Published in final edited form as:

Phys Med Biol. 2015 December 21; 60(24): 9515–9536. doi:10.1088/0031-9155/60/24/9515.

Prospective regularization design in prior-image-based reconstruction

Hao Dang, Jeffrey H. Siewerdsen, and J. Webster Stayman

Department of Biomedical Engineering, Johns Hopkins University, Baltimore MD 21205

Abstract

Prior-image-based reconstruction (PIBR) methods leveraging patient-specific anatomical information from previous imaging studies and/or sequences have demonstrated dramatic improvements in dose utilization and image quality for low-fidelity data. However, a proper balance of information from the prior images and information from the measurements is required (e.g., through careful tuning of regularization parameters). Inappropriate selection of reconstruction parameters can lead to detrimental effects including false structures and failure to improve image quality. Traditional methods based on heuristics are subject to error and sub-optimal solutions, while exhaustive searches require a large number of computationally intensive image reconstructions. In this work, we propose a novel method that prospectively estimates the optimal amount of prior image information for accurate admission of specific anatomical changes in PIBR without performing full image reconstructions. This method leverages an analytical approximation to the implicitly defined PIBR estimator, and introduces a predictive performance metric leveraging this analytical form and knowledge of a particular presumed anatomical change whose accurate reconstruction is sought. Additionally, since model-based PIBR approaches tend to be space-variant, a spatially varying prior image strength map is proposed to optimally admit changes everywhere in the image (eliminating the need to know change locations *a priori*). Studies were conducted in both an ellipse phantom and a realistic thorax phantom emulating a lung nodule surveillance scenario. The proposed method demonstrated accurate estimation of the optimal prior image strength while achieving a substantial computational speedup (about a factor of 20) compared to traditional exhaustive search. Moreover, the use of the proposed prior strength map in PIBR demonstrated accurate reconstruction of anatomical changes without foreknowledge of change locations in phantoms where the optimal parameters vary spatially by an order of magnitude or more. In a series of studies designed to explore potential unknowns associated with accurate PIBR, optimal prior image strength was found to vary with attenuation differences associated with anatomical change but exhibited only small variations as a function of the shape and size of the change. The results suggest that, given a target change attenuation, prospective patient-, change-, and data-specific customization of the prior image strength can be performed to ensure reliable reconstruction of specific anatomical changes.

1. Introduction

Sequential imaging studies are conducted in many clinical scenarios including disease monitoring (Hasegawa *et al* 2000, Caillot *et al* 2001), image-guided interventional surgeries (Siewerdsen *et al* 2009, Cleary and Peters 2010, Navab *et al* 2010, Dang *et al* 2012), and image-guided radiotherapy (Xing *et al* 2006, Jaffray *et al* 2002, Meeks *et al* 2005, Pouliot *et*

et al 2005). Although previously acquired images in a sequential study share a great amount of patient-specific anatomical information, traditional reconstruction methods tend to treat each scan in isolation and neglect the valuable information in prior images from previous scans. As a result, sequential studies are traditionally conducted through a series of *complete* acquisitions and the cumulative radiation dose can raise concern for both the patient and the radiation/surgical staff.

The ability to integrate prior images in image reconstruction has been recognized as a valuable tool in recent years. A number of prior-image-based reconstruction (PIBR) methods have been developed that enable substantial dose reduction (e.g., by sparse sampling and/or reduced x-ray fluence) while maintaining good image quality. One popular PIBR method is Prior Image Constrained Compressed Sensing (PICCS) (Chen *et al* 2008). PICCS encourages sparse differences between the reconstruction of the current anatomy and a prior image using an L^1 minimization subject to a data consistency constraint. Subsequent developments have integrated data fit terms that integrate statistical noise models within the PICCS framework (Lauzier and Chen 2013) and Prior Image Registration, Penalized-Likelihood Estimation (PIRPLE) (Stayman *et al* 2011, 2013). The latter approach employs a rigid registration of the prior image to accommodate changes in positioning between the prior image and current anatomy. The dPIRPLE variant (Dang *et al* 2013, 2014b) leverages deformable registration to accommodate nonrigid motion. Additionally, Pourmorteza *et al.* developed a penalized-likelihood method that reconstructs the difference image (with respect to a prior image) directly from the measurements (Pourmorteza *et al* 2015) while Abbas *et al.* leveraged a total-variation minimization algorithm to reconstruct the difference image from projection data difference between two successive scans (Abbas *et al* 2013). Xu *et al.* developed a penalized weighted least-squares method that penalizes differences between image patches in the reconstructed image based on prior image information (Xu and Tsui 2013). All of these approaches have demonstrated that inclusion of prior image information has great potential to dramatically reduce data fidelity requirements while maintaining or improving image quality.

Despite the benefits of PIBR, all of the methods need to answer a key question: To what extent should prior image information be used to achieve accurate image reconstruction? That is, how much information should come from the prior image and how much should come from the measurements. Such balance is typically controlled via regularization parameters, such as α/λ in PICCS with statistical weights (Lauzier and Chen 2013), β_P in PIRPLE and dPIRPLE (Stayman *et al* 2013, Dang *et al* 2014b), β_M in the Reconstruction of Difference method (Pourmorteza *et al* 2015), and β in image reconstruction that penalizes differences between image patches (Xu and Tsui 2013). Inappropriate selection of the prior image strength can lead to poor reconstruction. Specifically, the use of too little prior image information fails to produce a significant imaging benefit with low fidelity data, while the use of too much prior image information can force the reconstructed image to simply replicate the prior image, potentially obscuring anatomical changes and producing false structures. Further complicating the balance is that the optimal prior image strength can vary across different patients, anatomical changes (e.g., attenuation, shape, and size), acquisition geometry, x-ray techniques, and image reconstruction parameters (e.g., voxel size, volume size).

Traditional methods for choosing the proper prior image strength include exhaustive searches and heuristics/look-up tables. Exhaustive search involves performing a large number of image reconstructions with different regularization parameter values and choosing a value corresponding to an image that optimizes a certain image quality metric; however, this method can be extremely time-consuming, since each reconstruction requires iterative solution. Heuristics and look-up tables do not require image reconstruction and therefore require much less time, but they can be subject to error and suboptimal solutions due to the aforementioned variations of the optimal strength across imaging studies.

In this work, we propose a novel method that prospectively estimates the optimal prior image strength for PIBR without heuristics or exhaustive search. This method leverages an analytical approximation to PIBR objective functions containing non-quadratic penalties (Stayman *et al* 2012). A predictive performance metric is introduced that utilizes the approximate analytical solution and a specification of an anticipated change (i.e., an anatomical change for which accurate reconstruction is to be ensured). This performance metric is, in turn, used to estimate the optimal prior image strength. Additionally, because optimal prior image strength is dependent on the location of anatomical change (Dang *et al* 2014a), a spatially varying prior strength map is proposed to optimally admit changes everywhere in the image. Thus, proposed space-variant design can ensure accurate reconstructions without *a priori* knowledge of the change location.

The proposed methodology is investigated both in an ellipse phantom and in a realistic thorax phantom emulating a lung nodule surveillance scenario. Performance is compared with traditional exhaustive searches and the optimality of the space-variant design is explored. Additionally, the dependence of optimal prior image strength on different properties of the change (i.e., attenuation, shape, size) is also explored.

2. Methods

2.1 Regularization design in prior-image-based reconstruction (PIBR)

While a number of different PIBR approaches exist, the investigations below focused specifically on the PIRPLE methodology (Stayman *et al* 2013). We expect that the basic framework can be extended to other approaches including PICCS with statistical weighting (Lauzier and Chen 2013), image reconstruction using non-local prior functions (Xu and Tsui 2013), and reconstruction of difference using prior images (Pourmorteza *et al* 2015). The PIRPLE method employs a penalized-likelihood estimation (PLE) framework which incorporates patient-specific prior image information through a regularization term. The objective function of PIRPLE contains three terms: 1) a data fidelity term that uses the measured data based on measurement statistics; 2) an image roughness regularization term that enforces local smoothness and/or edge-preservation in the image estimate; and 3) a prior image regularization term that enforces similarity of the image estimate to a prior image while allowing sparse differences (i.e., anatomical changes) between the two. The prior image is simultaneously registered to the current patient anatomy in either a rigid fashion (Stayman *et al* 2013) or a nonrigid fashion (Dang *et al* 2013, 2014b) to account for patient motion between the previous scan and the current scan. For this work we consider the PIRPLE objective function without registration:

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

which includes the (Poisson) log-likelihood $L(y; \mu)$ parameterized by the image estimate μ (a $N_\mu \times 1$ vector of linear attenuation coefficients) and the measurements y (a $N_y \times 1$ vector).

The image roughness regularization term $\beta_R \|\Psi_R \mu\|_{p_R}^{p_R}$ includes a sparsifying operator Ψ_R , a modified p -norm operator $\|\cdot\|_{p_R}^{p_R}$, and an image roughness strength parameter β_R . In this work, Ψ_R computes first-order neighborhood pairwise voxel differences. The prior image regularization term $\beta_P \|\Psi_P (\mu - \mu_P)\|_{p_P}^{p_P}$ encourages similarity between the reconstruction and prior with the option of an additional sparsifying transform Ψ_P . In this work, Ψ_P is chosen to be the identity matrix since the anatomical changes between the prior image μ_P and the current anatomy are already sparse. In situations where the anatomical changes are less sparse (e.g., in the case of a much larger anatomical change), a sparsifying transform other than the identity matrix would be encouraged. Additionally, to avoid nondifferentiable penalties, both p -norms (i.e., $\|\cdot\|_p^p$) are modified to be quadratic within the region $[-\delta, \delta]$. In this work, we focus on $p_R = 1$ because L^1 norm penalty function or its variants (e.g., the Huber loss function (Huber 1981)) has been shown to encourage edge preservation and achieve improved noise-resolution tradeoff in the reconstructed image compared to quadratic penalty function (Thibault *et al* 2007, Wang *et al* 2014). We use $p_P = 1$ in this work because L^1 norm penalty function or its variants has been shown to encourage similarity of the reconstructed image to the prior image but also allow for sparse differences (Chen *et al* 2008, Stayman *et al* 2013). Similar ideas of using an L^1 norm constraint for sparse solution recovery can be found in compressed sensing theory (Candès *et al* 2006). The prior image regularization term is weighted by a prior image strength parameter β_P .

Accurate reconstruction of anatomical changes in PIBR requires a proper balance between the information from the prior image and the measurements. In PIRPLE, this balance is governed primarily by the prior image strength parameter β_P . A larger β_P leads to the use of more information from the prior image in the current image reconstruction, while smaller β_P restricts the amount of information from the prior image. Accurate reconstruction of anatomical changes in PIRPLE is also affected by the image roughness strength β_R through the control of image smoothness. Previous work (Stayman *et al* 2013) suggests that the optimization of these two parameters is often separable, suggesting that the two parameters can be selected independently. Specifically, previous work (Stayman *et al* 2013) found that the optimal value of β_P is nearly independent of β_R when β_R is low. Therefore, a 1D optimization over β_P can be first performed with a low β_R to estimate the optimal β_P , which is followed by another 1D optimization over β_R with the optimal β_P . The optimal β_P and β_R estimated from this pair of separate 1D optimization using root mean square error (RMSE) from a truth image as the metric has been found to closely match the optimal β_P and β_R estimated from an exhaustive 2D optimization (Stayman *et al* 2013). In this work, we focus on investigating the selection of the optimal prior image strength β_P in the context of a fixed image roughness strength β_R , which is the first step of these two 1D optimizations.

2.2 Proposed method for regularization design in PIBR

In this section, an approximate analytical solution is derived for the PIRPLE method. The selection of a key parameter in approximating the analytical solution – the operating point – is discussed, along with the introduction of the concept of prospective regularization design in PIBR. Finally, a predictive performance metric is proposed, which uses the approximate analytical solution to estimate the optimal prior image strength for a given anatomical change.

2.2.1 Approximate analytical solution of PIBR—The objective function of PIRPLE in Eq. (1) does not generally admit a closed-form solution because of the nonlinearities of both the log-likelihood function and the non-quadratic regularization. In (Stayman *et al* 2013, Dang *et al* 2013, 2014b), this objective function is solved iteratively using optimization approaches. However, it is possible to derive an *approximate* closed-form solution of the objective. Such approximation does not necessarily avoid a full iterative solution for the desired reconstruction of subsequent imaging studies, but may suffice for prospective regularization design. In this work, we derive an approximate closed-form solution by substituting each non-quadratic term in Eq. (1) with a quadratic term.

First, the data fidelity term can be approximated by a weighted least-squares term using a second-order Taylor approximation of the log-likelihood function (Sauer and Bouman 1993). The simplified objective function can be written as:

$$\hat{\mu} \approx \arg \min_{\mu} \|\mathbf{A}\mu - \mathbf{l}\|_{\mathbf{W}}^2 + \beta_R \|\Psi_R \mu\|_{p_R}^{p_R} + \beta_P \|\Psi_P (\mu - \mu_P)\|_{p_P}^{p_P} \quad (2a)$$

$$l_i = -\log(y_i/g_i), \quad \mathbf{W} = \mathbf{D}\{y\} \quad (2b)$$

where $\mathbf{A}$ denotes a $N_y \times N_\mu$ system matrix representing the linear projection operator (and $\mathbf{A}^T$ denotes the matched backprojection operator), $\mathbf{l}$ is a $N_y \times 1$ vector of line integrals computed from the measurements using a log transformation, $\mathbf{g}$ is a $N_y \times 1$ vector of measurement-dependent gains (e.g., number of photons when no object is present), the subscript i denotes the i th element of the subscripted vector, $\mathbf{W}$ is a diagonal weighting matrix, and $\mathbf{D}\{\cdot\}$ is an operator that places a vector on the main diagonal of a matrix. We note that in some cases the Gaussian assumption of the penalized weighted least-squares data fidelity term is preferred over the nonlinear Poisson likelihood. For example, the diagonal weighting $\mathbf{W}$ can be modified to accommodate various data corrections that modify the noise structure (Dang *et al* 2015b, 2015a).

Second, non-quadratic regularization (image roughness regularization and prior image regularization) terms can also be approximated by quadratic terms. Using the modified L^1 norm discussed above and choosing a proper operating point about which we wish to form approximate penalty estimates, the modified norm can be approximated by a quadratic function. We may approximate the scalar function applied to each element of the vector argument of the modified norm as follows:

$$f(x_i) \approx g(x_i) = \kappa_i(\tau_i) x_i^2 \quad (3a)$$

$$f(x_i) = \begin{cases} \frac{1}{2\delta} x_i^2 & |x_i| < \delta \\ |x_i| - \frac{\delta}{2} & |x_i| \geq \delta \end{cases}, \quad g(x_i) = \kappa_i(\tau_i) x_i^2 = \begin{cases} \frac{1}{2\delta} x_i^2 & |\tau_i| \geq \delta \\ \frac{|\tau_i| - \delta/2}{\tau_i^2} x_i^2 & |\tau_i| < \delta \end{cases} \quad (3b)$$

where f denotes the modified penalty for the $p = 1$ scenario which is equivalent to the Huber loss function (Huber 1981). The function f has a scalar input x_i and may be approximated with the quadratic function g . The function g includes a coefficient κ_i that is computed as a function of the operating point τ_i . As shown in Eq. (3b), when τ_i is chosen to be within the quadratic neighborhood $([-\delta, \delta])$ of the function f , the values of the function f and g are exactly matched for any x_i within the quadratic neighborhood. This scenario is illustrated in Fig. 1(a). When τ_i is chosen to be outside the quadratic neighborhood, the values of the function f and g are exactly matched at $x_i = \tau_i$ and remain close to each other around $x_i = \tau_i$ as shown in Fig. 1(b). Such quadratic approximation of the non-quadratic regularization yields two important observations. First, for each input element of f , a separate operating point will be chosen, and a separate parabola will be constructed indicating a location-dependent approximation. Second, since the approximation is most accurate around the operating point, it is desirable to select an operating point that is equal to or close to the value at which the penalty is evaluated.

Using the quadratic approximation in Eq. (3), the entire modified L^1 norm can be approximated and represented in the following matrix form:

$$\|x\|_1^1 = \sum_i f(x_i) \approx \sum_i \kappa_i(\tau_i) x_i^2 = (x)^T \mathbf{D} \{\kappa(\tau)\} (x) \quad (4)$$

Where $\|x\|_1^1$ denotes the modified L^1 norm of a vector x , and κ and τ denote a vector of coefficients and operating points, respectively. One can then apply Eq. (4) to the two non-quadratic regularization terms in Eq. (2a) to yield two quadratic terms:

$$\hat{\mu} \approx \arg \min_{\mu} \|\mathbf{A}\mu - l\|_{\mathbf{W}}^2 + \beta_R (\Psi_R \mu)^T \mathbf{D}_R \Psi_R \mu + \beta_P [\Psi_P (\mu - \mu_P)]^T \mathbf{D}_P \Psi_P (\mu - \mu_P) \quad (5a)$$

$$\mathbf{D}_R = \mathbf{D} \{\kappa(\Psi_R \tilde{\mu})\}, \quad \mathbf{D}_P = \mathbf{D} \{\kappa(\Psi_P (\tilde{\mu} - \mu_P))\} \quad (5b)$$

For each regularization term, an operating point must be defined in the quadratic approximations. Ideally, this operating point should be an image volume close to the solution of the original objective function. Presuming we have an image estimate $\tilde{\mu}$, this leads to an operating point $\Psi_R \tilde{\mu}$ for the image roughness regularization and $\Psi_P (\tilde{\mu} - \mu_P)$ for the prior image regularization, as shown in Eq. (5b). When the image estimate $\tilde{\mu}$ is chosen to be close to the PIRPLE solution $\hat{\mu}_{\text{PIRPLE}}$, such selection of the operation point is expected to yield accurate approximation of the regularization terms evaluated at $\hat{\mu}_{\text{PIRPLE}}$. Details on choosing a specific image estimate $\tilde{\mu}$ will be discussed in the next section.

The objective function in Eq. (5a) now contains three quadratic terms thereby admitting a closed-form solution, which can be written as:

$$\hat{\mu} \approx \left(\mathbf{A}^T \mathbf{W} \mathbf{A} + \beta_R \Psi_R^T \mathbf{D}_R \Psi_R + \beta_P \Psi_P^T \mathbf{D}_P \Psi_P \right)^{-1} \left(\mathbf{A}^T \mathbf{W} \mathbf{l} + \beta_P \Psi_P^T \mathbf{D}_P \Psi_P \mu_P \right) \quad (6)$$

This closed-form solution will be used later in Sec. 2.2.3 to estimate the optimal prior image strength.

2.2.2 Selection of an operating point—Proper selection of the operating points $\Psi_R \tilde{\mu}$ and $\Psi_P (\tilde{\mu} - \mu_P)$ in Eq. (5b) is important for accurate approximation of the actual solution of PIRPLE. Ideally, one should use the PIRPLE solution as $\tilde{\mu}$ so that the value of the approximate quadratic function exactly matches that of the Huber loss function. We refer to this as the Ideal estimate, whose approximate solution of PIRPLE is denoted $\hat{\mu}_{\text{Ideal}}$. While this estimate requires a full PIRPLE reconstruction, which is computationally expensive and supersedes the need for an approximate solution, the Ideal estimate is useful in investigating the accuracy of the underlying quadratic approximation of the non-quadratic regularizations.

A practical operating point is one that can be used for prospective regularization design without having performed the reconstruction. In sequential imaging studies, one often has general knowledge of the anticipated change (or perhaps the kind of changes one might wish to see) in the subsequent scan. For example, in pulmonary nodule surveillance, clinicians may have some knowledge of the nodule's attenuation, size, and shape in a follow-up scan based on the progress of the disease and based on the appearance of the nodule in a previous scan. Similar situations can be found in image-guided radiation therapy, where clinicians may anticipate the location of a tumor or tissue at risk in the current scan based on its location in previous scans. Likewise in image-guided procedures where specific tissue volumes are targeted for resection or specific implants are introduced into the patient, and preoperative scans provide the basis for prior image information.

While knowledge of possible change in the image volume generally includes varying levels of certainty about the specific attenuation, location, shape, and size, it is convenient to start with the assumption that the change is known. Similarly, for prospective regularization design, one might presume that a particular change is present in order to ensure that the regularization design is appropriate should the actual change be present. In this case, one could form an image estimate $\tilde{\mu}$ as the sum of a prior image μ_P and some presumed change μ_C :

$$\tilde{\mu} = \mu_P + \mu_C \quad (7)$$

We refer to this method as the P+C estimate, with the approximate PIRPLE solution based on (6) denoted as $\hat{\mu}_{P+C}$. Since the image estimate $\tilde{\mu}$ in the P+C estimate does not use the PIRPLE solution, a full image reconstruction is not required in this method. Moreover, one could use the approximate solution from the first P+C estimate (denoted as $\hat{\mu}_{P+C}^1$) as the image estimate $\tilde{\mu}$ to perform another P+C estimate, and repeat n times to get the n^{th} iteration P+C estimate (denoted as $\hat{\mu}_{P+C}^n$). Iterating on the P+C estimate is expected to yield better

selection of the operating point, thereby resulting in more accurate approximation of the PIRPLE solution, though such a procedure will increase the computation time associated with the estimation. All the estimation methods mentioned in this section along with the full PIRPLE reconstruction method are summarized in Table 1.

2.2.3 Predictive performance metric—Previous work (Stayman *et al* 2013, Dang *et al* 2014b) considered optimal prior image strength β_P based on minimizing the RMSE of a PIRPLE-reconstructed image with respect to a truth image in a region of interest (ROI) containing the change. In this work, we propose a predictive performance metric which uses the approximate analytical solution introduced in the previous section to estimate a value of β_P that optimally admits a given anatomical change in the image reconstruction. The metric can be written as:

$$\hat{\beta}_P = \arg \min_{\beta_P} \|\mu_P + \mu_C - \hat{\mu}(\beta_P)\|_S \quad (8)$$

which computes the RMSE between the approximate solution and the sum of a prior image and some presumed change in a ROI that contains the change. This ROI is referred to as the Change ROI and is represented using a binary mask $\mathbf{S}$ in Eq. (8).

We choose a simple scheme to minimize the metric in Eq. (8) by evaluating at different β_P with regular spacing and choosing the β_P that yields the minimal metric value. More sophisticated searches for the optimal β_P value are the subject of future work. Since the system matrix $\mathbf{A}$ is typically very large and is not computed explicitly, the inverse operation in Eq. (6) is solved in practice using a linear conjugate gradient (CG) method. Specifically, a system of linear equations are formed in a matrix form as $Zx = b$, in which matrix Z corresponds to the term $\mathbf{A}^T \mathbf{W} \mathbf{A} + \beta_R \Psi_R^T \mathbf{D}_R \Psi_R + \beta_P \Psi_P^T \mathbf{D}_P \Psi_P$ in Eq. (6), vector b corresponds to the term $\mathbf{A}^T \mathbf{W} l + \beta_P \Psi_P^T \mathbf{D}_P \Psi_P \mu_P$ in Eq. (6), and the vector x computed by a linear CG method corresponds to the image estimate $\hat{\mu}$ in Eq. (6). Since computing an approximate solution is much faster than doing a full image reconstruction, the proposed method is expected to be much more efficient than a traditional exhaustive search that requires full image reconstructions.

2.3 Spatially varying prior image strength map

The optimal prior image strength for a given anatomical change has been found to be shift-variant from previous work (Dang *et al* 2014a). That is, optimal strength designed at one location is not necessarily optimal for other locations even for identical anatomical changes. Thus, this shift-variance introduces a challenge for regularization design when the location of the change is not known *a priori*. To address this problem, a spatially varying β_P map is proposed, similar to other intentionally spatially varying regularization approaches such as for enforcing uniform resolution (Stayman and Fessler 2000) and optimization of task-based detectability (Gang *et al* 2014). Specifically, the approach performs individual optimizations of β_P at every possible location (given a specific presumed anatomical change positioned at each location) to form a β_P map that optimally admits change everywhere. Although repeating the proposed regularization design method at every location requires significantly

more computation time than for only one location, such designs can be performed at any time after the prior image scan and before the subsequent scan. For example, surveilling a solitary pulmonary nodule involves a time between scans of several months (Tan *et al* 2003), leaving ample time between any two adjacent scans for performing such design. Similarly in fractionated radiotherapy, there is typically several days between the planning scan and the first day of treatment, and one day between subsequent scans at each fraction of treatment. In practice, such design can also be accelerated by estimating the optimal β_P at each grid point and interpolating the results across the image. The design of a spatially varying β_P map is first demonstrated in the ellipse phantom in Sec. 3.3 and then used to reconstruct a solitary pulmonary nodule in the thorax phantom in Sec. 3.5.

2.4 Digital phantoms and simulation studies

Two digital phantoms are used in this work and are shown in Fig. 2. The ellipse phantom in Fig. 2(a) consists of three components: 1) a background ellipse (major axis 41 cm, minor axis 32.4 cm) with attenuation (0.021 mm^{-1}) similar to soft tissue; 2) a dense circular insert with attenuation (0.041 mm^{-1}) comparable to bone; 3) and a low-density circular insert with attenuation (0.001 mm^{-1}) close to air. These components together comprise the anatomical information in the previous scan. For subsequent scans, an anatomical change (a small disc) was introduced in one or both of the two locations shown in Fig. 2(a). The ellipse phantom is used in Sec. 3.3 to demonstrate the shift-variance of the optimal β_P and the design of a spatially varying β_P map, and used in Sec. 3.4 to study the dependence of the optimal β_P on three properties of an anatomical change (attenuation, shape, and size). In all of these studies, a system geometry was chosen that contained a 150 cm source-to-detector distance, 122 cm source-to-axis distance, and $0.556 \times 0.556 \text{ mm}^2$ detector pixel sizes. Only 20 projections equally distributed over 190° were acquired in the subsequent scan, representing a factor of 18 exposure reduction as compared to a typical complete scan (360 projections over 360°). The 20 projections were simulated using 10^6 photons per detector pixel with Poisson noise.

The thorax phantom in Fig. 2(b) was generated from an axial slice of a CT scan of a cadaver torso. We emulated a lung nodule surveillance scenario in which a uniform disc emulating a lung nodule (not present in the prior image scan but present in the subsequent scan) was placed in either the center or the periphery of the lung as shown in Fig. 2(b). The thorax phantom is used in Sec. 3.1 and 3.2 to evaluate the accuracy of the approximate analytical solution and predictive performance metric, and used in Sec. 3.5 to evaluate the design of a spatially varying β_P map in a lung nodule surveillance scenario. The projections were generated using the same system geometry and x-ray technique as for the ellipse phantom, except for a lower number of photons per detector pixel (10^5).

All PIRPLE reconstructions and image estimates used $340 \times 420 \times 1$ voxels for the ellipse phantom and $260 \times 300 \times 1$ voxels for the thorax phantom, both with $1 \times 1 \times 1 \text{ mm}^3$ voxel sizes. The Change ROI in the predictive performance metric was set to a circular region with a radius of 30 voxels (large enough to cover the change) as illustrated by the dashed circles in Fig. 2. It was found that changing the size of the Change ROI (still large enough to cover the change) did not change the results reported below; therefore, the size of the Change ROI

was fixed in this study. The size of the quadratic neighborhood δ in the Huber loss function was set to be 10^{-4} mm^{-1} as in (Dang *et al* 2014b). The image roughness strength β_R was fixed at 10^2 in the ellipse phantom experiments and 10 in the thorax phantom experiments.

2.5 Computational complexity and implementation

The computational complexity of both the proposed method (using P+C estimate) and the traditional exhaustive search are primarily determined by the total number of projection operations (forward projections and backprojections). Both methods use the same number of projection operations every iteration – CG iteration for the proposed method and SQS iteration for the exhaustive search (assuming pre-computed curvatures in SQS). However, the proposed method tends to require fewer iterations, because the objective function of the P+C estimate is quadratic and thereby easier to solve than the objective function in PIRPLE reconstruction which is not quadratic (not even guaranteed to be concave). The computation time of both methods are compared in the Results Section.

Both the PIRPLE reconstruction and image estimates were implemented in Matlab (The Mathworks, Natick MA), with the projection operations executed on GPU using CUDA-based libraries. The projection operations were implemented based on separable footprints (Long *et al* 2010). All experiments were performed on a workstation equipped with one GeForce GTX TITAN (Nvidia, Santa Clara CA) graphics card.

3. Results

3.1 Evaluation of approximate analytical solution

The proposed approximate analytical solution was evaluated in the thorax phantom by introducing a uniform disc emulating a lung nodule in the center of the lung as shown in Fig. 2(b). The uniform disc had 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 values for a solid solitary pulmonary nodule (Tan *et al* 2003). In this experiment, we assumed that the actual change could be exactly anticipated in the prospective regularization design and thereby set the presumed change used in the P+C estimate to be the same as the actual change. Figure 4 illustrates the PIRPLE reconstructions and image estimates using a number of β_P values. In the top row, the nodule reconstructed by PIRPLE exhibited low resolution and a high level of noise when using a very low β_P (due to high angular undersampling in the projection data) while it began to shrink or even disappear when using a very high β_P , demonstrating the importance of using appropriate prior image strength in PIBR. The Ideal estimate exhibited a high level of agreement with the PIRPLE reconstruction for all β_P , suggesting high accuracy in the quadratic approximation of the nonquadratic regularization given an ideal operating point. However, the Ideal estimate still requires a full PIRPLE reconstruction, which supersedes the need for an approximate solution. The P+C estimate did not use PIRPLE reconstruction results but exhibited some level of agreement with the PIRPLE reconstruction after one iteration of the P+C estimate and a high level of agreement after five iterations of the P+C estimate. These results suggest the possibility of approximating PIRPLE results without performing full image reconstructions.

The computation time of traditional exhaustive search (PIRPLE reconstruction) and our proposed method (P+C estimate) were also compared. For this dataset, PIRPLE reconstruction required 2000 SQS iterations to obtain a nearly converged image (RMSE less than 1 HU compared to a PIRPLE image formed using more than 20000 SQS iterations), while the P+C estimate only required 500 CG iterations to obtain a nearly converged image (similarly, RMSE less than 1 HU compared to a P+C estimate using more than 5000 CG iterations). The total computation time was ~180 seconds for a PIRPLE reconstruction and ~10 seconds for a P+C estimate. This suggests that our proposed method can reduce the computation time required for finding the optimal prior image strength by over an order of magnitude, compared to traditional exhaustive search.

3.2 Evaluation of predictive performance metric

The qualitative comparisons from the previous section are substantiated with quantitative measures in this section. Specifically, Figure 4 illustrates the evaluation of the proposed predictive performance metric at different β_P (with a uniform log spacing of $10^{0.1}$) using either PIRPLE reconstruction or one of the image estimate methods and the previously described experimental setup. We found that the proposed metric using PIRPLE reconstruction exhibited a single well-defined minimum in the range of β_P (10^0 to 10^7). The minimizer ($\beta_P = 10^{3.3}$) from this method was used as the ground truth for the optimal β_P . The metric using the Ideal estimate closely approximated the results using PIRPLE reconstructions, especially in the region where β_P was greater than 10^3 , and estimated the same optimum as the ground truth. This is consistent with the qualitative results in the previous section. When a (single iteration) P+C estimate was used, the metric plot yielded a similar shape and predicted a minimizer ($10^{3.4}$) very close to the ground truth. Using five iterations of the P+C estimate moved the curve much closer to the curve of PIRPLE reconstruction and predicted the same optimum as the ground truth. These results suggest that the proposed method can yield the same or very similar optimal β_P while having a significant computational advantage over traditional exhaustive search. The well-defined optimum in all plots – including the P+C estimates – suggests that direct minimization of the metric in Eq. (8) using more sophisticated optimization approaches may be possible and offer additional computational speedups.

3.3 Location dependence of regularization design in PIBR and evaluation of spatially varying prior image strength map

Thus far, we have presumed that the location of the anatomical change is known exactly. It remains a key question whether an optimal β_P designed for the presumed location remains optimal at other locations. To answer this question, the same change was introduced to either of the two locations in the ellipse phantom as shown in Fig. 2(a). The change was a disc with 0.05 mm^{-1} attenuation and 10 mm radius. Figure 5 shows the PIRPLE reconstructions and the predictive performance metric values at various β_P for the two locations. A clear difference in the optimal β_P for the two locations can be seen both in the PIRPLE reconstructions (optimal images outlined in black box) and in the metric curves. The optimal β_P was $10^{3.2}$ for Location I and $10^{4.5}$ for Location II, which differed over an order of magnitude. As a result, suboptimal reconstruction of the change could be seen at each of the two locations when using a β_P optimized for the other location, as seen in Fig. 5(a).

Interestingly, the optimal β_P was lower in Location I (higher attenuation area) and higher in Location II (lower attenuation area). This may be interpreted by recognizing the role of β_P in balancing the data fidelity term with the prior image regularization term in the PIRPLE objective function. The rays that go through Location II tend to have higher fluence than those that go through Location I, thereby leading to a larger value of the data fidelity term at Location II than at Location I. To maintain the same optimal balance between the two terms, a higher β_P is then needed at Location II than at Location I. Note also that the proposed method (using the P+C estimate after one iteration) predicted the same optimal β_P as the ground truth (using PIRPLE reconstruction) at both locations, again demonstrating the accuracy of the proposed method in estimating the optimal β_P .

Because one may not generally know where a change might occur, we also investigated the location-dependent prior image strength. Specifically, a spatially varying β_P map for the circular change mentioned above was generated by estimating the optimal β_P at each grid point of the image (with a spacing of 20 voxels in each dimension and a total number of 229 grid points) and interpolating across the image using radial basis functions (Powell 1987). Figure 6(a) shows a ground truth β_P map whose optimal β_P at each grid point was estimated using traditional exhaustive search. Spatial variations of the optimal β_P for the same change are seen throughout the image. Figure 6(b) shows a β_P map whose optimal β_P at each grid point was estimated by the proposed method. This map shows good agreement with ground truth, (RMSE = $10^{0.18}$) compared to the range of the optimal β_P within each map ($10^{3.2} \sim 10^{4.6}$). The total computation time for a β_P map was ~ 15 hours using traditional exhaustive search and below 1 hour using the proposed method, demonstrating the computational advantage of prospective design.

Figure 6(c) shows PIRPLE reconstructions of the same change at both locations using either a scalar β_P or a spatially varying β_P map. PIRPLE reconstruction using a scalar β_P optimized for Location I ($\hat{\beta}_P^I = 10^{3.2}$) resulted in accurate reconstruction of the change in Location I but streaks and low resolution for the change in Location II (especially visible in difference images). Similarly, PIRPLE reconstruction using a scalar β_P optimized for Location II ($\hat{\beta}_P^I = 10^{4.5}$) lead to accurate reconstruction of the change in Location II but a change with incorrect size in Location I. Whereas PIRPLE reconstruction using the spatially varying β_P map in Fig. 6(b) resulted in accurate reconstruction of both changes.

3.4 Attenuation, shape, and size dependence of regularization design in PIBR

The dependence of the optimal β_P on other properties of the anatomical change (besides location) including attenuation, shape, and size was also investigated. The ellipse phantom in Fig. 2(a) was used, and a change was introduced to Location II in the subsequent scan. The value of only one of the three properties was varied at a time. The proposed metric was evaluated at different β_P with a uniform log spacing of $10^{0.01}$. Since the optimal β_P estimated by the traditional exhaustive search and the proposed method were almost the same, only the optimal β_P estimated by the traditional exhaustive search are shown below.

First, the attenuation of the change was varied from 0.004 mm^{-1} to 0.060 mm^{-1} (i.e., -800 HU to 2000 HU assuming 0.02 mm^{-1} water attenuation) with an increment of 0.002 mm^{-1} ,

covering a broad range of possible changes – e.g. low-attenuating pulmonary ground-glass nodules (about -700 HU) (Funama *et al* 2009) to high-attenuating bones. The shape and size of the change were fixed to a 10 mm radius disc. Figure 7(a) shows the estimated optimal β_P as a function of the attenuation of the change. Note that the optimal β_P increased consistently as the attenuation increased, and the rate of the increase was higher for low attenuation changes (e.g., a change in soft tissue) and lower for high attenuation changes (e.g., a change in bone). This indicates a strong dependence of the optimal β_P on the attenuation of the change especially for low attenuation changes. The difference in the dependence of the optimal β_P between low attenuation changes and high attenuation changes may be explained by recognizing the effect of the use of too large β_P – that is, the use of too large β_P will enforce the reconstructed image to simply replicate the prior image, which prevents the change from being reconstructed in the image. Compared to high attenuation changes, low attenuation changes tend to be more vulnerable to such effect because they are more similar to the prior image. Therefore, their optimal β_P values have a stronger dependence on the attenuation of the change than high attenuation changes.

Shape of the change was varied to simulate different levels of morphologic irregularity (e.g. tumor speculation). Specifically, shape was varied by modeling anatomical changes with a shape whose radius varied as a function of angle using a sinusoid plus a constant. The amplitude of the sinusoid was varied from 0 mm (a circular disc) to 9 mm (highly spiculated) with an increment of 1 mm. Change attenuation was fixed at 0.02 mm^{-1} (e.g., soft tissue), and the size was fixed such that the area of the change was $770 \pm 1 \text{ mm}^2$ for every selected amplitude (This was achieved by tuning the mean radius constant in the shape model). Figure 7(b) shows that optimal β_P values exhibited only small variations as the amplitude of the sinusoidal contour increased, indicating a low dependence of optimal β_P on the shape of the change.

Lastly, the dependence on the size of the change was studied by varying the radius of a circular change from 3 mm to 20 mm with an increment of 1 mm. The attenuation was fixed to be 0.02 mm^{-1} (e.g., soft tissue). Figure 7(c) shows that the optimal β_P exhibited only small variations as the radius of the change increased, indicating a weak dependence of the optimal β_P on the size of the change.

The experiments on the shape and size of the change were also performed with respect to a high attenuation change (0.05 mm^{-1} attenuation), in which similarly weak dependence was observed. The results in this section together suggest that, when performing prospective regularization design, one may need to make sure that the attenuation of the presumed change is consistent with that of the actual change while such consistency may not need to be strictly enforced for the shape or size of the presumed change.

3.5 Evaluation of regularization design in lung nodule surveillance

Prospective regularization design was applied in a lung nodule surveillance scenario with a solitary pulmonary nodule. The nodule was not present in a previous baseline exam but is present 520 in a subsequent exam. The thorax phantom in Fig. 2(b) was used in this study and with a 6 mm radius nodule and an attenuation of 0.021 mm^{-1} (i.e., 50 HU). Presuming an unknown location, a spatially varying β_P map was generated using the method described

in Sec. 2.3. Fig. 8(a) shows a ground truth β_P map (generated using traditional exhaustive search) which exhibited a shift-variant optimal β_P within each side of the thoracic cavity as well as between sides. In this specific case, the optimal β_P was higher in the right cavity than in the left cavity, which is due to the asymmetry in the anatomy between the two sides. For example, the heart and trachea were not exactly centered and the lung in the right cavity had collapsed. The optimal β_P spanned almost an order of magnitude across the image ($10^{2.8} \sim 10^{3.6}$). Figure 8(b) shows a β_P map generated using the proposed method (including P+C estimate after one iteration), exhibiting good agreement with the ground truth β_P map and at the same time achieving about $\times 20$ reduction in computation time.

The spatially varying β_P map in Fig. 8(b) was then used in PIRPLE to reconstruct the actual solitary pulmonary nodule, which was simulated in the periphery of the right lung in the subsequent scan as shown in Fig. 2(b). The optimal β_P for a nodule at this location was found to be $10^{3.5}$ (using uniform log spacing of $10^{0.1}$). Figure 9 (a–d) shows a ROI of the current anatomy and images reconstructed by FBP, PIRPLE using a suboptimal scalar β_P , and PIRPLE using the spatially varying β_P map. The suboptimal scalar β_P was chosen to be $10^{2.8}$, which was optimal for the same nodule at the posterior of the left lung but was not optimal for the true change location. FBP image exhibited a high level of streaks and noise as a result of severe angular undersampling and lack of support from prior image information (RMSE = $43.7 \times 10^{-4} \text{ mm}^{-1}$). PIRPLE image using a suboptimal scalar β_P substantially reduced the streaks and noise, but still exhibited apparent error in the reconstructed nodule especially in the boundary of the nodule (error more pronounced the difference image) (RMSE = $9.2 \times 10^{-4} \text{ mm}^{-1}$). Finally, PIRPLE image using the β_P map exhibited excellent reconstruction of the nodule and the lowest error among all three methods (RMSE = $2.8 \times 10^{-4} \text{ mm}^{-1}$).

3.6 Evaluation of regularization design in a nodule disappearance scenario

The proposed regularization design method was also evaluated in a scenario in which a nodule was present in the prior image but not present in the current anatomy. This scenario is common in lung nodule surveillance, since benign nodules found in a previous exam can often be naturally resolved by the body before a follow-up exam is performed. Similar types of anatomical change can also be found in radiation therapy, in which tumor shrinks or disappears in follow-up exams after successful treatment. A simulation study was carried out on the thorax phantom, which was the same as the study described in Sec. 3.1 except that the lung nodule was present in the prior image but not in the current anatomy. Figure 10(a) illustrates the truth image and PIRPLE reconstructions using a number of β_P values. PIRPLE exhibited higher spatial resolution and a lower level of noise as β_P increased from $10^{1.5}$ to $10^{3.5}$, but exhibited features (false positive) as β_P kept increasing, indicating the importance of using proper prior image strength. Note that the false positive nodule information appeared first at the edge of the nodule and then at the interior of the nodule in the PIRPLE image. Fig. 10(b) shows that the proposed metric using PIRPLE reconstruction still exhibited a single well-defined minimum (ground truth) in the nodule disappearance scenario. Moreover, a (single iteration) P+C estimate was still able to predict a minimizer very close to the ground truth. These results suggest that accurate estimation of the optimal

prior image strength with substantial computational speedup can also be achieved in the nodule disappearance scenario.

4. Conclusions and Discussion

We have proposed a novel method that prospectively estimates the optimal prior image strength for accurate reconstruction of anatomical changes in PIBR. The approach uses an analytical approximation of PIBR objective functions and a predictive performance metric that leverages knowledge of a presumed change to estimate prior image strength that ensures accurate reconstruction of the change. The proposed prospective regularization strategy yields accurate estimates of the optimal prior image strength and substantially reduces computational time (by a factor of 20) compared to traditional exhaustive search.

A spatially varying prior image strength map was also introduced which optimally admits a presumed change everywhere in the image and eliminates the need to know the location of the change *a priori*. Optimal prior image strength was found to vary by at least an order of magnitude throughout the volume in phantom studies, indicating the potential importance of the spatially varying design for optimal performance. The optimal prior image strength was found to vary significantly with the attenuation difference associated with the anatomical change but was relatively insensitive to the shape and size of the change, suggesting accurate specification of change attenuation is important in regularization design in PIBR. Optimal penalty maps were found to improve the accuracy of lung nodule PIBR over uniformly penalized reconstructions. These results suggest great potential for the proposed method to provide prospective patient-, change-, and data-specific customization of the prior image strength to ensure reliable reconstruction of specific anatomical changes. While we have concentrated on the PIRPLE approach in this paper, one might form analogous regularization strength maps for other approaches that require a balance between current imaging data and prior image data. This includes other prior-image-based reconstruction (e.g. PICCS) as well as reconstruction of difference approaches (Pourmorteza *et al* 2015, Abbas *et al* 2013).

While the work reported in this paper provides a general strategy for prospective regularization design in PIBR, there are a number of potential developments that could further increase the utility of the underlying methodology. First, the predictive performance metric was solved in this work by evaluating the metric at different β_P values with regular spacing. While faster than traditional exhaustive search, additional acceleration may be found via more sophisticated minimization methods (e.g. simplex method, etc.). Such directed searches will be more computationally efficient and would be increasingly important for computing spatially varying maps with larger fields-of-view.

A second topic that needs to be investigated is the incorporation of image registration into the regularization design. Patient motion is commonly present between scans (e.g., due to patient repositioning or respiratory/cardiac motion) and needs to be accommodated to ensure accurate use of the prior image information in PIBR. For the purpose of regularization design, one might adopt a two-step approach – first estimating the patient motion by performing one PIRPLE (for rigid motion) or dPIRPLE (for nonrigid motion) reconstruction

using a nominal β_P value or using a dedicated image registration method such as in (Otake *et al* 2012), and then performing the proposed regularization design method with a prior image that has been deformed for motion compensation. While one might use the regularization design proposed in this work using the presumed change model with a perfectly registered prior image, the sensitivity of regularization design to registration errors also needs to be assessed.

The predictive performance metric proposed in this work estimates the prior image strength that minimizes the RMSE of the approximate analytical solution from the prior image plus the anticipated change. While this metric can yield image reconstructions with overall high accuracy, other metrics that are sensitive to specific imaging tasks including detectability index and various observer models (Siewerdsen and Antonuk 1998, Barrett and Myers 2004, Tward *et al* 2007, Qi and Huesman 2001, Gang *et al* 2014) should also be considered. Such task-based metrics could be used to find optimal prior image strength for particular abnormalities. For example, detectability index (Gang *et al* 2014) may be computed using the approximate analytical solution to estimate the optimal prior image strength for detecting high-contrast, low-frequency lesions such as a solitary pulmonary nodule in lung nodule surveillance.

While this work has focused on anatomical change within a lung nodule surveillance scenario, there are many other potential applications of optimized PIBR regularization including accurate visualization of resections, device/implant placement, and monitoring of other treatments and interventions. While each of these potential applications deserves additional investigations into the specific challenges associated with PIBR in each area, the proposed methodology for balancing prior image information with measurement data is general. Optimization of this balance is critical for reliable reconstruction and this work represents an important step in providing a degree of robustness and controllability for PIBR approaches. Such reliability is a necessity for delivering on the huge potential of PIBR methods and finding more widespread clinical adoption.

Acknowledgments

This work was supported in part by an academic-industry partnership with Varian Medical Systems (Palo Alto, CA) and by the National Institutes of Health Grant No. 2R01-CA-112163.

References

- Abbas S, Min J, Cho S. Super-sparsely view-sampled cone-beam CT by incorporating prior data. *J Xray Sci Technol.* 2013; 21:71–83. [PubMed: 23507853]
- Barrett, HH.; Myers, KJ. *Foundations of Image Science.* New York: Wiley; 2004.
- Caillot D, Couaillier JF, Bernard a, Casasnovas O, Denning DW, Mannone L, Lopez J, Couillault G, Piard F, Vagner O, Guy H. Increasing volume and changing characteristics of invasive pulmonary aspergillosis on sequential thoracic computed tomography scans in patients with neutropenia. *J Clin Oncol.* 2001; 19:253–9. [PubMed: 11134220]
- Candès EJ, Romberg J, Tao T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans Inf Theory.* 2006; 52:489–509.
- Chen G-H, Tang J, Leng S. Prior Image Constrained Compressed Sensing (PICCS). *SPIE Medical Imaging.* 2008; 6856:685618.

- Cleary K, Peters TM. Image-Guided Interventions: Technology Review and Clinical Applications. *Annu Rev Biomed Eng.* 2010; 12:119–42. [PubMed: 20415592]
- Dang H, Otake Y, Schafer S, Stayman JW, Kleinszig G, Siewerdsen JH. Robust methods for automatic image-to-world registration in cone-beam CT interventional guidance. *Med Phys.* 2012; 39:6484–98. [PubMed: 23039683]
- Dang, H.; Siewerdsen, JH.; Stayman, JW. SPIE Medical Imaging. 2014a. Regularization design and control of change admission in prior-image-based reconstruction; p. 90330O
- Dang H, Stayman JW, Sisniega A, Xu J, Zbijewski W, Wang X, Foos DH, Aygun N, Koliatsos VE, Siewerdsen JH. Statistical reconstruction for cone-beam CT with a post-artifact-correction noise model: application to high-quality head imaging. *Phys Med Biol.* 2015a; 60:6153–75. [PubMed: 26225912]
- Dang H, Stayman JW, Sisniega A, Xu J, Zbijewski W, Yorkston J, Aygun N, Koliatsos V, Siewerdsen JH. Cone-beam CT of traumatic brain injury using statistical reconstruction with a post-artifact-correction noise model. *SPIE Medical Imaging.* 2015b; 9412:941207.
- Dang H, Wang aS, Sussman MS, Siewerdsen JH, Stayman JW. dPIRPLE: a joint estimation framework for deformable registration and penalized-likelihood CT image reconstruction using prior images. *Phys Med Biol.* 2014b; 59:4799–826. [PubMed: 25097144]
- Dang H, Wang AS, Zhao Z, Sussman MS, Siewerdsen JH, Stayman JW. Joint Estimation of Deformation and Penalized-Likelihood CT Reconstruction Using Previously Acquired Images. *International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine.* 2013:424–7.
- Funama Y, Awai K, Liu D, Oda S, Yanaga Y, Nakaura T, Kawanaka K, Shimamura M, Yamashita Y. Detection of nodules showing ground-glass opacity in the lungs at low-dose multidetector computed tomography: phantom and clinical study. *J Comput Assist Tomogr.* 2009; 33:49–53. [PubMed: 19188784]
- Gang GJ, Stayman JW, Zbijewski W, Siewerdsen JH. Task-based detectability in CT image reconstruction by filtered backprojection and penalized likelihood estimation. *Med Phys.* 2014; 41:081902. [PubMed: 25086533]
- Hasegawa M, Sone S, Takashima S, Li F, Yang ZG, Maruyama Y, Watanabe T. Growth rate of small lung cancers detected on mass CT screening. *Br J Radiol.* 2000; 73:1252–9. [PubMed: 11205667]
- Huber, PJ. *Robust Statistics.* New York: Wiley; 1981.
- Jaffray DA, Siewerdsen JH, Wong JW, Martinez AA. Flat-panel cone-beam computed tomography for image-guided radiation therapy. *Int J Radiat Oncol Biol Phys.* 2002; 53:1337–49. [PubMed: 12128137]
- Lauzier PT, Chen G-H. Characterization of statistical prior image constrained compressed sensing (PICCS): II. Application to dose reduction. *Med Phys.* 2013; 40:021902. [PubMed: 23387750]
- Long Y, Fessler JA, Balter JM. 3D forward and back-projection for X-ray CT using separable footprints. *IEEE Trans Image Process.* 2010; 29:1839–50.
- Meeks SL, Harmon JF, Langen KM, Willoughby TR, Wagner TH, Kupelian PA. Performance characterization of megavoltage computed tomography imaging on a helical tomotherapy unit. *Med Phys.* 2005; 32:2673–81. [PubMed: 16193798]
- Navab N, Heining S-M, Traub J. Camera augmented mobile C-arm (CAMC): Calibration, accuracy study, and clinical applications. *IEEE Trans Med Imaging.* 2010; 29:1412–23. [PubMed: 20659830]
- Otake Y, Schafer S, Stayman JW, Zbijewski W, Kleinszig G, Graumann R, Khanna AJ, Siewerdsen JH. Automatic localization of vertebral levels in x-ray fluoroscopy using 3D-2D registration: a tool to reduce wrong-site surgery. *Phys Med Biol.* 2012; 57:5485–508. [PubMed: 22864366]
- Pouliot J, Bani-Hashemi A, Josephine Chen, Svatos M, Ghelmansarai F, Mitschke M, Aubin M, Xia P, Morin O, Bucci K, Roach M, Hernandez P, Zheng Z, Hristov D, Verhey L. Low-dose megavoltage cone-beam CT for radiation therapy. *Int J Radiat Oncol Biol Phys.* 2005; 61:552–60. [PubMed: 15736320]
- Pourmorteza A, Dang H, Siewerdsen JH, Stayman JW. Reconstruction of Difference using Prior Images and a Penalized-Likelihood Framework. *International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine.* 2015:252–5.

- Powell, MJ. Algorithms for approximation. Oxford: Clarendon Press; 1987.
- Qi J, Huesman RH. Theoretical study of lesion detectability of MAP reconstruction using computer observers. *IEEE Trans Med Imaging*. 2001; 20:815–22. [PubMed: 11513032]
- Sauer K, Bouman C. A local update strategy for iterative reconstruction from projections. *IEEE Trans Signal Process*. 1993; 41:534–48.
- Siewerdsen JH, Antonuk LE. DQE and System Optimization for Indirect-Detection Flat-Panel Imagers in Diagnostic Radiology. *SPIE Medical Imaging*. 1998; 3336:546–55.
- Siewerdsen JH, Daly MJ, Chan H, Nithiananthan S, Hamming N, Brock KK, Irish JC. High-performance intraoperative cone-beam CT on a mobile C-arm: an integrated system for guidance of head and neck surgery. *SPIE Medical Imaging*. 2009; 7261:72610J.
- Stayman JW, Dang H, Ding Y, Siewerdsen JH. PIRPLE: a penalized-likelihood framework for incorporation of prior images in CT reconstruction. *Phys Med Biol*. 2013; 58:7563–82. [PubMed: 24107545]
- Stayman JW, Fessler JA. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Trans Med Imaging*. 2000; 19:601–15. [PubMed: 11026463]
- Stayman JW, Prince JL, Siewerdsen JH. Information Propagation in Prior-Image-Based Reconstruction. *The First International Conference on Image Formation in X-ray Computed Tomography*. 2012:334–8.
- Stayman JW, Zbijewski W, Otake Y, Uneri A, Schafer S, Lee J, Prince JL, Siewerdsen JH. Penalized-Likelihood Reconstruction for Sparse Data Acquisitions with Unregistered Prior Images and Compressed Sensing Penalties. *SPIE Medical Imaging*. 2011; 7961:79611L.
- Tan BB, Flaherty KR, Kazerooni EA, Iannettoni M. The solitary pulmonary nodule. *Chest*. 2003; 123:89S–96S. [PubMed: 12527568]
- Thibault J-B, Sauer KD, Bouman CA, Hsieh J. A three-dimensional statistical approach to improved image quality for multislice helical CT. *Med Phys*. 2007; 34:4526–44. [PubMed: 18072519]
- Tward DJ, Siewerdsen JH, Daly MJ, Richard S, Moseley DJ, Jaffray DA, Paul NS. Soft-tissue detectability in cone-beam CT: evaluation by 2AFC tests in relation to physical performance metrics. *Med Phys*. 2007; 34:4459–71. [PubMed: 18072510]
- Wang AS, Stayman JW, Otake Y, Kleinszig G, Vogt S, Gallia GL, Khanna AJ, Siewerdsen JH. Soft-tissue imaging with C-arm cone-beam CT using statistical reconstruction. *Phys Med Biol*. 2014; 59:1005–26. [PubMed: 24504126]
- Xing L, Thorndyke B, Schreiber E, Yang Y, Li TF, Kim GY, Luxton G, Koong A. Overview of image-guided radiation therapy. *Med Dosim*. 2006; 31:91–112. [PubMed: 16690451]
- Xu J, Tsui BMW. C-arm CT Image Reconstruction From Sparse Projections. *International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*. 2013:34–7.

Figure 1.

Approximation of the modified penalty function f (equivalent to the Huber loss function in this work) with a quadratic function g about an operating point τ_i . The operating point is selected either within (a) or outside of (b) the quadratic neighborhood $([-\delta, \delta])$ of the function f .

Figure 2.

(a) Ellipse phantom with attenuation (mm^{-1}): 0.021 (background ellipse), 0.041 (bright circular insert), and 0.001 (dark circular insert). An anatomical change is introduced in one or both of Location I and Location II in the subsequent scan. The dashed circles illustrate the Change ROIs used in the predictive performance metric. (b) Thorax phantom generated from an axial slice of a CT scan of a cadaver torso. A uniform circular lung nodule was introduced in the subsequent scan in the center or the periphery of the lung to emulate a lung nodule surveillance scenario.

Figure 3. PIRPLE reconstruction, Ideal estimate, P+C estimate (1 iteration), and P+C estimate (5 iterations) of a simulated circular solitary pulmonary nodule at various β_P . Note that too small β_P or too large β_P leads to little benefits or false features (false negative) in PIRPLE reconstruction. Only 20 projections equally distributed over 190° were used. (Grayscale window: $[0 \ 0.04] \text{ mm}^{-1}$.)

Figure 4.

Evaluation of the proposed predictive performance metric at different β_P (with a uniform log spacing of $10^{0.1}$) using either PIRPLE reconstruction or one of the image estimation methods. Note that all the methods yielded almost the same optimal β_P , while the proposed method (using P+C estimate) does not require full image reconstruction.

Figure 5.

Comparison of PIRPLE reconstructions and the prospective performance metric for the same change at two different locations. A clear difference in the optimal $430 \beta_P$ (over an order of magnitude) between the two locations can be seen both in the PIRPLE reconstructions and in the metric curves. These results motivate the design of a spatially varying β_P map. (Grayscale window: $[0.03 \ 0.052] \text{ mm}^{-1}$ for Location I and $[-0.02 \ 0.052] \text{ mm}^{-1}$ for Location II.)

Figure 6.

(a–b) Spatially varying β_P map generated using traditional exhaustive search (a) or the proposed method (b). The β_P map in (b) exhibits good agreement with the β_P map in (a). (c) PIRPLE reconstruction of the same change at both locations using a scalar β_P optimized for Location I ($\hat{\beta}_P^I=10^{3.2}$) a scalar β_P optimized for Location II ($\hat{\beta}_P^{II}=10^{4.5}$), and the spatially varying β_P map in (b). The parameter ε stands for the RMSE of the PIRPLE reconstruction with respect to the truth image. (Grayscale window: $[0.03 \ 0.052]$ mm⁻¹ for reconstructed images at Location I, $[-0.02 \ 0.052]$ mm⁻¹ for reconstructed images at Location II, and $[-0.01 \ 0.01]$ mm⁻¹ for all the difference images.)

Figure 7.

Illustration of the dependence of the optimal β_P on the attenuation (a), shape (b), and size (c) of the anatomical change. One anatomical change was introduced to Location II of the ellipse phantom in Fig. 2(a), and only one of the three properties of the change mentioned above was varied at a time. The optimal β_P was estimated by evaluating the proposed metric at different β_P with a uniform log spacing of $10^{0.01}$. A negative exponential function was fit to the data points in (a) to help illustrate the relationship in (a). (Grayscale window: $[0 \ 0.052] \text{ mm}^{-1}$.)

Figure 8.

Spatially varying β_P map generated using traditional exhaustive search (a) or the proposed method (b) for optimally admitting a solitary pulmonary nodule everywhere in both sides of the thorax cavity. The optimal β_P was estimated on an image grid with a spacing of 20 voxels in each dimension and then interpolated into a β_P map using radial basis functions. The β_P map in (b) exhibited good agreement with the β_P map in (a) while reducing the computation time by a factor of ~ 20 .

Figure 9.

Image reconstruction of a solitary pulmonary nodule which was not present in a baseline exam (e) but appeared in the periphery of the right lung in the subsequent exam (a). (b–d) A ROI of the images reconstructed by FBP, PIRPLE using a suboptimal scalar β_P and PIRPLE using the spatially varying β_P map in Fig. 8(b). (f–h) Difference image between each of the images in (b–d) and the current anatomy. The parameter ε stands for the RMSE of the difference image in (f–h). (Grayscale window: $[0 \ 0.04] \text{ mm}^{-1}$ for reconstructed images and $[-0.01 \ 0.01] \text{ mm}^{-1}$ for difference images.)

Figure 10.

(a) PIRPLE reconstruction of the current anatomy (labeled as “Truth”) at various β_P . A lung nodule was present in the prior image but not present in the current anatomy. Note that too small β_P or too large β_P leads to little benefits or false features (false positive) in PIRPLE reconstruction. (Grayscale window: $[0 \ 0.04] \text{ mm}^{-1}$) (b) Evaluation of the proposed metric at different β_P in the nodule disappearance scenario. Note that all the methods yielded almost the same optimal β_P , while the proposed method (P+C estimate) does not require full image reconstruction.

Table 1

Summary of different ways of selecting the operating point for the approximate analytical solution of PIRPLE. An operating point is not needed if doing a full PIRPLE reconstruction (first row), while it is needed and defined as the PIRPLE solution in Ideal estimate (second row), defined as the sum of a prior image and some presumed change in P+C estimate (third row), defined as the results from the $(n-1)^{\text{th}}$ iteration P+C estimate in the n^{th} iteration P+C estimate row).

Method Name	$\tilde{\mu}$ in the Operating Point	Output Image
PIRPLE Reconstruction	N/A	$\hat{\mu}_{\text{PIRPLE}}$
Ideal Estimate	$\hat{\mu}_{\text{PIRPLE}}$	$\hat{\mu}_{\text{Ideal}}$
P+C Estimate (1 iteration)	$\mu_P + \mu_C$	$\hat{\mu}_{P+C}^1$
P+C Estimate (n iterations)	$\hat{\mu}_{P+C}^{n-1}$	$\hat{\mu}_{P+C}^n$