

CoIL: Coordinate-Based Internal Learning for Tomographic Imaging

Yu Sun , Graduate Student Member, IEEE, Jiaming Liu , Graduate Student Member, IEEE, Mingyang Xie, Brendt Wohlberg , Senior Member, IEEE, and Ulugbek S. Kamilov , Senior Member, IEEE

Abstract—We propose *Coordinate-based Internal Learning (CoIL)* as a new deep-learning (DL) methodology for continuous representation of measurements. Unlike traditional DL methods that learn a mapping from the measurements to the desired image, CoIL trains a multilayer perceptron (MLP) to encode the complete measurement field by mapping the coordinates of the measurements to their responses. CoIL is a self-supervised method that requires no training examples besides the measurements of the test object itself. Once the MLP is trained, CoIL generates new measurements that can be used within most image reconstruction methods. We validate CoIL on sparse-view computed tomography using several widely-used reconstruction methods, including purely model-based methods and those based on DL. Our results demonstrate the ability of CoIL to consistently improve the performance of all the considered methods by providing high-fidelity measurement fields.

Index Terms—Inverse problems, image reconstruction, regularized inversion, plug-and-play priors, deep learning.

I. INTRODUCTION

THE PROBLEM of reconstructing an unknown image from a set of noisy measurements is fundamental to computational imaging. The task is traditionally formulated as an inverse problem and solved using model-based optimization by leveraging a forward model characterizing the imaging system and a regularizer imposing prior knowledge on the unknown image. There has been significant progress in developing sophisticated image priors, including those based on transform-domain sparsity, self-similarity, and dictionary learning [1]–[4].

Manuscript received February 10, 2021; revised July 18, 2021 and September 20, 2021; accepted October 24, 2021. Date of publication November 8, 2021; date of current version December 24, 2021. This work was supported in part by NSF Awards CCF-1813910 and CCF-2043134, and in part by the Laboratory Directed Research and Development Program of Los Alamos National Laboratory under Project 20200061DR. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Jinwei Gu. (Corresponding author: Ulugbek S. Kamilov.)

Yu Sun and Mingyang Xie are with the Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130 USA (e-mail: sun.yu@wustl.edu; ericxie@wustl.edu).

Jiaming Liu is with the Department of Electrical and Systems Engineering, Washington University in St. Louis, St. Louis, MO 63130 USA (e-mail: jiaming.liu@wustl.edu).

Brendt Wohlberg is with Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: brendt@ieee.org).

Ulugbek S. Kamilov is with the Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130 USA, and also with the Department of Electrical and Systems Engineering, Washington University in St. Louis, St. Louis, MO 63130 USA (e-mail: kamilov@wustl.edu).

Digital Object Identifier 10.1109/TCI.2021.3125564




Fig. 1. The conceptual illustration of CoIL in the context of sparse-view CT. A multilayer perceptron (MLP) is used to represent the full measurement field by learning to map the measurement coordinate (θ, l) to its response r . Visual examples compare the recovered images with and without CoIL for total variation (TV). CoIL is used to generate 360 views from the data consisting of 120 noisy views of 40 dB input SNR. The quantitative and visual results in this paper highlight the ability of CoIL to significantly improve the imaging quality for several widely-used image reconstruction methods.

There has been a considerable recent interest in *deep learning (DL)* based solutions to imaging inverse problems [5]–[8]. The traditional DL approach involves training a convolutional neural network (CNN) to directly perform a regularized inversion of the forward model by exploiting redundancies in a training dataset [9]–[12]. Model-based DL is an alternative to traditional DL that explicitly uses knowledge of the forward model by integrating a CNN into model-based optimization. Two widely-used approaches in this context are *plug-and-play priors (PnP)* [13] and *regularization by denoising (RED)* [14], which have been used with pre-trained deep denoisers to achieve excellent performance in a number of imaging tasks [15]–[25]. An alternative model-based DL approach is *deep unfolding*, which interprets the iterations of a model-based optimization algorithm as layers of a CNN and trains it end-to-end in a supervised fashion [26]–[36].

There has been a considerable amount of work on DL for imaging inverse problems, the unifying theme being that one can train a CNN over a dataset to represent a prior for an unknown image. In this paper, we take a fundamentally different approach by proposing a methodology for leveraging redundancy within the measurements of a single unknown image, thus requiring no training examples besides the test-input itself. Our

proposed *Coordinate-based Internal Learning (CoIL)* seeks to represent the full continuous measurement field by exploiting the internal information within the subsampled and noisy measurements. This single-shot scheme is beneficial for applications where fully-sampled measurements are unavailable. The core of CoIL is a *multilayer perceptron (MLP)* that maps the measurement coordinates to the corresponding sensor responses. The measurement coordinates are the parameters corresponding to the geometry of the imaging system that determine the response measured by the sensors. For example, in computed tomography (CT) two parameters characterizing the response are the view angle θ of the incoming ray beam and the spatial location l of the relevant detector on the sensor plane. By training an MLP on the coordinate-response pairs extracted from the measurements of a desired object, CoIL is able to build a continuous mapping from the coordinates to the sensor responses. Thus, the learned MLP corresponds to a neural representation of the full measurement field. By querying the MLP with the relevant coordinates, CoIL can generate the full field that can be used for image reconstruction. Fig. 1 provides a conceptual illustration of the CoIL methodology. Note that CoIL is not restricted to a specific image reconstruction method, but is compatible with a majority of methods including those based on model-based optimization or DL.

The main contributions of this paper are as follows:

- We propose CoIL as a novel imaging methodology that leverages coordinate-based neural representations for estimating high-fidelity measurement fields [37]–[39]. Unlike other recent coordinate-based methods that focus on representing the unknown object, CoIL represents the measurement field, which can be subsequently combined with other information sources during reconstruction.
- We propose a new MLP architecture that uses a linear mapping strategy for the input coordinates of the network. We observe that this input mapping strategy is effective for representing measurement fields and achieves better performance compared to *positional encoding* [37] and *random sampling* [40].
- We extensively validate our method in the context of sparse-view CT. We show that CoIL synergistically combines with a majority of widely-used methods by being able to generate high-fidelity full-view sinograms. In all our experiments the methods with CoIL consistently outperform the ones without it.

The rest of the paper is organized as follows. Section II reviews the background on inverse problems in computational imaging and coordinate-based learning. Section III introduces the technical details of the proposed framework and discusses integration of CoIL into several widely-used image reconstruction methods. Section IV presents the numerical validation of CoIL on sparse-view CT. Section V concludes the paper by discussing potential applications and limitations of CoIL.

II. BACKGROUND

In this section, we review background information related to CoIL. We introduce the imaging inverse problem and review several popular reconstruction methods. We also discuss sensor-domain DL models and the recent progress on internal learning.

A. The Inverse Problem in Imaging

Consider the linear measurement model

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}, \quad (1)$$

where the measurement operator $\mathbf{A} \in \mathbb{R}^{m \times n}$ characterizes the response of the imaging system and vector $\mathbf{e} \in \mathbb{R}^m$ represents the noise, which is often assumed to be additive white Gaussian (AWGN). The associated imaging inverse problem involves the reconstruction of the image $\mathbf{x} \in \mathbb{R}^n$ from the measurements $\mathbf{y} \in \mathbb{R}^m$. Due to ill-posedness, practical inverse problems are often formulated as regularized optimization

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}), \quad \text{with } f(\mathbf{x}) = g(\mathbf{x}) + h(\mathbf{x}), \quad (2)$$

where g is the data-fidelity term that quantifies the consistency of $\mathbf{x}$ with $\mathbf{y}$, and h is the regularizer that imposes some prior knowledge on $\mathbf{x}$. For instance, two widely used functions in the context of imaging are the ℓ_2 norm and total variation (TV)

$$g(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \quad \text{and} \quad h(\mathbf{x}) = \tau \|\mathbf{D}\mathbf{x}\|_1, \quad (3)$$

where $\tau > 0$ controls the regularization strength and $\mathbf{D}$ is the discrete gradient operator [1]. The nonsmoothness of the regularizer is a common occurrence in imaging, which precludes the use of the standard gradient descent algorithms.

The family of proximal methods are effective solvers for nonsmooth optimization problems of form (2). Two common algorithms are *fast iterative shrinkage/thresholding algorithm* (FISTA) [41] and *alternating direction method of multipliers* (ADMM) [42]–[44]. These algorithms rely on a mathematical concept known as the *proximal operator* [45], defined as

$$\text{prox}_{\mu h}(\mathbf{z}) := \arg \min_{\mathbf{x} \in \mathbb{R}^n} \left\{ \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 + \mu h(\mathbf{x}) \right\}, \quad (4)$$

to handle nonsmooth terms without differentiation. The parameter $\mu > 0$ in (4) balances the importance of the term h . Note that the proximal operator can be interpreted as a *maximum a posteriori* (MAP) denoiser for AWGN with variance μ .

B. Traditional Deep Learning Methods

DL has become very popular for imaging inverse problems [5]–[8] due to its excellent performance. Traditional DL methods first bring the measurements $\{\mathbf{y}_i\}_{i=1}^N$ to the image domain via backprojection and then use a deep CNN architecture, such as UNet [46], to map the resulting low-quality images $\{\tilde{\mathbf{x}}_i\}_{i=1}^N$ to high-quality images $\{\mathbf{x}_i\}_{i=1}^N$. Here, $N > 0$ denotes the total number of training examples. Typically, these CNNs are trained by minimizing a loss function

$$\ell(\psi) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\mathcal{F}_\psi(\tilde{\mathbf{x}}_i), \mathbf{x}_i), \quad (5)$$

where $\mathcal{F}_\psi$ denotes the network parametrized by ψ , and function $\mathcal{L}$ quantifies the discrepancy between $\mathcal{F}_\psi(\tilde{\mathbf{x}}_i)$ and $\mathbf{x}_i$. Popular choices for $\mathcal{L}$ include the ℓ_1 and ℓ_2 norms. Some other models consider different schemes that directly map $\{\mathbf{y}_i\}$ to reconstructed images $\{\mathbf{x}_i\}$. These methods often adopt hybrid CNN architectures that contain fully connected layers for learning either an approximation of the inverse $(\mathbf{A}\mathbf{A}^\top)^{-1}$ [47] or an inversion to some implicit image manifolds [48]. Nevertheless,

traditional DL methods do not explicitly impose consistency with respect to the forward model during image reconstruction.

C. Deep Denoising Priors

The family of denoising-driven approaches represents an alternative to traditional DL by combining iterative model-based algorithms with deep denoisers as priors. These methods draw inspiration from the equivalence between proximal operator and image denoiser. One popular framework is PnP, which generalizes proximal methods, such as FISTA and ADMM, by replacing the proximal operator with an arbitrary AWGN denoiser $\mathcal{D}_\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^n$, with $\sigma > 0$ controlling the denoising strength. This simple replacement enables PnP to use advanced denoisers, including those based on CNNs [49]–[51], for regularizing the inverse problem. The PnP algorithms have been shown to be effective in various imaging applications [52]–[56]. However, $\mathcal{D}_\sigma$ may not correspond to any explicit h in (2), in which case PnP loses its interpretation as optimization. Theoretical analysis of PnP has also been proposed [20]–[24], [57]–[59].

RED [14] is a related framework that uses the operator

$$\mathcal{H}(\mathbf{x}) = \tau(\mathbf{x} - \mathcal{D}_\sigma(\mathbf{x})), \quad (6)$$

within many kinds of iterative algorithms [14], [18], [25], [60]. RED with deep denoisers has been reported to be effective in image super-resolution [61], phase retrieval [62], and tomographic imaging [18]. It has been shown that when the denoiser $\mathcal{D}_\sigma$ is locally homogeneous and has a symmetric Jacobian [14], [63], $\mathcal{H}(\mathbf{x})$ corresponds to the gradient of the following regularizer

$$h(\mathbf{x}) = \frac{\tau}{2} \mathbf{x}^T (\mathbf{x} - \mathcal{D}_\sigma(\mathbf{x})). \quad (7)$$

RED has recently been theoretically analyzed for general denoisers that may not be associated with any explicit regularizer [18], [25], [60], [64].

D. Other Model-Based Deep Learning

Deep unfolding is another widely used model-based DL methodology, originally proposed in [65] for sparse coding. The central idea of deep unfolding is that one can unfold an iterative algorithm and train it end-to-end as a deep neural network [26]–[29]. This enables integration of the physical information into the architecture in the form of data-consistency blocks that are combined with trainable CNN regularizers [30]–[34]. By training the corresponding model-based network end-to-end, one obtains a regularizer optimized for a specific problem. Excellent performance of deep unfolding has been reported in a number of imaging applications [31], [66], and recent work has addressed the computational and memory complexity of training such networks [35], [36].

Another family of DL methods has used generative adversarial networks (GAN) for regularizing inverse problems [67]–[69]

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^z} \|\mathcal{A}\mathcal{G}(\mathbf{z}) - \mathbf{y}\|_2^2 \quad \text{and} \quad \hat{\mathbf{x}} = \mathcal{G}(\hat{\mathbf{z}}), \quad (8)$$

where $\mathcal{G}$ is a pre-trained GAN, and $\mathbf{z} \in \mathbb{R}^z$ is the encoding in the latent space. The optimization in (8) implicitly imposes the regularization by restricting the solution $\hat{\mathbf{x}}$ to the range of a GAN $\mathcal{G}$. By searching for the optimal encoding $\hat{\mathbf{z}}$, one can obtain an estimate $\hat{\mathbf{x}} = \mathcal{G}(\hat{\mathbf{z}})$ in the domain defined by GAN that

has the smallest distance to the true $\mathbf{x}$. The recovery properties under GANs have been analyzed in the context of compressive sensing [67]–[69]. For more information on this topic, see the recent review [70].

It is worth pointing out that CoIL is complementary to all these prior works, since it seeks to learn the measurement field given measurements of a single unknown image. As shown in Section IV, CoIL can be naturally combined with the majority of reconstruction algorithms used in computational imaging.

E. Deep Learning for Measurement Synthesis

An interest in developing DL-based approaches for measurement synthesis has recently emerged. An end-to-end scheme that used in image super-resolution is commonly adopted to first linearly interpolate the measurements to the same scale as that of the target measurements and then use a CNN to map the intermediate output to the final refined results [71], [72]. GANs have been employed to synthesize missing measurements that are corrupted by the metal artifacts [73]–[76]. The effectiveness of these approaches has been shown in different imaging modalities [77]–[79]. Nevertheless, most deep synthesis methods require a dataset of fully-sampled measurements for supervised training. A scan-specific CNN model that avoids training on a large dataset has been recently proposed [80], but still requires fully-sampled measurements of the object as ground truth. CoIL is fundamentally different from the existing methods for measurement synthesis since it learns a representation of the full measurement field from the measurements of an unknown object without any ground truth.

F. Deep Internal Learning

Deep internal learning explores the internal information of the test signal for learning a neural network prior without using any external data. One successful approach is to exploit the patch-wise similarity within images, leading to significant results for spatial and temporal super-resolution [81], [82]. Another widely adopted approach is *deep image prior (DIP)*, which optimizes a CNN to parameterize the reconstructed image [83]–[85]. *Coordinate-based neural representation* is a recent alternative that encodes a spatial field into the weights of a MLP, which is trained to map coordinates (e.g., (x, y, z)) to the pixel values (e.g., $[0, 1]$). It has been quite successful in computer vision and graphics, but has not been widely explored in computational imaging. The coordinate-based MLPs have been used to represent images [40], [86], scenes [37], [38], and three-dimensional (3D) shapes [37], [38], [87]–[89]. *Neural radiance field (NeRF)* [37] is a recent model that has significantly improved the representation power of MLPs by first expanding the input coordinates into a Fourier spectrum (see Section III-A for detailed discussion). Its formulation has been adapted for improving scene resolutions [39], dealing with multiple lightening conditions [90], removing occluders [38], and handling small deformations [91]. Although there have been some early attempts in representing medical images [40], the usage of coordinate-based representation has never been explored for representing measurement fields in tomographic imaging, where the measurements are not the direct samples of the signal of interest but rather the aggregations along the sampling rays.




Fig. 2. Illustration of the CoIL workflow for a tomographic imaging system with free parameters $\mathbf{v} \in \mathbb{R}^v$. First, a set of $N > 0$ measurements are acquired by the system under different realization of $\mathbf{v}$. Then, the coordinate-response pairs $\{(\mathbf{v}_i, r_i)\}_{i=1}^N$ are used to train a coordinate-based MLP $\mathcal{M}_\phi : \mathbf{v} \rightarrow r$ for encoding the full measurement field. Once the training is finished, the encoded field is extracted from $\mathcal{M}_\phi$ with an arbitrary resolution by querying the relevant coordinates. In the final stage, the CoIL field and the actual measurements are jointly used for image reconstruction using a user-defined method.




Fig. 3. Visualization of the coordinate-based MLP used in the CoIL methodology. The network $\mathcal{M}_\phi = \mathcal{N}_\phi \circ \gamma(\mathbf{v})$ is a concatenation of a single Fourier feature mapping (FFM) layer $\gamma(\mathbf{v})$ and a conventional MLP $\mathcal{N}_\phi$. As training on example pairs $\{(\mathbf{v}_i, r_i)\}_{i=1}^N$, $\mathcal{M}_\phi$ is able to learn a continuous mapping from a coordinate to its response. Hence, $\mathcal{M}_\phi$ becomes an implicit neural representation of the full-measurement field.

This work addresses this gap by proposing CoIL as a novel image reconstruction methodology that leverages an MLP to represent the measurement fields.

III. COORDINATE-BASED INTERNAL LEARNING

In this section, we present the details of the CoIL methodology that leverages the power of coordinate-based learning for addressing imaging inverse problems. Fig. 2 illustrates the general workflow of CoIL for a given imaging system. We first explain the proposed MLP network and then discuss its integration into several common image reconstruction methods.

A. Measurement-Field Encoding With MLP

The coordinate-based MLP is the central component of CoIL. The network can be expressed as

$$\mathcal{M}_\phi : \mathbf{v} \rightarrow r \quad \text{with} \quad \mathbf{v} \in \mathbb{R}^v, r \in \mathbb{R},$$

where $\mathbf{v}$ represents the coordinate in the given imaging system, and r is the corresponding sensor response. The network can be conceptually separated into two parts, where the first part is a single *Fourier feature mapping (FFM)* layer $\gamma(\mathbf{v})$ that is




Fig. 4. Eight 512×512 images from the scans of two patients in the AAPM human phantom dataset [99] were used for testing.

pre-defined before training, and the second part is a standard MLP $\mathcal{N}_\phi : \gamma(\mathbf{v}) \rightarrow r$ whose parameters ϕ needs to be optimized. A visual illustration of the complete network architecture is provided in Fig. 3. While the numerical studies presented in this paper focus on CT, CoIL is also applicable to other imaging modalities by simply changing the coordinate-response pairs in the MLP representation. For example, one can potentially




Fig. 5. Illustration of the benefit of including the Fourier feature mapping (FFM) layer into CoIL. We plot sinograms and their FBP reconstructions in the first and second row, respectively. The proposed FFM in CoIL is compared against *No FFM* strategy (which does not have any FFM layer), positional encoding (*Pos Enc*) [37], and random Gaussian sampling (*Random*) [40]. The four MLPs are used to generate 360 views from the $P = 120$ projections with $I = 40$ dB noise. Both sinograms and images are labeled with the SNR values with respect to the ground truth shown in the right-most column. The bounding boxes highlight areas of significant visual difference. This comparison shows the benefit of using the FFM layer with linear spacing in the Fourier space.

integrate CoIL into optical diffraction tomography (ODT) [92]–[94] by letting $\mathbf{v}$ denote the sensor location and the angle of the incident light and letting $\mathbf{r}$ have two elements corresponding to the real and imaginary components of the light-field.

1) *Fourier Feature Mapping*: Although neural networks are known to be universal function approximators [95], it has been found that standard MLPs perform poorly in representing high-frequency variations [37], [96]. In our experiments, we also observed such issues when we directly applied $\mathcal{N}_\phi$ to learning the mapping $\mathbf{v} \rightarrow \mathbf{r}$ (see *No FFM* in Fig. 5). In order to overcome the limitations of standard MLPs, we include the FFM layer to expand the input coordinate $\mathbf{v}$ as a combination of different frequency components

$$\gamma(\mathbf{v}) = \begin{pmatrix} \sin(k_1\pi\mathbf{v}), \cos(k_1\pi\mathbf{v}) \\ \vdots \\ \sin(k_L\pi\mathbf{v}), \cos(k_L\pi\mathbf{v}) \end{pmatrix}, \quad (9)$$

where $\sin$ and $\cos$ compute element-wise sinusoidal and cosinusoidal values, respectively, and $\{k_i\}_{i=1}^L$ determine the frequencies in the mapping. The FFM layer pre-defines the frequency components so that the network $\mathcal{N}_\phi$ can actively select the ones that are the most useful for encoding sensor responses by learning the weights in the first layer. By manipulating the coefficients k_i and total number of components $L > 0$, we can explicitly control the expanded spectrum and thus impose some implicit regularization. The FFM layer was first introduced in NeRF as *positional encoding* of spatial coordinates [37], and a follow-up work [40] has further explored its functionality by using a concept known as neural tangent kernels [97]. The original formulation of $\gamma(\mathbf{v})$ in [37] sets k_i as an exponential function $k_i = 2^{i-1}$ with $L = 10$. We discovered that the presence of very high-frequency components leads to the overfitting of the MLP to the noise in the measurements. We thus adopted a linear sampling $k_i = \pi i/2$ in the Fourier space that leads to a higher number of frequency components in the low-frequency regions. Our empirical results show that our strategy can effectively improve $\mathcal{M}_\phi$ in representing high-frequency variations and preventing overfitting to noise (see Fig. 5 for examples).

We also compare our linear-expansion strategy with the recent random-projection mapping [40]

$$\gamma(\mathbf{v}) = (\cos(2\pi\mathbf{B}\mathbf{v}), \sin(2\pi\mathbf{B}\mathbf{v})), \quad (10)$$

where $\mathbf{B} \in \mathbb{R}^{d \times v}$ is a random matrix with each entry sampled from i.i.d. Gaussian distribution $N(0, \sigma^2)$. Numerical evaluations in Section IV-B show that our formulation achieves better performance for higher noise level and fewer projections.

2) *MLP Architecture*: The network implementing $\mathcal{N}_\phi$ is composed of 17 fully-connected (FC) layers. The first 16 layers have 256 hidden neurons whose outputs are activated by the rectified linear unit (ReLU), while the last layer has 128 hidden neurons without any activation. We include 7 skip connections after every even-numbered (fewer than 16) FC layer to concatenate the original input of $\mathcal{N}_\phi$ with the intermediate outputs. The use of the skip connections has been shown to be beneficial for fast training [87] and better accuracy [88]. Note that although $\mathcal{M}_\phi$ is a fully connected network, its input corresponds to a single coordinate, which enables element-wise processing of all the measurements. Additional evaluations of several other architectures that can be used as alternatives to CoIL are provided in the supplement.

3) *Additional Implementation Details*: CoIL trains a separate MLP to represent the full measurement field for each test objects. This means that the training pairs $\{(\mathbf{v}_i, \mathbf{r}_i)\}_{i=1}^N$ are obtained by extracting the measurements of the test object only, without any training dataset. The network $\mathcal{M}_\phi$ is trained by using Adam [98] to minimize the standard ℓ_2 -norm loss

$$\ell(\psi) = \frac{1}{N} \sum_{i=1}^N \|\mathcal{M}_\phi(\mathbf{v}_i) - \mathbf{r}_i\|_2^2. \quad (11)$$

We implement a decreasing learning rate, which decays exponentially as the training epoch increases, to smooth our optimization. Although $\mathcal{M}_\phi$ is a MLP, the network has a significantly smaller size (≈ 4.2 MB on disk) compared to the standard UNet architecture (≈ 108 MB on disk) used in many DL-based models.

B. Image Reconstruction in CoIL

After training, one can generate an arbitrary number of measurements by querying $\mathcal{M}_\phi$ using the relevant coordinates. We will refer to the corresponding measurement field as *CoIL field*. The CoIL field can be easily integrated into the majority of image reconstruction methods. Here, we discuss the integration of CoIL into four widely-used methods.

1) *Linear Reconstruction*: Filtered backprojection (FBP) is a classic method for bringing the measurements into the image domain [100]. Since the CoIL field is essentially a set of measurements, we can directly feed the field as input to FBP for image reconstruction. A slightly different way to apply FBP is to form a combined input that includes both the original measurements and those generated by CoIL. The key benefit of the latter approach is that it directly uses the real data while also complementing it with CoIL measurements.

2) *Model-Based Optimization*: Model-based methods reconstruct images by solving optimization problems of form (2). The CoIL field can be incorporated into the formulation by including an additional “data-fidelity” term $\tilde{g}$ in the objective

$$f(\mathbf{x}) = \underbrace{(1 - \alpha)g(\mathbf{x}) + \alpha\tilde{g}(\mathbf{x})}_{\text{New data-fidelity}} + h(\mathbf{x}). \quad (12)$$

Here, the parameter $0 \leq \alpha \leq 1$ controls the tradeoff between the real data and the generated field. In practice, we can fine-tune the value of α to obtain a good balance between two terms. For example, consider the least-squares function

$$\tilde{g}(\mathbf{x}) = \frac{1}{2} \|\tilde{\mathbf{A}}\mathbf{x} - \mathcal{M}_\phi(\tilde{\mathbf{v}})\|_2^2, \quad (13)$$

where $\tilde{\mathbf{A}} \in \mathbb{R}^{m \times n}$ corresponds to the sampling geometry of the CoIL field, $\tilde{\mathbf{v}}$ represents all the query coordinates for the trained MLP $\mathcal{M}_\phi(\tilde{\mathbf{v}})$. Since the network is pre-trained, one can directly use any existing image regularizer and solve the optimization problem with a standard iterative algorithm, such as FISTA or ADMM.

3) *End-to-End DL Models*: As reviewed in Section II-B, most end-to-end DL models are trained to directly map the low-quality images $\{\tilde{\mathbf{x}}_i\}_{i=1}^N$ to the high-quality images $\{\mathbf{x}_i\}_{i=1}^N$, making them vulnerable to unseen outliers. For example, this adversely influences the performance of DL, when there is a mismatch between training and testing angles. CoIL can be used to address this issue by generating the measurement field corresponding to the same subsampling rate as the measurements used for training the DL model

$$\tilde{\mathbf{x}} = \mathcal{F}_\psi(\text{FBP}(\mathcal{M}_\phi(\tilde{\mathbf{v}}))), \quad (14)$$

where $\mathcal{F}_\psi$ denotes the pre-trained CNN. Alternatively, one can include the original test image in the input by averaging the $\tilde{\mathbf{x}}$ and $\text{FBP}(\mathcal{M}_\phi(\tilde{\mathbf{v}}))$ using a weight α

$$\hat{\mathbf{x}} = \mathcal{F}_\psi(\underbrace{(1 - \alpha)\tilde{\mathbf{x}} + \alpha\text{FBP}(\mathcal{M}_\phi(\tilde{\mathbf{v}}))}_{\text{Joint input}}). \quad (15)$$

This approach enables the usage of the learned measurements by MLP with the true measurements from the imaging system. Our results in Section IV show that this CoIL-based strategy achieves better results than training a DL model directly on the measurements.

4) *Denoising-Driven Approches*: PnP/RED algorithms can be interpreted as extensions of model-based algorithms balancing consistency with the measurements against deep denoising priors [22], [25]. Consider gradient-based RED (GM-RED)

$$\mathbf{x}^+ \leftarrow \mathbf{x} - \gamma [\nabla g(\mathbf{x}) + \tau(\mathbf{x} - \mathcal{D}_\sigma(\mathbf{x}))] \quad (16)$$

where $\gamma > 0$ is the stepsize, and ∇g is the gradient of the data-fidelity term. Similar to the modification of model-based optimization, one straightforward way to integrate CoIL into GM-RED is to include the gradient of $\tilde{g}$ as an extra term

$$\mathbf{x}^+ \leftarrow \mathbf{x} - \gamma \underbrace{[(1 - \alpha)\nabla g(\mathbf{x}) + \alpha\nabla\tilde{g}(\mathbf{x})]}_{\text{New data enforcement}} + \tau(\mathbf{x} - \mathcal{D}_\sigma(\mathbf{x})),$$

where the new update ensures the consistency with the real measurements as well as the field generated by CoIL, with α controlling the relative weighting. This idea is also applicable to PnP, for example, by integrating CoIL within PnP-FISTA

$$\mathbf{x}^+ \leftarrow \mathcal{D}_\sigma(\mathbf{s} - \gamma[(1 - \alpha)\nabla g(\mathbf{s}) + \alpha\nabla\tilde{g}(\mathbf{s})]) \quad (17a)$$

$$\mathbf{s}^+ \leftarrow \mathbf{x}^+ + ((q^+ - 1)/q^+)(\mathbf{x}^+ - \mathbf{x}) \quad (17b)$$

where the acceleration parameter $q > 0$ is updated as

$$q^+ \leftarrow \frac{1}{2} \left(1 + \sqrt{1 + 4q^2} \right).$$

In the next section, we will provide results highlighting the performance of CoIL in the context of all these algorithms.

IV. NUMERICAL VALIDATIONS

We numerically validate CoIL in the context of sparse-view CT. We first substantiate the effectiveness of the proposed form of FFM, and then demonstrate the benefits of using CoIL for image reconstruction. We consider four reconstruction methods, *FBP*, *FISTA-TV*, *GM-RED*, and *FBP-UNet*. FBP-UNet refers to the end-to-end model proposed in [9] and FISTA-TV refers to the TV regularized inversion implemented using FISTA. We integrate CoIL into these algorithm by including the parameter α as discussed in Section III-B.

A. Sparse View CT and Experimental Setup

Sparse view X-ray CT is an imaging modality that aims to reconstruct a tomographic image from few projections. In medical applications, it can significantly reduce the radiation dose and hence reduce the risk of radiation exposure. The reconstruction task in CT can be formulated as the linear inverse problem of form (1). In our simulations, we adopt the parallel beam geometry with a measurement operator $\mathbf{A}$ corresponding to the Radon transform.

We consider the experimental setting where the X-ray beam is emitted from the view angle $\theta \in [0, \pi]$ and its radiation attenuation is recorded by the detectors at different (normalized) sensor-plane locations $l \in [0, 1]$. Thus, the MLP is trained to map the location and angle (l, θ) to the corresponding response r . Fig. 4 visualizes eight 512×512 test images used in all experiments. These images are selected from the scans of two patients in the APPM human phantom dataset¹ [99], while the scans of

¹The 2016 NIH-AAPM-Mayo Clinic Low Dose CT Grand Challenge.

TABLE I
THE AVERAGE SNR OF THE SINOGRAMS GENERATED BY NO FFM, POS ENC,
GAUSSIAN AND CoIL IN THE SCENARIOS CORRESPONDING TO
 $P \times I = \{60, 90, 120\} \times \{30, 40, 50\}$

# Views (P)	Noise Level (I)	MLP Architectures			
		No FFM	Pos Enc	Gaussian	CoIL
60	30	33.95	15.25	35.52	37.34
	40	42.62	21.79	43.88	43.68
	50	46.33	23.92	48.26	48.41
90	30	34.93	23.68	37.08	38.56
	40	43.82	30.56	44.95	44.72
	50	48.08	35.34	50.05	50.50
120	30	36.24	22.91	39.68	39.34
	40	44.81	24.37	45.65	45.29
	50	49.68	26.52	51.06	51.59




Fig. 6. Quantitative evaluation of the CoIL field for different projection numbers (P) and noise levels (I). The plot is divided into three regions, corresponding to P equal to 60, 90, and 120, respectively. Within each region, the average SNR values of the generated sinograms are plotted against different input SNR values, which are also drawn by the dotted horizontal lines for better visualization. First, note how CoIL generally produces measurement fields of better SNR than the noise level in the measurements. Second, the figure highlights that the quality of the generated CoIL field improves as the number of views increases or the noise level decreases.

TABLE II
THE AVERAGE SNR OF THE SINOGRAMS GENERATED BY LINEAR, BICUBIC,
CONVNET [72], AND CoIL IN THE SCENARIOS CORRESPONDING TO
 $P \times I = \{60, 90, 120\} \times \{30, 40, 50\}$

# Views (P)	Noise Level (I)	Interpolation methods			CoIL
		Linear	Bicubic	ConvNet	
60	30	30.37	29.51	41.37	<u>37.34</u>
	40	35.21	34.82	46.49	<u>43.68</u>
	50	36.28	36.10	50.34	<u>48.41</u>
90	30	31.18	30.16	42.42	<u>38.56</u>
	40	37.95	37.36	48.44	<u>44.72</u>
	50	40.14	39.94	52.36	<u>50.50</u>
120	30	31.23	30.42	43.53	<u>39.34</u>
	40	39.36	38.80	49.62	<u>45.29</u>
	50	43.03	42.86	54.35	<u>51.59</u>




Fig. 7. SNR improvements due to CoIL for each reconstruction algorithm. The plot is divided into three regions, corresponding to 60, 90, and 120 projections, respectively. Within each region, the average SNR improvement is plotted against the reconstruction method. The vertical axis is in log-scale for better visualization. Note that CoIL consistently improves the average SNR values for all the considered algorithms in every scenario.

other patients are used for training the FBP-UNet and the deep denoiser in GM-RED. We implemented $\mathbf{A}$ and its adjoint $\mathbf{A}^T$ by using RayTransform from the Operator Discretization Library (ODL) [101], which allows fast computation using a GPU backend. We synthesized the test sinograms corresponding to $P \in \{60, 90, 120\}$ projection views, each further contaminated by three noise levels equivalent to the input signal-to-noise ratio (SNR) of $I \in \{30, 40, 50\}$ dB. SNR is also used as a metric to quantify the reconstruction quality

$$\text{SNR}(\hat{\mathbf{x}}, \mathbf{x}) \triangleq 20 \log_{10} \left(\frac{\|\mathbf{x}\|_2}{\|\mathbf{x} - \hat{\mathbf{x}}\|_2} \right). \quad (18)$$

We denote the SNR values averaged over all test images as *average SNR*.

For each test image, CoIL trains separate MLPs to represent its full measurement field in different scenarios $P \times I = \{60, 90, 120\} \times \{30, 40, 50\}$. We conducted all the experiments, as well as the training of all neural networks, on a machine equipped with an Intel Xeon Gold 6130 Processor and four Nvidia GeForce GTX 1080 Ti GPUs. It takes about 30 minutes to train one MLP on this machine using one GPU.

B. Effectiveness of the FFM Layer

We first evaluate the effectiveness of the proposed FFM layer used in the coordinate-based MLP. We trained and compared three networks where: (a) the FFM layer is not implemented (*No FFM*); (b) the FFM layer implements the positional encoding where $k_i = 2^{i-1}$ (*Pos Enc*); (c) the FFM layer implements random projection (*Gaussian*); and (d) the FFM layer implements the proposed linear expansion $k_i = (\pi i)/2$ (*CoIL*). In the simulations, we use these networks to generate the sinograms corresponding to 360 views. We set *Pos Enc/CoIL* to expand the input to $L = 10$ frequency components and *Random* to project the input to $d = 256$ features.

Table I summarizes the average SNR values of the sinograms generated by the three networks in all scenarios. Here, we use SNR as the quality metric in the sinogram space because it enables straightforward comparison with the original measurements whose noise level is characterized by input SNR.




Fig. 8. Visual illustration of reconstruction with and without CoIL using the several methods. CoIL generates measurement fields corresponding to 360 (for FBP, TV, and RED) and 90 (used for FBP-UNet) views from $P = 60$ measurements with $I = 40$ dB noise. Each image is labeled with its SNR value with respect to the ground truth displayed in the left-most column. The visual differences are highlighted in the bounding boxes using green arrows. Note how CoIL enables the recovery of certain details missing in the reconstructions without it.




Fig. 9. Visual illustration of reconstruction with and without CoIL using the several methods. CoIL generates measurement fields corresponding to 360 (for FBP, TV, and RED) and 135 (used for FBP-UNet) views from $P = 90$ measurements with $I = 40$ dB noise. Each image is labeled with its SNR value with respect to the ground truth displayed in the left-most column. The visual differences are highlighted in the bounding boxes using green arrows.

As shown in the table, *CoIL* consistently achieves significantly higher SNR values than both *No FFM* and *Pos Enc*. Our interpretation is that *No FFM* is unable to represent the high-frequency variations in the measurement field, while *Pos Enc* overfits to noise by containing too many high-frequency components. We

observe that the sampling pattern in *CoIL* better captures the nature of the measurements without overfitting to the noise. Furthermore, *CoIL* also performs better than *Random* in most scenarios. In particular, *CoIL* leads to about 2 dB improvements for fewer measurements at higher noise levels ($I = 30$




Fig. 10. Visual illustration of reconstruction with and without CoIL using the several methods. CoIL generates measurement fields corresponding to 360 (for FBP, TV, and RED) and 180 (used for FBP-UNet) views from $P = 120$ measurements with $I = 40$ dB noise. Each image is labeled with its SNR value with respect to the ground truth displayed in the left-most column. The visual differences are highlighted in the bounding boxes using green arrows. Visual examples reconstructed with and without CoIL using the considered methods.

and $P = 60$). To better illustrate the difference, a set of visual examples are presented in Fig. 5, which plots the sinograms and their FBP reconstructions obtained by each network for $P = 120$ and $I = 40$ dB. Specifically, *No FFM* is able to represent the general structure of the sinogram but fails in generating the details; *Pos Enc* produces strong artifacts in its sinogram due to its FFM layer. Both *Random* and *CoIL* succeed in both representing the high-frequency details and avoiding strong artifacts in the generated measurements. However, *CoIL* still quantitatively outperforms *Random* in terms of SNR in most scenarios. In particular, the SNR margin of CoIL is greater than 1.5 dB when the number of views is relatively limited ($P = \{60, 90\}$) and the noise level is high ($I = 30$). The improvement in the sinogram quality is also reflected in the SNR values obtained after FBP reconstruction. Note how *CoIL* significantly differs from other approaches in the regions highlighted by the bounding boxes.

We have also investigated the evolution of the sinogram quality for different number of views and noise levels. Fig. 6 plots the SNR of the sinograms obtained by CoIL against the input SNR ($I \in \{30, 40, 50\}$) for different number of views ($P \in \{60, 90, 120\}$). The three dotted horizontal lines in the figure highlight each I value. We first note that CoIL generates sinograms that generally have higher SNR than the noise level in the measurements. In particular, when $I = 30$ dB, the average SNR values are more than 7 dB higher for every P . This highlights the ability of CoIL to generate high-quality sinograms. The figure also demonstrates that the SNR values improve as the number of views increases or noise level decreases. This highlights that the quality of the CoIL fields can be improved by having more measurements or acquiring those that are less noisy.

C. Evaluation of Sinogram Interpolation

In this section, we evaluate the interpolation performance of CoIL relative to three interpolation methods: (1) Linear, (2) Bicubic, and (3) interpolation using a pre-trained CNN (ConvNet). ConvNet is implemented by following the work [72], where separate CNNs are trained to map the sinograms of $P = \{60, 90, 120\}$ views to that of 360 views. Note that the training of ConvNet requires a dataset of ground-truth sinograms, while CoIL is trained exclusively on a single under-sampled sinogram of the desired object itself.

Table II summarizes the average SNR values of the sinograms generated by each aforementioned algorithm with respect to the ground truth for all projection views (P) and input SNR values (I). The best SNR values and the results of CoIL are respectively indicated by bold and underlined fonts. Note how CoIL significantly outperforms the two conventional interpolation methods, *Linear* and *Bicubic*, in all scenarios, showing the benefit of using CoIL for capturing complex patterns in the measurements. Additionally, CoIL achieves reasonable results compared to ConvNet, even though, unlike ConvNet, CoIL is not a traditional supervised learning method. The performance gap between CoIL and ConvNet is the smallest when the input SNR is high, demonstrating that the benefit of ConvNet is primarily due to its ability to separate noise from the data, achieved via traditional supervised training.

D. Evaluation of Reconstruction Performance

We next highlight the benefit of CoIL for image reconstruction. We trained all our MLPs by using the FFM layer based on our linear expansion. We implemented FBP by using `fbp-op`

TABLE III
THE AVERAGE SNR VALUES OBTAINED WITH AND WITHOUT COIL BY USING FBP, FISTA-TV, GM-RED, AND FBP-UNET IN THE SCENARIOS CORRESPONDING TO $P \times I = \{60, 90, 120\} \times \{30, 40, 50\}$

# Views (P)	Noise Level (I)	without CoIL				with CoIL			
		FBP	FISTA-TV	GM-RED	FBP-UNet	FBP	FISTA-TV	GM-RED	FBP-UNet
60	30	0.15	22.66	22.77	23.44	19.45	22.81	23.01	24.17
	40	9.09	26.08	27.12	27.08	23.48	26.95	27.42	27.93
	50	14.25	29.37	30.75	29.52	24.99	29.78	30.88	30.54
90	30	1.95	23.32	23.37	24.43	20.15	23.58	23.64	25.20
	40	10.92	26.98	28.86	28.57	24.14	28.28	29.31	29.28
	50	16.07	30.76	31.71	31.84	25.69	31.78	32.19	32.23
120	30	3.21	23.79	24.00	24.68	20.63	24.08	24.39	25.62
	40	12.10	27.59	29.30	29.18	24.52	28.95	29.79	29.71
	50	17.07	31.53	32.31	32.80	26.02	32.89	33.02	33.31

from the ODL package under the default parameter setting. We used DnCNN [49] to build the deep denoiser within GM-RED. In every experiment, we selected the network achieving the highest SNR value from the ones corresponding to five noise levels $\sigma \in \{5, 10, 15, 20, 25\}$. For FBP, FISTA-TV, and GM-RED, CoIL generates the measurement field with 360 projection views from the test measurements. FBP-UNet corresponds to our own implementation of the architecture used in [9]. The network was trained in the usual supervised fashion by directly predicting the ground truth from the FBP reconstruction using the ℓ_2 -loss [11]. For FBP-UNet using CoIL, we trained separate CNN models on the datasets consisting of the measurements having $1.5 \times P = \{90, 135, 180\}$ projection views and used CoIL to generate additional measurements to achieve that number. As a baseline, we trained a separate FBP-UNet that directly predicts the ground truth from the test measurements. Note that the baseline networks correspond to the optimal performance that FBP-UNet can achieve for the test measurements without integrating CoIL. In order to stabilize FBP-UNet, we trained these networks using the data with random fluctuations in both projection views (± 15) and noise amount (± 5 dB).

Fig. 7 quantitatively evaluates the improvements in imaging quality due to CoIL for all the considered reconstruction algorithms. For each algorithm, we plot the difference between the SNR obtained with and without CoIL. One can clearly observe that CoIL leads to significant quality improvements for all the algorithms. Remarkably, for the higher noise level ($I = 30$ dB), the average improvement by CoIL can sometimes be as high as 20 dB for FBP. In general, CoIL leads to SNR improvements for all algorithms, including for model-based and DL-based methods. In particular, for $P = 60$ and $I = 50$ dB, FBP-UNet without CoIL achieves 29.52 dB, while FBP-UNet with CoIL achieves 30.54 dB, which is nearly 1 dB improvement. The exact numbers obtained by each algorithm are summarized in Table III. These results highlight that CoIL is able to accurately represent the measurement field and generate high-fidelity measurements that can be used to improve image reconstruction. It can be observed that the improvements due to CoIL are less dramatic for regularized inversion methods, which we attribute to the strong influence of the regularizer on the final reconstruction quality, even when combined with CoIL.

Fig. 8 presents visual comparisons of images reconstructed with and without CoIL for $P = 60$ and $I = 40$. Each image is labeled with its SNR with respect to the ground truth and the visual differences are highlighted by arrows in the bounding boxes. This comparison highlights visual improvements due to CoIL. For example, consider the visual differences for FBP-UNet, where one can clearly see additional visual details after integration of the CoIL field. The better reconstruction quality with CoIL is also reflected in the higher SNR values. Additional visual comparisons in Figs. 9 and 10 also highlight the benefit of image reconstruction with CoIL.

V. CONCLUSION

The CoIL methodology developed in this paper is a new computational-imaging approach that leverages coordinate-based neural representations. CoIL seeks to represent the full measurement field as a single MLP network trained to map the measurement coordinates to their sensor responses. This makes CoIL a self-supervised model that can be trained without any external dataset. Extensive empirical results presented here demonstrate the improvements due to CoIL in the context of sparse-view CT, highlighting its great potential to work synergistically with existing image reconstruction methods.

The applications of CoIL in this paper are focused on sparse-view CT. One can anticipate the use of CoIL in other imaging modalities where the measurements, viewed as a function of sensing coordinates, exhibit learnable patterns, structures, or features. For example, we envision applications of CoIL in seismic imaging and diffraction tomography, where the measurements exhibit structural patterns relative to transmitter/receiver locations. On the other hand, it is unclear if the current version of CoIL would be applicable to applications where measurements are less structured or randomized relative to sensor locations, which is a common situation in traditional compressive sensing based on random measurement matrices.

It is worth mentioning potential limitations of CoIL. One limitation is the computational overhead for training the MLP, which when integrated to an overall imaging pipeline can significantly reduce the speed of image formation. A partial solution to this issue is to parallelize the MLP training across multiple GPUs, thus reducing the CoIL training overhead. Another possible


limitation of CoIL is that, by generating additional measurements, it can increase the per-iteration complexity of image reconstruction algorithms. This suggests that one has to carefully balance the number of synthesized measurements to achieve the best imaging performance under the computational constraints. One potential workaround for this problem in large-scale measurement settings is to replace batch reconstruction algorithms with their stochastic/online counterparts, as discussed in [18], [23], [25], [60], [102].

REFERENCES

- [1] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992.
- [2] M. A. T. Figueiredo and R. D. Nowak, "Wavelet-based image estimation: An empirical Bayes approach using Jeffreys' noninformative prior," *IEEE Trans. Image Process.*, vol. 10, no. 9, pp. 1322–1331, Sep. 2001.
- [3] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, Dec. 2006.
- [4] A. Danielyan, V. Katkovnik, and K. Egiazarian, "BM3D frames and variational image deblurring," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1715–1728, Apr. 2012.
- [5] M. T. McCann, K. H. Jin, and M. Unser, "Convolutional neural networks for inverse problems in imaging: A review," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 85–95, Nov. 2017.
- [6] A. Lucas, M. Iliadis, R. Molina, and A. K. Katsaggelos, "Using deep neural networks for inverse problems in imaging: Beyond analytical methods," *IEEE Signal Process. Mag.*, vol. 35, no. 1, pp. 20–36, Jan. 2018.
- [7] 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.
- [8] G. Wang, J. C. Ye, and B. De Man, "Deep learning for tomographic image reconstruction," *Nature Mach. Intell.*, vol. 2, no. 12, pp. 737–748, 2020.
- [9] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4509–4522, Sep. 2017.
- [10] E. Kang, J. Min, and J. C. Ye, "A deep convolutional neural network using directional wavelets for low-dose X-ray CT reconstruction," *Med. Phys.*, vol. 44, no. 10, pp. e360–e375, 2017.
- [11] Y. Sun, Z. Xia, and U. S. Kamilov, "Efficient and accurate inversion of multiple scattering with deep learning," *Opt. Exp.*, vol. 26, no. 11, pp. 14678–14688, May 2018.
- [12] Y. Han and J. C. Ye, "Framing U-Net via deep convolutional framelets: Application to sparse-view CT," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1418–1429, Jun. 2018.
- [13] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Glob. Conf. Signal Process. Inf. Process. (GlobalSIP)*, Austin, TX, USA 2013, pp. 945–948.
- [14] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," *SIAM J. Imag. Sci.*, vol. 10, no. 4, pp. 1804–1844, 2017.
- [15] S. Ono, "Primal-dual plug-and-play image restoration," *IEEE Signal Process. Lett.*, vol. 24, no. 8, pp. 1108–1112, Aug. 2017.
- [16] U. S. Kamilov, H. Mansour, and B. Wohlberg, "A plug-and-play priors approach for solving nonlinear imaging inverse problems," *IEEE Signal Process. Lett.*, vol. 24, no. 12, pp. 1872–1876, Dec. 2017.
- [17] S. A. Bigdeli, M. Zwicker, P. Favaro, and M. Jin, "Deep mean-shift priors for image restoration," in *Proc. Adv. Neural Inf. Process. Syst.* vol. 30, 2017, pp. 763–772.
- [18] Z. Wu, Y. Sun, A. Matlock, J. Liu, L. Tian, and U. S. Kamilov, "SIMBA: Scalable inversion in optical tomography using deep denoising priors," *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1163–1175, Oct. 2020.
- [19] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. S. Kamilov, "RARE: Image reconstruction using deep priors learned without ground truth," *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1088–1099, Oct. 2020.
- [20] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-play ADMM for image restoration: Fixed-point convergence and applications," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 84–98, Mar. 2017.
- [21] G. T. Buzzard, S. H. Chan, S. Sreehari, and C. A. Bouman, "Plug-and-play unplugged: Optimization free reconstruction using consensus equilibrium," *SIAM J. Imag. Sci.*, vol. 11, no. 3, pp. 2001–2020, Sep. 2018.
- [22] E. K. Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, and W. Yin, "Plug-and-play methods provably converge with properly trained denoisers," *Proc. 36th Int. Conf. Mach. Learn. (ICML)*, vol. 97, pp. 5546–5557, 2019.
- [23] Y. Sun, B. Wohlberg, and U. S. Kamilov, "An online plug-and-play algorithm for regularized image reconstruction," *IEEE Trans. Comput. Imag.*, vol. 5, no. 3, pp. 395–408, Sep. 2019.
- [24] X. Xu, Y. Sun, J. Liu, B. Wohlberg, and U. S. Kamilov, "Provable convergence of plug-and-play priors with MMSE denoisers," *IEEE Signal Process. Lett.*, vol. 27, pp. 1280–1284, Jul. 2020.
- [25] Y. Sun, J. Liu, and U. S. Kamilov, "Block coordinate regularization by denoising," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 380–390.
- [26] J. Zhang and B. Ghanem, "ISTA-Net: Interpretable optimization-inspired deep network for image compressive sensing," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018, pp. 1828–1837.
- [27] Y. Yang, J. Sun, H. Li, and Z. Xu, "Deep ADMM-Net for compressive sensing MRI," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 10–18.
- [28] A. Hauptmann et al., "Model-based learning for accelerated, limited-view 3-D photoacoustic tomography," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1382–1393, Jun. 2018.
- [29] J. Adler and O. Öktem, "Learned primal-dual reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1322–1332, Jun. 2018.
- [30] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-based deep learning architecture for inverse problems," *IEEE Trans. Med. Imag.*, vol. 38, no. 2, pp. 394–405, Feb. 2019.
- [31] S. A. Hosseini, B. Yaman, S. Moeller, M. Hong, and M. Akcakaya, "Dense recurrent neural networks for accelerated MRI: History-cognizant unrolling of optimization algorithms," *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1280–1291, Oct. 2020.
- [32] I. Y. Chun, Z. Huang, H. Lim, and J. Fessler, "Momentum-Net: Fast and convergent iterative neural network for inverse problems," *IEEE Trans. Patt. Anal. Machine Intell.*, pp. 1–1, 2020.
- [33] B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akcakaya, "Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data," *Magn. Reson. Med.*, vol. 84, no. 6, pp. 3172–3191, Jul. 2020.
- [34] H. K. Aggarwal and M. Jacob, "J-MoDL: Joint model-based deep learning for optimized sampling and reconstruction," *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1151–1162, Oct. 2020.
- [35] M. Kellman et al., "Memory-efficient learning for large-scale computational imaging," *IEEE Trans. Comput. Imag.*, vol. 6, pp. 1403–1414, Sep. 2020.
- [36] J. Liu, Y. Sun, W. Gan, X. Xu, B. Wohlberg, and U. S. Kamilov, "SGD-Net: Efficient model-based deep learning with theoretical guarantees," *IEEE Trans. Comput. Imag.*, vol. 7, pp. 598–610, Jun. 2021.
- [37] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 405–421.
- [38] R. Martin-Brualla, N. Radwan, M. SM Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth, "NeRF in the wild: Neural radiance fields for unconstrained photo collections," 2020, pp. 7210–7219, *arXiv:2008.02268*.
- [39] K. Zhang, G. Riegler, N. Snively, and V. Koltun, "NeRF: Analyzing and improving neural radiance fields," 2020, *arXiv:2010.07492*.
- [40] M. Tancik et al., "Fourier features let networks learn high frequency functions in low dimensional domains," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 7537–7547.
- [41] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [42] R. Glowinski and A. Marroco, "Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de dirichlet non linéaires," *ESAIM: Math. Model. Numer. Anal. - Modélisation Mathématique et Analyse Numérique*, vol. 9, no. R2, pp. 41–76, 1975.
- [43] D. Gabay and B. Mercier, "A dual algorithm for the solution of nonlinear variational problems via finite element approximation," *Comput. Math. Appl.*, vol. 2, no. 1, pp. 17–40, 1976.
- [44] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jul. 2011.

- [45] J. J. Moreau, "Proximité et dualité dans un espace Hilbertien," *Bull. Soc. Math. France*, vol. 93, pp. 273–299, 1965.
- [46] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Med. Image Comput. Comput.- Assist. Interv. (MICCAI)*, 2015, pp. 234–241.
- [47] T. Würfl *et al.*, "Deep learning computed tomography: Learning projection-domain weights from image domain in limited angle problems," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1454–1463, Jun. 2018.
- [48] B. Zhu, J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen, "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, no. 7697, pp. 487–492, 2018.
- [49] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.
- [50] K. Zhang, W. Zuo, and L. Zhang, "FFDNet: Toward a fast and flexible solution for CNN-based image denoising," *IEEE Trans. Image Process.*, vol. 27, no. 9, pp. 4608–4622, Sep. 2018.
- [51] G. Song, Y. Sun, J. Liu, Z. Wang, and U. S. Kamilov, "A new recurrent plug-and-play prior based on the multiple self-similarity network," *IEEE Signal Process. Lett.*, vol. 27, no. 1, pp. 451–455, Feb. 2020.
- [52] S. Sreehari *et al.*, "Plug-and-play priors for bright field electron tomography and sparse interpolation," *IEEE Trans. Comput. Imag.*, vol. 2, no. 4, pp. 408–423, Dec. 2016.
- [53] K. Zhang, W. Zuo, S. Gu, and L. Zhang, "Learning deep CNN denoiser prior for image restoration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 3929–3938.
- [54] Y. Sun, S. Xu, Y. Li, L. Tian, B. Wohlberg, and U. S. Kamilov, "Regularized fourier ptychography using an online plug-and-play algorithm," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, 2019, pp. 7665–7669.
- [55] K. Zhang, W. Zuo, and L. Zhang, "Deep plug-and-play super-resolution for arbitrary blur kernels," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 1671–1681.
- [56] R. Ahmad *et al.*, "Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 105–116, Jan. 2020.
- [57] T. Meinhardt, M. Moeller, C. Hazirbas, and D. Cremers, "Learning proximal operators: Using denoising networks for regularizing inverse imaging problems," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 1799–1808.
- [58] T. Tirer and R. Giryes, "Image restoration by iterative denoising and backward projections," *IEEE Trans. Image Process.*, vol. 28, no. 3, pp. 1220–1234, Mar. 2019.
- [59] A. M. Teodoro, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "A convergent image fusion algorithm using scene-adapted Gaussian-mixture-based denoising," *IEEE Trans. Image Process.*, vol. 28, no. 1, pp. 451–463, Jan. 2019.
- [60] Y. Sun, J. Liu, Y. Sun, B. Wohlberg, and U. Kamilov, "Async-RED: A provably convergent asynchronous block parallel stochastic method using deep denoising priors," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2021.
- [61] G. Mataev, P. Milanfar, and M. Elad, "DeepRED: Deep image prior powered by RED," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops (ICCVW)*, 2019.
- [62] C. Metzler, P. Schniter, A. Veeraraghavan, and R. Baraniuk, "prDeep: Robust phase retrieval with a flexible deep network," in *Proc. 35th Int. Conf. Mach. Learn. (ICML)*, 2018, pp. 3501–3510.
- [63] E. T. Reehorst and P. Schniter, "Regularization by denoising: Clarifications and new interpretations," *IEEE Trans. Comput. Imag.*, vol. 5, no. 1, pp. 52–67, Mar. 2019.
- [64] R. Cohen, M. Elad, and P. Milanfar, "Regularization by denoising via fixed-point projection (RED-PRO)," *SIAM J. Imag. Sci.*, vol. 14, no. 3, pp. 1374–1406, 2021.
- [65] K. Gregor and Y. LeCun, "Learning fast approximation of sparse coding," in *Proc. 27th Int. Conf. Mach. Learn. (ICML)*, 2010, pp. 399–406.
- [66] S. Biswas, H. K. Aggarwal, and M. Jacob, "Dynamic MRI using model-based deep learning and STORM priors: MoDL-STORM," *Magn. Reson. Med.*, vol. 82, no. 1, pp. 485–494, Jul. 2019.
- [67] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed sensing using generative models," *Proc. 34th Int. Conf. Mach. Learn. (ICML)*, vol. 70, pp. 537–546, 2017.
- [68] V. Shah and C. Hegde, "Solving linear inverse problems using gan priors: An algorithm with provable guarantees," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, 2018, pp. 4609–4613.
- [69] A. Jalal, L. Liu, A. G. Dimakis, and C. Caramanis, "Robust compressed sensing using generative models," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 713–727.
- [70] N. Shlezinger, J. Whang, Y. C. Eldar, and A. G. Dimakis, "Model-Based Deep Learning," 2020, *arXiv:2012.08405*.
- [71] H. Lee, J. Lee, and S. Cho, "View-interpolation of sparsely sampled Sinogram using convolutional neural network," in *Medical Imaging 2017: Image Processing*, M. A. Styner and E. D. Angelini, Eds., vol. 10133, Bellingham, WA, USA: SPIE, 2017, pp. 617–624.
- [72] H. Lee, J. Lee, H. Kim, B. Cho, and S. Cho, "Deep-neural-network-based Sinogram synthesis for sparse-view CT image reconstruction," *IEEE Trans. Radiat. Plasma Med. Sci.*, vol. 3, no. 2, pp. 109–119, Mar. 2019.
- [73] R. Anirudh, H. Kim, J. J. Thiagarajan, K. Aditya Mohan, K. Champley, and T. Bremer, "Lose the views: Limited angle CT reconstruction via implicit Sinogram completion," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018, pp. 6343–6352.
- [74] M. U. Ghani and W. C. Karl, "Fast enhanced CT metal artifact reduction using data domain deep learning," *IEEE Trans. Comput. Imag.*, vol. 6, pp. 181–193, Aug. 2020.
- [75] M. U. Ghani and W. C. Karl, "Data and image prior integration for image reconstruction using consensus equilibrium," *IEEE Trans. Comput. Imag.*, vol. 7, pp. 297–308, 2021.
- [76] M. U. Ghani, "Data and image domain deep learning for computational imaging," Ph.D. dissertation, Boston Univ., Boston, MA, USA, 2021.
- [77] B. E. H. Claus, Y. Jin, L. A. Gjesteb, G. Wang, and B. De Man, "Metal-artifact reduction using deep-learning based Sinogram completion: Initial results," in *Proc. 14th Int. Meeting Fully Three-Dimensional Image Reconstruction Radiol. Nucl. Med.*, 2017, pp. 631–634.
- [78] Q. De Man *et al.*, "A two-dimensional feasibility study of deep learning-based feature detection and characterization directly from CT sinograms," *Med. Phys.*, vol. 46, no. 12, pp. 790–800, 2019.
- [79] Y. Han, L. Sunwoo, and J. C. Ye, "k-space deep learning for accelerated MRI," *IEEE Trans. Med. Imag.*, vol. 39, no. 2, pp. 377–386, Feb. 2020.
- [80] M. Akçakaya, S. Moeller, S. Weingärtner, and K. Uğurbil, "Scan-specific robust artificial-neural-networks for K-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging," *Magn. Reson. Med.*, vol. 81, no. 1, pp. 439–453, 2019.
- [81] A. Shocher, N. Cohen, and M. Irani, "'Zero-shot' super-resolution using deep internal learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018, pp. 3118–3126.
- [82] L. P. Zuckerman, E. Naor, G. Pisha, S. Bagon, and M. Irani, "Across scales and across dimensions: Temporal super-resolution using deep internal learning," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 52–68.
- [83] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Salt Lake City, UT, USA, 2018, pp. 9446–9454.
- [84] J. Liu, Y. Sun, X. Xu, and U. S. Kamilov, "Image restoration using total variation regularized deep image prior," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, 2019, pp. 7715–7719.
- [85] Y. Gandelsman, A. Shocher, and M. Irani, "'Double-DIP': Unsupervised image decomposition via coupled deep-image-priors," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 11026–11035.
- [86] V. Sitzmann, J. N. P. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2020, pp. 7462–7473.
- [87] Z. Chen and H. Zhang, "Learning implicit fields for generative shape modeling," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 5939–5948.
- [88] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 165–174.
- [89] V. Sitzmann, M. Zollhöfer, and G. Wetzstein, "Scene representation networks: Continuous 3D-structure-aware neural scene representations," *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 1121–1132, 2019.
- [90] P. P. Srinivasan, B. Deng, X. Zhang, M. Tancik, B. Mildenhall, and J. T. Barron, "NeRV: Neural reflectance and visibility fields for relighting and view synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2021, pp. 7495–7504.

- [91] K. Park *et al.*, “Nerfies: Deformable neural radiance fields,” in *Proc. IEEE Int. Conf. Comp. Vis. (ICCV)*, 2021.
- [92] Y. Sung, W. Choi, C. Fang-Yen, K. Badizadegan, R. R. Dasari, and M. S. Feld, “Optical diffraction tomography for high resolution live cell imaging,” *Opt. Exp.*, vol. 17, no. 1, pp. 266–277, Dec. 2009.
- [93] U. S. Kamilov *et al.*, “Learning approach to optical tomography,” *Optica*, vol. 2, no. 6, pp. 517–522, Jun. 2015.
- [94] T.-A. Pham *et al.*, “Versatile reconstruction framework for diffraction tomography with intensity measurements and multiple scattering,” *Opt. Exp.*, vol. 26, no. 3, pp. 2749–2763, Feb. 2018.
- [95] K. Hornik *et al.*, “Multilayer feedforward networks are universal approximators,” *Neural Netw.*, vol. 2, no. 5, pp. 359–366, 1989.
- [96] N. Rahaman *et al.*, “On the spectral bias of neural networks,” in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 5301–5310.
- [97] A. Jacot, F. Gabriel, and C. Hongler, “Neural tangent Kernel: Convergence and generalization in neural networks,” *Proc. Adv. Neural Inf. Process. Syst.*, vol. 31, pp. 8571–8580, 2018.
- [98] D. Kingma and J. Ba, “ADAM: A method for stochastic optimization,” in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2015.
- [99] C. McCollough, “TU-FG-207A-04: Overview of the low dose CT grand challenge,” *Med. Phys.*, vol. 43, no. 6, pp. 3759–3760, 2016.
- [100] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Piscataway, NJ, USA: IEEE, 1988.
- [101] J. Adler and Ö. Öktem, “Solving ill-posed inverse problems using iterative deep neural networks,” *Inverse Problems*, vol. 33, no. 12, 2017, Art. no. 124007.
- [102] Y. Sun, Z. Wu, X. Xu, B. Wohlberg, and U. S. Kamilov, “Scalable plug-and-play ADMM with convergence guarantees,” *IEEE Trans. Comput. Imag.*, vol. 7, pp. 849–863, Jul. 2021.



Yu Sun (Student Member, IEEE) received the B.Eng. degree in electronics and information from Sichuan University, Chengdu, China, in 2015, and the M.S. degree in data analytics and statistics in 2017 from Washington University in St. Louis, St. Louis, MO, USA, where he is currently working toward the Ph.D. degree with the Computational Imaging Group. During his Ph.D., he worked as an Intern with Nvidia Cooperation in 2021. His research interests include computational imaging, machine learning, deep learning, and optimization.



Jiaming Liu (Student Member, IEEE) received the B.Sc. degree in electronic and information engineering from the University of Electronic Science and Technology of China, Chengdu, China, in 2017, and the M.S. degree in electrical and engineering from Washington University in St. Louis, St. Louis, MO, USA.

He is currently working toward the Ph.D. degree with the Computational Imaging Group. He worked as an Intern with Los Alamos National Laboratory in 2021. His research interests include deep learning,

image processing, computational imaging and optimization.



Mingyang Xie received the B.Sc. degree in computer science from Washington University, in St. Louis, MO, USA, in 2021. He is currently working toward the Ph.D. degree in computer science with the University of Maryland, College Park, MD, USA. His research interests include computational imaging and deep learning.



Brendt Wohlberg (Senior Member, IEEE) received the B.Sc. (Hons.) degree in applied mathematics, and the M.Sc. (applied science) and Ph.D. degrees in electrical engineering from the University of Cape Town, Cape Town, South Africa, in 1990, 1993, and 1996, respectively. He is currently a Staff Scientist with Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM, USA. His primary research interests include signal and image processing inverse problems and computational imaging. He was a co-recipient of the 2020 SIAM Activity Group on Imaging Science Best Paper Prize. He was an Associate Editor for the IEEE TRANSACTIONS ON IMAGE PROCESSING from 2010 to 2014, and for the IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING from 2015 to 2017, and was the Chair of the Computational Imaging Special Interest Group (now the Computational Imaging Technical Committee) of the IEEE SIGNAL PROCESSING SOCIETY from 2015 to 2017. He is currently an Associate Editor for the SIAM Journal on Imaging Sciences, and the Editor-in-Chief of the IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING.



Ulugbek S. Kamilov (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees in communication systems, and the Ph.D. degree in electrical engineering from EPFL, Lausanne, Switzerland, in 2008, 2011, and 2015, respectively. He is an Assistant Professor and the Director of Computational Imaging Group (CIG), Washington University in St. Louis, MO, USA. From 2015 to 2017, he was a Research Scientist with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA. His primary research area focuses on computational imaging with

a focus on signal and image processing, machine learning, and large-scale optimization.

Dr. Kamilov is an Associate Editor of IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING and a member of Computational Imaging Technical Committee of the IEEE Signal Processing Society. He was a recipient of the IEEE Signal Processing Society's 2017 Best Paper Award.