
Generative Tomography Reconstruction

Matteo Ronchetti
University of Pisa
mtronchetti@gmail.com

Davide Bacciu
University of Pisa
bacciu@di.unipi.it

Abstract

We propose an end-to-end differentiable architecture for tomography reconstruction that directly maps a noisy sinogram into a denoised reconstruction. Compared to existing approaches our end-to-end architecture produces more accurate reconstructions while using less parameters and time. We also propose a generative model that, given a noisy sinogram, can sample realistic reconstructions. This generative model can be used as prior inside an iterative process that, by taking into consideration the physical model, can reduce artifacts and errors in the reconstructions.

1 Introduction

This work tackles a prototypical linear inverse problem: reconstruction of Computed Tomography (CT) images in the presence of large amount of noise. The goal of CT reconstruction is to infer an image x , given a noisy sinogram $y_\delta \sim n(Ax)$, where A is the discrete Radon transform [3] and $n(\cdot)$ is the noise distribution. Solving such a reconstruction problem is difficult due to the ill-conditioning of Radon transforms. It is also a problem of high practical relevance and impact, given the widespread use of CT algorithms in applications such as medical imaging [12, 22, 2], luggage scanning [11] and damage detection in concrete structures [7].

This work is centred around two original contributions. First, we design an end-to-end differentiable architecture that directly maps a noisy sinogram to a denoised reconstruction. By combining reconstruction and denoising functionalities into a single model, we are able to improve reconstruction quality while reducing inference time. Existing ML approaches to noisy CT reconstruction [8, 21, 10] use a convolutional neural network (CNN) to denoise a reconstruction computed by a non-adaptive algorithm. Contrary to that, we apply the learned reconstruction and denoising process directly to the sinogram. This allows the model to reduce noise before it is amplified by the reconstruction process. The second contribution leverages the differentiable architecture to propose a generative model that can sample reconstructions which are realistic and compatible with sensor readings. The reconstructions produced by the proposed generative model can be iteratively improved by taking into consideration the physical model of sinogram generation. This iterative algorithm can reduce reconstruction errors and biases, guaranteeing that the produced solution is compatible with what is measured by the sensors. GAN-based approaches in literature for solving linear inverse problems [20, 4, 15] compute reconstructions iteratively by improving an initial random guess. In contrast to that, our approach can generate good solutions even without iterating, making the iterative improvement process much faster. As a final note, we point out the generality of our approaches, which allow a straightforward application to other linear inverse problems, such as limited/sparse angle tomography, image super-resolution and image deblurring.

2 Sinogram-Based Tomography Reconstruction

Before delving into the details of the two contributions, we briefly summarize the noise model used to simulate sensor readings. We take into consideration shot and electronic noise, and quantization

errors as follows

$$z \sim \text{Pois}(\exp(s - y)) + N(0, \epsilon)$$

$$r = \text{clamp}(\text{round}(z/k), 0, 2^b - 1)$$

where $\exp(s)$ is the X-ray intensity, b is the number of bits used by the detector and k is a scaling parameter. We assume that, given readings r (that are integer values), sinograms are distributed as $y|r \sim \mathcal{N}(\mu(r), \text{diag}(\sigma^2(r)))$, where $\mu(\cdot)$ and $\sigma(\cdot)$ can be modelled by CNNs. We train $\mu(\cdot)$ and $\sigma(\cdot)$ networks to minimize the negative log-likelihood (refer to Appendix A for more details).

2.1 Reconstruction Model

Inspired by filtered back-projection (FBP) [9], which reconstructs x as $\tilde{x} = A^T f(y)$, where $f(\cdot)$ is a fixed linear filtration, we design our reconstruction model as $g(r) = g_2(A^T g_1(r))$ where both g_1 and g_2 are trainable neural networks. The Radon transform is ill-conditioned and FBP is well-known to greatly amplify noise. Our rationale is that by combining reconstruction and denoising, we allow the model to mitigate the reconstruction errors caused by the ill-conditioning of the Radon transform. Furthermore, since the output of $g_1(r)$ is not limited to be a single channel, it allows richer information flow from the sinogram to the image.

Briefly, our reconstruction model is composed of two U-nets [19] connected by Radon backprojection (multiplication by A^T). The first one (g_1) processes the sensor readings and outputs a multi-channel sinogram, each channel is transformed into an image using Radon backprojection and fed to the second network. The first layer of g_2 concatenates to its input (of shape (ch, 256, 256)) a 4 dimensional positional embedding (of shape (4, 256, 256)) which is used by the network to encode pixel coordinates. The down-sampling blocks of the U-net are composed of a strided convolution followed by a sequence of residual blocks [6]. Each U-net up-sampling block takes the output of the previous block, double its resolution using bilinear up-sampling, and concatenates it with the skip connection coming from the down-sampling block at the corresponding resolution. Following, the concatenation channels are weighted using channel attention [14] and inputted to a 1×1 convolution that shrinks their number. Finally they are processed by a sequence of residual blocks. A detailed account of the architecture is provided in Appendix B.

The key point in our convolutional scheme is that the Radon backprojection between the two U-nets is differentiable. Therefore the model discussed above can be trained end-to-end by jointly optimizing g_1 and g_2 to minimize an image reconstructions loss.

2.2 Generative Model

We design an end-to-end generator model that, given noisy sensor readings r , can sample different possible reconstructions. We show how, by combining it with a discriminator and the noise model, it can be used to iteratively improve reconstructions.

We build on the Wasserstein GAN (WGAN) [1] framework and train a generator G to minimize

$$\mathbb{E}_{r,z} \left[\|AG(r, z) - \mu(r)\|_{\sigma(r)}^2 + \lambda D(G(r, z)) \right] \quad \text{where} \quad \|v\|_{\sigma(r)}^2 \stackrel{\text{def}}{=} \sum_i \left(\frac{v_i}{\sigma_i(r)} \right)^2 \quad (1)$$

where z is uniformly distributed over the sphere and λ is a fixed regularization parameter. Simultaneously, we train a discriminator D to minimize

$$\mathbb{E}_x[D(x)] - \mathbb{E}_{r,z}[D(G(r, z))] \quad \text{subject to} \quad \|D\|_L \leq c$$

where the Lipschitz-norm constraint on D is enforced using Spectral Normalization [16].

The generator $G(r, z)$ is an end-to-end model following the architecture described in Section 2.1. To make the computation of $\nabla_z G(r, z)$ efficient, we feed z to the third last block of the g_2 network, so that back-propagating the gradient to z only requires to go through few layers. The structure of this modified U-net block is depicted in Figure 1 to highlight the additional input from z . The first layer of this block concatenates its inputs (by applying bilinear up-sampling to the ones coming from the previous block) and applies a 3×3 convolution followed by a PReLU activation [5]. The effect of z on the output should be modulated depending on the amount of noise and uncertainty in the reconstruction. Therefore the last 32 channels produced by the first layer of the block are fed to an




Figure 1: Structure of the modified U-net block used to insert z into the generator.

g_1	g_2	SSIM	Parameters	FPS
FBP	L	76.6%	$6.30 \cdot 10^6$	211
FBP	XL	76.7%	$9.37 \cdot 10^6$	178
FBP	XXL	76.9%	$1.56 \cdot 10^7$	121
XXS	S-64	76.7%	$2.63 \cdot 10^6$	239
XS	M	76.9%	$4.12 \cdot 10^6$	187
S	L	77.4%	$6.93 \cdot 10^6$	152

Figure 2: Comparison of denoising results using FBP based models (top) and end-to-end models (bottom).

exponential activation and used to scale the elements of z (bottom part in Figure 1). The remaining channels are then concatenated with the scaled version of z and fed to a sequence of residual blocks. The parameter z has shape $(32, 16, 16)$ and it is up-scaled to $(32, 64, 64)$ using bilinear interpolation.

If we train the generator using the loss described in (1) it learns to ignore z . To avoid this problem, we regularize the output of G by pushing the distribution of $AG(r, z)$ towards $\mathcal{N}(\mu(r), \text{diag}(\sigma^2(r)))$. During training we generate two different reconstructions for each noisy reading. Then, we fit a normal distribution with diagonal covariance on the two corresponding sinograms and compare it against $\mathcal{N}(\mu(r), \text{diag}(\sigma^2(r)))$ using the Kullback–Leibler divergence

$$y_1 = AG(r, z_1) \quad y_2 = AG(r, z_2)$$

$$\left\| \frac{y_1 + y_2}{2} - \mu(r) \right\|_{\sigma(r)}^2 + \sum_i \left[2 \log \left(\frac{\sigma(r)}{\sigma_p} \right) + \left(\frac{\sigma_p}{\sigma(r)} \right)^2 \right] \text{ where } \sigma_p^2 = \frac{(y_1 - y_2)^2}{2}.$$

Using this loss in place of the first term of (1) solves the mode collapse problem without reducing training stability.

As a final note, the reconstruction produced by our generator can be iteratively improved using projected gradient to minimize $\|AG(r, z) - \mu(r)\|_{\sigma(r)}^2 + \lambda D(G(r, z))$ subject to $\|z\| = 1$.

3 Experimental Assessment

We use images from the DeepLesion dataset¹, cropped to guarantee that pixels outside of the inscribed circle are zero and rescaled to 256×256 (see Appendix D for more details about preprocessing). We split train/validation/test by patient identifier and use multiple signal intensities (therefore noise levels) during training and testing. No augmentation is used, sensor readings are not fixed but randomly simulated before each training iteration.

We compare the reconstruction quality obtained by our end-to-end approach and an FBP model, also varying the dimensions of the U-net components g_1 and g_2 . We use conventional names (XXS, XS, S, S-64, M, L, XL, XXL) to refer to increasing U-net dimensions (refer to Appendix B for further details). All the models are trained for 10 epochs using the RAdam [13] optimizer, batch size of 16 and mixed precision training. The learning rate is exponentially increased from 0 to $3 \cdot 10^{-4}$ during the first 5000 batches, then halved every 80000 batches.

We measure the quality of reconstructions using structured similarity (SSIM) and compute inference speed as the number of frames per second (FPS) obtained with the largest possible batch size on a Tesla V100 GPU. Results are shown in Table 2, from which it can be noticed that end-to-end models achieve a more accurate reconstruction than FBP based models at comparable number of FPS. End-to-end models scale better, a +0.7% improvement in SSIM requires to increase the number of parameters 2.6 \times , compared to FBP models where an improvement of +0.3% costs 2.5 \times more parameters. Furthermore end-to-end models are much more parameter-efficient, given that their largest model is only slightly larger than the smallest FBP model.

¹<https://nihcc.app.box.com/v/DeepLesion>




Figure 3: Reconstruction results using FBP on the estimated sinogram $\mu(r)$ compared with 4 samples generated by the proposed GAN. Notice that, as signal increases, GAN reconstructions become more similar to one another.




Figure 4: Progress of our iterative algorithm on two candidate reconstructions

Lower signal intensities (therefore higher noise) are used to simulate readings for the training of our generative model. Figure 3 depicts reconstructions computed by FBP and sampled from our GAN. Notice that GAN reconstructions look realistic even with very high noise values and, as the noise decreases, converge to the correct reconstruction.

Figure 4 depicts the progress of the projected gradient algorithm on two candidate reconstructions using a learning rate of 10^{-4} . Notice in particular how the reconstruction of the spinal cord is improved by the iterative algorithm.

4 Conclusions and Future Work

We proposed an end-to-end architecture for tomography reconstruction characterized by higher quality and performance than FBP-based denoising models. The key point of our approach is the incorporation of the Radon transform as a differentiable component of the neural architecture, allowing to build the reconstruction directly from the sinogram. The end-to-end architecture can be used to design generative models capable of sampling multiple possible reconstructions. Also, we show that by taking into consideration the physical model, it is possible to iteratively improve the reconstructions computed by the generative model. Our work can be easily extended to sparse angle and limited angle tomography or to other linear inverse problems.

References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 214–223. JMLR.org, 2017.

- [2] W Baumeister. Electron tomography of molecules and cells. *Trends in Cell Biology*, 9(2):81–85, February 1999.
- [3] Gregory Beylkin. Discrete radon transform. *IEEE transactions on acoustics, speech, and signal processing*, 35(2):162–172, 1987.
- [4] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G. Dimakis. Compressed sensing using generative models. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 537–546, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [7] Kaarlo A. Heiskanen, Hong C. Rhim, and Paulo J.M. Monteiro. Computer simulations of limited angle tomography of reinforced concrete. *Cement and Concrete Research*, 21(4):625–634, July 1991.
- [8] Allard Adriaan Hendriksen, Daniel Maria Pelt, and Kees Joost Batenburg. Noise2inverse: Self-supervised deep convolutional denoising for tomography. *IEEE Transactions on Computational Imaging*, 2020.
- [9] A. C. Kak and Malcolm Slaney. *Principles of Computerized Tomographic Imaging*. IEEE Press, 1998.
- [10] Eunhee Kang, Junhong Min, and Jong Chul Ye. A deep convolutional neural network using directional wavelets for low-dose X-ray CT reconstruction. *Medical Physics*, 44(10):e360–e375, October 2017.
- [11] Seemeen Karimi, Harry Martz, and Pamela Cosman. Metal artifact reduction for CT-based luggage screening. *Journal of X-Ray Science and Technology*, 23(4):435–451, July 2015.
- [12] V Kolehmainen, S Siltanen, S J rvenp, J P Kaipio, P Koistinen, M Lassas, J Pirttil, and E Somersalo. Statistical inversion for medical x-ray tomography with few radiographs: II. Application to dental radiology. *Physics in Medicine and Biology*, 48(10):1465–1490, May 2003.
- [13] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
- [14] Yue Lu, Zhuqing Jiang, Guodong Ju, Liangheng Shen, and Aidong Men. Recursive multi-stage upscaling network with discriminative fusion for super-resolution. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, pages 574–579. IEEE, 2019.
- [15] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2445, 2020.
- [16] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- [17] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy,

- Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* 32, pages 8024–8035. Curran Associates, Inc., 2019.
- [18] Matteo Ronchetti. Torchradon: Fast differentiable routines for computed tomography. *arXiv preprint arXiv:2009.14788*, 2020.
- [19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [20] Viraj Shah and Chinmay Hegde. Solving linear inverse problems using gan priors: An algorithm with provable guarantees. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 4609–4613. IEEE, 2018.
- [21] Dufan Wu, Kyungsang Kim, Georges El Fakhri, and Quanzheng Li. A cascaded convolutional neural network for x-ray low-dose ct image denoising. *arXiv preprint arXiv:1705.04267*, 2017.
- [22] Yiheng Zhang, Heang-Ping Chan, Berkman Sahiner, Jun Wei, Mitchell M. Goodsitt, Lubomir M. Hadjiiski, Jun Ge, and Chuan Zhou. A comparative study of limited-angle cone-beam reconstruction methods for breast tomosynthesis: Limited-angle cone-beam reconstruction for breast tomosynthesis. *Medical Physics*, 33(10):3781–3795, September 2006.

A Noise Model

Without taking into consideration noise the ideal sensor readings would be $\exp(s - y)$, where $\exp(s)$ is the intensity of the X-ray emitted by the machine. In practice due to shot noise and electronic noise readings can be modelled as

$$z \sim \text{Pois}(\exp(s - y)) + \mathcal{N}(0, \epsilon).$$

In this setting noise depends on the signal intensity (the stronger the signal the weaker the noise) and on the variance of the normal distribution that emulates electronic noise. Notice that, because of the Poisson distribution, noise will not be uniform across the sinogram and will be higher where y is large.

In practice the analog value at each pixel of the sensor is converted into a digital integer value, this process creates quantization error. If readings are stored as b bits values then this conversion process can be modelled as

$$r = \text{clamp}(\text{round}(z/k), 0, 2^b - 1)$$

where k is a normalization value that depends on the scale of the measured signal.

Given sensor readings r (a matrix of integers in range $[0, 2^b)$) we would like to estimate the posterior probability over noiseless sinograms $\pi(y|r)$. We make the simplifying assumption of $y|r \sim \mathcal{N}(\mu(r), \text{diag}(\sigma^2(r)))$, where $\mu(\cdot)$ and $\sigma(\cdot)$ are CNNs. Because both the mean and variance are modelled by CNNs intra-pixel relationships can be modelled.

We first train $\mu(\cdot)$ to minimize

$$\mathbb{E}[\|y - \mu(r)\|_2^2]$$

and, after $\mu(\cdot)$ is trained, we use it to train $\sigma(\cdot)$ to minimize the negative log likelihood

$$\mathbb{E} \left[\frac{(y - \mu(r))^2}{\sigma^2(r)} + 2 \log(\sigma(r)) \right].$$

Finally we use knowledge distillation to combine $\mu(\cdot)$ and $\sigma(\cdot)$ into a single CNN.

B Model Architecture

The architecture of the end-to-end reconstruction model is depicted in Figure 5. Both g_1 and g_2 are instances of the U-net architecture [19]. Blocks that have more than one input, first up-sample each input to the largest resolution using bilinear up-sampling, then concatenate all the inputs on the




Figure 5: Architecture of the end-to-end model used for simultaneous reconstruction and denoising.

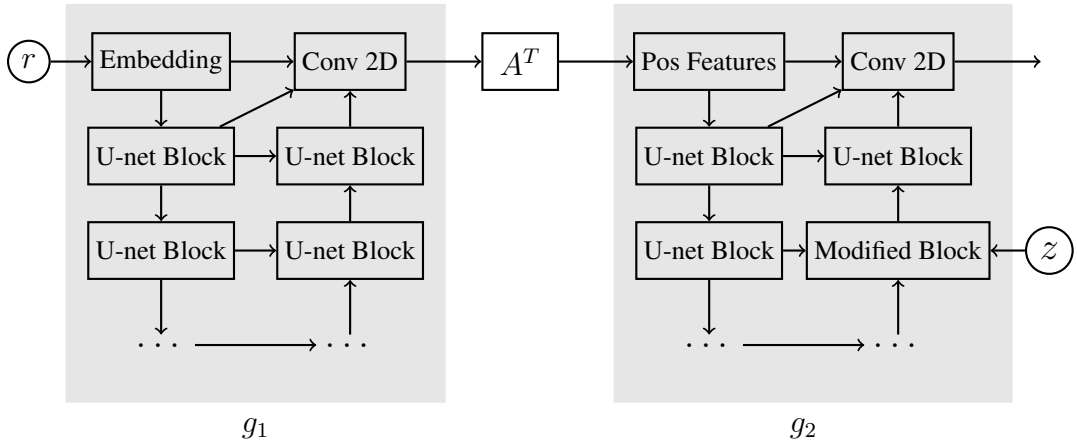


Figure 6: Architecture of the generator.

channel dimension.

The architecture of the proposed generator is shown in Figure 6 and is analogous to the end-to-end reconstruction model. The only difference is the modified block that is used to receive z (bottom right of Figure 6) which is described in Section 2.2.

The down-sampling blocks of the U-net are composed of a strided convolution followed by a sequence of residual blocks [6]. Each U-net up-sampling block takes the output of the previous block, double its resolution using bilinear up-sampling, and concatenates it with the skip connection coming from the down-sampling block at the corresponding resolution. Following, the concatenation channels are weighted using channel attention [14] and inputted to a 1×1 convolution that shrinks their number. Finally they are processed by a sequence of residual blocks.

We use conventional names (XXS, XS, S, S-64, M, L, XL, XXL) to refer to increasing U-net dimensions, the only difference between S and S-64 is the number of channels. The number of channels used by the residual blocks is doubled every two strided convolutions. The first block of XXS, XS, S uses 32 channels, while the first block of S-64, M, L, XL, XXL uses 64 channels. The number of residual blocks used inside each U-net block is shown in the following table:

In end-to-end models the latest convolution of g_1 reduces the number of channels before Radon backprojection. The smallest end-to-end model (which uses a XXS U-net for g_1 and a S-64 U-net for g_2) uses 16 channels for Radon backprojection, while the other two end-to-end models use 32 channels.

Resolution	XXS	XS	S	M	L	XL	XXL
256	1	1	1	2	2	3	5
128	1	1	1	1	2	3	5
64	2	1	1	1	2	3	5
32	-	2	2	1	2	3	5
16	-	-	2	1	2	3	5
8	-	-	-	2	3	5	9
16	-	-	-	2	3	4	6
32	-	-	2	2	3	4	6
64	-	2	2	2	3	4	6
128	2	2	2	2	3	4	6
256	1	2	2	2	4	4	6

C Motivation for the Generator Loss

A standard approach for solving inverse problems is to use regularization, this amounts to solving

$$\arg \min_x \|Ax - y\|_2^2 + \lambda D(x)$$

where $D(x)$ is a penalty term that promotes desired properties of x and λ is the regularization parameter.

First, we modify this formulation to take into consideration our sinogram estimation:

$$\arg \min_x \|Ax - \mu(r)\|_{\sigma(r)}^2 + \lambda D(x)$$

where

$$\|Ax - \mu(r)\|_{\sigma(r)}^2 \stackrel{\text{def}}{=} \sum \frac{(y - \mu(r))^2}{\sigma^2(r)}$$

With this formulation we are taking into consideration the non-uniform confidence in the estimated sinogram values. There is still the problem that we are directly optimizing over images, and we still have to define $D(x)$.

Instead of optimizing x we introduce a generator $G(r, z)$ and optimize over the free parameter z that lies on the $m + 1$ dimensional sphere:

$$\arg \min_{z \in S^m} \|AG(r, z) - \mu(r)\|_{\sigma(r)}^2 + \lambda D(G(r, z)).$$

D Dataset Preprocessing

We use the DeepLesion dataset (<https://nihcc.app.box.com/v/DeepLesion>) which is a large dataset of annotated tomographic images stored with 16 bit precision. We only use the images and don't consider the annotations for training our model. Images are preprocessed as follows:

1. Subtract 32768 from the pixel intensity to obtain the original Hounsfield unit (HU) values. HU is a measure of radiodensity where air has density -1000HU and water 0HU .
2. Create a binary version of the image by thresholding pixel with value larger than -9050HU (threshold is a bit higher than the density of air).
3. Find the smallest circle that contains all the pixels of the binarized image and to crop the corresponding square region from the image.
4. If the cropped region is smaller than 256×256 discard it, otherwise rescale to 256×256 and add to the dataset.

E Implementation Details

Training the model requires computation and differentiation of Radon forward and backward projections (corresponding respectively to products with A and A^T). These operations are linear and therefore differentiable, we use the implementation offered by the TorchRadon library [18] which is integrated with PyTorch [17] and allows the backpropagation of gradients.

F Additional Generated Reconstructions

The following figures contain additional images generated by our model using all the tested signal levels.

