

REGION OF INTEREST X-RAY COMPUTED TOMOGRAPHY VIA CORRECTED BACK PROJECTION

Michael T. McCann^{*†}, Laura Vilaclara^{*}, Michael Unser^{*}

^{*}Biomedical Imaging Group, EPFL, Lausanne, Switzerland

[†]Center for Biomedical Imaging Signal Processing Core, EPFL, Lausanne, Switzerland

ABSTRACT

We present a new method for region of interest (ROI) reconstruction from high-resolution, nontruncated data in parallel-ray X-ray computed tomography (CT). Many of the approaches to ROI CT reconstruction in the literature rely on a costly forward projection step to form an ROI-only sinogram. Our approach instead relies on a digital filtering implementation of the normal operator ($\mathbf{H}^T\mathbf{H}$) to compute a back projected version of the ROI-only sinogram that can be used directly for reconstruction, thus eliminating the forward projection step altogether. Results on three synthetic datasets with a variety of experimental conditions show that the method provides accuracy on par with a full reconstruction at a fraction of the time and memory cost.

Index Terms— X-ray tomography, computed tomography, region of interest tomography.

1. INTRODUCTION

In biomedical imaging applications, one is often interested in observing at high resolution a small region of interest (ROI) inside a larger field of view (FOV); for example an organ in a body or a lesion in a tissue. In the case of X-ray computed tomography (CT), ROI imaging is a challenge because each X-ray passing through the ROI also passes through the rest of the specimen. Any ROI reconstruction must account for this additional X-ray attenuation or it will induce significant artifacts (Figure 1).

In this paper, we present a new method for performing ROI reconstruction in the setting of parallel-ray X-ray CT. We assume that high-resolution projection data covering the entire object is available (thus, we are not in the case of truncated tomography as in [1], or of zoom-in tomography as in [2]) but that performing reconstruction of the entire object at this resolution would be prohibitively slow or impossible due to memory limitations. This situation arises in microtomography, where the size of the measured data (called the

This work was funded (in part) by the Center for Biomedical Imaging of the Geneva-Lausanne Universities and EPFL, as well as by the Foundations Leenaards and Louis-Jeantet. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (GA No. 692726 GlobalBioIm).

sinogram) and corresponding full FOV reconstruction can be on the order of $1,000^3$ voxels.

The main approach to X-ray CT ROI reconstruction (sometimes called *reprojection*) is to perform a reconstruction of the full FOV (usually at reduced resolution), forward project the region outside the ROI, then subtract it from the original sinogram. The ROI can then be reconstructed from this ROI-only sinogram. The method in [3] forms the full FOV reconstruction using filtered back projection (FBP); the authors of [4] provide a refinement of the method that deals better with noise. Similarly, the method in [5] uses iterative reconstructions and forward projection, but also considers truncation errors, and the method in [6] uses the same idea to reconstruct a high-resolution image in several pieces.

Our approach is novel in that it does not involve forward projection, but rather applies the normal operator, $\mathbf{H}^T\mathbf{H}$, to form an ROI-only back projection of the sinogram. The method relies on the fact that the normal operator takes the form of a discrete convolution in parallel-ray CT [7], making it essentially free compared to the cost of computing either a forward or back projection alone. (This fact also allows for iterative reconstructions without using forward and back projections.) Avoiding forward projection is especially important when using a high-accuracy forward model, where forward and back projection are costly.

2. METHODS

2.1. Problem Formulation

We begin with a continuous function, f , with a known non-zero support, $\Omega \in \mathbb{R}^d$, e.g. $\Omega = [-1, 1] \times [-1, 1]$. Given the X-ray measurements of f , the standard reconstruction problem is to estimate the values of the function on a grid covering Ω : $\mathcal{X} = \{\mathbf{k}T + \mathbf{x}_0 : \mathbf{k} \in \mathbb{Z}^d, \mathbf{k}T + \mathbf{x}_0 \in \Omega\}$, where T gives the sampling step (usually set to match the sampling step of the detector) and the vector $\mathbf{x}_0 \in \mathbb{R}^d$ specifies the center of the grid. In the ROI reconstruction problem, the goal is to estimate the values of f inside $\Omega_{\text{ROI}} \in \Omega$, and therefore on a restricted grid $\mathcal{X}_{\text{ROI}} = \{\mathbf{k}T + \mathbf{x}_{\text{ROI},0} : \mathbf{k} \in \mathbb{Z}^d, \mathbf{k}T + \mathbf{x}_{\text{ROI},0} \in \Omega_{\text{ROI}}\}$. The idea is that the ROI reconstruction has the same resolution (determined by T) as the full reconstruction, but




Fig. 1. Comparison of results for five X-ray CT ROI reconstruction methods, where the dashed line indicates the area over which the SNR is calculated. Naive reconstruction (c) has significant artifacts due to X-ray attenuation coming from outside the ROI. Reconstructing the full field of view and cropping (b) removes these artifacts, but is time and memory intensive. The proposed method (f) provides the same accuracy as the full reconstruction at a fraction of the computational cost.

smaller dimensions.

2.2. Forward Operators and Reconstruction

All of the ROI reconstruction methods we will present in the next sections rely on solving a convex optimization problem of the form $\arg \min_{\mathbf{c}} \|\mathbf{H}\mathbf{c} - \mathbf{g}\|^2 + R(\mathbf{c})$, where \mathbf{H} is a linear operator that produces a vector of X-ray measurements, \mathbf{g} , from a vectorized image, \mathbf{c} , and R is a regularization function. We use the notation $\mathbf{H}_{\mathcal{X}}$ to specify that the values in \mathbf{c} correspond to samples at the locations specified by the grid \mathcal{X} .

We compare two implementations for \mathbf{H} . The first is to use Matlab’s `radon` command for \mathbf{H} and `iradon` command (with no prefilter) for \mathbf{H}^T . These commands implement a pixel-driven approach to forward and back projection. They are fast but approximate; additionally, the commands are not strictly adjoints of one another. The second is based on a rigorous discretization of the X-ray transform using the separable sinc, described in detail in [8]. Briefly, we define the X-ray transform of a function, f , along the unit vector $\boldsymbol{\theta}$ and with offset \mathbf{y} as

$$\mathcal{P}_{\boldsymbol{\theta}} \{f\}(\mathbf{y}) = \int_{\mathbb{R}} f(t\boldsymbol{\theta} + \mathbf{P}_{\boldsymbol{\theta}^\perp}^T \mathbf{y}) dt, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^{d-1}$ is a coordinate on the detector plane and $\mathbf{P}_{\boldsymbol{\theta}^\perp}^T$ is a $d \times (d-1)$ matrix that has as its columns an orthogonal basis for the detector plane (orthogonal to $\boldsymbol{\theta}$). By expressing f as a sum of shifted sinc functions, we arrive at a discrete, implementable version of (1). The advantage of this forward model is that it is analytically correct; the downside is that it is significantly slower than approximate models because the underlying sinc function has infinite support and therefore each pixel in the space domain affects each pixel in the projection domain and vice versa.

We experiment with two variants of the regularized reconstruction problem. In the first, the regularizer $R(\mathbf{c})$ is infinity when \mathbf{c} has any negative elements and zero otherwise; we denote this as $I_{\geq 0}$. This problem is solved via 500 iterations

of FISTA [9]. In the second, the regularizer additionally includes the TV norm of \mathbf{c} weighted by a positive scalar λ , in which case the optimization proceeds via 100 iterations of ADMM [10] with 5 internal iterations. The parameter λ is not critical for the experiments here and is selected to give reasonable-looking reconstructions. We use the Inverse Problem Library [11] to perform the optimizations.

For both iterative reconstruction methods, the computational cost comes from computing the back projection $\mathbf{H}^T \mathbf{g}$ (once) and the normal operator $\mathbf{H}^T \mathbf{H} \mathbf{c}$ (each iteration). We implement the normal operator as a discrete convolution [7]. For the pixel-driven forward model, we find the corresponding convolution kernel by applying forward and backward projection to a Kronecker delta. For the sinc-driven forward model, we compute the kernel in closed form.

2.3. Comparison Methods

We now describe the four approaches to solving the ROI reconstruction problem that we will use for comparison with the proposed method.

In the **full reconstruction** approach, a full grid is defined with the same sampling step as the desired ROI, but which covers the entire object, $\mathcal{X} = \{\mathbf{k}T + \mathbf{x}_{\text{ROI},0} : \mathbf{k} \in \mathbb{Z}^d, \mathbf{k}T + \mathbf{x}_0 \in \Omega\}$. The entire field of view is reconstructed via $\mathbf{c}_{\text{full}} = \arg \min_{\mathbf{c}} \|\mathbf{H}_{\mathcal{X}} \mathbf{c} - \mathbf{g}\|^2 + R(\mathbf{c})$, and the part of \mathbf{c}_{full} that corresponds to the ROI is extracted. The full reconstruction provides an upper bound on how well we can hope to reconstruct the ROI, but we are interested in cases where the full reconstruction is either too slow or too memory-intensive to be used in practice.

The **naive reconstruction** ignores the contribution of the area outside the ROI. The ROI is reconstructed via $\mathbf{c}_{\text{ROI}} = \arg \min_{\mathbf{c}} \|\mathbf{H}_{\mathcal{X}_{\text{ROI}}} \mathbf{c} - \mathbf{g}\|^2 + R(\mathbf{c})$. The advantage here is that the inverse problem is much smaller than in the full reconstruction (only \mathbf{c}_{ROI} needs to fit in memory). The disadvantage is the reconstruction will contain serious artifacts as it attempts to explain the entire sinogram only using mass inside the ROI. Practically speaking, the naive reconstruction

can be useful, since the artifacts tend to be low frequency and qualitative information about the structure inside the ROI can be discerned.

The **interpolated reconstruction** reconstructs the entire FOV at a low resolution and interpolates to find the high resolution ROI. Specifically, a low resolution reconstruction grid that covers Ω : $\mathcal{X}_{LR} = \{\mathbf{k}T_{LR} + \mathbf{x}_0 : \mathbf{k} \in \mathbb{Z}^d, \mathbf{k}T_{LR} + \mathbf{x}_0 \in \Omega\}$ is defined such that $T_{LR} > T$ and therefore $|\mathcal{X}_{LR}| \approx |\mathcal{X}_{ROI}|$. A low resolution reconstruction is computed via $\mathbf{c}_{LR} = \arg \min_{\mathbf{c}} \|\mathbf{H}_{\mathcal{X}_{LR}} \mathbf{c} - \mathbf{g}\|^2 + R(\mathbf{c})$. Finally, \mathbf{c}_{LR} is interpolated to find \mathbf{c}_{ROI} . If speed is critical, the measurements themselves can also be downsampled (in space and/or angularly), along the lines of [12], because high-resolution measurements (sampling rate $1/T$) are not necessary for a low-resolution reconstruction (sampling rate $1/T_{LR}$). This approach has a similar computational cost to the naive reconstruction, but avoids the artifacts. The disadvantage is that the interpolation is unable to recover fine details in the ROI. (In our experiments, we interpolate with a separable cubic B-spline and downsample the sinogram in space so that the sampling rate for \mathbf{g} is approximately T_{LR} .)

The final comparison method uses **forward projection** to remove the effects of the region outside the ROI from the sinogram. \mathbf{c}_{LR} is reconstructed as in the interpolation approach. The ROI is zeroed out according to $\mathbf{c}_{LR, \setminus ROI}[\mathbf{k}] = 0$, when $\mathbf{k}T_{LR} + \mathbf{x}_0 \in \Omega_{ROI}$ and $\mathbf{c}_{LR, \setminus ROI}[\mathbf{k}] = \mathbf{c}_{LR}[\mathbf{k}]$, otherwise. Finally, a sinogram containing only measurements from the ROI is formed according to $\mathbf{g}_{ROI} = \mathbf{g} - \mathbf{H}_{\mathcal{X}_{LR}} \mathbf{c}_{LR, \setminus ROI}$ and the ROI is reconstructed via $\mathbf{c}_{ROI} = \arg \min_{\mathbf{c}} \|\mathbf{H}_{\mathcal{X}_{ROI}} \mathbf{c} - \mathbf{g}_{ROI}\|^2 + R(\mathbf{c})$. This method creates a high-resolution, artifact free reconstruction of \mathbf{c}_{ROI} at the cost of two small reconstructions and a single forward projection. This method is very similar to the forward projection-based methods in the literature [3–6] and serves as a state-of-the-art comparison.

We note that **FBP** can also be used for ROI reconstruction. We do not compare against it here because the advantages of iterative reconstruction over FBP are well documented in the literature, but note that it can give excellent results when the number of views is high and the noise is low.

2.4. Proposed Method

Our proposed method is similar in spirit to the forward-corrected reconstruction, with the key advantage that it avoids the expensive forward projection step by correcting the sinogram in the **back projection** domain. Just as for the forward projection approach, \mathbf{c}_{LR} is reconstructed and the ROI is zeroed out. The back projected ROI-only sinogram is computed via $\mathbf{H}_{\mathcal{X}_{ROI}}^T \mathbf{g}_{ROI} = \mathbf{H}_{\mathcal{X}_{ROI}}^T \mathbf{g} - \mathbf{H}_{\mathcal{X}_{ROI}}^T \mathbf{H}_{\mathcal{X}_{LR}} \mathbf{c}_{LR, \setminus ROI}$. This operation is fast because $\mathbf{H}_{\mathcal{X}_{ROI}}^T \mathbf{H}_{\mathcal{X}_{LR}}$ can be implemented as the discrete convolution $\mathbf{H}_{\mathcal{X}_{LR}}^T \mathbf{H}_{\mathcal{X}_{LR}}$ followed by interpolation. Finally, we note that, in general, $\mathbf{c}_{ROI} = \arg \min_{\mathbf{c}} \|\mathbf{H} \mathbf{c} - \mathbf{g}\|^2 + R(\mathbf{c})$ can be considered a function of $\mathbf{H}^T \mathbf{g}$ rather than of \mathbf{g} because \mathbf{g} does not appear alone in

the gradient. This fact allows us to reconstruct the ROI via $\mathbf{c}_{ROI} = \arg \min_{\mathbf{c}} \|\mathbf{H}_{\mathcal{X}_{ROI}} \mathbf{c} - \mathbf{g}_{ROI}\|^2 + R(\mathbf{c})$ without having to explicitly compute \mathbf{g}_{ROI} .




Fig. 2. Test images used in the experiments. Each image is 513×513 px, and the black squares mark ROIs of size 256×256 px and 128×128 px.

3. EXPERIMENTS

We compared the described methods by using them to perform ROI reconstructions from synthetic data under a variety of conditions. Though our experiments use 2D phantoms, all the methods generalize to the 3D case (fixed axis or not). To generate synthetic sinograms, we used three different phantoms (Figure 2):

The **filament** dataset consists of 400 generalized Kaiser-Bessel window (KBW) functions [13] positioned along a random walk. This simulated filament is located inside a single, large KBW. Projections are computed in closed form according to (1) and the definition of the KBW. The aim of this dataset is to provide a structured, high resolution object for reconstruction while remaining nearly bandlimited.

The **ellipse** dataset consists of 100 ellipses with random centers and orientations located within a large circle. Projections are computed in closed form according to (1). The aim of this dataset is to provide a very difficult ROI reconstruction problem with lots of density and high frequencies outside the ROI.

The **medical** dataset is based on a human CT image from the National Biomedical Imaging Archive website [14]. Starting with the original 513×513 px image, we apply a discrete forward model to simulate the sinogram.¹ The purpose of this dataset is simply to compare the methods on an image with realistic structures; it does not involve real data (raw measurements from a CT scanner) and suffers from inverse crime in the sense that the model used to create the measurements is the same as the one used to reconstruct.

For each dataset, we form a 513×513 px ground truth image (by sampling the continuously-defined function in the case of the filament and ellipse datasets) and

¹We chose to use the pixel-driven model to generate the synthetic sinogram; results obtained using the sinc-driven model were within 1 dB.

a 729 px by 513 views sinogram. We add random Gaussian noise to the sinogram so that its SNR is 40 or 60 dB: $\text{SNR}(g, \tilde{g}) = 10 \log_{10}(\sum_{\mathbf{k}} |g[\mathbf{k}]|^2 / \sum_{\mathbf{k}} |\tilde{g}[\mathbf{k}] - g[\mathbf{k}]|^2)$. We also define a 256 by 256 or 128 by 128 px ROI at the center of the ground truth image. We reconstruct this ROI via each of the methods described in Sections 2.3 and 2.4, using each combination of forward model and regularization described in Section 2.2. We record the SNR of the ROI reconstruction with respect to the ground truth ROI, ignoring pixels on the edge of the ROI to minimize the impact of edge effects. We also record the total time each method took to perform the reconstruction.

4. RESULTS

Selected accuracy results for ROI size 128 by 128 px are given in Table 1; for the full set quantitative results, see the reproducible research page for this paper [15]. Qualitative results for the medical dataset (ROI size 128 by 128, 40 dB sinogram, sinc operator) are given in Figure 1. The trend in accuracy between the methods is that the full reconstruction provides the highest SNR, with the FP and BP coming close to the same accuracy (within 1 dB). The interpolation-based reconstruction is slightly worse (and more blurred-looking). The naive reconstruction is uniformly the worst, though its performance improves when the ROI is size 256 by 256 px (not shown); this is logical because the strength of artifacts in the naive reconstruction should be proportional to the amount of mass outside the ROI. An exception to this trend is that the interpolation-based reconstruction can outperform all other methods when the sinogram is noisy and there is no regularization (ellipse-40 and medical-40 in Table 1). This happens only because the interpolation provides smoothing that is beneficial when noise is high; the advantage goes away when TV regularization is used.

Comparing the sinc- and pixel-based operators shows an advantage for the sinc-based operator on the filament and medical datasets when only nonnegativity is enforced. This advantage disappears when the sinogram is noisy and TV regularization is used; we hypothesize this is because the reconstruction error is dominated by noise rather than model inaccuracy. Comparing nonnegativity to TV regularization shows the expected trend: TV regularization dramatically improves the reconstructions in the high noise case. The filament dataset does not follow this trend, probably because it does not satisfy the piece-wise constant model assumed by TV regularization.

The timing results are given in Table 2. For ROIs of size 128 by 128 px with the pixel-driven forward model, the full reconstruction is around ten times slower than the naive and interpolation-based reconstructions and about five times slower than the forward and backward corrected reconstructions. For the pixel-driven operator, these numbers do not reflect the time savings of avoiding the forward projection

Table 1. SNR of ROI reconstruction (dB).

Experiment		full	naive	interp.	FP	BP
pixel	filament-40	29.6	13.3	20.3	29.5	29.5
	filament-60	40.8	13.4	20.3	40.8	40.8
	ellipse-40	16.1	-3.9	25.2	16.2	16.2
	$I_{\geq 0}$ ellipse-60	30.4	-3.8	27.1	31.4	31.4
	medical-40	18.8	-1.2	27.1	18.7	18.7
	medical-60	34.9	-1.1	28.5	34.8	34.9
sinc	filament-40	31.6	13.3	22.1	31.7	31.8
	filament-60	49.1	13.3	22.1	46.2	50.1
	$I_{\geq 0}$ ellipse-40	17.2	-3.9	26.4	18.1	18.1
	ellipse-60	27.4	-3.9	27.8	32.3	32.4
	medical-40	20.3	-1.2	28.5	20.6	20.6
	medical-60	35.6	-1.2	29.7	34.9	35.6
pixel TV	filament-40	28.9	13.2	21.6	28.9	28.9
	ellipse-40	30.0	-3.9	27.1	29.9	29.9
	medical-40	31.3	-1.2	28.6	31.2	31.2
sinc TV	filament-40	28.7	13.2	22.2	28.7	28.7
	ellipse-40	29.9	-3.9	27.0	29.8	29.9
	medical-40	31.2	-1.2	29.1	31.1	31.2

Table 2. Reconstruction time (s).

Operator	full	naive	interp.	FP	BP
pixel-128	22.6	2.2	2.7	4.8	4.6
sinc-128	37.3	4.6	4.6	82.3	8.7
pixel-256	22.5	7.4	5.9	13.4	13.2
sinc-256	36.1	12.3	10.6	600.0	22.8

because the time cost is dominated by the reconstruction. However, the difference is apparent for the sinc-driven FP method, which underscores why the proposed BP method is useful when using this forward model. Not reflected here is the fact that the full reconstruction requires storing the high-resolution reconstruction in memory—trivial here with such small images, but substantial when dealing with 3D reconstruction. Thus, the full method may not even be feasible in practice.

5. CONCLUSION

We have presented a new method for X-ray CT ROI reconstruction based on adjusting the back-projected sinogram. Relying on the fact that the normal operator in parallel-ray X-ray CT is a discrete convolution, the method avoids costly forward projection and provides the same accuracy as the state-of-the-art reprojected method. In future work, experiments on more realistic simulations (with Poisson or mixed Gaussian and Poisson noise) or real data could be performed, or the method could be extended to other modalities, such as fan and cone-beam CT or MRI.

6. REFERENCES

- [1] T. A. Bubba, D. Labate, G. Zanghirati, S. Bonettini, and B. Goossens, "Shearlet-based regularized ROI reconstruction in fan beam computed tomography," in *Wavelets and Sparsity XVI*, San Diego, California, Aug. 2015, vol. 9597, p. 95970K, International Society for Optics and Photonics.
- [2] Max Langer and Françoise Peyrin, "A wavelet algorithm for zoom-in tomography," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Rotterdam, Netherlands, Apr. 2010.
- [3] Andy Ziegler, Tim Nielsen, and Michael Grass, "Iterative reconstruction of a region of interest for transmission tomography," *Med. Phys.*, vol. 35, no. 4, pp. 1317–1327, Apr. 2008.
- [4] Felix K. Kopp, Radin A. Nasirudin, Kai Mei, Andreas Fehringer, Franz Pfeiffer, Ernst J. Rummeny, and Peter B. Nol, "Region of interest processing for iterative reconstruction in x-ray computed tomography," in *Medical Imaging 2015: Physics of Medical Imaging*, Orlando, Florida, Mar. 2015, vol. 9412, p. 94122E, International Society for Optics and Photonics.
- [5] Keisuke Yamakawa and Shinichi Kojima, "Two-step iterative reconstruction of region-of-interest with truncated projection in computed tomography," in *Medical Imaging 2014: Physics of Medical Imaging*, San Diego, California, Mar. 2014, vol. 9033, p. 90333B, International Society for Optics and Photonics.
- [6] Y. De Witte, J. Vlassenbroeck, and L. Van Hoorebeke, "A multiresolution approach to iterative reconstruction algorithms in X-ray computed tomography," *IEEE Trans. Image Process.*, vol. 19, no. 9, pp. 2419–2427, Sept. 2010.
- [7] Michael T. McCann, Masih Nilchian, Marco Stambanoni, and Michael Unser, "Fast 3D reconstruction method for differential phase contrast X-ray CT," *Opt. Express*, vol. 24, no. 13, pp. 14564, June 2016.
- [8] M. T. McCann and M. Unser, "High-Quality Parallel-Ray X-Ray CT Back Projection Using Optimized Interpolation," *IEEE Trans. Image Process.*, vol. 26, no. 10, pp. 4639–4647, Oct. 2017.
- [9] A. Beck and M. Teboulle, "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, Jan. 2009.
- [10] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.
- [11] Michael Unser, Emmanuel Soubies, Ferrol Soulez, Michael McCann, and Lorne Donati, "GlobalBioIm: A Unifying Computational Framework for Solving Inverse Problems," in *Imaging and Applied Optics 2017*, San Francisco, California, June 2017, p. CTu1B.1, Optical Society of America.
- [12] S. Basu and Y. Bresler, " $O(N^2 \log_2 N)$ filtered back-projection reconstruction algorithm for tomography," *IEEE Trans. Image Process.*, vol. 9, no. 10, pp. 1760–1773, Oct. 2000.
- [13] Robert M. Lewitt, "Multidimensional digital image representations using generalized Kaiser-Bessel window functions," *J. Opt. Soc. Am. A.*, vol. 7, no. 10, pp. 1834–1836, Oct. 1990.
- [14] "National Biomedical Imaging Archive (NBIA)," Sept. 2017, available at <https://imaging.nci.nih.gov/>.
- [15] Michael T. McCann, "Code and Supplementary Material for Region of Interest X-Ray Computed Tomography via Corrected Back Projection," Zenodo, Oct. 2017, available at <https://doi.org/10.5281/zenodo.1009068>.