

Supervised Reconstruction for Silhouette Tomography

Evan Bell^a, Michael T. McCann^b, Marc Klasky^b

^aMichigan State University; East Lansing, MI

^bTheoretical Division, Los Alamos National Laboratory; Los Alamos, NM

Abstract

In this paper, we introduce *silhouette tomography*, a novel formulation of X-ray computed tomography that relies only on the geometry of the imaging system. We formulate silhouette tomography mathematically and provide a simple method for obtaining a particular solution to the problem, assuming that any solution exists. We then propose a supervised reconstruction approach that uses a deep neural network to solve the silhouette tomography problem. We present experimental results on a synthetic dataset that demonstrate the effectiveness of the proposed method.

Introduction

X-ray computed tomography (CT) is an imaging modality that involves collecting X-ray projection images of the same object from many angles and using them to reconstruct a 3D image of the object. It has applications in medicine [20], nondestructive testing [11], and security [25]. When the number of projection angles is large and the scanner is well characterized in terms of its geometry, the spectrum of the X-rays used, and the spectral response of the detector, CT reconstruction may be accurately performed by either linear [2] methods or model-based iterative reconstruction [26, 12].

More recently, research on CT reconstruction has focused on so-called *extreme imaging*, wherein the projections are limited in number [14, 9], angular coverage [23], or signal-to-noise ratio [10, 16]. In these settings, heavy use of prior information—either in the form of regularization or machine learning—is needed to form an adequate reconstruction. Especially in the case of learning-based approaches to CT reconstruction, many methods have been developed and evaluated in a setting where both testing and training data are generated in silico and use exactly the same imaging model (sometimes called *inverse crime*). While this approach is useful at the proof-of-concept stage, it does not provide an accurate picture of how these methods would perform on real, experimental data.

In this work, we present *silhouette tomography (ST)*, a novel formulation of the X-ray CT reconstruction problem that sidesteps the need to accurately model a CT system in order to perform model-based reconstruction or generate training data for learning-based reconstruction. In ST, rather than relying on the grey-level values of the measured projections, we instead only use them to determine if an object of interest is present in each pixel (equivalently, along each ray). In doing so, we transform the CT reconstruction problem into one that is purely geometric, and therefore only requires knowledge of the geometry of the imaging system (e.g., the projection angles and the cone angle). This simplifies system calibration and increases the chance that a learning-based algorithm trained on simulation data will generalize well to real, experimental data. This approach has the added benefit of en-

abling reconstruction of objects in cluttered scenes and/or in the presence of highly-absorbing objects, provided that the object of interest can be accurately segmented in each projection view. On the other hand, choosing the ST formulation instead of the CT one generally leads to a less well-posed and therefore more challenging reconstruction problem.

Related work. In binary tomography [22, 21, 18, 8], the object to be reconstructed is binary, but—distinct from ST—the measurements are real-valued. In shadow tomography [15], the actual shadows cast by an object, rather than its X-ray projections, are used to determine its shape.

Outline. In the next section, we formulate the silhouette tomography reconstruction problem and define the maximal reconstruction, which provides a simple way of finding at least one solution to any ST problem. We then present a learning based-approach to ST reconstruction. Finally, we provide an experimental comparison of these methods and conclude.

Silhouette Tomography

The silhouette tomography problem is to recover the nonzero support of an object from its binarized X-ray projections. Mathematically, we represent the object as a binary vector $\mathbf{x} \in \{0, 1\}^N$, with either $N = N_1 \times N_2$ (height \times width, 2D problems) or $N = N_1 \times N_2 \times N_3$ (depth \times height \times width, 3D problems). Let $\mathbf{x}[n] = 1$ if the object occupies pixel/voxel n and $\mathbf{x}[n] = 0$ otherwise. The binarized X-ray projection measurements are $\mathbf{y} \in \{0, 1\}^M$, with $M = V \times M_1$ (number of views \times length, 2D problems) or $M = V \times M_1 \times M_2$ (number of views \times height \times width, 3D problems). Let $\mathbf{y}[m] = 1$ if the object occupies any of the pixels/voxels along ray m and $\mathbf{y}[m] = 0$ otherwise. We can relate the projections \mathbf{y} to the object \mathbf{x} via the silhouette tomography operator, $S: \mathbb{R}^N \rightarrow \{0, 1\}^M$, defined as

$$\mathbf{y} = S(\mathbf{x}) = T_{>0}(\mathbf{H}\mathbf{x}), \quad (1)$$

where $\mathbf{H} \in \mathbb{R}^{M \times N}$ is a linear X-ray transform operator with $\mathbf{H}[m, n] \geq 0$ and $T_{>0}$ is a thresholding operator defined by

$$T_{>0}(\mathbf{x})[n] = \begin{cases} 1 & \mathbf{x}[n] > 0; \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

In contrast to the linear tomography problem, the silhouette tomography problem generally admits multiple solutions even when the number of views is large. The linear tomography problem is to recover a real-valued object $\mathbf{x} \in \mathbb{R}^N$ from its real-valued measurements $\mathbf{y} \in \mathbb{R}^M$ given by

$$\mathbf{y} = \mathbf{H}\mathbf{x}, \quad (3)$$

where $\mathbf{H} \in \mathbb{R}^{M \times N}$ is again a linear X-ray transform operator. For this linear problem, the solution is unique when $M \geq N$ because

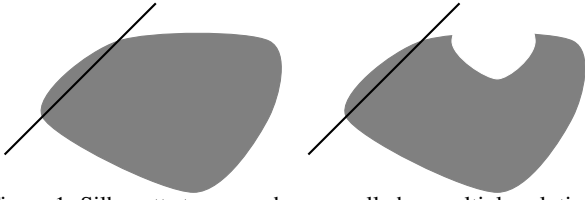


Figure 1: Silhouette tomography generally has multiple solutions. For example, each ray intersecting the left object also intersects the right one and each ray not intersecting the left object also does not intersect the right one.

\mathbf{H} can be assumed to have linearly independent rows. In silhouette tomography, we have strictly less information because the measurements are binary. While a theoretical analysis of its ill-posedness is beyond the scope of this paper, it is easy to come up with examples of shapes that cannot be distinguished by measurements of the form (1), see, e.g., Figure 1.

For any silhouette tomography problem, we now show that there is always one solution that is uniquely defined and easy to compute; we call it the maximal solution.

Definition 1 (Maximal solution). The *maximal solution*, $\mathbf{x}_{\max}(\mathbf{y})$, of (1) is the \mathbf{x} that has the maximal norm among all \mathbf{x} 's that satisfy (1),

$$\mathbf{x}_{\max}(\mathbf{y}) = \arg \max_{\mathbf{x}} \|\mathbf{x}\|_2^2 \quad \text{s.t.} \quad S(\mathbf{x}) = \mathbf{y}. \quad (4)$$

We will sometimes drop the argument (\mathbf{y}) when it is clear from context.

Theorem 1 (Computing the maximal solution). *When any solution to (1) exists, the maximal solution of (1) is given by*

$$\mathbf{x}_{\max}(\mathbf{y}) = \neg T_{>0}(\mathbf{H}^T(-\mathbf{y})), \quad (5)$$

where \neg denotes Boolean negation (i.e., swapping 0 and 1).

Proof. We first show that \mathbf{x}_{\max} satisfies (1). Equation (1) implies that $\mathbf{y}[m] = 1$ if and only if that there exists an n such that $\mathbf{H}[m, n] > 0$ and $\mathbf{x}[n] = 1$. Conversely, $\mathbf{y}[m] = 0$ if and only if for all n such that $\mathbf{H}[m, n] > 0$, $\mathbf{x}[n] = 0$. The maximal reconstruction has $\mathbf{x}_{\max}[n] = 1$ for all n except those that fall into the later case, i.e., it is only zero where it must be: by the construction (5), $\mathbf{x}_{\max}(\mathbf{y})[n] = 0$ if and only if $\mathbf{H}^T(-\mathbf{y})[n] > 0$, which implies that there exists m such that $\mathbf{y}[m] = 0$ and $\mathbf{H}[m, n] > 0$. Thus, if a solution to (1) exists, \mathbf{x}_{\max} is a solution to (1).

Next, we show that \mathbf{x}_{\max} has maximal norm among all solutions to (1). Suppose $\mathbf{z} \in \{0, 1\}^N$ with $\|\mathbf{z}\|_2^2 > \|\mathbf{x}_{\max}\|_2^2$. Then, there exists some n' such that $\mathbf{z}[n'] = 1$ but $\mathbf{x}_{\max}[n'] = 0$. Then, by the same reasoning as above, there exists an m' such that $\mathbf{y}[m'] = 0$ and $\mathbf{H}[m', n'] > 0$. We can then calculate $(\mathbf{H}\mathbf{z})[m'] = \sum_{j=1}^N \mathbf{H}[m', j]z[j] \geq \mathbf{H}[m', n']z[n'] > 0$. Thus $T_{>0}(\mathbf{H}\mathbf{z})[m'] = 1$, but $\mathbf{y}[m'] = 0$. Hence \mathbf{z} is not a solution to (1). This shows that \mathbf{x}_{\max} has the largest norm of any solution to (1), as claimed. \square

Preparing Data. We now describe how ordinary X-ray CT projection data can be processed so that it can be reconstructed using ST. In the simplest case wherein we are imaging a single

object, it suffices to simply threshold the projection images, creating binary masks indicating the presence or absence of the object at each pixel in each projection. If the goal is to reconstruct multiple objects, then each one should be segmented separately and one ST reconstruction performed for each. Automating this segmentation process is outside the scope of the current work.

Proposed Method: Supervised Reconstruction

We propose a deep learning-based approach to solving the ST reconstruction problem. The approach is based on using a database of paired objects, \mathbf{x} , and measurements, \mathbf{y} , along with a known linear X-ray transform operator, \mathbf{H} , to learn to reconstruct \mathbf{x} from \mathbf{y} . Thus, given D tuples of training data $(\mathbf{x}_d, \mathbf{y}_d, \mathbf{H})$, $d = 1, \dots, D$, we train our network weights according to

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \frac{1}{D} \sum_{d=1}^D \|\mathbf{x}_d - \mathbf{f}_{\boldsymbol{\theta}}(\mathbf{H}^T \mathbf{y}_d)\|_2^2, \quad (6)$$

where \mathbf{f} is a neural network with parameters $\boldsymbol{\theta}$.

Proposed Architecture. Because the ST reconstruction problem is highly ill-posed, the choice of \mathbf{f} is important for obtaining a useful solution to (6). The model we propose is a deep U-Net based on the architecture described in [24].

A complete schematic of the proposed architecture is shown in Figure 2. The proposed architecture has six down-sampling blocks and six up-sampling blocks. These blocks are inspired by the residual blocks used in BigGAN [3]. The internal structure of each of these blocks is shown in the figure. In each “down” block, the number of channels increases by 64, and the spatial dimension of the features decreases by a factor of two. In the “down” blocks, the first 3×3 convolution increases the number of channels. In each “up” block, the number of output channels decreases by 64 relative to the previous block, and the spatial dimension of the features increases by a factor of two. In the “up” blocks, the first 3×3 convolution decreases the number of channels. The network also includes skip connections, indicated by the gray dashed arrows in the figure. Features extracted in the down-sampling path are concatenated with those of the same size in the up-sampling path before each “up” block.

Batch normalization [7] is used throughout the network. This operation is indicated by “BN” in the figure. The sigmoid linear unit, also known as “SiLU” or “Swish,” is used as the activation function throughout the network, and is defined pointwise for $x \in \mathbb{R}$ by

$$\text{Swish}(x) = x \cdot \sigma(x) = \frac{x}{1 + e^{-x}}, \quad (7)$$

where σ is the standard logistic sigmoid function.

In total, our proposed architecture has approximately 22.3 million trainable parameters.

Experiments and Results

Dataset. We created a synthetic dataset to test the proposed approach to silhouette tomography. The dataset consisted of 2,036 3D objects (e.g., airplanes, birdhouses, bottles, ...) chosen at random from the ShapeNet dataset [4]. In order to generate binary ground truth volumes that were suitable for our experiments, we preprocessed each object in the following way: We

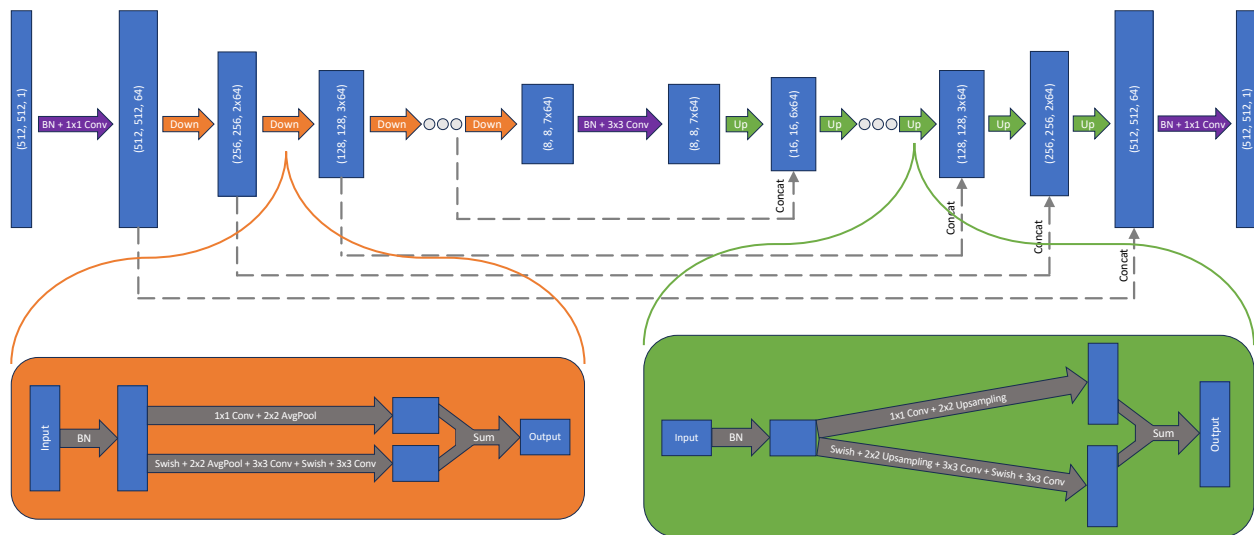


Figure 2: Schematic of the proposed neural network architecture. The internal structure of each downsampling and upsampling block is shown in the orange and green boxes respectively.

made the model watertight using code from [6]. Then, we randomly rotated the model using Trimesh [5] and voxelized on a $511 \times 511 \times 511$ grid using the method of [13]¹.

The dataset was then randomly split into 1,832 volumes for training, 102 volumes for validation, and 102 volumes for testing.

For each volume, we simulated X-ray measurements using parallel beam geometry with 8 equispaced views at angles in the range $[0, \pi)$. All X-ray projections and backprojections were simulated using the ASTRA toolbox [1]. With this choice of geometry, the reconstruction problem can be completely separated into 2D slices.

Most of the volumes in the dataset contained some slices with no object content. These empty slices were not used during training, validation, or testing. Once these slices were removed, there were 698 485 slices for training, 39 628 slices for validation, and 38 824 slices for testing. An object from the testing set and one corresponding slice is shown in Figure 3.

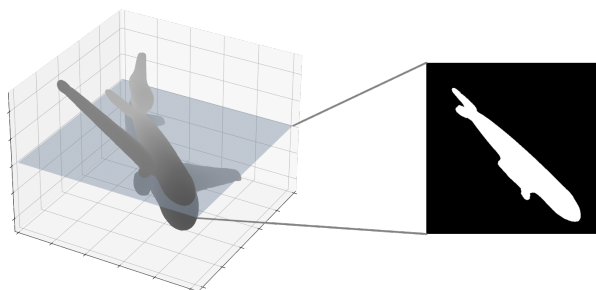


Figure 3: An object from the testing set and one slice from it.

Training details. Because the proposed architecture repeatedly down-samples its input by a factor of two, it is most naturally suited to an input size with many factors of two. To accommodate this, every input to the network was zero-padded to a size of 512×512 , and the output was cropped back to a size of 511×511 .

¹Available at https://github.com/nikw1/inside_mesh.

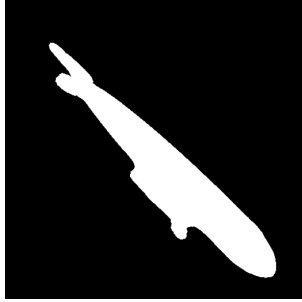
Method	MSE	PSNR	SSIM
maximal	0.0628	15.29	0.921
U-Net	0.0229	22.96	0.875
binarized U-Net	0.0283	22.02	0.952

Table 1: Quantitative comparison of maximal reconstruction and the proposed method on a silhouette tomography reconstruction problem. The best result in each row is indicated in **bold**.

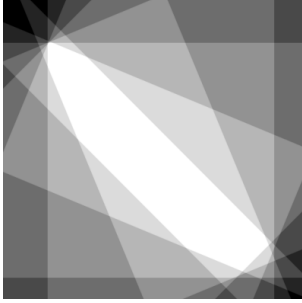
The training objective was given by (6). The model was implemented using PyTorch and trained using the Adam optimizer with a learning rate of 10^{-5} and a batch size of 8. The model was trained for a total of four passes over the training set (epochs), after which both the training and validation losses had empirically converged. The training took approximately 19 hours using eight NVIDIA GeForce RTX 2080 Ti GPUs.

Quantitative comparison. Table 1 gives a quantitative comparison of the reconstruction methods on the test set using three metrics: mean squared error (MSE), peak signal to noise ratio (PSNR), and the structural similarity index measure (SSIM). The numbers given in the table are the average for the three metrics across every slice of the test set. The SSIM was computed using the implementation in the TorchMetrics library with its default parameters [17]. The row “U-Net” gives metrics computed using the raw (floating point) output of the trained network. The row “binarized U-Net” gives metrics computed after thresholding the network’s output at a value of 0.5. There was one slice in the test set for which the thresholded neural network reconstruction had an MSE of 0, resulting in an infinite PSNR. The PSNR of this slice was discarded when computing the average PSNR over the test set.

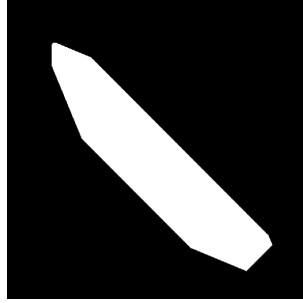
The neural network reconstruction significantly outperforms the maximal reconstruction in terms of both MSE and PSNR, both with and without the binarization. The maximal reconstruction outperforms the raw network output in terms of SSIM, but still falls short of the thresholded network output. This is likely be-



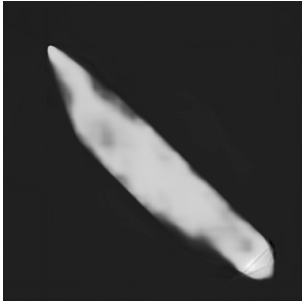
(a) ground truth \mathbf{x}



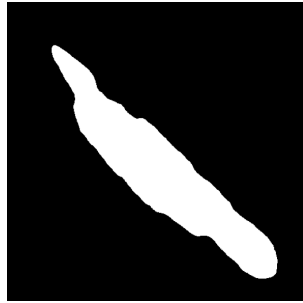
(b) network input $\mathbf{H}^T \mathbf{y}$



(c) maximal reconstruction $\mathbf{x}_{\max}(\mathbf{y})$



(d) NN reconstruction



(e) NN reconstruction (binarized)

Figure 4: Qualitative comparison of the proposed reconstruction methods on one slice of the test set.

cause SSIM penalizes differences in overall contrast between images, where the variance of an image is used as a proxy for its contrast [19]. Therefore the network is penalized for any uncertainty in its reconstructions, i.e. outputting pixel values between 0 and 1. Hence thresholding improves the SSIM of the supervised reconstruction, despite lowering its MSE and PSNR.

Qualitative comparison. Figure 4 gives a qualitative comparison of the maximal reconstruction and the learned reconstruction on a representative slice of the test set. It is evident that the maximal reconstruction significantly overestimates the object content of the slice, whereas the proposed method recovers some of the smooth structures present in the ground truth.

By reconstructing each slice of a 3D volume separately and stacking them, we can perform 3D reconstruction. Figure 5 gives a comparison of a full 3D volume of the test set reconstructed using both the maximal and proposed methods.

Comparison with linear tomography. For extra context, we include a comparison between the proposed supervised ap-

Method	MSE	PSNR	SSIM
U-Net	0.00123	34.20	0.982
binarized U-Net	0.00160	33.27	0.990

Table 2: Quantitative results for supervised linear tomography reconstruction using the proposed network architecture. The best result in each row is indicated in **bold**.

proach to silhouette tomography and a similar approach to linear tomography as formulated in (3). To do this, the proposed architecture was trained with training objective (6) with $\mathbf{y}_d = \mathbf{H}\mathbf{x}_d$. All training settings, including optimizer, batch size, learning rate, and length of training were identical to the model trained for silhouette tomography.

Quantitative results are given in Table 2. We again present metrics computed using both the floating point network output and the network output after thresholding at a value of 0.5. In this case, there were 74 slices in the test set for which the thresholded network output had an MSE of 0, giving an infinite PSNR. The PSNR of these slices was discarded when computing average PSNR on the test set.

The supervised model trained for linear tomography performs substantially better than the model trained for silhouette tomography across all metrics. This is unsurprising, because the ST problem is more ill-posed than the linear tomography problem. Moreover, we propose ST as a possible approach to tomography in settings where linear tomography is impractical. This experimental setup is ideal for linear tomography to succeed: the forward operator \mathbf{H} is known perfectly, the measurements \mathbf{y} are free of any additional noise, and the objects being imaged are completely uniform in material composition.

This comparison is valuable because it quantifies the size of the performance gap between linear and silhouette tomography under ideal conditions. Additionally, it demonstrates that the proposed neural network architecture is highly successful in the linear setting, which suggests that it is also a reasonable choice for silhouette tomography reconstruction.

Discussion and Conclusion

We introduced silhouette tomography (ST), a novel approach to X-ray CT that formulates the image reconstruction problem purely geometrically. Whereas both traditional model-based reconstruction and existing learning-based approaches to X-ray CT rely on an accurate model of the imaging system—either for reconstruction or to generate realistic training data—the ST approach requires only geometric information, i.e., the projection angles. Rather than reconstructing a 3D absorption map, ST reconstructs a binary map representing whether the object does or does not exist at each point in space. This type of reconstruction is particularly well-suited to nondestructive testing and security applications where object shapes are more important than density variations.

We mathematically formulated the ST problem and showed that when the problem has a solution, it always admits a unique solution that is maximal in norm. We then proposed a supervised deep learning approach for ST reconstruction, and demonstrated that this approach significantly outperforms the maximal solution across multiple reconstruction metrics. These results demonstrate that supervised learning is an effective method for ameliorating

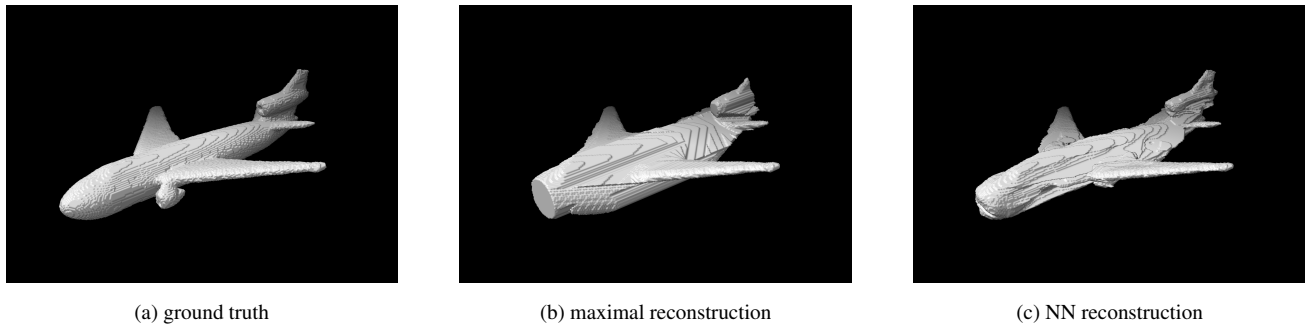


Figure 5: Visual comparison of maximal and supervised approaches for 3D object reconstruction.

the ill-posedness of the ST reconstruction problem, at least on our shape dataset.

Extensions to this work could involve adapting the network architecture to handle fully 3D reconstruction, performing experiments on real data, and developing automated methods to segment objects in X-ray projection images.

References

- [1] Wim van Aarle et al. “Fast and flexible X-ray tomography using the ASTRA toolbox”. In: *Opt. Express* 24.22 (Oct. 2016), pp. 25129–25147. DOI: 10.1364/OE.24.025129. URL: <https://opg.optica.org/oe/abstract.cfm?URI=oe-24-22-25129>.
- [2] R. N. Bracewell and A. C. Riddle. “Inversion of Fan-Beam Scans in Radio Astronomy”. en. In: *The Astrophysical Journal* 150 (Nov. 1967), p. 427. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/149346. (Visited on 07/20/2016).
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. *Large Scale GAN Training for High Fidelity Natural Image Synthesis*. 2019. arXiv: 1809.11096 [cs.LG].
- [4] Angel X. Chang et al. *ShapeNet: An Information-Rich 3D Model Repository*. Tech. rep. arXiv:1512.03012 [cs.GR]. Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [5] Dawson-Haggerty et al. *trimesh*. Version 3.2.0. Dec. 2019. URL: <https://trimesh.org/>.
- [6] Jingwei Huang, Hao Su, and Leonidas Guibas. *Robust Watertight Manifold Surface Generation Method for ShapeNet Models*. 2018. DOI: 10.48550/ARXIV.1802.01698. URL: <https://arxiv.org/abs/1802.01698>.
- [7] Sergey Ioffe and Christian Szegedy. “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift”. In: *Proceedings of the 32nd International Conference on Machine Learning*. Ed. by Francis Bach and David Blei. Vol. 37. Proceedings of Machine Learning Research. Lille, France: PMLR, July 2015, pp. 448–456. URL: <https://proceedings.mlr.press/v37/ioffe15.html>.
- [8] Ajinkya Kadu and Tristan van Leeuwen. “A Convex Formulation for Binary Tomography”. In: *IEEE Transactions on Computational Imaging* 6 (2020), pp. 1–11. DOI: 10.1109/tci.2019.2898333.
- [9] Hoyeon Lee, Jongha Lee, and Suengryong Cho. “View-interpolation of sparsely sampled sinogram using convolutional neural network”. In: vol. 10133. 2017, pp. 1013328–1013328–8. DOI: 10.1117/12.2254244. (Visited on 03/21/2017).
- [10] Zhipeng Li et al. “SUPER Learning: A Supervised-Unsupervised Framework for Low-Dose CT Image Reconstruction”. In: *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*. IEEE, Oct. 2019. DOI: 10.1109/iccvw.2019.00490.
- [11] E. Maire and P. J. Withers. “Quantitative X-ray tomography”. In: *International Materials Reviews* 59.1 (Jan. 2014), pp. 1–43. ISSN: 0950-6608. DOI: 10.1179/1743280413Y.0000000023. (Visited on 02/10/2016).
- [12] M. G. McGaffin and J. A. Fessler. “Alternating Dual Updates Algorithm for X-ray CT Reconstruction on the GPU”. In: *IEEE Transactions on Computational Imaging* 1.3 (Sept. 2015), pp. 186–199. ISSN: 2333-9403. DOI: 10.1109/TCI.2015.2479555.
- [13] Lars Mescheder et al. “Occupancy Networks: Learning 3D Reconstruction in Function Space”. In: *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. 2019.
- [14] Shanzhou Niu et al. “Sparse-view X-ray CT reconstruction via total generalized variation regularization”. In: *Phys. Med. Biol.* 59.12 (May 2014), pp. 2997–3017. DOI: 10.1088/0031-9155/59/12/2997.
- [15] Silvio Savarese. “Shape reconstruction from shadows and reflections”. en. PhD thesis. 2005. DOI: 10.7907/FH0V-3M10. URL: <https://resolver.caltech.edu/CaltechETD:etd-05242005-162056>.
- [16] Hongming Shan et al. “Competitive performance of a modularized deep neural network compared to commercial algorithms for low-dose CT image reconstruction”. In: *Nature Machine Intelligence* 1.6 (June 2019), pp. 269–276. DOI: 10.1038/s42256-019-0057-9.
- [17] *TorchMetrics - Measuring Reproducibility in PyTorch*. DOI: 10.21105/joss.0410. URL: <https://www.pytorchlightning.ai>.
- [18] László Varga, Péter Balázs, and Antal Nagy. “Projection Selection Dependency in Binary Tomography”. In: *Acta Cybernetica* 20.1 (2011), pp. 167–187. DOI: 10.14232/actacyb.20.1.2011.12.

- [19] Zhou Wang et al. “Image quality assessment: from error visibility to structural similarity”. In: *IEEE Transactions on Image Processing* 13.4 (2004), pp. 600–612. DOI: 10.1109/TIP.2003.819861.
- [20] Andrew Webb. *Introduction to Biomedical Imaging*. 1 edition. Hoboken, NJ: Wiley-IEEE Press, Dec. 2002. ISBN: 978-0-471-23766-2. DOI: 10.1118/1.1589017. URL: <http://eu.wiley.com/WileyCDA/WileyTitle/productCd-0471237663.html> (visited on 09/30/2015).
- [21] Stefan Weber et al. “A Benchmark Evaluation of Large-Scale Optimization Approaches to Binary Tomography”. In: *Discrete Geometry for Computer Imagery*. Springer Berlin Heidelberg, 2006, pp. 146–156. DOI: 10.1007/11907350_13.
- [22] Stefan Weber et al. “Binary Tomography by Iterating Linear Programs from Noisy Projections”. In: *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2004, pp. 38–51. DOI: 10.1007/978-3-540-30503-3_3.
- [23] Zenghui Wei et al. “A Joint Reconstruction and Segmentation Method for Limited-Angle X-Ray Tomography”. In: *IEEE Access* 6 (2018), pp. 7780–7791. DOI: 10.1109/access.2018.2800719.
- [24] Jay Whang et al. “Deblurring via Stochastic Refinement”. In: (2021). arXiv: 2112.02475 [cs.CV].
- [25] Jiajie Wu, Xianghua Xu, and Junyan Yang. “Object Detection and X-Ray Security Imaging: A Survey”. In: *IEEE Access* 11 (2023), pp. 45416–45441. DOI: 10.1109/ACCESS.2023.3273736.
- [26] Zhou Yu et al. “Fast model-based X-ray CT reconstruction using spatially nonhomogeneous ICD optimization”. In: *IEEE Transactions on Image Processing* 20.1 (Jan. 2011), pp. 161–175. ISSN: 1057-7149. DOI: 10.1109/TIP.2010.2058811.