Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2014 February 15; 9033: . doi:10.1117/12.2043781.

Regularization Design and Control of Change Admission in Prior-Image-based Reconstruction

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

Departments of Biomedical Engineering, Johns Hopkins University, Baltimore, MD USA 21205

Abstract

Nearly all reconstruction methods are controlled through various parameter selections. Traditionally, such parameters are used to specify a particular noise and resolution trade-off in the reconstructed image volumes. The introduction of reconstruction methods that incorporate prior image information has demonstrated dramatic improvements in dose utilization and image quality, but has complicated the selection of reconstruction parameters including those associated with balancing information used from prior images with that of the measurement data. While a noise-resolution tradeoff still exists, other potentially detrimental effects are possible with poor prior image parameter values including the possible introduction of false features and the failure to incorporate sufficient prior information to gain any improvements. Traditional parameter selection methods such as heuristics based on similar imaging scenarios are subject to error and suboptimal solutions while exhaustive searches can involve a large number of time-consuming iterative reconstructions. We propose a novel approach that prospectively determines optimal prior image regularization strength to accurately admit specific anatomical changes without performing full iterative reconstructions. This approach leverages analytical approximations to the implicitly defined prior image-based reconstruction solution and predictive metrics used to estimate imaging performance. The proposed method is investigated in phantom experiments and the shift-variance and data-dependence of optimal prior strength is explored. Optimal regularization based on the predictive approach is shown to agree well with traditional exhaustive reconstruction searches, while yielding substantial reductions in computation time. This suggests great potential of the proposed methodology in allowing for prospective patient-, data-, and change-specific customization of prior-image penalty strength to ensure accurate reconstruction of specific anatomical changes.

Keywords

Regularization; Prior-image reconstruction; Penalized-likelihood; CT reconstruction; Sparse data; Dose reduction; Shift variance; Noise-resolution tradeoff

I. INTRODUCTION

A variety of clinical scenarios involve sequential tomographic imaging studies performed on same patient including longitudinal surveillance of disease progression or therapy response

*web.stayman@jhu.edu; phone 1 410-955-1314; fax 410-955-1110.

(e.g. changes in tumor size) and providing up-to-date visualization of anatomical change throughout image-guided interventional procedures (e.g. tissue excisions). Traditional protocols require a series of complete tomographic acquisitions, raising concerns over accumulated radiation dose. Model-based reconstruction (MBR) with its ability to include sophisticated and accurate models of the acquisition physics and measurement noise along with regularization techniques, has demonstrated substantial improvement in tradeoff between radiation dose and image quality.¹ However, traditional regularization techniques in MBR tend to use only very general prior information such as edges and smoothness in the image,² neglecting a great amount of historical patient-specific anatomical information that could potentially be used in sequential imaging studies. In recent years, prior-image-based reconstruction (PIBR) has been developed that takes advantage of patient-specific prior images. For example, this idea of incorporating prior image was adopted in PICCS,³ which seeks sparse differences between the reconstruction of the current anatomy and the prior image using compressed sensing concepts and a linear equality constraint that enforces a match between the measurements and the reconstruction. Another PIBR approach that attempts to leverage both the advantages of statistical reconstruction and prior images is the PIRPLE approach.⁴ PIRPLE employs a MBR framework that integrates both a statistical objective function with noise model for data matching and a generalized prior-image regularization term. Inclusion of prior image in PIBR allows for dramatic reduction in the data fidelity requirements and has demonstrated good image quality under substantial undersampling and photon starvation, therefore permitting substantial dose reduction.^{4,5}

Despite these benefits, all PIBR methods require a key piece of information: To what extent should prior image information be used in the reconstruction process? This balance of information is typically controlled through a regularization parameter, and inappropriate selection of prior strength can lead to either introduction of false features or little image quality benefit. For example, too much reliance on the prior can force the reconstruction to appear too similar to the prior image, obscuring change; while too little prior image strength can result in relatively poor image quality (with coarse spatial resolution and/or increased noise). Further complicating the balance of prior information and the measurement data, is that the optimal strength will typically vary across different patients, acquisition geometries (including sparsity in the number of projection views and number of measurements), x-ray techniques, and the particular nature of the anatomical changes between the current anatomy and the prior image (e.g. size, contrast). Traditional methods for establishing prior image strength include heuristics and look-up tables based on reconstructions of similar anatomical data and exhaustive searches involving a large number of full reconstructions performed with different parameter values. While the former heuristic approach is subject to error and suboptimal solutions, the latter approach can be extremely time-consuming considering that each reconstruction requires iterative solution.

In this work, we propose a novel approach that efficiently determines an optimal prior image regularization strength without performing large number of full reconstructions. Specifically, the penalty strength is controlled to accurately admit specific anatomical changes from the prior image in the reconstruction. This control of *change admission* in PIBR leverages an analytical approximation to the solution of PIBR objective functions with

non-quadratic norms that was previously used⁶ to decompose portions of the reconstruction arising from prior image and from measurements. Given a prior image and a specific anatomical change for which accurate reconstructions are sought, we introduce a predictive metric (that utilizes the approximate analytical solution) and evaluate this metric to obtain estimates of optimal prior strength. The proposed method is investigated in a simple ellipse phantom and in a lung nodule surveillance scenario. Shift-variance and data-dependence of optimal penalty strength are explored, and the optimal regularization based on the predictive approach is shown to agree well with traditional exhaustive reconstruction and evaluation.

II. METHODS

PIBR with a Statistical Model for Measurements

Previous work by Stayman *et al.*⁴ has developed an objective function named PIRPLE that introduces an extra prior image regularization term into the standard penalized-likelihood estimation (PLE). The standard PLE uses a Poisson noise model and an image roughness penalty to enforce local smoothness and/or edge-preservation. The prior image term seeks to encourage similarity of current estimate μ with a previously acquired prior image μ_P . Without considering the simultaneous registration^{4,7}, the PIRPLE objective can be written as

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

where β_P is the prior image regularization strength that we want to estimate in this work, β_R is image roughness strength, Ψ_R and Ψ_P are sparsifying operators, y is measurements, and p_R and p_P are modified p -norm, which is quadratic within a small neighborhood $\pm\delta$ of zero and a shifted p -norm outside $\pm\delta$ such that the function and derivative match at $\pm\delta$.

Approximate Analytical Solution of PIBR

The estimator in (1) does not generally admit closed form solution and is typically solved iteratively using optimization approaches⁸. Nonlinearities of the likelihood function, the use of p -norms, and the lack of closed-form solution make direct analysis of this estimator difficult. Second-order Taylor approximation of the likelihood have previously been applied⁹ to simplify the data fit term. Specifically,

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

$$\mathbf{W} = \mathbf{D} \left\{ (y - r)^2 / y \right\}, \quad \hat{\mathbf{l}}_i = -\ln(y_i - r_i / b_i) \quad (3)$$

where $\mathbf{A}$ is the discretized projection operation. (Note that $\mathbf{A}^T$ represents the matched backprojection operation). $\mathbf{D}\{\cdot\}$ forms a diagonal matrix from a vector, r are (presumed known) scatter contributions, and b are pixel-dependent gain factors. While it is possible to write a closed-form solution to (2) when quadratic penalties ($p_P=2, p_R=2$) are used, this norm selection is undesirable because it tends to blend features in the measurements with those in prior image and blur imagery, as opposed to non-quadratic norms (i.e. lower p) that

tend to enforce sparse differences and enforce sharp edges. Consider the typical case of a modified 1-norm penalty, which uses a quadratic near the origin to enforce differentiability (also equivalent to Huber loss function). Assuming a suitable operating point τ near the solution is available, we may then approximate the modified 1-norm using a quadratic function written as follows and illustrated in Figure 1.

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

$$\kappa_i(\tau_i) = \begin{cases} 1/(2\delta) & |\tau_i| < \delta \\ |\tau_i - \delta/2|/\tau_i^2 & |\tau_i| \geq \delta \end{cases} \quad (5)$$

Applying the approximation of (4) and (5) in (2) yields the following closed-form approximate analytical solution of PIBR:

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

$$\mathbf{D}_R = \mathbf{D} \{\kappa(\Psi_R \tilde{\mu})\}, \quad \mathbf{D}_P = \mathbf{D} \{\kappa(\Psi_P (\tilde{\mu} - \mu_P))\} \quad (7)$$

where $\tilde{\mu}$ denotes an estimate of the reconstruction $\hat{\mu}_{PIRPLE}$. (While (6) is closed form, since the matrix $\mathbf{A}$ is very large and typically not computed explicitly, the inverse operation in (6) may be solved, in practice, using a linear Conjugate Gradient (CG) approach.) If $\hat{\mu}_{PIRPLE}$ is available, we can set $\tilde{\mu}$ to be $\hat{\mu}_{PIRPLE}$ and (4) is expected to be very accurate. We refer to the solution (6) using $\hat{\mu}_{PIRPLE}$ in the operating point as the “Ideal approximation”. However, not only is $\hat{\mu}$ relatively computationally expensive to compute; having $\hat{\mu}_{PIRPLE}$ supersedes the need for an approximate solution. The “Ideal approximation” is useful however, in investigating accuracy of the approximation. For prospective design purposes, a choice of operating point based on other estimates is more meaningful. For example, in this work, we will presume that we know the kind of change that is anticipated (or alternately, the kind of change we wish to admit accurately). That is, using the presumed change, denoted as μ_C , we can set $\tilde{\mu}$ to be $\mu_P + \mu_C$. We refer to the solution (6) using this operating point as “P+C approximation.” Since the approximation is based on a presumed change, no reconstructions are required in the evaluation of (6).

Metric for Prior-Image Regularization Design

In previous work¹⁰, we noted that parameter optimization could be performed piecewise - specifically, a 1D optimization over p_P with a very low p_R , followed by a 1D optimization over p_R would achieve a near optimal solution in terms of Root Mean Square Error (RMSE) relative to the true image. This suggests a similar set of 1D optimizations may be performed analytically. In this work, we focus on the 1D optimization over p_P . Using the analytic approximation above we can consider the predicted performance of the PIRPLE approach and optimize the prior image strength using the following objective function:

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

Equation (8) represents a predicted RMSE that can be written in terms of the prior image (μ_P) and anatomical change (μ_C). Moreover, we have allowed the optimization metric to include a binary weighting (S) and make the predicted RMSE a local evaluation (e.g., over the region containing the particular change). In this preliminary work, we choose to solve (8) by evaluating (8) at different β_P with regular spacing and identifying optimal β_P . More sophisticated routines, where one solves for an exact solution using standard optimization approaches (including root-finding based on the gradient of (8)) is also underway. Because computation of the closed-form solution is much faster than actual reconstruction, 1D optimization by this metric is much more efficient than by actual PIRPLE reconstructions.

III. RESULTS AND BREAKTHROUGH WORK

To evaluate the proposed methodology, we performed simulation studies on both a simple ellipse phantom and a chest phantom as shown in Figure 2. For the ellipse studies the prior image consisted of the background ellipse and the two disc regions of high and low attenuation, while change (e.g., present in subsequent scans) was represented with two small inserts with equal contrast at the center of each disc. In the chest phantom, an image without a lung nodule was used as prior image, and a nodule was introduced both digitally and physically. For Change I, a 1.2 cm-diameter disk was digitally added to simulate a nodule of uniform contrast, and the support region S of that nodule is known exactly. In Change II, the phantom was generated via an experiment where $\sim 1 \text{ cm}^3$ Petroleum jelly was physically injected to create a physical “nodule” of non-uniform contrast and irregular shape. In this case an inexact disc-shaped S was used centered at the tumor location with twice the in-plane area of the nodule. For both changes, the presumed changes μ_C were set to be the same as the actual changes. (For Change I, the actual change was known from the simulation; for Change II, the actual change was from a PLE image of a complete acquisition.) All studies used 20 projections over 190° with Poisson noise, representing $\times 18$ exposure reduction compared to a complete acquisition (360 projections over 360°). All the simulations were based on C-arm-type geometry with 120 cm source-to-detector, 60 cm source-to-axis, and 0.776 mm detector pixels for ellipse phantom studies, and 150 cm, 120 cm, 0.556 mm for chest phantom studies.

Figure 3(A) shows reconstructions using both exhaustive PIRPLE reconstruction (red) and the proposed analytical solution “P+C approximation” (green) for the two change positions. The degraded spatial resolution at low prior image strengths and the potential to obscure change at high prior image strengths is apparent in all reconstructions and approximations. In Figure 3(B), the proposed predictive RMSE metric is used to find optimal admission of these changes via evaluate the reconstruction accuracy at across β_P . Although the curves are not identical, the approximations are sufficiently similar to identify an optimal penalty strength. In both Location I and II, the optimal β_P evaluated on “P+C approximation” agrees well with the optimal β_P evaluated via PIRPLE reconstruction. The optimal β_P values in Figure 3(B) also agree qualitatively with Figure 3(A) image approximations (outlined in red/

green). Note that the optimal β_P is also shift-variant with lower β_P in the high-density region and higher β_P in the low-density region. Thus, one might use the proposed predictor to generate a shift-variant prior image strength map to optimally admit changes everywhere in the image.

Figure 4 illustrates PIRPLE reconstructions and image approximations using (4) for the chest phantom and the two cases of anatomical change. Again, the influence of too low a β_P (low resolution/noise) and too high a β_P (missing change) is evident. In this case, the ideal approximation, using $\hat{\mu}_{PIRPLE}$, is also presented. While these ideal approximations are very close to the true PIRPLE reconstruction, these approximations require multiple PIRPLE reconstructions whereas the P+C approximations are sufficiently accurate for optimization of penalty strength. In terms of computation time, one P+C approximation (12 seconds) is about $\times 60$ faster than one PIRPLE reconstruction (12 minutes) using GPU-accelerated forward projection and backprojection. Again there is a good match between qualitative and quantitative results. For Change I, an optimal β_P of $\sim 10^{3.5}$ was found with all quantitative measures. In the Change II case, since the nodule contour is not known exactly, we relax the definition of suspicious region S to a disk large enough to cover the expected nodule (twice the area of the actual nodule). Similar agreement was seen in this case, though the optimal β_P was $\sim 10^2$, in part due to the different nature of the anatomical change, but also due to the use of lower exposure data. These results suggest the proposed approach is capable of admitting additional errors from outside the actual change and robust against the particular probes used to optimize prior knowledge.

IV. CONCLUSION

We have introduced a novel approach that can efficiently determine optimal prior image regularization strength to admit specific anatomical changes in prior-image-based reconstruction (PIBR). This approach first approximates the PIBR result accurately by a closed-form analytical solution using approximations specifically designed for the non-quadratic penalties that are preferred in PIBR. Metrics utilizing this analytical solution and evaluating accuracy of local change rendering were developed to guide the selection of optimal prior strength and enforce admission of specific changes (e.g., those changes that are expected, desired, or whose accurate representation in the reconstruction are sought). Thus, the proposed methodology allows for prospective patient-, data-, and change-specific customization of how prior-image information is incorporated into the reconstruction.

ACKNOWLEDGEMENTS

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

- [1]. 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):4526–4544. [PubMed: 18072519]
- [2]. Lange K. Convergence of EM image reconstruction algorithms with Gibbs smoothing. *IEEE Trans. Med. Imag.* 1990; 9(4):439–446.

- [3]. 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):660–663. [PubMed: 18383687]
- [4]. 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. *Proc. SPIE Med. Imag.* 2011; 7961:79611L-1–79611L-6.
- [5]. Ding Y, Siewerdsen JH, Stayman JW. Incorporation of Noise and Prior Images in Penalized-Likelihood Reconstruction of Sparse Data. *Proc. SPIE Med. Imag.* 2012; 8313:831324-1–831324-7.
- [6]. Stayman, JW.; Prince, JL.; Siewerdsen, JH. Information propagation in prior-image-based reconstruction. *Proc. 2nd Intl. Conf. on Image Formation in X-ray CT*; 2012. p. 334-338.
- [7]. 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. *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.* 2013:424–427.
- [8]. Erdogan H, Fessler JA. Monotonic algorithms for transmission tomography. *IEEE Trans. Med. Imag.* 1999; 18(9):801–814.
- [9]. Sauer K, Bouman C. A Local Update Strategy for Iterative Reconstruction from Projections. *IEEE Trans. Signal Process.* 1993; 41(2):534–548.
- [10]. 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; 58:7563–7582. [PubMed: 24107545]

Figure 1. Approximation of the modified 1-norm with a quadratic function about an operating point τ_i .

Figure 2.

(Left) Ellipse phantom with attenuation (cm^{-1}): 0.02 (ellipse), 0.03 (left disk), 0.01 (right disk), 0.037 (left insert), 0.0017 (right insert). (Middle) Prior image of a cadaver chest. (Right column, top to bottom) Zoomed prior image, simulated nodule with uniform contrast (0.015 cm^{-1} attenuation) and known suspicious region-of-interest S , physically introduced nodule with non-uniform contrast and inexact known S .

Figure 3.

(A) Comparisons of PIRPLE reconstruction (red) and the predicted image using the P+C approximation (green) for the two image locations across regularization strengths. (B) Evaluation the proposed predictive performance RMSE metric for these reconstructions. PIRPLE parameters: $\Psi_P = \mathbf{I}$, $p_P = 1$, $p_R = 2$, $\beta_R = 10^5$, $\delta = 10^{-4}$ and 2000 iterations. P+C approximation: 100 CG iterations.

Figure 4.

(A-B) Chest Phantom - Change I: PIRPLE reconstruction, Ideal approximation, and P+C approximation at various β_p , and the RMSE curves. Note that too large β_p or too small β_p leads to false features (false negatives) or little benefits in PIRPLE reconstructions. (C-D) Change II reconstructions/approximations and RMSE metrics over β_p . Even with a mismatched region of interest, the optimal prior strength can be predicted successfully. PIRPLE parameters: $\psi_P=\mathbf{I}$, $p_P=1$, $\beta_R=0$, $\delta=10^{-4}$ and 2000 iterations. P+C and Ideal approximation: 50 CG iterations. Fluence in Change I case: 10^5 photons/detector; Change II: 10^4 photons/detector.