

# Sparse-View Cone Beam CT Reconstruction Using Data-Consistent Supervised and Adversarial Learning From Scarce Training Data

Anish Lahiri<sup>1b</sup>, *Member, IEEE*, Gabriel Maliakal<sup>1b</sup>, Marc L. Klasky, Jeffrey A. Fessler<sup>1b</sup>, *Fellow, IEEE*, and Saiprasad Ravishankar<sup>1b</sup>, *Senior Member, IEEE*

**Abstract**—Reconstruction of CT images from a limited set of projections through an object is important in several applications ranging from medical imaging to industrial settings. As the number of available projections decreases, traditional reconstruction techniques such as the FDK algorithm and model-based iterative reconstruction methods perform poorly. Recently, data-driven methods such as deep learning-based reconstruction have garnered a lot of attention in applications because they yield better performance when enough training data is available. However, even these methods have their limitations when there is a scarcity of available training data. This work focuses on image reconstruction in such settings, i.e., when both the number of available CT projections and the training data is extremely limited. We adopt a sequential reconstruction approach over several stages using an adversarially trained shallow network for ‘destreaking’ followed by a data-consistency update in each stage. To deal with the challenge of limited data, we use image subvolumes to train our method, and patch aggregation during testing. To deal with the computational challenge of learning on 3D datasets for 3D reconstruction, we use a hybrid 3D-to-2D mapping network for the ‘destreaking’ part. Comparisons to other methods over several test examples indicate that the proposed method has much potential, when both the number of projections and available training data are highly limited.

**Index Terms**—Sparse-views, computed tomography, machine learning, deep learning, image reconstruction.

## I. INTRODUCTION

COMPUTED Tomography (CT) is an important imaging modality across applications in medicine, industry, science

Manuscript received 26 January 2022; revised 13 July 2022 and 6 October 2022; accepted 7 October 2022. Date of publication 6 January 2023; date of current version 8 February 2023. This work was supported in part by the National Nuclear Security Administration and in part by the the Los Alamos National Laboratory. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Tobias Lasser. (*Corresponding author: Gabriel Maliakal.*)

Anish Lahiri and Jeffrey A. Fessler are with the Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, MI 48014 USA (e-mail: anishl@umich.edu; fessler@umich.edu).

Marc L. Klasky is with the Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: mklasky@lanl.gov).

Gabriel Maliakal is with the Department of Computational Mathematics, Science and Engineering, Michigan State University, East Lansing, MI 48824 USA, and also with the Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: maliakal@msu.edu).

Saiprasad Ravishankar is with the Department of Computational Mathematics, Science and Engineering, and the Department of Biomedical Engineering, Michigan State University, East Lansing, MI 48824 USA (e-mail: ravisha3@msu.edu).

Digital Object Identifier 10.1109/TCI.2022.3225680

and security. In this work, we develop an iterative machine learning-based approach for 3D cone beam CT reconstruction from very limited measurements or projections, and using limited training data. In the following, we first review some background in limited-view CT reconstruction before highlighting the contributions of this work.

## A. Background

Cone Beam CT (CBCT) is a CT-based technique that allows for three-dimensional imaging of an object using X-rays diverging from a source. In CBCT, an entire 3D image volume is reconstructed from a set of 2D projections through the corresponding object. These projections/measurements are obtained at different angles or ‘views’ around the object, and are collectively dubbed a *sinogram*. There are several approaches for the inverse problem of obtaining an image from these measurements. A classical method for this task is the analytical Feldkamp-Davis-Kress (FDK) algorithm [1]. More sophisticated methods for 2D or 3D reconstruction involve model-based reconstruction using iterative algorithms [2], [3], [4], and data-driven algorithms [5], [6].

Model-based image reconstruction (MBIR) or statistical image reconstruction (SIR) methods exploit sophisticated models for the physics of the imaging system and models for sensor and noise statistics as well as for the underlying object. These methods iteratively optimize for the underlying image based on the system forward model, measurement statistical model, and assumed prior for the underlying object [7], [8], [9], [10]. In particular, penalized weighted least squares (PWLS) approaches have been popular for CT image reconstruction that optimize a combination of a statistically weighted quadratic data-fidelity term (capturing the forward and noise model) and a regularizer penalty that captures prior information of the object [11]. MBIR methods have often used simple regularizers [12] such as edge-preserving regularization involving nonquadratic functions of differences between neighboring pixels [13] (implying image gradients may be sparse) or other improved regularizers [14], [15], [16], [17].

Within the class of data-driven approaches, dictionary learning [12] and deep learning based methods for reconstruction have gained popularity in recent years due to their demonstrated effectiveness in removing artifacts from images in a variety

of modalities, great flexibility and the availability of curated datasets for training [18], [19], [20].

However, in general applications, acquiring many projections or ‘views’ through the object may be undesirable or impossible. This constraint may be to reduce exposure to radiation in medical imaging applications, or due to only pre-set limited or sparse views being possible in industrial or security applications. Moreover, in dynamic imaging applications, where the object is changing while being imaged, we would also be limited to fewer views per temporal state, to prevent blurring. While total variation (TV)-based MBIR methods have been extensively applied to such sparse-view and sparse angle reconstruction problems [21], [22], [23], [24], [25], [26], other model-based CT reconstruction algorithms rely upon learned prior-based regularizers [27], [28], [29]. These methods often require many iterations to converge, leading to large runtimes, and also require careful selection of regularization parameters to obtain reasonable image quality trade-offs.

Deep learning algorithms have found considerable use in many problems ranging from artifact correction to image reconstruction by combining with model-based reconstruction [30], [31], [32], [33], [34], [35], [36], [37], [38]. Deep learning approaches could be supervised or unsupervised or mixed [39], [40] and include image-domain (denoising) methods, sensor-domain methods, AUTOMAP, as well as hybrid-domain methods (cf. reviews in [12], [41]). Hybrid-domain methods are gaining increasing interest and enforce data-consistency (i.e., the reconstruction should be consistent with the measurement model) during training and reconstruction to improve stability and performance. Deep learning methods often require large training data sets and long training times to work well. They may also struggle to generalize to data with novel features or that are obtained with different experimental settings.

When reconstructing 3D objects from extremely limited tomographic views or projections, many of the aforementioned approaches fail. The FDK algorithm yields reconstructions that are severely ridden with streak artifacts. While conventional iterative methods perform better, the quality of reconstructions leave a lot to be desired, and there is often poor bias-variance trade-offs. Deep learning-based approaches have the potential to perform better in this scenario, but still perform poorly when there is a scarcity of available data for training, such as in national security applications where experimental data is limited and accurate simulations are expensive [42]. While there are approaches that reconstruct from very limited projections, they either do not target 3D CBCT imaging [19], [25], [43], [44], or rely upon many paired training image volumes [18], [20], [45], [46], [47], [48].

### B. Contributions

This paper focuses on developing a method that improves the art of reconstruction of 3D objects from very sparse cone beam CT views, when there is extremely limited data for training the reconstruction model. The proposed reconstruction approach works across multiple stages, similar to an unrolled-loop

algorithm [12], where each stage consists of a shallow CNN block trained using a combined supervised and adversarial loss, followed by a data-consistency block. The adversarial component of our loss yields destreaked images that have more realistic texture. To mitigate the challenge of reduced training data, we reduce the scope of our learning to patches or image subvolumes. This approach allows us to provide several training examples from even a single training image. Destreaked patches are aggregated before data-consistency is applied to the whole volume. Furthermore, we prime our method using an edge-preserving regularized reconstruction as input.

We compare our methods to a variety of techniques including the FDK algorithm, edge-preserving regularized reconstruction, and deep CNN-based reconstruction without data-consistency. Simulation results suggest that the proposed method provides much better image quality than previous techniques with extremely limited (four or eight) views of 3D objects.

### C. Organization

The rest of this paper is organized as follows. Section II describes the proposed approach in detail. Section III explains our choices for various algorithm parameters as well as our experimental setup. Section IV presents the results of our comparisons to other algorithms as well as other experiments that offer insights into the process of our reconstruction. Section V elaborates upon these observations. Finally, Section VI states our conclusions and offers some avenues for future research.

## II. ALGORITHM & PROBLEM SETUP

Our proposed method for CBCT reconstruction focuses on addressing two primary challenges: a very limited number of available views, and limited number of available training objects. We address the former through a combination of three aspects: (1) using an edge-preserving regularized reconstruction [4] to initialize our iterative-type algorithm; (2) including an adversarial component to the training loss function for our learned destreaking networks (similar to generative adversarial networks or GANs); and (3) including data-consistency blocks that reinforce acquired measurements in the destreaked 3D reconstruction.

The problem of scarce training data is addressed primarily by two approaches. First, we split an entire image volume into patches in the form of overlapping subvolumes. Essentially, this step localizes the scope of CNN-based destreaking to a comparatively smaller neighborhood, while allowing us to generate many training examples from a single image volume. Second, we use a shallow destreaking CNN to avoid overfitting to the training data. To reduce computation time associated with multiple 3D convolutions and subsequent patch aggregation, the CNN is designed to map 3D subvolumes to 2D slices [49]. This approach enables using 3D contextual information for the destreaking task, while removing the need for patch averaging (of overlapping 3D patches) and associated artifacts during aggregation.

Many learning-based reconstruction approaches in applications like MRI work through end-to-end training, whereas such approaches are less practical in CT due to the complexity of

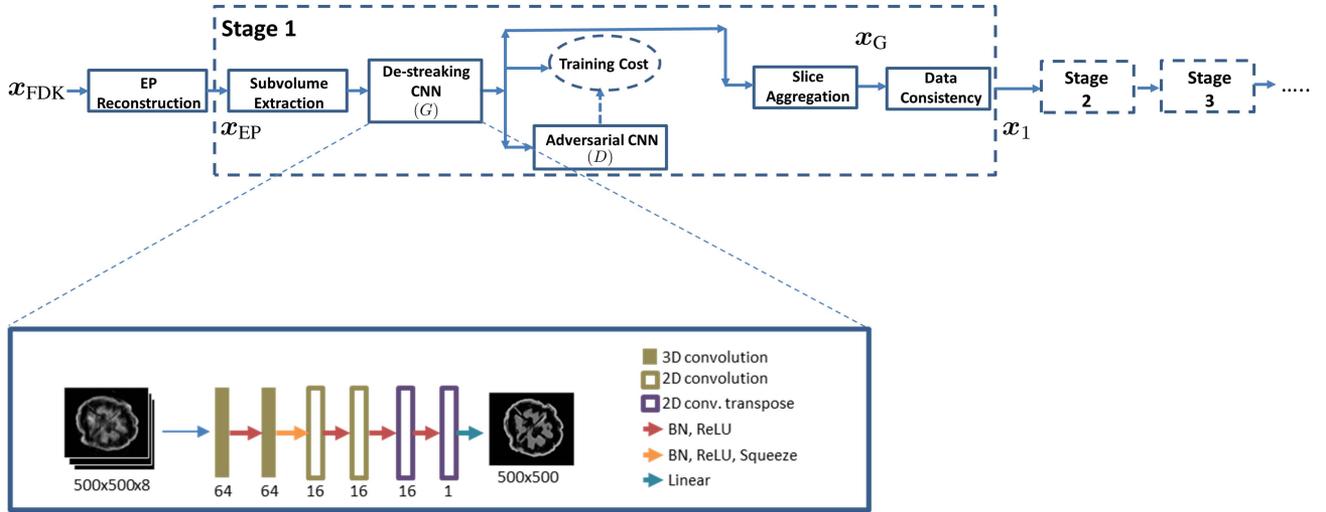


Fig. 1. Flow diagram depicting the overall pipeline of our algorithm, where  $\mathbf{x}_{\text{FDK}}$  is the FDK reconstruction,  $\mathbf{x}_{\text{EP}}$  is an edge-preserving regularized reconstruction,  $\mathbf{x}_{\text{G}}$  is the generator's output after slice aggregation, and  $\mathbf{x}_1$  is the output of the first stage.

the system matrix. Thus, our developed algorithm operates as a multi-stage greedy approach similar to works like [50], [51]. Each stage is composed of a CNN that maps each 3D subvolume with streaking artifacts to a clean 2D slice corresponding to the slice at the centre of the subvolume. Because the objects considered here have finite support, we treat the slices at the edge of the volume in the direction of aggregation are being empty, and set them to zero. One could use other boundary conditions for long objects [52]. Once an entire image volume has been aggregated from individual clean slices, this volume is passed through a data-consistency update to reinforce acquired measurements and reduce any ‘hallucinations’ introduced by the network. This output subvolume is then provided as input to the next stage. Fig. 1 depicts the process that is akin to algorithm unrolling [53]. As mentioned earlier, to reduce noise and streaks, the input to the first stage of our method is an edge-preserving regularized (iteratively obtained) reconstruction, using the regularizer in [4] and the algorithm in [54], which in turn is initialized with an FDK reconstruction for faster convergence.

We train the CNN parameters separately for each stage. The training loss for the destreaking CNN in each stage consists of a weighted combination of a masked mean squared error term calculated over a region of interest using the ground truth training image slices, and an adversarial input from another CNN that acts as a discriminator for the output of the destreaking CNN (also specific to the stage). Adversarial training often is posed as a min-max optimization problem [55], but the practical implementation involves alternating between updating generator (destreaking) network  $G$  parameters  $\phi^k$  and discriminator network  $D$  parameters  $\theta^k$ , with more frequent updates of the generator parameters. Our approach to updating the generator weights for the  $k$ th stage is mathematically expressed as:

$$\hat{\theta}^k = \arg \min_{\theta} -\lambda \mathbb{E} \left[ D_{\phi^k}(G_{\theta}(\mathcal{P}_3 \mathbf{x}_{k-1})) \right] + \mathbb{E} \left[ \left\| \mathcal{P}_3^{2,\text{mid}} \mathbf{x}_{\text{GT}} - G_{\theta}(\mathcal{P}_3 \mathbf{x}_{k-1}) \right\|_2^2 \right], \quad (1)$$

where  $\mathbf{x}_{k-1}$  is the output of the  $(k-1)$ th stage of our algorithm,  $\mathbf{x}_0$  is set to be  $\mathbf{x}_{\text{EP}}$  (the edge-preserving regularized reconstruction),  $\mathbf{x}_{\text{GT}}$  is the ground truth,  $\lambda$  is a regularization parameter that varies as the weights of  $G^k$  are updated (see Section III-B),  $\mathcal{P}_3$  is a 3D patch or subvolume extraction operator, and  $\mathcal{P}_3^{2,\text{mid}}$  is an operator that extracts the 2D central slice from an image subvolume, where the position of the subvolume is determined by  $\mathcal{P}_3$ . We restricted the reconstructions to a region in the image volume containing the object of interest. The expectation  $\mathbb{E}$  is taken over the set of training examples.

Our approach to updating the the discriminator network parameters is likewise given as:

$$\hat{\phi}^k = \arg \min_{\phi} \mathbb{E} \left[ (D_{\phi}(G_{\theta^k}(\mathcal{P}_3 \mathbf{x}_{k-1})) - 0)^2 \right] + \mathbb{E} \left[ (D_{\phi}(\mathcal{P}_3^{2,\text{mid}} \mathbf{x}_{\text{GT}}) - 1)^2 \right]. \quad (2)$$

The data-consistency update involves seeking an image that is consistent with the acquired measurements while still being ‘close’ to the slice-aggregated destreaked image. The optimization problem for this step is framed as:

$$\mathbf{x}_k = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \beta \|\mathbf{x} - \mathbf{x}_{k,G}\|_2^2, \quad (3)$$

where  $\mathbf{A}$  is the CBCT system matrix, implemented with the separable-footprint projector [56],  $\mathbf{y}$  denotes the projections or acquired measurements,  $\beta > 0$  is a regularization parameter, and  $\mathbf{x}_{k,G}$  is the output of the generator after slice aggregation at the  $k$ th stage. We used an ordinary least-squares (LS) data-fit term rather than a weighted LS (WLS) term because the focus here is on sparse views rather than low-dose imaging, but the method generalizes directly to the WLS case. We used 50 conjugate gradient (CG) iterations to minimize (3).

TABLE I

COMPARISON OF THE PERFORMANCE OF OUR PROPOSED METHOD AGAINST FDK RECONSTRUCTION, EDGE-PRESERVING (EP) REGULARIZED RECONSTRUCTION, A SINGLE STAGE OF CNN-BASED DESTREAKING WITHOUT DATA-CONSISTENCY (USING THE SAME ARCHITECTURE AS IN OUR METHOD), A 2.5D DnCNN [49] WITHOUT DATA CONSISTENCY, OUR PROPOSED APPROACH WITH 1 STAGE BUT WITH DESTREAKING CNN TRAINED FOR 160 EPOCHS AND OUR PROPOSED MULTISTAGE RECONSTRUCTION ALGORITHM FOR 8 ACQUIRED PROJECTIONS WITH 4 STAGES WITH DESTREAKING CNN TRAINED FOR 40 EPOCHS EACH. THE METRICS USED FOR PERFORMANCE ARE THE NORMALIZED MEAN ABSOLUTE ERROR (NMAE) AND THE NORMALIZED HIGH-FREQUENCY ERROR NORM (NHFEN)

Walnut #	8 views											
	FDK recon.		EP recon.		CNN destreaking		2.5D DnCNN		Proposed 1 stage		Proposed 4 stage	
	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN
1	0.77	0.90	0.45	0.58	0.40	0.59	0.38	0.69	0.33	0.62	<b>0.26</b>	<b>0.54</b>
2	0.77	0.88	0.45	0.57	0.38	0.58	0.37	0.70	0.32	0.63	<b>0.25</b>	<b>0.53</b>
3	0.79	0.90	0.49	0.61	0.42	0.62	0.41	0.74	0.38	0.68	<b>0.30</b>	<b>0.58</b>
4	0.79	0.96	0.45	0.62	0.39	0.62	0.38	0.74	0.34	0.67	<b>0.27</b>	<b>0.58</b>
5	0.82	0.98	0.48	0.65	0.41	0.66	0.39	0.78	0.35	0.72	<b>0.27</b>	<b>0.61</b>

### III. METHODS

#### A. Dataset and Experimental Settings

To train and test our method, we examined two diverse problems. The first problem is a standard sparse-view image reconstruction problem in which we utilized the publicly available 3D walnut CT dataset [57]. To study the ability to learn from very limited data, we used a single walnut for training our method, and tested our algorithm on 5 different walnuts. Furthermore, extremely limited data with 8 or 4 views/projections through the walnuts were used in training and testing the network. Separate networks were trained for reconstructing image volumes from 4 and 8 views, respectively. These CBCT views were generated using the MIRT [58] package, and were equally spaced over 360 degrees. The distance from source to detector was set to be 20 cm, the distance from the object to the detector was 4.08 cm, and each projection view was  $150 \times 150$  pixels of size  $\approx 0.4 \text{ mm}^2$ . Our second problem examines [59] the reconstruction of complex hydrodynamic phenomena using the computational fluid dynamics software CTH [60]. Specifically, we examined the ability to perform reconstructions from a single hydrodynamic training simulation using only a single temporal slice and 8 projections. To illustrate the variation of the voxel values, Fig. 8 provides a few one dimensional profiles of different simulations across the central row of the central slice.

The ability of the algorithm to reconstruct hydrodynamic simulations generated with different hydrodynamic parameters was tested using two simulations, one referred to as simulation 12 (see Figs. 9 and 10) that was generated using one set of parameters at an early time and another referred to as simulation 7 (see Figs. 11 and 12) that was generated using a different set of parameters and at a later time in its simulation. We also used two grossly different volumes to train the network to examine the ability of the network to generalize from different training cases. The CBCT data here were generated in the same manner as the walnut data. The 2.5D Deep CNN and our proposed method were both trained on the volume within a bounding box containing the outermost shell; this outer shell can be seen in Figs. 9–12. For comparison, all reconstruction metrics were reported in a  $130 \times 130 \times 130$  bounding box containing the inner topological structures in the volume. This box is shown in the zoomed in parts of images shown in Figs. 9–11.

Because the CBCT system simulated here has a small cone angle (almost parallel beam), 8 views over  $360^\circ$  probably has only a bit more information than 4 views over  $180^\circ$ . The image volume for each walnut was  $501 \times 501 \times 501$ . The dimensions of each voxel were approx.  $0.12 \times 0.12 \times 0.12 \text{ mm}^3$ , whereas the dimensions of the hydrodynamic simulations were  $448 \times 448 \times 448$  with voxel size  $0.025 \times 0.025 \times 0.025 \text{ cm}^3$ .

#### B. Hyperparameters and Network Architectures

The subvolume size for our scheme was chosen to be  $500 \times 500 \times 8$ . The parameter  $\beta$  was chosen to be 1, and  $\lambda$  was changed dynamically as the network weights were updated according to  $10^{\lfloor \log(r) \rfloor}$ , where  $r = \mathbb{E}[\|\mathcal{P}_3^{\text{mid}} \mathbf{x}_{\text{GT}} - G_\theta(\mathcal{P}_3 \mathbf{x}_{k-1})\|_2^2]$ . This was done to maintain balance between the  $\ell_2$  and adversarial loss components during generator training. The number of stages in our method was set to 4, and the networks in each stage were trained for 40 epochs. The weights for the discriminator were updated once for every 10 times the weights of the destreaking CNN were updated.

Fig. 1 depicts the generator architecture we used. The kernel size for 3D convolutions was  $(3 \times 3 \times 3)$ , and  $(3 \times 3)$  for the 2D convolutions. The discriminator was akin to a classifier with two convolutional layers and with 8 filters in each convolutional layer with  $(3 \times 3)$  kernel size with stride 1 followed by fully connected layers with (1152,8,8) nodes respectively, with a sigmoid activation at the final output to constrain the output to be between 0 and 1. The batch size during training the destreaking CNNs was 6. Training the destreaking CNN in each stage of our algorithm for walnut data took approx. 10 hours on 3 NVIDIA Quadro RTX 5000 GPUs, while at test time, each walnut volume required 7 minutes to reconstruct with a batch size of 3 on two of the same GPUs. The data consistency update required an additional 3 minutes on a workstation with Intel(R) Xeon(R) Silver 4214 CPU @ 2.20 GHz with 48 cores.

#### C. Compared Methods

To assess the performance of our method, we used 4 stages of our proposed algorithm to achieve a compromise between image quality and runtime. We compared the output for all 5 test walnuts to the conventional FDK reconstruction, an edge-preserving (EP) regularized (MBIR) reconstruction [4], as well as the slice-aggregated output from a single stage of our destreaking

TABLE II  
COMPARISON OF THE PERFORMANCE OF OUR PROPOSED METHOD AGAINST FDK RECONSTRUCTION, EDGE-PRESERVING (EP) REGULARIZED RECONSTRUCTION, A SINGLE STAGE OF CNN-BASED DESTREAKING WITHOUT DATA-CONSISTENCY (USING THE SAME ARCHITECTURE AS IN OUR METHOD), A 2.5D DnCNN [49] WITHOUT DATA CONSISTENCY AND OUR PROPOSED MULTISTAGE RECONSTRUCTION ALGORITHM 4 ACQUIRED PROJECTIONS. THE METRICS USED FOR PERFORMANCE ARE THE NORMALIZED MEAN ABSOLUTE ERROR (NMAE) AND THE NORMALIZED HIGH-FREQUENCY ERROR NORM (NHFEN)

Walnut #	4 views									
	FDK recon.		EP recon.		CNN destreaking		2.5D DnCNN		Proposed	
	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN
1	1.06	1.08	0.65	0.61	0.65	0.62	0.62	0.71	<b>0.48</b>	<b>0.62</b>
2	1.14	1.07	0.68	0.60	0.67	0.61	0.64	0.72	<b>0.51</b>	<b>0.62</b>
3	1.28	1.21	0.66	0.64	0.66	0.65	0.66	0.75	<b>0.52</b>	<b>0.66</b>
4	1.33	1.32	0.66	0.64	0.65	0.65	0.64	0.78	<b>0.52</b>	<b>0.66</b>
5	1.31	1.33	0.68	0.69	0.67	0.68	0.65	0.81	<b>0.53</b>	<b>0.71</b>

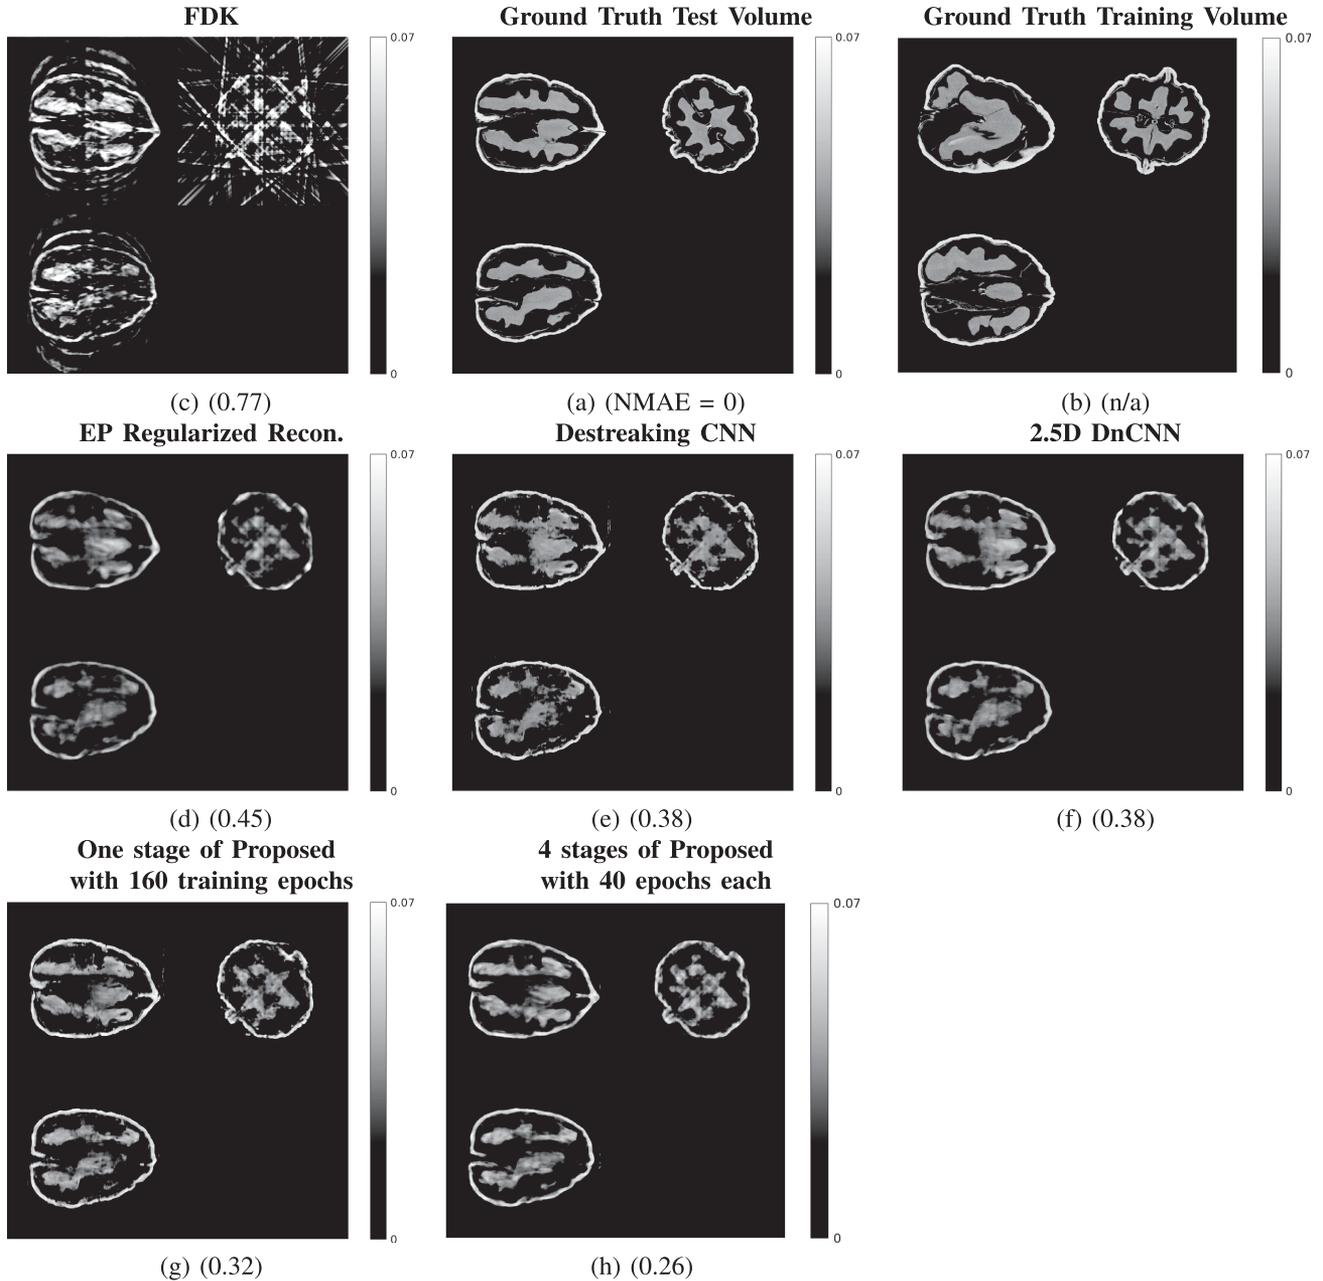


Fig. 2. Comparison of the quality of reconstruction of our proposed algorithm (h) for walnut 2 (8 views) in Table I to various reference methods. Each subfigure depicts slices through the center of the walnut volume in three different directions (or sagittal, coronal and transverse orientations). The normalized mean absolute errors have also been shown underneath each subfigure. The central slices corresponding to the ground truth training walnut volume have also been shown in (b).

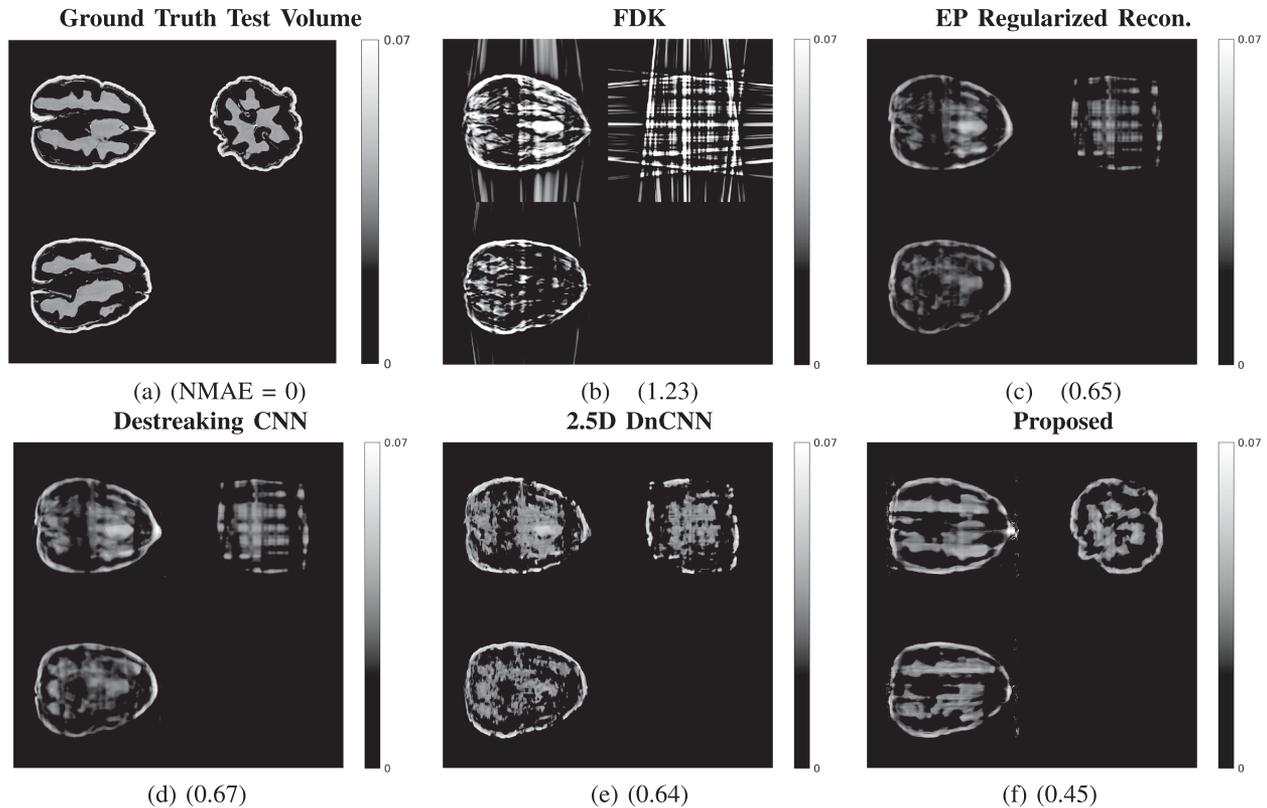


Fig. 3. Comparison of the quality of reconstruction of our proposed algorithm (f) for walnut 2 (4 views) in Table II to various reference methods. Each subfigure depicts slices through the center of the walnut volume in the sagittal, coronal and transverse orientations. The normalized mean absolute errors have also been shown underneath each subfigure.

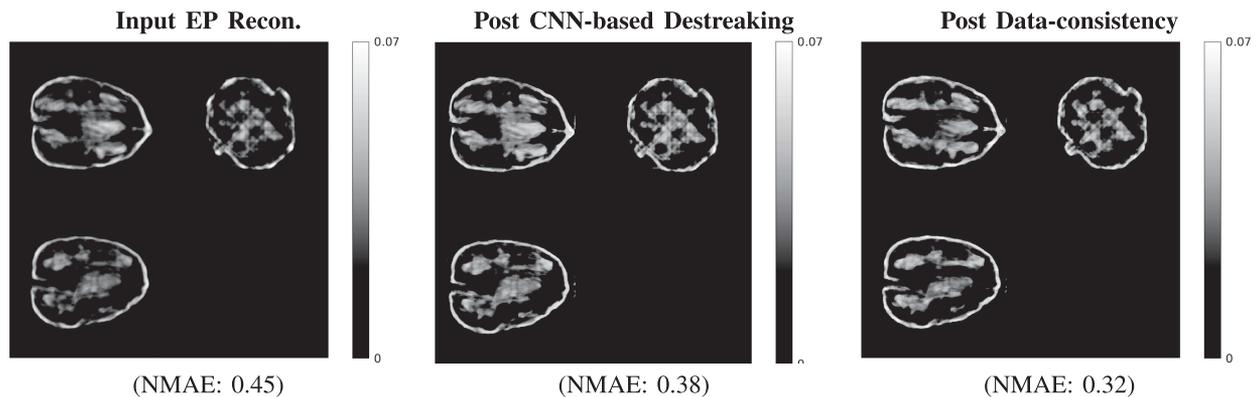


Fig. 4. Central slices corresponding to all three orientations of test walnut 2 (8 views) after various steps during a single stage (Stage 1) of our proposed algorithm.

CNN without data consistency. We also compared our proposed method to a deep 2.5D (residual) CNN whose architecture is as in [49] and without data consistency. This network had 16 hidden convolutional layers, whose input was a subvolume of three slices, and its output was a single slice trained to be the residual needed to be added to the middle slice of the EP reconstruction input to get the full-views image. The network was trained with a learning rate of 0.001 for Walnuts and 0.0002 for hydrodynamic simulations, both using the Adam optimizer with a learning rate scheduler that multiplied the current learning rate by 0.1 every 50 epochs. The network was trained for 100 epochs. A batch size of

4 was used for Walnut dataset whereas a batch size of 6 was used for hydrodynamic simulations. This network was trained using masked mean squared error between output residual and ground truth-based residual, where the mask was the bounding box containing the edges of the training volume. Ground truth residuals were obtained by subtracting the EP slice from full-views or ground truth image. We trained this model with and without data augmentation that involved random elastic deformations, random rotations, and flips. However, we observed that data augmentation did not improve the performance on the unseen test volumes for both the datasets. Data augmentation increased

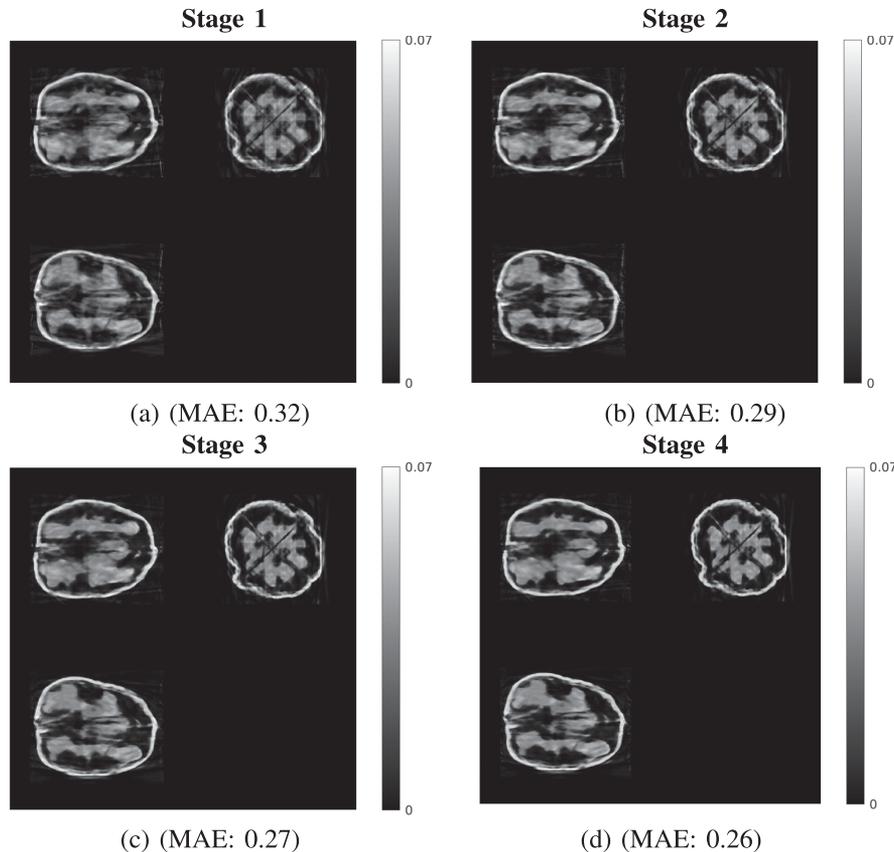


Fig. 5. Central slices through the reconstructions of walnut 1 (8 views) in all three orientations at the end of each of the 4 stages of our algorithm (post data-consistency). It is evident that the quality of the reconstruction progressively improves across the stages. The normalized mean absolute error (NMAE) for each figure is also provided underneath.

the training time, but it still gave similar performance to without data augmentation. We therefore simply show the latter results. Finally, we also compare the robustness of our proposed four stage method by testing with smaller object to detector and source to detector distances, wherein cone beam artifacts are more prevalent and different than in the training set.<sup>1</sup>

#### D. Performance Metrics

We primarily used the normalized mean absolute error (NMAE) as a metric for evaluating the performance of various methods. For the walnut data, the error is evaluated over the voxels within the region-of-interest (ROI) of a three-dimensional mask obtained by dilating a ground truth segmentation of the walnut being reconstructed. The masked region includes all voxels within the shell of the walnut. The NMAE normalization used the mean intensity of the ground truth voxels within this mask. Essentially,  $E_{\text{NMAE}}(x_{GT}, x_O, \mathcal{M}) = \|\mathcal{M} \odot (x_{GT} - x_O)\|_1 / \|\mathcal{M} \odot x_{GT}\|_1$ , where  $x_{GT}$  is the ground truth image volume,  $x_O$  is the reconstruction whose quality is being evaluated, and  $\mathcal{M}$  is a binary mask specific to the test volume (e.g., which excludes any pixel not within the dilation of the outer boundary

for walnuts).<sup>2</sup> For computing metrics for hydrodynamic data, we let the mask cover a central region of the volume having the predominant object dynamics. The masks were applied so that a cubical region of size  $130 \times 130 \times 130$  encompasses the central topological features. This was done to evaluate the reconstructions of this interior region which had complex features of interest.

Another metric that is used for comparison in our work is the normalized high-frequency error norm or NHFEN [61]. We computed the HFEN for every slice of the reconstructed walnut as the  $\ell_2$  norm of the difference of masked edges (obtained through a high-pass filtering) between the input and reference images. The masking is done similarly as described earlier. A Laplacian of Gaussian (LoG) filter was used as the edge detector. The kernel size was set to  $15 \times 15$ , with a standard deviation of 1.5 pixels. The normalization was performed over the high frequency components of the ground truth image over the masked ROI. Mathematically, this metric is calculated as  $E_{\text{NHFEN}}(x_{GT}, x_O, \mathcal{M}) = \frac{1}{N} \sum_i \|\mathcal{H}(\mathcal{M}[:, :, i]) \odot x_{GT}[:, :, i] - \mathcal{H}(\mathcal{M}[:, :, i]) \odot x_O[:, :, i]\|_2 / \|\mathcal{H}(\mathcal{M}[:, :, i]) \odot x_{GT}[:, :, i]\|_2$ , where  $\mathcal{H}$  denotes the LoG filter described earlier,  $i$  indexes the slices of the image volume in the  $z$  direction, where  $N$  is the total number of slices in that direction, and the other

<sup>1</sup>Implementation of proposed and compared methods are available in <https://github.com/gtm2122/SparseViewCT-TCI>

<sup>2</sup>These are obtained by a histogram-based thresholding of the corresponding ground truth volumes for the test walnuts.

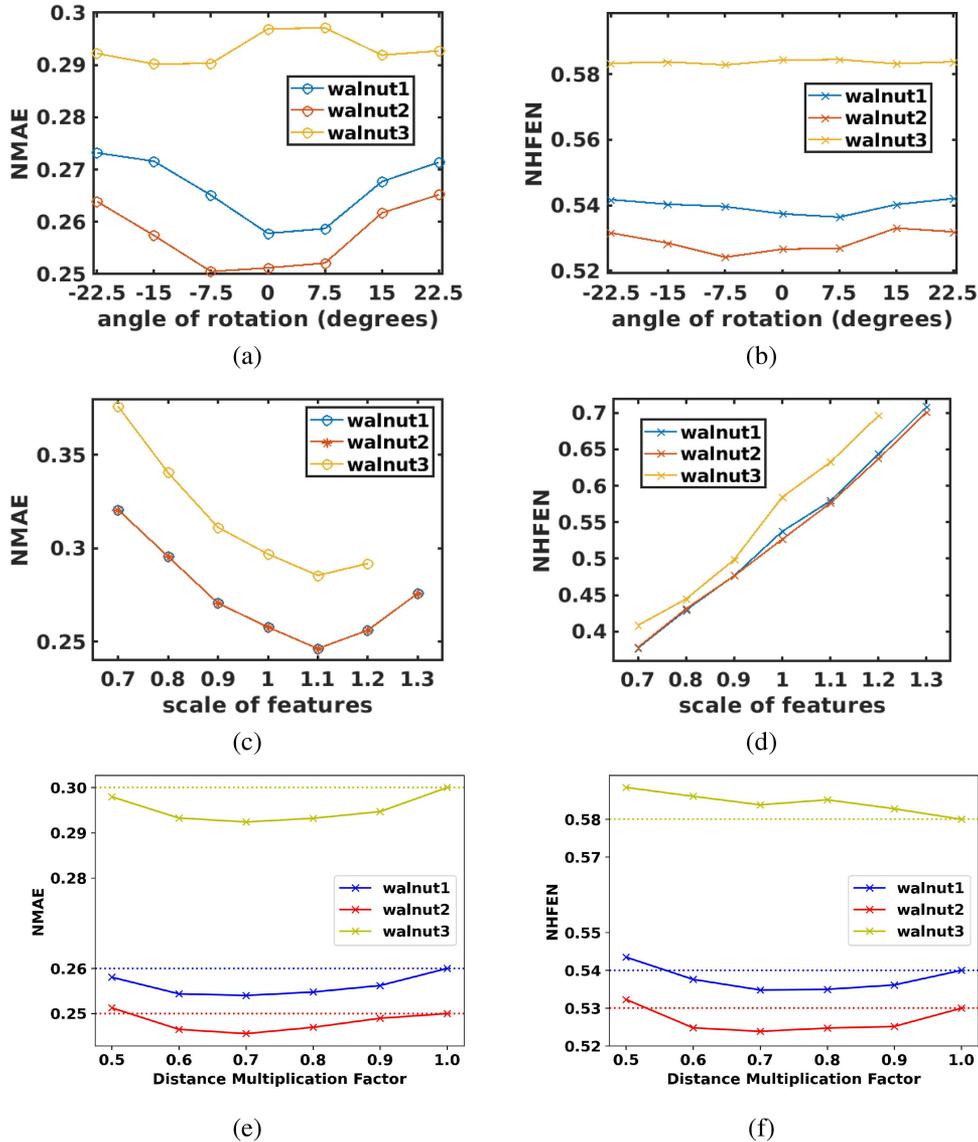


Fig. 6. (a) and (b) show the effects of rotating the test walnuts (by changing the angle of acquired projections) on the quality of the reconstructions for 8 acquired views, while (c) and (d) capture the effect of varying the scale of features of the test walnut compared to the scale used in training. (e) and (f) show the effect of different cone-beam artifacts by simultaneously multiplying source to detector and object to detector distances by a multiplication factor. (a), (c), and (e) depict the normalized MAE metric while (b), (d), and (f) depict the normalized HFEN error metric for the respective experiments.

symbols have their usual meaning, as described previously. An advantage of using such normalized metrics is that it allows for the evaluation of the reconstruction quality only in areas of interest in the volume, disregarding the effect of empty spaces around it.

## IV. RESULTS

### A. Walnut Dataset Reconstruction Comparisons

Table I and Table II compare the reconstruction performance of various methods (including our own) described in the previous sections. The proposed approach substantially improves the NMAE and NHFEN compared to the reference methods for 8 and 4 acquired projections, respectively. As expected, the quality of reconstructions using 4 acquired projections was worse than when 8 projections were acquired for reconstruction.

Additionally, to better test the advantage of having multiple stages in our method, we trained the destreaking network in Stage 1 for 4 times the epochs as our multi-stage scheme, and followed up the network with data-consistency once. This method was compared against our proposed approach where each destreaking CNN was trained for 40 epochs followed by a data consistency step for a total of 4 stages. Table I shows this comparison in terms of reconstruction metrics. Fig. 2 shows a visual comparison of reconstruction quality for 8 views with the previously mentioned methods and our 4 stage approach. The proposed algorithm provides significantly higher quality reconstructions than the other methods. This is particularly evident in the extent to which our algorithm is able to restore the finer features of the walnuts, and has fewer artifacts.

Fig. 3 compares the reconstructions from 4 views. The quality of the reconstruction is poorer compared to that using 8 views,

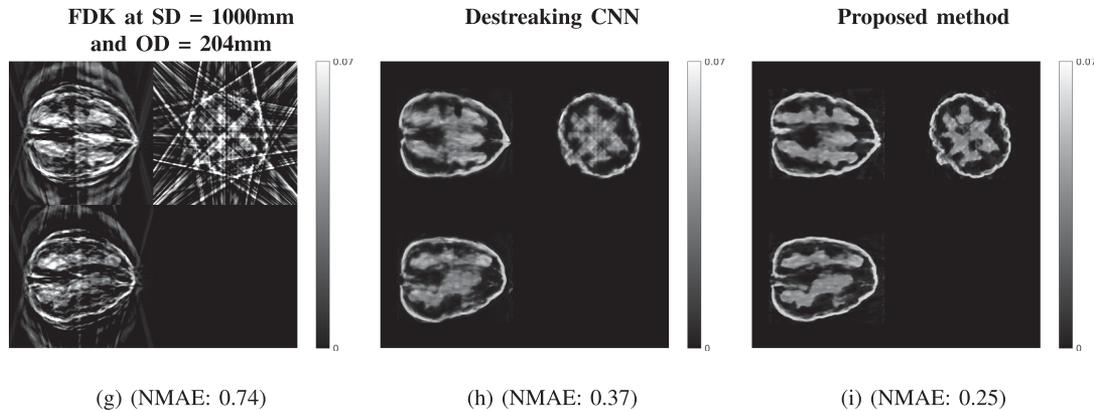


Fig. 7. Comparison of the quality of reconstruction of our proposed algorithm at half the source to detector and object to detector distances as used during training. Reconstructions are shown for walnut 2 (8 views) with (a) FDK, (b) Destreaking CNN, and (c) the proposed method. The cone beam artifacts are more prevalent in (a). Each subfigure depicts slices through the center of the walnut volume in three different orientations. In the image titles, OD stands for Object to Detector Distance and SD stands for Source to Detector distance.

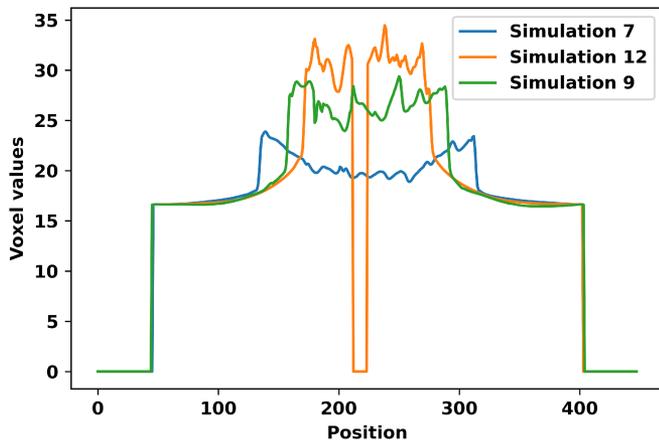


Fig. 8. 1-D Profile of central row of central slice of three different Hydrodynamic Simulations.

though the proposed approach still visibly outperforms the other methods.

Fig. 4 shows the results of the intermediate steps of the first stage of our reconstruction for walnut 2 (8 views) in the test dataset. While the patch-based destreaking compensates for the blurring introduced by the EP-regularized reconstruction and ‘fills-in’ details in the reconstructed volume, the data-consistency plays a key role in mitigating hallucinations introduced by the CNN, and reinforces image features that are consistent with the acquired measurements.

Fig. 5 shows walnut 1 from our test dataset being progressively reconstructed from 8 projections across the stages of our algorithm; as the stages progress, more features are restored in the reconstructed walnut, until the improvements become incremental. The residual streaking artifacts outside the walnut are mitigated in the reconstructions from the third and fourth stages.

Fig. 6(a) and (b) show how the quality of the reconstructions from 8 acquired views using our method varies when the orientation of the test walnuts is changed, in comparison to training-time. This is achieved by changing the position of

the acquired (equidistant) projections, as this is akin to rotating the test walnut. For this purpose, essentially the position of the first acquired projection is shifted by a specified angle. For the three test walnuts shown in the figure, the angle of rotation was changed between  $+22.5^\circ$  and  $-22.5^\circ$  in intervals of  $7.5^\circ$ . Both the normalized mean absolute error (NMAE) and the normalized high-frequency error norm (NHFEN) were used as quality metrics in this experiment. Our method seems to be fairly robust to rotated test-data.

We also studied the effect of varying the scale of the features of test walnuts on reconstructions from our method (from 8 projections), which was still trained on a single scale. For this purpose, we used the three walnuts which were used in the previous experiment. From Fig. 6(c) and (d), it is evident that while there are differences in the NMAE and NHFEN across scales, our method holds up reasonably well across multi-scale test data. Particularly, we notice that while higher-frequency features are better reconstructed when the test walnuts are scaled down, the best NMAE is observed when the scale of the test walnut is 1.1 times its original scale. We surmise that this happens because at a scale of 1.1, the scale of features of the test walnuts matches that of the training walnut very well. (We excluded scale 1.3 for walnut 3 because at that scale, the walnut exceeded the  $501 \times 501 \times 501$  voxel grid.)

To test the robustness of the proposed approach with respect to some variations in the acquisition geometry (different cone-beam artifacts), we simulated projections by varying the source to detector and object to detector distances simultaneously in our setup, such that the ratio of the two remained roughly the same. We used the same network for 8 views as in the previous cases and tested it in these modified setups, where only the data-consistency module used the correct forward model whereas the CNNs were fixed (to test their generalizability). Fig. 6(e) and (f) show the variation of reconstruction metrics when the source to detector and object to detector distances are multiplied by a multiplication factor. Fig. 7 also shows an example of reconstructions at 1000 mm source to detector distance and 204 mm object to detector distance (i.e., at multiplication factor of 0.5). The results show that the reconstruction quality varies

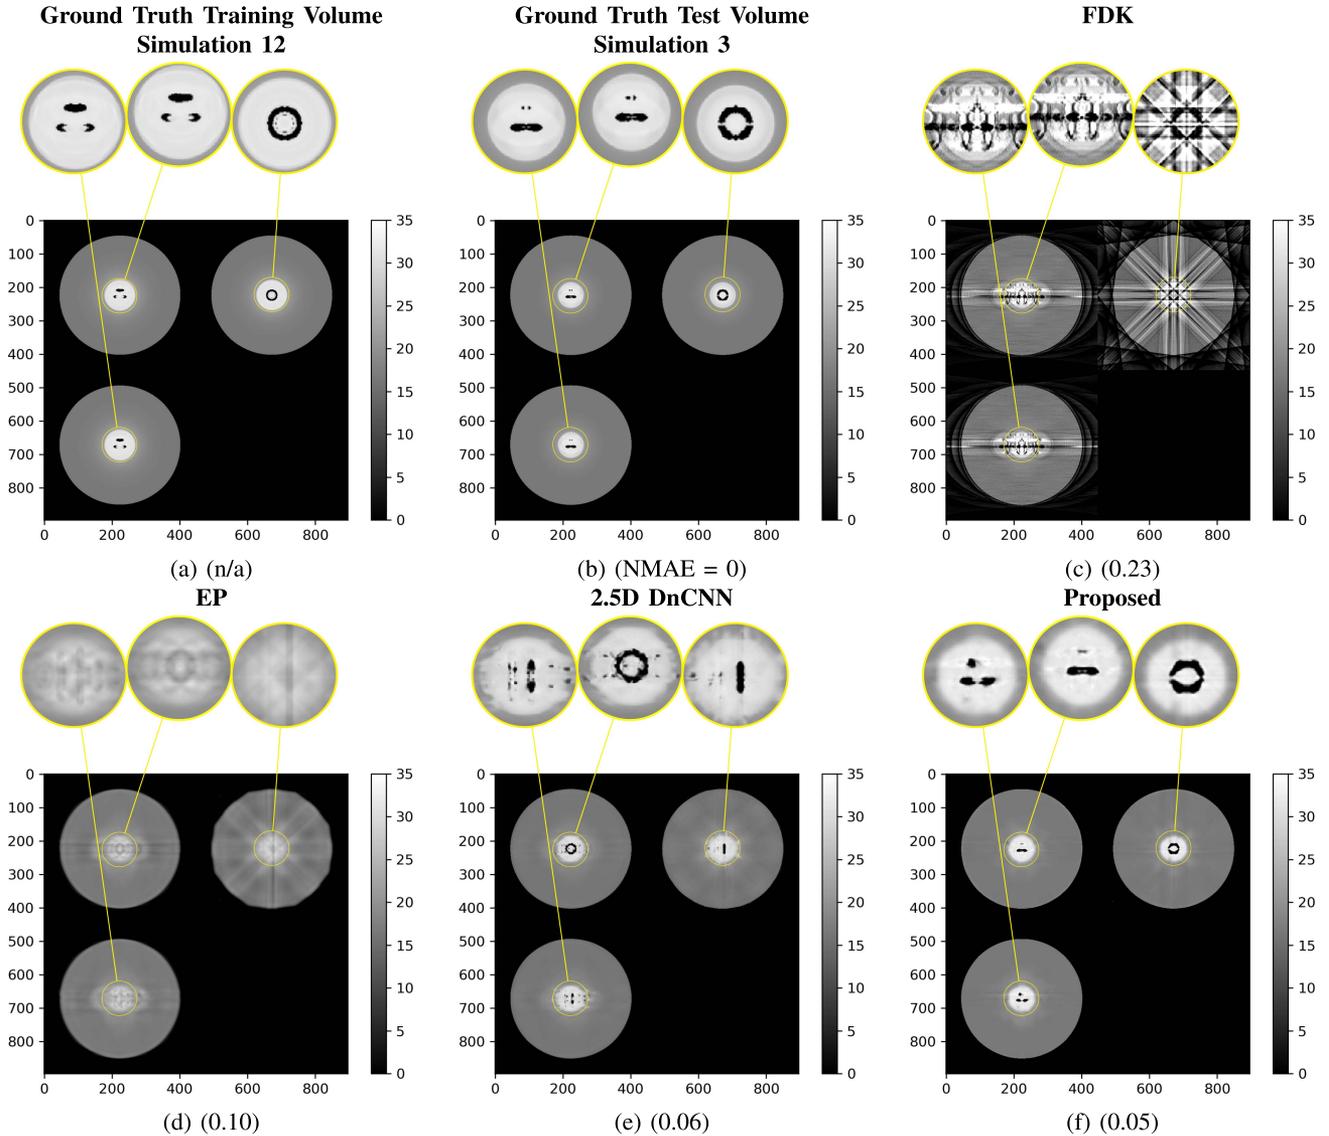


Fig. 9. Comparison of different methods for reconstructing hydrodynamic data from 8 views. All methods were trained on the simulation shown in (a). Zoom-ins of the central portion of each image are provided. The proposed method clearly outperforms the other schemes.

only somewhat with changes in the artifacts (i.e., with more cone beam type artifacts), indicating a level of robustness for the proposed approach.

### B. Hydrodynamic Simulation Reconstruction Comparisons

Table III displays a similar comparison with two versions of our proposed model, each trained on a different hydrodynamic simulation using the same acquisition geometry as used for the walnuts in Table I. There were 12 simulations in total generated using 4 set of simulation parameters and each were sampled at different timepoints. Out of these 12, simulations 7 and 12 were used for training. These were generated using different sets of parameters and the training volumes corresponded to two different time points in the simulations. The results show that the proposed method outperforms the compared methods across a variety of test cases in both training scenarios. Fig. 9

compares reconstructions of a test simulation by various methods mentioned above; our proposed method clearly outperformed all compared methods. Fig. 10 compares reconstructions by our proposed method against 2.5D DnCNN [49] for a case where the latter showed a comparable reconstruction to our proposed method in terms of NMAE. Despite the similar NMAE, the image features are still visually better for our proposed network. We see similar observations for Figs. 11 and 12, where the reconstruction is visually better for our proposed method when trained on Simulation 7.

## V. DISCUSSION

Our observations in Section IV indicate that the proposed physics-aware learning-based approach for limited-view CBCT reconstruction is able to improve upon the quality of reconstructions yielded not only by the FDK algorithm, but also

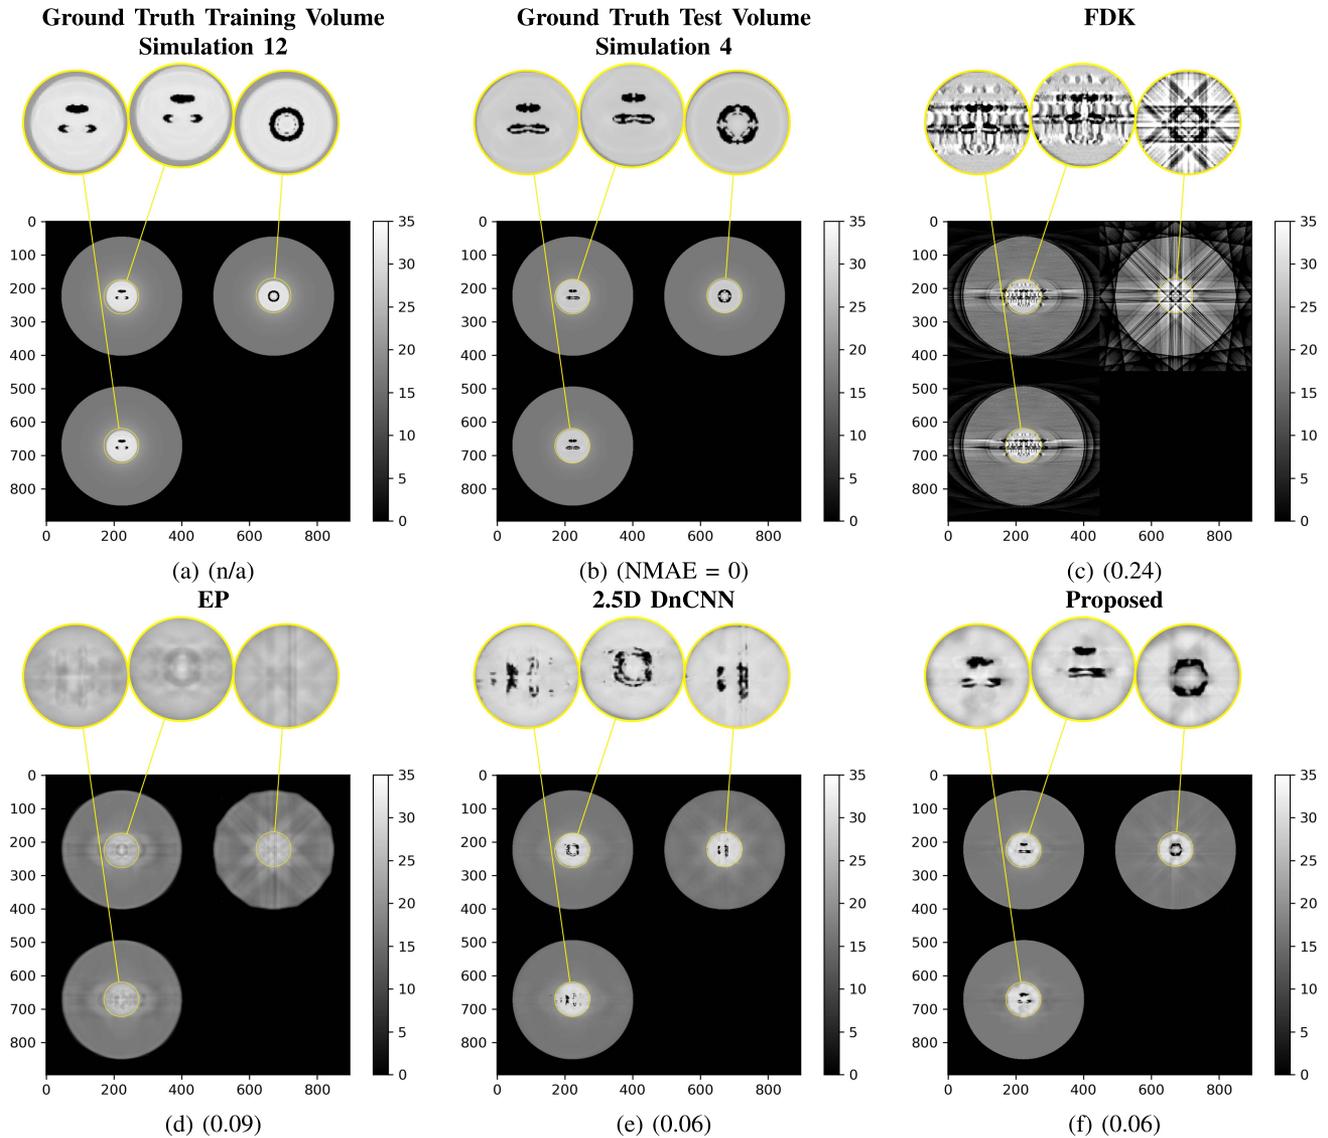


Fig. 10. Comparison of different methods for reconstructing hydrodynamic data from 8 views. All methods were trained on the simulation shown in (a). Zoom-ins of the central portion of each image are provided. The proposed method clearly outperforms the other schemes.

traditional prior-based iterative reconstruction, which serves a crucial role in initializing our algorithm, and ensures that the subsequent CNN-based destreaking is afforded a reasonable mapping to learn when trained with ground truth images as targets. Furthermore, in Fig. 4, we also demonstrated the importance of data-consistency when reconstructing from limited measurements. The data consistency update could easily be the most crucial step in our algorithm because it corrects for hallucinations introduced by the CNN-based destreaking step, which are very likely to occur given the limited availability of training data. This allows our algorithm to be repeated for several stages, which is key to the improvement provided by the proposed method over a simple image domain CNN-based denoiser. The generalizability of our approach is evident from the consistent improvements yielded by our approach in Table I and Table II, which is a consequence of using subvolumes or patches in training our method as well as the relative shallowness

of the CNNs used therein. These measures effectively reduce the chances of overfitting when training on extremely limited full-view data. The chosen patch size effectively allowed the learned destreaking CNN to utilize three-dimensional context while also reducing: (1) the total number of overlapping subvolumes that are forward propagated through the destreaking CNN at test time and (2) the memory requirements for training the network. The hybrid 3D to 2D mapping also plays an important role by reducing the computational and time demands of our algorithm, and by removing any chances of artifacts that may be introduced during (conventional overlapping) patch aggregation.

We noted that the final reconstruction quality for some test walnuts was better than others, and also that the final reconstruction quality from our algorithm depended on the quality of the initial EP reconstruction. We think this may be because the acquired views (equidistant over  $360^\circ$ ) may not be the best choice for higher fidelity reconstructions for some walnuts,

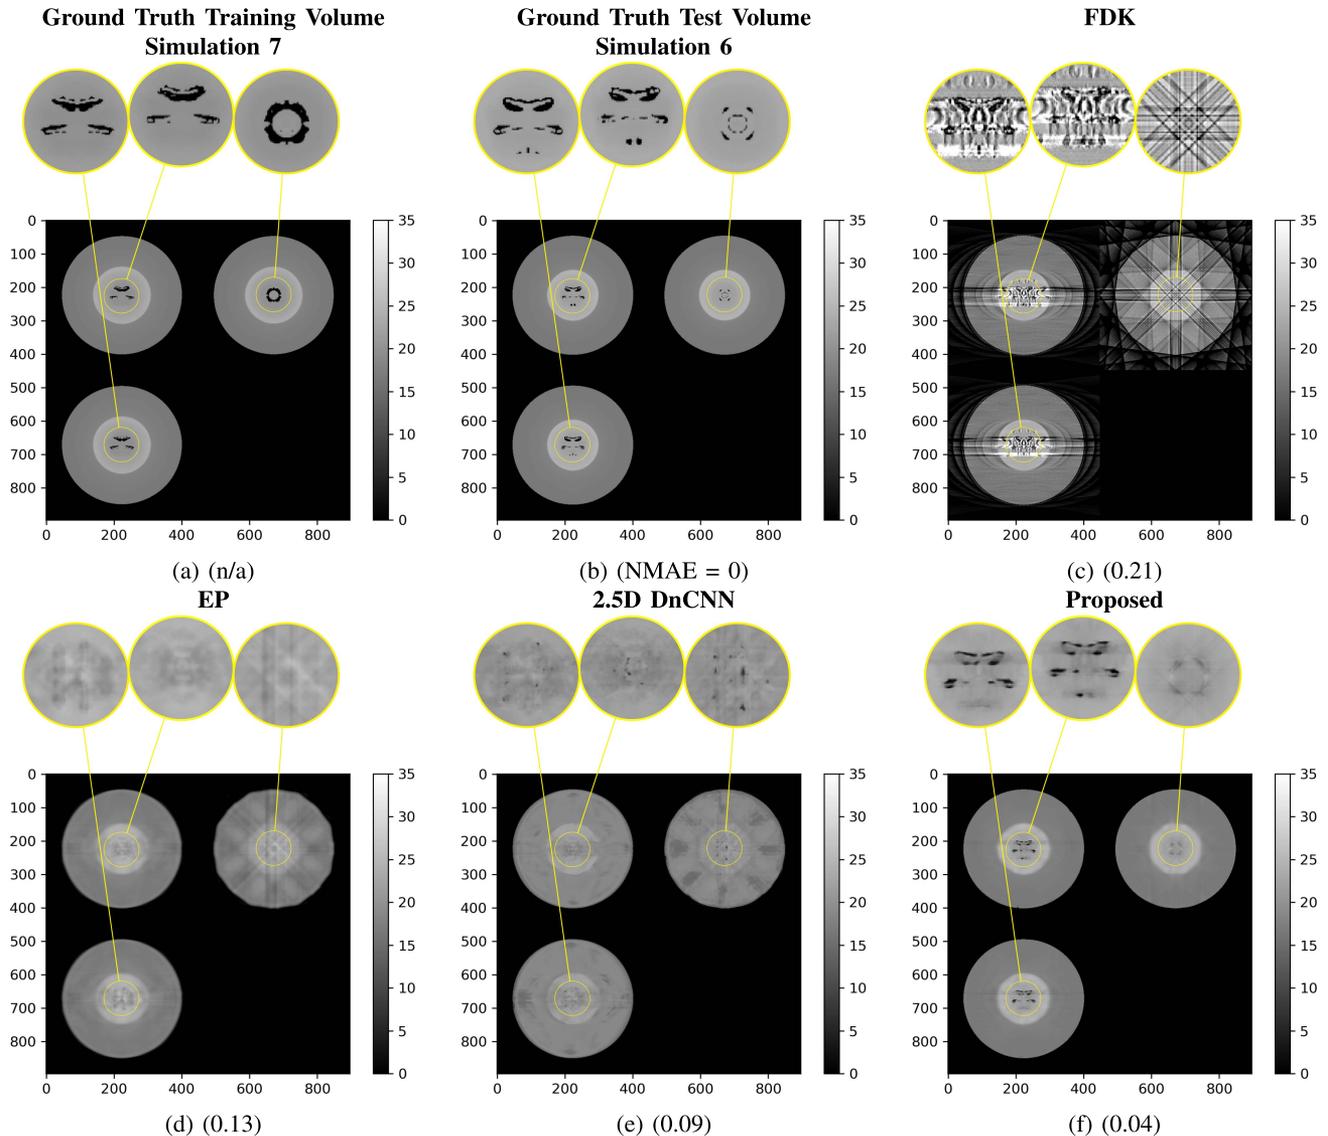


Fig. 11. Comparison of different methods for reconstructing hydrodynamic data from 8 views. All methods were trained on the simulation shown in (a). Zoom-ins of the central portion of each image are provided. The proposed method clearly outperforms the other approaches.

and because these walnuts may be less similar to the training walnut.

In comparison to 2.5D DnCNN reconstructions, the proposed approach was substantially better at reconstructing walnuts even with 4 views, where the differences are very noticeable, and the 2.5D DnCNN failed to reconstruct not only the internal features but also the boundary of the test walnut.

The proposed approach also showed robustness to different cone-beam artifacts caused by varying the source-detector and object-detector distances. Despite the initial FDK reconstruction having different amounts of cone-beam artifacts, the final reconstructions, in terms of both NMAE and HFEN metrics, hardly differed as shown in Fig. 7 and from the appearance of reconstructed features.

The proposed approach displayed robustness to reconstruction of hydrodynamic simulations. As seen in Fig. 12, despite being trained on a hydrodynamic simulation with very different

topological features in the center, our algorithm was able to reconstruct the edges of a test hydrodynamic simulation with very different features, whereas The 2.5D DnCNN failed to capture this despite having similar reconstruction NMAE values. Fig. 10 is another example that demonstrates the proposed method's ability to generalize better than the 2.5D DnCNN.

Regarding the performance of our reconstructions in terms of the NHFEN metric, we observe a lack of sensitivity with respect to other compared techniques for both 8 and 4 views. We surmise that this is because neither of the networks was trained for improved NHFEN, i.e., a NHFEN term was not part of the training loss.

While the experiments with rotated projections and varying scale of features showcased some robustness of our proposed approach to both rotation and up/downscaling of features— it also brought to light the necessity for data augmentation for improved performance. The consistent degradation of the high

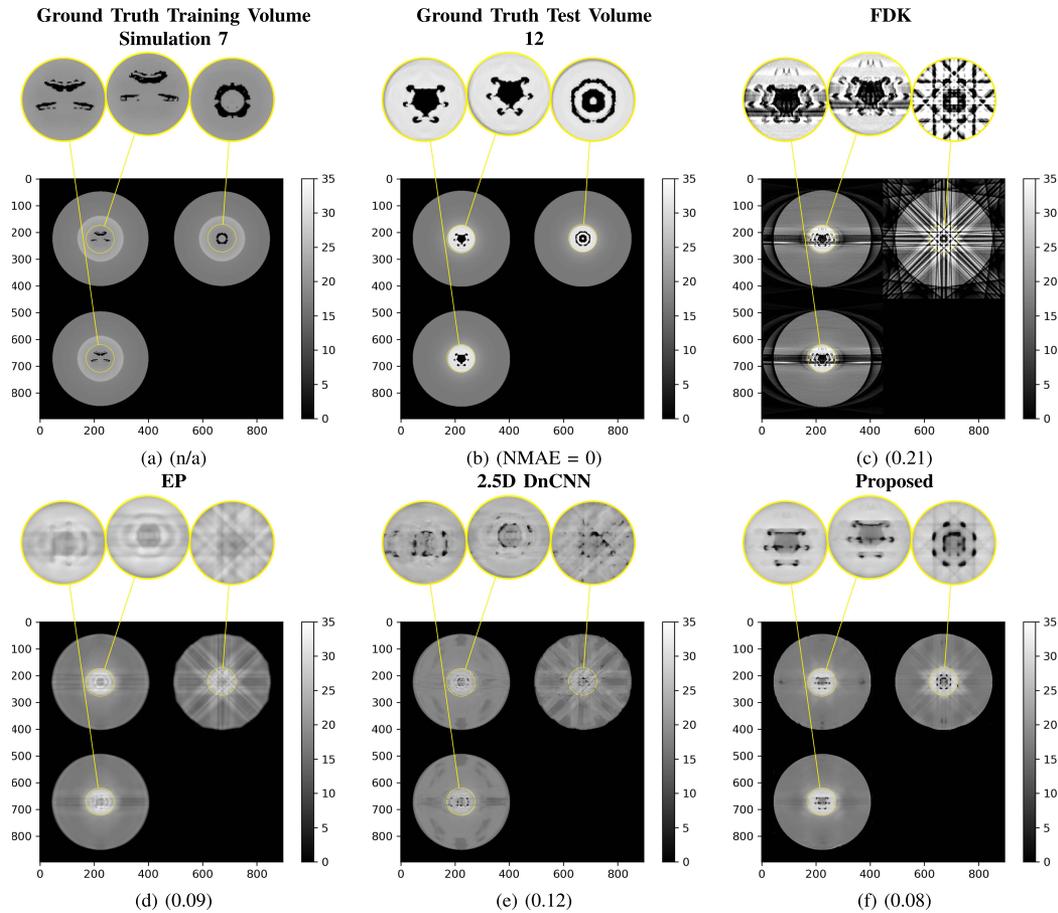


Fig. 12. Comparison of different methods for reconstructing hydrodynamic data from 8 views. All methods were trained on the simulation shown in (a). Zoom-ins of the central portion of each image are provided. The proposed method clearly outperforms the other approaches.

TABLE III

COMPARISON OF THE PERFORMANCE OF OUR PROPOSED METHOD FOR 8 VIEWS AGAINST FDK RECONSTRUCTION, EDGE-PRESERVING (EP) REGULARIZED RECONSTRUCTION AND 2.5D DNCNN [49]. THE TRAINING WAS DONE ON TWO DIFFERENT 3D HYDRODYNAMIC SIMULATIONS. ALL SCORES ARE REPORTED ON A  $130 \times 130 \times 130$  PIXEL VOLUME AT THE CENTER OF EACH RECONSTRUCTION.

Trained on Simulation 12										
Simulation #	FDK recon.		EP recon.		CNN destreaking		2.5D DnCNN		Proposed	
	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN
1	0.17	1.07	0.06	0.33	0.11	0.34	0.08	0.34	<b>0.05</b>	<b>0.31</b>
2	0.21	1.17	0.08	0.37	0.09	0.32	0.07	0.46	<b>0.05</b>	<b>0.3</b>
3	0.23	1.31	0.1	0.43	0.07	0.39	0.06	0.55	<b>0.05</b>	<b>0.37</b>
4	0.24	1.43	0.09	0.49	0.08	0.46	0.06	0.6	<b>0.06</b>	<b>0.45</b>
5	0.22	1.47	0.09	0.54	0.1	0.53	0.11	0.61	<b>0.09</b>	<b>0.52</b>
6	0.21	1.38	0.09	0.52	0.08	0.52	0.15	0.63	<b>0.07</b>	<b>0.5</b>
8	0.19	1.31	0.09	0.51	0.08	0.51	0.12	0.55	<b>0.08</b>	<b>0.49</b>
7	0.2	1.31	0.08	0.47	0.06	0.47	0.14	0.49	<b>0.05</b>	<b>0.45</b>
9	0.23	1.25	0.09	0.41	0.06	0.33	0.1	0.69	<b>0.04</b>	<b>0.31</b>
10	0.23	1.54	0.09	0.55	0.07	0.53	0.08	0.68	<b>0.07</b>	<b>0.51</b>
11	0.24	1.39	0.1	0.48	0.07	0.41	0.05	0.62	<b>0.05</b>	<b>0.4</b>
Trained on different Simulation 7										
Simulation #	FDK recon.		EP recon.		CNN destreaking		2.5D DnCNN		Proposed	
	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN	NMAE	NHFEN
1	0.17	1.07	0.06	0.33	0.1	0.47	0.08	0.41	<b>0.08</b>	<b>0.46</b>
2	0.21	1.17	0.08	0.37	0.09	0.34	0.1	0.46	<b>0.07</b>	<b>0.33</b>
3	0.23	1.31	0.1	0.43	0.09	0.42	0.12	0.57	<b>0.08</b>	<b>0.4</b>
4	0.24	1.43	0.09	0.49	0.09	0.46	0.13	0.58	<b>0.07</b>	<b>0.45</b>
5	0.22	1.47	0.09	0.54	0.07	0.48	0.14	0.57	<b>0.05</b>	<b>0.47</b>
6	0.21	1.38	0.09	0.52	0.05	0.44	0.13	0.53	<b>0.04</b>	<b>0.43</b>
8	0.19	1.31	0.09	0.51	0.09	0.46	0.14	0.55	<b>0.05</b>	<b>0.44</b>
9	0.23	1.25	0.09	0.41	0.1	0.39	0.11	0.57	<b>0.08</b>	<b>0.37</b>
10	0.23	1.54	0.09	0.55	0.1	0.52	0.14	0.63	<b>0.07</b>	<b>0.5</b>
11	0.24	1.39	0.1	0.48	0.1	0.45	0.13	0.58	<b>0.08</b>	<b>0.43</b>
12	0.21	1.27	0.09	0.43	0.1	0.4	0.12	0.54	<b>0.08</b>	<b>0.39</b>

frequency features (observed through the NHFEN metric) in reconstructions with increase in scale was interesting to note, and its contrast with the trends in NMAE points towards a sharpness/fidelity trade-off that may need more investigation.

However, it is key to realize that the improved performance in the NHFEN metric at smaller scales using our method is likely due to fewer edges that need to be reproduced in the image, and does not suggest that reconstructions of increasingly sharper quality may be obtained by continually shrinking the scale of walnuts.

## VI. CONCLUSIONS AND FUTURE WORK

This paper developed a method to provide high quality reconstructions from extremely limited CBCT projections and scarce training data. The key features of our approach were the multi-stage approach of alternating between learning-based de-streaking and data consistency and the use of subvolume-based learning and shallower (adversarially trained) CNNs to combat over-fitting. In the future, we will focus on extending our method to dynamic imaging applications in CT, as well as being able to jointly segment and reconstruct three-dimensional objects. Towards this end, we are also interested in finding a better metric than the mean absolute error to assess the fidelity and quality of our reconstructions for specific tasks.

## REFERENCES

- [1] L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone beam algorithm," *J. Opt. Soc. Amer. A*, vol. 1, no. 6, pp. 612–619, Jun. 1984.
- [2] J.-B. Thibault, K. Sauer, C. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multi-slice helical CT," *Med. Phys.*, vol. 34, no. 11, pp. 4526–4544, Nov. 2007.
- [3] J. Xu and B. M. W. Tsui, "Interior and sparse-view image reconstruction using a mixed region and voxel-based ML-EM algorithm," *IEEE Trans. Nuc. Sci.*, vol. 59, no. 5, pp. 1997–2007, Oct. 2012.
- [4] J. H. Cho and J. A. Fessler, "Regularization designs for uniform spatial resolution and noise properties in statistical image reconstruction for 3D X-ray CT," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 678–689, Feb. 2015.
- [5] G. Wang, "A perspective on deep imaging," *IEEE Access*, vol. 4, pp. 8914–8924, 2016.
- [6] I. Y. Chun, X. Zheng, Y. Long, and J. A. Fessler, "Sparse-view X-ray CT reconstruction using  $\ell_1$  regularization with learned sparsifying transform," in *Proc. Int. Meeting Fully 3D Image Reconstruction Radiol. Nucl. Med.*, 2017, pp. 115–119.
- [7] I. A. Elbakri and J. A. Fessler, "Statistical image reconstruction for polyenergetic X-ray computed tomography," *IEEE Trans. Med. Imag.*, vol. 21, no. 2, pp. 89–99, Feb. 2002.
- [8] K. Sauer and C. Bouman, "A local update strategy for iterative reconstruction from projections," *IEEE Trans. Signal Process.*, vol. 41, no. 2, pp. 534–548, Feb. 1993.
- [9] J. A. Fessler, "Statistical image reconstruction methods for transmission tomography," in *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis*, M. Sonka and J. M. Fitzpatrick, Eds., Bellingham, WA, USA: SPIE, 2000, pp. 1–70.
- [10] J.-B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh, "A recursive filter for noise reduction in statistical iterative tomographic imaging," *Proc. SPIE*, vol. 6065, 2006, Art. no. 60650X.
- [11] A. R. De Pierro, "On the relation between the ISRA and the EM algorithm for positron emission tomography," *IEEE Trans. Med. Imag.*, vol. 12, no. 2, pp. 328–333, Jun. 1993.
- [12] S. Ravishankar, J. C. Ye, and J. A. Fessler, "Image reconstruction: From sparsity to data-adaptive methods and machine learning," *Proc. IEEE*, vol. 108, no. 1, pp. 86–109, Jan. 2020.
- [13] S. Ahn et al., "Quantitative comparison of OSEM and penalized likelihood image reconstruction using relative difference penalties for clinical PET," *Phys. Med. Biol.*, vol. 60, no. 15, pp. 5733–5752, Aug. 2015.
- [14] W. Yu, C. Wang, X. Nie, M. Huang, and L. Wu, "Image reconstruction for few-view computed tomography based on  $\ell_0$  sparse regularization," *Procedia Comput. Sci.*, vol. 107, pp. 808–813, 2017.
- [15] C. Xu, B. Yang, F. Guo, W. Zheng, and P. Poignet, "Sparse-view CBCT reconstruction via weighted Schatten P-norm minimization," *Opt. Exp.*, vol. 28, no. 24, pp. 35469–35482, Nov. 2020.
- [16] X. Zhang, Y. Zhou, W. Zhang, J. Sun, and J. Zhao, "Adaptive prior patch size based sparse-view CT reconstruction algorithm," in *Proc. IEEE 17th Int. Symp. Biomed. Imag.*, 2020, pp. 1–4.
- [17] Y. Kim, H. Kudo, K. Chigita, and S. Lian, "Image reconstruction in sparse-view CT using improved nonlocal total variation regularization," *Proc. SPIE*, vol. 11113, 2019, pp. 329–337.
- [18] R. Anirudh, H. Kim, J. J. Thiagarajan, K. A. Mohan, K. Champley, and T. Bremer, "Lose the views: Limited angle CT reconstruction via implicit sinogram completion," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6343–6352.
- [19] Y. Han, J. Kang, and J. C. Ye, "Deep learning reconstruction for 9-view dual energy CT baggage scanner," in *Proc. 5th Int. Meeting Image Formation X-ray CT*, 2018, pp. 407–410.
- [20] H. Kim, R. Anirudh, K. A. Mohan, and K. Champley, "Extreme few-view CT reconstruction using deep inference," 2019, *arXiv:1910.05375*.
- [21] E. Y. Sidky, K.-M. Kao, and X. Pan, "Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT," *J. X-Ray Sci. Technol.*, vol. 14, no. 2, pp. 119–139, 2006. [Online]. Available: <http://iospress.metapress.com/content/ljduv1cll3f9e2br/> (not free)
- [22] H. Yu and G. Wang, "Compressed sensing based interior tomography," *Phys. Med. Biol.*, vol. 54, no. 9, pp. 2791–2806, May 2009.
- [23] J. Bian et al., "Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam CT," *Phys. Med. Biol.*, vol. 55, no. 22, Nov. 2010, Art. no. 6575.
- [24] S. Ramani and J. A. Fessler, "A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction," *IEEE Trans. Med. Imag.*, vol. 31, no. 3, pp. 677–688, Mar. 2012.
- [25] G. T. Herman and R. Davidi, "Image reconstruction from a small number of projections," *Inverse Problems*, vol. 24, no. 4, Aug. 2008, Art. no. 045011.
- [26] G.-H. Chen, J. Tang, and S. Leng, "Prior image constrained compressed sensing (PICCS): A method to accurately reconstruct dynamic CT images from highly undersampled projection data sets," *Med. Phys.*, vol. 35, no. 2, pp. 660–663, Feb. 2008.
- [27] L. Pfister and Y. Bresler, "Model-based iterative tomographic reconstruction with adaptive sparsifying transforms," *Proc. SPIE*, vol. 9020, 2014, Art. no. 90200H.
- [28] C. Zhang, T. Zhang, M. Li, C. Peng, Z. Liu, and J. Zheng, "Low-dose CT reconstruction via L1 dictionary learning regularization using iteratively reweighted least-squares," *Biomed. Eng. Online*, vol. 15, no. 1, pp. 1–21, 2016.
- [29] X. Zheng, S. Ravishankar, Y. Long, and J. A. Fessler, "PWLS-ULTRA: An efficient clustering and learning-based approach for low-dose 3D CT image reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1498–510, Jun. 2018.
- [30] H. Chen et al., "LEARN: Learned experts' Assessment-based reconstruction network for sparse-data CT," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1333–1347, Jun. 2018.
- [31] D. Wu, K. Kim, G. E. Fakhri, and Q. Li, "Iterative low-dose CT reconstruction with priors trained by artificial neural network," *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2479–2486, Dec. 2017.
- [32] I. Y. Chun, H. Lim, Z. Huang, and J. A. Fessler, "Fast and convergent iterative image recovery using trained convolutional neural networks," in *Proc. IEEE Allerton Conf. Commun., Control, Comput.*, 2018, pp. 155–159.
- [33] I. Y. Chun, Z. Huang, H. Lim, and J. A. Fessler, "Momentum-Net: Fast and convergent iterative neural network for inverse problems," 2019. [Online]. Available: <http://arxiv.org/abs/1907.11818>
- [34] B. Zhou, S. K. Zhou, J. S. Duncan, and C. Liu, "Limited view tomographic reconstruction using a cascaded residual dense spatial-channel attention network with projection data fidelity layer," *IEEE Trans. Med. Imag.*, vol. 40, no. 7, pp. 1792–1804, Jul. 2021.
- [35] H. Gupta, K. H. Jin, H. Q. Nguyen, M. T. McCann, and M. Unser, "CNN-based projected gradient descent for consistent CT image reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1440–1453, Jun. 2018.
- [36] Y. Huang, A. Preuhs, G. Lauritsch, M. Manhart, X. Huang, and A. Maier, "Data consistent artifact reduction for limited angle tomography with deep learning prior," in *Proc. Int. Workshop Mach. Learn. Med. Image Reconstruction*, Cham, Switzerland: Springer, 2019, pp. 101–112.

- [37] H. Chen et al., "Low-dose CT with a residual encoder-decoder convolutional neural network," *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2524–2535, Dec. 2017.
- [38] X. Yang et al., "Low-dose X-ray tomography through a deep convolutional neural network," *Sci. Rep.*, vol. 8, no. 1, pp. 1–13, 2018.
- [39] S. Ye, Z. Li, M. T. McCann, Y. Long, and S. Ravishankar, "Unified supervised-unsupervised (SUPER) learning for X-ray CT image reconstruction," *IEEE Trans. Med. Imag.*, vol. 40, no. 11, pp. 2986–3001, Nov. 2021.
- [40] A. Lahiri, G. Wang, S. Ravishankar, and J. A. Fessler, "Blind primed supervised (BLIPS) learning for MR image reconstruction," *IEEE Trans. Med. Imag.*, vol. 40, no. 11, pp. 3113–3124, Nov. 2021.
- [41] G. Ongie, A. Jalal, C. A. Metzler, R. G. Baraniuk, A. G. Dimakis, and R. Willett, "Deep learning techniques for inverse problems in imaging," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 1, pp. 39–56, May 2020.
- [42] M. Hossain et al., "High-precision inversion of dynamic radiography using hydrodynamic features," 2021. [Online]. Available: <http://arxiv.org/abs/2112.01627>
- [43] E. T. Quinto, "Exterior and limited-angle tomography in non-destructive evaluation," *Inverse Problems*, vol. 14, no. 2, 1998, Art. no. 339.
- [44] S. Guan, A. A. Khan, S. Sikdar, and P. V. Chitnis, "Limited-view and sparse photoacoustic tomography for neuroimaging with deep learning," *Sci. Rep.*, vol. 10, no. 1, pp. 1–12, 2020.
- [45] M. Xiang, "Deep learning-based reconstruction of volumetric CT images of vertebrae from a single view X-ray image" Ph.D. dissertation, Univ. Wisconsin Milwaukee, Milwaukee, WI, USA, 2020. [Online]. Available: <https://dc.uwm.edu/etd/2629>
- [46] E. Eulig et al., "Deep learning-based reconstruction of interventional tools and devices from four X-ray projections for tomographic interventional guidance," *Med. Phys.*, vol. 48, no. 10, pp. 5837–5850, 2021.
- [47] X. Ying, H. Guo, K. Ma, J. Wu, Z. Weng, and Y. Zheng, "X2CT-GAN: Reconstructing CT from biplanar X-rays with generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 10611–10620.
- [48] J. C. Montoya, C. Zhang, Y. Li, K. Li, and G.-H. Chen, "Reconstruction of three-dimensional tomographic patient models for radiation dose modulation in CT from two scout views using deep learning," *Med. Phys.*, vol. 49, no. 2, pp. 901–916, 2022.
- [49] A. Ziabari, D. H. Ye, S. Srivastava, K. D. Sauer, J. Thibault, and C. A. Bouman, "2.5D deep learning for CT image reconstruction using a multi-GPU implementation," in *Proc. IEEE 52nd Asilomar Conf. Signals, Syst., Comput.*, 2018, pp. 2044–2049.
- [50] H. Lim, I. Y. Chun, Y. K. Dewaraja, and J. A. Fessler, "Improved low-count quantitative PET reconstruction with a variational neural network," *IEEE Trans. Med. Imag.*, vol. 39, no. 11, pp. 3512–3522, Nov. 2020.
- [51] G. Corda-D'Incan, J. A. Schnabel, and A. J. Reader, "Memory-efficient training for fully unrolled deep learned PET image reconstruction with iteration-dependent targets," *IEEE Trans. Radiat. Plasma Med. Sci.*, vol. 6, no. 5, pp. 552–563, May 2022.
- [52] M. Magnusson, P.-E. Danielsson, and J. Sunnegardh, "Handling of long objects in iterative improvement of nonexact reconstruction in helical cone-beam CT," *IEEE Trans. Med. Imag.*, vol. 25, no. 7, pp. 935–940, Jul. 2006.
- [53] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 18–44, Mar. 2021.
- [54] H. Nien and J. A. Fessler, "Relaxed linearized algorithm for faster X-ray CT image reconstruction," in *Proc. Int. Meeting Fully 3D Image Reconstruction Radiol. Nuc. Med.*, 2015, pp. 260–263. [Online]. Available: <http://web.eecs.umich.edu/fessler/papers/files/proc/15/web/nien-15-rla.pdf>
- [55] I. J. Goodfellow et al., "Generative adversarial networks," 2014. [Online]. Available: <http://arxiv.org/abs/1406.2661>
- [56] Y. Long, J. A. Fessler, and J. M. Balter, "3D forward and back-projection for X-ray CT using separable footprints," *IEEE Trans. Med. Imag.*, vol. 29, no. 11, pp. 1839–1850, Nov. 2010.
- [57] H. D. Sarkissian, F. Lucka, M. Eijnatten, G. Colacicco, S. B. Coban, and K. J. Batenburg, "A cone-beam X-ray computed tomography data collection designed for machine learning," *Sci. Data*, vol. 6, no. 1, pp. 1–8, 2019.
- [58] J. A. Fessler, "Michigan image reconstruction toolbox (MIRT) for matlab," 2016. [Online]. Available: <http://web.eecs.umich.edu/fessler/irt/irt>
- [59] Z. Huang, M. Klasky, T. Wilcox, and S. Ravishankar, "Physics-driven learning of wasserstein GAN for density reconstruction in dynamic tomography," *Appl. Opt.*, vol. 61, no. 10, pp. 2805–2817, Apr. 2022. [Online]. Available: <http://opg.optica.org/ao/abstract.cfm?URI=ao-61-10-2805>
- [60] E. S. Hertel Jr and G. I. Kerley, "CTH EOS package: Introductory tutorial," Sandia Nat. Lab., Tech. Rep. SAND98-0945, 1998.
- [61] S. Ravishankar and Y. Bresler, "MR image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1028–1041, May 2011.



**Anish Lahiri** (Member, IEEE) received the B.Tech degree in electronics and communication engineering from the West Bengal University of Technology, Kolkata, India, in 2014, and the Ph.D. degree in electrical and computer engineering from the University of Michigan, Ann Arbor, MI, USA. He was a Postdoctoral Research Fellow of electrical and computer engineering with the University of Michigan. His research interests include machine learning for inverse problems, model-based medical image reconstruction, mathematical optimization for signal

processing, combining deep learning and model-based imaging for improved explainability, and generalization across signal processing applications.



**Gabriel Maliakal** received the B.Tech. from the National Institute of Technology, Surathkal, India, in 2015, and the M.S. degree in electrical engineering from Columbia University, New York, NY, USA, in 2017. He is currently working toward the Ph.D. degree (2nd year) with the Department of Computational Mathematics, Science and Engineering, Michigan State University, East Lansing, MI, USA, from Fall 2021. His research interest include machine learning and optimization for inverse problems in imaging particularly image reconstruction and segmentation from sparse measurements, and reinforcement learning.



**Marc L. Klasky** is currently a Senior Scientist with Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM, USA. He is developing new algorithms in tomographic analysis for the past 15 years to support analysis of images acquired with the Dual Axis Hydrodynamic Radiographic Test Facility. He has made major contributions in a number of diverse areas in computational physics including the development of inversion methods for spectral analysis of Compton spectrometers, scatter treatment in radiographic analysis, resistive magneto-hydrodynamics, corrosion science, computational chemistry, and dynamic tomography. His research interest include the development of machine learning algorithms to enable parameter estimation in 3-D hydrodynamics simulation codes using sparse tomographic projections and the development of reduced order models to describe the evolution of hydrodynamic features.



**Jeffrey A. Fessler** (Fellow, IEEE) received the BSEE degree from Purdue University, West Lafayette, IN, USA, in 1985, the MSEE degree in 1986, M.S. degree in statistics in 1989, and Ph.D. degree in electrical engineering in 1990 from Stanford University, Stanford, CA, USA. He is currently the William L. Root Professor with EECS, University of Michigan, Ann Arbor, MI, USA. From 1985 to 1988, he was the National Science Foundation Graduate Fellow with Stanford University. Since then, he has been with the University of Michigan. From 1991 to 1992, he was

the Department of Energy Alexander Hollaender Postdoctoral Fellow with the Division of Nuclear Medicine. From 1993 to 1995, he was an Assistant Professor of nuclear medicine and the bioengineering program. He is also a Professor with the Departments of Electrical Engineering and Computer Science, Radiology, and Biomedical Engineering. He was the recipient of the Francois Erbsmann Award for his IPMI93 presentation, Edward Hoffman Medical Imaging Scientist Award in 2013, and IEEE EMBS Technical Achievement Award in 2016. He was an Associate Editor for the IEEE TRANSACTIONS ON MEDICAL IMAGING, IEEE SIGNAL PROCESSING LETTERS, IEEE TRANSACTIONS ON IMAGE PROCESSING, IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING (T-CI), and is an Associate Editor for *SIAM Journal on Imaging Science* and Senior AE for T-CI. He has chaired the IEEE T-MI Steering Committee and the ISBI Steering Committee. He was Co-chair of the 1997 SPIE conference on Image Reconstruction and Restoration, Technical Program Co-Chair of the 2002 IEEE International Symposium on Biomedical Imaging (ISBI), and General Chair of ISBI 2007. His research interests include statistical aspects of imaging problems, and he has supervised doctoral research in PET, SPECT, X-ray CT, MRI, and optical imaging problems.



**Saiprasad Ravishankar** (Senior Member, IEEE) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology Madras, Chennai, India, in 2008, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Illinois at Urbana-Champaign, Urbana-Champaign, IL, USA, in 2010 and 2014, respectively. He is currently an Assistant Professor with the Departments of Computational Mathematics, Science and Engineering, and Biomedical Engineering, Michigan State University, Michigan, USA. He was

an Adjunct Lecturer and a Postdoctoral Research Associate with the University of Illinois at Urbana-Champaign. Since August 2015, he was a Postdoc with the Department of Electrical Engineering and Computer Science with the University of Michigan, and then a Postdoc Research Associate with the Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM, USA, from August 2018 to February 2019. His research interests include signal and image processing, biomedical and computational imaging, machine learning, inverse problems, and large-scale data processing and optimization. He was the recipient of the IEEE Signal Processing Society Young Author Best Paper Award in 2016, Best Student Paper Award at the IEEE International Symposium on Biomedical Imaging (ISBI) 2018 for the paper he co-authored, and other papers were award finalists at the IEEE International Workshop on Machine Learning for Signal Processing (MLSP) 2017 and ISBI 2020. He is currently the Member of the IEEE Computational Imaging Technical Committee. He has organized several special sessions or workshops on computational imaging and machine learning themes including at the Institute for Mathematics and its Applications (IMA), IEEE Image, Video, and Multidimensional Signal Processing (IVMSP) Workshop 2016, MLSP 2017, ISBI 2018, and the International Conference on Computer Vision (ICCV) 2019 and 2021.