A Self-Supervised Approach to Reconstruction in Sparse X-Ray Computed Tomography

Rey Mendoza[*1], **Minh Nguyen**[*1], **Judith Weng Zhu**[*1],
Vincent Dumont[2], **Talita Perciano**[2], **Juliane Mueller**[2,3], **Vidya Ganapati**[1,2]
[1] Swarthmore College, [2] Lawrence Berkeley National Laboratory,
[3] National Renewable Energy Laboratory

Abstract

Computed tomography has propelled scientific advances in fields from biology to materials science. This technology allows for the elucidation of 3-dimensional internal structure by the attenuation of x-rays through an object at different rotations relative to the beam. By imaging 2-dimensional projections, a 3-dimensional object can be reconstructed through a computational algorithm. Imaging at a greater number of rotation angles allows for improved reconstruction. However, taking more measurements increases the x-ray dose and may cause sample damage. Deep neural networks have been used to transform sparse 2-D projection measurements to a 3-D reconstruction by training on a dataset of known similar objects. However, obtaining high-quality object reconstructions for the training dataset requires high x-ray dose measurements that can destroy or alter the specimen before imaging is complete. This becomes a chicken-and-egg problem: high-quality reconstructions cannot be generated without deep learning, and the deep neural network cannot be learned without the reconstructions. This work develops and validates a self-supervised probabilistic deep learning technique, the physics-informed variational autoencoder, to solve this problem. A dataset consisting solely of sparse projection measurements from each object is used to jointly reconstruct all objects of the set. This approach has the potential to allow visualization of fragile samples with x-ray computed tomography. We release our code for reproducing our results at: https://github.com/vganapati/CT_PVAE.

1 Introduction & Related Work

1.1 Computed Tomography

The power of x-rays to penetrate through many materials has allowed for advances in scientific understanding. In particular, computed tomography has been used in fields such as biology [1], medicine [2], materials science [3], and geoscience [4] to elucidate the internal structure of samples. Computed tomography works by measuring the attenuation of an x-ray beam through an object at different rotations relative to the beam. From the projection images created by parallel beams at each rotation, a 3-dimensional object can be reconstructed with a computational algorithm. Imaging at many rotation angles allows for improved reconstruction, however, fragile samples can only be imaged under limited x-ray dose before damage and structural changes occur.

1.2 Reconstruction Algorithms

A 3-dimensional object can be reconstructed from projection images through direct algorithms such as filtered backpropagation or iterative algorithms [5, 6]. In recent years, algorithms using deep learning

---

[*]Equal contribution. Correspondence to vidyag@berkeley.edu.

have come into prominence [5, 7–12]. In these approaches, the parameters of the deep neural network are found by training on large datasets of known input-output pairs. The computational burden of training is high and a training dataset is required, which can be generated through high-quality experimental data and traditional algorithms. However, once trained, a relatively quick forward pass through the network can generate the reconstruction [5]. Deep learning has also shown the potential to improve upon conventional methods when there are a limited number of projection images, as a prior is learned from the data and embedded into the neural network parameters [5, 7].

1.3 Deep Learning without Ground Truth

While deep learning methods for computed tomography are promising, these approaches require a training dataset of reconstructions. High-quality reconstructions need measurements at many rotation angles, and this high radiation dose may destroy or alter fragile specimens during the data acquisition process. A chicken-and-egg problem presents itself: the deep learning algorithm needs reconstructions, and the reconstructions need the deep learning algorithm.

An approach to solve this problem involves using training datasets consisting only of sparse angle measurements for every object of the set. The intuition is that these sparse measurements taken together can provide enough information to infer a prior distribution of the objects. One method is to train generative adversarial networks to create high-quality reconstructions by rewarding reconstructions that have corresponding sparse measurements that lie in the training distribution of sparse measurements [13–17]. Another method is to train directly with low-dose measurements, penalizing reconstructions by the discrepancy between simulated low-dose and true low-dose measurements [18–21]. Finally, other works use multiple low-dose measurements of each object in training [22, 23], though this may be unsuitable for fragile samples that cannot withstand repeated measurement.

1.4 Probabilistic Formulation

In all deep learning methods without ground truth surveyed, the reconstruction is given as a point estimate. In this work, we aim to find the posterior probability distribution $P(O|M)$, where O is the object being reconstructed and $M = [M_1, M_2, ..., M_n]$ are the measurements, allowing the uncertainty of the prediction to be modeled. We aim to lower the total number of measurements n to minimize the x-ray dose. We assume that we have a set of m objects $\{O_1, O_2, ..., O_m\}$, sampled from some distribution $P(O)$, and we aim to reconstruct all objects in the set. For each of the m objects, we are allowed n measurements. Each set of measurements for an object j, $M_j = [M_{j1}, M_{j2}, ..., M_{jn}]$ is obtained with chosen rotation angles $p_j = [p_{j1}, p_{j2}, ..., p_{jn}]$. We assume that the forward model physics $P(M|O; p) = P(M|O)$ is known; in computed tomography the forward model is the Radon transform with Poisson noise. For every object O, we aim to find the posterior distribution $P(O|M) = \frac{P(M|O)P(O)}{P(M)}$. This work tackles the following problems that arise in finding the posterior: (1) construction of the prior $P(O)$ with no directly observed O, only measurements M on each object of the set, (2) calculating $P(O|M)$ in a tractable manner.

2 Methods

2.1 Physics-Informed Variational Autoencoder

We create a framework for posterior estimation that is inspired by the mathematics of the variational autoencoder [24, 25]. In a variational autoencoder, the goal is to learn how to generate new examples, sampled from the same underlying probability distribution as a training dataset of objects. To accomplish this task, a latent random variable z is created that describes the space on a lower-dimensional manifold. A deep neural network defines a function (the "decoder") from a sample of z to a probability distribution $P(O|z)$. The parameters of the deep neural network are optimized to maximize the probability of generating the objects of the training dataset.

In this work, we aim to find the posterior probability distribution $P(O|M)$, where O is the object being reconstructed and M are the projection measurements. In our case, we only have a dataset of noisy measurements M and no ground truth objects O. However, the forward model, $P(M|O)$, is

known. Thus, instead of maximizing the probability of generating O, we can maximize the probability of generating M, a formulation we call the "physics-informed variational autoencoder."
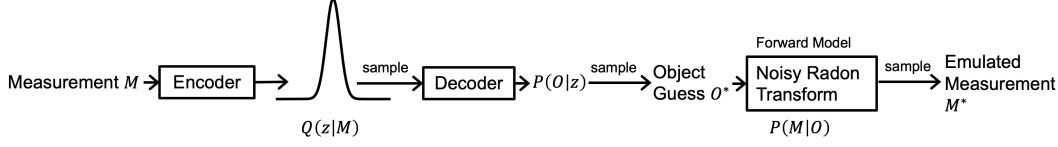

Figure 1: The developed physics-informed variational autoencoder.

We aim to maximize $P(M) = \int \int P(M|O)P(O|z)P(z)dOdz$. To compute this integral in a computationally tractable manner, we can approximate with sampled values. As in [24, 25], for most randomly sampled values of z and O, the probability $P(M|O, z)$ is close to zero, causing poor scaling of sampled estimates. Similar to variational autoencoders, our framework solves this problem by estimating the parameters of $P(z|M)$ by processing the measurements M using a function with trainable parameters (called the "encoder," also formulated as a deep neural network). The estimate of $P(z|M)$ is denoted $Q(z|M)$. The Kullback–Leibler divergence between the distributions is given by $D[Q(z|M)||P(z|M)] = E_{z \sim Q}[\log Q(z|M) - \log P(z|M)]$. We also have, by Bayes' Theorem, $\log P(z|M) = \log P(M|z) + \log P(z) - \log P(M)$. Combining the expressions yields:

$$\log P(M) - D[Q(z|M)||P(z|M)] = E_{z \sim Q}\left[\int P(M|O)P(O|z)dO\right] - D\left[\log Q(z|M)||\log P(z)\right].$$

The first term on the right side of this expression can be estimated with sampled values. As Kullback–Leibler divergence is always ≥ 0 and reaches 0 when $Q(z|M) = P(z|M)$, maximizing the right side (defined here as the $-$loss) during training causes $P(M)$ to be maximized while forcing $Q(z|M)$ towards $P(z|M)$. In contrast to a conventional variational autoencoder, we do *not* attempt to use this formulation to synthesize arbitrary objects O by sampling $P(z)$ directly. This framework only attempts reconstruction on the training examples themselves. After training, $P(O|M)$ for every object can be sampled by first sampling $Q(z|M)$ then $P(O|z)$; see Fig. 1. Crucially, unlike most data-driven approaches to reconstruction in computational imaging, the developed framework assumes that no ground truth dataset of objects O is available.

In this work, the encoder and decoder of Fig. 1 take the same basic shape as a U-Net [26]. The input to the encoder is an object reconstruction created with a standard algorithm from the sparse set of projections, using the TOMOPY Python package [27], as well as specification of the projection angles used. The skip connections to the decoder are parametrized as Gaussian distributions, and sampled to determine the latent variables z. The likelihood distribution $P(M|O)$ is given by the Radon transform with the addition of Poisson noise.

3   Results & Discussion

We prototype and validate the physics-informed variational autoencoder with synthetic datasets. Code to reproduce our results are at `https://github.com/vganapati/CT_PVAE`. We utilize 2-D objects and 1-D projections without loss of generality; a 3-D object in x-ray computed tomography is reconstructed as a series of 2-D axial slices.

3.1   Simple Toy Example

Our first dataset consists of two unique objects. Each object can be represented by 2×2 pixels, and projections can be taken at rotation angles of 0 and $\frac{\pi}{2}$ radians; see Fig. 2. The projection of both the objects at $\frac{\pi}{2}$ is the same, but the projections at 0 radians are different. We assume that $P(O = O_1) = P(O = O_2) = \frac{1}{2}$. We sample $P(O)$, and then take a noisy measurement $M = [M_1]$ at one rotation angle p (i.e. $n = 1$ and $P(p = 0) = P(p = \frac{\pi}{2}) = \frac{1}{2}$). We complete this sampling and measurement procedure $m = 1\,024$ times. From all the measurements, and no knowledge of the prior $P(O)$, we aim to determine the posterior $P(O|M)$ for each M by training the physics-informed variational autoencoder. Training on an NVIDIA GeForce
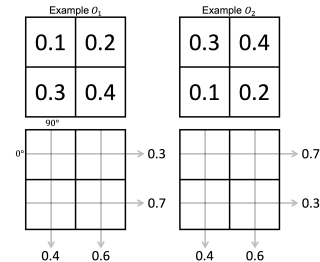

Figure 2: Toy dataset consisting of 2 object types.

3

(a) O_1 measured at $\frac{\pi}{2}$ radians.
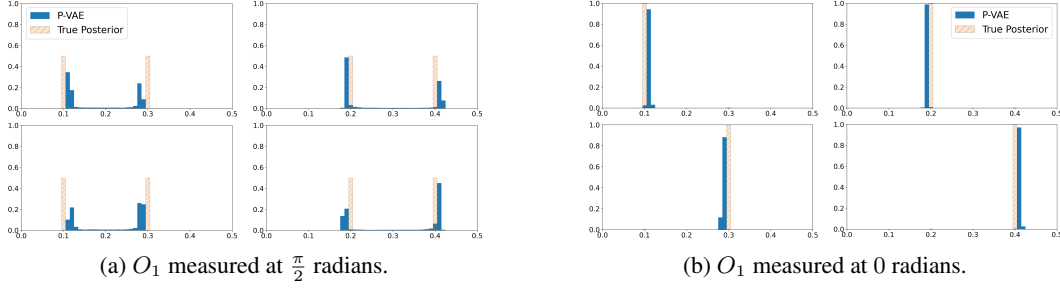
(b) O_1 measured at 0 radians.

Figure 3: Marginal posterior probabilities for each pixel from measurements on the toy dataset; measurements taken on object O_1 of the dataset. See Fig. 5 in the Appendix for posteriors on O_2.

GPU takes 47 minutes. Training parameters are available in the released documentation and code. The true posterior is compared with the posterior from 20 000 samples of the trained physics-informed variational autoencoder, showing good qualitative agreement; see Fig. 3. The marginal posterior for each of the four pixels is compared for four different example measurements. With measurements taken at the rotation angle $\frac{\pi}{2}$, there is approximately equal probability of the pixel values from the two object types. However, with measurements taken at 0 radians, the posterior is nearly a delta function at the true pixel value.

3.2 Synthetic Foam Dataset

We generate the second dataset of 1 000 foam images of 128×128 pixels through the Python package XDESIGN [28], and projections are taken at 180 equally-spaced angles. The image of the projections is known as the sinogram, and we emulate sparse sinograms by removing 160 of the projection angles either uniformly or randomly; see Fig. 6 in the Appendix for example images and corresponding sinograms of the dataset. We jointly reconstruct all images of this dataset by training the physics-informed variational autoencoder. Training on an NVIDIA Tesla V100 takes 2.2 hours. Training parameters are available in the released code and documentation.

Fig. 4 depicts example objects from the synthetic foam dataset, reconstructed with the Fourier grid reconstruction algorithm (`gridrec`) [29]. Reconstructions were performed with two other standard algorithms, the simultaneous algebraic reconstruction technique [27] and the Total Variation reconstruction technique [30]; however `gridrec` outperformed these techniques on image quality metrics of structural similarity (SSIM), mean-squared error (MSE), and peak signal-to-noise (PSNR). The reconstructions with `gridrec` on the sparse sinograms are compared with the reconstructions from the physics-informed variational autoencoder (P-VAE); see Fig. 4. We observe that the P-VAE improves over `gridrec` for the sparse sinograms, and there is a slight improvement by randomly selecting the angles in every example over uniform sampling; see Fig. 7 in the Appendix for results averaged over the entire 1 000 object dataset and aggregated over 10 independent trials.
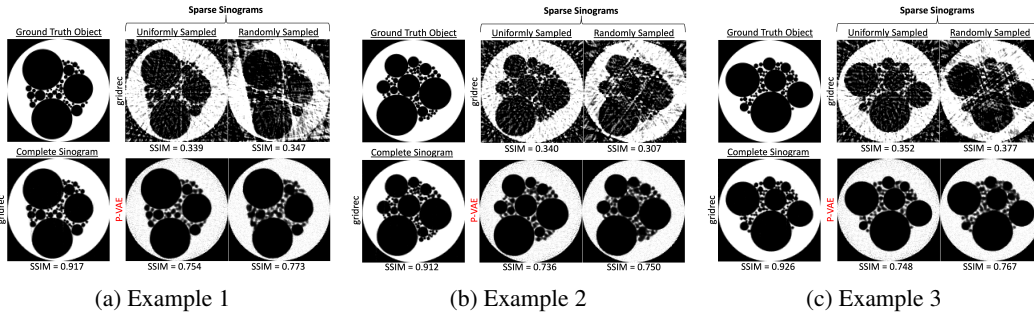

(a) Example 1   (b) Example 2   (c) Example 3

Figure 4: Foam dataset reconstruction results.

4 Conclusions & Limitations

We present a novel data-driven, self-supervised algorithm for image reconstruction in computed tomography and validate on synthetic datasets. Our method is probabilistic, allowing for quantification of uncertainty in applications, and self-supervised, requiring no ground truth or reference training examples. Though initial results are promising, our work is limited in that we only consider highly

synthetic data. Real experimental data may pose additional problems such as model mismatch that need to be considered in the construction of the likelihood $P(M|O)$. In experiments, thousands of rotation angles may be collected with measurements on the order of 100 GB per object [6]. An important direction for future work is high memory management techniques for training.

5 Broader Impact

Computed tomography uses x-rays to probe the 3-dimensional internal structure of objects, and is widely used in many fields such as medicine, biology, and materials science. However, x-rays can cause sample damage and deformation during the measurement process, limiting the use cases of computed tomography. Much effort has been put into sparse computed tomography, where a limited number of measurements are used to measure and reconstruct an object. This work proposes a novel algorithm for reconstruction in sparse computed tomography, with the potential positive broader impact of making visualization of radiation-damage prone objects possible. However, practitioners should be careful in applying these methods as there are no theoretical guarantees thus far. Application of this method could result in reconstruction artifacts that may have a different appearance than artifacts from conventional methods, evading detection.

References

[1] L. D. Wise, C. T. Winkelmann, B. Dogdas, and A. Bagchi, "Micro-computed tomography imaging and analysis in developmental biology and toxicology," *Birth Defects Research Part C: Embryo Today: Reviews*, vol. 99, no. 2, pp. 71–82, 2013. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/bdrc.21033

[2] G. D. Rubin, "Computed Tomography: Revolutionizing the Practice of Medicine for 40 Years," *Radiology*, vol. 273, no. 2S, pp. S45–S74, Nov. 2014. [Online]. Available: http://pubs.rsna.org/doi/10.1148/radiol.14141356

[3] S. Garcea, Y. Wang, and P. Withers, "X-ray computed tomography of polymer composites," *Composites Science and Technology*, vol. 156, pp. 305–319, Mar. 2018. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0266353817312460

[4] V. Cnudde and M. N. Boone, "High-resolution X-ray computed tomography in geosciences: A review of the current technology and applications," *Earth-Science Reviews*, vol. 123, pp. 1–17, Aug. 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S001282521300069X

[5] D. Pelt, J. Batenburg, and J. Sethian, "Improving tomographic reconstruction from limited data using Mixed-Scale Dense convolutional neural networks," *Journal of Imaging*, vol. 4, no. 11, pp. 128–128, Oct. 2018. [Online]. Available: https://ir.cwi.nl/pub/28279

[6] S. Marchesini, A. Trivedi, P. Enfedaque, T. Perciano, and D. Parkinson, "Sparse Matrix-Based HPC Tomography," in *Computational Science – ICCS 2020*, ser. Lecture Notes in Computer Science, V. V. Krzhizhanovskaya, G. Závodszky, M. H. Lees, J. J. Dongarra, P. M. A. Sloot, S. Brissos, and J. Teixeira, Eds. Cham: Springer International Publishing, 2020, pp. 248–261.

[7] D. Y. Parkinson, D. M. Pelt, T. Perciano, D. Ushizima, H. Krishnan, H. S. Barnard, A. A. MacDowell, and J. Sethian, "Machine learning for micro-tomography," in *Developments in X-Ray Tomography XI*, vol. 10391. International Society for Optics and Photonics, Sep. 2017, p. 103910J. [Online]. Available: https://doi.org/10.1117/12.2274731

[8] D. Ayyagari, N. Ramesh, D. Yatsenko, T. Tasdizen, and C. Atria, "Image reconstruction using priors from deep learning," in *Medical Imaging 2018: Image Processing*, vol. 10574. International Society for Optics and Photonics, Mar. 2018, p. 105740H. [Online]. Available: https://doi.org/10.1117/12.2293766

[9] Y. Huang, T. Würfl, K. Breininger, L. Liu, G. Lauritsch, and A. Maier, "Some Investigations on Robustness of Deep Learning in Limited Angle Tomography," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, ser. Lecture Notes in Computer Science, A. F. Frangi, J. A. Schnabel, C. Davatzikos, C. Alberola-López, and G. Fichtinger, Eds. Cham: Springer International Publishing, 2018, pp. 145–153.

[10] S. Bazrafkan, V. Van Nieuwenhove, J. Soons, J. De Beenhouwer, and J. Sijbers, "Deep Learning Based Computed Tomography Whys and Wherefores," *arXiv:1904.03908 [cs, eess]*, Apr. 2019. [Online]. Available: http://arxiv.org/abs/1904.03908

[11] L. Fu and B. D. Man, "A hierarchical approach to deep learning and its application to tomographic reconstruction," in *15th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, vol. 11072. International Society for Optics and Photonics, May 2019, p. 1107202. [Online]. Available: https://doi.org/10.1117/12.2534615

[12] J. He and J. Ma, "Radon inversion via deep learning," in *Medical Imaging 2019: Physics of Medical Imaging*, vol. 10948. International Society for Optics and Photonics, Mar. 2019, p. 1094810. [Online]. Available: https://doi.org/10.1117/12.2511643

[13] M. Kabkab, P. Samangouei, and R. Chellappa, "Task-Aware Compressed Sensing with Generative Adversarial Networks," *arXiv:1802.01284 [cs, stat]*, Feb. 2018. [Online]. Available: http://arxiv.org/abs/1802.01284

[14] A. Bora, E. Price, and A. G. Dimakis, "AmbientGAN: Generative models from lossy measurements," in *International Conference on Learning Representations*, 2018. [Online]. Available: https://openreview.net/forum?id=Hy7fDog0b

[15] S. Kuanar, V. Athitsos, D. Mahapatra, K. Rao, Z. Akhtar, and D. Dasgupta, "Low Dose Abdominal CT Image Reconstruction: An Unsupervised Learning Based Approach," in *2019 IEEE International Conference on Image Processing (ICIP)*. Taipei, Taiwan: IEEE, Sep. 2019, pp. 1351–1355. [Online]. Available: https://ieeexplore.ieee.org/document/8803037/

[16] E. K. Cole, F. Ong, S. S. Vasanawala, and J. M. Pauly, "Fast unsupervised mri reconstruction without fully-sampled ground truth data using generative adversarial networks," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, October 2021, pp. 3988–3997.

[17] S. Xu, S. Zeng, and J. Romberg, "Fast Compressive Sensing Recovery Using Generative Models with Structured Latent Variables," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Brighton, United Kingdom: IEEE, May 2019, pp. 2967–2971. [Online]. Available: https://ieeexplore.ieee.org/document/8683641/

[18] J. I. Tamir, S. X. Yu, and M. Lustig, "Unsupervised Deep Basis Pursuit: Learning inverse problems without ground-truth data," *arXiv:1910.13110 [eess]*, Feb. 2020, arXiv: 1910.13110. [Online]. Available: http://arxiv.org/abs/1910.13110

[19] W. Gan, C. Eldeniz, J. Liu, S. Chen, H. An, and U. S. Kamilov, "Image Reconstruction for MRI using Deep CNN Priors Trained without Groundtruth," in *2020 54th Asilomar Conference on Signals, Systems, and Computers*. Pacific Grove, CA, USA: IEEE, Nov. 2020, pp. 475–479. [Online]. Available: https://ieeexplore.ieee.org/document/9443403/

[20] M. Zhussip, S. Soltanayev, and S. Y. Chun, "Training deep learning based image denoisers from undersampled measurements without ground truth and without image prior," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[21] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. S. Kamilov, "RARE: Image Reconstruction Using Deep Priors Learned Without Groundtruth," *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1088–1099, Oct. 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9103213/

[22] Z. Xia and A. Chakrabarti, "Training Image Estimators without Image Ground-Truth," *arXiv:1906.05775 [cs]*, Oct. 2019, arXiv: 1906.05775. [Online]. Available: http://arxiv.org/abs/1906.05775

[23] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, "Noise2Noise: Learning Image Restoration without Clean Data," *arXiv:1803.04189 [cs, stat]*, Oct. 2018, arXiv: 1803.04189. [Online]. Available: http://arxiv.org/abs/1803.04189

[24] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv:1312.6114 [stats]*, 2014.

[25] C. Doersch, "Tutorial on variational autoencoders," *arXiv:1606.05908 [stats]*, 2021.

[26] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *arXiv:1505.04597 [cs]*, May 2015. [Online]. Available: http://arxiv.org/abs/1505.04597

[27] D. Gürsoy, F. De Carlo, X. Xiao, and C. Jacobsen, "TomoPy: a framework for the analysis of synchrotron tomographic data," *Journal of Synchrotron Radiation*, vol. 21, no. 5, pp. 1188–1193, Sep. 2014. [Online]. Available: http://scripts.iucr.org/cgi-bin/paper?S1600577514013939

[28] "XDesign documentation." [Online]. Available: https://xdesign.readthedocs.io/en/stable/

[29] B. A. Dowd, G. H. Campbell, R. B. Marr, V. V. Nagarkar, S. V. Tipnis, L. Axe, and D. P. Siddons, "Developments in synchrotron x-ray computed microtomography at the National Synchrotron Light Source," in *SPIE's International Symposium on Optical Science, Engineering, and Instrumentation*, Denver, CO, Sep. 1999, pp. 224–236. [Online]. Available: https://doi.org/10.1117/12.363725

[30] A. Chambolle and T. Pock, "A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging," *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 120–145, May 2011. [Online]. Available: http://link.springer.com/10.1007/s10851-010-0251-1

Checklist

1. For all authors...

    (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]

    (b) Did you describe the limitations of your work? [Yes] [See Section 4.]

    (c) Did you discuss any potential negative societal impacts of your work? [Yes] [See Section 5.]

    (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...

    (a) Did you state the full set of assumptions of all theoretical results? [N/A]

    (b) Did you include complete proofs of all theoretical results? [N/A]

3. If you ran experiments...

    (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] [We include the code and instructions to reproduce the results here:
    `https://github.com/vganapati/CT_PVAE.`]

    (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] [Training details are available in the code documentation at `https://github.com/vganapati/CT_PVAE.`]

    (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] [See Fig. 7 in the Appendix.]

    (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] [See Section 3.]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

    (a) If your work uses existing assets, did you cite the creators? [N/A]

    (b) Did you mention the license of the assets? [N/A]

    (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] [We release code for our developed framework.]

    (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

    (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

    (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

    (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

    (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Appendix

(a) O_2 measured at $\frac{\pi}{2}$ radians.

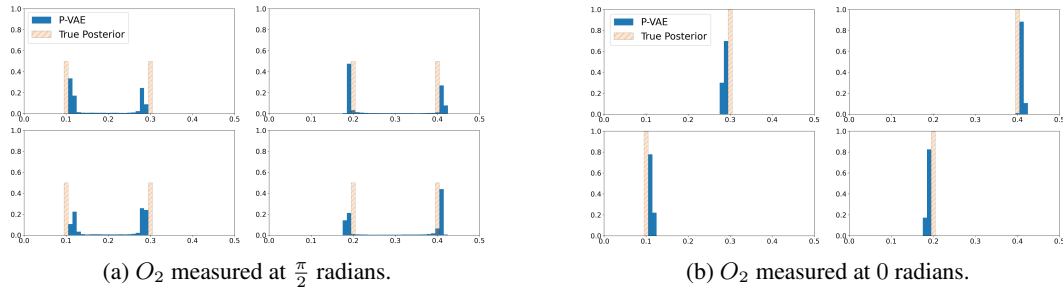(b) O_2 measured at 0 radians.

Figure 5: Marginal posterior probabilities for each pixel from measurements on the toy dataset; measurements taken on object O_2 of the dataset.
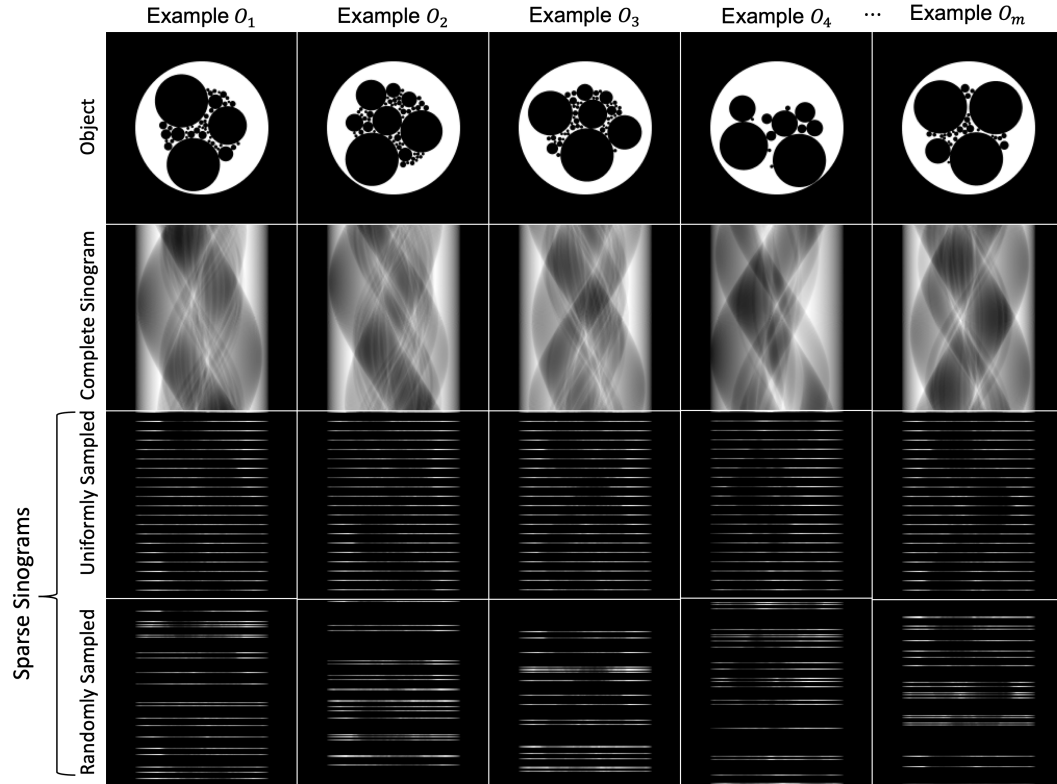

Figure 6: Object examples generated for the foam image dataset. Complete sinograms are taken at 180 equally spaced rotation angles. Uniformly sparse sinograms are taken at 20 equally spaced angles and randomly sampled sparse sinograms are taken at 20 randomly chosen angles; these angles are re-sampled for each object of the dataset.
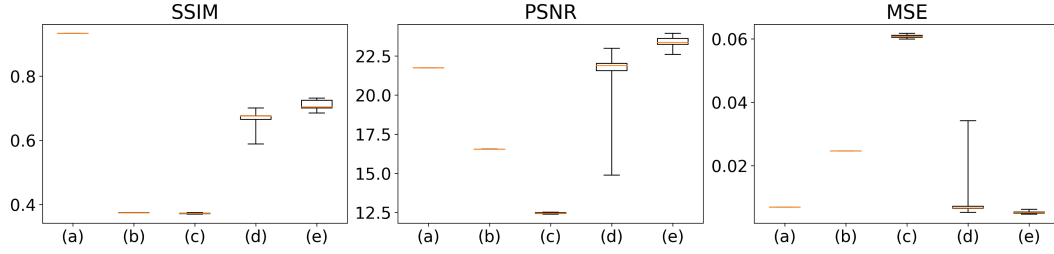
Figure 7: Image quality metrics of structural similarity (SSIM), peak signal-to-noise (PSNR), and mean-squared error (MSE) for reconstruction algorithms applied to the synthetic foam dataset, averaged over the entire dataset for 10 independent optimization trials. Reconstructions from (a) gridrec with the complete sinogram of 180 rotation angles, (b) gridrec with the sparse sinogram of 20 angles, uniformly spaced, (c) gridrec with the sparse sinogram of 20 angles, randomly spaced, with each example having different angles chosen, (d) the physics-informed variational autoencoder with the sparse sinogram of 20 angles, uniformly spaced, and (e) the physics-informed variational autoencoder with the sparse sinogram of 20 angles, randomly spaced.