{th} element of the vector argument. Note that this regularization design depends on the change (including its contrast and location), the geometry (\mathbf{A}), and the measurement statistics (\mathbf{W}).

D. Shift-variant prior image penalty design

One might also use (9) to design a shift-variant $\beta_{P,j}^*$ map that enforces a specific change contrast for all locations:

$$\forall_j \beta_{P,j}^* = (1 - \gamma) [\mathbf{A}^T \mathbf{W} \mathbf{A} \Delta \mu(j)]_j \quad (10)$$

This is potentially time consuming operation due to repeat projections and backprojections. However, as in [19], we can recognize that when applied to a compact change, $\mathbf{A}^T \mathbf{W} \mathbf{A} \approx \mathbf{\Lambda} \mathbf{A}^T \mathbf{A} \mathbf{\Lambda}$ where $\mathbf{\Lambda} = \mathbf{D}\{c\}$ is a diagonal matrix of aggregate certainties based on the data with

$$c_j = \sqrt{\frac{\sum_i a_{ij}^2 \gamma_i}{\sum_i a_{ij}^2}}. \quad (11)$$

When $\mathbf{A}^T \mathbf{A}$ is shift-invariant (e.g., for an evenly sampled tomographic system), this permits additional simplifications to (10). Specifically:

$$\forall_j \beta_{P,j}^* = (1 - \gamma) [\mathbf{\Lambda} \mathbf{A}^T \mathbf{A} \mathbf{\Lambda} \Delta \mu(j)]_j$$

For a (spatially) compact change, and since c_j are smooth

$$\forall_j \beta_{P,j}^* = (1 - \gamma) [\mathbf{\Lambda} \mathbf{A}^T \mathbf{A} c_j \Delta \mu(j)]_j$$

Similarly, since we are only interested in the j^{th} element of the right-handside.

$$\forall_j \beta_{P,j}^* = (1 - \gamma) c_j^2 [\mathbf{A}^T \mathbf{A} \Delta \mu(j)]_j \quad (12)$$

If, for different j , $\Delta \mu(j)$ are just shifted versions of the same function, then $[\mathbf{A}^T \mathbf{A} \Delta \mu(j)]_j$ is the same value for all j , and needs only to be computed once. Thus, an entire shift-variant $\beta_{P,j}^*$ map may be computed very efficiently.

E. Phantoms and simulation studies

Two phantoms were used in our study. The ellipse phantom in Fig. 2(a) consists of three attenuation regions and forms the prior image. For a subsequent scan, anatomical changes (small discs with the same contrast to the background) are introduced in two locations as shown in Fig. 2(b). For simulation studies, a system geometry was chosen with a 150 cm source-to-detector




Figure 2. Ellipse phantom: (a) prior image; (b) subsequent scan with two anatomical changes (indicated by yellow circles); (c) FBP reconstruction of subsequent scan from the simulated low-dose projection data.




Figure 3. Torso phantom: (a) prior image generated from an axial slice of a CT scan of a cadaver; (b) subsequent scan with two simulated lung nodules (indicated by yellow circles); (c) FBP reconstruction of subsequent scan from the simulated low-dose projection data.




Figure 4. PIRPLE reconstruction for the ellipse phantom from simulated low-dose projection data, using $\beta_R = 10^{2.5}$ and five different β_P values.

distance, 122 cm source-to-axis distance, and $0.556 \times 0.556 \text{ mm}^2$ detector bin sizes. 90 projections equally distributed over 360° were acquired in the subsequent scan, and the 90 projections were simulated using 10^5 photons per (bare-beam) detector bin with Poisson noise. Figure 2(c) shows a filtered back projection (FBP) reconstruction of the simulated low-dose projection data. The display window for all the images in this paper is $[0, 0.04] \text{ mm}^{-1}$.

The torso phantom in Fig. 3(a) was generated from an axial slice of a CT scan of a cadaver. We emulated a lung nodule surveillance scenario in which two uniform discs emulating lung nodules (not present in the prior image but present in the subsequent scan) were placed in the lung as shown in Fig. 3(b). The uniform disc has a radius of 6 mm and attenuation of 0.021 mm^{-1} (i.e., 50 HU assuming 0.02 mm^{-1} water attenuation), which is typical value for a solid solitary pulmonary nodule [20]. The projections were generated using the same system geometry as for the ellipse phantom, but with only 30 projections equally distributed over 360° and 10^4 photons per detector bin with Poisson noise. The FBP reconstructed image from the simulated low-dose projection data is illustrated in Fig. 3(c).

For both phantoms we perform an exhaustive sweep of the regularization parameter β_P to find the actual contrast of the nodule changes in each phantom. We compare this baseline truth with the prediction given in (9). Similarly, we investigate the design of a shift-variant penalty as described in (12) and evaluate its ability to provide uniform admission of change.

III. EXPERIMENTAL RESULTS

A. Ellipse phantom investigations

We performed a series of PIRPLE reconstruction in Eq. (2) with the iterative OS-SPS algorithm, using $\beta_R = 10^{2.5}$ and β_P varies from 10^1 to 10^5 with a $10^{0.05}$ step size. A sampling of these reconstruction results are shown in Figure 4. It can be observed that: (1) when β_P is too small, the PIRPLE reconstruction still has streak artifacts; when β_P is too large, the PIRPLE reconstruction closely resembles the prior image and anatomical changes do not appear; (2) there is an intermediate β_P value where right disc appears while the left disc disappears – suggesting significant location dependence.

Computing average the intensity of voxels in each disc, two curves of disc intensity versus prior image regularization




Figure 5. Plot of disc average intensity versus prior image regularization strength for the ellipse phantom for two different change locations. Our proposed analytical prediction approach shows very high accuracy as compared with the exhaustive search method. The $\gamma = 1/2$ case is indicated with a dotted line.




Figure 6. PIRPLE reconstruction with the shift-variant certainty-based penalty for the ellipse phantom from simulated low-dose projection data, using $\beta_R = 10^{-0.5}$ and five different β_P values. Note the uniform reconstruction of change between the two locations regardless of β_P value.

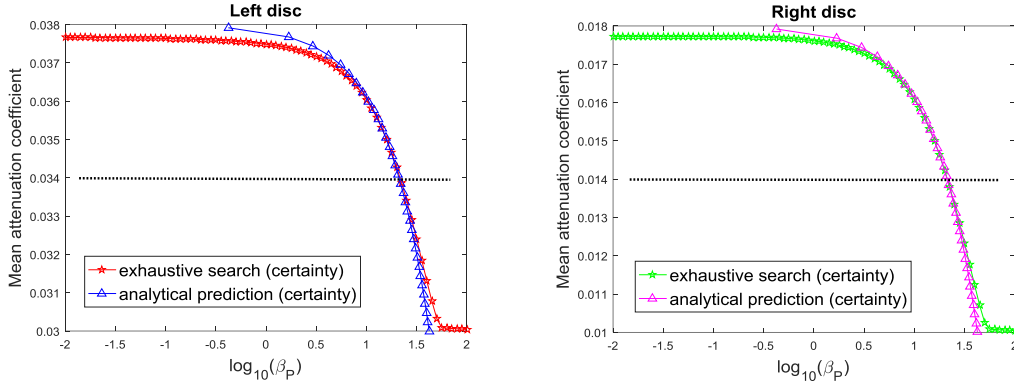


Figure 7. Plot of disc average intensity versus prior image regularization strength for the ellipse phantom when the certainty-based shift-variant penalty is applied. Again, the predictor shows good agreement with the exhaustive search. The shift-variant penalty enforces a common transition region between the two locations (centered at $\beta_P^* \approx 10^{1.3}$). Moreover the two curves have the same general form indicating uniform reconstruction of change for the two positions across β_P^* .

strength β_P can be plotted for two discs, as shown in Figure 5. Analytic predictions of the change contrast are also shown in the same set of axes. Predictions are produced by varying γ from 0 to 0.99 with a 0.03 step size and computing β_P^* for each disc, two curves of change acceptance ratio r versus prior image regularization strength can also be plotted for the two discs, as shown in Figure 5. The analytic predictions closely match the actual exhaustive search results. There is increasing deviation for values farther away from $\gamma = 1/2$. This is not unexpected since the analysis has focused on the transition region and the plateau regions represent increasingly nonlinear behavior of the PIRPLE algorithm. Note the relative shift between the

two curves indicating location-dependence. Specifically, the transition region for the left disc is at $\beta_P^* \approx 10^{3.4}$ versus $\beta_P^* \approx 10^{4.3}$ for the right disc.

The same set of experiments was reproduced for the ellipse phantom using shift-variant certainty penalty. (A certainty penalty was also used for the roughness penalty with $\beta_R = 10^{-0.5}$.) Note the range of penalty strengths was modified to account for the certainty weighting. PIRPLE reconstructions can show less location dependence effect, as illustrated in Figure 6. The two discs always appear and disappear simultaneously in PIRPLE reconstructions. Reconstructed attenuation versus penalty strength curves in Figure 7 again show good agreement between




Figure 8. PIRPLE reconstruction for the torso phantom from simulated low-dose projection data, using $\beta_R = 10^{2.5}$ and five different β_P values.




Figure 9. Plot of nodule average intensity versus prior image regularization strength for the torso phantom for the shift-invariant penalty.




Figure 10. PIRPLE reconstruction with the shift-variant certainty-based penalty for the torso phantom from simulated low-dose projection data, using $\beta_R = 10^{-0.5}$ and five different β_p values. Note the uniform reconstruction of change between the two locations regardless of β_p value.




Figure 11. Plot of nodule average intensity versus prior image regularization strength for the torso phantom for the shift-variant penalty is applied.

prediction and rote evaluation. These curves also indicate the improved uniformity with curves for both the left and right discs exhibiting a transition region centered at $\beta_p^* \approx 10^{1.3}$.

B. Torso phantom investigations

Both the shift-invariant and shift-variant regularization approaches were also applied in the torso phantom. The shift-invariant results are summarized in Figures 8 and 9, and the shift-variant results are shown in Figures 10 and 11. One can see similar behavior across these studies as in the more simple ellipse phantom. The agreement between the predictor and the exhaustive search is in good agreement. The anthropomorphic phantom shows shift-variant performance for the shift-invariant penalty suggesting the shift-variant penalty design may play an important role in clinically pertinent data sets. Like the ellipse phantom, the shift-variant design increases uniformity in the reconstruction of change at different points in the image volume.

IV. DISCUSSION AND CONCLUSION

We have proposed an analytical approach that relates prior image regularization strength to levels of bias associated with specific anatomical changes in PIRPLE reconstruction. These methods are direct and do not require any iterative reconstructions, and take into account the location and contrast of the change to be admitted in the reconstruction; the system geometry; and the measurement statistics. We have also demonstrated a computationally efficient method for specifying a shift-variant penalty that increases the uniformity of bias associated with the prior image. While the particular levels of acceptable bias are likely tied to specific imaging tasks, we have provided a general framework for predicting this bias, and selecting a regularization parameter to achieve that level. This level of control

is critically important for PIBR approaches that have the potential to obscure true features through over-regularization. This work provides the tools to prevent this in a prospective and quantitative fashion.

The development of the predictors in this work follows a similar path to the work performed in [15] which considered other PIBR methods (e.g., PICCS) using the same underlying approach. We expect that this work can be similarly generalized for other PIBR approaches.

ACKNOWLEDGMENT

This work was supported in part by an academic-industry partnership with Elekta AB.

REFERENCES

- [1] H. Yu, S. Zhao, E. A. Hoffman, and G. Wang, "Ultra-low Dose Lung CT Perfusion Regularized by a Previous Scan," *Acad. Radiol.*, vol. 16, no. 3, pp. 363–373, 2009.
- [2] B. Nett, J. Tang, B. Aagaard-Kienitz, H. Rowley, and G.-H. Chen, "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 Medical Imaging*, vol. 7258, p. 725803, 2009.
- [3] J. W. Stayman, W. Zbijewski, Y. Otake, A. Uneri, S. Schafer, J. Lee, J. L. Prince, and J. H. Siewerdsen, "Penalized-likelihood reconstruction for sparse data acquisitions with unregistered prior images and compressed sensing penalties," *Proc. SPIE Medical Imaging*, vol. 7961, p. 79611L, 2011.
- [4] J. Ma, J. Huang, Q. Feng, H. Zhang, H. Lu, Z. Liang, and W. Chen, "Low-dose computed tomography image restoration using previous normal-dose scan.," *Med. Phys.*, vol. 38, no. 10, pp. 5713–31, Oct. 2011.
- [5] H. Lee, L. Xing, R. Davidi, R. Li, J. Qian, and R. Lee, "Improved compressed sensing-based cone-beam CT reconstruction using adaptive prior image constraints," *Phys. Med. Biol.*, vol. 57, no. 8, pp. 2287–2307, Apr. 2012.
- [6] W. Xu and K. Mueller, "Efficient low-dose CT artifact mitigation using an artifact-matched prior scan.," *Med. Phys.*, vol. 39, no. 8, pp. 4748–60, Aug. 2012.
- [7] J. Xu and B. Tsui, "C-arm CT image reconstruction from sparse

- projections,” *International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, vol. 5, pp. 34–7, 2013.
- [8] J. W. Stayman, H. Dang, Y. Ding, and J. H. Siewerdsen, “PIRPLE: a penalized-likelihood framework for incorporation of prior images in CT reconstruction,” *Phys. Med. Biol.*, vol. 58, no. 21, pp. 7563–82, Nov. 2013.
- [9] H. Zhang, J. Huang, J. Ma, Z. Bian, Q. Feng, H. Lu, Z. Liang, and W. Chen, “Iterative reconstruction for x-ray computed tomography using prior-image induced nonlocal regularization,” *IEEE Trans. Biomed. Eng.*, vol. 61, no. 9, pp. 2367–78, Sep. 2014.
- [10] H. Dang, A. S. Wang, M. S. Sussman, J. H. Siewerdsen, and J. W. Stayman, “dPIRPLE: a joint estimation framework for deformable registration and penalized-likelihood CT image reconstruction using prior images,” *Phys. Med. Biol.*, vol. 59, no. 17, pp. 4799–826, Sep. 2014.
- [11] H. Zhang, H. Han, J. Wang, J. Ma, Y. Liu, W. Moore, and Z. Liang, “Deriving adaptive MRF coefficients from previous normal-dose CT scan for low-dose image reconstruction via penalized weighted least-squares minimization,” *Med. Phys.*, vol. 41, no. 4, p. 41916, Apr. 2014.
- [12] H. Zhang, H. Han, Z. Liang, Y. Hu, Y. Liu, W. Moore, J. Ma, and H. Lu, “Extracting Information From Previous Full-Dose CT Scan for Knowledge-Based Bayesian Reconstruction of Current Low-Dose CT Images,” *IEEE Trans. Med. Imaging*, vol. 35, no. 3, pp. 860–70, Mar. 2016.
- [13] A. Pourmorteza, H. Dang, J. H. Siewerdsen, and J. W. Stayman, “Reconstruction of difference in sequential CT studies using penalized likelihood estimation,” *Phys. Med. Biol.*, vol. 61, no. 5, pp. 1986–2002, Mar. 2016.
- [14] M. Wu, Q. Yang, A. Maier, and R. Fahrig, “Approximate path seeking for statistical iterative reconstruction,” *Proc. SPIE Medical Imaging*, vol. 9412, p. 94121D, 2015.
- [15] J. W. Stayman, J. L. Prince, and J. H. Siewerdsen, “Information Propagation in Prior-Image-Based Reconstruction,” *Conf. proceedings. Int. Conf. Image Form. X-Ray Comput. Tomogr.*, vol. 2012, pp. 334–338, 2012.
- [16] H. Dang, J. H. Siewerdsen, and J. W. Stayman, “Prospective regularization design in prior-image-based reconstruction,” *Phys. Med. Biol.*, vol. 60, no. 24, pp. 9515–36, Dec. 2015.
- [17] H. Erdoğan and J. A. Fessler, “Ordered subsets algorithms for transmission tomography,” *Phys. Med. Biol.*, vol. 44, no. 11, pp. 2835–51, Nov. 1999.
- [18] J. A. Fessler and W. L. Rogers, “Resolution properties of regularized image reconstruction methods,” *Technical Report No. 297*, 1995.
- [19] J. A. Fessler and W. L. Rogers, “Spatial resolution properties of penalized-likelihood image reconstruction: space-invariant tomographs,” *IEEE Trans. Image Process.*, vol. 5, no. 9, pp. 1346–58, Jan. 1996.
- [20] B. B. Tan, K. R. Flaherty, E. A. Kazerooni, M. D. Iannettoni, and American College of Chest Physicians, “The solitary pulmonary nodule,” *Chest*, vol. 123, no. 1 Suppl, p. 89S–96S, Jan. 2003.