Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2018 March ; 10573: . doi:10.1117/12.2293135.

Prospective Image Quality Analysis and Control for Prior-Image-Based Reconstruction of Low-Dose CT

Hao Zhang¹, Grace J. Gang¹, Hao Dang¹, Marc S. Sussman², Cheng Ting Lin³, Jeffrey H. Siewerdsen¹, and J. Webster Stayman^{1,*}

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

²Department of Surgery, Johns Hopkins University, Baltimore, MD 21205, USA

³Department of Radiology and Radiologic Science, Johns Hopkins University, Baltimore, MD 21205, USA

Abstract

Purpose—Prior-image-based reconstruction (PIBR) is a powerful tool for low-dose CT, however, the nonlinear behavior of such approaches are generally difficult to predict and control. Similarly, traditional image quality metrics do not capture potential biases exhibited in PIBR images. In this work, we identify a new bias metric and construct an analytical framework for prospectively predicting and controlling the relationship between prior image regularization strength and this bias in a reliable and quantitative fashion.

Methods—Bias associated with prior image regularization in PIBR can be described as the fraction of actual contrast change (between the prior image and current anatomy) that appears in the reconstruction. Using local approximation of the nonlinear PIBR objective, we develop an analytical relationship between local regularization, fractional contrast reconstructed, and true contrast change. This analytic tool allows prediction bias properties in a reconstructed PIBR image and includes the dependencies on the data acquisition, patient anatomy and change, and reconstruction parameters. Predictions are leveraged to provide reliable and repeatable image properties for varying data fidelity in simulation and physical cadaver experiments.

Results—The proposed analytical approach permits accurate prediction of reconstructed contrast relative to a gold standard based on exhaustive search based on numerous iterative reconstructions. The framework is used to control regularization parameters to enforce consistent change reconstructions over varying fluence levels and varying numbers of projection angles – enabling bias properties that are less location- and acquisition-dependent.

Conclusions—While PIBR methods have demonstrated a substantial ability for dose reduction, image properties associated with those images have been difficult to express and quantify using traditional metrics. The novel framework presented in this work not only quantifies this bias in an intuitive fashion, but it gives a way to predict and control the bias. Reliable and predictable reconstruction methods are a requirement for clinical imaging systems and the proposed framework is an important step translating PIBR methods to clinical application.

I. INTRODUCTION

Model-based iterative reconstruction (MBIR) has enabled substantial improvements in radiation dose vs. image quality tradeoffs as compared to traditional filtered back-projection (FBP) method. These improvements have been driven by improved system models and the integration of additional prior knowledge. Prior-image-based reconstruction (PIBR) has been explored in research settings to use highly patient-specific prior knowledge for even greater improvements.^{1–8} For example, in sequential CT studies such as lung nodule surveillance, a high-quality patient-specific prior image can be incorporated into the reconstruction of subsequent data acquisitions to achieve order-of-magnitude exposure reduction.^{5,8} However, there are major challenges with the application of PIBR. For example, while integration of prior image knowledge can have a dramatic effect on the apparent image quality of the reconstruction, traditional image quality metrics like spatial resolution do not capture biases associated with the reconstruction. This complicates the selection of reconstruction parameters for PIBR prior image regularization strength which controls the balance between information from current measurements and information from the prior image. Too little prior information yields little to no improvement over “ordinary” MBIR, while too much prior information can obscure important features (by forcing PIBR to resemble the prior image too closely).^{9,10} While exhaustive parameter searches can be used to establish regularization strength, this process is time consuming (requiring many reconstructions) and may not generalize for different imaging scenarios – including different fluence levels, sampling, patient size, anatomy, etc.^{9,10} The ability to reliably understand and control the performance of PIBR in an intuitive and quantifiable fashion is a major hurdle to their clinical application.

In this work, we identify a particular metric for the form of bias found in PIBR (related to the contrast of a feature that differs between the prior image and the current anatomy) and develop an analytical framework for prospectively predicting and controlling PIBR performance. While we focus on the prior image registration, penalized likelihood estimation (PIRPLE) method,^{4,5} we expect the same strategy can be extended to other PIBR methods. We apply the analytic framework in both simulation and physical experiments to demonstrate reliable prediction and control of performance at varying fluence levels and levels of sampling.

II. METHODS AND MATERIALS

2.1 PIRPLE framework and closed-form approximation

The PIRPLE method has previously been proposed^{4,5} and may be written as:

$$\hat{\mu} = \arg \max L(\mu; y) - \beta_R \|\Psi\mu\|_1 - \beta_P \|\mu - T\mu_P\|_1 \quad (1)$$

The objective includes: 1) a data fidelity term based on marginal log-likelihoods, L ; 2) a regularization term that enforces smoothness (e.g., Ψ applies a spatial gradient); and 3) a prior image regularization term that encourages similarity but permits sparse differences between the reconstruction μ and a prior image μ_P . Prior image registration is applied

through transformation T ; and the regularization parameters β_R and β_P control the relative strengths of the roughness and the prior-image penalties, respectively. Previous analysis^{9,11} (ignoring registration) has found an approximate closed-form solution to the PIRPLE objective:

$$\begin{aligned}\hat{\mu} &\approx \arg \min \|\mathbf{A}\mu - l\|_{\mathbf{W}}^2 + \beta_R \|\Psi_R \mu\|_{\mathbf{D}_R}^2 + \beta_P \|\mu - \mu_P\|_{\mathbf{D}_P}^2 \quad \mathbf{D}_R = \mathbf{D}\{\kappa(\Psi_R \tilde{\mu})\}; \mathbf{D}_P = \mathbf{D}\{\kappa(\tilde{\mu} - \mu_P)\} \\ &= (\mathbf{A}^T \mathbf{W} \mathbf{A} + \beta_R \Psi_R^T \mathbf{D}_R \Psi_R + \beta_P \mathbf{D}_P)^{-1} (\mathbf{A}^T \mathbf{W} l + \beta_P \mathbf{D}_P \mu_P) \quad \kappa(t) = 1/|t|\end{aligned}\quad (2)$$

where $\mathbf{W} = \mathbf{D}\{y\}$ is a diagonal matrix of statistical weights, and l is an estimate of the line integrals. The diagonal matrices $\mathbf{D}_R$ and $\mathbf{D}_P$ are related to quadratic approximations to the ℓ_1 norm penalties about an operating point $\tilde{\mu}$.

2.2 Prospective selection of β_P in PIRPLE

The amount of prior image information to integrate in reconstruction is controlled through β_P . Too little prior information (small β_P) yields few improvements for PIBR, while too much prior information (large β_P) can lead to a PIBR result too similar to the prior image – obscuring or misrepresenting features. Due to the ℓ_1 penalties such problems often occur abruptly with changes in β_P . For example, see Fig. 1 for typical behavior for PIBR where the contrast of a lung nodule, found in the current CT data but not found in a prior image, is plotted as a function of β_P . Such behavior suggests that PIBR would be most easy to analyze in the transition region (i.e., not in the plateaus) centered at one-half the true contrast of the change. This suggests we can apply (2) using an operating point $\tilde{\mu}$ in the transition region. Consider a reconstructed change ($\hat{\mu}$) is reconstructed with fractional contrast γ such that $\hat{\mu} = \mu_P + \gamma \mu$ for $0 < \gamma < 1$, and there exists a unique β_P^* that would achieve that nodule contrast. Furthermore, if we temporarily ignore the roughness penalty and consider the simplified $\beta_R = 0$ scenario, we may use (2) to write:

$$\mu_P + \gamma \Delta \mu \approx (\mathbf{A}^T \mathbf{W} \mathbf{A} + \beta_P \mathbf{D}_P)^{-1} (\mathbf{A}^T \mathbf{W} l + \beta_P \mathbf{D}_P \mu_P) \quad (3)$$

Since $l \approx \mathbf{A}(\mu_P + \mu)$, solving for β_P : $\beta_P l \rightarrow (1 - \gamma) \mathbf{D}^{-1} \{\text{sign}(\mu)\} \mathbf{A}^T \mathbf{W} \mathbf{A} \mu$. Thus, for a nonnegative change $\mu(j)$ centered at location j , this suggests the following relationship between γ and a local regularization $\beta_{P,j}^*$:

$$\forall_j \beta_{P,j}^* = (1 - \gamma) [\mathbf{A}^T \mathbf{W} \mathbf{A} \Delta \mu(j)]_j \quad (4)$$

This relationship serves both as an analytic tool for predicting bias properties in PIRPLE images (fractional contrast reconstructed, γ) as well as a tool for prospectively designing the

penalty, $\beta_{P,j}^*$, to admit user-specified levels of contrast. Note that (4) encompasses the geometry ($\mathbf{A}$), data- ($\mathbf{W} = \mathbf{D}\{y\}$), and location-dependence (j) of the bias properties in PIRPLE. As in the Fessler and Rogers,¹² when applied to a compact change, $\mathbf{A}^T \mathbf{W} \mathbf{A} \approx \mathbf{\Lambda} \mathbf{A}^T \mathbf{A} \mathbf{\Lambda}$ where $\mathbf{\Lambda} = \mathbf{D}\{c\}$ is a diagonal matrix of aggregate certainties based on the data. Thus, Eq. (4) can be rewritten as:

$$\forall_j \beta_{P,j}^* = (1 - \gamma) c_j^2 [\mathbf{A}^T \mathbf{\Lambda} \mathbf{\Delta} \mu(j)]_j \quad (5)$$

suggesting a *shift-variant* regularization $\beta_{P,j}^* \sim 1/c_j^2$ for location-independent behavior (for shift-invariant $\mathbf{A}^T \mathbf{A}$). The validity of (4) and (5) is explored in two studies.

2.3 Simulation phantom study

A lung phantom in Fig. 2a was generated from an axial slice of a CBCT scan of a cadaver with two irregular inserts digitally added to mimic nodules. We emulated a lung nodule surveillance scenario in which two nodules shrank (Fig. 2b) in a subsequent scan. We explored the relations in (4) and (5), varying exposure in the follow-up scan with 90 projections over 360° with 10000, 6000 and 3000 incident photons per detector pixel with Poisson noise. We compare the expressions above with “gold standard” contrasts measured via exhaustive PIRPLE reconstruction (100 iterations, 10 subsets of OS-SPS) over a range of β_P values.

2.4 Physical cadaver study

A CBCT testbench was used to scan a cadaver torso (Fig. 3a). A baseline scan was first acquired with 100 kVp, 1.25 mAs/projection, and 360 projections over 360° (referred to as standard exposure). This reconstructed image volume of the cadaver served as the prior image for PIRPLE reconstruction subsequent low-dose scans. To emulate a lung nodule surveillance scenario in which a suspicious nodule is found in the follow-up study, ~1 cm³ petroleum jelly was injected into the lung of the cadaver to mimic a solid solitary pulmonary nodule. Using an x-ray technique of 100 kVp and 0.6 mAs/projection, the cadaver was scanned with 360 projections over 360°. Since there was substantial deformation between the prior and follow-up scans, deformable registration¹³ was applied as the $\mathbf{T}$ transform in (1). Finally, we extracted some projection angles from the full 360 projections to simulate different low exposure levels.

III. RESULTS

3.1 Simulation study results

Sample reconstructions of the simulated phantom are shown in Fig. 4 for three different shift-invariant regularizations of $\beta_P = 10^3$, $10^{4.4}$, and 10^5 . These reconstructions illustrate the shift-variant behavior of regularization since the reconstructed contrast is location-dependent – in particular the variable contrast in the left and right nodules for the $\beta_P = 10^{4.4}$ is evident. Similarly, the dangers of too high a regularization ($\beta_P = 10^5$) are illustrated when the reconstruction duplicates the prior image.

We used (4) to produce estimate reconstructed contrast curves over a range of values. Specifically, we varied γ in (4) from 0 to 0.99 with a 0.03 step size and computed the average $\beta_{P,j}^*$ inside the nodules. We compare these plots with a series of PIRPLE reconstructions $\beta_R=10^3$ and β_P varying from 10^1 to 10^5 with a $10^{0.1}$ step size and computing the average intensity of pixels inside the nodules. The resulting curves for three different fluence levels are illustrated in Fig. 5 (the points in these curves were intentionally downsampled for better visualization). We see that the predictor in (4) is matched well with the curves produced by exhaustive evaluation. Similarly, the shift-variant and fluence-dependent behavior of the bias is captured. Individual points corresponding to the Fig. 4 reconstructions (orange triangles) echo these observations.

Applying the shift-variant penalty suggested after (5) results in a new set of curves as shown in Fig. 6. The curves suggest increased shift-invariance and fluence-independence for the two nodules and regularization-bias properties. Again, predicted and exhaustive evaluations are in good agreement for all the cases.

3.2 Outcomes for the cadaver study

We extracted 30, 60, 90, 180 and 360 (evenly spaced) projection angles from the follow-up CBCT scan of the cadaver to explore the performance of the predictor in (4) under variable angular sampling conditions (an alternate method to achieve different low exposure levels). Again, we varied γ in (4) from 0 to 0.99 with a 0.03 step size and computed the average $\beta_{P,j}^*$ inside the nodule. PIRPLE reconstructions using $\beta_R=10^3$ and a shift-invariant β_P varying from $10^{2.5}$ to $10^{6.5}$ with a $10^{0.1}$ step size were performed and the average intensity of pixels inside the nodule was computed. The resulting curves for the predictor and the exhaustive search for the five different angular samplings are illustrated in Fig. 7. There is a good match between the predicted performance and the exhaustive search. The dependence on angular sampling is captured by the predictor.

We note that the certainty-based penalty following (5) would not lead to sampling-independent performance; however, the predictor may be used to tabulate a relationship between β_P and a specific fractional reconstruction contrast level γ . Such a tabulation is illustrated in Fig. 8 for the $\gamma=0.5$ scenario. As one might expect, a smaller regularization strength is required to achieve the same performance with fewer projections.

IV. DISCUSSION AND CONCLUSION

We have proposed and validated a mathematical framework for predicting the relationship between the prior image regularization parameter and PIBR reconstruction bias. The experimental results with simulation and physical cadaver datasets illustrate that the proposed method has high accuracy as compared with the exhaustive search method. This framework permits prospective design of the regularization strength permitting quantitative and reliable reconstruction of features of a specific contrast (relative to what is in the prior image). These tools will facilitate more robust and intuitive PIBR which is critical for clinical translation and use of the huge potential for dose reduction in these sophisticated approaches.

Acknowledgments

This research was supported in part by an academic-industry partnership with Elekta AB (Stockholm, Sweden), NIH grant R21 CA219608, and an AAPM research seed funding grant.

References

1. Chen G-H, 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]
2. Yu H, Zhao S, Hoffman EA, Wang G. Ultra-low dose lung CT perfusion regularized by a previous scan. *Acad Radiol.* 2009; 16(3):363–373. [PubMed: 19201366]
3. Lee H, Xing L, Davidi R, Li R, Qian J, Lee R. Improved compressed sensing-based cone-beam CT reconstruction using adaptive prior image constraints. *Phys Med Biol.* 2012; 57(8):2287–2307. [PubMed: 22460008]
4. 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. [PubMed: 24107545]
5. Dang H, Wang AS, Sussman MS, Siewerdsen JH, Stayman JW. dPIRPLE: a joint estimation framework for deformable registration and penalized-likelihood CT image reconstruction using prior images. *Phys Med Biol.* 2014; 59(17):4799–4826. [PubMed: 25097144]
6. Zhang H, Han H, Liang Z, Hu Y, Liu Y, Moore W, Ma J, Lu H. Extracting information from previous full-dose CT scan for knowledge-based Bayesian reconstruction of current low-dose CT images. *IEEE Trans Med Imaging.* 2016; 35(3):860–870. [PubMed: 26561284]
7. Pourmorteza, A., Dang, H., Siewerdsen, JH., Stayman, JW. *Phys Med Biol.* Vol. 61. IOP Publishing; 2016. Reconstruction of difference in sequential CT studies using penalized likelihood estimation; p. 1986-2002.
8. Zhang H, Ma J, Wang J, Moore W, Liang Z. Assessment of prior image induced nonlocal means regularization for low-dose CT reconstruction: change in anatomy. *Med Phys.* 2017; 44(9):e264–e278. [PubMed: 28901622]
9. Dang H, Siewerdsen JH, Stayman JW. Prospective regularization design in prior-image-based reconstruction. *Phys Med Biol.* 2015; 60(24):9515–9536. [PubMed: 26606653]
10. Zhang H, Dang H, Gang GJ, Stayman JW. Prospective Regularization Analysis and Design for Prior-Image-Based Reconstruction of X-ray CT. *Int Meet Fully Three-Dimensional Image Reconstr Radiol Nucl Med.* 2017; 14:417–423.
11. Stayman JW, Prince JL, Siewerdsen JH. Information Propagation in Prior-Image-Based Reconstruction. *Int Conf Image Form X-Ray Comput Tomogr.* 2012; 2012:334–338.
12. Fessler JA, Rogers WL. Spatial resolution properties of penalized-likelihood image reconstruction: space-invariant tomographs. *IEEE Trans Image Process.* 1996; 5(9):1346–1358. [PubMed: 18285223]
13. 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 Imaging.* 1999; 18(8):712–721. [PubMed: 10534053]

Figure 1.

The reconstructed change, $\hat{\mu}$, vs. prior image penalty strength, β_P . We have observed that the contrast is reliably reproduced up to a certain β_P (first plateau, green region), then the change abruptly disappears (orange region), not appearing with higher β_P (red region).

Figure 2.

Digital lung phantom: (a) prior image generated from an axial slice of a CBCT scan of a cadaver and two simulated lung nodules (indicated by yellow circles); (b) subsequent follow-up scan with nodule shrinkage.

Figure 3.

(a) Cadaver torso; (b) prior image of the cadaver; (c) current anatomy after petroleum jelly injection and deformation; (4) difference between (b) and (c), indicating deformation between two scans.

Figure 4.

Reconstructions of the simulated phantom data (6×10^3 photons/pixel) at three different regularization values. Shift-variant bias is evident in the $\beta_P = 10^{4.4}$ case. When regularization is too high ($\beta_P = 10^5$), the reconstruction looks exactly like the prior image (Fig. 2a).

Figure 5.

The relationship between regularization and reconstructed nodule contrast using exhaustive search and analytical prediction over a range of exposures (1×10^4 , 6×10^3 , 3×10^3 photons/pixel). Individual points on the 6×10^3 photons/pixel curves that are associated with the Figure 4 reconstructions are indicated with an orange triangle ($\blacktriangle$). Shift-variant bias is evident with different reconstructed contrasts for the two nodules at the same regularization level, as is the dependence on fluence level. However, there is good agreement between the exhaustive search and the predicted contrast.

Figure 6.

Illustration of the shift-variant penalty design in the simulated phantom. Note that this penalty results in both location- and fluence-independent performance.

Figure 7.

Plot of nodule average intensity versus prior image regularization strength for the cadaver at five angular samplings. The proposed analytical prediction approach shows very high accuracy as compared with the exhaustive search method in all cases.

Figure 8.

The relationship between the number of projection views and β_P^* values at $\gamma=1/2$. Our proposed analytical prediction approach shows small estimation errors as compared with the exhaustive search method for all five cases.