

Non-homogeneous ICD Optimization for Targeted Reconstruction of Volumetric CT

Zhou Yu^a, Jean-Baptiste Thibault^b, Charles A. Bouman^a, Ken D. Sauer^c, and Jiang Hsieh^b

^aSchool of Electrical Engineering, Purdue University, West Lafayette, IN 47907-0501;

^bApplied Science Laboratory, GE Healthcare Technologies, W-1180, 3000 N Grandview Blvd, Waukesha, WI 53188;

^cDepartment of Electrical Engineering, 275 Fitzpatrick, University of Notre Dame, Notre Dame, IN 46556-5637

ABSTRACT

Medical imaging typically requires the reconstruction of a limited region of interest (ROI) to obtain a high resolution image of the anatomy of interest. Although targeted reconstruction is straightforward for analytical reconstruction methods, it is more complicated for statistical iterative techniques, which must reconstruct all objects in the field of view (FOV) to account for all sources of attenuation along the ray paths from x-ray source to detector. A brute force approach would require the reconstruction of the full field of view in high-resolution, but with prohibitive computational cost. In this paper, we propose a multi-resolution approach to accelerate targeted iterative reconstruction using the non-homogeneous ICD (NH-ICD) algorithm. NH-ICD aims at speeding up convergence of the coordinate descent algorithm by selecting preferentially those voxels most in need of updating. To further optimize ROI reconstruction, we use a multi-resolution approach which combines three separate improvements. First, we introduce the modified weighted NH-ICD algorithm, which weights the pixel selection criteria according to the position of the voxel relative to the ROI to speed up convergence within the ROI. Second, we propose a simple correction to the error sinogram to correct for inconsistencies between resolutions when the forward model is not scale invariant. Finally, we leverage the flexibility of the ICD algorithm to add selected edge pixels outside the ROI to the ROI reconstruction in order to minimize transition artifacts at the ROI boundary. Experiments on clinical data illustrate how each component of the method improves convergence speed and image quality.

Keywords: Computed tomography, iterative reconstruction, coordinate descent, targeted reconstruction, multi-resolution

1. INTRODUCTION

Statistical reconstruction algorithms for computed tomography have been developed in recent years and shown to improve image quality in comparison to conventional reconstruction algorithms.¹ In typical clinical applications, targeted reconstructions are used to provide a high-resolution image of a region of interest (ROI). Although targeting is relatively easy for direct reconstruction methods, it is less straightforward for iterative reconstruction (IR): IR generally requires optimizing a cost function formed over the full line integrals forming the scan, and therefore must reconstruct the full object within the field of view. However, computational cost for the iterative reconstruction of all the voxels within the full field of view in high-resolution would be prohibitive.

Multi-resolution approaches to iterative reconstruction for medical imaging applications have been widely investigated. Previous work^{2,3} have shown that iterating over a coarse grid followed by fine grid reconstruction can effectively reduce the computation time. More recently, recursive multi-grid techniques⁴ have also been

Further author information: (Send correspondence to J.B.T.)

Z.Y.: E-mail: yuz@purdue.edu, Telephone: 1 765 496 1263

J.B.T.: E-mail: jean-baptiste.thibault@med.ge.com, Telephone: 1 262 312 7404

C.A.B.: E-mail: bouman@purdue.edu, Telephone: 1 765 494 0340

K.D.S.: E-mail: sauer@nd.edu, Telephone: 1 574 631 6999

J.H.: E-mail: jiang.hsieh@med.ge.com, Telephone: 1 262 312 7635

proposed. In general, multi-resolution methods are well suited to targeted reconstruction since resolution changes with the field of view.⁵

In this paper, we propose a multi-resolution framework for the Non-Homogeneous Iterative Coordinate Descent (NH-ICD) algorithm⁶ for ROI reconstruction. First, we use a weighted NH-ICD (WNH-ICD) algorithm, which modifies the pixel selection criterion according to the voxel location relative to the ROI, in order to focus primarily on updating the voxels with strong correlation with the ROI, and achieve faster ROI convergence. Second, we avoid a costly full forward projection in between resolutions with a simple sinogram correction method when the forward model is not scale-invariant. Finally, we include some edge voxels outside ROI in addition to the ROI voxels in the high-resolution reconstruction in order to reduce artifacts near the ROI boundary.

The paper is organized as follows. Section 2 gives a brief review of the NH-ICD algorithm. Section 3 describes the proposed methods for multi-resolution targeted iterative reconstruction. In section 4, we illustrate with experiments on clinical data how convergence speed and image quality are improved.

2. NON-HOMOGENEOUS ICD ALGORITHM FOR VOLUMETRIC CT RECONSTRUCTION

Let y be the vector of tomographic measurements and x be the vector of image volume elements. In the Bayesian approach for statistical reconstruction, we model the data acquisition by the conditional probability density function $p(y|x)$, and the image by the prior density $p(x)$. The Maximum *a posteriori* estimate is given by maximizing the *a posteriori* density function $p(x|y)$ which leads to the following minimization problem:

$$\hat{x} = \arg \min_{x \geq 0} \left\{ \frac{1}{2} (y - Ax)^T D (y - Ax) + U(x) \right\}, \quad (1)$$

where A is the forward system matrix, and D is a diagonal weighting matrix which reflects variations in the credibility of data.⁷ The term $U(x)$ penalizes large variations in the image domain, while preserving edge characteristics. We use a Markov random field (MRF) model, where $U(x)$ takes the form of the logarithm of an *a priori* density, as the sum of potential functions $\rho(x_j - x_k)$, where x_j and x_k are neighboring voxels.^{1,8}

The iterative coordinate descent (ICD) algorithm⁹ has been successfully applied to solving this optimization problem and demonstrates relatively fast convergence when initialized with the FBP reconstruction. It decompose the multi-dimensional optimization problem into a sequence of one-dimensional steps, each resulting in a single voxel update. A full conventional ICD iteration updates each voxel in the image volume once. For each selected voxel, the reduced one-dimensional cost function derived from the error sinogram $e = Ax - y$ is minimized to compute a new voxel value, and the error sinogram is updated to reflect the changes in the image x .

In volumetric helical CT, define the coordinate system as follows. Let z be the axis about which the gantry rotates, and let the x - y plane be the plane that is orthogonal to z with y pointing upward. We use the notation $x_{i,j,k}$ to denote a voxel, where the indices i , j , and k corresponds to the voxel's x , y , and z coordinates. In our implementation, a single ICD update consists of first selecting a pixel location corresponding to a (i, j) pair, and then updating all the voxels along z corresponding to that (i, j) location. We define the pixel line (i, j) as the set of all voxels with the same (x, y) location (i, j) .

The non-homogeneous ICD algorithm⁶ is a variation of the ICD algorithm, which provides faster convergence speed by updating pixels non-uniformly across the reconstructed image. For example, the NH-ICD algorithm may repeatedly update a pixel line before all other pixel lines have been visited. To do this, we design a pixel selection criteria (PSC) which preferentially updates pixel lines that are likely to be far from converged, and we introduce a pixel update scheduler (PUS) which provides a method to order pixel lines for update based on the PSC. It updates the pixel lines with large PSC more often, while still ensuring that every pixel line is regularly visited. The NH-ICD algorithm alternates between homogeneous and non-homogeneous steps. Homogeneous steps are similar to a conventional ICD iteration, which visits all the pixel lines once in a random order, whereas non-homogeneous steps comprise several sub-iterations. In each non-homogeneous sub-iteration, the top 5% of pixel lines with the largest PSC are selected and then updated in sequence.




Figure 1. The multi-resolution framework of our targeted reconstruction algorithm.

3. NON-HOMOGENEOUS ICD OPTIMIZATION FOR TARGETED ITERATIVE RECONSTRUCTION

The multi-resolution framework of our targeted reconstruction algorithm is shown in Figure 1, where $x^{(0)}$ is the initial FBP image, $x^{(l)}$ is the low-resolution reconstructed image, $x^{(h)}$ denotes the high-resolution image interpolated from $x^{(l)}$, and $x_u^{(h)}$ is the final reconstructed image of the ROI in high-resolution; $e^{(l)}$ and $e^{(h)}$ are the error sinograms corresponding to $x^{(l)}$ and $x^{(h)}$ separately. In our approach, we first perform a low-resolution reconstruction that updates all the voxels in the full field of view encompassing all the objects measured by the CT system. We use the FBP images as initial conditions in the low-resolution reconstruction because they can be generated quickly and provide reasonable first estimates of the scanned object. Next, we perform a high-resolution reconstruction using the low-resolution images as initial conditions. The high-resolution reconstruction updates pixels in a set Ω formed by all the pixels within the ROI, to which we add selected edge pixels outside the ROI. Both the low-resolution and high-resolution reconstructions use the WNH-ICD algorithm to improve ROI convergence speed. In the transition between low and high resolutions, we perform an edge detection on the low-resolution image, and threshold it to form the set Ω . Low resolution images are interpolated to the desired resolution using a 4 point bi-cubic interpolation.¹⁰ Finally, an error sinogram correction reduces the mismatch due to the change in resolution when the forward model is not scale-invariant.

3.1. Weighted NH-ICD Algorithm

The motivation behind the WNH-ICD algorithm is to focus computation primarily on updating the voxels with strong correlation with the ROI (which includes the ROI itself). Updating preferentially the ROI and its surrounding areas leads to fast optimization and is sufficient to reach convergence quickly within the ROI while the remaining portions of the image may not be fully converged, to the extent that those errors do not significantly affect the ROI. The weighted NH-ICD algorithm assigns a weight $w(i, j)$ to each pixel line (i, j) as a function of each pixel's correlation to the ROI. This weighting function is applied to the NH-ICD algorithm, so that in the homogeneous steps, the frequency that a pixel location is updated is proportional to the weighting function; and in the non-homogeneous steps, the pixel selection criterion is proportional to the product of the pixel's update value and the weight.

Our study of the sensitivity of a single voxel's ICD update due to changes in other voxels has led to the development of an empirical method to compute the weighting function. Let u be an image with all the ROI voxels set to 1 and the rest voxels set to 0. First, u is forward projected into the sinogram domain. This sinogram is then back projected into the image domain. Let v be the back-projected image, that is, $v = A^T A u$. Finally, the sum of the back projected image volume along the z axis is taken to form the 2D weighting function. However, this empirical method is computationally costly, as each step of forward and back projection requires a full pass through the projection data and the image volume. Instead, we use the following analytical function to fit the

empirical weighting function.

$$\tilde{w}(r) = \begin{cases} \ln \frac{r+R+\beta R}{r-R+\beta R} & r > R \\ \ln \frac{(r+R+\beta R)(R-r+\beta R)}{\beta^2 R^2} & r \leq R \end{cases}, \quad (2)$$

where we assume the ROI is a circular region with radius R , r is the distance to the center of the ROI, and β is a constant that is used to fit the empirical weighting function. Then the weighting function for each pixel (i, j) is computed by

$$w(i, j) = \tilde{w}(r_{ij}), \quad (3)$$

where r_{ij} is the distance from the pixel location (i, j) to the center of the ROI in millimeters.

This weighting function is applied to both the non-homogeneous and the homogeneous steps in the NH-ICD algorithm to speed up convergence within the ROI. In the non-homogeneous steps, its use is straightforward. To compute the PSC, we first form a 2D update map $u^{(m)}$, in which $u^{(m)}(i, j)$ is the sum of the magnitudes of the last update changes along the pixel line (i, j) . Then we multiply the update map with the weighting function before computing the pixel selection criterion. The values of the PSC, which we denote by the function $S^{(m)}(i, j)$, are computed by applying a 2D low pass filter on the weighted update map to reduce high frequency noise, that is,

$$S^{(m)}(i, j) = \sum_{p=-2}^2 \sum_{q=-2}^2 w(i, j) u^{(m)}(i, j) h(i-p, j-q), \quad (4)$$

where the filter kernel h is a 5 by 5 Hamming window. For the homogeneous steps, we use a distributed pixel selection (DPS) method to guarantee that the frequency that each pixel is visited is proportional to its weight $w(i, j)$. To do this, we divide the homogeneous steps into several sub-steps. In each sub-step, we generate a 2D random image $t(i, j)$, and the scaled weights $\gamma w(i, j)$ are then thresholded by $t(i, j)$, the pixels above the threshold are selected to update in this sub-step.

3.2. Reducing Artifacts In The ROI

Iterative reconstruction of an ROI without updating voxels outside the ROI may result in image artifacts due to the propagation of error from the outside to the inside of the ROI. In particular, high frequency features such as organ boundaries and other high contrast transitions outside the ROI reconstructed at lower resolution may introduce errors into a high-resolution ROI image. Final image quality can therefore be improved by including selected edge pixels into the reconstruction mask, in addition to the ROI pixels. Handling a non-compact mask may not be trivial for projection-based algorithms such as Ordered Subsets or Conjugate Gradients, but is natural for ICD, which therefore appears well equipped for ROI reconstruction. In our implementation, we first sum the image volume along the z axis for simplification into a 2D image s , onto which a Sobel edge detector is applied. Both the horizontal and vertical components of the Sobel detection are combined, and only the top 5% of the pixels with the largest magnitudes are kept. Finally, morphological closing is used to fill in the small gaps remaining in the edge map. The resulting edge map is then combined with the pixels in the original ROI to form the set Ω , which is updated in the high-resolution reconstruction.

3.3. Sinogram correction method

In the multi-resolution framework, let $A^{(l)}$ and $A^{(h)}$ be the forward projection operators corresponding to the low-resolution image $x^{(l)}$ and high-resolution image $x^{(h)}$, respectively, and let I_l^h denote the interpolation operator for the image, that is, $x^{(h)} = I_l^h x^{(l)}$. When the forward model is not scale invariant, $A^{(l)} x^{(l)} \neq A^{(h)} I_l^h x^{(l)}$. Although the difference between $A^{(l)} x^{(l)}$ and $A^{(h)} I_l^h x^{(l)}$ may be relatively small compared to the measurement data y , the error sinogram must be corrected for this difference before starting the high-resolution reconstruction in order to avoid artifacts. An initial full forward projection of the high-resolution image prior to the beginning of the high-resolution reconstruction would address this inconsistency, but only at the cost of significant extra




Figure 2. Lower abdomen MAP reconstruction from $64 \times 0.625\text{mm}$ helical CT data with pitch 1.375:1. (a) Low-resolution 700mm reconstruction with voxel size $1.25 \times 1.25 \times 0.625\text{mm}$; (b) High-resolution reconstruction in 319mm field of view and voxel size $0.623 \times 0.623 \times 0.625\text{mm}$.

computation. In this section, we present a fast and simple approximate correction method to avoid a full forward projection.

Let $x^{(l)} = [x_u^{(l)}, x_v^{(l)}]^T$, where $x_u^{(l)}$ contains all the voxels in Ω , and $x_v^{(l)}$ comprises the remaining voxels. Similarly, let $x^{(h)} = [x_u^{(h)}, x_v^{(h)}]^T$. The error sinogram is then corrected as follows:

1. Subtract the forward projection of $x_u^{(l)}$ from the error sinogram, that is, $e \leftarrow e^{(l)} - A^{(l)}[x_u^{(l)}, 0]^T$
2. Add the forward projection of the interpolated targeted voxels $x_u^{(h)}$ to the error sinogram, that is, $e^{(h)} \leftarrow e + A^{(h)}[x_u^{(h)}, 0]^T$

The corrected error sinogram is then $e^{(h)} = A^{(h)}[x_u^{(h)}, 0]^T + A^{(l)}[0, x_v^{(l)}]^T - y$. This is equivalent to a multi-resolution forward projection where the set Ω is forward projected in high-resolution, and the rest of the image is forward projected in low-resolution. Since the low-resolution error sinogram remains available at the end of the low-resolution reconstruction, it is reused to save computation at the beginning of the high-resolution reconstruction.

4. EXPERIMENT RESULTS

This targeted reconstruction approach is now illustrated with clinical CT data. We consider a $64 \times 0.625\text{mm}$ helical scan with pitch 1.375:1 acquired on a GE Lightspeed VCT scanner. The size of the bore is 700mm , and targeted area is of radius 319mm . The reconstructed images are shown in Figure 2, where (a) represents the low-resolution reconstruction with voxel size $1.25 \times 1.25 \times 0.625\text{mm}$ in 700mm field of view, and (b) shows the high-resolution ROI reconstruction with voxel size $0.623 \times 0.623 \times 0.623\text{mm}$ in 319mm field of view.

We compare three methods for targeted reconstruction. The first method is the direct high-resolution reconstruction of the images in full field of view, which reconstructs a volume of size $1124 \times 1124 \times 709$ voxels to result in a 512×512 ROI in 319mm diameter. The second is a basic multi-resolution method, which uses the conventional NH-ICD algorithm instead of the WNH-ICD algorithm. Moreover, in the basic multi-resolution algorithm, a full forward projection of the high-resolution image is performed to initialize the error sinogram for the high-resolution reconstruction, and only the pixels in the ROI are updated in the high-resolution reconstruction. The third method is the improved multi-resolution method which is proposed in this paper.

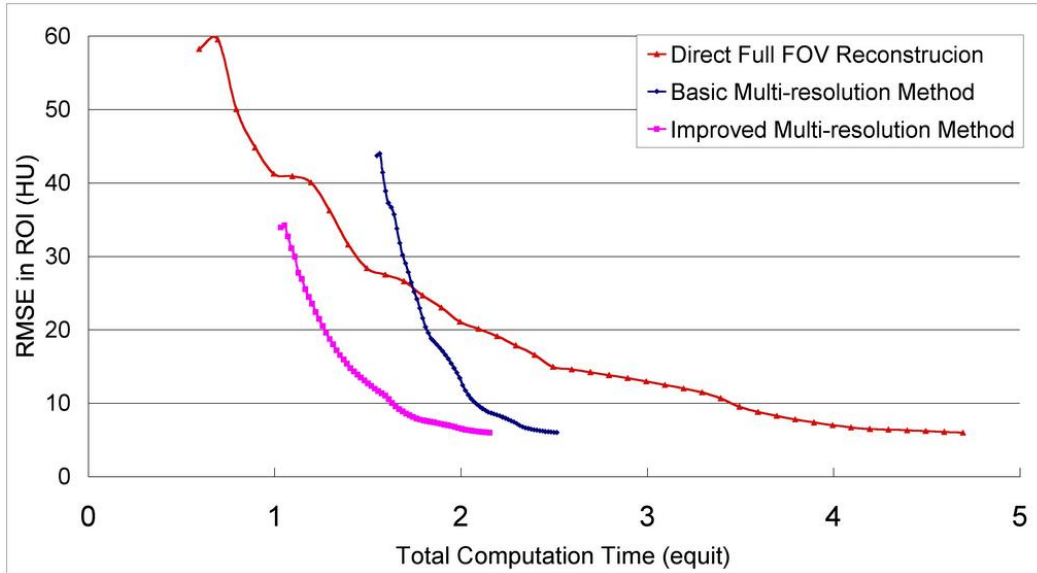


Figure 3. Comparison of computation time between the direct full FOV reconstruction and the basic and improved multi-resolution methods. The unit of the time axis is the equit corresponding to the direct full FOV reconstruction method. The image quality is measured by the RMSE in the ROI.

To compare the speed of the proposed multi-resolution method versus the direct high-resolution reconstruction method, we must first define separate measures of image quality and computational cost. Image quality is measured here by the root mean squared error (RMSE) of the ROI voxels in the reconstructed image, computed against a reference image x^* which we choose as the NH-ICD reconstruction obtained after 10 full iterations in full field of view. The RMSE is then:

$$RMSE = \sqrt{\frac{1}{|U|} \sum_{i \in U} |x_i - x_i^*|^2},$$

where U is the set of ROI voxels and $|U|$ denotes the number of elements in U . To measure the computational cost, we define the unit “equit” (or equivalent iteration) as the time required to update all the voxels in the full FOV at the desired high resolution. So for the conventional ICD algorithm, 1 equit is the time required to perform 1 full iteration of direct full FOV reconstruction. All plots in Figure 3 include the total computation time including all initialization and iterative updating operations. This is important since error sinogram initialization and correction at both high and low resolutions can require significant computational overhead.

Figure 3 shows the convergence plots of the three methods. For a fixed image quality level of RMSE equals 13 HU, the computation time for the first method is 3 equits, while it is only 1.95 equits and 1.5 equits for the basic and improved multi-resolution method. This represents a significant computational saving. In this example, the computation time is reduced by around 35% and 50% using the basic and improved multi-resolution methods.

The ROI weighting function for this example is illustrated in Figure 4 (a). It is normalized for the maximum of the weights to be 1. Figure 4 (b) shows the pixels selected in the first two consecutive homogeneous steps. Magnitudes represent the number of visits per pixel location in these steps. Figure 4 (b) resembles the weighting function shown in (a), which confirms that the voxel sampling frequency is proportional to the ROI weighting. Moreover, 99% pixels are sampled in the first 8 sub-steps, avoiding lack of convergence behavior which could arise if a significant portion of voxels outside the ROI would just not be visited at all.

To compare the image quality, Figure 5 (a) and (b) shows the reconstructed image of the basic and improved multi-resolution methods. Both images are of RMSE = 13 HU. Figure 5 (c) and (d) shows the the error image $\tilde{x} - x^*$ corresponding to (a) and (b) in a Hounsfield window of [-50, 50] HU. By comparing these images, we can find that at the same RMSE level, the errors in (d) appears more uniformly distributed than in (c). Especially, the shadings close to the boundary of the ROI in (c) is reduced in (d).




Figure 4. (a) Analytical ROI weighting function computed for the targeted area; (b) Pixels selected in the first 8 sub-steps of WNH-ICD. Magnitudes represent the number of visits to each pixel.

5. CONCLUSION

In this paper, we present a multi-resolution method for targeted iterative reconstruction. The WNH-ICD method provides faster convergence behavior in the ROI. The sinogram correction method is a simple but effective method to provide the initial error sinogram for the high-resolution reconstruction. By updating the selected edge pixels outside the ROI, we can improve the image quality inside the ROI. Overall, the proposed method can achieve faster speed than the direct full FOV reconstruction without mitigating the image quality.

REFERENCES

1. J.-B. Thibault, K. D. Sauer, C. A. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multi-slice helical CT," *Med. Phys.* **34**(11), pp. 4526–4544, 2007.
2. M. V. Ranganath, A. P. Dhawan, and N. Mullani, "A multigrid expectation maximization reconstruction algorithm for positron emission tomography," *IEEE Trans. on Medical Imaging* **7**, pp. 273–278, December 1988.
3. T. S. Pan and A. E. Yagle, "Numerical study of multigrid implementations of some iterative image reconstruction algorithms," *IEEE Trans. on Medical Imaging* **10**, pp. 572–588, December 1991.
4. S. Oh, A. B. Milstein, C. A. Bouman, and K. J. Webb, "A general framework for nonlinear multigrid inversion," *IEEE Trans. on Image Processing* **14**, pp. 125–140, January 2005.
5. C. W. Stearns, R. M. Manjeshwar, and S. D. Wollenweber, "An efficient algorithm for targeted reconstruction of tomographic data," in *Proc. of IEEE Nucl. Sci. Symp. and Med. Imaging Conf.*, (San Diego, CA), October 29–November 4 2006.
6. Z. Yu, J.-B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh, "Non-homogeneous updates for the iterative coordinate descent algorithm," in *Proceedings of the SPIE/IS&T Symposium on Computational Imaging V*, (San Jose, CA), Jan. 28 - Feb. 1 2007.
7. J.-B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh, "A recursive filter for noise reduction in statistical tomographic imaging," in *Proceedings of the SPIE/IS&T Symposium on Computational Imaging IV*, **6065**, (San Jose, CA), Jan. 16-18 2006.
8. J.-B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh, "New edge-preserving regularization for statistical reconstruction of clinical CT data," in *Proceedings of the Radiological Society of North America (RSNA), SST14*, (Chicago, IL), Nov. 26 - Dec. 1 2006.
9. C. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *IEEE Trans. on Image Processing* **5**, pp. 480–492, March 1996.




Figure 5. Reconstructed image of $RMSE = 13$ HU using (a) basic multi-resolution method and (b) improved multi-resolution method displayed in Hounsfield window of $[-200, 200]$ HU. (c) and (d) shows the error images corresponding to (a) and (b) in a Hounsfield window of $[-50, 50]$ HU.

10. R. G. Keys, "Cubic Convolution Interpolation for Digital Image Processing," *IEEE Transactions on Acoustics Speech and Signal Processing* **29**, pp. 1153–1160, 1981.