Reconstruction of Difference using Prior Images and a Penalized-Likelihood Framework

Amir Pourmorteza, Hao Dang, Jeffrey H. Siewerdsen, J. Webster Stayman
Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD

***Abstract*—Prior-image-based reconstruction (PBIR) algorithms provide great improvements in the trade-off between radiation exposure and image quality, especially in sequential imaging studies where high-fidelity prior images are available. Most PBIR methods incorporate the prior information through a penalty term in the objective function. In many sequential studies the primary goal is to estimate the *difference* between the prior image and the current anatomy. Here we introduce a method called reconstruction of difference (RoD) to reconstruct the difference image directly from the measurements. This method relies on prior image data, but unlike conventional PBIR, the prior information is integrated in the data consistency term. We investigated the performance of RoD under noisy and sparsely sampled projections in local and global acquisition scenarios using simulated data and test-bench experiments. RMS errors were compared to a standard penalized likelihood (PL) algorithm. The RoD outperformed standard PL in noisy and truncated data, and in simulation and in test-bench data reconstructions. The performance of the global and truncated RoD were comparable. This suggests computational speedups and dose reduction are possible through the use of the local acquisition and RoD.**

***Index Terms*—Prior image, model-based reconstruction, statistical reconstruction, computed tomography, cone-beam CT.**

I. Introduction

Prior-image-based reconstruction (PIBR) methods have found increasing use in a number of computed tomography (CT) applications[1]–[4]. The ability to include patient-specific anatomical knowledge (typically via prior images) in the image formation process has yielded dramatic improvements in the trade-off between image quality and radiation exposure. Prior image approaches tend to be particularly helpful in overcoming scenarios with low data fidelity including sparse/subsampled data and noisy projections from low exposure acquisitions.

There are many clinical scenarios where high-fidelity prior images are potentially available that may be used for reconstruction of subsequent data acquisitions. Examples of such sequential imaging studies include: 1) Longitudinal cases where anatomical change is monitored in disease progression or treatment (e.g. tumor growth/shrinkage in image-guided radiation therapy (IGRT)); 2) image-guided surgery (IGS) cases where preoperative diagnostic/planning are available but interoperative CT is used for up-to-date anatomical data for image-guidance; 3) imaging of periodic motion (e.g. the contraction of the heart or the motion of a joint); and 4) dynamic studies including perfusion.

A number of PBIR methods have been developed in recent years including prior image constrained compressed sensing (PICCS) [1], statistical PICCS [2], and prior image registration penalized-likelihood estimation [4]. These techniques vary somewhat in their data consistency term, and in the presence and type of noise model that is presumed for the underlying data. However, the methods are similar in that they seek to reconstruct the current anatomy and integrate prior image data through a constraint or penalty term. Specifically, these methods tend to use transformed differences between the current reconstruction estimate and the prior data and a metric (typically an ℓ_1 norm) that encourages sparsity in the transformed domain.

In many sequential imaging studies, the primary objective is to find the *difference* between the prior image and subsequent scans. For example, in monitoring tumor growth, visualizing and quantifying changes in tumor size is critical to treatment. In perfusion imaging, it is often differences before and after the administration of contrast that are important. Such difference images often greatly enhance the visualization of important features (e.g. 3D digital subtraction angiography).

Such a focus on differences between the current anatomy and a prior scan suggest an alternate reconstruction framework where it is the difference image that is modeled and reconstructed directly. Lee et al. [5] proposed a method that sought to reconstruct difference via penalized likelihood (PL) reconstruction based on the subtraction of prior and current projection data. In this work, we propose a modified data consistency term based on a likelihood function that models the relationship between the difference image and the observed data. Again, this approach relies on prior image data, but its integration is now in the data consistency term in contrast to conventional prior image methods.

There are a number of potential advantages to this reconstruction of difference (RoD) approach. Not only does RoD change the primary output of the reconstruction to be the difference image, but regularization and control of image properties is instead related to the change (difference) image as opposed to the current anatomy. Additionally, in many clinical cases including cardiac function, image-guided surgery (IGS), and image-guided radiation therapy (IGRT), change is limited to a relatively small volume of interest (VOI). This permits a dramatic reduction in the support size for reconstruction, facilitating computational speedups, memory reductions, and (if desired) truncated, limited FOV data acquisitions. In contrast, traditional model-based approaches require a region of support covering the entire anatomy in the axial plane.

In this paper we introduce the RoD framework and investigate its performance under noisy and sparsely sampled projections in local (VOI) and global reconstruction scenarios, and in simulated data and test-bench experiments.

II. Methods

A. Model-based Reconstruction of Difference

An overview of the RoD imaging chain is illustrated in Figure 1. We seek to use noisy measurements, y, and prior image data μ^p, to reconstruct a difference image, μ^Δ. This difference image, itself, may be of interest, or one can use the prior image and the difference image to compute the current anatomy. In general applications, a registration step may be required for the prior image data. This could potentially be performed as in [6] in a staged approach based on a coarse reconstruction of the current measurements, or in a more sophisticated joint estimation scheme as in [4]. For simplicity in this paper, we focus on a well pre-registered prior image scenario.

B. Forward Model and Objective Function

A standard mean measurement model can be described as

$$\bar{y}_i = I_{i0}\exp(-[\mathbf{A}\mu]_i), \quad (1)$$

where I_{i0} is the number of unattenuated photons, μ is vector of attenuation coefficients representing the current anatomy, $\mathbf{A}$ is the system matrix, and $[\mathbf{A}\mu]_i$ is the line integral associated with the i^{th} measurement. We presume that y_i is independent and Poisson distributed.

Recognizing that the current image is the sum of prior image and difference image, we can rewrite (1) as

$$\bar{y}_i = I_{i0} \cdot \exp(-[\mathbf{A}(\mu^p + \mu^\Delta)]_i), \quad (2)$$

$$= I_{i0}\exp(-[\mathbf{A}\mu^p]_i) \cdot \exp(-[\mathbf{A}\mu^\Delta]_i). \quad (3)$$

Since μ^p is known, we may combine scaling terms into a single gain parameter:

$$g_{i0} = I_{i0} \cdot \exp(-[\mathbf{A}\mu^p]_i), \quad (4)$$

such that

$$\bar{y}_i = g_{i0} \cdot \exp(-[\mathbf{A}\mu^\Delta]_i). \quad (5)$$

Equation (5) is convenient since it reduces the difference forward model to the same form as the traditional forward model in (1). This formulation will permit the use of standard reconstruction algorithms.

We choose the following penalized-likelihood objective for the reconstruction of the difference image:

$$\hat{\mu}^\Delta = \arg\min_{\mu \in \mathbb{R}} \{-L(y;\mu^\Delta) + \beta_R R(\mu^\Delta) + \beta_M M(\mu^\Delta)\}, \quad (6)$$

where $R(\mu)$ is a traditional roughness penalty that mitigates artifacts and noise by encouraging smooth solutions with a strength controlled by parameter β_R. In this paper, we focus on $R(\mu) = \|\Psi_R \mu\|_1$ with Ψ_R denoting a local difference for voxels over a 1st order neighborhood[1]. A second penalty term, $M(\mu) = \|\mu\|_1$, with strength given by β_M, ensures that the difference itself is sparse – i.e., local and relatively small.

At first glance this 2nd penalty appears redundant with the traditional roughness penalty, but both are important for control of noise and the amount of prior information. While the first term is fairly intuitive, controlling a standard noise/resolution tradeoff, the 2nd term is more complex. Increasing β_M increasingly penalizes the magnitude of the difference image and therefore the contribution of the prior information in image formation. A large value of β_M forces the difference image to zero, (no difference from the prior image, and the image is reconstructed solely from the prior data) and a small β_M allows for large differences and therefore less prior information is incorporated in the reconstruction.

Note that the difference image (μ^Δ) denotes the change in attenuation coefficients between scans; therefore it could be positive or negative. As a result, traditional nonnegativity constraints on the reconstruction are not applied.

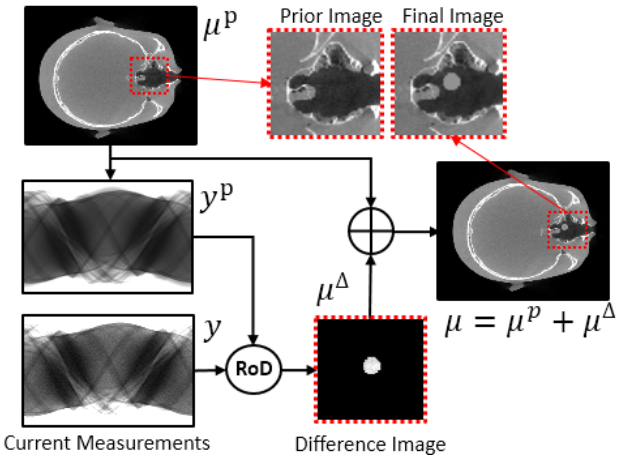

Figure 1- The flowchart of reconstruction of difference (RoD) method. Prior and final images are shown on the top right. The current measurements were created by placing a spherical object ("tumor") in the nasal cavity of the phantom. The dashed square indicates the region of interest used in local reconstruction and RMSE calculation.

1) Practical Implementation

To solve (6) an ordered subset separable paraboloidal surrogate (OS-SPS) [7] approach was chosen to minimize the PL cost function. This algorithm is highly parallelizable and was implemented in Matlab (The Mathworks, Natick MA), and we used CUDA-based libraries to perform the optimized projection operations on GPUs.

C. Experimental Design

Four experiments were conducted: 1) A regularization investigation using simulated data where various penalty strengths were investigated and optimized; 2) An investigation of the "local" VOI capabilities of RoD; 3) A performance study of RoD as compared with PL over a range of fluence levels and sampling; and 4) Sample cone-beam CT (CBCT) reconstructions using FBP, PL, and RoD. For simulation studies, we used a relatively high-dose CBCT scan (100 kVp, 453 mAs, and 720 projections over 360°) to produce a high fidelity prior image volume of an anthropomorphic head phantom. The ground truth current image ($\bar{\mu}$) was created by adding a small spherical "tumor" to the prior image as seen in Figure 1 (diameter: 21 voxels or 10.5 mm, and attenuation coefficient = 0.02 cm^{-1}). Simulated projection data were computed for 720 angles over 360°.

1) Regularization Investigation

To study general trends in optimal penalty strength, an exhaustive 2D sweep was performed and RoD images were estimated. Root-mean-square error (RMSE) between the RoD image and the original noiseless fully sampled simulated image

[1] Strictly speaking, we approximate the ℓ_1 norm using a Huber penalty with a small δ parameter which controls the location of the transition between the central quadratic neighborhood and the linear portions of the penalty. In this work $\delta = 10^{-4}$ for all cases.

($\bar{\mu}$) was used as a performance metric. RSME was calculated in a 100×100 voxel VOI around the tumor, as shown in Figure 1. We chose a large enough ROI to include the bone structures in the background in addition to the air and soft tissue of the nasal cavity.

2) *Local vs Global Performance*

We compared regularization and reconstruction performance of the global vs local scenarios to investigate the potential for RoD to use small support VOIs. Presuming that the anatomical change is local, we may also acquire the current image data by collimating the x-ray source to the VOI. We simulated this by selecting a 100×100 voxel region and truncating the projections not intersecting the VOI (Figure 1). The prior image was acquired without truncation. For comparison, RMSE was calculated over the same anatomical VOI in both local and global reconstructions.

3) *Performance as a Function of Data Fidelity*

To investigate RoD performance at varying fluence, different levels of Poisson noise were simulated for noisy acquisitions with I_0 ranging from 10^2 to 10^5 (with projection subsampled by 4). Similarly, to investigate the dependence on data sparsity, the number of projections was subsampled by a factor of 1, 2, 4, 8, 16, 30, and 45 at a typical fluence level ($I_0 = 10^4$). For these studies, optimal regularization parameters were determined by exhaustive searches and the local strategy was used.

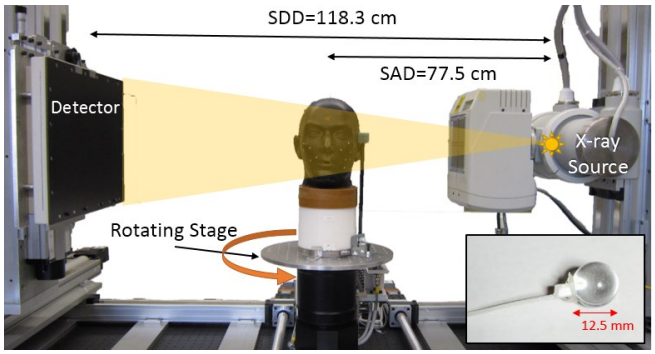

Figure 2-The CBCT test-bench with stationary flat panel detector (left) and x-ray tube (right), and the anthropomorphic head phantom on a rotating stage. The inset shows the acrylic sphere placed in the nasal cavity of the phantom to mimic a tumor.

4) *Test-Bench CBCT Reconstructions*

RoD was applied to physical CBCT data from a test-bench consisting of an x-ray source (Varian Rad-94), a flat-panel detector (Varian PaxScan 4030CB, 40×30 cm, at 0.388 mm pixel pitch after 2×2 binning), and a motion control system (Parker Hannifin), as shown in Figure 2. The source-to-detector (SDD) and source-to-axis distance (SAD) were set to 118 and 77.5 cm, respectively emulating a C-arm geometry. All images were reconstructed at 0.5 mm isotropic resolution.

We scanned an anthropomorphic head phantom to acquire high-fidelity projections (100 kVp, 453 mAs, 720 projections over 360°). We obtained scans with and without an acrylic sphere (12.5 mm size) placed in the nasal cavity mimicking a tumor (Figure 2). The phantom was not moved between scans. A PL reconstruction of the "no tumor" acquisition was used as the prior image.

Noisier follow-up scans (equivalent fluence=10^4 and 180 views) were formed by adding noise to a subset of the "with tumor" projections to form the current data (y). A PL reconstruction of the high-fidelity data (with tumor) acquisition was used as the reference for error calculation. RMSE was calculated over a 100×100 region around the acrylic sphere.

III. Results

A. Investigation of Regularization

RoD was used to reconstruct the simulated images (I_0=10^4, 180 views) using various penalty strength values swept linearly in the exponent with a $10^{1/2}$ step size, from 10^0 to 10^5. The penalty settings that resulted in the lowest RMSE were chosen as the optimal setting (Figure 3). For this specific experiment, the optimal values were $\beta_M = \beta_R = 10^{1.5}$. Basic trends are apparent: Noise increases for low β_R and becomes oversmoothed for high β_R (making the change eventually disappear). Similarly, high β_M eliminates all change in the reconstruction, while low β_M is subject to increased noise.

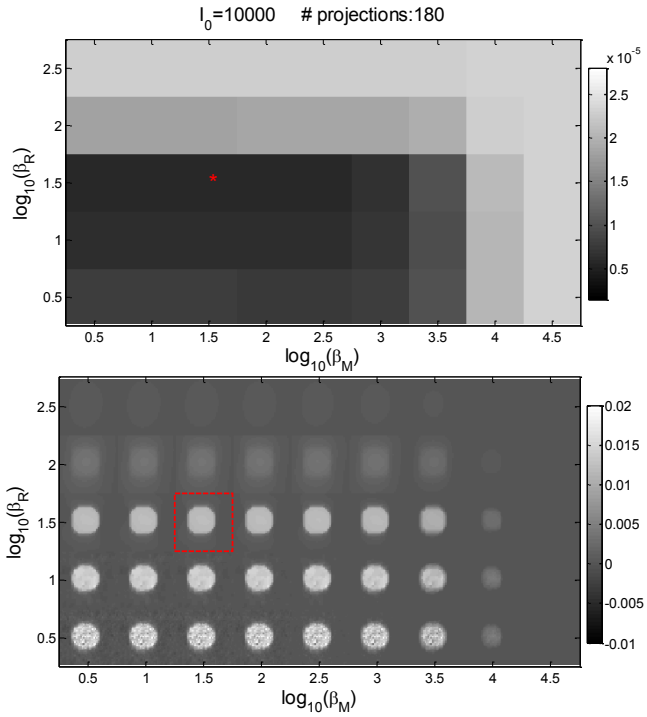

Figure 3-Results of regularization investigation. Top: RMSE for results of the exhaustive search for penalty coefficients. The asterisk denotes the point with optimal β values. Bottom: Zoomed VOI illustrating the difference reconstructions associated with each regularization strength pair.

B. Global vs. Local Reconstruction of Difference

For the simulated images mentioned above, we performed two separate exhaustive searches to find the optimal penalty coefficients in RoD: 1) local (collimated) acquisition and reconstruction, and 2) global acquisition and reconstruction. The optimal coefficients resulted in the following RMSE for global RoD (5.93×10^{-6} mm^{-1}) and local RoD (3.45×10^{-6} mm^{-1}).

We performed a similar analysis using the test-bench data as explained next in section D. Figure 4 shows the RMSE vs penalty coefficient maps for the local and global test-bench acquisitions. The optimal coefficients resulted in the following RMSE values for global RoD (1.17×10^{-5} mm^{-1}) and local RoD (1.24×10^{-5} mm^{-1}). In both simulation and test-bench results, the

penalty coefficient maps for the local and global RoD showed similar trends. The results indicate that the performance of the global and local reconstructions were comparable. For the next two sections the local reconstruction was used to speed up the search.

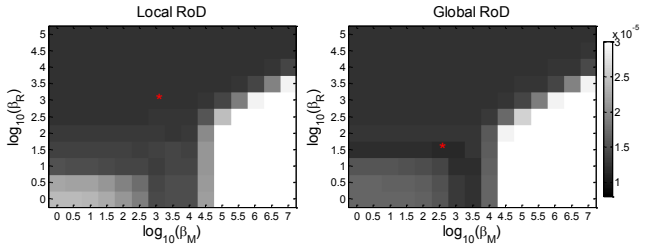

Figure 4- Exhaustive search for the optimal regularization coefficients in global vs local acquisition of test-bench data.

C. *RoD Performance with Varying Data Fidelity*

The performance of RoD with varying fluence and varying sparsity is summarized in Figure 5. RMSE for both RoD and PL is shown. In both cases, RoD outperforms PL reconstruction. As expected increased fluence and reduced sparsity decreases the RMSE in all cases; however, RoD RMSE appears to plateau in performance at lower levels of sparsity.

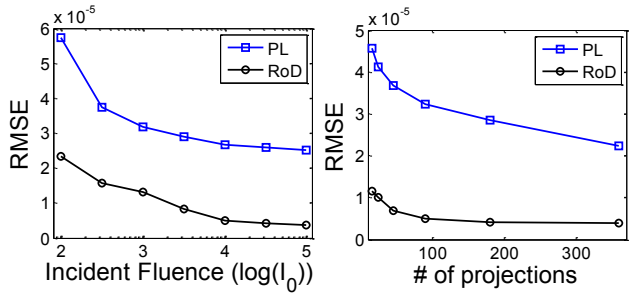

Figure 5- Effects of data fidelity. Left: RMSE (mm^{-1}) measured as a function of incident photon fluence (for 180 projections). Right: RMSE as a function of the number of projections (fluence level $I_0=10^4$). In each case, images were reconstructed with optimal penalty coefficients.

D. *Test bench data results*

CBCT test-bench data were reconstructed using the global and local RoD as well as PL. Using the coefficients determined in Figure 4, the images in Figure 6 show sample results of 3D reconstruction of the bench data. The RMSE was (1.2, 1.5, 1.8) $\times 10^{-6}$ mm^{-1} for RoD, PL, and FBP reconstructions respectively.

IV. Discussion

In this paper we have introduced a new reconstruction framework that focuses on integrating prior images into the data fidelity term of the reconstruction. The resulting reconstruction of difference approach estimates the change between the prior and subsequent scan – which may be of more interest than the current anatomy in some applications, though the context of surrounding anatomy may also be formed by adding the difference to the registered prior image data. The RoD reconstruction algorithm outperformed standard PL in noisy and truncated data, and in simulation and in test-bench data.

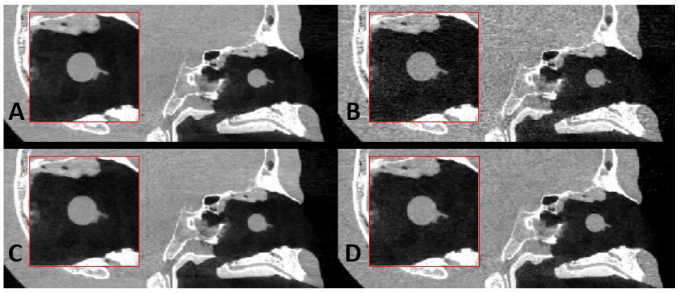

Figure 6- Coronal slices of sample results of 3D reconstruction of test-bench data for original (A: $I_0=2\times10^4$, FBP) and noisy (B-D: $I_0=5000$) acquisitions. B: FBP, C: Penalized Likelihood, D: Reconstruction of Difference (RoD)

In both real and simulated datasets, the performance of the local RoD was comparable to the global acquisition and reconstruction. Thus, there is the possibility for both computational speedups and dose reduction through the use of the local approach. The methodology proposed here may be directly applicable in some scenarios (e.g. cardiac imaging) where prior image data is already well registered to the current anatomy. However, more general application (e.g. in IGS and IGRT) will require proper rigid or deformable registration of the prior images which falls within the proposed framework in a manner analogous to previously reported methods for prior-image-registered reconstruction[3], [6]. These topics are the subject of ongoing and future work.

Acknowledgment

This work supported in part by an academic-industry partnership with Elekta Oncology (Stockholm, Sweden). The authors thank Dr. Junghoon Lee and Dr. John Wong (Department of Radiation Oncology, Johns Hopkins University) for collaboration and valuable conversations.

References

[1] G.-H. Chen, J. Tang, and S. Leng, "Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets," *Med. Phys.*, vol. 35, no. 2, pp. 660–663, 2008.

[2] P. T. Lauzier and G.-H. Chen, "Characterization of statistical prior image constrained compressed sensing (PICCS): II. Application to dose reduction," *Med. Phys.*, vol. 40, no. 2, p. 21902, 2013.

[3] H. Dang, A. S. Wang, M. S. Sussman, J. H. Siewerdsen, and J. W. Stayman, "dPIRPLE: a joint estimation framework for deformable registration and penalized-likelihood CT image reconstruction using prior images," *Phys. Med. Biol.*, vol. 59, no. 17, p. 4799, 2014.

[4] J. W. Stayman, H. Dang, Y. Ding, and J. H. Siewerdsen, "PIRPLE: A Penalized-Likelihood Framework for Incorporation of Prior Images in CT Reconstruction," *Phys. Med. Biol.*, 2013.

[5] J. Lee, J. W. Stayman, Y. Otake, S. Schafer, W. Zbijewski, A. J. Khanna, J. L. Prince, and J. H. Siewerdsen, "Volume-of-change cone-beam CT for image-guided surgery," *Phys. Med. Biol.*, vol. 57, no. 15, p. 4969, 2012.

[6] B. Nett, J. Tang, B. Aagaard-Kienitz, H. Rowley, and G.-H. Chen, "Low radiation dose C-arm cone-beam CT based on prior image constrained compressed sensing (PICCS): including compensation for image volume mismatch between multiple data acquisitions," in *SPIE Medical Imaging*, 2009, p. 725803.

[7] H. Erdogan and J. A. Fessler, "Ordered subsets algorithms for transmission tomography," *Phys. Med. Biol.*, vol. 44, no. 11, pp. 2835–2851, Nov. 1999.