

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

Hao Zhang, Grace J. Gang, Hao Dang, and J. Webster Stayman*, *Senior Member, IEEE*

Abstract—Prior-image-based reconstruction (PIBR) methods have demonstrated great potential for radiation dose reduction in computed tomography (CT) applications. PIBR methods take advantage of shared anatomical information between sequential scans by incorporating a patient-specific prior image into the reconstruction objective function, often as a form of regularization. However, one major challenge with PIBR methods is how to optimally determine the prior image regularization strength which balances anatomical information from the prior image with data fitting to the current measurements. Too little prior information yields limited improvements over traditional model-based iterative reconstruction (MBIR), while too much prior information can force anatomical features from the prior image not supported by the measurement data, concealing true anatomical changes. In this work, we develop quantitative measures of the bias associated with PIBR. This bias exhibits as a fractional reconstructed contrast of the difference between the prior image and current anatomy, which is quite different from traditional reconstruction biases which are typically quantified in terms of spatial resolution or artifacts. We have derived an analytical relationship between the PIBR bias and prior image regularization strength and illustrated how this relationship can be used as a predictive tool to prospectively determine prior image regularization strength to admit specific kinds of anatomical change in the reconstruction. Because bias is dependent on local statistics, we further generalized shift-variant prior image penalties which permit uniform (shift-invariant) admission of anatomical changes across the imaging field-of-view (FOV). We validated the mathematical framework in phantom studies and compared bias predictions with estimates based on brute force exhaustive evaluation using numerous iterative reconstructions across regularization values. The experimental results demonstrate that the proposed analytical approach can predict the bias-regularization relationship accurately, allowing for prospective determination of the prior image regularization strength in PIBR. Thus, the proposed approach provides an important tool for controlling image quality of PIBR methods in a reliable, robust, and efficient fashion.

I. INTRODUCTION

SEQUENTIAL CT studies are prescribed in many clinical applications including lung nodule surveillance, wherein multiple scans are performed over a period of time to identify phenotypic changes for indeterminate nodules or to find new nodules [1], [2]. The cumulative radiation dose of sequential CT studies is a concern for both clinical practitioners and patients. Model-based iterative reconstruction (MBIR) has demonstrated improved image quality for low-dose CT lung cancer screening in the clinic, as compared to the analytical filtered back-projection (FBP) methods. [3], [4] The improved trade-off between dose

and image quality from MBIR arises from more accurate modeling of projection data and measurement statistics as well as the integration of additional prior knowledge. [5]–[7] However, most MBIR treats each scan in sequence of CT studies as an individual isolated data set, neglecting a great amount of shared anatomical information between study time points.

Prior-image-based reconstruction (PIBR) provides an opportunity to use the same accurate system modeling of MBIR while incorporating a high-quality patient-specific prior image into the reconstruction of subsequent low-dose scans for additional advantages. Several studies have suggested that various PIBR methods can further reduce data fidelity requirements or improve image quality for low-dose CT [8]–[21]. For example, the prior image constrained compressed sensing (PICCS) algorithm [10] has been used in interventional imaging [11] and cardiac imaging [12] to reduce angular sampling requirements. The prior image registration, penalized likelihood estimation (PIRPLE) approach uses rigid [13] or deformable registration [14] within a joint registration-reconstruction framework to incorporate information from prior images that are not already aligned with the current data. The Reconstruction of Difference (RoD) method [16] has been used to directly estimate contrast enhancement by using a non-contrast enhanced image as a prior image to facilitate dose reduction in perfusion imaging. [17] All these methods use compressed sensing strategies like ℓ_1 -norms (or similar norms) to enforce sparsity in a domain based on the difference between prior image and current anatomy. Other PIBR methods have alternatively used approaches based on dictionary learning [7], nonlocal means [18], [19], and Markov random fields. [20], [21]

While PIBR methods are powerful tool for noise reduction, one challenge with them is how to control the amount of information integrated from the prior image so that one reliably obtains consistent and artifact-free reconstructions. In particular, too little integration prior information would yield little benefit over traditional MBIR, while too much prior information can force the PIBR image to resemble the prior image too closely – potentially obscuring features and concealing important anatomical changes. The balance between prior image information and current measurements is typically controlled by some form of regularization parameter associated with the objective function. Proper tuning of this parameter is critical to achieving good imaging performance.

Manuscript received May 14, 2018; accepted June 11, 2018. This research was supported in part by an academic-industry partnership with Elekta AB (Stockholm, Sweden), NIH grant R21 CA219608, and an AAPM research seed funding grant. H. Zhang, G. J. Gang, H. Dang, and J. W. Stayman are with the Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21205 USA (telephone: 410-955-1314, e-mail: web.stayman@jhu.edu). Copyright (c) 2017 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.




Fig. 1. An illustration of the challenges and opportunities of using PIBR. The reference true current anatomy is shown at left. The remaining images show PIRPLE reconstruction across a range of regularization strengths (β_p) of 2 mAs follow-up data where an emulated pulmonary nodule has appeared. If β_p is too low, image quality suffers and if β_p is too high, the reconstruction reproduces the prior image which has no nodule.

The challenges and opportunities that PIBR affords are illustrated in Fig. 1. This figure illustrates the use of the PIRPLE approach in lung nodule surveillance wherein a high quality prior CT image volume (no nodule) is used to improve the reconstruction of a very low exposure (2 mAs total) CT acquisition (with nodule). [14] With proper selection of the regularization parameter ($\beta_p=10^3$) a significant exposure reduction over current low-dose clinical protocols (i.e., ~20-30 mAs [22]) is achieved. However, too little prior image regularization and the solution reduces to classic MBIR with lower image quality. Too much regularization and the solution replicates the prior image.

In practice, one can determine the regularization parameter via a trial-and-error process by performing a series of image reconstructions with different regularization parameter values (i.e., exhaustive search). Then, one can select the parameter value whose corresponding reconstruction optimizes a certain quality metrics relative to a known truth image [13], [14] or simply has the best visual appearance [18]. However, there are several challenges with such an exhaustive search approach. First, the optima found via exhaustive search do not necessarily generalize to other imaging scenarios. That is, PIBR methods, as a class within MBIR, generally possess object-, data-, and location-dependent image properties.[23] Thus, optimal parameters for a particular object, x-ray technique, or anatomical location do not necessarily translate when the diagnostic target or technique is changed. [24] This suggests that an independent optimization may need to be performed for each scenario. One can attempt this by performing an exhaustive search; however, this can be extremely time-consuming due to a large number of iterative reconstructions. In MBIR efficient parameter sweeps have been effected using a path seeking algorithm [25] to generate a series of images with different regularization strengths via a single iterative reconstruction. Such methods could potentially be applied to PIBR.

Even with efficient searches, there is a more fundamental problem of how to assess image quality in PIBR without a known truth image. Ordinary MBIR methods require parameter tuning that is also potentially complex. However, classic MBIR approaches tend to exhibit a familiar noise-bias tradeoffs where bias can be measured using metrics of spatial resolution. Bias in PIBR is more complex with traditional spatial resolution and noise tradeoffs as well as biasing of specific features and anatomical structures from the prior image. Thus, due to the classical MBIR dependencies on data, object, and location, as well

as additional dependency on the prior image, and lack of suitable image quality metrics that fully encompass performance, PIBR parameter selection remains a challenge.

Several research efforts has been devoted to analyze and understand PIBR performance as well as the parameter selection. Lee et al. [26] developed an adaptive PICCS method which could detect mismatched regions between the prior image and the current anatomy during reconstruction process to avoid adverse effects of prior image. In [27], an image domain decomposition called information source mapping was proposed wherein a PIBR image separated into two additive elements comprised of features supported by measurement data and features from the prior image. Such maps may guide the selection of regularization parameters; however, the decomposition requires (computationally expensive) conjugate gradient (CG) estimation of each component and the information gained is retrospective – requiring initial reconstructions. Related work in [24] used similar analysis of PIBR and to approximate reconstruction behavior for specific changes between the prior image and the current data. That is, given specification of a specific feature difference between the prior image and the current data, an efficient parameter search was applied to guarantee accurate reconstruction of that feature. Although this methodology offered a substantial speedup (20-fold) over a traditional exhaustive search approach, the computational load is still high.

In this work, we identify a new bias metric based on the reconstructed contrast of specific changes between the prior image and the current data. We focus analysis on a specific PIBR method – PIRPLE [13] – to find a closed-form analytic expression relating prior image regularization strength and specific levels of bias. This analysis enables prediction of the bias-regularization relationship in PIBR images without explicit iterative reconstruction in advance. We use this relationship for prospective prediction and control of image properties for consistent behavior across different imaging conditions. Similarly, because the analysis is location-dependent, we explore shift-variant regularization strategies to enforce shift-invariant image properties across the imaging FOV. The proposed approach is validated in two simulation studies to illustrate its ability to accurately predict and control the PIBR performance. A preliminary derivation and application of this idea was presented in [28]. This work is a more thorough introduction, development, and validation of the proposed image property prediction and prospective regularization design methodology.

II. METHODS AND MATERIALS

A. PIRPLE framework

The PIRPLE approach [13] is a PIBR method that integrates a patient-specific prior image into a statistically motivated objective function. Specifically, the PIRPLE objective function can be written as:

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

The first term L denotes the data fidelity term which is a log-likelihood function derived from a Poisson noise model, and y is a vector of x-ray measurements over detector elements and projection angles. The second term is a standard roughness penalty where Ψ_R denotes a pairwise voxel difference operator (e.g., first order neighbors). The third term is a prior image penalty term that encourages similarity between the prior image, μ_P and the reconstruction $\hat{\mu}$. Because the prior image is not necessarily aligned with the current data, a transformation operator \mathbf{T} may be applied to permit rigid [13] or deformable [14] registration. An additional sparsifying operator Ψ_P may be applied; however, in many cases the anatomical changes between the prior image and current anatomy is already sparse. Leveraging a total variation roughness penalty and a convex, sparsifying prior image penalty, one typically chooses $p_R = p_P = 1$. (In practice, one can also approximate a l_1 norm using a Huber function with a small δ value). [13], [14]. Two regularization parameters β_R and β_P control the relative strengths of the roughness and prior image penalties, respectively.

For analysis in this work, we will adopt the following simplified PIRPLE objective with l_1 norms and without registration and without the additional sparsifying operator Ψ_P :

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

The objective function in (2) can be solved iteratively with various optimization algorithms. For reconstructions in this work, we choose the ordered subsets separable paraboloidal surrogates (OS-SPS) [29] algorithm because of its ability for parallelizable image updates.

B. Closed-form approximation of the PIRPLE objective

It is difficult to analyze the estimator in (2) directly due to the nonlinearities of both log-likelihood function and two regularizations. Previous research efforts have found approximations for the estimator which admit closed-form solution [24], [27]. Specifically, the log-likelihood function in (2) can be approximated by a weighted least-squares term using a second-order Taylor expansion [5]. Similarly, the l_1 norm (or a Huber function approximation) in (2) can also be approximated by a quadratic function [24], [27]. Consider the following set of expressions:

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

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

The l_1 norm of a vector x , $\|x\|_1$, is approximately equal to the Huber function, $f(x_i)$, and may be approximated with a quadratic




Fig. 2. Approximation of the Huber function f with a quadratic function g at an operating point τ_i for two different situations. Finding a suitable operating point τ_i is critical for the approximation.

function $g(x_i)$ as illustrated in Fig. 2. Specifically, we choose an operating point, τ_i , about which to form a weighted quadratic approximation $\kappa(\tau_i)x_i^2$ that is matched in value at $x_i = \tau_i$. The quadratic fit is made individually for each element (i) in the vector. The entire vector expression may be written compactly using a diagonal weighting $\mathbf{D}\{\kappa(\tau)\}$, that is a function of the set of all operating points about which the approximation is made.

Thus, the objective function in (2) can be approximated with a sum of three weighted 2-norms:

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

where the vector of line integral measurements, l , and is obtained from log-transformation of normalized measurement data y . The system matrix, \mathbf{A} , represents the linear projection operator. Each 2-norm in (4) is weighted differently with $\mathbf{W} = \mathbf{D}\{y\}$, a diagonal matrix with the measurements y along the diagonal; $\mathbf{D}_R = \mathbf{D}\{\kappa(\Psi_R \tilde{\mu})\}$ a weighting associated with the roughness penalty and $\mathbf{D}_P = \mathbf{D}\{\kappa(\tilde{\mu} - \mu_P)\}$ a weighting associated with the prior image penalty.

Note that the operating point, τ , for the quadratic approximations to the regularization terms is written as a function of the argument of the original norm and a particular image estimate $\tilde{\mu}$. For the approximation in (3) to be accurate, $\tilde{\mu}$ should be chosen to be close to the PIRPLE solution $\hat{\mu}$, so that the weighted 2-norm penalties in (4) match the l_1 norm penalties in (2) at the particular image volume of interest.

Since all terms in (4) have a quadratic form, one may write the following closed-form solution:

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

The approximation in (5) is critical and enables the following analytic development.

C. Determination of β_P in a PIRPLE reconstruction

The regularization parameter β_P in (2) controls amount of information from the prior image to include in the PIRPLE reconstruction. A small β_P constrains the amount of information integrated from the prior image and yields limited or no improvements over standard MBIR. In contrast, a large β_P encourages strong similarity with the prior image potentially obscuring true anatomical changes. Understanding the balance between too little and too much prior information is the key to determining β_P .

Previous studies [24] analyzing PIRPLE performance have been undertaken using combinations of exhaustive reconstruction and approximations for computational efficiency. These studies considered how the change between the prior image and




Fig. 3. The reconstructed change intensity ($\Delta\hat{\mu}$) versus prior image regularization strength. In previous work, we have observed that the contrast is reliably reproduced up to a certain β_P (first plateau, left region), then the contrast abruptly diminishes (middle region), and disappears with higher β_P (right region).

the current anatomy was reflected in the PIRPLE reconstruction. In particular, the relationship between the size, shape, and contrast of the anatomical change was quantified. Contrast was found to play a major role, and when anatomical changes were poorly represented, it was with reduced contrast. Consider observations within the context of a lung nodule surveillance scenario where PIRPLE is used to image a nodule which was not present in a prior CT exam but that appears in a follow-up study. We presume a high quality pre-registered prior image and a low exposure follow-up scan such that the only anatomical change between scans is the single nodule. As regularization strength β_P is varied, the reconstructed contrast of the nodule also varies with the general trends illustrated in Fig. 3. Specifically, for low β_P there is a broad area of reliable reconstruction of the nodule contrast. For intermediate β_P there is an abrupt transition where the contrast of the nodule drops quickly, and then for high β_P the nodule disappears completely.

This nonlinear behavior of PIRPLE suggests a few things. First, the bias associated with a prior image penalty can be described as the fraction of actual change contrast that appears in the PIRPLE reconstruction. Second, analysis of this nonlinear behavior should be focused on the abrupt transition region since there is no one-to-one correspondence between the reconstructed contrast and β_P value in the flat plateau regions.

Thus, mathematically, we may identify the transition region as where a reconstructed change ($\Delta\hat{\mu}$) is estimated with half of the true contrast change, e.g. $\hat{\mu} = \mu_P + \Delta\mu/2$. Similarly, we may find the corresponding β_P^* that would achieve the half contrast condition. While the regularization strength that achieves a half contrast reconstruction is useful in identifying the transition region between “good” and “bad” reconstructions, one may, more generally, consider where the PIBR achieves an arbitrary fractional contrast with $\hat{\mu} = \mu_P + \gamma\Delta\mu$ ($0 < \gamma < 1$). The acceptable level of fractional contrast, γ , is a design parameter that must be decided based on clinical/diagnostic task. In this work, we will provide illustrative design examples for identifying the transition region where $\gamma = 1/2$, recognizing this may not be most appropriate selection for particular diagnostic tasks. Although the roughness penalty is also an important element of PIRPLE reconstruction, consider the simpler case where regularization is dominated by the prior image data (e.g., $\beta_R \approx 0$). Based on the closed-form solution in (5) and seeking

the reconstruction with fractional reconstructed contrast, γ , we may write the following relation:

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

Recognizing that the line integral estimates should be $l \approx \mathbf{A}(\mu_P + \Delta\mu)$, one can substitute l in (6) and make the following manipulations: premultiply both sides to eliminate the matrix inverse, cancel terms containing μ_P , and isolate β_P term, so that

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

With $\mathbf{D}_P = \mathbf{D}\{\kappa(\tilde{\mu} - \mu_P)\}$ and assuming that $\delta \approx 0$, we choose an approximation point $\tilde{\mu} = \mu_P + \gamma\Delta\mu$ (the same fractional contrast reconstruction target $\hat{\mu}$) to obtain:

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

Thus, (7) yields the following analytical relationship between the prior image regularization and reconstructed contrast bias:

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

If one wants to guarantee a change contrast (γ) in the PIRPLE reconstruction, (8) suggests that we can quickly compute a specific β_P^* that achieves that goal. Since (8) is a vector equation, we expect that the prior image regularization design only holds locally for the particular change $\Delta\mu$. If we consider a nonnegative change $\Delta\mu(j)$ to be centered at location j (a negative change adds a minus to the right-hand side), we can use the following simplified relationship in a design objective:

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

which includes dependence on the imaging geometry (\mathbf{A}), data-dependence (through $\mathbf{W} = \mathbf{D}\{y\}$), change-dependence ($\Delta\mu$), and location-dependence (j) of the bias properties in PIRPLE. We note that (9) is a vector equation in which we are seeking a scalar optima. There is unlikely to be scalar value that satisfies (9) for all vector elements; however, it is most important that the equation is satisfied in the neighborhood of the change (near j). Thus, we adopt a scalar design objective based on a local least-squares fitting of β_P^* over all the voxels within the support of the change.

To use (9) for regularization design, one must have some kind of foreknowledge of the change, $\Delta\mu$. In previous work [24], design was performed by testing specific change stimuli in conjunction with the prior image. That is, if one can specify a presumed (or possible) change that one is seeking in the image, regularization can be designed to enforce a specific fractional contrast reconstruction. Note that, the change needs not actually be present for design. Regularization may be designed such that if the change is present, then it will be reliably reconstructed. Thus, this basic framework can be used to ensure reliable reconstruction of a specific anatomical change regardless of its location.

Other options exist for specification of $\Delta\mu$. For example, if one has no general knowledge of the anticipated change, one can rely directly on measurement data. That is, presuming a registered prior and current anatomy, $\mathbf{A}\Delta\mu \approx l - l_P$, where l is a vector of line integrals from current low-dose measurements and l_P is a vector of line integrals of the prior image. Thus, it is

possible to evaluate (9) from prior and measurement data directly by replacing $\mathbf{A}\Delta\mu$ with $l - l_p$. Lastly, since previous work [24] found that regularization was most highly dependent on contrast and more weakly dependent on shape and size, one can potentially perform designs using a generic change and clinically relevant contrast difference. In this work, we focus on specific presumed changes for the design process, but also consider one experiment where the line integral estimates from the data is used.

D. Shift-variant prior image penalty design

If $\Delta\mu(j)$ are just shifted versions of the same function, one can use (9) to design a shift-variant $\beta_{p,j}^*$ map that enforces a specific fractional change contrast for all locations across the field-of-view (FOV). Mathematically,

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

where $[\cdot]_j$ returns the j^{th} element of the vector argument. Thus, (10) may be used with a modified reconstruction objective that replaces the scalar regularization strength with a vector of location-dependent weights $\beta_{p,j}^*$. That is, the third term of (2) will become:

$$\|\mu - \mu_p\|_{\mathbf{D}\{\beta_{p,j}^*\}}^1 = \sum_j \beta_{p,j}^* |\mu - \mu_p|_j \quad (11)$$

to allow local control of the regularization strength.

The penalty design represented by (10) may be used to eliminate location-dependence and make PIBR reconstruction bias associated with a specific change between the prior image and the current anatomy uniform across the FOV. That is, for a given change, $\Delta\mu$, that change will be reconstructed with the same fractional contrast regardless of the location in the object. For similar lesions that have appeared on a uniform background relative to the prior image, all lesions would have the same absolute bias. While the penalty design represented by (10) may be used to eliminate location-dependence and make bias in anatomical change uniform across the FOV; this kind of design might be computationally inefficient due to repeated projections and backprojections for different locations. We may leverage previous observations that the fisher information matrix ($\mathbf{A}^T \mathbf{W} \mathbf{A}$) is concentrated near its diagonal elements. This permits the approximation $\mathbf{A}^T \mathbf{W} \mathbf{A} \approx \mathbf{\Lambda} \mathbf{A}^T \mathbf{A} \mathbf{\Lambda}$ when applied locally (e.g., the change $\Delta\mu(j)$ is relatively compact), where $\mathbf{\Lambda} = \mathbf{D}\{c\}$ is a diagonal matrix of aggregate certainties. [23] The expression for those certainties is based on the measurement data with

$$c_j = \sqrt{\sum_i [\mathbf{A}]_{ij}^2 y_i / \sum_i [\mathbf{A}]_{ij}^2}. \quad (12)$$

Therefore, (10) can be approximated as

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

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

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

Similarly, since we are only interested in the j^{th} element of the right-hand side, we can obtain:

$$\forall_j \beta_{p,j}^* \approx (1 - \gamma)c_j^2 [\mathbf{A}^T \mathbf{A} \Delta\mu(j)]_j \quad (13)$$

For an evenly sampled tomographic system such as CT, the action of $\mathbf{A}^T \mathbf{A}$ is nearly shift-invariant across the FOV, so $[\mathbf{A}^T \mathbf{A} \Delta\mu(j)]_j$ is the same value for all j and needs only to be computed once. Thus, we can obtain the c_j map from a single

computation (on the order of two backprojections). If the change $\Delta\mu(j)$ does not significantly change the statistical weights in \mathbf{W} , then the certainty map may be computed using (12) to obtain a shift-variant $\beta_{p,j}^*$ map very efficiently.

Under the above conditions we may eliminate the location-dependence of the regularization strength by integrating the certainty-based weighting from (12) directly into the PIRPLE objective as:

$$\hat{\mu} = \arg \max L(y; \mu) - \beta_R \|\Psi_R \mu\|_{\mathbf{D}_C}^1 - \beta_P \|\mu - \mu_p\|_{\mathbf{D}_C}^1 \quad (14)$$

where $\mathbf{D}_C = \mathbf{D}\{c^2\}$ is a weighting matrix of squared aggregate certainties along the diagonal. (In this case, the certainty weight is also applied to the roughness penalty to encourage more uniform smoothing as suggested in previous work [23]). Following previous analysis the regularization parameter may be found efficiently according to

$$\beta_P^* = (1 - \gamma)[\mathbf{A}^T \mathbf{A} \Delta\mu(j)]_j. \quad (15)$$

Note that this expression is location-dependent only in terms of the change, $\Delta\mu$. Thus, if $\Delta\mu$ differs only by a spatial shift (j), the modified objective function that always includes the certainty weights permits specification of a uniform *scalar* penalty strength β_P^* . We also note that the above development relies on an intrinsic response, $\mathbf{A}^T \mathbf{A}$, that is nearly shift-invariant across the FOV. This is expected for 2D axial CT geometries and permits use of the aggregate certainty approximation in (12). For short scans and cone-beam geometries (particularly near the edge of the FOV), these approximations may be less accurate. However, Cho and Fessler [30] proposed more sophisticated system response approximations that are more appropriate for these asymmetric geometries. Thus, while this work focuses on the 2D axial scenario, it should be straightforward to extend computationally efficient penalty design to cone-beam and short-scan fan-beam CT.

E. Phantoms and simulation studies

The validity of the previous analysis and the relationship between reconstructed change contrast and regularization is explored in phantom studies. Two sets of phantoms are used and




Fig. 4. Ellipse phantom: (a) prior image; (b) subsequent scan with two anatomical changes (indicated by yellow circles).




Fig. 5. Torso phantom: (a) prior image generated from an axial slice of a CT scan of a cadaver and two simulated lung nodules (indicated by yellow circles); (b) subsequent scan with nodule shrinkage.

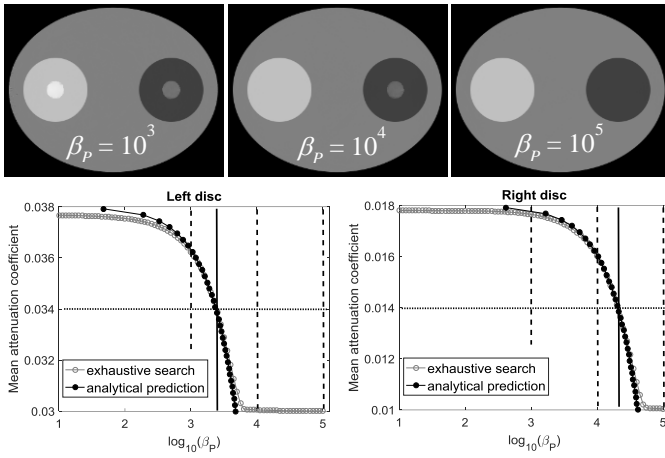


Fig. 6. Plot of disc average intensity versus prior image regularization strength for the ellipse phantom for two different change locations. Our proposed analytical prediction approach shows very high accuracy as compared with the exhaustive search method. The $\gamma = 1/2$ case is indicated with a horizontal dotted line and the corresponding regularization strength with a solid vertical line. Dotted vertical lines indicate the regularization values corresponding to the above reconstructions at $\beta_p = 10^3$, 10^4 , and 10^5 .

are shown in Fig. 4 and Fig. 5. The display window for all the images in this paper is $[0, 0.04] \text{ mm}^{-1}$ unless otherwise stated. The ellipse phantom in Fig. 4(a) consists of three attenuation regions (a background ellipse with attenuation 0.02 mm^{-1} , a dense circular insert with attenuation 0.03 mm^{-1} , and a low-density circular insert with attenuation 0.01 mm^{-1}) and is used as the prior image. For the subsequent scan, two anatomical changes (small discs with the same contrast of 0.008 mm^{-1} relative to the background) are introduced to the prior image, as shown in Fig. 4(b). The second phantom in Fig. 5(a) was generated from an axial slice of a CT scan of a cadaver torso. Two stimuli were digitally placed in the lung region to mimic irregular pulmonary nodules. A lung nodule surveillance scenario was emulated in which both nodules shrank in the subsequent scan (i.e., negative change), as shown in Fig. 5(b). The lung nodules in both images have attenuation of 0.021 mm^{-1} (typical for solid solitary pulmonary nodules [31]).

For simulation studies, a system geometry of 150 cm source-to-detector distance, 122 cm source-to-axis distance, and $0.556 \times 0.556 \text{ mm}^2$ detector bin sizes, was chosen to generate projection data for subsequent scans. Unless otherwise indicated the following parameters served as defaults: A total of 90 projections equally spaced over 360° were acquired; Poisson noise was added to the data using 10^5 and 10^4 incident photons per detector bin for the ellipse and torso phantoms, respectively; and analytical predictions used the true $\Delta\mu$ stimuli.

E.1. Exhaustive search vs. analytical prediction

For both phantoms, we compute ground truth values of the reconstructed contrast of changes via individual PIRPLE reconstructions (100 iterations and 10 subsets of OS-SPS) of the subsequent scan over an exhaustive sweep of the regularization parameter β_p . The validity of the analytical predictions is assessed via comparisons of this (ground truth) baseline with the analytical prediction given by (10) or the computationally efficient approximation in (13). These comparisons are made for the two change locations in each phantom.

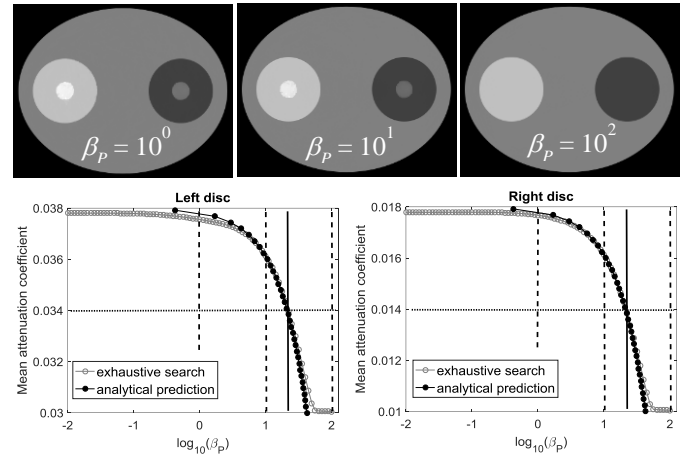


Fig. 7. Plot of disc average intensity versus prior image regularization strength for the ellipse phantom when the certainty-based shift-variant penalty is applied. Again, the predictor shows good agreement with the exhaustive search. The shift-variant penalty enforces a common transition region between the two locations (indicated with a vertical solid line centered at $\beta_p \approx 10^{1.5}$).

E.2. Factors affecting analytical prediction accuracy

There are a number of potential factors that will influence the accuracy of the analytical prediction. Three specific experiments using the ellipse phantom were undertaken to explore the following effects. First, the analytical prediction was derived under the assumption of $\beta_R = 0$ for mathematical convenience. In practice, one always applies the standard roughness penalty to help control the noise-resolution trade-off. Thus, we conduct a specific experiment using a range of β_R , comparing analytical prediction with exhaustive evaluation in the ellipse phantom. Second, we consider the case where noisy line integrals, $l - l_p$, are used to estimate $\mathbf{A}\Delta\mu$. To demonstrate the applicability of this approach and potential limitation due to noise in measurement data, analytical predictions using the line integral are compared with that using presumed change over a range of noise levels. Lastly, since contrast of the anatomical change plays a central role in image quality of the reconstructed PIRPLE image, we consider an additional experiment wherein the contrast of the stimulus is varied to illustrate low contrast predictive performance.

E.3. Quantitative evaluation of shift-variant penalty design

To quantify the performance of shift-variant penalty more completely, studies are conducted for a more dense sampling of points within the two phantoms for the case of a shift-variant penalty design. Specifically, the difference between brute force exhaustive search using reconstruction, the design given by (10) and the efficient design of (13) are compared for the specific choice of $\gamma = 1/2$.

E.4. Regularization design for varying acquisition protocols

Lastly, two additional studies are conducted using the anthropomorphic torso phantom in which data acquisition strategies are varied and regularization design is conducted. Specifically, since optimal regularization is known to be data-dependent, we consider a variable number of projection angles and a variable incident fluence (x-ray exposure) and whether the regularization design approach can be used to control for the variable acquisition strategy.

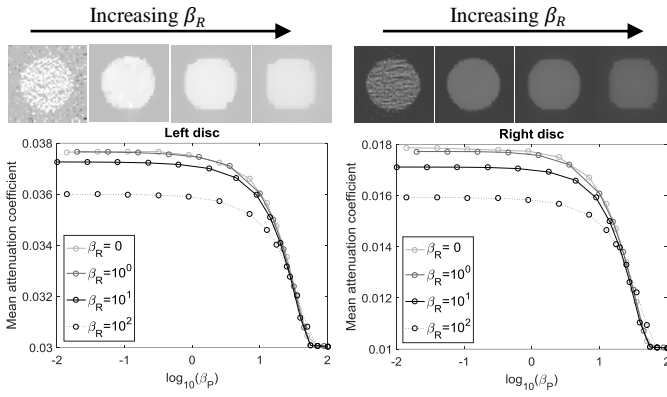


Fig. 8. The curves of exhaustive search using four different β_R values for the ellipse phantom when the shift-variant penalty are applied. When β_R is not too large, the curves nearly overlap with each other. When β_R is too large, the PIRPLE reconstructed images are over-smoothed, and there is additional bias resulting in underestimation of the contrast. Sample PIRPLE reconstructions of the two discs using $\beta_P = 10^0$ and four different β_R values ($0, 10^0, 10^1, 10^2$) are also shown.

III. EXPERIMENTAL RESULTS

A. Ellipse phantom investigations

A.1. Exhaustive search vs. analytical prediction

A summary of the simulation studies for PIRPLE reconstruction using standard shift-invariant regularization is shown in Fig. 6. Following the objective in (2) and using a fixed $\beta_R = 10^{2.5}$ and PIRPLE reconstructions with varying β_P (from 10^1 to 10^5 with a $10^{0.05}$ step size) were computed. Three sample reconstruction at $\beta_P = 10^3, 10^4$, and 10^5 are shown. Note when β_P is too large, the PIRPLE reconstruction replicates the prior image and true anatomical changes (i.e., two discs) do not appear. Also, there is an intermediate β_P where the right disc appears but the left disc disappears – suggesting significant location dependence of bias effects. In particular, we see that there is increased bias (a smaller γ fraction) in the dense portion of the phantom, where the stimuli has disappeared. This is consistent with the statistical model downweighting measurement data for that region due to increased density and increased noise, and more reliance on prior information in the objective function (e.g., increased similarity to the prior image) in this region.

The average intensity of voxels inside each disc for all PIRPLE reconstructions is plotted versus prior image regularization strength β_P for each disc below the sample reconstructions. A separate plot is shown for the left and right stimuli. Based on (10), we also varied γ from 0 to 0.99 with a 0.03 step size and computed the resulting β_P^* for each disc. The fractional reconstructed contrast of the change, γ , is related to the reconstructed attenuation (by adding the attenuation of the prior image) and curve of predicted stimuli attenuation versus prior image regularization strength was also plotted for each disc.

The analytical predictions match well with the exhaustive search results in the transition region, indicating a high accuracy of the proposed analytical prediction method. There was slightly increasing deviation between the brute force calculation and the predictions for values farther away from $\gamma = 1/2$, but this is expected because the analysis has focused on the transition region. We also observe a relative shift between the curves




Fig. 9. Analytical predictions of the bias-regularization curves using the line integral versus the presumed change approach. When the noise level is relatively low, the two predictions nearly overlap with each other. When noise levels become high, we see a systematic bias wherein the line integral approach curves are biased toward the left (lower β_P).

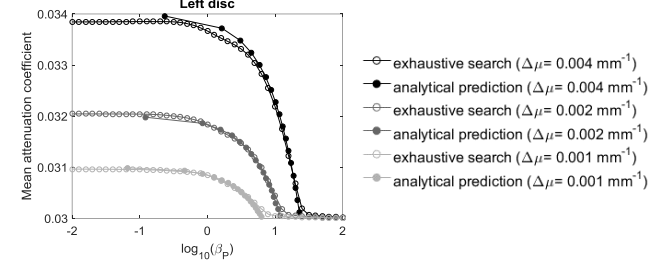


Fig. 10. Plot of left disc average intensity versus prior image regularization strength for the ellipse phantom with lower contrast disc stimuli of 0.004 mm^{-1} , 0.002 mm^{-1} , 0.001 mm^{-1} . Note that predictions remain accurate across stimulus contrast levels.

of the two stimuli. Specifically, the center of transition region for the left disc is at $\beta_P^* \approx 10^{3.4}$ but $\beta_P^* \approx 10^{4.3}$ for the right disc (identified with a solid vertical line).

The same set of experiments for the ellipse phantom using the shift-variant certainty penalty given by (14) was also conducted. Due to the certainty weighting, the range of penalty strengths was modified accordingly. Specifically, we employed a fixed $\beta_R = 10^0$ and varied β_P (from 10^{-2} to 10^2 with a $10^{0.05}$ step size) for PIRPLE reconstructions. Again, three sample reconstructions are shown. With the shift-variant penalty design, one can see that the two stimuli are reconstructed with similar fractional contrast showing less location dependence of bias. Reconstructed attenuation versus prior image regularization strength curves again show good agreement between prediction and exhaustive search. These curves also indicate the improved uniformity for the left and right discs, with curves exhibiting a transition region centered both at $\beta_P^* \approx 10^{1.3}$.

A.2. Factors affecting analytical prediction accuracy

To investigate the importance of β_R in the context of the derived relationship between β_P and fractional change reconstruction, we repeated the shift-variant experiment summarized in Fig. 7 with varying β_R . Specifically, the predicted analytic relationship is compared to the exhaustive reconstruction-based search for $\beta_R = 0, 10^0, 10^1$ and 10^2 . The results of this investigation are shown in Fig. 8. Note that the curves of the exhaustive search approach using $\beta_R \leq 10^0$ nearly overlap with each other. However, when β_R is larger, the curve for exhaustive search approach shifts downwards. Noting the associated reconstructions of the left and right stimuli, larger β_R , particularly for $\beta_R \geq 10^2$, the stimuli are noticeably over-smoothed and ex-

hibit decreased contrast due to blur with the surrounding material. Arguably, such a β_R is undesirable and should be avoided. Thus, the analytic framework of this paper still provides a prospective way to choose β_P since predictive curves will match for β_R choices that are not oversmoothing. Moreover, note that the location of the transition region in Fig. 8 does not change with β_R . This echoes previous work suggesting that the two regularization parameters may be optimized independently. And, since β_R controls a more familiar noise-resolution trade-off, such tuning may be more straightforward.

While this paper largely focuses on the presumed change framework in [24], we have identified the measured line integral approach where $\mathbf{A}\Delta\mu \approx l - l_P$ may be used in the design. To study efficacy of this approach at different noise levels, we varied incident fluence to 1×10^5 , 3×10^4 , and 1×10^4 photons/detector (for a fixed 90 projection angles), and compared the performance of these two analytical predictions (See Fig. 9). Note that the predictions overlap closely for higher fluence scenarios, with increasing bias at the lower fluence levels.

To investigate the effect of stimulus contrast on the accuracy of analytical prediction, we performed additional studies with decreased contrast for the disc stimuli. Specifically, we used a change contrast of 0.004 mm^{-1} , 0.002 mm^{-1} , and 0.001 mm^{-1} , and repeated the Section III.A.1 study using the shift-variant certainty penalty. The resulting bias-regularizations curves for the analytical prediction and exhaustive search for the left disc are shown in Fig. 10. The analytical predictions still match well with the exhaustive search results at these three contrast levels, suggesting that our predictive framework is applicable across contrast levels.

A.3. Quantitative evaluation of shift-variant penalty design

While the previous investigations of one or two stimuli are useful to illustrate the relationship between regularization and reconstruction bias, studies were limited to two location probes. We consider a more exhaustive spatial investigation for the shift-variant penalty design. Specifically, we investigated the location dependence of prior image regularization strength by adding a constant change (a disc with radius of 15 mm and attenuation of 0.008 mm^{-1}) to the prior image in Fig. 11(a) at many different locations. Fig. 11(b) shows an example of the constant change at one specific location in an FBP reconstruction of low-dose projection data. While the change is a constant, the absolute attenuation of the stimuli varies by location, depending on the attenuation of corresponding background. Fig. 11(c) illustrates the constant change at all 25 locations that were investigated. Note that in any given dataset there is only one stimulus at one specific location. The same low-dose simulations, projection geometry, etc. was used as in the previous ellipse examples.

For all of the locations, we chose a shift-variant penalty design goal to enforce $\gamma = 1/2$ for all points in the image volume. We performed that design using three approaches: 1) Gold standard, brute-force evaluation using exhaustive PIRPLE reconstruction with $\beta_R = 10^{2.5}$ and β_P varied from 10^1 to 10^5 with a fine $10^{0.01}$ step size. 2) Exhaustive predictive evaluation using evaluation of (10) for all 25 locations (j). And, 3) a computationally efficient shift-variant penalty specification based on




Fig. 11. Ellipse phantom: (a) prior image; (b) current low-dose image with one anatomical change (indicated by yellow circle) at a specific location; (c) one constant anatomical change at 25 different locations.




Fig. 12. Estimation of $\beta_{P,j}^*$ at $\gamma = 1/2$ by: (a) exhaustive search; (b) analytical prediction in Eq. (10); (c) analytical prediction in Eq. (13). The $\beta_{P,j}^*$ values at the 25 locations were obtained explicitly, and the maps were then generated by interpolation using radial basis functions.

(13) and the direct calculation of aggregate certainties. For the latter calculation, we used the center stimulus (indicated with green circle in Fig. 11(c)) to approximate the $\mathbf{A}^T \mathbf{A} \Delta\mu(j)$ term for all locations. To fill in penalty values for all points in the image volume, we used radial basis function [24] interpolation and extrapolation with the 25 computed samples.

Computed regularization maps (all $\beta_{P,j}^*$) for each strategy are shown in Fig. 12. The $\log_{10}(\beta_{P,j}^*)$ values are shown since the action of regularization is in the exponent for log-likelihood-based objective functions (like PIRPLE). The similarity between the three maps again demonstrates the high accuracy of the analytical prediction methods. Using the $\log_{10}(\beta_P^*)$ values from exhaustive search as ground truth, the root mean squared error (RMSE) for the analytical prediction in (10) and (13) for the 25 locations was computed as $\text{RMSE}=0.0165$ and $\text{RMSE}=0.1071$, respectively suggesting high accuracy in the design. While the increased computational efficiency of the (13) comes with a slight degradation in RMSE, we note that a $\log_{10}(\beta_{P,j}^*)$ shift of ~ 0.1 still localizes the transition region well. If one seeks to establish a specific level of bias associated with a change, this suggests the potential for a $\sim 10\%$ error (on average). If greater accuracy in fractional contrast reconstruction is required, the more computationally expensive approaches may be required. However, very high accuracy is usually not important for most applications. Previous work [24] has shown that there is a wide range of β_P that improves image quality throughout the image volume. The problematic aspect of regularization design is crossing into the transition region and beyond where features are grossly misrepresented or nonexistent. Thus, for many applications, it is sufficient to identify the location of the transition region and to keep β_P sufficiently low to avoid having important changes in the transition region. Localization of the exponent of β_P within ~ 0.1 easily establishes this goal since β_P is typically optimized on the scale of whole or half exponents.

The form of the shift-variant penalty design shown in Fig. 12 shows decreased regularization strength in the high density regions and increased strength in the low density region. This is

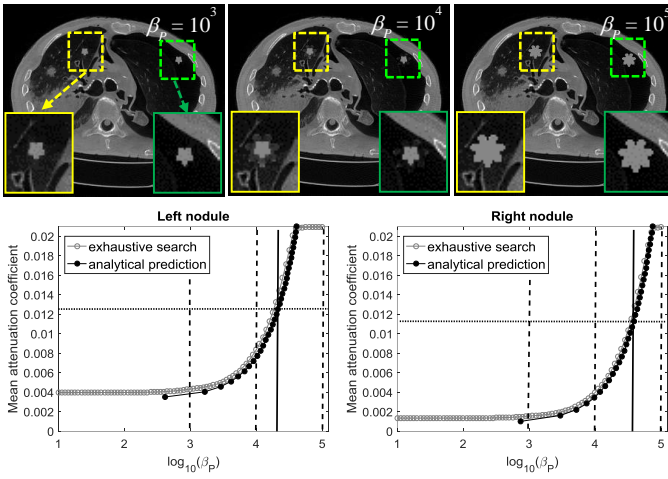


Fig. 13. Plot of nodule average intensity versus prior image regularization strength for the torso phantom for the shift-invariant penalty. The $\gamma = 1/2$ case is indicated with a dotted line.

consistent with previous observations with a uniform shift-invariant penalty where increased bias was observed in high-density regions and lower bias in low-density regions. The designed penalty (which seeks uniform bias in the change between prior and current anatomy) modifies the local regularization strength accordingly.

B. Torso phantom investigations

B.1. Exhaustive search vs. analytical prediction

While the ellipse phantom is illustrative of general dependencies on local density, it is not particularly realistic in terms of anatomy. Thus, we have conducted similar investigations in a torso phantom. We emulate a lung nodule surveillance scenario with irregular lesions and shrinkage over time (in contrast to the “simple” appearance of symmetric stimuli in the ellipse phantom). As in the previous investigation, both the shift-invariant and shift-variant regularization approaches were applied for the torso phantom. Since we emulated nodule shrinkage (i.e., negative change), a minus sign was added to the right-hand side of (10) and (13) for analytical prediction of β_p^* .

The shift-invariant results are summarized in Fig. 13. As before representative reconstructions are shown for three regularization values ($\beta_p = 10^3, 10^4$, and 10^5) and corresponding regularization bias curves formed via exhaustive reconstruction and analytical prediction are shown. Note that the change between the prior and current data is an irregular shape with a hollow center (no change at the center of the lesion). Average attenuation values in the reconstruction are computed over this irregular region. The intermediate reconstruction at $\beta_p = 10^4$ shows both the fractional representation of attenuation in the change region as well as location dependence with less bias in the right (green box) stimulus. Note that the regularization-bias are inverted as compared with Fig. 6 and Fig. 7, because of negative changes. Again, the curves accurately capture the bias in the reconstructed images as well as the location-dependence as indicated by the shift in transition region between the two lesion locations. Moreover, the analytical predictor and the exhaustive search remain in good agreement.

The shift-variant penalty design of (13) is applied to the torso phantom and summarized in Fig. 14. The results are consistent

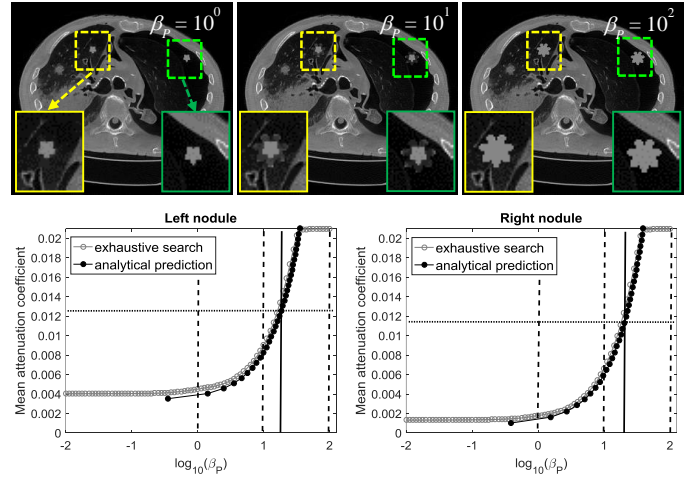


Fig. 14. Plot of nodule average intensity versus prior image regularization strength for the torso phantom for the shift-variant penalty. The $\gamma = 1/2$ case is indicated with a dotted line.

with previous observations and the shift-variant penalty design make the bias associated with the two lesion locations more uniform (see both the reconstructions and the shift in the transition region to a common $\beta_p^* \sim 10^{1.3}$). Thus, not only can object-dependence be observed in an anthropomorphic phantom, but the design strategy can be used to eliminate bias nonuniformity in anatomical change.

B.2. Quantitative evaluation of shift-variant penalty design

We also considered a more exhaustive investigation of location dependence of shift-variant prior image regularization strength for the torso phantom. Toward this end, we emulated a similar nodule shrinkage at 13 different locations within the lung parenchyma. As before, while a number of locations are investigated only a single lesion is emulated in a particular dataset. Fig. 15(a) and Fig. 15(b) illustrate an example of the nodule shrinkage at one specific location, showing the prior image and current low-dose FBP reconstruction, respectively. Fig. 15(c) illustrates all the 13 nodules simultaneously. The same low-dose protocol (i.e., 90 projection views and 10^4 incident photons per detector bin) was applied, and the regularization design goal targeted $\gamma = 1/2$. Three design strategies: 1)




Fig. 15. Torso phantom: (a) One example of a prior image; (b) One example of a current low-dose FBP image with nodule shrinkage; (c) The prior image with all the 13 nodules, indicating the probe nodule locations.




Fig. 16. Estimation of $\log_{10}(\beta_{p,j}^*)$ at $\gamma = 1/2$ for the torso phantom by: (a) exhaustive search; (b) analytical prediction in (10); (c) analytical prediction in (13). The $\log_{10}(\beta_{p,j}^*)$ values at the 13 locations were obtained explicitly, and the maps were then generated via radial basis function interpolation.

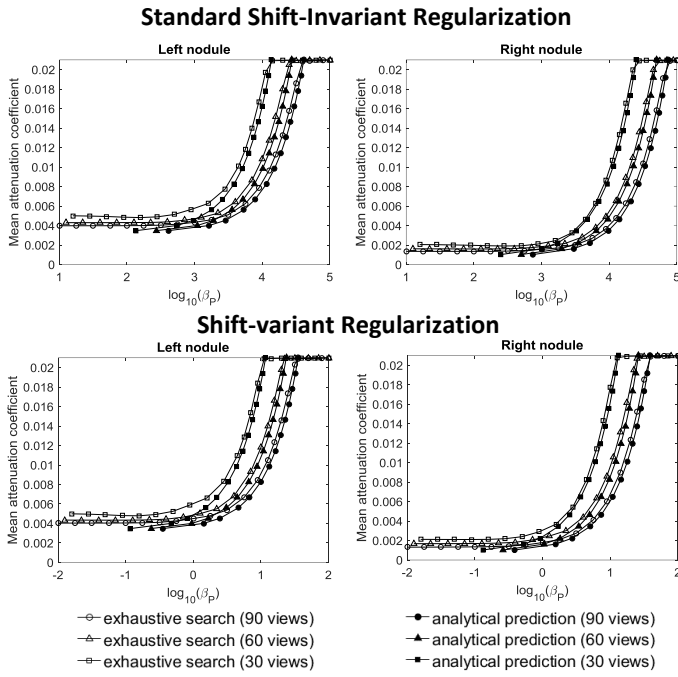


Fig. 17. Varied angular sampling investigation. Exhaustive search and analytical predictions for different numbers of projection angles for the torso phantom, using a standard penalty (top) and shift-variant penalty design (bottom).

exhaustive search via reconstruction; 2) analytical prediction using (10); and 3) certainty-based design were performed. For the latter design, we used the nodule change at *location 9* to approximate the $\mathbf{A}^T \mathbf{A} \Delta \mu(j)$ term for all locations, although the actual nodule change for different locations actually varies slightly since the lung background is not uniform.

Regularization maps containing $\log_{10}(\beta_{p,j}^*)$ values are shown in Fig. 16 using radial basis function interpolation and extrapolation for all points in the lung tissue. (Areas outside the lung are excluded since nodule shrinkage is not particularly well defined in these locations). Using the $\log_{10}(\beta_{p,j}^*)$ values from exhaustive search as the ground truth, the RMSE by the analytical predictions in (10) and (13) are RMSE=0.0473 and RMSE=0.0838, respectively. We note that the shift-variant $\log_{10}(\beta_{p,j}^*)$ maps retain high similarity between the exhaustive search and analytical prediction methods.

As previously observed, there is a general trend of lower regularization strength in high density regions (e.g., left portion of the image where lung edema is present). Though because the location-dependence is based on measurements through that location (e.g., the net attenuation/aggregate certainty for that location), there is also a general trend for decreased $\beta_{p,j}^*$ near higher density regions (e.g. mediastinum) and toward the center of the patient.

B.3. Regularization design for varying acquisition protocols

Previous experiments have investigated the bias dependencies on anatomy and location. In this section, we explore variations in the data acquisition in two experiments using the torso phantom in Fig. 5 with only two stimulus locations. The regularization-bias curves by exhaustive search and proposed analytical prediction were plotted using the same approach as described in Section III.B.1.

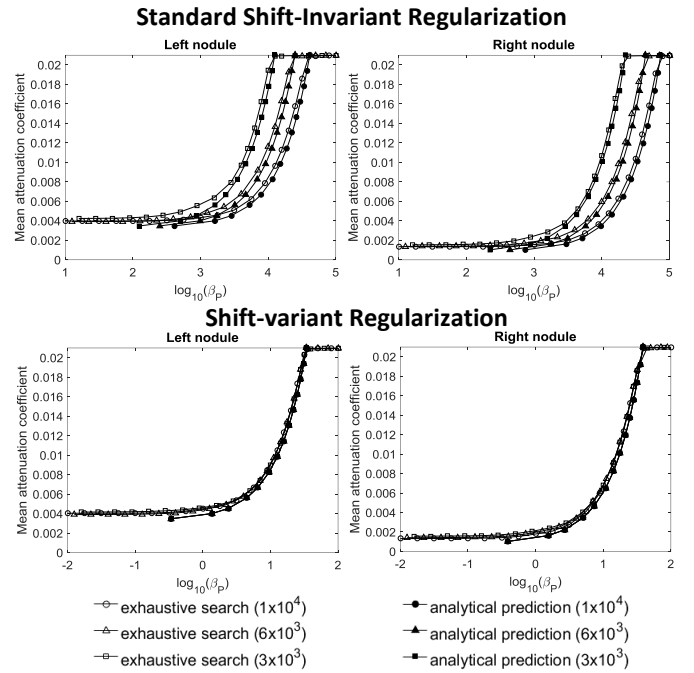


Fig. 18. Varied fluence investigation. Exhaustive search and analytical predictions for different fluence levels for the torso phantom, using standard penalty (top) and shift-variant penalty design (bottom).

The first experiment (Fig. 17) investigated variable angular sampling. We considered a total number of projection angles of 30, 60, and 90 (with fixed fluence at 10^4 photons/detector). (Smaller numbers of projection angles have been proposed as a dose reduction strategy that works well with ℓ_1 -norm-based regularization.) Both traditional shift-invariant β_P and shift-variant design are applied and compared to brute force evaluation via reconstruction. Again, we note good agreement and a reduction of location-dependence with the shift-variant design. We also find that one generally does need to adjust regularization strength for varying numbers of angles. Specifically, fewer angles requires lower β_P ; however, the location of the transition region remains predictable even for small numbers of angles.

A second experiment (Fig. 18) used varied fluence of 3×10^3 , 6×10^3 , and 1×10^4 photons/detector (for a fixed 90 projection angles). In the traditional regularization results with the original PIRPLE objective in (2), dependence on location and fluence levels are apparent, but there is good agreement between predictions and exhaustive search. For the shift-variant design achieved through implementation of the modified objective function of (14) with certainty-based weights, not only is location-dependence reduced by moving the regularization-bias curves to a common transition point, but the fluence-dependence is also eliminated – with all curves having the same transition point regardless of fluence.

IV. DISCUSSION AND CONCLUSION

We have proposed and evaluated an analytical approach that quantifies and predicts bias associated with prior image regularization strength in PIRPLE reconstruction. The bias is described as the fraction of the actual change between the prior image and the current anatomy that appears in the reconstruction. We have demonstrated the validity of the predictor with

comparisons to a brute force evaluation gold standard, and we have used the derived mathematical relationship between regularization and bias to drive reconstructions with uniform bias properties for reconstructed anatomical changes. Computationally efficient versions of the regularization design were also developed and demonstrated to be accurate. There are several merits with the presented analytical prediction approach. It does not require explicit iterative reconstructions. The expression takes various dependencies into account including the system geometry, sampling, fluence levels, measurement statistics, and the location and contrast of the anatomical changes – permitting reliable regularization control across those dependencies. This work provides a general framework to predict and control PIRPLE performance, which fills the current gap in understanding PIBR methods. We believe that the introduction of quantitative metrics of PIBR bias and predictive methods for bias estimation is critical for producing reliable PIBR images, which, in turn, is also critical for the eventual clinical translation of these powerful approaches.

While this work is an important stepping stone in PIBR analysis and control, there are many aspects that require additional investigation in future studies. In particular, the analytic framework developed here identifies the particular bias associated with PIBR (partial reconstruction of contrast changes) and establishes a relation between regularization strength and that bias. We note that the framework suggests that any amount of regularization results in some reconstruction bias. The specific level of bias resulting from regularization design is left to the user. The degree of acceptable bias is likely highly dependent on the clinical task. For example, detection tasks may be more tolerant of bias than quantification tasks. Similarly, this work concentrates only on bias and not the noise aspects of PIBR. Both bias and noise are likely to play an important role in the particular trade-offs of PIBR and require additional investigations.

The analytic framework makes the simplification that the standard roughness penalty plays a small role in the reconstruction. Despite illustrating that there is a range of β_R over which this assumption is a good one, that premise breaks down for larger β_R . In future studies we seek a unified regularization analysis (e.g. factoring the two regularization strategies) that will permit image quality predictions including both regularization terms.

Current work also ignores the effect of registration between the prior image and current anatomy. Registration is critical for many clinical applications including lung nodule surveillance due to different patient positioning, inhalation states, patient weight gain/loss, etc. While the studies presented in this work apply for well-registered prior images, future studies will need to assess the impact of misregistration. The interaction between registration and reconstruction is potentially complex since some changes (e.g. nodule growth) could potentially be modeled either as an anatomical change (in attenuation) or as a deformation of the prior image anatomy, in addition to “simple” misregistration errors. This raises a number of additional questions about the correct way to balance registration and recon-

struction in the context of PIBR. However, we note that reconstruction and registration can often be decoupled (as with PIRPLE), permitting different registration algorithms to be tested. This permits investigation of methods that explicitly forbid transformations that could be interpreted as nodule growth as well as studies where such constraints are not applied. We plan on addressing these registration issues in ongoing and future studies including clinical datasets with realistic patient motion and anatomical change.

Finally, our analysis in this study was focused on the PIRPLE method and lung nodule surveillance application. Nevertheless, the proposed methodology is general and uses similar mathematical techniques as were used in [27] which permitted analysis of other PIBR methods like PICCS. Naturally, this work has potential application beyond lung nodule surveillance including cardiac sequences, [12] perfusion, [17] angiography, [32] and cross-modality prior images. [33] For example, planning CT images may be used to improve onboard CBCT imaging in radiation therapy [33]. Prospective regularization design is significantly more complicated in this scenario, due to increased scatter fractions in CBCT data, attenuation differences between the planning CT and CBCT datasets, as well as positioning difference and organ deformations between scans [33]. Investigation of prospective regularization design for cross-platform applications is an important topic for future consideration.

REFERENCES

- [1] S. J. Swensen, J. R. Jett, T. E. Hartman, D. E. Midthun, J. A. Sloan, A.-M. Sykes, G. L. Aughenbaugh, and M. A. Clemens, “Lung Cancer Screening with CT: Mayo Clinic Experience,” *Radiology*, vol. 226, no. 3, pp. 756–761, Mar. 2003.
- [2] H. MacMahon, J. H. M. Austin, G. Gamsu, C. J. Herold, J. R. Jett, D. P. Naidich, E. F. Patz, S. J. Swensen, and Fleischner Society, “Guidelines for management of small pulmonary nodules detected on CT scans: a statement from the Fleischner Society,” *Radiology*, vol. 237, no. 2, pp. 395–400, Nov. 2005.
- [3] A. Neroladaki, D. Botsikas, S. Boudabbous, C. D. Becker, and X. Montet, “Computed tomography of the chest with model-based iterative reconstruction using a radiation exposure similar to chest X-ray examination: preliminary observations,” *Eur. Radiol.*, vol. 23, no. 2, pp. 360–366, Feb. 2013.
- [4] V. Vardhanabhuti, R. J. Loader, G. R. Mitchell, R. D. Riordan, and C. A. Roobottom, “Image quality assessment of standard- and low-dose chest CT using filtered back projection, adaptive statistical iterative reconstruction, and novel model-based iterative reconstruction algorithms,” *AJR. Am. J. Roentgenol.*, vol. 200, no. 3, pp. 545–52, Mar. 2013.
- [5] K. Sauer and C. Bouman, “A local update strategy for iterative reconstruction from projections,” *IEEE Trans. Signal Process.*, vol. 41, no. 2, pp. 534–548, 1993.
- [6] J.-B. Thibault, K. D. Sauer, C. A. Bouman, and J. Hsieh, “A three-dimensional statistical approach to improved image quality for multislice helical CT,” *Med. Phys.*, vol. 34, no. 11, p. 4526, 2007.
- [7] Q. Xu, H. Yu, X. Mou, L. Zhang, J. Hsieh, and G. Wang, “Low-dose X-ray CT reconstruction via dictionary learning,” *IEEE Trans. Med. Imaging*, vol. 31, no. 9, pp. 1682–97, Sep. 2012.
- [8] S. Cho, E. Pearson, E. Y. Sidky, J. Bian, C. A. Pelizzari, and X. Pan, “Prior-image-based few-view cone beam CT for applications to daily scan in image-guided radiation therapy: preliminary study,” *Proc. SPIE*, vol. 7258, p. 72581U, Feb. 2009.
- [9] 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.
- [10] G.-H. Chen, J. Tang, and S. Leng, “Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets,” *Med. Phys.*, vol. 35, no. 2, pp. 660–3, Feb. 2008.
- [11] 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 acquisitions,” *SPIE Med. Imaging*, vol. 7258, no. December, pp. 725803-1–12, 2009.

- [12] G.-H. Chen, P. Theriault-Lauzier, J. Tang, B. Nett, S. Leng, J. Zambelli, Z. Qi, N. Bevens, A. Raval, S. Reeder, and H. Rowley, “Time-Resolved Interventional Cardiac C-arm Cone-Beam CT: An Application of the PICCS Algorithm,” *IEEE Trans. Med. Imaging*, vol. 31, no. 4, pp. 907–923, Apr. 2012.
- [13] 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, p. 7563, 2013.
- [14] 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.
- [15] T. Heußner, M. Brehm, L. Ritschl, S. Sawall, and M. Kachelrieß, “Prior-based artifact correction (PBAC) in computed tomography,” *Med. Phys.*, vol. 41, no. 2, p. 21906, Jan. 2014.
- [16] 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, 2016.
- [17] S. Seyyedi, E. Liapi, T. Lasser, R. Ivkov, R. Hatwar, and J. Stayman, “Evaluation of Low-Dose CT Perfusion for the Liver using Reconstruction of Difference,” in *International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, 2017, pp. 343–347.
- [18] 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.
- [19] H. Zhang, J. Ma, J. Wang, W. Moore, and Z. Liang, “Assessment of prior image induced nonlocal means regularization for low-dose CT reconstruction : change in anatomy,” *Med. Phys.*, vol. 44, no. 9, pp. e264–e278, 2017.
- [20] 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, 2014.
- [21] 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.
- [22] S. Itoh, M. Ikeda, S. Arahata, T. Kodaira, T. Isomura, T. Kato, K. Yamakawa, K. Maruyama, and T. Ishigaki, “Lung Cancer Screening: Minimum Tube Current Required for Helical CT,” *Radiology*, vol. 215, no. 1, pp. 175–183, Apr. 2000.
- [23] 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.
- [24] 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.
- [25] M. Wu, Q. Yang, A. Maier, and R. Fahrig, “Approximate path seeking for statistical iterative reconstruction,” *Proc. SPIE*, vol. 9412, p. 94121D, Mar. 2015.
- [26] 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.
- [27] J. W. Stayman, J. L. Prince, and J. H. Siewerdsen, “Information Propagation in Prior-Image-Based Reconstruction,” in *International Conference on Image Formation in X-Ray Computed Tomography*, 2012, vol. 2012, pp. 334–338.
- [28] H. Zhang, H. Dang, G. J. Gang, and J. W. Stayman, “Prospective Regularization Analysis and Design for Prior-Image-Based Reconstruction of X-ray CT,” in *International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, 2017, vol. 14, pp. 417–423.
- [29] H. Erdogan and J. A. Fessler, “Ordered subsets algorithms for transmission tomography,” *Phys. Med. Biol.*, vol. 44, no. 11, pp. 2835–51, Nov. 1999.
- [30] J. H. Cho and J. A. Fessler, “Regularization Designs for Uniform Spatial Resolution and Noise Properties in Statistical Image Reconstruction for 3-D X-ray CT,” *IEEE Trans. Med. Imaging*, vol. 34, no. 2, pp. 678–689, Feb. 2015.
- [31] 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.
- [32] P. Wu, J. W. Stayman, M. Mow, W. Zbijewski, A. Sisniega, N. Aygun,

- R. Stevens, V. Koliatsos, D. Foos, X. Wang, and J. H. Siewerdsen, “Evaluation of the Reconstruction-of-Difference (RoD) Algorithm for Cone-Beam CT Neuro-Angiography,” *Proc. SPIE*, vol. 10573, p. 105731R, 2018.
- [33] H. Zhang, G. J. Gang, J. Lee, J. Wong, and J. W. Stayman, “Integration of prior CT into CBCT reconstruction for improved image quality via reconstruction of difference: first patient studies,” *Proc. SPIE*, vol. 10132, p. 1013211, 2017.