Published in final edited form as:

Phys Med Biol. 2014 September 7; 59(17): 4799–4826. doi:10.1088/0031-9155/59/17/4799.

dPIRPLE: A Joint Estimation Framework for Deformable Registration and Penalized-Likelihood CT Image Reconstruction using Prior Images

H. Dang¹, A. S. Wang¹, Marc S. Sussman², J. H. Siewerdsen¹, and J. W. Stayman¹

¹Department of Biomedical Engineering, Johns Hopkins University, Baltimore MD 21205

²Department of Surgery, Johns Hopkins University, Baltimore MD 21224

Abstract

Sequential imaging studies are conducted in many clinical scenarios. Prior images from previous studies contain a great deal of patient-specific anatomical information and can be used in conjunction with subsequent imaging acquisitions to maintain image quality while enabling radiation dose reduction (e.g., through sparse angular sampling, reduction in fluence, etc.). However, patient motion between images in such sequences results in misregistration between the prior image and current anatomy. Existing prior-image-based approaches often include only a simple rigid registration step that can be insufficient for capturing complex anatomical motion, introducing detrimental effects in subsequent image reconstruction. In this work, we propose a joint framework that estimates the 3D deformation between an unregistered prior image and the current anatomy (based on a subsequent data acquisition) and reconstructs the current anatomical image using a model-based reconstruction approach that includes regularization based on the deformed prior image. This framework is referred to as deformable prior image registration, penalized-likelihood estimation (dPIRPLE). Central to this framework is the inclusion of a 3D B-spline-based free-form-deformation model into the joint registration-reconstruction objective function. The proposed framework is solved using a maximization strategy whereby alternating updates to the registration parameters and image estimates are applied allowing for improvements in both the registration and reconstruction throughout the optimization process. Cadaver experiments were conducted on a cone-beam CT testbench emulating a lung nodule surveillance scenario. Superior reconstruction accuracy and image quality were demonstrated using the dPIRPLE algorithm as compared to more traditional reconstruction methods including filtered backprojection, penalized-likelihood estimation (PLE), prior image penalized-likelihood estimation (PIPLE) without registration, and prior image penalized-likelihood estimation with rigid registration of a prior image (PIRPLE) over a wide range of sampling sparsity and exposure levels.

1. Introduction

There are many clinical scenarios in which sequential imaging studies are conducted. For example, in diagnostic radiology, patients with suspicious lung nodules or cancer often

undergo a series of longitudinal chest x-ray CT scans to access the tumor growth rate (Hasegawa *et al.*, 2000) and/or monitor the therapy response (Werner-Wasik *et al.*, 2001). In image-guided interventions, pre-operative CT or MRI studies are often acquired, and multiple intra-operative C-arm Cone Beam CT (CBCT) studies may be acquired for up-to-date tracking, registration, and visualization. (Cleary and Peters, 2010; Siewerdsen *et al.*, 2009; Schafer *et al.*, 2011; Dang *et al.*, 2012; Navab *et al.*, 2010) Traditionally, such sequential studies are conducted through a series of *complete* acquisitions, and the cumulative radiation dose in such studies raises concerns for both the patient and the surgical staff. To address this problem, numerous methods for reducing radiation dose have been developed. Two straightforward methods are illustrated in Figure 1: 1) reduction in the amount of data acquired (i.e. sparse sampling), and 2) reduction of the x-ray technique (e.g., reduced tube current). However, such dose reduction techniques in the absence of countermeasures to mitigate noise and/or artifacts tend to experience a reduction in overall image quality of reconstructed images. For example, sparse sampling makes the reconstruction problem more ill-posed and leads to streak artifacts in the reconstructed image, while lowering fluence increases quantum noise in the projection data and image reconstruction.

Model-based reconstruction (MBR) such as penalized-likelihood estimation (PLE) has demonstrated improved trade-offs between radiation dose and image quality over traditional direct reconstruction approaches. (Thibault *et al.*, 2007; Evans *et al.*, 2011; Wang *et al.*, 2014) Such improvements arise from two sources. First, MBR allows the incorporation of accurate models for a variety of physical detection processes associated with quantum noise, detector blur, Figure 1: Illustration of two major techniques for dose reduction and the motivation system geometry, and spectral effects (Nuyts *et al.*, 2013) into the reconstruction process permitting more efficient use of the (noisy) measurements. Second, MBR allows the addition of other sources of information to enforce or encourage reconstructions with desirable image properties. For example, image reconstruction can be regularized using assumptions on image smoothness (Lange, 1990), edge-preservation (Vogel and Oman, 1996; Sidky and Pan, 2008), and self-similarity of the object via dictionary learning (Qiong *et al.*, 2012; Yang *et al.*, 2012). However, additional sources of information used in MBR usually entail only very general properties of the reconstructed image and typically do not incorporate patient-specific prior information. In sequential imaging studies, patient-specific prior knowledge, such as a prior image acquired in a previous imaging study, contains a tremendous amount of patient-specific anatomical information that could potentially be leveraged in MBR, presenting increased opportunities for dose reduction.

The importance of prior images in image reconstruction has been recognized in recent years in a number of prior-image-based reconstruction (PIBR) approaches. In Prior Image Constrained Compressed Sensing (PICCS) (Chen *et al.*, 2008), Chen *et al.* formed an objective function that seeks sparse differences between the reconstruction of current anatomy and a prior image, using compressed sensing concepts whereby a sparse signal can be recovered by L^1 minimization under certain assumptions (Candes *et al.*, 2006). The original formulation of PICCS has also been modified with the inclusion of statistical weightings (Lauzier *et al.*, 2013) and can be applied in situations where the forward model

can be transformed into a linear relationship between the image volume and processed line integrals. In contrast, (Stayman *et al.*, 2013) presented a PIBR approach termed Prior Image Registration, Penalized-Likelihood Estimation (PIRPLE), employing a MBR framework that integrates: 1) a statistical objective function with nonlinear forward model and noise model for the unprocessed measurements; and 2) a generalized regularization term based on a prior image. This framework permits flexibility in the selection of the forward model and the noise model and does not necessitate a linearizable forward model. In addition to direct use of a patient-specific prior image in the reconstruction objective function, prior images have also been utilized indirectly, such as in computing the weights of a roughness penalty. (Xu and Tsui, 2013)

Although the use of prior images in PIBR dramatically reduces the data fidelity requirements and demonstrates good image quality under conditions of substantial downsampling and photon starvation (Stayman *et al.*, 2011; Ding *et al.*, 2012), a critical aspect of effectively using prior images is the ability to compensate for deformation between the prior image and subsequent image acquisitions. Such deformation is typically caused by patient motion between the baseline scan (i.e., the scan that forms the prior image) and the follow-up scan (i.e., the scan that acquires measurements of current anatomy). Potential sources of mismatch between scans include re-positioning of a patient and acquisition during differing physiological states (e.g., acquisitions at different phases of respiratory or cardiac motion). If these mismatches are not compensated or are incorrectly compensated, incorrect information from the misaligned prior image will be injected into the image reconstruction. Moreover, PIBR without registration cannot differentiate between *true anatomical changes* (e.g. tumor growth, bone drill-out, tissue ablation, etc.) and *changes due to motion*. Ambiguity between these two types of changes in the PIBR results would make true anatomical change difficult to recognize and may introduce false anatomical changes.

Compensation of deformation was typically not considered in the PICCS implementations for cardiac imaging (Chen *et al.*, 2008) and dose reduction (Lauzier *et al.*, 2013) or in (Xu and Tsui, 2013), because in these applications a prior image was formed from a full set of current measurements. Nett *et al.* adopted a staged approach using a preregistration of the prior image followed by PICCS reconstruction in the context of C-arm interventional imaging where patient motion is commonly seen between the baseline and follow-up scans. (Nett *et al.*, 2009) Specifically, in this work the prior image was registered to a filtered back-projection (FBP) of sparse measurements of the current anatomy using a rigid 6 degree-of-freedom (DOF) transformation. However, the 6 DOF registration may not accurately capture the non-rigid nature of organ motion, and the accuracy of the staged registration will be limited by the image quality of initial sparse FBP reconstruction. Similarly, the initial PIRPLE implementation (Stayman *et al.*, 2013) entailed a rigid registration of the prior, and although the registration estimate was joint (not staged) with the PL reconstruction estimate, the 6 DOF pose estimate does not resolve deformations between the baseline and follow-up scans. Other methods have been developed that deformably register a prior image acquired from a planning CT with sparse projections from a subsequent CBCT. (Zeng *et al.*, 2007; Ren *et al.*, 2012; Wang and Gu, 2013) However, in these methods, the follow-up data were

used to estimate patient motion but not to reconstruct new images. Deformable registration of a prior image has also been used recently for artifact correction. (Heußner *et al.*, 2014)

A *joint* estimation of both the deformation and image reconstruction may be used to overcome the limitations of staged registration. The idea of joint estimation has been widely studied in MBR in many modalities. (Gilland *et al.*, 2002; Chun and Fessler, 2009; Fessler, 2010; Yang *et al.*, 2013; Taguchi *et al.*, 2007) For example, in cardiac gated ECT, Gilland *et al.* designed an objective function that jointly estimated cardiac images at two different frames and the cardiac motion between the two frames (Gilland *et al.*, 2002), thereby solving both deformation and attenuation together using a conjugate gradient method. In PET, Fessler used an objective function that jointly estimated a single image (at a specific time point) and a series of deformation fields (used to match motions at other time points), achieving a joint solution by alternately updating the deformation and attenuation parameters. (Fessler, 2010) In digital breast tomosynthesis, Yang *et al.* presented a framework that coupled the limited-angle reconstruction of the breast and registration of two temporal breast images and showed that the joint method gave superior results with respect to reconstruction fidelity and registration accuracy than traditional sequential method. (Yang *et al.*, 2013) Despite the varied use of joint estimation in MBR, the joint approach has received less attention in PIBR.

In this paper, we present a model-based approach that incorporates a Poisson noise model and a high-quality patient-specific prior image to reconstruct images from sparse and/or noisy measurements. Deformation between the baseline and follow-up scan introduced by patient motion is estimated jointly through a cubic B-spline-based free-form deformation (FFD) model. By extension of the PIRPLE framework, we refer to this approach as deformable Prior Image Registration, Penalized-Likelihood Estimation (dPIRPLE). An introduction to the approach and initial evaluation using rigid registration in 2D simulation and phantom studies were presented in (Stayman *et al.*, 2013), and preliminary investigation of incorporating 3D deformable registration was presented in (Dang *et al.*, 2013). In this work, we present a detailed description of the dPIRPLE algorithm and an alternating maximization strategy for solving dPIRPLE. Additionally, we analyze the convergence properties and the scheduling of registration / reconstruction updates, and we perform qualitative and quantitative comparisons of reconstruction results from dPIRPLE and other algorithms under various conditions of data sparsity and exposure levels.

2. Methods

2.1 Forward model and Penalized Likelihood Estimation (PLE)

For transmission tomography, we choose to model the mean measurements of x-ray photons $\bar{y}$ after attenuation by an object using the following forward model under the assumption of a monochromatic x-ray beam:

$$\bar{y} = \mathbf{D} \{g\} \exp(-l), l = \mathbf{A}\mu \quad (1)$$

where g denotes a $N \times 1$ vector of measurement-dependent gains (i.e., measurements of photons when no object is present), $\mathbf{D}$ denotes an operator converting a vector to a diagonal

matrix, l is a $N \times 1$ vector of line integrals, μ is a $M \times 1$ vector of linear attenuation coefficients of a discretized object, and $\mathbf{A}$ is a $N \times M$ matrix representing the linear projection operator. N and M are the total number of measurements and the total number of voxels in the reconstruction, respectively. In this work, the system matrix is not precomputed. The action of $\mathbf{A}$ (projection) and $\mathbf{A}^T$ (backprojection) is computed functionally using on-the-fly calculations. The specific projector set is a modified form (Wang *et al.*, 2014) of the separable footprint projector set in (Long *et al.*, 2010).

One may adopt an arbitrary noise model to account for the variations between mean measurements $\bar{y}_i$ and individual noisy measurement y_i ¹. Here we choose a Poisson model which yields the following log-likelihood function (ignoring the constant term)

$$\log L(y; \mu) = \sum_{i=1}^N h_i([\mathbf{A}\mu]_i) \quad h_i(l_i) = \sum_{i=1}^N y_i \log(g_i \exp(-l_i)) - g_i \exp(-l_i) \quad (2)$$

where h_i represents a marginal log-likelihood for each measurement and $[\cdot]_i$ denotes an operator that selects the i^{th} element of a vector. Because direct maximization of the log-likelihood function can result in noisy images, penalty terms, $R(\mu)$, are often added to encourage desirable image properties. So-called penalized-likelihood estimation (PLE) can be written in the following general form

$$\hat{\mu} = \underset{\mu}{\operatorname{argmax}} \log L(y; \mu) - R(\mu), \quad R(\mu) = \beta_R \|\Psi_R \mu\|_{p_R}^{p_R} \quad (3)$$

In this case, we have chosen a specific form of the penalty, which includes the operator Ψ_R , a p -norm metric (with $p = p_R$ and an exponent p_R), and a scalar regularization strength β_R . This general form of penalty allows for varied control of image properties and has a number of different interpretations including penalties on roughness (Lange, 1990), total variation (Vogel and Oman, 1996), and other decompositions of the image. (E.g., one may select Ψ_R to be an arbitrary sparsifying transform, particularly for $p_R = 1$, which encourages sparse solutions in the transformed domain.) For this penalty, we will concentrate on a Ψ_R operator that computes the pairwise difference between voxels in a first-order neighborhood around each voxel as in traditional roughness penalties.

Among popular approaches for solving PLE, we chose the separable paraboloidal surrogates (SPS) approach (Erdogan and Fessler, 1999) that allows for highly parallelizable image updates. To satisfy the five conditions (Erdogan and Fessler, 1999) of finding a parabolic surrogate in SPS approach, we introduced a modified p -norm (Stayman *et al.*, 2013) which replaced a δ -neighborhood about the origin in the traditional p -norm with a quadratic function. This modification ensures the differentiability of the p -norm operator at the origin. Both the function values and the derivatives match at the transition point $\pm\delta$ after the modification. Throughout this work, we will use either $p = 1$ or $p = 2$, and the modified p -norm becomes equivalent to the Huber loss function (Huber, 1981) or the standard L^2 norm,

¹The subscript i denotes the i^{th} element of the subscripted vector.

respectively. In both cases, the modified p -norm can be easily shown to satisfy the conditions for application of SPS.

2.2. Deformable Prior Image Registration, Penalized-Likelihood Estimation (dPIRPLE)

Prior images typically contain a great deal of patient-specific anatomical information. Thus, such images have great potential for regularization of the reconstruction problem and consequent dose reduction. One specific way to do this is to modify (3) with an additional penalty term (referred to as the *prior image penalty*), which encourages similarity between the estimated image and the prior image by penalizing their differences. However, one must recognize that changes between a prior image and the current patient anatomy can be introduced in two possible ways: 1) *true anatomical changes*, which are the result of disease progression or surgical interventions (e.g. tumor growth, bone drill-out); and 2) *changes due to motion*, which are caused by patient motion between scans (e.g. patient re-positioning, respiratory or cardiac motion). Thus, directly enforcing similarity to a prior image without the compensation of *changes due to motion* will limit the utility of the prior image and potentially provide incorrect information due to the misregistration. In other words, “efficient” or “complete” extraction of information from a prior image is only possible when proper registration is used to eliminate mismatches due to motion. Moreover, recognizing that the accuracy of an initial one-time registration procedure (e.g., staged registration followed by reconstruction) is limited by registration of low-fidelity data, we choose to incorporate registration parameters into a *joint* image reconstruction and registration objective function - solving for both registration and reconstruction parameters. This joint approach is expected to achieve improved results since updated image estimates can help to refine registration estimates and vice versa. We refer to this approach as deformable Prior Image Registration, Penalized-Likelihood Estimation (dPIRPLE), and its objective function can be written as

$$\{\hat{\mu}, \hat{\lambda}\} = \underset{\mu, \lambda}{\operatorname{argmax}} \Phi(\mu, \lambda) = \underset{\mu, \lambda}{\operatorname{argmax}} \log L(y; \mu) - \beta_R \|\Psi_R \mu\|_{p_R}^{p_R} - \beta_P \|\Psi_P (\mu - \mathbf{W}(\lambda) \mu_P)\|_{p_P}^{p_P} \quad (4)$$

where the last term denotes the prior image penalty with μ_P denoting the prior image. The prior image penalty term in (4) has a distinct set of parameters (β_P and p_P), which allows control of this penalization independent from that of the roughness penalty. Varying β_P controls the strength of prior image information in the reconstruction. The parameter p_P also affects the balance of prior information relative to other terms and can be freely chosen to penalize the difference image using different norms. For example, $p_P=2$ tends to apply larger penalties for greater differences, thereby enforcing smooth differences and blending features in the prior image and current measurements. In contrast, $p_P=1$ tends to apply relatively smaller penalties for greater differences and therefore allowing or encouraging large and sparse differences. Similar ideas can be found in compressed sensing theory that recovers sparse solutions by using an L^1 norm constraint. (Candes *et al.*, 2006) This norm selection is important to minimize bias in regions where change occurs between the prior and current scans. Specifically, in this study, we use $p_P=1$ which encourages similarity between the current reconstruction and prior image data but allows for potentially large but sparse differences. With proper selection of β_P , this penalty is thereby expected to allow the sparse

anatomical change (e.g. lung nodule) to appear in the reconstructed image and not to significantly bias the reconstructed image towards the prior data in the region of anatomical change. In this case the operator Ψ_P can be interpreted as a sparsifying operator. Options for this operator include an image gradient operator (e.g., finite differences on local voxels) or the identity matrix (if image changes are already sparse).

In (4), $\mathbf{W}(\lambda)$ represents a deformation operator that is a function of the deformation parameters λ . While there exist many methods to deformably register a 3D volume, we choose cubic B-spline-based FFD (Rueckert, 1999) as the deformation model. Specifically, we can write

$$\mathbf{W}_x(\lambda) = x + \sum_{x_i \in N_x} \lambda_i \beta \left(\frac{x - x_i}{\sigma} \right) \quad (5)$$

where $\beta(\cdot)$ is the tensor product of cubic B-spline functions, x_i are the control points, σ is the control point spacing, λ_i are the B-spline coefficient vectors (i.e., control point displacements), and N_x is the set of control points within the B-spline support of x . Cubic B-spline FFD has several advantageous properties as a deformation model. (Holden, 2008) First, it has low-dimensional parameterizations and local support, which reduces computational complexity and allows highly localized deformation to be modeled. Second, it provides C^2 continuity at the knots, allowing gradient-based optimization approaches to be used. Third, the FFD grids can be constructed hierarchically to allow the registration to be performed using morphological pyramids, reducing the susceptibility to local minima during optimization.

The estimator in (4) is a general estimator, but we can identify a number of specific forms. For example, if $\mathbf{W}(\lambda)$ in (4) is replaced with an identity matrix, the objective function will not consider any deformation when incorporating the prior image. We refer to this approach as Prior Image, Penalized-Likelihood Estimation (PIPLE). If $\mathbf{W}(\lambda)$ is replaced with a rigid transformation (3 translation and 3 rotation parameters), we refer to this approach as rigid PIRPLE. The general form in which $\mathbf{W}$ represents a FFD is termed dPIRPLE.

2.3. Strategy for solving dPIRPLE

The implicitly defined dPIRPLE approach of (4) requires a strategy for finding both the attenuation and the deformation parameters. One straightforward method might be to combine both types of parameters and solve by a general “off-the-shelf” gradient-based optimization approach. However, methods that are not tomography-specific or registration-specific, may be exceptionally time-consuming due to the large scale parameter space and further complicated by local minima due to the non-convexity in FFD registration. As such, we choose to use an alternating maximization approach that (a) maximizes the dPIRPLE objective with respect to attenuation parameters with fixed registration (referred to as “image update”) using a tomography-specific optimizer and (b) maximizes over registration parameters with fixed attenuation (referred to as “registration update”) using a FFD-specific algorithm. This alternating approach between image and registration updates can be expected to handle the large scale parameter space and local minima better than general optimization approaches.

The flow chart in Figure 2 illustrates this alternating approach. After initializing the parameter pair with $(\mu^{[0,0]}, \lambda^{[0,0]})$,² a number of registration updates (S_1) are applied to λ in a registration block, followed by a number of image reconstruction updates (T_1) applied to μ in an image block. (The superscripts denote the “outer loop” alternation number and the “inner loop” number of image or registration updates applied.) After a first alternation, S_2 registration updates are applied to λ in a new registration block to start another alternation. Starting the workflow with a registration block is preferred since patient motion can then be at least partially compensated before any use of the prior image.

In the image block, since the deformation parameters are fixed, the dPIRPLE objective function becomes dependent only on the attenuation parameters and therefore is equivalent to a standard PLE with a prior image penalty without registration. Since the prior image penalty shares the same structure as the image roughness penalty, it can be easily shown that the prior image penalty also satisfies the conditions in (Erdogan and Fessler, 1999). Therefore, the objective function in the image block can be optimized by SPS as summarized in Table 1. We use the optimum curvature as discussed in (Erdogan and Fessler, 1999) for the surrogate of the likelihood term in this study.

In the registration block, since the deformation parameters only appear in the prior image penalty term, for registration updates, the dPIRPLE objective function can be reduced to only this term. The objective function can be further transformed into a minimization of image differences after applying the deformation operation, which is essentially a standard image registration problem with a modified p -norm operator as similarity metric as shown below

$$\hat{\lambda} = \underset{\lambda}{\operatorname{argmax}} \Phi(\mu, \lambda) = \underset{\lambda}{\operatorname{argmax}} -\beta_p \|\Psi_p(\mu - \mathbf{W}(\lambda) \mu_p)\|_{p_p}^{p_p} = \underset{\lambda}{\operatorname{argmax}} \|\Psi_p(\mu - \mathbf{W}(\lambda) \mu_p)\|_{p_p}^{p_p} \quad (6)$$

For example, when $p_p=2$, the objective function is equivalent to the common Sum of Squared Differences (SSD) similarity metric for registration. Therefore, the objective function in the registration block can be solved using any number of existing methods for deformable registration with minor modifications. Moreover, morphological pyramids can be used as part of the update strategy to prevent local minima in the registration update.

Although keeping the objective functions strictly the same between the two alternating updates is preferred in general, we have found that differing p_p values in each scenario can provide better convergence behavior in the registration update while still encouraging sparse differences in the image update. Specifically, we use a higher value of p_p in the registration update than in image update. While a deterministic approach (limited-memory BFGS) was used previously as the optimization algorithm in registration update (Dang *et al.*, 2013), in this study we used the Adaptive Stochastic Gradient Descent (ASGD) approach (Klein *et al.*, 2009). This algorithm employs a strategy of randomly sampling a subset of image voxels and achieves a substantial reduction in computation time per iteration, while keeping favorable convergence properties. The gradient of the objective function used in ASGD approach can be written as

²Superscripts with square bracket $[i, j]$ indicate the i^{th} alternation and the j^{th} update in current i^{th} alternation.

$$[\nabla_{\lambda} \Phi(\mu, \lambda)]_i = \frac{\partial}{\partial \lambda_i} \Phi(\mu, \lambda) = -\beta_p \left[\Psi_p \dot{\mathbf{W}}_i(\lambda) \mu_p \right]^T \dot{f}_p(\Psi_p(\mu - \mathbf{W}(\lambda) \mu_p)) \quad (7)$$

where $\dot{f}_p$ denotes the derivative of the modified p -norm function in prior image penalty on each element of the operand. $\dot{\mathbf{W}}_i(\lambda)$ denotes the derivative of the deformation operator with respect to the i^{th} deformation parameter.

Table 1 presents the pseudocode for the alternating maximization approach. In the registration block, the 1st (outer) loop corresponds to the morphological pyramids and the 2nd (inner) loop corresponds to the ASGD updates. The details of the random sampling and step size in ASGD updates can be found in (Klein *et al.*, 2009). In the image block, $\dot{h}_i$ denotes the derivatives of the marginal log-likelihoods, $\dot{f}_R$ is the derivatives of the modified p -norm function in image roughness penalty, c_i is the optimum curvature of the marginal log-likelihoods, $\omega_f(t)$ is the curvature of the penalty function defined as $\omega_f(t) = \dot{f}(t)/t$, K is the number of neighboring voxels used to penalize the j^{th} voxel, and $[\cdot]_+$ is the non-negativity constraint. The image update equation (9) in table 1 is a modified form of (10) in (Stayman *et al.*, 2013) that includes a deformation operator with fixed λ on the prior image. The convergence properties and the alternating maximization schedules of the proposed algorithm are discussed in section 3.2.

2.4. Computational complexity and implementation

The computational complexity of dPIRPLE is divided between the image update and the registration update. In the image update, we may characterize the complexity by the number of projection operations needed (forward projections and backprojections), which are the dominant factors for the computation time. Every SPS update requires 3 projection operations if optimum curvature is used. As such, the execution of the dPIRPLE algorithm with respect to image reconstruction in table 1 requires $3 \times \sum_{z=1}^Z T_z$ back/projections, where Z denotes the number of alternations. In the registration update, computation time stems from four main sources: 1) computing the gradient of the reduced objective function; 2) evaluating the reduced objective function when computing the step size; 3) computing the Jacobian matrix (of the deformation operator $\mathbf{W}$ over deformation parameters λ) (Klein *et al.*, 2009) at the beginning of each level of the pyramid; and 4) warping and interpolating the moving image at the end of each registration block. The use of the ASGD approach substantially reduces the computation time from the first two sources by randomly sampling a subset of the image voxels to compute the gradient and computing the step size adaptively based on the gradient information instead of function evaluation. As such, the execution of the dPIRPLE algorithm with respect to registration in Table 1 is dominated by Z Jacobian calculations at each of the 4 levels of the pyramid and Z warping and interpolation operations.

We implemented dPIRPLE in Matlab (The Mathworks, Natick, MA), while computationally intensive functions were calculated using optimized C++ libraries. Specifically, the

projection operations in the image update used CUDA-based libraries, and the registration update used the image registration toolbox *Elastix*. (Klein *et al.*, 2010)

We compared the performance of dPIRPLE with a number of other approaches including filtered-backprojection (FBP), PLE, PIPE, and rigid PIRPLE. All the iterative approaches utilized matched separable footprint projection operations, which is a modified form (Wang *et al.*, 2014) of the separable footprint technique in (Long *et al.*, 2010), and FBP used a voxel-driven interpolating backprojection.

2.5 Cadaver experiments

Experiments were carried out on a cadaver torso on a CBCT test-bench. This study emulated a clinical scenario in which an initial diagnostic image has been acquired and a period of time has elapsed and a subsequent follow-up image is necessary. The initial image volume serves as a patient-specific prior image that is used to improve image quality and reduce the required radiation dose in subsequent follow-up scans. Specifically, we emulated a lung nodule surveillance scenario in which a suspicious nodule is imaged in a follow-up study to determine growth rates. In the cadaver experiments, a baseline scan was first acquired using a sufficiently large number of projections, to form a high-quality patient-specific prior image. Then, a follow-up scan was acquired with a substantial reduction in the radiation dose through a reduction in either the number of projections or the exposure per projection. The patient-specific prior image was then used in conjunction with the follow-up scan to reconstruct current anatomy and to detect a newly formed nodule in the presence of patient motion between scans.

The CBCT system consisted an x-ray source (DU694 in an EA10 housing; Dunlee, Aurora, IL), a flat-panel detector (PaxScan 4343CB with 1536×1536 pixels at 0.278 mm pixel pitch after 2×2 binning; Varian Medical Systems, Palo Alto, CA), and a motion control system (Parker Hannifin, OH), as shown in Figure 3(a). While the bench offers a wide range of source-detector motions, in this study, a system geometry was chosen that contains a 150 cm source-to-detector distance and 120 cm source-to-axis distance. All images were reconstructed with $260 \times 300 \times 330$ voxels and $1 \times 1 \times 1 \text{ mm}^3$ voxel size. The system geometry was calibrated using a separate scan following the calibration method described in (Cho *et al.*, 2005).

A baseline scan was first acquired with 360 projections over 360° (referred to as fully sampled dataset), 100 kVp and 1.25 mAs/projection (referred to as standard exposure). A PLE reconstruction ($p_R=2$, $\beta_R=10^6$) of this dataset was used as a patient-specific prior image, as shown in Figure 4. Approximately 1 cm^3 petroleum jelly ($\sim 0.013 \text{ mm}^{-1}$ attenuation) was then injected into the right lung of the cadaver by a thoracic surgeon, as shown in Figure 3(b). Many types of motion were imparted to the cadaver during this procedure, including the deformation of the abdominal soft tissues, flexing of the spine, and contraction of the chest wall. Two follow-up scans were acquired after the procedure: one with the same fully sampled protocol as used in the baseline scan and one with 0.1 mAs/projection (referred to as the low exposure dataset). A PLE reconstruction ($p_R=2$, $\beta_R=10^6$) of the fully sampled follow-up scan was used as the “ground truth” (and is referred to as the current anatomy).

Various sparse datasets were investigated through retrospectively selecting a reduced number of projections from fully sampled data.

To determine the robustness of the deformable registration method with regard to patient motion, a *simulation* study was conducted. In this study, a large number of instances of random rigid motion were applied to the high-quality prior image. These misregistered prior images were used as inputs to the dPIRPLE reconstruction and registration performance after a single registration was assessed. Performance analysis after a single registration leads to a somewhat conservative estimate of the capture range of patient motion since subsequent alternations have the potential to recover from a poor first registration. However, focusing on a single registration allowed for an assessment of a relatively large ensemble of misregistrations and recovery after a poor first registration was generally not observed.

The first registration is comprised of an initial rigid registration step followed by the deformable registration step described in section 2.2. The rigid registration uses Mutual Information (Mattes *et al.*, 2001) as the similarity metric and ASGD as the optimization. (The initial rigid registration was not used in the results above.) We included a performance analysis after each step of this first registration in our exploration of capture range. Random rigid misregistration was generated as perturbations from the nominal pose of the prior image. The nominal pose corresponded to the optimal rigid alignment of the prior image and the current anatomy, derived from an accurate rigid registration of the prior image and the fully sampled current anatomy reconstruction. Perturbations in each of the six parameters of rigid motion (i.e. translations in X, Y, and Z and rotations about X, Y, and Z) were described by Gaussian probability distributions with *three* standard derivations in each distribution as 100 mm in X (Anterior-Posterior), 100 mm in Y (Left-Right), 200 mm in Z (Superior-Interior), 10 degrees about X, 10 degrees about Y, and 10 degrees about Z. This perturbation model, adapted from (Otake *et al.*, 2012), was intended to emulate realistic variability in clinical setup of the patient in a diagnostic CT scanner. Each misregistration was quantified by computing the mean displacement averaged over all the voxels in the entire FOV of the CBCT.

2.6 Evaluation methods

To assess the accuracy and image quality of the reconstructed images, reconstructions were compared to the ground truth current anatomy using three metrics, including global root mean square error (RMSE), local RMSE, and structural similarity (SSIM) index (Wang *et al.*, 2004).

The RMSE between two images u and v is:

$$RMSE = \sqrt{\frac{1}{M} \sum_{i=1}^M (u_i - v_i)^2} \quad (10)$$

where M is the number of voxels in either image. The RMSE values have the same units as the image values (i.e. linear attenuation coefficients), that is, mm^{-1} . The global RMSE of a reconstructed image is computed over the entire field of view (FOV) of the CBCT, which

reflects the overall accuracy of the reconstructed image. The local RMSE was computed in a region of interest (ROI, $100 \times 100 \times 100$ voxels) that included the nodule and the adjacent soft tissue in the lung, reflecting reconstruction accuracy in close proximity to the nodule.

SSIM is a good complementary metric to RMSE, since it relates to the perceptual quality of an image, whereas the latter relates to quantitative accuracy. The SSIM is defined as a linear combination of luminance (l), contrast (c), and structure (s):

$$SSIM(u, v) = l(u, v) \cdot c(u, v) \cdot s(u, v) \quad (11)$$

where

$$l(u, v) = \left(2 \bar{u} \bar{v} + C_1 \right) / \left(\bar{u}^2 + \bar{v}^2 + C_1 \right) \quad (12a)$$

$$c(u, v) = (2\sigma_u \sigma_v + C_2) / (\sigma_u^2 + \sigma_v^2 + C_2) \quad (12b)$$

$$s(u, v) = (\sigma_{uv} + C_3) / (\sigma_u \sigma_v + C_3) \quad (12c)$$

and $\bar{u}$ and $\bar{v}$ denote the mean voxel values, σ_u and σ_v are standard deviations, and σ_{uv} is the sample covariance. The constants C_1 , C_2 , and C_3 were chosen as in (Wang *et al.*, 2004) to prevent instability in the computation of these three similarity measures. We computed SSIM in a ROI ($16 \times 16 \times 16$ voxels) about the nodule and the adjacent tissues.

To assess registration accuracy, the error vectors of the deformation field were generated from subtracting this field with the “true” field, which was approximated by registering the prior image with current anatomy as estimated by the fully sampled dataset. The deformation field vectors, the error vectors, and the error vector magnitudes were all visualized. Quantitatively, the average magnitude of the error vectors was computed over the object region (no air voxels) and compared to the average magnitude of the deformation field vectors. Additionally, both the difference images between the deformed prior image and the “true” current anatomy and the RMSE of the difference images were computed.

3. RESULTS

The experimental results are organized into three parts. First, parameter selection, convergence properties, and the alternating maximization schedule for dPIRPLE were investigated. All these results used the same sparsely sampled dataset (20 projections equally spaced over 190°) with a fixed exposure (1.25 mAs per projection), representing an 18-fold exposure reduction over a fully sampled dataset. Second, the reconstruction results of dPIRPLE using the same dataset were then compared to FBP and other iterative approaches (PLE, rigid PIRPLE, single-registration dPIRPLE). Third, the reconstruction results were evaluated at different measurement sparsity and exposure levels to investigate the robustness of dPIRPLE in preserving image quality and its limits in enabling exposure reduction.

3.1 Parameter Selection for dPIRPLE

Table 2 summarizes the major parameters used in dPIRPLE in this study. Variables are divided into two categories: 1) parameters that define the dPIRPLE objective function (i.e. first 7 parameters); and 2) parameters involved in solving the dPIRPLE objective function (i.e. last 3 parameters).

In these studies, $p_R=1$ was selected for its edge-preserving effect in the reconstructed image, and Ψ_P was chosen as the identity matrix since $\mu \cdot \mathbf{W}(\lambda) \mu_P$ is already very sparse for the lung nodule growth scenario. In the alternating optimization algorithm, a lower $p_P (=1)$ was used in the image update to encourage sparse differences, while a higher $p_P (=2)$ was used in the registration update to obtain better convergence. A four-level morphological pyramid was used in all the registrations with a control point spacing of 10 voxels at the final pyramid and a downsampling scheme of 8, 4, 2, and 1.

Previous work in (Stayman *et al.*, 2013) found that the optimal values of the three regularization parameters δ , β_P , and β_R in the PIRPLE objective function vary with differing datasets including differences in x-ray technique, system geometry, object, number of projections and volume size/sampling. Therefore parameter values need to be carefully selected for each dataset to achieve optimal image quality. That study showed that δ needs to be below a certain threshold relative to the expected attenuation differences in the reconstruction (and values smaller than this threshold have little effect on reconstruction quality). In the current study, we found $\delta=10^{-4} \text{ mm}^{-1}$ was small enough for all the datasets. For each dataset, optimal β_P and β_R were then selected by performing a pair of one-dimensional parameter sweeps as described in (Stayman *et al.*, 2013), that is, optimizing first over β_P while holding β_R fixed at an initially small value, and then optimizing over β_R at the optimal β_P , using RMSE as the metric. It has been observed in this study that as long as the initial β_R is chosen to be a *small* value and within the wide range of proper starting points (e.g. 10^0 to around $10^{3.5}$ in Figure 5), this pair of 1D sweeps would always reach close to the optimum. (More details on the double 1D sweep can be found in the sensitivity to parameter selection section in (Stayman *et al.*, 2013).) This pair of 1D sweeps (c.f., a full 2D sweep over β_P and β_R) is computationally convenient and takes advantage of the typical shape of $\text{RMSE}(\beta_P, \beta_R)$ as illustrated in Fig. 5 (dashed lines). The optimal (β_P, β_R) resulting from the double 1D sweep was nearly identical to the true 2D optimum (asterisk). This result demonstrates that the method proposed in (Stayman *et al.*, 2013) is also useful with respect to a deformable registration estimate (previously demonstrated with respect to rigid registration).

The last three parameters in Table 2 are important in determining the schedule of the proposed alternating maximization approach, and the choice of those parameters is investigated in section 3.3. A PLE image was found to give excellent initialization of the attenuation parameters in dPIRPLE. This choice provided a better starting point than a flat (zero) image, helping to speed convergence in the image update. Moreover, the PLE image also provided a feature rich dataset for estimating the deformation as compared to the prior image, which is critical in the first registration block. In contrast, initializing with a flat image can be detrimental to registration and is not particularly good for reconstruction. The

deformation parameters of dPIRPLE were initialized with all zeros in the first level of the morphological pyramid in the first registration block. For fair comparison, all prior-image-based approaches (PIPLE, rigid PIRPLE, and dPIRPLE) used the same initialization and δ , except that β_R and β_P were optimized separately for each approach and each dataset. PLE was initialized with a FBP image and used the same δ , while β_R was optimized for each dataset.

3.2 Convergence properties of dPIRPLE

Before evaluating the convergence properties of dPIRPLE with respect to registration, we first measured the accuracy of the final registration estimate. Figure 6 (a-b) shows the final deformation field in the axial and sagittal views, which clearly depict the movement of the body toward the posterior and the right side, as well as the anterior displacement of the heart. Note that each vector in the deformation field is plotted in the reverse direction from the physical deformation (i.e., pointing from a voxel in the current anatomy to the corresponding voxel in the prior image). Also note that the visualized deformation field vectors has a subsampled density (with a factor of 15) and a scaled vector length (with a factor of 2) from the original deformation field vectors to help better visual the field. Figure 6 (c-d) shows the error vectors generated from subtracting this final deformation field with the “true” field. Note that within the object the average magnitude of the error vectors is 0.77 mm, smaller than the 1 mm voxel size and much smaller than the average magnitude of the vectors in the estimated deformation field (6.56 mm). Error vectors with larger magnitudes outside the object are attributed to noise in air. To further visualize the deformation field error, the magnitudes of the error vectors were plotted in Figure 6 (e-f). Small to moderate errors (~0-2.5 mm) are evident within the object, such as the superior medial region in (e) and the inferior right region in (f), corresponding to errors in estimating the deformation of vasculature in the mediastinum and a slight mismatch in the displacement of the heart, respectively.

The convergence properties of dPIRPLE in registration were evaluated by examining the progression of the registration residual error at different alternations (i.e., outer loop iterations) as shown in Figure 7. The registration residual error was computed using the difference image between the deformed prior image and the current anatomy. It can be seen that the majority of mismatch was compensated after the 1st alternation of registration, which substantially prevented incorrect structures from being injected into subsequent image updates. The moving image used in this 1st registration was the initialization image in dPIRPLE, that is, a PLE image reconstructed using the sparsely sampled dataset (shown later in Figure 10, third row). However, remaining differences continued to be reduced between the 1st and the 20th alternation, especially in the areas of the ribs and primary bronchi, demonstrating the importance of the joint estimation. Note that the bright spot at $z=1$ and $z=20$ was not registration error but was created due to the introduced nodule which is inherent in the difference image between deformed prior and current anatomy. The local RMSE between the deformed prior image and current anatomy was $76.5 \times 10^{-4} \text{mm}^{-1}$ at $z=0$, $29.3 \times 10^{-4} \text{mm}^{-1}$ at $z=1$, $24.6 \times 10^{-4} \text{mm}^{-1}$ at $z=20$, and $22.6 \times 10^{-4} \text{mm}^{-1}$ using the “true” registration. These values reflect the same qualitative improvements as can be seen

throughout iterations providing additional indication of the advantage of the joint estimation approach.

We also quantified the convergence of dPIRPLE by plotting its objective function difference versus iteration in Figure 8(a). The objective function value at the solution, denoted as $\Phi^{[\infty]}$, was approximated by the value at the end of all 20 alternations ($Z=20$) including 20 registration blocks and 20 image blocks. Note that the cumulative number of image updates n labeled on the horizontal axis was a measurement of all the SPS updates spanning over alternations ($n=\sum_{z=1}^Z T_z$). In each image block a number of T_z ($=50$) SPS updates were performed, followed by a registration block with a number of S_z ($=1000$) registration updates in each level of the pyramid, followed by the next image block, etc. Note that the objective function values in Figure 8(a) were computed using (4) with $p_P=1$ (no change between image and registration blocks). Figure 8(a) shows that the objective function increased monotonically within every image block due to the monotonicity of the SPS approach. The increase in objective function values was greatest in the 1st registration block and exhibited smaller improvement in subsequent registration blocks, consistent with the progression of the residual registration error in Figure 7.

We also evaluated the convergence of dPIRPLE in comparison to other iterative approaches using local RMSE. As shown in Figure 8(b), PLE quickly reduced the RMSE but plateaued at a relatively high RMSE. PIPLE did not substantially reduce the RMSE due to the mismatched prior image. In contrast, both rigid PIRPLE and dPIRPLE saw steady reduction throughout the iterations, while dPIRPLE showed a higher reduction rate than rigid PIRPLE. Note that the RMSE was compared at the same iteration rather than at the same computation time. dPIRPLE and rigid PIRPLE with $T_z=50$ required 77.5% and 4.7% more computation time than PIPLE due to the registration step, while PLE required 15.5% less time than PIPLE due to the lack of a prior image penalty. The amount of time needed for the deformable registration in dPIRPLE is discussed in the following section.

3.3 Alternating maximization schedule

The alternating maximization schedule is determined by three parameters: the maximum number of alternations in the optimization (Z), the number of ASGD updates per pyramid level in the z^{th} alternation (S_z), and the number of SPS updates in the z^{th} alternation (T_z). $Z=20$ was found to be sufficient to produce a nearly converged dPIRPLE image with little or no change in the image after $Z=20$. While the proposed alternating algorithm allows S_z and T_z to vary during alternations, we only considered the case of constant S and T in this study. (Finding optimal values of S_z and T_z at each alternation to achieve fastest convergence is the subject of future work.) Within each level of the pyramid, since the registration updates are dominated by the Jacobian calculation, one can perform many registration updates after the Jacobian calculation without substantially increasing the computation time. In this work, $S=1000$ was selected since larger S leads to marginal changes in the deformation field. In contrast, the number T has more influence in determining the computation time of the dPIRPLE technique. Increasing T means more SPS updates in one image block and relatively less frequent registration. If registration computation time were not an issue, it is arguably preferable to do registration updates more frequently so that the deformation

estimation is more accurate at each image update. However, because multi-resolution deformable registration does have significant computational cost, there exists a trade-off between computation time and registration update frequency (and therefore reconstruction accuracy), primarily determined by T .

To investigate this trade-off, Figure 9 depicts the local RMSE after n cumulative image updates for various T values. Note that $T=1000$ means one registration followed by 1000 SPS updates, corresponding to the simple staged estimation scenario with one registration followed by reconstruction. The $T=1000$ case shows significantly higher local RMSE, again demonstrating the importance of the joint estimation. In the zoomed region, $T=50$, 25, and 5 all exhibit similarly low local RMSE at the 1000th iteration. However, since $T=50$ requires the least computation time among those three choices (as shown in table 3), $T=50$ was selected as the optimal T in these studies. With more relaxed time constraints, one might choose different optimal T values according to different clinical requirements. The computation times in Table 3 were measured on a high performance workstation equipped with two Intel Xeon E5-2600 processors and one Nvidia GeForce GTX 680 graphics card.

3.4 Reconstruction results

We compared the reconstruction results of dPIRPLE with other approaches as well as the “true” current anatomy, as shown in Figure 10. All of these results used the same sparsely sampled dataset (20 projections equally spaced over 190°) with a fixed exposure (1.25 mAs per projection), representing an 18-fold exposure reduction over a fully sampled dataset. All iterative approaches (PLE, rigid PIRPLE, single-registration dPIRPLE, and dPIRPLE) used 1000 iterations (i.e. 1000 SPS updates) to generate nearly converged images. FBP exhibited substantial artifacts from sparse sampling which made reliable nodule detection very difficult. PLE reduced the artifacts, but strong regularization resulted in relatively low spatial resolution which made nodule detection and volume changes difficult to assess. Rigid PIRPLE appeared to overcome ambiguous structures in areas of more rigid motion such as the ribs, spine, and primary bronchi, but mismatch in the more deformable areas resulted in errors in reconstruction of the nodule. Single-registration dPIRPLE corresponds to a scenario of one deformable registration followed by 1000 SPS updates (i.e. $T=1000$ case in section 3.3). Despite fairly accurate nodule reconstruction, incorrect anatomy can be seen in a few locations across the anatomy such as inside the dashed circle, mainly caused by registration residual errors after the 1st registration. In contrast, dPIRPLE presented a highly accurate estimate of the true anatomy and the nodule.

3.5 Performance with varying sparsity

To investigate the performance of dPIRPLE in comparison with other iterative approaches and varying levels of projection sparsity, subsets of data from a full set of 360 projections over 360° were selected to form sparse datasets of 200, 100, 40, 20, and 10 projections in a short scan orbit that covers 200°. Moreover, the standard exposure (1.25 mAs/projection) used in section 3.1-3.4 was replaced with a low exposure (0.1 mAs/projection) acquisition to further challenge dPIRPLE at every sparsity level. For each approach and each dataset, an optimal β_R (PLE) or pair of (β_P, β_R) (PIRPLE, rigid PIRPLE, and dPIRPLE) was chosen using

the aforementioned 1D parameter sweep to find the lowest global RMSE and SSIM. If the two metrics resulted in different optima, an average value between the two optima was used.

Figure 11 summarizes the reconstructed images for each approach over a range of sparse sampling. While all approaches exhibit accurate estimates at 200 projections (i.e., one projection per degree covering 180° plus fan angle), as the sparsity increases, PLE, PIPEL, and rigid PIRPLE all exhibit apparent reductions in image quality. In PLE, as the sparsity increases, both the spatial resolution and the contrast over the entire image drop quickly. Visualization of the nodule becomes less accurate and hard to identify when the number of projections drops to 20 or below. In PIPEL, with the addition of patient-specific prior information, both the contrast and spatial resolution are better maintained compared to PLE at each level of sparsity. However, due to a lack of registration, anatomical mismatches are readily apparent. As the sparsity increases, the nodule becomes more and more distorted due to the reduced information provided by the measurements, and the nodule largely disappears at 10 projections. The joint registration in rigid PIRPLE helps to reduce a number of mismatches associated and improves overall image quality compared to PIPEL at each level of sparsity. However, errors are still visible in areas of more deformable anatomy. In particular, artifacts and a loss of spatial resolution are evident at soft tissue and airway boundaries in the lung. Moreover, the accuracy in reconstruction of the nodule is degraded as the sparsity increases. The remaining mismatches due to deformation are further corrected in dPIRPLE, and a high degree of accuracy is maintained over a broad range in sparsity. At the lowest number of projections (10), dPIRPLE was still somewhat able to preserve the nodule shape and achieve much better nodule reconstruction compared to other approaches, although a reduction in the nodule contrast can be seen.

The accuracy of the reconstructed images was quantified using global RMSE and SSIM. Figure 12 (a) shows that RMSE exhibited a clear monotonic dependence on the sparsity for all iterative approaches. Moreover, the sensitivity to sampling was substantially lower for dPIRPLE than for other iterative approaches. Note that the use of a prior image without registration (PIPEL) resulted in even higher error than the case without a prior image (PLE), illustrating the importance of proper registration. The SSIM results in Figure 12 (b) show similar trends as in Figure 12 (a), and a notable difference in the sensitivity to sampling is observed for dPIRPLE in comparison to the other approaches. A possible reason is that the SSIM metric in Figure 12 (b) only considers the nodule area instead of the entire image and the residual registration errors in the other approaches have more influence on the image quality in the vicinity of the nodule. Figure 12 also shows better image estimation in dPIRPLE than PLE at 200 and 100 projections, indicating that incorporating a patient-specific prior image can increase imaging performance even at moderate or full sampling.

3.6 Capture range of the deformable registration in dPIRPLE

Previous sections demonstrated accurate registration from the deformable registration method used in dPIRPLE for the specific patient motion introduced in the cadaver experiment. A *simulation* study was also conducted to determine the robustness (i.e. capture range) of the deformable registration method with regard to patient motion. Figure 13 shows the RMSE between the registered prior image and the current anatomy as a function of the

mean displacement after both the initial rigid registration and after both rigid and deformable registration. The RMSE remained consistently small when the mean displacement was less than 125 mm and became relatively large when the mean displacement was over 125 mm. Even though the deformable registration is able to decrease the RMSE substantially across the entire range, those cases with $\text{RMSE} > 30 \text{ mm}^{-1}$ do not provide clinically useful image reconstructions. Thus, deformable registration does not compensate for a poor rigid registration and the delineation between success and failure cases lies at a mean displacement of 125 mm (indicated by the dashed line). While the 125 mm capture range may be appropriate for many clinical tasks, this range may be extended with a more robust initialization of the rigid registration (Otake *et al.*, 2012).

4. Discussion

This paper introduced a model-based framework, dPIRPLE, that jointly: 1) estimates 3D deformation between a high-quality patient-specific prior image and current anatomy based on subsequent data acquisitions, and 2) estimates a reconstructed image from noisy measurements using the deformed prior image. The proposed framework is solved using an alternating maximization strategy, and both the optimization schedule and parameter selection were investigated. Cadaver experiments were conducted emulating a lung nodule surveillance scenario, and the dPIRPLE algorithm demonstrated superior reconstruction accuracy and image quality compared to FBP, PLE reconstruction, PLE reconstruction with an unregistered prior image (PIPLE), and PLE reconstruction with rigid registration of a prior image (PIRPLE) over a wide range of sampling sparsity and exposure levels.

While dPIRPLE provides a general framework that allows for different kinds of deformable registration in the registration block, a specific registration method utilizing a cubic B-spline-based FFD model and ASGD algorithm was selected in this work. This particular registration method was able to capture a large portion of the deformation in the very first registration block and continued to be iteratively refined with additional alternations of the maximization algorithm. These alternations were found to be essential in minimizing residual errors and maximizing imaging performance. However, despite the success of the dPIRPLE approach using B-spline-based FFD, some residual errors were observed even after a large number of alternations. These residual mismatches could possibly be resolved using more sophisticated registration methods, and is the subject of future work. For example, the deformation field estimation might be further refined by applying the Demons algorithm and its variants (Vercauteren *et al.*, 2009; Nithiananthan *et al.*, 2009; Nithiananthan *et al.* 2012). The capture range of the deformable registration in dPIRPLE was characterized by adapting a large number of random rigid motions to the prior image. A more realistic way may be to perform cadaver or phantom studies with increasing levels of non-rigid motion and is the subject of future work.

Investigation of the schedule for dPIRPLE alternating maximization demonstrated the trade-off between the reconstruction image accuracy and the computation time. In this paper the trade-off was primarily determined by the parameter T , the number of image updates in one image block. However, this work also presumed a constant number of image updates (T) and registration updates (S) per block during alternations. The more generalized scenario where

T_z and S_z may vary at the z^{th} alternation may potentially provide greater opportunity to balance accuracy versus computation time. For example, since more dramatic changes in the image estimate typically happen in earlier image updates, one might start with relatively low T_i for the up-to-date deformation estimate and gradually increase T_i to reduce the computation time. Such studies are the subject of ongoing and future work.

The use of the modified p -norm in the prior image penalty term allowed additional flexibility to encourage or enforce differences from the prior image in the reconstruction. In this paper, investigations used a strategy where norms were mismatched between the image update (a modified L^1 norm) and the registration update (L^2 norm). This approach was found to be advantageous in improving convergence and avoiding local minima. However, despite the desirable convergence properties seen in practice with this method, this norm mismatch leads to solutions that do not strictly maximize the original objective (either the L^1 or L^2 form). Although beyond the scope of this paper, this issue may be handled by applying methods similar to graduated non-convexity optimization (Blake and Zisserman, 1987). For example, one could start the first registration block with $p_p=2$ and gradually reduce the value of p_p in subsequent registration blocks to $p_p=1$. In this fashion, the desirable convergence properties observed above could be combined with a stricter convergence in a single objective.

Proper selection of regularization parameters is important for all model-based reconstruction approaches including dPIRPLE. Incorrect selection of regularization parameters, especially prior image penalty strength β_p , will not only reduce reconstructed image quality but have the potential to bias the reconstruction towards the prior data (i.e. to show no change). The studies in this paper employed the same method as in (Stayman *et al.*, 2013) to select optimal regularization parameters and demonstrated that knowledgeable 1D parameter sweeps in β_p and β_R individually can be used to identify optimal dPIRPLE parameters accurately and much more efficiently than an explicit, computationally intense 2D parameter sweep. When a truth image is not provided, one could still perform the basic search method with visual assessment instead of using a RMSE metric. These approaches still require several reconstructions to perform the optimization sweep. One method that can be used to help reduce this computational burden leverages analytical approximations to the implicitly defined estimator along with predictive performance metrics to determine the optimal penalty strength (Dang *et al.*, 2014) and is the subject of ongoing work. This method enables efficient selection of proper regularization parameters especially prior image penalty strength β_p , ensuring minimal bias of the reconstructed image towards the prior data in regions where change is present between the current and prior anatomy. Moreover, although the penalty strength β_p and β_R are treated as scalars in this work, the approach in (Dang *et al.*, 2014) is sufficiently flexible to allow the use of a spatially varying map of β_p or β_R to design and customize a space-variant penalty strength. This is similar to other intentionally space-variant regularization approaches used in PLE to enforce spatially uniform resolution (Stayman and Fessler, 2000) and optimization of task-based detectability (Gang *et al.*, 2014). In the context of prior-image-based reconstruction, space-variant penalty design could enforce more uniform inclusion of prior image information or intentionally nonuniform designs based on where anatomical change is found.

One very important question associated with all prior-image-based reconstruction methods concerns the veracity of features found in the reconstruction. In these studies, false structures were observed in reconstructions where registration was absent or insufficiently accurate (e.g., doubling of anatomy in extreme circumstances). Moreover, despite these obvious flaws in the image, the apparent image quality (e.g., sharpness) of the image can remain high. Thus, even though such image defects were not observed with dPIRPLE, it would be valuable to have an assessment that identifies potential false structure. In (Stayman *et al.*, 2012) a methodology was proposed whereby the reconstruction is decomposed into portions individually attributable to the data and the prior image. Such an approach could potentially be applied in dPIRPLE to ensure that false structures are not present in reconstructions.

While this paper focused on the lung nodule surveillance application, the general framework is applicable to many other scenarios where patient-specific prior images, imaging sequences, or longitudinal studies are available, including image-guided radiation therapy, post-operative treatment assessment, and monitoring of patients in the intensive care unit. Many of these applications have more stringent timing constraints and will require refinements and more computationally efficient implementations of the dPIRPLE approach. For example, a fully GPU-based implementation of dPIRPLE is expected to accelerate parallelizable operations (the current implementation only uses the GPU for projection operations and requires frequent and inefficient data transfers between CPU and GPU). Employing GPU for registration in dPIRPLE alone could potentially reduce each registration update from current minutes to seconds. (Gu *et al.*, 2010) Furthermore, combining ordered subsets and momentum methods has recently been shown to substantially accelerate the x-ray CT image reconstruction (Kim *et al.*, 2013). Recent experiments (Wang *et al.*, 2014) have suggested that standard non-sparse reconstructions for interventional systems can be computed in a few minutes. With such improvements, we expect that the dPIRPLE approach will be valuable in a wide range of imaging scenarios and in different anatomical sites. Moreover, extensions to the approach whereby prior image data is available from alternate imaging modalities can further generalize the application of dPIRPLE. In image-guided surgery, for example, dPIRPLE could be adapted to using a preoperative diagnostic CT prior image to improve image quality and reduce dose in intraoperative CBCT. A similar scenario can be envisioned for image-guided radiotherapy. dPIRPLE may also be applied to dual-energy CT in which a high-fidelity high-energy image is used to improve low-fidelity low-energy image with certain mechanism to prevent from perturbing the accuracy of the attenuation coefficients. Even more varied approaches with cross-modality registration are potentially possible using MRI and nuclear imaging studies. Such extensions would bring a wider spectrum of interesting challenges to the dPIRPLE framework, such as other forms of reduced fidelity data (e.g., alternate sparse samplings) and the need for novel similarity metrics for inclusion of cross-modality prior images. To address the potential additional difficulties in inconsistent intensity values, possible solutions may include matching the intensity values of the prior image to those of the current anatomy such as in (Nithiananthan *et al.*, 2011), employing a similarity metric that does not require consistent intensity values such as a mutual-information-based metric (Maes *et al.*, 1997), or exploiting feature-based registration methods such as registering the surface point

sets of both images (Reaungamornrat *et al.*, 2013). Another potential difficulty lies in that motion between scans is likely to be much larger than the one introduced in the cadaver experiment for some applications, for example when a prior image is acquired months before. Section 3.6 has demonstrated that dPIRPLE has a fairly wide capture range thereby being able to compensate fairly large motion. For even larger motion, more advanced methods are needed that would tolerate large initial misregistration. For pulmonary imaging, this might include approaches that exploit airway bifurcation structures and lung surface meshes as in (Uneri *et al.*, 2013).

Acknowledgments

The authors would like to thank Yoshito Otake for valuable assistance with the CUDA-based reconstruction libraries, Zhe Zhao for the assistance in the cadaver experiments. We thank Sungwon Yoon, Richard Colbeth, and Edward Shapiro (Varian Medical Systems, Inc.) for valuable discussion and collaboration and students, staff, and faculty in the I-STAR Lab (www.jhu.edu/istar) at Johns Hopkins University for assistance and useful discussion. This work was supported in part by research grants from Varian Medical Systems, Inc. and the National Institute of Health (2R01-CA-112164).

References

- Blake, A.; Zisserman, A. Visual reconstruction. MIT Press; Cambridge, Mass: 1987.
- Candès EJ, Romberg J, Tao T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory.* 2006; 52(2):489–509.
- Chen GH, Tang J, Leng S. Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets. *Med. Phys.* 2008; 35(2):660–3. [PubMed: 18383687]
- Cho YB, Moseley DJ, Siewerdsen JH, Jaffray DA. An accurate technique for complete geometric calibration of cone-beam computed tomography systems. *Med. Phys.* 2005; 32(4):968–83. [PubMed: 15895580]
- Chun SY, Fessler JA. Joint image reconstruction and nonrigid motion estimation with a simple penalty that encourages local invertibility. *Proc. SPIE Medical Imaging.* 2009; 7258:72580U.
- 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(10):6484–98. [PubMed: 23039683]
- Dang H, Wang A, Zhao Z, Sussman M, Siewerdsen JH, Stayman JW. Joint estimation of deformation and penalized-likelihood CT reconstruction using previously acquired images. *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.* 2013:424–7.
- Dang H, Siewerdsen JH, Stayman JW. Regularization Design and Control of Change Admission in Prior-Image-based Reconstruction. *Proc. SPIE Medical Imaging.* 2014; 9033:90330O.
- Ding Y, Siewerdsen JH, Stayman JW. Incorporation of noise and prior images in penalized-likelihood reconstruction of sparse data. *Proc. SPIE Medical Imaging.* 2012; 8313:831324-1–7.
- Erdogan H, Fessler JA. Monotonic algorithms for transmission tomography. *IEEE Trans. Med. Imag.* 1999; 18(9):801–14.
- Evans JD, Polite DG, Whiting BR, O’Sullivan JA, Williamson JF. Noise-resolution tradeoffs in x-ray CT imaging: A comparison of penalized alternating minimization and filtered backprojection algorithms. *Med. Phys.* 2011; 38(3):1444–58. [PubMed: 21520856]
- Fessler JA. Optimization transfer approach to joint registration / reconstruction for motion-compensated image reconstruction. *Proc. IEEE Int. Symp. Biomed. Imaging.* 2010:596–9.
- Gang GJ, Stayman JW, Zbijewski W, Siewerdsen JH. Non-Stationary Noise and Detectability Index in CT Image Reconstruction by Filtered Backprojection and Penalized Likelihood Estimation. *Med. Phys.* 2014 under revision.

- Gu X, Pan H, Liang Y, Castillo R, Yang D, Choi D, Castillo E, Majumdar A, Guerrero T, Jiang S. Implementation and evaluation of various demons deformable image registration algorithms on a GPU. *Phys. Med. Biol.* 2010; 55(1):207–19. [PubMed: 20009197]
- Gilland DR, Mair BA, Bowsher JE, Jaszczak RJ. Simultaneous reconstruction and motion estimation for gated cardiac ECT. *IEEE Trans. Nucl. Sci.* 2002; 49(5):2344–9.
- 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. *Brit. J. Radiol.* 2000; 73(876):1252–9. [PubMed: 11205667]
- Heußner T, Brehm M, Ritschl L, Sawall S, Kachelrieß M. Prior-based artifact correction (PBAC) in computed tomography. *Med. Phys.* 2014; 41(2):021906. [PubMed: 24506627]
- Holden M. A review of geometric transformations for nonrigid body registration. *IEEE Trans. Medical Imaging.* 2008; 27(1):111–28.
- Huber, PJ. Robust statistics. Wiley; New York: 1981.
- Lange K. Convergence of EM image reconstruction algorithms with Gibbs smoothing. *IEEE Trans. Med. Imaging.* 1990; 9(4):439–46. [PubMed: 18222791]
- Lauzier PT, Chen GH. Characterization of statistical prior image constrained compressed sensing (PICCS): II. Application to dose reduction. *Med. Phys.* 2013; 40(2):021902. [PubMed: 23387750]
- Long Y, Fessler JA, Balter JM. 3D forward and back-projection for X-ray CT using separable footprints. *IEEE Trans. Med. Imag.* 2010; 29(11):1839–50.
- Kim D, Ramani S, Fessler JA. Accelerating X-ray CT ordered subsets image reconstruction with Nesterov's first-order methods. *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.* 2013:22–5.
- Klein S, Pluim JPW, Staring M, Viergever MA. Adaptive stochastic gradient descent optimization for image registration. *Int. J. Comput. Vis.* 2009; 81(3):227–39.
- Klein S, Staring M, Murphy K, Viergever MA, Pluim JP. Elastix: a toolbox for intensity-based medical image registration. *IEEE Trans. Med. Imag.* 2010; 29(1):196–205.
- Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P. Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imag.* 1997; 16(2):187–98.
- Navab N, Heining SM, Traub J. Camera augmented mobile C-arm (CAMC): Calibration, accuracy study, and clinical applications. *IEEE Trans. Med. Imag.* 2010; 29(7):1412–23.
- Nett B, Tang J, Aagaard-Kienitz B, Rowley H, Chen GH. 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. *Proc. SPIE Medical Imaging.* 2009; 7258:725803-1–12.
- Nithianathan S, Brock KK, Daly MJ, Chan H, Irish JC, Siewerdsen JH. Demons deformable registration for CBCT-guided procedures in the head and neck: convergence and accuracy. *Med. Phys.* 2009; 36(10):4755–64. [PubMed: 19928106]
- Nithianathan S, Schafer S, Uneri A, Mirota DJ, Stayman JW, Zbijewski W, Brock KK, Daly MJ, Chan H, Irish JC, Siewerdsen JH. Demons deformable registration of CT and cone-beam CT using an iterative intensity matching approach. *Med. Phys.* 2011; 38(4):1785–98. [PubMed: 21626913]
- Nithianathan S, Schafer S, Mirota DJ, Stayman JW, Zbijewski W, Reh DD, Gallia GL, Siewerdsen JH. Extra-dimensional Demons: A method for incorporating missing tissue in deformable image registration. *Med. Phys.* 2012; 39(9):5718–31. [PubMed: 22957637]
- Nuyts J, De Man B, Fessler JA, Zbijewski W, Beekman FJ. Modelling the physics in the iterative reconstruction for transmission computed tomography. *Phys. Med. Biol.* 2013; 58(12):R63. [PubMed: 23739261]
- Qiong X, Yu H, Mou X, Zhang L, Hsieh J, Wang G. Low-Dose X-Ray CT Reconstruction via Dictionary Learning. *IEEE Trans. Med. Imag.* 2012; 31(9):1682–97.
- 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(17):5485–508. [PubMed: 22864366]
- Reaungamornrat S, Liu WP, Wang AS, Otake Y, Nithianathan S, Uneri A, Schafer S, Tryggestad E, Richmon J, Sorger JM, Siewerdsen JH, Taylor RH. Deformable registration for cone-beam CT-

- guided transoral robotic base of tongue surgery. *Phys. Med. Biol.* 2013; 58(14):4951–79. [PubMed: 23807549]
- Ren L, Chetty IJ, Zhang J, Jin JY, Wu QJ, Yan H, Brizel DM, Lee WR, Movsas B, Yin FF. Development and clinical evaluation of a three-dimensional cone-beam computed tomography estimation method using a deformation field map. *Int. J. Radiat. Oncol. Biol. Phys.* 2012; 82:1584–93. [PubMed: 21477945]
- Rueckert D, Sonoda LI, Hayes C, Hill DLG, Leach MO, Hawkes DJ. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imag.* 1999; 18(8):712–21.
- Schafer S, Nithiananthan S, Mirota DJ, Uneri A, Stayman JW, Zbijewski W, Schmidgunst C, Kleinszig G, Khanna AJ, Siewerdsen JH. Mobile C-arm cone-beam CT for guidance of spine surgery: Image quality, radiation dose, and integration with interventional guidance. *Med. Phys.* 2011; 38(8):4563–74. [PubMed: 21928628]
- Sidky EY, Pan X. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Phys. Med. Biol.* 2008; 53(17):4777–807. [PubMed: 18701771]
- 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. *Proc. SPIE Visualization and Image-Guided Procedures.* 2009; 7261:72610J-1–72610J-8.
- Stayman JW, Fessler JA. Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction. *IEEE Trans. Med. Imag.* 2000; 19(6):601–15.
- Stayman JW, Zbijewski W, Otake Y, Schafer S, Lee J, Prince JL, Siewerdsen JH. Penalized-likelihood reconstruction for sparse data acquisitions with unregistered prior images and compressed sensing penalties. *Proc. SPIE Medical Imaging.* 2011; 7961:79611L-1–6.
- Stayman JW, Prince JL, Siewerdsen JH. Information propagation in prior-image-based reconstruction. *Proceedings of the 2nd International Conference on Image Formation in X-ray CT.* 2012:334–8.
- 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(21):7563–82. [PubMed: 24107545]
- Taguchi K, Sun Z, Segars WP, Fishman EK, Tsui BM. Image-domain motion compensated time resolved 4D cardiac CT. *Proc. SPIE Medical Imaging.* 2007; 6510:651016.
- Thibault JB, Sauer KD, Bouman CA, Hsieh J. A three-dimensional statistical approach to improved image quality for multislice helical CT. *Med. Phys.* 2007; 34(11):4526–44. [PubMed: 18072519]
- Uneri A, Nithiananthan S, Schafer S, Otake Y, Stayman JW, Kleinszig G, Sussman MS, Prince JL, Siewerdsen JH. Deformable registration of the inflated and deflated lung in cone-beam CT-guided thoracic surgery: Initial investigation of a combined model- and image-driven approach. *Med. Phys.* 2013; 40(1):017501. [PubMed: 23298134]
- Vercauteren T, Pennec X, Perchant A, Ayache N. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage.* 2009; 45(1):S61–S72. [PubMed: 19041946]
- Vogel CR, Oman ME. Iterative methods for total variation denoising. *SIAM J. Sci. Comput.* 1996; 17(1):227–38.
- 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(4):1005. [PubMed: 24504126]
- Wang AS, Stayman JW, Otake Y, Kleinszig G, Vogt S, Siewerdsen JH. Nesterov's Method for Accelerated Penalized-Likelihood Statistical Reconstruction for C-arm Cone-Beam CT. *Proc. Third Intl. Mtg. on image formation in X-ray computed tomography.* 2014
- Wang J, Gu X. High-quality four-dimensional cone-beam CT by deforming prior images. *Phys. Med. Biol.* 2013; 58(2):231–46. [PubMed: 23257113]
- Wang Z, Bovik AC, Sheikh HR, Simoncelli EP. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.* 2004; 13(4):600–12. [PubMed: 15376593]
- Werner-Wasik M, Xiao Y, Pequignot E, Curran WJ, Hauck W. Assessment of lung cancer response after nonoperative therapy: tumor diameter, bidimensional product, and volume. A serial CT scan-based study. *Int. J. Radiat. Oncol. Biol. Phys.* 2001; 51(1):56–61. [PubMed: 11516851]

- Xu J, Tsui BM. C-arm CT image reconstruction from sparse projections. Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med. 2013:34–7.
- Yang G, Hipwell JH, Hawkes DJ, Arridge SR. Numerical Methods for Coupled Reconstruction and Registration in Digital Breast Tomosynthesis. Annals of the BMVA. 2013; 2012(x):1–29.
- Yang L, Zhao J, Wang G. Few-view image reconstruction with dual dictionaries. Phys. Med. Biol. 2012; 57(1):173–89. [PubMed: 22155989]
- Zeng R, Fessler JA, Balter JM. Estimating 3-D respiratory motion from orbiting views by tomographic image registration. IEEE Trans. Med. Imag. 2007; 26(2):153–63.

Figure 1.

Two potential methods of dose reduction: sparse sampling (fewer projections acquired in the scan) and reduced fluence (lower dose per projection).

Figure 2.

Flow chart for solving the dPIRPLE objective function using an alternating maximization approach.

Figure 3.

(a) Experimental setup on the CBCT test-bench for cadaver experiments. (b) Simulation of lung tumor growth via petroleum jelly injection into the cadaver lung. A semiopaque rendering of a generic skeleton is overlaid on the photograph to illustrate the anatomical position of the cadaver.

Figure 4.

Three views of the patient-specific prior image, formed by PLE ($p_R=2$, $\beta_R=10^6$) reconstruction of the fully sampled dataset before the injection. Two zoomed-in regions in each view correspond to prior image (left) and current anatomy (right).

Figure 5.

Example optimization of (β_p, β_R) from a pair of 1D parameter sweeps (dashed lines) compared to the true 2D optimum (asterisk). Results here used the global RMSE as a metric; however, very similar results are seen when using local RMSE.

Figure 6.

(a, b) Deformation field estimated by dPIRPLE using 20 projections. The image below the deformation field is a merged image from the prior image and the current anatomy. The yellow and blue color correspond to positive and negative image value differences between the prior image and the current anatomy, while the original greyscale is used in the region where the image value differences are small. (c, d) Error vectors between deformation field estimated by dPIRPLE and the “true” deformation field approximated by registering the prior image with current anatomy. (e, f) Plot of the magnitude of error vectors (with a trace of the body outline superimposed).

Figure 7. Axial and sagittal views of a sequence of registration residual errors in dPIRPLE using 20 projections over 190°. The value z reflects the estimation at the z^{th} alternation of registration updates

Figure 8.

(a) The dPIRPLE objective function difference analyzed as a function of the cumulative number of image updates, $T_z=50$. (b) Local RMSE versus iteration number for all approaches.

Figure 9. Local RMSE computed versus the cumulative number of image updates for dPIRPLE at various T values.

Figure 10.

Reconstruction results of FBP, PLE ($\beta_R=10^3$), rigid PIRPLE ($\beta_P=10^4$, $\beta_R=10^{3.5}$), single-registration dPIRPLE ($\beta_P=10^4$, $\beta_R=10^{3.5}$), and dPIRPLE ($\beta_P=10^4$, $\beta_R=10^{3.5}$) with 20 projections and 1.25 mAs/projection. The local RMSE of the reconstruction result is displayed in each case. Note the accurate estimate of nodule and lung structures (e.g. airways) in dPIRPLE.

Figure 11.

PLE, PIPEL, rigid PIRPLE and dPIRPLE reconstruction images at different levels of projection sparsity at low exposure (0.1 mAs/projection).

Figure 12.

Global image accuracy [(a) RMSE] and local image quality [(b) SSIM] for PLE, PIPLE, rigid PIRPLE, and dPIRPLE reconstruction computed as a function of sparsity at low exposure (0.1 mAs/projection).

Figure 13.

RMSE between the registered prior image and the current anatomy as a function of the mean displacement after both the initial rigid registration and after both rigid and deformable registration. Each vertical pair of data points correspond to one instance of random rigid motion that was added to the prior image before registration. The dashed line indicates the delineation between success and failure cases at a mean displacement of 125 mm.

Table 1
Pseudocode for the dPIRPLE algorithm

Input $\mu^{[0,0]}, \lambda^{[0,0]}$

for $z = 1$ to max_alternations (Z)

% Registration Update Block

for each image pyramid

for $s = 1$ to number_of_ASGD_updates (S_z)

 Randomly sample a subset of image voxels

 Approximate gradient $\tilde{g}_s = \nabla_{\lambda} \Phi_{\mu} \left(\mu^{[z-1, T_{z-1}]}, \lambda^{[z, s-1]} \right)$ using sampled voxels and eq. (7)

 Compute an adaptive step size γ_s

$$\lambda^{[z, s]} = \lambda^{[z, s-1]} - \gamma_s \tilde{g}_s \quad (8)$$

end

end

% Image Update Block

for $t = 1$ to number_of_SPS_updates (T_z)

for $j = 1$ to number_of_voxels

$$\mu_j^{[z, t]} = \left[\mu_j^{[z, t-1]} + \frac{\left[\sum_{i=1}^N \mathbf{A}_{ij} \dot{h}_i \left(\left[\mathbf{A} \mu^{[z, t-1]} \right]_i \right) - \beta_R \sum_{k=1}^K [\Psi_R]_{kj} \dot{f}_R \left(\left[\Psi_R \mu^{[z, t-1]} \right]_k \right) - \beta_P \sum_{k=1}^K [\Psi_P]_{kj} \dot{f}_P \left(\left[\Psi_P \mu^{[z, t-1]} \right]_k \right) \right]}{\left[\sum_{i=1}^N [\mathbf{A}]_{ij}^2 c_i \left(\left[\mathbf{A} \mu^{[z, t-1]} \right]_i \right) + \beta_R \sum_{k=1}^K [\Psi_R]_{kj}^2 \omega_{f_R} \left(\left[\Psi_R \mu^{[z, t-1]} \right]_k \right) + \beta_P \sum_{k=1}^K [\Psi_P]_{kj}^2 \omega_{f_P} \left(\left[\Psi_P \mu^{[z, t-1]} \right]_k \right) \right]} \right]$$

end

end

end

return $\mu^{[Z, T_Z]}, \lambda^{[Z, S_Z]}$

Table 2
Summary of major parameters in dPIRPLE

Symbol	Name	Nominal values or range
Ψ_R	Sparsifying operator in image roughness penalty	First order difference operator
p_R	Modified p -norm in image roughness penalty	1
β_R	Penalty strength in image roughness penalty	$10^2 \sim 10^{3.5}$
Ψ_P	Sparsifying operator in prior image penalty	Identity matrix
p_P	Modified p -norm in prior image penalty	1 (Image update) 2 (Registration update)
β_P	Penalty strength in prior image penalty	$10^{2.5} \sim 10^4$
δ	Size of quadratic region in modified p -norm	10^{-4} mm^{-1}
σ	B-spline control point spacing	10 voxels
Z	Maximum number of alternations in optimization	20
$\{S_1, S_2, \dots, S_Z\}$	Number of ASGD updates per pyramid	$\{1000, 1000, \dots, 1000\}$
$\{T_1, T_2, \dots, T_Z\}$	Number of SPS updates per alternation	$\{50, 50, \dots, 50\}$

Table 3
Computation time at various settings of T

T	Registration Time (min)	Reconstruction Time (min)	Total Time (h)
5	1100	142	20.7
25	220	142	6
50	110	142	4.2
100	55	142	3.3
1000	5.5	142	2.5