

Slice-By-Slice Versus Fully 3D Reconstruction for Parallel-Beam X-ray Microtomography

Michael T. McCann^{1,2,*}, Masih Nilchian¹, and Michael Unser¹

¹Biomedical Imaging Group, EPFL, Lausanne, Switzerland.

²Center for Biomedical Imaging, Signal Processing Core, EPFL, Lausanne, Switzerland.

*michael.mccann@epfl.ch

Abstract: We present a 3D iterative reconstruction method for microtomography, where iterative methods have previously been limited by their high computational cost. Our preliminary results show that it improves reconstruction quality over the traditional slice-by-slice method.

OCIS codes: (100.3010) Image reconstruction techniques; (110.7440) X-ray imaging.

1. Introduction

Iterative methods including regularization have been shown to produce high-quality X-ray reconstructions from fewer and noisier projections than direct methods [1], however, in microtomography (including synchrotron-based tomography [2] and electron tomography), iterative methods have so far been limited to being performed slice by slice (e.g., [3, 4]). This is partly because in the parallel-beam, fixed-axis setting, each X-ray travels only within a single slice, and thus the 3D problem naturally decomposes into a set of 2D problems. But it is also because the large number of measurements collected in these modalities make full 3D iterative reconstruction prohibitively costly.

In this work, we present a computational framework for performing iterative parallel-beam microtomography reconstruction in full 3D, allowing for a broader class of regularizers, including 3D total variation (TV) or Hessian Schatten-norm regularization [5]. This enables the reconstruction to achieve a more favorable trade-off between dose reduction and reconstruction quality than that of slice-by-slice methods. To make the method computationally feasible, we uncover discrete convolutions so that iterations occur at the cost of a 3D filtering operation (an approach which was developed by [6] in the 2D case). In the following sections, we briefly formulate the X-ray tomography reconstruction problem, present the fully 3D method, and show preliminary results comparing it to the slice-by-slice method on a 3D analytical phantom.

2. Problem Formulation

We use the standard definition of the X-ray transform as the line integrals of a function of 2D or 3D space, f . We then perform a generalized discretization, representing f by a sum of shifted kernels: $f(\mathbf{x}) = \sum_{\mathbf{k}} \mathbf{c}[\mathbf{k}] \varphi(\mathbf{x} - \Lambda_{\mathbf{x}} \mathbf{k})$, where $\Lambda_{\mathbf{x}}$ contains the sampling step in each of d dimensions. Common choices for φ in the literature include voxels, splines [7, 8], or other *blobs*. Discretization allows us to define discrete versions of the forward, adjoint, and normal X-ray operators, \mathbf{H} , \mathbf{H}^T , and $\mathbf{H}^T \mathbf{H}$, respectively. The last two of these are needed for our iterative reconstruction algorithms.

3. Reconstruction Methods

Both the slice-by-slice and fully 3D methods are based on the discrete optimization problem,

$$\mathbf{c}^* = \arg \min_{\mathbf{c}} \|\mathbf{g} - \mathbf{H}\mathbf{c}\|_2 + \lambda \Psi(\mathbf{c}), \quad (1)$$

where \mathbf{g} is a sinogram containing X-ray measurements at discrete locations, Ψ is a regularization function, and λ controls the regularization strength. We solve (1) using the now-ubiquitous alternating direction method of multipliers (ADMM) [9].

3.1. Fully 3D

In the fully 3D case, we fix the discretization function, ϕ , to be an isotropic 3D generalized Kaiser-Bessel window (KBW) [10] with $m = 2$, $\alpha = 10.45$, and a spatial support equal to 4 times the sampling step in space. (We pick these values based on our analysis in [11], which shows that they improve the approximation properties of the kernel.) For the regularization term $\Psi(\mathbf{c})$, we use TV, which is the ℓ_1 norm of the magnitude of the gradient vector (approximated with finite differences) at each location in \mathbf{c} . This choice of regularization means that the values of \mathbf{c} are linked in all three dimensions, which is what makes the reconstruction fully 3D. With the kernel and regularizer thus fixed, the optimization proceeds via standard ADMM iterations (in our case, conjugate gradient, followed by the proximal map of Ψ , followed by an update of Lagrange multipliers).

We note that the main bottleneck of this scheme is computing $\mathbf{H}^T \mathbf{g}$ and $\mathbf{H}^T \mathbf{H} \mathbf{c}$. Our algorithm relies on reexpressing these as discrete convolutions, allowing iterations to occur at FFT-cost after a one-time precomputation. We will give the details of these accelerated algorithms for the adjoint and normal operators in future work.

3.2. Slice-by-slice

To frame slice-by-slice reconstruction in the same framework, we set the discretization kernel to be an isotropic KBW in the $\mathbf{x}_0 \mathbf{x}_1$ plane multiplied by a box in the \mathbf{x}_2 direction. As a result, the data fit term in (1) separates along \mathbf{x}_2 such that each row of \mathbf{g} corresponds to only a single slice of \mathbf{c} . We then set the regularization term $\Psi(\mathbf{c})$ to be the 2D TV in the $\mathbf{x}_0 \mathbf{x}_1$ plane. Thus the entire optimization problem splits into a set of 2D problems, which we solve independently via standard ADMM iterations.

The main differences between the slice-by-slice and fully 3D methods are therefore that (1) the regularization in the slice-by-slice method is computed only within and not between slices and (2) the discretization kernel in the slice-by-slice method is anisotropic.

4. Experimental Validation

We now compare the slice-by-slice and fully 3D methods using data from an analytical phantom comprised of ellipsoids. Specifically, we use nine of the ellipsoids from the 3D version of the Shepp-Logan phantom [12]: the ten corresponding to the original 2D Shepp-Logan phantom minus the skull. We generate a sinogram from this phantom by computing X-ray projections analytically on a 601×401 grid for 1,1201 angles spaced evenly from 0 to π , inclusive, around the \mathbf{x}_2 (vertical) axis. We then add zero-mean Gaussian noise to the sinogram so that the signal-to-noise ratio (SNR) is -10dB. This phantom is piecewise constant and therefore a good fit for TV regularization.

To evaluate the robustness of the methods to missing projection angles, we create downsampled versions of the sinogram containing 1201, 301, 101, 51, 26, and 13 projections. For each of these sinograms and for each method, we compute a $161 \times 161 \times 101$ reconstruction for a range of 8 different regularization strengths, $\lambda = 0, a^2, a^4, \dots, a^7$ with a chosen experimentally to cover the range between too little and too much regularization. We compute an SNR for each of these reconstructions using the analytical values of the phantom as the oracle and report the SNR for the unregularized and the best of the regularized reconstructions.

The results are shown in Figure 1. For both methods decreasing the number of views reduces the SNR and regularization greatly improves the SNR. When the number of views is high, the regularized methods perform similarly well, however, in the unregularized case, the fully 3D method gives an improvement of 3.5dB over the slice-by-slice. This underscores that the difference between the methods is not only the 3D regularization, but also the different discretization kernels. In this example, the isotropic discretization kernel of the fully 3D method offers beneficial smoothing that helps remove noise in the absence of regularization.

As the number of views decreases, the gap between the regularized methods increases, which suggests that the fully 3D regularization better handles low-view artifacts. Qualitatively, the difference is most apparent in the $\mathbf{x}_1 \mathbf{x}_2$ cross section, where the slice-by-slice method shows significant variability between slices.

The results suggest that fully 3D reconstruction is superior to slice-by-slice reconstruction, even when the problem naturally decomposes into a set of slices. This difference rests on the fact that fully 3D reconstruction allows the regularization to use information from vertically adjacent slices to reduce noise. Fully 3D reconstruction is a good choice when noise is high and the number of views is low, which corresponds to the low X-ray dose regime. When the desired reconstruction does not fit in memory or when the reconstruction cannot benefit from fully 3D regularization, slice-by-slice reconstruction remains the better alternative. In our future work, we will validate this method on real datasets, including differential phase contrast measurements, and provide the details of our algorithms for accelerating the reconstruction.

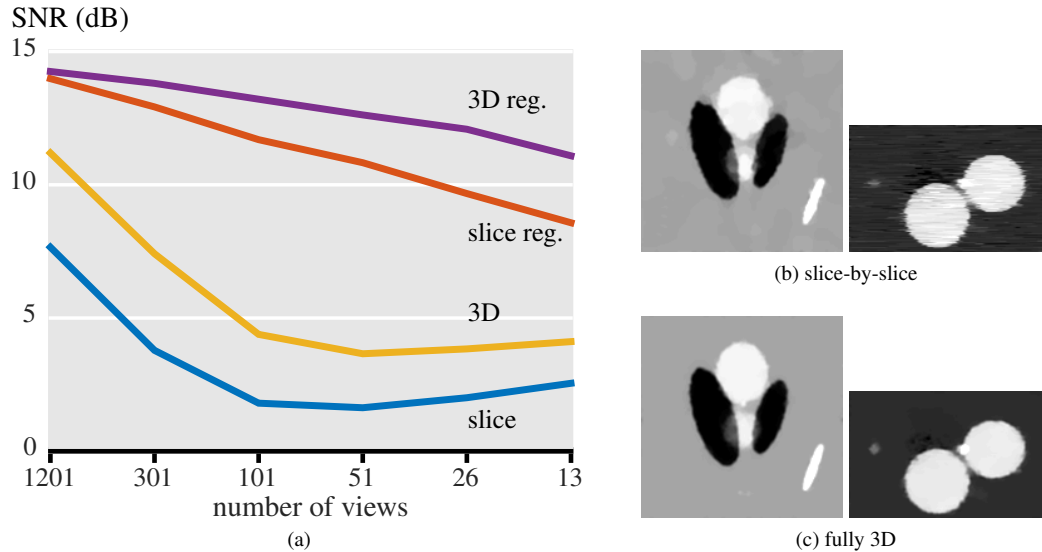


Fig. 1. (a) Reconstruction quality versus number of views for the slice-by-slice and fully 3D reconstruction methods. (b and c) Representative slices for each with 101 views in the x_0x_1 (left) and x_1x_2 (right) cross sections. Slice-by-slice reconstruction results in jitter between slices.

References

1. M. Beister, D. Kolditz, and W. A. Kalender, "Iterative reconstruction methods in X-ray CT," *Physica Med.* **28**, 94–108 (2012).
2. M. Stampanoni, G. Borchert, P. Wyss, R. Abela, B. Patterson, S. Hunt, D. Vermeulen, and P. Regsegger, "High resolution X-ray detector for synchrotron-based microtomography," *Nucl. Instrum. Methods Phys. Res., Sect. A* **491**, 291–301 (2002).
3. Q. Xu, E. Y. Sidky, X. Pan, M. Stampanoni, P. Modregger, and M. A. Anastasio, "Investigation of discrete imaging models and iterative image reconstruction in differential X-ray phase-contrast tomography," *Opt. Express* **20**, 10,724–10,749 (2012).
4. M. Nilchian, C. Vonesch, P. Modregger, M. Stampanoni, and M. Unser, "Fast iterative reconstruction of differential phase contrast X-ray tomograms," *Opt. Express* **21**, 5511–5528 (2013).
5. S. Lefkimiatis, J. P. Ward, and M. Unser, "Hessian Schatten-Norm Regularization for Linear Inverse Problems," *IEEE Trans. Image Process.* **22**, 1873–1888 (2013).
6. A. Delaney and Y. Bresler, "A fast and accurate Fourier algorithm for iterative parallel-beam tomography," *IEEE Trans. Image Process.* **5**, 740–753 (1996).
7. A. Entezari, M. Nilchian, and M. Unser, "A Box Spline Calculus for the Discretization of Computed Tomography Reconstruction Problems," *IEEE Trans. Med. Imaging* **31**, 1532–1541 (2012).
8. F. Momey, L. Denis, C. Burnier, E. Thiebaut, J.-M. Becker, and L. Desbat, "Spline Driven: High Accuracy Projectors for Tomographic Reconstruction From Few Projections," *IEEE Trans. Image Process.* **24**, 4715–4725 (2015).
9. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Found. Trends Mach. Learn.* **3**, 1–122 (2011).
10. R. M. Lewitt, "Multidimensional digital image representations using generalized Kaiser-Bessel window functions," *J. Opt. Soc. Am. A.* **7**, 1834–1836 (1990).
11. M. Nilchian, J. Ward, C. Vonesch, and M. Unser, "Optimized Kaiser-Bessel Window Functions for Computed Tomography," *IEEE Trans. Image Process.* **24**, 3826–3833 (2015).
12. L. A. Shepp, "Computerized tomography and nuclear magnetic resonance." *J. Comput. Assisted Tomogr.* **4**, 94–107 (1980).